

End-to-end Model for a Flow in the Internet

Muthaiah Venkatachalam* Jasleen Kaur Harrick M. Vin

Technical Report TR-01-32
Department of Computer Sciences
University of Texas at Austin
TAY 2.124, Austin, TX 78712-1188, USA

Aug 1, 2001

Abstract — We develop an end-to-end model for packet inter-arrival times of flows in the Internet. Our model illustrates that under asymptotic conditions of high network utilization, (1) packet inter-arrival times of flows become heavy-tailed after they traverse a large number of hops, and (2) the aggregate of such flows is a long-range dependent self-similar process. We validate our model through extensive simulations and derive the region of applicability—with respect to non-asymptotic settings of network utilization, heterogeneity in traffic composition, and number of hops that a flow traverses—of our model. We find that for flows that contribute significant amount ($\approx 10\%$) to the cumulative bit rate of the aggregate, the observations yielded by the model hold at relatively low levels of network utilization (30-40%), small number of network hops (4-6), and moderate levels of heterogeneity in packet sizes and bandwidth requirements of flows.

1 Introduction

Simplicity of implementation is central to the design of scalable, high-performance networks. Hence, high-performance routers have traditionally used First-In-First-Out (FIFO) packet scheduling algorithm. FIFO scheduling, unfortunately, does not provide isolation across flows; bursty traffic arrival from a flow can affect significantly the performance—delay, jitter, and loss—observed by other flows.

A number of analytical models have been proposed for characterizing delay and jitter performance of flows in FIFO networks [5, 17, 16, 20, 23, 22, 25]. However, primarily due to the complexity of the analysis involved, most of these studies do not provide closed-form characterization of end-to-end delay and jitter properties (see Section 9 for a detailed discussion on related work).

The aim of this paper is to provide a closed-form end-to-end analysis for flows in FIFO networks. For the remainder of this paper, we refer to a *flow* as a traffic stream, all packets of which traverse the same network path from the source to the destination. A wide variety of traffic streams—such as label-switched paths (LSP) in Multi-protocol Label Switching (MPLS) networks, Virtual Private Network (VPN) tunnels, and Virtual Circuits (VC) in ATM networks—fit this definition of a flow. Even in the current Internet, TCP, UDP and other best-effort traffic streams can be approximated as flows as long as route changes are not frequent.

We model the inter-arrival time between packets at the destination for individual flows under asymptotic conditions of network utilization and number of hops. Our model yields the following key insights.

- The variance in inter-arrival times for individual flows tends to infinity in the asymptotic case, indicating that flows become *heavy-tailed* after traversing a network of FIFO routers.

*Currently at Intel Corporation

- The aggregation of flows becomes a long-range dependent self-similar process.

A particularly interesting and deceptively simple type of flow is one that is shaped to a *constant bit-rate* (CBR)—with constant packet inter-arrival times and packet sizes—at the source. It is widely believed that (1) shaping flows to CBR at the source and (2) ensuring that the bandwidth available at each node to the CBR aggregate is at least as large as the cumulative rate requirement of the CBR flows results in a satisfactory end-to-end delay and jitter. The above philosophy, in fact, is the basis of the Virtual Leased Line service model proposed in the context of Differentiated Services (DiffServ) networks [13, 19]. We demonstrate that, at high network utilization, even CBR flows become heavy tailed after traversing a few hops. Hence, shaping flows to CBR benefits only in networks that operate at low levels of utilization. We further show that at low network utilization, shaping sources to constant packet sizes is sufficient to reduce the asymptotic burstiness; shaping sources to constant inter-arrival times is not necessary.

We validate our model through extensive simulations and derive the region of applicability—with respect to non-asymptotic network and traffic parameters—of the model. We find that although the model predicts the behavior of flows in the asymptotic case, for flows that contribute significant amount ($\approx 10\%$) to the cumulative bit rate of the aggregate, the observations yielded by the model hold at relatively low levels of network utilization (30-40%), small number of network hops (4-6), and moderate levels of heterogeneity in the packet sizes and bandwidth requirements of the flows. We argue that, in many instances, a flow—such as one that represents a VPN tunnel between two sites of an enterprise—does occupy significant fraction (10-20%) of the bit rate allocated for VPN services at routers. Our simulation results predict that such flows will become heavy-tailed at relatively low levels of network utilization and after traversing a small number of network hops; further, all of the micro-flows that get multiplexed onto such tunnels will themselves become heavy-tailed.

Finally, we point out several important implications of carrying heavy-tailed flows and long-range dependent self-similar aggregates in a network. We argue that a client receiving a heavy-tailed multimedia flow either perceives poor performance or requires large amount of buffers to smooth out the heavy tails. Long-range dependence of traffic aggregates, on the other hand, has several implications on network queuing delays and packet losses. We conclude that operating a network so as to maintain its traffic outside the heavy-tailed and long-range dependence domain is highly desirable.

The rest of the paper is organized as follows. We define related concepts in Section 2. In Section 3, we formulate the problem of end-to-end analysis of packet inter-arrival times of flows. In Sections 4 and 5, we present the models for individual flows and flow aggregates, respectively. We discuss the benefits, if any, of shaping sources to CBR in Section 6. In Section 7, we describe our simulation setup and present simulation results. We discuss the impact of heavy-tailedness and long-range dependence on the design of networks in Section 8. We discuss the related work in Section 9, and summarize our conclusions in Section 10.

2 Background

In this section, we first summarize the concepts of self-similarity, long-range dependence and heavy-tailed flow and then discuss the tests used to detect the presence of these properties in a time series.

2.1 Concepts

Self-similarity: Consider a time series $Y(t)$. $Y(t)$ is *self-similar* iff,

$$Y(t) \sim m^{-H} Y(mt)$$

That is, $Y(t)$ is identically distributed as $Y(mt)$ with a scaling factor m^{-H} , where $m > 0$ and $H \in [0, 1]$ is known as the *Hurst* parameter.

Let $X(t)$ define increments on $Y(t)$: $X(t) = Y(t+1) - Y(t)$. If $Y(t)$ has stationary increments, then it can be shown that $X \sim m^{1-H} X^m$, where $X^m = (X(0) + X(1) \dots X(m)) / m$. When $H > 1/2$, it can be

shown that the variance of X^m dies out at a sub-linear rate with increase in m :

$$\text{Var}(X^m) = m^{2H-2} \text{Var}(X)$$

Self-similar processes are often *long-range dependent* – we define this concept next.

Long-range Dependence: *Long-range dependence* of a process $\{X_i\}$ refers to the property that its auto-covariance function does not sum up to a finite quantity. Let

$$\begin{aligned} E[X_i] &= \mu \\ \text{Var}(X_i) &= \sigma^2 \\ \rho_i &= E[(X_i - \mu)(X_0 - \mu)]/\sigma^2 \end{aligned}$$

Then X is long-range dependent *iff*:

$$\sum_i \rho_i = \infty$$

The most commonly encountered form of such long-range dependent ρ_k in communication networks is hyperbolic and is given by:

$$\rho_k = ck^{2H-2}, k \rightarrow \infty$$

where $H > 1/2$ and c is a constant. This form of ρ_k corresponds to that of a second-order self-similar process [26].

