Stochastic Scheduling in In-Forest Networks*

Zhen LIU†
INRIA Centre Sophia Antipolis
2004 Route des Lucioles
06560 Valbonne, France

Don TOWSLEY
Dept. of Computer & Information Science
University of Massachusetts
Amherst, MA 01003, U.S.A.

First draft in April 1992, Revised in January 1993

Abstract

In this paper we study the extremal properties of several scheduling policies in an in-forest network consisting of multi-server queues. Each customer has a due date and we assume that service times at the different queues form mutually independent sequences of independent and identically distributed random variables independent of the arrival times and due dates. Furthermore, the network is assumed to consist of a mixture of nodes, some of which only permit non-preemptive service policies whereas the others permit preemptive resume policies. In the case of non-preemptive queues, service times may be generally distributed if there is only one server; otherwise the service times are required to be increasing in likelihood ratio (ILR). In the case of preemptive queues, service times are restricted to exponential distributions. Using stochastic majorizations and partial orders on permutations, we establish that, within the class of work conserving service policies, the stochastically smallest due date (SSDD) and the stochastically largest due date (SLDD) policies minimize and maximize, respectively, the vector of the customer latenesses of the first n customers in the sense of the Schur-convex order and some weaker orders, provided the due dates are comparable in some stochastic sense. It then follows that the first come first serve (FCFS) and last come first serve (LCFS) policies minimize and maximize, respectively, the vector of the response times of the first n customers in the sense of the Schur-convex order. We also show that the FCFS and LCFS policies minimize and maximize, respectively, the vector of customer end-to-end delays in the sense of the strong stochastic order. Extensions to the class of non-idling policies and to the stationary regime are also given.

Keywords : Queueing System, Service Discipline, Sample Path Analysis, Stochastic Ordering, Majorization, Permutation Ordering, Lateness, Response Time, End-to-End Delay.
To appear in *Advances of Applied Probability.*

*This work was supported in part by theNational Science Foundation under grants ASC 88-8802764, NCR-9116183.
†The work of this author was also partially supported by CEC DG-XIII under the ESPRIT-BRA grant QMIPS.

1 Introduction

Consider a network of $\cdot/GI/s$ queues having an in-forest topology. The arrival times of the customers are arbitrary, whereas the service times at each queue are independent and identically distributed (i.i.d.) random variables (r.v.'s). The sequences of service times are mutually independent and are independent of the arrival times and due dates. At each queue, a customer can be served by any of the servers. These servers are identical and have the same service rate, say 1. Finally, there are due dates associated with the customers. We study the effect of different service policies on customer response time, end-to-end delay (defined as the sum of its sojourn time and its resequencing delay, see Section 2 for a precise definition) and lateness (defined to be the difference between the customer completion time and its due date).

Many papers have studied the effects that service policies have on the performance of a single $G/GI/s$ queue. It has been shown by various authors that the first come first serve (FCFS) policy minimizes the stationary waiting times in the sense of convex ordering when the scheduling policies are non-preemptive and use only the information on the distribution of service times, see Kingman [14], Vasicek [27], Foss [10, 11], Wolff [29, 30] and Daley [8]. When service times have an Erlang distribution, preemptions are allowed and there is a single server $s = 1$, Shantikumar and Sumita [23] showed that the FCFS policy minimizes the stationary waiting times in the sense of increasing convex ordering. This last result was generalized by Hirayama and Kijima [12] and Chang and Yao [6] to the case when the service time distribution is of Increasing Failure Rate (IFR) type.

Several papers have studied the effect that different scheduling policies have on the customer lateness. The optimality of the shortest due date (SDD) policy was first established in [25] for queues in tandem in the sense of convex ordering. This result was then generalized to a class of parallel processing systems [2]. Finally, the optimality of stochastic versions of SDD has been established for the $G/GI/1$ queue [6] and $G/GI/s$ queue [17], and also for the $G/M/s$ queue when customers have hard deadlines (Note that under hard deadline, a customer leaves the system either when it finishes service or when its due date occurs) [26].

Another important property of queueing systems is overtaking and resequencing (cf. Kleinrock et al. [15], Whitt [28], Baccelli et al. [1, 3]). Whitt [28] analyzed the number of customers overtaken by an arbitrary customer for $GI/M/s$ and $M/GI/s$ models with FCFS service policy. Iliadis and Lien [13] studied the resequencing delay for two heterogeneous servers under threshold-type scheduling.

In this paper, we compare different scheduling policies in in-forest queueing networks consisting of $\cdot/GI/s$ queues with delay dependent customer behavior. A queue can have *i)* either a single server with general service time distribution and nonpreemptive service, *ii)* or multiple servers with increasing in likelihood ratio (ILR) service time distribution and nonpreemptive service, *iii)* or multiple servers with exponential service time distribution and preemptive service. When the due dates are comparable in some stochastic sense, we show that, the stochastically smallest due date (SSDD) and the stochastically largest due date (SLDD) policies minimize and maximize, respectively, the vector of the customer latenesses of the first n customers in the sense of the Schur-convex order and some weaker orders. ¿From this we conclude that the first come first serve (FCFS) and last come first serve (LCFS) policies minimize and maximize, respectively, the vector of the response times of the first n customers in the sense of the Schur-convex order. We also show that the FCFS and LCFS policies minimize and maximize, respectively, the vector of customer end-to-end delays in the sense of the strong stochastic order. Here the FCFS, LCFS, SDD, LDD, SSDD, SLDD policies are nonpreemptive if preemptions are prohibited. Otherwise, they are preemptive resume policies.

Our primary emphasis is on the class of non-idling policies. However, we describe extensions to the larger class of idling policies, where the FCFS policy is shown to be optimal for various performance measures in the case that all queues permit preemptive resume policies and all service times are exponentially distributed. All of these extremal properties are first analyzed in the transient regime, and then extended to the stationary regime.

Our results are obtained through sample path analysis. We use notions of majorization and stochastic orders. We also develop some properties associated with permutation orderings which were first introduced in Baccelli, Liu and Towsley [2]. These permutation orderings and their properties turn out to be crucial in establishing the main results of the paper.

The paper is organized as follows. In the next section, we define in a more precise way the model as well as the notation and assumptions. Section 3 presents some preliminaries on stochastic comparisons. The main results are formally stated in Section 4 and the proofs are contained in Section 5. A discussion of the stationary regime and other generalizations. as well as counterexamples are provided in Section 6.

2 Notation and Assumptions

2.1 Model Description

Consider an acyclic network $\mathcal{K}$ with $K \geq 1$ nodes. Node i, $1 \leq i \leq K$, consists of a waiting queue of infinite capacity and one or more identical servers. The network (or more precisely, its underlying graph) has an *in-forest* structure, viz., a node has at most one successor. The nodes having no predecessors are referred to as the leaves of the in-forest, and those having no successors as the roots. A customer served at a non-root node will later be served at the successor of the node. Without loss of generality, we assume that the nodes are labeled in such a way that node i is a predecessor of node j implies that $i < j$. Figure 1 illustrates an example of an in-forest.

The service times of the queues in the system form mutually independent r.v.'s. For each queue, the service times are i.i.d. r.v.'s. We distinguish between three types of nodes.

- a *type 1* node which contains a single server with general service time distribution and non-

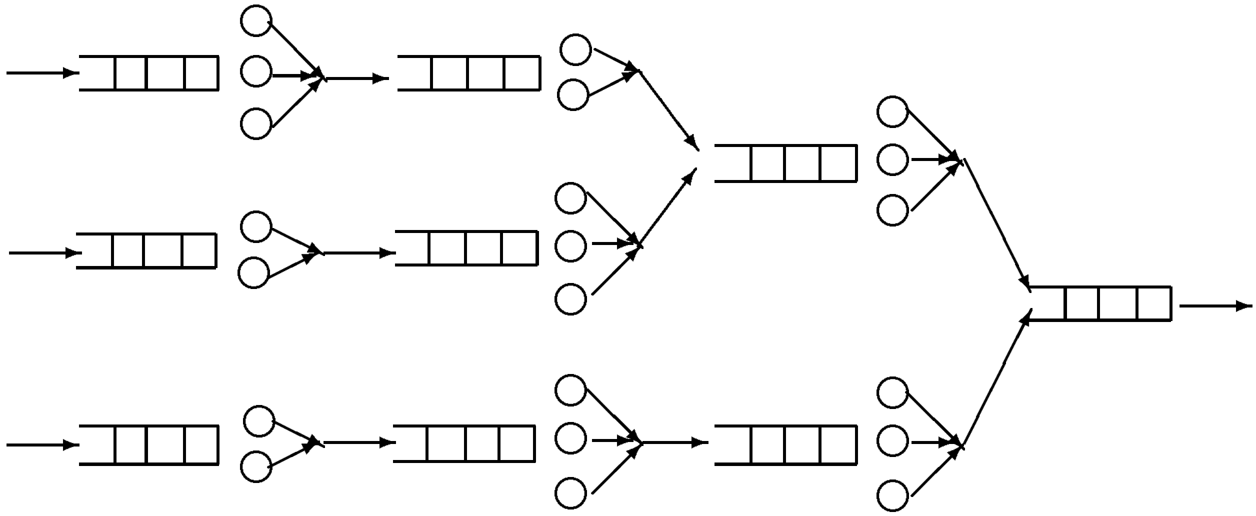

Figure 1: Example of an in-forest.

preemptive service,

- a *type 2* node which contains multiple servers with increasing likelihood ratio (ILR) service time distribution and nonpreemptive service,
- a *type 3* node which contains multiple servers with exponential service time distribution and preemptive (resume) service.

