12th GI/ITG CONFERENCE ON MEASURING, MODELLING AND
EVALUATION OF COMPUTER AND COMMUNICATION SYSTEMS

3rd POLISH-GERMAN TELETRAFFIC SYMPOSIUM

Elastic Token Bucket - A Traffic Characterization for Time-Limited Bursty Traffic

Michael Menth, Jens Milbrandt, and Florian Zeiger
Department of Distributed Systems, Institute of Computer Science, University of Würzburg
Am Hubland, D-97074 Würzburg, Germany
Phone: (+49) 931-8886655, Fax: (+49) 931-8886632
Email:{menth|milbrandt|zeiger}@informatik.uni-wuerzburg.de

Abstract

In this paper, we present the *Elastic Token Bucket* (ETB) as a dynamic traffic characterization for traffic shaping purposes, i.e. for traffic policing and spacing. We describe the ETB and study the impact of its parameters on the resulting packet delay and maximum throughput. In contrast to simple *Constant Bit Rate* (CBR) or *Token Bucket* (TB) based solutions, economy of scale can be achieved regarding queuing time and required buffer space when many ETB-regulated traffic sources are multiplexed.

1 INTRODUCTION

Traffic Shaping (TS) in combination with *Admission Control* (AC) is a feasible means for providing *Quality of Service* (QoS) in *Internet Protocol* (IP) networks. Entering a network at a single ingress, data traffic from multiple sources is first shaped and then multiplexed to be transported through the network as an aggregated stream (cf. Figure 1). While multiplexing a set of traffic streams, we face packet and flow scale congestion problems [1] leading to an increased packet delay induced by queueing at the multiplexer. To avoid large buffers for long multiplexer queues and to reduce the resulting packet delay, different TS schemes have been implemented.

Using *Constant Bit Rate* (CBR) TS, a source is characterized by a maximum bitrate. All traffic exceeding this rate is strictly marked as out-of-profile and it is dropped or delayed. This approach is not appropriate for bursty traffic.

Using conventional bucket-based TS [2, 3, 4], a source is characterized by a mean rate and a maximum burst size. This scheme allows for elastic traffic. However, if all sources send with their mean rate, eventually generate some bursts, and continue sending with mean rate, these burst cumulate in the multiplexer and lead to a flow scale congestion that can be temporally unlimited. The resulting multiplexer queue length affects the traffic of all sources, even of those not contributing to the congestion.

The traffic descriptions of TB-based TS mechanisms are very static. Furthermore, it is not possible to determine
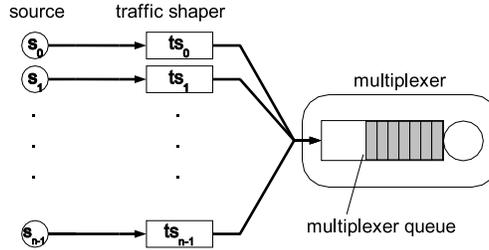
Figure 1: Multiplexing of shaped traffic streams

1. the waiting time distribution for the multiplexing of packets and

2. the multiplexing gain with regard to the required multiplexer buffer size

without further knowledge of the actual source behavior. We therefore introduce the *Elastic Token Bucket* (ETB), a TS mechanism that individually shapes the data streams such that they fulfill certain good-natured traffic characteristics. As a consequence, multiplexing ETB-regulated traffic sources results in a shorter average multiplexer queue length, requires less buffer space and reduces the packet delay. The ETB allows for temporally limited bursts.

The remainder of this work is structured as follows. In Section 2, some basics of TS are introduced and the ETB mechanism is explained. Section 3 is dedicated to the performance evaluation of the ETB. Section 4 summarizes and concludes this work.

2   CONVENTIONAL AND NEW TRAFFIC SHAPING MECHANISMS

In this section, we review well-known traffic characteristics and present the new *Elastic Token Bucket* (ETB).

