UNITÉ DE RECHERCHE
INRIA-SOPHIA ANTIPOLIS

Institut National
de Recherche
en Informatique
et en Automatique

-Domaine de Voluceau
Rocquencourt
B.P.105
78153 Le Chesnay Cedex
France
Tél.: (1) 39 63 55 11

Rapports de Recherche

1 9 9 2

ème

anniversaire

N° 1719

Programme 1

*Architectures parallèles, Bases de données,
Réseaux et Systèmes distribués*

STOCHASTIC SCHEDULING IN IN-FOREST NETWORKS

Zhen LIU
Don TOWSLEY

Juillet 1992

★ R R - 1 7 1 9 ★

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.

April 6, 1992

*This work was supported in part by the National Science Foundation under grant ASC 88-8802764

Stochastic Scheduling in In-Forest Networks

Abstract

In this paper we study the extremal properties of several scheduling policies in a queueing network consisting of multi-server queues where the topology is that of an in-forest. Associated with each customer is a due date. 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 provided 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 be exponentially distributed random variables.

Using stochastic majorization and a partial order on permutations, we establish that, within the class of work conserving service policies, the smallest due date (SDD) and largest due date (LDD) policies minimize and maximize, respectively, the vector of the customer latenesses of the first n customers in the sense of the Schur-convex order. 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. If due dates are only known stochastically, then we establish that the stochastic SDD and the stochastic LDD policies minimize and maximize the vector of latenesses according to a weaker ordering. 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. When idling policies are under consideration, the FCFS policy is shown to be optimal for various performance metrics in the case that all queues permit preemptive resume policies and all service times are exponentially distributed. Extensions of these results to the stationary regime are also given.

Keywords : Queueing System, Service Discipline, Stochastic Ordering, Sample Path Analysis, majorization, permutation ordering, Lateness, Response Time, End-to-End Delay.

Ordonnancement stochastique dans les réseaux arborescents

Résumé

Dans cet article nous étudions les propriétés extrémales de plusieurs politiques d'ordonnancement dans un réseau de files constituées de multiples serveurs dont la topologie est une arborescence (*in-forest*). Une date d'échéance est associée à chaque client. Nous supposons que les durées de service sur les différentes files forment des suites mutuellement indépendantes des variables aléatoires indépendantes et identiquement distribuées, et qu'elles sont indépendantes des dates d'arrivées et des dates d'échéance. Le réseau peut avoir différents types de nœuds. Certains d'entre eux sont limités à des disciplines de service non préemptives, tandis que d'autres permettent la préemption de service. Dans le cas de service non préemptif, les temps de service peuvent avoir une distribution générale s'il y a un seul serveur dans la file, et une distribution ILR s'il y a plusieurs serveurs. Dans le cas de service préemptif, les temps de service sont restreints à la distribution exponentielle.

En utilisant la Schur majoration et un ordre partiel sur les permutations, nous établissons que dans la classe des politiques conservatives, la politique Première Echéance d'Abord (PEA) (resp. la politique Dernière Echéance d'Abord (DEA)) minimise (resp. maximise) le vecteur des retards des n premiers clients dans le sens de l'ordre Schur convexe. Nous en concluons que la politique Premier Arrivé Premier Servi (PAPS) (resp. la politique Dernier Arrivé Premier Servi (DAPS)) minimise (resp. maximise) le vecteur des temps de réponses des n premiers clients dans le sens de l'ordre Schur convexe. Nous prouvons aussi que la politique PAPS (resp. DAPS) minimise (resp. maximise) le vecteur des délais de bout en bout dans le sens de l'ordre stochastique fort. Quand les politiques sont non conservatives, nous démontrons que la politique PAPS est optimale pour diverses mesures de performances, pourvu que les temps de service soient exponentiellement distribués. Des extensions de ces résultats au régime stationnaire sont aussi données.

Mots-Clés : Files d'attente, Discipline de service, Ordre stochastique, Analyse trajectorielle, Schur majoration. Retard, Temps de réponse, Délai de bout en bout.

1 Introduction

Consider a network of $G/GI/s$ 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, also known as soft real-time deadlines, associated with the customers.

In this paper we study the effect of different service policies on customer response time, end-to-end delay (defined as the difference between the maximum completion time of the customer and all customers that arrived prior to the customer in question and its arrival time), 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 [26], Foss [10, 11], Wolff [28, 29] and Daley [8]. When service times have an Erlang distribution, preemptions are allowed and there is a single server $s = 1$, Shantikumar and Sumita [21] 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 job lateness. The optimality of the Shortest Due Date first (SDD) policy was first established in [24] 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 a stochastic version of SDD has been established for the $G/GI/s$ queue [17] and for the $G/M/s$ queue when customers have hard real-time deadlines [25].

Another important property of a queueing system is overtaking and resequencing (cf. Kleinrock et al. [15], Whitt [27], Baccelli et al. [1, 3]). Whitt [27] 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 con-

sisting of GI/s queues with delay dependent customer behavior. We assume that each node falls into one of the following categories.

- the node has one server, does not allow service preemptions, and the associated service times have an arbitrary distribution,
- the node has two or more servers, does not allow preemptions, and the associated service times are increasing in likelihood ratio (ILR), and
- the node has one or more servers, allows service preemptions, and the service times are exponentially distributed random variables.

Under these assumptions, we establish

- When due dates are known, the smallest due date (SDD) and largest due date (LDD) policies minimize and maximize, respectively, the vector of customer latenesses in the sense of the Schur-convex ordering. If the due dates are set to the arrival times, this translates into the result that first come first serve (FCFS) and last come first serve (LCFS) respectively minimize and maximize the vector of response times in the sense of Schur-convex order.
- If due dates are only known stochastically, then we establish that the stochastic SDD (SSDD) and the stochastic LDD (SLDD) policies minimize and maximize the vector of latenesses according to a weaker ordering that will be defined later.
- Last, we establish that FCFS and LCFS minimize and maximize, respectively, the vector of end-to-end delays in the sense of stochastic ordering.

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 metrics 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.

Figure 1: Example of an in-forest.

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 Sections 5 and 6. A discussion of the stationary regime and other generalizations as well as counterexamples are provided in Section 7.

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 job executed at a non-root node will later be executed 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.

We distinguish between three types of nodes.

- a *type 1* node contains a single server which does not allow preemptions,
- a *type 2* node which contains two or more servers which do not allow preemptions, and
- a *type 3* node which contains one or more servers which do allow preemptions. Service is always resumed from the point of preemption.

The n -th customer, also referred to as customer n , arrives in the system at time a_n , $n = 1, 2, \dots$, $a_1 < a_2 < \dots < a_n < \dots$. When a customer arrives in the system, it enters one of the leaf queues. When a customer enters a queue, it waits for service at 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 $1, 2, \dots, n-1$ 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 $\mathcal{N} = \{1, 2, \dots, N\}$ 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 , $i \in \mathcal{K}$.

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.

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 service time information 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 $\Psi_{ni} \subset \Psi$ the class of non-idling policies. Among the well-known extremal policies, there are First Come First Serve (FCFS) and Last Come First Serve (LCFS) policies. Consider node i . If it is either of type 1 or type 2, then FCFS (resp. LCFS) behaves in the following way. When there is an available server in a queue, FCFS (resp. LCFS) schedules the oldest customer, i.e., the customer with the smallest arrival time to the system (resp., the youngest customer, i.e., the customer with the largest arrival time to the system). If the node is of type 3, then preemptions are allowed. Thus, FCFS (resp. LCFS) ensures that the oldest (resp. youngest) customers are always in service. Note that the FCFS and LCFS policies defined above 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. Note also that the FCFS and LCFS policies thus defined are non-idling.

