

Multiscale Queuing Analysis

Vinay J. Ribeiro,[†] Rudolf H. Riedi,^{†‡} Richard G. Baraniuk[†]

Department of Electrical and Computer Engineering[†] and Department of Statistics[‡]
Rice University, 6100 Main Street, Houston, TX 77005, USA

Abstract— We develop a new approach to queuing analysis for an infinite-length queue with constant service rate fed by any traffic process. Our approach provides important theoretical results, is easy to implement in practice, and is particularly relevant to queues fed with long-range-dependent (LRD) traffic. We use traffic statistics at only a small fixed set of time scales and develop three approximations for the tail queue probability. These are non-asymptotic, that is they apply to any finite queue threshold. Simulations with LRD traffic models and Internet traces demonstrate their accuracy. Besides non-asymptotic error bounds, asymptotic decay rates, and error bounds for the approximations, we prove an optimality property of exponential time scales. Simulations reveal that the second-order correlation structure of traffic by itself does not determine queuing behavior and that the tails of traffic marginals at different time scales have a strong impact on queuing.

I. INTRODUCTION

Maintaining low packet queuing delays and jitter at router queues is critical for the viability of real-time streaming media applications for telephony, telemedicine, video-conferencing, economic transactions etc. As a result admission control [2], network provisioning [3] and other techniques for reducing queuing delay have gained importance. These applications typically require accurate predictions of queuing delays in order to succeed. We model a router queue as an infinite length queue with constant service rate [4] and study the probability that the queue size Q exceeds a threshold b , $\mathbb{P}\{Q > b\}$, also called the *tail queue probability*.

We can predict $\mathbb{P}\{Q > b\}$ in several ways. First, we can model network traffic using different processes (also called *traffic models*) and use any exact formula for $\mathbb{P}\{Q > b\}$ that is available. Second, in case exact results are unavailable for a particular process we can employ analytical results that only approximate $\mathbb{P}\{Q > b\}$, which we call *queuing approximations*. Third, if modeling traffic with a standard random process is cumbersome or inadequate then we can predict $\mathbb{P}\{Q > b\}$ directly from measured traffic statistics. In such a scenario it is desirable to use few traffic statistics in order to reduce data acquisition and computational requirements.

In this paper we develop a novel approach to queuing analysis called the *multiscale queuing analysis* that ad-

dresses the second and third scenarios mentioned above. It uses few traffic statistics and is practical.

While our analysis is relevant to any traffic process we focus on processes with non-summable correlations (called long-range-dependence (LRD)) since LRD is a ubiquitous property of real-traffic [5]. Classical Poisson and Markov queuing techniques do not apply to the queuing analysis of LRD traffic which creates the need for new analytical tools. Up to now exact formulas for the queuing delay of LRD processes, other than for asymptotically large delays [6–8], have not been found and we are thus forced to use approximations.

To date, most approximations for the tail queue probability of queues fed with LRD processes have been based on the notion of the *critical time scale* [6, 7, 9–12]. Given a queue size threshold b , the critical time scale is the most likely amount of time it takes for the queue to fill up beyond b . While the critical time scale is a powerful theoretical tool, computing it directly from empirical measurements is impractical because this requires traffic statistics at *all* time scales.

By using traffic statistics at only a finite set of time scales, θ , our approach provides three practical approximations for $\mathbb{P}\{Q > b\}$: the *max* approximation, the *product* approximation, and the *sum* approximation. These have several important features:

- they apply to any finite queue threshold b , that is, they are non-asymptotic;
- they apply to any traffic model including non-stationary ones; and
- they are simple to employ because they require traffic statistics only at few time scales θ .

We prove numerous non-asymptotic error bounds, large-queue asymptotic results, and other bounds for the three approximations for different LRD traffic models including fractional Brownian motion (fBm), the wavelet-domain independent Gaussian model (WIG), and the multifractal wavelet model (MWM). We also compare the different approximations through numerical experiments.

It is of significant practical importance to determine an appropriate candidate for the time-scale set θ at which to collect statistics. The choice of θ involves a trade-off between the accuracy of the max approximation and the requirements for computation and data acquisition. A sparse θ decreases the accuracy of the max approximation but simultaneously requires the computation of traffic statistics and data acquisition at fewer time scales.

We prove that the choice of *exponential* time scales for

Supported by NSF grants ANI-9979465, ANI-0099148, and ANI-0338856, DoE SciDAC grant DE-FC02-01ER25462, DARPA/AFRL grant F30602-00-2-0557, and the Texas Instruments Leadership University program.
Email: {vinay, riedi, richb}@rice.edu; web: spin.rice.edu. A conference version of this paper with preliminary results was presented at INFOCOM 2000 [1].

θ is *optimal* with respect to this trade-off for a queue with fBm input traffic. A significant advantage of exponential time scales is their sparsity; just a few exponential time scales span a wide range. This result thus strongly recommends the use of traffic statistics at exponential time scales in queuing applications.

Traffic models such as fBm and the WIG suffice to capture the queuing behavior of traffic in Gaussian scenarios which can occur with traffic aggregation as on backbone links [11, 13]. However they do not perform as well in non-Gaussian traffic scenarios. We term the distributions of traffic at different time scales as *marginals*. Through simulations with the WIG and the MWM which have different marginals we demonstrate the strong impact of marginals on queuing. This result supports similar findings in [12, 14].

Our main contributions are thus (i) a novel multiscale approach to queuing analysis that provides practical queuing approximations, (ii) optimality, error bounds, and convergence results related to the approximations for various traffic models, and (iii) the demonstration that marginals can strongly influence queuing behavior. Additional results can be found in [15].

Paper Organization: Section II reviews previous work on the critical time scale. In Section III we present the multiscale queuing analysis of the paper and derive the various queuing approximations. Section IV describes the fBm, the WIG, and the MWM traffic models. In Section V we prove the optimality of exponential time scales for fBm. Section VI proves that large buffer asymptotic results and Section VII proves bounding results for the different queuing approximations. Section VIII demonstrates the accuracy of the approximations through simulations with Internet and synthetic model data and also demonstrates the impact of marginals on queuing. We conclude in Section IX.

Table I provides a list of symbolic notation to help the reviewers verify the mathematical proofs. It is not a part of the paper. More detailed versions of the proofs presented here can be found in [15].

II. REVIEW OF CRITICAL TIME SCALE ANALYSIS

In this section we review previous work on the critical time scale queuing analysis to set the stage for our multiscale queuing analysis in subsequent sections.

A. Queue size as a multiscale function

Consider a continuous-time fluid queue with constant service rate c with traffic process X_τ , $\tau \in \mathbb{R}$ as input. We refer to

$$K_\tau[t] := \int_{\tau-t}^{\tau} X_\omega d\omega \quad (1)$$

as the traffic process at *time scale* t . To avoid notational ambiguity we occasionally add superscripts such as in $K_\tau^{\{X\}}[t]$ to identify the traffic process. For the ease of notation we drop the subscript τ for all time-invariant quantities.

Assuming that the queue was empty at some time instant prior to τ , the *queue size* Q_τ equals the difference between

the total traffic that arrived at the queue and the total traffic serviced since the time instant the queue was last empty. This is succinctly captured by Reich's formula [16]

$$Q_\tau := \sup_{t>0} (K_\tau[t] - ct). \quad (2)$$

It is easily shown that the supremum in (2) always occurs at a value of t such that $\tau - t$ is the instant the queue was last empty. We address the requirement of an empty queue prior to τ with mathematical rigor in Section IV-D.

A key interpretation of (2) is that Q_τ equals a function of $K_\tau[t]$, the traffic process at all time scales t . The question arises as to whether or not we can accurately approximate $\mathbb{P}\{Q > b\}$ using the distribution of $K_\tau[t]$ at a *single* time scale t .

B. Critical time scale queuing approximation

The most widely proposed approximations of $\mathbb{P}\{Q > b\}$ for queues fed by LRD traffic are indeed based on a single time scale called the *critical time scale* [6, 7, 9–12], defined as

$$\lambda_\tau(b) := \arg \sup_{t>0} \mathbb{P}\{K_\tau[t] - ct > b\}. \quad (3)$$

We term the associated queue tail approximation the *critical time scale approximation*, which is defined as

$$\begin{aligned} C_\tau(b) &:= \sup_{t>0} \mathbb{P}\{K_\tau[t] - ct > b\} \\ &= \mathbb{P}\{K_\tau[\lambda_\tau(b)] - c\lambda_\tau(b) > b\}. \end{aligned} \quad (4)$$

Clearly $C_\tau(b)$ is a lower bound of $\mathbb{P}\{Q_\tau > b\}$ since by (2) $K_\tau[\lambda_\tau(b)] - c\lambda_\tau(b) \leq Q_\tau$; thus

$$C_\tau(b) \leq \mathbb{P}\{Q_\tau > b\}. \quad (5)$$

Earlier work based on large deviation theory has shown that $C_\tau(b)$ has the same log-asymptotic decay as $\mathbb{P}\{Q_\tau > b\}$ when $b \rightarrow \infty$ for a large class of input traffic processes including fBm [6, 7]. As the simulations in Section VIII demonstrate, $C_\tau(b)$ is also a good approximation for $\mathbb{P}\{Q_\tau > b\}$ for any finite b for fBm-fed queues. The intuition for the accuracy of $C_\tau(b)$ is that “*rare events occur in the most likely way.*” In other words given that $\{Q_\tau > b\}$ is a rare event, if the queue size is conditioned to fill up greater than b then it does so in time $\lambda_\tau(b)$ in which this is most likely, that is conditioned on $\{Q_\tau > b\}$ we have that Q_τ is approximately equal to $K[\lambda_\tau(b)] - c\lambda_\tau(b)$.

While the critical time scale is a powerful tool that has advanced the state-of-the-art in queuing theory, using it in practice is not straightforward. First, consider the problem of computing $C_\tau(b)$ for a queue fed with an arbitrary process, solely from empirical traffic measurements. From (4) we see that we require the distribution of $K_\tau[t]$ for all possible t . This is infeasible to obtain empirically. Similar problems of computation may arise when $C_\tau(b)$ is given by an analytical representation of the distribution of $K_\tau[t]$ for all t . Second, say that we wish to compute the critical time scale approximation when two independent processes X and Y are multiplexed and input to a queue. Such a

scenario often arises in admission control [2] and network provisioning [3]. Obtaining $C_\tau^{\{X+Y\}}(b)$ directly from the statistics of X and Y is again fraught with problems similar to those mentioned above.

III. MULTISCALE QUEUEING APPROXIMATIONS

In this section we develop three new queuing approximations that use traffic statistics only at a fixed finite set of time scales $\theta \subset \mathbb{R}_+$. By restricting analysis to the time scales θ , we overcome the practical problems associated with using the critical time scale approximation described above. Note that while some of our theoretical results are for countably infinite sets θ , in practice we always employ a truncated, finite set θ when computing the queuing approximations. We typically choose the set θ to span the range of time scales in which we expect the critical time scale $\lambda(b)$ to lie, for values of b relevant to a particular application. We now present the three queuing approximations.

