NANOG 26
October 27-29, 2002
Eugene, OR

Traffic Characteristics and Network Planning

Global Crossing

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

Agenda

- A methodology to analyze your traffic, and apply the results to the planning process
- Practical approach (Do try this at home!)
- 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 (single) 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**

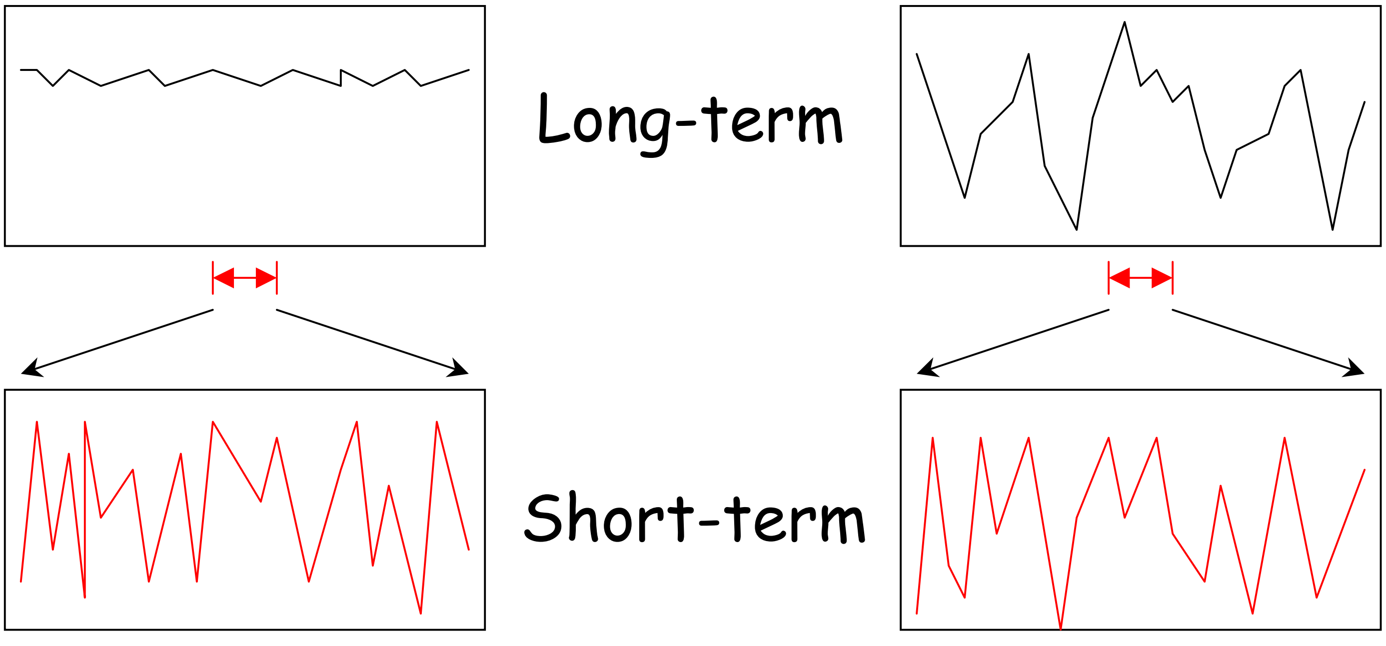

Network Planning Framework

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

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

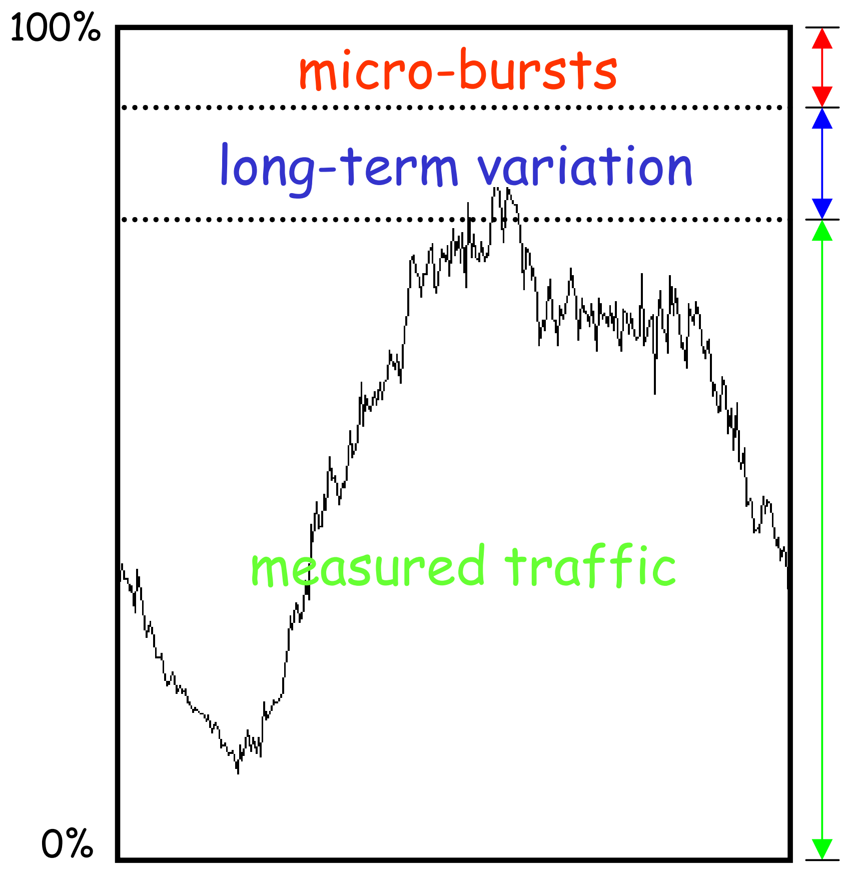

Telephony Traffic

(inter-city on 6/3/2002)

