



**TECHNICAL
RESEARCH
REPORT**

SRC TR 88-95

**Interchange Arguments In
Stochastic Scheduling**

by

P. Nain, P. Tsoucas, J. Walrand

SYSTEMS RESEARCH CENTER

UNIVERSITY OF MARYLAND

COLLEGE PARK, MARYLAND 20742

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE 28 NOV 1988		2. REPORT TYPE		3. DATES COVERED 00-00-1988 to 00-00-1988	
4. TITLE AND SUBTITLE Interchange arguments in stochastic scheduling				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) University of Maryland, Systems Research Center, College Park, MD, 20742				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT see report					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES 16	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

Interchange arguments in stochastic scheduling

*Philippe Nain*¹, *Pantelis Tsoucas*²
and
Jean Walrand

INRIA- Sophia Antipolis
Av. E. Hugues
06565 Valbonne Cedex
France

Systems Research Center
University of Maryland
College Park, MD 20742
USA

and

Department of Electrical Engineering and Computer Sciences
and Electronics Research Laboratory
University of California
Berkeley, CA 94720
USA

ABSTRACT

Interchange arguments are applied to establish the optimality of priority list policies in three problems. First, we prove that in a multi-class tandem of two $M/1$ queues it is always optimal in the second node to serve according to the "c μ " rule. The result holds more generally if the first node is replaced by a multi-class network consisting of $M/1$ queues with Bernoulli routing. Next, for scheduling a single server in a multi-class node with feedback, a simplified proof of Klimov's result is given. From it follows the optimality of the index rule among idling policies for general service time distributions, and among preemptive policies when the service time distributions are exponential. Lastly, we consider the problem of minimizing the blocking in a communication link with lossy channels and exponential holding times.

November 28, 1988

¹ Work performed while this author was visiting the Systems Research Center, University of Maryland, College Park, supported by the Office of Naval Research through grant number N00014-84-K-0614.

² The work of this author was supported in part by the NSF Engineering Research Centers Program #NSFD-CDR-88-03012.

Interchange arguments in stochastic scheduling

*Philippe Nain*¹, *Pantelis Tsoucas*²
and
Jean Walrand

INRIA- Sophia Antipolis
Av. E. Hugues
06565 Valbonne Cedex
France

Systems Research Center
University of Maryland
College Park, MD 20742
USA

and

Department of Electrical Engineering and Computer Sciences
and Electronics Research Laboratory
University of California
Berkeley, CA 94720
USA

1. Introduction

This paper has two main aims. The first is to demonstrate the use of interchange arguments in proving optimality properties and the second is to obtain new results in stochastic scheduling. The main idea of our arguments appears in Varaiya et al [14] where it is used in order to obtain the optimality of index rules in multi-armed bandit problems. There, the objective is to maximize the expected total discounted reward. We use variations of this idea together with path-wise coupling techniques.

We first apply an interchange argument in Section 2 to partially characterize the optimal policy for scheduling two servers in a tandem of two nodes with M different classes of customers with exponential service times. The result, motivated by Ross and Yao [11], is that the optimal policy in the second node is a "c μ " rule. This is an easy extension of the results of Baras et al [2] and Buyukkoc et al [3]. The result can be extended to the case where the first node is a network consisting of $M/1$ queues with Bernoulli routing.

Next, in Section 3 the problem of Klimov [7] is considered. A single server is to be scheduled in a network of $M/GI/1$ nodes. The objective is to minimize the expected long term average cost. It has been shown in [14] that this problem is equivalent to a multi-armed bandit problem. Our argument provides a simple proof of the result in [7], that the nonpre-emptive

¹ Work performed while this author was visiting the Systems Research Center, University of Maryland, College Park, supported by the Office of Naval Research through grant number N00014-84-K-0614.

² The work of this author was supported in part by the NSF Engineering Research Centers Program #NSFD-CDR-88-03012.

nonidling optimal policy is a priority rule. We also establish the optimality of that rule among idling policies. The priorities are determined and for the case where the service distributions are exponential we show that the same priority rule is optimal among pre-emptive policies. Remarkably, the optimal policy does not depend on the arrival rate. Our proof provides some insight into this fact.

