A Comparison of BCN and [B|F]ECN

Davide Bergamasco

IEEE 802.1 Interim Meeting
Monterey, CA
January 26th, 2007

Goals

- Conduct a simulation-based comparative analysis of BCN and [B|F]ECN
- Try to reproduce some of the results shown on Prof. Jain's Nov '06 presentation

  http://www.ieee802.org/1/files/public/docs2006/au-jain-ecn-20061115.pdf
- Not drawing any conclusions, just stating facts

Part I: BCN vs. BECN

Symmetric Topology

- Topology & Workload as per Prof. Jain's presentation

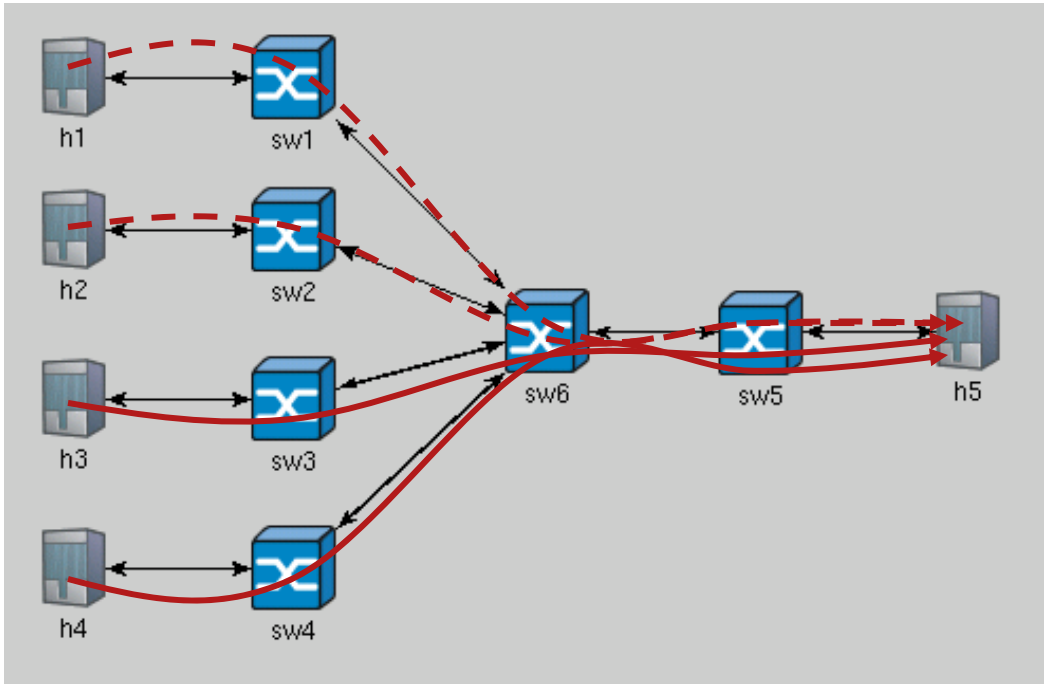

- Traffic pattern
  - Point-to-point from h1-4 to h5
  - Load: 100%
  - H1 and H2 on-off sources (T_{on} = T_{off} = 20 ms)

Symmetric Topology

- Shared-memory output-buffered switch
  - 16 Ports
  - 150 KB of space per output port
- Global Pause enabled on shared buffer
  - Assert threshold 140 KB
  - De-assert threshold 130 KB
- BCN parameters as per Baseline Scenario
  - W = 2
  - Qeq = 375 64-byte pages (24 KB)
  - Gi = 5.3333 x 10-1
  - Gd = 2.6667 x 10-4
  - Sampling rate = 1% (150 KB)
  - No BCN-Max or BCN(0,0)
  - No Over-sampling
  - No self increase
- BECN parameters as per Prof. Jain's presentation
  - Measurement interval = 30 µs
  - Qeq = 375
  - Queue Control Function: Hyperbolic
    - a = 1.05
    - b = 1.2
    - c = 0.5
- Simulation duration 100 ms

Symmetric Topology: Throughput

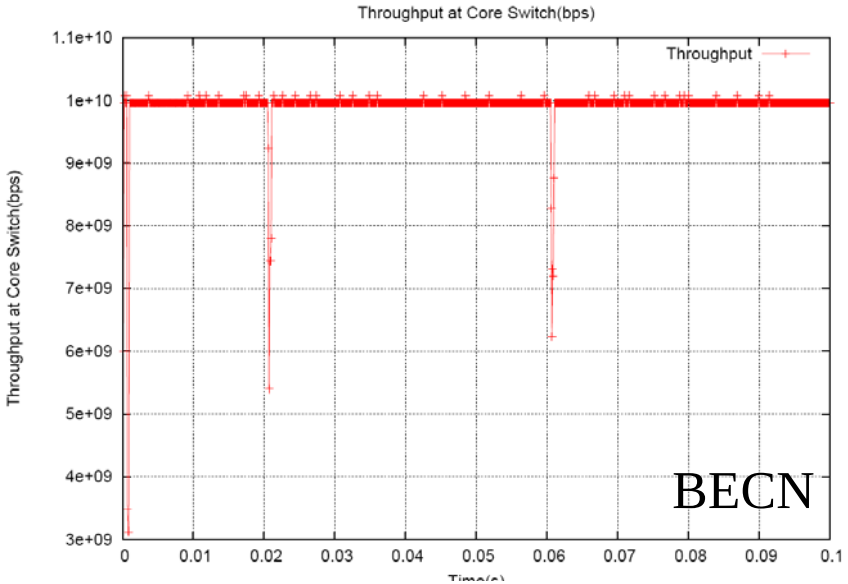

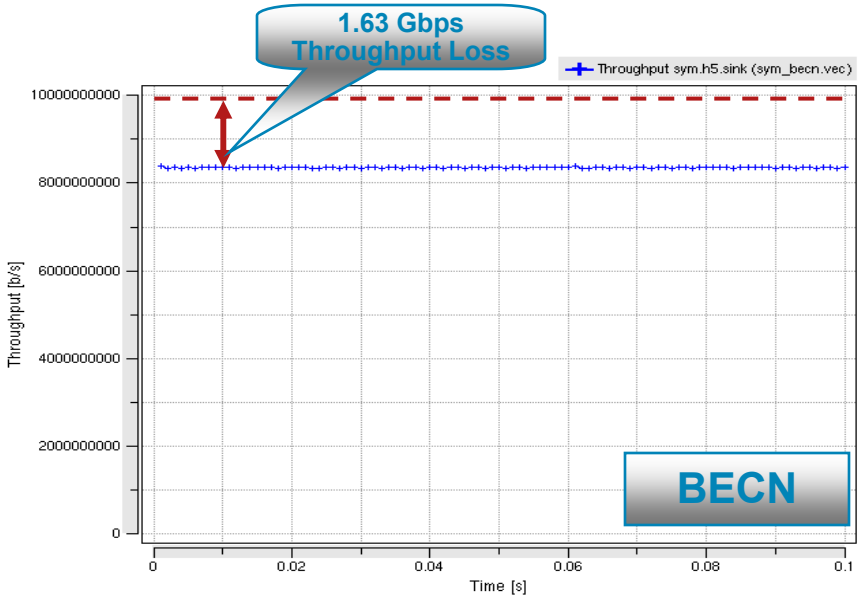

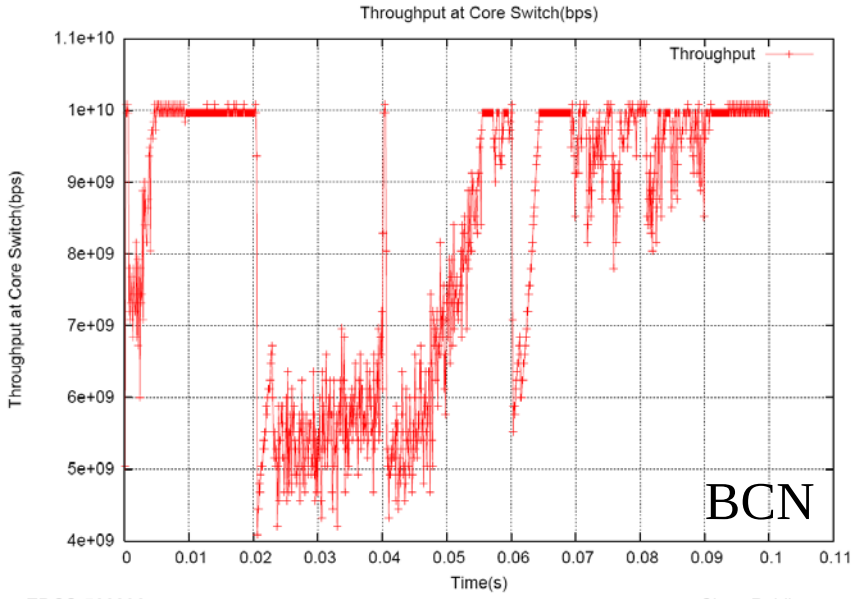

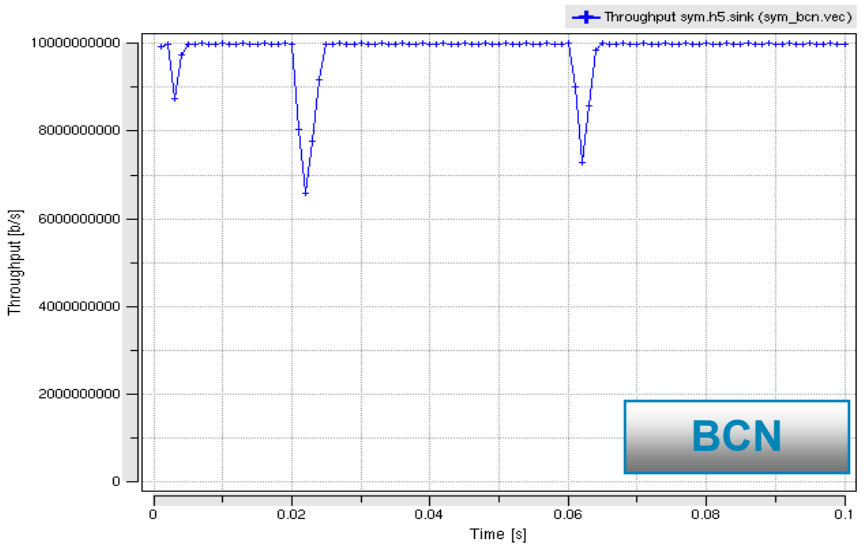

Symmetric Topology: Overhead

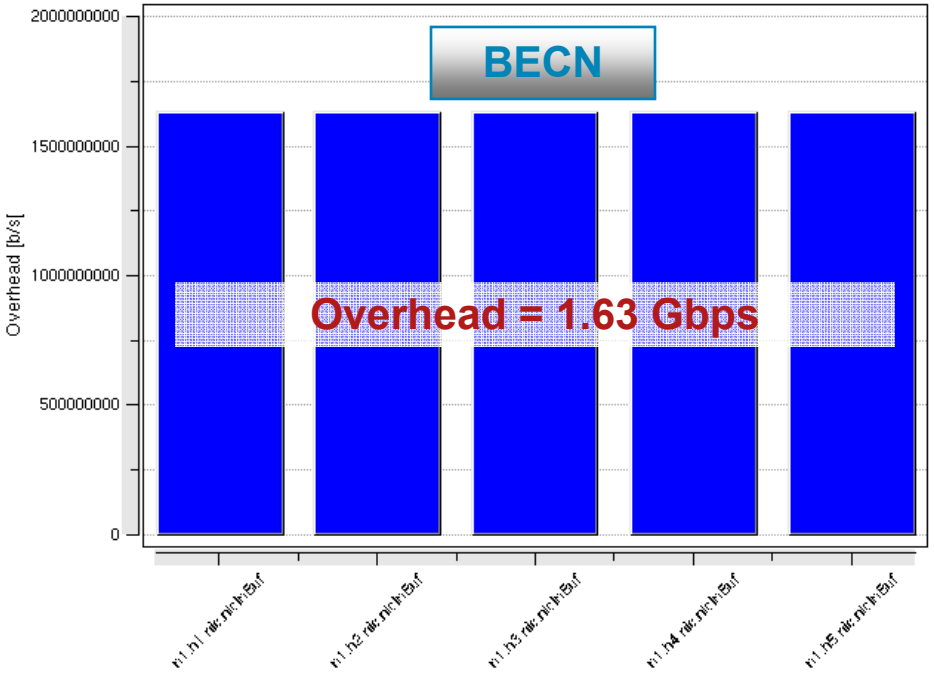

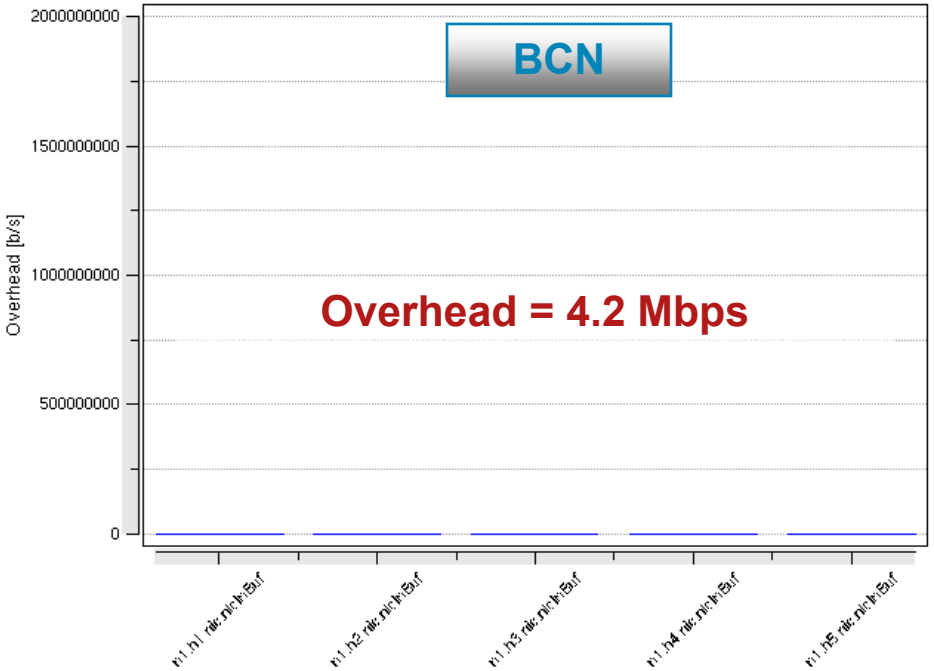

- ECN: Signaling overhead is responsible for loss of throughput
- BCN: Overhead is approx **3 orders** of magnitude smaller