Heavy-tailed Flows: A flow is said to be *heavy-tailed* *iff*

$$P(J > x) = c(x)x^{-\alpha}, x \rightarrow \infty, 0 < \alpha < 2 \quad (1)$$

where the random variable J denotes the packet inter-arrival time of the flow and $c(x)$ is a slowly varying function as $x \rightarrow \infty$. Slow variation can be stated as:

$$\lim_{x \rightarrow \infty} \frac{c(tx)}{c(x)} = 1, \forall t > 0$$

For all practical purposes, $c(x)$ can be approximated by a constant. It can be shown that $\text{Var}(J) = \infty$. If $\alpha < 1$, it can also be shown that $E[J] = \infty$. The association between heavy-tailedness and long-range dependent self-similarity is discussed in Section 5.

2.2 Tests for Heavy Tails and Long-range Dependent Self-similarity

The literature contains the following tests for identifying heavy-tailedness and long-range dependent self-similarity.

1. *LLCD Test:* The log-log complementary distribution (LLCD) test is used to detect the presence of heavy tails in a source. The complementary distribution of the packet inter-arrival time ($\text{Prob}(J > x)$) is plotted on a log-log scale. From Equation (1), it can be seen that for a heavy-tailed source, after some value of x , the distribution should look like a straight line with slope $-\alpha$.
2. *Hill Test:* Let U_1, U_2, \dots, U_n denote the packet inter-arrival times of a given flow in ascending order. The *Hill* function, $hill(k)$, for the flow is defined by:

$$hill(k) = \left[\sum_{i=0}^{i=k-1} k^{-1} \log \frac{U_{n-i}}{U_{n-k}} \right]^{-1}$$

A source is said to be heavy-tailed with parameter α if $hill(k)$ stabilizes to α after some value of k .

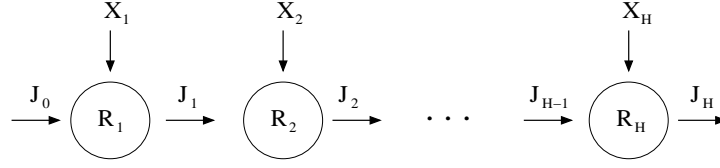


Figure 1: Network Model

3. *Variance-time test*: This test determines if a process $\{X_i\}$ is self-similar and long-range dependent. The logarithm of the variance of the aggregated (averaged) process $\{X_i^m\}$ is plotted versus the logarithm of the aggregation level, m . A least-squares line is fitted on the data, and the slope $(2H - 2)$ provides an estimate of H . If the slope is greater than -1 , then $H > 1/2$, indicating that the data is self-similar and long-range dependent.

3 Problem Formulation

Consider a network that carries flows between source-destination pairs. In this paper, a flow is considered to be any traffic stream, all packets of which traverse the same path, from the source to the destination, through the network. We assume all flows in the network are independent of each other. This is a reasonable assumption in large-scale networks with diverse routing. Let $\{R_1, \dots, R_H\}$ denote the routers along a path of length H through this network. We consider networks where each router R_h employs FIFO scheduling to arbitrate access to the link bandwidth. Very large router buffers are assumed which guarantee no loss of packets. Let μ_h be the utilization of the output link of router R_h .

Divide the time-axis into *very small* slots. Consider a flow f (referred to as *tagged flow*) that traverses the above path. Let the random variable J_h represent the inter-packet separation (IPS) of this flow as its packets leave router R_h (see Figure 1). Let J_0 represent the inter-packet separation of the flow at the source. We assume that the mean EJ_0 and the variance $Var(J_0)$ for the inter-packet separation of the tagged flow at the source are finite.

Let the random variable X_h^j denote the time it takes to service the cross-traffic that enters router R_h during the j^{th} time-slot after the arrival of a packet of flow f . Then the following relationship exists between J_{h-1} , J_h and X_h^j :

$$J_h = I + X_h^1 + \dots + X_h^{J_{h-1}-I} \quad (2)$$

where, I represents the index of the last time-slot after the arrival of a tagged packet but prior to the arrival of the next tagged packet, at which the queue becomes empty. We assume that X_h^j is independent of J_{h-1} – which is a reasonable assumption in a network with diverse routing and a large number of input links carrying independent cross-traffic.

Note that $EX_h^j = \mu_h$. For simplicity of analysis we assume that the cross-traffic characteristics are uniform over all routers. Henceforth, for simplicity of exposition, we shall therefore denote X_h^j as X^j and μ_h as μ . In the steady state, let $X_h^j \sim X$. Then $\mu = EX$. Let $\sigma^2 = Var(X)$.

Let ρ_k denote the auto-covariance in the cross-traffic that arrives at time instances separated by k time-slots, that is

$$\rho_k = E[(X^j - \mu)(X^{j+k} - \mu)]/\sigma^2$$

Notice that with the assumed independence of flows, the auto-covariance of the cross-traffic at a particular lag k is approximately the average of the auto-covariances at the same lag of the various cross-traffic flows present. Further, for most flows, the auto-covariance is a decreasing function of k . Hence, if the cross-traffic aggregates a large number of flows, then it is safe to assume that ρ_k is a strictly monotonically decreasing function of k . The following lemma states how this assumption leads to a property of ρ_k that we use in our analysis.

Lemma 1 If ρ_k is a strictly monotonically decreasing function, then $\sum_{k=1}^m \rho_k$ is a sub-linear function in m .

Proof: $\sum_{k=1}^m \rho_k$ cannot be super-linear due to the monotonically decreasing property. Suppose it is linear, that is, let $\sum_{k=1}^m \rho_k = c_1 m + c_2$ for some $c_1, c_2 \in R, c_1 > 0$. Differencing both sides w.r.t. m , we obtain: $\rho_{m-1} = c_1$, which violates the strictly monotonically decreasing property. ■

Our objective in this paper is to model J_h under these traffic and network conditions.

4 Model for a Tagged Flow

Let $S_j = X^1 + \dots + X^j$ and $N = J_h - I$. Then

$$\begin{aligned} J_{h+1} &= I + X^1 + \dots + X^{J_h - I} \\ &= I + S_N \end{aligned}$$

It is easy to see that X and I (and therefore X and N) are not independent. Due to multiplexing of traffic from various links, the IPS of the tagged flow J_h can be assumed to be independent of the cross-traffic intensity X and therefore I . Then,

$$EJ_{h+1} = EI + ES_N \quad (3)$$

$$Var(J_{h+1}) = Var(S_N) + Var(I) + 2Cov(S_N, I) \quad (4)$$

Further,

$$\begin{aligned} ES_N &= \sum_n E[X^1 + \dots + X^n | N = n] p(n) \\ &= \sum_n n E[X | N = n] p(n) \\ &= E[NX] \\ &= E[(J_h - I)X] \end{aligned} \quad (5)$$

$$\begin{aligned} Var(S_N) &= E[S_N - ES_N]^2 \\ &= E[S_N - NE[X|N] + NE[X|N] - ES_N]^2 \\ &= E[S_N - NE[X|N]]^2 + E[NE[X|N] - ES_N]^2 \\ &\quad + 2E[(S_N - NE[X|N])(NE[X|N] - ES_N)] \end{aligned} \quad (6)$$

The first term in (6) simplifies to:

$$\begin{aligned} &E[S_N - NE[X|N]]^2 \\ &= \sum_n E \left[\sum_{j=1}^n (X^j - E[X|N]) \right]^2 p(n) \\ &= E[(N + 2(N-1)\rho_1^{(N)} + \dots + 2\rho_{N-1}^{(N)})Var(X|N)] \end{aligned} \quad (7)$$

Here, $\rho_i^{(N)} = E[(X^{i+j} - E[X|N])(X^j - E[X|N])]/Var(X|N)$. By the assumption that $\rho_i^{(N)}$ is a strictly monotonically decreasing function of i , and using Lemma 1, we can show that (7) does not contain any $Var(N)$ term.