Finally, in Section 4 we consider a problem of stochastic scheduling that does not fall in the framework of multi-armed bandit problems. Calls arrive at a communication link where N channels are available. There are probabilities of immediate loss associated with each channel and a successful call occupies the channel for an exponential amount of time. If the holding times are all independent and identically distributed, Anantharam et al [1] show that the time to reach the state where all channels are full is independent of the placement policy used. We provide a simple proof of this result and further prove that in the case where the holding times are not identically distributed, the time to reach the full state is stochastically maximized by assigning calls to the free channel with the shortest holding time.

2. Server scheduling in a multi-class network

Consider two $M/1$ queues in tandem with M classes of jobs. Calls arrive at the first node at deterministic time instants $\{a_k\}_k$. The service rates at the first node are $\{\mu_i\}_{i=1}^M$ and there are associated holding costs denoted by $\{c_i\}_{i=1}^M$. In the second node the corresponding service rates are $\{\nu_i\}_{i=1}^M$ and the holding costs are $\{d_i\}_{i=1}^M$. Let $\mathbf{x}_t = (x_t^1, \dots, x_t^M)$ (respectively $\mathbf{y}_t = (y_t^1, \dots, y_t^M)$) be the vector of class populations in node 1 (respectively node 2). Assume that $d_1\nu_1 \geq d_2\nu_2 \geq \dots \geq d_M\nu_M$. A pre-emptive server allocation policy is a function

$$\pi : (\mathbf{x}_t, \mathbf{y}_t) \rightarrow (\pi_t^1(\mathbf{x}_t, \mathbf{y}_t), \pi_t^2(\mathbf{x}_t, \mathbf{y}_t)) \in \{1, 2, \dots, M\}^2$$

The objective is to minimize over π the expected discounted cost incurred in the interval $[0, T]$ given by

$$J(\pi, T) = E \left[\int_0^T e^{-\beta t} \left(\sum_{i=1}^M c_i x_t^i + \sum_{i=1}^M d_i y_t^i \right) dt \right] \quad (2.1)$$

The following result shows that a " $d\nu$ " policy is optimal for the second node. It is an extension of results in Baras et al [2] and Buyukkoc et al [3] who consider a single node. It also provides an extension to results in Ross and Yao [11] who consider multi-server scheduling in a network.

Theorem 2.1: In node 2 the optimal policy always serves job i , among the ones present in the queue, for which the quantity $d_i\nu_i$ is maximum.

Proof: The virtual service process of an exponential server with rate μ is a Poisson point process with parameter μ . A point of this process is a service completion if the queue is non-empty. Let $\{t_n^i\}$ (respectively $\{s_n^i\}$) be the points of the virtual service process for class $i \in \{1, \dots, M\}$ in node 1 (respectively in node 2). We only need consider policies π switching at times

$$\{t_n\} = \{a_n\} \bigcup_{i=1}^M \{t_n^i\} \bigcup_{i=1}^M \{s_n^i\}.$$

For $T \geq 0$, condition on the number of points of the process $\{t_n\}$ in the interval $[0, T]$, say $0 \leq t_1 \leq t_2 \leq \dots \leq t_k \leq T \leq t_{k+1} \leq \dots$. Optimality will be proved by induction on k . The result is trivially true for $k=0$ and assume that it holds for $k=1, \dots, K$. We will prove that the result remains true for $k=K+1$.

By the optimality principle and the induction hypothesis $\pi_2(\cdot, \cdot)$ has to follow the " $d\nu$ " rule at times $t_1, t_2, \dots, t_K$ and suppose that $\pi_0^2 = i$ while $y_0^j > 0$ with $i > j$. Then policy π cannot be optimal because it can be improved as follows. Denote by t_σ the first time when $\pi_{t_\sigma}^2 = j$ and define policy $\bar{\pi}$ as follows.

$$\begin{aligned}\bar{\pi}_t^1 &= \pi_t^1, \quad t \geq 0, \\ \bar{\pi}_0^2 &= j, \\ \bar{\pi}_t^2 &= \pi_t^2, \quad t = t_1, \dots, t_{\sigma-1}, \\ \bar{\pi}_{t_\sigma}^2 &= i, \\ \bar{\pi}_t^2 &= \pi_t^2, \quad t = t_{\sigma+1}, t_{\sigma+2}, \dots\end{aligned}$$

