

# Multi-Layer Protocol Tracing in a GPRS Network

Andrei Gurtov, Matti Passoja, Olli Aalto, Mika Raitola  
 Cellular Systems Development  
 Sonera  
 Helsinki, Finland

**Abstract**—This paper presents a performance evaluation of GPRS accomplished by combination of measurement at the end hosts and tracing inside the network. The multi-layer tracing approach allows not only observing, but also understanding the network performance. With end-to-end measurements we assess data rates, latency, and buffering experienced by users in a live GPRS network. Comparing the results to our previous measurements shows a notable improvement in the network and terminals over past two years. Mobility tests while driving in the urban environment quantify the interval, duration and data loss caused by cell reselections. In the test lab, multi-layer tracing of radio, link and transport protocols gives a closer picture of GPRS performance. For instance, TCP interacts inefficiently with resource allocation at the RLC layer and fragmentation at the LLC layer. Finally, we illustrate delay spikes and data losses during a cell reselection by tracing of signaling messages during a cell update and routing area update procedures.

**Keywords**—performance, wireless, link layer, TCP, protocol

## I. INTRODUCTION

General Packet Radio Service (GPRS) [1] is a packet-switched wireless wide area network being deployed worldwide. Performance evaluation of GPRS is an active research area. In particular, TCP performance [3], buffering [4], scheduling [9] and mobility procedures have been studied through analytical analysis and simulation. We present a performance study based on measurement data collected in live and test GPRS networks [11]. Some of the issues presented in this paper are affected by implementation details and the standardization status of the available network and terminal equipment. Therefore, we do not claim that the results apply for *all* GPRS users in general.

In the first part of the paper we evaluate performance of data transmission in live GPRS network from the end user point of view. Uplink and downlink throughput, round trip time, buffer size, delay spikes and data losses are characterized

by end-to-end measurements. Comparing with our earlier results [2], data rates, latency and reliability are notably improved. For instance, the maximum downlink user throughput increased from 27 kbps to 43 kbps. As expected, we have not detected error losses at the transport layer due to the reliable link layer protocol in GPRS. Instead, we find delay spikes and bursty losses due to mobility procedures.

Sufficient buffering is crucial to achieve efficient multiplexing of bursty user traffic over the radio link. However, our buffer measurement suggests that a GPRS network may overbuffer user data. Too large buffers cause negative effects such as high round trip time and delivery of stale data when the user “clicks from page to page” [6]. From the TCP point of view, the optimal buffer size should be slightly above the bandwidth-delay product of the network [10].

A powerful multi-layer tracing methodology is introduced in [6] to study the GSM data transmission. We perform multi-layer tracing in a GPRS test network. We found that if the frame size at the logical link layer is not configured to match with a typical IP packet size, small fragments are inefficiently transmitted. Furthermore, we confirm that the radio resource allocation can interact inefficiently with TCP [7] and, on the other hand, that competing error recovery between radio link and TCP is uncommon [3]. At the radio layer, we observed unnecessary retransmissions due to trade-offs in the acknowledgment and retransmission policy.

During mobility tests we measure frequency, duration and data loss of cell reselections in GPRS. In the test lab we trace the signaling procedures and end-to-end TCP behavior. In the test lab the duration of cell reselections is typically below 5 s, there as in the live network it may take up to twenty seconds.

The rest of the paper is organized as follows. In Section 2 the relevant aspects of data transmission over GPRS are presented. Section 3 describes the configuration of our GPRS




Figure 1. Buffering of user data in GPRS.




Figure 2. The cell update procedure in GPRS.




Figure 3. Measurement setup in the test lab.

Vendor	Model	Timeslots (UL+DL)	Multislot class	Interface
Ericsson	T39m	1+3	4	Serial
Ericsson	R520m	1+3	4	Serial
Nokia	N8310	2+3	6	IrDA
Motorola	T280	1+4	8	USB

Figure 4. Mobile terminals used for measurements.

network and setup of measurement. Section 4 reports end-to-end and tracing results. Finally, Section 5 sums up the main findings and outlines the future work.

## II. THE GPRS NETWORK

Figure 1 and 3 illustrate the data transmission path of GPRS. The relevant network elements for us are the Mobile Station (MS), Base Transceiver Station (BTS), Base Station Controller (BSC), Serving GPRS Support Node (SGSN) and Gateway GPRS Support Node (GGSN). BSC handles the medium access and radio resource scheduling, as well as data transmission toward MS over the Abis interface. SGSN handles mobility and controls the data flow toward BSC over the Gb interface. GGSN provides connectivity to external packet networks. A firewall shields the GPRS network from the rest of the Internet. A detailed overview of the GPRS system can be found in [1].

