A SOURCE TRAFFIC MODEL AND ITS TRANSIENT ANALYSIS FOR NETWORK CONTROL

N.G. DUFFIELD
AT&T Laboratories
Rm A175, 180 Park Avenue
Florham Park, NJ 07932-0971
duffield@research.att.com

W.WHITT
AT&T Laboratories
Rm A117, 180 Park Avenue
Florham Park, NJ 07932-0971
wow@research.att.com

Stochastic Models 14 (1998) 51–78

ABSTRACT

Traffic measurements from communication networks have shown that network traffic is quite complex, exhibiting phenomena such as long-tail probability distributions, long-range dependence and self similarity. Thus, in order to design and control new communication networks, we are motivated to consider new source traffic models and new ways to analyze network performance when there are many independent sources each with traffic that can be described by such models. We present a candidate source traffic model: The required bandwidth (arrival rate) as a function of time for each source is represented as the sum of two stochastic processes: (1) a macroscopic (longer-time-scale) level process and (2) a microscopic (shorter-time-scale) within-level variation process. We let the level process be a finite-state semi-Markov process (SMP), allowing general (possibly long-tail) level holding-time distributions, and we let the within-level variation process be a zero-mean piecewise-stationary process. An important special case is the traditional on-off traffic model, where the on and off times are allowed to have long-tail probability distributions. (Then there is no within-level variation process.) To cope with the added model complexity, we suggest making design and control decisions based on the likelihood that aggregate demand (the input rate from a set of sources) will exceed capacity (the maximum possible output rate), using a specification of the sources and their source traffic models to predict demand. This approach to model analysis avoids the customary queueing detail (focus on buffer content and overflow). The presence of substantial variation in a long time scale suggests that the current network state can be very useful for predicting network behavior over shorter time scales. Thus we propose using transient analysis for control, conditioning on the level and elapsed level holding time of each source. In doing so, we exploit asymptotics associated with multiplexing a large number of sources. A conditional law of large numbers supports approximating the future aggregate demand conditional on current state information by its conditional mean value. The conditional aggregate mean can be expressed compactly in terms of its Laplace transform and efficiently calculated by numerical transform inversion. When level holding times are long relative to the times of interest for control, we propose a single-transition approximation that can be computed without numerical transform inversion.

Key Words: source traffic model, admission control, congestion control, overload control, transient analysis, deterministic fluid approximation, long-tail distributions, Laplace transforms, numerical transform inversion, statistical multiplexing, value of information

1. Introduction

It is a pleasure to participate in this special issue honoring Marcel Neuts. One of the highlights of Marcel's very productive career is his systematic study of markov chains of GI/M/1 and M/G/1 type, which is well set forth in Neuts [29]. These structured Markov chains are generalizations of the embedded Markov chains in the classical GI/M/1 and M/G/1 queueing models. These structured Markov chains allow us to model queues with very complex traffic (arrival and service processes) and yet still be able to compute the performance measures of interest. They take us well beyond the basic Poisson modeling found in elementary textbooks.

The need for more flexible performance models is supported by recent traffic measurements in communication networks. The traffic measurements have shown that network traffic is quite complex, exhibiting phenomena such as long-tail probability distributions, long-range dependence and self-similarity; e.g., see Cáceres, Danzig, Jamin and Mitzel [6], Leland, Taqqu, Willinger, and Wilson [25], Paxson and Floyd [30], Crovella and Bestavros [11] and Feldmann [15].

Even with the greatly improved modeling and analysis capabilities provided by structured Markov chains, we are faced with a serious problem in developing source traffic models that can be both realistically fit to data and successfully analyzed.

In this paper, we attempt to address this problem in two ways: First, we propose a specific traffic source model and, second, we propose ways to analyze systems with many independent sources each represented by this model (allowing different parameters for different sources). In order to capture important traffic features, we allow the model to be relatively complex; e.g., we allow periods of high demand in each source to have long durations; i.e., these periods may have long-tail (also known as heavy-tail or fat-tail) distributions such as the Pareto distribution. Such long-tail distributions are known to cause self-similarity in aggregate traffic; see Willinger, Taqqu, Sherman and Wilson [34].

On the other hand, to achieve the required analyzability with this added model complexity, we propose a simplified kind of analysis. In particular, we avoid the customary queueing detail (focus on buffer content and overflow) and instead focus only on the probability that aggregate demand (the input rate from a collection of sources) exceeds capacity (the maximum possible output rate) at any time. That is, we propose an offered-load or infinite-capacity approximation; see Duffield and Whitt [12], Jennings, Mandelbaum, Massey and Whitt [23] and Leung, Massey and Whitt [26] for related work in this direction. (This approach also can generate approximations describing loss and delay with finite capacity; e.g., see [23] and Section 5 of [12].)

A key to being able to analyze the full system with sources represented by our general source traffic model is exploiting asymptotics associated with multiplexing a large number of sources. The ever-increasing network bandwidth implies that more and more sources will be able to be multiplexed. As the scale increases, describing the detailed behavior of all sources becomes prohibitively difficult, but fortunately it becomes easier to describe the aggregate, because the large numbers produce statistical regularity. As the size increases, the aggregate demand can be well described by laws of large numbers, central limit theorems and large deviation principles. In particular, our principal approximation in this paper is to approximate the aggregate demand by its mean; that step is justified by the law of large numbers.

We have in mind two problems: First, we want to do capacity planning and, second, we want to do real-time connection admission control and congestion control. The present paper is devoted to the second problem. However, in both cases, we want to determine whether any candidate capacity is adequate to meet the aggregate demand associated with a set of sources. In both cases, we represent the aggregate demand simply as the sum of the bandwidth requirements of all sources. In forming this sum, we regard the bandwidth processes of the different sources as probabilistically independent.

The performance analysis for capacity planning is coarser, involving a longer time scale. Hence, for capacity planning we propose a *steady-state* analysis. When we consider connection admission control and congestion control, we suggest focusing on a shorter time scale. We are still concerned with the relatively long time scale of connections, or the times between scene changes in video, instead of the shorter time scales of cells or bursts, but admission control and congestion control are sufficiently short-term that we propose focusing on the *transient* behavior of the aggregate demand process.

To a large extent, the transient analysis for control that we propose is also supported by the structure revealed by the traffic measurements. The presence of long-range dependence and long-tail holding-time distributions tends to increase the opportunity to exploit the demand history to predict the demand in the not-to-distant future. For example, for an exponentially distributed lifetime in progress, knowledge of the elapsed lifetime does not help predict the remaining lifetime, by the lack-of-memory property of the exponential distribution. In contrast, for a long-tail lifetime distribution such as the Pareto distribution, a long elapsed lifetime is a very strong indication of a long remaining lifetime. Our approach using transient analysis is designed to exploit this property.

In fact, even for capacity planning, the transient analysis plays an important role. The transient analysis determines how long it takes to recover from rare congestion events, as we discussed in our previous paper [12]. If the recovery time from overload is relatively long, then we may elect to provide extra capacity (or reduce demand) so that overload becomes less likely. In this paper we only discuss transient analysis for control, but in a companion paper [13] we discuss transient as well as steady-state analyses with the source model introduced here for capacity planning.