Then simple algebraic manipulation shows that

$$J(\pi, K+1) - J(\bar{\pi}, K+1) \geq 0,$$

if

$$p(j, t_k) d_j - p(i, t_k) d_i \geq 0, \quad k=1, \sigma,$$

where

$$p(l, t_k) = \Pr \{t_k \in \{s_n^l\}\}, \quad l=i, j, \quad k=1, \sigma.$$

It is easy to verify that because $\{a_n\}$ is a deterministic process and processes $\{t_n^i\}$, $\{s_n^i\}$ are Poisson,

$$\frac{p(j, t_k)}{p(i, t_k)} = \frac{\nu_j}{\nu_i}, \quad k=1, \sigma.$$

Therefore, $J(\pi, K+1) - J(\bar{\pi}, K+1) \geq 0$ since $d_j \nu_j \geq d_i \nu_i$ and π cannot be optimal. A similar argument shows that a policy that idles in node 2 at time 0 cannot be optimal. Note that policies π and $\bar{\pi}$ are not feasible because they are allowed to switch at all the points of process $\{t_n\}$, some of which are not observable. The above argument shows that the "dν" rule is optimal among all such policies. Yet, the "dν" rule is a feasible policy and is optimal among feasible policies as well.

□

Remarks 2.1

(a) In the above proof, policies π and $\bar{\pi}$ result in identical arrivals for node 2. The proof, except for the embedding, is identical to the one in Buyukkoc et al [7].

(b) The result stated in Theorem 2.1 remains true if node 1 is replaced by a network of $M/1$ queues with Bernoulli routing (see Figure 2.1), and the cost function in (2.1) is modified in the obvious way.

3. Klimov's problem

3.1. The problem

The following situation was considered in Klimov [7]. There are N queues. The service times are independent and have the distribution function $G_i(t)$ in queue i ($1 \leq i \leq N$). Customers arrive as an independent Poisson process with rate λ and are assigned to queue i with probability p_i . Write $\mathbf{p} = (p_1, \dots, p_N)$. Upon service completion in queue i , a customer is sent to queue j with probability p_{ij} , and leaves the network with probability $p_{i0} = 1 - \sum_{j=1}^N p_{ij}$, independently of the state of the network. There is a single server that is allocated to one of the nodes at a time, in a nonpre-emptive way.

Assumptions

1. The matrix $P = [p_{ij}, 1 \leq i, j \leq N]$ is such that every customer eventually leaves, i.e., $P^n \rightarrow 0$ as $n \rightarrow \infty$. In particular, this implies that $(I - P)$ is invertible.

2. It is also assumed that $\int_0^\infty t dG_i(t) =: \beta_i < \infty$, for $1 \leq i \leq N$.

3. Finally, one assumes that $\lambda p[I - R]^{-1}\beta < 1$, where $\beta = (\beta_1, \dots, \beta_N)^T$ ($(\cdot)^T$ denotes transposition).

Denote by Z_t^i the number of customers at time $t \geq 0$ in queue $i \in \{1, \dots, N\}$ and let $\mathbf{Z}_t = (Z_t^1, \dots, Z_t^N)$. Fix $c_i \geq 0$ for $1 \leq i \leq N$ and such that $\sum_{i=1}^N c_i = 1$. For a given server allocation policy π , one defines the average waiting cost per unit of time as

$$J(\pi) = \liminf_{T \rightarrow \infty} \frac{1}{T} E \int_0^T \sum_{i=1}^N c_i Z_t^i dt. \quad (3.1)$$

A policy is said to be *admissible* if it is non-idling, nonpre-emptive and nonanticipative. Non-idling means that the server is idle only when the system is empty. Nonanticipative means that the decision to allocate the server to queue i at time $t \geq 0$ is based on the evolution of the network up to time t . Under Assumptions 1,2 and 3 the system is ergodic under any non-idling policy (see Section 3.2).

A policy is optimal if it minimizes the cost (1.1) over all the admissible policies. The problem is to find an optimal policy.

Outline

We give a simple proof and provide two extensions of the result in [7]. Specifically, we show that a priority list policy that serves the non-empty node with the highest priority is optimal. Remarkably, the priorities do not depend on the parameters of the arrival process. In Section 3.2 we discuss the effects of idling in the simple case of two nodes with no feedback. Some auxiliary calculations are performed in Section 3.3 and are used subsequently in Section 3.4 to derive a priority index for each node. The optimality results extend to the case of pre-emptive policies for nodes where the service distributions are exponential.

