Effects of long RTT – A case for RTT adaptivity

Cyriel Minkenberg & Mitch Gusat
IBM Research GmbH, Zurich
November 1, 2007

Baseline Input-Generated Hotspot

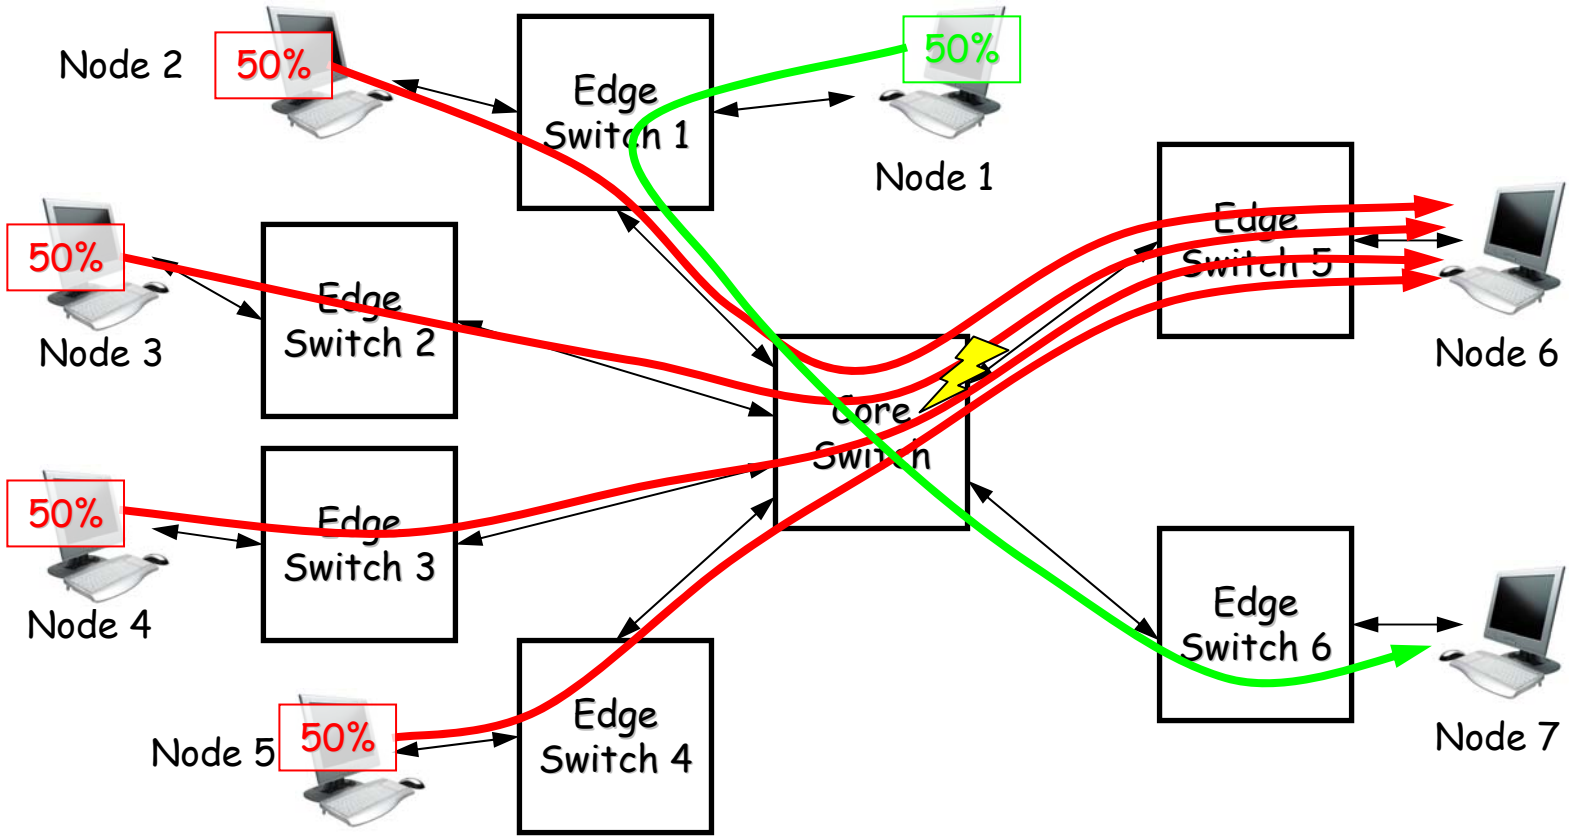

- Four culprit flows of 5 Gb/s each from nodes 2, 3, 4, 5 to node 6 (hotspot)
- One victim flows of 5 Gb/s from node 1 to node 7
- Fair allocation provides 2.5 Gb/s to all culprits and 5 Gb/s to the victim

Simulation Parameters

- Baseline scenario, $t_{hotspot}$ = 0.1 – 1.0 s
- M = 1.5 MB/port
- Unlimited input buffers
- Q_{eq} = 375 KB
- P_s = 1% (QCN: 1-10%)
- QCN active increase
  - to_thresh = packet_size / p_sample
  - Ri = 12 Mb/s
- Drift enabled: 4 Mb/s every 20 ms
- ECM_MAX enabled, Q_{mc} = 1.5 MB
- No ECM_(0,0), no PAUSE
- Per-link RTT = [2 us, 20 us, 50 us, 200 us, 500 us, 1ms]
  - Note that RP ⇔ CP RTT = 2*link RTT
- 8-bit quantization
- W = [2, 8, 32]