The second term in (6) simplifies to:

$$\begin{aligned} &E[NE[X|N] - ES_N]^2 \\ &= E[NE[X|N] - E[NE[X|N]]]^2 \\ &= Var(NE[X|N]) \\ &= Var((J_h - I)E[X|I]) \\ &= Var(J_h)E[E[X|I]^2] + EJ_h Var(E[X|I]) \\ &\quad + Var(IE[X|I]) - 2Cov(J_h E[X|I], IE[X|I]) \end{aligned} \quad (8)$$

The third term in (6) simplifies to:

$$\begin{aligned}
& E[(S_N - NE[X|N])(NE[X|N] - ES_N)] \\
&= \sum_n E \left[\sum_{i=1}^n X^i - NE[X|N] \right] (nE[X|N] - ES_N)p(n) \\
&= \sum_n E \left[\sum_{i=1}^n (X^i - E[X|N]) \right] (nE[X|N] - ES_N)p(n) \\
&= 0
\end{aligned} \tag{9}$$

From Equations (6), (7), (8) and (9), we have:

$$\begin{aligned}
& Var(S_N) \\
&= Var(J_h)E[E[X|I]^2] + EJ_h Var(E[X|I]) + \\
&\quad Var(IE[X|I]) - 2Cov(J_h E[X|I], IE[X|I]) + \\
&\quad E[(N + 2(N-1)\rho_1^{(N)} + \dots + 2\rho_{N-1}^{(N)})Var(X|N)]
\end{aligned} \tag{10}$$

From Equations (10) and (4), we have:

$$\begin{aligned}
& Var(J_{h+1}) \\
&= Var(J_h)E[E[X|I]^2] + EJ_h Var(E[X|I]) \\
&\quad + Var(IE[X|I]) - 2Cov(J_h E[X|I], IE[X|I]) \\
&\quad + E[(N + 2(N-1)\rho_1^{(N)} + \dots + 2\rho_{N-1}^{(N)})Var(X|N)] \\
&\quad + Var(I) + 2Cov(S_N, I)
\end{aligned} \tag{11}$$

For the limiting distribution $J = \lim_{h \rightarrow \infty} J_h$, we then have

$$Var(J) = \frac{f(X, I, J)}{1 - E[E[X|I]^2]} \tag{12}$$

where $f(X, I, J)$ is the collection of all but the first term in the right-hand side of (11).

Consider the behavior of (12) as the utilization asymptotically approaches 1. As $EX \rightarrow 1$, we have $EI \rightarrow 0$. But $I \geq 0$ always. Therefore, $I \rightarrow 0$. In this limiting scenario, several terms in $f(X, I, J)$ go to zero. We then have:

$$\begin{aligned}
& \lim_{EX \rightarrow 1} Var(J) \\
&= \lim_{EX \rightarrow 1} \frac{E[(J + 2(J-1)\rho_1^{(J)} + \dots + 2\rho_{J-1}^{(J)})Var(X)]}{1 - (EX)^2} \\
&= \infty
\end{aligned}$$

Thus, the variance of the tagged flow grows in an unbounded manner as it traverses a large number of hops H in the network with utilization μ arbitrarily close to one. From Equations (3) and (5), we get

$$\begin{aligned}
\lim_{EX \rightarrow 1} EJ_{h+1} &= \lim_{EX \rightarrow 1} EI + ES_{J_h - I} \\
&= \lim_{EX \rightarrow 1} EI + E[(J_h - I)X] \\
&= EJ_h
\end{aligned}$$

Thus, we see that

$$\lim_{EX \rightarrow 1} EJ = EJ_0$$

Thus, the limiting distribution of the inter-packet separation of the tagged flow has a finite mean and infinite variance under asymptotic conditions. Lemma 2 shows that when the marginal distribution of such a random variable J is well defined, it is heavy-tailed.

Lemma 2 If $J > 0$, $EJ < \infty$ and $Var(J) = \infty$ then,

$$P(J > x) = c(x)x^{-\alpha}, \quad x \rightarrow \infty$$

where $1 < \alpha < 2$ and $c(x)$ is slowly varying.

Proof: Consider

$$V_\beta(x) = \int_{x_0}^x y^\beta p_J(y) dy \quad 0 < x_0 < x < \infty$$

where, $p_J(y)$ is the probability density function for J . Observe that, for y , the inter-packet separation represented in time slots:

$$V_\beta(x) \leq V_{\beta+\delta}(x) \quad \forall \delta > 0 \quad (13)$$

Since $EJ < \infty$ and $Var(J) = \infty$, we have

$$\lim_{x \rightarrow \infty} V_1(x) < \infty \quad \text{and} \quad \lim_{x \rightarrow \infty} V_2(x) = \infty \quad (14)$$

From (13) and (14), we can conclude:

$$\begin{aligned} \exists \beta : \quad & 1 < \beta \leq 2, \quad \lim_{x \rightarrow \infty} V_\beta(x) = \infty \quad \text{and} \\ & \forall \delta > 0, \quad \lim_{x \rightarrow \infty} V_{\beta-\delta}(x) < \infty \end{aligned} \quad (15)$$

In [4](p.331,Th. 8.1.2) put $\alpha = 0$ and use the converse implication. i.e, let,

$$\lim_{x \rightarrow \infty} \frac{x^\beta \int_x^\infty p_J(y) dy}{V_\beta(x)} = \gamma$$

We assume that $p_J(y)$ is well-defined so that γ exists and $0 \leq \gamma \leq \infty$. From [4] and (15), we then have

$$\begin{aligned} P(J > x) &= c(x)x^{-\alpha}, \quad x \rightarrow \infty \\ \alpha &= \frac{\beta}{1 + \gamma} \\ \alpha &\in [0, \beta] \end{aligned} \quad (16)$$

where, $c(x)$ is slowly varying. From (16), $EJ < \infty$, and $Var(J) = \infty$, it can be shown that we must have $1 < \alpha < 2$. Thus,

$$P(J > x) = c(x)x^{-\alpha}, \quad x \rightarrow \infty, \quad 1 < \alpha < 2 \quad (17)$$

■

Hence the inter-packet separation of individual flows becomes heavy-tailed in the limit.

