A simple proof of the optimality of the best N-policy in the M/G/1 queueing control problem with removable server

Talman, A.J.J.

Published in:
Statistica Neerlandica

Publication date:
1979

A simple proof of the optimality of the best N-policy in the M/G/1 queueing control problem with removable server

by A. J. J. TALMAN *

Abstract An M/G/1 queueing system with removable server is considered. The following costs are incurred: a holding cost of h per unit time per customer in the system, a cost of $r_1(r_2)$ per unit time when the service mechanism is on (off) and a fixed cost of $K_1(K_2)$ for turning it on (off).

In this paper we shall give a very simple proof for the well known and intuitively obvious fact that the best N-policy is optimal for the average cost criterion among the class of all policies by first proving that the average cost formula of that policy and its relative cost function satisfy the optimality equation for the average cost criterion. The optimality of the best N-policy is then an immediate consequence.

1 Introduction

In this paper we consider the famous queueing control model with a removable server. This problem, first studied by YADIN and NAOR [9] and subsequently by a large number of authors, has many applications. Generally spoken a server is a mechanism that performs an operation on units called customers. For example, the customers could be people arriving at a ticket-office and the server the man who sells the tickets, or the server represents a production facility whereas the customers are the orders for the product.

In many practical situations the server may be turned off at any service-completion epoch of a customer and it may be turned on only at an arrival epoch of a customer. If the operational costs of the server are considerable, the manager can prefer to keep the server closed down during certain periods, especially when there are only a few customers in the system. Formally, we consider the following model. Customers arrive at a single server station in accordance with a Poisson process with rate λ. Each customer requires an amount of service where the service times of the customers are independent random variables having a common probability distribution function $F(.)$ with finite first moment μ and finite second moment $\mu^{(2)}$. It is assumed that $F(0) = 0$ and $\varrho < 1$ where $\varrho = \lambda\mu$. Each customer will be serviced. The service mechanism may be either turned on or off, where service will be provided only when the server is on. The control of the system will be based on the queue size. The following costs are considered. There is fixed cost of $K_1(K_2)$ for turning the server on (off) and there is a holding cost of $h > 0$ per customer per unit time. Further there is an operating cost at rate $r_1(r_2)$ when the server is on (off).

We shall assume that the system can be only controlled at the arrival epochs of the customers when the server is off and at the service-completion epochs when the server is on. At these epochs an action which is based on the queue size must be taken to

* Vrije Universiteit, Amsterdam.

change the current position of the server or not. A policy for controlling the system is any rule for choosing actions at these decision epochs, where a policy is said to be stationary if the action it chooses only depends on the current state of the system. We are interested in a policy that minimizes the average cost. For this so-called M/G/1 queueing problem with removable server, the intuitively appealing N-policies will be used in practice, cf. YADIN and NAOR [9] for further motivation. An N-policy is a stationary policy which turns the server off when the system is empty and turns the server on when the number of customers waiting in the system is equal to or larger than the value N, where N is a positive integer.

YADIN and NAOR [9] derived a formula for the long-run average cost of the N-policy. In a sequence of papers (BELL [1], HEYMAN [2], LIPPMAN [3] and SOBEL [6]) rather complicated proofs were given of the intuitively obvious fact that the best N-policy is optimal among the class of all policies for the average cost criterion. In this paper we shall provide a very simple and didactic proof which yields an even stronger result, cf. also TIJMS [8]. We shall show that the optimality equation for the average cost criterion has a solution by constructing this solution explicitly. More precisely, we shall prove that the average cost formula and the relative cost function of the best N-policy satisfy the optimality equation for the average cost criterion. As an immediate consequence we then find that this policy is optimal among the class of all policies.

2 Preliminaries

Clearly, the model may be regarded as a two action semi-Markov decision process with states $0, 1', 1, 2', 2, \ldots$ where state i means that there are i customers in the system and the server is on, $i = 0, 1, 2, \ldots$, whereas state i' means that there are i customers waiting in the system and the server is off, $i = 1, 2, \ldots$. Action 0 will mean "turn or keep the server off" and action 1 "turn or keep the server on". Hence, the state space S is the set $\{0, 1', 1, 2', 2, \ldots\}$ and the action space for the state s is the set $\{0, 1\}$, for all $s \in S$.

