SMALL AGGREGATIONS OF *ON/OFF* TRAFFIC SOURCES

R. Geist, J. Martin and J. Westall
Department of Computer Science
Clemson University
Clemson, SC 29634–0974
email: rmg, jmarty, westall@cs.clemson.edu

ABSTRACT
A simple and effective technique for synthesis of network traffic is evaluated, both in simulation and in a real network testbed. The technique, aggregation of a small number of *on/off* sources with heavy-tailed state holding times is seen to provide traffic that has characteristics consistent with long-range dependence and offers a convenient method for fitting mean, variance, and Hurst parameter of target workloads.

KEY WORDS
Traffic models, self-similarity, *on/off* sources.

1 Introduction

Several studies [1, 2, 7] have argued that real world network traffic can be reasonably well-modeled as an aggregation of *on/off* processes in which the holding times in the *on* and/or *off* states have heavy-tailed distributions. These models are commonly characterized as follows:

- Each source generates 1 packet per unit time in the *on* state.
- Each source generates no packets in the off state.
- State holding times are i.i.d. with means μ_{on} and μ_{off}.
- State holding times have heavy-tailed distributions in which $Prob(T_{on/off} < t) \approx 1 - (\beta_{on/off}/t)^{\alpha} \quad t > \beta_{on/off}$ and $1 < \alpha < 2$.

A powerful theorem of Taqqu, Willinger, and Sherman [9] established the connection between aggregated *on/off* sources and self-similarity of the aggregate traffic per unit time. The general theorem requires neither $\alpha_{on} = \alpha_{off}$ nor that state holding times have the Pareto distribution but, for simplicity, those are assumed in this discussion. Let $W_{M,T}(j)$ = the number of packets generated by M sources in the j^{th} interval of length T packet times. As $M \to \infty$ and $T \to \infty$ the aggregate process, $W_{M,T}(j)$, converges to:

$$TM\frac{\mu_{on}}{\mu_{on} + \mu_{off}} + T^H\sqrt{M}\sigma_{lim}FGN_H(j). \qquad (1)$$

The first term represents the steady-state component of the load. The time-varying component is weighted fractional Gaussian noise. The T^H factor arises because the standard deviation of any exactly self-similar process with Hurst parameter H aggregated over an interval of length T increases by a factor of T^H. The $\sqrt{M}$ term reflects the fact that the standard deviation of the sum of M i.i.d. processes having standard deviation σ is $\sqrt{M}\sigma$.

The impact of α and β on the standard deviation of the aggregate is carried by the σ_{lim} factor:

$$\sigma^2_{lim} = \frac{2(\mu^2_{off}\beta^{\alpha}_{on} + \mu^2_{on}\beta^{\alpha}_{off})}{(\mu_{on} + \mu_{off})^3(\alpha - 1)(2 - \alpha)(3 - \alpha)}. \qquad (2)$$

$FGN_H(j)$ is normalized ($\mu = 0, \sigma = 1$) fractional Gaussian noise with Hurst parameter

$$H = \frac{3 - \alpha}{2}.$$

When the holding times in the *on/off* states have the Pareto distribution, $Prob[T < t] = 1 - (\beta/t)^{\alpha}$, the mean holding times are given by

$$\mu_k = \frac{\beta_k \alpha}{\alpha - 1}, \quad k \in \{on, off\},$$

and the minimum holding time in state k is β_k.

The results of some empirical studies of large numbers of aggregated sources are presented in a related paper of Taqqu, Willinger, Sherman, and Wilson [11]. Nevertheless, no systematic study of the effect of varying M, T, and the length of the time series $W_{M,T}(j)$ was reported.

Wells [10], using simulation, carried out a systematic study for $\alpha_{on} = \alpha_{off}$ and $\mu_{on} = \mu_{off} = 1.0$. He showed that when $M >= 32$, $T >= 64$ and $W_{M,T}(j)$ consists of at least 10^5 observations, the sample autocorrelation of $W_{M,T}(j)$ is statistically consistent with the autocorrelation of a self-similar process having Hurst parameter $H = (3 - \alpha)/2$. He also showed that for $0.6 < H < 0.9$ [1] the observed variances were also consistent with those predicted by the theory. Wells did not extensively study heterogeneous state holding times extensively, and the results he obtained were inconclusive.