Let J^* be the limiting distribution of the inter-arrival time (time between the arrival of the first bits of two consecutive packets of the tagged flow). Let P represent the distribution of the time it takes to transmit a tagged packet (for simplicity of analysis, all link capacities are assumed to be equal). Then,

$$\begin{aligned} J^* &= J + P \\ Var J^* &= Var J + Var P + 2E[JP] - 2EJE P \end{aligned} \quad (18)$$

Therefore, if $EP < \infty$ and $Var J \rightarrow \infty$, then $Var J^* \rightarrow \infty$.

5 Aggregation of Flows

In [18, 27] it has been shown that the superposition of heavy-tailed *on-off* flows, with strictly alternating off- and on-periods, converges to Fractional Brownian Motion (FBM) as the number of flows grows large. Our model shows that individual flows become heavy-tailed under asymptotic conditions. The cross-traffic, which is a superposition of such flows, therefore converges to FBM. FBM is a self-similar process with Hurst parameter

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

Since $1 < \alpha < 2$ in (17), it can be seen that for the cross-traffic X , $1 > H > 1/2$. Thus the cross-traffic X is a long-range dependent FBM process.

The auto-covariance of such a cross-traffic ρ_k , is [18]:

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

6 Does it Help to Shape Sources to CBR?

We have seen that in the asymptotic limits, traffic arrival at a node in the network becomes very bursty. It is widely believed that such burstiness can be eliminated or reduced by shaping individual flows to a constant bit-rate (CBR) at the edge of the network, and ensuring that the sum of bit-rates of CBR flows at any node in the network does not exceed the available bandwidth. In fact, this conjecture is at the heart of the design of the Virtual Leased Line service model proposed in the context of Differentiated Services networks [13, 19]. Unfortunately, several recent studies have shown that the lack-of-burstiness property of CBR flows is not maintained once the CBR flows traverse a network of FIFO routers [10, 24]. In what follows, we analyze the CBR tagged traffic case using the model developed in Section 4. For the analysis, we define the *increasing convex ordering* of two random variables, \geq_{icx} as follows:

Definition 1 $X \geq_{icx} Y$ iff for all increasing convex functions $h(\cdot)$, $E[h(X)] \geq E[h(Y)]$.

Lemma 3 If $EX = EY$ and $VarY = 0$, then $X \geq_{icx} Y$.

Proof: Let $h(\cdot)$ be a convex function. Then, $Eh(X) = \sum_x h(x)p(x) \geq h(\sum_x xp(x)) = h(EX) = h(EY) = Eh(Y)$. Hence, $X \geq_{icx} Y$. ■

From (18) and (12) we have:

$$VarJ^* = \frac{f(X, I, J)}{1 - E[E[X|I]^2]} + VarP + 2E[JP] - 2EJEP$$

Notice that $E[E[X|I]] = EX$ and $Var(EX) = 0$. Applying Lemma 3, we therefore get $E[X|I] \geq_{icx} EX$. Since the square function is increasing and convex, this implies that $E[E[X|I]^2] \geq (EX)^2$. It follows that

$$VarJ^* \geq \frac{f(X, I, J)}{1 - (EX)^2} + VarP + 2E[JP] - 2EJEP \quad (19)$$

For a CBR flow, $VarP = 0$ and $VarJ_0^* = 0$. Therefore, for CBR flows, we derive the following observations from (19):

- When EX is large, the first term in (19) dominates the right hand side. $VarP$ affects the asymptotic inter-arrival time only if EX is small. Hence, shaping sources to CBR is beneficial only in networks that operate at low levels of utilization.
- $VarJ_0^*$ does not play a role in the asymptotic burstiness as long as it is finite, but $VarP$ does. It follows that at low network utilization, shaping sources to constant packet sizes is sufficient to reduce the asymptotic burstiness; shaping sources to constant bit rate is not necessary.

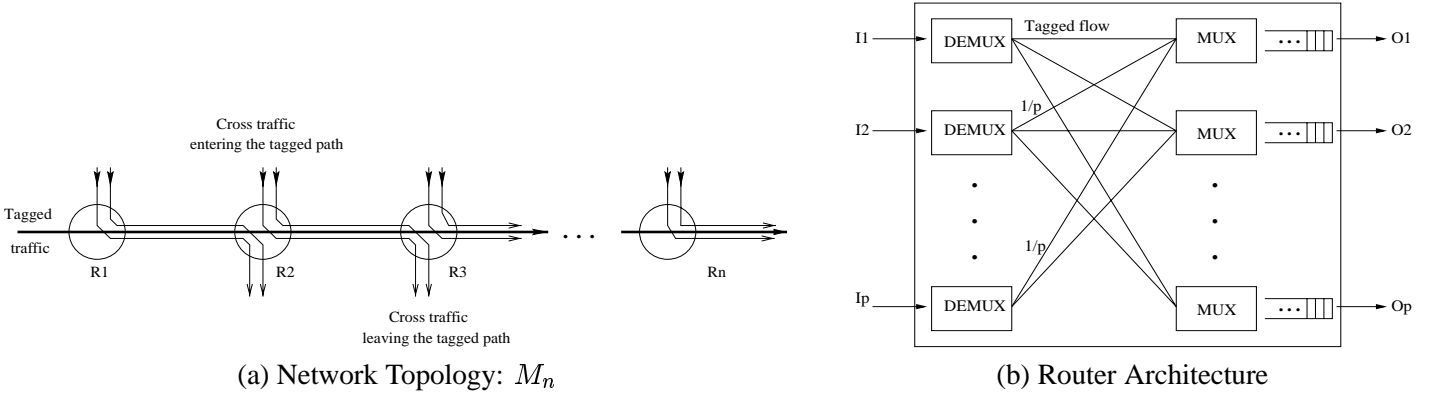


Figure 2: Simulation environment: network topology and router architecture

7 Experimental Validation and Analysis

We conduct a simulation study of networks with traffic inputs, with two objectives: (1) to validate under asymptotic conditions, the model developed in Sections 4 and 5; and (2) to identify the set of non-asymptotic network settings (for various traffic compositions) where the predicted heavy-tailedness of individual flow becomes visible. To conduct these experiments we have developed a network simulator using CSIM [2].

7.1 Simulation Environment

7.1.1 Network Topology

For our experiments, we consider a linear, multi-hop network topology (see Figure 2(a)). This network model is fairly general and has been used in literature [11, 15, 16, 28]. Let M_n denote a linear, multi-hop network topology with n routers, and let R_i ($i \in [1, n]$) denote the i^{th} router in the topology. Given such a topology, we are interested in the end-to-end performance of the *tagged flow*—a flow that enters the network topology at router R_1 and traverses the multi-hop network topology M_n . While traversing the network, the tagged flow interacts with aggregates of flows (referred to as the *cross traffic*) that enter and depart the network at each router along the path.