Symmetric Topology: Queue Length

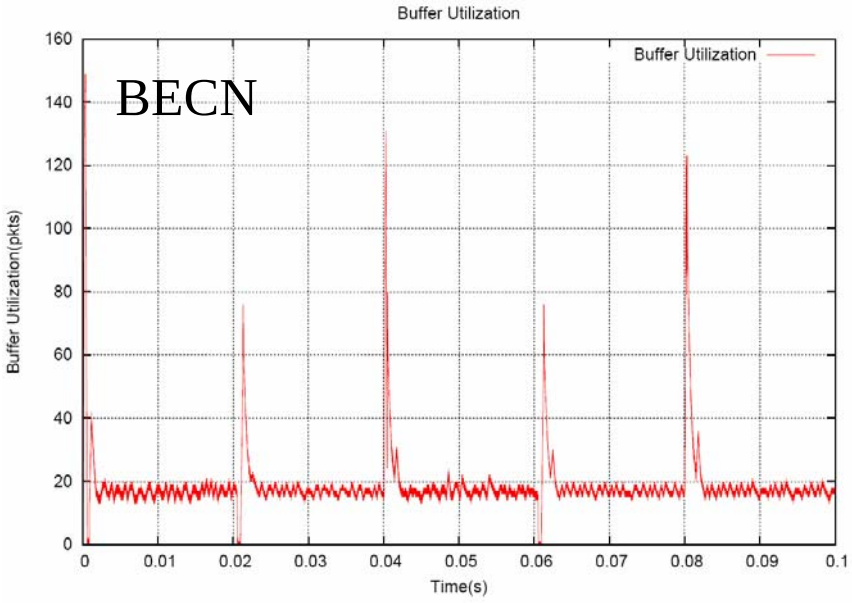

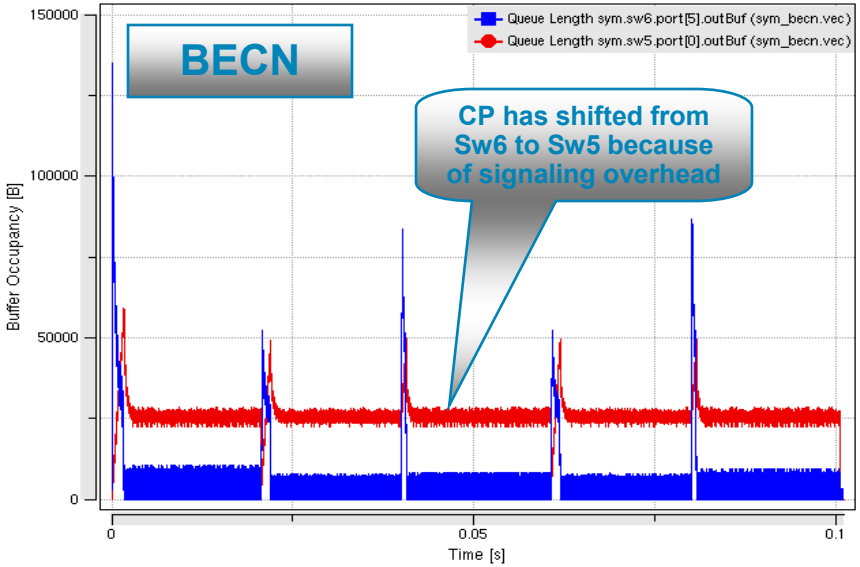

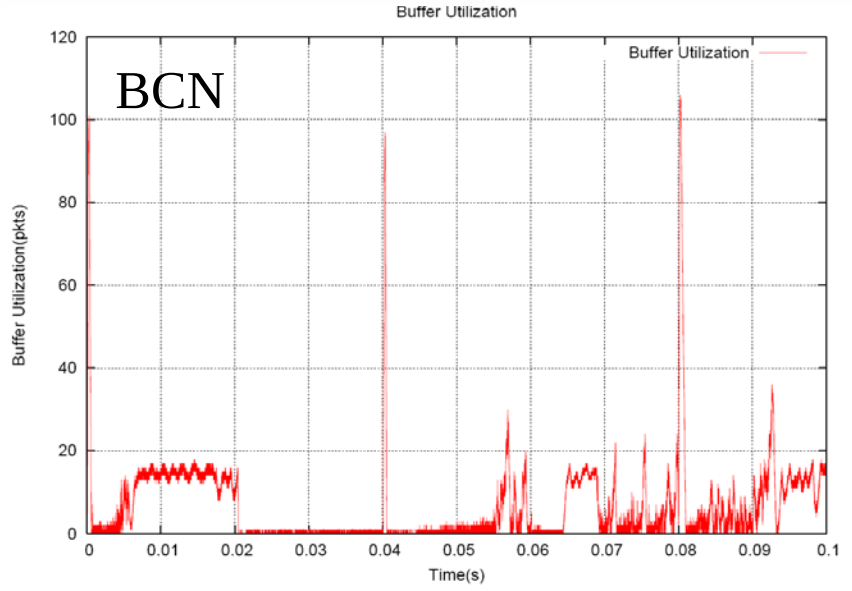

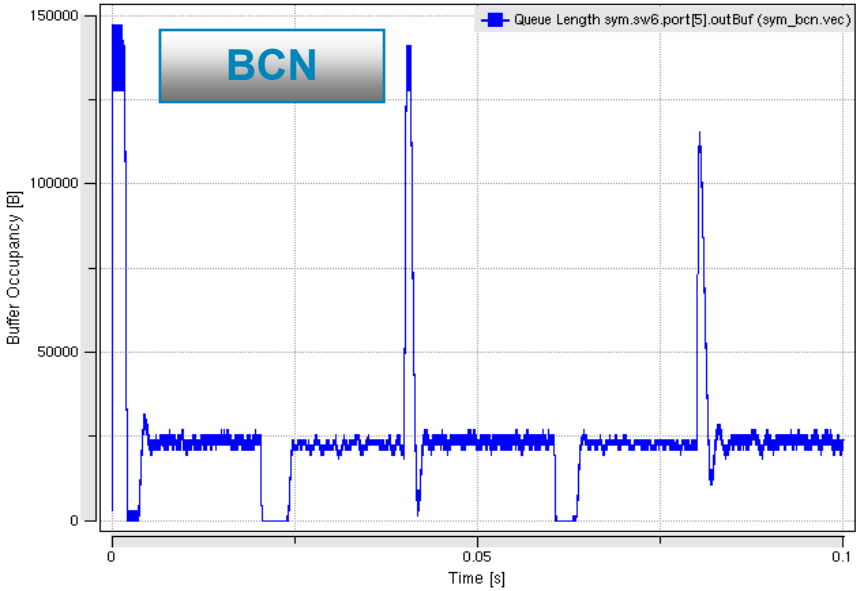

Symmetric Topology: Fairness

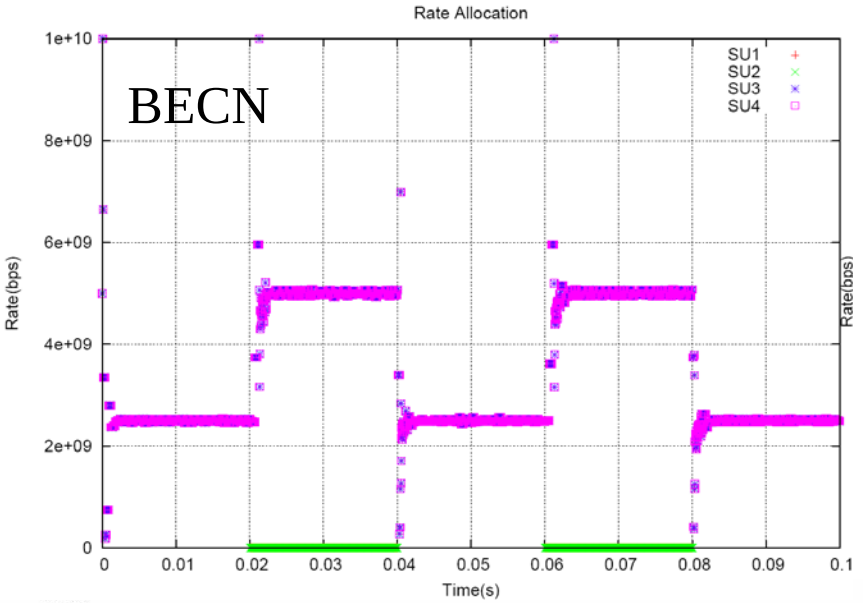

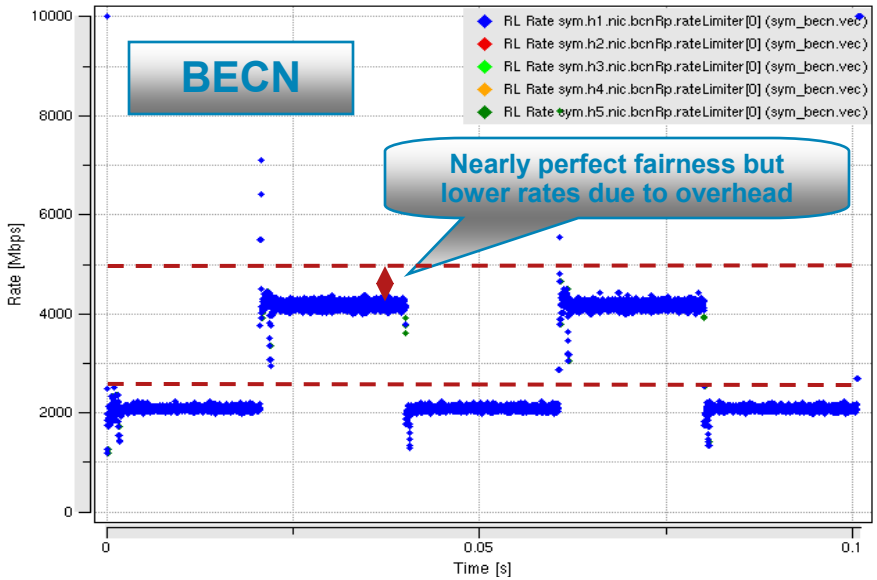

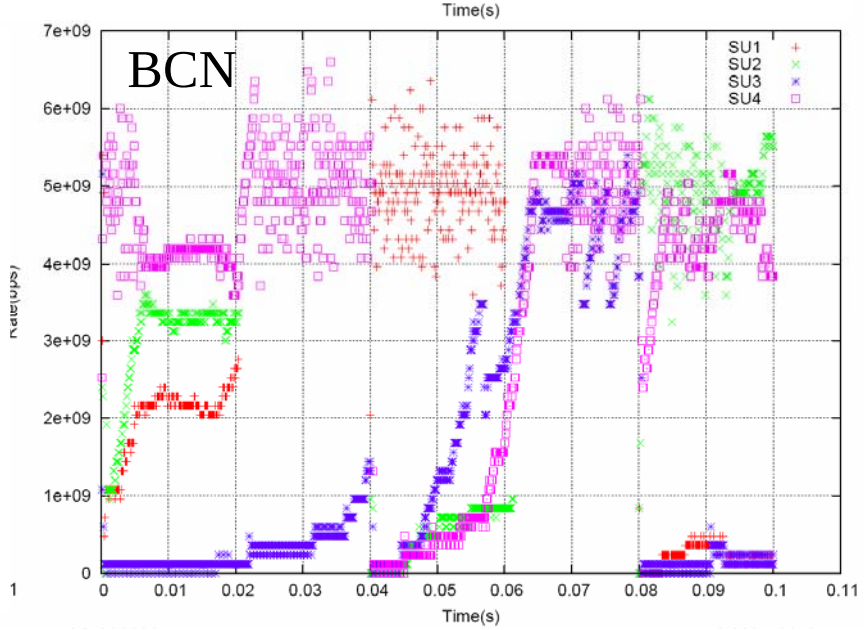

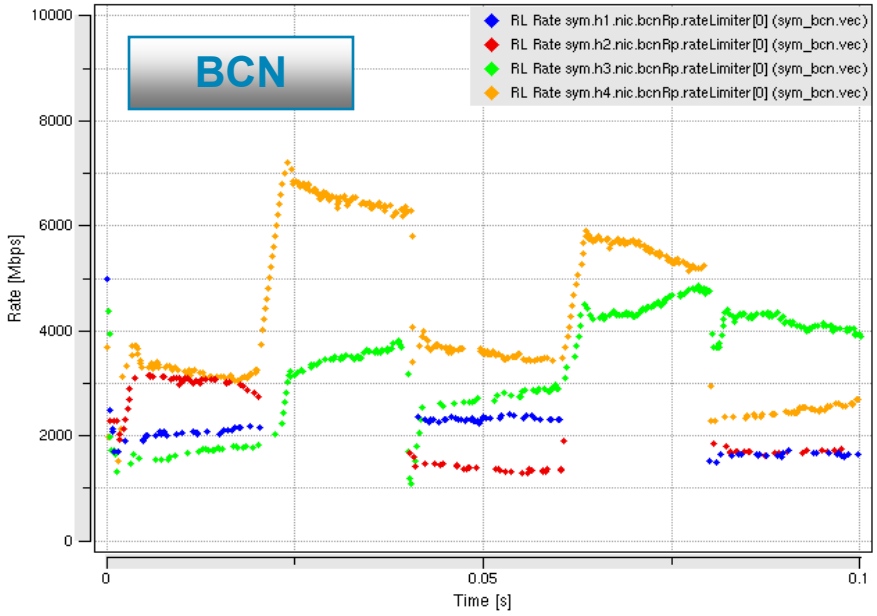

Symmetric Topology: Observations

- BECN seems to be trading **throughput** for **fairness**
- Nearly perfect fairness is achieved through a large amount of broadcast signaling traffic
- Such overhead consumes sizeable bandwidth resulting in reduced throughput
- What if BECN overhead is reduced to match BCN overhead?

Symmetric Topology: Reducing Overhead 400x

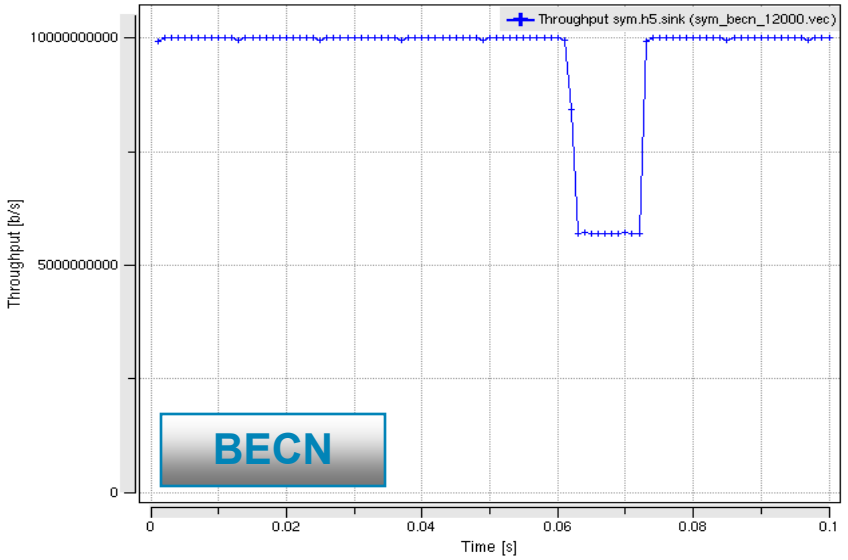

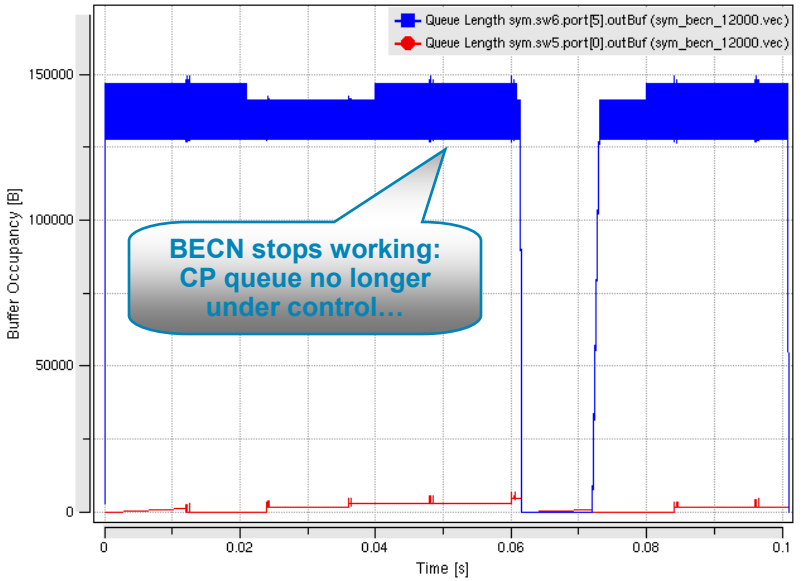

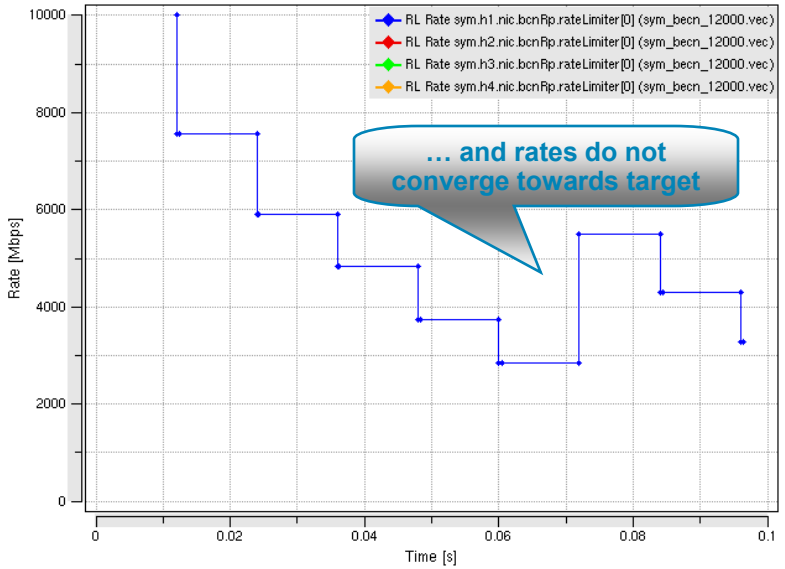

Measurement Interval = 12 ms

Aggregate Overhead = 4.07 Mbps

Asymmetric Topology

- Topology & Workload as per Prof. Jain's presentation

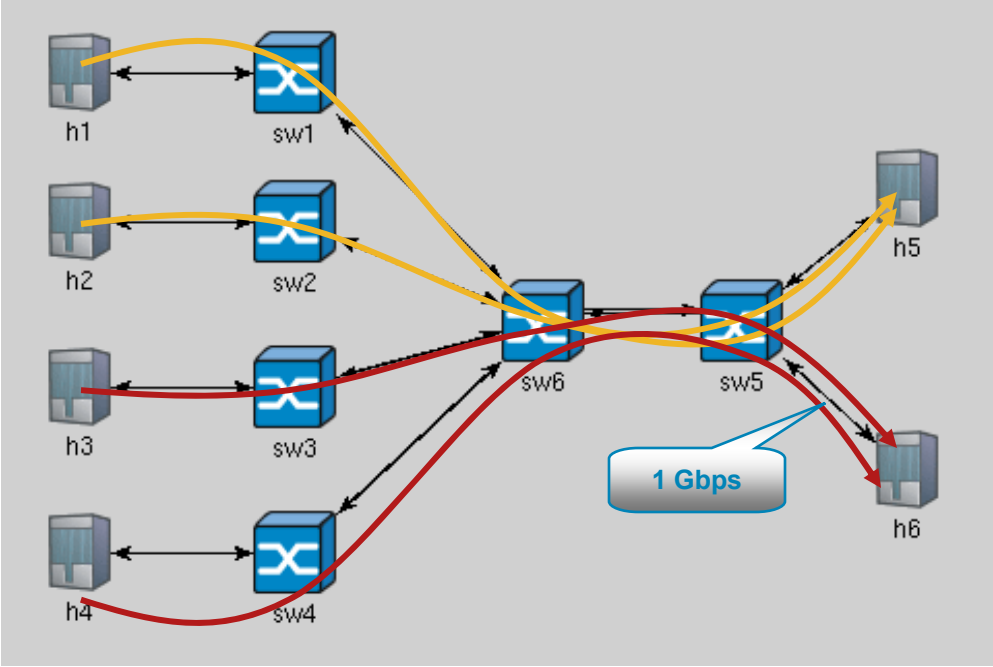

- Traffic pattern
  - Point-to-point from h1-2 to h5
  - Point-to-point from h3-4 to h6
  - Load: 100%

Asymmetric Topology

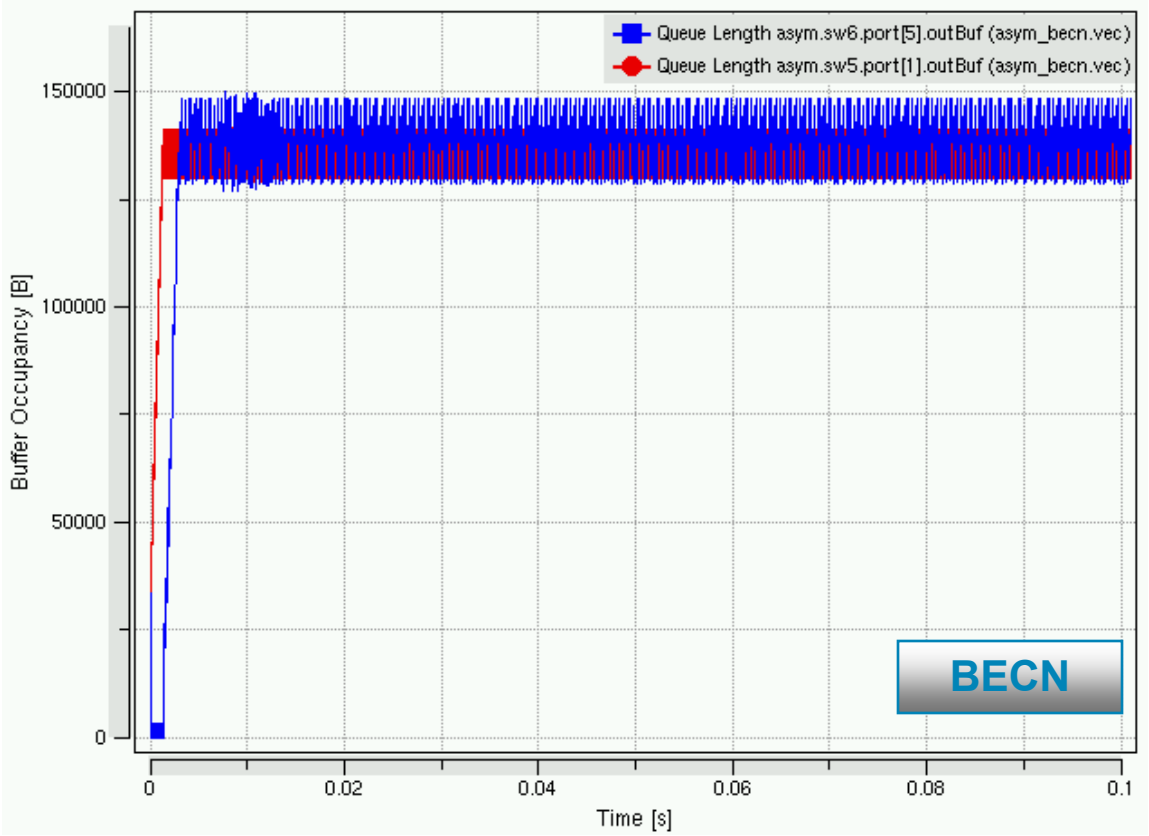

- Original BECN experiment is not reproducible
  - With T = 30 µs signaling overhead (1.63 Gbps) exceeds capacity of 1 Gbps link
  - Congestion spreading and associated collapse occurs
- Need to increase T to 300 µs

