

# Optimizing the End-to-End Performance of Reliable Flows over Wireless Links

Reiner Ludwig  
Ericsson Research  
Herzogenrath, Germany

Almudena Konrad, Anthony D. Joseph, Randy H. Katz  
Computer Science Division  
University of California at Berkeley

## Abstract

We present the results of a performance evaluation of link layer error recovery over wireless links for reliable flows. Our analysis is based on a case study of the circuit-switched data service implemented in GSM. We collected a large set of block erasure traces measuring in different real-world radio environments, with both stationary and mobile hosts. The analysis of those traces shows that the throughput on this wireless channel can be increased by using a larger (fixed) frame size for the reliable link layer protocol. This yields an improvement of up to 25 percent when the channel quality is good and 18 percent even under poor radio conditions. Our results suggest that adaptive frame length control could further increase the channel throughput. We further show that pure end-to-end error recovery fails as a general solution to optimize throughput when wireless links form parts of the end-to-end path. In many cases, it leads to decreased end-to-end throughput, an unfair load on best-effort networks and a waste of valuable radio resources. In fact, we show that link layer error recovery over wireless links is essential for reliable flows to avoid these problems. Finally, we discuss link and transport layer error control mechanisms and their interactions with end-to-end congestion control schemes. For reliable flows, we argue in favor of highly persistent error recovery and lossless handover schemes implemented at the link layer.

## 1. Introduction

The Internet is evolving to become *the* communication medium of the future. It will not be long before the last circuit switch is taken out of service and virtually all people-to-people, people-to-machine, and machine-to-machine communication are carried end-to-end in IP (Internet Protocol) [38] packets. The tremendous recent growth of the Internet in terms of connected hosts is only matched by the similar tremendous growth of cellular telephone subscribers. While most hosts on today's Internet are still wired, the next *big* wave of hosts has yet to hit the Internet. We believe that the predominant Internet access of the future will be wireless. Not only every cellular phone, but every *thing* that communicates will have: (1) an IP protocol stack and (2) a wireless network interface.

It is well known that the performance of reliable transport protocols such as TCP (Transmission Control Protocol) [39] may degrade when the end-to-end path includes wireless links. This is due to non-congestion related packet losses on the wireless link, causing a TCP sender to underestimate its share of bandwidth along the path. However, related work has mostly focused on the problem that wireless links cause for the congestion control scheme used in most

implementations of TCP [1]. Employing a link layer error recovery scheme over the wireless link removes this problem. Furthermore, for TCP our previous work shows that problems resulting from competition between end-to-end and link layer error recovery do not exist. In [31] we show this for the wireless network examined in this paper. For other wireless networks, related work [5], [14], [34] comes to the same conclusion. In [32] we provide a general analysis of this subject quantifying the high degree of conservatism implemented in TCP's retransmission timer [48].

Motivated by this result, we study the impact of link layer frame sizes on application layer throughput and the consumption of radio resources. We then quantify the benefit of link layer error recovery by comparing it against the performance of pure end-to-end error recovery.

The key premise for our analysis is that we assume a model of a network-limited bulk data transfer based on a reliable flow (e.g., based on TCP). In addition, we assume a *flow-adaptive* link layer implementation [30], [32]. A flow-adaptive link layer error recovery scheme distinguishes among unreliable and reliable flows to control its retransmission persistency. This ensures that link layer error recovery does not interfere with delay-sensitive (usually unreliable) flows (e.g., based on UDP (User Datagram Protocol) [37]). The attractiveness of this solution over approaches that require access to transport layer headers in the network (e.g., [2], [3], [4], [9], [13], [21], [28]), is

- its independence from transport (or higher) layer protocol semantics making it a "non-TCP-specific" solution,
- the possibility of co-existence with any form of network layer encryption as proposed in [27], and
- the fact that no per-flow state needs to be maintained in the network making this solution more scalable.

The concept of flow-adaptive link layer implementations was first introduced in [30]. There we proposed to use the protocol identifier field in the IP header for the purpose of distinguishing among flow types. This solution is limited, because not every UDP-based flow is delay-sensitive as some application layer protocols build end-to-end reliability on top of UDP [20], [44], [47]. In [32] we therefore propose to encode the "reliability" QoS (Quality of Service) requirement in the IP header's differentiated service field [6].

The analysis presented in this paper is based on a case study of the circuit-switched data service implemented in GSM (Global System for Mobile communications). Our measurement-based approach gave us the unique opportunity to capture the aggregate of real-world effects such as noise, interference, fading, and shadowing. This is a key advantage as unrealistic assumptions about the error characteristics of a wireless channel can completely change the results of a performance analysis and lead to a non-optimal design decision. For wireless systems it is therefore particularly important that prototypes are developed early in the design process so that measurement-based performance studies can be carried out.

The rest of this paper is organized as follows: Section 2 provides background on the circuit-switched data service implemented in GSM; Section 3 describes the measurement platform we developed to collect block erasure traces, and explains the goals, assumptions, and the methodology of our trace-based analysis; Section 4 presents and discusses our measurement results; Section 5 discusses error control performance from an end-to-end systems view; and Section 6 closes with our conclusions and plans for future research.

## 2. Circuit-Switched Data in GSM

GSM implements several error control techniques, including adaptive power control, frequency hopping, Forward Error Correction (FEC), and interleaving. In addition, the Circuit-Switched Data service (CSD) provides an optional reliable link layer protocol called Radio Link Protocol. Using Figure 1, we briefly describe the latter three control schemes as implemented for the GSM-CSD. More details can be found in [35].




Figure 1: Error control in GSM Circuit-Switched Data.

GSM is a TDMA-based (Time Division Multiple Access) circuit-switched network. At call-setup time a mobile terminal is assigned a user data channel, defined as the tuple [carrier frequency number, slot number]. The slot cycle time is 5 milliseconds on average, allowing 114 bits to be transmitted in each slot which yields a gross data rate of 22.8 kbit/s. The fundamental transmission unit in GSM is a *data block* (or simply *block*). The size of an FEC encoded data block is 456 bits (the payload of 4 slots). In GSM-CSD, the size of an unencoded data block is 240 bits resulting in a data rate of 12 kbit/s (240 bits every 20 ms).

