

F-RTO: An Enhanced Recovery Algorithm for TCP Retransmission Timeouts

Pasi Sarolahti^{*}
Nokia Research Center
pasi.sarolahti@nokia.com

Markku Kojo, Kimmo Raatikainen
University of Helsinki
Department of Computer Science
{markku.kojo,kimmo.raatikainen}@cs.helsinki.fi

ABSTRACT

Spurious TCP retransmission timeouts (RTOs) have been reported to be a problem on network paths involving links that are prone to sudden delays due to various reasons. Especially many wireless network technologies contain such links. Spurious retransmission timeouts often cause unnecessary retransmission of several segments, which is harmful for TCP performance. Recent proposals for avoiding unnecessary retransmissions after a spurious RTO require use of TCP options which must be implemented and enabled at both ends of the connection. We introduce a new TCP sender algorithm for recovery after a retransmission timeout and show that unnecessary retransmissions can be avoided without TCP options. The algorithm effectively avoids the unnecessary retransmissions after a spurious retransmission timeout, improving the TCP performance considerably. The algorithm is friendly towards other TCP connections, because it follows the congestion control principles and injects packets to the network at same rate as a conventional TCP sender. We implemented the algorithm and compared its performance to conventional TCP and Eifel TCP when RTOs occurred either due to sudden delays or due to packet losses. The results show that our algorithm either improves performance or gives similar throughput as the other TCP variants evaluated in different test cases.

1. INTRODUCTION

In the recent years the variety of Internet links with different properties has increased dramatically. The high speed networks have reached Gigabit rates, whereas the increasing number of mobile wireless access networks have introduced a prolific number of mobile hosts attached to the Internet through a slow, wireless links. Moreover, the challenging characteristics of wireless links, in particular high packet loss rate or delays due to various reasons such as link-layer retransmissions or hand-offs between the points of attach-

^{*}This work was initiated while the author was at the University of Helsinki, Department of Computer Science.

ment to the Internet, have introduced a large set of problems for the Internet transport protocols.

The TCP protocol [23] is the dominant Internet transport protocol and its congestion control algorithms [4] are essential for the stability of the Internet. Because these algorithms have a strong effect on TCP performance, finding solutions to improve TCP performance in various challenging conditions has yielded a large number of studies. For example, a taxonomy of the different solutions for improving TCP performance over slow wireless links can be found in [21]. The traditional problem regarding the use of TCP over wireless links or other challenging channels has concerned TCP congestion control. If a packet is lost, TCP interprets it as an indication of congestion and a TCP sender needs to reduce its transmission rate. Hence, TCP performance deteriorates with increasing packet loss rate. If the packet loss occurred due to corruption, reducing the TCP transmission rate is, however, the wrong action to take.

TCP uses the *fast retransmit* mechanism [4] to trigger retransmissions after receiving three successive duplicate acknowledgements (ACKs). If for a certain time period TCP sender does not receive ACKs that acknowledge new data, the TCP retransmission timer expires as a backoff mechanism. When the retransmission timer expires, the TCP sender retransmits the first unacknowledged segment assuming it was lost in the network. Because a retransmission timeout (RTO) can be an indication of severe congestion in the network, the TCP sender resets its congestion window to one segment and starts increasing it according to the slow start algorithm. However, if the RTO occurs spuriously and there still are segments outstanding in the network, a false slow start is harmful for the potentially congested network as it injects extra segments to the network at increasing rate.

Since wireless networks are often subject to high packet loss rate due to corruption or hand-offs, reliable link-layer protocols are widely employed with wireless links. The link-layer receiver often aims to deliver the packets to the upper protocol layers in order, which implies that the later arriving packets are blocked until the head of the queue arrives successfully. Due to the strict link-layer ordering, the communication end points observe a pause in packet delivery that can cause a spurious TCP RTO instead of getting out-of-order packets that could result in a false fast retransmit instead. Either way, interaction between TCP retransmission mechanisms and link-layer recovery can cause poor performance.

Wireless links may also suffer from link outages that cause persistent data loss for a period of time. If the link outage lasts long enough, it triggers the TCP RTO at the sender which then retransmits the unacknowledged TCP segments. However, if the link layer protocol is highly persistent in its retransmissions, it is able to deliver the original packets to the TCP receiver once the link outage is finished. In this case the TCP RTO may also be triggered spuriously. Other potential reasons for sudden delays that have been reported to trigger spurious RTOs include a delay due to tedious actions required to complete a hand-off or re-routing of packets to the new serving access point after the hand-off, arrival of competing traffic on a shared link with low bandwidth, and a sudden bandwidth degradation due to reduced resources on a wireless channel [13, 15].

In this paper we focus on attacking the TCP performance problems resulting from unnecessary retransmissions that originate from spurious RTOs. The possible solutions for alleviating poor performance can roughly be divided in two categories. One alternative is to avoid the RTOs in the first place by changing the algorithm used for the RTO calculation. Different constants and granularities applied to the standard algorithm documented in [22] have been studied [3]. In addition, totally new algorithms for setting the RTO timer have been suggested (e.g. [19]). However, we believe it is very difficult to come up with an algorithm that results in a good performance in various different network environments.

Another way to mitigate the performance penalty due to spurious retransmission timeouts is to change the TCP sender behavior after a timeout. In particular, there are two symptoms that typically follow the spurious retransmission timeout when conventional TCP is used:

- The TCP sender often transmits several segments unnecessarily, because after the RTO and first retransmission the cumulative acknowledgements for the original transmissions appear at the TCP sender one at a time, triggering further unnecessary retransmissions. Often a full TCP window is retransmitted quite unnecessarily. In particular, this is very costly for slow links.
- The TCP sender unnecessarily reduces TCP congestion window to one segment. This, in turn, is costly for high-bandwidth links as it takes a long time for the sender to reopen the window.

In this paper we present a new algorithm, called *Forward RTO-Recovery (F-RTO)*, for a TCP sender to recover after a retransmission timeout. Although the main motivation of the algorithm is to recover efficiently from a spurious RTO, we require it to achieve similar performance with the conventional RTO recovery in other situations where RTO may occur. Our approach requires modification only at the TCP sender, while adhering to the TCP congestion control principles. Moreover, the F-RTO recovery algorithm does not require use of any TCP options or additional bits in the TCP header, unlike other recent suggestions for avoiding unnecessary retransmissions, for example Eifel TCP [18]. The

F-RTO algorithm uses a set of simple rules for avoiding unnecessary retransmissions after a spurious RTO. When the first acknowledgements arrive after retransmitting the segment for which the RTO expired, the F-RTO sender does not immediately continue with retransmissions like the conventional RTO recovery does, but it first checks if the acknowledgements advance the window to determine whether it needs to retransmit, or whether it can continue sending new data. F-RTO can be considered somewhat similar to the *Limited Transmit* algorithm [2], but applied to the RTO recovery.