Asymmetric Topology

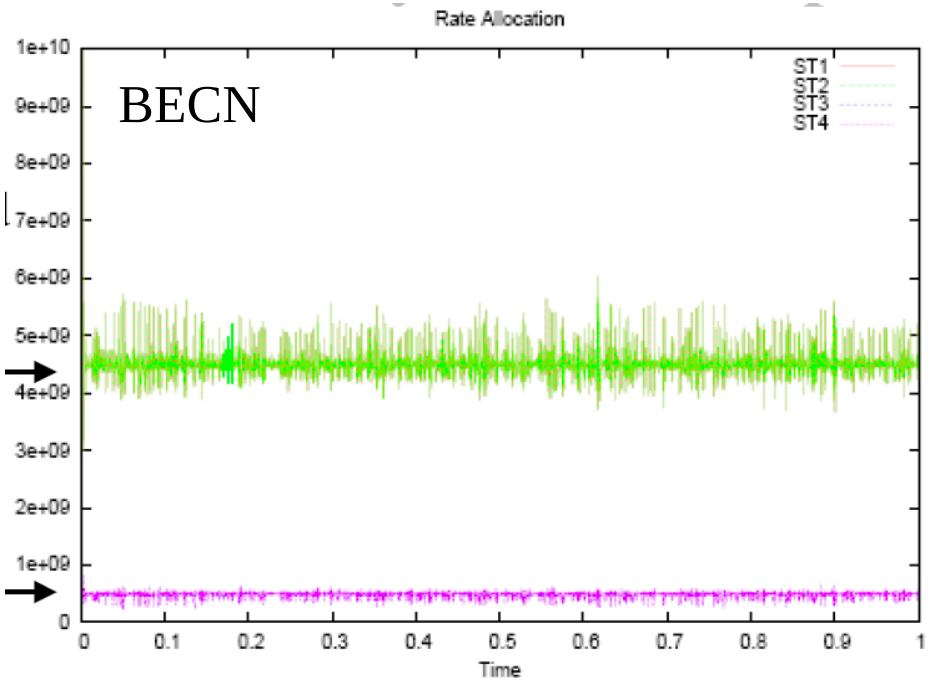

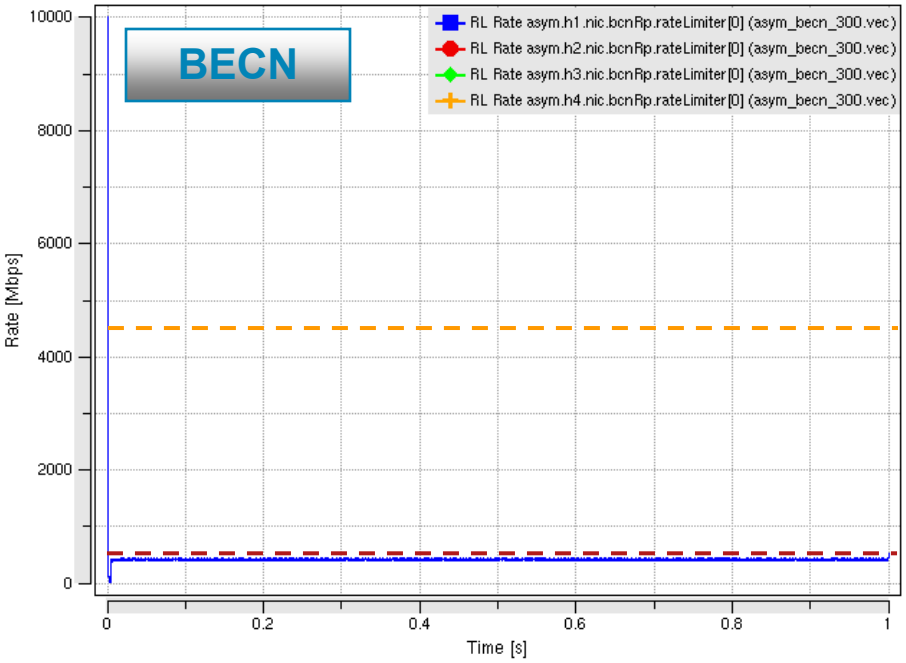

- Even increasing T to 300 µs the original experiment is still not reproducible
  - 4 flows settle around 418 Mbps [(1000 Mbps – 163 Mbps OH) / 2 = 418 Mbps]
- Why?
  - BECN rate is broadcasted by CPs to all RPs
  - The most severe CP dominates all RPs
  - Broadcasting feedback **does not work** in presence of multiple hotspots

Asymmetric Topology: Throughput

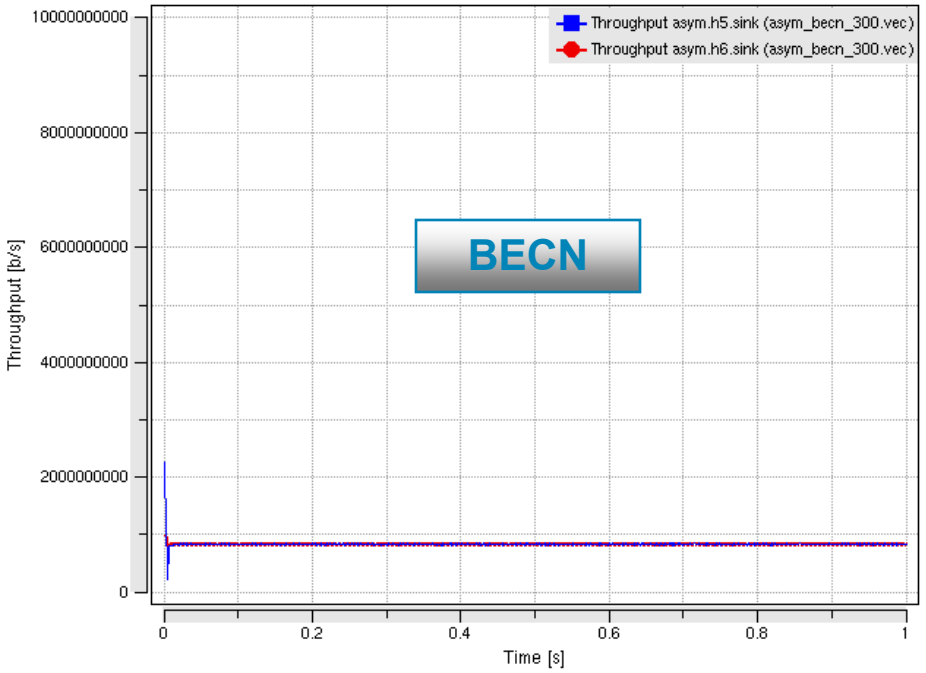

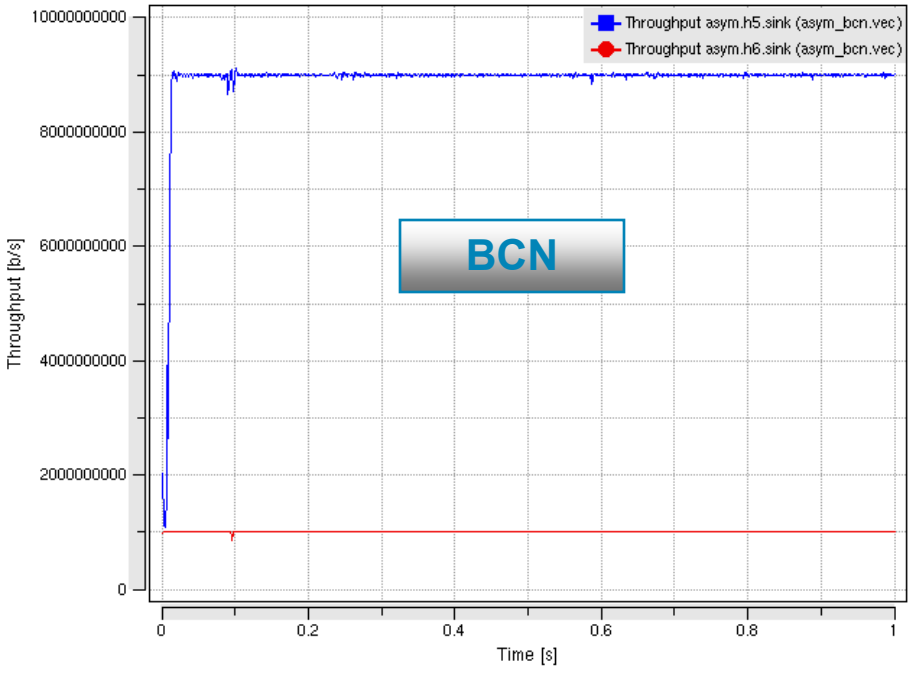

BECN Aggregate Throughput

2 * 836 Mbps = 1.672 Gbps

BCN Aggregate Throughput

~ 10 Gbps

Asymmetric Topology: Queue Length

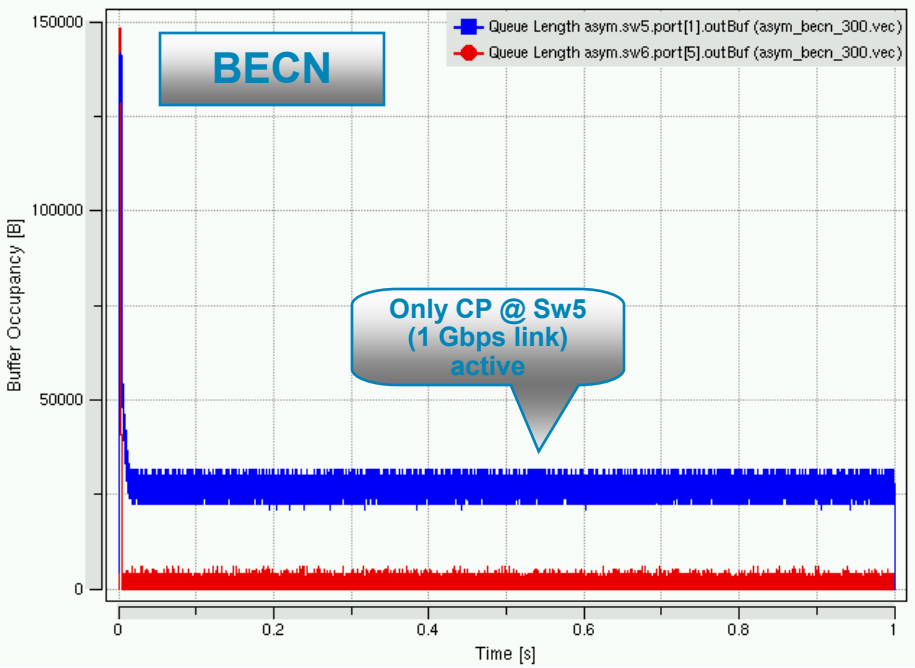

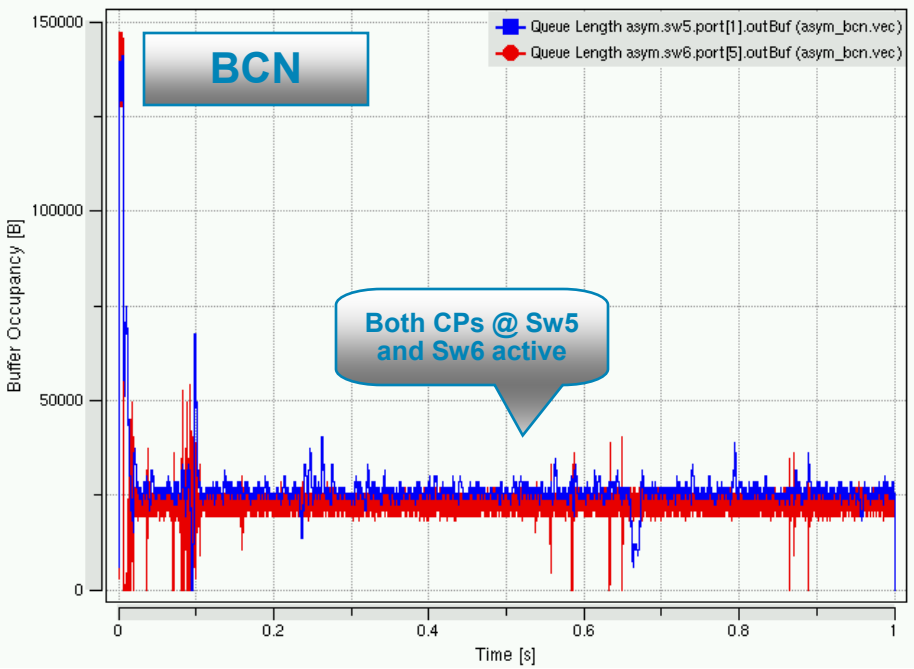

Asymmetric Topology: BCN Rates

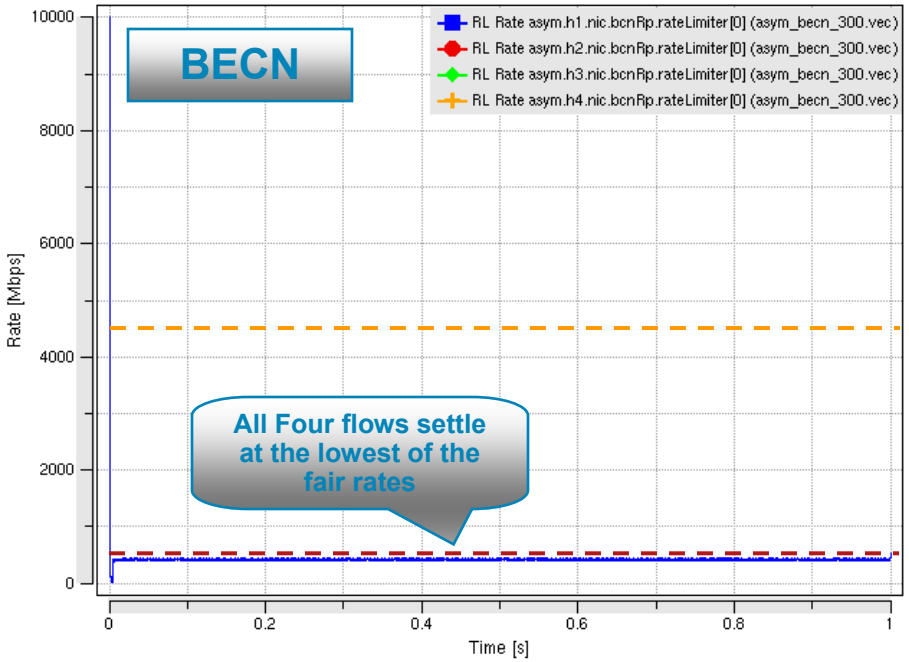

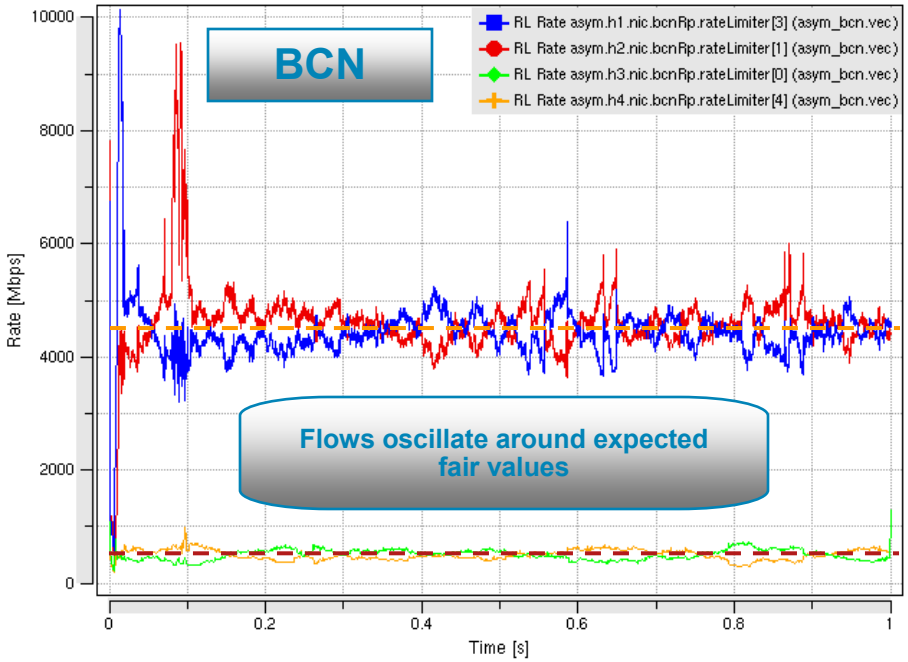

Part II: BCN vs. FECN

Parking Lot Topology

- Topology & Workload as per Prof. Jain's presentation

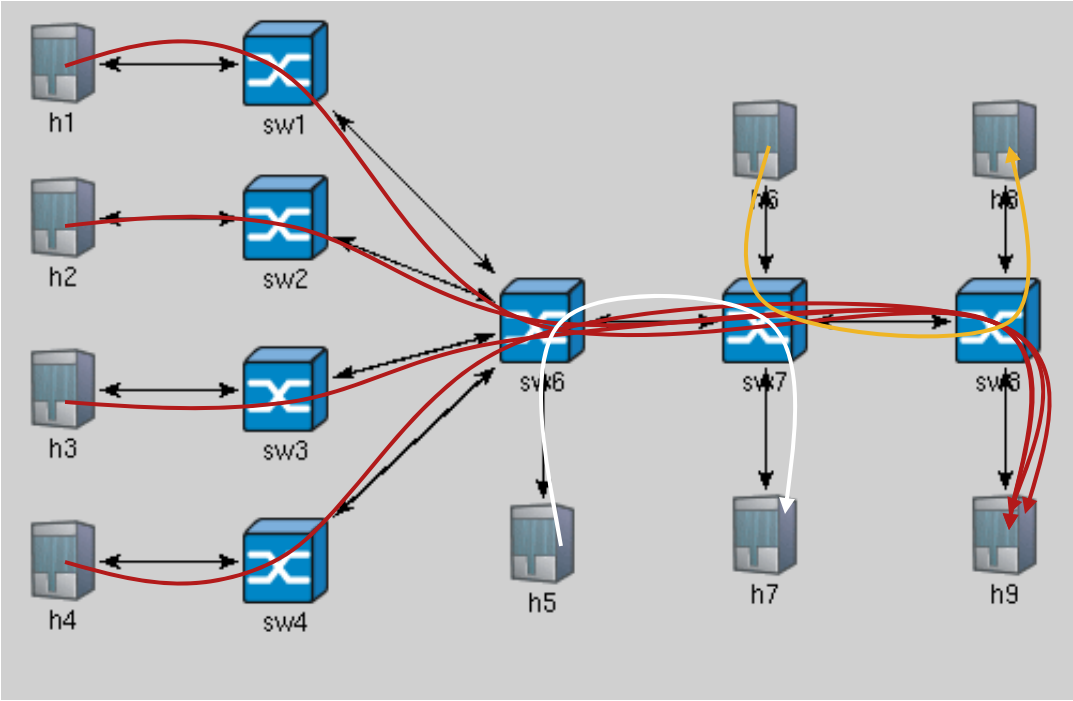

- Traffic pattern
  - Point-to-point from h1-4 to h9
  - Point-to-point from h5 to h7
  - Point-to-point from h6 to h8
  - Load: 100%

