

Modeling Wireless Links for Transport Protocols

Andrei Gurtov^{*}
University of Helsinki
gurtov@cs.helsinki.fi

Sally Floyd
ICIR
floyd@icir.org

ABSTRACT

Wireless links have intrinsic characteristics that affect the performance of transport protocols; these include variable bandwidth, corruption, channel allocation delays, and asymmetry. In this paper we review simulation models for cellular, WLAN and satellite links used in the design of transport protocols, and consider the interplay between wireless links and transport. We argue that the design and evaluation of transport protocols can be improved by providing easily available models of wireless links that strike a balance between realism, generality, and detail.

There is an ongoing tussle between wireless link design and transport protocol design, with papers about how wireless link designers should take into account the dynamics of TCP, and other papers about how TCP and other transport protocols can be designed or modified for better performance over current wireless link technologies. In this paper we consider how appropriate models for wireless links can help in this tussle, and in the general design and evaluation of transport protocols over wireless links¹.

1. INTRODUCTION

Wireless links are an important part of the Internet today; in the future the number of wireless or mobile hosts is likely to exceed the number of fixed hosts in the Internet. Internet access using wireless LANs and cellular links is growing particularly quickly. As an example, in Japan over 40 million users are accessing data services from mobile phones [56]. The main classes of wireless links are wireless LANs, wide-area cellular links, and satellite links. In this paper we focus on simulation models of these wireless links for unicast transport protocols. We do not consider ad-hoc and sensor wireless networks, as this is a separate modeling area on its own.

It is well known that the presence of a wireless link can significantly

^{*}This paper was written when the author was visiting ICIR.

¹This work was supported in part by National Science Foundation under Grant No. 0230921. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

affect the performance of end-to-end transport protocols [61]. Congestion control in today's Internet is based on an assumption that almost all packet losses result from congestion [41]. Packet losses on wireless links that are from corruption rather than congestion violate this assumption. While many link technologies include FEC (Forward Error Correction) and local retransmission for addressing corruption at the link layer, these mechanisms can introduce their own complications [15]. A high variability of bandwidth and delay on wireless links can reduce the performance of transport protocols (including those transport protocols that consider increased delay as an indication of congestion). Furthermore, because cellular and satellite links have high latency, connections traversing them can be penalized by the bias of TCP congestion control toward flows with high RTT (Round-Trip Time) [66]. Mechanisms for reducing the bit error rate of cellular links, such as interleaving and FEC, are the main sources of high latency on cellular links. For some link technologies, large link-level queues contribute to unnecessarily high delay [51, 33].

There is a clear benefit of using the same protocol stack for fixed and wireless links [40]. It enables interoperability between users with an adjacent wireless link and the rest of the Internet. Therefore, existing and new transport protocols should be designed for good performance over both wireless and fixed links. At the same time, wireless links should be designed to minimize negative effects on transport protocols [43].

There is a large body of work on the evaluation and design of transport protocols over wireless links [38, 39, 7]. However, the diverse set of models makes it difficult to compare results from different studies. One goal of this paper is to make it easier for researchers unfamiliar with details of wireless links to understand their impact on transport protocols, and to use wireless link models in their research. We have included our scenarios into ns-2 [62, 30] so that they are accessible to the wider Internet community.

The assumptions made about wireless link characteristics can strongly affect the evaluation of transport protocols. As an example, let's consider the issue of an abrupt increase in delay triggering a spurious timeout in TCP. Link-level error recovery was believed as one possible source of this form of delay variation [18]. However, a measurement study of a GSM radio link in poor radio conditions found that spurious TCP timeouts are rare [51]. This is because the delay variation had a pattern of moderate jitter in inter-packet arrivals rather than sharp delay spikes. This moderate jitter results in an inflated TCP retransmit timer, preventing spurious timeouts while prolonging loss recovery. In contrast, a measurement and simulation study found that handovers introduce a sharper pattern

of delay jitter and can indeed trigger spurious timeouts [29]. Therefore, when evaluating a transport protocol using simulation, it is important to select parameters that accurately reflect real-world scenarios.

In modeling wireless links for transport protocols, the researcher would like a model that is as simple as possible for the given purpose, but no simpler [25, 34]. Mechanisms such as link-level retransmissions and inter-system mobility are complex and difficult to completely represent in a model. The purpose of the models discussed in this paper is only to evaluate the effect of link-level mechanisms on end-to-end transport protocols and higher layers. In this case, simply changing link characteristics and introducing packet losses or delays to traffic are often sufficient for understanding transport protocol performance in the presence of the modeled wireless links.

Models of wireless links are used not only for the design and evaluation of transport protocols, but also for the consideration of new mechanisms of implicit or explicit communication between link layers and transport protocols [20]. This is discussed in more detail later in the paper.

The rest of the paper is organized as follows. Section 2 summarizes wireless technologies considered in this paper and our assumptions about topologies, traffic, and performance metrics used in evaluating transport protocols. Section 3 shows why modeling matters when considering transport protocol performance over wireless links. In Section 4 we discuss detailed models for corruption, delay variation, reordering, channel allocation, bandwidth variation, and asymmetry. Section 5 discusses modeling of buffering in wireless networks. Section 6 describes effects and modeling of handovers. While up to this point the paper only considers modeling of existing wireless links, in Section 7 we discuss the interplay between designing wireless links and designing transport protocols. Section 8 presents conclusions and plans for future work.

2. ESSENTIAL ASPECTS OF MODELS

We start this section by describing the types of wireless links considered in the rest of the paper. For each link type, we describe the link characteristics, and discuss relevant topologies, traffic models, and performance metrics.

2.1 Types of Wireless Links

In this paper we consider the most commonly used types of wireless links: cellular, WLAN and satellite.

Cellular. Most common cellular links are provided today by GPRS (General Packet Radio Service) and CDMA2000 systems, and in the future possibly by UMTS (Universal Mobile Telecommunication System) [63]. The bandwidth of such links is in the range of 0.01-1 Mbps, with high one-way latency of 0.1-0.5 seconds. The coverage radius of a single cell varies from several hundred meters in urban areas up to 30 km in rural areas. In this paper we use GPRS as an example of a cellular link. A GPRS link typically has 40 kbps bandwidth and 400 ms latency in downlink and 10 kbps, 200 ms in uplink.

Because of the challenging radio propagation environments that cellular links face, they are typically heavily protected by forward error correction and link-layer retransmissions [51]. Furthermore, due to high link round-trip times, acquiring a channel access can cause considerable delays. Every packet may require a new chan-

nel allocation. In addition to low bandwidth, battery power preservation is a major challenge for transport protocols. These factors are discussed in detail in Section 4.

Wireless LANs. The most commonly used WLAN today is IEEE 802.11b with bandwidth of 2-11 Mbps [63]. In general, WLANs have a low latency of 3-100 ms and bandwidth in the range of 1-50 Mbps. WLAN uplink and downlink channels are not independent as in cellular or satellite, but compete with each other for shared bandwidth. The coverage radius of a single base station varies from tens to hundreds of meters.

The link error control of 802.11b is tightly coupled with the MAC mechanism. There are at most three retransmission attempts per data frame [54]. Packet fragmentation is supported for higher efficiency of error recovery, but it is not commonly used.