We implemented the F-RTO algorithm in Linux OS and compared its performance to conventional TCP and Eifel TCP in different scenarios where RTOs occurred either due to sudden delays or due to packet losses. The results show that when RTOs are spurious, the F-RTO algorithm significantly improves TCP performance compared to conventional RTO recovery and performs slightly better than Eifel TCP. When the RTOs are due to packet losses, F-RTO yields similar throughput as conventional TCP while Eifel TCP has performance problems.

The rest of the paper is organized as follows. In Section 2 we describe the problem caused by unnecessary retransmission timeouts in greater detail and discuss the related work with suggested solutions to the problem. In Section 3 we describe the F-RTO algorithm for making forward transmissions after RTO. We continue by giving some examples of the F-RTO algorithm behavior in different situations involving RTOs in Section 4. In Section 5 we describe the experiments made with F-RTO in different network environments and the results of the experiments. Finally, we give some thoughts on the future work on F-RTO in Section 6 and conclude our work in Section 7.

2. RELATED WORK ON SPURIOUS RETRANSMISSION TIMEOUTS

Most of the retransmissions in current TCP implementations are expected to be triggered by duplicate ACKs. A TCP retransmission timeout is considered as a fallback mechanism for the cases when retransmissions cannot be triggered by duplicate ACKs. More specifically, a RTO-triggered retransmission is needed when a retransmission is lost, or when nearly a whole window of data is lost, thus making it impossible for the receiver to generate enough duplicate ACKs for triggering TCP fast retransmit. Under these assumptions, retransmitting the unacknowledged segments immediately after the RTO is likely to be the most efficient way of recovering.

With the conventional RTO recovery the TCP sender retransmits the first unacknowledged segment, sets the congestion window to one segment and the *slow start threshold (ssthresh)* to half of the number of currently outstanding segments, when the RTO expires. After this the sender continues in slow-start, increasing the congestion window by one segment on each ACK that advances the window and retransmitting the next unacknowledged segments allowed by the congestion window. In effect the sender retransmits an exponentially increasing number of segments on each subsequent round-trip time until the congestion window reaches *ssthresh* and the sender enters congestion avoidance [4]. This

allows efficient recovery of lost segments especially if several segments were lost within the same window.

However, if no segments were lost but the retransmission timer expires spuriously, the segments retransmitted in the slow-start are sent unnecessarily. Particularly, this phenomenon is very possible with the various wireless access network technologies that are prone to sudden delay spikes. Figure 1 shows a time-sequence diagram of a TCP transfer, when a 3-second delay occurs on the link. The retransmission timer expires because of the delay, spuriously triggering the RTO recovery and unnecessary retransmission of all unacknowledged segments. This happens because after the delay the ACKs for the original segments arrive at the sender one at the time but too late, because the TCP sender has already entered the RTO recovery. Therefore, each of the ACKs trigger the retransmission of segments for which the original ACKs will arrive after awhile. This continues until the whole window of segments is eventually unnecessarily retransmitted. Furthermore, because a full window of retransmitted segments arrive unnecessarily at the receiver, it generates duplicate ACKs for these out-of-order segments. Later on, the duplicate ACKs unnecessarily trigger fast retransmit at the sender.




Figure 1: A delay triggers spurious retransmission.

There is no known way to prevent the retransmission timeout from expiring because of a sudden delay. However, by having additional information in the TCP segments, the unnecessary retransmissions following the spurious RTO can be avoided. The *Eifel algorithm* [18] suggests that the TCP sender indicates whether a segment is transmitted for the first time, or whether it is a retransmission. When this information is echoed back in the acknowledgement, the sender can determine whether the original segment arrived at the receiver and declare the retransmission either correct or spurious action. Based on this knowledge, the sender either retransmits the unacknowledged segments in the conventional way, assuming the RTO was triggered by a segment loss, or reverts the recent changes on the congestion control parameters and continue with transmitting new data. The latter alternative is likely to be the correct action to take when the original segment was acknowledged after the RTO, indicat-

ing that the RTO was spurious.

The Eifel algorithm suggests using either the TCP timestamps option [7] or two of the reserved bits in the TCP header for distinguishing the original transmissions from retransmissions. Using the reserved bits in the TCP header requires modification to TCP at both ends. The TCP timestamps option is deployed on some Internet hosts¹, but in order to take advantage of Eifel, timestamps option would need to be deployed at both ends of the TCP connection. Given that the sudden delays are often a problem on wireless links with low bandwidth, including timestamps in each TCP segment increases the TCP header overhead and makes the communication inefficient. Moreover, the TCP timestamps are not supported in the current TCP/IP header compression specifications [14, 9].

Instead of distinguishing the ACKs of the original transmissions from the ACKs of the retransmissions at the TCP sender, the receiver can indicate whether it received a segment that had arrived earlier. The *Duplicate SACK (DSACK)* enhancement [11] suggests to use the first SACK block to indicate duplicate segments arriving at the receiver. This alternative has its benefits over the Eifel algorithm presented above, because the SACK option is being more widely deployed than the TCP timestamps [1], and the SACK blocks are appended to the TCP headers only when necessary. However, if the unnecessary retransmissions occurred due to spurious RTO caused by a sudden delay, the acknowledgements with the DSACK information arrive at the sender only after the acknowledgements of the original segments. Therefore, the unnecessary retransmissions following the spurious RTO cannot be avoided by using DSACK. Instead, the suggested recovery algorithm using DSACK can only revert the congestion control parameters to the state preceding the spurious retransmission [5]. Both ends of the TCP connection need to be aware of the DSACK extension in order to take advantage of it.

3. F-RTO ALGORITHM

The F-RTO algorithm affects the TCP sender behavior only after a retransmission timeout, otherwise the behavior is similar to the conventional TCP. The guideline behind F-RTO is, that an RTO either indicates a loss, or it is caused by an excessive delay in packet delivery while there still are outstanding segments in flight. If the RTO was due to delay, i.e. the RTO was spurious, acknowledgements for non-retransmitted segments sent before the RTO should arrive at the sender after the RTO occurred. If no such segments arrive, the RTO is concluded to be non-spurious and the conventional RTO recovery with go-back-N retransmissions should take place at the TCP sender.

To implement the principle described above, an F-RTO sender acts as follows: if the first ACK arriving after a RTO-triggered retransmission advances the window, transmit two new segments instead of continuing retransmissions. If also the second incoming acknowledgement advances the window, RTO is likely to be spurious, because the second ACK is triggered by an originally transmitted segment that

¹A study on use of the different TCP options indicates that 15 % of the WWW clients connected to a WWW server on the Internet used TCP timestamps in the early 2000s [1].

has not been retransmitted after the RTO. If either one of the two acknowledgements after RTO is a duplicate ACK, the sender continues retransmissions similarly to the conventional RTO recovery algorithm.

When the retransmission timer expires, the F-RTO algorithm takes the following steps at the TCP sender. In the algorithm description below we use `SND.UNA` to indicate the first unacknowledged segment.

