

Link Layer Enhancements for TCP/IP over GSM

Reiner Ludwig

Ericsson Eurolab Deutschland GmbH
Germany

Bela Rathonyi

Ericsson Mobile Communications AB
Sweden

Abstract – This paper has two main contributions. First the extremely high latency of the GSM link is revealed which through measurements is determined to have a magnitude usually only known from satellite links. This greatly impacts the link configuration time of packet framing protocols for serial links like the Point-to-Point Protocol (PPP). The key idea of the proposed solution - Quickstart-PPP - is to break the strict sequential order of the different signalling phases while ensuring that protocol semantics are not violated and interoperability with existing PPP implementations is preserved. As a result the link configuration delay can be completely eliminated. The second contribution uncovers a fundamental problem which is that the GSM circuit-switched data service is not capable of satisfying different reliability requirements simultaneously. More precisely it is not possible to use applications requiring the GSM link to operate in reliable mode together with loss-tolerant but delay-sensitive applications at the same time. The developed solution - Link Sniffer - suggests a mechanism that can be added to the implementation of GSM's reliable link layer protocol which allows to "sniff" on packet headers to determine the reliability mode to be used. The key advantage of this solution is that the link layer protocol itself does not have to be changed.

I. INTRODUCTION

Two major technologies are driving the information society of the late 90s: cellular telephony and the Internet. While both developments have taken place independent of each other in the past, manufacturers and operators of cellular networks are showing increasing interest in combining both technologies to provide wide-area cellular Internet access. Today it is already commonplace to see users "dial-in" to the Internet via wide-area cellular by connecting their laptops and palmtops to a mobile phone. In the near future, integrated devices will become available turning mobile phones running the TCP/IP stack or likewise palmtops equipped with cellular radios into regular Internet hosts. We believe that in the future a considerable fraction of the overall number of Internet hosts will be wireless devices.

In this paper we specifically look at GSM (Global System for Mobile communications), today undoubtedly the most successful digital cellular telephony system. While significant work is currently being put into providing higher bandwidths in GSM other challenges have yet not been addressed. Two problems that exist when using GSM as the access network to the Internet and that are not related to bandwidth are highlighted. We reveal the extremely high latency of the GSM link which through a series of measurements was determined to have a magnitude comparable to satellite links. This obviously

greatly impacts the performance of interactive traffic as will be shown in the particular case of the link configuration phase of the Point-to-Point Protocol (PPP). We suggest a solution, named *Quickstart-PPP*, which significantly reduces the delay for the PPP link configuration phase. Subsequently we uncover a fundamental problem that exists in GSM today when applications requiring reliable transmission and loss-tolerating but delay-sensitive applications are run on a mobile host at the same time. Addressing this problem we developed the *Link Sniffer* which makes the GSM link layer "flow-adaptive" and solves the problem without changing the underlying protocol itself.

The interested reader has several options when looking for literature on GSM in general. An excellent description including the various components that have been standardized for the so-called data bearer services can be found in [1]. However, those readers that are only interested in GSM's data capabilities, may have a hard time finding the right material. In Section II we therefore provide a detailed description of data capabilities provided GSM today and in the near future. In Section III we describe the methods used to determine the latency of the GSM link, report on the results and then develop the proposed Quickstart-PPP mechanism. The Link Sniffer mechanism is described in Section IV followed by the current implementation status presented in Section V. Future work is discussed in Section VI.

II. DATA SUPPORT IN GSM

Unlike with earlier analogue cellular systems, data services are an integral part of a GSM network and are equally supported together with ordinary voice services. This section outlines the architecture and mechanisms implemented in GSM for circuit-switched data as available today and also outlines the key features that will soon be deployed to better support circuit-switched Internet access via GSM. The descriptions given in this section pertain to all three "flavours" of GSM, namely GSM900, DCS1800, and PCS1900 as generally speaking these 3 systems are identical but running in different frequency bands.

Although this paper explicitly only deals with circuit-switched data the other data services supported in GSM today and in the future are mentioned here for completeness. The Short Message Service (SMS) is a 2-way paging mechanism and Unstructured Supplementary Service Data (USSD) is a service that allows text strings to be sent directly to the display of a GSM mobile phone. Both services are available in GSM

today but neither of these low-bandwidth services (< 600 bit/s) really qualify as bearers for TCP/IP in part because they do not provide a continuous duplex connection. Nevertheless, they can be used for Internet access through gateways supporting interactive applications which are based on the exchange of small amounts of data. Yet, probably the most powerful wireless Internet access technology to be available in the wide-area in almost every country in the world by the turn of the millennium is GPRS (General Packet Radio Service). GPRS is a packet-switched data service which has been standardized for GSM. Key advantages for GPRS users are that they are always on-line, can dynamically allocate bandwidth also in an asymmetric fashion on up- and downlink and pay per transmitted/received data volume. The benefits for the GSM operators offering GPRS are highly efficient and cost-effective use radio spectrum and network resources. More detail on GPRS can be found in [4].

