Simulation Modeling of Self-Similarity in Network Traffic

Robert Geist

Kevin Spicer

James Westall

Department of Computer Science

Clemson University

Clemson, SC 29634

http://www.cs.clemson.edu/~ westall/homepage.html

Outline

- Introduction to network traffic modeling
- Long-range dependence in stochastic processes
- Fast synthesis of self-similar processes
- Performance comparison of the sythesizers
- Modeling real traffic with fGn
- Conclusion

Network Performance Models

Objectives include prediction of...

- Queuing delays
- Packet loss rates

Experienced by network traffic at...

- Switching elements
- Transmission elements

As a function of...

- Switching or transmission element utilization
- Distributional characteristics of the arrival process

Poisson Arrival Models

On the positive side...

- Support analytic models of queuing behavior

On the not so positive side...

- Arrivals per unit time are necessarily I.I.D.
- Variance is equal to the mean

Characteristics of Real Network Traffic

(Packet arrivals per unit time)

Distributional characteristics...

- Variance $>>$ Mean
- Implications...
  - Arrival counts $>>$ mean are not rare..
  - Neither are arrival counts $<<$ mean.

Correlational characteristics...

- Successive arrival counts are *not* independent
- Evidence of *long-range dependence* is common
- Implications of long-range dependence
  - Arrival counts $>>$ mean occur in bursts
  - So do arrival counts $<<$ mean.

Modeling the Behavior of Real Traffic

Characteristics of real traffic imply...

- Poisson models are not adequate
- Simulation models are attractive

Simulation may be driven by...

- Captured traces of real traffic
- Synthetic arrivals

Synthetic traffic is attractive because...

- Need to capture and store traces is eliminated
- Exploration of entire design space is expedited

Synthesis of Simulated Network Traffic

One-step approach...

- Direct synthesis of interarrival times

Two-step approach...

- One procedure synthesizes
  - Arrival counts per unit time
- Another synthesizes
  - Interarrival times from arrival counts
- Facilitates the modeling of correlational effects

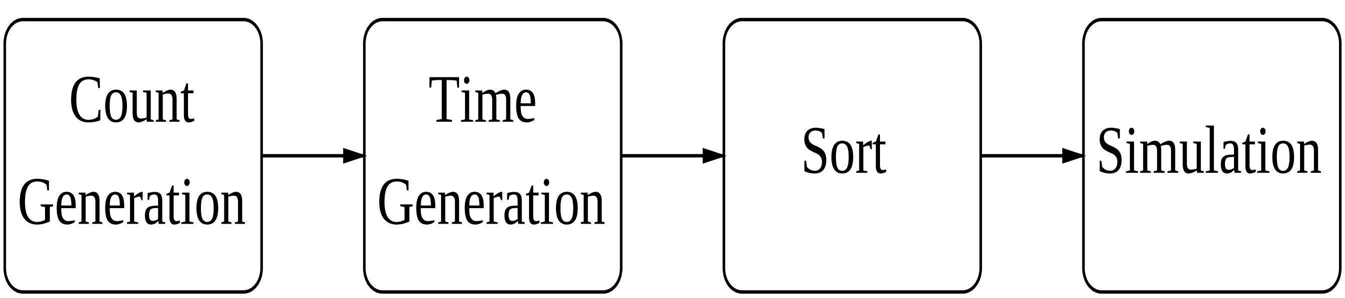

Synthesis of Arrival Counts

Objective:
Generate a time-series of integers representative of the

- distributional and
- correlational

characteristics of the number of packet arrivals per unit time of a real or hypothetical workload.

Capturing distributional aspects...

- Sample from an empirical distribution
- Match (two) moments of target distribution

Capturing correlational aspects...

- requires further discussion!

Self-Similarity and Long-Range Dependence

Let $X(t)$ be a stochastic process defined on a discrete set of equally-spaced values $\{t_i : i = 0, 1, \ldots\}$.

The increment process $Z(t)$ is defined by

$$Z(t_i) = X(t_i) - X(t_{i-1})$$

The autocorrelation of the increment process $Z(t)$,

$$r(k) = E[(Z(t_j) - \mu)(Z(t_{j+k}) - \mu)]/\sigma^2,$$

characterizes the long-range dependence of $X(t)$. $X(t)$ is said to be

- long-range dependent if
  $\sum_{k=0}^{\infty} r(k) = \infty$
- asymptotically self-similar if
  $r(k) \propto k^{(2H-2)} L(k)$
- exactly self-similar if
  $r(k) = (1/2)[(k+1)^{2H} - 2k^{2H} + (k-1)^{2H}]$

The parameter H, $0.5 < H < 1$, is called the Hurst Parameter.

Exactly Self-Similar Processes

If $X(t)$ is exactly self-similar, then

- $X(wt)$ is equal in distribution to $w^H X(t)$ for all values of t and $w > 0$.
- Equivalently stated, the "speeded up" process is a scaled copy of the original.

