TSV working group
Internet Draft
Document: draft-morita-tsvwg-pps-01.txt
Naotaka MORITA
NTT Corporation
Gunnar KARLSSON
KTH
Expires: April 2004
October 2003

Framework of Priority Promotion Scheme

Status of this Memo

Abstract

The Priority Promotion Scheme (PPS) is a new scheme for traffic control; more specifically, PPS involves applying a kind of admission control to achieve end-to-end QoS for a series of packets on a packet-based network. The main targets are interactive multimedia services such as VoIP, video chat, and video conferencing. The scheme is based on end-to-end measurement of network resources by end systems. Before a session is established or even during a session, the source end system senses, measures, or probes the availability of network resources by sending out packets with priority one level lower than that of normal packets. The result is modification of the DiffServ Code Point (DSCP) value of the succeeding IP packets: the priority is raised or promoted to firmly establish the session, lowered to leave resources with existing sessions, or otherwise adjusted so that the amount of packets does not exceed the available capacity. The network, i.e., output links of the routers or L2 switches is only assumed to support the per-class form of priority

control that accompanies the DiffServ architecture. Having all end systems follow the above behavior achieves end-to-end QoS without the maintenance of per-flow state in each item of network equipment.

This document describes the reasons for the end-to-end measurement-based approach and the general network architecture of PPS.

Conventions used in this document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC-2119 [2].

Table of Contents

1. Introduction

Emerging services such as VoIP, video chat, and video conferencing require session-based QoS. A number of schemes for providing the required QoS control have been put forward, but they either require per-flow management of routers within the network or handle the provision of QoS on a per-class basis, which requires the allocation of large amounts of resources. In this document, a framework for a new QoS scheme is proposed. The scheme is suitable for session-based interactive multimedia and adds less complexity to the network than previous approaches, while delivering per-flow QoS.
Karlsson [3] [4] originally proposed the basic concept. Here, we clarify the requirements for routers, introduce enhancements to session control using SIP, and show some alternative ways to implement the required monitoring of end-system behavior. We refer to this scheme as the "Priority Promotion Scheme".

One of the key functions of the Priority Promotion Scheme is the behavior of routers. We introduce the MF-PHB (Measurable Forwarding Per Hop Behavior) as a new per-hop behavior that provides the required functionality. Whether or not MF-PHB is feasible on given items of existing equipment will have to be verified. This framework is intended as a guide for device manufacturers, network administrators, and operators who need a way to provide QoS for interactive multimedia services. It is not intended, in its current state, for use by the majority of networks in the Internet. We make this proposal now because we feel that the only way to achieve a long-term solution for inter-domain QoS is to start putting intra-domain solutions into practice and then incrementally expand the scope of the work as more experience in deployment is gained.

In this document, we introduce a framework for Priority Promotion. We describe the target service category, which we refer to as "interactive multimedia services", in section 2. In section 3, we explain our motivation in focusing on an end-system-oriented measurement-based approach. The basic procedures of the Priority Promotion Scheme are then explained in section 4. In section 5, specific variant applications of the Priority Promotion Scheme are presented to show the scheme's potential. The feasibility of a measurement-based approach is presented in the appendix to this document and section 6 states why the arguments in the appendix are applicable to the PPS. The functional architecture of the scheme is described in section 7. Finally, the requirements for individual functional entities are summarized in section 8. MF-PHB (Measurable Forwarding) that is necessary to realize PPS is defined in [5] and the verification scenarios of MF-PHB is in [6].

2. The target service type - interactive multimedia

The major targets of the Priority Promotion Scheme are multimedia and interactive communications services provided through software tools running on PCs and operated by human beings. We call such services interactive multimedia (IMM) services. Typical examples of IMM are VoIP, video chat, and video conferencing. Several characteristics differentiate IMM services from existing data services. Web browsing and, in many cases, file retrieval are based on client/server models and the data transfers speeds required are not in general very high. In contrast to this, IMM services are any-to-any and require relatively high speeds in the range from less than 1 Mbps to several Mbps. These IMM-inherent characteristics may cause large fluctuations in traffic patterns and may not be predictable in advance.

