

SAMPLE PATH METHODS IN THE CONTROL OF QUEUES

Zhen LIU^{1*}, Philippe NAIN^{1†} and Don TOWSLEY^{2‡}

¹INRIA, B.P. 93, 06902 Sophia Antipolis Cedex, France

²Department of Computer Science
University of Massachusetts, Amherst, MA 01003, USA

Published in: **QUESTA 21 (1995) 293-335**

Abstract

Sample path methods are now among the most used techniques in the control of queueing systems. However, due to the lack of mathematical formalism, they may appear to be non-rigorous and even sometimes mysterious. The goal of this paper is threefold: to provide a general mathematical setting, to survey the most popular sample path methods including forward induction, backward induction and interchange arguments, and to illustrate our approach through the study of a number of classical scheduling and routing optimization problems arising in queueing theory.

Keywords: Control, scheduling, mathematical formalism, sample path arguments, queueing system, discrete event system, stochastic comparison.

*Z. Liu was supported in part by the CEC DG XIII under the ESPRIT BRA grant QMIPS.

†P. Nain was supported in part by NSF under grant NCR-9116183 and by the CEC DG XIII under the ESPRIT BRA grant QMIPS.

‡D. Towsley was supported in part by NSF under grant NCR-9116183.

1 Introduction

There is increasing interest in the subject of the control of queueing systems. This is due in large part to the wide applicability of results in this area to the control of computers and communication systems. Consequently, a number of methodologies have been developed to study these problems. These include queueing theory, discrete event simulation, Markov decision processes, and sample path analysis. The first two of these methodologies are solely concerned with evaluating the performance of control policies whereas the last two are concerned with either determining and characterizing the optimal control policies or obtaining qualitative properties of these policies.

In this paper we focus on the last of these four methodologies, sample path analysis. As the name indicates, sample path analysis aims at comparing, sample path by sample path, stochastic processes defined on a common probability space, so that the “optimal behavior” of the system at hand can be identified. Because of its power and wide-range applicability sample path analysis has been used for almost two decades to solve optimization problems in the control of queues. In spite of this, no general formalism has been proposed and statements are often not proven in a rigorous way. Instead, helpful, but not entirely satisfactory, intuitive explanations are provided. In contrast with the solid mathematical theory of Markov decision processes (see e.g., Ross [27], Schal [28]), sample path methods may appear to be less rigorous, not to say mysterious to the unfamiliar reader! More importantly, this lack of formalism opens the door to mistakes of various kinds as the problems to be studied become more and more intricate and the arguments more and more involved. Examples of such mistakes have been reported in Hordijk and Koole [15] and Towsley and Sparaggis [36].

In this paper we provide a theoretical framework for a rigorous use of sample path methods. In order to do this, we propose a general mathematical setting that can be used to address various control problems arising in queues and, more generally, in the area of discrete event dynamical systems. This formalism is based on recent works by Liu et al. [17] and Nain and Towsley [26] on optimal service disciplines in multi-class single-server systems. We demonstrate the versatility of this approach by applying it to a number of classical scheduling and routing optimization problems in queues.

The proof techniques we use belong to three general classes: forward induction, backward induction and interchange arguments. When applicable, forward induction provides a simple way of obtaining very strong sample path comparisons between an arbitrary control or policy and the optimal policy. In order to be applied, it is necessary to identify the hypothesized optimal policy and to derive equations that capture the time-behavior of the system under any policy. The proof then consists of a forward induction argument over times of all distinct events, establishing an ordering between the hypothesized optimal policy and the arbitrary policy. Backward induction appears to be more involved than forward induction. It requires more structural properties of the hypothesized optimal policy. The approach based on interchange arguments is generally the most delicate of the three approaches since it relies on the comparison of policies obtained through the interchange of control decisions and/or the order of some events in the input sequence (typically, arrival times of customers, service times, routings). When this happens, it is then essential to verify that the statistics of the input data have not been changed and that inadmissible control actions (see definition in Section 2)

are not generated. It is worth noting that all three proof techniques are used in the Markov decision process framework and are not unique to the sample path framework.

For each of these sample path methods, particular care is required when setting up the problem. This includes choosing an appropriate state description, defining the underlying probability space, writing down the dynamics of the system, choosing the statistical assumptions to be placed on the input data, defining the class of admissible controls, selecting the cost function, etc. There are typically several ways to describe the stochastic evolution of a system, not all of which permit a comparison between two given policies. Once policies have been compared on a pathwise basis, one can then compare the expectations of functions of one or more of the state variables. In general, this is done by invoking the property that the statistics of the input data have not been modified and by resorting to stochastic ordering (in particular, integral ordering) and stochastic comparison techniques (see Section 3).

Sample path analysis is usually performed in the transient regime of the system. In order to compare policies in steady state, one can either combine transient comparison results with a limiting argument based on convergence assumptions, or work directly in the stationary regime. Both approaches are illustrated in Section 6.

The paper is organized as follows. In Section 2, we introduce the mathematical setting that will be used throughout the paper. In Section 3, we define various deterministic and stochastic partial orderings useful for deriving optimality properties of control policies. We also describe useful properties related to these orderings. In Section 4, we illustrate the use of forward induction in establishing the optimality of control policies for queueing systems. Two applications of the forward induction argument are considered: in the first one we establish the optimality of the join the shortest queue policy for a queueing system with identical servers; in the second one we solve a scheduling problem for a multiclass G/G/1 queue and use this result to show the optimality of the μc -rule for the multiclass G/M/1 queue. In Section 5, we illustrate the use of backward induction in establishing the optimality of the routing policy that routes arriving customers to the queue with the smallest workload for a system of identical servers, each with its own queue. The routing problems considered in Sections 4 and 5 differ in the information available to the decision-maker. Finally, in Section 6, we illustrate the use of interchange arguments by revisiting two optimization problems: in the first one, we discuss the extremal properties of the First-Come-First-Served (FCFS) and Last-Come-First-Served (LCFS) service discipline with respect to customer response times; the second one is the well-known Klimov's problem. Reviews and discussions of related studies are given at the end of each section.

Because some proofs are very similar, some of them are omitted for the sake of exposition. However, each important concept is illustrated within at least one proof. This includes proofs that intermediate policies are admissible and that the statistics of the input data have not been modified after the interchange of events. The proofs of these kinds of assertions are subtle and sometimes tedious. However, they are important since they usually show why a policy is optimal within a particular class of policies, and illustrate the need for particular statistical assumptions on the input data that are required for the claimed result to be true.

2 The Mathematical Setting

In this section we outline the mathematical setting that will be used throughout this paper.

Let (Ω, \mathcal{F}, P) be a probability space. Throughout this paper Ω will be the canonical space of all standard queueing input sequences such as, customer arrival times, customer service requirements, customer routings, etc. Let ϕ be the identity mapping on Ω .

A typical example is $\Omega = \prod_{n=1}^{\infty} \Omega_n$ with $\Omega_n = \mathbb{R}_+ \times \mathbb{R}_+$, where a generic element in Ω is written in the form $\omega = (\omega_1^0, \omega_1^1, \dots, \omega_n^0, \omega_n^1, \dots)$. It will be convenient to define the coordinate processes $(\phi_n^m)_{m,n}$ on Ω as $\phi_n^m(\omega) = \omega_n^m$ for all $n = 1, 2, \dots, m = 0, 1$, so that $\phi = (\phi_1^0, \phi_1^1, \dots, \phi_n^0, \phi_n^1, \dots)$. Here, ϕ_n^0 and ϕ_n^1 may represent the arrival time and the service requirement of the n -th customer in a queueing system, respectively.

We shall consider stochastic processes $(X_n)_n$ of the form

$$X_n = f_n(\phi, (X_i)_{i=1}^{n-1}, (U_i)_{i=1}^{n-1}) \in S_n, \quad n = 2, 3, \dots, \quad (2.1)$$

where the initial state $X_1 \in S_1$ is known, $U_i \in A$ for $i = 1, \dots, n-1$, and f_n is a measurable mapping from $\Omega \times \prod_{i=1}^{n-1} S_i \times A^{n-1}$ into S_n , respectively. Here, $u = (U_n)_n$ is the *control* and A is the action set.

To illustrate this point let us return to the queueing system example. Define

$$X_{n+1} = \max(0, X_n + \mathbf{1}(U_n = 1) \phi_n^1 - (\phi_{n+1}^0 - \phi_n^0)), \quad n = 1, 2, \dots,$$

with $X_1 = 0$ and $U_n \in \{0, 1\}$. Then, X_n represents the workload at the arrival time of the n -th customer in a FCFS single-server queue with customer inter-arrival times $(\phi_{n+1}^0 - \phi_n^0)_n$ and customer service requirements $(\phi_n^1)_n$, where customers may be denied access to the queue upon arrival (customer n is rejected if $U_n = 0$).

Associated with each system will be a collection of sets $(H_n)_n$, a collection $(A(k))_k$ of subsets of the action set A , a sequence of measurable mappings $(g_n)_n$, $g_n : \Omega \times \prod_{i=1}^n S_i \times A^{n-1} \rightarrow H_n$, and a sequence of measurable mappings $(\alpha_n)_n$, $\alpha_n : H_n \rightarrow \{0, 1, \dots\}$. We shall say that a control $u = (U_n)_n$ is *admissible* if, for every $n = 1, 2, \dots$,

$$U_n = \pi_n(Y_n) \quad (2.2)$$

$$U_n \in A(\alpha_n(Y_n)), \quad (2.3)$$

with $Y_n := g_n(\phi, (X_i)_{i=1}^n, (U_i)_{i=1}^{n-1})$, where π_n is a measurable mapping from $H_n \rightarrow A$ for every $n = 1, 2, \dots$

Condition (2.2) reflects the fact that the decision-maker may only have partial information (i.e. Y_n) on the system at a decision epoch, whereas (2.3) gives the admissible actions at a decision epoch. In the following we will identify the set of all admissible controls for the problems at hand with the

set of all mappings $\pi := (\pi_n)_n$, hereafter referred to as *admissible policies* with a slight abuse of terminology, associated with each admissible control, and we will use the superscript π to indicate the enforced control (i.e., X_n^π , U_n^π and Y_n^π).

Coming back to the example of the single-server queue, let $A = A(1) = \{0, 1\}$ and $A(0) = \{1\}$. Assume that $Y_n^\pi = \left((X_i^\pi)_{i=1}^n, (U_i^\pi)_{i=1}^{n-1} \right)$ for all $n \geq 1$, and $\alpha_n(Y_n^\pi) = \mathbf{1}(X_n^\pi > 0)$. Thus, in this case the decision to admit or to reject a new customer is only based on the past and current state of the workloads (and not, for instance, on future arrival times or future service times) as well as on the past decisions. Also note from the constraints placed on U_1^π, U_2^π, \dots that a customer finding an empty system will always be admitted.

We conclude this section with a discussion on randomized policies. Randomized policies may easily be taken care of within this setting by enlarging the state-space Ω as shown in Nain and Towsley [26]. In this paper we shall not consider randomized policies for the sake of clarity. Another reason for this choice is that there will exist optimal policies that are non-randomized in all optimization problems that we shall address.

A few words about some notation to be used throughout this text: \mathbf{R} (resp. \mathbf{R}_+ , \mathbf{N}) will denote the set of all real numbers (resp. nonnegative real numbers, non-negative integers). For any X -valued mappings f and g , the notation $f \equiv g$ will stand for $f(x) = g(x)$ for all $x \in X$.

3 Comparison Methods

In this section, we introduce some partial orderings on \mathbf{R}^n , $n \geq 1$, that are used for deriving our main results. We begin with orderings between deterministic vectors and then conclude with stochastic orderings between random vectors.

Let $x, y \in \mathbf{R}^n$. The elementwise ordering $x_i \leq y_i$, $i = 1, \dots, n$, is denoted by $x \leq y$. For any $x = (x_1, x_2, \dots, x_n) \in \mathbf{R}^n$, let $x_{[1]} \geq \dots \geq x_{[n]}$ denote the components of x in decreasing order. Let Θ be the set of permutations on $\{1, 2, \dots, n\}$. For any vector $x \in \mathbf{R}^n$ and for every permutation $\theta \in \Theta$, define $x_\theta = (x_{\theta(1)}, \dots, x_{\theta(n)})$.

Definition 3.1 Consider a function $f : \mathbf{R}^n \rightarrow \mathbf{R}$.

- (i) f is increasing (decreasing) if $x \leq y$ implies $f(x) \leq (\geq) f(y)$. Here increasing (decreasing) is taken to mean non-decreasing (non-increasing);
- (ii) f is convex (concave) if, for all $\alpha \in [0, 1]$ and $x, y \in \mathbf{R}^n$, $z = \alpha x + (1 - \alpha)y$ implies $f(z) \leq (\geq) \alpha f(x) + (1 - \alpha)f(y)$;
- (iii) f is symmetric if $f(x) = f(y)$ for any $y = x_\theta$, where $\theta \in \Theta$;
- (iv) f is convex (concave) in each variable if

$$f(z_1, \dots, \alpha x_i + (1 - \alpha)y_i, \dots, z_n) \leq (\geq) \alpha f(z_1, \dots, x_i, \dots, z_n) + (1 - \alpha)f(z_1, \dots, y_i, \dots, z_n)$$

for all $\alpha \in [0, 1]$, $i = 1, \dots, n$;

(v) f is L -subadditive (L -superadditive) if

$$\begin{aligned} f(x_1, \dots, x_i + \epsilon_1, \dots, x_j + \epsilon_2, \dots, x_n) &+ f(x_1, \dots, x_i, \dots, x_j, \dots, x_n) \leq (\geq) \\ f(x_1, \dots, x_i + \epsilon_1, \dots, x_j, \dots, x_n) &+ f(x_1, \dots, x_i, \dots, x_j + \epsilon_2, \dots, x_n). \end{aligned}$$

The following deterministic partial orderings between vectors are useful in sample path arguments.

Definition 3.2 Let $x, y \in \mathbb{R}^n$.

(i) x is majorized by y ($x \prec y$) if $\sum_{i=1}^k x_{[i]} \leq \sum_{i=1}^k y_{[i]}$, $k = 1, \dots, n-1$ and $\sum_{i=1}^n x_{[i]} = \sum_{i=1}^n y_{[i]}$;

(ii) x is weakly submajorized by y ($x \prec_w y$) if $\sum_{i=1}^k x_{[i]} \leq \sum_{i=1}^k y_{[i]}$, $k = 1, \dots, n$;

(iii) x is weakly supermajorized by y ($x \prec^w y$) if $\sum_{i=k}^n x_{[i]} \geq \sum_{i=k}^n y_{[i]}$, $k = 1, \dots, n$;

(iv) x is smaller than y in the partial sum sense ($x \prec_{ps} y$) if $\sum_{i=1}^k x_i \leq \sum_{i=1}^k y_i$, $k = 1, \dots, n$.

We introduce now a partial ordering on Θ which will be of use later on. Define the binary relation \mathcal{R} on Θ as: $\theta' \mathcal{R} \theta$ if there exist a pair of integers $j < k$, such that

$$\theta(j) > \theta(k), \quad \theta'(j) = \theta(k), \quad \theta'(k) = \theta(j), \quad \theta'(i) = \theta(i), \quad i \neq j, i \neq k. \quad (3.1)$$

Define now a partial order \prec_{Θ} on Θ as the transitive closure of \mathcal{R} , defined by:

1. $\theta' \prec_{\Theta} \theta$ if $\theta' \mathcal{R} \theta$;
2. $\theta' \prec_{\Theta} \theta$ if there exists θ'' such that $\theta' \prec_{\Theta} \theta''$ and $\theta'' \prec_{\Theta} \theta$.

Lemma 3.1 Let $x_1 \leq x_2 \leq \dots \leq x_n$ and $y_1 \leq y_2 \leq \dots \leq y_n$ be real numbers. If $\theta' \prec_{\Theta} \theta$, then

$$(y_{\theta'} - x) \prec (y_{\theta} - x). \quad (3.2)$$

This is a special case of Lemma 3.4 in Baccelli et al. [1].

We will find the following corollary useful later on.

Corollary 3.1 Let $x_1 \leq x_2 \leq \dots \leq x_n$ and $y_1 \leq y_2 \leq \dots \leq y_n$ be real numbers. Then, for all $\theta \in \Theta$

$$(y - x) \prec (y_{\theta} - x). \quad (3.3)$$

This is a special case of Corollary 3.2 in Day [5].

