Washington University in St. Louis

Washington University Open Scholarship

All Computer Science and Engineering Research

Computer Science and Engineering

Report Number: WUCSE-2005-4

2005-01-28

Smartacking: Improving TCP Performance from the Receiving End

Daniel K. Blandford, Sally A. Goldman, Sergey Gorinsky, Yan Zhou, and Daniel R. Dooly

We present smartacking, a technique that improves performance of Transmission Control Protocol (TCP) via adaptive generation of acknowledgments (ACKs) at the receiver. When the bottleneck link is underutilized, the receiver transmits an ACK for each delivered data segment and thereby allows the connection to acquire the available capacity promptly. When the bottleneck link is at its capacity, the smartacking receiver sends ACKs with a lower frequency reducing the control traffic overhead and slowing down the congestion window growth to utilize the network capacity more effectively. To promote quick deployment of the technique, our primary implementation of smartacking modifies only... **Read complete abstract on page 2.**

Follow this and additional works at: https://openscholarship.wustl.edu/cse_research

Part of the Computer Engineering Commons, and the Computer Sciences Commons

Recommended Citation

Blandford, Daniel K.; Goldman, Sally A.; Gorinsky, Sergey; Zhou, Yan; and Dooly, Daniel R., "Smartacking: Improving TCP Performance from the Receiving End" Report Number: WUCSE-2005-4 (2005). *All Computer Science and Engineering Research.*
https://openscholarship.wustl.edu/cse_research/956

Department of Computer Science & Engineering - Washington University in St. Louis
Campus Box 1045 - St. Louis, MO - 63130 - ph: (314) 935-6160.

This technical report is available at Washington University Open Scholarship: https://openscholarship.wustl.edu/cse_research/956

Smartacking: Improving TCP Performance from the Receiving End

Daniel K. Blandford, Sally A. Goldman, Sergey Gorinsky, Yan Zhou, and Daniel R. Dooly

Complete Abstract:

We present smartacking, a technique that improves performance of Transmission Control Protocol (TCP) via adaptive generation of acknowledgments (ACKs) at the receiver. When the bottleneck link is underutilized, the receiver transmits an ACK for each delivered data segment and thereby allows the connection to acquire the available capacity promptly. When the bottleneck link is at its capacity, the smartacking receiver sends ACKs with a lower frequency reducing the control traffic overhead and slowing down the congestion window growth to utilize the network capacity more effectively. To promote quick deployment of the technique, our primary implementation of smartacking modifies only the receiver. This implementation estimates the sender's congestion window using a novel algorithm of independent interest. We also consider different implementations of smartacking where the receiver relies on explicit assistance from the sender or network. Our experiments for a wide variety of settings show that TCP performance can substantially benefit from smartacking, especially in environments with low levels of connection multiplexing on bottleneck links. Whereas our extensive evaluation reveals no scenarios where the technique undermines the overall performance, we believe that smartacking represents a promising direction for enhancing TCP.

Washington University in St. Louis

School of Engineering & Applied Science

Smartacking: Improving TCP Performance from the Receiving End*

Daniel K. Blandford Sally A. Goldman Sergey Gorinsky Yan Zhou Daniel R. Dooly

Technical Report WUCSE-2005-4
Department of Computer Science and Engineering
Washington University, St. Louis, MO 63130, USA

January 28, 2005

Abstract

We present *smartacking*, a technique that improves performance of Transmission Control Protocol (TCP) via adaptive generation of acknowledgments (ACKs) at the receiver. When the bottleneck link is underutilized, the receiver transmits an ACK for each delivered data segment and thereby allows the connection to acquire the available capacity promptly. When the bottleneck link is at its capacity, the smartacking receiver sends ACKs with a lower frequency reducing the control traffic overhead and slowing down the congestion window growth to utilize the network capacity more effectively. To promote quick deployment of the technique, our primary implementation of smartacking modifies only the receiver. This implementation estimates the sender's congestion window using a novel algorithm of independent interest. We also consider different implementations of smartacking where the receiver relies on explicit assistance from the sender or network. Our experiments for a wide variety of settings show that TCP performance can substantially benefit from smartacking, especially in environments with low levels of connection multiplexing on bottleneck links. Whereas our extensive evaluation reveals no scenarios where the technique undermines the overall performance, we believe that smartacking represents a promising direction for enhancing TCP.

1 Introduction

Internet hosts support efficient and fair sharing of traversed network links by participating in congestion control protocols that regulate the amount of transmitted data in response to the observed network performance. In particular, Internet applications routinely rely on Transmission Control Protocol (TCP) [4] which establishes a connection between two communicating end hosts and enforces a dynamic *congestion window* as an upper limit on the amount of sent data that the receiver has not yet acknowledged. Initially, the congestion window is small and hence prevents the sender from injecting a lot of data into the network. The window increases when a timely acknowledgment (ACK) confirms data delivery. However, if the stream of ACKs indicates loss, the TCP connection reduces the congestion window and thereby curbs the transmission.

TCP is an end-to-end protocol that does not require any network help beyond best-effort unreliable delivery of sent data segments and their ACKs. Whereas the end-to-end property improves the protocol

*This work was performed at Washington University and supported in part by NSF Grant CCR-9734940. Daniel K. Blandford (dkb@andrew.cmu.edu) is currently with the Department of Computer Science, Carnegie Mellon University, Pittsburgh, PA 15213, USA. Sally A. Goldman (sg@cs.wustl.edu) and Sergey Gorinsky (gorinsky@wustl.edu) are at the Department of Computer Science and Engineering, Washington University, St. Louis, MO 63130, USA. Yan Zhou (zhou@cis.usouthal.edu) is with the School of Computer and Information Sciences, University of South Alabama, Mobile, AL 36688, USA. Daniel R. Dooly (ddooly@siue.edu) is with the Department of Computer Science at Southern Illinois University, Edwardsville, IL 62026, USA.

extensibility, achieving the efficient and fair network usage exclusively from the end points is challenging. In TCP, the sender carries most of this burden by performing a variety of computations, e.g., estimating RTT (round-trip time), maintaining a retransmission timer, counting duplicate ACKs, and adjusting the congestion window.

In contrast to the sender, the receiving end of the TCP connection has a simple role: after delivery of a data segment, the receiver generates an ACK. The receiver can delay transmitting the ACK for up to half a second but has to send at least one ACK for each two delivered data segments [7]. Lowering the frequency of ACKs reduces the protocol processing overhead and frees network resources for communicating data in the opposite direction. Delayed ACKs have been found to be particularly beneficial in asymmetric networks [5]. On the other hand, the delayed acknowledgment mechanism can disrupt the RTT estimation and slow down the growth of the congestion window.

In this paper, we present *smartacking*, a technique for adaptive generation of ACKs at the receiver. Smartacking not only reduces the amount of control traffic in the network but also makes data transmission over TCP more efficient. The algorithm is derived from an observation that the gap between delivered segments becomes close to uniform when traffic exhausts the capacity of the bottleneck link. Note that it is the receiver – not the sender – that can monitor the inter-segment arrival time (ISAT) to detect the link saturation. When the bottleneck link is at its capacity, the smartacking receiver sends ACKs with a lower frequency reducing the control traffic overhead and slowing down the congestion window growth. On the other hand, when the bottleneck link is underutilized, the receiver transmits an ACK for each delivered data segment and thereby preserves the rate at which TCP acquires the available capacity. Our experiments confirm that smartacking helps TCP to utilize the network capacity more effectively, including in topologies with asymmetric data flows and high bandwidth-delay products. We also show that smartacking TCP interacts fairly with standard TCP traffic.

Whereas proposals for improving TCP performance usually modify the sender, our implementation of the smartacking technique changes only the receiver. It would be also possible to implement smartacking with some support from the sender. For example, while our receiver-only solution incorporates a novel mechanism enabling the receiver to estimate the congestion window maintained by the sender, an alternative implementation could have made the sender communicate the congestion window to the receiver explicitly. However, we deliberately chose to avoid any modifications at the sender for the following two reasons. First, pursuing the receiver-only approach allowed us not only to demonstrate its feasibility but also to develop mechanisms of independent relevance, e.g., to ACC (Acknowledgment Congestion Control) [5] and other schemes where the receiver must know the congestion window. Second, altering a protocol exclusively at one end gives the new version a higher chance to enjoy widespread deployment. For instance, although TCP SACK [18] is technically superior to TCP NewReno [15], the latter has been deployed substantially more widely because of requiring changes only at one end whereas TCP SACK mandates modifications at both sending and receiving ends [20]. We hope that the receiver-only property of our solution will promote prompt adoption of smartacking by widely deployed TCP implementations.

The rest of the paper is organized as follows. Section 2 gives a brief overview of related work. Section 3 presents smartacking, our technique for adaptive generation of ACKs at the TCP receiver. Section 4 describes the methodology behind our evaluation of smartacking. Section 5 reports experimental results. Finally, Section 6 concludes the paper with a summary.

2 Related Work

TCP is not a rigid design and has evolved substantially, with congestion control becoming one of the most important additions. Since then numerous extensions – including Reno [16], SACK [18], Vegas [9], and NewReno [15] – have been proposed for TCP congestion control. However, the diverse extensions employ

the same tool to regulate transmission: a *congestion window* serves as an upper limit on the amount of transmitted data that the receiver has not yet acknowledged. Although the congestion window measures data in bytes, TCP passes data to the network in segments that do not exceed one MSS (maximum segment size). Among the current implementations of TCP, 536 and 1460 bytes are the most common MSS values. Initially, the congestion window allows only a small amount of unacknowledged data. The typical initial size of the window equals one MSS even though some TCP extensions initialize the congestion window to two MSSs [3, 21, 25]. The sender increases the window when timely acknowledgments confirm data delivery. However, if the stream of ACKs indicates loss, the sender decreases the congestion window. The goal of the window adjustment algorithm is to stabilize the transmission at a level that supports fair and efficient utilization of the network.

The window adjustment algorithm in the most widely deployed TCP versions operates in two modes – *slow start* and *congestion avoidance*. In the slow-start mode, the sender increases the congestion window by one MSS per non-duplicate ACK. When the bottleneck link is underutilized, the slow start doubles the window every RTT and thereby enables the connection to acquire the available capacity promptly. In the congestion-avoidance mode, a non-duplicate ACK increases the congestion window by an inverse of its current value; these adjustments grow the window by approximately one MSS per RTT and are supposed to supply convergence to the fair share of the bottleneck link capacity. Every TCP connection starts in the slow-start mode. Upon inference of congestion, the sender sets a threshold to the half of the current congestion window. When a triple duplicate ACK serves as the congestion indication, the sender reduces the window to the threshold and switches to the congestion-avoidance mode. If the congestion is inferred from a retransmission timeout, the sender reduces the window to one MSS and reenters the slow-start mode. The sender switches from the post-congestion slow start to the congestion-avoidance mode when the growing window reaches the threshold. Whereas the sender also performs a variety of other computations (e.g., RTT estimation), the receiver's participation in the window adjustment is limited to transmitting an ACK upon delivery of a data segment.

The idea to delay ACKs at the receiver is far from being new. Even before TCP was enhanced with congestion control, Clark proposed postponing an ACK as a means to reduce TCP processing overhead and release network resources [10]. Later, delayed ACKs were found to be particularly beneficial in asymmetric networks [5]. In modern TCP versions, the receiver does not delay an ACK beyond half a second and transmits at least one ACK for each two delivered data segments [7].

Lowering the frequency of ACKs can result in a slower growth of the congestion window because the sender increases the window in response to non-duplicate ACKs. In the slow-start mode, reducing the window growth interferes with the objective of acquiring the available capacity promptly and is therefore undesirable. ABC (Appropriate Byte Counting) [1] and SABC (Scaled Appropriate Byte Counting) [2] are instantiations of a *byte counting* technique where the sender increases the congestion window in response to acknowledged bytes rather than acknowledged segments; ACKs that confirm delivery of more data trigger larger increases in the congestion window. Byte counting does not change the receiving end of the connection.

ACC (Acknowledgment Congestion Control) [5] is an alternative design that modifies both the sender and receiver. ACC also relies on support from network routers: when the router detects congestion on its output link, the router sets the ECN (Explicit Congestion Notification) [23] bit in the headers of packets forwarded to the link. The receiver checks the headers of delivered packets for set ECN bits and uses these observations to adjust a delay factor d. Possible values of d vary from 1 to a maximum determined by the congestion window which the sender communicates to the receiver explicitly. The receiver doubles d upon receiving a packet with the set ECN bit. While no such packets arrive, the receiver decreases the delay factor by one per RTT. In ACC, the receiver transmits one ACK for every d delivered data packets.

Related to the improvement of congestion control from the receiving end is an issue of receiver misbehavior [12, 24]. The receiver of a TCP connection can misbehave by providing feedback incorrectly in order

to trick the sender into transmitting data at an unfairly high rate. In particular, it has been shown that by generating ACKs more frequently than prescribed by the protocol, the misbehaving receiver can substantially increase the reliable throughput of the connection at the expense of competing traffic [24]. In contrast to such misbehaving receivers, the receiver that follows smartacking as described in Section 3 generates ACKs *less* frequently with an objective of benefiting the overall efficiency and fairness of the network usage. Our experiments in Section 5 confirm that smartacking improves fairness of TCP.

3 Smartacking

In this section, we present *smartacking*, our technique for adaptive generation of ACKs at the TCP receiver. First, Section 3.1 discusses the main ideas behind smartacking. Then, Section 3.2 describes our implementation of the technique.

