

# Measurement Based Performance Analysis of Internet over Satellite

Navid Ehsan, Mingyan Liu, Rod Ragland

## Abstract

Performance enhancing proxies (PEPs) are widely used as a solution to improve the performance of TCP over high delay-bandwidth product links and links with high error probability. In this paper we analyze the performance of using TCP connection splitting in combination with web caches via traces obtained from a commercial satellite system. We examine the resulting performance gain under different scenarios, including the effect of caching, congestion, random loss and file sizes. We show, by analyzing our measurements, that the performance gain from using splitting is highly sensitive to random losses (for small files) and the number of simultaneous connections, and that such sensitivity is alleviated by caching. On the other hand, the use of a splitting proxy enhances the value of web caching in that cache hits result in much more significant performance improvement over cache misses when TCP splitting is used. We also compared the performance of using different versions of HTTP over such a system.

**keywords:** performance enhancing proxy, connection splitting, cache, TCP performance, satellite network.

## I. INTRODUCTION

The performance of TCP over heterogeneous connections such as those including satellite and wireless links has been extensively studied for the past few years. Proposed performance enhancing techniques can roughly be categorized into link layer solutions (see for example [1], [2]), end-to-end solutions where the end-to-end semantic of TCP is maintained (see for example [3], [4], [5], [6], [7], [8], [9], [10]) and non end-to-end solutions where the end-to-end semantic is violated (see for example [11], [12], [13]). Various link layer and end-to-end approaches can be quite effective for connections over wireless links through improved error correction, local retransmission and distinguishing congestion losses from link failure losses for TCP. In a connection that incorporates a satellite link on the other hand, the main bottleneck in TCP performance is due to the large delay-bandwidth product nature of the satellite link

---

N. Ehsan and M. Liu are with the Electrical Engineering and Computer Science Department, University of Michigan, Ann Arbor. E-mail: {nehsan, mingyan}@eecs.umich.edu. R. Ragland is with Hughes Network Systems, Inc. This work is partially supported by Hughes Network Systems.

(for an overview of research on TCP over satellite see for example [14]). Over such a link the normal TCP window dynamics results in significant latency before the channel is fully utilized. This problem cannot be effectively solved simply by improving the satellite channel quality, or by using large initial window size. This is because a connection using the satellite link typically also has a terrestrial part, thus using large window end-to-end could affect the performance and fairness of the terrestrial part of the connection.

One typical non end-to-end solution that has been adopted by many satellite data communication service providers is the TCP connection splitting technique. The idea behind this technique is to segregate the end-to-end connection into segments so that each can be optimized separately, and in particular so that the TCP window over the satellite segment can be opened up faster. This involves placing at the Network Operating Center (NOC) a splitting proxy that acknowledges end user packets on behalf of the remote server and acknowledges the remote server on behalf of the end user (assuming the end user is directly connected to the satellite down link). This proxy operates at the TCP level and therefore breaks the end-to-end semantic of TCP. It is also not compatible with IPSec [15]. The benefit however, is that (1) splitting a connection results in shortened TCP feedback control loops, so that the remote server receives ACKs much sooner from the NOC (proxy) than from the end user, and therefore its window can quickly ramp up; and (2) by segregating an end-to-end connection into segments, each segment is enhanced separately. The proxy can use large initial window over the satellite link, which can either be the last hop or somewhere in the middle of the end-to-end connection.

The performance gain of using a splitting proxy has been reported in [16] as a result of simulation study and in [11] from experimental study. In [17] and [18] some analysis was presented to quantitatively study the benefit of such a proxy.

We observe that in a typical satellite system TCP connection splitting is not the only performance enhancing technique that is commonly used. Web caching is also widely implemented to reduce latency by pushing contents closer to the end users. Moreover, whenever there is a cache miss, the cache in effect “breaks” the server-client transfer into two separate connections. This is because the cache opens up a connection to the remote server and starts downloading the file to the cache (for cacheable objects) while forwarding packets to the client at the same time. Although this happens at the application layer, it results in a server-cache connection and a cache-client connection, and thus has a very similar effect on the end-to-end transport layer performance as that of a TCP splitting proxy as will be shown more clearly in the next section. Our focus in this study is the the combined performance implication of using both

types of enhancements - TCP splitting and web caching. By taking measurements from a live Internet over satellite system, we hope to gain more insights into the use of proxy as a solution to such systems, and more importantly to apply such understanding to system level design issues. Our measurements are obtained via repeated downloads of selected files (either pure text or with embedded objects).

Our main results are as follows:

(1) Connection splitting enhances the value of caching. When splitting is not used whether there is a cache hit or cache miss generates almost identical performance (in terms of delay and throughput). When splitting is used, a hit at the cache results in much higher throughput.

(2) The throughput of a split connection is more sensitive to increased congestion delay than an end-to-end connection. Having simultaneous connections will cause the performance of splitting to decrease much faster than the end-to-end case.

(3) The throughput of a split connection is highly sensitive to packet losses for small file transfers. Split connection provides much better *guarantee of performance improvement* for a large file than for a small file. By guarantee of performance improvement we mean the probability that the split connection results in higher throughput (or lower delay) than the end-to-end connection. This probability is defined as the relative frequency of such events in our measurements. This guarantee is improved when there is a cache hit.

(4) The performance gain of using connection splitting is reduced as the number of embedded objects in a file increases. In addition, connection splitting is no substitute for persistent connection. If a splitting proxy is used, it is important that persistent connection is also used between the client and the proxy. Non-persistent connection even when connection splitting is used can result in much worse performance than an end-to-end persistent connection.

It's worth pointing out that TCP connection splitting is also often called TCP spoofing. Strictly speaking splitting refers to breaking up a connection and spoofing refers to faking an address. They are often related because in splitting a connection a transparent proxy typically spoofs the end points' addresses. Since our interest is in the end-to-end performance as a result of split connections (either at the transport layer or at the application layer), we will limit ourselves to the term connection splitting in this paper.

The rest of the paper is organized as follows. We describe the system configuration and our experiment methodology in the next section. We then present the measurements and our observation/explanation in Section 3. These results are organized into four subsections, where we examine the effect of file size, number of simultaneous connections, congestion and random losses, and number of embedded objects,

respectively. Conclusions are given in Section 4.