Other important characteristics of IMM services are the QoS requirements: that is, the requirements for bandwidth guarantees and short delays. The latter is because of the real-time nature of these services. The former is because typical codecs are sensitive to fluctuations in bandwidth, which lead to degradation of the QoS. While several codecs adjust their information rates to suit the available bandwidth, they impose higher processing loads on the end systems; this approach also necessarily incurs noticeable and possibly annoying fluctuation in the perceived quality. This implies that once a session has been established, the bandwidth has to be guaranteed until the end of the session. In other words, the session should not be established unless the required bandwidth is available. Note that one desirable extended interpretation of this concept is to allow increases, but never decreases, in the bandwidth available to a session. That is, improvement is acceptable but deterioration is not. This is why we have included "promotion" in the name of the scheme.

Finally, a session of an IMM service is set up on-demand and may last for time of the order of minutes to tens of minutes.

When we take the above-described characteristics and requirements of IMM into account, we see that explicit admission control on a per-flow basis is necessary. A common argument is that simple over-provisioning is capable of meeting these requirements. As was stated above, however, IMM combines the characteristics of relatively large bandwidth requirements and strict QoS needs in general with unpredictable traffic patterns. Therefore, we need a form of session-based admission control to deliver QoS for IMM services.

It should be emphasized that admission control has a completely different goal from the existing TCP core functionality. The goal of admission control is to provide bandwidth guarantees with the appropriate QoS for a certain maximum number of sessions. For example, if the network is able to carry 100 Mbps and 100 users request sessions with guarantees of 1 Mbps, nearly 100 sessions should be established. If 1000 users request the same 1-Mbps guarantees, only around 100 sessions should be established. This is quite different from existing data services provided through the TCP. The idea of the TCP is to share network resources in a "fair" manner among the sessions requested at any time. If the network is able to carry 100 Mbps and 100 users request sessions, 100 sessions should be established, each with roughly 1 Mbps throughput. If 1000 users request sessions, all 1000 should be established, each with a throughput around 0.1 Mbps. This is not suitable for IMM services.

The SIP provides one suitable way to control IMM services. Although we focus on the SIP in this description, session-control protocols for the PPS are not restricted in this way.

The application of a QoS policy which includes differentiation based on the identity of the callers or callees in sessions has to be studied as a separate issue. Issues include competition between VIP calls and ordinary calls, or between preferential calls and ordinary calls in times of disaster. If such a policy that caters for such situations is to be applied along with simple admission control based on resource availability, policy credential information from the SIP or another signaling method may have to be incorporate into the PPS framework.

3. Motivation for the focus on an end-system-oriented measurement-based approach

As IP-based networks proliferate, overall network configurations become increasingly complex. In terms of bandwidth available in the access network, DSL alone includes many variants. 12-Mbps ADSL is quite popular in Japan and higher-speed ADSL services will be deployed in the near future, but the actual throughput is completely dependent on conditions such as the distance from the central office and interference among the lines.

Another point is the variations in the network configurations of customers, including broadband routers. The broadband routers initially offered for use with higher-speed access lines may not be capable of providing the same maximum throughput as is stated in the catalogue. A customer's PC may impose similar restrictions. Furthermore, wireless access introduces further complications in terms of the access environment. The network to which the customer is connected adds a lot of variables.

In such a complicated situation, end-to-end guarantees of QoS are difficult to achieve and the role of the end system becomes more important, because only the end system is able to see the actual conditions of communication. In the Priority Promotion Scheme, the end systems measure, monitor, or probe levels of network resources so that they are able, if possible, to set up and maintain media streams with required levels of QoS. We focus on an end-to-end approach because only the end systems are able to judge the overall relevant network situation.

We refer to the terminal points of the media stream, i.e. PCs or residential gateways and routers, as end systems.

4. Basic procedure for the Priority Promotion Scheme

The Priority Promotion Scheme (PPS) is a new scheme for traffic control; specifically, the PPS achieves end-to-end QoS for interactive multimedia services by exercising admission control for series of packets on a packet-based network. The scheme is based on end-to-end measurement of network resources through coordination of the end systems.