Satellite. In general, satellite links are characterized by high latency in the range of 50-300 ms and bandwidth of 0.01-50 Mbps. Today, satellite links are mostly provided by a fixed GEO satellite. Such links typically have a latency of 270 ms, downlink bandwidth of 40 Mbps, and uplink bandwidth of 1 Mbps [35]. There is tremendous variation in capacity provided by satellite links; the uplink bandwidth might be only 64 kbps for VSAT terminals. Modern satellite links are generally error-free except for occasional fades. A single GEO satellite can cover an entire continent.

Another type of satellite link is provided by constellations of mobile LEO satellites such as Iridium or Globalstar. When a satellite moves out of coverage, the link is handed off to another satellite. A study of TCP Westwood [52] simulates a LEO link with 100 ms latency, 1 Mbps bandwidth, a 0.1% loss rate, a handover interval of 4 seconds and a duration of 100 ms. All packets in flight are lost during a handover.

2.2 Topologies

The performance of transport protocols can be affected by the number and location of wireless links in the path. Latencies and error loss rates of multiple wireless links add up, making loss recovery and congestion control more difficult [66]. Possible topologies include a single wireless link located at the end or in the middle of the path, or multiple wireless links located at each end or in the middle.

A common scenario is that of a wireless link used as a last-hop link to a host. It is natural that mobile users deploy a wireless link to connect to the rest of the Internet.

It is not uncommon for there to be a second wireless link connecting a mobile phone to the laptop. Infrared and Bluetooth links are used for this purpose. Except for increased latency, such links do not typically affect transport protocols. Such links are typically used in close range with a good quality signal. However, transport protocols are affected in abnormal conditions when an infrared link is misaligned or a Bluetooth link experiences interference from WLAN.

When multiple wireless links are present in a path, it is possible to have packet losses due to interference between these links [17]. To avoid wasting wireless bandwidth, flow control is often used between multiple wireless links.

Currently, scenarios with wireless links at both ends of the path are

not common. This is partly due to security concerns with allowing someone to connect to a mobile host. For applications involving two mobile hosts, such as instant messaging, there are two transport connections via a server in the fixed network.

In the future, telephone calls may be implemented using VoIP. Then, a scenario with two wireless links at both ends will be more common. For VoIP a resource reservation is often made to emulate a circuit-switched link, with UDP used as the transport protocol. However, there has been some interest in the design of transport protocols for those VoIP calls that are carried as best-effort traffic over the public Internet [45, 24], and for this, a topology model with a wireless link at each end could be of interest.

Satellite links are often used as an ISP link to remote areas such as Hawaii. VSAT satellite links are also used in Africa and elsewhere to connect directly to the U.S. or Europe [53]. There is a substantial industry in the U.S. supplying VSAT services to, for example, gas stations for credit card verification. Satellite links can be also used for home access (e.g. DIRECWAY), with bandwidth scaled down by ten compared to the ISP case, and with lower levels of statistical multiplexing. When multiple satellite links are present in a path, for example due to inter-satellite forwarding in a constellation, they can generally be represented as a single link because inter-satellite links are seldom congested [35].

WLAN links are almost always used as a last-hop wireless link. Occasionally, a WLAN is used as a point-to-point link to connect two LANs. From a performance viewpoint, it generally doesn't matter whether users connect to the WLAN directly, or through a LAN [13]. However, in some cases, such as exploring communication between the transport and link layers, it can matter whether the model has a user directed attached to the WLAN, or has a LAN between the user and the WLAN.

2.3 Wireless Traffic

A traffic model can have significant effects on simulation results. In particular, drawing conclusions on performance from one-way bulk TCP transfers is often misleading [66]. Optimizations that enhance TCP performance in one direction can fail for popular bi-directional applications such as web browsing, or when multiple wireless links are present in the path.

Typically, mobile users transfer more data downlink than uplink. Thus, a common assumption is that the receiver is adjacent to the wireless link. Although wireless links are used for generic Internet access, the traffic mix carried over them may not be the same as on fixed links. We conjecture that, due to the low speed and high cost of cellular links, users are not likely to run certain types of applications, such as peer-to-peer file sharing. Measurements of GPRS traffic [44] showed that over 95% of all sessions transfer less than a megabyte of data, which is clearly below a typical object size for peer-to-peer applications [28].

Cellular users may mostly be visiting servers specifically set up for mobile services. This is an intrinsic consequence of display limitations of small communication devices that cannot show the contents of the "wired" Internet in a readable form. Such servers can use custom applications such as Wireless Application Protocol (WAP) that are uncommon in the wired Internet. WAP traffic differs considerably from HTTP traffic [60]. A wireless traffic model has been developed for choosing among proposals of cellular technologies [1, Section B.1.2.2].

Cellular links can be used for special-purpose data communications such as machine-to-machine control interface. According to the ETSI (European Telecommunications Standards Institute) ad-hoc evaluation guidelines for GPRS, traffic models have been implemented for the generation of E-mail (FUNET), fleet management (Mobitex), and railway traffic [12].

In general, there has been insufficient measurements of cellular data traffic. This is due to user privacy concerns and the high cost of performing measurements. Because modern cellular links resemble to a large extent a modem dial-up link, measurement studies of such links [49] are sometimes used as an estimate for cellular link traffic.

The traffic mix on a WLAN link is similar to that of a typical LAN, but with fewer large flows due to the lower bandwidth [5]. Traffic on a satellite link resembles traffic of other ISP links, perhaps with less interactive traffic due to the high latency of the link.

2.4 Performance Metrics

When evaluating transport protocols over wireless links, we first look at the usual metrics for evaluating performance. These include throughput, delay, fairness and dynamics.

There are a number of metrics that are of particular importance for wireless links. One important metric is goodput, defined as the fraction of useful data from all data delivered. The goodput affects the energy efficiency of the transport protocol in using the battery power of the mobile terminal. High goodput also means an efficient use of the radio spectrum and lower interference to other users. Last but not least, mobile users are often charged based on the amount of data they send over a wireless link, so improved goodput results in reduced expense.

3. ON THE NEED FOR BETTER MODELS

In this section we consider the need for better models in evaluating transport protocol performance over wireless links, and discuss the ways that the choice of the models can critically affect the results of the research.

One problem is the use of models that are not realistic. The use of unrealistic models includes the following:

- Running TCP over a wireless link with 40% packet loss rate due to errors [61]. This scenario appears impractical because most existing wireless links use local error recovery. Section 4.1 discusses models for packet corruption in wireless links.
- Evaluating Active Queue Management for wireless links by focusing on the start-up behavior of TFRC [21, 8]. This is an impractical scenario because TFRC is designed for long-lived flows such as streaming media, and therefore is not likely to be dominated by its slow-start behavior. Section 2.3 briefly discussed traffic scenarios for wireless links.
- Distinguishing corruption losses from congestion losses by making tight assumptions about inter-packet delays [9]. Heuristics based on simulation scenarios with perfectly-spaced packet inter-arrival times do not necessarily transfer well to the high variability of the real world.
- Modeling the TCP retransmission timer as a constant [18]. This study argues that spurious timeouts due to link layer




Figure 1: TCP over a duplex link of 1 Mbps (left), a duplex link of 0.65 Mbps (middle), and a detailed WLAN model at 1 Mbps (right).




Figure 2: Two reverse TCP flows over a 0.65 Mbps duplex link (left), an Ethernet link at 1 Mbps (middle), and a detailed WLAN model at 1 Mbps (right).

retransmissions is a severe problem but does not properly account for inflation of the transmit timer due to delay jitter.

