TCP Profile for Next Generation Wireless Networks

Andrei V. Gurtov, Dr. Iouri A. Bogoiavlenski

Department of Computer Science, University of Petrozavodsk

Lenin St., 33, Petrozavodsk, Republic of Karelia, 185640, Russia

E-mail: {ybgv, gurtov}@cs.karelia.ru

Abstract

TCP remains to be the dominant transport protocol in the Internet providing reliable end-to-end data delivery. Performance problems of TCP in wireless networks are well known. In this paper, we describe the link characteristics of 2.5G and 3G wireless networks. Furthermore, a profile of TCP options is presented that in most cases provides good TCP performance in 2.5G and 3G wireless networks. In particular, the effect of the TCP timestamp option is considered.

1 Introduction

Recent measurement studies [20] show that TCP maintains its position as the dominant transport protocol in the Internet. TCP is a stable, mature, and probably the most thoroughly tested protocol of its kind. Nevertheless, there are some special cases where TCP could still be improved.

This paper proposes a profile of such techniques, particularly effective for use with 2.5G and 3G wireless networks. These configuration options are commonly found in modern TCP stacks, and are widely available IETF standards-track mechanisms that the community has judged to be safe on the general Internet. The TCP profile has been approved in the IETF as a best current practice RFC [10].

2 2.5G and 3G Link Characteristics

Link layer characteristics of 2.5G/3G networks have significant effects on TCP performance. In this section we present various aspects of link characteristics unique to the 2.5G/3G networks [8].

2.1 Latency

The latency of 2.5G/3G links is high mostly due to the extensive processing required at the physical layer of those networks, such as for FEC and interleaving, and due to transmission delays in the radio access network (including link-level retransmissions). A typical RTT varies between a few hundred milliseconds and one second. The associated radio channels suffer from difficult propagation environments. Hence, powerful but complex physical layer techniques need to be applied to provide high capacity in a wide coverage area in a resource efficient way. Hopefully, rapid improvements in all areas of wireless networks ranging from radio layer techniques over signal processing to system architecture will ultimately also lead to reduced delays in 3G wireless systems.

2.2 Data Rates

The main incentives for transition from 2G to 2.5G to 3G are the increase in voice capacity and in data rates for the users. 2.5G systems have data rates of 10-20 kbps in uplink and 10-40 kbps in downlink. Initial 3G systems are expected to have bit rates around 64 kbps in uplink and 384 kbps in downlink. Considering the resulting bandwidth-delay product (BDP) of around 1-5 KB for 2.5G and 8-50 KB for 3G, 2.5G links can be considered LTNs (Long Thin Networks [21]), and 3G links approach LFNs (Long Fat Networks [12], as exemplified by some satellite networks [2]).

For good TCP performance both LFNs and LTNs require maintaining a large enough window of outstanding data. For LFNs, utilizing the available network bandwidth is of particular concern. LTNs need a sufficiently large window for efficient loss recovery. In particular, the fast retransmit algorithm cannot be triggered if the window is less than four segments. This leads to a lengthy recovery through retransmission timeouts. The Limited Transmit algorithm [1] helps avoid the deleterious effects of timeouts on connections with small windows. Nevertheless, making full use of the SACK [19] information for loss recovery in both LFNs and LTNs may require twice the window otherwise sufficient to utilize the available bandwidth.

Data rates are dynamic due to effects from other users and from mobility. Arriving and departing users can reduce or increase the available bandwidth in a cell. Increasing the distance from the base station decreases the link bandwidth due to reduced link quality. Finally, by simply moving into another cell the user can experience a sudden change in available bandwidth. For example, if upon changing cells a connection experiences a sudden increase in available bandwidth, it can underutilize it, because during congestion avoidance TCP increases the sending rate slowly. Changing from a fast to a slow cell normally is handled well by TCP due to the self-clocking property. However, a sudden increase in RTT in this case can cause a spurious TCP timeout. In addition, a large TCP window used in the fast cell can create congestion resulting in overbuffering in the slow cell.

2.3 Asymmetry

2.5G/3G systems may run asymmetric uplink and downlink data rates. The uplink data rate is limited by battery power consumption and complexity limitations of mobile terminals. However, the asymmetry does not exceed 3-6 times, and can be tolerated by TCP without the need for techniques like ACK congestion control or ACK filtering [5]. Accordingly, this document does not include recommendations meant for such highly asymmetric networks.

2.4 Delay Spikes

A delay spike is a sudden increase in the latency of the communication path. 2.5G/3G links are likely to experience delay spikes exceeding the typical RTT by several times due to the following reasons.

