Impact of GPRS buffering on TCP performance

A. Wennström, A. Brunstrom and J. Rendón

In GPRS networks, excessive buffering has a negative effect on TCP as the round trip times become very long. Measurements with different buffer settings indicate that the queueing delay can be reduced by orders of magnitude with a smaller buffer, without significantly degrading TCP throughput. The measurements are conducted in a GPRS testbed consisting of real network nodes.

Introduction: The General Packet Radio Service (GPRS) is a packet-oriented extension to the Global System for Mobile Communications (GSM) that supplies higher data rates and sharing of radio resources. The performance of the Internet protocols in GPRS is dependent not only on the quality of the radio link, but also on buffering in the GPRS network. According to [1–3], the over-sized network buffers used in many commercial networks may result in very long round trip times (RTTs) for the Transmission Control Protocol (TCP) when the buffers in the cellular network are filled. Suggested remedies [1, 3] typically require infrastructure changes.

In the work reported in this Letter we experimentally investigated the impact of GPRS buffering on TCP. In contrast to related work, [1–3], we experimented with different buffer settings and evaluated the effect on TCP behaviour. Our work is based on measurements in a GPRS testbed consisting of real network nodes with real protocol implementations and an emulated radio environment.

Experimental setup: The GPRS testbed used for the measurements consists of a laptop connected to a GPRS terminal, a GPRS network and a server. The GPRS terminal accesses the GPRS network over an emulated radio interface, which is connected to a base transceiver station (BTS) and a base station controller (BSC) which in turn is interconnected to a serving GPRS support node (SGSN) and a gateway GPRS support node (GGSN). The radio resources are more efficiently used than in GSM, since the GPRS users share packet data channels (PDCHs) that correspond to GSM time slots.

The buffer management in the GPRS network is co-ordinated between the SGSN and the BSC. The rate of the data transmitted from the SGSN to the BSC in the downlink is controlled by the base station subsystem GPRS Protocol (BSSGP). The BSSGP buffers in the SGSN and in the BSC may be considered as one logical BSSGP buffer. After a maximum time determined by the BSSGP buffer setting, data is discarded from the SGSN or the BSC buffer, depending on where data is resided. In the presented experiments, the end hosts use TCP to recover discarded data. During measurements, data is captured at the end hosts and at the Abis interface, between the BSC and the BTS.

Table 1: Parameter settings

Parameter	Value
Operating system	Linux 2.4.2-2
TCP options	SACK and timestamps
Max rwnd	64 KB
MSS	1460 bytes
PDCHs	1, 2, 3
RLC/MAC	acknowledged mode
LLC	unacknowledged mode
Coding scheme	CS-2
Signal level	−80 dBm
C/I	20 dB

The measurements are conducted for two settings of the BSSGP buffer, 63 and 5 s. A BSSGP buffer setting of 63 s is commonly used in commercial networks, and 5 s is chosen because this is a timeout value used for other timers in GPRS, e.g. on the radio link control/medium access control (RLC/MAC) layer. In each measurement data is transmitted over one TCP connection for 15 min. Further details on the parameters used in the measurements are provided in Table 1. The TCP options used follow the recommendations in [4]. The settings of the coding scheme, the RLC/MAC and the logical link control (LLC) correspond to settings commonly used in commercial networks.

Results: The results depend both on the buffer setting and on the number of PDCHs used. As illustrated by Fig. 1, TCP throughput remains almost the same for both buffer settings, but, as expected, it increases with an increasing number of PDCHs, from around 10 kbit/s for one PDCH to around 30 kbit/s for three PDCHs. The maximum RTT, on the other hand, is strongly related to the buffer setting, which is shown in Fig. 2. A BSSGP buffer of 5 s reduces the maximum RTT with 80%, from 48 to 9 s, for one PDCH and with 50%, from about 17 to 8 s, for three PDCHs.

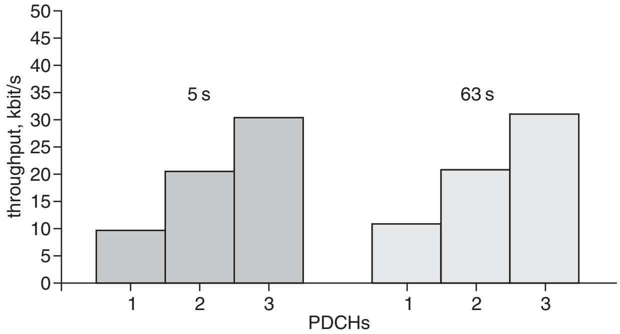

Fig. 1 *BSSGP buffer setting and PDCHs against throughput*

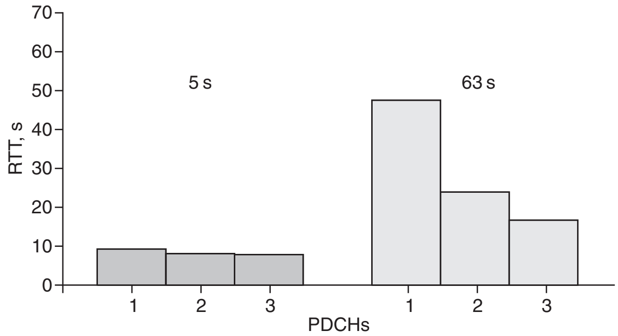

Fig. 2 *BSSGP buffer setting and PDCHs against maximum RTT*

The 63 s buffer results in an extremely long RTT because of the excessive buffering in the BSSGP buffer, which is illustrated in the RTT graph in Fig. 3. On the x-axis, time is indicated in minutes and seconds relative to the first segment. The RTT values are indicated in milliseconds on the y-axis. The RTT starts at 1 s and then increases up to 48 s, as more packets are injected into the network. After four min, TCP reduces its transmission rate due to a sporadic data loss, which gives the BSSGP buffer time to empty. The RTT is reduced to 1 s before it starts to increase again.

