

Traffic Engineering for VBR Video with Long-Range Dependence

K.R. Krishnan
Bellcore
MCC-1C-353R
445 South Street
Morristown, NJ 07960-6438, USA
Tel: 201-829-2891
Fax: 201-829-2645
e-mail: krk@bellcore.com

Gopalakrishnan Meempat
Bellcore
NVC-3X-311
331 Newman Springs Road
Red Bank, NJ 07701-5699, USA
Tel: 908-758-4034
Fax: 908-758-4371
e-mail: gopal@bellcore.com

ABSTRACT

Recent studies at Bellcore have shown that traffic rates in high-speed data networks exhibit 'long-range dependence' in their correlation, characterized by the so-called Hurst parameter H assuming values in the range $0.5 < H < 1$. This corresponds to correlation that has a smaller asymptotic rate of decay than in Markovian processes, for which $H=0.5$. Simulations suggest that link-engineering based on Markovian traffic models could, in some situations, underestimate the bandwidth requirements of the actual traffic. However, another Bellcore study presents results in which a Markovian traffic model has been successful in determining bandwidth requirements for Variable-Bit-Rate video sources, which have $H > 0.5$.

We report here the results of an investigation undertaken to bridge the gap in our understanding of the results of these two studies, and to analyze the effect of the Hurst parameter on link-engineering for ATM traffic. Our investigation has shown that a curious 'scaling' property of streams with long-range dependence, derived in another recent Bellcore study, helps to bridge the gap by accounting for the results of both studies. The analysis suggests that a high value of H is not, in itself, a reason to suppose that Markovian traffic models will lead to under-engineering of bandwidth on ATM links.

Keywords

ATM, long-range dependence, multiplexing, effective bandwidth

1. Introduction

The starting point of the investigation reported here was the need to bridge a gap in our understanding of the results reported from two studies, Heyman et al (1992) and Leland et al (1993), on traffic streams that exhibit the property of 'long-range dependence' in the behavior of the auto-correlation of their instantaneous arrival rates.

On the one hand, the study by Leland, Taqqu, Willinger, and Wilson (1993) established, on the basis of extensive traffic measurements, that traffic streams in high-speed networks exhibit 'long-range' dependence in their auto-correlation, characterized by an auto-correlation with a slower asymptotic decay than in standard Markovian traffic models. This property of 'persistent' correlation was described by the *Hurst* parameter H of the stream, with $H = 0.5$ for Markovian streams and $H > 0.5$ for streams with long-range dependence. It was further shown, both by analysis [Norros (1994)] and simulation [Erramilli et al (1994a)], that engineering formulae (such as effective bandwidth for connections) based on Markovian models could, in *some cases*, result in under-engineering.

On the other hand, in studies of the bandwidth needed on an ATM link for carrying Variable Bit Rate (VBR) video-traffic, Heyman, Tabatabai, and Lakshman (1992) and Elwalid, Heyman, Lakshman, and Mitra (1995) have shown that, *even though the Hurst parameter of the VBR stream has been determined to be about 0.7* [Beran et al (1995)] (and thus corresponds to the case of long-range dependence), effective bandwidth formulae derived for Markovian models are, nevertheless, successful in determining the bandwidth required to support the VBR streams.

Clearly, a better understanding of the significance of the Hurst parameter is needed in order to be able to explain and account for the results of both studies. In the course of our investigation of ATM link-

engineering for traffic streams with long-range dependence, we have shown that a curious 'scaling' property of streams with long-range dependence, derived in recent work by Krishnan (1995), helps to bridge the gap between the results of the two studies, and also suggests that a high value of H is not, in itself, a reason to suppose that Markovian traffic models will lead to under-engineering of bandwidth on ATM links.