1. A long delay spike can occur during link layer recovery from a link outage due to temporal loss of radio coverage, for example, while driving into a tunnel or within an elevator.

2. During a handover the mobile terminal and the new base station must exchange messages and perform some other time-consuming actions before data can be transmitted in a new cell.

3. Many wide area wireless networks provide seamless mobility by internally re-routing packets from the old to the new base station which may cause extra delay.

4. Blocking by high-priority traffic may occur when an arriving circuit-switched call or higher priority data temporarily preempts the radio channel. This happens because most current terminals are not able to handle a voice call and a data connection simultaneously and suspend the data connection in this case.

5. Additionally, a scheduler in the radio network can suspend a low-priority data transfer to give the radio channel to higher priority users.

Delay spikes can cause spurious TCP timeouts, unnecessary retransmissions and a multiplicative decrease in the congestion window size.

2.5 Packet Loss Due to Corruption

Even in the face of a high probability of physical layer frame errors, 2.5G/3G systems have a low rate of packet losses thanks to link-level retransmissions. Justification for link layer ARQ is discussed in [6, 13, 18]. In general, link layer ARQ and FEC can provide a packet service with a negligibly small probability of undetected errors (failures of the link CRC), and a low level of loss (non-delivery) for the upper layer traffic, e.g., IP. The loss rate of IP packets is low due to the ARQ, but the recovery at the link layer appears as delay jitter to the higher layers lengthening the computed RTO value.

2.6 Intersystem Handovers

In the initial phase of deployment, 3G systems will be used as a "hot spot" technology in high population areas, while 2.5G systems will provide lower speed data service elsewhere. This creates an environment where a mobile user can roam between 2.5G and 3G networks while keeping ongoing TCP connections. The inter-system handover is likely to trigger a high delay spike, and can result in data loss. Additional problems arise because of context transfer, which is out of scope of this document, but is being addressed elsewhere in the IETF in activities addressing seamless mobility [14].

Intersystem handovers can adversely affect ongoing TCP connections since features may only be negotiated at connection establishment and cannot be changed later. After an intersystem handover, the network characteristics may be radically different, and, in fact, may be negatively affected by the initial configuration. This point argues against premature optimization by the TCP implementation.

2.7 Bandwidth Oscillation

Given the limited RF spectrum, satisfying the high data rate needs of 2.5G/3G wireless systems requires dynamic resource sharing among concurrent data users. Various scheduling mechanisms can be deployed in order to maximize resource utilization. If multiple users wish to transfer large amounts of data at the same time, the scheduler may have to repeatedly allocate and de-allocate resources for each user. We refer to periodic allocation and release of high-speed channels as Bandwidth Oscillation. Bandwidth Oscillation effects such as spurious retransmissions were identified elsewhere (e.g., [17]) as factors that degrade throughput.

There are research studies [15, 23], which show that in some cases Bandwidth Oscillation can be the single most important factor in reducing throughput. For fixed TCP parameters the achievable throughput depends on the pattern of resource allocation. When the frequency of resource allocation and de-allocation is sufficiently high, there is no throughput degradation. However, increasing the frequency of resource allocation/de-allocation may come at the expense of increased signaling, and, therefore, may not be desirable. Standards for 3G wireless technologies provide mechanisms that can be used to combat the adverse effects of Bandwidth Oscillation. It is the consensus of the PILC Working Group that the best approach for avoiding adverse effects of Bandwidth Oscillation is proper wireless sub-network design [13].

3 TCP Profile for Wireless Overlay Networks

Operators having control over handset configuration, such as NTT DoCoMo, as well as standardization organizations, such as WAP Forum, who wish to adapt TCP for use in wireless networks would benefit from a wireless TCP profile. IETF RFC3481 [10] defines and motivates the use of state-of-the-art standard-track TCP features found in modern TCP stacks. These TCP features are widely available, can be used safely in the Internet, and include:

- a large initial window [3],
- a window scale option [12],
- Limited Transmit algorithm [1],
- discovery of the path Maximum Transfer Unit (MTU),
- Selective Acknowledgments (SACK) [19],
- Explicit Congestion Notification (ECN) [22],
- a timestamp option [12], and
- disabling TCP/IP header compression which is not robust to packet losses.

We found that the TCP timestamp option increases the accuracy of RTT measurements in bandwidth-limited networks and decreases the likelihood of spurious timeouts [7]. This is different from previous work that did not regard the timestamp option to be useful in the Internet [4]. As an example, Figure 1(a) shows that TCP underestimates the end-to-end delay and experiences a spurious timeout during fast recovery. The graph is obtained using the ns-2 simulator. When timestamps are enabled in Figure 1(b), there is no timeout and throughput is higher. Timestamps enable the TCP sender to have a more accurate estimate of RTT, because all acknowledgments, including those for retransmitted segments, can be used for taking a RTT sample.

