Multi-Scaling Models of TCP/IP and Sub-Frame VBR Video Traffic

Ashok Erramilli, Onuttom Narayan, Arnold Neidhardt, and Iraj Saniee

Abstract: **Recent measurement and simulation studies have revealed that wide area network traffic displays complex statistical characteristics–possibly *multifractal scaling*–on fine timescales, in addition to the well-known property of *self-similar scaling* on coarser timescales. In this paper we investigate the performance and network engineering significance of these fine timescale features using measured TCP and MPEG2 video traces, queueing simulations and analytical arguments. We demonstrate that the fine timescale features can affect performance substantially at low and intermediate utilizations, while the longer timescale self-similarity is important at intermediate and high utilizations. We relate the fine timescale structure in the measured TCP traces to flow controls, and show that UDP traffic–which is not flow controlled–lacks such fine timescale structure. Likewise we relate the fine timescale structure in video MPEG2 traces to sub-frame encoding. We show that it is possibly to construct a relatively parsimonious *multifractal cascade* model of fine timescale features that matches the queueing performance of both the TCP and video traces. We outline an analytical method to estimate performance for traffic that is self-similar on coarse timescales and multi-fractal on fine timescales, and show that the engineering problem of setting *safe operating points* for planning or admission controls can be significantly influenced by fine timescale fluctuations in network traffic. The work reported here can be used to model the relevant characteristics of wide area traffic across a full range of engineering timescales, and can be the basis of more accurate network performance analysis and engineering.**

Index Terms: **Self-similarity and long range dependence, performance and queueing analysis, TCP, UDP, MPEG2 traffic modeling, multifractals.**

I. INTRODUCTION

It is now generally accepted that appropriately aggregated wide area network packet network traffic exhibits self-similar scaling over a wide range of timescales [1]–[7]. More recent measurement studies have indicated that beyond self-similarity, wide area TCP traffic shows more complex statistical characteristics–perhaps multi-fractal scaling–on finer timescales. From the viewpoint of network engineering, such observations of self-similar or multi-fractal scaling are of interest only to the extent that these features can impact traffic performance. The performance and traffic engineering implications of self-similarity have already been demonstrated (see for example [8], [9]). The object of this paper is to examine the performance and network engineering significance of the fine timescale features that are "beyond self-similarity," to gain an understanding of situations when they are consequential, to verify multifractal models for these features in such situations, and to develop analysis methods based on these descriptions.

A second objective of this paper is to investigate the application of the same modeling methodology, based on multi-fractal scaling, to model traffic features of VBR video traffic below a frame level. Although there has been considerable research on video traffic (for a representative, but far from exhaustive list see [3], [10]–[12]), much of the prior literature (see [3], [11]) is concerned with modeling fluctuations in video traffic at and above the frame level. However, as has been argued [10], performance can be very often determined by fluctuations at the cell level. One would thus expect characterizations at the slice level to be more accurate than the frame level for many queueing scenarios.

Fine timescale, or large frequency, features of a traffic stream are typically lost when individual streams are aggregated in the core of a network, where as has been shown, the long time correlations and low frequencies dominate the overall performance. However, at the edge of the network, where call admission control and conformance tests are typically performed, these fine time-scale characteristics of the individual stream become relevant, irrespective of the subsequent aggregation in the network core.

Typically, fine timescale features in traffic are less robust than coarse timescale fluctuations: They can be readily modified through buffering or shaping, and are more sensitive to flow control parameters, content, coding scheme and coding parameters. For this reason, our focus is on a structural methodology that can capture the full range of fine timescale features observed in network traffic: TCP flow control induced fluctuations below the timescales of TCP segment round-trip times, and video sub-frame characteristics that are determined by the coding schemes. These are discussed next.

A. Beyond Self-similarity: Observed Fine Timescale Features

Given the demonstrated significance of long timescale correlations on network performance, a self-similar traffic model, Fractional Brownian Motion (FBM), has been proposed as a

Manuscript received November 23, 2000; approved for publication by Bo Li, Division III Editor, November 13, 2001.

A. Erramilli is with Qnetworx, 1119 Campus Drive West, Morganville, NJ 07751, USA, e-mail: ashok@qnetworx.com.

O. Narayan is with the Physics Department, University of California, Santa Cruz, CA 95064, e-mail: narayan@physics.ucsc.edu.

A. Neidhardt is with Telcordia Technologies Inc., 331 Newman Springs Road, Red Bank, NJ 07701, USA, e-mail: arnie@research.telcordia.com.

I. Saniee is with Bell Laboratories, Lucent Technologies, 600 Mountain Avenue, Murray Hill, NJ 07974, USA, e-mail: iis@research.bell-labs.com.

This work is supported in part by NSF grant NCR-9628067.

parsimonious and tractable model of packet traffic [13]. In this model, the number of arrivals (e.g., packets, cells, bytes) in an observation interval $(0, t)$ can be represented by a self-similar fluctuation about its expectation:

$$A(t) = mt + \sqrt{am}Z(t), \quad (1)$$

where $A(t)$ denotes the number of arrivals up to time t, m is the mean arrival rate, a is the peakedness parameter, and $Z(t)$ is a standard (mean zero, variance t^{2H}) FBM. The fluctuations observed over different timescales are self-similar in that $Z(\alpha t)$ is distributed as $\alpha^H Z(t)$, where H is the Hurst parameter characterizing the self-similarity.

In networking terms, FBM is an accurate model of traffic arrivals under the following conditions: i) the traffic is aggregated from a large number of low activity independent users,whose peak rates are small relative to link capacity, so that the ii) the timescales of interest (i.e., those that largely determine queueing behavior) fall within the scaling region, and iii) the impact of flow controls is not significant [4], [14]–[17]. The simple self-similar model can break down, for example, on finer timescales over which protocol interactions (e.g., TCP) or encoding schemes (e.g., MPEG2) can determine packet or cell transmission patterns.

For the case of TCP traffic, recent measurement studies [18], [19] show that low-aggregate, broadband, wide area TCP/IP traffic departs from FBM at time scales shorter than a lower cutoff in two key respects. Firstly, the marginal distribution of traffic counts is clearly non-Gaussian. Thus a purely second order description—in terms of the mean and second moment—is not *complete* (which is to say that two streams with the same second order descriptions can give rise to very different queueing behavior). Secondly, the scaling exponent of the variance on timescales shorter than the lower cutoff is smaller than the asymptotic exponent. As a refinement of these two deviations from FBM, it is claimed [18], [19] that wide area TCP/IP traffic below the lower cutoff exhibits multi-fractal scaling, with different moments of the traffic process showing scaling characterized by distinct exponents (in contrast, the scaling behavior in FBM is characterized by only a single exponent, namely, the Hurst parameter). Analysis of detailed simulations of TCP also support these observations [20], [21]. The lower cutoff is empirically observed [18] to be of the order of the round-trip time of a TCP segment.

The case of VBR video traffic presents special challenges to the modeler. For one, video traffic shows complex temporal structure, with characteristic features on short (sub-frame), intermediate (1–100 frames) and long timescales (> 100 frames). Much of the earlier work in this area has focused on the last two regimes. Secondly, traffic characteristics can depend sensitively on the coding scheme (e.g., H.261, MPEG) as well as specific parameters employed in the coding (as with MPEG2). We use MPEG2 video as representative of VBR video traffic, and consider just one aspect in its modeling: traffic fluctuations at the slice, or the sub-frame level. Typically, traffic on the sub-frame level is highly bursty and irregular, and as demonstrated in this paper, these features can be adequately described using multifractal scaling.

While it is inevitable that on the shortest timescales *any* packet traffic must be a non-Gaussian process, the important feature in wide area TCP/IP and VBR video traffic is that, unlike earlier work on LAN traffic [8], this is found to have a significant impact on network performance [22]. Thus a compact representation of short time scale features, either using multifractals or alternative models [23], are useful in understanding the queueing delays of WAN traffic.

B. Scope and Summary of Work

We investigate the fine timescale features in wide area TCP/IP and MPEG2 video traffic, using a series of simulation experiments and an analysis of the problem of setting network operating points. We use the TCP packets from a one-hour trace due to Digital Equipment Corporation, gathered at Digital's primary Internet access point [24]. The analysis reported here was conducted on *multiple* 6-minute samples of the 1 hour TCP trace, but for clarity of presentation, we report on a single 6-minute trace in this paper. For video traffic, the MPEG2 trace we study consists of 6 minutes of gymnastics at the 1996 Atlanta Olympics containing 280,000 time slices of duration 4/3 ms each, with a minimum, mean and maximum of 15,815 and 5,514 bytes per time slice respectively [25]. Traces of varying durations from other events at the Olympics were also investigated, and show the same essential features. While our primary focus in this paper are the traffic sensitive features within a single stream, it is understood that network traffic consists of (perhaps) limited aggregates of such individual streams.

We verify that both the TCP/IP and video traces show the same fine timescale deviations from FBM as reported in [18], [19] for TCP/IP traffic, on the basis of a multi-fractal scaling analysis. There is a crossover from the fine timescale to coarser timescale behavior at approximately the single frame level for video traffic; for TCP, a similar crossover is observed in the data at $\sim$ 512 ms.[1] The impact on performance of the fine timescale features is demonstrated in both cases by aggregating the traffic over time to the level of the crossover time and interpolating down to the finer timescales in various ways. Using a multi-fractal cascade to interpolate the fine timescale structure appears to work adequately. We also perform shuffling experiments, similar to those in [8] for LAN traffic, in order to show that low-frequency features remain important.