The fraction of time that a source spends in the *on* state is $\mu_{on}/(\mu_{on} + \mu_{off})$. For the aggregate load to be constrained to a level that is sustainable in a real network it is

[1] There are singularities in the predicted variance as $\alpha \to 1$ *and* $\alpha \to 2$.

T	μ	$\sigma(\mu)$	σ	$\sigma(\sigma)$	*skew*	$\sigma(skew)$
200	400.18	0.8145	124.65	0.4478	0.310	0.0000
400	800.35	1.6261	218.19	1.0361	0.270	0.0000
800	1600.70	3.2513	378.85	2.4162	0.240	0.0000
1600	3201.40	6.5083	652.63	5.6333	0.202	0.0042
3200	6402.80	13.0144	1116.73	13.1196	0.173	0.0048
6400	12805.59	26.0283	1900.06	31.1833	0.150	0.0000
12800	25611.19	52.0568	3219.86	73.9831	0.124	0.0052

Table 1. Observed statistics for 20 sources

necessary that $\mu_{on} << \mu_{off}$. Therefore, an understanding of the behavior of such distributions is important.

In the remainder of this paper we examine the distributional and correlational characteristics of small aggregations of *on/off* sources with heterogeneous state holding times. In section 2, using discrete event simulation, we show that, both the first order (distributional) and second order (correlational) statistics of the simulated processes converge rapidly to those of the limiting distribution. In section 3, our focus changes from simulation to real network traffic. Using a dedicated network testbed, we show that, under low network loads, the characteristics of TCP traffic are consistent with a sample from the distributions observed with simulation. We then show that as load is increased the distributional and correlational characteristics of the aggregate traffic process change in ways that might not be expected. The conclusion follows in section 4.

2 Simulated *on/off* sources

In this section we describe the distributional and correlational characteristics of small aggregations of *on/off* processes. These results are obtained via discrete event simulation because the problem is analytically intractable and run lengths in excess of one billion simulated packet times are required to obtain repeatable values of the sample autocorrelations for lags in the range 1-256.

2.1 Simulation methods and parameters

For the results reported here, each simulated packet source is an *on/off* process with holding times having the Pareto distribution with $\alpha_{on} = \alpha_{off} = 1.5$, $\beta_{on} = 4$, and $\beta_{off} = 36$. Thus, $\mu_{on} = 12$ packet times, and $\mu_{off} = 108$ packet times. The number of sources simulated is $M = \{10, 20, 40, 80\}$. For each number of sources simulated, ten independent replications of length 1,440,000,000 packet times are conducted. In each simulation, whenever the total number of sources in the *on* state changes, the simulator writes a record of the form $\{t_i, n_i\}$ where t_i is the simulated time and n_i is the number of sources now in the *on* state.

The 40 files created by the simulation are then reprocessed creating new files of integer valued time series in which the values, $\{W_{M,T}(j)\}$, represent the number of packets generated by the M sources in the j^{th} interval of length T packet times. Aggregations of length $T = \{200, 400, 1600, 3200, 6400, 12800\}$ are used, yielding a total of 240 of these time series.

Each of these 240 time series is then processed by a statistical analyzer which reports the sample mean, standard deviation, coefficient of skewness, and the first 256 lags of the sample autocorrelation. Because of the ten independent replications of the original simulations, there is a set of ten of these statistics for each unique $\{M, T\}$ pair. In the final step, the sample means and standard deviations of each of the 24 sets of the ten observations are computed.

$T \backslash M$	10	20	40	80
200	0.5236	0.5241	0.5252	0.5243
400	0.5448	0.5455	0.5466	0.5455
800	0.5623	0.5632	0.5646	0.5632
1600	0.5757	0.5768	0.5788	0.5769
3200	0.5854	0.5869	0.5896	0.5869
6400	0.5915	0.5938	0.5977	0.5937
12800	0.5949	0.5983	0.6036	0.5980

Table 2. Observed $\hat{\sigma}_{lim}$ by aggregation level

$T \backslash M$	10	20	40	80
200	0.440	0.310	0.220	0.160
400	0.386	0.270	0.190	0.142
800	0.333	0.240	0.170	0.120
1600	0.289	0.202	0.142	0.101
3200	0.246	0.173	0.120	0.090
6400	0.210	0.150	0.108	0.073
12800	0.177	0.124	0.090	0.060