The n-th ($n \geq 1$) customer, also referred to as customer n, arrives in the system at time a_n, $a_1 < a_2 < \cdots < a_n < \cdots$. We associate with customer n, $n \geq 1$, a due date, denoted by d_n. Let $u_n = d_n - a_n$ be the relative due date of customer n. Both d_n and u_n are (not necessarily positive) real numbers.

When a customer arrives in the system, it enters one of the leaf queues. When a customer enters a queue, it waits for service in the queue. The service time is a random variable whose distribution depends only on the identity of the queue. After having been served by one of the servers associated with that queue, it is routed to the successor queue.

There is a resequencing buffer with infinite capacity in the system. When a customer leaves a root node, it enters the resequencing buffer. A customer, say n, can leave the resequencing buffer (and, thus the system), if and only if all of the customers $1, 2, \cdots, n-1$ have already left this buffer. The resequencing buffer is also assumed to be of infinite size.

In order to analyze the transient behavior of various scheduling policies, we arbitrarily fix $N \geq 1$ as the number of total arrivals. Let $\mathcal{N} = \{1, 2, \cdots, N\}$ be the set of customers that arrive to the system. Denote by $\mathcal{N}_i \subseteq \mathcal{N}$ the set of customers routing through node i, $i \in \mathcal{K}$.

The sequences of arrival times $\mathcal{A} = \{a_n\}_{n=1}^{\infty}$ and of due dates $\mathcal{D} = \{d_n\}_{n=1}^{\infty}$ are independent of the service times, but are otherwise arbitrary. In particular, they can be deterministic sequences.

2.2 Scheduling Policies

A scheduling policy determines the time at which a particular customer is to be served in a queue. If the node is either of type 1 or 2, then the policy is restricted to be non-preemptive, i.e., no customer is ever interrupted and removed from a server while it is in the middle of service. If the node is of type 3, then the policy may be preemptive resume where the service of a customer is resumed at the point at which it was preempted. The policy is called non-idling or work conserving if no server is allowed to remain idle whenever there is a customer waiting in the queue.

Throughout this paper we assume that the scheduling policies cannot use any information on service time other than that regarding the distribution of the service times. This assumption implies that the shortest remaining processing time policy is not under consideration. We also assume that the scheduling policies are not anticipative in the sense that a decision can never use information on future arrivals.

Denote by Ψ the class of (possibly idling) policies that fulfill the above assumptions and and $\Psi_{ni} \subset \Psi$ the class of non-idling policies. Whether or not these policies permit preemptions depends on the type of node. Among the well-known extremal policies, there are FCFS and LCFS policies, which serve the customers according to their arrival dates $\mathcal{A} = \{a_n\}_{n=1}^{\infty}$. Note that the FCFS and LCFS policies thus defined are "global" in the sense that the times when the customers arrive in the system are used instead of the times when the customers arrive at that queue. By convention, we will assume that the FCFS and LCFS policies are non-idling. When the due dates are comparable in some stochastic ordering sense $\leq_{lr}$, $\leq_{hr}$ or $\leq_{st}$ (see the definitions below), which is the case when they are known or when they are unknown but the relative due dates are i.i.d. random variables (with ILR distribution if $\leq_{lr}$ is under consideration), we define the SSDD policies and the SLDD policies to be such that as soon as there is an available server, the customer waiting in the queue with the stochastically smallest and largest due dates, respectively, is assigned to the server. Again, by definition, SSDD and SLDD policies are non-idling.

Observe that when the relative due dates are i.i.d. random variables and are independent of the arrival and service times, and also are unknown a priori, then the SSDD and SLDD policies coincide with the FCFS and LCFS policies, respectively. When the due dates are known, the SSDD and SLDD policies coincide with the SDD and LDD policies, respectively.

2.3 Performance Measures

Let $\pi \in \Psi$ be an arbitrary scheduling policy. Denote by $c_n(\pi)$ the random variable (in $I\!R^+$) of the completion time of customer n at one of the root nodes. Denote by $R_n(\pi)$ and $L_n(\pi)$ the response time and the lateness of customer n under $\pi \in \Psi$, respectively, defined by $R_n(\pi) = c_n(\pi) - a_n$,

$L_n(\pi) = c_n(\pi) - d_n$. Denote by $D_n(\pi)$ the end-to-end delay of customer n under $\pi \in \Psi$, i.e., $D_n(\pi) = \max_{1\le l\le n} c_l(\pi) - a_n$. Let $\boldsymbol{R}(\pi) = (R_1(\pi), \cdots, R_N(\pi))$, $\boldsymbol{L}(\pi) = (L_1(\pi), \cdots, L_N(\pi))$, and $\boldsymbol{D}(\pi) = (D_1(\pi), \cdots, D_N(\pi))$.

3 Preliminaries on Stochastic Orderings

In this section, we first review some concepts of stochastic orderings. The interested reader is referred to [24, 22, 19, 6] for various properties and applications of these notions. The last subsection develops some properties associated with permutation orderings which will be crucial in establishing the main results of the paper.

Throughout this paper, the inequality $\le$ between two vectors is understood to be componentwise. Increasingness and decreasingness are used in the non-strict sense. Stochastic (partial) orders stand for stochastic (partial) preorders.

3.1 Stochastic Orderings of Random Variables

Definition 3.1 *Let $X, Y \in I\!R$ be two random variables. The random variable X is smaller than Y in the sense of strong stochastic ordering (resp. convex ordering, increasing convex ordering and decreasing convex ordering), denoted by $X \le_{st} Y$ (resp. $X \le_{cx} Y$, $X \le_{icx} Y$ and $X \le_{dcx} Y$) if the inequality $E[f(X)] \le E[f(Y)]$ holds for all increasing (resp. convex, increasing and convex, decreasing and convex) functions $f : I\!R \to I\!R$, provided the expectations exist.*

Note that such stochastic orderings can be defined for random vectors in a similar way. Clearly, the strong stochastic ordering $\le_{st}$ implies the increasing convex ordering $\le_{icx}$. The following two orderings are stronger than $\le_{st}$.

Definition 3.2 *Let $X, Y \in I\!R$ be two random variables with density functions f_X and f_Y respectively, and distribution functions F_X and F_Y respectively. The random variable X is smaller than Y in the likelihood ratio ordering (resp. hazard rate ordering), denoted by $X \le_{lr} Y$ (resp. $X \le_{hr} Y$), if for any x and any $a > 0$, $f_X(x)/f_X(x+a) \ge f_Y(x)/f_Y(x+a)$ (resp. for any x, $f_X(x)/[1 - F_X(x)] \ge f_Y(x)/[1 - F_Y(x)]$).*

When X and Y are discrete random variables, the orderings $\le_{lr}$ and $\le_{hr}$ are similarly defined with density functions replaced by the probability distributions.

It is known that the likelihood ratio ordering implies the hazard rate ordering, which in turn implies the stochastic ordering, i.e., $X \le_{lr} Y \;\Rightarrow\; X \le_{hr} Y \;\Rightarrow\; X \le_{st} Y$, where the symbol $\Rightarrow$ stands for the implication. Note that two constants are comparable in the sense of $\le_{lr}$ ordering.

A number of our results will require that service times have distributions with increasing likelihood ratio (ILR).

Definition 3.3 *The random variable $X \in I\!R^+$ is said to be increasing in likelihood ratio (ILR) if for all $0 \le s \le t$, $X_s \ge_{lr} X_t$, where X_t is the remaining lifetime of X from t on, given that it exceeds t.*

A random variable is increasing in likelihood ratio (ILR) iff its density function is log-concave (or, Polya frequency of order 2). Examples of random variables that are ILR include those with the following densities, *i)* Gamma: $f(x) = \lambda e^{-\lambda x}(\lambda x)^{\alpha-1}/\Gamma(\alpha)$, $\alpha > 1$ and *ii)* Weibull: $f(x) = \alpha\lambda(\lambda x)^{\alpha-1}e^{-(\lambda x)^{\alpha}}$, $\alpha \ge 1$.

3.2 Stochastic Majorizations

Define now the notion of majorization. Let $\boldsymbol{x}, \boldsymbol{y} \in I\!R^n$ be two real vectors.

Definition 3.4 *Vector $\boldsymbol{x}$ is said to be majorized by vector $\boldsymbol{y}$ (written $\boldsymbol{x} \prec \boldsymbol{y}$) if*

$$\sum_{i=1}^{k} x_{[i]} \le \sum_{i=1}^{k} y_{[i]}, \quad k = 1, \cdots, n-1, \qquad \text{and} \qquad \sum_{i=1}^{n} x_{[i]} = \sum_{i=1}^{n} y_{[i]},$$

where the notation $x_{[i]}$ is taken to be the i-th largest element of $\boldsymbol{x}$.

Definition 3.5 *A function $f : I\!R^n \to I\!R$ is Schur-convex if for all $\boldsymbol{x}, \boldsymbol{y} \in I\!R^n$ such that $\boldsymbol{x} \prec \boldsymbol{y}$,*

$$f(\boldsymbol{x}) \le f(\boldsymbol{y}).$$

A function $f : I\!R^n \to I\!R$ is L-subadditive if for all $\epsilon_1, \epsilon_2 \ge 0$ and $1 \le i < j \le n$,

$$f(\boldsymbol{x}) + f(\boldsymbol{x} + \epsilon_1 \boldsymbol{e}_i + \epsilon_1 \boldsymbol{e}_j) \le f(\boldsymbol{x} + \epsilon_1 \boldsymbol{e}_i) + f(\boldsymbol{x} + \epsilon_1 \boldsymbol{e}_j),$$

where $\mathbf{e}_i$ and $\mathbf{e}_j$ denote the i-th and j-th unit vector.

