Available online at www.sciencedirect.com
ScienceDirect
Performance Evaluation 64 (2007) 1121–1136
PERFORMANCE EVALUATION
An International Journal
www.elsevier.com/locate/peva

SQF: A slowdown queueing fairness measure☆

Benjamin Avi-Itzhak[a], Eli Brosh[b,*], Hanoch Levy[c]

[a] *RUTCOR, Rutgers University, New Brunswick, NJ 08854-8003, USA*
[b] *Department of Computer Science, Columbia University, New-York, NY 10027-7003, USA*
[c] *School of Computer Science, Tel Aviv University, Tel Aviv 69978, Israel*

Available online 27 June 2007

Abstract

Expected slowdown has been proposed as a *criterion* to evaluate queue fairness. In this work we examine how the constant slowdown principle can be used as a basis for a queueing *fairness measure*. We propose the Slowdown Queueing Fairness (SQF) measure based on the principle that customers' waiting time should be proportional to their service time. We analyze its properties and examine how they react to both seniority and service requirements. We also examine whether its behavior fits intuition. Its values for a variety of single-server scheduling policies as well as for multi-server architectures are derived.

Keywords: Queue; Fairness; Slowdown

1. Introduction

Fairness is an important factor in queues operation and scheduling (e.g. [1–3]) in computer-based applications, such as Web servers and call centers, as well as in other real life applications. While the subject has hardly been touched for decades, in recent years its importance has been recognized and has received much attention, especially in the computer performance community (e.g., [4–11]). For further motivation, applications, and review of approaches we refer the interested reader to [12]. A central issue in this area is the quest for a uniform fairness metric that can be used to evaluate the fairness of various systems and scheduling in order to compare them, fairness-wise, to each other.

The slowdown of a job is defined as $S(x) = T(x)/x$, where $T(x)$ is the job's sojourn time in the queueing system and x is the job's processing (service) time. Expected slowdown has been mentioned in the context of fairness as early as [13], where the analysis of the $M/G/1$ Processor Sharing (PS) policy yielded that the expected slowdown is the same for all values of x, namely $E[S(x)] = 1/(1-\rho)$. Slowdown has received much attention in recent years (e.g., [14–17]), in particular it has been recognized as an underlying factor for evaluating fairness, adhering to the *proportionality principle*, namely to the belief that it is fair that the sojourn time of a job would be proportional to its size (those who demand more from the system should wait more). As such, its expected value served as the basis for forming a fairness criterion in [4] where it is shown that no $M/G/1$ scheduling policy can attain the same constant

☆ This work was supported in part by the Israeli Ministry of Science and Technology, grant number 380-801, and by EURO-NGI network of excellence.

* Corresponding author. Tel.: +1 646 775 6087.
E-mail addresses: aviitzha@rutcor.rutgers.edu (B. Avi-Itzhak), elibrosh@cs.columbia.edu (E. Brosh), hanoch@cs.tau.ac.il (H. Levy).

0166-5316/$ - see front matter
doi:10.1016/j.peva.2007.06.018

expected slowdown lower than $1/(1-\rho)$ for all values of x. This leads to the fairness criterion stating that a policy is fair if $E[S(x)] \leq 1/(1-\rho)$ for all x. This leads to the classification of $M/G/1$ disciplines into three classes: "always fair", "sometimes fair" or "never fair". The advantage of the criterion is that it yielded to analytic treatment in the $M/G/1$ case and to classifying a wide set of $M/G/1$ policies into one of the three classes. Its disadvantages is that it is not a scaled measure (and thus cannot distinguish between policies within each of the three classes), that it does not address the system temporality (job seniority) and that it was proposed only to single processor $M/G/1$ systems.

The purpose of this work is to start with the proportionality principle as the underlying belief and systematically construct an effective scaled fairness measure, based on it. The idea for developing such measure was raised first in [12]. To achieve this objective we propose a new measure called *Slowdown Queueing Fairness* (SQF). SQF can be viewed as bridging the gap between the slowdown expectation criterion (which focuses on service requirements) and the 'natural' waiting time variance measure (which focuses on job seniority). The measure has several desired properties; it accounts for both system temporality and service requirements and yields to exact analysis over a wide range of systems, as described below.

We start our work (Section 2) by first focusing on the slowdown expectation fairness criterion [4], where the authors identified P-LCFS and PS as the "always fair" policies and posed an open question of whether there are other $M/G/1$ policies which are "always fair". We describe an indefinitely large class of policies which are "always fair", thus answering the question. We further demonstrate that the policies within this class can differ drastically from each other in their deviation from a fixed slowdown, raising the need for a measure based on the slowdown principle.

We then (Section 3) construct the slowdown fairness measure (SQF) and derive its basic properties. SQF is based on the proportionality fairness principle which can be stated as: The slowdown of all customers must be the same. We define the discriminations of an individual customer as $W(x) - px$, where $W(x)$ and x are the waiting time and service requirement (job size) of the customer, respectively and p is a constant which may vary from one system to another. The discrimination thus expresses the magnitude of the deviation from the absolutely fair proportional waiting time, $W(x) = px$, and may be positive, negative or zero. The exact measure, based on customers discriminations, is then defined for specific sample paths as well as for systems in steady state. Further in the section we study the properties of SQF and show that the measure is appropriately sensitive to both seniority differences and size differences between jobs using simple sample path-based tests. We also provide bounds on SQF and study the feasibility of obtaining zero unfairness.

We then (Section 4) turn to the analysis of SQF in steady state. We demonstrate that SQF yields to analysis for a wide spectrum of systems, which is important for its practical potential use. To demonstrate this we analyze both single-server systems and multi-server systems.

Addressing single-server systems with non-preemptive non-size-based policies we show that for $G/GI/1$ seniority preserving schedules are more fair than those which do not preserve seniority, and specifically, First-Come–First-Served (FCFS) is more fair than Non-Preemptive Random Order of Service (NP-ROS) which is more fair than Non-Preemptive Last-Come–First-Served (NP-LCFS). Addressing single-server systems with preemptive policies we show that PS is more fair than Preemptive Last-Come–First-Served (P-LCFS). We further derive closed form expressions for the SQF fairness values for a number of $M/M/1$ and $M/GI/1$ scheduling policies.

We then turn to multi-server systems. We note that most of the research on fairness measures of recent years has been limited to yield results for single-server systems (except for RAQFM that has been applied to multi-queue systems [18]) and thus the results on this subject are limited. Addressing this subject we look at several 'classical' multi-queue and multi-server systems ($M/M/m$, $M/M/1$ and m parallel $M/M/1$ systems) and compare their fairness values. We are able to provide a full (closed-form) comparison of these systems, which fits intuitive reasoning. Further, an interesting result of the analysis is that for $M/M/m$ the system fairness monotonically increases with the number of servers (and reaches optimal fairness when m reaches ∞). Thus, under SQF, many slow processors are more fair than a few fast processors.

Lastly, in Section 5 we discuss the results and evaluate the measure by examining its values for the various scheduling policies and in Section 6 we compare SQF to other approaches.

2. The use of expected slowdown for fairness

[4] defines an $M/G/1$ policy to be 'always fair' if for all loads and all service-time distributions, $E[S(x)] \leq 1/(1-\rho)$, $x \geq 0$. Accordingly, the P-LCFS policy (under which the processor is fully devoted to the most recent

arrival) and the PS policy (under which the processor is evenly shared among all jobs in the system) are classified as always fair. The authors also pose as an open question whether there exist other 'always fair' $M/G/1$ policies. In the following we describe an indefinitely large class of policies that are 'always fair', thus addressing the open question. We also show that, although they all satisfy $E[S(x)] \leq 1/(1-\rho)$, the differences between them with respect to slowdown variability can be very drastic, implying that a more sensitive treatment of slowdown might be useful for measuring fairness.

Consider the class of $M/G/1$ time-sharing queues with finite number of service positions [19] described below: "The queue is ordered and there are r service positions and an unlimited number of waiting positions. When there are n jobs in the system they are in positions $1, 2, \ldots, n$ and a proportion $\phi(i, n)$ of the service rate is directed at the job in position i, $(i = 1, 2, \ldots, \min(r, n))$ where $\sum_{i=1}^{\min(r,n)} \phi(i, n) = 1, n = 1, 2, \ldots$. A newly arrived job which encounters n in the system, $n = 0, 1, 2, \ldots$, is assigned to service position i with probability $\phi(i, n+1)$. The other n jobs present in the system at that time are re-ordered in accordance to some arbitrary service-independent rule. (A re-ordering rule is service independent if it involves permutations based only on the knowledge of the number of jobs, n, and their positions in the queue). When the processing of a job is completed it departs, and the remaining jobs are instantaneously re-ordered in accordance to some arbitrary service-independent rule. Preemptions due to newly arriving jobs or due to re-ordering do not result in loss and when a preempted job re-enters service its processing resumes from the point of the most recent interruption".