When vectors are restricted to take nonnegative integer values, $x, y \in \mathbf{N}^n$, we have the following operations that preserve \prec . Let $\beta_x, \chi_x \in \Theta$ be two permutations such that

$$x_{\beta_x(i)} < x_{\beta_x(i+1)}, \quad \text{or} \quad x_{\beta_x(i)} = x_{\beta_x(i+1)} \quad \text{and} \quad \beta_x(i) < \beta_x(i+1),$$

$$x_{\chi_x(i)} > x_{\chi_x(i+1)}, \quad \text{or} \quad x_{\chi_x(i)} = x_{\chi_x(i+1)} \quad \text{and} \quad \chi_x(i) > \chi_x(i+1).$$

In words, $\beta_x(i)$ (resp. $\chi_x(i)$) gives the index of the i -th smallest (resp. largest) component of x . For example, if $x = (3, 5, 6, 3, 5, 2)$, then $\beta_x = (6, 1, 4, 2, 5, 3)$ and $\chi_x = (3, 5, 2, 4, 1, 6)$.

Let $\mathbf{e}_i \in \mathbf{N}^n$ be the vector whose components are zero except the i -th component which is one.

Define $A_i(x) = x + \mathbf{e}_{\chi_x(i)}$ the vector that is obtained by adding one to the i -th largest component of x , and $D_i(x) = x - \mathbf{1}_{(x_{\chi_x(i)} > 0)} \mathbf{e}_{\chi_x(i)}$ the vector that is obtained by subtracting one from the i -th largest component of x if this component is strictly positive. These operations would correspond to the arrival or departure of a customer, respectively, at the i -th largest queue in a queueing system.

Lemma 3.2 For $x, y \in \mathbf{N}$,

- (i) if $x \prec (\prec_w) y$, then $A_i(x) \prec (\prec_w) A_j(y)$, for $1 \leq j \leq i \leq n$;
- (ii) if $x \prec_w y$, then $D_i(x) \prec_w D_j(y)$, for $1 \leq i \leq j \leq n$;
- (iii) if $x \prec_w y$, then $D_i(x) \prec_w y$, for $1 \leq i \leq n$;
- (iv) if $x \prec_w y$, then $x \prec_w A_j(y)$, for $1 \leq j \leq n$.

Definition 3.3 A function $f : \mathbf{R}^n \rightarrow \mathbf{R}$ is Schur convex if $f(x) \leq f(y)$ for all $x \prec y$.

The following result can be found in Marshall and Olkin [22, p. 59].

Lemma 3.3 A function $f : \mathbf{R}^n \rightarrow \mathbf{R}$ is increasing (resp. decreasing) and Schur convex if and only if $f(x) \leq f(y)$ for all $x \prec_w y$ (resp. $x \prec^w y$).

We are interested in the following classes of functions.

$$\mathcal{C}_1 = \{f : \mathbf{R}^n \rightarrow \mathbf{R} \mid f \text{ Schur convex} \};$$

$$\mathcal{C}_2 = \{f : \mathbf{R}^n \rightarrow \mathbf{R} \mid f \text{ convex and symmetric} \};$$

$$\mathcal{C}_3 = \{f : \mathbf{R}^n \rightarrow \mathbf{R} \mid f(x) = \sum_{i=1}^n g(x_i) \text{ for some convex function } g : \mathbf{R} \rightarrow \mathbf{R}\};$$

$$\mathcal{C}_4 = \{f : \mathbf{R}^n \rightarrow \mathbf{R} \mid f(x) = \sum_{i=1}^j x_{[i]} \text{ or } f(x) = -\sum_{i=j}^n x_{[i]} \text{ for some } j = 1, \dots, n\};$$

$$\mathcal{C}_5 = \{f : \mathbf{R}^n \rightarrow \mathbf{R} \mid f \text{ symmetric, L - subadditive, and convex in each variable} \}.$$

The classes of functions \mathcal{C}_i , $i = 1, 2, 3, 4$, were introduced in Marshall and Olkin [22] and class \mathcal{C}_5 in Chang [4]. Let $\mathcal{C}_i^\uparrow = \{f \mid f \in \mathcal{C}_i \text{ and } f \text{ increasing}\}$ and $\mathcal{C}_i^\downarrow = \{f \mid f \in \mathcal{C}_i \text{ and } f \text{ decreasing}\}$, $i = 1, 2, \dots, 5$.

The following relationships exist among these classes of functions: $\mathcal{C}_3 \subset \mathcal{C}_2 \subset \mathcal{C}_1$, $\mathcal{C}_3 \subset \mathcal{C}_5 \subset \mathcal{C}_1$.

We introduce the following partial orderings for random vectors.

Definition 3.4 *Let $X, Y \in \mathbb{R}^n$ be two random variables.*

(i) *X is stochastically smaller than Y ($X \leq_{st} Y$) if $E[f(X)] \leq E[f(Y)]$ for all increasing $f : \mathbb{R}^n \rightarrow \mathbb{R}$;*

(ii) *X is smaller than Y in the sense of convex ordering ($X \leq_{cx} Y$) if $E[f(X)] \leq E[f(Y)]$ for all convex $f : \mathbb{R}^n \rightarrow \mathbb{R}$;*

(iii) *X is smaller than Y in the sense of increasing (decreasing) convex ordering ($X \leq_{icx} (\leq_{dcx}) Y$) if $E[f(X)] \leq E[f(Y)]$ for all increasing (decreasing) convex $f : \mathbb{R}^n \rightarrow \mathbb{R}$;*

(iv) *X is smaller than Y in the sense of E_i ($E_i^\uparrow, E_i^\downarrow$) ordering ($X \leq_{E_i} (\leq_{E_i^\uparrow}, \leq_{E_i^\downarrow}) Y$) if $E[f(X)] \leq E[f(Y)]$ for all $f \in \mathcal{C}_i$ ($\mathcal{C}_i^\uparrow, \mathcal{C}_i^\downarrow$).*

These notions of the stochastic ordering on the random variables (or vectors) can be generalized to stochastic processes. Let X_t and Y_t be two stochastic processes, $t \in \mathbb{R}_+$. The process X_t is said to be smaller than Y_t in the sense of $\leq_{\mathcal{F}}$ for a certain class of functions \mathcal{F} , denoted by $\{X(t)\} \leq_{\mathcal{F}} \{Y(t)\}$, if

$$\forall n \in \mathbb{N}, \quad \forall (t_1, t_2, \dots, t_n) \in \mathbb{R}_+^n, \quad (X(t_1), X(t_2), \dots, X(t_n)) \leq_{\mathcal{F}} (Y(t_1), Y(t_2), \dots, Y(t_n)).$$

These could be the class of increasing functions (\leq_{st}), convex functions (\leq_{cx}), etc.

The E_4 ordering has received little attention in the area of control of queues because it does not exhibit useful closure properties. Although E_2 results exist in the literature regarding control of queues, they have not been obtained via sample path arguments. Instead, they are usually obtained via the following closure property.

Proposition 3.1 *Let $X^{(1)}, X^{(2)} \in \mathbb{R}^n$ and $Y^{(1)}, Y^{(2)} \in \mathbb{R}^n$ be pairs of independent random vectors with exchangeable components. If $X^{(i)} \leq_{E_2} (\leq_{E_2^\uparrow}) Y^{(i)}$, $i = 1, 2$, then $X^{(1)} + X^{(2)} \leq_{E_2} (\leq_{E_2^\uparrow}) Y^{(1)} + Y^{(2)}$.*

This is a useful property because \leq_{E_1} is not closed under convolution. We will show an example where an E_1^\uparrow ordering will be established via sample path arguments and the above closure property coupled with the fact that $\mathcal{C}_2 \subset \mathcal{C}_1$ will yield an E_2 ordering (Section 6).

Some of our results will require that random variables (r.v.'s) have distributions with increasing failure (hazard) rate (IFR) or increasing likelihood ratio (ILR). A r.v. $X \in \mathbb{R}_+$ has IFR (or has an

IFR distribution) if and only if

$$X_s \geq_{st} X_t, \quad 0 \leq s \leq t,$$

where X_t is the remaining life of X which has reached the age of t .

In order to define the ILR distribution, we first define the likelihood ratio ordering. Let $X, Y \in \mathbb{R}_+$ be two continuous nonnegative random variables with density functions f_X and f_Y respectively. A random variable X is smaller than a random variable Y in the sense of likelihood ratio ordering ($X \leq_{lr} Y$) if

$$f_Y(x)/f_X(x) \leq f_Y(y)/f_X(y), \quad 0 \leq x \leq y.$$

One of the properties of the likelihood ratio ordering is that it implies the strong stochastic ordering, i.e.,

$$X \leq_{lr} Y \Rightarrow X \leq_{st} Y.$$

The random variable $X \in \mathbb{R}_+$ is said to be increasing in likelihood ratio (ILR) (or has an ILR distribution) if

$$X_s \geq_{lr} X_t, \quad 0 \leq s \leq t,$$

where X_t is the remaining lifetime of X which has reached the age of t . A random variable is ILR if and only if its density function is log-concave (or, Polya frequency of order 2). One example is the r.v. with the gamma distribution with shape parameter greater than or equal to 1.

The likelihood ratio ordering can also be defined for discrete random variables. We say that $X \leq_{lr} Y$ if $P(Y = x)/P(X = x)$ increases in x . Last, an ILR random variable has an IFR distribution.

4 Forward Induction

In this section we will illustrate the use of forward induction in establishing the optimality of control policies for queueing systems. When it is applicable, forward induction provides a simple way of obtaining very strong sample path comparisons between an arbitrary policy and the optimal policy. In order to be applied, it is necessary to identify the hypothesized optimal policy and to derive equations that describe the behavior of the system under any policy over time. The proof then consists of a forward induction argument over the times of all distinct events establishing an ordering between the hypothesized optimal policy and the arbitrary policy. As we shall see, care must be taken in defining the probabilistic structure of the system. Typically there is more than one way of describing the evolution of a system but not all of them make possible a comparison between the optimal policy and the arbitrary policy.

We shall consider two applications of forward induction in this section, one establishing the optimality of the join the shortest queue (SQ) policy for a system of identical servers, each with its own queue; the other establishing the optimality of the μc rule for the multiclass G/M/1 queue. The first application will demonstrate some of the subtleties in choosing the appropriate probabilistic structure whereas the second application will demonstrate the application of sample path arguments, even in the case where the optimality result is in expectation.

Optimality of the Join the Shortest Queue Rule

As a first application, consider K identical servers, each with its own infinite capacity queue. We shall assume without loss of generality that the first customer enters an empty system. Customers arrive according to an arbitrary arrival process and, at the time of their arrival, are routed to one of the K queues. Customer service times are independently and identically exponentially distributed r.v.'s with mean $1/\mu$. Last, each queue is served by a non-idling scheduling policy. The objective is to determine the routing policy that minimizes some cost function of the joint queue lengths. We will show that the policy that routes a customer to the server with the shortest queue (SQ) of customers minimizes a rich set of cost functions.

We present a probabilistic description of the system.

Let $\Omega_n^k = \mathbb{R}_+$, $k \in \{0, 1, \dots, K\}$, and let $(\Omega, \mathcal{F}) = \left(\left(\prod_{n=1}^{\infty} \Omega_n^k \right)^{K+1}, \sigma \left(\left(\prod_{n=1}^{\infty} \Omega_n^k \right)^{K+1} \right) \right)$. For every $\omega = ((\omega_1^0, \omega_2^0, \dots), (\omega_1^1, \omega_2^1, \dots), \dots, (\omega_1^K, \omega_2^K, \dots)) \in \Omega$, define the coordinate processes

$$\phi_n^k(\omega) = \omega_n^k,$$

for all $n = 1, 2, \dots$, $k = 0, 1, \dots, K$.

Let P be any probability measure on (Ω, \mathcal{F}) such that $(\phi_n^1)_n, \dots, (\phi_n^K)_n$ are independent Poisson processes, each with rate μ , independent of $\phi^0 := (\phi_n^0)_n$, and such that $\phi_1^0 < \phi_2^0 < \dots$ a.s. and $\lim_{n \rightarrow \infty} \phi_n^0 = \infty$ a.s.

For every $n = 1, 2, \dots$, ϕ_n^0 will represent the arrival time of the n -th customer in the system. Note from the definition of P that the arrival process ϕ^0 is arbitrary as long as it is independent of the Poisson processes. The interpretation of the Poisson processes $(\phi_n^k)_k$, $k = 1, 2, \dots, K$, is given below.

It is useful to merge $\phi^0, \phi^1, \dots, \phi^K$ into a single ordered sequence of event times which will be referred to as $t_1 < t_2 < \dots$. Note that we can assume without loss of generality that at each time epoch t_n , there is only one event, as $(\phi_n^1)_n, \dots, (\phi_n^K)_n$ are independent Poisson processes. As it is necessary to identify the type of event that occurs at time t_n , let $e_n \in \{0, 1, \dots, K\}$ where $e_n = k$ if $t_n = \phi_l^k$ for some $l \leq n$, $k \in \{0, 1, \dots, K\}$. Here, $e_n = 0$ corresponds to an arrival.

Recall the definition of $D_i(x)$, $x_{[i]}$, \mathbf{e}_k given in Section 3.

Introduce the processes $(Q_n^\pi)_{n=0}^\infty$, $Q_n^\pi \in \mathbb{N}^K$, as

$$Q_0^\pi = 0 \tag{4.1}$$

$$Q_{n+1}^\pi = \mathbf{1}(e_n > 0) D_{e_n}(Q_n^\pi) + \mathbf{1}(e_n = 0) (\mathbf{e}_{U_n^\pi} + Q_n^\pi) \quad n = 1, 2, \dots, \tag{4.2}$$

where the admissible controls $(U_n^\pi)_n$ are in the form

$$U_n^\pi = \pi_n \left(\phi^0, (Q_i^\pi)_{i=1}^n, (U_i^\pi)_{i=1}^{n-1} \right), \quad \text{for } n = 1, 2, \dots \tag{4.3}$$

In (4.3), $(\pi_n)_n := \pi$ is any collection of measurable mappings $\pi_n : \mathbb{R}_+^N \times (\mathbb{N}^K)^n \times \{1, 2, \dots, K\}^{n-1}$ into $\{1, \dots, K\}$. Let Π be the collection of all such mappings. As mentioned in Section 2, we will identify the set of all admissible controls to the set Π .

It is worth observing from (4.2) that the value of U_n^π is irrelevant when $e_n > 0$. Therefore, although decisions are made at all times t_1, t_2, \dots , only those occurring at arrival times will influence the process $(Q_n^\pi)_n$.

From (4.1)-(4.3) we define the continuous-time process $(Q^\pi(t), t \geq 0)$ as

$$Q^\pi(t) = \sum_{n=1}^{\infty} Q^\pi(t_n) \mathbf{1}(t_n \leq t < t_{n+1}), \quad \forall t \geq 0. \quad (4.4)$$

We claim that the distribution of the process $(Q^\pi(t), t \geq 0)$ is the same as the distribution of the process $(\hat{Q}^\pi(t), t \geq 0)$, where $\hat{Q}^\pi(t)$ is the queue-length vector at time t in the original queueing system, when at time t_n the routing decision \hat{U}_n^π (if it applies, namely, if $e_n = 0$) is given by $\hat{U}_n^\pi = \pi_n(\phi^0, (\hat{Q}_i^\pi)_{i=1}^n, (\hat{U}_i^\pi)_{i=1}^{n-1})$, for every $n = 1, 2, \dots$

We shall not give a detailed proof of this equivalence. The proof relies on the fact that because of the exponential service times, $(\phi_n^k)_n$ in (4.2) is the *virtual departure process* of the k -th largest queue under any routing policy, for every $k = 1, 2, \dots, K$.

We shall denote by Π the set of all admissible routing policies. Let $\arg \min_{1 \leq j \leq K} x_j$ denotes the *smallest* index i in $\{1, 2, \dots, K\}$ such that $x_i \leq x_j$ for all $j = 1, 2, \dots, K$.

Let $\rho \in \Pi$ denote the policy that always routes customers to the shortest queue, namely,

$$\rho_n(\phi^0, (Q_i^\rho)_{i=1}^n, (U_i^\rho)_{i=1}^{n-1}) = \arg \min_{1 \leq k \leq K} Q_{k,n}^\rho, \quad (4.5)$$

where $Q_{k,n}^\rho$ is the k -th component of the vector Q_n^ρ .

We are now in the position to state and prove the following result.

Theorem 4.1 *The SQ policy ρ is the optimal policy out of the class of policies Π in the sense that it minimizes the joint queue length process in the sense of E_1^\uparrow , in other words,*

$$Q^\rho(t) \leq_{E_1^\uparrow} Q^\pi(t), \quad (4.6)$$

$\forall t \geq 0$ and $\pi \in \Pi$.

