Update on Effects of Delay on BCN for Symmetric Topology w/ Single Hot Spot Scenario

Bruce Kwan & Jin Ding

March 13-14, 2007

IEEE 802.1Qau Plenary Meeting in Orlando, FL

Overview

- Goals & System Parameters
- Key Observations
- Effects of Large Control Loop Delay
- Effects of Reducing Qeq
- Conclusions

BROADCOM
Connecting everything®

Goals

- Characterize performance across the target range of control loop delay as specified in the IEEE 802.1Qau PAR

- March 14, 2006 PAR
  - "The bandwidth-delay product limit is expected to be in the region of 1-5Mbits (*100-500us* control loop delay for 10Gbps network) and simulation and analysis will verify performance characteristics up to the advertised bandwidth-delay product."

- Additions over Previous Presentation*
  - Added control loop delay of 500us case
  - Adjustments to derivative weight (w) to stabilize queue behavior
  - Effects of reducing Qeq from 150KBytes down to 24KBytes

***http://www.ieee802.org/1/files/public/docs2007/au-kwan-ding-bcn-effects-of-delay-02152007.pdf**

BROADCOM
Connecting everything

Base System Parameters

- No PAUSE
  - Goal is to characterize BCN behavior without PAUSE
- Switch Parameters
  - Buffer Size (B)
    - 600Kbytes/Port.
  - Egress Port Discard Threshold
    - 600kbytes
- BCN Parameters
  - Frame Sampling
    - Frames are periodically sampled (on avg) every 75KB (2%)
  - W = 2
  - Qeq = B/4 or 24kbytes
  - Ru = 1Mbps
  - Gi (Initial)
    - Computed as (Linerate/10) * [1/((1+2*W)*Q_eq)]
    - Same as in baseline
  - Gd (Initial)
    - Computed as 0.5*1/((1+2*W)*Q_eq)
    - Same as in baseline
  - BCN(Max) Enabled
  - Other BCN Enhancements
    - No Oversampling
    - No BCN(0,0)
    - No Self Increase

BROADCOM
Connecting everything

Symmetric Topology Single HS – Non Bursty (Similar to Required Scenario #5)

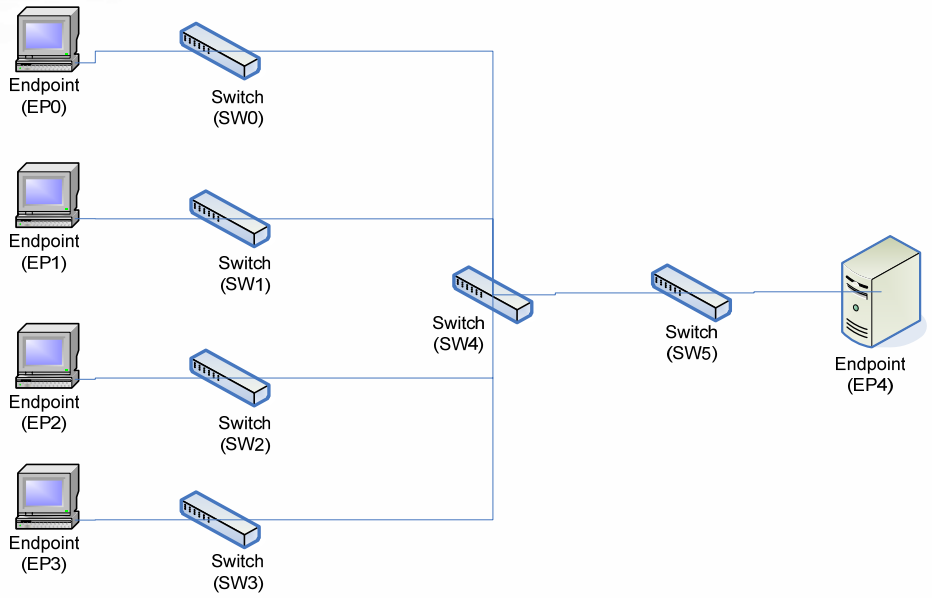

- Symmetric Topology Single HS
  - Link speed : 10Gbps for all links
- Traffic Pattern
  - Traffic Type: 100% UDP (or Raw Ethernet) Traffic
  - Destination Distribution: EP0-EP3 send to EP4
  - Frame Size Distribution: Fixed length (1500 bytes) frames
  - Arrival Distribution: Bernoulli temporal distribution
  - Offered Load/Endpoint = 50%
- Control Loop Delay is between the source Endpoints and the Congestion Point

BROADCOM
Connecting everything

Overview

- Goals & System Parameters
- Key Observations
- Effects of Large Control Loop Delay
- Effects of Reducing Qeq
- Conclusions

BROADCOM
Connecting everything®

Key Observations

- Adjustments to derivative weight ‘w’ enhances performance under large control loop delay conditions.

