Slope Based Discard: A Buffer Management Scheme for 3G Links Supporting TCP Traffic

Juan J. Alcaraz, Fernando Cerdán
Department of Information Technologies and Communications
Polytechnic University of Cartagena
Plaza del Hospital, 1, 30202 Cartagena, SPAIN
34 968326544

{juan.alcaraz, fernando.cerdan}@upct.es

ABSTRACT
This paper introduces Slope Based Discard (SBD), a novel Active Queue Management (AQM) scheme that improves TCP performance over 3G wireless links. Its main objective is to prevent buffer overflow at the downlink buffer in the Radio Network Controller, where the flows of a single user are multiplexed. In contrast to Random Early Detection (RED), SBD is a deterministic algorithm, and its discarding policy is based on the observation of the filling rate of the buffer. By means of extensive simulation experiments, we disclose the influence of each SBD parameter on TCP goodput and latency. Compared to RED, which parameter configuration is also addresed, SBD shows similar goodput figures but lower delays. We describe other characteristics that make SBD a better choice than RED for 3G links, like its greater capability to prevent buffer overflow, and its easier implementation.

Categories and Subject Descriptors
C.2.1 [**Network Architecture and Design**]: Wireless communication.

C.2.2 [**Network Protocols**]: Applications – TCP/IP.

General Terms
Algorithms, Performance.

Keywords
TCP over 3G links, Radio Link Control (RLC).

1. INTRODUCTION
Third generation cellular networks (3G) are expected to be an important part of the Internet. Many Internet applications like e-mail, web surfing and downloading applications rely on TCP for the end-to-end transport. TCP is a reliable, connection oriented transport protocol. One of its main tasks is the congestion control of the Internet hosts. Current implementations of TCP are unable to distinguish packet discards due to congestion from packet losses due to propagation errors. This causes a severe reduction in TCP performance when a wireless link is present in the end-to-end path [1].

In order to adapt 3G links to TCP traffic, the 3GPP standards provide a link layer protocol, the Radio Link Control (RLC), which offers a reliable service when configured in Acknowledge Mode (RLC AM) [2]. The benefit of a reliable radio bearer for TCP connections over wireless links is argued in previous studies, e.g. [1, 3]. However, several characteristics of 3G radio bearers like high and variable latency and buffer overflow of the downlink buffers [4, 5], have undesired effects on TCP performance.

There are several approaches to enhance the TCP–RLC interaction. Some of them propose the modification of the TCP protocol itself [6] or the introduction of an intermediate proxy [3, 5]. Many of these solutions are not practical in the short term because they may need changes on the protocol stack at the end users, and some proposals break the end-to-end semantics of TCP.

A less developed approach is the application of Active Queue Management (AQM) techniques at the downlink RLC buffers. This solution may improve TCP performance over 3G links [7] with a small change at the Radio Network Controler (RNC) nodes. One of the most extended AQM mechanisms in the Internet routers is Random Early Detection [8]. RED can be adapted to the particularities of 3G links [7]. However [9] argues that the complex parameter configuration of RED in 3G links is a drawback that deterministic approaches may overcome.

This paper describes a novel deterministic AQM algorithm called Slope Based Discard (SBD) especially suitable to the characteristics of 3G radio bearers.

We also contribute providing further insight into the parameter configuration of RED in a 3G link. Our evaluation, based on extensive simulations, shows the effect of SBD as well as RED parameters on TCP performance. Based on these results, and on some detailed examples, we disclose the reasons that make SBD more suitable for RLC than RED.

The rest of the paper is organized as follows. Section 2 describes the characteristics of the 3G radio bearer that degrade TCP performance. Section 3 addresses related works. Section 4 explains the SBD algorithm in detail. Section 5 provides a brief description of the simulation environment and summarizes the parameter configuration guidelines derived from the simulation results. Section 6 evaluates RED, which operation and

IWCMC'06, July 3–6, 2006, Vancouver, British Columbia, Canada.

peformance is compared to those of SBD. The paper concludes in section 7 giving a brief outline about future improvements.

2. PROBLEM DESCRIPTION

Previous experimental [10] and simulation [11] results provide a clear view of the characteristics of 3G wireless links. The behaviour of the link buffer occupancy has shown a great impact on TCP performance.

