Backpressure-based Control Protocols: Design and Computational Aspects

D.I. Miretskiy
University of Twente
Enschede, The Netherlands
Email: d.miretskiy@math.utwente.nl

W.R.W. Scheinhardt
University of Twente
Enschede, The Netherlands
Email: w.r.w.scheinhardt@utwente.nl

M.R.H. Mandjes
University of Amsterdam
Amsterdam, The Netherlands
Email: mmandjes@science.uva.nl

***Abstract*—Congestion control in packet-based networks is often realized by feedback protocols. In this paper we assess their performance under a back-pressure mechanism that has been proposed and standardized for Ethernet metropolitan networks. In such a mechanism the service rate of an upstream queue is reduced when the downstream queue is congested, in order to protect the downstream queue. We study a Markovian model that captures the essentials of the protocol, but at the same time allows for numerical analysis. We first derive explicit results for the stability condition of the model (which turns out to be nontrivial). Then we present logarithmic estimates of the probability of buffer overflow in the second queue, which are subsequentially used when devising an efficient simulation procedure based on importance sampling. We conclude the paper by presenting a number of numerical results, and some general design guidelines.**

I. Introduction

A variety of protocols has been proposed for purposes of congestion control in packet networks. A well-known mechanism is *random early detection*, as proposed in e.g. [4], which has been intensively studied. Here, packets are randomly dropped when the buffer content exceeds some threshold, so as to notify the user about incipient congestion.

Similar feedback-based mechanisms have been proposed and standardized for congestion control in Ethernet metropolitan networks. The backpressure scheme defined in IEEE 802.3x [11], is intended to provide flow control on a hop-by-hop basis. For ease restricting ourselves to a two-node network (but the principle applies to networks of any size), it works roughly as follows. Traffic enters the first queue ('upstream queue'), and after being served it is routed to a second queue ('downstream queue'), where it is served and leaves the network. When the buffer content of the downstream queue exceeds some threshold, the service rate at the first queue is reduced, so as to protect the downstream queue. The upstream queue resumes serving at its orginal rate when the buffer content of the downstream queue is below the threshold again.

It is clear that it is of crucial importance to find simple, yet adequate design rules that determine suitable values for the buffer thresholds and service rates. A problem is that there are hardly any explicit results known for models of this type. Even under Markovian assumptions, the buffer content distribution cannot be given in closed form – just rough characteristics are known, see e.g. [13]. The present paper presents performance evaluation techniques that enable numerical analysis, and in addition we present a number of insights that are useful when designing such a backpressure system.

More specifically, our contributions are the following:

A. We determine the stability condition of the backpressure-based model. In other words, we find a condition on the arrival rate, service rates, and threshold value such that the queue does not systematically fill. This condition is rather surprising (and very different from the common '$\varrho < 1$ conditions').

B. Main performance metrics are the (end-to-end) delay, the throughput, and the probability of overflow in the second queue (as that was the queue that was intended to be protected). From a computational point of view, the latter metric is the most challenging. It is noted that it is not straightforward to estimate this probability through simulation, because of the rare event involved. We propose an advanced simulation algorithm, that manages to reliably estimate the probability of interest (which is the probability of overflow in the downstream queue before the system becomes empty, starting off from any given state) in a reasonable amount of time. This simulation procedure is based on importance sampling (IS). The method borrows elements from algorithms developed earlier for related (tandem) systems [2], [7].

The main idea behind IS is to simulate a system under a new probability measure which guarantees more frequent occurrence of the rare event of interest. To obtain an unbiased estimate, the output of the simulations is corrected by so-called likelihood ratios. The challenge is to construct a 'good' new measure. An often-used notion in this respect is that of asymptotic efficiency (or: asymptotic optimality) which essentially means that the variance of the estimator behaves approximately as the square of its first moment. When this is not the case, the estimator may even have infinite variance. We refer to [5] for more background on IS.

In more detail, the contribution is that we present a simple and efficient IS implementation for simulating the overflow probability in the slow-down model. On the one hand it is as easy to implement as the scheme in [2], or even as that in [7], but on the other hand it performs comparably in terms of computational demand. On the other hand it allows any given starting state, while extensive numerical experimentation

shows that it inherits the efficiency properties of [9]. In the analysis it turns out that three regimes (in terms of the decay rates of the queues) need to be distinguished.

C. The last part of the paper is devoted to a study of the performance aspects of our system. In particular, we quantify the effect of the value of the threshold. It clearly affects the mean end-to-end delay as well as the overflow probability of the downstream queue, and it is not clear *a priori* what value is optimal. Therefore we construct a cost function (including both performance metrics), and use it to develop a procedure that is capable of finding the optimal threshold value. The numerical values for the overflow probability are obtained by using the IS algorithm mentioned under item B.

This paper is structured as follows. The next section contains a model description, the stability result, and background on IS. In Section III we propose an IS scheme, and argue why it has attractive efficiency properties. In Section IV we demonstrate how to develop procedures for setting the parameter values (focusing on the value of the threshold), and we provide supporting numerical results. We draw conclusions in Section V.

II. Model and Preliminaries

A. Model and Stability

In this section we describe the system, which is a series of two stations (or: *queues*), in detail. Jobs enter the system at the upstream queue according to a Poisson process with rate λ, and after being served they are forwarded to the downstream queue; after service in this second queue, they leave the network. Service times at the second station are exponential with parameter μ_2 all the time, but the service speed at the first queue depends on the content of the second queue. Normally, service times at the first station are exponential with parameter μ_1, but if the number of jobs in the second queue is larger than some prespecified threshold m – the so-called *slow-down threshold* – then the service times are exponential with parameter μ_1^+, where $\mu_1^+ < \mu_1$. When the system 'stabilizes' and the number of jobs in the second queue is again strictly below the slow-down threshold, the rate of the first station returns to its original value μ_1.

For convenience we choose the parameters such that $\lambda + \mu_1 + \mu_2 = 1$, without loss of generality (in fact, we 'rescale our time unit'), and hence $\lambda + \mu_1^+ + \mu_2 < 1$. As in [9], we assume (as an approximation) the waiting rooms at both stations to be infinitely large. We define the discrete-time joint queue-length process $Q_j = (Q_{1,j}, Q_{2,j})$, where $Q_{i,j}$ is the number of jobs at node i after the j-th transition. We define the possible jump directions of the process Q_j via vectors $v_0 = (1,0)$, $v_1 = (-1,1)$ and $v_2 = (0,-1)$ with corresponding jump rates λ, μ_1 (or μ_1^+) and μ_2 (where it is noted that v_k is impossible if queue k is empty – then take $v_k = (0,0)$ instead).

A first question is under what condition this process is stable – clearly for design purposes such a criterion is of crucial importance. Interestingly, the resulting criterion is substantially more involved than the usual '$\varrho < 1$ conditions'. Define $\psi := \mu_1/\mu_2$ and $\psi^+ := \mu_1^+/\mu_2$.

Theorem 2.1: Case I: $\mu_1^+ < \mu_2$. The network is stable if

$$\lambda < \frac{\mu_1(1-\psi^m)(1-\psi^+) + \mu_1^+\psi^m(1-\psi)}{(1-\psi^m)(1-\psi^+) + \psi^m(1-\psi)}.$$

Case II: $\mu_1^+ \geq \mu_2$. The network is stable if $\lambda < \mu_2$.

