Zurich Hotspot Benchmark
- Simulation Results -

Single-Hop Output-Generated Hotspot

Cyriel Minkenberg and Mitch Gusat

December 8, 2006

Simulation parameters

- Scenario
  - Single-hop output-generated hotspot (OG-HS)
    - OG-HS definition: step impulse of width [0.1-0.5] s
  - Uniform loading from all nodes
    - Uniform = sending at same rate to all nodes except self
  - Load = 85%
- Network
  - N = 16; M = 600 KB/port
  - 1) Shared memory
    - PAUSE applied to all ports simultaneously based on global high/low watermarks
    - $watermark_{high}$ = N*(M – rtt*bw)
    - $watermark_{low}$ = $watermark_{high}$ / 2
  - 2) Partitioned memory per input
    - Deadlock prevention
    - PAUSE applied on a per input basis based on local high/low watermarks
    - $watermark_{high}$ = M – rtt*bw
    - $watermark_{low}$ = $watermark_{high}$ / 2
- BCN modelled as quadratic lag system
- Params used:
  - W = 2.0 (baseline)
  - G_i = 6.6667*10-4
  - G_d = 1.6667*10-6
  - Q_{eq} = 150 KB (= M/4)
  - P_{sample} = 2%
  - R_u = R_{min} = 10 Mb/s
  - No BCN(0,0) or BCN_MAX