The specific model we propose has two components. For each source, the required bandwidth (input rate) as a function of time $\{B(t) : t \geq 0\}$ is represented as the sum of two stochastic processes: (1) a macroscopic (longer-time-scale) *level process* $\{L(t) : t \geq 0\}$ and (2) a microscopic (shorter-time-scale) *within-level variation process* $\{W(t) : t \geq 0\}$, i.e.,

$$B(t) = L(t) + W(t) , \quad t \geq 0 . \tag{1.1}$$

These processes are continuous-time processes, and thus are (general) stochastic fluid models, giving a flow rate at each time t. We let the macroscopic level process $\{L(t) : t \geq 0\}$ be a *semi-Markov process* (SMP) as in Chapter 10 of Çinlar [10] or Chapter 9 of Heyman and Sobel [22]; i.e., the level process is constant except for jumps, with the jump transitions governed by a Markov process, while the level holding times (times between jumps) are allowed to have general distributions depending on the originating level and the next level. Given a transition from level j to level k, the holding time in level j has cumulative distribution function (cdf) F_{jk}. Conditional on the sequence of successive levels, the holding times are mutually independent. To be consistent with traffic measurements, we allow the holding-time cdf's F_{jk} to have long tails.

We assume that the within-level variation process $\{W(t) : t \geq 0\}$ is a zero-mean piecewise-stationary process. During each holding-time interval in a level, the within-level variation process is an independent segment of a zero-mean stationary process, with the distribution of each segment being allowed to depend on the level. We allow the distribution of the stationary process segment to depend on the level, because it is natural for the variation about any level to vary from level to level. As a concrete example, the within-level variation process segments could be chosen to be segments from independent stationary Ornstein-Uhlenbeck (O-U) diffusion processes. Then the two O-U parameters (drift and diffusion parameters) for each segment can be a function of the level. Alternatively, the stationary segments could be increments of fractional Brownian motion. In choosing within-level processes for our model, we do not necessarily require that the nonnegativity $B(t) \geq 0$ hold. However, we only exploit a limited partial characterization of the within-level variation process. Fortunately, it turns out that the fine structure of the within-level variation process plays no role in our analysis. Indeed, that is one of our main conclusions.

Our purpose in this paper is to show how this source model can be used to determine the likelihood that aggregate demand will exceed capacity at future times, so that network control can be performed. We intend to discuss the important problem of model fitting in a subsequent paper. Nevertheless, it seems worthwhile to briefly consider how the level process might be fit to data. One way to fit the level process to data exploits a window length $w \equiv w(L)$ and a threshold $\tau \equiv \tau(L)$ for each level (both positive numbers). Assume that we start in some level L_0. We then look at the observed (empirical) input rates over successive windows (i.e., the number of arrivals divided by the window length). The successive windows could be jumping windows $((k-1)w, kw]$

for positive integers k) or sliding windows ($(t, t+w]$ for positive real t). We construct the level process sample path by staying in the initial level until the first time the observed input rate differs from the initial level by more than the threshold $\tau(L_0)$. At that time the level is made equal to this new average input rate falling outside the interval $[L_0 - \tau(L_0), L_0 + \tau(L_0)]$. And we proceed from there recursively.

The specified construction shows how the level process $\{L(t) : t \geq 0\}$ can be fit to data, i.e., a sample path of $\{B(t) : t \geq 0\}$. Then the within-level variation process can be *defined* by taking (1.1) as the definition. Of course, with this empirical construction, there is no guarantee that the within-level variation process will have either mean zero or the required stationarity. The additional properties are approximations in the spirit of the construction. As a refinement to the construction above, we can redefine the level between transitions as the average rate there. Then the within-level process will have zero mean.

In fact, in several examples of processes which we envisage modeling by these methods, there will only be the level process. First, the level process may be some smoothed functional of a raw bandwidth process. This is the case with algorithms for smoothing stored video by converting into piecewise constant rate segments in some optimal manner subject to buffering and delay constraints; see Salehi, Zhang, Kurose and Towsley [32]. With such smoothing, the input rate will directly be a level process as we have defined it. Alternatively, the level process may stem from rate reservation over the period between level-shifts, rather than the bandwidth actually used. This would be the case for Renegotiated Constant Bit Rate (RCBR) of Grossglauser, Keshav and Tse [21]. In this situation we act as if the reservation level is the actual demand, and thus again have a level process.

Thus, an important special case of our model is the level process alone, and much of our analysis is focused only on the level process. An important special case of the level process is the two-valued on-off process, which has positive rate in one level and zero rate in the other. When the off times are relatively long compared to on times, superpositions of independent and identically distributed (i.i.d.) on-off sources are well approximated by the $M/G/\infty$ model, which we analyzed in our previous paper [12]. The holding-time distributions in the $M/G/\infty$ model is an average of the on-time distributions; the Poisson arrival rate λ is chosen to make the offered loads match; e.g., if m_1 and m_2 are the mean off and on times for each of n homogeneous sources, then the offered loads in the two models are λm_2 and $nm_2/(m_1 + m_2)$. It is well known that as

$n \to \infty$ with the total offered load held fixed that the M/G/∞ approximation is asymptotically correct. Hence in [12] we already analyzed an important special case of the model introduced here.

Our level process model is also a generalization of Markov-modulated fluid models such as in Anick, Mitra and Sondhi [4], Elwalid and Mitra [14] and Roberts [31]. First, the multi-level level process here can apply to individual sources, not just the aggregate. Second, in the level process here we allow the level holding-time distribution to be non-exponential. Traffic measurements indicate that the level holding-time distributions should often be far from exponential. Moreover, as shown in Section 8 of [12], when the level holding-time distribution has a long-tail, we are able to strongly exploit the age of the level holding times for prediction.

The remainder of this paper is devoted to showing how to do transient analysis with the source traffic model. As in our previous paper [12], we suggest focusing on the future time-dependent mean conditional on the present state. The present state of each level process consists of the level and age (elapsed holding time in that level). Because of the anticipated large number of sources, the actual bandwidth process should be closely approximated by its mean, by the law of large numbers (LLN). As in [12], the conditional mean can be thought of as a deterministic fluid approximation; e.g., see Chen and Mandelbaum [7]. Since the within-level variation process has mean zero, the within-level variation process has no effect upon this conditional mean. Hence, the conditional mean of the aggregate bandwidth process is just the sum of the conditional means of the component level processes.

Unlike the more elementary M/G/∞ model considered in [12], however, the conditional mean here is not available in closed form. Nevertheless, we show that the Laplace transform of the conditional mean aggregate demand can be expressed concisely, so that the conditional mean itself can be very efficiently computed by numerically inverting its Laplace transform; see Section 3. To carry out the inversion, we use the Fourier-series method in Abate and Whitt [1, 2] (the algorithm Euler exploiting Euler summation), but other alternative methods could be used, such as the Gaver-Stehfest and Post-Widder algorithms reviewed in [1] or the Laguerre-series method in [3]. The inversion of the conditional mean here is remarkably tractable, being analogous to the inversion of the Laplace transforms of the time-dependent mean of reflected Brownian motion and the renewal function (with respect to time), discussed in Sections 10 and 13 of [1].

With our approach, we are able to treat much much larger models than we can for queueing models. For example the inversion algorithms to compute steady-state and transient distributions in the BMAP/G/1 queue in Choudhury, Lucantoni and Whitt [8] and Lucantoni, Choudhury and Whitt [27] are effective for only about 100 environment states in the batch Markovian arrival process (BMAP). That restriction permits 100 homogeneous (one class) Markovian on-off sources or 10 each of two classes of homogeneous Markovian on-off sources. In contrast, here we can easily treat 1000 heterogeneous non-Markovian on-off sources.

Indeed, we can avoid the inversion entirely and treat much larger models if we can assume that the level holding times are relatively long compared to the times of interest for control. Under that condition, we can apply a single-transition approximation, which amounts to assuming that the Markov chain is absorbing after one transition. Then the conditional mean is directly expressible in terms of the level holding-time distributions; see Section 4. Simplified inversion algorithms can also be developed by considering only a few transitions.