An m-aggregated increment process $Z(t_n)^{(m)}$ is formed by summing adjacent, length m, blocks of the increment process:

$$Z(t_n)^{(m)} = \sum_{j=nm}^{j=m(n+1)-1} Z(t_j)$$

If $X(t)$ is exactly self-similar, then

- $Z(t)^{(m)}$ inherits $Z(t)$'s autocorrelation $r(k)$.
- $V[Z(t)^{(m)}] = m^{2H} V[Z(t)]$.

Fractional Brownian motion (fBm) is a widely studied exactly self-similar process.

Fractional Brownian Motion

An fBm process $X(t)$ is a continuous stochastic process having the following characteristics:

- $X(0) = 0$
- the increment processes,
  $Z(t_k - t_j) = X(t_k) - X(t_j)$,
  have a Gaussian distribution for all $t_k > t_j$.
- $E[Z(t_k - t_j)] = 0$
- $V[Z(t_k - t_j)] = K(t_k - t_j)^{2H}$

If $X(t)$ satisfies these conditions,

- $V[w^H X(t)] = w^{2H} V[X(t)] = w^{2H} K t^{2H}$.
- $V[X(wt)] = K(wt)^{2H}$.

Since $w^H X(t)$ and $X(wt)$ are both Gaussian, exact self-similarity of fBm follows.

fBm and Synthesis of Network Traffic

If $X(t)$ is an fBm process, the increment process, $Z(t)$, is known as fractional Gaussian noise (fGn).

Suppose $X(t_i)$ defined on $\{t_i : i = 0, 1, \ldots\}$ is a *sample* from the fBm process $X(t)$ and $Z(t_i) = X(t_{i+1}) - X(t_i)$ has sample mean μ_z and variance ${\sigma_z}^2$

An arrival process of mean μ_a and variance ${\sigma_a}^2$ can be obtained by the H preserving linear transform:

$$A(t_i) = ({\sigma_a}^2 / {\sigma_z}^2)(Z(t_i) - \mu_z) + \mu_a$$

Therefore, a mechanism for synthesizing a sample, $X(t_i)$, from a target fBm process or a sample $Z(t_i)$ from the corresponding fGn process, yields synthetic traffic in which arrivals per unit time have:

- A Gaussian distribution
- Any desired mean and variance
- The target Hurst parameter, H.

Fast Synthesis of Approximate fBm

useful in many contexts:

techniques:

- midpoint displacement
- successive random additions
- floating proportionality

Midpoint Displacement

Specify values of $X(t)$ on interval $[0, T]$ at successive mid-points:

- $X(0) = 0$
- $X(T)$ is Gaussian, mean 0, $V[X(T)] = \sigma^2 T^{2H}$
- $X(T/2) = (X(0) + X(T))/2 + D$

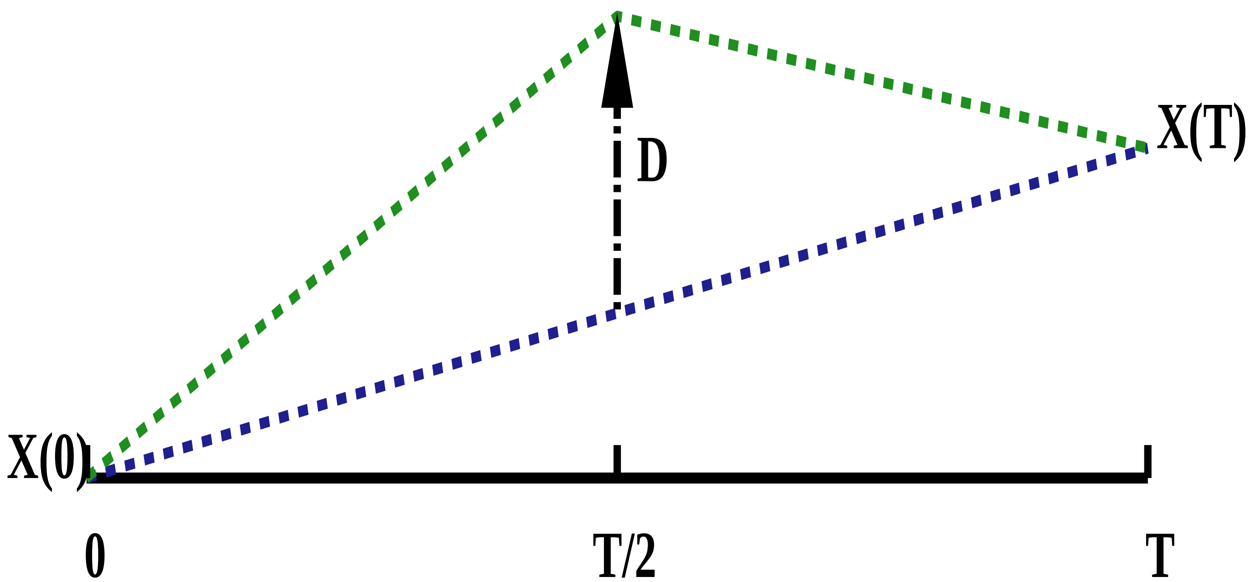

D is Gaussian, mean 0, variance chosen so that:

$$V[X(T/2) - X(0)] = \sigma^2 (T/2)^{2H}$$

Midpoint Displacement

(continued)