Interleaving is a technique that is used in combination with FEC to combat burst errors. Instead of transmitting a data block in four consecutive slots, it is divided into smaller fragments. Fragments from different data blocks are then interleaved before transmission. The interleaving scheme chosen for GSM-CSD, interleaves a single

data block over 22 TDMA slots. A few of these smaller fragments can be completely corrupted while the corresponding data block can still be reconstructed by the FEC decoder. The disadvantage of this deep interleaving is that it introduces a significant one-way latency of approximately 90 ms<sup>1</sup>. This high latency can have a significant adverse effect on interactive protocols [30].

The Radio Link Protocol (RLP) [15] is a full duplex logical link layer protocol. RLP uses selective reject and checkpointing for error recovery. The RLP frame size is fixed at 240 bits aligned to the above mentioned FEC coder. RLP introduces an overhead of 48 bits per RLP frame yielding a user data rate of 9.6 kbit/s in the ideal case (no retransmissions)<sup>2</sup>. RLP transports user data as a transparent byte stream (i.e., RLP does not “know” about IP packets). RLP loses data when the link is reset, e.g., after a maximum number of retransmissions of a single frame has been reached. Although, this can have a severe impact on higher layer protocol performance, it rarely happens under “normal” radio conditions [31].

## 3. Analysing Block Erasure Traces

In this Section, we describe the measurement platform we developed to collect *block erasure traces*. We have used this measurement platform in [31] to study the interactions between TCP and RLP. We then explain the goals, assumptions, and the methodology of our trace-based analysis.

### 3.1 What is a Block Erasure Trace?

The error characteristics of a wireless channel over a certain period of time can be captured by a bit error trace. A bit error trace contains information about whether a particular bit was transmitted correctly. The average Bit Error Rate (BER) is the first-order metric commonly used to describe the trace. The same approach can be applied on block level instead of on bit level. Hence, a block erasure trace contains information about whether a particular data block was transmitted correctly and the Block Erasure Rate (BLER) denotes the average rate at which block erasures occur in the trace.

The error characteristics we have measured are only valid for the particular FEC and interleaving scheme implemented in GSM-CSD (see Section 2). Nevertheless, we believe that our results presented in Section 4 provide new insights into how this widely deployed system can be optimized and suggests techniques that can be used to design future wireless links.

### 3.2 Measurement Platform

Our measurement platform is depicted in Figure 2. A single-hop network running the Point-to-Point Protocol (PPP) [43] connects the mobile to a fixed host which terminates the circuit-switched GSM connection. Various tools can be used to generate traffic on the link (e.g., `ping` as described in [45]). To collect block erasure traces, we have ported the RLP protocol implementation of a commercially available GSM data PC-Card to BSDi3.0 UNIX. In addition, we developed a protocol monitor for RLP which we call RLPDUMP. It logs (among other RLP events) whether a received block could be correctly recovered by the FEC decoder. This is possible because every RLP frame corresponds to an FEC encoded data

1. Note, that voice is treated differently in GSM. Unencoded voice data blocks have a size of 260 bits and the interleaving depth is 8 slots.
2. Note, that the transparent (not running RLP) GSM-CSD service introduces a wasteful overhead of modem control information that also reduces the user data rate to 9.6 kbit/s.

block (see Section 2). Thus, a received block suffered an erasure whenever the corresponding RLP frame has a frame checksum error. We then generated bulk data traffic and used RLPDUMP to capture the corresponding block erasure trace.




Figure 2: The measurement platform.

We have only performed measurements in commercially deployed GSM networks, where the network-side of RLP was not accessible. Thus, we could only collect downlink block erasure traces. Nevertheless, this allowed us to understand the GSM-CSD channel error characteristics to a degree that was sufficient for our analysis. Additional uplink block erasure traces would not have changed our conclusions.

### 3.3 Analysis Goals and Assumptions

Our goal is to evaluate the performance of the following two protocol design alternatives for a reliable data transfer over a path that includes a wireless link:

- End-to-end error recovery complemented with flow-adaptive link layer error recovery running over the wireless link.
- Pure end-to-end error recovery.

In general, “pure end-to-end” implies that no transport layer state is maintained in the network and that no assumptions are made about the existence of dedicated support from the link layer (e.g., error recovery) or the network layer (e.g., cell handover indications, as proposed in [10]). Nevertheless, throughout this paper, when we use the term “pure end-to-end error recovery”, we imply that the wireless link is *not* protected by link layer error recovery.

We perform the evaluation of the two alternatives through a case study of the GSM-CSD wireless link. First, we investigate the impact of changing the (fixed) RLP frame size on application layer throughput and the consumption of radio resources (e.g., spectrum and transmission power). We then quantify the benefits of link layer error recovery by comparing it against the performance of pure end-to-end error recovery, as proposed in [42].

The performance difference between the two protocol design alternatives depends on the wireless channel’s time varying error characteristics versus the channel’s packet<sup>3</sup> transmission delay. This is

3. Comprising the transport layer segment, the transport and network layer headers, and the packet framing overhead (e.g., PPP [43]) required on this channel. Note, that this does *not* include RLP overhead.

sketched in Figure 3, where “burst error” denotes time intervals during which data in transit is corrupted to the extent that it cannot be recovered at the destination. With respect to GSM-CSD, a burst error corresponds to a series of back-to-back block erasures where the channel is error-free before and after that series. A wireless channel’s error characteristic can be described by the length of burst errors and their correlation expressing the degree of clustering. Link layer error recovery is less effective on wireless links where the length of burst errors is large compared to the packet transmission delay (see “Channel 1” in Figure 3). In this case, pure end-to-end error recovery often yields higher throughput results by saving link layer protocol overhead. Another case is sketched with “Channel 2” in Figure 3 where the length of burst errors is small compared to the packet transmission delay and the burst errors often occur isolated. In this case, the link layer overhead is likely to be amortized when the “right” frame size is chosen. Studying this trade-off requires a realistic characterisation of the wireless channel and motivates our measurement-based analysis approach further outlined in Section 3.4.




Figure 3: Two different channel error characteristics.

