Predicting short-transfer latency from TCP arcana: extended version

Martin Arlitt, Balachander Krishnamurthy[1], Jeffrey C. Mogul
Internet Systems and Storage Laboratory
HP Laboratories Palo Alto
HPL-2005-137
September 30, 2005*

TCP latency, network performance

In some contexts it may be useful to predict the latency for short TCP transfers. For example, a Web server could automatically tailor its content depending on the network path to each client, or an "opportunistic networking" application could improve its scheduling of data transfers.

Several techniques have been proposed to predict the latency of short TCP transfers based on online measurements of characteristics of the current TCP connection, or of recent connections from the same client. We analyze the predictive abilities of these techniques using traces from a variety of Web servers, and show that they can achieve useful accuracy in many, but not all, cases. We also show that a previously-described model for predicting short-transfer TCP latency can be improved with a simple modification. Ours is the first trace-based analysis that evaluates these prediction techniques across diverse user communities.

* Internal Accession Date Only
[1] AT&T Labs – Research, Florham Park, NJ 07932
This is an extended version of a paper original published in the Proceedings of the Internet Measurement Conference, 19-21 October 2005, Berkeley, CA, USA

Approved for External Publication

Predicting short-transfer latency from TCP arcana: extended version

Martin Arlitt
HP Labs/University of Calgary
Palo Alto, CA 94304
Martin.Arlitt@hp.com

Balachander Krishnamurthy
AT&T Labs–Research
Florham Park, NJ 07932
bala@research.att.com

Jeffrey C. Mogul
HP Labs
Palo Alto, CA 94304
Jeff.Mogul@hp.com

Abstract

In some contexts it may be useful to predict the latency for short TCP transfers. For example, a Web server could automatically tailor its content depending on the network path to each client, or an "opportunistic networking" application could improve its scheduling of data transfers.

Several techniques have been proposed to predict the latency of short TCP transfers based on online measurements of characteristics of the current TCP connection, or of recent connections from the same client. We analyze the predictive abilities of these techniques using traces from a variety of Web servers, and show that they can achieve useful accuracy in many, but not all, cases. We also show that a previously-described model for predicting short-transfer TCP latency can be improved with a simple modification. Ours is the first trace-based analysis that evaluates these prediction techniques across diverse user communities.

1 Introduction

It is often useful to predict the latency (i.e., duration) of a short TCP transfer before deciding when or whether to initiate it. Network bandwidths, round-trip times (RTTs), and loss rates vary over many orders of magnitude, and so the transfer latency for a given data item can vary similarly.

Examples where such predictions might be useful include:

- a Web server could automatically select between "low-bandwidth" and "high-bandwidth" versions of content, with the aim of keeping the user's download latency below a threshold [11, 20].
- A Web server using shortest-remaining-processing-time (SRPT) scheduling [23] could better predict overall response times if it can predict network transfer latency, which in many cases is the primary contributor to response time.
- An application using *opportunistic networking* [25] might choose to schedule which data to send based on an estimate of the duration of a transfer opportunity and predictions of which data items can make the most effective use of that opportunity.

There are several possible ways to define "short" TCP transfers. Models for TCP performance typically distinguish between long flows which have achieved steady state, and short flows which do not last long enough to leave the initial slow-start phase. Alternatively, one could define short in terms of an arbitrary threshold on transfer length. While defining "short" in terms of slow-start behavior is less arbitrary, it is also less predictable (because the duration of slow-start depends on unpredictable factors such as cross traffic and packet loss), and so for this paper we use a definition based on transfer length. Similarly, while transfer length could be defined in terms of the number of data packets sent, this also depends

on unpredictable factors such as MTU discovery and the interactions between application buffering and socket-level buffering. So, for simplicity, in this paper we define "short" in terms of the number of bytes transferred.

Several techniques have previously been proposed for automated prediction of the transfer time for a short TCP transfer. Some of these techniques glean their input parameters from characteristics of TCP connections, such as round-trip time (RTT) or congestion window size (cwnd), that are not normally exposed to the server application. We call these characteristics *TCP arcana*. These parameters can then be used in a previously-described model for predicting short-transfer latency [2]. Other techniques use observations of the actual latency for past transfers to the same client (or to clients in a similar location of the network), and assume that past performance is a good predictor of future performance.

In this paper, we use packet-level traces captured near a variety of real Web servers to evaluate the ability of techniques based on both TCP arcana and historical observations to predict short transfer latencies. We show that the previously-described model does not quite fit the observations, but that a simple modification to the model greatly improves the fit. We also describe an experiment suggesting (based on a limited data set) that RTT observations could be used to discriminate, with modest accuracy, between dialup and non-dialup paths.

1.1 Related work

Our work complements previous work on predicting the *throughput* obtained by *long* TCP transfers. He *et al.* [9] characterized these techniques as either formula-based or history-based; our TCP arcana approach is formula-based.

Lakshminarayanan and Padmanabhan [14] briefly discussed the relationship between RTT and TCP throughput, but for transfer lengths of 100KB, since their paper focuses on peer-to-peer systems. They found a poor correlation between latency and throughput, which is not surprising, because for long transfers the formula-based method requires knowledge of packet loss rates, which they did not measure. They did remark that "latency may in fact be a good predictor of throughput when dial-up hosts ... are included," which agrees with the results we present in Section 5.

Hall *et al.* [8] studied the effect of early packet loss on Web page download times. As with our work, they point out that download times are not always correlated with path bandwidth. They focused on packet losses occurring early enough in the TCP connection that "neither client nor server has seen enough packets to establish a usable round trip time estimation," which leads to excessively long retransmission timeouts.

2 Latency prediction techniques

We start with the assumption that an application wishing to predict the latency of a short transfer must do so as early as possible, before any data has been transferred. We also assume that prediction is being done at the server end of a connection that was initiated by a client; although the approaches could be extended to client-side prediction, we have no data to evaluate that scenario.

We examine two prediction approaches in this paper:

- The **initial-RTT** approach: The server's first possible measurement of the connection RTT is provided by the interval between its initial SYN|ACK packet and the client's subsequent ACK. For short transfers, this RTT measurement is often sufficient to predict subsequent data transfer latency to this client. This approach was first proposed by Mogul and Brakmo [19] and discussed in [20]. We describe it further in Section 2.1.
- The **recent-transfers** approach: A server can predict the data transfer bandwidth to a given request

based on recently measured transfer bandwidths to the same client. This approach, in the context of Web servers, was proposed in [11]; we describe it further in Section 2.2.

2.1 Prediction from initial RTTs

Suppose one wants to predict the transfer latency, for a response of a given length over a specific HTTP connection, with no prior information about the client and the network path, and before having to make the very first decision about what content to send to the client. Let us assume that we do not want the server to generate extra network traffic or cause extra delays. What information could one glean from the TCP connection before it is too late?

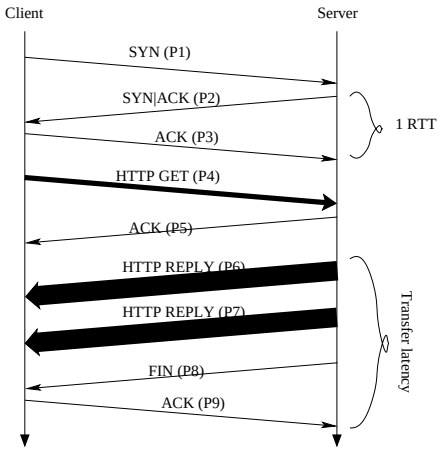

Figure 1: Timeline: typical HTTP connection

Figure 1 shows a timeline for the packets sent over a typical non-persistent HTTP connection. (We assume that the client TCP implementation does not allow the client application to send data until after the 3-way handshake; this is true of most common stacks.) In this timeline, the server has to make its decision immediately after seeing the GET-bearing packet (P_4) from the client.

It might be possible to infer network path characteristics from the relative timing of the client's first ACK-only (P_3) and GET (P_4) packets, using a packet-pair method [13]. However, the initial-RTT predictor instead uses the path's RTT, as measured between the server's SYN|ACK packet (P_2) and the client's subsequent ACK-only packet (P_3). Since these two packets are both near-minimum length, they provide a direct measurement of RTT, in the absence of packet loss.

Why might this RTT be a useful predictor of transfer latency?

- Many last-hop network technologies impose both high delay and low bandwidth. For example, dialup modems almost always add about 100ms to the RTT [4, 5] and usually limit bandwidth to under 56Kb/s. If we observe an RTT much lower than 100ms, we can infer that the path does not involve a modem. (See Section 5 for quantitative evidence.) A similar inference might be made about some (perhaps not all) popular low-bandwidth wireless media.
- Even when the end-to-end bandwidth is large, the total transfer time for short responses depends mostly on the RTT. (Therefore, an HTTP request header indicating client connection speed would not reliably predict latency for such transfers.)

Cardwell *et al.* [2] showed that for transfers smaller than the limiting window size, the expected latency

to transfer d segments via TCP, when there are no packet losses, is *approximated* by

$$E[latency] = RTT \cdot log_{\gamma}(\frac{d(\gamma-1)}{w_1} + 1) \quad (1)$$

where

- γ depends on the client's delayed-ACK policy; reasonable values are 1.5 or 2 (see [2] for details).
- w_1 depends on the server's initial value for `cwnd`; reasonable values are 2, 3, or 4 (see [2] for details).
- $d = \lceil \frac{len}{MSS} \rceil$
- len is the number of bytes sent.
- MSS is the TCP maximum segment size for the connection.

Note that median Web response sizes (we use the definition of "response" from the HTTP specification [6]) are typically smaller than the limiting window size; see Section 3.4.

End-to-end bandwidth limits and packet losses can only increase this latency. In other words, if we know the RTT and response size, then we can predict a lower bound for the transfer latency.

We would like to use the RTT to predict the transfer latency as soon as possible. Therefore, the first time a server sees a request from a given client, it has only one RTT measurement to use for this purpose. But if the client returns again, which RTT measurement should the server use for its prediction? It could use the most recent measurement (that is, from the current connection), as this is the freshest; it could use the mean of all measurements, to deal with noise; it could use an exponentially smoothed mean, to reduce noise while favoring fresh values; it could use the minimum measurement, to account for variable queueing delays; or it could use the maximum measurement, to be conservative.

"Most recent," which requires no per-client state, is the simplest to implement, and this is the only variant we have evaluated.

2.2 Prediction from previous transfers

Krishnamurthy and Wills originally described the notion of using measurements from previous transfers to estimate the connectivity of clients [11]. A prime motivation of this work was to retain poorly connected clients, who might avoid a Web site if its pages take too long to download. Better connected clients could be presented enhanced versions of the pages.

This approach is largely passive: it examines server logs to measure the inter-arrival time between base-object (HTML) requests and the requests for the first and last embedded objects, typically images. Exponentially smoothed means of these measurements are then used to classify clients. A network-aware clustering scheme [10] was used as an initial classification mechanism, if a client had not been seen before but another client from the same cluster had already used the site. Krishnamurthy and Wills used a diverse collection of server logs from multiple sites to evaluate the design, and Krishnamurthy *et al.* presented an implementation [12], using a modified version of the Apache server, to test the impact of various server actions on clients with different connectivity.

The recent-transfers approach that we study in this paper is a simplification of the Krishnamurthy and Wills design. Because their measurements use Web server logs, this gave them enough information about page structure to investigate the algorithm's ability to predict the download time for an entire page, including embedded objects. We have not extracted object-relationship information from our packet traces, so we only evaluated per-response latency, rather than per-page latency. On the other hand, most server logs provide timing information with one-second resolution, which means that a log-based evaluation cannot provide the fine-grained timing resolution that we got from our packet traces.

2.3 Defining transfer latency

We have so far been vague about defining "transfer latency." One might define this as the time between the departure of the first response byte from the server and the arrival of the last response byte at the client. However, without perfect clock synchronization and packet traces made at every host involved, this duration is impossible to measure.

For this paper, we define transfer latency as the time between the departure of the first response byte from the server and the arrival *at the server* of the acknowledgment of the last response byte. (Figure 1 depicts this interval for the case of a non-persistent connection.) This tends to inflate our latency measurement by approximately RTT/2, but because path delays can be asymmetric we do not attempt to correct for that inflation. We are effectively measuring an upper bound on the transfer latency.

2.4 Predicting at the client

This paper focuses on making latency predictions at the server end of a connection. We believe that the techniques we describe should be usable at the client end. For example, Figure 1 shows how the server can obtain an RTT sample from the timestamps of the SYN|ACK packet (P_2) and the ACK-only packet (P_3). But the client can also get an early RTT sample, from the timestamps of the initial SYN (P_1) and the SYN|ACK (P_2).

Similarly, a client could maintain historical information to drive a recent-transfers predictor. While a single client would not have sufficient history with respect to many servers, a pool of clients might jointly obtain enough history (as in the throughput-oriented SPAND system [24]). Such an approach would probably work best if the clients in the pool were located near to each other, in terms of network topology.

Unfortunately, our server-centric traces do not allow us to evaluate client-based latency prediction.

3 Methodology

We followed this overall methodology:

- **Step 1:** collect packet traces near a variety of Web servers with different and diverse user populations.
- **Step 2:** extract the necessary connection parameters, including client IDs, from these raw traces to create intermediate traces.
- **Step 3:** evaluate the predictors using simple simulator(s) driven from the intermediate traces.