Motivated by these results, we obtain analytical estimates for the viable operating point for a network, under the assumption that its traffic can be represented by an FBM process on long timescales, crossing over below a lower cutoff to the FBM behavior to a purely multi-fractal cascade.[2] These estimates agree with the simulation result that one has to operate at a *reduced* utilization due to burstiness at fine timescale features in the traffic, and show how this reduction depends on network parame-

[1]This value, close to the granularity of timers used in many TCP implementations, may suggest that the crossover is an artifact. However, for the data analyzed here, the timestamps were specified to μs and had an accuracy of $\sim$ 1ms, which eliminates this possibility. Moreover, such a crossover time has also been observed in other studies, with its value depending on the network environment: for instance, in [18], it is of the order of one second. In detailed TCP simulations using fine grained timers, the crossover time has been observed to vary with the round-trip time.

[2]For video traffic, there is also the intermediate scale, from the single frame to the group of frames (GOF) level, to be taken care of.

ters. Another important (but not surprising) finding is that multi-fractal scaling exponents alone are *not* sufficient to describe the traffic: one needs to specify the *magnitude* of the fluctuations in (different moments of) the traffic. We fix these by assuming that the behavior below the lower cutoff should satisfy the "boundary conditions" imposed by the behavior above it.

C. Origins of Fine Timescale Traffic Features

As mentioned earlier, the fine timescale features arise due to very different mechanisms in TCP and VBR video traffic. For TCP, where the data is obtained from direct measurements on a network, it has been argued [18] that the fine timescale behavior is due to the manner in which TCP regulates the load offered by each individual source to the network. The fact that the transition between the fine and coarse time behavior occurs at a timescale of the order of the round-trip time of a TCP segment lends support to this scenario. As additional evidence, we shall show later in this paper that when TCP flow controls are *not* present (as is the case with UDP transport), WAN traffic *reverts to simple FBM behavior.* On the other hand, for MPEG2 video the fine timescale structure exists in the manner in which frames are encoded, and arises independent of any network interactions. A common descriptive framework for these two very different mechanisms might appear unlikely. However, in a queueing context, both these mechanisms generate a highly irregular distribution of packet/ cell/ byte counts over a time interval, and as we demonstrate in this paper, multi-fractal cascades mimic such irregular distributions sufficiently well for estimating performance.

There is however one important difference in the way these models should be interpreted for TCP and VBR video traffic. It should be stressed that this paper is primarily concerned with "open loop" characterizations of TCP traffic (such as self similar or multi-fractal descriptions), in that it considers the effect of observed characteristics on network congestion, but does not explicitly address the effects of network congestion on the offered traffic. Indisputably, a closed loop model that captures these interactions is necessary for predictive purposes. The multi-fractal cascades studied here are an important intermediate step to achieving this goal because to some extent they mimic the effects of TCP flow control. Thus the open loop characterizations studied here should be viewed as a necessary intermediate step to achieving the long term objectives of a phenomenological model that completely captures the interactions between network state and traffic flows. Of course, this is not an issue for real-time VBR video traffic.

D. Outline

The rest of this paper is organized as follows: Section II provides a short background on mono-fractal and multi-fractal scaling. Section III characterizes the data, showing a transition point between fine and coarse time scale behavior in the case of TCP and MPEG2 traffic, with the fine time scale fluctuations described reasonably well as a multi-fractal cascade. Section IV describes experiments, in which the traffic traces are aggregated to this transition point, and then interpolated down to fine timescales in a variety of ways, with the simulated performance compared to that of the original in each case. Section V, obtains analytical performance estimates for a simple model which is multi-fractal on fine timescales and FBM on coarser timescales. Finally, Section VI summarizes our conclusions with suggestions for further work.

II. MULTI-FRACTAL SCALING

Since the main focus of this paper is to investigate the usefulness and implications of a multi-fractal description of traffic, we summarize some of the basic concepts associated with multi-fractals. Consider a traffic arrival process $A(t)$, defined as the total traffic that arrives in the interval $[0, t)$. The associated increment process $X_\Delta(i)$, is defined by

$$X_\Delta(i) = A(i\Delta) - A((i-1)\Delta). \tag{2}$$

The basic scaling hypothesis is that the moments of the increment process behave as:

$$\sum_i X_\Delta(i)^q \sim C(q)\Delta^{-\tau(q)} \qquad \text{for} \quad \Delta \to 0. \tag{3}$$

In practice, the scaling hypothesis can be said to be reasonable if the above behavior is satisfied over a *range* of timescales (for the processes considered in this paper, these would apply to the fine timescales). In general the *structure* function $\tau(q)$ defined above – if it exists – will be a decreasing and nonlinear function of q. Notice that $\tau(q = 0)$ must be equal to unity. When $\tau(q)$ is linear in q, the scaling behavior is said to be *mono-fractal*, and a single scaling exponent completely determines the scaling behavior of all the moments of the traffic arrival process.

Multi-fractal scaling can also be described in terms of an α-spectrum $f(\alpha)$ if certain regularity conditions are satisfied, in which case $f(\alpha)$ is defined by

$$f(\alpha) = \inf_q \; [\alpha q + \tau(q)] \tag{4}$$

with q ranging over $(-\infty, \infty)$ in the infimum. Under the regularity conditions, the α-spectrum can be interpreted as the Hausdorff dimension of the set of points on the time axis with the same local exponent α. If the regularity conditions are not satisfied, then at least three separate α-spectra can be defined [26].

When the scaling is mono-fractal, so that $\tau(q)$ is linear in q (and $\tau(0) = 1$), it is easy to see that $f(\alpha)$ has a value of unity at a single point, and $-\infty$ for all other α. That is to say, the same exponent α describes the scaling behavior of the traffic arrival process for all time t. In practical terms, multi-fractal scaling is indicated by the presence of a broad α-spectrum.

One of the standard techniques to generate multi-fractal scaling on finer timescales is the *cascade* construction. In the semi-random version of this approach, a coarse timescale count over an interval is distributed over finer timescales by assigning a fraction p to one half of the interval and the remainder to the other half. The choice of which of the two intervals to assign the fraction p is done randomly. The distribution process is repeated for a number of stages, and the result is marked by extreme irregularity over a range of fine timescales. There are numerous variations in the cascade construction [18]. For instance, the

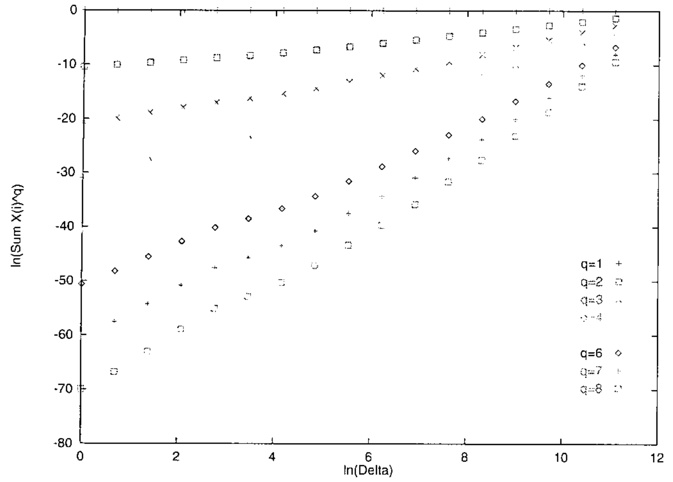

Fig. 1. Log-log plot for TCP traffic of various moments of the aggregated increment process, $\sum_i X_\Delta(i)^q$, as a function of the size of the aggregation interval Δ. Natural logarithms have been used. X_Δ is measured as a fraction of the total byte count in the entire trace, so that the logarithms of all the moments approach zero at the right end of the plot. Δ is measured in ms.

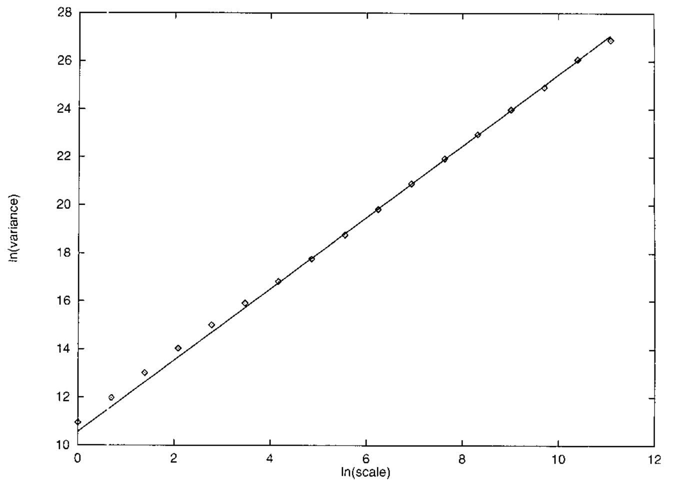

Fig. 2. Log-log plot for TCP traffic of the variance in the byte counts per time interval, as a function of the size of the time interval. Natural logarithms have been used. The straight line above approximately 512 ms corresponds to the low frequency FBM behavior of the traffic fluctuations.