Define the following classes of functions

- $\mathcal{C}_1$ ($\mathcal{C}_1^{\uparrow}, \mathcal{C}_1^{\downarrow}$) — the class of (increasing, decreasing) Schur-convex functions,
- $\mathcal{C}_2$ ($\mathcal{C}_2^{\uparrow}, \mathcal{C}_2^{\downarrow}$) — the class of (increasing, decreasing) symmetric and convex functions,
- $\mathcal{C}_3$ ($\mathcal{C}_3^{\uparrow}, \mathcal{C}_3^{\downarrow}$) — the class of functions of the form $f(\boldsymbol{x}) = \sum_{i=1}^{n} g(x_i)$, where g is (increasing, decreasing) convex,

- $\mathcal{C}_5$ ($\mathcal{C}_5^{\uparrow}$,$\mathcal{C}_5^{\downarrow}$) — the class of functions that are (increasing, decreasing) symmetric, L-subadditive and convex in each variable,

- $\mathcal{C}_{5-icx}$ ($\mathcal{C}_{5-icx}^{\uparrow}$, $\mathcal{C}_{5-icx}^{\downarrow}$) — the class of functions of the form $g \circ h$, where $h \in \mathcal{C}_5$ ($\mathcal{C}_5^{\uparrow}$,$\mathcal{C}_5^{\downarrow}$) and g is increasing and convex.

Definition 3.6 *Let $\boldsymbol{X}$ and $\boldsymbol{Y}$ be two random vectors in $I\!R^n$. We define the following stochastic orderings between these r.v.'s for $i = 1, 2, 3, 5, 5$-icx :*

$$\boldsymbol{X} \leq_{E_i} \boldsymbol{Y}, \quad \textit{if} \quad E[f(\boldsymbol{X})] \leq E[f(\boldsymbol{Y})], \quad \forall f \in \mathcal{C}_i,$$

provided the expectations exist.

The orderings $\leq_{E_i^{\uparrow}}$ and $\leq_{E_i^{\downarrow}}$ can be defined in a similar way, $i = 1, 2, 3, 5, 5$-*icx*.

In the literature, the orderings $\leq_{E_1}$, $\leq_{E_2}$ and $\leq_{E_3}$ are also referred to as the Schur-convex ordering, convex symmetric ordering and separable convex ordering, respectively. It is known that $\leq_{E_1} \Rightarrow \leq_{E_2} \Rightarrow \leq_{E_3}$ and $\leq_{E_1} \Rightarrow \leq_{E_{5-icx}} \Rightarrow \leq_{E_5} \Rightarrow \leq_{E_3}$.

3.3 Partial Orderings on Permutations

In this subsection, we extend the notion of partial ordering on permutations which was first introduced in Baccelli, Liu and Towsley [2]. Unless otherwisely stated, the results presented in this subsection are proved in [18].

Let $n \geq 1$ be an arbitrary integer, and Γ the set of permutations on $\{1, 2, \cdots, n\}$. For any two real vectors $\boldsymbol{x}, \boldsymbol{y} \in I\!R^n$, denote by $\boldsymbol{y}_\gamma = (y_{\gamma(1)}, \cdots, y_{\gamma(n)})$ where $\gamma \in \Gamma$, and $\boldsymbol{y} - \boldsymbol{x} = (y_1 - x_1, \cdots, y_n - x_n)$. Let $\boldsymbol{X}, \boldsymbol{Y} \in I\!R^n$ be two random vectors, and $\leq_a \in \{\leq_{lr}, \leq_{hr}, \leq_{st}\}$ be some ordering on the real or random vectors. Assume that the components of $\boldsymbol{X}$ are comparable in the sense of $\leq_a$, i.e., for any $i \neq j$, either $X_i \leq_a X_j$ or $X_i \geq_a X_j$.

Define the binary relation $\mathcal{B}^a_{\boldsymbol{X}}$ on the symmetric group Γ as:

Definition 3.7 *$\gamma' \mathcal{B}^a_{\boldsymbol{X}} \gamma$ if $\gamma' = \gamma$ or if there exist a pair of integers j, k, such that*

$$X_j \leq_a X_k, \quad \gamma(j) > \gamma(k), \quad \gamma'(j) = \gamma(k), \quad \gamma'(k) = \gamma(j), \quad \gamma'(i) = \gamma(i), \quad i \neq j, \ i \neq k, \qquad (3.1)$$

that is, γ' is the same as γ except that $\gamma(j)$ and $\gamma(k)$ are interchanged when $X_j \leq_a X_k$.

Define now a partial order $\prec^a_{\boldsymbol{X}}$ on Γ as the transitive closure of $\mathcal{B}^a_{\boldsymbol{X}}$:

1. $\gamma' \prec^a_{\boldsymbol{X}} \gamma$ if $\gamma' \mathcal{B}^a_{\boldsymbol{X}} \gamma$.

2. $\gamma' \prec_{\boldsymbol{X}} \gamma$ if there exists γ'' such that $\gamma' \prec^a_{\boldsymbol{X}} \gamma''$ and $\gamma'' \prec^a_{\boldsymbol{X}} \gamma$.

Note that $\prec^a_{\boldsymbol{X}}$ and $\prec^a_{\boldsymbol{X}'}$ define the same partial order provided

$$\forall i, j : X_i \leq_a X_j \quad \text{iff} \quad X_i' \leq_a X_j'.$$

Note also that when $\boldsymbol{X}$ is a constant equal to $\boldsymbol{x}$, the partial order $\prec^a_{\boldsymbol{X}}$ coincides with the deterministic partial order $\prec_{\boldsymbol{x}}$ that was first introduced in [2]. In such a case, the following result holds.

Lemma 3.1 *Assume that $x_1 \leq x_2 \leq \cdots \leq x_n$ and that $y_1 \leq y_2 \leq \cdots \leq y_n$. If $\gamma' \prec_{\boldsymbol{x}} \gamma$, then*

$$\max_{1 \leq i \leq m} y_{\gamma'(i)} \leq \max_{1 \leq i \leq m} y_{\gamma(i)}, \qquad 1 \leq m \leq n. \tag{3.2}$$

In general, the permutations are random (vectors) for which the following relations are valid.

Lemma 3.2 *Let $\boldsymbol{Y} \in I\!R^n$ be a random vector such that $Y_1 \leq Y_2 \leq \cdots \leq Y_n$, a.s. Assume further that $\boldsymbol{X}$ is independent of $\boldsymbol{Y}$ and of the random permutations $\gamma', \gamma \in \Gamma$. If $\gamma' \prec^{st}_{\boldsymbol{X}} \gamma$, then*

$$(\boldsymbol{Y}_{\gamma'} - \boldsymbol{X}) \leq_{E_3} (\boldsymbol{Y}_{\gamma} - \boldsymbol{X}). \tag{3.3}$$

When the components of $\boldsymbol{X}$ are mutually independent, we obtain the following property from Theorem 3.2 of [6].

Lemma 3.3 *Let $\boldsymbol{Y} \in I\!R^n$ be a random vector such that $Y_1 \leq Y_2 \leq \cdots \leq Y_n$, a.s. Let $\boldsymbol{X}$ be a random vector which has mutually independent components. Assume further that $\boldsymbol{X}$ is independent of $\boldsymbol{Y}$ and of the random permutations $\gamma', \gamma \in \Gamma$. If $\gamma' \prec^a_{\{X_i\}} \gamma$, then*

$$(\boldsymbol{Y}_{\gamma'} - \boldsymbol{X}) \leq_b (\boldsymbol{Y}_{\gamma} - \boldsymbol{X}), \tag{3.4}$$

where $(a, b) \in \{(lr, E_1), (hr, E_{5-icx}), (st, E_5)\}$.

As a consequence of the above lemma, if $\boldsymbol{X}$ is equal to a constant $\boldsymbol{x}$, then $(\boldsymbol{Y}_{\gamma'} - \boldsymbol{x}) \leq_{E_1} (\boldsymbol{Y}_{\gamma} - \boldsymbol{x})$, provided $\gamma' \prec_{\boldsymbol{x}} \gamma$, *a.s.*

Consider now the partial ordering on the "merging" of the permutations. Let N_1, N_2 be a partition of $\{1, 2, \cdots, n\}$: $N_1 \bigcup N_2 = \{1, 2, \cdots, n\}$, $N_1 \bigcap N_2 = \emptyset$. For $k = 1, 2$, let $\boldsymbol{X}_{N_k} = \{X_i\}_{i \in N_k}$. Denote by Γ_1 and Γ_2 the sets of permutations on N_1 and N_2, respectively. Let $\gamma_1 \in \Gamma_1$ and $\gamma_2 \in \Gamma_2$. Define the merging of the permutations γ_1, γ_2 as $\gamma = (\gamma_1, \gamma_2)$ such that

$$\gamma(i) = \begin{cases} \gamma_1(i), & i \in N_1; \\ \gamma_2(i), & i \in N_2. \end{cases}$$

Lemma 3.4 *Let $\boldsymbol{X} \in I\!R^n$ be a random vector whose components are comparable in the sense of $\leq_a \in \{\leq_{lr}, \leq_{hr}, \leq_{st}\}$. Let $\gamma_1, \gamma_1' \in \Gamma_1$, $\gamma_2, \gamma_2' \in \Gamma_2$. If $\gamma_1' \prec^a_{\boldsymbol{X}_{N_1}} \gamma_1$ and $\gamma_2' \prec^a_{\boldsymbol{X}_{N_2}} \gamma_2$, then*

$$\gamma' = (\gamma_1', \gamma_2') \prec^a_{\boldsymbol{X}} (\gamma_1, \gamma_2) = \gamma.$$

4 Main Results

In this section, we present the extremal properties of some scheduling policies in the multi-server in-forest network $\mathcal{K}$. The proofs of these results are provided in the next section. We focus on two classes of policies, the class of non-idling policies and the class of idling policies. We provide various extremal properties of the policies FCFS, LCFS, SSDD and SLDD that are applied to all of the nodes in the network.

