PROVIDING ENHANCED QOS CAPABILITIES IN BBNS NETWORKS

November 29, 1994

Subir Varma
Leo Georgiadis, Roch Guerin, Kumar Sivarajan

Emerging Technologies
Networking Systems Architecture
Research Triangle Park, North Carolina
919-254-7493

Advanced Networking Laboratory
IBM T.J. Watson Research Center
Hawthorne, New York

ATM QOS Parameters

- Cell Loss Ratio (CLR)
- Cell Transfer Delay (CTD)
- Cell Delay Variation (CDV)

QOS Capabilities in BBNS

- Reserves BW in the network based on (PCR, SCR, MBS) values.
- Provides a single CLR for all sessions on a link
- Provides worst case end-to-end delay and jitter guarantees

Question: Do we need to do more?

We will make a case for providing tighter end-to-end delay and jitter guarantees

Proposal

Enhance the QOS capabilities in BBNS by,

- Providing tighter end-to-end delay and jitter guarantees
- Providing a finer granularity for these guarantees
- Decouple these guarantees from network path (no. of hops)
- Guarantee losslessness (if desired) for these connections

The mechanisms to implement the enhanced QOS capabilities will co-exist with the present BW management mechanisms

Who Will Use the Enhanced QOS Capabilities

Real-time applications can be classified into the following two categories:

- Intolerant and rigid real-time applications (IRA): They need a hard end-to-end delay guarantee to function.
- Tolerant and adaptive real-time applications (TRA): They can live with soft end-to-end delay guarantees. They can adjust the size of their play-out buffer at the receiver to adapt to network congestion

The IETF (RFC 1633) is discussing a similar classification

Are our present QOS capabilities sufficient to support IRA and TRA applications?

We will argue that the enhanced QOS capabilities are necessary for IRAs, while the present BW management framework is suitable for TRAs and data traffic.

Our fundamental claim:

Using our proposed scheme, it is possible to achieve higher link utilization for delay constrained traffic (as compared to the present scheme)

Reasons

- In order to keep end to end delays small, we constrain the link buffers to a small size. However this means that Equivalent Capacity is close to PCR, thus leading to low utilization
- To remedy this situation, we may increase, the link buffer sizes, but this leads to larger end-to-end delay bounds, thus limiting the traffic to those streams with a larger delay constraint

Example

The following parameters were obtained by processing a real video trace

PCR = 10 mbps, SCR = 6.8 mbps, MBS = 2476 cells

Buffer Size	Equivalent Capacity	Delay Guarantee
16 Kbtes	9.82 mbps	0.82 ms
100 Kbytes	9.09 mbps	5.16 ms
1 Mbyte	7.38 mbps	51.63 ms
2 Mbytes	7.13 mbps	103.2 ms
10 Mbytes	6.92 mbps	516.1 ms

Our Solution

- Use the present RT-2 class for adaptive and tolerant applications. Do not use worst case end-to-end delay as a criteria for admission or path selection. Instead, make real time estimates of the end-to-end delay, and adjust the play-out threshold of the adaptive applications based on that. This will greatly increase the link utilization
- Introduce a new class, say RT-0, at a higher priority than RT-2, that uses our new schemes to ensure end-to-end delay and jitter. Intolerant and rigid applications would use this new class

Value add of Tighter Delay Guarantees

- Provides state of the art networking support for intolerant and rigid real-time applications
- Higher network utilization is achieved for all real-time applications, whether of the rigid or adaptive type
- Providing tight delay control does not mean that peak BW has to be reserved.
- Losslessness capability
- Makes possible a single delay and jitter guarantee to a flow with multiple destinations
- Since it decouples QOS attributes from network path, it will be possible to re-route a flow over a different number of hops, without changing the delay and jitter guarantees
- Removes uncertainty about traffic characteristics in the interior of the network
- It will position BBNS better as a best of breed sub-network technology for IP networks because the IETF is also in the process of defining real-time support

Techniques for Guaranteeing Delays

- Weighted Fair Queueing (WFQ) Based Technique
- Per Node Rate Regulation (PNRR) Based Technique. The scheduler can be any one that can guarantee delay bounds. Examples are:
  - Non Pre-emptive Earliest Deadline First (N-PEDF)
  - Static Priority (SP)
- PNRR based techniques with NPEDF scheduling has a larger schedubality region as compared to WFQ based techniques

Cost of Enhanced QOS Support: Hardware Support

- Each session needs to specify (σ, ρ) and D
- Each switch needs a GCRA per session to regulate that sessions traffic, as well as an NPEDF scheduler to the outgoing link. To reduce the implementation costs, only a finite number of deadlines say K, may be provided at each node

  Discussions are underway with PRIZMA switch developers in Zurich, who have indicated that providing per session GCRA is feasible

- Structures needed to implement EDF are:

  1. A linked list of size K that orders the HOL packets of the queues that correspond to the K classes
  2. Timestamps for the packets that join the queues

Cost of Enhanced QOS Support: Communication and Computation

- Need new fields in the topology database for the $(\sum\sigma_i, \sum\rho_i)$ information
- The computation cost consists of:
  1. For a given (σ, ρ) session, the min. delay that can be guaranteed at a link, say D_m
  2. Path computation with D_m as weights
  3. The allocation of the slack $D - \sum D_m$ to the nodes on the optimal path

Providing Jitter Guarantees

- Present technique for obtaining a jitter bound too conservative
- Strict jitter guarantees required for certain application such as MPEG-2 Transport Stream

How to Guarantee Jitter

- A information field with K bits is exchanged between nodes
- If the end-to-end delay guarantee is D and the transmission + propagation delay is D_P, then K bits in the header would guarantee jitter J, where

$$J = \frac{D - D_P}{2^K}$$

Open Issues

- Path selection algorithms
- Taking PCR into account in the NPEDF analysis
- If lossless transmission is guaranteed then the buffer requirements at each node may be too large. Can source smoothing be used to reduce the buffer requirement?
- Can source smoothing be used to control network delay?

Conclusion

- BBNS networks provide the best of breed QOS capabilities in the present generation of high speed networks
- It is important that we maintain this lead over the rest of the industry by continuing to enhance BBNS by providing better QOS support