Stochastic Fluid Flow Models for Determining Optimal Switching Thresholds

V. Aggarwal, N. Gautam*, S.R.T. Kumara
Harold & Inge Marcus Department of Industrial and Manufacturing Engineering,
The Pennsylvania State University,
310 Leonhard Building,
University Park, PA 16802
{vineet,ngautam,skumara}@psu.edu

and

M. Greaves
DARPA Tactical Technology Office,
3701 Fairfax Drive
Arlington, VA 22203-1714
mgreaves@darpa.mil

Abstract

This paper is motivated by the problem of capturing and releasing the CPU by a routine software application in order to accommodate other non-routine requests that need the CPU. Specifically, we consider a network of distributed software agents where each agent is assigned with routine tasks that need to be processed by a CPU. The CPU also receives requests from other processes running on the machine. The problem is to select an optimal threshold on the workload of the agent so that the agent releases the CPU and recaptures it from time-to-time based on its workload.

In order to do that, we use a stochastic fluid-flow model with two buffers, one for the agent that runs the routine tasks and the other for the remaining non-routine jobs at the CPU. Input to the two buffers are from on-off sources and the processor switches between the two buffers using a threshold-based policy. We develop analytical approximations for the buffer content distribution and determine the Quality of Service (QoS) experienced by the two sources of traffic. We use the QoS performance measures to formulate and solve an optimization problem to design an optimal threshold value.

Keywords : Quality of Service, multi-class traffic, fluid queues, multi-agent systems.

*Corresponding Author

1 Introduction

Consider a distributed agent architecture where each software agent resides on a workstation and shares the CPU with other jobs running on that workstation. We assume that the CPU is always available for processing tasks of the agent. Since the agent receives tasks only in spurts, it sometimes could share the CPU with other applications running on the workstation. Based on the workload of the agent, it decides when to release and when to capture the CPU. We study the following control policy. When the workload becomes zero, the agent gives up the CPU and waits until the workload reaches a threshold "a" when it recaptures the CPU. The main aim of this paper is to obtain optimal values for a so that performance measures for both the agent and the other jobs on the CPU are optimized.

Typically, such problems are modeled analytically as multi-class queueing systems where each class generates traffic into its own buffer and a single server empties out the buffers in some fashion. The two main types of analysis employed in the literature are: (i) determining the best policy for the server to empty out the different buffers with fairly simplistic objective functions; (ii) for a given policy, determining the performance measures experienced by the system and fine-tuning the policy under very complex objective functions. Problems such as load balancing (see Mirchandany et al [22], Kostin et al [18] and Tantawi and Towsley [26]) and armed bandits (see Darce et al [9] and Whittle [28]), fall under the former category. Polling systems (see Takagi [25], Daganzo [8], Boxma [7], and the references therein) and its variants including this paper, fall under the latter case.

Polling systems and its variants have been well-studied in the literature especially for discrete arrival systems for various policies. In Kesidis [17], policies such as Packetized General Processor Sharing Mechanism and Weighted Round Robin Mechanism are considered for example. Some of the other polling policies studied are full-service exhaustive policy and gated policy. A comprehensive treatment of polling systems for discrete arrivals can be found in Takagi [25] and Daganzo [8]. There is also a significant amount of literature on policies based on static priority. However for fluid arrivals (where customers or packets enter queues continuously one behind the other) the literature is currently in the developing stages.

Since tasks arrive in bursts at the agent, we use fluid models to analyze the system performance. In particular, we use a two-class fluid queueing system where the agent traffic belongs to one class and the rest of the traffic belong to another class. The justification for using fluid-flow models is explained in Section 2. We analyze the traffic flow by approximating it by fluids, following the large literature using fluid-flow models for computer and communication systems (see Anick et al [1], and, Elwalid and Mitra [10]).

Several policies for scheduling multi-class fluid traffic at a node have been analyzed using fluid models for design and performance of computer and communication networks. The threshold policy described in this paper is relatively less studied especially under fluid traffic conditions. Policies such as timed round-robin and static priority are well established (see Gautam and Kulkarni [14] for a comprehensive list of references). Under timed round robin policy (a variation of polling), the scheduler serves the buffers in a fixed cyclical fashion. The static priority service policy is a special case (when $a = 0$) of the threshold policy considered in this paper. Narayanan and Kulkarni [23] analyze multi-class fluid models that use static priority service policy. They develop the marginal buffer-content distributions for each class of fluid. Zhang [29] analyzes the joint distribution of the buffer contents of each class under static priority service policy. Elwalid and Mitra [12] develop a large-deviations based approach to evaluate the buffer content distributions for the static priority service policy.

For the multi-class traffic performance analysis, the concept of effective bandwidths and its applications to solving QoS problems is well established. Gibbens and Hunt [15], Kesidis et al [16], Elwalid and Mitra [11], Kulkarni [21], Choudhury et al [4], and Whitt [27], discuss the concept of effective bandwidths. We use effective bandwidths and their extensions developed in Palmowski and Rolski [24] and Gautam et al [13].

The paper is organized as follows. In Section 2 we describe the problem setting and define the notation used in this paper. In Section 3 we recapitulate some preliminary results in terms of single-class, single-buffer fluid models. We begin our analysis for the case when the switch-over times between the different classes of traffic is zero in Section 4. Then in Section 5, we consider the case when we have a non-zero switch-over time. Numerical results of our work are presented in Section 6. Analysis involving generalizations in terms of number and type of sources are included in Section 7. We describe some concluding remarks and directions for future work in Section 8.

2 Problem Description

Distributed systems of dynamic heterogeneous components without centralized control have emerged in a very wide range of applications. The motivation for this paper comes from a large scale distributed multi-agent system for military logistics called "Cognitive Agent Architecture" (COUGAAR see http://www.cougaar.org). The COUGAAR system comprises of several software agents that work in a cooperative manner in order to accomplish tasks for military logistics planning operations. The system must be able to dynamically adapt to the changes in its environment and reallocate essential processing, if necessary. It must be done in a distributed manner with no centralized control or human intervention. We focus on a society of agents where all the agents are distributed over different machines (or computers). At each machine, there is a single Central Processing Unit (CPU) that processes tasks submitted not only by the agent on the machine but also by other applications running on the machine.

Especially under wartime conditions when the resources are limited (such as a set of computers on a ship or air-craft), it becomes crucial for the software agent to share the CPU with other, possibly critical, applications. In that light, our objective is to adaptively control when the software agent captures the CPU and when the agent releases it for other applications. Other approaches that have been used in load-balancing to improve utilization and sharing of resources apply techniques at the kernel (operating system) or the hardware level. These techniques have a limitation in that they cannot be reconfigured easily by the end-user to adapt to changing load characteristics or different objectives set for the system.

2.1 Scenario

We consider the scenario of a CPU that receives jobs from a software agent as well as other applications. Assume there are two buffers, one for the agent and one for the rest of the processes. Tasks arrive at the two buffers in bursts. The CPU processes jobs at rate c when it is serving either buffers. When the agent buffer becomes empty, the agent gives up the CPU and waits until the workload reaches a threshold "a". At this time it recaptures the CPU and gets its tasks processed at rate c. There are two target values B_1 and B_2 for the agent buffer and the other buffer respectively, such that it is undesirable to exceed those quantities. They could be thought of as buffer sizes, if so desired. The objective is to obtain an optimal a value so that a weighted sum of the probability of exceeding limits B_1 and B_2 in buffers 1 and 2 respectively is minimized.

2.2 Modeling traffic as on-off fluids

We collected timestamps for over 60,000 tasks that were received by an agent in the COUGAAR system submitted to a CPU. Figure 1 illustrates the instances when the agent receives tasks. Clearly, the bursty nature of the task arrival process can be modeled as an on-off fluid. Although only a sample is presented, the task arrivals at all the agents can be modeled as on-off fluids. In addition, even the non-agent jobs, that other applications submit arrive in bursts and therefore fluid models are appropriate. We therefore use fluid models for analyzing the performance of such systems.

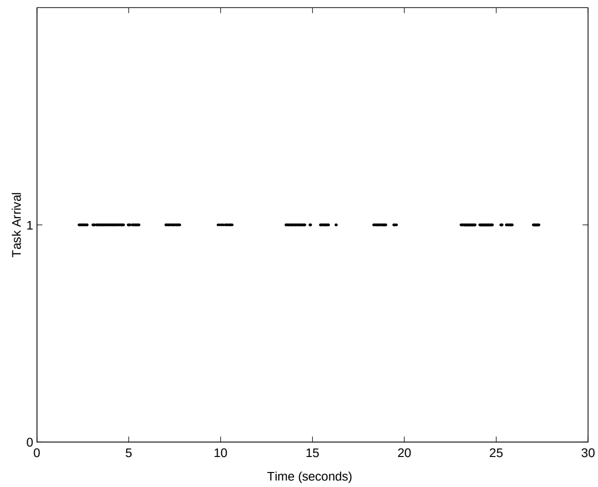

Figure 1: Tasks submitted by an agent

2.3 Model: Two-buffer fluid-flow model of the system

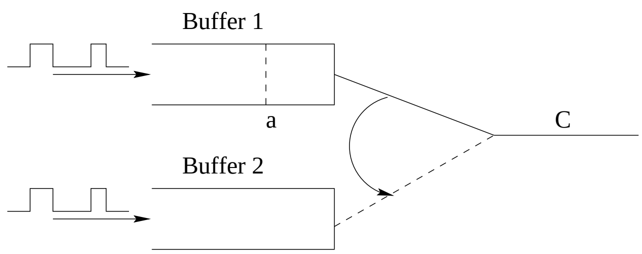

Figure 2: Two-Buffer System

We model the above scenario as a two-buffer fluid-flow system as shown in Figure 2. Both buffers are of infinite capacity. For $j = 1, 2$, fluid enters buffer j according to an alternating on-off process such that for an exponentially distributed time (with mean $1/\alpha_j$) fluid enters continuously at rate r_j and then no fluid enters for another exponentially distributed time (with mean $1/\beta_j$). When the off-time ends, another on-time starts, and so on. Let $X_j(t)$ be the amount of fluid in buffer j (for $j = 1, 2$) at time t. A scheduler alternates between buffers 1 and 2 while draining out fluid continuously at rate c. Without loss of generality assume that $r_1 > c$ and $r_2 > c$. Other conditions, such as stability will be derived subsequently. The policy adopted by the scheduler is as follows: as soon as buffer 1 becomes empty (i.e. $X_1(t) = 0$) the scheduler switches from buffer 1 to buffer 2. When the buffer contents in buffer 1 reaches a (i.e. $X_1(t) = a$), the scheduler switches back from buffer 2 to buffer 1. We denote a as the *threshold* for buffer 1. Note that the scheduler's policy is dependent on buffer 1 only. That means even if buffer 2 is empty (i.e. $X_2(t) = 0$), as long as buffer 1 has less than a (i.e. $X_1(t) < a$), the scheduler does not switch back to buffer 1. In addition, there is a constant switch-over time θ that is incurred every time the scheduler switches from one buffer to the other. Figure 3 illustrates a sample path of the buffer content processes $X_1(t)$ and $X_2(t)$. The numerical values used to generate this sample path are given in Section 6.

