End-to-end Analysis of Packet Inter-arrival Times of CBR Flows

Paper No: 248

Abstract

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

We validate our model through simulations and examine the region of applicability—with respect to network utilization, heterogeneity of CBR flows, and number of hops that a CBR flow traverses—of our model. We find that although the model predicts the behavior of CBR flows in the asymptotic case, for flows that contribute significant amount ($\geq$ 10%) to the cumulative bit rate of the CBR aggregate, the observations yielded by the model hold at relatively low levels of network utilization (40-50%), small number of network hops (6-8), and moderate levels of heterogeneity in the packet sizes and bandwidth requirements of the CBR flows.

1 Introduction

Simplicity of implementation is a central principle for designing scalable, higher-performance networks. High-performance routers traditionally implement First-In-First-Out (FIFO) packet scheduling and Drop Tail buffer management algorithms. FIFO Scheduling, unfortunately, does not provide any isolation to flows; bursty traffic arrival from a flow can significantly affect the performance—delay, packet inter-arrival times, and loss—observed by other flows. One possible solution to this problem is to eliminate bursts within a flow by shaping traffic to *constant bit rate (CBR)* at the network edge. It has been conjectured that if flows are shaped to CBR at the source or at the ingress routers and if the bandwidth available to the aggregate of such CBR flows is at least as large as the cumulative rate requirement of the CBR flows, then FIFO scheduling of packets is sufficient. This sufficiency is generally defined in terms of low end-to-end delay and small variance in packet inter-arrival times at the destinations.

The above philosophy is exemplified by the Virtual Leased Line service model proposed in the context of Differentiated Services (DiffServ) networks [15, 22]. The VLL service model guarantees a reserved rate to a flow. Flows requesting this service are shaped to constant bit rate (CBR) flows, and the packets of the flow are then marked as belonging to a particular service class (*expedited forwarding* service class [15]) by appropriately setting the Type-of-Service (ToS) byte in the IP header of the packet [21]. The routers transmit packets belonging to this service class in the FIFO order. To ensure that the rate guarantees of the flows can be met, the routers employ appropriate scheduling and admission control algorithm to ensure that on any given link, the rate available for this service class is at least the sum of the reserved rates of flows that request the Virtual Leased Line service. Note that this service does not guarantee an upper bound on delay; however, it is expected to closely resemble the service offered by a leased line.

The design and the viability of the Virtual Leased Line and other such service models are based on the conjecture of sufficiency of FIFO scheduling for CBR flows. 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 [12, 30]. Further, many frameworks for analyzing CBR flows that traverse a network have been proposed [17, 18, 26, 31], but no analysis exists that yields any realistic insights on the nature of CBR flows at destinations. This illustrates the need for characterizing the properties—specifically the packet inter-arrival times at destinations—of a CBR flow in a network of FIFO routers; however, the literature contains no such close-form models.

In this paper, we model the inter-arrival time for packets of individual CBR flows as well as CBR aggregates under asymptotic conditions of network utilization and number of hops that a flow may traverse. Our model yields the following key insights.

- The variance in the packet inter-arrival times of individual CBR flow tends to infinity in the asymptotic case, indicating that CBR flows become *heavy-tailed* after traversing a network of FIFO routers.
- The aggregation of CBR flows becomes self-similar, long-range dependent Fractional Brownian Motion (FBM) [34, 20]. It is interesting to note here that it has been proved in [6] that *any* background traffic can be modeled as FBM, though it may not be long-range dependent.

The attractive features of our model are its closed-form conclusions and the simple and intuitive analysis.

We validate our model through extensive simulations and derive the region of applicability—with respect to network and traffic parameters—of the model. We find that although the model predicts the behavior of CBR flows in the asymptotic case, for flows that contribute significant amount ($\geq 10\%$) to the cumulative bit rate of the CBR aggregate, the observations yielded by the model hold at relatively low levels of network utilization (40-50%), small number of network hops (6-8), and moderate levels of heterogeneity in the packet sizes and bandwidth requirements of the CBR flows.

The rest of the paper is organized as follows. In Section 2 we define the concepts of self-similarity, long-range dependence and heavy tailedness and describe the tests to detect these characteristics. In Section 3, we formate the problem of end-to-end analysis of packet inter-arrival times of CBR flows. In Sections 4 and 5, we present the models for individual CBR flows and aggregates of CBR flows, respectively. In Section 6, we describe our simulation setup and present simulation results. In Section 7, we discuss the impact of heavy-tailedness, self-similarity and long-range dependence on the design of networks. We outline related work on modeling CBR traffic in Section 8 and finally summarize our work in Section 9.

2 Background

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

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 $H \in [0, 1]$ is known as the *Hurst* parameter, and $m > 0$.

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 $X^m = (X(0) + X(1) \ldots X(m)) \ / \ m$. 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:

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

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

$$E[X_i] = \mu$$

$$Var(X_i) = \sigma^2$$

$$\rho_i = E[(X_i - \mu)(X_0 - \mu)]/\sigma^2$$

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:

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

where $H > 1/2$. This form of ρ_k corresponds to that of a second-order self-similar process [33].

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

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

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

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

For all practical purposes, $c(x)$ can be approximated by a constant. It can be shown that

$$Var(J) = \infty$$

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.

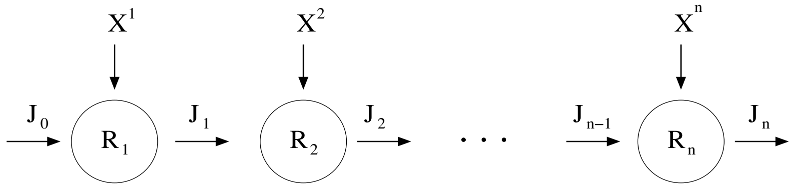

Figure 1: Network Model

2.1 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 ($Prob(J > x)$) 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 $U_1, U_2, \ldots, U_n$ denote the packet inter-arrival times of a given flow in ascending order. The *Hill* function, $hill(k)$, for the flow is defined by:

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

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

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 it demonstrates that $H > 1/2$, indicating that the data is self-similar and long-range dependent.

3 Problem Formulation

Consider a network that carries traffic shaped to constant bit-rates at the sources. Let $\{R_1, \ldots, R_N\}$ denote the routers along a path of length N through this network. We consider networks where each router R_n employs FIFO scheduling to arbitrate access to the link bandwidth. Very large router buffers are assumed which guarantee no loss of packets. Let μ_n be the utilization of the output link of router R_n.