2.1   Methods of Traffic Shaping

The application of TS requires a traffic description. There are basically two different approaches for TS: *Policing* and *Spacing*. A policer tests the data stream for conformance regarding its traffic description and drops all non-conform packets to achieve conformance. Moreover, policing can also be used to achieve almost equal bandwidth allocation for a set of competing flows [5]. The disadvantage of this method is packet loss and the subsequently necessary retransmission of dropped packets. A spacer achieves conformance of a data stream according to its traffic description by delaying non-conform packets. The disadvantage of this method is packet delay. For instance in *Asynchronous Transfer Mode* (ATM) [6], traffic leaving a network is spaced to avoid packet loss by the ingress policer of the neighboring network [7, 8, 9, 10].

2.2 Conventional Traffic Description

With *Constant Bit Rate* (CBR), traffic must not exceed a specified maximum bitrate R. Given the length L of a sent packet, the minimum temporal distance to the next packet is L/R. This method is appropriate for smooth traffic streams but provokes significant loss or dealy of packets when it is used for policing or shaping of bursty traffic. Setting R to a larger value alleviates that problem but increases transmission costs.

The *Token Bucket* (TB) traffic description [3] consists of a single bucket with size B. Initially, the bucket fill level $b(t)$ equals the bucket size B, i.e. the bucket is full of tokens. If the bucket is not full ($b(t) < B$), tokens are refilled at a constant rate r_{tb}. A packet of length L can only be sent at time t if there are enough tokens in the bucket, i.e. if $\min(b(t) + (t - t_{la}) \cdot r_{tb}, B) \geq L$, where t_{ls} denotes the time of the last packet arrival. Otherwise the whole packet is tagged as non-conform. The token refill rate r_{tb} determines the long-term transmission rate while the bucket size B limits the maximum burst size of the data stream. This type of traffic description allows for bursty traffic and is used e.g. in ATM to control *Peak Cell Rate* (PCR) and *Cell Delay Variation Tolerance* (CDVT). The *Leaky Bucket* (LB) [2] is a similar mechanism providing the same results but based on a different idea.

The *Dual Token Bucket* (DTB) consists of two separated TBs monitoring concurrently similar traffic characteristics on different time scales. The DTB is used, e.g., in ATM networks for monitoring PCR, CDVT, *Sustainable Cell Rate* (SCR), and *Maximum Burst Size* (MBS) of variable bitrate traffic. While the first bucket monitors the parameters PCR and CDVT on a short time scale, the second bucket monitors SCR and MBS on a long time scale. The parameters of both TBs can be individually set to control the traffic as desired.

2.3 Elastic Token Bucket

The *Elastic Token Bucket* (ETB) consists of two TBs, a *Real Token Bucket* (RTB) and a *Virtual Token Bucket* (VTB). The RTB works like a conventional TB. Its bucket size is B_r, its refill rate is r_r, and its fill level is called b_r. The VTB is used to "punish" a source if its data stream exceeds the mean rate. The VTB has a bucket size B_v, a token refill rate r_v, and a fill level named b_v.

The RTB and the VTB are coupled by a mathematical function and this idea causes the elastic behavior of the ETB. Specifically, the token refill rate r_v of the VTB depends on the fill level b_r of the RTB. This makes the ETB sensitive to bursty traffic.

A packet of length L can only be sent if there are enough tokens in both TBs, i.e. $b_r \geq L$ and $b_v \geq L$. If consecutive bursts decrease the fill level of the RTB substantially, the refill rate r_v of the VTB is decreased such that the VTB runs empty after a while. Then, the reduced refill rate r_v also throttles the rate of the ETB. The RTB can thus recover from the bursts which also increases the refill rate r_v and the overall ETB rate again.

In the following, we parameterize the throttle function that with angle α (cf. Figure 2). The figure shows the relation between the fill level b_r of the RTB and the refill rate r_v of the VTB that is depicted as a value normalized by the constant refill rate r_r of the RTB. For increasing α ($0° \leq \alpha \leq 90°$), the throttle function becomes more stringent, i.e. the sources are slowed down earlier and stronger.
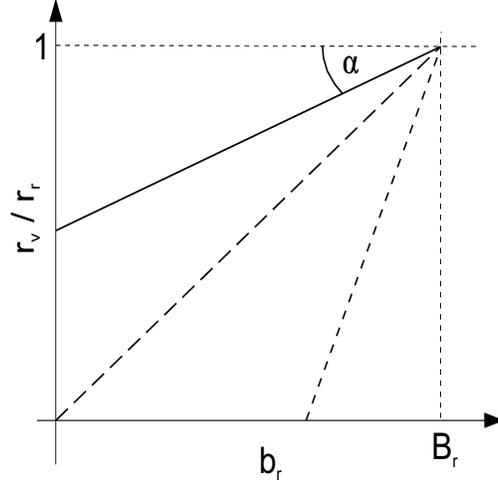
Figure 2: Throttle function parameterized with α