When multiple transitions are relevant, the numerical inversion remains a viable alternative. An essential initial step in the transform inversion, though, is obtaining the Laplace transform of the conditional residual level holding times in each state given the age for each source. Unfortunately, in general, it is not straightforward to compute the Laplace transform of a conditional distribution given the Laplace transform of a unconditional distribution. Moreover, the Laplace transform of the unconditional distribution is unknown for some unconditional distributions of interest, such as the Weibull distribution. Nevertheless, we obtain practical solutions to these problems. First, we exploit the fact that many classes of distributions satisfy a closure property with respect to conditioning on the age. A conditional phase-type distribution of order m, obtained by conditioning a phase-type distribution of order m upon its age is again a phase-type distribution of order m; see p. 45 of Neuts [28]. Similarly, conditional hyperexponential, Pareto, uniform and deterministic distributions are of the same form. Second, we can obtain suitable Laplace transforms for distributions such as the Weibull and Pareto distributions by approximating these distributions by distributions for which Laplace transforms are available. For the Weibull, Pareto and other decreasing-failure-rate (DFR) distributions, we can use hyperexponential distributions using the fitting algorithm in Feldmann and Whitt [16], which was shown there to yield remarkably good fits. For more general phase-type distributions, we can use the expectation-maximization (EM) algorithm, as in Asmussen, Nerman and Olsson [5] and

Turin [33]. Since long-tail distributions tend to have few finite moments, fitting by moment matching as in Johnson and Taaffe [24] tends to be ineffective in this context.

2. Transient Analysis for Control

The key idea for control is to observe the system state and exploit it to predict the future demand. In the context of our source model, the relevant system state is the level and elapsed level holding time in that level for each source. We assume that we do not observe and, thus do not use, the state of the within-level variation process. Since the within-level variation process is intended to capture "higher frequency" variations, it is natural to neglect it in predicting the future (assuming the future time of interest is suitably distant, roughly in the time scale of the level-process transitions). The elapsed level holding time is very important to predict the remaining holding time when the holding-time distribution is not nearly exponential. When the holding-time distribution is decreasing failure rate (DFR) with a long tail, then a large elapsed holding time means that a large remaining holding time is very likely; see Section 6 below and Section 8 of [12].

Conditional on the specified state information, we can compute the probability that each source will be in each possible level at any time in the future, from which we can calculate the conditional mean and variance of the aggregate required bandwidth by adding. The Lindeberg-Feller central limit theorem (CLT) for non-identically-distributed summands can be applied to generate a normal approximation; see p. 262 of Feller [17]. There are technical regularity conditions in the CLT, which we assume are satisfied. A sufficient condition is for all the summands to be uniformly bounded and for the sum of the variances to diverge as the number of sources increases; see Example (e) on p. 264 of [17]. In practice, this extra condition cannot be directly verified because it is on the asymptotic behavior as the number of sources increases, while we only have finitely many sources in an application. Practically, it suffices to check that the aggregate is not dominated by only a few sources.

To state the conditional CLT, let 0 represent current time. Let $B(t)$ and $I(t)$ denote the aggregate required bandwidth and state information at time t, respectively. Let $(B(t)/I(0))$ denote the random variable with the conditional distribution of $B(t)$ given the information $I(0)$, here regarded as known and deterministic. Let $N(m, \sigma^2)$ denote a normally distributed random variable with mean m and variance σ^2. Let $\Rightarrow$ denote convergence in distribution. The conditional CLT states that, for fixed $t > 0$,

$$\frac{(B(t)|I(0)) - E(B(t)|I(0))}{\sqrt{\mathrm{Var}(B(t)|I(0))}} \Rightarrow N(0,1) \tag{2.1}$$

as the number of sources gets large. As a corollary, we obtain the *(weak) law of large numbers* (LLN) stating that

$$\frac{(B(t)|I(0))}{E(B(t)|I(0))} \Rightarrow 1 \tag{2.2}$$

as the number of sources gets large. The LLN can also be obtained under more general conditions (without necessarily having the CLT); see Theorem 5.2.3 on p. 111 of Chung [9] for necessary and sufficient conditions.

Since the conditional mean alone tends to be very descriptive, we use the LLN approximation

$$(B(t)|I(0)) \approx E(B(t)|I(0)) \ . \tag{2.3}$$

We will show that the conditional mean in (2.3) can be efficiently computed, so that it can be used for real-time control. From (2.1), we see that the error in the approximation (2.3) is approximately characterized by the conditional standard deviation $\sqrt{\mathrm{Var}(B(t)|I(0))}$. We also will show how to compute this conditional standard deviation, although the required computation is more difficult. If there are n sources that are roughly homogeneous, than the conditional standard deviation will be $O(\sqrt{n})$, while the conditional mean is $O(n)$.

We anticipate that the conditional standard deviation $\sqrt{\mathrm{Var}(B(t)|I(0))}$ will increase with t because $\mathrm{Var}(L(0)|I(0)) = 0$, so that the accuracy of approximation (2.3) should decrease with t. For larger t, a convenient rough estimate of the conditional variance $\mathrm{Var}(B(t)|I(0))$ is the steady-state variance V in [13]. However, for the times relevant for control (involving relatively few transitions in each level process), we anticipate that the conditional standard deviation $\sqrt{\mathrm{Var}(B(t)|I(0))}$ will be substantially less than $\sqrt{V}$.

Given that our approximation is the conditional mean, and given that our state information does not include the state of the within-level variation process, the within-level variation process plays no role because it has zero mean. Let i index the source. Since the required bandwidths need not have integer values, we index the level by the integer j, $1 \le j \le J_i$, and indicate the associated required bandwidths in the level by b^i_j. Hence, instead of (1.1), the required bandwidth for source i can be expressed as

$$B^i(t) = b^i_{L^i(t)} + W_{L^i(t)}(t) \ , \quad t \ge 0 \ . \tag{2.4}$$

Let $P^{(i)}_{jk}(t|x)$ be the probability that the source-i level process is in level k at time t given that at time 0 it was in level j and had been so for a period x (i.e., the age or elapsed level holding time at time 0 is x). If $\mathbf{j} \equiv (j_1, \ldots, j_n)$ and $\mathbf{x} \equiv (x_1, \ldots, x_n)$ are the vectors of levels and ages of the n source level processes at time 0, then the *state information* is $I(0) = (\mathbf{j}, \mathbf{x}) = (j_1, \ldots, j_n; x_1, \ldots, x_n)$ and the *conditional aggregate mean* is

$$E(B(t)|I(0)) \equiv M(t|\mathbf{j}, \mathbf{x}) = \sum_{i=1}^{n} \sum_{k_i=1}^{J_i} P^{(i)}_{j_i k_i}(t|x_i) b^i_{k_i} \ . \tag{2.5}$$

From (2.5), we see that we need to compute the conditional distribution of the level, i.e., the probabilities $P^{(i)}_{jk}(t|x)$, for each source i. In this section we show how to compute these conditional probabilities. By similar reasoning, it is possible to describe the transition probabilities of the entire Markov process (elapsed holding time plus level). To proceed, we consider a single source and assume that its required bandwidth process is a semi-Markov process (SMP). (We now have no within-level variation process.) We now omit the superscript i. Let $L(t)$ and $B(t)$ be the level and required bandwidth, respectively, at time t as in (2.4). The process $\{L(t) : t \geq 0\}$ is assumed to be an SMP, while the process $\{B(t) : t \geq 0\}$ is a function of an SMP, i.e., $B(t) = b_{L(t)}$, where b_j is the required bandwidth in level j. If $b_j \neq b_k$ for $j \neq k$, then $\{B(t) : t \geq 0\}$ itself is an SMP, but if $b_j = b_k$ for some $j \neq k$, then typically $\{B(t) : t \geq 0\}$ is not an SMP.

