AD642277

OPTIMAL OPERATING POLICIES FOR STOCHASTIC SERVICE SYSTEMS

by

Daniel Paul Heyman

OPERATIONS RESEARCH CENTER

COLLEGE OF ENGINEERING

UNIVERSITY OF CALIFORNIA-BERKELEY

OPTIMAL OPERATING POLICIES FOR STOCHASTIC SERVICE SYSTEMS

by

Daniel Paul Heyman
Operations Research Center
University of California, Berkeley

October 1966   ORC 66-31

ABSTRACT

We consider the economic behavior of a queueing system,
operating under a specified linear cost structure,
in which the server may be turned on and off. Optimal
policies for turning the server on and off are derived
for differing assumptions about discounting of future
costs, length of the planning horizon, the form of
the arrival stream, and the number of servers.

The costs imposed are: a server start-up cost, a
server shut-down cost, a cost per unit time when the
server is turned off, a cost per unit time when the
server is turned on, and a holding cost per unit time
spent in the system for each customer. We prove that
for the single server queue there is a stationary
optimal policy of the form: Turn the server on when
n customers are present, and turn it off when the
system is empty.

For the undiscounted, infinite horizon problem with
Poisson arrivals, an exact expression for the cost
rate as a function of n and a closed form expression
for the optimal value of n is derived; bounds are
obtained for the cost rate and optimal policy when
the inter-arrival time distribution is allowed to be
any member of the class of IFR distributions. When
future costs are discounted, we obtain an equation for
the expected discounted cost as a function of n and
the interest rate, and prove that for small interest
rates the optimal discounted policy is approximately
the optimal undiscounted policy. The recursion relationship
to find the optimal (nonstationary) policy for finite
horizons is developed, concluding our results for single
server systems, each channel is restricted to have an
exponential service time distribution (possibly with
different rate), and the arrivals form a Poisson
process. One server is always turned on and the other,
the spare machine, can be turned on and off at arbitrary

times; we show that the stationary optimal policy for
undiscounted costs has the form: Turn the spare
machine on when n customers are present, and turn it
off when m customers are in the system, with
$m \leq 1$. We then derive equations for finding the
optimal values of m and n ; queue disciplines where
customers may be switched from one server to the other
are also considered.

CONTENTS

Chapter I

INTRODUCTION AND PRELIMINARY RESULTS

The similarity between queueing and inventory models has long been recognized; inventory analysis generally includes an explicit cost structure and a solution for optimal policies, but researchers in queueing theory have been more interested in the underlying probabilistic structure. Our research is directed towards finding optimal operating policies for a queueing system with a linear cost structure, with emphasis on models with Poisson arrivals.

1.1 Definitions

Two generic terms will be used throughout, server and customer. A server is a mechanism that performs an operation on units fed into it; these units are referred to as customers. Thus, a server can represent a production line and the customers can represent orders for the product, or the customer could be people arriving at a ticket window, and the server the ticket vendor. When the customer is being processed by the server, he is said to be in service, and the time he spends in service is called his service time. While one customer is in service, other waiting customers are in queue, that is, they are present but have not yet been served; the length of time a customer waits in queue is called his queueing or waiting time, and the queueing time plus the service time of a customer is called his life time. The system is the queue and the server, so the number of customers in the system is the number of customers in queue plus the number of customers in service.

The server may not be allowed to serve arriving customers, i.e., it

may be turned off. Time intervals when the server is turned off will be called dormant periods and during these periods the servers is said to be dormant. When the server is turned on it is running, and time intervals when the server is running are called running periods. Intervals when customers are not being served are idle periods, they occur when no customers are present and/or when the server is dormant; busy periods are time intervals when customers are being served. This definition of busy period corresponds to the usual queueing terminology, but our definition of idle period includes the additional time when customers are present, and the server is dormant. A busy cycle is a consecutive busy and idle period.


Figure 1. A Typical Time Interval

These definitions are easily understood by referring to Figure 1. The interval $[t_2, t_3]$ is a dormant period, (t_3, t_4) is a running period, $(t_1, t_3]$ is an idle period, $(t_3, t_4]$ is a busy period, and during (t_1, t_2) the server may be running or dormant.

The economics of system operation is influenced by the various costs

involved. When the server is dormant, costs for power, heat, maintenance, etc. may still be incurred on a per unit time basis; these are called dormant costs. Activating and deactivating the server may involve power surges, equipment, or manpower charges; the associated costs are called start-up and shut-down costs respectively, and fixed costs collectively. When the server is running, attendant, fuel, and other costs may be charged in addition to the dormant cost; the sum of these costs is the running cost. These four costs are the operating costs of the server, but they may not represent the entire cost picture of an operation; for example in a service such as aircraft repair, the airplanes are not productive during their stay in the repair depot, and this lost time represents a cost to the owner. Thus, we must consider a penalty for delaying the customer in the system, this penalty is called a holding cost.

1.2 Assumptions and Notation

Assumptions for a single server model are:

a) The arrival stream - customers arrive singly in a renewal process; the times between successive arrivals $\{\alpha_i, i=1, 2, ...\}$ have distribution function (d.f.) $A(t)$, $t \geq 0$. For most of our results we will require $A(t) = 1-e^{-\lambda t}$.

b) The service mechanism - customers are processed individually in their order of arrival; the service times $\{\sigma_i, i=1, 2, ...\}$ are independent, identically distributed, non-negative random variables with d.f. $B(t)$, $t \geq 0$.

c) The cost structure -

    i) The dormant cost rate is r_1 [$/hr.];

    ii) the running cost rate is r_2 [$/hr.], $r_2 \geq r_1$;

iii) the start-up cost is R_1[$];

iv) the shut-down cost is R_2[$];

v) the holding cost is h[$/customer-hr.].

All cost coefficients are assumed non-negative and finite, and to avoid trioiality h is positive. Future costs may, or may not be, discounted.

d) **The decision process** - The server may be turned on (or left on) at any point during an idle period, and it may be turned off at service completion epochs during a busy period, (note that this precludes deactivating the server during the service time of a customer).

Every possible policy of turning the server on off during the operating horizon for the the queueing system leads to a different operating cost. The central problem of this work is to find the optimal policy - i.e., the policy which will minimize the total cost of the operating horizon.

Throughout the paper Laplace-Stieltjes transforms (LST'S) of distribution functions will be denoted by a tilde, e.g., $\tilde{A}(s) = \int_0^\infty e^{-st} dA(t)$; moments of d.f.'s will be denoted by a ν with a first subscript indicating the moment and a second subscript indicating the d.f., e.g., the first moment of A(t) is ν_{1A} and the second moment of B(t) is ν_{2B}. The quantity $\rho = \lambda\nu_{1b}$ is the system utilization factor, always assumed to be less than unity.

1.3 Previous Results Used

When the arrivals are a Poisson process and the server is always running, the system forms an M/G/1 queue, a stochastic process that has been studied extensively (see, for example, reference 4). We will often refer to the following results:

At equilibrium, the expected number of customers in the system is given by the Pollaczek-Khinchine equation

(1) $$L = \rho + \frac{\rho^2 (c_B^2 + 1)}{2(1-\rho)},$$

where c_B^2 is the coefficient of variation of $B(t)$.

The length of a busy period, γ, has d.f. $G(t)$, and

(2) $$\tilde{G}(s) = \tilde{B}[s + \lambda - \lambda \tilde{G}(s)].$$

The number of customers served during a busy period, η, and

$$\tilde{F}(z) = \sum_{k=0}^{\infty} z^k P(\eta = k) \text{ satisfies}$$

(3) $$\tilde{F}(z) = z\tilde{B}[\lambda - \lambda \tilde{F}(z)].$$

The probability that the server is busy at an arbitrary point in time is

(4) $$P_B = \rho$$

1.4 The $M_\eta/G/1$ Queue

In the next chapter it will be shown that with Poisson arrivals and the assumptions of Section 1.2 there is an optimal operating policy of the form:

(0)   Turn the server on when n customers are present,

            and then turn it off when the system is empty.

The queueing process formed by this arrival and service pattern will be called an $M_n/G/1$ queue (when the subscript n is set equal to one, we have the ordinary $M/G/1$ queue). The basic properties of this type of queueing system are given by the following theorems:

Theorem 1: The idle period ζ, in a $M_n/G/1$ queue has d.f. $A_n(t)$ and

$$\tilde{A}_n(s) = \left[\frac{\lambda}{s+\lambda}\right]^n.$$

Proof: An idle period will be formed by $n \geq 1$ independent, identically distributed inter-arrival times so $\tilde{A}_n(s) = [\tilde{A}(s)]^n = \left[\frac{\lambda}{s+\lambda}\right]^n$, QED.

Theorem 2: The busy period γ, in an $M_n/G/1$ queue has d.f. $G_n(t)$ and

$$\tilde{G}_n(s) = [\tilde{G}(s)]^n, \quad n \geq 1, \quad \tilde{G}_0(s) = \tilde{G}(s).$$

Proof: Let T_n be the time to serve the n customer present when the busy period begins, and k be the number of customers that arrive during T_n. $P(T_n \leq t) = [B(t)]^{*n} = B_n(t)$ because the service times are independent and identically distributed random variables, and

$$E(z^k) = \int_0^\infty e^{-\lambda(1-z)t} dB_n(t) = [\tilde{B}(\lambda-\lambda z)]^n.$$

$$G_n(t|k,T_n) = T_n * [G(t)]^{*k}$$

$$\tilde{G}_n(s|k,T_n) = e^{-sT_n}[\tilde{G}(s)]^*$$

$$\tilde{G}_n(s|T_n) = \exp\{-T_n[s + \lambda-\lambda\tilde{G}(s)]\}$$

$$\therefore \tilde{G}_n(s) = \{\tilde{B}[s + \lambda-\lambda\tilde{G}(s)]\}^n = [\tilde{G}(s)]^n, \text{ QED.}$$

An immediate corollary is the $\nu_{1G_n} = \dfrac{n\nu_{1B}}{1-\rho}$ since

$$\nu_{1G_n} = \frac{d}{ds}\tilde{G}_n(s)\Big|_{s=0} = n[\tilde{G}(s)]^{n-1}\frac{d\tilde{G}(s)}{ds}\Big|_{s=0} = n\nu_{1G} = \frac{n\nu_{1B}}{1-\rho}.$$

Theorem 3: The number of customers served in a busy period of an $M_n/G/1$
queue has probability generating function $\tilde{F}_n(z) = [\tilde{F}(z)]^n$,
and $\nu_{1F_n} = n\nu_{1F}$, $n \geq 1$.

Proof: This theorem can be proved in a manner analagous to Theorem 2; we
present a more intuitive proof to help in understanding our latter
results.

Instead of serving the original n customers first, serve any
one of them and then all the customers that arrive during his
service, then those that arrive during these services, etc. until
only the n-1 original customers remain in the queue. The number
served up until this point is the number that would be served during
the busy period of an M/G/1 queue, which has p.g.f. $\hat{F}(z)$. Pro-
ceeding in the same manner with the remaining n-1 customers, they
each generate an M/G/1 busy period, and each busy period is inde-
pendent and identically distributed. Thus, $\hat{F}_n(z) = [\hat{F}(z)]^n$ and
$\nu_{1F} = n\nu_{1F}$, QED.

Theorem 4: The asymptotic probability that the server in an $M_n/G/1$ queue
is idle is $P_I = 1-\rho$.

Proof: Since the arrivals are Poisson, the sequence of idle and busy periods
forms an alternating renewal process, and the asymptotic probability
of being in an idle period is $P_I = \dfrac{E(\zeta)}{E(\zeta) + E(\beta)}$. Therefore, for
$n \geq 1$

$$P_1 = \frac{n(\frac{1}{\lambda})}{n(\frac{1}{\lambda} + \frac{\nu_{1B}}{1-\rho})} = \frac{1}{1 + \frac{\rho}{1-\rho}} = 1-\rho$$

and the cases $n=0$ and $n=1$ are the same, QED.

The importance of this theorem is that for all policies of the form (0), the fraction of time the server is busy is the same.

The next four chapters are devoted to finding optimal operating policies for an $M_n/G/1$ queueing system with different assumptions about discounting of total costs and the length of the planning horizon. In Chapter III, some bounds for a $G_n/G/1$ queue will be obtained. The last chapter contains a two-channel model in which the service time distributions are restricted to be exponential.

Chapter II

THE EXISTENCE OF STATIONARY OPTIMAL POLICIES

FOR INFINITE HORIZON PROBLEMS

In Chapters III and IV we will find stationary policies that minimize certain long-term objective functions; the purpose of this chapter is to show that there are no non-stationary policies that give a lower value of the objective function. The method of proof will be by formulating the decision process as a dynamic programming problem, and proving that this problem has a stationary optimal solution.

2.1 Dynamic Programming Formulation

Let[1] S be the set of states of a process and A be the set of acts available; when the process is in state $s \in S$ and the act $a \in A$ is chosen, the process moves to a new state $s' \in S$, where s' is chosen according to some probability distribution depending on s and a, and gives a transition reward $r(s, a, s')$. A policy π specifies which act to choose, at <u>all</u> decision points, as a function of the history of the process. A stationary policy can be represented as a function, f, from S into A such that whenever the process is in state s, the act $a = f(s)$ is chosen. Thus, a stationary policy is independent of the history of the process, except as summarized in its current state.

The total income from a policy is the sum of the transition rewards when that policy is used; when the transition rewards are random variables, the expected income is the sum of the expected transition rewards. If

---

[1] This description of dynamic programming is due to Blackwell, reference 2.