3.1 General Technique

In the slow-start mode, the sender of a TCP connection commonly injects a burst of data segments into the network. After the burst passes through the bottleneck link of the connection, gaps between the segments increase, and the burst spreads out. Eventually, the congestion window contains a large enough number of segments to utilize the bottleneck link fully. At this point, the segments arriving to the receiver are spread out the most over RTT of the connection. Also, since the bottleneck link reaches its capacity, continuing the aggressive growth of the congestion window does not improve the link utilization.

Note that it is the receiver – not the sender – that can monitor gaps between the arriving segments to determine the link saturation. This observation leads us to propose *smartacking*, a technique for improving TCP performance from the receiving end. In smartacking TCP, the receiver generates ACKs depending on its measurements of gaps between delivered data segments. When the bottleneck link is underutilized, the receiver transmits one ACK per segment and thereby preserves the rate at which TCP acquires the available capacity. When the bottleneck link is at its capacity, the receiver sends ACKs less frequently reducing the control traffic overhead and slowing down the congestion window growth.

For how long should the receiver delay ACKs when transmitting them with a lower frequency? Delaying an ACK creates a potential danger of underutilizing the bottleneck link. If the ACK is delayed for too long, the sender stops transmitting after the amount of unacknowledged data covers the congestion window, and subsequently the bottleneck link runs out of packets to forward. Hence, smartacking strives to generate ACKs with the lowest frequency that keeps the bottleneck link fully utilized. To achieve this goal, the receiver estimates time `LastSegmentTime` when it will receive the last segment allowed by the congestion window. Then, the receiver transmits the ACK with delay $\texttt{LastSegmentTime} - \texttt{RTT}$ so that the sender will receive the ACK just before exhausting the congestion window.

To estimate `LastSegmentTime`, the receiver maintains the following three variables that measure data in maximum-size segments: (1) `CwndEst` is an estimate of the current congestion window, (2) `LastReceived` records the last received data, and (3) `LastACKed` denotes the last acknowledged data. Then, $\texttt{CwndEst} + \texttt{LastACKed} - \texttt{LastReceived}$ represents the number of additional segments that the receiver expects from the current congestion window. The receiver also computes `ISAT`, an average of inter-segment arrival times. Using `ISAT`, the receiver estimates the arrival time of the last segment from the current congestion window as:

$$\texttt{LastSegmentTime} \leftarrow (\texttt{CwndEst} + \texttt{LastACKed} - \texttt{LastReceived}) \cdot \texttt{ISAT}. \tag{1}$$

When $\texttt{LastSegmentTime} < \texttt{RTT}$, the network capacity is underutilized, and the receiver transmits an ACK immediately. If $\texttt{LastSegmentTime} > \texttt{RTT}$, the receiver delays the ACK for up to $\texttt{LastSegmentTime} - \texttt{RTT}$ but no longer than the maximum delay allowed by TCP.

3.2 Receiver-Only Implementation

In addition to the inter-segment arrival time ISAT observable at the receiver, implementing the technique requires knowledge of the congestion window and round-trip time, i.e., information that is readily available at the sender. Hence, it might seem reasonable to implement smartacking in a distributed manner where the receiver sets its CwndEst and RTT variables to values communicated explicitly by the sender. We, however, deliberately pursue a receiver-only implementation of smartacking. This choice is chiefly due to deployment considerations. Experience shows that TCP extensions that upgrade only one of the communicating ends are more likely to enjoy wide adoption than those extensions that require changes at both ends. For example, both SACK [18] and NewReno [15] have been proposed for addressing multiple losses within the congestion window; although SACK is an earlier and technically superior solution than NewReno, the latter is the most widely deployed design today because NewReno upgrades only the sender whereas SACK modifies both the sender and receiver [20]. We hope that our receiver-only implementation will help smartacking to find wide deployment.

To implement smartacking with no support from the sender, the receiver can obtain most of the needed information either straightforwardly or using well-known estimation mechanisms. For example, the receiver has direct knowledge of LastACKed and LastReceived. In our implementation, the smartacking receiver computes ISAT as an exponentially weighted moving average (EWMA) with a gain of 0.25. The averaging ignores measurements triggered by the first segment arrival from any window. Also, the receiver ignores up to two consecutively measured values that are at least three times larger than the current ISAT. To compute RTT, the receiver reuses the standard TCP mechanism for RTT estimation; when the receiver is not transmitting its own data to the sender, the receiver estimates RTT by sending keep-alive messages once per second.

Estimating the congestion window is the only truly novel challenge for our receiver-only implementation of smartacking. Furthermore, since different extensions of TCP adjust the congestion window at the sender differently, the congestion window estimation at the receiver might need to be extension-specific as well. Figure 1 describes our mechanism for estimating the congestion window in NewReno. The presented procedure EstimateCwnd has three phases corresponding to the initial slow-start, post-congestion slow-start, and congestion-avoidance modes of the sender. During the initial slow-start phase, the receiver tracks the congestion window using the variable StartingCwnd. This variable is initially set to one segment and is incremented every time when the receiver transmits a non-duplicate ACK. After sending the third duplicate ACK, the receiver leaves the initial slow-start phase and marks the transition by setting StartingCwnd to -1. Subsequently, the boolean variable PastSsthresh determines the current phase of the receiver: true PastSsthresh corresponds to the congestion-avoidance phase, and false PastSsthresh denotes the post-congestion slow-start phase. In the congestion-avoidance phase, the receiver tracks the congestion window using variable CwndEst and increases this variable by $1/$CwndEst when transmitting a non-duplicate ACK. To initialize CwndEst when switching to the congestion-avoidance phase as well as to estimate the congestion window in the post-congestion slow-start phase, the receiver maintains a timer expiring once per RTT. When the timer expires at the end of a one-RTT interval, the receiver records the observed number of delivered in-order segments and transmitted non-duplicate ACKs in variables NumSegmentsThisInterval and NumACKsThisInterval respectively. Also, NumSegmentsLastInterval and NumACKsLastInterval save respectively the previous values of NumSegmentsThisInterval and NumACKsThisInterval recorded after the timeout one RTT earlier. Then, the receiver computes variable CwndIfSlowStart as a sum of NumSegmentsLastInterval, NumACKsLastInterval, and NumACKsThisInterval to estimate the congestion window in the post-congestion slow-start phase.

Our novel mechanism for the congestion window estimation plays an important role in our implementation of smartacking. However, this mechanism is also independently valuable because of its potential usefulness to other designs – such as ACC [5] – where the receiver must know the congestion window.

```
int EstimateCwnd(){
   if (StartingCwnd > -1)
      return StartingCwnd;
   if(PastSsthresh)
      return (int) CwndEst;
   Total = NumSegmentsLastInterval + NumACKsLastInterval + NumACKsThisInterval;
   if (Total > SsthreshEst)
      return max(SsthreshEst,NumSegmentsLastInterval) + 2;
   else return Total;
}
```

When sending a non-duplicate ACK, execute the following:

```
   NumACKsThisInterval++;
   if (StartingCwnd > -1)      // initial slow-start phase
      StartingCwnd++;
   if (CwndEst > 0)
      CwndEst += 1.0/CwndEst;
```

Whenever an in-order segment arrives, execute the following:

```
   NumSegmentsThisInterval++;
```

Once per RTT, execute the following:

```
   if (!PastSsthresh
        && CwndIfSlowStart > NumSegmentsThisInterval + 2
        && NumACKsLastInterval > 2
        && !SeenLossRecently){   // switching to the congestion-avoidance phase
     PastSsthresh = True;
     CwndEst= NumSegmentsThisInterval + 1;
   }
   if (NumSegmentsThisInterval  == 0 && NumSegmentsLastInterval == 0)
      SeenLossThisInterval = True;
   if (SeenLossThisInterval && !SeenLossRecently){
      SsthreshEst = max((int)(EstimateCwnd()/2), 2);
      StartingCwnd = -1;
      PastSsthresh = False;
   }
   if (NumSegmentsThisInterval > 1)
      SeenLossRecently = SeenLossThisInterval;
   SeenLossThisInterval = False;
   CwndIfSlowStart = NumSegmentsLastInterval + NumACKsLastInterval
                   + NumACKsThisInterval;
   NumSegmentsLastInterval = NumSegmentsThisInterval;
   NumACKsLastInterval = NumACKsThisInterval;
   NumSegmentsThisInterval = NumACKsThisInterval = 0;
```

Figure 1: Estimating the congestion window for TCP NewReno.

4 Evaluation Methodology

We evaluate the impact of smartacking on TCP performance in a variety of network configurations. We also compare smartacking with alternative approaches. After Section 4.1 describes the evaluated schemes, Section 4.2 presents our experimental setup. Then, Section 4.3 discusses metrics for assessing the performance.

4.1 Evaluated Designs

In Section 3.2, we presented our implementation of the smartacking technique for TCP NewReno. We refer to this implementation as *NewReno with smartacking* and compare it with schemes that can be classified into two categories: existing designs and different implementations of smartacking.

Existing Designs. Since *NewReno* [15] is the most widely deployed implementation of TCP today and serves as a basis for our smartacking implementation, we include NewReno in our studies. This design employs delayed ACKs as described in RFC 1122 [7]. In addition, we report results for *AckEvery*, a NewReno implementation that acknowledges every segment immediately. To compare smartacking with other proposals for overcoming the negative impact of delayed ACKs on the congestion window growth, we also evaluate *ACC* and *ABC*. Although we also conducted experiments with SABC, we do not report these results here because of their similarity to the results reported for ABC.

Different Implementations of Smartacking. Due to deployment considerations, we deliberately pursued a receiver-only solution while implementing the smartacking technique in Section 3.2. On the other hand, support from the sender or network can lead to more effective implementations of smartacking. In our experiments, we quantify potential benefits from such support mechanisms.

If the bottleneck link of a TCP connection is fully utilized, the smartacking receiver of the connection expects from arriving segments to be uniformly distributed over RTT. This ideal scenario happens when the connection is the only source of packets on the bottleneck link. However, when the saturated link carries other traffic as well, the receiver can observe a burstier distribution of the arriving segments: instead of being spread out over RTT, the segments from the congestion window can clump within a smaller portion of RTT. Consequently, the receiver can transmit ACKs sooner than theoretical principles of smartacking prescribe. To alleviate the distortion, the sender and routers can employ mechanisms that improve packet mixing on shared bottleneck links.

Pacing is a sender-based mechanism that smooths the bursty pattern of TCP transmission by distributing the sent segments over RTT [26]. Pacing helps not only with improved packet mixing on shared links but also in networks with extremely low reverse-path capacities. Since a burst of segment arrivals can cause a burst of ACKs, the bursty transmission of data segments can congest the low-capacity reverse path with ACKs. Ensued queueing of ACKs increases the RTT estimate and shortens the delays of ACKs at the smartacking receiver. Pacing addresses this problem by reducing the possibility of the reverse-path congestion. However, spreading of data segments over the whole RTT would interfere with operation of smartacking: since the smartacking receiver detects the availability of the bottleneck capacity via gaps between successive congestion windows, the spread of segments over the whole RTT would always trick the receiver into perceiving the bottleneck link as saturated even when the link is actually underutilized. Hence, we implement pacing by distributing the transmitted segments from a congestion window evenly over an interval equal to 0.9 RTT. We refer to the enhanced implementation as *NewReno with smartacking and pacing*.

Fair queueing of packets at the bottleneck link is a router-based mechanism that also can improve packet mixing [6, 11]. To assess the impact of fair queueing, we consider a simple round-robin algorithm that allocates a separate queue per flow and visits the queues in a round-robin manner to select a packet for transmission to the link. We use *NewReno with smartacking and fair queueing* to denote results achieved by our smartacking implementation in networks with round-robin scheduling.

An alternative to modifying the sender or routers is to implement smartacking differently at the receiver. Although the receiver cannot affect packet mixing directly, computations of ACK delays at the receiver can account for distortions in the distribution of arriving segments. Hence, we extend smartacking by introducing a parameter α to calculate the ACK delay as $\texttt{LastSegmentTime} - \alpha \cdot \texttt{RTT}$ rather than $\texttt{LastSegmentTime} - \texttt{RTT}$. When the arriving segments spread over the whole RTT, the receiver should use $\alpha = 1$ and compute the ACK delay in the same way as our smartacking implementation in Section 3.2. However, when the bottleneck link is fully utilized but the segments still clump within a smaller portion of RTT, the receiver should keep the value of α between 0 and 1 to balance the underestimation of `LastSegmentTime`. Setting α to 0 corresponds to the scenario when the receiver transmits the ACK upon arrival of the last segment from the congestion window. We implement the following algorithm for adjustment of α. Initially, the receiver sets α to 1 and does not reduce its value until generating a third duplicate ACK. Subsequently, the receiver multiplies α by 0.9 whenever the congestion window estimate decreases. The receiver also maintains two auxiliary parameters β and γ initialized to 1 and 0 respectively. Parameter β takes real values between 0 and 1 and acts with respect to α in the same way as the threshold acts toward the congestion window at the sender. Parameter γ is an integer reflecting the frequency of ACKs confirming a single segment. Whenever the receiver transmits an ACK, γ is adjusted: if the ACK confirms a single segment, the receiver increments γ; otherwise, the receiver decrements γ. When γ becomes negative with its absolute value exceeding $4 \cdot \texttt{CwndEst}$ (i.e., when most ACKs confirm multiple segments), the receiver resets γ to 0 and updates α as follows: if $\beta > \alpha$, then the receiver increases α by 0.05; otherwise, the receiver decreases α by 0.01. We refer to this enhanced implementation as *NewReno with extended smartacking*.