Proof. Let $\pi \in \Pi$ be an arbitrary routing policy. Condition on a single sample path $\omega \in \Omega$. We will establish that

$$Q^\rho(t) \prec_w Q^\pi(t), \quad t \geq 0, \quad (4.7)$$

by induction on the event times.

Basis step. Because the first customer always enters an empty system (4.6) trivially holds for $0 \leq t < t_1$.

Induction step. Assume that (4.7) holds for $0 \leq t < t_n$. Application of Lemma 3.2(i) with $i = K$ and $j \leq K$ in the case of an arrival and Lemma 3.2(ii) with $i = j = e_n$ in the case of a service completion yields (4.7) for $t = t_n$ (here $i = j$ when $e_n > 0$ because of the construction (4.2) that ensures that at time t_n a departure, if any, will occur in the k -th largest queue both in the system governed by ρ and in the system governed by π ; this is the justification for the construction (4.2)).

As there is no change in the process $(Q^\pi(t), t \geq 0)$ in the interval $[t_n, t_{n+1})$, (4.7) holds for $0 \leq t < t_{n+1}$.

It follows that $E[f(Q^\rho(t))] \leq E[f(Q^\pi(t))]$ for all $t \geq 0$ and for all increasing and Schur convex functions f which in turn yields (4.6), thus completing the proof. \blacksquare

This result was first established in Ephremides et al. [6].

Using the same type of arguments, this result has been extended in a number of different directions. For example, if the queues associated with the servers have finite capacities (which need not be identical), then the *join the Shortest Non-full Queue* (SNQ) policy has been shown to be optimal in the sense that

$$\begin{aligned} \{L^{SNQ}(t)\}_{t \geq 0} &\leq_{st} \{L^\pi(t)\}_{t \geq 0} \\ Q^{SNQ}(t) &\leq_{E_1^\dagger} Q^\pi(t), \end{aligned}$$

for all policies π that know the arrival times of all customers, all past queue length information and are required to route a customer to a non-full queue if one exists. Here $\{L^\pi(t)\}$ denotes the loss process for this system operating under policy π . Details of this and related finite capacity routing problems can be found in Sparaggis et al. [31] and in Towsley et al. [37].

In Ephremides et al. [6], besides stating and proving a result similar to (4.6), the authors also established the optimality of the *Round Robin* policy (RR) among the class of policies that use no information on the state of the queues, under the assumption of i.i.d. exponential service times, arbitrary arrivals and identical initial queue lengths at all of the servers. The trick here is to find a probabilistic description of the system appropriate for proving the result. In that paper, the authors described a system with an arbitrary arrival process and with a Poisson process to describe service completions. Each event in this process signals the *simultaneous* departure of customers *at all non-empty queues*. Because routing policies are unable to make any use of information regarding departures from servers, this probabilistic description does not affect the behavior of a policy. Although there is no equivalence between the joint queue lengths in the system defined above and the system of interest, the *marginal* queue lengths at each queue have the same statistical behavior in both systems.

Let $\tilde{Q}^\pi(t)$ denote the queue length vector for the system described above. It is proven in Ephremides

et al. [6] that the individual queue length statistics for this system are identical to those for the original system, hence

$$\tilde{Q}_k^\pi(t) =_{st} \hat{Q}_k^\pi(t), \quad 1 \leq k \leq s, t \geq 0.$$

The proof that the RR routing policy is the optimum policy consists of using the weak majorization arguments of Theorem 4.1 to establish

$$\tilde{Q}^{RR}(t) \leq_{E_1^\uparrow} \tilde{Q}^\pi(t), \quad \forall t \geq 0, \quad (4.8)$$

which has as its consequence,

$$\hat{Q}^{RR}(t) \leq_{E_3^\uparrow} \hat{Q}^\pi(t), \quad \forall t \geq 0. \quad (4.9)$$

This result has been generalized in two ways. First, RR has been shown to minimize the process of the number of losses at all times when all buffers have the same finite capacity and service times are exponentially distributed (see Sparaggis et al. [32]). Second, RR has been shown to be optimal in the case of infinite capacity buffers and i.i.d. service times with a distribution having increasing failure rate (see Liu and Towsley [21]). This last result of [21] has recently been extended to the case of i.i.d. service times with general distribution by Liu and Righter [18].

Similar ideas have been applied to the problem of scheduling n different classes of customers to a single server (see Towsley et al. [35] and Sparaggis et al. [30]). The latter study establishes a duality result between routing problems and scheduling problems.

Last, Chang [4] has used forward induction to establish the optimality of the balanced Bernoulli routing policy for identical servers and general service times out of the class of Markovian routing policies with positive correlation. This result is interesting because it is based on an $\leq_{E_5^\uparrow}$ ordering.

Optimality of the μc -rule

We provide an additional example of the usefulness of the forward induction argument. There are K classes of jobs to be processed on a single server, equipped with a buffer of infinite capacity. Each arriving job carries with it a random service requirement and leaves the system as soon as it completes service.

At every arrival and service completion epoch a decision-maker has to decide what class of customers is given access to the server until the next decision epoch. Within classes the service discipline is irrelevant.

The objective is to find an allocation policy of the server to the classes of customers that minimizes some weighted cost function, namely, a weighted sum of the remaining service requirements and a weighted sum of the expected number of jobs in the different classes.

We now give a more precise description of the model. Let $\Omega_k = \mathbb{R}_+^N \times \mathbb{R}_+^N$ for $k = 1, 2, \dots, K$, and let $(\Omega, \mathcal{F}) = \left(\prod_{k=1}^K \Omega_k, \sigma \left(\prod_{k=1}^K \Omega_k \right) \right)$. For every $\omega = ((\omega_1^0, \omega_1^1), \dots, (\omega_K^0, \omega_K^1))$ in Ω with

$\omega_k^m = (\omega_{k,1}^m, \omega_{k,2}^m, \dots)$ for $m = 0, 1$, define the coordinate processes

$$\phi_{k,n}^m(\omega) = \omega_{k,n}^m,$$

for all $n = 1, 2, \dots, k = 1, 2, \dots, K, m = 0, 1$. For every $n = 1, 2, \dots, k = 1, 2, \dots, K$, $\phi_{k,n}^0$ and $\phi_{k,n}^1$ are the arrival time and the service requirement, respectively, of the n -th customer of class k . Let $\phi = ((\phi_{k,n}^0, \phi_{k,n}^1)_{n=1}^\infty)_{k=1}^K$ be the identity mapping on Ω . Also define $\phi^0 = (\phi_{k,n}^0)_{k,n}$ to be the sequence of all arrival times.

In the following it will be convenient to think of this system as a system with K queues and a single server where customers of class k join queue k upon arrival, for $k = 1, 2, \dots, K$.

Let P be any probability measure on (Ω, \mathcal{F}) such that $0 \leq \phi_{k,1}^0 \leq \phi_{k,2}^0 \leq \dots$ a.s. and $\lim_{n \rightarrow \infty} \phi_{k,n}^0 = \infty$ a.s. for every $k = 1, 2, \dots$, and $0 < \phi_{k,n}^1 < \infty$ a.s. for all $n = 1, 2, \dots, k = 1, 2, \dots, K$.

Here, an admissible policy π , hereafter referred to as an admissible *scheduling* policy, is a collection of measurable mappings,

$$\pi_n : \mathbb{R}_+^N \times (\mathbb{N}^K)^n \times (\mathbb{R}_+^K)^n \times \{1, 2, \dots, K\}^{n-1} \rightarrow \{1, 2, \dots, K\},$$

such that $\pi_n(x, (q_{k,1})_{k=1}^K, \dots, (q_{k,n})_{k=1}^K, (v_{k,1})_{k=1}^K, \dots, (v_{k,n})_{k=1}^K, (u_i)_{i=1}^{n-1})$ is the allocation of the server at the n -th decision epoch given that $\phi = x$, the number of customers (resp. total unfinished work) in queue k ($k = 1, 2, \dots, K$) at the i -th decision epoch was $q_{k,i}$ (resp. $v_{k,i}$) for $i = 1, 2, \dots, n-1$, the current number of customers (resp. total unfinished work) in queue k is $q_{k,n}$ (resp. $v_{k,n}$), and the previous decisions are u_1, \dots, u_{n-1} . The decision k indicates that the server is allocated to queue k until the next decision epoch.

We shall denote by Π the set of all admissible scheduling policies. Let us assume that the first customer enters an empty system. Then, given a policy $\pi \in \Pi$, it is easy to construct stochastic processes representing the number of customers in queue k at time t (denoted as $Q_k^\pi(t)$), the unfinished work in queue k at time t (denoted as $V_k^\pi(t)$), the occurrence of the n -th decision epoch (denoted as τ_n^π) and the decision made at the n -th decision epoch (denoted as U_n^π). Observe, in particular, that

$$U_n^\pi = \pi_n \left(\phi, (Q_k^\pi(\tau_1^\pi))_{k=1}^K, \dots, (Q_k^\pi(\tau_n^\pi))_{k=1}^K, (V_k^\pi(\tau_1^\pi))_{k=1}^K, \dots, (V_k^\pi(\tau_n^\pi))_{k=1}^K, (U_i^\pi)_{i=1}^{n-1} \right),$$

for all $n = 1, 2, \dots$, so that at any decision epoch the decision-maker can use information regarding future arrivals and future service times, as well as service times of customers presently in the system, the number of customers (resp. unfinished work) in each queue, and the previous decisions. Also note that idling policies are allowed since the server may be allocated to an empty queue, and that the decision is irrelevant when the system is found empty just after a service completion. Clearly, the information available at a decision epoch exhibits some redundancy since past and current queue-lengths and workloads can be determined if one knows ϕ and the past decisions. This redundant information is introduced for future use (e.g., see (4.10)).

Consider the cost function $\sum_{k=1}^K r_k V_k^\pi(t)$ where the weights $(r_k)_{k=1}^K$ satisfy $r_1 \geq r_2 \geq \dots \geq r_K \geq 0$. Denote by $\gamma \in \Pi$ the policy that always allocates the server to the non-empty queue with the highest weight, namely,

$$U_n^\gamma = \gamma_n \left(\phi, (Q_k^\gamma(\tau_1^\gamma))_{k=1}^K, \dots, (Q_k^\gamma(\tau_n^\gamma))_{k=1}^K, (V_k^\gamma(\tau_1^\gamma))_{k=1}^K, \dots, (V_k^\gamma(\tau_n^\gamma))_{k=1}^K, (U_i^\gamma)_{i=1}^{n-1} \right) = k, \quad (4.10)$$

if $\sum_{i=1}^{k-1} Q_i^\pi(\tau_n^\pi) = 0$ and if $Q_k^\pi(\tau_n^\pi) > 0$.

Define $V^\pi(t) = (V_1^\pi(t), \dots, V_K^\pi(t))$ for all $t \geq 0$. We have the following result:

Theorem 4.2

$$\sum_{k=1}^K r_k V_k^\gamma(t) \leq \sum_{k=1}^K r_k V_k^\pi(t). \quad (4.11)$$

for all $t \geq 0$ and $\pi \in \Pi$.

Theorem 4.2 follows from the following two lemmas, the first of which is proven in Nain and Towsley [26]:

Lemma 4.1 Let (a_1, \dots, a_K) and (b_1, \dots, b_K) be \mathbb{R}_+^K -valued vectors such that $\sum_{i=1}^k a_i \leq \sum_{i=1}^k b_i$ for $k = 1, 2, \dots, K$. Then,

$$\sum_{k=1}^K r_k a_k \leq \sum_{k=1}^K r_k b_k. \quad (4.12)$$

Lemma 4.2

$$V^\gamma(t) \prec_{ps} V^\pi(t), \quad (4.13)$$

for all $t \geq 0$ and for all $\pi \in \Pi$.

Proof. Let $\pi \in \Pi$ be an arbitrary policy. Condition on a single path $\omega \in \Omega$. Let $(t_n)_{n=1}^\infty$, $0 \leq t_1 < t_2 < \dots$ be the sequence resulting from the superposition of the K arrival processes $(\phi_{n,k}^0)_{n,k}$, of the K departure processes in the system governed by policy γ , and of the K departure processes in the system governed by policy π (simultaneous events are allowed). Observe that $\lim_{n \rightarrow \infty} t_n = +\infty$ a.s. thanks to the definition of P .

The proof is by induction on the times of events.

Basis step. Because the first customer enters an empty system (4.13) trivially holds for $0 \leq t \leq t_1$.

Induction step. Assume that (4.13) holds for $0 \leq t \leq t_n$ and let us show that it is still true for $t_n < t \leq t_{n+1}$. There are two steps.

Step 1: $t_n < t < t_{n+1}$.

If $\sum_{i=1}^K V_i^\gamma(t_n) = 0$ then (4.13) clearly holds for $t_n < t < t_{n+1}$.

Consider the case that $\sum_{i=1}^K V_i^\gamma(t_n) > 0$ and let $l = \min\{i = 1, 2, \dots, K : V_i^\gamma(t_n) > 0\}$. By the definition of γ and of the sequence $(t_n)_{n=1}^\infty$ we have

$$(V_1^\gamma(t), \dots, V_K^\gamma(t)) = (0, \dots, 0, V_l^\gamma(t_n) - (t - t_n), V_{l+1}^\gamma(t_n), \dots, V_K^\gamma(t_n)). \quad (4.14)$$

For $k = 1, 2, \dots, l-1$, it is seen from (4.14) that

$$0 = \sum_{i=1}^k V_i^\gamma(t) \leq \sum_{i=1}^k V_i^\pi(t).$$

On the other hand, we have for $k = l, l+1, \dots, K$, cf. (4.14),

$$\sum_{i=1}^k V_i^\gamma(t) = \sum_{i=1}^k V_i^\gamma(t_n) - (t - t_n) \leq \sum_{i=1}^k V_i^\pi(t_n) - (t - t_n) \leq \sum_{i=1}^k V_i^\pi(t),$$

where the first inequality follows from the induction hypothesis.

Step 2: $t = t_{n+1}$.

Clearly, for $\delta = \gamma$ and $\delta = \pi$

$$V_i^\delta(t_{n+1}) = V_i^\delta(t_{n+1}^-) + \sum_{l=1}^{\infty} \phi_{l,i}^1 \mathbf{1}(\phi_{l,i}^0 = t_{n+1}),$$

for $i = 1, 2, \dots, K$. Inequality (4.13) at time t_{n+1} then follows from step 1. ■

We now show how the above sample path property can be used to establish the optimality of the μc rule for the G/M/1 queue. More precisely, we now address the minimization of the cost function $\sum_{k=1}^K c_k E[Q_k^\pi(t)]$ where c_k 's are arbitrary nonnegative constants. We restrict the analysis to policies in Π that do not know future service requirements. In other words, a policy $\pi = (\pi_n)_n$ is now an admissible scheduling policy if

$$U_n^\pi = \pi_n \left(\phi^0, (Q_k^\pi(\tau_1^\pi))_{k=1}^K, \dots, (Q_k^\pi(\tau_n^\pi))_{k=1}^K, (U_i^\pi)_{i=1}^{n-1} \right),$$

where π_n is a measurable mapping from $\pi_n : \mathbf{R}_+^N \times (\mathbf{N}^K)^n \times \{1, \dots, K\}^{n-1} \rightarrow \{1, 2, \dots, K\}$ for every $n = 1, 2, \dots$. Let Γ be the set of all such policies. Observe that the policy γ belongs to Γ (see (4.10)).

In addition to the conditions placed on the probability measure P , we shall assume that P must be such that the sequences of service requirements $(\phi_{1,n}^1)_n, \dots, (\phi_{K,n}^1)_n$ in queues $1, \dots, K$, respectively, are mutually independent i.i.d. sequences of r.v.'s such that $P(\phi_{k,n}^1 \leq x) = 1 - \exp(-\mu_k x)$ (exponential service requirements), further independent of the arrival process ϕ^0 .

Theorem 4.3 *Assume that $\mu_1 c_1 \geq \mu_2 c_2 \geq \dots \geq \mu_K c_K \geq 0$. Then, the scheduling policy in Γ that always allocates the server to the non-empty queue with highest $\mu_k c_k$ (the so-called μc -rule) minimizes the cost function $\sum_{k=1}^K c_k E[Q_k^\pi(t)]$ over all policies in Γ , for all $t \geq 0$.*

Proof. An argument based on Wald's lemma yields $E[Q_k^\pi(t)] = \mu_k E[V_k^\pi(t)]$ for $k = 1, 2, \dots, K$, $t \geq 0$, when $\pi \in \Gamma$. This combined with Theorem 4.2 (with $r_k = \mu_k c_k$, $k = 1, 2, \dots, K$) yields the desired result. ■

Theorem 4.3 says that the μc -rule is optimal out of the policies that may know future arrival times, past and current queue-lengths and past decisions, but not present and future service requirements in a multiclass G/M/1 queue. Proofs of the optimality of the μc -rule based on interchange arguments (see Section 6) have been obtained by Baras et al. [2], Buyukkoc et al. [3], and Nain [24].