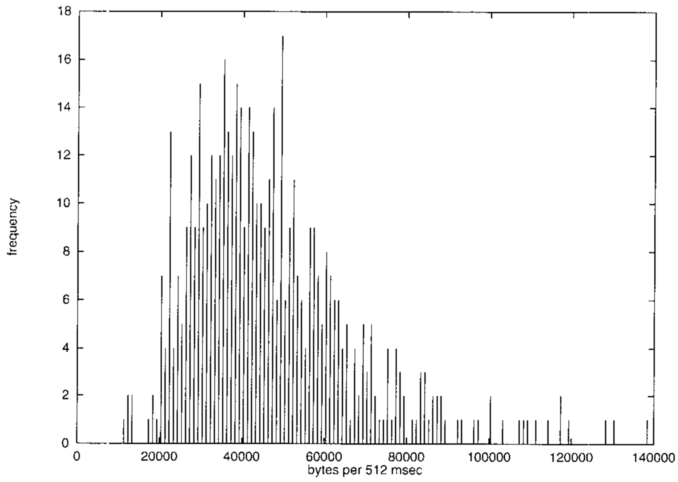

Fig. 3. Marginals of the TCP traffic aggregated to 512 ms, showing roughly Gaussian behavior.

construction can also be strictly deterministic (i.e., the traffic is distributed in the fixed ratio $p : 1 - p$ at every stage).

III. SCALING PROPERTIES OF DATA

Fig. 1 shows a log-log plot of the various moments of the TCP data. The moments of the increment process are plotted as a function of the size of the aggregation interval Δ, in accordance with Eq. (3).

There is a visible change in the behavior of all the different moments at $\Delta \sim 512$ ms, with the slopes of the curves changing at around this point. Below 512 ms, one can fit the different moments in Fig. 1 to straight lines of different slopes. At timescales coarser than 512 ms, the traffic is approximately FBM. This is shown in Fig. 2 and Fig. 3, where beyond 512 ms the traffic is approximately Gaussian with self-similar fluctuations. It appears that the scaling in the variance assumes its asymptotic behavior before the marginal distribution of counts "looks" fully Gaussian.

For a pure FBM process, the traffic would have been self-similar (and Gaussian) on all timescales. Conversely, for a pure multifractal, the traffic would have looked self-similar but non-Gaussian on all timescales. As Figs. 1–3 demonstrate, neither of these descriptions is satisfactory. There are two separate regimes, separated by $\Delta \sim 512$ ms. Our strategy is to model the traffic as FBM on coarse timescales, and interpolate down to fine timescales using a multifractal description. The linearity of the plots in Fig. 1 in the fine timescale regime conforms to Eq. (3). The extent of this regime is approximately the same as has been observed with other traffic measurements in [18]. The slopes in the fine timescale regime, which define $\tau(q)$, are plotted against q in Fig. 4, and show clear nonlinearity, invalidating a mono-fractal description in this regime. By comparison, the familiar FBM arrival process can be said to exhibit mono-fractal scaling, when the fluctuations about the mean arrival rate (instead of the arrival rate itself) are considered: $\tau(q)$ (for non-negative even q) is equal to $1 - qH$, where H is the Hurst parameter characterizing the FBM.

Note that there is a slight difference in perspective between the multi-fractal approach at fine timescales, which directly models the traffic rate, and the FBM approach at coarse timescales, which models the *fluctuations* from the mean in the traffic rate. For *all* traffic, at very long timescales Δ, $X_\Delta(i)$ will have a distribution sharply peaked around its mean. Correspondingly, $\sum_i X_\Delta^q(i)$ scales as Δ^{q-1} in this regime. Despite the trivial behavior of the moments of the traffic, the fluctuations still show long-range correlations. This, however, has been well studied, and is not relevant for the fine timescale structure which is the focus of this paper.

Fig. 4 also shows the structure function $\tau(q)$ for a semi-random cascade (defined in the previous section) on fine timescales, with the parameter p chosen to match the $\tau(q)$ obtained for the data most closely. This cascade construction is used in the next section where the queueing behavior for TCP

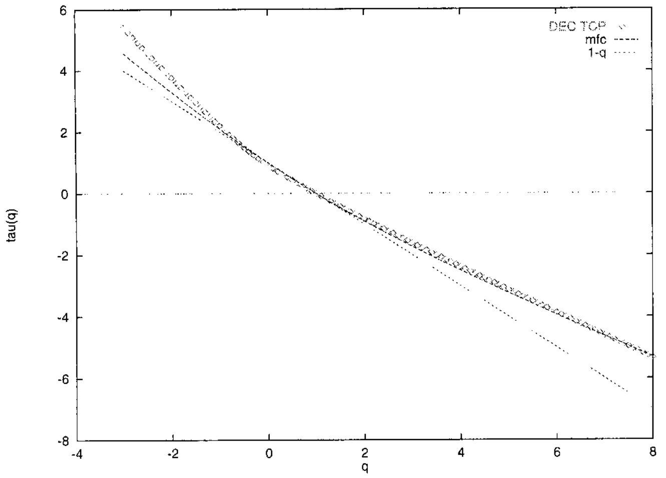

Fig. 4. Plot of the scaling function $\tau(q)$, defined in Eq. (3), for various values of the moment q. Here $\tau(q)$ is obtained from the part of Fig. 1 below 512 ms. The curve DEC TCP is for the original measured TCP data, while the curve mfc is for a semi-random cascade with $p = 0.65$, chosen to visually match the DEC TCP curve as closely as possible. The straight line $1 - q$ is drawn to show the curvature of the $\tau(q)$ curve.

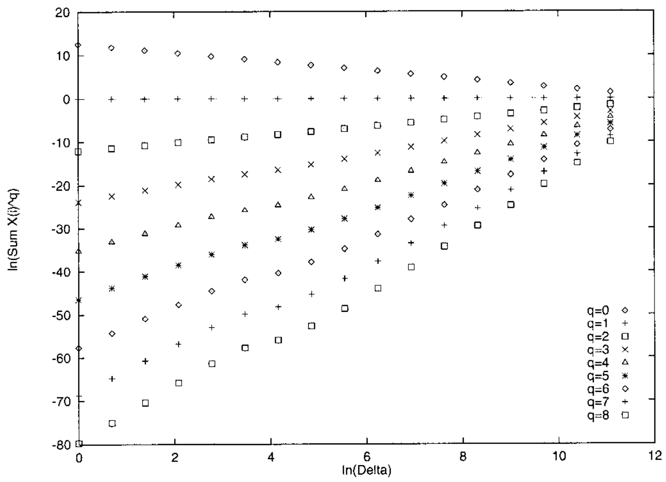

Fig. 5. Log-log plot for MPEG2 video traffic of various moments of the aggregated increment process $\sum_i X^q_\Delta(i)$ as a function of the aggregation interval Δ. Natural logarithms have been used. X_Δ is measured as a fraction of the total byte count in the entire trace. Δ is measured in units of the time interval between video slices, $1/750$ s.

traffic is analyzed.

We now perform a similar analysis for the MPEG2 video data. Fig. 5 is a log-log plot of the different moments. There is a transition from a short time regime at approximately 40 ms (or 30 slices, which corresponds to ~ 3.5 in the logarithmic time axis in Fig. 5). Interestingly, this is the time duration of a frame, suggesting that the fine timescale behavior is characteristic of how MPEG2 distributes data between slices within a single frame.[3] In the fine timescale regime, the curves in the log-log plot of Fig. 5 are all linear, although admittedly over a not very large range of Δ, showing that Eq. (3) holds in this regime.

[3]The same threshold value of 40 ms was observed in several MPEG2 traces that are characterized by a frame rate of 25 f/s, showing the connection between the observed fine timescale regime and sub-frame characteristics.

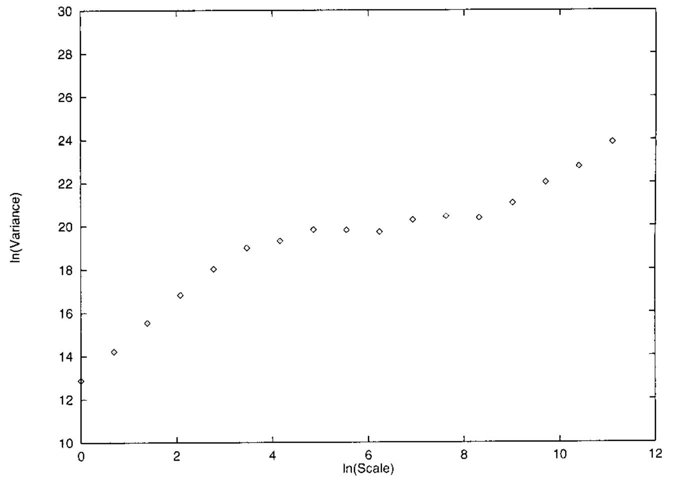

Fig. 6. Log-log plot for video traffic of the variance in the byte counts per time interval, as a function of the size of the time interval. Natural logarithms have been used. Three different regimes are seen.

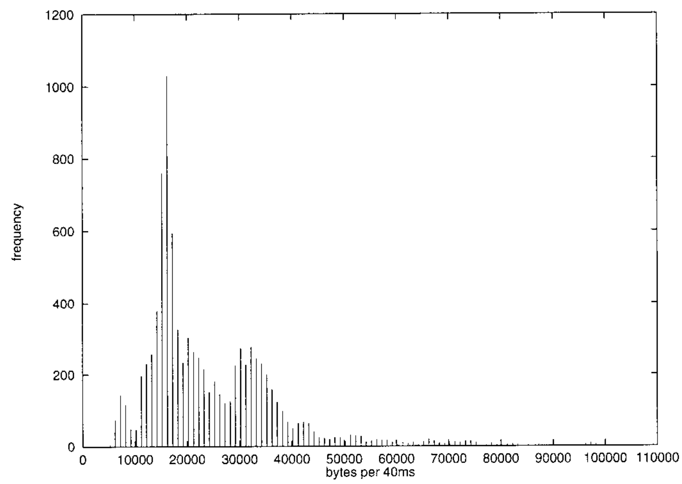

Fig. 7. Marginals of the video traffic aggregated to 40 ms, showing that clear non-Gaussian behavior exists even at this level. The unnormalized probabilities are plotted, with bin sizes of 1000 bytes. Higher levels of aggregation eventually yield Gaussian marginals.