$\mathcal{Q}(\pi,t)$ denotes the expected income earned by policy π at time t, and continuous discounting with interest rate β is employed, then

$I(\pi,\beta) = \int_0^\infty e^{-\beta t} d\mathcal{Q}(\pi,t)$ is the expected discounted income over an infinite

horizon under policy π. The gain rate of policy π is $C(\pi) \overset{\Delta}{=} \lim_{T\to\infty} \frac{1}{T} \int_0^T d\mathcal{Q}(\pi,t)$.

When all the rewards represent costs, the terms total cost and asymptotic

cost rate may be substituted for income and gain rate respectively.

This formulation can be applied to the problems considered in this

thesis in the following manner.

Let A=(1,2) where action 1 is "turn the server off" and action 2 is

"turn the server on"; the state of the system is the pair (n,k) where n is

the number of customers in the system and k indicates if the server is run-

ning or dormant and can be written as the denumerable set

$S = \{0,0',1,1',2,2', \ldots\}$ where the prime indicates the server is running.

The probability law that selects the next state and transition time is:

$s' = s + 1$ if action 1 is used, the transition time is α with d.f. $1-e^{-\lambda t}$;

if action 2 is used, s' will equal $s-1 + \delta(s) + \Psi$, where Ψ is the number

of customers that arrive in a service interval, $\delta(.)$ is the unit pulse at

the origin, the transition time is τ, which has d.f. $B(t)$. The cost of

each transition is the sum of the operating costs for the server and the

holding costs of the customers present during the transition.

Notice that if we allowed the instants when customers arrive during a

busy period to be decision points, a probability distribution for s'

depending only on s and would not be obtained unless $B(t) = 1-e^{-\mu t}$, and

the decision points during an idle period must be restricted to arrival

epochs when the arrival stream is not a Poisson process.

2.2 Problems with Discounted Rewards

When an interest rate $\beta > 0$ is used to discount future rewards, we have

<u>Theorem 1</u> (Blackwell)[2]: If the state space S is a Borel subset of some complete separable metric space, the action space A is finite, and the reward function $r(\cdot)$ is a Baire function on SxAxS, then there is a stationary optimal policy. That is, if $I(\pi,s,\beta)$ is the expected discounted income using policy π and starting in state s, there is a stationary policy π^* such that $I(\pi^*,s,\beta) \geq I(\pi,s,\beta)$ for all policies π and all initial states $s \in S$.

Blackwell proved this theorem when transitions from s to s' occured at times t=1, 2, ... and $r(s,\alpha,s')$ is deterministic. If the inter-transition times are random variables and the reward from a transition depends on this time, the theorem and proof are valid when $r(s,\alpha,s')$ is replaced by its expected value $\bar{r}(s,\alpha,s')$ given that the transition time distributions are such that the length of the experiment t, and the number of transitions n, have the property

$$t \to \infty \Longleftrightarrow n \to \infty \text{ a.s.}$$

If the transition times are not degenerate at zero or defective, a sufficient condition for this property to hold is that there is a finite number, k, of transition time distribution functions, since:

a) If τ_j denotes j-th transition time, the elapsed time after n transitions is $t = \sum_{j=1}^{n} \tau_j$, and $n < \infty$ implies $t < \infty$ a.s., so $t \to \infty \Rightarrow n \to \infty$ a.s. by forming the contra-positive statement.

b) after n transitions at least one of the distributions, $F_1(t)$ say,

must have been the probability law for a minimum of $\frac{n}{k} = n'$ transitions. The elapsed time for these transitions is $t' = \sum\limits_{i=1}^{n'} \tau_i$, where $\tau_i \sim F_1(t)$, and $n \to \infty \Rightarrow n' \to \infty \Rightarrow t' \to \infty$ a.s. $\Rightarrow t \Rightarrow \infty$ a.s.

Since our problem has only two transition time distributions, $A(t)$ and $B(t)$, we conclude <u>there is an optimal stationary policy when future costs are continuously discounted</u>.

2.3 Problems Without Discounting

For the single channel stochastic server system with $\beta=0$, we establish

Theorem 2: There exists a stationary policy f that maximizes the reward rate, and f is independent of the initial state.

Proof: First we show that there is a stationary optimal policy[1]. From the preceding theorem, if a positive interest rate were used, for each initial state there would be a stationary policy $f_s(\beta)$ that would earn the maximum expected discounted reward $I[f_s(\beta)] = I(\pi^*, s, \beta)$.

In Section III it's shown that the holding costs are proportional to β^{-2} and the operating costs are proportional to β^{-1}. This implies that for small interest reates the server must be turned on eventually, i.e., there is a finite bound on n^*. Thus, there is an interest rate β_0, a sequence $\mathcal{B} = \{\beta_i\}_{i=0}^{\infty} \to 0^+$ with $\beta_1 \leq \beta_0$, and a finite set, $\mathcal{F}_s = \{f_s : f_s = f_s(\beta_i)\}$, of stationary

---

[1] This part of the proof is based on the work of Fox, reference 8.

optimal policies. Since for every $\beta_i \epsilon \mathcal{B}$ there's a corresponding f_s, at least one member of $\mathcal{F}_s$ must appear infinitely often in the sequence $\{f_s(\beta_i)\}_{i=1}^{\infty}$. This implies that there is a subsequence of $\mathcal{B}, \mathcal{B}' = \{\beta_j\}_{j=1}^{\infty} \rightarrow 0^+$ and a stationary policy f_s that is optimal for all interest rates $\beta_j \epsilon \mathcal{B}'$.

Using standard Abelian and Tauberian theorems (reference Van der Pol and Bremmer [15] pg. 122),

$$C(\pi) = \lim_{\beta \to 0^+} \beta \int e^{-\beta t} d \mathcal{C}(\pi_s, t, 0)$$

$$= \lim_{\beta \to 0^+} \beta I(\pi_s)$$

$$\leq \lim_{j \to \infty} \beta_j I(f_s) = C(f_s).$$

The result that the optimal stationary policy is independent of the starting state follows from the fact that introducing any finite cost on each of the first finite number of transitions won't alter the asymptotic cost rate. If the service system starts in some state $i > 0$, it will incur only a finite expected cost before there are zero customers in the queue and the server is idle, so $C(f_i) = C(f_o)$. QED.

The above theorem was proved for the particular case of interest rather than as a general dynamic programming result because it is not true in general. Suppose the state space is $S = (0, 1, 2, \ldots)$, and in each state s the alternatives are: 1) go to $s+1$ in one time unit and receive no payment, and 2) go to state $s=0$ in s time units and receive $s - \frac{1}{s}$. Any stationary policy of going to state zero when the process reached state

n would have a rate of $\frac{1}{2}(1 - \frac{1}{n^3})$, and a greater reward rate can be achieved. Furthermore, when the rewards are discounted there <u>will</u> be a stationary optimal policy from theorem 1.

Now that the existence of stationary optimal policies for both models has been established, it is possible to specify the form of such policies..

Theorem 3: Under a stationary optimal policy either: (1) the server will always be dormant, (2) it will always be running, or (3) it will be turned on when there are n^* customers waiting before a dormant server and turned off when the system enters an idle period.

Proof: The first case arises only when discounting is used and the holding cost is small compared to the operating costs, and corresponds to using action 1 in the unprimed states so that the primed states are never reached.

If this is not the form of the policy, there must be a state s where action 2 is optimal (this state need not be unique as illustrated in example 3 of Section 3.4), so $s=n^*$. It remains to be shown that there is no state s' in which action 1 is optimal when process is in state k'. Since this state will be reached with probability one, the process will alternate between k and s customers in the system, once the server is turned on, which is equivalent to incurring the cost of reaching state s and then holding k customers on the side and turning the system on when s-k are present and off when zero are present, a strategy that is dominated by serving the k customers (possibly after a certain finite time τ) and then turning the system on at s-k and off at zero. If $k \geq s$ and action 1 is used in state k', the equivalent cost is achieved by

reaching state s and then holding s customers while turning the
server on at k-s and off at zero, which is dominated by serving the
s customers after holding them no more than a finite time, and then
turning the server on at k-s and off at zero. QED.

This theorem applies only to the allowable decision points, specifically,
only to those points when the server is busy and a service has just been
completed, otherwise it may not be valid. For example, consider a service
time distribution with $P(\sigma) = \left\{ \begin{array}{ll} .999 & \sigma=1 \\ .001 & \sigma=1000 \end{array} \right\}$. Suppose we observe the system
2 time units after a service has begun. We know that this customer will not
complete service for another 998 time units; if the interest rate is high
enough it will be better to turn the system off and incur only the idle
cost for a while, and have the running costs charged when their present value
is lower.

When $\beta=0$, the server will <u>never</u> be turned off when customers are present
because this will only introduce additional on-off costs.

Chapter III

THE UNDISCOUNTED INFINITE HORIZON MODEL

In this chapter we will present two ways to calculate the asymptotic cost rate (see Section 2.1) when a stationary policy of activating the server when n customers are present is employed; we then derive the stationary optimal policy. In Section 3.5 we discuss the effects of a proportional increase in the service and running cost rates; in Section 3.6 bounds on the cost rate and optimal policy are obtained when the inter-arrival times are drawn from an IFR distribution.

3.1 Derivation of the Asymptotic Cost Rate, Method I

Since the server must be turned on eventually, we only need to consider policies where the server is always running, or is turned on when $n \geq 1$ customers are present and off when zero customers are in the system. We will derive the cost rate for the latter form first.

The sequence of busy cycles forms a renewal process, and a basic result of renewal theory is that the asymptotic cost rate is the expected cost per cycle divided by the expected cycle time. The following results are immediate:

(1) The expected length of a busy cycle is $n(\frac{1}{\lambda} + \frac{\nu_B}{1-\rho})$,

(2) The expected dormant time cost in a busy cycle is $\frac{n}{\lambda r_1}$,

(3) The expected running time cost in a busy cycle is $\frac{n\nu_B}{1-\rho} r_2$,

(4) The total set-up cost in a busy cycle is $R_1 + R_2$.

To calculate the expected holding costs we need to know the expected

value of the sum of the life-times of the customers served in a busy period; call this number $W_T^{(n)}$. Consider an M/G/1 queueing system in which the server is always running, and number the busy periods by i and let η_i be the number of customers served in busy period i. The average life time of a customer served in the i-th busy period is $W_i = \frac{1}{\eta_i} (\omega_1^{(i)} + \omega_2^{(i)} + \ldots + \omega_{\eta_i}^{(i)})$, where ω_j^i represents the life time of the j-th customer served during busy period i. Thus, $W_i N_i = \sum_{j=1}^{\eta_i} \omega_j^i$ and multiplying and dividing the right-hand side by $\sum_{i=1}^{n} \eta_i$, dividing both sides by n, rearranging terms and taking limits

$$(1) \qquad \lim_{n \to \infty} \frac{1}{n} \sum_{i=1}^{n} W_i \eta_i = \lim_{n \to \infty} \frac{\sum_{i=1}^{n} \sum_{j=1}^{\eta_i} \omega_j^i}{\sum_{i=1}^{n} \eta_i} \lim_{n \to \infty} \frac{1}{n} \sum_{i=1}^{n} \eta_i.$$

Since busy periods are independent and identically distributed, the left-hand side of (1) is $W_T^{(1)}$ a.s. and the second limit on the right is $E(\eta_i)$ a.s. by the strong law of large numbers. John D. C. Little [11] showed the remaining limit is the expected life time of a customer W, so $W_T = WE(\eta)$.

This result refers to a busy period started by one customer. If n customers start the busy period, the total wait is the same if we serve the first customer in the queue first, then those customers who arrive during this service, then the customers who arrive during these services, and so forth until the busy time generated by this customer is completed, at which time the second of the original customers starts his service, and the process is repeated until the queue is empty.

Since the arrivals are Poisson, the total waiting time generated by each progenitor are independent and have a mean value of W_T. While the first progenitor and his descendents are being served $n-1$ progenitors are waiting, $n-2$ progenitors wait during the busy time generated by the second progenitor, and so forth. This contributed $\frac{n(n-1)}{2} \nu_{1G}$ to the expected total waiting time. During the idle period the queue is building up to n customers, so $\frac{n(n-1)}{2\lambda}$ is added to the expected total waiting time. Adding the parts we have

$$W_T^{(n)} = \left[\frac{W\nu_{1G}}{\nu_{1B}} + \frac{n-1}{2} \left(\nu_{1G} + \frac{1}{\lambda} \right) \right] n.$$

Assembling all the parts, the cost rate $C(n)$ is given by

$$C(n) = \frac{R_1 + R_2 + \frac{nr_1}{\lambda} + nr_1\nu_{1G} + h\left[\frac{W}{1-\rho} + \frac{n-1}{2} \left(\nu_{1G} + \frac{1}{\lambda} \right) \right] n}{n\left(\frac{1}{\lambda} + \nu_{1G} \right)}$$

which simplifies to

$$C(n) = r_1 + (r_2 - r_1)\rho + h\lambda W + \frac{\lambda(1-\rho)(R_1 + R_2)}{n} + h\frac{n-1}{2}.$$

Using the relationship $L = \lambda W$, we have finally,

$$(2) \quad C(n) = r_1 + (r_2 - r_1)\rho + h\left(L + \frac{n-1}{2} \right) + \frac{\lambda(1-\rho)(R_1 + R_2)}{n}.$$

3.2 Derivation of the Asymptotic Cost Rate, Method II

A second method of deriving equation (2) is to observe:

(1) the running cost of the server can be considered as an increment of $r_2 - r_1$ over a constant cost of r_1, and that this increment is incurred a fraction ρ of the time, independent of n, since ρ is the probability that the server is operating at an arbitrary point in time (Theorem 1-4)

(2) $R_1 + R_2$ is incurred during each busy cycle and from the elementary renewal theorem* the number of busy cycles per unit time is

$$\frac{1}{n} \left(\frac{1}{\lambda} + \nu_{1G} \right)^{-1} = \frac{\lambda(1-\rho)}{n}; \text{ and}$$

(3) the holding costs per unit time are proportional to average number of customers in the system.