## II. SYSTEM DESCRIPTION

Our measurements are taken from a commercial satellite system that uses a geo-stationary (GEO) satellite (Ku band) for forward data transfer (from the Network Operating Center (NOC) to the client/end host) and a regular phone line as the return channel (from the client/end host to the NOC via an ISP), as shown in Figure 1. Available bandwidth on the forward channel is up to 24 Mbps and the one way propagation delay of the satellite link is roughly 250 ms (however due to NOC configuration the maximum throughput we were ever able to achieve per client was 300-400Kbytes/sec). The return link has 4 KHz bandwidth.




Fig. 1. Main components of the satellite system

The TCP connection splitting proxy (which we will simply call proxy in the following) is implemented on a Hybrid Gateway (HGW) located in the NOC. The end host can choose to either enable or disable the proxy. When the proxy is disabled, a packet from the server to the end host passes through the HGW as if passing through a normal router, and goes directly to a Satellite Gateway (SGW) connected to the satellite uplink. When the proxy is enabled, it breaks up the server-client end-to-end connection in two, and pre-acknowledges one on behalf of the other in the following way, as illustrated in Figure 2(a).

During the connection establishment, the proxy simply forwards the SYN and SYNACK packets. Then the server starts sending data to the client upon receipt of the file request. Upon receiving a data packet from the server, the proxy sends an ACK back to the server. Since the proxy is transparent by spoofing the client's address, the server takes the ACK as an indication that the data being ACKed has been received

successfully by the client, and therefore moves on to the next window and so on. Since the proxy is located closer (than the client) to the server, this results in a much shorter round-trip time (RTT) seen by the server and thus enables the server to reach a much higher send rate. At the same time the proxy is maintaining a separate connection with the client by forwarding data packets to the client, waiting for ACKs to come back from the client, and then releasing more data packets to the client. All packets received from the server are stored at the proxy. When an ACK is received from the client, data being ACKed is purged from the unacknowledged buffer. Otherwise the proxy retransmits (via duplicate ACKs or timeouts) from its retransmission buffer (local retransmission) in contrast with the end-to-end case where each retransmission comes from the remote server.

The web caches are also located in the NOC. Regardless of whether the splitting proxy is enabled or disabled, when an HTTP request is received by the proxy, it first goes through the cache before being sent to the server. For cacheable content, if a fresh copy of the requested file is located in the cache (a *hit*), the file is delivered to the client directly from the cache without going to the remote server. If the requested file is not found in the cache (a *miss*), the cache will open up a connection to the remote server to fetch the file, as shown in Figure 2(b). This server-cache connection is concurrent with the cache-client connection in that as soon as the cache starts receiving data from the server (via the server-cache connection), it will transfer it to the client (via the cache-client connection). Thus in the case of a miss, the cache effectively handles two connections that constitute the end-to-end connection between the server and the client. In terms of data transfer, this is very similar to a splitting proxy. (However, with a cache this takes place at the application layer so the server sees a connection with the cache rather than being pre-acknowledged by the cache.) Figure 2 compares the packet flow in the case of a splitting proxy and in the case of a cache miss. Except for the connection establishment process, the data transfer essentially proceeds in an identical manner (note that this figure does not show processing delay). Consequently the splitting proxy together with the cache results in an end-to-end connection split twice upon a cache miss and once upon a cache hit, as shown in Figures 3(a) and 3(b), respectively. Figures 3(c) and 3(d) illustrates cache hit and cache miss, respectively, when the proxy is disabled. Figures 3(e) and 3(f) illustrates the cases where connections bypass the cache with the splitting proxy enabled and disabled, respectively.

Important parameters of our system are as follows. The client is running Windows ME that uses TCP SACK [19] with delayed acknowledgements (one ACK for every two received packets). Our web server is running Linux Redhat 7.1. Because of the high asymmetry in link speeds between the forward and return paths, we also use ACK filtering [20] at the client and send one out of every four ACKs. Thus each

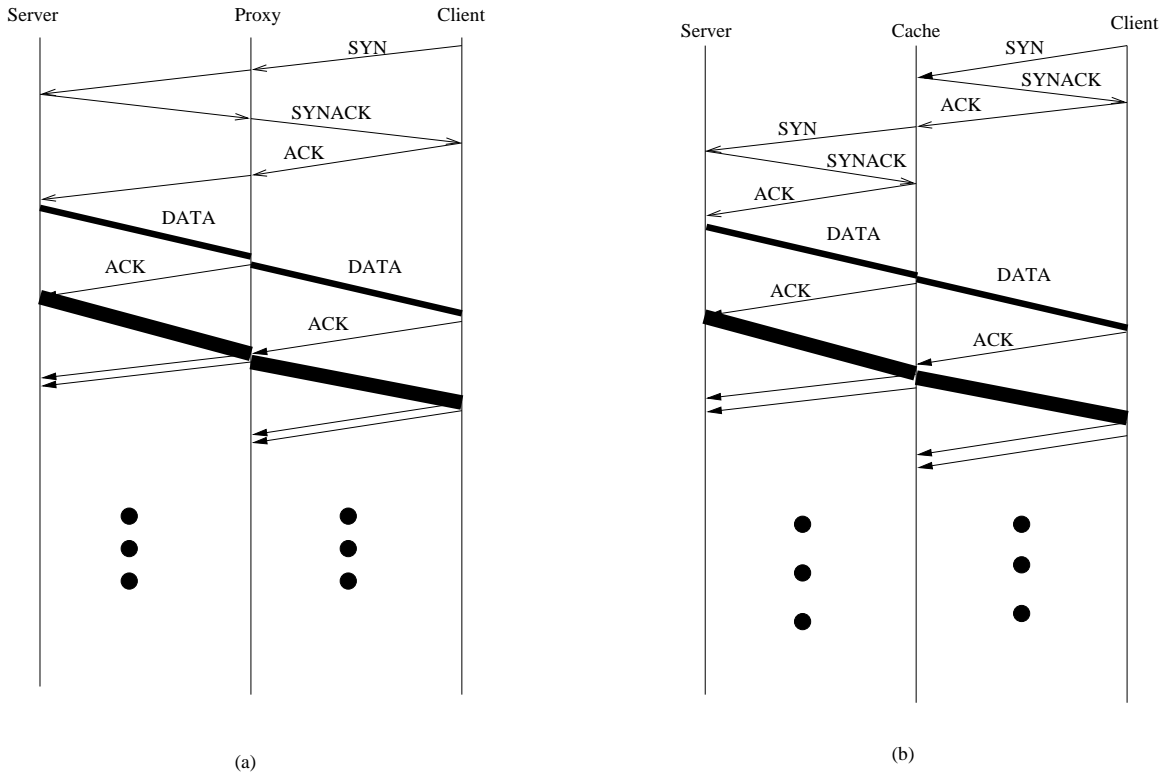


Fig. 2. Splitting a connection at (a) a proxy and (b) a cache (miss)