As mentioned in the previous section, for video traffic the fine and coarse time regimes are separated by an *intermediate* time regime, extending roughly from the single frame to the GOF level. Since this is not very clear in Fig. 5, in Fig. 6 we show a variance-time plot for the video traffic. The variance in the byte counts per time interval, as a function of the size of the time interval, is shown. The fine timescale regime, which is the focus of this paper, is seen in this figure. Also seen is a long time regime, where one obtains the long range dependence that seems ubiquitous for packet traffic (with a Hurst parameter H slightly greater than 0.6.) An intermediate regime separating these two is also clearly visible. The approximately flat nature of the variance here implies anticorrelations between the different frames at this level. This is known to be the case with MPEG2: I frames, containing all the information in a frame, occur at regular intervals, separated by the much more numerous P and B frames, which are more compactly coded and therefore shorter. One can also plot the marginals of the video traffic aggregated

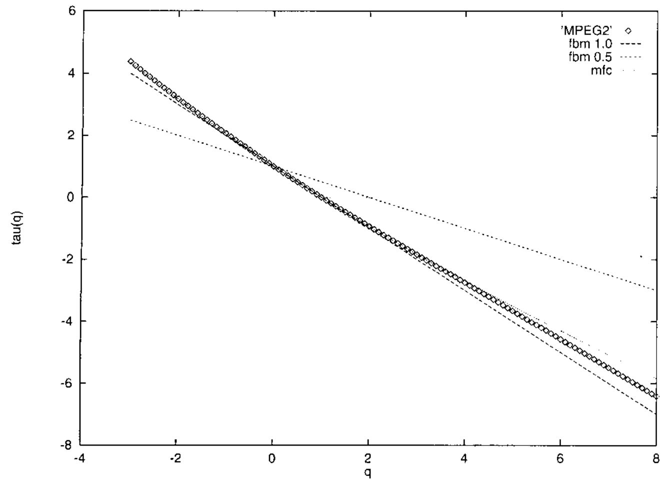

Fig. 8. The structure function, $\tau(q)$ plotted against q for the original video trace (MPEG2). Also shown (mfc) is $\tau(q)$ with $p = 0.6$, chosen to visually match the MPEG2 curve most closely. Linear structure functions for FBM with $H = 1.0$ (fbm 1.0) and $H = 0.5$ (fbm 0.5) are also shown.

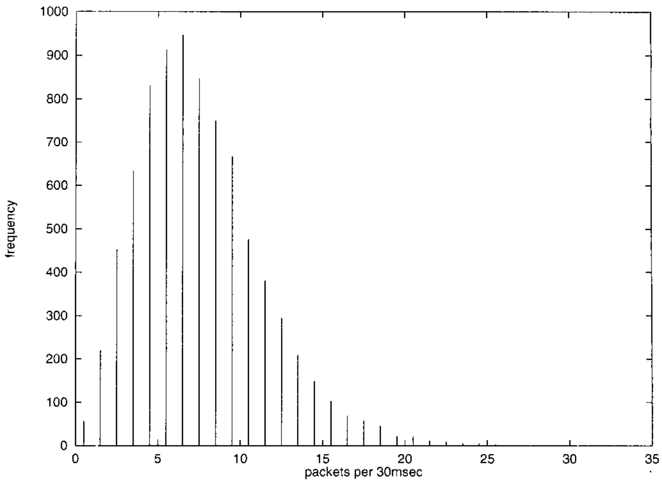

Fig. 9. Marginals of UDP traffic aggregated to 30ms, showing approximately Gaussian behavior.

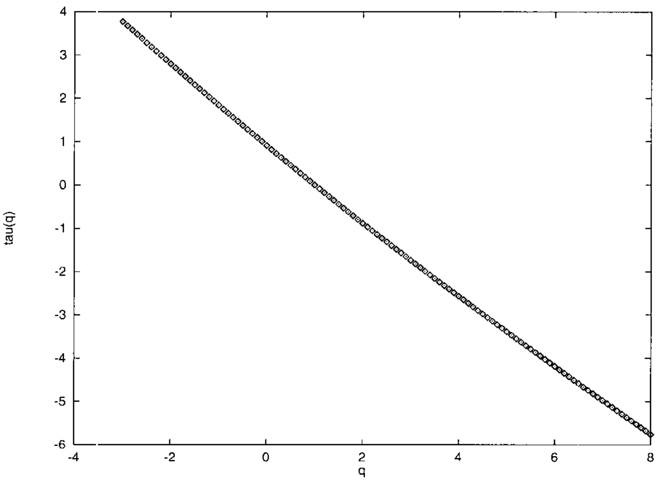

Fig. 10. The structure function $\tau(q)$ plotted as a function of q for UDP traffic, showing linear behavior.

to 40 ms (the edge of the fine timescale regime). This is shown in Fig. 7. Unlike Fig. 3, we see that the marginals at the edge of the fine timescale regime are still far from being Gaussian. It is only upon aggregation to the GOF level that approximately Gaussian marginals are obtained.

The description that we use for video traffic is therefore slightly different from TCP/IP traffic: a long time FBM regime above the GOF level, an intermediate time regime from the single frame to the GOF level, and a short time sub-frame regime. We do not attempt to model the first two regimes in this paper, but instead take the byte count at the frame level from measurements. We seek to model the third regime by interpolating down from the frame level with a multifractal description.

Fig. 8 is a plot of the structure function $\tau(q)$ as a function of q, obtained from the slopes in the fine timescale regime of the plots in Fig. 5. As was the case with the TCP data, the curve for $\tau(q)$ shows nonlinearity. Fig. 8 also shows the structure function for a multifractal cascade chosen to match $\tau(q)$ for the video data most closely. This is used in the next section of the paper; a deterministic cascade is used for the video traffic, though the semi-random construction is perhaps equally applicable.[4]

We have seen that both WAN TCP and MPEG2 video traffic have complex structure at fine timescales that goes beyond the FBM paradigm. For TCP traffic, the fine timescale structure is believed to be imposed on the input traffic by TCP flow controls. This is indirectly supported by the observation [18] that this fine timescale structure extends up to a timescale of the order of the network round trip time. Ideally, one would like to construct a faithful representation of TCP operating on a variety of model networks to demonstrate that feedback causes fine timescale structure. Here we take the simpler approach of repeating the analysis for TCP traffic on UDP traffic measured on the same network.

UDP traffic from the same DEC network [24] was analyzed for fine timescale structure. Fig. 9 shows that even at the 30ms level the marginals of the traffic can be reasonably modeled by a Gaussian distribution, in contrast to TCP. As a further check for non-trivial structure, moments of the UDP traffic distribution were calculated over timescales from 1 ms to $\sim$ 1 minute, and $\tau(q)$ computed as a function of q from these. As shown in Fig. 10, $\tau(q)$ is a linear function of q. Thus for UDP traffic, without any flow control feedback, the input traffic is adequately described by FBM, as was seen earlier for LAN traffic. Fine timescale structure is therefore not caused by differences in offered traffic input for LANs and WANs, but rather by how TCP flow controls regulate the offered traffic. It is important to appreciate that the mere *existence* of TCP is not enough to generate significant fine timescale structure; network parameters can mask this aspect of TCP flow controls. For instance, the earlier LAN data also consisted of TCP traffic, but the round trip time was not much greater than the packet transmission time, thus eliminating the short-time regime.

A cascade construction can mimic the effects of flow controls observed in wide area TCP traffic over fine timescales in the following sense. In the absence of flow controls, traffic sources

[4] A deterministic cascade is used since the fine timescale structure for video comes from encoding rather than (random) network feedback.

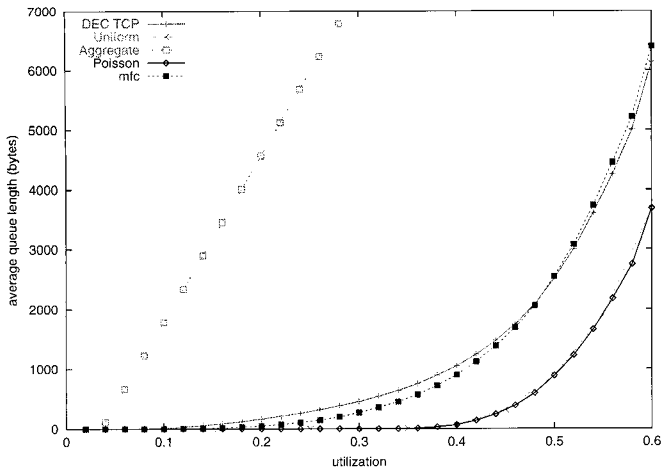

Fig. 11. The average queue length in bytes as a function of utilization for the measured DEC trace (DEC TCP) and various modifications thereof. The modifications considered are a smoothed version with all fluctuations below 512 ms removed (Uniform), a worst-case version where traffic is lumped at the beginning of each 512 ms interval (Aggregate), and Poisson (Poisson) and multifractal cascade (mfc) interpolations.