Parking Lot Topology

- Shared-memory output-buffered switch
  - 16 Ports
  - 150 KB of space per output port
- Global Pause enabled on shared buffer
  - Assert threshold 140 KB
  - De-assert threshold 130 KB
- BCN parameters as per Baseline Scenario
  - W = 2
  - Qeq = 375 64-byte pages (24 KB)
  - Gi = 5.3333 x 10-1
  - Gd = 2.6667 x 10-4
  - Sampling rate = 1% (150 KB)
  - No BCN-Max or BCN(0,0)
  - No Over-sampling
- FECN parameters as per Prof. Jain's presentation
  - Measurement interval = 30 μs
  - Qeq = 375
  - Queue Control Function: Hyperbolic
    - a = 1.05
    - b = 1.2
    - c = 0.5
  - Tagging Rate = 100%

Simulation duration 100 ms

Parking Lot Topology: FECN

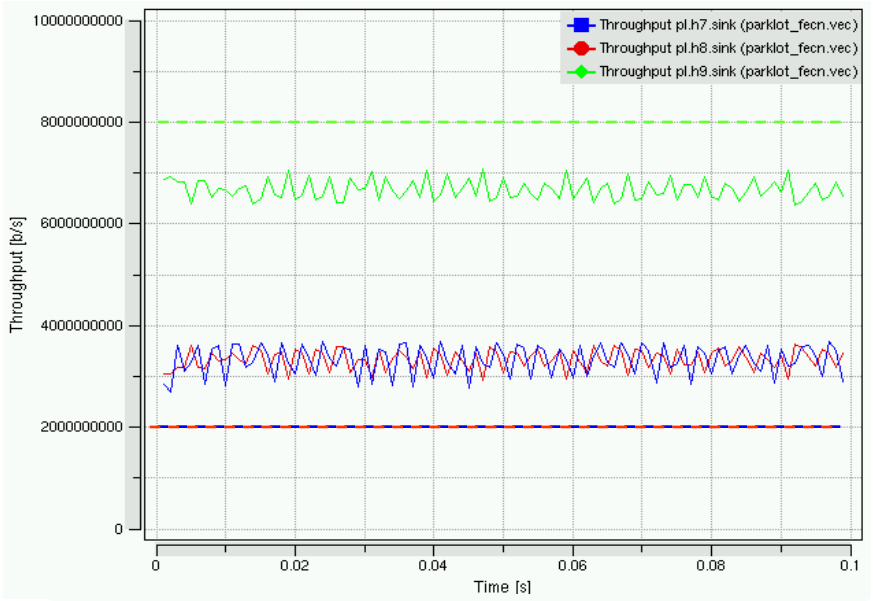

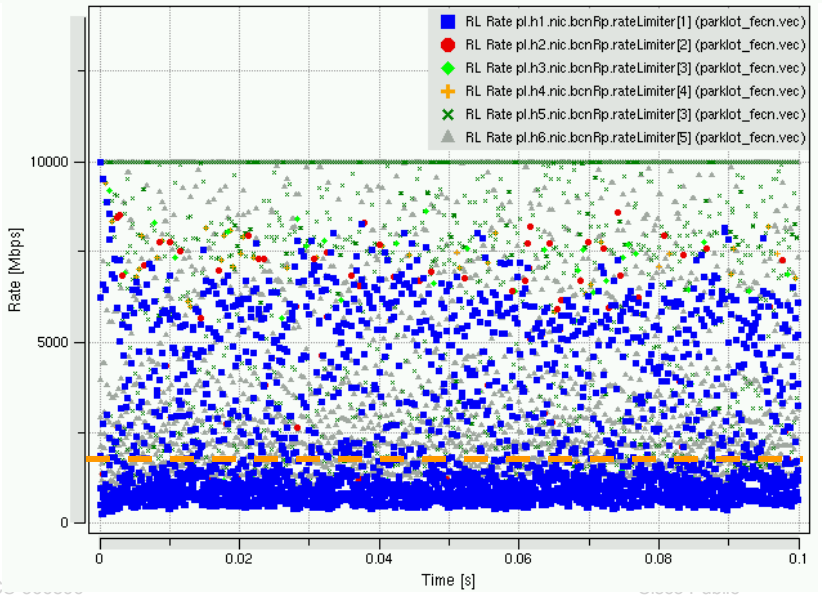

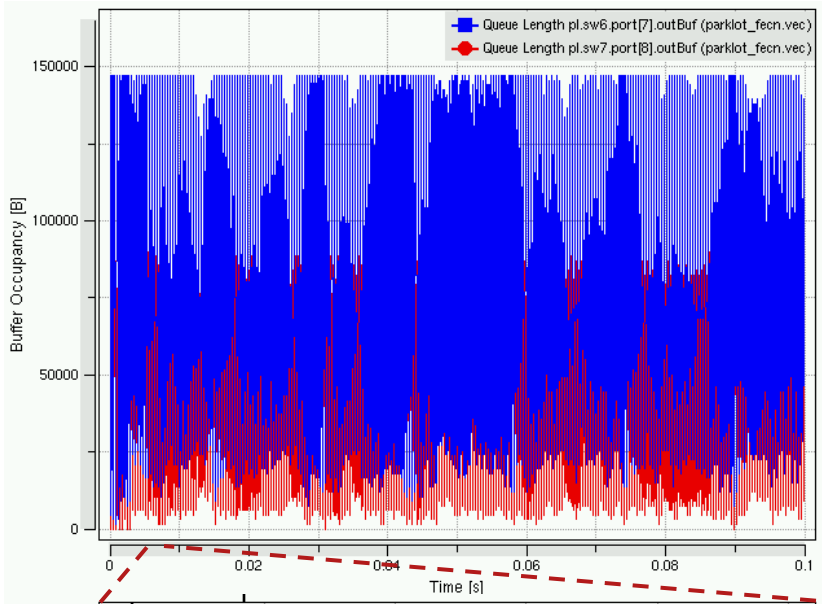

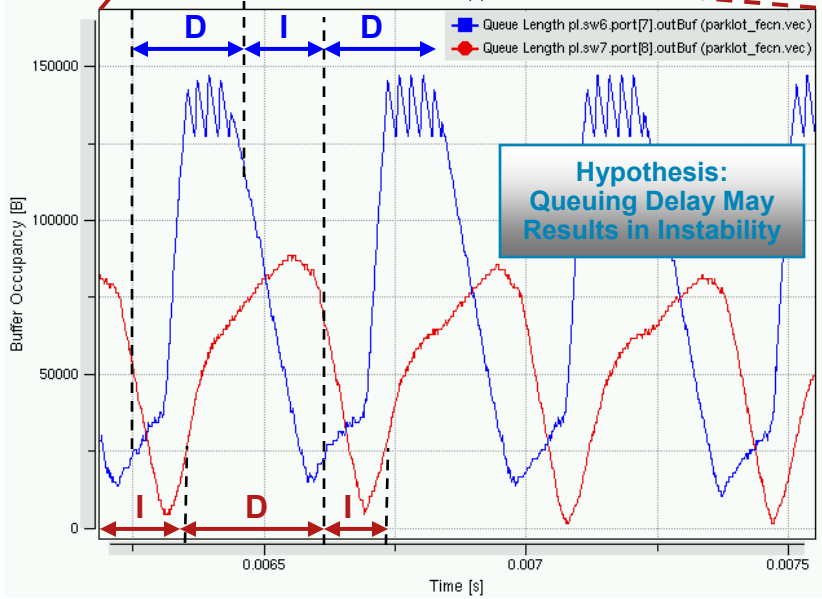

Parking Lot Topology: FECN

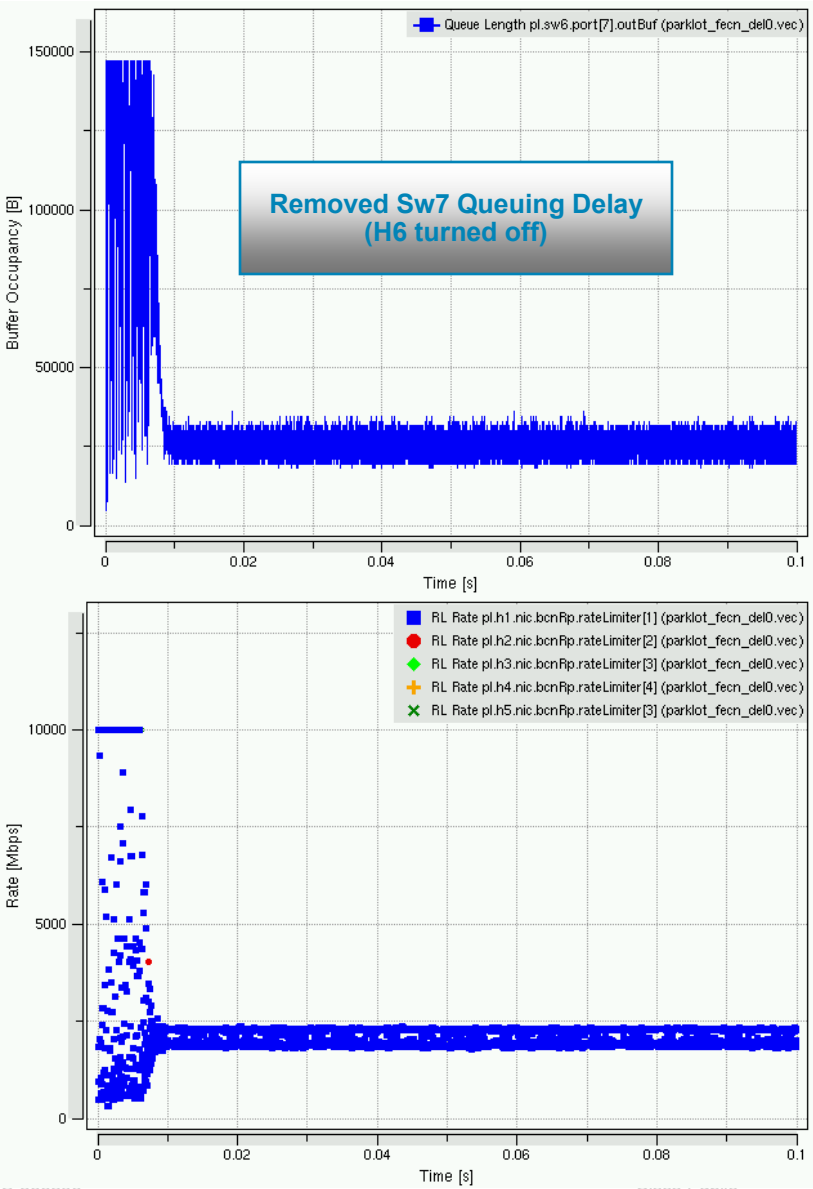

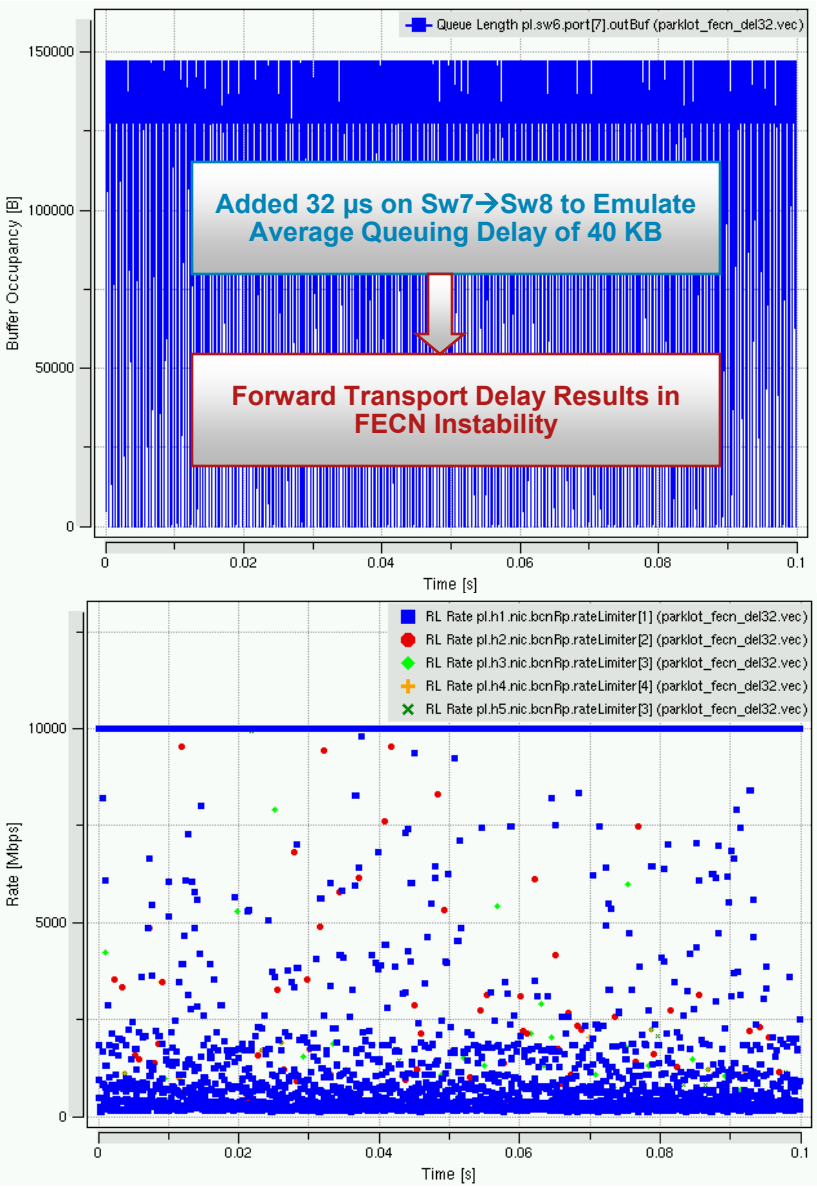

Parking Lot Topology: BCN + Delay

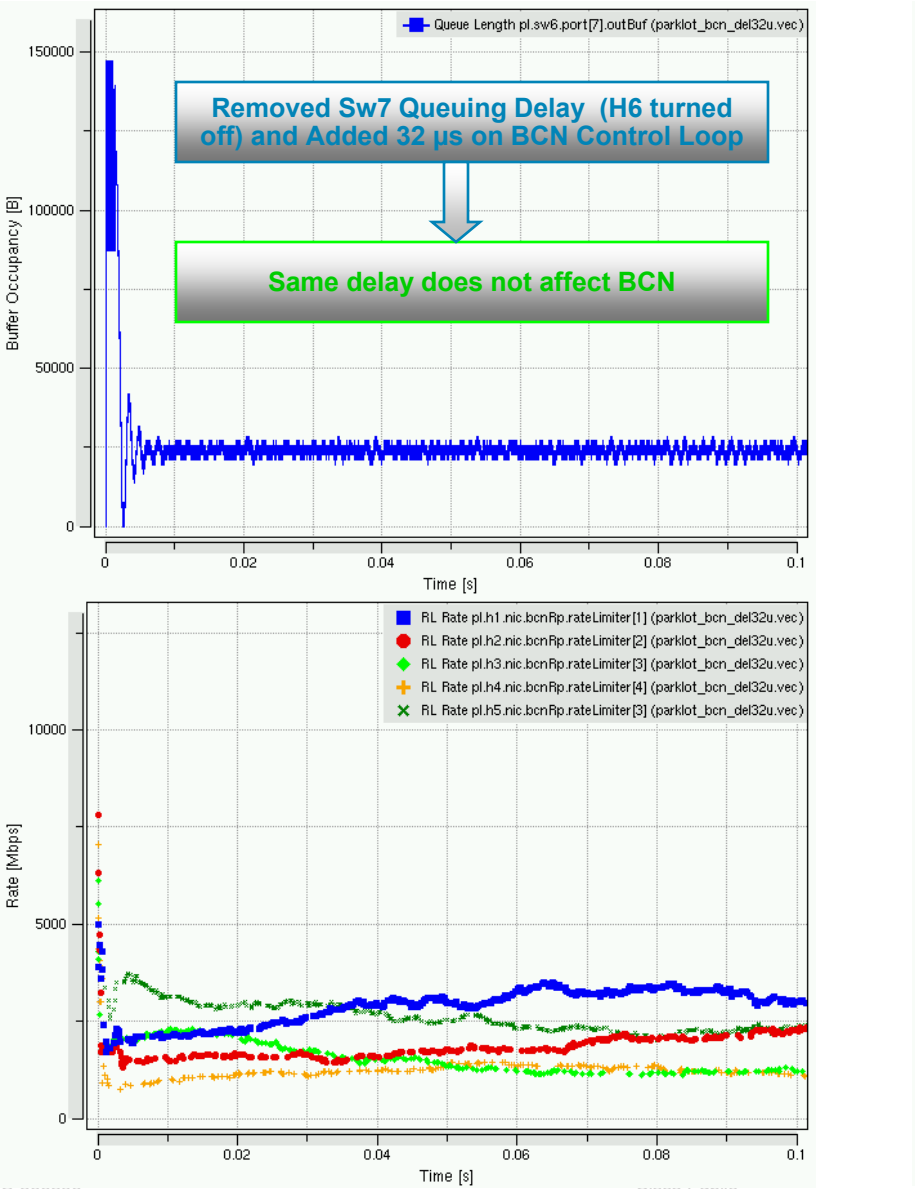

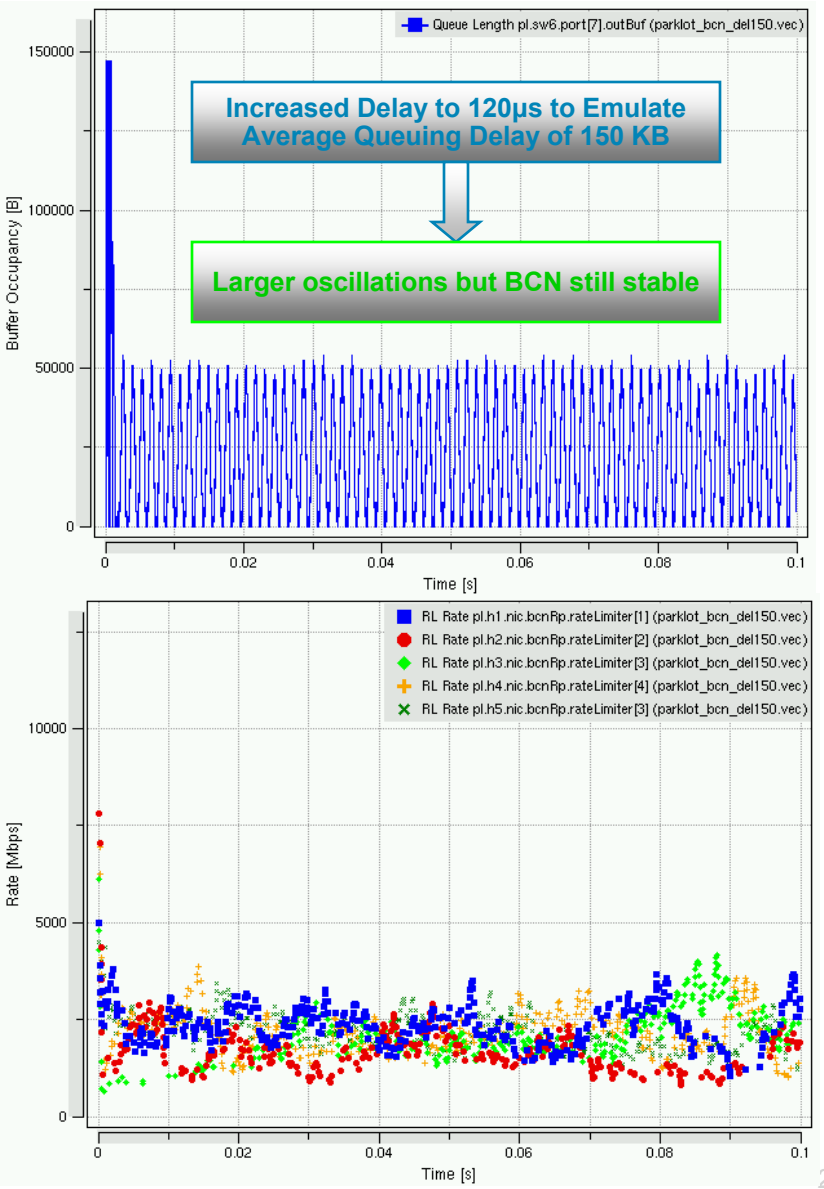

Parking Lot Topology: BCN + Delay

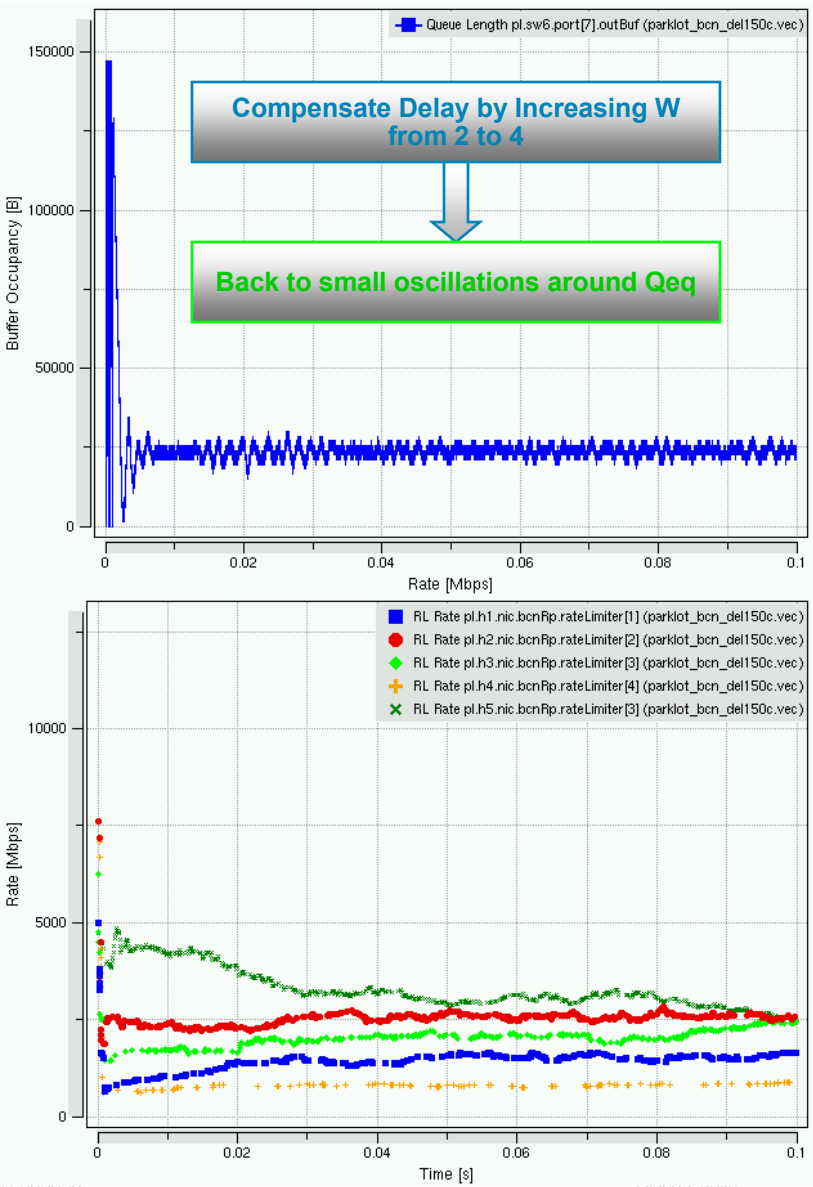

- BCN allows for delay compensation by adjusting the weight of the derivative component *W*

Symmetric Topology

- Topology & Workload as per Prof. Jain's presentation

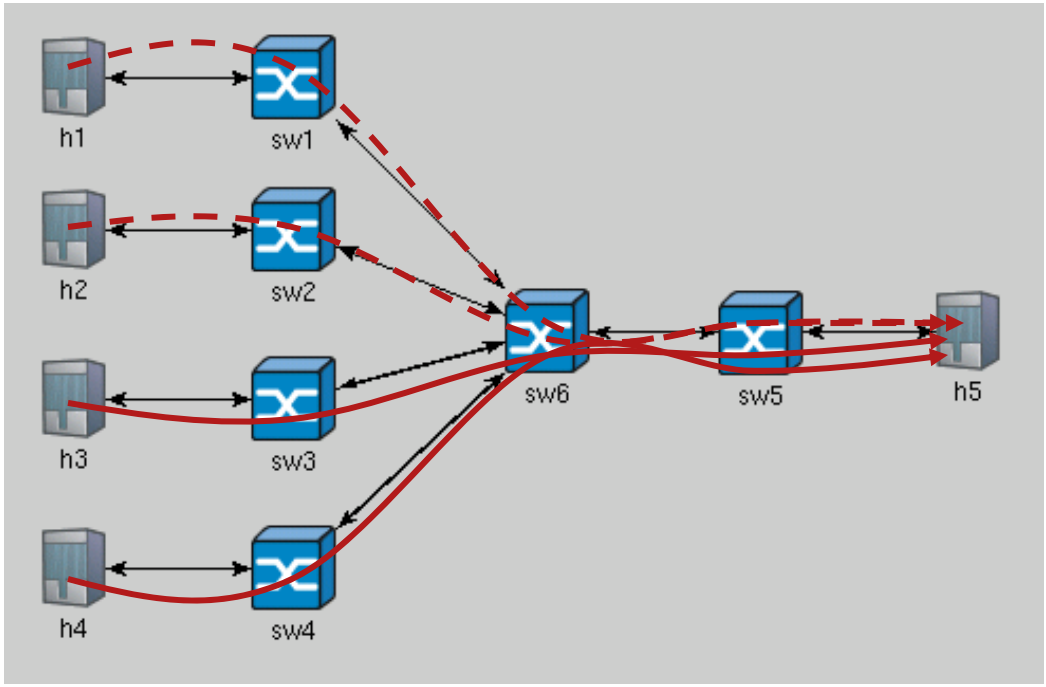

- Traffic pattern
  - Point-to-point from h1-4 to h5
  - Load: 100%
  - H1 and H2 on-off sources (T_{on} = T_{off} = 20 ms)