Proof: It is obvious that $\lambda < \mu_2$ is necessary, but not sufficient for stability. We deal with both cases separately.

Case I: $\mu_1^+ < \mu_2$. The proof relies on techniques from the theory of Quasi-Birth-and-Death (QBD) processes [6]; we aim to prove the positive recurrence of the discrete-time process Q_j. From now on we will treat Q_j as a discrete-time QBD with $Q_{1,j}$ and $Q_{2,j}$ being the level and the phase, respectively; note that this is not a 'standard QBD', as the number of phases is infinite.

We now introduce some QBD related notation. Let M_0, M_1 and M_2 be $(n+1)\times(n+1)$ dimensional matrices, with n being the number of phases (either finite or infinite). M_0 represents an increase in level (new job arrives to the system), M_1 no change in level (job leaves the system) and M_2 an decrease in level (job is forwarded from the first queue to the second one); the precise definitions of these matrices were given in [13] for our model. [6, Thm. 7.2.3] now states that if the number of phases is *finite*, then the QBD is positive recurrent if $\boldsymbol{\alpha} M_0 \boldsymbol{e} < \boldsymbol{\alpha} M_2 \boldsymbol{e}$, where the vector $\boldsymbol{\alpha}$ is the solution to $\boldsymbol{\alpha} M = 0$, with $M := M_0 + M_1 + M_2$; $\boldsymbol{e}$ is an all-1 vector.

Application of this result (which is not legitimate in our case, due to the infinite number of phases!) would yield that the QBD is positive recurrent if

$$\sum_{i=0}^{m-1} \alpha_0 \psi^i (\lambda - \mu_1) + \sum_{i=m}^{\infty} \alpha_0 \psi^m (\psi^+)^{i-m} (\lambda - \mu_1^+) < 0, \quad (1)$$

where

$$\alpha_i = \begin{cases} \alpha_0 \psi^i, & i < m \\ \alpha_0 \psi^m (\psi^+)^{i-m}, & i \geq m \end{cases};$$

$$\alpha_0 = \left(\sum_{i=0}^{m-1} \psi^i + \psi^m \sum_{i=m}^{\infty} (\psi^+)^{i-m} \right)^{-1},$$

from which the first statement of the theorem would follow. There is a counterpart of [6, Thm. 7.2.3] that *does* deal with an infinite number of phases, though: [12, Thm. 5] states that the QBD with infinite number of phases is positive recurrent if $\boldsymbol{\alpha} M_0 \boldsymbol{e} < \boldsymbol{\alpha} M_2 \boldsymbol{e}$ provided that $\bar{M} = M_1 + M_2$. Here $\bar{M}$ is an infinite-dimensional matrix that describes the behavior of the phase-process of the QBD at level 0, see again [13] for its precise form. The condition $\bar{M} = M_1 + M_2$ effectively means that the phase process in level 0 is the same as for any other level. Obviously, this requirement fails in our case. In order to be able to apply [12, Thm. 5] we modify the QBD Q_j in order to satisfy the condition condition $\bar{M} = M_1 + M_2$:

$$v_1 = (0,1), \text{ when } Q_{1,j} = 0,$$

where the other transition vectors remain unchanged. Now we can conclude that this modified process is stable if (1) holds.

An elementary inspection yields that the cycle time (i.e., the number of transitions it takes the discrete-time process to reach the origin) of the modified QBD is *stochastically larger* than the cycle time of the original QBD, and hence stability of the original QBD is implied by the stability of the modified QBD.

Case II: $\mu_2 \leq \mu_1^+$. Here we cannot apply the reasoning mentioned above since the distribution α_i does not exist when $\mu_2 \leq \mu_1^+$. However, in this case, stability can be proven rather straightforwardly. Clearly, the expected cycle length in this case can be bounded from above by the expected cycle length for $m = 0$, due to elementary coupling arguments. The latter value is finite, since it corresponds to the mean busy cycle length of the two-node Jackson tandem with parameters $(\lambda, \mu_1^+, \mu_2)$, which is stable under $\lambda < \mu_2 \leq \mu_1^+$. ■

We remark that somewhat related results for the slow-down system with a *finite* second buffer were reported in [13, Thm. 15]. Also, interestingly, the slow-down system can be stable *even when* $\lambda > \mu_1^+$*!* The intuition behind this is as follows. Consider the case when both $\lambda > \mu_1^+$ and the condition in Theorem 2.1 hold true. The content of the first queue typically increases when the number of jobs in the second queue is above m. However, it stays finite because the content of the second queue tends to decrease and the system returns to its normal state in which the number of jobs in the first queue tends to decrease. It is also worthwhile to mention that when the slow down threshold m is 0 or ∞, the criterion mentioned above reduces to the standard stability condition for the ordinary two-node Jackson tandem network.

B. State Space and System Dynamics

In Section III we focus on estimating the probability of reaching some high level B in the second queue before it returns to the origin, starting from any given state; this subsection describes a number of notions that are useful with this goal in mind. From now on we let the threshold m scale linearly with B that is, $m \equiv \theta B$ for some $\theta \in (0, 1)$. In terms of the scaled process $(X_1, X_2) = (Q_1/B, Q_2/B)$, we analyze the probability that its second coordinate attains the value 1 before reaching the origin. Note that an advantage of this scaling is that, in order to analyze this probability, we may use the state space $[0, \infty) \times [0, 1]$ (for any value of B!). We introduce the following subsets of the state space, with $x := (x_1, x_2)$:

$$
\begin{aligned}
D &:= \{x : x_1 > 0, 0 < x_2 < \theta\}, & \partial_2 &:= \{(x_1, 0) : x_1 > 0\},\\
D^+ &:= \{x : x_1 > 0, \theta \leq x_2 < 1\}, & \partial_\theta &:= \{(x_1, \theta) : x_1 \geq 0\},\\
\partial_1^+ &:= \{(0, x_2) : x_2 \in [\theta, 1)\}, & \partial_e &:= \{(x_1, 1) : x_1 \geq 0\},\\
& & \partial_1 &:= \{(0, x_2) : x_2 > 0\}.
\end{aligned}
$$

The full state space is $\bar{D} \cup \bar{D}^+$, where $\bar{D} := D \cup \partial_\theta \cup (\partial_1 \setminus \partial_1^+) \cup \partial_2$ and $\bar{D}^+ := D^+ \cup \partial_e \cup \partial_1^+ \cup \partial_\theta$. Recall that the transition v_k is impossible when queue k is empty, i.e., when $X_j \in \partial_k$. We modify the process X_j to deal with this by allowing 'self-loop transitions' in the following way (see also Figure 1): for $k = 1, 2$,

$$
\begin{aligned}
&\mathbb{P}(X_{j+1} = X_j | X_j \in \partial_k \setminus \partial_1^+) = \mu_k, \\
&\mathbb{P}(X_{j+1} = X_j | X_j \in \partial_1^+) = \mu_1^+ / (\lambda + \mu_1^+ + \mu_2).
\end{aligned} \tag{2}
$$

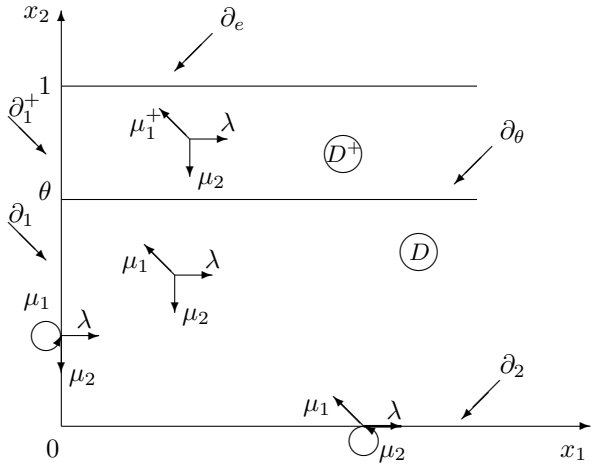

Fig. 1. State space and transition structure for the scaled process X_j.