The key premise for our analysis is a model of a network-limited bulk data transfer based on a reliable flow (e.g., TCP-based). A *flow* is a single instance of an application-to-application flow of packets which is identified by source address, source port, destination address, destination port and protocol identifier [6]. The end-to-end throughput provided by a *network-limited* flow is limited by the path’s bottleneck link. Consequently, to compare throughput among the two alternatives, we assume that the GSM-CSD wireless link is the path’s bottleneck. The requirements of applications using *reliable flows* are simple: transfer the application layer data object as fast as possible *but* reliably, i.e., the transfer fails if the data object is corrupted when received by the destination. This translates into similarly simple QoS requirements for reliable flows: maximize throughput while the per packet delay is (almost) irrelevant<sup>4</sup>.

We perform a best-case analysis that assumes that the bulk data transfer always fully utilizes the wireless bottleneck link. We use the term *utilization* as defined in [31]: a link is fully utilized if it never runs idle and never transmits a packet/frame which had already been successfully transmitted. The latter can happen in TCP, which exhibits go-back-N behavior after spurious timeouts [32]. Concerning link layer error recovery, this implies (1) the use of a selective reject based protocol, like RLP; and (2) fully-persistent retransmissions (i.e., a large value for “maximum number of retransmissions” only reached when the link effectively becomes disconnected). It also requires the use of large enough windows to allow the link layer sender to always fully utilize the link. This

4. In theory, it would not matter in a file transfer if the first packet reached the destination last. What usually matters is that the file transfer is completed in the shortest amount of time. In practice, e.g., transport layer receiver buffers required for packet re-sequencing place a limit on the maximum per packet delay that is tolerable without affecting performance. This limit is nevertheless low.

avoids the *stalled window* condition [32]. Concerning a network-limited, TCP-based flow and because of the congestion control policies required by [1], this implies that the maximum number of packets that the flow may queue at the bottleneck link before a packet is dropped, exceeds the pipe capacity. The *pipe capacity* is the minimum amount of data a flow needs to have in transit to fully utilize its share of bandwidth at the bottleneck link. The best-case assumption ignores interactions with end-to-end congestion control schemes that may lead to an under-utilization of the bottleneck link. For TCP over RLP, this is valid [31]. For pure end-to-end error recovery, however, this is unrealistic as discussed in Section 4.3 and Section 5. Nevertheless, a best-case study indicates the theoretical maximum application layer throughput that pure end-to-end error recovery could provide. Moreover, the best-case application layer throughput, as defined here, directly translates into radio resource consumption. For example, if transport layer sender A only achieves half the throughput that sender B achieves, it is using twice as much radio resources.

### 3.4 Analysis Methodology

Following the approach suggested in [36], we are not interested in identifying those physical link factors that caused measured block erasures. Rather, we are interested in the aggregate result captured by block erasure traces. We have collected block erasure traces for over 500 minutes of “air-time” and distinguish between measurements where the host was stationary versus mobile when driving in a car. All stationary measurements were taken in the exact same location. We were not able to log internal receiver signal strength measurements from the mobile phone to correlate them with the block erasure traces. Instead, we read the mobile phone’s visual signal level indicator ranging from 1-5. The following three radio environments were chosen:

- A. Stationary in an area with good receiver signal strength (3-4): 258 minutes.
- B. Stationary in an area with poor receiver signal strength (1-2): 215 minutes.
- C. Mobile in an area with mediocre receiver signal strength (2-4): 44 minutes.

Clearly, the size of an RLP frame does not need to match the size of an unencoded data block. Any multiple of the size of an unencoded data block would have been a valid design choice. In fact a multiple of 2 has been chosen for the new version of RLP [16] in the next generation of GSM-CSD, which also uses weaker FEC [17]. Larger frames introduce less relative overhead per frame, but an entire frame has to be retransmitted even if only a single data block incurs an erasure. Applying our technique of *retrace analysis*, we study this trade-off using the collected block erasure traces. Based on a given block erasure trace and a given bulk data transfer size, retrace analysis reverse-engineers the value of target metrics (e.g., channel throughput or number of retransmissions). It emulates RLP while assuming a particular fixed frame size and fixed per frame overhead. We then iterate the retrace analysis over a range of RLP frame sizes, defined as multiples of the data block size. We can thereby, e.g., find the frame size that maximizes the bulk data throughput for a particular block erasure trace.

We use different block erasure traces for our analysis. *Trace\_A* is a concatenation of all block erasure traces we collected in environment A. Likewise, *trace\_B* and *trace\_C* are concatenations of all block erasure traces we collected in environment B and C, respectively. We then choose an appropriate bulk data size to cover the

entire trace (e.g., for *trace\_B* a size corresponding to a transmission time of 215 minutes was chosen). Once the retrace analysis reaches the end of a trace, it wraps around to its beginning. In addition, we investigate the impact of error burstiness, i.e., the extent to which the distribution of block erasures within a trace influences our results. For that purpose, we artificially generated three more “non-bursty” block erasure traces, *trace\_A\_even*, *trace\_B\_even* and *trace\_C\_even*. These have the same BLER as the corresponding real traces, but with an even block erasure distribution, i.e., those traces have single and isolated block erasures with a constant distance to each other.

## 4. Measurement Results

In this Section, we show that the throughput of the GSM-CSD channel can be improved by up to 25 percent by increasing the (fixed) RLP frame size. Our results also suggest that techniques like adaptive frame length control and adaptive FEC are worth further exploration for increasing channel throughput. Furthermore, we argue why pure end-to-end error recovery fails to optimize end-to-end performance on this wireless link.

### 4.1 Block Erasure Rates and Burstiness

Deriving the overall BLERs for *trace\_A*, *trace\_B* and *trace\_C* would have delivered little useful information. Instead, we also capture how “fast” the BLER changes over time in a given radio environment. We therefore divide each trace into one minute *sub-traces* and treat each of those independently.




Figure 4: Measured BLERs.