Table 3. Observed $skewness$ by aggregation level

2.2 Distributional characteristics

Characteristics of the empirical distributions for M equal 20 sources are shown in Table 1. Each row of the table represents a different aggregation level, T. The column labeled μ contains the mean number of simulated packets created per T packet times over the 10 replications. Since each simulated source is active 10% of the time, 20 sources are expected to generate 0.1 x 200 x 20 = 400 packets when T is 200 and, as the aggregation level doubles, so does the expected number of packets generated. The column labeled $\sigma(\mu)$ is the sample standard deviation of the 10 observations of μ. The coefficient of variation, $\sigma(\mu)/\mu$, is consistent at 0.002 over all aggregation levels.

We now turn our attention to the sample standard deviations of the simulated processes. As $M, T \to \infty$, it has been shown that $\sigma \to T^H \sqrt{M} \sigma_{lim}$[9]. Thus, if this convergence is sufficiently fast that $20, 200 \approx \infty$, then, since $\sigma_{lim} = 0.62$ for $\alpha = 1.5$, $\beta_{on} = 4$, and $\beta_{off} = 36$, it should be the case that at aggregation level 200, $\sigma = 200^{0.75} \times \sqrt{20} \times 0.62 = 147.4$. The fact that observed value, 125, is approximately 15% too small is an indicator that the aggregation interval $T = 200$ is not sufficiently larger than $\mu_{off} = 108$. Under these conditions it is common for a single holding time in the *off* state to encompass multiple aggregation intervals thus reducing the variability of the aggregated process.

A useful way to compare standard deviations across all numbers of sources and aggregation levels is to factor out the dependence upon M and T by considering the sample value of $\hat{\sigma}_{lim} = \sigma/(T^H \sqrt{M})$. These values are shown in Table 2. All samples of $\hat{\sigma}_{lim}$ are smaller than the limiting value 0.62, but are converging toward it as T increases. For $T = 12800$, observed values are only 4% below the limiting value. Surprisingly, for all values of T measured, no dependence of statistical significance on M exists.

The last two columns of Table 1 characterize skewness of the sample distributions. Since the limiting distribution is fractional Gaussian noise, as $M, T \to \infty$, it is necessary that the sample coefficient of skewness, $\frac{1}{n\sigma^3} \sum_{j=1}^{n} (x_j - \mu)^3 \to 0$. It can be seen in Table 3 that the skewness decreases as both M and T increase. In practice, positive skewness is pervasive in real-world packet counts [4], and thus the positive skew produced by 10 or 20 sources may actually be more realistic than that produced by 80 sources.

2.3 Correlational characteristics

In this section we consider the stochastic dependence exhibited by the simulated traffic sources. Although this dependence is sometimes reduced to a single number, the Hurst parameter, we feel that the sample autocorrelation function yields more insight. For a finite sample of N observations

$$\hat{\rho}(k) = \frac{1}{N-k} \sum_{j=1}^{N-k} (X_j - \hat{\mu})(X_{j+k} - \hat{\mu})/\hat{\sigma}^2, \quad (3)$$

where $\hat{\mu}$ and $\hat{\sigma}^2$ represent the sample mean and variance, is an asymptotically unbiased estimator of the autocorrelation, $\rho(k)$ [5]. Since the autocorrelation of an exactly

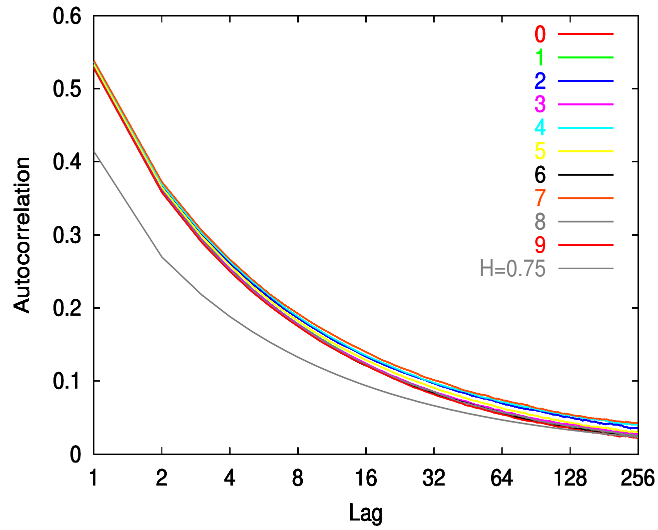

Figure 1. Sample autocorrelation: M = 20, T = 200

self-similar process with Hurst parameter H is given by[8]

$$\rho(k) = (1/2)((k+1)^{2H} - 2k^{2H} + (k-1)^{2H}),$$

it is easy to compare the stochastic dependence obtained in the samples with the theoretical limiting behavior. In the graphs that follow, the line labeled $H = 0.75$ is the autocorrelation of an exactly self-similar process with $H = 0.75$.

