Measurement of packet networks, e.g. the internet

John Schormans (EE)

Ben Parker (SMS) (*next speaker in this joint talk*)
and
Steven Gilmour (SMS Head of the Statistics Research Group and Director for the QMUL Centre for Statistics)

All-packet networks …

- … are now standard.
- Regardless of information type (voice, video, web-surfing, TV on demand etc) all the data is aggregated into packets:

and transmitted across the network as a series of discrete chunks of data.

- Guarantees are provided on the reliable and timely delivery of your packets.

Example - Sprint SLA

http://www.sprintworldwide.com/english/solutions/sla/

(Packet delay)

Committed Metrics

	Network Latency	Network Packet Loss
United States	55 ms	0.30%
Trans-Atlantic	95 ms	0.30%
Intra-Europe	45 ms	0.30%
United States to Japan	130 ms	0.30%
Hawaii to Continental United States	85 ms	0.30%
United States to Australia	210 ms	0.30%
United States to Hong Kong	190 ms	0.30%
United States to Singapore	300 ms	0.30%
India to North America	350 ms	0.70%

NTT IP networking SLA (http://us.ntt.net/support/sla/network/) last visited 29/12/2007

- “ We guarantee:
- The NTT Com Global IP Network is free of Network Outages 100% of the time.
- Latency will not exceed 50 milliseconds for the North American Network, 90 milliseconds for the Trans-Atlantic Network and 130 milliseconds or less for the Trans-Pacific Network.
- Packet Loss will not exceed 0.1%.
- Average jitter will not exceed 250 microseconds and will not exceed 10 milliseconds more than 0.1% of the time. “