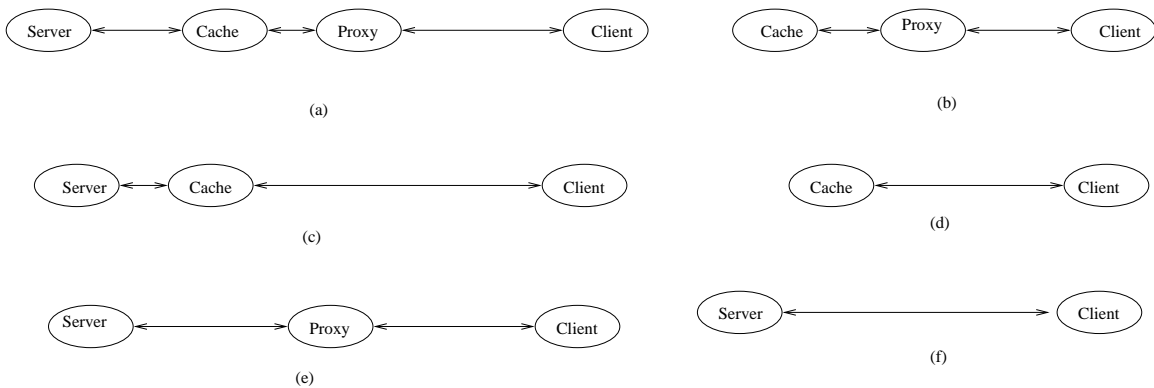


Fig. 3. Experiment scenarios (a) Cache miss, proxy enabled; (b) Cache hit, proxy enabled; (c) Cache miss, proxy disabled; (d) Cache hit, proxy disabled; (e) No cache, proxy enabled; (f) No cache, proxy disabled

ACK received at the server (assuming no loss) in general represents eight packets. Byte counting instead of ACK counting is used at the server as well as the proxy, so that the lower rate of ACKs does not reduce the data sending rate. The splitting proxy uses an initial window size of 64 Kbytes for the proxy-client connection over the satellite whenever enabled.

Our study consists of six main scenarios: splitting enabled or disabled with cache hit or cache miss and the option of whether to bypass the cache or not, as shown in Figure 3. Whether to use the file in the cache or not is controlled by a *no-cache pragma* [21], [22] sent in the request header. When set, this parameter tells the cache not to use the cached copy even if there is one and to get the most updated version from the server. Whether the connection splitting proxy used or not is controlled by the end hosts. We have two end hosts, one of which has the proxy option enabled and the other one has the option disabled. For comparison purposes, we always run experiments on these two hosts simultaneously. We download files from a dedicated web server onto both hosts repeatedly for durations of 1-2 hours per measurement point (measurements over this period are averaged into one data point).

All measurements are taken from a live commercial system with varying amounts of customer traffic. Our connections and experiment traffic go through a dedicated HGW and therefore a dedicated splitting proxy that is not shared by other traffic. However our connections does share the cache access, the Satellite Gateway and the satellite transponder with through traffic. Such a setup results in both controllable and uncontrollable factors as we will point out when we discuss our results in the next section.

The performance metrics we use in this study are the file transfer latency (using HTTP) and throughput. We define latency seen by the client as the time between when the SYN request is sent and the time when FINACK is sent by the client. For files with multiple embedded objects, this is defined as the time between when the first SYN request is sent and the time when the last FINACK is sent. Throughput is defined as file size divided by latency. Files used in this study are listed in Table I. Files 15-18 contain equal-sized objects and the total size of each (base page plus all the embedded objects) is approximately 64 KBytes.

### III. RESULTS AND ANALYSIS

#### A. Effect of Varying File Sizes

In this subsection we compare the performance under scenarios described in Figures 3(a)-3(d). We download files onto both the splitting enabled and the splitting disabled hosts repeatedly over a one-hour period and measured their throughput. We define the *Gain Of Splitting* (GoS) as

$$\text{GoS} = \frac{\text{Throughput}_{\text{splitting}} - \text{Throughput}_{\text{end-to-end}}}{\text{Throughput}_{\text{end-to-end}}}$$

TABLE I  
FILES USED FOR MEASUREMENTS

| File No. | Type      | Size (Bytes) | File No. | Type                | Size (Bytes) |
|----------|-----------|--------------|----------|---------------------|--------------|
| 1        | full text | 2459         | 10       | full text           | 77030        |
| 2        | full text | 11033        | 11       | full text           | 93141        |
| 3        | full text | 16565        | 12       | full text           | 120579       |
| 4        | full text | 21027        | 13       | full text           | 152740       |
| 5        | full text | 28927        | 14       | full text           | 167914       |
| 6        | full text | 37155        | 15       | 3 embedded objects  | 64K          |
| 7        | full text | 42513        | 16       | 7 embedded objects  | 64K          |
| 8        | full text | 53069        | 17       | 18 embedded objects | 64K          |
| 9        | full text | 65331        | 18       | 19 embedded objects | 64K          |

Thus larger GoS means higher throughput gain from using the splitting proxy. A negative GoS means that the end-to-end connection has a higher throughput (or smaller latency) than the split connection. Figure 4 shows the GoS of Files 2,3,4,9 and 12 in the cache hit and the cache miss cases. Each point is an average over the one-hour period.




Fig. 4. Comparing GoS for two cases: cache hit and cache miss

This result confirms some earlier reports, see for example [16], [18], that splitting provides better performance gain for larger files. This is due to the fact that the time spent in handshake has a bigger portion in the overall latency for smaller files and this time is not reducible by using splitting. At the same time larger files benefit more from the large initial window size (for example a 1kbyte file and a 64kbyte



file can both be transmitted in the first window and experience almost the same latency, if file transfer from source to the proxy is fast enough).

