RPR Priority Issues

Jon Schuringa, Günter Remsak

November 12, 2002

Introduction

- Guaranteed Bandwidth
  - Is only guaranteed if the STQ is large enough to hold as much traffic as could be received while adding local traffic and while waiting for upstream stations to back off in response to a congestion message from the local station.
- Example:

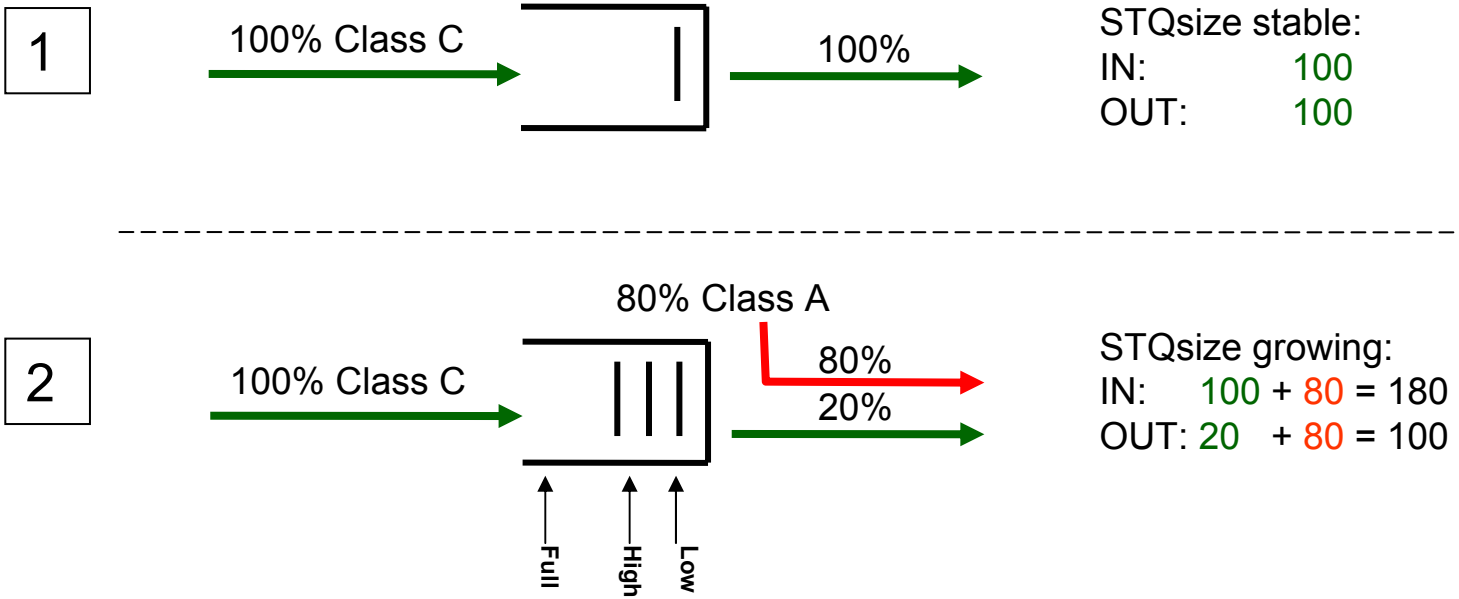

The incoming traffic must be reduced to 20% before the full threshold is reached, otherwise class A cannot get its 80%. (Packets on the ring may not be dropped)

