Multimedia Streaming Using Multiple TCP Connections

Sunand Tullimas, Thinh Nguyen, Richard Edgecombe
Oregon State University
Sen-ching Cheung
University of Kentucky

In recent years, multimedia applications over the Internet become increasingly popular. However, packet loss, delay, and time-varying bandwidth of the Internet have remained the major problems for multimedia streaming applications. As such, a number of approaches, including network infrastructure and protocol, source and channel coding have been proposed to either overcome or alleviate these drawbacks of the Internet. In this paper, we propose the MultiTCP system, a receiver-driven, TCP-based system for multimedia streaming over the Internet. Our proposed algorithm aims at providing resilience against SHORT TERM insufficient bandwidth by using MULTIPLE TCP connections for the same application. Our proposed system enables the application to achieve and control the desired sending rate during congested periods, which cannot be achieved using traditional TCP. Finally, our proposed system is implemented at the application layer, and hence, no kernel modification to TCP is necessary. We analyze the proposed system, and present simulation and experimental results to demonstrate its advantages over the traditional single TCP based approach.

Categories and Subject Descriptors: C.2 [**Computer-Communication Networks**]: General

General Terms: Design, Algorithm

Additional Key Words and Phrases: Multimedia streaming

1. INTRODUCTION

In recent years, there have been an explosive growth of multimedia applications over the Internet. All major news networks such as ABC and NBC now provide news with accompanying video clips. Several companies, such as MovieFlix [MovieFlix], also offer video on demand to broadband subscribers. However the quality of videos being streamed over the Internet is often low due to insufficient bandwidth, packet loss, and delay. To view a DVD quality video from an on demand video service, a customer must download either the entire video or a large portion of the video before starting to playback in order to avoid pauses caused by insufficient bandwidth during a streaming session. Thus, many techniques have been proposed to enable efficient multimedia streaming over the Internet. The source coding community has proposed scalable video [Tan and Zakhor 1999][Reyes et al. 2000], error-resilient coding, and multiple description [Reibman 2002] for efficient video streaming over the best-effort networks such as the Internet. A scalable video bit stream is coded in such a way to enable the server to easily and efficiently adapt the video bit rate to the current available bandwidth. Error-resilient coding and multiple description are aimed at improving the quality of the video in the presence of packet loss and long delay caused by retransmission. Channel coding techniques are also used to mitigate long delay for real-time applications such as video conferencing or IP-

telephony [Ma and Zarki 1998]. The main disadvantages of these approaches are first, specialized codecs are required and second, their performances are highly affected by the network traffic conditions.

From a network infrastructure perspective, Differentiated Services [Blake et al. 1998][Wang 2001] and Integrated Services [White 1997][Wang 2001] have been proposed to improve the quality of multimedia applications by providing preferential treatments to various applications based on their bandwidth, loss, and delay requirements. More recently, path diversity architectures that combine multiple paths and either source or channel coding have been proposed to provide larger bandwidth, and to combat packet loss efficiently [Nguyen and Zakhor 2004][Apostolopoulos 2001][Apostolopoulos 2002]. Nonetheless, these approaches cannot be easily deployed as they require significant changes in the network infrastructure.

The most straightforward approach is to transmit standard-based multimedia via existing IP protocols. The two most popular choices are TCP and UDP. A single TCP connection is not suitable for multimedia transmission because its congestion control may cause a large fluctuation in the sending rate. Unlike TCP, an UDP-based application is able to set the desired sending rate. If the network is not too much congested, the UDP throughput at the receiver would approximately equal to the sending rate. Since the ability to control the sending rate is essential to interactive and live streaming applications, majority of multimedia streaming systems use UDP as the basic building block for sending packets over the Internet. However, UDP is not a congestion aware protocol since it does not reduce its sending rate in presence of network congestion, and therefore potentially results in a congestion collapse. Congestion collapse occurs when a router drops a large number of packets due to its inability to handle a large amount of traffic from many senders at the same time. TCP-Friendly Rate Control Protocol (TFRC) has been proposed for multimedia streaming, using UDP however it also incorporates TCP-like congestion control mechanism [Floyd et al. 2000]. Another drawback of using UDP is its lack of reliable transmission and hence the application must deal with the packet loss.

Based on these drawbacks of UDP, we propose a new receiver-driven, TCP-based system for multimedia streaming over the Internet. The first version of this work was published in [Nguyen and Cheung 2005]. In particular, our proposed system, called *MultiTCP*, is aimed at providing resilience against *short-term* insufficient bandwidth by using *multiple* TCP connections for the same application. Furthermore, our system enables the application to achieve and control the sending rate during congested period, which in many cases, cannot be achieved using a single TCP connection. Finally, our proposed system is implemented at the application layer, and hence, no kernel modification to TCP is necessary. We note that our proposed MultiTCP is not designed for interactive applications such as video conferencing applications. The main difficulties in designing interactive applications over the Internet such as video conferencing are not only sufficient bandwidth, but also the stringent delay requirement. If a packet loss occurs, there may not be time to retransmit the lost packet due to long round trip time, e.g., 100 ms for two typical computers in the North America. Therefore, FEC or appropriate concealment schemes at the receiver are often used in conjunction with UDP to avoid

retransmissions. That said, our system is most appropriate for live and non-live video streaming applications in which, the user can tolerate some initial delay.

The rest of the paper is organized as follows. In Section 2, we describe other related works that utilize multiple network connections. In Section 3, we describe the two major drawbacks of using TCP for multimedia streaming: short-term insufficient bandwidth and lack of precise rate control. These drawbacks motivate the use of multiple TCP connections in our proposed system, which is described in Section 4. In Section 5, we demonstrate the performance of our system through NS simulations [Information Sciences Institute] and actual Internet experiments. Finally, we summarize our contributions in Section 6.

2. RELATED WORK

Our work has its root in the earlier work done by Crowcroft [Crowcroft and P.Oeschlin 1998]. In this work, the receiver window size is adjusted to achieve weighted proportional fair sharing web flows in order to provide end-to-end differentiated services. In similar spirit, the authors in [Semke et al. 1998] propose a technique for automatic tuning of receiver window size in order to increase the throughput of TCP. Our work is similar to these work in the sense that our algorithm also employs the receiver window size to achieve the desired throughput using multiple TCP connections for a multimedia streaming session.

There have been previous work on using multiple network connections to transfer data. For example, *path diversity* multimedia streaming framework [Apostolopoulos 2001][Apostolopoulos 2002][Nguyen and Zakhor 2004] provide multiple connections on different path for the same application. These work focus on either efficient source or channel coding techniques in conjunction with sending packets over multiple approximately independent paths. On the other hand, our work aims to increase and maintain the available throughput using multiple TCP connections on a single path.

In a similar approach, Chen et al. use multiple connections on a single path to improve throughput of a wired-to-wireless streaming video session [Chen and Zakhor 2004][Chen and Zakhor 2005]. This work focuses on obtaining maximum possible throughput and is based on TFRC [Floyd and Fall 1999] rather than TCP. In subsequent work [Chen and Zakhor 2006], Chen et al. use multiple TCP connections to maximize the throughput for wireless network. On the other hand, our work focuses on eliminating *short term* throughput reduction of TCP due to burst traffic and providing precise rate control for the application. As such, the analysis and rate control mechanism in our paper are different from those in [Chen and Zakhor 2004][Chen and Zakhor 2005][Chen and Zakhor 2006].