4.1 Results for Non-Idling Policies

Consider first the latenesses of the customers. We will assume that the due dates are comparable in the sense of $\leq_a \in \{\leq_{lr}, \leq_{hr}, \leq_{st}\}$. Note that if the due dates are known (so that they are deterministic and are comparable in $\leq$), they are comparable in $\leq_{lr}$. Therefore, all the results which hold under the assumption that the due dates are comparable in $\leq_{lr}$ remain valid in the known due date case.

Theorem 4.1 *If for any fixed sequence of arrival times $\mathcal{A} = \{a_n\}_{n=1}^N$, the due dates are stochastically comparable in the sense of $\leq_{st}$, then the SSDD (resp. SLDD) policy applied to all of the nodes minimizes (resp. maximizes) the vector of latenesses within the class Ψ_{ni} in the sense of $\leq_{E_3}$ ordering,*

$$\forall \pi \in \Psi_{ni}: \qquad \boldsymbol{L}(SSDD) \leq_{E_3} \boldsymbol{L}(\pi) \leq_{E_3} \boldsymbol{L}(SLDD). \tag{4.1}$$

Remark: If the due dates are unknown a priori, and if the relative due dates are stochastically increasing random variables in the sense of $\leq_{st}$, (i.e., $u_n \leq_{st} u_{n+1}$, for all $n \geq 1$), independent of the arrival times (e.g., the relative due dates are i.i.d.), then the SSDD and SLDD policies coincide with the FCFS and LCFS policies. Consequently, relation (4.1) still holds for this model when SSDD and SLDD are replaced by FCFS and LCFS, respectively.

Theorem 4.2 *If for any fixed sequence of arrival times $\mathcal{A} = \{a_n\}_{n=1}^N$, the due dates are mutually independent and are stochastically comparable in the sense of $\leq_a \in \{\leq_{lr}, \leq_{hr}, \leq_{st}\}$, then the SSDD (resp. SLDD) policy applied to all the nodes minimizes (resp. maximizes) the vector of latenesses within the class Ψ_{ni} in the sense of $\leq_b$,*

$$\forall \pi \in \Psi_{ni}: \qquad \boldsymbol{L}(SSDD) \leq_b \boldsymbol{L}(\pi) \leq_b \boldsymbol{L}(SLDD), \tag{4.2}$$

where $(a, b) \in \{(lr, E_1), (hr, E_{5-icx}), (st, E_5)\}$.

Setting the due dates to the arrival times in Theorem 4.2 implies the following extremal properties of the FCFS and LCFS policies on response times.

Corollary 4.1 *The FCFS (resp. LCFS) policy applied to all the nodes minimizes (resp. maximizes) the vector of response times within the class Ψ_{ni} in the sense of $\leq_{E_1}$ ordering,*

$$\forall \pi \in \Psi_{ni}: \qquad \boldsymbol{R}(FCFS) \leq_{E_1} \boldsymbol{R}(\pi) \leq_{E_1} \boldsymbol{R}(LCFS). \tag{4.3}$$

Owing to Corollary 4.1, a stronger relation than (4.1) and (4.2) can be obtained in the case that the relative due dates are i.i.d. random variables which are independent of the arrival times. Indeed, $\boldsymbol{L}(\pi)$ can be rewritten as

$$\boldsymbol{L}(\pi) = (R_1(\pi) - u_1, R_2(\pi) - u_2, \cdots, R_N(\pi) - u_N),$$

where $u_1, \cdots, u_N$ are i.i.d random relative due dates being independent of the response times. Appealing to the closure (under convolution) property of the $\leq_{E_2}$ ordering [19, Proposition F.6, p. 314] and making use of Corollary 4.1 readily yield the following:

Corollary 4.2 *If the due dates are unknown a priori, and if the relative due dates are i.i.d. random variables which are independent of the arrival times, then the FCFS (resp. LCFS) policy applied to all the nodes minimizes (resp. maximizes) the vector of latenesses within the class Ψ_{ni} in the sense of $\leq_{E_2}$ ordering,*

$$\forall \pi \in \Psi_{ni}: \qquad \boldsymbol{L}(FCFS) \leq_{E_2} \boldsymbol{L}(\pi) \leq_{E_2} \boldsymbol{L}(LCFS). \tag{4.4}$$

Consider now the end-to-end delays in the network.

Theorem 4.3 *The FCFS (resp. LCFS) policy applied to all of the nodes minimizes (resp. maximizes) the end-to-end delays within the class Ψ_{ni} in the sense of stochastic ordering $\leq_{st}$,*

$$\forall \pi \in \Psi_{ni}: \qquad \boldsymbol{D}(FCFS) \leq_{st} \boldsymbol{D}(\pi) \leq_{st} \boldsymbol{D}(LCFS). \tag{4.5}$$

4.2 Results for Idling Policies

If we consider the class of idling policies, then we have the following results analogous to those corresponding to the class of non-idling policies. However, they are limited to in-forests that consist exclusively of type 3 nodes whose service times are exponential r.v.'s.

Theorem 4.4 *Assume that $\mathcal{K}$ is an in-forest consisting exclusively of type 3 nodes. If for any fixed sequence of arrival times $\mathcal{A} = \{a_n\}_{n=1}^{N}$, the due dates are stochastically comparable in the sense of $\leq_{st}$, then the SSDD and SLDD policies are extremal with respect to the latenesses:*

$$\forall \pi \in \Psi: \qquad \boldsymbol{L}(SSDD) \leq_{E_3^{\uparrow}} \boldsymbol{L}(\pi) \leq_{E_3^{\downarrow}} \boldsymbol{L}(SLDD).$$

Theorem 4.5 *Assume that $\mathcal{K}$ is an in-forest consisting solely of type 3 nodes. If for any fixed sequence of arrival times $\mathcal{A} = \{a_n\}_{n=1}^N$, the due dates are mutually independent and are stochastically comparable in the sense of $\leq_a \in \{\leq_{lr}, \leq_{hr}, \leq_{st}\}$, then the SSDD and the SLDD policies are extremal with respect to the latenesses:*

$$\forall \pi \in \Psi : \qquad \boldsymbol{L}(SSDD) \leq_{b\uparrow} \boldsymbol{L}(\pi) \leq_{b\downarrow} \boldsymbol{L}(SLDD),$$

where $(a, b) \in \{(lr, E_1), (hr, E_{5-icx}), (st, E_5)\}$.

Corollary 4.3 *Assume that $\mathcal{K}$ is an in-forest consisting solely of type 3 nodes. Then the FCFS and LCFS policies are extremal with respect to the response times:*

$$\forall \pi \in \Psi : \qquad \boldsymbol{R}(FCFS) \leq_{E_1^\uparrow} \boldsymbol{R}(\pi) \leq_{E_1^\downarrow} \boldsymbol{R}(LCFS).$$

Corollary 4.4 *Assume that $\mathcal{K}$ is an in-forest consisting solely of type 3 nodes. If the due dates are unknown a priori, and if the relative due dates are i.i.d. random variables which are independent of the arrival times, Then the FCFS and LCFS policies are extremal with respect to the latenesses:*

$$\forall \pi \in \Psi : \qquad \boldsymbol{L}(FCFS) \leq_{E_2^\uparrow} \boldsymbol{L}(\pi) \leq_{E_2^\downarrow} \boldsymbol{L}(LCFS).$$

Theorem 4.6 *Assume that $\mathcal{K}$ is an in-forest consisting of type 3 nodes alone. Then the FCFS policy is extremal with respect to the end-to-end delays:*

$$\forall \pi \in \Psi : \qquad \boldsymbol{D}(FCFS) \leq_{st} \boldsymbol{D}(\pi).$$

5 Proofs of the Main Results

We begin this section by considering the scheduling problem in a single node. We derive properties that allow us to propagate permutation orderings from an input mapping to an output mapping for each of the three types of nodes defined earlier. These properties are then used in the second subsection to prove the main results.

5.1 Scheduling in a Single Node

Consider a single-queue multi-server model which is either type 1 or type 2 or type 3. The customers of this queue form a subset $\mathcal{N}_0$ of the set of customers in the network $\mathcal{N}_0$, i.e., $\mathcal{N}_0 \subseteq \mathcal{N}$. Customer $i \in \mathcal{N}_0$ arrives at the queue at (random) time $\tilde{a}_i$. The sequence of arrival times $\{\tilde{a}_i\}_{i \in \mathcal{N}_0}$ is disordered, i.e., it is not necessarily increasing in i. For any fixed $\{\tilde{a}_i\}_{i \in \mathcal{N}_0}$, there is a permutation γ on $\mathcal{N}_0$ such that the sequence $\{\tilde{a}_{\gamma(i)}\}_{i \in \mathcal{N}_0}$ is increasing in i. The permutation γ is referred to as the input mapping of the queue.

Let π be an arbitrary scheduling policy applied to that queue. Denote by σ_i^π the service time of customer $i \in \mathcal{N}_0$ under policy π. The random variables σ_i^π, $i \in \mathcal{N}_0$, are i.i.d. having distribution B which is independent of π. The service times are independent of the arrival times $\{\tilde{a}_i\}_{i\in\mathcal{N}_0}$, the input mapping γ and the due dates $\{d_i\}_{i\in\mathcal{N}_0} \equiv \boldsymbol{d}_{\mathcal{N}_0}$. Let s_i^π and t_i^π be the times when customer $i \in \mathcal{N}_0$ starts and completes its service at one of the servers under π, respectively. Let ψ_π^γ and θ_π^γ be the permutations on $\mathcal{N}_0$ such that $\{s^\pi_{\psi_\pi^\gamma(i)}\}_{i\in\mathcal{N}_0}$ and $\{t^\pi_{\theta_\pi^\gamma(i)}\}_{i\in\mathcal{N}_0}$ are increasing in i. The permutation ψ_π^γ and θ_π^γ are referred to as the scheduling mapping and the output mapping of the queue under policy π, respectively.