A. Circuit-Switched Data Today

Figure 1 shows the basic components involved for circuit-switched data transmission in GSM. Network components required for network management and the databases required for mobility management, user profile handling, and equipment authentication are not shown for simplicity. In very broad terms one could say that a GSM network is an ISDN network with base stations added to it to provide a wireless interface. In reality it is of course by far more complex than that but the core transport network is the same: switches are interconnected hierarchically, network internal signalling is based on Signalling System No.7 and traffic trunks carry voice and data in multiples of 64 kb/s channels. A detailed description of the GSM system can be found in [1]. In the following quite a few acronyms taken from the ETSI standards will be introduced, however, we have tried to avoid them as much as possible throughout this paper.

A mobile host - a laptop or palmtop - is connected to the GSM network using a GSM mobile phone and a device (e.g. a PCMCIA card) running the required adaptation. In the ETSI standards the latter two are called Mobile Station (MS) and Terminal Adaptation Function (TAF), respectively. Note that unlike with first generation cellular systems the TAF is not a modem. The modem resides in the network; more precisely in the Interworking Function (IWF) of the Mobile Switching Centre (MSC). An MSC is a backbone telephone switch that routes circuits within the GSM network and also serves as a gateway to the fixed telephone network (ISDN or PSTN). The radio interface is provided by a Base Transceiver Station (BTS) which together with other BTSs is controlled by one Base Station Controller (BSC) which in turn is bundled with other BSCs by one MSC. The GSM radio interface itself is based on frequency division of up- and downlink and each carrier comprising a pair of one up- and one downlink frequency is time-shared among 8 users. At call-setup the network assigns a carrier number and a time-slot number to the mobile

phone to be used throughout the entire call. Up- and downlink are offset by 3 time-slots so that a mobile phone never has to transmit and receive at the same time which reduces its complexity considerably.

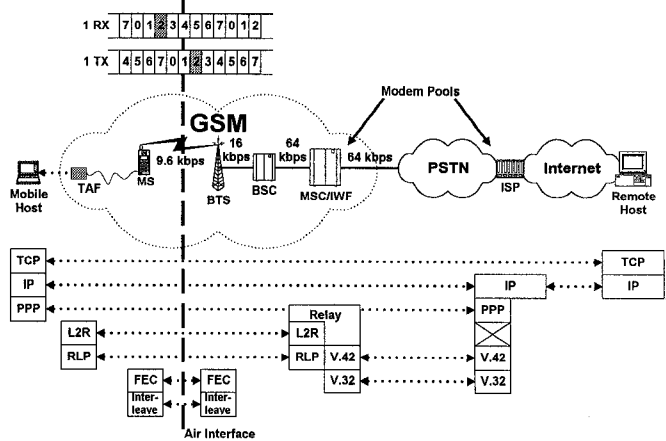


Fig. 1. TCP/IP over circuit-switched data in GSM.

The standards define an overwhelming set of different so-called circuit-switched data bearer services - data services for short. For each data service, a unique identifier is defined that the mobile phone has to signal to the network at call-setup using dedicated control channels in order for the network to allocate and configure the correct resources. Data rates between 300 and 9600 kb/s are specified which can either be transmitted asynchronous or synchronous. In addition the network has to distinguish whether data calls have to be switched into the PSTN or the ISDN and allocate either a modem or a rate adapter (Figure 1 shows the PSTN case). Also the mobile phone can request to either run a fully reliable link layer protocol called Radio Link Protocol (RLP) or not to do so. This is also referred to as operating the non-transparent or the transparent data service, respectively. Summing over all these possible combinations (sync./async., PSTN/ISDN and RLP/noRLP) a total of 8 different data services are specified just alone for the data rate of 9.6 kb/s. As the Internet access determines whether PSTN or ISDN will have to be used as transit network and assuming a user requests the highest possible data rate of 9.6 kb/s this boils down to 2 different choices for the user: with or without RLP. This will be further discussed in Section IV.