In this context, "priority" means priority or precedence at the packet level as represented by the DiffServ Code Point (DSCP) in the IP layer. If we apply the PPS in Layer 2, the priority is represented by the user_priority field specified in 802.1D and Q. If MPLS is used as an underlying transport, EXP field corresponds to the code.

4.1 Basic procedure for end systems

PPS largely relies on end-system behavior for sending the probe packets, which test the availability of network resources, and for decisions on whether or not the succeeding (higher priority) packets can in fact be sent.

Before a session is established and even, under certain conditions, during sessions, the source-end system senses, measures, or probes to detect the availability of network resources. This is done by sending packets with priority one level lower than that of the non-probe packets, i.e. those for established streams. Probe packets are given lower priority so that existing flows of packets are maintained and packet loss is confined to the probe packets; this gives a sharper focus to the loss characteristics.

Criteria for successful receipt at the destination-end system can include loss, delay, and delay jitter. The authors believe that loss will usually be the crucial parameter, but are willing to enlarge the scope of measurement to include the other two characteristics.

The conditions of receipt determine how the DSCP value for the succeeding IP packets is adjusted: the priority is raised or promoted to firmly establish the session, lowered to leave resources with

existing sessions, or otherwise adjusted to control the amount of packets such that the traffic fits into the available capacity.

The RTCP can be used to carry the report from the destination end system. Whether or not the probing packets can carry real media data depends on the required duration of measurement. If measurement will take more than a couple of seconds, the probe packets should carry real media so that the customer does not have to wait for completion of the measurement period.

4.2 Router behavior

The PPS in principle requires that the network, i.e. each output link of a router or Layer 2 switch, support per-class priority control. Prioritization allows the end systems to measure remaining resources without affecting existing streams. In addition to the simple priority control required by the PPS in itself, existing classes (Per-Hop Behaviors or PHBs) such as EF, AF, and BE should be supported. That is, we have to implement an extension to the DiffServ architecture. To clarify the requirements specific to the PPS, we propose Measurable Forwarding as a new PHB (MF-PHB). A detailed description of the MF-PHB has already been given [5]. Whether or not current DiffServ implementations are capable of supporting this new PHB for the PPS without elaboration of the queue configuration is not clear. However, having all end systems behave in the way described above and all network elements implement the MF-PHB ensures that the end-to-end QoS is achieved without having to maintain per-flow states in individual items of network equipment.

A great advantage of the PPS is that it avoids persistent contention among real-time streams. Note that we are talking about scheduling priority in the DiffServ scheduler as opposed to a policy perspective on call control preference or drop preference in a common queue.

4.3 Variation of measurement-based mechanisms

Measurement-based approaches have many basic variants. Any of the end systems - the media proxy or home gateway, the edge router at the ingress point of the network, or the border gateway - might be assigned the role of measurement and decision entity.

The items for measurement from which we identify the remaining bandwidth are packet loss and/or delay. Explicit congestion notification initiated by the network may also provide supplementary information.

For the sake of simplicity, we would like to focus on an approach that is 1) end-system oriented, 2) loss-rate-based, 3) includes no mechanism for explicit indication from the network.

As we have previously noted, the above concept is not new. It was originally proposed by Karlsson as probe-based admission control (PBAC) [3][4]. Based on Karlsson's proposal, we would like to extend the measurement-based approach to allow for various service models, to clarify the behavior required of routers, and to take into account monitoring of the correctness of end-terminal behavior.

4.4 Monitoring of terminal behavior

How we monitor, check, or audit the behavior of end systems is an important issue for a commercial service. Since the Priority Promotion Scheme is strongly reliant on the behavior of end systems, incorrect behavior, whether accidental or intentional, will affect the QoS for other customers.

Here, the items to be monitored include whether or not flows have been given permission to enter or access the network, whether flows are at the correct priority level, and whether flows are at the bit rates indicated by probing or signaled by SIP. These are the behaviors in the direction from source to destination. The behavior in the direction from the destination to the source should also be correct, and feedback reports on e.g. correctness of the conditions of receipt might be included to monitor this. Furthermore, the source behavior in response to such reports should be correct in terms of not promoting priority when the report indicates bad conditions. One of the benefits of the PPS is the allocation of resource-management functions to the end systems, since this reduces the burden on the network. If we implement functions of the kind just described to monitor the correctness of the behavior of end-systems, however, we place another burden on the network. There is a tradeoff between the extent to which we should protect the network and the costs of doing so.

