Optimal Buffer Sharing

Israel Cidon, *Senior Member, IEEE,* Leonidas Georgiadis, *Senior Member, IEEE,*
Roch Guérin, *Senior Member, IEEE,* and Asad Khamisy

Abstract—We address the problem of designing optimal buffer management policies in shared memory switches when packets already accepted in the switch can be dropped (pushed-out). Our goal is to maximize the overall throughput, or equivalently to minimize the overall loss probability in the system. For a system with two output ports, we prove that the optimal policy is of push-out with threshold type (POT). The same result holds if the optimality criterion is the weighted sum of the port loss probabilities. For this system, we also give an approximate method for the calculation of the optimal threshold, which we conjecture to be asymptotically correct. For the N-ported system, the optimal policy is not known in general, but we show that for a symmetric system (equal traffic on all ports) it consists of always accepting arrivals when the buffer is not full, and dropping one from the longest queue to accommodate the new arrival when the buffer is full. Numerical results are provided which reveal an interesting and somewhat unexpected phenomenon. While the overall improvement in loss probability of the optimal POT policy over the optimal coordinate-convex policy is not very significant, the loss probability of an *individual* output port remains approximately constant as the load on the other port varies and the optimal POT policy is applied, a property not shared by the optimal coordinate-convex policy.

I. INTRODUCTION

SHARED-MEMORY fast packet switches consist of a single large memory where packets arriving from all inputs are stored while they wait before being transmitted on their respective output(s). While this design presents a number of technical challenges, in particular memory access and speed, the sharing of a single memory by all input and output ports offers numerous advantages. One of them is improved buffer efficiency, which translates into smaller memory sizes to satisfy a given loss probability requirement. However, despite this greater efficiency, losses remain unavoidable and it is, therefore, still of interest to understand how they can be minimized. Furthermore, sharing of the memory also introduces new potential problems as individual inputs can now affect the performance seen by others. In this paper, we focus on identifying how to best share the memory between the system ports, so that overall system throughput is maximized.

There has been a number of prior works which have addressed this problem. In particular, Kamoun and Kleinrock [3]

analyzed several sharing schemes, namely, complete sharing (CS) in which an arriving packet is accepted if any storage space is available, complete partitioning (CP) in which the entire storage is permanently partitioned among the output ports, sharing with maximum queue lengths (SMXQ) in which a limit on the number of buffers allocated to each output port is imposed, sharing with a minimum allocation (SMA) in which a minimum number of buffers is always reserved for each output port and the remaining buffers are shared between all output ports, and sharing with a maximum queue and minimum allocation (SMQMA) which is a combination of the SMXQ and SMA schemes. Their study assumed independent Poisson arrivals and exponential service times and they obtained closed form expressions for the probability distribution of the buffer occupancy, based on the fact that it has a well-known product form solution. From their numerical examples, they showed that sharing can improve performance especially when little storage is available, but that some restrictions should be imposed to avoid throughput degradation in asymmetric systems.

The existence and the structure of an optimal sharing policy (in the sense of minimum packet loss or maximum throughput) was then first investigated by Foschini and Gopinath [2]. They considered optimality within the class of policies that never drop a packet once they admit it in the buffer, and have coordinate-convex state space Ω (if $x \in \Omega$, then $(x_1, x_2, \cdots, (x_i - 1)^+, \cdots, x_N) \in \Omega$ for all $i = 1, \cdots, N$). These policies, referred to as coordinate-convex policies, include the policies of [3]. For a switch with two output ports and independent Poisson arrivals and service times they proved that the optimal coordinate-convex policy is to limit the queue length of output port $i, i = 1, 2$ to some fixed level m_i, such that $m_1 + m_2 \geq B$, where B is the buffer size. For more than two ports they conjectured that the optimal policy is *simple* (see definition in [2]). Their proofs were based on the fact that the probability distribution of the buffer occupancy has a product form solution.

Wei *et al.* [8], suggested a sharing policy which allows for the dropping of accepted packets, and therefore does not belong to the class of coordinate-convex policies. According to this policy (named, drop-on-demand, or DoD), an arriving packet is always accepted if there is an empty buffer. If a packet destined for output port i arrives and finds the buffer full and output port l has more packets in the shared-memory than any other ports, the following action is taken: if $i = l$, the arriving packet is dropped; if $i \neq l$, the arriving packet joins the buffer and one port l packet is dropped. In general, policies which can accept an arriving packet by dropping

another packet from the system are known as *push-out* policies (see, e.g., [4]). Push-out policies include coordinate-convex policies (never push-out a packet) as well as the DoD policy. In [8], numerical examples were provided showing that the DoD policy yields better throughput and lower packet losses than either the CS and CP policies. However, as we shall show, this policy is optimal only for symmetric systems.

In this paper we consider a model similar to the one of [3]. The buffer size is denoted by B, and the arrival and service processes of type i (destined to output i) packets are Poisson and exponential with rates λ_i and μ_i, respectively. Upon arrival of a packet the system can decide to either accept the packet, or reject it, or accept it and drop another packet from the system. In other words, we include pushout policies and our goal is to determine the policy which maximizes the overall throughput, or equivalently minimizes the overall loss probability.

For a two-ported switch, we prove that the optimal policy is of *push-out with threshold* type (POT), i.e., whenever the buffer is nonfull, the arrival should be accepted, and whenever it is full, an arrival from type $i, i = 1, 2$, is accepted and a type $\bar{i}$ (the other type) packet is pushed-out if the number of type i packets is below some threshold k_i^* (where $k_1^* + k_2^* = B$). The same result is true if the optimality criterion is the weighted sum of the port loss probabilities. For $\mu_1 = \mu_2$ and $\lambda_1 \geq \lambda_2$ we also show that $k_1^* \leq B/2$. In general, the determination of the threshold k_1^* is computationally intensive, but for the two-ported system we develop a simple and yet reasonably accurate heuristic to obtain its value. The results for the two-ported system establish the nonoptimality of DoD for asymmetric systems.

For the symmetric N-ported system with identical arrival rates and identical transmission rates, we show that the optimal policy is to accept an arrival whenever the buffer is nonfull, or the queue corresponding to the type of the arriving packet is not the largest; in the second case a packet from the longest queue is dropped. This establishes the optimality of DoD for the N-ported symmetric system. The proofs of the results for both the two-ported and N-ported systems are based on the theory of Markov decision processes.

The behavior of the optimal policies are then investigated for the two-ported case by means of numerical examples, which reveal an interesting and somewhat unexpected phenomenon. While the overall improvement in loss probability of the optimal POT policy over the optimal coordinate-convex policy is found to be relatively minor, a significant difference is observed when focusing on the loss probability of an *individual* output port. The use of the optimal POT policy results in an approximately constant loss probability on a given port as the load on the other varies. In contrast, significant variations can be observed with the optimal coordinate-convex policy. The insensitivity of individual losses is clearly a desirable feature, but nevertheless surprising given the global nature (overall throughput) of our optimization. For the two-ported system, we also investigate a heuristic method for determining the threshold of the optimal POT policy, which based on the numerical results obtained is conjectured to be asymptotically correct as the buffer size B increases.

Numerical comparisons further show that the approximation is very good for most practical scenarios.

The paper is organized as follows. In Section II we introduce the system model and provide the formulation of the optimization problem. In Section III we investigate the structure of the optimal policy. Section IV is devoted to numerical comparisons between the performance of the optimal POT and coordinate-convex policies. Finally, the appendix provides proofs of lemmas used in Section III.

II. THE MODEL AND PROBLEM FORMULATION

The system consists of a buffer shared by packets destined to any of N output ports. Packets are said to be of type $i, 1 \leq i \leq N$, if they are destined to port i. Type $i, 1 \leq i \leq N$, packets arrive to the buffer according to a Poisson process with rate λ_i, and are transmitted by output port i with a transmission time which is exponentially distributed with rate μ_i. We assume that $\lambda_i, \mu_i < \infty$ so that only a finite number of transitions can occur in any finite interval of time, and that packet inter-arrival and transmission times from all sources are mutually independent. The total buffer size is taken to be B packets, and a packet occupies its buffer until it has been completely transmitted.

Our goal is to determine how the B buffers are "best" shared among packets of different types, so that the overall system throughput is maximized. This amounts to identifying rules that specify when and how packets of different types are allowed to occupy a space in the shared buffer. In this paper, acceptable rules include accepting or rejecting an arriving packet as well as discarding (pushing-out) an already stored packet to accommodate an arriving one. Because the state of the system can be represented by a Markov chain, the rules or policy governing the sharing of the buffer can be expressed as a continuous time Markov decision process. Decision epochs correspond to arrivals and departures from the system, where at each epoch, a decision is made as to whether the next arriving packet should be accepted, rejected, or accepted by pushing-out another packet from the system. The decision may depend on the type of the arriving packet. Next, we proceed with a precise formulation of this process.