- Using older deprecated TCP versions. Although vast majority of TCPs in the Internet use either SACK or NewReno for loss recovery, many simulation studies still use Reno [66] or even Tahoe TCP [37]. Possibly, the use of old TCP versions is due to the lack of support for modern TCPs in some network simulators.
- Modeling WLAN as a duplex link [36, 48, 9]. While in some cases this might be appropriate, there are also scenarios with bi-directional data where modeling WLAN as a duplex link could give misleading results.

As an example of the problems of modeling WLAN as a duplex link, Figure 1 compares a single, long-lived TCP-SACK connection over a duplex link and over a detailed WLAN model. Clearly, when a duplex link is used with the same bandwidth as WLAN link rate, simulations produce overly high TCP throughput. It is a well-known fact that quoted link-layer bandwidth typically corresponds to 50-80% throughput at the transport layer. When a duplex link of 65% of WLAN link-rate is used (middle graph), both models produce similar throughput. In a detailed WLAN model the packet transmission times are variable, while they are constant over a duplex link. In many cases it is feasible to omit this detail.

Figure 2 compares two single, long-lived TCP connections in reverse directions over a duplex link, an Ethernet model, and over a detailed WLAN model. As expected, using a duplex link in this case produces higher throughput than using a detailed WLAN model. Using an Ethernet model of a shared link of the same bandwidth as a detailed WLAN model gives more conservative results. To confirm that, we ran experiments changing the starting and stopping times of TCP flows. Generally, using a duplex link of 0.65

Mbps in these simulations gives a total TCP throughput that is one-third higher than using a detailed model of a 1 Mbps WLAN link. Using a model of an Ethernet link at 1 Mbps gives a total TCP throughput that is one-third lower than with the detailed WLAN model of the same bandwidth. One possible reason for using an Ethernet link to represent a WLAN link in the model is significantly less complexity required to set up an Ethernet link in the current ns-2 version.

A second problem occurs when the models are realistic, but explore only a small corner of the parameter space. Examples include the following:

- Evaluating Active Queue Management (AQM) for wireless links using a single flow [21, 58]. While it is true that in wireless networks the level of statistical multiplexing is often lower than in the rest of the Internet, AQM proposals should also be evaluated with a richer traffic mix than a single flow.
- Assuming that the wireless link is the only bottleneck and not trying the opposite case [58]. Even if this assumption may hold in most cases, there are certainly cases where the congestion will be elsewhere in the network. A mechanism designed based on assumptions of the wireless link as the only bottleneck could show severe misbehavior in other environments, possibly causing persistent congestion.
- One proposed mechanism for channel allocation in WLANs is to allocate a channel for reverse traffic immediately after a packet is transmitted in the forward direction [65]. The purpose would be to avoid collisions of TCP ACKs with data packets. In the simulations used to evaluate this proposal, there is an assumption that there will always be a packet available for transmission in the reverse direction. This may

not be a valid assumption, because of delayed ACKs and of possible time delays between a TCP receiving a data packet and generating an ACK.

A third problem is of models that are overly realistic, in that they focus on the transient implementation flaws of real world systems, and ignore the inherent underlying dynamics:

- Taking a bad implementation of GPRS buffering as a fundamental characteristic [14]. While simulation models should stay close to the real world, sometimes we want them to behave differently. It is because we know that something in the real world is faulty, but is difficult to fix. One such example is gigantic buffers in cellular networks.

A fourth problem with the use of models is a lack of reproducibility, when the researcher does not make the model specifications easily available.

- Not giving enough detail (e.g. using 5% loss rate but not saying if it applies in both directions) [9]. Lack of detail in model specification and parameters make it difficult to other researchers to understand and reproduce the results.

4. MODELING LINK CHARACTERISTICS

In this section we discuss specific characteristics of wireless links, including error losses, delay and bandwidth variation, packet re-ordering, on-demand resource allocation, and asymmetry. For each characteristic, we discuss its effect on transport protocols, and considerations for modeling. We draw upon the considerable body of knowledge in this area, including the work in the IETF to document the effects of link characteristics on transport protocols [40, 4].

4.1 Error Losses and Corruption

Effect on transport protocols: The negative effect of non-congestion-related losses on transport protocols is well-known; packet losses due to errors can incorrectly reduce the sending rates of transport protocols [7]. For reliable transport protocols, such as TCP, bursty losses can also trigger lengthy retransmission timeouts. On the other hand, bursty losses sometimes allow for higher throughput than an equivalent number of non-bursty losses, because a burst of losses is treated as a single congestion event by TCP.

Presence in current and future wireless links: Error losses are not a main concern for modern wireless links because of widespread use of FEC and link layer retransmissions. However, current link layers perform only a limited number of retransmission attempts (typically three) per data block. Therefore, in very poor radio conditions, error losses are still possible. For example, a measurement study of 802.11b WLAN reports packet losses under poor radio conditions [54]. Mobility is a more common source of error losses in wireless networks. Handovers between base stations of the same network or between different networks both cause packet losses [33, 59].

How to model: Error losses can be modeled by dropping packets according to a per-packet, per-bit or time-based loss probability. Several studies [20, 9] model error losses using a uniform per-packet error loss probability from 0 to 1. To reflect the bursty nature

of wireless errors, a Gilbert model with erroneous and error-free states is typically used. More complex approaches can model non-stationary distributions of errors and delays [47, 46]. An important modeling issue is whether losses occur independently in both directions or, which is more realistic, in a correlated way [66].

If losses are due to data corruption on the link, then the error module should be placed at the end of the link to reflect the use of bandwidth by dropped packets. However, if losses are due to handovers, drops should occur before the link, because such packets do not use link bandwidth.

4.2 Delay Variation

Effect on transport protocols: A sudden increase in the link delay (a delay spike) can have a negative impact on transport protocols. Transport protocols use timers for loss recovery (TCP) and for congestion control (TCP, TFRC) that are set according to the path RTT. Abrupt delay spikes can trigger spurious timeouts that cause unnecessary retransmissions and false congestion control [50]. Persistent delay variation can inflate the retransmission timeout, increasing the wait before a dropped packet is recovered.

Some experimental delay-based congestion control mechanisms use end-to-end delay to control the transmission rate of the sender. These include TCP Vegas [11], Fast TCP, TCP Westwood² [52], and to some extent TFRC [23]. Delay-based congestion control mechanisms assume that an increase in the end-to-end delay indicates packet queuing and growing congestion in the network. When delay variation occurs for reasons other than queuing delay, delay-based congestion control mechanisms can unnecessarily reduce the sending rate.

In many cases delay variation in wireless links does indicate congestion, for example, when the link layer retransmits lost data. However, there are cases when packets are delivered in a burst after a delay ends (due to the head-of-line blocking) and there is no congestion.

Rate-based protocols can perform badly in the presence of abrupt delay variation, compared to window-based protocols, when the increased delay is not due to increased propagation delay. For window-based protocols, the sending rate is limited by ACK-clocking. Rate-based protocols keep the same sending rate after a change in the round-trip time, possibly creating unnecessary congestion.

Presence in current and future wireless links: Delay spikes in wireless links can occur due to link-layer error recovery, handovers, or scheduling [29]. Error recovery using link-layer retransmissions typically introduces delay jitter without sharp spikes [51]. An exception can be a sudden change in radio conditions, such as entering a tunnel. Handovers cause delay spikes of up to several seconds and may occur once per minute when driving in an urban area [33]. Scheduling can cause delay spikes, for example, when circuit-switched calls take away radio channels. Also, there are radio schedulers that give the channel to the user with the currently-best radio conditions [15]. Section 7.3 discusses some of the approaches for reducing the delay variation introduced by wireless technologies. Delay variation can occur only in one direction, for instance if a path is asymmetric, but we expect both directions to be affected in most cases.