Next, we introduce the stopping time τ_B^s, which is the first time that the process X_j hits level 1, starting from state $s = (s_1, s_2)$, without visits to the origin: with $X_0 = s$,

$$
\tau_B^s = \inf\{k : X_k \in \partial_e, X_j \neq 0 \text{ for } j = 1, \ldots, k-1\}, \tag{3}
$$

and we define $\tau_B^s = \infty$ if X_j hits the origin before ∂_e. We let $I(A_B^s)$ be the indicator of the event $\{\tau_B^s < \infty\}$ for the scaled sample path $(X_j, j = 0, \ldots \ : X_0 = s)$. Thus we can write the probability of our interest as

$$
p_B^s = \mathbb{E} I(A_B^s) = \mathbb{P}(\tau_B^s < \infty). \tag{4}
$$

It is clear that estimating the probability p_B^s through direct, naïve, simulations is not feasible when B grows large, because of the rarity of the event involved. We therefore have to use some alternative techniques to obtain a reliable estimator. In Section III we focus on importance sampling, which we will now describe briefly.

C. Background on Importance Sampling

To estimate p_B^s, IS generates samples under a new probability measure $\mathbb{Q}$, with respect to which $\mathbb{P}$ is absolutely continuous. The probability p_B^s can, in self-evident notation, now alternatively be expressed as

$$
p_B^s = \mathbb{E}[I] = \mathbb{E}^{\mathbb{Q}}[LI], \tag{5}
$$

where I is an indicator function and L is the likelihood ratio of a realization ('path') ω:

$$
L = \frac{\mathrm{d}\mathbb{P}}{\mathrm{d}\mathbb{Q}}(\omega). \tag{6}
$$

After n replications we obtain a family of observations $(L_i, I_i), i = 1, \ldots, n$ and are able to construct the unbiased

estimator of p_B^s by $n^{-1} \cdot \sum_{i=1}^n L_i I_i$. We conclude this subsection by recalling that an IS scheme (corresponding to measure Q) is said to be *asymptotically efficient* if

$$\liminf_{B\to\infty} \frac{\log \mathbb{E}^{\mathbb{Q}}[L^2 I]}{\log \mathbb{E}^{\mathbb{Q}}[LI]} \geq 2. \tag{7}$$

If the probability of $\{A_B^s\}$ decays exponentially in B, i.e., $B^{-1} \log p_B^s$ tends (for B large) to some value in $(-\infty, 0)$, then, due to (5) and $\mathbb{E}^{\mathbb{Q}}[L^2 I] = \mathbb{E}[LI]$, (7) simplifies to

$$\limsup_{B\to\infty} \frac{1}{B} \log \mathbb{E}[LI] \leq 2 \lim_{B\to\infty} \frac{1}{B} \log p_B^s. \tag{8}$$

III. Importance Sampling

In [9] we have developed an asymptotically efficient IS-based method for estimating p_B^s, but this has, from a practical point of view, important drawbacks: the new measure is state-dependent, and needs to be recomputed at every transition (amounting to jointly solving two cubic equations), thus severely limiting the efficiency gain. In this section we present a new measure that is still state-dependent, but its computation is substantially less demanding, as it requires just a few cubic equations to be solved. As we will see, the speed-up of this new scheme is still substantial.

We give a detailed description of the IS scheme for the case $\mu_2 < \mu_1^+ < \mu_1$; for the other cases (i.e., $\mu_1^+ < \mu_2 \leq \mu_1$ and $\mu_1^+ \leq \mu_1 < \mu_2$), we just present the results. Throughout this section we fix the starting state s and assume it is situated below the slow-down threshold, i.e., $s \in \bar{D}$, which is evidently the most interesting case.

A. IS Scheme for Case $\mu_2 < \mu_1^+ < \mu_1$. At first recall from [9] the most probable path to overflow and the pair of new measures $(\tilde{\lambda}, \tilde{\mu}_1, \tilde{\mu}_2)$ and $(\tilde{\lambda}^+, \tilde{\mu}_1^+, \tilde{\mu}_2^+)$ that will ensure that any sample path will follow the optimal trajectory with high probability. To this end we assign a 'cost' to any path. Minimizing this we obtain the optimal trajectory and corresponding new measure, see [9]. To ease the exposition on the new measures we divide the state space as it is shown in Figure 2, where $\alpha_1 := (\mu_1 - \mu_2)/(\mu_1 - \lambda)$ and $\alpha_1^+ := (\mu_1^+ - \mu_2)/(\mu_1^+ - \lambda)$. This figure also provides some examples of the most probable overflow trajectories. We are particularly interested in the partition of the bottom part of the state space, i.e., in A_1, A_2 and A_3. The new measures for $s \in A_1 \cup A_3$ are not difficult, as we will see in (17). However, to find the optimal new measure for $s \in A_2$ one first needs to jointly solve

$$\begin{cases} \lambda^{(\text{line})} = \mu_1^{(\text{line})} + \frac{\kappa^{(\text{line})} - s_1}{\theta - s_2}(\mu_1^{(\text{line})} - \mu_2^{(\text{line})}) \\ \lambda^{(\text{line})} + \mu_1^{(\text{line})} + \mu_2^{(\text{line})} = \lambda + \mu_1 + \mu_2 \\ \lambda^{(\text{line})} \mu_1^{(\text{line})} \mu_2^{(\text{line})} = \lambda \mu_1 \mu_2 \\ \lambda^{(\text{line})} \leq \mu_1^{(\text{line})} \text{ and } \mu_1^{(\text{line})} > \mu_2^{(\text{line})} \\ \lambda^{(\text{line})}, \mu_1^{(\text{line})}, \mu_2^{(\text{line})} > 0 \end{cases} \tag{9}$$

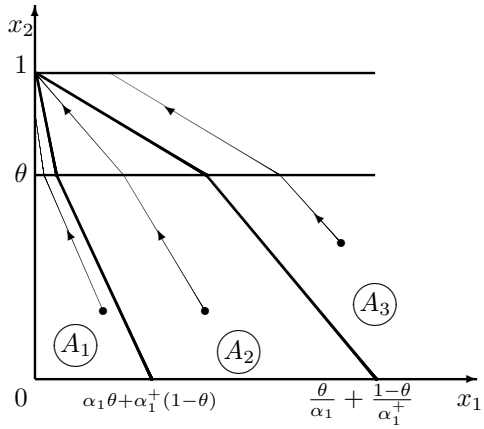

Fig. 2. Partition of $\bar{D} \cup \bar{D}^+$ when $\mu_2 < \mu_1^+ < \mu_1$.