- Choosing a value for ‘w’ across varying network control loop delays appears more challenging when Qeq is small.

BROADCOM
Connecting everything®

Overview

- Goals & System Parameters
- Key Observations
- Effects of Large Control Loop Delay
- Effects of Reducing Qeq
- Conclusions

BROADCOM
Connecting everything®

Symmetric Topology Single HS – Non Bursty
Effects of Control Loop Delay (Queue Size @ CP)

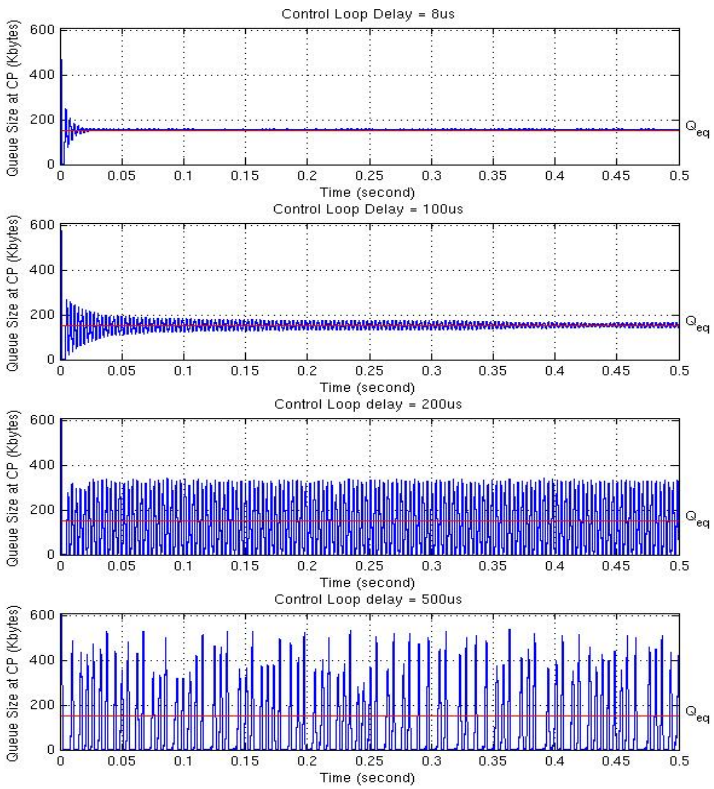

Ctrl Loop Delay	# of drops	Throughput
8us	6	9.992Gbps
100us	106	9.982Gbps
200us	185	9.775Gbps
500us	419	8.056Gbps

Mild throughput degradation and increase in frame drops as control loop delay increases. Derivative weight ‘w’ is kept at 2 in all of these experiments. At 200us and above, underutilization begins to occur.

BROADCOM
Connecting everything

Symmetric Topology Single HS – Non Bursty Effects of Control Loop Delay (Throughput)