Although the prediction mechanisms analyzed in this paper are not necessarily specific to Web traffic, we limited our trace-based study to Web traffic because we have not obtained significant and diverse traces of other short-transfer traffic. It might be useful to capture traffic near busy e-mail servers to get another relevant data set, since e-mail transfers also tend to be short [7, 17].

Given that we are defining "short" TCP transfers in terms of the number of data bytes sent, we analyzed three plausible thresholds: 8K bytes, 16K bytes, and 32K bytes; this paper focuses on the 32K byte threshold. (The response-size distributions in Figure 3 support this choice.)

3.1 Trace sets

We collected trace sets from several different environments, all in North America. For reasons of confidentiality, we identify these sets using short names:

- **C2:** Collected on a corporate network
- **U2,U3,U4:** Collected at a University
- **R2:** Collected at a corporate research lab

In all cases, the traces were collected on the public Internet (not on an Intranet) and were collected relatively near exactly one publicly-accessible Web server.

We collected full-packet traces, using tcpdump, and limited the traces to include only TCP connections to or from the local Web server.

While we wanted to collect traces covering an entire week at each site, storage limits and other restrictions meant that we had to collect a series of shorter traces. In order to cover representative periods over the course of a week (May 3–9, 2004), we chose to gather traces for two to four hours each day: 9:00AM-11:00AM Monday, Wednesday, and Friday; 2:00PM-4:00PM Tuesday and Thursday; and 10:00AM-2:00PM Saturday and Sunday (all are local times with respect to the trace site: MST for C2, MDT for U2, and PDT for R2). We additionally gathered two 24-hour (midnight to midnight) traces at the University: U3 on Thursday, Aug. 26, 2004, and U4 on Tuesday, Aug. 31, 2004.

3.2 Are these traces representative?

We certainly would prefer to have traces from a diverse sample of servers, clients, and network paths, but this is not necessary to validate our approach. Our goal is not to predict the latencies seen by all client-server pairs in the Internet, but to find a method for a given server to predict the latencies that it itself (and only itself) will encounter in the near future.

It is true that some servers or client populations might differ so much from the ones in our traces that our results do not apply. Although logistical and privacy constraints prevent us from exploring a wider set of traces, our analysis tools are available at *http://bro-ids.org/bro-contrib/network-analysis/akm-imc05/* so that others can test our analyses on their own traces.

The results in Section 4.6 imply that our equation-based predictor works well for some sites and not so well for others. One could use our trace-based methodology to discover if a server's response latencies are sufficiently predictable before deciding to implement prediction-based adaptation at that server.

3.3 Trace analysis tools

We start by processing the raw (full-packet binary) traces to generate one tuple per HTTP request/response exchange. Rather than write a new program to process the raw traces, we took advantage of *Bro*, a powerful tool originally meant for network intrusion detection [21]. Bro includes a *policy script interpreter* for scripts written in Bro's custom scripting language, which allowed us to do this processing with a relatively simple policy script – about 800 lines, including comments. We currently use version 0.8a74 of Bro.

Bro reduces the network stream into a series of higher level events. Our policy script defines handlers for the relevant events. We identify four analysis states for a TCP connection: **not_established, timing_SYN_ACK, established,** and **error_has_occurred**. We also use four analysis states for each HTTP transaction: **waiting_for_reply, waiting_for_end_of_reply, waiting_for_ack_of_reply,** and **transaction_complete**. (Our script follows existing Bro practice of using the term “reply” in lieu of “response” for state names.)

Progression through these states occurs as follows. When the client's SYN packet is received, a data structure is created to retain information on the connection, which starts in the **not_established** state. When the corresponding SYN|ACK packet is received from the server, the modeled connection enters the **timing_SYN_ACK** state, and then to the **established** state when the client acknowledges the SYN|ACK.

We then wait for **http_request()** events to occur on that connection. When a request is received, a data structure is created to retain information on that HTTP transaction, which starts in the **waiting_for_reply** transaction state. On an **http_reply()** event, that state becomes **waiting_for_end_of_reply**. Once the server has finished sending the response, the transaction state is set to **waiting_for_ack_of_reply**. Once the entire HTTP response has been acknowledged by the client, that state is set to **transaction_complete**. This

design allows our script to properly handle persistent and pipelined HTTP connections.

Our analysis uses an additional state, **error_has_occurred**, which is used, for example, when a TCP connection is reset, or when a packet is missing, causing a gap in the TCP data. All subsequent packets on a connection in an **error_has_occurred** state are ignored, although RTT and bandwidth estimates are still recorded for all HTTP transactions that completed on the connection before the error occurred.

For each successfully completed and successfully traced HTTP request/response exchange, the script generates one tuple that includes the timestamp of the arrival time of the client's acknowledgement of all outstanding response data; the client's IP address; the response's length, content-type, and status code; the position of the response in a persistent connection (if any); and estimates of the initial RTT, the MSS, the response transfer latency, and the response transfer bandwidth. The latency is estimated as described in Section 2.3, and the bandwidth estimate is then computed from the latency estimate and the length.

These tuples form an intermediate trace, convenient for further analysis and several orders of magnitude smaller than the original raw packet trace. For almost all of our subsequent analysis, we examine *only* responses with status code = 200, since these are the only ones that should always carry full-length bodies.

3.3.1 Proxies and robots

Most Web servers receive requests from multi-client proxy servers, and from robots such as search-engine crawlers; both kinds of clients tend to make more frequent requests than single-human clients. Requests from proxies and robots skew the reference stream to make the average connection's bandwidth more predictable, which could bias our results in favor of our prediction mechanisms.

We therefore "pruned" our traces to remove apparent proxies and robots (identified using a separate Bro script); we then analyzed both the pruned and unpruned traces.

In order to avoid tedious, error-prone, and privacy-disrupting techniques for distinguishing robots and proxies, we tested a few heuristics to automatically detect such clients:

- Any HTTP request including a `Via` header probably comes from a proxy. The converse is not true; some proxies do not insert `Via` headers.
- Any request including a `From` header probably comes from a robot. Not all robots insert `From` headers.
- If a given client IP address generates requests with several different `User-Agent` headers during a short interval, it is probably a proxy server with multiple clients that use more than one browser. It could also be a dynamic IP address that has been reassigned to a different client, so the time scale affects the accuracy of this heuristic. We ignore `User-Agent: contype` headers, since this is an artifact of a particular browser [16, 18].

The results of these tests revealed that the `From` header is not widely used, but it is a reasonable method for identifying robots in our traces. Our test results also indicated that simply excluding all clients that issued a `Via` or `User-Agent` header would result in excessive pruning.

An analysis of the `Via` headers suggested that components such as personal firewalls also add this header to HTTP requests. As a result, we decided to only prune clients that include a `Via` header that can be automatically identified as a multi-client proxy: for example, those added by a Squid, NetApp NetCache, or Inktomi Traffic-Server proxy.

We adopted a similar approach for pruning clients that sent multiple different `User-Agent` headers. First, we require that the `User-Agent` headers be from well-known browsers (e.g., IE or Mozilla). These browsers typically form the `User-Agent` header in a very structured format. If we cannot identify the type

Table 1: Overall trace characteristics

			All HTTP status codes					status code = 200		
Trace name	Total Conns.	Total Clients	Total Resp. bytes	Total Resps.	mean resp. size (bytes)	mean req. rate	peak req. rate	Total Resp. bytes	Total Resps.	mean resp. size (bytes)
C2	323141	17627	3502M	1221961	3005	2.3/sec	193/sec	3376M	576887	6136
C2p (pruned)	281375	16671	3169M	1132030	2935	2.1/sec	181/sec	3053M	533582	5999
R2	33286	7730	1679M	50067	35154	0.1/sec	35/sec	1359M	40011	35616
R2p (pruned)	23296	6732	1319M	38454	35960	0.1/sec	31/sec	1042M	31413	34766
U2	261531	36170	5154M	909442	5942	1.7/sec	169/sec	4632M	580715	8363
U2p (pruned)	203055	33705	4191M	744181	5904	1.4/sec	152/sec	3754M	479892	8202
U3	278617	29843	5724M	987787	6076	11.4/sec	125/sec	5261M	637380	8655
U3p (pruned)	197820	26697	4288M	756994	5939	8.8/sec	117/sec	3940M	491497	8405
U4	326345	32047	6800M	1182049	6032	13.7/sec	139/sec	6255M	763545	8589
U4p (pruned)	230589	28628	5104M	902996	5926	10.5/sec	139/sec	4689M	588954	8347

of browser, the browser version, and the client OS, we do not use the header in the analysis. If we then see requests from two different browsers, browser versions, or client OSs coming from the same IP address in the limited duration of the trace, we consider this to be a proxy, and exclude that client from the pruned trace.

We opted to err (slightly) on the side of excessive pruning, rather than striving for accuracy, in order to reduce the chances of biasing our results in favor of our predictors. In Section 7.1, we discuss how an actual server implementation might detect proxies and robots, since the criteria could be different in that setting.

3.4 Overall trace characteristics

Table 1 shows various aggregate statistics for each trace set, to provide some context for the rest of the results. For reasons of space, we omit day-by-day statistics for C2, R2, and U2; these show the usual daily variations in load, although C2 and R2 peak on the weekend, while U2 peaks during the work week. The table also shows totals for the pruned versions of each trace set. Finally, the table shows total response bytes, response count, and mean response size for just the status-200 responses on which most subsequent analyses are based.

We add "p" to the names of trace sets that have been pruned (e.g., a pruned version of trace set "C2" is named "C2p"). Pruning reduces the number of clients by 5% (for trace C2) to 13% (for R2); the number of HTTP responses by 7% (for C2) to 23% (for R2, U3, and U4); and the peak request rate by 6% (for C2) to 11% (for R2).

The mean values in Table 1 do not convey the whole story. In Figures 2, 3, and 4, respectively, we plot cumulative distributions for request rate, response size, and latency for status-200 responses (These plots exclude the U3 and U4 traces, since these CDFs are nearly identical to those for the U2 trace; Figure 4 also excludes C2p and U2p, since these CDFs are nearly identical to those for the unpruned traces.)

Figure 5 shows a histogram of request counts per client address. One might expect that our pruning would change these distributions, if the proxies and robots removed by pruning do indeed generate an unusually high number of requests. However, the results in Figure 5 do not strongly support this expectation. We are not sure if this reflects a flaw in our pruning approach, or simply that most proxies and robots do not visit the traced sites very often.

The three traces in Figure 3 show quite different response size distributions. The responses in trace C2 seem somewhat smaller than has typically been reported for Web traces; the responses in trace R2 are a lot larger. (These differences also appear in the mean response sizes in Table 1.) Trace R2 is unusual, in

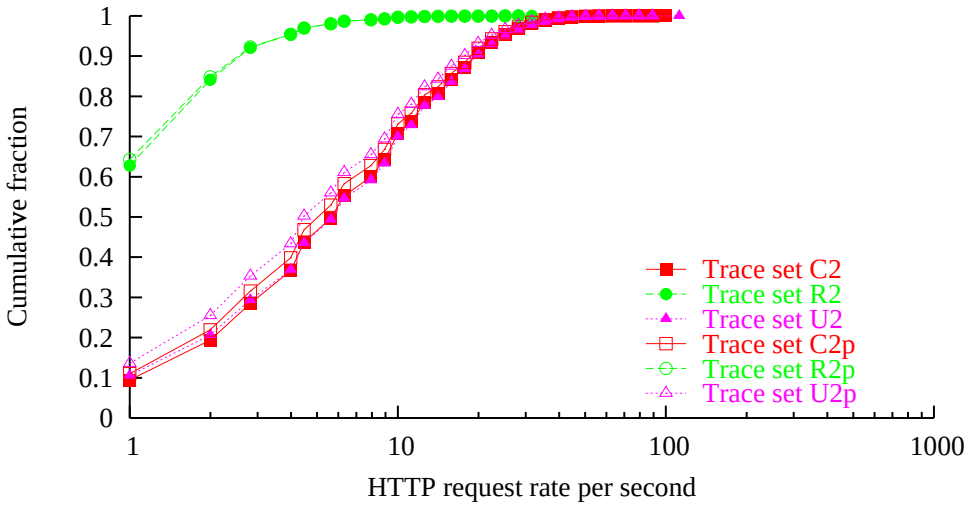

Figure 2: CDF of request rates

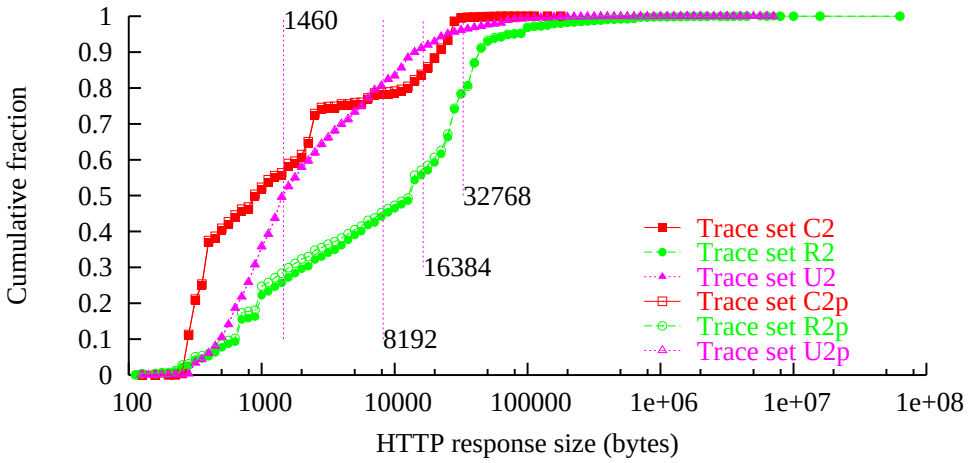

Figure 3: CDF of status-200 response sizes

part, because many users of the site download entire technical reports, which tend to be much larger than individual HTML or embedded-image files.