### A. The GPRS Protocol Stack

Figure 1 shows the protocol stack of the user data transmission plane of GPRS. The Radio Link Control (RLC) protocol provides acknowledged or unacknowledged data transfer between MS and BSC in uplink (UL) and downlink (DL) directions. The Logical Link Control (LLC) protocol provides acknowledged and unacknowledged mode between MS and SGSN. The Base Station Subsystem GPRS (BSSGP) protocol controls the data flow between BSC and SGSN. Finally, GPRS Tunneling Protocol (GTP) encapsulates user packets for delivery between SGSN and GGSN.

Medium Access Control (MAC) manages sharing of radio resources among multiple users. MS can utilize several radio timeslots simultaneously to increase the data rate and decrease the transmission latency. The multislot class of MS determines the maximum number of timeslots in uplink and downlink. Before transmitting user data, MS must activate a Temporal Block Flow (TBF) toward BSC. MSs contend on an ALOHA-style random access channel to receive a resource allocation from the network. Optionally, a second stage is used for extending the assignment if MS is not satisfied with allocated resources.

RLC operates on small (20 to 50 bytes) blocks of user data that are encoded with one of a coding schemes (CS-1 to CS-4 for basic GPRS) to provide Forward Error Correction (FEC). RLC uses wrapping sequence numbers for blocks in the range of 0-127 and a sliding window of 64 blocks. In an acknowledged mode, a bitmap of received block is used to

retransmit missing blocks. In contrary to TCP, RLC ACKs are sent on a separate control channel and cannot be piggybacked onto the reverse data traffic. BSC controls the acknowledgment frequency by sending ACKs in downlink or polling MS for ACKs in uplink. As we will see in Section IV, frequency of ACKs and retransmission policy at the sender are important to avoid unnecessary retransmissions in RLC. A situation when too many outstanding unacknowledged blocks prevent advancing of the sliding window should be avoided. Then, the window is *stalled*, and no new data blocks can be transmitted.

LLC provides a retransmission capability between MS and SGSN, and is supposed to recover losses caused by mobility. However, most GPRS networks nowadays operate in the unacknowledged LLC mode. LLC fragments and reassembles user packets if they exceed the maximum size. It can be configured up to 1556 bytes.

### B. Buffering of User Data in GPRS

Figure 1 illustrates buffering of user data in GPRS. Buffering is performed at multiple protocol layers, but a corresponding buffer is used only if the protocol operates in the acknowledged mode. In our measurements, reliable RLC and unreliable LLC modes are used, thus the only enabled buffers are at the RLC and BSSGP layer. In downlink, LLC frames are stored in the BSSGP buffer in SGSN prior to transmission over the Gb interface. Although the buffer is located in SGSN, it is controlled by the BSSGP function in BSC. This enables BSC to adjust the data flow rate from SGSN in order to match it with available radio resources to prevent an overflow of the RLC buffer. Therefore, BSSGP buffer can be seen as an extension of the RLC buffer. The RLC buffer size is 64-128 RLC blocks or up to 6 kilobytes of the user data. Using multiple timeslots the content of the RLC buffer can be transmitted in a few hundreds ms. Therefore, multiple retransmissions can easily stall the window. This problem is corrected in Enhanced GPRS where the RLC buffer size can be 64 – 1024 blocks [5].

### C. Cell Reselection

In the release 97 GPRS the mobile terminal selects the serving cell. This is different from circuit-switched GSM data where a handover is controlled by the network. In the simplest case when the user changes the serving cell while staying in the same routing area, a cell update procedure is performed. A *routing area* is a group of cell arranged together to balance between signaling overhead and positioning of MS. When the

new cell belongs to a different routing area, the cell reselection involves more signaling, especially if GSM-specific location information is updated as well. Finally, the most complicated case concerns an inter-SGSN handover. However, it is expected to be a rare event and therefore we have not measured it.

Figure 2 shows signaling required to accomplish a cell update. First, MS makes a cell reselection decision based on tracking signal power of surrounding cells. After synchronizing at the frequency in the new cell, MS starts a random access procedure to acquire radio resources. Then, MS starts transmitting data in the new cell. When SGSN receives an LLC frame with a new cell identity, it internally updates the MS location. Finally, SGSN signals the old cell using BSSGP protocol to release any resource reservations for MS and discard buffered data.

