

Robust TCP Congestion Recovery

Haining Wang[†]

[†]Computer Science Department
College of William and Mary
Williamsburg, VA 23187
hnw@cs.wm.edu

Kang G. Shin[‡]

[‡]EECS Department
The University of Michigan
Ann Arbor, MI 48109
kgshin@eecs.umich.edu

Abstract

This paper presents a new robust TCP congestion-recovery scheme to (1) handle bursty packet losses while preserving the self-clocking capability; (2) detect a TCP connection’s new equilibrium during congestion recovery, thus improving both link utilization and effective throughput; and (3) make the TCP behavior during congestion recovery very close to that during congestion avoidance, thus “extending” the performance model for congestion avoidance to that for TCP loss recovery. Moreover, the new recovery scheme requires only a slight modification to the sender side of TCP implementation, thus making it widely deployable.

The performance of the proposed scheme is evaluated for scenarios with many TCP flows under the drop-tail and RED gateways in the presence of bursty packet losses. The evaluation results show that the new scheme achieves at least as much performance improvements as TCP SACK and consistently outperforms TCP New-Reno. Furthermore, its steady-state TCP behavior is close to the ideal TCP congestion behavior. Since the proposed scheme does not require selective acknowledgments nor receiver modifications, its implementation is simpler than SACK, and only the servers in the Internet need to be modified while keeping intact millions of clients scattered in the Internet.

Keywords — Packet loss recovery, TCP, congestion control, performance evaluation

1 Introduction

The Transmission Control Protocol (TCP) [36] is a reliable, unicast data-transfer protocol used widely for numerous applications, including electronic mail, file transfer, remote login and WWW. The large-scale deployment of TCP in the Internet is due mainly to its robustness in heterogeneous networking environments. The congestion avoidance and control mechanisms of TCP initiated by Jacobson [19] have made significant impacts on the performance and behavior of the Internet [35, 40], and have been playing a critical role in achieving the robustness of the Internet. As the Internet continues to grow rapidly in size and scope, the increasing demand for network resources has increased packet-loss rate in the Internet, and bursty packet losses are reported to be common [31, 35]. Providing a robust congestion-recovery mechanism is, therefore, an important and difficult task for TCP implementation.

The popularity of BSD4.3 makes its TCP implementation — TCP Reno — be widely deployed. TCP Reno includes slow-start and congestion avoidance [19], as well as fast retransmit and fast recovery [20, 21, 38]. However, TCP Reno is shown to perform poorly in recovering from bursty losses within a

window of data packets [10, 27]. When multiple packets within the same window are lost, the fast-recovery algorithm treats each packet loss in a window as an independent congestion signal, thus halving the congestion window multiple times. The TCP Reno's drastic reduction of congestion window size, plus its over-estimation of data packets in flight, inhibits the transmission of new data packets, losing its self-clocking ability. A retransmission timeout is triggered and then slow-start begins to recover from packet losses, causing a substantial performance degradation. Recent Internet measurements [6] show that the majority of timeouts in TCP Reno are caused by bursty packet losses. Also, it is observed [17] that the performance gain of TCP Vegas [7] over TCP Reno is due mainly to TCP Vegas's new techniques for slow-start and congestion recovery, which are closely related to reduction and recovery of bursty packet losses, not the innovative congestion-avoidance mechanism in TCP Vegas. Thus, a robust TCP mechanism that can recover from bursty packet losses without causing timeouts is key to achieve high TCP performance.

Several enhancements to TCP Reno's congestion-recovery mechanism have been proposed, including the modified fast recovery in TCP New-Reno [18], SACK and FACK TCP [27, 28] for recovering from multiple packet losses within the same window of data. While SACK and FACK TCP can efficiently recover from multiple packet losses in a window of data packets, they add significant complexity to TCP implementation, both at the sender and receiver sides. The main weakness of SACK and FACK is that they require "cooperative" receivers. Considering hundreds of millions of clients scattered in the Internet, this requirement makes them practically unattractive for large-scale deployment in the Internet.

The latest characterization of the TCP behaviors in major web servers [32] indicates the growing number of web servers, which run newer versions of Linux and Solaris operating systems, using TCP New-Reno. However, although TCP New-Reno does not require selective acknowledgments and has better capability to recover from multiple packet losses in a window of data than TCP Reno, its ability to keep the "flywheel" of ACKs and data packets and prevent the loss of self-clocking, depends on the TCP window size at the time when the first packet loss is detected, as well as on the number of packets lost within a window. TCP New-Reno's ability to recover from packet losses is limited by its inherent weaknesses, including:

- In TCP New-Reno, the number of new data packets sent out per round-trip time (RTT) decreases exponentially due to its policy "one new data packet is sent out upon receipt of two duplicate ACKs" during the entire congestion-recovery period. Since TCP New-Reno can only recover from one dropped packet per RTT, this rapid decrease will eventually stop the flow of returning ACKs (hence, loss of self-clocking), and a coarse timeout will follow.
- During congestion recovery, TCP New-Reno only passively recovers from the dropped packets. The exponentially-decreased amount of data transmitted during each RTT lowers link utilization even if it does not cause the loss of self-clocking.
- TCP New-Reno cannot detect further data losses that might occur to the new data packets

sent out during congestion recovery. It has to resort to another trigger of fast retransmit or a retransmission timeout to detect such packet losses.

To reduce coarse timeouts and improve the effectiveness of fast retransmit under a tiny window condition, *right-edge* recovery [5] has been proposed, in which “one new data packet is sent out upon receipt of each duplicate ACK, instead of two duplicate ACKs.” Similarly, Lin and Kung [25] proposed that a new data packet be generated upon each arrival of first two duplicate ACKs. They retain TCP aggressiveness as before when there is no network congestion. However, the packet conservation rule [19] does not apply when congestion occurs. TCP aggressiveness should be reduced in order to drain the congestion from the network. These transmitted packets on the verge of detection of a packet loss — indicating network congestion — may add more fuel to the “fire” at the congested bottleneck. Also, these enhancements cannot quickly detect further data losses during congestion recovery either. To reduce the occurrence of bursty losses from a window of data, Smooth-start [42] has been proposed as an optimization of the Slow-start algorithm, which is orthogonal to the enhanced recovery schemes.

In this paper, we propose a robust TCP congestion recovery — called *Robust Recovery* (RR) — algorithm to make a TCP flow more robust to bursty packet losses. The key features of RR include:

- The amount of data in flight is accurately measured, since congestion window size (*cwnd*) overestimates the number of packets in flight during congestion recovery, stalling data transmission.
- RR treats bursty packet losses within a window as a single congestion signal. Like TCP New-Reno, RR exponentially backs off the sending rate after detecting the first packet loss within a window. However, the exponential decrease in the amount of data injected into the network does not last for the entire recovery period. The exponential decrease is applied only during the first RTT of the recovery period, which is consistent with the treatment of single congestion signal.
- By keeping track of the number of new data packets arrived at the TCP receiver in the previous RTT, the TCP sender can detect any further data loss very quickly without triggering fast retransmit or retransmission timeouts. Upon detection of a further data loss, the TCP sender linearly shrinks the pipe size and extends the exit point of RR.
- After the exponential back-off that happens at the first RTT during congestion recovery, as long as no further packet losses are detected, the TCP sender linearly increases the amount of new data transmitted during each RTT while recovering the dropped packet which is indicated by the arrival of a new partial ACK ¹. During this period, similarly to the *right-edge* recovery [5], a new data packet is transmitted upon receipt of each duplicate acknowledgment.
- Congestion recovery is seamlessly switched to congestion avoidance when all outstanding data packets at the beginning of the last RTT have been acknowledged. The big ACK problem, which

¹A partial ACK acknowledges some but not all of the outstanding packets at the start of the previous RTT.

causes bursty packet transmissions upon exit of congestion recovery, is eliminated without using the “maxburst” parameter in TCP New-Reno and SACK.

In addition to recovering the dropped packets, the RR algorithm finds the new equilibrium of a TCP connection during congestion recovery, so as to achieve higher link utilization while recovering the lost packets. Also, RR makes the TCP behavior during congestion recovery very similar to that during congestion avoidance, thereby enabling the performance model for TCP congestion avoidance [29] to represent that for TCP congestion recovery. This allows for accurate prediction of the TCP-consumed bandwidth even without using selective acknowledgments. Moreover, it can quickly detect further packet losses and recover from such losses without triggering fast retransmit or retransmission timeouts.

The performance benefits of the RR scheme are demonstrated via extensive simulation experiments with the *ns-2* [30]. Our simulation results show that the proposed scheme achieves at least as much performance improvements as TCP SACK and consistently outperforms TCP New-Reno. Furthermore, since it requires neither selective acknowledgments nor receiver modifications, its implementation and deployment is much simpler than that of TCP SACK, and only the servers in the Internet need to be modified slightly, while keeping intact millions of TCP clients scattered in the Internet.

One characteristic of TCP is its dependency on the returning ACKs as the trigger of data transmission and congestion window growth. Similarly, RR relies on the returning duplicate ACKs to maintain self-clocking during congestion recovery. We will elaborate on the effect of ACK losses on TCP congestion recovery in Section 3.3. RR also handles retransmission losses by using timeouts, as is usually done.