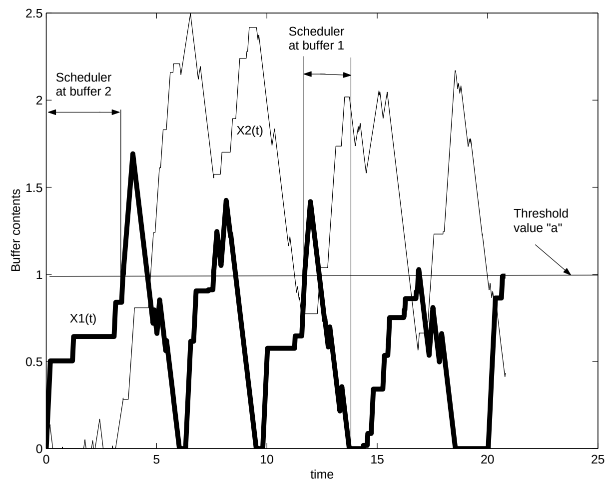

Figure 3: A sample path of buffercontent process

2.4 Objective

For the model described in Section 2.3 above, the objective of this paper is to:

- derive analytical expressions or approximations for the steady-state buffer content distributions (queue length distributions in fluid queues), $\lim_{t\to\infty} P\{X_1(t) > x\}$ and $\lim_{t\to\infty} P\{X_2(t) > x\}$. Then use the analytical expressions for the infinite buffers to approximate the buffer overflow probabilities as

$$\epsilon_1 = \lim_{t\to\infty} P\{X_1(t) > B_1\} \quad \text{and} \quad \epsilon_2 = \lim_{t\to\infty} P\{X_2(t) > B_2\}$$

  where B_1 and B_2 are the actual buffer sizes of the two buffers.

- solve an optimization problem to design a threshold a so that the weighted sum of QoS measures ϵ_1 and ϵ_2 are minimized, i.e., $\min_a \{w_1\epsilon_1 + w_2\epsilon_2\}$.

3 Preliminaries

We first present some preliminary results from the literature before addressing the objectives described in Section 2.4. Consider a single infinite-sized buffer (with constant output capacity c) that

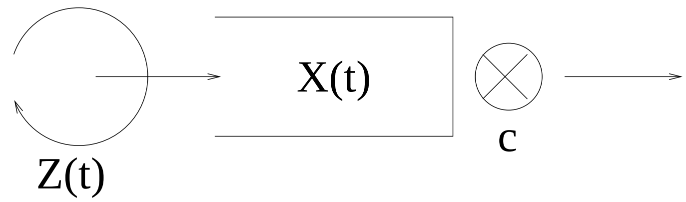

Figure 4: Single Buffer Fluid Model

admits traffic from a source driven by a random environment process $\{Z(t), t \geq 0\}$ (see Figure 4). At time t, the source generates fluid at rate $r(Z(t))$. Let $X(t)$ be the amount of fluid in the buffer at time t. We present some preliminary results from studies on such single buffer fluid models.

3.1 Effective bandwidths

Assume that the environment process $\{Z(t), t \geq 0\}$ is a stationary and ergodic process satisfying the Gärtner-Ellis conditions (see Kesidis et al [16]). Then, for a given v ($v > 0$), the *effective bandwidth* of the source is

$$eb(v) = \lim_{t \to \infty} \frac{1}{vt} \log E \left\{ \exp \left(v \int_0^t r(Z(t)) dt \right) \right\}. \tag{1}$$

When the $\{Z(t), t \geq 0\}$ process can be modeled as certain special stochastic processes, Kesidis et al [16], Elwalid and Mitra [11] and Kulkarni [21] illustrate how to compute $eb(v)$ in those cases. We consider two special cases that would be used later in this paper.

1. *General on-off source:* Consider a source modulated by a two-state (on and off) process $\{Z(t), t \geq 0\}$ that alternates between on and off states. The random amount of time the process spends in the on state (called *on-times*) has cdf $U(\cdot)$ with mean τ_U and the corresponding *off-time* cdf is $D(\cdot)$ with mean off time τ_D. Fluid is generated continuously at rate r during the on state and at rate 0 during the off state. The effective bandwidth of this source is given by (see Kulkarni [21]) the unique solution to

$$\tilde{U}(v\, eb(v) - rv)\ \tilde{D}(v\, eb(v)) = 1 \tag{2}$$

   where $\tilde{U}(\cdot)$ and $\tilde{D}(\cdot)$ are the Laplace Stieltjes transforms (LSTs) of $U(\cdot)$ and $D(\cdot)$ respectively.

2. *Exponential on-off source:* When $U(t) = 1 - e^{-\alpha t}$ and $D(t) = 1 - e^{-\beta t}$ for $t \geq 0$, we call the source an exponential on-off source with on-time parameter α, off-time parameter β and rate r. The effective bandwidth of the exponential on-off source is (see Elwalid and Mitra [11])

$$eb(v) = \frac{rv - \alpha - \beta + \sqrt{(rv - \alpha - \beta)^2 + 4\beta rv}}{2v}. \tag{3}$$

3.2 Bounds on limiting distribution of buffer content

Bounds on the limiting distribution of the buffer content process $\{X(t), t \geq 0\}$ when $\{Z(t), t \geq 0\}$ is a semi-Markov process is derived in Gautam et al [13]. We now present the case for a general on-off source described in Section 3.1. The effective bandwidth of a source, $eb(v)$ is an increasing function of v (evident from Equation (1)). Also, if the environment process $\{Z(t), t \geq 0\}$ is a stationary and ergodic process satisfying the Gärtner-Ellis conditions (see Kesidis et al [16]), then

$$\lim_{t\to\infty} E[r(Z(t))] \leq eb(v) \leq \lim_{t\to\infty} \max\{r(Z(t))\}.$$

In fact as v is increased from 0 to ∞, $eb(v)$ increases from the lower bound to the upper bound in the above equation. Therefore, if c is such that $\lim_{t\to\infty} E[r(Z(t))] < c \leq \lim_{t\to\infty} \max\{r(Z(t))\}$, then there exists a unique solution η that satisfies the following equation:

$$eb(\eta) = c. \tag{4}$$

The limiting distribution of the buffer contents are bounded as (see Gautam et al [13])

$$C_* e^{-\eta x} \leq \lim_{t\to\infty} P\{X(t) > x\} \leq C^* e^{-\eta x}, \tag{5}$$

where

$$C^* = \frac{\tilde{U}(-\eta(r-c)) - 1}{\tau_U + \tau_D} \frac{r}{c(r-c)\eta \inf_x \left\{ \frac{\int_x^\infty e^{\eta(r-c)(y-x)} dU(y)}{1-U(x)} \right\}} \tag{6}$$

and

$$C_* = \frac{\tilde{U}(-\eta(r-c)) - 1}{\tau_U + \tau_D} \frac{r}{c(r-c)\eta \sup_x \left\{ \frac{\int_x^\infty e^{\eta(r-c)(y-x)} dU(y)}{1-U(x)} \right\}}. \tag{7}$$

For the special case when $\{Z(t), t \geq 0\}$ corresponds to an exponential on-off source (see Section 3.1 for notation), the upper and lower bounds in Equation (5) become equal. Thus the limiting distribution of the buffer content is

$$\lim_{t\to\infty} P\{X(t) > x\} = \frac{r\beta}{c(\alpha+\beta)} e^{-\eta x}, \tag{8}$$

where

$$\eta = \frac{c\alpha + c\beta - \beta r}{c(r-c)}.$$

When traffic from two independent general on-off sources are multiplexed into a single infinite capacity buffer, we derive bounds for the limiting distribution of the buffer contents as follows:

$$K1_* K2_* e^{-\eta x} \leq \lim_{t\to\infty} P\{X(t) > x\} \leq K1^* K2^* e^{-\eta x}, \tag{9}$$

where η is the solution to

$$eb_1(\eta) + eb_2(\eta) = c, \tag{10}$$

$$K1^* = \left[\frac{\tilde{U_1}(-\eta(r_1 - eb_1(\eta))) - 1}{\tau_U^1 + \tau_D^1} \frac{r_1}{eb_1(\eta)(r_1 - eb_1(\eta))\eta \inf_x \left\{ \frac{\int_x^\infty e^{\eta(r_1 - eb_1(\eta))(y-x)} dU_1(y)}{1 - U_1(x)} \right\}} \right],$$

$$K2^* = \left[\frac{\tilde{U_2}(-\eta(r_2 - eb_2(\eta))) - 1}{\tau_U^2 + \tau_D^2} \frac{r_2}{eb_2(\eta)(r_2 - eb_2(\eta))\eta \inf_x \left\{ \frac{\int_x^\infty e^{\eta(r_2 - eb_2(\eta))(y-x)} dU_2(y)}{1 - U_2(x)} \right\}} \right],$$

$$K1_* = \left[\frac{\tilde{U_1}(-\eta(r_1 - eb_1(\eta))) - 1}{\tau_U^1 + \tau_D^1} \frac{r_1}{eb_1(\eta)(r_1 - eb_1(\eta))\eta \sup_x \left\{ \frac{\int_x^\infty e^{\eta(r_1 - eb_1(\eta))(y-x)} dU_1(y)}{1 - U_1(x)} \right\}} \right],$$

and

$$K2_* = \left[\frac{\tilde{U_2}(-\eta(r_2 - eb_2(\eta))) - 1}{\tau_U^2 + \tau_D^2} \frac{r_2}{eb_2(\eta)(r_2 - eb_2(\eta))\eta \sup_x \left\{ \frac{\int_x^\infty e^{\eta(r_2 - eb_2(\eta))(y-x)} dU_2(y)}{1 - U_2(x)} \right\}} \right].$$

