Virtual Dropping for Endpoint Admission Control

Torsten Braun, Matthias Scheidegger, Marco Studer
University of Bern, Institute of Computer Science and Applied Mathematics
Neubrückstrasse 10, 3012 Bern, Switzerland
{*braun*|*mscheid*|*studer*}@*iam.unibe.ch*

Abstract

Endpoint Admission Control is a scalable QoS mechanism that relies on measuring the amount of lost or marked packets of a probing stream before allowing user data flows to enter the network. The existing approaches can be classified based on whether or not they use a separate traffic class for probing traffic, and whether or not they rely on a marking mechanism. Approaches that use a separate traffic class and rely on a marking mechanism are called out-of-band marking approaches and have proven to be most reliable but are also hard to deploy. As an alternative for out-of-band marking the approach of virtual dropping has been proposed, which discards packets that would have been marked otherwise, based on a virtual queue algorithm, thus removing the requirement of a common marking scheme. This paper compares the behavior of virtual dropping and out-of-band marking. Both mechanisms do not behave exactly the same, but differ slightly.

Keywords

Communication Systems, Quality of Service, Admission Control, Probing

1. Introduction

Differentiated services [1] provide scalable mechanisms for Quality-of-Service (QoS) support in the Internet, as various results have shown [2–4]. However, the issue of admission control has always been an open one. A well-known approach is the use of so-called bandwidth brokers that control the reservations inside network areas assigned to them (usually ISP networks). Bandwidth brokers communicate with each other to set up reservations if an end-to-end path passes through several administrative areas. However, the scalability of centralized approaches such as bandwidth brokers for admission control and management of routers is a problem. Fully distributed mechanisms—such as endpoint admission control—do not have this limitation and have therefore been in the focus of recent work [8–16].

In endpoint admission control (EAC) endpoints always send a stream of probing packets to the peer before starting user data transmission. If the probing procedure indicates sufficient QoS on the path to the peer the endpoint proceeds with data transmission. Otherwise, it must back off. If all the endpoints (such as end systems or edge routers) in the network follow this procedure, high quality of service can be achieved. No strict quality of service guarantees are made, however. Instead, the user or the application can assume an upper bound for packet loss probability. The resulting system supports a soft real-time service similar to Controlled Load [5].

Among the EAC approaches, those belonging to the category of "out-of-band marking" achieve the lowest packet loss probabilities. However, they require a common marking scheme in the network, which may be a hurdle for deployment. Virtual dropping, as proposed in [8], is an alternative to out-of-band marking that is expected to perform nearly identical but does not rely on any marking mechanism. We have investigated whether or not virtual dropping really achieves the expected results, using simulation scenarios similar to [8] in order to obtain comparable results.

2. Related Work

2.1 Endpoint Admission Control

Most of the several existing variants of endpoint admission control can be coarsely classified by two criteria: First, probing traffic can be either transmitted with the same priority or service class as regular data traffic (this is called in-band probing), or it can be transmitted with a different one (out-of-band probing). Second, the admission control procedure can either simply rely on the probing stream's loss ratio, or it can require a marking mechanism in the network (e.g. based on the explicit congestion notification (ECN) bit in the IP header [7]). These criteria result in the four categories in-band dropping, out-of-band dropping, in-band marking, and out-of-band marking [8]. In all these variants the user data transmission is only admitted if the fraction of lost/marked packets stays below a certain threshold.

Of course this categorization does not fully describe all possible approaches. Elek, Karlsson et al. [10] for example propose an out-of-band dropping mechansim that exceeds the QoS performance of the basic approach by using forward error correction to compensate the remaining small probability of packet loss. Their architecture also distinguishes three traffic classes (instead of two), two for controlled traffic and probing traffic, and one for best effort to support legacy applications. Kelly, Key et al. [12] present an in-band marking approach where marking is performed using the same virtual queue mechanism used in this paper. The user is also charged a small amount of money for every forwarded marked packet in order to create additional incentive to adhere to the admission control procedure. Bianchi, Capone et al. [11] use a rather special technique to restrict the impact of probing packets on the admitted flows, called Probing Packet Lifetime (PLT). The PLT limits the time a probing packet may wait in a queue. If the PLT expires the router drops the packet. This allows for much finer analysis of network load than using drop-from-tail as it is effectively a way to restrict the maximum queue length along the path, but it also requires non-standard queuing disciplines in the routers, which may be a hurdle for deployment.