When the customers are associated with the known due dates, there exist the Smallest Due Date first (SDD) and the Largest Due Date first (LDD) policies, which assign customers according to their due dates. Whether or not these policies permit preemptions depends on the type of node.

When the due dates are comparable in the stochastic ordering sense $\leq_{st}$ (see the definition 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, we define the Stochastically Smallest Due Date first (SSDD) policies and Stochastically Largest Due Date first (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 date are known, the SSDD and SLDD policies coincide with the SDD and LDD policies, respectively.

2.3 Statistical Assumptions

Throughout this paper, we will assume that

- 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 whose distribution may be:
 - arbitrary if the queue is of type 1,
 - increasing in likelihood ratio (cf. Section 3.2) if the queue is of type 2, and

– exponential if the queue is of type 3.

- 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.

2.4 Performance Metrics

Let $\pi \in \Psi$ be an arbitrary scheduling policy. Denote by $b_n(\pi)$ the random variable (in $\mathbb{R}^+$) of the time of the service commencement of customer n at one of the leaf nodes, and $c_n(\pi)$ the random variable (in $\mathbb{R}^+$) 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

$$R_n(\pi) = c_n(\pi) - a_n, \quad (2.1)$$

$$L_n(\pi) = c_n(\pi) - d_n. \quad (2.2)$$

Denote by $D_n(\pi)$ the end-to-end delay of customer n under $\pi \in \Psi$, i.e.,

$$D_n(\pi) = \max_{1 \leq l \leq n} c_l(\pi) - a_n. \quad (2.3)$$

Let

$$\mathbf{R}(\pi) = (R_1(\pi), \dots, R_N(\pi)),$$

$$\mathbf{L}(\pi) = (L_1(\pi), \dots, L_N(\pi)),$$

$$\mathbf{D}(\pi) = (D_1(\pi), \dots, D_N(\pi)).$$

The purpose of this paper is to find extremal policies that minimize or maximize these performance measures in some stochastic semi-partial ordering sense. Hence, we define the partial orderings of interest to us in the next section.

3 Preliminaries on Stochastic Orderings

Throughout this paper, the inequality $\leq$ between two vectors is understood to be component-wise. Increasingness and decreasingness stand for non-decreasingness and non-increasingness, respectively.

The first two subsections review some concepts in stochastic orderings. The last subsection develops some properties associated with permutation orderings which will be crucial in establishing the main results of the paper.

3.1 Integral Orderings

Let $\mathbf{X}, \mathbf{Y} \in \mathbb{R}^n$ be two random vectors.

Definition 3.1 *The random vector $\mathbf{X}$ is stochastically smaller than the random vector $\mathbf{Y}$ in the sense of*

*strong stochastic ordering ($\mathbf{X} \leq_{st} \mathbf{Y}$), iff $E[f(\mathbf{X})] \leq E[f(\mathbf{Y})]$, $\forall$ increasing $f : \mathbb{R}^n \rightarrow \mathbb{R}$,
convex ordering ($\mathbf{X} \leq_{cx} \mathbf{Y}$), iff $E[f(\mathbf{X})] \leq E[f(\mathbf{Y})]$, $\forall$ convex $f : \mathbb{R}^n \rightarrow \mathbb{R}$,
increasing convex ordering ($\mathbf{X} \leq_{icx} \mathbf{Y}$), iff $E[f(\mathbf{X})] \leq E[f(\mathbf{Y})]$, $\forall$ increasing convex $f : \mathbb{R}^n \rightarrow \mathbb{R}$,
decreasing convex ordering ($\mathbf{X} \leq_{dcx} \mathbf{Y}$), iff $E[f(\mathbf{X})] \leq E[f(\mathbf{Y})]$, $\forall$ decreasing convex $f : \mathbb{R}^n \rightarrow \mathbb{R}$,*

respectively, provided the expectations exist.

In what follows, $=_{st}$ denotes equality in distribution. The following lemma is due to Strassen [23]:

Lemma 3.1 *Two random vectors $\mathbf{X}$ and $\mathbf{Y}$ satisfy $\mathbf{X} \leq_{st} \mathbf{Y}$ if and only if there exist two random vectors $\tilde{\mathbf{X}}$ and $\tilde{\mathbf{Y}}$ defined on a common probability space such that $\mathbf{X} =_{st} \tilde{\mathbf{X}}$, $\mathbf{Y} =_{st} \tilde{\mathbf{Y}}$, and $\tilde{\mathbf{X}} \leq \tilde{\mathbf{Y}}$ componentwise almost surely (a.s.).*

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

Definition 3.2 *Vector $\mathbf{x}$ is said to be majorized by vector $\mathbf{y}$ (written $\mathbf{x} \prec \mathbf{y}$) iff*

$$\sum_{i=1}^k x_{[i]} \leq \sum_{i=1}^k y_{[i]}, \quad k = 1, \dots, n-1, \quad (3.1)$$

$$\sum_{i=1}^n x_{[i]} = \sum_{i=1}^n y_{[i]}, \quad (3.2)$$

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

Definition 3.3 *A function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is Schur-convex if for all $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$,*

$$\mathbf{x} \prec \mathbf{y} \Rightarrow f(\mathbf{x}) \leq f(\mathbf{y}).$$

We define the following classes of functions

- $C_1 (C_1^\uparrow, C_1^\downarrow)$ — the class of (increasing, decreasing) Schur-convex functions,
- $C_2 (C_2^\uparrow, C_2^\downarrow)$ — the class of (increasing, decreasing) symmetric and convex functions,
- $C_3 (C_3^\uparrow, C_3^\downarrow)$ — the class of functions of the form $f(\mathbf{X}) = \sum_{i=1}^n f(x_i)$ where f is (increasing, decreasing) convex.

Definition 3.4 Let $\mathbf{X}$ and $\mathbf{Y}$ be two random vectors in $\mathbb{R}^n$. We define the following stochastic orderings between these r.v.'s

$$\begin{aligned} \mathbf{X} \leq_E \mathbf{Y}, \quad & \text{iff} \quad E[f(\mathbf{X})] \leq E[f(\mathbf{Y})], \quad \forall f \in C_i, \quad i = 1, 2, 3, \\ \mathbf{X} \leq_{E_i^\uparrow} \mathbf{Y}, \quad & \text{iff} \quad E[f(\mathbf{X})] \leq E[f(\mathbf{Y})], \quad \forall f \in C_i^\uparrow, \quad i = 1, 2, 3, \\ \mathbf{X} \leq_{E_i^\downarrow} \mathbf{Y}, \quad & \text{iff} \quad E[f(\mathbf{X})] \leq E[f(\mathbf{Y})], \quad \forall f \in C_i^\downarrow, \quad i = 1, 2, 3, \end{aligned}$$

provided the expectations exist.

Various properties concerning these orderings can be found in [18]. From the definition of these orderings and the fact that any symmetric and convex function is Schur-convex (cf. [18, Proposition C.2, p. 67]), we can easily see that

Lemma 3.2 The following implication relations between the stochastic orderings hold:

$$\begin{array}{ccccc} \geq_{st} & \Leftarrow & =_{st} & \Rightarrow & \leq_{st} \\ \Downarrow & & \Downarrow & & \Downarrow \\ \leq_{E_1^\downarrow} & \Leftarrow & \leq_{E_1} & \Rightarrow & \leq_{E_1^\uparrow} \\ \Downarrow & & \Downarrow & & \Downarrow \\ \leq_{E_2^\downarrow} & \Leftarrow & \leq_{E_2} & \Rightarrow & \leq_{E_2^\uparrow} \\ \Downarrow & & \Downarrow & & \Downarrow \\ \leq_{E_3^\downarrow} & \Leftarrow & \leq_{E_3} & \Rightarrow & \leq_{E_3^\uparrow} \end{array}$$

The relations still hold when $\leq_{E_1}$, $\leq_{E_1^\uparrow}$, and $\leq_{E_1^\downarrow}$ are replaced by $\leq_{cx}$, $\leq_{icx}$, and $\leq_{dcx}$ respectively.

