

University of Würzburg
Institute of Computer Science
Research Report Series

**Performance of Experience-Based
Admission Control in the Presence
of Traffic Changes**

Jens Milbrandt, Michael Menth, Jan Junker

Report No. 381

August 2006

University of Würzburg
Institute of Computer Science
Chair of Distributed Systems
Würzburg, Germany
{milbrandt,menth,junker}@informatik.uni-wuerzburg.de

Performance of Experience-Based Admission Control in the Presence of Traffic Changes

Jens Milbrandt, Michael Menth, Jan Junker

University of Würzburg

Institute of Computer Science

Chair of Distributed Systems

Würzburg, Germany

{milbrandt,menth,junker}@informatik.uni-wuerzburg.de

Abstract

This paper investigates the transient behavior of *experience-based admission control* (EBAC) in case of traffic changes. EBAC is a robust and resource-efficient admission control (AC) mechanism used for reservation overbooking of link capacities in packet-based networks. Recent analyses gave a proof of concept for EBAC and showed its efficiency and robustness through steady state simulation on a single link carrying traffic with constant properties. The contribution of this paper is an examination of the memory from which EBAC gains its experience and which strongly influences the behavior of EBAC in both, stationary and non-stationary state. For the latter, we investigate the transient behavior of the EBAC mechanism through simulation of strong traffic changes which are characterized by either a sudden decrease or increase of the traffic intensity. Our results show that the transient behavior of EBAC partly depends on its tunable memory and that it copes well with even strongly changing traffic characteristics.

1 Introduction

Internet service providers operating next generation networks (NGNs) are supposed to offer quality of service (QoS) to their customers. As packet-based Internet protocol (IP) technology becomes more and more the basis of these networks, QoS in terms of limited packet loss, packet delay, and jitter is required to support real-time services. There are two fundamentally different methods to implement QoS: capacity overprovisioning (CO) and admission control (AC). With CO the network has so much capacity that congestion becomes very unlikely [1, 2], but this also implies that its utilization is very low even in the busy hour. Although CO is basically simple, it requires traffic forecasts and capacity provisioning must be done on a medium or long time scale.

Parts of this work were funded by the Bundesministerium für Bildung und Forschung of the Federal Republic of Germany (Förderkennzeichen 01AK045) and Siemens AG, Munich, Germany. The authors alone are responsible for the content of the paper.

In contrast, AC works on a smaller time scale. It grants access to flows with QoS requirements if the network load is sufficiently low and rejects excessive flow requests to shelter the network from overload before critical situations can occur. QoS is thus realized by flow blocking during overload situations. Compared to CO, AC requires less capacity and yields better resource utilization at the expense of more signaling, coordination and state management [3, 4, 5] especially in the context of a network-wide AC. In this work, we focus on link-based AC, i.e., on methods that protect a single link against overload. These methods are usually extended for application in entire networks. There are various approaches towards AC that can coarsely be classified into parameter-based AC (PBAC) [6, 7, 8], measurement-based AC (MBAC) [9, 10, 11, 12, 13, 14], and derivatives thereof.

PBAC, also known as (a priori) traffic-descriptor-based AC, is an approach appropriate for guaranteed network services [15], i.e., for traffic with imperative QoS requirements. It relies solely on traffic descriptors that are signaled by a source/application and describe the traffic characteristics of a flow such as the mean and the peak rate together with token bucket parameters. If an admission request succeeds, bandwidth is reserved and exclusively dedicated to the new flow. As a consequence, PBAC is often inefficient regarding its resource utilization since the traffic descriptors usually overestimate the actual rate to avoid traffic delay and loss due to spacing or policing. With PBAC, traffic is limited either by deterministic worst case considerations like network calculus [7, 8] or by stochastic approaches such as effective bandwidth [16]. In addition, PBAC calculations for heterogeneous traffic mixes can be very complex.

MBAC, in contrast, is an AC method adequate for controlled load network services [17], i.e., for traffic with less stringent QoS requirements. It measures the current link or network load in real-time and takes an estimate of the new flow to make the admission decision. The determination of traffic characteristics is thus shifted from a source/application to the network and the specified traffic descriptor, e.g. the peak rate, can be very simple. Most MBAC methods presented in the literature are link-oriented. They measure the current load on a link [11, 12, 13] while others perform measurements of individual flows [9]. Other MBAC approaches, also known as endpoint or probe-based AC [18, 19, 20], work on end-to-end measurements in terms of active or passive probing. All those MBAC methods take advantage of real-time measurements and admit traffic as long as enough capacity is available. The disadvantage of MBAC is its sensitivity to measurement accuracy and its susceptibility to traffic estimation errors which can occur, e.g., during QoS attacks, i.e. when admitted traffic flows are "silent" at the moment and congest the link/network later by simultaneously sending at high bitrate.

