

R U T C O R  
R E S E A R C H  
R E P O R T

SQF: A SLOWDOWN QUEUEING  
FAIRNESS MEASURE

Benjamin Avi-Itzhak <sup>a</sup>      Hanoch Levy <sup>b</sup>

Eli Brosh <sup>c</sup>

RRR 2-2006, JANUARY 2006

RUTCOR  
Rutgers Center for  
Operations Research  
Rutgers University  
640 Bartholomew Road  
Piscataway, New Jersey  
08854-8003  
Telephone:      732-445-3804  
Telefax:        732-445-5472  
Email:      rrr@rutcor.rutgers.edu  
<http://rutcor.rutgers.edu/~rrr>

---

<sup>a</sup>RUTCOR, Rutgers University, New Brunswick, NJ, USA, avi-  
itzha@rutcor.rutgers.edu

<sup>b</sup>School of Computer Science, Tel-Aviv University, Tel-Aviv, Israel,  
hanoch@cs.tau.ac.il

<sup>c</sup>Department of Computer Science, Columbia University, New-York, NY,  
USA, elibrosh@cs.columbia.edu

# RUTCOR RESEARCH REPORT

RRR 2-2006, JANUARY 2006

## SQF: A SLOWDOWN QUEUEING FAIRNESS MEASURE

**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 are derived.

# 1 Introduction

Fairness is an important factor in queues operation and scheduling (e.g. [6], [9], [7]). In a recent work [3], surveying proposed approaches to defining and measuring queueing fairness and discussing their underlying principles, the authors suggest development of a fairness measure based on the proportionality principle, namely, waiting time of a job should be in proportion to the servers' time provided to it. It is intuitively appealing to require that customers who get more will also wait more. In this paper we propose and analyze a proportionality fairness measure based on discriminations of individual customers. The discrimination of a customer is defined 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 is thus measured by the magnitude of the deviation from the absolutely fair waiting time  $W(x) = px$ , which may be positive, negative or zero. Denoting the sojourn time of the customer by  $T(x)$ , where  $T(x) = W(x) + x$ , its slowdown is defined as  $S(x) = T(x)/x$ . Requiring that  $W(x) = px$  is the same as requiring that  $S(x) = c$ , where  $c = p + 1$ , i.e. the proportionality fairness principle can be stated as: The slowdown of all customers must be the same.

Expected slowdown has been mentioned in the context of fairness as early as [5], 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  $E(S(x)) = 1/(1-\rho)$ . In [12] it is shown that no M/G/1 scheduling policy can attain a same constant slowdown lower than  $1/(1-\rho)$  for all values of  $x$ . The authors of [12] used this result to suggest a criterion of fairness stating that a policy is fair if for all  $x$   $E(S(x)) \leq 1/(1-\rho)$ . This leads to the classification of M/G/1 disciplines into "always fair", "sometimes fair" or "never fair".

A few other approaches have been proposed in recent years for quantitatively measuring fairness of queueing systems. These are [11] [2], [10] and [8], where the first two can be described as seniority based evaluation, the third as seniority and service requirement based measure, and fourth as resource allocation based measure.

The structure of this paper is as follows: In Section 2 we discuss the mean slowdown approach and explain its limitations. In Section 3 we construct the slowdown fairness measure (SQF) and derive its basic properties. In Section 4 we analyze GI/GI/1 systems under a

variety of scheduling policies and derive for them their SQF values. For the specific cases of M/GI/1 systems closed form expressions are given for a number of scheduling policies. In Section 5 we discuss the results and evaluate the measure by examining its values for the various scheduling policies.

## 2 Expected Slowdown and Fairness

[12] proposes to use the expected slowdown as a fairness criterion for M/G/1 queues. Specifically, it defines an M/G/1 policy to be 'always fair' if for every  $x$ ,  $E(S(x)) \leq 1/(1 - \rho)$ . Accordingly the authors classify the preemptive-last-come-first-served (P-LCFS) and processor-sharing (PS) policies as always unfair. 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 ([1]) 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  $1, 2, \dots, n$  and a proportion  $\phi(i, n)$  of the service rate is directed at the job in position  $i$ , ( $i = 1, 2, \dots, \min(r, n)$ ) where  $\sum_{i=1}^{\min(r, n)} \phi(i, n) = 1$ ,  $n = 1, 2, \dots$ . A newly arrived job which encounters  $n$  in the system,  $n = 0, 1, 2, \dots$ , 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 reordering 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 [14], that all members of this class possess the equilibrium properties

