

Toward A Scalable and Robust QCN

M. Gusat, C. Minkenberg, R. Luijten
IBM ZRL GmbH

with G. Roeck and R. Jain

Nov. 06, 2007

Outline

- • Preliminary comparison of Q-ECM, QCN, QCN-SONAR
- QCN w/ non-negligible RTT and Adaptive Sampling
 1. analytical model
 2. simulation validation
- QCN+ : Robust QCN
 1. architecture, operation, key features
 2. selected results
- Conclusions and Future Work

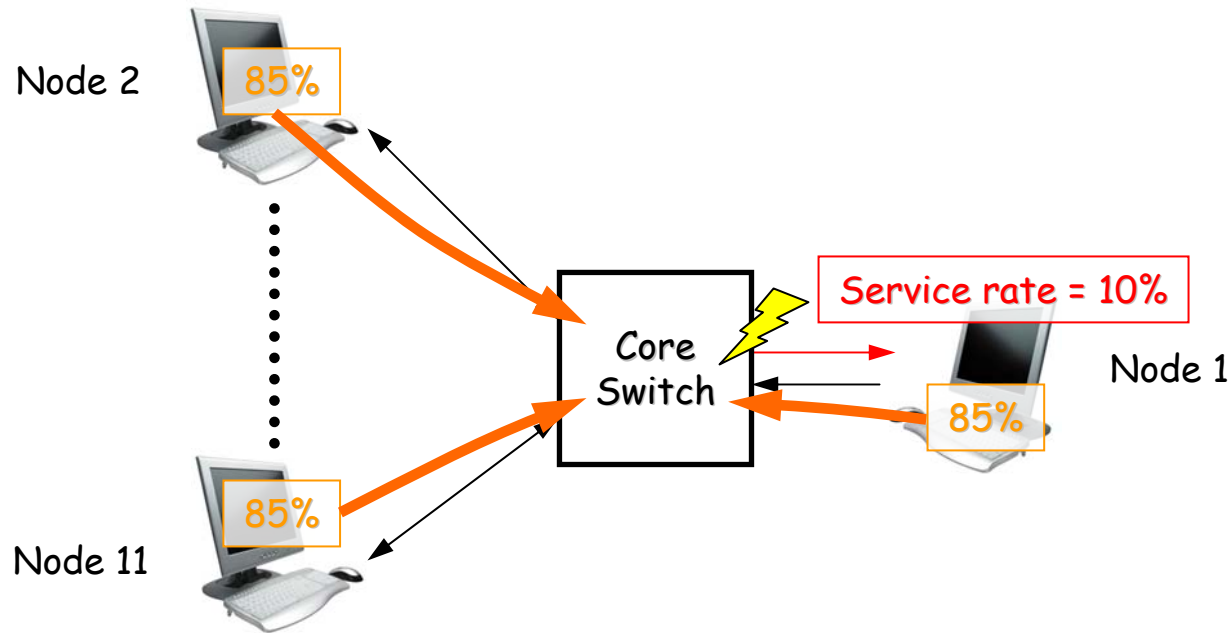
Congestion Management Parameters

- Common settings
 - $W = 2$
 - $Q_{eq} = M / 4$
 - $R_{unit} = R_{min} = 1 \text{ Mb/s}$
 - $G_d = 0.5 / 255$
 - $G_i = 0.1 * (R_{link} / R_{unit}) / 255$
 - ECM_{MAX} enabled, $Q_{mc} = M$
 - $ECM_{0,0}$ disabled
 - Drift enabled: 4 Mb/s increase every 20 ms
 - 8-bit quantization
- ECM-Q
 - ECM but with quantized feedback; separate quantization of Q_{off} and Q_{delta} with respect Q_{eq}
 - $P_{sample} = 1\%$
- QCN
 - $P_{sample} = [1, 10] \%$
 - QCN active increase
 - $to_thresh = packet_size / P_{sample}$
 - $R_i = 12 \text{ Mb/s}$
 - No Fb-hat, no cycle shrinking
- QCN-SONAR
 - WP timer = 10 ms
 - WE timer = 2 ms
 - SF timer = 0.5 ms
 - "not congested" threshold = 15000 B
 - "not congested" interval = 10 ms
 - No FR1 target-rate adjustment

Switch & Adapter Parameters

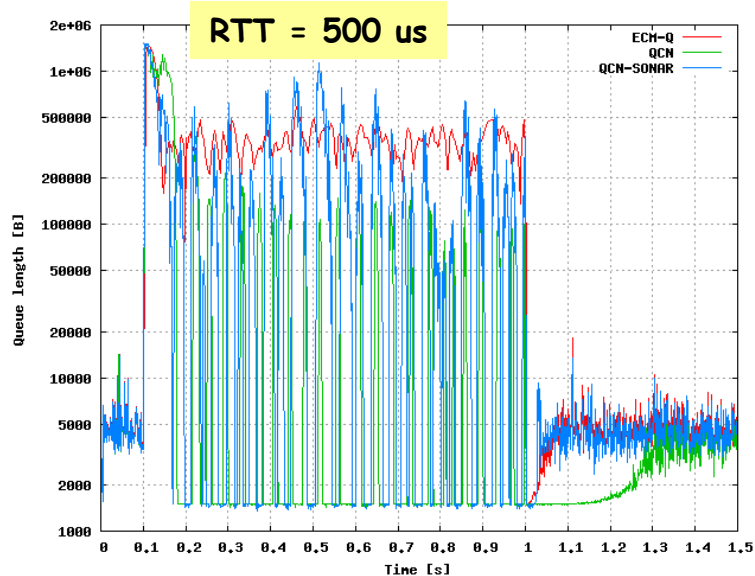
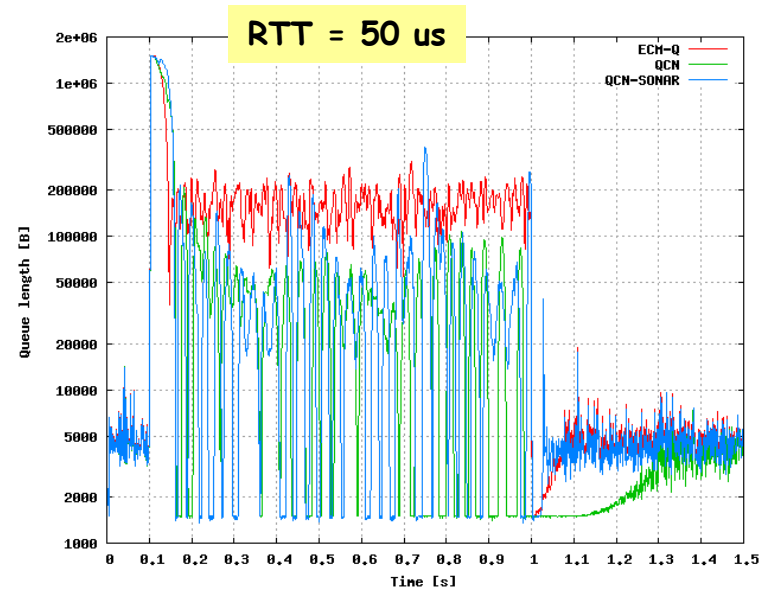

- Switch parameters
 - $M = 150 \text{ KB}$ or 1.5 MB per port
 - Dedicated per input, shared across all outputs
 - Configurable OQ limit; frames are dropped when OQ length exceeds limit (disabled)
 - PAUSE enabled or disabled
 - Applied on a per input basis based on local high/low watermarks
 - $\text{watermark}_{\text{high}} = M - \text{rtt} * \text{bw}$ KB
 - $\text{watermark}_{\text{low}} = M - \text{rtt} * \text{bw} - 10 \text{ KB}$
 - If disabled, frames dropped when input partition full
- Adapter parameters
 - Virtual output queuing, round-robin VOQ service
 - Input buffer size $IB = 1.5 \text{ MB}$, partitioned per VOQ
 - Drop when VOQ full
 - Output buffer size $OB = 1.5 \text{ MB}$
 - No limit on number of rate limiters
 - PAUSE enabled
 - $\text{watermark}_{\text{high}} = OB - \text{rtt} * \text{bw}$ KB
 - $\text{watermark}_{\text{low}} = \text{watermark}_{\text{high}} - 10 \text{ KB}$

Output-Generated Single-Hop Hotspot






- All nodes: Uniform destination distribution, load = 85% (8.5 Gb/s)
- Node 1 service rate = 10%
- One congestion point
 - Hotspot degree = 10
 - All flows affected