3G links employs per-user buffering. Flows going to a single mobile terminal share a single buffer. Therefore, radio bearers are expected to multiplex a number of simultaneous connections ranging form 1 to 4 TCP flows [12]. At a reliable RLC layer, the upper layer packets will be stored in the downlink buffer until they are fully acknowledged by the receiver side. The consequence is that, as described in [5, 7, 11] frame losses in the downlink channel result in higher link layer buffer occupancy at the RLC network side. Considering that current RLC implementations use drop-tail buffers, the buffer may overflow causing consecutive packet losses. This situation is especially harmful in the first stages of a TCP connection (slow start) and has a higher impact in TCP Reno, which can only recover from consecutive losses with a Retransmission TimeOut (RTO). An RTO reduces TCP transmission window to one, causing the highest reduction of the source rate.

2.1 Overbuffering

The buffer should be large enough to avoid frequent overflow. However, excessive queueing cause some additional problems [5] like Round Trip Time (RTT) inflation, unfairness between competing flows and viscous web surfing.

Figure 1 illustrates the end-to-end goodput and delay of TCP over a radio bearer for different RLC buffer sizes and a number of flows ranging from 1 to 4 (the callouts of this figure are shared by

Table 1 Simulation Parameters

TCP parameters	Setting
Maximum TCP/IP packet size	1500 bytes
Maximum allowed window	64 Kbytes
Initial window	1
Wired Network Round Trip Delay	200 ms
3G link parameters	**Setting**
PDU payload size	320 bits
PDUs per TTI	12
TTI (Transm. Time Interval)	10 ms
RB nominal bit rate	384 Kbit/s
Transmission window	1024 PDUs
maxDAT	10
In-order-delivery	True
Status Prohibit Timer	60 ms
Missing PDU detection	true
Poll Timer	60 ms
Wireless Round Trip Delay	50 ms
Normalized doppler frequency	0,01
Poll window	50 %
Last PDU in buffer Poll	yes
Last retransmitted PDU Poll	yes

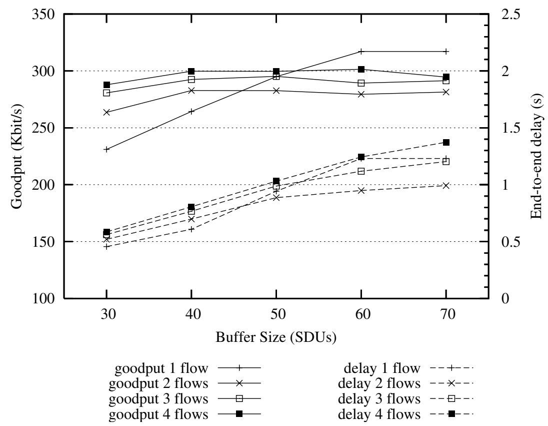

Figure 1. TCP performance over a 384 kbit/s radio bearer

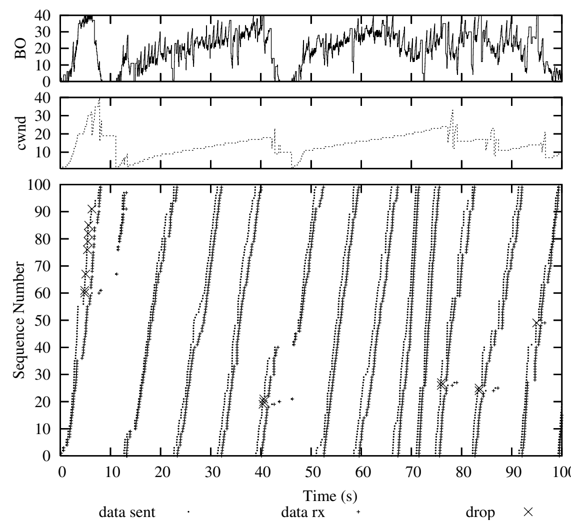

Figure 2. Trace of a TCP connection over RLC

the rest of the figures in this paper). The goodput represents the successfully received packets at the receiver and the delay measures the transfer time of a packet in the downlink direction, at the TCP layer. The buffer size is given in RLC Service Data Units (SDU) of 1500 bytes. Table 1 shows the parameter configuration for the RLC and TCP protocols. The RLC parameters were set according to the optimizing considerations described in [7, 11]. As expected, Figure 1 reveals that a larger buffer benefits the goodput performance but an oversized buffer increases the latency.