Figure 3 includes three vertical lines indicating the 8K byte, 16K byte, and 32K byte thresholds. Note that 8K is below the median size for R2, but above the median size for C2 and U2, but the median for all traces is well below 32K bytes.

Figure 4 shows that response durations are significantly longer in the R2 trace than in the others, possibly because of the longer response sizes in R2.

We calculated, for each distinct client, a mean bandwidth across all transfers for that client. Figure 6 shows the distributions; the pruned traces had similar distributions and are not shown. Trace C2 has a much larger fraction of low-bandwidth users than R2 or U2. The apparent slight excess of high-bandwidth clients in R2 might result from the larger responses in R2; larger transfers generally increase TCP's efficiency at using available bandwidth.

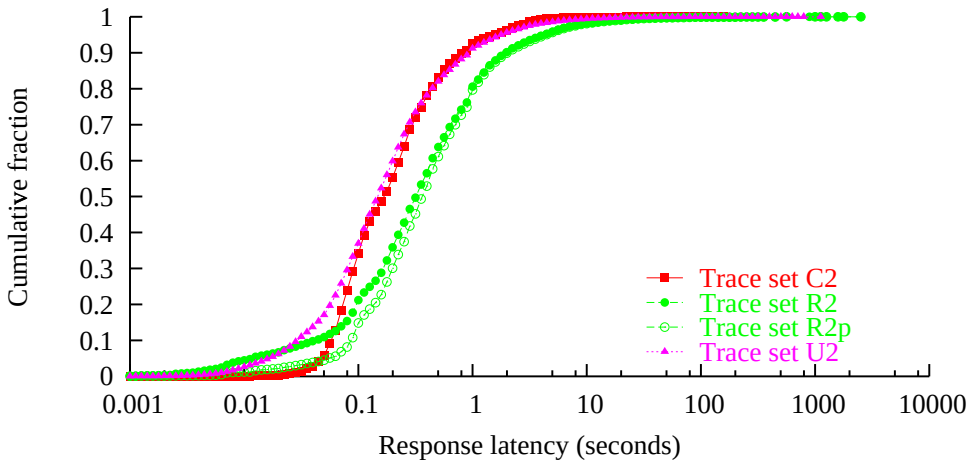

Figure 4: CDF of status-200 response latencies

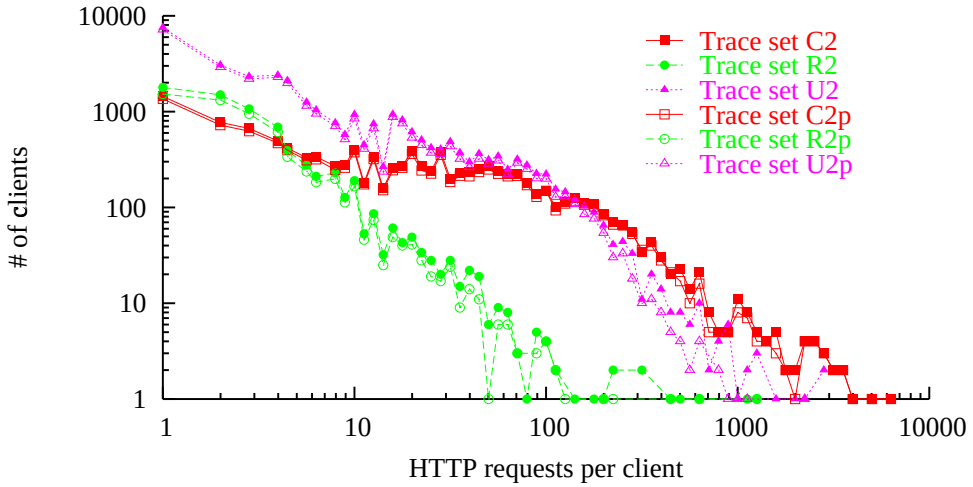

Figure 5: Histogram of request counts per client address

We also looked at the distribution of the TCP Maximum Segment Size (MSS) values in our traces. In trace R2, virtually all of the MSS values were at or close to the standard Ethernet limit (about 1460 bytes); in traces C2 and U2, about 95% of the MSS values were near the limit, with the rest mostly close to 512 bytes. Figure 3 shows a vertical line at 1460 bytes, indicating approximately where the dominant MSS value lies on the response-size distribution.

3.5 Trace anomalies

The monitoring architectures available to us differed at each of the collection sites. For example, at one of the sites port mirroring was used to copy packets from a monitored link to the mirrored link. At another site, separate links were tapped, one for packets bound for the Web server, the second for packets sent by the server. These monitoring infrastructures are subject to a variety of measurement errors:

- Port mirroring multiplexes bidirectional traffic from the monitored link onto the unidirectional mirror link. This can cause packets to appear in the trace in a different order than they arrived on the

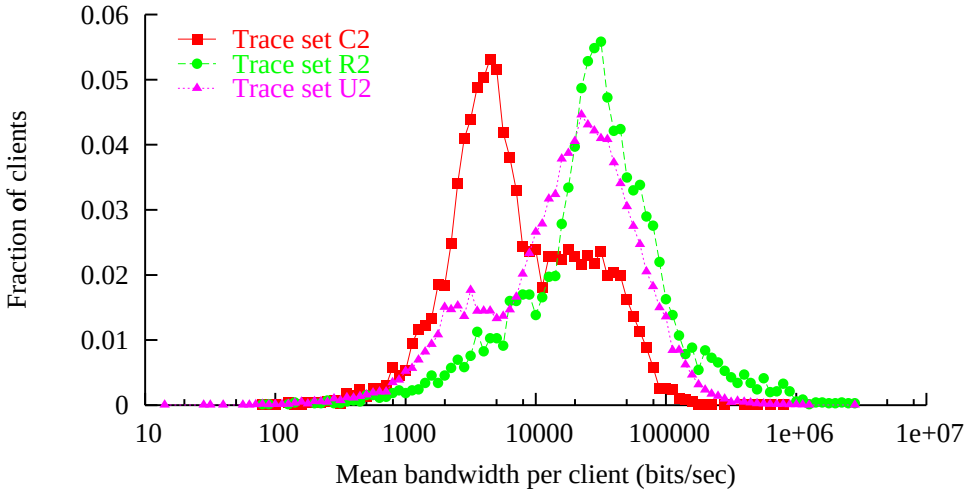

Figure 6: PDF of mean bandwidth per client

monitored link. Such reordering typically affects packets that occurred close together in time. For example, in the U2 trace, 10% of connections had the SYN and SYN|ACK packets in reverse order. Our Bro script corrects for this.

- Port mirroring temporarily buffers packets from the monitored link until they can be sent over the mirrored link. This buffer can overflow, causing packets to be dropped.
- Several of our environments have multiple network links that transfer packets to or from the Web server. Since we could not monitor all of these links, we did not capture all of the HTTP request/response transactions. In some cases we capture only half of the transaction (about 48% of the connections are affected by this in one trace).
- Ideally, a traced packet would be timestamped at the precise instant it arrives. However, trace-collection systems buffer packets at least briefly (often in several places) before attaching a timestamp, and packets are often collected at several nearby points (e.g., two packet monitors on both members of a pair of simplex links), which introduces timestamp errors due to imperfect clock synchronization. Erroneous timestamps could cause errors in our analysis by affecting either or both of our RTT estimates and our latency estimates.

Table 2: Packet loss rates

Trace name	Total packets	Total Conns.	Measurement system lost pkts.	Retransmitted packets	Conns. w/ retransmitted packets	Conns. w/no pkts in one direction
C2	40474900	1182499	17017 (0.04%)	114911 (0.3%)	53906 (4.6%)	572052 (48.4%)
R2	2824548	43023	1238 (0.04%)	27140 (1.0%)	4478 (10.4%)	460 (1.1%)
U2	11335406	313462	5611 (0.05%)	104318 (0.9%)	26815 (8.6%)	17107 (5.5%)
U3	11924978	328038	2093 (0.02%)	89178 (0.7%)	26371 (8.0%)	14975 (4.6%)
U4	14393790	384558	5265 (0.04%)	110541 (0.8%)	30638 (8.0%)	18602 (4.8%)

We estimated the number of packets lost within our measurement system by watching for gaps in the TCP sequence numbers. This could overestimate losses (e.g., due to reordered packets) but the estimates,

as reported in Table 2, are quite low.

Table 2 also shows our estimates (based on a separate Bro script) for packet retransmission rates on the path between client and server, implied by packets that cover part of the TCP sequence space we have already seen. Retransmissions normally reflect packet losses in the Internet, which would invalidate the model used in equation 1. Knowing these rates could help understand where the initial-RTT approach is applicable.

Note that Table 1 only includes connections with at least one complete HTTP response, while Table 2 includes all connections, including those that end in errors. We were only able to use 28% of the connections listed in Table 2 for C2, partly because we only saw packets in one direction for 48% of the connections. Our analysis script failed to reconstruct another ≈19% of the C2 connections due to gaps in the traced TCP data, possibly due to unknown problems in the monitoring infrastructure.

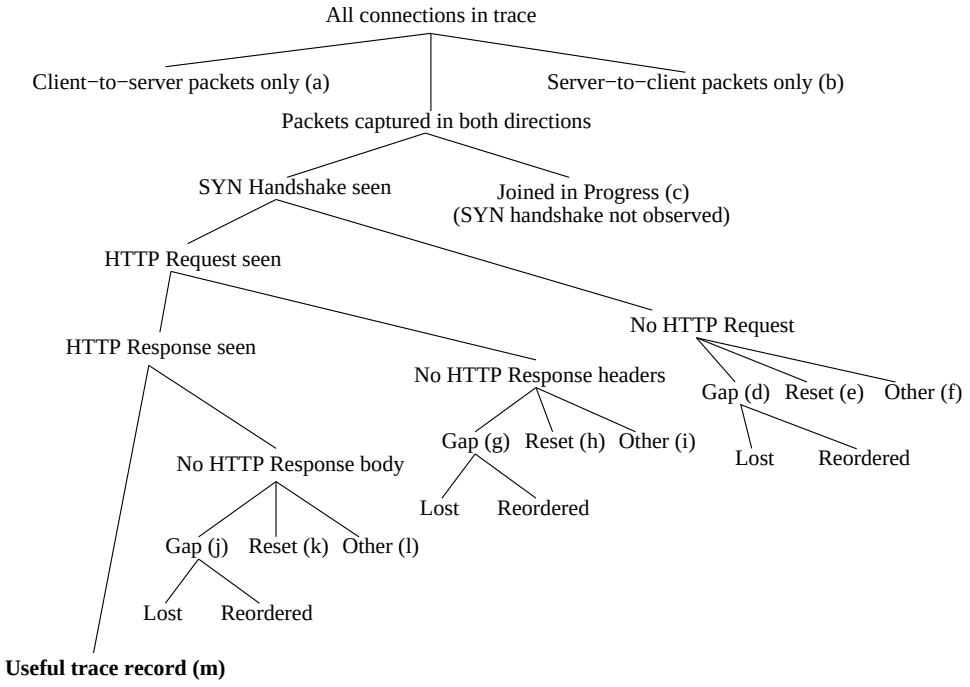

Figure 7: Classification tree for HTTP transactions in traces

Figure 7 illustrates the many ways in which we can fail to reconstruct a complete HTTP request-response transaction from our traces. We sometimes only capture packets in one direction (client-to-server or server-to-client) on a connection. If we capture packets in both directions, we might fail to observe the SYN exchange (perhaps because the connection started before the trace did). We might fail to see the HTTP request message, either because of a gap in the packet stream, a TCP Reset, or some other reason. If we see the request, we might still fail to see the HTTP response headers, or the HTTP response body, for similar reasons.

Table 3 quantifies these problems for each of the traces. Note that many of the tree nodes in Figure 7 are labelled with lower-case letters in parentheses; these labels are also shown for rows in Table 3. Major contributors to our failure to reconstruct complete HTTP request-response transactions are shown in bold. These include, as mentioned above, the large number of connections in trace C2 where we saw only server-to-client packets (row (b)), and those where we failed to see the response due to a gap in the trace (row

Table 3: Classification of reasons for transaction reconstruction failures

Class	Node label	C2	R2	U2	U3	U4
Client-to-server packets only	(a)	0.79 %	0.39 %	3.44 %	3.26 %	3.42 %
Server-to-client packets only	(b)	**47.81** %	0.34 %	2.65 %	1.65 %	1.84 %
Joined in progress	(c)	0.11 %	1.32 %	0.27 %	0.24 %	0.15 %
No HTTP Request						
Gap	(d)	0.65 %	0.00 %	0.08 %	0.01 %	0.02 %
Reset	(e)	0.37 %	0.66 %	1.43 %	1.35 %	1.49 %
Other	(f)	1.87 %	0.27 %	0.45 %	0.47 %	0.47 %
No HTTP Response headers						
Gap	(g)	**19.09** %	0.00 %	0.02 %	0.00 %	0.02 %
Reset	(h)	0.28 %	0.01 %	0.14 %	0.12 %	0.14 %
Other	(i)	0.87 %	0.06 %	0.05 %	0.05 %	0.04 %
No HTTP Response Body						
Gap	(j)	0.05 %	0.06 %	0.07 %	0.01 %	0.01 %
Reset	(k)	0.10 %	8.32 %	3.58 %	3.55 %	3.22 %
Other	(l)	0.00 %	0.00 %	0.00 %	0.00 %	0.00 %
Useful trace record	(m)	28.01 %	88.57 %	87.82 %	89.29 %	89.18 %

(g)).