Symmetric Topology: FECN Tagging Rate = 100%

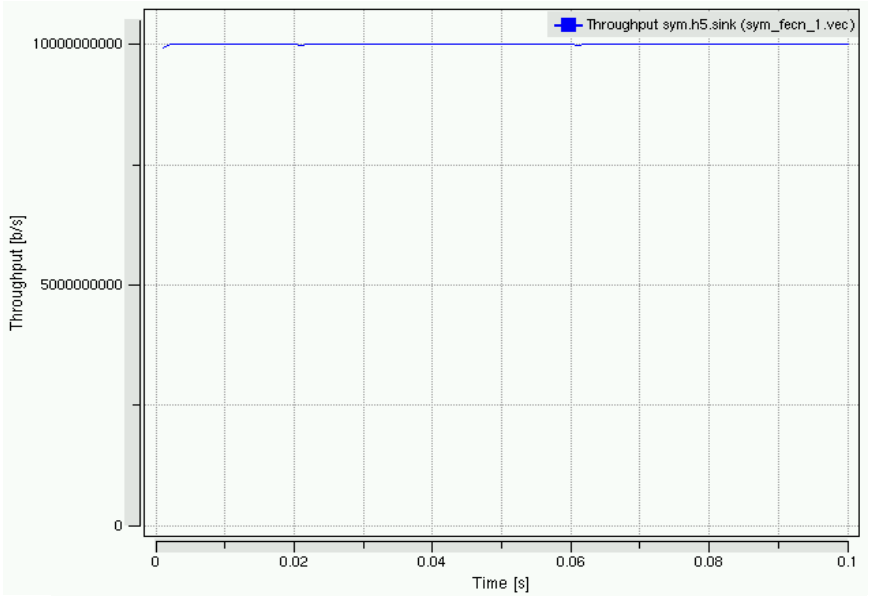

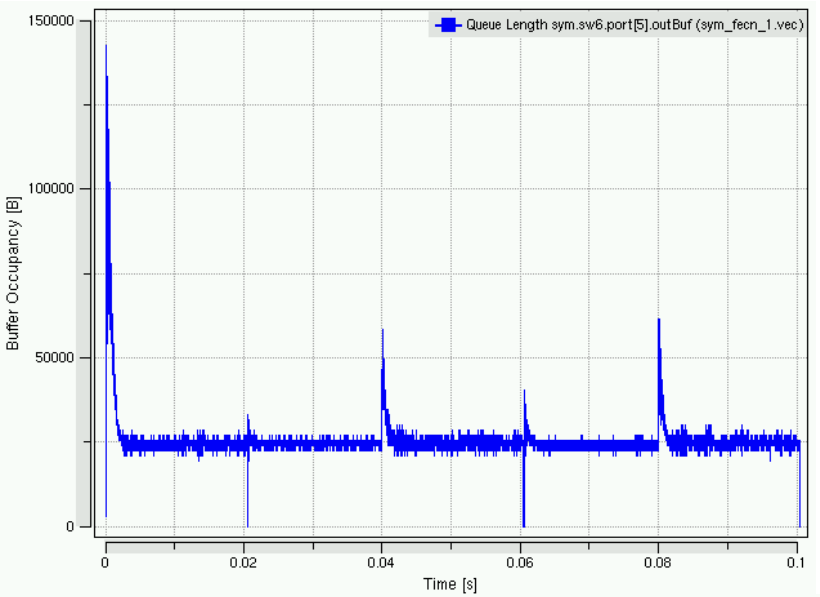

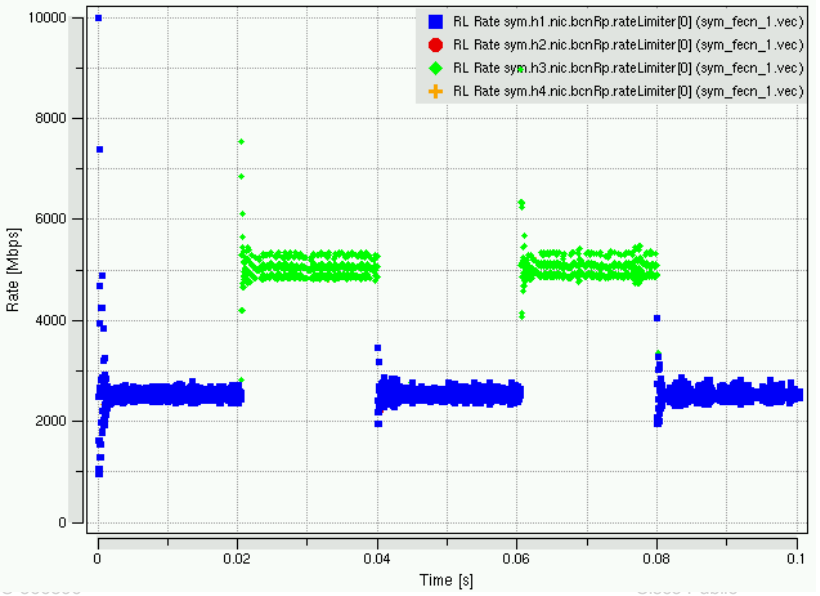

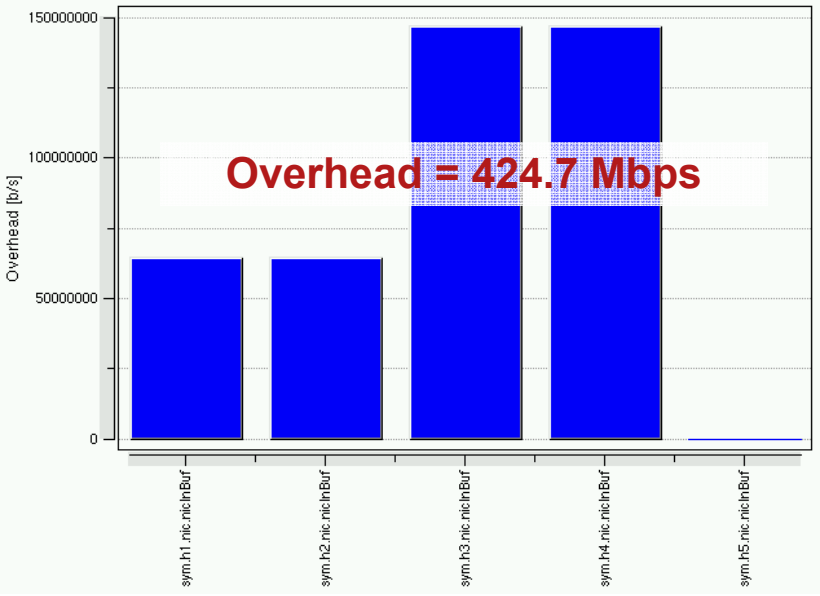

Symmetric Topology: FECN Tagging Rate = 10%

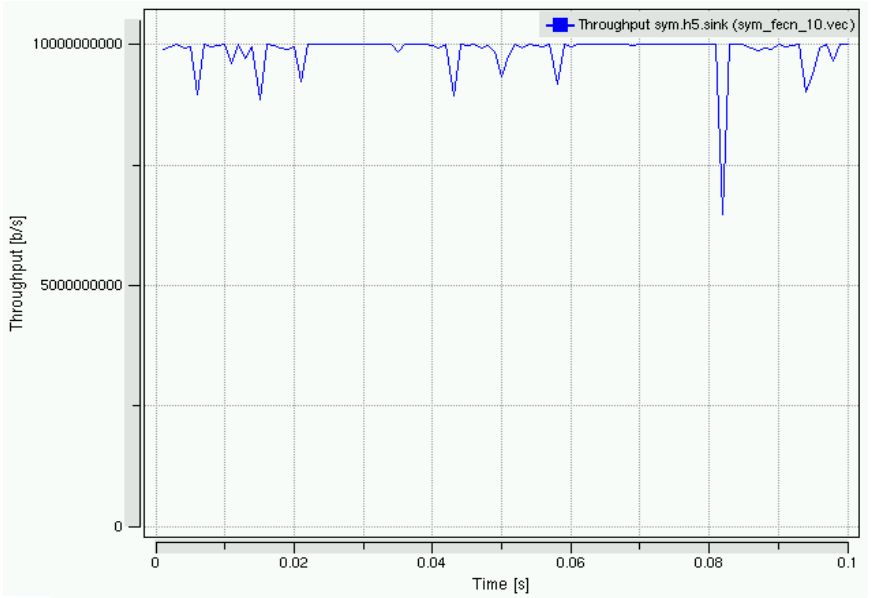

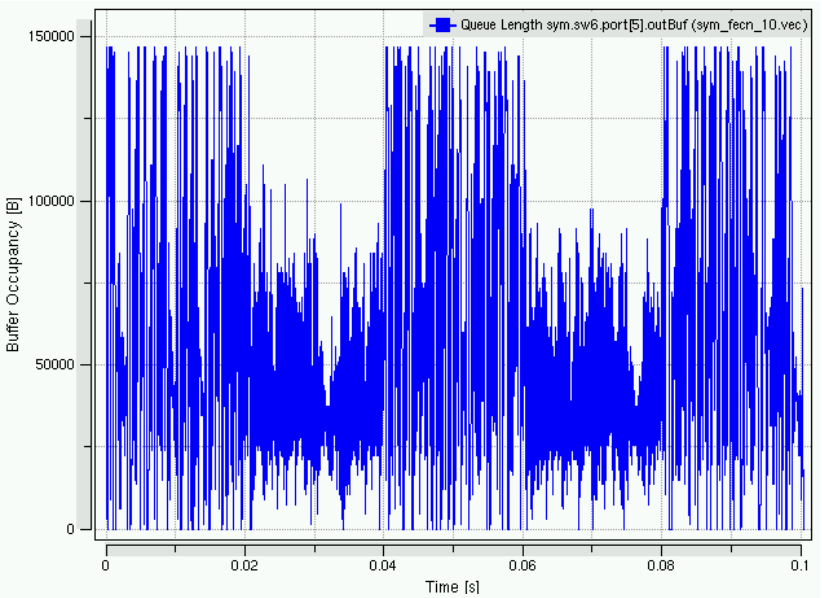

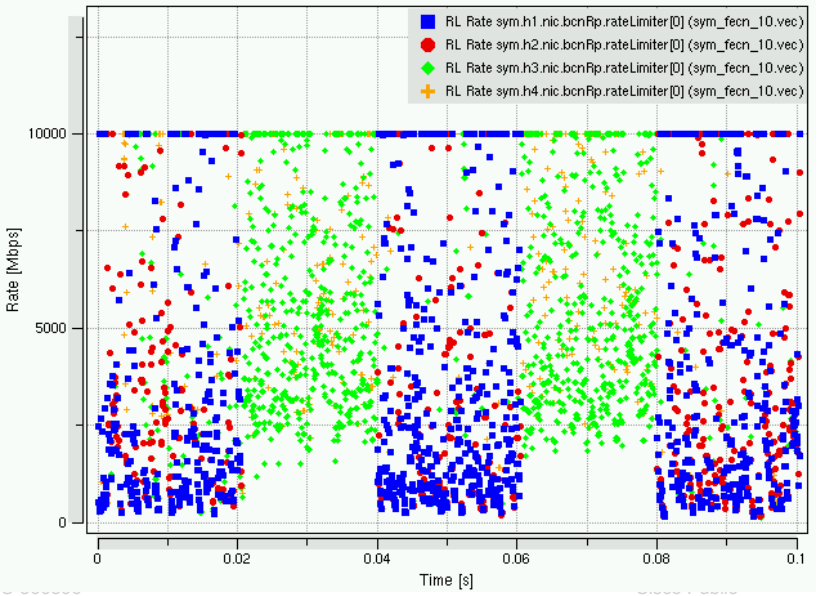

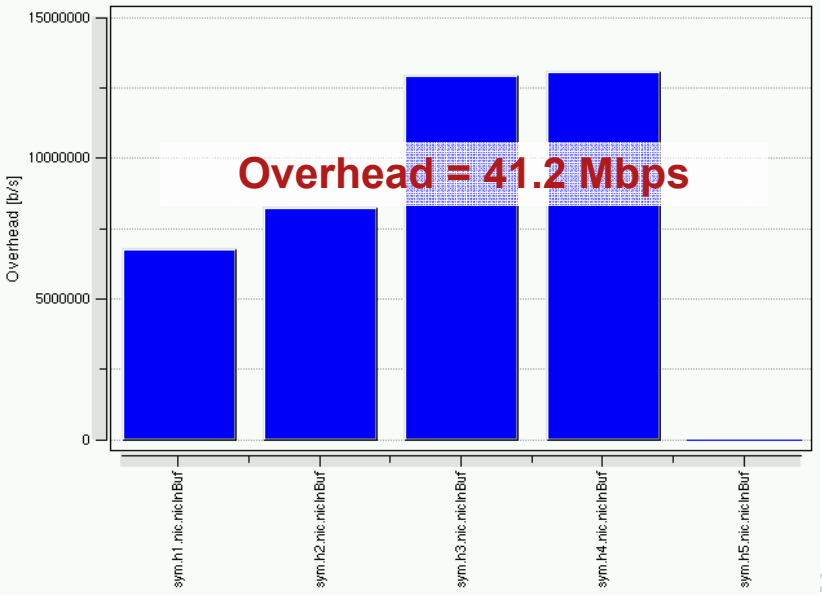

Symmetric Topology: FECN Tagging Rate = 5%

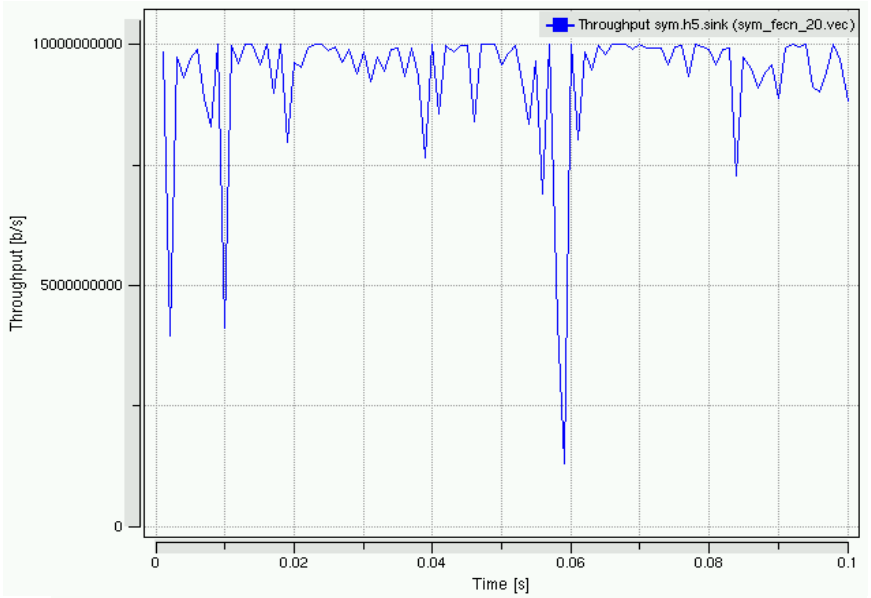

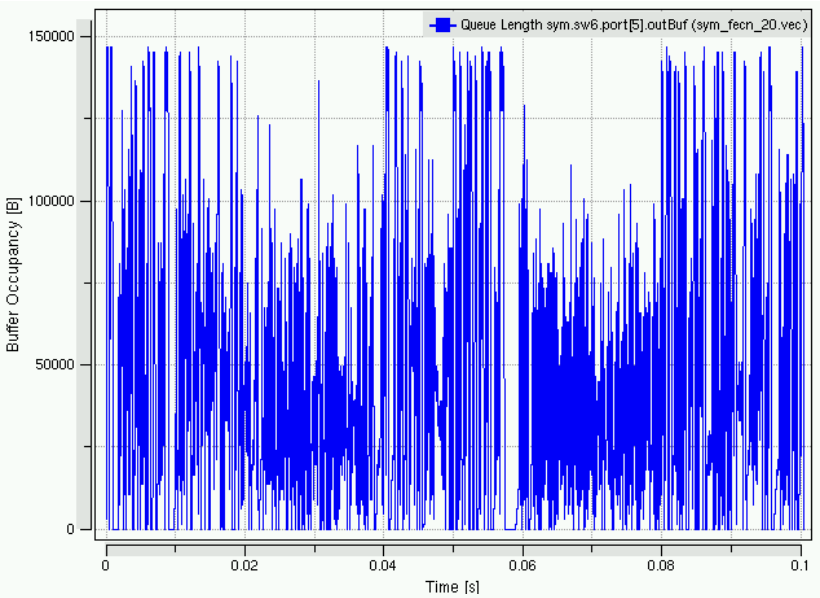

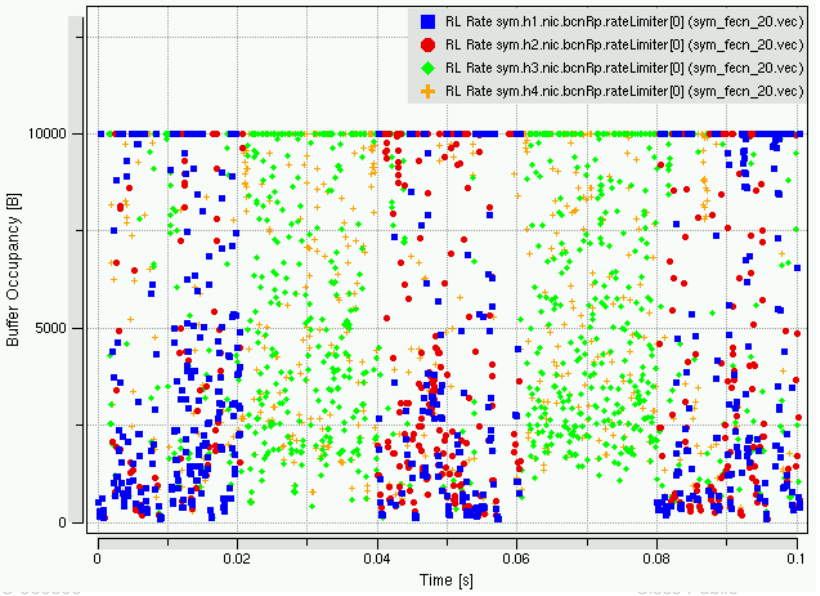

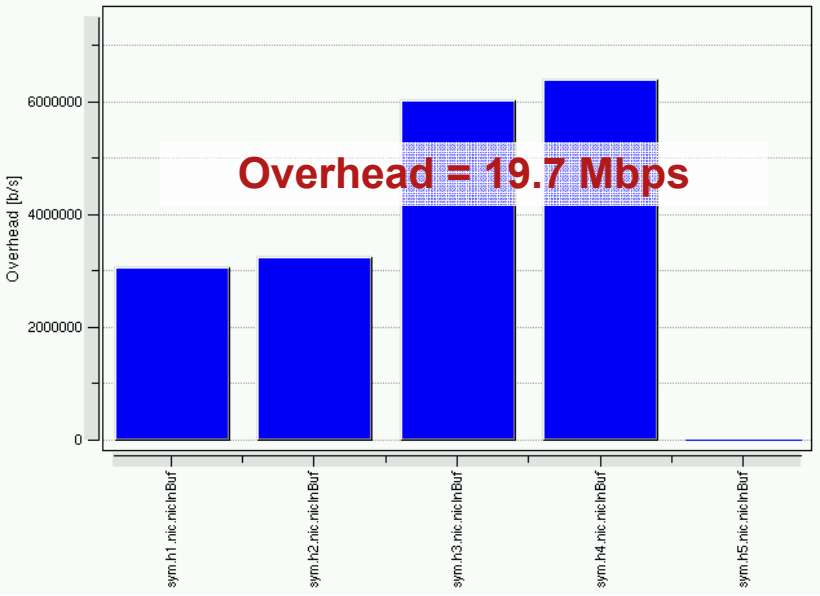

Symmetric Topology: FECN Tagging Rate = 1%