1. *When the retransmission timer expires, retransmit the segment that triggered the timeout.* As required by the TCP congestion control specifications, the `ssthresh` is adjusted to half of the number of currently outstanding segments. However, the congestion window is not yet set to one segment, but the sender waits for the next two acknowledgements before deciding on what to do with the congestion window.
2. *When the first acknowledgement after RTO arrives at the sender, the sender chooses the following actions depending on whether the ACK advances the window or whether it is a duplicate ACK.*
 - (a) *If the acknowledgement advances `SND.UNA`, transmit up to two new (previously unsent) segments.* This is the main point in which the F-RTO algorithm differs from the conventional way of recovering from RTO. After transmitting the two new segments, the congestion window size is set to have the same value as `ssthresh`. In effect this reduces the transmission rate of the sender to half of the transmission rate before the RTO. At this point the TCP sender has transmitted a total of three segments after the RTO, similarly to the conventional recovery algorithm. If transmitting two new segments is not possible due to advertised window limitation, or because there is no more data to send, the sender may transmit only one segment. If no new data can be transmitted, the TCP sender follows the conventional RTO recovery algorithm and starts retransmitting the unacknowledged data using slow start.
 - (b) *If the acknowledgement is duplicate ACK, set the congestion window to one segment and proceed with the conventional RTO recovery.* Two new segments are not transmitted in this case, because the conventional RTO recovery algorithm would not transmit anything at this point either. Instead, the F-RTO sender continues with slow start and performs similarly to the conventional TCP sender in retransmitting the unacknowledged segments. Step 3 of the F-RTO algorithm is not entered in this case. A common reason for executing this branch is the loss of a segment, in which case the segments injected by the sender before the RTO may still trigger duplicate ACKs that arrive at the sender after the RTO.
3. *When the second acknowledgement after the RTO arrives, either continue transmitting new data, or start retransmitting with the slow start algorithm, depending on whether new data was acknowledged.*

- (a) *If the acknowledgement advances `SND.UNA`, continue transmitting new data following the congestion avoidance algorithm.* Because the TCP sender has retransmitted only one segment after the RTO, this acknowledgement indicates that an originally transmitted segment has arrived at the receiver. This is regarded as a strong indication of a spurious RTO. However, since the TCP sender cannot surely know at this point whether the segment that triggered the RTO was actually lost, adjusting the congestion control parameters after the RTO is the conservative action. From this point on, the TCP sender continues as in the normal congestion avoidance.

If this algorithm branch is taken, the TCP sender ignores the `send_high` variable that indicates the highest sequence number transmitted so far [10]. The `send_high` variable was proposed as a “bugfix” for avoiding unnecessary multiple fast retransmits when RTO expires during fast recovery with NewReno TCP. As the sender has not retransmitted other segments but the one that triggered RTO, the problem addressed by the *bugfix* cannot occur. Therefore, if there are duplicate ACKs arriving at the sender after the RTO, they are likely to indicate a packet loss, hence fast retransmit should be used to allow efficient recovery. Alternatively, if there are not enough duplicate ACKs arriving at the sender after a packet loss, the retransmission timer expires another time and the sender enters step 1 of this algorithm to detect whether the new RTO is spurious.

- (b) *If the acknowledgement is a duplicate ACK, set congestion window to three segments, continue with the slow start algorithm retransmitting unacknowledged segments.* The duplicate ACK indicates that at least one segment other than the segment that triggered RTO is lost in the last window of data. There is no sufficient evidence that any of the segments was delayed. Therefore, the sender proceeds with retransmissions similarly to the conventional RTO recovery algorithm, with the `send_high` variable stored when the retransmission timer expired to avoid unnecessary fast retransmits.

If either one of the two acknowledgements arriving after the RTO is a duplicate ACK, the algorithm is safe, because it reverts back to the conventional retransmissions and adjusts the congestion window appropriately. However, the validity of the algorithm when the two first acknowledgements advance `SND.UNA` is worth discussing. As described above, this indicates that at least one segment was delayed. If the next segments in the window were also delayed, for example being blocked by the first delayed segment, the algorithm performs as intended, as we will show in Section 4. If the next segments would not have been delayed, they would have arrived before the delayed segment and triggered duplicate ACKs. We will discuss the F-RTO behavior under packet reordering in more detail in Section 4.

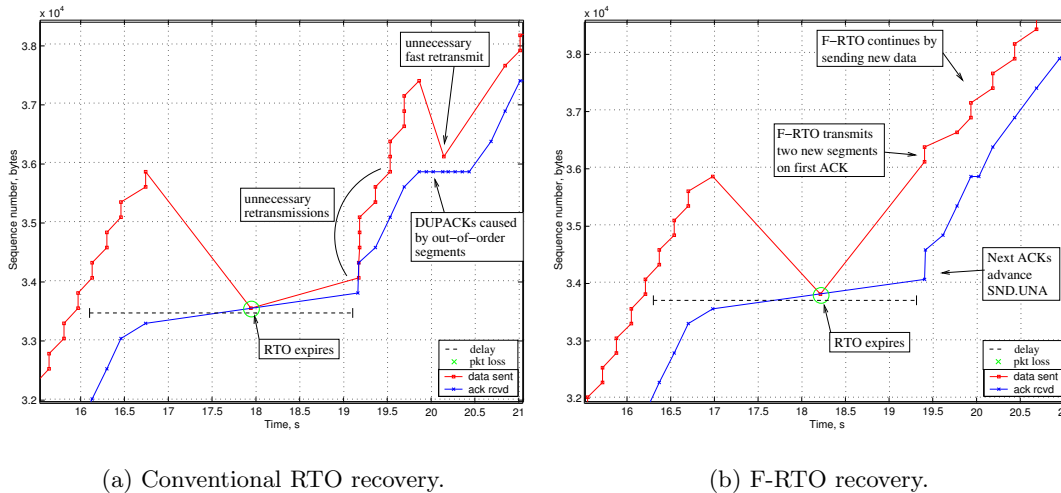


Figure 2: Comparison of the conventional RTO and F-RTO after an excessive delay.

When algorithm branch (3a) is taken, the sender does not reduce the congestion window to one segment, but halves it to the level of `ssthresh`. Because the sender does not enter slow start, it increases congestion window only once in a round-trip time after RTO, and therefore is slightly more conservative than the conventional recovery algorithm. In fact, if the segment that triggered RTO was not lost, the correct behavior would have been to not decrease the congestion window at all. If the DSACK option is in use, the sender can detect whether the retransmission was unnecessary, and revert the last adjustments on the congestion control parameters in such a case. In general, it is possible to separate the detection of a spurious RTO from the actions taken as congestion control response, and employ a different response alternative than what was described above. We will discuss some of the suggested response alternatives that could be applied with F-RTO-based detection in Section 6.

If an ACK acknowledges the whole outstanding window up to the highest transmitted segment at algorithm branch (2a), the TCP sender should not declare the RTO spurious, but follow the conventional TCP behavior. A common case of this is that the RTO was caused due to lost retransmission, and the rest of the window was successfully delivered to the receiver before the RTO occurred. In this case the ACK following the RTO acknowledges all of the outstanding window, and the F-RTO algorithm as described above could end up in algorithm branch (3a) that is meant to be applied in the case of a spurious RTO. With some kind of burst avoidance mechanism and a conservative congestion control after spurious RTOs this exception is not necessarily needed, because then the TCP behavior stays within the congestion control rules, as it was in our case when using the Linux TCP implementation. However, considering the use of F-RTO as a generic detection mechanism for spurious RTOs the additional rule given above is recommended.