and

$$\begin{cases} \lambda^{+(\text{line})} = \mu_1^{+(\text{line})} - \frac{\kappa^{(\text{line})}}{1-\theta}(\mu_1^{+(\text{line})} - \mu_2^{+(\text{line})}) \\ \lambda^{+(\text{line})} + \mu_1^{+(\text{line})} + \mu_2^{+(\text{line})} = \lambda + \mu_1^+ + \mu_2 \\ \lambda^{+(\text{line})} \mu_1^{+(\text{line})} \mu_2^{+(\text{line})} = \lambda \mu_1^+ \mu_2 \\ \lambda^{+(\text{line})} \leq \mu_1^{+(\text{line})} \text{ and } \mu_1^{+(\text{line})} > \mu_2^{+(\text{line})} \\ \lambda^{+(\text{line})}, \mu_1^{+(\text{line})}, \mu_2^{+(\text{line})} > 0 \end{cases} \tag{10}$$

with the condition

$$\kappa^{(\text{line})} := s_1 - \frac{\mu_1^{(\text{line})} - \lambda^{(\text{line})}}{\mu_1^{(\text{line})} - \mu_2^{(\text{line})}}(\theta - s_2) = \frac{\mu_1^{+(\text{line})} - \lambda^{+(\text{line})}}{\mu_1^{+(\text{line})} - \mu_2^{+(\text{line})}}(1 - \theta). \tag{11}$$

It can be verified that this amounts to solving two coupled cubic equations. The superscripts (line) and +(line) indicate that the solution is in fact the optimal change of measure to reach the state $(0, 1)$ following a concatenation of two straight line starting in s, with intersection in $(\kappa^{(\text{line})}, \theta)$.

Now we can define the (overall) optimal new measures below and above the slow-down threshold, which depend only on the starting state s. The new measure below the slow-down threshold, as given through $(\tilde{\lambda}, \tilde{\mu}_1, \tilde{\mu}_2)$, is

$$(\tilde{\lambda}, \tilde{\mu}_1, \tilde{\mu}_2) = \begin{cases} (\mu_2, \mu_1, \lambda) & \text{if } s \in A_1, \\ (\lambda^{(\text{line})}, \mu_1^{(\text{line})}, \mu_2^{(\text{line})}) & \text{if } s \in A_2, \\ (\lambda, \mu_1, \mu_2) & \text{if } s \in A_3. \end{cases} \tag{12}$$

Above the slow-down threshold the new measure, given through $(\tilde{\lambda}^+, \tilde{\mu}_1^+, \tilde{\mu}_2^+)$, is

$$(\tilde{\lambda}^+, \tilde{\mu}_1^+, \tilde{\mu}_2^+) = \begin{cases} (\mu_2, \mu_1^+, \lambda) & \text{if } s \in A_1, \\ (\lambda^{+(\text{line})}, \mu_1^{+(\text{line})}, \mu_2^{+(\text{line})}) & \text{if } s \in A_2, \\ (\lambda, \mu_1^+, \mu_2) & \text{if } s \in A_3. \end{cases} \tag{13}$$

Now, let us define $\gamma^s(x)$ to be the residual cost of moving from state x to ∂_e along the path to overflow that started in s:

$$\gamma^s(x) := \begin{cases} \gamma_1^s(x) + \gamma_2^s(\kappa^\star, \theta) & \text{if } x \in \bar{D}, \\ \gamma_2^s(x) & \text{if } x \in \bar{D}^+, \end{cases} \tag{14}$$

with

$$\gamma_1^s(x) := -(x_1 - \kappa^\star) \log \frac{\tilde{\lambda}}{\lambda} - (\theta - x_2) \log \frac{\tilde{\mu}_2}{\mu_2}, \quad \text{if } x \in \bar{D} \tag{15}$$

being the minimal cost of the bottom part of the path to overflow and

$$\gamma_2^s(x) := -x_1 \log \frac{\tilde{\lambda}^+}{\lambda} - (1-x_2)\log\frac{\tilde{\mu}_2^+}{\mu_2}, \quad \text{if } x \in \bar{D}^+ \quad (16)$$

being the minimal cost of the top part of the optimal path to overflow; the optimal crossing state $(\kappa^\star, \theta)$ is as follows

$$\kappa^\star := \begin{cases} \max(0, s_1 - \alpha_1(\theta - s_2)) & \text{if } s \in A_1, \\ \kappa^{(\text{line})} & \text{if } s \in A_2, \\ s_1 - (\theta - s_2)/\alpha_1 & \text{if } s \in A_3. \end{cases} \quad (17)$$

Note that $(\tilde{\lambda}, \tilde{\mu}_1, \tilde{\mu}_2)$, $(\tilde{\lambda}^+, \tilde{\mu}_1^+, \tilde{\mu}_2^+)$ and $\kappa^\star$ (given by (12), (13) and (17) respectively) are fixed, i.e., they only depend on the fixed initial state s, and not on the current state x (as was the case in [9]). It is also important to note that $\gamma^s(s)$, the total cost of moving from the starting state s to the exit boundary ∂_e, equals the decay rate of p_B^s, i.e., $B^{-1} \cdot \log p_B^s \to -\gamma^s(s)$, see Theorem 3.2 in [9].

Notice that the function $\gamma^s(x)$ is piecewise-linear in x, since the new tilde-measure, i.e., $(\tilde{\lambda}, \tilde{\mu}_1, \tilde{\mu}_2)$ and $(\tilde{\lambda}^+, \tilde{\mu}_1^+, \tilde{\mu}_2^+)$, depends only on the fixed initial state s, and not on the current state x. This is the main difference with the new measure studied in [9], where we used the optimal new measure for each current state x with its cost $\gamma^x(x)$. Therefore a pair of cubic equations (corresponding to (9)–(11) with s replaced by x) had to be solved for *each* state x in the sample path. In our current approach, computation of the new measure requires the solution of (9)–(11) just once.

It is known, e.g. from our previous research [7], that the new measures $(\tilde{\lambda}, \tilde{\mu}_1, \tilde{\mu}_2)$ and $(\tilde{\lambda}^+, \tilde{\mu}_1^+, \tilde{\mu}_2^+)$, which make a sample path 'on average' follow the optimal trajectory to the rare set, is not necessarily asymptotically efficient; this is due to the possibility of several visits to the horizontal axis, which inflate the likelihood ratio, cf. [1], [10]. In order to resolve this, we first introduce the 'hat'-measure $(\hat{\lambda}, \hat{\mu}_1, \hat{\mu}_2)$ and $(\hat{\lambda}^+, \hat{\mu}_1^+, \hat{\mu}_2^+)$ as in [3], to be used when the current state is on or near the horizontal axis, through

$$(\tilde{\lambda}, \mu_1\lambda/\tilde{\lambda}, \mu_2) \text{ and } (\tilde{\lambda}^+, \mu_1^+\lambda/\tilde{\lambda}^+, \mu_2^+). \quad (18)$$

The main idea behind it is to make the likelihood ratios of the loops around the horizontal axis not greater than 1 (by ensuring $\hat{\mu}_2 = \mu_2$).