5 Backward Induction

In this section, we illustrate the use of backward induction in establishing the optimality of the routing policy that routes arriving customers to the queue with the smallest workload for a system of identical servers, each with its own queue. The routing problem considered here differs from the one analyzed in Section 4 in that the decision-maker has now information on the workloads.

Consider a queueing system with $K \geq 1$ servers labeled $1, 2, \dots, K$, each server being equipped with its own infinite-capacity waiting queue. When a customer arrives in the system, it is immediately routed to one of the K queues. The routing decisions are made based on the current workload of the queues. The customer arrival times are arbitrary and the customer service times are i.i.d. r.v.'s which are independent of the arrival times.

We will assume that the system is initially empty. However, the result of this section holds for arbitrary initial workload, provided it is independent of the sequence of service times. Similarly, we will not deal with randomized policies for sake of simplicity. Nevertheless, using the technique presented in Nain and Towsley [26] one can easily extend our result to the larger class of policies that include randomized policies.

We give below a precise description of the underlying probability space for this model. Introduce $\Omega_n := \mathbf{R}_+ \times \mathbf{R}_+$ and let $(\Omega, \mathcal{F}) := (\prod_{n=1}^\infty \Omega_n, \sigma(\prod_{n=1}^\infty \Omega_n))$. For every $\omega = (\omega_0, \omega_1, \dots) \in \Omega$ with $\omega_n = (\omega_n^0, \omega_n^1) \in \Omega_n$, define the coordinate processes

$$\phi_n^m(\omega) = \omega_n^m, \tag{5.1}$$

for all $n = 0, 1, \dots$, $m = 1, 2$. Observe from (5.1) that the mapping $\phi : \Omega \rightarrow \Omega$ defined as $\phi = (\phi_1^0, \phi_1^1, \phi_2^0, \phi_2^1, \dots, \phi_n^0, \phi_n^1, \dots)$ is nothing but the identity mapping on Ω . In the following we shall use the shorthand $\phi = (\phi_n^m)_{n,m}$ and $\phi^m = (\phi_n^m)_n$.

The components of the input process ϕ will receive the following interpretation: for every $n =$

$1, 2, \dots, \phi_n^0$ and ϕ_n^1 are the the arrival time in the system and the service time, respectively, of the n -th customer.

Let P be any probability measure on (Ω, \mathcal{F}) such that the service times form a sequence of i.i.d. random variables with arbitrary distribution and that the arrival process is independent of the service times but otherwise arbitrary. We will assume without loss of generality that the system is empty when the first customer arrives.

Here, an admissible policy π , hereafter referred to as an admissible *routing* policy, is a collection of measurable mappings $(\pi_n)_{n=1}^\infty$,

$$\pi_n : \mathbb{R}_+^N \times (\mathbb{R}_+^K)^n \times \mathbb{R}_+^{n-1} \times \{1, 2, \dots, K\}^{n-1} \rightarrow \{1, 2, \dots, K\},$$

such that $\pi_n(x^0, (w_i)_{i=1}^n, (x_i^1)_{i=1}^{n-1}, (u_i)_{i=1}^{n-1})$ is the queue to which the n -th customer is routed to given that $\phi^0 = x^0$, the vector of the workloads at the i -th arrival epoch was w_i for $i = 1, 2, \dots, n-1$, the current vector of the workloads is w_n , the service times of the $(n-1)$ -st customers are x_1^1, \dots, x_{n-1}^1 , and the previous routing decisions were u_1, \dots, u_{n-1} . We shall denote by Π the set of all admissible routing policies.

Given $\pi \in \Pi$, we have the following evolution equations for the workload in the system:

$$W_{n+1,k}^\pi = \left(W_{n,k}^\pi + \phi_n^1 \mathbf{1}(U_n^\pi = k) - (\phi_{n+1}^0 - \phi_n^0) \right)^+ \quad k = 1, 2, \dots, K, \quad n = 1, 2, \dots, \quad (5.2)$$

with $(x)^+ := \max(0, x)$, where $\forall k, W_{1,k}^\pi = 0$, and

$$U_n^\pi = \pi_n \left(\phi^0, (W_i^\pi)_{i=1}^n, (\phi_i^1)_{i=1}^{n-1}, (U_i^\pi)_{i=1}^{n-1} \right), \quad (5.3)$$

for all $n = 1, 2, \dots$. In (5.2) $W_{n,k}^\pi$ is the workload in queue k *just before* the n -th arrival epoch, and U_n^π is the queue to which this customer is routed to, under policy $\pi \in \Pi$.

As in the example of the optimality of the μc -rule (cf. Section 4), it is worth noting that the information available at a decision epoch exhibits some redundancy since the workload in queue k just before the arrival of the n -th customer may clearly be determined from the arrival times ϕ^0 and from the sequence of service times $(\phi_i^1)_{i=1}^{n-1}$ as shown in (5.2). The reason for introducing explicitly the past and current workloads in the definition of an admissible policy will become apparent soon (see (5.5)).

Define $W_n^\pi = (W_1^\pi, \dots, W_K^\pi)$ to be the vector of the workloads just before the n -th arrival epoch. Rewriting (5.2) in vector form yields

$$W_{n+1}^\pi = \left(W_n^\pi + \mathbf{e}_{U_n^\pi} \phi_n^1 - \mathbf{i} (\phi_{n+1}^0 - \phi_n^0) \right)^+, \quad n = 1, 2, \dots, \quad (5.4)$$

with $\mathbf{i} := (1, 1, \dots, 1)$, where the notation $(\mathbf{v})^+$ stands for $((v_1)^+, \dots, (v_K)^+)$ if $\mathbf{v} = (v_1, \dots, v_K) \in \mathbb{R}^K$ (see the definition of $\mathbf{e}_k := (0, \dots, 0, 1, 0, \dots, 0)$ in Section 3). Recall that $\arg \min_{1 \leq j \leq K} x_j$ denotes the *smallest* index i in $\{1, 2, \dots, K\}$ such that $x_i \leq x_j$ for all $j = 1, 2, \dots, K$.

Let $\rho \in \Pi$ denote the Smallest Workload (SW) policy that routes every customer to the queue with the smallest workload, namely,

$$\rho_n \left(x^0, (v_i)_{i=1}^n, (x_i^1)_{i=1}^{n-1}, (u_i)_{i=1}^{n-1} \right) = \arg \min_{1 \leq k \leq K} v_{n,k}, \quad n = 1, 2, \dots, \quad (5.5)$$

where $x^0 \in \mathbb{R}_+^N$, $v_i = (v_{i,1}, \dots, v_{i,K}) \in \mathbb{R}_+^K$, $x_i^1 \in \mathbb{R}_+$ and $u_i \in \{1, 2, \dots, K\}$.

Observe from the definition of $\arg \min$ that if two queues or more have the smallest workload then the customer is routed to the queue with the smallest index among these queues.

The following result is due to Foss [8, 9, 11]. The proof we give here is based on that in Foss [11].

Theorem 5.1 *The policy ρ is optimal in the sense that*

$$W_n^\rho \leq_{\mathbb{E}_1^\uparrow} W_n^\pi, \quad (5.6)$$

for all $n = 1, 2, \dots$ and for all $\pi \in \Pi$.

Proof. Relation (5.6) trivially holds for $n = 1$ since $W_1^\rho = W_1^\pi = 0$. Let $n \geq 2$ be a fixed, arbitrary, integer.

Let $\pi \in \Pi$ be an arbitrary policy and define $\pi_{(l,m)} = (\pi_l, \pi_{l+1}, \dots, \pi_m)$ for all $1 \leq l \leq m$.

For any $k \in \{1, 2, \dots, n-2\}$, define $\pi^{(k)} \in \Pi$ to be the policy that makes the same decisions as π at the first k decision epochs and then switches to the SW rule at decision epochs $k+1, \dots, n-1$, namely,

$$\pi_{(1,n-1)}^{(k)} = \pi_{(1,k)} \rho_{(k+1,n-1)}. \quad (5.7)$$

What $\pi^{(k)}$ does after the $(n-1)$ -st decision epoch is irrelevant.

We shall use a backward induction argument to show that for $0 \leq k \leq n-2$,

$$W_n^{\pi^{(k)}} \leq_{\mathbb{E}_1^\uparrow} W_n^\pi. \quad (5.8)$$

Basis step: $k = n-2$.

Since $\pi_j^{(n-2)} \equiv \pi_j$ for $j = 1, 2, \dots, n-2$, it is clear that for all $\omega \in \Omega$

$$W_{n-1}^{\pi^{(n-2)}}(\omega) = W_{n-1}^\pi(\omega). \quad (5.9)$$

Since

$$U_{n-1}^{\pi^{(n-2)}} = U_{n-1}^\rho = \arg \min_{1 \leq j \leq K} w_{n-1}^\rho = \arg \min_{1 \leq j \leq K} w_{n-1}^{\pi^{(n-2)}},$$

we obtain from (5.4) that

$$\begin{aligned}
W_n^{\pi^{(n-2)}} &= \left(W_{n-1}^{\pi^{(n-2)}} + \mathbf{e}_{U_{n-1}^{\pi^{(n-2)}}} \phi_{n-1}^1 - \mathbf{i}(\phi_n^0 - \phi_{n-1}^0) \right)^+ \\
&= \left(W_{n-1}^{\pi} + \mathbf{e}_{U_{n-1}^{\rho}} \phi_{n-1}^1 - \mathbf{i}(\phi_n^0 - \phi_{n-1}^0) \right)^+ \\
&\prec \left(W_{n-1}^{\pi} + \mathbf{e}_{U_{n-1}^{\pi}} \phi_{n-1}^1 - \mathbf{i}(\phi_n^0 - \phi_{n-1}^0) \right)^+ \\
&= W_n^{\pi},
\end{aligned}$$

where the ordering \prec is defined in Section 3 (see Definition 3.2).

Therefore, relation (5.8) holds for $k = n - 2$.

Inductive step: assume that for a given l , $1 \leq l \leq n - 2$, (5.8) holds for all $l \leq k \leq n - 2$. We will prove that (5.8) still holds for $k = l - 1$.

Define Π_n to be the set of all finite collections of mappings $(\pi_m)_{m=1}^n$ such that (5.3) holds for $m = 1, 2, \dots, n$. We shall use the following lemma whose proof will be provided later.

Lemma 5.1 *Let $n \geq 2$ be an arbitrary positive integer, and let $\pi \in \Pi$ be an arbitrary policy. Then, for any arbitrary $k \in \{1, 2, \dots, n - 2\}$, there exists a policy $\gamma^{(k)} \in \Pi_{n-1}$ such that*

$$\gamma_{(1,k)}^{(k)} \equiv \pi_{(1,k-1)} \rho_k, \quad (5.10)$$

and

$$W_n^{\gamma^{(k)}} \leq_{E_1^\dagger} W_n^{\pi^{(k)}}. \quad (5.11)$$

Thus, by applying Lemma 5.1 to $\pi^{(l)}$ entails the existence of a policy $\gamma^{(l)} \in \Pi_{n-1}$ such that

$$\gamma_{(1,l)}^{(l)} \equiv \pi_{(1,l-1)}^{(l)} \rho_l \equiv \pi_{(1,l-1)} \rho_l,$$

and

$$W_n^{\gamma^{(l)}} \leq_{E_1^\dagger} W_n^{\pi^{(l)}}. \quad (5.12)$$

Applying the inductive assumption to policy $\gamma^{(l)}$ implies that there exists a policy $\nu^{(l)} \in \Pi_{n-1}$ such that $\nu^{(l)} \equiv \gamma_{(1,l)}^{(l)} \rho_{(l+1,n-1)}$ and that

$$W_n^{\nu^{(l)}} \leq_{E_1^\dagger} W_n^{\gamma^{(l)}}. \quad (5.13)$$

Since

$$\nu^{(l)} \equiv \gamma_{(1,l)}^{(l)} \rho_{(l+1,n-1)} \equiv \pi_{(1,l-1)} \rho_l \rho_{(l+1,n-1)} \equiv \pi^{(l-1)},$$

we obtain from relations (5.12) and (5.13) that

$$W_n^{\pi^{(l-1)}} = W_n^{\nu^{(l)}} \leq_{E_1^\dagger} W_n^{\gamma^{(l)}} \leq_{E_1^\dagger} W_n^{\pi^{(l)}} \leq_{E_1^\dagger} W_n^\pi,$$

where the last inequality uses again the inductive assumption. This completes the induction.

The proof is completed by noting that $\pi_{(1,n-1)}^{(0)} = \rho_{(1,n-1)}$. ■

Proof of Lemma 5.1. Let π be an arbitrary policy in Π and let $k \in \{1, 2, \dots, n-2\}$ be a fixed integer. Recall the definition of $\arg \min$ given earlier in this section. Let $r = U_k^\pi$ and $s = \arg \min_{1 \leq j \leq K} W_{kj}^\pi$. Here r is the queue where the k -th customer is routed to and s is the queue with the smallest workload at the k -decision epoch, in the system governed by π .

Define

$$\widehat{\Omega} = \left\{ \omega \in \Omega \mid W_{k,r(\omega)}^\pi - (\phi_{k+1}^0(\omega) - \phi_k^0(\omega)) < 0 \text{ or } r(\omega) = s(\omega) \right\}.$$

We construct a new input process $\psi = (\psi_n^m)_{n,m}$ on Ω such that on $\Omega - \widehat{\Omega}$ the service times of customers k and $k+1$ are interchanged. More precisely, for all $\omega \in \Omega$,

$$\begin{aligned} \psi^0(\omega) &= \phi^0(\omega) \\ \psi_i^1(\omega) &= \phi_i^1(\omega) \quad i \geq 1, \quad i \neq k, \quad i \neq k+1 \\ \psi_k^1(\omega) &= \phi_k^1(\omega) \mathbf{1}(\omega \in \widehat{\Omega}) + \phi_{k+1}^1(\omega) \mathbf{1}(\omega \in \Omega - \widehat{\Omega}) \\ \psi_{k+1}^1(\omega) &= \phi_{k+1}^1(\omega) \mathbf{1}(\omega \in \Omega - \widehat{\Omega}) + \phi_k^1(\omega) \mathbf{1}(\omega \in \widehat{\Omega}). \end{aligned}$$

The key observation is that both processes ϕ and ψ have the same distribution. The proof, given in the Appendix, relies on the fact that policies in Π do not use information on future service times, including the service time of arriving customers.

We now define a new policy $\gamma^{(k)} \in \Pi_{n-1}$ as follows. First, let

$$\gamma_i^{(k)} \equiv \pi_i^{(k)}, \quad \text{for } i = 1, 2, \dots, k-1 \tag{5.14}$$

$$\gamma_i^{(k)} \equiv \rho_i. \tag{5.15}$$

so that $\gamma^{(k)}$ follows π during the $(k-1)$ -st decision epochs and follows the SW rule at the k -th decision epoch.

Since $\phi^0 \equiv \psi^0$ and $\phi_i^1 \equiv \psi_i^1$ for $i = 1, 2, \dots, k-1$, it is worth observing from (5.14) and (5.15) that for all $\omega \in \Omega$

$$W_i^{\gamma^{(k)}}(\psi(\omega)) = W_i^\pi(\phi(\omega)), \quad i = 1, 2, \dots, k. \tag{5.16}$$

We now complete the definition of policy $\gamma^{(k)}$. Define the r.v.'s $(U_j)_{j=k+1}^{n-1}$ as

$$U_{k+1}(\psi(\omega)) = \begin{cases} U_{k+1}^{\pi^{(k)}}(\omega) & \text{if } r(\omega) = s(\omega) \\ r(\omega) & \text{otherwise,} \end{cases} \quad (5.17)$$

and, for $j = k + 2, \dots, n - 1$,

$$U_j(\psi(\omega)) = \begin{cases} U_j^{\pi^{(k)}}(\omega), & \text{if } r(\omega) = s(\omega) \text{ or if} \\ & W_{ks}^{\gamma^{(k)}}(\psi(\omega)) \geq \psi_{k+1}^0(\omega) - \psi_k^0(\omega) \\ r(\omega) \mathbf{1}(U_j^{\pi^{(k)}}(\omega) = s(\omega)) \\ + s(\omega) \mathbf{1}(U_j^{\pi^{(k)}}(\omega) = r(\omega)) \\ + U_j^{\pi^{(k)}}(\omega) \mathbf{1}(U_j^{\pi^{(k)}}(\omega) \notin \{r(\omega), s(\omega)\}), & \text{otherwise,} \end{cases} \quad (5.18)$$

for all $\omega \in \Omega$.