3.2. The busy period

Decomposition

Convention: The set of nodes is partitioned into $\mathbf{n} = \{1, \dots, n\}$ and $\mathbf{n}^c = \{n+1, \dots, N\}$. Assume that while the nodes in $\mathbf{n}$ are not all empty the server serves according to the priority rule $1 > 2 > \dots > n$.

Notation: For a matrix M and sets of natural numbers A and B , M_{AB} denotes the matrix $\{M_{ij}\}_{i \in A, j \in B}$. Similar notation for vectors has the obvious meaning.

Definition: Let $B^{(\mathbf{n})}$ be the time it takes to empty nodes $\mathbf{n}$, i.e.,

$$B^{(\mathbf{n})} = \inf\{t > 0 \mid \mathbf{Z}_n(t) = 0\}$$

We represent the queueing process in the system during a busy period as a collection of trees (see Feller [4]). A job arriving in node i with service requirement S_j is represented as a node of type i with weight S_j . The children of each node are jobs arriving in the system while the customer is being served. Each job initially present in the system is the root of a tree and the length of a busy period is the sum of the weights of these trees. It is shown in [13] that under Assumptions 2 and 3 one has $E[B^{(\mathbf{n})}] < \infty$ for all $\mathbf{n}$. It is easy to see by this construction that

Fact: The random variable $B^{(\mathbf{n})}$ does not depend on the order of service.

Furthermore, one obtains a decomposition for the mean of a busy period. Let $\mathbf{e}_i$ be the i th unit vector in $(-\infty, \infty)^N$ and for any $\mathbf{Z}(0)$ write $\mathbf{Z}(0) = m_1 \mathbf{e}_1 + \dots + m_N \mathbf{e}_N$. Then,

$$E[B^{(\mathbf{n})} \mid \mathbf{Z}(0)] = \sum_{i=1}^n m_i E[B^{(\mathbf{n})} \mid \mathbf{e}_i]. \quad (3.2)$$

3.3. Auxiliary calculations

Probabilities of transition

In this section we calculate the transition probabilities of a customer exiting the set of nodes $\mathbf{n}$, i.e., we are interested in the quantity

$r_{ij}^{(\mathbf{n})}$: Indicator function of the event that node $j \in \mathbf{n}^c$ is the first node not in $\mathbf{n}$ that a job visits starting in node i . The probabilities are given by

$$p_{ij}^{(\mathbf{n})} = E[r_{ij}^{(\mathbf{n})}]. \quad (3.3)$$

Lemma 3.1: The probabilities defined in (3.3) above are given, in matrix form by

$$P_{\mathbf{n}^c \mathbf{n}^c}^{(\mathbf{n})} = P_{\mathbf{n}^c \mathbf{n}^c} + P_{\mathbf{n}^c \mathbf{n}}(I - P_{\mathbf{nn}})^{-1}P_{\mathbf{n} \mathbf{n}^c} \quad (3.4)$$

Proof: For $i \in \mathbf{n}^c$ and $l \in \mathbf{n}$ we write the following first step equations.

$$p_{ij}^{(\mathbf{n})} = p_{ij} + \sum_{l \in \mathbf{n}} p_{il} p_{lj}^{(\mathbf{n})} \quad (3.5)$$

$$p_{lj}^{(\mathbf{n})} = \sum_{l' \in \mathbf{n}} p_{ll'} p_{l'j}^{(\mathbf{n})} + p_{lj} \quad (3.6)$$

The result is obtained by writing the above equations in matrix form and solving (3.5) for $P_{\mathbf{n}^c \mathbf{n}^c}^{(\mathbf{n})}$.

□

Expected sojourn times

We now calculate the expected sojourn time for each passage of a customer through the set of nodes $\mathbf{n}$. For this define

$S_j^{(\mathbf{n})}$: Total amount of service that a job receives until it exits $\mathbf{n}$ having started at node $j \in \mathbf{n}^c$.