- repeat for sub-intervals $[0, T/2]$ and $[T/2, T]$
- continue to arbitrary level
- different displacements required at each level:

$$V[D_i] = \sigma^2 T^{2H} \left[\frac{1 - 2^{2H-2}}{(2^{2H})^{i+1}} \right]$$

for level $i \geq 0$ guarantees that:

$$V[X(T/2^{i+1}) - X(0)] = \sigma^2 (T/2^{i+1})^{2H}$$

Oops: after two subdivisions

$$V[X(T) - X(T/4)] \neq \sigma^2 T^{2H} (3/4)^{2H}$$

unless $H = 1/2$

Successive Random Additions

- midpoint displacement yields visual artifacts
- approximate fBm is not “believable”
- variation: displace all points on each iteration, not just the new midpoints

Specifically:

- let $D_i(t)$ = i.i.d. displacement for level i
- let $X_i(t)$ = constructed process at level i, e.g.,

$$\begin{aligned} X_1(0) &= X_0(0) + D_0(0) \\ X_1(T/2) &= (X_0(0) + X_0(T))/2 + D_0(T/2) \\ X_1(T) &= X_0(T) + D_0(T) \end{aligned}$$

  $V[D_0]$: so that $V[X_1(T/2) - X_1(0)] = \sigma^2 (T/2)^{2H}$

- $V[D_i]$ = half of that for std. midpoint displacement

Oops: $V[X_{+\infty}(T) - X_{+\infty}(0)] \neq \sigma^2 T^{2H}$

Floating Proportionality

use successive random additions, but change proportionality constant per level (not fixed at σ^2)

- let $K_0 = 1$ and for level $i \geq 0$ define

$$
\begin{aligned}
V[D_i] &= (K_i/2)T^{2H}\left[\frac{1-2^{2H-2}}{(2^{2H})^{i+1}-1}\right] \\
K_{i+1} &= K_i + 2V[D_i]/T^{2H}
\end{aligned}
$$

- this guarantees:

$$
\begin{aligned}
V[X_i(T) - X_i(0)] &= K_i T^{2H} \\
V[X_i(T/2^i) - X_i(0)] &= K_i (T/2^i)^{2H}
\end{aligned}
$$

- K_{final} can be rescaled to match target σ^2

Direct synthesis of fGn

(Method developed by V. Paxson)

The power spectrum of an fGn process is given by:

$$f(\lambda;H) = A(\lambda,H)(|\lambda|^{-2H-1} + B(\lambda,H))$$

$$A(\lambda;H) = 2sin((\pi H)\Gamma(2H+1)(1-cos(\lambda))$$

$$B(\lambda;H) = \sum_{j=1}^{\infty}((2\pi j+\lambda)^{2H-1} + (2\pi j-\lambda)^{2H-1})$$

Transform space synthesis of fGn

- Construct $\{f_1, \ldots, f_{n/2}\}$ where $f_j = f(2\pi j/n;H)$
- Multiply the f_j's by independent exponential r.v.'s with mean 1
- Construct $\{z_1, .. \; z_{n/2}\}$, where $z_j = \sqrt{f_j}e^{i\theta}$ and θ is a uniform r.v. in $(0,2\pi)$
- For $n/2 < j < n$ let $z_j = \bar{z}_{n/2-j}$
- Apply the inverse D.F.T. to $\{z_j\}$.

Evaluation of the Three Synthesizers

(Hurst parameter matching)

Results for matching a target $H = 0.80$

	fBm	Pax	rmd
R/S	0.796 0.017	0.798 0.013	0.790 0.013
V/T	0.767 0.037	0.775 0.029	0.767 0.034
Whit	0.745 0.009	0.807 0.004	0.807 0.004

Table 1: μ_H and σ_H for 40 runs.

Evaluation of the Three Synthesizers

(Sample autocorrelation)

An empirical look that the sample autocorrelation..

For N observations

$$\hat{r}(k) = \frac{1}{N-k} \sum_{j=1}^{N-k} (Z_j - \mu)(Z_{j+k} - \mu)/\sigma^2$$

is an asymptotically unbiased estimator of $r(k)$

Here N = 32768 and each point represents the mean of 64 runs.