Let $x(n) = (x_1(n), x_2(n), \cdots, x_N(n))$ be the state of the system at decision epoch $n = 0, 1, \cdots$, where $x_i(n), 1 \leq i \leq N$, denotes the number of type i packets in the system at decision epoch n. Let $X = \{0, 1, \cdots, B\}^N$ be the state space of the system. Define the following operators to denote the rejection, acceptance and push-out of packets, respectively (upon arrival epoch)

$$P_r(x) = x$$
$$P_{a_i}(x) = x + e_i, \quad x \in \text{Dom}(P_{a_i}), \quad 1 \leq i \leq N \qquad (1)$$
$$P_{p_i^j}(x) = x + e_i - e_j, \quad x \in \text{Dom}(P_{p_i^j}), \quad 1 \leq i \neq j \leq N$$

where, $e_i, 1 \leq i \leq N$, denotes the vector with all components zero except the ith which is equal 1, $\text{Dom}(P_{a_i}) \overset{\Delta}{=} \{\sum_{j=1}^N x_j < B\}$ and $\text{Dom}(P_{p_i^j}) \overset{\Delta}{=} \{x_j > 0\}$.

Let $U = \{u = (u_1, u_2, \cdots, u_N) : u_i \in \{r, a_i, p_i^j, 1 \leq j \neq i \leq N\}\}$ be the set of possible decisions, and $U(x) = \{u \in$

$U : x \in \cap_{i=1}^{N} \text{Dom}(P_{u_i})\}$ be the set of all admissible actions when the system state is x. The specification of the continuous time Markov decision process is then completed once we have defined the length of time between successive decision epochs and the transition probability function. When the system is in state x, the length of time until the next decision epoch is an exponential random variable with transition rate $\sum_{i=1}^{N}(\lambda_i + \mu_i \mathbf{1}\{x_i > 0\})$, and the transition probability to the next state is given by

$$\Pr[x(n+1) = P_{u_i}(x)|x(n) = x, u(n) = u]$$
$$= \frac{\lambda_i}{\sum_{i=1}^{N}(\lambda_i + \mu_i \mathbf{1}\{x_i > 0\})}, \qquad 1 \le i \le N$$
$$\Pr[x(n+1) = D_i(x)|x(n) = x, u(n) = u]$$
$$= \frac{\mu_i \mathbf{1}\{x_i > 0\}}{\sum_{i=1}^{N}(\lambda_i + \mu_i \mathbf{1}\{x_i > 0\})}, \qquad 1 \le i \le N$$
$$(2)$$

where $D_i(x)$ indicates a departure of a type i packet when the system is in state x, i.e., $D_i(x) = x - e_i$.

As mentioned earlier, the objective of the optimization is to minimize the overall loss probability, or equivalently to maximize the overall throughput of the system. That is, denoting by $n(T)$ the number of decision epochs up to time T, we are interested in the policy that for any initial state $x(0) = x$, maximizes

$$\lim_{T\to\infty} \inf \frac{1}{T} E\left(\sum_{k=1}^{n(T)} \sum_{i=1}^{N} d_i(k)|x(0) = x\right) \qquad (3)$$

where $d_i(k) = 1$ if at the kth decision epoch a departure of type i occurs and $d_i(k) = 0$ otherwise. The above expression can be generalized to a *weighted* average gain of the form

$$\lim_{T\to\infty} \inf \frac{1}{T} E\left(\sum_{k=1}^{n(T)} \sum_{i=1}^{N} w_i d_i(k)|x(0) = x\right),$$
$$w_i \ge 0, 1 \le i \le N$$

and the proof of Section III on the structure of the optimal policy actually goes through for this more general case. However, in order to keep notations simple we will assume that $w_i = 1, 1 \le i \le N$ in the rest of the paper.

The above continuous time problem needs to be translated into an equivalent discrete time optimization problem, and we use uniformization (see [5], [7], [1, sect. 6.7]) for that purpose. In our case, uniformization amounts to introducing, when the system is in one of the states $x \in \{x_i = 0,$ for at least one $i\}$, *fictitious* departure epochs of type i which occur with rate $\mu_i \mathbf{1}\{x_i = 0\}$. The gain of a fictitious departure is taken to be zero. This modification does not alter the stochastic behavior of the process and results in identical transition rates, equal to $\sum_{i=1}^{N}(\lambda_i + \mu_i)$, independent of the state of the system and the decision taken. The transition probabilities in the new system are now given by

$$\Pr[x(n+1) = P_{u_i}(x)|x(n) = x, u(n) = u]$$
$$= \frac{\lambda_i}{\sum_{i=1}^{N}(\lambda_i + \mu_i)}, \qquad 1 \le i \le N$$

$$\Pr[x(n+1) = D_i(x)|x(n) = x, u(n) = u]$$
$$= \frac{\mu_i}{\sum_{i=1}^{N}(\lambda_i + \mu_i)}, \qquad 1 \le i \le N \qquad (4)$$

and the equivalent discrete time optimization objective is to find a policy that for any initial state $x(0) = x$, maximizes

$$\lim_{n\to\infty} \inf \frac{1}{n} E\left(\sum_{k=1}^{n} \sum_{i=1}^{N} d_i(k)|x(0) = x\right)$$

where $d_i(n) = 1$ if at the nth decision epoch a *real* (nonfictitious) departure of type i occurs and $d_i(n) = 0$ otherwise. In the next section, we consider the problem of identifying the policy that maximizes this gain function.

III. THE OPTIMAL POLICY

To avoid cumbersome notation, in the following we assume without loss of generality that the arrival and departure rates are normalized so that, $\sum_{i=1}^{N}(\lambda_i + \mu_i) = 1$. A standard approach to solving the kind of optimization problem we consider, is to first work with the discounted gain criterion [6]

$$E \sum_{n=1}^{\infty} \beta^n \sum_{i=1}^{N} d_i(n) \qquad (5)$$

where $0 < \beta < 1$ is a discount factor.

In our system, since the gains $d_i(n)$ are bounded, it is known [6] that there is always an optimal stationary policy to the discounted problem. Therefore, we limit our investigations to this class of policies. A stationary policy π is a function $\pi : X \to U$ with $\pi(x) \in U(x)$ for every $x \in X$, and such that under π the decision $u = \pi(x)$ is always taken whenever the system is in state x. In order to carry out our investigation of the optimal policy, we need to introduce the Banach space $\mathcal{F}$ of all bounded real functions $f : X \to R$ with norm $\| \cdot \|$ given by $\|f\| = \max_{x \in X}|f(x)|$.

For any stationary policy π we then define $T_{\pi} : \mathcal{F} \to \mathcal{F}$ (the dynamic programming operator) by

$$(T_\pi f)(x) = \sum_{i=1}^{N} \mu_i \mathbf{1}\{x_i > 0\}$$
$$+ \sum_{i=1}^{N} \beta \mu_i f(D_i(x)) + \sum_{i=1}^{N} \beta \lambda_i f(P_{u_i}(x)) \qquad (6)$$

where, $\pi(x) = (u_1, u_2, \cdots, u_N)$. The first term in the right hand expression is the one-step gain for the system under policy π, while the rest of the terms correspond to the gain that incurs after the first step. Based on (6) we then define the operator T by $(Tf)(x) = \max_{\pi}(T_\pi f)(x)$ for all x. If we denote the optimal gain starting from state x by $J_\beta(x)$, the following results are well-known (see, e.g., [6]).

- For every $x \in X$, $J_\beta(x) = TJ_\beta(x)$.
- For any $f \in \mathcal{F}$, $\lim_{n\to\infty} T^{(n)} f(x) = J_\beta(x)$ for every $x \in X$, where $T^{(n)}$ is the $n-$ fold composition of the operator T.
- A stationary policy π is optimal iff $J_\beta(x) = T_\pi J_\beta(x)$ for every $x \in X$.

In the next sections, we rely on these results to identify the optimal policy for our system. We focus first on the two-ported system for which we show that the optimal policy is of POT type in the general case of different arrival and service rates on each port. We also prove some interesting properties of this policy in some cases and propose a simple heuristic to compute the optimal threshold value. The accuracy of this approximation is later evaluated in Section IV. The N-ported system is treated next, but only for the symmetric case for which the optimal policy is identified.