It can be verified, see [20], that all the members of this class possess the equilibrium properties of symmetric queues. Thus, for all members of this class the expected delay and response times of a job requiring x units of service time are given respectively by $E[W(x)] = x\rho/(1-\rho)$ and $E[T(x)] = x/(1-\rho)$. By the criterion proposed in [4] all the members of this class are always fair. We note that if $\phi(i, n) = 1/i$, $i = 1, 2, \min(r, n)$, we get the PS queue when $r \to \infty$ and an indefinitely large class of preemptive queues for each finite r. In particular, in the case of $r = 1$ we get the P-LCFS queue, where a displaced customer is always placed in position 2 and all customers in the waiting positions are moved one position back. Let the conditional delay be denoted by $W^f(x)$ for P-LCFS, by $W^l(x)$ for the same discipline where the displaced customer is placed at the end of the line, and by W^∞ for the PS case. [19] shows that $\mathrm{Var}[W^f(x)] \geq \mathrm{Var}[W^l(x)]$ and $\mathrm{Var}[W^f(x)] \geq \mathrm{Var}[W^\infty(x)]$. The authors of [19] conjecture that the variance of the conditional delay is decreasing in r and therefore $\mathrm{Var}[W(x)] \geq \mathrm{Var}[W^\infty(x)]$, where $W(x)$ is the conditional waiting time in any arbitrary policy in this class. The implication of this conjecture, regarding the fairness measure proposed in this paper, is discussed later on in Section 4.2.3.

To demonstrate the magnitude of the differences between the policies, one can take the closed form expression derived in [19] for exponentially distributed service times $\frac{\mathrm{Var}[W^f(x)]-\mathrm{Var}[W^l(x)]}{\mathrm{Var}[W^f(x)]} = \frac{\rho[1-\mathrm{e}^{-\lambda(1-\rho)x}]}{\lambda(1-\rho)x} \geq 0$. For a fixed value of x, this expression increases with ρ and goes to ∞ when $\rho \to 1$. For the exponential case they also show that $\mathrm{Var}[W^l(x)] \geq \mathrm{Var}(W^\infty(x))$. The variance reduction obtained when preempted jobs are placed at the end of the line, or when PS is used, is very significant when the traffic intensity is high. For example, if $\lambda = 0.9$, $\rho = 0.9$ and $x = 0.5$ one gets $\mathrm{Var}[W^f(0.5)] = 90.4\mathrm{Var}(W^\infty(0.5)) = 18.5\mathrm{Var}(W^l(0.5))$. If the discrimination is defined as $W(x) - x\rho/(1-\rho)$ and the variability of the discrimination reflects the unfairness, then for customers requiring service time of 0.5 the P-LCFS is 90.4 times more unfair than the PS schedule, and the relation above implies that this ratio can become infinitely large. Thus, the expected slowdown might not fully reflect fairness and a more sensitive measure might be called for.

3. The slowdown queueing fairness (SQF) measure: Definition and properties

Assume an arbitrary queueing system where customers $C_1, C_2, \ldots$ (interchangeably called jobs $J_1, J_2, \ldots$) arrive at the system in that order. The arrival time of C_i is a_i, its service requirement (job size) is x_i and its departure time (which depends on the scheduling policy) is d_i. The sojourn time of C_i is given by $T_i = d_i - a_i$, its waiting time is $W_i = T_i - x_i$, and its *slowdown* is $S_i = T_i/x_i$.

To derive unfairness for a given scenario (a sample path or part of it), consider a finite path consisting of N jobs $J_1, \ldots J_N$. Then absolute fairness is obtained if $T_i = cx_i$, $1 \leq i \leq N$ for some constant c.

The *individual discrimination* of J_i is defined as $D_i = T_i - cx_i$, namely the deviation of its sojourn time from its absolutely fair sojourn time. If $D_i = 0$ (its slowdown is exactly identical to the constant c), J_i is not discriminated. If $D_i < 0$, its sojourn time is greater than cx_i and J_i is *negatively* discriminated; if $D_i > 0$, it is positively discriminated. We note that $D_i = W_i - px_i$ where $p = c - 1$ is an equivalent definition for D_i. Discrimination within a group is assumed to have a zero sum property, i.e., positively discriminating a customer must result in negatively discriminating other customers and vice versa. Individual discrimination in the group must therefore sum to zero (see [3] for a discussion of the basic principles of fairness). The appropriate value of c is obtained from solving

$$\sum_{i=1}^{N} D_i = \sum_{i=1}^{N} T_i - \sum_{i=1}^{N} c x_i = 0, \tag{1}$$

which leads to:

$$c = \left(\sum_{i=1}^{N} T_i\right) \Big/ \left(\sum_{i=1}^{N} x_i\right) = \frac{\bar{T}}{\bar{x}}. \tag{2}$$

Since the mean discrimination of the scenario is zero, we define its unfairness as the variance of the individual discriminations (equalling the second moment):

$$\frac{1}{N}\sum_{i=1}^{N} D_i^2 = \frac{1}{N}\sum_{i=1}^{N} (T_i - c x_i)^2. \tag{3}$$

Having defined individual discrimination and unfairness of a scenario we can now define the unfairness of a system, in particular for a system in equilibrium. Assume that service times are distributed as a random variable B with pdf $b(x)$ and moments $b_k = E[B^k]$.

Let T and W denote the equilibrium sojourn time and waiting time, respectively, and let $T(x)$ and $W(x)$ denote the (conditional) sojourn time and the (conditional) waiting time for a customer whose service time is x, respectively. $S(x) = T(x)/x$ is the *slowdown* of job of size x.

The requirement for a zero sum leads to an equation similar to Eq. (1):

$$\int_0^\infty c x b(x)\mathrm{d}x - \int_0^\infty E[T(x)] b(x)\mathrm{d}x = 0 \tag{4}$$

yielding

$$c = \frac{E[T]}{b_1}, \qquad p = \frac{E[W]}{b_1}. \tag{5}$$

Remark 3.1. It is worthwhile noting that the value of the proportionality coefficient c, as given by (5), ensures that the SQF of any given queueing system is determined by comparing it to an 'absolutely fair' system of the same level of efficiency (efficiency being traditionally measured by the expected sojourn time). Though the 'absolutely fair' system may not be feasible at all, it is an important property which separates unfairness from efficiency. A system may be highly efficient and highly unfair, and vice versa.

We define the slowdown unfairness measure as the limit of (3):

$$\mathrm{SQF} = \lim_{N\to\infty} \sum_{i=1}^{N} (c x_i - T_i)^2 = \int_0^\infty E[T(x) - cx]^2 b(x)\mathrm{d}x = \int_0^\infty E[W(x) - px]^2 b(x)\mathrm{d}x, \tag{6}$$

from which we obtain, after taking the limit and substitution of (5), the form

$$\begin{aligned}\mathrm{SQF} &= E[T^2] + \frac{b_2}{{b_1}^2} E^2[T] - \frac{2E[T]}{b_1}\int_0^\infty x E[T(x)] b(x)\mathrm{d}x \\ &= E[W^2] + \frac{b_2}{{b_1}^2} E^2[W] - \frac{2E[W]}{b_1}\int_0^\infty x E[W(x)] b(x)\mathrm{d}x. \end{aligned} \tag{7}$$

The first two terms in (7) are determined by the first and second moments of the sojourn time, which are known for many queueing systems. The third term $\int_0^\infty x E[T(x)] b(x)\mathrm{d}x$, however, is not as known and may pose tractability problems in some cases. Thus, in some cases, the derivation of SQF may require the conduction of a new analysis not given in the literature. It is important to note that SQF is based on individual customer discriminations and thus accounts for job seniority. This is in contrast to the slowdown expectation criterion, which is based on expected slowdown values, and thus cannot address seniority.

Nonetheless, the following observation shows that there exists a class of systems (characterized by a simple linear expression for $E[T(x)]$), for which the first two moments of the sojourn time suffice to derive the SQF value of the system:

Observation 3.1. *If for a given system $E[T(x)] = \alpha x + \beta$ for some constants α and β, then the SQF value can be computed solely based on $E[T]$ and $E[T^2]$ and the first two service-time moments as follows:*

$$\mathrm{SQF} = E[T^2] + \frac{b_2}{{b_1}^2} E^2[T] - 2E[T]\left(\alpha \frac{b_2}{b_1} + \beta\right). \tag{8}$$

Examples of systems that fall into this category are: $M/G/1$ symmetric queues (e.g., P-LCFS, Preemptive Random Order of Service (P-ROS), and PS), for which $\alpha = 1/(1-\rho)$, $\beta = 0$ (as explained in Section 2), an $M/G/1$ system with FCFS service, for which $\alpha = 1$, $\beta = E[W]$, and other more complicated systems where the waiting time is independent of the service time (e.g., an $M/G/2$ system with FIFO queueing). Under P-ROS, customers are served in random order: whenever a service is completed or an arrival occurs, the next customer to be served is selected uniformly at random from the customers present, if any.