The sample autocorrelations of all ten replications for $M = 20$ and $T = 200$ are shown in Figure 1. Note that the x-axis scale is logarithmic and that continuous lines are used to represent the discrete autocorrelation. It can be seen in this graph that the variablilty in the values obtained in the 10 independent replications grows with increasing lag. The maximum difference is 0.018 at lag 256. For small lags the correlations are significantly greater than would be expected for an exactly self-similar process. This is another consequence of T being sufficiently small that *on/off* epochs of a given source regularly span multiple, adjacent aggregation intervals.

Figure 2 shows the effect of extending the length of the aggregation period, T. Each point plotted here corresponds to the mean of the 10 values plotted in Figure 1. It can be seen that for small lags, as the value of T increases the sample autocorrelations converge to the target $\rho(k)$. For large lags and values of T, the sample autocorrelation can be seen to cross the target line. By running the simulations for even larger numbers of packet times we have found that for large lags and values of T, the sample autocorrelations increase with the length of the run, and thus the values shown do not represent the characteristics of the underlying distribution.

It was previously observed that, for the number of sources considered here, the sample standard deviation of the number of simulated packets generated per unit time is

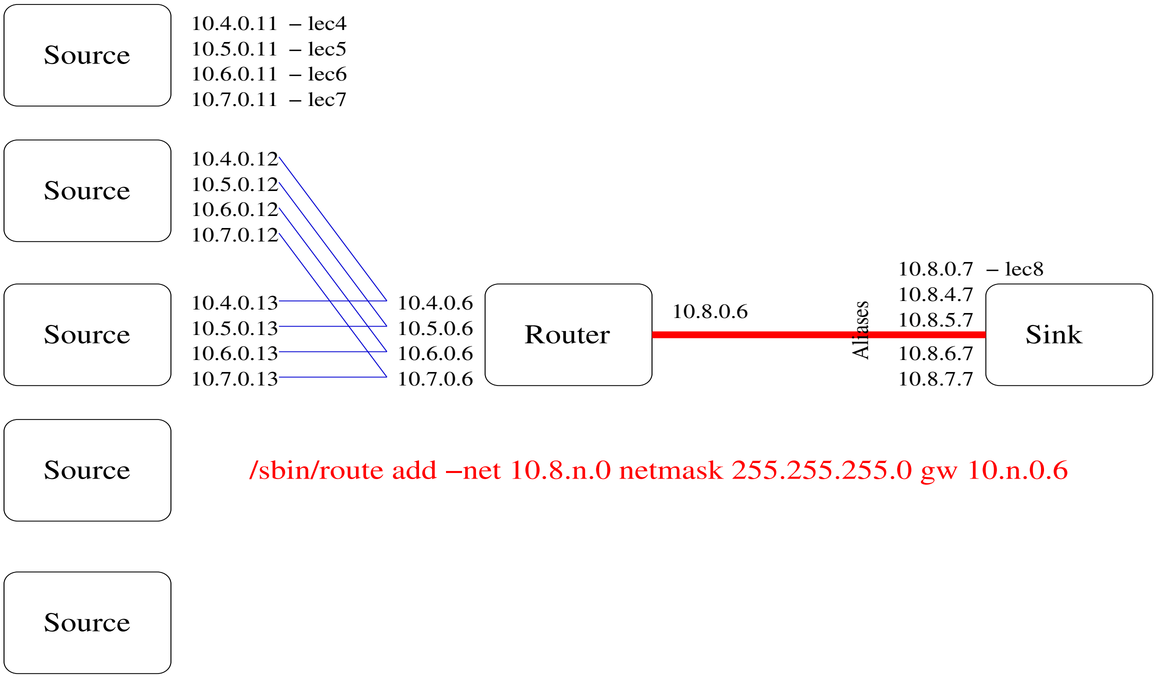

Figure 4. Network Configuration

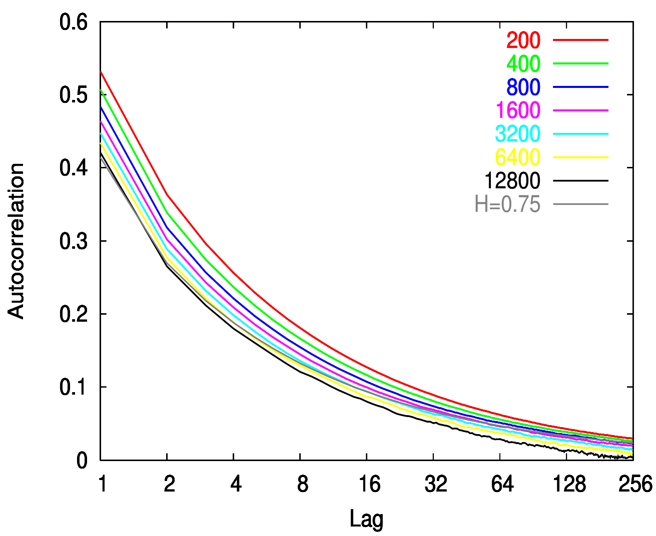

Figure 2. Sample autocorrelation: M = 20