The site of monitoring is another issue we face in designing the network. One solution is to install checking mechanisms of the kind described above in every edge router and have them monitor every session. This is perfect in terms of protecting the network from all kinds of incorrect behavior, but would cost too much.

Another practical solution is to introduce two-stage monitoring of end-system behavior. The intention here is to classify items for monitoring as either primary or secondary and having them checked at the appropriate places. Primary monitoring may be implemented at the edge routers and is triggered by session initiation. Secondary monitoring might be done by a dedicated media-monitoring server. The primary monitor checks every PPS-controlled media stream it handles. Examples of items to check include whether the flow has been given permission to enter the network, whether the flow rate is no greater than the probed bit rate, and the correctness of the usage of the DSCPs. The secondary monitor checks the details of end-system behavior. Whether or not the two monitoring stages are really used will depend on the specific network environment, but both should be available to allow flexibility in implementation.

4.5 Accommodation of variable-bit-rate sources

Any measurement-based form of admission control is more suitable with constant bit rate (CBR) sources than with variable bit rate (VBR) sources. CBR sources to which silence suppression is not applied are often used in public voice communications in Japan. For interactive multimedia, on the other hand, it is important that we take VBR into account.

Another approach is possible, relying on declared traffic parameters and deterministic capacity allocation rather than results of measurement. The admission control system gets the declared parameters, estimates the equivalent bandwidth, and then judges whether or not admission is possible. The drawbacks here are the difficulty of deriving truly representative parameters for each of the many popular codecs and of estimating the total required bandwidth when a new flow is offered.

VBR has quite different implications for a measurement-based approach such as PPS. PPS requires no parameters, no estimation, and no calculation. In addition, utilization of bandwidth is ideal because measurement is of actual traffic. There is, however, a trade off. The PPS depends on the usage of resources at the time of measurement. Measurement for a particular session may occur when the flows already present are at relatively low rates. The new session may then suffer loss of QoS when the volume of flows returns to typical levels.

The tuning of the PPS to support VBR sources thus has to reflect statistical variation, which can be done by probing over a longer time or by sending the probing packets at a higher rate than the non-probing packets. A new (elastic) mode of PHB provides a way of avoiding such mechanisms and is introduced in the definition of the MF-PHB[5].

Investigations with VBR sources including ON/OFF source have already been done by Prof. Karlsson as is indicated by the Appendix of the document.

5. Service models provided by the PPS

The Priority Promotion Scheme can be viewed as a kind of admission control. However, it is not limited to the kind of connection/session admission control we imagine if we think of the

legacy telephone network. The probing can even be handled by the media packets themselves. In this section, we examine the possible service models provided by the PPS.

5.1 Admission control

Admission control alone is suitable for conventional service models such as legacy switched services. The measurement is simply used for admission control when the session is established. If the trial fails, the session is not established. The user may retry, but the terminal behavior does not specify the extent to which this is possible. PPS is quite effective in this role as long as the duration of probing is less than a couple of seconds.

5.2 Quality improvement

The case of PPS where the media packets are used for probing is particularly applicable to quality improvement. The source starts by sending media packets at probe level. If the conditions of receipt are poor, the source stops sending the media packets at probe level, and recommences sending them as packets of another class. After a while, the source returns to probing; if this succeeds, the packets are sent as packets of the higher (non-probing) MF-PHB class.

5.3 Available bit rate

In the available-bit-rate service model, the transmitter uses the information on network conditions received in response to probing to estimate the actual available bandwidth, selects the closest bandwidth lower than the available bandwidth, and then sends the media at the higher MF-PHB priority level. The transmission may be made to fit the available bit rate by sending the video data with less size or resolution than was originally desired or sending speech data alone rather than a mix of video and speech. The quality of the session is then maintained.

