Bell Labs Internet Traffic Research

Jin Cao, Bill Cleveland, and Don Sun {cao, wsc, dxsun}@bell-labs.com

Measurement, analysis, and modeling of Internet traffic for device design, network performance monitoring, and traffic engineering

http://cm.bell-labs.com/stat/InternetTraffic

50% of Bell Labs NMS project (Anwar Elwalid will describe other 50%)

Collaborators

- Bell Labs: Dong Lin, Francis Zane, and Anwar Elwalid
- Waikato and NLANR: Joerg Micheel
- NCNI/NCREN: Dan Stevenson and Ilia Baldine (DARPA Helios Project)
- UNC Computer Science: Don Smith and Michele Weigle
- Rice: Rui Castro

Overview of Research Program (DARPA NMS Project)

S-Net

- an open source UNIX software system for comprehensive statistical analysis of IP packet header databases

Packet header databases

- 1 terabyte of compressed headers from 20 Internet interfaces, 10 mbps to 2.5 gbps

Statistical modeling

- effect of magnitude of statistical multiplexing on statistical properties of Internet traffic variables
- FSD and FSD-MA(1): statistical models for packet sizes and inter-arrivals for open-loop generation of packet traffic
- PackMime: models for application connection variables (e.g., HTTP 1.1)

Monitoring and performance

- local network firewall
- performance by application

Device engineering

- user behavior and the performance of access devices
- packet sampling on routers and switches

Traffic engineering

- bandwidth allocation: how much traffic can a link of x mbps carry and maintain a fixed QoS criterion
- input to MPLS traffic optimization algorithms (Anwar Elwalid's presentation)

Our Environment for Empirical Study: S-Net and Databases

S-Net (DARPA Helios Project)

- an open source UNIX software system for comprehensive statistical analysis of IP packet header databases
- based on S system for graphics and data analysis
- implemented as computing environment distributed across network
- extensible
- dataflow management tools
  - raw packet header traces
  - to Unix flat files
  - to S objects

Packet header database

- 1 terabyte of compressed headers
- 20 Internet interfaces
- 10 mbps to 2.5 gbps
- trace sources: NLANR, U. of Waikato DAG group, NCNI/NCREN, and Bell Labs
- Two databases
  - cluster of linux PCs at Bell Labs
  - cluster of linux PCs at NCNI facility

Statistical Multiplexing

Packet traffic on a link

- the packets of simultaneous active connections are intermingled on the link
- “statistical multiplexing” in the packet network literature
- “superposition” in the statistical point-process literature

Connection load

- measure of magnitude of multiplexing
- at a given time: the number of active connections
- over an interval
  - the average number over the times in the interval
  - sum of the connection durations divided by the interval length

What is the effect of the connection load on statistical properties of three link packet traffic variables: packet inter-arrivals, packet sizes, and packet counts?

Results: The Effect of Increased Multiplexing

Starting at low connection loads with little congestion on the link: three traffic variables are long-range dependent

As load increases, attraction to Poisson arrivals and independent sizes

Sizes and inter-arrivals always retain LRD components, but the component variances go to zero as the connection load increases

Variability of the counts declines like one over the square root of the connection load

Autocorrelation of counts is stable so LRD remains, fueled by remaining but diminishing component of LRD in inter-arrivals

Results based on theory, empirical study, and simulation

If the link speed is high enough, the traffic can get quite close to Poisson and independent before the push-back begins in force

The OC48 DAG measurements we have studied are very accurate ($\pm$ 50 nanosec) and very close to Poisson and independent

Results: The Effect of Increased Multiplexing

Push-back on the attraction to Poisson and independence is caused by queueing on the link-input router or further upstream

Not Altered

- attraction of sizes to independence (and therefore dissipation of LRD)
- dissipation of LRD of inter-arrivals

Altered in Theory But Not Yet Observed

- decline in variability of the counts like one over square root of connection load
- stability of autocorrelation of the counts

Altered

- attraction of the marginal distribution of the inter-arrivals to exponential is altered to attraction to marginal of packet size divided by link speed
- attraction of inter-arrivals to independence altered to attraction to very short-range dependence

FSD (Fractional Sum-Difference) Time Series Models

Simple time series models fit the statistical variability of sizes and inter-arrivals account for the multiplexing effects

Can be used to generate open-loop packet arrivals and sizes for simulation studies