Hot queue length, PAUSE disabled, M = 150 KB/port



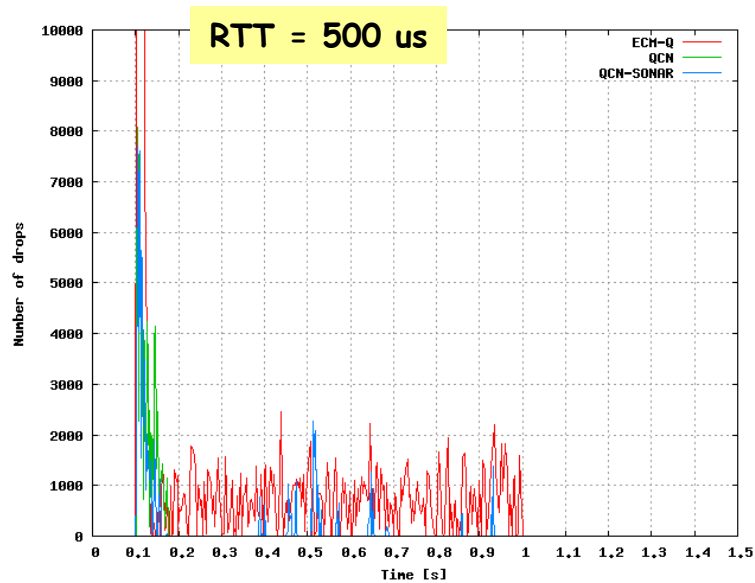
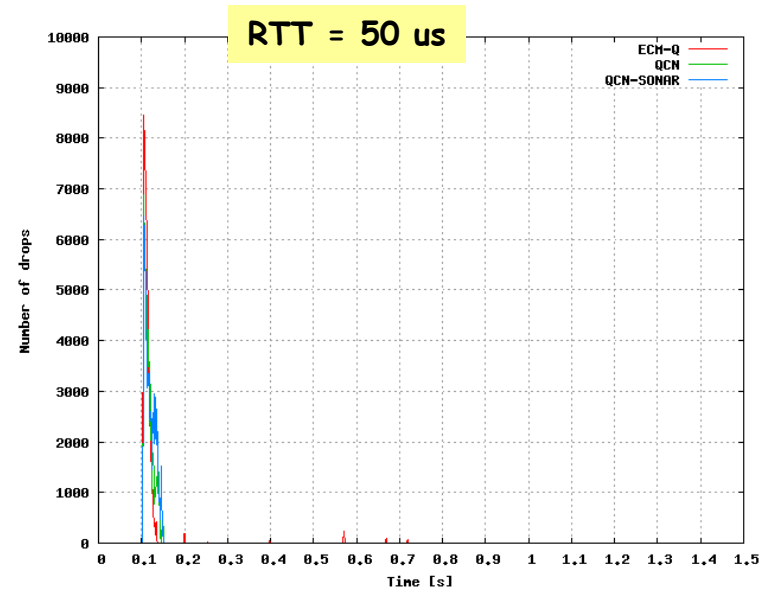

- Initial reactions similarly fast
- QCN and SONAR oscillating much stronger than ECM
- ECM steady state queue length too high
 - Related to only issuing Fb when $Q > Q_{eq}$? (suppresses spurious RL creation)

Aggregate throughput, PAUSE disabled, M = 150 KB/port



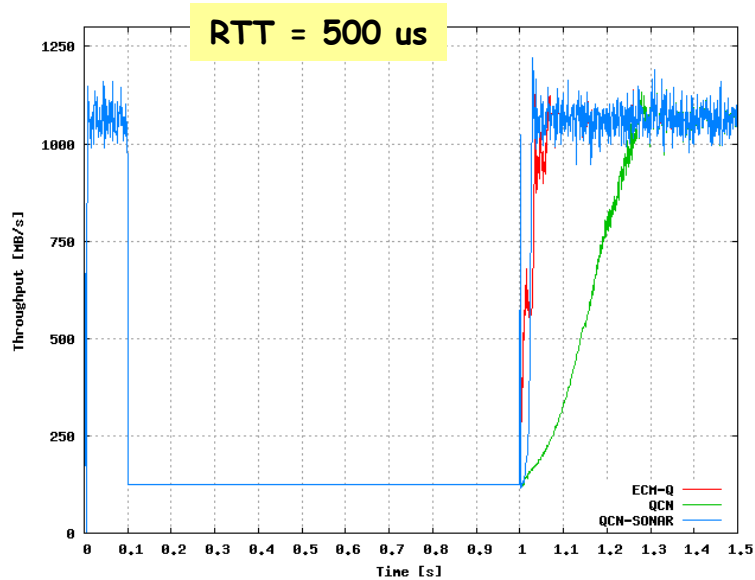


- Aggregate throughput recovers quickly in all cases
- ECM exhibits some underutilization because of drops (see next slide)

Frame drops, PAUSE disabled, M = 150 KB/port



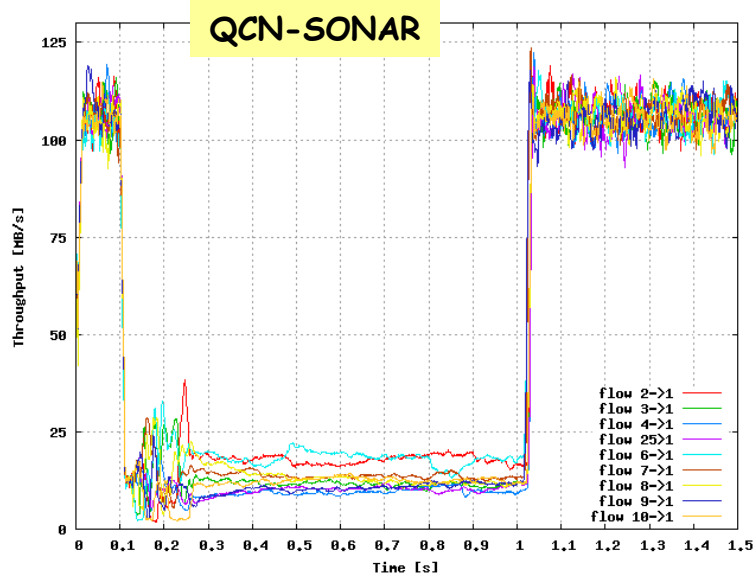

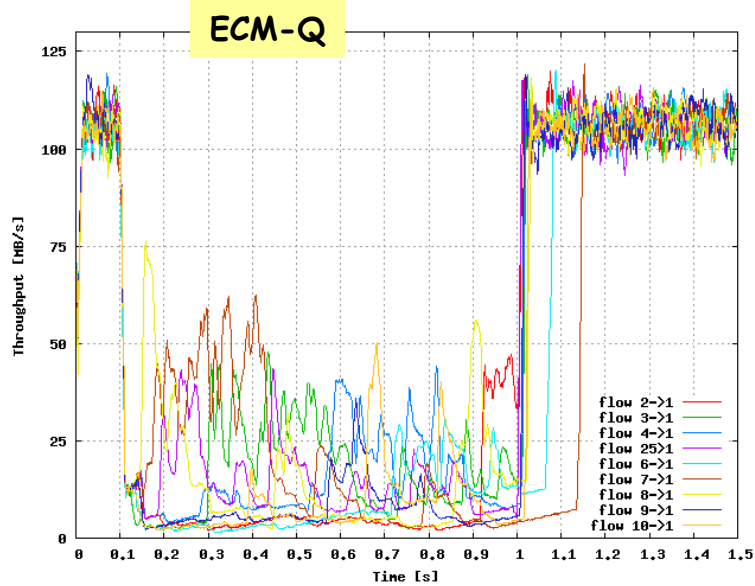
- RTT = 500 us: ECM has more drops (queue too long)