Figure 2 shows the trace of a TCP connection over a 384 kbit/s radio bearer multiplexing 2 TCP flows. The curve at the top shows the RLC buffer occupancy. The buffer size is limited to 40 SDUs. The TCP congestion window (*cwnd*) is shown below, and the curve at the bottom shows the sequence number of the packets when they are sent, received and dropped.

At the first stages of the connection, multiple packets are dropped due to buffer overflow, causing an RTO. The buffer is drained because both connections reduce their rate, with the consequent underutilization of the resources. After that, the rate is recovered slowly and the overbuffering appears again, causing high delay and additional consecutive packet losses.

3. RELATED WORK

AQM techniques were developed to avoid buffer overflow at Internet routers, but now they are considered as a feasible strategy to enhance TCP performance over 3G links. The reason is that these techniques are able to maintain the buffer occupancy around certain level, thus avoiding consecutive packet losses and reducing the delay jitter. One of the most extended AQM schemes for Internet routers is Random Early Discard (RED) [8]. Reference [7] shows how RED can be adapted to 3G links.

4. PROPOSAL

We propose a new AQM algorithm for the downlink RLC buffer in 3G networks, named Slope Based Discard (SBD).

The SBD algorithm is based on the following ideas:

1. A packet discard is a congestion signal directed to the TCP sender side that takes a certain amount of time, T_S, to arrive to the TCP source (see Figure 3). The rate reduction is perceived at the buffer after the propagation time, T_f, of the fixed (wired) network.
2. The instantaneous quality of the wireless channel has a direct impact on the RLC buffer occupancy process. Given a constant transmission rate at the source, a higher frame loss ratio in the wireless channel imposes a faster buffer filling rate, because the buffer is drained more slowly.
3. The discarding policy is driven by the buffer filling rate, r. In normal operation, whenever r exceeds a critical value, r_c a packet is dropped. The buffer occupancy level determines the value of r_c.
4. The value of r_c represents the filling rate that, if sustained, would fill the buffer entirely before the rate reduction caused by a packet discard is perceived at the buffer.
5. After a packet drop, additional packet discarding should be avoided until the rate reduction at the TCP source can be noticed at the RLC buffer. In Figure 3, this reaction time equals $T_S + T_f$. Only after this period it is possible to know if the rate reduction is enough to prevent buffer overflow.
6. The packet chosen for discard will be as close as possible to the front of the queue, in order to reduce the reaction time. Additionally, the algorithm should not discard a packet if its transmission over the RLC link has already started. Otherwise, upon a packet discard, the RLC would start the signalling procedure required to synchronize RLC sender and receiver sides [2]. Consequently, our protocol discards the first packet whose transmission has not started, thus reducing complexity and avoiding changes in the 3GPP specification itself

The following parameters control the SBD algorithm:

- *minth* is the occupancy level above which packets can be dropped.
- *maxth* is the maximum occupancy allowed in the buffer.
- T_r is the estimated reaction time.
- α is used for the estimation of the slope of the buffer occupancy process curve. Instead of doing the measure over fixed periods of time, we propose to use occupancy intervals of α bits. Thus, the algorithm controls the time used to reach the threshold (th_n) of this interval. This method is easier to implement and allows a continuous control over the instantaneous queue size.

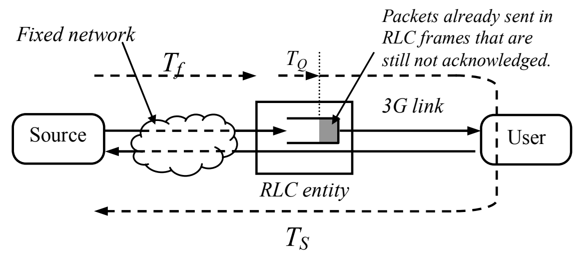

Figure 3. Schematic diagram of the end-to-end connection with the delay times of each section.