A. Max approximation

In analogy to the queue size formula and the critical time scale (see (2) and (3)) we define

$$Q_\tau^{[\theta]} := \sup_{t \in \theta} (K_\tau[t] - ct) \quad (6)$$

and

$$\lambda_\tau^{[\theta]}(b) := \arg \sup_{t \in \theta} \mathbb{P}\{K_\tau[t] - ct > b\}. \quad (7)$$

This leads to the *max* approximation

$$\begin{aligned} M_\tau^{[\theta]}(b) &:= \sup_{t \in \theta} \mathbb{P}\{K_\tau[t] - ct > b\} \\ &= \mathbb{P}\left\{K_\tau\left[\lambda_\tau^{[\theta]}\right] - c\lambda_\tau^{[\theta]} > b\right\}. \end{aligned} \quad (8)$$

Comparing (4) and (8) we see that the max approximation is similar to the critical time scale approximation with the difference that the supremum is taken over a finite set in (8) instead of over all time scales as in (4). From (4), (5), and (8) we have the bounds

$$M_\tau^{[\theta]}(b) \leq C_\tau(b) \leq \mathbb{P}\{Q_\tau > b\}. \quad (9)$$

We note from (2) and (6) that

$$Q_\tau = Q_\tau^{[\mathbb{R}^+]} \geq Q_\tau^{[\theta]} \quad (10)$$

and from (6), (8) and (10) that

$$M_\tau^{[\theta]}(b) \leq \mathbb{P}\left\{Q_\tau^{[\theta]} > b\right\} \leq \mathbb{P}\{Q_\tau > b\}. \quad (11)$$

The max approximation provides a practical replacement for $C_\tau(b)$. Since the max approximation requires estimates of $\mathbb{P}\{K_\tau[t] - ct > b\}$ only for $t \in \theta$, the difficulties associated with computing $C_\tau(b)$ as we described earlier do not arise. First, consider the problem of obtaining the max approximation from empirical measurements. We simply compute histograms of the traffic at time scales $t \in \theta$ and then estimate $\mathbb{P}\{K_\tau[t] - ct > b\}$. Second, consider the

problem of computing the max approximation when two independent processes X and Y are multiplexed and input to a queue. By simply convolving the distributions of $K_\tau^{\{X\}}[t]$ and $K_\tau^{\{Y\}}[t]$ for $t \in \theta$ we obtain the corresponding distributions of $K_\tau^{\{X+Y\}}[t]$ which immediately give the max approximation.

B. Product and sum approximations

Two additional approximations of $\mathbb{P}\{Q_\tau > b\}$ based on the set of time scales θ are the *product* approximation

$$P_\tau^{[\theta]}(b) := 1 - \prod_{t \in \theta} \mathbb{P}\{K_\tau[t] - ct < b\} \quad (12)$$

and the *sum* approximation

$$S_\tau^{[\theta]}(b) := \sum_{t \in \theta} \mathbb{P}\{K_\tau[t] - ct > b\}. \quad (13)$$

Note that the product approximation equals $\mathbb{P}\left\{Q_\tau^{[\theta]} > b\right\}$ if the events $\{K_\tau[t] - ct > b\}$, $t \in \theta$, are independent,¹ and that the sum approximation equals $\mathbb{P}\left\{Q_\tau^{[\theta]} > b\right\}$ if the same events are mutually exclusive.

C. Intuition for the accuracy of the approximations

The max, product, and sum approximations inherit the accuracy of the critical-time scale approximation while being practical. If there exists an element of θ close enough to the critical time-scale then $M_\tau^{[\theta]}(b)$ will be close to $C_\tau(b)$ (see (4) and (8)). Moreover, if a single probability term dominates the summation in (13), then the product and sum approximations will closely approximate $M_\tau^{[\theta]}(b)$ and hence $C_\tau(b)$. Simulations demonstrate that the product and sum approximations are often closer to $\mathbb{P}\{Q_\tau > b\}$ than the max approximation.

In subsequent sections we study several issues related to the three approximations. Of particular interest is the optimal choice of θ in terms of size that gives a max approximation of certain accuracy. This is the topic of Section V. In Sections VI and VII we compare the max, product, and sum approximations against $\mathbb{P}\{Q_\tau > b\}$ as well as against $\mathbb{P}\left\{Q_\tau^{[\theta]} > b\right\}$ for queues with input from different traffic models.

IV. TRAFFIC MODELS

This section describes three traffic models that we focus on in this paper. These have been shown to model different aspects of real Internet traffic well.

A. Fractional Brownian motion

Fractional Brownian motion (fBm) is the unique Gaussian process with stationary increments and the following scaling property for all $a > 0$, $\tau \in \mathbb{R}$ [17]

$$B_{a\tau} \stackrel{d}{=} a^H B_\tau. \quad (14)$$

¹If events E_i , $i \in \mathbb{N}$, are independent then so are their complements.




Fig. 1. Multiscale tree representation of a traffic trace. Nodes at each horizontal level in the tree correspond to the sum (aggregates) of the process in non-overlapping blocks of sizes of powers of two, with lower levels corresponding to smaller block sizes. Each node is the sum of its two child nodes.

The symbol “ $\stackrel{d}{=}$ ” denotes equality in distribution. While its increment process *fractional Gaussian noise* (fGn) is stationary, fBm is itself non-stationary by definition. Denote the stochastic differential of B_τ as $\Delta_\tau B$. We denote fGn by

$$G_\tau[t] := K_\tau^{\{\Delta_\tau B\}}[t] = B_\tau - B_{\tau-t}. \quad (15)$$

While it is difficult to define $\Delta_\tau B$ rigorously, its aggregate $K_\tau^{\{\Delta_\tau B\}}[t]$ is well defined. Often one is interested only in the time series $\{G_{it'}[t']\}_{i \in \mathbb{Z}}$ with t' a constant time lag. From (14) and (15) we have that

$$K_\tau^{\{\Delta_\tau B\}}[t] \stackrel{d}{=} B_t \stackrel{d}{=} t^H B_1, \quad (16)$$

thus

$$\text{var}(G_{it'}[t']) = \text{var}\left(K^{\{\Delta_\tau B\}}[t']\right) = \sigma^2(t')^{2H} \quad (17)$$

where $\sigma^2 = \text{var}(B_1)$. When $1/2 < H < 1$, fGn possesses LRD. We use “var”, \mathbb{E} and “cov” to denote *variance*, *expectation*, and *covariance* respectively.

B. Wavelet-domain independent Gaussian (WIG) model

The WIG is a Gaussian model that is able to approximate fBm and fGn as well as processes with a more general scaling than (14) and (17). It uses a multiscale tree to model traffic over the time interval $[0, T]$ [18,19]. The nodes $V_{j,k}$ on the multiscale tree correspond to the total traffic in the time interval $[k2^{-j}T, (k+1)2^{-j}T]$, $k = 0, \dots, 2^j - 1$ (see Fig. 1).

Starting at node $V_{j,k}$, the WIG models its two child nodes $V_{j+1,2k}$ and $V_{j+1,2k+1}$ using independent *additive* random innovations $Z_{j,k}$ through

$$\begin{aligned} V_{j+1,2k} &= (V_{j,k} + Z_{j,k})/2, \\ V_{j+1,2k+1} &= (V_{j,k} - Z_{j,k})/2. \end{aligned} \quad (18)$$

In practice one uses a WIG tree of finite depth n to obtain a discrete-time process $V_{n,k}$. The $Z_{j,k}$ have the same variance within each scale j , thus guaranteeing that $V_{n,k}$ is a first-order stationary process. The tree-root $V_{0,0}$ and $Z_{j,k}$ are Gaussian which ensures the Gaussianity of all tree nodes.

To *fit* a traffic model means to choose its parameters either to match key statistics of observed traffic or to ensure

that the model has certain prespecified statistical properties. Fitting the WIG involves choosing its parameters to obtain a required variance progression of $\text{var}(V_{j,k})$. The WIG can provide a Gaussian approximation for *any* stationary discrete-time process X , that is the WIG can be fit to obtain

$$\text{var}(V_{n-j,k}) = \text{var}\left(K^{\{X\}}[2^j]\right). \quad (19)$$

We will refer to a WIG model for which (19) holds as a “WIG model of X ” in the rest of the paper. The WIG has been shown to capture the queuing behavior of Gaussian-like traffic well [18].

C. Multifractal wavelet model (MWM)

The MWM is a *non-Gaussian* model based on a multiscale tree that allows a general scaling behavior of the variance of tree nodes [20]. Unlike the WIG, it ensures positivity at all time scales, an intrinsic property of real data traffic which is often ill approximated by Gaussian models. Setting $V_{0,0} \geq 0$ the MWM uses independent *multiplicative* innovations $U_{j,k} \in [0, 1]$ to model the two children of node $V_{j,k}$ through

$$\begin{aligned} V_{j+1,2k} &:= V_{j,k} U_{j,k}, \\ V_{j+1,2k+1} &:= V_{j,k} (1 - U_{j,k}). \end{aligned} \quad (20)$$

Because the product of independent random variables converges to a log-normal distribution by the central limit theorem, the nodes $V_{j,k}$ become approximately log-normal with increasing j .

Following [20], we model the $U_{j,k}$ ’s as symmetric beta random variables

$$U_{j,k} \sim \beta(p_j, p_j), \quad p_j \geq 0 \quad (21)$$

and $V_{0,0}$ as

$$V_{0,0} \sim \varrho U_{-1} \quad (22)$$

with $\varrho \geq 0$ a constant and $U_{-1} \sim \beta(p_{-1}, p_{-1})$. The tree node $V_{j,k}$ is thus the product of several independent beta random variables. Using Fan’s result [21], we approximate the distribution of $V_{j,k}$ as another beta distribution with known parameters in order to compute different queuing approximations for the MWM.

Fitting the MWM involves choosing its parameters to obtain a required variance progression of $\text{var}(V_{j,k})$. The MWM can model *any* stationary discrete-time process X with *positive* autocovariance in the sense of (19). It has been shown to capture the queuing behavior of certain heavy-tailed, non-Gaussian traffic well [22].

While the WIG and MWM models are first-order stationary they are not second-order stationary. This is apparent from Fig. 1. Observe that $V_{j+2,4k}$ and $V_{j+2,4k+1}$ have the same parent node while $V_{j+2,4k+1}$ and $V_{j+2,4k+2}$ do not. Thus the correlation of $V_{j+2,4k+1}$ with its two neighbors, $V_{j+2,4k}$ and $V_{j+2,4k+2}$, are different. Both models however have a time-averaged correlation structure that is close to the stationary process X that they model (see [19, 20] for details).

D. Queuing analysis setup for fBm, WIG, and MWM

We now state precisely the queuing setup for the fBm, WIG, and MWM models that we analyze in subsequent sections. We set the initial queue size to be empty to satisfy the sufficient condition for (2) to hold (see Section II-A).

In this paper all queuing results for queues with fBm input correspond to a continuous-time queue with service rate c , initial value $Q_0 := 0$, and $K_\tau[t] = K_\tau^{\{\Delta_\tau B+m\}}[t] = B_\tau - B_{\tau-t} + mt$. We have

$$\begin{aligned} \mathbb{P}\{Q_\tau > b\} &= \mathbb{P}\left\{\sup_{0 < t < \tau} (K_\tau[t] - ct) > b\right\} \\ &\xrightarrow{\tau \rightarrow \infty} \mathbb{P}\left\{\sup_{t > 0} (K_0[t] - ct) > b\right\} \\ &=: \mathbb{P}\{Q_\infty > b\} \end{aligned} \quad (23)$$

where the limit holds because of stationarity of fBm increments and the Lemma 9 (in Appendix A). We assume that $\hat{c} := c - m > 0$ and study the quantity $\mathbb{P}\{Q_\infty > b\}$ as defined in (23).

All queuing results for the WIG and MWM correspond to a discrete-time queue initialized to $Q_0 := 0$ which evolves according to

$$Q_{\tau+1} = \max(Q_\tau + V_{n,\tau} - \tilde{c}^{(n)}, 0) \quad (24)$$

for $\tau = 0, 1, \dots, 2^n - 1$ with $\tilde{c}^{(n)} = cT2^{-n}$.

Defining $K_\tau[t] := \sum_{k=\tau-t}^{\tau-1} V_{n,k}$, for $t = 1, 2, \dots, \tau$ and $\tau = 1, 2, \dots, 2^n$, and $K_\tau[0] := 0$ we have

$$Q_\tau := \max_{t=0,1,\dots,\tau} \left(K_\tau[t] - \tilde{c}^{(n)}t \right). \quad (25)$$

We assume that

$$\mathbb{E}(V_{n,k}) < \tilde{c}^{(n)} \quad (26)$$

and study $\mathbb{P}\{Q_\tau > b\}$ which is a time-varying quantity. For the WIG and MWM models, $K_\tau[t]$ is only defined for $t = 0, 1, \dots, \tau$. Thus we define $M_\tau^{[\theta]}(b)$, $P_\tau^{[\theta]}(b)$, and $S_\tau^{[\theta]}(b)$ as in (8), (12), and (13) except that we replace θ by $\theta \cap \{0, 1, \dots, \tau\}$.