There are also hybrid AC approaches, e.g. [21], which combine techniques of PBAC and MBAC in a single AC framework to join the advantages of both AC schemes, i.e., they try to improve the network resource utilization while keeping the packet loss ratio below certain limits. However, these hybrid approaches still rely on real-time measurements or use rather static traffic estimators.

To the best of our knowledge, our method called *experience-based admission control* (EBAC) is the first hybrid AC approach that makes traffic measurements without real-time requirements and uses historical information about previously admitted traf-

fic to make current admission decisions. The EBAC method was first introduced in [22]. With EBAC, a new flow is admitted to a link at time t if its peak rate together with the peak rates of already admitted flows does not exceed the link capacity multiplied by an overbooking factor $\varphi(t)$. The overbooking factor is calculated based on the reservation utilization of the admitted flows in the past. Hence, this method relies on experience. EBAC also requires traffic measurements to compute the reservation utilization. However, these measurements do not have real-time requirements and influence the admission decision only indirectly. A proof of concept for EBAC is given in [23] by simulations and corresponding waiting time analyses of the admitted traffic. In particular, EBAC has been investigated in steady state for traffic with rather static characteristics. MBAC methods are susceptible to traffic changes. Therefore, we investigate in this paper the behavior of EBAC in case of sudden traffic changes and its impact on the EBAC-controlled traffic.

The remainder of this paper is organized as follows. In Section 2, we briefly review the EBAC concept. Section 3 describes our simulation design and the applied traffic model and summarizes related results from previous work. In Section 4, we investigate the behavior of EBAC in case of sudden traffic changes. Section 5 summarizes this work and gives a conclusion.

2 Experience-Based Admission Control (EBAC)

In this section, we briefly review the EBAC concept with emphasis on the EBAC memory which holds the experience used to make AC decisions.

The idea of EBAC is briefly described as follows. An AC entity limits the access to a link l with capacity $c(l)$ and records all admitted flows $f \in \mathcal{F}(t)$ at time t together with their requested peak rates $\{r(f) : f \in \mathcal{F}(t)\}$. When a new flow f_{new} arrives, it requests a reservation for its peak rate $r(f_{new})$. If

$$r(f_{new}) + \sum_{f \in \mathcal{F}(t)} r(f) \leq c(l) \cdot \varphi(t) \cdot \rho_{max} \quad (1)$$

holds, admission is granted and f_{new} joins $\mathcal{F}(t)$. If flows terminate, they are removed from $\mathcal{F}(t)$. The experience-based overbooking factor $\varphi(t)$ is calculated by statistical analysis and indicates how much more bandwidth than $c(l)$ can be safely allocated for reservations. The maximum link utilization threshold ρ_{max} limits the traffic admission such that the expected packet delay W exceeds an upper delay threshold W_{max} only with probability p_W . Details regarding its computation are described in [23].

For the calculation of the overbooking factor $\varphi(t)$, we define the reserved bandwidth of all flows as $R(t) = \sum_{f \in \mathcal{F}(t)} r(f)$ while $C(t)$ denotes the unknown mean rate of the traffic aggregate $\mathcal{F}(t)$. EBAC makes traffic measurements $M(t)$ on the link and collects a time statistic for the reservation utilization $U(t) = M(t)/R(t)$. $U_p(t)$ is the p_u -percentile of the empirical distribution of U and the reciprocal of this percentile is the overbooking factor $\varphi(t) = 1/U_p(t)$. For the computation of the overbooking factor $\varphi(t)$ the following three functional components of the EBAC system are required:

1. *Measurement Process for $M(t)$* : To obtain $M(t)$, we use disjoint interval measurements such that for a time interval I_i with length Δ_i , the measured rate $M_i = \frac{\Gamma_i}{\Delta_i}$ is determined by metering the traffic volume Γ_i sent during I_i .
2. *Statistic Collection $P(t, U)$* : The reservation aggregate $R(t)$ is known from the AC process. The utilizations $U(t)$ are sampled in constant time intervals and are stored as hits in bins of a time-dependent histogram $P(t, U)$. The time-dependent utilization quantile $U_p(t)$ can be derived from $P(t, U)$ by

$$U_p(t) = \min_u \{u : P(t, U \leq u) \geq p_u\}. \quad (2)$$

3. *Statistic Aging Process*: If the traffic characteristics of the traffic aggregate $\mathcal{F}(t)$ change over time, the statistic collection $P(t, U)$ must forget obsolete data to increase the importance of the properties of the new traffic mix. Therefore, we record new samples by incrementing the respective bins by 1 and devalue the contents of the histogram bins in regular devaluation intervals I_d by a constant devaluation factor f_d .

The histogram $P(t, U)$ implements the EBAC memory and the statistic aging process makes this memory forget about reservation utilizations in the past. The devaluation interval I_d and factor f_d yield typical half-life periods T_H after which collected values have lost half of their importance in the histogram. Therefore, we have $\frac{1}{2} = f_d^{T_H/I_d}$ and define the EBAC memory length based on the half-life period

$$T_H(I_d, f_d) = I_d \cdot \frac{-\ln(2)}{\ln(f_d)} \quad (3)$$

of the importance of reservation utilization values stored in the histogram. Various combinations of parameters (I_d, f_d) can lead to the same half-life period.