The average number of customers in the system is $L^{(n)}$, which can be expresses as

$$L^{(n)} = [L^{(n)} | \text{Server Busy}]P(\text{Server Busy}) + [L^{(n)} | \text{Server Idle}]P(\text{Server Idle}).$$

The average number of customers in the system when the server is idle is $\frac{n-1}{2}$, and from the arguments used to derive $w_T^{(n)}$ it's easy to see that the average number of customers in the system when the server is busy is $L^{(1)} + \frac{n-1}{2}$. Therefore,

---

* See Barlow and Proschan [1], Theorem 2.5, for a precise statement and proof.

(3) $L^{(n)} = (L^{(1)} + \frac{n-1}{2})\rho + \frac{n-1}{2}(1-\rho) = L + \frac{n-1}{2}$,

and thus we obtain

$$C(n) = r_1 + (r_2 - r_1)\rho + h(L + \frac{n-1}{2}) + \frac{(R_1 + R_2)\lambda(1-\rho)}{n},$$

which is equation (2).

It's important to note that equation (2) only holds when $n \geq 1$; if the server is always running, there is no idle time, start-up, or shut-down costs, so the cost rate is

(4) $C(0) = r_2 + hL$.

3.3 Determining the Optimal Policy

Setting $\frac{dC(n)}{dn} = 0$ we obtain

(5) $n^* = \sqrt{\frac{2\lambda(R_1 + R_2)(1-\rho)}{n}}$,

(6) $C(n^*) = r_1 + (r_2 - r_1)\rho + h(L - \frac{1}{2}) + [2\lambda(R_1 + R_2)h(1-\rho)]^{\frac{1}{2}}$..

Since $\frac{d^2 C(n)}{dn^2} > 0$, n^* gives the unique minimum value of $C(n)$ for $1 < n < \infty$. The reason that n^* doesn't depend on r_1 and r_2 is because for __all values__ of $n \geq 1$, the fraction of time the server is busy is ρ (Theorem 1-4); the optimal value will only balance the increase in holding costs with the decrease in set-up costs as   is increased. However, the optimal policy depends on r_1 and r_2 because equation (4) may show a cost rate lower than the one given by equation (6).

The optimal policy must be given by an integral value of n, call it n^{opt}. If n^* is bigger than one but is not an integer, the best integer value of n is one of the integers surrounding n^*; therefore, n^{opt} is either one of these integers or zero. The decision is made by evaluating the cost rate given by each of the three candidates.

If $0 \leq n^* \leq 1$, either $n^{opt} = 0$ or $n^{opt} = 1$; since $C(0) < C(1)$ when $(r_2-r_1)(1-\rho) - (R_1 + R_2)\lambda(1-\rho) < 0$, or

$$(7) \qquad\qquad r_2-r_1 < \lambda(R_1 + R_2)$$

so the choice is easily made.

When $n^*=0$, either $\lambda=0$ or $R_1 + R_2 = 0$ so $n^{opt} = 0$. Summarizing, n^{opt} is either 0, $[n^*]$, or $[n^* + 1]$.

3.4 Numerical Examples

Example 1: $\rho=\frac{1}{2}$, $\lambda=1$, $h=1$, $R_1 + R_2 = 4$, $r_1=1$, $r_2=4$

Then $n^* = \ 1 \times 4 \times \frac{1}{2} = 2$, and

$C(2) = 1 + 3 \times \frac{1}{2} + (L + 1) + 4 \times \frac{1}{2} \div 2 = 4\frac{1}{2} + L$, and

$C(0) = 4 + L$.

Therefore, the optimal policy is to leave the server on all the time.

Example 2: $\rho=\frac{1}{2}$, $\lambda=1$, $h=1$, $R_1 + R_2 = 5$, $r_1=1$, $r_2=6$

Then $n^* = \sqrt{\frac{1}{2} \times 2 \times 5} = \ 5 \approx 2.2$, and

$$C(0) = 6 + L,$$

$$C(2) = 1 + \tfrac{1}{2} + (L + \tfrac{1}{2}) + 5 \times \tfrac{1}{2} \div 2 = 5\tfrac{1}{4} + L,$$

$$C(3) = 1 + \tfrac{1}{2} + (L + 1) + 5 \times \tfrac{1}{2} \div 3 = 5\tfrac{1}{3} + L.$$

Therefore, the optimal policy is to turn the server on when 2 customers are present.

Example 3: $p = \tfrac{1}{2}$, $\lambda = 1$, $h = 2$, $R_1 + R_2 = 4$, $r_1 = 1$, $r_2 = 5$

Then $n^* = \sqrt{2 \times \tfrac{1}{2} \times 4} = 2$, and

$$C(0) = 5 + 2L,$$

$$C(1) = 1 + 2 + 2L + \tfrac{1}{2} \times 4 = 5 + 2L,$$

$$C(2) = 1 + 2 + 2(L + \tfrac{1}{2}) + \tfrac{1}{2} \times 4 \div 2 = 5 + 2L.$$

Therefore, there are 3 optimal policies.

3.5 A Control Model

We have been considering the service time as a random variable σ with d.f. $B_\sigma(t)$. The random varialbe $\sigma' = \theta\sigma$, which is σ mapped onto a different time scale, has the properties:

(1) $B_{\sigma'}(t) = B_\sigma(\tfrac{t}{\theta})$,

(2) $E(\sigma') = \theta E(\sigma)$

(3) $c_{\sigma'}^2 = c_\sigma^2$.

When $\theta < 1$, σ' is stochastically smaller than σ so this can be interpreted as an improvement in the service facility; similarly $\theta > 1$ represents

a decrease in the capability of the server.

Suppose we can change σ to σ' with a corresponding change in running cost from r_2 to $r_2' = \dfrac{r_2}{\theta}$, with all other costs remaining the same. In the undiscounted case, the optimal policy with respect to opening the server becomes

$$(8) \qquad n^*(\theta) = \sqrt{\frac{2(R_1 + R_2)\lambda(1-\theta\rho)}{h}},$$

since $\rho' = \lambda E(\theta\sigma) = \theta\rho$. Since the optimal policy for serving customers is uniquely determined by equation (1), the cost rate can be expressed as a function of θ only, VIZ.

$$(9) \quad C(\theta) = r_1(1-\theta\rho) + r_2\rho + h\left[L(\theta) + \frac{n^*(\theta)-1}{2}\right] + (R_1 + R_2)\frac{1-\theta\rho}{n^*(\theta)}$$

where $L(\theta) = \theta\rho + \dfrac{\theta^2\rho^2(C_B^2 + 1)}{2(1-\theta\rho)}$ from the Pollaczek-Khinchin equation.

One can find the optimal value of θ by investigating the roots of

$$(10) \quad \frac{dC(\theta)}{d\theta} = h\left[\rho + \frac{\theta\rho^2(C_B^2 + 1)(2-\theta\rho)}{2(1-\theta\rho)^2}\right] - \rho[2(1-\theta\rho)]^{-\frac{1}{2}}[\lambda h(R_1 + R_2)]^{\frac{1}{2}} + r_1 = 0.$$

A solution to equation (10) can't be given in closed form; however, one can obtain the conditions for which it's desirable to slow the server down. The initial control setting is $\theta^0 = 1$; $C(\theta)$ decreases as θ increases when

$$\frac{dC(\theta)}{d\theta}\Big|_{\theta=1} = h\left[\rho + \frac{\rho^2(2-\rho)(C_B^2 + 1)}{2(1-\rho)}\right] - \rho(\lambda(R_1 + R_2)h)^{\frac{1}{2}}[2(1-\rho)]^{-\frac{1}{2}} - r_1\rho < 0.$$

Simplifying, we have the local condition for slowing the server down,

$$
(11) \qquad \left[\frac{\lambda(R_1 + R_2)}{2h(1-\rho)}\right]^{\frac{1}{2}} + \frac{r_1}{h} > 1 + \frac{\rho(2-\rho)(C_B^2 + 1)}{2(1-\rho)};
$$

if the reverse inequality holds, the server should be speeded up.

The qualitative reasoning that explains the form of this result is that increasing the service time will lengthen the busy periods and reduce the fraction of time the server is idle, which is desirable if $R_1 + R_2$ or r_1 is large. It also has the effect of increasing the mean number of customers in the system, which is undesirable if h is the dominant cost.

3.6 General Arrival Distributions

Equation (2) is not valid for G/G/l queues because $\nu_{1G_n} \neq n\nu_{1G}$; however, we expect that equation (5) will be approximately correct when ρ is close to unity, giving a policy that in a highly utilized system the server should always be available, unless the holding costs are negligible. For arbitrary values of ρ we can obtain bounds on the cost rate and optimal policy if the arrival distribution is suitably restricted. Generalizing the arrival distribution requires that the decision points during an idle period be restricted to arrival epochs, as pointed out in Section 2.1.

A class of distributions that has been given much attention in reliability models is the class of IFR distributions, where IFR stands for increasing failure rate. The IFR assumption for $A(t)$ corresponds to arrival patterns where the probability that a customer arrives in the next time increment increases as the interval since the last arrival epoch increases; imposing the IFR assumption on the arrival stream is natural for models where the customers represent failed pieces of equipment and the service

mechanism is the repair facility. Since the Erlang distributions and many other common distributions are IFR, assuming $A(t)$ is IFR is a non-parametric way of analyzing a large family of common queueing models.

A distribution $A(t)$ is IFR if

$$\frac{A(t + x) - A(t)}{A^c(t)}$$

is increasing in t for all $t > 0$; the properties of these distributions are discussed in Barlow and Proschan [1].

Marshall [12] found bounds for the expected number of customers in the system and the expected length of the busy cycle for the $G_n/G/1$ queue we shall consider. When $A(t)$ is IFR and $\rho < 1$, he obtained:

$$(12) \qquad A \leq L < + \tfrac{1}{2}, \quad A = \rho + \frac{\rho[c_A^2-1] + \rho^2[c_B^2 + 1]}{2(1-\rho)} + \frac{n-1}{2}$$

where c_A^2 and c_B^2 are the coefficients of variation of $A(t)$ and $B(t)$ respectively, and

$$(13) \qquad \frac{(1-\rho)(n-\rho)}{\lambda} \leq E(X) \leq \frac{n(1-\rho)}{\lambda}$$

where $E(X)$ is the expected length of a busy cycle.

Since the cost rate is highest when L assumes its largest value and $E(X)$ its smallest, and the minimum cost rate is achieved when these conditions are reversed, the cost rate is bounded for all n by

$$r_1(1-\rho) + r_2\rho + h\lambda A + \frac{(R_1 + R_2)\lambda}{n(1-\rho)} \lesseqgtr C(0) < r_1(1-\rho)$$

(14)

$$+ r_2\rho + h(\lambda A + \tfrac{1}{2}) + \frac{(R_1 + R_2)\lambda}{(1-\rho)(n-\rho)}, \quad n \geq 1$$

$$r_2 + \left[\rho + \frac{\rho[c_A^2 - 1] + \rho^2[c_B^2 + 1]}{2(1-\rho)}\right] \leq C(0) < r_2$$

$$+ h\left[\rho + \frac{\rho[c_A^2 - 1] + \rho^2[c_B^2 + 1]}{2(1-\rho)} + \tfrac{1}{2}\right].$$

The width of the bounding interval is easily calculated to be

$\dfrac{(R_1 + R_2)\lambda\rho}{n(1-\rho)(n-\rho)} + \dfrac{h}{2}$, $n \geq 1$, which decreases as n increases and ρ decreases.

Let $C_u(n)$ and $C_l(n)$ be the upper and lower bounds; simple differentiation reveals that both are convex, so $C(n)$ is bounded by an interval as shown in Figure 1.

Figure 1. Bounds for the Cost Rate of the IFR$_n$/G/1 Queues.

Setting the derivative of $C_u(n)$ with respect to n equal to zero yields

$$(15) \qquad n_u^* = \sqrt{\frac{2\lambda(R_1 + R_2)}{h(1-\rho)}} + \rho > n^*,$$

where n^* is the optimal policy when the arrivals are Poisson (equation 5), and

$$(16) \qquad C_u(n_u^*) = r_1 + (r_2 - r_1)\rho + h\left[\frac{3}{2}\rho + \frac{\rho[c_A^2 - 1] + \rho^2[c_B^2 + 1]}{2(1-\rho)}\right]$$

$$+ \left[\frac{2\lambda(R_1 + R_2)h}{1-\rho}\right]^{\frac{1}{2}}.$$

Since n_u^* obtains the minimum value of the maximum cost rate, it is the minimax strategy and $C_u(n_u^*)$ is the upper bound for the optimal cost rate. In a similar manner, one finds that $C_l(n)$ is minimized by

$$(17) \qquad n_l^* = \sqrt{\frac{2(R_1 + R_2)\lambda}{h(1-\rho)}} > n^*$$

and

$$(18) \qquad C_l(n_l^*) = r_1 + (r_2 - r_1)\rho + h\left[\rho + \frac{\rho[c_A^2 - 1] + \rho^2[c_B^2 + 1]}{2(1-\rho)} - \frac{1}{2}\right]$$

$$+ \left[\frac{2\lambda(R_1 + R_2)h}{1-\rho}\right]^{\frac{1}{2}}.$$

Since $C_l(n_l^*)$ is the lowest possible cost rate and

$$(19) \qquad C_u(n_u^*) - C_l(n_l^*) = \frac{h(1+\rho)}{2} < h,$$

the optimal cost rate is bounded within $\frac{h}{2}$ of its value if $A(t)$ was completely specified.