The following lemma is another application of Strassen's Theorem (cf. [23]) to the semi-partial ordering $\prec$. A proof was provided in [18, Theorem B.1, p. 483].

Lemma 3.3 *Two random vectors X and Y satisfy $X \leq_{E_1} Y$ if and only if there exist two random variables $\widetilde{X}$ and $\widetilde{Y}$ defined on a common probability space such that $X =_{st} \widetilde{X}$, $Y =_{st} \widetilde{Y}$, and $\widetilde{X} \prec \widetilde{Y}$ a.s.*

3.2 Likelihood Ratio Ordering

A number of our results will require that service times have distributions with increasing likelihood ratio (ILR). Consequently, we briefly introduce this class of random variables before presenting the main results of the paper. Let $X, Y \in \mathbb{R}^+$ be two continuous nonnegative random variables with density functions f_X and f_Y respectively.

Definition 3.5 *The random variable X is smaller than the random variable Y in the sense of likelihood ratio ($X \leq_{lr} Y$) 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 ordering, i.e., $X \leq_{lr} Y \Rightarrow X \leq_{st} Y$.

Definition 3.6 *The random variable $X \in \mathbb{R}^+$ is said to be increasing in likelihood ratio (ILR) if*

$$c_1 + X \leq_{lr} c_2 + X, \quad 0 \leq c_1 \leq c_2,$$

or equivalently

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

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: $f(x) = \alpha \lambda (\lambda x)^{\alpha-1} e^{-(\lambda x)^\alpha}$, $\alpha \geq 1$.

The likelihood ratio ordering can also be defined for discrete random variables that are defined over the same set of values. We say that $X \leq_{lr} Y$ if $P(Y = x)/P(X = x)$ increases in x .

3.3 Partial Orderings on Permutations

Let $n \geq 1$ be an arbitrary integer, and Γ the set of permutations on $\{1, 2, \dots, n\}$. For any two real vectors $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$, denote by $\mathbf{y}_\gamma = (y_{\gamma(1)}, \dots, y_{\gamma(n)})$ where $\gamma \in \Gamma$, and $\mathbf{y} - \mathbf{x} = (y_1 - x_1, \dots, y_n - x_n)$.

We define a partial ordering on Γ which was first introduced in Baccelli, Liu and Towsley [2]. Assume that $\{x_i\}_{i=1}^n$ is a sequence of mutually different real numbers. Define the binary relation $B_{\{x_i\}}$ on the symmetric group Γ as:

Definition 3.7 $\gamma' B_{\{x_i\}} \gamma$ if $\gamma' = \gamma$ or if there exist a pair of integers j, k , such that

$$x_j < x_k, \quad \gamma(j) > \gamma(k), \quad \gamma'(j) = \gamma(k), \quad \gamma'(k) = \gamma(j), \quad \gamma'(i) = \gamma(i), \quad i \neq j, i \neq k. \quad (3.3)$$

Define now a partial order $\prec_{\{x_i\}}$ on Γ as the transitive closure of $B_{\{x_i\}}$:

1. $\gamma' \prec_{\{x_i\}} \gamma$ if $\gamma' B_{\{x_i\}} \gamma$.
2. $\gamma' \prec_{\{x_i\}} \gamma$ if there exists γ'' such that $\gamma' \prec_{\{x_i\}} \gamma''$ and $\gamma'' \prec_{\{x_i\}} \gamma$.

Note that $\prec_{\{x_i\}}$ and $\prec_{\{x'_i\}}$ define the same partial order provided

$$\forall i, j : x_i < x_j \quad \text{iff} \quad x'_i < x'_j.$$

Lemma 3.4 (Baccelli, Liu and Towsley [2]) Assume that $y_1 \leq y_2 \leq \dots \leq y_n$. If $\gamma' \prec_{\{x_i\}} \gamma$, then

$$(\mathbf{y}_{\gamma'} - \mathbf{x}) \prec (\mathbf{y}_\gamma - \mathbf{x}). \quad (3.4)$$

Lemma 3.5 Let $\mathbf{Y} \in \mathbb{R}^n$ be a random vector such that

$$Y_1 \leq Y_2 \leq \dots \leq Y_n, \quad \text{a.s.}$$

If, on a common probability space, $\gamma' \prec_{\{x_i\}} \gamma$, a.s., then

$$(\mathbf{Y}_{\gamma'} - \mathbf{x}) \leq_{E_1} (\mathbf{Y}_\gamma - \mathbf{x}). \quad (3.5)$$

Proof. Applying Lemmas 3.4 and 3.3 immediately yields the desired inequality. ■

Lemma 3.6 Assume that $x_1 \leq x_2 \leq \dots \leq x_n$ and that $y_1 \leq y_2 \leq \dots \leq y_n$. If $\gamma' \prec_{\{x_i\}} \gamma$, then

$$\max_{1 \leq i \leq m} y_{\gamma'(i)} \leq \max_{1 \leq i \leq m} y_{\gamma(i)}, \quad 1 \leq m \leq n. \quad (3.6)$$

Proof. We need only consider the case that γ' differs from γ in two positions. Specifically, assume that

$$x_j < x_k, \quad \gamma(j) > \gamma(k), \quad \gamma'(j) = \gamma(k), \quad \gamma'(k) = \gamma(j), \quad \gamma'(i) = \gamma(i), \quad i \neq j, i \neq k.$$

Note that under the assumption of this lemma, $x_j < x_k$ iff $j < k$. Thus,

$$\begin{aligned} \forall 1 \leq m \leq j-1: \quad & \max_{1 \leq i \leq m} y_{\gamma'(i)} = \max_{1 \leq i \leq m} y_{\gamma(i)}, \\ \forall k \leq m \leq n: \quad & \max_{1 \leq i \leq m} y_{\gamma'(i)} = \max_{1 \leq i \leq m} y_{\gamma(i)}, \\ \forall j \leq m \leq k-1: \quad & \max_{1 \leq i \leq m} y_{\gamma'(i)} = \max \left(y_{\gamma'(j)}, \max_{1 \leq i \leq m, i \neq j} y_{\gamma'(i)} \right) \\ & = \max \left(y_{\gamma(k)}, \max_{1 \leq i \leq m, i \neq j} y_{\gamma(i)} \right) \\ & \leq \max \left(y_{\gamma(j)}, \max_{1 \leq i \leq m, i \neq j} y_{\gamma(i)} \right) \\ & = \max_{1 \leq i \leq m} y_{\gamma(i)}. \end{aligned}$$

■

Lemma 3.7 Assume that $x_1 \leq x_2 \leq \dots \leq x_n$. Let $Y \in \mathbb{R}^n$ be a random vector such that

$$Y_1 \leq Y_2 \leq \dots \leq Y_n, \quad \text{a.s.}$$

If, on a common probability space, $\gamma' \prec_{\{x_i\}} \gamma$, a.s., then for all $1 \leq m \leq n$,

$$\left\{ \max_{1 \leq i \leq m} Y_{\gamma'(i)} \right\}_{m=1}^n \leq_{st} \left\{ \max_{1 \leq i \leq m} Y_{\gamma(i)} \right\}_{m=1}^n. \quad (3.7)$$