Scenario 1

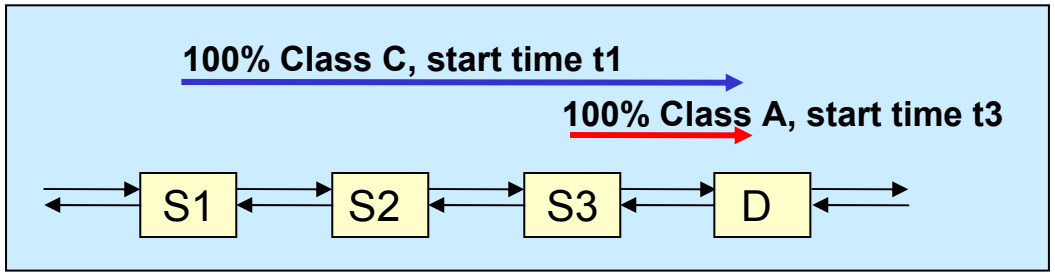

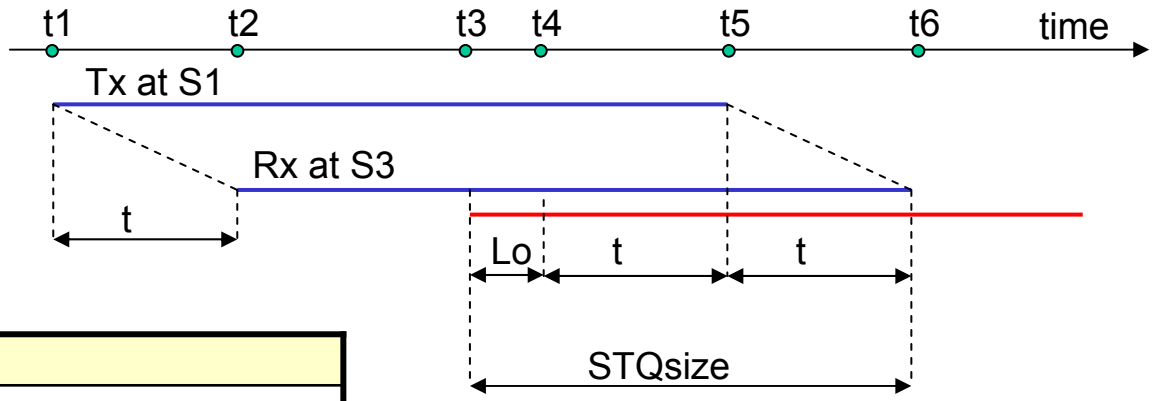

Time	Event
t1	Start transmission class C traffic at S1
t2	First byte arrives at station S3
t3	Start of class A1 transmission at S3, STQ is filling up
t4	Low threshold reached at S3, fairness message sent
t5	Fairness message received at S1, transmission stopped
t6	Last byte received at S3 from S1