From the diagram it is easy to see that the values n_1 and n_2 that satisfy $C_1(n_1) = C_1(n_2) = C_u(n_u^*)$ are the bounds on the optimal policy. Using equations (14) and (16) one obtains

$$(20a) \quad n_1 = n_u^* + 1-2\rho - \left[(1-\rho)^2 + 4\left[\frac{-2\lambda(R_1 + R_2)(1-\rho)}{h}\right]^{\frac{1}{2}}\right]^{\frac{1}{2}}$$

$$(20b) \quad n_2 = n_u^* + 1-2\rho + \left[(1-\rho)^2 + 4\left[\frac{-2\lambda(R_1 + R_2)(1-\rho)}{h}\right]^{\frac{1}{2}}\right]^{\frac{1}{2}}$$

$$(20c) \quad n_2 - n_1 \approx 2\sqrt{n^*},$$

where n^* is given by equation (5). Since small values of h cause large values of n^*, equations (19) and (20c) imply that as the bounds on the optimal cost rate get better, the bounds on the optimal policy will get worse. This means that large variations in the policy will have small effects on the expected cost rate.

3.7 Conclusions

From equations (6), (16), and (18) we see that the expected cost rate can be reduced by decreasing the variability of the service and inter-arrival time; thus, _ceterus paribus_, the queue with constant inter-arrival and service times will have the smallest cost rate.

If we assume a Poisson arrival stream but $A(t)$ is actually IFR, substituting equation (5) into the equation for $c_u(n)$, we find

$$(20) \quad C_u(n^*) - C_u(n_u^*) = \frac{\lambda(R_1 + R_2)}{(1-\rho)(n^*-\rho)} - \frac{(1 + \rho)h}{2}(n^*-1).$$

Since small operating costs and large holding costs tend to give the policies n^{opt} equal zero or one for both Poisson and IFR arrivals, the error given by equation (2) will be large only when n^* is large. Thus, when the utilization and holding costs are large and the operating costs are small, the loss in generality in assuming Poisson arrivals is not very critical.

Chapter IV

THE DISCOUNTED INFINITE HORIZON MODEL

In this chapter we will find the expected present value (discounted total return) of a stationary policy that turns the server on when n customers are present, and turns it off when the system is empty. The analysis in this chapter is more complicated than that of Chapter III because explicit account must be taken of the epochs when costs are incurred.

4.1 Some Properties of Laplace-Stieltjes Transforms

We will use the following properties of Laplace-Stieltjes transforms extensively in this section:

<u>Property 1</u>: If a cost R is incurred τ time units from now, and τ is a random variable with d.f. $F(t)$, then the expected present value of this cost is $R\widetilde{F}(\beta)$.

<u>Proof</u>: Assume that $\tau=t$, then the discounted cost is $Re^{-\beta t}$, unconditioning on $\tau=t$ we have the expected cost is $\int_0^\infty Re^{-\beta t}dF(t) = R\widetilde{F}(\beta)$, QED.

<u>Property 2</u>: If a cost rate of r dollars per unit time is incurred until τ units from now, and τ is a random variable with d.f. $F(t)$, then the expected present value of the cost is $\frac{r}{\beta}[1-\widetilde{F}(\beta)]$.

<u>Proof</u>: Assume $\tau > t$, the present value of the cost at time t is $re^{-\beta t}$. The probability that $\tau > t$ is $1-F(t)$ so the expected present value of the cost at time t is $\int_0^\infty re^{-\beta t}[1-F(t)]dt = \frac{r}{\beta} - r\int_0^\infty e^{-\beta t}F(t)dt = \frac{r}{\beta}[1-\widetilde{F}(\beta)]$, QED.

<u>Property 3</u>:   If $F(t)$ is a distribution function of a non-negative random variable that is not degenerate at the origin, then if $\widetilde{F}(\beta) =$

$$\int_0^\infty e^{-\beta t} dF(t) \text{ exists, } 0 < \widetilde{F}(\beta) < 1 \text{ for } \beta > 0.$$

<u>Proof</u>: $\widetilde{F}(\beta) > 0$ since $e^{-\beta t} > 0$ and $dF(t) \geq 0$ but not identically zero.

$$\widetilde{F}(\beta) < 1 \text{ since } \left| \int_0^\infty e^{-\beta t} dF(t) \right| \leq \left| e^{-\beta t} \right| \left| \int_0^\infty dF(t) \right| < 1.$$

4.2 Calculation of the Expected Total Cost

Using these results, we first find the expected cost, $\mathcal{C}(n;\beta)$, for $n \geq 1$, starting with zero customers in the queue. The first two properties show that in the first busy cycle:

(1)   the expected start-up cost is $R_1[\widetilde{A}(\beta)]^n = R_1 (\frac{\lambda}{\beta+\lambda})^n \triangleq c_1$

(2)   the expected shut-down cost is $R_2[\widetilde{A}(\beta)\widetilde{G}(\beta)]^n = R_2 (\frac{\lambda}{\beta+\lambda})^n [\widetilde{G}(\beta)]^n \triangleq c_2$

(3)   the expected idle time cost is $\frac{r_1}{\beta} [1-(\widetilde{A}(\beta))^n] = \frac{r_1}{\beta} \left[1-(\frac{\lambda}{s+\lambda})^n \right] \triangleq c_3$

(4)   the expected running cost is $\frac{r_2}{\beta} \left[1-[\widetilde{G}(\beta)]^n \right] [\widetilde{A}(\beta)]^n = \frac{r_2}{\beta} (\frac{\lambda}{\beta+\lambda})^n$

$$\left[1-(\widetilde{G}(\beta))^n \right] \triangleq c_4$$

The LST of a busy cycle is $\widetilde{H}_n(\beta) = [\widetilde{A}(\beta)\widetilde{G}(\beta)]^n$, so the expected discounted operating costs for the entire planning period is

(1) $\Omega(n;\beta) = c + cH_n(\beta) + c[H_n(\beta)]^2 + \ldots = \dfrac{c}{1-H_n(\beta)}$

where $c = c_1 + c_2 + c_3 + c_4$, and $0 < H_n(\beta) < 1$ from property 3.

The holding costs will be calculated in two parts. When the queue size is increasing from zero to n during an idle period, one customer waits for n-1 arrivals which cost $\frac{h}{\beta} \widetilde{A}(\beta) \left[1 - [\widetilde{A}(\beta)]^{n-1} \right]$ since he starts waiting one inter-arrival time from the beginning of the busy cycle. The second customer arrives after two inter-arrival intervals and waits for n-2 more arrivals, so the holding cost is $\frac{h}{\beta} [\widetilde{A}(\beta)]^2 \left[1 - [\widetilde{A}(\beta)]^{n-2} \right]$; the costs of the remaining customers is computed in the same way, and the total discounted cost is

$$\iota_q(n;\beta) = \frac{h}{\beta} \left\{ \widetilde{A}(\beta)[1 - (\widetilde{A}(\beta))^{n-1}] + [\widetilde{A}(\beta)]^2 [\widetilde{A}(\beta))^{n-2}] + \ldots \right.$$

$$+ (\widetilde{A}(\beta))^{n-1}[1 - \widetilde{A}(\beta)]$$

$$= \frac{h}{\beta} \left[\frac{\widetilde{A}(\beta) - \widetilde{A}(\beta))^n}{1 - \widetilde{A}(\beta)} - (n-1)(\widetilde{A}(\beta))^n \right]$$

(2)
$$\iota_q(n;\beta) = \frac{h}{\beta} \frac{\lambda(\beta+\lambda)^n - \lambda^n(\lambda+n\beta)}{\beta(\beta+\lambda)^n}$$

The expected discounted total cost over an infinite horizon is

(3)
$$\iota_q(n;\beta) = \frac{\iota_q(n;\beta)}{1 - \widetilde{H}_n(\beta)} = \frac{h}{\beta^2} \frac{\lambda(\beta+\lambda)^n - \lambda^n(\lambda+n\beta)}{(\beta+\lambda)^n - [\lambda\widetilde{G}(\beta)]^n}.$$

The remaining costs are due to customers who are waiting while the server is busy. If $\iota_s(n;\tau)$ is the expected number of customers in the system at time τ, and τ is a point in the first busy period, the expected

discounted cost of holding these customers until $\tau + d\tau$ is $he^{-\beta\tau}\iota_s(n,\tau)d\tau$

and the cost for the busy period is $\iota_s(n;\beta) = h\int_0^\infty e^{-\beta\tau}\iota_s(n;\tau)d\tau$, where τ is

a random variable with d.f. $G_n(t)$.

Let θ equal the service time of the first customer served in a busy

period and k be the number of arrivals during θ, then

$$(4) \qquad \iota_s(1;\beta|\theta,k) = h\int_0^\infty e^{-\beta\tau}\iota_s(1;\tau|\theta,k)d\tau + e^{-\beta\theta}\iota_s(k;\beta)$$

$$= h\int_0^\theta e^{-\beta t}(1 + kt)dt + e^{-\beta\theta}\iota_s(k;\beta).$$

Changing the order of service as in Section 3, and using properties

one and two, we find that

$$\iota_s(n;\beta) = \iota_s(1;\beta)[1 + \widetilde{G}(\beta) + (\widetilde{G}(\beta))^2 + \ldots + (\widetilde{G}(\beta))^{n-1}]$$

$$(5) \qquad + h\frac{n-1}{\beta}[1-\widetilde{G}(\beta)] + h\frac{n-2}{\beta}[1-\widetilde{G}(\beta)]\widetilde{G}(\beta) + \ldots + h\frac{1}{\beta}[1-\widetilde{G}(\beta)]$$

$$[\widetilde{G}(\beta)]^{n-2} = \iota_s(1;\beta)\frac{1-[\widetilde{G}(\beta)]^n}{1-\widetilde{G}(\beta)} + \frac{h}{\beta}\left[n - \frac{1-[\widetilde{G}(\beta)]^n}{1-\widetilde{G}(\beta)}\right].$$

Substituting equation (5) into equation (4) and unconditioning on k,

$$\iota_s(1;\beta|\theta) = h\int_0^\theta e^{-\beta t}(1+\lambda t)dt + e^{-\beta\theta}\left\{\iota_s(1;\beta)\frac{1-\exp[-\lambda\theta(1-\widetilde{G}(\beta))]}{1-\widetilde{G}(\beta)}\right.$$

$$\left.+ \frac{h}{\beta}\left[\lambda\theta - \frac{1-\exp[-\lambda\theta(1-\widetilde{G}(\beta))]}{1-\widetilde{G}(\beta)}\right]\right\};$$

integrating the first term by parts and simplifying the second term gives

(6) $\iota_s(1;\beta|\theta) = -\dfrac{\beta+\lambda}{\beta^2}(e^{-\beta\theta}-1)h + \dfrac{1-\exp[-\lambda\theta(1-\widetilde{G}(\beta))]}{1-\widetilde{G}(\beta)}[\iota_s(1;\beta) - \dfrac{h}{\beta}]e^{-\beta\theta}.$

Unconditioning on θ and recalling $\widetilde{G}(\beta) = \widetilde{B}[\beta+\lambda-\lambda\widetilde{G}(\beta)]$,

$$\iota_s(1;\beta) = \dfrac{\beta+\lambda}{\beta^2}h + [\iota_s(1;\beta) - \dfrac{h}{\beta}]\dfrac{\widetilde{B}(\beta)-\widetilde{G}(\beta)}{1-\widetilde{G}(\beta)}$$

from which one obtains

(7) $$\iota_s(1;\beta) = \dfrac{\beta+\lambda}{\beta^2}[1-\widetilde{G}(\beta)]h + \dfrac{\widetilde{G}(\beta)-\widetilde{B}(\beta)}{\beta[1-\widetilde{B}(\beta)]}h.$$

Substituting equation (7) into equation (5), the expected cost of a busy cycle is seen to be

(8) $$\iota_s(n;\beta) = \dfrac{[1-(\widetilde{G}(\beta))^n][\lambda-(\lambda+\beta)\widetilde{B}(\beta)]h}{\beta^2[1-\widetilde{B}(\beta)]} + \dfrac{nh}{\beta};$$

when $h=1$, $\iota_s(n;\beta)$ is the Laplace transform of the mean number of customers in the system during a busy period[*]. The expected costs for the entire horizon is

(9) $L_s(n;\beta) = \dfrac{\iota_s(n;\beta)}{1-H_n(\beta)} = \dfrac{[1-(\widetilde{G}(\beta))^n][\lambda-(\lambda+\beta)\widetilde{B}(\beta)]h}{\beta^2[1-\widetilde{B}(\beta)][1-(\frac{\lambda\widetilde{G}(\beta)}{\lambda+\beta})^n]} + \dfrac{nh}{\beta[1-(\frac{\lambda\widetilde{G}(\beta)}{\lambda+\beta})^n]}.$

Thus, the total expected discounted cost function is, finally:

---

[*] This transform can be computed from the results of Gaver [9] but this derivation is more direct.

(10) $\quad \mathcal{C}(n;\beta) = \Omega(n;\beta) + L_q(n;\beta) + L_s(n;\beta)$

$$= \frac{\beta R_1 + \beta[G(\beta)]^n R_2 + (r_2 - r_1)[1 - (\widetilde{G}(\beta))^n]}{\beta\left[\frac{(\lambda+\beta)\widetilde{G}(\beta)}{\lambda}\right]} + \frac{r_1}{\beta}$$

$$+ \frac{h}{\beta^2} \frac{\lambda(\lambda+\beta)^n - \lambda^n(\lambda+n\beta)}{(\beta+\lambda)^n - [\lambda\widetilde{G}(\beta)]^n}$$

$$+ \frac{[1-(G(\beta))^n][\lambda - (\lambda+\beta)\widetilde{B}(\beta)]h}{\beta^2[1-\widetilde{B}(\beta)]\left[1-(\frac{\lambda\widetilde{G}(\beta)}{\lambda+\beta})^n\right]}.$$

In principle, one can find the minimizing value of r, $n^*(\beta)$ by classical calculus methods; however, a closed form solution for the usual first order conditions doesn't exist. Equation (10) contains terms that are linear in n and terms in which n is an exponent, hence, the derivative of $\mathcal{C}(n;\beta)$ with respect to n will also contain terms of these forms. This implies that $\frac{d}{dn}\mathcal{C}(n;\beta) = 0$ is an implicit equation in n, so there is no closed form expression for its roots.