Hot queue length, QCN, P_{sample_base} = 1%

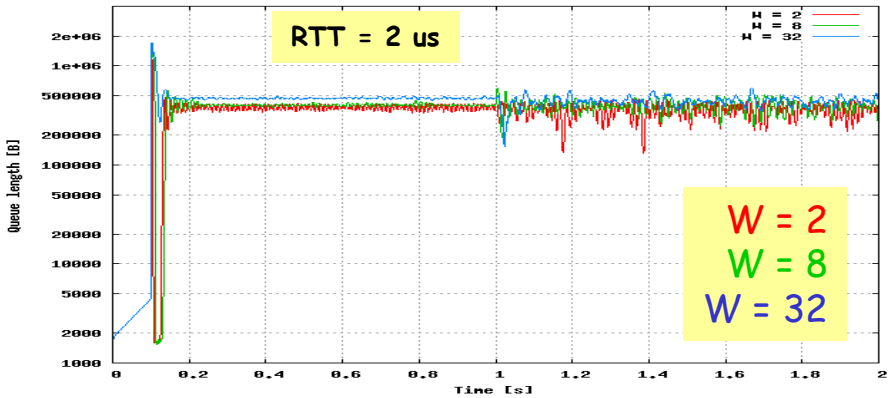

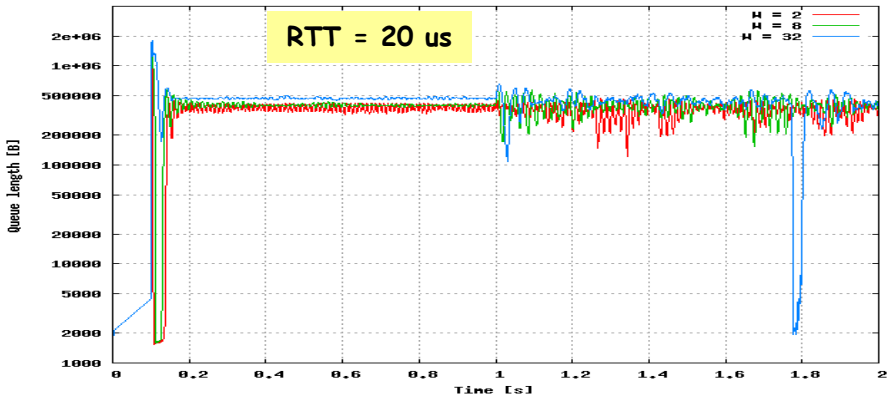

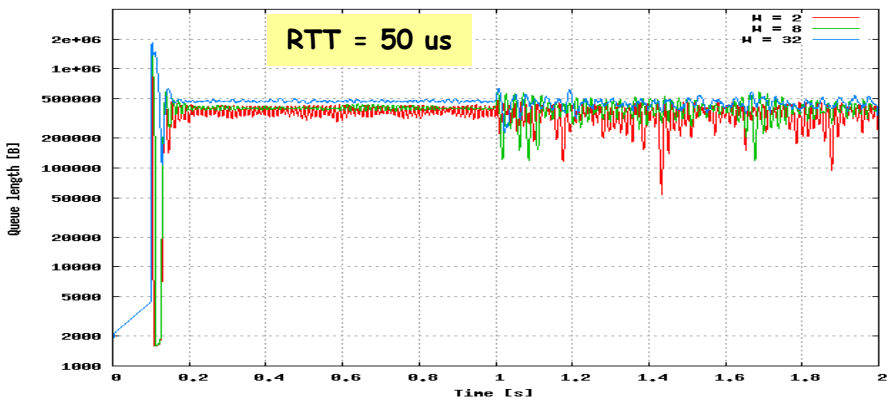

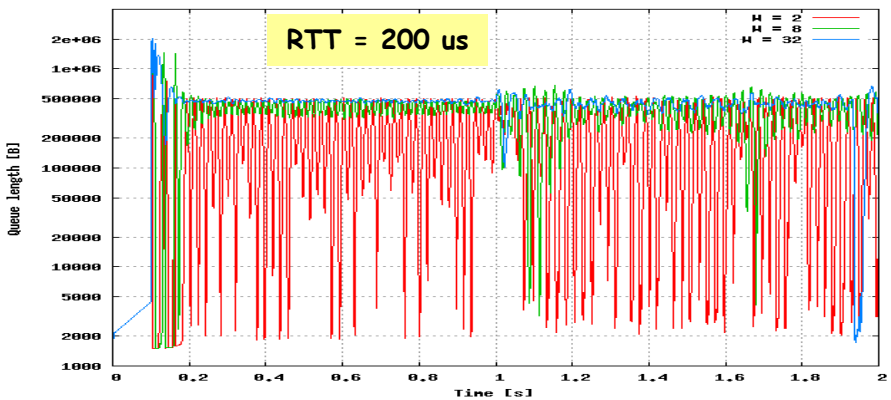

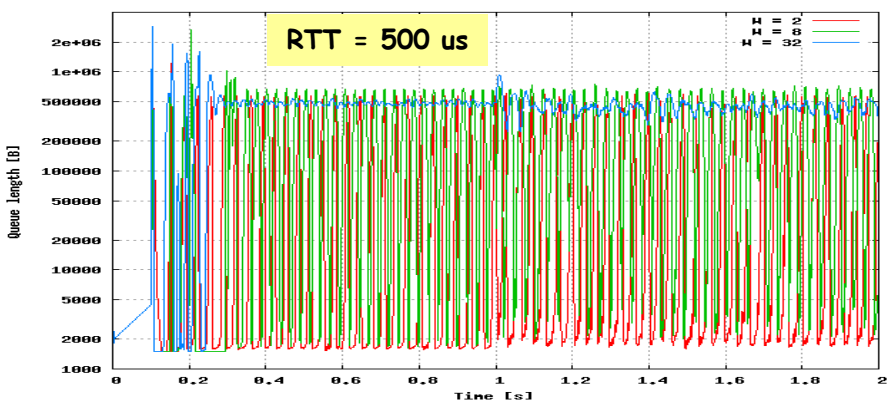