We model each router in this network as having p input ports (I_1, \dots, I_p) and p output ports (O_1, \dots, O_p). The network topology M_n consists of n routers such that, for all i ($1 \leq i \leq n - 1$), the output port O_1 of router R_i is connected to the input port I_1 of router R_{i+1} . The tagged flow enters the network through port I_1 of router R_1 . Through each port I_2, \dots, I_p of router R_i , aggregates enter the network, and $1/p$ of each of these aggregates are routed to output port O_1 (see Figure 2(b)). In addition, the tagged flow entering input port I_1 of R_i is routed to its output port O_1 . Thus, for each router, the traffic routed to the output port O_1 consists of: (1) The tagged traffic (entering the router from port I_1); and (2) $1/p$ of the flows entering from input ports I_2, \dots, I_p . All of the remaining traffic entering through the input ports is routed to output ports, O_2, \dots, O_p .

The above topology ensures that: (1) the tagged flow that enters the network at router R_1 is routed all the way through the multi-hop network M_n , and (2) the cross traffic entering the network at router R_i ($i \in [1, n]$) interferes with the transmission of the tagged traffic for a single hop, and leaves the network at router R_{i+1} . This topology facilitates experimentation with different compositions of the cross traffic and different network depths. We have conducted experiments for p ranging from 8 to 32. We present results for experiments with $p = 8$; these results also hold for higher values of p .

7.1.2 Traffic Sources

We consider two kinds of traffic flows.

- *CBR flows*: CBR is one of the more popular traffic models. Flows with CBR transmissions have been used commonly while transporting voice and video across packet networks, in ATM networks, and in the Virtual Leased Line service in DiffServ networks.
- *Non-CBR flows*: Most traffic sources in the current Internet are not CBR in nature. In fact, the following empirical evidence suggests that a sub-exponential distribution of packet inter-arrival times may model sources in the current Internet well.
 - WWW traffic forms a major portion of traffic on the Internet. The inactive *off* times that form the tail of the inter-arrival times of the WWW traffic are shown to have a sub-exponential nature [7].
 - *telnet* packet inter-arrival times have been shown to have a sub-exponential distribution [21].

For these traffic sources, α is chosen to be 3, in order to ensure that the variance of the inter-arrival times at the source is finite. *We thus ensure that the input traffic to the network is not heavy-tailed.*

7.1.3 Modeling Heterogeneity in Cross Traffic

The extent to which cross traffic entering each router affects the characteristics of tagged flows depends on the *burstiness* of the cross traffic. Burstiness in cross-traffic results from: (1) super-positioning of flows, (2) traffic distortions that result from flows traversing through multiple routers in a network, and (3) heterogeneity in the average inter-arrival times across various flows.

To reasonably approximate traffic distortions, we model the cross traffic entering at each router in the network as consisting of an aggregate of two types of flows: (1) flows that are at the beginning of their routes or have traversed through a *small* (M_1 or M_2) number of routers, and (2) flows that are at the end of their routes or have traversed through a *large* (M_{20}) number of hops. This model closely approximates the current Internet where each backbone router is a small number of hops away from some set of hosts while being far away from some others.

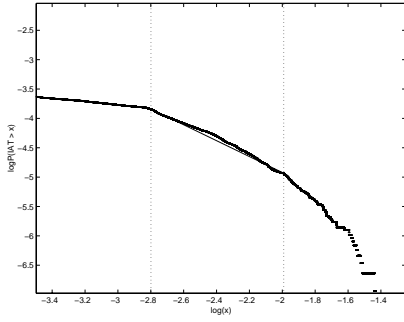
To capture heterogeneity, we consider two classes of cross-traffic flows—flows with *large* and *small* average packet inter-arrival times at the source. We quantify the heterogeneity in the cross-traffic flows in terms of *inter-arrival time ratio* (IATR), which refers to the ratio of the average packet inter-arrival times at the source for these two classes. We derive the two flow classes in three ways.

1. *CBR flows with heterogeneous packet sizes*: We consider CBR flows with same bandwidth requirement but with packet sizes chosen uniformly from two intervals.
2. *CBR flows with heterogeneous bit rates*: We consider flows with fixed packet sizes but with bandwidth requirements chosen uniformly from two intervals.
3. *Flows with heterogeneous sub-exponential inter-arrival times*: The average inter-arrival times of the flows are chosen uniformly from two intervals. Packet sizes are chosen according to an empirically-derived distribution [1].

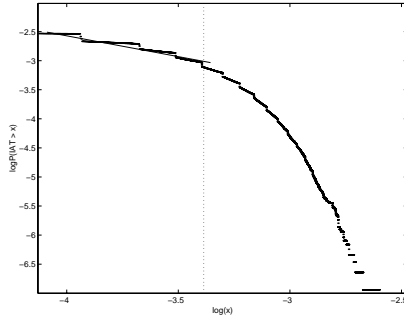
7.2 Model Validation Under Near-asymptotic Conditions

To validate the model presented in Sections 4 and 5, we simulate a network setting where the tagged flow traverses a 20 hop topology with 40 Mbps links operating at 90% utilization. We select a tagged flow with 100 Byte packet size and 4 Mbps bandwidth requirement. We consider cross traffic with an IATR of 500.

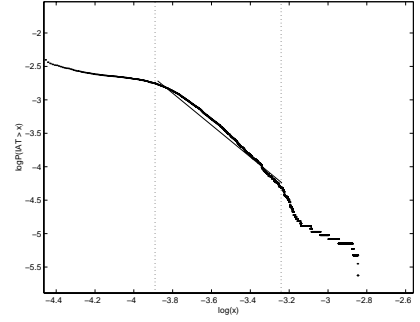
Figure 3 shows the *LLCD* plots of the inter-packet times of the tagged flow after it has traversed the network with the three traffic settings. The slope (i.e., the value of α) of the selected region on all of the *LLCD* plots is between 1 and 2. Figure 4 plots the *Hill* function for the inter-arrival times of the tagged flow obtained for the three traffic settings. The figures illustrate that the *Hill* plots stabilize to a value between 1 and 2, indicating that the values of α



