Bandwidth Dimensioning in Packet-based Aggregation Networks

Wolfram Lautenschläger, Alcatel-Lucent Bell Labs

Wolfgang Frohberg, Alcatel-Lucent Optics Division

Networks 2008, Budapest, 2 October 2008

Agenda

1. Introduction: Dimensioning of packet link capacities
2. Theory of application stream collisions
3. Experiments and findings
4. Conclusions

Dimensioning of packet link capacities

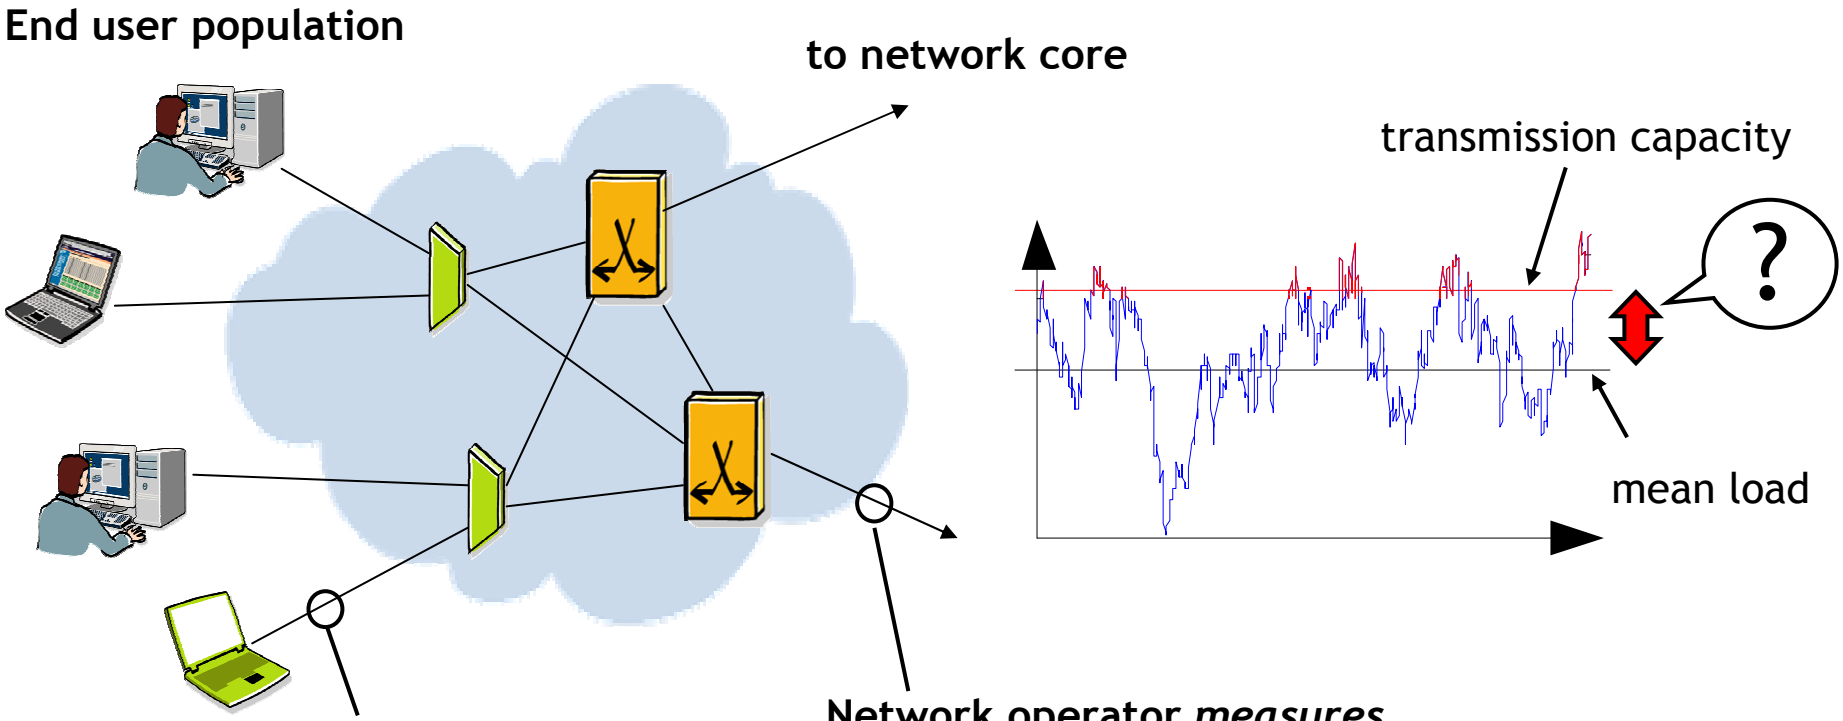

Network operator *knows*
→ the link capacity of individuals
→ but not their future link load

Network operator *measures*
the mean load
→ rather invariant
for large population
→ predictable in future