²Westwood uses the minimum RTT of the path to set the congestion window after a packet loss, but does not use RTT for detecting congestion.

How to model: Delay spikes can be modeled by suspending the data transmission on a simulated link [32]. Delay variation due to handovers in a cellular network can have the interval of 40-400 seconds with duration of 3-15 seconds [33]. There is no sufficient measurement available of delay jitter introduced by link retransmissions.

4.3 Packet Reordering

Wireless links can introduce packet reordering during link-level error recovery using retransmissions. Out-of-order delivery is shown to be beneficial for real-time applications by decreasing delay, especially in case of multiple wireless links in the path [66].

Effect on transport protocols: Significant packet reordering can incorrectly trigger packet retransmissions and congestion control responses for TCP and related transport protocols [10, 50].

Presence in current and future wireless links: To our knowledge, WLAN links do not introduce reordering. Cellular links include an option for out-of-order delivery. Allowing reordering is attractive on satellite links with a high bandwidth-delay product, because it reduces the per-packet delay for other traffic on the link. However, reordering is not widely enabled in practice to our knowledge. Out-of-order transmission may be particularly useful for unreliable real-time traffic; reliable traffic is more tolerant of the head-of-line blocking delay caused by in-order delivery. Tradeoffs in link design between delay variation, reordering, and reliability are discussed further in Section 7.

How to model: 1) By swapping two packets in a queue at a given time [10]. The disadvantage of this method is that a queue must have formed for the reordering to happen. 2) Delaying one packet for a given time, and letting other packets pass it.

Because current wireless links generally do not allow reordering, there is little measurement data to guide modeling of reordering in wireless links. A study of handover mechanisms [37] shows reordering of 20 packets passing over 15 outstanding packets.

4.4 On-Demand Resource Allocation

Wireless links often allocate channels based on the availability of user traffic. For instance, in GPRS a radio channel is allocated when data arrives toward the user, and released when the queue size falls below a certain threshold [63, 33]. This section discusses modeling of on-demand resource allocation. Packet Reservation Multiple Access (PRMA) [27] from GPRS is used as an example.

Effect on transport protocols: On-demand resource allocation causes delay variation that depends on traffic patterns. The effects of delay variation on transport protocols have been discussed in a previous section. As an example, the channel allocation delay can trigger spurious TCP timeouts [55]. In addition, the increased delay due to on-demand resource allocation translates to increased delay to the user.

Presence in current and future wireless links: A typical GPRS network requires a 200 ms delay to allocate a channel for the uplink and 80 ms delay for the downlink. For WLAN and satellite links, a new data burst triggers MAC contention. However, subsequent data can often be transmitted without delay, due to the capture effect.

Figure 3 shows measurements of uplink and downlink delay jitter for a GPRS link [2]. The left graph shows pings sent at 0.5 sec-

ond intervals. The uplink delay is highly variable, because while some packets trigger a channel allocation, other packets get a “free ride” on an existing channel. In the downlink, the channel is kept allocated all the time, thus there is no variability for the downlink delay. The right graph shows pings sent at 5 second intervals. The uplink delay is higher than for 0.5 second pings, but is less variable because every packet triggers a channel allocation. The downlink delay is higher with 5 second pings for the same reason.

How to model: On-demand channel allocation can be modeled by introducing an additional delay when a packet arrives to a queue that has been empty longer than the channel hold time. The delay value represents the channel allocation delay, and the hold time represents the duration of channel holding after transmitting a data packet. Detailed simulations of channel allocation in GPRS are also available [64].

Figure 4 shows a simulation of channel allocation for a GPRS link. The model uses a uniform distribution for channel allocation delay separately for the uplink and downlink. Although the simulation results are not exactly the same as the measurements in Figure 3, they are reasonably close. The mean and variance of the uplink and downlink latency are similar in the measured and simulated results. The largest divergence is present for the uplink latency with pings sent at a 0.5 second interval. Measurement samples are evenly distributed within the 220-800 ms region. In simulation, samples concentrate on two levels of 300 and 800 ms that correspond to active and inactive channel states. Still, evaluating a transport protocol with this model should produce more relevant results than with a link having a constant delay.

4.5 Bandwidth Variation

Periodic bandwidth variation or bandwidth oscillation in wireless links occurs when a wireless scheduler assigns a high speed channel to a user for a limited time [67].

Effect on transport protocols: When the link bandwidth decreases abruptly, spurious timeouts can be triggered in TCP due to increased inter-packet delay. Periods of low link bandwidth can also result in congestion, while periods of high bandwidth could result in underutilization of that link.

Presence in current and future wireless links: Bandwidth oscillation can occur in CDMA2000 and UMTS (with High Speed Downlink Packet Access) links for certain configurations, possibly only theoretical [67]. It is an open question to what extent transport protocols can be designed to optimize performance in the presence of bandwidth oscillation, and to what extent wireless link-level mechanisms for bandwidth variation can be designed to take into account the transport protocols that will run over those links [40].

How to model: Bandwidth oscillation can be modeled by changing the bandwidth of a link. A study of CDMA2000 [67] uses periods of high and low bandwidth of 0.02-20 seconds with bandwidth changing between 163 and 9.6 kbps. Another study models bandwidth variation using a sine pattern [19]. In practice bandwidth variation is likely to be limited.

Packet reordering can occur when the link bandwidth changes, due to packets sent after the increase in link bandwidth arriving before packets transmitted earlier over the low bandwidth link. An algorithm has been proposed to prevent packet reordering resulting from bandwidth oscillation [15].





Figure 3: Measured one-way latency of a GPRS link at a 0.5 second (left) and 5 second (right) interval.





Figure 4: Simulated one-way latency of a GPRS link at a 0.5 second (left) and 5 second (right) interval.

4.6 Asymmetry in Bandwidth and Latency

In wireless links, bandwidth and latency are often different in the uplink and downlink directions.

Effect on transport protocols: Asymmetric bandwidth can cause congestion for TCP ACKs [6], with the bandwidth available for ACK packets limiting throughput in the forward direction.

Presence in current and future wireless links: Cellular links have moderate bandwidth asymmetry, with factors of 2 to 5. The downlink latency can be twice that of the uplink [33]. There is no asymmetry in WLAN links. Satellite links are often significantly asymmetric in bandwidth and moderately asymmetric in latency, if a terrestrial link is used as a reverse path.

How to model: Asymmetric bandwidth and latency can be easily modeled in simulations by configuring the uplink and downlink with different parameters for bandwidth and latency. A study of transport over satellite links uses a downlink bandwidth that is 10-100 times higher than the uplink bandwidth [35].

5. QUEUE MANAGEMENT

Effect on transport protocols: Queue management has a strong impact on performance by accommodating bursty traffic while controlling queuing delay. For buffers with Drop-Tail queue management, the buffer size affects the performance of transport protocols.

An overly large buffer causes high queuing delay and poor loss recovery due to long retransmission timeouts. An overly small buffer can result in low link utilization and frequent packet drops. Regardless of the buffer size, Drop-Tail queue management can result in bursts of packet drops, with multiple packet drops from a single traffic burst. In the extreme case, these bursts of packet drops can trigger retransmission timeouts in TCP. For buffers with Active Queue Management (AQM) mechanisms, the details of the AQM mechanism affect the loss rate and delay seen by transport protocols.