4.2 Experimental Setup

In this section, we discuss network topologies, traffic patterns, and protocol settings in our experiments. Since we expect that smartacking provides most of its benefits in environments with small degrees of flow multiplexing on bottleneck links, we use few flows in most of the experiments. However, we also conduct some experiments with a large number of flows to study the impact of smartacking in networks with high levels of multiplexing. Also, although we follow suggestions from Floyd and Jacobson [13] and focus on settings where data flows traverse bottleneck links in both directions, we also consider some simpler settings with no data flows on the reverse path. We examine each TCP extension discussed in Section 4.1 in a bulk-transfer scenario where the sender communicates a large file containing ten thousand data segments. Packets that carry the data segments are 1000-byte long. The size of packets carrying ACKs is 40 bytes. The maximum value for an ACK delay is set to 200 ms (since our results for the maximum ACK delay of 100 ms are similar, we do not report them here). The start times of connections stagger by 100 ms. We set the advertised window of each receiver so that it never limits the congestion window at the sender. In some experiments, we add on-off UDP (User Datagram Protocol) [22] cross traffic to study the reaction to sharp changes in the available network capacity.

In most of our experiments, we use routers with *Droptail* buffer management because of their prevalence in the modern Internet. The size of a Droptail buffer is usually set to the half of the bandwidth-delay product. However, we also report results for our experiments with different buffer sizes. Finally, we briefly report on a small portion of results from our extensive experiments with networks of *RED* (Random Early Detection) routers [8, 14] where RED parameters are configured as $\min_{\text{th}} = 5$, $\max_{\text{th}} = 15$, $w_q = 0.002$, and $\max_p = 0.1$.

We conduct the experiments in two general network topologies. Figure 2a depicts a *Dumbbell topology* common in congestion control studies: each flow traverses three hops, and the link that connects the only two routers in the topology is the middle hop for each flow. For the Dumbbell topology, we examine a number of configurations of increasing complexity:

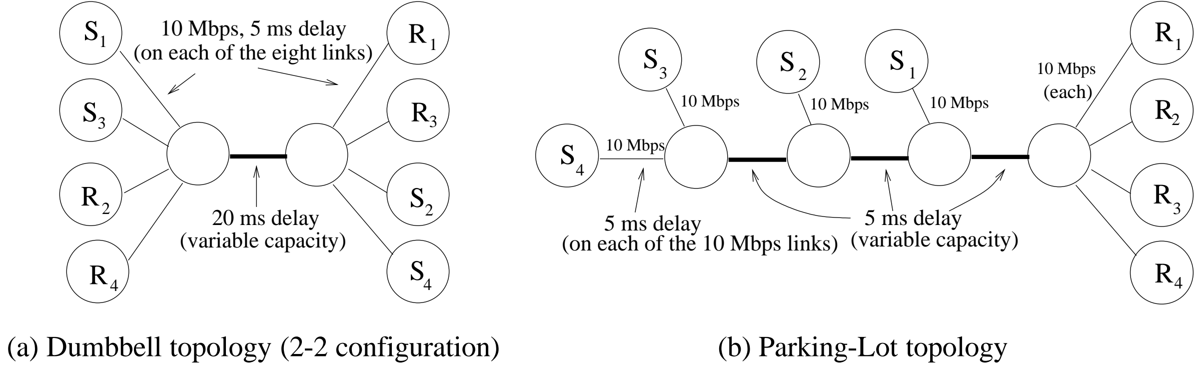

(a) Dumbbell topology (2-2 configuration)

(b) Parking-Lot topology

Figure 2: Network topologies in our experiments.

- A *1-0 configuration* is the simplest. It has symmetric links, one TCP connection with a one-way propagation delay of 50 ms, and no data traffic on the reverse path.
- To create an *Asymmetric configuration*, we adopt the setup from Balakrishnan, Padmanabhan, and Katz [5]. Each of the edge links has a capacity of 10 Mbps and propagation delay of 1 ms in both directions. However, the middle link of the Dumbbell topology is asymmetric. It has a capacity of 10 Mbps and propagation delay of 5 ms along the data path of the only TCP connection in the configuration. In the reverse direction, the link has a propagation delay of 50 ms and capacity that varies between 5 Kbps and 30 Kbps.
- A *2-2 configuration* is symmetric. Four TCP connections transmit data over the bottleneck link in opposite directions so that each router forwards to the bottleneck link a mix of ACKs and data packets. The first connection starts sending its data at time 0. Another connection starts its transmission in the same direction 200 ms later. The two connections that send data in the reverse direction start at time 100 ms and 300 ms respectively. The middle bottleneck link has a propagation delay of 20 ms and variable capacity. Each of the edge links has a capacity of 10 Mbps and propagation delay of 5 ms. Hence, the round-trip propagation time for the connections is 60 ms.
- A *High-Multiplexing configuration* replaces every connection in the 2-2 configuration with 50 parallel TCP connections. Hence, 100 connections communicate data in one direction of the bottleneck link, and the other 100 connections deliver data in the opposite direction.
- A *Different-RTT configuration* is a variation of the 2-2 configuration. In each direction, one of the two connections has an increased propagation RTT of 120 ms. The propagation RTT for the other connection remains 60 ms.
- A *TCP-UDP configuration* replaces one of the two TCP connections in each direction of the 2-2 configuration with an on-off UDP flow that changes its mode of operation once per 5 seconds. When the UDP flow is on, it injects packets into the network at a constant rate equal to the half of the bottleneck link capacity.

Figure 2b presents a *Parking-Lot topology* which is commonly used to assess fairness of congestion control protocols. Our configuration of the topology has four TCP connections and three bottleneck links. S_i and R_i refer respectively to the sender and receiver of connection i. S_1, S_2, S_3, and S_4 start transmitting their files at time 0, 100 ms, 200 ms, and 300 ms respectively.

We perform all the experiments in our own simulator. To validate the simulator, we have repeated some of the experiments in NS-2 [19] and received results that are indistinguishable from those reported below.

4.3 Performance Metrics

Completion time of a TCP connection, defined as the amount of time between the transmission of the first data segment and delivery of the last ACK, is the main measure of performance in our studies. Completion time also serves as a basis for other long-term performance metrics. For example, *relative completion time* is computed as the ratio of the completion times for a pair of compared connections. To measure fairness of network sharing, we use a *fairness index* [17]:

$$F(t_1, t_2, \ldots, t_n) = \frac{\left(\sum_{i=1}^{n} t_i\right)^2}{n \cdot \sum_{i=1}^{n} t_i^2} \tag{2}$$

where n is the number of connections, and t_i denotes the completion time of connection i. In addition to the long-term metrics, we monitor dynamics of TCP connections on shorter timescales and record the congestion window and threshold at the sender, congestion window estimate at the receiver, transmission and arrival times of data segments and ACKs.

5 Experimental Results

This section presents results from our experimental evaluation of smartacking. First, we explore the impact of smartacking on dynamics of TCP communications in Section 5.1. Then, Section 5.2 shifts our attention to long-term performance in networks with Droptail routers. Section 5.3 reports results for networks of RED routers. Section 5.4 studies performance of smartacking in the Asymmetric configuration. Section 5.5 quantifies the reaction of the examined protocols to sharp changes in the available capacity. Section 5.6 evaluates smartacking in network configurations with high levels of connection multiplexing on bottleneck links. Finally, Section 5.7 investigates the impact of smartacking on fairness of network sharing.

5.1 Impact of Smartacking on Dynamics of TCP Communications

Before evaluating the widely deployed TCP designs and their extensions in terms of long-term performance, it is important to understand how short-term dynamics of the schemes affect the cumulative metrics. In this section, we use *timeline diagrams* to reflect the behavior of a TCP connection over a range of shorter timescales, starting from the under-RTT timescale to intervals that capture a sequence of transitions between the congestion control modes of the connection. In each of our timeline diagrams, a horizontal band represents interactions between the sender and receiver over a fixed period. Whereas the top edge of the band corresponds to the sender, the bottom edge denotes the receiver. Time advances from the left to the right. Every timeline diagram consists of 14 horizontal bands stacked in their temporal order: the left edge of the top band corresponds to time 0, the right edge of this band represents the same time as the left edge of the band immediately below, and so on until the right edge of the bottom band shows the last time reflected in the diagram. While gray lines that descend to the right within a band represent communications of data segments, black lines that rise to the right denote ACKs. Consequently, darkness of a band region indicates intensity of communications during the corresponding time interval. We annotate each timeline diagram with the name of the examined protocol and also report what percentage of the file is communicated by the end of the timeline.

To show the impact of smartacking, Figure 3 presents two timeline diagrams for the 1-0 configuration with the bottleneck link capacity of 9 Mbps. The diagram in Figure 3a demonstrates that NewReno stalls after delivering the first data segment because the receiver delays acknowledging the segment for the maximum duration allowed. When the transmission resumes, the sender doubles the congestion window each

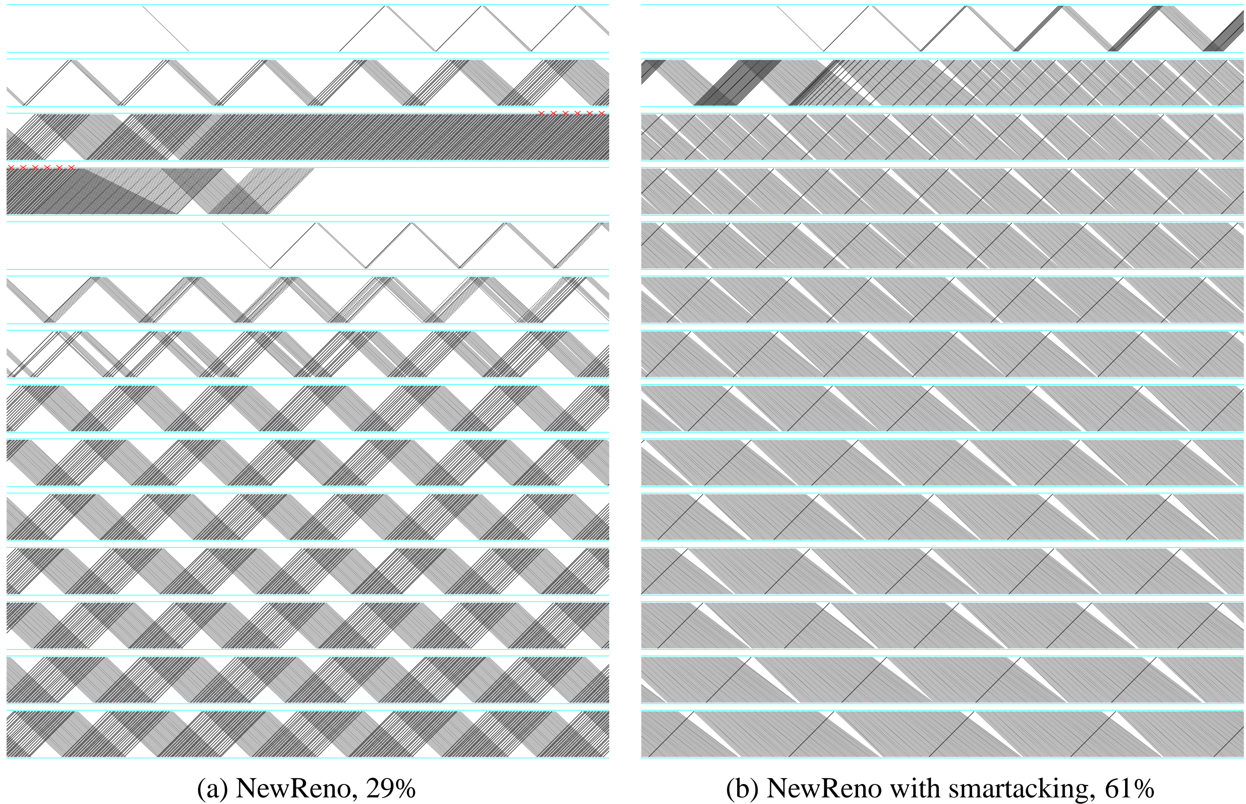

(a) NewReno, 29%

(b) NewReno with smartacking, 61%

Figure 3: Timeline diagrams for the 1-0 configuration with the bottleneck link capacity of 9 Mbps.

RTT until a burst of data segments saturates the bottleneck link and causes losses. After a retransmission timeout ends the ensued second stall, NewReno reduces the window to one MSS and then goes through the post-congestion slow start into congestion avoidance. However, the connection fails to recapture the bottleneck capacity and communicates only 29% of the file by the end of the timeline. As Figure 3b shows, NewReno with smartacking achieves a higher throughput of 61%. Smartacking improves the throughput partly by avoiding both stalls. Upon receiving the first data segment or whenever the bottleneck link is underutilized, the smartacking receiver transmits an ACK immediately and does not slow down the growth of the congestion window. As the transmission approaches the bottleneck link capacity, NewReno with smartacking reduces the frequency of ACKs and converges to a pattern where the receiver provides the sender with one ACK per window so that – as the narrow white triangles in Figure 3b indicate – data from a new window starts arriving to the bottleneck link router before the link runs out of packets to forward. In this phase, the connection grows the congestion window slowly and induces no losses.

We repeat the above experiments in the 1-0 configuration with the bottleneck link capacity of 5 Mbps. For this setting where loss rates are higher, Figures 4a and 4b also show that smartacking reduces the reverse-path traffic and improves the connection throughput. Albeit, the improvement is less dramatic: from 23% to 27%. Figure 4c confirms a known impact of pacing on NewReno – the smoother transmission slightly reduces the throughput. Figure 4d demonstrates the same effect of pacing on NewReno with smartacking. However, NewReno with smartacking and pacing still outperforms plain NewReno.