Let $A(t)$ be the age of the level holding time at time t. We are interested in calculating the transition probabilities

$$P_{jk}(t|x) \equiv P(L(t) = k|L(0) = j, A(0) = x) \tag{2.6}$$

as a function of j, k, x, and t. The state information at time 0 is the pair (j, x). The transition probabilities in (2.6) can be calculated by standard conditioning arguments. In analogy to delayed renewal processes, we are treating a delayed SMP. The transition probabilities satisfy a Markov renewal equation that can be solved explicitly in the transform domain. To express the result, let P be the transition matrix of the DTMC governing level transitions and let $F_{jk}(t)$ be the holding-time cdf given that there is a transition from level j to level k. For simplicity, we assume that $F^c_{jk}(t) > 0$ for all j, k, and t, so that all positive x can be level holding times. Let $P(t|x)$ be the matrix with elements $P_{jk}(t|x)$ and let $\hat{P}(s|x)$ be the Laplace transform (LT) of $P(t|x)$, i.e., the matrix with elements that are the Laplace transforms of $P_{jk}(t|x)$ with respect to time, i.e.,

$$\hat{P}_{jk}(s|x) = \int_0^\infty e^{-st} P_{jk}(t|x) dt. \tag{2.7}$$

We will derive an expression for $\hat{P}(s|x)$. For this purpose, let G_j be the holding-time cdf in level j, unconditional on the next level, i.e.,

$$G_j(x) = \sum_k P_{jk} F_{jk}(x) \ . \tag{2.8}$$

For any cdf G, let G^c be the complementary cdf, i.e. $G^c(x) = 1 - G(x)$. Also let

$$H_{jk}(t|x) = \frac{P_{jk} F_{jk}(t+x)}{G_j^c(x)} \quad \text{and} \quad G_j(t|x) = \sum_k H_{jk}(t|x) \tag{2.9}$$

for G_j in (2.8). Then let $\hat{h}_{jk}(s|x)$ and $\hat{g}_j(s|x)$ be the associated Laplace-Stieltjes transforms (LSTs), i.e.,

$$\hat{h}_{jk}(s|x) = \int_0^\infty e^{-st} dH_{jk}(t|x) \quad \text{and} \quad \hat{g}_j(s|x) = \int_0^\infty e^{-st} dG_j(t|x) \ . \tag{2.10}$$

Let $\hat{h}(s|x)$ be the matrix with elements $\hat{h}_{jk}(s|x)$. Let $\hat{q}(s)$ be the matrix with elements $\hat{Q}_{jk}(s)$, where

$$Q_{jk}(t) = P_{jk} F_{jk}(t) \quad \text{and} \quad \hat{q}_{jk}(s) = \int_0^\infty e^{-st} dQ_{jk}(t) \ . \tag{2.11}$$

Let $\hat{D}(s|x)$ be the diagonal matrix with diagonal elements

$$\hat{D}_{jj}(x|x) \equiv [1 - \hat{g}_j(s|x)]/s \tag{2.12}$$

Let $\hat{D}(s)$ be the diagonal matrix with diagonal elements

$$\hat{D}_{jj}(s) \equiv [1 - \hat{g}_j(s)]/s \ , \tag{2.13}$$

where $\hat{g}_j(s)$ is the LST of the cdf G_j in (2.8).

Theorem 2.1 *The transient probabilities for a single SMP source have the matrix of Laplace transforms*

$$\hat{P}(s|x) = \hat{D}(s|x) + \hat{h}(s|x)\hat{P}(s|0) \ , \tag{2.14}$$

where

$$\hat{P}(s|0) = (I - \hat{q}(s))^{-1} \hat{D}(s). \tag{2.15}$$

Proof. In the time domain, condition on the first transition. For $j \neq k$,

$$P_{jk}(t|x) = \sum_l \int_0^t dH_{jl}(u|x) P_{lk}(t-u|0) \ ,$$

so that

$$\hat{P}_{jk}(s|x) = \sum_l \hat{h}_{jl}(s|x)\hat{P}_{lk}(s|0),$$

while

$$P_{jj}(t|x) = G_j^c(t|x) + \sum_l \int_0^t dH_{jl}(u|x) P_{lj}(t-u|0) \ ,$$

so that

$$\hat{P}_{jj}(s|x) = \frac{1-\hat{g}_j(s|x)}{s} + \sum_l h_{jl}(s|x)\hat{P}_{lj}(s|0).$$

Hence, (2.14) holds. However, $P(t|0)$ satisfies a Markov renewal equation, as in Section 10.3 of Çinlar [10], i.e., for $j \neq k$,

$$P_{jk}(t|0) = \sum_l \int_0^t dQ_{jl}(u) P_{lk}(t-u|0)$$

and

$$P_{jj}(t|0) = G_j^c(t) + \sum_l \int_0^\infty dQ_{jl}(u) P_{lj}(t-u|0) \ ,$$

so that

$$P(t|0) = D(t) + Q(t) * P(t|0)$$

where $*$ denotes convolution, and (2.15) holds. ∎

Remark 2.1. Note that as $x \to 0$, $\hat{D}(s|x) \to \hat{D}(s)$ and $\hat{h}(s|x) \to \hat{q}(s)$, so that $\hat{P}(s|x) \to \hat{P}(s|0)$ in (3.9) and (3.10), because

$$\begin{aligned}\hat{P}(s|x) &\to \hat{D}(s) + \hat{q}(s)(1-q(s))^{-1}\hat{D}(s)\\ &= \hat{D}(s) + \hat{q}(s)\sum_{n=0}^{\infty} q(s)^n \hat{D}(s)\\ &= (I-\hat{q}(s))^{-1}\hat{D}(s).\end{aligned}$$

To compute the LT $\hat{P}(s|0)$, we only need the LSTs $\hat{f}_{jk}(s)$ and $\hat{g}_j(s)$ associated with the basic holding-time cdf's F_{jk} and G_j. However, to compute $\hat{P}(s|x)$, we also need to compute $\hat{D}(s|x)$ and $\hat{h}(s|x)$, which require computing the LSTs of the *conditional* cdf's $H_{jk}(t|x)$ and $G_j(t|x)$ in (2.9). We will show how to compute these conditional LSTs for a large class of holding-time distributions in Section 5.

If the number of levels is not too large, then it will not be difficult to compute the required matrix inverse $(I-q(s))^{-1}$ for all required s. Note that, because of the probability structure, the inverse is well defined for all complex s with $\mathrm{Re}(s) > 0$. To illustrate with an important simple example, we next give the explicit formula for an on-off source.