A. The Two-Ported System ($N = 2$)

1) Optimal Policy Derivation: In order to determine the optimal policy for the two-ported system, it is necessary to specify enough of its properties so that it is fully characterized. The basic approach we employ to identify these properties is the value iteration method (see [6]). It is based on the fact that the optimal value function can be shown to obey a certain property simply by showing that if this property holds for a function $f \in \mathcal{F}$, then it continues to hold for the function Tf (which also belongs to $\mathcal{F}$). The main difficulties in using this approach are in initially "guessing" the properties of the optimal policy, and in selecting appropriate auxiliary properties for the function f which are usually necessary to prove that the desired properties hold.

The dynamic programming operator for the two-ported system is

$$Tf(k_1, k_2) = \mu_1 1\{k_1 > 0\} + \mu_2 1\{k_2 > 0\}$$
$$+ \beta\mu_1 f((k_1 - 1)^+, k_2)$$
$$+ \beta\mu_2 f(k_1, (k_2 - 1)^+)$$
$$+ \beta\lambda_1 \Phi_1(k_1, k_2) + \beta\lambda_2 \Phi_2(k_1, k_2) \quad (7)$$

where

$$\Phi_1(k_1, k_2) = \max\{f(k_1, k_2), f(k_1 + 1, k_2 - 1)\} \text{ if } k_1$$
$$+ k_2 = B$$
$$\Phi_1(k_1, k_2) = \max\{f(k_1, k_2), f(k_1 + 1, k_2),$$
$$f(k_1 + 1, k_2 - 1)\} \text{ if } k_1 + k_2 < B \quad (8)$$

and

$$\Phi_2(k_1, k_2) = \max\{f(k_1, k_2), f(k_1 - 1, k_2 + 1)\} \text{ if } k_1$$
$$+ k_2 = B$$
$$\Phi_2(k_1, k_2) = \max\{f(k_1, k_2), f(k_1, k_2 + 1),$$
$$f(k_1 - 1, k_2 + 1)\} \text{ if } k_1 + k_2 < B \quad (9)$$

where we define $\Phi_i(k_1, k_2) \triangleq 0$ for $k_1, k_2 < 0$ or $k_1, k_2 > B$.

Lemma 1: The optimal value function $J_\beta(k_1, k_2)$ has the following properties:

1) Monotonicity and boundedness in k_1: $0 \leq J_\beta(k_1 + 1, k_2) - J_\beta(k_1, k_2) \leq 1, 0 \leq k_1 \leq B - 1$.
2) Monotonicity and boundedness in k_2: $0 \leq J_\beta(k_1, k_2 + 1) - J_\beta(k_1, k_2) \leq 1, 0 \leq k_2 \leq B - 1$.
3) Concavity along k_1: $J_\beta(k_1 + 1, k_2) - J_\beta(k_1, k_2) \leq J_\beta(k_1, k_2) - J_\beta(k_1 - 1, k_2), 1 \leq k_1 \leq B - 1$.

4) Concavity along k_2: $J_\beta(k_1, k_2 + 1) - J_\beta(k_1, k_2) \leq J_\beta(k_1, k_2) - J_\beta(k_1, k_2 - 1), 1 \leq k_2 \leq B - 1$.
5) Concavity along the line $k_1 + k_2 = b, 2 \leq b \leq B$: $J_\beta(k_1 + 1, k_2 - 1) - J_\beta(k_1, k_2) \leq J_\beta(k_1, k_2) - J_\beta(k_1 - 1, k_2 + 1), 1 \leq k_1 \leq B - 1, 1 \leq k_2 \leq B - 1$.

Proof: Consider a function $f \in \mathcal{F}$ which satisfies Properties 1–5 of Lemma 1 (for example the function "0"). In the appendix, we show that the function Tf also satisfies these properties. Using the fact that $\lim_{n\to\infty} T^{(n)} f(x) = J_\beta(x)$ for every $x \in X$, the function $J_\beta(\cdot)$ then satisfies these properties as well. □

Properties 3–4 are only auxiliary properties needed for the proof of Property 5, and the combination of Properties 1, 2, and 5 of Lemma 1 allows to directly establish the following proposition.

Proposition 1: The optimal policy for the discounted gain problem is of type POT.

Proof: Properties 1–2 of Lemma 1 imply that an arrival from either type should be accepted if an empty buffer is available. From Property 5, the optimal value function is concave on the line $k_1 + k_2 = B$ and hence it has either a unique maximum or two maxima at consecutive points on this line. The fact that $J_\beta(x)$ satisfies the dynamic programming equation $J_\beta(x) = TJ_\beta(x)$, implies that the value of k_1 (or k_2) corresponding to the maximum (in case two maxima exist either one can be chosen) is the optimal threshold in the POT policy. □

Lemma 1 and Proposition 1 can be shown to hold for the weighted average gain criterion with weights $0 \leq w_1, w_2 \leq 1, w_1 + w_2 = 1$. In this case, the differences in Properties 1 and 2 of Lemma 1 should be upper bounded by w_1 and w_2 instead of 1, respectively, and the proofs follow as in the appendix.

2) The Two-Ported System with Equal Transmission Rates ($\mu_1 = \mu_2$): In this section we establish an interesting albeit intuitive property of the optimal policy for the two-ported system when the transmission rate on both output ports are equal, $\mu_1 = \mu_2 \triangleq \mu$. For that system, we show that when $\lambda_1 \geq \lambda_2$, the optimal threshold satisfies, $k_1^* \leq B/2$. The proof of this property relies again on the value iteration method.

The dynamic programming operator for this system is now

$$Tf(k_1, k_2) = \mu[1\{k_1 > 0\} + 1\{k_2 > 0\}]$$
$$+ \beta\mu[f((k_1 - 1)^+, k_2) + f(k_1, (k_2 - 1)^+)]$$
$$+ \beta\lambda_1 \Phi_1(k_1, k_2) + \beta\lambda_2 \Phi_2(k_1, k_2) \quad (10)$$

where, $\Phi_i(\cdot), i = 1, 2$, were defined in (8)–(9).

Lemma 2: If $\lambda_1 \geq \lambda_2$, the optimal value function $J_\beta(k_1, k_2)$ has the following property

$$J_\beta(k_1, k_2) \geq J_\beta(k_2, k_1), k_2 \geq k_1$$

Proof: See the appendix. □

The above lemma essentially states that when $\lambda_1 \geq \lambda_2$, it is preferable to have more packets of type 2 than of type 1 in the system. When combined with Proposition 1, it directly gives the desired result stated in the following Proposition.

Proposition 2: If $\lambda_1 \geq \lambda_2$, the optimal policy for the discounted gain is POT with threshold $k_1^* \leq B/2$.

Proof: From Proposition 1, the optimal policy for the discounted gain problem is POT. From Lemma 2, we know that

$$J_\beta(k_1, B - k_1) \geq J_\beta(B - k_1, k_1), k_1 \leq \frac{B}{2}$$

which, together with the concavity of the optimal value function $J_\beta(k_1, k_2)$ on the line $k_1 + k_2 = B$ (Property 5 of Lemma 1), implies that the maximum of the optimal value function on this line occurs in the range $k_1 \leq B/2$. $\square$

This result confirms the intuition that the higher arrival rate of type 1 packets implies that they can be pushed out more often.

3) Threshold Design for the Two-Ported System: The two previous sections established that the optimal policy for the two-ported system is of type POT. However, the value of the optimal threshold was not explicitly identified. There are a number of possible approaches to obtain the value of the optimal threshold, but typically they are computationally intensive and for large values of the buffer size may even be infeasible. In this section we propose a heuristic approximation to compute the value of the optimal threshold k_1^*. The accuracy of this approximation will be assessed numerically in Section IV-B.

Using the value iteration method, it is easy to check that an upper bound to the solution of the equation $J_\beta(k_1, k_2) = T J_\beta(k_1, k_2)$ is the function

$$f_\beta(k_1, k_2) = d - c_1 \alpha_1^{k_1} - c_2 \alpha_2^{k_2}$$

where $d = (\mu_1 + \mu_2)/(1 - \beta)$, and α_i, $i = 1, 2$, is the unique solution in $(0, 1)$ of the quadratic equation

$$(1 - \beta + \mu_i\beta + \lambda_i\beta)\alpha_i = \lambda_i\beta\alpha_i^2 + \mu_i\beta \qquad (11)$$

and

$$c_i = \frac{\mu_i}{1 - \beta(1 - \lambda_i(1 - \alpha_i))}. \qquad (12)$$

In fact, the function $f_\beta(k_1, k_2)$ corresponds to the discounted long term throughput when the buffer size is infinite. We are interested in determining the value k_β^* that maximizes $J_\beta(k, B - k)$ for $k = 0, \cdots, B$. Our proposed approximation consists of assuming that k_β^* approximately maximizes $f_\beta(k, B - k)$, $k = 0, \cdots, B$ as well. We expect this to be especially true for large buffers since the optimal gain is then closer to $f_\beta(\cdot)$.