It is shown in the Appendix that there exist mappings $(\tilde{\gamma}_j)_{j=k+1}^{n-1}$, $\tilde{\gamma}_j : \mathbb{R}_+^N \times \mathbb{R}_+^{j-1} \rightarrow \{1, 2, \dots, K\}$ such that

$$U_j = \tilde{\gamma}_j(\psi^0, (\psi_i^1)_{i=1}^{j-1}), \quad j = k + 1, 2, \dots, n - 1. \quad (5.19)$$

The definition of $\gamma^{(k)}$ is then completed by letting $\gamma_j^{(k)} = \tilde{\gamma}_j$ for $j = k + 1, \dots, n - 1$.

We now compare $W_i^{\gamma^{(k)}}(\psi(\omega))$ to $W_i^{\pi^{(k)}}(\phi(\omega))$ for $i = 1, 2, \dots, n$. We have already observed (see (5.16)) that $W_i^{\gamma^{(k)}}(\psi(\omega)) = W_i^{\pi^{(k)}}(\phi(\omega))$ for $i = 1, 2, \dots, k$. On the other hand, it is easily seen from the definition of $\gamma^{(k)}$ and ψ that for all $\omega \in \Omega$,

$$W_{ij}^{\gamma^{(k)}}(\psi(\omega)) = W_{ij}^{\pi^{(k)}}(\phi(\omega)) \quad i = k + 1, \dots, n, \quad j \in \{1, 2, \dots, K\} - \{s(\omega), r(\omega)\}. \quad (5.20)$$

Consider now the workloads in queues r and s . Let $\Omega' = \{\omega \mid r(\omega) \neq s(\omega)\}$. It is clear that for all $\omega \in \Omega - \Omega'$, $W_i^{\gamma^{(k)}}(\psi(\omega)) = W_i^{\pi^{(k)}}(\phi(\omega))$ for $i = k + 1, \dots, n$. Therefore, on $\Omega - \Omega'$, the workloads are identical in both systems (one by $\pi^{(k)}$ with input process ϕ , the other by $\gamma^{(k)}$ with input process ψ):

$$W_i^{\gamma^{(k)}}(\psi(\omega)) = W_i^{\pi^{(k)}}(\phi(\omega)), \quad \omega \in \Omega - \Omega', \quad i = 1, 2, \dots, n. \quad (5.21)$$

Consider now Ω' . Using the evolution equation (5.2) together with the definition of $\gamma^{(k)}$ (cf. (5.15) and (5.17)) and ψ , we obtain that on $\widehat{\Omega} \cap \Omega'$,

$$W_{k+2, s(\omega)}^{\gamma^{(k)}}(\psi(\omega)) = \left(\left(W_{k, s(\omega)}^{\gamma^{(k)}}(\psi(\omega)) + \psi_k^1(\omega) - (\psi_{k+1}^0(\omega) - \psi_k^0(\omega)) \right)^+ - (\psi_{k+2}^0(\omega) - \psi_{k+1}^0(\omega)) \right)^+$$

$$\begin{aligned}
&= \left(\left(W_{k,s(\omega)}^{\pi^{(k)}}(\phi(\omega)) + \phi_k^1(\omega) - (\phi_{k+1}^0(\omega) - \phi_k^0(\omega)) \right)^+ - (\phi_{k+2}^0(\omega) - \phi_{k+1}^0(\omega)) \right)^+ \\
&\leq \left(\left(W_{k,r(\omega)}^{\pi^{(k)}}(\phi(\omega)) + \phi_k^1(\omega) - (\phi_{k+1}^0(\omega) - \phi_k^0(\omega)) \right)^+ - (\phi_{k+2}^0(\omega) - \phi_{k+1}^0(\omega)) \right)^+ \\
&= W_{k+2,r(\omega)}^{\pi^{(k)}}(\phi(\omega)), \tag{5.22}
\end{aligned}$$

and

$$\begin{aligned}
W_{k+2,r(\omega)}^{\gamma^{(k)}}(\psi(\omega)) &= (\psi_{k+1}^1(\omega) - (\psi_{k+2}^0(\omega) - \psi_{k+1}^0(\omega)))^+ \\
&= (\phi_{k+1}^1(\omega) - (\phi_{k+2}^0(\omega) - \phi_{k+1}^0(\omega)))^+ \\
&= W_{k+2,s(\omega)}^{\pi^{(k)}}(\phi(\omega)). \tag{5.23}
\end{aligned}$$

Similarly, on $(\Omega - \widehat{\Omega}) \cap \Omega'$,

$$\begin{aligned}
W_{k+2,s(\omega)}^{\gamma^{(k)}}(\psi(\omega)) &= \left(\left(W_{k,s(\omega)}^{\gamma^{(k)}}(\psi(\omega)) + \psi_k^1(\omega) - (\psi_{k+1}^0(\omega) - \psi_k^0(\omega)) \right)^+ - (\psi_{k+2}^0(\omega) - \psi_{k+1}^0(\omega)) \right)^+ \\
&= \left(\left(W_{k,s(\omega)}^{\pi^{(k)}}(\phi(\omega)) + \phi_{k+1}^1(\omega) - (\phi_{k+1}^0(\omega) - \phi_k^0(\omega)) \right)^+ - (\phi_{k+2}^0(\omega) - \phi_{k+1}^0(\omega)) \right)^+ \\
&\leq \left(\left(W_{k,s(\omega)}^{\pi^{(k)}}(\phi(\omega)) - (\phi_{k+1}^k(\omega) - \phi_k^0(\omega)) \right)^+ + \phi_{k+1}^1(\omega) - (\phi_{k+2}^0(\omega) - \phi_{k+1}^0(\omega)) \right)^+ \\
&= W_{k+2,s(\omega)}^{\pi^{(k)}}(\phi(\omega)), \tag{5.24}
\end{aligned}$$

since $(x + y)^+ \leq x^+ + y$ if $y \geq 0$, and

$$\begin{aligned}
W_{k+2,r(\omega)}^{\gamma^{(k)}}(\psi(\omega)) &= \left(\left(W_{k,r(\omega)}^{\gamma^{(k)}}(\psi(\omega)) - (\psi_{k+1}^0(\omega) - \psi_k^0(\omega)) \right)^+ + \psi_{k+1}^1(\omega) - (\psi_{k+2}^0(\omega) - \psi_{k+1}^0(\omega)) \right)^+ \\
&= \left(\left(W_{k,r(\omega)}^{\pi^{(k)}}(\phi(\omega)) - (\phi_{k+1}^0(\omega) - \phi_k^0(\omega)) \right)^+ + \phi_k^1(\omega) - (\phi_{k+2}^0(\omega) - \phi_{k+1}^0(\omega)) \right)^+ \\
&= \left(W_{k,r(\omega)}^{\pi^{(k)}}(\phi(\omega)) - (\phi_{k+1}^0(\omega) - \phi_k^0(\omega)) + \phi_k^1(\omega) - (\phi_{k+2}^0(\omega) - \phi_{k+1}^0(\omega)) \right)^+ \\
&= \left(\left(W_{k,r(\omega)}^{\pi^{(k)}}(\phi(\omega)) + \phi_k^1(\omega) - (\phi_{k+1}^0(\omega) - \phi_k^0(\omega)) \right)^+ - (\phi_{k+2}^0(\omega) - \phi_{k+1}^0(\omega)) \right)^+ \\
&= W_{k+2,r(\omega)}^{\pi^{(k)}}(\phi(\omega)), \tag{5.25}
\end{aligned}$$

where the third and fourth equalities follow from the definition of the set $\Omega - \widehat{\Omega}$.

Using now (5.21) and (5.22)-(5.25), together with the definition of $\gamma_i^{(k)}$ when $i \geq k + 2$ (cf. (5.18)), we obtain for all $k + 2 \leq i \leq n$ that if $W_{k,r(\omega)}^{\gamma^{(k)}}(\phi(\omega)) \geq \phi_{k+1}^0(\omega) - \phi_k^0(\omega)$ or $r(\omega) = s(\omega)$, then

$$W_{i,s(\omega)}^{\gamma^{(k)}}(\psi(\omega)) \leq W_{i,s(\omega)}^{\pi^{(k)}}(\phi(\omega)) \tag{5.26}$$

$$W_{i,r(\omega)}^{\gamma^{(k)}}(\psi(\omega)) \leq W_{i,r(\omega)}^{\pi^{(k)}}(\phi(\omega)), \tag{5.27}$$

otherwise

$$W_{i,s(\omega)}^{\gamma^{(k)}}(\psi(\omega)) \leq W_{i,r(\omega)}^{\pi^{(k)}}(\phi(\omega)) \quad (5.28)$$

$$W_{i,r(\omega)}^{\gamma^{(k)}}(\psi(\omega)) \leq W_{i,s(\omega)}^{\pi^{(k)}}(\phi(\omega)). \quad (5.29)$$

Therefore, cf. (5.20), (5.26)-(5.29),

$$W_n^{\gamma^{(k)}}(\psi(\omega)) \prec_w W_n^{\pi^{(k)}}(\phi(\omega)), \quad \forall \omega \in \Omega. \quad (5.30)$$

Relation (5.30) together with the equivalence in law of the processes ϕ and ψ and the property that ϕ is the identity mapping on Ω readily imply (5.11). \blacksquare

Note that if the waiting buffers of the different queues in the above routing problem is considered to be merged together, then the SW policy is identical to the FCFS discipline. Related results are discussed in the next section.

6 Interchange Arguments

In this section, we illustrate the use of interchange arguments by revisiting two optimization problems: in the first one, we discuss the extremal properties of the FCFS and LCFS service discipline with respect to customer response times; the second one is the well-known Klimov's problem.

Extremal Properties of FCFS and LCFS

We consider the effect that non-idling, non-preemptive scheduling policies have on customer response times. In particular, we illustrate how interchange arguments can be used to show that FCFS minimizes, in a sense to be made precise, the variability in customer response times. A similar interchange argument can be used to show the reverse is true with LCFS.

Let $\Omega_n = \mathbb{R}_+ \times \mathbb{R}_+$ and $(\Omega, \mathcal{F}) = (\prod_{n=1}^N \Omega_n, \sigma(\prod_{n=1}^N \Omega_n))$. As before, for every $\omega = (\omega_1, \dots, \omega_N) \in \Omega$ with $\omega_n = (\omega_n^0, \omega_n^1)$, we define the coordinate processes $\phi_n^k(\omega) = \omega_n^k$, $n = 1, 2, \dots$, $k = 0, 1$. Here ϕ_n^0 and ϕ_n^1 are the arrival and service times, respectively, of the n -th customer.

Let P be any probability measure on (Ω, \mathcal{F}) such that the service times form a sequence of i.i.d. random variables having arbitrary distribution and the arrival process is independent of the service times but otherwise arbitrary. Note that this describes a system in which there are exactly N arrivals.

Introduce the stochastic processes

$$I_n^\pi = \sum_{k=1}^N k \mathbf{1}(U_k^\pi = n)$$

$$\begin{aligned}
t_n^\pi &= \max\{t_{n-1}^\pi + s_{n-1}, \phi_n^0\} \\
c_n^\pi &= t_n^\pi + \phi_{U_n^\pi}^1 \\
s_n^\pi &= \phi_{U_n^\pi}^1 \\
\mathcal{S}_n^\pi &= \mathcal{S}_{n-1}^\pi - \{\pi_{n-1}\} + \{i \mid t_{n-1}^\pi < \phi_i^0 \leq t_n^\pi\} \\
r_n^\pi &= t_{I_n^\pi}^\pi + \phi_n^1 - \phi_n^0,
\end{aligned}$$

for $n = 1, 2, \dots, N$, where the admissible controls $(U_n^\pi)_n$ are given by

$$U_n^\pi = \pi_n(\phi^0, (s_i)_{i=1}^{n-1}, \mathcal{S}_n^\pi, (U_i^\pi)_{i=1}^{n-1}) \in \mathcal{S}_n. \quad (6.1)$$

In (6.1), $(\pi_n)_n := \pi$ is any collection of measurable mappings $\pi_n : \mathbf{R}_+^N \times \mathbf{R}_+^{n-1} \times \mathcal{P}(N) \times (\mathbf{N}^K)^n$ into $\{1, 2, \dots, N\}$, where $\mathcal{P}(N)$ is the set of all nonempty subsets of $\{1, 2, \dots, N\}$. Let Π be the collection of all such mappings. As usual we will identify the set of all admissible controls to the set Π .

These processes receive the following interpretation: under policy π , U_n^π is the identity of the n -th scheduled customer, I_n^π is the position in which customer n is scheduled, t_n^π is the time of the n -th decision, c_n^π is the n -th departure time, s_n^π is the service time of the n -th scheduled customer, \mathcal{S}_n^π the set of customers in the queue at the n -th decision epoch (note that $U_l^\pi \notin \mathcal{S}_n^\pi$ for $l = 1, 2, \dots, n-1$), and r_n^π is the response time of the n -th customer.

We will use R_N^π to denote $(r_1^\pi, \dots, r_N^\pi)$ and C_N^π to denote $(c_1^\pi, \dots, c_N^\pi)$.

Let $\rho \in \Pi$ denote the FCFS policy and let $\gamma \in \Pi$ be the LCFS policy. They are defined as

$$\begin{aligned}
\rho_n(\phi^0, (s_i)_{i=1}^{n-1}, \mathcal{S}, (\rho_i)_{i=1}^{n-1}) &= \min\{i : i \in \mathcal{S}\} \\
\gamma_n(\phi^0, (s_i)_{i=1}^{n-1}, \mathcal{S}, (\rho_i)_{i=1}^{n-1}) &= \max\{i : i \in \mathcal{S}\}, \quad \forall \mathcal{S} \in \mathcal{P}(N).
\end{aligned}$$

We have the following result

Theorem 6.1 *Policies ρ and γ are extremal policies within the class Π in the sense that for all $\pi \in \Pi$, $N = 1, 2, \dots$,*

$$R_N^\rho \leq_{E_1} R_N^\pi \leq_{E_1} R_N^\gamma. \quad (6.2)$$

Proof. We focus on the FCFS policy, ρ . Let $\pi \in \Pi$ be an arbitrary policy. We construct a new process $\psi = (\psi_n^m)_{m,n}$ on Ω as follows,

$$\begin{aligned}
\psi^0 &\equiv \phi^0 \\
\psi_{U_n^\pi}^1 &\equiv \phi_n^1, \quad n = 1, \dots, N.
\end{aligned}$$

This corresponds to an interchange of service times among customers so that, under FCFS, the n -th customer receives the same service as the customer corresponding to the n -th departure under π .

Using similar arguments as in the proof of Lemma 5.1, it can be shown that ϕ and ψ are equivalent. This is a consequence of the fact that service times form an i.i.d. sequence independent of the arrival times and the fact that admissible policies in Π do not use information on service times of customers that have not been served.

It follows from the construction of ψ that $c^\rho(\psi(\omega)) = c^\pi(\omega)$ for all $\omega \in \Omega$. Condition on $\omega \in \Omega$. Define the permutation γ on $\{1, \dots, N\}$ as $\gamma(n) = U_n^\pi$, $n = 1, \dots, N$. An application of corollary 3.1 yields

$$\begin{aligned} R_N^\rho(\omega) &= c_N^\pi(\omega) - \phi^0(\omega) \\ &\prec (c_N^\pi(\omega))_\gamma - \phi^0(\omega) \\ &= R_N^\pi(\omega). \end{aligned}$$

Removing the conditioning on ω and using the equivalence in distribution of ϕ and ψ yields the desired result. The proof of the relation for LCFS proceeds in a similar vein relying on Lemma 3.1. ■

Observe that the assertions of Theorem 6.1 can also be shown by forward induction. In this case, the probability space should be constructed in such a way that the n -th service time is independent of the identity of the customer. The interested reader is invited to write a complete proof.