Thus, an indirect method for calculating $n^*(\beta)$ is required. For example, a computer solution using successive approximations could be used; for small values of β, one could approximate $\frac{d}{dn}\mathcal{C}(n;\beta)$ by a Maclauren's series expansion.

In the discounted case both endpoints must also be considered as possible minimizing values of n. When n=0, R_1 is charged when the first customer arrives, the running cost is always charged, and there are no customers who wait for the server to be turned on, so

(11) $\mathcal{C}(0;\beta) = \dfrac{R_1\lambda}{\beta+\lambda} + \dfrac{r_2}{\beta} + h \left\{ \dfrac{\lambda}{\beta^2} + \dfrac{\lambda\widetilde{B}(\beta)[1-\widetilde{G}(\beta)]}{\beta[s+\lambda-\lambda\widetilde{G}(\beta)][1-\widetilde{B}(\beta)]} \right\}$

The interpretation of $n = \infty$ is that the server is never turned on. In this case, the only costs incurred are holding costs in the queue, and

(12) $\mathcal{C}(0;\beta) = \lim_{n\to\infty} L_q(n;\beta) = \dfrac{h\lambda}{\beta^2}.$

We note that the undiscounted cost rate [Equation (2), Chapter III] depended only on the first and second moments of the service distribution, while $(n;\beta)$ in (10) depends on the knowledge of the entire distribution, through $\widetilde{B}(\beta)$ and $\widetilde{G}(\beta)$.

4.3 Limiting Results When the Interest Rate Vanishes

As β vanishes, the expected cost given by $\mathcal{C}(n;\beta)$ approaches infinity; an important result is that as β approaches zero, the discounted model gives the same cost rate as the undiscounted model.

Theorem 1: $\displaystyle\lim_{\beta\to 0^+} \beta\mathcal{C}(n;\beta) = C(n)$, where $C(n)$ is given by equation (III-2).

Proof: Recall that $\mathcal{C}(n;\beta)$ and $C(n)$ were defined in Chapter II as:

$\mathcal{C}(n;\beta) = \displaystyle\int_0^\infty e^{-\beta t} d\mathcal{C}(n,t),\ C(n) = \lim_{T\to\infty} \frac{1}{T}\int_0^T d\mathcal{C}(n,t),$ where $C(n,t)$ is

the cumulative costs incurred at time t using a policy that turns the server on when n customers are present. Since $\mathcal{C}(n,t)$ is a monotone function because all the costs are non-negative, we can apply a standard Tauberian theorem for LST's and

$$\underset{\beta \to 0^+}{\text{Lim}} \beta \, \mathcal{C}(n;\beta) = \underset{\beta \to 0^+}{\text{Lim}} \beta \int_0^\infty e^{-\beta t} d\mathcal{C}(n;t) = \underset{T \to \infty}{\text{Lim}} \frac{1}{T} \int_0^\infty d\mathcal{C}(n,t)$$

$$= C(n), \text{ QED}.$$

Theorem 1 doesn't <u>necessarily</u> imply that $n^*(\beta) \to n^*$ as $\beta \to 0$ because examples have bee.: constructed for other models (see for example [2]) where the discounted cost function satisfies Theorem 1 but the policy for small interest rates is not "close" to the undiscounted policy. However, for our problem, it <u>is</u> possible to show that:

<u>Theorem 2</u>: $n^*(\beta) \to n^*$ as $\beta \to 0$.

<u>Proof</u>: Since n is treated as a continuous variable over the range $[1,\infty)$ the derivatives $\frac{d}{dn}\mathcal{C}(n;\beta)$ and $\frac{d}{dn} C(n)$ exist over $(1,\infty)$, and

$$\underset{\beta \to 0}{\lim} \beta \, \mathcal{C}(n;\beta) = C(n),$$

(13) $$\underset{\beta \to 0}{\lim} \beta \frac{d}{dn} \mathcal{C}(n;\beta) = \frac{d}{dn} \underset{\beta \to 0}{\lim} \beta\mathcal{C}(n;\beta) = \frac{d}{dn} C(n).$$

The meaning of equation (13) is that there exists a small positive number which may depend on β and n, $\epsilon(\beta,n)$ say, such that

(14) $$\beta \frac{d}{dn} \mathcal{C}(n;\beta) = \frac{d}{dn} C(n) + \epsilon(\beta,n)$$

and $\epsilon(\beta,n) \to 0$ as $\beta \to 0$. Pick a small value of β greater than zero, β^o, then equation (14) implies

$$\frac{d}{dn} \mathcal{C}(n;\beta^o) = \frac{1}{\beta^o} \left[\frac{d}{dn} C(n) + \epsilon(\beta^o,n) \right]$$

and when $n = n^*(\beta^0)$,

(15) $\quad \frac{d}{dn}\mathcal{C}(n,\beta^0)\big|_{n=n^*(\beta^0)} = \frac{1}{\beta}\left[\frac{d}{dn} C(n)\big|_{n=n^*(\beta^0)} + \epsilon(\beta^0,n)\right] = 0.$

Since $\frac{1}{\beta^0} \neq 0$, the term in square brackets must equal zero, which from equation (III-2) means

$$\frac{(R_1 + R_2)\lambda(1-\rho)}{[n^*(\beta^0)]^2} + \frac{h + 2\epsilon(\beta^0,n)}{2} = 0.$$

This implies

(16) $\quad n^*(\beta^0) = \sqrt{\frac{2\lambda(R_1 + R_2)(1-\rho)}{h + 2\epsilon(\beta^0,n)}} \rightarrow n^* \text{ as } \beta^0 \rightarrow 0.$

It remains to be shown that $n^*(\beta)$ is close to n^* for small values of β. We have

$$[n^*]^2 = \frac{2R\lambda(R_1 + R_2)(1-\rho)}{h} \quad \frac{h + 2\epsilon(\beta,n)}{h + 2\epsilon(\beta,n)} \quad \text{and}$$

$$[n^*(\beta)]^2 = \frac{2\lambda(R_1 + R_2)(1-\rho)}{h + 2\epsilon(\beta,n)} \quad \frac{h}{h}$$

from which

(17) $\quad n^*(\beta) = n^* \sqrt{1 - \frac{2\epsilon(\beta,n)}{h + 2\epsilon(\beta,n)}}$

The square root term can be made arbitrarily close to unity by picking β close enough to zero, QED.

Equation (17) shows that convergence is faster as h is larger. This is because for small values of h, n^* is very large, but if the interest rate is large enough $n^*(\beta)$ is infinite (i.e., never turn the server on), and as the interest rate is decreased, $n^*(\beta)$ decreases and the approximation of $n^*(\beta)$ by n^* becomes useful.

The operational implication of Theorem 2 is that for small values of β, one can use n^*, which is easy to compute, to estimate the integer value of $n(\beta)$ that minimizes $\mathcal{C}(n;\beta)$. As before, this integer value is one of the integers surrounding $n^*(\beta)$, if $n^*(\beta)$ is not an integer, because of

Theorem 3: For β sufficiently small, $\mathcal{C}(n;\beta)$ is convex in β.

Proof: Differentiating both sides of equation (5),

$$\lim_{\beta \to 0} \beta \, \frac{d^2}{dn^2} \, \mathcal{C}(n;\beta) = \frac{d^2}{dn^2} \, C(n)$$

which means there exists a function $\delta(\beta,n)$, such that for small values of β

(18) $\qquad \beta \, \dfrac{d^2}{dn^2} \mathcal{C}(n;\beta) = \dfrac{d^2}{dn^2} \, C(n) + \delta(\beta,n) = \dfrac{R\lambda(1-\rho)}{n^3} + \delta(\beta,n),$

and $\delta(\beta,n)$ can be made as small as desired by picking β sufficiently close to zero. In particular, since we don't know the sign of $\delta(\beta,n)$, pick β small $|\delta(\beta,n)| < \dfrac{R\lambda(1-\rho)}{n^3}$; this insures that the right hand side of equation (18) is positive, and since $\beta > 0$,

$\dfrac{d^2}{dn^2} \mathcal{C}(n;\beta) > 0$, QED.

4.4 Arbitrary Initial Conditions

The above development has assumed that the system is initially empty, and the discussion in Section II indicated that the stationary optimal policy may be dependent on the initial state. If there are m customers present initially, the total cost function, call it $\mathcal{C}_m(n;\beta)$, can be constructed from our present results.

Since the policy is assumed stationary when we reach a regular busy cycle, the server will not be turned on until n customers are present; if m < n, the server will be turned on after n-m arrivals; if m ≥ n, the server will be turned on immediately. Thus, if k = Max [0,m-n] is the first time, and n is the number present when the server is turned on thereafter, we can express the expected discounted cost function in terms of k and n, and then look for the minimizing values.

Since there will be zero customers in the system at the end of the first busy period, $\mathcal{C}_m(k,n;\beta)$ equals the cost of the first busy period plus the present value of $\mathcal{C}_o(n;\beta)$ started after k customers have arrived and the first busy period is completed, thus

(19) $\quad \mathcal{C}_m(k,n;\beta) = R_1\left(\frac{\lambda}{\lambda+\beta}\right)^k + R_2\left(\frac{\lambda}{\lambda+\beta}\right)^k [G(\beta)]^{m+k}$

$$+ \frac{r_1}{\beta}\left[1-\left(\frac{\lambda}{\lambda+\beta}\right)^k\right] + \frac{r_2}{\beta}\left(\frac{\lambda}{\beta+\lambda}\right)^k [1-(\tilde{G}(\beta))^{m+k}]$$

$$+ \frac{L_s(n;\beta)}{1-H_n(\beta)} + \left(\frac{\lambda}{\beta+\lambda}\right)^k [G(\beta)]^{k+n}\mathcal{C}_o(n;\beta).$$

Minimizing values must be found numerically or by some approximation technique.

Equation (19) indicates that the optimal policy will, in general, depend

on on. However, as $\beta \to 0^+$ $\beta_m(k,n;\beta) \to C(n)$, and the optimal policy becomes independent of m.

4.5 An Alternate Formulation

In many applications the value of h is not known, e.g., the holding cost of military equipment in a repair depot, and one might formulate the problem as minimizing operating costs subject to some operating constraints such as the expected life time of a customer. The next theorem shows that this minimization will be trivial.

<u>Theorem 4</u>: The operating cost $\Omega(n;\beta)$ is a strictly convex, monotone decreasing function of n.

<u>Proof</u>: When $\beta=0$ this theorem refers to the operating cost rate, which is
$$r_1 + (r_2 - r_1)\rho + \frac{\lambda(1-\rho)(R_1 + R_2)}{n}$$ which is obviously strictly convex and monotone decreasing in n (we ignore the completely trivial case where $R_1 = R_2 = 0$, in which case the function is constant). The case where $\beta > 0$ is more complicated. For notational simplicity, let $x = \widetilde{B}(\beta)$, $y = \widetilde{G}(\beta)$, and $z = xy$, so

$$(20) \qquad \Omega(n;\beta) = \frac{r_1}{\beta} + \frac{R_1 x^n + R_2 z^n}{1-z^n} + \frac{r_2 - r_1}{\beta} \frac{x^n - z^n}{1-z^n},$$

and $0 < x < 1$, $0 < y < 1$, $0 < z < x$, $0 < z < y$, with x^n, y^n, and z^n strictly convex monotone decreasing functions of n.

The first term of equation (20) is constant; the second term is obviously decreasing and convexity is shown by writing it as two terms $R_1 x^n (1 + z^n + z^{2n} + \ldots) + R_2 z^n (1 + z^n + z^{2n} + \ldots)$ which is

a non-negative sum of convex terms. The third term is a constant,

$\dfrac{r_2 - r_1}{\beta}$ multiplied by the function $\dfrac{x^n - z^n}{1 - z^n} = f(n)$.

The proof of the theorem is completed by proving $f(n)$ is monotone decreasing and strictly convex, which we do by induction. Let $n=1$; the first difference is

$$\frac{x-z}{1-z} - \frac{x^2-z^2}{1-z^2} = \frac{(1-z^2)(x-z)-(1-z)(x^2-z^2)}{(1-z)(1-z^2)} = \frac{(x-z)(1-z)(1-x)}{(1-z)(1-z^2)} > 0$$

since $0 < z < x$ or $y < 1$.

Assume $f(n)\downarrow$ for $n=N$, then the first difference must be negative,

$$0 < \frac{x^N-z^N}{1-z^N} - \frac{x^{N+1}-z^{N+1}}{1-z^{N+1}} = \frac{x^N[1-y^N-z^{N+1}-x+xy^{N+1}+xz^N]}{(1-z^N)(1-z^{N+1})}$$

$$= \frac{Ax^N}{(1-z^N)(1-z^{N+1})} \Rightarrow A > 0.$$

For $n=N+1$, the first difference is

$$\frac{x^{N+1}-z^{N+1}}{1-z^{N+1}} - \frac{x^{N+2}-z^{N+2}}{1-z^{N+2}} = \frac{x^{N+1}[1-y^{N+1}-z^{N+2}-x+xy^{N+2}+xz^{N+2}]}{(1-z^{N+1})(1-z^{N+2})}$$

$$= \frac{Bx^{N+1}}{(1-z^{N+1})(1-z^{N+2})}.$$

We have

$$B - A = y^N(1-y) + z^{N+1}(1-z) + xy^{N+1}(y-1) + xz^N(z-1)$$

$$= (1-y)y(1-z) + (1-z)(z^{N+1} - xz^N)$$

$$= (1-z)y^N(1-y)(1-x^{N+1}) > 0,$$

so $A > 0 \Rightarrow B > 0 \Rightarrow \dfrac{x^{N+1} - z^{N+1}}{1 - z^{N+1}} > \dfrac{x^{N+2} - z^{N+2}}{1 - z^{N+2}}$, and by induction $f(n)$ is strictly decreasing.