An alternative approach for defining a slowdown fairness is to measure discrimination by the relative deviation from proportionality $c_1 - T(x)/x$. This principle leads to deriving the constant via $c_1 = \int_0^\infty b(x)\frac{E[T(x)]}{x}\mathrm{d}x$ and hence yield an unfairness measure $\int_0^\infty b(x)E[(S(x)-c_1)^2]\mathrm{d}x$. While this measure is more biased in favor of small jobs than SQF, it seems, however, to be very hard to compute. The feasibility of such a metric remains as an open subject for future research. Note that c_1 is not necessarily equal to c.

3.1. Bounds on SQF

Manipulation of the SQF expression yields:

$$\begin{aligned}\mathrm{SQF} &= \int_0^\infty E\left[(T(x) - E[T(x)] + E[T(x)] - cx)^2\right] b(x)\mathrm{d}x \\ &= \int_0^\infty \left(\mathrm{Var}[T(x)] + (E[T(x)] - cx)^2\right) b(x)\mathrm{d}x.\end{aligned} \tag{9}$$

Using the above relation we can derive bounds on the unfairness value

$$\mathrm{Var}[T] \le \mathrm{SQF} \le E[T^2] + \frac{b_2}{{b_1}^2} E^2[T]. \tag{10}$$

The lower bound results from (9) and $(E[T(x)] - cx)^2 \ge 0$, and the upper bound from (7). Note that the lower bound is tight since it can be achieved for symmetric queues.

3.2. Sensitivity to seniority and size

We study the sensitivity of the SQF measure to both seniority differences and service requirement differences by examining how it operates in simple-to-understand scenarios and showing that it is very intuitive in these cases. Specifically, we show that under SQF when two jobs have identical service times it is more fair to complete the service of the more senior job earlier than vice versa, and if the arrival times are identical it is more fair to complete the service of the shorter job earlier. Similar properties were studied for other fairness measures in [12].

Theorem 3.1. *Let jobs J_i and J_j arrive at $a_i < a_j$ and obey $x_i = x_j$. Let π be a scheduling policy where the service of J_i is completed before that of J_j and let π' be identical to π except for interchanging the service schedule[1] of J_i and J_j, assuming that the service schedule is interchangeable. Then the SQF unfairness values of these schedules, $f(\pi)$ and $f(\pi')$, obey $f(\pi) < f(\pi')$.*

---

[1] This means that time slices allocated by the server to the jobs are exchanged.

Proof. First note that the constant c (Eq. (5)) is the same under both π and π' since the exchanged jobs have equal service times. This implies that the individual discrimination inflicted on all jobs other than J_i and J_j is the same under both policies. Thus the change in the unfairness value is

$$f(\pi) - f(\pi') = \frac{1}{n}((t_i - x_i c)^2 + (t_j - x_j c)^2 - (t_i + d_j - d_i - x_i c)^2 - (t_j - d_j - d_i - x_j c)^2)$$

where t_i and d_i are, respectively, the sojourn time and departure epoch of J_i according to schedule π. Using the service-time equality $x_i = x_j$ the last equation can be simplified to $\frac{2(d_j - d_i)(a_i - a_j)}{n} < 0$, implying that $f(\pi) < f(\pi')$ and that π is more 'fair' than π'. □

From the above proof it can be clearly seen that the unfairness difference in this scenario is given by $\frac{2(d_j - d_i)(a_i - a_j)}{n}$ and thus is monotonic in the seniority difference (e.g., in $a_j - a_i$). The higher the seniority difference is, the more unfair it is to serve the less senior job J_j first; the lower it is, the less unfair it is to serve J_j first.

Corollary 3.1. *The change in the unfairness value due to the interchange $\pi \to \pi'$ is monotonic in the seniority difference of the exchanged jobs.*

Theorem 3.2. *Consider jobs J_i, J_j where $a_i = a_j$ and $x_i < x_j$. Let π be a schedule where the service completion times obey $d_i < d_j$. Let π' be obtained from π by keeping the service completion of all other jobs the same and interchanging the completion times of J_i, J_j: $d'_i = d_j$, $d'_j = d_i$. Then $f(\pi) < f(\pi')$.*

Proof. Observe that the sum of completion times remains unchanged under both π and π' and thus the constant c and the discrimination inflicted on all jobs except of J_i and J_j are the same under both policies. We thus have $f(\pi) - f(\pi') = \frac{1}{n}((t_i - x_i c)^2 + (t_j - x_j c)^2 - (t_j - x_i c)^2 - (t_i - x_j c)^2)$ where t_i is the sojourn time of J_i under π. This expression can be simplified to $\frac{-2c(x_i - x_j)(t_i - t_j)}{n} < 0$, implying that $f(\pi) < f(\pi')$ and that π is more 'fair' than π'. □

Similarly to Corollary 3.1, the unfairness difference in this scenario is given by $\frac{-2c(x_i - x_j)(t_i - t_j)}{n}$ and thus is monotonic in the service-time difference; that is, the higher the service-time difference $x_j - x_i$ is, the more unfair it is to serve the larger job first.

Corollary 3.2. *The change in the unfairness value due to $\pi \to \pi'$ is monotonic in the service-time difference of the exchanged jobs.*

3.3. Optimal scheduling and zero unfairness

First, observe that in a system with infinite processing capacity a trivial resource allocation (scheduling) can achieve a zero SQF value:

Observation 3.2. *In an $M/G/\infty$ system the discrimination of every job is zero and therefore system's SQF value is zero also.*

Next, we examine the feasibility of obtaining zero unfairness for single-server systems. For comparison note that two unfairness measures proposed in the past, do have zero unfairness policies. For the measure proposed in [6], which focuses on job seniority (disregarding job size), FCFS always achieves minimal unfairness (which can be pegged to zero). For the measure proposed in [9], PS always achieves zero unfairness (which is optimal).

We use sample path analysis (see Section 3) to study the system conditions that lead to a zero unfairness value.

Lemma 3.1. *For a system with a finite number of servers and a non-zero probability of busy periods of more than one customer, there exist sample paths for which zero unfairness cannot be achieved by any non-preemptive policy.*

Proof. Let x_1 be the size of the first job served in a busy period consisting of more than one job, and let S be the sum of sizes of all other jobs in the busy period. From (2) and the definition of D_i we have $D_i = (x_1 - x_1(x_1 + S + WT)/(x_1 + S))^2$, where WT is the sum of waiting times of all other jobs. This quantity equals zero only when WT is zero, which is not feasible. □

If preemption is allowed then in certain cases one can come up with schedules leading to zero unfairness.

Lemma 3.2. *Consider a single-server busy period consisting of N jobs, $J_1, \ldots, J_N$ with service times $x_1, \ldots, x_N$, and arrival times $a_1, \ldots, a_N$, respectively. Assume the busy period starts at a_1 and let $\hat{x} = \sum_{i=1}^{N} x_i$. Let π be a scheduling policy that achieves a zero unfairness. Let J_k, $1 \le k \le N$ be a job for which* $\min_i \frac{a_1+\hat{x}-a_i}{x_i} = \frac{a_1+\hat{x}-a_k}{x_k}$, $i = 1, \ldots, N$. *Let π be a policy that achieves a zero SQF value. Then:* (i) *π must schedule J_k to complete service last at the busy period, that is at $a_1 + \hat{x}$.* (ii) *The completion time of any other job J_j, $j \ne k$, should be* $a_j + \frac{a_1+\hat{x}-a_k}{x_k} x_j$.

Suppose J_k is not scheduled to terminate at $a_1+\hat{x}$. Then $T_k < a_1+\hat{x}-a_k$. Assume J_j is the job that completes service at the end of the busy period, then $T_j = a_1 + \hat{x} - a_j$. Now, from the condition of the lemma $\frac{a_1+\hat{x}-a_k}{x_k} < \frac{a_1+\hat{x}-a_j}{x_j}$, and thus $T_k/x_k < T_j/x_j$. Hence, regardless of the value of c it is impossible that both $T_k - cx_k = 0$ and $T_j - cx_j = 0$, from which the proof of (i) follows. Claim (ii) follows since we must have $T_j = cx_j$ and $c = \frac{a_1+\hat{x}-a_k}{x_k}$. □

A simple example of how to construct a zero unfairness schedule is given as follows. Suppose we have a single-server with a preemptive non-idling service policy and a sample path consisting of 3 jobs, J_1, J_2, and J_3 with sizes 3, 7, 10, respectively, that arrived at time zero. From Lemma 3.2, $T_3 = 20$. Since the discrimination of J_3 should be zero, we have $c = T_3/x_3 = 20/10$. Therefore, we could obtain a zero unfairness if we serve x_1 and x_2 at time intervals (3, 6) and (7, 14), respectively, and serve x_3 at time intervals (0, 3), (6, 7), (14, 20). Note that high fairness (zero unfairness) does not necessarily imply that the system is efficient, namely, that the waiting times are low. As stated before (Remark 3.1), fairness and efficiency can be viewed as seperate performance entities, where the optimization of one may hurt the other, and vice versa.