Divide the time-axis into *very small* slots. Consider a CBR flow f that traverses the above path, and let the packet inter-arrival time of this flow at the source be T time-slots. Let the random variable J_n represent the packet inter-arrival time of this flow as its packets leave router R_n. Let J_0 represent the packet inter-arrival time of the flow at the source. Since the flow is shaped to CBR at the source, $J_0 = T$.

Let the random variable X_n denote the time it takes to service the cross-traffic that enters router R_n in any given time-slot. We assume that μ_n is large enough so that the router queue is never empty during any given time-slot.

Then the following relationship exists between J_{n-1}, J_n and X_n:

$$J_n = X_n^1 + \ldots + X_n^{J_{n-1}} \tag{2}$$

where X_n^j is the amount of service in the form of cross-traffic that enters router R_n during the j^{th} time-slot after the arrival of a packet of the flow f. In steady state the distribution of X_n^j is stationary in j. We assume that X_n^j is independent of J_{n-1} – which is a reasonable assumption in a network with diverse routing and a large number of input links carrying cross-traffic.

Note that $EX_n^j = \mu_n$. For simplicity of analysis we assume that the cross-traffic characteristics are uniform over all routers, i.e., let $X_n^j \sim X$, and $\mu_n = \mu$, for $1 \leq n \leq N$. Let σ^2 denote $Var(X)$.

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

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

Notice that with independent cross-traffic 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 CBR 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: Assume the contrary. That is, $\sum_{k=1}^{m} \rho_k$ is either a linear or a super-linear function in m.

$\sum_{k=1}^{m} \rho_k$ cannot be super-linear due to the monotonically decreasing property. Therefore, 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_n under these traffic and network conditions.

4 Model for single CBR flow

Let $P(J_{n-1} = j)$ stand for the marginal distribution of J_{n-1}. From (2) we have:

$$\begin{aligned}
EJ_n &= \sum_{j \geq 0} E[X_n^1 + X_n^2 + .. + X_n^j]\, P(J_{n-1} = j) \\
&= \sum_{j \geq 0} j\, EX\, P(J_{n-1} = j) \\
&= EX\, EJ_{n-1} && (3) \\
&= \mu\, EJ_{n-1} && (4)
\end{aligned}$$

$$\begin{aligned}
Var(J_n) &= E[J_n - EJ_n]^2 \\
&= E[J_n - J_{n-1}EX + J_{n-1}EX - EJ_n]^2 \\
&= E[J_n - J_{n-1}EX]^2 + E[J_{n-1}EX - EJ_n]^2 + 2E[(J_n - J_{n-1}EX)(J_{n-1}EX - EJ_n)] && (5)
\end{aligned}$$

Consider the first term in the right hand side of Equation (5):

$$\begin{aligned}
E[J_n - J_{n-1}EX]^2 &= E[(X_n^1 - EX) + .. + (X_n^{J_{n-1}} - EX)]^2 \\
&= \sum_{j\geq 0} E[(X_n^1 - EX) + .. + (X_n^j - EX)]^2 \; P(J_{n-1} = j) \\
&= \sum_{j\geq 0} (jE[X - EX]^2 + 2(j-1)E[(X_n^1 - EX)(X_n^2 - EX)] + \ldots \\
&\qquad\qquad +2E[(X_n^1 - EX)(EX_n^j - EX)]) \; P(J_{n-1} = j) \\
&= \sigma^2 \sum_{j\geq 0} ((j + 2(j-1)\rho_1 + \ldots + 2\rho_{j-1}) \; P(J_{n-1} = j)) \\
&= \sigma^2 \; E[J_{n-1} + 2(J_{n-1} - 1)\rho_1 + 2(J_{n-1} - 2)\rho_2 + \ldots + 2\rho_{(J_{n-1}-1)}] \qquad (6)
\end{aligned}$$

Now,

$$(J_{n-1} - 1)\rho_1 + (J_{n-1} - 2)\rho_2 + .. + \rho_{(J_{n-1}-1)} \;\; \leq \;\; (J_{n-1} - 1) \sum_{k=1}^{(J_{n-1}-1)} \rho_k \qquad (7)$$

From Lemma 1, we have: $\sum_{k=1}^{m} \rho_k$ is a sub-linear function of m. Using this and (7), we can conclude that $E[J_n - J_{n-1}EX]^2$ does not contain a $Var(J_{n-1})$ term.

The second term in Equation (5) can be simplified as:

$$\begin{aligned}
E[J_{n-1}EX - EJ_n]^2 &= E[J_{n-1}EX - EXEJ_{n-1}]^2 \\
&= (EX)^2 \; E[J_{n-1} - EJ_{n-1}]^2 \\
&= \mu^2 \; Var(J_{n-1}) \qquad (8)
\end{aligned}$$

The third term in Equation (5) simplifies to:

$$\begin{aligned}
E[(J_n - J_{n-1}EX)(J_{n-1}EX - EJ_n)] &= E[(J_n - J_{n-1}EX)(J_{n-1}EX - EJ_{n-1}EX)] \\
&= \mu \sum_{j\geq 0} (j - EJ_{n-1}) \; E[(X_n^1 - EX) + .. + (X_n^j - EX)] \; P(J_{n-1} = j) \\
&= 0 \qquad (9)
\end{aligned}$$

From Equations (5), (6), (8), (9), we get:

$$Var(J_n) \;\; = \;\; \mu^2 Var(J_{n-1}) + \sigma^2 E[J_{n-1} + 2(J_{n-1} - 1)\rho_1 + .. + 2\rho_{J_{n-1}-1}]$$

Let $\lim_{n\to\infty} J_n = J$, i.e., J is the limiting distribution for J_n. Then, we have:

$$Var(J) = \sigma^2 E[J + 2(J-1)\rho_1 + .. + 2\rho_{J-1}]/(1 - \mu^2)$$

We see that $Var(J) \to_{\mu\to 1} \infty$. Thus the variance of the tagged CBR flow grows in an unbounded manner as it traverses a very large number of hops N in the network with utilization μ arbitrarily close to one.

From repeated application of Equation (4) we get

$$\begin{aligned} EJ &= Lim_{n\to\infty} EJ_n \\ &= Lim_{n\to\infty} Lim_{\mu\to 1} \mu^n T \\ &= T \end{aligned}$$

Thus, the limiting distribution of packet inter-arrival time of the tagged flow has a finite mean and infinite variance. If we make the time-slots we consider infinitesimally small so as to approximate a continuous time model, then Lemma 2 shows that if the marginal distribution of such a random variable J is well defined then it is heavy-tailed.