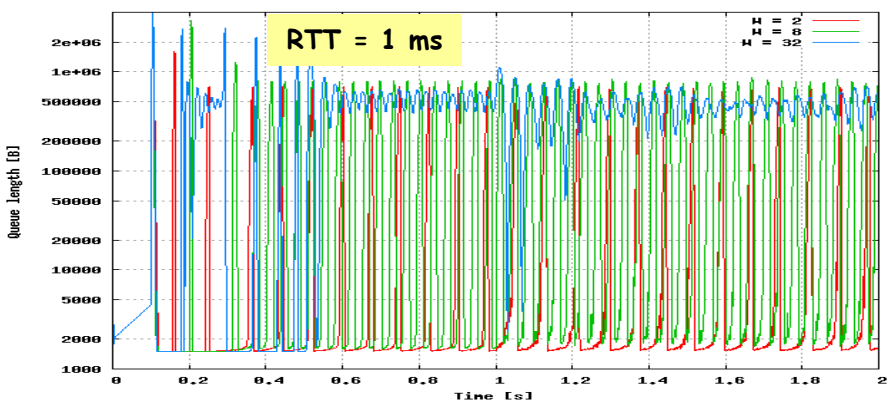

Hot port throughput, QCN, P_{sample_base} = 1%

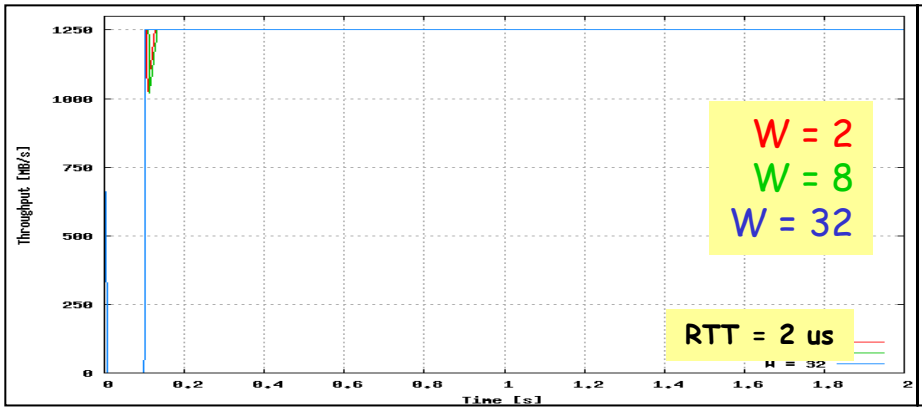

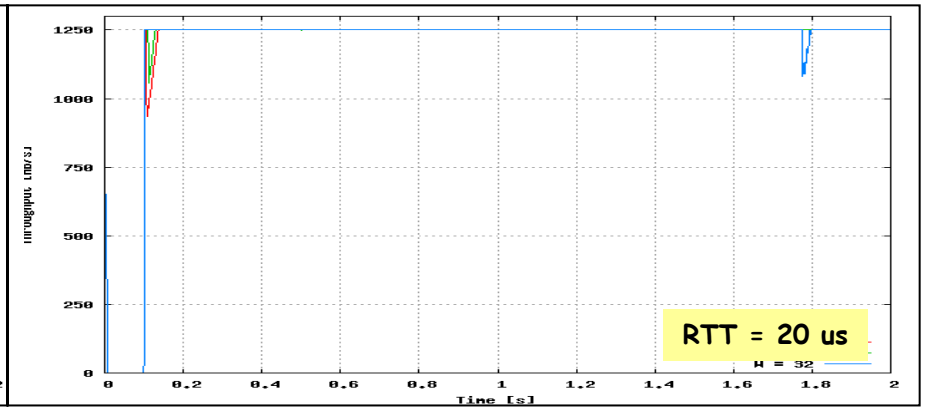

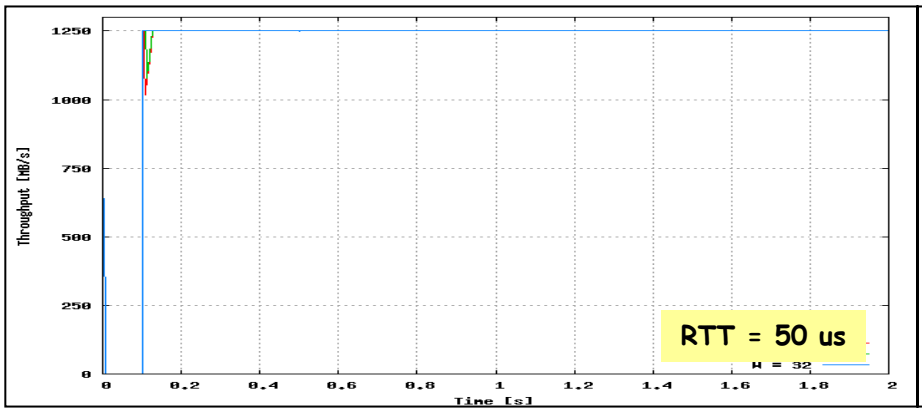

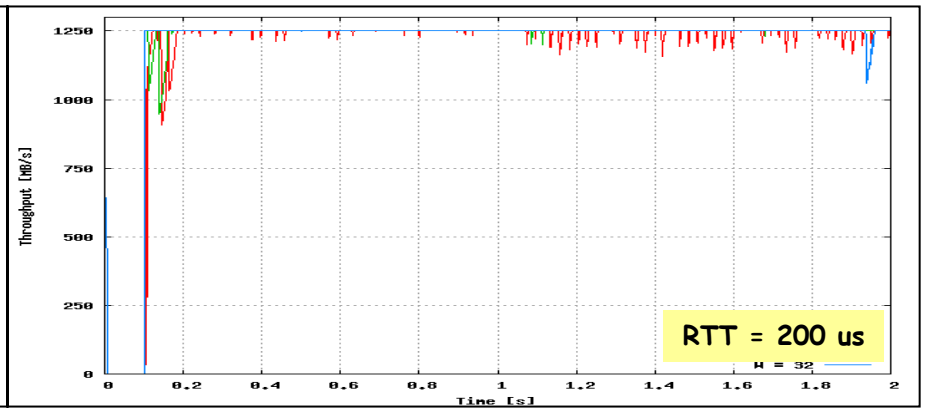