Lemma 5.1 *Let the arrival times $\{\tilde{a}_i\}_{i\in\mathcal{N}_0}$ be fixed. Assume there are two input mappings γ and γ'. If $\gamma' \prec^a_{\boldsymbol{d}_{\mathcal{N}_0}} \gamma$ for some $a \in \{lr, hr, st\}$, then, for all $\pi \in \Psi_{ni}$, there exists a probability space such that*

$$\{t^\pi_{\theta^{\gamma'}_{SSDD}(i)}\}_{i\in\mathcal{N}_0} = \{t^\pi_{\theta^\gamma_\pi(i)}\}_{i\in\mathcal{N}_0} \quad a.s. \qquad and \qquad \theta^{\gamma'}_{SSDD} \prec^a_{\boldsymbol{d}_{\mathcal{N}_0}} \theta^\gamma_\pi \quad a.s. \tag{5.1}$$

There also exists a probability space such that

$$\{t^\pi_{\theta^\gamma_{SLDD}(i)}\}_{i\in\mathcal{N}_0} = \{t^\pi_{\theta^{\gamma'}_\pi(i)}\}_{i\in\mathcal{N}_0} \quad a.s. \qquad and \qquad \theta^{\gamma'}_\pi \prec^a_{\boldsymbol{d}_{\mathcal{N}_0}} \theta^\gamma_{SLDD} \quad a.s. \tag{5.2}$$

Proof. We consider a type 2 node first. We will focus on relation (5.1). The proof of (5.2) can be shown in an analogous way.

For notational simplification, we assume, without loss of generality, that $\mathcal{N}_0 = \mathcal{N} = \{1, 2, \cdots, N\}$, and $\boldsymbol{d}_{\mathcal{N}_0} = \boldsymbol{d}$. Thus the mapping $\gamma(i)$ (resp. $\psi^\gamma_\pi(i)$, $\theta^\gamma_\pi(i)$) can be interpreted as the index of the i-th arrived (resp. scheduled, completed) customer.

The proof is based on a sample path interchange argument. From an arbitrary policy π defined on γ, we will construct a (finite) series of policies such that the final policy is SSDD defined on γ' and that each new policy improves the previous one in the sense of output mapping. Each policy defines a system with customer arrival times, service times, scheduling times and completion times. We will construct these systems on a common probability space in such a way that the arrival times in all these systems are coupled. The service times of a system will be defined as a function of the scheduling decisions and service times in the previous system.

More precisely, given the arrival times and the input mapping γ, we fix the service times and compute the completion times of the customers under policy π. We will describe a procedure of assigning service times to the customers under SSDD so that *i)* they are the same in law as under π, *ii)* they generate the same completion times for this sample path, and *iii)* $\theta^{\gamma'}_{SSDD} \prec^a_{\boldsymbol{d}} \theta^\gamma_\pi$. This will be done in two steps.

Claim 1: *For the given input mapping γ, there exists a service time assignment such that*

$$\{t^\pi_{\theta^\gamma_{SSDD}(i)}\}_{i\in\mathcal{N}} = \{t^\pi_{\theta^\gamma_\pi(i)}\}_{i\in\mathcal{N}} \quad a.s. \qquad and \qquad \theta^\gamma_{SSDD} \prec^a_{\boldsymbol{d}} \theta^\gamma_\pi \quad a.s. \tag{5.3}$$

Proof of Claim 1. We use an interchange argument to prove (5.3) where we modify π one scheduling decision at a time until SSDD is produced.

Consider the first scheduling point at which π deviates from SSDD. Let this be scheduling decision m. Suppose that π schedules customer j whereas there is another customer k in the queue such that $\psi_\pi^\gamma(m) = j$ and $\psi_\pi^\gamma(n) = k$, $m < n$ and $d_j \geq_a d_k$. We will construct a new policy ρ which differs from π only in that the scheduling decisions of j and k are switched:

$$\psi_\rho^\gamma(m) = k, \quad \psi_\rho^\gamma(n) = j, \qquad \psi_\rho^\gamma(l) = \psi_\pi^\gamma(l), \ \ \forall l \neq m, n.$$

Assign the same service times of customers i, $i \neq j, k$, under ρ as under π: $\sigma_i^\rho = \sigma_i^\pi$. Let us focus on the service times that customers j and k will be assigned under ρ. We will use the property that the service time distribution is ILR to construct the service times σ_j^ρ and σ_k^ρ in such a way that the completion times of j and k are either switched or not. Thus, the sequences of scheduling times and completion times under ρ, $\{t^\rho_{\theta_\rho^\gamma(i)}\}_{i\in\mathcal{N}}$ and $\{s^\rho_{\psi_\rho^\gamma(i)}\}_{i\in\mathcal{N}}$, are identical to those of π, $\{t^\pi_{\theta_\pi^\gamma(i)}\}_{i\in\mathcal{N}}$ and $\{s^\pi_{\psi_\pi^\gamma(i)}\}_{i\in\mathcal{N}}$. Moreover, the completion times of j and k will never be switched if this will create a situation where k departs after j under ρ and j departs after k under π. In order to do so, the two customers may not receive the same amount of service times under ρ as under π.

Assume that under π, customers j and k are scheduled at times s_j and s_k and that they complete at times t_j and t_k. Let $\Delta = s_k - s_j > 0$. Figure 2 illustrates several possibilities.

Assume that the service times are continuous random variables (the discrete case can be analyzed analogously). Denote by $f_\sigma(x)$ the density function of service time σ at point x. We define the service times for j and k under ρ as follows,

$$\begin{aligned}(\sigma_k^\rho, \sigma_j^\rho) \;=\;& \mathbf{1}(t_j \leq s_k)(\sigma_j^\pi, \sigma_k^\pi) \\ &+ \mathbf{1}(s_k < t_j \leq t_k)\left[U(\sigma_j^\pi, \sigma_k^\pi, \Delta)(\sigma_j^\pi, \sigma_k^\pi) + (1 - U(\sigma_j^\pi, \sigma_k^\pi, \Delta))(\Delta + \sigma_k^\pi, \sigma_j^\pi - \Delta)\right] \\ &+ \mathbf{1}(t_k < t_j)(\Delta + \sigma_k^\pi, \sigma_j^\pi - \Delta) \qquad (5.4)\end{aligned}$$

where $U(a, b, \Delta)$ is a Bernoulli r.v. with probability distribution $\Pr[U(a, b, \Delta) = 1] = p(a, b, \Delta)$, $\Pr[U(a, b, \Delta) = 0] = 1 - p(a, b, \Delta)$, where

$$\begin{aligned}p(a, b, \Delta) &= 1 - \frac{f_{\sigma|\sigma>\Delta}(b + \Delta) f_\sigma(a - \Delta)}{f_{\sigma|\sigma>\Delta}(a) f_\sigma(b)} \\ &= 1 - \frac{f_\sigma(b + \Delta) f_\sigma(a - \Delta)}{f_\sigma(a) f_\sigma(b)}. \qquad (5.5)\end{aligned}$$

This is defined only for the case that $a \geq \Delta > 0$. Furthermore the ILR assumption guarantees that $0 \leq p(a, b, \Delta) \leq 1$.