Ctrl Loop Delay

8us

100us

200us

500us

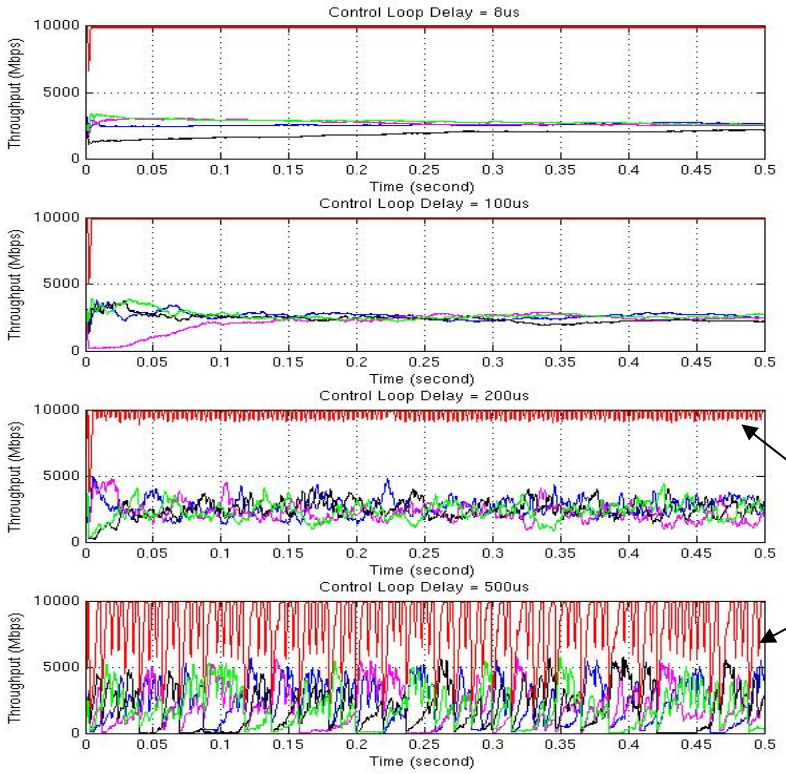

Underutilization occurs as control loop delay increases higher than 200us.

Adjustments to Derivative Weight
Effects of Control Loop Delay (Queue Size @ CP)

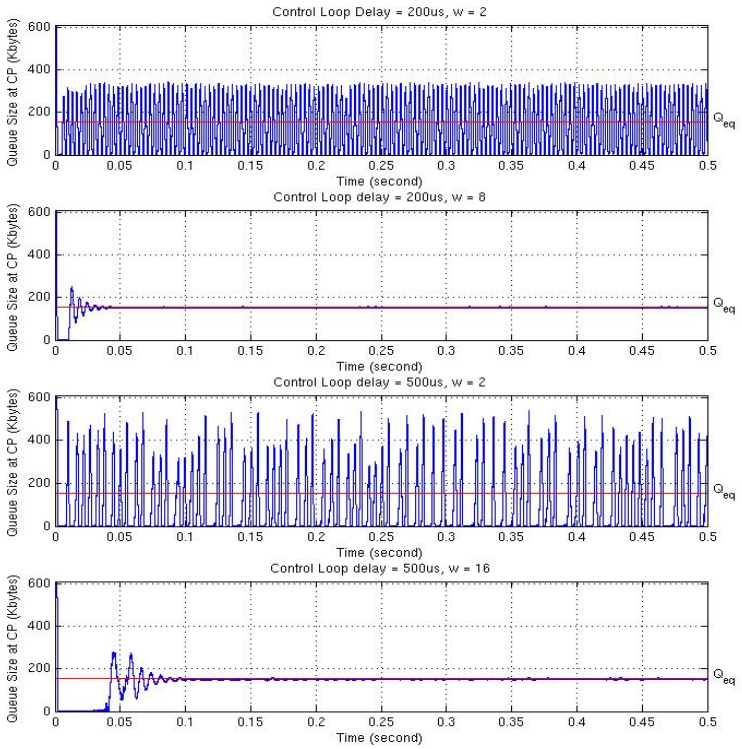

Ctrl Loop Delay	w	# of drop	Throughput
200us	2	185	9.775Gbps
200us	8	204	9.939Gbps
500us	2	419	8.056Gbps
500us	16	451	9.728Gbps

Adjustments to derivative weight (w) enable BCN to better control the queue and avoid underutilization.

BROADCOM
Connecting everything

Adjustments to Derivative Weight
Effects of Control Loop Delay (Throughput)