Hot port throughput, PAUSE disabled, M = 150 KB/port



- No underutilization issues on hot port
- Recovery time: SONAR < ECM < QCN

Per-flow throughput, PAUSE disabled, RTT = 5 us



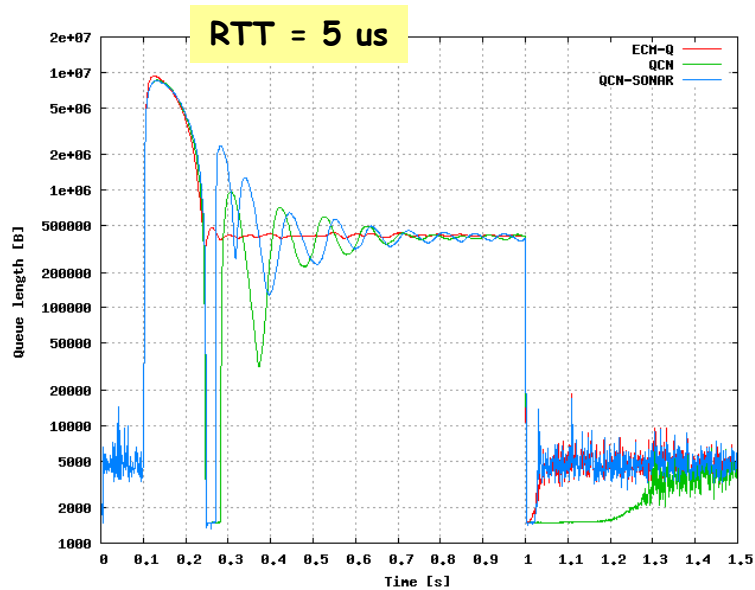


- QCN and SONAR exhibit decent fairness
- ECM exhibits strong rate fluctuations

PAUSE

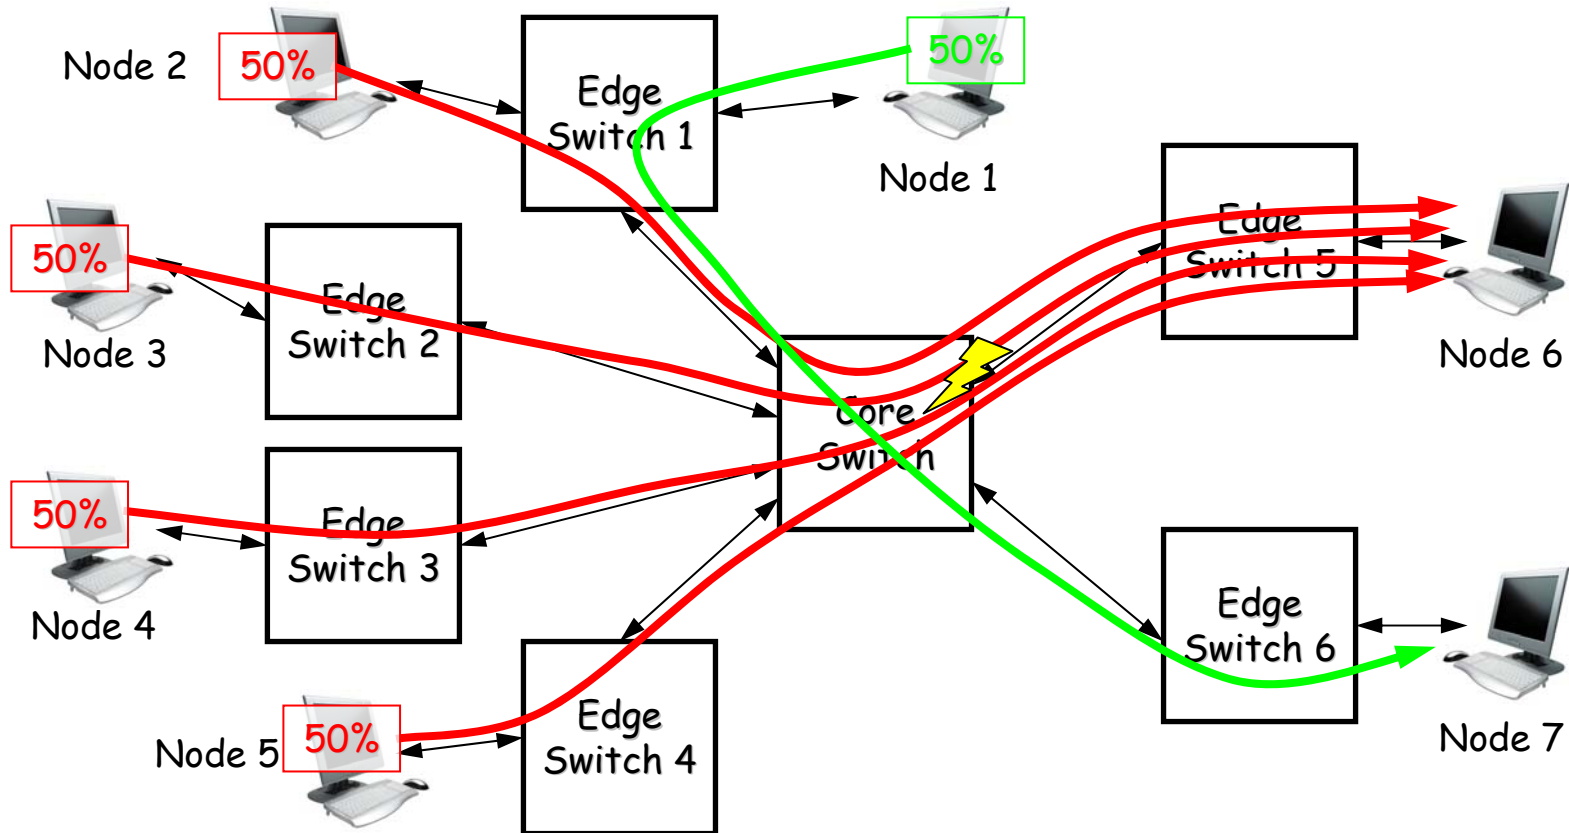
- PAUSE enabled
- Switch and adapter memory size and thresholds configured to ensure losslessness
 - $M = 1.5 \text{ MB/port}$ ($\text{RTT} = 500 \text{ us} \rightarrow 2 * 625 \text{ KB}$ of headroom needed)
- No OQ limit \rightarrow no drops

Hot queue length, PAUSE enabled, $M = 1.5$ MB/port



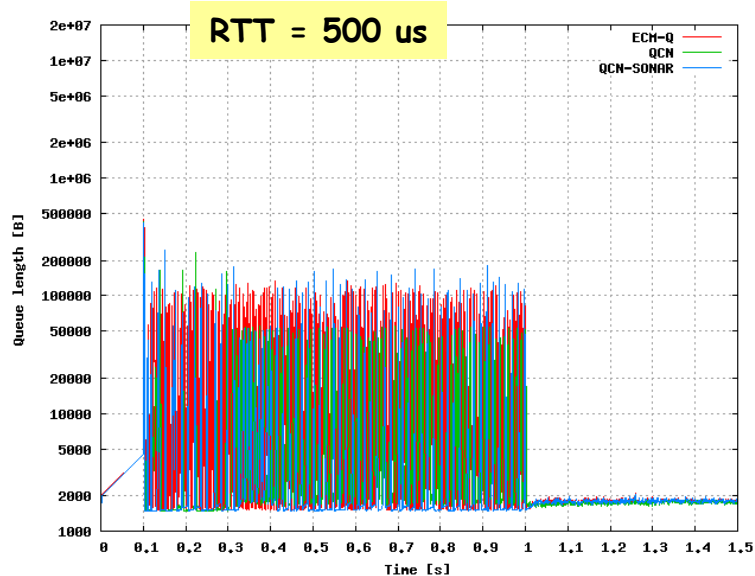


- Initial reactions similarly fast
- ECM queue length stable, but steady state level again too high
- QCN and SONAR larger oscillations, but lower steady state level
- With large M , drift is important for QCN/SONAR (see backup)

Baseline Input-Generated Hotspot



- Four culprit flows of 5 Gb/s each from nodes 2, 3, 4, 5 to node 6 (hotspot)
- One victim flows of 5 Gb/s from node 1 to node 7
- Fair allocation provides 2.5 Gb/s to all culprits and 5 Gb/s to the victim


Hot queue length, PAUSE disabled, $M = 150$ KB/port




- RTT = 5, 50 us: Exemplary behavior for all schemes
- RTT = 500 us: Very strong oscillations