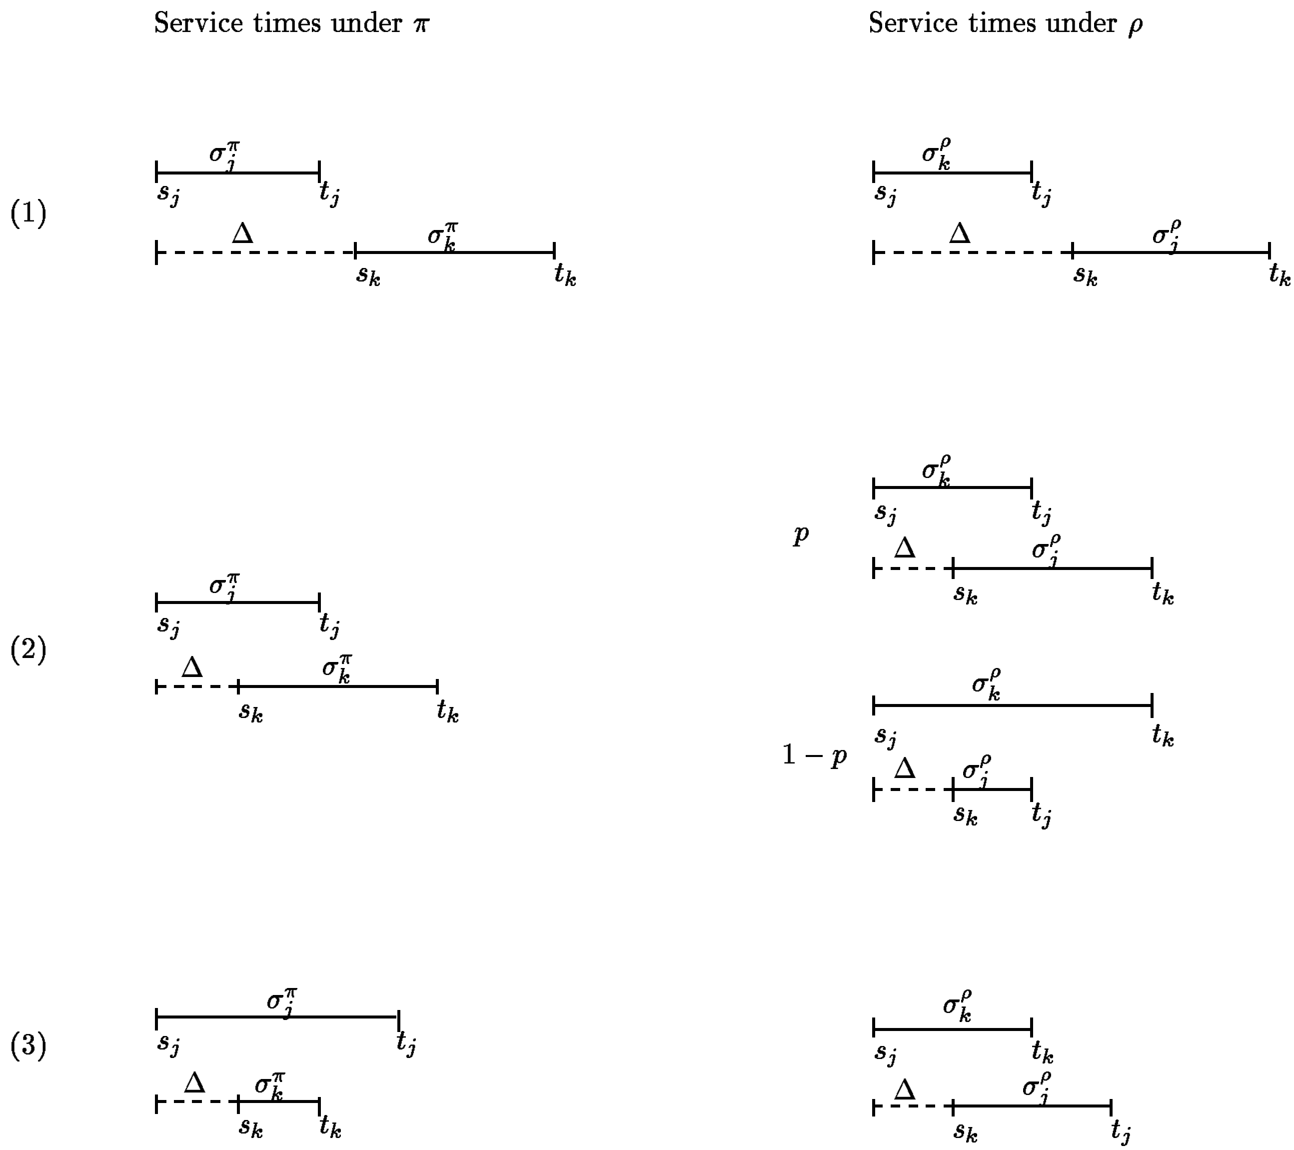

Figure 2: Construction of Service Times under Policy ρ

It is easy to see that this construction either switches the completion times of j and k or not. Therefore, the sequences of the scheduling times and of the completion times of ρ are unchanged. Moreover, customer k departs after j under ρ only if k departs after j under π. Hence, if the order of the completions of the j and k under ρ is the same as under π, then $\theta_\rho^\gamma = \theta_\pi^\gamma$. Otherwise, $\theta_\rho^\gamma(j) = \theta_\pi^\gamma(k)$ and $\theta_\rho^\gamma(k) = \theta_\pi^\gamma(j)$. In any case, we have

$$\theta_\rho^\gamma \prec_{\mathbf{d}}^a \theta_\pi^\gamma.$$

We now need to show that the random variables in $\{\sigma_i^\rho\}_{i \in \mathcal{N}}$ are i.i.d. This is done by evaluating the joint density function for the service times $(\sigma_k^\rho, \sigma_j^\rho)$. Clearly, in the case that $x \le \Delta$ (cf. case (1) in Figure 2), $f_{\sigma_k^\rho, \sigma_j^\rho}(x,y) = f_\sigma(x) f_\sigma(y)$. In the case that $\Delta < x \le y + \Delta$, we could have $\sigma_k^\rho = x$ and $\sigma_j^\rho = y$ if either $\sigma_j^\pi = y + \Delta$ and $\sigma_k^\pi = x - \Delta$ (cf. case (3) in Figure 2) or if $\sigma_j^\rho = x$ and $\sigma_k^\rho = y$ and $U(a,b,\Delta) = 1$ (cf. first subcase of (2) in Figure 2). Thus,

$$\begin{aligned} f_{\sigma_k^\rho, \sigma_j^\rho}(x,y) &= f_\sigma(x) f_\sigma(y) p(y,x,\Delta) + f_\sigma(\Delta + y) f_\sigma(x - \Delta) \\ &= f_\sigma(x) f_\sigma(y) \left[1 - \frac{f_\sigma(\Delta + y) f_\sigma(x - \Delta)}{f_\sigma(x) f_\sigma(y)} \right] + f_\sigma(\Delta + y) f_\sigma(x - \Delta) \\ &= f_\sigma(x) f_\sigma(y). \end{aligned}$$

Finally, in the case of $y + \Delta < x$, to obtain $\sigma_k^\rho = x$ and $\sigma_j^\rho = y$, it must be that $\sigma_j^\pi = x - \Delta$ and $\sigma_k^\pi = y + \Delta$ and $U(a,b,\Delta) = 0$ (cf. second subcase of (2) in Figure 2), so that

$$\begin{aligned} f_{\sigma_k^\rho, \sigma_j^\rho}(x,y) &= f_\sigma(\Delta + y) f_\sigma(x - \Delta) \frac{f_\sigma(\Delta + y) f_\sigma(x - \Delta)}{f_\sigma(x) f_\sigma(y)} \\ &= f_\sigma(x) f_\sigma(y). \end{aligned}$$

The necessity that the service times have a distribution with ILR should be obvious from this construction. Therefore we conclude that for all $(x,y) \in I\!R^{+2}$,

$$f_{\sigma_j^\rho, \sigma_k^\rho}(x,y) = f_\sigma(x) f_\sigma(y) = f_{\sigma_j^\pi, \sigma_k^\pi}(x,y).$$

The i.i.d. assumption on the random variables in $\{\sigma_i^\pi\}_{i \in \mathcal{N}}$ readily implies that the random variables in $\{\sigma_i^\rho\}_{i \in \mathcal{N}}$ are i.i.d.

This interchange argument can be repeated until SSDD is produced for which we have relation (5.3). ∎

Claim 2. *Given two input mappings γ' and γ with $\gamma' \prec_{\mathbf{d}}^a \gamma$, there exists a service time assignment such that*

$$\{t^{SSDD}_{\theta^{\gamma'}_{SSDD}(i)}\}_{i \in \mathcal{N}} = \{t^{SSDD}_{\theta^{\gamma}_{SSDD}(i)}\}_{i \in \mathcal{N}} \quad a.s. \qquad and \qquad \theta^{\gamma'}_{SSDD} \prec_{\mathbf{d}}^a \theta^{\gamma}_{SSDD} \quad a.s. \tag{5.6}$$

Proof of Claim 2. If the input mappings γ and γ' are the same, then we are done. We assume that γ and γ' differ only in two positions. The general case can be shown by iterating the interchange procedure that will be performed later in the section. Let j and k be the customers whose arrival times are switched in the inputs γ and γ':

$$d_j \geq_a d_k, \quad \gamma(j) < \gamma(k), \quad \gamma'(j) = \gamma(k), \quad \gamma'(k) = \gamma(j), \quad \gamma'(i) = \gamma(i), \;\; i \neq j, \; i \neq k.$$

Fix the service times of policy SSDD for the input mapping γ and determine the sequences of scheduling times and completion times of the SSDD policy associated with γ: $\{s^{SSDD}_{\psi^{\gamma}_{SSDD}(i)}\}_{i\in\mathcal{N}}$ and $\{t^{SSDD}_{\theta^{\gamma}_{SSDD}(i)}\}_{i\in\mathcal{N}}$.

Let ρ be a policy associated with input mapping γ'. Assume first that $\psi^{\gamma}_{SSDD}(j) > \psi^{\gamma}_{SSDD}(k)$. Define $\psi^{\gamma'}_{\rho} = \psi^{\gamma}_{SSDD}$ and let $\sigma^{\rho}_i = \sigma^{SSDD}_i$ for all $i \in \mathcal{N}$. Clearly ρ is feasible (in that a customer is never served before its arrival) with respect to the input mapping γ'. Moreover,

$$\{s^{\rho}_{\psi^{\gamma'}_{\rho}(i)}\}_{i\in\mathcal{N}} = \{s^{SSDD}_{\psi^{\gamma}_{SSDD}(i)}\}_{i\in\mathcal{N}}, \qquad \{t^{\rho}_{\theta^{\gamma'}_{\rho}(i)}\}_{i\in\mathcal{N}} = \{t^{SSDD}_{\theta^{\gamma}_{SSDD}(i)}\}_{i\in\mathcal{N}}, \qquad \theta^{\gamma'}_{\rho} \prec^{a}_{\mathbf{d}} \theta^{\gamma}_{SSDD} \tag{5.7}$$

(actually, $\theta^{\gamma'}_{\rho} = \theta^{\gamma}_{SSDD}$).

Assume now that $\psi^{\gamma}_{SSDD}(j) < \psi^{\gamma}_{SSDD}(k)$. Define

$$\psi^{\gamma'}_{\rho}(j) = \psi^{\gamma}_{SSDD}(k), \quad \psi^{\gamma'}_{\rho}(k) = \psi^{\gamma}_{SSDD}(j), \quad \psi^{\gamma'}_{\rho}(i) = \psi^{\gamma}_{SSDD}(i), \quad \forall i \in \mathcal{N} - \{j,k\}.$$

The service times under ρ are defined to be identical to those under SSDD associated with γ for customers other than j and k: $\sigma^{\rho}_i = \sigma^{SSDD}_i$ for all $i \in \mathcal{N} - \{j,k\}$. The service times of customers j and k under ρ is defined by (5.4) with $\pi = SSDD$. where s_j, s_k, t_j, t_k are the scheduling times and completion times of customers j, k, respectively, under policy SSDD associated with input mapping γ, and where $\Delta = s_k - s_j$.