Branch (3a) can also be taken when a segment loss is immediately followed by a delay. In this case the retransmission

triggered by the timer was not unnecessarily made. Arrival of new acknowledgements after the RTO indicates that there was only one segment loss in addition to the excessive delay that triggered the RTO. Therefore, we consider that reducing the congestion window to half of its previous size is an adequate action at this point, because a similar action is taken when TCP sender enters fast recovery.

4. DISCUSSION OF F-RTO BEHAVIOR IN SPECIFIC SCENARIOS

In this section we discuss the different reasons that may cause the RTO to expire and study the different scenarios after a RTO has expired due to these reasons. We compare the packet traces produced using the conventional RTO recovery and using F-RTO, and discuss the differences of the two recovery methods. *Selective Acknowledgements (SACK)* [20] and *limited transmit* [2] TCP enhancements are used in the examples presented in this section, since SACK can be considered rather widely deployed today, and limited transmit is a sender side modification that can be implemented with F-RTO to further improve the TCP performance. However, the F-RTO algorithm does not require either of these enhancements to be present.

4.1 Sudden delays

Recovering efficiently from spurious retransmission timeouts is the main motivation of the F-RTO algorithm. Figure 2 compares the packet traces of the conventional RTO recovery and F-RTO. Figure 2(a) shows that the conventional recovery method eventually retransmits the whole window of segments unnecessarily, since the acknowledgements of the originally transmitted segments arrive at the sender after the RTO. When the retransmissions arrive at the receiver, it generates a duplicate ACK for each arriving retransmission, thus causing an unnecessary fast retransmit at the TCP sender.

Figure 2(b) shows that F-RTO avoids the unnecessary re-




Figure 3: Comparison of the conventional RTO and F-RTO after a lost retransmission.

transmissions following the spurious RTO. The first acknowledgement arriving at the sender after the RTO advances `SND.UNA`, and the sender transmits two previously unsent segments. The second ACK arriving after the RTO acknowledges two originally transmitted delayed segments, hence the sender continues transmitting new data. However, since the congestion window was reduced after the RTO, the sender waits for a few acknowledgements without sending new segments to balance the number of packets in flight towards the present congestion window size.

4.2 Lost retransmission

A common reason for triggering TCP RTO is the loss of a retransmitted segment. Once a segment has been retransmitted, it can only be retransmitted again after the RTO expires. Figure 3 compares the packet traces of the conventional RTO recovery with the traces of the F-RTO recovery when a retransmitted segment is lost and it is retransmitted again as triggered by RTO. One can notice that the behavior of the conventional RTO recovery and the F-RTO recovery is similar. In the scenario shown, both variants get to transmit two new segments after the RTO, and then proceed with transmitting new data.

Figure 3(a) shows that when the RTO retransmission arrives at the receiver, it acknowledges the whole window, and the conventional TCP sender can proceed with sending new data in slow start. In the presented case the F-RTO recovery shown in Figure 3(b) differs from the conventional recovery only by not entering slow start after the RTO. Because the next ACK arriving at the sender after the RTO acknowledges all outstanding packets, that is, advances `SND.UNA`, the F-RTO sender transmits new segments using congestion avoidance. Instead of setting the congestion window to one segment, F-RTO decreases it to half of its previous size. As one can see, the practical difference between the recovery alternatives is negligible because the number of outstanding packets was rather small when the first packet loss occurred in the presented scenario.

Using congestion avoidance instead of slow start after the F-RTO recovery does not limit the TCP performance in cases where the number of outstanding segments is larger than in the example above. However, because F-RTO sets the congestion window to half of its previous size when the next acknowledgements advance `SND.UNA`, and on the other hand, because we require using burst avoidance with F-RTO, the conventional RTO recovery algorithm and F-RTO result in similar performance. In our implementation the burst avoidance method decreases the congestion window to allow transmitting at most three segments for the first incoming ACK. If the congestion window size is reduced below the slow start threshold, the sender uses slow start in adjusting the congestion window when the next acknowledgements arrive.

4.3 Burst losses

Because losses of several successive packets can result in a retransmission timeout, it is interesting to compare the F-RTO behavior with the conventional RTO recovery in such a case. Figure 4 compares the packet trace of the conventional recovery after a RTO caused by a window of lost segments with the packet trace of the F-RTO recovery. One can see from Figure 4(a) that the segment retransmitted after the second RTO is successfully acknowledged, after which the TCP sender retransmits the rest of the lost segments in slow start.

Figure 4(b) shows a similar scenario with a F-RTO sender. When the segment retransmitted due to RTO is acknowledged, the F-RTO sender transmits two new segments. Because several other segments were dropped in the last window, the two new segments trigger duplicate ACKs. As given by the F-RTO algorithm, the arrival of the duplicate ACK as the second acknowledgement following the RTO makes the sender retransmit unacknowledged segments in slow start. When the second acknowledgement after the RTO arrives, the sender has a congestion window of three segments, similarly to the conventional RTO recovery after

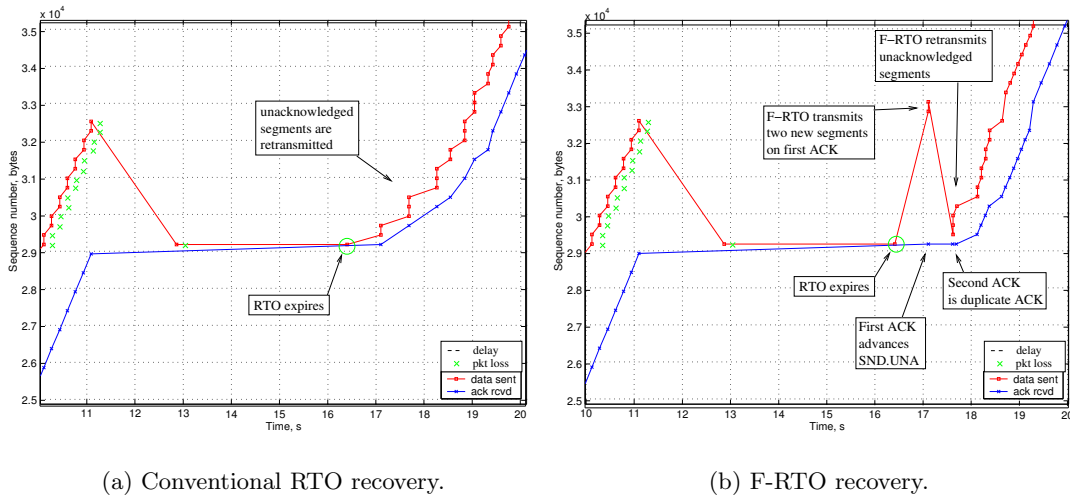


Figure 4: Comparison of the conventional RTO and F-RTO after a burst loss.

two round-trip times. From this point on the congestion window is increased according to the standard TCP congestion control specifications. More generally, if there are any packets lost in the last window of data in addition to the one that triggers RTO, the F-RTO sender enters slow start and retransmits the unacknowledged segments, because the two new segments transmitted after the RTO would trigger duplicate ACKs at the receiver.

