



the economics of network control

How full is full?

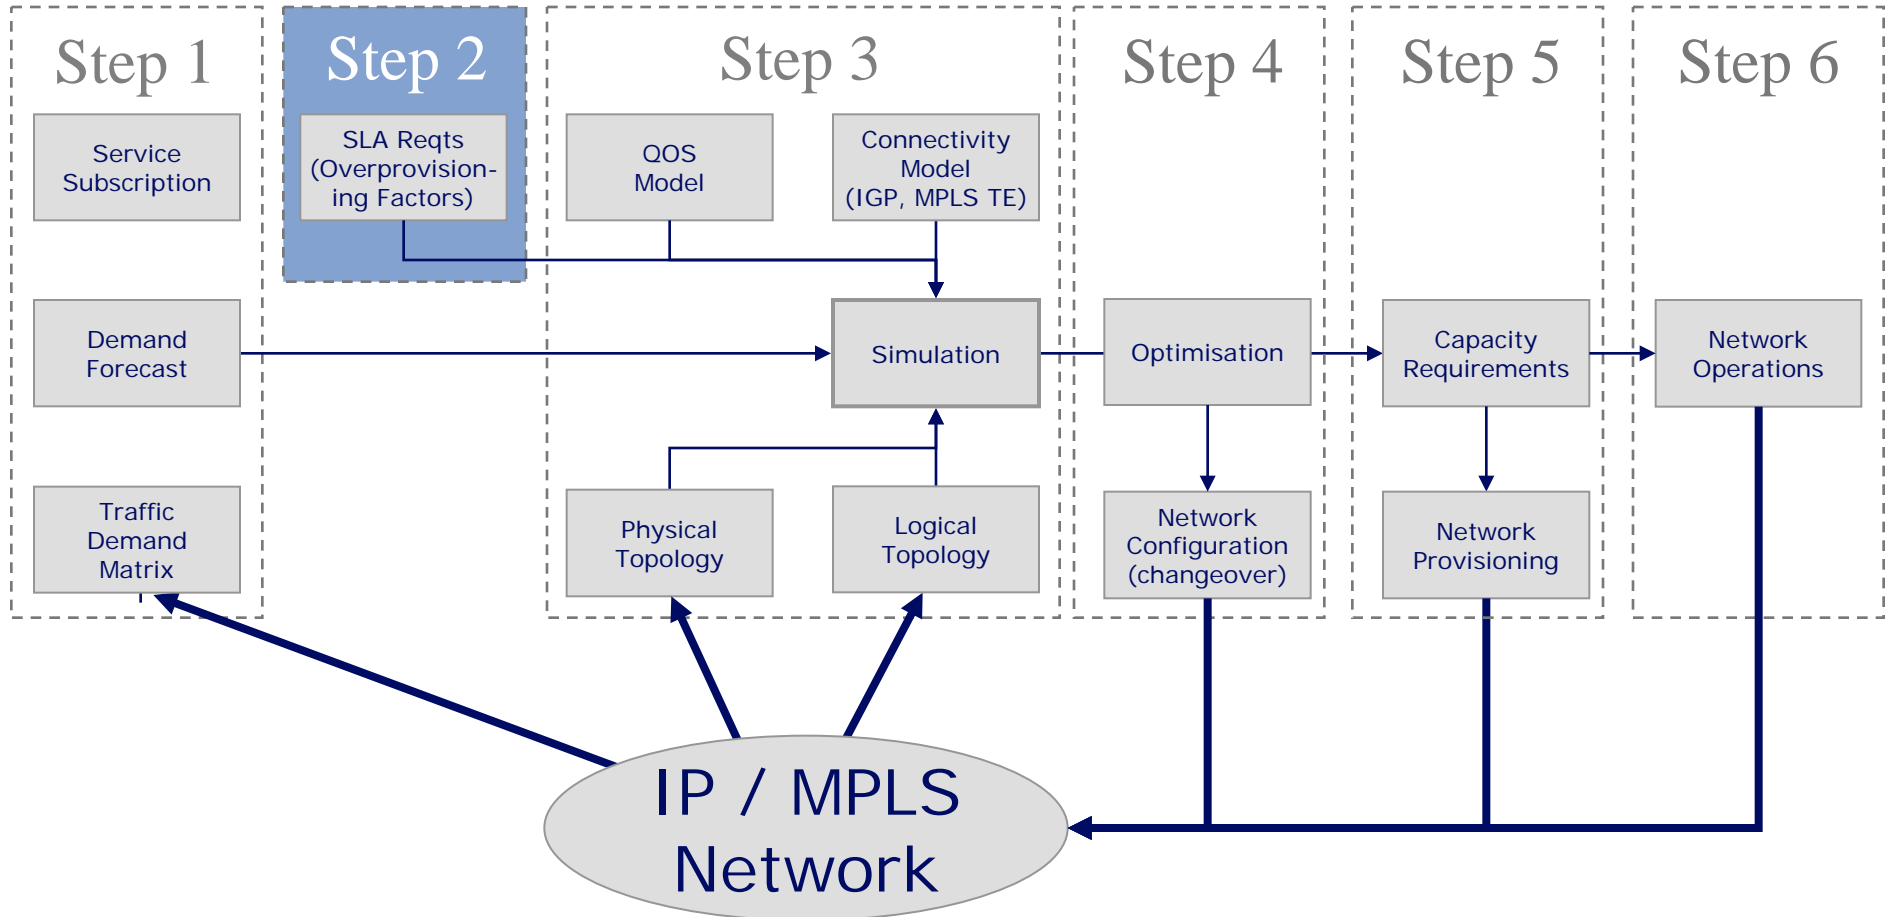
Engineer your network to meet SLA's

Thomas Telkamp

DENOG, Nov. 5th 2009
Frankfurt/Main, Germany

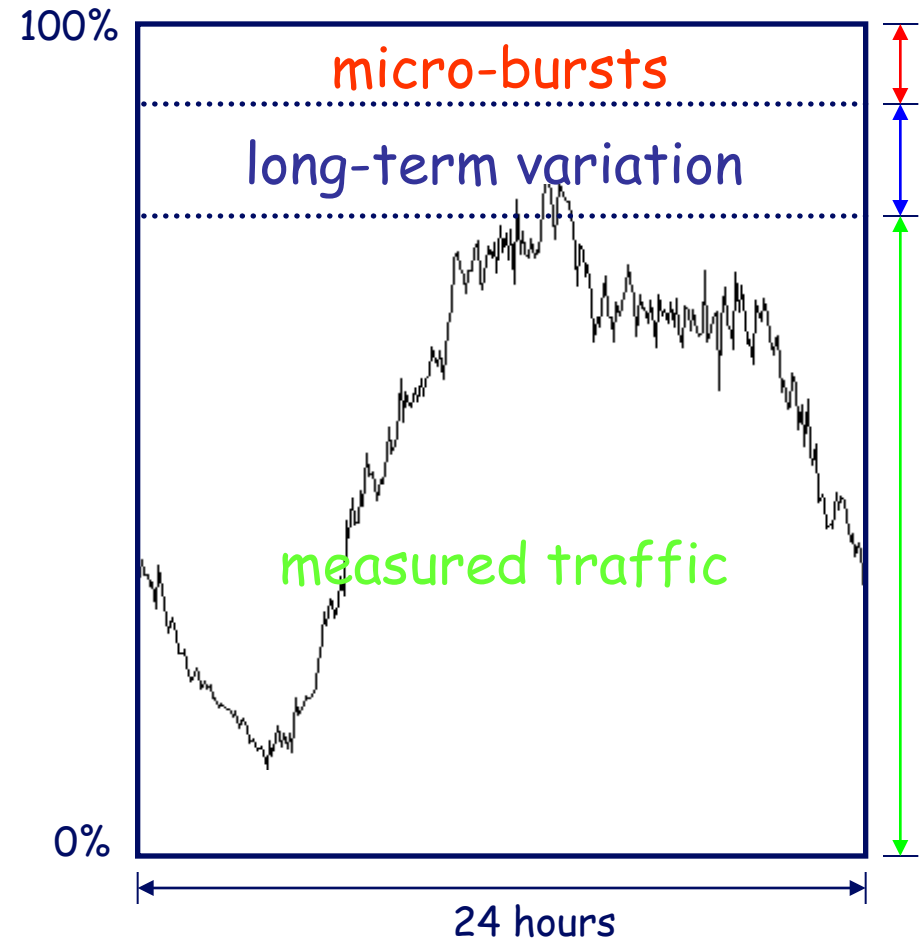
Network Planning Methodology

2. The relationship between SLAs and network planning targets ...




Demand Characterization

- Long-Term
 - Measured Traffic
 - E.g. P95 (day/week)
 - “unforeseen” events and growth
- Short-Term
 - Critical scale for queuing
 - Determine over-provisioning factor that will prevent queue buildup against micro-bursts





Relevant Timescales

- Long-Term: > 5 minutes
- Short-Term: < 5 minutes



Fiber Tap (Gigabit Ethernet)






Raw Results

30 sec of data, 1ms scale

- Mean = 950 Mbps
- Max. = 2033 Mbps
- Min. = 509 Mbps


- 95-percentile: 1183 Mbps
- 5-percentile: 737 Mbps

- (around 250 packets per 1ms interval)