The remainder of this paper is organized as follows. Section 2 briefly surveys the current TCP congestion-recovery mechanisms, putting RR in a comparative perspective. Section 3 describes the proposed RR algorithm, and Section 4 presents its performance evaluation results in the presence of drop-tail and RED gateways. Section 5 assesses the fitness of the proposed algorithm to the ideal congestion-avoidance model. Section 6 discusses the incremental deployability of the proposed RR algorithm along with the TCP. Finally, the paper concludes with Section 7.

2 Current TCP Congestion-Recovery Mechanisms

The evolution of TCP congestion-control schemes is marked by the progress of packet-loss detection and recovery. The same packet-loss detection mechanism is employed by TCP Tahoe, Reno, New-Reno and SACK, and is triggered either by fast retransmit after receiving three duplicate ACKs, or by a retransmission timeout. Moreover, the TCP behaviors after a retransmission timeout for all of these schemes are similar: retransmit the timed-out packet, set *ssthresh* to a half of *cwnd*, then reduce *cwnd* to 1 and re-enter the slow-start phase. However, these TCP versions differ in the way of recovering from packet losses after a fast retransmit.

In TCP Tahoe, the recovery behavior after triggering a fast retransmit is the same as that after a retransmission timeout. After retransmitting the first lost packet, the TCP connection reduces *cwnd*

to 1 and enters the slow-start phase. Since it may take many RTTs to reach the connection's new equilibrium, link bandwidth will be under-utilized during these RTTs. Another shortcoming of such congestion recovery is that it may retransmit those packets that have already been successfully delivered. The advantage of TCP Tahoe is that it is more robust than Reno and even New-Reno if a string of packets are lost due to heavy network congestion.

In TCP Reno, fast recovery is included; following a fast retransmit, instead of entering the slow-start phase, the TCP sender reduces *cwnd* by half and uses additional incoming duplicate ACKs to inflate *cwnd*. After half a window of duplicate ACKs are received, the sender transmits a new packet for each additional duplicate ACK it received. If only a single packet is dropped from a window of data, then the ACK for the first retransmission will complete the recovery and enter the congestion-avoidance phase, and hence, no retransmission timeout occurs.

However, if multiple packets are dropped from the same window, the ACK for the first retransmission covers only the outstanding data packets that were transmitted before the second packet loss in the window. This partial ACK detects the second packet loss in the window. Unfortunately, TCP Reno fails to utilize partial ACKs for packet-loss recovery. Instead, as soon as the TCP sender receives the partial ACK for the first retransmission, it will prematurely end the recovery and deflate *cwnd* back to *ssthresh*. In most cases TCP Reno can only recover from one or two packet losses without losing self-clocking.

Recognizing the packet-loss signal indicated by a partial ACK, in TCP New-Reno, upon arrival of a partial ACK, the sender retransmits the packet immediately following the partial ACK without taking TCP out of the fast-recovery phase. During the recovery, upon receipt of two duplicate ACKs, a new data packet will be transmitted. The TCP sender does not exit the fast-recovery phase until all of the outstanding packets at the start of fast recovery have been acknowledged.

In TCP SACK, the selective acknowledgment option is implemented. The SACK option field contains two or three SACK blocks. When non-contiguous data is received, each SACK block in the duplicate ACK reports a non-contiguous set of data that has been received and queued at the receiver. Using selective acknowledgments, the TCP sender can recover from multiple packet losses during one RTT. The disadvantage of TCP SACK is that it requires to change implementations of *both* the sender and receiver sides.

TCP Vegas uses more accurate RTT estimation to quickly detect and retransmit a dropped packet without having to wait for 3 duplicate ACKs when a duplicate ACK is received or when the first or second non-duplicate ACK is received after retransmitting a packet. More importantly, TCP Vegas avoids reducing the congestion window multiple times when bursty losses occur within a window of data, which is shown to be the main source of TCP performance gain during the recovery period, rather than its more aggressive fast retransmission policy [17].

The more sophisticated congestion control mechanisms for unicast applications, which include the recently-proposed equation-based congestion control [15] and XCP [22], have the ability of agile flow-rate adaption to achieve high link utilization and fair bandwidth allocation. The transmission rate is

determined by the estimated loss event rate, but the loss-recovery behavior within each loss event is similar to that of TCP New-Reno.

3 Robust TCP Congestion Recovery

To recover bursty packet losses within a window while preserving the self-clocking capability, we propose a new TCP congestion-recovery (RR) algorithm. In RR, the TCP sender not only recovers from packet losses, but also finds the connection’s new equilibrium during congestion recovery. It also makes the recovery behavior of bursty losses within a window of data very close to an ideal congestion-avoidance behavior in which only a single packet loss within a window of data occurs periodically. RR is detailed in the next subsections.

3.1 Accurate Estimation of Data in Flight

One of the key problems with current congestion recovery schemes is that congestion window size ($cwnd$), which represents the outstanding packets at the sender side, is still used as the control “pedal” during congestion recovery.

During congestion recovery, the outstanding packets at the sender side can be divided into three groups: *active*, *dormant*, and *dropped*. The active group is the set of data packets that are in transit, which also include the retransmitted packets. The dormant group is the set of data packets that were transmitted during the past RTTs and have already arrived and queued at the receiver, but have not yet been acknowledged. Actually, each dormant packet has caused the receiver to send a duplicate ACK to the sender. The dropped group is the set of data packets that were lost during the past RTTs. Clearly, the outstanding packets at the sender side as a whole do not represent the data packets in the pipe. Only the active group, in which data packets are in “flight,” represents the data packets in the pipe, since the packets in dormant and dropped groups have left the network either normally or abnormally in the previous RTTs, and do not consume network resources any longer.

Thus, as a measure of the outstanding packets at the sender side, $cwnd$ over-estimates the number of packets in the pipe and is no longer adequate for transmission control during congestion recovery. A new state variable $actnum$ is thus introduced to measure the amount of data in the pipe at each RTT of congestion recovery. During congestion recovery, $actnum$ plays the usual role of $cwnd$ as the means to provide congestion control at the sender side. Once the congestion-recovery phase ends, the congestion-control responsibility is returned to $cwnd$.

A similar variable $pipe$ [10] has been proposed in TCP SACK, which counts the number of outstanding packets in the path, not at the sender side. However, the role of congestion control is still played by $cwnd$. The TCP sender can transmit a data packet only when $pipe < cwnd$. The variable $pipe$ just passively estimates the number of outstanding packets in the path. By contrast, $actnum$ not only represents the number of outstanding packets in the path but also controls the data-transmission rate.

Table 1: TCP parameters in congestion recovery

Name	Meaning of state variables
<i>ndup</i>	the number of duplicate ACKs received.
<i>snd.una</i>	the sequence number of the first byte of unacknowledged data.
<i>snd.nxt</i>	the sequence number of the first byte of unseqnt data.

In each RTT during congestion recovery, *actnum* linearly grows or shrinks according to the network condition.

3.2 Description of Robust Recovery

Several TCP parameters [38] are used to describe RR as listed in Table 1. For a better description of RR, we divide it into two sub-phases — *retreat* and *probe* — and three transient states — *entrance*, *exit*, and *occurrence* of a further data loss. Figure 1 shows a high-level organization of the RR algorithm. The detailed flowchart of the algorithm is given in Figure 2. The state variables referred in this figure are explained in Table 2.

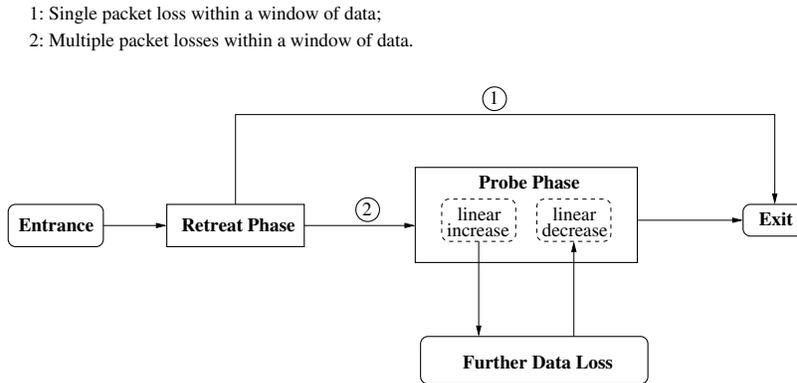


Figure 1: The structure of the RR algorithm

As with the fast recovery, RR is triggered by a fast retransmit. However, in RR *cwnd* remains unchanged until the end of congestion recovery, as it is not used for congestion control in RR. At the very beginning of each RTT during congestion recovery, the state variable *ndup* is initialized to 0. As duplicate ACKs arrive within one RTT, *ndup* measures the number of duplicate ACKs received by the sender.

Note upon the arrival of an out-of-sequence data packet at the receiver, the delayed acknowledgment mechanism is off; the receiver immediately sends out an ACK for each received out-of-sequence data packet. The RR algorithm utilizes this to detect further data losses in the probe sub-phase.

