TCP over Wireless Networks

1 Introduction

Wireless networks are becoming more widely deployed and more often used to access services in the Internet. Internet technology has been successful in providing services to users in fixed networks. In wireless networks, on the other hand, the performance of the Internet protocols has been reported to be much lower than in fixed networks . The main reason for the performance degradation is that the Transmission Control Protocol (TCP) works less efficiently in wireless networks. This problem is important, since TCP is used by many popular Internet applications, such as e-mail, web browsing, and remote login. TCP was designed for networks with wired links and stationary hosts. In these networks, data is lost mainly due to congestion. TCP interprets all data loss as congestion in the network, and in case of data loss TCP slows down its transmission rate in order to reduce the congestion. In a wireless network, it is no longer appropriate to assume that most losses are caused by congestion. Data loss is often caused by the relatively low quality of the wireless link. Terminal mobility, which is supported by many wireless networks, may also result in data loss. If data gets lost for some other reason than congestion, then performance is unnecessarily degraded as TCP reduces its transmission rate in response to the loss.

There are many proposals in the literature on how to optimize TCP performance in wireless networks. The main idea, shared between the optimization proposals, is that TCP should only reduce its transmission rate in case of congestion, not if data is lost for other reasons. Some of the optimization proposals are primarily concerned with problems related to the quality of the wireless link, and yet others try to enhance TCP performance in respect to mobility. Many optimization proposals are evaluated in a simulated environment that is based on general assumptions of a wireless link, such as a high probability of transmission errors, low data rates, and long delays. In comparison to fixed networks, all wireless networks may seem to have the same properties. However, if wireless networks are compared to each other instead, then there are important differences.

we try to evaluate optimization proposals in relation to currently deployed wireless networks. This, however, turned out to be a more challenging task than expected. Some proposals are clearly intended to solve problems that are inherent to a specific type of wireless network, while other proposals might work efficiently in many networks. First, an overview is provided of some wireless networks. The focus is placed on properties that are important for data transmission. Then, some ideas on how to improve TCP performance in wireless networks are presented. The applicability of the optimization proposals in different wireless networks is also discussed. The optimizations presented in this report are primarily aimed at problems related to the quality of the wireless link and to mobility. Communication over satellite links is not considered, since many of the problems involved are not applicable to other wireless networks. Most of the optimization approaches presented do not require mobility on the IP level, such as Mobile IP . It is assumed that mobility is handled by the wireless network.

2 Wireless Networks

A wireless network consists of mobile stations and various intermediate nodes. Some type of intermediate node is required to connect a wireless network with a wireline network. A cellular telephony network, for example, is connected to a wireline network by an inter-working unit, and a wireless local area network (WLAN) is interconnected with an access point (also called base station).

Wireless links are not as robust as wireline links, since the radio quality may vary considerably over time, the bandwidth is usually lower, and transmission errors occur more frequently. Sending signals over an omnidirectional radio based medium gives rise to more errors than in a guided medium such as fiber or coax. Signal strength weakens with the distance between the mobile station and the base station, and radio waves bounce off objects, giving rise to interference and multi-path effects.

In order to shield upper protocol layers from transmission errors both error correction, interleaving and retransmissions can be used at lower layers. In many wireless networks, the data link layer performs error recovery according to some automatic repeat request (ARQ) protocol. In [18] link ARQ protocols are categorized according to the level of reliability provided to upper layers. An ARQ protocol is defined as perfectly persistent or reliable, if it retransmits frames until they are acknowledged or, after a very large number of retransmission attempts, disconnects the link and notifies upper layers. If the maximum number of retransmissions (or the maximum time spent retransmitting a frame) is limited, to tens of retransmissions, then the ARQ protocol is defined as highly persistent or highly reliable. A low persistent or partially reliable ARQ protocol, on the other hand, retransmits a frame 2-5 times before it gives up and transmits the next frame instead.

In this section, we present a selection of wireless networks that provide data services and support user mobility. We begin with WLANs. Then wireless wide are networks (WWANs) are described. Finally, some conclusions are presented.

2.1 Wireless LANs

WLANs are standardized both by IEEE and by ETSI. The standards cover the physical layer and the medium access control (MAC) protocol used in the lower part of the data link layer. On top of the MAC protocol, a logical link control (LLC) protocol, such as IEEE 802.2, is typically used.

The data rates are much higher than in wireless wide area networks. In comparison to WWANs, mobility is more or less limited, depending on the technology used, e.g. infrared WLANs hardly support mobility at all. Today, WLANs with data rates up to 108Mbps are commercially available. The cost for this higher data rate is that the terminals must be close to the access point, e.g. 50-100 meters. As IEEE 802.11 is the dominant standard for WLANs, we have chosen to describe the IEEE 802.11 standard in some more detail below.

2.1.1 IEEE 802.11

An IEEE 802.11 WLAN [48] (operating in infrastructure mode) consists of one or more access points (APs) interconnected by a distribution system (typically a wired LAN). The coverage area of an AP is called a basic service set. The basic service sets of many interconnected APs may form an extended service set. Mobility between service sets is supported, since handover is performed between APs. Handover is initiated by the mobile station. It takes between 60 and 400ms for a handover to complete, depending on the network interface card [36].

In the IEEE 802.11 standard, the algorithm used for medium access control is carrier sense multiple access with collision avoidance (CSMA/CA). CSMA/CA is similar to the access control used in wireline LANs, CSMA with collision detection (CSMA/CD). The number of collisions can be reduced by fragmenting the frames into smaller units before they are transmitted over the air. However, fragmentation is rarely used [22]. In a WLAN, collision detection cannot be used as a sign of unsuccessful transmission as in a wireline LAN, since a station in a WLAN is unable to simultaneously transmit and listen to the same channel. Instead medium access for unicast communication in a WLAN relies on positive acknowledgments. Successfully transmitted frames are acknowledged and if no acknowledgment arrives the sender retransmits the frame. After a few, typically three, retransmissions the frame is discarded. For delay sensitive traffic, IEEE 802.11 also standardizes contention-free frame transfer. A sender reserves the wireless link and does not have to contend for the channel.

2.2 Wireless WANs

The first generation of cellular telephony networks were based on analog technology for the radio interface. In comparison to the second generation (2G) of cellular telephony networks, the first generation networks are less suitable for data traffic, since lower bandwidth, poorer radio quality, and less security are provided.

The 2G systems are digital and based on either time division multiple access (TDMA) or code division multiple access (CDMA). In TDMA networks, a mobile station can only listen to one base station at a time. As a result, there is a short loss of connectivity during handover, as a mobile station moves from one cell to another. A handover which results in a short interruption is called a hard handover. In CDMA networks, it is possible to support soft handover in which connectivity is maintained during the handover. In a soft handover, interruption due to handover is avoided, since the mobile station may communicate with both the old and the new base station during the handover. However, soft handover cannot always be applied, e.g. if the base stations are unsynchronized or use different frequency bands.

In a 2G system, an interworking function (IWF) in a mobile switching center (MSC) handles inter-networking between the wireless network and the fixed telephony network. The bit error rate (BER) achieved after applying various techniques at the physical layer (e.g. channel coding, error correction, and interleaving) is sufficiently low for telephony. For data services, on the other hand, additional error recovery is usually required, since most data services are more sensitive to bit errors than telephony services. The reliability is enhanced for data services by using a radio link protocol that applies ARQ.