Subscripts and superscripts "1" and "2" are used in the above relations to denote the corresponding values for the two different sources. Note that Artiges and Nain [2] also obtain exponential bounds for multiplexing non-homogeneous multiclass Markovian on-off sources. Their upper bounds are similar to those reported above.

3.3 First passage times in fluid flow models

Let T be the time when the buffer contents $X(t)$ for the first time reaches zero. Therefore

$$T = \inf\{t > 0 \; : \; X(t) = 0\}.$$

Let $\{Z(t), t \geq 0\}$ be an exponential on-off source (with parameters defined in Section 3.1) such that if $Z(t) = 0$, the source is off at time t and if $Z(t) = 1$, the source is on. Define the first passage time distribution

$$G_x(t) = P\{T \leq t | X(0) = x, Z(0) = 1\}.$$

The LST of $G_x(t)$ can be derived from Narayanan and Kulkarni [23] and Gautam et al [13] as

$$\tilde{G_x}(w) = \begin{cases} \frac{w + \beta + cs_0(w)}{\beta} e^{x \, s_0(w)} & \text{if } w \geq w^* \\ \infty & \text{otherwise,} \end{cases} \tag{11}$$

where $w^* = (2\sqrt{c\alpha\beta(r-c)} - r\beta - c\alpha - c\beta)/r$, $s_0(w) = \frac{-b - \sqrt{b^2 + 4w(w+\alpha+\beta)c(r-c)}}{2c(r-c)}$ and $b = (r-2c)w + (r-c)\beta - c\alpha$.

3.4 Steady state distribution for fluid flow models

Consider a single buffer fluid model described in Figure 4 with the only exception that the output capacity is not a constant, but a stochastic and time-varying quantity $\mu(Z(t))$. Recollect that $r(Z(t))$ is the rate at which fluid enters the buffer at time t and $\{Z(t), t \geq 0\}$ is the environment process. Kulkarni and Rolski [19] use the analysis in Borovkov [3] p. 24 to prove the existence of a unique steady state distribution for the above fluid model. They state that, *if $\{Z(t), t \geq 0\}$ is stationary (in strict sense) ergodic and $E[R(Z(t))] - E[\mu(Z(t))] < 0$, then for each $x \geq 0$, there exists a finite (a.s.) random variable X^* such that*

$$\lim_{t\to\infty} P(X(t) > x) = P(X^* > x).$$

4 Analysis: Zero Switch-over Times

Consider the two-buffer fluid model in Section 2.3. In this section we derive the buffer content distributions $\lim_{t\to\infty} P\{X_1(t) > x\}$ and $\lim_{t\to\infty} P\{X_2(t) > x\}$ for the zero switch-over time case ($\theta = 0$). In Section 5 we will extend these results to the $\theta > 0$ case by making suitable approximations.

4.1 Time between switching

Under the zero switch-over time case, we first derive expressions for the time the scheduler spends in each of the buffers. For $j = 1, 2$, let T_j be the time spent by the scheduler serving buffer j before switching to the other buffer. Mathematically we denote

$$T_1 = \inf\{t > 0 \ : \ X_1(t) = 0 | X_1(0) = a\}$$

and

$$T_2 = \inf\{t > 0 \ : \ X_1(t) = a | X_1(0) = 0\}.$$

We now study the distributions of T_1 and T_2, as well as derive expressions for their moments.

4.1.1 Distribution of time spent on buffer 1

Let $O_1(t)$ be the cdf of the random variable T_1 such that

$$O_1(t) = P\{T_1 \leq t\}.$$

Define $\tilde{O_1}(w)$ as the LST of $O_1(t)$ such that

$$\tilde{O_1}(w) = E[e^{-wT_1}].$$

Theorem 1 *The LST $\tilde{O_1}(w)$ is*

$$\tilde{O_1}(w) = \begin{cases} \frac{w+\beta_1+cs_0(w)}{\beta_1} e^{a\, s_0(w)} & \text{if } w \geq w^* \\ \infty & \text{otherwise,} \end{cases} \tag{12}$$

where

$$w^* = (2\sqrt{c\alpha_1\beta_1(r_1 - c)} - r_1\beta_1 - c\alpha_1 - c\beta_1)/r_1, \tag{13}$$

$$s_0(w) = \frac{-b - \sqrt{b^2 + 4w(w + \alpha_1 + \beta_1)c(r_1 - c)}}{2c(r_1 - c)} \tag{14}$$

and $b = (r_1 - 2c)w + (r_1 - c)\beta_1 - c\alpha_1$.

Proof. Recall that due to the definition of T_1, the source is "on" initially, so we are in the situation described in Section 3.3. Therefore, by directly substituting for the first passage time expressions in Section 3.3 and Equation (11), we obtain the above expressions. ■

4.1.2 Distribution of time spent on buffer 2

Let $O_2(t)$ be the cdf of the random variable T_2 (defined in Section 4.1) such that

$$O_2(t) = P\{T_2 \leq t\}.$$

Define $\tilde{O_2}(s)$ as the LST of $O_2(t)$ such that

$$\tilde{O_2}(s) = E[e^{-sT_2}].$$

Theorem 2 *The LST $\tilde{O_2}(s)$ is*

$$\tilde{O_2}(s) = \frac{\beta_1}{\beta_1 + s} e^{-a\frac{\alpha_1 s + \beta_1 s + s^2}{r_1 s + r_1\beta_1}}. \tag{15}$$

Proof. During time T_2, the scheduler serves only buffer 2. Therefore the buffer 1's contents $X_1(t)$ is non-decreasing. Clearly,

$$O_2(t) = P\{T_2 \leq t\} = P\{X_1(t) > a | X_1(0) = 0\}. \tag{16}$$

For all $t \in [0, \infty)$, let $Z_1(t) = 0$ denote that source 1 is off and $Z_1(t) = 1$ denote that source 1 is on at time t. Define for $i = 0, 1$

$$H_i(x, t) = P\{X_1(t) \leq x, Z_1(t) = i\}.$$

Also define the vector $H(x, t) = [H_0(x, t) \;\; H_1(x, t)]$. By a similar argument in Anick et al [1] we can show that $H(x, t)$ satisfies the following partial differential equation

$$\frac{\partial H(x,t)}{\partial t} + \frac{\partial H(x,t)}{\partial x} R = H(x,t)Q \tag{17}$$

with initial conditions $H_0(x,0) = 1$ and $H_1(x,0) = 0$, where

$$R = \begin{bmatrix} 0 & 0 \\ 0 & r_1 \end{bmatrix} \text{ and } Q = \begin{bmatrix} -\beta_1 & \beta_1 \\ \alpha_1 & -\alpha_1 \end{bmatrix}.$$

Now, taking the Laplace transform of Equation (17) with respect to t, we get

$$sH^*(x,s) - H(x,0) + \frac{\partial H^*(x,s)}{\partial x} R = H^*(x,s)Q. \tag{18}$$

Due to the initial conditions below Equation (17), Equation (18) reduces to

$$sH^*(x,s) - [1 \quad 0] + \frac{\partial H^*(x,s)}{\partial x} R = H^*(x,s)Q.$$

Taking the LST of the above equation with respect to x yields

$$s\tilde{H}^*(w,s) - [1 \quad 0] + w\tilde{H}^*(w,s)R - wH^*(0,s)R = \tilde{H}^*(w,s)Q.$$

Since $P\{X_1(t) \leq 0, Z_1(t) = 1\} = 0$, we have $H^*(0,s) = [H_0^*(0,s) \quad 0]$ and therefore $wH^*(0,s)R = [0 \quad 0]$. Hence the above equation reduces to

$$\tilde{H}^*(w,s) = [1 \quad 0][sI + wR - Q]^{-1}.$$

Plugging in for R and Q, and taking the inverse of the matrix yields

$$\tilde{H}^*(w,s) = \frac{1}{wr_1(s+\beta_1) + \alpha_1 s + \beta_1 s + s^2}[s + wr_1 + \alpha_1 \qquad \beta_1]. \tag{19}$$

However,

$$\tilde{O}_2(s) = 1 - sH_0^*(a,s) - sH_1^*(a,s).$$

Therefore inverting the transform in Equation (19) with respect to w, and substituting in the above equation, yields the desired result. ■

4.1.3 Average time between switching

Having derived expressions for the distributions of T_1 and T_2, we now address the issue of computing their moments. In particular we derive a closed-form algebraic expression for the mean of both T_1 and T_2 in the next theorem.