V. OPTIMALITY OF EXPONENTIAL TIME SCALES FOR THE MAX APPROXIMATION OF AN FBm QUEUE

Comparing (4) and (8) we see that the more dense θ is in \mathbb{R}_+ , the closer the max approximation is to $C_\tau(b)$. However, we simultaneously have to acquire data at more time scales and the max computational cost increases (see (8)). Ideally we prefer a set θ that optimally balances this trade-off in accuracy versus computational cost. In this section we prove certain optimality properties of sets of exponential time scales,

$$\theta_\alpha := \{\alpha^k : k \in \mathbb{Z}\}, \quad \alpha > 1. \quad (27)$$

More precisely, for a queue with fBm input we first obtain a non-asymptotic bound on the error of $M^{[\theta_\alpha]}(b)$ in approximating $C(b)$. This bound proves that $M^{[\theta_\alpha]}(b)$ accurately approximates $C(b)$ for a wide range of α . Second, we prove that θ_α is the most sparse of all sets θ that satisfy a particular accuracy criterion for $M^{[\theta]}(b)$.

A. Accuracy of $M^{[\theta_\alpha]}(b)$

Consider a queue fed by fBm traffic as described in Section IV-D. Then for $t > 0$, using (14) it is easily shown that [9]

$$\mathbb{P}\{K[t] - ct > b\} = \Phi(g(b, t)) \quad (28)$$

where

$$g(b, t) := \frac{b + \hat{c}t}{\sigma t^H} = \frac{b + (c - m)t}{\sigma t^H} \quad (29)$$

and Φ is the complementary cumulative distribution function of a zero mean unit variance Gaussian random variable. From (4) and (8) we have

$$C(b) = \sup_{t > 0} \Phi(g(b, t)) = \Phi\left(\inf_{t > 0} g(b, t)\right) \quad (30)$$

and

$$M^{[\theta]}(b) = \sup_{t \in \theta} \Phi(g(b, t)) = \Phi\left(\inf_{t \in \theta} g(b, t)\right). \quad (31)$$

We characterize the accuracy of $M^{[\theta]}(b)$ in terms of the following metric:

$$h_\theta := \sup_{b > 0} \frac{\inf_{t \in \theta} g(b, t)}{\inf_{t > 0} g(b, t)}. \quad (32)$$

Intuitively, the closer h_θ is to 1 the tighter we can bound the error of $M^{[\theta]}(b)$ in approximating $C(b)$.

The following theorem states the remarkable fact that h_θ is solely a function of the largest *ratio* of consecutive time scales in θ and does not depend on any other property of θ . In addition h_θ is not a function of the traffic model and queue parameters m , σ , and c .

Theorem 1: *Let $\theta = \{t_k\}_{k \in \mathbb{Z}}$ be a countable set of time scales such that*

$$\sup_k t_k = \infty, \quad \text{and} \quad \inf_k t_k = 0. \quad (33)$$

Assuming $t_{k-1} < t_k$, denote the largest ratio of consecutive time scales by

$$d_\theta := \sup_k \frac{t_k}{t_{k-1}}. \quad (34)$$

Assume $d_\theta < \infty$, $0 < H < 1$ and $\hat{c} > 0$. Then the accuracy metric of θ is given by

$$h_\theta = \zeta(d_\theta, H) := \frac{(d_\theta - 1)H^H(1 - H)^{1-H}}{(d_\theta - d_\theta^H)^{1-H}(d_\theta^H - 1)^H}. \quad (35)$$

The proof is in Appendix A.

According to Theorem 1, to set the accuracy metric h_θ we have only to choose the largest ratio of consecutive time scales d_θ appropriately. From (32) and (34) note that $h_\theta \geq 1$ and $d_\theta > 1$, $\forall \theta$. In Appendix A we prove that h_θ is closer to 1 for values of d_θ closer to 1.

We can use Theorem 1 to obtain the maximum error of $M^{[\theta_\alpha]}(b)$ in approximating $C(b)$ for all possible fBm traffic processes satisfying $\hat{c} > 0$.




Fig. 2. (a) $\zeta^*(\alpha)$ vs. α . For a large range of α , $\zeta^*(\alpha)$ is close to 1. (b) Comparison of $M^{[\theta_\alpha]}(b)$ to its upper bound $C(b)$, for different values of α . Observe that the lower bound of $M^{[\theta_\alpha]}(b)$ is almost identical to the upper bound $C(b)$.

Corollary 2: For a queue with fBm input traffic with $\hat{c} > 0$

$$\Phi(\zeta^*(\alpha)\Phi^{-1}(C(b))) \leq M^{[\theta_\alpha]}(b) \leq C(b), \quad (36)$$

where $\zeta^*(\alpha) := \max_{H \in (0,1)} \zeta(\alpha, H)$. Equivalently

$$M^{[\theta_\alpha]}(b) \leq C(b) \leq \Phi\left(\frac{1}{\zeta^*(\alpha)}\Phi^{-1}(M^{[\theta_\alpha]}(b))\right). \quad (37)$$

Proof: Note that by the construction of θ_α (see (27)), $d_{\theta_\alpha} = \alpha$. Thus from (9), (30) – (32), and (35) we have

$$\begin{aligned} C(b) &\geq M^{[\theta_\alpha]}(b) = \Phi\left(\inf_{t \in \theta_\alpha} g(b, t)\right) \\ &\geq \Phi\left(h_{\theta_\alpha} \inf_{t > 0} g(b, t)\right) \\ &= \Phi(\zeta(\alpha, H)\Phi^{-1}(C(b))) \\ &\geq \Phi(\zeta^*(\alpha) \cdot \Phi^{-1}(C(b))). \end{aligned} \quad (38)$$

□

In Fig. 2(a) we see that the plot of $\zeta^*(\alpha)$ versus α , which we obtained numerically, is close to 1 for a large range of values of α . As a result the lower bound of $M^{[\theta_\alpha]}(b)$ from (36) for different values of α is close to $C(b)$ as depicted in Fig. 2(b). In fact $M^{[\theta_\alpha]}(b)$ is almost identical to $C(b)$ when $C(b) > 10^{-8}$. Thus $M^{[\theta_\alpha]}(b)$ is for all practical purposes as accurate as $C(b)$ in approximating $\mathbb{P}\{Q_\infty > b\}$. We confirm the accuracy of $M^{[\theta_\alpha]}(b)$ through simulations in Section VIII.

B. Optimality of exponential time scales θ_α

Given a range of time scales \mathcal{T} we wish to find that time-scale set which is the most sparse (i.e. has the fewest elements) in \mathcal{T} while guaranteeing a certain accuracy of $M^{[\theta]}(b)$. The next theorem proves that for arbitrary \mathcal{T} , θ_α is the most sparse of all sets θ that have accuracy metric h_θ less than a given threshold.

Theorem 3: Let $\theta = \{t_k\}_{k \in \mathbb{Z}}$ satisfy (33). Let $\mathcal{A}_\mathcal{T}(\theta)$ denote the number of elements of θ that lie in a range of time scales $\mathcal{T} := (\underline{t}, \bar{t})$, $0 < \underline{t} < \bar{t}$. Denote by $\Gamma(\alpha)$ the set of all time-scale sets θ for which accuracy metric $h_\theta \leq \zeta(\alpha, H)$ and define the generalized exponential time scales as

$$\theta_{\alpha, \nu} := \{\nu \alpha^k : k \in \mathbb{Z}\} \quad (39)$$

where $\nu > 0$. Then for arbitrary \mathcal{T} and $\forall \nu$, we have $\theta_{\alpha, \nu} \in \Gamma(\alpha)$ and

$$\mathcal{A}_\mathcal{T}(\theta_{\alpha, \nu}) \leq 1 + \min_{\theta \in \Gamma(\alpha)} \mathcal{A}_\mathcal{T}(\theta). \quad (40)$$

Moreover there exists $\xi > 0$ such that

$$\mathcal{A}_\mathcal{T}(\theta_{\alpha, \xi}) = \min_{\theta \in \Gamma(\alpha)} \mathcal{A}_\mathcal{T}(\theta). \quad (41)$$

The proof is in Appendix A.

Theorem 3 is a direct consequence of the fact that the accuracy metric h_θ increases with the largest ratio of consecutive time scales d_θ . Thus $h_\theta = \zeta(d_\theta, H) \leq \zeta(\alpha, H)$ if and only if $d_\theta \leq \alpha$. Since the ratio of all consecutive time scale elements in θ_α equals the maximum allowed value of α , θ_α is the most sparse among all sets θ with accuracy metric less than $\zeta(\alpha, H)$.

VI. ASYMPTOTICS FOR FBM QUEUES

In this section, for a queue with fBm input, we study the accuracy of the max, the product, and the sum approximations of $\mathbb{P}\{Q_\infty > b\}$ for asymptotically large queue thresholds, that is as $b \rightarrow \infty$. While asymptotic queuing results are not always directly applicable to scenarios with finite queues, they often provide powerful intuition for network design [6–8, 23, 24].

We begin with some terminology. If $\lim_{b \rightarrow \infty} \Omega(b)/\Theta(b) = 1$ we say that Ω and Θ have the same *asymptotic decay* and denote it by $\Omega(b) \simeq \Theta(b)$. If $\log \Omega(b) \simeq \log \Theta(b)$ we say that Ω has the same *log-asymptotic decay* as Θ . Under the assumption that $\Omega(b) \rightarrow 0$ it is easily shown that an asymptotic decay implies a log-asymptotic decay, that is,

$$\Omega(b) \simeq \Theta(b) \left(\Leftrightarrow \frac{\Omega(b)}{\Theta(b)} \rightarrow 1 \right) \Rightarrow \log \Omega(b) \simeq \log \Theta(b) \quad (42)$$

but not vice versa.

A. Related work

Research on the asymptotic queuing behavior of fBm traffic has produced many enlightening results over the years. Large deviation principles reveal that $\mathbb{P}\{Q_\infty > b\}$ and $C(b)$ have the same log-asymptotic decay (see [6, 7])

$$\log \mathbb{P}\{Q_\infty > b\} \simeq \log C(b) \simeq -\frac{\eta b^{-(2-2H)}}{2} \quad (43)$$

where $\eta > 0$ is a constant depending on the traffic parameters and independent of b . More recent results show that for fBm $\mathbb{P}\{Q_\infty > b\}$ has a Weibull asymptotic decay [8, 23, 24]

$$\mathbb{P}\{Q_\infty > b\} \simeq \vartheta b^{(1-H)(1-2H)/H} e^{-\eta b^{2-2H}/2}, \quad (44)$$

where $\vartheta > 0$ is a constant independent of b . When fBm possesses LRD ($1/2 < H < 1$), this Weibull decay is slower than the exponential decay for a queue fed with short-range dependent traffic, for example fBm with $H = 1/2$ [5].

We call Θ an *asymptotic upper bound* of Ω if $\lim_{b \rightarrow \infty} \Omega(b)/\Theta(b) = 0$. From (44) we obtain that

$e^{-\eta b^{2-2H}/2}$ is an asymptotic upper bound of $\mathbb{P}\{Q_\infty > b\}$ when $1/2 < H < 1$, since

$$\lim_{b \rightarrow \infty} \frac{b^{(1-H)(1-2H)/H} e^{-\eta b^{2-2H}/2}}{e^{-\eta b^{2-2H}/2}} = 0. \quad (45)$$

This asymptotic upper bound was derived as the *maximum variance* approximation in [11]. For a detailed discussion on large queue asymptotics of LRD traffic see Chs. 4 to 11 of [25] and the references therein.

B. Asymptotic decay of multiscale queuing approximations

We now compare the log-asymptotic and asymptotic decay rates of the max, the product, and the sum approximations with that of $\mathbb{P}\{Q_\infty > b\}$. We only consider the case $\theta = \theta_\alpha$. The next theorem summarizes our results.

