ISMA Oct '02
October 7-11, 2002
Lorentz Center, Leiden, The Netherlands

Traffic Characteristics and Network Planning

Global Crossing®

Thomas Telkamp
Director Network Architecture
Global Crossing Telecommunications, Inc.
telkamp@gblx.net

What to expect?

- A methodology to analyze your traffic, and apply the results to the planning process
- Practical approach
- An example from Global Crossing's network
- *BUT*, your network might be different in:
  - Scale
  - SLA's
  - Applications
  - Etc…

QoS in Backbone Networks

- Requirements are:
  - low delay
  - low jitter
  - low packet loss
- Common practice in backbone networks is overprovisioning:
  - Enough capacity in the network to meet demands
  - In peak times, and under failure conditions
- Prevent significant queue buildup

QoS in Backbone Networks

- The overprovisioning approach is effective
  - See Packet Design presentation at NANOG 22 [1]
- But capital is limited today…
- Can we do better than the rules-of-thumb:
  - "upgrade at 40% or 50% utilization"
  - "maximum 75% utilization under failure"
- Is aggregated traffic well-behaved enough to do "tight" capacity planning?

Related work: Opposite views (!)

- **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]

Opposite views

- **M/M/1 queuing formula**
- **Self-Similarity**

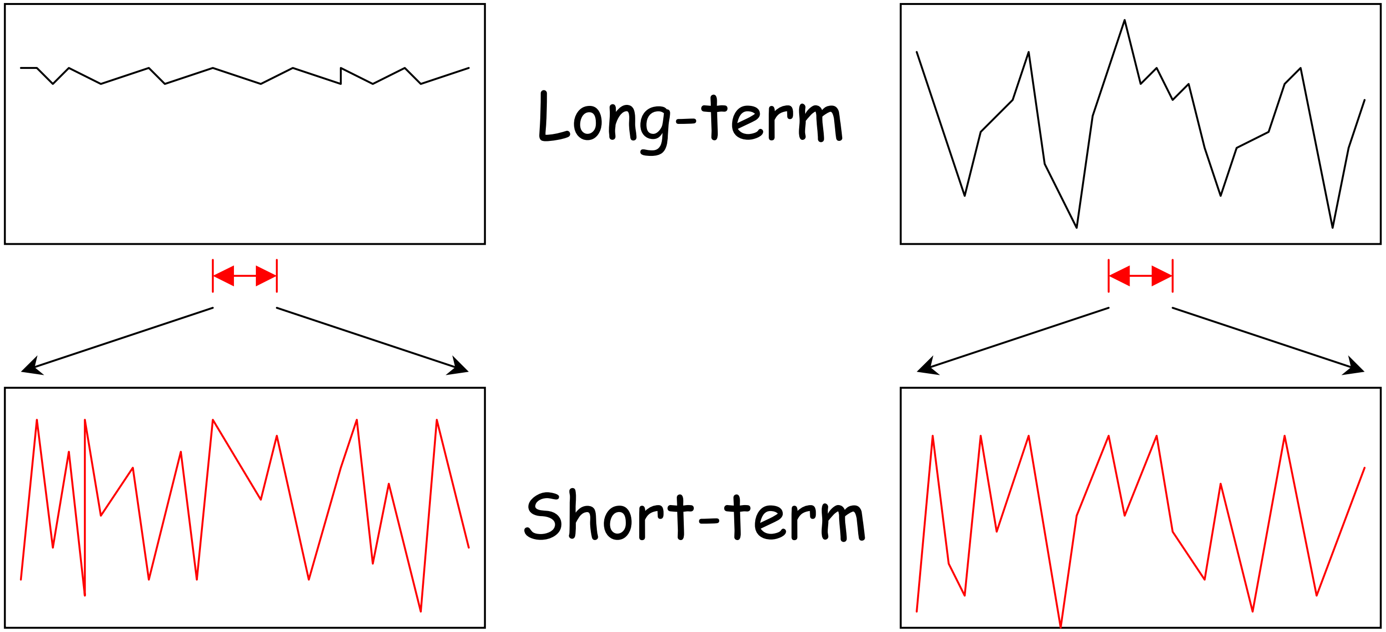

Long-term

Short-term

Network Planning Framework

- <u>Demand Characterization</u>
  - Long-term: days/weeks timeframe
  - Short-term: dynamics at sub-5-min timescale
- <u>Failure Analysis</u>
  - Determine failure scenarios and SRLG's
- <u>Simulation and Optimization</u>
  - Determine minimum capacity deployment to meet objectives under normal and failure conditions

Demand Characterization

- <u>Long-term</u>
  - Robust estimation of 5-minute peak values
    - E.g. 95-percentile over day or week
  - Estimate "unforeseen" events
  - Calculate growth rate
- <u>Short-term</u>
  - Critical scale for queuing (1ms)
  - Determine overprovisioning factor that will prevent queue buildup against micro-bursts

Demand Characterization

- Measured Traffic
  - P95 (day/week)
- Long-term variation
  - P95 to peak ratio
  - "unforeseen" events
- Micro-bursts
  - Short timescale traffic dynamics

- But let's first take a look at the telephony world…

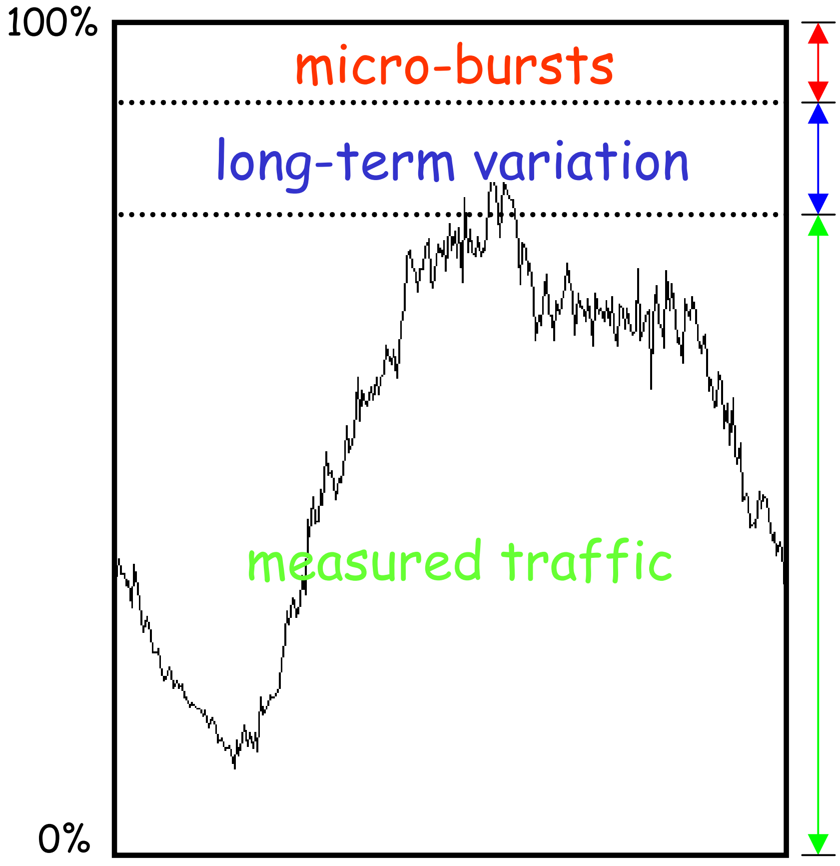

Telephony Traffic

(inter-city on 6/3/2002)