Slightly weaker results hold for the $G/GI/s$ queue. Let w_n^π denote the time that the n -th arriving customer spends in the queue, prior to service, under policy $\pi \in \Pi$, and let $W_N^\pi = (w_1^\pi, \dots, w_N^\pi)$. The proof of the previous theorem can be adapted, with little modification to prove

$$W_N^\rho \leq_{E_1} W_N^\pi \leq_{E_1} W_N^\gamma, \quad \forall \pi \in \Pi. \quad (6.3)$$

The following weaker result has also been established for response time vectors (see Liu and Towsley [19]).

$$R_N^\rho \leq_{E_2} R_N^\pi \leq_{E_2} R_N^\gamma, \quad \forall \pi \in \Pi. \quad (6.4)$$

Proof of (6.4). We have

$$R_N^\pi = (w_1^\pi + s_1^\pi, \dots, w_N^\pi + s_N^\pi),$$

where s_1^π, \dots, s_N^π are exchangeable rv's. If $f : \mathbb{R}^N \rightarrow \mathbb{R}$ is a symmetric convex function, then $g : \mathbb{R}^N \rightarrow \mathbb{R}$ defined as

$$g(x_1, \dots, x_N) = E[f(x_1 + s_1^\pi, \dots, x_N + s_N^\pi)],$$

is also symmetric convex (see Marshall and Olkin [22, Proposition B.5, p. 289]). Since any symmetric convex function is a Schur-convex function and

$$E[f(r_1^\pi, \dots, r_N^\pi)] = E[g(w_1^\pi + s_1^\pi, \dots, w_N^\pi + s_N^\pi)],$$

the relation (6.4) follows from (6.3). ■

The reason for presenting this argument is because it was not based on a sample path argument. We, (and others, see Chang [4]) have observed that the E_2 ordering is not amenable to such arguments. Instead, such an ordering is typically established by invoking the property that E_2 is closed under convolution (under suitable assumptions). This is the idea behind the proof given above.

In the case of the $G/D/s$ queue, the result on response times can be extended to an E_1 ordering using the argument presented in theorem 6.1 (see Towsley and Baccelli [34] and Liu and Towsley [19] for details). This is because customers depart in the same order as they begin service. A more intricate argument based on a sample path construction has been used to extend the result to the $G/ILR/s$ queue as well (see Liu and Towsley [20] for details).

The previous results can be extended to the stationary regime, provided that it exists. Recall that the sequence of r.v.'s $X_n \in \mathbb{R}$, $n \geq 1$ converge to the r.v. X for the class of Borel mappings $\mathcal{F} : \mathbb{R} \rightarrow \mathbb{R}$ in the Cesaro sense (see Feller [7, p. 249]) if

$$\lim_{N \rightarrow \infty} \frac{1}{N} \sum_{i=1}^N E[f(X_i)] = E[f(X)].$$

Lemma 6.1 *Let $\{X_i\}_{i=1}^{\infty}$ and $\{Y_i\}_{i=1}^{\infty}$ be two sequences of r.v.'s such that*

$$(X_1, \dots, X_N) \leq_{E_3} (Y_1, \dots, Y_N), \quad \forall N \geq 1.$$

If the sequences $\{X_i\}_{i=1}^{\infty}$ and $\{Y_i\}_{i=1}^{\infty}$ converge weakly to X and Y with respect to the class of convex functions in the Cesaro sense as n goes to ∞ , then

$$X \leq_{cx} Y.$$

Proof. See Liu and Towsley [19]. ■

This last lemma implies that the stochastic orderings \leq_{E_1} and \leq_{E_2} established for the vectors of waiting times and response times, respectively, reduce to the stochastic ordering \leq_{cx} on the corresponding stationary performance metrics W^π and R^π , respectively, provided that the weak convergence assumptions are satisfied. Results on stationary variables had been established previously in Vasicek [39], Shanthikumar and Sumita [29], Liu and Towsley [19]. They can also be obtained from Theorem 5.1.

If we consider the class of non-preemptive policies that allow idling, then the results pertaining to the optimality of FCFS policy hold [19], provided that E_1^\uparrow , E_2^\uparrow and icx replace E_1 , E_2 , and cx in the above results. If we consider the class of preemptive non-idling policies Π_p , then we have the following result, Liu and Towsley [19].

Theorem 6.2 *Consider a $G/M/s$ system with n arrivals. If the service times are independent of the arrival process, then, for all $N \geq 1$*

$$R_N^\rho \leq_{E_1} R_N^\pi \leq_{E_1} R_N^\gamma, \quad \forall \pi \in \Pi_p. \tag{6.5}$$

Proof. A sketch of the proof is provided here. The complete proof is found in Liu and Towsley [19]. First, we can restrict ourselves to policies that preempt only at arrival times and service completion times. This is a consequence of the memoryless property of the exponential distribution.

Second, we associate a Poisson process with each of the s servers. The events correspond to times at which service completions may occur. In particular, a service completion occurs if there is a customer assigned to that server; otherwise there is no service completion. A customer assigned to service at an idle server always receives, as its service time, the remaining time until the next event. The memoryless property guarantees that the received service time is an exponentially distributed r.v. with the proper parameter.

Now, fix the arrival times, the service completion events, pick an arbitrary policy π and focus on ρ (FCFS). We establish the following relation,

$$R_N^\rho \prec R_N^\pi.$$

If we consider the sequence of decision epochs (arrival and service completion times), an interchange argument can be used to show that the policy π^1 that differs from π in that it makes one less non-FCFS decision yields a response time vector that is majorized by the response time vector associated with π . This step can then be repeated until the FCFS policy is produced. Removal of the conditioning on arrival times and service events yields the desired result for FCFS. A similar argument yields the result for LCFS (γ). ■

In the case of the $G/IFR/1$ queue, combining arguments in Hirayama and Kijima [13] with those presented above yield

$$R_N^\rho(n) \leq_{E_1^\dagger} R_N^\pi(n), \quad \forall \pi \in \Sigma_p.$$

This does not extend to the $G/IFR/s$ system for $s > 1$ as it is easy to construct a counterexample in the case of deterministic service times, (see Liu and Towsley [19] for details).

Last, these results have been extended to in-forests Liu and Towsley [20] and to a class of parallel processing systems modeled by fork/join queueing networks (see Baccelli et al. [1]). Similar results have also been obtained for systems in which customers have deadlines (see Baccelli et al. [1] and Liu and Towsley [19, 20]).

Klimov's Problem

The following situation was considered in Klimov [16]. There are K infinite capacity FIFO queues. The service requirements are mutually independent r.v.'s and have p.d.f. $B_k(x)$ in queue k ($k = 1, 2, \dots, K$). Customers arrive according to an independent Poisson process with rate λ and are assigned to queue k with probability p_k ($\sum_{k=1}^K p_k = 1$). There is a *nonidling* single server that is allocated to one of the queues at a time, in a *nonpreemptive* way. In other words, a new allocation of the server to a queue may only occur at a departure epoch if the system is nonempty, or when a customer enters an empty system. Upon his service completion in queue j a customer is sent to

queue k with the probability p_{jk} , and leaves the system with the probability $p_{j0} = 1 - \sum_{k=1}^N p_{jk}$, independently of the state of the system.

We assume that (1) the routing matrix $P := (p_{jk})_{1 \leq j, k \leq K}$ is such that every customer eventually leaves the system (which implies that $I - P$ is invertible, where I is the identity matrix; cf. Klimov [16, Lemma 3]), (2) $\int_0^\infty x dB_k(x) := b_k < \infty$ and $\int_0^\infty x^2 dB_k(x) := b_k^{(2)} < \infty$ for $k = 1, 2, \dots, K$, and (3) $\lambda p (I - P)^{-1} b^T < 1$, where $p = (p_1, p_2, \dots, p_K)$ and $b = (b_1, b_2, \dots, b_K)^T$ (here $(\cdot)^T$ denotes transposition). It is shown in Klimov [16] that (3) is the stability condition under any nonidling service policy.

Let $c_k \geq 0$ ($k = 1, 2, \dots, K$) be the holding cost per unit of time and per customer in queue k . The objective is to find a server allocation policy that minimizes the long-run average expected cost incurred over an infinite horizon.

We now give a more precise description of the underlying probability space.

Introduce $\Omega_n := \mathbf{R}_+ \times \mathbf{N} \times (\{1, 2, \dots, K\} \times \mathbf{R}_+)^{\mathbf{N}}$ and let $(\Omega, \mathcal{F}) := (\prod_{n=1}^\infty \Omega_n, \sigma(\prod_{n=1}^\infty \Omega_n))$. For every $\omega = (\omega_1, \omega_2, \dots) \in \Omega$ with $\omega_n = (\omega_n^0, \omega_n^1, \dots) \in \Omega_n$, define the coordinate processes

$$\phi_n^m(\omega) = \omega_n^m, \quad (6.6)$$

for all $n = 1, 2, \dots$, $m = 0, 1, \dots$. Observe from (6.6) that the mapping $\phi : \Omega \rightarrow \Omega$ defined as $\phi(\omega) = (\phi_1^0(\omega), \phi_1^1(\omega), \dots, \phi_2^0(\omega), \phi_2^1(\omega), \dots, \phi_n^0(\omega), \phi_n^1(\omega), \dots)$ is the identity mapping on Ω . As usual, we shall use the shorthand $\phi = (\phi_n^m)_{m,n}$.

The components of the process ϕ will receive the following interpretation: for every $n = 1, 2, \dots$, ϕ_n^0 is the arrival time in the system of customer n , ϕ_n^1 is the total *number* of visits of customer n to a queue, ϕ_n^2 is the queue to which customer n is routed upon arrival in the system, ϕ_n^{2j+2} , $j = 1, 2, \dots, \phi_n^1 - 1$, is the queue to which customer n is routed upon completion of his j -th service, and ϕ_n^{2j+1} , $j = 1, 2, \dots, \phi_n^1$, is the service required by customer n at his j -th visit to a queue.

Let P be any probability measure on (Ω, \mathcal{F}) such that the aforementioned statistical assumptions are satisfied (Poisson arrivals, independent service times, etc.) and such that the first customer arrives at time 0 (i.e., $\phi_1^0 = 0$ a.s.).

A policy π is a collection of measurable mappings $(\pi_n)_n$, $\pi_n : (\mathbf{N}^K)^n \rightarrow \{1, 2, \dots, K\}$, such that for every $k = 1, 2, \dots, K$,

$$\pi_n(x_1, \dots, x_n) \neq k \quad \text{if} \quad x_{k,n} = 0, \quad (6.7)$$

where $x_i = (x_{1,i}, \dots, x_{K,i})$ for $i = 1, 2, \dots, n$. The left-hand side of (6.7) gives the index of the n -th queue to be visited by the server given that the current number of customers in queue k is $x_{k,n}$ for $k = 1, 2, \dots, K$, and the number of customers in queue k at the i -th decision epoch was $x_{k,i}$ for $k = 1, 2, \dots, K$, $i = 1, 2, \dots, n - 1$. The constraint in the right-hand side of (6.7) indicates that the server cannot be allocated to an empty queue (nonidling policy).

We say that a policy $\pi \in \Pi$ is stationary if there exists a measurable mapping $f : \mathbf{N}^K \rightarrow \{1, 2, \dots, K\}$ such that, for every $n = 1, 2, \dots$, $\pi_n(x_1, \dots, x_n) = f(x_n)$ for all $(x_1, \dots, x_n) \in (\mathbf{N}^K)^n$.

We will assume that the first customer enters an empty system. Thus, given a policy π in Π and an input sequence ϕ , it is easy to construct on (Ω, \mathcal{F}, P) the following stochastic processes:

- t_n^π , the time of the n -th decision, $n = 1, 2, \dots$;
- $Q_k^\pi(t)$, the number of customers in queue k at time $t \geq 0$ for $k = 1, 2, \dots, K$;
- $q_n^\pi := (q_{1,n}^\pi, \dots, q_{N,n}^\pi)$, where $q_{k,n}^\pi := Q_k^\pi(t_n^\pi)$ is the number of customers in queue k at the n -th decision epoch, for $n=1,2,\dots, k=1,2,\dots,K$.

Let Π be the set of all policies such that for every $\pi \in \Pi$ the process $\{(Q_1^\pi(t), \dots, Q_K^\pi(t)), t \geq 0\}$ is a regenerative process with respect to the times where a customer enters an empty system. For instance, any stationary policy belongs to Π . Also any policy that may only know the queue-length history in the *current* busy cycle and that follows the same policy in each busy cycle belongs to Π .

We want to find a policy in Π that minimizes the long-run average expected cost

$$H^\pi := \liminf_{T \rightarrow \infty} E \left[\frac{1}{T} \int_0^T \sum_{k=1}^K c_k Q_k^\pi(t) dt \right]. \quad (6.8)$$

Because of the statistical assumptions we have placed on the model, and because of the definition of the set Π it is easily seen from renewal theory that minimizing H^π is equivalent to minimizing the cost $J^\pi := E \left[\int_0^Z \sum_{k=1}^K c_k Q_k^\pi(t) dt \right]$, where Z is the duration of a busy period under any policy $\pi \in \Pi$. Note that Z does not depend on $\pi \in \Pi$ since the server cannot idle under policies in Π .

Also observe that $J^\pi < \infty$ for every $\pi \in \Pi$ since

$$J^\pi \leq \left(\max_{1 \leq k \leq K} c_k \right) \left(E[Z^2] \right)^{1/2} \left(E[M^2] \right)^{1/2}, \quad (6.9)$$

by Cauchy-Schwartz inequality, where M is the number of customers served during a busy period. The proof that $E[Z^2]$ and $E[M^2]$ are both finite is left to the reader (hint: compute the first two moments of the total service time required by a customer and apply Theorem 1 in Ghahramani and Wolff [12]).

Define $\mathbf{k} := \{1, 2, \dots, k\}$ and $\mathbf{k}^c := \{k+1, k+2, \dots, K\}$. Assume that the queues are numbered as follows: assign number 1 to the queue that maximizes the quantity

$$\frac{c_i - \sum_{j=1}^K p_{ij} c_j}{b_i}, \quad i = 1, 2, \dots, K.$$

Then, recursively for $k = 1, 2, \dots, K - 1$ assign the number $k + 1$ to the queue in \mathbf{k}^c that maximizes the quantity

$$\frac{c_i - \sum_{j \in \mathbf{k}^c} p_{ij}^{(k)} c_j}{b_i + T_i^{(k)}},$$

where \mathbf{k} is the set of nodes $\{1, 2, \dots, k\}$ in the new numbering and where $T_i^{(k)}$ is the expected sojourn of a job in \mathbf{k} having started at queue $i \in \mathbf{k}^c$. A closed-form expression for $T_i^{(k)}$ is given in Nain et al. [25].

The following result is due to Klimov [16]. The proof we give is a variant of the proof given in Nain et. al [25]. Similar arguments can be found in [10, 38, 42], where some generalizations of the model were considered in [10].

Theorem 6.3 *The policy $\pi^* \in \Pi$ that always assigns the server to the nonempty queue with the smallest number is optimal among all policies in Π .*

Proof. Let π be any policy in Π different from π^* . Since π^* is stationary, we shall substitute the notation $\pi_n^*(x_1, \dots, x_n)$ for $\pi^*(x_n)$ with a slight abuse of notation.

Define $\pi^l = (\pi_n^l)_n$ to be the policy that follows policy π for the first l steps and then π^* afterwards. In other words,

$$\pi_n^l(x_1, \dots, x_n) = \begin{cases} \pi_n(x_1, \dots, x_n) & \text{for } n = 1, 2, \dots, l \\ \pi_n^*(x_n) & \text{for } n = l + 1, l + 2, \dots, \end{cases} \quad (6.10)$$

for all $(x_1, \dots, x_n) \in (\mathbb{N}^K)^n$, $n = 1, 2, \dots$

Assume that

$$J(\pi^{l-1}) \leq J(\pi^l), \quad (6.11)$$

for $l = 1, 2, \dots$. Then, $J^{\pi^*} \leq J^\pi$.

Indeed, (6.11) together with the fact that $\pi^0 \equiv \pi^*$, yields

$$J^{\pi^*} \leq J^{\pi^l}, \quad (6.12)$$

for $l = 1, 2, \dots$. On the other hand, we see from the inequalities $\int_0^Z \sum_{k=1}^K c_k Q_k^{\pi^l}(t) dt \leq MZ \max_{1 \leq k \leq K} c_k$ a.s. (M has been defined earlier in this section), $E[Z^2] < \infty$, $E[M^2] < \infty$, and from $\lim_{l \rightarrow \infty} \pi^l = \pi$ that the bounded convergence theorem applies to J^{π^l} to yield $\lim_{l \rightarrow \infty} J^{\pi^l} = J^\pi$. Combining this result together with (6.12) gives $J^{\pi^*} \leq J^\pi$.