STQsize = 2*t*cap + low_threshold
= 2*t*cap + 1/8*STQsize
STQsize= 16/7*t*cap

Scenario 2

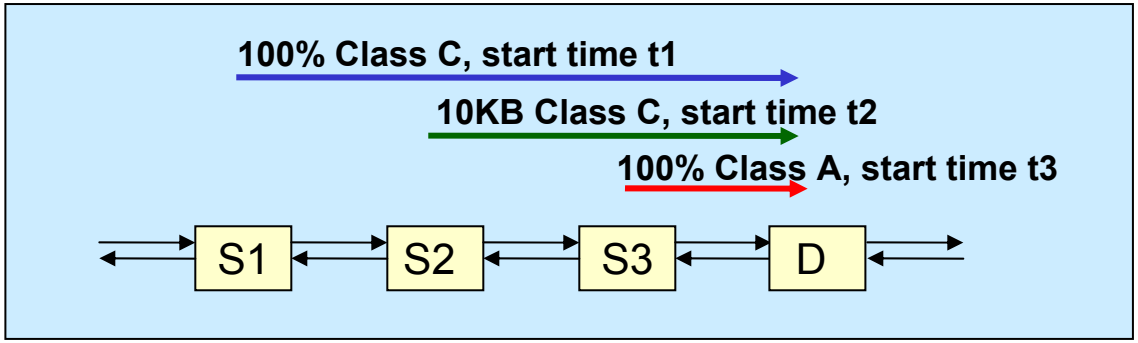

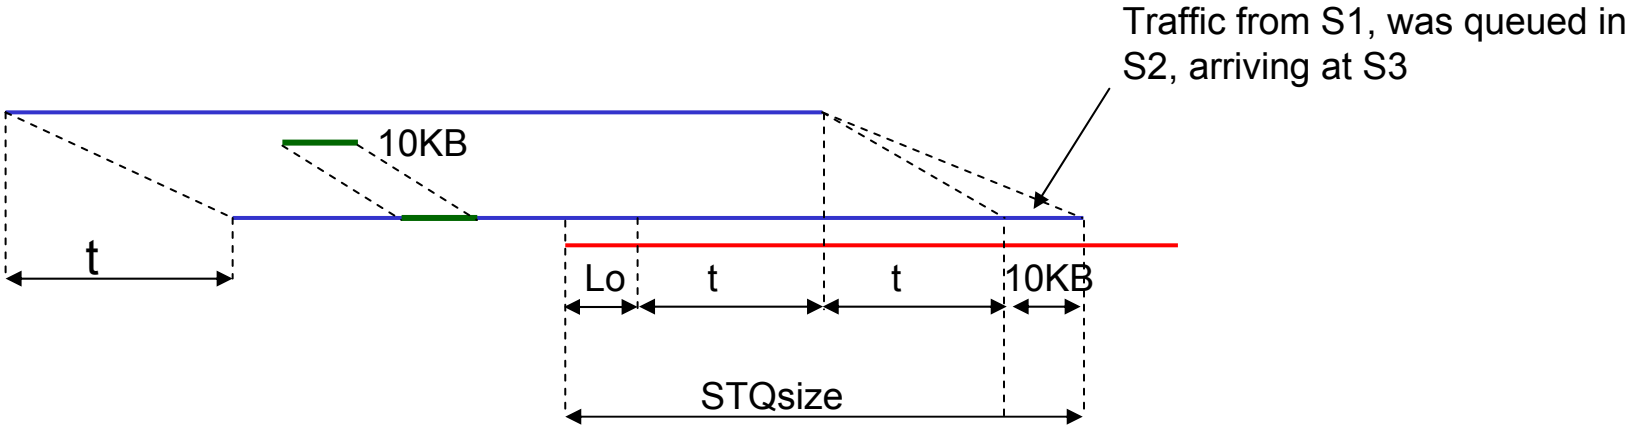

New formula:

STQsize = 2*t*cap + low_threshold * N

= 2*t*cap + 1/8*STQsize * N

STQsize = 16*t*cap / (8 – N)

N is the number of STQ's on the path

Low_threshold = STQsize / 8

Minimum STQsize

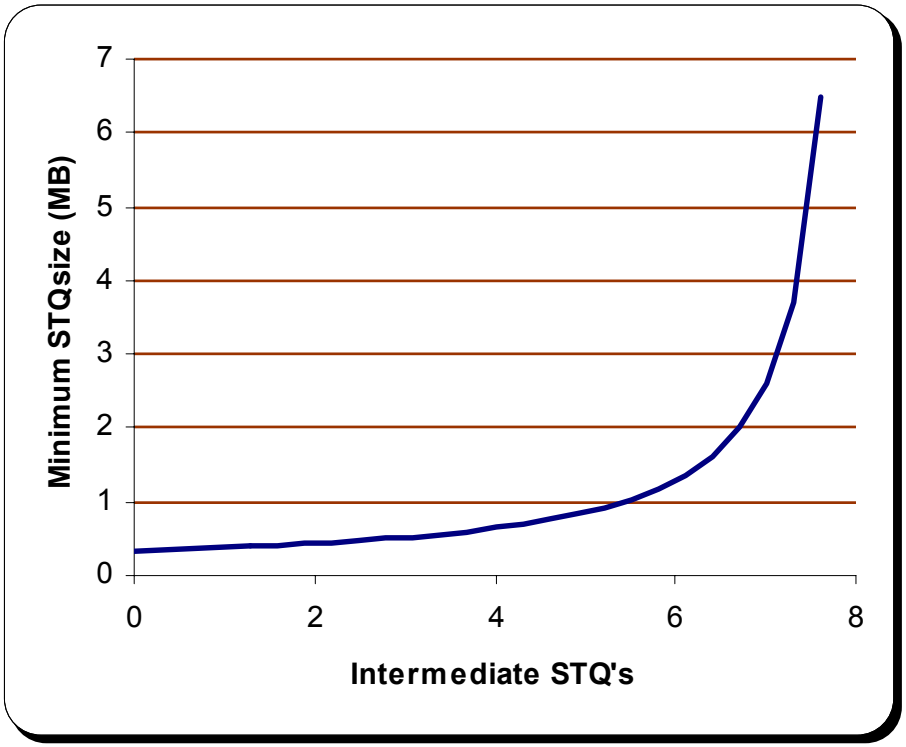

- STQsize = ∞ for N = 8

OC-192, 10 km distance between active sources (t=0.13 msec.)
low_threshold = STQsize / 8

What does it mean:

Scenarios exist where the current draft is unable to give bandwidth guarantees for N≥8, no matter how large the STQ is!