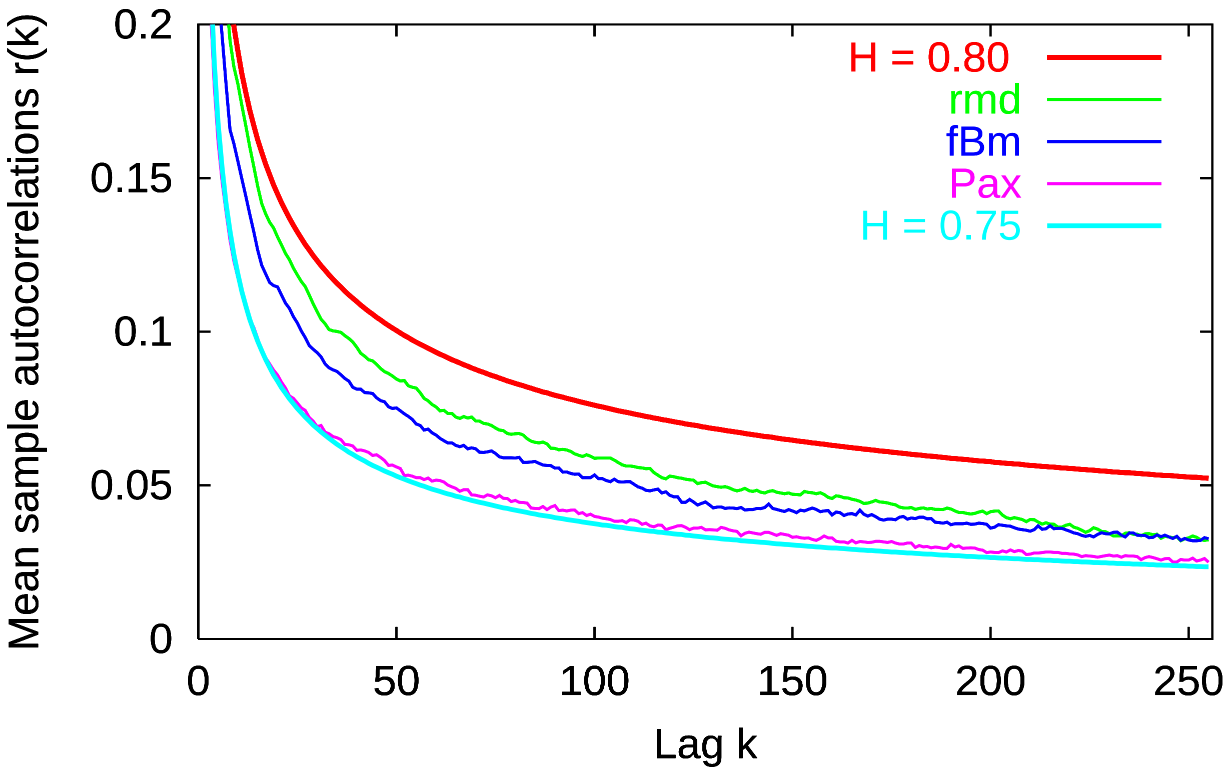

Evaluation of the Three Synthesizers

(Simulated system performance)

Configuration and parameterization:

- G/D/1 FCFS simulated queuing system
- Parameters: (fGn, 100, 625, 0.80)
- Simulated run time: N = 32768 seconds
- Simulated transactions: 3.27 million

ρ	MD1	fBm	rmd	Pax	pxe
.20	0.23	0.23	0.23	0.23	0.29
.25	0.29	0.30	0.30	0.30	0.42
.30	0.36	0.39	0.39	0.39	0.60
.35	0.44	0.48	0.48	0.48	0.85
.40	0.53	0.60	0.60	0.60	1.22
.45	0.63	0.77	0.77	0.77	1.74
.50	0.75	1.00	1.00	1.01	2.46
.55	0.89	1.35	1.35	1.36	3.51
.60	1.05	1.88	1.89	1.90	5.12
.65	1.25	2.77	2.76	2.79	7.62
Mean		2.754	2.788	2.746	
StdDev		0.038	0.020	0.037	

Table 2: Mean population as a function of ρ.

Modeling Real Traffic

Two publicly available traffic archives were studied

- BCR: the BC-pAug89 LAN traffic archive
- DEC: the TCP portion of the wide-area dec-pkt-1 trace
- Both span roughly one hour of real time
- Both contain sub-millisecond level arrival times.

Parameter	BCR	DEC
μ	318	598
σ	114	193
σ/μ	0.36	0.32
H	0.845	0.84

Table 3: Workload parameters (arrivals/sec).

Modeling Real Traffic

Observation:

- Aggregation results are consistent with self-similarity

Resolution	Mean	True Var	Pred Var	Pred::True
0.125	39.77	477		
0.25	79.55	1526	1537	1.00
0.5	159.08	4348	4925	1.13
1.0	318.17	13112	14029	1.06
2.0	636.34	42341	42308	0.99
4.0	1272.71	135645	136617	1.00

Table 4: BCR aggregation variances: measured and expected.

