Buffer Sizing: “Layer 8+” Aspects

Simon Leinen, SWITCH simon.leinen@switch.ch

Buffer Sizing workshop, Stanford, 5 March 2019

Economics

- Buffer sizes used for product differentiation
  - Between vendors (our products are better because buffers are deeper)
  - Between product classes (you want to use “carrier-grade” equipment for your important network, not the cheap stuff)

Economics

- What do operators optimize for?
  - Wide variety of goals:
    - High utilization (less than in past? Under what assumptions about traffic load?)
    - ~~Single-stream TCP throughput over long RTTs (I2 Land Speed Record~~) maybe still relevant?
    - Throughput
    - Losslessness
    - 99%ile xyz completion times
    - Median …
    - “Fairness” (but what kind?)
    - …

Risk aversion

- Ops want assurance new recommendations don't make things worse
  - How much of my "too big" buffers am I using today?
  - Only practical way to find out: Artificially limit queue sizes, look for drops
- Need a kind of "MRTG for buffer occupancy"?
  - Simple but *ubiquitous* tool to look at historical buffer utilization
    - All chipsets seem to have *some* support to sample buffer fill state in data path
    - Routers & switches often fail to give operators access to these

Incentives in research vs. deployment

- "The *X* proposal mostly works well" is not novel/exciting
- "The *X* proposal falls down under some conditions" is better
- "Counter-proposal *Y* works better than *X* in these cases" even better
- Possible result:
  - Vibrant research area
  - Confused operators
  - Nothing gets deployed in practice (cf. AQM work)
- Let dust settle before pushing towards production deployment?

What could help deployment?

- Tools to provide transparency (see above)
- Socialization of experimental results in operator communities
- …a battle cry (“bufferbloat”)?