Having introduced the 'tilde-measure' and the 'hat-measure', we are now ready to define the measure $(\bar{\lambda}(x), \bar{\mu}_1(x), \bar{\mu}_2(x))$, of which we will prove asymptotic efficiency, and which is a combination of the two measures defined above and the original measure:

$$\begin{array}{lll} \bar{\lambda}(x) = & \tilde{\lambda}^{\rho_1}\hat{\lambda}^{\rho_2}\lambda^{\rho_3}M(x), & \text{if } x \in \bar{D}, \\ \bar{\mu}_1(x) = & \tilde{\mu}_1^{\rho_1}\hat{\mu}_1^{\rho_2}\mu_1^{\rho_3}M(x), & \text{if } x \in \bar{D}, \\ \bar{\mu}_2(x) = & \tilde{\mu}_2^{\rho_1}\hat{\mu}_2^{\rho_2}\mu_2^{\rho_3}M(x), & \text{if } x \in \bar{D}, \\ \bar{\lambda}^+(x) = & (\tilde{\lambda}^+)^{\rho_1}(\hat{\lambda}^+)^{\rho_2}(\lambda)^{\rho_3}M^+(x), & \text{if } x \in \bar{D}^+, \\ \bar{\mu}_1^+(x) = & (\tilde{\mu}_1^+)^{\rho_1}(\hat{\mu}_1^+)^{\rho_2}(\mu_1^+)^{\rho_3}M^+(x), & \text{if } x \in \bar{D}^+, \\ \bar{\mu}_2^+(x) = & (\tilde{\mu}_2^+)^{\rho_1}(\hat{\mu}_2^+)^{\rho_2}(\mu_2)^{\rho_3}M^+(x), & \text{if } x \in \bar{D}^+, \end{array} \quad (19)$$

where $M(x)$ and $M^+(x)$ are normalization functions, and the $\rho_i(x)$ are *weights*, cf. [2], [3], given by

$$\begin{aligned} \rho_1(x) &= N(x) \cdot e^{-\frac{2\gamma^s(x_1,x_2)-\delta}{\epsilon}}, \\ \rho_2(x) &= N(x) \cdot e^{-\frac{2\gamma^s(x_1, \frac{\delta}{2}\log\frac{\mu_2}{\lambda})-\delta}{\epsilon}}, \\ \rho_3(x) &= N(x) \cdot e^{-\frac{2\gamma^s(0,0)-\delta}{\epsilon}}. \end{aligned} \quad (20)$$

Here $N(x)$ is a normalization function and δ and ϵ are small positive numbers. We mention again that this measure is, albeit state-dependent, of low computational complexity, as it does not require to solve cubic systems for any point along the path (except s). There are interesting relations between this new measure, and the scheme proposed in [2].

B. IS Scheme for Case $\mu_1^+ \le \mu_2 < \mu_1$. Here we present the IS scheme for the case when $\mu_1^+ \le \mu_2 < \mu_1$. At first we provide the pair of new measures $(\tilde{\lambda}, \tilde{\mu}_1, \tilde{\mu}_2)$ and $(\tilde{\lambda}^+, \tilde{\mu}_1^+, \tilde{\mu}_2^+)$ under which virtually any sample path follows the most probable trajectory with high probability. The bottom part of the state space $\bar{D}$ is divided in subspaces B_i as depicted in Fig. 3, where $\alpha_2^+ := (\mu_2 - \mu_1^+)/(\mu_2 - \lambda)$ and α_1 has been introduced in the previous subsection, cf. [9]. The new measure below m, $(\tilde{\lambda}, \tilde{\mu}_1, \tilde{\mu}_2)$, is as follows:

$$(\tilde{\lambda}, \tilde{\mu}_1, \tilde{\mu}_2) = \begin{cases} (\mu_2, \mu_1, \lambda), & \text{if } s \in B_1, \\ (\lambda^{(\text{line})}, \mu_1^{(\text{line})}, \mu_2^{(\text{line})}), & \text{if } s \in B_2, \\ (\lambda, \mu_1, \mu_2), & \text{if } s \in B_3, \end{cases} \quad (21)$$

whereas above m we have

$$(\tilde{\lambda}^+, \tilde{\mu}_1^+, \tilde{\mu}_2^+) = \begin{cases} (\sqrt{\frac{\lambda\mu_1^+}{z^+}}, \sqrt{\frac{\lambda\mu_1^+}{z^+}}, \mu_2 z^+), & \text{if } s \in B_1 \\ (\lambda^{+(\text{line})}, \mu_1^{+(\text{line})}, \mu_2^{+(\text{line})}), & \text{if } s \in B_2, \\ (\lambda, \mu_2, \mu_1^+), & \text{if } s \in B_3, \end{cases} \quad (22)$$

where z^+ is the solution (unique in $(0,1)$) to

$$\lambda + \mu_1^+ + \mu_2(1 - z^+) = 2\sqrt{\frac{\lambda\mu_1^+}{z^+}}. \quad (23)$$

The optimal crossing state $(\kappa^\star, \theta)$ is now given by

$$\kappa^\star := \begin{cases} 0 & \text{if } s \in B_1, \\ \kappa^{(\text{line})} & \text{if } s \in B_2, \\ s_1 - (\theta - s_2)/\alpha_1 & \text{if } s \in B_3. \end{cases} \quad (24)$$

The function $\gamma^s(x)$ is defined by (14) with (15), (16) and (24). The total cost of the path $\gamma^s(s)$ again is the decay rate of p_B^s, i.e., $B^{-1} \cdot \log p_B^s \to -\gamma^s(s)$, see [9, Thm. 3.2]. The new state-dependent measures $(\bar{\lambda}(x), \bar{\mu}_1(x), \bar{\mu}_2(x))$ and $(\bar{\lambda}^+(x), \bar{\mu}_1^+(x), \bar{\mu}_2^+(x))$ are given by (19), where the 'hat'-measures and the weights $\rho_i(x)$ are defined by (18) and (20), respectively.

C. IS Scheme for Case $\mu_1^+ < \mu_1 \le \mu_2$. This case is the most difficult case. We partition the state space as in Fig. 4. To see whether a starting state s belongs to C_1 or C_2, one needs to jointly solve (9)–(11). Then s belongs to C_1 if and only if $f(s) \le 0$, where

$$f(s) = \theta \log \frac{\mu_2}{\lambda} + (\theta - 1)\log q - s_1 \log \frac{\mu_1}{\lambda} - \gamma^s(s),$$

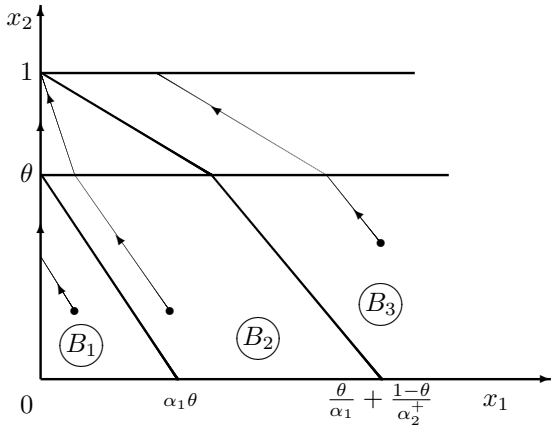

Fig. 3. Partition of $\bar{D}\cup\bar{D}^+$ when $\mu_1^+\le\mu_2<\mu_1$.