Example 2.1. Suppose that we have an on-off source, i.e., so that there are two states with transition probabilities $P_{12} = P_{21} = 1$ and holding time cdf's G_1 and G_2. From (2.9) or by direct calculation,

$$\begin{aligned} \hat{P}(s|0) \equiv \begin{pmatrix} \hat{P}_{11}(s|0) & \hat{P}_{12}(s|0) \\ \hat{P}_{21}(s|0) & \hat{P}_{22}(s|0) \end{pmatrix} &= (I-\hat{q}(s))^{-1}\hat{D}(s) \\ &= \frac{1}{s(1-\hat{g}_1(s)\hat{g}_2(s))} \begin{pmatrix} 1-\hat{g}_1(s) & \hat{g}_1(s)(1-\hat{g}_2(s)) \\ \hat{g}_2(s)(1-\hat{g}_1(s)) & 1-\hat{g}_2(s) \end{pmatrix} . \end{aligned} \tag{2.16}$$

Suppose that the levels are labeled so that the initial level is 1. Note that all transitions from level 1 are to level 2. Hence when considering the matrix $\hat{h}(s|x)$ in (2.10) it suffices to consider only the element $\hat{h}_{12}(s|x)$. Since

$$H_{12}^c(t|x) = G_1^c(t|x) = \frac{G_1^c(t+x)}{G_1^c(x)} , \tag{2.17}$$

$$\hat{h}_{12}(s|x) = \hat{g}_1(s|x) = \int_0^\infty e^{-st} dG_1(t|x) . \tag{2.18}$$

Since $P_{11}(t|x) = 1 - P_{12}(t|x)$, it suffices to calculate only $P_{12}(t|x)$. Hence, in this context

$$\hat{P}_{12}(s|x) = \frac{\hat{g}_1(s|x)(1-\hat{g}_2(s))}{s(1-\hat{g}_1(s)\hat{g}_2(s))} . \tag{2.19}$$

We now determine the mean, second moment, and variance of the bandwidth process of a general multi-level source as a function of time. It is elementary that

$$m_j(t|x) = E(B(t)|L(0)=j, A(0)=x) = \sum_k P_{jk}(t|x) b_k \tag{2.20}$$

$$s_j(t|x) = E(B(t)^2|L(0)=j, A(0)=x) = \sum_k P_{jk}(t|x) b_k^2 \tag{2.21}$$

$$v_j(t|x) = \mathrm{Var}(B(t)|L(0)=j, A(0)=x) = s_j(t|x) - m_j(t|x)^2 . \tag{2.22}$$

We can calculate $m_j(t|x)$ and $s_j(t|x)$ by a single inversion of their Laplace transforms, using

$$\hat{m}_j(s|x) \equiv \int_0^\infty e^{-st} m_j(t|x)dt = \sum_k P_{jk}(s|x)b_k \tag{2.23}$$

and

$$\hat{s}_j(s|x) = \sum_k \hat{P}_{jk}(s|x)b_k^2. \tag{2.24}$$

To properly account for the within-level variation process when it is present, we should add its variance in level j, say $w_j(t,x)$, to $v_j(t,x)$, but we need make no change to the mean $M_j(t,x)$. We anticipate that $w_j(t,x)$ will be much less than $v_j(t,x)$, so that $w_j(t,x)$ can be omitted, but it could be included.

Finally, we consider the aggregate bandwidth associated with n sources. Again let a superscript i index the sources. The conditional aggregate mean and variance are

$$M(t|\mathbf{j},\mathbf{x}) \equiv E(B(t)|I(0)) = \sum_{i=1}^n m_{j_i}^i(t|x_i) \tag{2.25}$$

and

$$V(t|\mathbf{j},\mathbf{x}) \equiv \mathrm{Var}(B(t)|I(0)) = \sum_{i=1}^n [v_{j_i}^i(t|x_i) + w_{j_i}^i(t,x_i)] \ , \tag{2.26}$$

where $\mathbf{j} = (j_1,\ldots,j_n)$ is the vector of levels and $\mathbf{x} = (x_1,\ldots,x_n)$ is the vector of elapsed holding times for the n sources with the single-source means and variances as in (2.20) and (2.22).

It is significant that we can calculate the conditional aggregate mean at any time t by performing a single inversion. We summarize this elementary but important consequence as a theorem.

Theorem 2.2 *The Laplace transform of the n-source conditional mean aggregate required bandwidth as a function of time is*

$$\hat{M}(s|\mathbf{j},\mathbf{x}) \equiv \int_0^\infty e^{-st} M(t|\mathbf{j},\mathbf{x})dt = \sum_{i=1}^n \sum_{k_i=1}^{J_i} \hat{P}_{j_i k_i}^{(i)}(s|x_i) b_{k_i} \ , \tag{2.27}$$

where the single-source transform $\hat{P}_{j_i k_i}^{(i)}(s|x_i)$ is given in Theorem 2.1.

Unlike the aggregate mean, for the aggregate variance we evidently need to perform n separate inversions to calculate $v_{j_i}^i(t|x_i)$ for each i and then add to calculate $V(t|\mathbf{j},\mathbf{x})$ in (2.26). Hence, we suggest calculating only the conditional mean on line for control, and occasionally calculating the conditional variance off line to evaluate the accuracy of the conditional mean.

3. The One-Transition and Two-Transition Approximations

The most complicated part of the conditional aggregate mean transform $\hat{M}(s|\mathbf{j},\mathbf{x})$ in (2.27) is the matrix inverse $(I-\hat{q}(s))^{-1}$ in the transform of the single-source transition probability in (2.15). Since the matrix inverse calculation can be a computational burden when the number of levels is large, it is natural to seek approximations which avoid this matrix inverse. We describe such approximations in this section.

The matrix inverse $(I-q(s))^{-1}$ is a compact representation for the series $\sum_{n=0}^{\infty} q(s)^n$. For $P(t|x)$, it captures the possibility of any number of transitions up to time t. However, if the levels are relatively long in the time scale relevant for control, then the mean for times t of interest will only be significantly affected by a very few transitions. Indeed, often only a single transition need be considered, and that is the main approximation we propose here.

The single-transition approximation is obtained by making the Markov chain absorbing after one transition. Hence, the single-transition approximation is simply

$$P_{jk}(t|x) \approx H_{jk}(t|x) \ , \quad j \neq k \ , \tag{3.1}$$

and

$$P_{jj}(t|x) \approx G_j^c(t|x) + H_{jj}(t|x) \tag{3.2}$$

for $H_{jk}(t|x)$ in (2.9) and $G_j(t|x)$ in (2.9). From (3.1) and (3.2) we see that no inversion is needed.

Alternatively, we can develop a two-transition approximation. (Extensions to higher numbers are straightforward.) Modifying the proof of Theorem 2.1 is a straightforward manner, we obtain

$$P_{jk}(t|x) = \int_0^t G_k^c(t-u) dH_{jk}(u|x) + \sum_{\ell} \int_0^t P_{\ell k} F_{lk}(t-u) dH_{j\ell}(u|x) \tag{3.3}$$

for $j \neq k$ and

$$P_{jj}(t|x) = G_j^c(t|x) + \sum_{\ell} \int_0^t P_{\ell j} F_{lj}(t-u) dH_{j\ell}(u|x) \ . \tag{3.4}$$

Expressed in the form of transforms, (3.3) and (3.4) become

$$\hat{P}_{jk}(s|x) = \hat{h}_{jk}(s|x) \frac{(1-\hat{g}_k(s))}{s} + \sum_{\ell} \hat{h}_{j\ell}(s|x) P_{\ell k} \frac{\hat{f}_{lk}(s)}{s} \tag{3.5}$$

for $j \neq k$ and

$$\hat{P}_{jj}(s|x) = \frac{1 - \hat{g}_j(s|x)}{s} + \sum_{\ell} \hat{h}_{j\ell}(s|x) P_{lj} \frac{\hat{f}_{lj}(s)}{s} \ . \tag{3.6}$$

Numerical inversion can easily be applied with (3.5) and (3.6). However, since the time-domain formulas (3.3) and (3.4) involve single convolution integrals, numerical computation of (3.3) and (3.4) in the time domain is also a feasible alternative. Moreover, if the underlying distributions have special structure, then the integrals in (3.3) and (3.4) can be calculated analytically. For example, analytical integration can easily be done when all holding-time distributions are hyperexponential.

