TR-87-15

Extension of the Optimality of the Threshold Policy in Heterogeneous Multiserver Queuing Systems

by

Ioannis Viniotis
Anthony Ephremides

revised
Oct. '86

EXTENSION OF THE OPTIMALITY OF THE THRESHOLD POLICY IN HETEROGENEOUS MULTISERVER QUEUEING SYSTEMS

by

Ioannis Viniotis , Anthony Ephremides

University of Maryland

College Park, MD 20742

ABSTRACT

We extend the validity of some results on the optimal control of two-server queueing models with service times of unequal distribution, operating in continuous or discrete time. The distribution of arrivals can be arbitrary subject to some conditions. Both discounted and long run average costs are considered. Dynamic programming and probabilistic arguments are the key tools used to establish the assertion that the optimal policy is of threshold type, i.e. the slower server should be utilized only when the queue length exceeds a certain threshold value.

I. INTRODUCTION AND BACKGROUND

The queueing system considered is shown in figure 1. The motivation for studying such a system comes from problems of dynamic routing in computer systems or communication networks. For example, the system may model a node in a communication network where the "customers" represent messages and the "servers" represent communication lines (with different delays) over which messages are sent. Customers arrive in a single stream (not necessarily Poisson) and immediately join the queue. The queue is served by two **nonidentical** servers. We shall assume that the service times either have the same exponential distribution with different means or have entirely different distributions (one exponential and one r-stage Erlangian, again with different means). We shall impose certain conditions on the mean service and arrival times, to ensure stability of the system. We focus primarily on the discrete time case, since it provides a more natural framework in a number of digital communications and computer system applications.

Two cost criteria are considered: discounted and average number of customers in the system. Our objective is to choose the control actions to minimize these costs. In the average cost case, this is equivalent to minimizing the mean time a customer spends in the system.

Multiserver queueing systems are of interest in applications such as multi-processor systems as well as virtual-circuit-switching communication networks. Most of the results obtained so far rely on the existence of optimal stationary policies for fairly general systems as established by Lippman [1]. Larsen [2] has considered a multiserver system with different rate exponential servers and Poisson arrivals. He conjectured the optimality of threshold policies and provided a detailed analysis of their performance. Lin and Kumar [3] proved Larsen's conjecture, using a dynamic programming argument. They also gave a closed form expression for the value of the optimal threshold. Walrand [4] gave a simpler proof of the same result, using a probabilistic argument. The latter approach however, does not facilitate the simple calculation of the threshold.

Hajek [5] considered the case of two interacting nonidentical service stations. His framework is quite general, but somewhat different from ours. He showed optimality of switching type policies. He generalized and unified several earlier results. Bell [6] analyzed related multiserver systems under different cost criteria. Kumar and Walrand [7] studied individually optimal controls of multiserver queueing systems, with general arrival and/or service processes. Hahne [8] , Tsitsiklis [9] and Seidmann and Schweitzer [10] studied the dynamic routing of customers among multiple servers in queueing systems arising in manufacturing networks. Optimality of threshold type policies has been shown in a number of related models ([11], [12]). Here we generalize the earlier work in two directions: we obtain results for a discrete time framework and we assume fairly general arrival processes and service time distributions for the continuous time framework. We use well known techniques (dynamic programming and probabilistic arguments) to establish our results.

The paper is organized as follows: in section II the control problem for the discrete time model is formulated in detail. The continuous time models, namely the $E_r/M/2$ and $M/M,E_r/2$ models, can be simply converted into equivalent discrete time problems ([15]). Equivalence means that for infinite horizon cost criteria the optimal policies for the two formulations coincide. Since the differences in the two formulations are not essential, they will not be presented here in detail. For a full description of the continuous time case, see [13]. The optimality of the threshold policy, for the discounted cost criterion, is shown in section III, using a dynamic programming argument. In section IV, we discuss the average cost criterion for the more general $GI/M/2$ model with different rate servers and with arrival statistics satisfying certain conditions. It is claimed that if an optimal policy exists and is unique, it will be a threshold type policy for any fixed residual interarrival time.

II. CONTROL PROBLEM FORMULATION

Consider the queueing system shown in figure 2. The system operates in discrete time. Arrivals to the system form a Bernoulli stream, with probability p of an arrival in a given time slot. The service time at the i-th server is geometrically distributed with parameter q_i, i.e. the probability that the customer completes service in a given time slot is q_i, when the customer is being served by server i, $i = 1, 2$.Without loss of generality, we assume that $q_1 > q_2$.

Let $x = (x_0, x_1, x_2)$ be the state of the system, where x_0 is the number of the customers in queue and $x_i = 1$ or 0 depending on whether server i , $i = 1, 2$ is busy or not. The state space of the system is $X = \{0, 1, 2, ...\} \times \{0, 1\} \times \{0, 1\}$. Let $\{E_i\}$, $i = 0, ..., 7$ denote the events of possible combinations of arrivals and service completions that may occur during one time slot. Define the operators $Q_i : X \to X$, $i = 0, ..., 7$ as follows:

$$Q_0(x_0, x_1, x_2) = (x_0, x_1, x_2)$$

$$Q_1(x_0, x_1, x_2) = (x_0, x_1, (x_2 - 1)^+)$$

$$Q_2(x_0, x_1, x_2) = (x_0, (x_1 - 1)^+, x_2)$$

$$Q_3(x_0, x_1, x_2) = (x_0, (x_1 - 1)^+, (x_2 - 1)^+)$$

$$Q_4(x_0, x_1, x_2) = (x_0 + 1, x_1, x_2)$$

$$Q_5(x_0, x_1, x_2) = (x_0 + 1, x_1, (x_2 - 1)^+)$$

$$Q_6(x_0, x_1, x_2) = (x_0 + 1, (x_1 - 1)^+, x_2)$$

$$Q_7(x_0, x_1, x_2) = (x_0 + 1, (x_1 - 1)^+, (x_2 - 1)^+)$$

where $n^+ = max(n, 0)$. These operators describe the effects of the events $\{E_i\}$ on the system state.