The R2 trace includes a somewhat elevated number of connections that appeared in one of our traces while the connections were already in progress (row (c)). We believe these events result from half-closed connections where neither the client application nor the server's TCP stack ever time out. (This server's TCP stack does not appear to time out connections in the FIN_WAIT_2 state [1].)

4 Predictions based on initial RTT: results

In this section, we summarize the results of our experiments on techniques to predict transfer latency using the initial RTT. We address these questions:

1. Does RTT *per se* correlate well with latency?
2. How well does equation 1 predict latency?
3. Can we improve on equation 1?
4. What is the effect of modem compression?
5. How sensitive are the predictions to parameter choices?

There is no single way to define what it means for a latency predictor to provide "good" predictions. We evaluate prediction methods using several criteria, including the correlation between predicted and measured latencies, and the mean and median of the difference between the actual and predicted latencies.

4.1 Does RTT itself correlate with latency?

Perhaps it is unnecessary to invoke the full complexity of equation 1 to predict latency from RTT. To investigate this, we examined the correlation between RTT *per se* and either bandwidth or latency.

For example, Figure 8 shows a scatter plot of bandwidth vs. initial RTT, for all status-200 responses in trace C2. (In order to avoid oversaturating our scatter plots, we randomly sampled the actual data in each plot; the sampling ratios are shown in the figures.) The graph shows an apparent weak correlation between initial RTT and transfer bandwidth. Corresponding scatter plots for R2, U2, U3, and U4 show

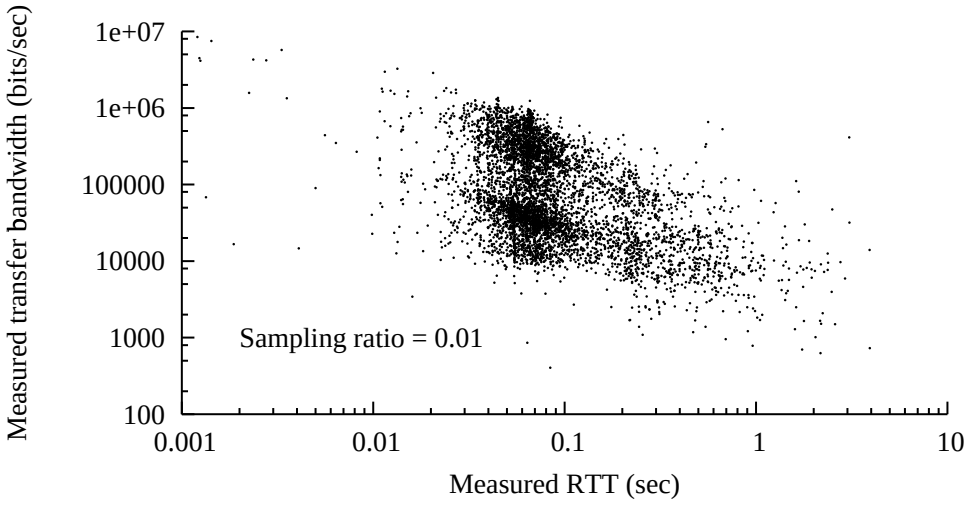

Figure 8: Scatter plot of bandwidth vs. RTT, trace C2

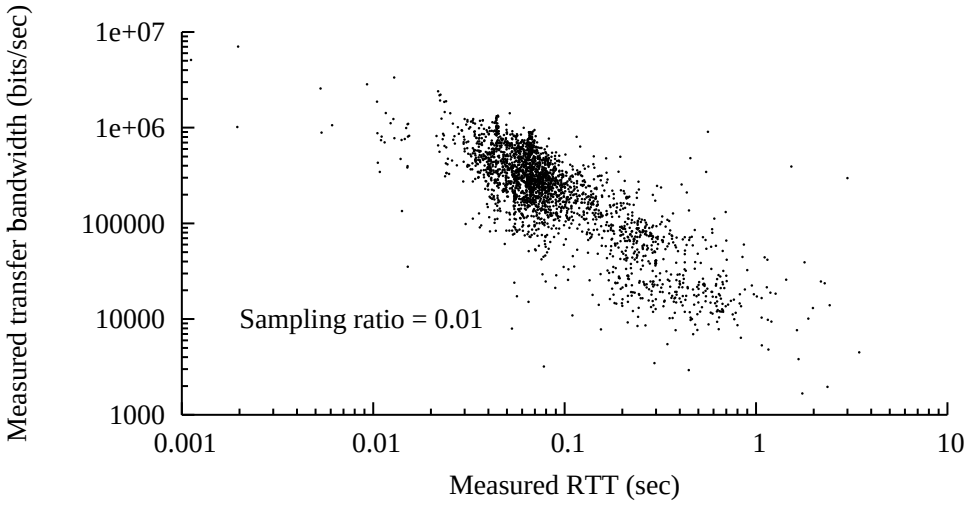

Figure 9: BW vs. RTT, trace C2, 1 MSS $<$ length $<$ 32KB

even weaker correlations.

We found a stronger correlation if we focused on transfer lengths above one MSS and below 32K bytes, as shown in Figure 9. Our technique for measuring latency is probably least accurate for responses below one MSS (i.e., those sent in just one packet). Also, single-packet responses may suffer excess apparent delay (as measured by when the server receives the final ACK) because of delayed acknowledgment at the client. In our subsequent analyses, we exclude responses with lengths of one MSS or less because of these measurement difficulties. The 32KB threshold represents one plausible choice for defining a "short" transfer.

For a more quantified evaluation of this simplistic approach, we did a statistical analysis using a simple R [22] program. The results are shown in Table 4(a) and (b), for lengths limited to 8K and 32K bytes, respectively.

Table 4: Correlations: RTT vs. either bandwidth or latency

Trace name	Samples included	Correlation w/bandwidth	Correlation w/latency
C2	140234 (24.3%)	-0.352	0.511
C2p	129661 (24.3%)	-0.370	0.508
R2	7500 (18.7%)	-0.112	0.364
R2p	5519 (17.6%)	-0.054	0.418
U2	218280 (37.6%)	-0.163	0.448
U2p	181180 (37.8%)	-0.178	0.458
U3	234591 (36.8%)	-0.181	0.421
U3p	181276 (36.9%)	-0.228	0.427
U4	283993 (37.2%)	-0.179	0.364
U4p	219472 (37.3%)	-0.233	0.411

(a) 1 MSS $<$ length $<$ 8KB

Trace name	Samples included	Correlation w/bandwidth	Correlation w/latency
C2	261931 (45.4%)	-0.325	0.426
C2p	238948 (44.8%)	-0.339	0.426
R2	20546 (51.4%)	-0.154	0.348
R2p	15407 (49.0%)	-0.080	0.340
U2	312090 (53.7%)	-0.165	0.392
U2p	258049 (53.8%)	-0.179	0.401
U3	336443 (52.8%)	-0.162	0.263
U3p	259028 (52.7%)	-0.215	0.276
U4	414209 (54.2%)	-0.167	0.287
U4p	320613 (54.4%)	-0.215	0.343

(b) 1 MSS $<$ length $<$ 32KB

The tables show rows for both pruned and unpruned versions of the five basic traces. We included only status-200 responses whose length was at least one MSS; the "samples included" column shows that count for each trace. The last two columns show the computed correlation between initial RTT and either transfer bandwidth or transfer latency. (The bandwidth correlations are negative, because this is an inverse relationship.)

For the data set including response lengths up to 32K bytes, none of these correlations exceeds 0.426, and many are much lower. If we limit the response lengths to 8K bytes, the correlations improve, but this also eliminates most of the samples.

We tried excluding samples with an initial RTT value above some quantile, on the theory that high RTTs correlate with lossy network paths; this slightly improves RTT vs. bandwidth correlations (for example, excluding records with an RTT above 281 msec reduces the number of 32K-or-shorter samples for R2 by 10%, and improves that correlation from -0.154 to -0.302) but it actually worsens the latency correlations (for the same example, from 0.348 to 0.214).

Note that, contrary to our expectation that traces pruned of proxies and robots would be less predictable, in Table 4 this seems true only for the R2 trace; in general, pruning seems to slightly improve predictability. In fact, while we present results for both pruned and unpruned traces throughout the paper, we see no consistent difference in predictability.

4.2 Does equation 1 predict latency?

Although we did not expect RTT to correlate well with latency, we might expect better results from the sophisticated model derived by Cardwell *et al.* [2]. They validated their model (equation 1 is a simplified version) using HTTP transfers over the Internet, but apparently used only "well-connected" clients and so did not probe its utility for poorly-connected clients. They also used RTT estimates that included more samples than just each connection's initial RTT.

We therefore analyzed the ability of equation 1 to predict transfer bandwidths and latencies using only the initial RTT, and with the belief that our traces include some poorly-connected clients.

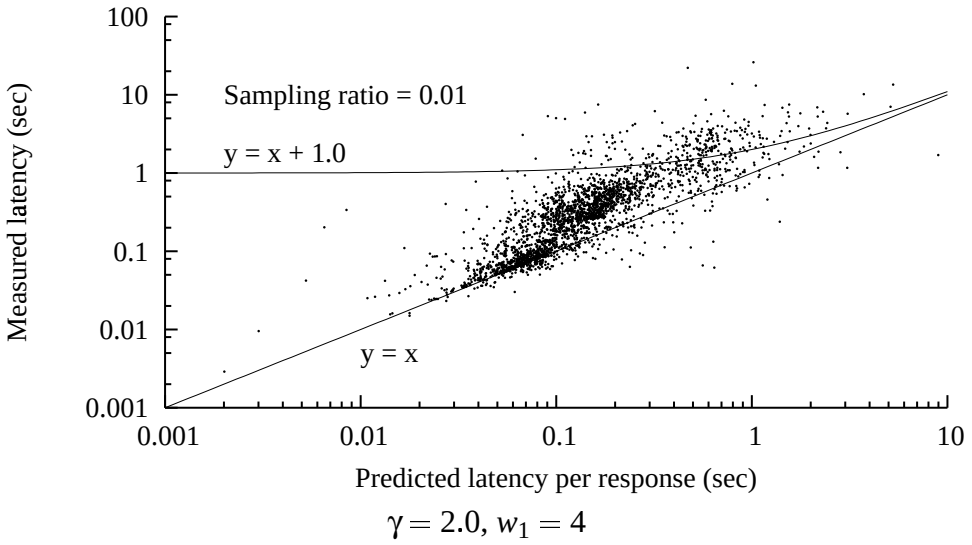

Figure 10: Real vs. predicted latency, trace C2

Figure 10 shows an example scatter plot of measured latency vs. predicted latency, for trace C2. Again, we include only status-200 responses at least one MSS in length. We have superimposed two curves on the plot. (Since this is a log-log plot, most linear equations result in curved lines.) Any point above the line $y = x$ represents an underprediction of latency; underpredictions are generally worse than overpredictions, if (for example) we want to avoid exposing Web users to unexpectedly long downloads. Most of the points in the plot are above that line, but most are below the curve $y = x + 1.0sec$, implying that most of the overpredictions (in this example) are less than 1 sec in excess. However, a significant number are many seconds too high.

We extended our R program to compute statistics for the predictive ability of equation 1. For each status-200 trace record with a length between one MSS and 32K bytes, we used the equation to predict a latency, and then compared this to the latency recorded in the trace record. We then computed the correlation between the actual and predicted latencies. We also computed a residual error value, as the difference between the actual and predicted latencies. Table 5 shows the results from this analysis, using $\gamma = 1.5$ and $w_1 = 1$, a parameter assignment that worked fairly well across all five traces.

In Table 5, the median residuals are always negative, implying that equation 1 overestimates the transfer latency more often than it underestimates it. However, the mean residuals are always positive, because the equation's underestimates are more wrong (in absolute terms) than its overestimates. The samples in Figure 10 generally follow a line with a steeper slope than $y = x$, suggesting that equation 1 especially underestimates higher latencies.

Table 5: Quality of predictions based on equation 1

Trace name	Samples included	Correlation w/latency	Median residual	Mean residual
C2	261931 (45.4%)	0.581	-0.017	0.164
C2p	238948 (44.8%)	0.584	-0.015	0.176
R2	20546 (51.4%)	0.416	-0.058	0.261
R2p	15407 (49.0%)	0.421	-0.078	0.272
U2	312090 (53.7%)	0.502	-0.022	0.110
U2p	258049 (53.8%)	0.519	-0.024	0.124
U3	336443 (52.8%)	0.334	-0.018	0.152
U3p	259028 (52.7%)	0.353	-0.016	0.156
U4	414209 (54.2%)	0.354	-0.013	0.141
U4p	320613 (54.4%)	0.425	-0.010	0.136

Residual values are measured in seconds; 1 MSS $<$ length $<$ 32KB

One possible reason is that, for lower-bandwidth links, RTT depends on packet size. For a typical 56Kb/s modem link, a SYN packet will see an RTT somewhat above 100 msec, while a 1500 byte data packet will see an RTT several times larger. This effect could cause equation 1 to underestimate transfer latencies.