Traffic Distribution Histogram (1ms scale)

- Fits normal probability distribution very well (Std. dev. = 138 Mbps)
- No Heavy-Tails
- Suggests small overprovisioning factor



Autocorrelation Lag Plot (1ms scale)

- Scatterplot for consecutive samples
- Are periods of high usage followed by other periods of high usage?
- Autocorrelation at 1ms is 0.13 (=uncorrelated)






Utilization

30 sec of data, 1 sec scale

- Mean = 950 Mbps
- Max. = 978 Mbps
- Min. = 927 Mbps

- Clearly longer derivations from the mean
- High autocorrelation at 1 sec. (0.68)



Relevant Timescales

Behavior

White Noise

LRD

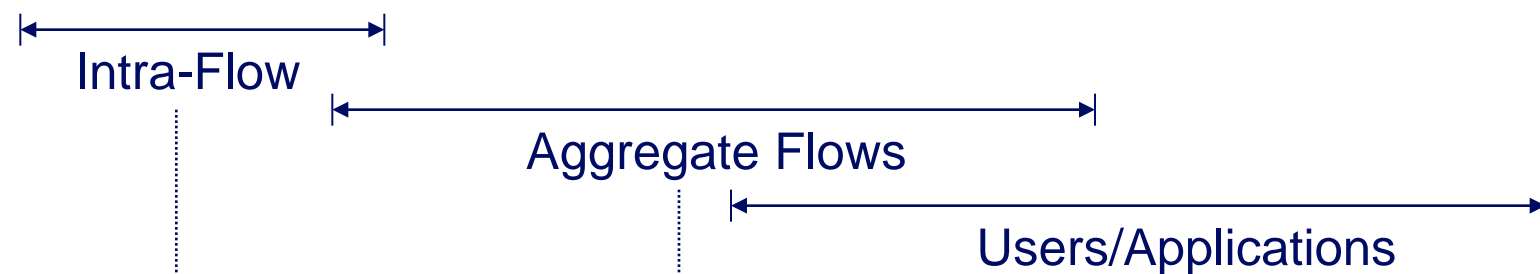
Smooth



Timescale



Dynamics



Characteristics

TCP (<RTT)

Flow Sizes/Durations

Diurnal variation

Traffic Behavior

- < 150ms: Gaussian White Noise
 - I.I.D, Normal distribution, no correlation
 - Short-Range Dependent ($H=0.5$)
- 150ms – 30sec: Long-Range Dependence
 - Non-summable autocorrelation
 - Statistically Self-Similar
 - $X =_d m^{1-H} X^{(m)}$, "Clustering", Fractional Brownian Motion (fBM)
- > 30 sec: Smooth
 - Degenerate case: autocorrelation = 1
 - Almost constant, diurnal variation


- M/M/1 queuing formula
- Markovian
 - Poisson-process
 - Infinite number of sources
- “Circuits can be operated at over 99% utilization, with delay and jitter well below 1ms” [2] [3]
- Self-Similarity
- Traffic is bursty at many or all timescales
- “Scale-invariant burstiness (i.e. self-similarity) introduces new complexities into optimization of network performance and makes the task of providing QoS together with achieving high utilization difficult” [4]



Telephony Traffic

(inter-city on 6/3/2002)