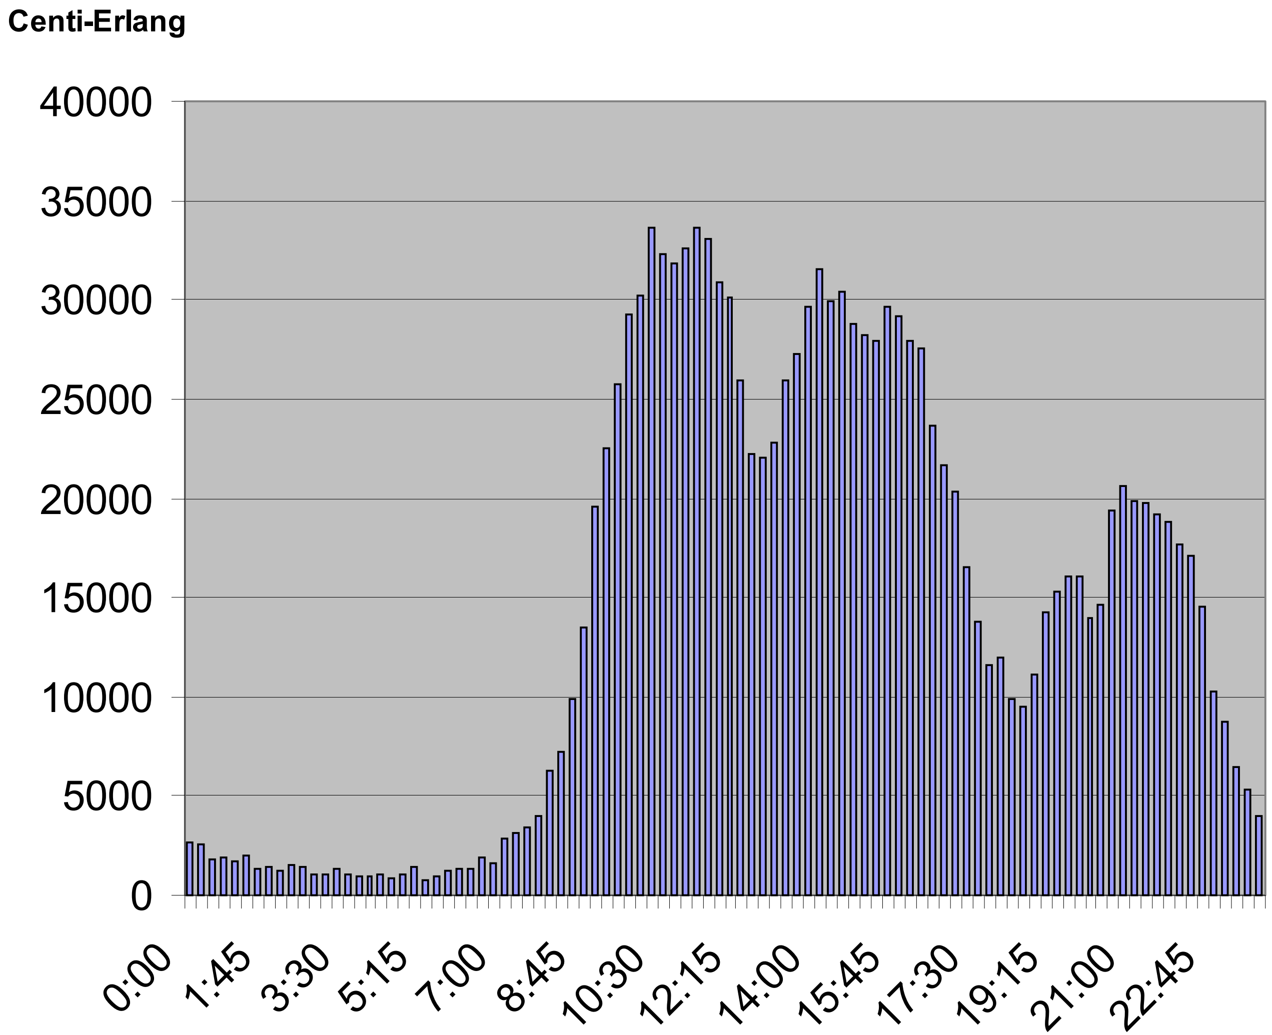

Voice Capacity Planning

- 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
- "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) <u>Empirically derive guidelines</u>

Long-term Traffic Characterization

- Investigate burstiness of 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