2. Question to be Investigated

The study by Heyman et al (1992) makes use of an actual trace of the traffic generated by a single source during a video-conference call of 32.3 minutes duration. The original trace consists of 48500 frames, each frame being 40 milliseconds in duration. The traffic data consist of a sequence of 48500 integers, corresponding to the numbers of cells in the frames. The frame-size sequence is shown to correspond to a stationary sequence. The mean frame-size is 130.3 cells, the maximum frame-size is 629 cells, and the variance of frame-size is 5536.9. With a frame interval of 40 milliseconds and with 64 bytes per cell, the VBR stream has thus a mean rate of 1.668 Mb/s and a peak rate of 8.05 Mb/s. The marginal distribution of frame-size is shown to be well-approximated by a gamma distribution with a scale parameter of 0.02353 and a shape parameter of 3.066. In simulations to determine the performance obtained with the VBR video-source, we have made use of the subsequence consisting of the first 45000 frames of data, which corresponds to a 30-minute video-call.

Beran, Sherman, Taqqu, and Willinger (1995) have analyzed this VBR data stream and have concluded that it exhibits long-range dependence, with a Hurst parameter of $H=0.7$, considerably larger than the value 0.5 applicable for Markovian streams. However, Heyman et al (1992) have constructed a Markovian model, called the Discrete Auto-Regressive (DAR) model, that provides a good match both to the distribution of frame-size and to the auto-correlation between successive frame-sizes, extending over about 150 frame-lags.

Elwalid et al (1995) extend the effective bandwidth theory developed for Markovian traffic streams to the DAR model of the VBR stream, to determine the number of independent, simultaneous VBR sources that can be supported with a cell-loss rate of 1.0×10^{-6} by an ATM link with a bandwidth of 45 Mbits/sec. Event simulations are carried out using time-shifted versions of the actual 30-minute data stream to confirm that the bandwidth of 45 Mbits/sec can, in fact, support the number of sources determined on the basis of the Markovian DAR model of the VBR trace.

Thus, the example considered by Elwalid et al (1995) is at least one instance in which link-engineering based on Markovian models turns out to be adequate for traffic known to possess long-range dependence in its correlations. The question arises, therefore, whether there are situations in which effective bandwidth determined from Markovian models will be adequate for traffic with long-range dependence. The goal of this paper is to investigate this question and present the answers, though incomplete, that have emerged from our investigation.

3. Preliminary Matters

As noted earlier, the actual VBR data available to Heyman et al (1992) consists of the traffic generated by a single video-source of 45000 frames. In order to consider the combination of multiple *independent* streams, Heyman et al (1992) use phase-shifted versions of the original 45000-frame stream, (regarded as a cyclic sequence), to represent independent traffic streams. The phase-shifts themselves were chosen in a random manner over the span of 45000 frames; the idea behind this procedure is that the correlation between X_k , the size of frame k , and X_{k+m} , the size of frame $(k+m)$ becomes negligible for large enough values of the lag m , and thus the stream $\{X_k\}$ and the stream $\{X_{k+m}\}$ are effectively uncorrelated over a large span. In fact, a minimum value for the phase-separation could be selected by examining the decay of the correlation in the original data stream with the lag m , and requiring all phase-separations to be larger than a value of m for which the correlation becomes negligible. Generating independent

realizations of the DAR *representation* of the original trace, however, is quite straightforward, with the use of different non-overlapping sections taken from a long stream generated by a pseudo-random seed.

3.1 Cell-Loss Rate Criterion for Link-Engineering

It is common in papers on ATM to consider extremely low cell-loss rates, of the order of 1.0e-6 to 1.0e-9, as the performance criterion to be met on an ATM link. It is often difficult to ascertain, by means of simulation over a limited duration, whether such low, long-term loss-rates are being met or not. For example, the VBR traffic data for the single source considered here contains a total of 5,818,500 cells in its 45000 frames over its 30-minute duration. A loss-rate of 1.0e-6 would correspond to the loss of *just 6 cells* over the entire 30-minute duration. Since losses often occur in bursts, and not at a steady rate over the finite duration of interest, the position occupied by a single burst in the arrival stream over the 30-minute duration could make the difference between meeting or violating the grade-of-service. In particular, in such cases, there can be no reasonable confidence levels associated with the results of simulation.