Corollary 3.3. *If all Jobs arrive at the same time then a necessary condition for obtaining zero unfairness scheduling would be to schedule a job with the largest service time as the last job to complete service.*

However, Lemma 3.2 provides only a necessary (and not sufficient) condition for obtaining a zero unfairness.

Theorem 3.3. *For any single-server system under preemptive policy, there exists a sample path for which no zero unfairness schedule is possible.*

Proof. Consider a busy period consisting of three jobs with sizes 1, 1, 10 that arrive simultaneously at time zero. From Lemma 3.2 it follows that the job of size 10 must be scheduled to complete last and thus $c = 1.2$. This implies that each one of the other two jobs should complete service at time 1.2, which is not feasible. □

Observe that in all the zero unfairness computations above the constant c relies on complete sample path knowledge and thus requires the scheduler to be anticipative, that is, base its decisions on future events (arrival times and service times of future arrivals). This leads to the following theorem:

Theorem 3.4. *Consider the set of all $G/G/1$ sample paths for which there exists a schedule of zero unfairness. Within this set there exist paths for which a zero unfairness schedule must be anticipative.*

Proof. Consider a non-anticipative scheduler denoted by π, i.e., π uses only past history to make scheduling decisions. Consider a sample path consisting of two jobs, J_1, J_2 with service times 1 and 2, respectively, arriving together at time $t = 0$. Let us examine the schedule generated by π till $t = 2$. We distinguish between two cases: (a) If J_1 was scheduled to complete service before time 2, then an arrival of a new job J_3 with unit service time at $t = 2$ will result in non-zero unfairness since according to Lemma 3.2 (applied on the sample path consisting of J_1, J_2 and J_3) we should have $T_1 = 2$. (b) If J_1 was scheduled to complete service at time 2 or later, then π will results in non-zero unfairness since according to Lemma 3.2 we should have $T_1 = 1.5$. □

4. Slowdown fairness analysis

Below we examine a variety of common queueing schedules and structures and evaluate their fairness via the SQF measure. We start with results for general systems, then analyze single-server single-queue systems and then turn to multi-server systems.

4.1. General systems

We consider a general system where the waiting time of a job is *independent* of its service time. In particular this holds for $GI/GI/1$ or $GI/GI/m$ systems with non-preemptive non-size-based scheduling policies. From (7) we get that the unfairness is given by

$$\mathrm{SQF} = \int_{x=0}^{\infty} E\left[\left(W(x) - x\frac{E[W]}{b_1}\right)^2\right] b(x)\mathrm{d}x = E[W^2] - \frac{2E[W]E[E[W(x)x]]}{b_1} + \frac{b_2 E^2[W]}{{b_1}^2}$$
$$= E[W^2] + E^2[W]\left(\frac{b_2}{{b_1}^2} - 2\right) = E^2[W](\gamma_W^2 + \gamma_S^2), \tag{11}$$

where γ_W^2 and γ_S^2 are the coefficients of variation of the waiting and service times, respectively. Note that the integral yields this result since the waiting and service times are independent and therefore $E[E[W(x)]] = E[W]$.

We note that when service times are constant, $\gamma_S = 0$ and therefore SQF $= \mathrm{Var}[T]$. This result is in agreement with seniority-based fairness measure developed in [6]. When service times are exponentially distributed, $\gamma_S = 1$ and (11) reduces to the second moment of the waiting time, i.e., SQF $= E[W^2]$.

4.2. Single-server single-queue analysis

In the following we consider single-server single-queue systems and examine several scheduling policies. We divide the analysis into non-size-based and size-based scheduling policies under general service times and then provide specific results for $M/M/1$ and $M/D/1$ systems.

4.2.1. Non-size-based non-preemptive policies

For a single-server system all non-size-based non-preemptive service orders have the same mean waiting time [21]. Hence, the following theorem follows from (11):

Theorem 4.1. *The SQF relative fairness ranking of non-preemptive and non-service-time-based (non-size-based) policies in a single-server system is determined only by the second moment of the waiting time. The unfairness value is increasing linearly in the second moment of the waiting time.*

Furthermore, it is well known that for a single-server system with a non-size-based non-preemptive service policy, the variance of the waiting time is smallest when the service policy is FCFS [21], and largest when the service policy is LCFS [22]. Therefore, we have that:

Corollary 4.1. *The unfairness value in a $G/GI/1$ system is minimized under an FCFS policy and maximized under an LCFS policy.*

Next we are interested in the relative fairness of various scheduling policies in this family of policies. We show that for $G/GI/1$, preferential service to more senior customers, who are served adjacently, always increases fairness.

Lemma 4.1. *Consider a $G/GI/1$ system. Let π be a non-preemptive non-service-time-based scheduling policy. Let J_i and J_j be two jobs which concurrently reside in the system under the scheduling policy π, and for which $a_i < a_j$. Let X_i and X_j be random variables denoting their service times, respectively. Assume that π schedules J_i and J_j consecutively, J_i ahead of J_j. Let π' be a scheduling policy that is identical to π except that the scheduling of J_i and J_j is done on reverse order. Then* $\mathrm{SQF}^{\pi} < \mathrm{SQF}^{\pi'}$.

Proof. Let P be a finite sample path (realization) that consists of the arrivals and service times, $a_1, \ldots a_n$ and $x_i, \ldots, x_n$ where $1 \le i < j \le n$. Thus the value of X_i is given by x_i and that of X_j by x_j. Let P' be a realization identical to P except for having the service times of J_i and J_j interchanged, that is $X_i = x_j$ and $X_j = x_i$. Due to the service times being i.i.d the probability density of P and P' are the same. Let ${w^i}_P(\pi)$ denote the waiting time experiences by J_i in P under π. Similarly define the waiting for all paths (P, P') of both jobs (J_i, J_j) under scheduling (π, π').

Let $a_j - a_i = \Delta > 0$. Then we have ${w^j}_P(\pi) = w_j^P(\pi) + x_i - \Delta$, ${w^i}_{P'}(\pi) = w_j^P(\pi)$, and $w_j^{P'}(\pi) = w_j^P(\pi) + x_j - \Delta$. Similarly for π' we have: ${w^i}_P(\pi') = w_j^P(\pi) - \Delta$, ${w^j}_P(\pi') = w_j^P(\pi) + x_j$, ${w^i}_{P'}(\pi') = w_j^P(\pi) - \Delta$,

$w^j_{P'}(\pi') = w^P_j(\pi) + x_i$. It is easy to see that the sum of these variables is the same for π and π', thus $E[W]$ is not affected by the scheduling policy. Thus, from Eq. (11) we have only to examine $E[W^2]$. The contribution of these terms to $E[W^2]$ under π is given by: $(w^P_j(\pi))^2 + (w^P_j(\pi) + x_i - \Delta)^2 + (w^P_j(\pi))^2 + (w^P_j(\pi) + x_j - \Delta)^2$, and their contribution under π' is given by: $(w^P_j(\pi) - \Delta)^2 + (w^P_j(\pi) + x_j)^2 + (w^P_j(\pi) - \Delta)^2 + (w^P_j(\pi) + x_i)^2$, and the difference between these contributions is $-2\Delta(x_i + x_j) < 0$. Since this holds for every sample path, $E[W^2]$ under π is smaller than that under π'. □

The relative fairness ranking of FCFS, NP-ROS (a random order of service policy under which the job in service is not preempted by new arrivals) and NP-LCFS is given below.

Corollary 4.2. *The unfairness values under SQF for FCFS, NP-ROS and NP-LCFS in a $G/GI/1$ system obey* $\text{SQF}^{\text{FCFS}} < \text{SQF}^{\text{NP-ROS}} < \text{SQF}^{\text{NP-LCFS}}$.

We note that this ordering is in agreement with intuition. Note also that one can propose a variety of policies that lie between FCFS and NP-LCFS, for example Second-Come–First-Served that selects for service the second most senior customer in the system whenever there are two or more customers in the system. The fairness ranking of such policies can be done directly from Lemma 4.1.

As an example we bring the special case of $M/GI/1$ systems in steady state, where we explicitly derive closed form expressions for the SQF values of three policies, FCFS, NP-LCFS, and NP-ROS. The second moment of the waiting time of these policies is given in [23]: $E[W^2]^{\text{FCFS}} = 2E^2[W] + \frac{\lambda b_3}{3(1-\rho)}$, $E[W^2]^{\text{NP-LCFS}} = \frac{1}{1-\rho}E[W^2]^{\text{FCFS}}$, and $E[W^2]^{\text{NP-ROS}} = \frac{1}{1-\rho/2}E[W^2]^{\text{FCFS}}$. Plugging the second moment expressions into (11) yields that

$$\begin{aligned}
\text{SQF}^{\text{FCFS}} &= \frac{\lambda b_3}{3(1-\rho)} + \frac{\lambda^2 b_2^2}{4(1-\rho)^2}\frac{b_2}{b_1^2}\\
\text{SQF}^{\text{NP-ROS}} &= \frac{\lambda b_3}{3(1-\rho)(1-\rho/2)} + \frac{\lambda^2 b_2^2}{4(1-\rho)^2}\left(\frac{b_2}{{b_1}^2} + \frac{2\rho}{2-\rho}\right)\\
\text{SQF}^{\text{NP-LCFS}} &= \frac{\lambda b_3}{3(1-\rho)^2} + \frac{\lambda^2 b_2^2}{4(1-\rho)^2}\left(\frac{b_2}{{b_1}^2} + \frac{2\rho}{1-\rho}\right).
\end{aligned} \tag{12}$$

4.2.2. Size-based, non-preemptive policies

In this section we complete our discussion of the SQF value for non-preemptive policies. Our analysis is based on observing that $\lambda \int_0^\infty xE[W(x)]b(x)\mathrm{d}x$, the scaled cross-correlation term in (7), is the average amount of unfinished work in the queue and that this quantity is known for any work-conserving $M/G/1$ system. Using this observation we show that for all non-preemptive policies in a $M/G/1$ system, the SQF formula reduces to the same expression.