Define now the following "action" operators, that describe the assignment of customers to the servers:

$$P_h(x_0, x_1, x_2) = (x_0, x_1, x_2) \quad , \quad domP_h = X$$

$$P_1(x_0, x_1, x_2) = (x_0 - 1, 1, x_2) \quad , \quad domP_1 = \{x \in X : x_0 \geq 1, x_1 = 0\}$$

$$P_2(x_0, x_1, x_2) = (x_0 - 1, x_1, 1) \quad , \quad domP_2 = \{x \in X : x_0 \geq 1, x_2 = 0\}$$

$$P_b(x_0, x_1, x_2) = (x_0 - 2, 1, 1) \quad , \quad domP_b = \{x \in X : x_0 \geq 2, x_1 = x_2 = 0\}$$

The P-operators will represent the dispatcher's decisions: P_h denotes holding a customer, while P_1,P_2,P_b denote customer forwarding to server 1,2 or both respectively. Let, therefore, h, 1, 2, b represent mnemonically the above actions.

In each time slot one of eight possible events occurs and the dispatcher may take one of the above mentioned four actions. Let $u \triangleq (u_0, u_1, ..., u_7)$ where $u_i \in \{h, 1, 2, b\}$, $i = 0, ..., 7$ denote a control action and U denote the set of all controls. Let $U(x) \triangleq \{u \in U : Q_i(x) \in \ dom\ P_{u_i},\ i = 0, ..., 7\}$ denote the set of admissible controls when the system state is x.

Our goal is to choose the control actions $u(t)$ so as to minimize

$$E \sum_{t=0}^{\infty} | x(t) | \beta^t$$

where $| x(t) | = x_0 + x_1 + x_2$ denotes the number of customers in system at time t, and β is a discount factor $0 < \beta < 1$.

Let $J^\beta(x)$ denote the optimum cost function, as a function of the initial state $x(0) = x$. It is well known that there always exists an optimal policy which is stationary. Here by a stationary policy we mean any function $\pi : X \to U$ subject to the constraint $\pi(x) \in U(x) \quad \forall x \in X$. When π is adopted as a policy, $u = \pi(x)$ is applied whenever the system is in state x.

For any stationary policy π, let T_π be the dynamic programming operator

$$(T_\pi f)(x) = | x | + \sum_{i=0}^{7} \beta Pr(E_i) f(P_{u_i} Q_i x)$$

where $Pr(E_i)$ is the probability of occurence of event E_i.

It is a standard result in dynamic programming that

$$J^\beta(x) = \min_{u \in U(x)} [\ | x | + \sum_{i=0}^{7} \beta Pr(E_i) J^\beta(P_{u_i} Q_i x)\ ] \tag{A}$$

$$\lim_{n \to \infty} T^{(n)} f(x) = J^\beta(x) \tag{B}$$

where T is the dynamic programming operator, i.e.

$$(Tf)(x) = \min_{\pi} (T_\pi f)(x)$$

and $T^{(n)}f \triangleq T^{(n-1)}(Tf)$.

III. OPTIMALITY OF THRESHOLD POLICIES

We are now going to show that the optimal policy is of threshold type, i.e. the slower server should be utilized only when the number of customers waiting in queue exceeds a certain threshold. This threshold may be $+\infty$ which simply means that the slower server may never be used.

Our goal is to show that the following theorem (theorem 5 of [3]) is true. The proof follows readily from lemmas 1-4.

Theorem: i) There exists an optimal stationary policy which is of threshold type with a threshold $m^* \leq \infty$.

ii) If $J^\beta_{t_i}(x) < J^\beta_{t_{i+1}}(x)$ for some $x \in X$, then $m^* \leq i$.

We first prove some properties of the optimal cost function in lemma 1. Lemma 2 states the (intuitively obvious) fact that the faster server should be always kept busy, whenever this is possible. Lemma 3 states the fact that the policy iteration algorithm produces as its limit an optimal policy. Finally, in the crucial lemma 4 we show that when we apply the policy iteration algorithm to a threshold policy, we get as an improvement a threshold policy again. Moreover the new policy's threshold is at most 1 unit more than the threshold of the original policy. All these facts combined, prove the validity of the theorem stated above.

Since the proofs parallel those of [3], they will not be shown here in full details. For a complete description, see [13].

Lemma 1. The optimal cost function satisfies the following properties:

i) $J^\beta(P_1 x) \leq J^\beta(P_h x) \quad , \quad \forall x \in domP_1$

ii) $J^\beta(P_1 x) \leq J^\beta(P_2 x) \quad , \quad \forall x \in domP_1 \bigcap domP_2$

Proof: We have to show that

$$| P_1 x | + \sum_{i=0}^{7} \beta Pr(E_i) \min_{u_i} J^\beta(P_{u_i} Q_i P_1 x) \ \leq | P_h x | + \sum_{i=0}^{7} \beta Pr(E_i) \min_{u_i} J^\beta(P_{u_i} Q_i P_h x) \quad (C)$$

$$| P_1 x | + \sum_{i=0}^{7} \beta Pr(E_i) \min_{u_i} J^\beta(P_{u_i} Q_i P_1 x) \ \leq | P_2 x | + \sum_{i=0}^{7} \beta Pr(E_i) \min_{u_i} J^\beta(P_{u_i} Q_i P_2 x) \quad (D)$$

Consider any function f which satisfies

$$f(P_1 x) \leq f(P_h x) \tag{A1}$$

$$f(P_1 x) \leq f(P_2 x) \tag{A2}$$

$$f(x) \leq f(y) \qquad x \leq y \tag{A3}$$

It can be easily seen (by comparing corresponding terms) that the differences $Tf(P_1x) - Tf(P_hx)$, $Tf(P_1x) - Tf(P_2x)$,$Tf(x) - Tf(y)$ for $x \leq y$ are all nonpositive. Thus Tf satisfies A1-A3 as well. Inductively $T^{(n)}f$ satisfies A1-A3 and thus $J^\beta = \lim_{n\to\infty} T^{(n)}f$ possesses the desired properties (C) and (D).