3 EBAC Performance Simulation

In this section, we first present the simulation design of EBAC on a single link and the traffic model we used on the flow and packet scale level. Afterwards, we summarize recent simulation results of EBAC in steady state.

3.1 Simulation Design

The design of our simulation is shown in Fig. 1. Different types of traffic *source generators* produce flow requests that are admitted or rejected by the *admission control* entity. To make an admission decision, this entity takes the overbooking factor $\varphi(t)$ into account. In turn, it provides information regarding the reservations $R(t)$ to the *EBAC system* and yields flow blocking probabilities $p_b(t)$. For each admitted source, a *traffic generator* is instantiated to produce a packet flow that is shaped to its contractually

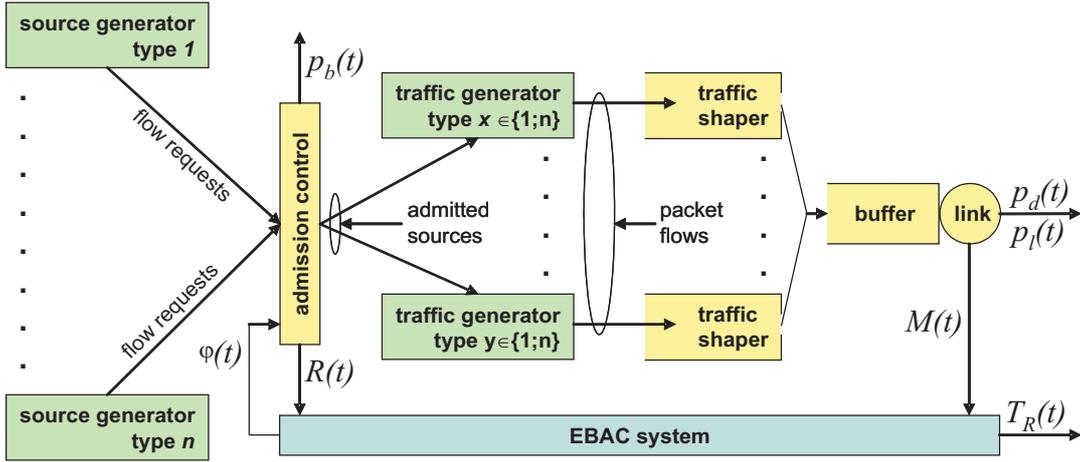


Figure 1: Simulation design for EBAC in steady and transient state.

defined peak rate. Traffic flows leaving the *traffic shapers* are then multiplexed on the buffered link with capacity $c(l)$. The link provides information regarding the measured traffic $M(t)$ to the EBAC system and yields packet delay probabilities $p_d(t)$ and packet loss probabilities $p_l(t)$. The primary performance measure of our non-stationary EBAC simulations is the overall response time T_R , i.e., the time span required by the EBAC system to fully adapt the overbooking factor to a new traffic situation.

3.2 Traffic Model

In our simulations, the admission-controlled traffic is modelled on two levels, namely the flow scale level and the packet scale level. While the flow level controls the inter-arrival times of flow requests and the holding times of admitted flows, the packet level defines the inter-arrival times and the sizes of packets within a single flow.

3.2.1 Flow Level Model

On the flow level, we distinguish different traffic source types, each associated with a characteristic peak-to-mean rate ratio (PMRR, defined below) and corresponding to a source generator type in Fig. 1. The inter-arrival time of flow requests and the holding time of admitted flows both follow a Poisson model [24], i.e., new flows arrive with rate λ_f and the duration of a flow is controlled by rate μ_f , where $1/\mu_f$ denotes the mean flow holding time. For the non-stationary EBAC simulations, the source-type specific parameters λ_f vary over time which directly impacts the load and the composition of the traffic on the link. Provided that no blocking occurs, the overall offered load $a_f = \lambda_f/\mu_f$ is the average number of simultaneous flows measured in Erlang. If not mentioned differently, the holding time of a flow is exponentially distributed with a mean of $1/\mu_f = 90$ s and the traffic load is set to $a_f \geq 1.0$ such that the EBAC-controlled link is saturated with flow requests. The latter assumption allows for an investigation of the EBAC performance under heavy traffic load where some flow requests are rejected.