One additional interesting observation from this comparison is that when a file is in the cache, the splitting gain is much higher. This can be more clearly seen in Figures 5 (a) and (b), where we compare the file transfer throughput separately for the cache hit and cache miss cases.

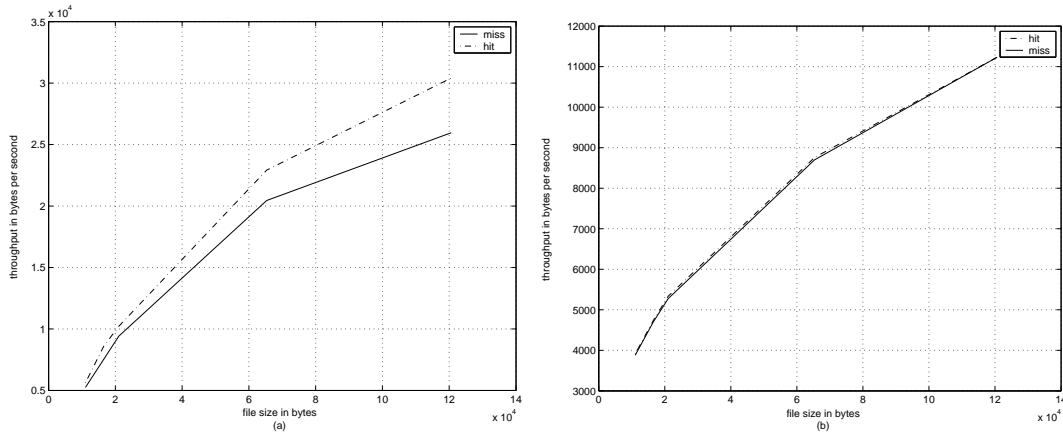


Fig. 5. Comparing throughput for cache hit and cache miss with (a) splitting enabled and (b) splitting disabled

It becomes obvious, by comparing the two, that the use of splitting at the hybrid gateway enhances the value of caching, i.e., when splitting is used, having the file in the cache provides significant increase in throughput over the case where the file has to be retrieved remotely from the server. In addition, this improvement increases as the file size increases. On the other hand, when splitting is disabled, whether the file is in the cache or not makes little difference.

The reason lies in the following. Consider the case where the connection is not split by the proxy. Assuming there is a cache miss (Figure 3(c)), since the cache-client connection is much slower (as a result of higher propagation delay) than the server-cache connection, by the time the cache-client connection gets to the first few windows, the entire file could be available in the cache (i.e. the server-cache connection is completed). As an example consider 1Kbyte packet size and a 7Kbyte file. Assume that the round trip time of the server-cache connection is 50msec and the round trip time of the cache-client connection to be 500msec. As the first packet arrives at the cache it is immediately sent to the client. It takes 2 more round trip times for the whole file to be available in the cache (100 msec), but by this time the first packet has not even reached the client yet. By the time the first acknowledgment gets to the cache from client, the file will be completely available in the cache, so a cache hit and a cache miss have about the same latency and throughput.

Therefore having the file locally in the cache provides very limited benefit. Intuitively when an end-to-end connection is split in two, the slower segment (as a result of higher propagation delay, smaller initial window size, or higher losses) will dominate the end-to-end performance. In this case the cache-client segment is much slower than the server-cache segment, and clearly dominates the end-to-end performance. Having the file locally in the cache has the effect of “speeding up” the server-cache connection, i.e. this connection is completely eliminated. However since the overall performance is governed by the cache-client connection, whether the server-cache connection is a bit faster or not does not matter much, as shown in Figure 5(b).

Now consider the case where the connection splitting proxy is enabled. Splitting the connection at the gateway results in either three or two segments of an end-to-end connection (Figures 3(a) and 3(b)). As we have just discussed, if the proxy only splits the connection, then the server-cache connection and the cache-proxy connection would still be much faster than the satellite link and therefore the proxy-client connection would again dominate the overall performance. However, in addition to splitting the proxy also opens up the window size over the satellite link much faster by using an initial window size of 64 KBytes and thus bypassing the slow-start stage of normal TCP window evolution. This means that the satellite link is now comparable to or even faster than the server-cache and cache-proxy connections in terms of throughput. For instance, for a file smaller than 64 KBytes, the entire file fits into the very first window. Therefore the transfer of the file is constrained by how fast the proxy receives rather than how fast the proxy can send since the window size would be perceived as “unlimited” for such a file. Thus having the file in the cache (much closer to the proxy) would enable the proxy to receive much faster than having to fetch the file remotely, and results in higher throughput and lower latency. This result highlights the importance of optimizing different segments of a split connection. More importantly, such optimization has to be done in a way to reduce asymmetry between the segments, e.g., to bring the slower link faster, which in this case roughly corresponds to using a large initial window size over the satellite link.

### *B. Effect of simultaneous connections*

Results from the previous subsection were obtained by having 10 simultaneous connections to the HGW from each client (5 to the cache and 5 to the remote server). We showed that receiving faster at the proxy results in higher throughput. However, the proxy’s sending rate can be constrained by the number of connections it is handling. That is, although each connection has a large initial window size, this large window does not get fully utilized fast enough due to the number of simultaneous connections. In this

subsection we compare the performance as a result of different number of simultaneous connections from the same host.

Our measurements are taken as follows. First we simultaneously download 6 files, Files 1, 2, 4, 7, 10, and 12, repeatedly over a period of one hour from each host. Then we repeat this process with 9 simultaneous downloads, for Files 1, 2, 3, 4, 6, 7, 10, 11, and 12, and 14 simultaneous downloads, for Files 1-14. In this experiment, connections do not go through the cache and files are directly originated from the remote server, and the same experiment is run on both the proxy-enabled host and the proxy-disabled host (corresponding to Figures 3(e) and 3(f)). Figure 6 shows the throughput of split connection and end-to-end connection under these three scenarios and Figure 7 compares their GoS.