Theorem 4: *For a queue with fBm input traffic with parameters \hat{c} , σ , and H , define*

$$b_k := \alpha^k \hat{c}(1-H)/H, \quad k \in \mathbb{Z}, \quad (46)$$

where $\alpha > 1$ is arbitrary. Then the max, the product, and the sum approximations have the same log-asymptotic decay as $\mathbb{P}\{Q^{[\theta_\alpha]} > b_k\}$ and $\mathbb{P}\{Q_\infty > b_k\}$; that is as $b_k \rightarrow \infty$ we have

$$\begin{aligned} \log M^{[\theta_\alpha]}(b_k) &\simeq \log P^{[\theta_\alpha]}(b_k) \simeq \log S^{[\theta_\alpha]}(b_k) \\ &\simeq \log \mathbb{P}\{Q^{[\theta_\alpha]} > b_k\} \simeq \log \mathbb{P}\{Q_\infty > b_k\}. \end{aligned} \quad (47)$$

Moreover the max, the product, and the sum approximations all have the same asymptotic decay as $\mathbb{P}\{Q^{[\theta_\alpha]} > b_k\}$: as $b_k \rightarrow \infty$ we have

$$M^{[\theta_\alpha]}(b_k) \simeq P^{[\theta_\alpha]}(b_k) \simeq S^{[\theta_\alpha]}(b_k) \simeq \mathbb{P}\{Q^{[\theta_\alpha]} > b_k\}. \quad (48)$$

However

$$\lim_{k \rightarrow \infty} \frac{\mathbb{P}\{Q^{[\theta_\alpha]} > b_k\}}{\mathbb{P}\{Q_\infty > b_k\}} = 0. \quad (49)$$

The proof is in Appendix A.

Theorem 4 reveals the strengths and limitations of using traffic statistics only at exponential time scales θ_α to capture queuing behavior. Recall from (2) and (6) that $Q^{[\theta_\alpha]}$ approximates the queue size Q using traffic only at time scales $t \in \theta_\alpha$. From (47) we see that θ_α is dense enough in \mathbb{R}_+ to ensure that $\mathbb{P}\{Q^{[\theta_\alpha]} > b_k\}$ and $\mathbb{P}\{Q_\infty > b_k\}$ have the same log-asymptotic decays for a particular unbounded increasing sequence of queue sizes b_k . However, θ_α is not dense enough to ensure that $\mathbb{P}\{Q^{[\theta_\alpha]} > b_k\}$ and $\mathbb{P}\{Q_\infty > b_k\}$ have the same asymptotic decay.

We also observe from (48) that the max, the product, and the sum approximations have the same asymptotic decay as $\mathbb{P}\{Q^{[\theta_\alpha]} > b_k\}$. As a result they have the same log-asymptotic decay but different asymptotic decay as $\mathbb{P}\{Q_\infty > b_k\}$.

We next present non-asymptotic results comparing the different queuing approximations to $\mathbb{P}\{Q^{[\theta]} > b\}$.

VII. BOUNDS FOR THE QUEUING APPROXIMATIONS

The knowledge of whether or not a queuing approximation is an upper or lower bound of $\mathbb{P}\{Q > b\}$ aids different applications. For example if we provision the queue service rate such that the critical time scale approximation $C(b)$ equals 10^{-6} , then we must expect the actual tail queue probability $\mathbb{P}\{Q > b\}$ to exceed 10^{-6} since $C(b)$ lower bounds $\mathbb{P}\{Q > b\}$ (see (5)). If $C(b)$ is an accurate approximation of $\mathbb{P}\{Q > b\}$ to an order of magnitude, as our simulations with fBm traffic in Section VIII affirm, then we would effectively be provisioning for $\mathbb{P}\{Q > b\} < 10^{-5}$. If we replace the lower bound $C(b)$ by an approximation that is an upper bound of $\mathbb{P}\{Q > b\}$, then $\mathbb{P}\{Q > b\}$ is guaranteed to be less than 10^{-6} .

In this section we prove bounding results for the max, the product, and the sum approximations, which we compare to $\mathbb{P}\{Q^{[\theta]} > b\}$ rather than $\mathbb{P}\{Q > b\}$. Note from (10) that lower bounds of $\mathbb{P}\{Q^{[\theta]} > b\}$ are also lower bounds of $\mathbb{P}\{Q > b\}$. While the queuing approximations that are upper bounds of $\mathbb{P}\{Q^{[\theta]} > b\}$ are not necessarily upper bounds of $\mathbb{P}\{Q > b\}$, they approximate $\mathbb{P}\{Q > b\}$ well as we show through simulations in Section VIII.

A. Bounds for general input traffic processes

We first state a general result that holds for a queue fed by *any* traffic random process and then present model-specific results.

Lemma 5: *For a discrete or continuous-time queue of infinite size, with an arbitrary input traffic process and constant service rate*

$$M_\tau^{[\theta]}(b) \leq \mathbb{P}\{Q_\tau^{[\theta]} > b\} \leq S_\tau^{[\theta]}(b) \quad (50)$$

and

$$M_\tau^{[\theta]}(b) \leq P_\tau^{[\theta]}(b) \leq S_\tau^{[\theta]}(b), \quad (51)$$

where θ is any countable subset of \mathbb{R}_+ .

The proof is in Appendix A.

From Lemma 5 we see that max and sum approximations are always lower and upper bounds respectively of both $\mathbb{P}\{Q^{[\theta]} > b\}$ and the product approximation. In the rest of this section we compare the product approximation to $\mathbb{P}\{Q^{[\theta]} > b\}$.

Our results establish that for queues fed with fBm, WIG, or MWM input traffic, the product approximation is also an upper bound of $\mathbb{P}\{Q^{[\theta]} > b\}$, like the sum approximation. For these three models, from Lemma 5 we then have

$$M_\tau^{[\theta]}(b) \leq \mathbb{P}\{Q^{[\theta]} > b\} \leq P_\tau^{[\theta]}(b) \leq S_\tau^{[\theta]}(b), \quad (52)$$

implying that the product approximation is a closer upper bound of $\mathbb{P}\{Q^{[\theta]} > b\}$ than the sum approximation.² We note from (12) and (13) that the product approximation has the added advantage that it is guaranteed to be less than or equal to 1 unlike the sum approximation. In fact it can be shown that the product approximation is *strictly* less than 1 for queues with fBm, WIG, and MWM traffic as input [15].

²We prove (52) for the WIG and MWM only for $\theta = \theta_2$.

B. Product approximation bounds for Gaussian traffic

For queues fed with traffic from a large class of Gaussian processes, including fBm, $P_\tau^{[\theta]}(b)$ is an upper bound of $\mathbb{P}\{Q_\tau^{[\theta]} > b\}$ as claimed in (52).

Theorem 6: Consider a Gaussian traffic process X_τ as input to an infinite buffer queue with constant service rate (discrete or continuous-time). If $\text{cov}(K_\tau[t], K_\tau[r]) \geq 0$ for all $t, r \in \theta$ then

$$\mathbb{P}\{Q_\tau^{[\theta]} > b\} \leq P_\tau^{[\theta]}(b), \quad (53)$$

where θ is any countable subset of \mathbb{R}_+ .

The proof is in Appendix A.

Note that fBm satisfies the requirements of Theorem 6 since for all $t, r \geq 0$ and $0 < H < 1$

$$\text{cov}\left(K^{\{dB\}}[t], K^{\{dB\}}[r]\right) = \frac{1}{2}\left(t^{2H} + r^{2H} - |t - r|^{2H}\right) \geq 0. \quad (54)$$

C. Product approximation bounds for WIG and MWM traffic

Recall from Section IV-B that the WIG and MWM are non-stationary traffic models. As a consequence $P_\tau^{[\theta_2]}(b)$ changes with time location τ . We first compare $P_\tau^{[\theta_2]}(b)$ to $\mathbb{P}\{Q_\tau^{[\theta_2]} > b\}$ for $\tau = 2^n$, that is at the final time instant of the tree process, and then at all other time instants τ . We denote the final time instant 2^n by “end”.

Theorem 7: For the WIG and MWM with arbitrary model parameters

$$\mathbb{P}\{Q_{\text{end}}^{[\theta_2]} > b\} \leq P_{\text{end}}^{[\theta_2]}(b) \quad \forall b > 0. \quad (55)$$

The proof is in Appendix B.

Theorem 7 states that $P_\tau^{[\theta_2]}(b)$ is an upper bound of $\mathbb{P}\{Q_\tau^{[\theta_2]} > b\}$ at the final time instant for the WIG and the MWM for arbitrary model parameters. The only ingredient of the proof of Theorem 7 is the fact that the quantities $K_{\text{end}}[2^j]$, $j = 1, 2, \dots, n$ that determine $P_{\text{end}}^{[\theta_2]}(b)$ are nodes along the right edge of the tree and hence are related through independent innovations (see Fig. 1). Since this fact is true for arbitrary model parameters, so is (55).

Generalizing the proof of Theorem 7 so that (55) holds for all time instants τ is not straightforward because the quantities $K_\tau[2^j]$, $j = 1, 2, \dots, n$ are not always tree nodes for arbitrary τ and are hence not related through independent innovations as the quantities $K_{\text{end}}[2^j]$, $j = 1, 2, \dots, n$ are. However, for a WIG model of fGn we can extend (55) to all τ as stated next.

Theorem 8: For the WIG model of fGn

$$\mathbb{P}\{Q_\tau^{[\theta_\alpha]} > b\} \leq P_\tau^{[\theta_2]}(b) \leq P_{\text{end}}^{[\theta_2]}(b), \quad \forall \tau. \quad (56)$$

As a consequence

$$\frac{1}{2^n} \sum_{\tau=1}^{2^n} \mathbb{P}\{Q_\tau^{[\theta_2]} > b\} \leq P_{\text{end}}^{[\theta_2]}(b). \quad (57)$$




Fig. 3. Comparison of the max, the product and the sum approximations to $\mathbb{P}\{Q > b\}$ for fGn traffic with parameters $H = 0.8$, link capacity 10Mbps and $\sigma = 8 \times 10^5$ bits. For different utilizations in (a) and (b), the product and sum approximations are close to $\mathbb{P}\{Q > b\}$ for a wide range of queue thresholds b . The max approximation is a lower bound of $\mathbb{P}\{Q > b\}$ and is accurate to an order of magnitude.

The proof is in Appendix B.

Theorem 8 reveals that for a WIG model of fGn $P_{\text{end}}^{[\theta_2]}(b)$ is an upper bound of the time average of $\mathbb{P}\{Q_\tau^{[\theta_2]} > b\}$. The same result holds for a large class of WIG models of which the WIG model of fGn is a special case. The proof can be found in [15].

Earlier work on the queuing behavior of the WIG model of fGn proved that the time average of the tail queue probability $\mathbb{P}\{Q_\tau > b\}$ has the same log-asymptotic behavior as that of fGn [19]. In contrast Theorem 8 holds for any fixed queue threshold b .

We demonstrate through simulations in Section VIII that $P_{\text{end}}^{[\theta_2]}(b)$ approximates the time average of $\mathbb{P}\{Q_\tau > b\}$ well for a large range of queue sizes b for both the WIG and the MWM.

VIII. SIMULATIONS

In this section we demonstrate the accuracy of the max, the product, and sum approximations of $\mathbb{P}\{Q > b\}$ through simulations with fGn, WIG, and MWM synthetic traces as well as with video and measured Internet traces. We also demonstrate that the tails of multiscale marginals of traffic have a significant impact on queuing in certain scenarios by comparing the queuing behavior of the WIG and MWM models with that of measured Internet traffic. We restrict our attention to exponential time-scales with $\alpha = 2$, that is $\theta = \theta_2$. All error bars in the plots correspond to 95% confidence intervals.

A. Comparison of queuing approximations for fGn traffic

In earlier sections we theoretically compared the max, product, and sum approximations to $\mathbb{P}\{Q > b\}$ for an fBm-fed queue. Recall from (9) that the max approximation $M^{[\theta_2]}(b)$ is a lower bound of $\mathbb{P}\{Q > b\}$. While we have not proved an explicit relationship between $\mathbb{P}\{Q > b\}$ and the product and sum approximations, we proved that they are upper bounds of $\mathbb{P}\{Q^{[\theta_2]} > b\}$ which is itself a lower bound of $\mathbb{P}\{Q > b\}$ (see (52) and (10)). We now compare the different approximations of $\mathbb{P}\{Q > b\}$ through simulations with fGn traffic.

