Strategies for an FDL based feed-back buffer with QoS differentiation

Chris Develder
Mario Pickavet
Piet Demeester

Dept. of Information Technology (INTEC)
Ghent University, Belgium

Outline

- Node architecture
- Operation of the switch
- Simulation set-up
- Choosing a buffer configuration
- Choosing a buffer strategy
- The cost of service differentiation
- Conclusion

Node Architecture

- Node in core OPS network (backbone)
- Switch functionality:
  - fixed length packets, slotted operation
  - WDM ports
  - fully non-blocking switching matrix (SOA based)
  - wavelength conversion to solve contention
  - FDLs, int. multiple of slot, to provide buffering
- Concept:

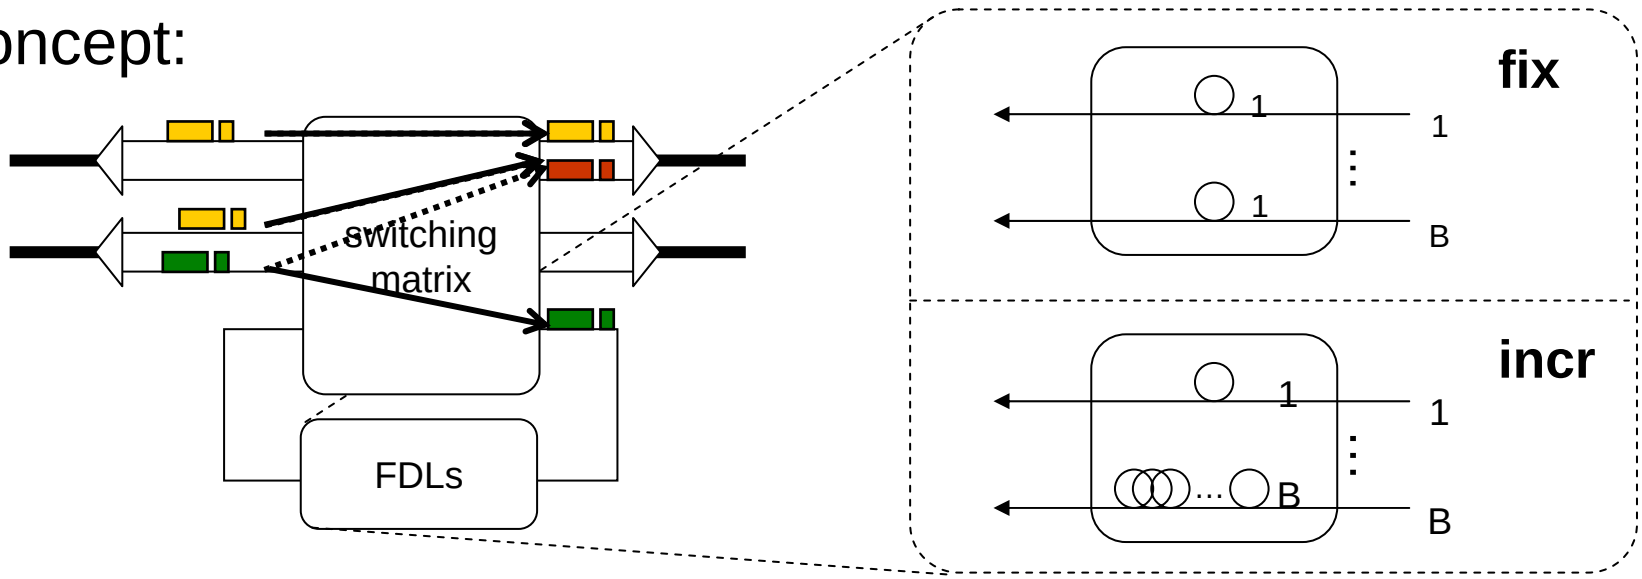

Operation of the switch

- Scheduling: each timeslot:
  - **(0)** collect packets (from inputs + FDLs) per destination output port
  - **(1)** select packets for **forwarding** along outgoing fibres; collect remaining (excess) packets
  - **(2)** elect packets for **buffering**; drop remaining packets