(a) CBR: heterogeneous packet sizes

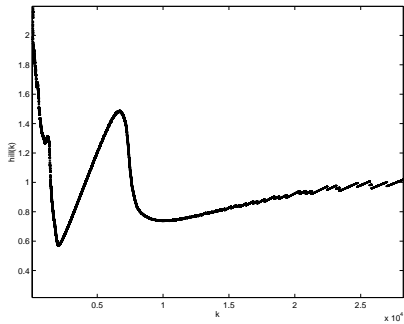


(b) CBR: heterogeneous bit-rates

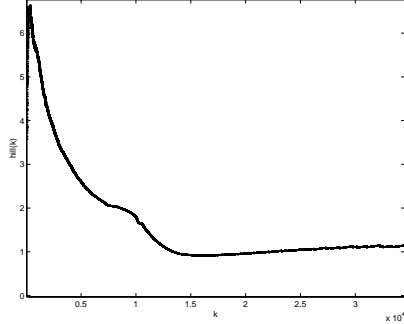


(c) Sub-exponential-Empirical

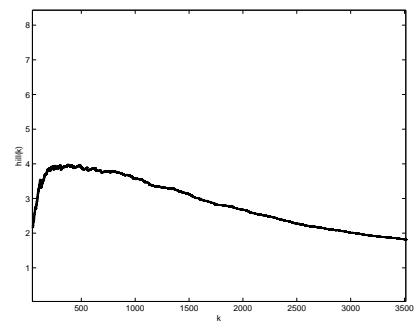
Figure 3: LLCD plots



(a) CBR: heterogeneous packet sizes

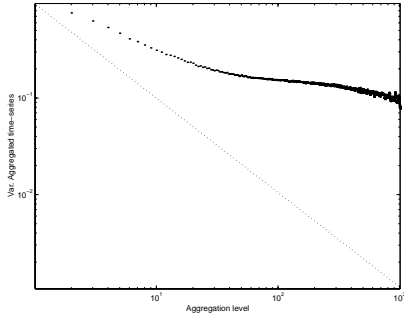


(b) CBR: heterogeneous bit-rates

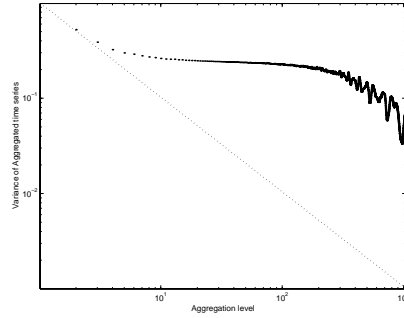


(c) Sub-exponential-Empirical

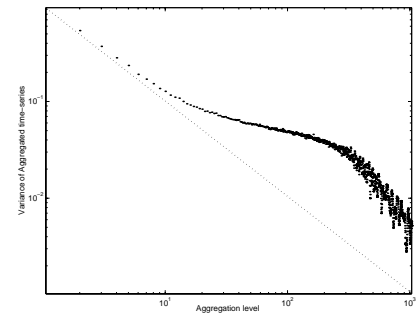
Figure 4: Hill plots



(a) CBR: Heterogeneous packet sizes

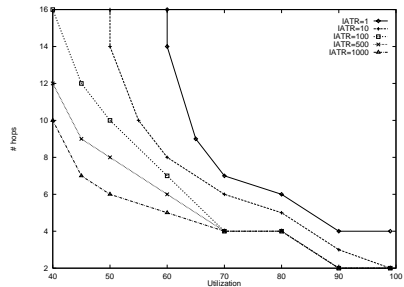


(b) CBR: Heterogeneous bit rates

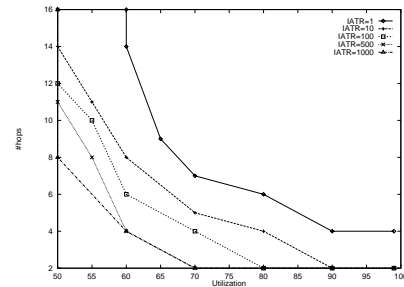


(c) Sub-exponential-Empirical

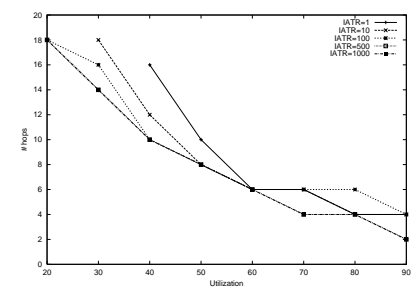
Figure 5: Variance-time plot for cross traffic



(a) CBR: Heterogeneous packet sizes



(b) CBR: Heterogeneous bit rates



(c) Sub-exponential-Empirical

Figure 6: Region of applicability of the model with heterogeneity of cross-traffic

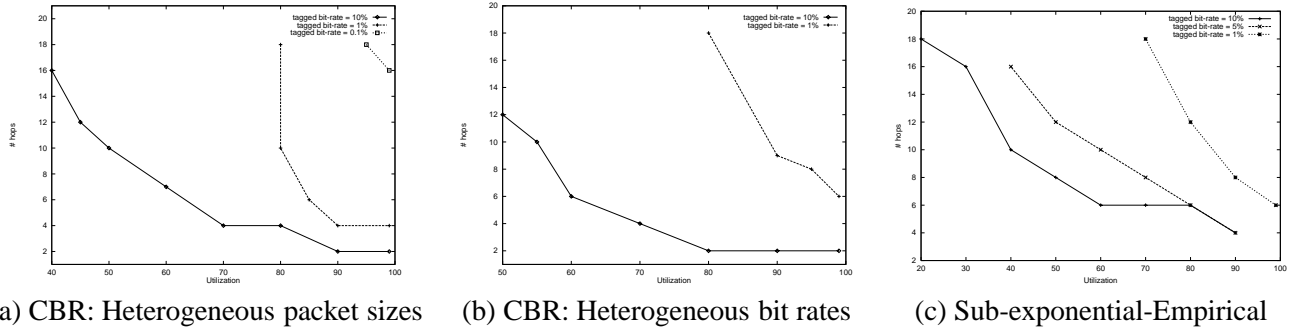


Figure 7: Region of applicability of the model with different bit-rates of tagged flow

are between 1 and 2 as the model predicts. These graphs demonstrate the heavy-tailedness of the packet inter-arrival times of traffic, and thereby validate the predictions of our model.

Figure 5 depicts the variance-time plots of the cross-traffic, for all the three traffic settings. The slope of all of the plots is greater than -1 (for comparison, a line with slope of -1 is drawn), indicating long-range dependence of aggregates.

7.3 Non-asymptotic Region of Applicability

We now use the following *dimensions* of experimental settings to identify the non-asymptotic *region of applicability* of the model in practice.

Network Utilization: To identify the network utilization levels above which flows exhibit heavy-tailed behavior, we conduct experiments with link utilization varying from 20% to 99%.

Network Depth: The greater the number of hops a flow traverses, the larger is the likelihood for its inter-arrival time to approach the limiting distribution, J . We evaluate the effect of increasing the number of hops—from 2 to 20—on the validity of the model for individual flows.

IATR: The greater the heterogeneity in the cross-traffic, the greater is its interference with the tagged flow. We evaluate the effect of this parameter by varying the IATR of the cross-traffic from 1 to 1000.

We conduct experiments with network settings defined by each combination of specific parameter values for utilization, network depth, and IATR. Any network setting that yields a heavy-tailed tagged flow (with $\alpha \in (1, 2)$) at the destination is said to belong to the *region of applicability* of the model.

We vary the IATR of the cross-traffic for each of the three different cross-traffic settings. For the three respective traffic settings, Figures 6(a), 6(b) and 6(c) plot the *minimum* utilization level versus the *minimum* number of router hops at that utilization necessary for observing $\alpha \in (1, 2)$ for the tagged flow. In each of these experiments, the tagged flow occupies 10% of the total traffic sharing a link. For each value of IATR, the area above and to the right of its corresponding curve identifies the *region of applicability* of the model. The following conclusions can be drawn from these figures:

1. The greater the heterogeneity in the cross-traffic, the larger is the region of applicability of the model (i.e., tagged flows show heavy-tailed behavior with $\alpha \in (1, 2)$ for a larger set of network settings).
2. For same value of IATR, heterogeneity in packet-sizes results in a slightly larger *region of applicability* of the model than heterogeneity in bit-rates.
3. More importantly, these figures demonstrate that flows can become heavy-tailed in large heterogeneous networks running at relatively low utilization levels (e.g., 20%), or after traversing as few as 2 hops for networks running at high utilization levels.