Figure 4 summarizes the BLERs that we have determined in this manner. The BLERs for the sub-traces of *trace\_A* are not shown because we found *trace\_A* to be almost free of block erasures: over 96 percent of all sub-traces do not have a single block erasure and the remaining ones have a BLER of less than 0.0025. This result shows how strongly the GSM-CSD channel is protected by FEC and interleaving, leaving little error recovery work for RLP. This is especially striking because radio environment A was far from ideal, as it only provided a receiver signal strength of 3-4. Many radio environments provide a maximum receiver signal strength of 5. This suggests that a weaker FEC scheme and/or a larger RLP frame size would increase channel throughput. The results for *trace\_B* and *trace\_C* are similar but different from the results for *trace\_A*. In these, over 30 percent of all sub-traces have no single block erasure or a BLER of less than 0.01. But overall the BLERs vary considerably and can be as high as 0.28. These large variations take place over time scales of one minute, which corresponds to 3000 RLP

frames. This is “slow” enough to make adaptive error control schemes applicable even within the same radio environment. This is important because otherwise such schemes would only be effective if the mobile user changed location to a different radio environment. The reason is that adaptive error control schemes only adapt with a certain latency, which depends on the delay required to feedback channel state information. In future work, we will study the potential of adaptive frame length control (e.g., proposed in [14] and [29]) to increase channel throughput. This decision is partly driven by our measurement-based analysis approach and the fact that we are currently not able to implement schemes like adaptive FEC (e.g., standardized in [18] and [19]) in our measurement platform.




Figure 5: Burst error length distribution.

Figure 5 shows the cumulative distribution function for the burst error lengths, i.e., the number of consecutive blocks that suffered an erasure, for *trace\_B* and *trace\_C*. There was no point in showing this for *trace\_A*, as it was basically error-free. Over 50 percent of burst errors are only 1 or 2 blocks long. Longer error bursts are more common when the mobile host is stationary, e.g., in *trace\_B* less than 5 percent of all error bursts are larger than 26 blocks, whereas in *trace\_C* this number drops to 18. This comparison is valid as the BLERs of both traces are of the same order. As discussed in Section 3.3, the distributions shown in Figure 5 do not sufficiently describe the wireless channel’s error characteristic. They do not show whether the burst errors occur in clusters or isolated, i.e., the correlation between error bursts is not captured. In the following Section, we show how the (fixed) frame size that maximizes channel throughput can be used to quantify this correlation.

#### 4.2 Error Burstiness Allows for Larger Frames

In this Section, we determine the fixed RLP frame size that maximizes channel throughput in the radio environments *A*, *B*, and *C*. This also indicates maximum throughput improvement that adaptive frame length control could yield. The implementation complexity of such techniques must be justified with substantial performance improvements. Thus, if the margin for improvements was too small, it would not be worth studying algorithms for adaptive frame length control in GSM.

We perform the retrace analysis described in Section 3.3 leading to the results shown in Figure 6. An optimal frame size of 1410 bytes yields a throughput of 1423 bytes/s for *trace\_A* and a frame size of 210 bytes maximizes throughput to 1295 bytes/s for *trace\_C*. The results for *trace\_C* are close to those of *trace\_B*. The gradual performance improvements in the case of *trace\_A* rapidly decrease

above a frame size of 210 bytes. A frame size of 210 bytes still yields a throughput of 1392 bytes/s. This indicates that for an adaptive frame length control algorithm, it would be sufficient to adapt the frame size in a range of about 30-210 bytes.

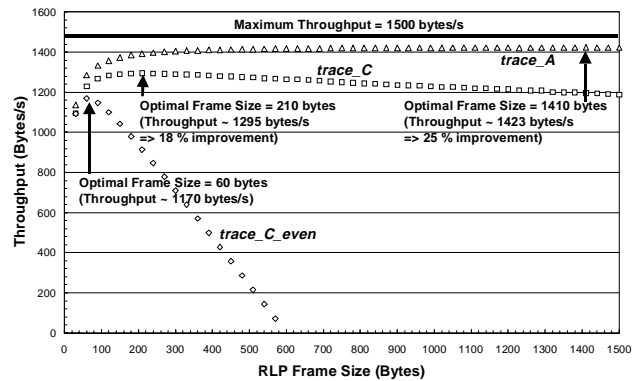


Figure 6: Throughput versus frame size.

A key result of our analysis is that the (fixed) frame size chosen for RLP was overly conservative. Increasing it to 210 bytes improves the channel throughput by at least 18 and up to 23 percent depending on the radio environment<sup>5</sup>. This still leaves a (theoretical) margin of potential throughput improvement of 8-16 percent for adaptive frame length control, depending on the radio environment. We were not able to verify which studies led to the decision to standardize an RLP frame size of 30 bytes [15]. However, our results show that they must have been based on an unrealistic error model of the GSM-CSD radio channel. This highlights the importance of measurement-based analysis of protocol performance over wireless links.

Another result is that the error burstiness on the GSM-CSD channel allows for larger frame sizes than if block erasures are not bursty. This effect can be seen by comparing the graphs *trace\_C* and *trace\_C\_even* in Figure 6. The retrace analysis for *trace\_C\_even* yields an optimal frame size of only 60 bytes (*trace\_B/trace\_B\_even* is similar). One could view the quotient of the optimal frame size for an error trace (bit error trace or block erasure trace) and the corresponding “\_even” trace as the *burst error factor*. The closer a trace’s burst error factor is to 1, the less the corresponding channel exhibited error burstiness. Note, that the burst error factor also depends on the per frame overhead chosen for the retrace analysis. To eliminate this dependency, one could base the definition of the burst error factor on a retrace analysis that assumes a per frame overhead of zero.

#### 4.3 Problems of Pure End-to-End Error Recovery

Based on *trace\_C*, we perform the best-case analysis described in Section 3.3 using TCP [39] as an example of a pure end-to-end error recovery protocol. For that purpose we repeat the retrace analysis assuming a per *MTU* (*Maximum Transmission Unit*) [45] overhead of 47 bytes (20 bytes TCP header, 20 bytes IP header, and 7 bytes of PPP overhead). The retrace analysis shows that the end-to-end throughput is maximized with an *MTU* size of 690 bytes. The reason for the difference with RLP is the larger overhead per transmis-

5. For example, for *trace\_A*, the retrace analysis yields a throughput of 1392 bytes/s for a frame size of 210 bytes and a throughput of 1138 bytes/s for a frame size of 30 bytes/s. For *trace\_B* and *trace\_C*, these frame sizes yield a throughput of 1295 bytes/s and 1096 bytes/s, respectively.

sion unit. The first row of Table 1 shows the result for commonly used MTU sizes. The second row shows the end-to-end throughput that is achieved when running RLP with a frame size of 210 bytes, providing a channel throughput of 1295 bytes/s (see Figure 6).