Lemma 2: Whenever the faster server is idle, it is optimal to start serving a customer, if one is waiting for service.

Proof: See [3].

The importance of this lemma is that it enables us to restrict the set of admissible controls to

$$\overline{U}(x) = \{u \in U(x) : (P_{u_i} Q_i x)_1 = 1 \ , \ i = 4-7 \ ; \ (P_{u_i} Q_i x)_1 = 1 \ , \ i = 0-3 \quad \text{if} \quad x_0 \geq 1\}$$

i.e. to a set of controls where the faster server is always kept busy. Thus the only decision we need to make is whether to utilize server 2 or not.

Lemma 3. Let $\{\pi_n\}_0^\infty$ denote a sequence of policies generated by the policy iteration algorithm. Let $\pi^*(x) = \lim_{k\to\infty} \pi_{n_k}(x)$, $\forall x \in X$. Then π^* is optimal.

Proof: See [3].

Define now the following operator $F_m : X \to X$

$$F_m(x) = \begin{cases} P_h(x) & \text{if} \quad x \in \{x_0 = 0\} \bigcup \{x_1 = 1, x_2 = 1\} \bigcup \{x_1 = 1, x_2 = 0, x_0 \leq m\} \\ P_1(x) & \text{if} \quad x \in \{x_1 = 0, x_2 = 1, x_0 \geq 1\} \bigcup \{x_1 = 0, x_2 = 0, 1 \leq x_0 \leq m+1\} \\ P_2(x) & \text{if} \quad x \in \{x_1 = 1, x_2 = 0, x_0 > m\} \\ P_b(x) & \text{if} \quad x \in \{x_1 = 0, x_2 = 0, x_0 > m+1\} \end{cases}$$

Then a threshold policy t_m, with threshold m (i.e. a policy which utilizes the idle server 2 iff the number of customers in queue is strictly larger than m) is defined as

$$t_m(x) = (F_m(Q_0 x), F_m(Q_1 x), ..., F_m(Q_7 x))$$

If T_{t_m} denotes the dynamic programming operator associated with t_m, and $J^{\beta}_{t_m}(x)$ denotes the (β-discounted) cost obtained by using t_m, then

$$T_{t_m} J^{\beta}_{t_m}(x) = J^{\beta}_{t_m}(x) = | x | + \sum_{i=0}^{7} \beta Pr(E_i) J^{\beta}_{t_m}(F_m Q_i x)$$

Lemma 4: For any finite $i \geq 0$ there exists a j, $0 \leq j \leq i+1$ such that $T_{t_j} J^{\beta}_{t_i} = T J^{\beta}_{t_i}$. That is, there exists a policy t_j which has a threshold not larger than $i+1$ and which achieves the minimum cost over all policies. Thus the policy iteration algorithm, when applied to a threshold policy, produces a threshold policy again.

Proof: Define the cost differences

$$h_0 = J^{\beta}_{t_i}(0,1,0) - J^{\beta}_{t_i}(0,0,1)$$

$$h_k = J^{\beta}_{t_i}(k,1,0) - J^{\beta}_{t_i}(k-1,1,1) \qquad k \geq 1$$

and examine when the differences change sign.Note that if $h_l \geq 0$, $\forall\ l \geq j$ and $h_l < 0$, $\forall\ l \leq j-1$,then the new policy will be a threshold policy with threshold $j-1$. It is shown in [13] that for $i \geq 3$

$$h_k \geq 0 \quad , \quad k \geq i+2 \tag{1}$$

$$h_k \geq \beta Pr(E_2) h_{k-1} \quad , \quad k = i+1 \tag{2}$$

$$\beta Pr(E_2)\{h_k - h_{k-1}\} \geq -(1 - \beta Pr(E_2) - \beta Pr(E_6)) h_k \quad , \quad k = i \tag{3}$$

$$\beta Pr(E_2)\{h_k - h_{k-1}\} \geq \beta Pr(E_4)\{h_{k+1} - h_k\} -$$

$$(1 - \beta Pr(E_0) - \beta Pr(E_2) - \beta Pr(E_4) - \beta Pr(E_6)) h_k \ , \ 1 \leq k \leq i-1 \tag{4}$$

Similar relations hold for the special cases $i = 0, 1, 2$. The detailed analysis of relations (1) - (4) is carried over in [13], where it is shown that the differences do change sign for some $0 \leq j \leq i+1$. Thus the proof of the lemma is complete.

Note that in principle, the value of the optimal threshold could be computed using the policy iteration algorithm: Suppose the optimal policy is t_j for some j. Then if we start the policy iteration algorithm with policy t_0, we reach the optimal policy in exactly j iterations, since each iteration will increase the threshold by one. However, we would have to compute $J_{t_i}^{\beta}(x)$, which is very difficult, if not impossible. The situation is somewhat different in the case of the average cost criterion, where a numerical calculation of the optimal cost function is possible and thus the threshold value can be computed or at least approximated.