3 PERFORMANCE EVALUATION OF THE ELASTIC TOKEN BUCKET (ETB)

In our performance studies, we consider ETB-regulated traffic sources. The bucket sizes B_r and B_v are always set equal. First, the general behavior of the ETB is analyzed by means of a single source TS scenario. Then, the multiplexing of multiple ETB-regulated traffic sources is examined. We finally show the impact of the ETB on the economy of scale regarding the multiplexer buffer space.

3.1 Analysis of the ETB Behavior

We consider saturated traffic sources, i.e. they are only limited by the ETB mechanism. At time $t = 0$, the source generates a burst of maximum size and continues sending as many packets as possible, i.e. with rate r_r. After the ETB accepts the burst, all tokens in the RTB and the VTB are exhausted and the ETB output rate is reduced to a minimum that depends on the parameter α of the throttle function (cf. Figure 2). In Figure 3 the development of the ETB output rate over time is illustrated for parameters $\alpha \in \{0°, 30°, 60°\}$. For $\alpha = 0°$, there is no throttle effect and the ETB equals a normal TB, i.e. after the burst, the ETB output rate is r_r. For more stringent throttle functions with angles $\alpha = 30°$ and $\alpha = 60°$, we observe a definite temporary reduction of the ETB output rate. The recover time R of the ETB, i.e. the time until the traffic source can send again with rate r_r, depends on the parameter α. In Figure 4, the recover time R is plotted against the angle α. The recover time R is expressed in B_r/r_r for comparison reasons. The recover time R decreases for higher values of α, i.e. for more stringent throttle functions.
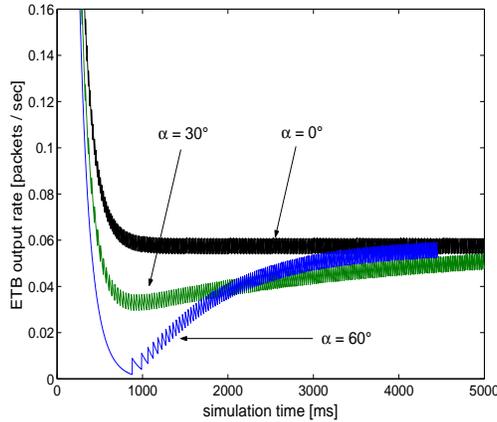
Figure 3: ETB rate reduction depending on α

3.2 Analysis of Multiplexed ETB-Controlled Traffic

For analyzing the impact of the ETB on traffic multiplexing, we refer to the simulation scenario illustrated in Figure 1. For our analysis, each source generates a data stream with hyper-exponentially distributed packet inter-arrival times. The average rate r_{avg} of this stream is set to 95% of the token refill rate r_r of the RTB, i.e. $r_{avg} = 0,95 \cdot r_r$. Each stream is shaped using an ETB spacer and all n data streams are finally multiplexed onto a common traffic trunk with rate $r_{tt} = n \cdot r_r$. For the simulation, we set parameter $n = 15$. The input queue buffers for the ETBs and the multiplexer are set unlimited such that the number of packets in the ETB and multiplexer queues can be monitored.