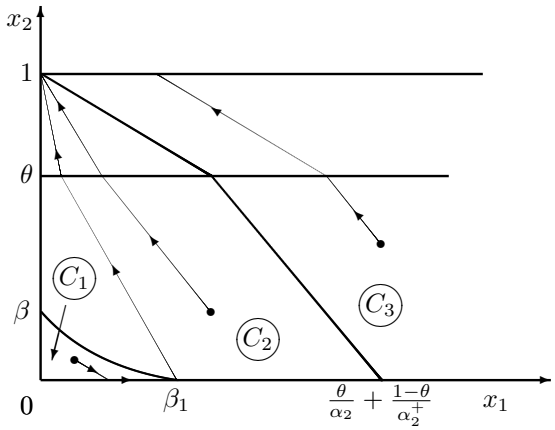

Fig. 4. Partition of $\bar{D}\cup\bar{D}^+$ when $\mu_1^+<\mu_1\le\mu_2$.

with q being the solution (unique in $(0,\mu_1^+\lambda/\mu_1^2)$) to

$$\mu_1\mu_2q^2+\mu_1(\mu_1-\lambda-\mu_1^+-\mu_2)q+\lambda\mu_1^+=0, \qquad (25)$$

see also [7], [9]. The constants β and β_1 are solutions to $f(0,s_2)=0$ and $f(s_1,0)=0$ respectively; $\alpha_2:=(\mu_2-\mu_1)/(\mu_2-\lambda)$.

- At first let us assume that $s\in C_2\cup C_3$. Then we define the new measure below the slow-down threshold

$$(\tilde\lambda,\tilde\mu_1,\tilde\mu_2)=\begin{cases}(\lambda^{(\text{line})},\mu_1^{(\text{line})},\mu_2^{(\text{line})}), & \text{if } s\in C_2,\\ (\lambda,\mu_2,\mu_1), & \text{if } s\in C_3,\end{cases} \qquad (26)$$

and above it

$$(\tilde\lambda^+,\tilde\mu_1^+,\tilde\mu_2^+)=\begin{cases}(\lambda^{+(\text{line})},\mu_1^{+(\text{line})},\mu_2^{+(\text{line})}), & \text{if } s\in C_2,\\ (\lambda,\mu_2,\mu_1^+), & \text{if } s\in C_3,\end{cases} \qquad (27)$$

The optimal crossing state $(\kappa^\star,\theta)$ is now given by

$$\kappa^\star:=\begin{cases}\kappa^{(\text{line})} & \text{if } s\in C_2,\\ s_1-(\theta-s_2)/\alpha_2 & \text{if } s\in C_3.\end{cases} \qquad (28)$$

The function $\gamma^s(x)$ is again defined by (14) with (15), (16) and (28). The total costs of the path, $\gamma^s(s)$, is the decay rate of p_B^s, i.e., $B^{-1}\cdot\log p_B^s\to-\gamma^s(s)$, see [9, Thm. 3.2]. The new measures $(\bar\lambda(x),\bar\mu_1(x),\bar\mu_2(x))$ and $(\bar\lambda^+(x),\bar\mu_1^+(x),\bar\mu_2^+(x))$ are given by (19), where the 'hat'-measures and the weights $\rho_i(x)$ are defined by (18) and (20), respectively.

- Let us now proceed to the case $s\in C_1$. In this case the optimal new measure below the slow-down threshold at the k-th transition is

$$(\tilde\lambda,\tilde\mu_1,\tilde\mu_2)= \qquad (29)$$
$$(\mu_1,\lambda I_{\{k<\tau^\star\}}+\mu_2 I_{\{k\ge\tau^\star\}},\mu_2 I_{\{k<\tau^\star\}}+\lambda I_{\{k\ge\tau^\star\}}),$$

where $\tau^\star:=\min\{k: X_k\in C_3 \text{ and } X_{2,k}=0\}$; see also the optimal trajectory that starts in C_1, Figure 4. Above it the new measure is defined by

$$(\tilde\lambda^+,\tilde\mu_1^+,\tilde\mu_2^+)=(\mu_1,\lambda\mu_1^+/q\mu_1,q\mu_2), \qquad (30)$$

where q is defined by (25).
It is important, that the residual cost of the bottom part of the path to overflow, namely $\gamma_1(x)$, for this case is different from the one defined by (15). Here it is

$$\gamma_1^s(x)=\theta\log\frac{\mu_2}{\lambda}-(x_1-\kappa^\star)\log\frac{\tilde\lambda}{\lambda}+x_2\log\frac{\tilde\mu_2}{\mu_2}, \qquad (31)$$

where

$$\kappa^\star=(1-\theta)(\lambda\mu_1^+-q\mu_1^2)/(\lambda\mu_1^+-q^2\mu_1\mu_2). \qquad (32)$$

For this case $\gamma_2^s(x)$ is defined by (16), and $\gamma^s(x)$ again by (14) with (31), (16) and (32). As before, $\gamma^s(s)$ is the decay rate of p_B^s, i.e., $B^{-1}\cdot\log p_B^s\to-\gamma^s(s)$. The new state-dependent measures $(\bar\lambda(x),\bar\mu_1(x),\bar\mu_2(x))$ and $(\bar\lambda^+(x),\bar\mu_1^+(x),\bar\mu_2^+(x))$ are given by (19), where the 'hat'-measures and the weights $\rho_i(x)$ are defined by (18) and (20) respectively.

For convenience we summarize the resulting IS scheme for the different cases.

- When $\mu_2<\mu_1^+<\mu_1$ one needs to
  1. define the 'primary' new measures (12) and (13);
  2. define 'hat'-measures (18);
  3. define weights $\rho_i(x)$ by (20) based on (14)–(17);
  4. apply (19).
- When $\mu_1^+\le\mu_2<\mu_1$, the same procedure is followed, only replacing the 'primary' new measures (12) and (13) by (21) and (22) in step 1; and (17) by (24) in step 3.
- When $\mu_1^+<\mu_1\le\mu_2$ and $s\in C_2\cup C_3$, we follow the same algorithm, only with the 'primary' new measure replaced by (26) and (27) in step 1; and using (28) instead of (17) in step 3.
- When $\mu_1^+<\mu_1\le\mu_2$ and $s\in C_1$, again the same procedure is followed, this time replacing the 'primary' new measure by that in (29) and (30); we also replace (15) by (31) and (17) by (32) in step 3.

In [9] we proved asymptotic efficiency of the fully state-dependent IS scheme, i.e., when the 'primary' new measure was dependent on the current state. Analyzing the simplified

IS scheme (19) we have to deal with the additional complication that the discontinuity of $\gamma^s(x)$ around the slow-down threshold, see e.g., (14)–(16), (12) and (13). We do conjecture, though, that the scheme is asymptotic efficient; for specific cases we proved this in [8].