Furthermore, there are also measurement-based admission control architectures that differ strongly from the above concepts. One such approach is the one by Cetinakaya and Knightly [17]. Here, the measuring node is the egress router on the re-

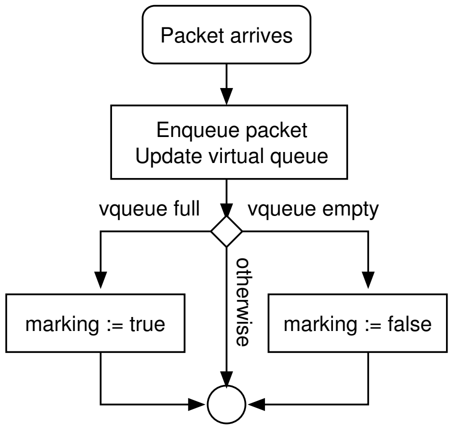

Figure 1: Algorithm for packet arrivals with a virtual queue

mote end of the connection, which allows for measurement-based admission control without any active probing, based only on flow characteristic descriptions provided by the source nodes. This requires that nodes must be able to characterize their flows in advance, however.

For marking, it has been proposed to use virtual queues with a size much smaller than the corresponding real queue [12]. Virtual queues serve as indicators for imminent congestion in the network. They are implemented as byte counters simulating a queue with significantly lower buffer capacity and bandwidth than the real queue. Packets put to the real queue are also put to the virtual queue. When the virtual queue overflows, packets subsequently arriving at the real queue get marked until the virtual queue is completely emptied again. The diagram in Figure 1 illustrates this. The `marking` variable used in the diagram determines whether or not to mark packets at a given point in time. Marking could also be done with a high-water-mark algorithm that marks packets if the queue length exceeds a certain threshold at the time the packet leaves the queue. However, while the high-water-mark algorithm is cheaper in terms of memory and computation it tends to react rather slowly to imminent congestion. Virtual queues have shown to do better in this respect and are thus good candidates for further research.

2.2 Virtual Dropping

EAC based on marking has proven to provide more accurate admission control decisions than EAC based on dropping. This is due to the active nature of congestion marks: If an endpoint receives a single marked packet it immediately knows that the network is close to a congestion state. Without marking, congestion can only be detected with a certain delay and with reduced accuracy. However, the disadvantage of marking is that it may be an obstacle for deployment—marking requires that intermediate routers and probing endpoints support a common marking scheme. This

Table 1 Traffic sources

Source	EXP1	EXP2	POO1
Burst Rate	256 kbps	1024 kbps	256 kbps
On Time	500 ms	125 ms	500 ms
Off Time	500 ms	875 ms	500 ms
Avg. Rate	128 kbps	128 kbps	128 kbps
Beta	—	—	1.2

is still relevant today since many small problems such as compatibility issues with firewalls and missing support in legacy operating systems on many edge nodes, and even servers, persist. [6] contains a list of some related articles and other references related to ECN problems.

As an alternative to marking, Breslau et al. [8] proposed an approach called *virtual dropping*, which was predicted to perform nearly identical to out-of-band marking—the category that provides the best QoS—without the need for any marking mechanism. Similarly to many marking mechanisms it uses the virtual queue approach to detect imminent congestion. However, probing packets that would have been marked with out-of-band marking are now dropped. Regular data packets are not affected by this rule. This approach improves deployability since the network nodes do not need to implement a common marking scheme, although it still requires the routers to implement virtual queue functionality. Since different implementations of virtual queues on varying router platforms would not significantly affect the system this is a comparatively small constraint. Furthermore, while the accuracy of admission control decisions suffers in the presence of legacy routers without support for virtual queues, the functioning of the system as a whole will not be disturbed by them.