Theorem 3 *The mean of T_1 and T_2 are given by*

$$E[T_1] = \frac{r_1 + a(\alpha_1 + \beta_1)}{c\alpha_1 + c\beta_1 - r_1\beta_1} \tag{20}$$

$$E[T_2] = \frac{r_1 + a(\alpha_1 + \beta_1)}{r_1\beta_1} \tag{21}$$

Proof. Using the relations $E[T_1] = -\frac{d\tilde{O_1}(w)}{dw}$ at $w = 0$ and $E[T_2] = -\frac{d\tilde{O_2}(s)}{ds}$ at $s = 0$, the theorem can be proved. ∎

Remark 1 *The ratio $E[T_1]/E[T_2]$ is independent of a. This indicates that no matter what the threshold is, the ratio of time spent by the scheduler serving buffers 1 and 2 remains the same.*

Before computing the performance measures, the important question to ask is when do the limiting distributions exist for the buffer contents $X_1(t)$ and $X_2(t)$. The following theorem describes the condition under which the limiting distributions exist.

Theorem 4 *Limiting distributions for the buffer contents $X_1(t)$ and $X_2(t)$ exist when*

$$\frac{r_1\beta_1}{\alpha_1+\beta_1} + \frac{r_2\beta_2}{\alpha_2+\beta_2} < c. \tag{22}$$

Proof. Using the notation and results in Section 3.4, in order to show that steady distribution for a fluid queue exists, it is enough to show that (see Kulkarni and Rolski [19] and Borovkov [3] p. 24) (i) the environment process $\{Z(t), t \geq 0\}$ is stationary ergodic, and (ii) the mean input rate is smaller than the mean output capacity, i.e., $E[R(Z(t))] - E[\mu(Z(t))] < 0$.

Clearly the environment processes driving the two fluid sources, $Z_1(t)$ and $Z_2(t)$ are stationary ergodic. Therefore it is enough to find the conditions when the mean input rate is less than the available output capacity to determine when steady state distributions for $X_1(t)$ and $X_2(t)$ exist. In other words, if

$$\frac{r_1\beta_1}{\alpha_1+\beta_1} < c, \tag{23}$$

steady state distribution for $X_1(t)$ exists. Also, steady-state distribution for $X_2(t)$ exists if the mean input rate is less than the available output capacity, i.e.

$$\frac{r_2\beta_2}{\alpha_2+\beta_2} < c\frac{E[T_2]}{E[T_1]+E[T_2]} + 0\frac{E[T_1]}{E[T_1]+E[T_2]} = c - \frac{r_1\beta_1}{\alpha_1+\beta_1}. \tag{24}$$

Since (24) is more constraining than (23), it is the necessary condition for the existence of limiting distributions for buffers 1 and 2. ∎

4.2 Buffer 2 analysis

We begin by studying the buffer content process of buffer 2. In particular, our aim is to derive the limiting buffer-content distribution as our main performance measure. If we consider buffer 2 in isolation, its input is from an exponential on-off source but the output capacity alternates between c (for T_2 time) and 0 (for T_1 time). However the effective-bandwidth approximation and the bounds detailed in Section 3 assume that the output channel capacity is a constant. In order to utilize those techniques, we need to first transform our model into an appropriate one with a constant output channel capacity. The next section outlines a procedure for this.

4.2.1 Compensating source

Consider a single-buffer fluid model for buffer 2 with a constant output channel capacity c whose input is generated by the original exponential on-off source and a fictitious compensating source. The compensating source is such that it stays on for T_1 time units and off for T_2 time units. When the compensating source is on, it generates fluid at rate c and when it is off it generates fluid at rate 0. Note that the compensating source is independent of the original source. Clearly, the dynamics of the buffer-content process (of buffer 2) remain unchanged for this transformed single-buffer-fluid model with two input sources (including the compensating source) and constant output capacity c.

4.2.2 Limiting distribution of buffer contents

Using the above model and the analysis in Section 3.2, we can derive the limiting distribution of the buffer contents of buffer 2. We first obtain the effective bandwidth of source 2 (the original source into buffer 2) using Equation (3) as

$$eb_2(v) = \frac{r_2 v - \alpha_2 - \beta_2 + \sqrt{(r_2 v - \alpha_2 - \beta_2)^2 + 4\beta_2 r_2 v}}{2v}.$$

Using Equation (2) we can derive the effective bandwidth of the compensating source ($eb_0(v)$) as the unique solution to

$$\tilde{O_1}(v\, eb_0(v) - cv)\, \tilde{O_2}(v\, eb_0(v)) = 1$$

where $\tilde{O_1}(\cdot)$ and $\tilde{O_2}(\cdot)$ can be obtained from Equations (12) and (15) respectively. Using Equation (10) we obtain η as the solution to

$$eb_0(\eta) + eb_2(\eta) = c. \tag{25}$$

Theorem 5 *The limiting distribution of the contents of buffer 2 is bounded as*

$$LB \le \lim_{t\to\infty} P\{X_2(t) > x\} \le UB \tag{26}$$

where

$$LB = \frac{r_2\beta_2}{eb_2(\eta)(\alpha_2+\beta_2)} \frac{1-\tilde{O_2}(\eta eb_0(\eta))}{E(T_1)+E(T_2)} \frac{c}{eb_0(\eta)(c-eb_0(\eta))\eta} e^{-\eta x},$$

$$UB = \frac{r_2\beta_2}{eb_2(\eta)(\alpha_2+\beta_2)} \frac{1/\tilde{O}_2(\eta eb_0(\eta)) - 1}{E(T_1)+E(T_2)} \frac{c}{eb_0(\eta)(c-eb_0(\eta))\eta} e^{-\eta x}$$

and η is the solution to Equation (25).

Proof. Using Equation (9), making the appropriate variable substitutions, then going over the algebra and taking the limits will yield the bounds in Inequality (26). ■

4.3 Buffer 1 analysis

We now derive the limiting distribution of the buffer contents of buffer 1. Note that the scheduler is never idle when it is serving buffer 1. This makes it impossible to use a compensating source and study the buffer contents similar to buffer 2. Therefore we resort to alternate techniques as detailed in the following sections.

4.3.1 SMP model

For a value of x greater than the threshold a, we can model the system as a semi-Markov process (SMP) and derive the limiting distribution $\lim_{t\to\infty} P\{X_1(t) > x\}$. Fix x ensuring that $x > a$. For all $t \in [0, \infty)$, let $Z_1(t) = 0$ denote that source 1 is off and $Z_1(t) = 1$ denote that source 1 is on at time t. Consider the following four states of the system:

$$\begin{aligned}
Y(t) = 1 &\quad \text{if} \quad X_1(t) = 0 \text{ and } Z_1(t) = 0,\\
Y(t) = 2 &\quad \text{if} \quad X_1(t) = a \text{ and } Z_1(t) = 1,\\
Y(t) = 3 &\quad \text{if} \quad X_1(t) = x \text{ and } Z_1(t) = 1,\\
Y(t) = 4 &\quad \text{if} \quad X_1(t) = x \text{ and } Z_1(t) = 0,\\
Y(t) = 0 &\quad \quad \text{otherwise.}
\end{aligned}$$

Let S_n be the n^{th} time the system reaches one of the four states 1, 2, 3 or 4. Also let $Y_n = Y(S_n)$.

Theorem 6 *The sequence of bivariate random variables $\{(Y_n, S_n), n \geq 0\}$ is a Markov renewal sequence.*

Proof. Directly follows from the definition of Markov renewal sequences (see Kulkarni [20]). ■

4.3.2 Approximate model

The greatest difficulty is in obtaining the kernel of the Markov renewal sequence defined in the previous section. Therefore we resort to approximations here. Consider the setting in the previous section where we fix x such that $x > a$. If x is suitably larger than a, then experimentally we have found that once the buffer contents $X_1(t)$ falls to a from above (source 1 will be off at this time), there is a very small probability that the buffer contents will exceed a before reaching zero. This is also evident from the sample path shown in Figure 3. Under that restrictive framework, we divide a scheduler service cycle into 3 parts:

1. V_1*: the time when the scheduler is serving buffer 2.* This is also the time when buffer 1 climbs from 0 to a. From Equation (21) we get

$$E[V_1] = \frac{r_1 + a(\alpha_1 + \beta_1)}{r_1\beta_1}.$$

2. V_2*: the time when the scheduler is serving buffer 1 and the buffer content of buffer 1 goes above* a *and returns back to* a*.* This is like the first passage time for the buffer when it leaves zero and returns to zero. The expected time is

$$E[V_2] = \frac{r_1}{c\alpha_1 + c\beta_1 - r_1\beta_1}.$$

The probability that the buffer content will exceed x in this stage is $\dfrac{r_1\beta_1}{c(\alpha_1+\beta_1)}e^{-\eta(x-a)}$, where

$$\eta = \frac{c\alpha_1 + c\beta_1 - \beta_1 r_1}{c(r_1 - c)}. \tag{27}$$

3. V_3*: the time when the scheduler is serving buffer 1 and the buffer content of buffer 1 declines from* a *to zero.* Using the above equation and Equation (20), the expected time is

$$E[V_3] = \frac{a(\alpha_1 + \beta_1)}{c\alpha_1 + c\beta_1 - r_1\beta_1}.$$

Assuming that a fraction $\dfrac{r_1\beta_1}{c(\alpha_1+\beta_1)}e^{-\eta(x-a)}$ of the time V_2, $X_1(t)$ is above x, we obtain the following approximation for the limiting distribution of the buffer contents:

$$\lim_{t\to\infty} P\{X_1(t) > x\} \approx \frac{\dfrac{r_1\beta_1}{c(\alpha_1+\beta_1)}E[V_2]e^{-\eta(x-a)}}{E[V_1] + E[V_2] + E[V_3]}. \tag{28}$$

We conducted simulation experiments to validate this approximation. Figure 5 shows that the approximation is valid even when the switch-over times are positive.