RLP is an HDLC-derived [5] protocol using selective-ARQ for which the frame size (30 bytes) has been optimized for the GSM radio link. Figure 1 shows the case where RLP is operated to provide a reliable radio link. The network side of RLP is terminated in the MSC/IWF which greatly simplifies link management in the face of handover. This way RLP remains terminated in the MSC/IWF serving at call-setup independent of which type of handover has to be performed while the circuit-switched connection is established. Even in the event of inter-MSC-handover no RLP state has to be transferred as GSM mobility management is based on the so-called anchor-MSC concept [1]. Thus, from the perspective of an ISP (Internet Service

Provider) the GSM link appears just like any other dial-up connection. Note that above the FEC (Forward Error Correction) Layer 12 kb/s are available but still a data rate of 1200 bytes/s is provided to the higher layers with or without RLP (!) in an ideal radio situation. This might seem surprising - and it is - because one would expect that the overhead introduced by RLP (20%) could be used for user data in the transparent case. However, it turns out that in the transparent case these 20% are wasted for an overkill of modem control signals. An additional protocol called the L2R (Layer 2 Relay) protocol is used in the non-transparent services for flow control, framing and communicating status control signals between the TAF and the IWF. Figure 1 shows the commonly used setup using the standard modem protocols V.42/V.32 towards the PSTN which in fact constitutes the bottleneck in this case as the V.42 overhead limits the data rate available for IP traffic to roughly 1150 bytes/sec¹. Given this architecture a mobile host then uses a standard serial link protocol like the Point-to-Point Protocol (PPP) [14] to connect to the Internet via GSM and a dial-in access fabric of an Inter-/Intranet Service Provider (ISP).




Fig. 2. Transmission of IP packets over GSM radio.

On the GSM radio link itself, channel coding (Forward Error Correction - FEC) and interleaving are techniques that have been implemented to combat bit errors which in a radio environment typically occur in bursts. As described earlier a mobile phone gets a channel - defined as the tuple (carrier #, slot #) - assigned at call-setup for data transmission. The slot cycle is 5 ms on average and 114 bits can be transmitted in each slot allocated to the mobile phone resulting in a gross data rate of 22.8 kb/s. Channel coding adds redundancy bits per data block (e.g. RLP frame) which then only leaves a net data rate of 12 kb/s to RLP or the transparent counterpart. Interleaving is a technique which in combination with channel coding is used to combat burst errors. This means that data blocks are not transmitted as a whole but are segmented into smaller pieces. Pieces of different data blocks are then interleaved before transmission. The benefit is that a few of these smaller pieces can get completely corrupted while a single data block can still be re-

1. The GSM specification 09.07 allows the use of any standard modem. A V.32bis modem at both ends would e.g. not impose this extra overhead thus providing the full 1200 bytes/s. However, today only V.32 modems are used in most GSM networks.

constructed by the channel decoder. The drawback of this approach is that it introduces additional latency as a data block now takes up more room in the time dimension. For the 9.6 kb/s data service in GSM one RLP frame of 240 bits is channel coded to 456 bits which would fit exactly into 4 slots or 20 ms. Instead these 456 bits are interleaved over 22 slots or roughly 110 ms which is outlined in Figure 2.

B. Enhanced Circuit-Switched Data

Phase 1 GSM systems which have been deployed in the early 90s and the Phase 2 systems which are operational today have mostly focused on voice services. Until lately GSM data services have only played a minor role and in fact GSM operators have reported that in the past only a negligible fraction of their overall traffic volume has been data. With the explosive growth of the Internet, however, this has changed dramatically and the current evolution of GSM - the so-called GSM Phase 2+ - is clearly driven by data and improved GSM Internet access. This section summarizes those changes that directly impact circuit-switched data.

The Direct-IP access solution integrates the Internet/Intranet access fabric into the GSM network, i.e. it terminates PPP connections in the GSM network (see Figure 6). Although simple, this architecture brings a number of advantages for both the GSM operator and the user. A GSM operator saves charges to the fixed network operator for using the fixed network as transit to an ISP. Instead the GSM operator could offer complete packages (e.g. "free Internet access with every GSM subscription") by teaming-up with an ISP to which all IP traffic would get routed or else the GSM operator could become an ISP. The latter would open the door for a number of what telcos call value-added services. Such services could e.g. include PPP-tunneling to corporate Intranets reducing costs for long distance circuit connections or location-dependent Web services could be offered as the GSM network always knows the whereabouts of its subscribers. The mobile user on the other hand might benefit from possibly better tariffs for wireless Internet access and will also benefit from shorter call-setup times. Today call-setup times for circuit-switched data excluding the establishment of the PPP link are in the range of 20 - 30 seconds when going through PSTN and 4 - 10 seconds when going through ISDN. The reduced call-setup times when using a Direct-IP Access Node stem from the fact that modems or rate adapters and the handshaking for these kinds of equipment are no longer involved in the transmission chain.