T_{put} & Q_{length} - Shared memory: With and w/o BCN

Aggregate T_{put}

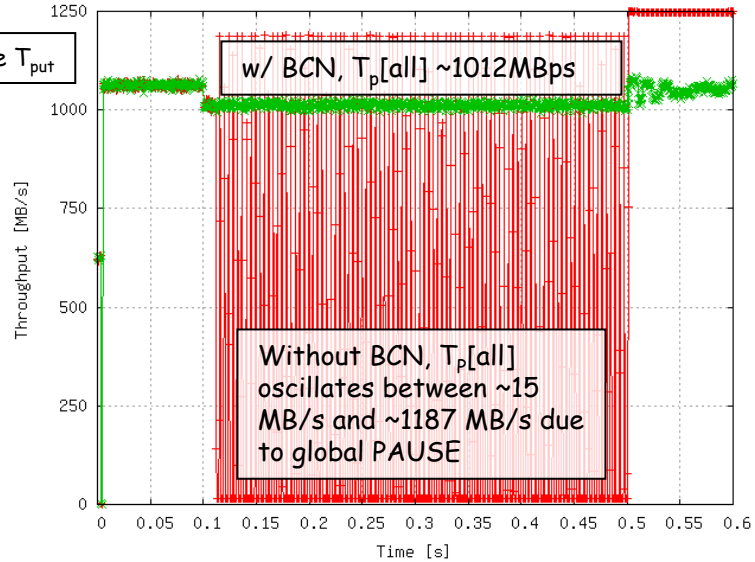

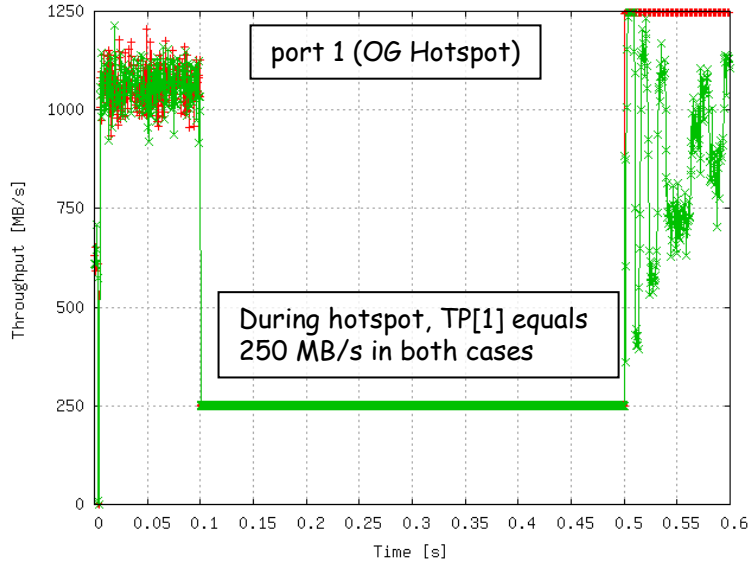

port 1

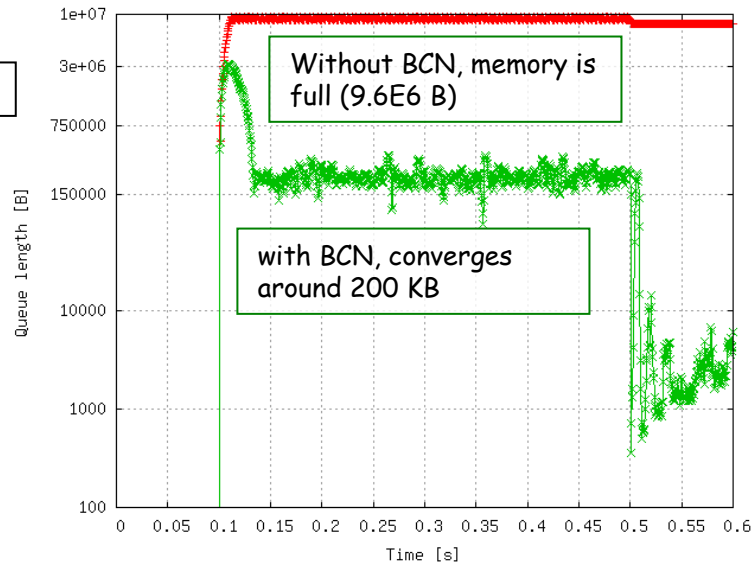