Example 3.1. To illustrate how the two approximations compare to the exact conditional mean, we give a numerical example. We consider a single source with four levels. The transitions move cyclically through the levels: $P_{12} = P_{23} = P_{34} = P_{41} = 1$. The level holding-time ccdf's are:

$$\begin{aligned} G_1^c(t) &= 0.5e^{-10t} + 0.5e^{-0.1t} \ , \quad G_2^c(t) = e^{-0.1t} \ , \\ G_3^c(t) &= 0.9e^{-2t} + 0.1e^{-0.1t} \ , \quad G_4^c(t) = e^{-0.1t} \ , \ t \geq 0 \ . \end{aligned}$$

The level bandwidths are $b_1 = b_3 = 100$ and $b_2 = b_4 = 0$. Suppose that we start in level 1 with an age of 8. From the form of $G_1^c(t)$, we see that the conditional level-1 holding-time ccdf $G_1^c(t|x)$ is then approximately $e^{-0.1t}$. Hence the first two mean level holding times are approximately 10. Hence we might consider the one-transition and two-transition approximations in the interval $[0, 10]$. The two approximations are compared to the exact value of the conditional mean in Figure 1. (All are computed by numerical transform inversion.) The approximations are very good up to $t = 1$ or 2, but they start to degrade by $t = 10$. The two-transition approximation performs not so well for larger t because the actual holding time in level 3 is likely to be quite short. More generally, our experience is that the one-transition and two-transition approximations tend to perform quite satisfactorily if the mean level holding times in the first few levels are substantially *larger* than the times t of interest. In this example the approximations are quite good in the interval $[0, 1]$.

Figure 1: A comparison of the one-transition and two-transition approximations with the exact conditional mean aggregate demand in Example 3.1.

4. Computing Laplace-Stieltjes Transforms of Level Holding-Time Distributions

In order to apply numerical transform inversion to compute the transforms $\hat{P}(s|x)$ in Theorem 2.1 and $\hat{M}(s|\mathbf{j},\mathbf{x})$ in Theorem 2.2, or the two-transition approximation in (3.5) and (3.6), we need to be able to calculate all component transforms. This means that we need to be able to calculate the LSTs $\hat{g}_j(s)$, $\hat{g}_j(s|x)$ and $\hat{h}_{jk}(s|x)$.

Since some distributions do not have convenient LST's, we suggest approximating such distributions by distributions with convenient LSTs. For this purpose, Feldmann and Whitt [16] present an algorithm for approximating decreasing failure rate (DFR) distributions, including long-tail distributions such as Weibull and Pareto distributions, by hyperexponential (H_k) distributions. More generally, the EM algorithm can be used to approximate arbitrary distributions by a phase-type (PH) distribution, as in Asmussen, Nerman and Olsson [5]. Since phase-type distributions are dense in the space of all probability distributions, this should offer enough candidates. Previous algorithms for fitting phase-type distributions by moment matching are described in Johnson and Taaffe [24] and references therein, but these algorithms tend to be ineffective for long-tail distributions with few (e.g., 0 or 1) finite moments.

An H_k complementary cdf (ccdf) has the form

$$G^c(t) = \sum_{i=1}^{k} p_i e^{-\lambda_i t}, t \geq 0. \tag{4.1}$$

The associated H_k cdf G has LST

$$\hat{g}(s) \equiv \int_0^\infty e^{-st} dG(t) = \sum_{i=1}^{k} p_i \lambda_i / (\lambda_i + s) \ . \tag{4.2}$$

Phase-type distributions are discussed in Chapter 2 of Neuts [28]. In the notation there, a phase-type ($PH \equiv PH(m)$) distribution is characterized by an m-dimensional vector α and an $m \times m$ matrix T. With that notation, the ccdf is

$$G^c(t) = \alpha e^{Tt} u \ , \quad t \geq 0, \tag{4.3}$$

where $u = (1, 1, \ldots, 1)^t$. The LST of the associated cdf G is

$$\hat{g}(s) = \alpha_{m+1} + \alpha(sI - T)^{-1}(-Tu) \tag{4.4}$$

where $\alpha_{m+1} = 1 - \alpha u$.

If we require that α be a probability vector, then the PH distribution does not have the atom α_{m+1} at the origin. Then the PH distribution has a probability density function (pdf) and thus will tend to cause no problems in the numerical inversion algorithms (see [1]). An H_k distribution always has a pdf (provided that we exclude $\lambda_i = \infty$ from (4.1)).

It is significant that, when a holding-time cdf G is H_k or PH, the conditional remaining holding-time cdf has the same form. Similar results hold for the Pareto and uniform distributions. We state the result as a theorem, but omit the elementary proof.

Theorem 4.1 *(a) If G is H_k as in (4.1), then $G(\cdot|x)$ is H_k with*

$$G^c(t|x) = \sum_{i=1}^{k} p_i(x) e^{-\lambda_i t} \ , \quad t \geq 0 \ , \tag{4.5}$$

where

$$p_i(x) = \frac{p_i e^{-\lambda_i x}}{\sum_{j=1}^{k} p_j e^{-\lambda_j x}} \ , x \geq 0 \ . \tag{4.6}$$

(b) If G is PH as in (4.3), then $G(\cdot|x)$ is PH with

$$G^c(t|x) = \alpha(x) e^{Tt} u, \ t \geq 0, \tag{4.7}$$

where

$$\alpha(x) = \frac{\alpha e^{Tx}}{\alpha e^{Tx} u} \ , \quad x \geq 0 \ . \tag{4.8}$$

(c) If G is Pareto of the form

$$G^c(t) = (1 + bt)^{-a} \ , t \geq 0 \ , \tag{4.9}$$

for $a > 0$, then $G(\cdot|x)$ is Pareto with

$$G^c(t|x) = (1 + b(x)t)^{-a} \ , t \geq 0, \tag{4.10}$$

where

$$b(x) = \frac{b}{1+bx} \ . \tag{4.11}$$

(*d*) *If G is uniform on* $[a,b]$ *with ccdf*

$$G^c(t) = \frac{b-t}{b-a} \ , \quad a \le t \le b \ , \tag{4.12}$$

then $G^(\cdot |x)$ *is uniform on* $[x,b]$ *for* $a < x < b$*, i.e.,*

$$G^c(t|x) = \frac{b-t}{b-x} \ , \quad x \le t \le b \ . \tag{4.13}$$

By parts (a) and (c), we could approximate a Pareto by an H_k and condition in either order. Since conditioning with respect to the age x with a Pareto (a,b) distribution corresponds simply to a rescaling of time, this conditioning applied to the approximating H_k ccdf in (4.5) would lead to

$$G^c(t|x) = \sum_{i=1}^{k} p_i e^{-\lambda_i(x)t} \ , \quad t \ge 0 \ , \tag{4.14}$$

where

$$\lambda_i(x) = \lambda_i/(1+bx) \ , \quad x \ge 0 \ ; \tag{4.15}$$

i.e., we would keep the H_k form but adjust the exponential rates λ_i instead of the probability weights p_i as in (4.5) and (4.6). For greater accuracy in the approximation, it is usually better to condition first and then approximate. However, when x is not too large and enough exponential components are used, the order does not matter. This was confirmed using the H_{13} fit to the Pareto (2,2, 0.83) in Section 6 of Feldmann and Whitt [16].