For any policy π define (cf. ROSS [5]) $\varphi_\pi(s)$ is the long run average expected cost for initial state s, $s \in S$.

For any class C of policies a policy $\pi^* \in C$ is said to be average cost optimal within the class C of policies if and only if $\varphi_{\pi^*}(s) \leqslant \varphi_\pi(s)$ for all $s \in S$ and all $\pi \in C$.

For a given N-policy we now introduce the functions $t_N(s)$ and $k_N(s)$ which play an important role in our analysis.

Define

$t_N(s)$ = the expectation of the epoch of the first return to state 0, given that the initial state is s and that the N-policy is used

and

$k_N(s)$ = the expected cost incurred up to the epoch of the first return to state 0, given that the initial state is s and that the N-policy is used.

It is well known that for the standard M/G/1 queueing model the expected length of one busy cycle equals

$$\alpha = \mu(1-\varrho)^{-1}$$

and that the total expected time spent by customers in the system during one busy period is (cf. e.g. [7])

$$\beta = \mu(1-\varrho)^{-1} + \tfrac{1}{2}\lambda\mu^{(2)}(1-\varrho)^{-2}.$$

Suppose now that at epoch 0 a service starts with i customers present. Then any of these i customers in fact generates a busy period, so that the time until the system becomes empty can be seen as the sum of i independent busy periods each starting with one customer. It is now easily seen that for all $i \geqslant 1$

$$t_N(i) = \alpha i \quad \text{and} \quad k_N(i) = h\left\{i\beta + \sum_{k=1}^{i-1} k\alpha\right\} + r_1 t_N(i) = h\{i\beta + \tfrac{1}{2}i(i-1)\alpha\} + r_1\alpha i.$$

Since the expected time between two arrivals is equal to λ^{-1} and the system is kept closed until the Nth customer arrives,

$$t_N(i') = (N-i)\lambda^{-1} + t_N(N) \quad \text{and} \quad k_N(i') = h\sum_{k=i}^{N-1} k\lambda^{-1} + r_2(N-i)\lambda^{-1} + K_1 + k_N(N)$$
$$\text{for} \quad i = 1, \dots, N-1.$$

Clearly

$$t_N(i') = t_N(i) \quad \text{and} \quad k_N(i') = K_1 + k_N(i) \quad \text{for} \quad i \geqslant N,$$

and

$$t_N(0) = \lambda^{-1} + t_N(1') \quad \text{and} \quad t_N(0) = K_2 + r_2\lambda^{-1} + k_N(1').$$

Observe that for any policy the state 0 is regenerative since $\varrho < 1$. The interval between two epochs at which the state is 0 is called a cycle. From the theory of regenerative processes (e.g. Theorem 7.5 in Ross [5]) it follows that for any N-policy $\varphi_\pi(s)$ exists and is independent of s. Denoting the average cost of the N-policy therefore by $g(N)$, we then have (see Ross [5])

$$g(N) = k_N(0)/t_N(0),$$

i.e. the longrun average cost of an N-policy is equal to its average expected cost during one cycle, which results in the following well-known formula (see Yadin and Naor [9])

$$g(N) = K\lambda(1-\varrho)N^{-1} + \varrho r_1 + (1-\varrho)r_2 + \tfrac{1}{2}h(N-1) + h\varrho + \tfrac{1}{2}h\lambda^2\mu^{(2)}(1-\varrho)^{-1}$$

where K equals $K_1 + K_2$. Remember that N is a positive integer. Since $g(N)$ is convex and $g(N) \to \infty$ as $N \to \infty$, we van easily find a positive integer N^* such that $g(N^*) \leqslant g(N)$ for all $N = 1, 2, \ldots$. Clearly, $N^* = m+1$, where m is the entier of $(2K\lambda(1-\varrho)/h)^{\frac{1}{2}}$, if $m = 0$ or $g(m+1) \leqslant g(m)$. Otherwise, $N^* = m$.