As an intermediate step toward the third generation of cellular networks (3G), the 2.5G systems provide higher data rates and a packet-switched service for data traffic. Due to packet-switching, many more users may share the available resources than in 2G, since a channel is only temporarily assigned to a mobile station. Instead of reserving channels for longer time periods, the mobile stations contend for medium access when they have data to transmit. Channel access is controlled by a MAC protocol.

Third generation cellular networks (3G) are currently being deployed. The ITU standard for 3G wireless communications, the International Mobile Telecommunications-2000 (IMT-2000) provides a framework for 3G systems. In contrast to the 2G systems, 3G is already from the beginning designed with support for high-speed data. The average data rates are in the range 64-384kbps. In the future, peak data rates between 2 and 20Mbps are expected. In the late 1990s, CDMA was chosen as the predominate physical access technology for 3G. CDMA and packet-switching provide for a higher capacity than in circuit-switched systems based on TDMA technology. The introduction of 3G is made in steps which implies that the first systems are based on 2G and 2.5G. The Universal Mobile Telecommunications System (UMTS), also called wide-band CDMA (WCDMA) is the evolution of TDMA based systems, such as GSM/GPRS and D-AMPS (IS-136). A similar development of CDMA based systems (IS-95) is defined in the cdma2000 specification. In China yet another variant of 3G is defined, time division synchronous CDMA (TD-SCDMA).

In the next sections, some examples of WWANs are presented. The relation between the presented WWANs are illustrated in Figure 1, as introduced above. First, we give an overview of the development of the Global System for Mobile Communications (GSM) to UMTS. Then, the Advanced Mobile Phone Service (AMPS) and its development are described. Finally, the development of IS-95 from 2G to 3G is outlined.

2.2.1 Development from GSM to UMTS

GSM

GSM [37] is a TDMA based digital cellular network for circuit-switched voice and data transmission. A mobile station accesses the Internet via a base transceiver station and a base station controller. An IWF in an MSC serves as the interface between GSM and the fixed telephone network which in turn is connected to the Internet. In GSM, a mobile assisted handover scheme is applied [28]. The mobile station performs measurements of the radio quality and reports this information back to the network which in turn uses the information to decide if a handover should be initiated. A handover takes about one second to complete [28]. A mobile station communicates with one base station at a time which implies that a short interruption occurs due to handover.

The data service available in GSM provides transparent and non-transparent modes of operation. The non-transparent mode uses an radio link protocol (RLP) [16] between the mobile station and the IWF. The retransmission scheme is a selective repeat ARQ based on positive acknowledgments. The frames size is 240 bits in single-link mode and 576 bits in multi-link mode. Frames are retransmitted in response to a reject, a selective reject or a status report request that the sender requests from the receiver if no acknowledgment has arrived after timeout interval.

Depending on the RLP implementation, the link is either reset or disconnected after a limited number of repeated retransmissions of the same frame. The maximum number of retransmissions is configurable with a default value of 6. An implementation that resets the link provides a partially or highly reliable service, depending on the maximum number or retransmissions [18]. At link reset, the RLP entities empty their buffers and further error recovery of lost data is silently left to higher layers. A fully reliable RLP, on the other hand, disconnects the link and notifies its users [29]. Data is not, as in the reset case, silently discarded by the link layer.

GPRS

General Packet Radio Service (GPRS) [17, 9, 42, 43] is a packet oriented extension to GSM which provides for higher data rates and more efficient network utilization compared to circuit-switched GSM. GSM channels, each corresponding to a time slot, may either be reserved for GPRS or dynamically allocated when required. If no channels are exclusively reserved, a GPRS transfer may be interrupted due to preemption by traffic with higher priority, such as circuit-switched GSM. The data rate depends on the coding scheme. A mobile station may receive data on up to eight GSM channels, but most terminals today are not capable of receiving data on more than four channels. The maximum data rate for the most commonly used coding scheme (CS-2) is 13.4kbps per channel.

The new functions required for GPRS are provided by GPRS support nodes which are connected by an IP-based GPRS backbone. Gateway GPRS support nodes (GGSNs) serve as interfaces between the GPRS network and other packet data networks such as the Internet. Serving GPRS support nodes (SGSNs) are responsible for routing packets to the right base station subsystem.

An LLC protocol between the mobile station and an SGSN provides acknowledged and unacknowledged modes of operation. The maximum LLC frame size is 1556 bytes. A radio link control (RLC) protocol between the mobile station and the base station system fragments the LLC frames if necessary to fit into blocks suitable for transmission over the radio interface. The RLC frame size is 256 bits for CS-2. RLC also supports acknowledged and unacknowledged modes of operation. The MAC protocol used for channel access is a slotted ALOHA reservation protocol.

UMTS

In UMTS [49], as compared to GPRS, functionality has moved from the core network to the radio access network, the UMTS terrestrial radio access network (UTRAN). The radio resource management, for example, is placed closer to the mobile station, in the radio network controller (RNC) instead of in the SGSN. No LLC protocol is therefore required between the MS and the SGSN. This gives a lower protocol overhead. The RLC protocol in UMTS supports transparent, unacknowledged and acknowledged modes of operation. The transparent mode provides a byte stream service to the layer above. The overhead is lower than in other modes, since no RLC header is added. Error correction is performed only in acknowledged mode. A sliding window based ARQ scheme uses selective repeat to request retransmission. The RLC block size is configurable with typical values of around 320 bits. Retransmissions are triggered by acknowledgments called status report messages that are transmitted by the receiver. The status report messages serve as both positive and negative acknowledgments, i.e. received frames are cumulatively acknowledged and missing blocks are indicated. The RLC receiver transmits a status report message in three cases: to request retransmission of erroneous RLC blocks, in response to poll messages from the sender, and at certain times as a periodic event. Retransmissions have a higher priority than new data blocks. RLC delivers data in order and removes duplicates before delivery. The reliability is low to high depending on the the maximum number of retransmissions, which can be set to a value of up to 40.

2.2.2 Development from AMPS to D-AMPS and CDPD

The analog system used in North America, AMPS, has been developed into the digital advanced mobile phone service (D-AMPS), a digital 2G system, and into the cellular digital packet data (CDPD) system, a packet data overlay network on top of the analog AMPS system.

D-AMPS

D-AMPS [45] is a TDMA based system, which is similar to GSM. As in GSM, a mobile assisted handover scheme is applied [28]. Also the RLP is similar to the one used in GSM. In contrast to GSM, a selective reject may be used to request retransmission of multiple frames. The size of an RLP frame is 256 bits, and each frame has a sequence number in the range 0-127 and a 16-bit checksum. The maximum data rate is 9.6 kbps. The 2.5G development of D-AMPS is also similar to GSM, since D-AMPS is extended to provide GPRS.

CDPD

CDPD [45, 42] is an infrastructure for packet-switched data that is implemented as an overlay network to the analog AMPS system. CDPD uses idle voice channels to transmit data and if no channel is available packets will be delayed until a channel becomes free. The data rate achieved after error coding is 9.6 kbps. As analog cellular networks are being replaced, the packet operations provided by CDPD will be taken over by second and third generation cellular networks.