Proof. The result follows from Lemmas 3.6 and 3.1. ■

Define now another partial ordering on Γ with respect to a sequence of random variables. Assume that $\{X_i\}_{i=1}^n$ is a sequence of stochastically comparable random variables, i.e., for all

$1 \leq i, j \leq n$, either $X_i \leq_{st} X_j$ or $X_j \leq_{st} X_i$. Define the binary relation $B_{\{X_i\}}^{st}$ on the symmetric group Γ as: $\gamma' B_{\{X_i\}}^{st} \gamma$ if there exist a pair of integers j, k , such that

$$X_j \leq_{st} X_k, \quad \gamma(j) > \gamma(k), \quad \gamma'(j) = \gamma(k), \quad \gamma'(k) = \gamma(j), \quad \gamma'(i) = \gamma(i), \quad i \neq j, i \neq k. \quad (3.8)$$

Define now a partial order $\prec_{\{X_i\}}^{st}$ on Γ as the transitive closure of $B_{\{X_i\}}^{st}$:

1. $\gamma' \prec_{\{X_i\}}^{st} \gamma$ if $\gamma' B_{\{X_i\}}^{st} \gamma$.
2. $\gamma' \prec_{\{X_i\}}^{st} \gamma$ if there exists γ'' such that $\gamma' \prec_{\{X_i\}}^{st} \gamma''$ and $\gamma'' \prec_{\{X_i\}}^{st} \gamma$.

Note that $\prec_{\{X_i\}}^{st}$ and $\prec_{\{X'_i\}}^{st}$ define the same partial order provided

$$\forall i, j : X_i \leq_{st} X_j \quad \text{iff} \quad X'_i \leq_{st} X'_j.$$

Lemma 3.8 *Let $Y \in IR^n$ be a random vector such that*

$$Y_1 \leq Y_2 \leq \dots \leq Y_n, \quad a.s.$$

Assume further that $\{X_i\}_{i=1}^n$ is a sequence of stochastically comparable random variables, and that $\{X_i\}$ is independent of Y and of the random permutations $\gamma', \gamma \in \Gamma$. If $\gamma' \prec_{\{X_i\}}^{st} \gamma$, then

$$(Y_{\gamma'} - X) \leq_{E_3} (Y_{\gamma} - X). \quad (3.9)$$

Proof. It suffices to consider the case $\gamma' B_{\{X_i\}}^{st} \gamma$. The general case can be shown using transitivity of $\leq_{E_3}$. Let j, k be the pair of integers fulfilling (3.8):

$$X_j \leq_{st} X_k, \quad \gamma(j) > \gamma(k), \quad \gamma'(j) = \gamma(k), \quad \gamma'(k) = \gamma(j), \quad \gamma'(i) = \gamma(i), \quad i \neq j, i \neq k.$$

Applying Strassen's Theorem to the random variables X_j, X_k entails that there are two random variables $\hat{X}_j$ and $\hat{X}_k$ on a probability space such that

$$\hat{X}_j =_{st} X_j, \quad \hat{X}_k =_{st} X_k, \quad \text{and} \quad \hat{X}_j \leq \hat{X}_k \quad a.s.,$$

which implies that

$$\gamma' B_{\{\hat{X}_i\}} \gamma \quad a.s.$$

Since X is independent of Y and γ', γ , we have on the same probability space that $Y_1 \leq Y_2 \leq \dots \leq Y_n$, a.s.. Owing to Lemma 3.4, we obtain that

$$(Y_{\gamma'} - \widehat{X}) \prec (Y_{\gamma} - \widehat{X}), \quad \text{a.s.},$$

which implies that for all convex function $f : \mathbb{R} \rightarrow \mathbb{R}$,

$$\sum_{i=1}^n f(Y_{\gamma'(i)} - \widehat{X}_i) \leq \sum_{i=1}^n f(Y_{\gamma(i)} - \widehat{X}_i) \quad \text{a.s.}$$

Thus,

$$\sum_{i=1}^n Ef(Y_{\gamma'(i)} - X_i) = \sum_{i=1}^n Ef(Y_{\gamma'(i)} - \widehat{X}_i) \leq \sum_{i=1}^n Ef(Y_{\gamma(i)} - \widehat{X}_i) = \sum_{i=1}^n Ef(Y_{\gamma(i)} - X_i).$$

■

Consider now the partial ordering on the “merging” of the permutations. Let N_1, N_2 be a partition of $\{1, 2, \dots, n\}$:

$$N_1 \cup N_2 = \{1, 2, \dots, n\}, \quad N_1 \cap N_2 = \emptyset.$$

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 3.9 *Let $\{x_i\}_{i=1}^n$ be a sequence of mutually different real numbers. Let $\gamma_1, \gamma'_1 \in \Gamma_1$, $\gamma_2, \gamma'_2 \in \Gamma_2$. If $\gamma'_1 \prec_{\{x_i\}_{i \in N_1}} \gamma_1$ and $\gamma'_2 \prec_{\{x_i\}_{i \in N_2}} \gamma_2$, then*

$$\gamma' = (\gamma'_1, \gamma'_2) \prec_{\{x_i\}_{i=1}^n} (\gamma_1, \gamma_2) = \gamma.$$

Proof.

$$(\gamma'_1, \gamma'_2) \prec_{\{x_i\}_{i=1}^n} (\gamma_1, \gamma_2) \prec_{\{x_i\}_{i=1}^n} (\gamma_1, \gamma_2).$$

■

Lemma 3.10 *Let $\{X_i\}_{i=1}^n$ be a sequence of stochastically comparable random variables. Let $\gamma_1, \gamma'_1 \in \Gamma_1$, $\gamma_2, \gamma'_2 \in \Gamma_2$. If $\gamma'_1 \prec_{\{X_i\}_{i \in N_1}}^{st} \gamma_1$ and $\gamma'_2 \prec_{\{X_i\}_{i \in N_2}}^{st} \gamma_2$, then*

$$\gamma' = (\gamma'_1, \gamma'_2) \prec_{\{X_i\}_{i=1}^n}^{st} (\gamma_1, \gamma_2) = \gamma.$$

Proof.

$$(\gamma'_1, \gamma'_2) \prec_{\{X_i\}_{i=1}^n}^{st} (\gamma_1, \gamma'_2) \prec_{\{X_i\}_{i=1}^n}^{st} (\gamma_1, \gamma_2).$$

■

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 two sections. 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, SDD, LDD, 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. Two cases will be investigated depending on the information on due dates available to the scheduler:

1. Exact due dates are known;
2. Exact due dates are not necessarily known, but the due dates are stochastically comparable in the sense of $\leq_{st}$, viz., for any $m, n \geq 1$, either $d_m \leq_{st} d_n$ or $d_n \leq_{st} d_m$.

Theorem 4.1 *If the due dates are known, then the SDD (resp. LDD) policy applied to all the nodes minimizes (resp. maximizes) the vector of latenesses within the class Ψ_{ni} in the sense of $\leq_{E_1}$ ordering,*

$$\forall \pi \in \Psi_{ni} : \quad L(SDD) \leq_{E_1} L(\pi) \leq_{E_1} L(LDD). \quad (4.1)$$

Theorem 4.2 *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.

$$\forall \pi \in \Psi_{ni} : \quad L(SSDD) \leq_{E_3} L(\pi) \leq_{E_3} L(SLDD). \quad (4.2)$$

Remark: If the due dates are unknown a priori, and if the relative due dates are stochastically increasing random variables (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 (4.2) still holds when SSDD and SLDD are replaced by FCFS and LCFS, respectively.

