Worst-case latency in 802.1Qav Ethernet bridges
(v1 – very preliminary)

Michael Johas Teener

March 16, 2008

Notes

- This calculation is only for Class A
  - I want to make sure we understand the limits on a “2ms” latency network
  - Once we understand that, then I’ll add the Class B traffic to the analysis
- The parameters to be explored include:
  - Network topology (number of bridges and number of ports on each bridge)
  - Stream packet limitations (max packet size)
- For NOW all links are only 100Mbit/sec
- All shapers are as described in Qav 0.3

Input, output & methodology

- The input parameters to be explored include:
  - Network topology (number of bridges and number of ports on each bridge)
  - Stream packet limitations (max packet size)
- Output is worst case delay
- Looking only at first order effects
  - mention will be made of 2^{nd} order effects that are being ignored for now

Talker model

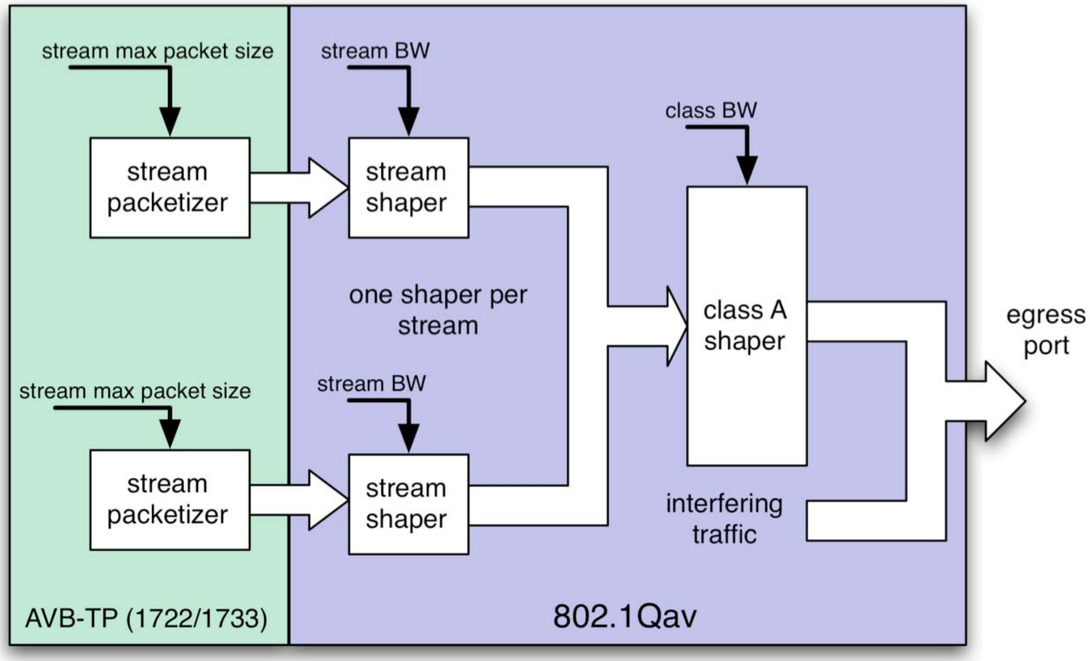

- Talker consists of transport protocol packetizers feeding into stream shapers feeding into class shaper
- Stream shapers have infinite “sendSlope”
- Sum of all stream’s “idleSlope” is the class “idleSlope”
  - SRP bandwidth allocation is “idleSlope”