Presence in current and future wireless links: Buffering in wireless networks can be complex, with several layers of buffers and flow control between them. Often the link buffers at least the link's bandwidth-delay product worth of data to use for local error recovery with retransmissions. Data can also be buffered for compression, fragmentation and reassembly.

Cellular wireless links often have per-user buffering; flows going to a single mobile terminal share a single buffer. The level of statistical multiplexing for such links can be low, because low bandwidth limits the number of concurrent flows that a user can have with tolerable performance. Overbuffering is common for cellular links [51, 33].

When the level of statistical multiplexing is low and the buffer size

is small, existing AQM algorithms may not perform sufficiently well. New AQM algorithms have been proposed specifically for cellular links [58, 21]. The main problem in such scenarios is preventing the slow start overshoot of TCP connections. The overshoot happens because TCP detects congestion up to one RTT after filling the buffer and the first packet drop. The TCP's sending rate at this point can be twice the available bandwidth of the path, generating many packet drops.

RED is designed to allow bursts of packets to go through the queue. However, at least with a default set of parameters, RED does not prevent the slow start overshoot. However, in simulations where the queue is not empty when a TCP connection is started, RED can perform well. Furthermore, the addition of ECN can improve performance [57].

How to model: Using a Drop-Tail buffer with a configurable maximum size in packets is a plausible approach to model current cellular and WLAN links. For modeling satellite or future cellular and WLAN links, using RED may be more appropriate. These recommendations are based on the authors' experience that current wireless networks use Drop-Tail buffers, but that the design of future networks includes active queue management based on RED. Simulations of satellite and WLAN links should create a moderate level of statistical multiplexing in the buffer.

6. EFFECTS OF MOBILITY

Mobility is an intrinsic property for most wireless links. Mobility presents a major challenge to transport protocols through the packet losses and delay introduced by handovers. Modeling the effects of mobility implies the use of mobility models [42] (e.g., linear, ping-pong), models of network architecture (coverage of individual cells) and of network or link-layer mobility protocols [37] (e.g., Mobile IP). Intersystem handovers are especially challenging for transport protocols. In addition to delays and losses, they can also cause a significant change in link bandwidth and latency.

Effect on transport protocols: A sudden change in link characteristics affects end-to-end congestion control in transport protocols. In particular, slowly-responsive congestion control, such as TFRC, can require considerable time to adapt to changes. In a simulation in the upper right graph in Figure 5, a TFRC flow underutilizes a fast link (WLAN) after a handover from a slow link (GPRS). However, when this experiment is made in real networks, TFRC is able to increase the rate much faster. The reason for this mismatch between simulation and measurement is in part due to differences between the experimental and simulation TFRC implementations, and in part due to the inevitable effects of small changes on transport protocol performance, as described below.

A modest change in the parameters of a simulation or experiment often results in a large change in performance. The simulations in the lower left of Figure 5 differ from those in the upper right only in having larger and probably more realistic buffer sizes. In the lower left, the flow does not have any losses until the 51st second, and thus is in slow start when the handover occurs at time 30. We believe the same scenario occurs in the experiment in the upper left graph, because there are no losses in this experiment. For the simulation in the upper right graph, losses occur at 7 and 10 seconds, kicking the TFRC sender out of slow-start. If the TFRC flow is in slow-start when the handover occurs, this should allow it to increase its sending rate much faster.

As another example of ways that seemingly-small changes in simulation scenarios can have dramatic changes in performance, the simulation in the lower right differs from that in the lower left only in using three receiver reports per round-trip time, instead of the default of one. With the higher feedback frequency, the sender receives more up-to-date information about the receive rate, and as a result increases its sending rate faster than the sender in the lower left. (The TFRC sender is limited to sending at most twice as fast as the receive rate reported by the receiver.) The faster increase in the sending rate generates losses earlier in this simulation, at time 11-16 seconds, and the sender has to slow down. In addition, there are multiple packet drops over a number of round-trip times, resulting in a persistent history of congestion for the sender to maintain. This significantly slows down the sender's rate increase after the handover to the fast WLAN link at time 30.

Presence in current and future wireless links: No existing or emerging wireless access technology provides the best available coverage, bandwidth and cost. Therefore, intersystem handovers appear to be a fundamental property of future wireless networks. The amount of effort invested in making transport protocols perform well in the presence of intersystem handovers is likely to depend on the frequency of their occurrence. It is unclear if scenarios with frequent intersystem handovers will be common in the future. One possible scenario could be Infostations [26], where high-speed connectivity is provided in small coverage areas e.g. on a highway.

How to model: If the researcher is interested in details of an intersystem handover implementation, then the underlying mechanisms would have to be adequately modeled. The topology should include the new and old link, home and foreign agents and Mobile IP registration messages. However, if the researcher is interested only in effect of intersystem handovers on transport protocols, simply changing the link characteristics is sufficient. The link bandwidth, latency and buffer size can be instantly changed in ns-2. When the buffer size is changed, the default ns-2 behavior is not to discard packets above the new limit. This models a situation when data is transferred from the old to the new link so that no losses occur. An alternate behavior available in ns-2 is to discard excess packets when the buffer size is reduced. This models a situation when a buffer on the new link is not sufficient to accommodate all data from the old link.

7. CHANGING WIRELESS LINKS OR TRANSPORT PROTOCOLS?

It is important that designers of the future wireless networks and transport protocols take each other's concerns into consideration. It is not necessarily that wireless characteristics come first as a given and transport protocols have to adapt to wireless characteristics, or vice versa. There is an interplay between the two.

In some cases, wireless link level technologies have been designed to take into account the performance requirements of the current TCP standard, such as TCP's poor performance in the case of corrupted packets or of significant reordering. In other cases, the discussion has been more about the adaptation of transport protocols to the needs of wireless links. In this paper we are not arguing that transport protocols should necessarily be modified to adapt to wireless technologies, or that instead wireless technologies should be designed to meet the expectations of current transport protocols. Instead, we trying to consider the models that researchers can use in investigating these questions and others about transport protocol performance over wireless links.