A further possible application of this approach is to send media data at the full rate but only assign the higher MF-PHB priority to the core part of the flow, which fits the available bit rate; the other parts are sent but assigned to another class. This approach should work well with hierarchical coding (in MPEG for example, I frames would be sent with high priority and P or B frames with low priority).

5.4 Bit-rate increase

This is an extension to the available-bit-rate service model. If initial probing indicated that the requested bit rate is not available, the source sends at the lower rate than requested but retries probing from time to time. When the requested rate becomes

available, the source starts sending media packets at the requested rate.

6. The feasibility of probe-based admission control

Karlsson has already investigated the characteristics of probe-based admission control (PBAC). Although the overall system architecture of PBAC is slightly different from the PPS, the basic dynamics are the same and the analysis of PBAC is applicable to the PPS. A summary of the analysis is thus given in the Appendix of this document.

7. Functional architecture of the Priority Promotion Scheme

Figure 1 shows the functional architecture of the Priority Promotion Scheme. The main functional elements are the two end systems, i.e. the source and destination, the source-side edge router, the core routers, the SIP proxy, and the media-monitoring server.

```
                   SIP proxy      (Media-monitoring server)
                    |------|            |------|
          /---------|      |------------|      |
         /          |------|            |------|
        /              |               //
       /               |              //
  |------|         |------|      |------|      |------|      |------|
  |      |=========| Edge |======| Core |======| Edge |======|      |
  |------|         |------|      |------|      |------|      |------|
End system                                                 End system
(Source)                                                 (Destination)
```

Figure 1. Functional architecture of the Priority Promotion Scheme

8. Requirements of the Priority Promotion Scheme

In this section, we describe the requirements for the various functional entities.

8.1 Routers

Although the end systems play an important role in the Priority Promotion Scheme, the scheme places a few other requirements on the network. Specifically, the queuing mechanism or PHB (per-hop behavior) for the PPS creates new requirements for network elements. The Priority Promotion Scheme is intended to work with the existing Diffserv PHBs, as was indicated in the introduction. However, to clearly explain how the scheme would be implemented in this context, we have to define a new PHB. We refer to this as measurable forwarding (MF). The essential requirements for MF are as follows.

- MF has two sub-classes, MF-High (MF-H) and MF-Middle (MF-M).
- MF-H and MF-M share the same capacity.
- MF-H takes priority over MF-M.

In other words, we have a total amount of MF-H and MF-M traffic as a limit rather than separate limits for the two sub-classes. However, since MF-M traffic will always defer to MF-H traffic, MF-M traffic may experience markedly higher levels of jitter and loss than MF-H, while one would expect MF-H traffic to experience very low levels of jitter and loss.

Another view of MF is that, if a given amount of MF-M traffic for a particular stream passes through a router, at least the same amount of MF-H traffic for that stream must also be able to pass through. In the absence of other DiffServ classes, configuring existing commercially available routers to implement the MF-PHB should be feasible. Further requirements are as follows.

1) The MF must co-exist with other PHBs, such as the EF, AF, and BE. Existing implementations may not be capable of satisfying this extended requirement.
2) MF should take priority over AF and BE. This is because the target services are IMM services, where real-time variations in traffic characteristics are crucially important.

The more detailed definition of MF-PHB and scenarios for its verification are available in [5][6].

8.2 End systems

The transmitter should send trial packets before or at the beginning of a session.

The receiver should record the results of trial-packet reception and report this information to the transmitter.
The RTCP would be the best candidate to handle reporting of the results of reception. Some improvements might be necessary to reduce the measurement period and to make quick decisions. Actually, the minimum measurement period is the key factor that determines the usability of the Priority Promotion Scheme. This determines whether or not the scheme is applicable to admission control, as was described in section 5.

The transmitter then decides on the next action.
- If the conditions of reception are good, the transmitter sends the remaining packets with the higher priority.
- If the conditions are not good, the transmitter gives up sending monitor packets and either 1) sends the remaining packets with

another class such as BE, 2) stops sending any media data and, after a while, starts sending monitoring packets again, or 3) terminates the session.

According to the service models described in section 5, further actions are necessary.