Simulation (1)

- 16 Stations, 622Mbps, 1500 km distance, 500 bytes packets
- STQsize = 2MB
- Scenario is similar to scenario 2, slide 4

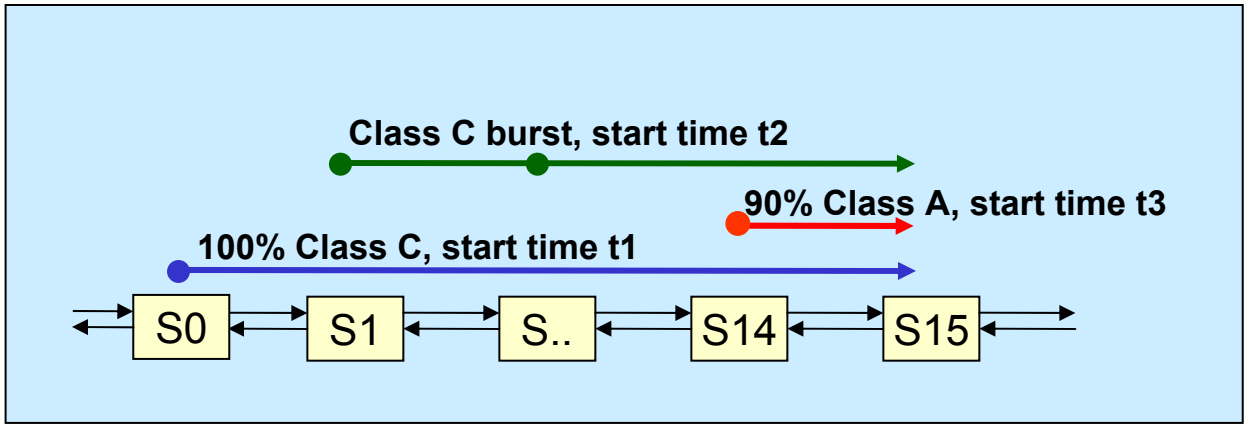

Time	Value (sec.)
t1	0.01
t2	0.02
t3	0.03

Simulation (2)