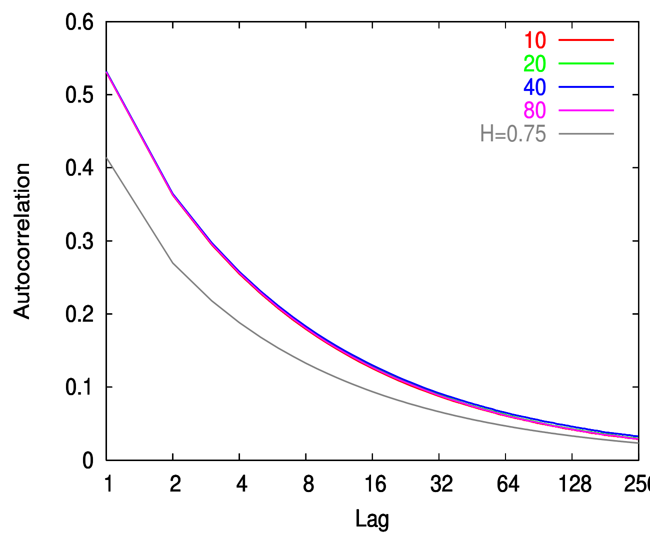

Figure 3. Mean sample autocorrelation: M = {10,20,40,80}, T = 200

not dependent on the number of sources in a statistically significant way. Figure 3 shows that this also applies to the sample autocorrelation. Here the lines representing the mean sample autocorrelations for $T = 200$ and $M = 10$, 20, 40, and 80 sources are virtually indistinguishable.

2.4 Three parameter fitting

These results indicate that using a relatively small number of sources M, it is generally possible to identify values of α, β_{on}, and β_{off}, that produce useful fits for a target network load parsimoniously parameterized as (μ, σ, H). Here μ and σ refer to the number of packet transmissions per some unit time and H characterizes the long-range dependence.

Fitting can be done as follows. It was previously shown that $\alpha = 3 - 2H$. Let T the number of packets per unit time that each *individual* source generates in the *on* state and M is the number of sources. Since

β_{on}	β_{off}	T	σ
2	3	50	68.193
2	8	100	98.257
2	18	200	123.935
2	38	400	149.387
2	78	800	177.115
4	6	50	81.095
4	16	100	116.847
4	36	200	147.385
4	76	400	177.653
4	156	800	210.627

Table 4. Standard Deviations

Router			Sink		
μ	σ	$skew$	μ	σ	$skew$
393.73	124.19	0.32	393.72	123.99	0.31
393.68	123.60	0.31	393.59	122.77	0.31
402.94	129.24	0.32	402.25	129.74	0.32
400.94	127.15	0.32	400.74	126.97	0.32
402.80	127.86	0.32	402.75	128.03	0.32
402.63	127.33	0.32	401.59	126.52	0.32
395.96	122.80	0.31	395.89	123.60	0.31
405.98	135.11	0.33	405.75	135.55	0.33
408.72	133.82	0.33	408.68	133.61	0.33
420.89	137.12	0.33	419.69	137.27	0.33