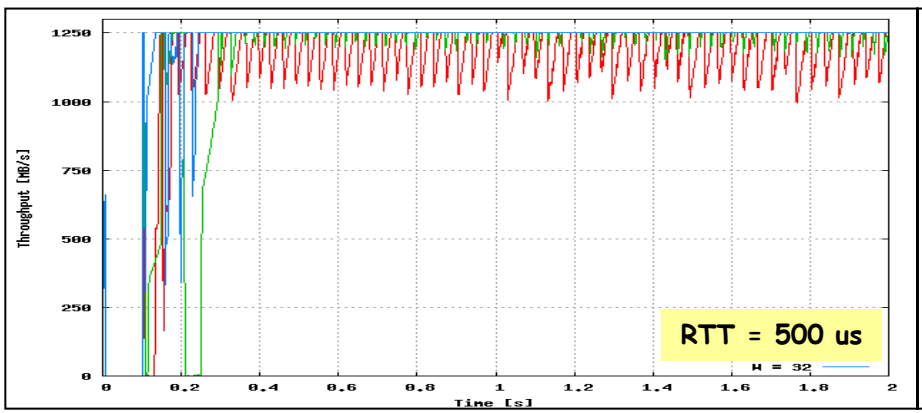

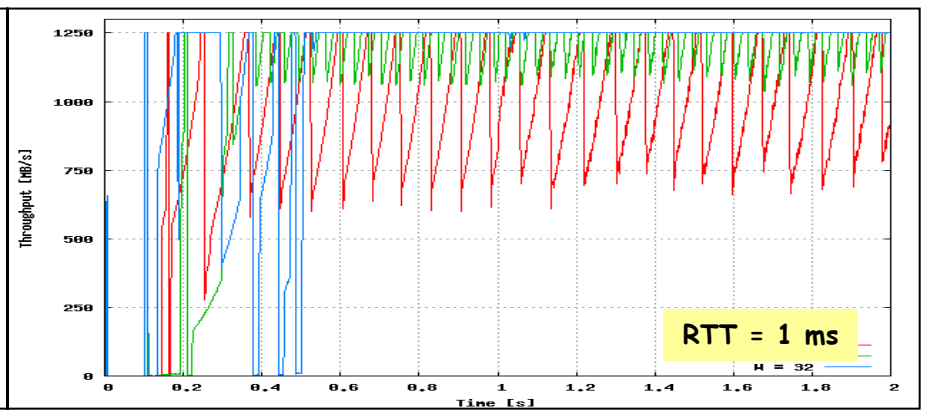

Hot queue length, QCN, Link RTT = 500 us

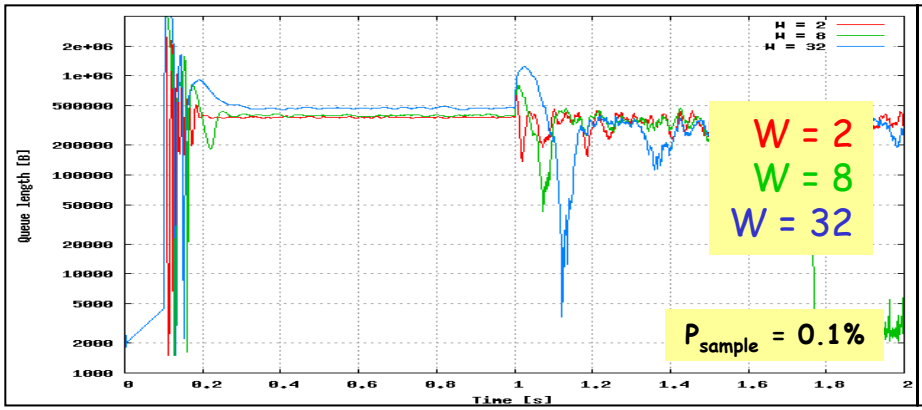

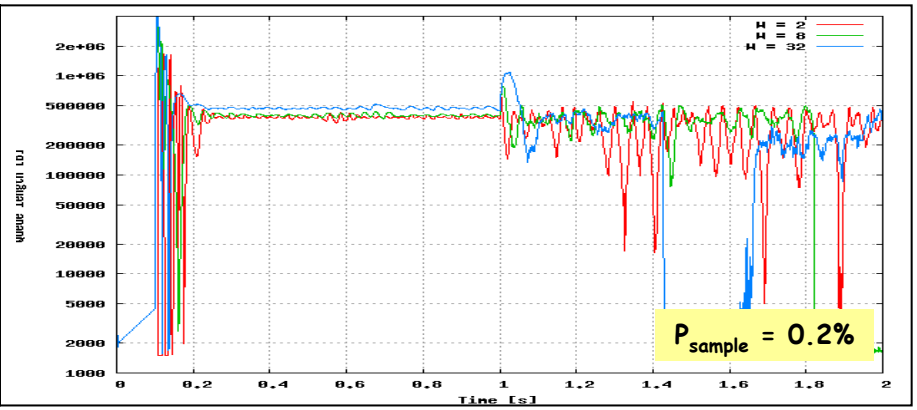

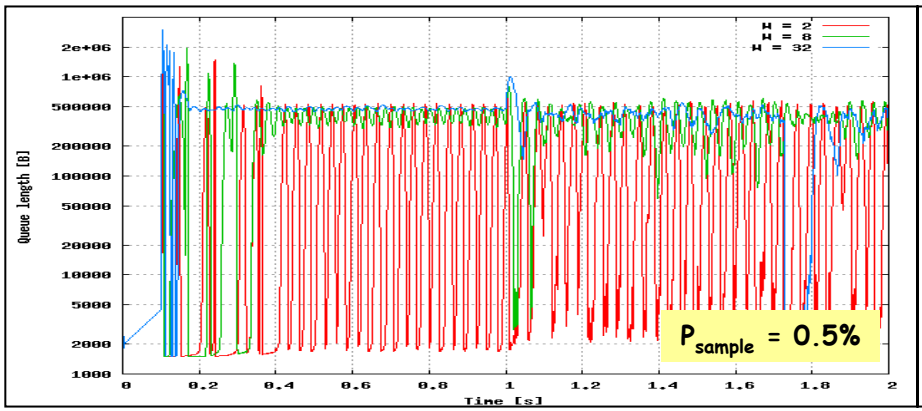

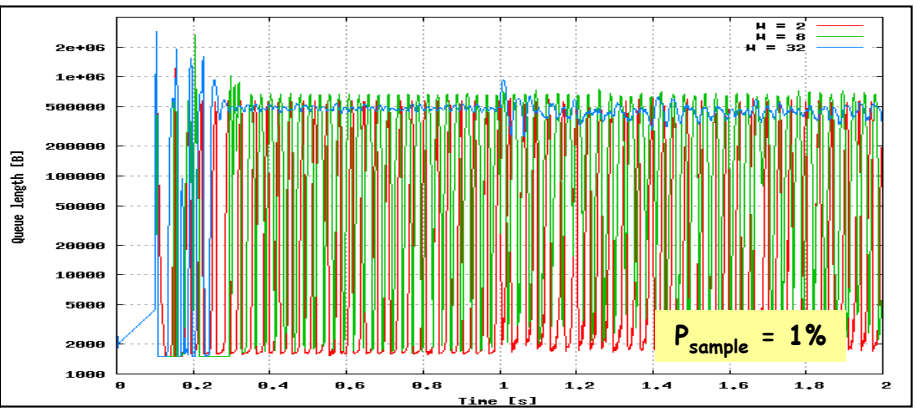