A. The $M/M,E_r/2$, $E_r/M/2$ models

We shall now briefly formulate the control problem and discuss the optimality of the threshold policy for the models $M/M,E_r/2$ and $E_r/M/2$. Again for a detailed treatment of the models see [13].

1. The $M/M,E_r/2$ model

Consider the queueing system shown in figure 3. Customers arrive in a single Poisson stream of rate λ. Server 1 is exponential, with rate μ_1. Server 2 is r-stage Erlangian with rate (per stage) μ_2. We shall assume that $\mu_1 > \mu_2$. Note that Erlangian service-time models the case of general consecutive service "stages", with the restriction that a customer cannot enter the first stage of service until the preceding one completes the last stage. In a distributed database system, this could be the case if the customer (i.e. a transaction to be processed) requests and locks either 1 or r resources. This can happen when the system operates in a two phase commitment protocol to achieve atomicity of transactions. It is useful to study such a model for another reason as well: we can approximate an arbitrary service distribution by Erlangian ones, choosing r appropriately. When $r = 1$, this model reduces to the $M/M/2$ model studied in [2] and [3]. We have not been able to successfully analyze either the $M/E_r,M/2$ model, i.e. the one in which the faster server is Erlangian, or the somewhat more general $M/M,E_r/2$ model with $\mu_1 > \frac{\mu_2}{r}$.

As we have mentioned before, the continuous time problem can be converted to a discrete time equivalent one. Let $x = (x_0, x_1, x_2)$ denote the state of the system just before an arrival or a departure occurs, (i.e. at t^-). We may define arrival, departure and assignment operators A, D_1, D_2, P_h, P_1, P_2, P_b, respectively, in a slightly modified way than it was done in [3]. Here we have to take into account that a departure from stage i at the slower server 2 means transition to the next stage $i-1$, while departure from stage 1 is the "real" departure of the customer (i.e. completion of his entire service).

The proof of lemmas 1-4 has a structure similar to those in [3], with the understanding that now $X = \{0,1,2,\ldots\} \times \{0,1\} \times \{0,1,\ldots,r\}$ and that the domains of the various operators and the state transitions become more complicated. Lemmas 2 and 3 remain unchanged, as does the proof of the main result, namely theorem 5 in [3].

Since all the necessary details can be found in [13], we shall only describe the set of threshold policies in this model and outline the proof and its differences with respect to lemma 4, which is the core of the proof of the main result. Notice that the dynamic programming equation for this model becomes

$$J^\beta(x) \;=\; \min_{u \in U(x)} \Big[\, |x| \;+\; \beta\lambda J^\beta(P_{u_0}Ax) \;+\; \beta\mu_1 J^\beta(P_{u_1}D_1x) \;+\; \beta\mu_2 J^\beta(P_{u_2}D_2x) \Big]$$

where $u = (u_0, u_1, u_2)$, $u_i \in \{h, 1, 2, b\}$ for $i = 0,1,2$ is the control action and $|x|$ denotes the number of customers in the system.

Let us define by $F_m : X \to X$ the following operator:

$$F_m(x) \;=\; \begin{cases} P_h(x) & \text{if} \quad x \in \{x_0 = 0\} \cup \{x_1 = 1, x_2 \neq 0\} \cup \{x_1 = 1, x_2 = 0, x_0 \leq m\} \\ P_1(x) & \text{if} \quad x \in \{x_1 = 0, x_2 \neq 0, x_0 \geq 1\} \cup \{x_1 = 0, x_2 = 0, 1 \leq x_0 \leq m+1\} \\ P_2(x) & \text{if} \quad x \in \{x_1 = 1, x_2 = 0, x_0 > m\} \\ P_b(x) & \text{if} \quad x \in \{x_1 = 0, x_2 = 0, x_0 > m+1\} \end{cases}$$

Then a threshold policy t_m, with threshold m, is defined as

$$t_m(x) \;=\; (F_m(Ax), F_m(D_1x), F_m(D_2x))$$

and if $J^\beta_{t_m}(x)$ denotes its cost, we have

$$T_{t_m} J^\beta_{t_m}(x) = J^\beta_{t_m}(x) = |x| + \beta\lambda J^\beta_{t_m}(F_mAx) + \beta\mu_1 J^\beta_{t_m}(F_mD_1x) + \beta\mu_2 J^\beta_{t_m}(F_mD_2x)$$

Lemma 4: For any finite $i \geq 0$, there exists some j such that $0 \leq j \leq i+1$ and

$$T_{t_j} J^\beta_{t_i} \;=\; T J^\beta_{t_i}$$

Proof: Once again, we define the cost differences

$$h_0 \;=\; J^\beta_{t_i}(0,1,0) \;-\; J^\beta_{t_i}(0,0,r)$$

$$h_k \;=\; J^\beta_{t_i}(k,1,0) \;-\; J^\beta_{t_i}(k-1,1,r)$$

and determine when the differences change sign. It is shown in [13] that for $i \geq 3,\ k \geq 2$

$$h_k \geq 0 \ , \ \ k \geq i+2$$

$$h_k \geq \beta\mu_1 h_{k-1} \ , \ \ k = i+1$$

$$h_k \geq \frac{\beta\mu_1}{1-\beta+\beta\mu_1+\beta\lambda} h_{k-1} \ , \ \ k = i$$

$$-(1-\beta)h_k + \beta\lambda(h_{k+1}-h_k) \leq \beta\mu_1(h_k - h_{k-1}) \ , \ \ 2 \leq k \leq i-1$$

and that similar expressions hold for the boundary cases $k=1,\ i=0,1,2$. It was also shown that $\{h_k\}$ changes sign for some $j \leq i+1$ and thus t_j is also a threshold policy.