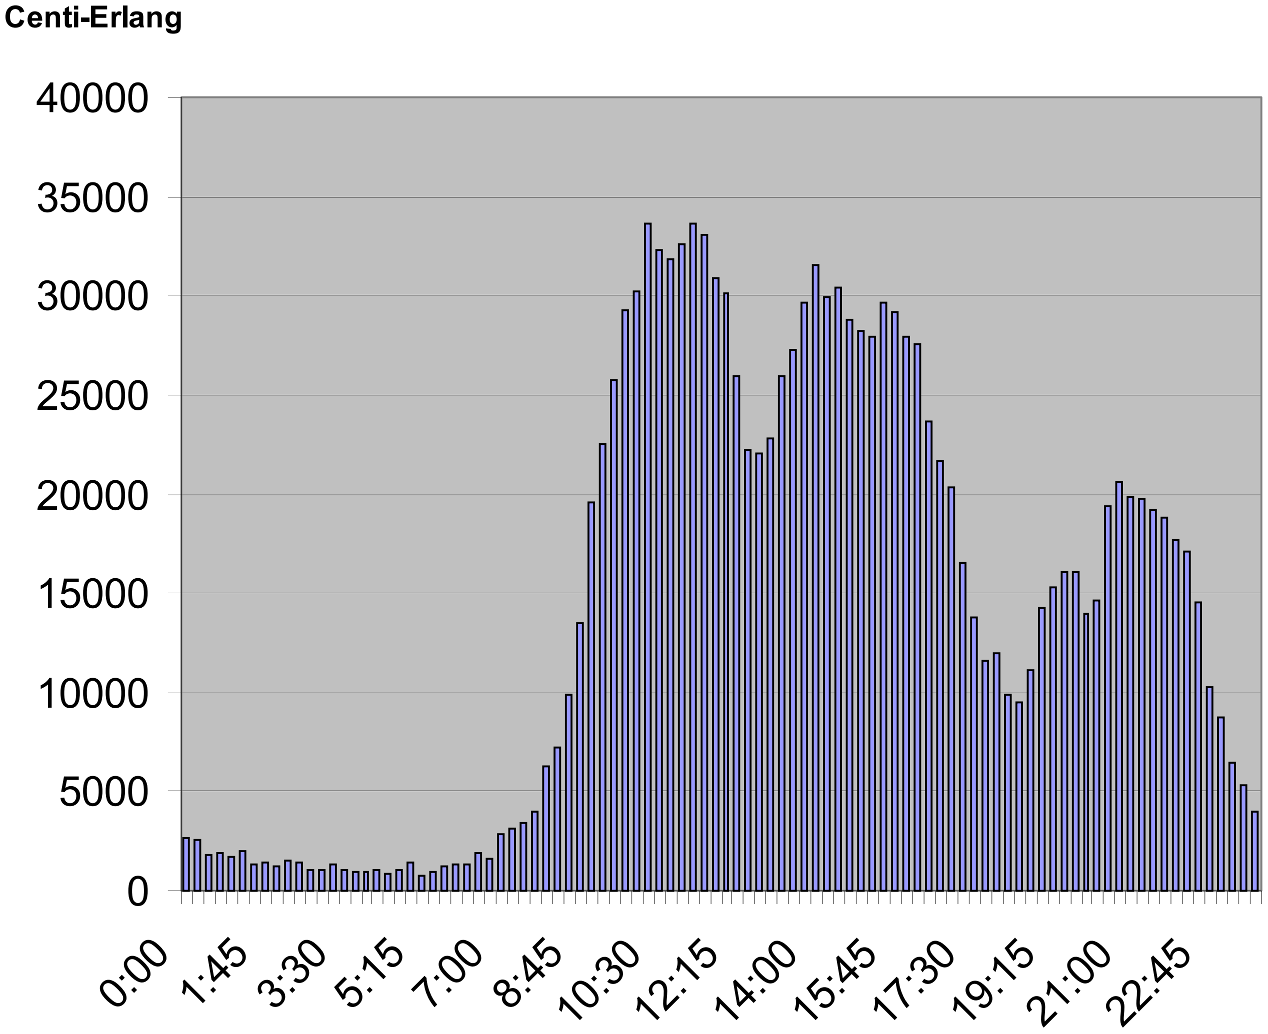

Voice Capacity Planning (Some) Assumptions

- Erlang B:
  - Call arrivals are random (Poisson)
  - Blocked calls are cleared
- Extended Erlang B:
  - Includes a retry percentage
- Erlang C
  - Blocked calls are queued (*"your call is very important to us, blah, blah…"*)

Voice Capacity Planning Example

- 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
- "Overprovisioning" for 1% blocking: 7.3%
- "Overprovisioning" for 0.1% blocking: 13.9%

IP Capacity Planning

- Measurement data
  - E.g. 5-min average utilization
- Performance objectives
  - E.g. packet loss < 0.1%, jitter < 20ms
  - 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) Experimentally derive guidelines

Long-term Traffic Characterization

- Investigate burstiness in 5-min measurements over days/weeks
- Bursty traffic: peaks are very large compared to average
  - I.e. the distribution is Heavy-Tailed
  - Mean and 95-percentile do not represent the traffic very well
  - Planning becomes very difficult
- Collect (SNMP) and analyze network data
  - Traffic Matrix via NetFlow or MPLS mesh

High- vs Low-Bandwidth Demands

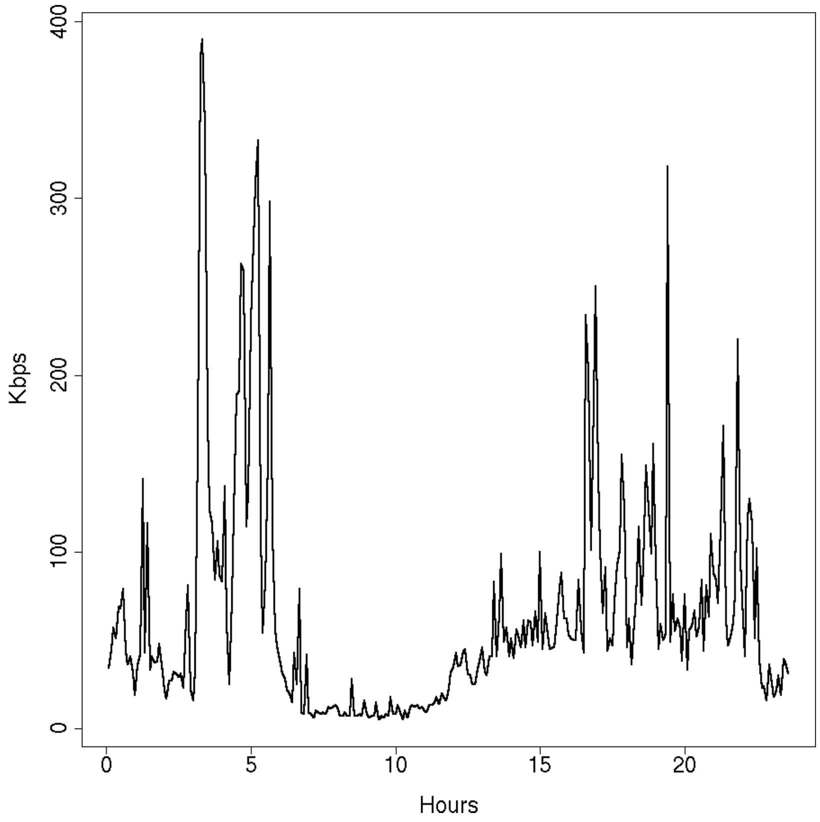

Cleveland -> Denver
Mean=64Kbps, Max=380Kbps
P95=201Kbps, alpha=1.8 (tail index)