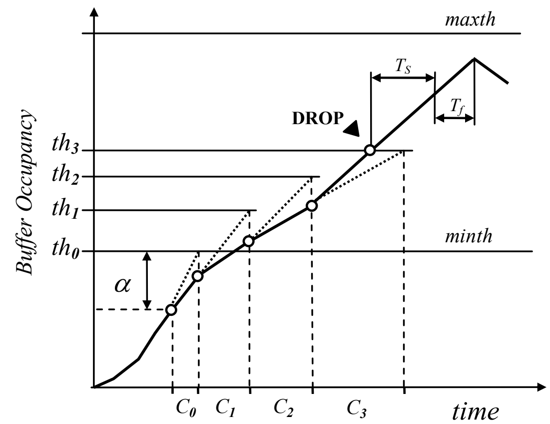

Figure 4. Occupancy curve of a buffer implementing SBD.

```
for each packet arrival
    update BO
    if timer Cn on
        if BO ≥ thn
            drop packet
            deactivate timer Cn
            start timer Tr
    if timer Cn off
        if BO ≥ minth - α
            calculate Cn
            timer Cn ← Cn
            thn ← BO + α
            start timer Cn

upon timer (Cn or Tr) expiration
    if BO ≥ minth - α
        calculate Cn
        timer Cn ← Cn
        thn ← BO + α
        start timer Cn
```

Figure 5. Pseudocode of the SBD algorithm

4.1 Description of the Algorithm

Figure 4 shows a graphical example of the algorithm. For the sake of clarity, the curve that represents the buffer occupancy (*BO*) is always increasing until a certain point, where it starts to decrease.

In reality, there is a fast oscillation around the average occupancy caused by instantaneous arrival and departure of packets.

When BO is higher or equal to $minth - \alpha$, the algorithm calculates the time C_n that it will take the buffer to store α additional bits at a filling rate equal to r_c ($C_n = \alpha / r_c$). The threshold level for the measuring interval is also determined ($th_n = BO + \alpha$). The value of r_c depends linearly on th_n according to its definition and expressed in (1).

$$r_c = \frac{maxth - th_n}{T_r} \tag{1}$$

A timer for C_n is started. If the timer expires and BO is below the threshold th_n, then the actual filling rate is lower than r_c, and no packet is dropped. If th_n is reached before the expiration of the timer, the current filling rate surpasses r_c and therefore a packet will be discarded.

In Figure 4, the dotted segments starting at each measuring interval represent the buffer filling at the critical rate (critical curve). In the intervals C_0, C_1 and C_2, no packet is discarded because the BO curve is below the critical curve. In contrast, in the C_3 period, BO is above r_c. Hence, the threshold th_3 is reached before the timer C_3 expires. When th_3 is reached, a packet is dropped. The monitoring and discarding algorithm is deactivated for a period T_r, avoiding consecutive packet discards. A description of the algorithm in pseudocode can be seen in Figure 5.

5. PARAMETER CONFIGURATION

In most buffer management schemes, parameter configuration is a key issue, and SBD is not an exception. In order to know the influence of each parameter in the overall performance, we have performed extensive simulation experiments. In many Internet applications, e.g. downloading and web surfing, goodput and latency are an accurate measure of the user perceived QoS.

5.1 Simulation Environment

The simulation environment employed in this research has been developed in OMNeT++ [13] and includes a complete implementation of TCP and RLC protocols. The simulation scenario consists of one or several TCP sources connected to their respective receivers in a mobile equipment. The end-to-end connection consists of two sections, the wired network and the radio bearer. The wired network comprises the Internet and the 3G core network. The radio bearer has a round trip time of 50 ms [3] and a bidirectional nominal rate of 384 kbit/s, representing the bottleneck link, which is the situation expected in most cases [12]. The wired network is modeled with a 1 Mb/s link with a round trip delay of 200 ms [3].

The wireless channel generates error bursts according to the model described in [14] where the Doppler frequency, f_d, of the user equipment determines the average burst length. Lower f_d causes longer bursts of errors. It is usual to employ the normalized Doppler frequency, equal to the product of f_d and the radio frame duration (10 ms).

In order to obtain more realistic results, the error probability in our model is the same in the uplink and in the downlink direction. The frame loss ratio is 10%, a typical UMTS design value [3].

The simulation results exposed in this paper are obtained averaging 20 runs per sample. Each run is a 60 seconds download session. The radius of the confidence interval is less than 3% of the magnitude of the averaged sample, with a confidence degree of 90% according to a t-student distribution.

The TCP flavour employed is TCP Reno, one of the most extended in the Internet [12]. RLC and TCP configuration is shown in Table 1.

5.2 Simulation Results

In our performance simulation analysis, multiple parameter combinations of SBD were tested, in order to disclose the efect of each one in the end-to-end performance. The following values were evaluated:

- Distance between *maxth* and *minth* (Δ): 10, 15, 20, 25, 30 and 35 SDUs. (*maxth* is equal to the buffer size)
- T_r: from 50 ms to 900 ms in steps of 50 ms.