Simulation setup: The simulations use fGn traces with Hurst parameter $H = 0.8$, $t' = 10^{-4}$ s, and standard deviation at the 1s time-scale $\sigma = 8 \times 10^5$ bits. We generate the traces using the method described in [26]. We vary the mean rate of the traces to obtain different utilizations. The traces are input to a discrete-time queue of infinite length that operates at a time granularity of 10^{-4} s and has service rate 10Mbps. Because this time granularity is small we expect the queuing behavior of this discrete-time fGn-fed queue to closely resemble the continuous-time fBm-fed queue we have analyzed in detail in previous sections.

We estimate $\mathbb{P}\{Q > b\}$ for each simulation run as the fraction of time for which the queue size exceeds b . To eliminate transients we only make estimates using queue sizes during the second half of the simulation. The plots of $\mathbb{P}\{Q > b\}$ correspond to the mean obtained from 300 simulation runs. Each run uses a trace of length 2^{19} data points corresponding to a 52s simulation time.

Simulation results: The simulation results for two different utilizations are depicted in Fig. 3. We obtain the various queuing approximations using (8), (12), and (13) by choosing $\theta = \{t', 2t', \dots, 2^{20}t'\}$ which is equivalent to θ_2 truncated to lie within a fixed range of time-scales. Observe that in all cases $M^{[\theta_2]}(b)$ is a lower bound of $\mathbb{P}\{Q > b\}$ as predicted by (9). We also see that $M^{[\theta_2]}(b)$ is within an order of magnitude of $\mathbb{P}\{Q > b\}$ for a wide range of values of $\mathbb{P}\{Q > b\}$ ($\in [10^{-6}, 1]$). We conclude that $C(b)$ lying between $M^{[\theta_2]}(b)$ and $\mathbb{P}\{Q > b\}$ (see (9)) is also within an order of magnitude of $\mathbb{P}\{Q > b\}$ for the same range of $\mathbb{P}\{Q > b\}$.

From Fig. 3 observe that the product and sum approximations (which are almost identical) accurately track $\mathbb{P}\{Q > b\}$ for a wide range of queue sizes b and are better approximations than the max approximation in general. However unlike the max approximation, which is a guaranteed lower bound of $\mathbb{P}\{Q > b\}$, these two approximations do not bound $\mathbb{P}\{Q > b\}$ from above or from below and in fact intersect it at some point. Call the queue threshold at which the product approximation and $\mathbb{P}\{Q > b\}$ intersect b' . We observe that in all cases the product approximation is greater than $\mathbb{P}\{Q > b\}$ at $b = 0$ and for $b > b'$ is always less than $\mathbb{P}\{Q > b\}$. Thus for $b > b'$ the product approximation lies between the max approximation and $\mathbb{P}\{Q > b\}$, thus guaranteeing that it is a better approximation than the max approximation. The sum approximation has a similar behavior.

B. Impact of multiscale marginals on queuing: WIG vs. MWM

The impact of different traffic statistics on queuing has been extensively studied. Several studies have debated the importance of LRD for queuing [9, 28–31]. LRD is however only a function of the asymptotic second-order correlation structure of traffic (or equivalently the *variance* of traffic at multiple time scales).

In this section we move beyond second-order statistics and demonstrate the importance of the *tails* of traffic marginals at different time scales on queuing. We do so

by comparing the queuing behavior of the WIG and MWM processes with video and Internet WAN traces through simulations. Recall from Section IV that both the WIG and the MWM can capture a wide range of second-order correlation structures. The WIG and MWM however differ in their marginal characteristics: the WIG process is Gaussian whereas the MWM process is non-Gaussian. We interpret our results using the product approximation and the results of [12] which studied the influence of link utilization on queuing.

Traces: The two traces we use are AUCK, which contains the number of bytes per 2ms of recorded WAN traffic (mostly TCP packets) [27] and VIDEO, which consists of 15 video clips multiplexed with random starting points [32]. The finest time-scale in VIDEO corresponds to 2.77ms, 1/15 the duration of a single frame. The mean rates of AUCK and VIDEO are 1.456Mbps and 53.8Mbps, respectively. AUCK contains 1.8×10^6 data points and VIDEO 2^{18} . The Hurst parameter of AUCK obtained from the variance-time plot using time-scales 512ms to 262.144s is $H = 0.86$. For VIDEO, we find $H = 0.84$ using time-scales 354ms to 90.76s. From Fig. 4 and Fig. 5 observe that AUCK has a strongly non-Gaussian marginal while VIDEO's marginal resembles a Gaussian distribution.

Simulation results: We fit the WIG and MWM to the real data and then generated synthetic traces from the models. We then compared the queuing behavior of the synthesized WIG and MWM traces with that of the real data when they are input to a FIFO queue of infinite length. The plots of $\mathbb{P}\{Q > b\}$ correspond to the mean obtained from 1000 simulation runs.

We first present results for high link utilizations ($> 70\%$). Observe from Figs. 6(a) and (b), where we used the WAN traffic trace AUCK, that the real and synthetic traces exhibit asymptotic Weibullian tail queue probabilities, in agreement with the theoretical findings for LRD traffic (compare (44)). However, apart from this asymptotic match, the MWM is much closer to the queuing behavior of the real trace. The link capacity we use is 2Mbps, resulting in a utilization of 72%.

In the experiments with VIDEO (see Figs. 6(c) and (d)), which is much closer to a Gaussian process than AUCK, we observe that both the WIG and MWM closely match the correct queuing behavior. This confirms the influence of marginals and also reassures us that the MWM is flexible enough to model Gaussian traffic. Gaussian-like traffic, which must be positive, necessarily has a mean at least comparable to its standard deviation. Since for a large mean to standard deviation ratio the lognormal and Gaussian distributions resemble each other closely (see Fig. 5), the approximately lognormal MWM is suitable for Gaussian traffic [20]. The link capacity we use is 69Mbps, which corresponds to a utilization of 77%.

In the case of lower link utilizations ($< 50\%$) from Fig. 7 we see that the MWM outperforms the WIG for both AUCK and VIDEO traces to a greater extent than in the high utilization case.

For both the MWM and WIG we observe that the prod-




Fig. 4. Histograms of the bytes-per-time processes at time-scale 2ms for (a) wide-area traffic at the University of Auckland (trace AUCK) [27], (b) one realization of the WIG model, and (c) one realization of the MWM. Note the large probability mass over negative values for the WIG model.




Fig. 5. Histograms of the bytes-per-time processes at time-scale 2.77ms for (a) video traffic formed by multiplexing 15 video traces (trace VIDEO), (b) one realization of the WIG model, and (c) one realization of the MWM. Note that the MWM matches the marginal of the video traffic better than the WIG; however, the video traffic is more Gaussian than the AUCK traffic.




Fig. 6. Queuing performance of real data traces and synthetic WIG and MWM traces at high utilization. In (b), we observe that the MWM synthesis matches the queuing behavior of the AUCK data closely, while in (a) the WIG synthesis is not as close. In (c) and (d), we observe that both the WIG and the MWM match the queuing behavior of VIDEO. We also observe that the product approximation ($P^{[\theta_2]}(b)$) is close to the empirical queuing behavior for both synthetic traffic loads (both WIG and MWM) and that it performs better than the max approximation, $M^{[\theta_2]}(b)$.

uct approximation is close to $\mathbb{P}\{Q > b\}$ (see Figs. 6 and 7). The max approximation is within an order of magnitude of $\mathbb{P}\{Q > b\}$.

Interpreting results using the product approximation: Accepting the product approximation $P^{[\theta_2]}(b)$ as a close approximation to the actual tail queue probabilities, a closer look at (12) unravels how the marginals affect queue sizes. For traffic with heavier tailed marginals, the terms $\mathbb{P}\{K[2^i] < b + c2^i\}$ are smaller and the product approxi-

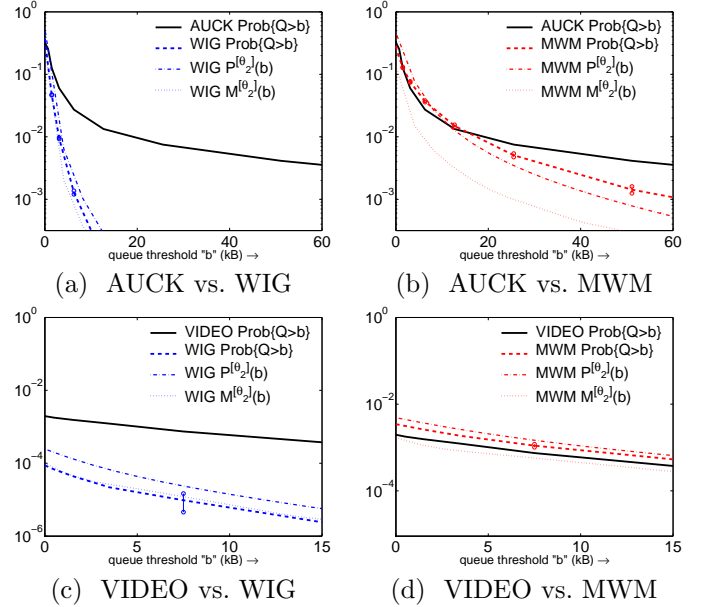


Fig. 7. Queuing performance of real data traces and synthetic WIG and MWM traces at low utilization. The MWM outperforms the WIG even more than at higher utilizations.

mation is larger. Since the approximately lognormal MWM marginals are more heavy tailed than the Gaussian WIG marginals, the MWM has a larger product approximation than the WIG.

In the case of VIDEO, which shows marginals much closer to Gaussian (see Fig. 5), both the WIG and MWM perform similarly in terms of capturing the tail queue probability at a high utilization, while at a low utilization the MWM outperforms the WIG. This result is easily explained using the finding in [12] that fine time-scale statistics influence queuing more than coarse time-scale statistics at low utilizations. Since fine time-scale marginals of VIDEO are more non-Gaussian than coarse time-scale marginals, obviously the MWM performs better than the WIG at low

utilizations.

IX. CONCLUSIONS

While this paper makes several new contributions to the field of queuing theory, many important problems still remain unsolved. First, we have ignored the case of finite length queues where packet drops occur. Our results thus are more useful in predicting packet queuing delays rather than packet losses.

Second, our analysis is for open-loop traffic models while real Internet traffic is mainly composed of closed-loop TCP traffic. Closed-loop traffic reacts to changes in network conditions unlike open-loop traffic. For example, unlike open-loop traffic, closed-loop traffic will reduce its offered load if a bottleneck link speed is reduced [33]. Thus one must use open-loop queuing results with caution in the Internet while ensuring that one does not affect network properties (delay and loss) which can influence the TCP traffic significantly. Possible applications of open-loop models are for Internet backbone provisioning (for very low delay/loss ISPs) [3], network inference schemes [34], and obviously in networks dominated by open-loop traffic like certain UDP streaming applications.

APPENDIX A

Proof of Theorem 1

We prove the theorem in four steps.

Step I: Determine $\inf_{t>0} g(b, t)$ for a fixed value $b > 0$.

From (29) we obtain the partial derivative of $g(b, t)$ with respect to t :

$$\frac{\partial g(b, t)}{\partial t} = \frac{t^H \hat{c} - (b + \hat{c}t)Ht^{H-1}}{\sigma t^{2H}} = \frac{\hat{c}t(1-H) - bH}{\sigma t^{1+H}}. \quad (58)$$

Thus $g(b, t)$ is minimized at $t = \lambda(b)$ where

$$\lambda(b) = \frac{bH}{\hat{c}(1-H)}. \quad (59)$$

In addition $g(b, t)$ is non-decreasing with t as we move away from $t = \lambda(b)$. Clearly

$$\inf_{t>0} g(b, t) = g(b, \lambda(b)) = \frac{b + \hat{c} \frac{bH}{\hat{c}(1-H)}}{\sigma \left(\frac{bH}{\hat{c}(1-H)} \right)^H} = \frac{b^{1-H} \hat{c}^H}{\sigma H^H (1-H)^{1-H}}. \quad (60)$$

Note that $\lambda(b)$ is indeed the critical time scale defined by (3).

Step II: Determine $\zeta(b) := \inf_{t \in \theta} g(b, t) / g(b, \lambda(b))$ for fixed b .

Observe from (33) that the sequence $\{t_k\}_{k \in \mathbb{Z}}$ extends from 0 to ∞ . Thus there must exist an $l \in \mathbb{Z}$ such that $\lambda(b) \in [t_{l-1}, t_l]$.