Setting the due dates to the arrival times in Theorem 4.1 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,*

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

Owing to Corollary 4.1, a stronger relation than (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, $L(\pi)$ can be rewritten as

$$L(\pi) = (R_1(\pi) - u_1, R_2(\pi) - u_2, \dots, R_N(\pi) - u_N),$$

where $u_1, \dots, u_N$ 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 [18, Proposition F.6, p. 314] and making use of Lemma 3.2 and 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,*

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

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 $\leq_{st}$,*

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

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 of type 3 nodes. If the due dates are known, then the SDD and LDD policies are extremal with respect to the latenesses:*

$$\forall \pi \in \Psi : \quad L(SDD) \leq_{E_1^1} L(\pi) \leq_{E_1^1} L(LDD).$$

Theorem 4.5 *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:*

$$\forall \pi \in \Psi : \quad L(SSDD) \leq_{E_3^1} L(\pi) \leq_{E_3^1} L(SLDD).$$

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:*

$$\forall \pi \in \Psi : \quad R(FCFS) \leq_{E_1^1} R(\pi) \leq_{E_1^1} R(LCFS).$$

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:*

$$\forall \pi \in \Psi : \quad L(FCFS) \leq_{E_2^1} L(\pi) \leq_{E_2^1} L(LCFS).$$

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:*

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

5 Proofs of the Results for Non-Idling Policies

We begin this section by considering the scheduling problem at 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 at the end of the section to prove the main results for non-idling policies.

5.1 Scheduling in a Single Node

Consider a single-queue multi-server model. The customers of this queue form a subset $\mathcal{N}_0$ of the set of customers in the network $\mathcal{N}$, i.e., $\mathcal{N}_0 \subseteq \mathcal{N}$. Customer $i \in \mathcal{N}_0$ arrives at the queue at (random) time $\bar{a}_i$. The sequence of arrival times $\{\bar{a}_i\}_{i \in \mathcal{N}_0}$ is disordered, i.e., it is not necessarily increasing in i . For any fixed $\{\bar{a}_i\}_{i \in \mathcal{N}_0}$, there is a permutation γ on $\mathcal{N}_0$ such that the sequence $\{\bar{a}_{\gamma(i)}\}_{i \in \mathcal{N}_0}$ 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 . 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 $\{s_{\psi_\pi^\gamma(i)}^\pi\}_{i \in \mathcal{N}_0}$ and $\{t_{\theta_\pi^\gamma(i)}^\pi\}_{i \in \mathcal{N}_0}$ 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.

In order to establish these properties in the case of a type 3 node, we rely on the following lemma.

Lemma 5.1 *Consider a type 3 node with i.i.d service times with an exponential distribution. For any function of the random variables R , L , and D , there exists an optimal policy that minimizes or maximizes the expectation of this function whose decision points at each node occur only at the arrival times of that node and all but the last service completion time of that node.*

In other words, preemptions and new customer assignments at each type 3 node occur only at the instants of customer arrivals to that node and at the instants of all but the last service completion of that node. This fact results from the memoryless property of the exponential distribution. Indeed, between these instants the state represented by the existing customers and their remaining service times does not change. The formal proof is left to the interested reader. Throughout this section, we will confine ourselves to scheduling policies exhibiting the property stated in Lemma 5.1.

Lemma 5.2 *Let the arrival times $\{\bar{a}_i\}_{i \in \mathcal{N}_0}$ and the due dates $\{d_i\}_{i \in \mathcal{N}_0}$ be fixed. Assume there are two input mappings γ and γ' . Assume further that the service times are i.i.d. random variables having a distribution that depends on the type of node,*

- arbitrary distribution for a type 1 node,

- an ILR distribution for a type 2 node,
- an exponential distribution for a type 3 node.

If $\gamma' \prec_{\{d_i\}_{i \in \mathcal{N}_0}} \gamma$, then, for all $\pi \in \Psi_{ni}$, there exists a probability space such that

$$\{t_{\theta_{SDD}^\pi(i)}^\pi\}_{i \in \mathcal{N}_0} = \{t_{\theta_\pi^\pi(i)}^\pi\}_{i \in \mathcal{N}_0} \text{ a.s.} \quad \text{and} \quad \theta_{SDD}^{\gamma'} \prec_{\{d_i\}_{i \in \mathcal{N}_0}} \theta_\pi^\gamma \text{ a.s.} \quad (5.1)$$

There also exists a probability space such that

$$\{t_{\theta_{LDD}^\pi(i)}^\pi\}_{i \in \mathcal{N}_0} = \{t_{\theta_\pi^{\gamma'}(i)}^\pi\}_{i \in \mathcal{N}_0} \text{ a.s.} \quad \text{and} \quad \theta_\pi^{\gamma'} \prec_{\{d_i\}_{i \in \mathcal{N}_0}} \theta_{LDD}^\gamma \text{ a.s.} \quad (5.2)$$

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 $\mathcal{N}_0 = \mathcal{N} = \{1, 2, \dots, N\}$. Thus the mapping $\gamma(i)$ (resp. $\psi_\pi^\gamma(i)$, $\theta_\pi^\gamma(i)$) can be interpreted as the index of the $\gamma(i)$ -th arrived (resp. $\psi_\pi^\gamma(i)$ -th scheduled, $\theta_\pi^\gamma(i)$ -th 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 SDD 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 SDD 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) $\theta_{SDD}^{\gamma'} \prec_{\{d_i\}_{i \in \mathcal{N}}} \theta_\pi^\gamma$. This will be done in two steps.

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

$$\{t_{\theta_{SDD}^\pi(i)}^\pi\}_{i \in \mathcal{N}} = \{t_{\theta_\pi^\pi(i)}^\pi\}_{i \in \mathcal{N}} \text{ a.s.} \quad \text{and} \quad \theta_{SDD}^{\gamma'} \prec_{\{d_i\}_{i \in \mathcal{N}}} \theta_\pi^\gamma \text{ a.s.} \quad (5.3)$$

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

Consider the first scheduling point at which π deviates from SDD. 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 $d_j > d_k$. We will construct a new policy ρ which differs from π only in that the scheduling of j and k are switched:

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

Assign the same service times of customers i , $i \neq j, k$, under ρ as under π :

$$\sigma_i^\rho = \sigma_i^\pi, \quad \forall i \neq j, k.$$

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 , $\{t_{\psi_\rho^\gamma(i)}^\rho\}_{i \in \mathcal{N}}$ and $\{s_{\psi_\rho^\gamma(i)}^\rho\}_{i \in \mathcal{N}}$, are identical to those of π , $\{t_{\psi_\pi^\gamma(i)}^\pi\}_{i \in \mathcal{N}}$ and $\{s_{\psi_\pi^\gamma(i)}^\pi\}_{i \in \mathcal{N}}$. 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,

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

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

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

Service times under π

Service times under ρ

Figure 2: Construction of Service Times under Policy ρ

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

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

$$\theta_\rho^\gamma \prec_{\{d_i\}_{i \in \mathcal{N}}} \theta_\pi^\gamma.$$

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 $(\sigma_j^\rho, \sigma_k^\rho)$. In the case that $x \leq \Delta$

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

In the case that $\Delta < x \leq y + \Delta$ we have

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

Finally, in the case of $y + \Delta < x$,

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

The necessity that the service times have a distribution with ILR should be obvious from this construction.