Since the actual VBR traffic stream available to us is limited to the 30-minute trace, we have used a loss-rate of 1.0e-3 as the engineering criterion in some of the results presented below. This loss-rate corresponds to the loss of about 5800 cells per source over the 30-minute duration, which is about 45 times the average frame-size of the stream.

4. Effect of Long-Range Dependence on Performance

4.1 Effective Bandwidth and Multiplexing Gain for the Zero-Buffer Case

We use the following general definition for G , the multiplexing gain achieved when n independent streams are combined for transmission on a single transmission link:

First, consider a single traffic source, and let

R_p = the peak rate of the source,

R_m = the mean rate of the source,

Given a buffer-size B and a desired quality-of-service q (cell-loss rate, cell-delay and jitter, etc), let

C_s = link bandwidth needed to meet desired quality of service for the single source.

We know that $R_m \leq C_s \leq R_p$. The statistical variations in the transmitting rate of the traffic source may enable us to meet the desired quality-of-service with a link bandwidth smaller than the peak rate R_p . The statistical gain achieved for a single source with the use of C_s as the effective bandwidth is defined to be

$$G_s = R_p / C_s \tag{1}$$

We now consider multiplexing the realizations of n such independent sources for transmission on a single link. The independence in the statistical variations of the n individual sources offers the potential for achieving a further reduction in the bandwidth required for the combined stream. Let

\bar{C}_n = link bandwidth needed to meet desired quality of service for the multiplexed stream of n sources

Then, we define the total multiplexing gain G as follows:

$$\begin{aligned}
 G &= \frac{nR_p}{\bar{C}_n} & (2) \\
 &= \left[\frac{nC_s}{\bar{C}_n} \right] \left[\frac{R_p}{C_s} \right] \\
 &= \left[\frac{nC_s}{\bar{C}_n} \right] G_s \\
 &= [\text{multiplexing gain across sources}] [\text{statistical gain in a single source}]
 \end{aligned}$$

Thus, we see that the total multiplexing gain due to the independent combination of identical sources can be resolved into two factors, one expressing the efficiency of statistical multiplexing of independent streams (gain *across* sources) and the other expressing the advantage gained (by means of buffering) from the statistical rate variations *within* a source.

One can obtain a conservative estimate of the bandwidth needed for a traffic stream by supposing that delayed cells cannot be buffered (i.e. $B=0$), and thus insisting that the link-bandwidth \bar{C} be chosen such that the probability of the rate r of the stream ever exceeding \bar{C} is smaller than the probability of loss specified for the quality-of-service. Clearly, this calculation for a stationary process makes use only of the marginal distribution of the rate r , and is unaffected by the properties of the auto-correlation, and thus by the presence or absence of 'long-range dependence'. Heyman et al (1992) have shown that the distribution of frame-sizes in the video-stream (a measure of the average cell-arrival rate over the duration of a frame) is well-represented by a gamma distribution with a scale parameter $\lambda = 0.02353$ and a shape parameter $s = 3.066$. For a single source, the use of the above gamma distribution of frame-size gives a link-bandwidth $C = 10.5$ Mb/sec for the zero-buffer case, when the probability of cell-loss is not to exceed 1.0×10^{-6} . Of course, since the single source has the *known* peak-rate of 8.05 Mb/sec, the assumed marginal distribution is not quite accurate for this case, and we use just the peak-rate. However, when one considers the multiplexing of 15 independent streams, as in Heyman et al (1992), the assumption of the gamma distribution for the single source enables us to conclude that the frame-size distribution for the multiplexed stream is again a gamma distribution, with a scale parameter of λ and a shape parameter of $15s$. For a loss-probability of 1.0×10^{-6} , the link-bandwidth for the zero-buffer case is easily determined to be 46.6 Mb/sec, which is much smaller than 15 times 8.05 Mb/sec (the bandwidth determined for a single source). In fact, the multiplexing gain G for this case is $15 \times 8.05 / 46.6 = 2.59$. Thus, even when the effective bandwidth is based entirely on the statistics of the arrival rate, with no consideration of buffers, and *even when it is necessary to allocate the peak rate for a single source* (as in this example), it is still possible to realize significant multiplexing gain when multiple *independent* sources are combined.

