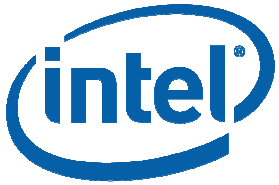

Baseline Scenario Simulation Results - PAUSE

Tanmay Gupta

19th Oct 2006

Topology & Workloads

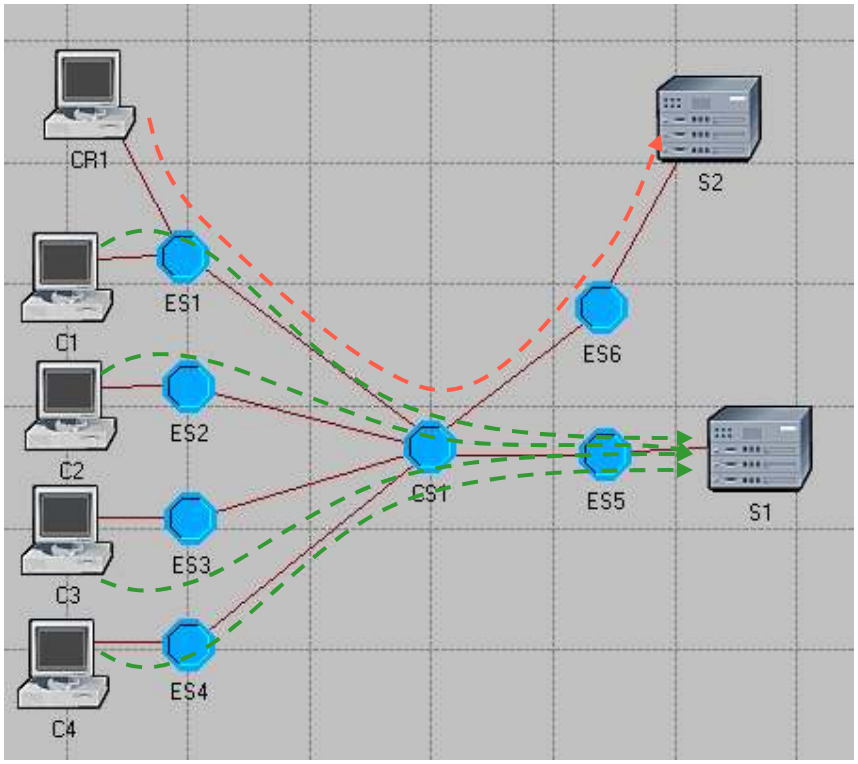

- All links 10Gbps
- Output buffered Switch with 150KB/port
- 150KB of buffering in Host, but traffic source stops after memory is full (no drops)
- Latency
  - Switch = 1us
  - Each link = 0.5us
  - Host response time = 2uS
- Sources C1, C2, C3, C4 sending ~4.8Gbps of UDP data to S1
- Reference Source CR1 sending ~4Gbps of UDP data to S2
- 1500 byte fixed payload size
- Bernoulli temporal arrival distribution
- Total run time = 100ms
- All sources start at 5ms
- 2 sources stop at 85ms

Parameters

- PAUSE parameters:
  - XON/XOFF threshold sets towards the top of the switch output port buffer
    - XOFF threshold = 136,192 bytes
    - XON threshold = 123,392 bytes
  - Global PAUSE
    - XOFF is sent to all ports except the current port when buffer >= XOFF threshold
    - XON is sent to all ports except the current port when buffer falls below XON threshold
- BCN parameters:
  - Qeq = 375 64 byte pages
  - Sampling interval = 150KB +- 20KB
  - W = 2
  - Gd = 5.3 * 10 ^-1
  - Gi = 2.6 * 10^-4

Scenarios

1. no_cm
   - No flow/congestion control mechanism (frames are dropped as a result of congestion)
2. pause
   - 802.3x PAUSE flow control only
3. bcn
   - BCN only
4. bcn_pause
   - BCN as well as 802.3x PAUSE