Centi-Erlang



Voice Capacity Allocation

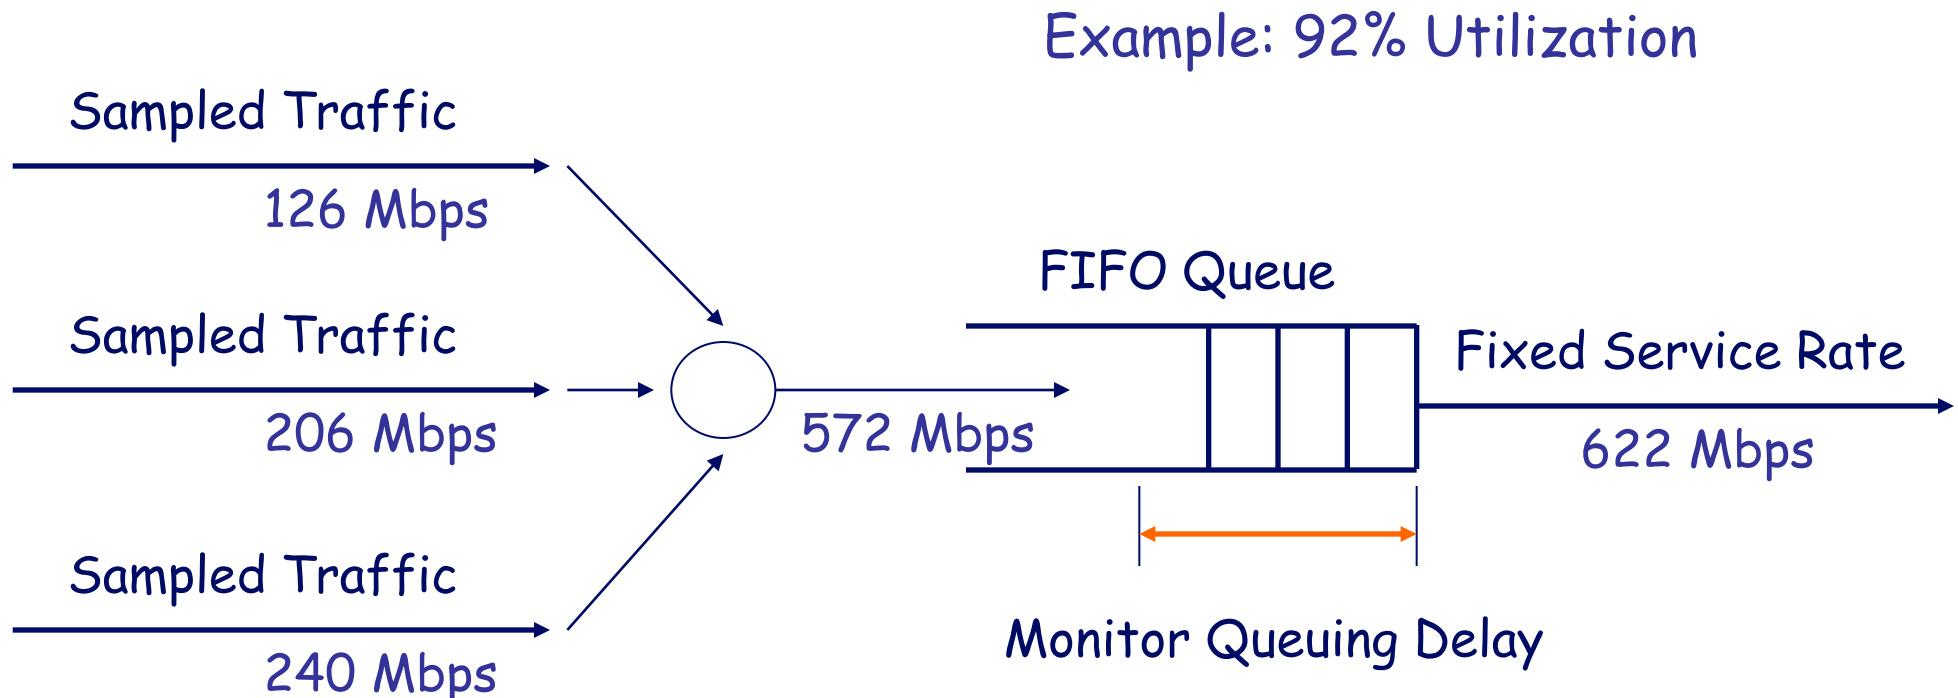
- Erlang traffic model(s)
- 1 Erlang = 1 hour of calls
 - Average numbers of calls in an hour
- Busy Hour Traffic: about 330 Erlang
- Erlang B formula (for 330 Erlang):
 - Blocking 1% -> 354 lines required
 - Blocking 0.1% -> 376 lines required
- Reverse calculation:
 - 100 lines available, blocking 0.1%:
 - 75 Erlang (=average of 75 calls in an hour)

IP Capacity Allocation

- Measurement data
 - E.g. 5-min average utilization
- Performance objectives
 - E.g. packet loss = 0%, jitter < 10ms
 - End-to-end: convert to per-hop objective
- But we don't have an "Erlang formula"...
- Two paths towards a solution:
 - 1) Model the traffic, and fit parameters
 - 2) Empirically derive guidelines


Queuing Simulation

- Feed multiplexed sampled traffic data into FIFO queue
- Measure amount of traffic that violates the delay bound