The effective bandwidth of 46.6 Mb/sec for the zero-buffer case for 15 sources corresponds quite closely to the result obtained by Elwalid et al (1995) that 15 sources can be supported with a link-bandwidth of 45 Mb/sec with a buffer of 5000 cells (which is about 39 times the mean frame-size). It thus appears that, for the VBR streams studied by Heyman et al (1992), for small and moderate buffer-sizes, *the multiplexing gain is essentially achieved 'across' sources, rather than from the statistical rate variations within a source*. For the zero-buffer case, this component of multiplexing gain is essentially determined by the marginal distribution of the cell-arrival rate, and is uninfluenced by the presence or absence of 'long-range dependence' in the auto-correlation. Thus, the VBR data-stream studied by Heyman et al (1992) is an example in which significant multiplexing gain is achieved for multiple sources in the presence of long-range dependence.

4.2 The DAR Approximation for VBR Traffic

We now come to the central question of our investigation: if, in general, long-range dependence in a traffic stream is important for queue-performance, as suggested by the results of Erramilli et al (1994a), how do we explain the successful use of Markovian DAR models for VBR video-traffic, in Heyman et al (1992) and Elwalid et al (1995), for the determination of effective link-bandwidth for a specified cell-loss rate?

We examine below a few of the factors which we believe to be pertinent to this question.

4.2.1 The Correlation-Structure in the VBR Video-Stream Studied

As noted earlier, Beran et al (1992) exhibits long-range dependence, with a Hurst parameter $H=0.7$. The Hurst parameter, however, is an *asymptotic* property of the auto-correlation. For traffic streams that are only *approximately* self-similar [Leland et al (1993)], (unlike fractional Brownian motion [Norros (1994)], which is *exactly* self-similar), the behavior of auto-correlation for small time-lags could be quite different from its eventual behavior. In fact, for the VBR traffic stream in question, it is shown in Heyman et al (1992) that the auto-correlation of frame-size, *over an initial range of about 150 frame-lags*, is very accurately matched by the exponential decay associated with the DAR model. On the other hand, the fact that the *initial* behavior of the auto-correlation of the VBR stream is quite different from its *asymptotic* behavior is confirmed by the 'variance-time' plot for this stream appearing in Beran et al (1995).

Thus, the video-stream under study is consistent with a Markovian model for auto-correlation over the span of the first 150 frame-lags; the same property has also been shown to hold for other video-streams studied by Heyman et al in (1994). A lag of 150 frames corresponds to a time-lag of 6 seconds. Thus, in situations where 'memory' beyond 6 seconds has no significant effect on the buffer-queue, one expects that it is safe to ignore the asymptotic behavior of the auto-correlation. However, it is hard to define precisely what such 'safe' situations are, except to note that low levels of utilization correspond to small buffer backlogs, and thus to 'safe' conditions.

4.2.2 A Scaling Property of Streams with Long-Range Dependence

In recent work by Krishnan (1995), it has been shown that the fractional Brownian model [Norros (1994)] for traffic with long-range dependence exhibits the following property: given two such streams with the same mean data-arrival rate λ , the same *fixed* (variance/mean) ratio α for arrivals per unit time, and two different Hurst parameters H_1 and H_2 , with $H_1 < H_2$, then, for a given buffer-level b and a given buffer-overflow probability p (in an infinite buffer), there exists a value λ^* of λ such that for $\lambda > \lambda^*$, the stream with the higher Hurst parameter requires *smaller* bandwidth to meet the specified grade-of-service. Thus, there appears to be a 'crossover' in the relative performance of the two streams with increasing load (for a fixed variance/mean ratio). This behavior has been confirmed in recent simulations of queue-behavior with traffic generated from the fractional Brownian model [Erramilli (1994b)].