The value of α is equal to the size of 5 maximum-sized SDUs (1500 bytes), which was found to be an optimum value and a compromise value between fast detection of congestion and excessive sensitivity to occupancy oscillations. The capacity of the buffer is 40 SDUs, enough to prevent buffer overflow, keeping a low delay.

For the sake of generality, the number of multiplexed flows ranged form 1 to 4. From these experiments we derive the following conclusions regarding each parameter:

1) T_r: This parameter has a direct impact on the discarding policy. A higher T_r implies a more aggressive discarding policy, because, according to (1), r_c is reduced, and consequently, the buffer occupancy and the delay decrease. Figure 6(a) shows the effect of T_r on TCP performance for *minth* = 25 SDUs.

2) minth: The effect of this value is tied to T_r. Higher values of T_r result in a greater system sensitivity to *minth*. The goodput and the delay are lower for lower values of *minth*. Figure 6(b) shows the effect of T_r for *minth* = 20 SDUs.

The effect of SBD parameters on TCP performance is predictable and sustained in a wide range of values. This makes SBD suitable for automatic configuration. The RLC layer is able to estimate the wireless round trip time, which is a significant component of T_r. The value of *minth* should be above the Bandwidth Delay Product (BDP) of the path, which can be roughly estimated multiplying T_r by the nominal radio bearer rate. This value assures that the goodput does not fall too much below the optimum, especially when the RLC serves only one flow.

5.3 Implementation Issues

One of the main advantages of SBD compared to other RED-like mechanisms is that it does not need to generate random numbers to compute the discarding probability because of its deterministic operation. This reduces the computational cost of the algorithm, and makes it more feasible for its implementation at the RLC level where the buffering is done in a per-user basis.

The maintenance of a new timer, C_n, does not add too much complexity to the RLC operation, which already handles several timers, e.g. *Poll Timer* and *Status Prohibit Timer* [2]. The RLC

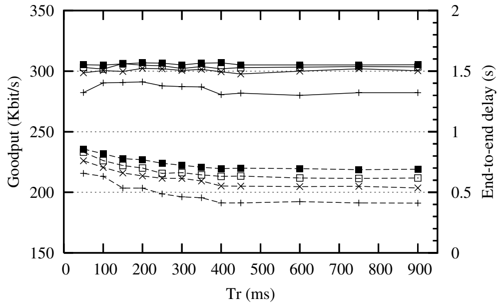

Figure 6(a). SBD: Effect of T_r on TCP (*minth* = 25)

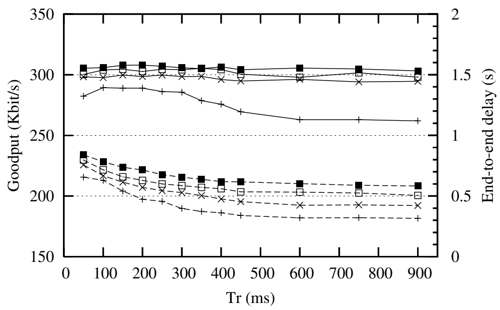

Figure 6(b). SBD: Effect of T_r on TCP (*minth* = 20)

can synchronize C_n to the Transmission Time Interval (TTI) which is equivalent to a clock signal with a granularity of 10 ms (similar to that of other RLC timers).

The protocol could be further simplified in terms of implementation with the use of pre-calculated C_n values. Therefore, the buffer could be considered as divided in gaps, each one having a corresponding C_n value. In our simulations, the number of C_n values equals the number of SDUs fitting in the buffer space delimited by *minth* and *maxth*.

6. PERFORMANCE COMPARISON

In order to measure the benefits of SBD compared to RED, we have also performed extensive tests to find the optimum parameter configuration for RED in the RLC buffer. The following values were evaluated:

- *maxth* – *minth* (Δ): 10, 15, 20, 25 and 30 SDUs.
- *maxp*: from 0.005 to 1.

It should be mentioned that although RED is extensively used for Internet routers, the effect of its parameter configuration in a wireless link layer is still not fully known. According to [7, 9], an RED algorithm at the RLC buffer should not average the buffer occupancy. By using the instantaneous queue size instead, the reaction of the algorithm to sudden increases is faster and, additionally, simplifies its configuration and reduces computing load.