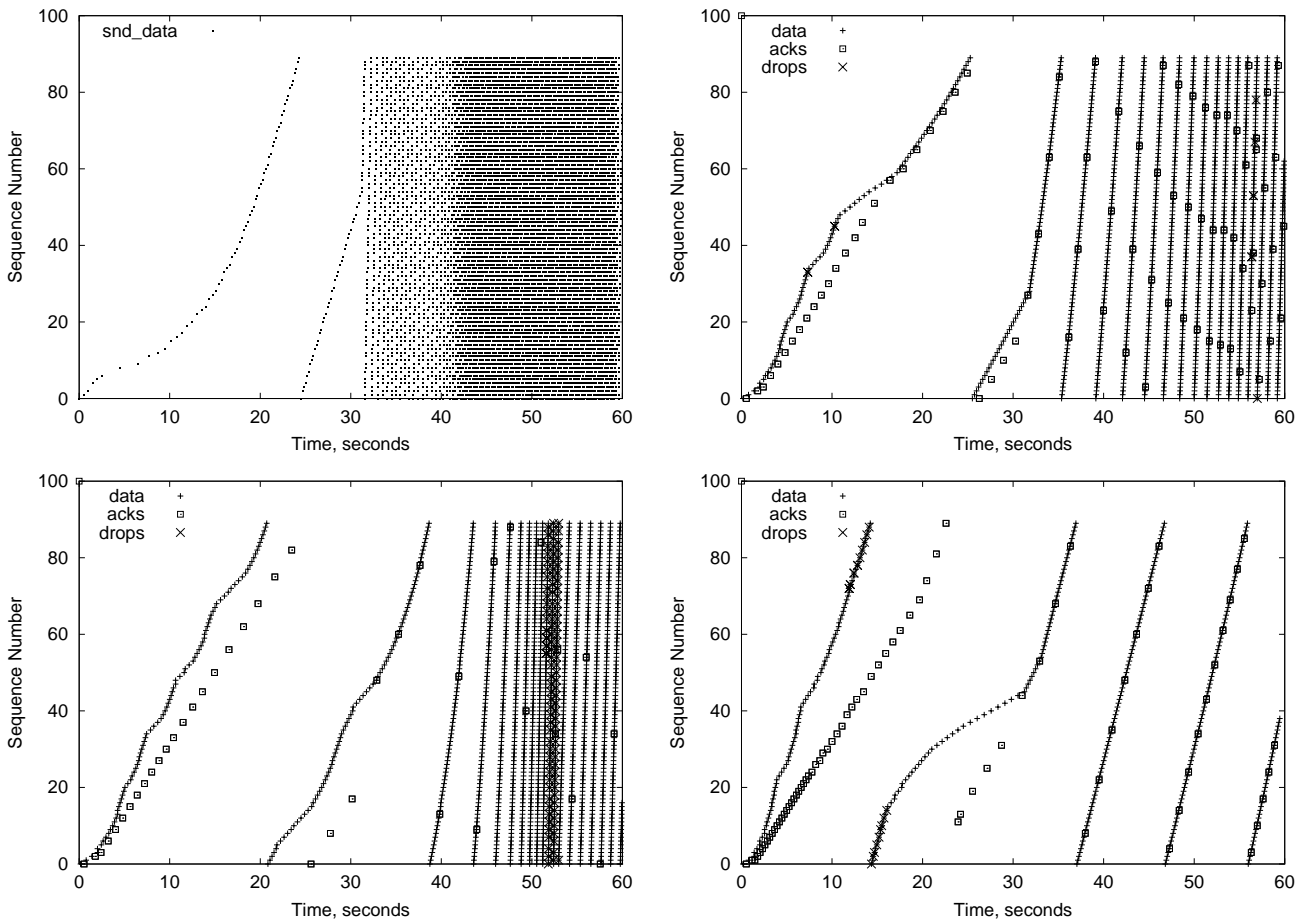


Figure 5: Dynamics of a TFRC flow after a handover from GPRS to WLAN. Handovers occur on the 30th second. Measurements (upper left); Simulations with small buffers (upper right); Simulations with larger buffers (lower left); Simulations with larger buffers and three TFRC receiver reports per RTT (lower right).

Sometimes it is the case for wireless links that their characteristics are fundamental and unadjustable. For example, the high latency of a satellite link is due to the limited speed of light. In contrast, the high latency of terrestrial wireless links is not fundamental, but could often be reduced by changing their design. Such design choices can be explored for the good performance of Internet transport protocols.

Similarly, while some characteristics of transport protocols are fundamental, other characteristics, such as the poor performance in the presence of reordering, are not fundamental, but could be improved in new transport protocols, or in modified versions of the current ones.

Adding to the complications of understanding the interactions between transport protocols and wireless technologies is the reality of the current deployment of wireless link technologies and of transport protocols. Wireless specifications are complex and often implemented in hardware, making changes difficult. Similarly, transport protocols are often deployed on a large number of hosts, and future modifications of transport protocols are often limited by backward compatibility. Models are needed to support researchers trying to understand the performance issues of currently-deployed transport protocols over currently-deployed wireless technologies,

as well as to support researchers designing new technologies for the future.

Below we discuss several issues where there is not a straightforward answer of transport protocols adapting to the characteristics of link-level technologies, or vice-versa, but instead a more complex interplay between transport protocols and link-level technologies.

7.1 Bit Errors on Wireless Links

Because current transport protocols (e.g., TCP) perform poorly over links with bit errors, the recommended approach [3, 16, 43] has been for link layers to incorporate link-level mechanisms such as FEC or link-level retransmission to deal with corrupted packets. However, there are two complications with this approach.

(1) Sometimes corrupted packets cannot be repaired at the link-level in a timely fashion. If the link level abandons its attempts at link-level retransmissions, TCP and related transport protocols (SCTP, DCCP) take the loss as an indication of congestion. If mechanisms were available for explicit corruption notification from the link layer to the transport end nodes, transport protocols could correctly interpret the loss as corruption instead of congestion, and respond appropriately [20].

If the link level is overly-persistent in retransmission attempts, then the TCP sender might already have invoked a retransmission timeout and retransmitted the packet. Similarly, real-time connections with playback times might no longer be interested in the retransmitted packet if it arrives after a long delay from many link-level retransmissions.

(2) Some applications such as streaming audio can tolerate a small level of bit errors in the data payload, and would prefer to receive packets with corrupted payloads rather than have them dropped in the network. For example, the development of transport protocols tolerant of occasional bit errors in the data payload (e.g., using partial checksums, as in the proposal for UDP-Lite) could enable the development of link-level technologies for IP traffic where different tradeoffs are made for different flows about responding to bit errors in the payload.

7.2 Reordering

Because TCP performs poorly over paths with significant reordering, the recommended approach [43] has been for routers and link layers to avoid introducing reordering between packets in a flow as much as possible. However, sometimes packets are reordered in the network in spite of the goal of not introducing reordering. There are several possible responses to the reordering that currently exists: to do nothing; to modify transport protocols to be more robust to reordering [10]; and/or to correct the sources of reordering in the network.

Because it is possible to make current transport protocols considerably more robust to reordering, the question then arises of whether this could enable future developments in routing, router technologies, or link-level protocols. For example, if the transport protocol was robust to reordering, wireless links could use link-level retransmissions or FEC mechanisms that minimized delay by introducing reordering between packets of a flow.

7.3 Delay Variation

The past approach to delay variation has been to understand the ways that transport protocols respond poorly to delay variation, and to propose changes to transport protocols to make performance more robust. While there are some environments where delay variation is fundamental [22], it is also feasible to design link layers in ways that limit the delay variation seen by transport protocols. For example, many current link layers perform only a small number of retransmissions for a corrupted packet, to prevent a large delay spike. Improved handover mechanisms are being designed that reduce the delay spikes due to handovers. Similarly, existing radio schedulers could be configured to keep delay variation within reasonable bounds. Thus, while it can be useful to make transport protocols more robust to delay variation [32], there is also much that is being done to reduce the delay variation introduced by link level technologies.

7.4 Communication Between Link-layers and Transport End-nodes

Our experience is that wireless links are capable of adapting their characteristics so that transport protocols perform well. However, currently link layers have limited information on the transport capabilities of flows.

Cross-layer communication could be helpful for future generations of transport protocols and wireless links. For example, if a trans-

port protocol could inform the link layer that it is tolerant to packet reordering, the link layer could avoid the unnecessary overhead of ordered packet delivery. Using IP options [31] or a DiffServ Codepoint in the IP header seem to be the most feasible ways to implement such approaches. (It would be a poor design for the router to access headers above IP, and changing the format of the IPv4 or IPv6 header is not realistic.)