3.2.2 Packet Level Model

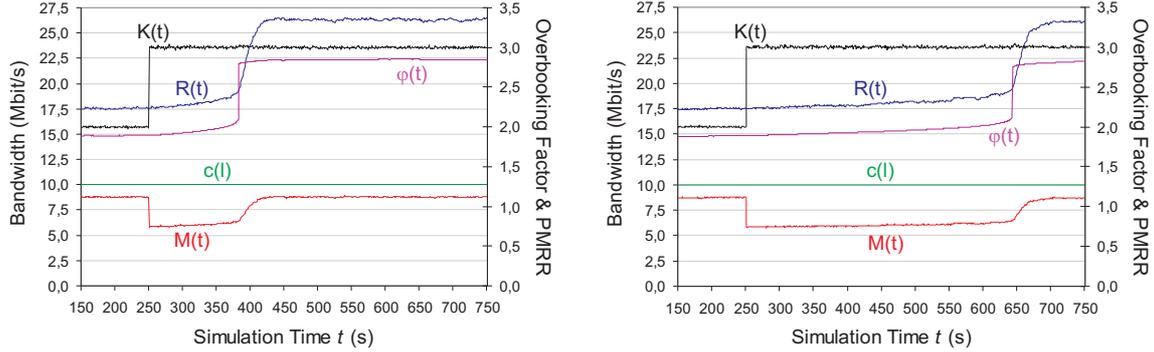
On the packet level model, we abstract from the wide diversity of packet characteristics induced by applications and different transmission protocols. Since we are interested in the basic understanding of the transient behavior of EBAC, we abstract from real traffic patterns and define a flow of consecutive data packets simply by a packet size and a packet inter-arrival time distribution. Both contribute to the rate variability within a flow that is produced by a traffic generator in Fig. 1. To keep things simple, we assume a fixed packet size and use a Poisson arrival process to model a packet inter-arrival time distribution with rate λ_p . We are aware of the fact that Poisson is not a suitable model to simulate Internet traffic on the packet level [25]. We therefore do not take it unconditioned, but use it for the generation of packet streams that are subsequently policed by peak-rate traffic shapers (cf. Fig. 1). The properties of the flows are primarily determined by the configuration of these shapers. In practice, the peak rate $r(f)$ of a flow f is limited by an application or a network element and the mean rate $c(f)$ is often unknown. In our simulations, however, the mean rate is known a priori and we control the rate of flow f by its peak-to-mean rate ratio $k(f) = r(f)/c(f)$. Analogously, $K(t) = R(t)/C(t)$ is the peak-to-mean rate ratio of the entire traffic aggregate $\mathcal{F}(t)$ at time t . It is a natural upper limit for the achievable overbooking factor $\varphi(t)$.

3.3 EBAC in Steady State

The intrinsic idea of EBAC is the exploitation of the peak-to-mean rate ratio $K(t)$ of the traffic aggregate admitted to the link. In [23], we simulated EBAC on a single link with regard to its behavior in steady state, i.e., when the properties of the traffic aggregate were rather static. These simulations provided a first proof of concept for EBAC. We showed for different peak-to-mean rate ratios that EBAC achieves a high degree of resource utilization through overbooking while packet loss and packet delay are well limited. Further simulation results allowed us to give recommendations for the EBAC parameters such as measurement interval length and reservation utilization percentile to obtain appropriate overbooking factors $\varphi(t)$. They furthermore showed that the EBAC mechanism is robust against traffic variability in terms of packet size and inter-arrival time distribution as well as to correlations thereof.

4 EBAC Performance under Traffic Changes

This section discusses the transient behavior of EBAC when the characteristics of the EBAC-controlled traffic change significantly. We investigate the response time T_R of EBAC to provide a new appropriate overbooking factor $\varphi(t)$ after a decrease or increase of the traffic intensity. We consider sudden changes of the traffic characteristics to have worst case scenarios and to obtain upper bounds on the response time. We simulate them with only a single type of traffic flows since only the properties of the entire admitted traffic aggregate are of interest for the calculation of the overbooking factor.



(a) Short memory with half-life period $T_H = 20$ s (b) Long memory with half-life period $T_H = 60$ s

Figure 2: Impact of the EBAC memory on the time-dependent overbooking performance.

4.1 Decrease of the Traffic Intensity

We first investigate the change of the traffic intensity from a high to a low value which corresponds to an increase of the peak-to-mean rate ratio $K(t)$ of the entire traffic aggregate, i.e., all currently and future admitted traffic sources simultaneously reduce their sending rate. For all simulation experiments, we use a link capacity of $c(l) = 10$ Mbit/s and a reservation utilization quantile $p_u = 99\%$. On the flow level, we set the traffic characteristics $\lambda_f = \frac{1}{750 \text{ ms}}$ and $\mu_f = \frac{1}{90 \text{ s}}$ to offer enough traffic to the link, i.e., the link is saturated with traffic. At simulation time $t_0 = 250$ s, the peak-to-mean rate ratio suddenly increases from $K(t) = 2$ to $K(t) = 3$, i.e., all traffic sources slow down and use a lower packet arrival rate λ_p .