Writing $f(n) = \dfrac{1 - (\frac{x}{z})^n}{1 - z^{-n}} = \dfrac{1 - y^{-n}}{1 - z^{-n}}$ and exponentiating, we see $f(n)$ is convex if $\exp(1 - y^{-n}) \exp(z^{-n} - 1) = \exp(z^{-n} - y^{-n}) = g(n)$ is convex. Considering n as a continuous variable, $n \geq 1$, we can calculate

$$\frac{d}{dn^2} g(n) = \left[(1 \ z)^2(z^{-n} + z^{-2n}) + (1 \ y)^2(y^{-2n} - y^{-n}) - 21 \ (yz)(yz)^n \right] \exp(z^{-n} - y^{-n})$$

$$> 0$$

since each term is positive. Thus, $g(n)$ is convex if n can assume all values, equal to or greater than 1, which implies $g(n)$ is convex for n a positive integer by using embedding arguments, QED.

As a result of this theorem, the optimal value of n is the largest feasible value, subject to the operating constraint. If one takes the attitude that operating constraints are imposed because of a subjective value of the holding cost h, one can compute an operationally valid value for h by calculating what value it must assume so that minimizing $\mathcal{C}(n;\beta)$ satisfies the operating constraints.

4.6 Conclusions

In most applications, inter-arrival and service times will not be so
long that it is necessary to discount costs incurred at the end of the
interval, and the undiscounted policy presented in Chapter III is the one
that should be used. The discounted model can be used in investment problems
where one is interested in the expected present worth of the future costs
of a proposed system; given the parameters of a proposed system
$(\lambda, \rho, \nu_{2B}, R_1,$ etc.$)$, the policy for using the system will be n^{opt} from
Section 3.3, and equation (10) is used to calculate the expected present
cost of operation.

Chapter V

THE FINITE HORIZON MODEL

Where there is a finite horizon, T, several new difficulties arise. First, the optimal policy is generally non-stationary, so finding the best stationary policy is not sufficient; secondly, the random variables α and β are unbounded so a transition may not be completed; lastly, it may be optimal to turn the server off while there are customers in the system, so the optimal policy for the imbedded problem may not be optimal for the full problem. The last assertion is illustrated by an example: Suppose σ is constant at thirty minutes; twenty minutes remain until the end of the horizon; and there is a service completion with $k > 0$ customers are in queue. Since no more services can be completed, the best thing to do is turn the server off and save $r_2 - r_1$ per minute for the remaining ten minutes. These considerations make it very difficult to calculate optimal policies, but for large values of T good policies can be found by using asymptotic results.

5.1 The Recursion Formula for Optimal Policies

The boundary conditions that will be imposed at the end of the horizon are:

(1) A charge of Q[$/customer] for all customers left in the system [service uncompleted or not started].

(2) If the server is running, it must be turned off.

Define the state of the system as (i,j), where j is the number of customers present and $i=1$ indicates the server is dormant and $i=2$ means it's running. In each state, the possible actions are $\alpha = 1$, turn (or leave) the

server off, and $\alpha = 2$, turn (or leave) the server on. Let $y_{ij}(\alpha, t)$ be the expected cost of a transition (which may not be completed) leaving state (i, j), using act α, and with $t(0 \leq t \leq T)$ time units remaining; $v_{ij}(t)$ is the expected total cost starting with state (i, j) with a remaining horizon, t, and following an <u>optimal policy</u>.

Letting Ψ be the number of customers that arrive during a service interval, and ξ be the number of customers in the system after a service completion, for $0 \leq t \leq T$ we have the recursion relationships:

$$(1a) \qquad y_{1j}(1, t) = \int_0^t [(jh + r_1)\alpha + v_{1, j+1}(t-\alpha)]\lambda e^{-\lambda\alpha}d\alpha$$

$$+ [(jh + r_1)t + jQ]e^{-\lambda t},$$

$$(1b) \qquad y_{2j}(1, t) = R_2 + y_{1j}(1, t),$$

$$(1c) \qquad y_{1j}(2, t) = R_1 + \int_0^t [(jh + \frac{\Psi}{2}h + r_1 + r_2)\sigma + v_{2, j+\xi}(t-\sigma)]dB(\sigma)$$

$$+ [(jh + \frac{\lambda}{2} + r_1 + r_2)t + R_2]B^c(t),$$

$$(1d) \qquad y_{2j}(2, t) = y_{1j}(2, t) - R_1,$$

and $P(\Psi = k) = \int_0^t e^{-\lambda\sigma} \frac{(\lambda\sigma)^k}{k!} dB(\sigma)$, $\xi = j - 1 + \delta(j) + \Psi$.

From the principle of optimality, the optimal policy must satisfy the usual recursion relationship of dynamic programming [10]:

$$(2) \qquad v_{ij}(t) = \underset{1,2}{\text{Min}} [y_{ij}(1, t), y_{ij}(2, t)], \quad i = 1, 2; \quad 0 \leq t \leq T.$$

Which, in principle, can be built up for successive values of t, since the y_{ij} only depend on <u>prior</u> values of $v_{ij}(t)$.

5.2 Conclusions

There is no general way to obtain solutions to equation (2); discrete approximations of the continuous variable t can be made, and digital computing methods will obtain sufficiently accurate approximations of the integrals so the resulting policy will be optimal.

In example 3, Section 3.4, we showed that the infinite horizon problem may not have a unique optimal policy. This means that as $T \to \infty$, the optimal finite horizon policy may not approach the optimal infinite horizon policy as a limit. That is, for some large values of T one of the infinite horizon policies may be best, while for some other values of T another policy may be best.

Since the transitions from state to state are governed by an ergodic (irreducible and positive recurrent) Markov chain, for large values of t

$$(3) \qquad v_{ij}(t) \approx C(n)t + k_{ij}$$

where $C(n)$ is the cost rate for an infinite horizon (Chapter III) and k is a bias term that's independent of t but dependent on the policy and state (reference 10). This implies that an optimal policy for the infinite horizon problem will be nearly optimal for large finite horizons.

This conjecture may be tested by seeing if equation (3) and the optimal policy from Chapter III satisfy equation (2).

Chapter VI

A TWO-CHANNEL MODEL

Many production facilities have spare machines that are activated when the workload reaches a critical level; these machines are run in a parallel with normally used machines until the workload level is sufficiently reduced. The problem of specifying the workload levels that trigger activation and deactivation of the spare machines will be called the spare machine problem. Because simple probabilistic results can only be obtained when the service times in each channel have an exponential distribution, this particular service time distribution will be assumed.

6.1 Assumptions

The assumptions of the two-channel spare machine problem are:

(a) The arrival stream - Customers arrive in a Poisson process at rate λ, and form a single queue.

(b) The service mechanism - There are two servers, machine one and machine two. The service times in each machine are independent, exponential random variables at rates μ_1 and μ_2 respectively, with $0 < \mu_1, \mu_2 < \infty$ and $\lambda < \mu_1 + \mu_2$. Machine one is always running and machine two may be turned on and off arbitrarily. If machine two is turned off while processing a customer, that customer rejoins the queue, or possibly, directly into machine one; otherwise, customers are served in their order of arrival by any available running machine.

(c) The cost structure - The running cost rates are r_{21}[\$/time] and r_{22}[\$/time] for machines one and two, respectively. A fixed cost

of $R_1[\$]$ is charged when machine two is turned on, and the fixed cost of shut-down is $R_2[\$]$. A holding cost of $h[\$/customer-hr.]$ is charged during the lifetime of each customer. The costs are non-negative and $h > 0$ avoids triviality; future costs are not discounted.

(d) The decision problem - When should the spare machine be turned on and off to minimize the cost rate over an infinite horizon?

6.2 Stationary Optimal Policies

We will analyze this problem by imbedding it in a dynamic programming problem at arrival and departure epochs. This causes no loss of generality with respect to making decisions at arbitrary time instants because the assumptions on the arrival and service distributions imply that the time to the next event (arrival epoch or service completion) has an exponential distribution. Since the cost rate r_{21} is always charged, it will not effect the desirability of various policies, it is, however, part of the cost rate.

Define the state space to be $S = \{0, 0', 1, 1', \ldots, k, k', \ldots\}$, where k represents the number of customers in the system; an unprimed k means the spare machine is dormant, and a primed k indicates the spare machine is running. The action space is $A=(1,2)$, where action one is "turn (or keep) the spare machine off" and action two is "turn (or keep) the spare machine on".

Using the methods of Chapter II one can prove there is a stationary optimal policy that is independent of the initial state of the system. To derive the form of this policy we first prove two lemmas.

Lemma 1: When a stationary optimal policy is used, if the spare machine is turned (or left on) when n or more customers are present, it will not be turned off when there are more than n customers in the system, $n \geq 2$.

Proof: Assume a stationary policy that turns the spare machine on in state n and turns it off when n+k $(k \geq 0)$ customers are present, and assume this policy gives a cost rate $g = C(\pi)$ which is minimal. Obviously, this cannot be true when k=0, so we only have to consider $k \geq 1$. First we observe that as the horizon approaches infinity the number of transitions out of states n and $(n+k)'$ approach infinity with probability one, so that if an improvement can be made on the expected cost of a transition leaving state $(n+k)'$, the cost rate can be improved. Using act 2 in state n implies

(1)
$$\frac{r_{22}+nh}{\lambda+\mu_1+\mu_2} + g \leq \frac{nh}{\lambda+\mu_1} + g \Rightarrow \frac{r_{22}+nh}{\lambda+\mu_1+\mu_2} \leq \frac{nh}{\lambda+\mu_1},$$

and using act 1 in state $(n+k)'$ implies

(2)
$$R_2 + \frac{(n+k)h}{\lambda+\mu_1} + g \leq \frac{(n+k)h}{\lambda+\mu_1+\mu_2} + g \Rightarrow \frac{nh}{\lambda+\mu_1} \leq \frac{r_{22}+nh}{\lambda+\mu_1+\mu_2} + kh \frac{\mu_2}{(\lambda+\mu_1+\mu_2)(\lambda+\mu_1)}$$

which contradicts equation (1) for any $k \geq 1$.

Therefore, using act 2 in state $(n+k)'$ will lower the cost rate, and the lemma is proved. QED.

Lemma 2: When a stationary optimal policy is used, if the spare machine is turned (or left) on when n customers $(n \geq 2)$ are present, it will not be turned on when two or more customers are in the system.

Proof: As a consequence of lemma 1, we only need to consider turning the

spare machine off when there are m, $2 \le m < n$, customers in the

system. Assume it's optimal to use act 1 in state $(n-k)'$,

$1 \le k \le n-2$, then

(3) $$\frac{(n-k)h}{\lambda+\mu_1} - \frac{(n-k)h}{\lambda+\mu_1+\mu_2} \le \frac{r_{22}}{\lambda+\mu_1+\mu_2}.$$

Combining equations (1) and (3) we obtain

(4) $$kh\left[\frac{\mu_2}{(\lambda+\mu_1+\mu_2)(\lambda+\mu_1)}\right] \le 0$$

wich can't be satisfied for $k \ge 1$, so act one cannot be optimal in

state $m' = (n-k)'$, QED.

Theorem 1: A stationary optimal policy either (a) keeps the spare machine

on at all times, (b) never turns the spare machine on, or (c)

has the form: Turn the spare machine on when n customers are

in the system, and off when $m \le 1$ customers are in the system.

Proof: Lemmas 1 and 2 imply that $m \le 1$ when $n \ge 2$. When service interruptions

are prohibited, turning the spare machine on when n=1 cannot be

optimal because the running cost rate r_{22} is incurred without re-

ducing the expected holding costs; the only remaining policies are

those where n=0. Since turning the spare machine on when n=0, off

when $m'=1$, and on when n=2 is clearly not optimal, all possible

stationary policies have the required ferm, QED.

As a consequence of this theorem, when service interruptions are pro-

hibited there are only four forms the optimal stationary policy can take,

viz:

π_1 - Never turn the spare machine on.

π_2 - Leave the spare machine on all the time.

π_3 - Turn the spare machine on when $n \geq 2$ customers are present, and turn it off when the system becomes empty.

π_4 - Turn the spare machine on when $n \geq 2$ customers are present, and turn it off when one customer is left in the system and he's being served by machine one.

When interrupting service and switching the customer to the other server is permitted, three more policies, call switching policies, may be optimal. They are:

π_5 - Turn the spare machine on when $n \geq 2$ customers are present, and turn it off when it is serving the only customer in the system, [i.e., restart him in machine one].

π_6 - Turn the spare machine on when $n \geq 2$; if machine one is serving the only customer in the system and machine two is running, switch this customer to machine two. Turn the spare machine off when the system is empty.

π_7 - Turn the spare machine on when $n = 1$ and switch the customer in machine one to machine two. Turn the spare machine off when the system is empty.

Policy π_7 is a switching policy because we assume that a customer who arrives when machine one is idle and machine two is dormant, will enter machine one.

6.3 Calculation of the Cost-Rates when Switching is Prohibited

When policy π_1 is used, the system behaves like an M/M/1 queue with service rate μ_1, and the cost rate, obtained from equation (III-4), is

$$(5) \qquad C(\pi_1) = r_{21} + \frac{\lambda h}{\mu_1 - \lambda}, \quad \mu_1 > \lambda.$$

Since this policy yields an infinite cost rate when $\mu_1 \leq \lambda$, π_1 cannot be optimal unless $\mu_1 < \lambda$.