As seen in the proof of Claim 1, under such a construction, ρ is a feasible policy with respect to the input mapping γ' and (5.7) holds.

Consequently, in both cases policy ρ is a feasible with respect to input mapping γ' and, using the arguments of the proof of Claim 1, we can show that there is a service time assignment such that for the input mapping γ',

$$\{s^{SSDD}_{\psi^{\gamma'}_{SSDD}(i)}\}_{i\in\mathcal{N}} = \{s^{\rho}_{\psi^{\gamma'}_{\rho}(i)}\}_{i\in\mathcal{N}}, \qquad \{t^{SSDD}_{\theta^{\gamma'}_{SSDD}(i)}\}_{i\in\mathcal{N}} = \{t^{\rho}_{\theta^{\gamma'}_{\rho}(i)}\}_{i\in\mathcal{N}}, \qquad \theta^{\gamma'}_{SSDD} \prec^{a}_{\mathbf{d}} \theta^{\gamma'}_{\rho}. \tag{5.8}$$

The proof of the claim is concluded by combining relations (5.7) and (5.8). ■

The assertion of the lemma for a type 2 node is a consequence of the Claims 1 and 2.

The proof for type 3 node is analogous. The detailed proof can be found in [18] and is omitted here. In the case of a type 1 node, the proof is simpler since the assignment of service times to customers j and k under the constructed policy ρ is simply an interchange of their service times under π. Hence, the service times can be arbitrarily distributed r.v.'s and the proof proceeds in a similar way. ■

5.2 Proofs of the Main Results of the Non-Idling Case

Proof of Theorem 4.1. As in the proof of Lemma 5.1, we use a sample path interchange argument to prove the theorem. From an arbitrary policy π, we construct a (finite) series of policies such that the final policy is SSDD and each new policy improves the previous one in the sense of the permutation orderings on the output mappings on each node. We will consider these policies on a common probability space in such a way that the external arrival times in all these systems are coupled. The service times of a system will be defined as a function of the scheduling decisions and service times in the previous system.

For the given policy π, let $a_n^j(\pi)$ and $c_n^i(\pi)$ be the arrival and completion times of customer $n \in \mathcal{N}_i$ at node i. Let γ_π^i and θ_π^i be the input and output mappings of node i, $1 \le i \le K$.

Fix the customer arrival times $a_1, a_2, \cdots, a_N$, and the service times at all the nodes. Determine for each node i, $1 \le i \le K$, the customer arrival times $\{a_n^i(\pi)\}_{n\in\mathcal{N}_i}$, the input mapping γ_π^i, the customer completion times $\{c_n^i(\pi)\}_{n\in\mathcal{N}_i}$, and the output mapping θ_π^i.

We will show by induction on i that there is a service time assignment such that for all $1 \le i \le K$,

$$\gamma_{SSDD}^i \prec_{\boldsymbol{d}_{\mathcal{N}_i}}^{st} \gamma_\pi^i \;\; a.s., \qquad \{a^i_{\gamma^i_{SSDD}(n)}(SSDD)\}_{n\in\mathcal{N}_i} = \{a^i_{\gamma^i_\pi(n)}(\pi)\}_{n\in\mathcal{N}_i} \;\; a.s., \tag{5.9}$$

$$\theta_{SSDD}^i \prec_{\boldsymbol{d}_{\mathcal{N}_i}}^{st} \theta_\pi^i \;\; a.s., \qquad \{c^i_{\theta^i_{SSDD}(n)}(SSDD)\}_{n\in\mathcal{N}_i} = \{c^i_{\theta^i_\pi(n)}(\pi)\}_{n\in\mathcal{N}_i} \;\; a.s. \tag{5.10}$$

Recall that the nodes in the network are labeled in such a way that if node i is a predecessor of node j, then $i < j$.

Consider node $i = 1$, which is necessarily a leaf node. It is clear that

$$\gamma_{SSDD}^i = \gamma_\pi^i, \qquad \{a^i_{\gamma^i_{SSDD}(n)}(SSDD)\}_{n\in\mathcal{N}_i} = \{a^i_{\gamma^i_\pi(n)}(\pi)\}_{n\in\mathcal{N}_i},$$

so that (5.9) holds for $i = 1$. Appealing to Lemma 5.1 implies that (5.10) also holds for $i = 1$.

Assume that for some j, $2 \le j \le K$, relations (5.9) and (5.10) hold for all $i < j$. Since the sequence of arrival times of node j is the superposition of the sequences of the completion times of the predecessor nodes of j, we obtain from the inductive assumption and Lemma 3.4 that

$$\gamma_{SSDD}^j \prec_{\boldsymbol{d}_{\mathcal{N}_j}}^{st} \gamma_\pi^j, \qquad \{a^j_{\gamma^j_{SSDD}(n)}(SSDD)\}_{n\in\mathcal{N}_j} = \{a^j_{\gamma^j_\pi(n)}(\pi)\}_{n\in\mathcal{N}_j},$$

so that (5.9) holds for $i = j$. Applying Lemma 5.1 yields that (5.10) also holds for $i = j$, which completes the proof of relations (5.9) and (5.10).

Applying Lemma 3.2 yields

$$E[f(\boldsymbol{L}(SSDD))|\mathcal{A}] \leq E[f(\boldsymbol{L}(\pi))|\mathcal{A}], \quad \forall f \in \mathcal{C}_3.$$

Removing the conditioning on the arrival times implies $\boldsymbol{L}(SSDD) \leq_{E_3} \boldsymbol{L}(\pi)$. The proof of the inequality $\boldsymbol{L}(\pi) \leq_{E_3} \boldsymbol{L}(SLDD)$ is analogous. ∎

Proof of Theorem 4.2. The proof is similar to that of Theorem 4.1, and uses Lemma 3.3. ∎

Proof of Theorem 4.3. Let $q_n(\pi)$ denote the departure time of customer n under $\pi \in \Psi_{ni}$, i.e.,

$$q_n(\pi) = \max(c_n(\pi), q_{n-1}(\pi)) = \max_{1 \leq l \leq n} c_l(\pi),$$

and let $\boldsymbol{q}(\pi) = (q_1(\pi), \ldots, q_N(\pi))$. Fix the arrival times and replace the due dates in (5.10) of the proof of Theorem 4.1 with the arrival times. According to Lemma 3.1, $\boldsymbol{q}(FCFS) \leq \boldsymbol{q}(\pi)$ *a.s.*, or $\boldsymbol{D}(FCFS) \leq \boldsymbol{D}(\pi)$ *a.s.* Removing the conditioning on the arrival times and the service times yields $\boldsymbol{D}(FCFS) \leq_{st} \boldsymbol{D}(\pi)$. The proof of the relation $\boldsymbol{D}(\pi) \leq_{st} \boldsymbol{D}(LCFS)$ is similar, and is omitted. ∎

5.3 Proof of Idling Policy Results

In this subsection we focus on in-forests that consist solely of type 3 nodes so that the service times are exponential r.v.'s. In order to establish the results for idling policies, we present Lemma 5.2 below. The interested reader is referred to [18] for a detailed proof.

Lemma 5.2 *Assume that $\mathcal{K}$ is an in-forest consisting solely of type 3 nodes. For each policy $\pi \in \Psi$, there is a non-idling policy $\rho \in \Psi_{ni}$ such that $(c_1(\rho), \cdots, c_N(\rho)) \leq_{st} (c_1(\pi), \cdots, c_N(\pi))$, $\boldsymbol{L}(\rho) \leq_{st} \boldsymbol{L}(\pi)$ and $\boldsymbol{R}(\rho) \leq_{st} \boldsymbol{R}(\pi)$.*

It is now simple to see that Theorems 4.4, 4.5, and 4.6 are immediate consequences of Lemma 5.2 and the results of Section 4.1. The detailed proofs are omitted.

6 Concluding Remarks

In this paper, we studied the scheduling problems in in-forests with identical multiple servers at each node. We focused on the following performance measures: customer response time, lateness, and end-to-end delay. Various extremal properties were established for several simple policies using stochastic ordering techniques.

Another performance measure of interest is *customer tardiness* which is defined as $T_n = \max(0, L_n)$ for customer n. Observe that for $x \in I\!R$, the function $f(x) = \max(0, x)$ is increasing and convex, and that the composition of an increasing and (Schur) convex function with such a function f is still increasing and (Schur) convex. Moreover, it is readily checked that for all $h \in \mathcal{C}_5^{\uparrow}$, the composition $h \circ f$ is still in $\mathcal{C}_5^{\uparrow}$, where $f(x) = \max(0, x)$. Therefore, the stochastic orderings $\leq_{E_i}$ and $\leq_{E_i^{\uparrow}}$ obtained above on the vectors of latenesses $\boldsymbol{L}$ imply the stochastic orderings $\leq_{E_i^{\uparrow}}$ on the vectors of tardinesses $\boldsymbol{T} = (T_1, \cdots, T_N)$, $i = 1, 2, 3, 5, 5$-icx.