In [8] the authors claim that "one could easily achieve exactly the same results doing [...] virtual dropping instead of out-of-band marking," without evaluating this claim. This paper investigates whether or not virtual dropping really achieves the expected results.

3. Performance Evaluation of Virtual Dropping

3.1 Basic Simulation Parameters

For our simulations we used the same topologies and traffic parameters as in [8]. For reference, Tables 1 and 2 show the traffic sources used in this paper. In Section 3.2 we used a single-link topology, and in Section 3.3 we used a topology with multiple links. Virtual queue capacities were set to 90% of the corresponding real queues. We performed the evaluation with an own simulation framework [18] based on the

Table 2 Simulation scenarios

Fig.	Source	Gamma	Description
3(a)	EXP1	3.5	Basic scenario
3(b)	EXP1	1.0	Higher load
3(c)	EXP2	3.5	Bursty traffic
3(d)	Star Wars Trace	—	Video trace data
3(e)	POO1	3.5	Long-tailed on/off times

network simulator ns-2 that can support various EAC mechanisms, and we verified the results described in [8]. For comparison, Figure 2 shows the results of the basic experiment corresponding to [8]. As expected they are very similar with only a slight difference between the graphs for out-of-band dropping. Our simulation implementation of out-of-band dropping shows slightly greater sensitivity to different values of ε. In any case, this difference does not have any influence on the comparisons between out-of-band marking and virtual dropping presented here.

Admission-controlled traffic is modeled by a Poisson arrival process with average inter-arrival time $\gamma = 3.5\,\text{s}$. Flows are admitted if the measured loss rate is below a threshold ε. Threshold values of 1, 2, 3, 4, and 5% have been used.

Probing is performed by a slow start procedure and report interval times of 1 s. The slow start procedure is an optimized measurement procedure, which, unlike the measurement procedure normally used in endpoint admission control, does not send measurement traffic at a fixed rate. Instead, the slow start procedure starts with a low rate, which is doubled with each received report packet that indicates a loss ratio below the threshold ε. If the indicated loss ratio is greater than ε it immediately gives up the connection attempt. Otherwise, if the desired data rate has been successfully probed, the connection is admitted. Note that one report packet will usually contain the acknowledgments for several probe packets. This procedure reduces the risk of the system to enter a thrashing state. Thrashing is a state where a very high rate of flow arrivals effectively prevents the admission of any new flows since the large amount of probing traffic in the network makes it impossible to achieve packet loss ratios lower than the threshold when probing.

3.2 Simulation Results with Single Link Scenario

First, we used a single link (10 Mbps) scenario with many sources sharing a single congested link. Figure 3(a) shows the simulation results of virtual dropping and out-of-band marking in terms of packet loss probability of already admitted flows and the resulting utilization (i.e. the fraction of available bandwidth occupied by the admitted flows). The five points on each curve represent the five values of ε (1–5%) used, from the lower left to the top right. Utilization describes the ratio of admission-controlled

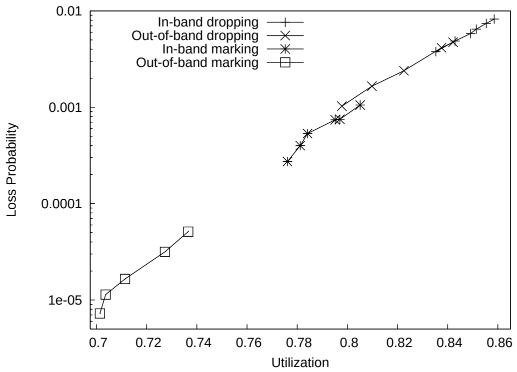

Figure 2: The four basic EAC approaches compared (basic scenario)

traffic and allocated resources. All flows have an exponential lifetime with an average of 300 s. The traffic sources used have exponential on/off times of 500 ms, burst rates of 256 kbps, and packet sizes of 125 bytes.