Let z_u be FSD time series

$$z_u = \sqrt{1-\theta}\; s_u + \sqrt{\theta}\; n_u,$$

where

s_u = LRD Gaussian time series with mean 0 and variance 1

n_u = Gaussian white noise mean 0 and variance 1

s_u and n_u are independent of one another

$$0 \leq \theta \leq 1$$

FSD Models

s_u is a fractional ARIMA model

$$(I - B)^d s_u = \epsilon_u + \epsilon_{u-1}$$

where

$$Bs_u = s_{u-1}$$

$$0 < d < 0.5$$

ϵ_u is Gaussian white noise with mean 0 and variance σ_ϵ^2

$$\sigma_\epsilon^2 = \frac{(1-d)\Gamma^2(1-d)}{2\Gamma(1-2d)}$$

Name "FSD" because

$$(I - B)^d z_j = \sqrt{1-\theta}\,(I + B)\epsilon_j + (I - B)^d n_j$$

If $\theta < 1$, z_u has LRD component, and if $\theta = 1$, z_u is white noise

FSD-MA(1)

Replace white-noise n_u by first order moving-average

$$n_u = \zeta_u + \beta\zeta_{u-1}$$

where ζ_u is Gaussian white noise with $E(\zeta_u) = 0$,
and $E(\zeta_u^2) = (1+\beta^2)^{-1}$ to make $E(n_u^2) = 1$