Hot port throughput, PAUSE disabled, $M = 150$ KB/port


RTT = 5 us



RTT = 50 us

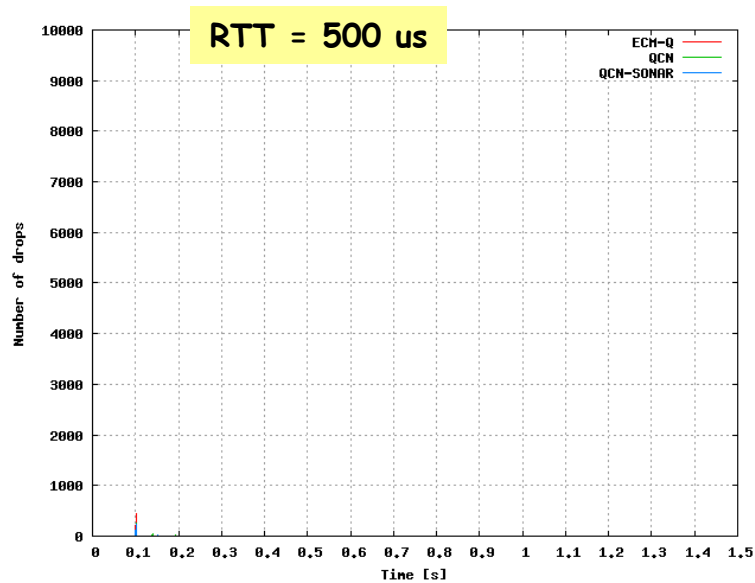

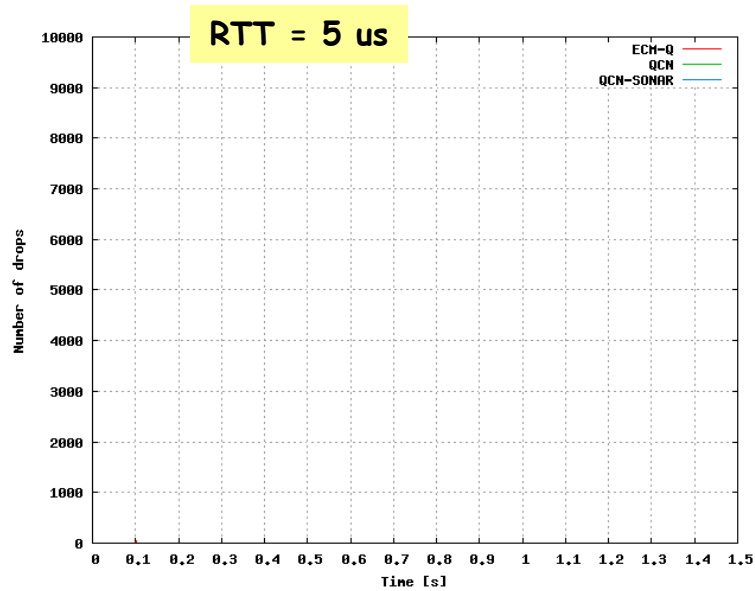


RTT = 500 us





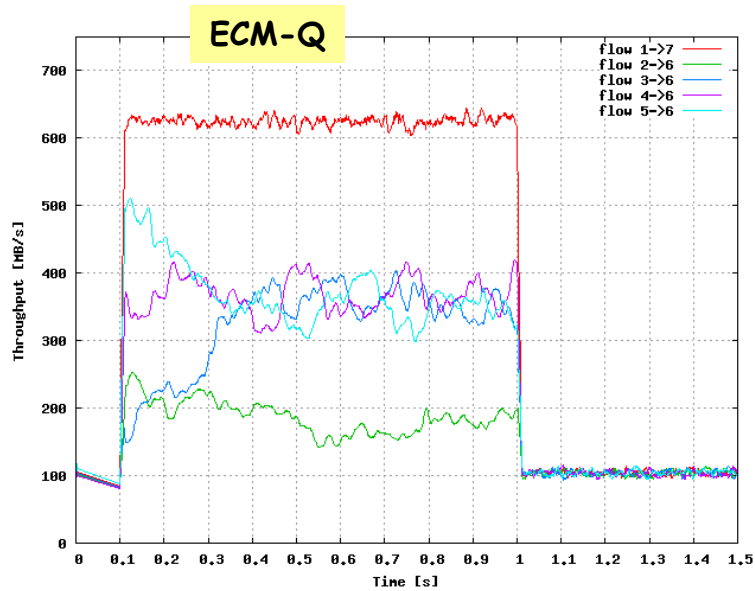
- RTT = 500 us: Queue oscillations cause underflow and loss of throughput
- SONAR suffers most

Frame drops, PAUSE disabled, M = 150 KB/port



- Practically no drops ... low 100s

Hot port throughput, PAUSE disabled, M = 150 KB/port, RTT = 5 us






- "Victim" flow not affected
- Poor fairness in all cases

PAUSE

- PAUSE enabled
- Switch and adapter memory size and thresholds configured to ensure losslessness
 - $M = 1.5 \text{ MB/port}$ ($\text{RTT} = 500 \text{ us} \rightarrow 2 * 625 \text{ KB}$ of headroom needed)
- No OQ limit \rightarrow no drops


Hot queue length, PAUSE enabled, $M = 1.5$ MB/port




- SONAR shows strong initial oscillations even with short RTT
- These are not related to PAUSE, but to the memory size (same thing happens when $M=1.5$ MB without PAUSE)

Hot port throughput, PAUSE enabled, $M = 1.5$ MB/port


RTT = 5 us



RTT = 50 us






RTT = 500 us



- RTT = 5, 50 us: SONAR exhibits underflow during transient
- RTT = 500 us: ECM exhibits underflow


Per-flow throughput, PAUSE enabled, M = 1.5 MB/port




- "Victim" flow not affected
- Poor fairness in all cases

Input-generated hotspot, 100 sources, RTT = 5 us


Hot queue length



Hot port throughput




Frame drops




- $M = 150$ KB/port; OQ limit = 1.5 MB, pause disabled
- Hot queue is stable for all schemes
- ECM has more drops, but less underrun

Input-generated hotspot, 100 sources, RTT = 50 us


Hot queue length



Hot port throughput




Frame drops




- $M = 150$ KB/port; OQ limit = 1.5 MB, pause disabled
- Hot queue reasonably stable for all schemes, but
 1. ECM too high
 2. QCN/SONAR very jittery
- ECM has more drops, significantly less underrun

Input-generated hotspot, 100 sources, RTT = 500 us


Hot queue length



Hot port throughput

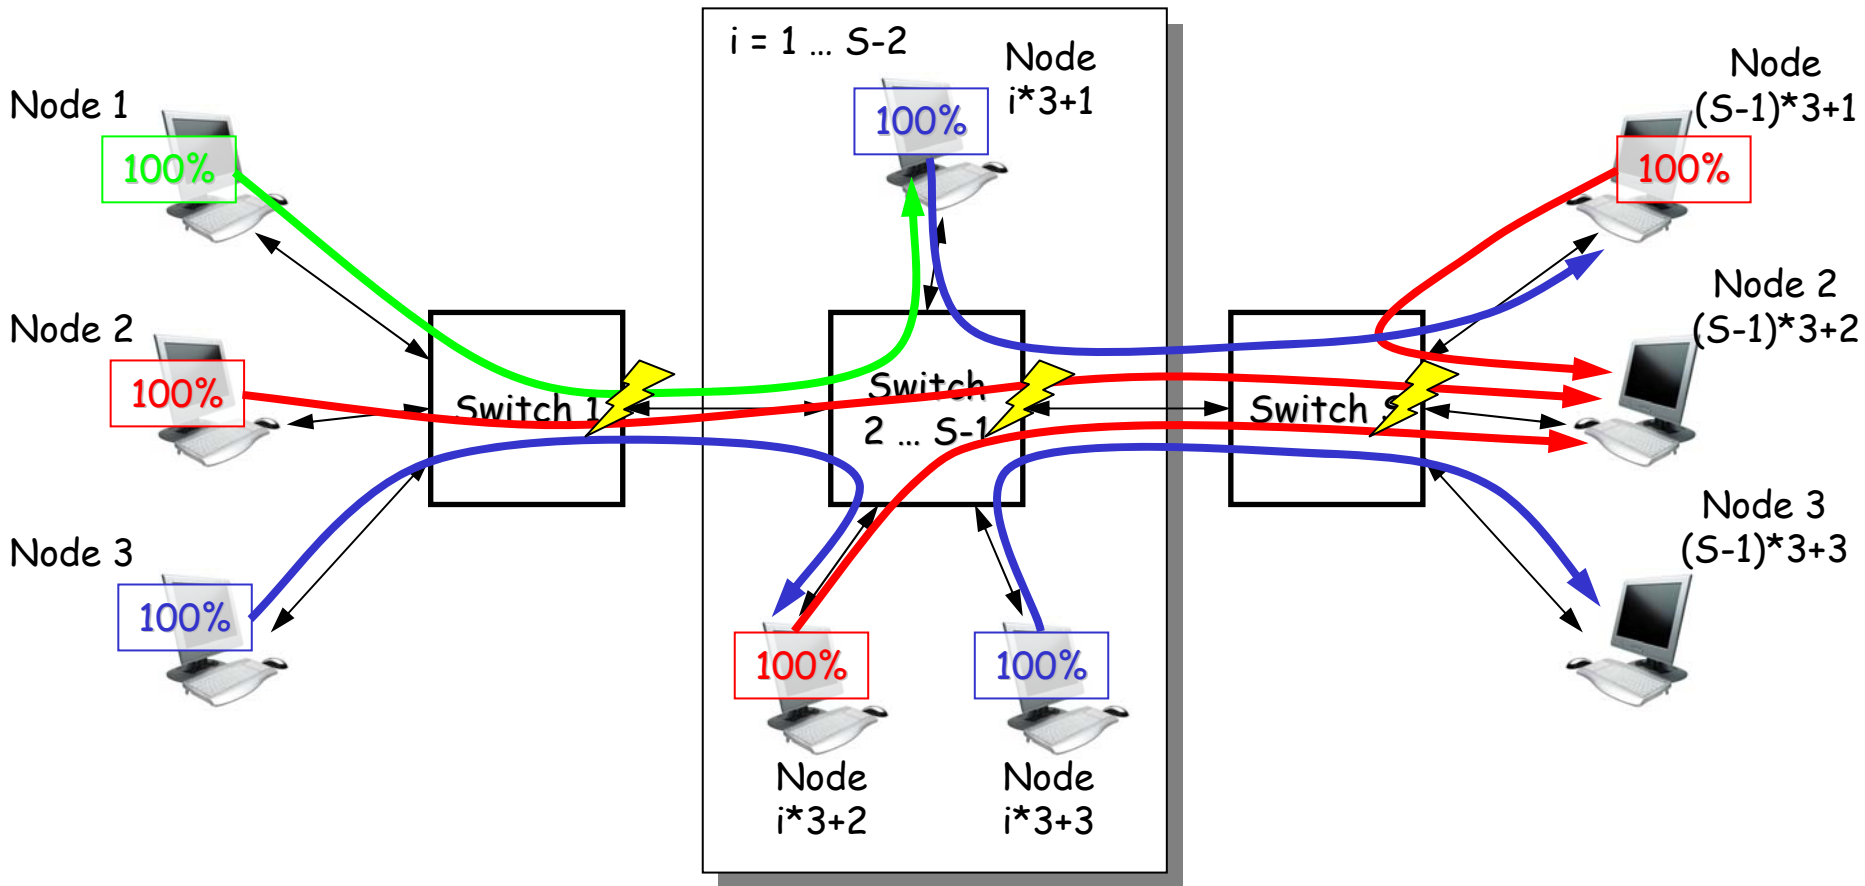


Frame drops



- $M = 150$ KB/port; OQ limit = 1.5 MB pause disabled
- Hot queue barely stable for all schemes
- Severe underrun for QCN and QCN-SONAR, ECM coping reasonably