If there are no further data losses, the threshold that determines the end of RR’s congestion recovery is the same as that of New-Reno and SACK, which is the sequence number of the first byte of unseqnt data when the first packet loss was detected by the fast retransmit. The congestion-recovery phase ends

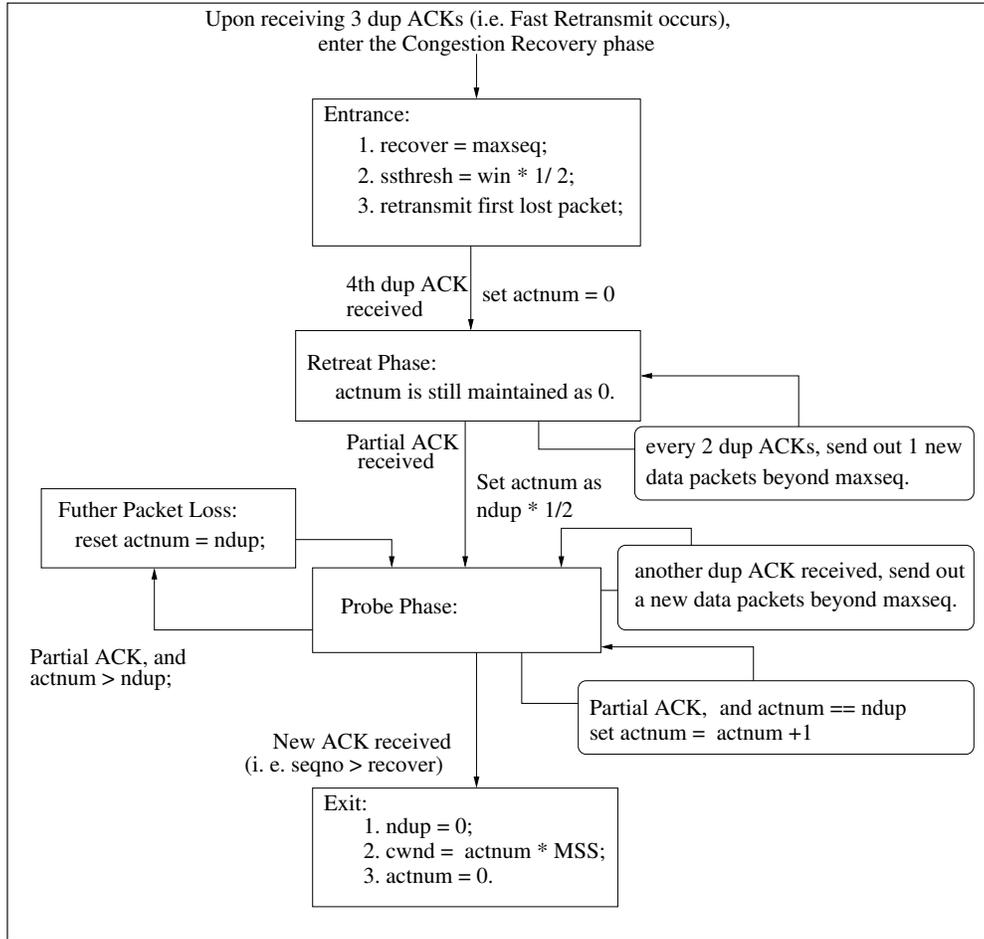


Figure 2: The flowchart of the RR algorithm

as soon as *snd.una* advances to, or beyond, this threshold, indicating that all outstanding data packets at the time of entering the congestion-recovery phase have been acknowledged.

Table 2: State variables in the description of robust recovery

Name	Meaning of state variables
maxseq	the highest sequence number sent so far.
seqno	the sequence number indicated by the currently received ACK.
recover	the highest sequence number sent before receiving the latest dup ACKs.

However, if any further data-packet loss is detected, RR advances the threshold used to detect the end of congestion recovery. In particular, the threshold is updated to the value of *snd.nxt* when a further data-packet loss was detected. The congestion-recovery phase ends when *snd.una* advances to, or beyond, this threshold, which may differ significantly from the value of *snd.nxt* when the congestion-recovery phase was entered.

3.2.1 Retreat Sub-phase

The retreat sub-phase only covers the first RTT of congestion recovery. In this sub-phase, the TCP sender exponentially backs off its transmission rate. Like in New-Reno, during the retreat sub-phase the sender injects one new data packet for every two duplicate ACKs received. So, the data-transmission rate per RTT is reduced at least by half. The end of the retreat sub-phase is marked by the arrival of the first non-duplicate ACK, i.e., an ACK for a higher-sequence numbered packet; if multiple packets had been lost within the same window, it must be the first partial ACK. During the retreat sub-phase, *actnum* remains to be zero.

When the first partial ACK (i.e., the end of the retreat sub-phase) is received, *actnum* assumes the role of congestion control for the rest of congestion recovery. The variable *actnum* is initially set to $ndup * 1/2$, which is the number of new data packets sent out during the retreat sub-phase.

As can be seen from the above description, the end of the first RTT during congestion recovery is pivotal; the role of congestion control at the sender is transferred from *cwnd* to *actnum*, and the retreat sub-phase ends while the probe sub-phase starts. Note that the first lost packet is recovered in the retreat sub-phase. If only a single packet within a window of data is lost, the TCP sender exits the congestion-recovery phase after the retreat sub-phase. However, if multiple packets in a window of data were lost, all but the first of the lost packets are recovered, one per RTT, in the probe sub-phase, in which the sender linearly adjusts the value of *actnum* according to the network condition. The sending TCP can distinguish the two sub-phases by testing if $actnum = 0$.

3.2.2 RTT in the Probe Sub-phase

If there are multiple packet losses within the same window, a key characteristic of the probe sub-phase is that each RTT is distinguished by the receipt of a new partial ACK. At the sender side, the end of the current RTT and the beginning of the next RTT are indicated by the receipt of a new partial ACK. Figure 3 illustrates this feature. Suppose that in a window of data, four packets are dropped and their sequence numbers are 4000, 5000, 7000 and 8000, respectively. The first loss is recovered in the retreat sub-phase, and the rest are recovered in the probe sub-phase, which are represented as 5, 7, 8 in Figure 3.

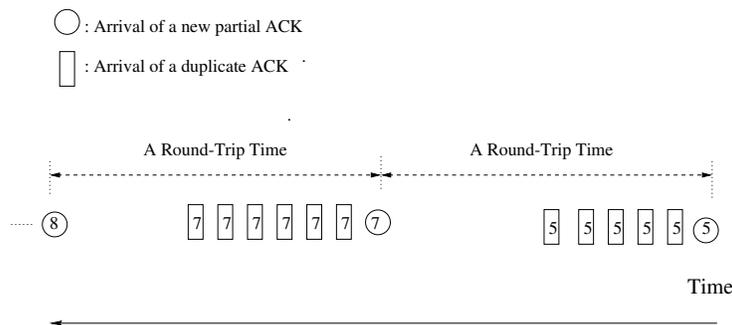


Figure 3: RTT in the probe sub-phase (sender's side)

A packet retransmission is triggered by the arrival of a partial ACK, and this retransmission is acknowledged via the next partial ACK if the packets are delivered in order. If there are no ACK losses, the state variable *ndup* represents the number of new data packets sent out during the last RTT that have been received, because during congestion recovery each received data packet triggers an ACK immediately.

In case of out-of-order delivery, a partial ACK would be generated by the arrival of a new data packet at the receiver, thereby generating one of duplicate ACKs upon arrival of the retransmitted packet. However, since the sender can treat the partial ACK as the ACK of the retransmission, *ndup* still measures the number of new data packets sent during the last RTT that have been received. So, out-of-order delivery does not skew the measurement of the number of new data packets sent during the last RTT that have been received. Unfortunately, this is not valid if there are ACK losses, and we will discuss it later in this section.

3.2.3 Probe Sub-phase

At the very beginning of each RTT of the probe sub-phase, the arrival of a partial ACK triggers an immediate retransmission. Upon receiving each duplicate ACK after this partial ACK, the TCP source sends a new data packet. If there are no further data and ACK losses, at the end of this RTT, the value of *ndup* should be equal to that of *actnum*. The reason for this is: *actnum* indicates the number of new data packets that were sent out during the last RTT; and *ndup* represents the number of new data packets sent out during the last RTT that have been received.

Given no ACK losses, at the end of each RTT, a further data loss can be detected by comparing *ndup* with *actnum*. The equality $ndup = actnum$ indicates no further data loss had occurred. However, if $ndup < actnum$, further data losses had occurred during the last RTT. The difference between *ndup* and *actnum* indicates the number of further data losses.

In case of no further data loss, the sender will increment *actnum* by 1 and send one more new data packet for every RTT until the end of congestion recovery, or until a further data loss is detected, which is similar to the congestion-avoidance algorithm. However, if further data losses are detected, *actnum* is reduced to the value of *ndup*. So, reduction of *actnum* linearly depends on the number of further data losses. Since each duplicate ACK triggers the transmission of a new data packet in this RTT, the reduced *actnum* still indicates the number of data packets in flight.

Figure 4 illustrates the linear adjustment of *actnum*. The rationale behind the linear back-off when further data losses are detected is two-fold. The first is to reduce the disturbance caused by ACK losses. The second is to avoid the drastic reduction of in-flight data since the sender exponentially backs off in the retreat sub-phase that happened only a few RTTs ago.