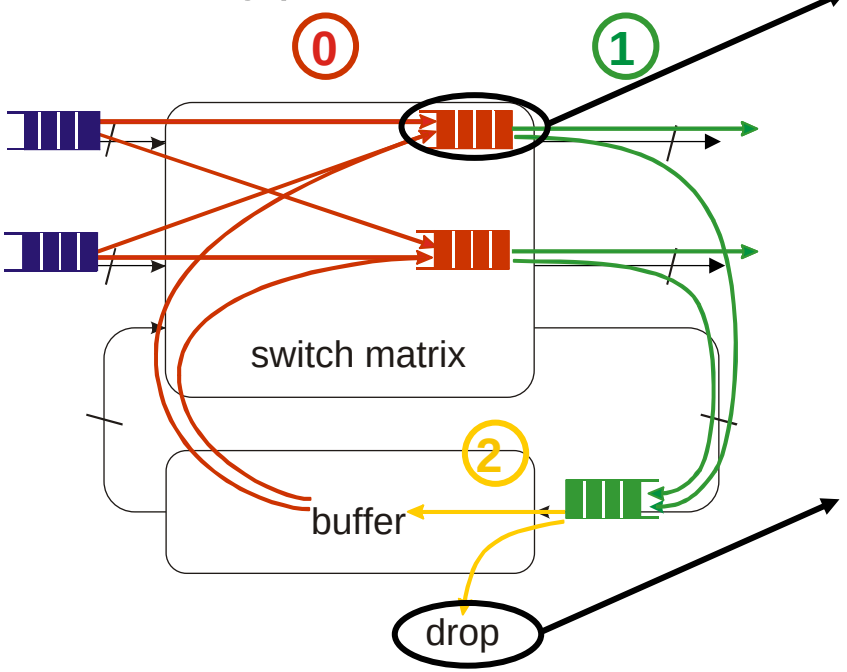

"priority queue":
1) first higher priority packets;
2) same priority: first "oldest"
3) same timestamp: random (uniform over same pri and tstamp)

tstamp = when packet enters switch

we want to evaluate performance of switch in terms of Packet Loss Rate (PLR)

Simulation set-up (1)

- F=6 input/output fibres
- W=32 wavelengths per fibre
- 256x256 switch: max. 64 buffer ports ($B \leq 64$)

- P=3 priority classes
  - 25% lowest pri 0
  - 25% middle pri 1
  - 50% highest pri 2

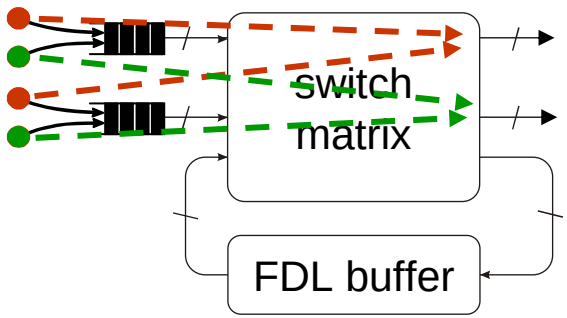

- traffic sources:
  - source generates traffic for certain (in,out)-pair and with fixed priority $\Rightarrow$ F·F·P=6·6·3=108 sources
  - uniform pattern: avg. amount of traffic is same for all (in,out)-pairs

Simulation set-up (2)

- Different traffic source types

- Poisson:
  - classical Poisson process (exp. distr. IATs)
- GeoOnOff:
  - bursty
  - strictly alternating On/Off periods, with packet every timeslot during On-period; no packets during Off-period
  - geometric distribution for period lengths
- ParetoOnOff:
  - bursty, self-similar
  - strictly alternating On/Off periods
  - Pareto distribution (heavy tailed) for period lengths

