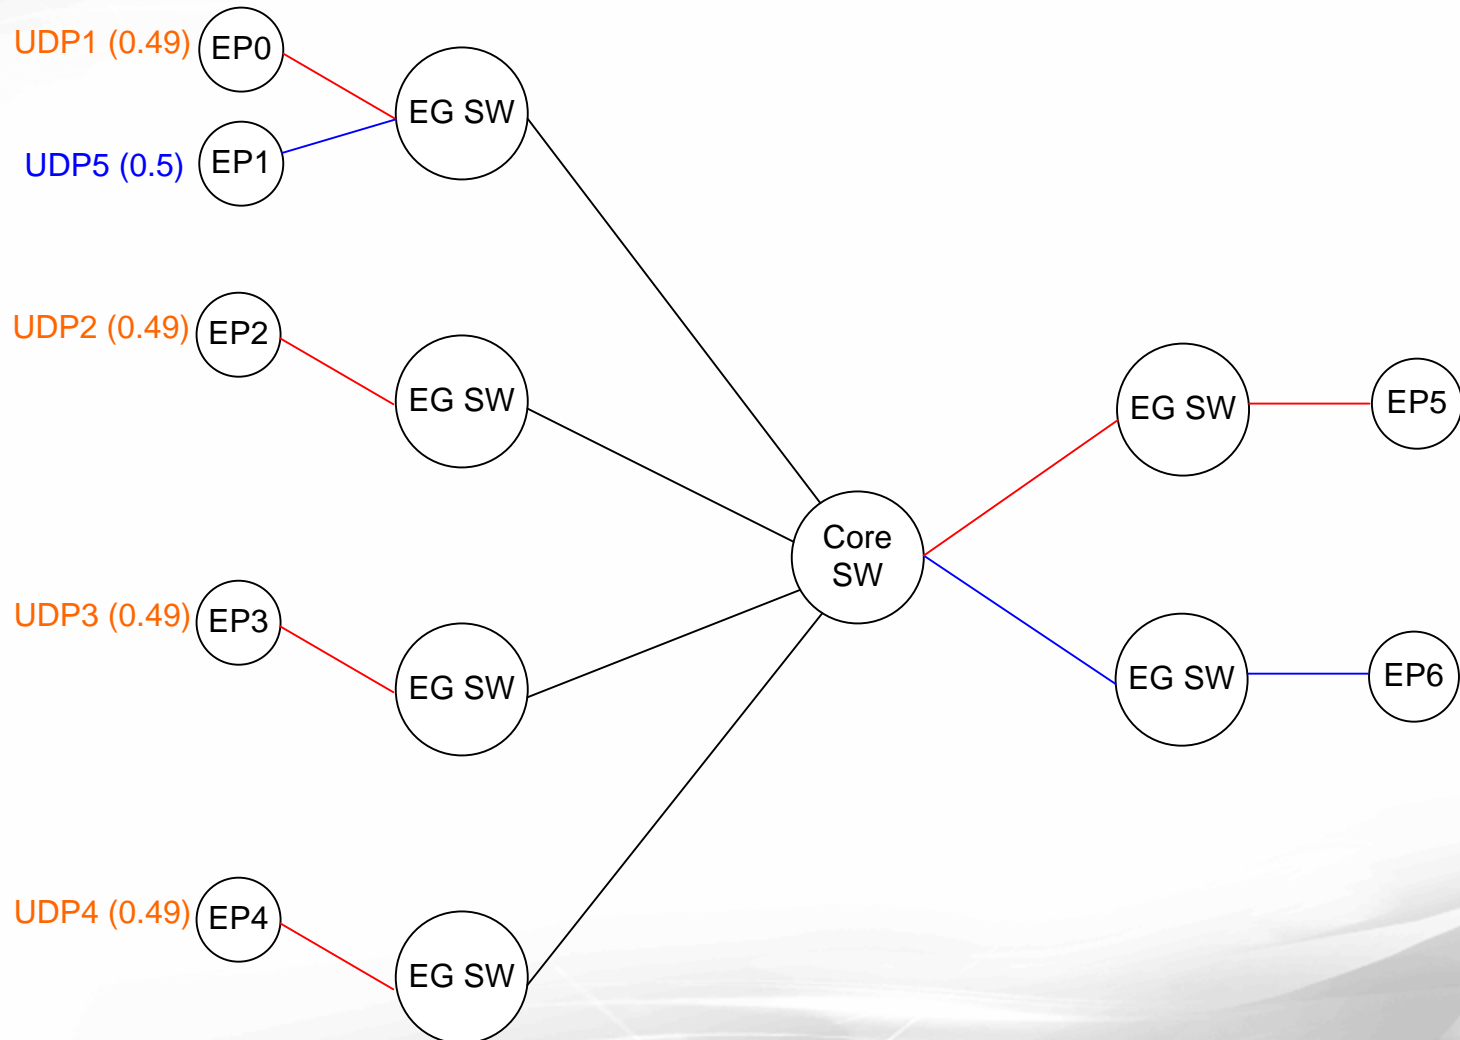




BCN + Pause Simulation Results

October 17, 2006

Topology



Workload

- Short Range, High-Speed Datacenter-like Network
 - Link Capacity = 10 Gbps
 - Egress Port Buffer Size = 150KB
 - Switch Latency = 1 us
 - Link Delay = 0.5 us
 - Endpoint response time = 1 us
- Traffic Type: 100% UDP (or Raw Ethernet) Traffic
- Destination Distribution:
 - EP0, EP2, EP3, EP4 send to EP5
 - EP1 sends to EP6
- Frame Size Distribution: Fixed length (1500 bytes) frames
- Arrival Distribution: Bernoulli temporal distribution
- Offered Load/Endpoint
 - 49% for EP0, EP2, EP3, EP4
 - 50% for EP1