When policy π_2 is used, the system performs like an M/M/2 queue with service rate μ_1 when one customer is in the system. Using well known formulae (see reference 4, Section 2.4), the expected number of customers present in this type of queue is given by

$$L = \frac{\rho_1}{(1-\rho)(1+\rho_1-\rho_1)}, \quad \rho = \frac{\lambda}{\mu_1+\mu_2}, \quad \rho_1 = \frac{\lambda}{\mu_1};$$

the cost rate is

$$(6) \qquad C(\pi_2) = r_{21} + r_{22} + \frac{\rho_1 h}{(1-\rho)(1+\rho_1-\rho)}.$$

The remaining two cost rates will be calculated using the policy evaluation routine of Markov-renewal programming [reference 10]. When either of the policies π_3 or π_4 are followed, a Markov-renewal process with additive costs during each transition epoches of the queueing process so that the cost rates of both processes are equal. Furthermore, theorem 1 implies this can be done in a manner such that the underlying Markov chain is ergodic (irreducible and positive recurrent) and the expected time for each transition is finite. For an infinite-time process operating under a

stationary policy, the system of equations

$$(7) \qquad v_i + v_i = y_i + \sum_{i=1}^{N} p_{ij} v_i, \qquad i=1, \ldots, N$$

$$v_0 \equiv 0,$$

where v_i is the relative value of state i, is the cost rate [previously called $C(\pi_i)$], v_i is the expected length of a transition leaving state i, y_i is the expected one-step cost of a transition from state i, and p_{ij} are the conditional transition probabilities of the underlying Markov chain, can be solved for g and v_i uniquely.

Policy π_3

When policy π_3 is used, the imbedding points are arrival and service epochs when only machine one is running, and turning machine two on is represented as a transition from state n to state 0. The corresponding Markov chain and conditional probabilities are

$$(8) \qquad P_{01}=1, \ P_{n0}=1, \ P_{ij} = \left\{ \begin{array}{ll} \dfrac{\mu_1}{\lambda+\mu_1} & j=i-1 \\ \dfrac{\lambda}{\lambda+\mu_1} & j=i+1 \end{array} \right\} \ i=2, \ldots, n-1,$$

and one easily finds that

$$(9) \qquad v_0 = \frac{1}{\lambda}, \ y_0 = \frac{r_{21}}{\lambda}$$

$$v_i = \frac{1}{\lambda+\mu_1}, \ y_i = \frac{r_{21}+ih}{\lambda+\mu_1} \qquad i=1, \ldots, n-1.$$

A transition from state n to state zero can be thought of as a transition from n to one followed by a transition from one to zero. The first part, $(n \to 1)$, behaves like a busy period of an $M_{n-1}/G/1$ queue with service

rate $\mu_1+\mu_2$; from the results of Chapter III we have that the expected

length of this interval $E(n \to 1)$, and the expected number of customers

present during this time $L(n \to 1)$, are given by

$$(10) \qquad E(n \to 1) = \frac{n-1}{\mu_1+\mu_2+\lambda}, \quad L(n \to 1) = \frac{\mu_1+\mu_2}{\mu_1+\mu_2-\lambda} + \frac{n}{2}.$$

The second part of the transition, $(1 \to 0)$, behaves like the busy

period of an M/M/2 queue where the service rate when one customer is in the

system is μ_1 with probability $\frac{\mu_1}{\mu_1+\mu_2}$, and is μ_2 with probability $\frac{\mu_1}{\mu_1+\mu_2}$. Let

E_1 and E_2 be the expected lengths of busy periods when their service rates

are μ_1 and μ_2 during the <u>first</u> time there is only one customer in service.

Conditioning on the first transition leaving state 1', we find that E_1 and

E_2 satisfy

$$(11) \qquad E_1 = \frac{1}{(\lambda+\mu_1)} + \frac{\lambda}{\lambda+\mu_1}\left[E(2 \to 1) + \frac{\mu_1}{\mu_1+\mu_2} E_2 + \frac{\mu_2}{\mu_1+\mu_2} E_1\right],$$

$$(12) \qquad E_2 = \frac{1}{(\lambda+\mu_2)} + \frac{\lambda}{\lambda+\mu_2}\left[E(2 \to 1) + \frac{\mu_1}{\mu_1+\mu_2} E_2 + \frac{\mu_2}{\mu_1+\mu_2} E_1\right];$$

substituting equation (12) into equation (11) we obtain

$$(13) \qquad E_1 = \frac{(\mu_1+\mu_2)^2[(\lambda+\mu_2) + \mu_2(\lambda+\mu_1+\mu_2)]}{(\lambda+\mu_1)(\lambda+\mu_1+\mu_2)[\mu_1(\mu_1+\mu_2)(\lambda+\mu_2) - \lambda(\mu_2)^2]}$$

$$+ \frac{\lambda(\mu_1+\mu_2)[(\lambda+\mu_2)(\mu_1+\mu_2) + \mu_2(\mu_1+\mu_2+\lambda)]}{(\lambda+\mu_1+\mu_2)[\mu_1(\mu_1+\mu_2)(\lambda+\mu_2) - \lambda(\mu_2)^2]} E(2 \to 1).$$

E_2 is given by a similar equation with the subscripts interchanged. Thus,

the expected time to go from state one to state zero is

$$(14) \qquad E(1 \to 0) = \frac{\mu_2}{\mu_1 + \mu_2} E_1 + \frac{\mu_1}{\mu_1 + \mu_2} E_2.$$

Letting L_j equal the time-average of the number of customers in the system when μ_j is the service rate the first time only one customer is in service, we find that L_1 and L_2 can be calculated by solving

$$(15) \qquad L_j = \frac{1 + \lambda\left[\frac{\mu_1 + \mu_2}{(\mu_1 + \mu_2 - \lambda)^2} + \frac{\mu_1}{\mu_1 + \mu_2} E_2 L_2 + \frac{\mu_2}{\mu_1 + \mu_2} E_1 L_1\right]}{(\lambda + \mu_j) E_j}, \quad j=1, 2$$

simultaneously, and the average number of customers in the system during a transition from state one to state zero is

$$(16) \qquad L(1 \to 0) = \frac{\frac{\mu_2}{\mu_1 + \mu_2} L_1 E_1 + \frac{\mu_1}{\mu_1 \mu_2} L_2 E_2}{E(1 \to 0)}.$$

From equations (10), (14), and (16) we obtain

$$(17) \qquad v_n = \frac{n-1}{\mu_1 + \mu_2 - \lambda} + E(1 \to 0),$$

$$y_n = R_1 + R_2 + (r_{21} + r_{22})v_n + h\left[\frac{1}{\mu_1 + \mu_2 - \lambda} + \frac{n}{2} + L(1 \to 0)\right]v_n.$$

The special form of the p_{ij} given in equations (8) allows us to write the first n-1 equations of the system (6) as the linear second order difference equation

$$(18) \qquad v_{i+2} - \frac{\lambda + \mu_1}{\lambda} v_{i+1} + \frac{\mu_1}{\lambda} v_i = \frac{g^{-r_{21} - ih}}{\lambda}, \quad i=0, \ldots, n-2,$$

with boundary conditions

$$(19) \qquad\qquad v_0 = 0, \quad v_1 = \frac{g^{-r}21}{\lambda}.$$

The solution to equation (18) is

$$(20) \quad v_i = d_1 + d_2 \left(\frac{2\mu_1}{\lambda}\right)^i + \left[\frac{r_{21}-g}{\mu_1-\lambda} + \frac{h(3\lambda-\mu_1)}{2(\mu_1-\lambda)^2}\right]i + \frac{hi^2}{2(\mu_1-\lambda)} \qquad i=2, \ldots, n$$

where the coefficients d_1 and d_2 are obtained from the boundary conditions (19).

Since we can find v_n in terms of g from equation (20), and $v_n = y_n - v_n$ from equation (7), the value of $g = C(\pi)$ can be obtained as a function of n and the parameters of the model. One can then find the optimal n <u>for this policy</u>.

<u>Policy π_4</u>

The values of v_i represent the relative value of being in state i when machine two is dormant; for policies π_4 and π_5 we are interested in the relative value of being in state i when machine two is running, denoted w_i, and particulary in w_1. Observe that $\dot{w}_i = R_2 + v_i$.

Let w_1' be the relative value of state 1 when machine two is running and busy, and w_1'' be the relative value when machine two is running but idle. Under policy π_4, the v_i satisfy equation (18) and the w_i are given by

$$(21a) \qquad\qquad w_0 = R_2$$

$$(21b) \qquad w_1' + \frac{g}{\lambda+\mu_2} = \frac{r_{22}+r_{11}+h}{\lambda+\mu_2} + \frac{\mu_2}{\lambda+\mu_2} w_0 + \frac{\lambda}{\lambda+\mu_2} w_2$$

$$(21c) \qquad\qquad w_1'' = R_2 + v_1$$

(21d) $\quad w_2 + \dfrac{g}{\lambda+\mu_1+\mu_2} = \dfrac{r_{22}+r_{21}+2h}{\lambda+\mu_1+\mu_2} + \dfrac{\mu_1+\mu_2}{\lambda+\mu_1+\mu_2}\left(\dfrac{\mu_1}{\mu_1+\mu_2} w_1'' + \dfrac{\mu_2}{\mu_1+\mu_2} w_1'\right)$

$$+ \dfrac{\lambda}{\lambda+\mu_1+\mu_2} w_3$$

(21d) $\quad w_i + \dfrac{g}{\lambda+\mu_1+\mu_2} = \dfrac{r_{22}+r_{21}+ih}{\lambda+\mu_1+\mu_2} + \dfrac{\mu_1+\mu_2}{\lambda+\mu_1+\mu_2} w_{i-1} + \dfrac{\lambda}{\lambda+\mu_1+\mu_2} w_{i+1}$$

$$i=3, 4, \ldots n-1.$$

Equation (21d) is a linear second order difference equation whose solution is

(22) $\quad w_i = q_1 + q_2\left(\dfrac{\mu_1+\mu_2}{\lambda}\right)^i + \left[\dfrac{r_{21}+r_{22}-g}{\mu_1+\mu_2-\lambda} + \dfrac{(3\lambda-\mu_1-\mu_2)h}{2(\mu_1+\mu_2-\lambda)^2}\right] i + \dfrac{hi^2}{2(\mu_1+\mu_2-\lambda)}$

$$i=4, 5, \ldots, n,$$

and the constants q_1 and q_2 are obtained by solving for w_2 and w_3 in terms of w_1' using equations (21a-d), and then using the relationship $w_n = v_n + R_2$ to calculate w_1'.

Under policy π_4, transitions leaving state n' enter state $1'$ in the imbedded Markov-renewal process; using equation (8) we obtain

(23) $\quad w_n + \dfrac{(n-1)}{\mu_1+\mu_2-\lambda} = \dfrac{n-1}{\mu_1+\mu_2-\lambda}\left[r_{21}+r_{22}+\left(\dfrac{1}{\mu_1+\mu_2-\lambda} + \dfrac{n}{2}\right)h\right] + \dfrac{\mu_1}{\mu_1+\mu_2} w_1'' + \dfrac{\mu_2}{\mu_1+\mu_2} w_1'.$

Using the values of w_n obtained from equations (22) and (23), an equation for the cost rate g can be obtained.

In summary, the cost rate and optimal policy of the form π_4 is found by:

(1)   Solving for the v_i ($i=0, \ldots, n$) from equations (19) and (20),

(2)   Solving for the w_i ($i=0, \ldots, n$) from equations (21) and (22).

(3) Equating the values of w_n obtained from equations (22) and (23), solve for the cost rate in terms of the decision variable n.

(4) Find the value of n $(n \geq 2)$ than minimizes the cost rate.

6.4 Calculation of the Cost-Rates for the Switching Policies

Policy π_5 is evaluated with the same equations as policy π_4 and the additional relationship $w_1' = R_2 + v_1$, which implies $w_1' = w_1'' = w_1$. This simplifies equation (21) to

(24a)
$$w_0 = R_2$$

(24b)
$$w_1 = R_2 + v_1$$

(24c) $$w_i + \frac{g}{\lambda+\mu_1+\mu_2} = \frac{r_{22}+r_{21}+ih}{\lambda+\mu_1+\mu_2} + \frac{\mu_1+\mu_2}{\lambda+\mu_1+\mu_2} w_{i-1} + \frac{\lambda}{\lambda+\mu_1+\mu_2} w_{i+1}$$

$$i=2, \ldots, n-1.$$

The solution to equation (24c) is given by equation (22), and the constants q_1 and q_2 are evaluated from the boundary conditions (24a) and (24b).

Equations (23) simplifies to

(25) $$w_n + \frac{(n-1)}{\mu_1+\mu_2-\lambda} = \frac{n-1}{\mu_1+\mu_2-\lambda}\left[r_{22}+r_{21}+(\frac{\mu_1+\mu_2}{\mu_1+\mu_2-\lambda} + \frac{n}{2})h\right] + w_1,$$

and an expression for g can be obtained using the values of w_n given by equations (24c) and (25).

Policy π_6 is solved in an analagous manner; the boundary condition

(24b) is replaced by

$$(24d) \qquad \omega_1 = v_1 = \frac{h + r_{21} + r_{22} - g}{\lambda + \mu_2} + \frac{\lambda}{\lambda + \mu_2} v_2,$$

where v_2 is given by equation (20).