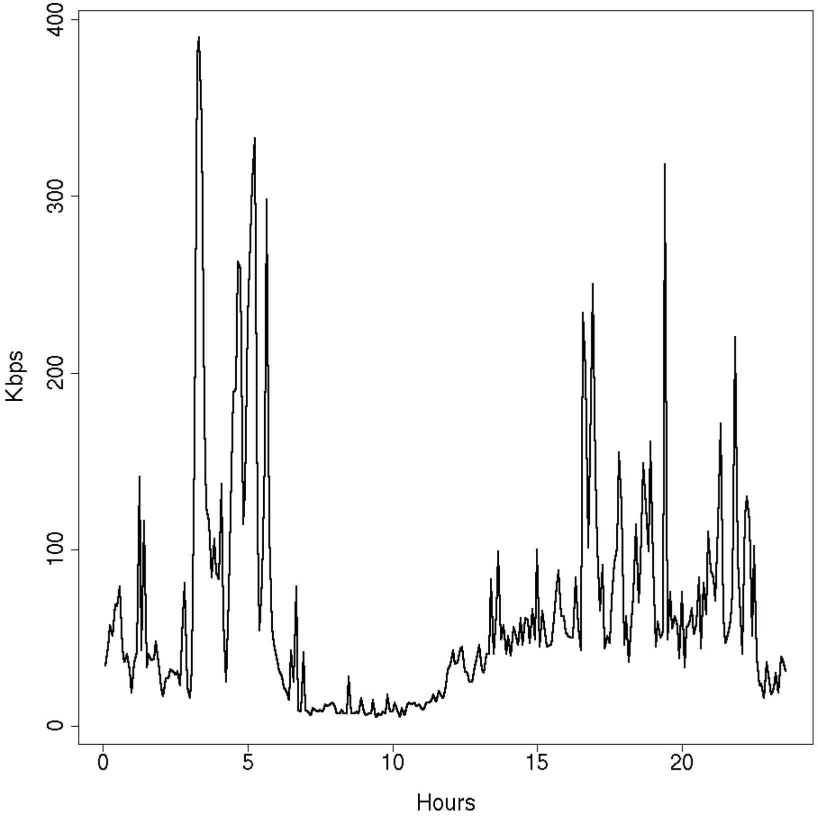

Cleveland -> Denver
Mean=64Kbps, Max=380Kbps
P95=201Kbps, Std. dev.=66Kbps
alpha=1.8 (tail index)

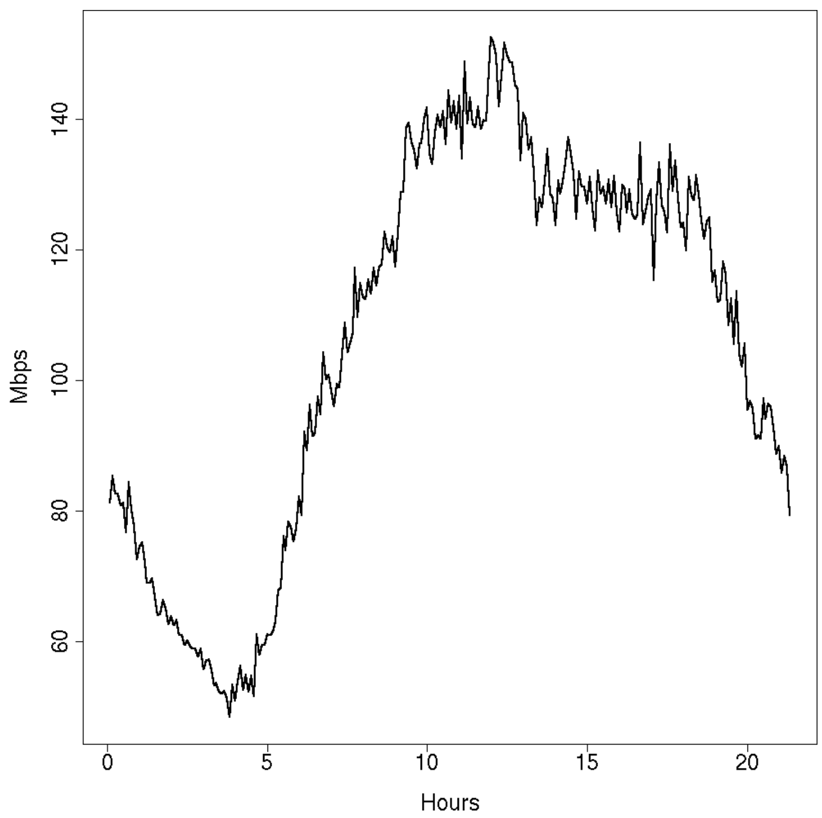

Washington D.C. -> Copenhagen
Mean=106Mbps, Max=152Mbps
P95=144Mbps, Std. dev=30Mbps
alpha=21 (tail index)