Once further data losses are detected, the exit of the congestion-recovery phase must advance to recover from them. Any further data loss during the congestion-recovery phase can be identified by a new partial ACK beyond the original exit, and recovered by the subsequent packet retransmission

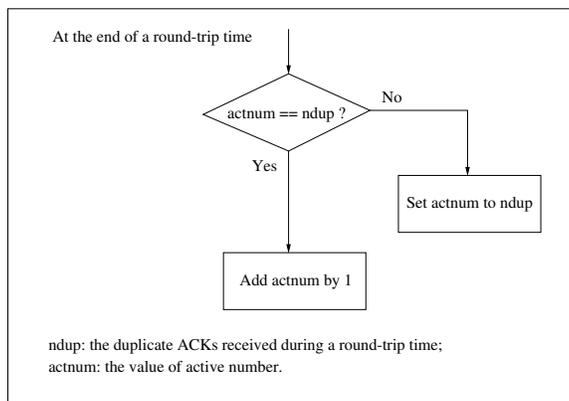


Figure 4: The linear adjustment of *actnum*

without waiting for two more duplicate ACKs. The exit of congestion recovery extends to the point where all outstanding data packets at the start of last RTT, instead of at the time of entering congestion recovery, have been acknowledged.

After the sender recovers from further data losses, it exits the congestion-recovery phase and enters the congestion-avoidance phase. At that time, the role of congestion control is transferred back to *cwnd*. The sender assigns the current value of *actnum* to *cwnd*. Since *cwnd* is measured in bytes, instead of packets, *cwnd* is set to $actnum \times MSS$. Then *actnum* is set to 0 again. Since the reset value of *cwnd* accurately measures the amount of data packets in flight, the arrival of the new ACK that takes the sender out of congestion recovery only triggers a new packet out, which observes the conservation of packets.

As shown in Figure 5, let's assume packet 6 has the highest sequence number at the start of last RTT, packets 7, 8, 9, 10, and 11 have been sent out during last RTT. Upon receipt of ACK for packet 7 that indicates that all data packets up to (including) 7 have been received at the receiver, the sender leaves the congestion-recovery phase for the congestion-avoidance phase. At the same time, there are still four outstanding data packets that have not yet been acknowledged, which are packets 8, 9, 10, and 11. So, resetting *cwnd* can only allow one more data packet (packet 12) to be sent out.

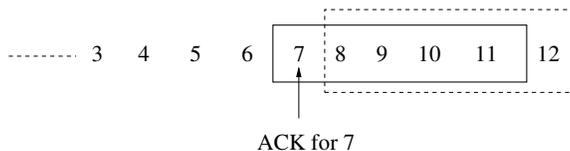


Figure 5: The exit of RR: observe the conservation of packets

The big ACK problem that causes sending packets in burst at the exit of congestion recovery has been eliminated. TCP New-Reno and SACK use a “maxburst” parameter to limit the number of packets that can be sent upon receipt of a single incoming ACK. However, it only limits burstiness but doesn't remove it. Also, it adversely affects bandwidth utilization if the bottleneck has drained all

packets; or it causes potential packet losses if the bottleneck is not yet back to its knee area.

3.3 Effect of ACK Loss upon Congestion Recovery

Since RR also relies on returning ACKs to inject new data into the network, loss of a string of ACKs will cause RR to lose its self-clocking. Note, however, that RR is more robust to ACK losses than New-Reno. Rare ACK losses cause only a slight negative effect upon congestion recovery. In the probe sub-phase, the transmission rate only linearly decreases when an ACK loss falsely signals a further data loss. Although TCP SACK is less vulnerable to ACK losses, it still has to time out if many ACKs are lost, or the ACK for a retransmission is lost, as shown in [8].

Even if there is no data packet loss in the forward path, significant ACK losses in the backward path can cause problems, such as the sender's increased burstiness, a slower growth of congestion window and a decrease in the effectiveness of fast retransmit. There have been some schemes proposed in [4] to resolve these problems. The solutions are to ease the side-effect of infrequent ACKs brought by ACK Congestion Control and ACK Filter [4].

TCP Santa-Cruz [34] has been proposed to be resilient to ACK losses, which decouples the growth of the congestion window from the number of returning ACKs. Also, it provided an efficient loss recovery via a ACK window that is similar to the bit vectors used in TCP SACK. However, the drawback of TCP Santa-Cruz is that it requires the modification of TCP header and a cooperative receiver, just like what TCP SACK requires.

Although data loss on the forward path and the ACK loss on the backward appear uncorrelated in the current Internet [35], we believe that if a fair share is given to each flow at the routers, the loss probability of an ACK packet should be much smaller than that of a data packet. Because the size of ACK packets is usually much smaller than that of data packets — except for those that piggyback other pieces of information — and hence an ACK-packet flow consumes much less network resources than a data-packet flow. A transport-aware IP router architecture [41] has been proposed to differentiate TCP control packets from data packets, which protects the TCP control packets from the congestion caused by data traffic.

3.4 The Sender's Mode

In current congestion control schemes, a sender could be in one of two modes: slow-start or congestion avoidance. In the slow-start mode, the sender exponentially increases the congestion window, whereas, in the congestion-avoidance mode, it linearly increases the congestion window. However, they have not addressed congestion recovery.

Actually, there should be three modes: slow-start, congestion avoidance, and congestion recovery. A sender is in the slow-start mode when it first starts transmitting data or immediately after a retransmission timeout. After the *cwnd* reaches *ssthresh* without any packet loss, the sender enters the congestion-avoidance mode. Once a fast retransmit is triggered, the sender will enter the congestion-

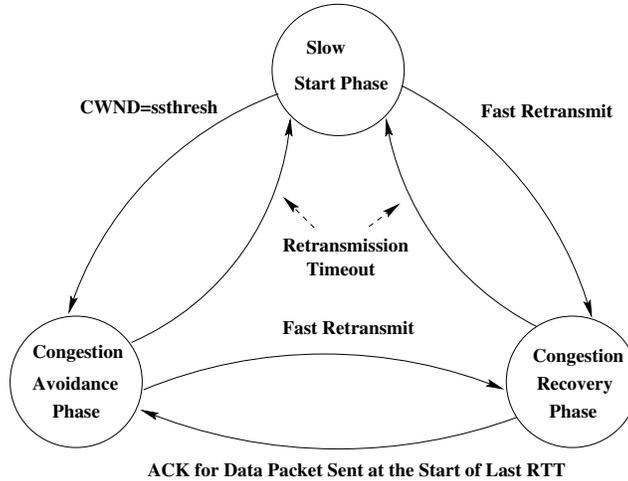


Figure 6: Transitions of a sender's mode in the TCP congestion control scheme

recovery mode irrespective of the mode it was in before. The receipt of a whole ACK moves the sender from the congestion-recovery mode to the congestion-avoidance mode. Figure 6 shows the transitions among these three modes.

4 Performance Evaluation of RR

We evaluated the RR algorithm using the *ns-2* [30]. Since New-Reno is known to perform much better than Reno in the presence of multiple packet losses, we focused on the performance comparison among RR, Tahoe, New-Reno, and SACK TCP. The performance evaluation is based on effective throughput, which is a commonly-used metric for end-to-end protocols. The effective throughput measures the actual “good” data transmission rate, where “good” data does not include retransmissions, dropped or duplicate data packets.

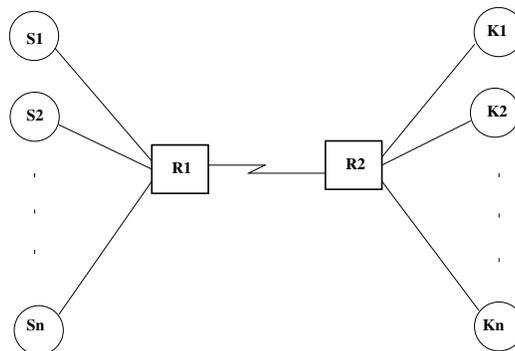


Figure 7: The network topology used for RR experiments

4.1 The Simulation Setup

The simulated network topology is shown in Figure 7, where S_i (K_i) represents a sending (receiving) host, $i = 1, \dots, n$. $R1$ and $R2$ represent two finite-buffer gateways. Different connections from S_i to K_i share the common bottleneck between $R1$ and $R2$. In our simulation experiments each data packet is 1000 bytes long and the size of an ACK packet is 40 bytes. The data traffic in our simulation is generated by FTP. The receiver sends an ACK for every data packet it received. The window size and buffer space at the gateways are measured in number of fixed-size packets, instead of bytes.

4.2 Behavior with Drop-Tail Gateways

The drop-tail gateway with FIFO queueing service has been widely deployed in the Internet, which schedules incoming packets in a FIFO manner and discards incoming packets when the buffer is full. The advantages of the drop-tail gateway are simple, scalable and easy to implement.