Consider the function

$$f(b, t) := \frac{g(b, t)}{g(b, \lambda(b))} = \frac{(b + \hat{c}t)(\sigma H^H (1-H)^{1-H})}{\sigma t^H (b^{1-H} \hat{c}^H)}. \quad (61)$$

Since $g(b, t)$ is non-decreasing as we move away from $t = \lambda(b)$, we must have that

$$\zeta(b) = \min\{f(b, t_{l-1}), f(b, t_l)\}. \quad (62)$$

Step III: Determine $\sup_{b \in A_l} \zeta(b)$ where

$$A_l := [\lambda^{-1}(t_{l-1}), \lambda^{-1}(t_l)], \quad (63)$$

and $\lambda^{-1}(t)$ is the inverse of $\lambda(b)$ given by

$$\lambda^{-1}(t) := \hat{c}t(1-H)/H. \quad (64)$$

By elementary calculus we obtain that $f(b, t_{l-1})$ monotonically increases with b when $b > \lambda^{-1}(t_{l-1})$. Also $f(b, t_l)$ monotonically decreases with increasing b when $b < \lambda^{-1}(t_l)$. If there exists $a_l \in A_l$ such that $f(a_l, t_{l-1}) = f(a_l, t_l)$, then $\zeta(b)$ must attain its supremum over A_l at this point (from (62)). Indeed such an a_l does exist. For the ease of notation we use $s_l := t_l/t_{l-1}$. Clearly $s_l > 1$ for all l . From (61) we obtain a_l as

$$a_l = \frac{\hat{c}t_{l-1}t_l(t_l^{H-1} - t_{l-1}^{H-1})}{t_l^H - t_{l-1}^H} = \frac{\hat{c}t_l}{s_l} \cdot \frac{s_l - s_l^H}{s_l^H - 1}. \quad (65)$$

As a result, after simplification

$$\sup_{b \in A_l} \zeta(b) = f(a_l, t_l) = \frac{(s_l - 1)H^H(1-H)^{1-H}}{(s_l^H - 1)^H(s_l - s_l^H)^{1-H}} = \zeta(s_l, H) \quad (66)$$

Step IV: Determine $h_\theta = \sup_{b \in \mathbb{R}_+} \zeta(b)$.

Claim ():* $\zeta(s_l, H)$ increases with s_l .

Proof of Claim ():* Note from (66) that $\zeta(s_l, H)$ equals $f(a_l, t_l)$. It is thus sufficient to prove that $f(a_l, t_l)$ increases with s_l . Without loss of generality we study how $f(a_l, t_l)$ changes by varying t_{l-1} keeping t_l fixed. Note that this is equivalent to varying s_l . We have from (65)

$$\frac{1}{\hat{c}t_l} \cdot \frac{\partial a_l}{\partial t_{l-1}} = \frac{t_{l-1}^{2H-1} s_l^{H-1} (Hs_l - s_l^H + (1-H))}{(t_l^H - t_{l-1}^H)^2}. \quad (67)$$

It is easily shown that the function $Hs_l - s_l^H + (1-H)$ equals 0 at $s_l = 1$ and has a positive derivative for $s_l > 1$. Thus $\frac{\partial a_l}{\partial t_{l-1}} > 0$ for all $s_l > 1$. Using this fact, the knowledge that $a_l < \lambda^{-1}(t_l)$, and the fact that $f(b, t_l)$ monotonically decreases for $b < \lambda^{-1}(t_l)$ we see that $f(a_l, t_l)$ decreases with increasing t_{l-1} , or equivalently it increases with increasing s_l . Claim (*) is thus proved.

From (33) and (64) we obtain that $\cup_l A_l = \mathbb{R}_+$. Exploiting the continuity of $\zeta(s_l, H)$ (see (66)) we then have

$$\sup_{b \in \mathbb{R}_+} \zeta(b) = \sup_l \zeta(s_l, H) = \zeta(\sup_l s_l, H) = \zeta(d_\theta, H). \quad (68)$$

□

Proof of Theorem 3

From (35) and the fact that $\zeta(s, H)$ is an increasing function of s (see Claim (*) in the previous proof) we have that $\theta \in \Gamma_\alpha$ if and only if $d_\theta \leq \alpha$. Since $d_{\theta_{\alpha, \nu}} = \alpha$ for all $\nu > 0$, we have $\theta_{\alpha, \nu} \in \Gamma(\alpha)$.

Consider $\underline{\theta} = \{w_k : k \in \mathbb{Z}\} \in \Gamma_\alpha$, where $w_k < w_{k+1} \forall k$, such that

$$\mathcal{A}_{\mathcal{T}}(\underline{\theta}) = \min_{\theta \in \Gamma(\alpha)} \mathcal{A}_{\mathcal{T}}(\theta). \quad (69)$$

Let w_i be the first element of $\underline{\theta}$ in (\underline{t}, \bar{t}) . Set $\xi = w_i/\alpha^i$. Consider the set

$$\theta_{\alpha, \xi} = \{y_k : y_k = \xi \alpha^k, k \in \mathbb{Z}\}. \quad (70)$$

Clearly $y_i = w_i$ and y_i must be the first element of $\theta_{\alpha, \xi}$ in (\underline{t}, \bar{t}) . Because $w_{k+1}/w_k \leq \alpha$ and $y_{k+1}/y_k = \alpha$, $\forall k$, we must have $w_k \leq y_k, \forall k \geq i$. Consequently $\mathcal{A}_{\mathcal{T}}(\theta_{\alpha, \xi}) \leq \mathcal{A}_{\mathcal{T}}(\underline{\theta})$ which proves (41).

It only remains to prove (40). We can write (\underline{t}, \bar{t}) as a union of the following $\mathcal{A}_{\mathcal{T}}(\theta_{\alpha, \xi}) + 1$ intervals: (\underline{t}, y_i) , $[y_i, y_{i+1})$, $[y_{i+1}, y_{i+2})$, \dots , $[y_{i+\mathcal{A}_{\mathcal{T}}(\theta_{\alpha, \xi})-2}, y_{i+\mathcal{A}_{\mathcal{T}}(\theta_{\alpha, \xi})-1})$, and $[y_{i+\mathcal{A}_{\mathcal{T}}(\theta_{\alpha, \xi})-1}, \bar{t})$. Note that the ratio of the supremum to infimum of each of these intervals is less than or equal to α . Consider $\theta_{\alpha, \nu}$ for arbitrary ν . Clearly by definition $\theta_{\alpha, \nu}$ can have at most one element in each of the these intervals. Thus (40) is proved. \square

Proof of Theorem 4: The proof relies on the following two claims.

Claim (a): $M^{[\theta_{\alpha}]}(b_k) \simeq S^{[\theta_{\alpha}]}(b_k)$.

Claim (b): $\lim_{b \rightarrow \infty} \frac{C(b)}{\mathbb{P}\{Q_{\infty} > b\}} = 0$.

From (51), (50), and Claim (a) we have (48). From (46) and (59) note that

$$\lambda(b_k) = \lambda^{[\theta_{\alpha}]}(b_k) = \alpha^k. \quad (71)$$

Thus (4) and (8) give

$$M^{[\theta_{\alpha}]}(b_k) = C(b_k), \quad \forall k. \quad (72)$$

From (43), (48), (72) and (42) we have (47). Finally Claim (b) combined with (48) gives (49).

We now prove the two claims. Recall the definition of $g(b, t)$ (see (29)). For the ease of notation we denote $g(b_k, \alpha^l)$ by $g_{k,l}$. From (60) and (71) we have

$$\inf_{t > 0} g(b_k, t) = g(b_k, \lambda(b_k)) = g(b_k, \alpha^k) = g_{k,k}. \quad (73)$$

Proof of Claim (a): From (8), (13), (28), and (73) we have

$$S^{[\theta_{\alpha}]}(b_k) = \sum_{l \in \mathbb{Z}} \Phi(g_{k,l}) \quad (74)$$

and

$$M^{[\theta_{\alpha}]}(b_k) = \sup_{l \in \mathbb{Z}} \Phi(g_{k,l}) = \Phi(g_{k,k}). \quad (75)$$

We now prove that the maximum term, $\Phi(g_{k,k})$, dominates the summation of (74). We note two properties of $\frac{g_{k,l}}{g_{k,k}}$. First, we have from (46) and (29)

$$\begin{aligned} \frac{g_{k,l}}{g_{k,k}} &= \frac{b_k + \widehat{c}\alpha^l}{\sigma\alpha^{lH}} \cdot \frac{\sigma\alpha^{kH}}{b_k + \widehat{c}\alpha^k} \\ &= (1-H)\alpha^{(k-l)H} + H\alpha^{(l-k)(1-H)} \\ &\geq \epsilon_H \alpha^{|l-k|\epsilon_H}, \end{aligned} \quad (76)$$

where $\epsilon_H = \min(H, 1-H)$.

Second, from (58) and (71) observe that $\frac{g_{k,l}}{g_{k,k}}$ monotonically increases with increasing l when $l \geq k$ and also with decreasing l when $l \leq k$. We then have $\forall l \neq k$

$$\frac{g_{k,l}}{g_{k,k}} \geq \min\left(\frac{g_{k,k+1}}{g_{k,k}}, \frac{g_{k,k-1}}{g_{k,k}}\right) =: \mathcal{I}_H > 1 \quad (77)$$

Now $g_{k,k}$ is an increasing unbounded function of k . From (74), and using the following estimates of Φ (see page 42 in [35]),

$$\left(1 - \frac{1}{\delta^2}\right) \frac{e^{-\delta^2/2}}{\delta\sqrt{2\pi}} \leq \Phi(\delta) \leq \frac{e^{-\delta^2/2}}{\delta\sqrt{2\pi}}. \quad (78)$$

we have

$$\begin{aligned} \Phi(g_{k,k}) &\leq S^{[\theta_k]}(b_k) = \Phi(g_{k,k}) + \sum_{l > k} \Phi(g_{k,l}) + \sum_{l < k} \Phi(g_{k,l}) \\ &\leq \Phi(g_{k,k}) \left(1 + \frac{2}{\epsilon_H} \cdot \frac{g_{k,k}^2 e^{-(T_H^2-1)g_{k,k}^2/2} \alpha^{-\epsilon_H}}{(g_{k,k}^2-1)(1-\alpha^{-\epsilon_H})}\right). \end{aligned} \quad (79)$$

From (77) and the fact that $g_{k,k} \xrightarrow{k \rightarrow \infty} \infty$ we have

$$\lim_{k \rightarrow \infty} \frac{S^{[\theta_k]}(b_k)}{\Phi(g_{k,k})} = 1, \quad (80)$$

which proves Claim (a).

Proof of Claim (b): From (78) observe that $\Phi(\delta) \simeq \frac{e^{-\delta^2/2}}{\delta\sqrt{2\pi}}$.

Set $\eta := \left(\frac{\widehat{c}^H}{\sigma H^H (1-H)^{1-H}}\right)^2$. From (30) and (60) we then have

$$C(b) = \Phi(b^{1-H} \eta^{1/2}) \simeq \frac{b^{-(1-H)}}{\eta^{1/2} \sqrt{2\pi}} e^{-b^{2-2H} \eta/2}. \quad (81)$$

When $1/2 < H < 1$ we have $0 < \frac{2H-1}{H} < 1$ which implies that

$$\lim_{b \rightarrow \infty} \frac{b^{-(1-H)}}{b^{-(1-H)(2H-1)/H}} = 0. \quad (82)$$

Claim (b) follows from (44), (81) and (82), and the theorem is proved. \square

Proof of Lemma 5: Consider a queue with constant service rate c bits per unit time. Clearly

$$\sup_{t \in \theta} \mathbb{P}\{\mathcal{E}[t]\} \leq \mathbb{P}\{\cup_{t \in \theta} \mathcal{E}[t]\} \leq \sum_{t \in \theta} \mathbb{P}\{\mathcal{E}[t]\}, \quad (83)$$

where $\mathcal{E}[t] := \{K_{\tau}[t] - ct > b\}$. From (6) we see that $\mathbb{P}\{Q^{[\theta]} > b\}$ is identical to $\mathbb{P}\{\cup_{t \in \theta} \mathcal{E}[t]\}$. Then (8), (13) and (83) give (50).

