

Delivering Excess Traffic with Delay Guarantees - An Extension to Expedited Forwarding

(extended abstract)

Johan Karlsson* Ulf Bodin* Andrej Brodnik* Andreas Nilsson*
Olov Schelén *

Abstract

This paper presents a new set of forwarding behaviors that fits rate-adaptive and delay-sensitive applications with limited loss tolerance. We consider an application to have limited loss tolerance if it needs loss-free forwarding of specific packets up to a certain rate. The new set of forwarding behaviors can be attractive for developing real-time applications for the Internet. In particular, such applications can be designed to use reserved forwarding capacity efficiently and compete for more forwarding capacity while being fair to best-effort traffic. To provide the new set of forwarding behaviors, we define a scheduling mechanism that can be implemented efficiently. Through simulations, we show that this mechanism supports the requirements of the defined forwarding behaviors.

1 Introduction

Real-time applications are becoming increasingly common in the Internet, e.g. video and voice over IP. Such applications need to present data to users with short delay (i.e., they are delay-sensitive). Real-time applications may also prefer low loss-rate since error resilience then can be sacrificed for better compression efficiency and higher quality. Applications optimized for low loss-rate are said to be intolerant in contrast to applications that are tolerant to packet loss.

Clearly, delay-sensitive and intolerant applications gain from guarantees on bounded delay and

loss. Such guarantees can be provided with the Expedited Forwarding (EF) per-hop behavior (PHB) [1]. EF is part of the Differentiated Services (DiffServ) architecture [2].

EF requires that traffic using the PHB is peak-rate limited. This can be achieved by associating users with service profiles. Then, traffic can be policed to fit these profiles as it enters DiffServ capable networks. EF traffic conforming to the peak-rate of each profile (i.e., the *conforming* traffic) is given loss-free and low delay forwarding, while EF traffic exceeding this peak-rate is dropped (i.e., the *excess* traffic).

EF fits delay-sensitive and intolerant applications that need a certain bit-rate, but do not gain from more forwarding capacity if available. Delay-sensitive and tolerant applications that need a minimal bit-rate and can gain from additional forwarding capacity are however not well supported by EF. Such applications adapt their sending rate as response to packet loss (i.e., they are rate-adaptive¹). Still, they may need loss-free forwarding of specific packets up to a certain rate (e.g., to maintain a minimal frame rate²). This kind of rate-adaptive and delay-sensitive applications can be said to have

¹Rate-adaptive video applications can, for example, be constructed by varying the frame rate or resolution, or by varying the quantization step size.

²In [3], an error resilient (to packet loss) and scalable compression method for video is presented. The compression method varies the sending rate by using different quantization steps and, when these steps cannot be reduced further, by sub-sampling the input sequence to reduce the frame rate. With loss-free forwarding up to a certain bit-rate, decreasing the frame rate can be avoided for this method. In addition, it might be possible to scarify error resilience for better compression efficiency and higher quality for the protected sub-stream.

*Department of Computer Science and Electrical Engineering, Luleå University of Technology, Luleå, Sweden. E-mail: k,uffe,brodnik,andreas,olov@sm.luth.se

limited loss tolerance.

In this paper, we define a new PHB group that fits rate-adaptive and delay-sensitive applications with limited loss tolerance. We name the new PHB group In-Time (IT). The IT PHB group consists of three PHBs; the IT-conforming PHB, the IT-excess PHB, and the IT-background PHB. It provides delay limited and in-order forwarding of conforming and excess packets. In addition, conforming traffic is guaranteed loss-free forwarding³. The PHB group aims at treating excess traffic and background traffic equally with regard to loss (i.e., excess traffic is given a loss-rate equal to or higher than the loss-rate given to background traffic). Hence, excess traffic and background traffic can compete on similar terms for available forwarding capacity.

Most real-time applications require data to be ordered before processing. Without ordered forwarding of conforming and excess packets such applications would need to buffer arriving packets to place data in order. Since this introduces delay, we consider the property of ordered forwarding in the IT PHB group important.

As the traditional best-effort forwarding service, the IT-background PHB provides no guarantees. It can, but does not need to, be equal to the default PHB [4], which is the best-effort service through a Differentiated Services compliant node. In this paper, we use the term *best-effort* for the IT-background PHB.

The forwarding behaviors provided by IT cannot be supported by any combination of the PHBs currently specified by the IETF. IT contributes thus with new differentiation properties. We believe these properties to be attractive in developing real-time applications for the Internet. With IT, a real-time application can be designed to both use reserved forwarding capacity efficiently (e.g., sacrifice error resilience for better compression efficiency and higher quality) and compete for more bandwidth while being fair to TCP⁴. Without IT, real-time applications need to use reserved forwarding capacity only (provided by EF) or operate without any upper bound on delay.