If $\beta = 0$, the moving-average component is white noise, so model is an FSD

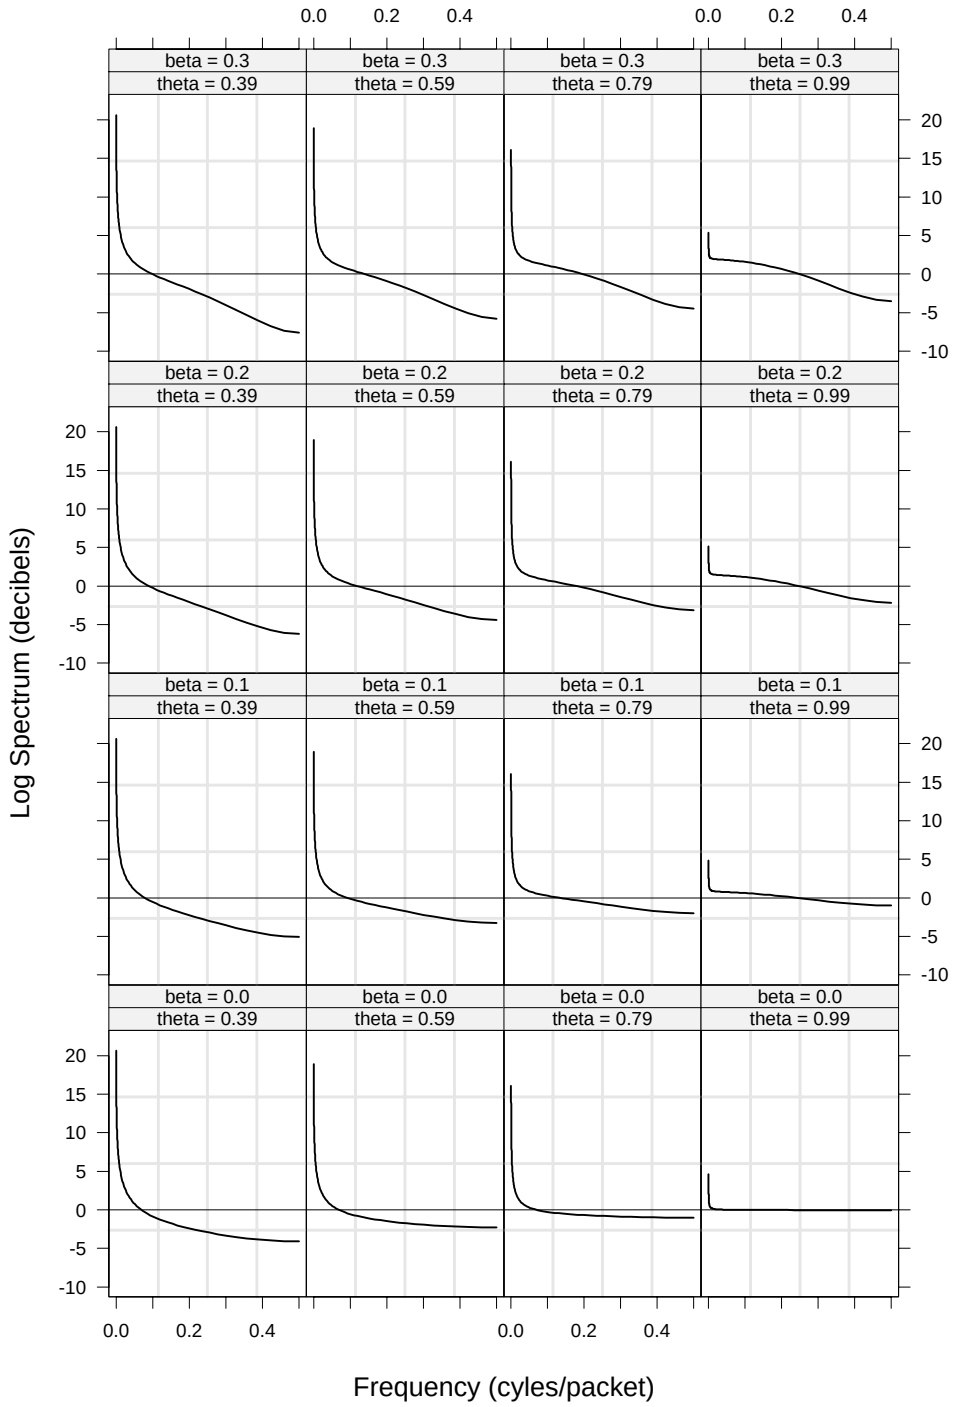
Power spectrum, $p(f)$, at frequency f

$$p_z(f) = (1-\theta)\sigma_\epsilon^2 \frac{4\cos^2(\pi f)}{\left(4\sin^2(\pi f)\right)^d} + \theta \frac{1+\beta^2+2\beta\cos(2\pi f)}{1+\beta^2}$$

so

$$p(f) \sim \sin^{-2d}(\pi f) \text{ as } f \to 0$$

Next, $p(f)$ is plotted against f for different values of θ and β


beta = 0.3
theta = 0.39
beta = 0.3
theta = 0.59
beta = 0.3
theta = 0.79
beta = 0.3
theta = 0.99
beta = 0.2
theta = 0.39
beta = 0.2
theta = 0.59
beta = 0.2
theta = 0.79
beta = 0.2
theta = 0.99
beta = 0.1
theta = 0.39
beta = 0.1
theta = 0.59
beta = 0.1
theta = 0.79
beta = 0.1
theta = 0.99
beta = 0.0
theta = 0.39
beta = 0.0
theta = 0.59
beta = 0.0
theta = 0.79
beta = 0.0
theta = 0.99
Log Spectrum (decibels)
Frequency (cyles/packet)
0.0
0.2
0.4
20
15
10
5
0
-5
-10

Modeling Sizes and Inter-Arrivals

Marginal distribution of t_j: Weibull with shape λ

New variables

- t_j^*: transform t_j (using Weibull model) to have Gaussian marginal, and then normalize to have mean 0 and variance 1
- q_j^*: normalize q_j to have mean 0 and variance 1

New variables both modeled as FSD-MA(1) time series

FSD Models for Sizes and Inter-Arrivals

As the connection load increases:

t_j (inter-arrivals): $\lambda \to 1$

q_j (sizes): marginal does not change appreciably as connection load increases on a link, but does change from link to link

q_j^* (normalized sizes):

- $\beta = 0$
- $\theta \to 1$

t_j^* (normalized, transformed inter-arrivals):

- $\beta = 0$ until upstream queueing becomes significant and then β becomes positive
- $\theta \to 1$

t_j^* & q_j^*: d is constant, equal to about 0.4

The models do not reproduce statistical characteristics at very low loads (about 5 mpbs and less) when TCP effects of individual connections have a significant effect on the variability, but above this value the model appear to fit well.

For generation of q_j, must transform FSD output to have marginal of q_j

NMS Project Study

FSD Models

The models have been put through extensive validation using a large number of visualization tools from S-Net and are close to use for Internet studies.

For convenience of use in generation, develop models that relate the FSD parameters to the connection load (or equivalently, the load in bits/sec)

Develop C code for generation and make publicly available.

To what extent can synthetic packet traffic generated by the models can serve as reliable substitute for live traces in open-loop studies?

Do the models serve as a mechanism for helping to explain Internet scaling effects uncovered by wavelet analyses of live data in the sense that the wavelet analyses show similar results when applied to synthetic traffic from the models?

Utilization

The multiplexing results show qualitatively that if we fix a buffer size and fix an amount traffic as measured by the connection load, then the link speed needed to achieve a fixed QoS criterion, such as 0.5% packet loss, results in a utilization that increases with the load.

Through

- extensive study of live traffic
- open-loop simulation with FSD models
- closed-loop study using NS and PackMime application models

study the bandwidth allocation problem: the traffic rate that a link of x mbps can carry and maintain a fixed QoS criterion.

Use results in MPLS traffic engineering studies using optimization methods (Anwar Elwalid's presentation).