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

$$P(J > x) = c(x)x^{-\rho},\ x \to \infty$$

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

Proof: Consider

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

Here, $p_J(y)$ is the probability density function for J. Observe that,

$$V_\beta(x) \le V_{\beta+\delta}(x) \qquad\qquad \forall \delta > 0 \tag{10}$$

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

$$Lim_{x\to\infty} V_1(x) < \infty \tag{11}$$

$$Lim_{x\to\infty} V_2(x) = \infty \tag{12}$$

From (10), (11),(12), we can conclude:

$$\exists \beta : 1 < \beta \le 2 \qquad Lim_{x\to\infty} V_\beta(x) = \infty \quad and \quad \forall \delta > 0,\ Lim_{x\to\infty} V_{\beta-\delta}(x) < \infty \tag{13}$$

In [3](p.331,Th. 8.1.2) put $\alpha = 0$ and use the converse implication. i.e, let,

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

We assume that $p_J(y)$ is well-defined so that γ exists and $0 \le \gamma \le \infty$.

From [3] and (13), we then have

$$\begin{aligned} P(J > x) &= c(x)x^{-\rho},\ x \to \infty \qquad (14) \\ \rho &= \frac{\beta}{1+\gamma} \\ \rho &\in [0, \beta] \end{aligned}$$

Here, $c(x)$ is slowly varying.

From (14), $EJ < \infty$, and $Var(J) = \infty$, it can be shown that we must have $1 < \rho < 2$. Thus,

$$P(J > x) = c(x)x^{-\rho},\ x \to \infty,\ 1 < \rho < 2$$

■

Hence the packet inter-arrival time of individual CBR flows becomes a heavy-tailed renewal process in the limit and:

$$P(J > y) = c(y)y^{-\alpha},\ y \to \infty,\ 1 < \alpha < 2$$

5 Aggregation of CBR flows

In this section we make use of existing results on the superposition of heavy-tailed renewal processes in order to deduce the characteristics of CBR aggregates on a link.

In [19] and [20] it has been found that a stream that consists of the superposition of heavy-tailed renewal flows and heavy-tailed *on-off* flows, respectively, converges to a Fractional Brownian Motion (FBM) process as the number of flows grows large. FBM is a self-similar process with Hurst parameter

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

Since $1 < \alpha < 2$ it can be seen that $1 > H > 1/2$. Thus the cross-traffic is long-range dependent also.

The increments of an FBM process are represented by Fractional Gaussian Noise (FGN). Therefore, we can approximate the CBR aggregate traffic X as:

$$X^j = \mu + G^j$$

where G^j is the zero mean FGN process. The properties of G^j are given by:

$$\begin{aligned} EG^j &= 0 \\ Var(G^j) &= \sigma^2 \\ E[G^jG^{j+k}] &= \sigma^2((k+1)^{2H} - 2k^{2H} + (k-1)^{2H})/2, k > 0 \end{aligned}$$

Here, H is the Hurst parameter of the FBM (or FGN) process.

The auto-covariances of the cross-traffic, ρ_k, can be computed as:

$$\begin{aligned} \rho_k &= E[(X^j - \mu)(X^{j+k} - \mu)]/\sigma^2 \\ &= E[G^jG^{j+k}]/\sigma^2 \\ &= ((k+1)^{2H} - 2k^{2H} + (k-1)^{2H})/2 \end{aligned}$$

6 Experimental Validation and Analysis

We conduct a simulation study of CBR networks with two objectives: (1) to validate the model developed in Sections 4 and 5; and (2) to identify network settings and traffic compositions where the predicted heavy-tailedness of individual CBR becomes visible. We conduct these experiments using our network simulator, designed using CSIM [1]. In what follows, we first describe our simulation setup and then describe the results of our simulations.

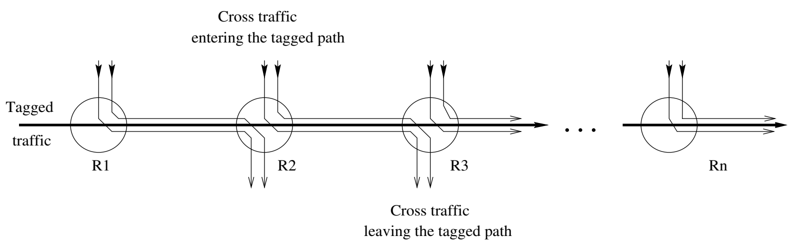

(a) Network Topology: M_n

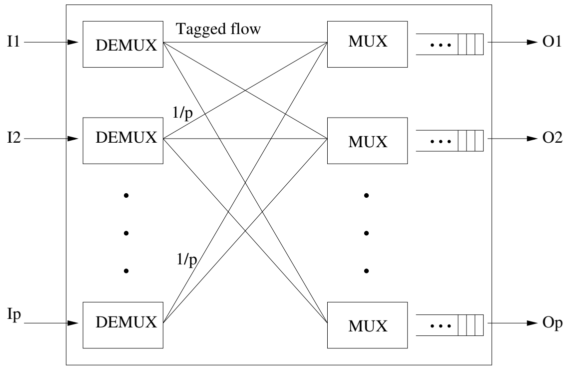

(b) Router Architecture

Figure 2: Simulation environment: network topology and router architecture

6.1 Simulation Environment

6.1.1 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 [13, 16, 17, 35]. 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 CBR 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 CBR 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 ($I_1, \ldots, I_p$) and p output ports ($O_1, \ldots, O_p$). 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 $I_2, \ldots, I_p$ of router R_i, CBR 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 $I_2, \ldots, I_p$. All of the remaining traffic entering through the input ports is routed to output ports, $O_2, \ldots, O_p$.

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.

6.1.2 Modeling Cross Traffic

The extent to which the cross traffic entering each router affects the characteristics of the CBR flows in the tagged traffic depends on the *burstiness* of the cross traffic. Note that, although we have assumed that each flow entering the network is shaped to CBR at the source, the aggregate of CBR flows entering each router may be bursty. This burstiness results from: (1) super-positioning of CBR flows, (2) the inherent heterogeneity in the packet inter-arrival times of the CBR flows, and (3) the traffic distortions that result from flows traversing through multiple routers in a network.