Figure 5 depicts the impact of smartacking on NewReno in the 2-2 configuration with the bottleneck link capacity of 20 Mbps. Similarly to the above scenarios with no data packets on the reverse path, plain NewReno stalls twice: (1) due to the delayed ACK of the first data segment, and (2) recovering from losses via a retransmission timeout. On the other hand, NewReno with smartacking acquires the available capacity

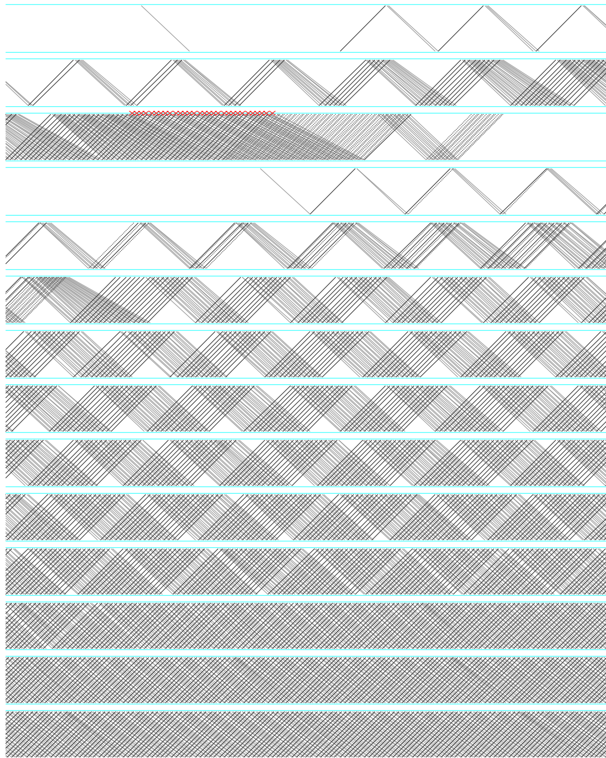

(a) NewReno, 23%

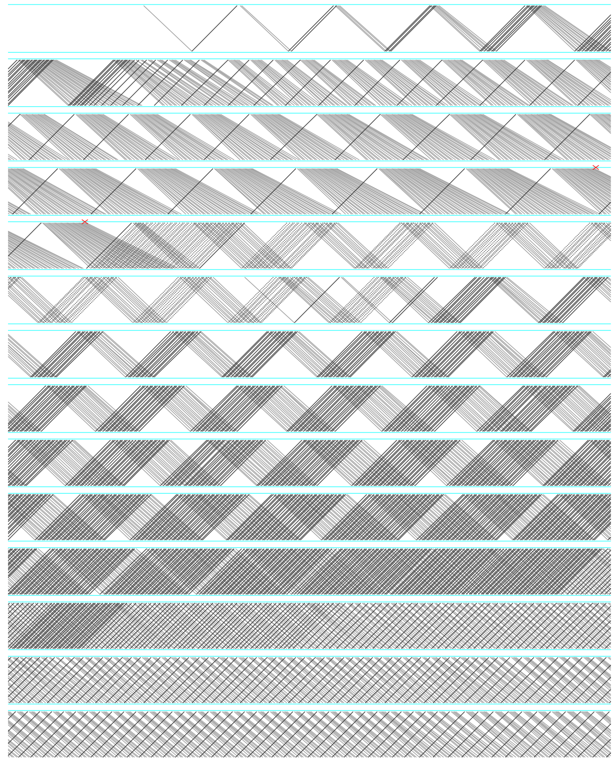

(b) NewReno with smartacking, 27%

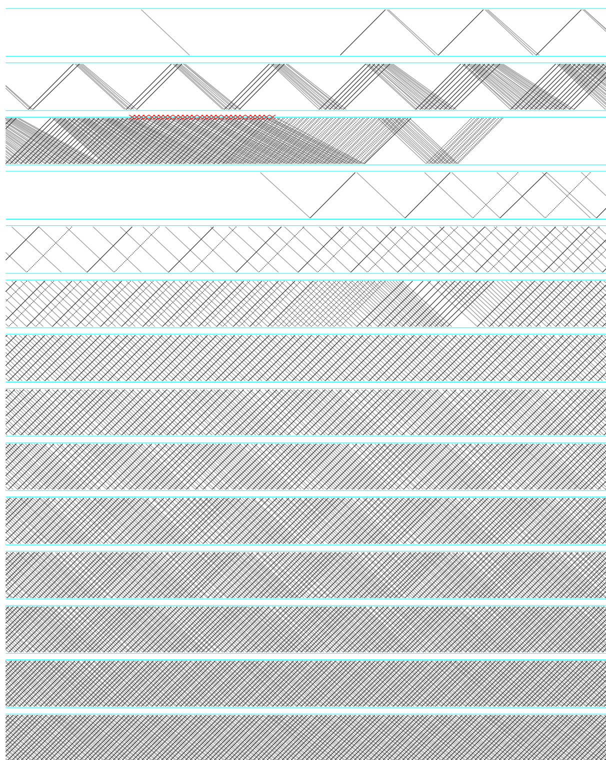

(c) NewReno with pacing, 22%

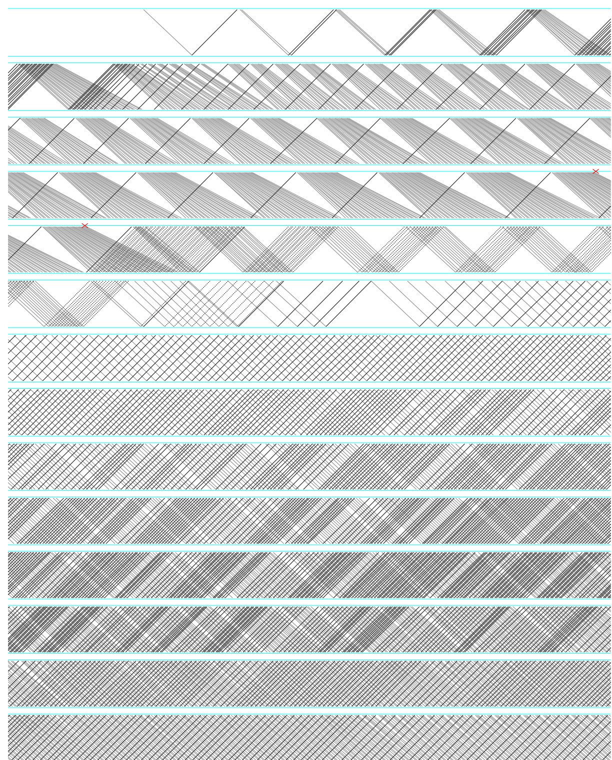

(d) NewReno with smartacking and pacing, 25%

Figure 4: Timeline diagrams for the 1-0 configuration with the bottleneck link capacity of 5 Mbps.

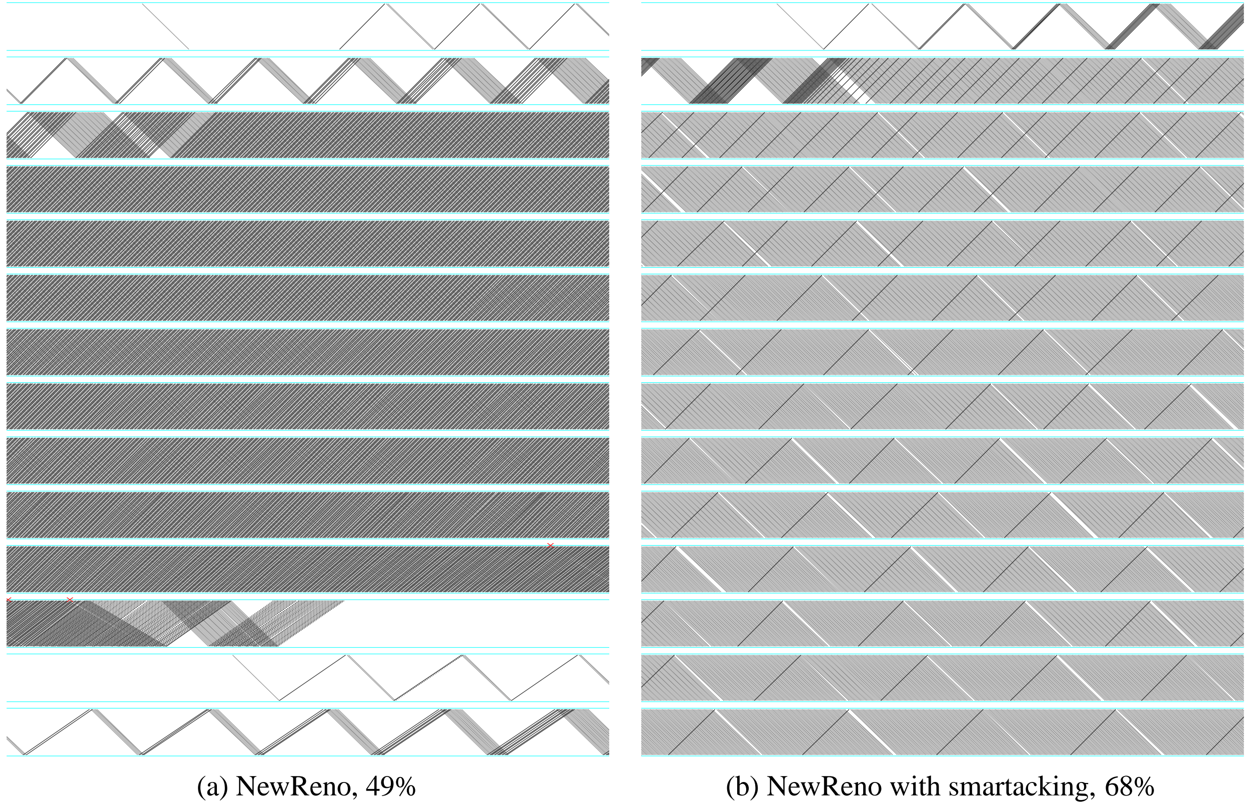

(a) NewReno, 49%

(b) NewReno with smartacking, 68%

Figure 5: Timeline diagrams for the 2-2 configuration with the bottleneck link capacity of 20 Mbps.

promptly and then maintains smooth data delivery with a low frequency of ACKs and no retransmission timeouts. The 49%-to-68% throughput increase quantifies the improvement in performance.

To provide insights into reasons for throughput improvements offered by smartacking, we trace the congestion window at the sender. Figure 6 reports the traced values for eight of the examined TCP extensions in the 1-0 configuration with the bottleneck link capacity of 5 Mbps. In the schemes with no smartacking (NewReno, NewReno with pacing, AckEvery, ABC, and ACC), the window in the congestion-avoidance mode grows at approximately the same rate regardless of the bottleneck link utilization. On the other hand, the designs with smartacking (NewReno with smartacking, NewReno with smartacking and pacing, and NewReno with extended smartacking) reduce the window growth when the bottleneck link gets saturated. The smoother increase prolongs the periods when the bottleneck link capacity is fully utilized. Since the area under the congestion window curve is closely correlated to the connection throughput, Figure 6 illustrates why the adjustment of the ACK frequency allows smartacking to improve the throughput. The graphs also trace our estimate of the congestion window at the receiver and demonstrate accuracy of the proposed estimation mechanism.

Figure 7 traces the congestion window in the 2-2 configuration for two values of the bottleneck link capacity (8 Mbps and 4 Mbps) and four TCP versions (NewReno, NewReno with pacing, NewReno with smartacking, and NewReno with smartacking and pacing). Despite the higher degree of connection multiplexing, estimation of the congestion window at the receiver remains precise. The graphs also show that pacing helps smartacking to smooth oscillations of the congestion window when the bottleneck link is fully utilized.