Therefore we conclude that for all $(x, y) \in \mathbb{R}^{+2}$,

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

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 SDD is produced for which we have relation (5.3). ■

Claim 2. Given two input mappings γ' and γ with $\gamma' \prec_{\{d_i\}_{i \in \mathcal{N}}} \gamma$, then, there exists a service time assignment such that

$$\{t_{\theta_{SDD}^{\gamma'}(i)}^{SDD}\}_{i \in \mathcal{N}} = \{t_{\theta_{SDD}^{\gamma}(i)}^{SDD}\}_{i \in \mathcal{N}} \text{ a.s.} \quad \text{and} \quad \theta_{SDD}^{\gamma'} \prec_{\{d_i\}_{i \in \mathcal{N}}} \theta_{SDD}^{\gamma} \text{ a.s.} \quad (5.6)$$

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 proof. Let j and k be the customers whose arrival times are switched in the inputs γ and γ' :

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

Fix the service times of policy SDD for the input mapping γ and determine the sequences of scheduling times and completion times of the SDD policy associated with γ : $\{s_{\psi_{SDD}^{\gamma}(i)}^{SDD}\}_{i \in \mathcal{N}}$ and $\{t_{\theta_{SDD}^{\gamma}(i)}^{SDD}\}_{i \in \mathcal{N}}$.

Let ρ be a policy associated with input mapping γ' . Assume first that $\psi_{SDD}^{\gamma}(j) > \psi_{SDD}^{\gamma}(k)$. Define $\psi_{\rho}^{\gamma'} = \psi_{SDD}^{\gamma}$. The service times under ρ are defined to be identical to that of SDD associated with γ : $\sigma_i^{\rho} = \sigma_i^{SDD}$ for all $i \in \mathcal{N}$. It is easy to see that ρ is a feasible policy (in the sense that a customer is never served before its arrival) with respect to the input mapping γ' . Moreover,

$$\{s_{\psi_{\rho}^{\gamma'}(i)}^{\rho}\}_{i \in \mathcal{N}} = \{s_{\psi_{SDD}^{\gamma}(i)}^{SDD}\}_{i \in \mathcal{N}}, \quad \{t_{\theta_{\rho}^{\gamma'}(i)}^{\rho}\}_{i \in \mathcal{N}} = \{t_{\theta_{SDD}^{\gamma}(i)}^{SDD}\}_{i \in \mathcal{N}}, \quad \theta_{\rho}^{\gamma'} = \theta_{SDD}^{\gamma}. \quad (5.7)$$

Assume now that $\psi_{SDD}^{\gamma}(j) < \psi_{SDD}^{\gamma}(k)$. Define

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

The service times under ρ are defined to be identical to that of SDD associated with γ for those customers that are different from j and k : $\sigma_i^{\rho} = \sigma_i^{SDD}$ for all $i \in \mathcal{N} - \{j, k\}$. The service times of customers j and k under ρ is defined analogously to (5.4):

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

where s_j, s_k, t_j, t_k are the scheduling times and completion times of customers j, k , respectively, under policy SDD associated with input mapping γ , and where $\Delta = s_k - s_j$. The symbol $U(a, b, \Delta)$ denotes a Bernoulli r.v. with probability distribution $\Pr[U(a, b, \Delta) = 0] = p(a, b, \Delta)$, $\Pr[U(a, b, \Delta) = 1] = 1 - p(a, b, \Delta)$ defined as in (5.5).

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 that

$$\{s_{\psi_{\rho'}^{(i)}}^\rho\}_{i \in \mathcal{N}} = \{s_{\psi_{SDD}^{(i)}}^{SDD}\}_{i \in \mathcal{N}}, \quad \{t_{\theta_{\rho'}^{(i)}}^\rho\}_{i \in \mathcal{N}} = \{t_{\theta_{SDD}^{(i)}}^{SDD}\}_{i \in \mathcal{N}}, \quad \theta_{\rho'}^{\gamma'} \prec_{\{d_i\}_{i \in \mathcal{N}}} \theta_{SDD}^{\gamma}. \quad (5.8)$$

Consequently, in both cases, policy ρ is a feasible with respect to input mapping γ' , and the following relation holds,

$$\{s_{\psi_{\rho'}^{(i)}}^\rho\}_{i \in \mathcal{N}} = \{s_{\psi_{SDD}^{(i)}}^{SDD}\}_{i \in \mathcal{N}}, \quad \{t_{\theta_{\rho'}^{(i)}}^\rho\}_{i \in \mathcal{N}} = \{t_{\theta_{SDD}^{(i)}}^{SDD}\}_{i \in \mathcal{N}}, \quad \theta_{\rho'}^{\gamma'} \prec_{\{d_i\}_{i \in \mathcal{N}}} \theta_{SDD}^{\gamma}. \quad (5.9)$$

Using now 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' ,

$$\{s_{\psi_{SDD}^{(i)}}^{SDD}\}_{i \in \mathcal{N}} = \{s_{\psi_{\rho'}^{(i)}}^{\rho'}\}_{i \in \mathcal{N}}, \quad \{t_{\theta_{SDD}^{(i)}}^{SDD}\}_{i \in \mathcal{N}} = \{t_{\theta_{\rho'}^{(i)}}^{\rho'}\}_{i \in \mathcal{N}}, \quad \theta_{SDD}^{\gamma'} \prec_{\{d_i\}_{i \in \mathcal{N}}} \theta_{\rho'}^{\gamma'}. \quad (5.10)$$

The proof of the claim can thus be concluded by combining relations (5.9) and (5.10). ■

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

In the case of a type 1 node, the proof is simpler as 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 fashion.

Focus now on a type 3 node. As before, assume that $\mathcal{N}_0 = \mathcal{N} = \{1, 2, \dots, N\}$. It again suffices to establish relations (5.3) and (5.6). As the arguments are similar, we provide them solely for (5.3).

We consider the system as if each server was continually serving customers. Whenever a service completion occurs and there is no customer assigned to that server, it corresponds to the completion of a fictitious customer. When a customer is assigned to a server, it is assigned a service time equal to the remainder of the service time already underway at that server. The exponential assumption guarantees that the customer service times are i.i.d. exponential r.v.'s.

Assume that the arrival times and service times are given. Let $0 = e_1 < e_2 < \dots < e_{2N}$ be the decision epochs (note that under the restrictions of Lemma 5.1, there are at most $2N$ decision epochs). Let m , $1 \leq m \leq 2N - 1$, be the first time that π does not follow the SDD rule, i.e., at time e_m , there exist customers j and k in the system with $d_j > d_k$ and that customer j is assigned to a server at time e_m but not customer k .

We shall construct a new policy ρ that violates the SDD rule at least one less time. The (residual) service times under ρ are, as under π , associated with the servers, and are the same as those under π . The decisions of ρ are defined as follows. For all $1 \leq n < m$, the decisions of ρ at time e_n is the same as π . At time e_m , ρ assigns customer k instead of j to a server and does the same assignment for all other customers. If at time e_{m+1} , customer j finishes under π , which implies that customer k finishes under ρ , then for all $m < n \leq 2N$, the assignment decisions of π at time e_n are the same as π except that when customer k is assigned to a server under π , customer j will be assigned to the server under ρ . Otherwise, if customer j does not finish under π at time e_{m+1} (nor does customer k under ρ), then for all $m < n \leq 2N$, the assignment decisions of π at time e_n are exactly the same as π (even for customers j and k).