The major limiting factor for wireless Internet access via GSM today is clearly the low bandwidth of 9.6 kb/s and even the upcoming 14.4 kb/s coding scheme does not help so much. ETSI has therefore standardized a new data service called High-Speed Circuit-Switched Data Service (HSCSD) that allows a mobile phone to aggregate a number of channels. One key feature of this service is that a mobile phone can choose from a number different configurations (see Figure 3) including asymmetric allocation of up- and downlink bandwidth. Ba-

sically everything that has been said about the single-slot service also applies to the multi-slot service. A few changes are necessary of course like e.g. an adapted RLP protocol is required but the general architecture is the same. The second key feature is that the slot allocation that has been chosen at call-setup can be changed during the call or data session, i.e. slots on up- and downlink can be allocated and de-allocated dynamically. Both the 9.6 kb/s and the 14.4 kb/s channel coding scheme will be available for this service but the interleaving scheme has remained unchanged from the single-slot service. Thus the latency problems discussed in Section III remain.




Fig. 3. Possible slot allocations for HSCSD.

Although the standards define configurations where up to 8 time slots are aggregated both on up- and downlink the initial introduction of this service will probably be limited to maximum 4 time slots downlink and 2 time slots uplink. This again is partly due to terminal complexity. Thus, this new data service will initially provide a downlink data rate of up to 64 kb/s in ideal radio conditions. As the GSM radio interface only allows allocation of up- and downlink time slots symmetrically, resources will be wasted in the case of the asymmetric allocations. It will depend upon the operator's tariffing schemes how this will in the end be charged to the GSM subscriber.

III. QUICKSTART-PPP

After the physical GSM link has been established there is still a considerable delay before the first IP packet can be sent from the mobile host. This delay is due to the Point-to-Point Protocol (PPP) [14] link configuration phase over the high-latency GSM connection. In this section we suggest a solution by which the PPP configuration time can be completely eliminated in cases where both peer protocol entities have been appropriately modified. In cases where one PPP peer is and the other is not altered the modified peer will fall back to conform with PPP. This change will not increase the usual PPP configuration time in these situations but ensures interoperability.

A. Latency on the GSM Link

The high latency of a GSM link has been mentioned before [10], [11]. However an exact quantification is still an open question. We were therefore interested in determining the round trip time (RTT) as defined in [6] of a GSM link and wanted to investigate the causes for this delay. In the case of GSM circuit-switched data today this cost of course includes the delay incurred over the fixed telephone connection. All measurements cited in this section were done in a GSM network operating in the 1900 MHz band in the San Francisco Bay Area. The setup shown in Figure 1 was installed for these

experiments with the exception that the remote host at the same time was also the dial-in server, i.e. the PSTN was terminated directly in the remote host and there was no Internet involved. Standard PCs running BSDi 3.0 UNIX were used as end-hosts together with commercial GSM terminal equipment (Ericsson CF388 mobile phone and DC23 terminal adapter). All measurements were done using 9.6 kb/s circuit-switched data connecting via PSTN to the remote host and for the reasons discussed in Section IV RLP was always enabled. The mobile host was stationary at all times. We ran simple experiments using the UNIX ping command with a ping payload size of 12 bytes which resulted in a total packet size of 47 bytes².

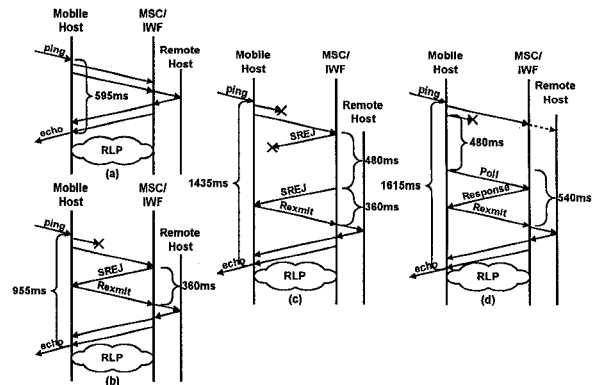


Fig. 4. Possible ping traces via GSM.

Several hundred ping measurements were done to get a sound statistical basis. Over 95 percent of all measurements resulted in an RTT of 595 ms with a standard deviation of less than 20 ms. A cluster of outliers resulted in an RTT of about 955 ms which suggested that in these cases a single RLP retransmission occurred and that the RTT on RLP layer is 360 ms (see (b) in Figure 4). At a later stage of our work we were able to confirm this with the trace collection platform described in Section V. In Figure 4 we have sketched four possible scenarios of what can happen on the RLP layer during a "ping" when initiated from the mobile host. Note that the payload of an RLP frame is 24 bytes so that two RLP frames have to be sent for the ping packet in uplink direction and also for the echo packet on the downlink. Note also that RLP uses a fixed retransmission timer recommended as 480 ms [3] for both lost selective reject messages and checkpoint recovery (see (c) and (d) in Figure 4 respectively).