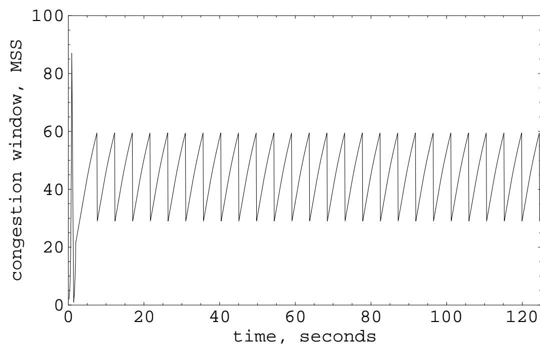

(a) NewReno

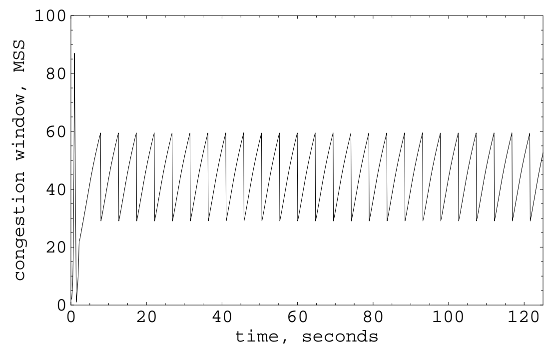

(b) NewReno with pacing

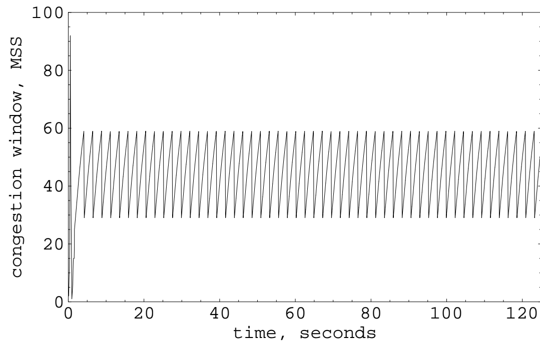

(c) AckEvery

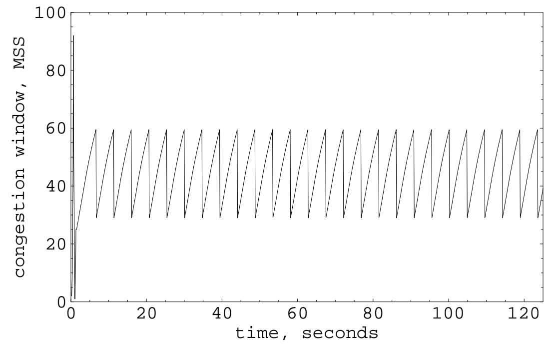

(d) ABC

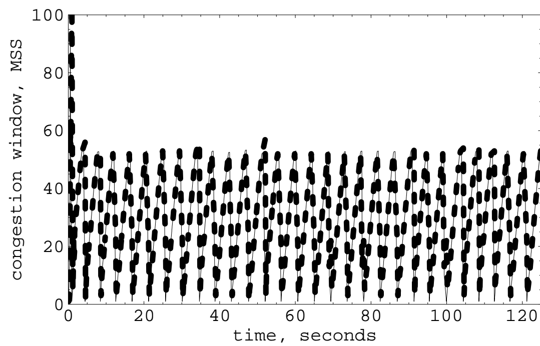

(e) ACC

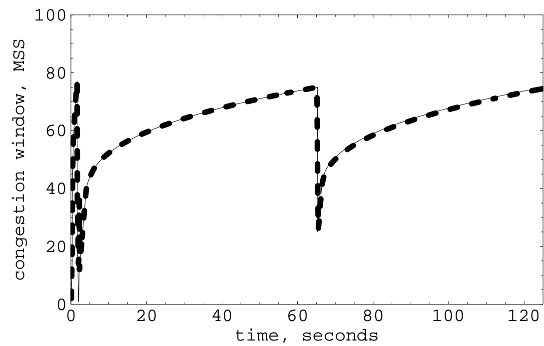

(f) NewReno with smartacking

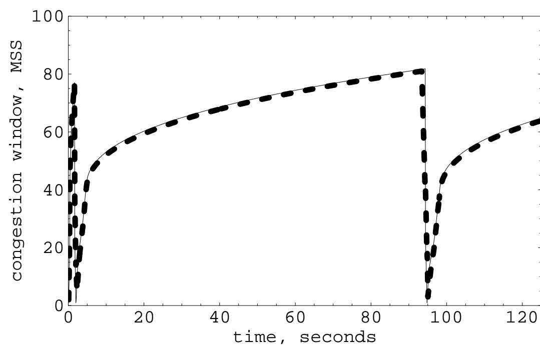

(g) NewReno with smartacking and pacing

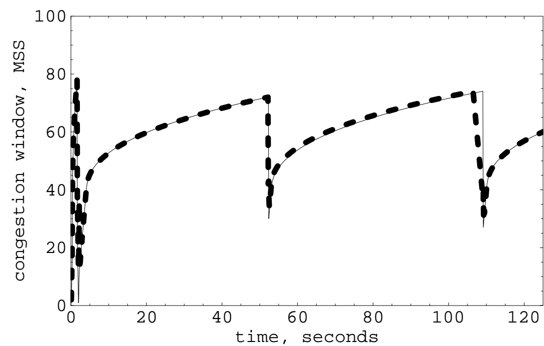

(h) NewReno with extended smartacking

Figure 6: Congestion windows at the sender (solid lines) and their estimates at the receiver (dotted lines) in the 1-0 configuration with the bottleneck link capacity of 5 Mbps.

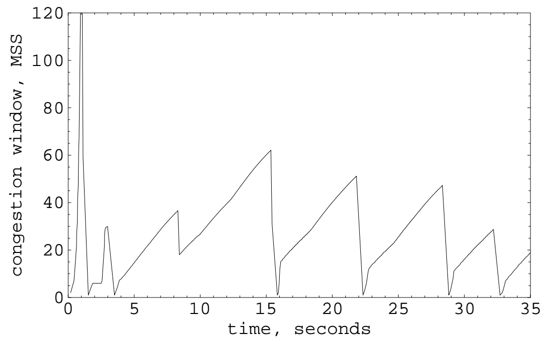

(a) NewReno, 8 Mbps

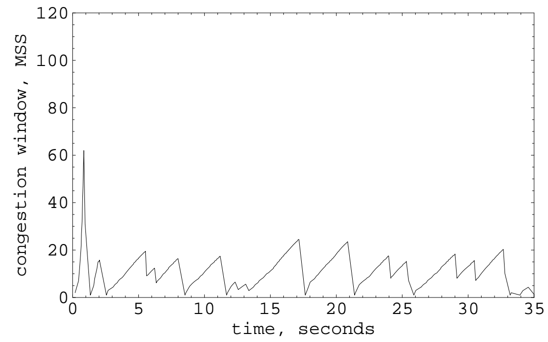

(b) NewReno, 4 Mbps

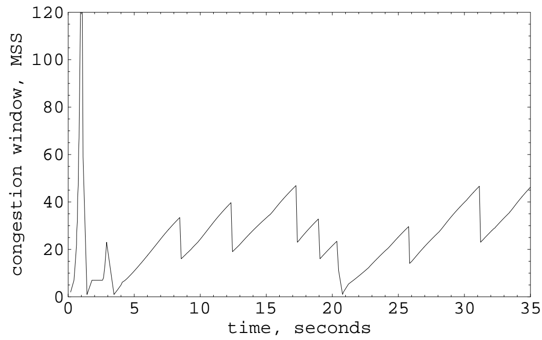

(c) NewReno with pacing, 8 Mbps

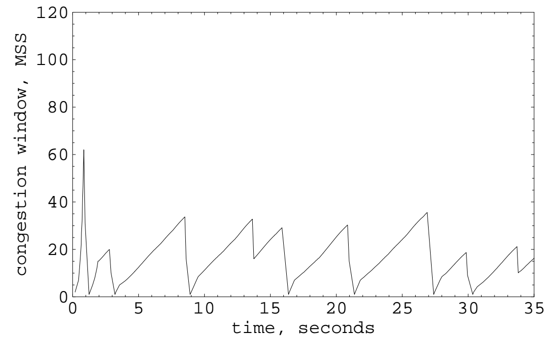

(d) NewReno with pacing, 4 Mbps

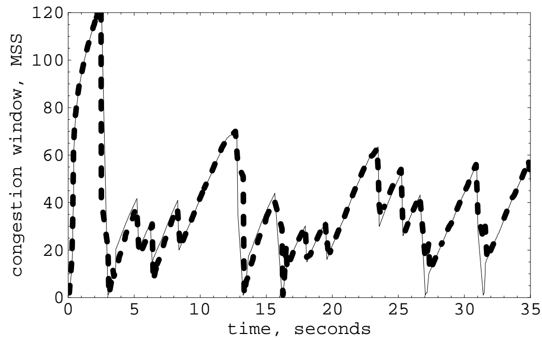

(e) NewReno with smartacking, 8 Mbps

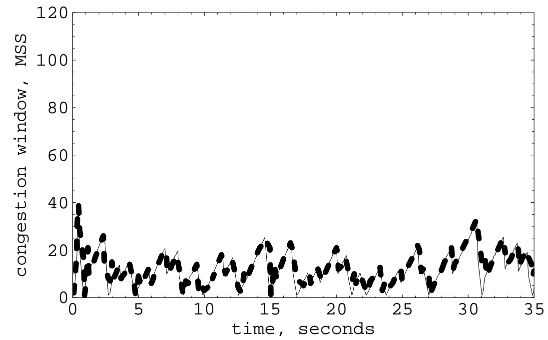

(f) NewReno with smartacking, 4 Mbps

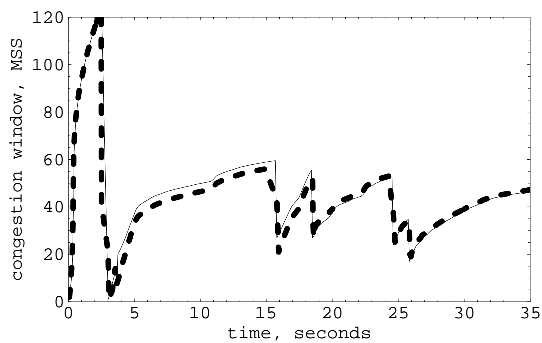

(g) NewReno with smartacking and pacing, 8 Mbps

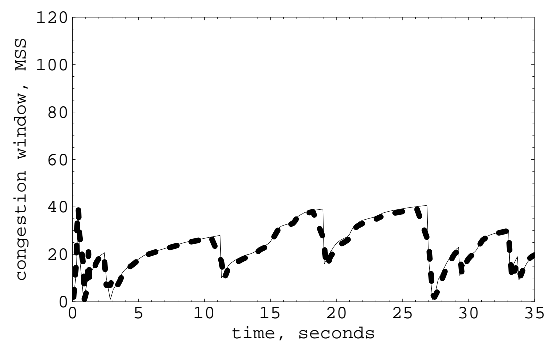

(h) NewReno with smartacking and pacing, 4 Mbps

Figure 7: Congestion windows at the sender (solid lines) and their estimates at the receiver (dotted lines) in the 2-2 configuration for two values of the bottleneck link capacity.

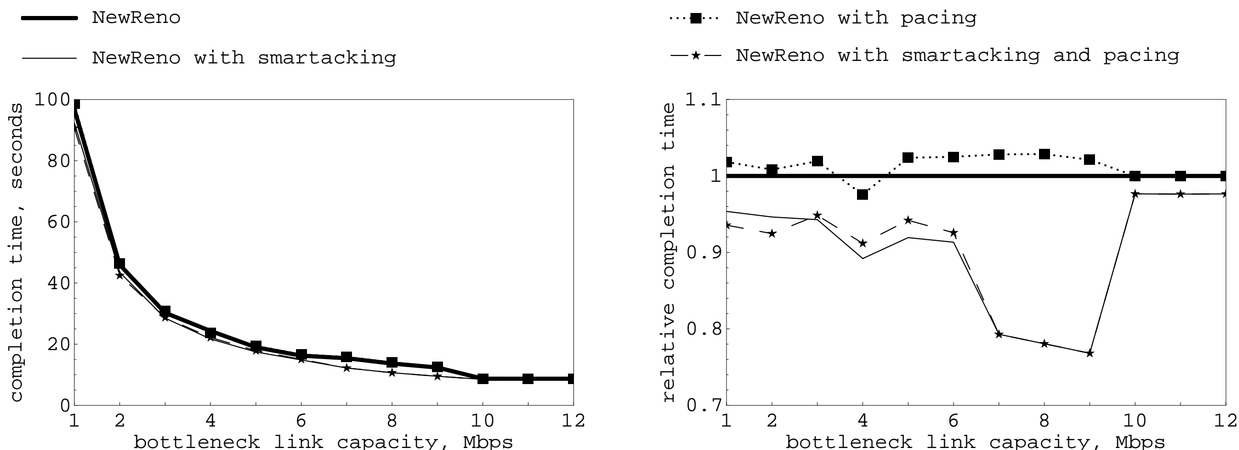

Figure 8: Completion times in the 1-0 configuration.

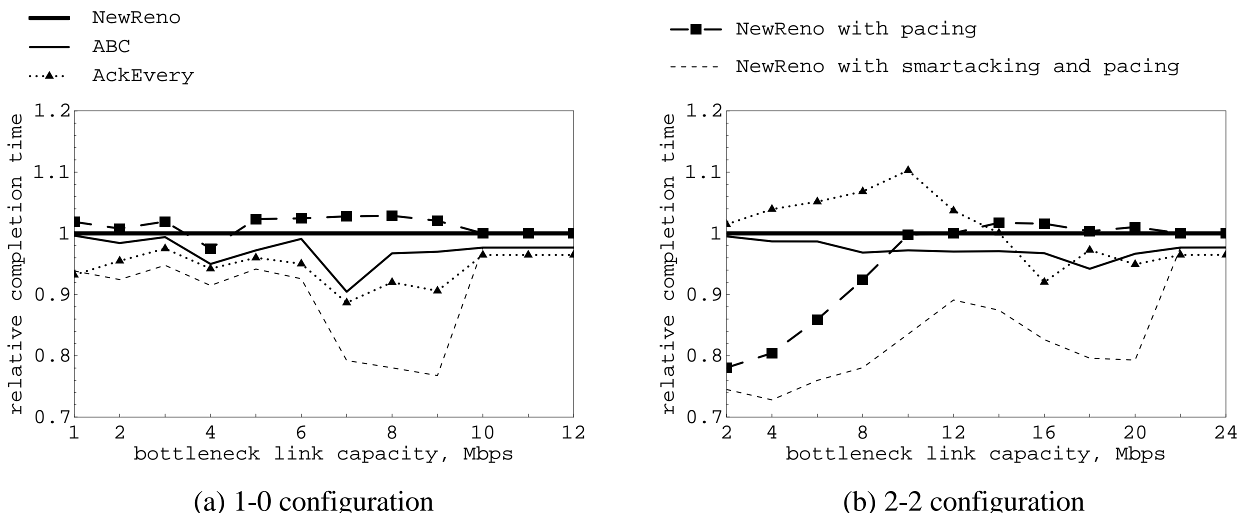

(a) 1-0 configuration

(b) 2-2 configuration

Figure 9: Comparison of smartacking with byte counting.