F-RTO has a side-effect of triggering an acknowledgement for every incoming retransmission at the TCP receiver, because the receiver is required to send an immediate ACK when it has out-of-order segments in its buffers [4]. However, we believe this detail does not increase the stress on the network significantly, since it only affects the TCP sender's transmission rate during the slow start.

4.4 Packet reordering

Packet reordering is a scenario worth discussing when evaluating the F-RTO behavior, although packet reordering does not usually cause the retransmission timer to expire. A more detailed study on the effect of packet reordering on TCP performance can be found in [5], hence we only discuss here how the F-RTO algorithm relates to packet reordering.

A delayed segment that arrives at the TCP receiver out-of-order appears as a hole in the sequence number space of incoming packets, thus having largely similar effects on the TCP behavior than a dropped packet. Packet reordering may cause fast retransmit, but if there are no retransmission timeouts involved, the F-RTO algorithm does not change the TCP behavior from the conventional recovery. A more interesting scenario arises if the RTO timer expires while packets arrive at the receiver out of order. If the out-of-order segments cause duplicate ACKs to arrive at the sender after the RTO, the F-RTO sender retransmits the unacknowledged packets, similarly to the conventional RTO algorithm. If the delayed packets trigger new acknowledgements that arrive at the sender just after the RTO, the F-

RTO sender proceeds with sending new data. This is likely to be the correct action, because the acknowledgements were triggered by a segment transmitted before the retransmission timeout.

5. PERFORMANCE ANALYSIS

In order to validate the discussion in Section 4, we made experiments in networks with characteristics similar to those that could be expected when communicating over a bottleneck wireless link to a fixed server in a nearby network. This is a typical environment where scenarios presented in Section 4 may occur. We compared the F-RTO performance to the performance achieved with the conventional RTO recovery, both with SACK TCP and with NewReno TCP. In addition, we conducted experiments with the Eifel algorithm.

5.1 Test Arrangements

The general test setup is illustrated in Figure 5. We emulate the wireless link and the last-hop router by using a real-time wireless network emulator [16]. The end hosts are Linux systems, in which we implemented the F-RTO algorithm. The fixed link is an isolated LAN that is connected to the remote host and to the network emulator.




Figure 5: Test setup.

We selected the link parameters to approximate the properties of a typical currently used wireless wide-area networking system, such as the *General Packet Radio Service (GPRS)* [8]. The emulated wireless link has a bandwidth of 28,800 bps and a propagation delay of 200 ms. The last-hop router has a router buffer for holding seven packets.

In addition to the router buffer, the emulated wireless link uses a link send buffer and a link receive buffer for both uplink and downlink traffic. Any link that provides a retransmission mechanism needs to have a certain amount of buffering capacity. The link send buffer holds packets that have not yet been acknowledged as received, and the link receive buffer collects out-of-order packets for delivering them for the upper layer receiver in the correct order. The link buffers have a size of 1776 bytes, which is large enough to cover the bandwidth-delay product of the link.

We made one set of experiments with a wireless link that does not drop packets, but randomly inflicts sudden delays for some packets. Another set of experiments was made using an unreliable link that drops random packets with given packet drop probabilities. Finally, experiments were conducted by having periods of persistent packet loss on the link. The link scenarios are listed below:

- **Sudden delays.** Since the primary motivation of the F-RTO algorithm is to improve the TCP performance when sudden delays cause spurious retransmission timeouts, we start by a scenario that involves sudden delays on the link. We explained the possible reasons for a sudden delay on the wireless access network in the introduction. Such a delay can occur, for example, due to loss burst with a link layer protocol providing highly persistent reliability. During the delay no packets are delivered to the receiver. In this scenario a packet is delayed with a probability of 0.02. The random delay lengths are exponentially distributed with a mean delay length of 3.5 seconds. Exponential distribution has been reported to characterize the length of the loss periods on a wireless link reasonably well [17]. Even though the link is reliable in this scenario, packet losses may occur due to congestion at the last-hop router.
- **Packet losses.** In this scenario a packet is randomly dropped by the link with given probability. This scenario models the case of an unreliable link layer over a lossy link. Therefore the packet delays on the link are fairly constant. We tested packet loss probabilities of 2 %, 5 %, and 10 %. The packet losses are uniformly distributed. The main purpose of these scenarios is to test the TCP performance when the retransmission timeouts occur due to lost segments. These timeouts occur mainly when retransmissions are lost, since lost original packets are usually recovered by fast retransmit.
- **Bursty losses.** This scenario is to model the effect of link outages when the link layer is not reliable and drops several successive packets. The link conditions are split into two distinct states. In a good state no packets are dropped at the link. When the link is in bad state, all packets in both directions are lost. The link layer does not retransmit any packets. The two states alternate randomly. The good state length is uniformly distributed between 0.1 seconds and 20 seconds. The bad state duration is exponentially distributed with a mean of 3.5 seconds.

In each of the scenarios presented above we test five TCP variants based on the TCP implementation of Linux kernel version 2.4.7 [24]. For the purposes of the experiments, we disabled the ratehalving algorithm used by default in the Linux TCP implementation, and made small modifications to implement the Eifel algorithm as it has been defined by its authors [18]². In addition, we modified the SACK loss recovery to behave similarly to the conservative algorithm recently published by the IETF [6]. Firstly, we test a SACK TCP [20] with the conventional RTO recovery, and with the F-RTO recovery. Secondly, we do experiments with a NewReno TCP [10] with both conventional and F-RTO recovery algorithms. Finally, we test a TCP variant using the TCP timestamp option both with the SACK TCP and with the NewReno TCP. This variant implements the Eifel algorithm based on the use of TCP timestamps. Our Eifel sender implementation continues transmitting new data and reverts the changes made on the congestion window and `ssthresh` when it detects a spurious timeout from the timestamps. The limited transmit algorithm is used with all TCP alternatives.

We use unidirectional 100 KB bulk transfers from the fixed end source to the mobile end sink as the workload. The data is transmitted using a single TCP connection using a maximum segment size of 256 bytes. A small segment size is recommended for slow links in order to achieve a better interactive response times [21], although this is a factor not significant in our tests. For each scenario and TCP variant the experiment is repeated 30 times.

5.2 Results

We present the results of the experiments by using box-plot diagrams. The diagrams compare the throughput of each TCP variant evaluated in the experimentation. The box-plot diagram shows the median throughput for the 30 replications with a horizontal line splitting the filled box. The lower and upper edge of the box represent the 1st and 3rd quartiles of the test results, respectively. The whiskers are drawn at the minimum and the maximum throughput measured with the TCP variant. On rare occasions some test runs were involved with a notably different number of RTOs than the majority of the tests due to randomness of the link events. Because the RTOs typically have a strong effect on the TCP performance, the minimum or maximum throughput values may appear to differ considerably from the results within the upper and lower quartiles in some cases.