Sampling by probing for delay and loss

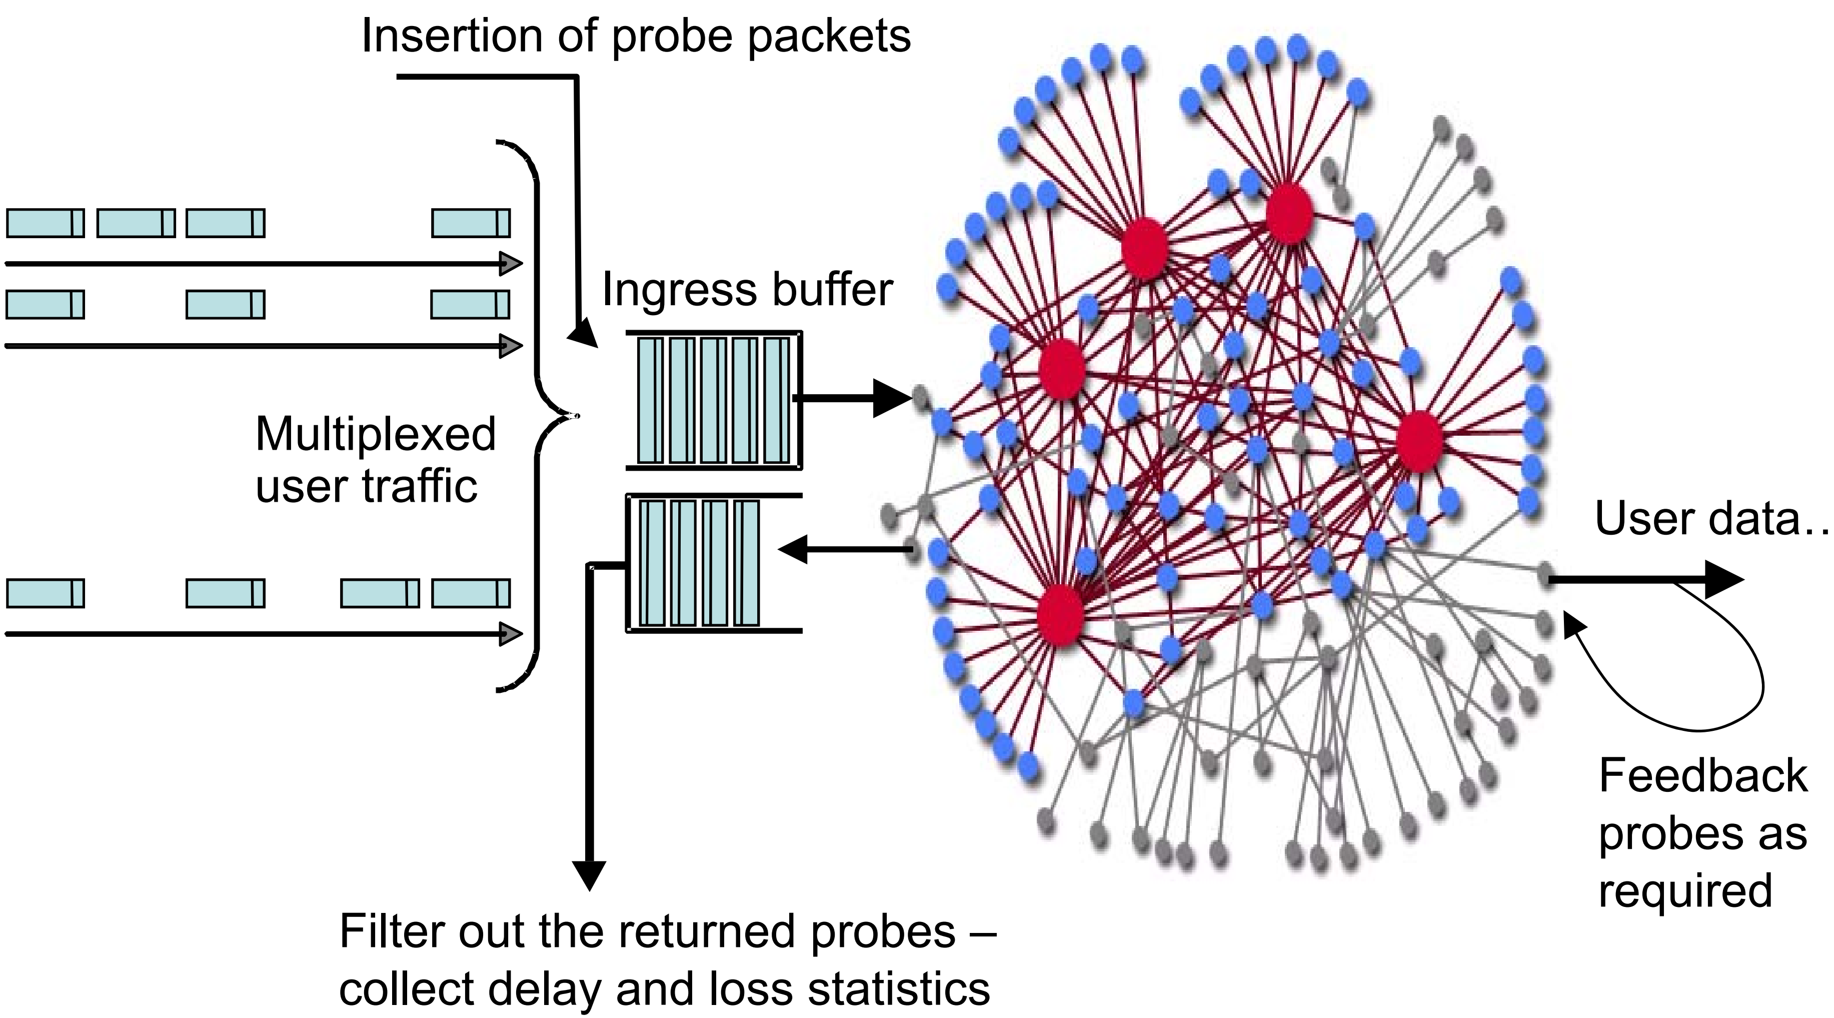

Original (interdisciplinary) problem

- … *How do we predict the accuracy available from measurements when we sample packet level performance using probe packets?*
- I had already done some early work on this (+Tijana Timotijevic when at VF, and Jonathan Pitts); this suggested that variances could be large, and so also sampling errors
- QMUL S&E has funded an interdisciplinary PhD for Ben with Steve as 1st supervisor and me as 2nd supervisor
- Ben's 3 years is over this Autumn; he's just presented his thesis abstract as a "Best Abstract" at this year's SIGCOMM in Washington.
- Also a provisional "accept" on a full IET journal paper that reports some of the work we discuss here.

Causes of packet delays

1. Packetisation – how long to fill the packet. Fixed for a link
2. Transmission – how long to put the bits out on the link (function of bitrate and packet size in bits). Fixed for a link
3. Propagation – function of the speed of signal transmission (often light, e.g. in a fibre optic cable). Mostly fixed (satellite / wireless complicates matters)

Relatively easy to measure

4. Queueing – function of queue length at packet arrival, and the queue scheduler in use. Variable, often with large variance, particularly at high load