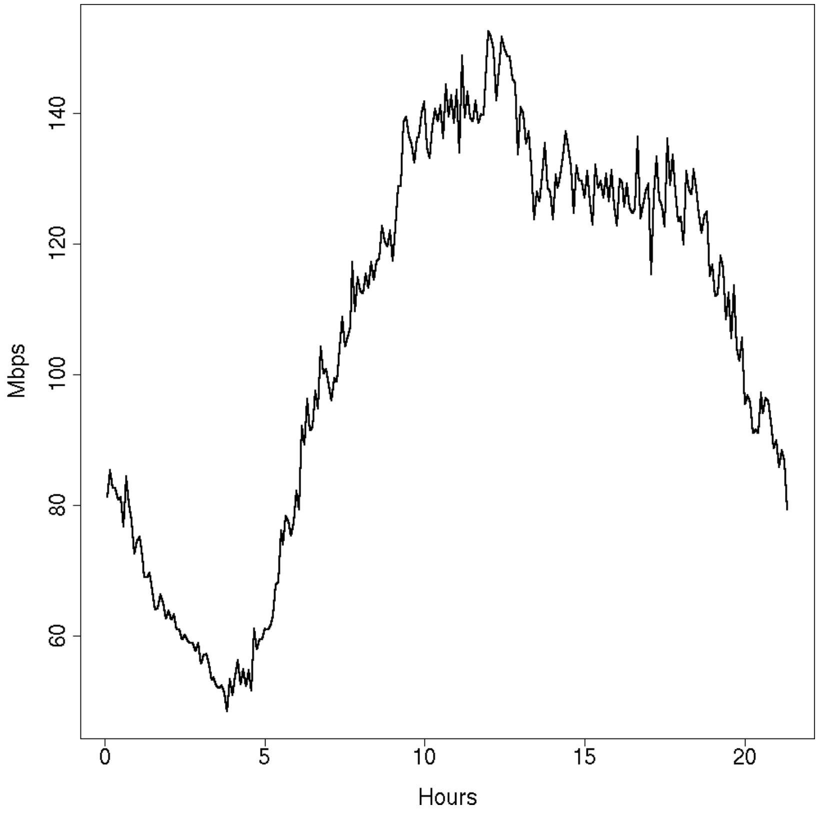

Washington D.C. -> Copenhagen
Mean=106Mbps, Max=152Mbps
P95=144Mbps, alpha=21 (tail index)

Variance vs Bandwidth

- Around 8200 demands between core routers
- Relative variance decreases with increasing bandwidth [5]
- Vertical red line is 0.5 Mbps
- High-bandwidth demands seem well-behaved
- 98% of traffic is carried by the demands larger than 0.5 Mbps

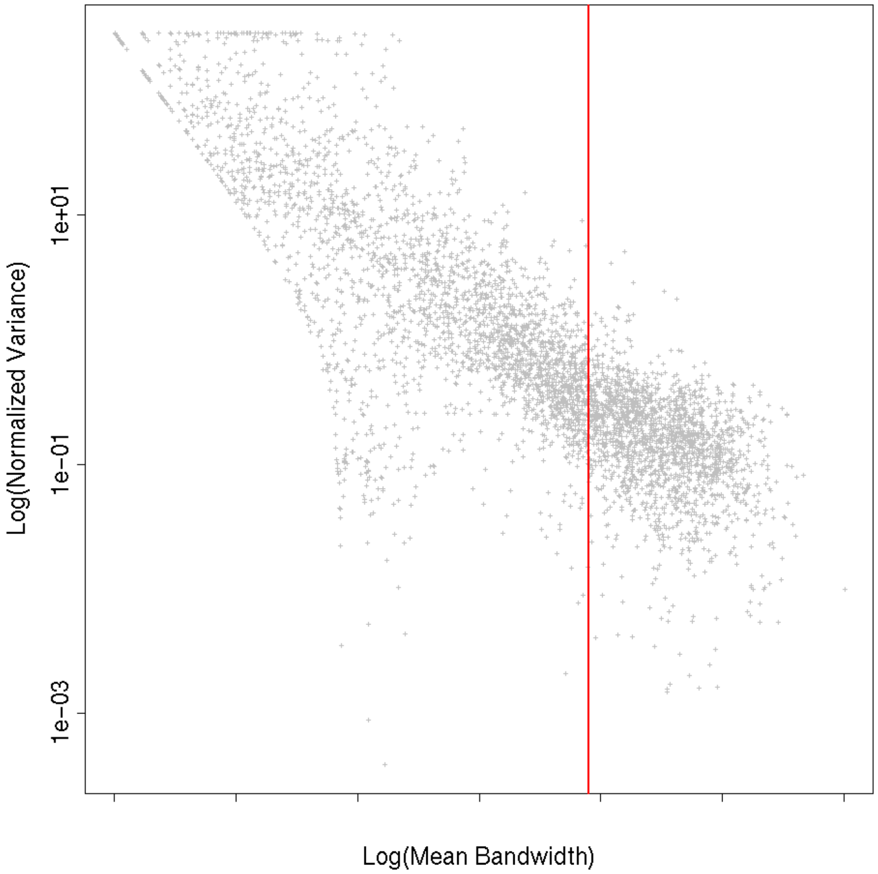

Short-term Traffic Characterization

- Investigate burstiness within 5-min intervals
- Measurements at critical timescale for queuing, like 1ms or 10ms
- Only at specific locations
  - Complex setup
  - A lot of data
- Analyze statistical properties

Fiber Tap (Gigabit Ethernet)