Both mechanisms, virtual dropping and out-of-band marking, behave rather similarly. They achieve loss probabilities of approximately 10^{-5} and utilization values between 70 and 74%. For comparison, classical in-band and out-of-band dropping achieve loss probabilities between 10^{-2} and 10^{-3}, at 80% to 86% utilization. While out-of-band marking achieves lower loss probability, virtual dropping has higher utilization.

The slight difference can be explained by two reasons: First, when a packet gets marked (out-of-band marking) it will still be forwarded. This increases the load on subsequent links, and thus the probability of congestion on these links, while a dropped packet has no impact on any subsequent links. Additionally, marked packets remain in the queue and thus influence the dropping probability of following packets. The number of dropped probes in virtual dropping is thus somewhat smaller than the respective number of marked probes in out-of-band marking.

The second reason for the difference between the two approaches is due to the fact that packet loss is detected by counting the gaps in the sequence numbers of the received probing flow. Thus, if the last n packets of a probing flow are all dropped, they will not be counted as being lost. With out-of-band marking these packets would be marked and would more probably reach their destination. Because of this, out-of-band marking again rejects a small number of flows that virtual dropping would accept. The obvious solution is to make the receiver know the maximum sequence number of the probing flow. With the standard measurement procedure this is easy to implement. However, with slow start probing, the maximum sequence number depends on the performance observed while probing.

Figures 3(a)–3(e) describe the behavior of both mechanisms under various load conditions. Figure 3(b) shows the behavior under high load with many arriving flows. The average inter-arrival time is γ=1 s. Since the other parameters remain the same,