2. The $E_r/M/2$ model

Consider now the queueing system shown in figure 4. Each customer arrives at the system in r stages, $r \geq 1$. A customer arrival is assumed complete (and thus the customer can be forwarded for service) if all r stages of arrival have been completed. The overall arrival rate is $\lambda = \frac{\lambda'}{r}$,where λ' is the arrival rate per stage. The two servers are exponential, with rates $\mu_1 > \mu_2$. Each server picks up a bulk of r stages (one customer) at a time. In a communication network this could be the case when the customer (a message) arrives in parts (r packets let's say) in a network node and is served as a whole when all parts have arrived. Also this model is a better approximation of a physical system than the Poisson arrival model.

The formulation of this problem is analogous to the previous ones with some differences in notation and proofs, which take into account the fact that now an arrival is completed only after r stage transitions. For example, x_0 now represents the number of stages of arrival in the system, where we understand that an "arrived" customer corresponds to r stages (see [13] for details). We shall only give here the threshold policy definition for this model, since it is that threshold definition that is primarily responsible for the difference between the models.

Define operator $F_m:\ X \to X$, where $X = \{0,1,2\ldots\} \times \{0,1\} \times \{0,1\}$ as

$$F_m(x) = \begin{cases} P_h(x) & \text{if } x \in \{x_0 < r\} \cup \{x_1 = 1, x_2 = 1\} \cup \{x_1 = 1, x_2 = 0,\ x_0 \leq mr + r - 1\} \\ P_1(x) & \text{if } x \in \{x_1 = 0, x_2 = 1, x_0 \geq r\} \cup \{x_1 = x_2 = 0, r \leq x_0 \leq (m+2)r - 1\} \\ P_2(x) & \text{if } x \in \{x_1 = 1, x_2 = 0, x_0 \geq (m+1)r\} \\ P_b(x) & \text{if } x \in \{x_1 = x_2 = 0, x_0 \geq (m+2)r\} \end{cases}$$

Now t_m, the threshold policy with threshold m is defined as

$$t_m(x) \;=\; (F_m(Ax), F_m(D_1x), F_m(D_2x))$$

Notice that if $r = 1$ the notation reduces to that for the $M/M/2$ model of [3], as expected. The results presented so far were first presented in part in [16]. For all these results, the case of the average cost criterion can be obtained by letting the discount factor $\beta \uparrow 1$ and then use corollary 3 of [1]. This would require (as in [3]) computation of $\lim_{\beta\uparrow 1} J^{\beta}_{t_i}(x) \;\triangleq\; J_{t_i}(x) \;=\; \sum_{x'\in X} p(x')|x'|$ where $p(x')$ is the steady state probability of state x' in the Markov chain obtained using policy t_i and starting from state x initially. However, an analytical computation of $p(x')$ in any of the models is not tractable, since it requires solution of polynomial equations of degree 3 or higher.

Since there is an inherent difficulty in calculating the value of the cost function, it is of interest to consider alternative approaches to showing the optimality of the threshold policy that may not involve this computation at all. Such an alternative method was proposed in [4]. Although we do not wish to repeat here the details of this alternative approach, we found that it is useful in extending the optimality results in the more interesting direction of fairly general arrival statistics. This is the subject of the next section.

IV. The $GI/M/2$ Model

The system considered is shown in figure 5. The motivation for studying this model, besides the generality it allows for arrival processes, is its potential usefulness in studying interconnections of service stations in networks. For example, customers arriving at a node are usually the output of another node and therefore the interarrival times are correlated with their message lengths, which makes the Poisson assumption a poor one. We assume that customers arrive in a single stream and immediately join a queue served by two exponential servers of rates $\mu_1 > \mu_2$. The mean arrival time is $\frac{1}{\lambda}$ and for stability reasons we assume that $0 < \lambda < \mu_1 + \mu_2$. We further assume that the interarrival interval lengths obey the following stochastic dominance property: if X_1, X_2 denote two such lengths, then $EX_1 \geq EX_2 \;=>\; Pr(X_1 \geq t) \geq Pr(X_2 \geq t)$, $\forall\, t \geq 0$, i.e. $X_1 \geq_{st} X_2$. Notice that this property is not very restrictive. For example, the case of Erlangian arrivals satisfies this property [14]. It is however crucial for the application of the probabilistic methodology proposed in [4].

We wish again to choose the control actions at departure and arrival instants to minimize the average cost

$$\liminf_{t\to\infty} \frac{1}{t} E \int_0^t |x_s| ds$$

where $|x_s|$ denotes the number of customers in the system at time s. Let again $x_s = (x^0, x^1, x^2)$,where x^0 is the number of customers in the queue and $x^i = 1$ or 0 depending on whether server i is busy or not.

Clearly the residual interarrival time at a decision instant should be part of the state of such a system. Furthermore, it is not clear whether the results of [1] about existence and uniqueness of an optimal policy hold for this model. Assuming that an optimal policy exists and is unique, we proceed to show that for a given and fixed residual interarrival time at any decision instant, the optimal policy is of threshold type. In general, the threshold value depends on the value of that residual time. Thus, in a sense, the overall optimal policy is of the "switching curve" variety in the state space of queue sizes and residual times. For the proof it suffices to show that the following lemma (lemma 3.2 of [4]) holds true. Then the result is an immediate consequence of corollary 1 of [1]. For the sake of brevity, we will only outline the proof. We follow the notation of [4].

Lemma : The optimal policy π has the following properties

1) $\pi(x^0,0,0) \neq (x^0,0,0)$ and $\pi(x^0,0,1) \neq (x^0,0,1)$, $\forall\, x^0 > 0$ i.e. the policy does not leave server 1 idle if the queue is not empty.