When policy π_7 is used, the system behaves like an M/M/2 queue where the service rate is μ_2 when one customer is in the system. For this type of queue, the probability that the system is empty P_0, and the expected number of customers present L, are given by [reference 4, Section 2.4].

$$(26) \qquad P_0 = \frac{1-\rho}{1+\rho_2+\rho}, \quad L = \frac{\rho_2}{(1-\rho)(1+\rho_2-\rho)},$$

where $\rho_2 = \frac{\lambda}{\mu_1}$ and $\rho = \frac{\lambda}{\mu_1 + \mu_2}$. Using the relationship

$$P_0 = \frac{E(\zeta)}{E(X)} = \frac{1}{\lambda E(X)} = \frac{1}{1 + \lambda E(\gamma)}$$

where $E(X)$ and $E(\gamma)$ are the expected lengths of a busy cycle and a busy period respectively, one obtains

$$(27) \qquad E(X) = \frac{1 + \rho_2 - \rho}{\lambda(1-\rho)}, \quad E(\gamma) = \frac{1}{\mu_2(1-\rho)}.$$

From equations (26) and (27), we obtain the cost rate of policy π_7,

$$(28) \qquad C(\pi_7) = r_{21} + \frac{h\rho_2}{(1-\rho)(1-\rho_2+\rho)} + \frac{r_{22}\rho_2 + (R_1 + R_2)\lambda(1-\rho)}{1+\rho_2-\rho}.$$

6.5 Numerical Examples

Example 1: $\lambda=1$, $\mu_1=\mu_2=1$, $r_{22}=5$, $r_{21}=2$, $h=10$, $R_1=R_2=0$

When switching is prohibited, a policy of the form π_4 which turns the space machine on when two customers are present is optimal, and the cost rate is $14\frac{11}{12}$ [$/hr.]. If switching were allowed, switching to the only customer from machine two to machine one, policy π_5, would lower the cost rate to $13\frac{2}{3}$ [$/hr.].

Example 2: $\lambda=1$, $\mu_1=1$, $\mu_2=2$, $r_{22}=5$, $r_{21}=2$, $h=10$, $R_1=0$, $R_2=5$

When switching is prohibited, a policy of the form π_3 is optimal; turning the spare machine on when two customers are present and turning it off when the system becomes empty gives a cost rate of 19 [$/hr.]. When switching is allowed, a policy of the form π_6 is optimal, and the cost rate is reduced to 12.6 [$/hr.].

6.6 Conclusions

Although closed form expressions are not obtained for the cost rates of the seven policies considered in this chapter, we are able to give qualitative relationships among the cost parameters that will indicate the optimal policy.

Policy π_1 will be used when h is small, $\lambda < \mu_1$, and the operating costs of the spare machine are large; while π_2 will be more advantageous when the fixed costs of the spare machine and the holding costs are large, and the running cost is small. Policy π_3 will be used to hedge against high waiting lines when all the costs are of the same relative importance. Policies π_4 and π_5 protect against high holding costs but incur the fixed costs

frequently; π_5 will preferred to π_4 when the running cost of the spare machine high.

Policies π_6 and π_7 tend to lower the cost rate when the spare machine is faster than the regular machine and its running cost rate is small. Policy π_7 will be preferred to policy π_6 when the fixed costs of the spare machine are large.

One should not be mislead about the difficulty of finding the cost rates for policies π_3, π_4, π_5, and π_6; numerical solutions will be easy to obtain since all the equations are linear in the cost rate. The major difficulty will be to calculate the optimal value of n since calculus methods lead to implicit equations for n^*.

SUMMARY

The aim of this thesis is to describe the economic behavior of a controllable system with a linear cost structure, and to find cost-minimizing policies for turning the server on and off. The costs considered are: a server start-up cost, a server shut-down cost, a cost per unit time when the server is turned off, a cost per unit time when the server is turned on, and a holding cost of waiting customers. Single-channel queues with Poisson arrivals and arbitrary service time distributions are emphasized; Chapter Six is devoted to a two-channel system with Poisson arrivals and exponentially-distributed service times.

In Chapter Two the decision process for a single-channel queue operating for an infinite horizon is formulated as a dynamic program. We show that when future costs are discounted, there exists a stationary optimal policy to minimize the expected total cost; we use this result to prove that when discounting is not used, there is a stationary optimal policy that minimizes the cost rate. We then prove that for both models, the stationary optimal policy has the form: Turn the server on when n customers are present, and turn it off when the system is empty.

Chapter Three deals with infinite-horizon, undiscounted models. For models where the arrival stream is a Poisson process, two methods for deriving the cost rate as a function of the decision variable n are presented, and the optimal value of n and the minimum cost rate are obtained. We consider decreasing the expected service time with an increase in the running cost of the server; an expression showing when the server should be speeded up, or slowed down, is given. When the inter-arrival time distribution is generalized to the class of IFR distributions, we obtain narrow

bounds for the optimal expected cost rate.

An equation for the expected discounted cost over an infinite horizon is derived in Chapter Four; we prove that for small interest rates, the undiscounted policy will be a good approximation to the optimal policy.

When the horizon is finite, the optimal policy is generally non-stationary. A recursion relationship to find the optimal policy when the horizon is small is presented, and the optimal undiscounted policy for the infinite horizon is shown to be a good approximation to the optimal policy for large horizons.

The last chapter is devoted to a two-channel service system, where each channel is restricted to have an exponentially distributed service time (possibly with different rates), and the arrivals form a Poisson process. One server is always turned on; the other, the spare machine, can be turned on and off at arbitrary times. Using a dynamic programming formulation of the decision process, we show that the stationary optimal policy for undiscounted costs has the form: Turn the spare machine on when n customers are present, and turn it off when m customers are in the system, with $m \leq 1$. We derive equations for finding the optimal value of m and n when service interruptions are prohibited; then we consider queue disciplines where customers may be switched, without delay or cost, from one server to the other.

In addition to being applicable to policy problems for existing systems, these models should be useful when comparing proposed investments in service system because they relate the parameters of the server to the cost of the facility. The most promising areas for future research appear to be: systems where arriving customers may not enter the queue if the waiting line is too large, systems where customers leave the waiting line if they have not been served after a given wait in queue, and processes with different types of customers.

REFERENCES

1. Barlow, R. E., and Frank Proschan, <u>Mathematical Theory of Reliability</u>, Wiley, New York, 1965.

2. Blackwell, David, "Discounted Dynamic Programming", <u>Annals of Mathematical Statistics 36</u>, 1965, pp. 226-235.

3. Cox, D. R., <u>Renewal Theory</u>, Wiley, New York, 1961.

4. Cox, D. R., and Walter Smith, <u>Queues</u>, Wiley, New York, 1961.

5. Denardo, E. V., "Contraction Mappings in the Theory Underlying Dynamic Programming", Memorandum RM-4755-PR, The RAND Corporation, Santa Monica, 1966.

6. Erlander, Sven, "The Remaining Busy Period for a Single Server Queue", <u>Operations Research 13</u>, 1965, pp. 734-746.

7. Feller, William, <u>An Introduction to Probability Theory and its Applications</u>, Vol. 1, 2nd ed., Wiley, New York, 1957.

8. Fox, "Markov Renewal Programming by Linear Fractional Programming", P-3257-1, The RAND Corporation, Santa Monica, 1966.

9. Gaver, D. P., Jr., "Imbedded Markov Chain Analysis of a Waiting-Line Problem in Continuous Time", <u>Annals of Mathematical Statistics 30</u>, 1959, pp. 698-720.

10. Jewell, William S., "Markov Renewal Programming I and II", <u>Operations Research 11</u>, 1963, pp. 938-971.

11. Little, J. D. C., "A Proof of the Queueing Formula: $L=\lambda W$", <u>Operations Research 9</u>, 1961, pp. 383-387.

12. Marshall, Kneale, <u>Some Inequalities for Single Server Queues</u>, Ph.D. Dissertation, College of Engineering, University of California, Berkeley, 1966.

13. Morse, P. M., <u>Queues, Inventories and Maintenance</u>, Wiley, New York, 1957.

14. Takacs, L., "The Transient Behavior of a Single-Server Queueing Process with a Poisson Input", <u>Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability 2</u>, 1961, pp. 535-567.

15. Van der Pol, Balth, and H. Bremmer, <u>Operational Calculus</u>, University Press, Cambridge, England, 1964.

16. Wolff, Ronald W., <u>Notes on Queueing Theory</u>, Department of Industrial Engineering and Operations Research, University of California, Berkeley, 1965.

DOCUMENT CONTROL DATA - R&D
(Security classification of title, body of abstract and indexing annotation must be entered when the overall report is classified)

1. ORIGINATING ACTIVITY (Corporate author)	2a. REPORT SECURITY CLASSIFICATION
University of California, Berkeley	Unclassified
	2b. GROUP

3. REPORT TITLE

OPTIMAL OPERATING POLICIES FOR STOCHASTIC SERVICE SYSTEMS

4. DESCRIPTIVE NOTES (Type of report and inclusive dates)

Research Report

5. AUTHOR(S) (Last name, first name, initial)

Heyman, Daniel, P.

6. REPORT DATE	7a. TOTAL NO. OF PAGES	7b. NO. OF REFS
October 1966	72	16

8a. CONTRACT OR GRANT NO.	9a. ORIGINATOR'S REPORT NUMBER(S)
Nonr-222(83)	
b. PROJECT NO.	ORC 66-31
NR 047 033	
c.	9b. OTHER REPORT NO(S) (Any other numbers that may be assigned this report)
Research Project No. RR 003-07-01	
d.	

10. AVAILABILITY/LIMITATION NOTICES

Available without limitation on dissemination

11. SUPPLEMENTARY NOTES	12. SPONSORING MILITARY ACTIVITY
	Mathematical Science Division

13. ABSTRACT

See 2nd page

DD FORM 1 JAN 64 **1473**

14. KEY WORDS	LINK A		LINK B		LINK C	
	ROLE	WT	ROLE	WT	ROLE	WT
queueing system						
economic behavior						
linear cost structure						

INSTRUCTIONS

1. **ORIGINATING ACTIVITY:** Enter the name and address of the contractor, subcontractor, grantee, Department of Defense activity or other organization (*corporate author*) issuing the report.

2a. **REPORT SECURITY CLASSIFICATION:** Enter the overall security classification of the report. Indicate whether "Restricted Data" is included. Marking is to be in accordance with appropriate security regulations.

2b. **GROUP:** Automatic downgrading is specified in DoD Directive 5200.10 and Armed Forces Industrial Manual. Enter the group number. Also, when applicable, show that optional markings have been used for Group 3 and Group 4 as authorized.

3. **REPORT TITLE:** Enter the complete report title in all capital letters. Titles in all cases should be unclassified. If a meaningful title cannot be selected without classification, show title classification in all capitals in parenthesis immediately following the title.

4. **DESCRIPTIVE NOTES:** If appropriate, enter the type of report, e.g., interim, progress, summary, annual, or final. Give the inclusive dates when a specific reporting period is covered.

5. **AUTHOR(S):** Enter the name(s) of author(s) as shown on or in the report. Enter last name, first name, middle initial. If military, show rank and branch of service. The name of the principal author is an absolute minimum requirement.

6. **REPORT DATE:** Enter the date of the report as day, month, year; or month, year. If more than one date appears on the report, use date of publication.

7a. **TOTAL NUMBER OF PAGES:** The total page count should follow normal pagination procedures, i.e., enter the number of pages containing information.

7b. **NUMBER OF REFERENCES:** Enter the total number of references cited in the report.

8a. **CONTRACT OR GRANT NUMBER:** If appropriate, enter the applicable number of the contract or grant under which the report was written.

8b, 8c, & 8d. **PROJECT NUMBER:** Enter the appropriate military department identification, such as project number, subproject number, system numbers, task number, etc.

9a. **ORIGINATOR'S REPORT NUMBER(S):** Enter the official report number by which the document will be identified and controlled by the originating activity. This number must be unique to this report.

9b. **OTHER REPORT NUMBER(S):** If the report has been assigned any other report numbers (*either by the originator or by the sponsor*), also enter this number(s).

10. **AVAILABILITY/LIMITATION NOTICES:** Enter any limitations on further dissemination of the report, other than those imposed by security classification, using standard statements such as:

   (1) "Qualified requesters may obtain copies of this report from DDC."

   (2) "Foreign announcement and dissemination of this report by DDC is not authorized."

   (3) "U. S. Government agencies may obtain copies of this report directly from DDC. Other qualified DDC users shall request through

                                               ."

   (4) "U. S. military agencies may obtain copies of this report directly from DDC. Other qualified users shall request through

                                               ."

   (5) "All distribution of this report is controlled. Qualified DDC users shall request through

                                             ."

If the report has been furnished to the Office of Technical Services, Department of Commerce, for sale to the public, indicate this fact and enter the price, if known.

11. **SUPPLEMENTARY NOTES:** Use for additional explanatory notes.

12. **SPONSORING MILITARY ACTIVITY:** Enter the name of the departmental project office or laboratory sponsoring (*paying for*) the research and development. Include address.

13. **ABSTRACT:** Enter an abstract giving a brief and factual summary of the document indicative of the report, even though it may also appear elsewhere in the body of the technical report. If additional space is required, a continuation sheet shall be attached.

It is highly desirable that the abstract of classified reports be unclassified. Each paragraph of the abstract shall end with an indication of the military security classification of the information in the paragraph, represented as (TS), (S), (C), or (U).

There is no limitation on the length of the abstract. However, the suggested length is from 150 to 225 words.

14. **KEY WORDS:** Key words are technically meaningful terms or short phrases that characterize a report and may be used as index entries for cataloging the report. Key words must be selected so that no security classification is required. Identifiers, such as equipment model designation, trade name, military project code name, geographic location, may be used as key words but will be followed by an indication of technical context. The assignment of links, roles, and weights is optional.

DD FORM 1 JAN 64 **1473** (BACK)