A NOTE ON THE HYPERCUBE MODEL

M.N. KATEHAKIS

Department of Applied Mathematics and Statistics, State University of New York at Stony Brook, Stony Brook, NY 11794, USA

Received September 1984

We consider a variation of the hypercube model in which there are N distinguishable servers and R types of customers. Customers that find all servers busy (blocked customers) are lost. When service times are exponentially distributed and customers arrive according to independent Poisson streams, we show that the policy which always assigns customers to the fastest available server minimizes the long-run average number of lost customers. Furthermore, we derive an upper bound for the blocking probability and the long-run average number of customers lost.

first passage times * Markov decision processes * alternating renewal process

1. Introduction

Consider a queueing system in which the customers can be classified into R distinct types. Arrivals occur in independent Poisson streams; the arrival rate for the ith customer type is λ_i. There are N distinguishable servers numbered 1 through N. The amount of service provided by server i is exponentially distributed with rate μ_i, independent of the type of customer being served. Exactly one server is assigned to each customer if a server is available. No preemption is allowed. Customers arriving when the system is saturated (all servers busy) are lost or handled by some external system. In this paper we show that the policy which assigns all arrivals to the fastest idle server minimizes the long-run average number of lost customers. Furthermore, we derive an upper bound for the steady state blocking probability for the system.

The above model is a variation of the hypercube model developed by Larson [6] to study problems of resource allocation in public systems. Jarvis [3] developed an iterative procedure for finding policies that minimize the long-run average cost for the case in which there are costs of the form $\sigma(i, n)$ when server i is assigned to a customer of type n and a penalty cost s_n when a customer of type n is lost. Katehakis and Levine [5] derived computational procedures for the Jarvis model for the case of heavy and light traffic intensities.

Derman et al. [2] considered the case of a single arrival class, with general iid interarrival times, and showed that the policy which always assigns customers to the fastest idle server stochastically minimizes the number of customers in the system. This approach of Derman et al. [2] can be used to establish the result of this paper. However, we present a different proof which essentially involves establishing that the conjectured optimal policy attains the solution to the functional equations of an equivalent Markov decision problem. A similar approach was taken in Katehakis and Derman [4] for an optimal repair allocation problem.

2. Notation and definitions

Since the servers are distinguishable, the state of the system can be represented by a vector $\boldsymbol{x} = (x_1, \ldots, x_N)$, where $x_i = 1$ if server i is busy and $x_i = 0$ if server i is idle. Thus, the hypercube $S = \{0, 1\}$ is the set of all possible states and $\boldsymbol{1} = (1, \ldots, 1)$ is the state in which all servers are busy.

0167-6377/85/$3.30

Given state $\boldsymbol{x} \in S$, define

$$B(\boldsymbol{x}) = \{ j\colon x_j = 1 \}, \quad \text{i.e., } B(\boldsymbol{x}) \text{ is the set of busy servers in state } \boldsymbol{x},$$
$$I(\boldsymbol{x}) = \{ j\colon x_j = 0 \}, \quad \text{i.e., } I(\boldsymbol{x}) \text{ is the set of idle servers in state } \boldsymbol{x},$$
$$m(\boldsymbol{x}) = \min I(\boldsymbol{x}),$$
$$(1_j, \boldsymbol{x}) = (x_1, \dots, x_{j-1}, 1, x_{j+1}, \dots, x_N) \quad \text{if} \quad j \in I(x),$$
$$(0_j, \boldsymbol{x}) = (x_1, \dots, x_{j-1}, 0, x_{j+1}, \dots, x_N) \quad \text{if} \quad j \in B(x),$$
$$|\boldsymbol{x}| = \sum_{j=1}^{N} x_j, \quad \mu(\boldsymbol{x}) = \sum_{k \in B(\boldsymbol{x})} \mu_k,$$
$$\mu = \sum_{k=1}^{N} \mu_k, \quad \lambda = \sum_{i=1}^{R} \lambda_i,$$
$$K(\boldsymbol{x}) = \{ \boldsymbol{a}/\boldsymbol{a} = (a_1, \dots, a_R),\ a_i \in I(\boldsymbol{x}) \}.$$

$K(\boldsymbol{x})$ denotes the set of all possible actions in state $\boldsymbol{x}$; $a_i = k$ $(i \leqslant k \leqslant N)$ means assign next arrival of type i to server k.