Buffer configuration

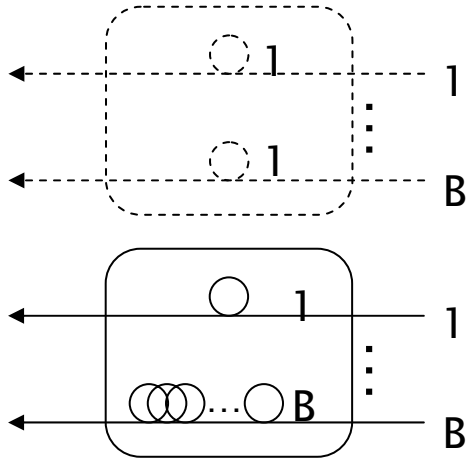

- **Increasing** FDL lengths give far lower PLRs (order of magnitude or more)
- **ParetoOnOff:** difference is limited (factor ~2) and doesn't vary much with increasing nr. buffer ports

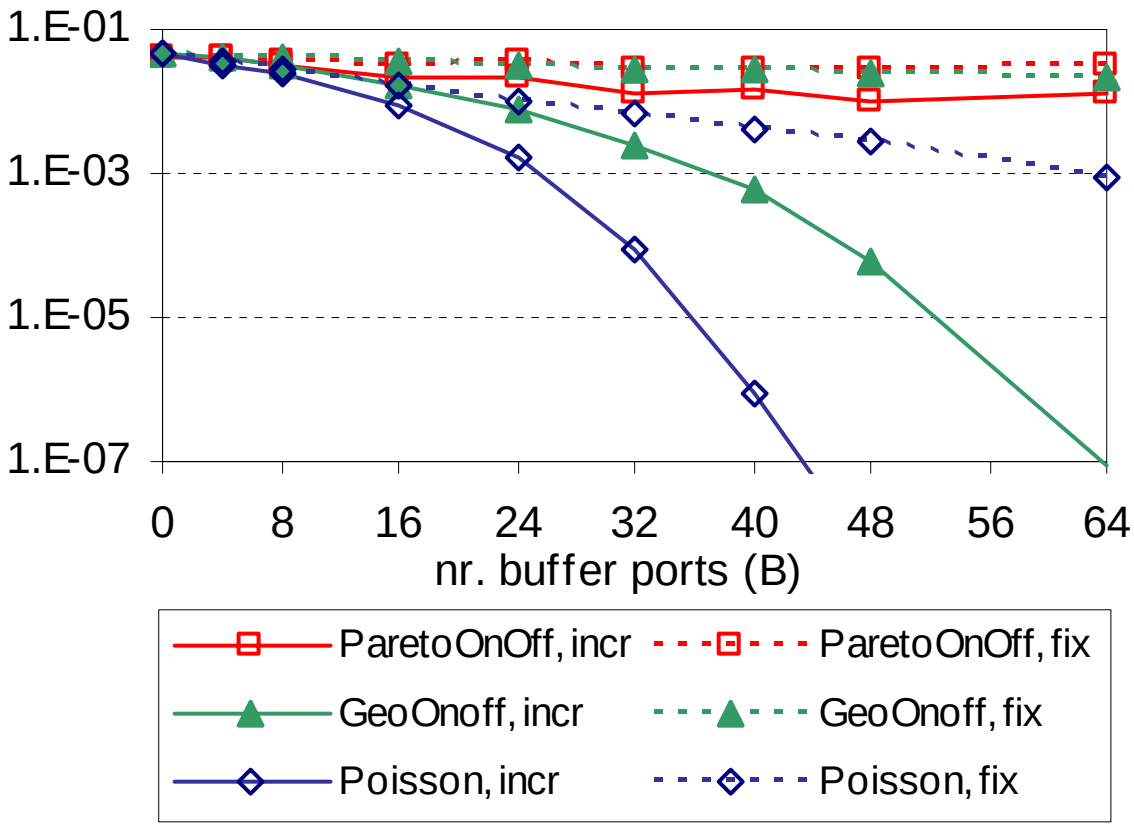

Buffer strategy (1)