We set

$$T_j^{(\mathbf{n})} = E[S_j^{(\mathbf{n})}] \quad (3.7)$$

Lemma 3.2: The expected sojourn times defined in (3.7) are given by

$$T_{\mathbf{n}^c}^{(\mathbf{n})} = P_{\mathbf{n}^c \mathbf{n}}(I - P_{\mathbf{nn}})^{-1}\beta_{\mathbf{n}} \quad (3.8)$$

Proof: For $j \in \mathbf{n}^c$ and $l \in \mathbf{n}$ first step equation give

$$T_j^{(\mathbf{n})} = \sum_{l \in \mathbf{n}} p_{jl} T_l^{(\mathbf{n})}$$

$$T_l^{(\mathbf{n})} = \beta_l + \sum_{l' \in \mathbf{n}} p_{ll'} T_{l'}^{(\mathbf{n})}$$

The proof then proceeds as in Lemma 3.1 above.

□

We next turn our attention to a quantity that will be important in the computation of priority indices in Section 3.5. Define

$R_j^{(\mathbf{n})}$: Total amount of time needed to clear the set of nodes $\mathbf{n}$ of the arrivals resulting from serving a job in node $j \in \mathbf{n}^c$ and through its sojourn in $\mathbf{n}$.

Lemma 3.3: The expectation of $R_j^{(\mathbf{n})}$ is given by

$$E[R_j^{(\mathbf{n})}] = \lambda(\beta_j + T_j^{(\mathbf{n})}) \sum_{i=1}^n p_i E[B^{(\mathbf{n})} | \mathbf{e}_i] \quad (3.9)$$

Proof: By $A(t)$ denote the number of arrivals in the interval $[0, t]$. The result then follows easily from relationship (3.2) and by noting that

$$\begin{aligned} E[R_j^{(n)}] &= E[E[B^{(n)} | A(S_j + S_j^{(n)})]] \\ &= \sum_{i=1}^n E[A_i(S_j + S_j^{(n)})] E[B^{(n)} | e_i] \\ &= \lambda(\beta_j + T_j^{(n)}) \sum_{i=1}^n p_i E[B^{(n)} | e_i] \end{aligned}$$

The last step follows from the fact that the arrival process is Poisson.

□

3.4. Optimality

The nonpre-emptive case

In this section we prove that, as mentioned in the Section 3.1, the policy that minimizes the cost defined in (3.1) is a priority list which we determine. It is clear that an optimal policy also minimizes the expected cost in each busy period of the system given by

$$J(\pi, B) = E \left[\int_0^B \mathbf{c} \cdot \mathbf{Z}(t) dt \right]$$

We give expressions for the priority index of each node.

First, the nodes $\{1, 2, \dots, N\}$ are renumbered as follows. Assign number 1 to the node that maximizes the quantity

$$\frac{c_i - \sum_k p_{ik} c_k}{\beta_i}, \quad i = 1, \dots, N. \quad (3.10)$$

Recursively, for $1 \leq n < N$, assign the number $n+1$ to the node $i \in \mathbf{n}^c$ that maximizes the quantity

$$\frac{c_i - \sum_{k \in \mathbf{n}^c} p_{ik}^{(n)} c_k}{\beta_i + T_i^{(n)}},$$

where $\mathbf{n}$ is the set of nodes $\{1, \dots, n\}$ in the new numbering. Denote by π the priority assignment list that corresponds to this ordering.

Theorem 3.1: Policy π is optimal among all nonpre-emptive, non-idling policies.

Proof: Let $J(\pi' \pi, B)(\mathbf{Z})$ be the cost incurred in a busy period starting from state $\mathbf{Z}$ and following policy π' in the first step and π thereafter. Then it suffices to prove that

$$J(\pi' \pi, B)(\mathbf{Z}) \geq J(\pi, B)(\mathbf{Z}), \text{ for all } \mathbf{Z} \in \{0, 1, \dots\}^N. \quad (3.11)$$