We now prove (6.11) which will conclude the proof.

Let $l \geq 1$ be an arbitrary integer. Assume that $\pi_l^l(x_1, \dots, x_l) \neq \pi^*(x_l)$ for some $(x_1, \dots, x_l) \in (\mathbf{N}^K)^l$ since otherwise (6.11) is trivially true from the definition of policy π^l .

On Ω define the r.v. θ to be the time when the l -th decision is made under policy π^l . Let $\hat{\Omega}$ be the subset of Ω such that $\theta(\omega) < Z(\omega)$ for every $\omega \in \hat{\Omega}$ and such that the nonempty queue with the smallest number *is not* selected by policy π^l at the l -th decision epoch for every path ω in $\hat{\Omega}$.

In other words,

$$\hat{\Omega} = \left\{ \omega \in \Omega : \theta(\omega) < Z(\omega) \text{ and } \pi_l^l \left(q_1^{\pi^l}(\omega), \dots, q_l^{\pi^l}(\omega) \right) \neq \min \left\{ k = 1, 2, \dots, K : q_{l,k}^{\pi^l}(\omega) \neq 0 \right\} \right\}.$$

On $\hat{\Omega}$ define two $\{1, 2, \dots, N\}$ -valued r.v.'s j_0 and i_0 , where j_0 is the nonempty queue with the smallest number at time θ , and i_0 is the queue that is selected by policy π^l at time θ . Observe that $j_0 < i_0$ from the definition of the set $\hat{\Omega}$ and from the definition of π^l . Let S_{i_0} (resp. S_{j_0}) be the service time of the customer at the head of queue i_0 (resp. j_0) at time θ , and call ϵ (resp. η) this customer.

On $\hat{\Omega}$ define the $[0, \infty)$ -valued r.v.'s

$$\begin{aligned} \rho &= \inf \left\{ s \geq S_{i_0} : \sum_{k=1}^{j_0-1} Q_k^{\pi^l}(\theta + s) = 0 \right\} \\ \sigma &= \inf \left\{ s \geq S_{j_0} : \sum_{k=1}^{j_0-1} Q_k^{\pi^l}(\theta + \rho + s) = 0 \right\}. \end{aligned}$$

With the above definitions we see that, on $\hat{\Omega}$, policy π^l serves ϵ in $(\theta, \theta + S_{i_0})$, empties queues in \mathbf{j}_0-1 in $(\theta + S_{i_0}, \theta + \rho)$, serves η in $(\theta + \rho, \theta + \rho + S_{j_0})$, and empties again queues in \mathbf{j}_0-1 in $(\theta + \rho + S_{i_0}, \theta + \rho + \sigma)$ (cf. Figure 1).

We introduce the following notation: for any finite sequence $(x_i)_i$ of nonnegative real numbers we shall denote by $x_{(i)}$ its i -th smallest element.

From the process ϕ we now construct a new process $\psi = (\psi_n^m)_{m,n}$ such that on $\hat{\Omega}$ the route and the successive service requirements along this route of a customer that enters a queue in \mathbf{j}_0-1 in $(\theta, \theta + \rho)$ and in $(\theta + \rho, \theta + \rho + \sigma)$, respectively, are interchanged. More precisely, we define from the sequence of arrival times $(\phi_n^0)_n$ a new sequence $(a_n)_n$ such that

$$a_n = \begin{cases} \phi_n^0 & \text{if } \phi_n^2 \notin \mathbf{j}_0-1 \text{ or if } \phi_n^0 \notin (\theta, \theta + \rho + \sigma) \\ \phi_n^0 + \sigma & \text{if } \phi_n^2 \in \mathbf{j}_0-1 \text{ and if } \phi_n^0 \in (\theta, \theta + \rho) \\ \phi_n^0 - \rho & \text{if } \phi_n^2 \in \mathbf{j}_0-1 \text{ and if } \phi_n^0 \in (\theta + \rho, \theta + \rho + \sigma). \end{cases}$$

On $\hat{\Omega}$, the process ψ is then obtained by letting

$$(\psi_n^m)_{m,n} = (a_{(n)}, \phi_{\alpha_n}^1, \phi_{\alpha_n}^2, \dots), \quad (6.13)$$

for every $k = 1, 2, \dots$, where α_n is a mapping from \mathbf{N} into \mathbf{N} that gives the position of the n -th element of the ordered sequence $(a_{(n)})$ in the original sequence $(a_n)_n$ (for instance, if $a_{(1)} = a_3$, $a_{(2)} = a_2$ and $a_{(3)} = a_1$ then $\alpha_1 = 3$, $\alpha_2 = 2$ and $\alpha_3 = 1$). We shall further assume that $\psi \equiv \phi$ on $\Omega - \hat{\Omega}$.

The key observation is that both processes ϕ and ψ have the same distribution. This result is a direct consequence of the statistical assumptions we have placed on the model and of the fact that policies in Π do not know the arrival times, service times and routes of the customers (see Section 5 and the Appendix for the proof of a similar assertion).

We now introduce a new policy $\pi^{(1)}$ such that

$$\pi_n^{(1)} \equiv \pi_n \quad \text{for } n = 1, 2, \dots, l-1 \quad (6.14)$$

$$\pi_l^{(1)} \equiv \pi^*, \quad (6.15)$$

and

$$\pi_n^{(1)} \left(q_1^{\pi^{(1)}}(\psi(\omega)), \dots, q_n^{\pi^{(1)}}(\psi(\omega)) \right) = \begin{cases} i_0(\omega) & \text{if } \sum_{k=1}^{j_0(\omega)-1} q_{r,n}^{\pi^{(1)}}(\psi(\omega)) \neq 0 \\ & \text{for } l+1 \leq r \leq n-1, \\ & \sum_{k=1}^{j_0(\omega)-1} q_{n,i}^{\pi^{(1)}}(\psi(\omega)) = 0, \\ & \text{and } \omega \in \hat{\Omega} \\ \pi^* \left(q_n^{\pi^{(1)}}(\psi(\omega)) \right) & \text{otherwise,} \end{cases}$$


for $n \geq l+1$ and $\omega \in \Omega$.

Observe that $\pi_n^{(1)}$ is well-defined as a function of $q_1^{\pi^{(1)}}(\psi(\omega)), \dots, q_n^{\pi^{(1)}}(\psi(\omega))$ when $n \geq l+1$ and $\omega \in \hat{\Omega}$ since $i_0(\omega) = \pi_l \left(q_1^{\pi^{(1)}}(\psi(\omega)), \dots, q_l^{\pi^{(1)}}(\psi(\omega)) \right)$ and $j_0(\omega) = \min\{s = 1, 2, \dots, K : q_{l,s}^{\pi^{(1)}}(\psi(\omega)) \neq 0\}$. This clearly follows from (6.14)-(6.15) together with the definition of ψ . This observation, together with the fact that $\pi^{(1)}$ is nonidling, ensures that $\pi^{(1)} \in \Pi$.

In words, policy $\pi^{(1)}$ follows π at steps $1, 2, \dots, l-1$ and then follows π^* afterwards except that it applies π_l the first time after the l -th decision epoch when queues in $\mathbf{j}_0 - \mathbf{1}$ become empty. To see more precisely how $\pi^{(1)}$ behaves, consider the input sequence $\psi(\omega)$. If $\omega \in \Omega - \hat{\Omega}$, then on the input sequence $\psi(\omega)$ policy $\pi^{(1)}$ behaves like policy π^l does on the input sequence $\phi(\omega) = \omega$.

Assume now that $\omega \in \hat{\Omega}$. Then, (cf. Figure 1) on the input sequence $\psi(\omega)$, $\pi^{(1)}$ will follow π in $(0, \theta(\omega))$, will serve η in $(\theta(\omega), \theta(\omega) + S_{j_0}(\omega))$, will empty queues in $\mathbf{j}_0 - \mathbf{1}$ in $(\theta(\omega) + S_{j_0}(\omega), \theta(\omega) + \sigma(\omega))$, will serve ϵ in $(\theta(\omega) + \sigma(\omega), \theta(\omega) + \sigma(\omega) + S_{i_0}(\omega))$, and will finally resume policy π^* afterwards

Policy π^l on $\hat{\Omega}$



Policy $\pi^{(1)}$ on $\psi(\hat{\Omega})$




Figure 1: Policies π^l and $\pi^{(1)}$.

(which implies, in particular, that queues in \mathbf{j}_0-1 will be emptied in $(\theta(\omega) + \sigma(\omega) + S_{i_0}(\omega), \theta(\omega) + \sigma(\omega) + \rho(\omega))$).

From the definition of policy $\pi^{(1)}$ it is seen that

$$Q_k^{\pi^l}(t, \omega) = Q_k^{\pi^{(1)}}(t, \omega), \quad \text{for } k = 1, 2, \dots, K, \quad (6.16)$$

for all $\omega \in \Omega$, $t \in [0, \theta(\omega)) \cup (\theta(\omega) + \sigma(\omega) + \rho(\omega), Z(\omega))$.

Therefore,

$$\begin{aligned} J^{\pi^l} - J^{\pi^{(1)}} &= E \left[\int_0^Z \sum_{k=1}^K c_k \left(Q_k^{\pi^l}(t) - Q_k^{\pi^{(1)}}(t) \right) dt \right] \\ &= E \left[\int_\theta^{\theta+\rho+\sigma} \sum_{k=1}^K c_k \left(Q_k^{\pi^l}(t) - Q_k^{\pi^{(1)}}(t) \right) dt \right] \quad \text{from (6.16)} \\ &= E \left[\int_\theta^{\theta+\rho+\sigma} \sum_{k=1}^K c_k \left(Q_k^{\pi^l}(t) - Q_k^{\pi^{(1)}}(t) \circ \psi \right) dt \right] \\ &= \sum_{i>j} P(i_0 = i, j_0 = j) f(i, j), \end{aligned} \quad (6.17)$$

where

$$f(i, j) := E \left[\int_{\theta}^{\theta+\rho+\sigma} \sum_{k=1}^K c_k \left(Q_k^{\pi^l}(t) - Q_k^{\pi^{(1)}}(t) \circ \psi \right) dt \mid i_0 = i, j_0 = j \right].$$

In (6.17) the notation $Q_k^{\pi^{(1)}}(t) \circ \psi$ stands for $\left(Q_k^{\pi^{(1)}}(t) \circ \psi \right) (\omega) = Q_k^{\pi^{(1)}}(t, \psi(\omega))$.

Notice that (6.17) follows from the identities $\theta(\omega) = \theta(\psi(\omega))$, $\rho(\omega) = \rho(\psi(\omega))$, $\sigma(\omega) = \sigma(\psi(\omega))$ that hold by construction of ψ , and from the property that both processes ϕ and ψ have the same distribution.

Define x_k (resp. y_k) to be the number of jobs that left queues in $\mathbf{j}_0 - \mathbf{1}$ in $(\theta + S_{i_0}, \theta + \rho)$ (resp. $(\theta + \rho + S_{j_0}, \theta + \rho + \sigma)$) under π^l to enter queue k for $k \in (\mathbf{j}_0 - \mathbf{1})^c$.

We have (cf. Figure 1)

$$f(i, j) = E \left[\left(c_j - \sum_{k \in (\mathbf{j}-\mathbf{1})^c} c_k y_k \right) \rho \mid i_0 = i, j_0 = j \right] - E \left[\left(c_i - \sum_{k \in (\mathbf{j}-\mathbf{1})^c} c_k x_k \right) \sigma \mid i_0 = i, j_0 = j \right]. \quad (6.18)$$

The right-hand side of (6.18) is the same as the right-hand side of equation (3.12) in Nain et al. [25]. Therefore, by following the arguments in [25] it is easily seen that $f(i, j) \geq 0$ if and only if

$$\frac{c_i - \sum_{k \in (\mathbf{j}-\mathbf{1})^c} p_{ik}^{(j-1)} c_k}{b_i + T_i^{(j-1)}} \leq \frac{c_j - \sum_{k \in (\mathbf{j}-\mathbf{1})^c} p_{jk}^{(j-1)} c_k}{b_j + T_j^{(j-1)}}.$$

This proves that $\pi^{(1)}$ performs better than π^l , that is,

$$J^{\pi^{(1)}} \leq J^{\pi^l}. \quad (6.19)$$

Define now θ to be the first time when $\pi^{(1)}$ does not follow π^* after the l -th decision epoch. Using the same procedure we can generate a new policy $\pi^{(2)}$ that follows π at steps $1, 2, \dots, l-1$, follows π^* at steps l and $l+1$, and such that $J(\pi^{(2)}) \leq J(\pi^{(1)})$. Repeating this argument gives us a sequence of policies $(\pi^{(n)})_n$ such that $\pi^{(n)}$ follows π at steps $1, 2, \dots, l-1$, follows π^* at steps $l, l+1, \dots, l+n-1$, and such that

$$J^{\pi^{(n)}} \leq J^{\pi^{(1)}}. \quad (6.20)$$

Since $\lim_{n \rightarrow \infty} \pi^{(n)} = \pi^{l-1}$ by construction of $\pi^{(n)}$, we see from the bounded convergence theorem that

$$J^{\pi^{l-1}} = \lim_{n \rightarrow \infty} J^{\pi^{(n)}}.$$

Combining this result with (6.19) and (6.20) yields $J^{\pi^{l-1}} \leq J^{\pi^l}$, which completes the proof. ■

7 Appendix

In this appendix, we first prove the equivalence in law of the process ϕ and ψ in Section 5. Then, we give a proof of (5.19).

Proof of the equivalence of ϕ and ψ . In order to show the stochastic equivalence of processes ϕ and ψ , we prove the following claims:

Claim 1: $(\psi_k^1(\omega), \psi_{k+1}^1(\omega))$ has the same distribution as $(\phi_k^1(\omega), \phi_{k+1}^1(\omega))$;

Claim 2: $\psi_k^1(\omega)$ and $\psi_{k+1}^1(\omega)$ are independent of the sequence $(\psi_i^1, i \neq k, i \neq k+1)$;

Claim 3: $\psi_k^1(\omega)$ and $\psi_{k+1}^1(\omega)$ are independent of the sequence ψ^0 .

These assertions clearly imply the stochastic equivalence of the processes ϕ and ψ in view of the assumption of mutual independence between ϕ^0 and ϕ^1 , and the i.i.d. assumption of the sequence ϕ^1 .