2) $\pi(x^0,0,0) \neq (x^0-1,0,1)$, $\forall\, x^0 > 0$ i.e. the optimal policy does not "prefer" the slower server to the faster one.

3) Let $A \triangleq \{m > 0 \mid \pi(m,1,0) = (m-1,1,1)\}$. Then if $y^0, z^0 \in A$ and $y^0 < x^0 < z^0$, $x^0 \in A$ as well, i.e. A is "compact".

4) There is no y^0 such that $x^0 \notin A$, $\forall\, x^0 > y^0$,i.e. A is not bounded from above.

5) There is a finite $x^0 = x^0(\alpha_0, \mu_1, \mu_2) \in A$, $\forall\, 0 < \alpha < \alpha_0$, $0 < \lambda < \mu_1 + \mu_2$, where $e^{-\alpha} = \beta$ denotes the discount factor.

Proof: The main idea is to show that the (optimal) policy π can be strictly improved if it does not satisfy one of properties (1)-(4). Properties 1, 2 and 3 can be proved exactly as in [4], since they do not depend at all on the arrival statistics. One way to show that A is not bounded from above is to show that if $x \notin A$, then one can find another integer $y(x) \in A$, such that $x < y(x) < \infty$. This guarantees that there exists no integer $x' > y(x)$, such that $x' \notin A$, since in that case $y(x) < x' < y(x')$ and $y(x), y(x') \in A$ would imply by property 3 that $x' \in A$.

To show that A contains at least one element x^0, it suffices to show that if $x \notin A$ then $x < x^0 < x + n_0$, where $n_0 < \infty$ and $\pi(x^0, 1, 0) = (x^0 - 1, 1, 1)$. Consequently, we have to determine n_0.

As we show in the Appendix, we can construct n_0 as in [4], and thus A contains at least one element.This completes the proof of the lemma and the theorem. The crucial part of the proof is the one discussed in the Appendix, where the assumption is made that the initial residual interarrival time has some fixed value R_s.

V. Conclusions

We have considered the problem of controlling a multiserver queueing system, in both continuous and discrete time. The distinguishing characteristic of the system is that the statistics of the severs are different. Arrivals to the system are assumed to be separated by intervals that are independent and have arbitrary distributions, subject to a stochastic ordering property. We have used a combination of dynamic programming and probabilistic arguments to establish different parts of the extension to earlier results.

Threshold estimation and performance analysis of suboptimal threshold policies merit further investigation, since the value of the optimal threshold is very difficult to compute exactly. Generalizations of the result by further relaxing the assumptions on the arrival and service statistics as well as by increasing the number of servers are also of interest.

Appendix

We prove here property (5) of the lemma in section IV. We consider a $GI/M/1$ queue, with service rate μ_1, and two possible arrival rates: 0 and $\mu = \mu_1 + \mu_2$. For a given initial queue length m, and an initial interarrival residual time R_s common for both cases, let τ be the first time that one of the queue sizes hits the value 0 or the value n, for some fixed $n > m$. Let then

$$\phi(m, n, \alpha) \triangleq E[1(\sigma \leq \tau) \int_\sigma^\tau e^{-\alpha t} dt \; - \; 1(\sigma > \tau) \int_0^\sigma e^{-\alpha t} dt] \tag{1}$$

where σ is exponentially distributed with rate μ_2.

Define τ_r^k as the first time the queue size hits k, when the arrival rate is r (for the fixed values of m and R_s). Obviously $\tau_0^k = \infty \ , \ \forall\ k > m$. Now define $\tau_\mu(k) \triangleq \min\{\tau_0^0, \tau_0^k, \tau_\mu^0, \tau_\mu^k\}$. Note that $\tau_\mu(k) = \min\{\tau_0^0, \tau_\mu^0, \tau_\mu^k\} = \min\{\tau_0^0, \tau_\mu^k\}$ since $\tau_0^0 \leq \tau_\mu^0$. Let, further, $\tau_\lambda(k) \triangleq \min\{\tau_\lambda^0, \tau_\lambda^k\}$.

Notice that $\phi(0,k,\alpha) < 0 \ , \ \forall\ k > 0,\ 0 < \alpha < \alpha_0$ and also that, as $\tau \to \infty$ (which happens when $m = \frac{1}{2}n \to \infty$), $\phi(m,n,\alpha) > 0$. Thus there exist integers $0 < m_0 < n_0 < \infty$ such that

$$\phi(m_0, n_0, \alpha) > 0 \ , \ \forall\, 0 < \alpha < \alpha_0 \tag{2}$$

Consider now a $GI/M/1$ queue with an intermediate arrival rate λ such that $0 < \lambda < \mu_1 + \mu_2$; let the initial queue length be m_0 as defined by (2) and the initial residual interarrival time have the same value R_s as considered above. Let $\tau_\lambda(n_0), \tau_\mu(n_0)$ be the n_0-crossing times for the queues with rates λ and μ respectively as defined above. We want to prove that $\tau_\lambda(n_0) \geq_{st} \tau_\mu(n_0)$; since ϕ increases as τ increases stochastically, the above ordering implies that $\phi(m_0, n_0, \alpha)$ is positive for any $0 < \lambda < \mu_1 + \mu_2$.