Another related work is Streaming Control Transmission Protocol (SCTP)[Internet Engineering Task Force 2000], designed to transport PSTN signaling messages over IP networks. SCTP allows user's messages to be delivered within multiple streams, but it is not clear how it can achieve the desired throughput in a congestion scenario. In addition, SCTP is a completely new protocol, as such the kernel of the end systems need to be modified.

There are also other work related to controlling TCP bandwidth. For example, the works in [Mehra and Zakhor 2003][P.Mehra and A.Zakhor 2005] focus on al-

locating bandwidth among flows with different priorities. This work assumes that the bottleneck is at the *last-mile* and that the required throughput for the desired application is achievable using a single TCP connection. On the other hand, our work does not assume the *last-mile* bottleneck, and the proposed *MultiTCP* system can achieve the desired throughput in variety of scenarios.

The work in [Dong et al. 2002] uses the receiver advertised window to limit the TCP video bandwidth in VPN link between video and proxy servers. The authors in [Liang et al. 2002] proposed an approach that leads to real-time applications that are responsive to network congestion, sharing the network resources fairly with other TCP applications.

3. DRAWBACKS OF TCP FOR MULTIMEDIA STREAMING

TCP is unsuitable for multimedia streaming due partly to its fluctuating throughput and its lack of precise rate control. TCP is designed for end-to-end reliability and fast congestion avoidance. To provide end-to-end reliability, a TCP sender retransmits the lost packets based on the packet acknowledgment from a TCP receiver. To react quickly to network congestion, TCP controls the sending rate based on a window-based congestion control which works as follows. The sender keeps track of a window of maximum number of unacknowledged packets, i.e., packets that have not been acknowledged by the receiver. In the steady state, the sender increases the window size W by $1/W$ upon successfully receiving an acknowledged packet, or equivalently, it increases the sending rate by one packet per round trip time. Upon encountering a loss, the window size is reduced by half, or equivalently, the sending rate is cut in half. In TCP, the receiver has the ability to set a maximum window size for the unacknowledged packets, hence imposing a maximum sending rate. Thus, in a non-congestion scenario, the application at the receiver can control the sending rate by setting the window size appropriately. On the other hand, during congestion, the actual throughput can be substantially low as the maximum window size may never be reached.

Based on the above discussion, we observe that a single packet loss can drop the TCP throughput abruptly and the low throughput lingers due to the slow increase of the window size. If there is a way to reduce this throughput reduction without modifying TCP, we can effectively provide higher throughput with proper congestion control and reliable transmission. In addition, if there is a way to control the TCP sending rate during congestion, then TCP can be made suitable for multimedia streaming. Unlike non real-time applications such as file transfer and email, precise control of sending rate is essential for interactive and live streaming applications due to several reasons. First, sending at too high a rate can cause buffer overflow in certain receivers with limited buffer such as mobile phones and PDAs. Second, sending at a rate lower than the coded bit rate results in pauses during a streaming session, unless a large buffer is accumulated before playback.

In the following section, we propose a system that can dynamically distribute streaming data over multiple TCP connections per application to achieve higher throughput and precise rate control. The control is performed entirely at the receiver side and thus, suitable for streaming applications where a single server may serve up to hundreds of receivers simultaneously.

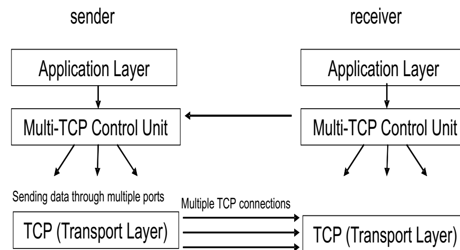

Fig. 1. *MultiTCP* system diagram.

4. MULTITCP OVERVIEW AND ANALYSIS

As mentioned in Section 3, the throughput reduction of TCP is attributed to the combination of (a) reduction of the sending rate by half upon detection of a loss event and (b) the slow increase of sending rate afterward or *congestion avoidance*. To alleviate this throughput reduction, one can modify TCP to (a) reduce the sending rate by a small factor other than half upon detection of a loss, or (b) speed up the congestion avoidance process, or (c) combine both (a) and (b). There are certain disadvantages associated with these approaches. First, these changes affect all TCP connections and must be performed by recompiling the OS kernel of the sender machine. Second, changing the decreasing multiplicative factor and the additive term in isolated machines may potentially lead to instability of TCP in a larger scale of the network. Third, it is not clear how these factors can be changed to dynamically control the sending rate.

As such, we propose a different approach: instead of using a traditional, single TCP connection, we use multiple TCP connections for a multimedia streaming application. Our approach does not require any modification to the existing TCP stack or kernel. Figure 1 shows a diagram of our proposed *MultiTCP* system. The MultiTCP control unit is implemented immediately below the application layer and above the transport layer at both the sender and the receiver. The MultiTCP control unit at the receiver receives the input specifications from streaming application which include the streaming rate. The MultiTCP control unit at the receiver measures the actual throughput and uses this information to control the rate precisely by using multiple TCP connections and dynamically changing receiver's window size for each connection. In the next two sections, we show how multiple TCP connections can mitigate the short term throughput reduction problem in a lightly loaded network and describe our mechanism to maintain the desired throughput in a congested network.

4.1 Alleviating Throughput Reduction In Lightly Loaded Network

In this section, we analyze the throughput reduction problem in a lightly loaded network and show how it can be alleviated by using multiple TCP connections.

When there is no congestion, the receiver can control the streaming rate in a single TCP connection quite accurately by setting the maximum the receiver's window

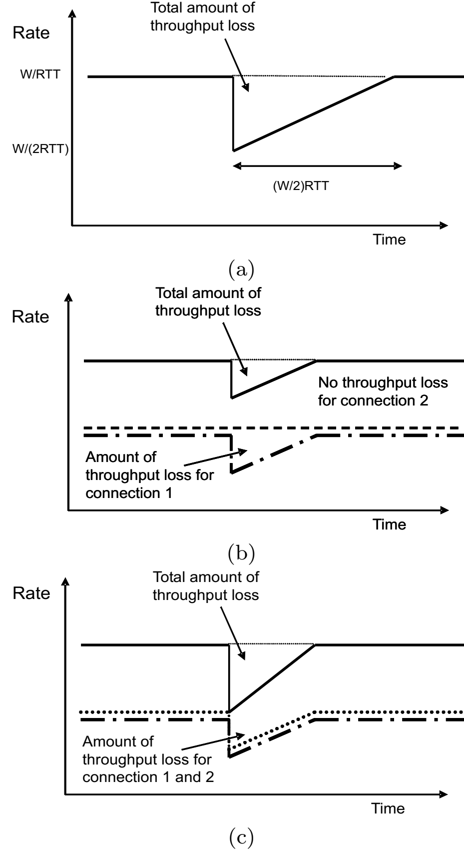

Fig. 2. Illustrations of throughput reduction for (a) one TCP connections with single loss; (b) two TCP connections with single loss; (c) two TCP connections with double losses.

size W_{max}. The effective throughput during this period is approximately equal to

$$T = \frac{W_{max} MTU}{RTT} \tag{1}$$

where RTT denotes the round trip time, including both propagation and queuing delay, between the sender and the receiver. MTU denotes the TCP maximum transfer unit, typically set at 1000 bytes. If a loss event occurs, the TCP sender instantly reduces its rate by half as shown in Figure 2(a). As a result, the area of the inverted triangular region in Figure 2(a) indicates the amount of data that would have been transmitted if there were no loss event. Thus, the amount of data reduction D equals to

$$D = (\frac{1}{2})(\frac{W_{max} MTU RTT}{2})(\frac{W_{max}}{2RTT}) = \frac{W_{max}^2 MTU}{8} \tag{2}$$

Note that the time it takes for the TCP window to increase from $W_{max}/2$ to W_{max} equals to $W_{max} RTT/2$ since the TCP window increases by one every round trip time. Clearly, if there are a burst of loss events during a streaming session, the

total throughput reduction can potentially be large enough to deplete the start up buffer, causing pauses in the playback.