Although the paper has focussed on transient results, it should be clear that similar extremal properties hold in the stationary regime. In particular, the stochastic orderings $\leq_{E_3}$, $\leq_{E_3^{\uparrow}}$ and $\leq_{E_3^{\downarrow}}$ (recall that $\leq_{E_1^{(\uparrow,\downarrow)}} \Rightarrow \leq_{E_2^{(\uparrow,\downarrow)}} \Rightarrow \leq_{E_3^{(\uparrow,\downarrow)}}$ and $\leq_{E_{5-icx}^{(\uparrow,\downarrow)}} \Rightarrow \leq_{E_5^{(\uparrow,\downarrow)}} \Rightarrow \leq_{E_3^{(\uparrow,\downarrow)}}$), established in this paper on the transient measures $\boldsymbol{R}(\pi)$ and $\boldsymbol{L}(\pi)$ reduce to the stochastic orderings $\leq_{cx}$, $\leq_{icx}$ and $\leq_{dcx}$ on the corresponding stationary random variables $R(\pi)$ and $L(\pi)$, provided that the sequences of random variables, $R_n(\pi)$ and $L_n(\pi)$ weakly converge to $R(\pi)$ and $L(\pi)$, for the classes of convex functions, increasing convex functions, and decreasing convex functions, respectively, in the Cesaro sense (cf. Feller [9, p. 249]). Similarly, the stochastic ordering, $\leq_{st}$, established on the transient measure $\boldsymbol{D}(\pi)$ reduces to a $\leq_{st}$ ordering on the stationary random variable $D(\pi)$ under the weak convergence assumption. The interested reader is referred to [2] for detailed proofs.

Most of the results obtained in this paper are new except for a single node. In this case, Theorem 4.2 was established in [6] for a type 1 node, and Theorem 4.4 was established in [17] for a type 3 node. Moreover, it was shown in [17] that (4.1) and a slightly weaker relation than (4.5) still hold for a single multi-server queue $G/GI/s$ ($s \geq 1$). For a single $G/IFR/1$ queue, where IFR stands for Increasing Failure Rate service distribution, Hirayama and Kijima [12] obtained that, within the class of *preemptive* policies, the FCFS policy is optimal for the $G/IFR/1$ queue:

$$\forall \pi \in \Psi : \qquad \boldsymbol{R}(FCFS) \leq_{E_3^{\uparrow}} \boldsymbol{R}(\pi). \tag{6.1}$$

A slightly stronger ordering was established in Chang and Yao [6]. However, such an ordering does not appear to generalize easily to other queueing models with IFR service times.

First, relation (6.1) does not hold for arbitrary multi-server queue with IFR service times, namely the $G/IFR/s$ ($s \geq 2$) model when service preemption is permitted, as illustrated by the following counterexample. Consider a $G/D/3$ queue with 10 arriving customers. The service times are all 5 and the arrival times are $0, 1, 2, 3, 4, 5, 6, 14, 14.001, 14.002$. Under FCFS, the average response time is 6.0998. The following preemptive schedule gives an average response time of 6.0. Server 1 serves customer 1 at $t = 0$, customer 2 at $t = 5$, customer 5 at $t = 7$, and customer 8 at $t = 14$. Server 2 serves customer 2 at $t = 1$, customer 4 at $t = 4$, customer 7 at $t = 9$, and customer

9 at $t = 14.001$. Server 3 serves customer 3 at $t = 2$, customer 6 at $t = 7$, and customer 10 at $t = 14.002$. Note that customer 2 was preempted at $t = 4$.

Relation (6.1) does not hold for tandem queues with single IFR servers either. Righter and Shanthikumar [21] provided a counterexample for two single-server queues in tandem. The first queue has an IFR service and the second one is deterministic. They constructed a policy such that the completion times at the second queue are stochastically smaller than those under FCFS policy. Therefore, the mean response time is smaller under that policy than under FCFS.

Nevertheless, if ILR service times (which are more restrictive than those of IFR) are assumed, we can prove the following:

Theorem 6.1 *In any tandem queueing network consisting of $\cdot/ILR/1$ queues, FCFS policy applied to all the queues is optimal within the class of preemptive policies in the sense that*

$$\forall \pi \in \Psi : \qquad \boldsymbol{R}(FCFS) \leq_{E_1^{\uparrow}} \boldsymbol{R}(\pi).$$

The above theorem follows from the facts that FCFS stochastically minimizes the vector of completion times in each queue [21] and that FCFS minimizes the input and output mappings of each queue in the sense of the permutation ordering. The detailed proof can be obtained by combining our arguments and those in [21].

Acknowledgments: The authors are grateful to Professor Rhonda Righter and the referee for their various constructive comments on both the contents and the style of the paper.

References

[1] F. Baccelli, E. Gelenbe, B. Plateau, "An End-to-End Approach to the Resequencing Problem", *Journal of the ACM*, Vol. 31, pp. 474-485, 1984.

[2] F. Baccelli, Z. Liu, D. Towsley, "Extremal Scheduling of Parallel Processing with and without Real-Time Constraints", to appear in the *Journal of the ACM.*

[3] F. Baccelli, A. M. Makowski, "Queueing Models for Systems with Synchronization Constraints", *Proceedings of the IEEE*, Vol. 77, Special Issue on Dynamics of Discrete Event Systems, pp. 138-161, Jan. 1989.

[4] R. Barlow, F. Proschan, *Statistical Theory of Reliability and Life Testing.* Holt, Rinehart and Winston, 1975.

[5] A. Brandt, P. Franken, B. Lisek, *Stationary Stochastic Models.* Akademic-Verlag, 1989.

[6] C. S. Chang, D. D. Yao, "Rearrangement, Majorization and Stochastic Scheduling", IBM Research Report RC 16250, 1990.

[7] P.W. Day, "Rearrangement Inequalities," *Can. J. Math.*, Vol. 24, 5, pp. 930-943, 1972.

[8] D. J. Daley, "Certain Optimality Properties of the First Come First Served Discipline for $G/G/s$ Queues", *Stochastic Processes and their Applications*, Vol. 25, pp. 301-308, 1987.

[9] W. Feller, *An Introduction to Probability Theory and Its Applications*, Vol. II, Second Edition, John Wiley & Sons, 1971.

[10] S. G. Foss, "Approximation of Multichannel Queueing Systems" (in Russian), *Sibirski Mat. Zh.*, Vol. 21, pp. 132-140, 1980. (Transl.: *Siberian Math. J.*, Vol. 21, pp. 851-857, 1980.)

[11] S. G. Foss, "Comparison of Servicing Strategies in Multichannel Queueing Systems" (in Russian), *Sibirski Math. Zh.*, Vol. 22, pp. 190-197, 1981. (Transl.: *Siberian Math. J.*, Vol. 22, pp. 142-147, 1981.)

[12] T. Hirayama, M. Kijima, "An Extremal Property of FIFO Discipline in $G/IFR/1$ Queues", *Adv. Appl. Prob.*, Vol. 21, 481-484, 1989.

[13] I. Iliadis, L. Lien, "Resequencing Delay for a Queueing System with Two Heterogeneous Servers Under a Threshold-Type Scheduling", *IEEE Trans. on Communications*, Vol. 36, pp. 692-702, 1988.

[14] J. F. C. Kingman, "Inequalities in the Theory of Queues," *J. Roy. Stat. Soc.*, ser. B, Vol. 32, pp. 102-110, 1970.

[15] L. Kleinrock, F. Kamoun, R. Muntz, "Queueing Analysis of the reordering issue in a distributed database concurrency control mechanism", *Proc. of the 2nd International Conference on Distributed Computing Systems*, Versailles, France, 1981.

[16] Z. Liu, P. Nain, "Optimal Scheduling in Some Multi-Queue Single-Server Systems." *IEEE Transactions on Automatic Control*, Vol. 37, pp. 247–252, 1992.

[17] Z. Liu, D. Towsley, "Effects of Service Disciplines in $G/G/s$ Queueing Systems", to appear in the *Annals of the Operations Research*, special issue on queueing networks.

[18] Z. Liu, D. Towsley, "Stochastic Scheduling in In-Forest Networks". COINS Technical Report, TR 92-27, 1992, INRIA Rapport de Recherche, No. 1719, 1992.

[19] A. W. Marshall, I. Olkin, *Inequalities: Theory of Majorization and Its Applications*, Academic Press, 1979.

[20] M. Pinedo, "Stochastic Scheduling with Release Dates and Due Dates." *Oper. Res.*, Vol. 31, No. 3, pp. 559-572, May-June 1983.

[21] R. Righter, J. G. Shanthikumar, "Extremal Properties of the FIFO Discipline in Queueing Networks", manuscript.

[22] S. Ross, *Stochastic Processes*, J. Wiley & Sons, 1983.

[23] J. G. Shantikumar, U. Sumita, "Convex Ordering of Sojourn Times in Single-Server Queues: Extremal Properties of FIFO and LIFO Service Disciplines." *J. Appl. Prob.*, Vol. 24, pp. 737-748, 1987.

[24] D. Stoyan, *Comparison Methods for Queues and Other Stochastic Models.* English translation (D.J. Daley editor), J. Wiley and Sons, New York, 1983.

[25] D. Towsley, F. Baccelli, "Comparison of Service Disciplines in a Tandem Queueing Network with Delay Dependent Customer Behavior". *Operations Research Letters*, **10**, 49-55, Feb. 1991.

[26] D. Towsley, S.S. Panwar, "Optimality of the Stochastic Earliest Deadline Policy for the $G/M/c$ Queue Serving Customers with Deadlines," COINS Technical Report 91-61, Univ. Massachusetts, Aug. 1991.

[27] O. A. Vasicek, "An Inequality for the Variance of Waiting Time Under a General Queueing Discipline." *Operations Research*, Vol. 25, pp. 879-884, 1977.

[28] W. Whitt, "The Amount of Overtaking in a Network of Queues," *Networks*, Vol. 14, pp. 411-426, 1984.

[29] R. W. Wolff, "An Upper Bound for Multi-Channel Queues," *J. Appl. Prob.*, Vol. 14, pp. 884-888, 1977.

[30] R. W. Wolff, "Upper Bounds on Work in System for Multichannel Queues," *J. Appl. Prob.*, Vol. 24, pp. 547-551, 1987.