EF is straightforward to implement in routers.

³Conforming traffic is, as with EF, peak-rate limited to the bit-rate allocated for IT in the network.

⁴TCP fairness for real-time applications can be achieved through TCP Friendly Rate Control (TFRC) [5].

A prioritized queue in addition to the best-effort queue is sufficient to support the EF PHB. IT, on the other hand, requires more sophisticated packet scheduling. Such scheduling mechanisms tend to consume considerably more clock cycles and memory than simpler scheduling mechanisms such as prioritized queuing. A carefully designed scheduling mechanism supporting IT can however be implemented efficiently. We present a scheduling mechanism for IT that consumes reasonable amounts of memory and have moderate processing overhead on common platforms (i.e., platforms with instruction sets like the Intel Pentium III that do operations such as bit-shifts, additions and subtractions efficiently, but that do operations such as multiplications, divisions and floating-point arithmetic less optimally).

The scheduling mechanism implementing IT has one configurable parameter only, i.e. the maximum allowed delay, d , for conforming and excess traffic. Conforming packets are delayed as much as possible to allow excess traffic and best-effort traffic to be forwarded. The loss-rate experienced by excess traffic decreases with this delay limit. We show, through simulations with the network simulator version 2 (NS-2) [6], that IT can be provided with our scheduling mechanism. With d set as low as 2 ms, excess traffic can be given a loss-rate close to the loss-rate experienced by best-effort traffic. Moreover, for low d , the delay is considerably lower for excess and conforming traffic than for best-effort traffic.

2 Summary

This paper presents a new per-hop behavior (PHB) group, In-Time (IT). IT fits rate-adaptive and delay-sensitive applications with limited loss tolerance. We consider an application to have limited loss tolerance if it needs loss-free forwarding of specific packets up to a certain rate.

IT cannot be supported by any combination of the PHBs currently specified by the IETF and contributes thus with new differentiation properties. We believe these properties to be attractive in developing real-time applications for the Internet. With IT, a real-time application can be designed to both use reserved forwarding capacity efficiently and compete for more bandwidth while being fair

to best-effort traffic. Without IT, real-time applications need to use reserved forwarding capacity only (provided by EF) or operate without any upper bound on delay.

To provide IT, we define a scheduling mechanism that can be implemented efficiently, the TICKET scheduler (see “Extended Expedited Forwarding: the In-Time PHB group” [7] for details). The implementation consumes reasonable amounts of memory and have moderate processing overhead on common platforms.

We evaluate the TICKET scheduler through simulations (not included for space reasons). These simulations show that the IT PHB group can be implemented with the TICKET scheduler. Excess traffic gets a useful amount of forwarding capacity even at low maximum allowed delays (i.e., 2 ms). The maximum delay experienced by conforming and excess traffic is always below the maximum allowed delay and considerable lower than the maximum delay of best-effort traffic for low maximum allowed delays.

The loss-rates are equal or higher for excess traffic compared to best-effort traffic (as required for the IT PHB group). Still, the loss-rate for excess traffic is close to the loss-rate of best-effort traffic. This similarity in loss-rates between excess and best-effort traffic suggests that excess traffic can be useful to rate-adaptive and delay-sensitive applications with limited loss tolerance.

References

- [1] Van Jacobson, Kathleen Nichols, and Kedar-nath Poduri, “An expedited forwarding phb,” IETF RFC 2598, June 1999.
- [2] S. Blake, D. Black, M. Carlsson, E. Davies, Z. Wang, and W. Weiss, “An architecture for differentiated services,” IETF RFC 2475, Dec. 1998.
- [3] Wai tian Tan and Avideh Zakhor, “Real-time internet video using error resilient scalable compression and TCP-friendly transport protocol,” *IEEE Transactions on Multimedia*, vol. 1, no. 2, pp. 172–186, 1999.
- [4] K. Nichols, S. Blake, F. Baker, and D. Black, “Definition of the differentiated service field (ds field) in the IPv4 and IPv6 headers,” IETF RFC 2474, Dec. 1998.
- [5] Sally Floyd, Mark Handley, Jitendra Padhye, and Jörg Widmer, “Equation-based congestion control for unicast applications,” in *Proceedings of ACM SIGCOMM 2000 Conference*, Stockholm, Sweden, Oct. 2000, vol. 30 (4) of *Computer Communications Review*, pp. 43–56, ACM SIGCOMM.
- [6] “NS (Network Simulator),” <http://www.isi.edu/nsnam/ns/>.
- [7] Johan Karlsson, Ulf A. Bodin, Andrej Brodnik, and Andreas Nilsson, “Extended expedited forwarding: In-time PHB group,” manuscript submitted to INFOCOM 2002, k@sm.luth.se, 2001.