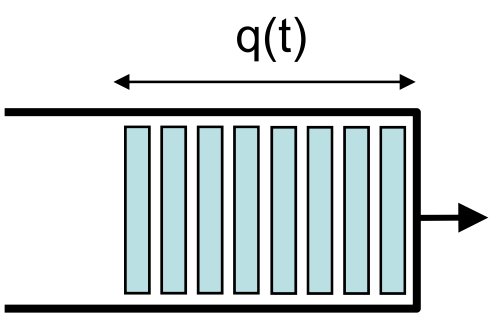

Variance also increases with the number of buffers in the end to end path

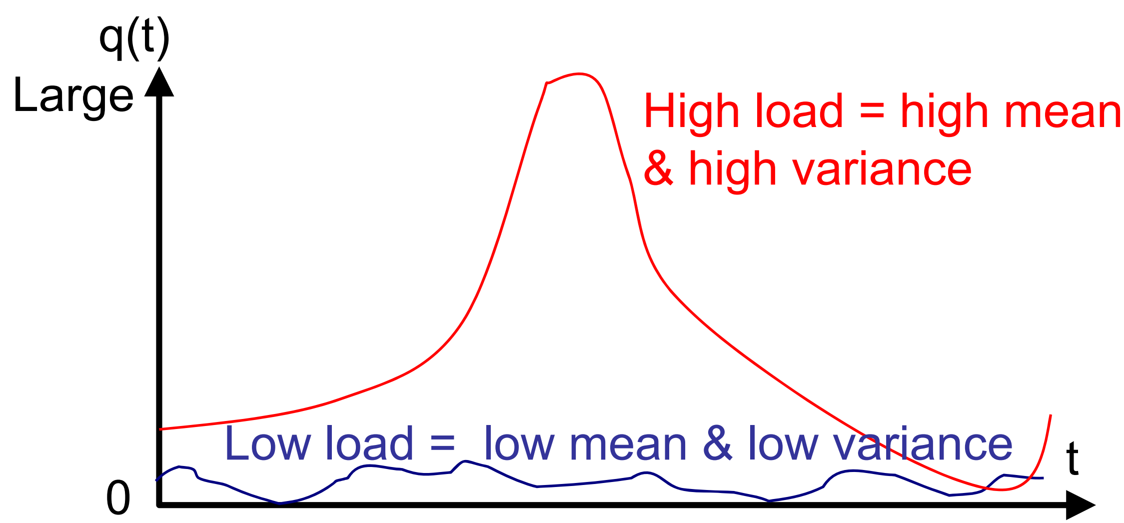

Example form an earlier student's work - active probe = 1/1000, 10 hops E2E

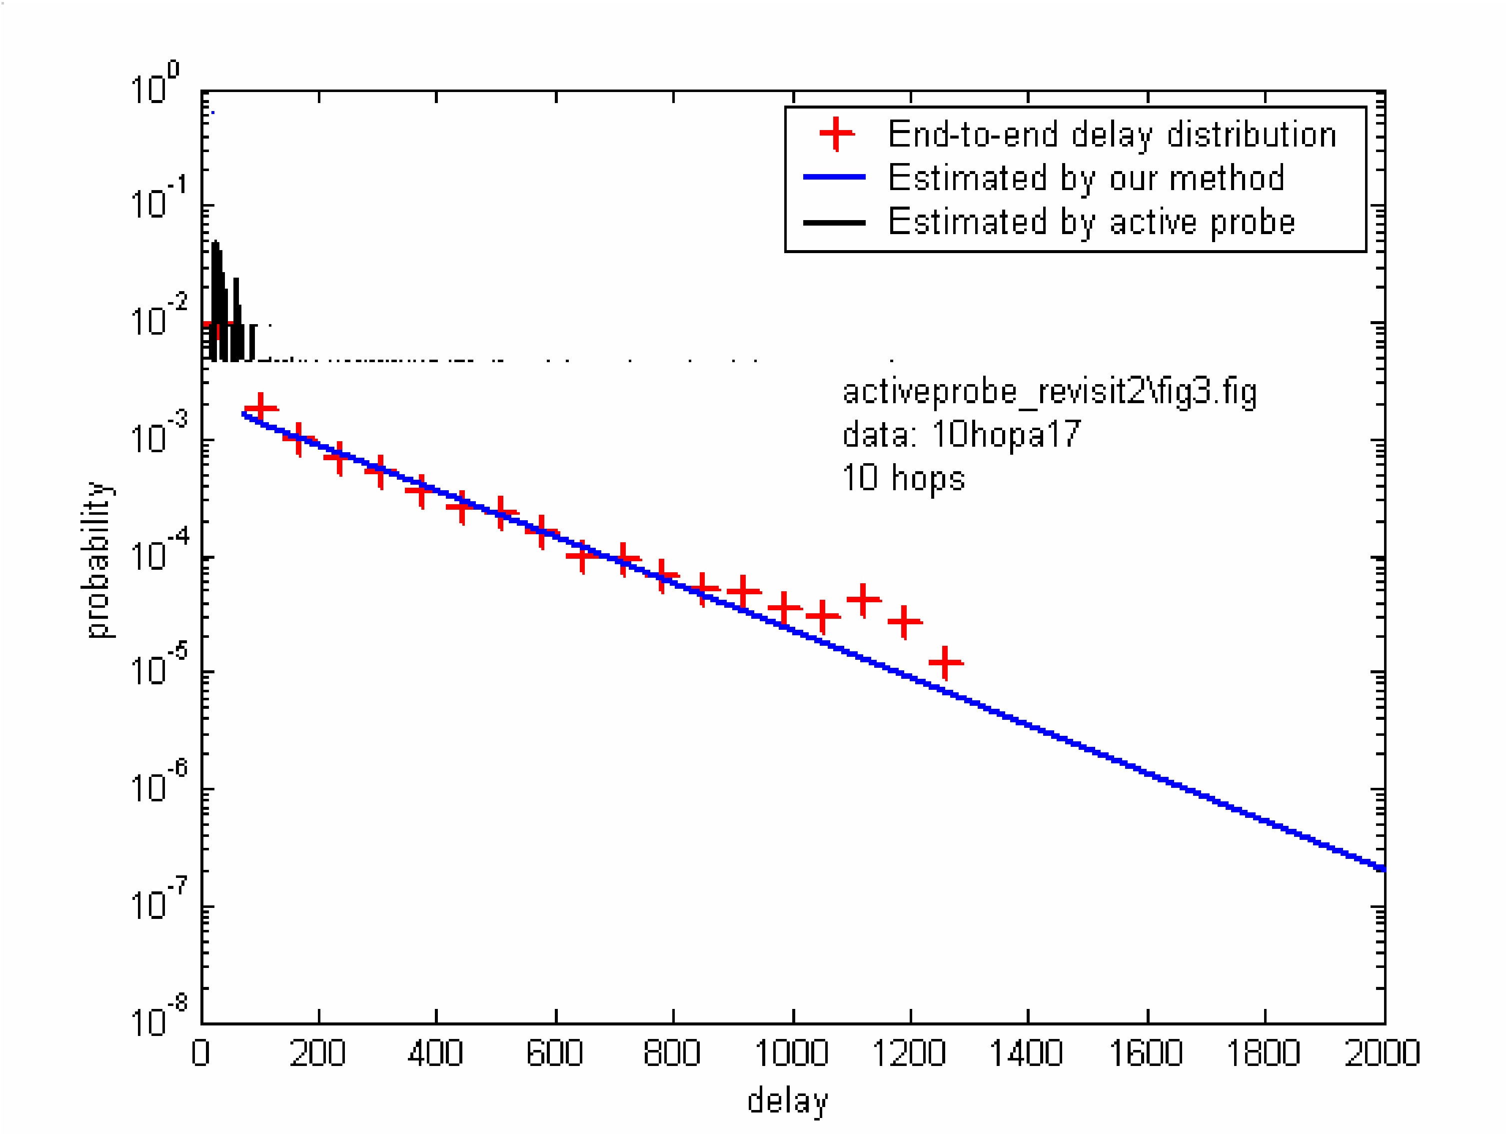

Causes of packet loss

1. Loss due to buffer overflow

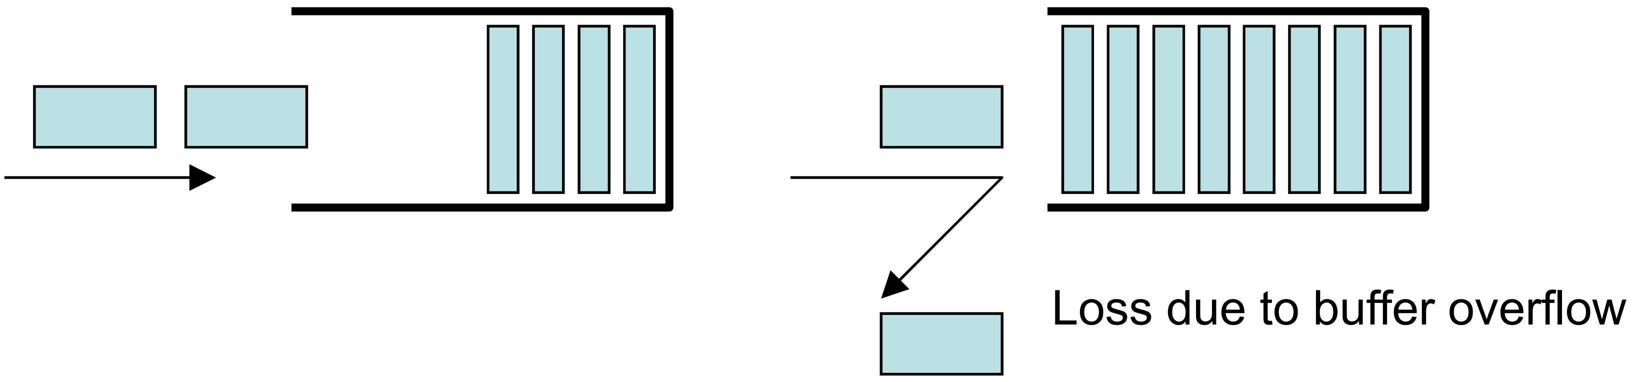

2. Loss due to broken link

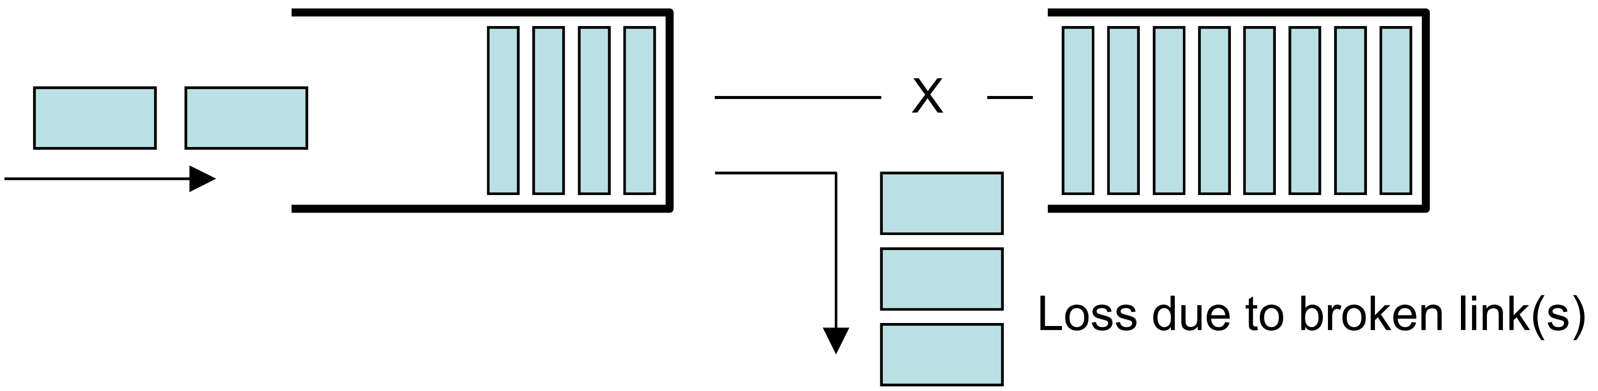

Either way, measurement is by noting the number of lost probe packets. Complicated by the fact that losses tend to be bursty and not simply randomly distributed.