2 Theory of Application Stream Collisions

Multiplexing of packet streams is different from Poisson traffic

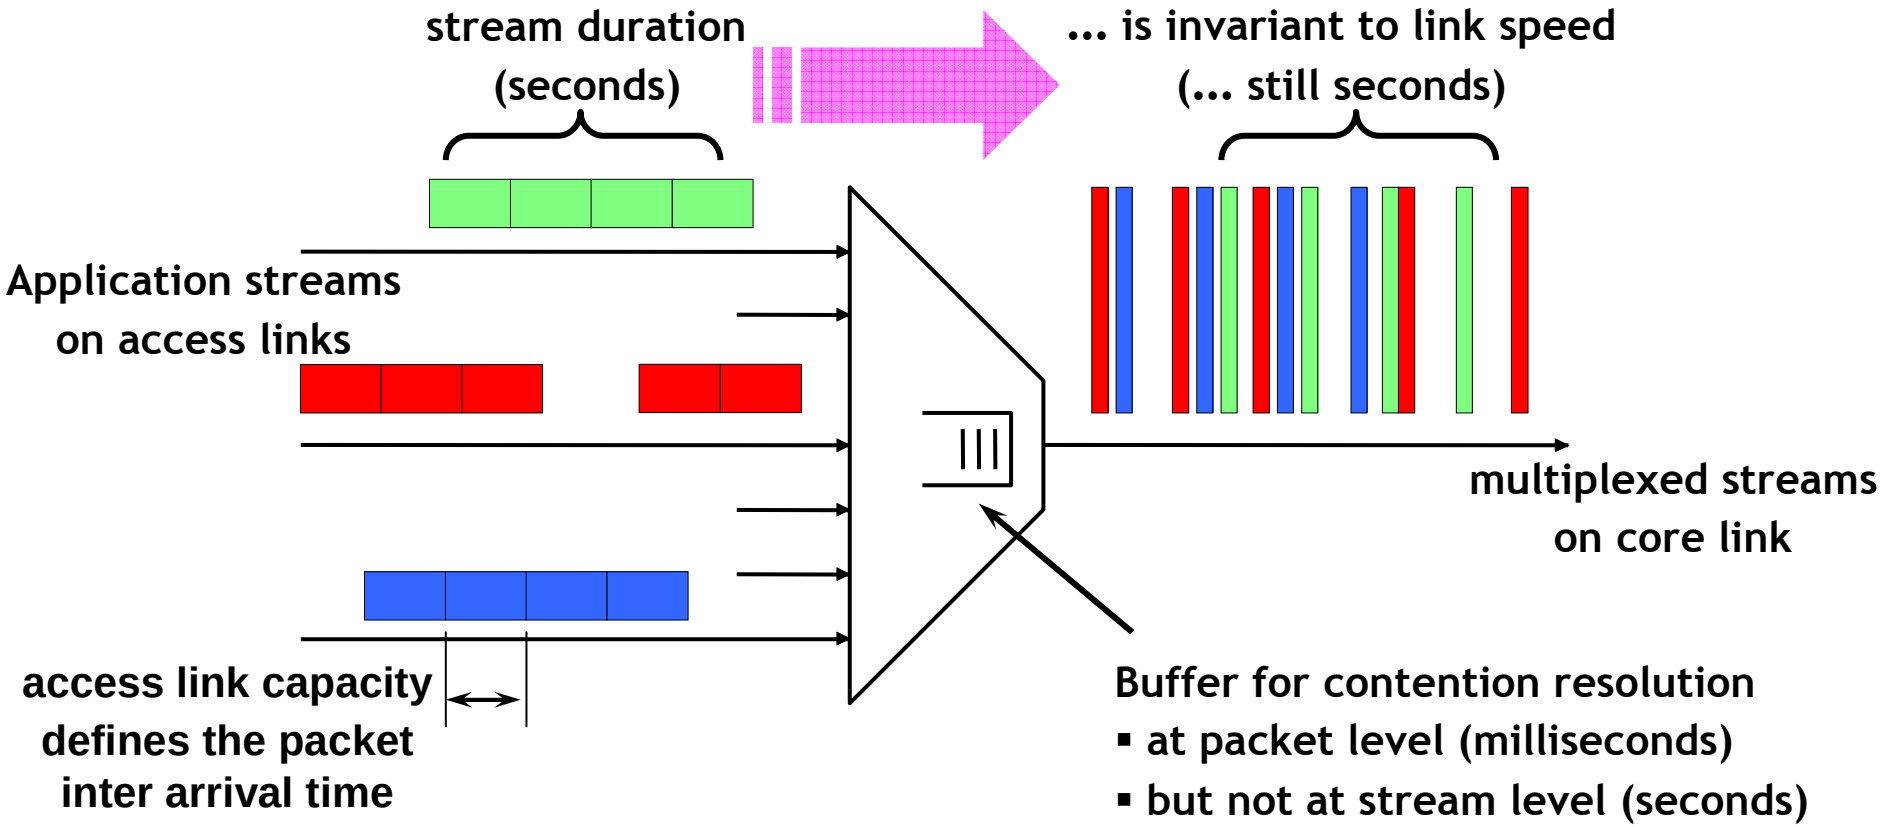

Link overflow = collision of streams
Overflow duration ≈ stream duration

Time scales in packet networks

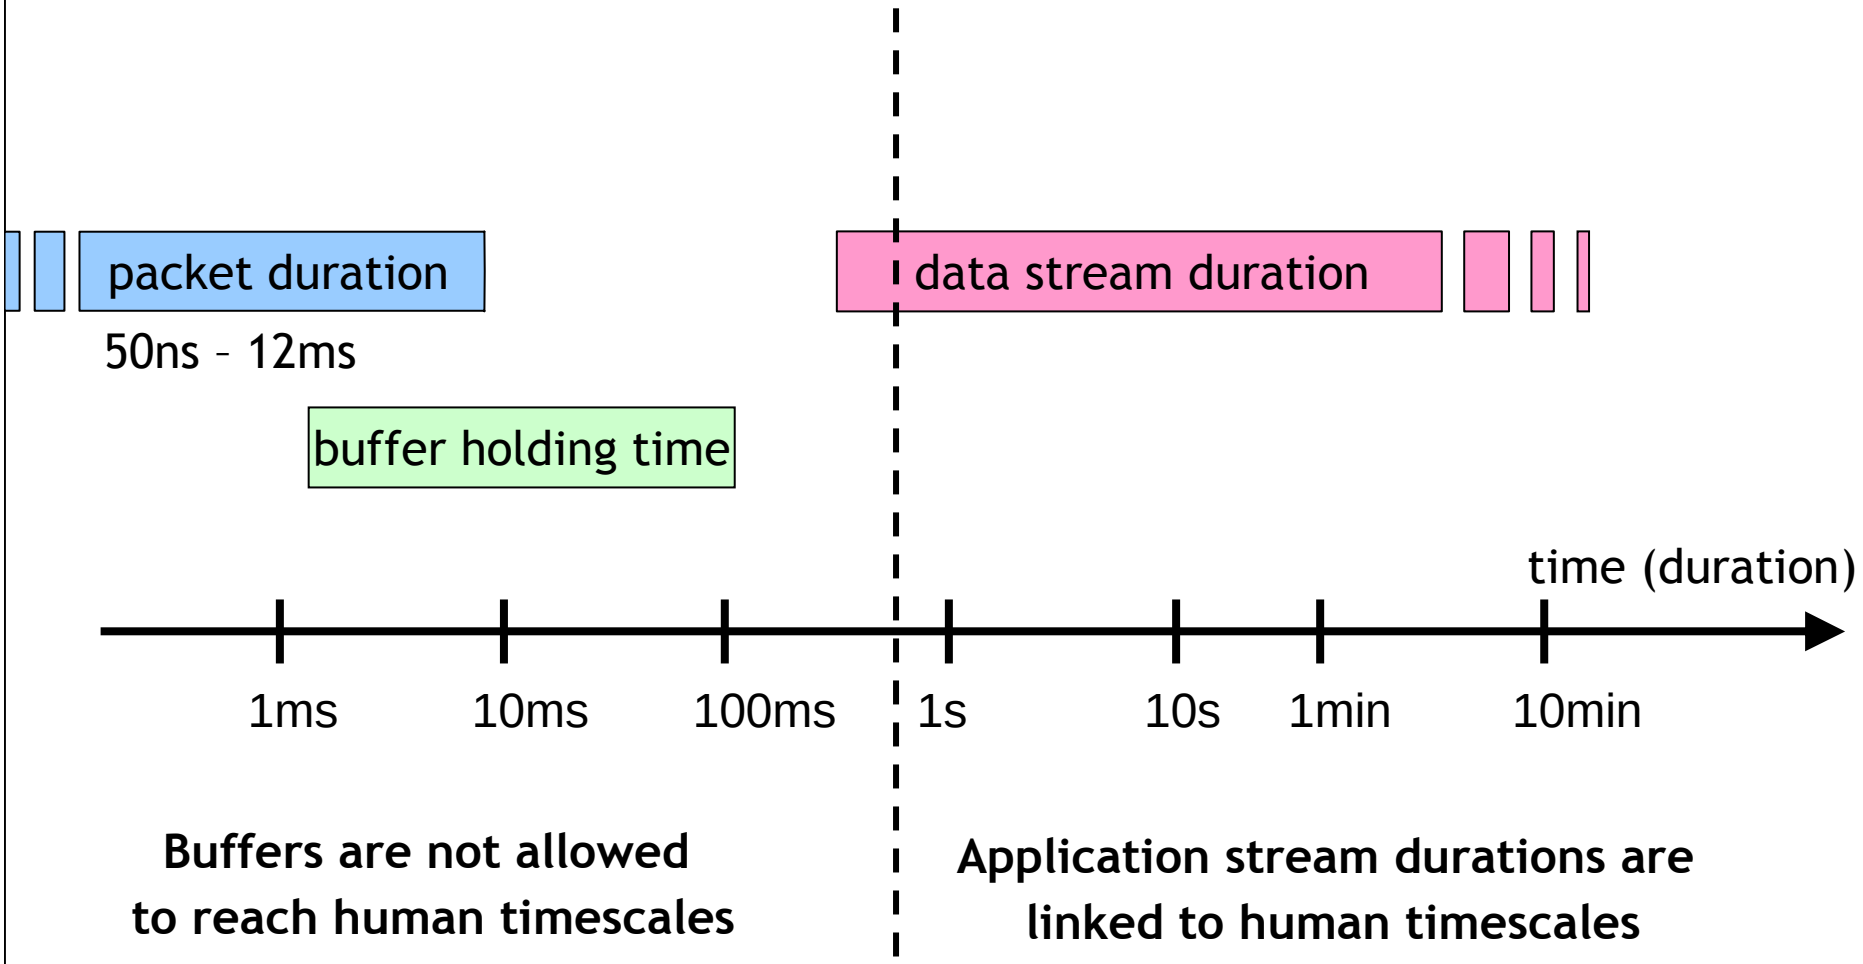