5.2 Long-Term Performance in Networks of Droptail Routers

For the 1-0 configuration, Figure 8a shows completion times of connections using NewReno, NewReno with pacing, NewReno with smartacking, or NewReno with smartacking and pacing. Figure 8b plots relative completion times computed as ratios of the connection completion times to the completion time of the NewReno connection. The implementations with smartacking perform consistently better than the schemes without smartacking, with up to 20% reduction in the completion time. Both smartacking implementations achieve similar performance without any substantial benefits from pacing.

Figure 9 compares smartacking (represented by NewReno with smartacking and pacing) with byte counting (represented by ABC) in the 1-0 and 2-2 configurations while NewReno, AckEvery, and NewReno with pacing form a background for the comparison. By increasing the rate of the window growth during the slow start, ABC typically provides a smaller completion time than NewReno. However, smartacking consistently yields a much larger reduction. This further improvement is due to the smoother window growth when the bottleneck link capacity becomes saturated.

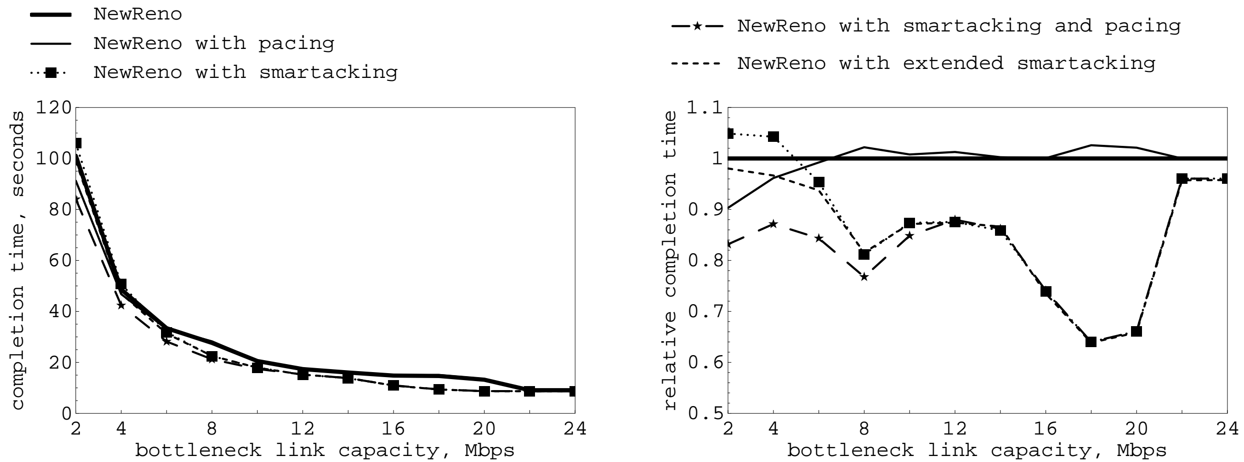

Figure 10: Completion times in the Different-RTT configuration.

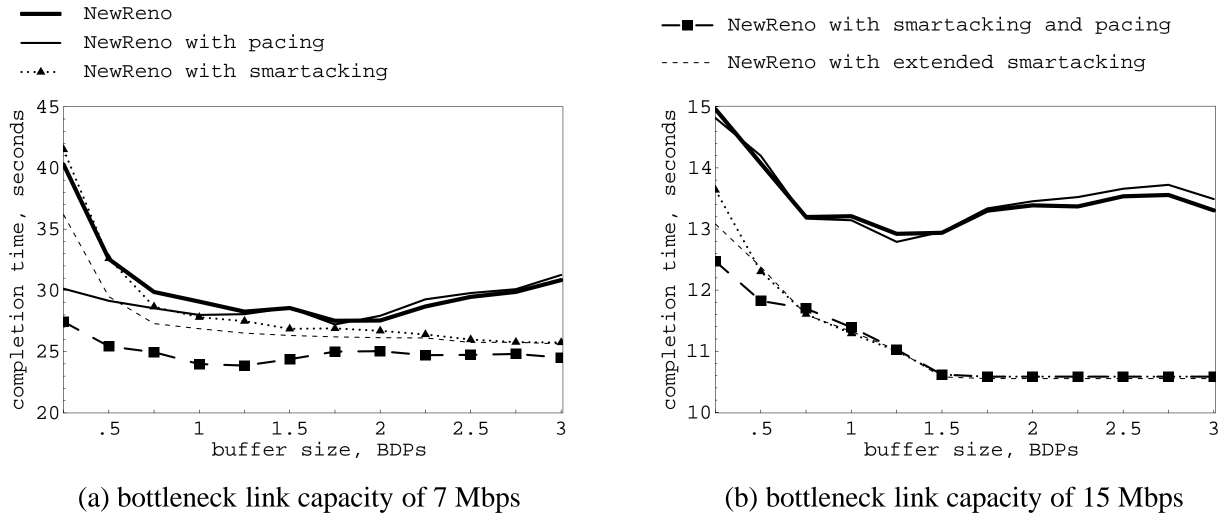

(a) bottleneck link capacity of 7 Mbps

(b) bottleneck link capacity of 15 Mbps

Figure 11: Impact of the buffer size of the bottleneck link on completion times in two 2-2 configurations with different capacities of the bottleneck link.

Figure 10 presents experimental results for the Different-RTT configuration. In general, the smartacking versions provide lower completion times. However, to achieve the improvement for smaller capacities of the bottleneck link shared by connections with different RTTs, the receiver needs the sender's assistance in the form of pacing or α adjustments.

It is well known that TCP performance depends on the size of the bottleneck link buffer. A common rule of thumb prescribes setting the buffer size to one bandwidth-delay product (BDP) [16]. To evaluate the impact of the buffer size on the benefits from smartacking, we conduct experiments in the 2-2 configurations where the bottleneck link has capacities of 7 Mbps and 15 Mbps, and its buffer size varies from 0.3 BDP to 3 BDP. Figure 11 confirms that the buffer size of 1 BDP provides the traditional TCP versions with near optimal performance. The optimal buffer size for the smartacking implementations appears to be somewhat larger. We attribute the increase in the optimal buffer size to extra delay imposed by smartacking on ACKs at the receiver. The graphs show that smartacking improves TCP performance over a wide range of examined

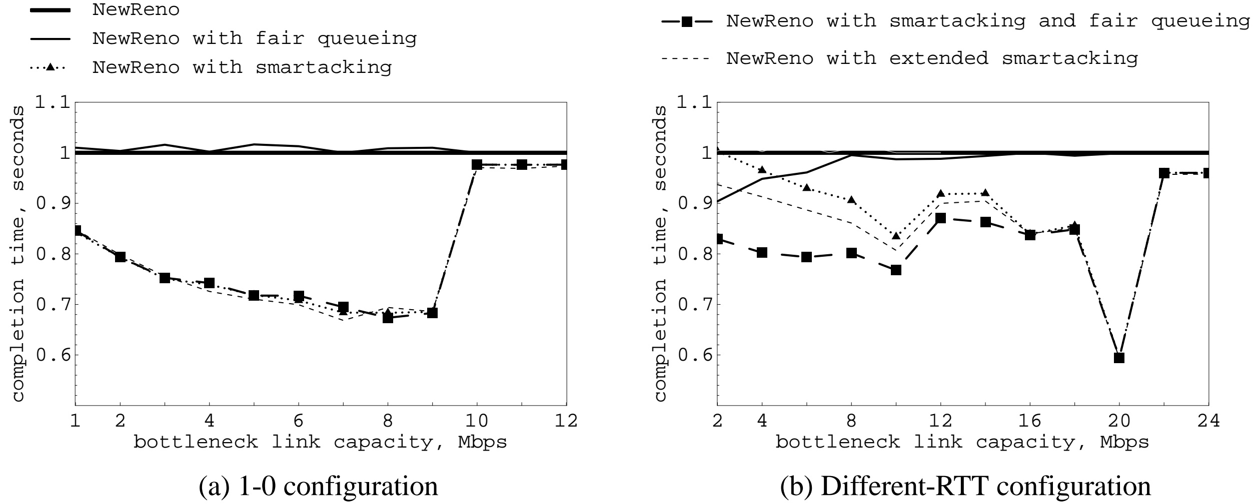

(a) 1-0 configuration (b) Different-RTT configuration

Figure 12: Completion times in networks with RED routers.

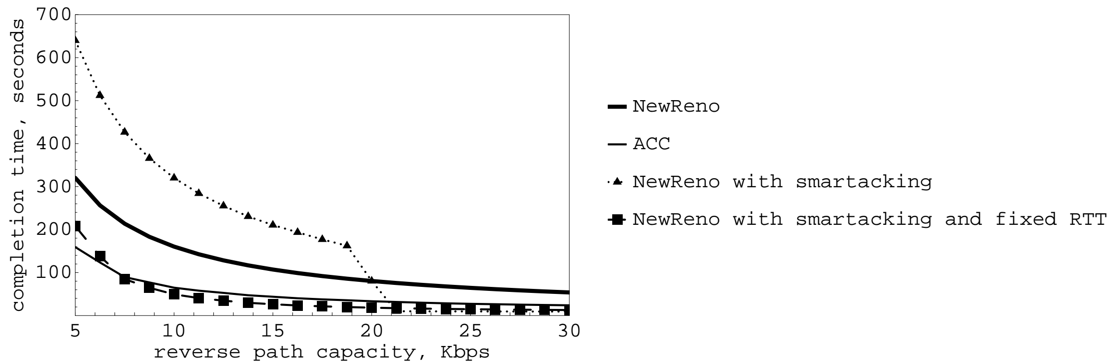

Figure 13: Completion times in the Asymmetric configuration.

buffer sizes. Furthermore, the relative performance improvement becomes higher as the buffer size of the bottleneck link increases.

5.3 Networks with RED Routers

Figure 12 presents some of our extensive experimental results for networks with RED routers. In general, smartacking TCP extensions (represented in our graphs by NewReno with smartacking, NewReno with extended smartacking, and NewReno with smartacking and fair queueing) consistently outperform schemes with no smartacking (represented by NewReno and NewReno with fair queueing). In the 1-0 configuration, all the implementations of smartacking behave similarly. However, in the Different-RTT configuration where two connections compete for the bottleneck link capacity in each direction, the smartacking extensions with advanced features – such as dynamic α or fair queueing – provide larger reductions in the completion time than the reductions offered by our receiver-only implementation of smartacking.

5.4 Asymmetric Networks

We compare performance of smartacking implementations and ACC in the Asymmetric configuration used originally in the studies proposing ACC. Our results confirm that ACC provides shorter completion times

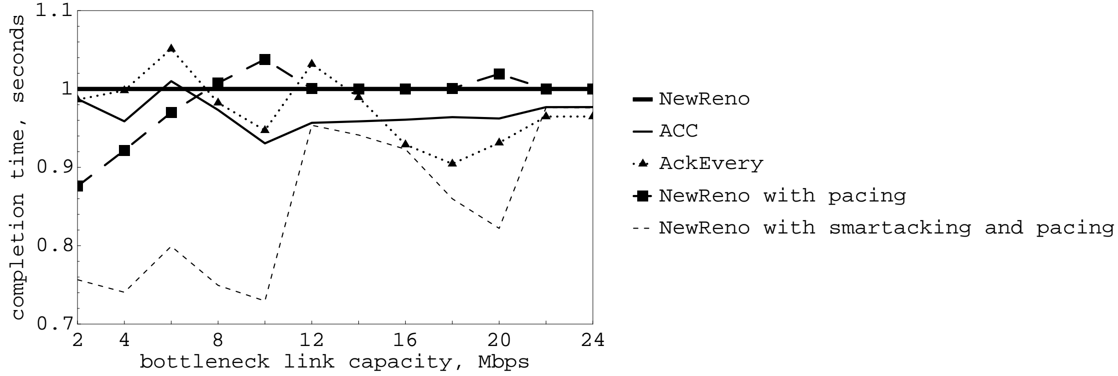

Figure 14: Completion times in the TCP-UDP configuration.

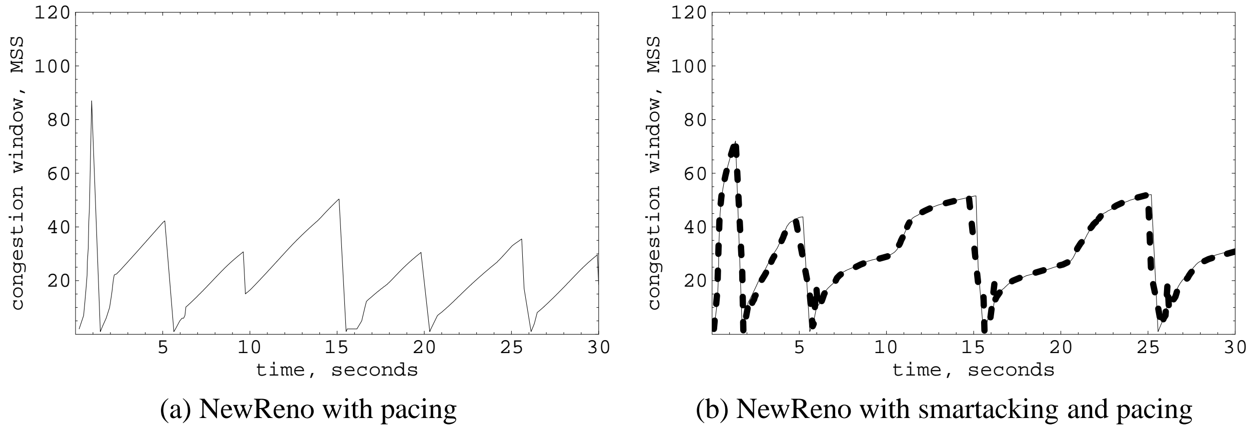

(a) NewReno with pacing

(b) NewReno with smartacking and pacing

Figure 15: Congestion windows at the sender (solid lines) and their estimates at the receiver (dotted lines) in the TCP-UDP configuration with the bottleneck link capacity of 5 Mbps.

than NewReno over the whole range of reverse-path capacities between 5 Kbps and 30 Kbps. We also observe that NewReno with smartacking and pacing reduces the completion times even further. However, Figure 13 demonstrates that smartacking performs less consistently without support from pacing. For reverse-path capacities of at least 21 Kbps, NewReno with smartacking yields lower completion times than ACC. On the other hand, even plain NewReno outperforms NewReno with smartacking when the reverse-path capacity is less than 20 Kbps. Our analysis links this degradation in performance to the RTT estimate at the receiver. When the reverse path has a low capacity, ACKs are frequent enough to saturate the capacity and increase the reverse-path queueing delay. Consequently, the receiver increases its RTT estimate and transmits one ACK per data segment arrival even when the forward-path capacity is fully utilized. To validate the above analysis, we repeat the experiment for NewReno with smartacking when the receiver does not change the RTT estimate after measuring it in the beginning of the connection. The new extension – denoted in Figure 13 as *NewReno with smartacking and fixed RTT* – behaves similarly to NewReno with smartacking and pacing. Hence, not only sender-dependent implementations such as the one with pacing but also receiver-only implementations of smartacking can consistently outperform ACC in asymmetric networks with an extremely low reverse-path capacity.