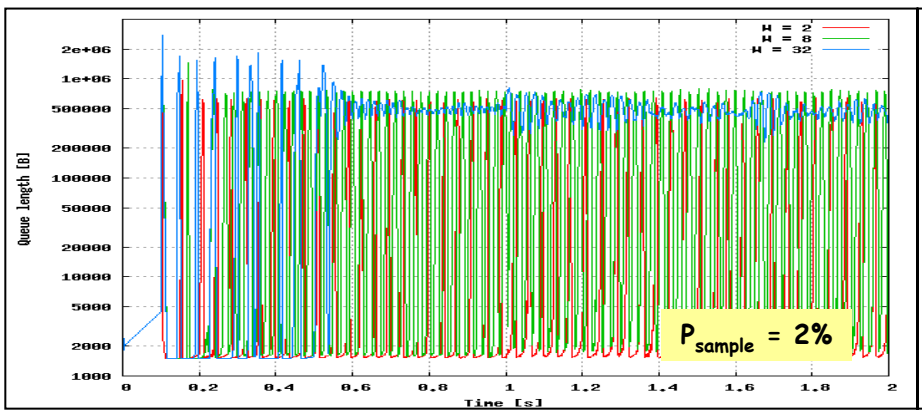

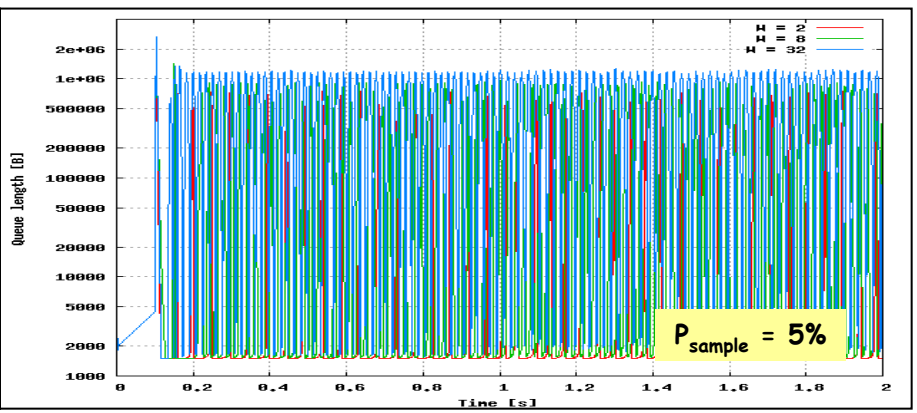

Hot queue length, QCN, P_{sample_base} = 1%, with RTT filtering

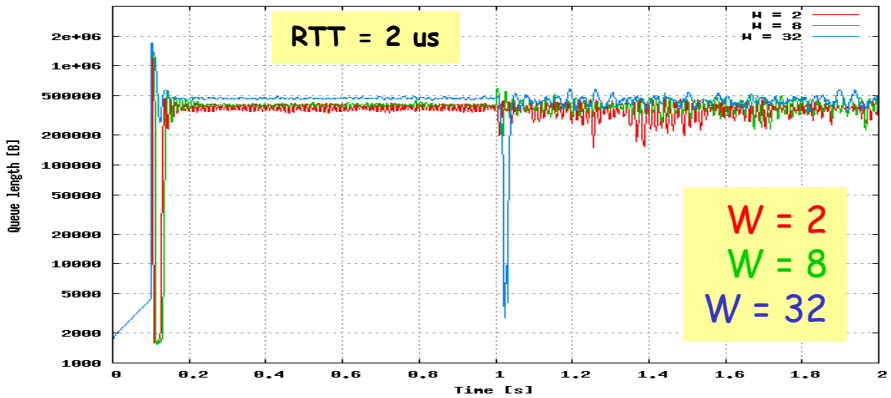

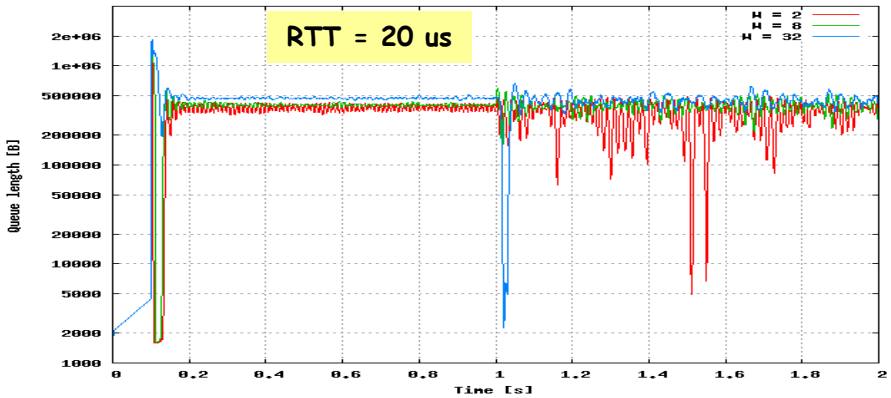

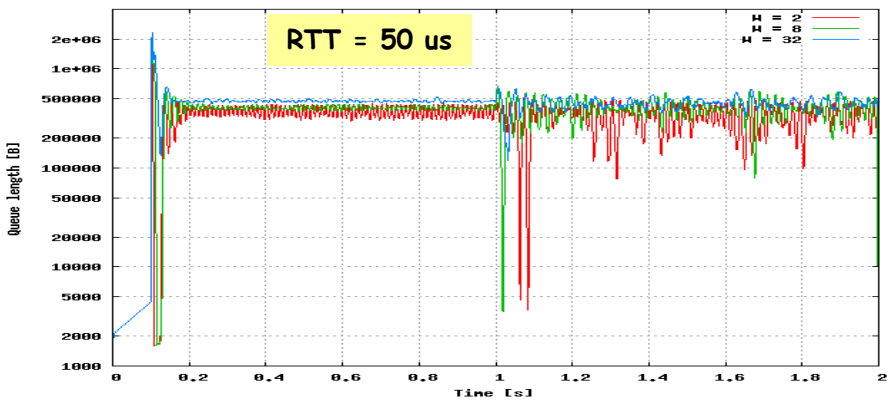