20-stage Hotspot




- $S = 20$ switches (stages); 3 nodes per switch
 - Node $i*3+1$ sends to node $(i+1)*3+1$, node $i*3+2$ sends to node $(S-1)*3+2$, node $i*3+3$ sends to node $(i+1)*3+2$ for $i = 0 \dots S-2$.
- 100% load from all nodes except $(S-1)*3+2$ and $(S-1)*3+3$
- Node $(N-1)*3+2$ receives traffic from S sources
- N hotspots, $M = 150$ KB/port, per-hop RTT = $[2, 20]$ us

Hotspot throughput, PAUSE disabled

Per-hop RTT = 2 us



Per-hop RTT = 20 us





- ECM exhibits minimal underrun
- Significant underrun for QCN and QCN-SONAR

Queue length , PAUSE disabled



Per-hop RTT = 2 us

Per-hop RTT = 20 us

ECM



QCN-SONAR





Hotspot throughput, PAUSE disabled, M = 150 KB/port



Per-hop RTT = 2 us

Per-hop RTT = 20 us

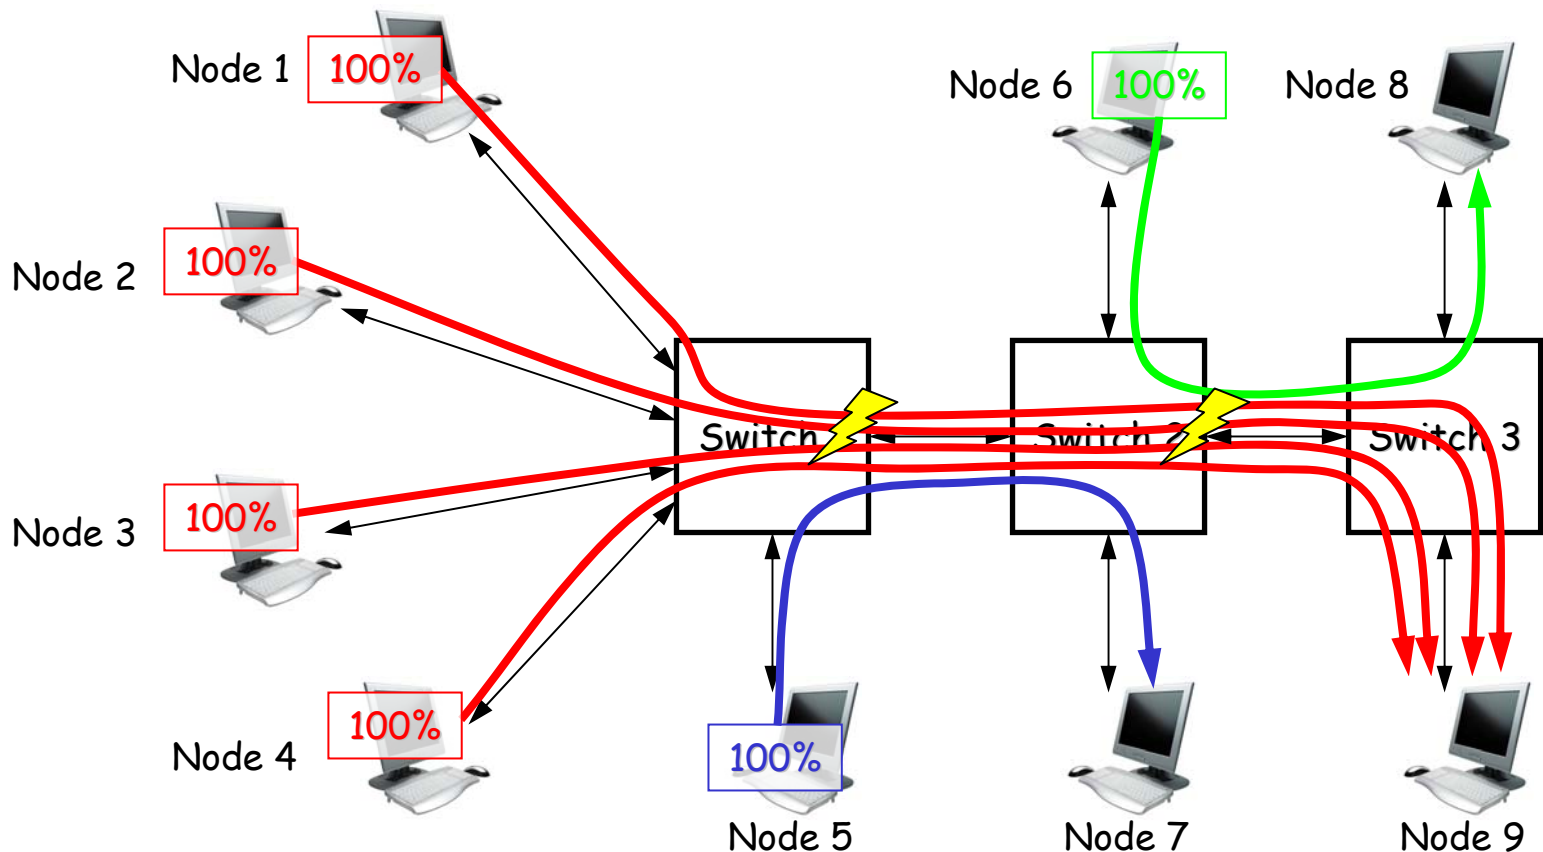
ECM



QCN-SONAR



Parking Lot Scenario





- Four hot flows of 10 Gb/s each from nodes 1, 2, 3, 4 to node 9 (hotspot)
- Two cold flows of 10 Gb/s from node 5 to 7 and 6 to 8
- Max-min fair allocation provides 2.0 Gb/s to **all** flows
- Proportionally fair allocation provides 1.67 Gb/s to all hot flows and 3.33 Gb/s to all cold flows
- **Pause disabled, M = 150 KB/port, ingress = egress buffer size = 1.5 MB, no OQ limit**

Per-flow throughput, PAUSE disabled, M = 150 KB/port



Per-hop RTT = 10 us

Per-hop RTT = 100 us

ECM



QCN-SONAR





Output queue lengths, PAUSE disabled, M = 150 KB/port



Per-hop RTT = 10 us

Per-hop RTT = 100 us

ECM



QCN-SONAR



Outline

- Preliminary comparison of Q-ECM, QCN, QCN-SONAR
- • QCN w/ non-negligible RTT and Adaptive Sampling
 1. Model
 2. Validation
- QCN+ : Robust QCN
 1. architecture, operation, key features
 2. sim results
- Conclusions and Future Work

QCN as a Control Loop w/ Lag (T) and Delay (τ)

What happens when delay exceeds the dominant lag?

- Delay fundamentally affects closed loop control. Critical when $T > \tau$
- QCN¹: load sensor model reduced to 1st order system w/ dominant lag (sampling time constant T) and non-negligible delay ($\tau = \text{RTT}_{e2e}$)

$$1. \quad QCN^1(s) = \frac{e^{-\tau s}}{1 + sT}$$

Note: QCN's control loop is a higher order system

$$2. \quad \tau = \tau_{\text{queue}} + \tau_{\text{transport}} \approx \tau_{\text{queue}},$$

Conservative assumption in datacenters

$$3. \quad T = 1/f_s = 1/(P_s * \lambda_{\text{aggr}}) = 1/(P_s * n * \lambda(t)), \quad n = \text{no. flows @ CP}, \lambda(t) = \text{rate}$$

$$1. \quad 0 < P_s \leq 1,$$





$$2. \quad \text{QCN: } P_s = 1..10\%$$

$$4. \quad (1 + 3) \Rightarrow QCN^1(s) = \frac{n \cdot \lambda \cdot e^{-\tau s}}{n \cdot \lambda + \frac{s}{P_s}}$$

Conflicting effect of “ $P_s : \tau$ ”-ratio during HS

- 1. $P_s \uparrow \Rightarrow T \downarrow$ (improved observability)
- 2. $\tau = \text{RTT}_{e2e} \uparrow$ (2-5 orders of magnitude)

Impact of variable sampling frequency P_s @ constant RTT (=19)





Why QCN's Adaptive Sampling Depends on RTT probing?