Queuing Simulation: Results



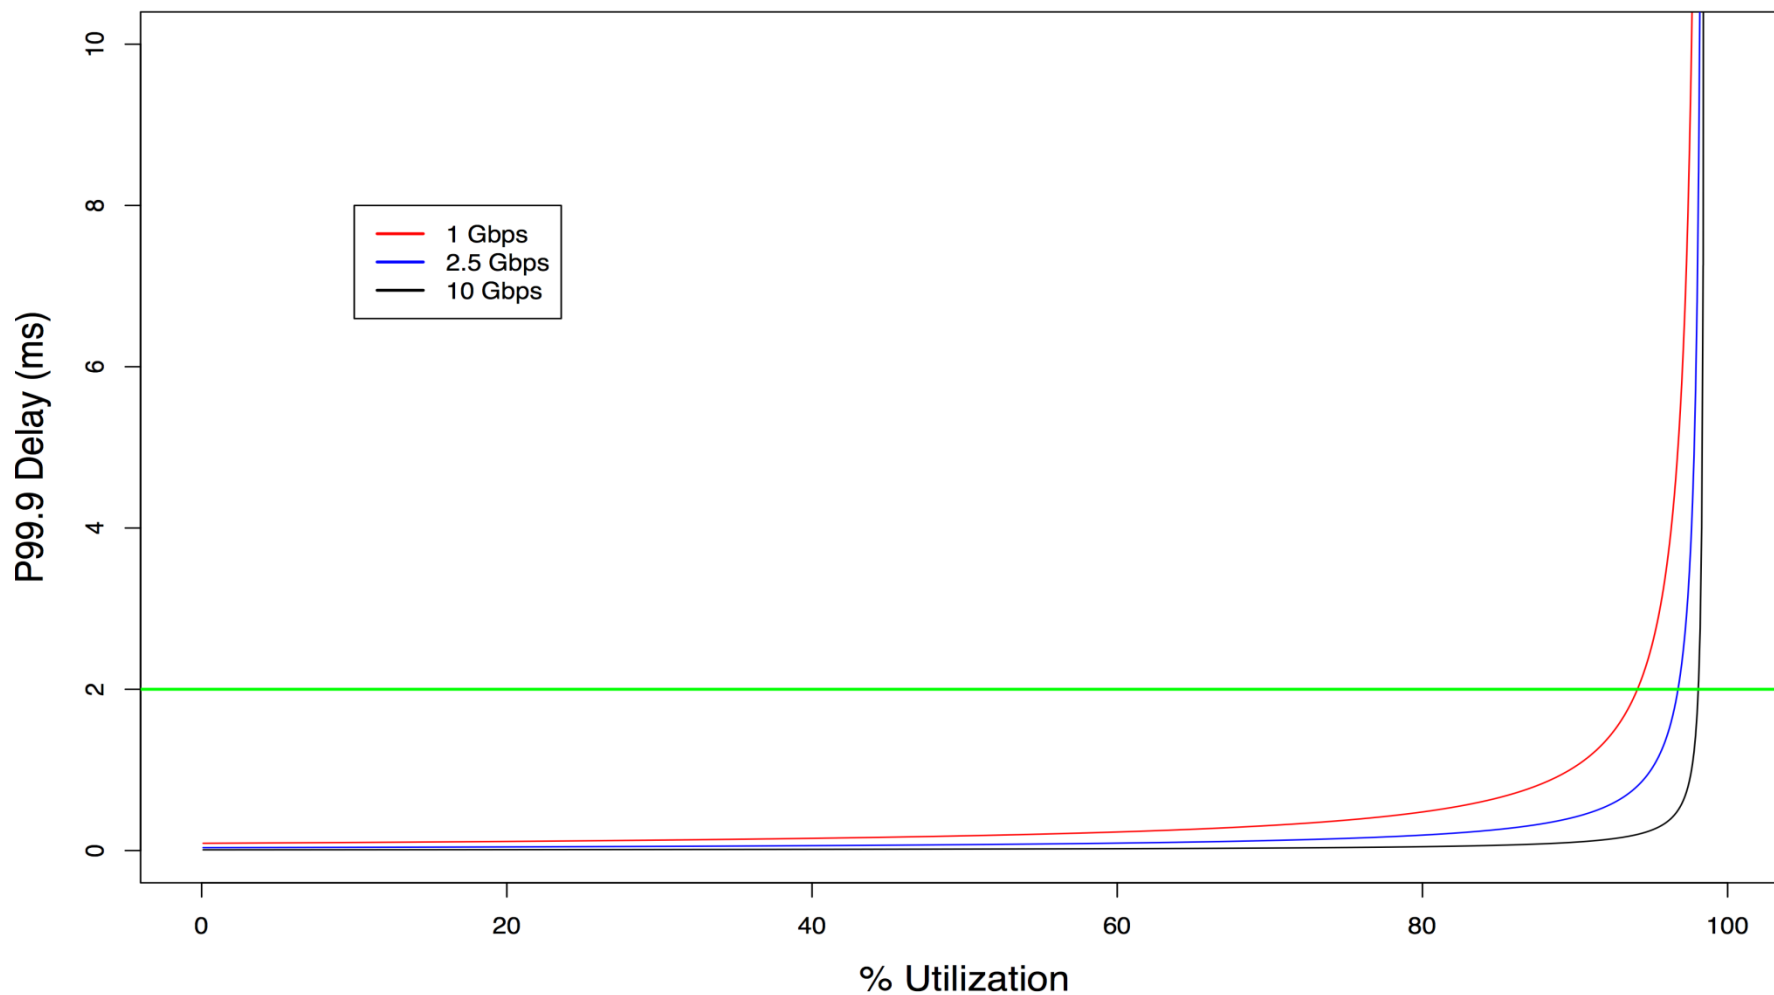


Scaled Delay Function (1G)






P99.9 Delay Function





P99.9 Delay Function



Cao, Cleveland, Sun


- Bandwidth Estimation for Best-Effort Internet Traffic
 - Jin Cao, William S. Cleveland, and Don X. Sun
 - Statist. Sci. Volume 19, Number 3 (2004), 518-543.
- Data:
 - BELL, AIX, MFN, NZIX
- Model depends on:
 - #connections, traffic rate, delay target
 - Avg. bandwidth per connection: 16384 bps
- Best-Effort Delay Formula:

$$\text{logit}_2(u) = o + (o_c + o_{\tau\delta}) \log_2(c) + o_{\tau\delta} \log_2(\gamma_b\delta) + o_\omega(-\log_2(-\log_2(\omega))).$$



Best-Effort Delay Model

P99.9, 2ms

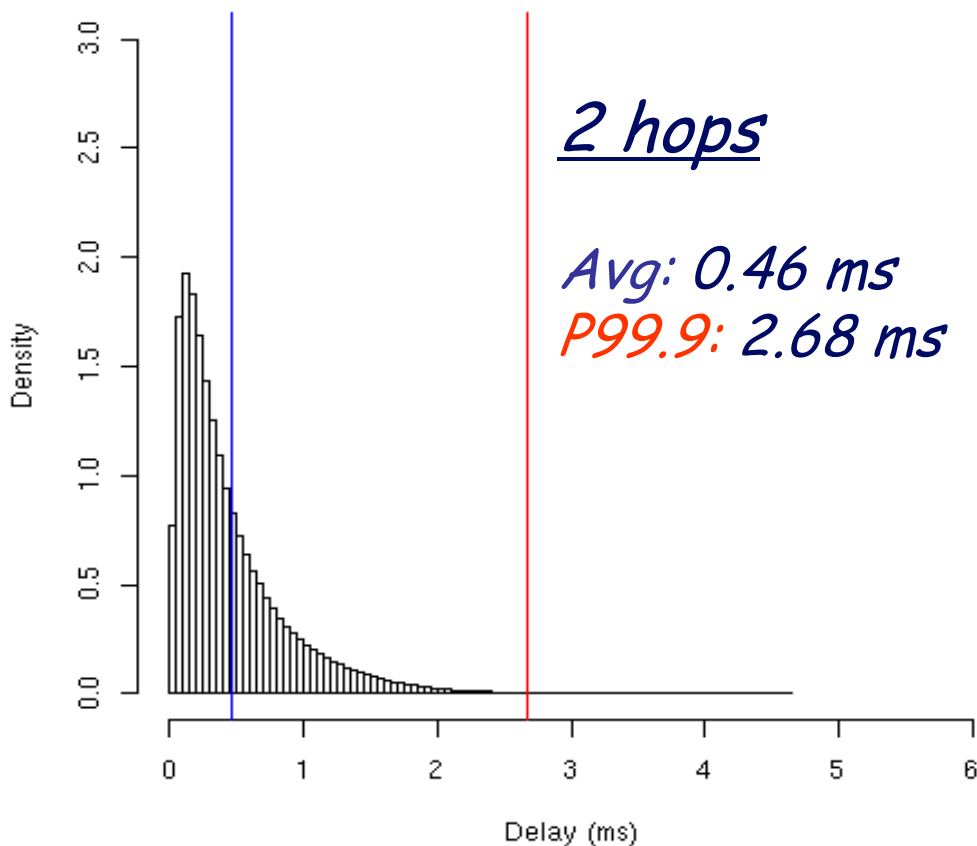



Queuing Simulation: Numeric Results

- 1 Gbps (Gigabit Ethernet)
 - 1 ms delay bound for 999 out of 1000 packets (99.9-percentile):
 - 90% maximum utilization
- 622 Mbps (STM-4c/OC-12c)
 - 1 ms delay bound for 999 out of 1000 packets (99.9-percentile):
 - 85% maximum utilization