Table 5. TCP / 2 hour samples / 24Mbps bottleneck

$\alpha = \alpha_{on} = \alpha_{off}$, $\beta_{on/off} \propto \mu_{on/off}$ and the aggregate packet rate is necessarily

$$\mu = MT\beta_{on}/(\beta_{on} + \beta_{off}).$$

Approximate integral solutions to this equation may be readily computed and enumerated as in Table 4. This table shows some feasible solutions and expected standard deviations for $H = 0.75$, 20 sources, and $\mu = 400$ packets per unit time. Fitting the target σ is via table lookup. It can be seen from the table that the value of the standard deviation is strongly correlated to the magnitude of T and β_{off} with the value of β_{on} being available as a tuning parameter.

3 Application to real network traffic

In this section we present an example in which the aggregation of a small number of *on/off* TCP sources is used to evaluate load-dependent effects on the distributional and correlational characteristics of the traffic process. The network testbed is an IP-over-ATM (LAN Emulation) network configured as shown in Figure 4. The seven network nodes shown are Linux PCs equipped with 155 Mbps Interphase 5515 NICs and running kernel 2.4.18. The systems are attached to IBM 8265 and IBM 8285 ATM switches with switch-to-switch links also running at 155 Mbps.

The Interphase 5515 device driver, developed at Clemson University, includes facilities for capture and logging of performance data and for the assignment of QoS attributes to VCC's carrying IP traffic. Each time a packet is received, queued for transmission, or completes transmission, a record containing the time, source IP address, destination IP address and current length of the transmit queue is logged. The QoS facility allows an administrator to bind a local IP address to one of the 5515's eight hardware Tx queues and to specify the outgoing ATM bit rate and the Tx buffer queue size. This QoS facility is used to provision the bottleneck link from Router to Sink and performance data is captured on both Router and Sink.

3.1 Workload characteristics

Traffic is generated by a custom TCP application. In the *on* state, the application sends full-sized segments at a specified rate. The packet rate in the *on* state, although bounded by the outgoing link speed, is *not* directly coupled to it. The application uses the 10 msec resolution timing services of Linux to clock segments out at approximately the specified rate.

In the study described here $M = 20$ sources. Each of the five source hosts runs four instances of the source application. As shown in Figure 4 each source has a dedicated ATM SVC to the router. These SVCs are provisioned at 10Mbps. Each source generates 200 packets per second in the *on* state. Thus, an aggregation factor of $T = 200$ corresponds to one second of real time. State holding times are Pareto with $\alpha = 1.5, \beta_{on} = 4, and \beta_{off} = 36$ as in the simulations, and the expected aggregate packet rate is 400 packets / second. Including network, link, AAL5, and ATM overhead each packet requires 1696 bytes (32 cells). The sustained load is 5.4Mbps. Using σ_{lim} from equation 2 yields a predicted standard deviation of 147, but, based upon the results of the simulation studies one would expect a value closer to 125. These parameters were selected to produce traffic whose distributional and correlational characteristics were similar to the widely studied Bellcore network traffic traces[6, 3].

3.2 Performance analysis

Each run consists of ten consecutive two hour tests. Since each source generates 200 packets/sec in the *on* state, two hours corresponds to 1,440,000 packet times (1,000 times shorter than the simulations.) As a baseline run the bit rate on the bottleneck link was set to 24Mbps. The distributional properties of the traffic arrival processes at Router and Sink are shown in table 5. Because of the relatively short run-lengths, the results are clearly more variable than were the simulations, but at a 23% utilization of the bottleneck link, there is no compelling evidence that TCP dynamics have altered the arrival process at the router or the