The above results hold for the case when the tagged flow occupies 10% of the link bandwidth. To investigate the *region of applicability* of the model when the tagged traffic occupies smaller fractions of total link bandwidth, we conduct the following sets of experiments. We fix the IATR of the cross-traffic to a value of 100. We vary the contribution of the tagged flow to the cumulative bit rate of flows from 0.01% to 10%. For each of these values, we plot the *minimum* utilization level versus the *minimum* number of router hops at that utilization necessary for observing $\alpha \in (1, 2)$ for the tagged flow. Figures 7(a), 7(b) and 7(c) illustrate that when the tagged flow contributes only a small fraction of the total bit rate of aggregate, $\alpha \in (1, 2)$ only at very high levels of utilization and only after the flow has traversed a large number of hops. In fact, when the tagged flow contributes just 0.01% of the bit rate of the aggregate, even after traversing a 20 hop, highly loaded network, the value of α does not belong to the desired interval.

8 Implications for Network Design

8.1 Impact of Heavy-tails

Heavy-tailed behavior of flows can be *contagious*. Consider for instance, a persistent HTTP session in which a client waits for the server response prior to issuing subsequent requests. If the response traffic streams from the server is heavy-tailed, future request generation from the client to the server also becomes heavy-tailed. Heavy-tailed behavior thus spreads from the server traffic streams to the client traffic streams.

To cope with the heavy tailed distribution of inter-arrival times of packets from servers to clients, multimedia clients either have to maintain a play-out buffer large enough to smooth out the variation in inter-arrival times, or have to play out incomplete data. The latter deteriorates client performance, whereas the former requires provisioning of very large buffers to smooth out the heavy tails. A client may, therefore, be willing to pay more for a network service that does not result in flows becoming heavy-tailed. At the server end, managing heavy-tailed behavior of incoming flows requires careful buffer management and request service procedures.

8.2 Impact of Long-range Dependence

Several researchers have studied the impact of long-range dependence on network queues using analytical models and simulations (see [8, 9, 12] and the references there-in). In [12], it has been found that if the input traffic to a network of queues has a large variance in service times, then the rate of approach to steady state can be so slow that the steady state performance may not be of engineering interest. This raises interesting questions on stability in the presence of such traffic. For network queues with a long-range dependent input, sharp increase in queuing delays at fairly low levels of utilization and slow decay of queue lengths have been observed in [8, 9]. This implies that an incremental improvement in loss performance requires a significant increase in buffer size. These raise important questions on source traffic control and buffer/link dimensioning for satisfactory loss performance.

Observe that long-range dependence can occur even if the network backbones are not highly utilized. It is often the case that edge routers run at high levels of utilization—hence, a flow may become heavy-tailed even after crossing just the edge routers. Aggregation of such heavy-tailed flows in the backbone would induce long-range dependence, even if the backbone network is not highly utilized.

In summary, long-range dependent traffic has severe implications on queuing performance. Hence, operating a network so as to maintain traffic outside the long-range dependence domain is highly desirable.

8.3 Routing of Flow Aggregates

Recall that *flows* in our model could themselves be aggregates of micro-flows that traverse the same path from the source (ingress edge router) to the destination (egress edge router)—examples include LSPs in MPLS networks, VPN tunnels, and VPs in ATM Networks. As mentioned in Section 1, such flows can easily occupy a significant fraction of the cumulative bit-rate (for example, around 10%). Our simulation results predict that such flows (that is, aggregates

of micro-flows) would become heavy-tailed at relatively low levels of network utilization and after traversing a small number of network hops. It can easily be shown that, as a result, individual micro-flows constituting these flows also become heavy-tailed.

To prevent such flows from becoming heavy-tailed, it may be desirable to create a large-capacity virtual tunnel (that may aggregate multiple micro-flows) between a source-destination pair as a collection of multiple lower bandwidth tunnels (each carrying a subset of micro-flows) that traverse through different network paths. The lower the bandwidth requirement of each of these tunnels, the smaller is the possibility of flows carried on these tunnels becoming heavy-tailed.

8.4 Considerations for Deploying VLL Service in DiffServ Networks

In a differentiated services network, the Virtual Leased Line (VLL) service will utilize the *Expedited Forwarding (EF)* per-hop-behavior (PHB). It has been suggested that a network could provide the low end-to-end delay requested by the VLL service either (1) by providing bandwidth guarantees—possibly through the use of fair queuing algorithms for bandwidth differentiation across flow aggregates—to the class of flows requesting the EF PHB; or (2) by providing higher priority to flows requesting EF PHB.

If a fair queuing algorithm is used for allocating link bandwidth among flow aggregates, then even if the class of traffic requesting EF PHB occupies a small fraction (e.g., 10-30%) of the total link bandwidth, the overall link utilization can be quite high (due to the presence of flows from other classes as well as best-effort TCP flows—which can utilize all of the available link bandwidth). Consequently, individual CBR flows requesting the VLL service could become heavy-tailed and the total VLL traffic could become long-range dependent.

If, on the other hand, EF traffic is provided higher priority over other flow aggregates, then since EF PHB occupies only a small fraction of the link bandwidth, individual CBR flows requesting VLL service are not likely to become heavy-tailed—provided that individual flows do not occupy significant fractions of the bandwidth allocated to the VLL service class. Hence, the CBR aggregates are likely to remain outside the domain of long-range dependence. This argument suggests that, from the perspective of EF traffic, priority queuing is a better implementation option than fair queuing. Unfortunately, however, with priority queuing, other flow aggregates as well as best-effort TCP flows perceive the network as one with dynamically changing bandwidth. This may affect the throughput and fairness observed by other flows. This may also introduce long-range dependence in the lower priority traffic if the utilization perceived by the lower priority traffic becomes high due to starvation.

9 Related Work

Models for inter-arrival time (Jitter) for flows in general and CBR flows in particular in the ATM environment (homogeneous on the packet size) can be found in [5, 17, 16, 20, 22, 23, 25]. In [17], the authors model a general renewal source with some background traffic and derive expressions for the packet inter-arrival time (IAT) of the renewal source after one hop. The IAT of the tagged renewal stream is expressed as the difference in queue lengths seen by two consecutive packets of the same flow. Expressions are derived for the IAT in the frequency domain under correlated and uncorrelated background traffic. In [16], the authors model the same network for multiple nodes. In [22], the authors model CBR flows with an aggregated CBR background for one hop of the network and extend it to the end-to-end case in [25] under high utilization. The main disadvantage of these models is that they are very complicated, computationally intensive and do not yield any insights into the nature of CBR flows. Further, these studies do not discuss the possibility of self-similarity and long-range dependence in such networks. Our work is different from that of [17, 16, 22, 25] in the following ways: Firstly, our analysis yields closed form results for the inter-arrival times of flows under high utilization. Secondly, it does not make any assumptions about the nature of the cross-traffic entering the network, unlike [17, 16] where the authors evaluate their model for specific cross-traffic models such as Markov Modulated processes.