	MTU 296 bytes	MTU 576 bytes	MTU 1500 bytes
Pure End-to-End (No Header Compr.)	1151	1219	1196
End-to-End with RLP (No Header Compr.)	↑5.2% 1094	↑2.4% 1191	↓4.9% 1255
End-to-End with RLP (With Header Compr.)	↓7.6% 1239	↓3.8% 1265	↓7.4% 1284

Table 1: Application layer throughput in bytes/s.

Pure end-to-end error recovery achieves a 2.4 and 5.2 percent increase in best-case application layer throughput for MTU sizes of 576 and 296 bytes, respectively. This shows that pure end-to-end error recovery consumes less radio resource for these MTU sizes as discussed in Section 3.3. However, even when TCP-SACK [33] is used, it is unlikely that the advantage in end-to-end throughput would be achieved in practice, due to interference with the end-to-end congestion control scheme commonly implemented in TCP [1]. The benefit of link layer error control becomes evident with larger MTU sizes (e.g., the commonly used 1500 bytes - see Table 1) and when IP header compression is used over the wireless link<sup>6</sup>.

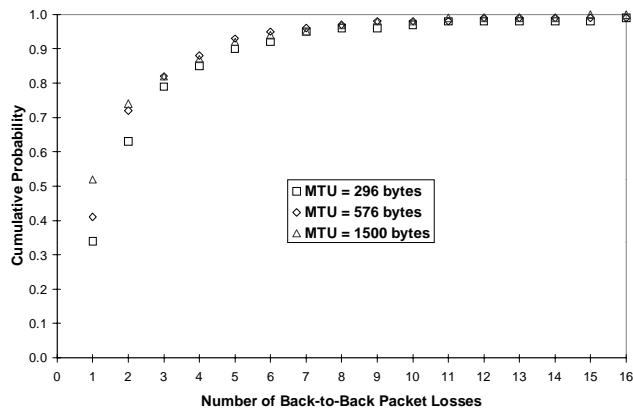


Figure 7: Distribution of back-to-back packet losses.

For pure end-to-end error recovery, IP header compression as defined in [12] and [23], are not an option<sup>7</sup>. The reason is that the mechanism described in [23] causes the loss of an entire window worth of packets for each packet lost *after* the compression point. While, the *Twice* algorithm proposed in [12] is more robust, it causes the same problem when two or more packets with com-

6. In this case, we assume that the TCP/IP header is compressed to 6 bytes. Although, compressed TCP/IP headers are typically 4 bytes long, a network-limited TCP connection drops one packet - in the ideal case - per congestion avoidance cycle. This causes one packet to be sent with a full header (40 bytes), and 2 packets - after the packet loss and after the retransmission - to be sent with a compressed header of 7 bytes. Given the bandwidth-delay product of a GSM-CSD link this leads to an average of about 6 bytes.  
7. Unfortunately, this is not an option in any case when network layer encryption spans across the link.

pressed headers are lost back-to-back. However, this is a likely event for the GSM-CSD wireless link (if not protected by RLP) as shown in Figure 7. The cumulative distribution of the number of back-to-back packet losses shows that 66, 59, and 48 percent of all such losses have a length of 2 or larger for an MTU of size 296, 576, and 1500 bytes, respectively. Alternatively, [12] also defines a “header request” mechanism but as our results show, link layer error recovery would be more appropriate on this wireless link.

One could argue in favor of pure end-to-end error recovery by requiring the wireless link’s MTU to be set to small values. Transport protocols like TCP could then use the *MSS option* (Maximum Segment Size) or *path MTU discovery* [45] to adapt the path’s MTU accordingly. However, that does not work when the link’s end points (e.g., the PPP peers) are not “aware” that the link includes a wireless segment as, e.g., in GSM-CSD (see Figure 2). Also, the path’s MTU cannot be re-negotiated during a connection in current transport protocols<sup>8</sup> when the wireless link’s error characteristics change. Link layer error recovery does not have these problems. It is independent of MTU sizes and also interworks well with IP header compression. Future systems favor link layer error recovery even more. Weaker FEC schemes are being deployed<sup>9</sup>, which further decrease the throughput optimal frame size on those wireless links. Also, the next generation of the IP protocol [11] requires a minimum MTU of 1280 bytes and recommends an MTU of 1500 bytes or more on links such as GSM-CSD.




Figure 8: Number of end-to-end retransmissions.

Another shortcoming of pure end-to-end error recovery is that each retransmission has to traverse the entire path. This is depicted in Figure 8 for *trace\_C*, showing the number of retransmissions (as fraction of the overall number of transmissions) that are required for a range of different MTU sizes. The commonly used MTU size of 1500 and 576 bytes would cause 18 and 12 percent retransmissions, respectively. Such flows impose an unfair load on a best-effort network, such as the Internet, and also on shared wireless access links. Apart from fairness, a higher fraction of retransmissions also decreases the end-to-end throughput if the corresponding packets had already traversed the bottleneck link regardless of where it is located in the path. This is a common situation when, e.g., data is downloaded from the Internet and the last-hop is an unreliable wireless link. End-to-end error recovery complemented with link layer

8. Implementing such a mechanism would also be a poor design choice as optimizing a link’s frame length is not an end-to-end issue.  
9. Weaker FEC schemes are used in the new GSM-CSD service [17] and the upcoming GSM packet-switched data service [19].

error recovery running over the wireless link “typically” does not require a single end-to-end retransmission [32].

## 5. Discussion: Error Control Performance

When addressing the problem of “reliable flows over wireless links” it is not sufficient to study a particular link as an isolated entity. Moreover, the entire end-to-end system including link and transport layer error control mechanisms and their interactions with end-to-end congestion control schemes has to be taken into account. We have deliberately ignored this for the best-case analysis presented in Section 3 and Section 4. In this Section we discuss this subject in detail.