In AMPS, network controlled handover is applied [28]. This type of handover is less efficient than the mobile assisted handover used in GSM and D-AMPS. The base stations must measure the radio quality of the mobile stations and transmit this information to other base stations and to the MSC. Handover may take up to ten seconds, since extensive signaling is required and the radio quality is measured infrequently.

Digital carrier sense multiple access (DSMA/CD) is the medium access control protocol used in CDPD. The protocol is similar to the carrier sense multiple access (CSMA/CA) protocol used in wireless LANs described above. A mobile station that has data to transmit must first check if a busy/idle flag bit is set on the forward channel before it can contend for the reverse channel. The mobile data link protocol (MDLP) manages logical links between a mobile end system and a mobile data intermediate system which is the node connecting the CDPD network with external packet oriented networks. The protocol supports both acknowledged and unacknowledged mode of operation. Between the network and link layer lies the subnetwork dependent convergence protocol (SNDCP). SNDCP segments network protocol data units which have a maximum size of 2048 bytes into blocks of 130 bytes (default value) that are passed to the link layer.

2.2.3 Development from IS-95 to cdma2000

IS-95

IS-95 is a 2G digital cellular system [14, 25, 26] based on CDMA. IS-95 and its later revision, IS-95B, are also called cdmaOne. The network architecture is similar to GSM and D-AMPS. A mobile station is connected to a base station and a base station controller, and an MSC interconnects the wireless network with external networks. The handover scheme is also similar to GSM and D-AMPS in that the mobile assists the network with measurements of the radio quality. In contrast to the TDMA systems, IS-95 CDMA supports soft handover [28], as described above.

For data transmission in IS-95 [26], a non-transparent mode of operation uses an RLP based on negative acknowledgments (NAKs). The retransmission scheme is less complex than the one used in GSM. The receiver requests a retransmission by sending NAKs to the sender. The NAK scheme is possible, since the system is synchronous and RLP frames are transmitted every 20ms. When there is no data to transmit, the sender transmits the sequence number of the last data frame in an idle frame (which requires less resources than a data frame). The receiver responds with NAKs if frames are missing. If a (1,2,3) NAK scheme is used (other schemes are possible), the receiver first transmits one NAK after a timeout period. If data is still missing after a second timeout period, then two NAKs are transmitted. Finally, after a third timeout period, three NAKs are transmitted. After the attempt with three NAK rounds for the same frame, RLP gives up and leaves further error recovery to higher layers.

The maximum data rate is 9.6kbps. The frame size is 171 or 266 bits, depending on the rate. Undetected errors may occur, since the frame checksum is relatively weak, only 8 or 12 bits per frame depending on the rate [26].

IS-95B

IS-95B is the 2.5G development of IS-95. Packet-switched data services are added to the IS-95 system by a software upgrade in the base stations and new user terminals. As in GPRS, up to eight channels may be assigned to one user which gives a maximum data rate of 115kbps. In practice, the achievable data rate is about 64kbps.

cdma2000

The cdma2000 specification is a 3G development of IS-95. In contrast to UMTS, it is possible to reuse the radio spectrum, bandwidth, and radio interface from preceding 2G and 2.5G, since these are already based on CDMA. The first step is cdma2000 1x which gives a higher capacity for both voice and data as compared to previous systems. The average throughput per user is 144kbps with instantaneous data rates of up to 307kbps. Even higher rates, up to 2-3Mbps, will be available with cdma2000 1xEV-DO (evolution for data only) and cdma2000 1xEV-DV (evolution for data and voice). The radio link protocol used over the air uses a selective repeat ARQ based on NAKs similar to the one applied in IS-95, described above. RLP performs a limited number of retransmissions before it gives up. Error detection and further error recovery are left to upper layers.

2.3 Conclusion

In comparison to fixed networks, wireless networks may seem to all have very similar properties. However, as presented above, a closer investigation shows that there are some important difference between various systems. Next, the presented wireless networks are summarized and contrasted with emphasize on characteristics that may affect upper protocol layers.

When the radio quality is low, data loss may occur over a wireless link due to transmission errors. The link level ARQ is capable of providing reliability is most cases, especially if the ARQ is highly persistent as in many cellular networks. In WLANs, on the other hand, data may be lost if the radio conditions are very poor, since only a few retransmission attempts are made. In cellular networks, handover is a more likely cause to data loss [22]. Data loss due to handover occurs particularly often if a user moves at high speed.

Handover typically results in delay and, in many cases, also in data loss. Delay is introduced, since it takes time to forward data to the new base station and to perform the handover procedure, e.g. signaling messages must be transmitted between the nodes involved in the handover. Data loss occurs if the old base station flushes its buffer instead of forwarding data to the new base station. Data loss due to handover can be avoided if a logical link protocol that performs ARQ is used on an additional link layer above the radio link protocol. Some examples are LLC in acknowledged mode in GPRS, and the logical link protocol used in CDPD. Even so, the handover is not transparent to upper layers, since recovery of lost data introduces delay.

Wireless networks have a long delay compared to wired networks, since transmission over a radio interface is slower than over a wired medium. Additional delay may be introduced due to processing on the physical layer and on the data link layer. Processing on the physical layer results in a constant delay. In cellular networks, processing on the physical layer (error correction and interleaving) is extensive and therefore gives a relatively long delay. On the data link layer, delay is increased by the use of MAC and ARQ. This delay is variable, since it depends on other users' activity (MAC) and on the radio conditions (ARQ). The delay variation experienced by upper protocol layers may become very large if the link layer ARQ is highly persistent as in GSM, GPRS and UMTS (depending on the configuration).

3 TCP

TCP [40] is a connection oriented transport protocol which provides a reliable byte stream to the application layer. Application data submitted to TCP is divided into protocol data units (PDUs) called segments, before transmission. Reliability is achieved since TCP uses an ARQ mechanism based on positive acknowledgments. Each byte is numbered and the number of the first byte in a segment is used as a sequence number in the TCP header. A receiver transmits a cumulative acknowledgment in response to an incoming segment which implies that many segments can be acknowledged at the same time.

TCP manages a retransmission timer which is started when a segment is transmitted. If the timer expires before the segment is acknowledged, then TCP retransmits the segment. The retransmission timeout value (RTO) is calculated dynamically based on measurements of the round trip time (RTT) [39], i.e. the time it takes from the transmission of a segment until the acknowledgment is received.

In October 1986 the Internet had its first congestion collapse. The end hosts transmitted more data than the routers were able to handle, and did not lower the transmission rate even though many packets were lost. Hence the congested state persisted in the routers. Since then TCP has been extended with mechanisms for congestion control [24]. Today all TCP implementations are required to use algorithms for congestion control, namely, slow start, congestion avoidance, fast retransmit, and fast recovery [3].

3.1 Slow Start and Congestion Avoidance

The purpose of slow start and congestion avoidance is to control the transmission rate in order to prevent congestion from occurring. TCP is described as a self-clocking protocol, since the transmission rate is determined by the rate of incoming acknowledgments. The sender only transmits segments when acknowledgments are received.