Bridge model

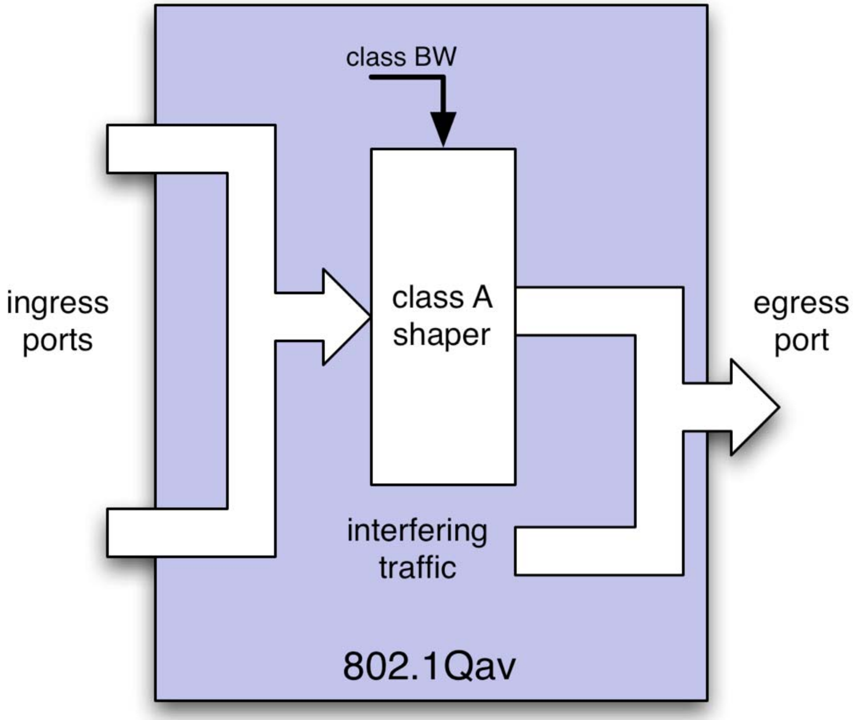

- Same as a talker with no stream shapers
  - conversely, a talker can be thought of as a bunch of single stream sources each with an infinitely fast link to a bridge

Talker delays

- Talker has just started to transmit a best effort frame of b bytes
- There are m streams, each with a max packet size of s_j bytes
- Egress port rate is e bytes/sec
- Delay is

$$\frac{\left(b + \sum_{i=1}^{m} s_i\right)}{e}$$

Bridge delays

- Bridge has just started to transmit a best effort frame of b bytes
- There are m ports, each routing class A traffic with a max packet size of s_j bytes through the egress port
- Egress port rate is e bytes/sec
- Delay is

$$\frac{\left(b + \sum_{i=1}^{m} s_i\right)}{e}$$

Network delays

- There are n bridges
  - so there are n+1 devices for queuing delays
- For each hop between devices there is no common stream
  - so it's possible for a stream to always be delayed by new interfering packets on each hop
- Delay is $$\sum_{j=0}^{n}\left(\frac{\left(b_j + \sum_{i=1}^{m} s_{ij}\right)}{e_j}\right)$$

Simplification

- all links are 100 Mbit/sec (e_j)
- worst case best effort interfering packet of 2048 bytes (b_j)
- all other class A packets are the same size (s_{ij})
- the talker launches m streams and each bridge has $m+1$ ports
- so, delay is

$$\frac{(n+1)(b+ms)}{e} = \frac{(n+1)(2048+ms)}{100}$$

2nd order effects

- Cumulative “bunching”?
  - I don’t think this is a problem since I’m forcing the interfering traffic on each hop to be uncorrelated (not following the measured stream on any other link)