5 Approximations: Positive Switch-over Times

Now we consider the case of positive switch-over times θ defined in Section 2.3. This means every time the scheduler switches from buffer 1 to buffer 2 or from buffer 2 to buffer 1, it incurs a time

$\theta > 0$ to switch. That means in every cycle there would be a time 2θ that the scheduler does not serve either buffers. The implications of this in terms of performance is that if the scheduler switches very often (for very small a values) then a larger proportion of time is wasted than when the scheduler switches less often. This brings us to the first question (before addressing performance issues): when do the limiting distributions of the buffer contents exist?

5.1 Necessary condition for existence of distributions

In order to derive the necessary condition for the existence of steady state distributions of the buffer content processes $X_1(t)$ and $X_2(t)$, we need to define m_θ, the fraction of time the scheduler spends in switching, in the long run. Using that we state the following theorem:

Theorem 7 *The two buffers in the system defined in Section 2.3 with $\theta > 0$ have a steady-state distribution if*

$$\frac{r_1\beta_1}{\alpha_1+\beta_1} + \frac{r_2\beta_2}{\alpha_2+\beta_2} < c(1-m_\theta). \tag{29}$$

Proof. The theorem can be proved similar to Theorem 4 with the understanding that only a fraction $(1-m_\theta)$ of time the capacity c is available to the entire system. ■

The difficulty is in obtaining a closed-form expression for m_θ. We obtain an approximate expression that works well when the switching time θ is much smaller than the up times or down times of source 1. An approximation for m_θ is

$$m_\theta \approx \frac{2\theta}{\frac{r_1+(a+r_1\theta)(\alpha_1+\beta_1)}{c\alpha_1+c\beta_1-r_1\beta_1} + \frac{r_1+a(\alpha_1+\beta_1)}{r_1\beta_1} + \theta},$$

where the numerator is the time spent in one cycle in switching, and in the denominator the term $\frac{r_1+(a+r_1\theta)(\alpha_1+\beta_1)}{c\alpha_1+c\beta_1-r_1\beta_1}$ is the average time for buffer 1 to empty from $(a+r_1\theta)$ which is the approximate level when the scheduler starts serving buffer 1, the term $\frac{r_1+a(\alpha_1+\beta_1)}{r_1\beta_1}$ is the average time the scheduler spends switching and serving buffer 2 before buffer 1 reaches a, and, the term θ denotes the time the scheduler takes to return to buffer 1. We now develop approximations for the QoS measures.

5.2 Buffer 2 approximations

The analysis proceeds very similar to that in Section 4.2. Consider the compensating source defined in Section 4.2.1. Now with $\theta > 0$, the compensating source is such that its on times are $T_1^{adj} + 2\theta$ and off times are $T_2 - \theta$, where T_2 is as defined in Section 4.1 and T_1^{adj} is the adjusted T_1 defined in Section 4.1 such that the first passage time starts at $(a + r_1\theta)$ instead of a. Therefore the LSTs of the on and off times are now respectively

$$\tilde{O_1}^{adj}(w) = E[e^{-w(T_1^{adj}+2\theta)}] = \begin{cases} \left(e^{-2w\theta}\right) \frac{w+\beta_1+cs_0(w)}{\beta_1} e^{(a+r_1\theta)\, s_0(w)} & \text{if } w \geq w^* \\ \infty & \text{otherwise.} \end{cases}$$

and

$$\tilde{O_2}^{adj}(s) = E[e^{-s(T_2-\theta)}] = \left(e^{s\theta}\right) \frac{\beta_1}{\beta_1 + s} e^{-a\frac{\alpha_1 s+\beta_1 s+s^2}{r_1 s+r_1\beta_1}}.$$

The effective bandwidth of the compensating source is the unique solution to

$$\tilde{O_1}^{adj}(v\, eb_0^{adj}(v) - cv)\, \tilde{O_2}^{adj}(v\, eb_0^{adj}(v)) = 1.$$

The next step is to solve for η in

$$eb_0^{adj}(\eta) + eb_2(\eta) = c.$$

Using Theorem 5 we can derive the distribution of the contents in buffer 2 as

$$LB^{adj} \leq \lim_{t\to\infty} P\{X_2(t) > x\} \leq UB^{adj},$$

where

$$LB^{adj} = \frac{r_2\beta_2}{eb_2(\eta)(\alpha_2+\beta_2)} \frac{1-\tilde{O_2}^{adj}(\eta eb_0^{adj}(\eta))}{E(T_1^{adj}+2\theta)+E(T_2-\theta)} \frac{c}{eb_0^{adj}(\eta)(c-eb_0^{adj}(\eta))\eta} e^{-\eta x},$$

and

$$UB^{adj} = \frac{r_2\beta_2}{eb_2(\eta)(\alpha_2+\beta_2)} \frac{1/\tilde{O_2}^{adj}(\eta eb_0^{adj}(\eta))-1}{E(T_1^{adj}+2\theta)+E(T_2-\theta)} \frac{c}{eb_0^{adj}(\eta)(c-eb_0^{adj}(\eta))\eta} e^{-\eta x}.$$

As an approximation for the limiting distribution of buffer contents of buffer 2, we use

$$\lim_{t\to\infty} P\{X_2(t) > x\} \approx (LB^{adj} + UB^{adj})/2. \tag{30}$$

5.3 Buffer 1 approximations

We proceed analyzing the approximations for buffer 1 in a fashion very similar to that in Section 4.3.2. The cycle is split into three parts, this time it is V_1^{adj}, V_2^{adj} and V_3^{adj}. Their expected values are approximated as

$$E[V_1^{adj}] = \frac{r_1 + (a+r_1\theta)(\alpha_1+\beta_1)}{r_1\beta_1},$$

$$E[V_2^{adj}] = \frac{r_1}{c\alpha_1 + c\beta_1 - r_1\beta_1}$$

and

$$E[V_3^{adj}] = \frac{(a+r_1\theta)(\alpha_1+\beta_1)}{c\alpha_1 + c\beta_1 - r_1\beta_1}.$$

Therefore the limiting distribution of the contents of buffer 1 can be approximated as

$$\lim_{t\to\infty} P\{X_1(t) > x\} \approx \frac{\dfrac{r_1\beta_1}{c(\alpha_1+\beta_1)} E[V_2^{adj}] e^{-\eta(x-a-r_1\theta)}}{E[V_1^{adj}] + E[V_2^{adj}] + E[V_3^{adj}]}, \tag{31}$$

where η is defined in Equation (27).

In Section 4.3.2 we derived an approximate expression for $\lim_{t\to\infty} P\{X_1(t) > x\}$ when $\theta = 0$ and here we made a further approximation to that expression. Therefore in order to test the accuracy of the expression in Equation (31), we resort to simulations. However the main concern with simulations is that the simulation results do not converge for computationally tractable number of runs and number of replications. Therefore we compare $\lim_{t\to\infty} E[X_1(t)]$, the steady-state average buffer contents for which the simulations converge. It turns out that even for $\lim_{t\to\infty} E[X_2(t)]$, the simulation does not converge due to very high variances. Obtaining the expected buffer contents in the intervals $E[V_1^{adj}]$, $E[V_2^{adj}]$ and $E[V_3^{adj}]$, and scaling them appropriately, we get

$$\lim_{t\to\infty} E[X_1(t)] \approx \frac{\frac{a}{2}E[V_1^{adj}] + \left\{\dfrac{r_1\beta_1}{\eta c(\alpha_1+\beta_1)} + a + r_1\theta\right\} E[V_2^{adj}] + \frac{a}{2}E[V_3^{adj}]}{E[V_1^{adj}] + E[V_2^{adj}] + E[V_3^{adj}]}, \tag{32}$$

where η is defined in Equation (27).

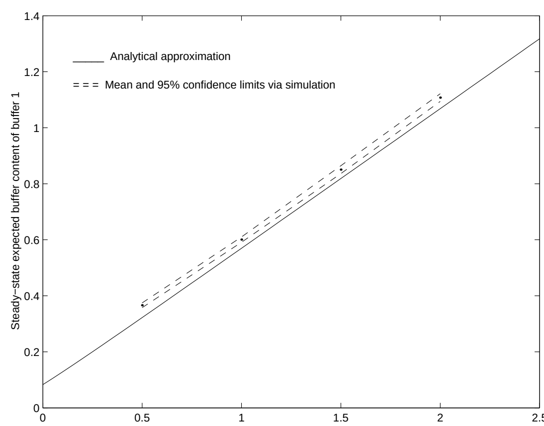

Figure 5: Testing the approximation against simulations

In Figure 5 we compare $\lim_{t\to\infty} E[X_1(t)]$ obtained from Equation (32) against simulations for various values of the threshold a. The solid line in Figure 5 is the analytical approximation and the dots represent the simulation estimates of $\lim_{t\to\infty} E[X_1(t)]$ for a values of 0.5, 1, 1.5 and 2. The two dashed lines around the dots represent the 95% confidence limits for the simulation estimates (i.e. the dots). For example, for $a = 1$, simulations yielded a mean buffer content of 0.6007 and a standard error of 0.0048 on the mean. The numerical values used as inputs to generate the figure are given in Section 6. Although the analytical approximation does not fall within the confidence interval of the simulations, it is clear from the figure that the expression in Equation (32) is very accurate. Although this does not prove that Equation (31) is accurate but leads us to believe that the approximation may be reasonable.

6 Results

We now present some numerical results based on our analysis in the previous sections. Since our main goal is to determine an optimal threshold a, in all our graphs the parameter that varies is a. The numerical values for other parameters are: $\beta_1 = 2$, $\alpha_1 = 8$, $r_1 = 2.645$, $\beta_2 = 3$, $\alpha_2 = 9$, $r_2 = 1.87$, $c = 1.06$ and $\theta = 0.0078125$, all in appropriate units. Although we assumed that the buffer sizes are infinite, for the QoS approximation for overflow probability we compute the probabilities that $X_1(t)$ and $X_2(t)$ exceed B_1 and B_2 respectively. For that purpose we chose $B_1 = 2.5$ and $B_2 = 8$. In the next section we present our performance-analysis results followed by results to obtain the optimal a in a latter section.