Cross-layer communication would also be investigated for transport protocols to inform link layers of the degree of delay-tolerance, or of a willingness to receive packets with bit errors in the data payload. Similarly, cross-layer communication can also be investigated for link-layers to inform transport end-nodes of bandwidth or delay changes from handovers, or of link-up conditions of links recovering from an outage.

8. CONCLUSIONS AND FUTURE WORK

In this paper we presented important aspects of modeling wireless links for the evaluation of transport protocols. For delay variation, bandwidth on demand and intersystem handovers we compared simulation models with measurement results. As an accompaniment to this paper, we have included a suite of simulation scenarios in ns-2 that could help in evaluating transport protocol performance over wireless links [30].

Future work is likely to include further investigation of the benefits and costs of explicit communication between link layers and transport protocols. An additional topic for further work is that of the effect of link layer mechanisms on protocols and applications above the link layer.

We note that wireless links are not the only links that can have unexpected interactions with transport protocols and higher layers. On Ethernet traffic in one direction can temporarily block traffic for the other direction because of the capture effect. (This of course does not apply to those “Ethernet” implementations that are point-to-point, with full-duplex links and a switch.) We plan to also explore simulation scenarios for evaluating transport protocol performance over Ethernet.

9. ACKNOWLEDGMENTS

We thank Lloyd Wood, Almudena Konrad, Merkourios Karaliopoulos, Tom Henderson, Guido Hiertz, Farid Khafizov, Saverio Mascolo, Yasuhiko Matsunaga, and the anonymous reviewers.

10. REFERENCES

- [1] 3GPP. TS 30.03: Selection procedures for the choice of radio transmission technologies of the UMTS, Aug. 2003.
- [2] O. Aalto. Measured delay performance of IP bearer over GPRS. Master’s thesis, University of Helsinki, Apr. 2003.
- [3] M. Allman, S. Dawkins, D. Glover, J. Griner, D. Tran, T. Henderson, J. Heidemann, J. Touch, H. Kruse, S. Ostermann, K. Scott, and J. Semke. Ongoing TCP research related to satellites. IETF RFC 2760, Feb. 2000.
- [4] M. Allman, D. Glover, and L. Sanchez. Enhancing TCP over satellite channels using standard mechanisms. IETF RFC 2488, Jan. 1999.
- [5] A. Balachandran, G. Voelker, P. Bahl, and P. Rangan. Characterizing user behavior and network performance in a

- public Wireless LAN. In *Proc. of ACM SIGMETRICS'02*, June 2002.
- [6] H. Balakrishnan, V. Padmanabhan, G. Fairhurst, and M. Sooriyabandara. TCP performance implications of network path asymmetry. IETF RFC 3449, Dec. 2002.
- [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 Trans. on Networking*, 5(6):756–769, Dec. 1997.
- [8] D. Beaufort, L. Fay, C. Samson, and A. Teil. Measured performance of TCP friendly rate control protocol over a 2.5G network. In *Proc. of the IEEE Vehicular Technology Conference (VTC'02 Fall)*, Sept. 2002.
- [9] S. Biaz and N. Vaidya. Discriminating congestion losses from wireless losses using inter-arrival times at the receiver. In *Proc. of IEEE Symposium ASSET'99*, Mar. 1999.
- [10] E. Blanton and M. Allman. On making TCP more robust to packet reordering. *ACM Computer Communication Review*, 32(1):20–30, Jan. 2002.
- [11] L. S. Brakmo and L. L. Peterson. TCP Vegas: End to end congestion avoidance on a global Internet. *IEEE Journal on Selected Areas in Communications*, 13(8):1465–1480, Oct. 1995.
- [12] G. Brasche and B. Walke. Concepts, services and protocols of the new GSM phase 2+ general packet radio service. *IEEE Communications Magazine*, 35(8):94–104, Aug. 1997.
- [13] S. Cen, P. Cosman, and G. Voelker. End-to-end differentiation of congestion and wireless losses. *IEEE/ACM Trans. on Networking*, 11(5):703–717, Oct. 2003.
- [14] R. Chakravorty, S. Katti, J. Crowcroft, and I. Pratt. Flow aggregation for enhanced TCP over wide area wireless. In *Proc. of IEEE INFOCOM'03*, Apr. 2003.
- [15] M. Chan and R. Ramjee. TCP/IP performance over 3G wireless links with rate and delay variation. In *Proc. of ACM MOBICOM'02*, Sept. 2002.
- [16] S. Dawkins, G. Montenegro, M. Kojo, V. Magret, and N. Vaidya. End-to-end performance implications of links with errors. IETF RFC 3155, Aug. 2001.
- [17] D. De Couto, D. Aguayo, B. A. Chambers, and R. Morris. Performance of multihop wireless networks: Shortest path is not enough. *ACM Computer Communication Review*, 33(1):83–88, Jan. 2003.
- [18] A. DeSimone, M. C. Chuah, and O.-C. Yue. Throughput performance of transport-layer protocols over wireless LANs. In *Proc. of IEEE GLOBECOM'93*, Dec. 1993.
- [19] D. Dutta and Y. Zhang. An active proxy based architecture for TCP in heterogeneous variable bandwidth networks. In *Proc. of IEEE Globecom'01*, Nov. 2001.
- [20] W. Eddy, S. Ostermann, and M. Allman. New techniques for making transport protocols robust to corruption-based loss. Submitted for publication, Jan. 2004.
- [21] H. Ekstrom and R. Ludwig. Queue management for TFRC-based traffic in 3G networks. In *Proc. of Workshop on Mobile and Wireless Networks (MWN)*, May 2003.
- [22] K. Fall. A delay-tolerant network architecture for challenged internets. In *Proc. of ACM SIGCOMM'03*, Aug. 2003.
- [23] S. Floyd, M. Handley, J. Padhye, and J. Widmer. Equation-based congestion control for unicast applications. In *Proc. of ACM SIGCOMM'00*, Aug. 2000.
- [24] S. Floyd and J. Kempf. IAB concerns regarding congestion control for voice traffic in the internet. Work in progress, draft-iab-congestion-01.txt, Oct. 2003.
- [25] S. Floyd and E. Kohler. Internet research needs better models. In *Proc. of HotNets-I*, Oct. 2002.
- [26] R. H. Frenkiel, B. R. Badrinath, J. Borrás, and R. D. Yates. The Infostations challenge: Balancing cost and ubiquity in delivering wireless data. *IEEE Personal Communications Magazine*, 7(2):66–71, Apr. 2000.
- [27] D. J. Goodman, R. A. Valenzuela, K. T. Gayliard, and B. Ramamoorthi. Packet reservation multiple access for local wireless communications. *IEEE Trans. on Communications*, 37(8):885–890, Aug. 1989.
- [28] K. P. Gummadi, R. J. Dunn, S. Saroiu, S. D. Gribble, H. M. Levy, and J. Zahorjan. Measurement, modeling, and analysis of a peer-to-peer file-sharing workload. In *Proc. of the 19th ACM symposium on operating systems principles*, Oct. 2003.
- [29] A. Gurtov. Effect of delays on TCP performance. In *Proc. of IFIP Personal Wireless Communications (PWC'01)*, Aug. 2001.
- [30] A. Gurtov. Simulation tests for modeling wireless links, Dec. 2003. Directory tcl/ex/wireless-scripts in the ns-2 simulator.
- [31] A. Gurtov and R. Ludwig. Lifetime packet discard for efficient real-time transport over cellular links. *ACM Mobile Computing & Communications Review*, 7(4):32–45, Oct. 2003.
- [32] A. Gurtov and R. Ludwig. Responding to spurious timeouts in TCP. In *Proc. of IEEE INFOCOM'03*, Apr. 2003.
- [33] A. Gurtov, M. Passoja, O. Aalto, and M. Raitola. Multi-layer protocol tracing in a GPRS network. In *Proc. of the IEEE Vehicular Technology Conference (VTC'02 Fall)*, Sept. 2002.
- [34] J. Heidemann, N. Bulusu, J. Elson, C. Intanagonwiwat, K. Lan, Y. Xu, W. Ye, D. Estrin, and R. Govindan. Effects of detail in wireless network simulation. In *Proc. of the SCS Multiconference on Distributed Simulation*, Jan. 2001.
- [35] T. Henderson and R. Katz. Transport protocols for Internet-compatible satellite networks. *IEEE Journal on Selected Areas in Communications*, 17(2):345–359, Feb. 1999.
- [36] H.-Y. Hsieh and R. Sivakumar. A transport layer approach for achieving aggregate bandwidths on multi-homed mobile hosts. In *Proc. of ACM MOBICOM'02*, Sept. 2002.
- [37] R. Hsieh and A. Seneviratne. A comparison of mechanisms for improving Mobile IP handoff latency for end-to-end TCP. In *Proc. of ACM MOBICOM'03*, Sept. 2003.