- Observations
 1. Whenever delay exceeds sampling lag the loop becomes unstable
 1. Hence the intrinsic conflict between increasing P_s and delay stability
 2. No clear trade-off is possible w/ RTT knowledge
 2. Sampling is aggregate @ CP, while F_b is per flow @ RP
 3. CP does not know RTT, nor "n" (# flows)
 4. Flooding RPs w/ bursts of outdated feedback requires adaptivity
 1. near RP's benefit directly from an increased P_s
 2. remote RP's don't... (must filter - decimation, Kalman)

see "Effects of long RTT [and P_s] on QCN"

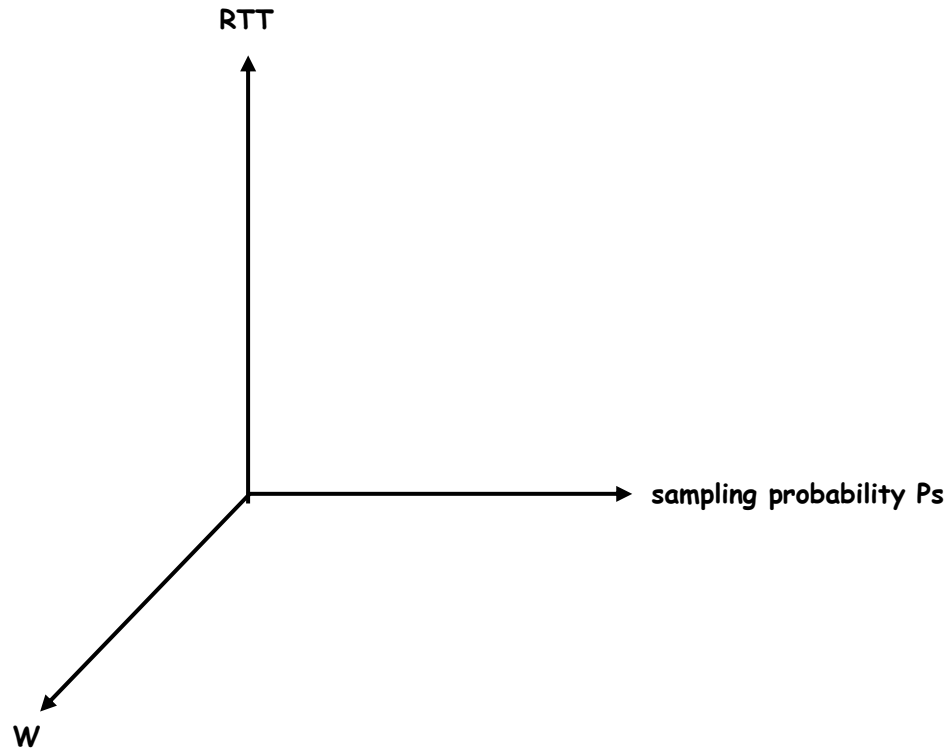
Impact of variable RTT @ constant sampling $1/P_s (= 2)$



Simulation Parameters




- Baseline scenario, $t_{\text{hotspot}} = 0.1 - 1.0 \text{ s}$
- $M = 1.5 \text{ MB/port}$
- Unlimited input buffers
- $Q_{\text{eq}} = 375 \text{ KB}$
- $P_s = 1\%$ (QCN: 1-10%)
- QCN active increase
 - $\text{to_thresh} = \text{packet_size} / \text{p_sample}$
 - $R_i = 12 \text{ Mb/s}$
- Drift enabled: 4 Mb/s every 20 ms
- ECM_MAX enabled, $Q_{\text{mc}} = 1.5 \text{ MB}$
- No ECM_(0,0), no PAUSE
- Per-link RTT = [2 us, 20 us, 50 us, 200 us, 500 us, 1ms]
 - Note that $\text{RP} \Leftrightarrow \text{CP RTT} = 2 * \text{link RTT}$
- 8-bit quantization
- $W = [2, 8, 32]$

What We Analyzed




Conflicting requirements on W

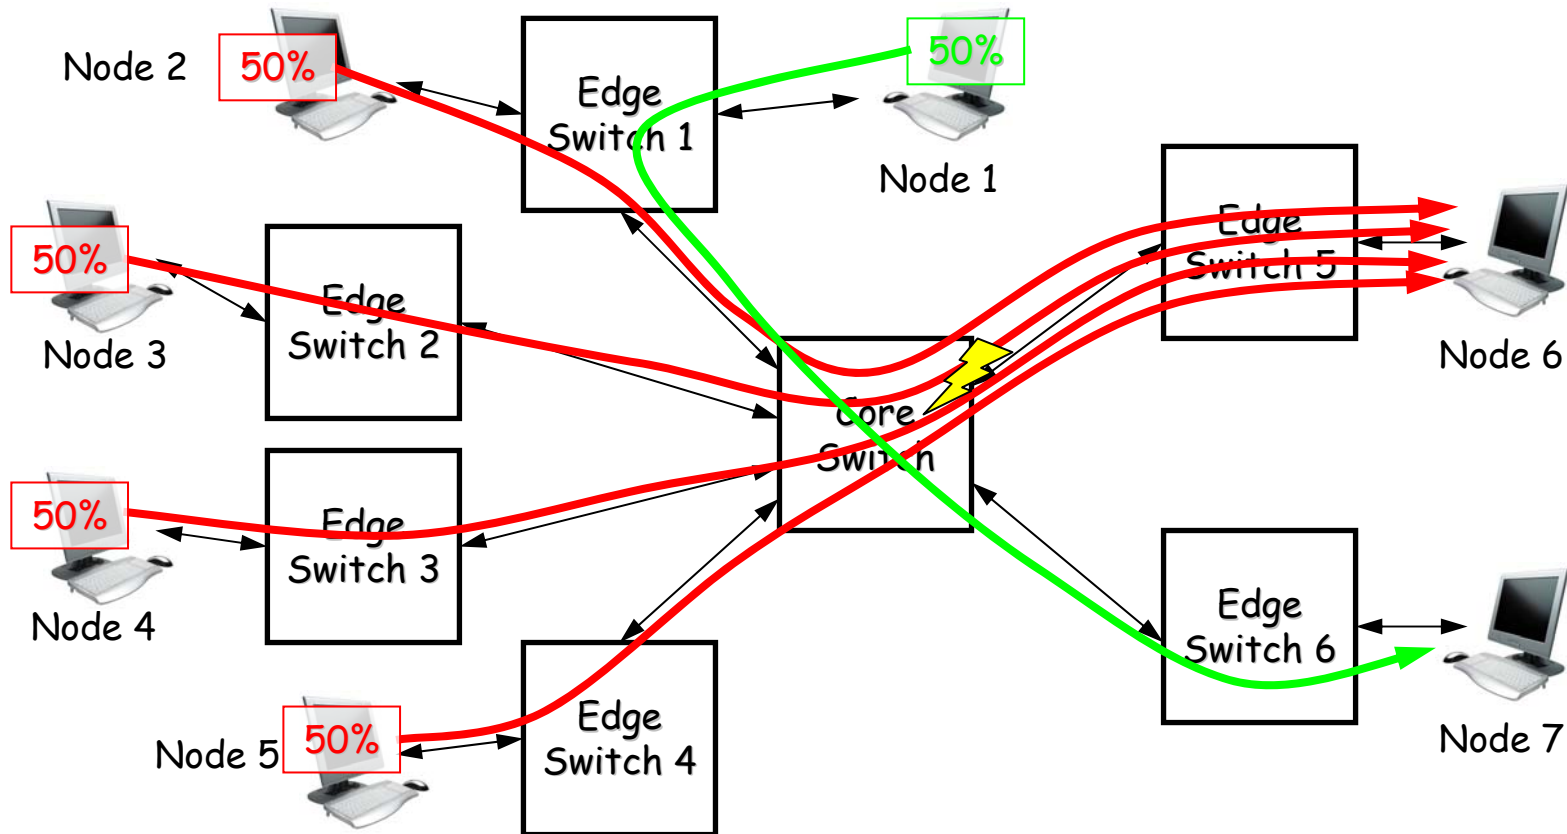
IG vs OG



Hot queue length, QCN, $P_{\text{sample_base}} = 1\%$, $\text{RTT} = 2 \text{ us}$









Baseline Input-Generated Hotspot









- Four culprit flows of 5 Gb/s each from nodes 2, 3, 4, 5 to node 6 (hotspot)
- One victim flows of 5 Gb/s from node 1 to node 7
- Fair allocation provides 2.5 Gb/s to all culprits and 5 Gb/s to the victim







Hot queue length, QCN, $P_{\text{sample_base}} = 1\%$, vary RTT









Hot port throughput, QCN, $P_{\text{sample_base}} = 1\%$







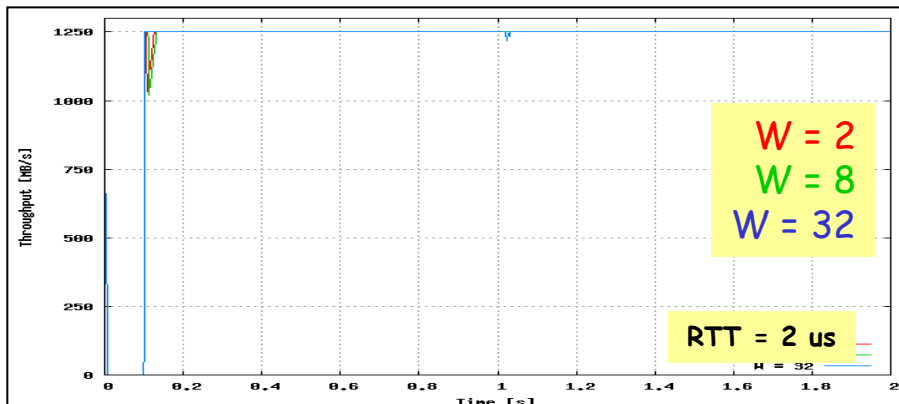
Hot queue length, QCN, Link RTT = 500 us, vary sampling









Hot queue length, QCN, $P_{\text{sample_base}} = 1\%$, with RTT filtering, vary RTT