The maximum of $f_\beta(k, B - k)$, $k = 0, \cdots, B$ is easily found to be either $\max(\lfloor x^* \rfloor, 0)$ or $\max(\lfloor x^* \rfloor + 1, 0)$, where

$$x^* = \frac{\ln(c_2 \ln \alpha_2/(c_1 \ln \alpha_1))}{\ln \alpha_1 + \ln \alpha_2} + \frac{\ln \alpha_2}{\ln \alpha_1 + \ln \alpha_2} B = C_1 + C_2 B. \qquad (13)$$

To find the form of the approximation for the average gain criterion (see Section III-C for details), we have to take the limit as $\beta \to 1$. Let $\rho_i \triangleq \lambda_i/\mu_i, i = 1, 2$. It can be seen that

$$\lim_{\beta \to 1} \alpha_i = \begin{cases} 1, & \text{if } \rho_i \leq 1 \\ \rho_i^{-1}, & \text{if } \rho_i > 1 \end{cases} \quad \text{and}$$

$$\lim_{\beta \to 1} c_i = \begin{cases} \infty, & \text{if } \rho_i \leq 1 \\ \frac{\mu_i}{\lambda_i - \mu_i}, & \text{if } \rho_i > 1. \end{cases}$$

Also, when $\rho_i \leq 1$

$$\alpha_i'(1) \triangleq \lim_{\beta \to 1} \frac{d\alpha_i}{d\beta} = \frac{1}{\mu_i - \lambda_i}.$$

We can, therefore, compute (except for the case $\rho_1 = \rho_2 = 1$ which requires that we apply L'Hospital's rule one more time)

$$\bar{C}_2 = \lim_{\beta \to 1} C_2 = \lim_{\beta \to 1} \frac{\ln \alpha_2}{\ln \alpha_1 + \ln \alpha_2}$$

$$= \begin{cases} 1, & \text{if } \rho_1 \leq 1, \rho_2 > 1 \\ \frac{\ln \rho_2}{\ln \rho_1 + \ln \rho_2}, & \text{if } \rho_1 > 1, \rho_2 > 1 \\ \frac{\mu_1 - \lambda_1}{\mu_1 - \lambda_1 + \mu_2 - \lambda_2}, & \text{if } \rho_1 < 1, \rho_2 < 1 \\ 0, & \text{if } \rho_1 > 1, \rho_2 \leq 1 \end{cases}$$

and similarly $\bar{C}_1$ can be computed.

The proposed approximation for the optimal threshold k_1^* consists then of selecting the closest integer to $x^* \triangleq \bar{C}_1 + \bar{C}_2 B$. Note that the main contribution for large B comes from $\bar{C}_2$. In fact, as will be seen in Section IV-B, numerical evidences suggest the following conjecture which, however, we were not able to prove

$$\lim_{B \to \infty} \frac{k_1^*}{B} = \bar{C}_2. \qquad (14)$$

B. The Symmetric N-Ported System

In this section, we consider a system with N identical output ports, i.e., the arrival and transmission rates are the same on all ports and denoted by λ and μ, respectively. For this special case, we show that the optimal policy is to accept all packets whenever the buffer is nonfull, and when the buffer is full to accept a packet only if the queue corresponding to its destination output port is not the largest. In the latter case, a packet from the largest queue is pushed-out to accommodate the new arrival.

The dynamic programming equation for this system becomes

$$Tf(\bar{k}) = \mu \sum_{i=1}^{N} 1\{k_i > 0\}$$

$$+ \beta\mu \sum_{i=1}^{N} f((\bar{k} - e_i)^+) + \beta\lambda \sum_{i=1}^{N} \Phi_i(\bar{k}) \qquad (15)$$

where

$$\Phi_i(\bar{k}) = \max_{1 \leq j \leq N} \{f(\bar{k} + e_i - e_j)\} \text{ if } \sum_{i=1}^{N} k_i = B$$

$$\Phi_i(\bar{k}) = \max_{1 \leq j \leq N} \{f(\bar{k} + e_i), f(\bar{k} + e_i - e_j)\} \text{ if } \sum_{i=1}^{N} k_i < B \qquad (16)$$

where, we define $\bar{k} \triangleq (k_1, \cdots, k_N)$, $(\bar{k} - e_i)^+ \triangleq (k_1, \cdots, (k_i - 1)^+, \cdots, k_N)$, and $\Phi_i(\bar{k}) \triangleq 0$ if $k_j < 0$ or $k_j > B$ for some $1 \leq j \leq N$.

As for the two-ported case, we first proceed to establish a number of key properties of the optimal value function which will enable us to characterize the optimal policy.

Lemma 3: The optimal value function $J_\beta(\overline{k})$ has the following properties:

1) Monotonicity and boundedness in k_i: $0 \le J_\beta(\overline{k} + e_i) - J_\beta(\overline{k}) \le 1, 0 \le k_i \le B - 1, 1 \le i \le N$.
2) Symmetry: $J_\beta(\overline{k}) = J_\beta(\pi(\overline{k}))$ for any permutation $\pi(\overline{k})$ of the vector $\overline{k}$.
3) Balancing: For $1 \le i, j \le N$, if $k_i \ge k_j$ then $J_\beta(\overline{k}) \ge J_\beta(\overline{k} + e_i - e_j)$, otherwise $J_\beta(\overline{k}) \le J_\beta(\overline{k} + e_i - e_j)$.
4) Drop from the longest queue: For $1 \le i, j, l \le N$, if $k_i < k_j < k_l$ then $J_\beta(\overline{k} + e_i - e_l) \ge J_\beta(\overline{k} + e_i - e_j)$.

Proof: See the appendix.   □

As before, the above properties can now be applied to characterize the optimal policy which we state in the following proposition.

Proposition 3: The optimal policy for the discounted gain problem is to accept a packet whenever the buffer is nonfull. When the buffer is full, a packet is accepted only if the queue corresponding to its destination output port is not the largest among all queues. The arriving packet is then accommodated by pushing-out a packet from the largest queue.

Proof: Property 1 of Lemma 3 implies that an arriving packet should be accepted if an empty buffer is available. Property 3 implies that an arrival to a full buffer should be accepted only if it is not destined to the longest output queue; in this case Property 4 implies that a packet from the longest output queue should be dropped in order to accommodate the arriving packet.   □

C. The Average Cost Problem

In this section, we establish that the usual conditions required to extend the solution of the discounted gain problem to the average gain problem are indeed satisfied. The state space of our system is finite, and for $\mu_i > 0, 1 \le i \le N$, the state $x = 0$ is accessible from every other state regardless of which stationary policy is used. Hence, the conditions of Corollary 2.5, [6, sect. V] are satisfied and the following properties hold:

1) There exists a bounded function $h(i)$ and a constant J (the optimal value of the average gain problem which doesn't depend on the initial state) which satisfy the average gain version of the optimality equation

$$J + h(\overline{k}) = \sum_{i=1}^{N} \mu_i 1\{k_i > 0\}$$
$$+ \sum_{i=1}^{N} \mu_i h((\overline{k} - e_i)^+) + \sum_{i=1}^{N} \lambda_i \Phi_i(\overline{k}) \quad (17)$$

where $\Phi_i(k)$ is defined in (16) with $h(\cdot)$ replacing $f(\cdot)$.
2) For some sequence $\beta_n \to 1, h(\overline{k}) = \lim_{n \to \infty}[J_{\beta_n}(\overline{k}) - J_{\beta_n}(0)]$.
3) $J = \lim_{\beta \to 1}(1 - \beta)J_\beta(0)$.

From Property 2 it follows that the function $h(\overline{k})$ has the same properties as the function $J_{\beta_n}(\overline{k})$ listed in Lemmas 1–3 for some sequence $\beta_n \to 1$. Equation (17) has a form similar

to the dynamic programming equation of the discounted gain problem and Propositions 1–3 can, therefore, be shown to also hold for the average gain problem.

IV. NUMERICAL RESULTS

In this section we investigate the performance of a shared memory switch operated under the optimal policy identified in this paper. Our focus is on the two-ported case for which more general results are available. For this system, we conduct two kinds of investigations: A comparison with the performance of the optimal coordinate-convex policy; An evaluation of the accuracy of the heuristic approximation proposed to compute the optimal threshold.