- Problem:
  - FDLs are shared over all output ports: use buffer as efficient as possible (only single copy in FDL)
  - choosing FDL length = deciding when it will re-enter the switch (and have another attempt at forwarding)
  - multiple packets will leave FDL buffer block at same time

- Need to choose “appropriate” FDL

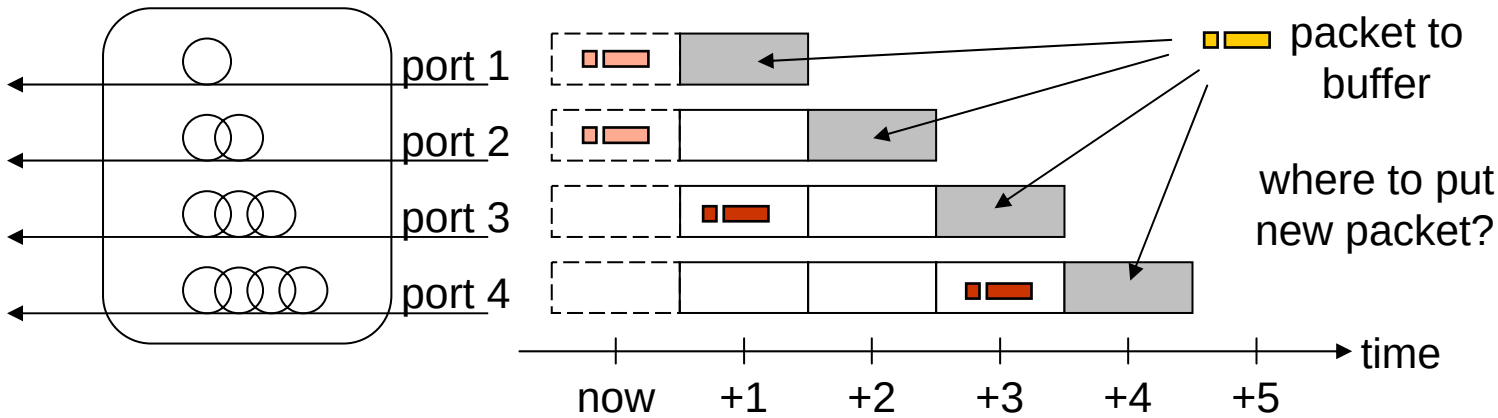

Buffer strategy (2)

- Comparison of 4 strategies:
  1) **MinDelay**: minimal delay (not intelligent, but fast)
  2) **NoOvr**: do not allow overload (buffer output will never contain more than W packets destined for same fibre)
  3) **AvoidOvr**: avoid overload iff possible; if not: use smallest delay
  4) **Balance**: use FDL with length L such that nr. of packets leaving buffer simultaneously (at now+L), is minimal