Table 3: Simulation configuration

Buffer size	8 packets
TCP maximum window size	30 packets
Bottleneck bandwidth	0.8 Mbps
Bottleneck link delay (1-way)	95 ms
Total delay of side-links (1-way)	5 ms
Side-links bandwidth	10 Mbps

The simulation parameters for RR with drop-tail gateways are summarized in Table 3. There are three TCP connections from S_i to K_i , $i = 1, 2, 3$. Only the first connection is shown in the graphs. The second and the third connections are included to obtain the desired packet loss pattern for the first connection, which only has a limited amount of data to send. Note the buffer size is set to achieve the desired packet loss pattern. If a larger buffer size is set, we can add more background traffic to achieve the same loss pattern. So, the TCP behaviors in each simulation experiment are deterministic, and do not change with different runs as long as the simulation setup and the background traffic remain unchanged.

The simulation results for scenarios with 3 and 6 lost packets within a window of data are plotted in Figure 8, where the effective throughput of the TCP connection during the congestion-recovery period is shown with different TCP recovery schemes.

The RR's effective throughput is significantly higher than that of Tahoe and New-Reno, and slightly higher than that of SACK. Its consistently better performance across 3_drop and 6_drop scenarios indicates RR's resilience to bursty losses within a window of data.

It is interesting to observe that Tahoe is more robust than New-Reno in case of high bursty losses, and achieves a higher effective throughput than New-Reno. This observation has been confirmed by

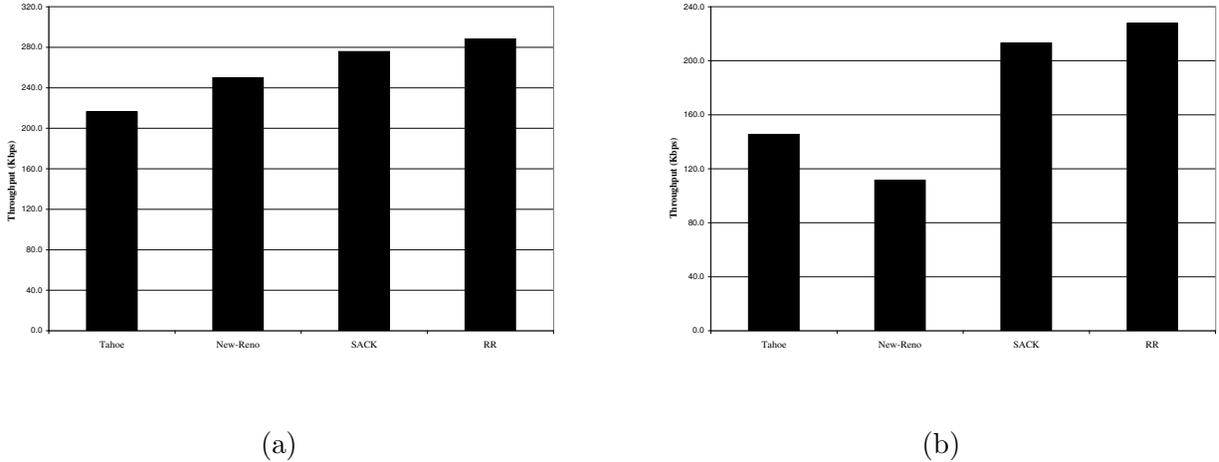


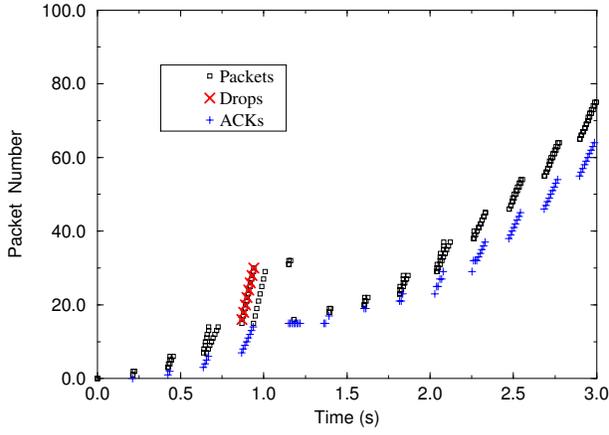
Figure 8: Effective Throughput (a) for a scenario with 3 packet losses; (b) for a scenario with 6 packet losses from a window of data

another recent study [2], which showed that in a wireless channel with correlated packet losses, the performance of New-Reno is very poor, and is worse than the performance of Tahoe even in case of moderately bursty losses.

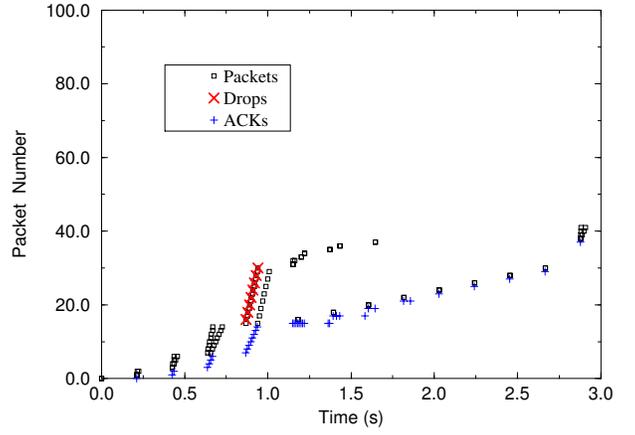
To clearly illustrate the dynamics of the TCP connection and tell the difference of recovery behavior between the current congestion-recovery mechanisms and RR, the standard TCP sequence number plots are used and shown in Figure 9. The graphs were generated by tracing packets entering and departing R1. For each figure, the x-axis represents the packet arrival or departure time in seconds while the y-axis represents the packet number. Packets are numbered starting with 0. Arrival and departure of each packet are marked by squares, and each dropped packet and the receipt of an ACK by “x” and “+,” respectively, in the figures.

As shown in Figure 9, a burst of packet losses within a window occurs between 0.86s and 0.95s. Eight packets are dropped when *cwnd* reaches 16 during the slow-start phase, i.e., a half of the window of packets are lost. The other half of the window of packets have arrived, and are then queued at the receiver, and each of them triggers the same ACK for packet 15, i.e., there are eight duplicate ACKs flowing back to TCP sender. The first ACK for packet 15 is received at time 1.15s. Since the TCP sender is still in the slow-start phase and unaware of the occurrence of packet losses, this ACK triggers two new packets out as each arrival of ACK increases *cwnd* by 1 in the slow-start phase. These two new packet transmissions can be easily identified in Figures 9 (a) and (c) at time 1.15s.

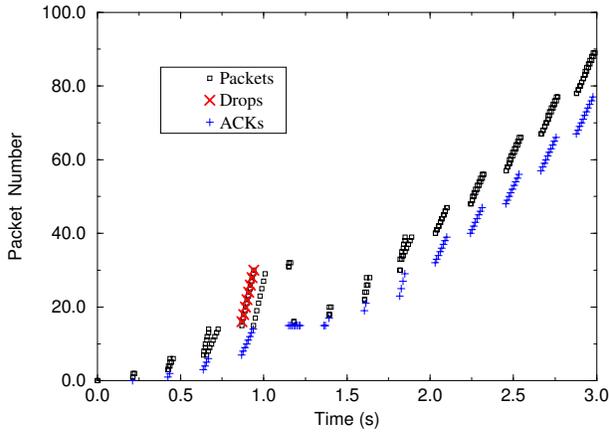
However, upon receipt of the subsequent three duplicate ACKs for packet 15, the TCP sender concludes that packet 16 has been dropped, and enters the congestion-recovery phase. As shown in Figure 9 (b) and (d), the rest of four duplicate ACKs for packet 15 trigger the transmission of two new data packets according to the “one new packet per two duplicate ACKs” policy in New-Reno and the RR’s retreat sub-phase. So, four new data packets are sent out during the first RTT of the



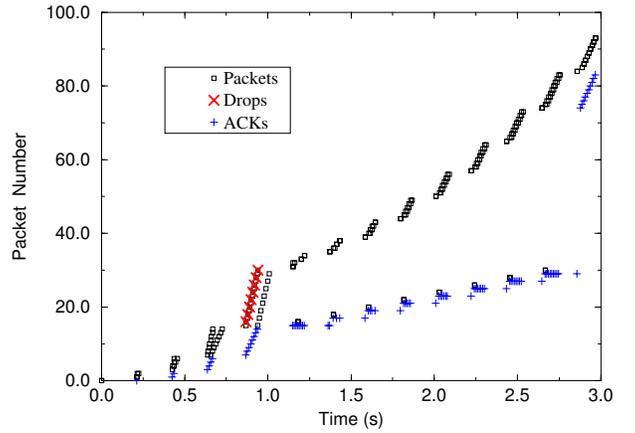
(a). Tahoe



(b). New-Reno



(c). SACK



(d). Robust Recovery

Figure 9: Traffic dynamics with bursty packet losses within a window of data packets

congestion-recovery phase in New-Reno and RR.

In the subsequent RTTs of congestion recovery, Figure 9 (b) clearly show the exponential reduction of data transmission in New-Reno, which significantly lowers link utilization even if it does not cause a coarse timeout. In contrast, RR linearly increases the amount of transmitted data per RTT while recovering the dropped packet. Since no further data loss occurs during the congestion-recovery phase, RR exits the congestion-recovery phase upon receipt of a new ACK at time 2.87s just like what New-Reno does, and seamlessly switches to the congestion-avoidance phase by setting $cwnd$ to the current value of $actnum$, 10.