Congested Link Throughput & Drops

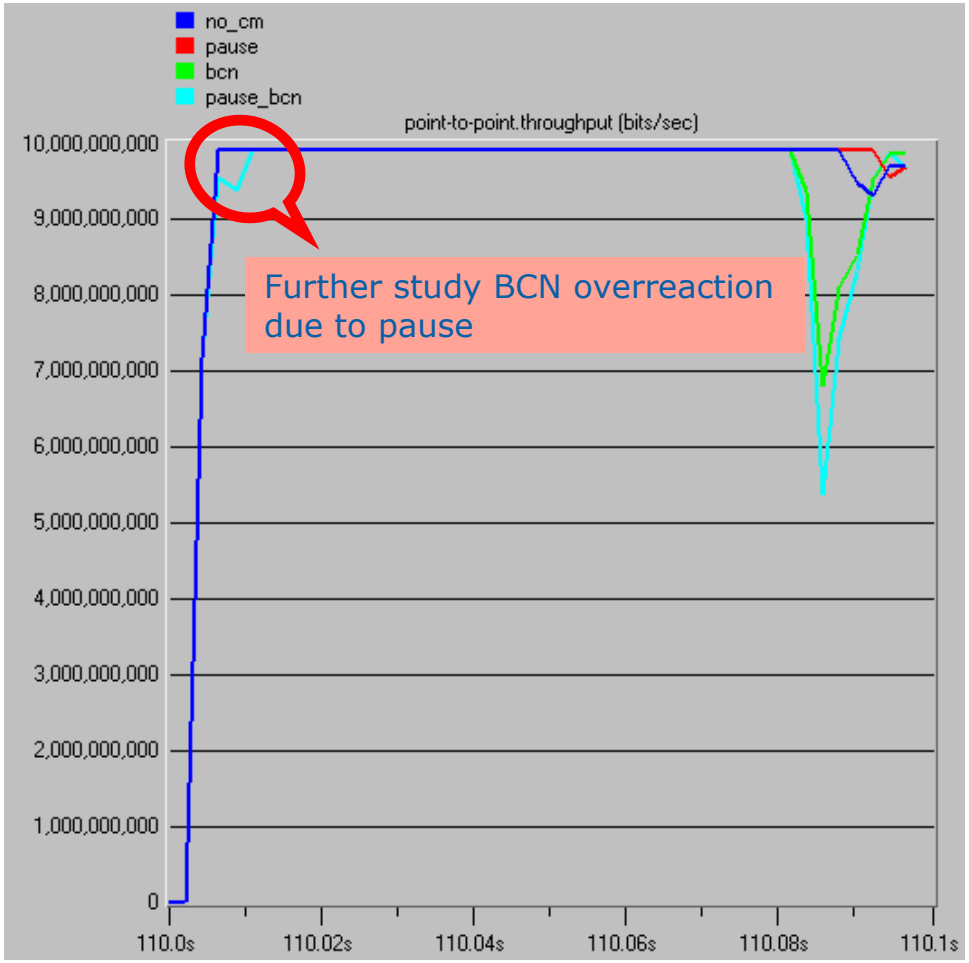

Drops:

No CM = 59,382

Pause = 0

BCN = 1,371 (initial only)

BCN Pause = 0

Pause eliminates drops
BCN reduces drops significantly as compared to no CM but cannot avoid drops
BCN+Pause eliminates drops