Since $\lambda > 0$, we have that $\tau_0^0 \leq \tau_\lambda^0$; so it is sufficient to establish the inequality $\tau_\mu^{n_0} \leq_{st} \tau_\lambda^{n_0}$ for the desired result. We have

$$\tau_\mu^{n_0} = \sum_{j=1}^{L_1} X_\mu^{(j)} + R_s \qquad ; \qquad \tau_\lambda^{n_0} = \sum_{j=1}^{L_2} X_\lambda^{(j)} + R_s \tag{3}$$

where $X_\mu^{(j)}$ and $X_\lambda^{(j)}$ denote the j-th interarrival interval length, when the arrival rate is μ or λ respectively. Clearly L_1 and L_2 are random and in general not equal.

Recall that we have assumed that the arrival process is such that $\mu > \lambda$ implies $X_\mu \leq_{st} X_\lambda$. Then from prop. 2.2.5 of [17] we have that for any L

$$Pr\left(\sum_{j=1}^{L} X_\lambda^{(j)} > \epsilon\right) \ \geq \ Pr\left(\sum_{j=1}^{L} X_\mu^{(j)} > \epsilon\right)$$

Now let $L = L_2$. It is easily seen that the $L_2 - th$ arrival in the queue with arrival rate μ will occur at a time τ_1 stochastically smaller than $\tau_\lambda^{n_0}$. Moreover, (see theorem 7 of [14]), it will find a queue size stochastically larger than n_0 , since the two systems have exactly the same departures. Thus $\tau_\mu^{n_0} \leq_{st} \tau_\lambda^{n_0}$ and $\tau_\lambda(n_0) \geq_{st} \tau_\mu(n_0)$. So there exist m_0, n_0 such that $\phi > 0$ for any $0 < \lambda < \mu_1 + \mu_2$. *

The rest of the argument goes exactly as in [4].

REFERENCES

[1] S. Lippman, "Semi-Markov decision processes with unbounded rewards," Man. Sci., Vol. 19, No. 7, pp. 717-731, 1973.

[2] R.L. Larsen, "Control of multiple exponential servers with application to computer systems," Ph.D. thesis, Univ. of MD, 1981.

[3] W. Lin, P.R. Kumar, "Optimal control of a queueing system with two heterogeneous servers," IEEE AC-29, No. 8, pp. 696-703, 1984.

[4] J. Walrand, "A note on 'Optimal control of a queueing system with tho heterogeneous servers'", Systems and Control Letters, Vol. 4, pp. 131-134, 1984.

[5] B. Hajek, "Optimal control of two interacting service stations," IEEE AC-29, No. 6, pp. 421-429, 1984.

[6] C.E. Bell, "Optimal operation of a $M/M/2$ queue with removable servers," Oper. Res. ,Vol. 28, No. 5, pp. 1189-1204, 1980.

[7] Y.C. Chow, W.H. Kohler, "Models for dynamic load balancing in a heterogeneous multiple processor system," IEEE C-28, pp. 354-361, 1979.

---

* Another possibility is to view $\tau_\lambda^{n_0}, \tau_\mu^{n_0}$ as first passage times from "state" m_0 to "state" n_0 [18]. Define $T_\nu(j, j+1)$ as the first passage time from "state" j to "state" $j+1$ when the arrival rate is ν. Then

$$\tau_\mu^{n_0} = \sum_{j=m_0}^{n_0-1} T_\mu(j, j+1) \qquad ; \qquad \tau_\lambda^{n_0} = \sum_{j=m_0}^{n_0-1} T_\lambda(j, j+1)$$

Since $T_\mu(j, j+1), T_\lambda(j, j+1)$ are independent, it would suffice to show that

$$T_\mu(j, j+1) \leq_{st} T_\lambda(j, j+1) \qquad . \tag{4}$$

However, each term in (4) is a sum of a random number of interarrival interval lengths, so this is equivalent to proving that $\tau_\mu^{n_0} \leq_{st} \tau_\lambda^{n_0}$.

[8] E.L. Hahne, "Dynamic routing in an unreliable manufacturing network with limited storage," MIT, Cambridge, Rep. LIDS-TH-1063, 1981.

[9] J.N. Tsitsiklis, "Convexity and characterization of optimal policies in a dynamic routing problem," MIT, Cambridge, Rep. LIDS-R-1178, 1982.

[10] A. Seidmann, P.J. Schweitzer, "Part selection policy for a flexible manufacturing cell feeding several production lines," IIE Trans. ,Vol. 16, 4, pp. 355-362, 1984.

[11] A.K. Agrawala et al., "A stochastic optimization algorithm minimizing expected flow times on uniform processors," IEEE C-33, pp. 351-356, 1984.

[12] C.E. Bell, "Optimal operation of a $M/M/2$ queue with removable servers," Oper. Res. ,Vol. 28, No. 5, pp. 1189-1204, 1980.

[13] I. Viniotis, "Optimal Control of a multiserver queueing system with heterogeneous servers," M.Sc. thesis, Univ. of MD, Elect. Eng. Dept., 1985.

[14] W. Whitt, "Comparing counting processes and queues," Adv. Appl. Prob., 13, pp. 207-220, 1981.

[15] S. Lippman, "Applying a new device in the optimization of exponential queueing systems," Oper. Res., Vol. 23, pp. 687-710, 1975.

[16] I. Viniotis, A. Ephremides, "An optimization problem in a multiserver queue," Proc. of the 23rd Annual Allerton Conf. on Control, Comm. and Comp., Urbana-Champaign, pp. 555-563.

[17] D. Stoyan, "Comparison methods for queues and other stochastic models," J. Wiley and Sons, 1983.