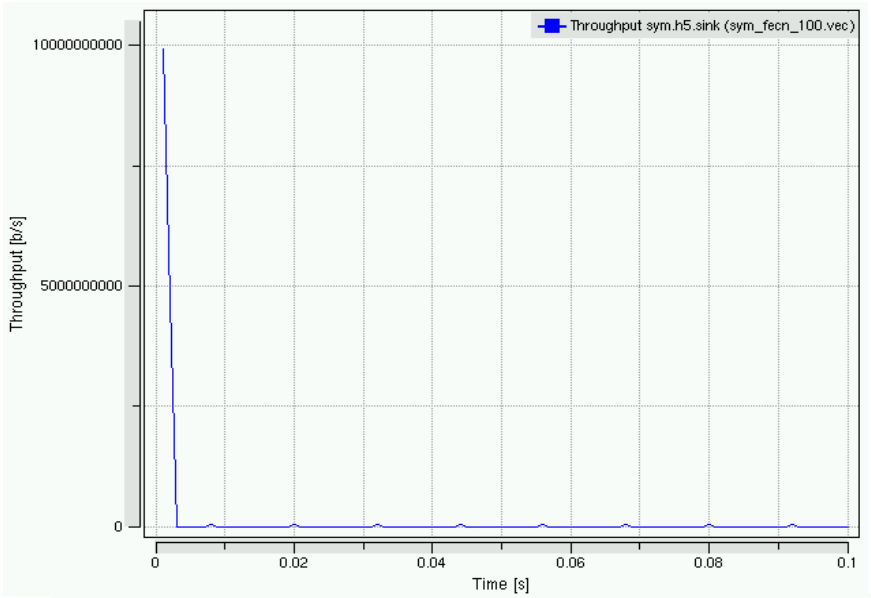

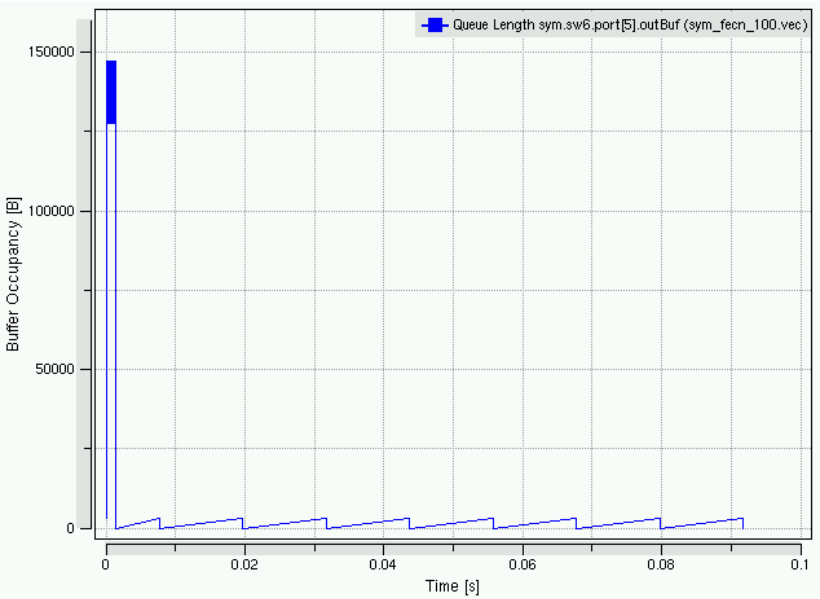

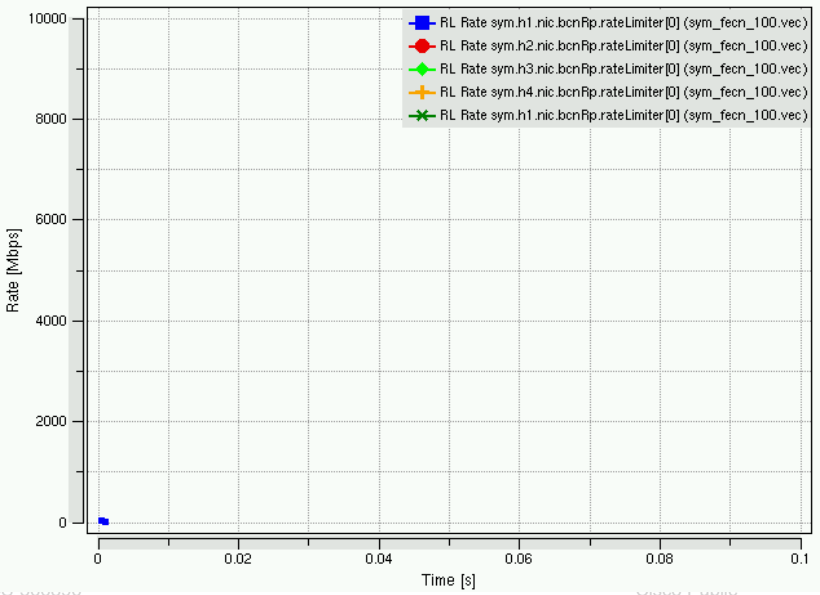

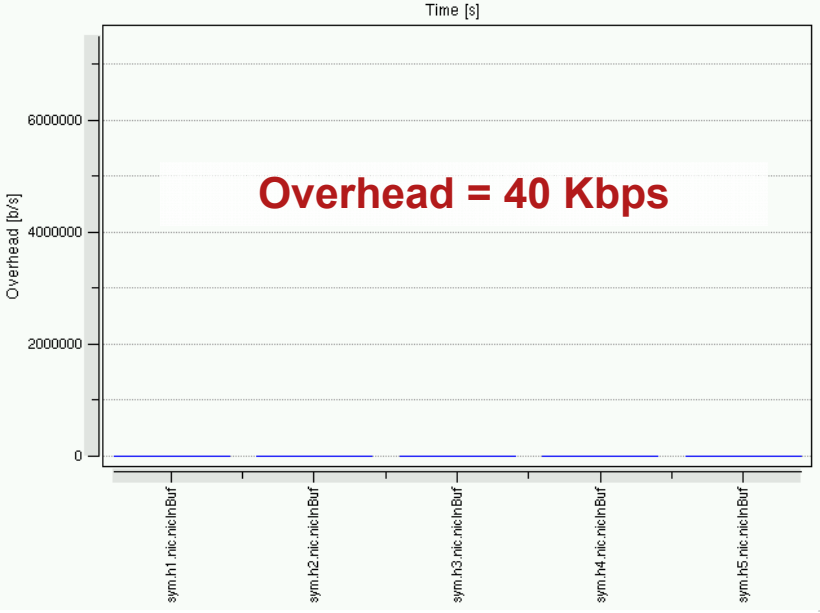

Symmetric Topology: BCN Sampling Rate = 1%

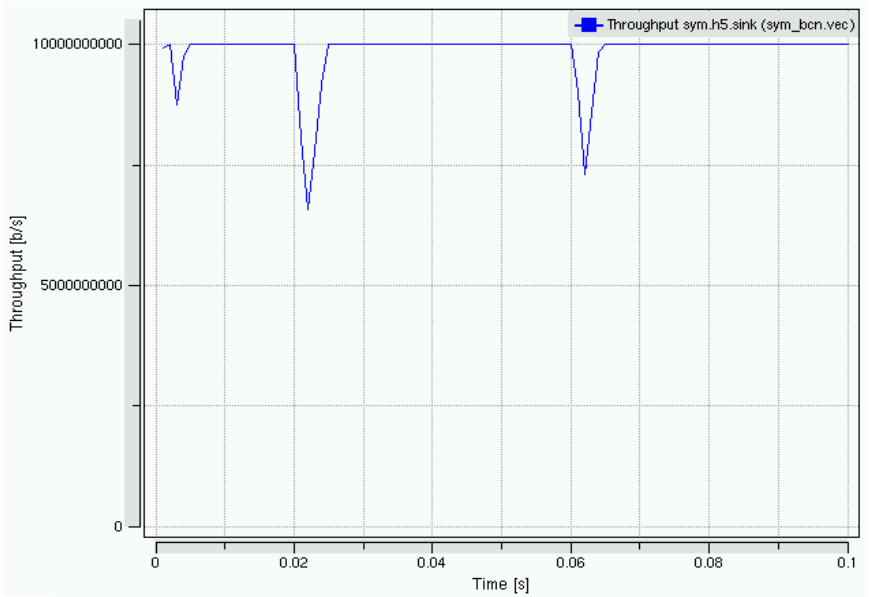

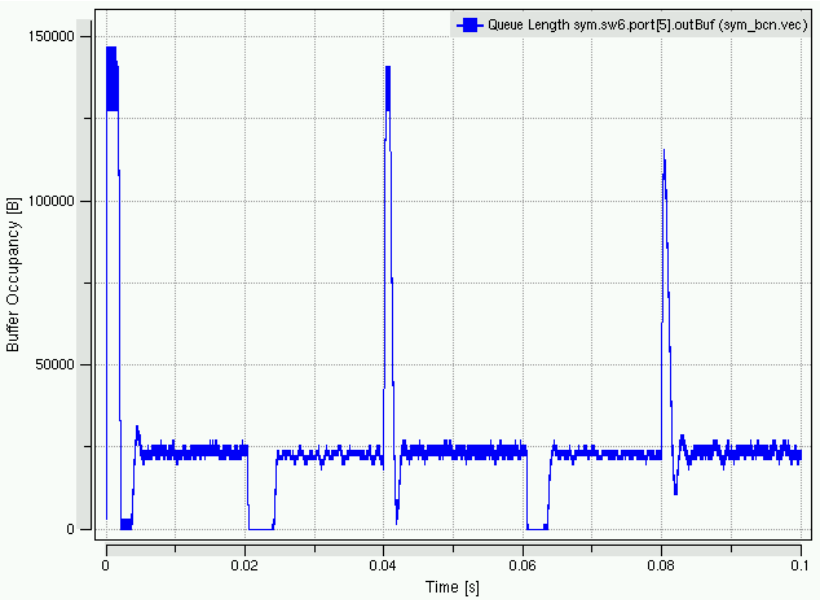

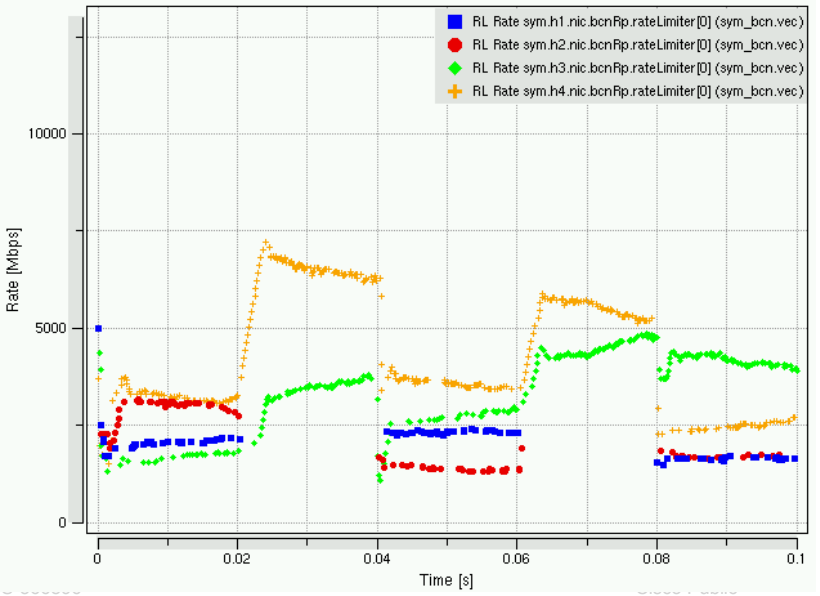

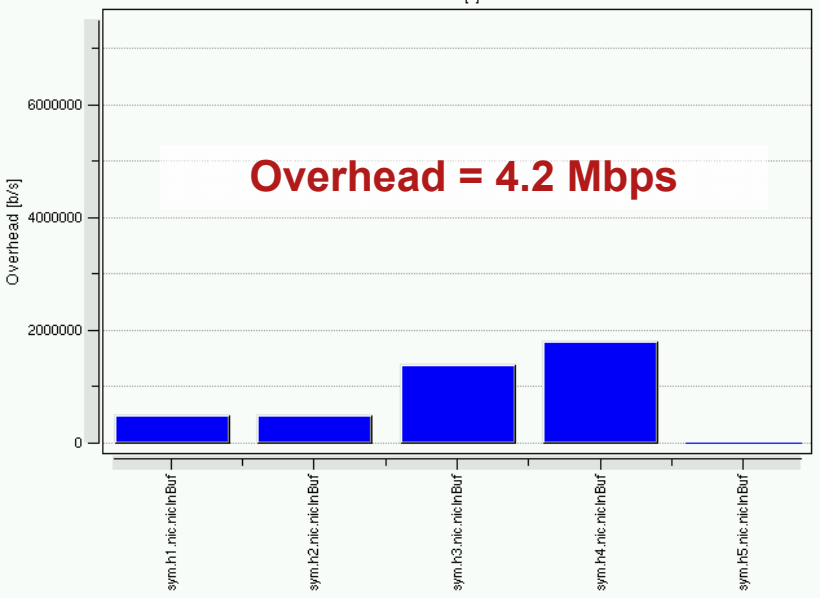

Backup I: BCN vs BECN

Symmetric Topology

- Topology & Workload as per Prof. Jain's presentation

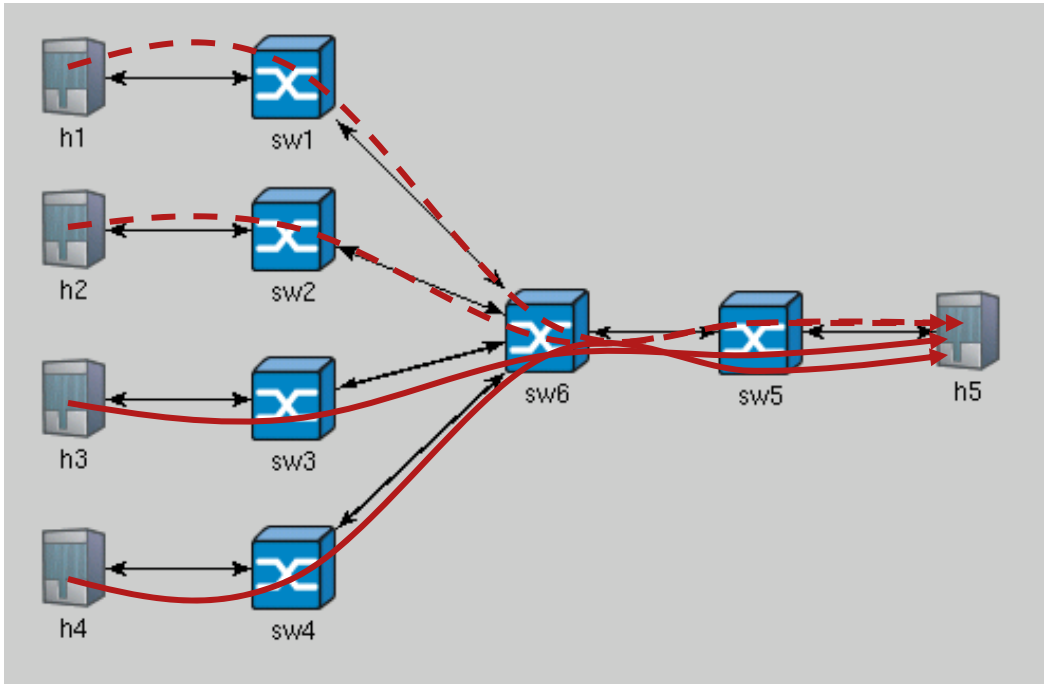

- Traffic pattern
  - Point-to-point from h1-4 to h5
  - Load: 100%
  - H1 and H2 on-off sources (T_{on} = T_{off} = 20 ms)