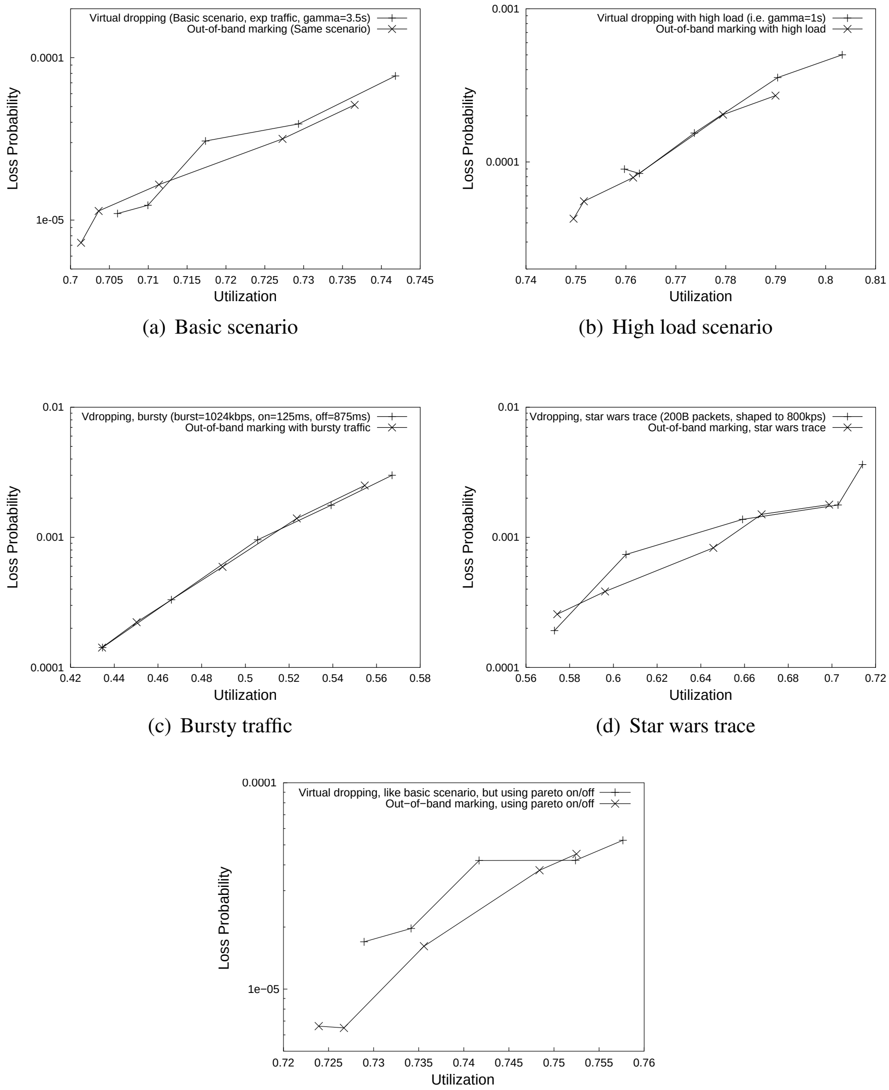

(a) Basic scenario

(b) High load scenario

(c) Bursty traffic

(d) Star wars trace

(e) Pareto

Figure 3: Virtual dropping and out-of-band marking

the offered load is 400% of the available bandwidth on average. In this case, the difference between virtual dropping and out-of-band marking is much smaller than in the original scenario shown in Figure 3(a). This can be explained by the fact that with high load most out-of-band marking packets are lost. The outcome of the out-of-band marking admission control procedure is therefore mainly affected by lost packets

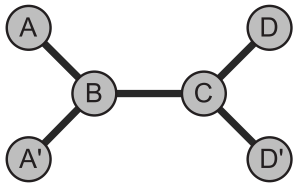

Figure 4: Multi-Link Scenario

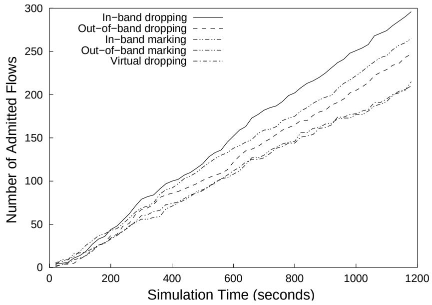

Figure 5: Number of accepted flows for each EAC mechanism

instead of marked ones, and it therefore resembles the outcome of virtual dropping. A similar behavior can be observed with the very bursty source shown in Figure 3(c), which uses exponential on/off traffic sources with burst rates of 1024 kbps, on times of 125 ms, and off times of 875 ms.

Figure 3(d) shows the results with star wars video sources and packet sizes of 200 bytes. The traffic source has been shaped to a token bucket size of 200 kbytes and a rate of 800 kbps. Figure 3(e) shows Pareto on/off traffic sources with 256 kbps burst rate and on/off times of 500 ms. Figure 3(c) shows smaller differences between virtual dropping and out-of-band marking than the other sources. The explanation for these small differences is the same as for the high load case.

3.3 Simulation Results with Multi-Link Scenario

Finally, we performed simulation measurements with a multi-link scenario as depicted in Figure 4. All links have a bandwidth of 10 Mbps, and 20 ms delay. The threshold value ε is 5%. We compare the results of virtual dropping not only with out-of-band marking but also with the other three combinations for endpoint admission control, in-band marking, in-band dropping, and out-of-band dropping. Flows exist between A and D as well as between A' and D'. The bottleneck between B and C causes more than 50% of the flows to be rejected (blocked).

Figure 5 shows the number of accepted flows dependent on time. Compared to in-band marking, in-band and out-of-band dropping, the virtual dropping and out-of-band marking mechanisms accept a rather low number of flows and seem to behave nearly identical. The blocking and the loss probabilities are given in Table 3. Again, there are slight differences between out-of-band marking and virtual dropping.