Note that (51) is equivalent to

$$\max_{k=1, \dots, l} (1 - a_k) \leq 1 - \prod_{k=1}^l a_k \leq \sum_{k=1}^l (1 - a_k), \quad (84)$$

for $0 \leq a_k \leq 1, k = 1, \dots, l$, which is elementary. \square

Lemma 9: (see [36]): If $E_1 \subset E_2 \dots$ and $E = \cup_i E_i$ then $\lim_{i \rightarrow \infty} \mathbb{P}\{E_i\} = \mathbb{P}\{E\}$. If $E_1 \supset E_2 \dots$ and $E = \cap_i E_i$ then $\lim_{i \rightarrow \infty} \mathbb{P}\{E_i\} = \mathbb{P}\{E\}$.

The following lemma helps us prove Theorem 6.

Lemma 10: (page 6 in [37]) Let $\Upsilon[t]$ and $\Psi[t]$, $t \in \theta$, be separable Gaussian random processes, where θ is a parameter set. If the following relations hold for their covariance functions:

$$\text{var}(\Upsilon[t]) = \text{var}(\Psi[t]), \quad \forall t \in \theta \quad (85)$$

$$\text{cov}(\Upsilon[t], \Upsilon[r]) \leq \text{cov}(\Psi[t], \Psi[r]), \quad \forall t, r \in \theta \quad (86)$$

plus their expected values are the same $\forall t$: then for any $x \in \mathbb{R}$

$$\mathbb{P}\left\{\sup_{t \in \theta} \Upsilon[t] < x\right\} \leq \mathbb{P}\left\{\sup_{t \in \theta} \Psi[t] < x\right\}. \quad (87)$$

Proof of Theorem 6

For $t \in \theta$ define independent Gaussian random variables $\Upsilon[t] \sim \mathcal{N}(\mathbb{E}(K_\tau[t] - tc), \text{var}(K_\tau[t]))$ and set $\Psi[t] := K_\tau[t] - tc$. Label the elements of θ as $\{t_k\}_{k \in \mathbb{Z}}$, which we assume satisfy (33). Using Lemma 9 with the events $E_i := \cap_{k=-i}^i \{\Upsilon[t_k] < b\}$ and $E := \cap_i E_i = \{\sup_{t \in \theta} \Upsilon[t] < b\}$ we have

$$P^{[\theta]}(b) = 1 - \prod_{t \in \theta} \mathbb{P}\{\Upsilon[t] < b\} = 1 - \mathbb{P}\left\{\sup_{t \in \theta} \Upsilon[t] < b\right\}. \quad (88)$$

Then the fact that $\sup_{t \in \theta} \Psi[t] = Q_\tau^{[\theta]}$ along with (87) and (88) prove the theorem. \square

APPENDIX B

Lemma 11: *Assume that the events W_i are of the form $W_i = \{I_i < \kappa_i\}$, where $I_i = R_0 + R_1 + \dots + R_i$ for $1 \leq i \leq n$ and where R_0, \dots, R_n are independent, otherwise arbitrary random variables. Then, for $1 \leq i \leq n$, we have*

$$\mathbb{P}\{W_i | W_{i-1}, \dots, W_0\} \geq \mathbb{P}\{W_i\}. \quad (89)$$

Proof: We first spell out some notation. By f_L and F_L we denote the probability density function (p.d.f) and cumulative distribution function (c.d.f), respectively, of a random variable L . Furthermore, we denote by $F_{L|E}(l)$ the c.d.f. of L conditioned on knowing the event E . For convenience, let us write $W_i := \{I_i < \kappa_i\}$ for short, and let us introduce the auxiliary random variables $Y_0 := L_0 := I_0 := R_0$,

$$Y_i := I_i | W_{i-1}, \dots, W_0 \quad \text{and} \quad L_i := I_i | W_i, \dots, W_0, \quad i \geq 1. \quad (90)$$

To prove the lemma, it is enough to show that

$$F_{Y_i}(r) \geq F_{L_i}(r) \quad (91)$$

$\forall r \in \mathbb{R}$ and $\forall i$ and then set $r = \kappa_i$.

We prove (91) by induction. First note that $F_{Y_0}(r) \geq F_{L_0}(r)$. Next, we assume that (91) holds for i and show that it holds also for $i+1$. Bayes' rule yields

$$F_{L_i}(r) = \begin{cases} \frac{F_{Y_i}(r)}{F_{Y_i}(\kappa_i)}, & \text{if } r \leq \kappa_i \\ 1, & \text{otherwise} \end{cases} \geq F_{Y_i}(r). \quad (92)$$

The key to the proof is to note that $Y_{i+1} = L_i + R_{i+1}$, where R_{i+1} is independent of L_j and hence of W_j for $j \leq i$. In short, R_{i+1} is independent of L_i . This fact, (91) and (92) allow us to write

$$\begin{aligned} F_{Y_{i+1}}(r) &= \mathbb{P}\{L_i + R_{i+1} < r\} \\ &= \int_{-\infty}^{\infty} F_{L_i}(r - r_{i+1}) f_{R_{i+1}}(r_{i+1}) dr_{i+1} \\ &\geq \int_{-\infty}^{\infty} F_{Y_i}(r - r_{i+1}) f_{R_{i+1}}(r_{i+1}) dr_{i+1} \\ &\geq \int_{-\infty}^{\infty} F_{L_i}(r - r_{i+1}) f_{R_{i+1}}(r_{i+1}) dr_{i+1} \\ &= F_{L_{i+1}}(r). \end{aligned} \quad (93)$$

This proves the claim by induction. \square

Proof of Theorem 7

Let us first show that Lemma 11 applies to the WIG and the MWM for the events $W_i = K_{\text{end}}[2^{n-i}] < b$. To this end we need only show that these W_i can be written in the appropriate form. Recall that we have $K_{\text{end}}[2^{n-i}] = V_{i, 2^{i-1}}$.

WIG: The WIG uses additive innovations $Z_{j,k}$ arranged on a tree as in Fig. 1. It is immediate from (18) that $K_{\text{end}}[2^{n-i}]$ becomes

$$K_{\text{end}}[2^{n-i}] = V_{i, 2^{i-1}} = 2^{-i} V_{0,0} - \sum_{j=0}^{i-1} 2^{j-i} Z_{j, 2^{j-1}}. \quad (94)$$

It suffices, thus, to set $\kappa_i = 2^i b + 2^n \tilde{c}^{(n)}$, $R_0 = V_{0,0}$ and $R_i = -2^{i-1} Z_{i-1, 2^{i-1-1}}$.

MWM: The MWM employs the same tree structure as the WIG, however, with multiplicative innovations $U_{j,k}$. Recalling (20), $K_{\text{end}}[2^{n-i}]$ becomes

$$K_{\text{end}}[2^{n-i}] = V_{i, 2^{i-1}} = V_{0,0} \prod_{j=0}^{i-1} (1 - U_j). \quad (95)$$

Taking logarithms, it is a simple task to write the events W_i in the required form, this time by setting $\kappa_i = \ln(b + 2^{n-i} \tilde{c}^{(n)})$, $R_0 = \ln(V_{0,0})$, and $R_i = \ln(1 - U_{i-1})$.

Using (89) we find

$$\begin{aligned} \mathbb{P}\{Q_{\text{end}}^{[\theta_2]} > b\} &= 1 - \mathbb{P}\{Q_{\text{end}}^{[\theta_2]} < b\} = 1 - \mathbb{P}\{\cap_{i=0}^n W_i\} \\ &= 1 - \mathbb{P}\{W_0\} \prod_{i=1}^n \mathbb{P}\{W_i | W_{i-1}, \dots, W_0\} \\ &\leq 1 - \prod_{i=0}^n \mathbb{P}\{W_i\} = P_{\text{end}}^{[\theta_2]}(b). \end{aligned} \quad (96)$$

\square

The following lemma helps prove Theorem 8.

Lemma 12: (Theorem 5 in [19]) For a WIG model of fGn with $1/2 < H < 1$

$$\text{var}(K_{\text{end}}[t]) \geq \text{var}(K_\tau[t]) \quad (97)$$

for $t = 1, 2, \dots, \tau$ and for $\tau = 1, \dots, 2^n$.

Proof of Theorem 8: Note that

$$\mathbb{P}\{K_\tau[t] - \tilde{c}^{(n)} t < b\} = 1 - \Phi\left(\frac{b + \tilde{c}^{(n)} t - \mathbb{E}(K_\tau[t])}{\sqrt{\text{var}(K_\tau[t])}}\right). \quad (98)$$

From (26) we have that

$$b + \tilde{c}^{(n)} t - \mathbb{E}(K_\tau[t]) > 0. \quad (99)$$

Since the process $V_{n,k}$ is first-order stationary, $\mathbb{E}(K_{\text{end}}[t]) = \mathbb{E}(K_\tau[t])$ for all τ and t . This fact along with Lemma 12, (98), and (99) then give

$$\mathbb{P}\{K_\tau[2^i] - \tilde{c}^{(n)} 2^i < b\} \leq \mathbb{P}\{K_{\text{end}}[2^i] - \tilde{c}^{(n)} 2^i < b\}, \quad (100)$$

$i = 0, \dots, \lfloor \log_2 \tau \rfloor$, $\tau = 1, \dots, 2^n$. We thus have

$$P_{\text{end}}^{[\theta_2]}(b) \geq P_{\tau}^{[\theta_2]}(b), \quad \forall \tau = 1, \dots, 2^n. \quad (101)$$

To complete the proof we show the following claim.
Claim 1: $\text{cov}(K_{\tau}[t], K_{\tau}[r]) \geq 0$ for $0 \leq t, r \leq \tau$.

It is easy to show that the covariance of any two arbitrary leaf nodes is positive for a WIG model of fGn with $1/2 < H < 1$, that is $\text{cov}(V_{n,k}, V_{n,l}) > 0 \forall k, l$. Because $K_{\tau}[t] = \sum_{k=\tau-t}^{\tau-1} V_{n,k}$ it follows that $\text{cov}(K_{\tau}[t], K_{\tau}[r])$ is a linear combination of covariances of leaf nodes with *positive* weights. This proves Claim 1.

Claim 1 and Theorem 6 give

$$P_{\tau}^{[\theta_2]}(b) \geq \mathbb{P}\left\{Q_{\tau}^{[\theta_2]} > b\right\}, \quad \forall \tau = 1, \dots, 2^n. \quad (102)$$