Figure 5 shows the average ETB buffer occupation of a single source and the average multiplexer queue length in packets depending on parameter α of the throttle function. For $\alpha = 0$, the ETB acts like a conventional TB. For higher values of $\alpha \rightarrow 90°$, the behavior of the ETB becomes similar to CBR TS. With increasing α, i.e. with a more stringent throttle function, the packet buffering is shifted from the multiplexer to the ETB. Concerning the two buffer sizes, the ETB allows for an α-dependent interpolation between the TB and the CBR TS mechanism. Therefore, higher values of α steadily reduce the impact of traffic of bad-natured sources on the overall packet delay induced by the multiplexer queueing. As a consequence, good-natured sources are not unnecessarily delayed.
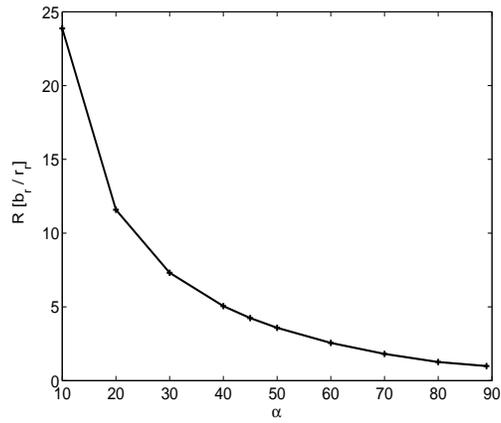
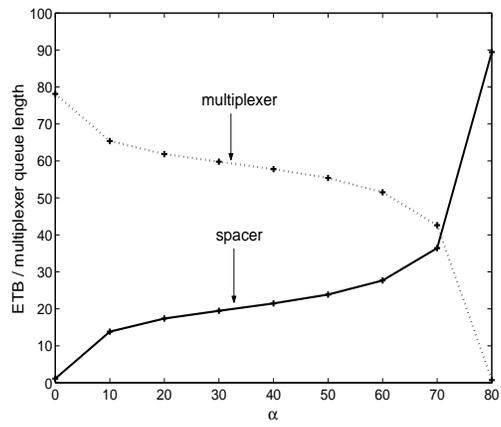Figure 4: ETB recover time depending on α

Figure 5: ETB and multiplexer queue length

3.3 Impact of the ETB on the Economy of Scale of Buffer Space

The simulation model used for studying the impact of the ETB on the economy of scale of the required multiplexer buffer capacity is still based on the scenario presented in Figure 1. A single simulation run is performed for each $\alpha \in \{0°, 20°, 40°, 60°, 80°\}$. In each simulation, $n = 100$ ETB-controlled data streams sending with rate r_r are multiplexed to a single output stream. The recover time R of the ETB is used as the interval in which the n sources emit a burst of maximum size. After emitting a burst all sources continue sending with full rate r_r. It is thus guaranteed that all ETBs and the multiplexer are continuously working under full load.

Figure 6 shows the complementary cumulative distribution function (CCDF) of the packet waiting time probability in the multiplexer. Here, the packet waiting time probability $P(W > t)$ is depicted depending on the normalized waiting time of a packet. The normalization factor is $1/W_{max}$, with W_{max} being the maximum packet waiting time in the multiplexer with buffer capacity $B_{mx} = n \cdot B_r$.
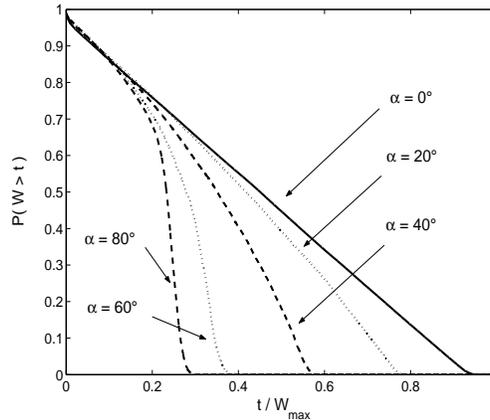

Figure 6: CCDF of the packet waiting time probability

The curves for the different values of α in Figure 6 admit the following conclusions. The probability of a high multiplexer buffer occupation - equal to high values of t/W_{max} - clearly decreases for larger values of α and thus for more stringent throttle functions. The resulting shorter multiplexer queue reduces the packet waiting time of all data streams. Further simulations reveal that the effect of multiplex gain becomes more definite, the more sources are multiplexed onto a single aggregate.