4.3 Can we improve on equation 1?

Given that equation 1 seems to systematically underestimate higher latencies, exactly the error that we want to avoid, we realized that we could modify the equation to reduce these errors.

We experimented with several modifications, including a linear multiplier, but one simple approach is:

```
function ModifiedEqnOne(RTT, MSS, Length,w1, γ, CompWeight)
   temp = EquationOne(RTT, MSS, Length, w1, γ);
   return(temp + (temp*temp*CompWeight));
```

That is, we "overpredict" by a term proportional to the square of the original prediction. This is a *heuristic*, not the result of rigorous theory.

We found by trial and error that a proportionality constant, or "compensation weight," $CompWeight = 2.25$ worked best for C2, but $CompWeight = 1.75$ worked better for R2 and U2, and $CompWeight = 1.25$ worked best for U3 and U4. For all traces, $\gamma = 2$ got the best results, and we set $w_1 = 4$ for C2 and U2, and $w_1 = 3$ for R2, U3, and U4. We discuss the sensitivity to these parameters in Section 4.5.

Figure 11 shows how the modified prediction algorithm systematically overpredicts at higher latencies, while not significantly changing the accuracy for lower latencies. (For example, in this figure, $CompWeight = 2.25$; if equation 1 predicts a latency of 0.100 seconds, the modified prediction will be 0.1225 seconds). However, even the modified algorithm significantly underpredicts a few samples; we do not believe we can avoid this, especially for connections that suffer packet loss (see Table 2).

Table 6 shows that the modifications to equation 1 generally worsen the correlations, compared to those in Table 5, but definitely improves the residuals – the median error is always less than 100 msec, and the mean error is less than 15 msec, except for traces U3p and U4p (our parameter choices were tuned for the unpruned traces).

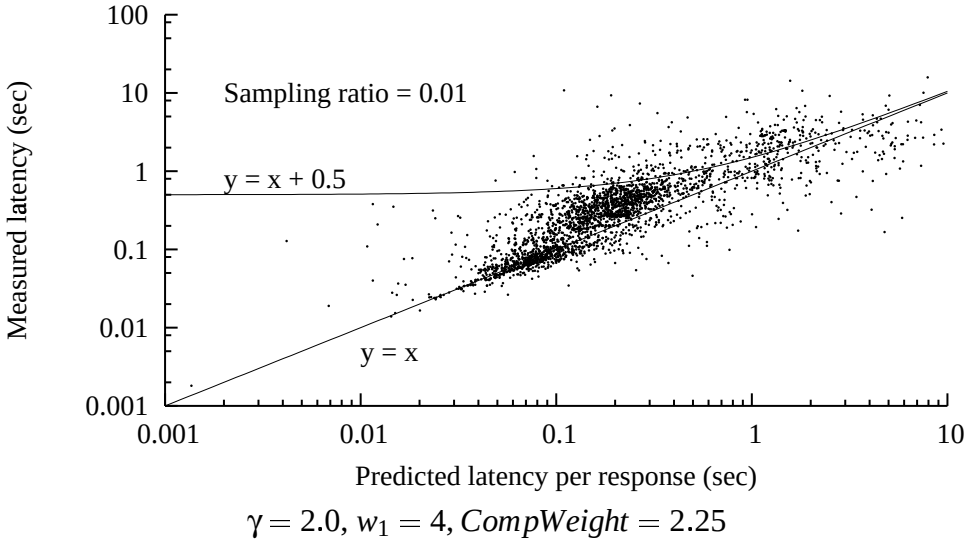

Figure 11: Modified prediction results, trace C2

Table 6: Predictions based on modified equation 1

Trace name	Samples included	Correlation w/latency	Median residual	Mean residual
C2	261931 (45.4%)	0.417	0.086	-0.002
C2p	238948 (44.8%)	0.423	0.092	-0.006
R2	20546 (51.4%)	0.278	0.015	0.002
R2p	15407 (49.0%)	0.311	0.019	0.013
U2	312090 (53.7%)	0.386	0.053	0.010
U2p	258049 (53.8%)	0.402	0.056	0.001
U3	336443 (52.8%)	0.271	0.034	0.011
U3p	259028 (52.7%)	0.302	0.036	-0.020
U4	414209 (54.2%)	0.279	0.035	0.003
U4p	320613 (54.4%)	0.337	0.038	-0.033

Residual values are measured in seconds; 1 MSS $<$ length $<$ 32KB

4.4 Text content and modem compression

Many people still use dialup modems. It has been observed that to accurately model path bandwidth, one must account for the compression typically done by modems [3]. However, most image Content-Types are already compressed, so this correction should only be done for text content-types.

HTTP responses normally carry a MIME Content-Type label, which allowed us to analyze trace subsets for "text/*" and "image/*" subsets. Table 7 shows the distribution of these coarse Content-Type distinctions for the traces.

We speculated that the latency-prediction model of equation 1, which incorporates the response length, could be further improved by reducing this length value when compression might be expected. (A server making predictions knows the Content-Types of the responses it plans to send. Some servers might use a compressed content-coding for text responses, which would obviate the need to correct predictions for those responses for modem compression. We found no such responses in our traces.)

Table 7: Counts and frequency of content-types (excluding some rarely-seen types)

Content-type	C2	R2	U2	U3	U4
Unknown	3 (0.00%)	26 (0.06%)	178 (0.03%)	157 (0.02%)	144 (0.02%)
TEXT/*	122426 (21.22%)	23139 (57.83%)	85180 (14.67%)	92108 (14.45%)	107958 (14.14%)
IMAGE/*	454458 (78.78%)	13424 (33.55%)	465160 (80.10%)	507330 (79.60%)	607520 (79.57%)
APPLICATION/*	0 (0.00%)	3410 (8.52%)	29733 (5.12%)	37581 (5.90%)	47765 (6.26%)
VIDEO/*	0 (0.00%)	4 (0.01%)	17 (0.00%)	10 (0.00%)	5 (0.00%)
AUDIO/*	0 (0.00%)	8 (0.02%)	446 (0.08%)	194 (0.03%)	140 (0.02%)

We cannot directly predict either the compression ratio (which varies among responses and among modems) nor can we reliably determine which clients in our traces used modems. Therefore, for feasibility of analysis our model assumes a constant compressibility factor for text responses, and we tested several plausible values for this factor. Also, we assumed that an RTT below 100 msec implied a non-modem connection, and RTTs above 100 msec implied the *possible* use of a modem. In a real system, information derived from the client address might identify modem-users more reliably. (In Section 5 we classify clients using hostnames; but this might add too much DNS-lookup delay to be effective for latency prediction. Even a caching system such as Rapid DNS [15] probably would not help with classification latency for the initial connection.)

Table 8: Predictions for text content-types only

Trace name	Samples included	Correlation w/latency	Median residual	Mean residual
C2	118217 (96.6%)	0.442	0.142	0.002
C2p	106120 (96.4%)	0.449	0.152	-0.003
R2	12558 (54.3%)	0.288	0.010	0.066
R2p	8760 (50.2%)	0.353	0.017	0.105
U2	70924 (83.3%)	0.292	0.100	0.073
U2p	56661 (83.0%)	0.302	0.110	0.066
U3	76714 (83.3%)	0.207	0.063	-0.021
U3p	56070 (83.2%)	0.198	0.072	-0.099
U4	90416 (83.8%)	0.281	0.065	-0.034
U4p	65708 (83.8%)	0.359	0.078	-0.122

Residual values are measured in seconds; 1 MSS $<$ length $<$ 32KB

Table 8 shows results for text content-types only, using the modified prediction algorithm based on equation 1, but without correcting for possible modem compression. We set $\gamma = 2.0$ for C2 and U2, and $\gamma = 1.5$ for R2, U3, and U4; $w_1 = 2$ for C2 and $w_1 = 3$ for the other traces; and $CompWeight = 1$ for all traces. (We have not tested a wide range of *CompWeight* values to see if text content-types would benefit from a different *CompWeight*.) Compared to the results for all content types (see Table 6), the residuals for text-only samples are generally higher.

Table 9 shows results for text content-types when we assumed that modems compress these by the factor shown in the third column. Note that for C2 and C2p, we got the best results using a compression factor of 1.0 – that is, without correcting for compression. For the other traces, correcting for compression did give

Table 9: Predictions for text with compression

Trace name	Samples included	Compression factor	Correlation w/latency	Median residual	Mean residual
C2	118217	1.0	0.442	0.142	0.002
C2p	106120	1.0	0.449	0.152	-0.003
R2	12558	4.0	0.281	0.013	0.002
R2p	8760	4.0	0.345	0.021	0.044
U2	70924	3.0	0.295	0.083	0.008
U2p	56661	3.0	0.306	0.096	-0.004
U3	76714	4.0	0.208	-0.002	0.001
U3p	56070	4.0	0.201	0.003	-0.063
U4	90416	4.0	0.277	-0.000	-0.011
U4p	65708	4.0	0.353	0.007	-0.083

Residual values are measured in seconds; 1 MSS $<$ length $<$ 32KB

better results. Here we set the other parameters as: $\gamma = 2$ (except for U3 and U4, where $\gamma = 1.5$ worked best), $w_1 = 1$ (except for C2, where $w_1 = 2$ worked best), and *CompWeight* $= 1.0$ (except for R2, where *CompWeight* $= 2.25$ worked best). We experimented with assuming that the path did not involve a modem (and thus should not be corrected for compression) if the initial RTT was under 100 msec, but for R2 and U2 it turned out that we got the best results when we assumed that all text responses should be corrected for compression.

Table 9 shows that, except for trace C2, correcting for modem compression improves the mean residuals over those in Table 8. We have not evaluated the use of compression factors other than integers between 1 and 4, and we did not evaluate a full range of *CompWeight* values for this section.

Table 10: Predictions for image content-types only

Trace name	Samples included	Correlation w/latency	Median residual	Mean residual
C2	143714 (31.6%)	0.332	0.040	0.066
C2p	132828 (31.4%)	0.329	0.041	0.067
R2	6259 (46.6%)	0.275	0.017	0.046
R2p	5096 (46.0%)	0.315	0.018	0.032
U2	216590 (46.6%)	0.445	0.040	0.023
U2p	181447 (46.9%)	0.459	0.041	0.019
U3	228097 (45.0%)	0.336	0.025	-0.008
U3p	178861 (45.2%)	0.377	0.026	-0.028
U4	283029 (46.6%)	0.308	0.026	-0.012
U4p	223721 (47.2%)	0.339	0.027	-0.035

Residual values are measured in seconds; 1 MSS $<$ length $<$ 32KB

Image content As shown in Table 7, image content-types dominate most of the traces, except for R2. Also, Web site designers are more likely to have choices between rich and simple content for image types than for text types. (Designers often include optional "Flash" animations, but we found almost no Flash

content in C2 and R2, and relatively little in U2, U3, and U4.) We therefore compared the predictability of transfer latencies for image content-types, but found no clear difference compared to the results for all content in general. Table 10 shows the results for image content-types, using the same γ, w_1, and *CompWeight* settings as used for Table 6. We have not evaluated whether a different set of parameter values would give better results for image content-types.

4.5 Sensitivity to parameters

How sensitive is prediction performance to the parameters γ, w_1, and *CompWeight*? That question can be framed in several ways: how do the results for one server vary with parameter values? If parameters are chosen based on traces from server X, do they work well for server Y? Are the optimal values constant over time? Do optimal parameter values depend on the performance metric? (We have not evaluated other similar questions, such as whether the optimal values are constant over client sub-population, content-type, or response length.)