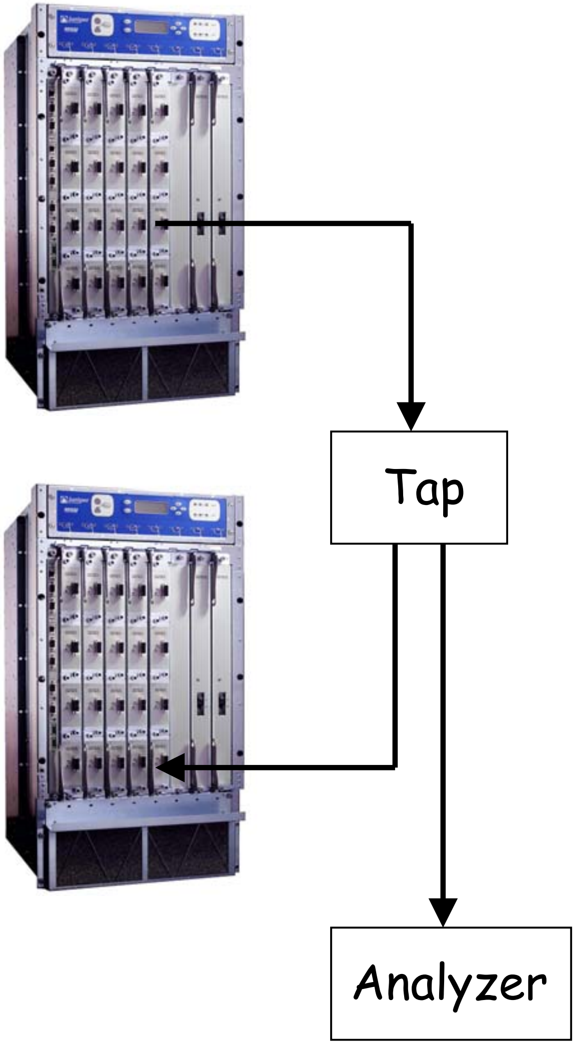

Tap

Analyzer

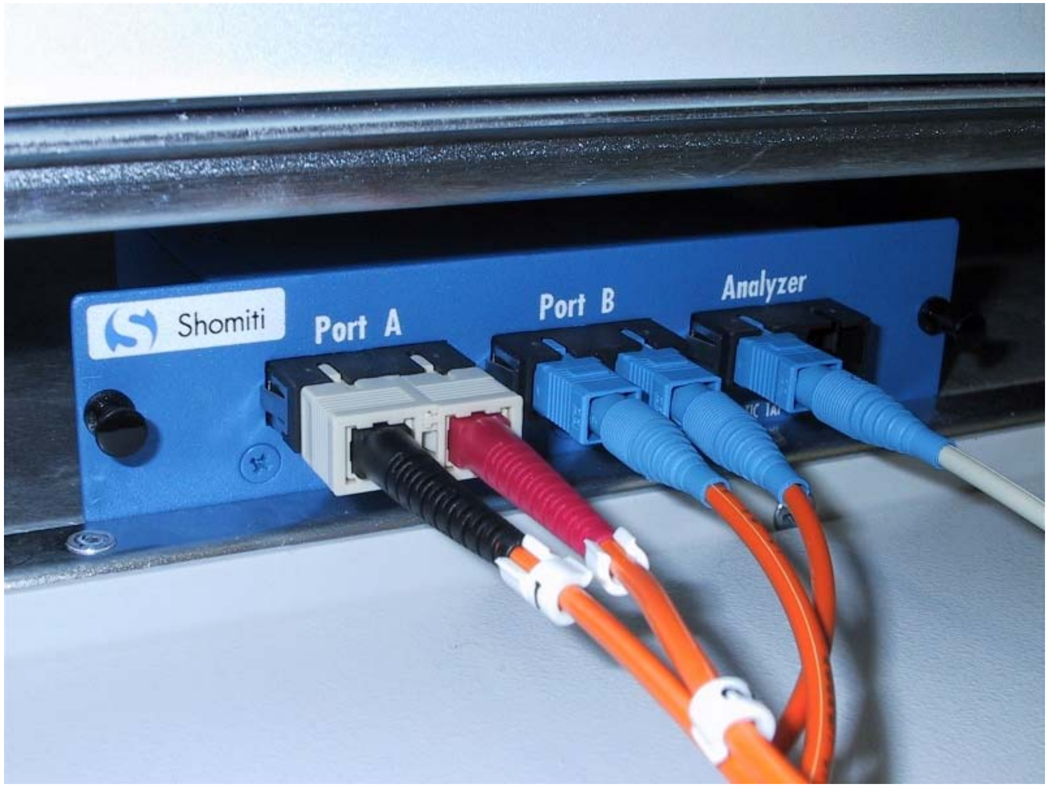

Raw Results
10 min. of data, 10ms scale

- Mean = 225 Mbps
- Max. = 342 Mbps
- Min. = 128 Mbps

- 95-percentile: 266 Mbps
- 5-percentile: 187 Mbps

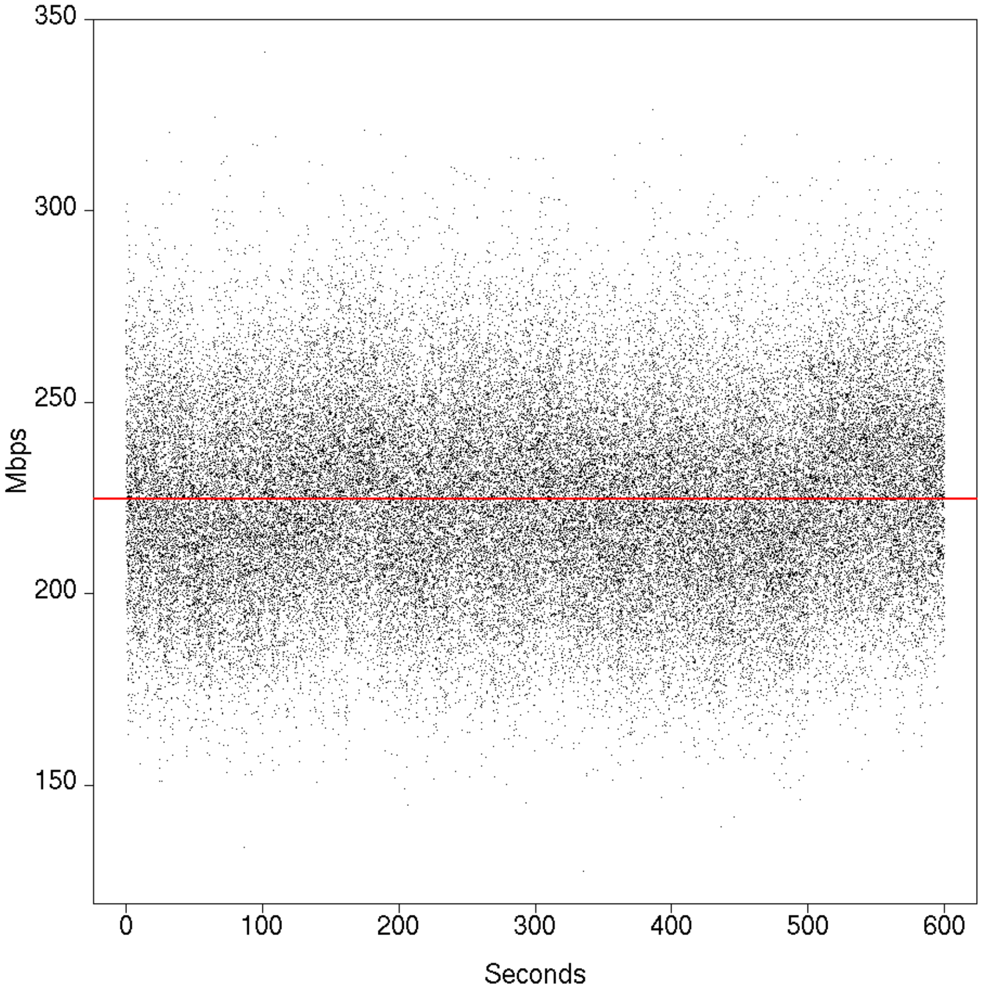

Traffic Distribution Histogram (10ms scale)

- Fits normal probability distribution very well (Std. dev. = 24 Mbps)
- No Heavy-Tails
- Suggests small overprovisioning factor

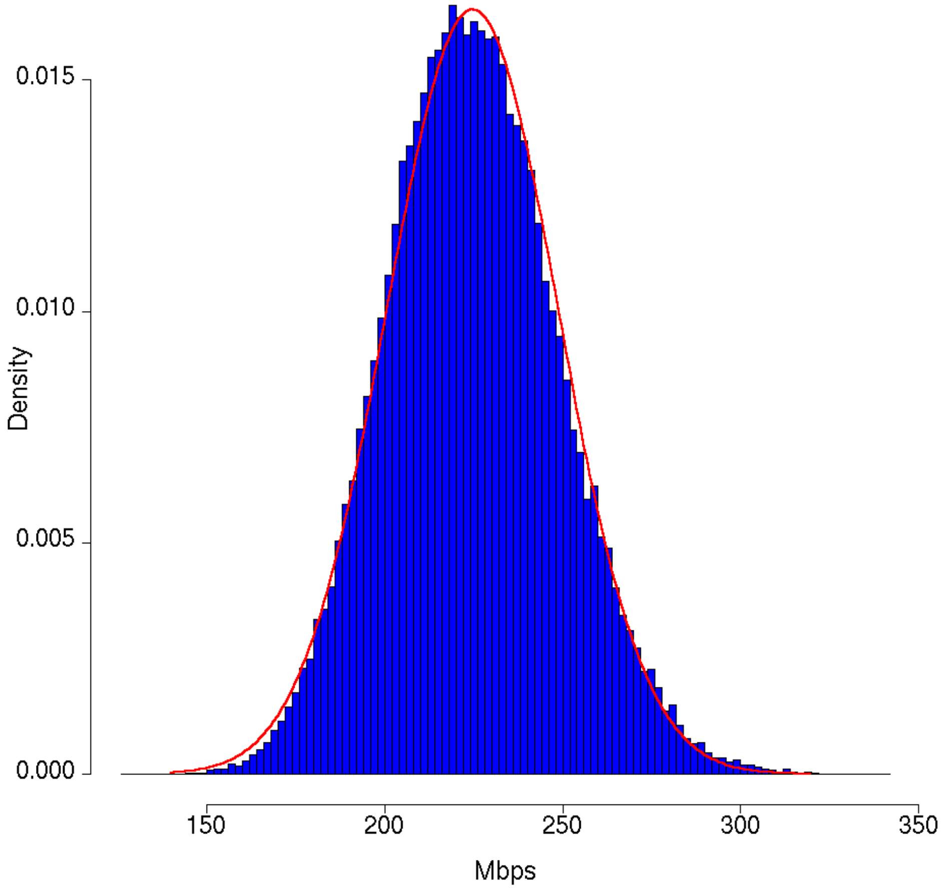

Autocorrelation
Lag Plot (10ms scale)

- Scatterplot for consecutive samples
- Are periods of high usage followed by other periods of high usage?