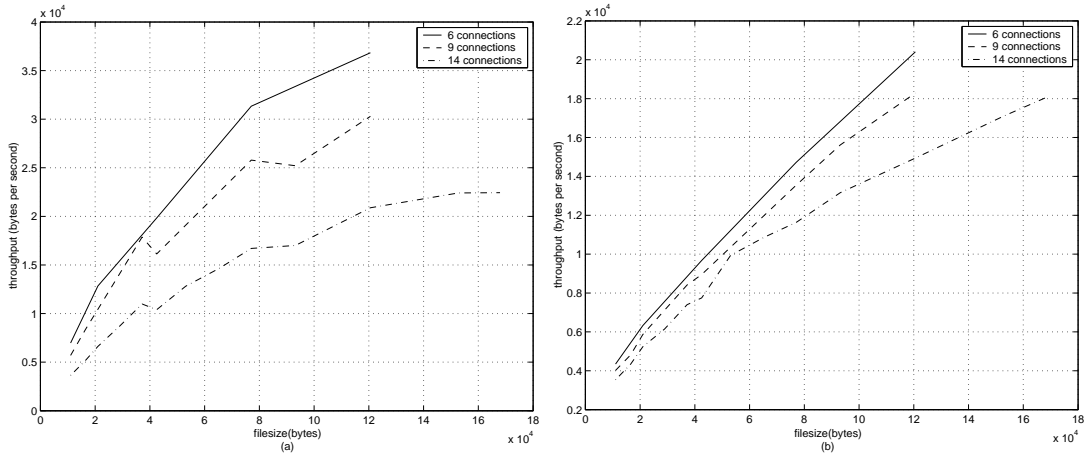


Fig. 6. Throughput for different number of simultaneous connections (a) Splitting enabled; (b) Splitting disabled




Fig. 7. GoS with different number of simultaneous connections

As can be seen from these results, connection splitting suffers more from higher number of simultane-

ous connections than end-to-end connections (although the gain is still positive). This is mostly due to the fact that with more simultaneous connections, the large initial window size of the split connections becomes less effective.

As mentioned earlier, the maximum overall throughput achieved per client was observed to be between 300-400 Kbytes/sec. Although each of the split connections has an initial window size of 64 KBytes, when there are multiple simultaneous connections, this 64 KBytes cannot be fully utilized all at once. The effective share of window size each connection actually realizes decreases with the increase in the number of simultaneous connections. For connections with the proxy disabled, the initial window size is much smaller than 64 KBytes, and therefore the impact of increased number of simultaneous connections is also much smaller.

Another reason for this could be due to the higher round trip time in end-to-end case than that of when splitting is used. Both end-to-end and split-connection packets experience the same queuing delay at the Satellite Gateway. It's well-known that the steady-state throughput of TCP is inversely proportional to RTT (see for example [23], [24]). For the split-connection case the round trip time ( $R$ ) is obviously smaller than the end-to-end case ( $R'$ ). Now we are adding a constant  $q$  (queuing delay) to the round trip time in both cases. Since  $R < R'$  we expect this constant to have a larger impact on the steady-state throughput of split connections.

We note that the results shown in Figures 6 and 7 are obtained over different time periods, and therefore could reflect different traffic load and congestion levels in the network. However, the same experiments were repeated several times and each time the results show the same trend with similar measurements. Therefore although random fluctuations in traffic load do exist in the network, the results we show here is typical and representative of the performance and change in performance under the given scenarios.

### *C. Effect of congestion and packet losses*

By examining the traces of each file download, we can determine the exact number of losses and retransmissions occurred per connection. However, such losses could involve both random and congestion losses, the distinction of which not directly available to us by only taking measurements at the end point. On the other hand, congestion and losses are highly correlated (cause and effect) with increased end-to-end delay, which is observable. In this section we illustrate the relationship between increased file transfer delay and the gain from using connection splitting. In doing so we attempt to understand the relationship between the splitting gain and congestion/losses.

First we repeatedly download a file directly from the server for 2 hours so that the resulting trace

may reflect a reasonable range of work load and congestion conditions in the network. We then sort the latency trace of the proxy-enabled connection in descending order, and reorder the proxy-disabled trace accordingly. These two cases corresponds to that illustrated in Figures 3(a) and 3(c). Figure 8 shows the reordered traces for Files 2 (11033 bytes) and 12 (120579 bytes). Figure 9 shows the GoS for these two files. It can be seen that the gain decreases as the latency of the proxy-enabled connection increases. This decrease is much steeper for the 11 Kbyte file and there is a sizeable portion of samples showing the proxy-enabled connections under-performing the proxy-disabled connections. This however is not observed in the case of the 120 Kbyte file. In the 120Kbyte file transfer, split connection always results in smaller latency (higher throughput) throughout the duration of our measurement, whereas in the 11Kbyte file transfer, there are some samples showing higher latency (smaller throughput) for the split connection in comparison to end-to-end. We define guarantee of performance improvement as  $\frac{\text{number of samples with } GoS > 0}{\text{total number of samples}}$ .




Fig. 8. Sorted latency traces in case of a cache miss for (a) File 12 (120K) and (b) File 2 (11K)

We then repeat the same experiment with the same files, but this time files are directly from the cache (corresponding to Figures 3(b) and 3(d)). Figures 10 and 11 show the latency and GoS in this case. Note that the correlation between the slowdowns of proxy-enabled and proxy disabled connections is not very obvious. This is due to the fact that the slowdowns are mostly caused by random packet losses rather than persistent congestion.

There are two main observations from these results. (1) The splitting gain decreases as the file transfer latency increases due to higher loss and/or congestion. (2) Whether the file is directly from the cache or from the remote server, the proxy-enabled connections experience worse performance (higher latency) than the proxy-disabled connections for small file transfers (e.g. 11 KByte), for a small portion of the




Fig. 9. GoS in case of a cache miss for (a) File 12 (120K) and (b) File 2 (11K)

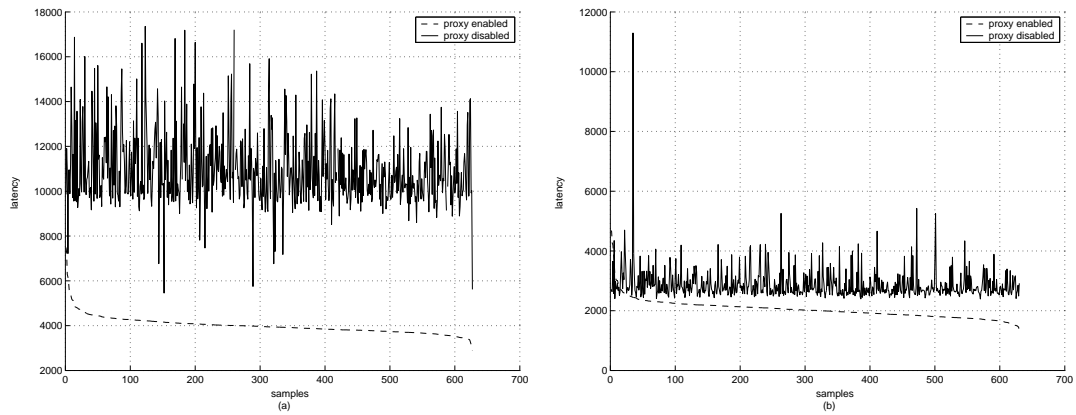


Fig. 10. Sorted latency traces in case of a cache hit for (a) File 12 (120K) and (b) File 2 (11K)