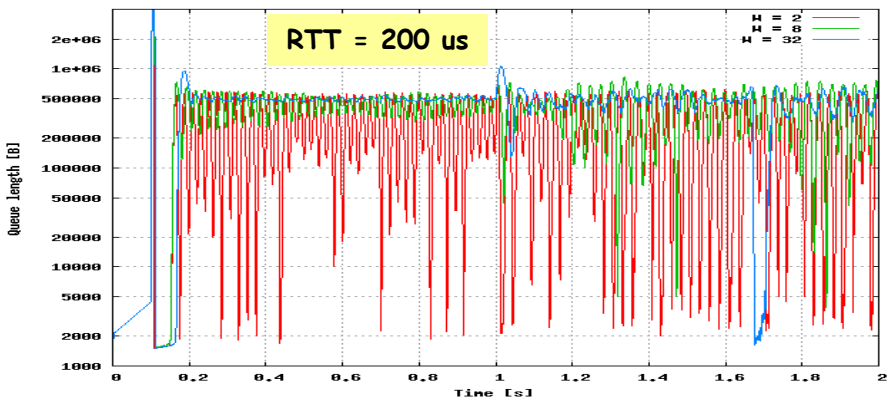

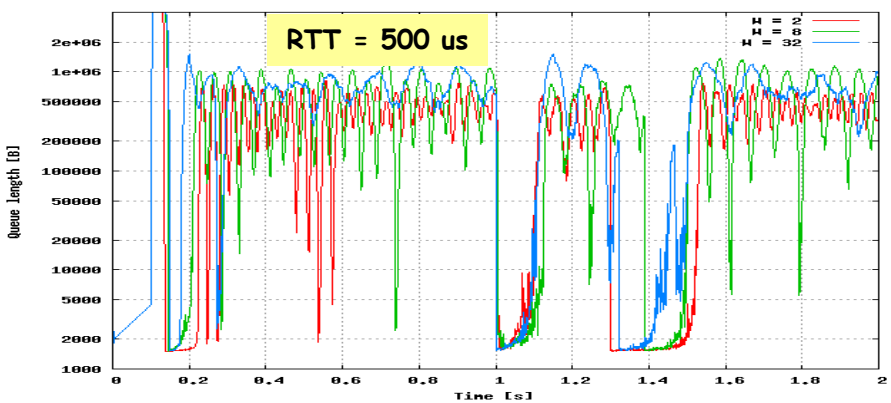

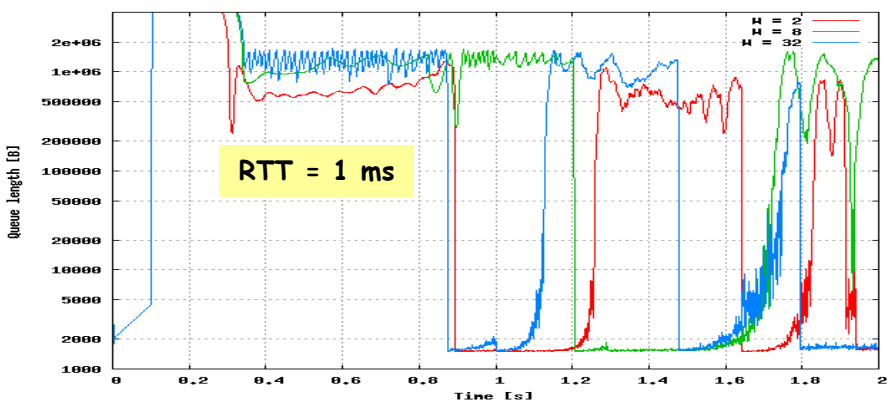

Hot port throughput, QCN, P_{sample_base} = 1%, with RTT filtering

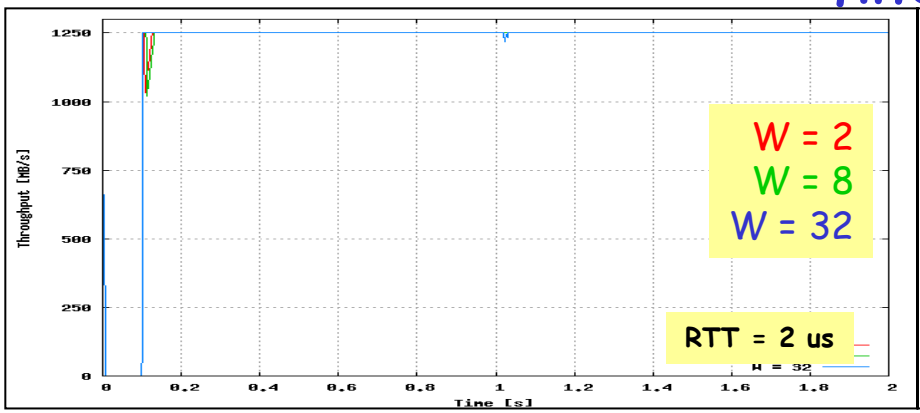