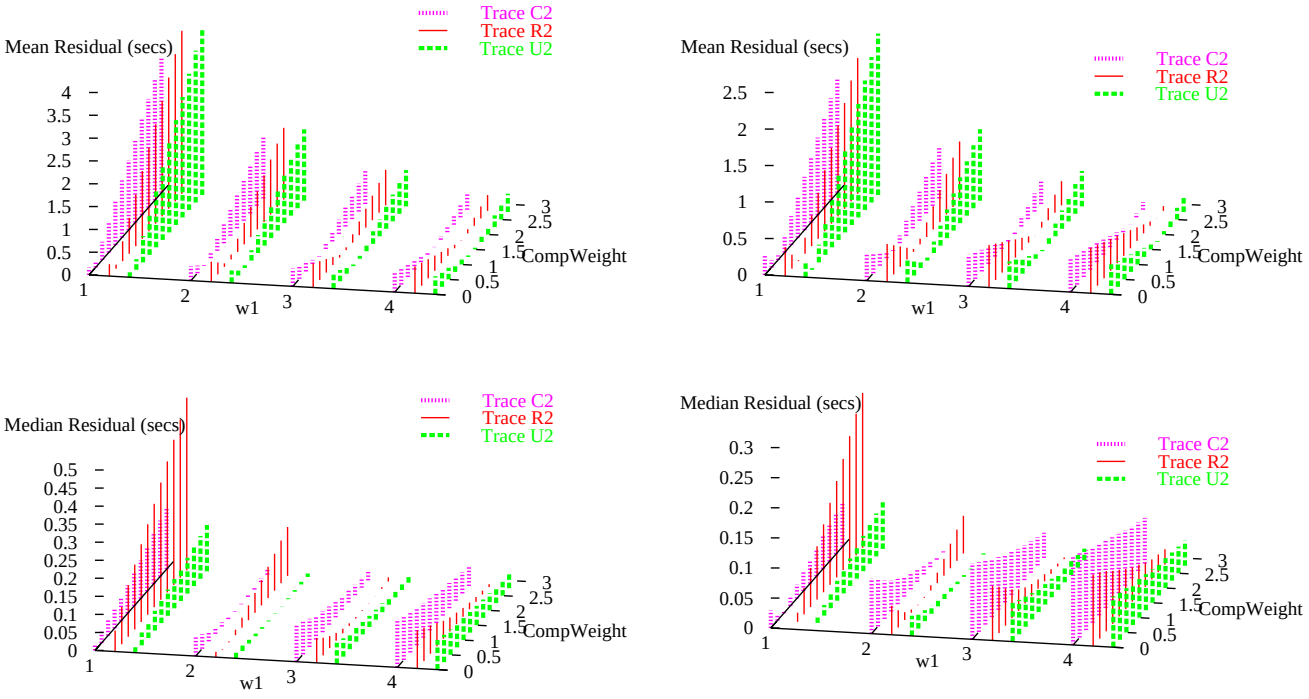

Mean and median for $\gamma = 1.5$

Mean and median for $\gamma = 2.0$

1 MSS $<$ length $<$ 32KB (note different y-axis scales)

Figure 12: Sensitivity of residual absolute values to parameters (smaller is better)

Figure 12 shows how the absolute values of the mean and median residuals vary with γ, w_1, and *CompWeight* for traces C2, R2, and U2. The optimal parameter choice depends on whether one wants to minimize the mean or the median; for example, for R2, $\gamma = 2.0$, $w_1 = 3$, and *CompWeight* $= 1.75$ yields an optimal mean of 1.5 msec (and a median of 15 msec). The median can be further reduced to 0.2 msec, but at the cost of increasing the mean to over half a second.

Figure 12 also shows how the optimal parameters vary across several traces. (Results for traces U3 and U4 are similar to those for U2, and are omitted to reduce clutter.) It appears that no single choice is optimal across all traces, although some choices yield relatively small mean and medians for many traces.

For example, $\gamma = 2$, $w_1 = 3$, and *CompWeight* $= 1.25$ yields optimal or near-optimal mean residuals for U2, U3, and U4, and decent results for C2.

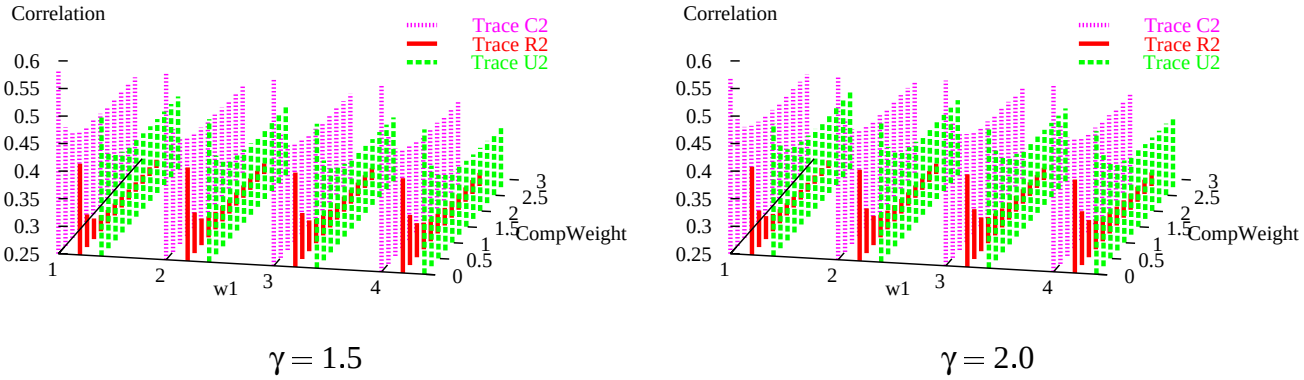

$\gamma = 1.5$ $\gamma = 2.0$

1 MSS $<$ length $<$ 32KB

Figure 13: Sensitivity of correlation values to parameters (larger is better)

Instead of optimizing the parameter choices for lowest mean or median residuals, one could optimize for highest correlation between predicted and actual latencies. Figure 13 shows how the correlations vary with γ, w_1, and *CompWeight* for traces C2, R2, and U2. As we mentioned in Section 4.3, a non-zero *CompWeight* improves the residuals but generally worsens the correlations; this effect on correlations is obvious in Figure 13. The correlations are not particular sensitive to changes in the other parameters, γ and w_1.

4.6 Training and testing on different data

The results we have presented so far used parameter choices "trained" on the same data sets as our results were tested on. Since any real prediction system requires advance training, we also evaluated predictions with training and testing on different data sets.

Our trace collection was not carefully designed in this regard; we have no pairs of data sets that are completely identical and adjacent in time. For the C2, R2, and U2 data sets, we chose the first three days as the training data set, and the last four days as the testing data set. However, because we collected data at different hours on each day, and because there are day-of-week differences between the training and testing sets (the testing sets includes two weekend days), we suspect that these pairs of data sets might not be sufficiently similar. We also used the U3 data set to train parameters that we then tested on the U4 data set; these two traces are more similar to each other.

Table 11 shows results for training vs. testing. We tested and trained with 96 parameter combinations, based on the two possible choices for γ, the four choices for w_1, and twelve equally-spaced choices for *CompWeight*. The **trained parameters** are those that minimize the absolute value of the mean **residual in training**. The columns under **testing results** show how the results using the trained parameters rank among all of the testing results, the mean residual when using those parameters, and the residual for the best possible parameter combination for the testing data.

These results suggest that the degree to which training can successfully select parameter values might vary significantly from site to site. Based on our traces, we would have had the most success making useful predictions at the University site (U3-U4), and the least success at the Research site (R2).

However, the difference in "trainability" that we observed might instead be the result of the much closer match between the U3 and U4 datasets, compared to the time-of-day and day-of-week discrepancies in

Table 11: Training and testing on different data

	Trained parameters			Residual	Testing results		
Trace name	γ	w_1	*CompWeight*	in training	Rank (of 96)	Mean residual w/those params.	Best residual
C2	2.0	4	2.50	-0.000	15	-0.098	-0.004
C2p	2.0	3	1.75	-0.004	12	-0.089	-0.002
R2	1.5	4	1.50	-0.004	20	0.136	0.000
R2p	1.5	3	1.00	0.003	16	0.125	0.003
U2	1.5	4	1.50	0.001	10	-0.072	0.012
U2p	2.0	2	0.75	-0.004	9	-0.081	-0.002
U3U4	2.0	2	0.75	-0.007	3	-0.013	0.003
U3U4p	2.0	1	0.25	0.000	2	-0.013	-0.010

Residual values are measured in seconds; 1 MSS $<$ length $<$ 32KB

Table 12: Training and testing on different data (Tuesday vs. Wednesday)

	Trained parameters			Residual	Testing results		
Trace name	γ	w_1	*CompWeight*	in training	Rank (of 96)	Mean residual w/those params.	Best residual
C2	1.5	4	1.75	-0.008	8	0.037	0.001
C2p	1.5	4	1.75	-0.001	7	0.033	-0.001
R2	2.0	2	1.25	0.000	15	0.093	0.003
R2p	2.0	4	2.50	0.009	24	0.169	0.000
U2	2.0	4	2.00	0.006	6	-0.028	-0.001
U2p	2.0	4	2.00	0.000	6	-0.035	-0.004

Residual values are measured in seconds; 1 MSS $<$ length $<$ 32KB

the other train/test comparisons. For C2, R2, and U2, we tried training just on one day (Tue., May 4, 2004) and testing on the next day; see Table 12 for the results. These show significantly better trainability (except for R2p, which was slightly worse) than in Table 11, which supports the need to match training and testing data sets more carefully.

4.7 Predictions vs. a threshold

Statistical analysis gives intuition into the numeric accuracy of our model, but it does not answer the question "will a server using this model make the right decisions?" We therefore tested the prediction accuracy by choosing a variety of latency thresholds, and then simulating the modified equation 1 model to see if it correctly predicted whether a measured value was above or below the threshold. Figure 14 shows the results, using the same γ, w_1, and *CompWeight* settings as used for Table 6.

In each graph in Figure 14, the x-axis shows a range of threshold values, and the y-axis shows the percentage of accurate predictions. The three curves show (1) the success rate for all predictions, (2) the success rate for predictions where the true (measured) value was above the threshold, and (3) the success rate for predictions where the true value was below the threshold.

As we argued earlier, a server probably wants to err on the side of sending smaller responses (to avoid forcing a user to wait), so we focus on the "true latency was above threshold" curve. For small thresholds, such predictions are good for all three traces (e.g., at a threshold of 100 msec, predictions range from 69% correct for U2 to 92% correct for R2). At higher thresholds, prediction quality drops (e.g., at a threshold

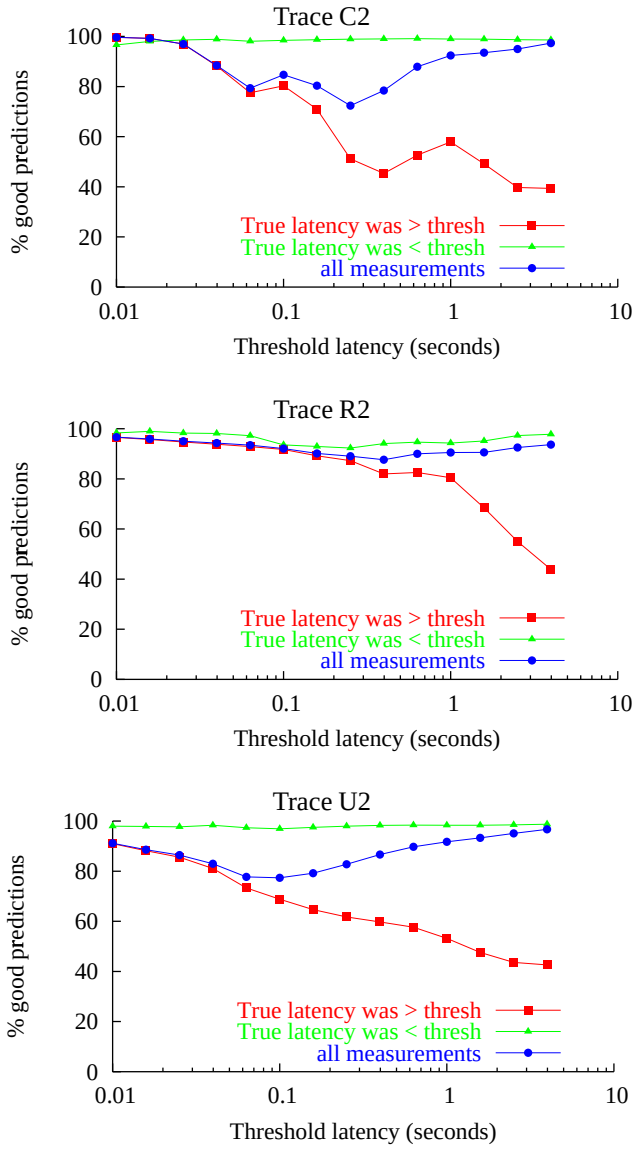

Figure 14: Predictions vs. threshold, modified eqn. 1

of 1 sec., predictions range from 53% correct for U2 to 80% correct for R2). However, the graphs show that certain thresholds yield lower accuracies, depending on the trace.

Predictions for the first contact with any given client are slightly more accurate than the predictions for all transfers, as can be seen by comparing the "true latency was above threshold" curves in Figure 14 with those in Figure 18.

Our approach almost always yields accurate predictions for responses whose true latency was below the threshold – always better than 92% correct in these three examples.

4.8 A server's decision algorithm

To understand how a server might use the initial-RTT approach in practice, Figure 15 presents pseudo-code for generating predictions. (This example is in the context of a Web server adapting its content based on predicted transfer latency, but the basic idea should apply to other contexts.) If the server has $N > 1$ choices of response length for a given request, it would invoke *PredictLatency* $N - 1$ times, starting with

the largest candidate and moving down in size, until it either finds one with a small-enough predicted latency, or has only one choice left. The first three arguments to the *PredictLatency* function (RTT, MSS, and client IP address) are known as soon as the connection is open. The last two (response content type and length) are specific to a candidate response that the server might send.

```
1.  function
        PredictLatency(RTT, MSS, ClientIP, ContentType, Length)

2.  if (ProbablyDialup(ClientIP, RTT)
          and (ContentType == TEXT)) then
3.    effectiveLength := Length/TextCompressionFactor;
4.  else
5.    effectiveLength := Length;
6.  end

7.  if (length > maxPredictableLength) then
8.    return(NO_PREDICTION); /* probably leaves slow-start */
9.  else if (length < MSS) then
10.   return(NO_PREDICTION); /* only one data packet to send */
11. end

12. return(ModifiedEqnOne(RTT, MSS, Length, w1, γ,
          CompWeight));
```

TextCompressionFactor is an estimate of the mean compression ratio for modems on text files;
CompWeight. w_1, and γ could themselves vary based on the server's observation of recent history, the ContentType, etc.