We were curious about this enormous³ RTT so we also ran "pings" to the same remote host using an asynchronous 9.6 kb/s landline modem, i.e. without GSM and only going via PSTN. This time the RTT turned out to be 200 ms. Taking out the transmission delay in both cases (94 bytes / 1150 bytes/s

2. 8 bytes ICMP header, 20 bytes IP header and 7 bytes PPP overhead.
3. As a comparison running the same "pings" between LAN-connected hosts via the Internet from west coast U.S.A. to Europe or Asia which involves multiple hops resulted in RTTs of 250 - 300 ms in the worst case.

with GSM; 94 bytes / 960 bytes/s for the async. modem) we arrived at a round trip latency of roughly 520 ms for the GSM link and 110 ms for the PSTN link. Consequently the round trip latency on the GSM link excluding the PSTN part is on the order of 410 ms. Now, going back to Figure 2 we see that the interleaving mechanism used for the circuit-switched GSM data service introduces a one-way latency of 90 ms - a big hurdle for real-time data - which does not include the latency introduced by the interleaving/de-interleaving operations at both ends itself. We were not able to track down where the remaining latency of 230 ms resides but assume that it is due to the mentioned interleaving/de-interleaving operations, the expensive rate transcoding operations required at the BTS/BSC and the BSC/MSC interfaces in both directions, and buffering in the MS/TAF and MSC/IWF. We do not expect that the round trip latency of about 410 ms will become less for the future High Speed Circuit-Switched Data Service described in Section II as none of the factors that introduce the latency like e.g. interleaving will be changed for that service. Thus, in the future the GSM circuit-switched data link will remain to be a high-latency link.

B. The Problem

Figure 5 shows the PPP configuration mechanism consisting of messages that are exchanged during the LCP phase in which among other things the maximum transmission unit for IP packets is negotiated, the optional authentication phase (PAP or CHAP) and the IPCP phase which is required e.g. for negotiating the use of TCP/IP header compression [17] and for IP address assignment⁴. For simplicity LCP- and IPCP-Requests initiating from the dial-in server are not shown. In reality both peers start transmitting these requests almost simultaneously as soon as the “physical link up” event gets indicated to them.

It is important to note that [14] requires that neither the PAP or CHAP phase nor the IPCP phase, nor the exchange of IP packets can start before the preceding phase has successfully concluded. This results in a minimum PPP link configuration time of 2 to 3.5 RTTs depending on whether authentication is required or not. In practice this configuration time is often much longer when parameter settings differ on both sides and negotiations require more round trips in the LCP and IPCP phases. Given the latency in today’s GSM networks when using an external ISP (500-700 ms) as outlined above this can easily result in configuration times of several seconds. But even with the Direct-IP Access Nodes mentioned in Section II being deployed with an expected call-setup time for the physical link of below 2 seconds and a latency of around 450 milliseconds the PPP configuration time will at least add another 1 - 2 seconds before the link is available for IP traffic. Thus, the

4. Link Control Protocol (LCP) [14], Password Authentication Protocol (PAP) [15], Challenge-Handshake Authentication Protocol (CHAP) [15], and Internet Protocol Control Protocol (IPCP) [16].

problem is the high PPP link configuration time due to the extensive handshaking over a high latency link.

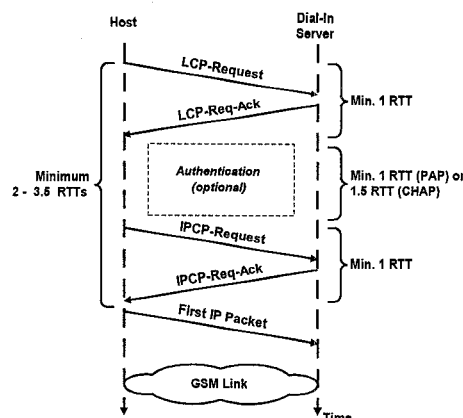


Fig. 5. PPP link configuration phase.

C. Suggested Solution

The general idea of the solution is to deliberately violate [14] by breaking the strict sequential order of the different negotiation phases (LCP, PAP or CHAP, and IPCP) and the exchange of IP packets. Instead we allow the concurrent exchange of all types of packets. However, it is important to point out that we do not require any protocol changes.

The challenge is to ensure that the defined protocol semantics are observed and that interoperability with existing PPP implementations is ensured. The following terminology is used in this section:

- Standard-PPP-Peer (S-Peer):

A PPP-Peer which only conforms to [14].

- Quickstart-PPP-Peer (Q-Peer):

A PPP-Peer which has an extended protocol state machine which allows it to conform to [14] if required, but also allows it to conform to the PPP protocol modifications as suggested here.