Ctrl Loop Delay	w
200us	2
200us	8
500us	2
500us	16

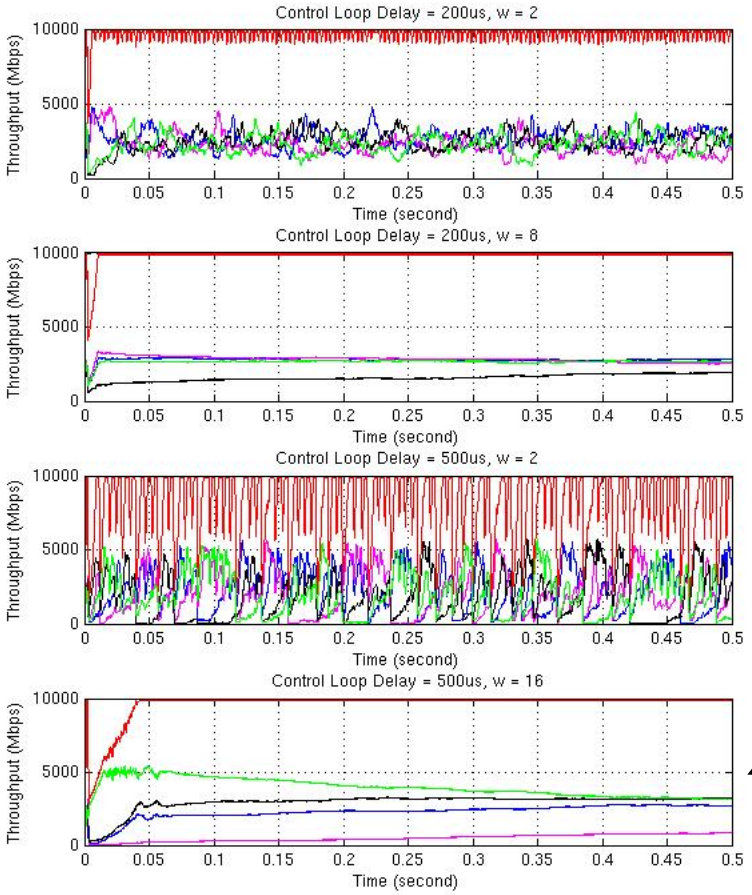

Increasing w helps improve the throughput during the steady state

BROADCOM
Connecting everything

Adjustments to Derivative Weight
Effects of Large Derivative Weight when Control Loop Delay is Small (Queue Size)

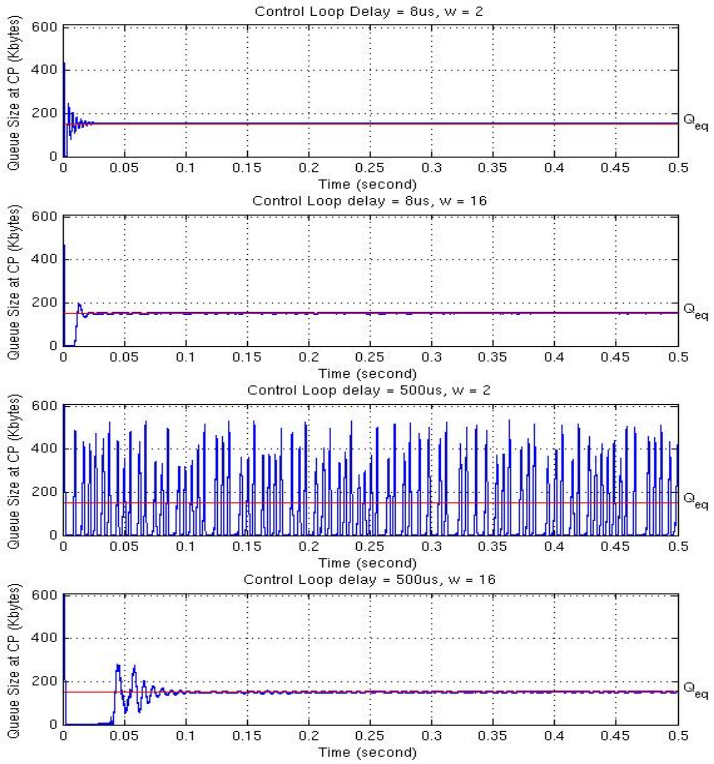

Ctrl Loop Delay	w	# of drop	Throughput
8us	2	6	9.992Gbps
8us	16	41	9.970Gbps
500us	2	419	8.056Gbps
500us	16	451	9.728Gbps

A ‘w’ of 16 works well when Qeq = 150kbytes for small and large control loop delays.