Variance vs Bandwidth

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

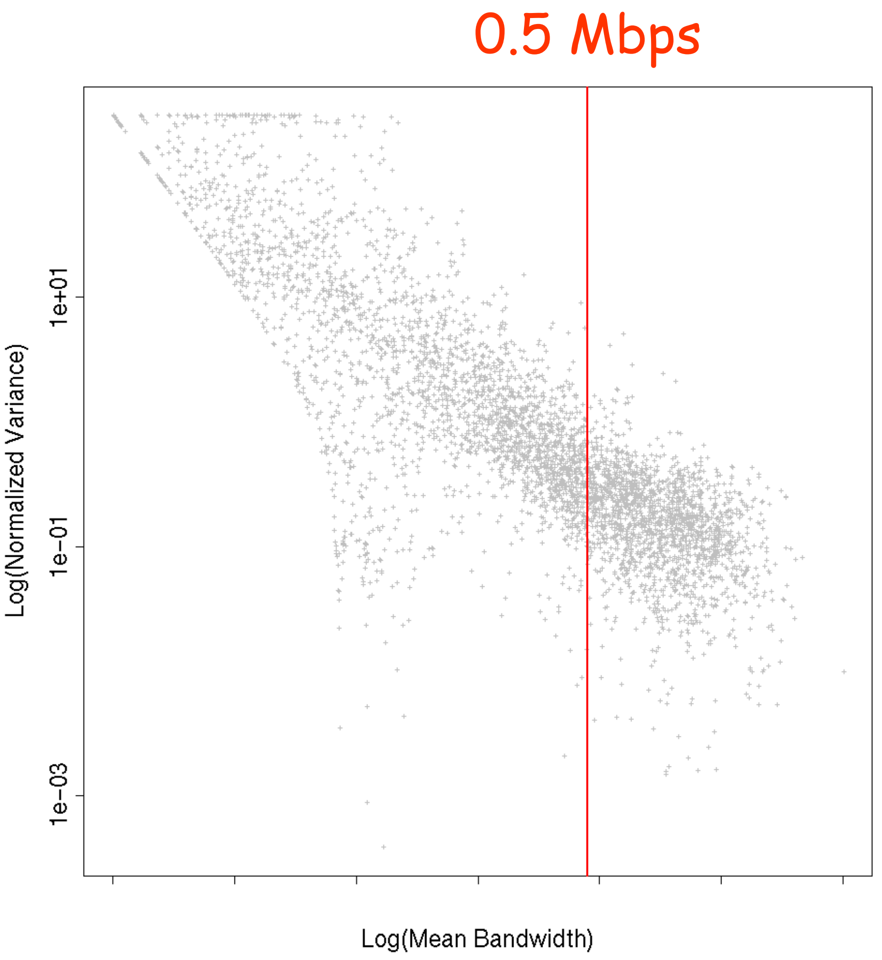

Short-term Traffic Characterization

- 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)