It suffices to consider the case where $\pi'(\mathbf{Z}) = i \neq \pi(\mathbf{Z}) = j$ with $i > j$. This implies that $\mathbf{Z} = (0, \dots, 0, Z^j, \dots, Z^i, *, \dots, *)$ with $Z^i Z^j > 0$. For simplicity consider that $\mathbf{Z}_0 = \mathbf{Z}$. To establish (3.11) define ρ to be the first time that policy $\pi' \pi$ serves node j . By ϵ (respectively ς) denote the job that was served in node i at time 0 (respectively in node j at time ρ). In the context of Section 3.2, ρ is the time it takes to clear the system of the descendants of job ϵ that have priority higher than j . Let l be the node in $(j-1)^c$ where job ϵ ends up after its sojourn in $j-1$ and let $\mathbf{x}$ be the vector of the rest of the descendants of ϵ after their sojourn in $j-1$. Then define $\rho + \sigma$ to be the time it takes to serve job ς in node j and clear the system of the descendants of ς that have priority higher than j . Similarly, let ς end up in $k \in (j-1)^c$ after its sojourn in $j-1$ and

let $\mathbf{y}$ be the vector of the rest of the descendants of ζ after their sojourn in $j-1$. For each sample path that is obtained by applying policy π' construct a sample path where π is followed until time σ , then any policy $\bar{\pi}$ such that $\bar{\pi}(\mathbf{Z}(\sigma)) = i$ is followed for one step, and π is resumed afterwards. Denote this policy by $\pi^{(\sigma)}\bar{\pi}\pi$. The arrival and service processes of jobs with priority higher than j are interchanged as in the construction of Section 3.2, i.e., the descendants of job ϵ with priority higher than j and the descendants of job ζ with priority higher than j are the same in both realizations. The arrival and service processes of jobs with priority lower than j are the same in both realizations. One then obtains (see Figure 3.1),

$$\begin{aligned} J(\pi' \pi, B)(\mathbf{Z}) - J(\pi^{(\sigma)}\bar{\pi}\pi, B)(\mathbf{Z}) &= \\ &= c_j(\beta_i + T_i^{(j-1)} + E[R_i^{(j-1)}]) + (\sum_{l \in (j-1)^c} p_{il}^{(j-1)} c_l + c \cdot E[\mathbf{x}])(\beta_j + T_j^{(j-1)} + E[R_j^{(j-1)}]) \\ &\quad - \{c_i(\beta_j + T_j^{(j-1)} + E[R_j^{(j-1)}]) + (\sum_{k \in (j-1)^c} p_{jk}^{(j-1)} c_k + c \cdot E[\mathbf{y}])(\beta_i + T_i^{(j-1)} + E[R_i^{(j-1)}])\}. \end{aligned} \quad (3.12)$$

To simplify this expression we need to determine $E[\mathbf{x}]$ and $E[\mathbf{y}]$ as functions of the system parameters. For this, denote by $a_m^{(j-1)}$ the expected number of jobs that enter node $m \in j-1$ during a busy cycle that starts with a job in node $l \in j-1$. From Section 3.2 one gets

$$c \cdot E[\mathbf{x}] = \sum_{i' \in (j-1)^c} c_{i'} \sum_{m \in j-1} p_{mi'}^{(j-1)} \sum_{l \in j-1} a_m^{(j-1)} \lambda p_l (\beta_i + T_i^{(j-1)})$$

From this, relation (3.9) and some rearrangement one gets that $J(\pi' \pi, B) - J(\pi^{(\sigma)}\bar{\pi}\pi, B) \geq 0$ if

$$\frac{c_j - \sum_k p_{jk}^{(j-1)} c_k}{\beta_j + T_j^{(j-1)}} \geq \frac{c_i - \sum_k p_{ik}^{(j-1)} c_k}{\beta_i + T_i^{(j-1)}}$$

To complete the proof one now argues that

$$J(\pi' \pi, B) \geq J(\pi^{(\sigma)m} \bar{\pi} \pi, B) \xrightarrow{m \rightarrow \infty} J(\pi, B) \quad (3.13)$$

by bounded convergence. Policy $\pi^{(\sigma)m} \bar{\pi} \pi$ is defined recursively as $\pi^{(\sigma)} \pi^{(\sigma)(m-1)} \bar{\pi} \pi$ for $m=1, 2, \dots$.

Suppose now that policy $\bar{\pi}$ idles for some amount of time, at state $\mathbf{Z}$. The above argument then shows that $\bar{\pi}$ can be improved and thus policy π is optimal among idling policies as well. □

Remark 3.2: The argument used in the above proof is a variation of an argument in Varaiya et al [14]. We have followed closely the notation in Weiss [16], where a similar argument appears. Our argument also gives a simple proof of the results in Foss [5] who considers a generalized version of Klimov's problem and obtains a corresponding index rule.