- [38] IETF. Performance implications of link characteristics working group, Jan. 2004.
<http://www.ietf.org/html.charters/pilc-charter.html>.
- [39] IETF. TCP over satellite working group, Jan. 2004.
<http://www.isi.edu/tcpsat/>.
- [40] H. Inamura, G. Montenegro, R. Ludwig, A. Gurtov, and F. Khafizov. TCP over second (2.5G) and third (3G) generation wireless networks. IETF RFC 3481 (BCP 71), Feb. 2003.
- [41] V. Jacobson. Congestion avoidance and control. In *Proc. of ACM SIGCOMM'88*, Aug. 1988.
- [42] A. Jardosh, E. Belding-Royer, K. Almeroth, and S. Suri. Towards realistic mobility models for mobile ad hoc networks. In *Proc. of ACM MOBICOM'03*, Sept. 2003.
- [43] P. Karn, C. Bormann, G. Fairhurst, D. Grossman, R. Ludwig, J. Mahdavi, G. Montenegro, J. Touch, and L. Wood. Advice for Internet subnetwork designers. Work in progress, draft-ietf-pilc-link-design-15.txt, Dec. 2003.
- [44] J. Kilpi. A portrait of a GPRS/GSM session. In *Proc. of the 18th International Teletraffic Congress*, Sept. 2003.
- [45] E. Kohler, M. Handley, and S. Floyd. Designing DCCP: Congestion control without reliability. Available at <http://www.icir.org/kohler/dccp/>, May 2003.
- [46] A. Konrad and A. D. Joseph. Choosing an accurate network path model. In *Proc. of ACM SIGMETRICS'03*, June 2003. Extended abstract.
- [47] A. Konrad, B. Y. Zhao, A. D. Joseph, and R. Ludwig. A Markov-based channel model algorithm for wireless networks. *ACM/Baltzer Wireless Networks*, 9(3):189–199, May 2003.
- [48] R. Krashinsky and H. Balakrishnan. Minimizing energy for wireless web access with bounded slowdown. In *Proc. of ACM MOBICOM'02*, Sept. 2002.
- [49] D. Loguinov and H. Radha. Measurement study of low-bitrate internet video streaming. In *Proc. of the First ACM SIGCOMM Internet Measurement Workshop (IMW-01)*, Nov. 2001.
- [50] R. Ludwig and R. H. Katz. The Eifel algorithm: Making TCP robust against spurious retransmissions. *ACM Computer Communication Review*, 30(1):30–36, Jan. 2000.
- [51] R. Ludwig, B. Rathonyi, A. Konrad, K. Oden, and A. Joseph. Multi-layer tracing of TCP over a reliable wireless link. In *Proc. of ACM SIGMETRICS'99*, May 1999.
- [52] S. Mascolo, C. Casetti, M. Gerla, M. Y. Sanadidi, and R. Wang. TCP Westwood: Bandwidth estimation for enhanced transport over wireless links. In *Proc. of ACM MOBICOM'01*, July 2001.
- [53] V. Mbarika, M. Jensen, and P. Meso. Cyberspace across Sub-Saharan Africa. *Communications of the ACM*, 45(12):17–21, Dec. 2002.
- [54] C. H. Nam, S. C. Liew, and C. P. Fu. An experimental study of ARQ protocol in 802.11b Wireless LAN. In *Proc. of Wireless Personal Multimedia Communications (WPMC'02)*, Oct. 2002.
- [55] J. Neale and A. Mohsen. Impact of CF-DAMA on TCP via satellite performance. In *Proc. of IEEE Globecom'01*, Nov. 2001.
- [56] NTT. DoCoMo subscriber growth, Jan. 2004.
<http://www.nttdocomo.com/companyinfo/subscriber.html>.
- [57] K. Ramakrishnan, S. Floyd, and D. Black. The addition of explicit congestion notification (ECN) to IP. IETF RFC 3168, Sept. 2001.
- [58] M. Sagfors, R. Ludwig, M. Meyer, and J. Peisa. Queue management for TCP traffic over 3G links. In *Proc. of the IEEE Wireless Communications and Networking Conference (WCNC'03)*, Mar. 2003.
- [59] M. Stemm and R. H. Katz. Vertical handoffs in wireless overlay networks. *ACM Mobile Networks and Applications*, 3(4):335–350, Dec. 1998.
- [60] P. Stuckmann, H. Finck, and T. Bahls. A WAP traffic model and its appliance for the performance analysis of WAP over GPRS. In *Proc. of the IEEE International Conference on Third Generation Wireless and Beyond (3Gwireless 2001)*, June 2001.
- [61] V. Tsaoussidis and I. Matta. Open issues on TCP for mobile computing. *Journal of Wireless Communications and Mobile Computing*, 2(1):3–20, Feb. 2002.
- [62] UCB/LBNL/VINT. The ns-2 network simulator, Aug. 2003.
<http://www.isi.edu/nsnam/ns/>.
- [63] B. Walke. *Mobile Radio Networks, Networking and Protocols (2. Ed.)*. Wiley & Sons, 2001.
- [64] H. Wiemann, A. Schieder, and H. Ekstrom. Enhanced TBF features in GERAN. In *Proc. of Wireless Personal Multimedia Communications (WPMC'01)*, Sept. 2001.
- [65] H. Wu, Y. Peng, K. Long, S. Cheng, and J. Ma. Performance of reliable transport protocol over IEEE 802.11 Wireless LAN: Analysis and enhancement. In *Proc. of IEEE INFOCOM'02*, June 2002.
- [66] G. Xylomenos. *Multi Service Link Layers: An Approach to Enhancing Internet Performance over Wireless Links*. PhD thesis, University of California at San Diego, 1999.
- [67] M. Yavuz and F. Khafizov. TCP over wireless links with variable bandwidth. In *Proc. of the IEEE Vehicular Technology Conference (VTC'02 Fall)*, Sept. 2002.