- Autocorrelation at 10ms is 0.16 (=uncorrelated)

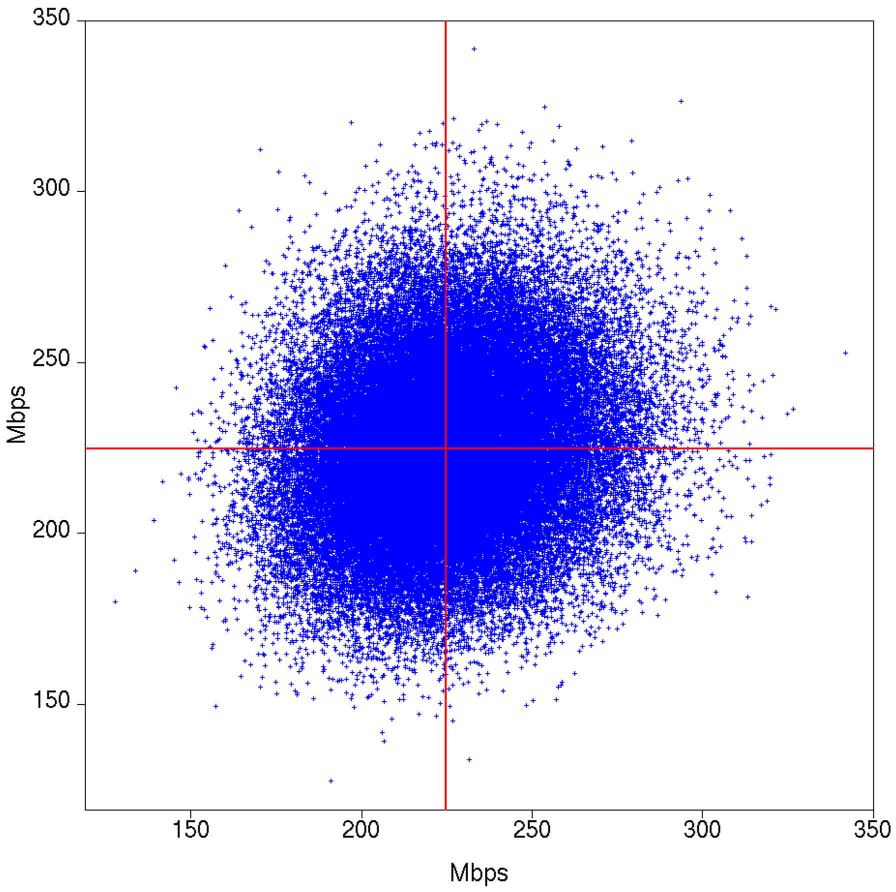

Utilization
10 min. of data, 10 sec. scale

- Mean = 225 Mbps
- Max. = 233 Mbps
- Min. = 214 Mbps

- Clearly longer derivations from the mean
- High autocorrelation at 10 sec. (0.65)

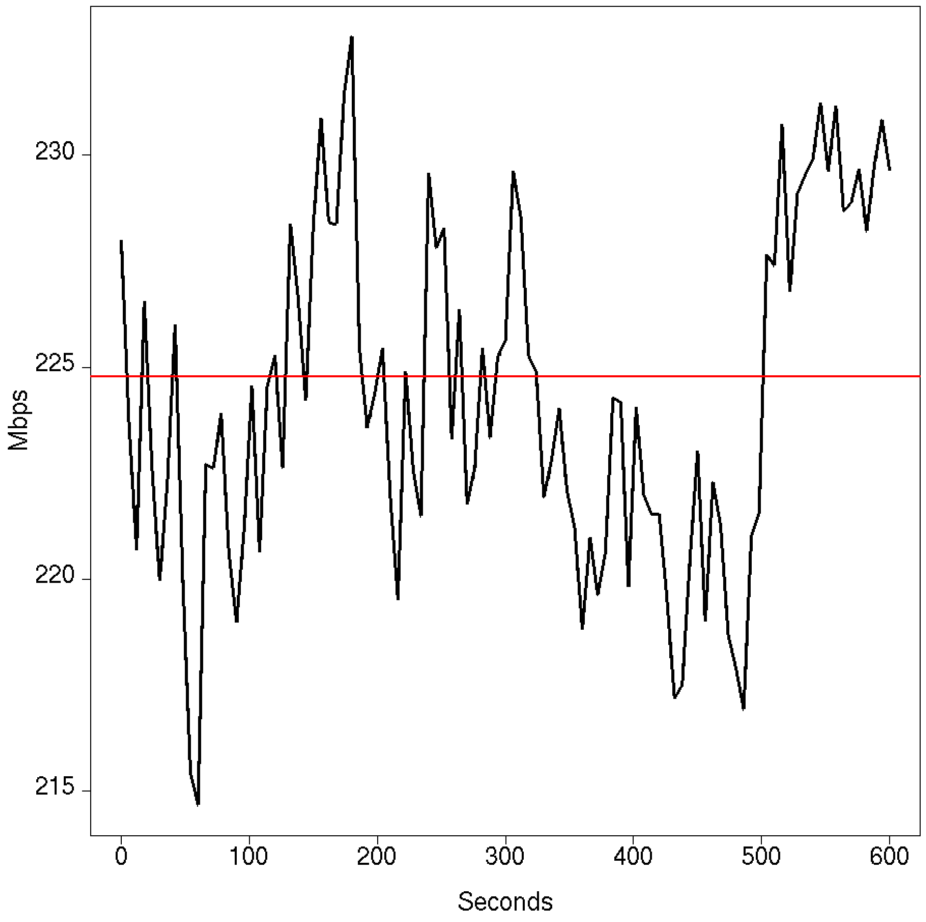

Queuing Simulation

- Feed sampled traffic data into FIFO queue (1ms)
- Fix *Service Rate* and max. *Queuing Delay*
- Measure amount of traffic that violates the delay bound
- Repeat for different Service Rates and Queuing Delays

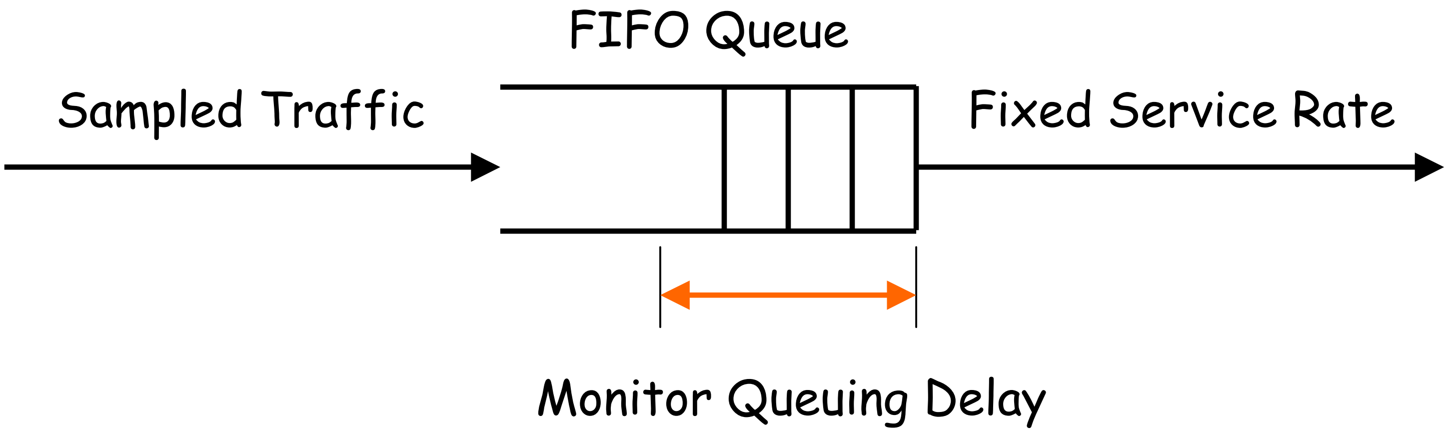

Bandwidth Requirement vs Delay Bound

- How much Bandwidth is needed to meet the Delay Bound for a certain percentage of the traffic?