Important observation:

- DEC resembles a fractional aggregation of BCR
- $\mu_{DEC}/\mu_{BCR} = 1.88$
- ${\sigma_{BCR}}^2 = 12996$
- ${\sigma_{BCR}}^{2^{2H}} = 12996 \times 1.88^{2\times 0.845} = 37796$
- ${\sigma_{DEC}}^2 = 37604$

Modeling Real Traffic

(Predicting population and drop rate)

Four representations of BCR and DEC are simulated

- True: arrival times from the traces (deterministic)
- Count: arrivals per second derived from True
- Emp: sampling from the distribution of Count
- fGn: fGn matching (μ, σ, H) of Count

For Count, Emp, and fGn, synthetic interarrival times have the truncated exponential density $f(x) = e^{-x}/(1-e^{-1})$

Implications of the results

- True $\approx$ Count $=>$ interarrival time synthesis is not perturbing results.
- True $\approx$ fGn $=>$ self-similarity is the major factor.
- True $\approx$ emp $=>$ distribution is the major factor
- True $\not\approx$ (fGn *or* emp) $=>$ ignore either at your peril!

Modeling Real Traffic

(BCR Results)

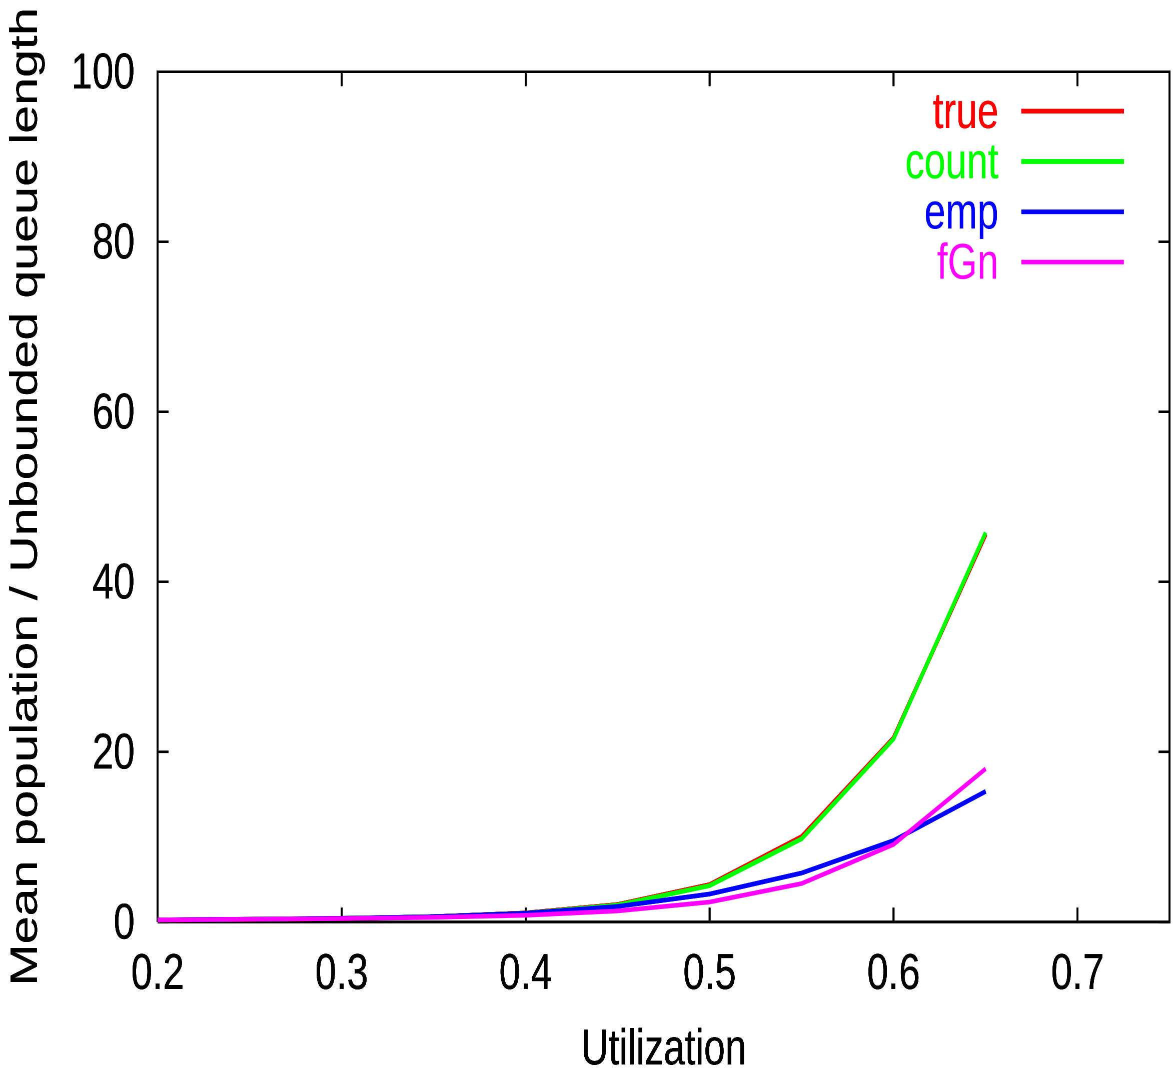

Analysis:

- “Count” workload matches “True” almost perfectly.
  $=>$ Interarrival time synthesis is satisfactory
- “Emp” and “fGn” match each other!
  $=>$ Neither captures essential characteristics

Modeling Real Traffic

(DEC Results)

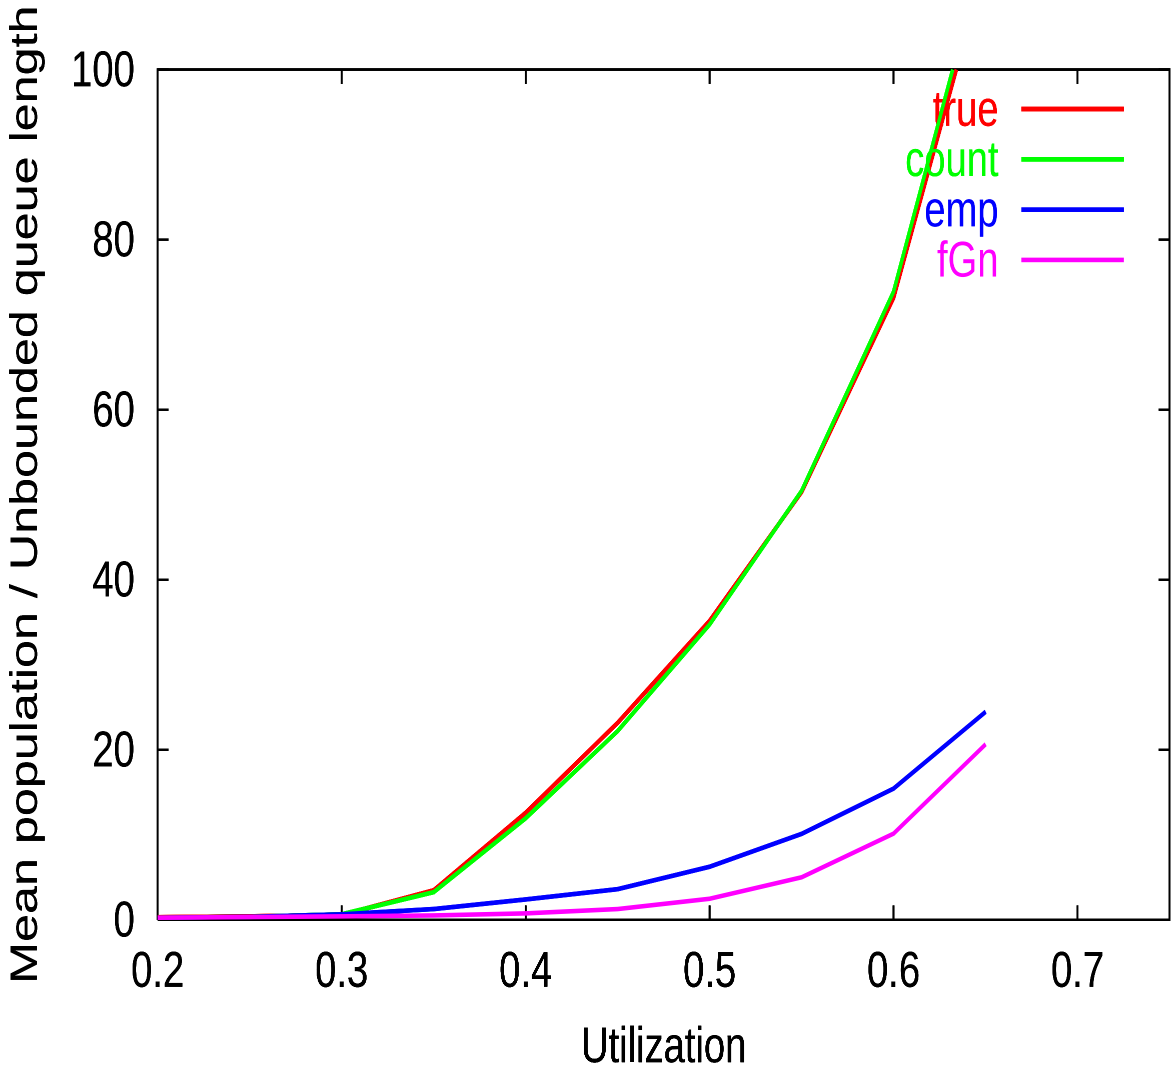

Analysis:

- “Count” again matches “True” almost perfectly.
- “Emp” does slightly better than “fGn”.
- “fGn(DEC)” matches “fGn(BCR)” almost perfectly.