To reasonably approximate this burstiness, we model the cross traffic entering at each router in the network as consisting of an aggregate of two types of CBR 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. Further, we consider cross-traffic CBR flows with heterogeneous packet sizes and bandwidth requirements. We capture the heterogeneity in these two dimensions in terms of the inter-arrival times of packets of CBR flows at the source; for flow f_i, the packet inter-arrival time at the source is given by l_i/r_i, where l_i and r_i denote the packet size and bit rate requirement of flow f_i.

For our simulations, we consider cross-traffic consisting of two classes of flows—CBR flows with *large* and *small* packet inter-arrival times at the source. We capture the heterogeneity in the cross-traffic flows in terms of *inter-arrival time ratio* (IATR), which refers to the ratio in the average packet inter-arrival time at the source for these 2 classes. We derive these two classes of flows as follows.

1. *Heterogeneity in packet sizes*: We consider flows with same bandwidth requirement but with packet sizes chosen uniformly from two different intervals.

2. *Heterogeneity in bit rates*: We consider flows with fixed packet sizes but with bandwidth requirements chosen uniformly from two different intervals.

These methods help determine the relative importance of heterogeneity in packet sizes and bit rate requirements among the CBR flows carried by a network on the validity of our results.

6.2 Model Validation Under Near-asymptotic Conditions

Recall that the model presented in Sections 4 and 5 predicts that under asymptotic conditions—i.e., at high level of network utilization and after traversing a large number of hops—individual CBR flows become heavy-tailed and CBR aggregates become self-similar. To validate the model, we simulate a network setting where the tagged CBR 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 IATR of 100.

Figures 3(a) and 3(b), respectively, plot the marginal distribution for the packet inter-arrival times of the tagged flow with the two settings of cross-traffic outlined above. As the figures illustrate, these distributions have multiple modes; further, the distributions are different from any well-known, analytically-tractable distributions. Hence, a jitter analysis of these flows based on modeling the marginal distribution of packet inter-arrival times does not seem feasible.

Figures 4(a) and 5(a), respectively, are the $LLCD$ plots of the inter-arrival times of the tagged flow after it has traversed the network with the two cross-traffic settings. The slope (i.e., the value of α) of the selected region on

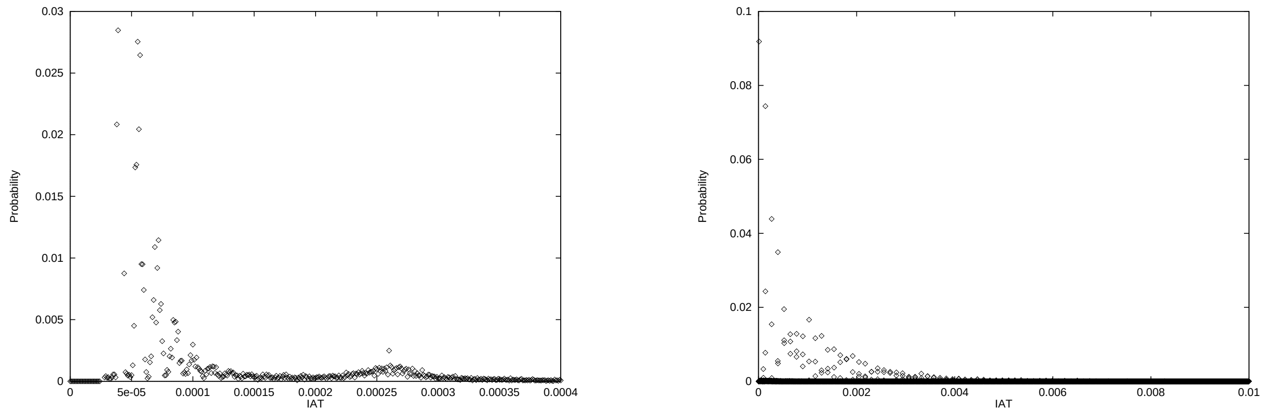

(a) Heterogeneous packet sizes

(b) Heterogeneous bit rates

Figure 3: Distribution of tagged IAT with heterogeneous cross-traffic

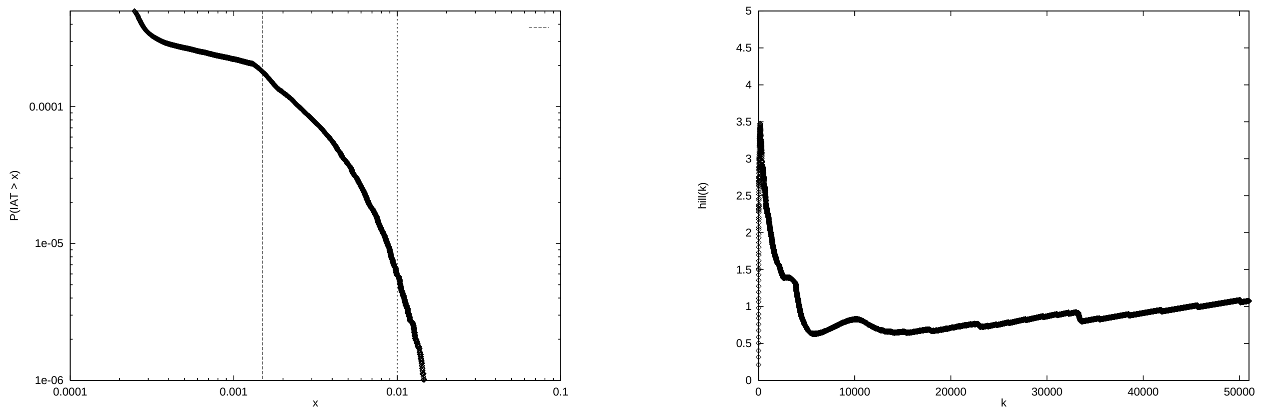

(a) LLCD Plot

(b) Hill Plot

Figure 4: Heavy-tail tests for tagged IAT with heterogeneous packet-size cross-traffic (100B tagged flow packets)