In addition to the box-plot diagrams we show with each scenario a table presenting the median values for connection elapsed time from sending the first SYN packet to receiving the last FIN acknowledgement at the sender, the number of packet losses, and the number of retransmitted segments of each TCP variant. If the number of retransmissions is higher than the number of lost packets, at least some of the retransmissions are made unnecessarily. On the other hand, the number of lost packets can be higher than the number

²The recent Linux kernels implement a timestamp-based detection algorithm similar to Eifel, but there are a couple of minor differences to the algorithm described in the original article.

Table 1: Results of the tests with sudden delays. The median values of 30 replications.

TCP Variant	Time (s) / 100 KB	Pkts Lost	Nr. of Rexmits
Eifel w/ SACK	77.94	11	16
F-RTO w/ SACK	76.23	4	12
Regular SACK	94.13	9	57
Regular NewReno	90.72	10	60
F-RTO w/ NewReno	75.18	6	13
Eifel w/ NewReno	79.21	11	19

of retransmissions, because lost acknowledgements do not necessarily trigger retransmissions.

Sudden delays

Figure 6 shows the box-plot diagrams of the throughput measured with different TCP variants. Additionally, Table 1 shows the median values for the connection statistics described above. The results show that using F-RTO improves performance over the conventional RTO recovery both with the SACK TCP and with the NewReno TCP. The number of unnecessary retransmissions with the F-RTO algorithm is considerably smaller than with the conventional RTO recovery algorithm, resulting in improved throughput with the F-RTO algorithm. There is no significant difference between the SACK TCP and the NewReno TCP, when RTOs are triggered by excessive delays.




Figure 6: TCP performance with different variants with excessive delays on the link.

The Eifel TCP avoids most of the unnecessary retransmissions similarly to the F-RTO algorithm. However, the Eifel sender reverts the congestion control parameters back to the values preceding the spurious RTO, and continues sending at the previous rate although the last-hop router could not drain the queue during the delay. Hence, Eifel typically has more packet losses due to congestion than F-RTO, resulting in a slightly lower throughput than F-RTO. This suggests that responding to the spurious RTO by directly reverting

Table 2: Results of the tests with packet errors. The median values of 30 replications.

TCP Variant	Time (s) / 100 KB	Pkts Lost	Nr. of Rexmits
Eifel w/ SACK	80.68	39	26
F-RTO w/ SACK	75.69	36	24
Regular SACK	76.18	36	22
Regular NewReno	82.38	36	26
F-RTO w/ NewReno	81.67	36	26
Eifel w/ NewReno	89.64	38	27

the congestion control parameters may be too aggressive action to take.

Packet losses

Figure 7 illustrates the throughput distribution with different TCP variants when the wireless link has a packet loss rate of 5%. The trend with the packet loss rates of 2% and 10% is similar: the performance of F-RTO is not different from the performance achieved with the conventional RTO recovery, regardless of whether SACK or NewReno TCP is used. In these tests the retransmission timeouts are usually due to lost retransmissions. After the TCP sender has successfully retransmitted the segment that triggered the RTO, it can usually proceed with transmitting new data. Table 2 shows that the number of retransmissions are similar with all TCP variants tested. As expected, the SACK TCP improves the performance over the NewReno TCP, since there are often multiple packet losses in one round-trip time, and SACK recovers more efficiently in such a case.




Figure 7: TCP performance with different variants with packet drop probability of 0.05.

Eifel TCP using SACK and TCP timestamps has a lower throughput than SACK TCP without timestamps. However, a closer examination of the TCP packet traces does not show any problems related to the Eifel algorithm. The difference is explained due to use of the TCP timestamps, which adds 12 bytes of overhead to each packet transmitted,

Table 3: Results of the tests with bursty losses on the link. The median values of 30 replications.

TCP Variant	Time (s) / 100 KB	Pkts Lost	Nr. of Rexmits
Eifel w/ SACK	123.89	50	45
F-RTO w/ SACK	64.94	42	40
Regular SACK	71.23	42	39
Regular NewReno	74.48	42	43
F-RTO w/ NewReno	67.80	39	43
Eifel w/ NewReno	80.03	42	42

resulting in approximately a 4 % increase in the number of packets to send with the small segment size we were using. By using a larger segment size the additional packet overhead would have had less effect on the results.

Bursty losses

Figure 8 shows that the TCP performance with F-RTO does not differ significantly from the performance with the conventional RTO recovery when there are link outages. As described in Section 4.3, the F-RTO sender transmits segments at a similar rate as the conventional RTO recovery, although it transmits two new segments before continuing retransmissions. The difference of whether to transmit the two new segments before or after the retransmissions, does not affect the throughput. Use of the SACK TCP does not notably improve the performance with bursty losses, especially if the losses trigger a retransmission timeout. After the RTO the TCP sender retransmits the unacknowledged segments in slow start, regardless of whether SACK TCP or NewReno TCP is used. Table 3 shows the median connection times and the retransmission statistics for the different TCP variants.




Figure 8: TCP performance with different variants with bursty losses on the link.

The test results show that Eifel TCP gives a clearly worse throughput than the conventional TCP when SACK TCP is used. Our experiments revealed a significant problem when

using TCP timestamps for detecting unnecessary retransmissions in Eifel TCP. We will describe the problem below.

When the link is in the bad state as in our link outage scenario, all packets are dropped for a period of time. Therefore, the latest cumulative acknowledgements generated by the receiver are also dropped by the link. This usually leads to a retransmission timeout and an unnecessary retransmission of a segment that had already arrived at the receiver, but for which the acknowledgement was lost. When this unnecessary retransmission arrives at the receiver, it appears as an out-of-order segment and generates a duplicate ACK carrying a timestamp of an earlier data segment³. Furthermore, because the earlier acknowledgements were lost during the link outage, the duplicate ACK appears as an acknowledgement for new data to the TCP sender. Therefore, the Eifel decision rules declare that the retransmission was spurious, although a number of data segments were lost in the last window.

The Eifel sender responds to the spurious retransmission indication by sending new data and reverting the congestion control variables. However, in the case described above the sender gets back duplicate ACKs because there were data segments missing. The sender enters fast recovery due to the duplicate ACKs and reduces the congestion window. At this point the sender stops sending data for a while to balance the number of outstanding packets to the congestion window size. Because the sender needs to wait for the halved congestion window's worth of acknowledgements to arrive before it can continue retransmitting, and on the other hand, many of the packets were dropped due to link outage, the pipe runs out of packets while the sender is waiting for incoming acknowledgements. Therefore, the Eifel sender has to wait for another RTO to continue the retransmissions for the rest of the lost segments. This leads to a significant degradation of throughput. Unlike SACK, the NewReno TCP ensures that a retransmission is made for each partial ACK. Therefore the Eifel sender often avoids the second RTO with NewReno.

The events presented above showed up very frequently in our experiments with bursty losses, which explains the poor throughput of Eifel TCP in these tests. The reported behavior is specific to TCP timestamps used as an indication of spurious retransmissions, and we do not believe it to show up, if some other mechanism was used for indicating spurious retransmissions instead of TCP timestamps⁴. Furthermore, our preliminary tests show, that if F-RTO recovery is combined with Eifel, the problem described above does not appear.