Shortcomings of the mobile-driven cell reselection are widely recognized and improvements are being standardized in 3GPP. A Network Controlled Cell Change (NCCC) will make the BSC responsible for a cell change decision. NCCC will eliminate unnecessary cell reselections currently performed by stationary MSs. A Network Assisted Cell Change (NACC) should reduce the delay and data losses seen by a moving MS.

### III. MEASUREMENT ARRANGEMENTS

At the time of measurements, Sonera's GPRS network is implementing the 3GPP release 97. Figure 4 lists mobile terminals used for measurements. The GPRS network was able to support the maximum number of time slots defined by the terminal multislot class. All terminals are forced to use CS-2 encoding, as it provides better throughput with only a small loss over CS-1 in error recovery. Usage of CS-3 and 4 is currently not possible due to capacity limitations at the Gb interface. The network uses unacknowledged LLC. The Van Jacobson header compression is disabled due to poor performance in presence of packet losses, high computing burden on the network and lack of support from terminals.

For end-to-end throughput measurements, we use a tool generating bulk transfers over TCP. For measuring latency, we use a standard ping program. The NetHawk tracing tool records data traffic and signaling messages at the Abis or Gb interface as shown in Figure 3. Using an engineering mode available in some terminals it is possible to see an identifier of the serving cell, as well as force a cell reselection to one of surrounding cells. Finally, we use tcpdump to record TCP traces at the end hosts.

### IV. MEASUREMENT RESULTS

#### A. Throughput, Latency and Buffering

Figure 5 and 6 show downlink and uplink throughput measured with four different terminals. In addition to minimum, average and maximum throughput observed over 40 replications, graphs also show the line rate computed based on the available number of timeslots and coding scheme, and the

maximum TCP throughput taking into account TCP/IP header overhead. The multislot class of the terminal chiefly determined the throughput; the data rate per slot was approximately the same for all terminals. In downlink the maximum measured value is 43 kbps achieved by T280 terminal using four timeslots. In uplink, the maximum value is 21 kbps by N8310 using two timeslots. This is a notable improvement over earlier measurements [2]; at that time the maximum downlink throughput was 20 kbps and only 7 kbps in uplink.

In general, setting a larger IP MTU at the end hosts resulted in higher throughput. For instance, increasing the MTU from 576 bytes up to 1480 bytes improves throughput by one percent due to reduced TCP/IP header overhead. However, the MTU of 1500 bytes gave slightly lower throughput than of 1480 bytes. Tracing at the Gb interface showed inefficient fragmentation at the LLC layer. The maximum LLC frame size was configured to 500 bytes, thus 1500-byte IP packets were fragmented into four frames, with the fourth frame only a few bytes long. Sending plenty of small frames reduces efficiency due to higher header overhead. In the acknowledged LLC mode, using a smaller frame size than 1556 bytes can be beneficial by achieving finer grain retransmissions. However, for the unacknowledged mode we do not see a compelling reason to reduce the maximum LLC size.

We measured RTT of a GPRS link using 32-byte pings. RTT varies depending on the terminal and a serving cell in the range of 500-1100 ms on the unloaded link. A typical value is 700 ms. The minimum RTT improved since our earlier measurements [2] approximately by 200 ms. Interestingly, we observed regular oscillations in RTT when every second ping gets roughly 100 ms higher RTT than the other. This effect seems to relate to radio resource allocation as explained in the next section.

To estimate the buffer size of the GPRS link, we started a bulk TCP transfer with a sufficiently large window (200KB) to overflow the bottleneck buffer. The amount of outstanding data when the first loss occurs reflects the size of a drop-tail buffer. Based on specifications, we expected to see an approximately 10 kilobytes buffer. However, measurements indicate the downlink buffer of 50 kilobytes. Apparently, GPRS implementations include additional backlog buffers not present in standards. The per-user buffer size for GPRS downlink is optimally 5-10 packets since the bandwidth-delay product of a GPRS link does not exceed 5 kilobytes.

The uplink buffer measurement indicated the buffer size in terminals in the range of 3 to 30 kilobytes. A terminal (not listed in Figure 4) having only a 3-kilobyte buffer showed throughput of one-third of other terminals for an uplink bulk transfer. The reason was in repeating TCP retransmission timeouts. TCP needs at least three buffers in the network to utilize the fast retransmit algorithm [10]. On the other hand, the buffer of 30 kilobyte is excessive as it allows for unacceptably high link round-trip time and unnecessary delivery of data from aborted TCP connections.