Blocking by buffer overflow in presence of application streams

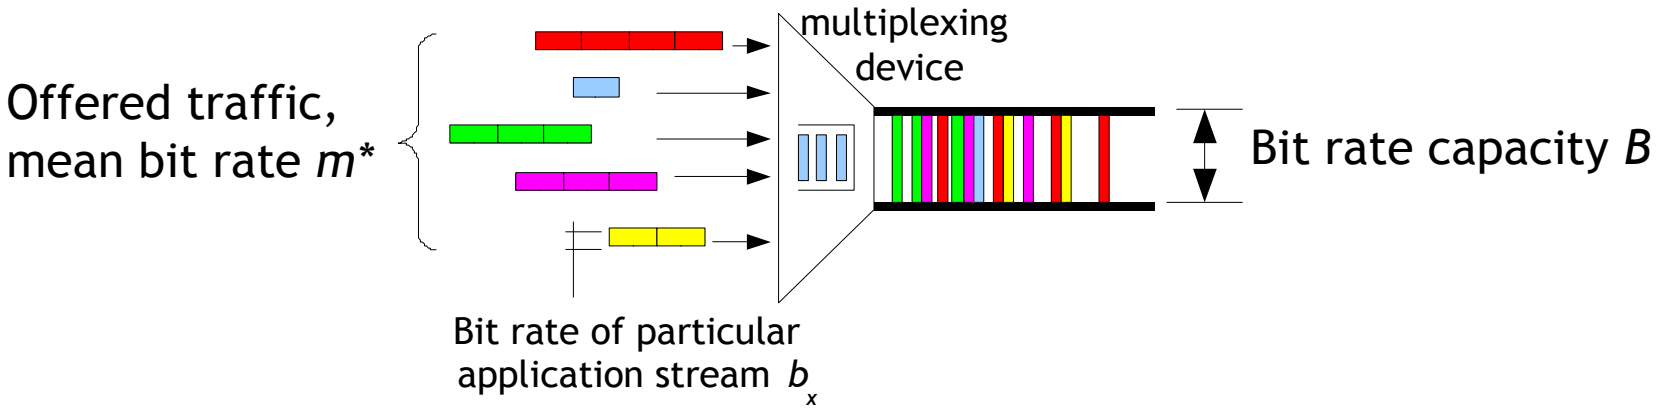

Fluid flow interpretation

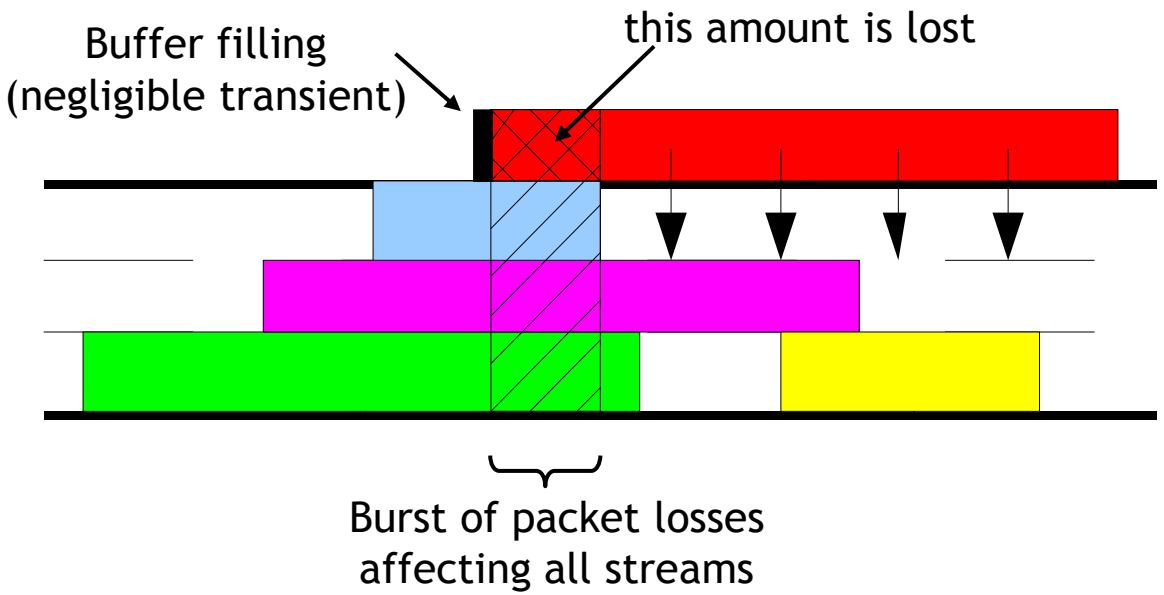

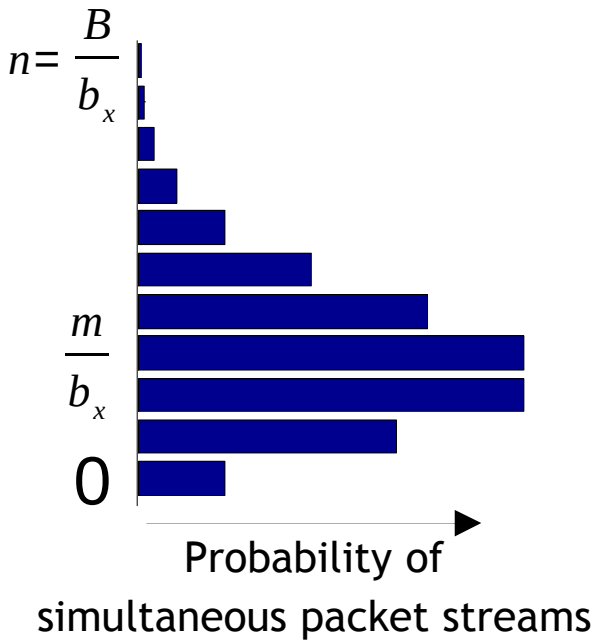

Numeric solution examples

$$P_{loss} = f_{B,b_x}(m)$$

with preset parameter:
B/b_x – speed up from access to core
(possible number of simultaneous streams)

$$B = f_{P_{loss},b_x}(m)$$

with preset parameters:
P_{loss} – permitted loss probability
b_x – installed access link capacity
= maximum application stream bit rate