Figure 7 shows TCP performance for *minth* = 25 SDUs in a buffer of 40 SDUs. Considering goodput performance, there is an optimum *maxp* value around 0.05. However, in this point the delay is higher than in the SBD buffer (e.g. for T_r = 350 ms). In RED, the delay is reduced using lower values of *minth*. However, as seen in Figure 8 where *maxp* = 0.05, the RED configuration achieving lower delay values has worse goodput figures than SBD configurations with similar delays.

The diffences in the operation of RED and SBD at an RLC buffer can be seen in figures 9 and 10. Figure 9 shows the trace of a TCP connection over a radio bearer multiplexing 2 TCP flows and a RED buffer with *maxp* = 0.05 and *minth* = 20. In Figure 10, the scenario is similar, but the buffer implements SBD with T_r = 250 ms and *minth* = 15.

The first difference is that RED is less capable to avoid the overflow at the first stages of the connection, when the sources are in the *Slow Start* state. The increment of the buffer occupancy is too steep for RED to react. SBD is specially designed to react against fast and sustained BO increments; therefore the overflow is avoided even in this state.

Once the sources are in *Congestion Avoidance*, RED is able to keep BO below *maxth*, avoiding additional overflows, and achieving better performance than the *Drop Tail* scheme considered in Section 2. However, *BO* reaches higher values with RED than with SBD, even with RED having a lower *minth* value. In consequence, the delay is also higher with RED. SBD is also more capable to keep the source rate arround a certain value, because *cwnd* oscillates with a lower and more constant period.

7. CONCLUSIONS

Our simulation experiments show that SBD is an effective AQM method for buffers at the RLC level in 3G networks. Compared to the drop-tail implementation of RLC, AQM mechanisms reduce the latency and improve the goodput of TCP connections, especially when the radio bearer multiplexes more than one TCP flow.

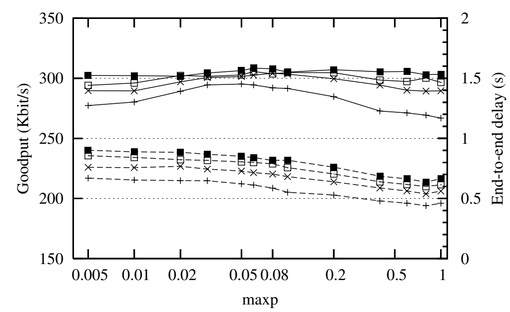

Figure 7. RED: Effect of *maxp* on TCP Performance

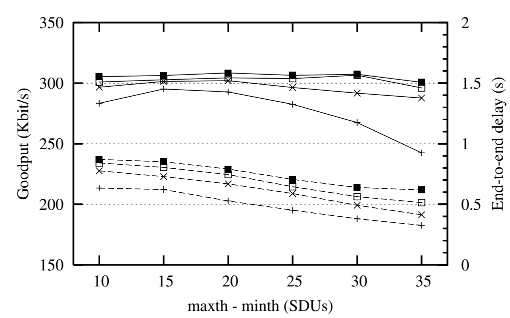

Figure 8. RED: Effect of *minth* on TCP Performance

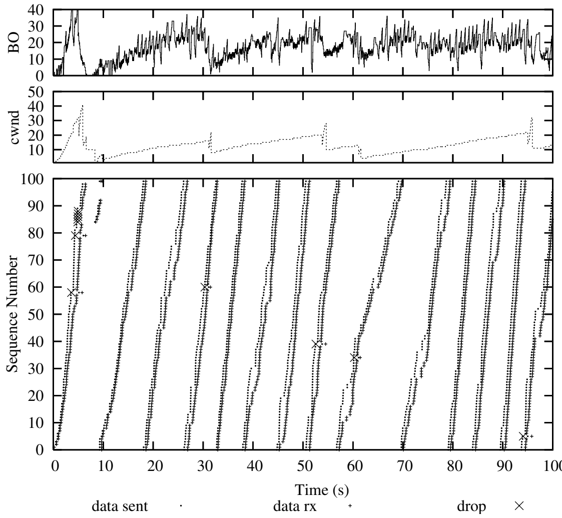

Figure 9. Trace of a TCP flow over RLC with RED.