Figure 5. Bulk TCP throughput in GPRS downlink.




Figure 6. Bulk TCP throughput in GPRS uplink.

### B. Tracing at the RLC Layer

Tracing at the Abis interface illustrates several interesting details on functioning of the RLC protocol. The first problem is allocation and release of TBF. According to release 97 specifications, TBF should be turned down immediately when the data buffer empties. Such a policy increases TCP RTT since every segment and ACK may trigger setup of a new TBF. Keeping TBF for longer periods has been suggested [7] and is reflected in Enhanced GPRS specifications [5]. The extended TBF release decreases RTT seen by TCP by more than a hundred milliseconds and reduces the signaling load. However, BSC is unaware if MS with an active TBF has any data to transmit and has to schedule also idle MSs, which wastes radio resources. Furthermore, the number of simultaneous TBFs is limited and postponing the TBF release can prevent data transmission by other MSs. Figure 7 shows an RLC trace of a downlink TCP transfer. In uplink, TCP ACKs are sent in groups of two on separate TBFs. The graph also shows an increase in TCP RTT caused by signaling to set up TBF for every TCP ACK in uplink. The network transmits dummy RLC blocks downlink to keep up TBF for instance at 11.5-12 s.

Another interesting case is premature retransmissions when there are unacknowledged blocks at the RLC layer but no new blocks to transmit. The RLC sender retransmits unacknowledged blocks in round robin until an ACK is received [8]. It can be seen in Figure 7 for instance at 10.7 s. On one hand, it increases the probability of data blocks to get through the radio link. On the other side, it may waste radio resources and battery power of MS. However, MS has no knowledge whether there is new data in BSC to be sent and therefore has to decode assigned timeslots anyway. It consumes the battery power as well. Avoiding such retransmissions when other users have data to transmit would prevent waste of radio resources.

At times the RLC sender retransmits lost blocks several times unnecessarily. It happens when the RLC receiver generates several ACKs before the first retransmission has arrived to it. Such ACKs will indicate the same lost segments and therefore can trigger unnecessary retransmissions at the RLC sender. A timer at the RLC receiver could prevent repeated retransmission of blocks approximately for one RLC RTT [8]. This gives enough time for the first retransmission to

arrive and be acknowledged. Alternatively, the receiver can generate ACKs less frequently.

According to specifications, BSC schedules ACKs in uplink and downlink “when needed” [5]. For instance, BSC in Figure 7 requests an ACK for every tenth block. Less frequent ACKs preserve radio resources and battery power, but increase probability of stalling the window. We suggest that BSC during uplink transfers sends a selective ACK immediately or shortly after a missing block is detected. On the other hand, when all blocks are received correctly, infrequent ACKs suffice. In downlink transfers, BSC can poll MS for ACKs more frequently when the link quality drops down.

### C. Mobility Measurements

We measured frequency, length and data loss of cell reselections in a live GPRS network. In the test network we recorded multi-layer traces of cell reselections.

Tests in the live network were performed while driving in downtown Helsinki. Cell reselections occurred at irregular intervals on the average every 40-70 s. The interval depends on the route and speed, but cell reselections occurring even in stationary conditions are not uncommon. Cell reselections suspended the data transfer by 3 to 15 s with most cases below 5 s. There were a few exceptions when a failed cell reselection made the link unusable for two minutes. By examining receiver TCP traces it is possible to calculate the number of lost segments during cell reselections. In downlink direction a large number and sometimes all outstanding packets were lost. However, we also observed cases where data segments were not lost but just delayed. In uplink, data loss was less common.

In the test network we traced LLC and TCP layers during cell reselections with and without a routing area update. At the LLC layer cell reselections cause a delay spike in uplink and downlink transfers of 2-4 s with cell update procedure only. The delay increases to 4-5 s when a routing area update is also triggered. We recorded signaling messages when a routing area update was performed between BSCs from two different vendors. The message exchange lasted typically less than two seconds.

During cell reselections about ten TCP segments were lost in downlink and none or one segment in uplink. The difference is due to the fact that in uplink transfers buffered data can be easily sent in the new cell there as in downlink direction data




Figure 7. Interactions of TCP with channel allocation at the RLC layer.

need to be transferred to a new cell, which is rarely done. Therefore, in downlink cell reselections are seen by TCP as loss bursts that can cause lengthy timeouts and underutilization of the radio link in a new cell. On the contrary, for uplink transfers cell reselections are seen by TCP as delay spikes that can cause spurious timeouts [6]. Figure 8 illustrates two cell reselections during a downlink transfer. At the LLC layer a visible break due to cell reselection is approximately 5 s. However, it takes 5 to 10 s. more for TCP to retransmit lost segments.