Now let us consider the case where two TCP connections are used for the same application. Since we want to keep the same total streaming rate W_{max}/RTT as in the case of one TCP connection, we set $W'_{max} = W_{max}/2$ for each of the two connections as illustrated in Figure 2(b). Assuming that only a single loss event happens in one of the connections, the total throughput reduction would be equal to

$$D^{'} = \frac{(W'^{2}_{max}MTU)}{8} = \frac{(W^{2}_{max}MTU)}{32} = \frac{D}{4} \tag{3}$$

Equation (3) shows that, for a single loss event, the throughput reduction of using two TCP connections is four times less than that of using a single TCP connection. Even in the case when there are simultaneously losses on both connections as indicated in Figure 2(c), the throughput reduction is half of that of the single TCP. In general, let N denote the number of TCP connections for the same application and n be the number of TCP connections that suffer simultaneous losses during short congestion period, the amount of throughput reduction equals to

$$D_N = \frac{nW^{2}_{max}MTU}{N^2} \tag{4}$$

As seen in Equation (4), the amount of throughput reduction is inversely proportional to the square of the number of TCP connections used. Hence, using a only small number of TCP connections can greatly improve the resilience against TCP throughput reduction in lightly loaded network.

4.2 Control Streaming Rate in a Congested Network

In the previous section, we discuss the throughput reduction problem in a lightly loaded network and show that using multiple TCP connections can alleviate the problem. In a lightly loaded network condition, one can set the desired throughput T_d by simply setting the receiver window $W_{max} = T_d RTT/MTU$. However, in a moderately or heavily congested network, the throughput of a TCP does not depend on W_{max}, instead, it is determined by the degree of congestion. This is due to the fact that in a non-congested network, i.e. without packet loss, TCP rate would increase additively until $W_{max}MTU/RTT$ is reached, after that the rate would remain approximately constant at $W_{max}MTU/RTT$. However, in a congested network, a loss event would most likely occur before the sending rate reaches its limit and cut the rate by half, resulting in a throughput lower than $W_{max}MTU/RTT$.

A straightforward method for achieving a higher throughput than the available TCP throughput would be to use multiple TCP connections for the same application. Using multiple TCP connections results in a larger share of the fair bandwidth. Hence, one may argue that this is unfair to other TCP connections. On the other hand, one can view this approach as a way of providing higher priority for streaming applications over other non time-sensitive applications under resource constraints. We also note that one can use UDP to achieve the desired throughput. However unlike UDP, using multiple TCP connections can provide (a) congestion control mechanism to avoid congestion collapse, and (b) automatic retransmission

of lost packets. Assuming multiple TCP connections are used, there are still issues associated with providing the desired throughput in a congested network.

In order to maintain a constant throughput during a congested period, one possible approach is to increase the number of TCP connections until the measured throughput exceeds the desired one. This approach has a few drawbacks. First, the total throughput may still exceed the desired throughput by a large amount since the sending rate of each additional TCP connection may be too high. Second, if only a small number of TCP connections are required to exceed the desired throughput, this technique may not be resilient to the sudden increase in traffic as analyzed in Section 4.1. A better approach is to use a larger number of TCP connections but adjust the receiver window size of each connection to precisely control the sending rate. It is undesirable to use too many TCP connections as they use up system resources and may further aggravate an already congested network. In this paper, we consider two algorithms: the first($AdjustWindowSize()$) algorithm maintains a fixed number of TCP connections while varying the size of the receiver windows of each TCP connection to achieve the desired throughput. The second algorithm dynamically changes the number of TCP connections based on the network congestion.

4.2.1 *AdjustWindowSize Algorithm.* Our *AdjustWindowSize* algorithm uses a fixed, default number of TCP connections and only varies receiver window size to obtain the desired throughput T_d. Hence, the inputs to the algorithm are the desired user's throughput T_d and the number of TCP connections. Below are the steps of our algorithm.

We now discuss each step of the algorithm in detail and show how to choose appropriate values for the parameters. In the initialization steps, we found empirically, $N = 5$ works well in many scenarios. If user does not specify the number of TCP connections, the default value is set to $N = 5$. In step 2, we assume that the network is not congested initially, hence the total expected throughput and the total receiver window size W_s would equal to T_d and $T_d RTT/MTU$ respectively. Note that the average RTT can be obtained easily at the receiver.

In the running steps, δ should be chosen to be several times the round trip time since the sender cannot respond to the receiver changing window size for at least one propagation delay, or approximately half of RTT. As a result, the receiver may not observe the change in throughput until several RTTs later. In most scenarios, we found that setting the measuring interval δ in the range of $(6RTT, 8RTT)$ works quite well in practice. In step 3, the algorithm tries to increase the throughput by increasing the window size of each connection equally when the measured throughput is still smaller than the desired throughput. In this step, we also place a restriction on the maximum total receiver window size. In the congestion state, increasing the receiver window size will not increase the throughput since a packet loss is most likely to happen and thus, reduces the window by half before the the size of the congestion window reaches the receiver window size. Therefore, increasing the receiver window size much larger than $W_{max} = \frac{4T_m RTT}{3MTU}$ [Kurose and Ross 2005]-the average congestion window size under congestion, would not increase the TCP throughput. However, if we let w_i increase without bound, there will be a spike in throughput once the network becomes less congested. To prevent unnec-

Algorithm 1 AdjustWindowSize($numConn$)

1: $N = numConn$ {Initialization steps}
2: $w_i = \frac{T_d RTT}{(MTU)N}$ for connection i
{Running steps: The actual throughput T_m is measured at at every δ second and the algorithm dynamically changes the window size based on the measured T_m.}
3: **if** $T_m < T_d$ and $\sum_i^N w_i \leq \frac{f T_d RTT}{MTU}$ **then**
4: $D_s = \lceil |T_d - T_m| RTT/MTU \rceil$
5: Sort the connections in the increasing order of w_i
6: **while** $D_s > 0$ **do**
7: $w_i = w_i + 1$
8: $D_s = D_s - 1$
9: $i = (i + 1) \mod N$
10: **end while**
11: **else if** $T_m > T_d + \lambda$ **then**
12: Sort the connections in the decreasing order of w_i
13: **while** $D_s > 0$ **do**
14: $w_i = w_i - 1$
15: $D_s = D_s - 1$
16: $i = (i + 1) \mod N$
17: **end while**
18: **else**
19: keep the receiver window size the same.
20: **end if**

essary throughput fluctuation, our algorithm limits the sum of receiver window size W_s to $f \frac{T_d RTT}{MTU}$ where $f > 4/3$ is used to control the height of the throughput spike. Larger and small values of f result in higher and lower throughput spikes respectively, as discussed later in Section 5.