Uncongested Link (innocent flow) Throughput

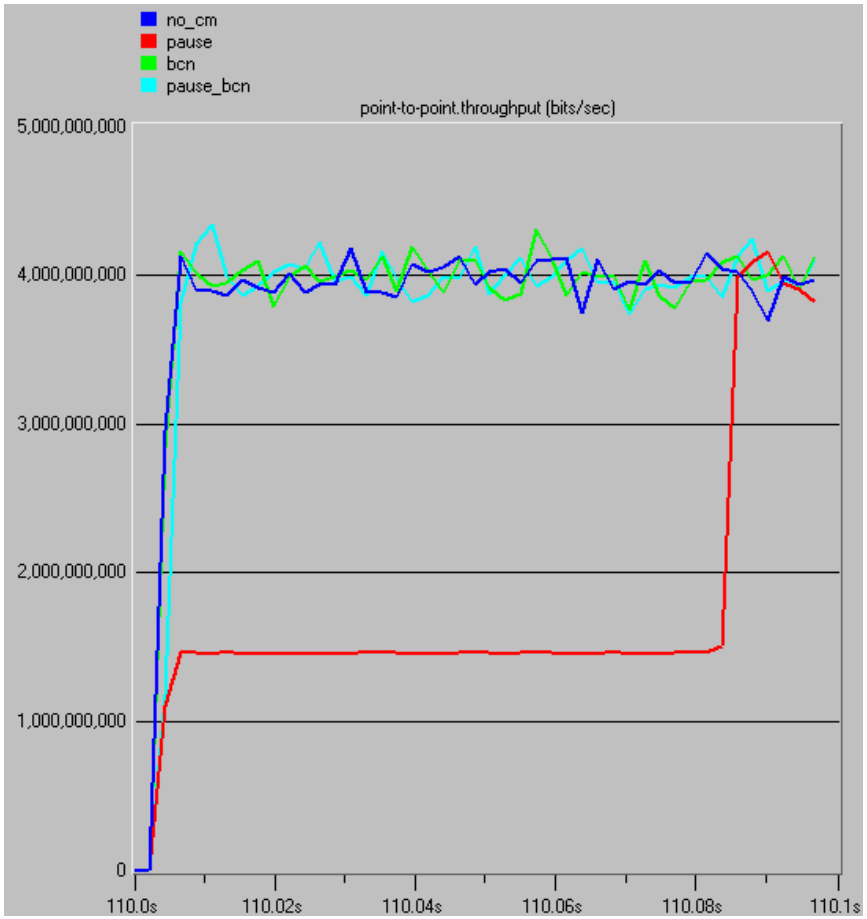

Pause results in significant congestion spreading
BCN+Pause minimizes congestion spreading

Core Switch Output Buffer Utilization (Congested Port)