- Mean BW
- 99%
- 99.9%
- 99.99%
- 99.9999%

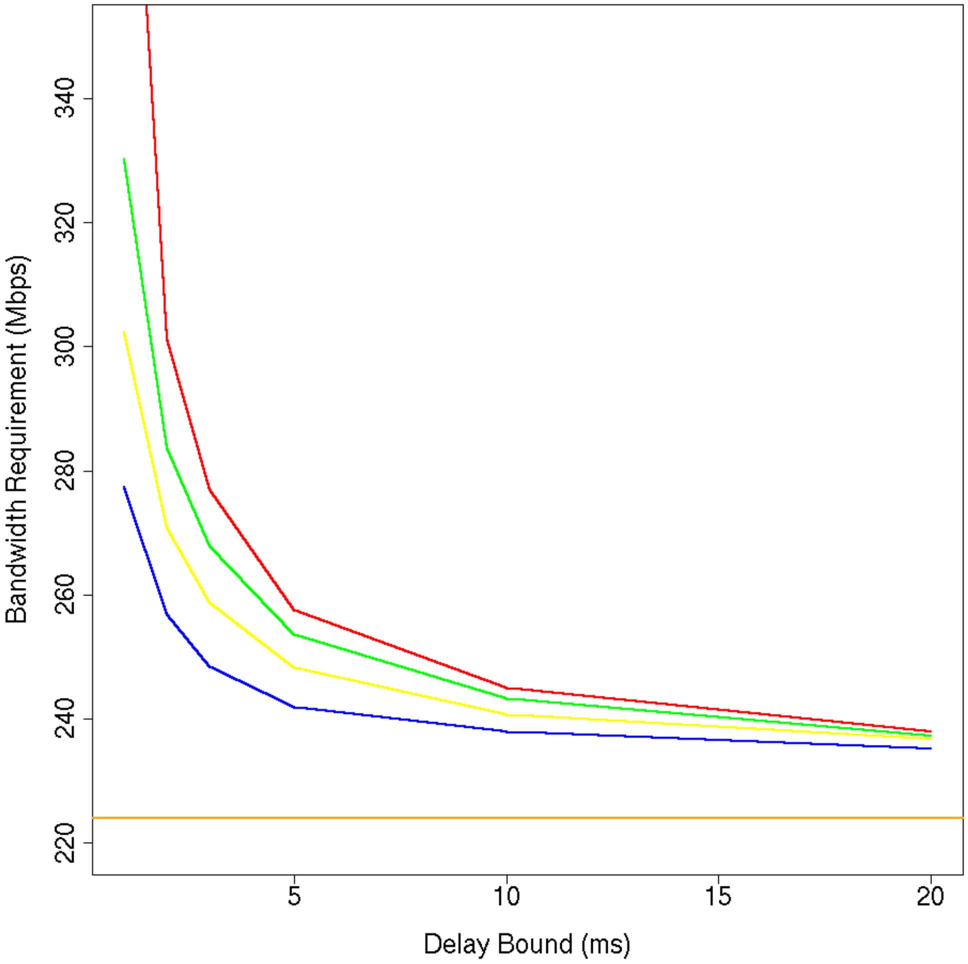

Bandwidth Requirements Numeric Results

- Example 1
  - 5ms delay bound
  - 99.9999% of the traffic (10^{-6})
  - BW required: 257 Mbps
  - “Overprovisioning”: 14%
- Example 2
  - 10ms delay bound
  - 99.9% of the traffic (10^{-3})
  - BW required: 241 Mbps
  - “Overprovisioning”: 7%

Bandwidth Requirements Numeric Results (draft)

- Synthesized data: 704Mbps
  - 5ms delay bound
  - 99.9999% of the traffic (10^{-6})
  - BW required: 755 Mbps
  - "Overprovisioning": 7.2%
- Synthesized data: 1228Mbps
  - 5ms delay bound
  - 99.9999% of the traffic (10^{-6})
  - BW required: 1271 Mbps
  - "Overprovisioning": 3.5%

Back to the Framework

- Demand Characterization
  - Long-term well-behaved traffic
  - Overprovisioning for short-term bursts can be experimentally derived
- How to use this for planning purposes?
- Failure Analysis
  - Determine failure scenarios
    - E.g. single link failures, routers, SRLG, etc…
- Input for simulation

Simulation

- Feed demands and overprovisioning factors into simulation tool
- Run simulation for normal and failure scenarios
- Optimize Capacity Deployment and Routing (IGP or MPLS based) to meet requirements

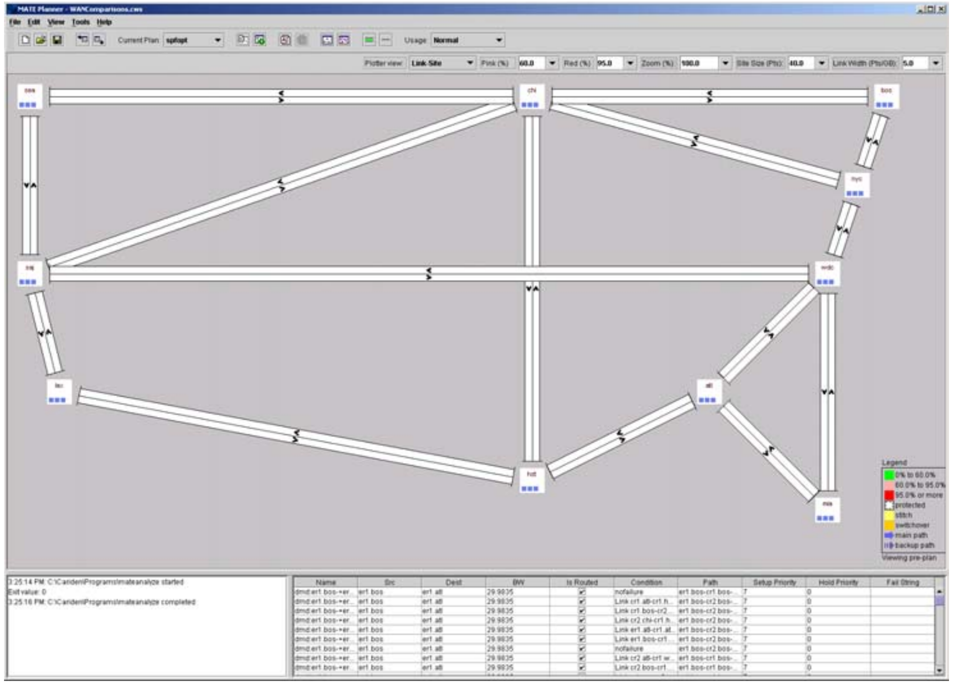

```
3:25:14 PM: C:\Cariden\Programs\mateanalyze started
Exit value: 0
3:25:16 PM: C:\Cariden\Programs\mateanalyze completed
```