5. The Value of Information

We can use the source model to investigate the value of information. We can consider how prediction is improved when we condition on, first, only the level and, second, on both level and age. The reference case is the steady-state mean

$$M = \sum_{i=1}^{n} m^i \quad \text{and } m^i = \sum b_j^i p_j^i \ , \tag{5.1}$$

where p_j^i is the steady-state probability, i.e., omitting the superscript i,

$$p_j = \frac{\pi_j^i m(G_j^i)}{\sum_k \pi_k^i m(G_k^i)} \ , \tag{5.2}$$

with π the steady-state vector of the Markov chain P ($\pi = \pi P$) and $m(G_j)$ the mean of G_j for G_j in (2.8). With the steady-state mean, there is no conditioning. Section 2 gives the formula for conditioning on both level and age. Now we give the formulas conditioning only on the level; i.e., we condition on the level, assuming that we are in steady-state. Again we omit the superscript i. Then the age in level j has the stationary-excess cdf

$$G_{je}(t) = \frac{1}{m(G_j)} \int_0^t G_j^c(u)du \ , \quad t \geq 0 \ . \tag{5.3}$$

Let $P_{jk}(t)$ be the probability of being in level k at time t conditional on being in level j in steady state at time 0. Let $\hat{P}_{jk}(s)$ be its Laplace transform. Let $m_j(t)$ be the conditional steady-state mean given level j at time 0 and let $\hat{m}_j(s)$ be its Laplace transform: Clearly

$$m_j(t) = \sum_{k=1}^{J} P_{jk}(t) b_k \quad \text{and} \quad \hat{m}_j(s) = \textstyle\sum_{k=1}^{J} \hat{P}_{jk}(s) b_k \ . \tag{5.4}$$

Hence, it suffices to calculate $\hat{P}_{jk}(s)$.

Theorem 5.1 *Assume that the level-holding-time cdf depends only on the originating level, i.e., $F_{jk}(t) = G_j(t)$. The steady-state transition probabilities conditional on the level for a single SMP source have the matrix of Laplace transforms*

$$\hat{P}(s) = \hat{D}_e(s) + \hat{g}_e(s)\hat{P}(s|0) \ , \tag{5.5}$$

where $\hat{P}(s|0)$ is the matrix in (2.15), $\hat{g}_e(s)$ is the matrix with elements

$$\hat{g}_{ejk}(s) = P_{jk}\hat{g}_{je}(s) = P_{jk}\frac{(1-\hat{g}_j(s))}{sm(G_j)} \ , \tag{5.6}$$

$\hat{D}_e(s)$ is the diagonal matrix with diagonal elements

$$\hat{D}_{ejj}(s) \equiv \frac{1-\hat{g}_{je}(s)}{s} = \frac{sm(G_j) - 1 + \hat{g}_j(s)}{s^2 m(G_j)} \ , \tag{5.7}$$

$\hat{g}_j(s)$ is the level-j holding-time LST and $\hat{g}_{je}(s)$ is the LST of its stationary-excess cdf in (5.3).

Proof. Modify the proof of Theorem 2.1, inserting $P_{jl}G_{je}(t)$ for $H_{jl}(t|x)$ and $G_{je}^c(t)$ for $G_j^c(t|x)$. ∎

Example 5.1. Consider the on-off source in Example 2.1. Paralleling (2.19), it suffices to calculate only $P_{12}(t)$. Its Laplace transform is

$$\hat{P}_{12}(s) = \frac{\hat{g}_{1e}(s)(1-\hat{g}_2(s))}{s(1-\hat{g}_1(s)\hat{g}_2(s))} \ . \tag{5.8}$$

For a given time of interest, the importance of knowing the level increases as the level holding times increase. For large holding times, steady state will be approached relatively slowly, so that transient descriptions are more useful. The importance of knowing the holding-time age as well as the level is greater when the holding-time distribution is far from exponential. By the lack of memory property of the exponential distribution, the age provides no information when the holding-time distribution is exponential. Given that we know the level but not the age, we act as if the age has the holding-time stationary-excess distribution. Then the remaining holding-time cdf is the average of the conditional holding-time cdf given an age x with respect to the stationary-excess distribution:

$$\int_0^\infty \left(\frac{G_j(t+x)}{G_j^c(x)}\right) dG_{je}(x) = \int_0^\infty \frac{G_j(t+x)dx}{m(G_j)} = G_{je}(t) \ . \tag{5.9}$$

Hence, knowing the age is more important when the conditional cdf $G(t|x)$ varies significantly with x. If we know only the level, then we get an appropriate average.

Example 5.2. To show the value of knowing the age, consider an on-off source with holding-time ccdf's

$$\begin{aligned} G_1^c(t) &= 0.01e^{-0.01t} + 0.1e^{-0.1t} + .89e^{-t} \\ G_2^c(t) &= e^{-t} \ , \quad t \geq 0 \ . \end{aligned}$$

Let the bandwidths be $b_1 = 100$ and $b_2 = 0$. Since $m(G_1) = 2.89$ and $m(G_2) = 1.00$, the steady-state mean is

$$EB(\infty) = \frac{100m(G_1)}{m(G_1)+m(G_2)} = 74.29 \ .$$

Let the initial level be 1. Since G_1 has an exponential component with mean 100, we anticipate the time to reach steady state to be between 100 and 1000. In Figure 2 we plot the conditional mean $m_1(t|x)$ for $x = 0.5$, 5.0 and 50.0, computed by numerical transform inversion. Figure 2 shows that the age plays a very important role.

Figure 2: The conditional mean aggregate demand as a function of the age of the holding time in level 1 for Example 5.2

6. Control

The source model makes it possible to investigate several different kinds of controls. As in Sections 4 and 9 of [12], we can consider admitting new sources or removing existing sources. We can also consider changing the levels of existing sources. For example, on-off sources in the on state might be turned off. We also can consider rate controls corresponding to changing the bandwidths assigned to the levels. With each of these controls, we can calculate the resulting conditional mean aggregate bandwidth to evaluate the performance of the control.

To illustrate, suppose that we wish to consider which of n candidate sources to serve over the time interval $[0,T]$. Suppose that source i earns a fixed revenue R_i plus a revenue rate r_i per unit of bandwidth per time. Also suppose that we want to keep the demand (total input rate) below a capacity c at all times. Then we can solve the following integer program. Let y_i be the decision variable, with $y_i = 1$ if source i is served and $y_i = 0$ if not. The integer program can be formulated as:

$$\max \sum_{i=1}^{n} y_i (R_i + r_i \int_0^T m_{j_i}^i (t|x_i) dt) \tag{6.1}$$

subject to

$$\sum_{i=1}^{n} y_i m_{j_i}^i (t_k|x_i) \le c \ , \quad 0 \le k \le K \ , \tag{6.2}$$

for a set of time points $0 = t_0 < t_1 < \ldots < t_K = T$. In (6.1) and (6.2) the level j_i and age x_i of source i are included because these are presumed to be known.

Our approach here also provides an approach to admission control when some sources book ahead; see Greenberg, Srikant and Whitt [19], Greenberg and Wischik [20] and references therein. If we stipulate that book-ahead (BA) customers specify their bandwidth requirements, then future capacity can be adjusted for BA calls. In particular, the future capacity can be regarded as a function $C(t)$, $t \ge 0$. We can then use the conditional mean to describe the demand associated with previously admitted instantaneous-request (IR) sources. We can then admit a new IR source if, after it has been admitted, the conditional mean remains below $C(t)$ for all t. Upon admission, each source is represented by our source model with appropriate parameters. Unlike [19], but like [20], the approach here allows us to treat heterogeneous IR sources.