- Masked PPP Frames:

PPP frames are used to encapsulate LCP, PAP, CHAP, IPCP, or IP packets. PPP discriminates among those through the protocol field contained in each PPP frame header. Masked PPP frames are PPP frames which carry an “invalid” PPP protocol field value, i.e. they are currently unused values which will only be recognized by Q-Peers but must be silently discarded when received by S-Peers [14]. This guarantees interoperability. Thus, the protocol field values for LCP, PAP, CHAP, IPCP, and IP will be mapped to pre-defined but unused values.

In order to guarantee interoperability with S-Peers a Q-Peer must start off by sending an initial set of back-to-back masked PPP frames preceded by a standard PPP frame encapsulating an LCP packet which would have been sent in conformance to [14]. A receiving S-Peer will silently discard the masked PPP frames and process the standard PPP frame. A receiving Q-

Peer must silently discard the standard PPP frame but instead process the masked PPP frames. The Q-Peer delays discarding the standard PPP frame in order to check whether in fact masked PPP frames follow. If no masked PPP frames follow it processes the standard PPP frame. From the first frames (standard or masked) a Q-Peer receives it will be able to infer whether it is communicating with an S-Peer or a Q-Peer. When communicating with an S-Peer it falls back to operate as an S-Peer.




Fig. 6. Quickstart-PPP in a GSM network.

All IP packets (e.g. a TCP connect request) sent by a Q-Peer until it knows whether it communicates with another Q-Peer or not will have to be buffered for potential retransmission. In case the remote peer was in fact an S-Peer all non-LCP packets it received will have been discarded [14] and thus need to be retransmitted. Otherwise those buffered IP packets have to be discarded as the remote Q-Peer will have accepted them. If authentication is required, IP packets received during the authentication phase by a Q-Peer must be buffered and must not be released to the IP layer before the authentication phase has succeeded. If authentication fails while an IP address has already been assigned and/or IP packets have already been received the link must be terminated, the assigned IP address must be de-allocated and any received IP packets must be discarded. If the sending Q-Peer on the mobile host initially does not have an IP address assigned it may nevertheless send IP packets. In that case the source address will be left open and will be inserted by the receiving Q-Peer.

Thus, with these modifications the mobile host can send IP packets instantly saving at least 2 - 3 RTTs (depending on whether authentication is used) of initial delay. Implemented at an ISP's dial-in server and in mobile hosts this would in today's GSM networks result in a gain of usually several seconds but at least 2 seconds. The solution would be even more attractive for the Direct-IP Access solution where call setup times of the physical connection are expected to be reduced to below 2 seconds. Quickstart-PPP will make it feasible to implement a pseudo-packet mode where the physical connection is closed when idle and re-established when data has to be transmitted.

IV. MAKING THE LINK FLOW-ADAPTIVE

In this section we uncover a fundamental problem which exists when delay-sensitive and reliable transmission is required in parallel over an established GSM link. We suggest a solution based on the RLP protocol which will not require protocol changes but a more intelligent implementation of the RLP protocol peers.

A. The Benefit of RLP

As the GSM specifications [2] allow bit error rates of up to 10^{-3} after channel coding, i.e. as seen by RLP (see Figure 2) a small RLP frame size (30 bytes) had to be chosen for optimal performance. Hence, this rules out the possibility of relying solely on transport layer error control for reliable end-to-end data transmission. This choice would be an option when disabling RLP and instead using the transparent GSM circuit-switched data service. Choosing a small maximum segment size for TCP would be possible in that case but the overhead on the GSM link even with TCP/IP header compression [17] would be prohibitively high. Also the amount of data that had to be retransmitted through the entire Internet - assuming downlink transmission and packet corruption on the GSM link - would create quite some degree of unfairness as it wastes bandwidth that could otherwise have been used by other connections in the Internet. Hence, we strongly believe that reliable link layer protocols should be used in combination with TCP under the radio conditions prevailing in GSM. In fact [12] confirms that TCP performance does not at all degrade when run over RLP unlike claimed in related work [11]. On the other side delay-sensitive e.g. UDP-based realtime applications will certainly prefer not to run RLP and use the transparent mode instead.

B. The Problem

As motivated above we believe that for reliable end-to-end data transmission over radio access links with error characteristics as found in GSM the use of a reliable link layer protocol like RLP is required in addition to end-to-end error recovery as e.g. provided by TCP. On the other hand delay-sensitive applications will most likely not tolerate the delay introduced by potential link layer retransmissions which in the case of GSM makes the transparent data service a more suitable choice. However, this raises a fundamental problem when TCP-based and UDP-based applications are supposed to be used at the same time as the transparent and non-transparent (using RLP) GSM circuit-switched data services cannot be operated in parallel. The user has to decide at call-setup time to use one of them. Nevertheless, this does not have to be this way as RLP does provide both a reliable (I-mode) and an unreliable mode (UI-mode) [3], but it is only the I-mode which is implemented by most manufacturers of GSM equipment today.