Finally, for any action $\boldsymbol{a} \in K(\boldsymbol{x})$, the total arrival rate to server j, $j \in I(\boldsymbol{x})$, is

$$\lambda_j(\boldsymbol{a}) = \sum_{i=1}^{R} \lambda_i \delta(a_i, j),$$

3. Optimality equations

Standard results in Markov decision theory (Derman [1]) show that there exists an optimal deterministic policy. Let Π denote the set of deterministic policies (i.e., policies that assign arrivals to servers as a deterministic function of the state of the system). When a policy $\pi \in \Pi$ is employed, it is easy to see that the evolution of the state of the system can be described by a continuous time finite state irreducible Markov chain $\{x^{\pi}(t),\ t \geqslant 0\}$ where $x_i^{\pi}(t) = 1$ if server i is busy at time t and 0 otherwise. Since customers are lost only when the system is in state $\boldsymbol{1}$, to minimize the long-run average number of customers lost is equivalent to minimize the ergodic probability of the system being in state $\boldsymbol{1}$. Let $e_{\pi}(\boldsymbol{x})$ denote the ergodic probability of the system being in state $\boldsymbol{x}$ and $\tau_{\pi}(x)$ denote the expected first passage time to state $\boldsymbol{1}$ from state $\boldsymbol{x}$, when a policy π is employed. Returns to state $\boldsymbol{1}$ generate an alternating renewal process. Thus employing an alternating renewal argument it is easy to show that $e_{\pi}(\boldsymbol{1})$ is given by (1).

$$e_{\pi}(\boldsymbol{1}) = \frac{1}{\mu}\left(1/\mu + \sum_{j=1}^{N} (\mu_j/\mu)\tau_{\pi}(0_j, \boldsymbol{1})\right)^{-1}. \tag{1}$$

Note that $e_{\pi}(\boldsymbol{1})$ is minimized by a policy that maximizes $\tau_{\pi}(\boldsymbol{x})$. We can now formally state the following.

Lemma 1. *A policy π^* minimizes the long-run average number of customers lost if and only if*

$$\tau_{\pi^*}(\boldsymbol{x}) \geq \tau_{\pi}(\boldsymbol{x}) \quad \forall \boldsymbol{x} \in S - \{\boldsymbol{1}\} \quad \forall \pi \in \Pi. \tag{2}$$

By conditioning on the first transition out of state $\boldsymbol{x}$, we see that for any $\pi \in \Pi$ the $\tau_{\pi}(\boldsymbol{x})$'s can be obtained as the unique solution to the following system of linear equations:

$$\tau_{\pi}(\boldsymbol{x}) = \frac{1}{\lambda + \mu(\boldsymbol{x})}\left(1 + \sum_{j \in I(\boldsymbol{x})} \lambda_j(\pi(\boldsymbol{x}))\tau_{\pi}(1_j, \boldsymbol{x}) + \sum_{i \in B(\boldsymbol{x})} \mu_i \tau_{\pi}(0_i, \boldsymbol{x})\right), \quad \boldsymbol{x} \neq \boldsymbol{1},$$
$$\tau_{\pi}(\boldsymbol{1}) = 0. \tag{3}$$

In (2) $\pi(x)$ denotes the action taken by policy π in state x; $\pi(x) \in K(x)$. A policy π^* is optimal (i.e., Lemma 1 holds) iff the associated first passage times satisfy the following functional equations (Derman [1]):

$$\tau_{\pi^*}(x) = \max_{a \in K(x)} \left\{ \frac{1}{\lambda + \mu(x)} \left(1 + \sum_{j \in I(x)} \lambda_j(a) \tau_{\pi^*}(1_j, x) = \sum_{i \in B(x)} \mu_i \tau_{\pi^*}(0_i, x) \right) \right\}, \quad x \neq \boldsymbol{1}. \tag{4}$$

Now it is easy to see that (4) are equivalent to the following:

$$\tau_{\pi^*}(x) = \max_{a \in K(x)} \left\{ 1 + (\mu - \mu(x)) \tau_{\pi^*}(x) + \sum_{j \in I(x)} \lambda_j(a) \tau_{\pi^*}(1_j, x) + \sum_{i \in B(x)} \mu_i \tau_{\pi^*}(0_i, x) \right\}, \tag{5}$$

where in (5) above we have assumed without any loss in generality that $\lambda + \mu = 1$. In the next section, we assume that $\mu_1 > \mu_2 > \cdots > \mu_N$ and show that for π^* defined by: $\pi^*(x) = (m(x), \ldots, m(x))$ (5) hold.

Proof of the optimality of π^*

For notational simplicity we drop the subscript π^* from $\tau_{\pi^*}(x)$.

We now state the main result.

Proposition 1. *If $\mu_1 > \mu_2 > \cdots > \mu_N$ then π^* is an optimal policy.*

Proof. Since $\lambda_j(\pi^*(x)) = 0$ if $j \neq m(x)$ to show that $\{\tau(x),\ x \in S,\ x \neq 1\}$ form a solution to (5), it suffices to show that

$$\tau(1_{m(x)}, x) \geqslant \tau(1_j, x) \quad \forall j \in I(x) \ \ j \neq m(x) \quad \forall x \in S \ \ x \neq \boldsymbol{1}. \tag{6}$$

A condition stronger than (6) is

$$\tau(1_k, x) \geqslant \tau(1_l, x) \quad \forall k, l \in I(x) \ \ k < l \quad \forall x \in S \ \ x \neq \boldsymbol{1}. \tag{7}$$

We next prove (7) as follows.