Combining (101) and (102) proves the theorem. \square

REFERENCES

- [1] V. Ribeiro, R. Riedi, M. S. Crouse, and R. G. Baraniuk, "Multi-scale queuing analysis of long-range-dependent network traffic," *IEEE INFOCOM*, pp. 1026–1035, Mar. 2000.
- [2] L. Breslau, S. Jamin, and S. Shenker, "Comments on the performance of measurement-based admission control," *IEEE INFOCOM*, pp. 1233–1242, Mar. 2000.
- [3] C. Fraleigh, F. Tobagi, and C. Diot, "Provisioning IP backbone networks to support delay-based service level agreements," *IEEE INFOCOM*, Apr. 2003.
- [4] K. Papagiannaki, S. Moon, C. Fraleigh, P. Thiran, F. Tobagi, and C. Diot, "Analysis of measured single-hop delay from an operational backbone network," *IEEE INFOCOM*, Jun. 2002.
- [5] W. Leland, M. Taqqu, W. Willinger, and D. Wilson, "On the self-similar nature of Ethernet traffic (extended version)," *IEEE/ACM Trans. Networking*, pp. 1–15, 1994.
- [6] I. Norros, "A storage model with self-similar input," *Queueing Systems*, vol. 16, pp. 387–396, 1994.
- [7] N. Duffield and N. O'Connell, "Large deviations and overflow probabilities for the general single-server queue, with applications," *Math. Proc. Camb. Phil. Soc.*, vol. 118, pp. 363–374, 1995.
- [8] J. Husler and V. Piterbarg, "Extremes of a certain class of gaussian processes," *Stochastic Processes and their Applications*, vol. 83, pp. 257–271, 1999.
- [9] A. L. Neidhardt and J. L. Wang, "The concept of relevant time scales and its application to queuing analysis of self-similar traffic," *ACM SIGMETRICS*, pp. 222–232, Mar. 1998.
- [10] M. Grossglauser and J.-C. Bolot, "On the relevance of long-range dependence in network traffic," *Computer Communication Review*, vol. 26, pp. 15–24, Oct. 1996.
- [11] J. Choe and N. B. Shroff, "Queueing analysis of high-speed multiplexers including long-range dependent arrival processes," *IEEE INFOCOM*, pp. 617–624, Mar. 1999.
- [12] A. Erramilli, O. Narayan, A. Neidhardt, and I. Sanjee, "Performance impacts of multi-scaling in wide area TCP/IP traffic," *IEEE INFOCOM*, pp. 352–359, Mar. 2000.
- [13] Z.-L. Zhang, V. J. Ribeiro, S. B. Moon, and C. Diot, "Small-time scaling behaviors of Internet backbone traffic: an empirical study," *IEEE INFOCOM*, 2003.
- [14] B. Hajek and L. He, "On variations of queue response for inputs with the same mean and autocorrelation function," *IEEE/ACM Trans. on Networking*, vol. 6, no. 5, pp. 588–598, 1998.
- [15] V. J. Ribeiro, R. Riedi, and R. Baraniuk, "Multiscale queuing analysis," *Technical Report*, 2004. Available at www.ece.rice.edu/~vinay/publications.html.
- [16] E. Reich, "On the Integrodifferential Equation of Takács I," *Annals of Mathematical Statistics*, vol. 29, pp. 563–570, 1958.
- [17] B. B. Mandelbrot and J. W. V. Ness, "Fractional Brownian Motions, Fractional Noises and Applications," *SIAM Review*, vol. 10, pp. 422–437, Oct. 1968.
- [18] S. Ma and C. Ji, "Modeling video traffic in the wavelet domain," *IEEE INFOCOM*, pp. 201–208, Mar. 1998.
- [19] S. Ma, "Network traffic modeling and analysis," *Ph.D. Thesis, Rensselaer Polytechnic Institute*, May 1998.
- [20] R. Riedi, M. S. Crouse, V. Ribeiro, and R. G. Baraniuk, "A multifractal wavelet model with application to TCP network traffic," *IEEE Trans. on Information Theory*, vol. 45, pp. 992–1018, 1999.
- [21] D.-Y. Fan, "The distribution of the product of independent beta variables," *Commun. Statist.-Theory Meth.*, vol. 20, no. 12, pp. 4043–4052, 1991.
- [22] V. Ribeiro, R. Riedi, M. S. Crouse, and R. G. Baraniuk, "Simulation of non-Gaussian long-range-dependent traffic using wavelets," *ACM SIGMETRICS*, pp. 1–12, May 1999.
- [23] O. Narayan, "Exact asymptotic queue length distribution for fractional Brownian traffic," *Advances in Performance Analysis*, vol. 1, no. 1, pp. 39–63, 1998.
- [24] L. Massoulié and A. Simonian, "Large buffer asymptotics for the queue with fBm input," *Journal of Applied Probability*, vol. 36, no. 3, pp. 894–906, 1999.
- [25] *Self-similar network traffic and performance evaluation*, Wiley Interscience, 2001. Eds. K. Park and W. Willinger.
- [26] M. Crouse and R. G. Baraniuk, "Fast, exact synthesis of Gaussian and non-Gaussian long-range dependent processes," *submitted to IEEE Trans. on Information Theory*, 1999.
- [27] NLNR, "Auckland-II trace archive." Available at <http://moat.nlanr.net/Traces/Kiwitraces/>. Trace 20000125-143640, corresponding to 3:11:28 hours of mostly TCP traffic.
- [28] A. Erramilli, O. Narayan, and W. Willinger, "Experimental queuing analysis with long-range dependent traffic," *IEEE/ACM Trans. on Networking*, vol. 4, pp. 209–223, Apr. 1996.
- [29] D. P. Heyman and T. V. Lakshman, "What are the implications of long-range dependence for VBR-video traffic engineering?," *IEEE/ACM Trans. on Networking*, vol. 4, pp. 301–317, June 1996.
- [30] M. Palekar and A. M. Makowski, "Tail probabilities for a multiplexer with self-similar traffic," *IEEE INFOCOM*, pp. 1452–1459, Mar. 1996.
- [31] B. K. Ryu and A. Elwalid, "The importance of long-range dependence of VBR video traffic in ATM traffic engineering: Myths and realities," *ACM SIGCOMM*, vol. 26, no. 4, pp. 3–14, 1996.
- [32] O. Rose, "Statistical properties of MPEG video traffic and their impact on traffic modeling in ATM systems," Tech. Rep. 101, University of Wuerzburg. Institute of Computer Science Research Report Series., Feb. 1995.
- [33] Y. Joo, V. Ribeiro, A. Feldmann, A. C. Gilbert, and W. Willinger, "TCP/IP traffic dynamics and network performance: A lesson in workload modeling, flow control, and trace-driven simulations," *Computer Communication Review*, vol. 31, pp. 25–37, Apr. 2001.
- [34] V. Ribeiro, M. Coates, R. Riedi, S. Sarvotham, B. Hendricks, and R. Baraniuk, "Multifractal cross-traffic estimation," *ITC Specialist Seminar on IP Traffic Measurement*, Sept. 2000.
- [35] R. J. Adler, *An Introduction to Continuity, Extrema, and Related Topics for General Gaussian Processes*. Institute of Mathematical Statistics Lecture Notes - Monograph Series, 1990.
- [36] P. Billingsley, *Probability and Measure*. Wiley Interscience, 1995.
- [37] V. I. Piterbarg, *Asymptotic methods in the theory of Gaussian processes and fields*. Translations of mathematical monographs, vol. 148, American Mathematical Society, 1996.

TABLE I
SYMBOLIC NOTATION

Symbol	Definition	Relevant Equation
\sim	is distributed as	–
\simeq	equal asymptotic decay	$\Omega(b) \simeq \Theta(b) \Leftrightarrow \lim_{b \rightarrow \infty} \Omega(b)/\Theta(b) = 1$
$\stackrel{d}{=}$	equality in distribution	–
$\mathbb{E}, \text{var}, \text{cov}$	expectation, variance, covariance	–
\mathbb{P}	probability of an event	–
$\mathbb{Z}, \mathbb{R}_+, \mathbb{N}$	integers, positive real numbers, natural numbers	–
τ	time variable	–
t	time-scale variable	–
c	continuous-time queue service rate	–
X_τ	traffic input to queue	–
$K_\tau^{\{X\}}[t]$	traffic at time scale t	$K_\tau^{\{X\}}[t] := \int_{\tau-t}^\tau X_\omega d\omega$
Q_τ	queue size at time τ	$Q_\tau := \sup_{t>0} (K_\tau[t] - ct)$
b	queue threshold	–
$\lambda_\tau(b)$	critical time scale	$\lambda_\tau(b) := \arg \sup_{t>0} \mathbb{P}\{K_\tau[t] - ct > b\}$
$C_\tau(b)$	critical time scale approximation of $\mathbb{P}\{Q > b\}$	$C_\tau(b) := \sup_{t>0} \mathbb{P}\{K_\tau[t] - ct > b\}$
θ	countable set of time scales	–
$Q_\tau^{[\theta]}$	approximation of Q using time scales $t \in \theta$	$Q_\tau^{[\theta]} := \sup_{t \in \theta} (K_\tau[t] - ct)$
$\lambda_\tau^{[\theta]}(b)$	approximation of $\lambda_\tau(b)$	$\lambda_\tau^{[\theta]}(b) := \arg \sup_{t \in \theta} \mathbb{P}\{K_\tau[t] - ct > b\}$
$M_\tau^{[\theta]}(b)$	max approximation of $\mathbb{P}\{Q > b\}$	$M_\tau^{[\theta]}(b) := \sup_{t \in \theta} \mathbb{P}\{K_\tau[t] - ct > b\}$
α	exponential parameter	$\alpha > 1$
θ_α	set of exponential time scales	$\theta_\alpha := \{\alpha^j : j \in \mathbb{Z}\}$
$P_\tau^{[\theta]}(b)$	product approximation of $\mathbb{P}\{Q > b\}$	$P_\tau^{[\theta]}(b) := 1 - \prod_{t \in \theta} \mathbb{P}\{K_\tau[t] - ct < b\}$
$S_\tau^{[\theta]}(b)$	sum approximation of $\mathbb{P}\{Q > b\}$	$S_\tau^{[\theta]}(b) := \sum_{t \in \theta} \mathbb{P}\{K_\tau[t] - ct > b\}$
β	Beta distribution	–
\mathcal{N}	Gaussian distribution	–
H	Hurst parameter	$0 < H < 1$
B_τ	fractional Brownian motion (fBm)	$B_{a\tau} \stackrel{d}{=} a^H B_\tau$
σ	fBm standard deviation	$\sigma^2 = \text{var}(B_1)$
m	mean rate of an fBm-fed queue	$K_\tau[t] := B_\tau - B_{\tau-t} + mt$
$G_\tau[t]$	fractional Gaussian noise	$G_\tau[t] := K_\tau^{\{\Delta\tau B\}}[t] = B_\tau - B_{\tau-t} - t$
\hat{c}	unused capacity of an fBm-fed queue	$\hat{c} = c - m$
t'	finest time scale for fGn	$\{G_{it'}[t']\}_{i \in \mathbb{Z}}$
$g(b, t)$	function relevant to an fBm-fed queue	$g(b, t) := (b + \hat{c}t)/\sigma t^H$
Φ	Gaussian complementary c.d.f.	–
d_θ	max. ratio of consecutive time scales	$d_\theta := \sup_k t_k/t_{k-1}, \theta = \{t_k\}_{k \in \mathbb{Z}}$
h_θ	metric for “goodness” of θ	$h_\theta = \sup_{b>0} \inf_{t \in \theta} g(b, t) / \inf_{t>0} g(b, t)$
$\zeta(\alpha, H)$	–	$\zeta(\alpha, H) = \frac{(\alpha-1)H^H(1-H)^{1-H}}{(\alpha-\alpha^H)^{1-H}(\alpha^H-1)^H}$
$\mathcal{A}(\theta)$	number of elements of θ in (\underline{t}, \bar{t})	–
ϵ_H	constant	$\epsilon_H := \min(H, 1-H)$
\mathcal{I}_H	constant	$\mathcal{I}_H := \min((1-H)\alpha^{-H} + H\alpha^{(1-H)}, (1-H)\alpha^H + H\alpha^{-(1-H)})$
$\Gamma(\alpha)$	set composed of time scale sets θ	$\Gamma(\alpha) = \{\theta : h_\theta \leq \zeta(\alpha, H)\}$
b_k	sequence of queue thresholds	$b_k = \alpha^k \hat{c}(1-H)/H, k \in \mathbb{Z}$
η, ϑ	constants	$\mathbb{P}\{Q_\infty > b\} \simeq \vartheta b^{(1-H)(1-2H)}/H e^{-\eta b^{2-2H}/2}$
$V_{j,k}$	tree node at depth j and location k	–
$Z_{j,k}$	WIG innovations	$V_{j+1,2k} = (V_{j,k} + Z_{j,k})/2$
n	tree model depth	–
T	tree model time interval T	–
$\bar{c}^{(n)}$	queue service rate for a tree model queue	$\bar{c}^{(n)} = cT2^{-n}$
$U_{j,k}$	MWM innovations	$V_{j+1,2k} = V_{j,k}U_{j,k}$
p_j	parameter of MWM innovation	$U_{j,k} \sim \beta(p_j, p_j), p_j \geq 0$
end	final time instant of tree process, 2^n	–
ϱ, U_{-1}	MWM root node parameters	$V_{0,0} \sim \varrho U_{-1}$
k, r, a_k, l, i	free variables	–

Note: For clarity we occasionally add a superscript enclosed in $\{ \}$ to denote which input traffic process a quantity corresponds to, for example $S_\tau^{\{X\}, [\theta]}(b)$ when the input process is X . Superscripts in square brackets “[]” denote sets of time scales. Superscripts in parenthesis “()” denote the tree-depth of the WIG and MWM models. For time-invariant quantities we drop subscript τ .