When we consider the multiplexing of *independent* sources of identical statistical characteristics, this result suggests that when enough streams are multiplexed, a high- H source requires less bandwidth than a low- H source. Of course, strictly speaking, the result holds only for fractional Brownian models, which are *exactly* self-similar, whereas the VBR video-source studied in Heyman et al (1992) is only asymptotically self-similar with $H=0.7$. Nevertheless, we decided to investigate whether the 'crossover' effect can be observed with the VBR source. The DAR model in Heyman et al (1992) for the VBR source, it should be noted, has the same mean load and variance of frame-size as the actual source, but has $H=0.5$, since it is a Markovian model.

In devising a satisfactory procedure to check for the crossover effect, we had to contend with the fact that the only data available for VBR traffic consisted of a *single*, 30-minute trace for a single source. To simulate the multiplexing of an ensemble of such independent sources, Heyman et al (1992) proposed the

use of *phase-shifted copies of the single realization*. Even though the combination of phase-shifted streams may be regarded, for some choices of phase-shifts, as an acceptable approximation to the multiplexing of independent sources, it is nevertheless true that the actual *performance* obtained with such a combined stream can be quite sensitive to the set of phase-shifts used. Therefore, it becomes necessary to consider simulations for various choices of phase-shifts in order to arrive at representative results.

Heyman et al (1992) also propose a more detailed Markov Chain (MC) model for VBR traffic which offers greater accuracy than the DAR model, *especially when the number of multiplexed sources is small*. For this reason, we repeated the above steps for the more elaborate MC model also, but, owing to space limitations, present only the DAR results.

Thus, we proceeded as follows in our experiment:

1. For various values of n in the range $1 \leq n \leq 20$, we multiplexed n DAR streams, generated by multiplexing phase-shifted realizations of the DAR model of the source, for *two different choices of phases*, both of which produced acceptable approximations to independent sources. For each value of n , and for each of the two choices of phase-shifts, we determined the effective bandwidth required for the corresponding multiplexed stream, to realize a buffer-overflow probability of $1.0e-3$ for a buffer-level of 5000 cells. In order to obtain a large enough sample for the estimation of buffer overflows, we chose to work with the overflow-probability of $1.0e-3$ rather than $1.0e-6$. *The effective bandwidth was determined by means of actual simulation* because the algorithm of Elwalid and Mitra (1993) used for the results in Elwalid et al (1995) was unavailable to us at the time this work was performed.
2. Using the effective bandwidth determined for n multiplexed DAR streams, we determined the performance that is realized when n VBR sources are multiplexed, with the streams generated, once again, by suitable phase-shifting of the single-source trace that is actually available.

The results (for the infinite-buffer case) for the VBR-DAR comparison are shown in Figures 1a-d and in Figures 2a-d, for the two choices of phases used for multiplexing n sources, for different values of n . Of course, Figures 1a and 2a, for $n = 1$, are the same. In Figures 1 and 2, we observe that the performance obtained is quite sensitive to the choice of phases, but, in both cases, when n (the number of multiplexed streams) is greater than 11, the performance of the actual VBR source ($H=0.7$) is better than that of the DAR model ($H=0.5$). For several values of n below 11 (not included in the Figures), the relative performance was found to depend on the choice of phases, and rather erratic.

On the basis of the results in Figures 1 and 2 (and similar results obtained with the MC model), it appears safe to conclude that the results are consistent with the property, proved for fractional Brownian Motion by Krishnan (1995), that for a large enough number of multiplexed streams, the performance of a high- H source becomes better than that of a low- H source.

We also present, in Table 1, a VBR-DAR performance comparison (for the phases corresponding to Figure 1) for the case of a *finite* buffer of 4500 cells, with the effective bandwidth determined for the multiplexed DAR streams for a loss-probability of $1.0e-3$. The results show that, for the finite-buffer case also, the performance of the VBR source is better than that of the DAR model for $n > 11$.