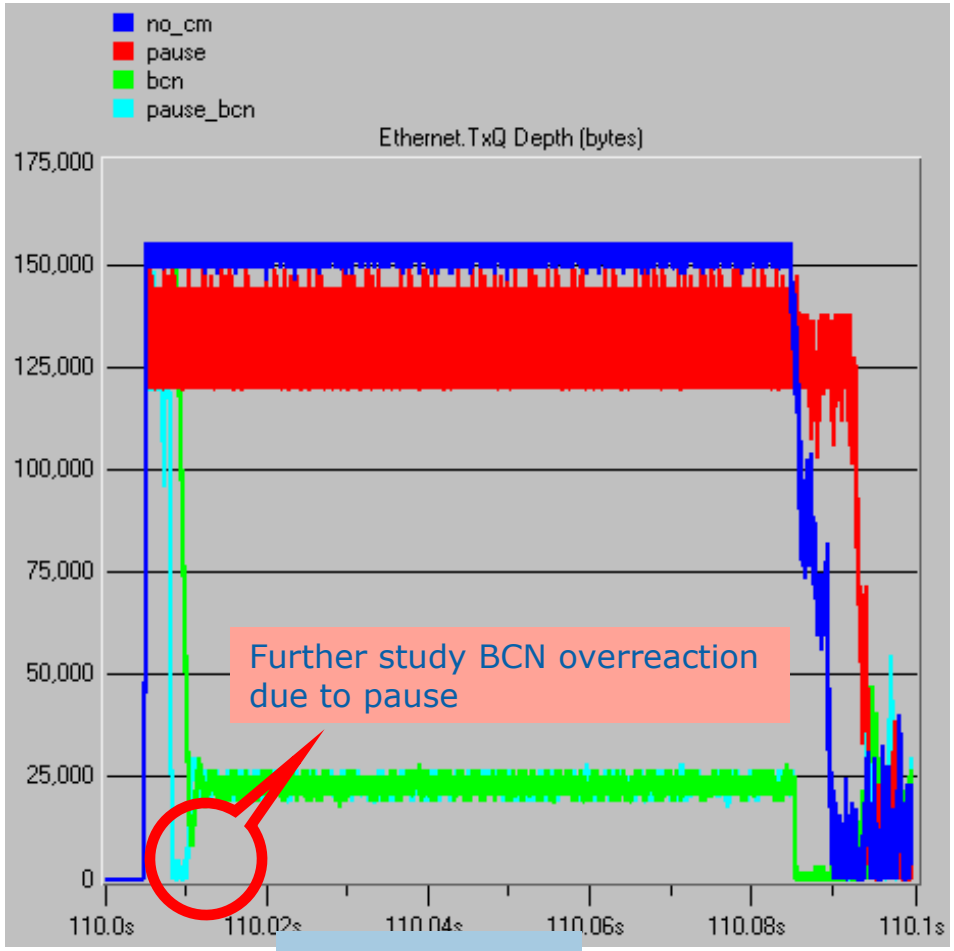

Overlaid

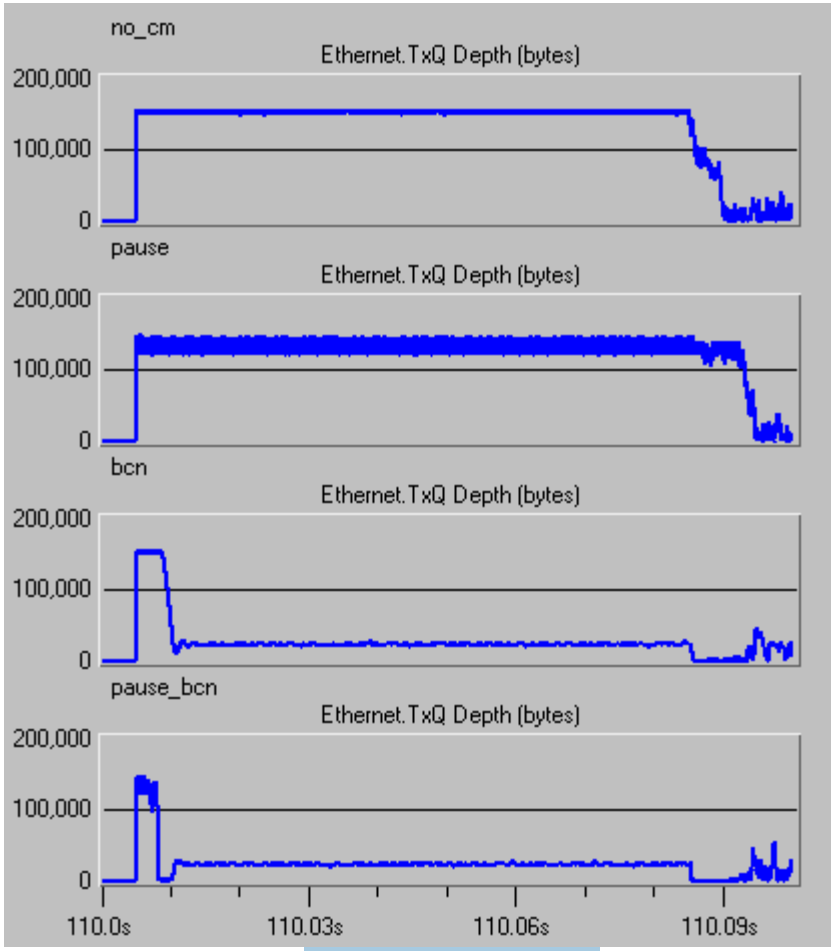

Stacked

Edge Switch Output Buffer Utilization

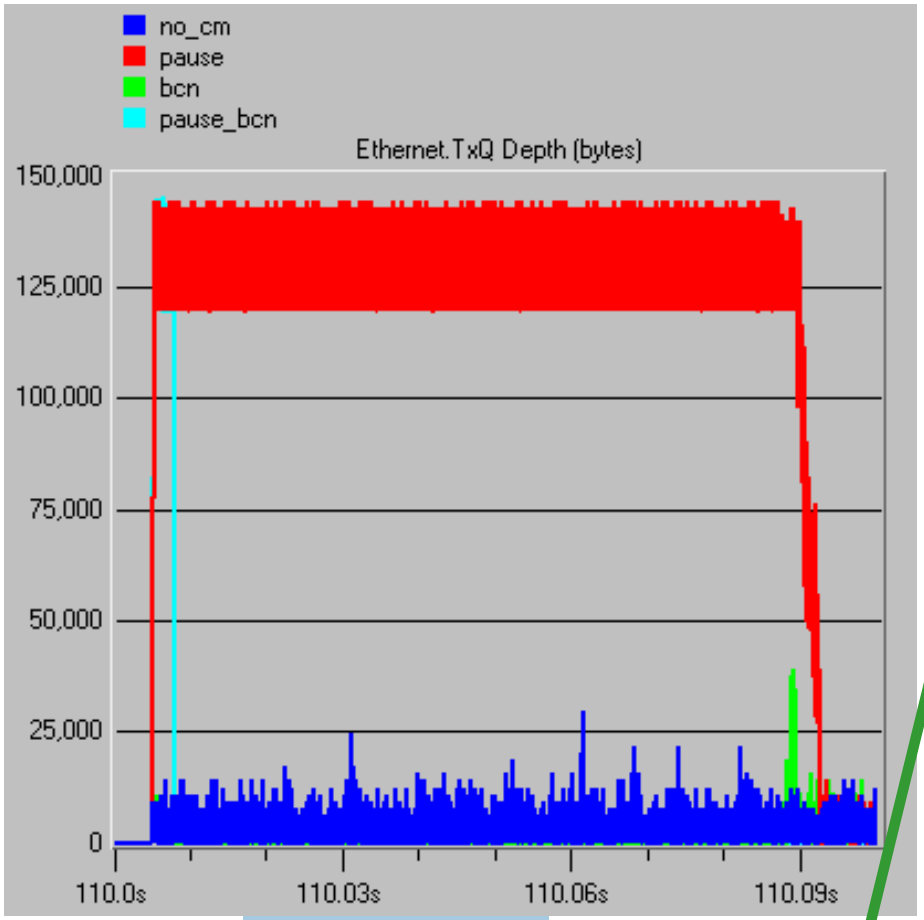

Overlaid

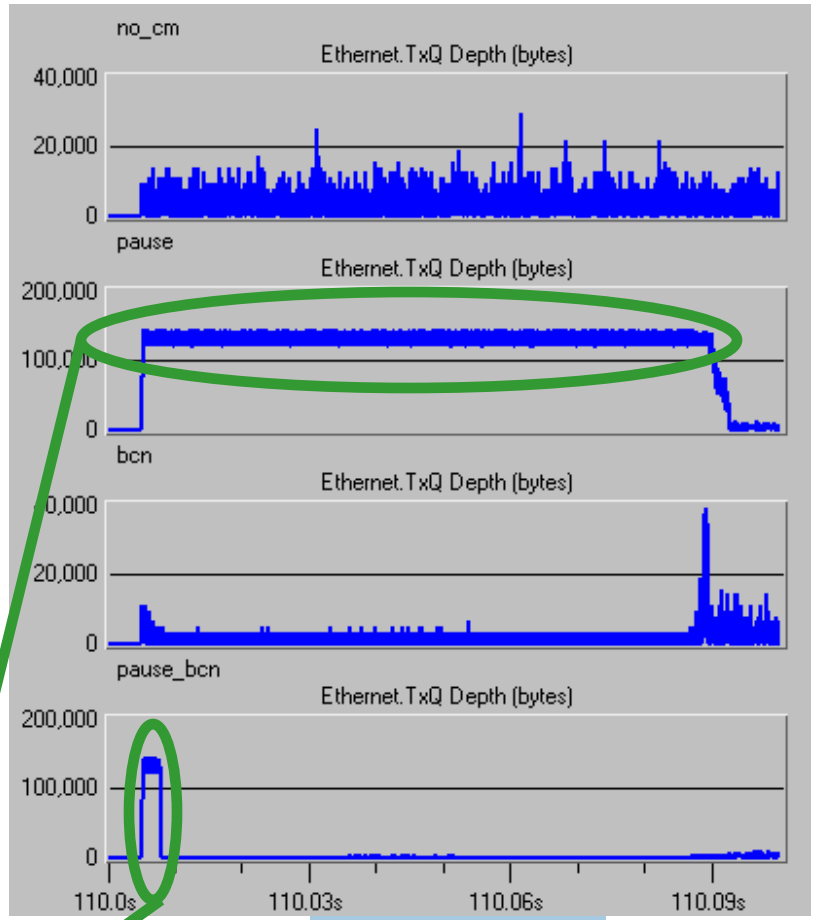

Stacked

congestion spreading

Conclusions

- PAUSE flow control provides no-drop behavior
  - However, persistent PAUSE results in head-of-line blocking and congestion spreading
- BCN reduces packet drops significantly
  - However, packet drops still happen during the transient phase
- BCN with PAUSE provides no-drop behavior
  - PAUSE is only triggered during the initial transient phase which limits the head-of-line blocking and congestion spreading
  - Need to further study BCN overreaction due to PAUSE