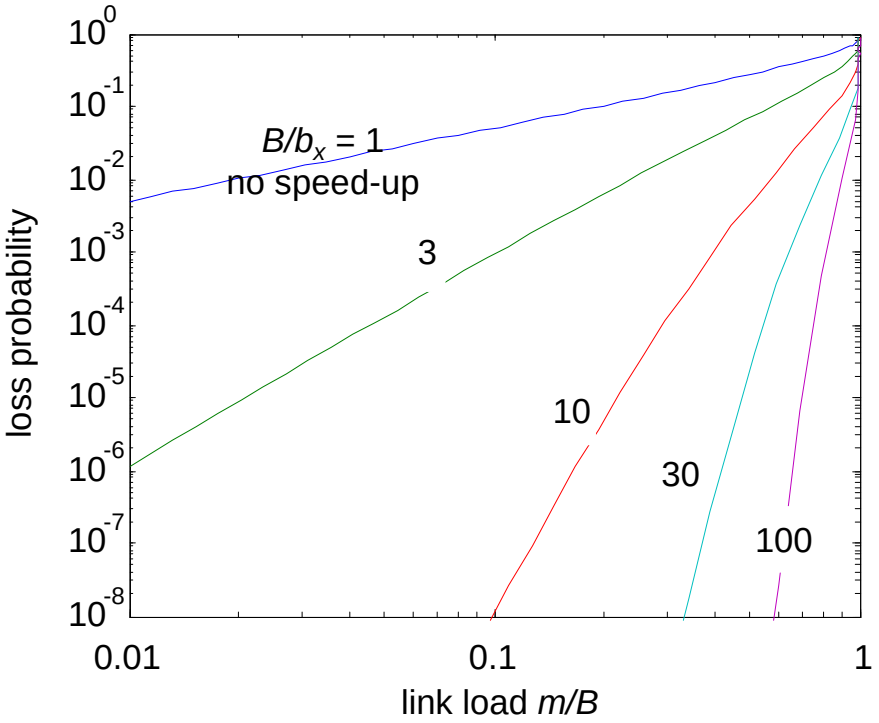

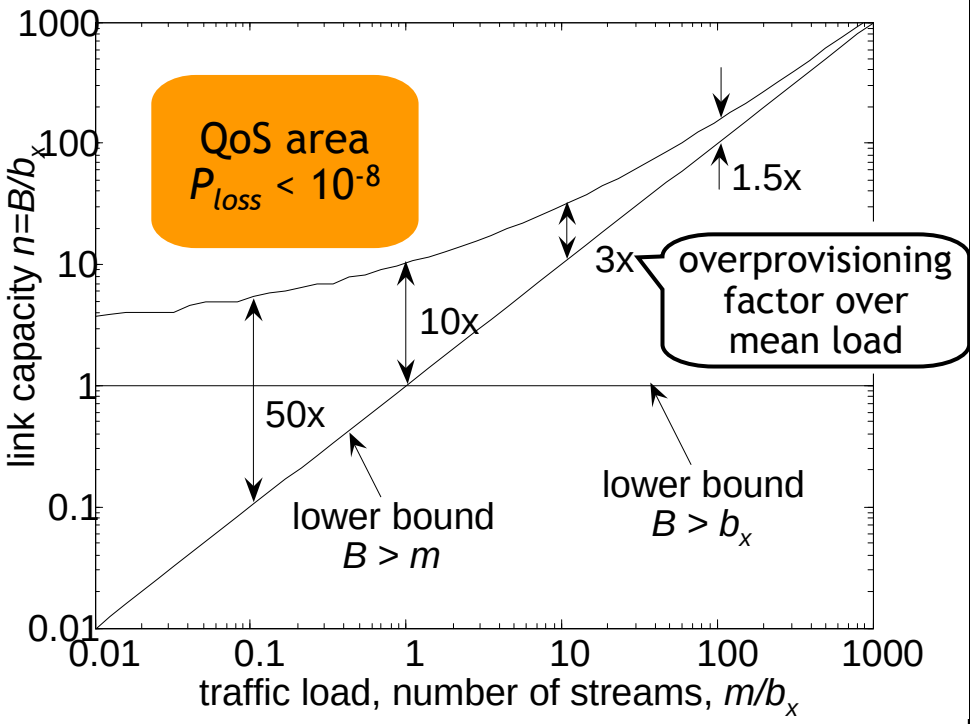

3 Experiments Findings

Experimental verification

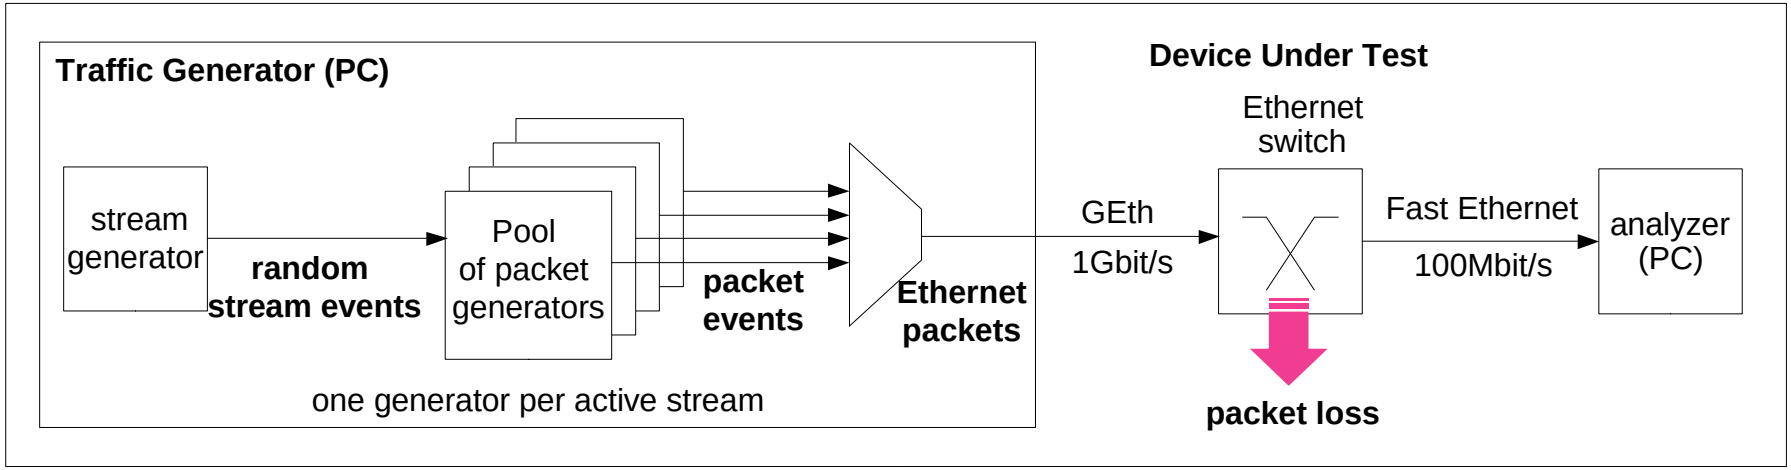

access link capacity	random stream duration		
	min	average	max
3Mbit/s	8ms	4s	1min 30s
100Mbit/s	0.24ms	120ms	2.6s
File size	3kB	1.5MB	33MB

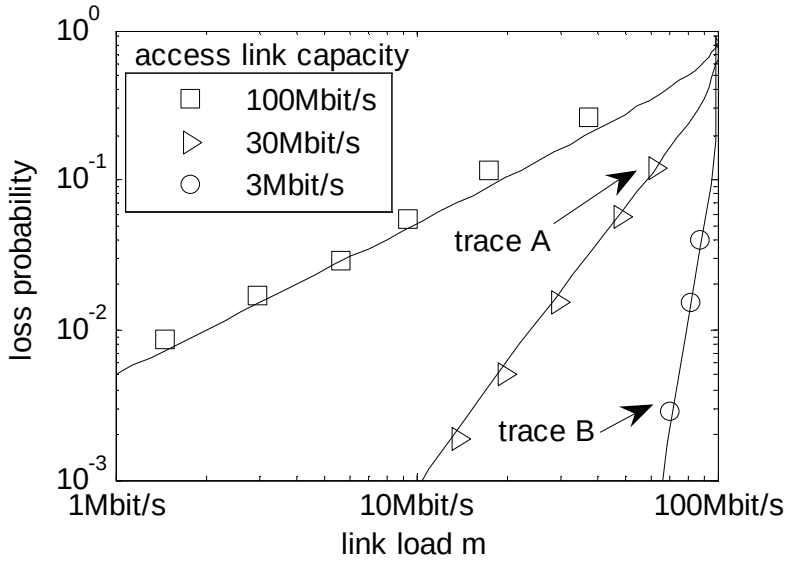

Impact of stream length
Comparison to pure Poisson traffic

- Experiment as before, but dedicated stream length
  - speed-up B/b_x = 30
  - link load = {0.7, 0.8, 0.9}
- Extended by theoretical values in range of undetectable low loss rates

➔ Most dramatic impact in the range of only a few packets.

➔ Nearly constant at longer streams

➔ The crucial question is:
  - Could real traffic contain a weighty fraction of single packet streams?

➔ What is obviously not a problem:
  - Results are insensible to longer streams.