Figures 2(a) and 2(b) illustrate simulations averaged over 50 runs for a short and a long EBAC memory. The sudden increase of the peak-to-mean rate ratio results in an immediate decrease of the consumed link bandwidth $M(t)$. As a consequence, the reservation utilization $U(t) = M(t)/R(t)$ also decreases. After a while, the histogram has collected enough low utilization values so that its 99%-percentile $U_p(t)$ decreases which, in turn, leads to a higher overbooking factor $\varphi(t) = 1/U_p(t)$. Hence, more traffic sources are admitted to the link and the reserved rate $R(t)$ rises. Finally, the EBAC system stabilizes again with an expected overbooking factor $\varphi(t) \approx 3$. The speed of the adaptation process is obviously influenced by the EBAC memory.

To measure the duration of the transient phase, i.e., the time until the overbooking factor reaches a new stable value, we calculate the difference between the peak-to-mean rate ratio $K(t)$ and the overbooking value $\varphi(t)$. If $K(t) - \varphi(t) < \varepsilon$, the transition between the two traffic scenarios is considered completed and the EBAC system is in steady state again. We therefore define the EBAC response time

$$T_R = \min \{t_i - t_0 : K(t_i) - \varphi(t_i) < \varepsilon \wedge t_i > t_0\} \quad (4)$$

and set the threshold $\varepsilon = 0.2$ in our simulations. This value is specific to our experiments but seems to be appropriate with regard to the asymptotic convergence of $\varphi(t)$ to $K(t)$.

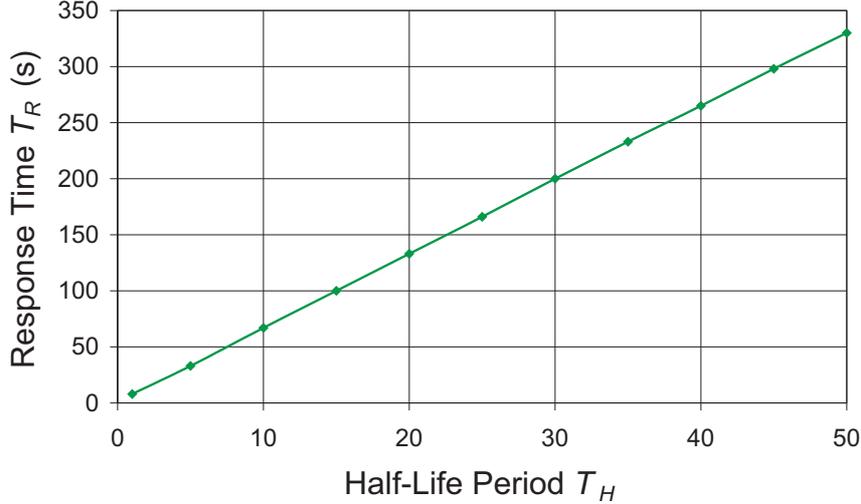


Figure 3: Correlation of EBAC memory with half-life period T_H and EBAC response time T_R for decreasing traffic intensity.

Due to the design of the EBAC mechanism, $\varphi(t) \leq K(t)$ always holds. The statistical significance of our results is assured by calculating the 95%-confidence intervals of the overbooking factor $\varphi(t)$ within 50 iterations of the simulation. As a result, the confidence intervals turn out to be so narrow that we omit them in Fig. 2 for the sake of clarity.

The different progressions of the overbooking factor $\varphi(t)$ in Fig. 2(a) and 2(b) show that in this experiment, the EBAC response time T_R strongly depends on the EBAC memory length represented by the half-life period T_H . To investigate the correlation between T_R and T_H , we perform a series of experiments with varying half-life periods and measure the EBAC response times. Figure 3 shows that there is an almost linear dependency between the EBAC response time T_R and the half-life period T_H of the EBAC memory.

4.2 Increase of the Traffic Intensity

We now change the traffic intensity from a low to a high value which corresponds to a decrease of the peak-to-mean rate ratio $K(t)$ of the entire traffic, i.e., all admitted and future traffic sources raise their sending rate simultaneously which corresponds to a collaborative QoS attack. In contrast to the previous experiment, the QoS is at risk since the link suddenly gets overloaded and the packet delay and flow blocking probabilities increase as could be expected for a QoS attack. To blind out the impact of the link buffer on the EBAC response time, we set its value to infinity.