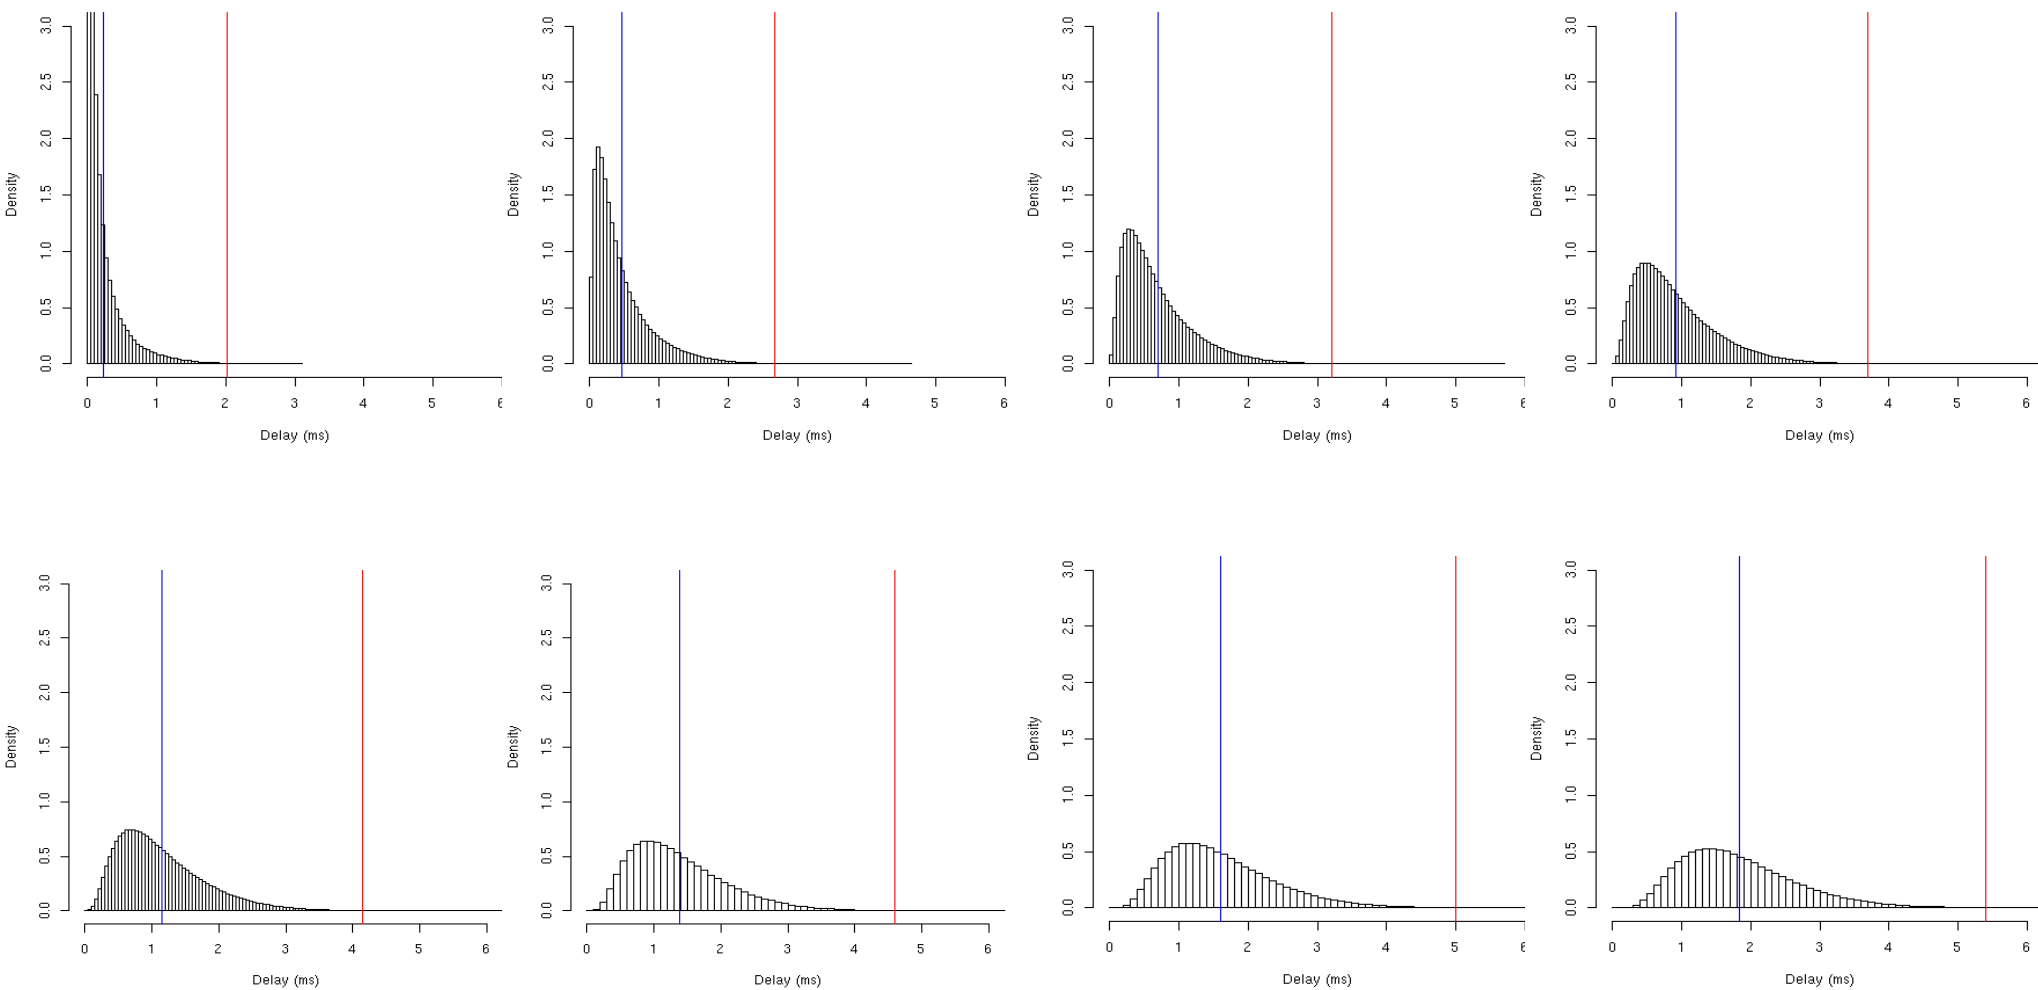


Multi-hop Queueing





Multi-hop Queueing (1-8 hops)

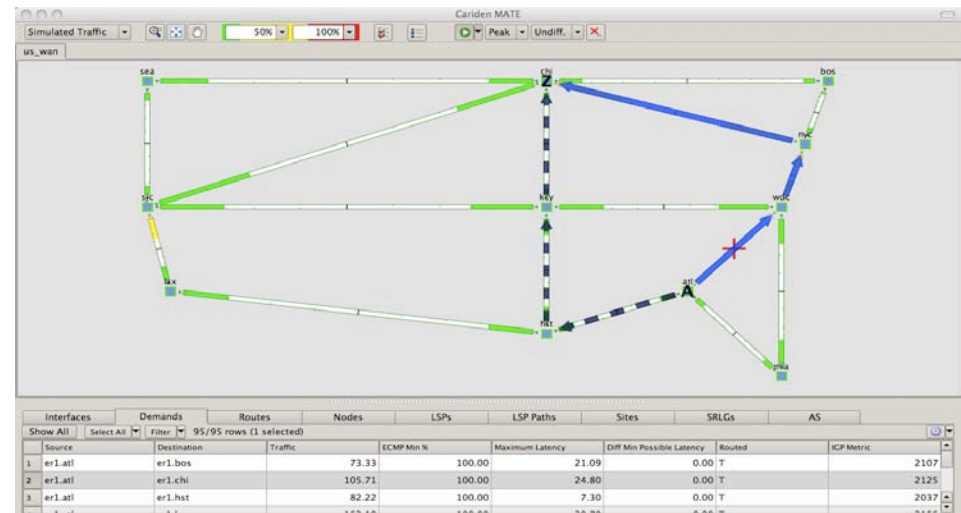


Queueing: Summary

- Queueing Simulation:
 - Gigabit Ethernet (backbone) link
 - overprovisioning percentage in the order of 10% is required to bound delay/jitter to less than 1 ms
 - Lower speeds (<1G)
 - overprovisioning factor is significant,
 - Higher speeds (2.5G/10G)
 - overprovisioning factor becomes very small
- P99.9 multi-hop delay/jitter is not additive

Future work

- Use more conservative delay formula
- Estimate avg. bandwidth per connection (NetFlow?)
- Investigate VoIP traffic (ongoing)
- Integrate numbers in end-to-end planning and engineering process
 - Include DiffServ
 - E.g.using Cariden MATE





the economics of network control

- Web: <http://www.cariden.com>
- Phone: +1 650 564 9200
- Fax: +1 650 564 9500
- Address: 888 Villa Street, Suite 500
Mountain View, CA 94041
USA