(e.g., those using UDP transport) transmit at their peak rates for as long as there is traffic to send. Such sources can therefore be modeled by ON/OFF sources which transmit at a peak rate in their ON (or burst) state. In contrast, a source using TCP transport will show more complex ON/OFF patterns within a burst. Thus single or limited aggregates of TCP sources will show highly irregular traffic distributions on fine timescales, and these are the features described by multi-scaling and multi-fractal cascades. Such irregular distributions also occur in video traffic, but due to MPEG2 coding rather than network flow controls.

IV. QUEUEING BEHAVIOR

We now turn to the queueing behavior of the measured traffic, and that of its variants discussed in the Introduction. The experiments described here are essentially open loop, and seek to compare the performance impacts of multi-scaling, vs asymptotic features observed in the traffic traces under a common scenario. We use the average queue length in bytes as a representative performance indicator because of its robustness, but the conclusions stated here also apply to tail probabilities (as indicated at the end of this section).

Fig. 11 shows a plot of average queue length in bytes vs. utilization for DEC TCP trace, obtained by running the trace through a simulated queue with capacity adjusted so that utilizations are varied from 10%–60%. The average queue lengths increase sharply above 55%, but it is interesting to note that the queueing backlogs are observed at utilizations as low as 20%.[5] In most standard queueing models (with the exception of batch arrival models), queueing backlogs are negligible at such low utilizations.

In order to determine the role of fine timescale features in the traffic, the trace is aggregated to the level of 512 ms (equal to the rough lower cutoff to the FBM regime obtained in the previous section), and the total byte count in any 512 ms interval is distributed uniformly over the interval. The result of this is the second curve in Fig. 11, and is seen to be far too optimistic. Fig. 11 also shows the queueing performance obtained with a modified trace in which a worst-case scenario is used to deal with the fine timescale features in the traffic. This is done by calculating the maximum per-millisecond rate in the traffic trace, aggregating the traffic over intervals of 512 ms, and assigning all the traffic in each 512 ms interval to the beginning of the interval at this maximum rate. This was suggested as a very conservative estimate in [12]. As can be seen from the figure, this approximation, while duly serving as an upper bound for the average queue length, is a gross overestimate. Thus while high frequency features are important in determining performance at low utilizations, it is worth trying to find better approximations than this worst-case estimate.

If fine timescale features had not been important, the difference between the upper and lower bound curves in Fig. 11 would have been insignificant. Thus while it is *a priori* clear that these curves *are* bounds, the fact that they do so poorly shows that fine timescale features cannot be ignored. (In fact, both of these approximations have been used extensively in the literature.) To see whether a simple interpolation scheme would adequately take care of fine timescale features, obviating the use of multifractals, Fig. 11 also shows the average queue length with the packets in each 512 ms interval distributed assuming Poisson arrivals; this is seen to be almost indistinguishable from the uniform distribution curve, demonstrating that much higher variability on fine timescales is needed.

We next consider a multi-fractal interpolation method to generate variability and non-Gaussian features on fine timescales. As mentioned in Section II, there are several variations in multi-fractal construction methods. For the TCP traffic, we use the semi-random multi-fractal cascade (SRMC). This has the advantage of simplicity, in that only the single parameter p needs to be specified. The parameter p was chosen so as to match the scaling function $\tau(q)$ (defined in Eq. (3)) obtained from Fig. 4 most closely. Fig. 4 (graphs labeled DEC TCP and mfc) shows the result for $p = 0.65$, which is the value we use. The average queue length obtained by simulating the trace generated by this SRMC interpolation is the fourth curve shown in Fig. 11. As can be seen, the trace closely matches the characteristic obtained with the original trace, indicating that it captures the essential fine timescale features: a variability greater than that predicted by FBM, coupled with distinctly non–Gaussian marginals[6]. The use of the cascade construction offers a parsimonious description of multi-fractal scaling, which is otherwise highly parameterized. While it is certainly possible to use more complicated multi-fractal interpolations - and matching the observed local and global scaling behavior in traces does require additional complexity in the construction [18] - the close match seen in Fig. 11 suggests that a simple construction may be adequate for estimating (first order) performance. Note that the multi-fractal interpolation procedure is versatile enough to capture the full regime of possibilities considered in Fig. 11, with $p = 0.5$ corresponding to the uniform distribution, and $p = 1$ resulting in

[5]Increasing the utilization beyond 60% continues the trend seen in the figure.

[6]The high variability and non–Gaussian marginals obtained with the SRMC construction can be verified directly

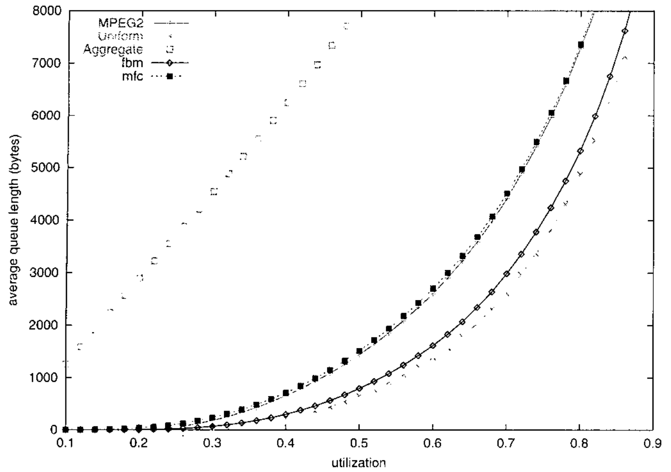

Fig. 12. The average queue length in bytes as a function of utilization for the measured video traffic (MPEG2), and various modifications thereof: i) a uniform interpolation (Uniform) of the aggregate traffic on timescales shorter than 40 ms ii) all the traffic in any 40 ms interval lumped at the beginning of the interval (Aggregate) iii) an fbm (fbm) interpolation below 40 ms and iv) a deterministic multifractal cascade (mfc) interpolation below 40 ms, with $p = 0.60$.

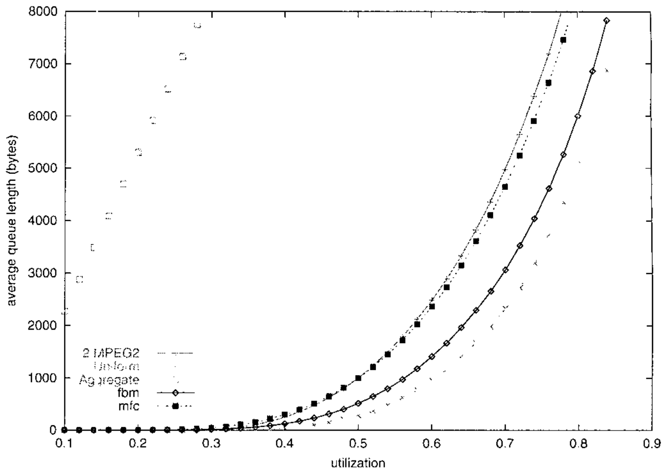

Fig. 13. The average queue length in bytes as a function of utilization for two superimposed MPEG2 streams, and various modifications thereof: i) a uniform interpolation (Uniform) of the aggregate traffic on timescales shorter than 40 ms ii) all the traffic in any 40 ms interval lumped at the beginning of the interval (Aggregate) iii) an fbm (fbm) interpolation below 40 ms and iv) a deterministic multifractal cascade (mfc) interpolation below 40 ms, with $p = 0.60$.

the batch construction.

Further, note that there is no contradiction in the fact that fine timescale structures which arise from flow control actions are most consequential at low utilizations. There is a misconception that flow controls are invoked or have impacts only under congestion at high utilizations. What causes the irregular distribution of packets over a round trip time is that the traffic source has to wait for acks; a high capacity TCP source will transmit a batch of packets (whose distribution is determined by the TCP window dynamics) and then wait for the rest of the round trip time for the acks to be returned. The "batchiness" arising from the TCP windowing contributes to significant queueing delays even at the lowest utilizations. In contrast, a UDP source will continue to transmit packets at its peak rate as long as there is data to transmit.

Similar experiments are performed for the MPEG2 video data. Fig. 12 shows a plot of average queue length in bytes vs. utilization for measured video traffic, obtained by running the trace through a simulated queue with capacity adjusted so that utilizations are varied from 10%-80%. Given that simulations are being done for a single video stream, a corresponding ATM networking scenario is: the VBR video stream is assigned to a single VC, and the capacity of the link corresponds to the bandwidth allocated to this VC in a per-VC queueing discipline. In most ATM switches, unallocated or unused capacity can also be used to serve this VC, so that the queueing backlogs studied here are upper bounds of the actual backlogs observed in a per-VC queueing system. The single stream simulation is also relevant for determining effective bandwidths, and for setting policing and shaping parameters. In this paper, the value of the single stream simulation is in identifying the statistical features of the traffic that determine performance. The average queue lengths increase sharply above 70%, and are significant even at the 50-60% utilization level. Interpreted as a per-VC backlog, average queues of the order of several kB are significant.

Interpolation schemes similar to those used for the TCP traffic in Fig. 11 are also tried for the video traffic. The trace is aggregated to the frame level, and the bytes in each frame are either distributed uniformly between the slices in the frame, or all assigned to the beginning of the frame. Both these interpolation schemes are used extensively in performance studies based on frame level measurements. The results of these are also shown in Fig. 12.

As was seen for the TCP data, the two interpolation schemes yield lower and upper bounds to the actual delay curve, but are poor estimates. To show that slightly more complicated interpolation schemes–without multifractals–that would be between these two bounds are not adequate, here we show the results for an FBM interpolation; this is not much better than the lower bound. Finally, the result of a cascade interpolation for the MPEG2 video traffic is shown in Fig. 12, except that as mentioned in the previous section, a deterministic cascade is used. The corresponding $\tau(q)$ curve is shown in Fig. 8. It can be seen that, as in the case of TCP traffic, the delay curve generated with the cascade interpolation agrees quite well with that obtained with the original data.