In a connection-less best-effort network like the Internet, senders have to interpret “packet loss” a signal of congestion [1], [8], [22], [40]. We refer to such flows as being *loss-responsive*. However, “packet loss” is not unambiguous. A packet can get lost because of packet drop due to congestion at a router or a host, or because of packet corruption due to a transmission error causing a *non-congestion related packet loss*. A sender is cannot discriminate among these, because packet corruption usually leads to a frame checksum error and subsequent discard of the packet at the link layer. Transmission errors beyond a certain rate inevitably lead to an underestimation of available bandwidth by the sender of a loss-responsive flow, reflected in reduced application layer throughput. This explains why wireless links are problematic: whereas transmission error rates on today’s wireline links can be neglected, this is not true for wireless links. In addition, when hosts are mobile, cell handovers may cause data loss and some wireless networks may in certain situations only provide intermittent connectivity. We view the latter two cases as “long” transmission errors that do not have to be treated different from “normal” transmission errors.

Error control performance is the strongest argument in favor of flow-adaptive link layer implementations [30], [32]. In Section 4, we show that link layer error recovery over wireless links is essential for reliable flows to optimize end-to-end performance (throughput and fairness) while minimizing radio resource consumption<sup>10</sup>. Implementing an optimal solution only from the end points of a path seems impossible; even if knowledge about the time varying error characteristics of each wireless link in the path was available. Flow-adaptive link layer implementations adapt their error control schemes to the individual QoS requirements of each flow sharing the link. For example, channel throughput can be aggressively maximized for a reliable flow as the per packet delay is less important. The flows’ QoS requirements are derived (only) from the IP headers, e.g., the proposed differentiated service field [6], and are made available to the link layer on a per packet basis. This is implemented by an appropriate network-layer/link-layer interface or the link layer itself inspecting each IP header. The latter violates the principle of “protocol layering”, but has the advantage that existing interface implementations need not be changed. Therefore, we promote flow-adaptive implementations of link layer error control schemes as what the end-to-end argument [41] calls “an incomplete version of the function provided by the communication system [that] may be useful as a performance enhancement”. We believe that carrying the application’s QoS requirements as part of the flow’s headers and accordingly adapting lower layer functions, such as error control, advances the discussion provided in Section 2.3 of [41].

---

10. We believe that a similar line of argument applies to unreliable but delay-sensitive loss-responsive flows.

There have been debates [14], [26], [31], [46] about how persistent link layer error recovery should be implemented. For non-flow-adaptive link layer implementations, retransmission persistency must be low to avoid interference with delay-sensitive flows. However, for senders (in this context we omit the prefix “link layer”) that are capable of discriminating reliable from delay-sensitive flows the question about how to treat reliable flows remains. The options are to implement either semi-reliable or fully-reliable link layer error recovery. A *semi-reliable sender* gives up after a few retransmission attempts, discards the corresponding packet, and resumes transmission with the next packet. This introduces a *retransmission delay* on the order of a few 100 milliseconds [26] or in more persistent implementations on the order of a few seconds, as in RLP (see Section 2). A *fully-reliable sender*, on the other hand, does not lose any packets even over long link outages, up to some conservative termination condition<sup>11</sup>. An upper limit for such a condition is the *MSL (Maximum Segment Lifetime)* [39] of 2 minutes<sup>12</sup> which also serves as an upper bound for the reassembly timeout after IP fragmentation [7]. We are not aware of the existence of such a reliable link layer protocol implementation.

The end-to-end argument [41] tells us that it is not worth the effort to implement “perfect” reliability at the link layer. Yet, our design should eliminate non-congestion related packet loss to avoid interference with end-to-end congestion control schemes. Implementing semi-reliable link layer error recovery is always a compromise that avoids this conflict by emphasizing end-to-end error recovery [26]. However, this approach has some fundamental problems. First, the sender has no way to decide *when* to “give up” and discard the packet to, e.g., stay within the bounds of TCP’s retransmission timer and/or to reduce the rate of non-congestion packet losses below a certain target rate. This is not feasible as it requires knowledge of the path’s round trip time, which cannot be known at the link layer. Therefore, a semi-reliable sender requires a channel with strong FEC in order to keep the rate of false congestion signals due to non-congestion related packet discards low. Together with the non-data-preserving property of semi-reliable link layer error recovery, this cannot yield optimal end-to-end throughput. Another fundamental problem occurs in case of temporary link outages, e.g., when a user temporarily roams into an area without wireless connectivity. In this case, all of the flow’s unacknowledged packets are eventually discarded by semi-reliable link layer sender. This causes an idle wait for a possibly backed-off transport layer retransmission timer to expire before the next packet is sent (up to 64 seconds [45]), while the link may already have become available again. If, on the other hand, packets are still queued at the wireless link, the end-to-end flow of data is re-started immediately after the link has become available. Therefore, we promote the implementation of fully-reliable link layer error recovery for reliable flows as it has none of these problems and guarantees that any loss at the link is due to congestion<sup>13</sup>. This is the right signal to give to senders of loss-responsive flows. In case of temporary link outages, this most likely causes a spurious timeout which in turn forces a go-back-N behavior in TCP [32], but (1) that is still better than the idle wait and

---

11. Note, that this has nothing to do with queue management techniques. Packets that are dropped by the network layer according to simple drop-tail or a more advanced active queue management scheme [8] are never handed to the link layer.

12. In theory, additional fully-reliable links could exist “further down” the path. Thus, a more conservative upper limit is to divide the MSL by the value of the *TTL (Time To Live)* [45] field in the IP header.

13. Apart from the more unlikely events of link layer error detection failures.

(2) can be avoided with the Eifel algorithm described below. For the same reasons that favor fully-reliable link layer error recovery for reliable flows, we promote that wireless networks that provide seamless terminal mobility implement mechanisms to support lossless intra- (and if possible also inter-) system cell handovers for data belonging to reliable flows.