5.5 Reaction to Changes in the Network Capacity

To evaluate smartacking when the available capacity changes quickly, we compare traditional TCP implementations (NewReno, AckEvery, and NewReno with pacing) to ACC and NewReno with smartacking and

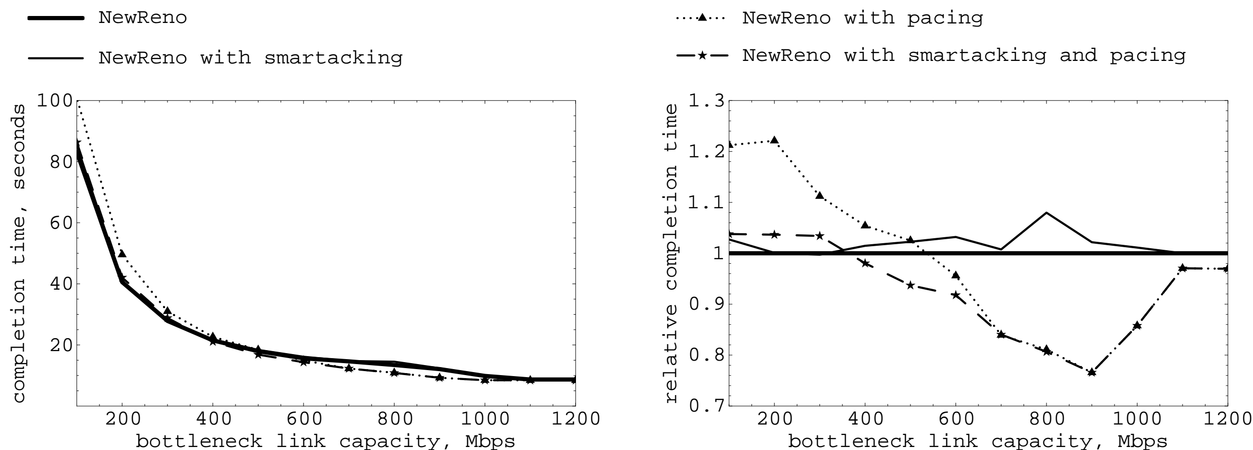

Figure 16: Completion times in the High-Multiplexing configuration.

pacing in the TCP-UDP configuration. Figure 14 shows that addition of smartacking reduces the completion time and consistently provides a lager improvement in performance than ACC.

Figure 15 illustrates reasons why smartacking is beneficial in networks with variable capacities. The graphs trace the congestion windows for NewReno with pacing and NewReno with smartacking and pacing in the TCP-UDP configuration where the bottleneck link capacity is 5 Mbps. When the UDP flow starts transmitting, the influx of its packets congests the network and causes losses. Both NewReno implementations notice the congestion, curb their transmission, and then probe for the new available capacity. Initially, the smartacking receiver resumes sending one ACK per data segment arrival. As the utilization of the bottleneck link returns to its capacity, the smartacking receiver decreases the ACK frequency, and this yields the desired reduction in the congestion window growth. However, when the UDP flow turns quiet and thereby releases the half of the bottleneck link capacity, the smartacking receiver reacts to the change by returning to the highest frequency of ACKs. Once again, as the bottleneck link gets saturated, NewReno with smartacking and pacing reduces the ACK frequency and congestion window growth. In contrast, NewReno with pacing notices neither the newly released capacity nor the approaching saturation of the bottleneck link and continues to grow the window in the same linear fashion. Our experiments also show that smartacking recognizes the availability of the new capacity much faster than ACC which takes up to sixteen RTTs to increase the ACK frequency to one ACK per data segment arrival.

5.6 Networks with High Levels of Connection Multiplexing on Bottleneck Links

So far, we experimented in environments where the degree of connection multiplexing on bottleneck links was low. What happens for higher levels of connection multiplexing? Figure 16 reports completion times for the High-Multiplexing configuration with 200 concurrent connections. The results are representative of all our experiments with high levels of connection multiplexing: smartacking provides no tangible improvements in efficiency of congestion control when the number of connections is large. This conclusion is not surprising. With many connections sharing the bottleneck link, inefficiencies of an individual connection are counterbalanced by other connections. Hence, smartacking has a smaller headroom for improving the efficiency. Furthermore, our ISAT-based mechanism for detecting available capacity becomes less precise. On the other hand, our experiments with smartacking in high-multiplexing settings reveal no negative impact substantial enough to offset the great benefits from smartacking in low-multiplexing environments.

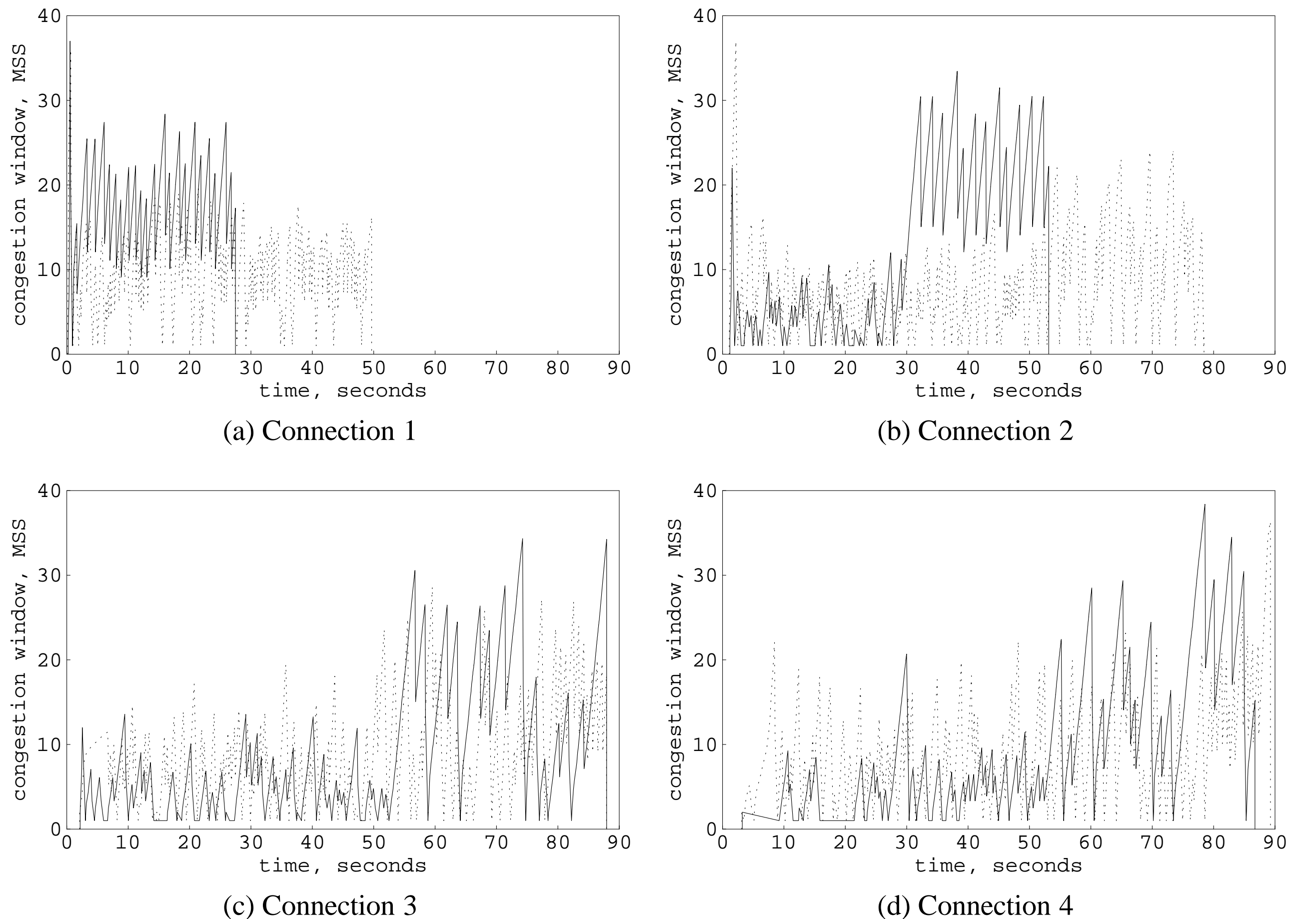

(a) Connection 1 (b) Connection 2

(c) Connection 3 (d) Connection 4

Figure 17: Congestion windows in the Parking-Lot topology with NewReno (solid lines) and NewReno with smartacking (dotted lines).

Furthermore, our results in Section 5.7 indicate that even configurations with many concurrent connections can benefit from smartacking because smartacking improves fairness of network sharing.

5.7 Fairness of Network Sharing

We study the impact of smartacking on fairness by considering first the issue of intra-protocol fairness and then shifting our attention to TCP-friendliness.

Figure 17 traces the congestion windows of the four connections in the Parking-Lot topology with Drop-tail routers and shared link capacities of 4 Mbps. Since the differences between the starting times of the connections are minor in comparison to the completion times, the last of the three shared links serves as a bottleneck for all the connections most of the time. Hence, under a fair allocation of the bottleneck link capacity, the connections should have similar completion times. However, traditional TCP extensions discriminate against connections with long RTTs. Figure 17 confirms this phenomenon for NewReno. Connection 1, which has the smallest RTT, grabs most of the bottleneck capacity and finishes much earlier than the three others. Having the second smallest RTT, connection 2 inherits the domination over the bottleneck link and finishes much before connections 3 and 4 which share the longest RTT. Then, the remaining connections 3 and 4 take turns in grabbing a larger portion of the bottleneck capacity. Figure 17 also demonstrates that smartacking helps NewReno to improve substantially the fairness of the bottleneck link sharing. The congestion windows of all the connections become more stable and similar to each other. Consequently, the completion times of the connections become less diverse as well.

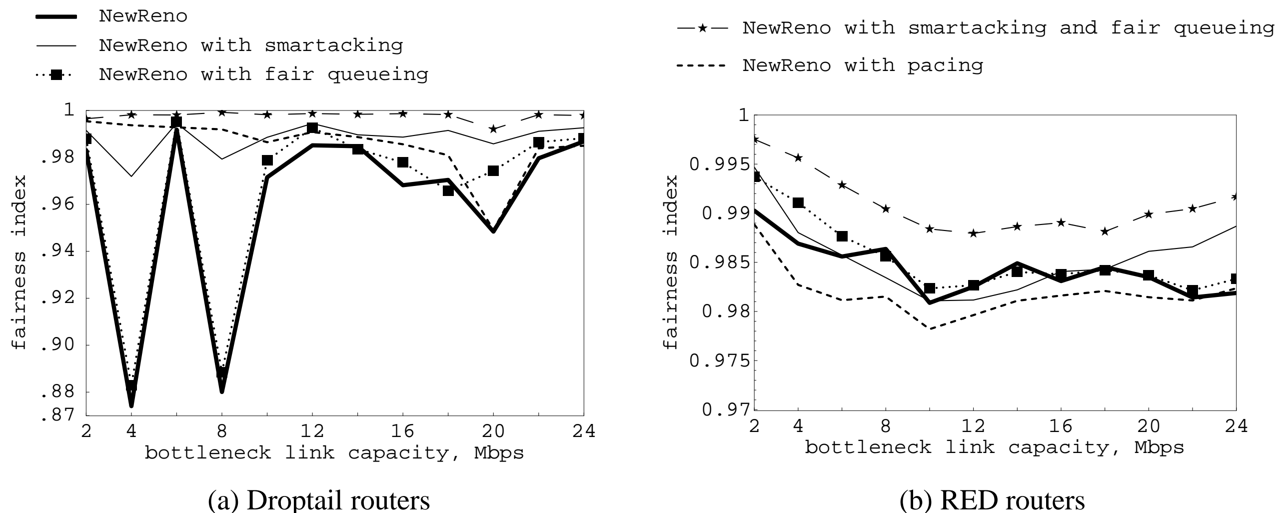

(a) Droptail routers

(b) RED routers

Figure 18: Intra-protocol fairness in the Parking-Lot topology.