C. Suggested Solution

As described in Section III PPP is used to transport packets of multiple protocols simultaneously over a single serial link. The PPP protocol field identifies which type of packet is encapsulated in a particular PPP frame. Also IP - which in turn is carried by PPP - can carry data of multiple protocols, e.g. TCP, UDP, and ICMP which are also distinguished by a protocol identifier in the IP header. The different flows that are eventually transported by PPP (TCP, UDP, ICMP, LCP, PAP, etc.) have different requirements concerning the trade-off of reliability versus delay. Consequently, either the I-mode or the UI-

mode of RLP is appropriate for certain flows. More specifically TCP, ICMP, and PPP signalling packets should be carried in I-mode whereas real-time data like UDP should be sent in UI mode. We therefore suggest a solution that allows both modes (I and UI) of RLP to be used simultaneously *without* specific control from higher layers. We call such a link layer “flow-adaptive” reflecting the fact that it can dynamically adapt itself to each flow’s requirements by inspecting the respective packet headers.

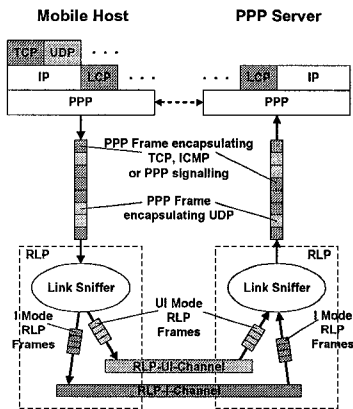


Fig. 7. Reliable and delay-sensitive data over RLP.

Currently the RLP sender handles the data it receives from the higher layer as a transparent byte stream and the protocol guarantees reliable in-order delivery of this byte stream when released by the RLP receiver. The solution is simply to make the RLP protocol implementation flow-adaptive by parsing for PPP frame delimiters and then to inspect the PPP frame headers and if required also the IP packet headers, i.e. to “sniff” on every PPP frame. In most cases it will be sufficient to just look for the protocol identifier field in the PPP frame header to determine the type of packet contained. In other cases it will be necessary to in addition inspect the IP header to discriminate among TCP, UDP and ICMP packets. Then by pre-configuration - or in the future by new signalling means between the RLP peers - the RLP peers can transmit 2 categories of higher layer protocols simultaneously. The category of protocols requiring reliability will be transmitted in I-mode and the category of protocols which tolerate losses but require low delay will be transmitted in UI-mode. This could be combined with priority-based scheduling e.g. giving UDP traffic higher priority. This mechanism is depicted in Figure 7.

V. IMPLEMENTATION

This section describes in progress implementation work which is done within the context of the ICEBERG project [18] at the University of California at Berkeley lead by Prof. Randy Katz and Prof. Anthony Joseph. The project is a cooperation with Ericsson Radio Systems AB (Sweden). The overall

project scope is to study the potentiality of an IP-based core network for future cellular systems.

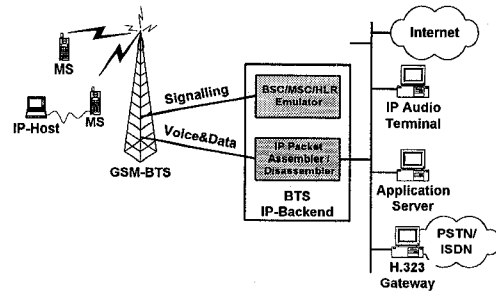


Fig. 8. The ICEBERG testbed.

As one of the first steps in this project we have implemented a testbed connecting a GSM base transceiver station (BTS) to an IP subnet. Eventually the testbed depicted in Figure 8 will serve multiple purposes according to the different research interests. The main components of the testbed is the BSC/MSC/HLR-Emulator that takes care of all SS7-based GSM signalling and the IP-PAD which converts between circuit-switched and IP-routed traffic.

With respect to the work described in this paper the testbed will be used for performance measurements. For that purpose and also for the on-going implementation of the Link Sniffer described in Section IV we have ported the RLP code to BSDi UNIX. In addition we have instrumented the code with a monitoring tool which logs the RLP sender and receiver state over time. This has been used in [12] to study interactions between the RLP and TCP error control mechanisms and also for trace collection to be used for trace replay in simulations [13]. The resulting trace collection platform is shown in Figure 9.