BROADCOM
Connecting everything

Adjustments to Derivative Weight
Effects of Large Derivative Weight when Control Loop Delay is Small (Throughput)

Ctrl Loop Delay	w
8us	2
8us	16
500us	2
500us	16

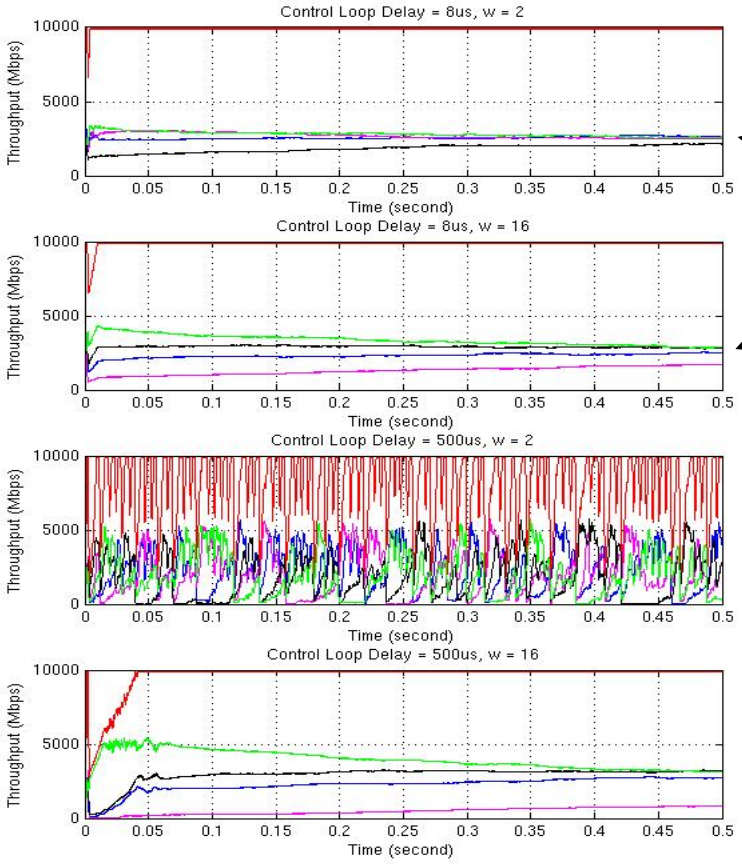

Fairness appears to degrade when derivative weight is increased.

BROADCOM
Connecting everything

Overview

- Goals & System Parameters
- Key Observations
- Effects of Large Control Loop Delay
- Effects of Reducing Qeq
- Conclusions

BROADCOM
Connecting everything®

Effects of Eeq
Effects of Control Loop Delay (Queue Size @ CP)

Qeq = 24kBytes

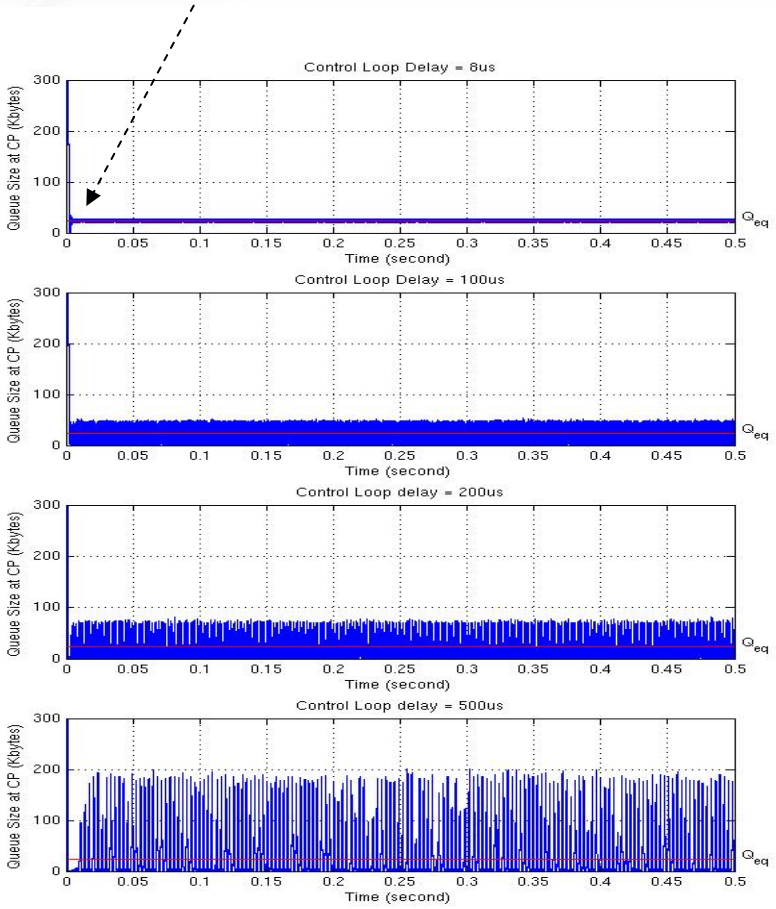

Ctrl Loop Delay	# of drops	Throughput
8us	0	10Gbps
100us	0	9.996Gbps
200us	0	9.608Gbps
500us	196	8.731Gbps

At 200us and above, underutilization begins to occur.