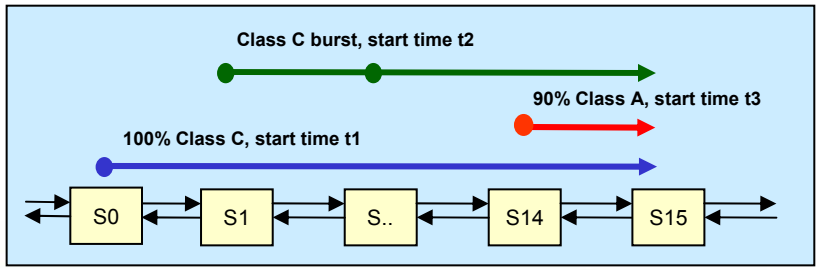

Class A on link 14 -> 15

Total throughput on link 14 -> 15

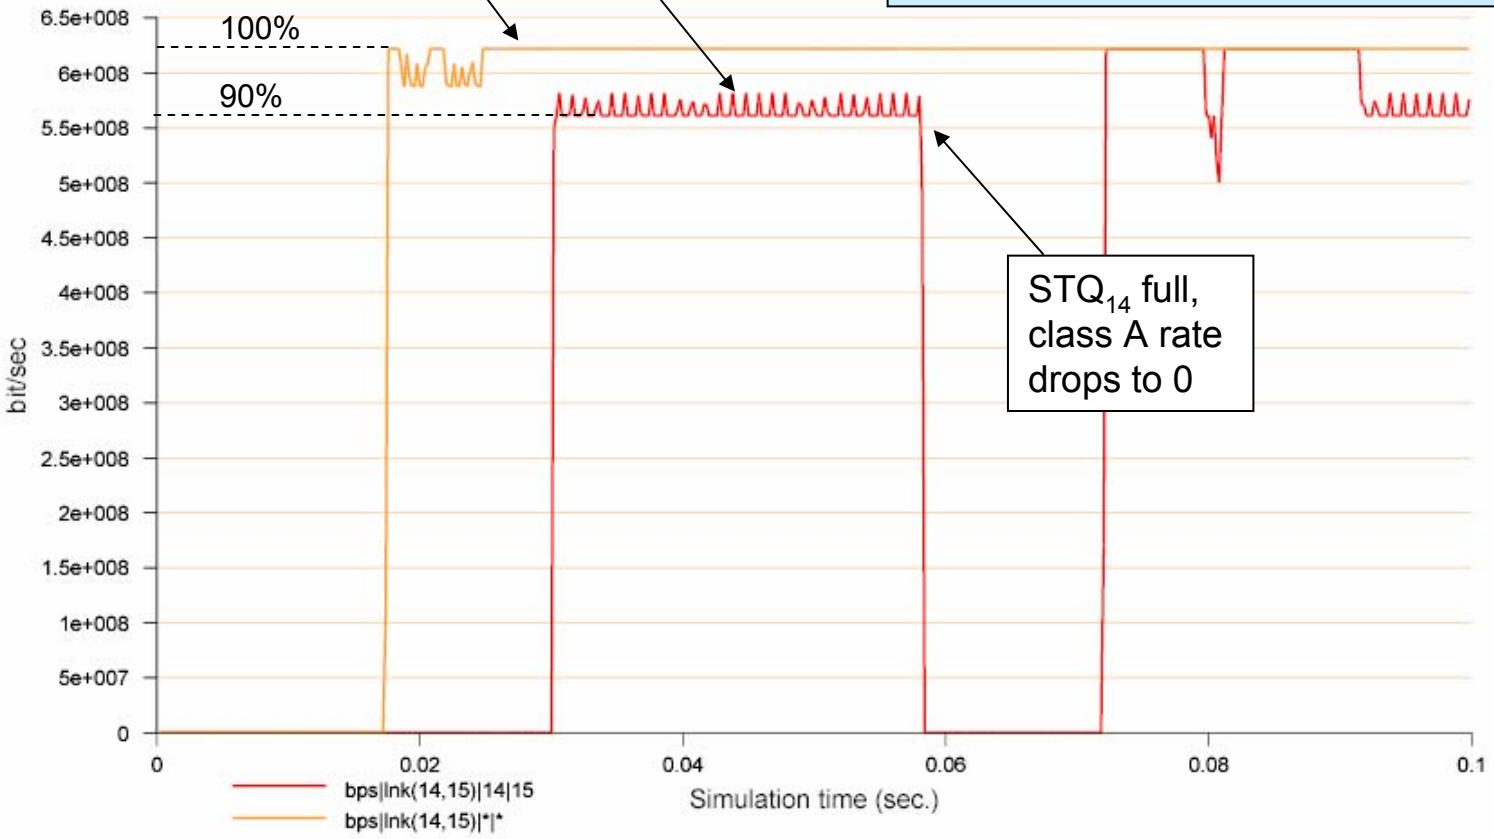

STQ_{14} full, class A rate drops to 0

Simulation (3)

- Different scenario:
  - Similar as before, but with different timing
  - First class A traffic, then class C