P_{sample} = 2%

No BCN

Throughput & queue length - Partitioned memory

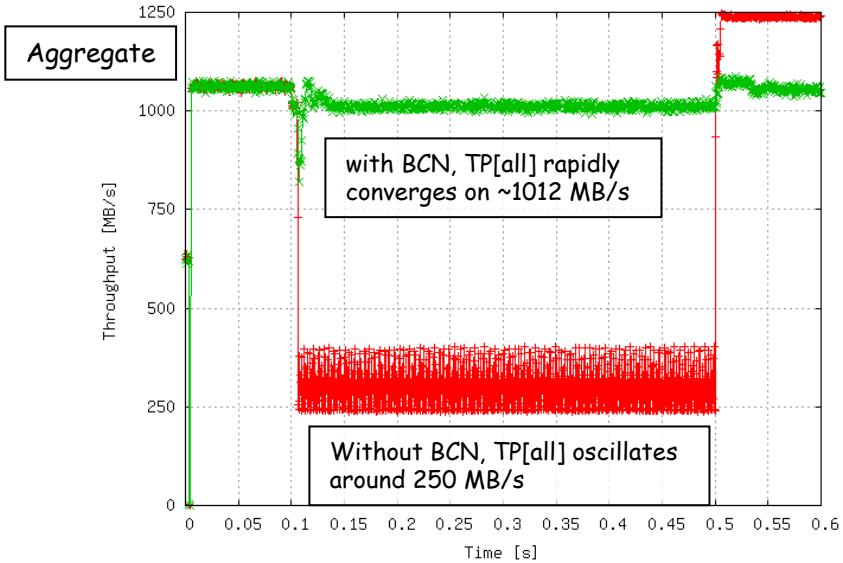

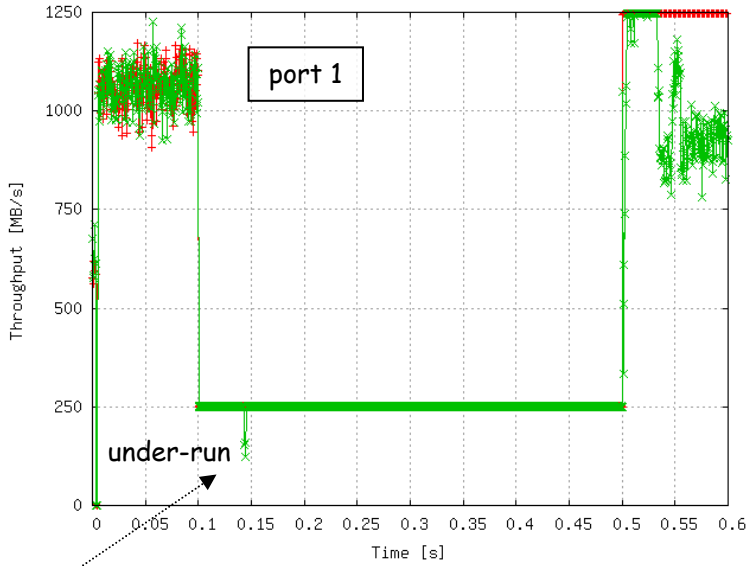

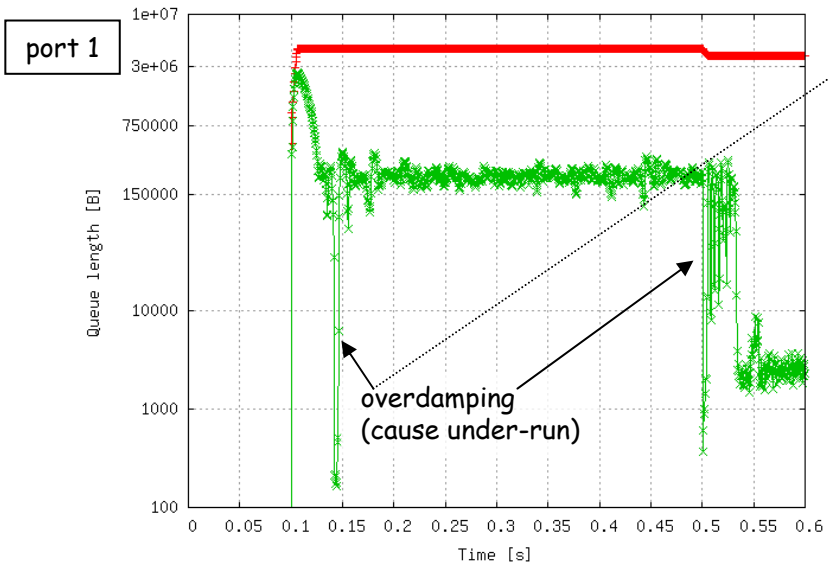