Next, step 4 to ste 9, the algorithm tries to increase the window size w_i for a subset of connections if the measured throughput is smaller than the desired throughput. There exists an optimal way for choosing a subset of connections for changing the window size and the corresponding increments in order to achieve the desired throughput. If the number of chosen connections for changing the window size and the corresponding window increments are small, then the time for achieving the desired throughput maybe longer than necessary. On the other hand, choosing too large a number of connections and increments may result in higher throughput than necessary. For example, assuming we have five TCP connections, each with RTT of 100 milliseconds, MTU equals to 1000 bytes, and the network is in non-congestion state, then changing the receiver window size of all the connections by one can result in a total change in throughput of $5(1000)/.1 = 50$ Kbytes per second. In a congested scenario, the change will not be that large, however, one may still want to control the throughput change to a certain granularity. To avoid these drawbacks, our algorithm chooses the number of connections for changing their window size and the corresponding increments based on the current difference between the desired and measured throughput.

The reasoning behind the algorithm is as follows. Consider the case when $T_m < T_d$. If there is no congestion, setting the sum of window size increments D_s from all the connections to $\lceil (T_d - T_m) RTT/MTU \rceil$ would result in a throughput increment of $T_d - T_m$, hence the desired throughput would be achieved. If there is a congestion, this total throughput increment would be smaller. However, subsequent rounds of window increment would allow the algorithm to reach the desired throughput. This method effectively produces a large or small total window increment at every sampled point based on a large or small difference between the measured and desired throughput, respectively. Steps 6 to 9 in the above algorithm ensure the total throughput increment is equally contributed by all the connections. On the other hand, if only one connection j is responsible for all the throughput, i.e. $w_i = 0$ for $j \neq i$, then we simply have a single connection whose throughput can be substantially reduced in a congested scenario. We note that using our algorithm, w_i's for different connections at any point in time differ from each other at most by one. The scenario where $T_m > T_d$ is similar.

4.2.2 *Adjust Number of Connections.* While using a fixed number connections, e.g. 5 works well in many scenarios, nevertheless, in a lightly loaded network, this number of connections maybe unnecessary large. In this section, we introduce a new algorithm that changes the number of connections dynamically depending upon the congestion in the network.

Our algorithm is based on the following observations. If the congestion in the network is reduced, the algorithm can reduce the number of connections and adjust the congestion window size of the remaining connections. Similarly, if the network congestion increases, the algorithm can increase the number of connections to obtain the desired throughput while avoid causing too much congestion. There are many indicators for the change in congestion, e.g. the change in RTT. In our algorithm we use the Congestion ratio. The congestion ratio is given by W_s/W_n where $W_s = \sum_i w_i$ and W_n denotes the window size of single connection when there is no congestion, i.e, $W_n = T_d RTT/MTU$. If the congestion ratio is one or less then we say that there is no congestion in the network. But if it increases over one then we say that there is congestion in the network. The algorithm for increasing and decreasing number of connections is based on the congestion ratio. In particular, we store the value of congestion ratio for every iteration in *prevcongR* (previous congestion ratio), and compare the current congestion ratio*prevconR*. If the current value is greater than *prevongR*, the congestion level must have increased. On the other hand, if the opposite is true, then the network congestion must have decreased. The below algorithm is used to dynamically adjust the number of connections. We now discuss each cases in the *AdjustNumConnections*. In the first case, the congestion ratio decreases by β, thus we can say that the congestion in the network reduces. If the measured throughput, T_m, is greater than the desired throughput, T_d, we reduce the number of connections being used for streaming by one. Note that we only decrease the number of connections when $T_m > T_d$. These steps aim to minimize the number of used connections when the network becomes less congested.

In the second case, the congestion ratio increases by β, thus we can say that the congestion in the network increases. If the current congestion ratio, W_s/W_n, is less

Algorithm 2 AdjustNumConnections()

1: **if** ($\frac{W_s}{W_n} < prevcongR - \beta$) and ($T_m > T_d$) **then**
2: $prevcongR = \frac{W_s}{W_n}$
3: stop one connection
4: **end if**{cf is a congestion factor to be discussed shortly}
5: **if** ($\frac{W_s}{W_n} > prevcongR + \beta$) and ($T_m < T_d$) and ($\frac{W_s}{W_n} < cf$) **then**
6: $prevcongR = \frac{W_s}{W_n}$
7: start one connection
8: **end if**
9: **if** $\frac{W_s}{W_n} \geq cf$ **then**
10: $prevcongR = \frac{W_s}{W_n}$
11: stop one connection
12: **end if**
13: return number of connections

than the congestion factor, cf, and if the measured throughput, T_m, is less than the desired throughput, T_d, we increase the number of connections being used for streaming by one. cf is a special factor which is used to determine if the network is slightly congested or severely congested. If the congestion factor is smaller than cf, the network is not severely congested, and the algorithm is allowed to start a new TCP connection.

In the third case, the current congestion ratio, is greater than or equal to the congestion factor, cf, we decrease the number of connections being used for streaming. This step 3 is used when the congestion in the network increases to a level that even by increasing the number of connections we cannot achieve the desired throughput, T_d. This means that the network is severely congested. So we have to reduce the network congestion by stopping as many connections as possible. If none of the three cases above are true, the the number of connections remain the same, and the $AdjustWindowSize$ algorithm is the only mechanism to control the desired bandwidth.

Once the appropriate number of connections has been determined, it is used as the input to the $AdjustWindowSize$ algorithm to control the desired throughput. Note that the value β is used to control the frequency of starting and stopping

Algorithm 3 VariedConnVariedWinSize()

1: $N = AdjustNumConn()$
2: $AdjustWindowSize(N)$

connections. If β is small, the fluctuation of throughput increases, and vice versa. In our Red Hat Linux implementation of the MultiTCP system, we note that setting $\beta = 1$ works well in most situations.

4.2.3 *Remarks on Sender.* At the sender, data is divided into packets of equal size. These packets are always sent in order. The MultiTCP system chooses the TCP connection to send the next packet in a round robin fashion. If a particular

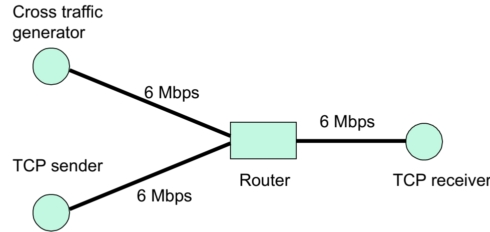

Fig. 3. *Simulation topology.*

TCP connection is chosen to send the next packet, but it is blocked due to TCP congestion mechanism, the MultiTCP system chooses the first available TCP connection in a round robin manner. For example, suppose there are 5 connections, denoted by TCP1 to TCP5. If none of TCP connection is blocked, packet 1 would be sent by TCP1, packet 2 by TCP2, and so on. If TCP1 is blocked, then TCP2 would send packet 1 and TCP3 would send packet 2, and so on. When it is TCP1's turn again and if TCP1 is not blocked, it would send packet 5. This is similar to socket striping technique in [Leigh et al. 2001].