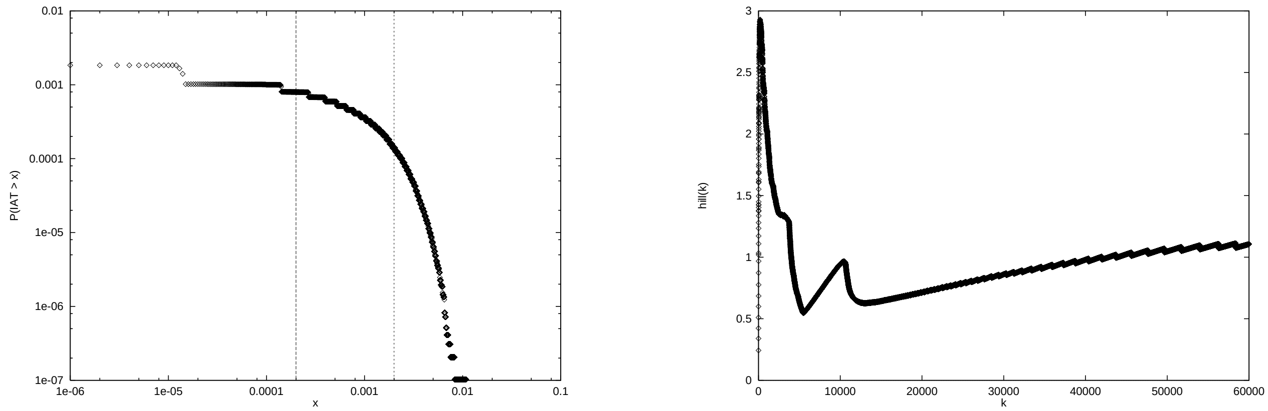

(a) LLCD Plot

(b) Hill Plot

Figure 5: Heavy-tail tests for tagged IAT with heterogeneous bit-rate cross-traffic (100B tagged flow packets)

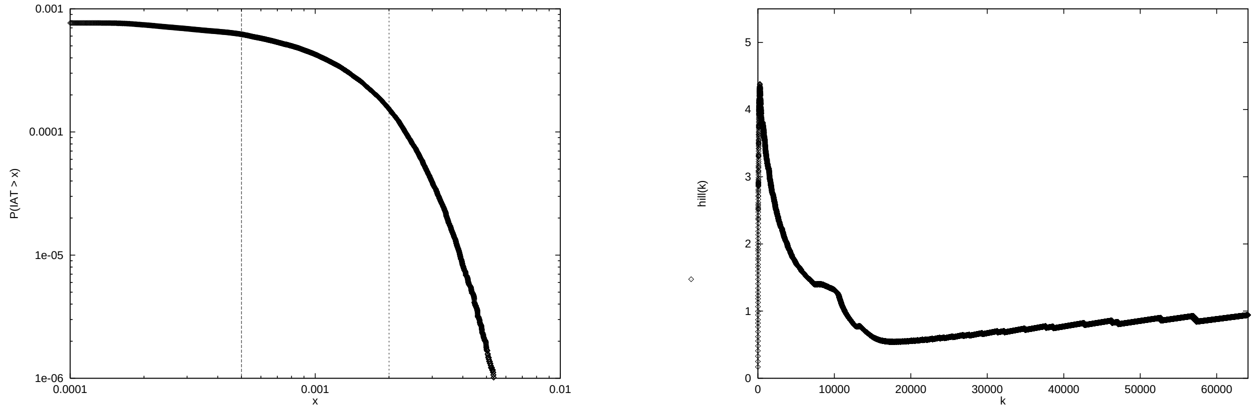

(a) LLCD Plot

(b) Hill Plot

Figure 6: Heavy-tail tests for tagged IAT with heterogeneous packet-size cross-traffic (500B tagged flow packets)

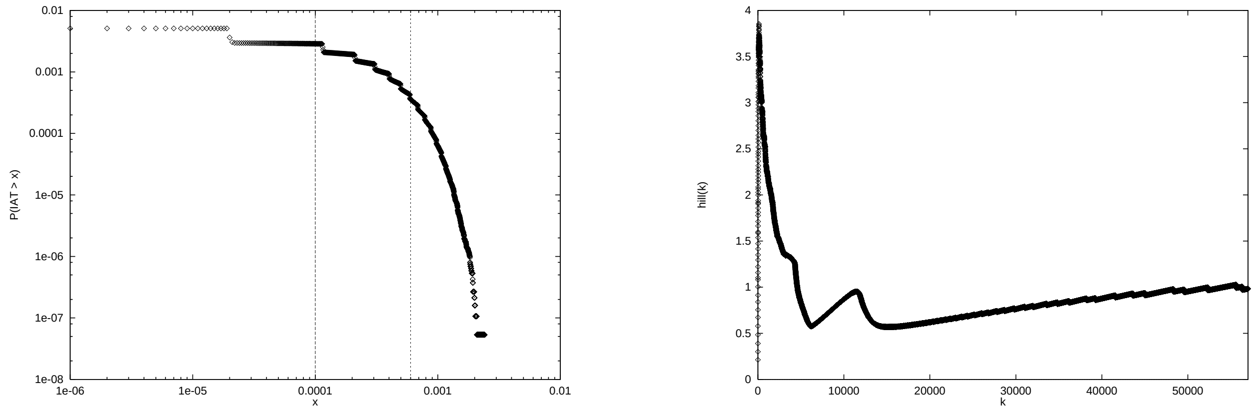

(a) LLCD Plot

(b) Hill Plot

Figure 7: Heavy-tail tests for tagged IAT with heterogeneous bit-rate cross-traffic (500B tagged flow packets)