IV. Computational Aspects and Design Issues

In this section we first study the efficiency gain achieved by applying IS (rather than naïve simulation), and then we address a number of design issues.

In Table I we present simulation results for four different parameter settings using the new measure defined in (19). Instead of performing a fixed number of simulation runs such as in much of the IS literature, we simulated until the relative error of the estimator reached some prespecified value, viz. 10^{-2} for the first three cases and $5 \cdot 10^{-2}$ for the last one. In the table we present 95% confidence intervals for p_B^s, the number of needed replications (# runs), the used machine time in seconds, and the number of 'succesful' replications (# succ.), i.e. the number of runs that resulted in buffer overflow.

We compare three values of the overflow level B; the value of δ was taken to be $\delta = -\frac{1}{3}\epsilon \log \epsilon$ and $\epsilon = 0.001$, as in [8]. We fixed the starting state s at the origin, since this is the most natural choice from a practical perspective. Clearly,

$(\lambda, \mu_1, \mu_1^+, \mu_2) = (0.1, 0.7, 0.15, 0.2)$, RE = 0.01				
B	p_B^s	# succ.	# runs	time
20	$3.79 \cdot 10^{-7} \pm 7.44 \cdot 10^{-9}$	15,576	28,332	0.4
50	$1.28 \cdot 10^{-16} \pm 2.52 \cdot 10^{-18}$	33,542	58,332	2
100	$3.54 \cdot 10^{-32} \pm 6.95 \cdot 10^{-34}$	56,982	109,992	8
$(\lambda, \mu_1, \mu_1^+, \mu_2) = (0.3, 0.36, 0.32, 0.34)$, RE = 0.01				
20	$5.63 \cdot 10^{-2} \pm 1.11 \cdot 10^{-4}$	39,496	91,596	2
50	$1.19 \cdot 10^{-3} \pm 2.33 \cdot 10^{-5}$	99,567	241,332	18
100	$1.63 \cdot 10^{-6} \pm 3.21 \cdot 10^{-8}$	128,864	320,120	49
$(\lambda, \mu_1, \mu_1^+, \mu_2) = (0.3, 0.36, 0.35, 0.34)$, RE = 0.01				
20	$5.86 \cdot 10^{-2} \pm 1.44 \cdot 10^{-4}$	32,283	76,169	2
50	$1.42 \cdot 10^{-3} \pm 2.79 \cdot 10^{-5}$	61,034	152,283	12
100	$2.64 \cdot 10^{-6} \pm 5.18 \cdot 10^{-8}$	113,527	279,196	47
$(\lambda, \mu_1, \mu_1^+, \mu_2) = (0.25, 0.35, 0.28, 0.4)$, RE = 0.05				
20	$1.11 \cdot 10^{-4} \pm 1.09 \cdot 10^{-5}$	45,685	83,436	2
50	$3.43 \cdot 10^{-11} \pm 3.36 \cdot 10^{-12}$	79,901	148,256	7
100	$5.72 \cdot 10^{-22} \pm 5.60 \cdot 10^{-23}$	235,502	439,006	42

TABLE I
Simulation results for $\theta = 0.8$

the IS scheme provides fast and reliable estimates. In some cases, especially when B grows large, the running times may be sensitive to the choice of ϵ and δ. In Table I we used relatively 'good' ϵ and δ. For instance, choosing $\epsilon = 0.01$ will lead to almost $9-$fold increase in the number of replications in the last line of the first part of Table I, i.e., for the case when $(\lambda, \mu_1, \mu_1^+, \mu_2) = (0.1, 0.7, 0.15, 0.2)$ and $B = 100$.

We also performed a few standard Monte Carlo simulations (i.e., without IS) for comparison, using the same relative error of 10^{-2} to assess the efficiency gain. For the parameter settings of the first part of Table I with $B = 20$, this took 4521 seconds ($\pm 5 \cdot 10^9$ runs). In the settings of the second part of Table I with $B = 50$ it took 118 seconds ($\pm 10^7$ runs).

To enable comparison with the state-independent scheme in [7] and the state-dependent scheme in [2], we also fixed the

$(\lambda, \mu_1, \mu_1^+, \mu_2) = (0.1, 0.7, 0.15, 0.2)$			
B	st.-ind., [7]	st.-dep., [2]	current
20	$1.49 \cdot 10^{-3}$	$2.63 \cdot 10^{-3}$	$3.54 \cdot 10^{-3}$
50	$2.06 \cdot 10^{-3}$	$7.87 \cdot 10^{-3}$	$8.00 \cdot 10^{-3}$
100	$2.75 \cdot 10^{-3}$	$19.71 \cdot 10^{-3}$	$17.01 \cdot 10^{-3}$
$(\lambda, \mu_1, \mu_1^+, \mu_2) = (0.3, 0.36, 0.32, 0.34)$			
20	$0.92 \cdot 10^{-3}$	$5.30 \cdot 10^{-3}$	$6.00 \cdot 10^{-3}$
50	$12.50 \cdot 10^{-3}$	$8.40 \cdot 10^{-3}$	$11.00 \cdot 10^{-3}$
100	$39.69 \cdot 10^{-3}$	$12.20 \cdot 10^{-3}$	$11.00 \cdot 10^{-3}$

TABLE II
Comparison of relative errors for three IS schemes

number of runs to be 10^6 and compared the relative errors, see Table II. Here, $s = (0, 0)$, $\theta = 0.8$, and in the state-dependent schemes $\epsilon = 0.03/\sqrt{B}$ and $\delta = -\epsilon \log \epsilon$. As can be expected, both state-dependent schemes provide good estimates (in terms of the relative errors), but the performance of the state-independent scheme strongly depends on the parameters.

We now demonstrate techniques that enable selection of a proper value for the slowdown threshold $m = \theta B$. A first *caveat* is the following. It is natural to expect that smaller θ will provide better protection of the second node and consequently smaller probability of overflow (with s being the origin), but this is not always the case. Indeed, numerical experiments show that, starting in $\theta = 1$, decreasing θ leads to a reduction of p_B^s. However, continuing to decrease θ, the probability of interest will start to *increase*. The same holds for the *stationary* probability of the process X_k, with the slow-down threshold θ, to be above level 1, denoted by $\pi^\theta(B)$, see Figs. 5 and 6. In these graphs we plotted $\pi^\theta(B)$ against θ for parameters $(\lambda, \mu_1, \mu_1^+, \mu_2) = (0.3, 0.36, 0.32, 0.34)$, with overflow levels $B = 20$ and $B = 50$.