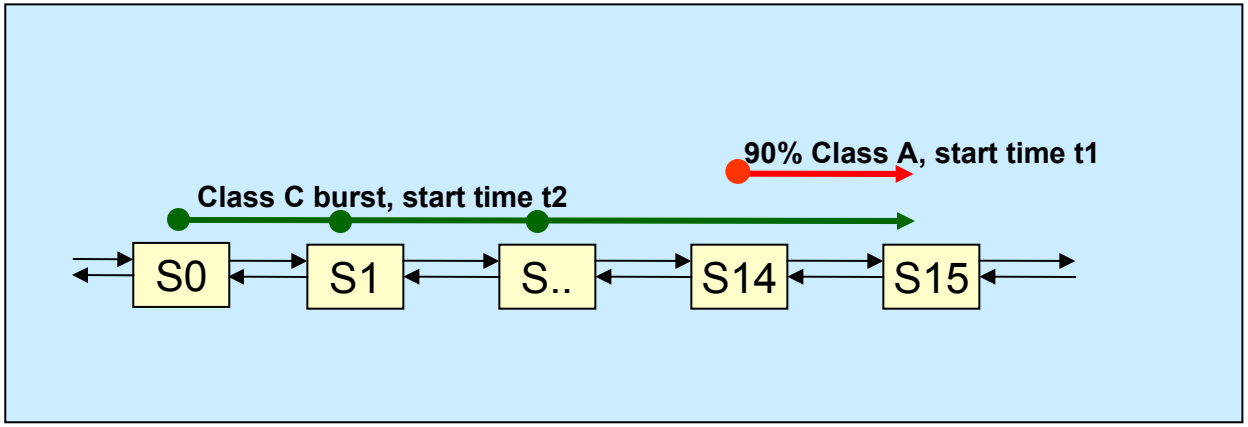

Time	Value (sec.)
t1	0.01
t2	0.03

Simulation (4)

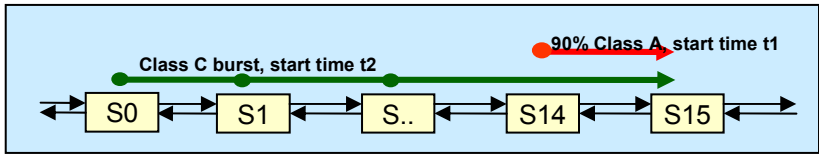

Class A on link 14 -> 15

Total throughput on link 14 -> 15

100%

90%

STQ_{14} full, class A rate drops to 0

bit/sec

6.5e+008
6e+008
5.5e+008
5e+008
4.5e+008
4e+008
3.5e+008
3e+008
2.5e+008
2e+008
1.5e+008
1e+008
5e+007
0

0
0.02
0.04
0.06
0.08
0.1

Simulation time (sec.)

bps|lnk(14,15)|14|15
bps|lnk(14,15)|*|*

Solutions?

- Limiting local traffic (Sum of class A,B and C)
  - E.g., if limit is set to 50% of the link capacity, guarantees can be made for $N < 16$
- Lower the low_threshold
  - Guaranteed traffic if: low_threshold ≤ STQsize / N
  - May cause bad link utilization
- Lower low_threshold, increase STQsize
  - Absolute value of low_threshold remains constant
  - Would require extremely large STQ‘s (higher cost and higher delay)