Figure 15: Pseudo-code for the decision algorithm

The function *ProbablyDialup*, not shown here, is a heuristic to guess whether a client is connected via a modem (which would probably compress text responses). It could simply assume that RTTs above 100 msec are from dialups, or it could use additional information based on the client's DNS name or AS (Autonomous System) number to identify likely dialups.

5 Detecting dialups

We speculated that a server could discriminate between dialups and non-dialups using clues from the client's "fully-qualified domain name" (FQDN). We obtained FQDNs for about 75% of the clients in the U4 trace, and then grouped them according to clues in the FQDNs that implied geography and network technology. Note that many could not be categorized by this method, and some categorizations are certainly wrong.

Table 13 shows how initial RTTs vary by geography and connection type. For the connections that we could categorize, at least 95% of "dialup" connections have RTTs above 100 msec, and most "cable" and "DSL" connections have RTTs below 200 msec. These results seem unaffected by further geographical subdivision, and support the hypothesis that a threshold RTT between 100 and 200 msec would discriminate fairly well between dialup and non-dialup connections. We do not know if these results apply to other

Table 13: RTTs by geography and connection type

Category	Connections	5th percentile	median	mean	95th percentile
By geography					
All	326359	0.008	0.069	**0.172**	**0.680**
N. America	35972	0.003	0.068	**0.124**	**0.436**
S. America	2372	**0.153**	**0.229**	**0.339**	**0.882**
Europe	12019	**0.131**	**0.169**	**0.262**	**0.717**
Asia-Pacific	9176	**0.165**	**0.267**	**0.373**	**0.885**
Africa	2027	**0.206**	**0.370**	**0.486**	**1.312**
”Dialup” in FQDN					
All	11478	**0.144**	**0.350**	**0.664**	**2.275**
Regional	5977	**0.133**	**0.336**	**0.697**	**2.477**
Canada	1205	**0.208**	**0.460**	**0.751**	**2.060**
US	575	**0.189**	**0.366**	**0.700**	**2.210**
Europe	566	**0.183**	**0.216**	**0.313**	**0.861**
”DSL” in FQDN					
All	59211	0.003	0.023	0.060	**0.210**
Local	1816	0.011	0.022	0.034	0.085
Regional	47600	0.009	0.018	0.032	0.079
US	1053	0.071	0.085	**0.117**	**0.249**
Europe	118	**0.148**	**0.162**	**0.178**	**0.313**
”Cable” in FQDN					
All	6599	0.039	0.077	**0.132**	**0.338**
Canada	2741	0.039	0.055	0.088	**0.222**
US	585	0.072	0.086	0.094	**0.127**
Europe	600	**0.143**	**0.155**	**0.176**	**0.244**

Times in seconds; **bold** entries are > 0.1 sec.

traces.

Previously, Wei *et al.* [26] classified access network types using interarrival times from packet pairs. They concluded that it is hard to distinguish between dialups and other “low-bandwidth, non-WLAN” connections. However, because they used one-way probes, they had no RTT information and so would not have seen the characteristic 100–200 msec RTT signature we identified.

6 Predictions from previous bandwidths: results

In this section, we compare how well prediction based on variants of equation 1 compares with predictions from the older recent-transfers approach. We address these questions:

1. How well can we predict latency from previous bandwidth measurements?
2. Does a combination of the two approaches improve on either individual predictor?

Note that the recent-transfers approach cannot specifically predict the latency for the very first transfer to a given client, because the server has no history for that client. This is a problem if the goal is to provide the best user experience for a client's initial contact with a Web site. For initial contacts, a server using the recent-transfers approach to predict latency has several options, including:

- Make no prediction.
- “Predict” the latency based on history across all previous clients; for example, use an exponentially

smoothed mean of all previous transfer bandwidths.

- Assume that clients with similar network locations, based on routing information, have similar bandwidths; if a new client belongs to "cluster" of clients with known bandwidths, use history from that cluster to make a prediction. Krishnamurthy and Wang [10] describe a technique to discover clusters of client IP addresses. Krishnamurthy and Wills [11] then showed, using a set of chosen Web pages with various characteristics, that clustering pays off in prediction accuracy improvements ranging up to about 50%. We speculate that this approach would also work for our traces.
- Use the initial-RTT technique to predict a client's first-contact latency, and use the recent-transfers technique to predict subsequent latencies for each client. We call this the *hybrid* technique.

We first analyze the purest form of recent-transfers (making no prediction for first-contact clients), and then consider the mean-of-all-clients and hybrid techniques.

6.1 Does previous bandwidth predict latency?

Table 14: Correlations: measured vs. recent bandwidths

Trace name	Samples included	Correlation with most recent bandwidth	Correlation with mean previous bandwidth	Correlation with weighted mean bandwidth
C2	262165 (45.4%)	0.674	0.742	**0.752**
C2p	238957 (44.8%)	0.658	0.732	**0.737**
R2	24163 (60.4%)	0.589	0.655	**0.666**
R2p	17741 (56.5%)	0.522	0.543	**0.579**
U2	310496 (53.5%)	0.527	0.651	**0.654**
U2p	254024 (52.9%)	0.437	**0.593**	0.561
U3	341968 (53.7%)	0.495	0.627	**0.638**
U3p	260470 (53.0%)	0.508	**0.659**	0.625
U4	421867 (55.3%)	0.521	**0.690**	0.647
U4p	323811 (55.0%)	0.551	**0.690**	0.656

Best correlation for each trace shown in **bold**

We did a statistical analysis of the prediction ability of several variants of the pure recent-tranfers technique. In each case, we made predictions and maintained history only for transfer lengths of at least one MSS. Table 14 shows the results. The first two columns show the trace name and the number of samples actually used in the analysis. The next three columns show the correlations between the bandwidth (not latency) in a trace record and, respectively, the most recent bandwidth for the same client, the mean of previous bandwidths for the client, and the exponential weighted mean $X_i = \alpha \cdot X_{i-1} + (1-\alpha)measurement_i$. We followed Krishnamurthy *et al.* [12] in using $\alpha = 0.7$, although other values might work better for specific traces.

These results suggest that some form of mean is the best variant for this prediction technique; although the choice between simple means and weighted means varies between traces, these always outperform predictions based on just the most previous transfer. Since Krishnamurthy *et al.* [12] preferred the weighted mean, we follow their lead for the rest of this paper.

Pruning the traces, as we had expected, does seem to decrease the predictability of bandwidth values, except for the U3 and U4 traces. This effect might be magnified for the recent-transfers technique, since (unlike the initial-RTT technique) it relies especially on intra-client predictability.

Table 15: Latency prediction via weighted mean bandwidth

Trace name	Samples included	Correlation w/latency	Median residual	Mean residual
C2	262165 (45.4%)	0.514	-0.042	-0.502
C2p	238957 (44.8%)	0.515	-0.046	-0.529
R2	24163 (60.4%)	0.525	-0.066	-4.100
R2p	17741 (56.5%)	0.560	-0.140	-5.213
U2	310496 (53.5%)	0.475	-0.028	-1.037
U2p	254024 (52.9%)	0.460	-0.033	-1.142
U3	341968 (53.7%)	0.330	-0.025	-1.138
U3p	260470 (53.0%)	0.374	-0.029	-1.288
U4	421867 (55.3%)	0.222	-0.021	-0.957
U4p	323811 (55.0%)	0.251	-0.024	-1.111

(a) 1 MSS $<$ length

Trace name	Samples included	Correlation w/latency	Median residual	Mean residual
C2	256943 (44.5%)	0.516	-0.038	-0.485
C2p	234160 (43.9%)	0.516	-0.043	-0.512
R2	17445 (43.6%)	0.317	-0.018	-0.779
R2p	12741 (40.6%)	0.272	-0.054	-0.959
U2	287709 (49.5%)	0.256	-0.020	-0.407
U2p	235481 (49.1%)	0.247	-0.024	-0.454
U3	314965 (49.4%)	0.447	-0.017	-0.300
U3p	239843 (48.8%)	0.484	-0.020	-0.336
U4	390981 (51.2%)	0.338	-0.015	-0.274
U4p	299905 (50.9%)	0.312	-0.017	-0.314

(a) 1 MSS $<$ length $<$ 32KB

Table 14 showed correlations between bandwidth measurements and predictions. To predict a response's latency, one can combine a bandwidth prediction with the known response length. Table 15 shows how well the weighted mean bandwidth technique predicts latencies. Table 15(a) includes responses with length at least one MSS; Table 15(b) excludes responses longer than 32 Kbytes. Because short responses and long responses may be limited by different parameters (RTT and bottleneck bandwidth, respectively), we hypothesized that it might not make sense to predict short-response latencies based on long-response history. Indeed, the residuals in Table 15(b) are always better than the corresponding values in Table 15(a), although the correlations are not always improved.

The correlations in Table 15(a) are better than those from the modified equation 1 as shown in Table 6, except for trace U4. However, the mean residuals in Table 15 are much larger in magnitude than in Table 6; it might be possible to correct the bandwidth-based predictor to fix this.

The previous-bandwidth approach consistently overpredicts latency, which in some applications might be better than underprediction. Figure 16 shows an example scatter plot, for R2. In the Web-server content adaptation application, excessive overprediction increases the chances that a well-connected user will fail to receive rich content, although this is less harmful than sending excessive content to a poorly-connected user.

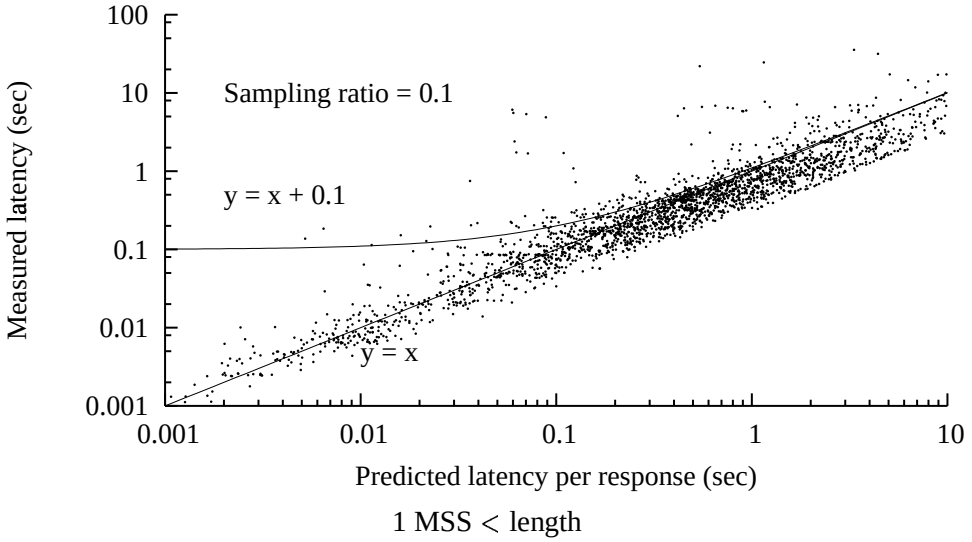

Figure 16: Real vs. bandwidth-predicted latency, trace R2

6.2 Predictions vs. a threshold

Just as we did for the initial-RTT approach, we evaluated the recent-transfers approach against a set of latency thresholds; Figure 17 shows the results for traces C2, R2, and U2. Compared to Figure 14, the recent-transfers approach is usually more accurate when the true latency is above the threshold (71% or better for C2, R2, and U2) but somewhat less accurate when the true latency is below the threshold.

However, our original goal was to predict latency for a client's first contact with a server. In this case, the recent-transfers approach must revert to a heuristic. We simulated a version using the exponentially smoothed mean ($\alpha = 0.7$) of all previous transfer bandwidths, and measured its accuracy for first-contact transfers only.

Figure 18 plots the results for first-contact transfers only, for all three traces, and for both the recent-transfers and initial-RTT approaches. This figure shows only the results for transfers whose true latency was above the threshold. For each of the traces, the recent-transfers is more accurate for thresholds below about 700 msecs, and the initial-RTT approach is more accurate for higher thresholds.

For first-contact transfers with true latency below the threshold, the initial-RTT predictor is correct at least 92% of the time, and usually more often, but the recent-transfers predictions are frequently wrong for thresholds below a few hundred msec.

6.3 Combining predictors

Given that the initial-RTT approach seems more accurate at predicting first-contact latencies, for many thresholds, than the recent-transfers approach, we speculated that a hybrid of the two predictors might yield the best results. This hybrid would use the modified equation 1 predictor for a client's first-contact transfer, and the smoothed mean of the client's previous bandwidths for its subsequent transfers.

We found that the overall (all-transfers) accuracy of this hybrid is nearly indistinguishable from the overall accuracy of the recent-transfers approach because, as the statistics in Table 1 imply, only a small fraction of transfers in our traces are first contacts.