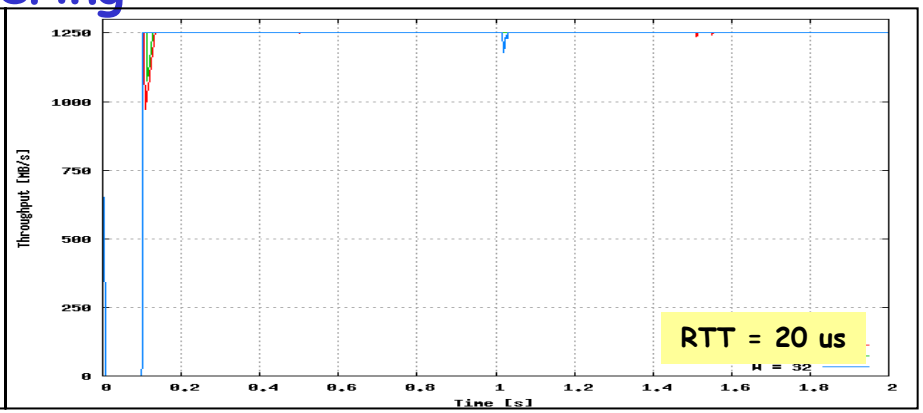

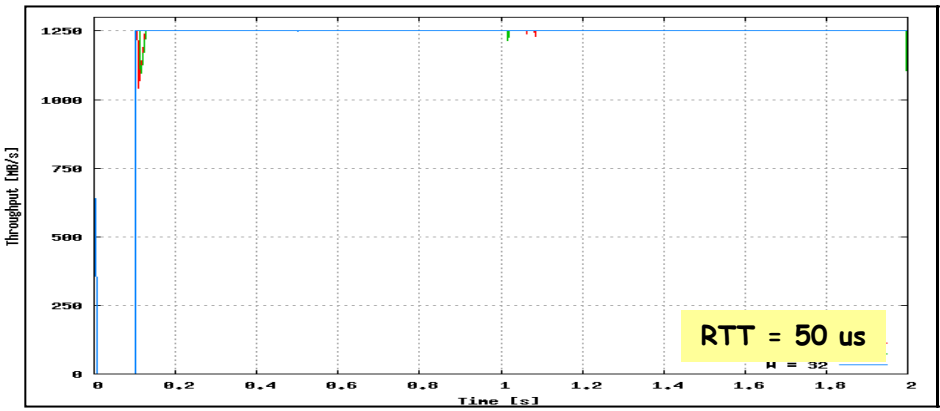

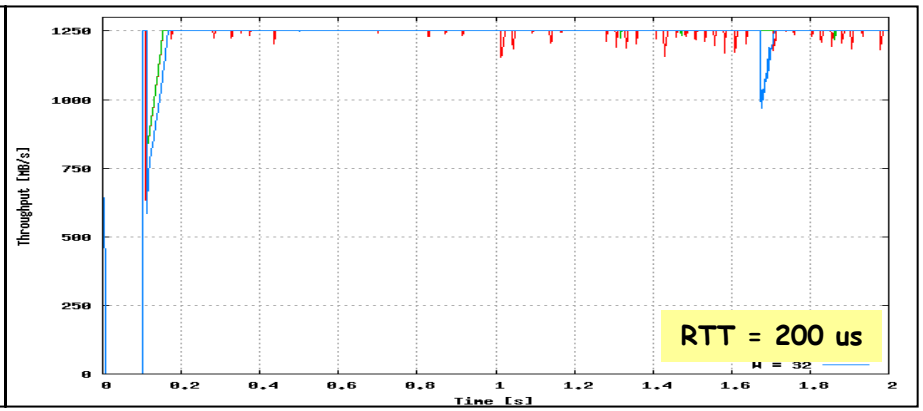

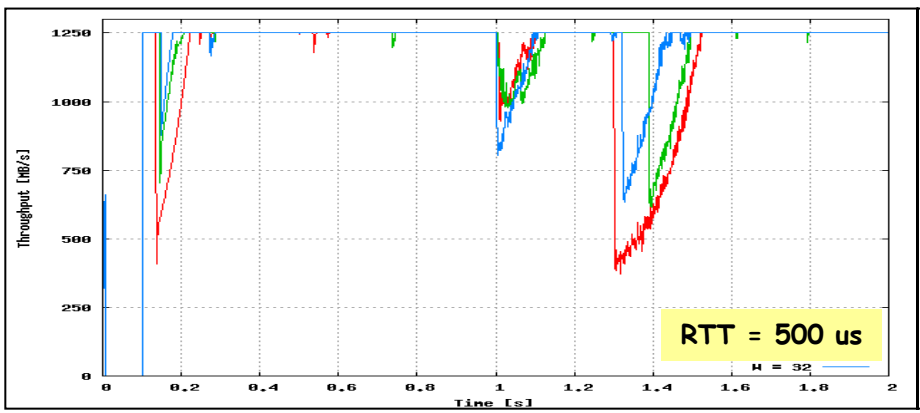

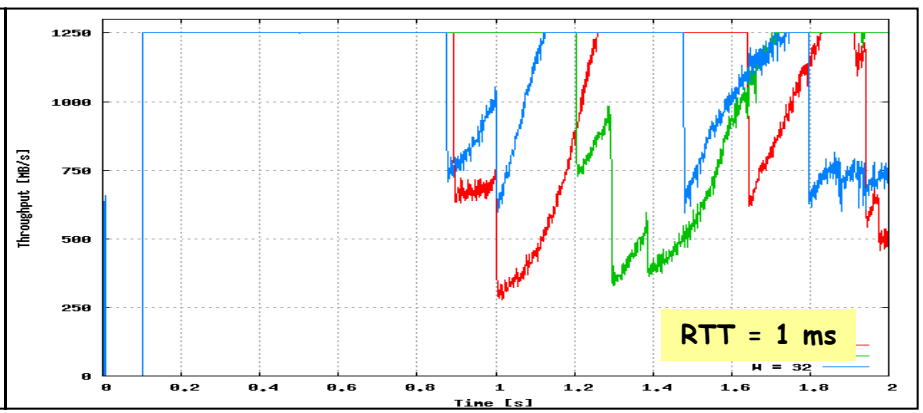

Hot queue length, ECM, P_{sample} = 1%

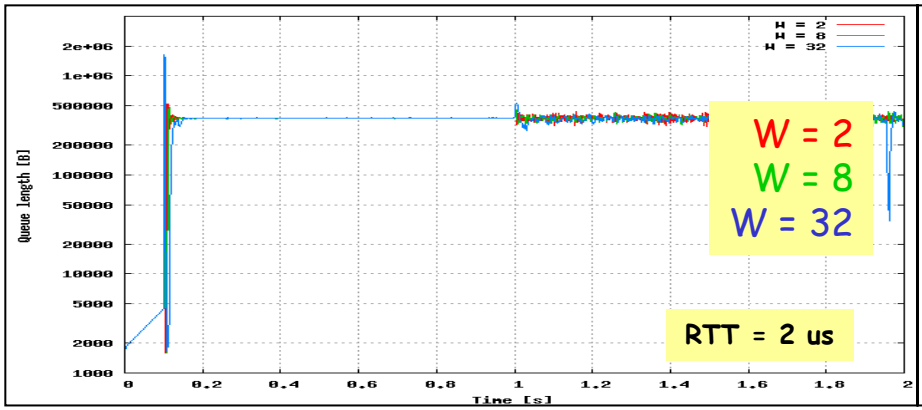