As we allow the server to be on when the system is empty, we also have to consider the policy π_0, which always prescribes action 1.

Under this policy we have the POLLACZEK-KHINTCHINE formula that the average number of customers in the system equals $\lambda(1-\varrho)\beta$, so the average cost of this policy is given by

$$g_0 = r_1 + h\lambda(1-\varrho)\beta. \tag{1}$$

Note that the policy which prescribes action 1 in all states i', with i sufficiently large, and never turns off, has the same average cost.

We shall first consider the case where $g(N^*) < g_0$, or equivalently the case where

$$r_1 - r_2 > K\lambda/N^* + \tfrac{1}{2}h(N^*-1)/(1-\varrho). \tag{2}$$

We shall prove that under this condition the N-policy is average cost optimal within a wide class of policies. To do this we need the so-called relative cost function $v(s)$, defined by

$$v(s) = k_{N^*}(s) - g^* t_{N^*}(s) \quad \text{for} \quad s \in S,$$

where $g^* = g(N^*)$, i.e. the relative cost function is the expected total costs minus the average costs incurred up to the end of the cycle given that the initial state is s and that the best N-policy is used. Notice that

$$v(s) = 0 \tag{3}$$

and

$$v(i') = K_1 + v(i) \quad \text{for} \quad i \geqslant N^*. \tag{4}$$

3 Analysis

Before proving that the average cost and the relative cost function of the N*-policy satisfy the optimality equation for the average cost criterion, we establish some properties of the function $v(s)$.

LEMMA. The function $v(s)$ has the following properties:

a. $v(i) = hi\mu + \frac{1}{2}h\lambda\mu^{(2)} + r_1\mu - g^*\mu + \mathrm{E}v(i-1+X)$ for $i \geqslant 1$

b. $v(i') = hi\lambda^{-1} + r_2\lambda^{-1} - g^*\lambda^{-1} + v((i+1)')$ for $i = 1, \ldots, N^*-1$

c. $v(0) = K_2 + r_2\lambda^{-1} - g^*\lambda^{-1} + v(1')$

d. $v(i) - K_2 \leqslant v(i') \leqslant v(i) + K_1$ for $i = 1, \ldots, N^*-1$

e. $v(i) - K_2 \leqslant v(i') \leqslant hi\lambda^{-1} + r_2\lambda^{-1} - g^*\lambda^{-1} + v((i+1)')$ for $i \leqslant N^*$

f. $v(0) < r_1\lambda^{-1} - g^*\lambda^{-1} + v(1)$

where X in part a) is distributed as the number of customers arriving during one service time.

PROOF. By substitution and after some simple calculations we obtain

$$v(i) = \tfrac{1}{2}h\alpha i^2 + h(1-\varrho)\beta i + (r_1 - r_2)\mu i - \tfrac{1}{2}h\alpha N^* i - K\varrho i/N^* \quad \text{for} \quad i \geqslant 0 \tag{5}$$

and

$$v(i') = -\tfrac{1}{2}h\lambda^{-1}i^2 + h(1-\varrho)\beta i + (r_1 - r_2)\mu i + \tfrac{1}{2}hN^*\lambda^{-1}i + K(1-\varrho)i/N^* - K_2 \quad \text{for} \quad i = 1, \ldots, N^* - 1. \tag{6}$$

The first three assertions of the lemma immediately follow from these expressions. Notice that $EX = \lambda\mu$ and $EX^2 = \lambda\mu + \lambda^2\mu^{(2)}$. To prove d and e, we first observe that $g(N^*) \leqslant g(i)$ for all $i \geqslant 1$ implies

$$(K/N^* - h\gamma i)(i - N^*) \leqslant 0 \quad \text{for all} \quad i \geqslant 1, \tag{7}$$

where for ease of notation $\gamma^{-1} = 2\lambda(1-\varrho)$. From (5) and (6) we get

$$v(i') - v(i) = Ki/N^* + \tfrac{1}{2}h\lambda^{-1}(1-\varrho)^{-1}i(N^* - i) - K_2 \quad \text{for} \quad i = 1, \ldots, N^* - 1.$$