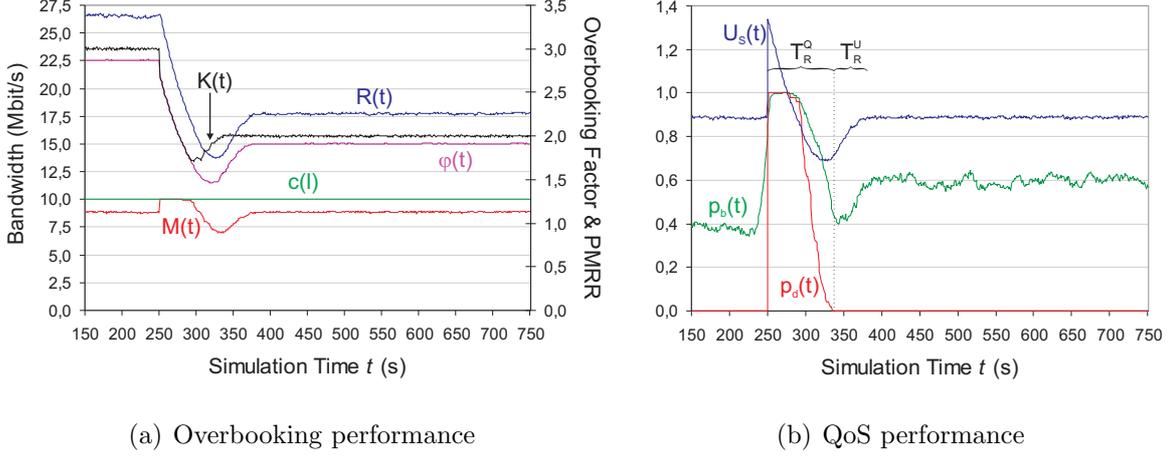


Figure 4: Time-dependant EBAC performance during a QoS attack for a short EBAC memory with half-life period $T_H = 5.76$ s.

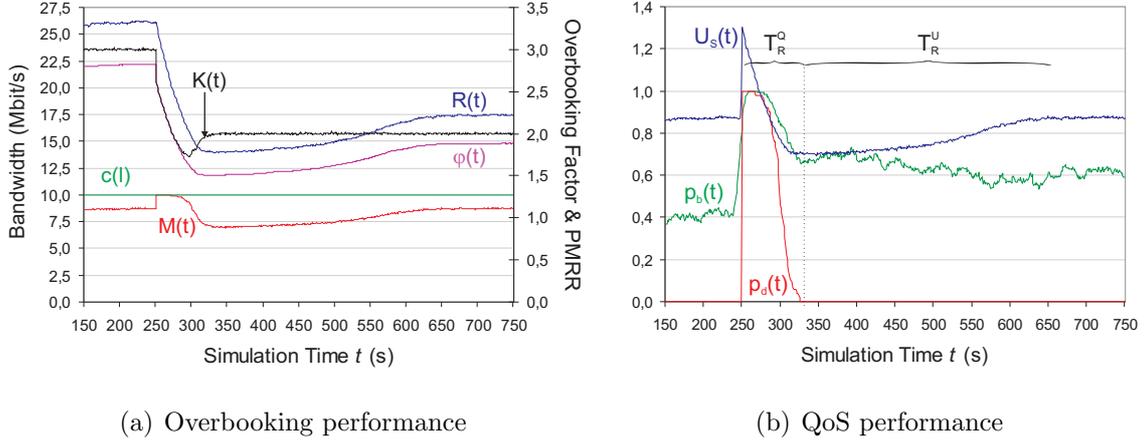


Figure 5: Time-dependant EBAC performance during a QoS attack for a long EBAC memory with half-life period $T_H = 65.79$ s.

The QoS attack experiment is designed analogously to the decrease of the traffic intensity in Sec. 4.1. The only difference is that the packet arrival rates λ_p of all flows are increased such that the peak-to-mean rate ratio decreases from $K(t) = 3$ to $K(t) = 2$. Figures 4 and 5 show the overbooking and QoS performance of the EBAC system for short and long EBAC memory, respectively.

At time $t_0 = 250$ ms the QoS attack starts. As the link becomes overutilized, the fill level of the link buffer increases and the probability for excessive packet delay $p_d(t) = P(\text{packet delay} > 50\text{ms})$ and the flow blocking probability $p_b(t)$ raise to 100% (cf. Fig. 4(b) and 5(b)). As another consequence, the overbooking factor $\varphi(t)$ decreases due to a rising reservation utilization quantile $U_p(t)$ and all newly arriving

flows are blocked by EBAC. As time goes by, some admitted flows expire and their reserved bandwidth is released. However, the overbooking factor $\varphi(t)$ is further decreased as long as the packet delay and the link load are high. Hence, the overbooking factor decreases below its target value of $\varphi(t) \approx 2$ (cf. Fig. 4(a) and 5(a)). When enough flows have expired, the link buffer empties and the QoS is restored as a result of the decreased overbooking factor. Figures 4(b) and 5(b) show that the QoS recovery duration T_R^Q is almost the same for short and long EBAC memory, respectively. After a certain time span T_R^U , the overestimated reservation utilizations in the histogram are faded out by statistic aging. Simultaneously, the overbooking factor $\varphi(t)$ and the link utilization $U_S(t)$ converge to stable values when the EBAC system reaches its steady state again. In contrast to Equ. 4, we now define the EBAC response time as

$$T_R = T_R^Q + T_R^U \quad (5)$$

where interval $T_R^Q = \min\{t_i - t_0 : p_d(t_i) = 0 \wedge t_i > t_0\}$ and time span $T_R^U = \min\{t_j - (t_0 + T_R^Q) : K(t_j) - \varphi(t_j) < \varepsilon \wedge t_j > t_0 + T_R^Q\}$. Our simulation results compiled in Fig. 6 show that the EBAC memory length represented by the half-life period T_H influences its overall response time T_R after a QoS attack. However, it does not influence the time T_R^Q that is required to recover the QoS. We conclude that the EBAC memory length influences the adaptation speed of the overbooking factor only for a decrease of the traffic intensity. Hence, the parameter T_H provides a means for configuring the conservativeness of the EBAC system.