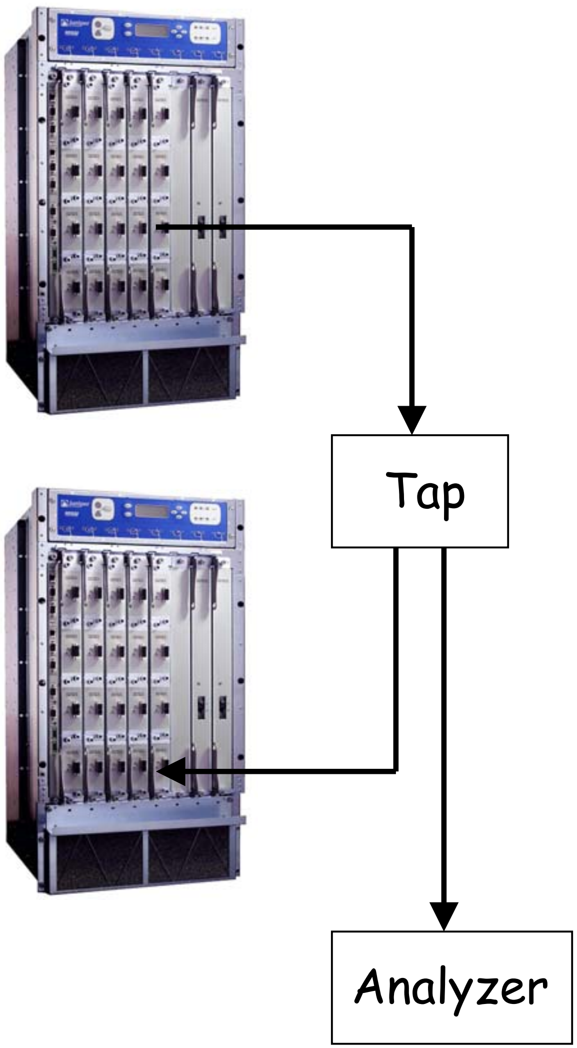

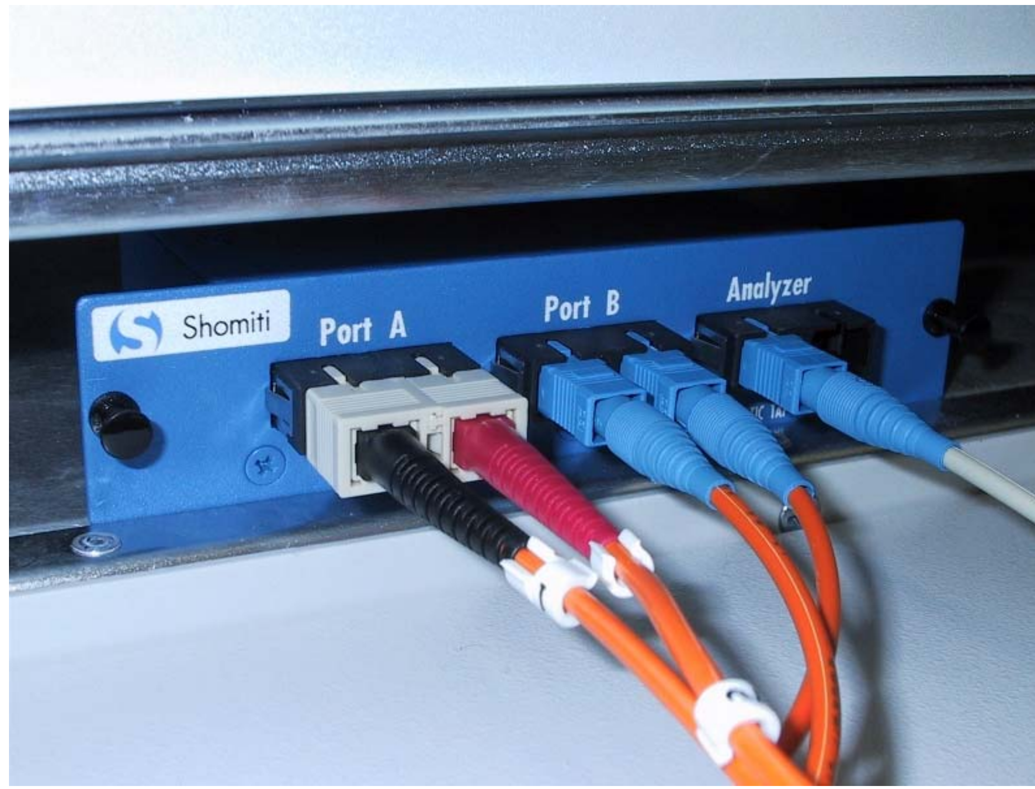

Raw Results
10 min. of data, 10ms scale

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

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

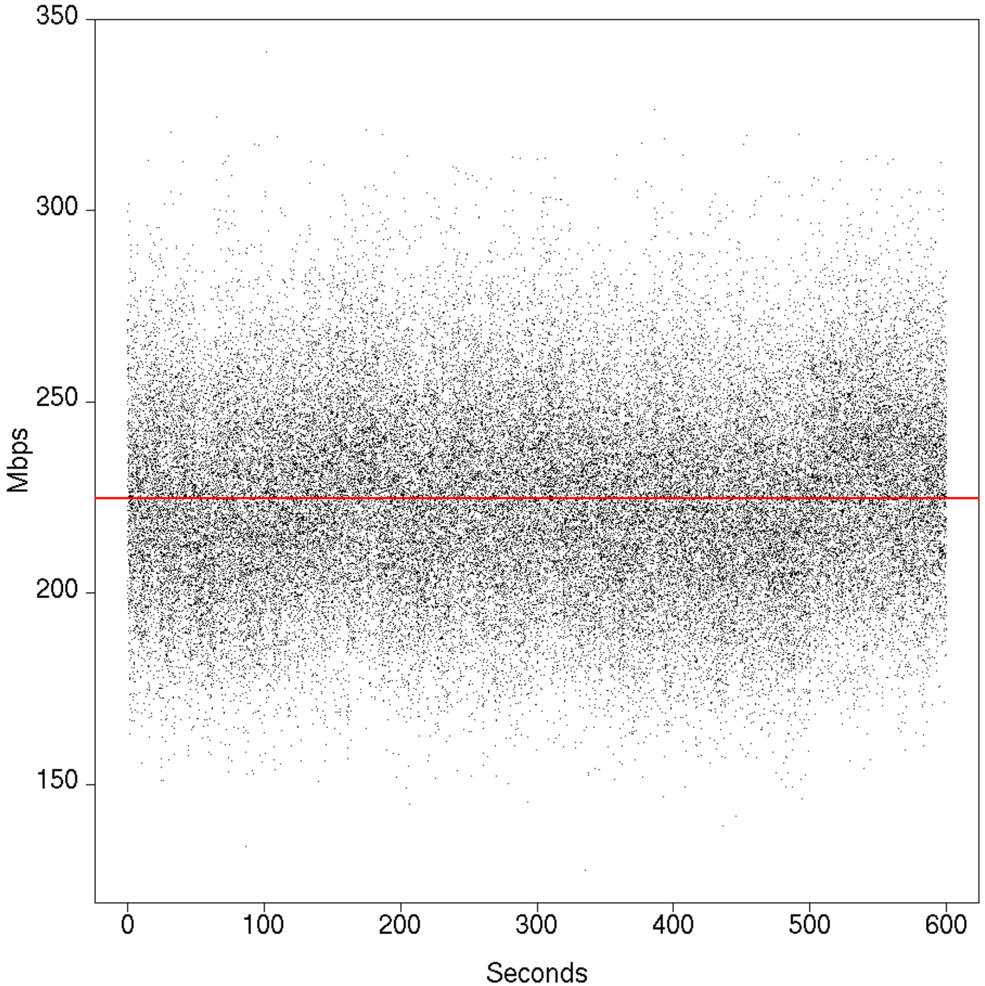

Traffic Distribution Histogram (10ms scale)