The proof of these claims uses the fact that admissible policies of Π do not use information on service times of customers to be routed to queues. This last fact together with the stochastic assumptions made on process ϕ imply that random variables ϕ_k^1 and ϕ_{k+1}^1 are independent of random variables W_k^π, ϕ^0, r and s . Therefore, for all $A, B \subset \mathbb{R}_+$,

$$\begin{aligned}
& P\left(\psi_k^1(\omega) \in A, \psi_{k+1}^1(\omega) \in B\right) \\
&= P\left(\psi_k^1(\omega) \in A, \psi_{k+1}^1(\omega) \in B, \omega \in \widehat{\Omega}\right) + P\left(\psi_k^1(\omega) \in A, \psi_{k+1}^1(\omega) \in B, \omega \in \Omega - \widehat{\Omega}\right) \\
&= P\left(\phi_k^1(\omega) \in A, \phi_{k+1}^1(\omega) \in B, \omega \in \widehat{\Omega}\right) + P\left(\phi_{k+1}^1(\omega) \in A, \phi_k^1(\omega) \in B, \omega \in \Omega - \widehat{\Omega}\right) \\
&= P\left(\phi_k^1(\omega) \in A, \phi_{k+1}^1(\omega) \in B\right) P\left(\omega \in \widehat{\Omega}\right) + P\left(\phi_{k+1}^1(\omega) \in A, \phi_k^1(\omega) \in B\right) P\left(\omega \in \Omega - \widehat{\Omega}\right) \\
&= P\left(\phi_k^1(\omega) \in A\right) P\left(\phi_{k+1}^1(\omega) \in B\right) P\left(\omega \in \widehat{\Omega}\right) + P\left(\phi_{k+1}^1(\omega) \in A\right) P\left(\phi_k^1(\omega) \in B\right) P\left(\omega \in \Omega - \widehat{\Omega}\right) \\
&= P\left(\phi_k^1(\omega) \in A\right) P\left(\phi_k^1(\omega) \in B\right) P\left(\omega \in \widehat{\Omega}\right) + P\left(\phi_k^1(\omega) \in A\right) P\left(\phi_k^1(\omega) \in B\right) P\left(\omega \in \Omega - \widehat{\Omega}\right) \\
&= P\left(\phi_k^1(\omega) \in A\right) P\left(\phi_k^1(\omega) \in B\right) \\
&= P\left(\phi_k^1(\omega) \in A, \phi_{k+1}^1(\omega) \in B\right).
\end{aligned}$$

Hence, Claim 1 holds. Similarly, for all $A, B \subset \mathbb{R}_+$, and $C \subset \mathbb{R}_+^{\mathbb{N}}$,

$$\begin{aligned}
& P\left(\psi_k^1(\omega) \in A, \psi_{k+1}^1(\omega) \in B, \psi^0 \in C\right) \\
&= P\left(\psi_k^1(\omega) \in A, \psi_{k+1}^1(\omega) \in B, \psi^0 \in C, \omega \in \widehat{\Omega}\right)
\end{aligned}$$

$$\begin{aligned}
& +P\left(\psi_k^1(\omega) \in A, \psi_{k+1}^1(\omega) \in B, \psi^0 \in C, \omega \in \Omega - \widehat{\Omega}\right) \\
= & P\left(\phi_k^1(\omega) \in A, \phi_{k+1}^1(\omega) \in B, \phi^0 \in C, \omega \in \widehat{\Omega}\right) \\
& + P\left(\phi_{k+1}^1(\omega) \in A, \phi_k^1(\omega) \in B, \phi^0 \in C, \omega \in \Omega - \widehat{\Omega}\right) \\
= & P\left(\phi_k^1(\omega) \in A, \phi_{k+1}^1(\omega) \in B\right) P\left(\phi^0 \in C, \omega \in \widehat{\Omega}\right) \\
& + P\left(\phi_{k+1}^1(\omega) \in A, \phi_k^1(\omega) \in B\right) P\left(\phi^0 \in C, \omega \in \Omega - \widehat{\Omega}\right) \\
= & P\left(\phi_k^1(\omega) \in A\right) P\left(\phi_{k+1}^1(\omega) \in B\right) P\left(\phi^0 \in C, \omega \in \widehat{\Omega}\right) \\
& + P\left(\phi_{k+1}^1(\omega) \in A\right) P\left(\phi_k^1(\omega) \in B\right) P\left(\phi^0 \in C, \omega \in \Omega - \widehat{\Omega}\right) \\
= & P\left(\phi_k^1(\omega) \in A\right) P\left(\phi_k^1(\omega) \in B\right) P\left(\phi^0 \in C, \omega \in \widehat{\Omega}\right) \\
& + P\left(\phi_k^1(\omega) \in A\right) P\left(\phi_k^1(\omega) \in B\right) P\left(\phi^0 \in C, \omega \in \Omega - \widehat{\Omega}\right) \\
= & P\left(\phi_k^1(\omega) \in A\right) P\left(\phi_k^1(\omega) \in B\right) P\left(\phi^0 \in C\right) \\
= & P\left(\phi_k^1(\omega) \in A, \phi_{k+1}^1(\omega) \in B, \psi^0 \in C\right).
\end{aligned}$$

Hence, Claim 3 holds. Claim 2 can be shown in an analogous way and the detailed proof is left to the interested reader. \blacksquare

Proof of (5.19).

Recall that policy $\pi^{(k)}$ makes SW decisions from the $(k+1)$ -st decision epoch. Thus, if $r \neq s$, i.e., if $\pi_k^{(k)}$ is not an SW decision, then necessarily $\pi_{k+1}^{(k)} = s$. In other words, $\pi_{k+1}^{(k)}$ is independent of ϕ_k^1 (the service time of customer n in process ϕ) under the condition $r \neq s$. Therefore, one can readily check from (5.17) that U_{k+1} is independent of ψ_{k+1}^1 .

More generally, we verify below that policy $\gamma^{(k)}$ thus defined is an admissible policy, $\gamma^{(k)} \in \Pi_{n-1}$. For this, we first note that the workloads of the queues can be described by the routing decisions and the input process ϕ . Indeed, it is easily seen from the evolution equation (5.2) that, given the initial workloads of the queueing system, the workload vector W_m^π is fully determined by the arrival times, service times of previously arrived customers, and the previous routing decisions. Therefore, the mapping π_m in (5.3) can be in fact defined as a mapping $\hat{\pi}_m : \mathbb{R}_+^N \times \mathbb{R}_+^{m-1} \times \{1, 2, \dots, K\}^{m-1} \rightarrow \{1, 2, \dots, K\}$ such that

$$\pi_m\left(\phi^0, (W_i)_{i=1}^m, (\phi_i^1)_{i=1}^{m-1}, (U_i^\pi)_{i=1}^{m-1}\right) = \hat{\pi}_m\left(\phi^0, (\phi_i^1)_{i=1}^{m-1}, (U_i^\pi)_{i=1}^{m-1}\right). \quad (7.1)$$

Using a simple induction implies that there is a mapping $\tilde{\pi}_m : \mathbb{R}_+^N \times \mathbb{R}_+^{m-1} \rightarrow \{1, 2, \dots, K\}$ such that (7.1) can be further written as

$$\pi_m\left(\phi^0, (W_i)_{i=1}^m, (\phi_i^1)_{i=1}^{m-1}, (U_i^\pi)_{i=1}^{m-1}\right) = \tilde{\pi}_m\left(\phi^0, (\phi_i^1)_{i=1}^{m-1}\right). \quad (7.2)$$

We now come back to the proof of the admissibility of the controls $(U_j)_{j=k+1}^{n-1}$ defined in (5.17)-(5.18).

Consider first U_{k+1} . We note that there exist mappings $f_k, g_k : \mathbf{R}_+^{\mathbf{N}} \times \mathbf{R}_+^{k-1} \rightarrow \{1, 2, \dots, K\}$ such that

$$r = f_k \left(\psi^0, (\psi_i^1)_{i=1}^{k-1} \right) \quad (7.3)$$

$$s = g_k \left(\psi^0, (\psi_i^1)_{i=1}^{k-1} \right). \quad (7.4)$$

Therefore, it follows from (5.17) that

$$\begin{aligned} U_{k+1} &= U_{k+1}^{\pi^{(k)}} \mathbf{1}(r = s) + r \mathbf{1}(r \neq s) \\ &= \tilde{\pi}_{k+1}^{(k)} \left(\phi^0, (\phi_i^1)_{i=1}^k \right) \mathbf{1}(r = s) + r \mathbf{1}(r \neq s) \\ &= \tilde{\pi}_{k+1}^{(k)} \left(\psi^0, (\psi_i^1)_{i=1}^k \right) \mathbf{1}(r = s) + r \mathbf{1}(r \neq s). \end{aligned}$$

Thus, in view of (7.3) and (7.4), there exists a mapping $\tilde{\gamma}_{k+1} : \mathbf{R}_+^{\mathbf{N}} \times \mathbf{R}_+^k \rightarrow \{1, 2, \dots, K\}$ such that

$$U_{k+1} = \tilde{\gamma}_{k+1} \left(\psi^0, (\psi_i^1)_{i=1}^k \right). \quad (7.5)$$

Similarly, it follows from (5.18), that for $k+2 \leq j \leq n-1$,

$$\begin{aligned} U_j &= \pi_j^{(k)} \mathbf{1} \left(W_{kr}^{\gamma^{(k)}} \geq \psi_{k+1}^0 - \psi_k^0 \text{ or } r = s \right) \\ &\quad + \left(r \mathbf{1}(U_j^{\pi^{(k)}} = s) + s \mathbf{1}(U_j^{\pi^{(k)}} = r) + U_j^{\pi^{(k)}} \mathbf{1}(U_j^{\pi^{(k)}} \notin \{r, s\}) \right) \mathbf{1} \left(W_{kr}^{\gamma^{(k)}} < \psi_{k+1}^0 - \psi_k^0, r \neq s \right). \end{aligned}$$

Since

$$U_j^{\pi^{(k)}} = \tilde{\pi}_j^{(k)} \left(\phi^0, (\phi_i^1)_{i=1}^{j-1} \right) = \tilde{\pi}_j^{(k)} \left(\psi^0, (\psi_i^1)_{i=1}^{j-1} \right),$$

for $k+2 \leq j \leq n-1$, we obtain a mapping $\tilde{\gamma}_j : \mathbf{R}_+^{\mathbf{N}} \times \mathbf{R}_+^{j-1} \rightarrow \{1, 2, \dots, K\}$ such that (5.19) holds.

Therefore, $(U_j)_{j=k+1}^{n-1}$ are admissible controls, with associated admissible policies $(\tilde{\gamma}_j)_{j=k+1}^{n-1}$. ■

Acknowledgements: The authors are grateful to the reviewer and Professor Rhonda Righter for useful comments on the first draft of the paper.

References

- [1] F. Baccelli, Z. Liu, and D. Towsley, “Extremal Scheduling of Parallel Processing with and without Real-Time Constraints”, *J. ACM*, **40**, 5, 1209 – 1237, 1993.
- [2] J. S. Baras, D.-J. Ma and A. M. Makowski, “K competing queues with geometric requirements and linear costs: the μc -rule is always optimal,” *Systems Control Lett.*, **6**, 173 – 180, 1985
- [3] C. Buyukkoc, P. Varaiya and J. Walrand, “The μc -rule revisited,” *Adv. Appl. Prob.*, **17**, 237 – 238, 1985.
- [4] C. S. Chang, “A New Ordering for Stochastic Majorization: Theory and Applications”, *Adv. Appl. Prob.*, **24**, 604 – 634, 1992.
- [5] P. W. Day, “Rearrangement Inequalities”, *Can. J. Monthly*, **24**, 5, 930 – 943, 1972.
- [6] A. Ephremides, P. Varaiya, and J. Walrand, “A Simple Dynamic Routing Problem”, *IEEE Trans. Aut. Contr.*, **25**, 1980.
- [7] W. Feller, *An Introduction to Probability Theory and Its Applications*, Vol. II. 2nd Ed., Wiley, 1971.
- [8] S. G. Foss, “Approximation of Multichannel Queueing Systems” (in Russian), *Sibirski Mat. Zh.*, **21**, 132 – 140, 1980. (Transl.: *Siberian Math. J.*, **21**, 851–857, 1980.)
- [9] S. G. Foss, “Comparison of Servicing Strategies in Multichannel Queueing Systems” (in Russian), *Sibirski Math. Zh.*, **22**, 190 – 197, 1981. (Transl.: *Siberian Math. J.*, **22**, 142 – 147, 1981.)
- [10] S. G. Foss, “Queues with Customers of Several Types”, in *Limit Theorems and Related Problems*, ed. A. A. Borovokov, Optimization Software, 348–377, 1984.
- [11] S. G. Foss, “Comparison of Service Disciplines in G/GI/m Queues”, INRIA Technical Report, No. 1097, 1989.
- [12] S. Ghahramani and R. W. Wolff, “A New Proof of Finite Moment Conditions for GI/G/1 Busy Periods”, *Queueing Systems*, **4**, 171 – 178, 1989.
- [13] T. Hirayama and M. Kijima, “An Extremal Property of FIFO Discipline in G/IFR/1 Queues”, *Adv. Appl. Prob.*, **21**, 481 – 484, 1989.
- [14] A. Hordijk and G. Koole, “On the Optimality of the Generalized Shortest Queue Policy”, *Prob. Eng. Inf. Sci.*, **4**, 477 – 487, 1991.
- [15] A. Hordijk, G. Koole, “The μc -Rule is not Optimal in the Second Node of the Tandem Queue: a Counterexample”, *Adv. Appl. Prob.*, **24**, 234-237, 1992.
- [16] G. P. Klimov, “Time-Sharing Service Systems I”, *Theory of Prob. and its Appl.*, **19**, 3, 532 – 551, 1974.

- [17] Z. Liu, P. Nain, D. Towsley, "On Optimal Polling Policies", *Queueing Systems*, **11**, 59 – 83, 1992.
- [18] Z. Liu, R. Righter, "Optimal Load Balancing on Distributed Homogeneous Unreliable Processors". Submitted for publication.
- [19] Z. Liu, D. Towsley, "Effects of Service Disciplines in $G/G/s$ Queueing Systems", *Ann. of Oper. Res.*, **48**, 401 – 429, 1994.
- [20] Z. Liu, D. Towsley, "Stochastic Scheduling in In-Forest Networks", *Adv. Appl. Prob.*, **26**, 222 – 241, 1994.
- [21] Z. Liu, D. Towsley, "Optimality of the Round Robin Routing Policy". *J. Appl. Prob.*, **31**, 466 – 475, 1994.
- [22] A. W. Marshall and I. Olkin, *Inequalities: Theory of Majorization and Its Applications*. Academic Press, 1979.
- [23] R. Menich and R. F. Serfozo, "Optimality of Routing and Servicing in Dependent Parallel Processing Systems", *Queueing Systems*, **9**, 403 – 418, 1991.
- [24] P. Nain, "Interchange arguments for classical scheduling problems in queues", *Syst. Contr. Lett.*, **12**, 177 – 184, 1989.
- [25] P. Nain, P. Tsoucas, and J. Walrand, "Interchange Arguments in Stochastic Scheduling", *J. Appl. Prob.*, **27**, 815 – 826, 1989.
- [26] P. Nain and D. Towsley, "Optimal Scheduling in a Machine with Stochastic Varying Processing Rate", *IEEE Trans. Aut. Contr.*, **39**, 9, 1853 – 1855, 1994.
- [27] S. M. Ross, *Applied Probability Models with Optimization Applications*. Holden-Day, San Francisco, 1970.
- [28] M. Schäl, "Conditions for optimality in dynamic programming and for the limit of n -stage optimal policies to be optimal," *Z. Wahrscheinlichkeitsth.*, **32**, 179 – 196, 1975.
- [29] J. G. Shanthikumar and U. Sumita, "Convex Ordering of Sojourn Times in Single-Server Queues: Extremal Properties of FIFO and LIFO Service Disciplines", *J. Appl. Prob.*, **24**, 737 – 748, 1987.
- [30] P. D. Sparaggis, C. G. Cassandras, and D. Towsley, "On the Duality Between Routing and Scheduling Systems with Finite Buffer Space", *IEEE Trans. Aut. Contr.*, **38**, 9, 1440 – 1446, 1993.
- [31] P. D. Sparaggis, D. Towsley, and C. G. Cassandras, "Sample Path Criteria for Weak Majorization", *Adv. Appl. Prob.*, **26**, 155 – 171, 1994.
- [32] P. D. Sparaggis, C. G. Cassandras, and D. Towsley, "Optimality of Static Routing Policies in Queueing Systems with Blocking". To appear in *IEEE Trans. Aut. Contr.*

- [33] D. Stoyan, *Comparison Methods for Queues and Other Stochastic Models*. J. Wiley & Sons, Berlin, 1983.
- [34] D. Towsley and F. Baccelli, “Comparison of Service Disciplines in Tandem Queueing Networks with Delay Dependent Customer Behavior”, *Oper. Res. Lett.*, **10**, 49 – 55, 1991.
- [35] D. Towsley, S. Fdida, and H. Santoso, “Flow Control Protocols for Interconnected High Speed Networks”, *Proc. of NATO Workshop on High Speed Networks*, Sophia-Antipolis, France, 1990.
- [36] D. Towsley, P. Sparaggis. “Optimal Routing in Systems with ILR Service Time Distributions”, Department of Computer Science Technical Report TR93-13, U. Massachusetts, 1993.
- [37] D. Towsley, P. Sparaggis, and C. G. Cassandras, “Optimal Routing and Buffer Allocation for a Class of Finite Capacity Queueing Systems”, *IEEE Trans. Aut. Contr.*, **37**, 1446 – 1451, 1992.
- [38] P. P. Varaiya, J. Walrand, C. Buyukkoc, “Extensions of the Multiarmed Bandit Problem: the Discounted Case”, *IEEE Trans. Aut. Contr.*, **30**, 426 – 439, 1985.
- [39] O. A. Vasicek, “An Inequality for the Variance of Waiting Time under a General Queueing Discipline”, *Oper. Res.*, **25**, 879 – 884, 1977.
- [40] J. Walrand, *An Introduction to Queueing Networks*. Prentice Hall, Englewoods Cliff, 1988.
- [41] R.R. Weber, “On the Optimal Assignment of Customers to Parallel Queue”, *J. Appl. Prob.*, **15**, 406 – 413, 1978.
- [42] G. Weiss, “Branching Bandit Processes”, *Prob. Eng. Inf. Sci.*, **2**, 269 – 278, 1988.
- [43] W. Whitt, “Deciding Which Queue to Join”, *Oper. Res.*, **34**, 55 – 62, 1986.
- [44] W. Winston, “Optimality of the Shortest Line Discipline”, *J. Appl. Prob.*, **14**, 181 – 189, 1977.