For the sake of completeness, we performed further QoS attack experiments to investigate the impact of the link buffer size, the mean flow holding time, and the link capacity on the transient behavior of EBAC. Details on these experiments are omitted due to the lack of space. We just summarize their results as follows: T_R^Q and T_R^U both depend on the buffer size B and the mean flow holding time $1/\mu_f$. Hence, larger buffers and longer flow holding times extend T_R . In contrast, the link capacity $c(l)$ has no effect on the overall EBAC response time T_R . The above statements hold for arbitrary settings of T_H .

5 Conclusion

In this paper, we investigated the transient behavior of experience-based admission control (EBAC) on a single link.

EBAC is a new link admission control paradigm [22] and represents a hybrid solution between parameter-based and measurement-based admission control. We briefly reviewed the EBAC system whose proof of concept was already given in [23] for traffic with stationary characteristics. We explained the simulation design and the traffic model used for the analyses of the transient behavior of EBAC. We finally simulated EBAC under extremely changing traffic conditions and showed the results which build the main contribution for this paper.

As EBAC partly relies on traffic measurements, it is susceptible to changes of the traffic characteristics. There are certain influencing parameters coupled with this problem. One

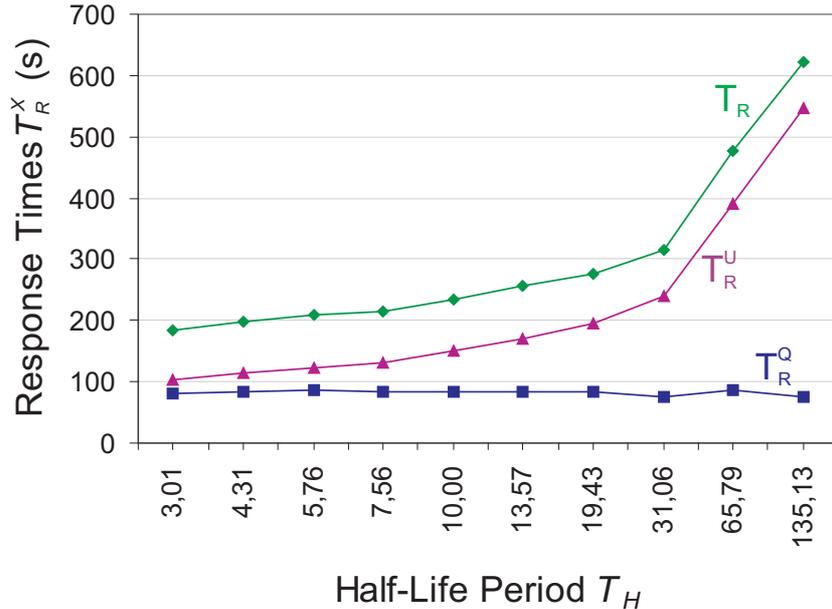


Figure 6: Correlation of EBAC memory with half-life period T_H and EBAC response times T_R^X for increasing traffic intensity.

of them is the length of the EBAC memory which has been defined by its half-life period T_H . We tested the impact of the EBAC memory on a sudden decrease and increase of the traffic intensity which has been expressed by the change of the peak-to-mean rate ratio of the simulated traffic flows. We showed that, for a changing traffic intensity, the response time T_R required to adapt the overbooking factor to the new traffic situation depends linearly on the half-life period T_H . For decreasing traffic intensity, the QoS of the traffic was not at risk. For a suddenly increasing traffic intensity, however, it was compromised for a certain time span T_R^Q which was less than the average flow holding time. Note that the respective experiment used an unlimited link buffer and investigated the performance of EBAC under very extreme traffic conditions that correspond to a collaborative and simultaneous QoS attack by all traffic sources.

Currently, we are working on an EBAC extension, called type-specific overbooking, that provides different overbooking factors for traffic from different applications. In future investigations, the performance of EBAC may be studied with real traffic traces and the concept may be extended to a network-wide scope.

References

- [1] R. Martin, M. Menth, and J. Charzinski, "Comparison of Link-by-Link Admission Control and Capacity Overprovisioning," in *Proc. of 19th International Teletraffic Congress (ITC19)*, (Beijing, China), Aug 2005.