4.3 Behavior with RED gateways

The drop-tail gateway has been shown to arbitrarily distribute packet losses among TCP connections, leading to global synchronization [43], and biasing against bursty connections. The Random Early Detection (RED) [14] gateway has been proposed to alleviate the problems of the drop-tail gateway. The RED gateway detects incipient congestion based on the computation of the average queue size, and randomly drops or marks incoming packets before its buffer is exhausted.

Although RED reduces the frequency of global synchronization and packet-loss rates and minimizes the bias against bursty connections, its performance strongly depends on the behavior of aggregate flows. It is in general difficult to configure a RED gateway into its ideal region as the aggregate flows change with time. Furthermore, RED does not guarantee avoidance of multiple packet losses in a window of data, especially under heavy network congestion. Therefore, given a widely-deployed active queue management mechanisms such as RED, robust TCP congestion recovery is still very important to TCP performance.

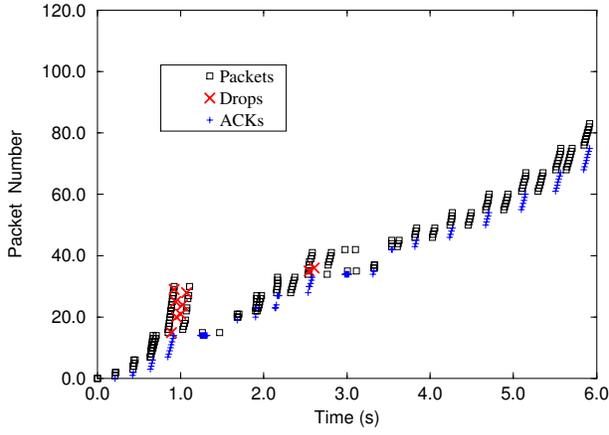
Table 4: RED parameter configuration

Minimum threshold	5 packets
Maximum threshold	20 packets
Maximum drop probability	0.02
Weight queue	0.002

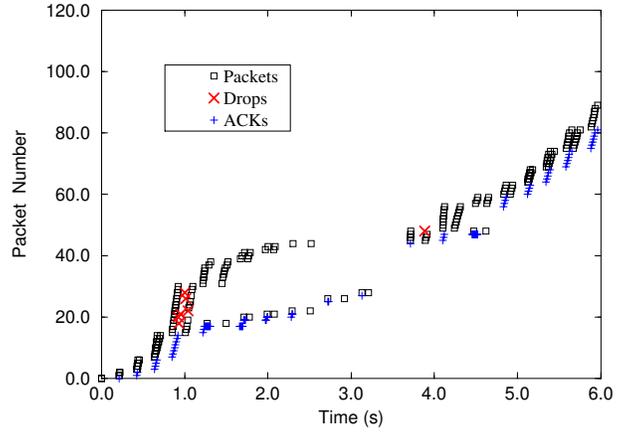
The simulation setup for the RED gateway experiments is similar to the previous one, except that RED gateways replace drop-tail gateways and the buffer size is set to 25 packets. The configuration of the RED parameters is summarized in Table 4. Instead of only 3 TCP flows in the drop-tail experiments, 10 TCP flows share the common bottleneck of 0.8 Mbps. The first five TCP flows start at time 0. Then, a new TCP flow starts every 0.5 second. The last one starts at time 2.5s. The duration of the simulation is 6 seconds. All TCP flows have an infinite amount of data to send during the simulation. The purpose of configuring this simulation environment is to generate heavy congestion at the RED gateway. Due to the random drops at gateways, the TCP behaviors in the RED experiments are no longer deterministic. However, except for the case of a random retransmission loss that rarely occurs, randomness does not effect the TCP behaviors in recovering from bursty losses.

In each simulation experiment, all TCP flows use the same congestion-recovery mechanism. Since they experience a similar recovery behavior, only the first one is shown in the graphs. Figure 10 shows the dynamics of the first TCP flow for different TCP congestion-recovery mechanisms. As with the drop-tail gateway, RR achieves the highest TCP effective throughput when the RED gateway is deployed. As expected, RR's effective throughput is significantly higher than that of Tahoe and New-Reno, and is clearly higher than that of SACK.

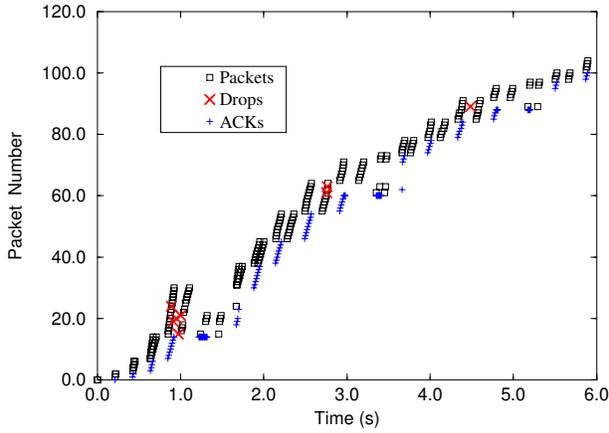
As shown in Figure 10, bursty packet losses occur after *cwnd* reaches 16. RR recovers the dropped



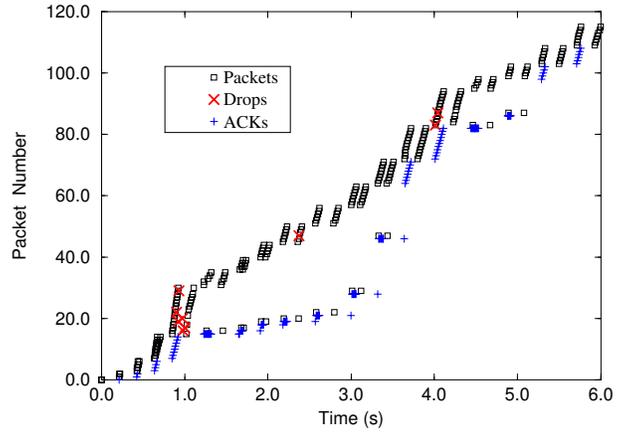
(a). Tahoe



(b). New-Reno



(c). SACK



(d). Robust Recovery

Figure 10: Simulation results for a scenario with RED gateways

packets during the next RTT while transmitting new data packets. During the recovery period, a further packet loss occurs at time 2.37s and the lost packet is then retransmitted around time 3.33s. Upon receipt of the new ACK for packet 64 at time 3.65s, the TCP sender leaves the congestion-recovery phase and enters congestion-avoidance phase. However, two packet losses within a window around time 4.0s make the TCP sender enter the congestion-recovery phase again. It will then switch back to congestion-avoidance when the new ACK for packet 98 is received at time 5.30s.

Figure 10 (b) clearly shows that in New-Reno, the exponential reduction in the new data transmission stops the flow of returning ACKs and stalls the transmission of new data packets. Before the receipt of the new ACK that takes the TCP sender out of the recovery phase, only a retransmissions and a new partial ACK flow around the pipe, which significantly degrades the link utilization. Although Tahoe is robust to bursty packet losses, entering slow-start to recover from further data losses, which happens at time 2.65s when two data packets within a window of data are dropped, substantially degrades the

TCP performance as shown in Figure 10 (a).

4.4 Discussion

In our simulation experiments, we assume that end hosts have unlimited buffer space and the sender always has data to send, i.e., a persistent source. This is not true in reality. If the sender does not have enough data to send or the free buffer space at end hosts is exhausted, no new data packets but retransmissions would be sent out. In this case, RR will behave like New-Reno as shown in Figure 10(b): retransmits a dropped packet, then stalls until it receives a new partial ACK; upon receipt of a new partial ACK, retransmits, and this cycle repeats until enough free buffer space becomes available at end hosts or until the end of transmission.

In most cases the buffer requirement at end hosts during loss recovery can be met by today's end hosts. Even if most packets in a large window of data are lost, like the scenario shown in Figure 9, the number of outstanding packets at an end host is still less than 60. If each data packet size is 1 MSS (maximum segment size) of 1K bytes each, the requirement of buffer space at an end host is at most 60 Kbytes, which is easy to meet with current technology. Moreover, the automatic TCP Buffer Tuning [37] alleviates the potential degradation caused by limited buffer space.

To recover from packet losses with non-persistent sources, two solutions have been proposed in [25]. One solution suggests that the sender should re-send the packets starting from the first unacknowledged packet. The other solution is to send header-only packets to keep the flywheel going.

5 Fitting the Square-Root Model

To analytically characterize the throughput of a TCP connection in steady-state as a function of packet-loss rate and RTT, a model has been proposed to describe the macroscopic behavior of the TCP congestion-avoidance algorithm [29]. It assumes that no retransmission timeouts exist with a persistent source and a sufficient receiver window. Under this assumption, the model gives an upper bound on the bandwidth of a TCP connection that can be achieved for a given random packet-loss rate. The estimation of an ideal achievable throughput is:

$$BW = \frac{MSS}{RTT} \frac{C}{\sqrt{p}}$$

where p is the packet-loss rate, and C is a constant that lumps several factors into one term, including the ACK strategy.