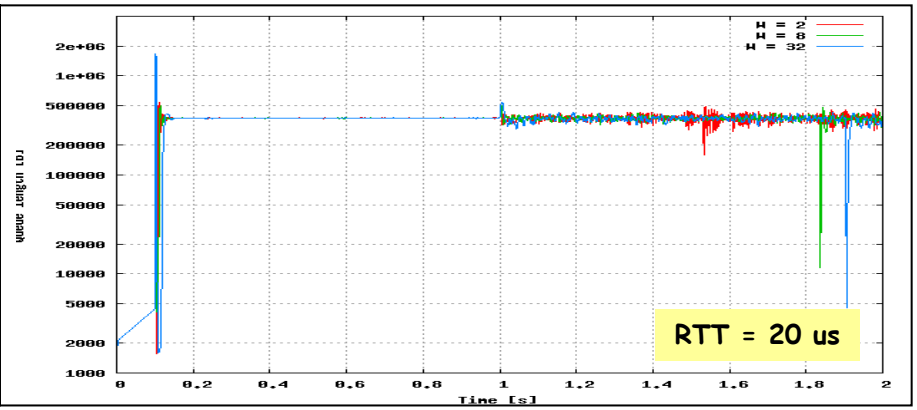

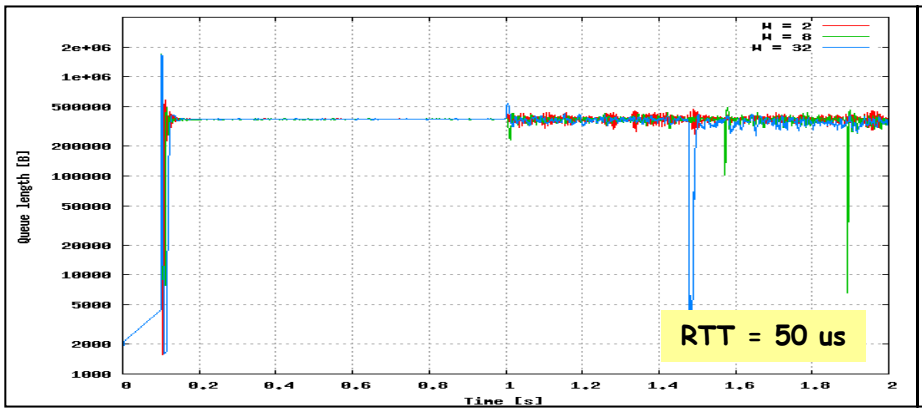

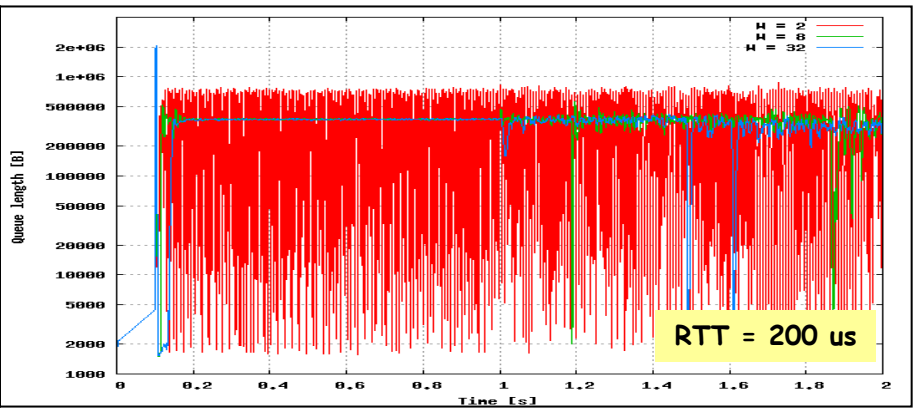

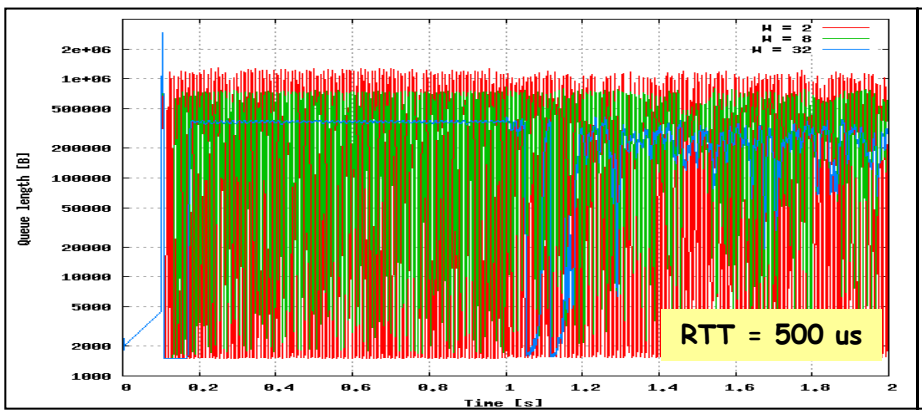

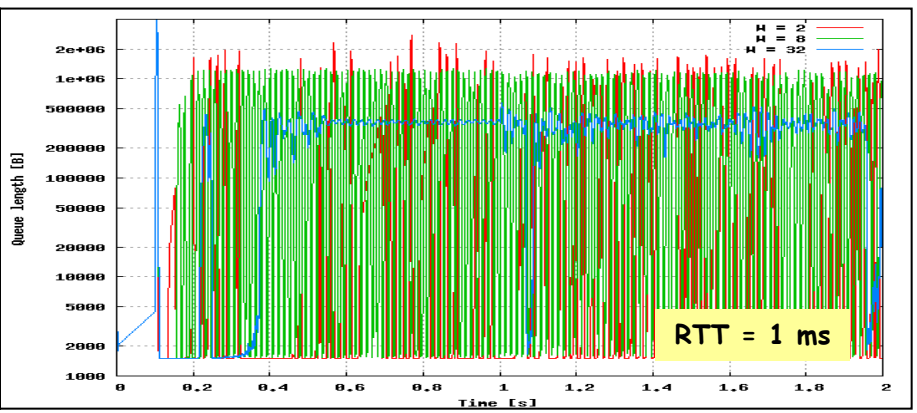

Hot port throughput, ECM, P_{sample} = 1%