Modeling Real Traffic

(BCR drops per 1000 packets)

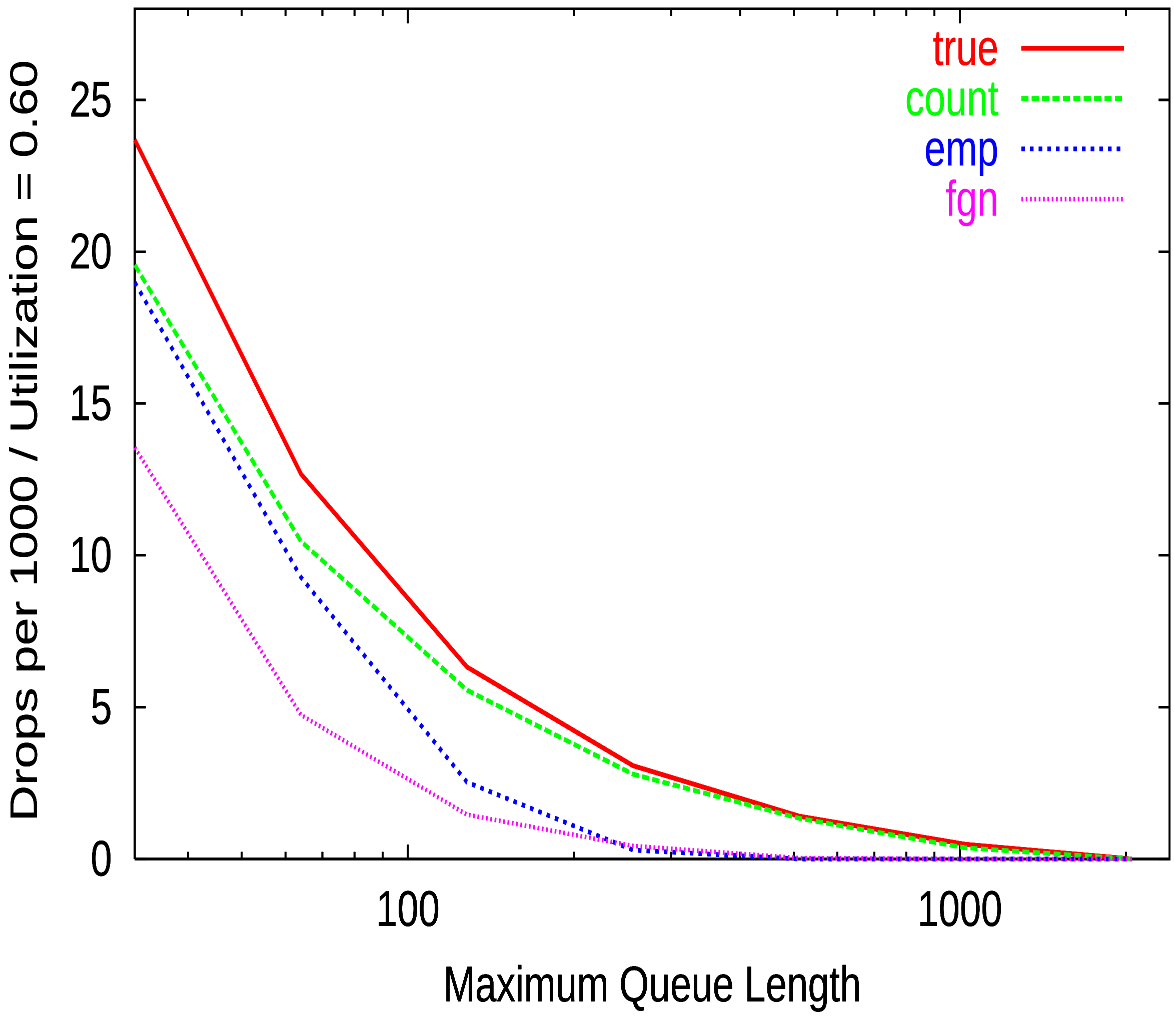

Analysis:

- Results are as expected for large capacity queues
- For small capacity queues “count” and “emp” match $=>$ This IS an artifact of interarrival time synthesis.