P_{sample} = 2%

No BCN

G_d sensitivity – Shared memory

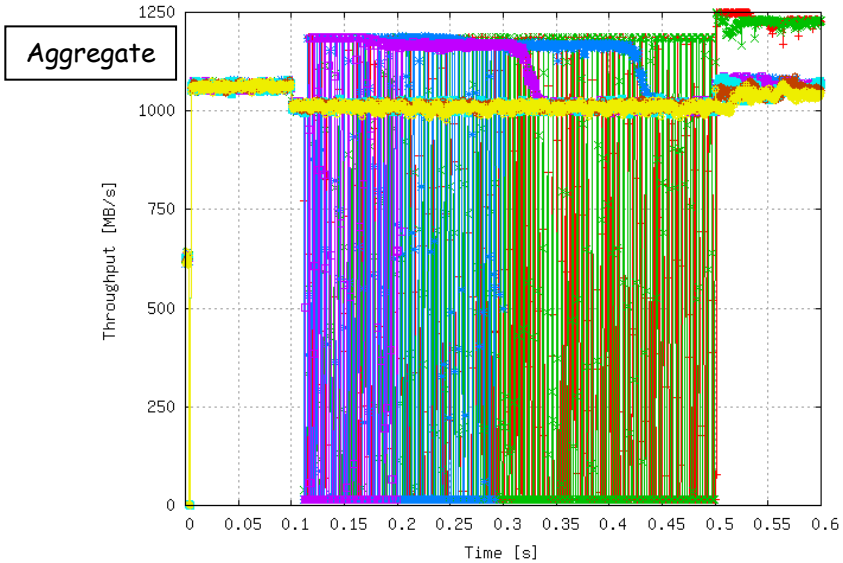

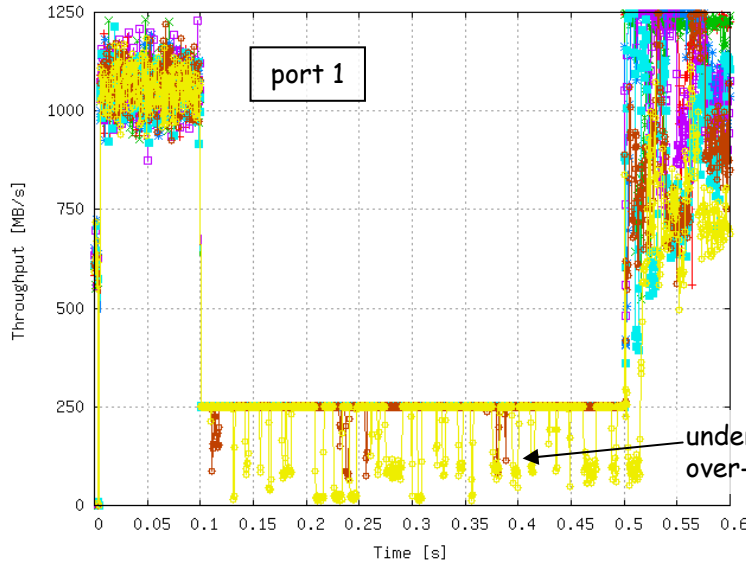

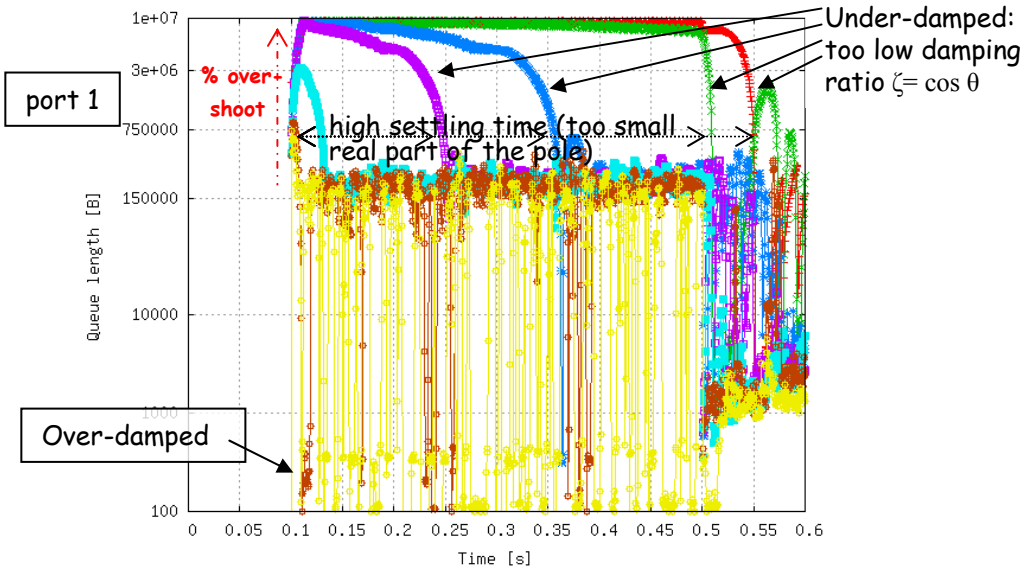

Damping range:

