An Advanced QoS Protocol for Mass Content

Thomas Dreibholz, Avril IJsselmuiden
University of Duisburg-Essen
Ellernstrasse 29
45326 Essen, Germany
{dreibh,avril}@exp-math.uni-essen.de

John L. Adams
British Telecom
Martlesham Heath
Suffolk, United Kingdom
john.l.adams@bt.com

Abstract

This paper presents a novel network device being located in network edge nodes. It provides a solution for QoS guarantees to certain flows on a congested link by focussing packet discard on selected flows. Unlike IntServ solutions like RSVP, our approach only requires minimal signalling and provides both efficiency and scalability. In this paper, we first describe the ideas of our QoS device and then provide first results from a fast-track simulation model implementing a lightweight version of our approach.

1. Introduction

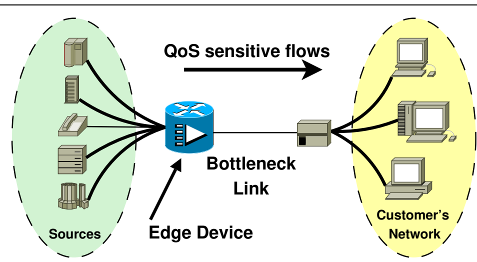

Figure 1. Broadband Scenario

With DSL technology becoming widespread, a rising amount of customers obtains access to high-speed Internet backbones. Such links do not only speed up existing applications but also make services like video and audio on demand possible. But unlike best-effort applications, such services have stricter QoS requirements; especially they need an assured bandwidth. We assume that core bandwidth is usually over-provided and only the link to the customer becomes the bottleneck (see figure 1). When multiple equal-priority flows exceed the DSL link's bandwidth, the quality of all flows suffers due to packet loss. Using RSVP to establish per-flow reservations would solve the problem but introduces complex signalling and a lack of scalability [1]. In this paper, we present a novel, simple and scalable approach to ensure QoS guarantees while only requiring a minimum effort on signalling. Our approach [3, 9] is based on a QoS device located inside an edge node before the bottleneck link. It is currently under consideration by the Standards Bodies ITU-T and ETSI [5–7].

2. Our QoS Device

The key idea of our QoS device is that in case of congestion is it better to focus packet discard on selected flows than to discard arbitrary packets. Then, only the selected flows would suffer from quality loss instead of all flows. This principle, applied at the ATM cell level, was first suggested in [10] and its value has been shown in other publications since [4, 8]. We apply it to our device by making the latest flow(s) the subject of discard (the default policy; arbitrary other schemes may be applied as well.). Therefore, our device has to know when a new flow starts, and has to maintain a record of it.

For the device to recognise the start of a flow, its sender is only required to send a *start packet*; no further signalling is necessary. In particular, the sender is not required to wait for any acknowledgement, it may start sending data immediately. The start packet contains flow identity information (i.e. packet header fields; see [2] for more information) to identify the data flow and an estimation of the flow's bandwidth. This information is recorded by our device.

The device maintains a window of flows that are vulnerable to packet discard in case of congestion, the *drop window*. As new flows start up, the flow moves through the drop window, until eventually it is removed from the window (by being overwritten by new entries), when one of the following conditions is satisfied: (1) The sum of the rates of the flows in the drop window, minus rate of the oldest flow, is greater than R, where R is a percentage of the link bandwidth. (2) The flow has been in the window for at least time t_{min}. (3) It has received at least p_{min} packets after the start packet.

When a flow identity is removed from the drop window, it becomes a guaranteed flow, except under extreme traffic conditions. Note, that the list of guaranteed flow identities is not stored. All packets not belonging to flows of the drop window are implicitly assumed to be guaranteed flows. This ensures the scalability of our device.

3. Experimental Testbed

For a proof of concept, we have created a fast-track simulation model using an Open Source simulation package under BSD licence written in Common LISP [11]. This model is being able to run with our QoS protocol turned on or off; if it is turned off, packets arriving are discarded indiscriminately in the case of congestion.

The example scenario has been set up with 5 traffic generators attached: flow 1 (multimedia, 256 Kbit/s, 30 fps, 1060 bytes/frame), flow 2 (voice, 64 Kbit/s, 80 fps, 100 bytes/frame), flow 3 (voice, 64 Kbit/s, 8 fps, 1000