In order to determine the performance of the optimal policy, we need to compute the packet loss probabilities of the system under this policy. A direct computation of state probabilities proved numerically rather unstable especially for large buffers (note that the number of states is $(B + 1)^2/2$, where B is the buffer size). Therefore, to compute the loss probabilities we used a technique based on the method of successive approximation for computing average gains presented in [1, sect. 7.2]. This technique requires certain properties from the Markov chain, but they are easily seen to be satisfied in our case. While slow to converge, this method is numerically very stable.

A. Comparison with Coordinate-Convex Policies

Since a coordinate-convex policy is a special case of a pushout policy, considering as cost the overall switch loss probability, the optimal pushout policy will have a smaller loss probability than the optimal coordinate convex policy. The price of this improvement is, however, a more complex implementation, and it is of interest to evaluate the trade-off between the gain in performance and the higher cost. In this section, we provide numerical results to help us compare the relative performance of the optimal coordinate-convex, π_c^*, and the optimal pushout, π_p^* policies. Fig. 1 gives the performance of the two policies for a buffer size of 50. In Fig. 1(a) and (b), the utilization ($\rho = \lambda/\mu$) of port 1 is kept constant at 0.6 while the utilization of port 2 varies from 0.6 to 1.9. In Fig. 1(c) and (d), the utilization of port 1 is kept constant at .8 while the utilization of port 2 varies again between 0.6 and 1.9.

We observe that the difference in overall loss probability between the two policies is not very significant. For the specific examples we consider, the maximum value of the ratio

$$\frac{\text{Loss probability of } \pi_c^*}{\text{Loss probability of } \pi_p^*}$$

is 1.18. However, when individual loss probabilities are considered, significant differences between the two policies are observed. Specifically, the loss probability of port 1 whose utilization is kept constant, is affected significantly by the variation of the utilization of port 2 under policy π_c^* and varies by up to seven orders of magnitude. Under policy π_p^*, however, the loss probability of port 1 never increases above what it experiences in the reference balanced case (equal port loads) by more than one order of magnitude as the utilization of port 2 varies. Furthermore, the better overall performance of π_p^* often also results in lower loss probabilities for both ports.

Buffer Size=50, Port 1 Utilization=.6

(a) (b)

Buffer Size=50, Port 1 Utilization=.8

(c) (d)
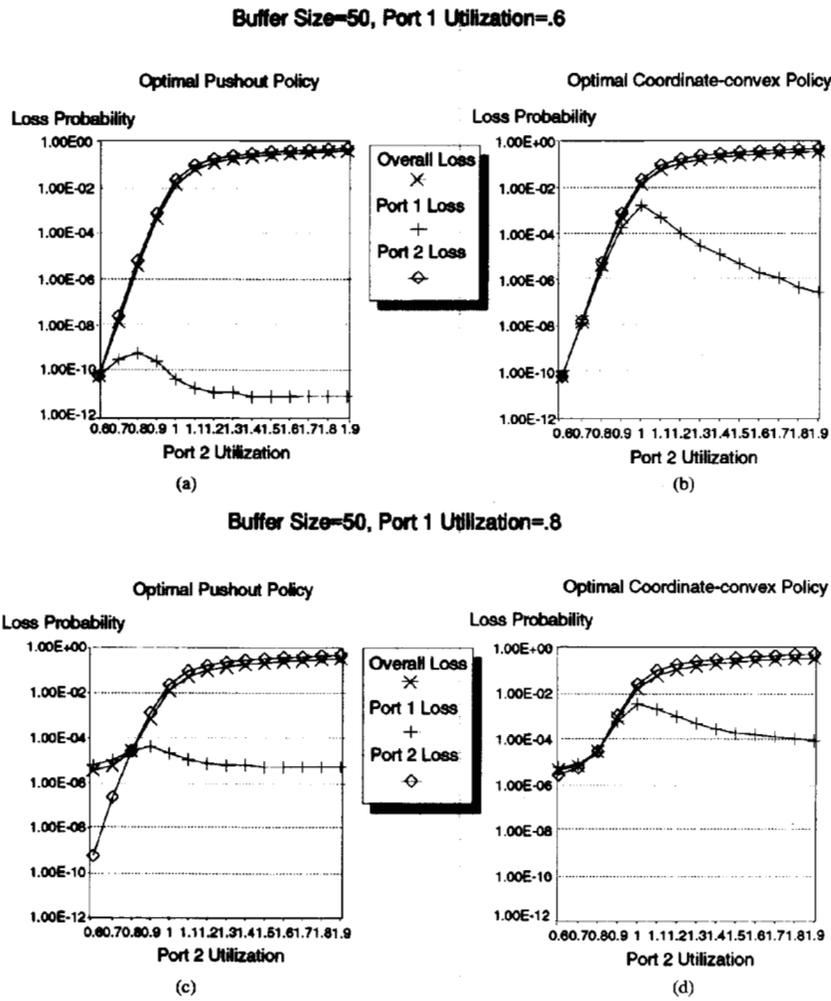
Fig. 1. Comparison of pushout and coordinate-convex policies.

Even in the few cases where the loss probability for the heavier loaded port is lower under π_c^* than under π_p^*, the difference remains minimal. Specifically, for the examples of Fig. 1, the loss probability for the heavier loaded port under π_c^* is never less than 65% of its value under π_p^*. Similar experiments were conducted for many other values of port utilizations and for buffer sizes up to 80, and similar behaviors were observed. Therefore, we concluded that an important advantage of π_p^* compared to π_c^* is that it effectively isolates the performance of a port with constant utilization from fluctuations in the utilization of the other port (assuming, the optimal policy is used in each case).

Fig. 2 illustrates another noteworthy feature of the pushout policies. In this figure we plot the total loss probability and the loss probabilities of the two ports when a POT policy, not necessarily the optimal, is employed. The buffer size in this case is 30 and the k_1 thresholds used by the POT policy are varied between 0 and 30. We observe that the overall loss varies little as the threshold changes. This implies that the performance of the optimal policy is not critically dependent on the choice of the optimal threshold. However, this is not true for individual loss probabilities at each port. As we have seen above it may be desirable, in addition to achieving a throughput close to the optimal, to also provide appropriate loss probabilities to each port. In this case a good approximation to the optimal threshold k_1^* is still required. As we shall see in the next subsection, the approximation proposed in Section III-A3 for the optimal threshold is quite good, especially when the port utilizations are smaller than 1.

B. Accuracy of the Approximation

Fig. 3 compares the values of the optimal threshold and the proposed approximate threshold for the parameters of Fig. 1.

We see that the approximate values are very close to the optimal when utilization of both ports is smaller than one. Numerical experimentation has also shown, that the approximation while somewhat less accurate, is still good when the utilization of both ports is larger than one. The approximation is less accurate when the utilization of one port is smaller than but close to 1, while the utilization of

Buffer Size = 30
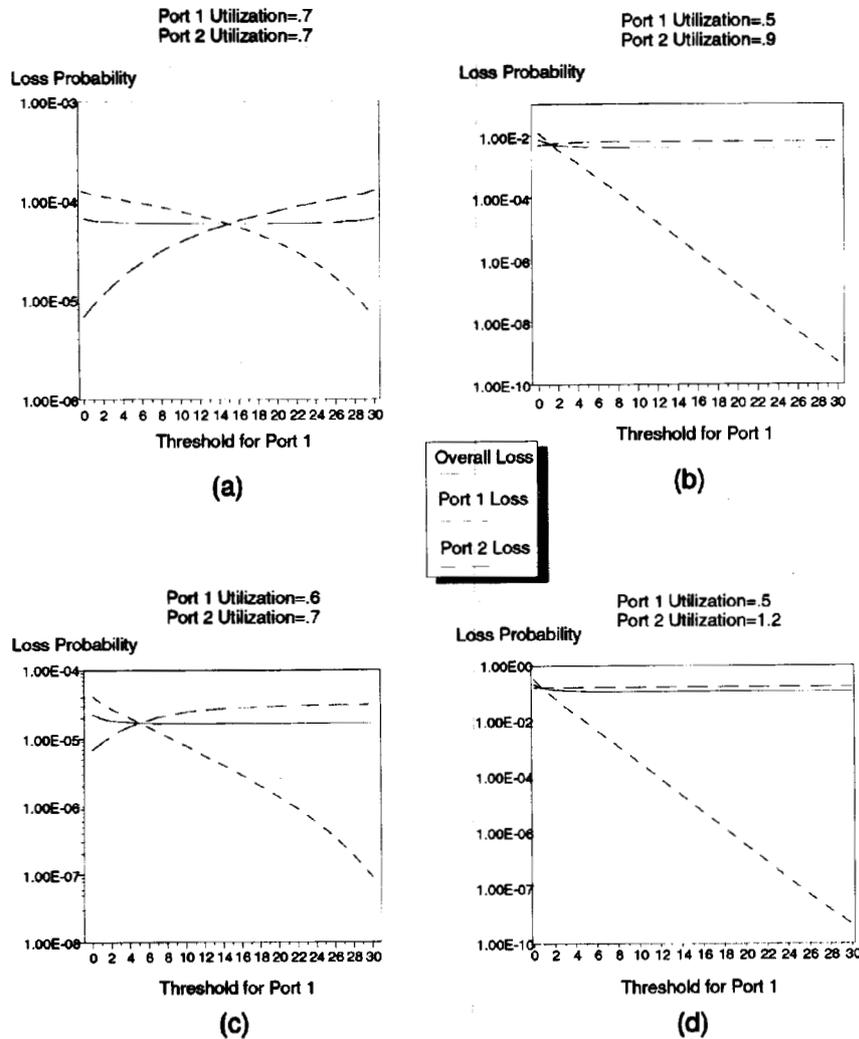
Fig. 2. Performance of a POT policy as the threshold varies.