It has been shown in [8, 29] that TCP SACK is much closer to the ideal congestion behavior than TCP Reno. Therefore, in this paper we only compare the fitness of RR with that of TCP SACK against the model. The simulation environment is the same as that in Section 4, except that the simulation length is 100 seconds. Only one TCP connection is active during the simulation, and its start-up phase is ignored. Artificial losses are introduced at the gateway *R1*. The uniform random packet-loss rate is

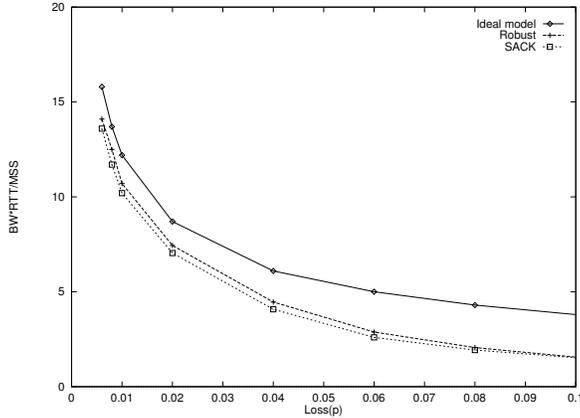


Figure 11: Comparison of fitness to the model

varied in each experiment, while the MSS and RTT are fixed. (MSS is set to 1000bytes and RTT is set to 200ms.) Since the receiver sends an ACK for every data packet received, C is set to $\sqrt{\frac{3}{2}}$.

The simulation results are plotted in Figure 11, where the x-axis represents the random packet-loss rate and the y-axis represents $BW * RTT/MSS$ indicating the window size of the TCP connection. From Figure 11, one can see that RR achieves the same level of fitness to the model as SACK does, and its window size is even slightly closer to the upper bound provided by the model.

In Figure 11, with the increase of packet-loss rate, the upper bound of window size decreases, and the behavior of both RR and SACK does not fit the model very well. This deviation is due to the timeouts triggered for the following two reasons.

- The possibility of retransmission loss increases with the increase of packet-loss rate, and the retransmission loss will lead to a retransmission timeout.
- The small window size reduces the effectiveness of fast retransmit, and hence recovery of some random losses has to resort to retransmission timeouts if three duplicate ACKs are not received.

To predict bandwidth more accurately, more sophisticated models[1, 33] have been proposed, which captures not only the behavior of fast retransmit but also the effect of retransmission timeouts upon throughput, and validated their models over a wider range of loss rates.

6 Fairness

It is clear that RR improves the effectiveness of TCP. Moreover, due to RR's fitness to the ideal TCP congestion model, the stability of TCP is enhanced and the fluctuation of traffic load on the network is reduced. One remaining concern is the fairness of RR.

Owing to its exponential back-off in the retreat sub-phase and linear adjustments in the probe sub-phase, RR strictly follows the AIMD rule [9] and is TCP-friendly [13]. It converges to the optimal point if competing TCP connections have same RTTs as shown in Figure 12, where bandwidth allocation

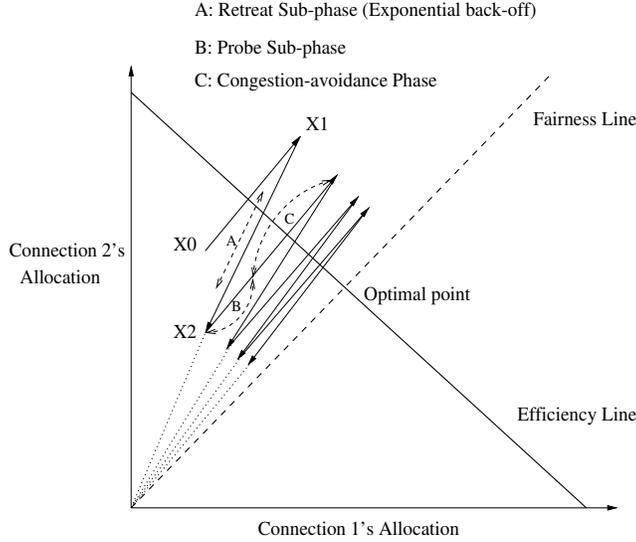


Figure 12: RR converges to the optimal point

to every pair of connections $\{x_1(t), x_2(t)\}$ is represented as a point (x_1, x_2) in a 2-dimensional space. In the figure, the horizontal axis represents bandwidth allocation to connection 1, and the vertical axis represents bandwidth allocation to connection 2. The linear adjustments in the probe sub-phase seamlessly link the congestion-recovery phase with the congestion-avoidance phase, which makes a better use of the bandwidth under-utilized by TCP New-Reno or TCP Reno.

However, to be an incrementally deployable TCP enhancement, RR must interoperate well, in terms of fairness, with existing TCP congestion-recovery strategies, especially TCP Reno. Given less timeouts and linear increase during the probe sub-phase, an obvious concern is that potential unfairness may result from RR when RR competes with other TCP connections whose implementation are TCP Reno.

The simulation environment used for testing fairness is similar to that of testing RR behavior with drop-tail gateways in Section 4.2, except that the buffer size at routers is set to 25 packets. The common bottleneck of 0.8 Mbps is shared by 20 TCP connections. All but one TCP connection have an infinite amount of data to send during the simulation, which are used as the background traffic load, and their starts are staggered. The first TCP connection starts at time 0, then a new TCP connection starts every 0.5 second. All background TCP connections have the same TCP implementation in each experiment. The targeted TCP connection, which has a 100 KBytes file to send from S_{20} to K_{20} , starts at 4.8 seconds. The transfer delay and packet loss rate of the targeted TCP connection are measured.

According to the different TCP implementations used by background TCP connections and the targeted TCP connection, the simulation experiments are categorized into four cases. The simulation results of these four cases are summarized in Table 5.

The simulation results show that when the targeted TCP Reno connection competes with other TCP connections that are implemented with RR, its transfer delay and packet loss rate are lower than those of competing with homogeneous connections that are also implemented with TCP Reno. The

Table 5: Performance of the targeted TCP connection

Case	Targeted	Background	Transfer Delay	Packet Loss Rate
I	Reno	Renos	33.3s	15%
II	Reno	RRs	30.1s	12%
III	RR	RRs	25.6s	13%
IV	RR	Renos	18.0s	11%

performance of the targeted TCP-Reno connection is improved when the background traffic is changed from Reno to RR, instead of being degraded. The reduced global synchronization and traffic fluctuation are the keys to the TCP improvement in Case 2 and 3, which is caused by the background traffic of RR. It demonstrates that RR is TCP-friendly and does not bias against a less aggressive TCP connection. Its fitness to the square-root model shown in Section 5 also confirms its TCP-friendliness [13].

As shown in Table 5, when a single RR competes with heterogeneous connections that are implemented with TCP Reno, it achieves shorter transfer delay and less packet loss rate. However, the high bandwidth consumption of RR does not result from taking the bandwidth away from TCP-Reno’s connections. In the homogeneous scenario, the consumed bandwidth of each TCP Reno connection is only 24 Kbps and the total bandwidth consumption by 20 TCP-Reno connections is 480 Kbps. The unused bandwidth at the bottleneck is 320 Kbps. In Case 4, the achieved bandwidth of RR is 44 Kbps, which is slightly higher than the fair share of 40 Kbps, but there is still a 300 Kbps unused bandwidth. Therefore, RR simply makes a better use of the bandwidth that is under-utilized by TCP Reno.

Besides TCP implementation, two other factors influence the fairness of TCP: round-trip delay and routers’ mechanism. The TCP has been observed to have a bias against the connections with longer delays [12, 24], since TCP does not provide a uniform increase in the rates of connections with different RTTs. The congestion window size is increased by one per RTT in the congestion-avoidance phase or is doubled per RTT in the slow-start phase, and hence, the connections with longer delays open their windows more slowly than those with shorter delays. To improve TCP fairness, the Constant-Rate (CR) policy [12] and the increase-by-K (IBK) [16] policy have been proposed.

In routers, where flows converge in and diverge out, one can find enough information to control sharing and fair allocation of the resources. Router mechanisms have been recognized as playing an important, or even necessary, role in the Internet congestion control [13]. As AVQ, Blue, CSFQ, FRED and REM [3, 11, 23, 26, 39] propose, a router will be able to make fair bandwidth allocations to TCP-friendly connections and identify misbehaving connections without significant overhead.

7 Conclusions

This paper proposes and evaluates a new congestion-recovery mechanism, called Robust Recovery (RR), to improve the performance of TCP flows in the presence of bursty packet losses. RR is evaluated by

simulation, and found to be able to recover from multiple packet losses in a window of data without any significant performance degradation. The key feature of RR is a *variable-length* recovery period, during which a sender recovers from packet losses, while also sending new packets to probe the new equilibrium of the TCP connection.