Name	Src	Dest	BW	Is Routed	Condition	Path	Setup Priority	Hold Priority	Fail String
dmd:er1.bos-+er...	er1.bos	er1.atl	29.9835	☑	nofailure	er1.bos-cr1.bos-...	7	0	
dmd:er1.bos-+er...	er1.bos	er1.atl	29.9835	☑	Link cr1.atl-cr1.h...	er1.bos-cr2.bos-...	7	0	
dmd:er1.bos-+er...	er1.bos	er1.atl	29.9835	☑	Link cr1.bos-cr2....	er1.bos-cr2.bos-...	7	0	
dmd:er1.bos-+er...	er1.bos	er1.atl	29.9835	☑	Link cr2.chi-cr1.h...	er1.bos-cr2.bos-...	7	0	
dmd:er1.bos-+er...	er1.bos	er1.atl	29.9835	☑	Link er1.atl-cr1.at...	er1.bos-cr2.bos-...	7	0	
dmd:er1.bos-+er...	er1.bos	er1.atl	29.9835	☑	Link er1.bos-cr1....	er1.bos-cr2.bos-...	7	0	
dmd:er1.bos-+er...	er1.bos	er1.atl	29.9835	☑	nofailure	er1.bos-cr2.bos-...	7	0	
dmd:er1.bos-+er...	er1.bos	er1.atl	29.9835	☑	Link cr2.atl-cr1.w...	er1.bos-cr1.bos-...	7	0	
dmd:er1.bos-+er...	er1.bos	er1.atl	29.9835	☑	Link cr2.bos-cr1....	er1.bos-cr1.bos-...	7	0	

- Tools like MATE (Cariden) and NPAT (WANDL)

How does Diff-Serv fit in this picture?

- All traffic in one class (no Diff-Serv) might require large overprovisioning factor for tight objectives (e.g. low delay/jitter for VoIP)
- Prioritizing that traffic (using a SPQ) would make the overprovisioning factor only applicable to that class
- The rest of the available bandwidth can be filled with less sensitive traffic
- But don't deploy too many classes…

Conclusions

- Not "Theory of Everything", but empirical approach
- Backbone traffic is well-behaved enough to do meaningful network planning, but is not completely "smooth"
- Need several small timescale measurements to cover various types and rates of traffic

What did we learn from this example?

- On a Gigabit Ethernet (backbone) link a 'considerable' overprovisioning percentage is required to bould delay/jitter to a few milliseconds (in the order of 5-10%), on top of your overprovisioning for failures
- There is a good reason to deploy DiffServ to take care of really sensitive/critical traffic

Extra Slides

Abry-Veitch Estimator

- Wavelet decomposition
  - Discrete wavelet transform
  - Time-scale wavelet domain
- Detail variance estimation
  - Coefficients squared and averaged over time (*u[j]*)
- Analysis using *Logscale Diagram*
  - Plot *log(u[j])* vs octave *j*
- LRD parameter estimation
  - *Hurst* paramater *H* from slope of plot