We recall at this stage that the video traffic data has an intermediate time regime beyond the fine timescale regime with complex cross–frame correlations. However, since the interpolation schemes discussed here work with frame–level aggregates, they preserve all the correlations or anticorrelations present at higher timescales without explicitly modeling them. A more ambitious approach might be to work with traffic aggregated to the group of frames (GOF) level (or to start with pure FBM traffic at this level, generated synthetically), interpolate to the frame level keeping in mind the manner in which MPEG2 distributes the different frames, and then proceed to a sub-frame level using multifractal (or other) approaches. (One might even consider doing the sub-frame interpolation differently for the different types of frames, although such an approach, if not strongly constrained by the details of MPEG2, can easily end up with too many free parameters to be of practical use.) We leave this for future work.

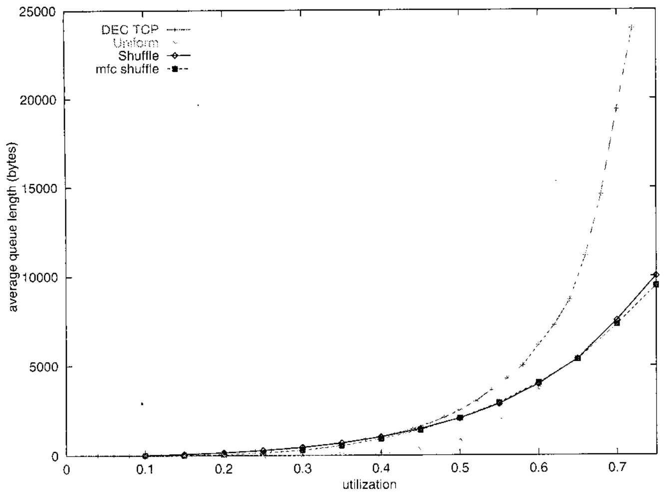

Fig. 14. The average queue length in bytes as a function of utilization for the DEC TCP trace with 512 ms intervals shuffled (Shuffle), thereby destroying all long time correlations, and for the measured DEC trace (DEC TCP). The uniform interpolation curve from Fig. 11 (Uniform) and a shuffled cascade trace (mfc shuffle) are also shown.

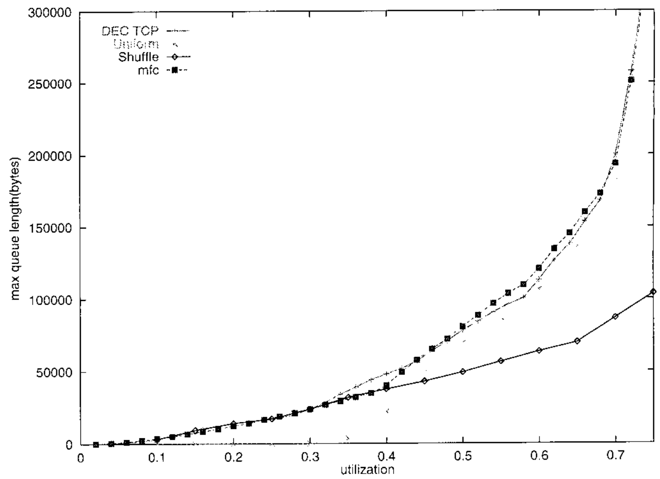

Fig. 15. The maximum queue length in bytes over a 5 minute time interval for TCP traffic, as a function of utilization. Curves corresponding to original traffic trace (DEC TCP), the trace with 512 ms blocks shuffled (Shuffle), and uniform (Uniform) and cascade (mfc) interpolations below 512 ms are shown.

For TCP traffic, this is not an issue, because the fine timescale regime connects directly to a long time Gaussian regime. A reasonable model would thus seem to consist of FBM on coarse timescales, matched to a fine timescale SRMC interpolation.

However, to see that the multi-scaling method is not sensitively dependent on byte variations due to a single MPEG2 trace, we combined two MPEG2 traces and did a similar analysis. The the average queue length is shown in Fig. 13. We observe the same behavior as that of a single MPEG2 trace, with the multifractal scaling showing the best fit.

So far we have focussed on the fine timescale features in the traffic. We now turn to the complementary question, of whether the long timescale correlations have any bearing on performance. We report results for TCP traffic only. We divide the DEC TCP trace into blocks of 512 ms and shuffle the blocks, while preserving the byte counts in each of the 1 ms subintervals of each block. This has the effect of destroying long time correlations and preserving fine timescale correlations. The average queue lengths for this shuffled trace and for the original DEC TCP trace are shown in Fig. 14. Note the significant deviation from the original queue length curve at utilizations above 0.5; this discrepancy becomes very large for utilizations above 0.65. Fig. 14 also reproduces the performance curve for uniform interpolation below 512 ms (from Fig. 11), demonstrating that long time and fine timescale correlations in the traffic affect performance in complementary regimes of heavy and light loading respectively. At intermediate utilization levels, where it would be reasonable to operate the network, *both* short and long time correlations are important. Finally, Fig. 14 shows the average queue length for a trace with blocks shuffled, and with the bytes within each block distributed according to a cascade. This curve is very close to the first curve in Fig. 14, emphasizing that the cascade construction is adequate for fine timescale features.

We conclude this section by briefly looking at the tail probability instead of the average for the queue length. This is more relevant to the analysis in the following section of this paper. Fig. 15 shows the maximum queue length achieved over a 5 minute interval, for the original TCP traffic trace and some of its modifications considered so far. It can be seen that the conclusions we have drawn from the average queue lengths about the regimes of importance of short and long time features remain valid.

V. PERFORMANCE ANALYSIS AND ENGINEERING

In the previous sections we observed the significant performance impacts of the fine timescale fluctuations in the TCP and video trace, and the fact these could be modeled reasonably by a multi-fractal cascade. We now examine whether such a description is analytically tractable to the extent that one can estimate relevant engineering metrics (such as queueing backlogs, safe operating points). In the remainder of this paper we show how one can analyze performance with a semi-random multifractal cascade description, and further consider its consequences for admissions control and capacity planning, using large-deviations theory.

Using standard large-deviations arguments [27], the loss probability at a link of capacity C and buffer size B that is driven by an arrival process $A(t)$ can be estimated by

$$\begin{aligned} P(Q > B) &\sim \sup_{t>0} P(A(t) > Ct + B) \\ &\approx \sup_{t>0} \exp[-\Lambda^*(Ct + B, t)] \end{aligned} \quad (5)$$

where Λ^* is the Legendre transform of the function Λ, defined as $\Lambda^*(x,t) = \sup_s (sx - \Lambda(s,t))$, and Λ is the logarithm of the moment-generating function of $A(t)$. If $A(t)$ is comprised of n independent streams $A_I(t)$, the function Λ can be expressed as [28]

$$\begin{aligned} \Lambda(s,t) &= n\Lambda_I(s,t) = n\log\left[E\big(\exp(sA_I(t))\big)\right] \\ &= n\log\left[\sum_{q=0}^{\infty} s^q E(A_I^q(t))/q!\right]. \end{aligned} \quad (6)$$

With a target loss probability of $\varepsilon = \exp[-\lambda]$, using Eq. (5) as an estimate of the loss probability yields the condition $\lambda \leq \inf_{t>0}\sup_s[s(Ct+B) - n\Lambda_I(s,t)]$. This can be inverted, to yield the maximum number N of sources that can be supported on the communication link:

$$N \sim \inf_{t>0} \sup_{s} \frac{s(Ct+B)-\lambda}{\Lambda_I(s,t)}. \quad (7)$$

Note that while estimates of performance measures such as loss rates can be far from their actual values using large-deviations-based methods, the corresponding engineering recommendations are typically more accurate. This is because near the operating point, small changes in traffic levels can have a large effect on the relevant performance measure; conversely, large discrepancies in performance estimates have nevertheless a small effect on the operating points.

Eq. (6) indicates analytically how multi-scaling is related to performance. First, the equations above show how performance is related to the moments of the arrival process at all timescales, especially at the performance-relevant timescale, which corresponds to the value of t that realizes the infimum in Eq. (7). Meanwhile, the basic multi-scaling relation Eq. (3) shows that on fine timescales, moments of arrivals over a timescale depend on the timescale through simple power laws. For a single traffic stream, these sums $\sum_i X_\Delta(i)^q$ are basically proportional to $E(A_I^q(\Delta))/\Delta$. These power laws allow one to express the time dependence of each of the various moments in the sum of Eq. (6) in terms of a scaling exponent and a constant pre-factor or intercept, rather than explicitly specifying a family of marginal distributions over a continuum of timescales. It should be stressed again that the multi-scaling relations apply only over fine timescales, so that one must resort to other methods (e.g., monofractal scaling) for coarser timescales. For the purposes of this discussion, we restrict ourselves to scenarios in which the performance-relevant timescale is within the fine-scale regime of the multi-scaling relations.

A multi-fractal characterization of the fine timescale behavior specifies the dependence of $E(A_I^q(\Delta))$ on Δ. As seen from Eq. (3), the family of functions $E(A_I^q(\Delta))$ (one function for each q) is replaced by $\tau(q)$ and $C(q)$. Although this is a great simplification, we still need τ and C for each q. This is a highly parameterized description, of limited use in practice. This motivates the development of more parsimonious constructions, such as the multi-fractal cascade discussed in Section III, that can generate fine-scale behavior from a given coarse-scale behavior in a way that depends on just a single parameter p. The cascade generator allows all the separate parameters $\tau(q)$ to be determined by the parameter p.