6.1 Buffer content distribution

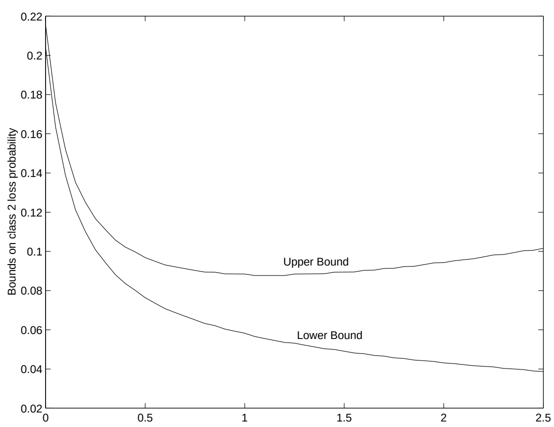

Figure 6: Bounds on the limiting distribution of buffer 2

We first consider buffer 2. For the numerical values described above, using the bounds LB^{adj} and UB^{adj} (described in Section 5.2), we obtain upper and lower bounds for $\epsilon_2 = \lim_{t\to\infty} P\{X_2(t) > B_2\}$. The bounds on the overflow probability ϵ_2 is plotted against various threshold values a in Figure 6.

An approximate expression for ϵ_2 is obtained using Equation (30). Using this approximation, ϵ_2 is plotted against various threshold values a in Figure 7. From the figure it is clear that the QoS for buffer 2 improves with increasing a.

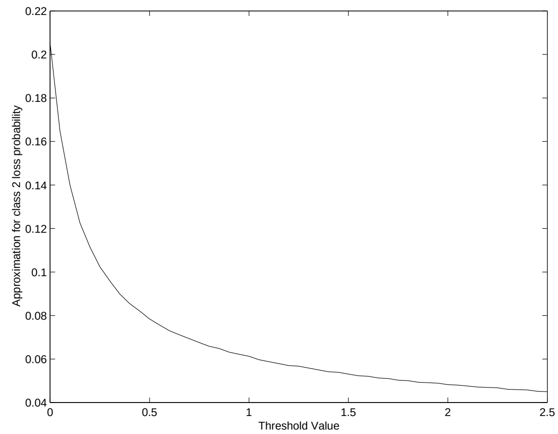

Figure 7: Approximation for limiting distribution of buffer 2

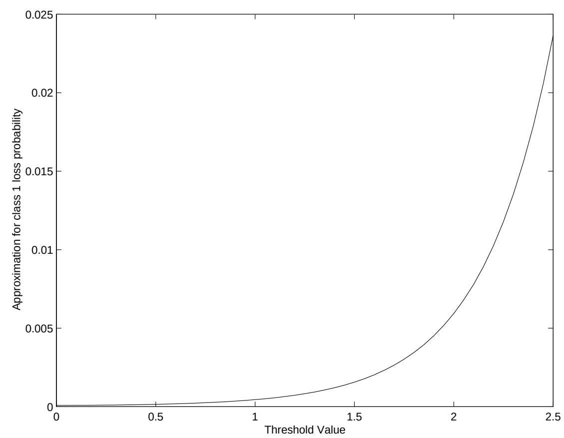

Figure 8: Approximation for limiting distribution of buffer 1

Next we consider buffer 1. We use the approximation in Equation (31) to obtain estimates for the buffer overflow probability $\epsilon_1 = \lim_{t \to \infty} P\{X_1(t) > B_1\}$. Figure 8 illustrates ϵ_1 for various values of threshold a, where it is clear that the QoS worsens for buffer 1 as a increases.

6.2 Optimization results

From Figures 7 and 8, it seems like ϵ_1 increases with a and ϵ_2 decreases with a. We would therefore like to obtain a trade-off between the performance experienced by the two classes of traffic. In particular we solve an optimization problem to select a that minimizes the weighted sum of overflow probabilities $\min_a \{w_1 \epsilon_1 + w_2 \epsilon_2\}$. We consider several values for weights $w_1 : w_2$ such that $w_1 > w_2$ since source 1 gets, in some sense, a priority on the scheduler. The objective function $w_1\epsilon_1 + w_2\epsilon_2$ is plotted against a in Figure 9 for three weight ratios $2:1$, $5:1$ and $10:1$. For example, we obtain the optimal solution for the case $w_1 : w_2 = 10 : 1$ as $a = 1.25$ which would keep the

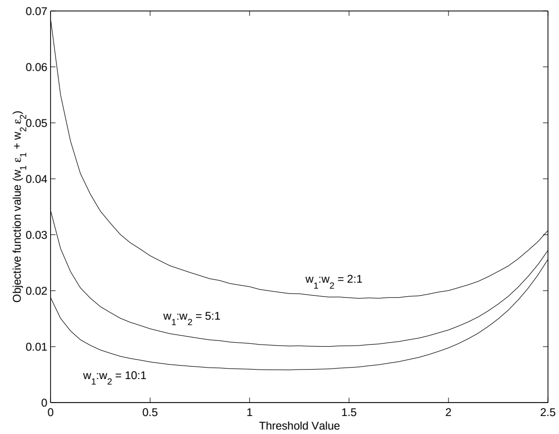

Figure 9: Optimal Solution for the Threshold Under Various Weights

objective function the lowest. The optimal solution can be found using several techniques such as the dichotomy search, golden search or complete enumeration.

7 Generalization: N general input sources for buffer 2

In this section we consider a generalization to the scenario studied thus far. The agent buffer (buffer 1 in Figure 10) and its input process i.e. exponential on-off source, remains unchanged. However we accommodate a variety of sources into buffer 2 (with possibly different traffic rates) to submit jobs to the CPU which processes them in an FCFS fashion. In particular, we assume that there are N general on-off sources that generate traffic into buffer 2 in the following manner: for $i = 1, 2, \ldots, N$, source i stays "on" for a random time with cdf $U_i(\cdot)$ and mean τ_U^i; "off" for a random time with cdf $D_i(\cdot)$ and mean τ_D^i; fluid is generated continuously at rate r_2^i during the on state and at rate 0 during the off state. The other notation and service scheduling policy are identical to that in Section 2.3. In particular, the threshold is a and the switching time is θ. The motivation for this generalization stems from the fact that although there is only one agent on a machine controlling the CPU, there could be several applications submitting jobs at different rates to buffer 2 in order to be processed by the CPU when the agent buffer relinquishes the CPU.

Based on the arguments in Section 5.1, we can state that the steady-state distributions for the buffer contents $X_1(t)$ and $X_2(t)$ exist if

$$\frac{r_1 \beta_1}{\alpha_1 + \beta_1} + \sum_{i=1}^{N} \frac{r_2^i \tau_U^i}{\tau_U^i + \tau_D^i} < c(1 - m_\theta)$$

where m_θ is defined in Section 5.1.

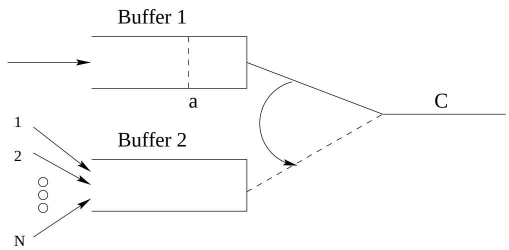

Figure 10: Two-Buffer System

Note that the scheduler's policy is dependent on buffer 1 only. In this generalization the $X_1(t)$ process is identical to that in the scenario defined in Section 2.3. So the analysis for buffer 1 is identical to that in Section 5.3 and we do not present it here. Also, since the scheduler's policy is dependent only on buffer 1, the compensating source defined in Section 5.2 for buffer 2 also remains unchanged. Therefore we only need to present the analysis for buffer 2 with the understanding that buffer 2 can be modeled as one with a constant channel capacity c but with an extra compensating source besides the N original sources.

The effective bandwidth, $eb_2^i(v)$, of source i (for $i = 1, 2, \ldots, N$) generating traffic into buffer 2 is given by (see Kulkarni [21]) the unique solution to

$$\tilde{U}_i(v\, eb_2^i(v) - r_2^i v)\; \tilde{D}_i(v\, eb_2^i(v)) = 1 \tag{33}$$

where $\tilde{U}_i(\cdot)$ and $\tilde{D}_i(\cdot)$ are the LSTs of $U_i(\cdot)$ and $D_i(\cdot)$ respectively. Recall that the effective bandwidth of the compensating source is $eb_0^{adj}(\eta)$, as defined in Section 5.2. The next step is to solve for η in

$$eb_0^{adj}(\eta) + \sum_{i=1}^{N} eb_2^i(\eta) = c.$$

Generalizing Equation (9) to accommodate the $N+1$ on-off sources (including the compensating source), we can derive the distribution of the contents in buffer 2 as

$$K_{0*}\left\{\prod_{i=1}^{N} K_{i*}\right\} e^{-\eta x} \leq \lim_{t\to\infty} P\{X_2(t) > x\} \leq K_0^*\left\{\prod_{i=1}^{N} K_i^*\right\} e^{-\eta x},$$

where

$$K_{0*} = \frac{1 - \tilde{O_2}^{adj}(\eta eb_0^{adj}(\eta))}{E(T_1^{adj} + 2\theta) + E(T_2 - \theta)} \frac{c}{eb_0^{adj}(\eta)(c - eb_0^{adj}(\eta))\eta},$$

$$K_0^* = \frac{1/\tilde{O_2}^{adj}(\eta eb_0^{adj}(\eta)) - 1}{E(T_1^{adj} + 2\theta) + E(T_2 - \theta)} \frac{c}{eb_0^{adj}(\eta)(c - eb_0^{adj}(\eta))\eta},$$

for $i = 1, 2, \ldots, N$,

$$K_{i*} = \left[\frac{\tilde{U}_i(-\eta(r_2^i - eb_2^i(\eta))) - 1}{\tau_U^i + \tau_D^i} \frac{r_2^i}{eb_2^i(\eta)(r_2^i - eb_2^i(\eta))\eta \sup_x \left\{ \frac{\int_x^\infty e^{\eta(r_2^i - eb_2^i(\eta))(y-x)} dU_i(y)}{1 - U_i(x)} \right\}} \right]$$

and

$$K_i^* = \left[\frac{\tilde{U}_i(-\eta(r_2^i - eb_2^i(\eta))) - 1}{\tau_U^i + \tau_D^i} \frac{r_2^i}{eb_2^i(\eta)(r_2^i - eb_2^i(\eta))\eta \inf_x \left\{ \frac{\int_x^\infty e^{\eta(r_2^i - eb_2^i(\eta))(y-x)} dU_i(y)}{1 - U_i(x)} \right\}} \right].$$

Note that the above extension is derived in Gautam et al [13]. In fact in Gautam et al [13], there is an expression for the bounds of the buffer content when the input sources are modulated by semi-Markov processes (and not just alternating renewal or on-off processes). Therefore it is possible to further generalize these results even further.