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

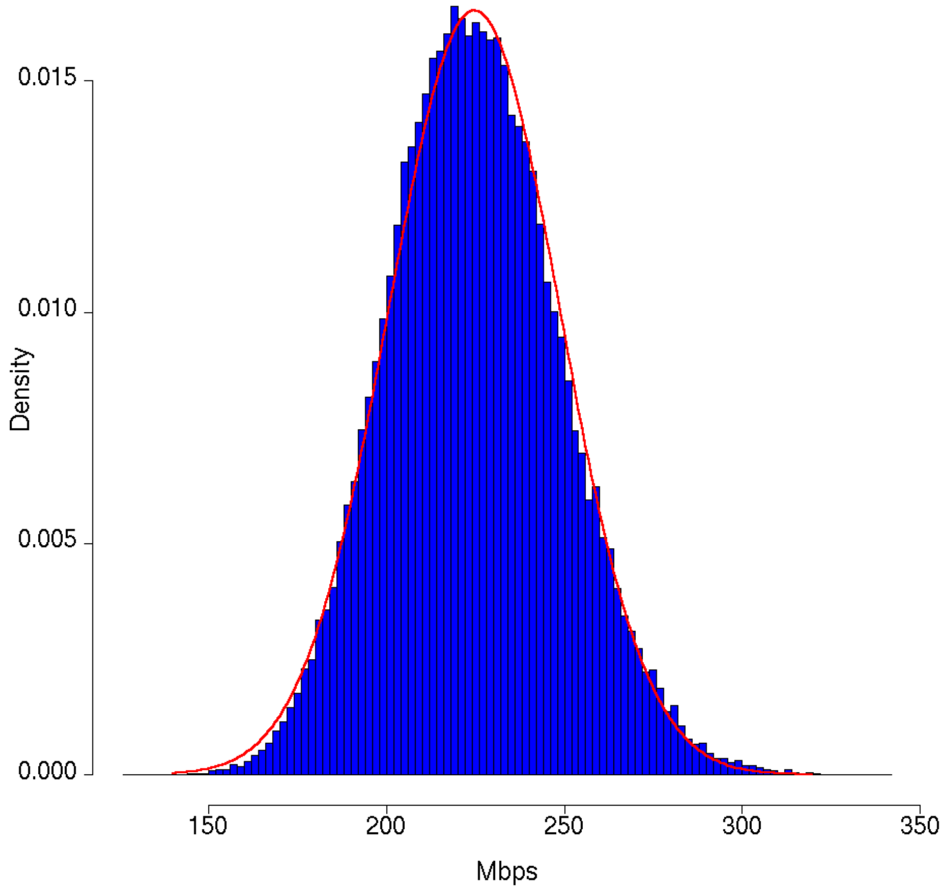

Autocorrelation
Lag Plot (10ms scale)

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

- Autocorrelation at 10ms is 0.16 (=uncorrelated)