Traditional compression protocols, such as Van Jacobson's header compression [11], do not work well with a high level of data losses in a wireless environment [16, p. 35]. The Robust Header Compression working group in IETF is developing new solutions suitable for wireless networks [9].

4 Conclusions

In this paper, we described challenging link characteristics of 2.5G and 3G wireless networks. When profiled with the state-of-the-art capabilities,

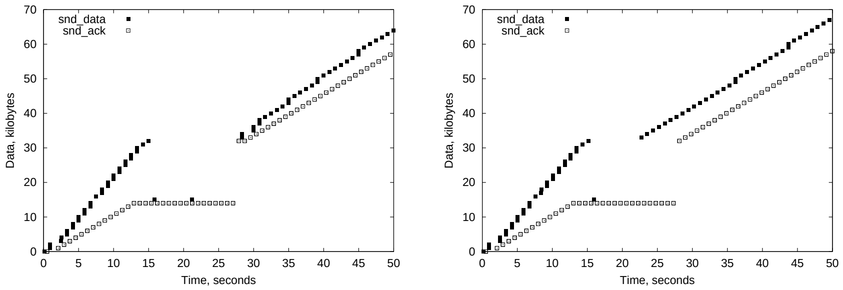

(a) Without timestamps

(b) With timestamps

Figure 1: SACK TCP without timestamps experiences a spurious timeout during fast recovery.

TCP is a well-suited protocol for reliable data transport over wireless links. We presented a set of TCP options suited for use in wireless networks. The timestamp TCP option is shown to improve the retransmit timer accuracy and help avoiding spurious timeouts in fast recovery.

Acknowledgments

We thank co-authors of RFC 3481 and other members of the IETF PILC working group for fruitful discussions.

References

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

[2] 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.

[3] M. Allman, S. Floyd, and C. Partridge. Increasing TCP's initial window. IETF RFC 3390, Oct. 2002.

[4] M. Allman and V. Paxson. On estimating end-to-end network path properties. In *Proc. of ACM SIGCOMM'99*, Aug. 1999.

[5] H. Balakrishnan, V. Padmanabhan, G. Fairhurst, and M. Sooriyabandara. TCP performance implications of network path asymmetry. IETF RFC 3449, Dec. 2002.

[6] G. Fairhurst and L. Wood. Advice to link designers on link Automatic Repeat reQuest (ARQ). IETF RFC 3366, Aug. 2002.

[7] A. Gurtov. Making TCP robust against delay spikes. Technical Report C-2001-53, University of Helsinki, Nov. 2001.

[8] 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.

[9] IETF. ROHC: Robust header compression, June 2003. http://www.ietf.org/html.charters/rohc-charter.html.

[10] 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.

[11] V. Jacobson. Compressing TCP/IP headers for low-speed serial links. IETF RFC 1144, Feb. 1990.

[12] V. Jacobson, R. Braden, and D. Borman. TCP extensions for high performance. IETF RFC 1323, May 1992.

[13] 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.

[14] J. Kempf. Problem description: Reasons for performing context transfers between nodes in an ip access network. IETF RFC 3374, Sept. 2002.

[15] F. Khafizov and M. Yavuz. Running TCP over IS-2000. In *Proc. of the IEEE International Conference on Communications (ICC'02)*, Apr. 2002.

[16] R. Ludwig. *Eliminating Inefficient Cross-Layer Interactions in Wireless Networking.* PhD thesis, Aachen University of Technology, Apr. 2000.

[17] 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.

[18] R. Ludwig, A. Konrad, A. D. Joseph, and R. H. Katz. Optimizing the end-to-end performance of reliable flows over wireless links. *ACM/Baltzer Wireless Networks*, 8(2):289–299, Mar. 2002.

[19] M. Mathis, J. Mahdavi, S. Floyd, and A. Romanow. TCP selective acknowledgement options. IETF RFC 2018, Oct. 1996.

[20] MAWI. Packet traces from WIDE backbone. Available at http://tracer.csl.sony.co.jp/mawi/, Aug. 2003.

[21] G. Montenegro, S. Dawkins, M. Kojo, V. Magret, and N. Vaidya. Long thin networks. IETF RFC 2757, Jan. 2000.

[22] K. Ramakrishnan, S. Floyd, and D. Black. The addition of explicit congestion notification (ECN) to IP. IETF RFC 3168, Sept. 2001.

[23] 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.