TCP estimates the available capacity in the network by gradually increasing the number of outstanding segments. The congestion window (cwnd) limits the amount of data the TCP sender can inject into the network. The initial value of the congestion window is between one and four segments [2]. The receiver window (rwnd) indicates the maximum number of bytes the receiver can accept. The value of the rwnd is advertised to the sender, since the receiver includes rwnd in the segments going back to the sender. At any moment, the amount of outstanding data is limited by the minimum of the cwnd and the rwnd.

In the slow start phase, the congestion window is increased by one segment for each acknowledgment received, which gives an exponential increase of the congestion window. Slow start is used for newly established connections and after retransmission

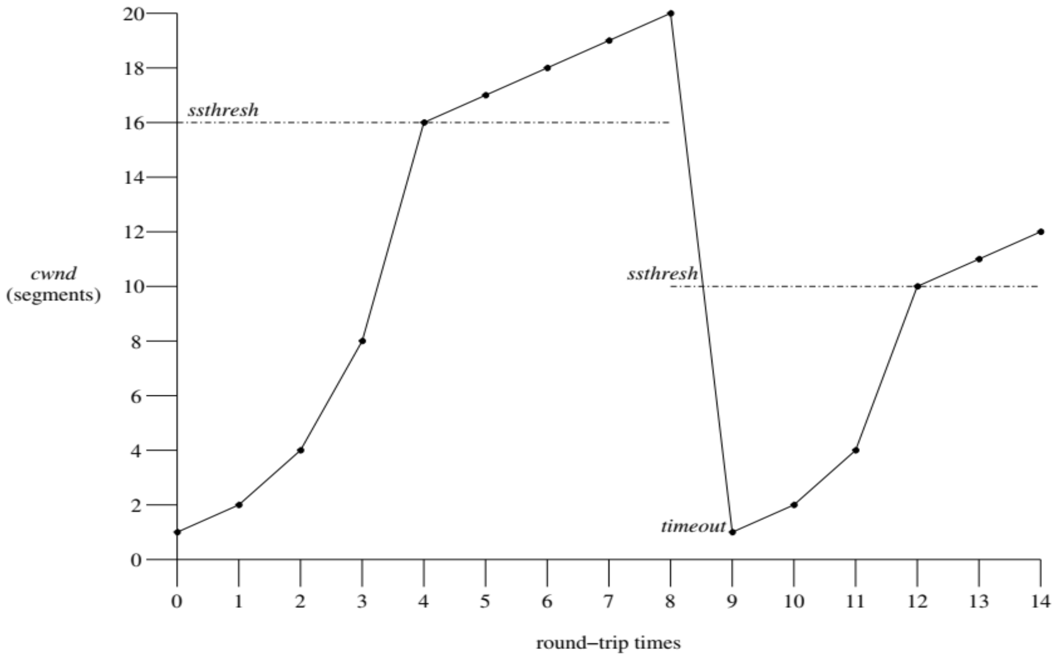

Figure 2: Slow start and congestion avoidance when timeout occurs.

due to timeout. The congestion window is increased until a timeout occurs or a threshold value (ssthresh) is reached. If a timeout occurs, then ssthresh is reduced to half the amount of outstanding data, the congestion window is reduced to one full-sized segment, and the slow start phase is entered again. If ssthresh is reached, then the slow start phase ends and congestion avoidance is entered instead. During the congestion avoidance phase, the congestion window is increased by one segment per round trip time, which gives a linear increase of the congestion window. Figure 2 illustrates how the congestion window is changed in slow start and in congestion avoidance.

As compared to the congestion collapse discussed above, these algorithms make TCP slow down when packets are lost. With less packets injected into the network, the load on the routers decreases and packets can flow through.

3.2 Fast Retransmit and Fast Recovery

The fast retransmit and fast recovery algorithms [3] allow TCP to detect data loss and perform error recovery before the transmission timer expires in some situations. The algorithms increase TCP performance, partly due to the earlier loss detection and re-transmission, partly since the transmission rate is not reduced as much as after timeout.

If a segment arrives out of order, the receiver transmits an acknowledgment for the last segment received in sequence. Since this segment already has been acknowledged once before, when it was first received, this subsequent acknowledgment is called a duplicate acknowledgment (dupack). After receiving three dupacks in a row, the sender concludes that unacknowledged data that was transmitted before the dupacked segment must have been lost. Data is retransmitted directly after the receipt of the third dupack (the fourth acknowledgment) even if the retransmission timer has not expired. After the retransmission, fast recovery is performed until all lost data is recovered and the sender receives an acknowledgment which covers new data. First, the ssthresh is reduced, as after timeout, to half the amount of outstanding data. and the congestion window is set a higher value, to three full-sized segments more than ssthresh. The additional three segments accounts for the three segments which triggered the receiver to transmit the dupacks. If more dupacks are received, then the congestion window is increased with one segment for each dupack, since each dupack indicates that one segment has left the network. When an acknowledgment for new data is received, then the congestion window is set to the same value as ssthresh. The effect of this adjustment of the congestion window after fast retransmit and fast recovery, is that TCP may enter congestion avoidance instead of slow start, as is done after timeout.

3.3 TCP Options

The performance of TCP may be enhanced by the use of optional features. Some of the commonly used options which are relevant for TCP in wireless networks are selective acknowledgments (SACK) [33], timestamps [23] and window scaling [23].

3.3.1 Selective Acknowledgments

The selective acknowledgment options (SACKs) [33] improves TCP performance, if multiple segments are lost in the same window. With SACK enabled, a receiver can acknowledge up to three non-continuous blocks of received bytes in the same acknowledgment. The sender then knows which segments are missing and can retransmit only those.

3.3.2 Timestamps

The timestamps option [23] provides an additional means to identify segments and their acknowledgments. A 12 byte timestamp is added to outgoing segments and the receiver adds the same timestamp to the acknowledgments going back to the sender. If the timestamps option is enabled, then the sender can sample the round trip time with a higher frequency, which gives a more accurate round trip time estimation. This is especially useful when using large windows, since the round trip time can be estimated more often than only once per window.

3.3.3 Window Scaling

The window scale option [23] can be used in order to utilize the network capacity between the sender and the receiver more efficiently. The bandwidth-delay product (a measure of the capacity) may be larger than the maximum value of the header field for the advertised receiver window (16 bits). This means that the transmission is limited by the advertised receiver window, although the network can transport more data. With the window scale option, a larger window can be used, since, it is possible to advertise receiver windows of 32 bits.

3.4 Other Mechanisms

3.4.1 Limited Transmit

Limited transmit [1] is a modification to the loss recovery algorithm in TCP. Without limited transmit, TCP segments are only retransmitted at timeout or when three dupacks are received, as discussed earlier. The limited transmit mechanism allows TCP to transmit a new segment already when the first dupack arrives. The arrival of a dupack indicates that one segment has reached the receiver and left the network. By transmitting a new segment, TCP checks if the network is congested or not. If the new segment reaches the receiver, it increases the probability of fast retransmit, since the receiver transmits a dupack in response to the new segment.

Limited transmit improves TCP performance when the transmission window is too small for fast retransmit and fast recovery to be triggered. For example, if the window is only two packets, the sender may only transmit two packets and therefore three dupacks can not be generated. Limited transmit allows for injection of more packets which may lead to more dupacks, which then trigger fast retransmit.