Probing for loss probability

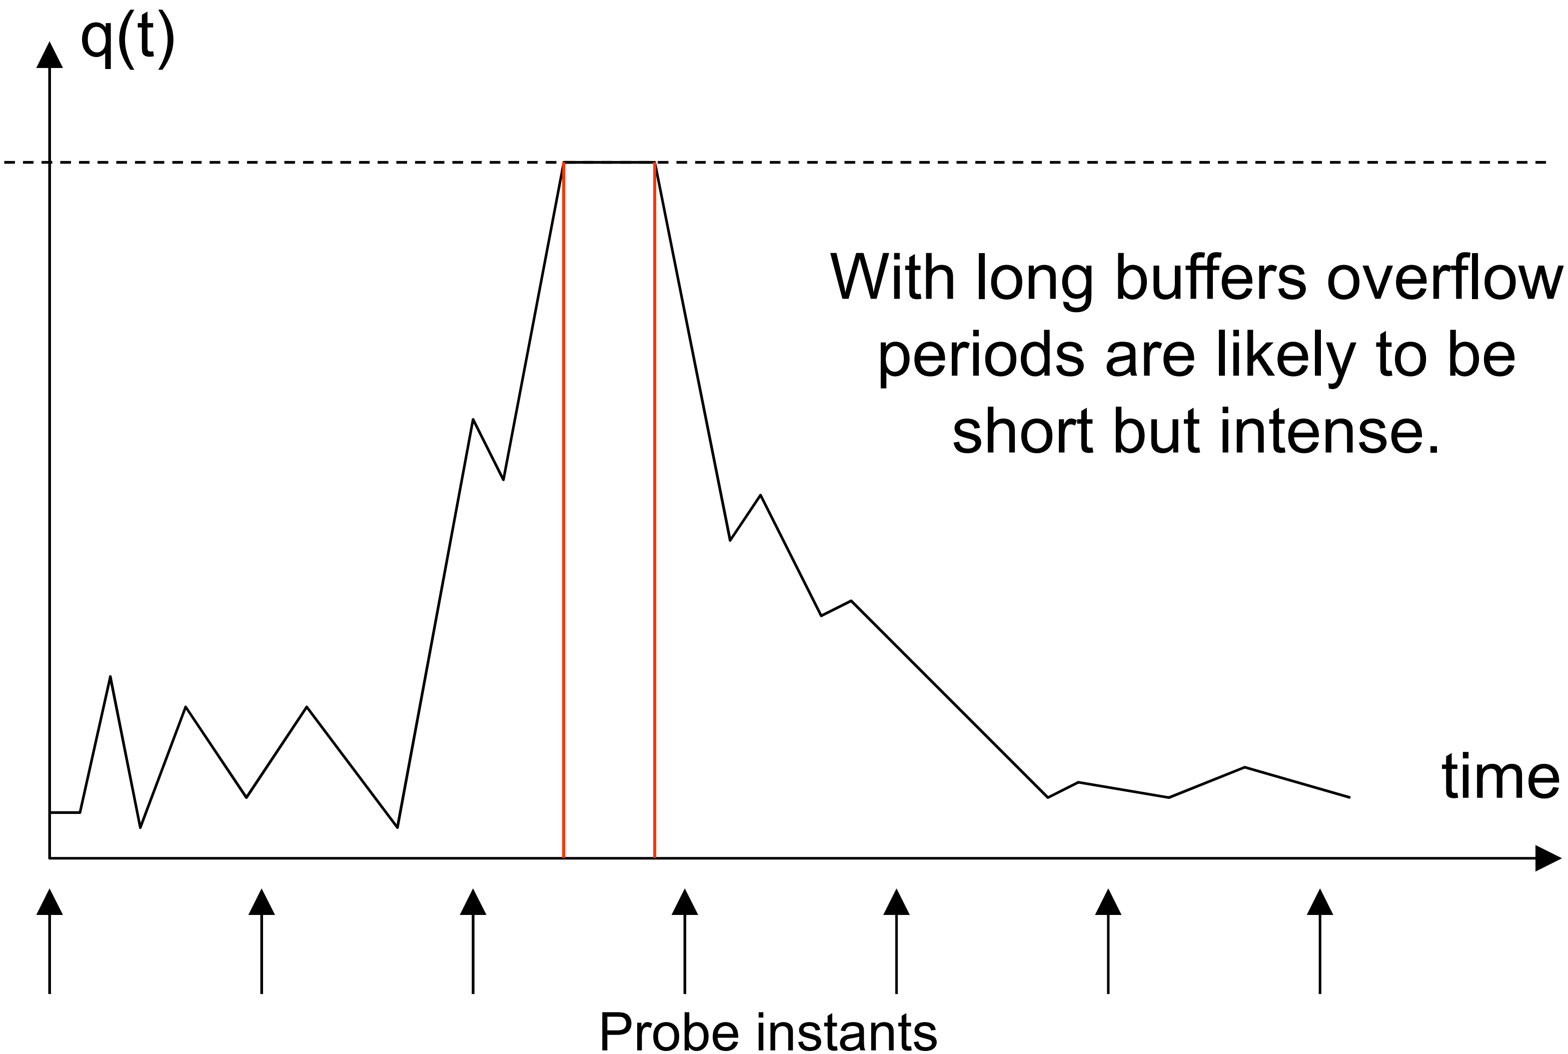

What affects the loss process due to overflow in e.g. a single buffer?