Theorem 4.2. *The unfairness of any size-based non-preemptive policy in an $M/G/1$ system is given by (11), namely,* $\text{SQF} = E[W^2] + E^2[W](\frac{b_2}{{b_1}^2} - 2)$.

Proof. The proof is based on the $M/G/1$ conservation law [13], which states that for any $M/G/1$ priority queue and any work-conserving non-preemptive ordering policy it must be that $\sum_{i=1}^k \lambda p_i W_i = \rho E[W]$, where k is the number of priority classes, and p_i and W_i are the portion and average waiting time of customers in priority i, respectively. Assume service times vary between y_0 and y_n. Define $n+1$ boundary points $y_0, \ldots, y_n$ such that $y_{i-1} < y_i$, $1 \le i \le n$ and assign all customers of service time (y_{i-1}, y_i) to class i. As $n \to \infty$, the conservation law yields that $\lambda \int_0^\infty xE[W(x)]b(x)\mathrm{d}x = \rho E[W]$. Plugging this expression into (7) gives the desired result. □

4.2.3. Preemptive policies

We analyze policies that can be set up as $M/G/1$ symmetric queues, as defined in Section 2. These policies include PS, P-LCFS, P-ROS, and a variety of policies that lie in between them. One such example is the *finite-processor-sharing*, where the processor is shared by all jobs when the number of jobs in the system does not exceed r (the number of service positions), and shared by only r of them and the remaining ones are kept waiting when the number of jobs present exceeds r. Recall from Section 2 that for all members of this class the expected sojourn time of a job

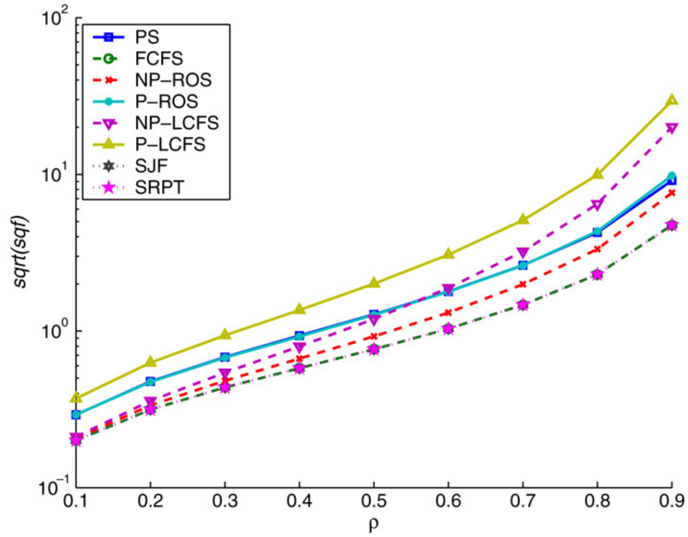

Fig. 1. The SQF unfairness of common scheduling policies in an $M/D/1$ system.

requiring x units of service is given by $E[T(x)] = \frac{x}{1-\rho}$. By substituting $\alpha = \frac{1}{1-\rho}$, $\beta = 0$ and $E[T] = \frac{b_1}{1-\rho}$ into (8) we obtain

$$\mathrm{SQF}^{SQ} = E[T^2] - \frac{b_2}{(1-\rho)^2} = \mathrm{Var}[T]. \tag{13}$$

Theorem 4.3. *The SQF relative fairness ranking of $M/G/1$ symmetric queue policies is determined only by the variance of the sojourn time. The SQF value is linear in* $\mathrm{Var}[T]$*, but affine in* $\mathrm{Var}[W]$*, since* $\mathrm{Var}[T] = \mathrm{Var}[W] + \mathrm{Var}[B]$.

Note the similarity of this result to that of Theorem 4.1 where in both cases the relative ranking follows the variance of the sojourn time. This follows from (8). Moreover, Theorem 4.3 implies that

Observation 4.1. *If the conditional waiting time of PS is smallest for each service-time requirement, as conjectured by* [19]*, then the unfairness value in a $M/G/1$ symmetric queue is minimized under a PS policy.*

In [19] it is shown that the unfairness of PS (which is of a rather complex form) is lower than that of P-LCFS (given as $\mathrm{SQF}^{\text{P-LCFS}} = \frac{\rho b_2}{(1-\rho)^3}$). If the conjecture of [19] is true, then also the PS unfairness is lower than that of P-ROS. Indeed, our calculations and simulations, shown in Figs. 1–3 of Section 5, conform with the conjecture.

The analysis of size-based policies, such as Shortest-Remaining-Processing-Time (SRPT) where at any point in time the server is processing the job with the shortest remaining processing time, is out of the scope of this paper and is a subject of ongoing work

4.2.4. Fairness in $M/D/1$ and $M/M/1$ models

Next we compare the fairness of the following policies, P-LCFS, PS, FCFS, NP-LCFS and NP-ROS for the $M/D/1$ and $M/M/1$ models.

The unfairness of non-preemptive policies is derived by substituting the second and third moments of the service time, namely $b_2 = 2(b_1)^2$, $b_3 = 6(b_1)^3$ for exponential service time and $b_2 = (b_1)^2$, $b_3 = (b_1)^3$ for deterministic service time, in (12). The unfairness for $M/M/1/PS$ is derived from the conditional sojourn time given in [13]: $\mathrm{Var}[T(x)]^{PS}_{M/M/1} = \frac{2\rho x}{(1-\rho)^3} - \frac{2\rho}{\mu(1-\rho)^4}\left(1 - e^{-\mu x(1-\rho)}\right)$. The unfairness for $M/D/1/PS$ is derived from the Laplace–Stieltjes transform of the sojourn time distribution given in [24]: $L_T(s) = E[e^{-sT}] = \frac{(1-\rho)(\lambda+s)^2 e^{-(\lambda+s)b_1}}{s^2+\lambda(s+s(1-\rho)+\lambda(1-\rho))\mathrm{e}^{-(\lambda+s)b_1}}$, i.e., $\mathrm{Var}[T]^{PS}_{M/D/1} = \left.\frac{\mathrm{d}^2 L_T(s)}{\mathrm{d}s}\right|_{s=0} - \left(-\left.\frac{\mathrm{d}L_T(s)}{\mathrm{d}s}\right|_{s=0}\right)^2$.

Table 1 provides a summary of the SQF values for the various policies considered. Thus, for $M/D/1$ the fairness ranking is $\mathrm{SQF}^{\text{P-LCFS}} > \mathrm{SQF}^{\text{NP-LCFS}} > \mathrm{SQF}^{\text{NP-ROS}} > \mathrm{SQF}^{\text{FCFS}}$ whereas for $M/M/1$ the ranking is $\mathrm{SQF}^{\text{P-LCFS}} = \mathrm{SQF}^{\text{NP-LCFS}} > \mathrm{SQF}^{\text{NP-ROS}} > \mathrm{SQF}^{\text{FCFS}}$ (recall that for both systems we have $\mathrm{SQF}^{\text{P-LCFS}} \geq \mathrm{SQF}^{PS}$).

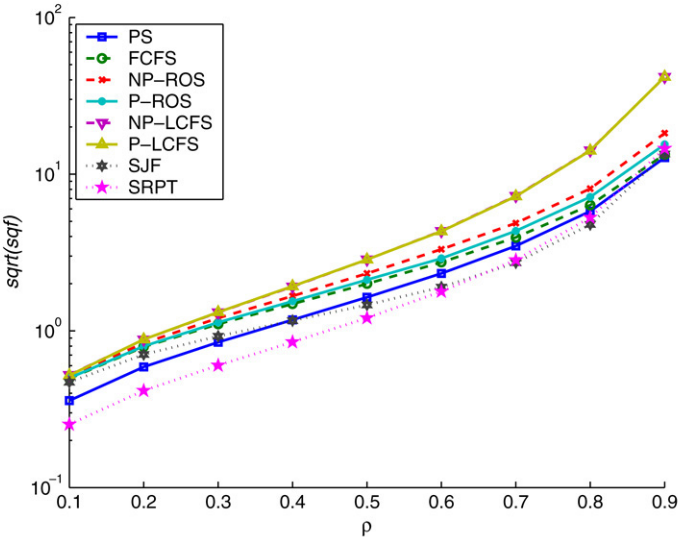

Fig. 2. The SQF unfairness of common scheduling policies in an $M/M/1$ system.