8 Conclusions and Extensions

In this paper we consider a 2-class fluid model with one buffer for each class of traffic. Traffic enters the two buffers according to an exponential on-off process. Traffic is emptied at rate c by a single scheduler that alternates between the two buffers. When buffer 1 becomes empty the scheduler switches to buffer 2 (after a switch-over time) and returns back to buffer 1 when the contents of buffer 1 exceeds a. Under this policy, we develop the steady state distribution of the contents of buffers 1 and 2. We use that to formulate and solve an optimization problem to design the optimal threshold a.

The next step is to use these analytical results in an implementation test-bed for agent task scheduling. In the future we will consider other threshold-based scheduling policies as well. Besides the threshold on the agent workload, we will also study thresholds on the CPU utilization. We will extend the results to a network of agents.

Appendix 1: Joint Distribution of Buffer Contents

A powerful technique to compute the joint distribution of queue lengths in discrete-customer queues (especially for 2-class or 2-buffer queues) is to formulate a Riemann-Hilbert boundary value problem. This technique is described in detail in Cohen and Boxma [5] and [6]. In fact, Boxma [7] considers a system of two identical queues, attended by a server who alternately serves a customer of each queue after spending a random switching time. In that paper, the generating function equation for the joint queue length distribution is formulated as a Riemann-Hilbert boundary value problem and solved.

Unlike discrete-customer queues, fluid queues have not enjoyed much success with Riemann-Hilbert boundary value problem formulations. Zhang [29] is one of the few papers that employs this technique for fluid queues, however the author does not use the term Riemann-Hilbert boundary value problem. The major differences (in terms of the setting) between what is described in Section 2.3 and Zhang [29] are: (i) the threshold a is zero and therefore the problem reduces to a static priority problem, (ii) there is no switching time, (iii) the input traffic to the two buffers are correlated. In Zhang [29], it is possible to conveniently analyze the processes $X_1(t)$ and $X_1(t) + X_2(t)$. However in this paper, due to the buffer-content dependent switching policy as well as the non-work-conserving service discipline, the analysis becomes complex and intractable. We now present an attempt of going over the analysis.

Recall the model in Section 2.3. In order to simplify the analysis, we assume that $\theta = 0$, i.e. switch over is instantaneous. We now present some additional notation. For $m = 1, 2$, $i = 0, 1$ and $j = 0, 1$,

$$
\begin{aligned}
Z_m(t) &= \begin{cases} 0 & \text{if source } m \text{ is off at time } t, \\ 1 & \text{if source } m \text{ is on at time } t, \end{cases} \\
Y(t) &= \begin{cases} 1 & \text{if the scheduler is serving buffer 1 at time } t, \\ 2 & \text{if the scheduler is serving buffer 2 at time } t, \end{cases} \\
H_{ij}^{(m)}(t; x_1, x_2) &= P\{X_1(t) \le x_1,\ X_2(t) \le x_2,\ Z_1(t) = i,\ Z_2(t) = j,\ Y(t) = m\}, \\
H^{(m)}(t; x_1, x_2) &= [H_{00}^{(m)}(t; x_1, x_2)\ \ H_{01}^{(m)}(t; x_1, x_2)\ \ H_{10}^{(m)}(t; x_1, x_2)\ \ H_{11}^{(m)}(t; x_1, x_2)], \\
D_1^{(1)} &= diag[-c\ \ -c\ \ r_1 - c\ \ r_1 - c], \\
D_1^{(2)} &= diag[0\ \ 0\ \ r_1\ \ r_1], \\
D_2^{(1)} &= diag[0\ \ r_2\ \ 0\ \ r_2], \\
D_2^{(2)} &= diag[-c\ \ r_2 - c\ \ -c\ \ r_2 - c], \\
Q &= \begin{bmatrix} -\beta_1 - \beta_2 & \beta_1 & \beta_2 & 0 \\ \alpha_2 & -\alpha_2 - \beta_1 & 0 & \beta_1 \\ \alpha_1 & 0 & -\alpha_1 - \beta_2 & \beta_2 \\ 0 & \alpha_1 & \alpha_2 & -\alpha_1 - \alpha_2 \end{bmatrix}.
\end{aligned}
$$

The vector $H^{(m)}(t; x_1, x_2)$ satisfies the following partial differential equation, for $m = 1, 2$,

$$
\frac{\partial}{\partial t} H^{(m)}(t; x_1, x_2) + \frac{\partial}{\partial x_1} H^{(m)}(t; x_1, x_2) D_1^{(m)} + \frac{\partial}{\partial x_2} H^{(m)}(t; x_1, x_2) D_2^{(m)} = H^{(m)}(t; x_1, x_2) Q,
$$

with the following boundary conditions:

$$\begin{aligned}
H_{ij}^{(1)}(t;0,x_2) &= 0 \quad \text{for all } i,j \\
H_{ij}^{(2)}(t;\infty,x_2) &= H_{ij}^{(2)}(t;a,x_2) \quad \text{for all } i,j \\
H_{i1}^{(m)}(t;x_1,0) &= 0 \quad \text{for all } i,m \\
H_{1j}^{(2)}(t;0,x_2) &= 0 \quad \text{for all } j
\end{aligned}$$

Unlike Zhang [29], here we cannot re-write the differential equation in steady-state by dropping the ∂t term. This is because here the corresponding processes are not ergodic (the $m = 1$ equation would render $X_2 \to \infty$ and viceversa). Our only hope is to analyze the transient case and compute the 3-dimensional Laplace transform

$$F(s;w_1,w_2) = \int_0^\infty \int_0^\infty \int_0^\infty e^{-st-w_1x_1-w_2x_2} H^{(m)}(t;x_1,x_2).$$

Then we need to couple the Laplace transform equations with analyzing the Markov Regenerative Process (MRGP) $\{(X_1(t), X_2(t), Z_1(t), Z_2(t), Y(t)), t \geq 0\}$ by considering the scheduler switching instants as Markov-renewal epochs. To combine the Riemann-Hilbert boundary value formulation with the MRGP formulation is analytically intractable. For this reason, we resort to approximation techniques.

Appendix 2: Random Switch-over Times Analysis

In Section 5, we considered the switch-over times to be positive constants. Here we extend the analysis to the case when the switch-over times are non-negative random variables. In particular, let the switch-over times be random with mean θ, cumulative distribution function $F_S(\cdot)$ and LST $\tilde{F}_S(\cdot)$. Assume that for $y < 0$, $F_S(y) = 0$. Let S_{12} be the random time for the scheduler to switch from buffer 1 to buffer 2. Likewise define S_{21}.

Before developing approximations for the marginal distribution of the buffer contents in steady state, we need to make some assumptions regarding the switch-over times:

1. The switch-over time must be much smaller than the “on” times for source 1. In other words, if $U_1 \sim \exp(\alpha_1)$ denotes the “on” times for source 1, then $P(S_{21} > U_1) \approx 0$.

2. The switch-over time must be much smaller than T_2, i.e., $P(S_{12} > T_2) \approx 0$.

The above assumptions are practical since the switch-over time is usually a very tiny amount of time. These assumptions are needed to keep the analysis tractable.

Define $T_1^{adj}(t)$ as the random time for buffer 1 to empty out from a level $a + r_1 t$ while the scheduler is serving that buffer continuously. Note that T_1^{adj} defined in Section 5.2 corresponds to $T_1^{adj}(\theta)$. Figure 11 pictorially depicts the buffer and scheduler dynamics that will be useful in our analysis.