  Note: if B≤W, then no overload is possible: 1,2,3 are equivalent

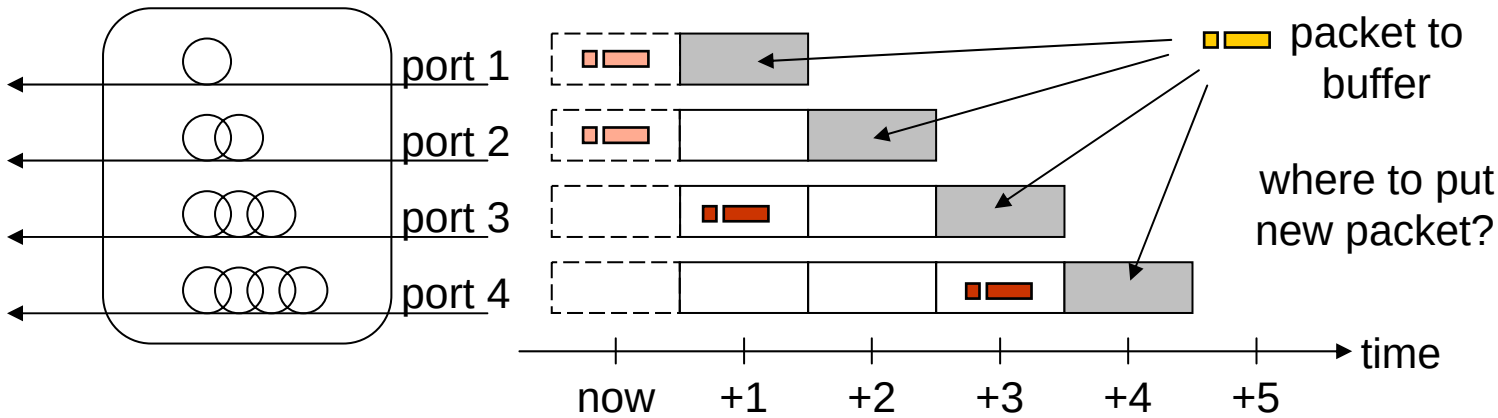

Buffer strategy (3)