Symmetric Topology: Throughput

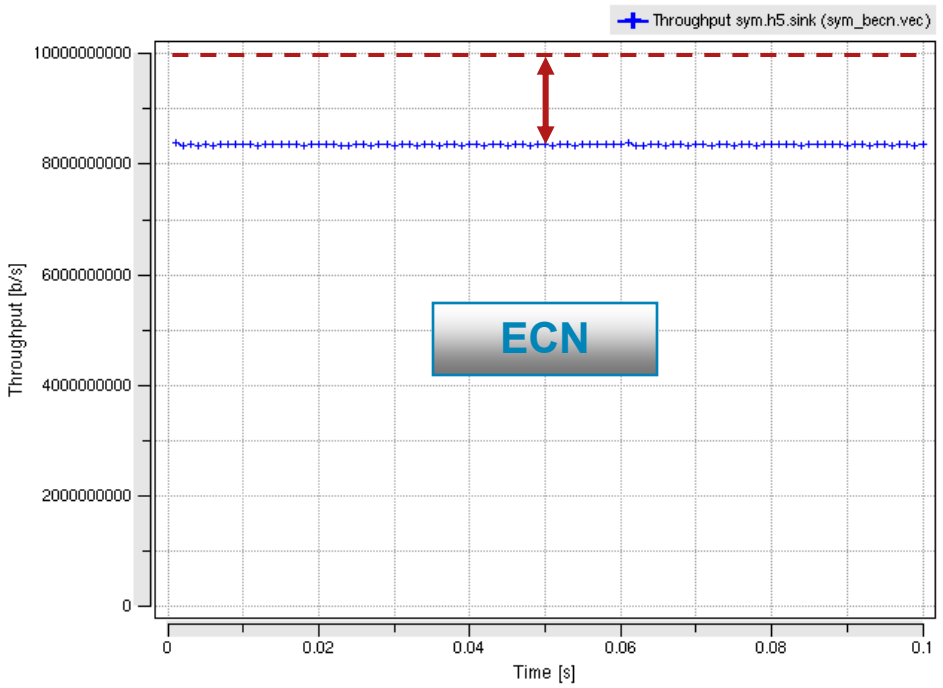

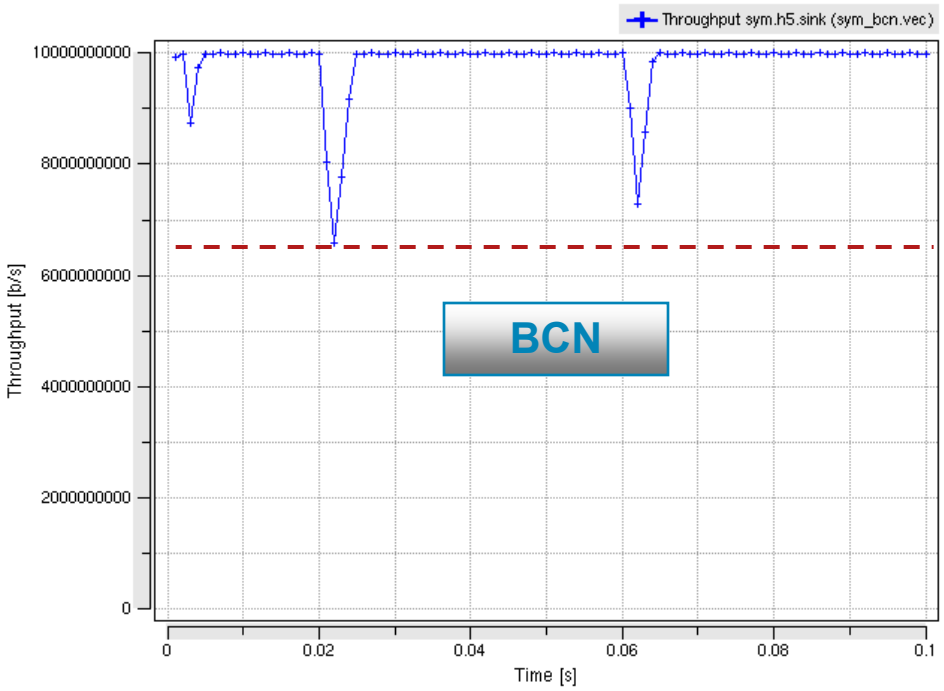

- ECN: Persistent throughput loss
- BCN: Some dips associated with H1/2 switching off

Symmetric Topology: Queue Length

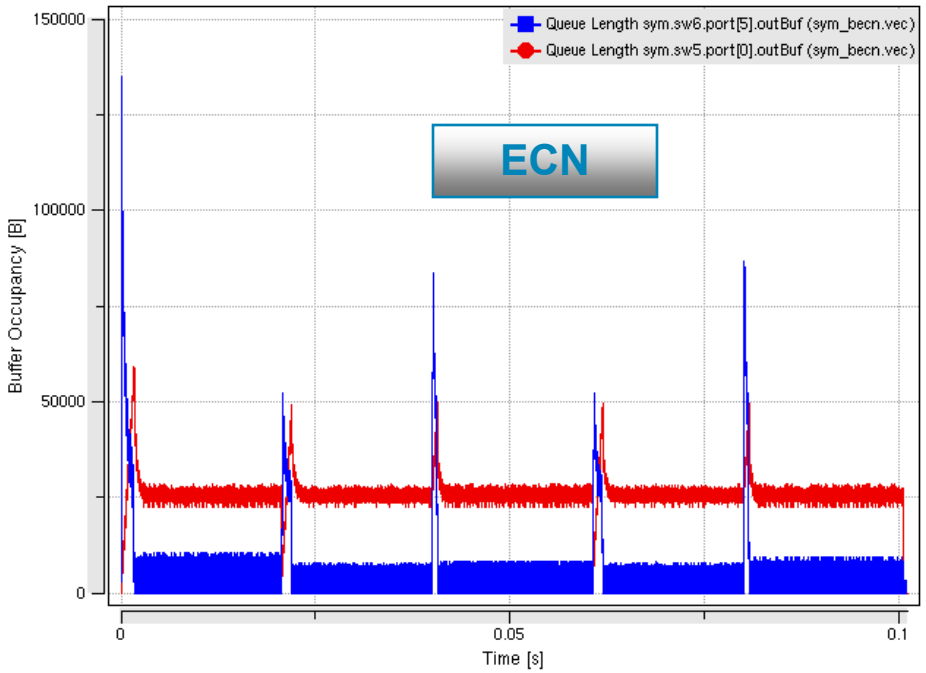

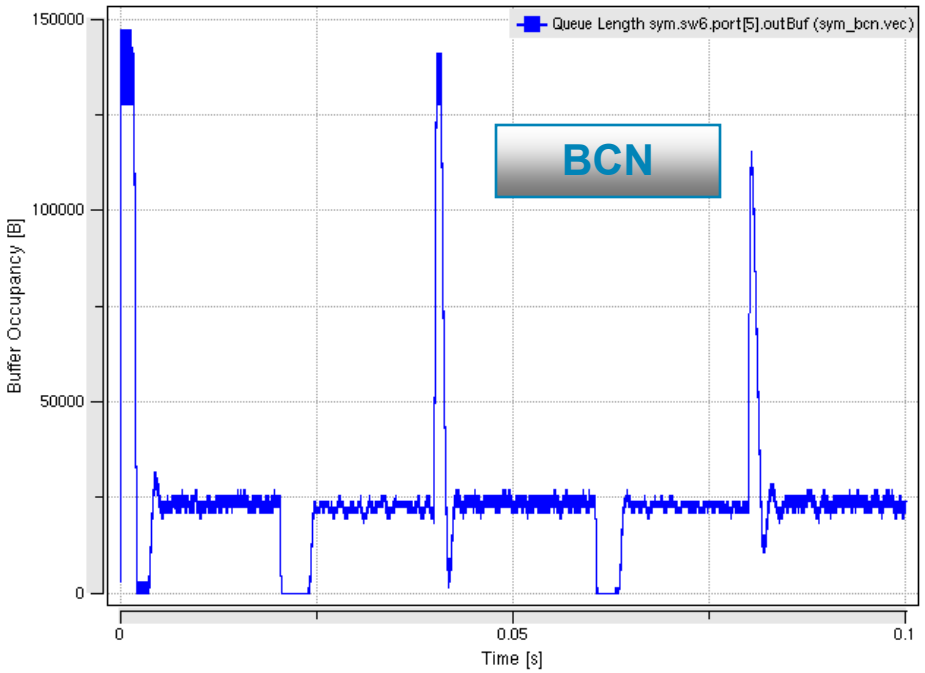

- ECN: CP has moved from core Sw to edge Sw because of signaling traffic
- BCN: peaks and valleys when H1/2 switch Off→On and On→Off

Symmetric Topology: Fairness

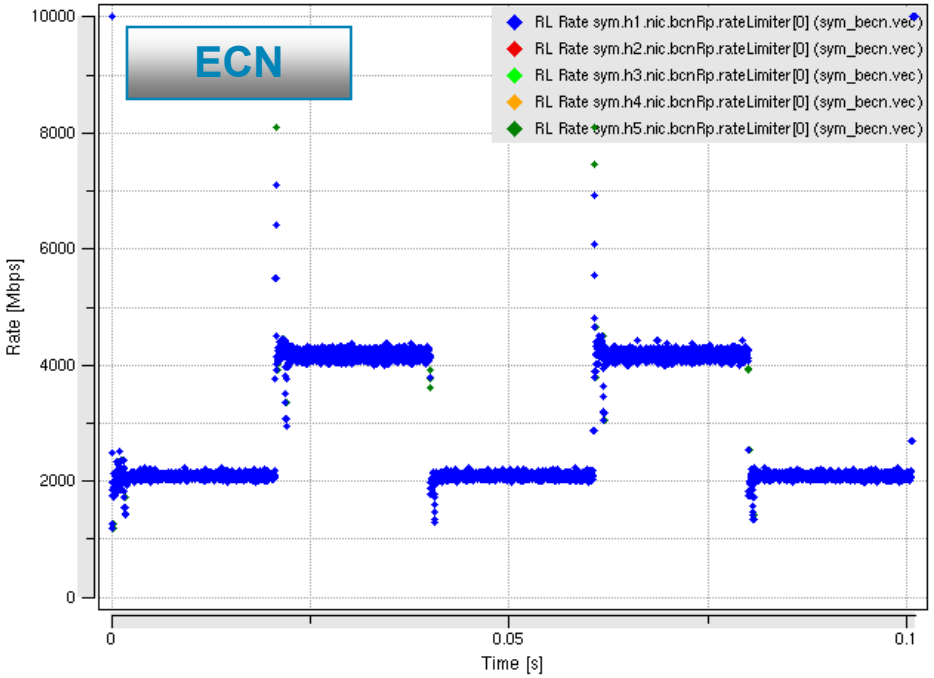

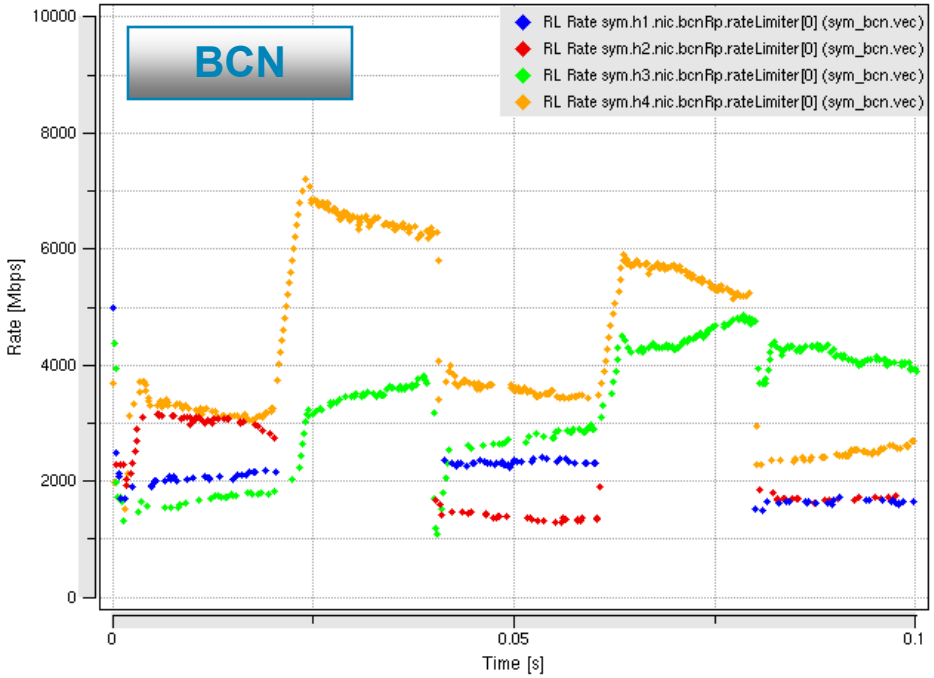

- ECN: Nearly perfect fairness and fast convergence
- BCN: Slower convergence to fairness