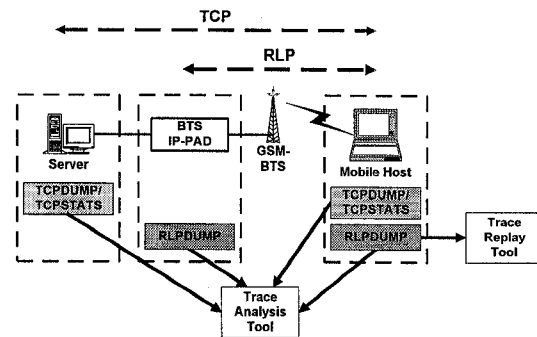


Fig. 9. The TCP/RLP trace collection platform.

VI. FUTURE WORK

After having developed solutions to two link layer problems we will extend our work to also look at higher layer protocols. More specifically we will investigate the interactions of TCP and link layer (e.g. RLP) error recovery. The goal is to find out through extensive measurements if at all and if yes under which circumstances and how frequent the existing race condition of competing retransmissions at both layers has a negative impact on TCP performance for bulk data transmissions. The

TCP/RLP trace collection platform shown in Figure 9 is actually a first step to start that analysis. The platform has the great advantage of allowing us to correlate the events on both layers to understand their interdependence. In parallel we have started to implement this setup in a simulation environment to be able to rapidly reproduce and analyse certain effects. This will help us to study such complex interactions and provide a basis to investigate the benefit of solutions proposed in [7]-[11] and compare them with potential alternatives. In particular we are interested in extending the notion of flow-adaptiveness which we have introduced with the Link Sniffer. The goal will be to offer more fine-grained differentiated service at the link layer by also adapting e.g. forward error correction schemes or power control to given per flow QoS requirements. The RLP traces collected during our measurements will be used for trace replay [13] in the simulator. Promising solutions will be implemented and evaluated through real world measurements. These will then be used to validate simulation results obtained beforehand.

ACKNOWLEDGMENT

Thanks to Almudena Ordonez and Kimberly Oden for collecting many of the ping measurements mentioned in Section III and their work on the trace analysis tool mentioned in Section V. Thanks to Keith Sklower for his help in getting the TCP/RLP trace collection platform running. Last but not least many thanks to Prof. Randy Katz and Prof. Anthony Joseph for the hospitality we enjoy at UCB and the great cooperative spirit in the ICEBERG project.

REFERENCES

- [1] Mouly M., Pautet M.-B., *The GSM System for Mobile Communications*, Cell & Sys, France 1992.
- [2] ETSI, *Quality of Service*, GSM Specification 02.08, Version 3.0.0, March 1990.
- [3] 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 Centre (BSS-MSC) interface*, GSM Specification 04.22, Version 5.0.0, December 1995.
- [4] Ludwig R., Turina D., *Link Layer Analysis of the General Packet Radio Service for GSM*, In Proceedings of ICUPC 97, October 1997.
- [5] Tanenbaum A., *Computer Networks*, 2nd Edition, Prentice Hall, 1989.
- [6] Stevens W.R., *TCP/IP Illustrated, Volume 1 (The Protocols)*, Addison Wesley, November 1994.
- [7] Bakre A., Badrinath B. R., *I-TCP: Indirect TCP for Mobile Hosts*, In Proceedings of ICDCS 95, May 1995.
- [8] 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.
- [9] Balakrishnan H., Padmanabhan V., Seshan S., Katz R. H., *A Comparison of Mechanisms for Improving TCP Performance over Wireless Links*, In Proceedings of ACM SIGCOMM 96.
- [10] Baucke S., *Leistungsbewertung und Optimierung von TCP für den Einsatz im Mobilfunknetz GSM*, Diploma Thesis, CS-Dept. 4, Aachen University of Technology, Germany, April 1997.
- [11] 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.
- [12] Ludwig R., Rathonyi B., Ordonez A., Oden K., Joseph A., *Multi-Layer Tracing of TCP over a Reliable Wireless Link*, submitted for publication.
- [13] Noble B. D., Satyanarayanan M., Nguyen G. T., Katz R. H., *Trace-Based Mobile Network Emulation*, In Proceedings of ACM SIGCOMM 97.
- [14] Simpson W., *The Point-to-Point Protocol*, RFC1661, July 1994.
- [15] Lloyd B., Simpson W., *PPP Authentication Protocols*, RFC1334, October 1992.
- [16] McGregor G., *The PPP Internet Protocol Control Protocol (IPCP)*, RFC1332, May 1992.
- [17] Jacobson V., *Compressing TCP/IP Headers for Low-Speed Serial Links*, RFC1144, February 1990.
- [18] The Iceberg Project. CS Division, EECS Department. University of California at Berkeley. Berkeley, CA. <http://iceberg.cs.berkeley.edu/>.