of symmetric queues as defined in [4]. 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  $E(T(x)) = E(W(x)) + x = x/(1-\rho)$ . By the criterion proposed in [12] all 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  $r \rightarrow \infty$  and an indefinitely large class of preemptive queues for each finite  $r$ . In particular, in the case of  $r = 1$  P-LCFS (called the stack by [4]) is the one where a displaced customer is always placed in position 2 and all customers in the waiting positions are moved one position back, and the one in position 2, if not empty, is always the one to enter service when a customer departs. 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  $W^\infty$  for the PS case. [1] shows that  $Var(W^f(x)) \geq Var(W^l(x))$  and  $Var(W^f(x)) \geq Var(W^\infty(x))$ . The authors of [1] conjecture that the variance of the conditional delay is decreasing in  $r$  and therefore  $Var(W(x)) \geq Var(W^\infty(x))$ , where  $W(x)$  is the conditional waiting time in any arbitrary policy in this class. For the case of exponentially distributed service times they show that

$$\frac{Var(W^f(x)) - Var(W^l(x))}{Var(W^f(x))} = \frac{\rho[1 - e^{-\lambda(1-\rho)x}]}{\lambda(1-\rho)x} \geq 0.$$

For a fixed value of  $x$ , this expression increases with  $\rho$  and goes to  $\infty$  when  $\rho \rightarrow 1$ . For the exponential case they also show that  $Var(W^l(x)) \geq Var(W^\infty(x))$ . 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  $x = 0.5$  we get  $Var(W^f(0.5)) = 90.4Var(W^\infty(0.5)) = 18.5Var(W^l(0.5))$ . 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 use of expected slowdown might not fully reflect fairness and a more sensitive measure might be called.

### 3 Definition of the Slowdown Queueing Fairness (SQF) Measure and its Properties

Assume an arbitrary queueing system where customers  $C_1, C_2, \dots$  (interchangeably called jobs  $J_1, J_2, \dots$ ) 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  $T_i = d_i - a_i$ , 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  $J_1, \dots, J_N$ . 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  $D_i = cx_i - T_i$ , 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. 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

$$\sum_{i=1}^N D_i = \sum_{i=1}^N cx_i - \sum_{i=1}^N T_i = 0, \quad (1)$$

which leads to:

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

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

$$\sum_{i=1}^N D_i^2 = \frac{1}{N} \sum_{i=1}^N (cx_i - T_i)^2. \quad (3)$$

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 now to an equation that is similar to Eq. (1):

$$\int cxb(x)dx - \int E(T(x))b(x)dx = 0. \quad (4)$$

Deriving the constant  $c$  under this setting is similar to the derivation in Eq. (2) and is derived from Eq. (4) leading to  $cb_1 = E(T)$ , and thus  $c = \frac{E(T)}{b_1}$ .

Therefore, similarly to Eq. (3), we define the slowdown unfairness measure as the second moment (variability) of  $T(x)$  around  $cx$ :

$$SQF = \int E[(T(x) - cx)^2]b(x) dx,$$

which, after substitution of  $c$  and term expansion, yields

$$\begin{aligned} SQF &= \int \left( E[T(x)^2] + \frac{x^2 E^2[T]}{b_1^2} - \frac{2xE[T(x)]E[T]}{b_1} \right) b(x) dx \\ &= E[T^2] + \frac{b_2}{b_1^2} E^2[T] - \frac{2E[T]}{b_1} \int xE[T(x)]b(x) dx. \end{aligned} \quad (5)$$

Note that SQF is based in its first two terms on the first and second moments of the sojourn time, which have been investigated thoroughly in the literature. However, it consists also on the term  $\int xE[T(x)]b(x) dx$ , which is perhaps new to the literature.

**Remark 3.1.** *Observe that slowdown unfairness can be expressed alternatively using the waiting time measure, which yields exactly the same structure:*

$$\begin{aligned} SQF &= \int E[(T(x) - x\frac{E[T]}{b_1})^2]b(x) dx = \int E[(x + W(x) - x(1 + \frac{E[W]}{b_1}))^2]b(x) dx \\ &= \int E[(W(x) - x\frac{E[W]}{b_1})^2]b(x) dx. \end{aligned} \quad (6)$$

Of course, this alternative expression is meaningful only in cases where service is not interruptible. Also, note that the constant  $c$  differs in this case from the constant of SQF.