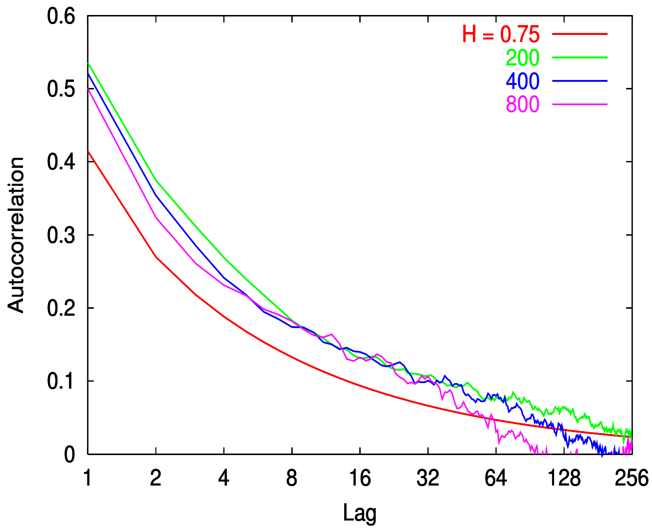

Figure 5. Sample autocorrelation: Bottleneck = 24Mpbs

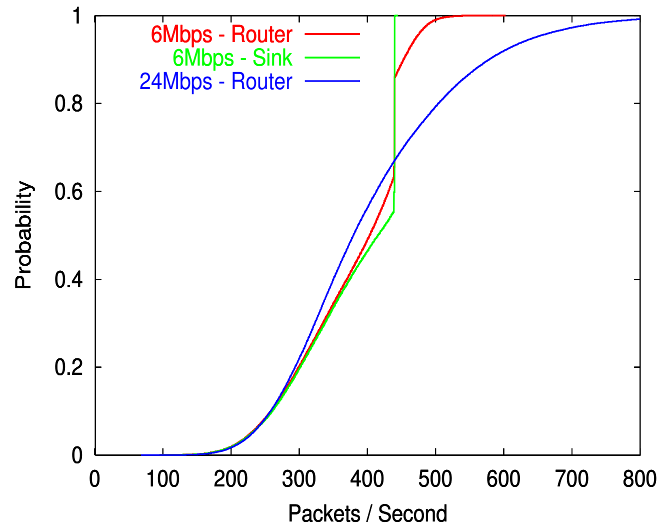

Figure 6. Sample CDFs

Router			Sink		
μ	σ	$skew$	μ	σ	$skew$
378.76	80.27	0.21	377.75	73.20	0.19
375.78	81.03	0.22	374.95	74.30	0.20
377.40	82.01	0.22	376.50	74.92	0.20
371.75	83.23	0.22	370.85	76.10	0.21
375.76	83.33	0.22	375.14	75.92	0.20
379.03	83.15	0.22	377.58	75.93	0.20
386.70	80.90	0.21	384.82	73.67	0.19
379.19	83.10	0.22	378.52	75.89	0.20
368.49	84.89	0.23	366.93	78.62	0.21
386.43	80.58	0.21	385.63	72.83	0.19

Table 6. TCP / 2 hour samples / 6Mbps bottleneck

sink from the distribution predicted by the simulation. The mean sample autocorrelation of the arrival process is shown in Figure 5. For low lags, the values obtained are consistent with the values obtained with simulation, but for large lags, particularly for aggregation levels $T = 400$ and 800, the sample values become erratic because of the shortness of the runs.

To evaluate the effects of congestion, we maintain the sustained application level offered load at 400 packets/second and incrementally reduce the bit rate on the bottleneck link to 6 Mbps. As the bottleneck bandwidth drops, two effects combine to significantly alter the distributional characteristics of the arrival process at both Router and Sink. The culmination of these effects can be seen in Table 6 and Figure 6. Sustained throughput drops only 5% to 380 packets/second which corresponds to 85% utilization of bottleneck link. The packet loss rate, $(\mu_{router} - \mu_{sink})$, remains negligible at about 1 packet/sec because the Tx buffer quota at Router is set to 256KB[2].

[2]Reducing the quota to 64KB increases the drop rate to approximately 12 packets/sec but does not significantly alter σ or the sample autocorrelation.

The major impact is upon σ which drops from 125 in the uncongested case to 82 at Router, and 74 at Sink. The reasons for this decrease are evident in Figure 6. The graphs show the sample cumulative distribution functions of packet arrivals per second. The distribution at Sink is necessarily clipped at 442 packets/sec because of the leaky bucket effect imposed by the 6Mbps link. This constrains the rate at which *acks* are fed back to the source and leads to the similar shape of the arrival distribution at Router.

Sample autocorrelations of the arrival processes under low and heavy load and aggregation level $T = 200$ are shown in Figure 7. Surprisingly, the load dependent effects are minimal. Correlation actually increases with load for lags between 1 and 32 and marginally decreases for higher lags. At lags greater than 32 the sample autocorrelation of the arrival processes at Router and Sink are indistinguishable.