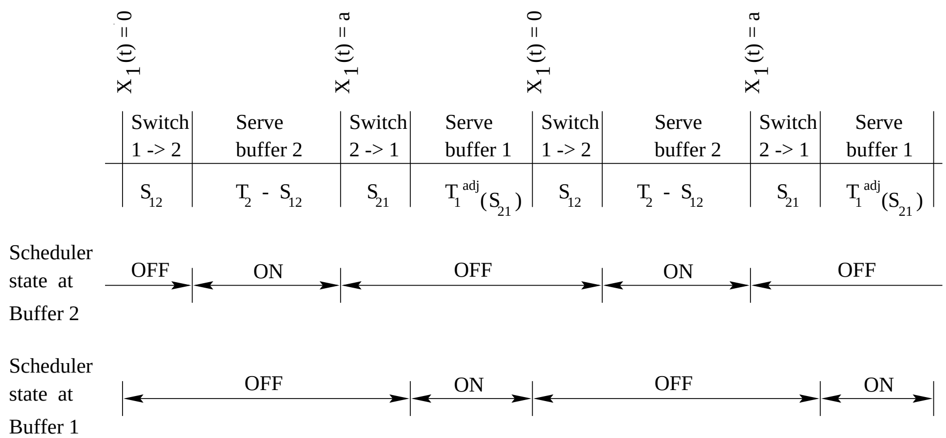

Figure 11: Scheduler Dynamics

Buffer 2 approximations

The analysis proceeds identical to that in Section 5.2. Consider the compensating source defined in Section 4.2.1. The compensating source is such that its on times are $T_1^{adj}(S_{21}) + S_{12} + S_{21}$ and off times are $T_2 - S_{12}$. By conditioning and unconditioning on the switch-over times, we can derive the LSTs of the on and off times respectively as

$$\tilde{O_1}^{adj}(w) = E[e^{-w(T_1^{adj}(S_{21})+S_{12}+S_{21})}] = \begin{cases} \tilde{F_S}(w)\frac{w+\beta_1+cs_0(w)}{\beta_1}e^{a\ s_0(w)}\tilde{F_S}(w - r_1 s_0(w)) & \text{if } w \geq w^* \\ \infty & \text{otherwise.} \end{cases}$$

and

$$\tilde{O_2}^{adj}(s) = E[e^{-s(T_2-S_{12})}] = \tilde{F_S}(-s)\frac{\beta_1}{\beta_1 + s}e^{-a\frac{\alpha_1 s+\beta_1 s+s^2}{r_1 s+r_1\beta_1}}.$$

The effective bandwidth of the compensating source is the unique solution to

$$\tilde{O_1}^{adj}(v\, eb_0^{adj}(v) - cv)\ \tilde{O_2}^{adj}(v\, eb_0^{adj}(v)) = 1.$$

The next step is to solve for η in

$$eb_0^{adj}(\eta) + eb_2(\eta) = c.$$

Using Theorem 5 we can derive the distribution of the contents in buffer 2 as

$$LB^{adj} \leq \lim_{t\to\infty} P\{X_2(t) > x\} \leq UB^{adj},$$

where

$$LB^{adj} = \frac{r_2\beta_2}{eb_2(\eta)(\alpha_2+\beta_2)} \frac{1-\tilde{O_2}^{adj}(\eta eb_0^{adj}(\eta))}{E(T_1^{adj}(S_{21})+S_{12}+S_{21})+E(T_2-S_{12})} \frac{c}{eb_0^{adj}(\eta)(c-eb_0^{adj}(\eta))\eta} e^{-\eta x},$$

and

$$UB^{adj} = \frac{r_2\beta_2}{eb_2(\eta)(\alpha_2+\beta_2)} \frac{1/\tilde{O_2}^{adj}(\eta eb_0^{adj}(\eta))-1}{E(T_1^{adj}(S_{21})+S_{12}+S_{21})+E(T_2-S_{12})} \frac{c}{eb_0^{adj}(\eta)(c-eb_0^{adj}(\eta))\eta} e^{-\eta x}.$$

Note that $E(T_1^{adj}(S_{21})+S_{12}+S_{21}) = E[T_1^{adj}(\theta)]+2\theta$ and $E(T_2-S_{12}) = E(T_2)-\theta$, where from Equation (20) we have $E[T_1^{adj}(\theta)] = \frac{r_1+(a+r_1\theta)(\alpha_1+\beta_1)}{c\alpha_1+c\beta_1-r_1\beta_1}$ and $E(T_2)$ is given in Equation (21).

Buffer 1 approximations

Proceeding in a similar fashion as in Section 5.3 we can show that

$$\lim_{t\to\infty} P\{X_1(t)>x\} \approx \frac{\frac{r_1\beta_1}{c(\alpha_1+\beta_1)} E[V_2^{adj}] e^{-\eta(x-a)} \tilde{F}_S(-\eta r_1)}{E[V_1^{adj}]+E[V_2^{adj}]+E[V_3^{adj}]},$$

where η is defined in Equation (27), and the terms $E[V_1^{adj}]$, $E[V_2^{adj}]$ and $E[V_3^{adj}]$ are defined in Section 5.3.

Acknowledgements

This work was partially supported by DARPA (Grant: MDA 972-01-1-0038). The authors thank the editor and the anonymous reviewers for their comments and suggestions that led to considerable improvements in the content and presentation of this paper.

References

[1] D. Anick, D. Mitra and M. M. Sondhi. *Stochastic Theory of a Data Handling System with Multiple Sources.* Bell System Tech. Journal, 61, 1871–1894, 1982.

[2] D. Artiges and P. Nain. *Upper and Lower Bounds for the Multiplexing of Multiclass Markovian On/Off Sources.* Performance Evaluation, 27 & 28, 673–698, 1996.

[3] A.A. Borovkov. *Stochastic Processes in Queueing Theory.* Springer, New York, 1976.

[4] G. L. Choudhury, D. M. Lucantoni and W. Whitt. *On the Effectiveness of Effective Bandwidths for Admission Control in ATM Networks.* Proc. of ITC-14, 411–420, Elsevier Science, 1994.

[5] J. W. Cohen and O. J. Boxma. *The $M/G/1$ Queue with Alternating Service Formulated as a Riemann-Hilbert Problem.* Performance '81, Proc. of the 8th Intl. Symp., ed. F.J. Kylstra, pp 181-199, North-Holland, Amsterdam, 1981.

[6] J. W. Cohen and O. J. Boxma. *Boundary Value Problems in Queueing System Analysis.* Mathematical Studies, 79, North-Holland, Amsterdam, 1983.

[7] O. J. Boxma. *Two Symmetric Queues with Alternating Service and Switching Times.* Performance '84, Proc. of the 10th Intl. Symp. on Models of Computer System Performance, ed. E. Gelenbe, pp 409-434, North-Holland, Amsterdam, 1984.

[8] C. F. Daganzo. *Some Properties Of Polling Systems.* Queueing Systems, Theory and Applications, 6, 137-154, 1990.

[9] M. Darce, K.D. Glazebrook and J. Nino-Mora. *The Achievable Region Approach to the Optimal Control of Stochastic Systems (with discussion).* J. Roy. Statist. Soc., B61, 747-791, 1999.

[10] A.I. Elwalid and D. Mitra. *Analysis and Design of Rate-Based Congestion Control of High Speed Networks, part I: Stochastic Fluid Models, Access Regulation.* Queueing Systems, Theory and Applications, Vol.9, 29–64, 1991.

[11] A.I. Elwalid and D. Mitra. *Effective Bandwidth of General Markovian Traffic Sources and Admission Control of High-speed Networks.* IEEE/ACM Transactions on Networking, 1(3), 329–343, June 1993.

[12] A.I. Elwalid and D. Mitra. *Analysis, Approximations and Admission Control of a Multi-service Multiplexing System with Priorities.* INFOCOM'95, 463–472, 1995.

[13] N. Gautam, V. G. Kulkarni, Z. Palmowski and T. Rolski. *Bounds for Fluid Models Driven by Semi-Markov Inputs.* Probab. in the Engr. and Informational Sciences, 13, 429–475, 1999.

[14] N. Gautam and V. G. Kulkarni. *Applications of SMP Bounds to Multi-Class Traffic in High-Speed Networks.* Queuing Systems: Theory and Applications, 351-380, Vol. 36, # 4, 2000.

[15] R. J. Gibbens and P.J. Hunt. *Effective bandwidths for the Multi-type UAS Channel.* Queueing Systems, Theory and Applications, 9, 17–28, 1991.

[16] G. Kesidis, J. Walrand, and C.S. Chang. *Effective bandwidths for Multiclass Markov Fluids and Other ATM sources.* IEEE/ACM Transactions on Networking, 1(4), 424–428, 1993.

[17] G. Kesidis. *ATM Network Performance.* Kluwer Academic Publishers, 1996.

[18] A.E. Kostin, I. Aybay and G. Oz. *A Randomized Contention-Based Load-Balancing Protocol for a Distributed Multiserver Queuing System.* IEEE Trans. on Parallel and Distributed Systems, Vol. 11, No. 12, Dec. 2000.

[19] V. G. Kulkarni, and T. Rolski. *Fluid Model Driven by an Ornstein-Ühlenbeck Process.* Probability in Engineering and Information Sciences, 8, 403–417, 1994.

[20] V. G. Kulkarni. *Modeling and Analysis of Stochastic Systems.* Texts in Statistical Science Series. Chapman and Hall, Ltd., London, 1995.

[21] V. G. Kulkarni. *Effective Bandwidths for Markov Regenerative Sources.* Queueing Systems, Theory and Applications, 24, 137–153, 1997.

[22] R. Mirchandany, D. Towsley and J.A. Stankovic. *Adaptive Load Sharing in Heterogeneous Distributed Systems.* Journal of Parallel and Distributed Computing, 331-346, Sept. 1990.

[23] A. Narayanan and V. G. Kulkarni. *First Passage Times in Fluid Models with an Application to Two-priority Fluid Systems.* Proceedings of the IEEE International Computer Performance and Dependability Symposium, 1996.

[24] Z. Palmowski and T. Rolski. *A Note on Martingale Inequalities For Fluid Models.* Statist. Probab. Lett., 31, No. 1, 13–21, 1996.

[25] H. Takagi. *Analysis Of Polling Systems.* MIT Press, Cambridge, MA, 1986.

[26] A.N. Tantawi and D. Towsley. *Optimal Static Load Balancing in Distributed Computer Systems.* J. ACM, vol 32, no. 2, 445-465, Apr. 1985.

[27] W. Whitt. *Tail Probabilities with Statistical Multiplexing and Effective Bandwidths for Multiclass Queues.* Telecommun. Syst., 2, 71–107, 1993.

[28] P. Whittle. *Restless Bandits: Activity Allocation in a Changing World.* J. Appl. Prob., A25, 287-298, 1988.

[29] J. Zhang. *Performance Study of Markov-modulated Fluid Flow Models with Priority Traffic.* INFOCOM'93, 10–17, 1993.