3.4.2 Increased Initial Window

In [2], it is proposed to allow an initial value of the congestion window of up to four segments, instead of up to two segments [3]. A larger initial value of the congestion window increases TCP performance, especially for connections that transmit a small amount of data. Fewer round trip times are required before the congestion window opens up and the total transmission time for a certain amount of data is shorter than it would with a smaller initial window.

3.5 TCP Variants

There are many implementations of TCP, some of which are considered as baseline TCP implementations. Three TCP implementations in the BSD operating system, named Tahoe, Reno and NewReno, are the ones most commonly referred to. TCP Tahoe is the original 4.4BSD implementation, including the congestion control scheme devised by Jacobson [24], i.e. the slow start and congestion avoidance algorithms mentioned earlier. The addition of the fast retransmit and fast recovery algorithms [3] to TCP Tahoe are called TCP Reno. TCP NewReno [19] further improves upon TCP Reno by changing some thresholds in the fast recovery algorithm and avoiding a scenario where multiple retransmits can occur.

However, as TCP additions and modifications are continuously proposed and implemented, TCP implementations are not that easily classified to these three original implementations. For example, the Linux TCP includes most standardized features, such as selective acknowledgments, timestamps, and window scaling, as well as more experimental features [44].

4 Problems with TCP in Wireless Networks

The performance of TCP is generally lower in wireless networks than in fixed. This is explained by the fact that TCP cannot distinguish problems that typically occur in wireless networks from congestion. The congestion control algorithms in TCP are based on the assumptions that data is lost mainly due to congestion and that data loss due to transmission errors is rare [24]. Therefore, data loss is interpreted as a signal of congestion in the network. Even in a wireless network, where data loss may not be related to congestion, data loss still signals congestion to the sender.

TCP segments may be lost if the radio conditions are poor and the link layer protocol provides a low reliability. After some retransmission attempts the link layer protocol gives up and leaves further error recovery to TCP. Handover events may also lead to data loss. A whole window of data may be lost due to handover. Data loss due to an unreliable link layer or a handover, may cause a timeout event followed by slow start or fast retransmit and fast recovery. In either case, the congestion control action taken by TCP is unnecessary. Directly after the loss event, the radio quality may become high again, and after handover data may be transmitted without problems to the new base station.

TCP may also misinterpret a sudden increase in the round trip time as data loss. If the delay is long enough for the retransmission timer to expire before an acknowledgment is received, then TCP misinterprets the delay as an indication of data loss due to congestion. The delayed data is unnecessarily retransmitted and TCP enters slow start. A highly variable round trip time can also lead to a large RTO, since the RTO is based both on estimates of the round trip time and on variations in the round trip time. If the RTO is large, then TCP reacts slowly to data loss. Variations in the round trip time can be caused by link level retransmissions of a wireless link. If the link layer frames that contain a TCP segment must be retransmitted because of a poor radio environment, then the whole segment is delayed. Round trip time variations may also be caused by handover or competing traffic. Queuing in routers, base stations, and other intermediate nodes may also lead to a long round trip time. A long round trip time may cause low throughput and underutilization of the network, since it takes a number of round trip times before the congestion window reaches the capacity of the network. TCP performance is degraded, especially for short lived flows, which transmits a small amount of data.

5 Proposed Optimizations

This section gives an overview of some optimizations that have been proposed to improve TCP performance over wireless networks. The proposed optimization are categorized into four groups: link layer, split connection, explicit notification, and end-to-end. Link layer optimizations for improved TCP performance are presented in Section 5.1. Section 5.2 describe how optimization at the transport layer can be achieved by splitting connections in an intermediate node between the wireless and the fixed network. In order to distinguish congestion related losses from losses in the wireless network explicit notifications can be used between an intermediate node and the end hosts. Examples of explicit notification approaches are given in Section 5.3. In Section 5.4, some end-to-end approaches are described. End-to-end approaches do not require any modifications of intermediate nodes, only of the end hosts.

5.1 Link Layer

The idea behind the proposals presented in this section is to improve TCP performance at the link layer. By using link level retransmissions locally over the wireless link, instead of end-to-end retransmissions, the probability of packet loss due to problems over the wireless part of the connection is decreased. The purpose of the proposals is to increase TCP performance over wireless links that provide low or no reliability.

5.1.1 Snoop

One of the first approaches to improve TCP performance over wireless links, Snoop, is presented in [8]. The Snoop scheme uses link layer retransmissions to improve the reliability of the wireless link and at the same time it actively tries to avoid unnecessary TCP retransmissions. Snoop is implemented as an agent in the base station. Snoop caches link layer frames and examines the contents of TCP headers, but it does not require TCP to run in the base station. A retransmission over the wireless link is triggered at the base station after a link layer timeout period or if a duplicate acknowledgment arrives from the mobile station, which is assumed to be the receiver. When the first duplicate acknowledgment arrives the base station retransmits the lost, or possibly out-of-sequence, frame. Further duplicate acknowledgments from the mobile station are dropped at the base station. This prevents the sender in the fixed network from performing fast retransmit and fast recovery while link layer retransmissions are performed over the wireless link.

5.1.2 WTCP (Ratnam and Matta)

WTCP [41] is similar to Snoop in that the base station performs retransmissions over the wireless link based on timeouts and duplicate acknowledgments. In addition, WTCP further improves TCP performance by using the timestamps option [23] and by adjusting the transmission window used over the wireless link. WTCP operates in the base station and the only modification required at the end hosts is that the timestamps option is enabled. The base station increments the timestamp in the TCP header for each retransmission that is required over the wireless link. This gives a more accurate RTT estimate. For acknowledgments from the mobile station, on the other hand, the base station leaves the timestamp unchanged. Furthermore, WTCP changes the upper limit of outstanding data that it allows over the wireless link. If a timeout occurs, then the upper limit of outstanding data is set to one segment, since the wireless link is assumed to be in a bad state. When acknowledgments come back to the base station again, the wireless link is assumed to be in a good state and the upper limit of outstanding data is set to the same value as the receiver window.

5.1.3 TULIP

In [38], the transport unaware link improvement protocol (TULIP) is proposed to improve TCP performance over half-duplex radio links, e.g. IEEE 802.11. TULIP is described as transport unaware, since it does not use information in the TCP header, but relies only on the protocol field in the IP header. TULIP provides a reliable service for TCP data segments, and an unreliable service for UDP and TCP acknowledgments. TULIP provides in-order delivery of packets, which prevents unnecessary fast retransmits.

5.1.4 Conclusion

The proposals presented above use link level retransmissions to minimize packet loss due to the wireless part of a connection. The link level proposals are intended for wireless links with very low persistent ARQ protocols. All the proposals preserve the end-to-end semantics of TCP, but TCP data and acknowledgments are required to pass through the same intermediate node. In the TCP-aware approaches, Snoop and WTCP, the base station is required to process information in the TCP header, which is based on the assumption that each TCP segment is encapsulated in a link layer frame. This is usually the case in WLANs, but not in most other wireless networks. If the link layer frame used over the radio interface is to small to encapsulate a TCP segment, as in many WWANs, then the link layer proposals could operate on the logical link control layer instead (provided that a TCP segment fits into a frame). Unfortunately, if TCP headers are encrypted [27], the TCP-aware approach will not work.