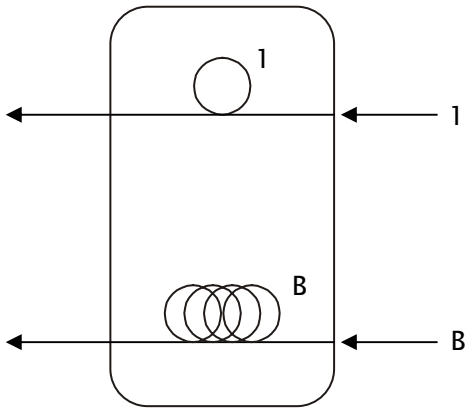

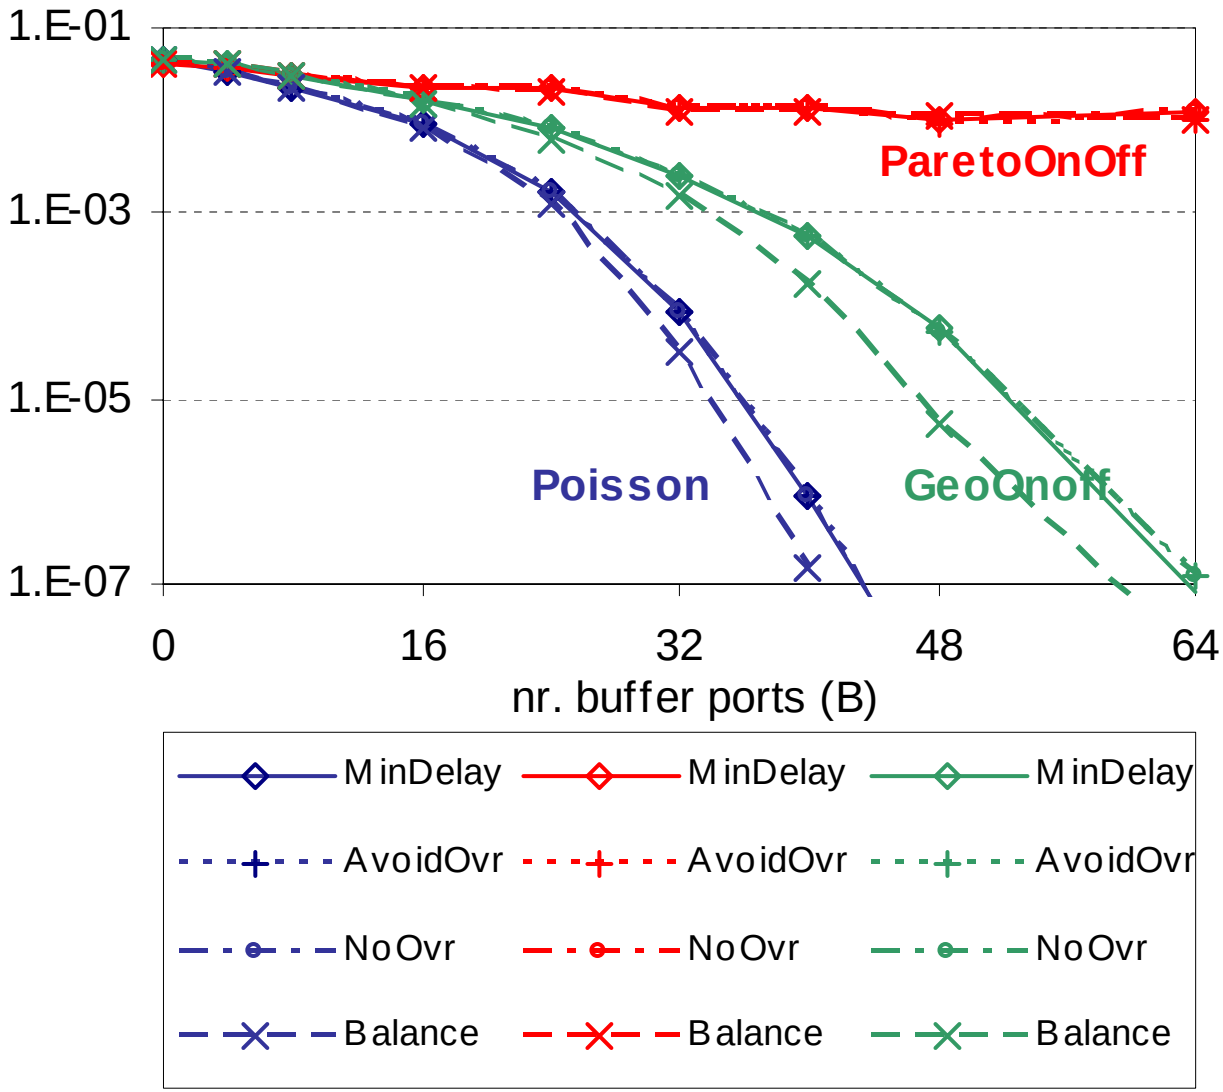

- **Balance** outperforms other strategies for GeoOnoff and Poisson
- **ParetoOnOff:** no matter what, PLR can't be reduced significantly

The cost of QoS (1)

- QoS:
  - Priority scheme achieves effective service differentiation
  - e.g. results for increasing FDL loops, Balance strategy, load=0.95