Elwalid et al (1995) consider 15 or more sources in their investigations, a number greater than the 'critical' point observed in the above results. (Their bandwidth calculations were made for a cell loss-rate of $1.0e-6$, while we used a loss-rate of $1.0e-3$, but this has a minor effect on the precise point of the crossover). In other words, the results suggest that *the results of Heyman et al (1992) and of Elwalid et al (1995) correspond to a regime in which the Markovian model (DAR or MC) for determining effective bandwidth provides a conservative estimate for the bandwidth needed for the actual VBR video-sources*. This is in agreement with the simulation results in Elwalid et al (1995), which show that more sources can

be supported than predicted by the DAR model, for the numbers of sources considered there.

TABLE 1: Loss Probabilities for Multiple Sources (Finite Buffers)

Number of sources	Bandwidth (Mbps)	Loss probability	
		VBR	DAR
1	4.27	2.195 e-3	3.308 e-4
5	14.50	5.776 e-4	1.141 e-4
8	21.62	3.336 e-4	5.421 e-5
10	24.25	4.017 e-4	3.955 e-5
11	27.72	1.506 e-5	2.557 e-5
15	35.68	0.0	3.199 e-5
20	44.30	0.0	2.819 e-5

4.2.3 Finite-Buffer Effects for Traffic with Long-Range Dependence

The above simulations demonstrating 'crossover' assumed infinite buffers (the case for which the crossover effect was proved by Krishnan (1995)), while Elwalid et al (1995) worked with finite buffers of various sizes in their simulations. With Markovian models, there is reason [Avi-Itzhak and Halfin (1994)] to expect that estimating the loss from a finite buffer by the corresponding overflow probability in an infinite buffer is a conservative approximation; i.e., the loss rate from a finite buffer of size B is smaller than the 'overflow' rate above level B in an infinite buffer. A comparison of the above simulation results with those of Table 1 shows a similar relation for long-range dependent traffic also. Moreover, recent research [Rao et al (1996)] suggests that finite buffers can, in certain cases, mitigate the ill-effects of long-range dependence on buffer queues, rendering Markovian approximations acceptable. These results will be discussed in a future publication.

5. Conclusion

Recent Bellcore studies [Leland et al (1993) and Erramilli et al (1994a)] have demonstrated that traffic streams in high-speed networks exhibit long-range dependence in their auto-correlation, characterized by correlation that decays at a smaller rate than in Markovian streams. It has also been shown [Erramilli et al (1994a)] that there are cases in which link-engineering based on Markovian models would underestimate the actual bandwidth requirements for streams with long-range dependence. However, there is also an example [Elwalid et al] of the successful application of Markovian models to determine the bandwidth requirements of VBR video-sources, which are known to exhibit long-range dependence. Thus, there is a need to bridge the gap in our understanding of the effects of long-range dependence on traffic engineering in order to account for both results. Spurred by this need, we undertook an investigation to gain a fuller understanding of the effect of the Hurst parameter H on link-engineering.

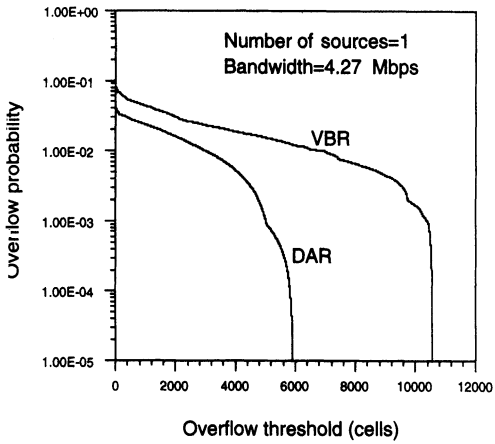
During the course of our investigation, Krishnan (1995) discovered the following curious 'scaling' property of streams with long-range dependence that helps to bridge the gap between the results of Leland et al (1993) and those of Heyman et al (1992) and Elwalid et al (1995): for *exactly* self-similar traffic streams described by the fractional Brownian model [Norros (1994)], when a sufficiently large number of sources are multiplexed, a high- H source requires less bandwidth than a low- H source [Krishnan (1995)]. From the results presented here, it appears that this effect, derived for the fractional Brownian model, is also at work in the VBR video sources considered by Heyman et al (1992). On the basis of these results, we have offered an explanation that accounts for the results of both Leland et al (1993) and Heyman et al (1992). The analysis suggests that a high value of H is not, in itself, a reason to suppose that Markovian traffic models will lead to under-engineering of bandwidth on ATM links.