Logscale Diagram
10 min. of data, 10ms samples

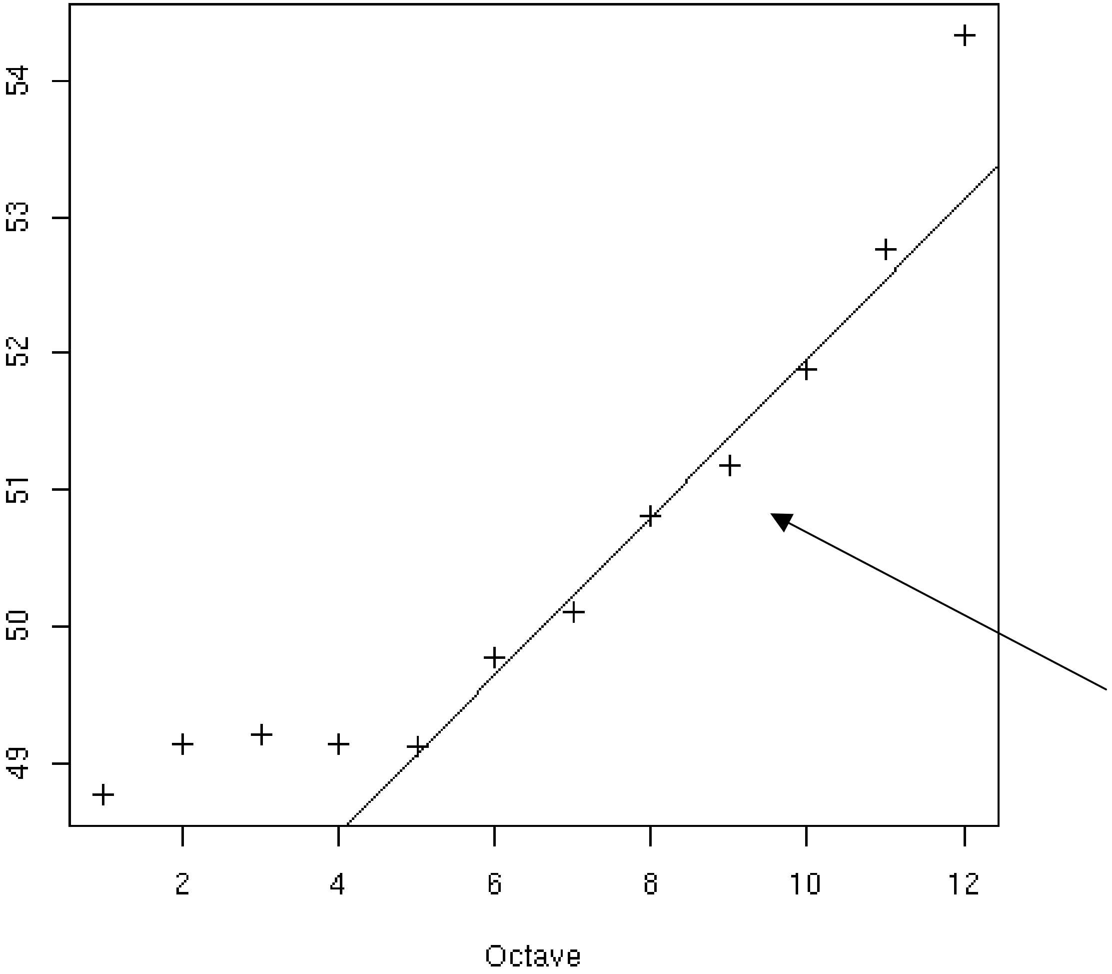

H=0.79

Logscale Diagram
60 min. of data, 1ms samples

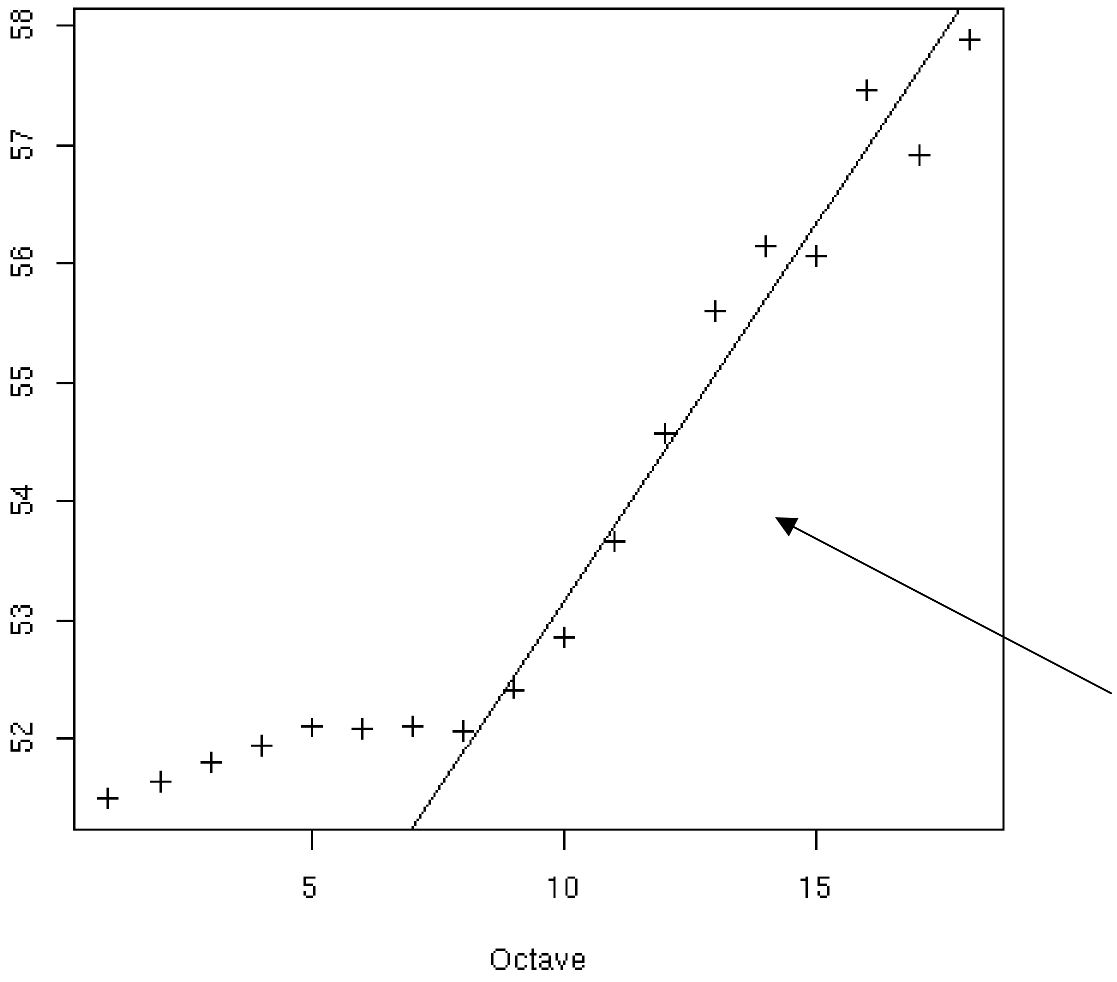

H=0.81

Geographical Traffic Profiles

- Does a world-wide network create utilization efficiencies because customers are distributed over several time zones?
- I.e. do Asian and European customer use the US network during non-peak hours?
- Yes… and No…
- Regional peaks overlap, around 3pm GMT
- Depends also on traffic ratios

Geographical Traffic Profiles

- Total
- N. America
- Europe
- Asia
- S. America

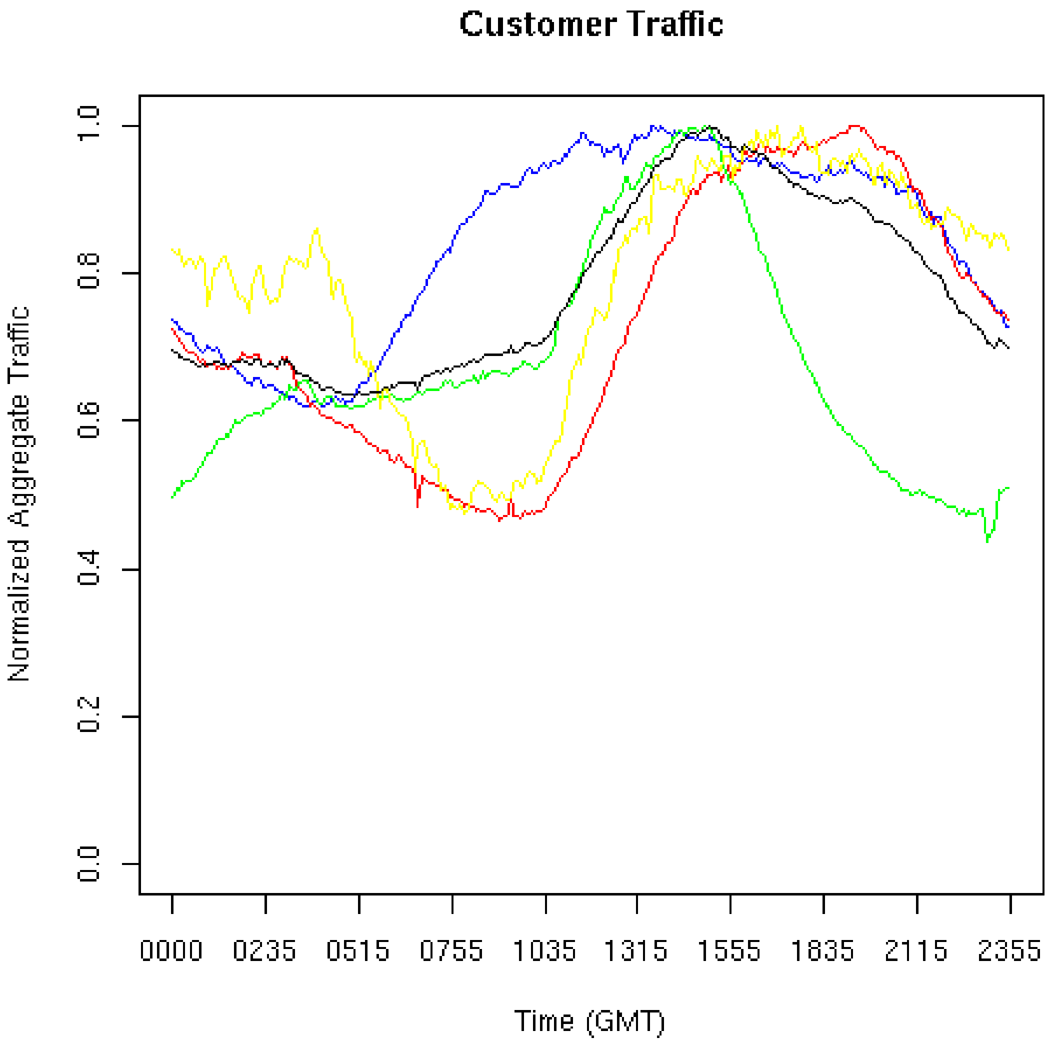

Acknowledgements

- Arman Maghbouleh (Cariden)
- Haobo Yu (Packet Design)
- Clarence Filsfils (Cisco)
- Fergal Toomey (Corvil)
- Richard Rensman (KPN)

- Upcoming Paper:
  Realizing QoS with Efficient Network Design,
  Steven Gordon, Arman Maghbouleh, Vishal
  Sharma, Thomas Telkamp

Questions?

telkamp@gblx.net

References

[1] Steve Casner, Cengiz Alaettinoglu and Chia-Chee Kuan, *A Fine-Grained View of High-Performance Networking*, NANOG 22 http://www.nanog.org/mtg-0105/casner.html

[2] Chris Liljenstolpe, *Design Issues in Next Generation Carrier Networks*, MPLS 2001 Conference

[3] Peter Lothberg, *A View of the Future: The IP-Only Internet*, NANOG 22, http://www.nanog.org/mtg-0105/lothberg.html

[4] Zafer Sahinoglu and Sirin Tekinay, *On Multimedia Networks: Self-Similar Traffic and Network Performance*, IEEE Communications Magazine, January 1999

[5] Robert Morris and Dong Lin, *Variance of Aggregated WebTraffic*, IEEE INFOCOM 2000, Tel Aviv, March 2000, pages 360-366