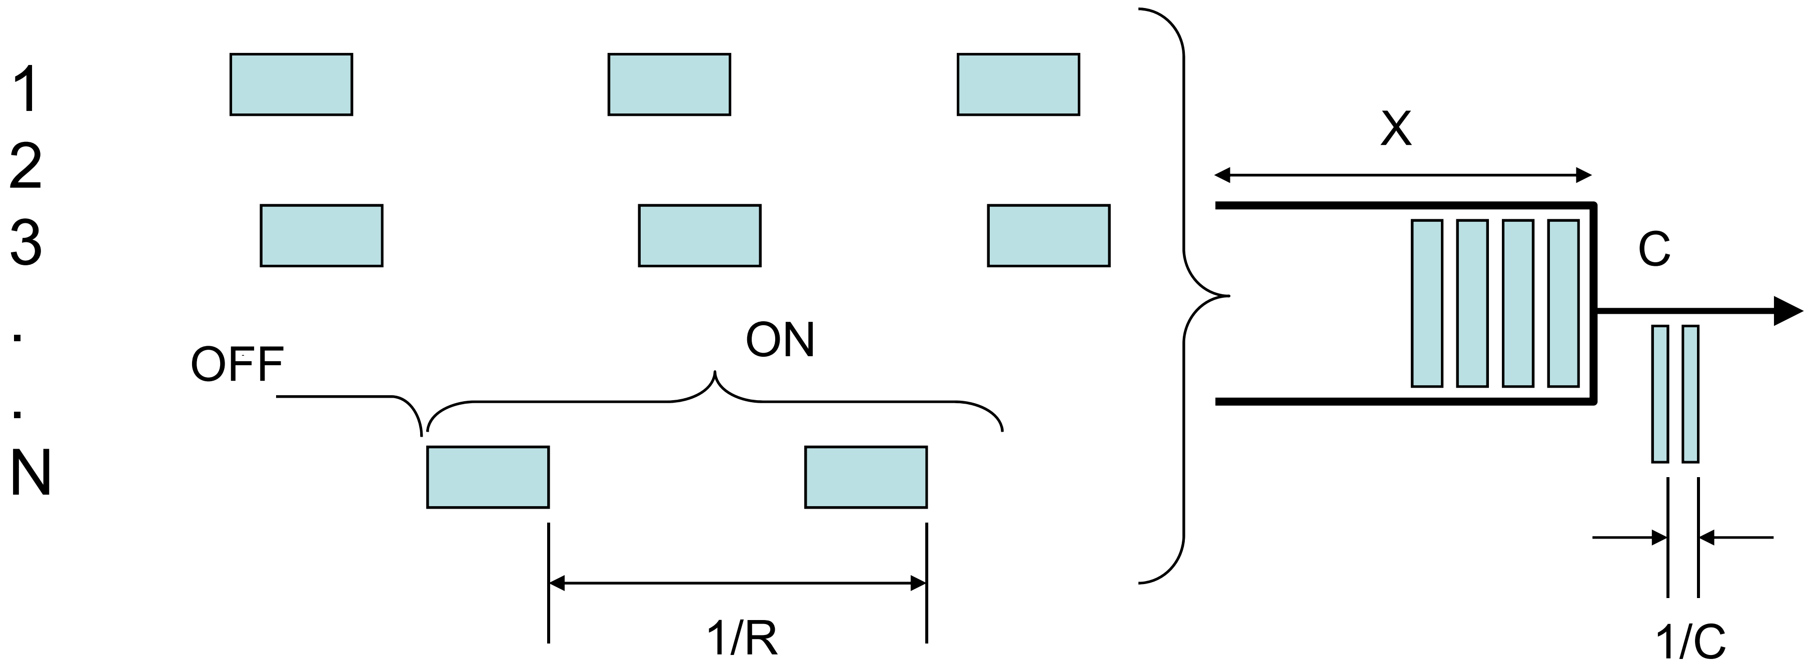

-The target loss probability, e.g. 10^-3, 10^-4 … the lower the loss probability the more samples required to measure it

-The buffer length, X … the longer the buffer the more widely spaced the overflow periods and the harder they are to measure accurately

-Ratio R/C … as R/C→1 the more samples needed

-The number of buffers in the end-to-end path, more is harder!

Three aspects of probing error

- Inject too few probes and the sampling error will be too large
- Inject too many and you may (& probably will) adversely affect the performance you are trying to measure (some people call this the "Heisenberg effect")
- Additionally – and this is less intuitive – there is an effect due to correlation between samples in measuring e.g. waiting times in queueing systems (if the wait for packet 'j' is long, then it is much more likely that the wait for packet 'j+1' will be long).
- Engineers I believe would not normally think of this last one (certainly Jonathan and I didn't, and we didn't see it anywhere else either)… it was pointed out by a mathematician – Matt Roughan – who was working at AT&T "Bell" Labs, and presented this at Intel Lab at Cambridge Uni

Can see probing as a "numerical experiment" … and experiments can be optimally designed via <u>statistical DOE</u>

- What probing pattern should we use – Poisson, deterministic, uniform (over a short period), or as a gamma process?
- What rate should we use? Cisco equipment always had default rates of e.g. 1 probe / sec, 1 probe / min …
- Should probes be single or bunched (I have known a large telco that was going to adopt a batch-probe method; also probing for "available bandwidth" tends to assume probes will be launched virtually back-to-back
- Can we do >1 metric at the same time? (e.g. probe for loss and delay together …)

QMUL expertise in DOE

- DOE has been widely applied to research in the Biological Sciences and many aspects of Industrial Engineering but it has not yet to networking.
- Prof Steven Gilmour heads The Statistics Group at QM – an international reputation for its work in DOE (large proportion of the UK's experts are based at Queen Mary)
- Recent award of a £470K EPSRC grant (PI: Gilmour), and by the Defra sponsored project
- DOE techniques have been very successfully applied in linear and static environments. Most work on nonlinear systems has so far assumed static processes, or at least deals only with static aspects of the processes.