The pre-emptive case

Assume now that the service time distributions are exponential at all the nodes and consider the coupling described above where π is now a pre-emptive policy following the same priority assignment list as in the nonpre-emptive case. One then sees that $J(\pi' \pi, B) - J(\pi^{(\sigma)}\bar{\pi}\pi, B)$ is the same as in (3.12). It follows that π is optimal among pre-emptive policies.

Remark 3.3: While this paper was under review, the paper of Lai and Ying [9] appeared. There, asymptotics of the "open bandit problem" are studied as the discount factor approaches 1. The above results are then derived. Our approach is simpler in that it does not rely on previous results on multi-armed bandits. At the same time, the results in [9] can be simply obtained by a direct argument similar to ours (see [8] and [16]).

4. A communication link model

In this section we demonstrate how interchange arguments can be employed in problems that do not fall in the multi-armed bandit framework.

The model

We consider the following model of a communication link. There are N channels to be used for the transmission of telephone calls. The calls arrive according to a deterministic sequence $\{a_k\}_k$. Each call is to be placed on one of the idle communication links, if one is available, and is lost otherwise. A call placed on link i is immediately lost with probability p_i and with probability $1-p_i$ it occupies the link for a period of time which is exponentially distributed with parameter μ_i . The system is described by the vector $\mathbf{Z} \in \{0,1\}^N$ where $Z^i = 1$ if a call is present at link i and $Z^i = 0$ otherwise. A placement policy is a function

$$\mathbf{u} : \mathbf{Z} \rightarrow \mathbf{u}(\mathbf{Z}) \in \{1, \dots, N\}.$$

such that $\mathbf{Z}^*(\mathbf{Z}) = 0$ if $\mathbf{Z} \neq (1, \dots, 1)$. That is, in state $\mathbf{Z}$ policy $\mathbf{u}$ will place the next arrival on link $\mathbf{u}(\mathbf{Z})$. We have restricted ourselves to deterministic policies. Our arguments however, easily extend to randomized policies.

Outline

First, the case where $\mu_i = \mu$, $i = 1, \dots, N$ is considered. We prove that from any state $\mathbf{Z}$ the time $T_{\mathbf{Z}}$ it takes to reach state $(1, \dots, 1) \in \{0,1\}^N$ has a distribution that is independent of the policy $\mathbf{u}$. This result was obtained by Smith [12] and Anantharam et al [1] by explicit computation of the moment generating function of $T_{\mathbf{Z}}$.

Next, for unequal μ_i 's, the problem of stochastically maximizing $T_{\mathbf{Z}}$ for any initial state $\mathbf{Z}$ is considered. We prove that the optimal policy always places calls on the free channel with the largest μ_i .

4.1. Invariance

In this subsection it will be assumed that $\mu_1 = \mu_2 = \dots = \mu_N = \mu$.

Theorem 4.1: For any $\mathbf{Z} \in \{0,1\}^N$, the distribution of $T_{\mathbf{Z}}$ does not depend on the policy $\mathbf{u}$.

Proof: We will use a stochastic variation of the argument used in Section 3.5. Let $\mathbf{u}$ be a priority list assigning calls in the order $1, \dots, N$. As in the proof of Theorem 3.1 it will suffice to show that

$$T_{\mathbf{Z}}(\mathbf{u}' \mathbf{u}) = T_{\mathbf{Z}}(\mathbf{u}) \quad (4.1)$$

for any policy $\mathbf{u}'$. To establish (4.1) it suffices to consider the case where $\mathbf{Z} \neq (1, \dots, 1)$ and $\mathbf{u}'(\mathbf{Z}) = i \neq \mathbf{u}(\mathbf{Z}) = j$ with $i > j$. This implies that $Z^l = 1$, $l = 1, \dots, j-1$, and $Z^i = Z^j = 0$.

Arguing again as in the proof of Theorem 3.1, we will establish an analogue of relationship (3.13). To this end, denote by a_o the first time that $\mathbf{u}' \mathbf{u}$ places a call on link j and assume for simplicity that $a_1 = 0$. Denote the virtual service processes of the links by $\{S_t^l\}_{l=1}^N$ and their points by $\{s_n^l\}_{n=1}^N$. Also, for $l = 1, \dots, N$, set $r_n^l = 1$ if the n th trial to engage link l is a success, and 0 otherwise.