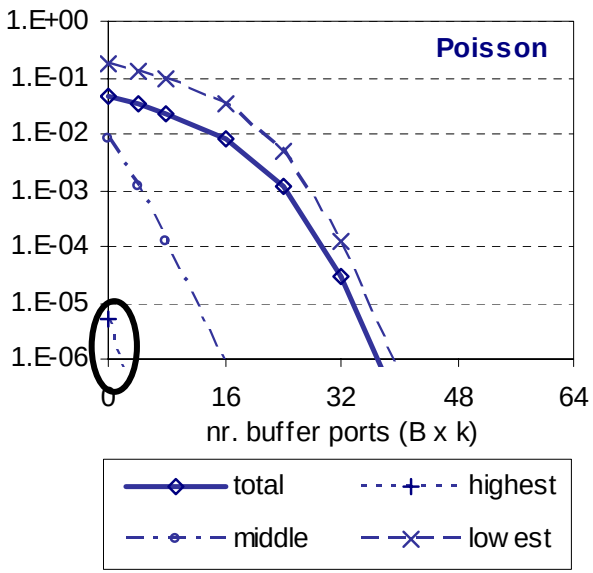

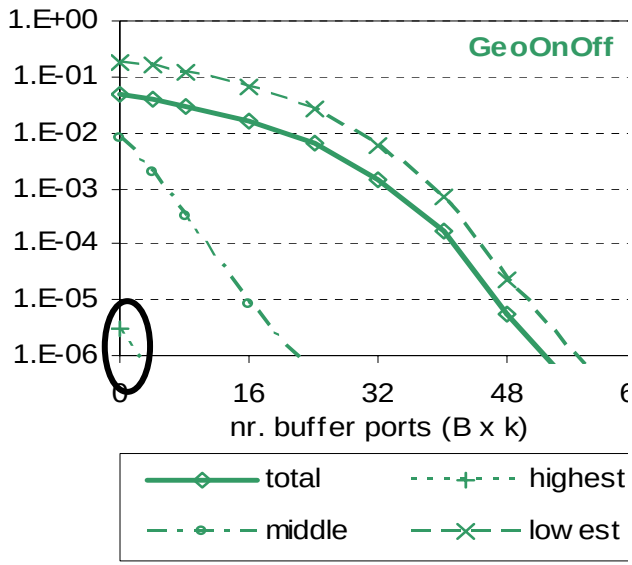

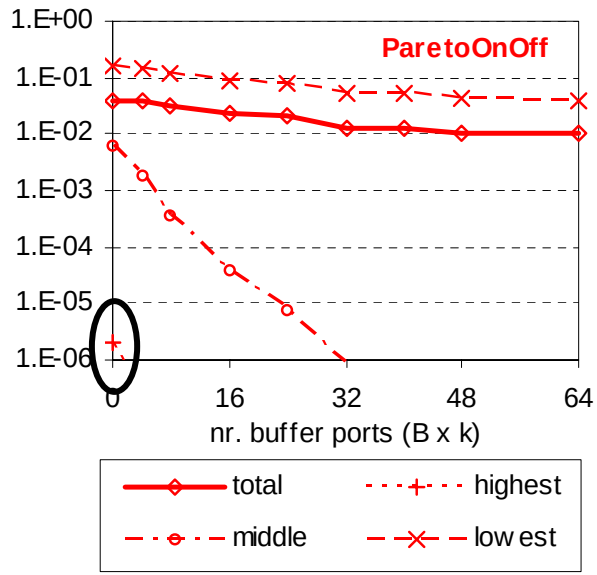

The cost of QoS (2)

- QoS:
  - Priority scheme achieves effective service differentiation
- But…
  - is there a penalty in terms of (higher) overall packet loss rate?
- Answer:
  - compare, using identical traffic pattern in either case, results of using priority scheme with results when ignoring priority info

The cost of QoS (3)

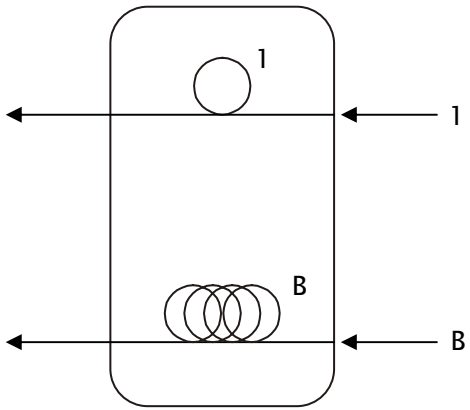

- **Poisson: no diff**
- **GeoOnOff: with prio is better;**
  (favour some output ports, thus spreading more in time; less correl. over larger timescales)
- **ParetoOnOff: no prio slightly better (?);**
  (correl. over very large timescales)