Ben's talk is about probing loss in an edge router modelled as a single queue

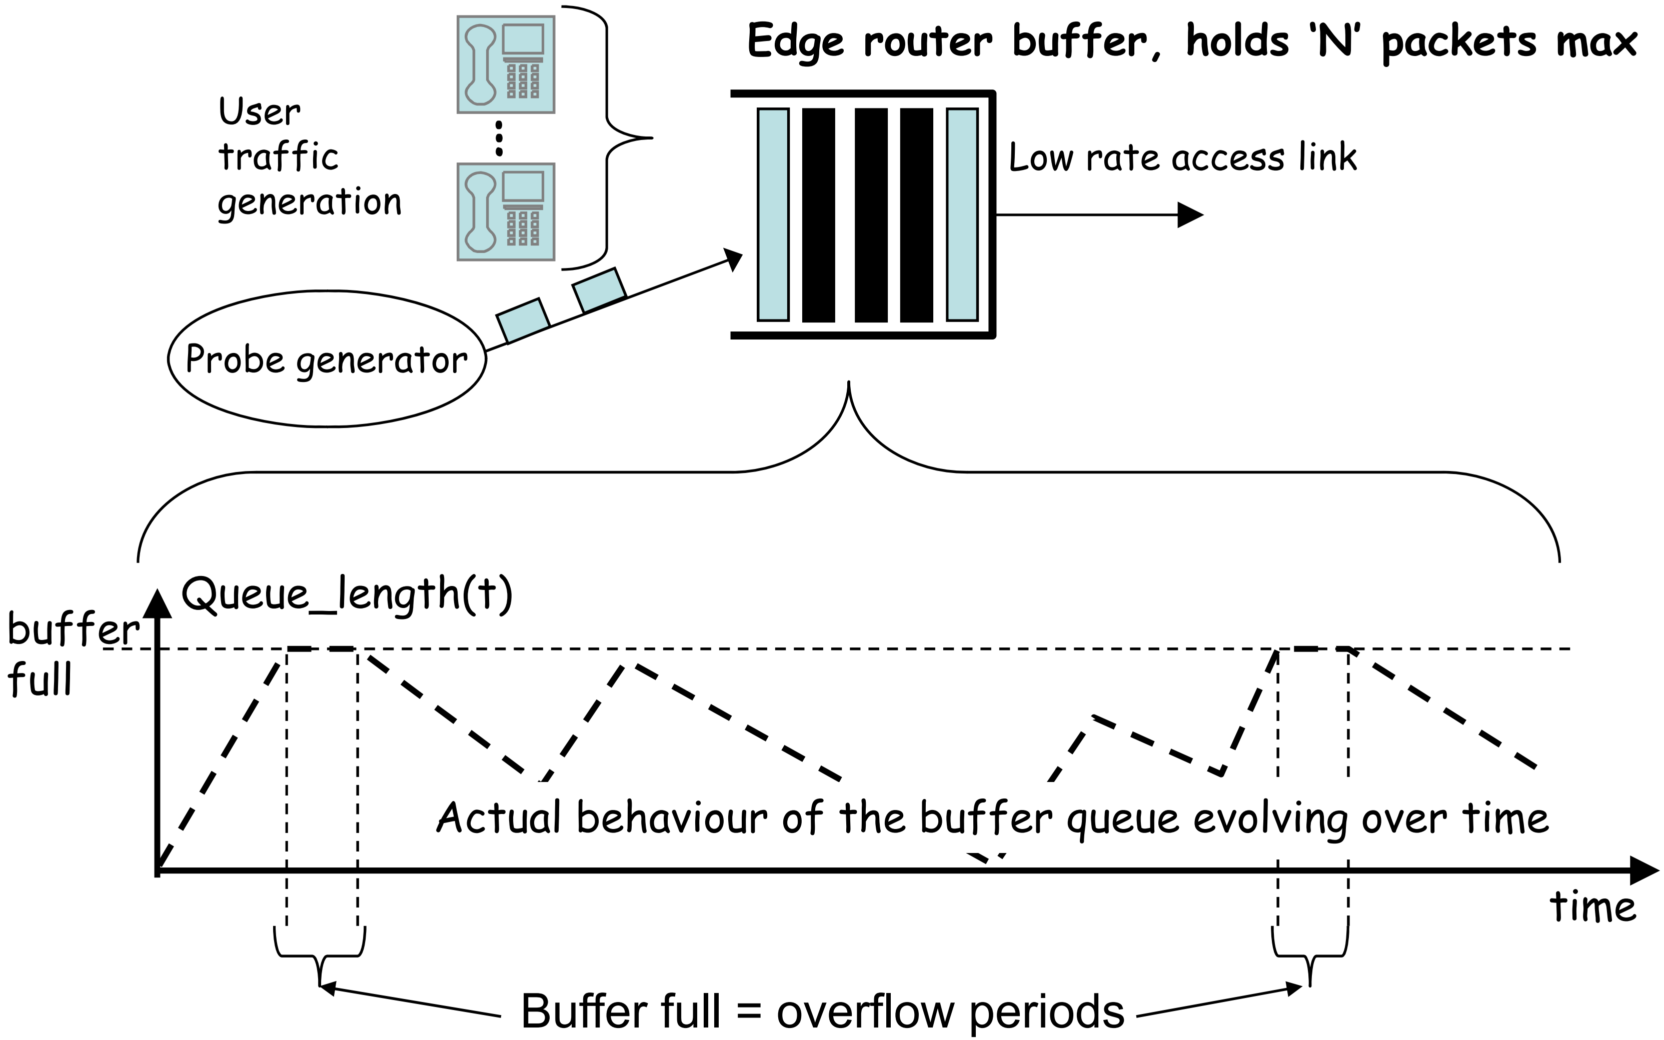

Representation of a packet queue through the full N states – too hard to form into a properly designed experiment (so far…)