Figure 18a reports the fairness index for five TCP extensions in the Parking-Lot topology with Droptail routers. Addition of smartacking to New Reno or NewReno with pacing improves the intra-protocol fairness consistently. As expected, NewReno with smartacking and fair queueing yields the highest fairness index among the three examined implementations of smartacking. Figure 18b shows the fairness index for the Parking-Lot topology with RED routers. By discarding packets probabilistically before the link buffer gets full, the RED router of the bottleneck link allows New Reno to raise its low fairness index in the scenarios where Droptail buffer management starves a connection by sending it into a series of retransmission timeouts. In general, switching to RED routers results in a smoother fairness index for each of the examined TCP extensions.

To evaluate TCP-friendliness of smartacking, we conduct experiments in the 2-2 configuration where three connections employ NewReno while the fourth connection uses either NewReno, or NewReno with pacing, or NewReno with smartacking, or NewReno with smartacking and pacing. Figure 19a reports completion times for the fourth connection as well as for the NewReno connection that delivers data in the same direction. Figure 19b shows that adoption of smartacking by the fourth connection not only does not harm the parallel NewReno connection but also helps the NewReno connection to reduce its completion time significantly.

6 Conclusion

We presented *smartacking*, a technique that improves performance of TCP via adaptive generation of ACKs at the receiver: when the bottleneck link is underutilized, the receiver transmits an ACK for each delivered data segment and thereby allows the connection to acquire the available capacity quickly; when the bottleneck link is at its capacity, the receiver sends ACKs with a lower frequency reducing the control traffic overhead and slowing down the congestion window growth to utilize the network capacity more effectively. Smartacking estimates availability of the network capacity by measuring the inter-segment arrival time (ISAT) at the receiver. Our experiments confirmed that this estimation mechanism operates precisely. In particular, when UDP or TCP cross traffic ceases, ISAT measurements promptly reflect the freed network capacity and boost the rate of ACKs; then, the triggered faster growth of the congestion window enables the connection to capture the released capacity aggressively. Also, since smartacking allows a new TCP

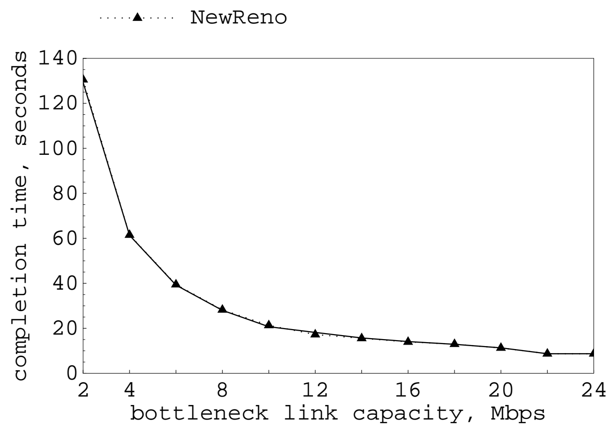

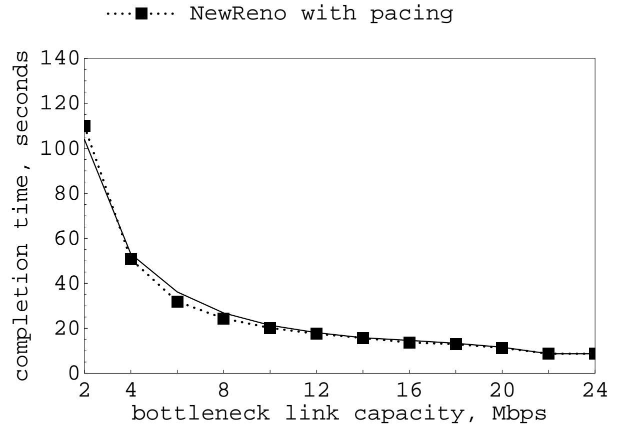

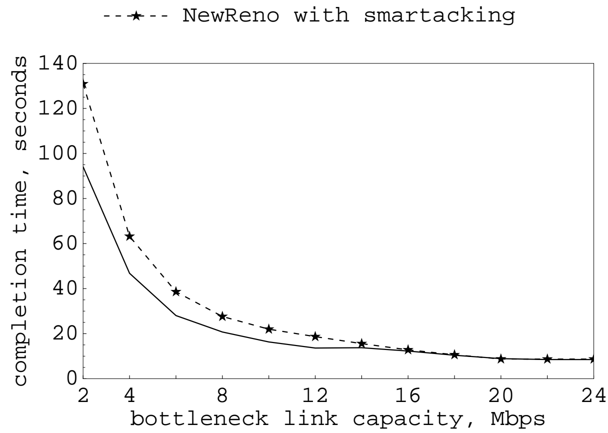

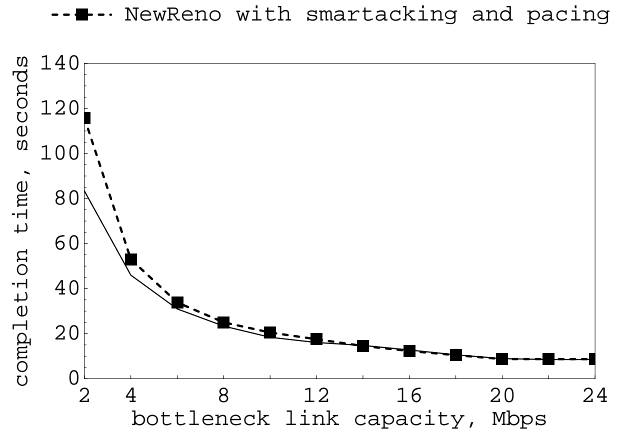

(a) completion times for a NewReno connection (solid line) and parallel connection (as identified above each graph)

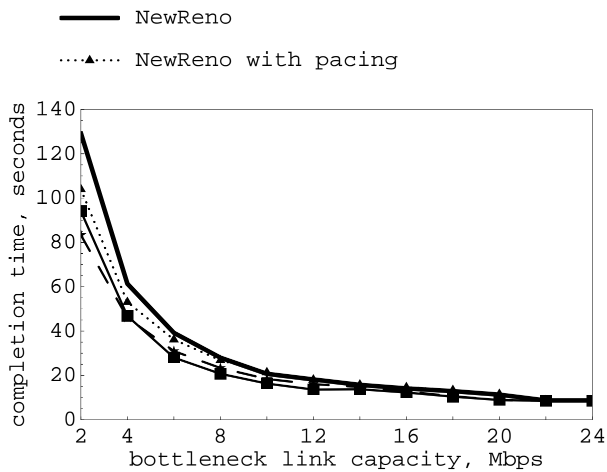

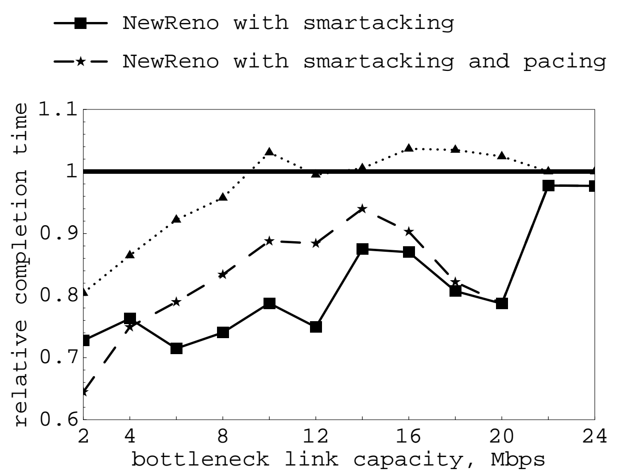

(b) completion times for the NewReno connection with different types of cross traffic

Figure 19: TCP-friendliness of smartacking in the 2-2 configuration.

connection to raise its congestion window quickly, the technique is particularly beneficial in networks with many short-lived connections.

To promote quick deployment of smartacking, our primary implementation of the technique modified only the receiver. This implementation estimates the sender's congestion window using a novel algorithm that has independent value, e.g., for ACC and other protocols where the receiver must know the congestion window. We also considered different implementations of smartacking where the sender or network provides the receiver with explicit assistance such as pacing or fair queueing.

Our experiments in a wide variety of settings showed that all the considered implementations of smartacking help TCP to be substantially more efficient in networks with low levels of connection multiplexing. In networks with high levels of connection multiplexing, efficiency gains from smartacking are negligible because of two reasons. First, the ISAT-based mechanism for detecting availability of the network capacity is less precise when the number of connections is large. Second, with a high level of connection multiplexing, TCP utilizes the bottleneck link quite efficiently even without smartacking. On the other hand, our experiments indicated that networks with many connections on bottleneck links can also benefit from smartacking because smartacking improves fairness of network sharing.

Based on our findings, we believe that smartacking represents a promising approach for improving TCP. However, additional extensive studies over real networks are needed before the technique becomes ready for wide deployment. Whereas this paper presented implementations of smartacking for TCP NewReno, another direction for future work is to implement smartacking for other TCP versions such as SACK. Also, it seems enticing to combine smartacking with ACK filtering [5] in order to derive a more effective protocol for asymmetric networks.

Acknowledgments

The authors would like to thank Sally Floyd, Guru Parulkar, Jon Turner, George Varghese, Marcel Waldvogel, and Ellen Zegura for their extremely valuable feedback.

References

[1] M. Allman. On the Generation and Use of TCP Acknowledgments. *ACM SIGCOMM Computer Communication Review*, 28(5):4–21, October 1998.

[2] M. Allman. TCP Byte Counting Refinements. *ACM SIGCOMM Computer Communication Review*, 29(3):14–22, July 1999.

[3] M. Allman, C. Hayes, and S. Ostermann. An Evaluation of TCP with Larger Initial Windows. *ACM SIGCOMM Computer Communication Review*, 28(3):41–52, July 1998.

[4] M. Allman, V. Paxson, and W. Stevens. TCP Congestion Control. RFC 2581, April 1999.

[5] H. Balakrishnan, V. Padmanabhan, and R. Katz. The Effects of Asymmetry on TCP Performance. In *Proceedings ACM/IEEE MobiCom 1997*, September 1997.

[6] J. Bennett and H. Zhang. Hierarchical Packet Fair Queueing Algorithms. *IEEE/ACM Transactions on Networking*, 5(5):675–689, October 1997.

[7] R. Braden. Requirements for Internet Hosts - Communication Layers. RFC 1122, October 1989.

[8] R. Braden, D. Clark, J. Crowcroft, B. Davie, S. Deering, D. Estrin, S. Floyd, V. Jacobson, G. Minshall, C. Partridge, L. Peterson, K.K. Ramakrishnan, S. Shenker, J. Wroclawski, and L. Zhang. Recommendations on Queue Management and Congestion Avoidance in the Internet. RFC 2309, April 1998.

[9] L.S. Brakmo, S.W. O'Malley, and L.L. Peterson. TCP Vegas: New Techniques for Congestion Detection and Avoidance. In *Proceedings ACM SIGCOMM 1994*, August 1994.

[10] D.D. Clark. Window and Acknowledgement Strategy in TCP. RFC 813, July 1982.

[11] A. Demers, S. Keshav, and S. Shenker. Analysis and Simulation of a Fair Queueing Algorithm. In *Proceedings ACM SIGCOMM 1989*, September 1989.

[12] D. Ely, N. Spring, D. Wetherall, S. Savage, and T. Anderson. Robust Congestion Signaling. In *Proceedings IEEE ICNP 2001*, November 2001.

[13] S. Floyd and V. Jacobson. On Traffic Phase Effects in Packet-Switched Gateways. *Internetworking: Research and Experience*, 3(3):115–156, September 1992.

[14] S. Floyd and V. Jacobson. Random Early Detection Gateways for Congestion Avoidance. *IEEE/ACM Transactions on Networking*, 1(4):397–413, August 1993.

[15] S. Floyd and T. Henderson. The NewReno Modification to TCP's Fast Recovery Algorithm. RFC 2582, April 1999.

[16] V. Jacobson. Modified TCP Congestion Control Algorithm. End2end-interest mailing list, April 1990.

[17] R. Jain. *The Art of Computer Systems Performance Analysis: Techniques for Experimental Design, Measurement, Simulation, and Modeling*. Wiley-Interscience, April 1991.

[18] M. Mathis, J. Mahdavi, S. Floyd, and A. Romanow. TCP Selective Acknowledgment Options. RFC 2018, October 1996.

[19] UCB/LBNL/VINT Network Simulator ns-2. http://www-mash.cs.berkeley.edu/ns, May 2004.

[20] P. Patel, A. Whitaker, D. Wetherall, J. Lepreau, and T. Stack. Upgrading Transport Protocols using Untrusted Mobile Code. In *Proceedings ACM SOSP 2003*, October 2003.

[21] K. Poduri and K. Nichols. Simulation Studies of Increased Initial TCP Window Size. RFC 2415, September 1998.

[22] J. Postel. User Datagram Protocol. RFC 768, October 1980.

[23] K.K. Ramakrishnan and S. Floyd. A Proposal to Add Explicit Congestion Notification (ECN) to IP. RFC 2481, January 1999.

[24] S. Savage, N. Cardwell, D. Wetherall, and T. Anderson. TCP Congestion Control with a Misbehaving Receiver. *ACM Computer Communications Review*, 29(5):71–78, October 1999.

[25] T. Shepard and C. Partridge. When TCP Starts Up With Four Packets Into Only Three Buffers. RFC 2416, September 1998.

[26] L. Zhang, S. Shenker, and D. Clark. Observations and Dynamics of a Congestion Control Algorithm: The Effects of Two-Way Traffic. In *Proceedings ACM SIGCOMM'91*, September 1991.