the other port is larger than but close to 1. Even in this case, as the buffer size increases, the approximation does become more accurate. This is illustrated in Fig. 4, where we plotted the ratio of the optimal to the approximate threshold as the buffer size increases. This was done for two sets of port utilizations. In the first set $(\rho_1, \rho_2) = (.8, 1.1)$, while in the second, $(\rho_1, \rho_2) = (.9, 1.1)$. We see that in both cases, while the approximation is poor for small buffers, it continually improves as the buffer size increases. This leads us to conjecture the correctness of the asymptotic behavior of the approximation expressed in (14).

APPENDIX
PROOFS OF LEMMAS

Proof of Lemma 1: Assume that Properties 1–5 of Lemma 1 hold for a function $f \in \mathcal{F}$.

Proof of Property 1:

Case 1 $(k_1 + k_2 \leq B - 2)$: From properties 1-2 of the function f, we have

$$Tf(k_1, k_2) = \mu_1 1\{k_1 > 0\} + \mu_2 1\{k_2 > 0\}$$
$$+ \beta\mu_1 f((k_1 - 1)^+, k_2) + \beta\mu_2 f(k_1, (k_2 - 1)^+)$$
$$+ \beta\lambda_1 f(k_1 + 1, k_2) + \beta\lambda_2 f(k_1, k_2 + 1) \quad (18)$$

$$Tf(k_1 + 1, k_2)$$
$$= \mu_1 + \mu_2 1\{k_2 > 0\} + \beta\mu_1 f(k_1, k_2)$$
$$+ \beta\mu_2 f(k_1 + 1, (k_2 - 1)^+)$$
$$+ \beta\lambda_1 f(k_1 + 2, k_2)$$
$$+ \beta\lambda_2 f(k_1 + 1, k_2 + 1). \quad (19)$$

Buffer Size = 50
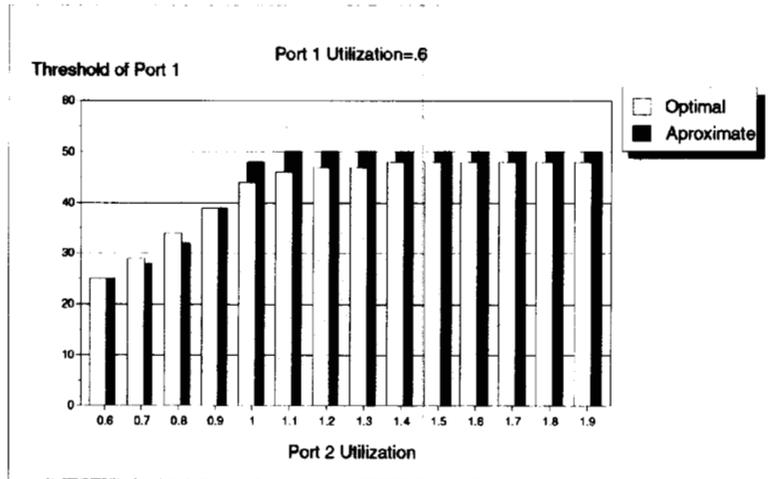
(a)
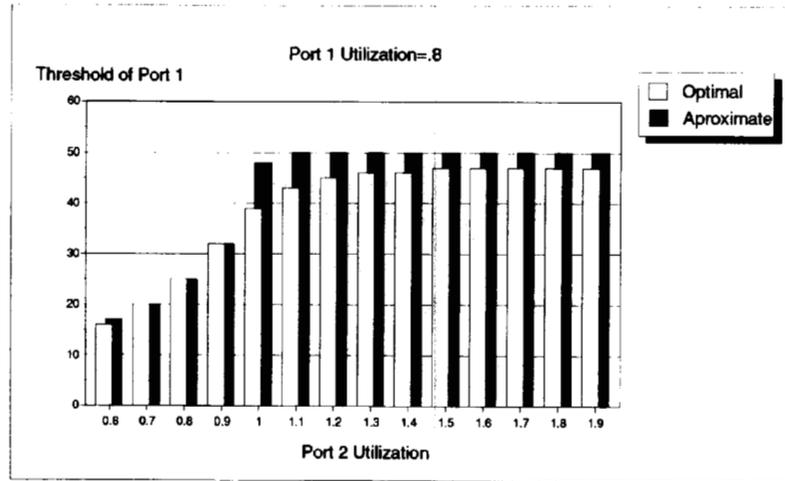
(b)

Fig. 3. Comparison of approximate and optimal threshold.

Comparing (18) and (19) term by term and using Properties 1–2 of the function f, we have that

$$\mu_1 1\{k_1 = 0\} \leq Tf(k_1 + 1, k_2) - Tf(k_1, k_2)$$
$$\leq \mu_1 1\{k_1 = 0\} + \beta(1 - \mu_1 1\{k_1 = 0\})$$

from which we have $0 \leq Tf(k_1 + 1, k_2) - Tf(k_1, k_2) \leq 1$.

Case 2 $(k_1 + k_2 = B - 1)$: From Properties 1–2 of the function f, we have

$$Tf(k_1 + 1, k_2) = \mu_1 + \mu_2 1\{k_2 > 0\} + \beta\mu_1 f(k_1, k_2)$$
$$+ \beta\mu_2 f(k_1 + 1, (k_2 - 1)^+)$$
$$+ \beta\lambda_1 \max\{f(k_1 + 1, k_2), f(k_1 + 2, k_2 - 1)\}$$
$$+ \beta\lambda_2 \max\{f(k_1 + 1, k_2), f(k_1, k_2 + 1)\}$$

$$(20)$$

and $Tf(k_1, k_2)$ is the same as (18). From Properties 1–2 of the function f, we have $0 \leq f(k_1 + 2, k_2 - 1) - f(k_1 + 1, k_2 - 1) \leq 1$ and $-1 \leq f(k_1 + 1, k_2 - 1) - f(k_1 + 1, k_2) \leq 0$, from which we have

$$0 \leq \max\{f(k_1 + 1, k_2), f(k_1 + 2, k_2 - 1)\} - f(k_1 + 1, k_2) \leq 1$$
$$(21)$$

and similarly we have

$$0 \leq \max\{f(k_1 + 1, k_2), f(k_1, k_2 + 1)\} - f(k_1, k_2 + 1) \leq 1.$$
$$(22)$$

Comparing (18) and (20) term by term and using Properties 1–2 of the function f and (21)–(22), we have $0 \leq Tf(k_1 + 1, k_2) - Tf(k_1, k_2) \leq 1$.

Proof of Property 2: similar to the proof of Property 1.