Figure 7 shows the quantiles of the packet waiting time distribution. Here, the 90%, 99%, 99.9%, and 99.99% quantile is shown for different settings of parameter α. The y-axis reflects the maximum packet waiting time normalized by W_{max} and indicates also the the respective quantile of the average occupation of multiplexer buffer space B_{mux}.

Figure 7 can also be interpreted in the following way. If the multiplexer buffer size B_{mux} is set to an α-corresponding value of one of the curves, it is smaller than required by the worst case, i.e.
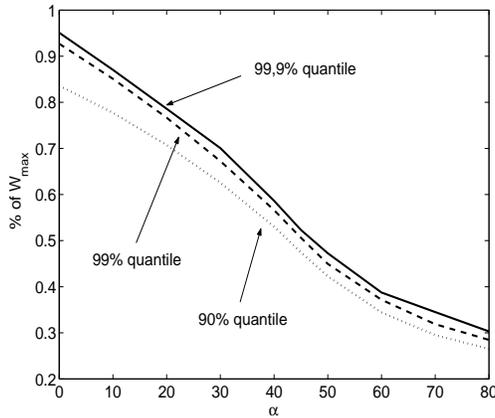
Figure 7: Quantiles of the packet waiting time distribution

$B_{mux} < n \cdot B_r$, but the 10%, 1%, 0.1%, and 0.01% packet loss can be guaranteed without knowing the actual traffic characteristics. The fact that increasing parameter α lowers the quantiles can be interpreted as multiplex gain.

4   CONCLUSION

We presented the *Elastic Token Bucket* (ETB) as a traffic description for policing and spacing purposes. The ETB enforces a rate reduction if a flow exceeds its mean rate extensively. This property is parameterizable such that the ETB can interpolate between a *Token Bucket* (TB) and a *Constant Bit Rate* (CBR) spacer with regard to spacer delay. When TB-regulated traffic sources are multiplexed onto a common trunk, the required multiplexer buffer is the sum of the bucket sizes of all TB-regulated flows. In contrast, the ETB allows for a realization of multiplexing gain, i.e. a certain multiplexer queue occupation will be exceeded only with a small probability. This is independent of the actual source characteristic and helps to reduce buffer space and packet waiting time in multiplexer devices. After all, the ETB is a smart mechanism to shape data flows such that they are better-behaved than TB-shaped streams and can thus contribute to the protection of well-behaved flows with *Quality of Service* (QoS) requirements.

REFERENCES

[1] J. Roberts, U. Mocci, and J. Virtamo, *Broadband Network Teletraffic - Final Report of Action COST 242*. Berlin, Heidelberg: Springer, 1996.

[2] J. Turner, "New Directions in Communications (or Which Way to the Information Age?)," in *IEEE Communications Magazine*, October 1986.

[3] A. S. Tanenbaum, *Computer Networks*. Prentice Hall PTR, August 2002.

[4] R. Steinmetz, *Multimedia Technologie: Grundlagen, Komponenten und Systeme*. Springer, Heildelberg, 1999.

[5] J. Kidambi, D. Ghosal, and B. Mukherjee, "Dynamic Token Bucket (DTB): A Fair Bandwidth Allocation Algorithm for High-Speed Networks," *Journal of High-Speed Networks*, 2000.

[6] O. Kyas and G. Crawford, *ATM Networks*. Prentice Hall PTR, May 2002.

[7] P. E. Boyer, F. M. Guillemin, M. J. Servel, and J.-P. Coudreuse, "Spacing Cells At The ATM Network Entry Points," *IEEE Network Magazine, Vol.6, No.5*, pp. 38–49, September 1992.

[8] F. Hübner and P. Tran-Gia, "Discrete-Time Analysis of Cell Spacing in ATM Systems," *Telecommunication Systems*, pp. 3:379–395, 1995.

[9] M. R. Frater and O. Rose, "Cell Loss Analysis of Broadband Switching Systems Carrying VBR Video," *Telecommunication Systems*, pp. 4:319–338, 1995.

[10] M. Ritter, "Performance Analysis of the Dual Cell Spacer in ATM Systems," in *Proc. IFIP 6th International Conference on High Performance Networking*, (Palma de Mallorca, Spain), pp. 345–362, September 1995.