Together with (7) this implies for $i = 1, \ldots, N^* - 1$

$$-K_2 \leqslant v(i') - v(i) = (K/N^* - \tfrac{1}{2}h\lambda^{-1}(1-\varrho)^{-1}i)(i - N^*) + K_1 \leqslant K_1$$

which proves assertion d. Note that by (4) the first unequality of e is trivial. To prove the other one, it follows from (4) and (5) that

$$\begin{aligned} v((i+1)') + hi\lambda^{-1} + r_2\lambda^{-1} - g^*\lambda^{-1} - v(i') &= h\gamma(2i + 1 - N^*) - K/N^* \\ &= h\gamma(i - N^*) + h\gamma(i+1) - K/N^* \quad \text{for} \quad i \geqslant N^*. \end{aligned}$$

Again using (7) this is positive, which proves assertion e. Finally, by (5) we have

$$v(1) + r_1\lambda^{-1} - g^*\lambda^{-1} = (r_1 - r_2)\lambda^{-1} - h\gamma(N^* - 1) - K/N^*$$

which is positive under condition (2), yielding the last assertion of the lemma by (3). So, the lemma has been proved.

We are now in a position to prove that g^* and $v(s)$ satisfy the average cost optimality equation. Let $T(s, a)$ denote the next state of the system when action a is taken in state s. Also let $\tau(s, a)$ be the expected time until the next decision epoch and $r(s, a)$ the expected cost incurred up to that epoch when action a is taken in state s.

THEOREM 1. The average cost of the N*-policy, g^*, and the relative cost function $v(s)$ satisfy

$$v(s) = \min_{a\in\{0,1\}} \{r(s,a) - g^*\tau(s,a) + \mathrm{E}[v(T(s,a))]\} \quad \text{for} \quad s \in S.$$

Moreover, for any state s the action prescribed by the N*-policy in state s minimizes the right side.

PROOF. Trivially, we have $\tau(s,0) = \lambda^{-1}$ for $s \in S$,

$$\tau(s,1) = \mu \quad \text{for} \quad s \in S/\{0\}, \quad \text{and} \quad \tau(0,1) = \lambda^{-1} + \mu.$$

Using the fact that, given a service time equal to t, any customer arriving in this time interval waits an average amount of $\frac{1}{2}t$ during this service, we find

$$r(i,1) = h\mu i + \tfrac{1}{2}h\lambda\mu^{(2)} + r_1\mu \quad \text{and} \quad r(i',1) = K_1 + r(i,1) \quad \text{for} \quad i \geqslant 1$$

$$r(i',0) = hi\lambda^{-1} + r_2\lambda^{-1} \quad \text{and} \quad r(i,0) = K_2 + r(i',0) \qquad \text{for} \quad i \geqslant 1$$

$$r(0,0) = K_2 + r_2\lambda^{-1} \quad \text{and} \quad r(0,1) = r_1\lambda^{-1} + r(1,1).$$

Letting again X be distributed as the number of customers arriving during one service time, we find that $T(i,0) = T(i',0) = (i+1)'$ and $T(i,1) = T(i',1) = i+X-1$ for $i \geqslant 1$ and $T(0,1) = X$. Hence, by assertion a of the lemma it suffices to prove that

$$v(i') = \min\{K_1 + v(i), (hi + r_2 - g^*)\lambda^{-1} + v((i+1)')\} \quad \text{for} \quad i \geqslant 1 \tag{8}$$

$$v(i) < K_2 + (hi + r_2 - g^*)\lambda^{-1} + v((i+1)') \quad \text{for} \quad i \geqslant 1$$

and

$$v(0) = \min\{r_1\lambda^{-1} - g^*\lambda^{-1} + v(1), K_2 + r_2\lambda^{-1} - g^*\lambda^{-1} + v(1')\} \tag{9}$$

where in (8) $v(i')$ equals the first (second) expression between brackets for $i \geqslant N^*$ ($i = 1, \ldots, N^* - 1$) and in (9) $v(0)$ equals the second expression between brackets. Using (4) the theorem immediately follows from the other assertions of the lemma.