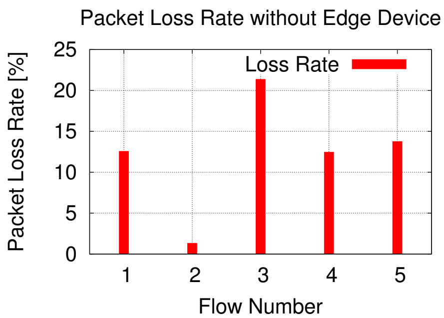

Figure 2. Packet loss, protocol turned off

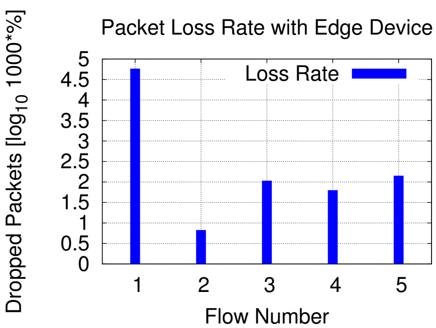

Figure 3. Packet loss, protocol turned on (%*1000 converted to log scale)

bytes/frame), flow 4 (multimedia, 256 Kbit/s, 30 fps, 1060 bytes/frame), flow 5 (multimedia, 512 Kbit/s, 50 fps, 1280 bytes/frame). Frame inter-arrival times deviate by +/- 25% (random, uniform distribution). The packet MTU is 1000 bytes, larger frames are segmented. The link bandwidth is 1 Mbit/s; the edge device's output buffer is 100000 bit (i.e. 100ms at 1 Mbit/s speed), the drop window size R is 250 Kbit/s.

4. Results

A number of experiments were performed, using confidence intervals of 95%. Our first experiment was performed with the protocol turned off, in order to verify that packet loss would be indiscriminate, and spread across all 5 flows.

Figure 2 shows the percentage loss rate of the total packets transmitted, for each flow; it is indiscriminate in that all 5 flows have lost packets. While the loss rates have values between 13% and 22% for flows 1 and 3-5, the rate for flow 2 is lower due to its small packet sizes (100 bytes instead of 1000 to 1280).

The next step was to run simulations with the protocol turned on. Flow 1 is the reference flow (in the drop window) and flows 2-5 were in the guaranteed area. Figure 3 shows the resulting percentage packet loss rate using a log_{10} scale.

Obviously the loss rate for the reference flow is highest: about 58%. For the other flows, the loss rate is about 0.1% or less. This clearly shows that even a very simple edge device protocol achieves a quite low packet loss rate (about 0.1%) for 4 of 5 flows. Clearly, in terms of customer satisfaction, four of five users notice no significant reduction in service quality. This is far better than all flows being unsatisfactory!

However, it should be noted that flows 2-5 do not receive exactly the same QoS as in a congestion-free network: the packet transmission is non - pre-emptive, i.e. when currently a drop window packet is being transmitted, it is not possible to suspend its transmission. Instead, new guaranteed area packets queue up and may fill up the queue, possibly causing a loss for guaranteed flows.

5. Conclusions

In this paper, an advanced QoS protocol for mass content has been described. The protocol is suitable for controlling congestion in network edge devices. A lightweight version of the protocol has been implemented in a LISP simulation model, and this was described, along with some results. The initial results show that even a lightweight protocol is better than using no protocol at all. The research now splits into 3 tracks: this model is further used to investigate security mechanisms; an Opnet model is used to implement the full device functionality and provide simulation results to the Standards Bodies, ITU-T and ETSI; and a third model is currently being developed as a student project, in OMNeT++.

References

[1] T. Dreibholz. Management of Layered Variable Bitrate Multimedia Streams over DiffServ with Apriori Knowledge. Masters thesis, University of Bonn, Institute for Computer Science, Feb 2001.

[2] T. Dreibholz. An IPv4 Flowlabel Option. Internet-Draft Version 03, University of Duisburg-Essen, Apr 2005. draft-dreibholz-ipv4-flowlabel-03.txt, work in progress.

[3] T. Dreibolz, A. Smith, and J. L. Adams. Realizing a scalable edge device to meet QoS requirements for real-time content delivered to IP broadband customers. In *Proceedings of the 10th IEEE International Conference on Telecommunications*, Papeete/French Polynesia, Feb 2003.

[4] S. Floyd and V. Jacobsen. Random early detection gateways for congestion avoidance. *IEEE/ACM Transactions on Networking*, 1(4):397–413, 1993.

[5] A. IJsselmuiden and J. L. Adams. Delivering QoS from remote content providers. British Telecom Contribution TISPAN 01(03)TD132, ETSI, September 2003.

[6] A. IJsselmuiden and J. L. Adams. Delivery of assured QoS content in NGNs. British Telecom Contribution D-,Q4,6,10,11,16,SG13, ITU-T, February 2004.

[7] A. J. IJsselmuiden and J. L. Adams. Proposal for a new IP transfer capability and future QoS studies. British Telecom Contribution D-,Q4,6,10,11,16,SG13, ITU-T, February 2004.

[8] K. Kawahara, K. Kitajima, T. Takine, and Y. Oie. Performance evaluation of selective cell discarding schemes in ATM networks. In *Proceedings of the IEEE Infocom '96*, pages 1954–1061, 1996.

[9] A. J. Smith and J. L. Adams. Packet discard control for broadband services. European Patent Application EP 01 30 5209, British Telecom, June 2001.

[10] A. J. Smith, J. L. Adams, C. J. Adams, and A. G. Tagg. Use of the Cell Los Priority Tagging Meachanism in ATM switches. In *Proceedings of the ICIE 1991*, Singapore, Dec 1991.

[11] M. Tüxen. LISP Simulation Package. http://sctp.fh-muenster.de/sim.