5. RESULTS

In this section, we show simulation results using NS and the results produced using actual network to demonstrate the effectiveness of our MultiTCP system in achieving the required throughput as compared to the traditional single TCP approach.

5.1 Results for $AdjustWindowSize$ Algorithm

5.1.1 *NS Simulation.* Our simulation setup consists of a sender, a receiver, and a traffic generator connected together through a router to form a dumb bell topology as shown in Figure 3. The bandwidth and propagation delay of each link in the topology are identical, and are set to 6 Mbps and 20 milliseconds, respectively. The sender streams 800 kbps video to the receiver continuously for a duration of 1000s, while the traffic generator generates cross traffic at different times by sending packets to the receiver using either long term TCPs or short bursts of UDPs. In particular, from time $t = 0$ to $t = 200$s, there is no cross traffic. From $t = 200$s to $t = 220$s and $t = 300$s to $t = 340$s, bursts of UDPs with rate of 5.5 Mbps are generated from the traffic generator node to the receiver. At $t = 500$s the traffic generator opens 15 TCPs connections to the receiver, and 5 additional TCP connections at $t = 750$s. We now consider this setup under three different scenarios: (a) the sender uses only one TCP connection to stream the video, while the receiver sets the receiver window size to 8, targeting at 800 kbps throughput, (b) the sender and the receiver use our MultiTCP system to stream the video with the number TCP connections limited to two, and (c) the sender and the receiver also use our proposed MultiTCP system, except the number of TCP connections are now set to five. Table I shows the parameters used in our MultiTCP system.

Figure 4(a) shows the throughput of three described scenarios. As seen, initially without congestion, using the traditional single TCP connection can control the throughput very well since setting the size of the receiver window to 8 achieves

Sampling interval δ	300 ms
Throughput smoothing factor α	0.9
Guarding threshold λ	7000 bytes
Throughput spike factor f	6

Table I. *Parameters used in MultiTCP system*

the desired throughput. However, when traffic bursts occur during the intervals $t = 200s$ to $t = 220s$ and $t = 300s$ to $t = 340s$, the throughput of using a single TCP connection reduces substantially to only about 600 kbps. For the same congested period, using two TCP connections results in higher throughput, approximately 730 kbps. On the other hand, using five TCP connections produces approximately the desired throughput, demonstrating that a larger number of TCP connections results in higher throughput resilience in the presence of misbehaved traffic such as UDP flows. These results agree with the analysis in Section 4.1. It is interesting to note that when using two TCP connections, there are spikes in the throughput immediately after the network is no longer congested at t = 221s and t = 341s. This phenomenon relates to the maximum receiver window size set during the congestion period. Recall that the algorithm keeps increasing the w_i until either (a) the measured throughput exceeds the desired throughput or (b) the sum of receiver window size $W_s = \sum_i w_i$ reaches $f\frac{T_d RTT}{MTU}$. In the simulation, using two TCP connections never achieves the desired throughput during the congested periods, hence the algorithm keeps increasing the w_i. When network is no longer congested, the W_s already accumulates to a large value. This causes the sender to send a large amount of data until the receiver reduces the window size to the correct value a few RTTs later. On the other hand, when using 5 TCP connections, the algorithm achieves the desired throughput during the congestion periods, as such W_s does not increase to a large value, resulting in a smaller throughput spike after the congestion vanishes. Next, when 15 cross traffic TCP connections start at t = 500s, the resulting throughput when using one and two TCP connections reduce to 700 kbps and 350 kbps, respectively. However, throughput when using 5 TCP connections stays approximately constant at 800 kbps. At t = 750s, 5 additional TCP connections start, throughput are further reduced for the one and two connection cases, but it remains constant for the five-connection case.

Figure 4(b) shows the average of the sum of window size W_s as a function of time. As seen, W_s increases and decreases appropriately to respond to network conditions. Note that using two connections, W_s increases to a larger value than when using 5 TCP connections during the intervals of active UDP traffic. This results in throughput spikes discussed earlier. Also, the average window size in the interval t = 500s to t = 750s is smaller than that of the interval t = 750s to t = 1000s, indicating that the algorithm responds appropriately by increasing the window size under a heavier load.

We now show the results when the cross traffic has different round trip time from that of the video traffic. In particular, the propagation delay between the router and traffic generator node is now set to 40 milliseconds. All cross traffic patterns stay the same as before. The new simulation shows the same basic results. As seen in Figure 5.1.1, the throughput of 5 connections is still higher than that of

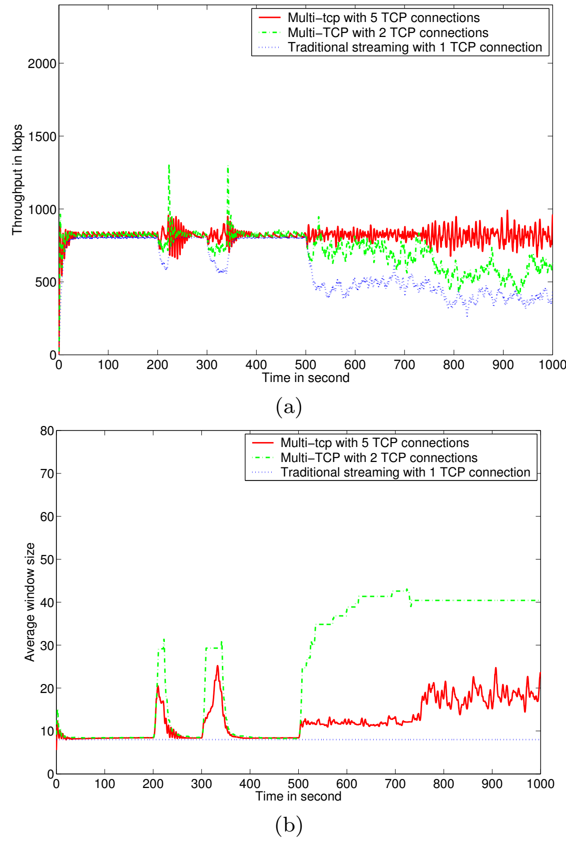

(b)

Fig. 4. (a) Resulted throughput and (b) average receiver window size when using 1, 2 and 5 TCP connections.

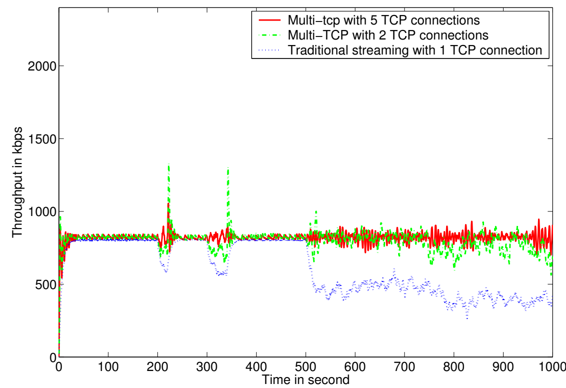

Fig. 5. Resulted throughput when using 1, 2 and 5 TCP connections with cross traffic having larger RTT than that of video traffic.