Now we are able to prove the average cost optimality of the N*-policy.

THEOREM 2. The N*-policy is average cost optimal within the class of policies π satisfying

$$\lim_{n\to\infty} \mathrm{E}_\pi[Y_n^2 | X_0 = s]/n = 0 \quad \text{for all} \quad s \in S \tag{10}$$

where Y_n denotes the number of customers in the system and X_n the state of the system at the n-th decision epoch.

PROOF. As under any policy the expected time between two decision epochs is not less than μ, i.e. $\mathrm{E}_\pi[\tau(X_n, a_n)|X_0 = s] \geqslant \mu$ for all $s \in S$ and all $n \geqslant 0$, we have from theorem 1 and the proof of theorem 2 in Ross [5] that the N*-policy is average cost optimal within the class of policies having the property

$$\lim_{n \to \infty} \mathrm{E}_\pi[v(X_n)|X_0 = s]/n = 0 \qquad \text{for all } s \in S$$

which is equivalent with (10).

COROLLARY. The N*-policy is average cost optimal within the class C_M of policies where, for any positive M, we define C_M to be the class of policies under which the server will be always on when the number of customers in the system is larger than the value M.

PROOF. Let N_n be the number of customers just after the service completion epoch of the n-th customer when policy π_0 is used and define $L(t)$ as the number of customers in the system at time t, $t \geqslant 0$, under the same policy. In Tijms [8] it is proved that $\lim_{n \to \infty} n^{-1}\, \mathrm{E}N_n^2 = 0$, which implies

$$\lim_t t^{-1}\mathrm{E}L^2(t) = 0.$$

However, for any policy $\pi \in C_M$ and initial state $s \in S$, we have with probability 1 that $Y_n \leqslant L(t_n) + M$, where t_n denotes the epoch of the n-th decision of policy π. The corollary now follows immediately from theorem 2.

Obviously, any stationary policy which never turns the server off except when the state of the system is 0, and which turns the server on in state $N^{*\prime}$ and in an infinite number of states not belonging to the set $\{1', 2', \ldots, (N^* - 1)'\}$, has the same average cost as the N*-policy. From this fact we obtain by a very simple argument that the N*-policy is also average cost optimal within the class of all stationary policies.

REMARK. If condition (2) is not fulfilled, it can be shown in just the same way that the policy π_0, which always prescribes action 1, is average cost optimal. Its average cost is given in (1).

Acknowledgement

I am indebted to Henk Tijms for his helpful discussions on this work.

References

[1] BELL, C. (1971), Characterization and computation of optimal policies for operating an M/G/1 queueing system with removable server, Operations Research 19, 208–218.

[2] HEYMAN, D. P. (1968], Optimal operating policies for M/G/1 queueing systems, Operations Research 16, 362–382.

[3] LIPPMAN, S. A. (1973), On dynamic programming with unbounded rewards, Working paper no. 212, W. M. S. I., University of California, Los Angeles, U.S.A.

[4] ROSS, S. M. (1970), Applied probability models with optimization applications, Holden-Day Inc., San Francisco, U.S.A.

[5] ROSS, S. M. (1970), Average cost semi-Markov decision processes, J. of Applied Probability 7, 3, 649–656.

[6] SOBEL, M. (1969), Optimal average cost policy for a queue with start-up and shut-down costs, Operation Research 17, 145–162.

[7] TIJMS, H. C. (1974), A control policy for a priority queue with removable server, Operation Research 22, 833–837.

[8] TIJMS, H. C. (1976), Optimal control of the workload in an M/G/1 queueing system with removable server, Mathem. Operationsforschung und Statistik 6, 933–943.

[9] YADIN, M. and P. NAOR (1963), Queueing systems with a removable service station, Operations Research Quarterly 14, 4, 393–405.