Eq. (6) also demonstrates that a multi-fractal characterization in terms of $\tau(q)$ or $f(\alpha)$ is only "half-complete" in that it provides the scaling exponents, but not the constant pre-factors $C(q)$ needed to evaluate the sum in (6). In analogy with an FBM model with parameters $\{m, a, H\}$, this is equivalent to providing the Hurst parameter H, but not the peakedness parameter a. A separate procedure must be used to fix the pre-factors $C(q)$. Based on the earlier observations of the adequacy of the FBM model for TCP data at coarse timescales, the procedure here is to adopt an FBM model as an approximation for these coarser scales. Such a procedure accomplishes two tasks: first it parsimoniously describes the coarse-scale behavior not modeled by the multi-fractal cascade; second, it provides a "boundary condition" or starting point from which we can compute the constant pre-factors.

To summarize, our approach is based on:

- finding analytical expressions for the scaling exponents of a multi-fractal cascade, which is feasible, at least for simple cascades;
- using the "boundary condition" of an FBM description over coarser timescales to determine the constant pre-factors.

This approach, of matching FBM at coarse timescales to a multi-fractal cascade at fine timescales, should be reasonable for modeling TCP traffic. For video, there is an intermediate time regime which has to be modeled as well.

Assuming that the cascade construction is applied at the level of individual streams, the moments of the MFC interpolated process over fine timescales are approximately given by:

$$E(A_{MFC}^q(t)) = E(A_{FBM}^q(\theta))(t/\theta)^{1-\tau(q)}, \quad (8)$$

where we assume that an FBM description is valid for timescales greater than θ, and a multi-fractal cascade is used below it. Eq. (8) is the condition of continuity of $E(A_I^q(t))$ across θ. In terms of the parameter p of the semi-random MFC, the function $\tau(q)$ is given by

$$\tau(q) = 1 + \log_2\left[\frac{p^q + (1-p)^q}{2}\right]. \quad (9)$$

The boundary condition at θ is obtained from

$$E(A_{FBM}^q(\theta)) = \sum_{k=0}^{k\leq q/2} \binom{q}{2k} (m\theta)^{q-2k}(am\theta^{2H})^k(2k-1)!! \quad (10)$$

where we have used the fact that $A_{FBM}(\theta) = m\theta + \sqrt{am}X_{FBM}(\theta)$, with $X_{FBM}(\theta)$ a normal random variable having zero mean and θ^{2H} variance.[7] Thus, in principle, given a description of the traffic in terms of a coarse timescale FBM model, with a multi-fractal cascade generator of fine timescale fluctuations, one can estimate several performance measures.

The values of the parameters for this analytical model are taken to match the observations for TCP traffic noted in the previous sections. First, a transition between self similarity and multifractality was found to occur at $\theta = 512$ ms. At coarser timescales, the traffic exhibited features of an FBM with $m = 92.0$, $a = 204$, and $H = 0.8$, in units of bytes and milliseconds. At finer timescales, the structure function matched that of a semi-random cascade with $p = 0.65$. Accordingly, these are the parameter values used in the analytical calculations.

We will now compare the prediction of Eq. (7) for the number of sources N of TCP flows that can be supported, using either an exactly self-similar model (FBM), or an asymptotically self-similar model whose fine timescale behavior is described by the multi-fractal cascade of Section III. In order to do this, we have

[7] $(2k-1)!!$ is the double factorial: $(2k-1)!! = (2k-1)(2k-3)\cdots 1$.

Table 1. Allowed occupancy on a link as a function of the traffic model and the link speed.

link speed	multifractal model	FBM model
T1 (1.54 Mb/s)	0.228	0.341
E1 (2.05 Mb/s)	0.273	0.396
Ethernet (10 Mb/s)	0.636	0.684
DS3 (43 Mb/s)	0.854	0.854
OC-1 (49.5 Mb/s)	0.866	0.866
OC-3 (150 Mb/s)	0.930	0.930
OC-6 (300 Mb/s)	0.954	0.954
OC-12 (601 Mb/s)	0.970	0.970

to choose specific values for various network parameters. The tolerance for losses is taken to be $\varepsilon = 10^{-4}$. We explore a range of link speeds from T1 to OC-12 (corresponding to payload rates ranging from 1.54 to 601 Mb/s). Regardless of the link speed, the buffer size is taken to be 50 ms, so $B = 50C$ bytes if the link speed C is expressed in bytes/ms.

Table 1 compares the multifractal and FBM models for the operating-point predictions they produce for these scenarios. The horizontal line through the table separates the cases in which the multi-fractal model predicts safe operating points that are significantly different from the FBM model. As can be seen from Table 1, when the link rate is relatively small, comparable to the peak rate of the stream whose traffic was characterized, the multifractal model shows the safe operating point to be at a significantly lower occupancy than would be suggested by the FBM model. For these scenarios, the relevant timescales of queuing are below the 512 ms threshold, and the deviations are explained by the non-Gaussian marginals and change in the scaling exponent below the threshold. As link speeds increase, the multi-scaling features become less relevant for two reasons. First, these scenarios correspond to the aggregation of a large number of sources so that the marginals of the aggregate are well represented by a Gaussian distribution. Secondly, the relevant timescales increase beyond the lower cutoff, as the buffer size is scaled proportionally to the link rate, and utilizations increase. Intuitively, the timescale, at which traffic fluctuations are most likely to overflow a buffer, should be proportional to the size $B/(C(1-U))$ of the buffer expressed as a time for the excess capacity $C(1-U)$ to serve the fluctuations, where U is the link's utilization. (For purely FBM arrivals, this intuition is accurate and the proportionality constant is $H/(1-H)$.) Based on this intuition, when buffer sizes are scaled to keep B/C constant, multiplexing gains should allow U to rise and hence the performance-relevant timescale to rise as well. Eventually, the performance-relevant timescale crosses from the fine timescales of the multi-scaling regime to the coarse scales of the self-similar regime.

Finally, we will briefly consider the impacts of closed loop, flow control feedback on these conclusions. The traffic offered by a source will decrease as network delays increase, so that the number of sources that can be supported for a given utilization will increase. Secondly, safe operating points should also increase. Our experience indicates that the utilizations at which open loop models predict poor performance (e.g., high losses) roughly coincide with regions of poor performance (e.g., low throughputs caused by excessive delays) predicted by closed loop models. An alternative to the open loop models studied here is to perform detailed network simulations based on TCP level source models that incorporate the details of TCP dynamics. However, in practice, from the viewpoint of network operations such detailed models are not usable as many of the parameters that describe TCP source behavior cannot be estimated from network level measurements. For this reason, IP flow level models (e.g., self similar or multi-scaling descriptions) that can be parameterized on the basis of network level measurements are useful. In fact, much of network engineering today is done on the basis of open loop analytical and simulation models of network elements. Models that additionally incorporate fine timescale fluctuations in traffic can be expected to provide more accurate engineering rules over a wider range of scenarios. Finally, it is expected that the open loop traffic description framework will still be valid with network congestion, but the values of some of the specific parameters (e.g., the cascade parameter p) will change, while FBM parameters associated with the asymptotic long-time behavior will be unaffected [8]. The longer term challenge is to develop phenomenological models that can predict these changes.

VI. CONCLUSIONS AND FURTHER WORK

In this paper we have investigated a structural modeling approach to describe the fine timescale behavior in wide area TCP and MPEG2 video traffic. This is a problem of engineering importance, as demonstrated by the performance impacts of these fine timescale fluctuations studied in this paper. The following summarizes our findings:

- Traffic has fine timescale features that cannot be described by simple Gaussian self-similar models earlier shown to be adequate for LAN traffic. This conclusion is based on the specific WAN TCP and MPEG video traces analyzed in this report, as well as many more similar traces not explicitly discussed here. The conclusion confirms recent studies on other WAN TCP traces [18], [19].
- These fine timescale features in the traffic adversely affect performance at low utilizations. Long time correlations, which were shown to be very important for LAN traffic performance [8], continue to affect performance at high utilizations.
- For WAN TCP and MPEG video, the fine timescale features fit reasonably to a multifractal cascade description, in accordance with recent work [18], [19]. In both cases, aggregating traffic to the "edge" of the fine timescale regime and interpolating down with a multifractal cascade matches performance (as measured by trace-driven simulations) very well.
- (a) For TCP, the complex behavior in the fine timescale regime is believed to be caused by TCP flow controls that operate within the round-trip time (RTT) of packet receipt acknowledgements [21]. This short time regime gives way to FBM behavior when aggregated to the RTT, reminiscent of LAN traffic.

  (b) For video traffic, the fine timescale behavior occurs at the sub-frame level, and is due to the intrinsic properties of MPEG coding scheme. Between the subframe and group of frames (GOF) levels, the coding cycles between different types of frames in the well-known pattern of MPEG2. This does not follow any simple scaling law, but has been described by Markovian models [11], [12]. For time scales greater than the GOF level, one sees FBM behavior.