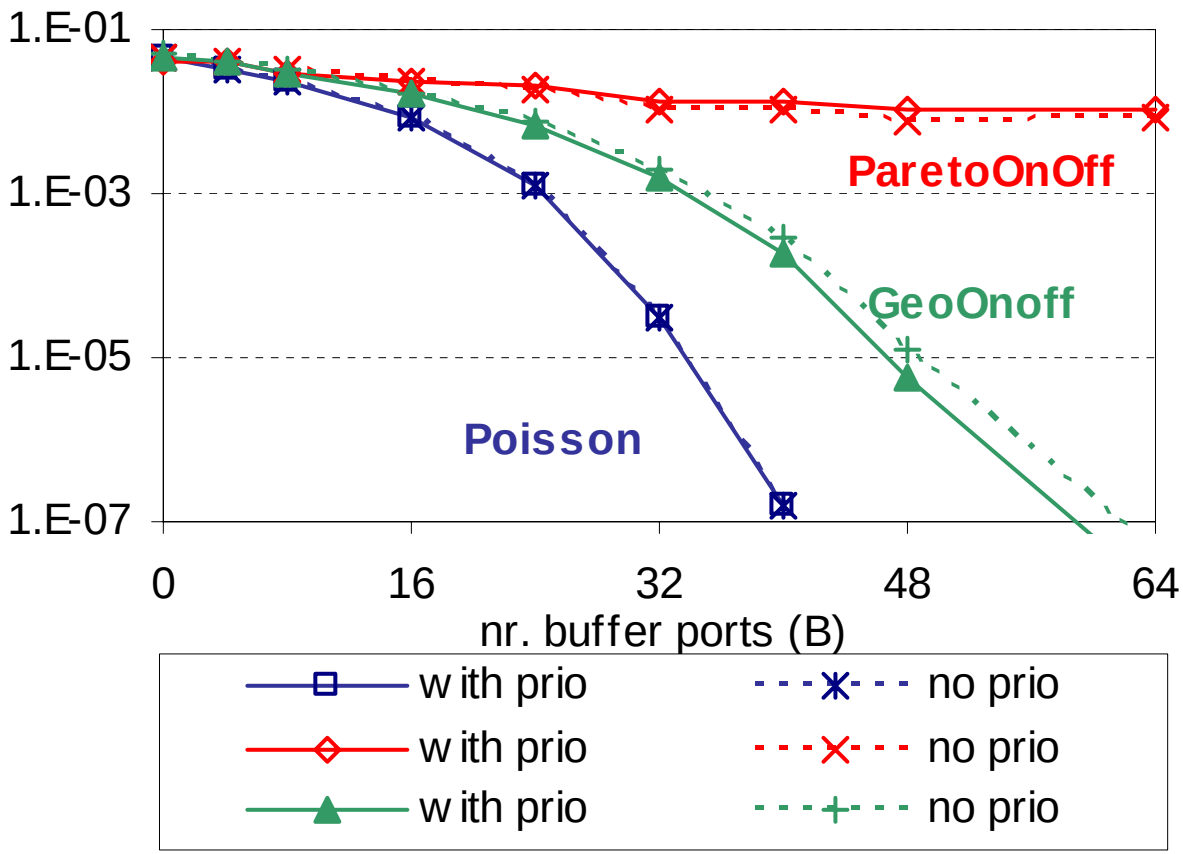

Conclusions

- ParetoOnOff traffic: minor PLR reduction through buffering
- simple priority scheme: good CoS separation;
  does not significantly increase overall PLR
- buffer structure: using different FDL lengths gives significantly better PLR performance for same switch fabric port count
- increasing FDL lengths: Balance strategy is best in terms of PLR performance

That’s all, folks!

… thanks for your attention

Buffer strategy (4)

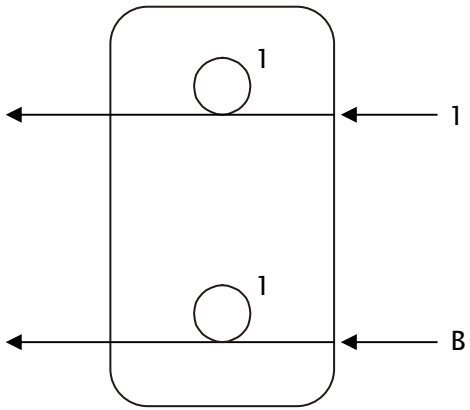

- Single FDL length: MinDelay, AvoidOvr and Balance are equivalent
- **No significant difference** for NoOvr (slightly better when positive correlation between overload in successive slots: On/Off; slightly worse for memoryless Poisson)