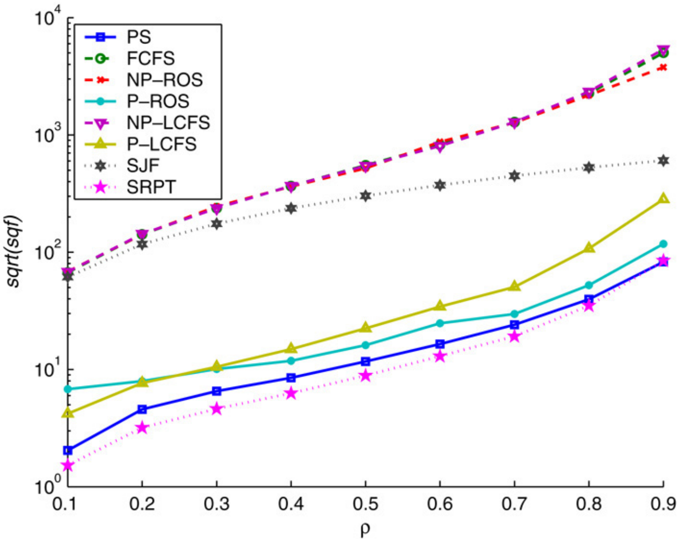

Fig. 3. The SQF unfairness of common scheduling policies in an $M/GI/1$ system with coefficient of variation 10.

4.3. Multi-server multi-queue analysis

Next we analyze multi-queue and multi-server systems. Note that the analysis of fairness in queueing systems has been limited, so far, mainly to single-server single-queue systems, with the exception of [18]. The results below demonstrate that SQF has good potential in dealing with such complex systems.

We consider two 'classical' queueing systems, the $m * M/M/1$, and $M/M/m$, both with arrival rate λ, total service rate $1/b_1$, and FCFS service policy. The first consists of m parallel queues each equipped with its own server; the (Poisson) arrival stream is split (using a random mechanism) into m Poisson streams of rate λ/m each; thus, the system consists of m independent $M/M/1$ queues each served by a server with mean service rate $\frac{1}{mb_1}$. The second system is an $M/M/m$ system with arrival rate λ and an average service time mb_1. Define $\rho = \lambda b_1$.

Operational question of interest is whether multiplicity of servers (processors) and/or queues improves system performance (efficiency and fairness). Analysis of these two systems and their comparison, from *efficiency* (mean waiting time) point of view, is provided in [13] Ch. 5.1. Their relative fairness values are calculated and compared below.

Table 1
SQF expressions for $M/M/1$ and $M/D/1$

	$M/D/1$	$M/M/1$
$\mathrm{SQF}^{\mathrm{FCFS}}$	$\frac{{b_1}^2\rho(4-\rho)}{12(1-\rho)^2}$	$\frac{2{b_1}^2\rho}{(1-\rho)^2}$
$\mathrm{SQF}^{\mathrm{NP\text{-}ROS}}$	$\frac{{b_1}^2\rho(8-2\rho+3\rho^2)}{12(2-\rho)(1-\rho)^2}$	$\frac{2{b_1}^2\rho}{(1-\rho/2)(1-\rho)^2}$
$\mathrm{SQF}^{\mathrm{NP\text{-}LCFS}}$	$\frac{{b_1}^2\rho(4-\rho+3\rho^2)}{12(1-\rho)^3}$	$\frac{2{b_1}^2\rho}{(1-\rho)^3}$
$\mathrm{SQF}^{\mathrm{P\text{-}LCFS}}$	$\frac{{b_1}^2\rho}{(1-\rho)^3}$	$\frac{2{b_1}^2\rho}{(1-\rho)^3}$
$\mathrm{SQF}^{\mathrm{PS}}$	$\frac{2e^{\rho}(-1+\rho)-\rho^2+2}{\lambda^2(1-\rho)^2}$	$\frac{2{b_1}^2\rho(1/b_1+\rho-2)}{(1-\rho)^3(\rho-2)}$

Recall that the SQF measure reduces to $E[W^2]$ in a non-preemptive system with exponential service-time distribution. The second moment of the waiting time $E[W^2]$ in an $M/M/m$ system is given by [25]: $2C(\lambda, m, m\rho)b_1^2/(1-\rho)$, where

$$C(\lambda, m, a) = \left(\sum_{i=0}^{m-1} \frac{a^i}{i!} + \frac{a^m}{m!}\frac{1}{1-\rho}\right)^{-1} \frac{a^m}{m!(1-\rho)} \tag{14}$$

is the probability of waiting in the system, known as the Erlang-C formula. Since the Erlang-C formula reduces to ρ for $M/M/1$, the relationship between the fairness values of the queueing systems is $\mathrm{SQF}(m * M/M/1) = m^2\mathrm{SQF}(M/M/1)$ and $\mathrm{SQF}(M/M/m) = \rho^{-1}C(\lambda, c, c\rho)\mathrm{SQF}(M/M/1)$. Therefore, an $M/M/1$ may be viewed as m^2 times more 'fair' than an $m * M/M/1$ system. This result fits intuition since the $M/M/1$ schedule preserves seniority (serving by order of arrival) while the $m * M/M/1$ allows serious order violations.

To address the question of how $M/M/1$ compares to $M/M/m$ in terms of their fairness values we observe that the scaling factor $\rho^{-1}C(\lambda, m, m\rho)$ has a probabilistic interpretation, it is the probability that a customer sees $m-1$ busy servers upon arrival to an $M/M/m$ system, and thus must be smaller than one (given a stable system and $m > 1$). This yields the relative fairness ranking of the systems for $m > 1$: $\mathrm{SQF}(M/M/m) < \mathrm{SQF}(M/M/1) < \mathrm{SQF}(m * M/M/1)$, i.e., $M/M/m$ is the most fair. This, result fits intuition well since an $M/M/m$, $m > 1$ queue allow customers with 'short' service requirements to complete service before customers with 'long' service requirements even when 'shorts' arrives after 'longs'.

We now show that, for any $M/M/m$, fairness increases in m.

Theorem 4.4. *Consider a set of $M/M/m$ systems, $m = 1, 2, \ldots$ all having arrival rate λ and a total service rate $\tilde{\mu}$; let μ be the service rate of a single server, thus $\mu = \tilde{\mu}/m$. Let $\rho = \lambda/\tilde{\mu}$ be the system utilization. The unfairness value associated with these systems decreases as the number of servers increases, i.e.,* $\mathrm{SQF}(M/M/m+1) < \mathrm{SQF}(M/M/m)$.

Proof. Recall that $\mathrm{SQF} = 2C(\lambda, m, m\rho)b_1^2/(1-\rho)$ in an $M/M/m$ system. The SQF value depends on the number of servers only through the factor $C(\lambda, m, m\rho)$. Thus it is sufficient to show that $C(\lambda, m, m\rho)$ is monotonically non-increasing in m. □

Lemma 4.2. *Let $S(\lambda, m, \tilde{\mu}/m)$ be an $M/M/m$ system with arrival rate λ and m servers, each operating at rate $\tilde{\mu}/m$. Let $C(\lambda, m, m\rho)$ be the probability of queueing in $S(\lambda, m, \mu/m)$, that is the Erlang-C formula as defined above. Then for $m = 1, 2, \ldots, C(\lambda, m, m\rho)$ is monotonically non-decreasing in m.*

Proof. Let $p_k^{(m)}$, $k = 0, \ldots, 1$ be the probability of having k customers in system $S(\lambda, m, \tilde{\mu}/m)$. Let $Q^{(m)} = \sum_{i=c}^{\infty} p_k^{(m)}$ be the probability of queueing in $S(\lambda, m, \tilde{\mu}/m)$ $(=C(\lambda, m, m\rho))$. Similarly let $N^{(m)} = \sum_{i=0}^{m-1} p_k^{(m)}$ be the probability of not queueing. For each system we will examine its Markov chain and compare the states whose probabilities are summed up by $Q^{(m)}$ to those summed up by $N^{(m)}$. It is easy to see that $Q^{(m)} = p_m^{(m)}/(1-\lambda/\tilde{\mu})$ for all these systems.