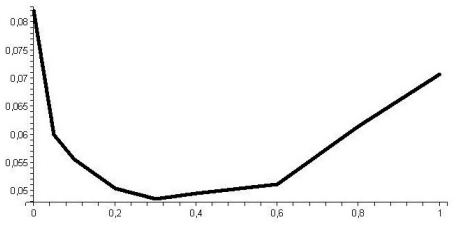

Fig. 5. $\pi^\theta(20)$ against θ

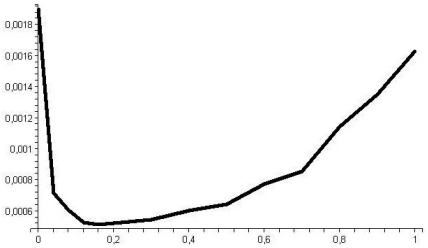

Fig. 6. $\pi^\theta(50)$ against θ

For the case of 'shifting bottlenecks', i.e., $\mu_1^+ < \mu_2 < \mu_1$, we now provide an explanation for the non-monotone behavior of $\pi^\theta(B)$. Clearly, $\pi^\theta(B)$ has decay rate $\theta \log \rho_2 + (1 - \theta) \log z^+$, for $\theta \in (0, 1]$ and z^+ from (23), see Section III;

ρ_2 is defined as λ/μ_2. However, when $\theta = 0$ the decay rate is just $\log \rho_2$, as the queue is then an ordinary tandem queue (without backpressure); in [7] it was shown that $z^+ < \rho_2$. This shows that the decay rate is discontinuous in $\theta = 0$, explaining that $\pi^0(B) > \pi^{0+}(B)$ (for B large).

The above type of justification for the observed non-monotonicity is valid only for the case that $\mu_1^+ \leq \mu_2 < \mu_1$. Another explanation for the decreasing nature of $\pi^\theta(B)$ for small θ is in the 'specific' behavior of the X_k around the origin. More precisely, consider the process X_k with starting state $s = (0, 1)$, and compare threshold levels $\theta = 0$ and $\theta = 1/B$. In the latter case, the first server operates at full speed only when the second queue is empty. It is not difficult to see that the probability of transition $(1, 0) \to (0, 1)$ is higher when $\theta = 1/B$; and the probability of the 'terminating' transition $(0, 1) \to (0, 0)$ is $\mu_2/(\lambda + \mu_2)$, which does not depend on θ. This means that the probability of overflow starting from the origin, p_B^0, is higher when $\theta = 0$ than when $\theta = 1/B$, even though we have 'more slow-down' in the first case. One can generalize this type of arguments for the other states around the origin and the other values of θ. These arguments, unlike the ones based on decay rates, hold for all parameter values.

We now demonstrate how to develop procedures for optimally choosing the value of the slow-down threshold. The primary role of the backpressure mechanism is to control the probability of some undesirable event, viz. overflow in the second buffer (expressed in terms of p_B^s). However, introducing server slow-down has a negative side effect: the expected sojourn time of a job is decreasing in θ. In order to find an optimal value of θ, one could, for given coefficient α and β, minimize the following (dis-)utility function

$$u(\theta, B) = -\alpha \log^{-1} p_B^0 + \beta S(\theta, B),$$

where $S(\theta, B)$ is the mean sojourn time of a job, α is the penalty for overflow and β is the cost for each job being in the system per unit time; we assume s is the origin. The α and β should be chosen by the service provider, and should reflect the Service Level Agreement (SLA) as agreed upon.

We present plots of the utility functions $u(\theta, 20)$ and $u(\theta, 50)$, with $\alpha = 10$ and $\beta = 3$, for a system with parameters $(\lambda, \mu_1, \mu_1^+, \mu_2) = (0.3, 0.36, 0.32, 0.34)$ in Figures 7 and 8 respectively.

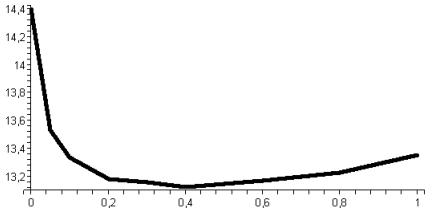

Fig. 7. Utility function against θ for $B = 20$

We observe that the optimal points are located close to the minimum of $\pi^\theta(B)$, see Figures 5 and 6, as one may expect.

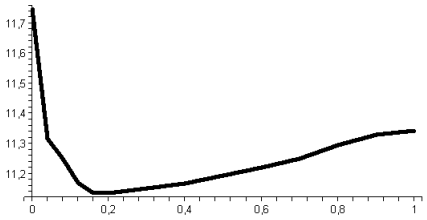

Fig. 8. Utility function against θ for $B = 50$

V. Conclusion

In this paper we analyzed a backpressure-based control mechanism. We first determined its (non-trivial) stability condition. Then we focused on efficient IS-based simulation techniques for estimating the probability of overflow in the downstream queue (which outperforms methods [9] developed earlier). It is noted that several aspects, which are not captured by the notion of asymptotic efficiency, play a crucial role in assessing the performance of this type of simulation techniques: it matters for instance very much whether a new measure requires computation of new transition rates 'on the fly', or whether these can be precomputed. These issues have been taken into account in the present paper. We then demonstrated how the evaluation techniques developed in this paper help in tuning the design parameters involved, specifically focusing on selecting an appropriate value for the slow-down threshold.

References

[1] P.T. de Boer. Analysis of state-independent importance-sampling measures for the two-node tandem queue. *ACM Transactions on Modeling and Computer Simulation*, 16(3):225–250, 2006.

[2] P. Dupuis, K. Leder, and H. Wang. Large deviations and importance sampling for a tandem network with slow-down. *Queueing Systems: Theory and Applications*, 57(2-3):71–83, 2007.

[3] P. Dupuis, A.D. Sezer, and H. Wang. Dynamic importance sampling for queueing networks. *Annals of Applied Probability*, 17(4):1306–1346, 2007.

[4] S. Floyd and V. Jacobson. Congestion gateways for packet networks. *IEEE/ACM Transactions on Networking*, 1:397–413, 1993.

[5] P. Heidelberger. Fast simulation of rare events in queueing and reliability models. *ACM Transactions on Modeling and Computer Simulation*, 5(1):43–85, 1995.

[6] G. Latouche and V. Ramaswami. *Introduction to matrix analytic methods in stochastic modelling*. SIAM, 1999.

[7] D.I. Miretskiy, W.R.W. Scheinhardt, and M.R.H. Mandjes. Efficient simulation of a tandem queue with server slow-down. *Simulation*, 83(11):751–767, 2007.

[8] D.I. Miretskiy, W.R.W. Scheinhardt, and M.R.H. Mandjes. Simple and efficient importance sampling scheme for a tandem queue with server slow-down. In *Proceedings of RESIM*, pages 38–50, Rennes, France, 2008.

[9] D.I. Miretskiy, W.R.W. Scheinhardt, and M.R.H. Mandjes. State-dependent importance sampling for a Slow-down tandem queue. *Submitted*, 2008. See also Memorandum 1879, Dept. of Applied Mathematics, University of Twente, *URL: http://eprints.eemcs.utwente.nl/13251/*.

[10] S. Parekh and J. Walrand. A quick simulation method for excessive backlogs in networks of queues. *IEEE Transactions on Automatic Control*, 34:54–66, 1989.

[11] IEEE standard 802.3. Carrier Sense Multiple Access with Collision Detection (CSMA/CD) access method and physical layer specification, annex 31b. 1998.

[12] R.L. Tweedie. Operator-geometric stationary distributions for Markov chains, with applications to queueing models. *Adv. in Appl. Probab.*, 14:368–391, 1982.

[13] N. D. van Foreest, M.R.H. Mandjes, J.C.W. van Ommeren, and W.R.W. Scheinhardt. A tandem queue with server slow-down and blocking. *Stochastic Models*, 21(2-3):695–724, 2005.