Modeling Real Traffic

(Empirical Distribution of BCR)

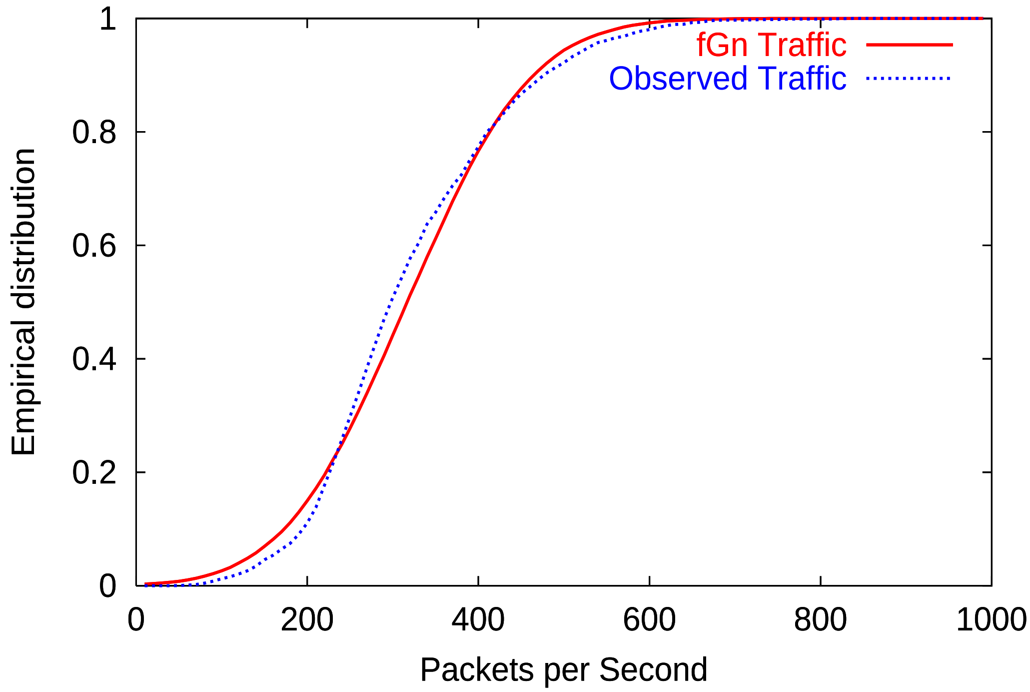

Analysis:

- The observed distribution is positively skewed
  - The median is less than the mean.
  - The left tail is shorter than the right.
  - Almost 60% of arrival counts are $<$ mean.
  - $=>$ The other 40% tend to be large.
- fGn is symmetric and thus unable to capture this effect.

Modeling Real Traffic

(Empirical Distribution of DEC)

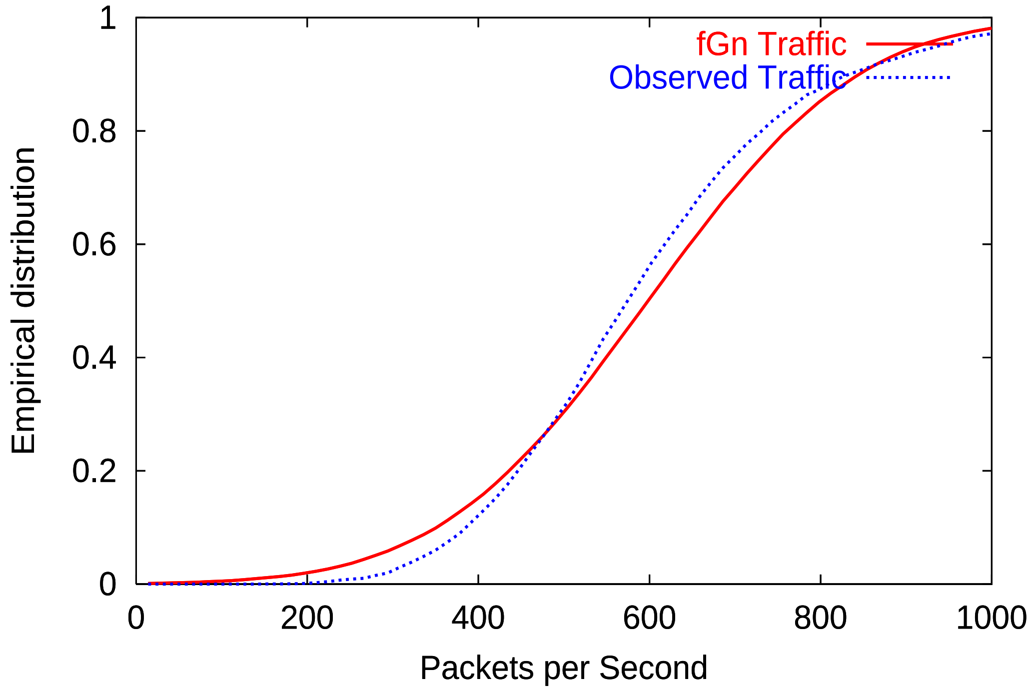

Analysis:

- Even more skewed than BCR

Conclusions

Positively skewed distributions are common.

Long-range dependence is also common.

Therefore:

- Sampling from empirical distributions and
- Using properly parameterized fGn

can both lead to significant underestimation of queuing delay and packet loss rate. However, the effect of the actual method used to generate the fGn is secondary and negligible.

Possible alternatives include:

- Exponentiation of fGn (the log-Normal distribution)
- Sample from the empirical distribution in a long-range dependent way. (IPDS - 2000)