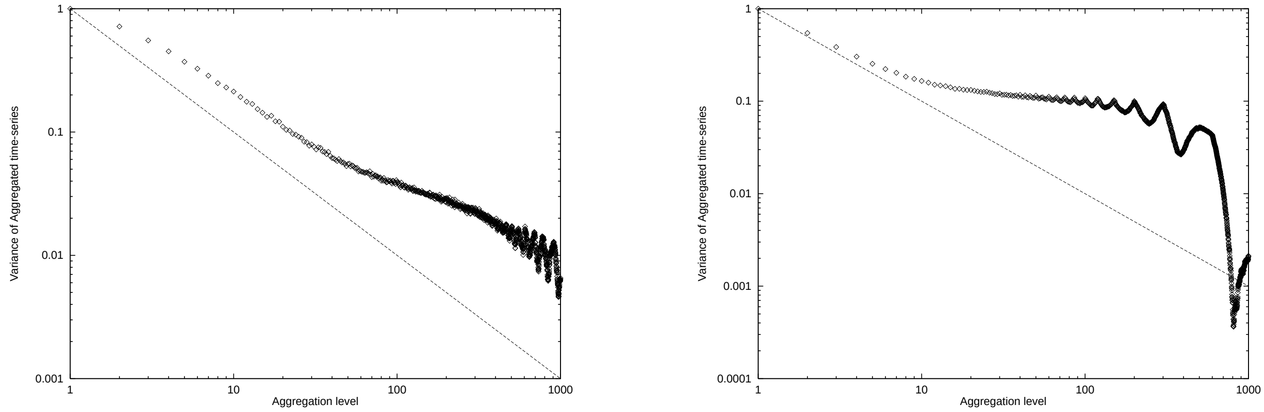

(a) Heterogeneous packet sizes

(b) Heterogeneous bit rates

Figure 8: Variance-time plot for cross traffic

both of the $LLCD$ plots is between 1 and 2. Figures 4(b) and 5(b) plot the $Hill$ function for the inter-arrival times of the tagged CBR flow obtained for the two cross-traffic settings. The figures illustrate that the $Hill$ plots stabilize to a value between 1 and 2, indicating that the values of α are between 1 and 2 as the model predicts. These graphs demonstrate the heavy-tailedness of the packet inter-arrival times of CBR traffic, and thereby validate the predictions of our model.

Figures 8(a) and 8(b) depict the variance-time plots of the cross-traffic, for both cross traffic settings. The slope of both of the plots is greater than -1 indicating long-range dependence of CBR aggregates.

We repeated the above experiment with a different tagged flow—with packet size of 500 Byte instead of 100 Byte. This increases the inter-arrival time of the tagged flow at the source. Figures 6(a), 6(b) and 7(a), 7(b) are the $LLCD$, $Hill$ plots for the inter-arrival time of the tagged CBR flow with the two cross-traffic settings. Notice that, unlike the previous experiment, the observed value of α is slightly smaller than unity. This indicates that, in the limit, even the average packet inter-arrival time should tend to infinity. In practice, however, for flows that transmit a large number of packets, the average inter-arrival time for packets at the destination is roughly equal to the inter-arrival time of packets at the source. This discrepancy can be explained as follows. Although the packet inter-arrival time is finite in practice, it has a large value when expressed in terms of system time slots (which are infinitesimally small). Hence, the packet inter-arrival times at the destination appear to be very large.

6.3 Region of Applicability

Having validated the model for relatively high levels of utilization and large number of network hops, we now use the following *dimensions* of experimental settings to identify the *region of applicability* of the model in practice.

Network Utilization: To identify the network utilization levels above which CBR 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 CBR flow traverses, the larger is the likelihood for its inter-packet spacing to distort from the fixed separation at the source. We evaluate the effect of increasing the number of hops—from 2 to 20—on the validity of the model for individual CBR flows.

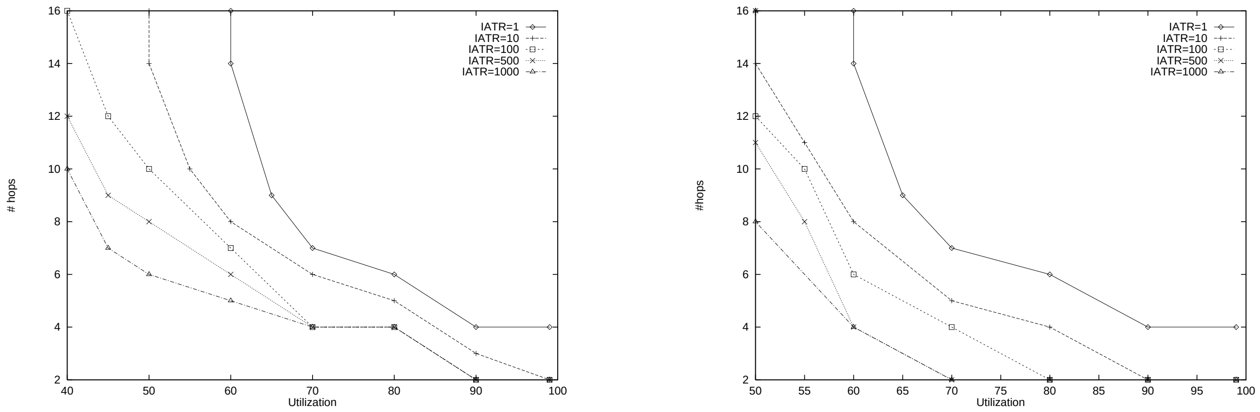

(a) Heterogeneous packet sizes (b) Heterogeneous bit rates

Figure 9: Region of applicability of the model with heterogeneous cross-traffic

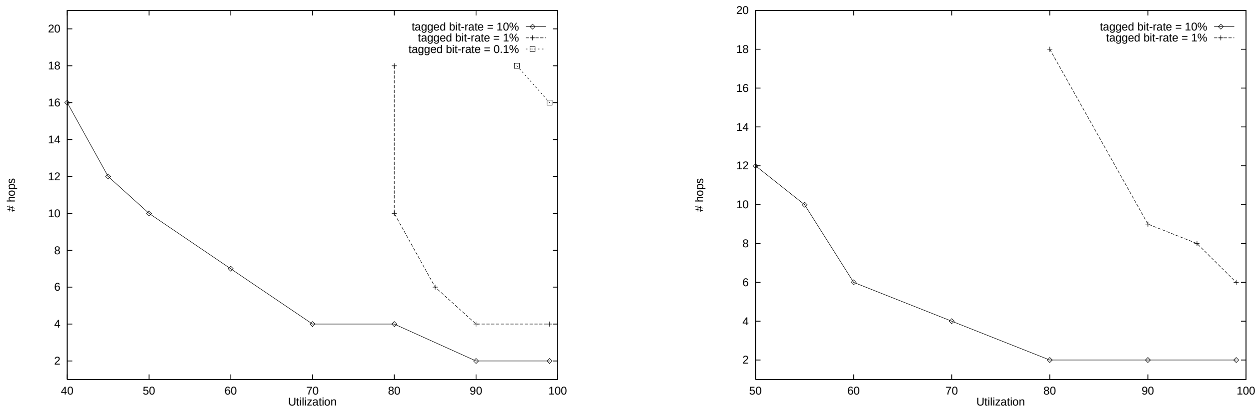

(a) Heterogeneous packet sizes (b) Heterogeneous bit rates

Figure 10: Region of applicability of the model with different bit-rates of tagged flow

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 heavy-tailed tagged flow (with $\alpha \in (1, 2)$) is said to belong to the *region of applicability* of the model.