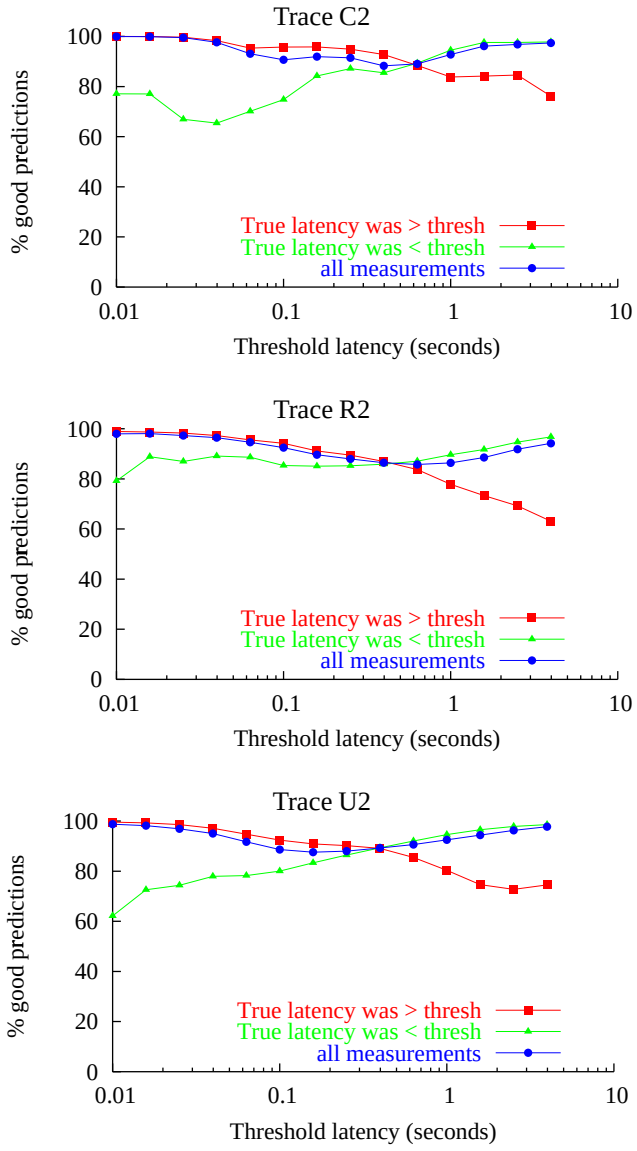

Figure 17: Predictions vs. threshold, using prev. bandwidths

7 Implementation issues

In this section, we cover a few issues related to implementing our prediction approaches in an actual Web server. Note that a version of the recent-transfers approach has previously been implemented [12].

7.1 Identifying proxies and robots

In Section 3.3.1, we explained how we pruned our traces to exclude likely proxies and robots. An actual implementation of our techniques would probably not apply them to requests from proxies, since the server-to-proxy bandwidth might be much higher than the proxy-to-client bandwidth. We would not want a Web server to select "fat" content to send via a well-connected proxy to a poorly-connected user. Robots, however, probably are almost all well-connected, and a server might not care if they receive a high-bandwidth variant.

How would a server decide not to send fat content to a proxy? In an ideal world, where all clients and proxies comply with the HTTP/1.1 specification [6]), the server could just look for `Via` headers to

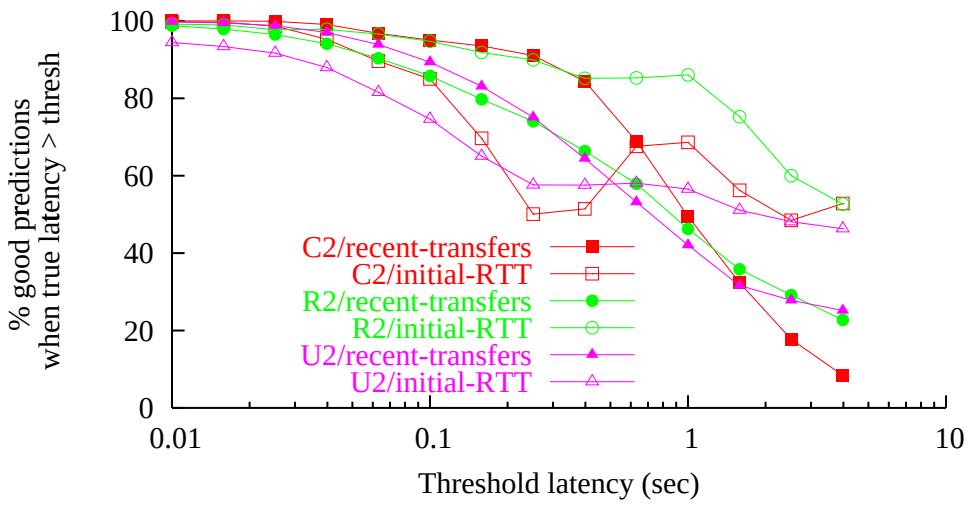

Figure 18: Predictions vs. threshold for first-contact transfers

identify proxies, although it might make an exception for a known set of personal firewalls (since if these proxies are well-connected, so are their end users). But the world is not ideal, and so a server might need to support other means of proxy detection. Krishnamurthy and Wang [10] describe several plausible techniques, including looking for multiple User-Agent fields from the same client IP address during a short time interval (as we did in Section 3.3.1), or inferring from the arrival rate of requests that they are multiplexed from several human sources.

In any case, we would not demand that a real server achieve perfect separation between proxies and other hosts. Not only is this likely to be impossible, but because our prediction techniques are not entirely accurate, it makes no sense to strive for perfection in deciding which clients to apply them to.

7.2 DHCP clients

One might suppose that a set of clients sharing a pool of IP addresses allocated by DHCP would confuse a server using the recent-transfer technique. That is, the server would apply, to a given IP address, predictions based on a previous user of that address. However, we believe this is not a problem in practice, because clients sharing a pool of DHCP-assigned addresses are likely to have similar connectivity. For example, they would all be using the same modem pool or wireless network.

7.3 API support

Both the initial-RTT technique and the recent-transfers technique require minor changes to a Web server application. The initial-RTT technique also requires minor kernel changes; it could also use the API proposed in [20], which is intended to provide portable access to TCP connection parameters, such as the RTT estimate.

8 Future work

We see numerous ways in which this work could be extended, including:

- Re-evaluating the recent-transfers technique using the address-clustering technique of Krishnamurthy and Wang [10] to keep history for clusters, not just individual addresses. This may provide useful predictions for previously unseen hosts that belong to clusters previously seen and identified.
- Evaluating the value of other “early” connection metrics as predictors. These include the apparent

bandwidth seen for the client's request message, or the interval between the server's SYN|ACK packet and the client's first request packet.

- Attempting to define a time window over which the prediction from past behavior remains valid.
- Attempting to use past history (perhaps with address clustering) to estimate the frequency of packet loss, and include that in a model-based predictor.
- Extending the analysis to traffic other than Web transfers, such as email, where a significant fraction are short TCP transfers.
- Evaluating the possibility for a Web server to self-adapt its choice of parameters (γ, w_1, *CompWeight*, and the assumed text-compression factor), rather than requiring these parameters to be established *a priori*. A server could include a self-adaptation module, which would continually update the parameters by training on all (or recent) previous requests.

One might also study certain human-factors issues, such as:

- How does user behavior change if the server is automatically adapting the content richness?
- Given a maximum threshold for overall page download latency, are existing sites respecting that threshold? If they are below that threshold, how much richer could their content be without violating the threshold?

9 Summary and conclusions

We conducted a study, based on traces from several different user communities, to demonstrate how well two different approaches can predict the latency of short TCP transfers. We found that by making a minor modification to a previously-described formula, we could greatly reduce its absolute prediction errors. We showed that predictions based on observation of past history generally yield better overall correlations than our formula-based predictor, but the formula-based predictor has lower mean prediction errors. We also show that the formula-based predictor could be improved to handle the specific case of text content, where modem-based compression can affect latency. Finally, we reported results from a study on the relationship between round-trip time and the use of modems, suggesting that this relationship might be exploited to improve prediction accuracy.

Acknowledgments

We would like to thank Vern Paxson, for his help with Bro and especially for writing the first draft of our Bro script and for improving Bro to meet our needs; Jerry Shan, for lots of help with statistical analysis; and Mike Rodriquez, Oliver Spatscheck, and others for their help in support of data collection. We thank Rick Jones, Zhuoqing Morley Mao, Dejan Milojicic, Mehul Shah, Chris Tuttle, Carey Williamson, and the anonymous reviewers for their helpful comments.

References

[1] Apache Software Foundation. Connections in the FIN_WAIT_2 state and Apache. http://httpd.apache.org/docs/1.3/misc/fin_wait_2.html.

[2] N. Cardwell, S. Savage, and T. Anderson. Modeling TCP latency. In *Proc INFOCOM (3)*, pages 1742–1751, Tel Aviv, Mar. 2000.

[3] Y.-C. Cheng, U. Hölzle, N. Cardwell, S. Savage, and G. M. Voelker. Monkey See, Monkey Do: A Tool for TCP Tracing and Replaying. In *Proc. USENIX Annual Tech. Conf.*, pages 87–98, Boston, MA, June 2004.

[4] S. Cheshire. Latency and the quest for interactivity, Nov. 1996. http://www.stuartcheshire.org/papers/LatencyQuest.ps.

[5] S. Cheshire. Latency survey results (for “It's the Latency, Stupid”).

http://www.stuartcheshire.org/rants/LatencyResults.html, 1996.

[6] R. Fielding, J. Gettys, J. Mogul, H. Frystyk, L. Masinter, P. Leach, and T. Berners-Lee. RFC 2616: Hypertext transfer protocol—HTTP/1.1, June 1999.

[7] L. Gomes, F. Castro, V. Almeida, J. Almeida, R. Almeida, and L. Bettencourt. Improving spam detection based on structural similarity. In *Proc. SRUTI*, pages 85–91, Cambridge, MA, July 2005.

[8] J. Hall, I. Pratt, I. Leslie, and A. Moore. The effect of early packet loss on Web page download times. In *Proc. Passive and Active Measurement Workshop*, La Jolla, CA, April 2003.

[9] Q. He, C. Dovrolis, and M. Ammar. On the Predictability of Large Transfer TCP Throughput. In *Proc. SIGCOMM*, Philadelphia, PA, Aug 2005.

[10] B. Krishnamurthy and J. Wang. On network-aware clustering of Web clients. In *SIGCOMM*, pages 97–110, Stockholm, Aug. 2000.

[11] B. Krishnamurthy and C. E. Wills. Improving Web performance by client characterization driven server adaptation. In *Proc. WWW2002*, pages 305–316, Honolulu, HI, May 2002.

[12] B. Krishnamurthy, C. E. Wills, Y. Zhang, and K. Vishwanath. Design, implementation, and evaluation of a client characterization driven web server. In *Proc. WWW2003*, pages 138–147, Budapest, May 2003.

[13] K. Lai and M. Baker. Nettimer: A tool for measuring bottleneck link bandwidth. In *Proc. USITS*, pages 123–134, San Francisco, CA, Mar 2001.

[14] K. Lakshminarayanan and V. N. Padmanabhan. Some findings on the network performance of broadband hosts. In *Proc. 3rd Internet Measurement Conf.*, pages 45–50, Oct. 2003.

[15] W. LeFebvre and K. Craig. Rapid Reverse DNS Lookups for Web Servers. In *Proc. USITS*, pages 233–242, Boulder, CO, Oct. 1999.

[16] Microsoft Corp. Knowledge Base Article 254337: Internet Explorer Sends Two GET Requests for the Same Data When the MIME Handler Is an ActiveX Control, 2000. http://support.microsoft.com/default.aspx?scid=kb;en-us;254337.

[17] Microsoft Corp. Exchange 2003 Design and Architecture at Microsoft. http://www.microsoft.com/technet/itsolutions/msit/deploy/ex03atwp.mspx, Aug 2003.

[18] Microsoft Corp. Knowledge Base Article 293792: Three GET Requests Are Sent When You Retrieve Plug-in Served Content, 2003. http://support.microsoft.com/default.aspx?scid=kb;en-us;293792.

[19] J. Mogul and L. Brakmo. Method for dynamically adjusting multimedia content of a web page by a server in accordance to network path characteristics between client and server. U.S. Patent 6,243,761, June 2001.

[20] J. Mogul, L. Brakmo, D. E. Lowell, D. Subhraveti, and J. Moore. Unveiling the transport API. In *Proc. 2nd Workshop on Hot Topics in Networks*, Cambridge, MA, November 2003.

[21] V. Paxson. Bro: A system for detecting network intruders in real-time. *Computer Networks*, 31(23-24):2435–2463, Dec. 1999.

[22] R Development Core Team. *R: A language and environment for statistical computing*. R Foundation for Statistical Computing, Vienna, Austria, 2003. ISBN 3-900051-00-3.

[23] B. Schroeder and M. Harchol-Balter. Web servers under overload: How scheduling can help. *ACM Trans. on Internet Techologies*, 6(1), Feb. 2006.

[24] S. Seshan, M. Stemm, and R. Katz. SPAND: Shared passive network performance discovery. In *Proc. USITS*, Monterey, CA, Dec. 1997.

[25] J. Su, A. Chin, A. Popivanova, A. Goel, and E. de Lara. User mobility for opportunistic ad-hoc networking. In *Proc. WMCSA*, pages 41–50, Lake District, UK, Dec 2004.

[26] W. Wei, B. Wang, C. Zhang, J. Kurose, and D. Towsley. Classification of Access Network Types: Ethernet, Wireless LAN, ADSL, Cable Modem or Dialup? In *Proc. IEEE Infocom*, pages 1060–1071, Miami, FL, March 2005.