Fig. 11. GoS in case of a cache hit for (a) File 12 (120K) and (b) File 2 (11K)

TABLE II

PERCENTAGE OF SAMPLES WHERE DISABLING THE PROXY OUTPERFORMS ENABLING THE PROXY

| File No. | Size (bytes) | % of GoS < 0 (cache hit) | % of GoS < 0 (cache miss) |
|----------|--------------|--------------------------|---------------------------|
| 1        | 11033        | 2.2                      | 3.49                      |
| 2        | 16565        | 0.79                     | 1.26                      |
| 3        | 21027        | 0.0                      | 0.48                      |
| 4        | 65331        | 0.16                     | 0.0                       |
| 5        | 120579       | 0.0                      | 0.0                       |

samples. This portion is reduced in the case of a cache hit. The same phenomenon was not observed for large file transfers (e.g., 120 KByte). The percentages of negative GoS for different files are listed in Table II (guarantee of performance improvement is one minus each term in the table). Note that the *average* performance of using the splitting proxy is still above that of disabling the proxy.

Since our connections go through a dedicated proxy, the fluctuation in file transfer latency as seen from these traces is mainly due to the fluctuation in work load, congestion and loss situations elsewhere along the path of the connection, i.e., from the server to the cache (in the case of a cache miss), from the cache to the proxy, and from the proxy to the end host. (Note that in general, connection splitting leading to worse performance can be caused by excessive congestion and delay at the proxy, which would only be experienced by split traffic, but not by end-to-end traffic. This can happen if the proxy handles many split connections at the TCP layer, while the end-to-end traffic simply goes through the IP queue and be unaffected. However, since we use a dedicated proxy the increase in delay and loss incurred by splitting is minimal.) The reason that connection splitting can result in higher latency for small files lies within the relationship between the reduction in latency due to splitting and the increase in latency due to losses in general, both as a function of file size. The reduction in latency by using the proxy is largely due to the fact that the end-to-end TCP is broken down into shorter feedback loops and that a large initial window size is used over the satellite. When the file is small, one more or one less lost packet during the transfer, or whether the loss occurs sooner or later can result in significant difference in latency. This difference may not be compensated by the benefit from using the splitting proxy when the file is relatively small since the transfer completes soon afterwards. However, as the file size increases, the benefit of splitting becomes more prominent (see our results in Section 3.A). In other words, it would take more packet losses for a split connection to perform at a similar level as a non-split connection, which can have very

TABLE III

THROUGHPUT OF FILES WITH DIFFERENT NUMBER OF EMBEDDED OBJECTS (IN THE CASE OF A CACHE MISS)

| No. of embedded objects | Proxy-enabled (bytes/sec) | Proxy-disabled (bytes/sec) | GoS  |
|-------------------------|---------------------------|----------------------------|------|
| 0                       | 32124                     | 12544                      | 1.56 |
| 7                       | 6027                      | 3663                       | 0.64 |
| 19                      | 2701                      | 2026                       | 0.33 |

small probability considering the fact that a long connection tends to reflect more average loss rate. When the file is located in the cache, one segment (server-cache) is eliminated from the path, thus reducing the probability of packet losses and consequently reducing the probability that a particular split connection experiences longer latency than a non-split connection due to different loss occurrences.

In summary, when the file size is large, split connections can sustain more losses than non-split connections and still yield shorter latency, thus provide higher probability of performance improvement. When the file is small, the split connection is more affected by packet losses and therefore the probability of performance improvement is lower.

#### D. Effect of embedded objects and persistent connection

So far all our measurements are taken from full text files transferred using HTTP. In this section we examine the effect of embedded objects in a file/page, and the effect of using persistent connection.

We first compare the latency for files with same size but different numbers of embedded objects. We repeatedly download Files 9, 15 and 16 over a two-hour period. File 9 is a text file, and Files 15 and 16 consist of equal-sized embedded objects. The total size of each of these three files is about 64 Kbytes. In downloading these files, HTTP/1.0 [21] is used between the end hosts and the proxy or the cache. The throughput of proxy-enabled and proxy-disabled transfers is shown for both the cache miss (corresponds to Figures 3(a) and 3(c)) and cache hit (corresponds to Figures 3(b) and 3(d)) cases, in Tables III and IV, respectively.

We see that when a file contains a larger number of embedded objects the GoS decreases. This result is expected considering our observations in Sections 3.A that the gain from using the splitting proxy decreases as the file size decreases. This is because the time spent in handshake has a bigger portion in the overall latency for smaller files. Thus if we break a large file into many small objects and open a new connection for each of these objects, we expect to see a lower performance gain.



TABLE IV

THROUGHPUT OF FILES WITH DIFFERENT NUMBER OF EMBEDDED OBJECTS (IN THE CASE OF A CACHE HIT)

| No. of embedded objects | Proxy-enabled (bytes/sec) | Proxy-disabled (bytes/sec) | GoS  |
|-------------------------|---------------------------|----------------------------|------|
| 0                       | 31451                     | 15899                      | 0.98 |
| 7                       | 6006                      | 3845                       | 0.56 |
| 19                      | 2772                      | 2110                       | 0.31 |

It should be noted that the measurements in Tables III and IV are taken over different time periods and thus numbers from different tables are not directly comparable.