Sampling interval δ	6*RTT ms
Throughput smoothing factor α	0.9
Guarding threshold λ	3000 bytes
Throughput spike factor f	10

Table II. *Parameters used in MultiTCP system*

two connections which, in turn is higher than that of one connection during the congestion periods. The throughput of two connections in this new scenario is slightly higher than that of the previous scenario during the congested period from $t = 500$s onward. This is due to a well known phenomena that TCP connection with shorter round trip times gets a larger share of bandwidth for the same loss rate. Since the round trip time of the video traffic is now shorter than that of the TCP cross traffic, using only two connections, the desired throughput of 800 kbps can be approximately achieved during the period from $t = 500$s to $t = 750$s, which is not achievable in previous scenario. So clearly, the number of connections to achieve the desired throughput depends on the competing traffic.

5.1.2 *Internet Experiments.* We have implemented our MultiTCP system on Redhat Linux with the options for using either fixed or varied number of TCP connections. Our MultiTCP system is mostly transparent to the application layer. The applications only need to specify the required bit rate. Every calls to the current network library are intercepted and the equivalent MultiTCP functions are invoked. Therefore, with minimal effort, any network program can be changed to employ the MultiTCP system. We now show the results from actual Internet experiments using Planet-Lab nodes [PlanetLab]. In this experiment our setup consists of a sender, planetlab1.netlab.uky.edu at University of Kentucky, and a receiver, planetlab2.een.orst.edu at Oregon State University. The sender streams the video to the receiver continuously for a duration of nearly 1000s, while the FTP connections generate cross traffic at the client at different times. From time $t = 0$ to $t = 400$s, there is no cross traffic. From around $t = 401$s to $t = 600$s cross traffic is generated using two FTP connections and after $t = 601$s two additional FTP connections are opened for generating more traffic at the receiver. We now consider performances under three different scenarios: (a) only one TCP connection is used to stream the video, while the receiver sets the receiver window size targeting at 480 kbps throughput, and (b) the sender and the receiver use the proposed MultiTCP system, using the $AdjustWindowSize$ algorithm with the number of TCP connections set to four (c) the sender and the receiver use the proposed MultiTCP system, using the $AdjustWindowSize$ algorithm with the number of TCP connections now set to eight. Table II shows the parameters used in our MultiTCP system. Figure 6(a) shows the throughput of three described scenarios. As seen, initially without congestion, using the traditional single TCP connection can control the throughput very well since setting the size of the receiver window achieves the desired throughput. So using either one, four or eight connections does not make any difference during the interval $t = 0s$ to $t = 400s$ as there is no cross traffic. However, when traffic bursts occur during the interval $t = 401s$ to $t = 600s$ and after $t = 601s$, the throughput of using a single TCP connection

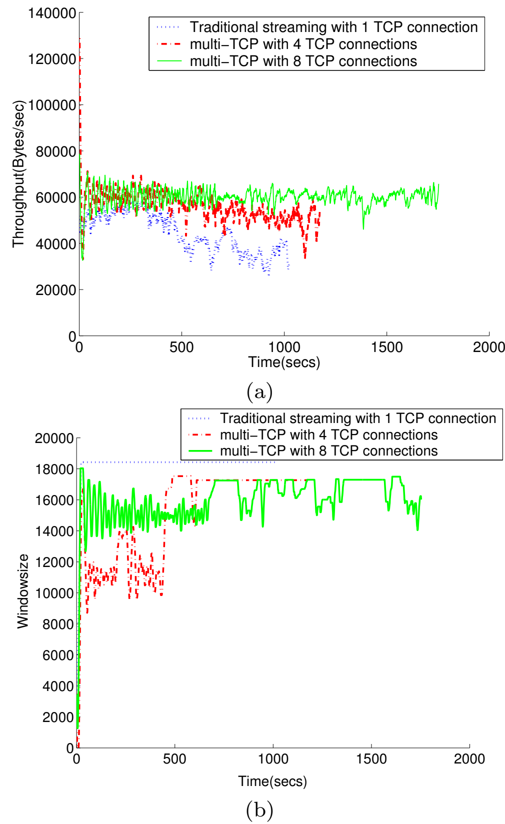

Fig. 6. (a) Resulted throughput, (b) receiver window size when using 1, 4 and 8 TCP connections with cross traffic.

reduces substantially from 480 kbps to 280 kbps. Using four TCP connections we were able to achieve the desired throughput during the interval $t = 401s$ to $t = 600s$. But after $t = 601$, the throughput reduced from 480 kbps to 400kbps. For the same congested period, using eight TCP connections results in higher throughput, which is the desired throughput.

Figure 6(b) shows the average of the sum of window size W_s as a function of time. As seen, W_s increases and decreases appropriately to respond to network conditions. Note that using single connection, W_s increases to a larger value than when using multiple TCP connections during the intervals of active FTP traffic. This results in throughput spikes discussed earlier. Also, the average window size in the interval $t = 300$s to $t = 600$s is smaller than that of the interval $t = 601$s to $t = 1000$s for multiple TCP connections, indicating that the algorithm responds appropriately by increasing the window size under a heavier load.

5.2 Results for the $VariedConnVariedWinSize$ Algorithm

Sampling interval δ	6*RTT ms
Throughput smoothing factor α	0 - 0.9
Guarding threshold λ	3000 bytes
Connection damping guarding threshold β	1
Congestion factor cf	5

Table III. *Parameters used in MultiTCP system*

5.2.1 *Internet Experiments.* We now show the results of the $VariedConnVariedWinSize$ algorithm. For this experiment our setup consists of a sender, a receiver, and fifteen FTP connections to generate cross traffic. The machine used as server is planetlab1.csres.utexas.edu which is located at University of Texas. The setup under two different scenarios is as follows: (a) the sender uses only one TCP connection to stream the video, while the receiver sets the receiver window size targeting at 800 kbps throughput, and (b) the sender and receiver use our MultiTCP system to stream the video. Table III shows the parameters used in our MultiTCP system.

Figure 7(a) shows the throughput of two described scenarios with cross traffic. During the interval $t = 0s$ to $t = 400s$ there is no cross traffic. However, when traffic bursts occur due to six FTP connections during the interval $t = 401s$ to $t = 600s$ and nine more FTP's after $t = 601s$, the throughput of using a single TCP connection reduces from 800 kbps to 720 kbps. For the same congested period, MultiTCP system is able to achieve the desired throughput. These results demonstrate that a larger number of TCP connections results in higher throughput resilience in the presence of misbehaved traffic. Figure 7(b) shows the average of the sum of window size W_s as a function of time. As explained before, W_s increases and decreases appropriately to respond to network conditions.

Figure 7(c) shows how the number of connections vary depending upon the congestion ratio for that particular scenario. When the window size increased the number of connections used for streaming also increased. This implies that when the congestion in the network increases and if we are not able to receive the desired throughput, the number of connections to be used for streaming are increasing. These results agree with the analysis in Section 4.1.