To evaluate $N^{(m)}$ we will express each of its terms relatively to $p_m^{(m)}$. To this end note that for state i, $i = 0, \ldots, m-1$, the arrival rate is λ and the departure rate is $i\tilde{\mu}/m$. Thus $p_i^{(m)} = p_m^{(m)} \frac{(\tilde{\mu})^{m-i}(i+1)\cdot\cdots\cdot m}{\lambda^{m-i}m^{m-i}}$ and

$N^{(m)} = p_m^{(m)} \sum_{i=0}^{m-1} \frac{(i+1)\cdot\cdots\cdot m}{\rho^{m-i}(m)^{m-i}}$. Next we show that $N^{(m)}/p_m^{(m)} \geq N^{(m-1)}/p_{m-1}^{(m-1)}$. To show this we note that $N^{(m)}$ consists of one more term (which is positive) and thus it is sufficient to show that each of the last $m-1$ terms of $N^{(m)}/p_m^{(m)}$ is greater than or equal to its counterpart term in $N^{(m-1)}/p_{m-1}^{(m-1)}$, that is that for $0 \leq i \leq m-1$, $\frac{p_i^{(m)}}{p_m^{(m)}} \geq \frac{p_{i-1}^{(m-1)}}{p_{m-1}^{(m-1)}}$. Namely, we have to show that $\frac{(i+1)\cdot\cdots\cdot m}{\rho^{m-i} m^{m-i}} \geq \frac{i\cdot\cdots\cdot(m-1)}{\rho^{m-i}(m-1)^{m-i}}$. Since each of these products consists of exactly $m-i$ factors, it is sufficient to show that $\frac{j}{\rho m} \geq \frac{j-1}{\rho(m-1)}$, for $j = i+1, \ldots, m$, which holds since $j \leq m$.

Finally, since we have $N^{(m)}/p_m^{(m)} \geq N^{(m-1)}/p_{m-1}^{(m-1)}$ and $Q^{(m)}/p_m^{(m)} = Q^{(m-1)}/p_{m-1}^{(m-1)}$, then combining this with $N^{(m)} + Q^{(m)} = N^{(m-1)} + Q^{(m-1)} = 1$, implies, $Q^{(m)} \leq Q^{(m-1)}$. □

5. Measure evaluation and policy comparison

In this section we examine the properties of SQF by evaluating its values for a variety of scheduling policies. We consider both non-size-based policies (FCFS, NP-ROS and NP-LCFS) and size-based (including preemptive) policies (PS, P-ROS, P-LCFS, Shortest-Job-First (SJF)) and SRPT. Policies whose SQF value is not available in closed form from the previous section (e.g., SRPT) are evaluated via a simulation.

An emphasis is given in the evaluation to examine how the measure reacts to the two underlying physical factors of relative seniority and relative service time. A good test for this, for systems under steady state, is to vary the service time variability: When it is small (or zero) service-time variations do not exist and thus the only factor affecting fairness is relative seniority. When the variability is large, service-time variations become more radical than seniority variations and then play a major role. We thus evaluate these values for an $M/GI/1$ system, considering three cases of service time variability: (i) No variability ($M/D/1$), (ii) medium variability ($M/M/1$) and (iii) very high variability ($M/GI/1$ with service-time coefficient of variation equaling 10, namely $b_1 = 1$ and $b_2 = 101$). Figs. 1–3 show the corresponding unfairness plots.

A representative comparison, in this context, is between the unfairness values of FCFS and P-LCFS, where the former emphasizes seniority and disregards service times completely, while the latter disregards seniority and emphasizes service times (implicit prioritization of short jobs). Examination of Eqs. (12) and (13) reveals that

$$\frac{\mathrm{SQF}^{\mathrm{FCFS}}}{\mathrm{SQF}^{\mathrm{P\text{-}LCFS}}} \approx \frac{\lambda b_2^2(1-\rho)}{4b_1^3}. \tag{15}$$

This implies that if service-time variability is not high, FCFS is more fair due to its emphasis on seniority. However, once service time variability gets very large, Eq. (15) reveals that FCFS is more unfair; this is due to the fact that serving a very large job ahead of a very short job becomes a dominant unfairness factor and then P-LCFS, which prioritizes small jobs, becomes more fair. Nevertheless, once the moments of service time are fixed (at some values), when utilization approaches unity we get $\lim_{\rho\to 1} \mathrm{SQF}^{\mathrm{FCFS}}/\mathrm{SQF}^{\mathrm{P\text{-}LCFS}} = 0$, namely FCFS is more fair. The reason for this is that at very high load the seniority violation introduced by P-LCFS increases drastically (some jobs are served immediately, while others wait a whole busy period), diminishing the "size unfairness" of FCFS. Thus, the unfairness values of SQF follow intuition quite closely.

It is interesting to examine how alternative fairness approaches treat the relative fairness of P-LCFS and FCFS. First, the criterion offered in [4] classifies P-LCFS as 'always fair' and FCFS as 'always unfair'. The reason is that the criterion accounts for the *mean* slowdown of each job size and thus its sensitivity to relative seniority is probably very small (if any). In contrast, the RAQFM measure offered in [9], whose unfairness values are depicted in Figs. 5 and 4, shows very similar behavior to that of SQF: Lower unfairness to P-LCFS for large variability service times, and lower unfairness to FCFS for low–medium variability service times and for very high loads.[2]

An evaluation of the fairness values for all policies yields the following observations:

(1) For all policies examined that the unfairness measure increases with service-time variability. One can observe it either from the expressions derived in Section 4 or from the figures. This growth is more radical for the non-size-based policies and milder for the size-based policies. This agrees with intuition since the latter policies do not let drastic size-driven discriminations to occur.

[2] At some values of the parameters SQF and RAQFM differ on the relative ranking of P-LCFS and FCFS. However, they do agree in their general behavior and when $\rho \to 1$ both measures consider P-LCFS to be more unfair.

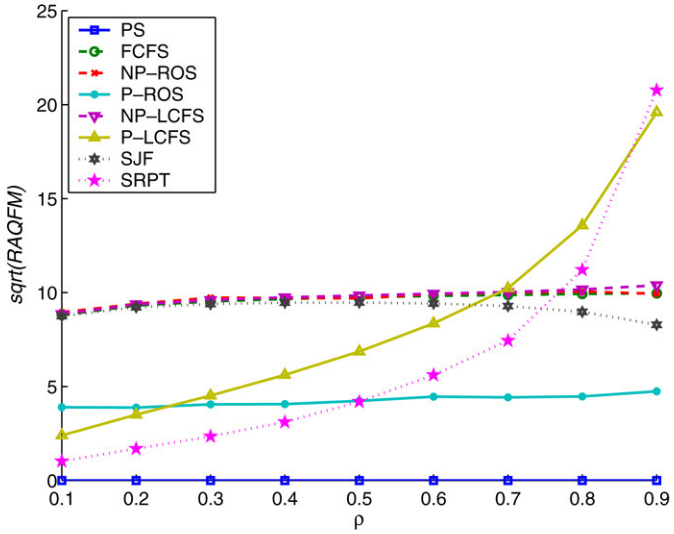

Fig. 4. The RAQFM unfairness of common scheduling policies in an $M/GI/1$ system with coefficient of variation 10.

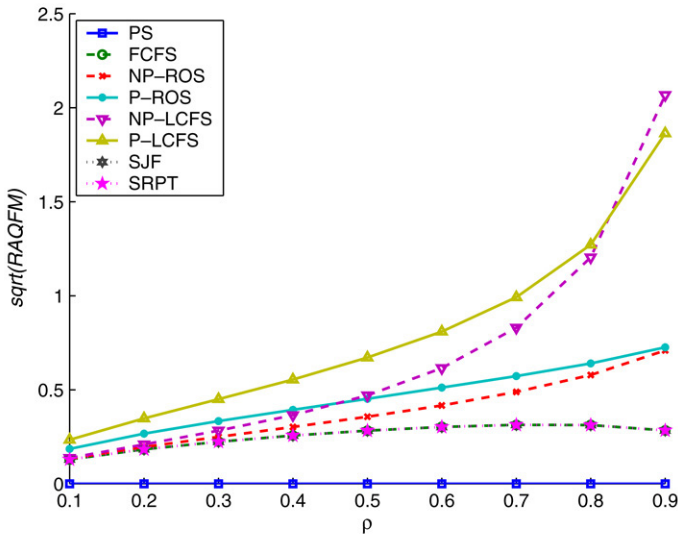

Fig. 5. The RAQFM unfairness of common scheduling policies in an $M/D/1$ system.

(2) For small variability (deterministic service times) the most fair policies are those that are based on seniority, that is FCFS and SRPT (which coincides with FCFS for deterministic service times). The most unfair policies are those that drastically violate seniority, namely P-LCFS and NP-LCFS.

(3) For medium variability P-LCFS remains the most unfair policy. SRPT and PS are the most fair policies.

(4) For high variability the non-size-based policies (FCFS, NP-LCFS, NP-ROS) become the most unfair policies. This results from the fact that very large jobs that enter service are not interrupted, and the short jobs that arrive during this period are all subject to very large values of slowdown (residual service time of the large job divided by service time of short job). Further, the variability across these small jobs is huge, since some of them experience the delay only of the large job while other experience, in addition, the delay due to many jobs that arrived during the busy period generated by the large job. Under high variability SRPT and PS remain the most fair policies, which is intuitive since they prioritize short jobs without introducing too much seniority violation.

(5) It is interesting to compare some of these results to those of [4] which used slowdown as a fairness criterion. Thus, SQF shares with that criterion the underlying principle that waiting times should be proportional to service times.