Next we explore the performance comparison between using HTTP/1.0 and HTTP/1.1 [22]. Web browsers are typically configured to use HTTP/1.0 in the proxy mode where a separate connection is established for each embedded object on a page. HTTP/1.0 generally results in large delay due to the time spent in hand-shaking and low throughput in the slow-start phase of TCP. HTTP/1.1 on the other hand opens up a *persistent* connection which is used to deliver both the base page and subsequent embedded objects. Latency is thus reduced since with a single connection for everything there is only one hand-shake procedure and one slow-start stage of TCP. There has been extensive studies on the performance of different versions of HTTP, see for example [25], [26].

Here we compare HTTP/1.0 when proxy is enabled vs. HTTP/1.1 when proxy is disabled. The connection setup of this part of our measurements corresponds to Figures 3(e) and 3(f), i.e., the connections do not go through the cache and that the connection is either end-to-end (proxy-disabled) or split into two (proxy-enabled). We repeatedly download Files 9,15,16 and 17 for over a period of 30 minutes and measure their throughput.

For comparison purposes we summarize the results obtained from different combination of splitting enabled/disabled and persistent connection enabled/disabled in Tables V, VI and VII (we did not compare HTTP/1.0 when proxy is disabled vs. HTTP/1.1 when proxy is enabled, as we see no special interest in this case). Again the results from different tables are not directly comparable with each other, as they are taken over different time periods which may reflect different work load situations in the network. Results within the same table are measurements over the same period and are thus comparable.

We see that overall the gain from using connection splitting decreases with increased number of embedded objects in a file (of the same size). This is consistent with our results earlier in this subsection. However when persistent connection is used, GoS decreases much faster than the case for non-persistent

TABLE V

COMPARISON BETWEEN SPLITTING ENABLED AND DISABLED WITH NON-PERSISTENT CONNECTION

| No. of embedded objects                                      | 0     | 3     | 7    | 18   |
|--|-------|-------|------|------|
| Throughput (Byte/sec) for non-persistent, splitting-disabled | 15325 | 5910  | 3864 | 2197 |
| Throughput (Byte/sec) for non-persistent, splitting-enabled  | 33130 | 10937 | 6409 | 3024 |
| GoS  | 1.16  | 0.85  | 0.66 | 0.38 |

TABLE VI

COMPARISON BETWEEN SPLITTING ENABLED AND DISABLED WITH PERSISTENT CONNECTION

| No. of embedded objects                                  | 0     | 3     | 7    | 18    |
|--|-------|-------|------|-------|
| Throughput (Byte/sec) for persistent, splitting-disabled | 12361 | 9902  | 8862 | 5239  |
| Throughput (Byte/sec) for persistent, splitting-enabled  | 26468 | 14636 | 9890 | 5136  |
| GoS  | 1.14  | 0.48  | 0.12 | -0.02 |

connections as shown in Table VI (note that so long as  $GoS > 0$  it means that split connection has a higher throughput than end-to-end connection). The decreasing throughput is due to the fact that the client requests each object individually (although within the same connection) after receiving the base page. Depending on the number of objects and the size of these objects, using end-to-end connection could achieve very similar performance to that of a split connection. Here we would like to note that although we did not provide any measurement for the case of persistent connection with pipelining, its effect can be predicted using Table VI. When pipelining is used, all the requests for embedded objects are sent at the same time after receiving the base page. This is similar to having one large embedded object instead of many small ones. Using Table VI (moving towards the left) we would expect pipelining to result in higher GoS.

Table VII compares proxy with non-persistent connection vs. persistent connection without proxy. We see that the gain from using connection splitting quickly becomes negative, i.e. split connection results in smaller throughput than end-to-end connection. This shows that using a persistent connection is far more important than connection splitting for a file that contains multiple embedded objects (for even a relatively small number) since the time spent in handshake in this case is significant over the satellite link. In other words, the gain from using connection splitting cannot substitute for the performance improvement achieved by having a persistent connection. Therefore in an Internet over satellite system,

TABLE VII

COMPARISON BETWEEN SPLITTING ENABLED WITH PERSISTENT CONNECTION AND DISABLED WITH  
NON-PERSISTENT CONNECTION

| No. of embedded objects                                     | 0     | 3    | 7     | 18    |
|---|-------|------|-------|-------|
| Throughput (Byte/sec) for persistent, splitting disabled    | 11194 | 9081 | 8244  | 4827  |
| Throughput (Byte/sec) for non-persistent, splitting enabled | 29977 | 9920 | 5813  | 2699  |
| GoS   | 1.68  | 0.09 | -0.29 | -0.44 |

configuring the browser to run the default HTTP/1.0 with a connection splitting proxy would defeat the purpose of having a proxy, unless all the files are pure text. It is thus crucial that a persistent connection is established between the client and the proxy in such a system.

#### IV. CONCLUSIONS AND DISCUSSIONS

In this paper we examined the gain from using a TCP connection splitting proxy along with web caching under various scenarios by taking measurements from a live commercial system. In these experiments, not all parameters are controllable, e.g., the bit error rate over the satellite link, the amounts of traffic and congestion in the network, etc. However, by properly isolating elements that contribute to the result, not only does our measurement confirm earlier studies on connection splitting, but it also revealed additional insights into the common performance enhancing techniques used for Internet over satellite.

To summarize, connection splitting is a valid approach to improve TCP throughput and reduce file transfer latency over the satellite. The performance gain increases as the file size increases, but decreases as the number of embedded objects in the file increases. This gain is also very sensitive to the number of simultaneous connections and to congestion and packet losses, especially for small files. On the one hand having a cache hit will alleviate such sensitivity, on the other hand using connection splitting proxy enhances the benefit of caching. We also showed that although connection splitting improves throughput, it is no substitute for persistent connection. The best performance is achieved by using both the splitting proxy and persistent connection (HTTP/1.1) between the proxy and the client.

By segregating an end-to-end connection into separate connection segments we have shortened TCP feedback control loops. Thus for each of these connection segments we have a smaller RTT and lower loss probability  $p$  comparing to the original end-to-end connection. Since the TCP throughput is inversely proportional to RTT and  $\sqrt{p}$  (see for example [23], [24]), having smaller connection segments naturally

results in higher throughput. However, these segments are not completely un-coupled, and the slowest segment will constrain the other segments and dominate the end-to-end performance. The proxy cannot forward client data it has not received from the server. Therefore if the proxy is constrained by slow receiving then increasing the rate at which proxy receives will improve the overall performance. This is why cache hits improve the throughput when the splitting proxy is enabled. On the other hand if we only use the cache, we will have a cache-client connection highly constrained by slow sending due to the large propagation delay. This is why having a large initial window size over the satellite link is so important. Overall this illustrates the importance of properly provisioning the proxy and the NOC, and the importance of optimizing each split segment to reduce the asymmetry among them so that one is not constrained by the other.