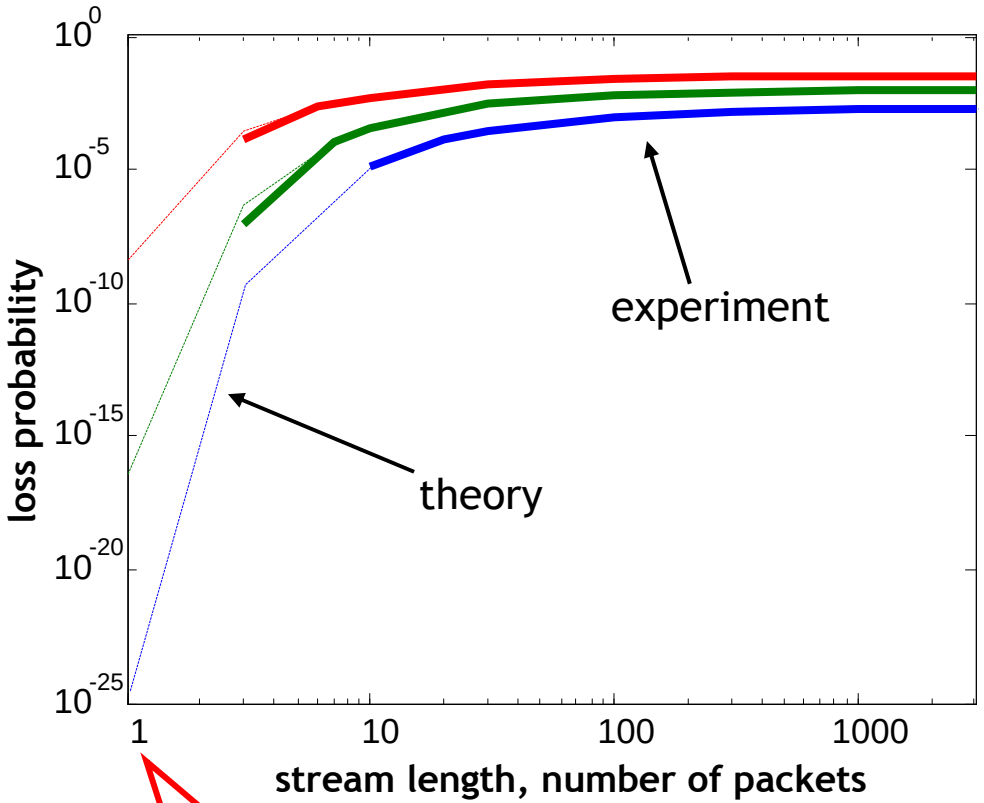

stream length = 1
➔ Poisson traffic

Traffic traces at minute scale hide the real fluctuations …

This is what really happens

What is visible in minute intervals

30Mbit/s access
13% loss

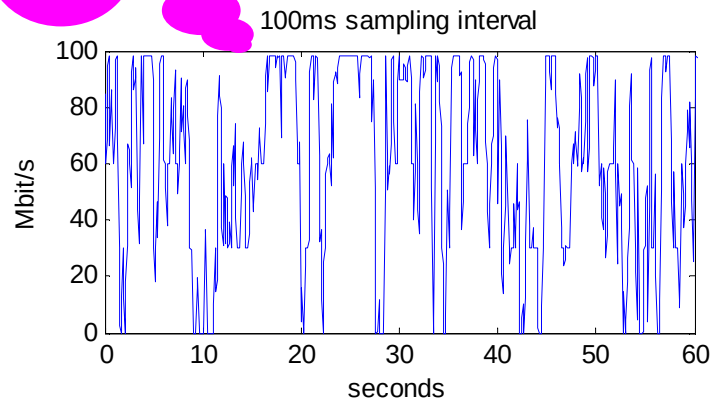

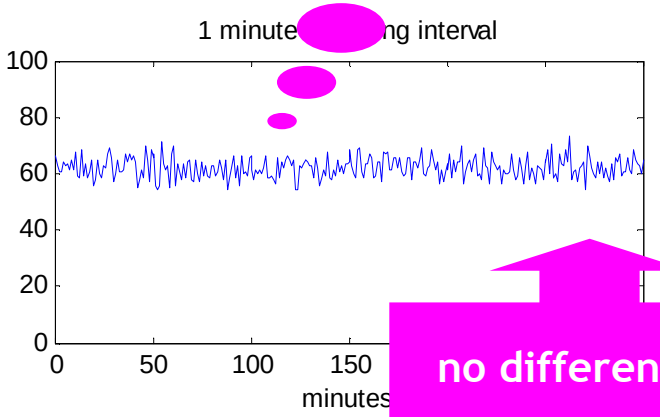

no difference?!

3Mbit/s access
0.3% loss

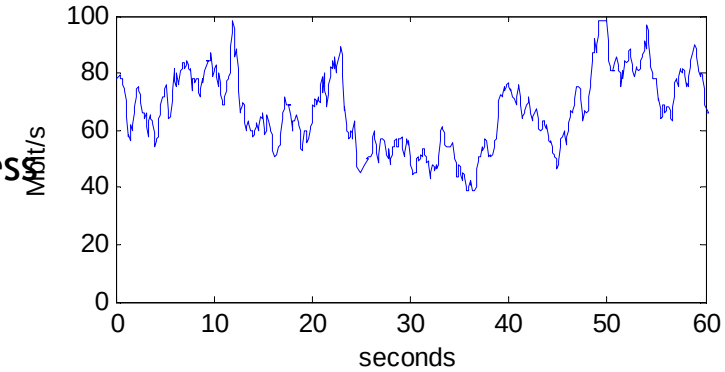

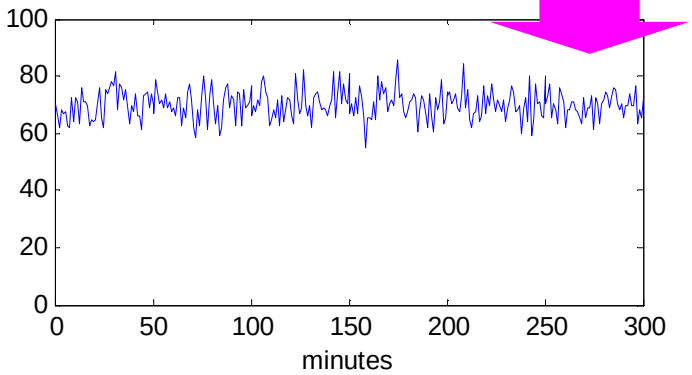

… why?