- [2] R. Martin, M. Menth, and J. Charzinski, “Comparison of Border-to-Border Budget Based Network Admission Control and Capacity Overprovisioning,” in *Proc. of 4th International IFIP Networking Conference*, (Waterloo, Canada), May 2005.
- [3] R. Braden, L. Zhang, S. Berson, S. Herzog, and S. Jamin, “RFC 2205: Resource ReSerVation Protocol (RSVP) – Version 1 Functional Specification,” Sep. 1997.
- [4] R. Yavatkar, D. Pendarakis, and R. Guerin, “RFC 2753: A Framework for Policy-Based Admission Control,” Jan. 2000.
- [5] R. Yavatkar, D. Hoffman, Y. Bernet, F. Baker, and M. Speer, “RFC 2814: SBM (Subnet Bandwidth Manager): A Protocol for RSVP-Based Admission Control over IEEE 802-Style Networks,” May 2000.
- [6] J. Wroclawski, “RFC 2210: The Use of RSVP with IETF Integrated Services,” Sep. 1997.
- [7] J. L. Boudec, “Application of Network Calculus to Guaranteed Service Networks,” *IEEE Transactions on Information Theory*, vol. 44, no. 3, May 1998.
- [8] M. Fidler and V. Sander, “A Parameter Based Admission Control for Differentiated Services Networks,” *Computer Networks*, vol. 44, no. 4, pp. 463–479, 2004.
- [9] R. Gibbens and F. Kelly, “Measurement-Based Connection Admission Control,” in *Proc. of 15th International Teletraffic Congress*, (Washington D. C., USA), Jun 1997.
- [10] S. Jamin, S. Shenker, and P. Danzig, “Comparison of Measurement-Based Call Admission Control Algorithms for Controlled-Load Service,” in *Proc. of IEEE Infocom 2000*, pp. 973–980, Mar 1997.
- [11] M. Grossglauser and D. Tse, “A Framework for Robust Measurement-Based Admission Control,” *IEEE Transactions on Networking*, vol. 7, no. 3, pp. 293–309, 1999.
- [12] L. Breslau, S. Jamin, and S. Shenker, “Comments on the Performance of Measurement-Based Admission Control Algorithms,” in *Proc. of IEEE Infocom*, Mar 2000.
- [13] M. Mandjes and M. van Uitert, “Transient Analysis of Traffic Generated by Bursty Sources, and its Application to Measurement-Based Admission Control,” *Telecommunication Systems*, vol. 15, no. 3-4, pp. 295–321, 2000.
- [14] J. Qiu and E. Knightly, “Measurement-Based Admission Control with Aggregate Traffic Envelopes,” *IEEE Transactions on Networking*, vol. 9, no. 2, pp. 199–210, 2001.

- [15] S. Shenker, C. Partridge, and R. Guerin, “RFC 2212: Specification of Guaranteed Quality of Service,” Sep. 1997.
- [16] J. Roberts, U. Mocci, and J. Virtamo, *Broadband Network Teletraffic - Final Report of Action COST 242*. Berlin, Heidelberg: Springer, 1996.
- [17] J. Wroclawski, “RFC 2211: Specification of the Controlled-Load Network Element Service,” Sep. 1997.
- [18] C. Cetinkaya and E. Knightly, “Egress Admission Control,” in *Proc. of IEEE Infocom 2000*, pp. 1471–1480, 2000.
- [19] V. Elek, G. Karlsson, and R. Rönngren, “Admission Control Based on End-to-End Measurements,” in *Proc. of IEEE Infocom 2000*, pp. 1233–1242, Mar 2000.
- [20] I. Más and G. Karlsson, “PBAC: Probe-Based Admission Control,” in *2nd International Workshop on Quality of future Internet Services (QofIS 2001)*, Sep 2001.
- [21] S. Georgoulas, P. Trimintzios, and G. Pavlou, “Joint Measurement- and Traffic Descriptor-Based Admission Control at Real-Time Traffic Aggregation Points,” in *Proc. of IEEE Int. Conference on Communications (ICC 2004), QoS and Performance Symposium*, (Paris, France), June 2004.
- [22] J. Milbrandt, M. Menth, and S. Oechsner, “EBAC - A Simple Admission Control Mechanism,” in *Proc. of 12th IEEE International Conference on Network Protocols (ICNP 2004)*, (Berlin, Germany), Oct 2004.
- [23] M. Menth, J. Milbrandt, and S. Oechsner, “Experience Based Admission Control (EBAC),” in *Proc. of 9th IEEE Symposium on Computers and Communications (ISCC 2004)*, (Alexandria, Egypt), Jun 2004.
- [24] A. M. Law and W. D. Kelton, *Simulation Modeling and Analysis*. Boston, USA: McGraw-Hill, 2000.
- [25] V. Paxson and S. Floyd, “Wide-Area Traffic: The Failure of Poisson Modeling,” *IEEE/ACM Transactions on Networking*, vol. 3, no. 3, pp. 226–244, June 1995.