[18] J. Keilson, "Markov chain models-Rarity and exponentiality," Springer Verlag, 1979

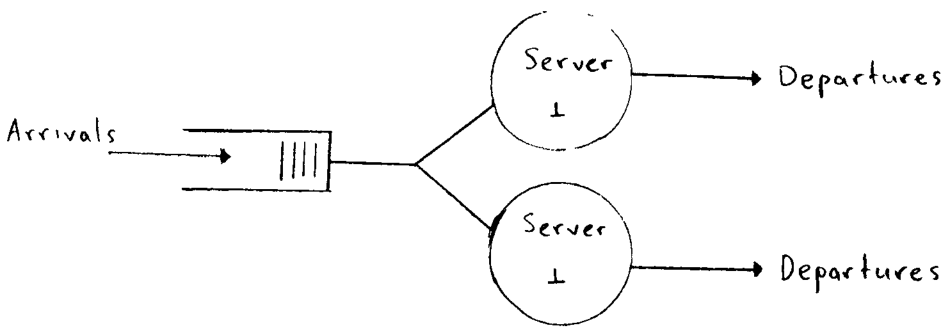

Figure 1

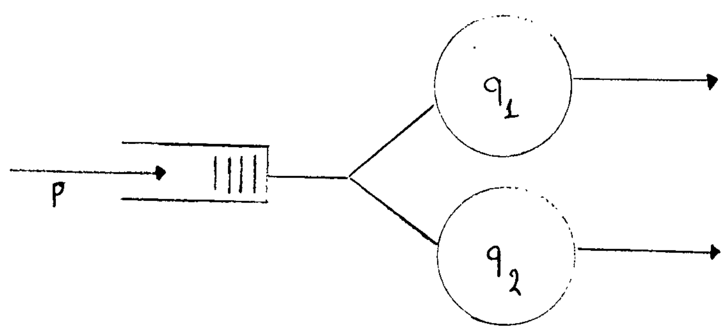

Figure 2

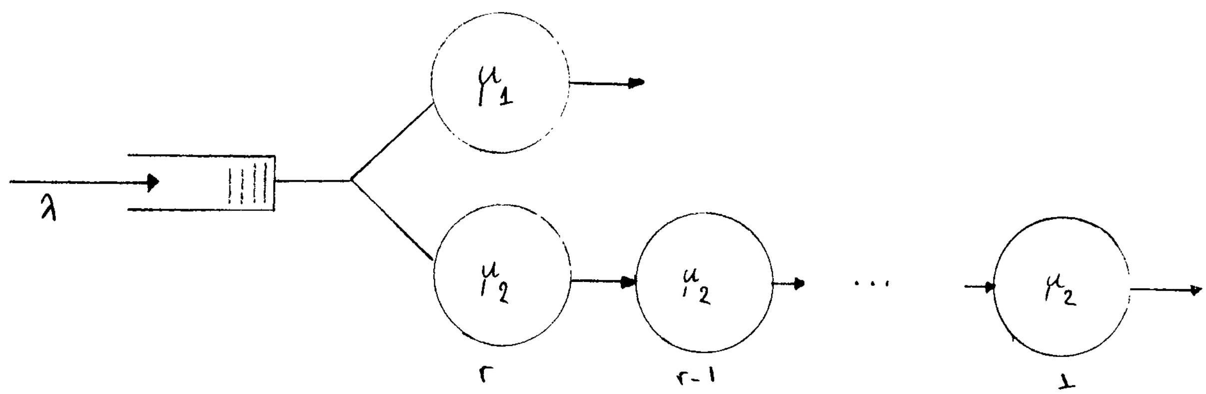

Figure 3

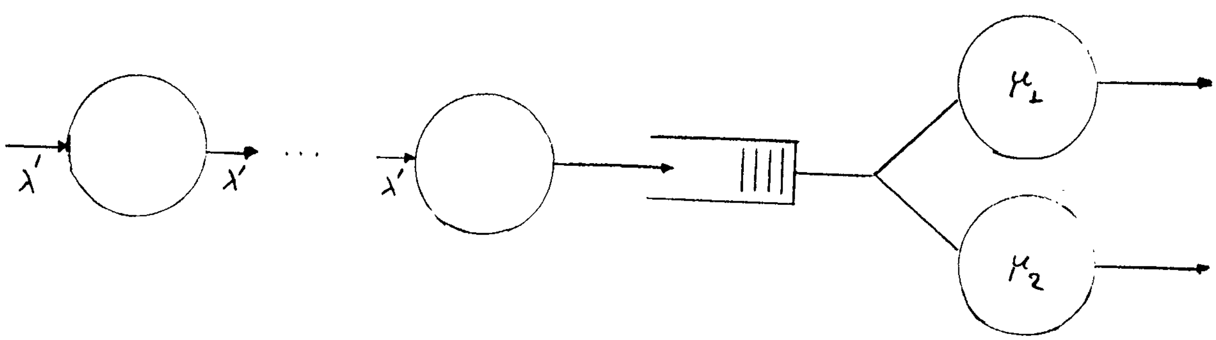

Figure 4

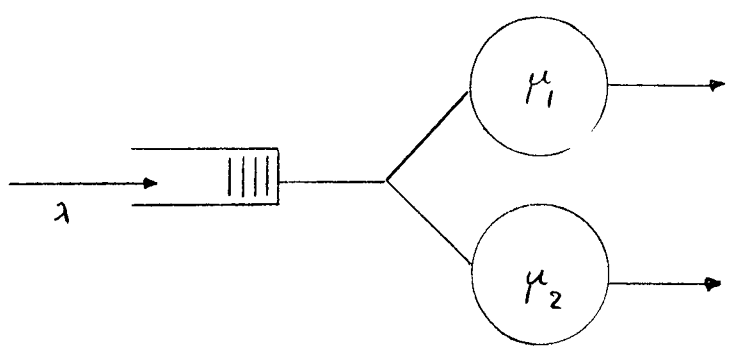

Figure 5