5.2 Split Connection

The idea of split connection approaches is to divide each TCP connection into two separate connections at an intermediate node. By splitting connections congestion related losses in the wired network are discriminated from wireless losses. Many split connection proposals use TCP, or modified versions of TCP, also over the wireless link. Any transport protocol that is suitable for wireless communication may be used over the wireless link. Another protocol stack may also be used over the wireless link, for example as is done in WAP 1.0 [51].

5.2.1 Indirect TCP

Indirect TCP (I-TCP) [4, 5] is intended primarily for WLANs with support for mobility at the IP level. An I-TCP connection between a mobile station and a fixed host is split at the base station into one connection over the wireless network and one over the fixed. When the mobile station moves and a handover occurs, TCP state is transferred to the new base station and the retransmission timer is reset. A consequence of this is that the fixed host will be less affected by delay caused by handover and link level retransmissions. Even though TCP is used also over the wireless link, the performance is improved as compared to having one connection all the way. This is explained by the fact that the round trip time becomes shorter. Therefore, the TCP sender in the base station can recover faster from data loss, especially if the connection over the fixed part experiences a non-negligible delay, but the wireless link still is the bottleneck.

5.2.2 MTCP and Selective Repeat Protocol

Two alternative split connection schemes are described in [52]. Both support mobility on top of a version of mobile IP. After a handover the TCP states of the connections are transmitted to the new base station. The first alternative, MTCP, uses standard TCP also over the wireless network. For the second alternative a selective repeat protocol (SRP) suitable for wireless communication is used instead of TCP between the base station and the mobile station. SRP is a UDP based protocol that performs flow and error control optimized for wireless links. Since SRP uses selective acknowledgments, it is possible to recover more than one segment each round trip time.

5.2.3 Conclusion

By splitting a TCP connection, data loss due to handover can be avoided if data and state information are forwarded to the new base station. Furthermore, with the split, it is possible to discriminate data loss over the wireless link from data loss that occurs over the fixed network. The proposals have similar disadvantages as the link layer schemes presented above. All data to and from the mobile station is required to pass the same base station (or some other intermediate node which splits the connection). TCP headers must be processed, which implies that the split connection schemes do not work with IPsec. In addition, the end-to-end semantics of TCP are not preserved. The presented proposals are intended for WLANs which support mobility on the IP level. However, split connection schemes may improve TCP performance also in other environments, since a shorter round trip time results is faster error recovery. An example of this is provided in [35], which evaluates a split connection proxy in WCDMA.

5.3 Explicit Notification

As TCP interprets all losses as congestion in the network, various explicit notification schemes have been proposed to enable the TCP sender to identify the real cause of a data loss. The idea is that an explicit notification is transmitted in order to inform the sender about data loss due to corruption in the wireless network.

5.3.1 ICMP Messaging

Instead of trying to hide problems due to a wireless link, [20] proposes that an explicit notification is transmitted to the TCP sender. The idea is that the intermediate node that connects the fixed and the wireless network has information about the wireless transmission success, and should be responsible for generating and transmitting explicit notifications. A new ICMP message type, ICMP-DEFER, is used for the explicit notification. If data is lost over the wireless network, an ICMP-DEFER is transmitted to the TCP sender, which postpones the expiration of the retransmission timer. This avoids conflicts between the link level retransmissions at the base station and end-to-end retransmissions. If a segment needs retransmission and an ICMP-DEFER has been received, then the cwnd is not reduced at the sender. The scheme is further improved by adding another ICMP message, ICMP-RETRANSMISSION. The base station transmits an ICMP-RETRANSMISSION to the TCP sender when the maximum number of retransmission attempts is reached and the packet is discarded.

5.3.2 Explicit Loss Notification

Explicit Loss Notification (ELN) [6] can be used if the mobile host is the sender. If data is lost over the wireless link, the base station transmits an ELN to the mobile station (the sender). The base station examines TCP headers and keeps track of the sequence numbers. A hole in the sequence space indicates that data must have been lost over the wireless link on its way from the mobile station to the base station. The fixed host (the receiver) transmits a dupack in response. The base station sets an ELN flag in the dupack header before the dupack is forwarded to the mobile station. The mobile station then retransmits the lost segment but does not reduce its congestion window.

5.3.3 Syndrome

Syndrome [13] is a modification of the base station to enable detection of packet loss over the wireless link. The base station counts the number of packets it relays, and includes this counter as a TCP option. This counter is called a ”syndrome”, and is used together with the sequence number to determine if a loss occurred in the wired or wireless part of the network. If there is a gap in the syndrome counter, this indicates that packets were lost on the wireless part. Gaps in sequence numbers, but not in the syndrome, indicate that segments must have been lost in the fixed network. Explicit loss notification is then used by the receiver to inform the sender about the loss.

5.3.4 Partial Acknowledgments

Another means to distinguish congestion from data loss due to a wireless link is to introduce new types of acknowledgments. The proposal described in [10] uses one additional acknowledgment, a partial acknowledgment that the base station transmits in response to data from a sender in the fixed network. Provided that no segments are lost, the sender receives two acknowledgments for each segment: a partial acknowledgment from the base station and a complete acknowledgment from the mobile host. If only the partial acknowledgment arrives, then the sender can conclude that data must have been lost over the wireless network and no congestion control action is required. If no acknowledgments arrive, then the most likely cause is data loss due to congestion.

A similar scheme, which introduces two new partial acknowledgments, is proposed in [15]. If the mobile station is the receiver, then the base station transmits a last hop acknowledgment in response to the fixed host. The scheme work also if the mobile station is the sender. In that case the base station transmits a first hop acknowledgment to the mobile station. As in the partial acknowledgment approach above, the sender receives two acknowledgments for each successfully transmitted segment, one from the receiver and one from the base station.

5.3.5 Conclusion

The explicit notification proposals have a different philosophy compared to most of the other proposals. The sender can distinguish congestion from data loss over the wireless network, since it receives information about the transmission status. This does not solve the problem with the higher unreliability of the wireless network, but the sender knows about it and is able to make a more informed decision. One disadvantage of the explicit notification schemes presented above is that they assume that the base station is TCP-aware. The only exception is ICMP messaging. Explicit notification assumes a low persistent link ARQ, such as in WLANs.

5.4 End-to-end

Optimizations may also be placed entirely in the end-points to avoid adding complexity to the network. The proposals in this section differ from each other, but the underlying idea that they have in common is to avoid congestion control action at the sender when data is lost for some other reason than congestion.

5.4.1 TCP Westwood

TCP Westwood [12] is a sender-side modification that improves the congestion control algorithm in TCP Reno. Instead of relying on slowly probing the available bandwidth until segments are dropped, as in standard TCP, the rate of incoming acknowledgments is used to determine the available bandwidth. TCP Westwood has been shown to give improvements over wireless links as well as over wired. Error recovery is faster than in standard TCP, since the bandwidth estimation is considered when values for cwnd and ssthresh are computed after data loss.