- For traffic that can be modeled by FBM at long timescales and a multifractal at fine timescales (as seems to be the case for TCP traffic) we obtain analytical performance estimates and engineering parameters, such as the number of sources that can be admitted given QoS parameters. The analytical treatment also shows that the usual multifractal description in terms of (a family of) scaling exponents is incomplete and must be supplemented by estimating prefactors/amplitudes. Unfortunately, this is an aspect of statistical inference that is currently overlooked. Much attention is paid to estimate scaling exponents, while what is really required is the joint estimation of scaling exponents and prefactors.
- (a) For WAN TCP, where the fine timescale features are believed to arise from network feedback, a more complete description would consist of modeling TCP feedback explicitly, and generating these features from the flow control in closed loop models. However, treating the fine timescale parameters as being given a priori, as we have done, is a useful first step in analysis.
  (b) For MPEG video, where the fine timescale features arise from the coding mechanism, it is not necessary to incorporate feedback. It would be useful to relate the fine timescale behavior directly to specific features of MPEG coding. This is a subject for further research, along with the modeling of the intermediate regime between the fine and coarse timescales.
- For WAN UDP traffic, which has neither complex coding nor network feedback effects, a simple FBM description is found to be adequate for both traffic characterization and engineering analysis.

There are numerous avenues to expand this work, including repeating these experiment with additional data and video traces (in progress); implementing the numerical methods to analyze traffic generated by multi-fractal cascades; and extensions to analyze and describe packet traffic over the full range of engineering timescales of interest. In the longer term, the objective is to develop robust, tractable, parsimonious and predictive network traffic models and management methods that are usable in practice.

ACKNOWLEDGMENTS

We thank Walter Willinger for useful discussions, and Amarnath Mukherjee [25] for kindly supplying us with the video data for this work. The TCP data was obtained from http://ita.ee.lbl.gov/html/contrib/DEC-PKT.html, thanks to the Digital Equipment Corporation.

REFERENCES

[1] K. Park and W. Willinger *Self-similar network traffic and performance evaluation,* For a review see (John Wiley, 2000).
[2] W. E. Leland *et al.*, "On the self-similar nature of ethernet traffic (Extended Version)", *IEEE/ACM Trans. Networking*, vol. 2, no. 1, pp. 1–15, Feb. 1994.
[3] J. Beran *et al.*, "Long-range dependence in variable-bit-rate video traffic," *IEEE Trans. Commun.*, vol. 43, no. 2–4, pp. 1566–1579, Apr. 1995.
[4] M. E. Crovella and A. Bestavros, "Self-similarity in world wide web traffic: Evidence and possible causes," *IEEE/ACM Trans. Networking*, vol. 5, no. 6, pp. 835–846, Dec. 1997.
[5] D. E. Duffy *et al.*, "Statistical analysis of CCSN/SS7 traffic data from working subnetworks," *IEEE J. Select. Areas Commun.*, vol. 12, no. 3, pp. 544–551, 1994.
[6] A. Erramilli and J. L. Wang, "Monitoring packet traffic levels," in *Proc. IEEE GLOBECOM*, pp. 274–280, San Francisco, CA, 1994.
[7] G. Iacovoni, V. Manca, and D. Vergni,"Single source TCP behaviour: A multifractal analysis," in *Proc. ICC/IEEE*, San Francisco, CA, 2000.
[8] A. Erramilli, O. Narayan, and W. Willinger, "Experimental queueing analysis with long-range dependent packet traffic," *IEEE/ACM Trans. Networking*, vol. 4, no. 2, pp. 209–223, Apr. 1996.
[9] K. R. Krishnan, A. L. Neidhardt and A. Erramilli, "Scaling analysis in traffic management of self-similar processes," in *Proc. ITC-15*, vol. 2, pp. 1087–96, June 1997.
[10] P. Pancha and M. El-Zarki, "MPEG coding for variable bit rate video transmission," *IEEE Commun.*, vol. 32, no. 5, pp 54–66, May 1994.
[11] D. P. Heyman and T. V. Lakshman, "Source models for VBR broadcast-video traffic," *IEEE/ACM Trans. Networking*, vol. 4, no. 1, pp. 40–48, Feb 1996.
[12] A. Elwalid *et al.*, "Fundamental bounds and approximations for ATM multiplexers with applications to video teleconferencing," *IEEE J. Select. Areas Commun.*, vol. 13, no. 6, pp. 1004–1016, Aug. 1995.
[13] I. Norros, "A storage model with self-similar input," *Queueing Systems*, vol. 16, no. 3-4, pp. 387–396, 1994.
[14] D. R. Cox, "Long-range dependence: A review," Statistics: An Appraisal, in *Proc 50th Anniversary Conf.*, ed by David and David, Iowa State University Press, 1984.
[15] T. G. Kurtz, "Limit Theorems for Workload Input Models," in Stochastic Networks: Theory and Applications, ed. F. P. Kelly, S. Zachary and I. Ziedins, Oxford University Press, 1996.
[16] M. S. Taqqu, W. Willinger, and R. Sherman, "Proof of a fundamental result in self-similar traffic modeling," *Computer Commun. Review*, vol. 27, pp. 5–23, 1997.
[17] W. Willinger, V. Paxson, and M. S. Taqqu, "Self-similarity and heavy tails: Structural modeling of network traffic," A Practical Guide to Heavy Tails: Statistical Techniques and Applications, R. Adler, R. Feldman, and M. S. Taqqu (Eds.) Birkhauser, Boston MA, 1998.
[18] A. Gilbert, A. Feldmann, and W. Willinger, "Data networks as cascades: Explaining the multifractal nature of Internet WAN traffic," in *Proc. ACM SIGCOMM*, pp. 42–55, 1998.
[19] R. H. Reidy and J. Levy Vehel, "TCP traffic is multifractal: A numerical study," INRIA Research Report no. 3129, Preprint 1997.
[20] I. Saniee, "Multiscaling in low-aggregate fast packet network traffic," Bellcore Report, Sept. 22, 1998 (available by request from Telcordia Technologies Inc.).
[21] A. Feldmann *et al.*, "Dynamics of IP traffic: A study of the role of variability and the impact of control," *Proc. ACM SIGCOMM*, pp. 301–313, 1999.
[22] A. Erramilli *et al.*, "Performance impacts of multi-scaling in wide area TCP-IP traffic," to be published in Infocom 2000.
[23] A. Erramilli *et al.*, unpublished.
[24] Trace due to Digital Equipment Corporation, collected by Jeff Mogul of Digital's Western Research Lab. This trace and others are available at http://ita.ee.lbl.gov/html/contrib/DEC-PKT.html. The TCP packets from the first one hour trace at this site were analyzed. The UDP trace traces analyzed also correspond to the first hour trace at this site.
[25] This data was supplied by Amarnath Mukherjee of Knoltex Corporation.
[26] K. J. Falconer, "Techniques in fractal geometry," Wiley, New York 1997.
[27] J. A. Bucklew, *Large Deviation Techniques in decision, simulation, and Estimation*, Wiley, New York, 1990.
[28] $\Lambda(s,t)/st$ is also referred to as the effective bandwidth for the traffic; see for instance F. Kelly, "Notes on effective bandwidths," in *Stochastic Networks: Theory and Applications*, ed. F. P. Kelly, S. Zachary and I. Ziedins, Oxford University Press, 1996.

Ashok Erramilli received his M.S. and Ph.D. degrees in Electrical Engineering from Yale University in 1981, and 1984, respectively. Prior to this, he obtained two bachelor's degrees, one in Physics from Poona University, India, in 1977, and a degree in Electrical Communications Engineering from the Indian Institute of Science in 1980. He joined ITT's Advanced Technology Center in 1983, where he developed congestion and flow control strategies for a high performance packet switch. He joined Bellcore's Traffic Analysis groups in 1987, where he worked on the management of traffic flows across the full range of communications networks, with emphasis on the characterization, engineering and control of packet traffic. In 1995, he became manager of Bellcore's network design and traffic analysis research group, and with overall responsibilities for providing Bellcore's client companies the traffic management methods needed to operate high-speed Frame Relay, ATM and IP networks. He founded QNetworx in 1997, offering research, consulting and software development services in the areas of network traffic management and control. He has been a Visiting Professor at the Royal Institute of Technology in Stockholm, Sweden between 1993-1996, and at the Royal Melbourne Institute of Technology in Melbourne, Australia (1999). He has authored over 35 papers in the field of network traffic.

Onuttom Narayan received his M.Sc degree in physics from the Indian Institute of Technology in 1987, and his Ph.D in physics from Princeton University in 1992. He was at Harvard University and Bell Laboratories from 1992 to 1995. He joined the University of California at Santa Cruz as an assistant professor in 1995, and became an associate professor in 1999. His interests are generally in statistical mechanics of disordered and nonequilibrium systems. He has several papers on network traffic.

Arnold Neidhardt is a Senior Scientist at the Network Design and Traffic Research Group at Telcordia technologies.

Iraj Saniee manages the Mathematics of Networks and Systems at the Mathematical Sciences Research Center at Bell Labs, Lucent technologies. He has been at Bell Labs since 1998 during which he has worked on a variety of modeling and optimization problems arising in performance analysis and optimization of next generation networks. These include multiscaling and limiting models of data traffic, optimization of optical networks and more recently scaling models for estimating parameters in optical networks. Prior to Bell Labs, he worked for 13 years at Bellcore as, Member of Technical Staff, Distinguished Member of Technical Staff, Senior Scientist and Director of the Network Design and Traffic Research Group, in the Network Technologies and Applies Research laboratories. He pioneered the development of modern versions of network design tools subsequently used by the Bellcore client companies that employed domain-specific optimization and heuristics for large-scale problems associated with dimensioning and survivability planning for SONET/SDH and ATM networks. He was the leader of the 1994 runner-up team for INFORMS Edelman International Prize in Operations Research and has published over 35 articles in IEEE, SIAM, and INFORMS journals and proceedings. He received his B.A. and M.A. in Math and Ph.D. in Operations Research all from Cambridge University.