BCN Parameters

- Qeq
 - 16 (1500-byte frames)
 - 375 * 64 byte pages
- Frame Sampling
 - Frames are sampled on average 150 KB received to the egress queue
 - Sampling Jitter for BCN: 20KB
- $W = 2$
- $G_i = 12.42$
 - Computed as $(\text{Linerate}/10) * [1/((1+2*W)*Q_eq)]$
 - $G_i = 5.3 \times 10^{-1} * (1500/64) = 12.42$
- $G_d = 6.09 \times 10^{-3}$
 - Computed as $1/2 * [1/((1+2*W)*Q_eq)]$
 - $G_d = 2.6 \times 10^{-4} * (1500/64) = 6.09 \times 10^{-3}$
- $R_u = 1 \text{ Mbps}$

BCN Detection & Global Pause

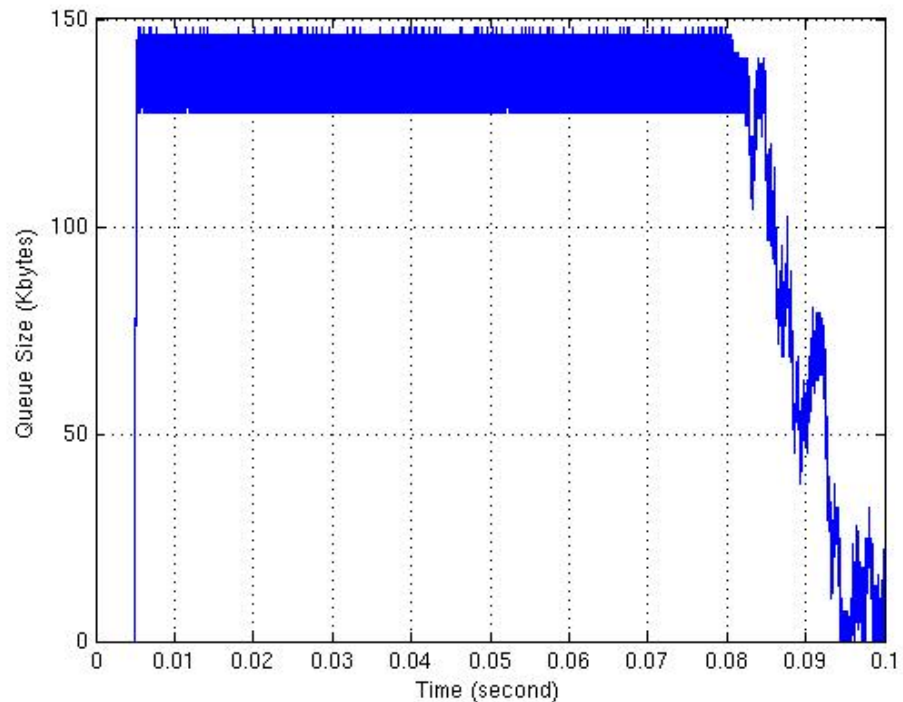
- BCN detection is enabled at CS
 - No BCN(0,0)
- Global Pause: send pause msg to each input port based on the output queue
 - Pause detection is enabled at CS and ES
 - CS Threshold
 - Xoff thresh = 140 KB
 - Xon thresh = 130 KB
 - ES Threshold
 - Xoff thresh = 130 KB
 - Xon thresh = 110 KB

Simulation

- Simulation
 - Duration: 100ms
 - $t = 5\text{ms}$ (source start)
 - $t = 80\text{ms}$ (EP3 and Ep4 stop traffic)
- Throughput collected for multiple UDP flows at 80ms.
- The CS queue size is collected.
- Fairness is also collected for multiple runs.

Only Global Pause Throughput and CS Queue


Only Pause	
UDP flow	Throughput at 80ms (Mbps)
EP0	1408.00
EP2	2857.92
EP3	2857.60
EP4	2861.76
EP1	1404.96



- No packet loss at CS and ES.
- Pause can affect the innocent flow at EP1.

Global Pause and BCN Throughput and CS Queue

Pause + BCN	
UDP flow	Throughput at 80ms (Mbps)
EP0	3511.32
EP2	1886.10
EP3	1951.23
EP4	2646.63
EP1	4824.16



- No packet loss at CS and ES
- BCN can alleviate the pause problem for the innocent flow at EP1.

Simulation Statistics

- Fairness Statistics of 4 UDP flows (EP0, EP2-4)
 - Error: % difference from target rate for each flow = $|(R_i - T)/T|$
 - R_i : rate of individual flows, T = target rate (2.5 Gbps), $N = 4$ (number of flows)
 - Root Mean Square Fairness:
$$\sqrt{\frac{\sum (\frac{R_i - T}{T})^2}{N}}$$
- Min, Mean, Max, and Standard Deviation of Fairness Index across different runs

Fairness Result: 20ms - 80ms

Pause	
# of Runs	RMS Fairness Index (Min, Mean, Max, Std)
25	(0.25, 0.25, 0.26, 0.005)
100	(0.25, 0.25, 0.26, 0.005)
200	(0.25, 0.25, 0.26, 0.005)
300	(0.25, 0.25, 0.26, 0.005)

Pause + BCN	
# of Runs	RMS Fairness Index (Min, Mean, Max, Std)
25	(0.09, 0.20, 0.32, 0.060)
100	(0.05, 0.18, 0.32, 0.062)
200	(0.02, 0.18, 0.36, 0.067)
300	(0.02, 0.18, 0.40, 0.066)

BCN improves fairness