Max network delays

#ports	16	8	4	2	1
max pkt	99	199	395	792	1500

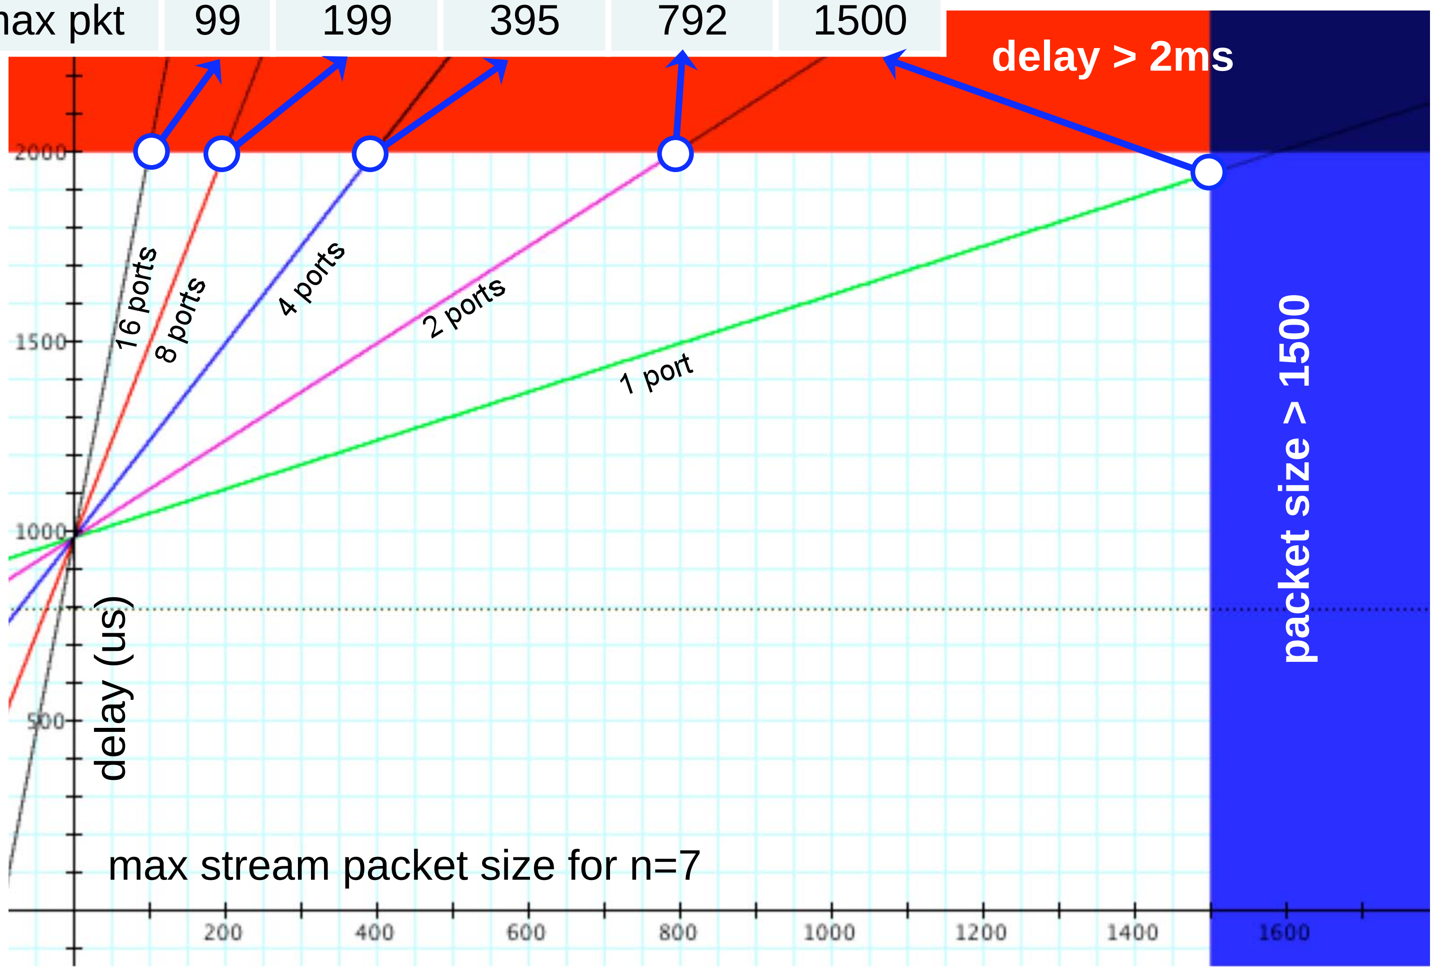

Conclusions

- Class 5 max packet size directly effects the latency, as does the number of bridges in a path from talker to listener, as does the number of ports on those bridges
  - We have been assuming 7 hops is a good limit for class A at 2ms max delay.
  - So we need to assume limits for the number of ports on the bridges and the max packet size
- For a 7 hop 100 Mbit/sec Ethernet configuration, we should perhaps assume 8 port bridges are a maximum
  - If so, then class 5 packets need to be no larger than about 200 bytes
- SRP *can* allow larger packets, but it will have to be ready to deny requests even when there is bandwidth available on a egress port
  - because the latency budget of “250 usec/bridge” is used up