6. Acknowledgement

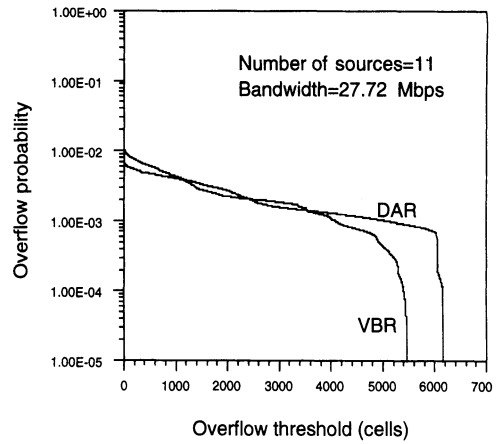
We thank A. Erramilli, D.P. Heyman, T.V. Lakshman, and W. Willinger for their helpful advice at several stages of this investigation.

7. REFERENCES

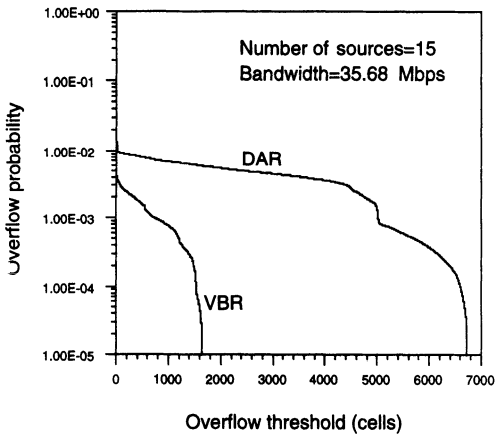
- Avi-Itzhak, B. and Halfin, S. (1994) A method for call admission control in ATM networks with heterogeneous sources, Bellcore Technical Report.
- Beran, J., Sherman, R., Taqqu, M.S. and Willinger, W. (1995) Long-range dependence in variable-bit-rate video traffic, *IEEE Transactions on Communications*, **43**, 1566-1579.
- Elwalid, A.I. and Mitra, D (1993) Effective bandwidth of general Markovian traffic sources and admission control of high speed networks, *IEEE/ACM Transactions on Networking*, **1**, 329-343.
- Elwalid, A.I., Heyman, D.P., Lakshman, T.V. and Mitra, D (1995) Fundamental bounds and approximations for ATM multiplexers with applications to video teleconferencing, *IEEE Journal on Selected Areas in Communications*, **13**, 1004-1016.
- Erramilli, A., Narayan, O. and Willinger, W. (1994a) Experimental queueing analysis with long-range dependent packet traffic, Bellcore Technical Report.
- Erramilli, A. (1994b), personal communication.
- Heyman, D.P., Tabatabai, A. and Lakshman, T.V. (1992) Statistical analysis and simulation study of video teleconference traffic in ATM Networks, *IEEE Transactions on Circuits and Systems for Video Technology*, **2**, 49-59.
- Heyman, D.P., Tabatabai, A., Lakshman, T.V. and Heecke, H. (1994) Source models for VBR broadcast-video traffic, *Proceedings of IEEE INFOCOM'94*, 664-671,
- Krishnan, K.R. (1995) A new class of performance results for fractional Brownian traffic model, Bellcore Technical Report.
- Leland, W.E., Taqqu, M.S., Willinger, W. and Wilson, D.V. (1993) On the self-similar nature of ethernet traffic, *Proceedings of ACM SIGCOMM'93*, 183-193.
- Norros, I. (1994) A storage model with self-similar input, *Queueing Systems*, **16**, 387-396.
- Rao, B.V., Krishnan, K.R. and Heyman, D.P. (1996), to be published.