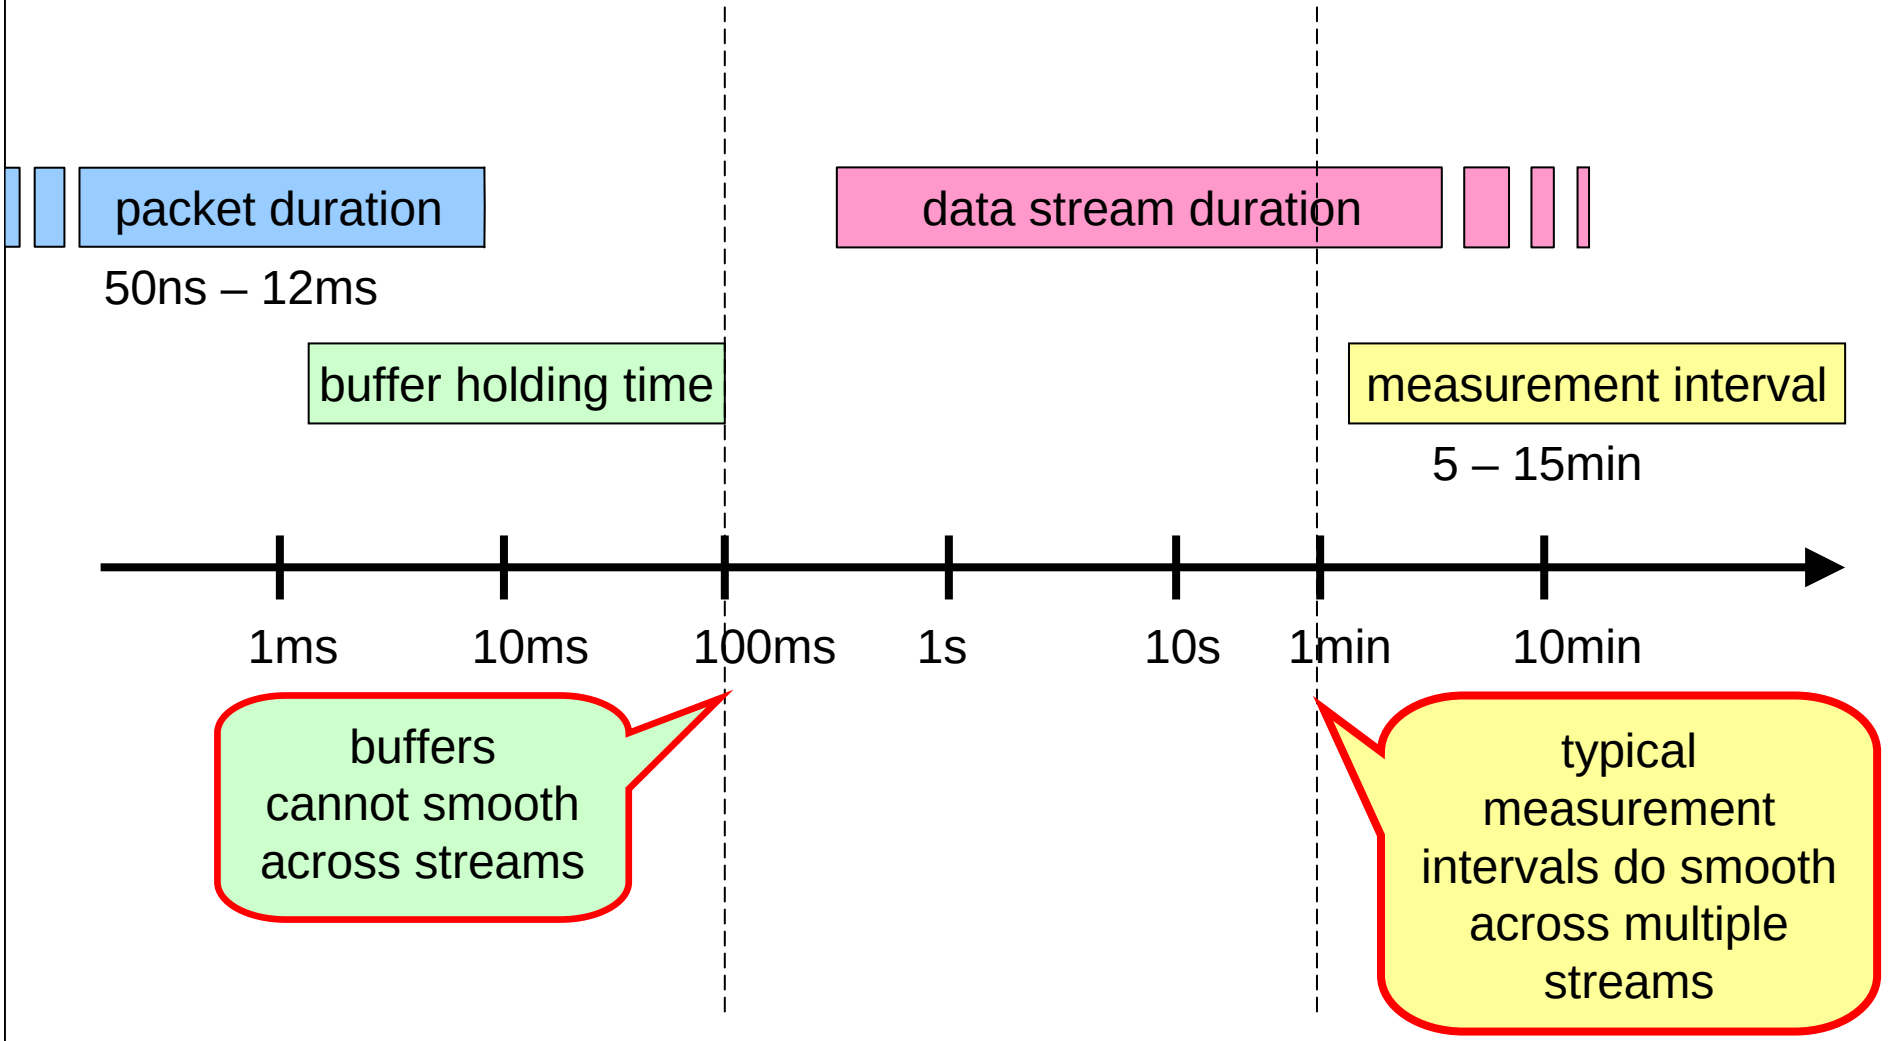

Error bursts
Long observation times required

➔ Finding: Losses are concentrated in bursts

- The overflow event is a collision of streams by its nature
- Burst size is in range of stream size

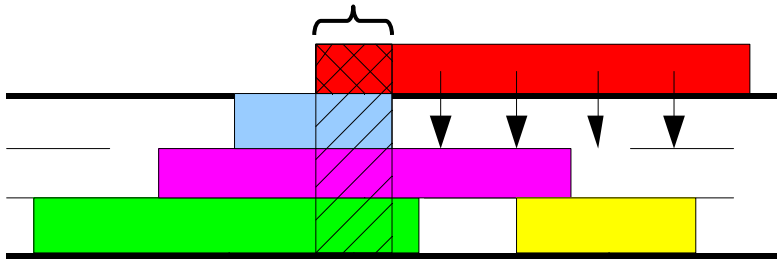

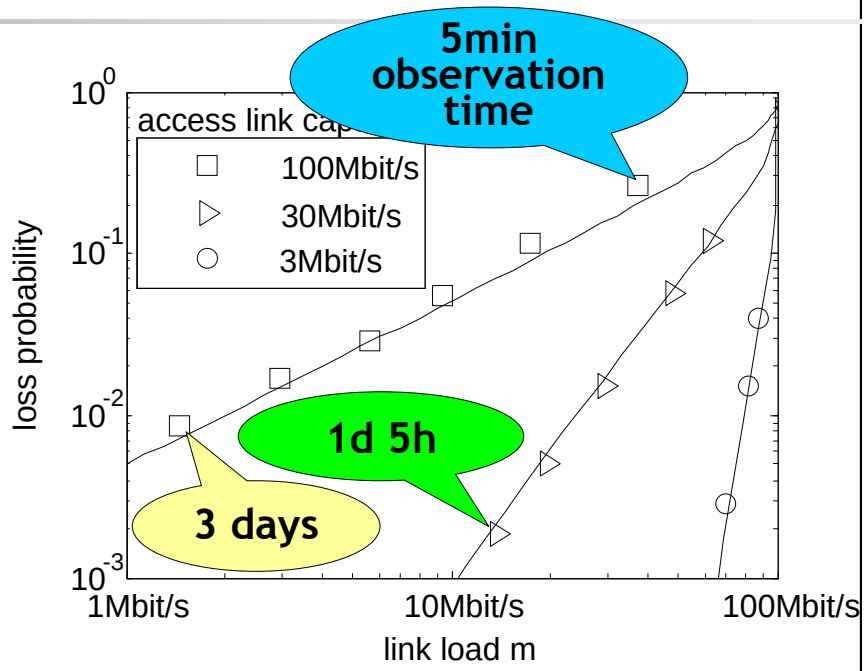

➔ Consequences:

- ➔ Statistical acquisition of loss rate requires huge samples / long observation time
- ➔ Customer experience might be different to what is commonly expected:
  - $P_{loss}=10^{-3}$ does not mean “one packet out of 1000 might get lost”,
  - but rather “one failing transmission within a week of network usage”.
- ➔ Calculation of TCP performance: Assumption of steady loss rate is questionable!