In both cases, one easily verifies (using the fact that $d_j > d_k$)

$$\{t_{\theta_\rho^*(i)}^\pi\}_{i \in \mathcal{N}} = \{t_{\theta_\pi^*(i)}^\pi\}_{i \in \mathcal{N}} \quad \text{a.s.} \quad \text{and} \quad \theta_\rho^\gamma \prec_{\{d_i\}_{i \in \mathcal{N}}} \theta_\pi^\gamma \quad \text{a.s.}$$

This can be performed repeatedly to produce a scheduling policy where the SDD rule is applied everywhere so that (5.3) holds. $\blacksquare$

Similarly, we can show

Lemma 5.3 *Let the arrival times $\{\tilde{a}_i\}_{i \in \mathcal{N}}$ be fixed. Assume there are two input mappings γ and γ' . Assume further that the service times are independent of the due dates $\{d_i\}_{i \in \mathcal{N}}$ and are i.i.d. random variables having a distribution that depends on the type of node,*

- arbitrary distribution for a type 1 node,
- an ILR distribution for a type 2 node,
- an exponential distribution for a type 3 node.

If $\gamma' \prec_{\{d_i\}_{i \in \mathcal{N}}}^{st} \gamma$, then, for all policy $\pi \in \Psi_{ni}$, there exists a probability space such that

$$\{t_{\theta_{SSDD}^*(i)}^{\pi'}\}_{i \in \mathcal{N}} = \{t_{\theta_\pi^*(i)}^\pi\}_{i \in \mathcal{N}} \quad \text{a.s.} \quad \text{and} \quad \theta_{SSDD}^{\gamma'} \prec_{\{d_i\}_{i \in \mathcal{N}}}^{st} \theta_\pi^\gamma \quad \text{a.s.} \quad (5.11)$$

There also exists a probability space such that

$$\{t_{\theta_{SLDD(i)}}^\pi\}_{i \in \mathcal{N}} = \{t_{\theta_{\pi'}^\pi(i)}^\pi\}_{i \in \mathcal{N}} \text{ a.s.} \quad \text{and} \quad \theta_{\pi'}^{\gamma'} \prec_{\{d_i\}_{i \in \mathcal{N}}}^{\text{st}} \theta_{SLDD}^\gamma \text{ a.s.} \quad (5.12)$$

5.2 Proofs of the Main Results of Non-Idling Case

Proof of Theorem 4.1. As in the proof of Lemma 5.2, we use a sample path interchange argument to prove the theorem. From an arbitrary policy π , We will construct a (finite) series of policies such that the final policy is SDD and that each new policy improves the previous one in the sense of 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 $a_n^i(\pi)$ 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 , $1 \leq i \leq K$.

Fix the customer arrival times $a_1, a_2, \dots, a_N$, and the service times at all the nodes. Determine for each node i , $1 \leq i \leq K$, 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 $1 \leq i \leq K$,

$$\gamma_{SDD}^i \prec_{\{d_n\}_{n \in \mathcal{N}_i}} \gamma_\pi^i \text{ a.s.}, \quad \{a_{\gamma_{SDD}^i(n)}^i(SDD)\}_{n \in \mathcal{N}_i} = \{a_{\gamma_\pi^i(n)}^i(\pi)\}_{n \in \mathcal{N}_i} \text{ a.s.}, \quad (5.13)$$

$$\theta_{SDD}^i \prec_{\{d_n\}_{n \in \mathcal{N}_i}} \theta_\pi^i \text{ a.s.}, \quad \{c_{\theta_{SDD}^i(n)}^i(SDD)\}_{n \in \mathcal{N}_i} = \{c_{\theta_\pi^i(n)}^i(\pi)\}_{n \in \mathcal{N}_i} \text{ a.s.} \quad (5.14)$$

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

$$\gamma_{SDD}^1 = \gamma_\pi^1, \quad \{a_{\gamma_{SDD}^1(n)}^1(SDD)\}_{n \in \mathcal{N}_1} = \{a_{\gamma_\pi^1(n)}^1(\pi)\}_{n \in \mathcal{N}_1},$$

so that (5.13) holds for $i = 1$. Appealing to Lemma 5.2 implies that (5.14) also holds for $i = 1$.

Assume that for some $2 \leq j \leq K$, relations (5.13) and (5.14) 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 3.9 that

$$\gamma_{SDD}^j \prec_{\{d_n\}_{n \in \mathcal{N}_j}} \gamma_\pi^j, \quad \{a_{\gamma_{SDD}^j(n)}^j(SDD)\}_{n \in \mathcal{N}_j} = \{a_{\gamma_\pi^j(n)}^j(\pi)\}_{n \in \mathcal{N}_j},$$

so that (5.13) holds for $i = j$. Applying Lemma 5.2 yields that (5.14) also holds for $i = j$, which completes the proof of relations (5.13) and (5.14).

Removing the conditioning on the arrival times and applying Lemma 3.5 yields

$$L(SDD) \leq_{E_1} L(\pi).$$

The proof of the inequality $L(\pi) \leq_{E_1} L(LDD)$ is analogous. ■

Proof of Theorem 4.2. The proof is similar to that of Theorem 4.1, and uses Lemmas 5.3, 3.10 and 3.8. ■

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 $q(\pi) = (q_1(\pi), \dots, q_N(\pi))$. Fix the arrival times and replace the due dates in (5.14) of the proof of Theorem 4.1 with the arrival times. According to Lemma 3.6,

$$q(FCFS) \leq q(\pi) \quad a.s.,$$

or

$$D(FCFS) \leq D(\pi) \quad a.s.$$

Removing the conditioning on the arrival times and the service times yields

$$D(FCFS) \leq_{st} D(\pi).$$

The proof of the relation $D(\pi) \leq_{st} D(LCFS)$ is similar, and is omitted. ■

6 Proof of Idling Policy Results

In this section we focus on in-forests that consist solely of type 3 nodes. In addition, we restrict ourselves to the case where the service times are exponential r.v.'s. In order to establish the results for idling policies, we establish Lemma 6.1 below. It will be clear that theorems 4.4, 4.5, and 4.6 are immediate consequences of Lemmas 6.1 and 3.2, and the results of Section 4.1. The detailed proofs are left to the interested reader.

Lemma 6.1 *Assume that service times are exponentially distributed and 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*

$$(c_1(\rho), \dots, c_N(\rho)) \leq_{st} (c_1(\pi), \dots, c_N(\pi)), \quad (6.1)$$

$$L(\rho) \leq_{st} L(\pi), \quad (6.2)$$

$$R(\rho) \leq_{st} R(\pi). \quad (6.3)$$

Proof. Let $a_n^i(\pi)$ and $c_n^i(\pi)$ be the arrival and departure times of customer $n \in \mathcal{N}_i$ at node i under π , respectively.

According to Lemma 5.1 we can restrict ourselves to the class of non-idling policies whose decision points at each node occur only at the arrival times of that node and all but the last service completion time of that node (cf. 5.1), i.e., preemptions and new customer assignments at each node occur only at the instants of customer arrivals to that node and at the instants of service completion of that node.

As in the last part of the proof of Lemma 5.2, we consider the system as if each server was continually serving customers. Whenever a (virtual) service completion occurs and there is no customer assigned to that server, it corresponds to the completion of a fictitious customer. When a customer is assigned to a server, it is assigned a service time equal to the remainder of the service time already underway at that server. The exponential assumption guarantees that the customer service times are i.i.d. exponential r.v.'s. Assume that the arrival times and service times are given. We show that in such a probability space, there are policies $\rho^1, \rho^2, \dots, \rho^K$, such that for all k , $1 \leq k \leq K$, policy ρ^k is non-idling at all the nodes $1, 2, \dots, k$, and that

$$\forall i, 1 \leq i \leq k: \{a_n^i(\rho)\}_{n \in \mathcal{N}_i} \leq \{a_n^i(\pi)\}_{n \in \mathcal{N}_i} \text{ a.s.}, \quad \{c_n^i(\rho)\}_{n \in \mathcal{N}_i} \leq \{c_n^i(\pi)\}_{n \in \mathcal{N}_i} \text{ a.s.}, \quad (6.4)$$

which will imply the assertions of the lemma in taking $\rho = \rho^K \in \Psi_{ni}$.