BROADCOM
Connecting everything

Symmetric Topology Single HS – Non Bursty Effects of Control Loop Delay (Throughput)

Ctrl Loop Delay

8us

100us

200us

500us

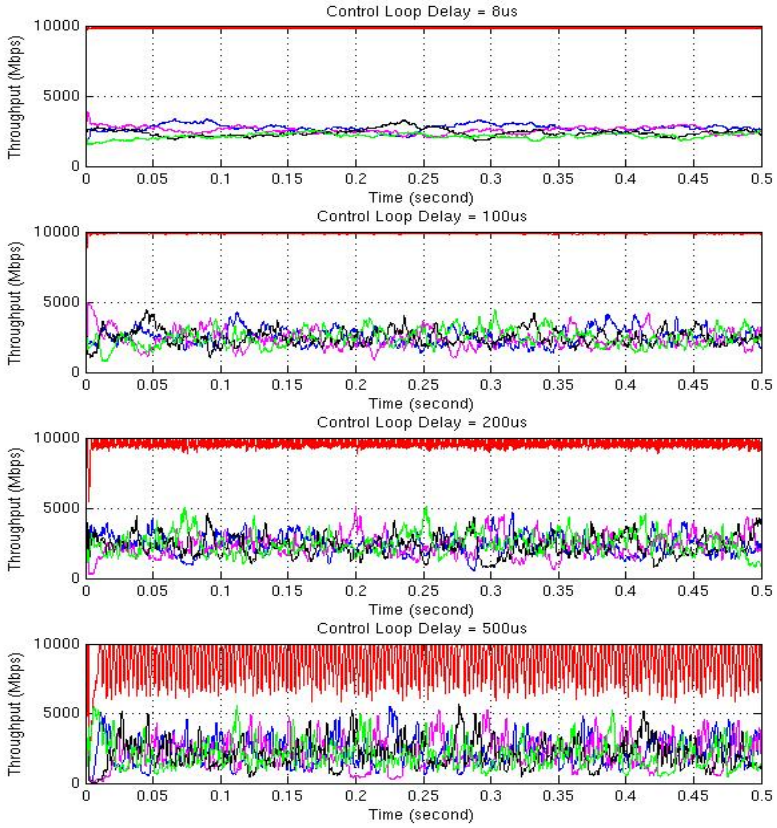

Adjustments to Derivative Weight
Effects of Control Loop Delay (Queue Size @ CP)