We vary the IATR by selecting different levels of heterogeneity in either the packet sizes or bit-rate requirements of CBR flows constituting the cross-traffic. For these two sets of experiments, Figures 9(a) and 9(b) plot the variation in 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 CBR flows in the cross-traffic, the larger is the region of applicability of the model (i.e., CBR tagged flows show heavy-tailed behavior with $\alpha \in (1, 2)$ for a larger set of network settings).

2. For the same value of IATR, heterogeneity in packet-sizes results in a slightly larger *region of applicability* of the model than heterogeneity in bit-rates.

More importantly, these figures demonstrate that CBR flows can become heavy-tailed in large heterogeneous networks running at relatively low levels of utilization (e.g., 40%), and after traversing as few as 2 hops for networks running at high levels of utilization.

The above results hold for the case when the tagged flow contributes to 10% of the total traffic sharing a link. Tagged flows that occupy lower shares of the total traffic have larger packet inter-arrival times at the source. As discussed in Section 6.2, the heavy-tailedness tests for these flows are more likely to yield a value of $\alpha < 1$. To investigate the *region of applicability* of the model under smaller fractions of total traffic occupied by the tagged flow, 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 CBR flows from 0.01% to 10%. For each of these values, we plot the variation in 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 10(a) and 10(b) illustrate that when the tagged flow contributes only a small fraction of the total bit rate of CBR 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 CBR aggregate, even after traversing a 20 hop, highly loaded network, the value of α does not belong to the desired interval.

7 Implications on Network Design

7.1 Impact of Long-range Dependence

In the previous sections, we have demonstrated that in a network of FIFO routers, at moderate to high utilization and after traversing a reasonably small number of hops, the packet inter-arrival times of a CBR flow becomes heavy-tailed; further, an aggregate of CBR flows can becomes a self-similar, long-range dependent FBM process. In this section, we examine the implications of such long-range dependent traffic on the network queues and performance.

Several researchers have studied the impact of long-range dependence on network queues using analytical models and simulations [23, 14, 32, 9, 25, 11]. In [23], the author models a queue with long-range dependent FBM input and concludes that the tail behavior is *Weibullian*, rather than exponential as is the case with Markovian traffic. For Weibullian distribution, the tail dies down slower than exponential. In [32], the authors consider queues with constant service times—such as ATM networks—and show that the tail behavior of such queues is hyperbolic provided the holding times of sessions entering the queues is heavy-tailed. In [14], the author considers traffic with very large service time variances (aggregation of such traffic sources lead to long-range dependence) as input to network of queues and concludes that the rate of approach to steady state performance can be so slow that the steady state performance may not be of engineering interest. This raises interesting questions on resource allocation and stability in the presence of such traffic.

The above mentioned analytical queuing models analyze the impact of long-range dependence in the asymptotic case; the authors either assume an aggregation of asymptotically large number of flows or infinite buffer size. In [9], the authors analyze the impact of self-similarity and long-range dependence in traffic on the network queues through trace-driven simulation, with real-life long-range dependent traffic traces. The authors observed a steep increase in queuing delays at fairly small network utilization. To isolate the cause of the delay increase, the authors eliminated short-range dependence from their traffic traces, and repeated the above experiment; they found that the delay behavior observed with this traffic was similar to what was observed with the original traffic trace. This emphasized the significance of impact of long-rage dependent traffic on network performance. Further, the authors observed hyperbolic decay of queue lengths. With hyperbolic decay of queue lengths, obtaining satisfactory loss performance requires very large buffer sizes.

In [11], the authors report that with long-range dependent traffic, an incremental improvement in loss performance requires a significant increase in buffer size. Hence, for a satisfactory loss performance, the authors suggest the use of small buffer sizes in conjunction with large link capacities and source traffic control. This raises important questions on source traffic control and buffer/link dimensioning for satisfactory loss performance.

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

7.2 Considerations for Deploying VLL Service in DiffServ Networks

As we had outlined in Section 1, the Virtual Leased Line (VLL) service and other similar services are being defined under the following conjecture: if flows are shaped to CBR at the source or at the ingress routers and if the bandwidth available to the aggregate of such CBR flows is at least as large as the cumulative rate requirement of CBR flows, then FIFO scheduling of packets is sufficient.

In a differentiated services network, the 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 could become heavy-tailed; further, if individual flows contribute significant amount to the cumulative bit rate of the flows requesting VLL service, then 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 are not likely to become heavy-tailed (and hence, the CBR aggregates will 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. Further, it can be shown that our model for analyzing the behavior of CBR flows generalizes to any renewal stream. Hence, if the network operates at high levels of utilization, it is quite likely that the flow aggregates of lower priority traffic may become long-range dependent.

The above discussion indicates that architecting and deploying differentiated services networks will require careful evaluation of mechanisms for providing service differentiation.

8 Related Work

Behavior of CBR flows has been analyzed using delay models and packet inter-arrival time (or jitter) models. Whereas delay models help compute the delay behavior of queues in the presence of homogeneous and heterogeneous CBR flows, models for predicting packet inter-arrival times determine the nature of output flows from such queues.

Various delay models for CBR traffic exist in the literature. In [8, 29], the authors analyze a queue with homogeneous CBR traffic input and analytically obtain the queue size distribution ($nD/D/1$ model). In [29], the authors also obtain the upper and lower bounds on the queue size for heterogeneous traffic ($\sum D_i/D/1$). Similar analysis can also be found in [2, 27, 5]. In [10] it is shown that bursts can be formed simply by superposing CBR streams under FIFO scheduling. In [13], the author studies a connection going through multiple hops and observes that the end-to-end delay distribution depends on the auto-covariance of the cross-traffic sharing the network links.

A number of studies have also conducted end-to-end delay analysis through simulations. In [30, 12], the authors examine the end-to-end delay and the buffer sizes needed for the restoration of CBR property at the destination using FIFO and WFQ scheduling policies. Further, in [12], the authors provide simple analytical models for estimating end-to-end delay of CBR sources based on Multi-class parametric decomposition method. [7] shows that the inter-arrival time of packets becomes more bursty when a flow interacts with cross traffic with highly variable packet sizes. However, this study focuses on non-CBR traffic, and is not obvious if the same result would hold for CBR networks. Furthermore, the analysis is carried out using a single cross traffic stream and a small network.