Now let us consider another experiment involving higher target bandwidth. For this experiment our setup consists of a sender, a receiver, and eight FTP connections to generate cross traffic. The machine used as server is planetlab1.cs.pitt.edu which is located at University of Pittsburgh. The setup under two different scenarios is as follows: (a) the sender uses only one TCP connection to stream the video, while the receiver sets the receiver window size targeting at 1.6 Mbps throughput, and (b) the sender and receiver use our MultiTCP system to stream the video. We used the same parameters as in previous experiment. Figure 8(a) shows the throughput of two described scenarios with cross traffic. During the interval $t = 0s$ to $t = 300s$ there is no cross traffic. Even though there is generated cross traffic during that period, the client was able to receive only 480 kbps [1]. On the other hand, when MultiTCP system is employed, the throughput increased to 1.6 Mbps and remain

[1] This is due to both the current actual traffic and the imposed bandwidth limit of the server per TCP connection

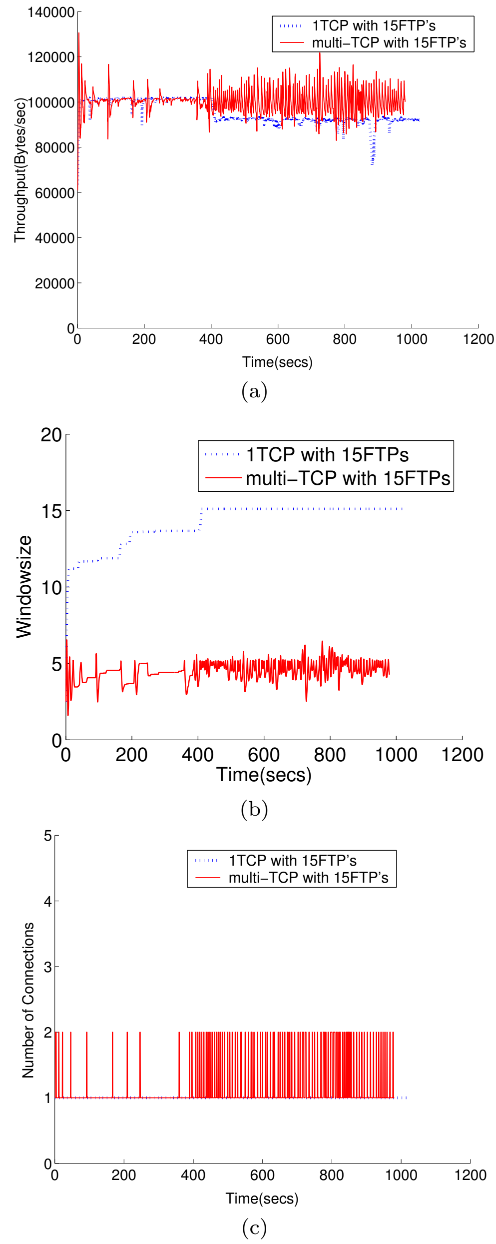

Fig. 7. (a) Throughputs and (b) variation of the receiver window size for the traditional and the MultiTCP streaming systems; (c) Number of connections used for streaming in MultiTCP system.

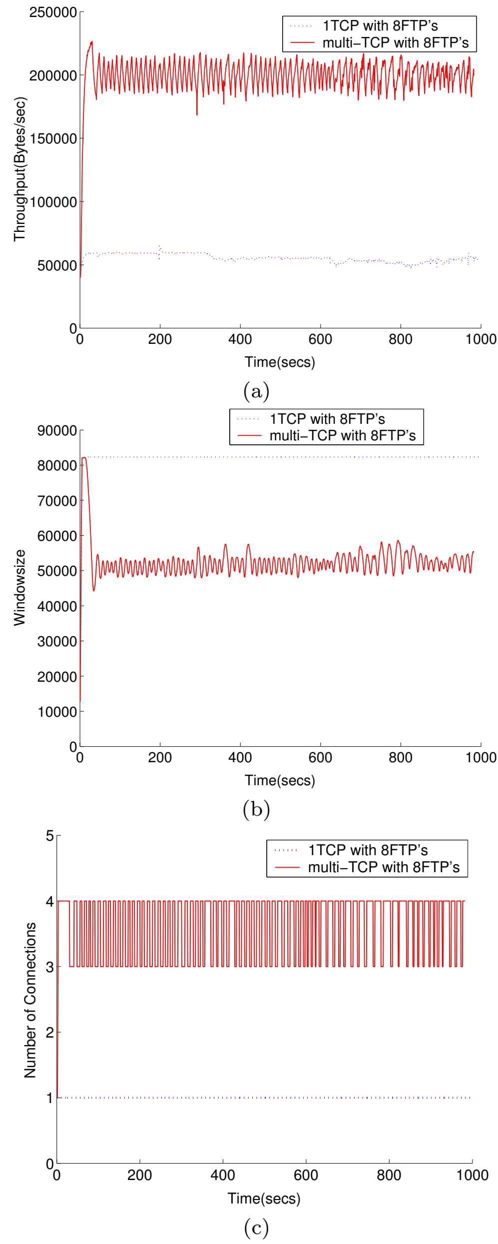

Fig. 8. (a) Throughputs and (b) variation of the receiver window size for the traditional and the MultiTCP streaming systems; (c) Number of connections used for streaming in MultiTCP system.

at this value. We then simulate heavy traffic by creating four FTP connections during the interval $t = 301s$ to $t = 600s$ and four more FTP's after $t = 601s$. During this time, the throughput of using a single TCP connection reduced further from 480 kbps to 400 kbps. For the same congested periods, the MultiTCP system still maintains the desired throughput at 1.6 Mbps.

Figure 8(b) shows the average of the sum of window size W_s as a function of time. As explained before, W_s increases and decreases appropriately to respond to network conditions. But initially as it has to get to a high throughput the window size went up to maximum and then got adjusted slowly. Figure 8(c) shows how the number of connections vary depending upon the congestion in the network.

These results also agree with the analysis in Section 4.1. Recall that the algorithm keeps increasing the w_i until either (a) the measured throughput exceeds the desired throughput or (b) the sum of receiver window size $W_s = \sum_i w_i$ reaches $f\frac{T_d RTT}{MTU}$. In the results we have shown, using single TCP connection never achieves the desired throughput during the congested periods, hence the algorithm keeps increasing the w_i and reaches maximum. When network is no longer congested, the W_s already accumulates to a large value. This causes the sender to send a large amount of data until the receiver reduces the window size to the correct value a few RTTs later. On the other hand, when using eight or five TCP connections, the algorithm achieves the desired throughput during the congestion periods, as such W_s does not increase to a large value, which results in a smaller throughput spike after the congestion vanishes.

5.3 Stream Buffer

In this section we present the performance of our MultiTCP system when prebuffering is used. Prebuffering is used to overcome temporary insufficient network bandwidth during the streaming session by allowing the video player to accumulate a large enough amount of data in a buffer before playing back. A larger buffer results in fewer number of stops during the video streaming session. At one extremity, if the amount of data in the buffer equals to the size of the entire video, then will be no stop during playback, but the user has to wait for the entire video to be downloaded.