Synchronization between the two directions of the media stream remains a subject for further study.

8.3 SIP proxies

In principle, SIP is not directly related to the Priority Promotion Scheme. However, for commercial applicability, the operator would have to be able to monitor the service subscription of the customer before establishing the call. Furthermore, if the edge router is capable of monitoring user streams, an SIP proxy can send commands to an edge router, requesting that it check on a particular end system's behavior.

The specific signaling sequence may depend on the selected service model.

If the policy is applied as was described in section 5, signaling is where the policy credentials are exchanged.

8.4 Edge routers

As noted above, in some networks an SIP server might be available and is able to instruct edge routers to monitor the behavior of end systems. An edge router might monitor the following items.

- Packet-transmission rates: the transmitter should not send packets at rates above the peak bit rate offered in the monitoring phase.
- Continuous sending of packets: if the transmitter pauses in the sending of packets, the other end systems overestimate the remaining network resources and incorrectly send higher-priority packets. Transmitters should thus not pause during sending.

8.5 Media monitoring servers

In addition to primary monitoring by the edge routers, more detailed monitoring may be required. The typical items to be monitored are as follows:
- the accuracy of packet-reception information from receivers, and the correctness of reactions of transmitters to this information; and
- if the received information indicates poor conditions, the transmitter stops sending high-priority packets; if a next trial is

allowed, a certain time interval should be maintained between the initial trial and the next trial.

9. Security Considerations

To be described.

10. IANA Considerations

To be described.

Acknowledgements

The authors would like to thank Fred Baker, David Oran, Glenn Reitsma and other technical experts at Cisco for some insightful suggestions.

Authors' Addresses

Naotaka Morita
Network Service Systems Laboratories
NTT Corporation
9-11, Midori-Cho 3-Chome,
Musashino-Shi, Tokyo
150-8585 Japan
E-mail: morita.naotaka@lab.ntt.co.jp

Gunnar KARLSSON
KTH, Royal Institute of Technology
Department of Microelectronics & Information Technology
Laboratory of Communication Networks
Isafjordsgatan 39
P.O.Box Electrum 229
SE-164 40 Kista, Sweden
E-mail: gk@imit.kth.se

References

1 Bradner, S., "The Internet Standards Process -- Revision 3", BCP 9, RFC 2026, October 1996.

2 Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

3 Karlsson, K., "Providing Quality for Internet Video Services," in Proc. of the CNIT/IEEE 10th International Tyrrhenian Workshop on Digital Communications, Ischia, Italy, September 15-18, 1998.

4 Fodor, (nee Elek), V., Karlsson, G., and Roenngren, R., "Admission Control Based on End-to-End Measurements," in Proc. IEEE INFOCOM, Tel-Aviv, Israel, March 26-30, 2000.

5 Morita, N., " Measurable Forwarding: A New per-Hop Behavior (PHB) ," Internet draft, October 2003.

6 Morita, N., " Verification scenarios for Measurable Forwarding PHB (Per-Hop Behavior)," Internet draft, October 2003.

Appendix: Probe-Based Admission Control (PBAC) - Current experimental results and obervations

1. System definitions

. Complete semantic definition of the probe-based admission control [A1, A2].

. Multicast application of PBAC [A3]. The quality of service scheme for multicast traffic is based on admission control for both senders and receivers. The admission control is well suited to multicast sessions with a single multimedia stream or with several layered streams.

. Simple security model to verify the end host identities and secure the probe phase and the admission decision [A4]. The scheme verifies the end user's identities and secures the transmission during the probing phase.

2. Analytical models

. Approximate mathematical model that relates probe and data packet loss rate, queue buffer sizes and achieved link utilization for the double queue system [A5]. The analysis is based on the following steps: First, computation of the probability of a single probe packet being successfully transmitted; second, computation of the acceptance probability as a binomial distribution; third, computation of the link utilization as a birth--death Markov chain; and fourth, computation of the data packet loss for a particular source type and the probe/data loss relationship.

. Numerical results with figures for probe packet loss probability, acceptance probability as a function of the load on the system, link utilization and data packet loss probabilities. The results agree with the simulations and prove that the considered probe--based admission control leads to a stable link utilization and has a clear upper bound on the packet loss probability.