4 Conclusions

Conclusions

- Congestion on aggregated packet links is a collision of application streams by its nature.
  - Inherently too long for buffers,
  - Stream collisions are causing bursts of packet losses interleaved with lossless periods.
- Visible fluctuations in traffic traces at minute intervals are not telling anything about the real losses.
  - Extrapolations with some arbitrary safety margin are close to irrelevant.
- Fluctuations on packet links are dominated by the installed end user access link capacities.
  - The decisive parameter for link performance is its speed-up wrt. the access links.
  - Increasing end user access capacity boosts the fluctuations, even at unchanged mean load
  - Operators can control the congestion probability by measuring the mean load combined with the knowledge of the access link capacities.

Conclusions

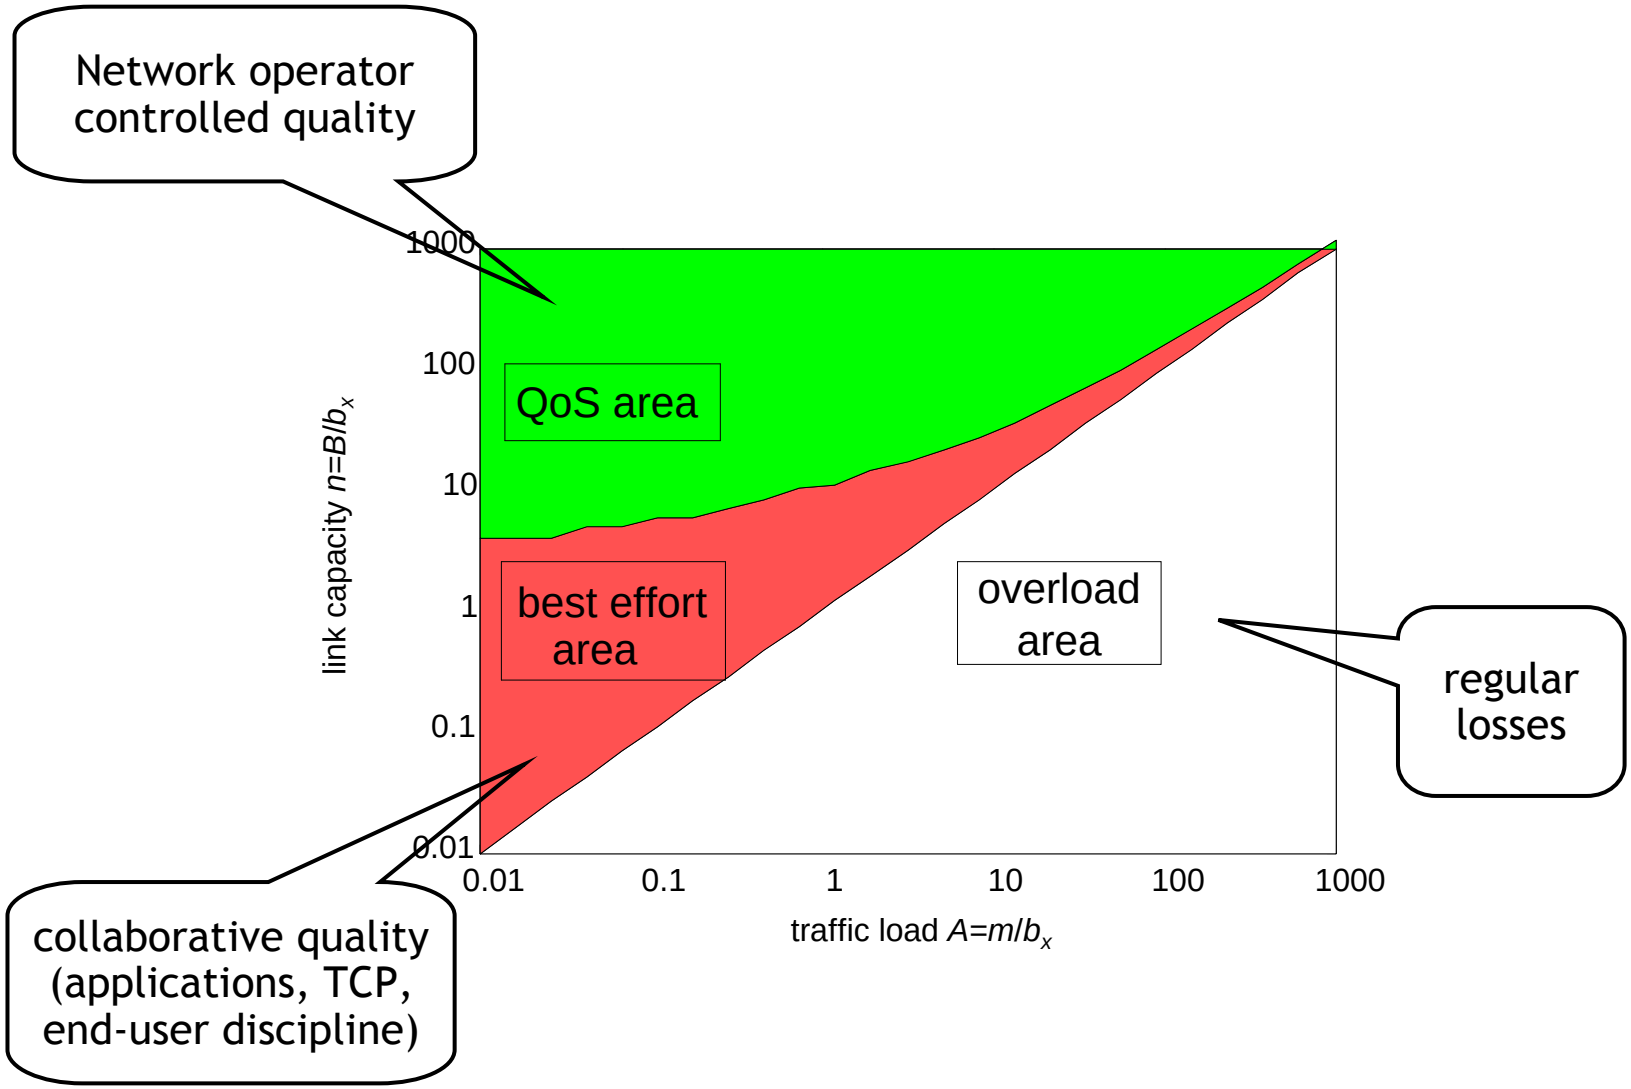

www.alcatel-lucent.com