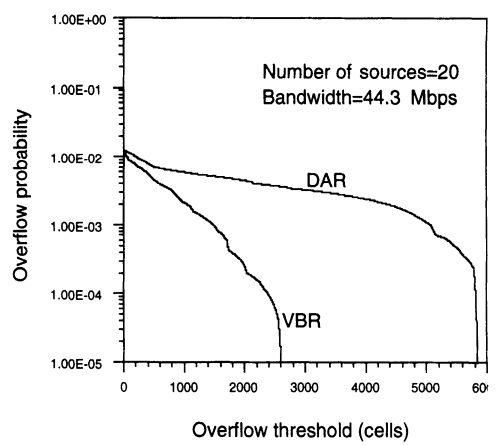
(a)



(b)

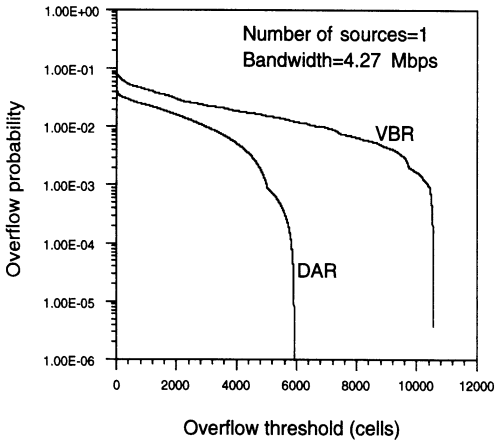


(c)

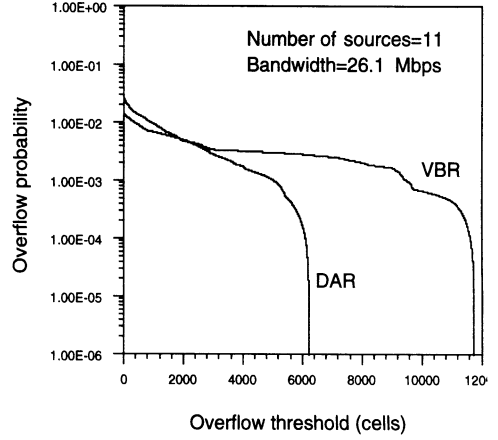


(d)

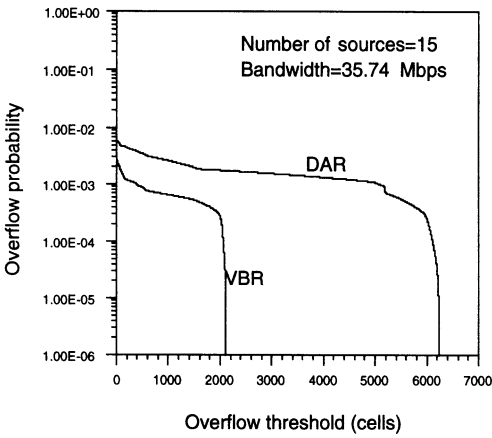
Figure 1 Buffer Distribution for Multiple Sources (Phase Sequence 1).



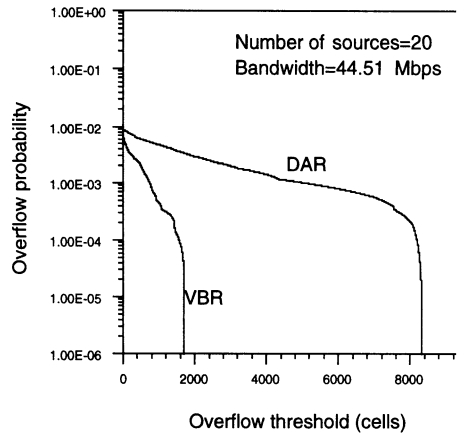
(a)



(b)



(c)



(d)

Figure 2 Buffer Distribution for Multiple Sources (Phase Sequence 2).