The proof is by induction on $k \in \mathcal{K}$. Consider first node $k = 1$.

Let $0 \leq e_1 < e_2 < \dots$ be the superposition of the sequences of arrival times and of the (virtual) service completion times at node 1. Define policy ρ^1 as follows:

- Policy ρ^1 make the same scheduling decisions as π at all the nodes $i \geq 2$.
- At node 1, for $m = 1, 2, \dots$,
 - If a customer $n \in \mathcal{N}_1$ is assigned to a server at e_m under π , and if this customer is not finished by e_m under ρ^1 , then n is also assigned to a server under ρ^1 .

- If there is a server such that no customer is assigned to at time e_m under π , or if the customer that is assigned to at e_m under π is already finished by time e_m under ρ^1 , then ρ^1 assigns a customer waiting in the queue, if any, to that server.

Under this construction, it is easy to see that ρ^1 is non-idling at node 1 and that

$$\{c_n^1(\rho^1)\}_{n \in \mathcal{N}_1} \leq \{c_n^1(\pi)\}_{n \in \mathcal{N}_1} \text{ a.s.},$$

so that relation (6.4) holds for ρ^1 .

Assume that for some $k \geq 2$, there is policy ρ^{k-1} that is non-idling at all the nodes $1, 2, \dots, k-1$, and that relation (6.4) holds for ρ^{k-1} .

Define policy ρ^k as follows. Policy ρ^k makes the same scheduling decisions as ρ^{k-1} at all the nodes $1, \dots, k-1$, and ρ^k makes the same scheduling decisions as π at all the nodes $k+1, k+2, \dots, K$. Consider the decisions made by ρ^k at node k .

If k is a leaf node in $\mathcal{K}$, then the relation

$$\{a_n^k(\rho^k)\}_{n \in \mathcal{N}_k} \leq \{a_n^k(\pi)\}_{n \in \mathcal{N}_k} \text{ a.s.},$$

trivially holds (in fact, the equality holds.). Assume that k is not a leaf node in $\mathcal{K}$. The arrival times at node k is the superposition of the completion times at its predecessor nodes. Therefore, the above inequality still holds, due to the induction assumption.

Now, policy ρ^k can be constructed from π by removing the idleness at node k in the same way as in the definition of ρ^1 for node 1, where, the sequence of decision epochs $0 \leq e_1 < e_2 < \dots$, should be understood as the superposition of the arrival times $\{a_n^k(\rho^k)\}_{n \in \mathcal{N}_k}$ and $\{a_n^k(\pi)\}_{n \in \mathcal{N}_k}$, and the (virtual) service completion times at node k . Hence,

$$\{c_n^k(\rho^k)\}_{n \in \mathcal{N}_k} \leq \{c_n^k(\pi)\}_{n \in \mathcal{N}_k} \text{ a.s.},$$

which, in addition to the induction assumption, imply that relation (6.4) holds for ρ^k .

We have thus completed the inductive step and have shown that (6.4) holds for all $1 \leq k \leq K$. Since $\rho = \rho^K$ is non-idling at all the nodes, we have thus completed the proof. ■

7 Concluding Remarks

In this paper, we studied the scheduling problems in in-forests with identical multiple servers at each node. We focused on those performance metrics as the response times, the latenesses,

and the end-to-end delays of the customers. Various extremal properties have been established for several simple policies using stochastic ordering techniques.

Although the focus of the paper has been on transient results, it should be clear that similar extremal properties hold in the stationary regime. In particular, the stochastic orderings $\leq_{E_3}$, $\leq_{E_3^!}$ and $\leq_{E_3^!}$ (recall that $\leq_{E_1} \Rightarrow \leq_{E_2} \Rightarrow \leq_{E_3}$, $\leq_{E_1^!} \Rightarrow \leq_{E_2^!} \Rightarrow \leq_{E_3^!}$, and $\leq_{E_1^!} \Rightarrow \leq_{E_2^!} \Rightarrow \leq_{E_3^!}$), established in this paper on the transient metrics $R(\pi)$ and $L(\pi)$ 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 metric $D(\pi)$ reduces to a $\leq_{st}$ ordering on the stationary random variables $D(\pi)$ 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, even for the case when $s = 1$. As far as the authors are concerned, the only more general result existing in the literature is the comparison for the response times in the single-queue single-server system with Increasing Failure Rate (IFR) service times. Indeed, Hirayama and Kijima [12] obtained that in the $G/IFR/1$ model, FCFS policy is optimal within the class of *preemptive* policies :

$$\forall \pi \in \Psi : \quad R(FCFS) \leq_{E_3^!} R(\pi). \quad (7.1)$$

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 (7.1) 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. Here is a 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 (7.1) does not hold for tandem queues with single IFR servers either. Righter and Shanthikumar [20] provided a counterexample with 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 7.1 *In any tandem queueing network consisting of n ILR/1 queues, FCFS policy applied to all the queues is optimal within the class of preemptive policies in the sense that*

$$\forall \pi \in \Psi: \quad R(\text{FCFS}) \leq_{E_1} R(\pi).$$

The above theorem follows from the facts that FCFS stochastically minimizes the vector of completion times in each queue [20] 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 [20].

In case a single multi-server queue is considered (i.e., $\mathcal{K}$ contains a single node), the assumption of ILR service times in Theorems 4.2 and 4.5 can be relaxed. The statements hold for arbitrary i.i.d. service times [17].

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*. Akademik-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", submitted.
- [18] A. W. Marshall, I. Olkin, *Inequalities: Theory of Majorization and Its Applications*, Academic Press, 1979.
- [19] M. Pinedo, "Stochastic Scheduling with Release Dates and Due Dates." *Oper. Res.*, Vol. 31, No. 3, pp. 559-572, May-June 1983.
- [20] R. Richter, J. G. Shanthikumar, "Extremal Properties of the FIFO Discipline in Queueing Networks", manuscript.
- [21] J. G. Shanthikumar, 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.
- [22] D. Stoyan, *Comparison Methods for Queues and Other Stochastic Models*. English translation (D.J. Daley editor), J. Wiley and Sons, New York, 1984.

- [23] V. Strassen, "The existence of Probability Measures with Given Marginals," *Ann. Math. Stat.*, Vol. 36, pp. 423-439, 1965.
- [24] 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.
- [25] 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.
- [26] O. A. Vasicek, "An Inequality for the Variance of Waiting Time Under a General Queueing Discipline." *Operations Research*, Vol. 25, pp. 879-884, 1977.
- [27] W. Whitt, "The Amount of Overtaking in a Network of Queues," *Networks*, Vol. 14, pp. 411-426, 1984.
- [28] R. W. Wolff, "An Upper Bound for Multi-Channel Queues," *J. Appl. Prob.*, Vol. 14, pp. 884-888, 1977.
- [29] R. W. Wolff, "Upper Bounds on Work in System for Multichannel Queues," *J. Appl. Prob.*, Vol. 24, pp. 547-551, 1987.

ISSN 0249 - 6399