Consider the following successive approximations procedure obtained from (5).

$$\tau^{(0)}(x) = 0 \quad \forall x \in S,$$

$$\tau^{(n+1)}(x) = \max_{a \in K(x)} \left\{ 1 + (\mu - \mu(x)) \tau^{(n)}(x) + \sum_{j \in I(x)} \lambda_j(a) \tau^{(n)}(1_j, x) + \sum_{i \in B(x)} \mu_i \tau^{(n)}(0_i, x) \right\},$$

$$\tau^{(n)}(\boldsymbol{1}) = 0. \tag{9}$$

Then, as $n \to \infty$, $\tau^{(n)}(x) \to \tau(x)\ \forall x \in S$; see, for example, Derman [1, p. 54]. Thus, to prove (7), it suffices to establish the following:

$$\tau^{(n)}(1_k, x) \geqslant \tau^{(n)}(1_l, x) \quad \forall k, l \in I(x)) \ \ k < l \ \ x \in S. \tag{10}$$

Now (10) can be established by induction as follows.

Assume that (10) hold for some n, then using this induction hypothesis we obtain

$$\tau^{(n+1)}(1_k, x) = 1 + (\mu - \mu(1_k, x)) \tau^{(n)}(1_k, x) + \lambda \tau^{(n)}(1_{m(1_k, x)}, x) + \sum_{i \in B(1_k, x)} \mu_i \tau^{(n)}(0_i, 1_k, x), \tag{11}$$

$$\tau^{(n+1)}(1_l, x) = 1 + (\mu - \mu(1_l, x)) \tau^{(n)}(1_l, x) + \lambda \tau^{(n)}(1_{m(1_l, x)}, x) + \sum_{i \in B(1_l, x)} \mu_i \tau^{(n)}(0_i, 1_l, x). \tag{12}$$

It is easy to establish using (11), (12) and the induction hypothesis that $\tau^{(n+1)}(1_k, x) \geqslant \tau^{(n+1)}(1_l, x)$. The proof is now complete since (10) hold for $n = 0$. □

5. An upper bound for the blocking probability

Even though it is difficult to obtain an exact expression for the steady state blocking probability $e_{\pi^*}(\boldsymbol{1})$, it is easy to obtain an upper bound $e_{\pi^0}(\boldsymbol{1})$ for it; $\lambda e_{\pi^0}(\boldsymbol{1})$ is then an upper bound for the long-run average number of lost customers.

Proposition 2

$$e_{\pi^*}(\boldsymbol{1}) \leqslant e_{\pi^0}(\boldsymbol{1}),$$

where

$$e_{\pi^0}(\boldsymbol{1}) = \left[1 + \sum_{i=1}^{N} \frac{\mu_i}{\lambda} + \sum_{i_1 < i_2} \frac{2!\mu_{i_1}\mu_{i_2}}{\lambda^2} + \sum_{i_1 < i_2 < i_3} \frac{3!\mu_{i_1}\mu_{i_2}\mu_{i_3}}{\lambda^3} + \cdots + \frac{N!\mu_1\mu_2 \cdots \mu_N}{\lambda^N}\right]. \tag{13}$$

Proof. Consider the suboptimal policy π^0 defined as follows. Whenever the system is in state $\boldsymbol{x}$, π^0 assigns the next arrival to any of the idle servers $I(\boldsymbol{x})$ with equal probability, $1/|I(\boldsymbol{x})|$. It is easy to see that the corresponding to π^0 Markov chain $\{\boldsymbol{x}_{\pi^0}(t), t \geqslant 0\}$ is time reversible. Thus, one can solve the detailed balance equations (see Ross [7, p. 156]) to show that $e_{\pi^0}(\boldsymbol{1})$ is as in (13). The result then follows from the fact that π^* is an optimal policy. □

References

[1] C. Derman, *Finite State Markovian Decision Processes*, Academic Press, New York, 1970.

[2] C. Derman, G. Lieberman and S.M. Ross, "On the optimal assignment of servers and repairmen", O.R.C. 78-22, University of California, Berkeley, CA, 1978.

[3] J.P. Jarvis, "Optimal assignments in a Markovian queueing system", *Comp. and Oper. Res.* **8**, 17–23 (1981).

[4] M.N. Katehakis and C. Derman, "Optimal repair allocation in a series system", *Mathematics of Operations Research* **9**(4), 615–623 (1984).

[5] M.N. Katehakis and A. Levine, "Allocation of distinguishable servers", *Comp. and Oper. Res.* (under review).

[6] R.C. Larson, "A hypercube queueing model for facility location and redistricting in urban emergency services", *Comp. and Oper. Res.* **1**, 67–95 (1974).

[7] S.M. Ross, *Stochastic Processes*, Wiley, New York, 1983.