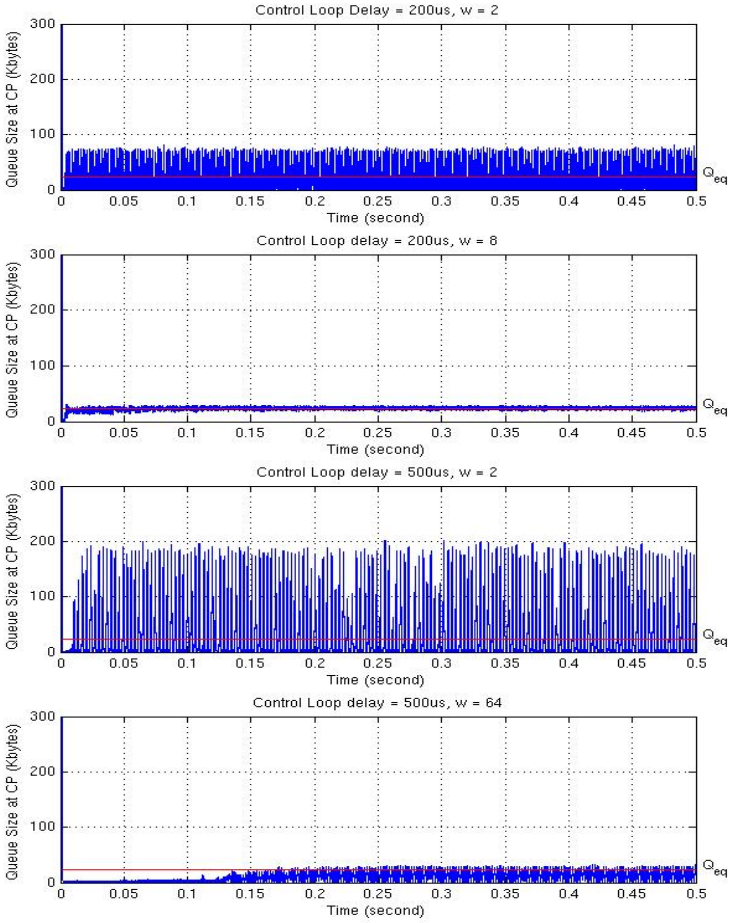

Ctrl Loop Delay	w	# of drop	Throughput
200us	2	0	9.608Gbps
200us	8	0	9.991Gbps
500us	2	196	8.731Gbps
500us	64	214	9.46Gbps

For the reduced Qeq case, identifying a derivative weight with no underutilization appeared to be challenging. Still need to see if other parameter modifications can aid with this (if needed).

BROADCOM
Connecting everything

Adjustments to Derivative Weight
Effects of Control Loop Delay (Throughput)

Ctrl Loop Delay	w
200us	2
200us	8
500us	2
500us	64

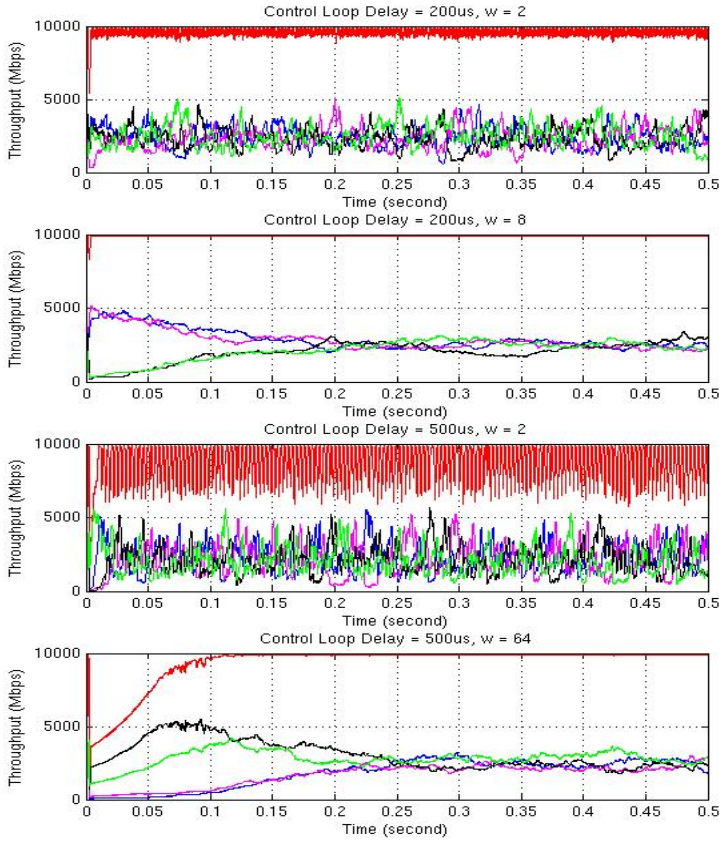

BROADCOM
Connecting everything

Adjustments to Derivative Weight
Effects of Large Derivative Weight when Control Loop Delay is Small (Queue Size)