5.4.2 WTCP (Sinha et al.)

The Wireless Transmission Control Protocol (WTCP) [46] is aimed at improving the performance in WWANs such as CDPD. As TCP Westwood, WTCP uses a rate based approach to control the transmission rate. Inter-packet separation at the sender and the receiver is used as the primary metric for transmission rate calculations. This means that WTCP is more resilient to non-congestion related packet loss, thereby improving performance in wireless networks.

5.4.3 TCP Real

TCP Real [53] is a rate based scheme in which the receiver controls the transmission rate at sender. The receiver uses changes in the rate of incoming segments to compute the congestion window the sender should use. If the rate of incoming segments is decreasing, then this is taken as an indication of an increasing load in the network and therefore the cwnd should be decreased accordingly. After data loss, the cwnd is adjusted to the network conditions sooner than in standard TCP, since the receiver includes estimates of the cwnd in the acknowledgments that go back to the sender.

5.4.4 Freeze-TCP

Freeze-TCP [21] is a mechanism to improve the performance of TCP in wireless environments in which handover frequently occurs. By exploiting the properties of the advertised receiver window, a TCP connection can be frozen. If the receiver sets the receiver window to zero, then the sender leaves its cwnd unchanged until the receiver advertises a new receiver window. Just before handover occurs, the mobile station freezes the TCP connection by advertising a receiver window of zero. This prevents segments from getting lost and unnecessary congestion control action to be taken by the sender. When the handover is completed, the receiver sends an acknowledgment which opens up the window again. The transmission then continues at the same rate as before.

5.4.5 Delayed Dupack

The delayed dupack scheme proposed in [50] is an end-to-end scheme that imitates Snoop. The third and subsequent dupacks are delayed while the base station performs link level retransmissions. This prevents the sender (fixed host) from taking congestion control action while the base station retransmits data over the wireless link. The third and subsequent dupacks are delayed for a predetermined time interval. During this time the receiver may receive the missing data and transmit a cumulative acknowledgment instead. The delayed dupacks are then discarded. If the missing segment is not received, then the delayed dupacks are instead released when the timer expires.

A disadvantage with the delayed dupacks scheme is that the dupacks are delayed also in case of congestion, and hence error recovery is delayed. Explicit loss notification to the receiver (ELNR) [34] is proposed as an enhancement to the delayed dupacks scheme. The mobile host delays dupacks unless it receives an ELNR which indicates that the network may be congested. The base station transmits an ELNR to the mobile host if it receives a packet out-of-sequence. Upon receipt of an ELNR the mobile host does not delay any dupacks. When the base station has received the missing packets, the base station transmits an explicit delayed dupack activation notification (EDDAN) to the mobile host. The EDDAN indicates that the network is no longer congested, and subsequent dupacks should be delayed until a new ELNR arrives.

5.4.6 Eifel

The Eifel algorithm described in [31, 32, 30] allows the sender to detect whether an already initiated error recovery action is necessary or not. When the first acknowledgment that covers previously unacknowledged data arrives, the sender can determine if this is an acknowledgment of the original segment or of a retransmission. The algorithm uses the timestamps option to match acknowledgments with segments. If the first acknowledgment is for the original segment, then the retransmission is spurious and there is no reason for the sender to reduce the transmission rate. The original packet is not lost due to congestion, but it has only been delayed before it arrived at the receiver.

5.4.7 Conclusion

The end-to-end proposals are based on various ideas. The rate based approaches try both to avoid congestion and to recover quickly from random losses over the wireless network. Freeze-TCP prevents data loss by making a pause in the data transfer during handover. The Eifel approach, on the other hand, limits performance degradation once delay has already been misinterpreted as a sign of congestion. End-to-end proposals are advantageous since the end-to-end semantics of TCP are preserved. The data transferred between the end hosts is not required to pass the same intermediate node. No additional processing is required in the network and IPsec can be used.

5.5 Summary

The optimization proposals presented above, as illustrated by many different types of approaches, have been suggested to improve TCP performance over wireless networks.

The link layer approaches use link layer retransmissions to increase the reliability provided to TCP. The requirements on the link layer service may vary depending on the application. This is, for example, considered in UMTS, since the radio link protocol allows configuration of many parameters, such as the maximum number of retransmissions.

There are many types of split connection proposals, only some of which are described in this report. Performance improvements are often significant, but at the cost of violating the end-to-end semantics of TCP. The intermediate node which divides the connection must process data up to the transport layer and all TCP segments belonging to the same connection are required to pass that node. The TCP-aware link layer proposals have similar limitations as the split connection proposals. However, the link layer proposals preserve the end-to-end semantics of TCP.

Most of the explicit notification proposals require TCP-awareness of the intermediate node that is responsible for transmitting explicit notifications. These proposals seem to be intended for wireless links with low reliability. A link ARQ with higher persistence would reduce or eliminate data loss due to an unreliable wireless link. However, another application of explicit notifications could be to signal to the TCP sender that a handover has been initiated.

End-to-end proposals are based on the idea that complexity belongs in the end hosts rather than in the network. The end-to-end semantics of TCP are preserved. Intermediate nodes do not have to be TCP-aware, as in many of the other proposals. The data transferred between the end hosts is not required to pass the same intermediate node.

References

[1] M. Allman, H. Balakrishnan, and S. Floyd. RFC 3042: Enhancing TCP's loss recovery using limited transmit, January 2001.

[2] M. Allman, S. Floyd, and C. Partridge. RFC 3390: Increasing TCP's initial window, August 2002.

[3] M. Allman, V. Paxson, and W. Stevens. RFC 2581: TCP congestion control, April 1999.

[4] A. Bakre and B. R. Badrinath. I-TCP: Indirect TCP for mobile hosts. *15th International Conference on Distributed Computing Systems*, May 1995.

[5] A. V. Bakre and B. Badrinath. Implementation and performance evaluation of Indirect TCP. *IEEE Transactions on Computers*, 46(3), March 1997.

[6] H. Balakrishnan and R. Katz. Explicit loss notification and wireless web performance. *Proceedings Globecom Internet Mini-Conference, Sydney, Australia*, November 1998.

[7] H. Balakrishnan, V. N. Padmanabhan, S. Seshan, and R. H. Katz. A comparison of mechanisms for improving TCP performance over wireless links. *IEEE/ACM Transactions on Networking*, 5(6):756–769, December 1997.

[8] H. Balakrishnan, S. Seshan, and R. H. Katz. Improving reliable transport and handoff performance in cellular wireless networks. *ACM Wireless Networks*, 1(4), December 1995.

[9] C. Bettstetter, H.-J. Vögel, and J. Eberspächer. GSM phase 2+ general packet radio service GPRS: Architecture, protocols, and air interface. *IEEE Communication Surveys*, 2(3), 1999.

[10] S. Biaz, M. Mehta, S. West, and N. H. Vaidya. TCP over wireless networks using multiple acknowledgements. *Technical Report 97-001, Computer Science, Texas A&M University*, January 1997.

[11] R. Cáceres and L. Iftode. Improving the performance of reliable transport protocols in mobile computing environments. *IEEE Journal on Selected Areas in Communications*, 13, June 1995.