7. Conclusions

We proposed a general source traffic model composed of a semi-Markov level process and a zero-mean piecewise-stationary within-level variation process. We approximated the aggregate demand from many sources by the conditional aggregate mean given level values and ages. We justified this deterministic fluid approximation by applying the law of large numbers. It is significant that the within-level variation process plays no role in this approximation. We showed that the conditional mean can be effectively computed using numerical transform inversion (Sections 2 and 4). When level holding times are relatively long compared to the times of interest, we can further approximate the conditional mean by looking at only one or two transitions, which produces even more elementary approximations (Section 3). We showed that the conditional standard deviation can be computed to evaluate the accuracy of the conditional-mean approximation. We showed how the model can be exploited to study the value of information (Section 5), and control (Section 6).

Even though our approach is to focus on offered load, unaltered by loss and delay associated with finite capacity, we can apply the conditional mean approximation in Section 2 to develop an approximation to describe loss and delay from a finite-capacity system, just as described in Section 5 of [12]. The idea is to approximate the bandwidth stochastic process by the deterministic conditional-mean process $E(B(t)|I(0))$, $t \geq 0$, and then describe the loss and delay associated with exceeding the capacity C. If there is no buffer, then the amount of loss associated with the rare event of hitting a level above C would be estimated by the integral of $E(B(t)|I(0)) - C$ over the region before and after $t = 0$ over which it is positive.

In the introduction we briefly discussed ways in which the source traffic model might be fit to data. Experiments with fitting are natural next steps. We believe that the source traffic model has the potential for realistically representing traffic, but clearly care is needed in the fitting. On the other hand, we hope that the performance predictions will be reasonably robust to the fitting.

References

[1] Abate, J. and Whitt, W. (1992) The Fourier-series method for inverting transforms of probability distributions, *Queueing Systems*, 10, 5–88.

[2] Abate, J. and Whitt, W. (1995) Numerical inversion of Laplace transforms of probability distributions, *OA, RS J. on Computing*, 7, 36–43.

[3] Abate, J., Choudhury, G. L. and Whitt, W. (1996) On the Laguerre method for numerically inverting Laplace transforms, *S, IN, FM, OR J. on Computing*, 8, 413-427.

[4] Anick, D., Mitra, D. and Sondhi, M. M. (1982), Stochastic theory of a data-handling system with multiple sources, *Bell System Tech. J.*, 61, 1871–1894.

[5] Asmussen, S., Nerman, O. and Olsson, M. (1996) Fitting phase-type distributions via the EM algorithm, *Scand. J. Statist.*, 23, 419–414.

[6] R. Cáceres, Danzig, P. G., Jamin, S. and Mitzel, D. J. (1991) Characteristics of wide-area TCP/IP conversations, *Computer Communications Review*, 21, 101–112.

[7] Chen, H. and Mandelbaum, A. (1991) Discrete flow networks: bottleneck analysis and fluid approximations, *Math. Oper. Res.*, 16, 408–446.

[8] Choudhury, G. L., Lucantoni, D. M. and Whitt, W. (1996) Squeezing the most out of M, AT, *IE, EE Trans. Commun.*, 44, 203–217.

[9] Chung, K. L., *Course, A in Probability Theory*, second edition, Academic Press, New York, 1974.

[10] E. Çinlar, *Introduction to Stochastic Processes*, Prentice-Hall, Englewood Cliffs, NJ, 1975.

[11] Crovella, M. E. and A. Bestavros Self-similarity in World Wide Web traffic – evidence and possible causes, *Proc. Sigmetrics, M, AC '96*, pp. 160–169, 1996.

[12] Duffield, N. G. and Whitt, W., Control and recovery from rare congestion events in a large multi-server system, *Queueing Systems*, 1997, to appear.

[13] Duffield, N. G. and Whitt, W. (1997) A source traffic model and its use in capacity planning, AT&T Labs, Florham Park, NJ

[14] Elwalid, A. I. and Mitra, D., Markovian arrival and service communication systems: spectral expansions, separability and Kronecker-product forms, in *Computation with Markov Chains*, Stewart, W. J. (ed.), Kluwer, Boston, pp. 507–546, 1995.

[15] Feldmann, A., Modeling characteristics of P, TC connections, AT&Laboratories, T , 1996.

[16] Feldmann, A. and Whitt, W. (1997) Fitting mixtures of exponentials to long-tail distributions to analyze network performance models, *Proc. Infocom, IE, EE '97*, 1098–1106.

[17] Feller, W., *Introduction, An to Probability Theory and its Applications*, vol. II, sec. ed., Wiley, New York, 1971.

[18] Gnedenko, B. V. and Kolmogorov, A. N., *Limit Distributions for Sums of Independent Random Variables* (translated from Russian), Addison-Wesley, Reading, MA, 1954.

[19] Greenberg, A. G., Srikant, R. and Whitt, W., Resource sharing for book-ahead and instantaneous-request calls. Submitted.

[20] Greenberg, A. G. and Wischik, D., Admission control for booking ahead shared resources. Submitted.

[21] Grossglauser, M., Keshav, S., and Tse, D., RA, CBR: simple and efficient service for multiple time-scale traffic, in *Proc. M, AC SIGCOMM*, pp. 219-230, 1995.

[22] Heyman, D. P. and Sobel, M. J. (1982), *Stochastic Models in Operations Research*, 1, McGraw-Hill, New York.

[23] Jennings, O. B., Mandelbaum, A., Massey, W. A. and Whitt, W. (1996) Server staffing to meet time-varying demand, *Management Science*, 42, 1383-1394.

[24] Johnson, M. A. and Taaffe, M. R. (1991) An investigation of phase-type distribution moment matching algorithms for use in queueing modesl, *Queueing Systems*, 8, 129–148.

[25] Leland, W. E., Taqqu, M. S., Willinger, W. and Wilson, D. V. (1994) On the self-similar nature of Ethernet traffic, *IEEE/ACM Transactions on Networking*, 2, 1–15.

[26] Leung, K. K., Massey, W. A. and Whitt, W. (1994) Traffic models for wireless communication networks, *IE, EE J. Sel. Areas Commun.*, 12, 1353–1364.

[27] Lucantoni, D. M., Choudhury, G. L. and Whitt, W. (1994) The transient BMAP/G/1 queue, *Stochastic Models*, 10, 145–182.

[28] Neuts, M. F., *Matrix-Geometric Solutions in Stochastic Models*, The Johns Hopkins University Press, Baltimore, 1981.

[29] Neuts, M. F., *Structured Stochastic Matrices of M/G/1 Type and Their Applications*, Marcel Dekker, New York, 1989.

[30] Paxson, V. and Floyd, S. (1995) Wide-area traffic: the failure of Poisson modeling, *IM, EE, E/Transactions, AC on Networking*, 3, 226–244.

[31] Roberts, J. W., *Performance Evaluation and Design of Multiservice Networks*, Cost 224 Project Final Report, Brussels, 1992.

[32] Salehi, J., Zhang, Z., Kurose, J., and Towsley, D., Supporting stored video: Reducing rate variability and end-to-end resource requirements through optimal smoothing, in *Proc. M, AC Sigmetrics*, pp. 222-231, 1996.

[33] Turin, W. (1996) Fitting stochastic automata via the EM algorithm, *Stochastic Models*, 12, 405–424.

[34] Willinger, W., Taqqu, M. S., Sherman, R. and Wilson, D. V. (1997) Self-similarity through high-variability: statistical analysis of Ethernet N, LA traffic at the source level, *IM, EE, E/Transactions, AC on Networking*, 5, 71–86.