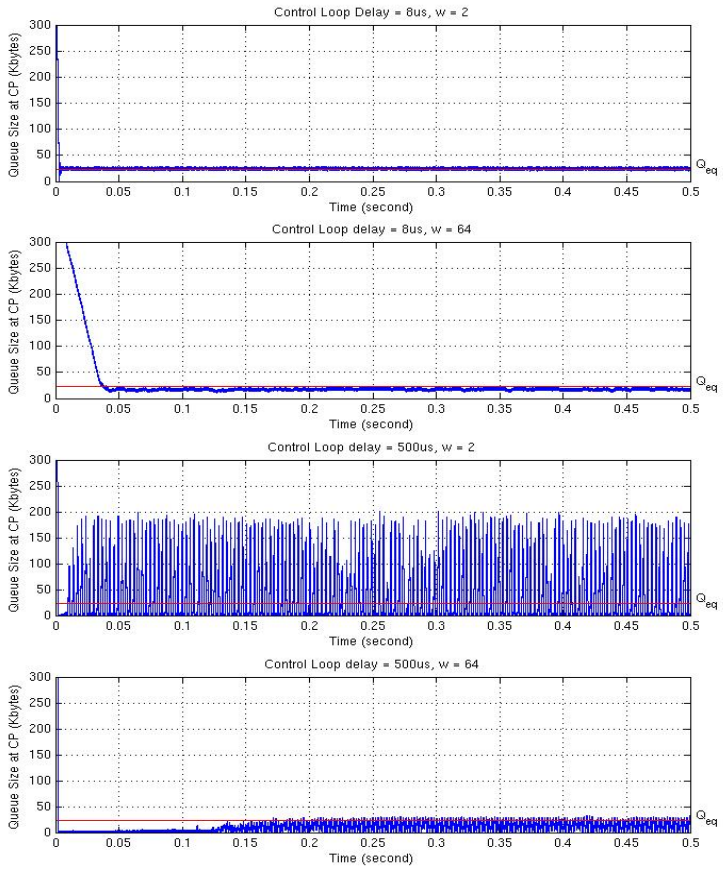

Ctrl Loop Delay	w	# of drop	Throughput
8us	2	0	9.999Gbps
8us	64	0	9.997Gbps
500us	2	196	8.731Gbps
500us	64	214	9.416Gbps

Using a high derivative weight places more emphasis on the Qdelta and can result in longer convergence times.

BROADCOM
Connecting everything

Adjustments to Derivative Weight
Effects of Large Derivative Weight when Control Loop Delay is Small (Throughput)

Ctrl Loop Delay	w
8us	2
8us	16
500us	2
500us	16

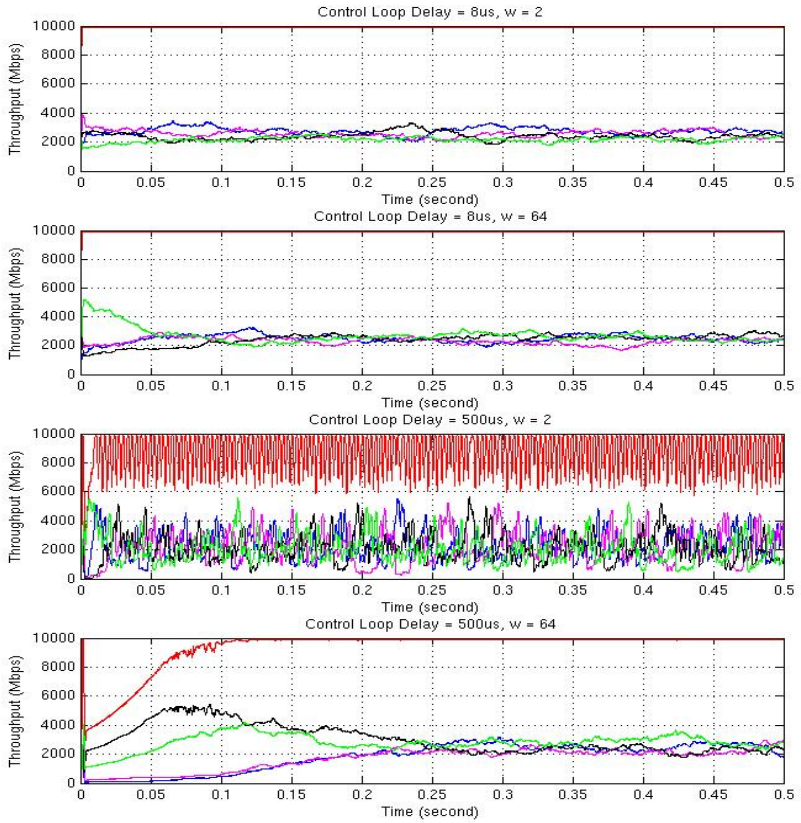

BROADCOM
Connecting everything

Overview

- Goals & System Parameters
- Key Observations
- Effects of Large Control Loop Delay
- Effects of Reducing Qeq
- Conclusions

BROADCOM
Connecting everything®

Conclusions

- Adjustments to derivative weight ‘w’ enhances performance under large control loop delay conditions.
- Identifying an appropriate derivative weight ‘w’ parameter across a network with varying control loop delay conditions appears more challenging when Qeq is small.

BROADCOM
Connecting everything®