3. Performance evaluation

All the performance figures have been obtained with the NS-2 simulator. Different source types and source rates have been used: sources with exponential and Pareto on--off holding times and traces of real MPEG-2 encoded videos, with peak rates from 64 kb/s to 10 Mb/s. The sources are listed in Table 1. The following issues have been investigated:

. Performance and comparison of the proposed queuing schemes for the controlled load service, a double queue system with two priorities and a single queue system with a discard threshold for probe packets [A2]. Both queue systems can be used with a proper buffer and threshold dimensioning.

. The validity of the assumption of a normal distribution of the probe packet loss for the admission decision [A2]. Histograms of the probe packet loss prove the assumption valid.

. Stress test with short sessions and sessions that keep silent for long periods of time [A2]. None of this special sessions have a serious effect unless they represent a substantial percentage of the link capacity (over 15 %). The performance of the system under heavy stress (many simultaneous probes or sessions that keep silent for periods of time longer than some probe lengths) is stable. In general, as the situation worsens, the admission control is conservative, allowing less ongoing sessions, but never failing to keep the data packet loss under the threshold for maximum session peak rates of less than 5% of the link capacity.

. Relationship between probe packet loss and session data loss for different source types and peak rates [A1, A2]. Basically all source types show between half to one order of magnitude difference. All the figures show that there is a nearly linear relationship between the probe and the data packet loss.

. Effect of multiple links scenarios with cross traffic [A1]. The simulations prove that the bottleneck link dominates the behavior.

. Blocking and data packet loss probabilities and their relation to the probe length and the location of a multicast receiver [A3]. The simulations prove that receivers in different branches of the multicast tree have different blocking probabilities, depending on the link loads on the different multicast branches.

. Performance evaluation of an implementation of the security model proposed in [A4] with commodity hardware, focusing in the trade off between security level and setup delay. The simple solution does not require any change in the network nodes, just a cryptographic interface in the access gateways and the end nodes.

Table 1: Parameters of the different test sources

Source	On Time	Off Time	Peak Rate
Exponential	20 and 325ms	35.5 and 650ms	64kb/s to 10Mb/s
Pareto (fi=1.5)	20 and 325ms	35.5 and 650ms	64kb/s to 10Mb/s
Mixed	20 and 325ms	35.5 and 650ms	64kb/s to 10Mb/s

Video Traces	360kb/s (64kb/s average)

4. On-going work

. Software implementation of PBAC for Linux. A library to provide the probing features is being developed, which will enable software generators or end applications to perform the probing before transmitting. The queuing system will be implemented using the QoS capabilities of the Linux kernel (iproute2 (1)).

. A possible policing and metering tool for PBAC is under investigation using Netramet (2).

References

[A1] Viktoria Elek, G. Karlsson, and R. Roenngren, "Admission control based on end-to-end measurements," in Proc. of the 19th Infocom, (Tel Aviv, Israel), pp. 623--630, IEEE, March 2000.

[A2] I. Mas Ivars and G. Karlsson, "PBAC: Probe--based admission control," in Proc. of QofIS 2001, vol. 2156 of LNCS, (Coimbra, Portugal), pp. 97--109, Springer, September 2001.

[A3] I. Mas Ivars, V. Fodor, and G. Karlsson, "Probe--based admission control for multicast," in Proc. of the 10th IWQoS, (Miami Beach, Florida), pp. 99--105, IEEE, May 2002.

[A4] M. Conte, I. Mas Ivars, V. Fodor, and G. Karlsson, "Policy enforcing for probe--based admission control," in Proc. of NTS 16, (Espoo, Finland), pp. 45--55, Helsinki University of Technology, August 2002.

[A5] I. Mas Ivars, V. Fodor, and G. Karlsson, "The performance of endpoint admission control based on packet loss," in Proc. of QofIS 2003, vol. 2856 of LNCS, (Stockholm, Sweden), Springer, October 2003.

(1) ftp://ftp.inr.ac.ru/ip-routing/
(2) http://www.auckland.ac.nz/net/NeTraMet/