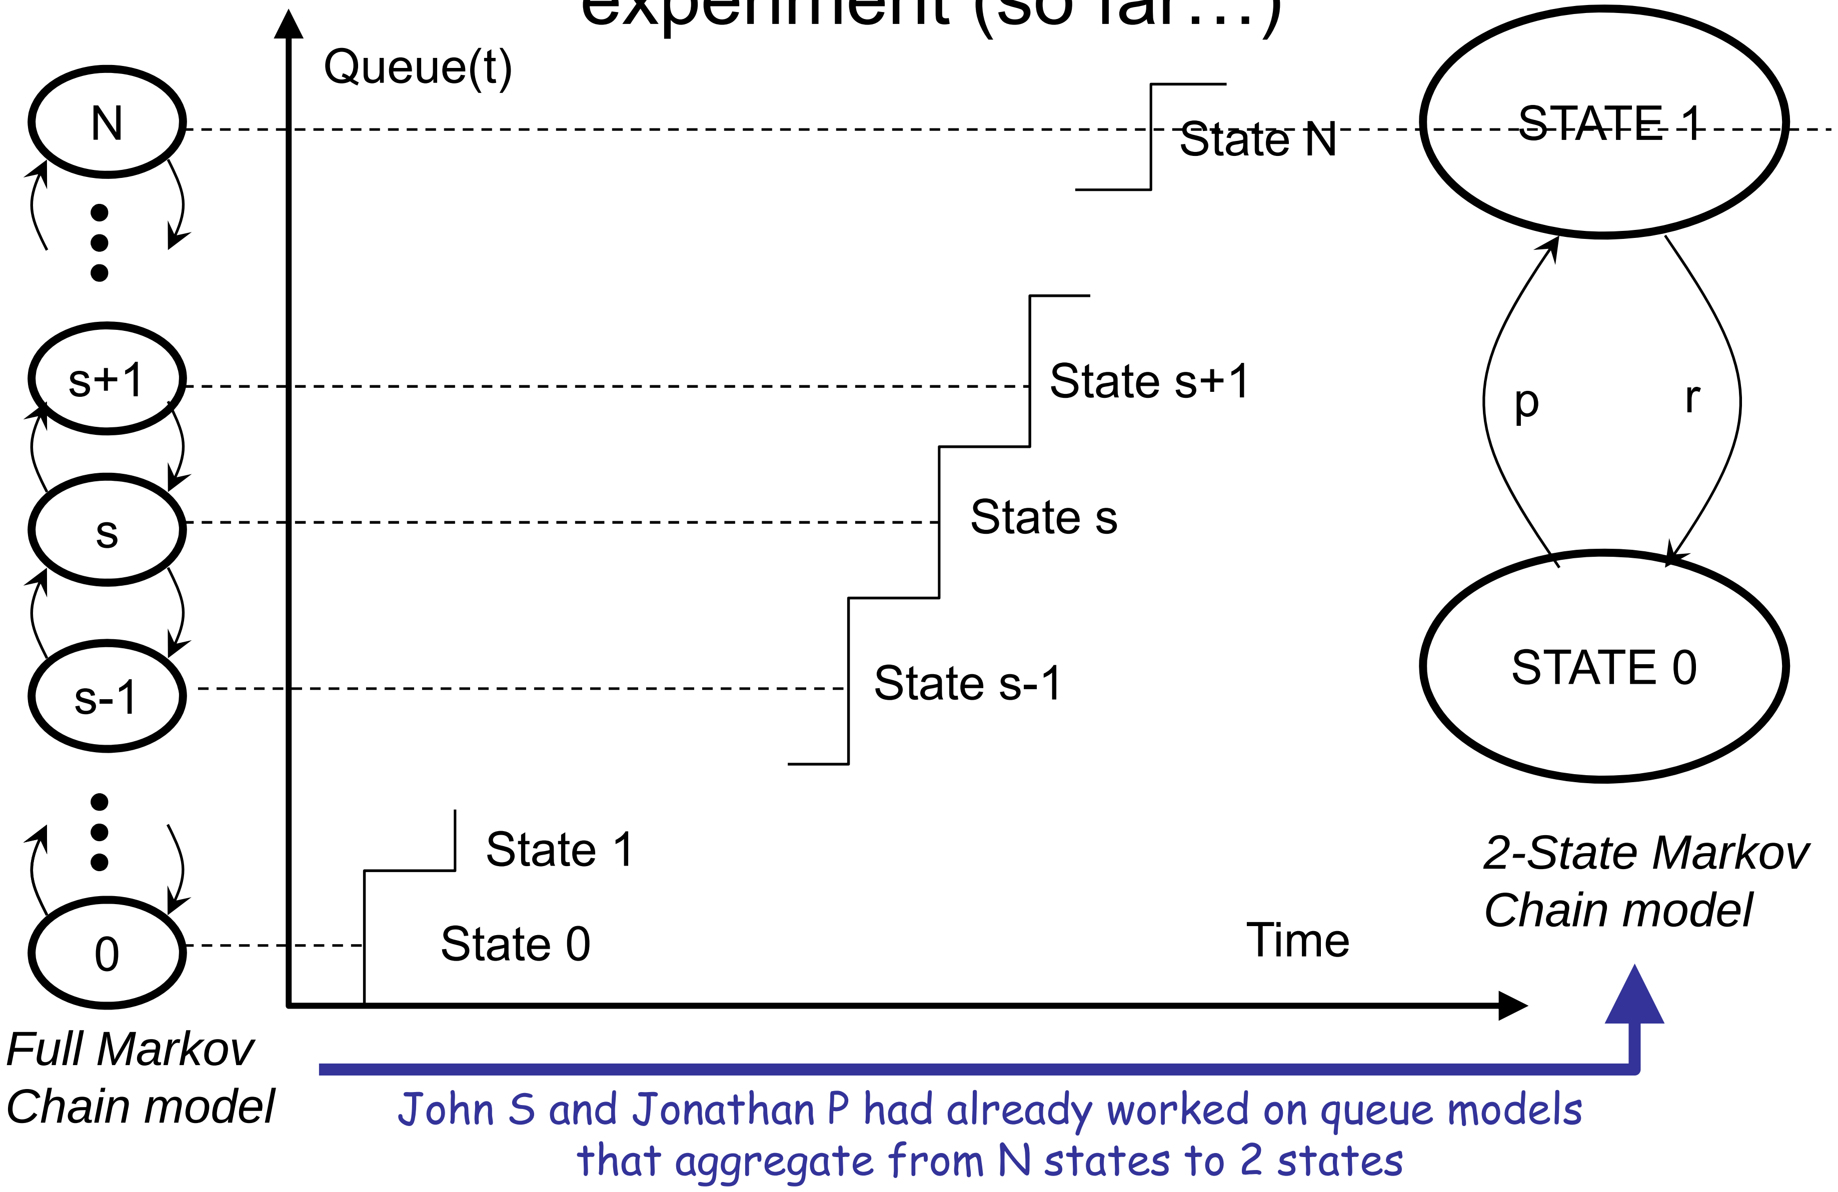

So this can be seen as …

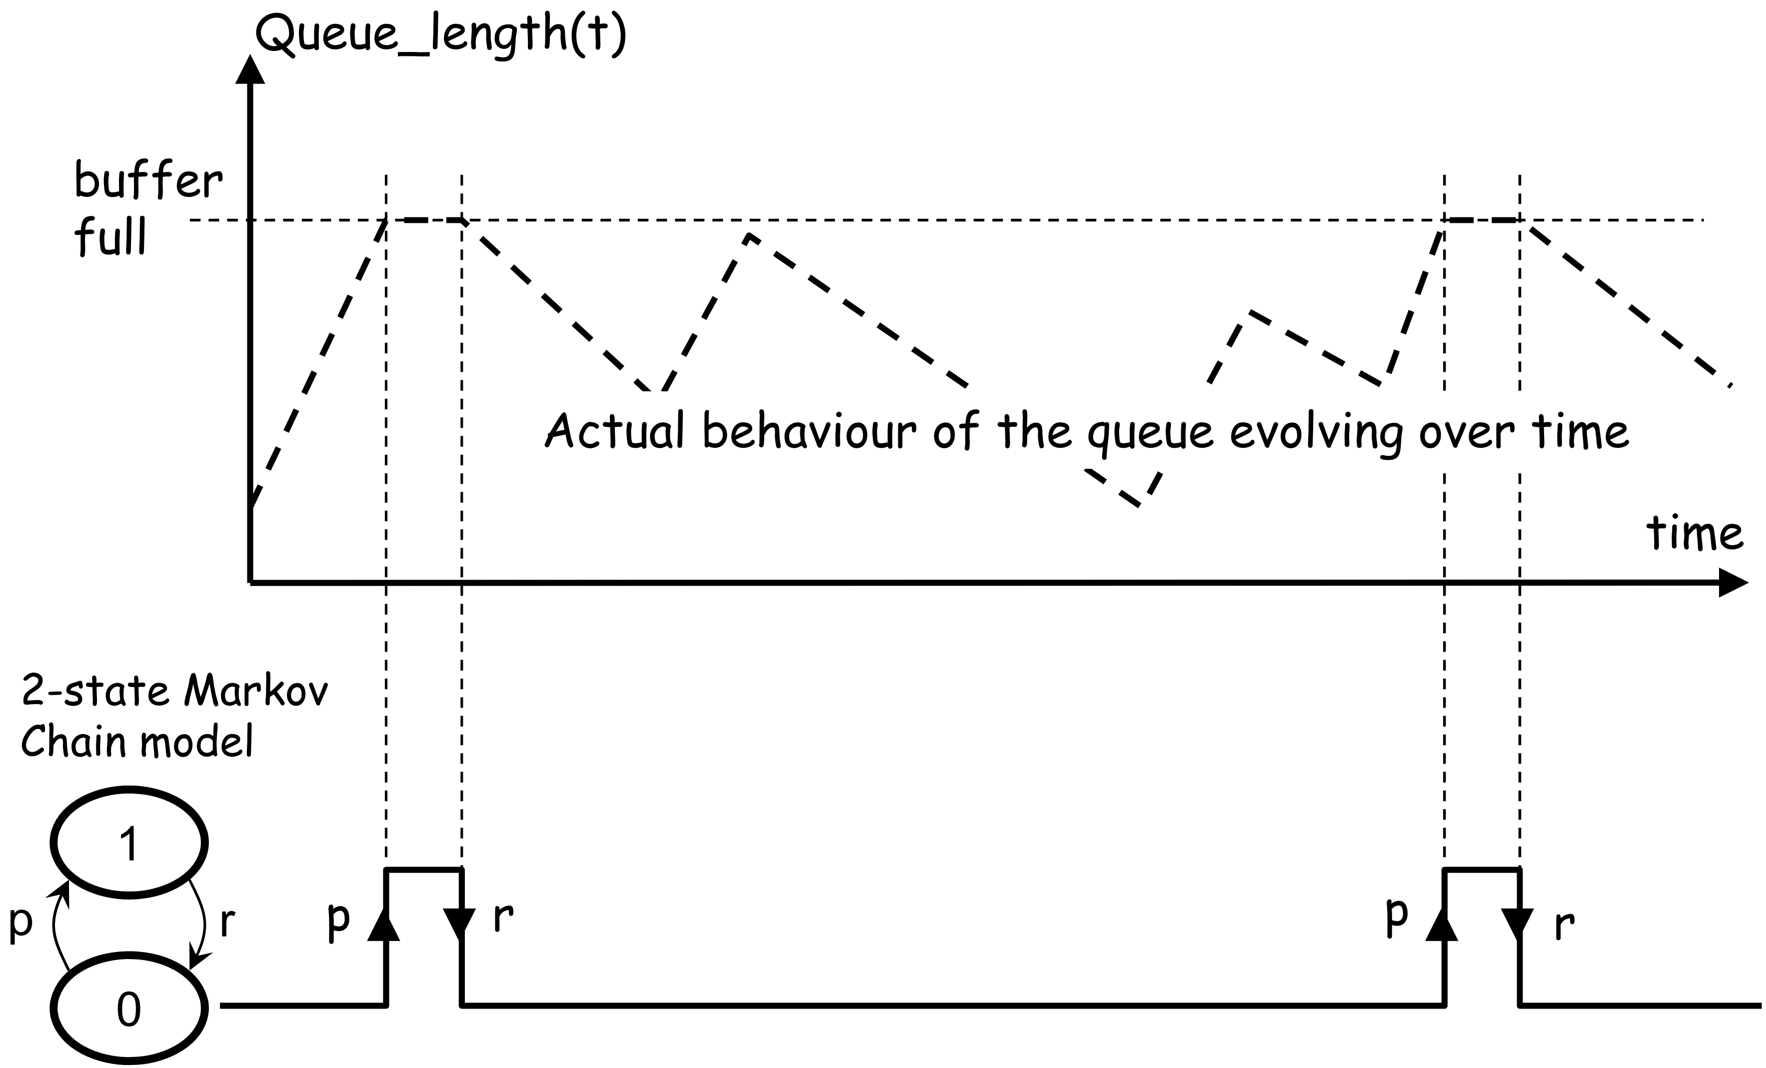

Relationship to probe packets - arrange as a “designable” experiment …