We found three implementation problems in the mechanism of cell reselection. First, one or two LLC frames sent in the old cell before a cell reselection were unnecessary retransmitted in the new cell due to a bug in MS. Second, SGSN sent the FLUSH-LL message to the *new* cell instead of the old cell causing waste of resources in the old cell and loss of data in the new cell. Third, dummy LLC frames caused an outage in the LLC data transmission. Dummy frames are required for completing the cell reselection in case when no data is available for transmission in MS.

## V. CONCLUSIONS AND FUTURE WORK

We measured performance of GPRS in stationary and mobile operation. The maximum downlink TCP throughput was 43 kbps and 21 kbps in uplink. The typical RTT of the unloaded link is around 0.7 s. We have estimated a 50 kilobytes downlink buffer available for a single GPRS user. It exceeds the optimal value by several times and allows for undesired effects such as inflated RTT and delivery of stale data. On the other hand, one terminal had only a two-packet buffer in the uplink direction and showed throughput of one-third of the normal. TCP needs at least three buffers per connection for efficient loss recovery.

By combining end-to-end tracing with tracing performed within the network, we observed several undesired cross-layer interactions between RLC, LLC and TCP. In particular, the slow start phase of TCP and delayed ACKs interact badly with radio resource allocation at the RLC layer. At the LLC layer, mismatch between the maximum size of data units in LLC and in TCP results into inefficient fragmentation. Finally, we showed situations when RLC retransmits data unnecessarily.

While driving in an urban area we observed cell reselections to occur roughly every minute and to last for five




Figure 8. Effect of cell reselections on TCP and LLC in a downlink transfer.

seconds. In downlink, most of outstanding data gets lost during a cell reselection. It takes 5 to 10 s. for TCP to recover lost data. In the uplink, cell reselections often do not cause data loss, but instead are seen as a delay spike by the upper layers. In this case, TCP can experience a spurious timeout and retransmit the outstanding data unnecessarily [12].

Our future work on GPRS will include testing network-controlled cell reselections, measuring throughput and battery consumption under varying radio conditions and network load. We also plan to evaluate performance of real-time streaming applications and quality of service mechanisms.

## ACKNOWLEDGEMENTS

We thank Roger Kalden, Michael Meyer, Janne Peisa, Hannes Ekström and Reiner Ludwig for useful comments that improved the paper significantly.

## REFERENCES

- [1] G. Brasche and B. Walke. Concepts, services and protocols of the new GSM phase 2+ general packet radio service. IEEE Communications Magazine, pages 94--104, August 1997.
- [2] J. Korhonen, O. Aalto, A. Gurtov, H. Laamanen, Measured Performance of GSM HSCSD and GPRS, IEEE ICC, June 2001.
- [3] M. Meyer, TCP Performance over GPRS, IEEE WCNC, September 1999.
- [4] Ho, J.; Zhu, Y.; Madhavapeddy, S., Throughput and buffer analysis for GSM General Packet Radio Service (GPRS), IEEE WCNC, vol. 3, 1999.
- [5] 3GPP TS 04.60 V8.6.0 Mobile Station (MS) - Base Station System (BSS) interface; Radio Link Control/ Medium Access Control (RLC/MAC) protocol (2000). Release 99.
- [6] R. Ludwig, B. Rathonyi, A. Konrad, K. Oden, and A. Joseph. Multi-layer tracing of TCP over a reliable wireless link. In Proceedings of the ACM SIGMETRICS, May 1999.
- [7] H. Wiemann, A. Schieder, H. Ekström, Enhanced TBF Features in GERAN, WPMC'01, Aalborg, Denmark, Septemeber 2001.
- [8] B. Walke, Mobile Radio Networks, Networking and Protocols (2. Ed.), Wiley & Sons, Chichester 2001.
- [9] J. Sau, C. Scholefield, Scheduling and quality of service in the General Packet Radio Service, IEEE ICUPC, 1998.
- [10] W. Stevens, TCP/IP Illustrated, Volume 1; The Protocols, Addison Wesley, 1995.
- [11] M. Passoja, Effects of Cell Reselection to the Performance of GPRS, Diploma Thesis, University of Oulu, September 2001.
- [12] A. Gurtov, R. Ludwig, Evaluating the Eifel Algorithm for TCP in a GPRS network, In Proceedings of European Wireless, Florence, Italy, February 2002.