5.3 Fairness towards conventional TCP

We expect the connections using the F-RTO algorithm to be friendly towards the TCP connections with conventional

³The specification for TCP round-trip time measurements [7] requires that the echoed timestamp should correspond to the most recent data segment that advanced the window

⁴Some TCP implementations do not strictly follow RFC 1323 by echoing the timestamp of a retransmitted segment arriving out-of-order at the receiver. Such implementation would have avoided the problem described here, but may be vulnerable to IP-spoofing attacks.

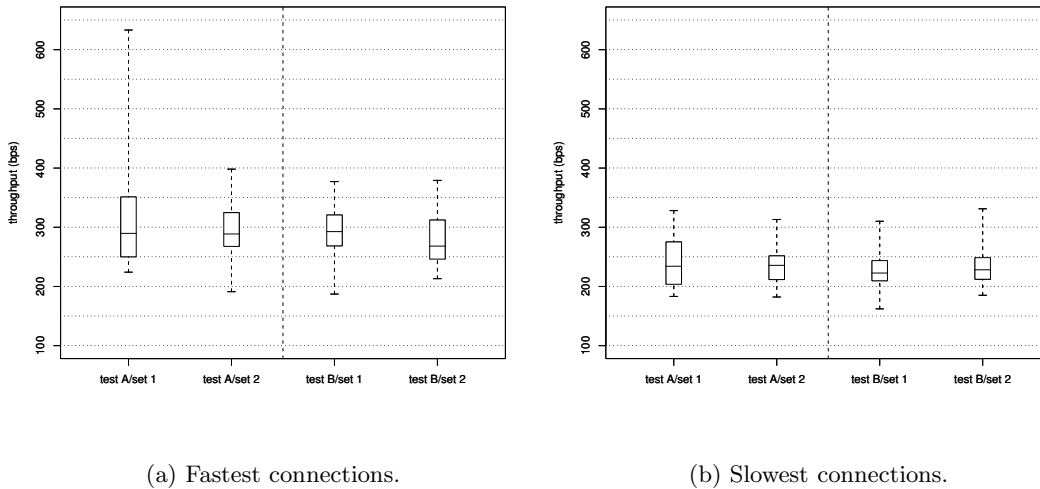


Figure 9: Effect of parallel connections on TCP performance. Test A includes three F-RTO connections, test B uses only conventional RTO recovery.

RTO recovery, because F-RTO is ACK-clocked and it transmits data at an equal rate as the conventional TCP. We back up this reasoning by conducting experiments that use six parallel bulk TCP connections as a workload over the bottleneck wireless link. The workload is separated in two connection sets having three connections each. The three TCP connections in connection set 1 are started at the same time, and the other three TCP connections in connection set 2 are started three seconds after the first connection set. The purpose of this study is to measure how much the connections in connection set 2 interfere with the data transfer in connection set 1. Especially, the effect of the new F-RTO connections on the ongoing TCP transmissions should not differ from the effect of conventional TCP connections.

The test setup with multiple TCP connections is similar to the setup presented earlier in Figure 5, with the exception that it consists of six TCP connections separated in two connection sets. Connection set 1 consists of three TCP connections that use the conventional RTO recovery. Connection set 2 has another three TCP connections, that use F-RTO in test A, and the conventional RTO recovery in test B. All connections transfer 50 KB of bulk data from the remote host to the mobile host. This experiment was made both with and without additional sudden delays on the link. As with the experiments described earlier, the bottleneck link bandwidth is 28,800 bps and the input queue length is 7 packets. Injecting packets from six bulk TCP connections on this kind of network results in severe congestion that causes a number of packet losses and RTOs triggered at the TCP sender. We repeated the experimentation 20 times.

For each connection set we measured the throughput of the TCP connection that was the last to finish its data transfer, i.e. the slowest connection of its connection set. This metric gives a coarse understanding about the fairness between the

TCP connections, because a low throughput of the slowest connection often indicates that the other connections have used a larger share of the common bandwidth. Correspondingly, a high throughput of the slowest connection indicates that the equality between the parallel connections is better. In addition, we present the throughput distribution of the fastest connections in the connection sets.

Figure 9 shows the results of tests with additional delays. Figure 9(a) shows the throughput distribution of the fastest connection for both connection sets in test A using F-RTO connections in connection set 2, and in test B using the conventional RTO recovery in all TCP connections. Figure 9(b) gives the throughput of the slowest connections in the connection sets. The box-plot diagrams show that the connection sets between test A and test B give similar performance. This indicates that the influence of the three new F-RTO connections on the existing TCP connections on the link is not different from the effect of starting three new conventional TCP connections. Similarly, the results of the experiments without random additional delays do not show significant difference between the test runs involving F-RTO connections and the test runs having only TCP connections with the conventional RTO recovery. The results support our reasoning of F-RTO being friendly towards the TCP connections with the conventional RTO recovery.

6. DISCUSSION AND FUTURE WORK

Two problem cases were identified concerning the F-RTO algorithm in Section 3, although they did not have meaningful effect on performance in our experiments. First, the presence of packet reordering can cause the F-RTO sender to enter the conventional RTO recovery with go-back-N retransmissions even if the RTO was spurious. Similarly, duplicate ACKs during TCP fast recovery often prevent the F-RTO algorithm from working. Second, if the sender does

not have new data to transmit, or the receiver's advertised window does not allow the sender to transmit new data in F-RTO algorithm step 2, the sender may not proceed to detect whether the RTO was spurious. We discuss both of these cases and possible solutions for them below.

A potential solution for enhancing F-RTO in the face of reordering comes with the availability of the TCP SACK option. By using the information in the SACK blocks after an RTO, the TCP sender can get acknowledgements for segments transmitted before the RTO and thus detect a spurious RTO according to the principles given in Section 3 even if there were duplicate ACKs arriving. Furthermore, availability of the SACK information makes it possible to better utilize the F-RTO algorithm during fast recovery periods. We intend to study the performance benefits of applying the SACK information in the F-RTO algorithm as a future work item.

In this paper we suggested that if a TCP sender cannot transmit new previously unsent data in F-RTO algorithm step 2, it should revert to the conventional RTO recovery and continue by retransmitting unacknowledged data. This is a straightforward solution and it can be considered to be a safe against unpredictable side-effects, since the sender reverts to the normal and well-known TCP behavior. However, in this situation a spurious RTO may go undetected, so it may be desirable to look for other solutions. Fortunately, there are alternatives that could be used to enable F-RTO in window-limited situations. When limited by the application or receiver's advertised window, a TCP sender could transmit data from the tail of the retransmission queue. The downside of this action is, that it would cause more potentially unnecessary retransmissions that are especially harmful in environments that are not prone to spurious RTOs. Another alternative worth considering when transmission is limited by the receiver's advertised window would be to simply transmit a segment with limited amount data beyond the advertised window. In the worst case the TCP receiver would just drop the segment, if it is incapable of storing it. However, we believe that in many cases the receiver would have acquired available buffer space for the last packet during the time the packet propagates through the network.