For each sample path of $(\mathbf{Z}_t)$ resulting from $\mathbf{u}' \mathbf{u}$ we construct a sample path of a process $(\bar{\mathbf{Z}}_t)$ as follows. Consider a policy that places a call on link j at time a_1 , follows policy $\mathbf{u}$ afterwards, and places a call on link i at time a_o , if $a_o < s_1^i$. Then, $\mathbf{Z}_{a_o} = \bar{\mathbf{Z}}_{a_o}$ and the paths of $(\mathbf{Z}_t)$ and $(\bar{\mathbf{Z}}_t)$ can be made to coincide from time a_o onward. This would be the obvious argument in the case where $\mu = 0$. It would suffice to let $\bar{r}_n^l = r_n^l$, $l = 1, \dots, N$, $n = 1, 2, \dots$.

On the other hand, if $s_1^i < a_o$, the paths of $(\mathbf{Z}_t)$ and $(\bar{\mathbf{Z}}_t)$ can again be made to coincide from time s_1^i onward. This is achieved in the construction of $(\bar{\mathbf{Z}}_t)$ by letting $\{S_t^j\}$ be the virtual service process at link j , i.e., $\bar{s}_1^j = s_1^i$. Then, $\mathbf{Z}_{s_1^i}^j = \bar{\mathbf{Z}}_{s_1^i}^j = 0$. The two cases are illustrated in Figures 4.1(a) and 4.1(b) respectively.

Formally, in the construction of process $(\bar{Z}_t)$ the arrival process remains $\{a_k\}$ and policy u is followed (recall $u(Z) = j$) until $\tau = a_o \wedge s_1^i$ with

$$\begin{aligned}\bar{S}_t^l &= S_t^l, l \neq i, j, t \geq 0, \\ \bar{r}_n^l &= r_n^l, l = 1, \dots, N, n = 1, 2, \dots\end{aligned}$$

$$\bar{S}_t^j = S_t^j, t \leq \tau.$$

If $\tau = s_1^i = \bar{s}_1^j$ then continue with u . Otherwise, i.e., if $\tau = a_o$, follow $\bar{u}(\bar{Z}_\tau) = i$ and then continue with u . Denote this composite policy by $\bar{u}$. In either case set

$$\bar{S}_t^i = S_t^i, \bar{S}_t^j = S_t^j, t > \tau.$$

We have thus obtained that

$$T_Z(u' u) \stackrel{st}{=} T_Z(u^{(\tau)} \bar{u}).$$

The proof now concludes as in Theorem 3.1. □

Remark 4.1: The model considered here is similar to the well known repairman model (see e.g. Nash and Weber [10]). Our methods should apply to that model as well. In particular, one should be able to simply obtain the results in Hirayama [6] where a related optimization problem is studied.

4.2. Optimality

In this subsection we assume that $\mu_1 \geq \mu_2 \geq \dots \geq \mu_N$ and that policy u assigns calls according to the priorities $1, \dots, N$.

Theorem 4.2: For any $Z \in \{0, 1\}^N$, T_Z is stochastically maximized by policy u .

Proof: Again, as in Theorem 3.1, we will show that

$$T_Z(u' u) \stackrel{st}{\leq} T_Z(u)$$

where $u'(Z) = i \neq u(Z) = j$ with $i > j$. With notation as in the proof of Theorem 4.1 one can check that using a similar construction for a process $(\bar{Z}_t)$ we have

$$T_Z(u' u) \leq T_Z(u^{(\tau)} \bar{u}), \text{ a.s.}, \quad (4.2)$$

on some probability space, where the stopping time τ remains to be specified.

The construction of $(\bar{Z}_t)$ only differs from the one in the proof of Theorem 4.1 in that $\bar{S}^j[0, \tau]$ is a superset of $S^i[0, \tau]$. This can be done since $\mu_j \geq \mu_i$. Also, note that in this case it suffices to define τ as

$$\tau = \inf\{t \mid Z_t \geq \bar{Z}_t\}$$

where the inequality is component-wise. □

5. Acknowledgements

We would like to thank Professor A. Makowski, and Mr. I. Lambadaris for useful discussions.




Figure 2.1






Figure 3.1



○ = $u'u$
● = $u\bar{u}$

(a)



○ = $u'u$
● = $u\bar{u}$

(b)

Figure 4.1