Hot port throughput, QCN, $P_{\text{sample_base}} = 1\%$, with RTT filtering




Hot queue length, classic ECM, $P_{\text{sample}} = 1\%$, vary RTT



Hot port throughput, classic ECM, $P_{\text{sample}} = 1\%$, vary RTT





Output-Generated Single-Hop Hotspot

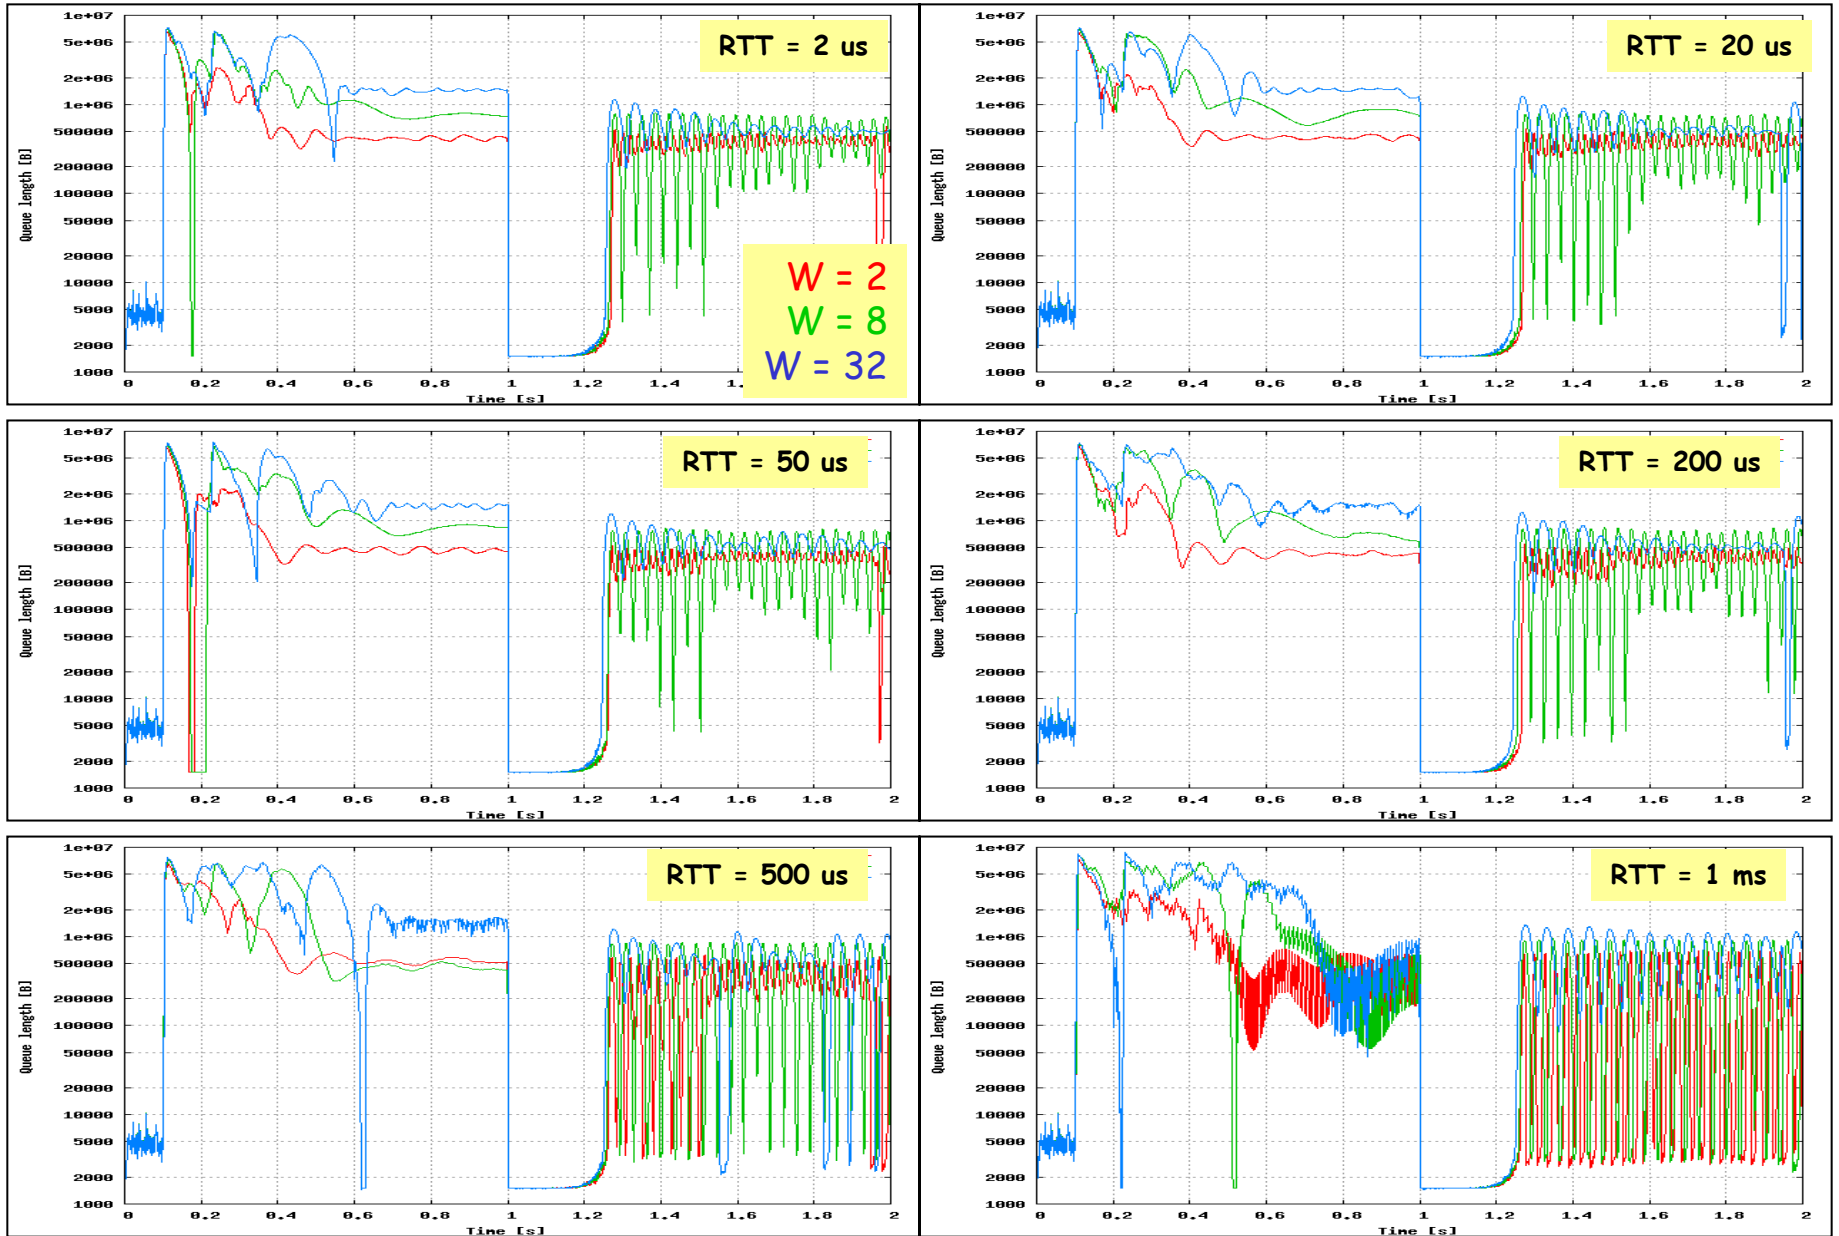


- All nodes: Uniform destination distribution, load = 85% (8.5 Gb/s)
- Node 1 service rate = 10%
- One congestion point
 - Hotspot degree = $N-1$
 - All flows affected