Comprehensive models for CBR packet inter-arrival time (Jitter) for ATM environment (homogeneous on the packet size) can be found in in [4, 18, 17, 24, 28, 26, 31]. In [18], 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 [17], the authors model the same network for multiple nodes. In [26], the authors model CBR flows with a aggregated CBR background for one hop of the network and extend it to the end-to-end case in [31] 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 [18, 17, 26, 31] in the following ways: Firstly, our analysis is simple and intuitive. Secondly, it does not make any assumptions about on the nature of the cross-traffic entering the network, unlike [18, 17] where the authors evaluate their model for specific cross-traffic models such as Markov Modulated processes. Thirdly, we provide a definitive answer to the question of what the CBR flows transform to under certain network conditions and demonstrate that an aggregation of such transformed flows can be modeled as long-range dependent FBM.

9 Concluding Remarks

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

We validate our model through simulations and examine the region of applicability—with respect to network utilization, heterogeneity of CBR flows, and number of hops that a CBR flow traverses—of our model. We find that although the model predicts the behavior of CBR flows in the asymptotic case, for flows that contribute significant amount ($\geq$ 10%) to the cumulative bit rate of the CBR aggregate, the observations yielded by the model hold at relatively low levels of network utilization (40-50%), small number of network hops (6-8), and moderate levels of heterogeneity in the packet sizes and bandwidth requirements of the CBR flows. We discuss the impact of our findings on the performance of networks that carry CBR flows as well as describe some considerations for deploying VLL service in DiffServ networks.

References

[1] CSIM18 - The Simulation Engine. http://www.mesquite.com.

[2] N. Bambos and J. Walrand. On Queues with Periodic Inputs. *Journal of Applied Probability*, 26:381–189, 1989.

[3] N.H. BIngam, C.M. Goldie, and J.L. Teugels. Regular Variation, Cambridge University Press. 1989.

[4] P. Boyer, F. Guillemin, M. Servel, and J.P. Coudreuse. Spacing Cells Protects and Enhances Utilization of ATM Network Links. *IEEE Network*, 6(5), Sept 1992.

[5] I. Cidon, R. Geurin, I. Kessler, and A. Khamisy. Analysis of Statistical Multiplexer with Generalized Periodic Source. *QUESTA*, 20:139–169, 1995.

[6] C. Courcoubetis, A. Dimakis, and G.D. Stamoulis. Traffic Equivalence and Substitution in a Multiplexer. In *IEEE INFOCOM*, 1999.

[7] A. DeSimone. Generating Burstiness in Networks: A Simulation Study of Correlation Effects in Networks of Queues. *ACM Computer Communication Review*, pages 24–31, 1991.

[8] L.G. Dron, G. Ramamurthy, and B. Sengupta. Delay Analysis of Continuos bit-rate traffic over an ATM network. *IEEE Journal of Selected Areas in Communications*, 9:402–407, April 1991.

[9] 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.

[10] S.J. Golestani. Congestion-Free Transmission of Real-Time Traffic in Packet Networks. In *Proceedings of IEEE INFOCOM' 90*, pages 527–536, June 1990.

[11] M. Grossglauser and J-C. Bolot. On the relevance of long-range dependence in network traffic. *IEEE/ACM Transactions on Networking*, 7(5), 1999.

[12] M. Grossglauser and S. Keshav. On CBR Service. In *Proceedings of INFOCOM'96*, pages 129–137, March 1996.

[13] 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.

[14] D.P. Heyman. Performance Implications of very large service-time variances. *Performance Evaluation*, May 2000.

[15] V. Jacobson, K. Nichols, and K. Poduri. An Expedited Forwarding PHB. August 1998. available as IETF draft(draft-ietf-diffserv-phb-ef-00.txt).

[16] 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.

[17] W. Matragi, C. Bisdikian, and K. Sohraby. Jitter Calculus in ATM Networks: Multiple Node Case. *IEEE INFOCOM '94*, June 1994.

[18] W. Matragi, C. Bisdikian, and K. Sohraby. Jitter Calculus in ATM Networks: Single Node Case. *IEEE INFOCOM '94*, June 1994.

[19] M.S.Eberlein and M.S.Taqqu. Dependence in Probability and statistics. 1986.

[20] M.S.Taqqu, W. Willinger, and R. Sherman. Proof of a fundamental result in self similar traffic modeling. *ACM Computer Communications Review*, 1997.

[21] K. Nichols, S. Blake, F. Baker, and D. Black. Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers. December 1998. Internet RFC 2474.

[22] 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.

[23] I. Norros. A Storage model with self-similar input. *QUESTA*, 16:387–396, 1994.

[24] 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.

[25] K. Park, G. Kim, and M. Crovella. On the Effect of Traffic Self-similarity on network performance. In *SPIE International Conference on Performance and Control of Network Systems*, 1997.

[26] A. Privalov and K. Sohraby. Per-stream Jitter Analysis in CBR ATM Multiplexors. *IEEE/ACM Transaction on Networking*, 6(2), 1998.

[27] G. Ramamurthy and B. Sengupta. Delay Analysis of the packet voice multiplexer by the $\sum D_i/D/1$ queue. *IEEE Transactions on Communications*, 39(7):1107–1114, July 1991.

[28] J.W. Roberts and F. Guillemin. Jitter in ATM Networks and its impact on peak-rate enforcement. *Performance Evaluation*, 16:35–48, 1992.

[29] J.W. Roberts and J.T. Virtamo. The Superposition of Periodic Cell Arrival Streams in an ATM Multiplexer. *IEEE Transactions on Communications*, 39(2):298–303, February 1991.

[30] J. Sahni, P. Goyal, and H. Vin. Scheduling CBR Flows: FIFO or Per-Flow Queueing? In *Proceedings of NOSSDAV'99*, June 1999.

[31] K. Sohraby and A. Privalov. End-to-End Jitter Analysis in Networks of Periodic Flows. In *IEEE INFOCOM*, volume 2, pages 575–583, 1999.

[32] B. Tysbakov and N. D. Georganas. Overflow Probability in an ATM queue with Self-Similar input traffic. In *IEEE ICC*, volume 2, pages 822–826, 1997.

[33] B. Tysbakov and N. D. Georganas. Self-Similar Processes in Communication Networks. *IEEE Transactions on Information Theory*, 44(5):1713–1725, 1998.

[34] 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.

[35] 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.