Under that criterion P-LCFS and PS are 'always fair' (interpreted as being the most fair out of all policies studied there); in contrast, SRPT is only 'sometimes fair' since at high loads $E[T(x)]/x > 1/(1-\rho)$ for large values of x. In contrast, SRPT becomes the most fair policy under SQF, which compares *all jobs* to each other, since it reacts to relative seniority better than the other policies. The reason for this difference is that the sensitivity of the criterion to relative seniority is small.

6. Comparison to other approaches

A few other approaches (in addition to [4]) have been proposed in recent years for quantitatively measuring fairness of queueing systems. These are [5–7] (and [8]) and [9], where the first two can be described as seniority-based evaluation, the third as seniority and service requirement-based measure, and the fourth as resource allocation-based measure. A comprehensive overview of alternative measures for *job fairness* is available in [12]. Other approaches, aiming at fairness of parallel job processing are presented in [10] and [11] (where the results are mainly simulative).

Compared to the approach of [4] the advantage of SQF is in providing a full scale measure and in reacting to job seniority; a disadvantage is that it is mathematically harder to compute. Compared to [5] and [6] its advantage is in reacting to job sizes. Compared to [9] its advantages are in simplicity of analysis and simpler applicability to large systems, while its disadvantage is in not admitting the locality of reference property [26] (i.e., the discrimination distribution within busy periods is not equal to the distribution across the whole population). The latter property follows immediately from the definition of the proportionality constant c given in (5), which is based on the sojourn times of the whole population. Thus, if one believes that for fairness purposes only discriminations within busy periods should be considered, SQF might yield non-precise values. Comparison to [7] is limited since the results under that measure have been somewhat limited.

In general, it seems that a major advantage of SQF is in its simple applicability to complex systems. This is a subject of an ongoing research.

References

[1] R.C. Larson, Perspective on queues: Social justice and the psychology of queueing, Operations Research 35 (1987) 895–905.
[2] M.H. Rothkopf, P. Rech, Perspectives on queues: Combining queues is not always beneficial, Operations Research 35 (1987) 906–909.
[3] A. Rafaeli, G. Barron, K. Haber, The effects of queue structure on attitudes, Journal of Service Research 5 (2) (2002) 125–139.
[4] A. Wierman, M. Harchol-Balter, Classifying scheduling policies with respect to unfairness in an $M/GI/1$, in: Proceedings of ACM Sigmetrics 2003 Conference on Measurement and Modeling of Computer Systems, San Diego, CA, 2003, pp. 238–249.
[5] E. Sherry-Gordon, New problems in queues: Social injustice and server production management, Ph.D. Thesis, MIT, May 1987.
[6] B. Avi-Itzhak, H. Levy, On measuring fairness in queues, Advances in Applied Probability 36 (3) (2004) 919–936.
[7] W. Sandmann, A discrimination frequency based queueing fairness measure with regard to job seniority and service requirement, in: R. Sabella (Ed.), Proceedings of the 1st Euro NGI Conference on Next Generation Internet Networks — Traffic Engineering, NGI2005, IEEE Computer Society Press, Rome, Italy, 2005.
[8] W. Sandmann, Analysis of a queueing fairness measure, in: Proceedings of the 13th GI/ITG Conference on Measurement, Modelling and Evaluation of Computer and Communication Systems, MMB2006, Nürnberg, Germany, 2006, pp. 219–231.
[9] D. Raz, H. Levy, B. Avi-Itzhak, A resource-allocation queueing fairness measure, in: Proceedings of Sigmetrics 2004/Performance 2004 Joint Conference on Measurement and Modeling of Computer Systems, New York, NY, 2004, pp. 130–141, Performance Evaluation Review 32 (1) (2004) 130–141.
[10] G. Sabin, P. Sadayappan, Unfairness metrics for space-sharing parallel job schedulers, in: Proceedings of the 11th Inernational Workshop on Job Scheduling Strategies for Parallel Processing, JSSPP 2005, Boston, MA, USA, 2005.
[11] S. Vasupongayya, S. Chiang, On job fairness in non-preemptive parallel job scheduling, in: Proceedings of The 17th IASTED International Conference on Parallel and Distributed Computing and Systems, PDCS 2005, Phoenix, AZ, USA, 2005.
[12] B. Avi-Itzhak, H. Levy, D. Raz, Quantifying fairness in queueing systems: Principles, approaches and applicability, Tech. Rep. RRR-25-2005, RUTCOR, Rutgers University, 2005 (submitted for publication).
[13] L. Kleinrock, Computer Applications, in: Queueing Systems, vol. 2, Wiley, 1976.
[14] N. Bansal, M. Harchol-Balter, Analysis of SRPT scheduling: Investigating unfairness, in: Proceedings of ACM Sigmetrics 2001 Conference on Measurement and Modeling of Computer Systems, 2001, pp. 279–290.
[15] M. Bender, S. Chakrabarti, S. Muthukrishnan, Flow and stretch metrics for scheduling continuous job streams, in: Proceedings of the 9th Annual ACM-SIAM Symposium on Discrete Algorithms, San Francisco, CA, 1998, pp. 270–279.
[16] M. Harchol-Balter, K. Sigman, A. Wierman, Asymptotic convergence of scheduling policies with respect to slowdown, Performance Evaluation 49 (1–4) (2002) 241–256.
[17] A. Wierman, M. Harchol-Balter, Classifying scheduling policies with respect to higher moments of conditional response time, in: Proceedings of ACM Sigmetrics 2005 Conference on Measurement and Modeling of Computer Systems, Banff, Alberta, Canada, 2005, pp. 229–239.

[18] D. Raz, B. Avi-Itzhak, H. Levy, Considerations of scheduling in multi-server and multi-queue systems, in: Proceedings of Valuetools 2006, Pisa, Italy, 2006.
[19] B. Avi-Itzhak, S. Halfin, Server sharing with a limited number of service positions and symmetric queues, Journal of Applied Probability 24 (1987) 990–1000.
[20] S. Yashkov, Properties of invariance of probabilistic models of adaptive scheduling in shared-use systems, Automatic Control and Computer Sciences 12 (1980) 56–62.
[21] J.F.C. Kingman, The effect of queue discipline on waiting time variance, Proceedings of the Cambridge Philosophical Society 58 (1962) 163–164.
[22] D.G. Tambouratzis, On a property of the variance of the waiting time of a queue, Journal of Applied Probability 5 (1968) 702–703.
[23] H. Takagi, Vacation and Priority Systems (Part 1), in: Queueing Analysis, A Foundation of Performance Evaluation, vol. 1, North-Holland, Amsterdam, The Netherlands, 1991.
[24] T. Ott, The sojourn-time distribution in the $M/G/1$ queue with processor sharing, Journal of Applied Probability (21) (1984) 360–378.
[25] J. Medhi, Stochastic models in queueing theory, 2nd edition, Academic Press Professional, Inc., 2003.
[26] H. Levy, D. Raz, B. Avi-Itzhak, Locality of reference and the use of sojourn time variance for queue fairness, Operations Research Letters 35 (2007) 311–318.

Benjamin Avi-Itzhak is currently a Professor II at Rutgers University (A Fellow of RUTCOR with a joint appointment in the Department of Management Science and Information Systems of the School of Business). He holds a B.Sc. degree in Mechanical Engineering, a B.Sc. in Industrial Engineering and Management and M.Sc. and D.Sc. degrees in Operations Research, all from Technion I.I.T. Dr. Avi-Itzhak served in professorial positions at Technion I.I.T., Cornell University, Case Institute of Technology, Hebrew University, Tel Aviv University, City University of NY, Haifa University and Stanford University and as a visiting scientist at Bell Labs and at Bellcore. His interests, as reflected in his publications, cover a wide spectrum of application areas, including electric power systems, telecommunications and data networks, computer systems, supply chain management, marine shipping, transportation, and service systems. He also served in administrative, academic and professional capacities and as a consultant and scientific advisor to numerous government, public and private sector organizations.

Eli Brosh received his B.Sc. and M.Sc. degrees from Tel Aviv University, Israel. He worked in the telecom industry as a systems engineer and architect. He is currently a Ph.D. student in the Department of Computer Science of Columbia University, USA. His research interests are design and analysis of algorithms for communication networks and performance evaluation of network systems.

Hanoch Levy received the B.A. degree in computer science with distinctions from the Technion, Israel Institute of Technology, Haifa, Israel, in 1980, and the M.Sc. and the Ph.D. degrees in computer science from the University of California at Los Angeles in 1982 and 1984, respectively. From 1984 to 1987, he was a Member of Technical Staff in the Department of Teletraffic Theory at AT&T Bell Laboratories. Since 1987, he has been with the School of Computer Science, Tel Aviv University, Tel Aviv, Israel. During 1992–1996, he was with the School of Business and RUTCOR, Rutgers University, New Brunswick, NJ, on leave from Tel Aviv University. His interests include computer communications networks, wireless networks, quality of service in high-speed networks, performance evaluation of computer systems and queueing theory.