Queue length and thruptut QCN w/o (0,0), vary w



Hot queue length, QCN, $P_{\text{sample_base}} = 1\%$, $w/(0,0)$, vary RTT



Hot queue length, ECM, $P_{\text{sample_base}} = 1\%$, w/ (0,0), vary RTT

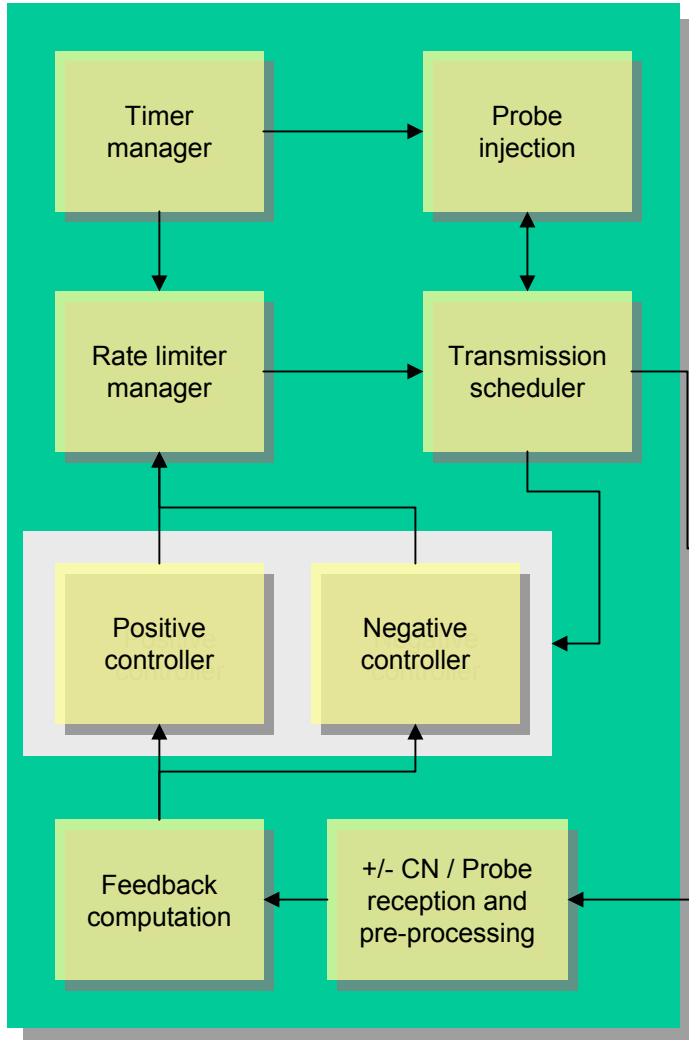


Outline

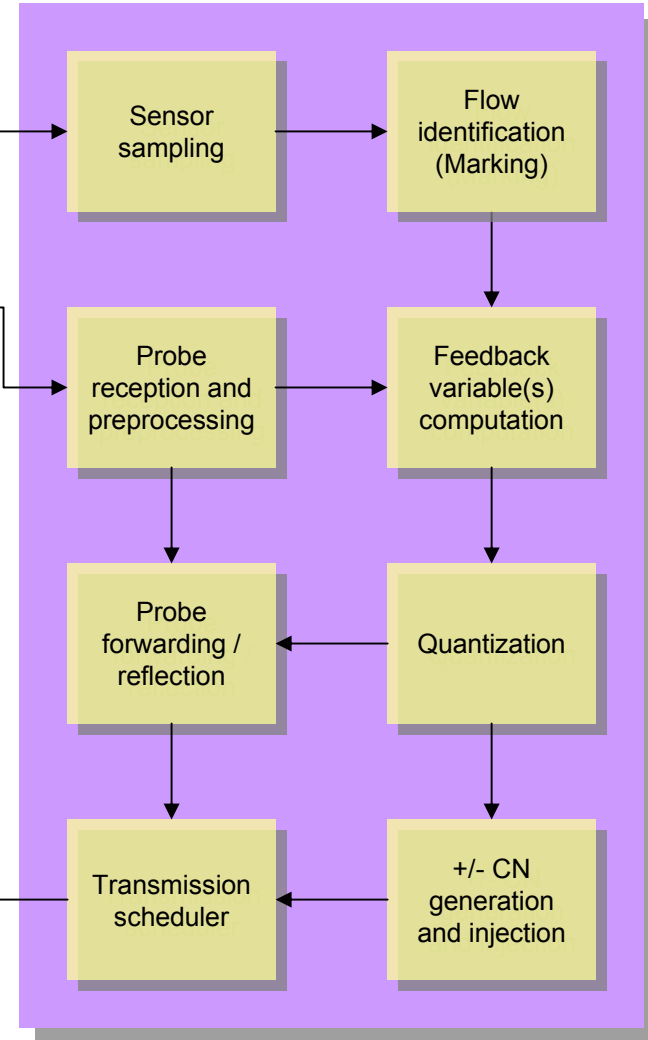
- QCN w/ non-negligible RTT and Adaptive Sampling
 1. analytical model
 2. simulation validation
- • QCN+ : Robust QCN
 1. architecture, operation, key features
 2. sim results
- Conclusions and Future Work

QCN+ Architecture

Reaction Point (RP)



Congestion Point (CP)




Sonar vs. QCN-SP (QCN+)




Throughput at Hotspot


ECM




QCN-SP, adjusted for RTT



QCN-Sonar




QCN-SP, adjusted for RTT and N



Sonar vs. QCN-SP (QCN+)



20-stage Hotspot: Throughput per switch



Conclusions and Future Work

1. Transport and sampling delays affect response and stability
2. Scalability and robustness are required across 1/10/100Gbps, for both small and large datacenters

QCN+ improves on QCN

- Directed probing to CP
 - provides scalability w/ size, topology and switch technology
 - adds a Fb loop which decouples dynamic response from stability
 - enables gain tuning
- Perfectible
 - Changes will affect fairness and/or dynamic response, not stability
 - Open to future improvement while ensuring backward compatibility
- Next steps
 - analyse QCN-Sonar-Fb99
 - quantitatively benchmark QCN+ vs. QCN-Sonar

Recommendation: Due Diligence

- Decide on which part to optimize and which is 'good enough'
 1. Sensor & feedback (load info, state vars, mandatory/optional)
 2. Fb loops: closed and open, probe/tag, signalling, overhead
 3. Rate control algorithms.
- L2 Tip: Invest in (1) => highest quality sensor (optimize)
=> tracking performance, dynamic response and stability
- Timers, counters and settings [p.3] introduce zeroes / poles in the loop => treat with care ... more extensive validation required
- Return to rigorous benchmarking => non-linearities and time-variance elicit ultimately 'proof' by discrete event simulations
 - Investigate modern TCP/REM CM tuned for DC: CU/BIC, XCP, FAST, RCP, HS-TCP...
 - performance w/ the .1au sensor ?
 - is the equilibrium fairness abandoned?

QCN+ : Result of The Stockholm Agreement

- Builds on .1au's collective experience and results
 1. QCN (baseline and fail-safe mode of QCN+)
 2. ECM
 3. FECN
 4. E2CM
 5. Probing.
- Probing to CP
 - maintains the 2pt. paradigm as requested
 - additional CL that decouples stability from dyn. response and provides delay robustness (time invariant scalability)
- Closed loop AQM feedback: -ve and +ve F_b (on demand thru probing).
 - improved dynamic response: faster tracking & high slope SRF




BACKUP

Gain Tuning r1.01




Algorithm implemented by G. Roeck in the current QCN+

1. RTT based loop gain control in RP
2. Accept one negative adjustment per RTT
3. Adjust TO_THRESHOLD based on RTT and current datarate
 1. Set ToThreshold to $\max(\text{TO_THRESHOLD}, \text{RTT} * 2 * \text{rate})$
4. Adjust W (and calculate Fb) based on RTT and current datarate
 1. $N = \langle \text{switch link capacity} \rangle / \langle \text{current rate} \rangle$
 2. $W = \text{baseW} + (\text{RTT} * \langle \text{factor} \rangle / N)$
5. Reduce positive loop gain based on RTT
6. For each ms of RTT, loop gain is reduced by 50%




Hot queue length, PAUSE disabled, drift disabled, M = 150 KB/port



Hot queue length, PAUSE enabled, drift disabled, $M = 1.5$ MB/port

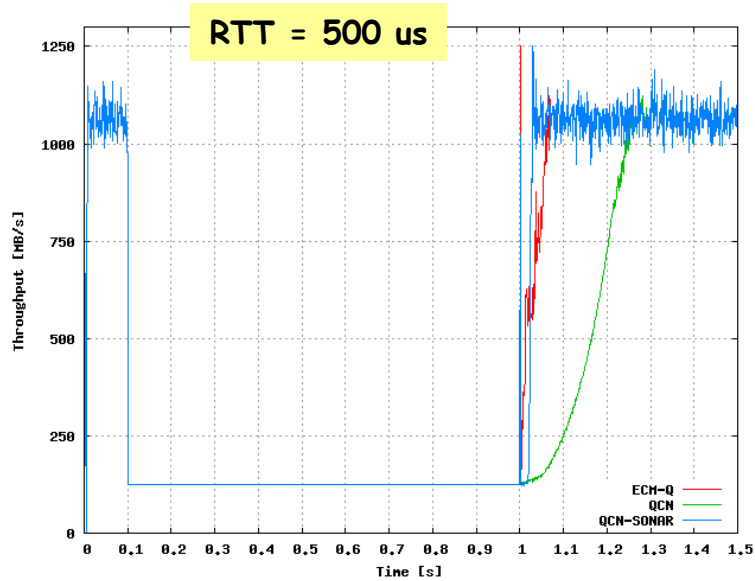




Aggregate throughput, PAUSE enabled, $M = 1.5$ MB/port






- tbd

Hot port throughput, PAUSE enabled, $M = 1.5$ MB/port



- tbd

Per-flow throughput, PAUSE enabled, RTT = 5 us




- tbd


20-stage hotspot frame drops, PAUSE disabled, M = 150 KB/port

ECM



Per-hop RTT = 2 us









Per-hop RTT = 20 us



QCN-SONAR



Hot queue length, classic ECM, $P_{\text{sample_base}} = 1\%$, w/ (0,0), vary RTT



Mandatory and Optional Behavioral Spec of QCN+

CP

1. MUST provide at least q and q' as distinct load sensor vars
 1. extensions w/ new sensors (rate etc.) MAY be possible
2. MUST implement at least fixed sampling freq.
3. MUST inject ECN w/ F_b based on marking conditions
4. MUST respond accordingly to probes [per type]

RP

1. MUST react to $-ECN$ by reducing rate [possibly directly to 0]
2. MAY react to $+ECN$ by increasing rate
 1. Upon a safety timer TO, RP MUST be able to obviously increase rate in Open Loop (without positive feedback)
3. MAY send directed and/or path probes
4. MUST respond accordingly to probes [per type]
 1. [DST is last chance to return a probe if no CP reacted before]