Ratio of Optimal to Approximate Thresholds
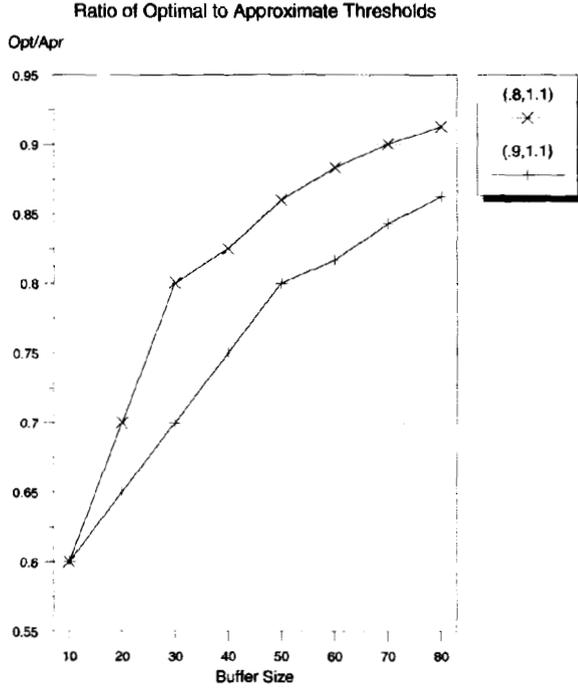
Fig. 4. Accuracy of the approximation as the buffer size increases.

Proof of Property 3: follows directly by comparing $Tf(k_1+1,k_2) - Tf(k_1,k_2)$ with $Tf(k_1,k_2) - Tf(k_1-1,k_2)$ term by term as in the proof of Property 1. The proof of Property 4 follows in a similar way.

Proof of Property 5:

Case 3 $(k_1 + k_2 \leq B - 1)$: From Properties 1–2 of the function f, we have

$$Tf(k_1 + 1, k_2 - 1) - Tf(k_1, k_2) = \mu_2(1\{k_2 > 1\} - 1)$$
$$+ \beta\mu_1(f(k_1, k_2 - 1) - f(k_1 - 1, k_2))$$
$$+ \beta\mu_2(f(k_1 + 1, (k_2 - 2)^+) - f(k_1, k_2 - 1))$$
$$+ \beta\lambda_1(f(k_1 + 2, k_2 - 1) - f(k_1 + 1, k_2))$$
$$+ \beta\lambda_2(f(k_1 + 1, k_2) - f(k_1, k_2 + 1)) \qquad (23)$$

$$Tf(k_1, k_2) - Tf(k_1 - 1, k_2 + 1) = \mu_1(1 - 1\{k_1 > 1\})$$
$$+ \beta\mu_1(f(k_1 - 1, k_2) - f((k_1 - 2)^+, k_2 + 1))$$
$$+ \beta\mu_2(f(k_1, k_2 - 1) - f(k_1 - 1, k_2))$$
$$+ \beta\lambda_1(f(k_1 + 1, k_2) - f(k_1, k_2 + 1))$$
$$+ \beta\lambda_2(f(k_1, k_2 + 1) - f(k_1 - 1, k_2 + 2)). \qquad (24)$$

For $k_1, k_2 \geq 2$ it follows directly by comparing (23) with (24) term by term and using Property 5 of the function f that the LHS of (23) is less or equal to the LHS of (24). For $k_1 = 1, k_2 \geq 2$, comparing (23) with (24) term by term we have that, the terms multiplied by $\beta\mu_2, \beta\lambda_1, \beta\lambda_2$ in (23) are less equal to the corresponding terms in (24) by Property 5. We still have to show that

$$\beta\mu_1(f(1, k_2 - 1) - f(0, k_2)) \leq \mu_1 + \beta\mu_1(f(0, k_2)$$
$$- f(0, k_2 + 1)). \qquad (25)$$

From Property 4 we have $f(0, k_2 + 1) - f(0, k_2) \leq f(0, k_2) - f(0, k_2 - 1)$, and from Property 1 we have $f(1, k_2 - 1) \leq 1 + f(0, k_2 - 1)$. From the last two inequalities we have $f(1, k_2 - 1) - f(0, k_2) \leq 1 + f(0, k_2) - f(0, k_2 + 1)$, from which the inequality (25) follows. The proof is similar for the case $k_1 \geq 2, k_2 = 1$. *case 2* $(k_1 + k_2 = B)$: The inequality $Tf(k_1 + 1, k_2 - 1) - Tf(k_1, k_2) \leq Tf(k_1, k_2) - Tf(k_1 - 1, k_2 + 1)$ is easily shown to hold by comparing these expressions term by term and using Property 5 of the function f on the lines $k_1 + k_2 = B - 1$ and $k_1 + k_2 = B$. For the terms involving λ_1, λ_2 the above inequality holds for all cases of the location of the maximum of the function f on the line $k_1 + k_2 = B$. $\square$

Proof of Lemma 2: Assume that Lemma 2 holds for a function $f \in \mathcal{F}$. *case 1* $(k_1 + k_2 \leq B - 1)$: From Properties 1–2 of the function f, we have

$$Tf(k_1, k_2) = \mu(1\{k_1 > 0\} + 1\{k_2 > 0\})$$
$$+ \beta\mu(f((k_1 - 1)^+, k_2) + f(k_1, (k_2 - 1)^+)$$
$$+ \beta\lambda_1 f(k_1 + 1, k_2) + \beta\lambda_2 f(k_1, k_2 + 1)$$
$$\qquad (26)$$

$$Tf(k_2, k_1) = \mu(1\{k_1 > 0\} + 1\{k_2 > 0\})$$
$$+ \beta\mu(f((k_2 - 1)^+, k_1) + f(k_2, (k_1 - 1)^+)$$
$$+ \beta\lambda_1 f(k_2 + 1, k_1) + \beta\lambda_2 f(k_2, k_1 + 1). $$
$$\qquad (27)$$

Comparing (26) with (27) term by term and using Lemma 2 for the function f, we have that $Tf(k_1, k_2) \geq Tf(k_2, k_1)$. *case 2* $(k_1 + k_2 = B)$: Here we have

$$Tf(k_1, k_2) = \mu(1\{k_1 > 0\} + 1\{k_2 > 0\})$$
$$+ \beta\mu(f((k_1 - 1)^+, k_2) + f(k_1, (k_2 - 1)^+)$$
$$+ \beta\lambda_1 \max\{f(k_1, k_2), f(k_1 + 1, k_2 - 1)\}$$
$$+ \beta\lambda_2 \max\{f(k_1, k_2), f(k_1 - 1, k_2 + 1)\}$$
$$\qquad (28)$$

$$Tf(k_2, k_1) = \mu(1\{k_1 > 0\} + 1\{k_2 > 0\})$$
$$+ \beta\mu(f((k_2 - 1)^+, k_1) + f(k_2, (k_1 - 1)^+)$$
$$+ \beta\lambda_1 f(k_2, k_1) + \beta\lambda_2 f(k_2 - 1, k_1 + 1)$$
$$\qquad (29)$$

where in the last two terms of (29) we used Lemma 2 and Property 5 of Lemma 1 (concavity on the line $k_1 + k_2 = B$) for the function f. Applying Lemma 2 twice to the function f, we obtain the following two inequalities

$$\max\{f(k_1, k_2), f(k_1 + 1, k_2 - 1)\} - f(k_2, k_1)$$
$$\geq f(k_1 + 1, k_2 - 1) - f(k_1, k_2)$$
$$\geq f(k_2 - 1, k_1 + 1)$$
$$- \max\{f(k_1, k_2), f(k_1 - 1, k_2 + 1)\}. \qquad (30)$$

From (30) and the inequality $\lambda_1 \geq \lambda_2$ we have $Tf(k_1, k_2) \geq Tf(k_2, k_1)$. $\square$

Proof of Lemma 3: Assume that Properties 1–4 of Lemma 3 hold for a function $f \in \mathcal{F}$.

Proof of Property 1: Similar to the proof of Property 1 of Lemma 1.

Proof of Property 2: Consider a state $\overline{k}$ and define the state $\overline{k}_{ij}$ as the state $\overline{k}$ with k_i and k_j interchanged. It is enough to show that $Tf(\overline{k}) = Tf(\overline{k}_{ij})$. Compare the dymanic programming equations (15) of $Tf(\overline{k})$ and $Tf(\overline{k}_{ij})$ term by term. Clearly, the first terms are equal. Using Property 2 $(f((\overline{k} - e_i)^+) = f((\overline{k}_{ij} - e_j)^+))$ it follows that the second terms are equal as well. Using Property 2 again one can easily show that the third terms are also equal.