The transport layer *Eifel* algorithm [32], increase end-to-end error recovery performance. When applied to TCP (*TCP-Eifel*) it uses the timestamp option [24] to eliminate the *retransmission ambiguity problem* [25]. It thereby avoids unnecessary duplicate packet transmissions caused by TCP's go-back-N behavior after spurious timeouts. In addition, it avoids unnecessary reductions of the flow's send rate by restoring the TCP sender's *congestion window* [22] after spurious timeouts and after spurious fast retransmits. The latter happens in the case of packet re-orderings beyond the *DUPACK threshold* [45]. We promote the deployment of the Eifel algorithm for TCP in particular, and in general for any reliable transport protocol. TCP-Eifel can be incrementally deployed as it is backwards compatible and does not change TCP's congestion control behavior [1]. In addition, we argue that an adaptive transport layer retransmission timer should not be tuned to prevent all spurious timeouts. Some wireless networks often only provide intermittent connectivity. Despite a highly conservative retransmission timer, e.g., implemented in TCP [32], spurious timeouts cannot be avoided in such an environment unless dedicated transport layer support is implemented in the network, as proposed in [9]. A transport layer retransmission timer which is too conservative has a negative impact on end-to-end performance whenever the sender has to resort to a (long) timeout to recover a lost packet. This affects interactive applications, but also bulk data transfers when the receiver's receive buffer cannot absorb any additional out-of-order packets. As a result, the sender is blocked from sending further packets. Instead, an adaptive transport layer retransmission timer should be "reasonably" conservative, while a sender should be able to detect spurious timeouts and react appropriately by using the Eifel algorithm.

Given a reliable wireless link, transport layer selective acknowledgements [33] have nothing to add apart from improving end-to-end performance in the case of burst packet losses that may be caused by congestion. Nevertheless, any transport layer protocol should be robust against such cases. Also, many legacy wireless networks do *not* provide reliable wireless links. Thus, for both reasons we promote the implementation of transport layer selective acknowledgements.

## 6. Conclusion and Future Work

In this paper, we presented the results of a performance evaluation of link layer error recovery over wireless links for reliable flows. The analysis is based on a case study of the circuit-switched data channel implemented in GSM. We show that the throughput on this channel can be increased by using a larger (fixed) link layer frame size. This yields an improvement of up to 25 percent when the channel quality is good and 18 percent even under poor radio conditions. Larger frame sizes are made possible due to the channel's error burstiness, a quantity we define as the burst error factor. These results highlight the importance of measurement-based analysis in wireless networks where protocol performance is highly dependent on the error characteristics of the wireless channel. Our results also suggest that techniques such as adaptive frame length control are worth further exploration for increasing channel throughput. This is a topic for our future research, as we plan to implement a measure-

ment-based adaptive frame length control scheme in our testbed. This will also require a study of interactions with adaptive FEC schemes as implemented in upcoming wireless systems such as the General Packet Radio Service (GPRS) and the Universal Mobile Telecommunications System (UMTS).

We show that pure end-to-end error recovery fails as a general solution for optimizing throughput when wireless links form parts of the end-to-end path. The fundamental problems are that the path's MTU is often too large to yield efficient error recovery, and that the path's end points are not capable of dynamically adapting their MTU to changing local error characteristics on (possibly multiple) wireless links. In many cases, this leads to decreased end-to-end throughput, an unfair load on a best-effort network, such as the Internet, and a waste of precious radio resources. In fact, we show that link layer error recovery over wireless links is essential for reliable flows to avoid these problems.

For reliable loss-responsive flows, we promote fully-reliable link layer error recovery and lossless handover mechanisms. This eliminates non-congestion related packet losses to the highest possible degree. Furthermore, a flow-adaptive implementation ensures that link layer error recovery does not interfere with delay-sensitive flows. The attractiveness of this solution is that it does not require access to transport layer headers in the network. The major advantages being the independence from transport (or higher) layer protocol semantics and the possibility of co-existence with any form of network layer encryption. As an optional, but complementary mechanism for reliable transport protocols, we promote the implementation of the Eifel algorithm. It increases end-to-end performance by avoiding unnecessary duplicate packet transmissions and unnecessary reductions of the flow's send rate. We believe that the combination of these mechanisms solves the problem of "reliable flows over wireless links" in the best possible way. Beyond that, we are currently experimenting with an implementation of TCP where the sender does not retransmit only 12 times [45] and then closes the connection, but instead retransmits until the *application* decides to close the connection. This is necessary in wireless networks where long disconnections may occur but are not critical for the application's operation.

In future work, we intend to focus on flow-adaptive link layers that optimize error control schemes for the QoS requirements of unreliable, but delay-sensitive flows. This will require that much more explicit information be provided in the IP header than is currently the case.

## Acknowledgments

We would like to thank Keith Sklower for helping us port the RLP code to UNIX and for lots of fruitful discussions. We thank Mikael Degermark for deepening our insights into IP header compression. We thank Tom Henderson, Michael Meyer, Bela Rathonyi, and the anonymous reviewers for comments on earlier versions of this paper.