Also, RR makes the TCP behavior during congestion recovery very close to that during congestion avoidance, thereby extending the performance model for TCP congestion avoidance to represent that for TCP congestion recovery as well. This allows for accurate prediction of the TCP-consumed bandwidth even if selective acknowledgments are not used. In addition to the performance advantages offered by RR, it offers an implementation advantage. In particular, RR is much simpler than SACK TCP, and does not require selective acknowledgments or receiver modification. Moreover, RR is shown to interoperate well, in terms of fairness, with existing TCP congestion-recovery strategies, making it possible to deploy RR incrementally in the Internet.

References

- [1] E. Altman, K. Avrachenkov, and C. Barakat, "A Stochastic Model of TCP/IP with Stationary Random Losses", *Proceedings of ACM SIGCOMM'2000*, Stockholm, Sweden, August 2000.
- [2] F. Anjum and L. Tassiulas, "On the Behavior of Different TCP Algorithms over a Wireless Channel with Correlated Packet Losses", *Proceedings of ACM SIGMETRICS'99*, Atlanta, GA, May 1999.
- [3] S. Athuraliya, V. H. Li, S. H. Low, and Q. Yin, "REM: Active Queue Management" *IEEE Network*, Vol. 15, No. 3, May/June 2001.
- [4] H. Balakrishnan, V. Padmanabhan, and R. H. Katz, "The Effects of Asymmetry on TCP Performance" *Proceedings of ACM/IEEE MOBICOM'97*, Budapest, Hungary, September 1997.
- [5] H. Balakrishnan, V. Padmanabhan, S. Seshan, M. Stemm, and R. Katz, "TCP Behavior of a Busy Internet Server: Analysis and Improvement", *Proceedings of IEEE INFOCOM'98*, San Francisco, CA, March 1998.
- [6] J. Bolliger, U. Hengartner, and T. Gross, "The Effectiveness of End-to-End Congestion Control Mechanisms" *Technical Report No. 313*, Dept. Computer Science, ETH Zürich, February 1999.
- [7] L. Brakmo, S. O'Malley, and L. Peterson, "TCP Vegas: New Technique for Congestion Detection and Avoidance", *Proceedings of ACM SIGCOMM'94*, London, UK, August 1994.
- [8] R. Bruyeron, B. Hemon, L. Zhang, "Experimentations with TCP Selective Acknowledgment", *ACM Computer Communication Review*, Vol. 28, No. 2, April 1998.
- [9] D. Chiu and R. Jain, "Analysis of the Increase and Decrease Algorithm for Congestion Avoidance in Computer Networks", *Computer Networks and ISDN Systems*, 17:1-14, 1989.
- [10] K. Fall and S. Floyd, "Simulation-based Comparisons of Tahoe, Reno, and SACK TCP", *ACM Computer Communication Review*, Vol. 26, No. 3, July 1996.
- [11] W. Feng, K. G. Shin, D. D. Kandlur, and D. Saha, "The BLUE active queue management algorithms", *IEEE/ACM Transactions on Networking*, Vol. 10, No. 4, August 2002.

- [12] S. Floyd, "Connections with Multiple Congested Gateways in Packet-Switched Networks, Part 1: One-Way Traffic" *ACM Computer Communication Review*, Vol. 21, No. 5, October 1991.
- [13] S. Floyd, and K. Fall, "Promoting the Use of End-to-End Congestion Control in the Internet", *IEEE/ACM Transactions on Networking*, Vol. 7, No. 4, August 1999.
- [14] S. Floyd and V. Jacobson, "Random Early Detection gateways for Congestion Avoidance", *IEEE/ACM Transactions on Networking*, Vol. 1, No. 4, August 1993.
- [15] S. Floyd, M. Handley, J. Padhye, and J. Widmer, "Equation-based Congestion Control for Unicast Applications", *Proceedings of ACM SIGCOMM'2000*, Stockholm, Sweden, August 2000.
- [16] T. Henderson, E. Sahouria, S. McCanne, and R. Katz, "Improving Fairness of TCP Congestion Avoidance", *Proceedings of IEEE GLOBECOM'98*, Sydney, Australia, Nov. 1998.
- [17] U. Hengartner, J. Bolliger, and T. Gross, "TCP Vegas Revisited", *Proceedings of IEEE INFO-COMM'2000*, Tel-Aviv, Israel, March 2000.
- [18] J. Hoe, "Improving the Start-up Behavior of a Congestion Control Scheme for TCP", *Proceedings of ACM SIGCOMM'96*, Stanford, CA, August 1996.
- [19] V. Jacobson, "Congestion Avoidance and Control", *Proceedings of ACM SIGCOMM'88*, Stanford, CA, pp. 314-329, August 1988.
- [20] V. Jacobson, fast retransmit. Message to the end2end-ietf mailing list, April 1990.
- [21] V. Jacobson, "Berkeley TCP Evolution from 4.3-Tahoe to 4.3-Reno", *Proceedings of the Eighteenth Internet Engineering Task Force*, 1990.
- [22] D. Katabi, M. Handley, and C. Rohrs, "Congestion Control for High Bandwidth-delay Product Networks", *Proceedings of ACM SIGCOMM'2002*, Pittsburgh, PA, August 2002.
- [23] S. Kunniyur and R. Srikant, "Analysis and Design of an Adaptive Virtual Queue", *Proceedings of ACM SIGCOMM'2001*, San Diego, CA, August 2001.
- [24] T. Lakshman and U. Madhow, "The performance of TCP/IP for Networks with High Bandwidth-Delay Products and Random Loss", *IEEE/ACM Transactions on Networking*, Vol.5, No. 3, June 1997.
- [25] D. Lin and H. T. Kung, "TCP Fast Recovery Strategies: Analysis and Improvements", *Proceedings of IEEE INFOCOM'98*, San Francisco, CA, March 1998.
- [26] D. Lin and R. Morris, "Dynamics of Random Early Detection" *Proceedings of ACM SIGCOMM'97*, Cannes, France, September 1997.
- [27] M. Mathis and J. Mahdavi, "Forward Acknowledgment (FACK): Refining TCP Congestion Control", *Proceedings of ACM SIGCOMM'96*, Stanford, CA, August 1996.
- [28] M. Mathis, J. Mahdavi, S. Floyd, and A. Romanow, "TCP Selective Acknowledgment Option", Internet Draft, work in progress, May 1996.
- [29] M. Mathis, J. Semke, J. Mahdavi and T. Ott, "The Macroscopic Behavior of the TCP Congestion Avoidance Algorithm", *ACM Computer Communication Review*, Vol. 27, No. 3, July 1997.
- [30] S. McCanne and S. Floyd, ns-2 Network Simulator. Obtain via: <http://www.isi.edu/nsnam/ns>.

- [31] S. McCreary and K. Claffy, “Trends in Wide Area IP Traffic Patterns — A View from Ames Internet Exchange”, *Proceedings of ITC’2000*, Monterey, CA, September 2000.
- [32] J. Padhye and S. Floyd, “On Inferring TCP Behavior” *Proceedings of ACM SIGCOMM’2001*, San Diego, CA, August 2001.
- [33] J. Padhye, V. Firoiu, D. Towsley, and J. Kurose, “Modeling TCP Throughput: A Simple Model and its Empirical Validation”, *Proceedings of ACM SIGCOMM’98*, Vancouver, Canada, pp. 303-314, September 1998.
- [34] C. Parsa and J.J. Garcia-Luna-Aceves, “Improving TCP Congestion Control Over Internets with Heterogeneous Transmission Media” *Proceedings of IEEE ICNP’99*, Toronto, Canada, November 1999.
- [35] V. Paxson, “End-to-End Internet Packet Dynamics”, *IEEE/ACM Transactions on Networking*, Vol.7, No.3, June 1999.
- [36] J. Postel, Transmission Control Protocol, Request for Comments 793, DDN Network Information Center, SRI International, September 1981.
- [37] J. Semke, J. Mahdavi, and M. Mathis, “Automatic TCP Buffer Tuning”, *Proceedings of ACM SIGCOMM’98*, Vancouver, Canada, September 1998.
- [38] W. Stevens, *TCP/IP Illustrated*, Volume 1. Addison-Wesley Publishing Company, 1994.
- [39] I. Stoica, S. Shenker and H. Zhang, “Core-Stateless Fair Queueing: Achieving Approximately Fair Bandwidth Allocations in High Speed Networks”, *Proceedings of SIGCOMM’98*, Vancouver, Canada, September 1998.
- [40] K. Thompson, G. J. Miller, and R. Wilder, ”Wide-Area Internet Traffic Patterns and Characteristics”, *IEEE Network*, Vol. 11, No. 6, November/December 1997.
- [41] H. Wang and K. G. Shin, “Layer-4 Service Differentiation and Resource Isolation”, *Proceedings of IEEE RTAS’2002*, San Jose, CA, September 2002.
- [42] H. Wang, H. Xin, D. Reeves, and K. G. Shin, “ A Simple Refinement of Slow-Start of TCP Congestion Control”, *Proceedings of IEEE Symposium on Computers and Communications’2000*, Antibes, France, July 2000.
- [43] L. Zhang, S. Shenker, and D. Clark, “Observations on the Dynamics of a Congestion Control Algorithm: The Effects of Two Way Traffic”, *Proceedings of ACM SIGCOMM’91*, Zürich, Switzerland, September 1991.