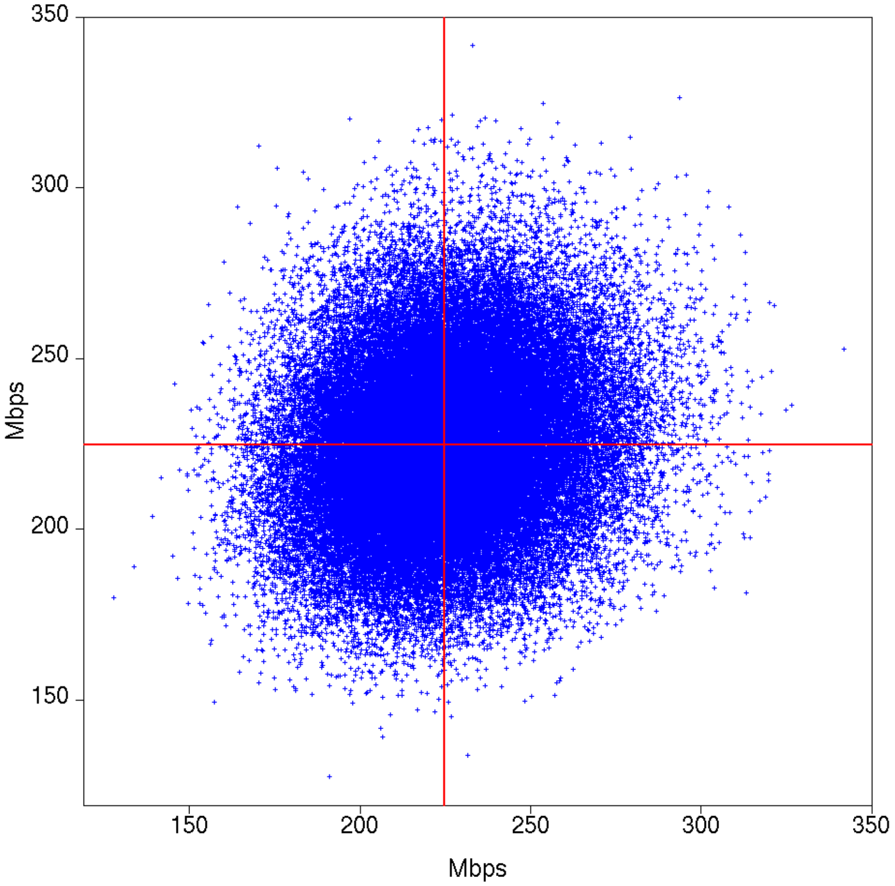

Utilization
10 min. of data, 10 sec. scale

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

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

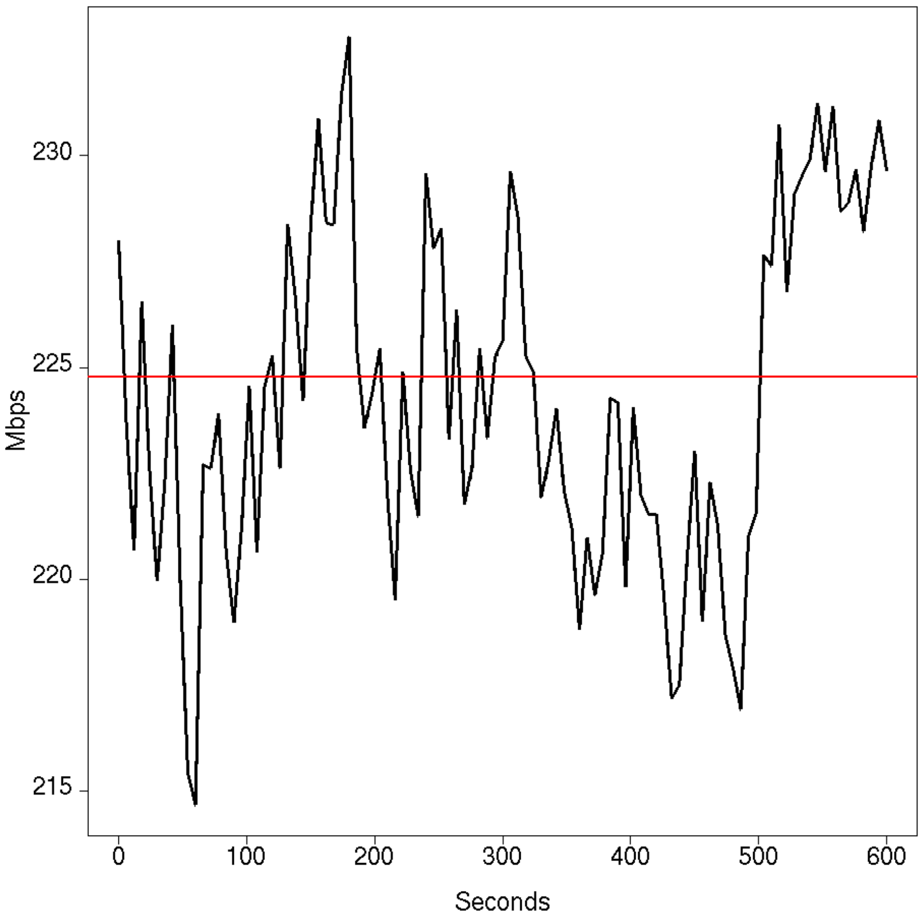

Queuing Simulation

- 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

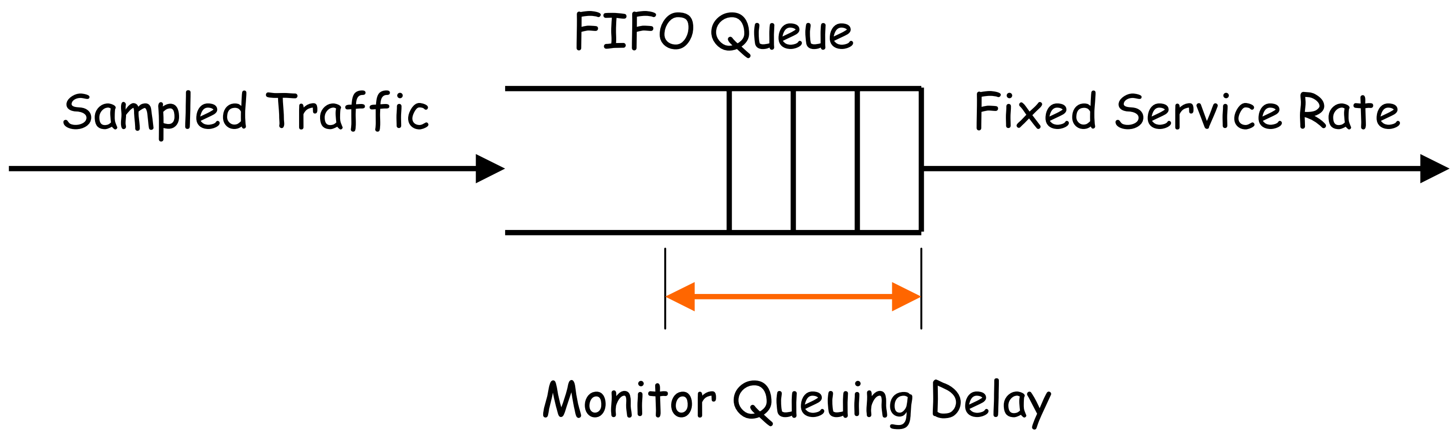

Bandwidth Requirement vs Delay Bound

- 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%

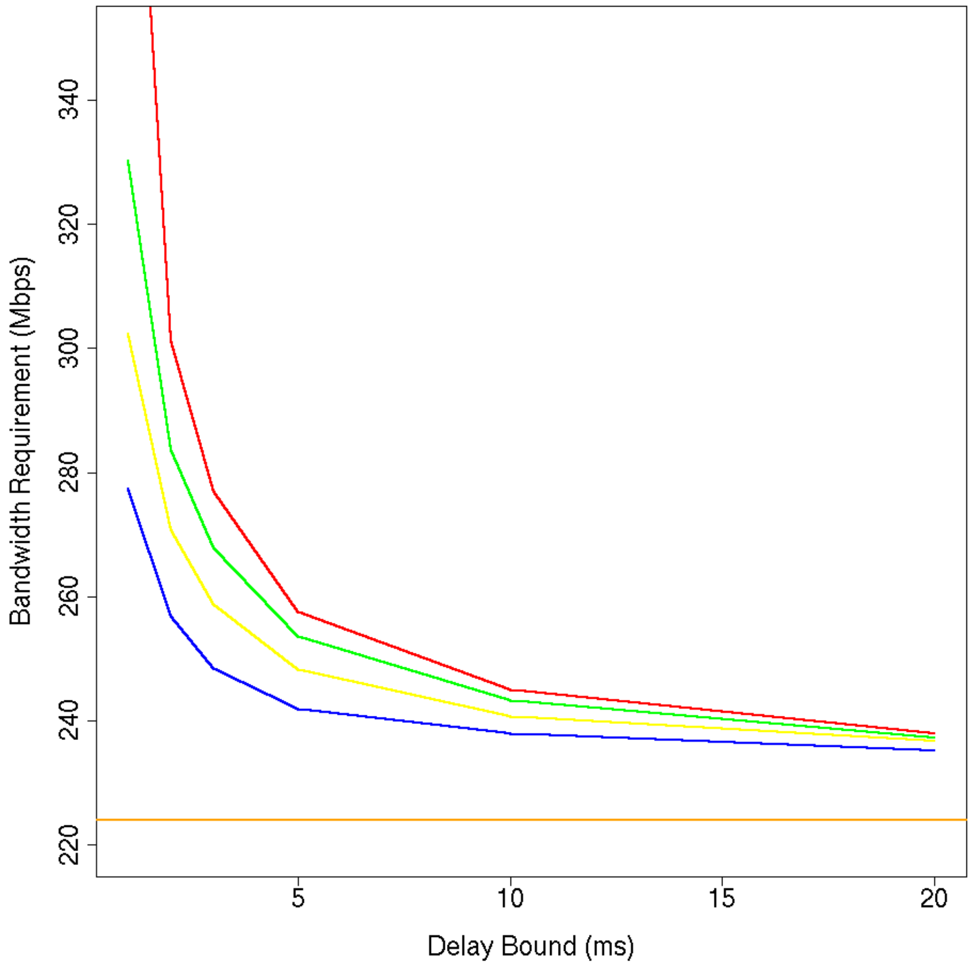

Bandwidth Requirements Numeric Results

- 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

- Synthesized data
- 5ms delay bound, for 99.9999% of the traffic (10^{-6})

- 647 Mbps: 10% overprovisioning
  - BW required: 715 Mbps
- 963 Mbps: 6% overprovisioning
  - BW required: 1021 Mbps
- 1.185 Gbps: 5% overprovisioning
  - BW required: 1.246 Gbps

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

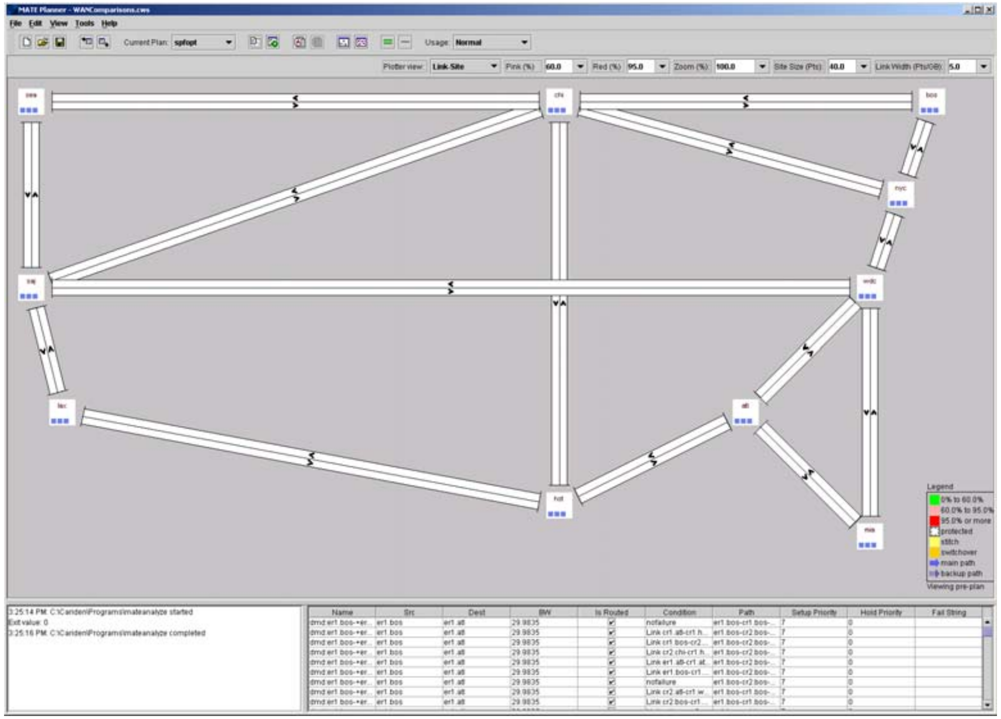

How does Diff-Serv fit in this picture?

- All traffic in one class (no Diff-Serv) requires overprovisioning factor that matches the tight objectives (e.g. low delay/jitter for VoIP) *for all traffic*
- Prioritizing that traffic (using a Strict Priority Queue) would make the overprovisioning factor only applicable to that class
- The rest of the available bandwidth can be filled with less sensitive traffic
- More complicated for WRR/MDDR queuing

General Conclusions

- Not "Theory of Everything", but empirical approach
- Backbone traffic is well-behaved enough to do meaningful network planning, but is not completely "smooth"
- Several small timescale measurements are needed to cover various types and rates of traffic
- An overprovisioning factor per link type/speed can be empirically derived

Conclusions from Example Data

- On a Gigabit Ethernet (backbone) link an overprovisioning percentage in the order of 5-10% is required to bound delay/jitter to less than 5 ms, on top of your overprovisioning for failures
- At lower speeds (<1G) this overprovisioning factor is significant, but at higher speeds (2.5G/10G) it becomes very small

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