Symmetric Topology: Reducing Overhead 10x

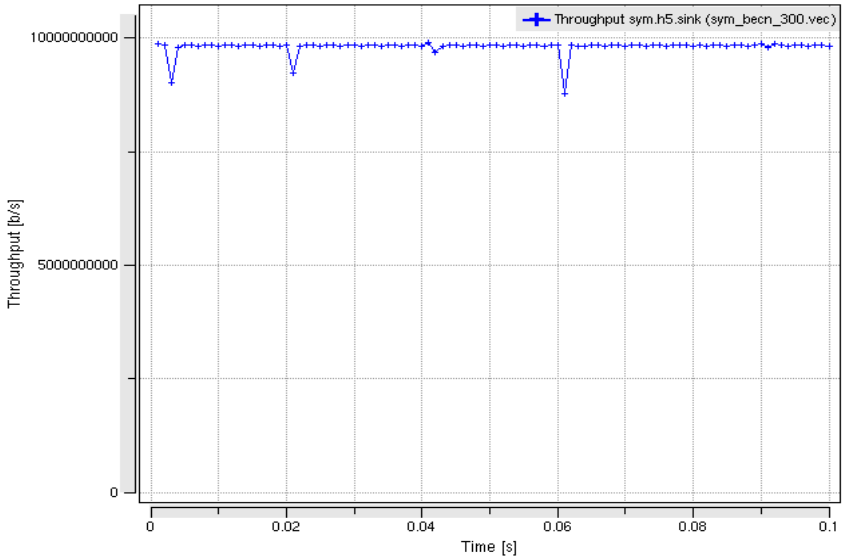

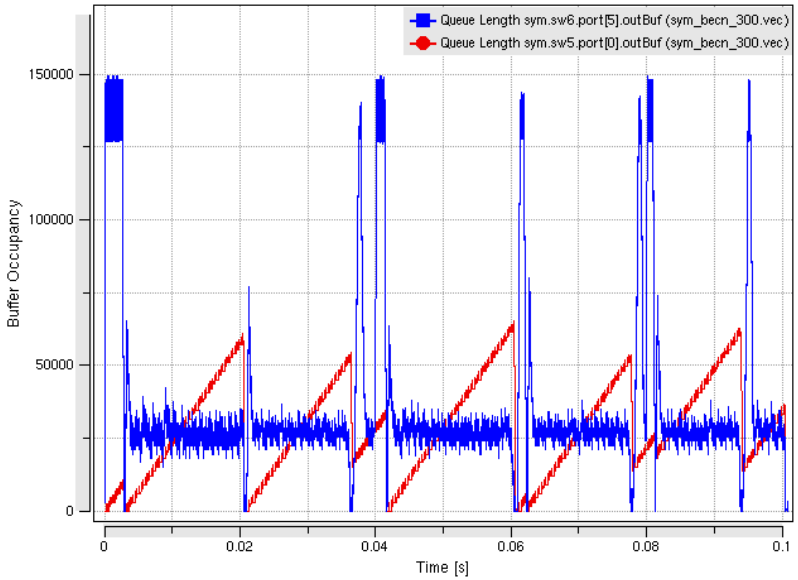

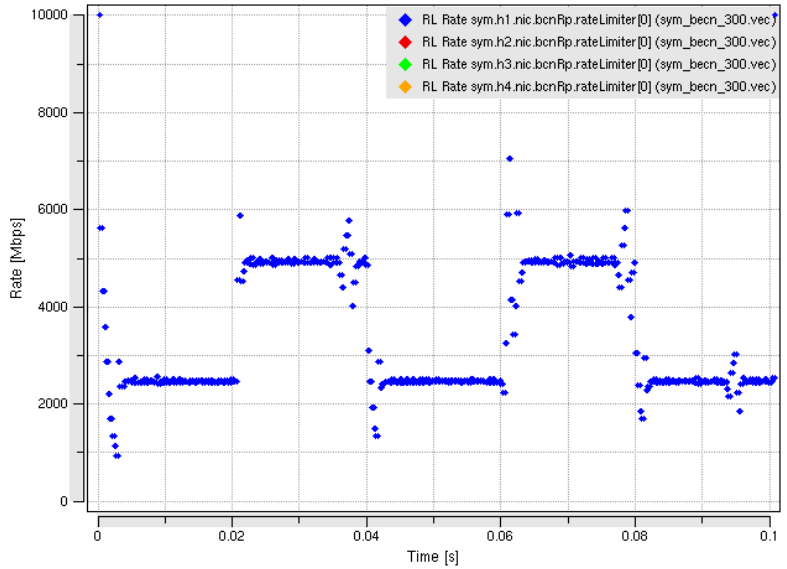

Measurement Interval = 300 µs

Aggr. Overhead = 163 Mbps

Symmetric Topology: Reducing Overhead 10x

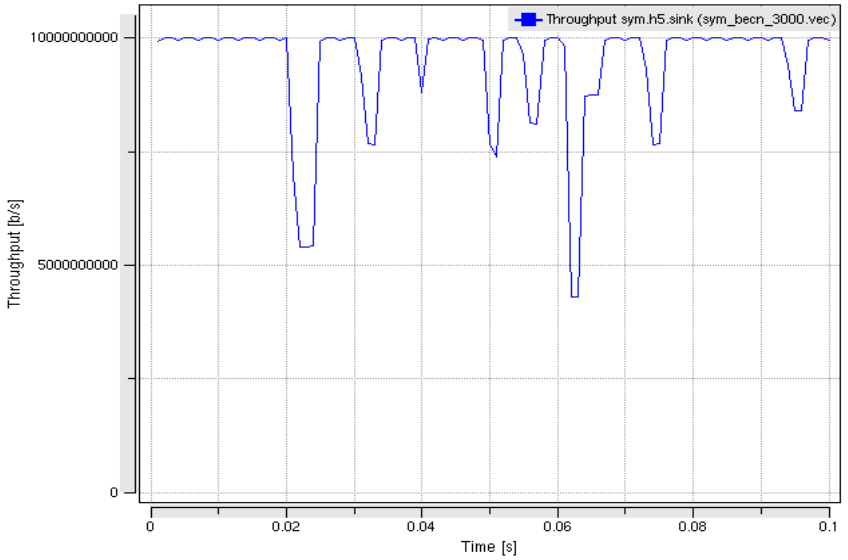

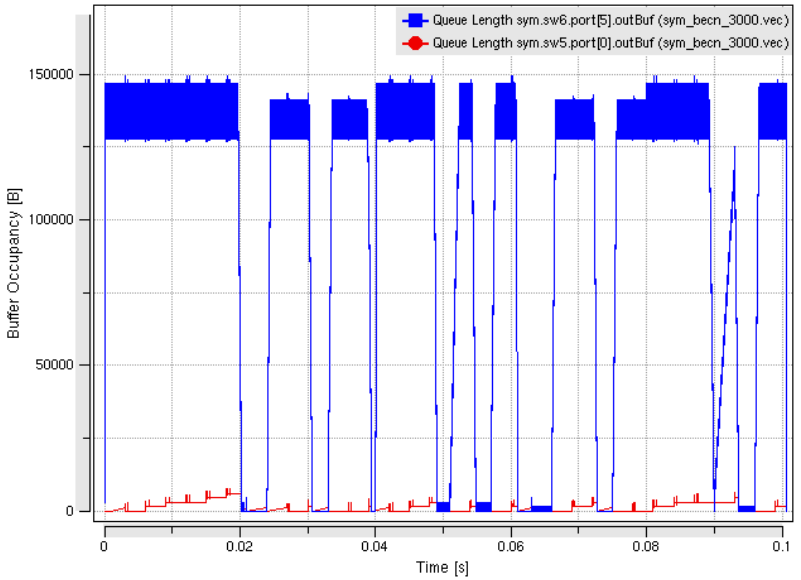

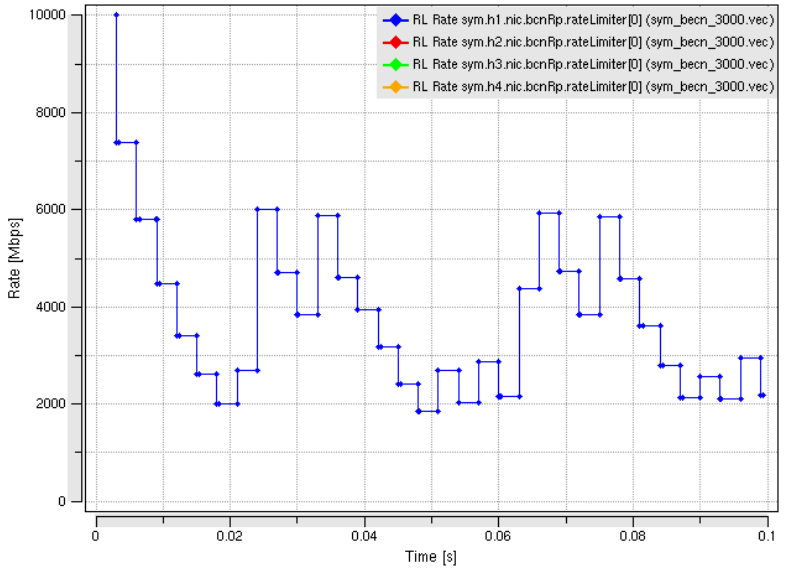

Measurement Interval = 3 ms

Aggr. Overhead = 16.3 Mbps

Parking Lot Topology

- Topology & Workload as per Prof. Jain's presentation

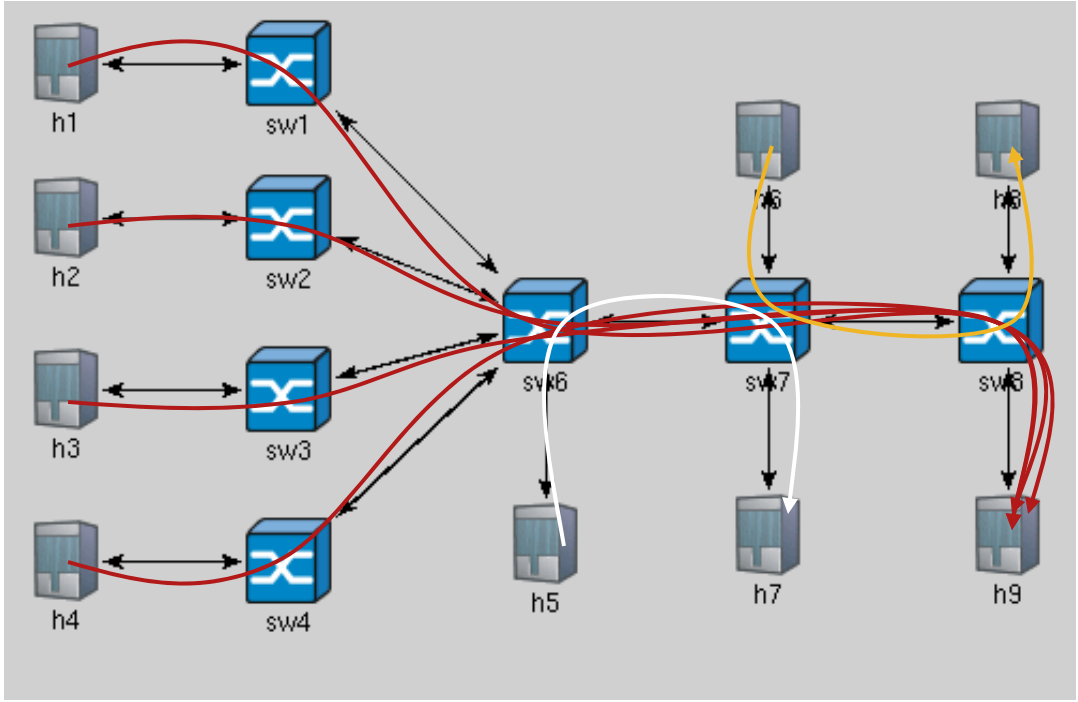

- Traffic pattern
  - Point-to-point from h1-4 to h9
  - Point-to-point from h5 to h7
  - Point-to-point from h6 to h8
  - Load: 100%

Parking Lot Topology

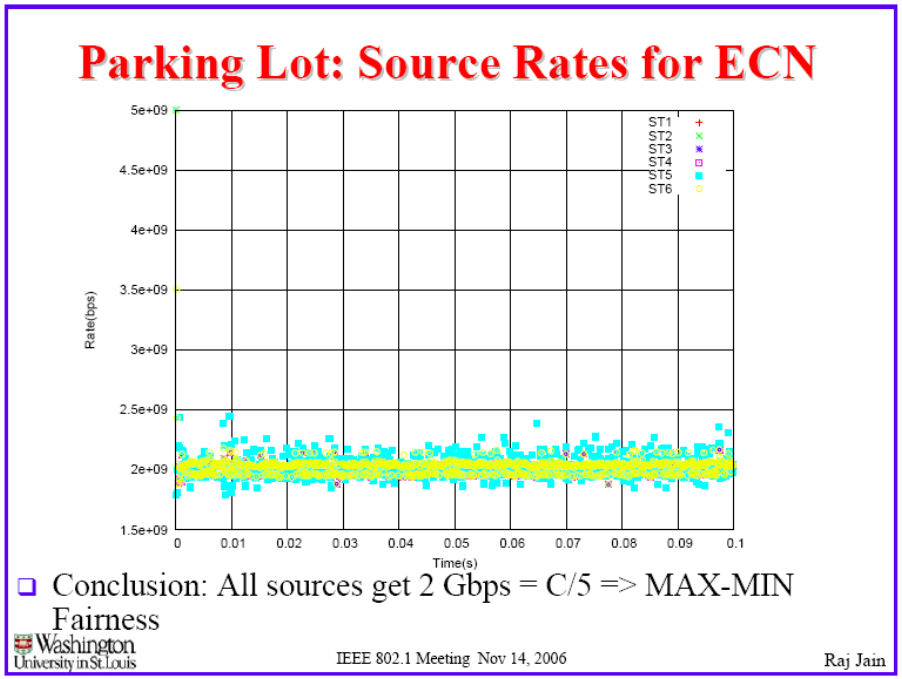

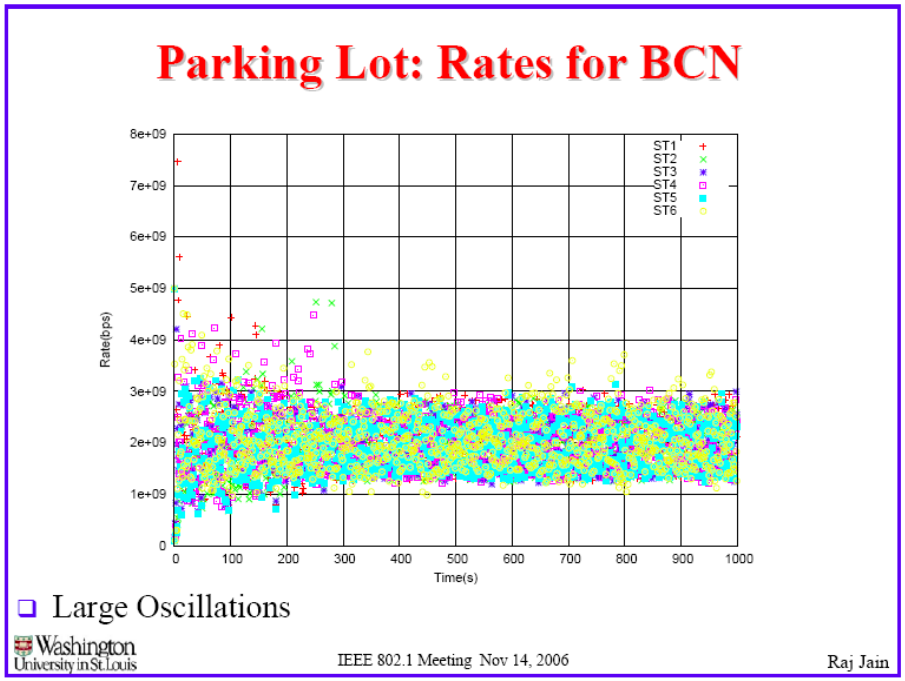

- Slide 34 seems to imply that BCN achieves proportional fairness while BECN achieves Max-min fairness
- However slides 35 and 36 show that both mechanisms achieve Max-min fairness

Parking Lot Topology

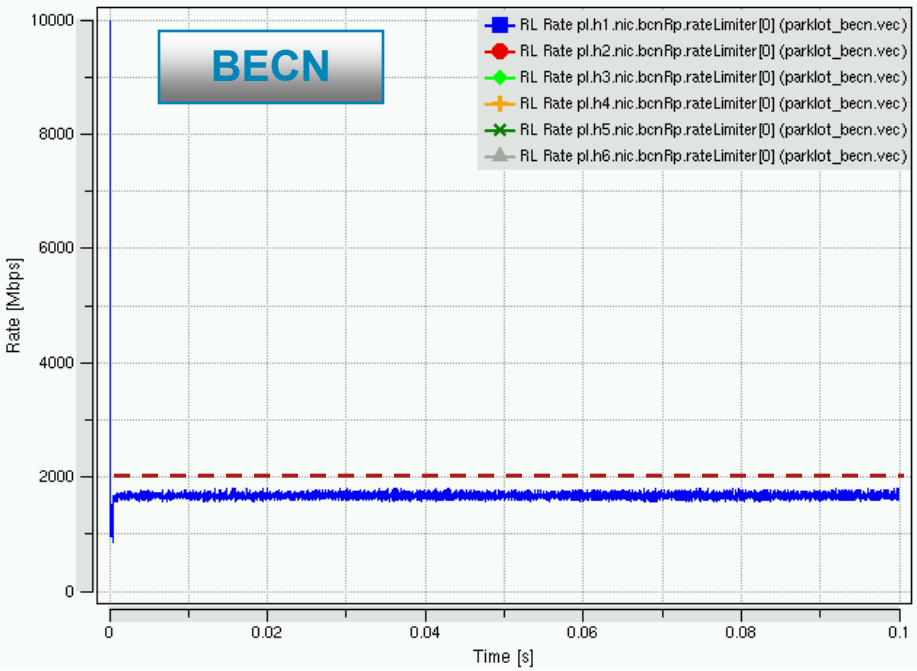

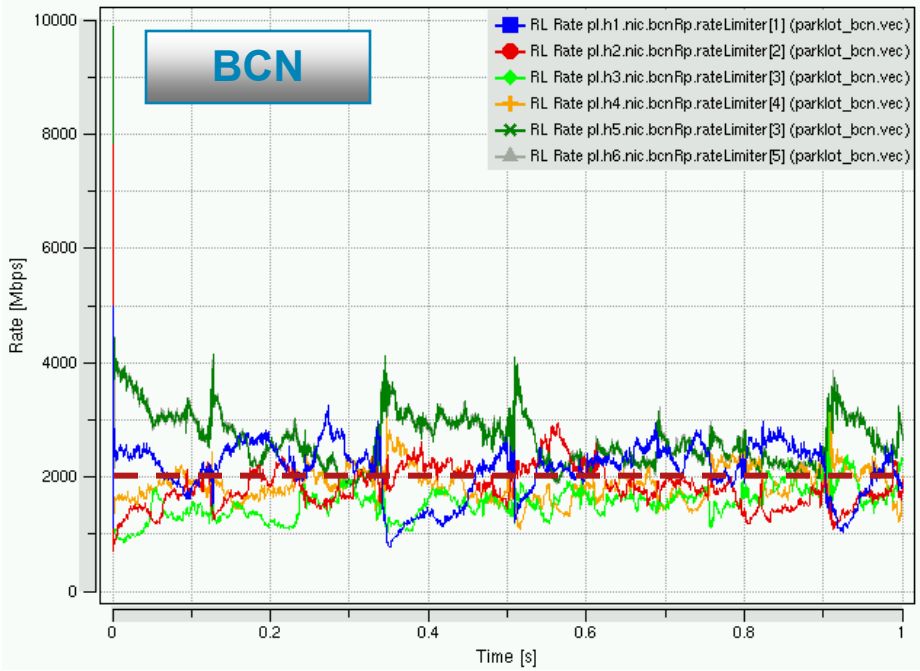

- Tradeoff between short term fairness and throughput
  - BECN: Perfect fairness, 1.63 Gbps throughput loss
  - BCN: Slower convergence to fairness, virtually no throughput loss