$G_{d0} = 6.6667*10^{-7}$

$G_d = 0.10*G_{d0}$

$G_d = 0.25*G_{d0}$

$G_d = 0.50*G_{d0}$

$G_d = 1.0*G_{d0}$

$G_d = 2.5*G_{d0}$

$G_d = 5.0*G_{d0}$

$G_d = 10.0*G_{d0}$

under

over

G_d sensitivity – Partitioned memory

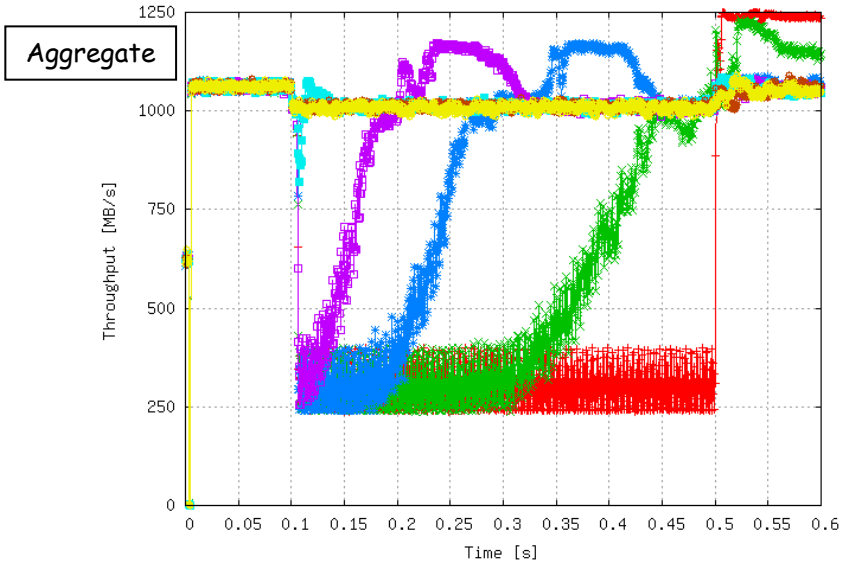

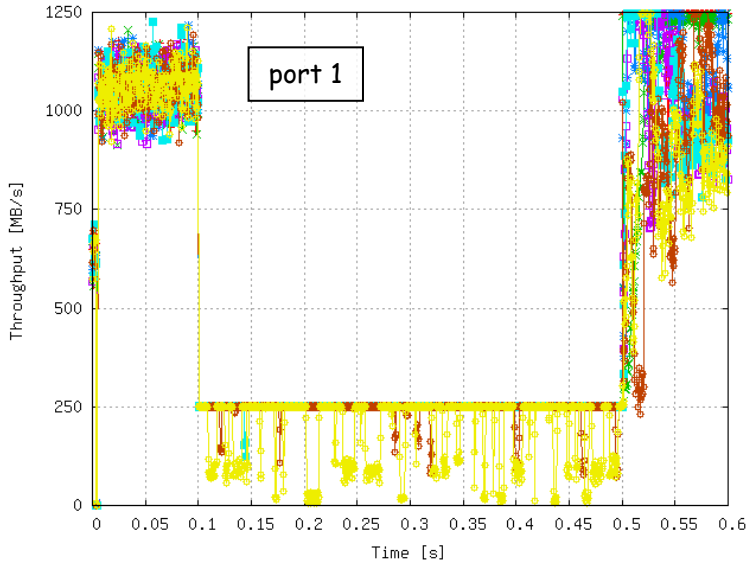

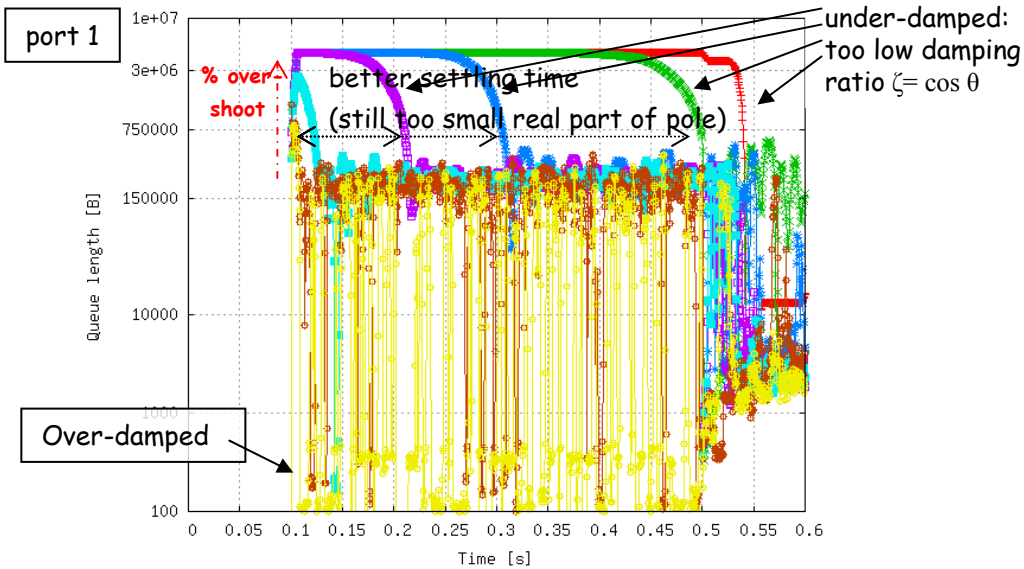

Damping range:

$G_{d0} = 6.6667*10^{-7}$

$G_d = 0.10*G_{d0}$

$G_d = 0.25*G_{d0}$

$G_d = 0.50*G_{d0}$

$G_d = 1.0*G_{d0}$

$G_d = 2.5*G_{d0}$

$G_d = 5.0*G_{d0}$

$G_d = 10.0*G_{d0}$

under

over

Conclusions

- Without BCN, overall performance is severely degraded
  - Hogging occurs with shared as well as partitioned memory
  - Mean aggregate throughput gated by hotspot throughput
- BCN is able to control the hotspot
  - OQ steady state length exceeds target
  - OG exposes sensitivity to G_d
    - G_d too low: Underdamping => Low response speed; overall throughput suffers because hogging is not sufficiently reduced
    - G_d too high: Overdamping => Excessive throttling; hotspot throughput suffers, queue length oscillates strongly

      => narrow range of optimal G_d
  - Shared memory has worse settling time and % overshoot than partitioned memory