The arrival process at Sink provides an excellent reminder that long-range dependence in an arrival process per se does not imply performance problems in a queuing system [4]. This arrival process could drive an outgoing link at 85% utilization with no queuing at all! It is only when long-range dependence is coupled with high variability that trouble arises.

4 Conclusion

We have shown useful three parameter fits of target network traffic loads may be produced using small aggregations of *on/off* sources. It might be argued that the traffic produced is not as "realistic" as traffic produced using Mah's model [7] or the *Surge* system of Barford and Crovella [1]. However, it can also be argued that fitting a target (μ, σ, H) with

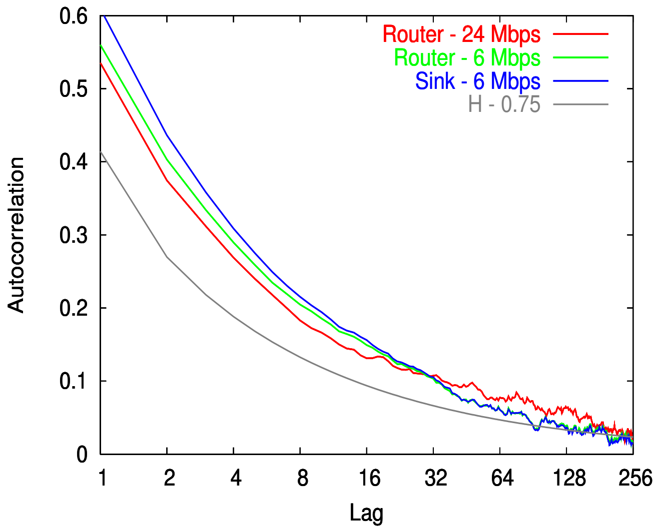

Figure 7. Sample autocorrelation: $T = 200$

these more realistic loads is a extremely difficult task!

Thus, small aggregations of the type described here should viewed as augmenting, not competing with, tools such *Surge*. As shown in section 3, they can provide useful insights into load dependent effects on traffic arrival processes. They are also ideal for performing systematic studies of the effect of varying a workload parameter such as σ or H over a range of values of interest.

References

[1] Paul Barford and Mark Crovella. Generating representative web workloads for network and server performance evaluation. In *Proceedings of Performance '98 / ACM SIGMETRICS '98*, pages 151–160, Madison, WI, July 1998.

[2] M. Christiansen, K. Jeffay, D. Ott, and F. Smith. Tuning red for web traffic. In *Proc. of ACM SIGCOMM '00*, pages 139–150, Stockholm, Sweeden, August 2000.

[3] R. Geist and J. Westall. Practical aspects of simulating systems having arrival processes with long-range dependence. In *Proc. of the 2000 Winter Simulation Conf. (WSC '00)*, pages 666–674, Orlando, FL, December 2000.

[4] R. Geist and J. Westall. Correlational and distributional effects in network traffic models: Extended version. *Performance Evaluation*, 44(1):121–138, April 2001.

[5] A. Kobayashi. *Modeling and Analysis*. Addison-Wesley, Reading, MA, 1978.

[6] W. Leland, M. Taqqu, W. Willinger, and D. Wilson. On the self-similar nature of Ethernet traffic. In Deepinder P. Sidhu, editor, *ACM SIGCOMM*, pages 183–193, San Francisco, California, 1993.

[7] Bruce A. Mah. An empirical model of HTTP network traffic. In *INFOCOM (2)*, pages 592–600, 1997.

[8] V. Paxson. Fast approximation of self-similar network traffic. *ACM SIGCOMM Computer Communication Review*, 27(5):5–18, 1997.

[9] M. Taqqu, W. Willinger, and R. Sherman. Proof of a fundamental result in self-similar traffic modeling. *Computer Communication Review*, 27(2):5–23, April 1997.

[10] P. Wells. Simulation of self-similar network traffic using high variance ON/OFF sources. *M.S. Thesis*, May 2002.

[11] W. Willinger, M. Taqqu, R. Sherman, and D. Wilson. Self-similarity through high-variability: Statistical analysis of Ethernet LAN traffic at the source level. *IEEE/ACM Transactions on Networking*, 5(1):71–86, February 1997.