In Section 3 we suggested to reduce the congestion window and slow start threshold to half of their previous values after a spurious RTO has been detected. We considered this a conservative enough in various possible scenarios. However, reducing the congestion control parameters may not be the best alternative in all situations, especially when the connection path has a high bandwidth-delay product. In the research and standardization forums there have been different suggestions from entering slow start after a spurious RTO to fully reverting the congestion control parameters to the state preceding the spurious RTO. Some alternatives have been evaluated with a set of simulated network parameters in related work [12], and we intend to compare different response alternatives as well under various network conditions in our future work.

7. CONCLUDING REMARKS

In this work we have shown that it is possible to avoid most of the unnecessary retransmissions following the spurious TCP retransmission timeouts without any additional information in the TCP packet headers. We presented the F-RTO algorithm that avoids the unnecessary retransmissions following the spurious RTO by determining based on the incoming acknowledgements whether to retransmit or continue sending new data. In addition, because the use of F-RTO algorithm effectively avoids unnecessary retransmits, it obviates the NewReno "bugfix" rule that disables fast retransmit during an RTO recovery. This allows more efficient recovery from packet losses in some scenarios. An F-RTO sender follows the conventional TCP congestion control principles by sending data at an equal rate as the conventional TCP and by being clocked by incoming acknowledgements. We showed by experiments that F-RTO improves the TCP performance when there are sudden delays on the link, and it yields competitive performance if the RTOs are caused because of other reasons than delays.

We compared F-RTO with the Eifel algorithm and concluded that their performance is similar in the majority of cases. The Eifel algorithm can perform better than F-RTO, if packet reordering or packet losses are present for the two next segments following the RTO. Eifel makes the detection of spurious RTO already on the first incoming ACK after RTO, whereas F-RTO is able to detect the spurious RTO after two acknowledgements have arrived. However, in most of the cases F-RTO avoids unnecessary retransmissions as successfully as Eifel does, and after one window has been transmitted, it has delivered the same amount of data than Eifel. On the other hand, while making the detection at the first incoming acknowledgement, the Eifel algorithm can end up to a false positive conclusion, if the outstanding acknowledgements and data segments have been lost in the same window due to a loss burst. The DSACK-based algorithms are not directly comparable with F-RTO, since they are not able to detect spurious retransmissions until one window of data has been transmitted and therefore cannot be used for avoiding unnecessary retransmissions.

So far we have verified the F-RTO algorithm by implementing it in the Linux OS and running experiments by emulating the expected behavior of the wireless link. Taking this approach makes it possible to study the F-RTO performance in a real network environment when the TCP traffic is generated by the commonly used network applications. We have contributed our implementation also to the Linux kernel development, and F-RTO is included in the Linux kernels starting from version 2.4.21, and in the development kernel versions 2.5.43 and later. Therefore, it is possible for the reader using Linux to try out the F-RTO algorithm by getting a recent version of the Linux kernel.

Acknowledgements

Authors are thankful for the fruitful discussions with a number of colleagues that have been helpful in developing and improving the F-RTO algorithm. We would also like to thank the reviewers for the detailed reading and useful feedback on the draft version of this paper.

8. REFERENCES

- [1] M. Allman. A Web Server's View of the Transport Layer. *ACM Computer Communication Review*, 30(5), October 2000.
- [2] M. Allman, H. Balakrishnan, and S. Floyd. Enhancing TCP's Loss Recovery Using Limited Transmit. RFC 3042, January 2001.
- [3] M. Allman and V. Paxson. On Estimating End-to-End Network Path Properties. In *Proceedings of ACM SIGCOMM '99*, September 1999.
- [4] M. Allman, V. Paxson, and W. Stevens. TCP Congestion Control. RFC 2581, April 1999.
- [5] E. Blanton and M. Allman. On Making TCP More Robust to Packet Reordering. *ACM Computer Communication Review*, 32(1), January 2002.
- [6] E. Blanton, M. Allman, K. Fall, and L. Wang. A Conservative Selective Acknowledgement (SACK)-based Loss Recovery Algorithm for TCP. RFC 3517, April 2003.
- [7] D. Borman, R. Braden, and V. Jacobson. TCP Extensions for High Performance. RFC 1323, May 1992.
- [8] J. Cai and D. J. Goodman. General packet radio service in GSM. *IEEE Communications Magazine*, pages 122–131, October 1997.
- [9] M. Degermark, B. Nordgren, and S. Pink. IP Header Compression. RFC 2507, February 1999.
- [10] S. Floyd and T. Henderson. The NewReno Modification to TCP's Fast Recovery Algorithm. RFC 2582, April 1999.
- [11] S. Floyd, J. Mahdavi, M. Mathis, and M. Podolsky. An Extension to the Selective Acknowledgment (SACK) Option for TCP. RFC 2883, July 2000.
- [12] A. Gurtov and R. Ludwig. Response to Spurious Retransmission Timeouts. In *Proceedings of IEEE Infocom 2003*.
- [13] A. Gurtov and R. Ludwig. Evaluating the Eifel Algorithm for TCP in a GPRS Network. In *Proceedings of European Wireless '02*, February 2002.
- [14] V. Jacobson. Compressing TCP/IP headers for low-speed serial links. RFC 1144, February 1990.
- [15] F. Khafizov and M. Yavuz. Running TCP over IS-2000. In *Proceedings of IEEE ICC 2002*, April 2002. To appear.
- [16] M. Kojo, A. Gurtov, J. Manner, P. Sarolahti, T. Alanko, and K. Raatikainen. Seawind: a Wireless Network Emulator. In *Proceedings of 11th GI/ITG Conference on Measuring, Modelling and Evaluation of Computer and Communication Systems*, pages 151–166, Aachen, Germany, September 2001. VDE Verlag.
- [17] A. Konrad, B.Y. Zhao, A. Joseph, and R. Ludwig. A Markov-Based Channel Model Algorithm for Wireless Networks. In *Proceedings of ACM MSWiM 2001*, pages 28–36, Rome, Italy, July 2001.
- [18] R. Ludwig and R. H. Katz. The Eifel Algorithm: Making TCP Robust Against Spurious Retransmissions. *ACM Computer Communication Review*, 30(1), January 2000.
- [19] R. Ludwig and K. Sklower. The Eifel Retransmission Timer. *ACM Computer Communication Review*, 30(3), July 2000.
- [20] M. Mathis, J. Mahdavi, S. Floyd, and A. Romanow. TCP Selective Acknowledgement Options. RFC 2018, October 1996.
- [21] G. Montenegro, S. Dawkins, M. Kojo, V. Magret, and N. Vaidya. Long Thin Networks. RFC 2757, January 2000.
- [22] V. Paxson and M. Allman. Computing TCP's Retransmission Timer. RFC 2988, November 2000.
- [23] J. Postel. Transmission Control Protocol. RFC 793, September 1981.
- [24] P. Sarolahti and A. Kuznetsov. Congestion Control in Linux TCP. In *Proceedings of Usenix 2002/Freenix Track*, pages 49–62, Monterey, CA, USA, June 2002.