[12] C. Casetti, M. Gerla, S. Mascolo, M. Y. Sanadidi, and R. Wang. TCP Westwood: Bandwidth estimation for enhanced transport over wireless links. *Proceedings of ACM Mobicom, Rome, Italy*, pages 287–297, July 2001.

[13] W.-P. Chen, Y.-C. Hsiao, J. C. Hou, Y. Ge, and M. P. Fitz. Syndrome: a light-weight approach to improving TCP performance in mobile wireless networks. *Wireless Communications and Mobile Computing*, (2):37–57, 2002.

[14] M. C. Chuah, B. Doshi, S. Dravida, R. Ejzak, and S. Nanda. Link layer retransmission schemes for circuit-mode data over the CDMA physical channel. *Mobile Networks and Applications*, 2(2):195 – 211, October 1997.

[15] J. A. Cobb and P. Agrawal. Congestion or corruption? A strategy for efficient wireless TCP sessions. *IEEE Symposium on Computers and Communications, Alexandria, Egypt*, pages 262–268, 1995.

[16] ETSI. Digital cellular telecommunications system (Phase 2+); Radio Link Protocol (RLP) for data and telematic services on the Mobile Station - Base Station System (MS - BSS) interface and the Base Station System - Mobile-services Switching Centre (BSS - MSC) interface (GSM 04.22 version 7.0.1 Release 1998). *ETSI*, 1999.

[17] ETSI. Digital cellular telecommunications system (Phase 2+); General Packet Radio Service (GPRS); Service description; Stage 2 (GSM 03.60 version 6.7.0 Release 1997). *ETSI*, 2000.

[18] G. Fairhurst and L. Wood. RFC 3366: Advice to link designers on link Automatic Repeat reQuest (ARQ), August 2002.

[19] S. Floyd. RFC 2582: The NewReno Modification to TCP's Fast Recovery Algorithm, April 1999.

[20] S. Goel and D. Sanghi. Improving TCP performance over wireless links. *Proceedings of TENCON 98 - IEEE Region Ten Conference on Global Connectivity in Energy, Computer Communication and Control*, 1998.

[21] T. Goff, J. Moronski, D. Phatak, and V. V. Gupta. Freeze-TCP: A true end-to-end enhancement mechanism for mobile environments. *Proceedings of IEEE INFOCOM*, March 2000.

[22] A. Gurtov and S. Floyd. Modeling wireless links for transport protocols. *To appear in ACM CCR*, 2004.

[23] V. Jacobson, R. Braden, and D. Borman. RFC 1323: TCP Extensions for High Performance, May 1992.

[24] V. Jacobson and M. J. Karels. Congestion Avoidance and Control. *Proceedings, SIGCOMM '88 Workshop*, pages 314–329, 1988.

[25] P. Karn. The Qualcomm CDMA digital cellular system. In *Proceedings of the USENIX Mobile and Location-Independent Computing Symposium*, pages 35–39, Aug. 1993.

[26] P. Karn. IP data services over CDMA digital cellular. *Mobile Computing and Communications Review*, 4(4):30–35, October 2000.

[27] S. Kent and R. Atkinson. RFC 2401: Security architecture for the internet protocol, November 1998.

[28] Y.-B. Lin and I. Chlamtac. *Wireless and Mobile Network Architectures*. John Wiley & Sons, Inc., 2000.

[29] R. Ludwig. Eliminating inefficient cross-layer interactions in wireless networking. *Ph.D. Thesis, Aachen University of Technology, Germany*, April 2000.

[30] R. Ludwig and A. Gurtov. The Eifel Response Algorithm for TCP, 2003.

[31] R. Ludwig and R. H. Katz. The Eifel algorithm: Making TCP robust against spurious retransmissions. *ACM Computer Communication Review*, 30(1):30–37, January 2000.

[32] R. Ludwig and M. Meyer. RFC 3522: The Eifel Detection Algorithm for TCP, 2003.

[33] M. Mathis, J. Mahdavi, S. Floyd, and A. Romanow. RFC 2018: TCP selective acknowledgment options, Oct. 1996.

[34] M. Mehta and N. H. Vaidya. Delayed duplicate acknowledgements: A proposal to improve performance of TCP on wireless links. *Tecnical Report (Texas A&M University)*, February 1998.

[35] M. Meyer, J. Sachs, and M. Holzke. Performance evaluation of a TCP proxy in WCDMA networks. *IEEE Wireless Communications*, 10(5):70 – 79, October 2003.

[36] A. Mishra, M. Shin, and W. Arbaugh. An empirical analysis of the IEEE 802.11 MAC layer handoff process. *SIGCOMM Comput. Commun. Rev.*, 33(2):93–102, 2003.

[37] M. Mouly and M.-B. Pautet. The GSM system for mobile communications. *Cell & Sys*, 1992.

[38] C. Parsa and J. Garcia-Luna-Aceves. Improving TCP performance over wireless networks at the link layer. *ACM Mobile Networks and Applications Journal, vol 5, issue 1.*, 2000.

[39] V. Paxson and M. Allman. RFC 2988: Computing TCP's Retransmission Timer, November 2000.

[40] J. Postel. RFC 793: Transmission control protocol, Sept. 1981.

[41] K. Ratnam and I. Matta. WTCP: An efficient transmission control protocol for networks with wireless links. *Proc. Third IEEE Symposium on Computers and Communications (ISCC '98), Athens, Greece*, June 1998.

[42] A. K. Salkintzis. A survey of mobile data networks. *IEEE Communication Surveys*, 2(3), 1999.

[43] B. Sarikaya. Packet mode in wireless networks: Overview of transition to third generation. *IEEE Communications Magazine*, September 2000.

[44] P. Sarolahti and A. Kuznetsov. Congestion control in Linux TCP. *USENIX 2002*, pages 49–62, June 2002.

[45] A. D. Simone and S. Nanda. Wireless data: systems, standards, service. *Wireless Networks*, 1(3):241–253, March 1995.

[46] P. Sinha, T. Nandagopal, N. Venkitaraman, R. Sivakumar, and V. Bharghavan. WTCP: A reliable transport protocol for wireless wide-area networks. *Wireless Networks*, 8(2):301–316, March 2002.

[47] W. Stallings. Mobile IP. *The Internet Protocol Journal*, 4(2):2–14, June 2001.

[48] W. Stallings. *Wireless Communications and Networks*. Prentice Hall, 2001.

[49] M. Taferner and E. Bonek. *Wireless Internet Access over GSM and UMTS*. Springer-Verlag, 2002.

[50] N. Vaidya, M. Mehta, C. Perkins, and G. Montenegro. Delayed duplicate acknowledgements: A TCP-unaware approach to improve performance of TCP over wireless. *Technical Report, Computer Science Dept., Texas A&M University*, February 1999.

[51] WAP Forum Ltd. Wireless Application Protocol - architecture specification. *http://www.wapforum.org*, page 20, April 1998.

[52] R. Yavatkar and N. Bhagwat. Improving end-to-end performance of TCP over mobile internetworks. *Workshop on Mobile Computing Systems and Applications (Mobile 94)*, December 1994.

[53] C. Zhang and V. Tsaoussidis. TCP Real: Improving real-time capabilities of TCP over heterogeneous networks. *Proceedings of the 11th IEEE/ACM NOSSDAV*, 2001.