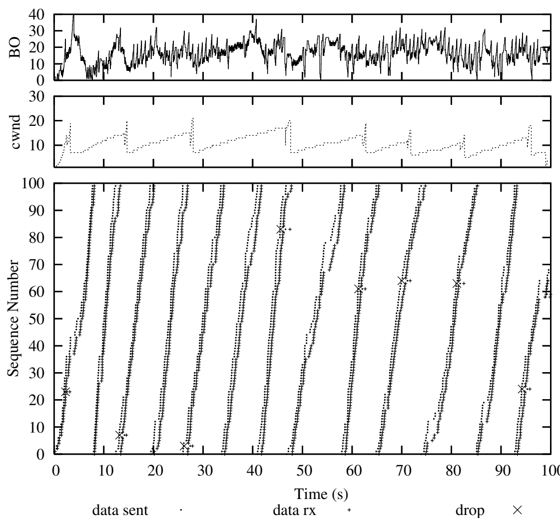

Figure 10. Trace of a TCP flow over RLC with SBD.

SBD is a deterministic AQM protocol which presents several advantages compared to RED: is more effective in reducing the delay, is easier to implement and configure at the RLC layer and provides more stability to the TCP connection. We provide some advice to reduce the computing load of SBD, which is a key issue, given that RLC operates in a per-user basis. From the simulation results, we have found that, in SBD, while one parameter, T_r, determines the aggressiveness of the protocol in reducing the delay, another parameter, *minth*, acts as a "security" limit, assuring that the goodput does not fall bellow certain threshold.

An interesting future line is the development of algorithms for dynamic configuration of SBD parameters. Following a cross-layer approach, these algorithms should consider changes in the 3G link and receive input data from upper layers, e.g. number of TCP flows and application requirements.

8. ACKNOWLEDGMENTS

This work was supported by the Spanish Inter Ministerial Sciece and Technology Commission under project TEC2005-08068-C04-01/TCM.

9. REFERENCES

[1] V. Tsaoussidis and I. Matta, "Open Issues on TCP for Mobile Computing", *Journal of Wireless Communications and Mobile Computing*, Wiley Academic Publishers, Issue 2, vol.2, Feb. 2002.

[2] 3GPP TS 25.322, "Radio Link Control (RLC) protocol specification", v. 6.4.0., Jun. 2005.

[3] M. Meyer, J. Sachs and M. Holzke, "Performance Evaluation of a TCP Proxy in WCDMA Networks", *IEEE Wireless Communication*, Oct. 2003, 70-79.

[4] H. Inamura et al., "TCP over Second (2.5G) and Third (3G) Generation Wireless Networks" IETF RFC 3481, Feb. 2003.

[5] R. Chakravorty, A. Clark and I. Pratt, "Optimizing Web Delivery over Wireless Links: Design, Implementation and Experiencies", *IEEE J. Select. Areas Commun.*, vol. 23, no. 2, Feb. 2005, 402-416.

[6] C. Barakat, E. Altman and W. Dabbous, "On TCP Performance in a Heterogeneous Network: A Survey" *IEEE Commun. Mag.*, vol. 38, no. 1, Jan. 2000, 40-46.

[7] J. J. Alcaraz, F. Cerdan and J. García-Haro, "Optimizing TCP and RLC Interaction in the UMTS Radio Access Network", *IEEE Network*, vol 20, no. 2, 56 - 64.

[8] S. Floyd and V. Jacobson, "Random early detection gateways for congestion avoidance," *IEEE/ACM Trans. Networking*, vol. 1, 397-413, Aug. 1993.

[9] M. Agfors, R. Ludwig, M. Meyer and J. Peisa, "Queue Management for TCP Traffic over 3G Links", in Proc. IEEE WCNC 2003, 1663 -68.

[10] P. Benko; G. Malicsko and A. Veres, "A large-scale, passive analysis of end-to-end TCP performance over GPRS" in *Proc*. *IEEE INFOCOM*, vol: 3, Mar. 2004, 1882 -92.

[11] M. Rossi, L. Scaranari and M. Zorzi, "On the UMTS RLC Parameters Setting and their Impact on Higher Layers Performance", in *Proc*. *IEEE 57th VTC*, vol. 3, Oct. 2003, 1827- 32.

[12] A. Gurtov, S. Floyd, "Modeling Wireless Links for Transport Protocols", *ACM SIGCOMM Computer Communication Review*, vol. 34, no. 2, Apr. 2004, 85-96.

[13] A. Varga, "The OMNeT++ Discrete Event Simulation System", in *Proc. European Simulation Multiconference*. June 2001.

[14] A. Chockalingam and M. Zorzi, "Wireless TCP Performance with Link Layer FEC/ARQ", in *Proc. IEEE ICC*, June.1999, 1212-16.