Our measurement and subsequent results are obtained from a live system, over which we have limited control. These results therefore may reflect the combined effect of a range of factors. We have reasoned the most prominent factor(s) and cause(s) for each of these results, but we strongly believe that deriving a detailed quantitative analysis will help explain the exact role of each factor and serve as good supplement to our measurement-based study. This is part of our current research.

#### REFERENCES

- [1] C. Parsa and J. J. Garcia-Luna-Aceves, "Improving TCP Performance Over Wireless Network at The Link Layer," *ACM Mobile Networks & Applications Journal*, 1999.
- [2] C. Parsa and J. J. Garcia-Luna-Aceves, "TULIP: A Link-Level Protocol for Improving TCP over Wireless Links," *Proc. IEEE WCNC'99*, pp. 1253–1257, 1999.
- [3] K. Ratnam and Ibrahim Matta, "WTCP: An Efficient Mechanism for Improving TCP Performance Over Wireless Links," *Proc. IEEE ISCC*, pp. 74–78, 1998.
- [4] K. Ratnam and Ibrahim Matta, "Effect of Local Retransmission at Wireless Access Points on The Round Trip Time Estimation of TCP," *Proc. 31st Annual Simulation Symp.*, pp. 150–156, 1998.
- [5] Liviu Iftode Ramón Cáceres, "Improving the Performance of Reliable Transport Protocol in Mobile Computing Environment," *IEEE J-SAC*, vol. 13, no. 5, pp. 850–857, 1995.
- [6] N. H. Vaidya, M. Mehta, C. Perkins, and G. Montenegro, "Delayed Duplicated Acknowledgments: A TCP-Unaware Approach to Improve Performance of TCP over Wireless," *Technical Report 99-003, TAMU*, 1999.
- [7] S. Biaz and N. H. Vaidya, "Discriminating Congestion Losses From Wireless Losses Using Inter-arrival Times at the Receiver," *Proc. IEEE ASSET*, pp. 10–17, 1999.
- [8] H. Balakrishnan, S. Seshan, E. Amir, and R. H. Katz, "Improving TCP/IP Performance Over Wireless Networks," *ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom'95)*, vol. 2, no. 11, 1995.
- [9] H. Balakrishnan, V. N. Padmanabhan, S. Seshan, and R. H. Katz, "A Comparison of Mechanisms for Improving TCP Performance over Wireless Links," *IEEE/ACM Trans. on Networking*, vol. 5, no. 6, pp. 756–769, 1997.

- [10] B.R. Badrinath and P. Sudame, "To Send or Not to Send: Implementing Deferred Transmissions in a Mobile Host," *Proc. IEEE ICDCS*, pp. 327–333, 1996.
- [11] V. G. Bharadwaj, "Improving TCP Performance over High-Bandwidth Geostationary Satellite Links," Tech. Rep. MS 99-12, Institute for Systems Research, University of Maryland, College Park, 1999, <http://http://www.isr.umd.edu/TechReports/ISR/1999/>.
- [12] A. Bakre and B.R. Badrinath, "I-TCP: Indirect TCP for Mobile Hosts," *Proc. IEEE ICDCS*, pp. 136–143, 1995.
- [13] A. V. Bakre and B. R. Badrinath, "Implementation and Performance Evaluation of Indirect TCP," *IEEE Transactions on Computers*, vol. 46, no. 3, pp. 260–278, 1997.
- [14] M. Allman, S. Dawkins, D. Glover, D. Tran, T. Henderson, J. Heidemann, J. Touch, H. Kruse, S. Ostermann, K. Scott, and J. Semke, "Ongoing TCP Research Related to Satellites," *IETF RFC 2760*, 2000.
- [15] M. Karir, "IPSEC and the Internet," Tech. Rep. MS 99-14, Institute for Systems Research, University of Maryland, College Park, 1999, <http://http://www.isr.umd.edu/TechReports/ISR/1999/>.
- [16] J. Ishac and M. Allman, "On the Performance of TCP Spoofing in Satellite Networks," *IEEE Milcom*, 2001.
- [17] P. Rodriguez, S. Sibal, and O. Spatscheck, "TPOT: Translucent Proxying of TCP," Tech. Rep., AT & T labs-Research and EURECOM Technical Report, 2000.
- [18] M. Liu and N. Ehsan, "Modeling and Analysis of TCP Enhancement over Heterogeneous Links," *International Conference on Networking (ICN'01)*, 2001.
- [19] M. Mathis, J. Mahdavi, S. Floyd, and A. Romanow, "TCP Selective Acknowledgement Options," *IETF RFC 2018*, 1996.
- [20] H. Balakrishnan, V. N. Padmanabhan, and R. H. Katz, "The Effects of Asymmetry on TCP Performance," *ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom'95)*, 1997.
- [21] T. Berners-Lee, R. Fielding, and H. Frystyk, "Hypertext Transfer Protocol – HTTP/1.0," *IETF RFC 1945*, 1995.
- [22] R. Fielding, J. Gettys, J. Mogul, H. Frystyc, and Berners-Lee, "Hypertext Transfer Protocol – HTTP/1.1," *IETF RFC 2068*, 1997.
- [23] J. Padhye, V. Firoiu, D. F. Towsley, and J. F. Kurose, "Modeling TCP Reno Performance: A Simple Model and Its Empirical Validation," *IEEE Trans. Networking*, vol. 8, no. 2, pp. 133–145, 2000.
- [24] T. V. Lakshman, U. Madhow, and B. Suter, "TCP Performance with Random Loss and Bidirectional Congestion," *IEEE Trans. Networking*, vol. 8, no. 5, pp. 541–555, 2000.
- [25] J. Heidemann, K. Obraczka, and J. Touch, "Modeling the Performance of HTTP over Several Transport Protocols," *IEEE/ACM Transactions on Networking*, vol. 5, no. 5, 1997.
- [26] H. F. Nielsen et al., "Network Performance Effects of HTTP/1.1, Cssi, and PNG," Tech. Rep., 1997, <http://www.w3.org/Protocols/HTTP/Performance/Pipeline.html>.