Proof of Property 3: Consider a state $\overline{k}$ and assume that the buffer is full, i.e., $\sum_{i=1}^{N} k_i = B$. Fix some $i, j (i \neq j)$, and assume that $k_i \geq k_j (\geq 1)$. We show that $Tf(\overline{k}) - Tf(\overline{k} + e_i - e_j) \geq 0$. From (15) we have that

$$
\begin{aligned}
Tf(\overline{k}) - Tf(\overline{k} + e_i - e_j) &= \mu[1 - 1\{k_j > 1\}] \\
&+ \beta\mu \sum_{\substack{l=1 \\ l \neq i,j}}^{N} [f((\overline{k} - e_l)^+) - f((\overline{k} + e_i - e_j - e_l)^+)] \\
&+ \beta\mu[f(\overline{k} - e_i) - f(\overline{k} - e_j)] \\
&+ \beta\mu[f(\overline{k} - e_j) - f((\overline{k} + e_i - 2e_j)^+)] \\
&+ \beta\lambda \sum_{l=1}^{N} [\max_{1 \leq p_1 \leq N}\{f(\overline{k} + e_l - e_{p_1})\} \\
&- \max_{1 \leq p_2 \leq N}\{f(\overline{k} + e_i - e_j + e_l - e_{p_2})\}].
\end{aligned}
\tag{31}
$$

Each of the second and the third terms of (31) is greater than or equal to zero (gez) by Properties 2 and 3. For $k_j > 1$ the first term of (31) vanishes and the fourth term is gez by Property 3. For $k_j = 1$ the first term equals μ and for the fourth term we have by Property 1 (boundedness) that it is greater or equal to $-\beta\mu$, hence the sum of the first and the fourth terms is greater than zero for $\beta < 1$. This shows that the sum of the first four terms is gez.

Then, we show that the last term of (31) is gez. We show that for each l the term that appears in the sum is gez. Fix l and denote by p_1^* and p_2^* the indices that give the maximum value in the first and the second maximum terms, respectively. We break ties in the maximum terms by choosing p_1^* and p_2^* such that $k_{p_1^*}$ and $k_{p_2^*}$ are the maximal components of the vectors $\overline{k}$ and $\overline{k} + e_i - e_j$, respectively (obviously, this convention doesn't affect the validity of the proof). Next, we consider all cases of the location of the maximum indices p_1^* and p_2^* and show that each term that appears in the sum of the last term of (31) is gez. First, consider the case $p_1^* = j$. Then, by Property 4 of the function f it must be that k_j is the maximal component of the vector $\overline{k}$ (there may be other components with the same value), and since $k_i \geq k_j$ we have that $k_i = k_j$. Then, applying Property 4 of the function f to the second maximum term we have that $p_2^* = i$, which implies that the two maximum terms are equal. Next, consider the case $p_1^* = i$ (and assume that $k_i > k_j$, otherwise we get the previous case). Then, by Property 4 of the function f it must be that k_i is the maximal component of the vector $\overline{k}$, and hence p_2^* is equal to

i (by Property 4 again). Then, the difference between the two maximum terms is gez for $i = l$ or $i, j \neq l$ by Property 3 of the function f and for $i \neq l = j$ by Properties 2 and 3 of the function f. Finally, we consider the case $p_1^* \neq i, j$ (and assume that $k_{p_1^*} > k_i$, otherwise we get the previous case). Then, by Property 4 of the function f it must be that $k_{p_1^*}$ is the maximal component of the vector $\overline{k}$, and hence p_2^* is equal to p_1^* (by Property 4). The proof then follows exactly as the previous case.

Consider the case $k_i < k_j$ and define $\hat{k} \triangleq \overline{k} + e_i - e_j$. Then, for $k_j > k_i + 1$ we have $Tf(\hat{k}) \geq Tf(\hat{k} + e_i - e_j)$ by the first part of Property 3, and for $k_j = k_i + 1$ we have $Tf(\hat{k}) = Tf(\hat{k} + e_i - e_j)$ by Property 2. This completes the proof of Property 3 for the case of full buffer. The proof for the case of a nonfull buffer is simpler and follows in a similar way.

Proof of Property 4: This property follows directly from Property 3.   □

REFERENCES

[1] D. M. Bertsekas, *Dynamic Programming, Deterministic and Stochastic Models.* Englewood Cliffs, NJ: Prentice-Hall, 1987.

[2] G. J. Foschini and B. Gopinath, "Sharing memory optimally," *IEEE Trans. Commun.*, vol. COM-31, no. 3, Mar. 1983.

[3] F. Kamoun and L. Kleinrock, "Analysis of shared finite storage in a computer network node environment under general traffic conditions," *IEEE Trans. Commun.*, vol. COM-28, no. 7, July 1980.

[4] H. Kroner, G. Hebuterne, P. Boyer, and A. Gravey, "Priority management in ATM switching nodes," *IEEE Trans. Commun.*, vol. 9, no. 3, pp. 418–427, Apr. 1991.

[5] S. Lippman, "Applying a new device in the optimization of exponential queueing systems," *Op. Res.*, vol. 23, pp. 687–710, 1975.

[6] S. Ross, *Introduction to Stochastic Dynamic Programming.* New York: Academic.

[7] R. F. Serfozo, "An equivalence between continous and discrete time Markov decision processes," *Op. Res.*, vol. 27, pp. 616–620, 1979.

[8] S. X. Wei, E. J. Coyle, and M. T. Hsiao, "An optimal buffer management policy for high-performance packet switching," in *Proc. IEEE GLOBECOM'91*, vol. 2, Dec. 1991, pp. 924–928.

Israel Cidon (M'85–SM'90) received the B.Sc. (summa cum laude) and the D.Sc. degrees from the Technion–Israel Institute of Technology, Haifa, in 1980 and 1984, respectively, both in electrical engineering.

From 1984 to 1985, he was with the faculty of the Electrical Engineering Department at the Technion. In 1985, he joined the IBM T. J. Watson Research Center, Yorktown Heights, NY, where he was a Research Staff Member and a Manager of the Network Architectures and Algorithms group, involved in various broadband networking projects such as the Paris/Planet Gigabit testbed, the Orbit Gigabit LAN, and the IBM Broad-Band Networking Services (BBNS) architecture. Recently, he joined Sun Microsystems Labs in Mountains View, CA, as Manager of High-Speed Networking, working in various ATM projects including Openet, an open ATM network control platform. Since 1990, he has also been with the Department of Electrical Engineering at the Technion. His research interests are in high-speed wide and local area networks, distributed network algorithms, network performance and wireless networks.

In 1989 and 1993, Dr. Cidon received the IBM Outstanding Innovation Award for his work on the PARIS high-speed network and topology update algorithms, respectively. He currently serves as an Editor for the IEEE/ACM TRANSACTIONS ON NETWORKING. Previously, he served as the Editor for Network Algorithms for the IEEE TRANSACTIONS ON COMMUNICATIONS and as a guest editor for Algorithmica.

Leonidas Georgiadis (M'86–SM'95) received the Diploma degree in electrical engineering from the University of Thessaloniki, Greece, in 1979, and the M.S. and Ph.D degrees from the University of Connecticut, Storrs, in 1981 and 1986, respectively, both in electrical engineering .

From 1981 to 1983 was with the Greek army. From 1986 to 1987, he was Research Assistant Professor at the University of Virginia, Charlottesville. In 1987, he joined the IBM T. J. Watson Research Center, Yorktown Heights, NY, where he is presently a Research Staff Member in the Broadband Networking group. His interests are include high-speed network design, scheduling, wireless communications, modeling and performance analysis.

In 1992, Dr. Georgiadis received the IBM Outstanding Innovation Award for his work on Goal-Oriented Workload Management for Multi-Class Systems. He is a member of IEEE Communications Society.

Asad Khamisy received the B.Sc. degree in computer engineering and the M.Sc. and D.Sc. degrees in electrical engineering from the Technion–Israel Institute of Technology, Haifa, in 1988, 1991, and 1994, respectively.

During 1994, he was on the faculty of the Electrical Engineering Department at the Technion. Currently, he is with Sun Microsystems Labs, Mountain View, CA. His research interests include high-speed local and wide area networks, and performance analysis.

Roch Guérin (S'85–M'86–SM'91) received the "Diplôme d'Ingénieur" from the École Nationale Supérieure des Télécommunications, Paris, France, in 1983, and the M.S. and Ph.D. degrees from the California Institute of Technology, Pasadena, in 1984 and 1986, respectively, both in electrical engineering.

Since August 1986 he has been with the IBM T. J. Watson Research Center, Yorktown Heights, NY, where he now manages the Network System Design group in the Broadband Networking department. His current research interests are in the areas of modeling, architecture, and Quality-of-Service issues in high-speed networks. In 1994, he received an IBM outstanding Innovation Award for his work on traffic management for the Networking Broadband Services architecture.

Dr. Guérin is a member of Sigma Xi and the IEEE Communications Society, and Editor for the IEEE/ACM TRANSACTIONS ON NETWORKING. He was an Editor for the IEEE TRANSACTIONS ON COMMUNICATIONS and the *IEEE Communications Magazine*.