4. Conclusions

We have evaluated and compared the virtual dropping and out-of-band marking approaches for endpoint admission control. The out-of band marking approach provides very good QoS, but it also requires a common marking scheme (like ECN) in the network, which may be an obstacle for deployment. Our results show that, apart

Table 3 Multi-Link Scenario

	Blocking Probability	**Loss Probability**
In-band dropping	0.54754	0.031000
Out-of-band dropping	0.60280	0.006090
In-band marking	0.59869	0.005110
Out-of-band marking	0.65222	0.000041
Virtual dropping	0.65334	0.000079

from a few minor differences, both approaches behave very similarly. Therefore, virtual dropping is a good replacement for out-of-band marking in places where the deployment of a common marking scheme is difficult.

References

[1] B. Carpenter and K. Nichols, “Differentiated Services in the Internet,” in *Proceedings of the IEEE*, vol. 90, pp. 1479–1494, Sep. 2002.

[2] Günther Stattenberger, Matthias Scheidegger, Torsten Braun, Marcus Brunner and Heinrich Stüttgen, “Performance Evaluation of a Linux DiffServ Implementation,” in *Computer Communications Journal*, Elsevier, August 2002.

[3] T. Ferrari, G. Pau, C. Raffaelli, “Priority Queuing Applied to Expedited Forwarding: A Measurement-based Analysis,” in *Proceedings of Quality of Future Internet Service*, QofIS 2000, pp. 167–181.

[4] F. Garcia, C. Chassot, A. Lozes, M. Diaz, P. Anelli, E. Lochin, “Conception, Implementation and Evaluation of a QoS-based Architecture for an IP Environment Supporting Differential Services,” *IDMS’2001 Interactive Distributed Multimedia Systems*, Lancaster, UK, September 4–7, 2001.

[5] J. Wroclawski, “Specification of the Controlled-Load Network Element Service,” Internet RFC 2211, 1997.

[6] ICSI Center for Internet Research, ECN Problems webpage, http://www.icir.org/floyd/ecnProblems.html.

[7] K. Ramakrishnan and S. Floyd, “A Proposal to add Explicit Congestion Notification (ECN) to IP,” *Request for Comments 2481*, Jan. 1999.

[8] L. Breslau, E. Knightly, S. Shenker, I. Stoica and H. Zhang, “Endpoint Admissin Control: Architectural Issues and Performance,” in *Proceedings of ACM SIGCOMM 2000*, pp. 57–69, Aug. 2000, Stockholm, Sweden.

[9] R. J. Gibbens and F. P. Kelly, “Distributed Connection Acceptance Control for a Connectionless Network,” *Teletraffic Congress*, Edinburgh, 1999.

[10] V. Elek, G. Karlsson and R. Rönngren, ‘Admission Control on End-to-End Measurements,” in *Proceedings of IEEE Infocom*, 2000.

[11] G. Bianchi, A. Capone and C. Petrioli, "Throughput Analysis of End-to-End Measurement-Based Admission Control in IP," in *Proceedings of IEEE Infocom*, 2000.

[12] F. P. Kelly, P. B. Key and S. Zachary, "Distributed Admission Control," in *IEEE Journal on Selected Areas in Communications*, Vol. 18, 2000.

[13] M. Reisslein, "Measurement-Based Admission Control for Bufferless Multiplexers," in *International Journal of Communication Systems*, Vol. 14, October 2001.

[14] G. Bianchi, F. Borgonovo and A. Capone, "Endpoint Admission Control with Delay Variation Measurements for QoS in IP Networks," in *ACM Sigcomm Computer Communications Review* Vol. 32, April 2002.

[15] R. Hill and H. Kung, "A DiffServ Enhanced Admission Control Scheme," in *Proceedings of IEEE Globecom 2001*, November 2001.

[16] Tom Kelly, "An ECN Probe-Based Connection Acceptance Control," in *Computer Communication Review*, Vol. 13, No. 3, July 2001.

[17] C. Cetinkaya, V. Kanodia and E. Knightly, "Scalable Services via Egress Admission Control," in *IEEE Transactions on Multimedia*, March 2001.

[18] M. Studer, *A Simulation Framework for Endpoint Admission Control*, Diploma Thesis (in German), Nov. 2002, University of Bern, Switzerland.