These are no real solutions

Limiting Local Traffic

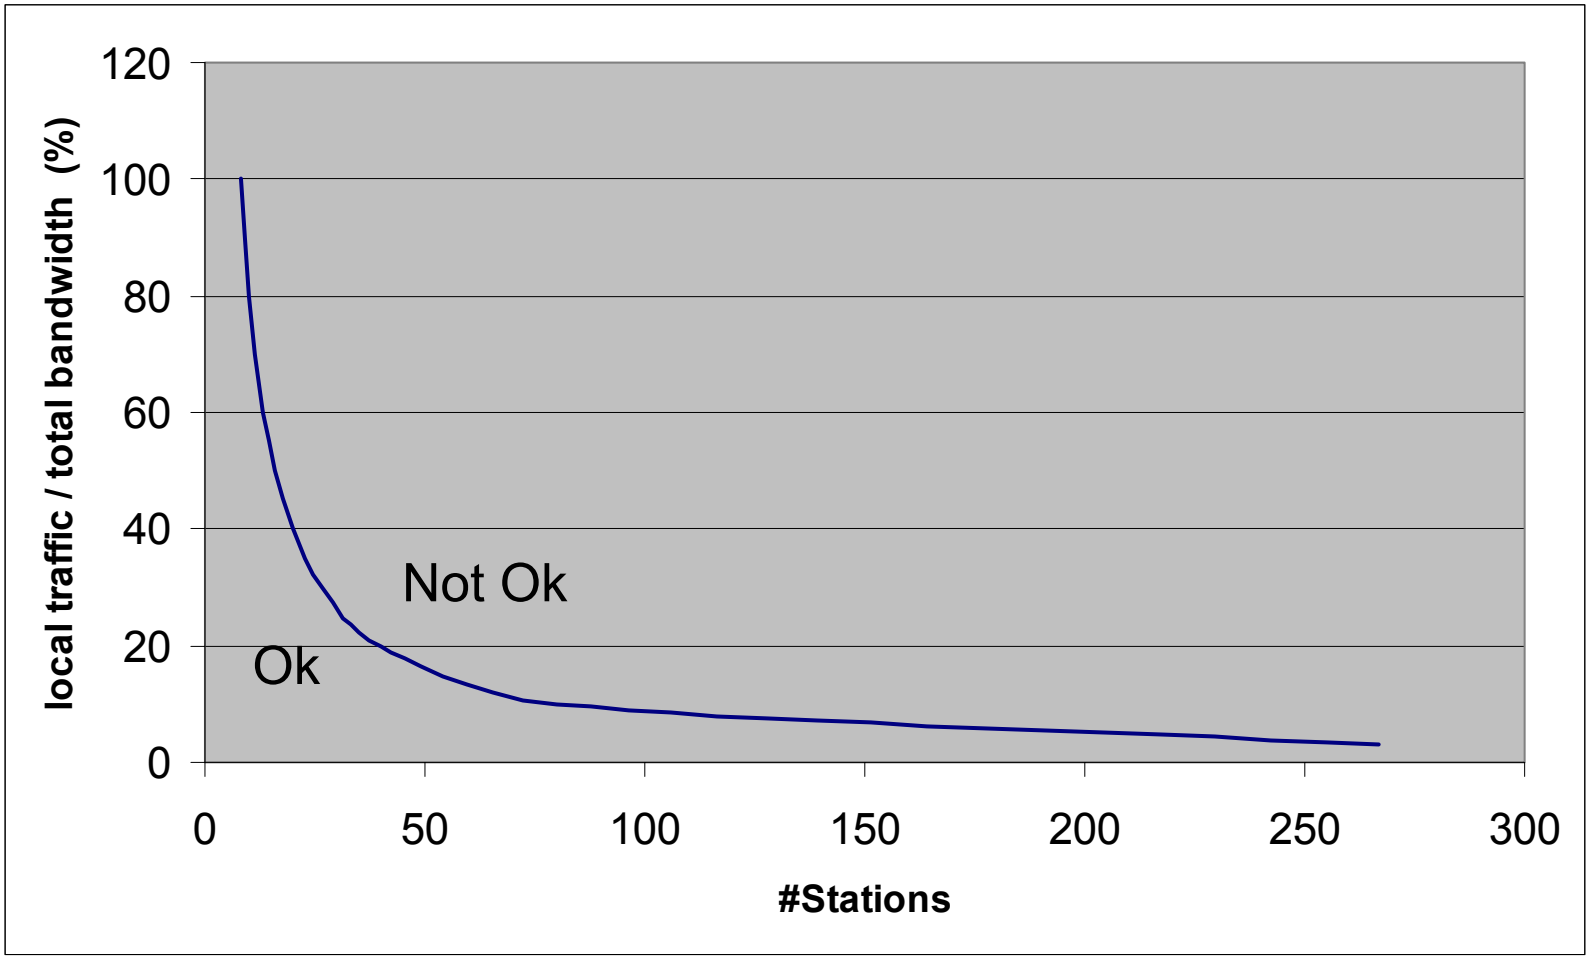

Lowering low_threshold

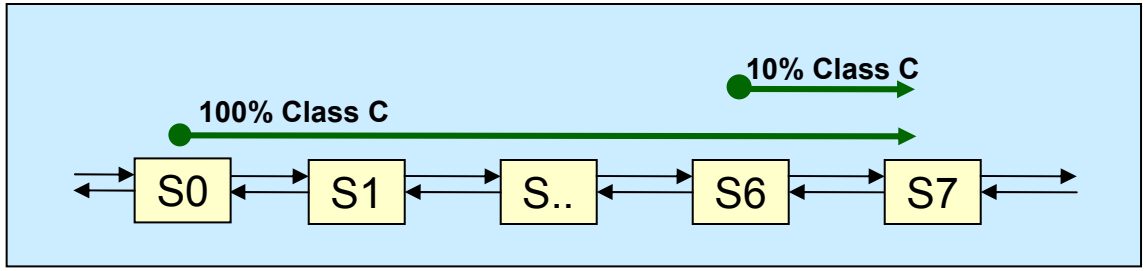

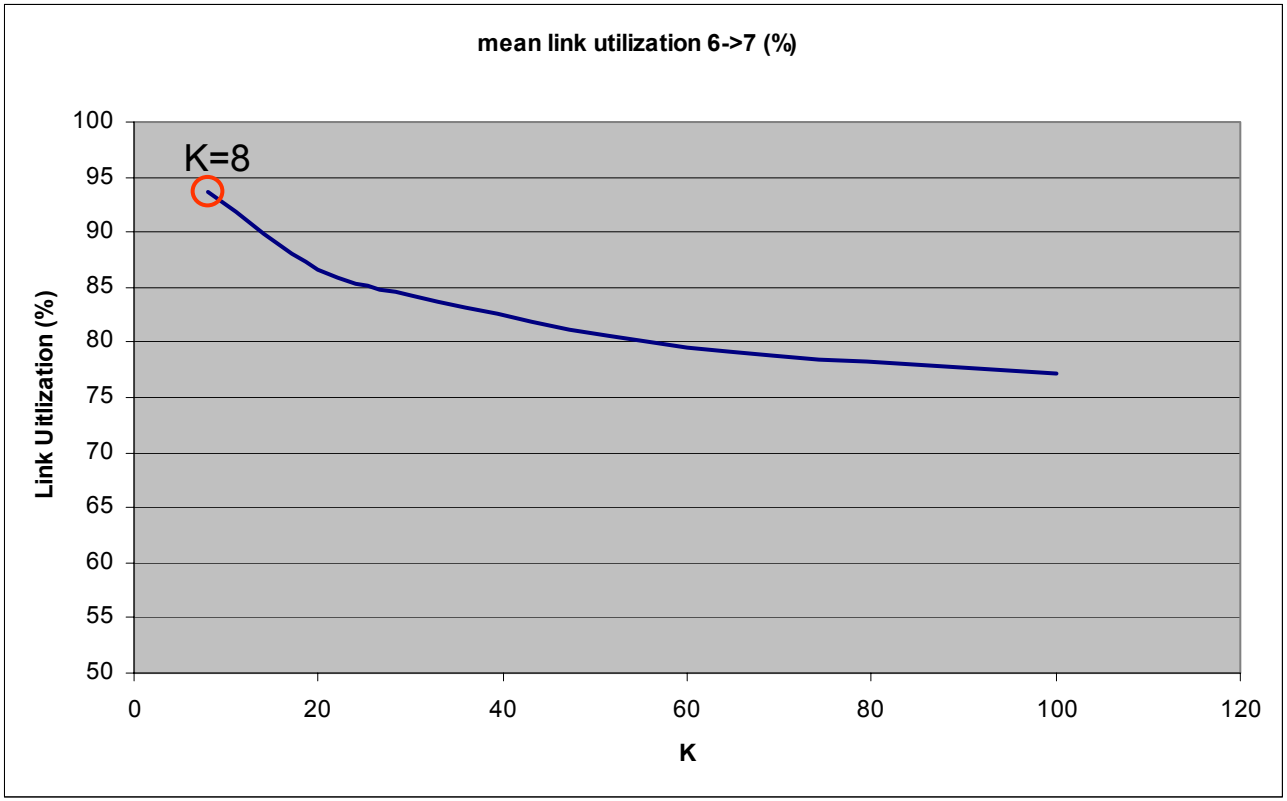

Low_threshold = STQsize/ K

Open Issues

- What if we have stations with (very) different STQ sizes on 1 ringlet?
- What if we have 255 stations?
- Worst case?
  - No, worse scenarios exist:
    - Intermediate STQ‘s can hold more than *low_threshold* bytes
    - On-off class A flow (?)
- Will it happen in practice?
  - More simulations with realistic traffic needed
  - It does not occur in SRP, because SRP does not provide guaranteed bandwidth. (And therefore uses different scheduling rules)
  - Example: Hub configuration with
    - flows to the hub creating a bottleneck (→ 100% link utilization between bottleneck and hub)
    - X other stations sending traffic to the hub (→ STQ’s fill)
    - A station wants to send Y% guaranteed traffic to the hub
      → potential problem for most X, Y combinations!

Conclusion

- The current draft can not give class A and B bandwidth guarantees for a large set of scenarios.
  - Increasing the STQsize does not help
  - Lowering the *low_threshold* would help, but has a negative impact on the link utilization
- Is there a good solution?
  - Yes!