## References

- [1] Allman, M., Paxson, V., Stevens, W., *TCP Congestion Control*, RFC 2581, April 1999.
- [2] Bakre, A., Badrinath, B. R., *I-TCP: Indirect TCP for Mobile Hosts*, In Proceedings of ICDCS 95, May 1995.
- [3] Balakrishnan, H., Seshan, S., Katz, R. H., Improving reliable transport and handoff performance in cellular wireless networks, *Wireless Networks*, Vol. 1, No. 4, pp. 469-481, December 1995.
- [4] Balakrishnan, H., Katz, R. H., *Explicit Loss Notification and Wireless Web Performance*, In Proceedings of IEEE GLOBECOM 98.
- [5] Bhagwat, P., Bhattacharya, P., Krishna, A., Tripathi, S. K., *Using channel state dependent packet scheduling to improve TCP throughput over wireless LANs*, *Wireless Networks*, Vol. 3, No. 1, pp. 91-102, 1997.
- [6] Blake, S. et al., *An Architecture for Differentiated Services*, RFC 2475, December 1998.
- [7] Braden, R., Requirements for Internet Hosts - Communication Layers, RFC 1122, October 1989.
- [8] Braden B., et al., *Recommendations on Queue Management and Congestion Avoidance in the Internet*, RFC 2309, April 1998.
- [9] Brown, K., Singh, S., *M-TCP: TCP for Mobile Cellular Networks*, *Computer Communications Review*, 27(5), October 1997.
- [10] Cáceres, R., Iftode, L., *Improving the Performance of Reliable Transport Protocols in Mobile Computing Environments*, *IEEE JSAC*, Vol. 13, No. 5, pp. 850-857, June 1995.
- [11] Deering, S., Hinden, R., *Internet Protocol, Version 6 (IPv6) Specification*, RFC 2460, December 1998.
- [12] Degermark, M., Nordgren, B., Pink, S., *IP Header Compression*, RFC 2507, February 1999.
- [13] Durst, R. C., Miller, G. J., Travis, E. J., *TCP Extensions for Space Communications*, In Proceedings of ACM MOBICOM 96.
- [14] Eckhardt, D. A., Steenkiste, P., *Improving Wireless LAN Performance via Adaptive Local Error Control*, In Proceedings of IEEE ICNP 98.
- [15] ETSI, *Radio Link Protocol for data and telematic services on the Mobile Station - Base Station System (MS-BSS) interface and the Base Station System - Mobile Switching Center (BSS-MSC) interface*, GSM Specification 04.22, Version 5.0.0, December 1995.
- [16] ETSI, *Digital cellular communications system (Phase 2+); Radio Link Protocol for data and telematic services on the Mobile Station - Base Station System (MS-BSS) interface and the Base Station System - Mobile Switching Center (BSS-MSC) interface*, GSM Specification 04.22, Version 6.1.0, November 1998.
- [17] ETSI, *Digital cellular communications system (Phase 2+); Rate adaption on the Mobile Station - Base Station System (MS - BSS) Interface*, GSM Specification 04.21, Version 7.0.0, October 1998.
- [18] ETSI, *Digital cellular communications system (Phase 2+); General Packet Radio Service (GPRS); Mobile Station (MS) Base Station System (BSS) interface; Radio Link Control / Medium Access Control (RLC/MAC) protocol*, GSM Specification 04.60, Version 6.1.0, August 1998.
- [19] ETSI, *Digital cellular communications system (GSM Radio Access Phase 3); Channel coding*, GSM Specification 05.03, Version 6.0.0, January 1998.
- [20] Floyd, S., Jacobson, V., Liu, C., McCanne, S., and Zhang, L., *A Reliable Multicast Framework for Light-weight Sessions and Application Level Framing*, *IEEE/ACM Transactions on Networking*, December 1997, Volume 5, Number 6, pp. 784-803.
- [21] Henderson, T. R., Katz, R. H., *Transport Protocols for Internet-Compatible Satellite Networks*, *IEEE JSAC*, Vol. 17, No. 2, pp. 326-344, February 1999.
- [22] Jacobson, V., *Congestion Avoidance and Control*, In Proceedings of ACM SIGCOMM 88.
- [23] Jacobson, V., *Compressing TCP/IP Headers for Low-Speed Serial Links*, RFC 1144, February 1990.
- [24] Jacobson, V., Braden, R., Borman, D., *TCP Extensions for High Performance*, RFC 1323, May 1992.
- [25] Karn, P., Partridge, C., *Improving Round-Trip Time Estimates in Reliable Transport Protocols*, In Proceedings of ACM SIGCOMM 87.
- [26] Karn, P., *The Qualcomm CDMA Digital Cellular System*, In Proceedings of the USENIX Mobile and Location-Independent Computing Symposium, USENIX Association, August 1993.
- [27] Kent, S., Atkinson, R., *Security Architecture for the Internet Protocol*, RFC 2401, November 1998.
- [28] Kojo, M., Raatikainen, K., Liljeberg, M., Kiiskinen, J., Alanko, T., *An Efficient Transport Service for Slow Wireless Telephone Links*, *IEEE JSAC*, Vol. 15, No. 7, pp. 1337-1348, September 1997.
- [29] Lettieri, P., Srivastava, M. B., *Adaptive Frame Length Control for Improving Wireless Link Throughput, Range, and Energy Efficiency*, In Proceedings of IEEE INFOCOM 98.
- [30] Ludwig, R., Rathonyi, B., *Link Layer Enhancements for TCP/IP over GSM*, In Proceedings of IEEE INFOCOM 99.
- [31] Ludwig, R., Rathonyi, B., Konrad, A., Oden, K., Joseph, A., *Multi-Layer Tracing of TCP over a Reliable Wireless Link*, In Proceedings of ACM SIGMETRICS 99.
- [32] Ludwig, R., *A Case for Flow-Adaptive Wireless Links*, Technical Report UCB//CSD-99-1053, University of California at Berkeley, May 1999.
- [33] Mathis, M., Mahdavi, J., Floyd, S., Romanow A., *TCP Selective Acknowledgement Options*, RFC 2018, October 1996.
- [34] Meyer, M., *TCP Performance over GPRS*, In Proceedings of IEEE WCNC 99.
- [35] Mouly, M., Pautet, M.-B., *The GSM System for Mobile Communications*, Cell & Sys, France 1992.
- [36] Noble, B. D., Satyanarayanan, M., Nguyen, G. T., Katz, R. H., *Trace-Based Mobile Network Emulation*, In Proceedings of ACM SIGCOMM 97.
- [37] Postel, J., User Datagram Protocol, RFC 768, August 1980.
- [38] Postel, J., Internet Protocol, RFC 791, September 1981.
- [39] Postel, J., Transmission Control Protocol, RFC793, September 1981.
- [40] Ramakrishnan, K. K., Floyd, S., *A Proposal to add Explicit Congestion Notification (ECN) to IP*, Internet Draft, Work in progress, January 1999.

- [41] Saltzer, J. H., Reed, D. P., Clark, D. D., *End-To-End Arguments in System Design*, ACM Transactions on Computer Systems, Vol. 2, No. 4, November 1984.
- [42] Samaraweera, N. K. G., Fairhurst, G., *Reinforcement of TCP Error Recovery for Wireless Communication*, Computer Communications Review, 28(2), April 1998.
- [43] Simpson, W., *The Point-to-Point Protocol*, RFC 1661, July 1994.
- [44] Srinivasan, R., *RPC: Remote Procedure Call Protocol Specification Version 2*, RFC 1831, August 1995.
- [45] Stevens, W. R., *TCP/IP Illustrated, Volume 1 (The Protocols)*, Addison Wesley, November 1994.
- [46] The Mailing List of the PILC Working Group (Performance Implications of Link Characteristics) of the Internet Engineering Task Force, <http://pilc.grc.nasa.gov/pilc/list/archive/>
- [47] Walsh, D., et al., *The overview of the sun network file system*, In Proceedings of Winter USENIX, January 1985.
- [48] Wright G. R., Stevens, W. R., *TCP/IP Illustrated, Volume 2 (The Implementation)*, Addison Wesley, January 1995.