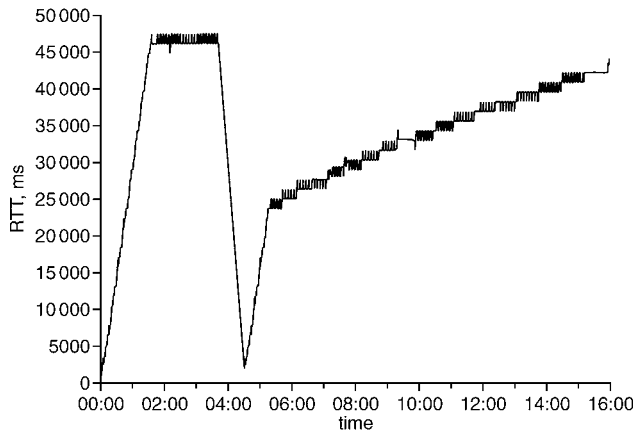

Fig. 3 *Round trip times over one PDCH with 63 s BSSGP buffer*

The 5 s buffer gives a shorter RTT, but at the cost of data loss due to BSSGP buffer overflows that occur at regular intervals. The RTT fluctuates between 2 and 9 s, as shown in Fig. 4. As the TCP sender increases its transmission rate, more data is buffered in the BSSGP buffer, and eventually data is discarded because it has been stored for the maximum time of 5 s. The TCP sender reduces its transmission rate in response to the data loss. The BSSGP buffer empties, and then the same sequence of events is repeated all over again. However, the retransmitted TCP segments do not waste any radio resources, since the lost data is discarded from the BSSGP buffer before it is transmitted over the radio link. Even though frequent data losses occur, the radio link is fully utilised most of the time, which explains the limited impact on TCP throughput.

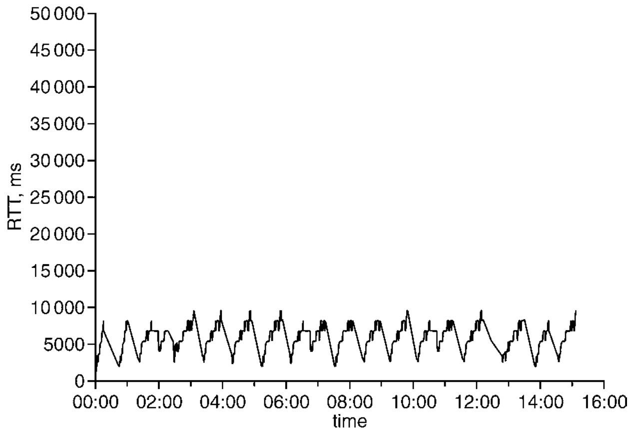

Fig. 4 *Round trip times over one PDCH with 5 s BSSGP buffer*

The long RTT caused by a BSSGP buffer setting of 63 s is not a problem for the throughput of the considered bulk transfer. From a user's perspective, on the other hand, a long RTT may be undesirable. If a user chooses to interrupt a transmission, e.g. by pressing the stop button in a web browser, then data buffered in the BSSGP buffer must first be transmitted before the user can receive any new data. Excessive buffering and long RTTs may also cause problems for interactive traffic if data is transmitted over both bulk and interactive connections simultaneously.

Conclusion: The impact of GPRS buffering on TCP is evaluated through measurements. The results indicate that if the BSSGP buffer is reduced, then the delay can be decreased significantly with almost no negative effect on TCP throughput. By adjusting the buffer setting, the delay can thus be decreased without any modification of infrastructure or end-user equipment.

24 June 2004
Electronics Letters online no: 20045967
doi: 10.1049/el:20045967

A. Wennström and A. Brunstrom (*Department of Computer Science, Karlstad University, SE-651 88 Karlstad, Sweden*)

E-mail: annika.wennstrom@kau.se

J. Rendón (*Technology Department, Pompeu Fabra University, Passeig de Circumvalació 8, 08003 Barcelona, Spain*)

References

1 Chakravorty, R., Cartwright, J., and Pratt, I.: 'Practical experience with TCP over GPRS'. IEEE GLOBECOM, Taipei, Taiwan, November 2002
2 Gurtov, A., Passoja, M., Aalto, O., and Raitola, M.: 'Multi-layer protocol tracing in a GPRS network'. Proc IEEE Vehicular Technology Conf. (Fall VTC2002), Vancouver, Canada, September 2002
3 Sågfors, M., Ludwig, R., Meyer, M., and Peisa, J.: 'Queue management for TCP traffic over 3G links'. Wireless Communications and Networking, New Orleans, LA, USA, March 2003
4 Inamura, H., Montenegro, G., Ludwig, R., Gurtov, A., and Khafizov, F.: *RFC 3481: TCP* over Second (2.5G) and Third (3G) Generation Wireless Networks, February 2003