We perform the following experiment. A 720kbps video is streamed for a duration of 200 seconds with a initial buffer of 10 seconds. At the same time, a random number of FTP connections sharing the same bandwidth (from 0 to 8) also start and stop at random times to emulate different levels of congestion. Once the video starts playing, it will only stop when there is no more data in the buffer. This situation happens when the playback rate is faster than the receiving rate for some period of time. Once the video player stops, it would stop for 5 seconds to accumulate data before playing back again. Using this scheme, Figure 9 shows the number of stops when using our $AdjustWindowSize$ algorithms with different number of connections. As seen, there are 12 stops when using single TCP connection, 1 stop for each of the 2, 3, 4 TCP connections and no stops when 5 TCP connections are used. On the other hand, when the $VariedConnectionVariedWinSize$ algorithm is used, there is no stops since the algorithm can achieve the desired throughput automatically by increasing and decreasing the number of connections.

These results demonstrate that our algorithm is able to achieve the desired

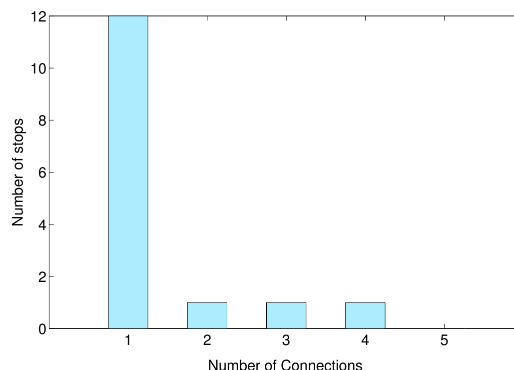

Fig. 9. Resulted number of stops when using 1, 2, 3, 4 and 5 TCP connections with cross traffic having same RTT as that of video traffic.

throughput and maintain precise rate control under a variety congested scenarios with competing UDP and TCP traffic. We should emphasize again that, the applications based on our system indeed obtain a larger share of the fair bandwidth. However, we believe that under limited network resources, time-sensitive applications like multimedia streaming should be treated preferentially as long as the performance of all other applications do not degrade significantly. Since our system uses TCP, congestion collapse is not likely to happen as in the case of using UDP when network is highly congested. In fact, DiffServ architecture uses the same principle by providing preferential treatment to high priority packets.

6. CONCLUSIONS

We conclude our paper with a summary of contributions. First, we propose and implement a receiver-driven, TCP-based system MultiTCP for multimedia streaming over the Internet using multiple TCP connections for the same applications. Second, our proposed system is able to provide resilience against short-term insufficient bandwidth due to traffic bursts. Third, our proposed system enables the application to control the sending rate in a congested scenario, which cannot be achieved using traditional TCP. Finally, our proposed system is implemented at the application layer, and hence, no kernel modification to TCP is necessary. The simulation and Internet results demonstrate that using our proposed system, the application can achieve the desired throughput in many scenarios, which cannot be achieved by traditional single TCP approach.

REFERENCES

Apostolopoulos, J. 2001. Reliable video communication over lossy packet networks using multiple state encoding and path diversity. In *Proceeding of The International Society for Optical Engineering (SPIE)*. Vol. 4310. 392–409.

Apostolopoulos, J. 2002. On multiple description streaming with content delivery networks. In *InfoComm*. Vol. 4310.

Blake, S., Black, D., Carson, M., Davis, E., Wang, Z., and Weiss, W. 1998. An architecture for differentiated services. In *RFC2475*.

Chen, M. and Zakhor, A. 2004. Rate control for streaming over wireless. In *INFOCOM*.

CHEN, M. AND ZAKHOR, A. 2005. Rate control for streaming video over wireless. *IEEE Wireless Communications 12,* 4 (August).

CHEN, M. AND ZAKHOR, A. 2006. Flow control over wireless network and application layer implementation. In *INFOCOM.*

CROWCROFT, J. AND P.OESCHLIN. 1998. Differentiated end-to-end internet services using weighted proportional fair sharing tcp.

DONG, Y., ROHIT, R., AND ZHANG, Z. 2002. A practical technique for supporting controlled quality assurance in video streaming across the internet. In *Packet Video.*

FLOYD, S. AND FALL, K. 1999. Promoting the use of end-to-end congestion control in the internet. *IEEE/ACM Transactions on Networking.*

FLOYD, S., HANDLEY, M., PADHYE, J., AND WIDMER, J. 2000. Equation-based congestion control for unicast application. In *Architectures and Protocols for Computer Communication.* 43–56.

Information Sciences Institute. *Network simulator.* Information Sciences Institute, http://www.isi.edu/nsnam/ns.

Internet Engineering Task Force 2000. *Stream Control Transmission Protocol.* Internet Engineering Task Force, RFC 1771.

KUROSE, J. AND ROSS, K. 2005. *Computer networking, a top down approach featuering the Internet, third edition.* Addison Wesley.

LEIGH, J., ANDD. SCHONFELD, O. Y., AND ANSARI, R. 2001. Adaptive networking for tele-immersion. In *Immersive Projection Techonology/Eurographics Virtual Environments Workshop(IPT/EGVE).*

LIANG, Y., SETTON, E., AND GIROD, B. 2002. Channel adaptive video streaming using packet path diversity and rate-distortion optimized reference picture selection. In *IEEE Fifth Workshop on Multimedia Signal Processing.*

MA, H. AND ZARKI, M. E. 1998. Broadcast/multicast mpeg-2 video over wireless channels using header redundancy fec strategies. In *Proceedings of The International Society for Optical Engineering (SPIE).* Vol. 3528. 69–80.

MEHRA, P. AND ZAKHOR, A. 2003. Receiver-driven bandwidth sharing for tcp. In *INFOCOM.* San Francisco.

MovieFlix. MovieFlix, http://www.movieflix.com/.

NGUYEN, T. AND CHEUNG, S. 2005. Multimedia streaming using mutiple tcp connections. In *IPCCC.*

NGUYEN, T. AND ZAKHOR, A. 2004. Multiple sender distributed video streaming. *IEEETransactions on Multimedia and Networking 6,* 2 (April), 315–326.

PlanetLab. PlanetLab, http://www.planet-lab.org.

P.MEHRA, C. V. AND A.ZAKHOR. 2005. Receiver-driven bandwidth sharing for tcp and its application to video streaming. *IEEE Transactions on Multimedia 7,* 4 (August).

REIBMAN, A. 2002. Optimizing multiple description video coders in a packet loss environment. In *Packet Video Workshop.*

REYES, G. D. L., REIBMAN, A., CHANG, S., AND CHUANG, J. 2000. Error-resilient transcoding for video over wireless channels. *IEEE Transactions on Multimedia 18*, 1063–1074.

SEMKE, J., MAHDAVI, J., AND MATHIS, M. 1998. Automatic tcp buffer tuning. In *SIGCOMM.*

TAN, W. AND ZAKHOR, A. 1999. Real-time internet video using error resilient scalable compression and tcp-friendly transport protocol. *IEEE Transactions on Multimedia 1*, 172–186.

WANG, Z. 2001. *Internet QoS, Architecture and Mechanism for Quality of Service.* Morgan Kaufmann Publishers.

WHITE, P. 1997. Rsvp and integrated services in the internet: A tutorial. *IEEE Communication Magazine*, 100–106.