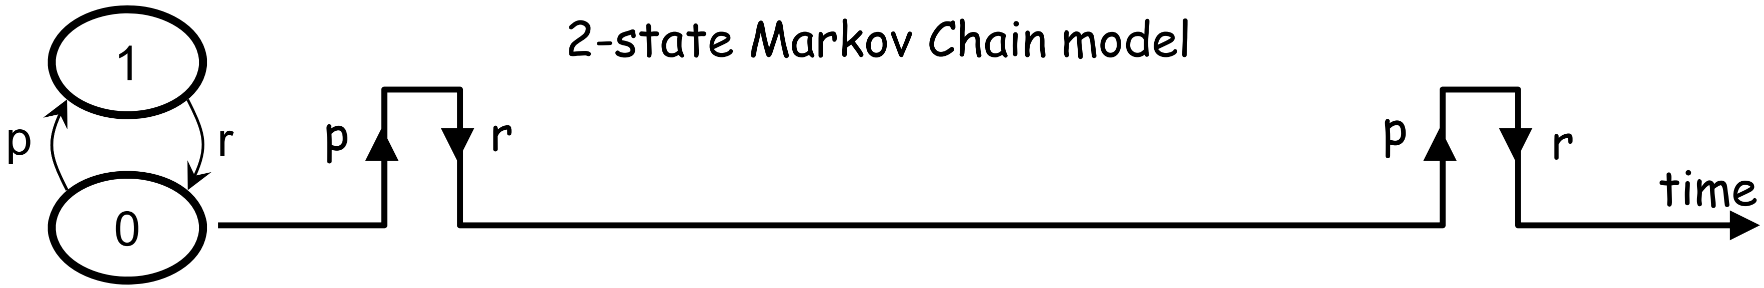

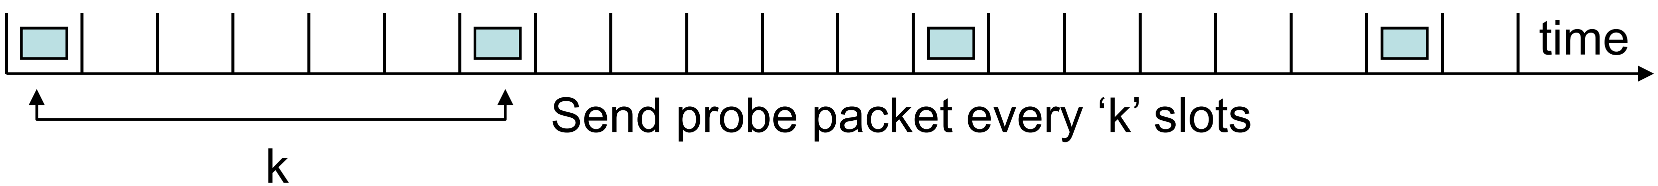

We can then pose the question – *what is the optimal ‘k’ ??*

This has not been addressed previously!