A more recent work on end-to-end jitter analysis can be found in [3], where it is shown that the jitter of a renewal stream could grow unbounded in the particular example network family the authors consider. Based on their observations, the authors propose a new definition called “Packet Scale Rate Guarantee” for the Expedited Forwarding PHB in DiffServ. In [6, 14], the authors describe a deterministic approach to the end-to-end analysis of renewal flows. They provide a bound on the end to end delays and buffer-sizes for an arbitrary network with FIFO scheduling; but do not discuss end-to-end jitter.

10 Concluding Remarks

In this paper, we develop a simple end-to-end model for packet inter-arrival times of flows in the Internet. Our model illustrates that at high network utilization levels: (1) the packet inter-arrival times of flows becomes heavy-tailed once the flows traverse through a large number of hops, and (2) the aggregate of such flows is a long-range dependent self-similar process.

We validate our model through simulations and examine the region of applicability—with respect to network utilization, heterogeneity of traffic composition, and number of hops that a flow traverses—of our model. We find that although the model predicts the behavior of flows in the asymptotic case, for flows that contribute significant amount ($\approx 10\%$) to the cumulative bit rate of the aggregate traffic, the observations yielded by the model hold at relatively low levels of network utilization (30-40%), small number of network hops (4-6), and moderate levels of heterogeneity in the cross-traffic composition. We point out several important implications of carrying heavy-tailed flows and long-range dependent self-similar aggregates in a network. We argue that a client receiving a heavy-tailed multimedia flow either perceives poor performance or requires large amount of buffers to smooth out the heavy tails. Long-range dependence of traffic aggregates, on the other hand, has several implications on network queuing delays and packet losses. Hence, we conclude that operating a network so as to maintain its traffic outside the heavy-tailed and long-range dependence domain is highly desirable.

References

- [1] <http://www.caida.org>.
- [2] CSIM18 - The Simulation Engine. <http://www.mesquite.com>.
- [3] J.C.R. Bennett, K. Benson, A. Charny, W.F. Courtney, and J.Y.LeBoudec. Delay Jitter Bounds and Packet Scale Rate Guarantee for Expedited Forwarding. In *IEEE INFOCOM*, volume 3, pages 1502–1509, 2001.
- [4] N.H. BINGAM, C.M. Goldie, and J.L. Teugels. Regular Variation, Cambridge University Press. 1989.
- [5] P. Boyer, F. Guillemin, M. Serval, and J.P. Coudreuse. Spacing Cells Protects and Enhances Utilization of ATM Network Links. *IEEE Network*, 6(5), Sept 1992.
- [6] I. Chlamtac, A. Farago, H. Zhang, and A. Fumagalli. A deterministic approach to end-to-end analysis of packet flows in connection-oriented networks. *IEEE/ACM Transaction on Networking*, 6(24):422–431, 1998.
- [7] Mark Crovella and Azer Bestavros. Self-Similarity in the World Wide Web Traffic: Evidence and Possible Causes. *IEEE/ACM Transactions on Networking*, 5(6):835–846, Dec. 1997.
- [8] A. Erramilli, O. Narayan, and W. Willinger. Experimental Queuing analysis with Long-range dependent packet traffic. *IEEE/ACM Transactions on Networking*, 4(2), April 1996.
- [9] M. Grossglauser and J-C. Bolot. On the relevance of long-range dependence in network traffic. *IEEE/ACM Transactions on Networking*, 7(5), 1999.
- [10] M. Grossglauser and S. Keshav. On CBR Service. In *Proceedings of INFOCOM’96*, pages 129–137, March 1996.
- [11] R. Gruenenfelder. A Correlation Based End-to-End Cell Queueing Delay Characterization in an ATM Network. In *Proc. Thirteenth International Teletraffic Congress (ITC-13)*, volume 15, pages 59–64, June 1991.

- [12] D.P. Heyman. Performance Implications of very large service-time variances. *Performance Evaluation*, May 2000.
- [13] V. Jacobson, K. Nichols, and K. Poduri. An Expedited Forwarding PHB. August 1998. available as IETF draft(draft-ietf-diffserv-phb-ef-00.txt).
- [14] J.Y.LeBoudec and G. Hebuterne. Comments on "A deterministic approach to end-to-end analysis of packet flows in connection-oriented networks". *IEEE/ACM Transaction on Networking*, 6(24):422–431, 1998.
- [15] J. Kurose. Open Issues and Challenges in Providing Quality of Service Guarantees in High-Speed Networks. *ACM Computer Communication Review*, 23:6–15, January 1993.
- [16] W. Matragi, C. Bisdikian, and K. Sohraby. Jitter Calculus in ATM Networks: Multiple Node Case. *IEEE INFOCOM '94*, June 1994.
- [17] W. Matragi, C. Bisdikian, and K. Sohraby. Jitter Calculus in ATM Networks: Single Node Case. *IEEE INFOCOM '94*, June 1994.
- [18] M.S.Taqqu, W. Willinger, and R. Sherman. Proof of a fundamental result in self similar traffic modeling. *ACM Computer Communications Review*, 1997.
- [19] K. Nichols, V. Jacobson, and L. Zhang. A Two-bit Differentiated Services Architecture for the Internet. November 1997. <ftp://ftp.ee.lbl.gov/papers/dsarch.pdf>.
- [20] Y. Ohba, M. Murata, and H. Miyahara. Analysis of interdeparture processes for bursty traffic in ATM networks. *IEEE Journal of Selected Areas in Communications*, 8(3), 1990.
- [21] V. Paxson and S. Floyd. Wide Area Traffic: The Failure of Poisson Modeling. *IEEE Transactions on Communications*, 3(3):226–244, June 1995.
- [22] A. Privalov and K. Sohraby. Per-stream Jitter Analysis in CBR ATM Multiplexors. *IEEE/ACM Transaction on Networking*, 6(2), 1998.
- [23] J.W. Roberts and F. Guillemin. Jitter in ATM Networks and its impact on peak-rate enforcement. *Performance Evaluation*, 16:35–48, 1992.
- [24] J. Sahni, P. Goyal, and H.M. Vin. Scheduling CBR Flows: FIFO or Per-Flow Queueing? In *Proceedings of NOSSDAV'99*, June 1999.
- [25] K. Sohraby and A. Privalov. End-to-End Jitter Analysis in Networks of Periodic Flows. In *IEEE INFOCOM*, volume 2, pages 575–583, 1999.
- [26] B. Tsybakov and N. D. Georganas. Self-Similar Processes in Communication Networks. *IEEE Transactions on Information Theory*, 44(5):1713–1725, 1998.
- [27] W. Willinger, M.S. Taqqu, R. Sherman, and D.V. Wilson. Self Similarity Through High Variability: Statistical Analysis of ethernet Traffic at the Source Level. *IEEE/ACM Transactions on Networking*, 5(1), 1997.
- [28] D. Yates, J.F. Kurose, D. Towsley, and M.G. Hluchyj. On per-session end-to-end delay distributions and the call admission problem for real-time applications with QOS requirements. In *Proceedings of ACM SIGCOMM*, pages 160–166, October 1993.