### 3.1 An Alternative Fairness Approach

An alternative approach for defining a slowdown fairness is to measure discrimination by  $c_1 - T(x)/x$  leading to deriving the constant via  $c_1 = \int b(x) \frac{T(x)}{x} dx$  which then leads to an unfairness measure  $\int b(x) E[(S(x) - c_1)^2] dx$ . This seems, however, to be very hard to compute. The feasibility of such a metrics remains as an open subject for future research. Note that  $c_1$  is not necessarily equal to  $c$ .

### 3.2 Bounds on SQF

Manipulation of the SQF expression yields:

$$\begin{aligned}
 SQF &= \int E [(T(x) - cx)^2] b(x) dx \\
 &= \int E [(T(x) - E[T(x)] + E[T(x)] - cx)^2] b(x) dx \\
 &= \int (Var[T(x)] + (E[T(x)] - cx)^2 + 2E [(T(x) - E[T(x)]) (E[T(x)] - cx)]) b(x) dx \\
 &= \int (Var[T(x)] + (E[T(x)] - cx)^2) b(x) dx, \tag{7}
 \end{aligned}$$

where the term  $E [(T(x) - E[T(x)]) (E[T(x)] - cx)]$  is eliminated in the third line since given the value  $x$ , the term  $E[T(x)] - cx$ .

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

$$\int Var[T(x)] b(x) dx \leq SQF \leq E[T^2] + \frac{b_2}{b_1^2} E^2[T]. \tag{8}$$

The lower bound results from (7) and  $(E[T(x)] - cx)^2 \geq 0$ , and the upper bound from Eq. (5). Note that the lower bound is tight since it can be achieved for symmetric queues.

### 3.3 Fairness Tests

We study the sensitivity of the SQF measure to both seniority differences and service requirement differences using a simple scenario (sample path) based tests. To this end recall the definition of unfairness given in Eq. (3).

Recall from [3] the definitions of the service requirement test and the seniority preference test:<sup>1</sup>

- *Strong service-requirement preference test:* Consider jobs  $J_i$  and  $J_j$  arriving at  $a_i = a_j$  and having their service times obey  $s_i < s_j$ . Let  $\pi$  be a scheduling policy where the service of  $J_i$  is completed before that of  $J_j$  and let  $\pi'$  be identical to  $\pi$  except for exchanging the service schedule of  $J_i$  and  $J_j$ . A fairness measure is said to satisfy the test if the fairness value it associates with  $\pi$  is higher than that it associates with  $\pi'$ .
- *Strong seniority preference test:* Consider jobs  $J_i$  and  $J_j$  arriving at  $a_i < a_j$  and obeying  $s_i = s_j$ . Let  $\pi$  and  $\pi'$  be defined as above. A fairness measure is said to satisfy the test if the fairness value it associates with  $\pi$  is higher than that it associates with  $\pi'$ .

A *Weak Service-requirement Preference Test* can be defined similarly to the strong one, where the requirement that  $a_i = a_j$  is replaced by a requirement that the arrival times of all jobs present are identical. In a similar manner a *Weak Seniority Preference Test* can be defined.

**Theorem 3.1.** *The SQF measure satisfies the strong seniority preference test.*

*Proof.* Assume that jobs  $J_i$  and  $J_j$  obey  $s_i = s_j$  and  $a_i < a_j$ . Let  $\pi$  and  $\pi'$  be defined as in the test definition. In the event that the scheduling is interruptible Let  $f(\pi)$  be the *unfairness value* associated by SQF with  $\pi$ . The constant  $c$  is the same under both  $\pi$  and  $\pi'$  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

$$f(\pi) - f(\pi') = \frac{1}{n} \left( (t_i - s_i c)^2 + (t_j - s_j c)^2 - (t_i + d_j - d_i - s_i c)^2 - (t_j - d_j - d_i - s_j c)^2 \right)$$

---

<sup>1</sup>Additional requirement (omitted to have a concise presentation) is that the evaluated jobs  $J_i$  and  $J_j$ , would reside concurrently in the system

where  $d_i$  is the departure epoch of  $J_i$  according to schedule  $\pi$ . Using the service time equality  $s_i = s_j$  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  $\pi$  is more 'fair' than  $\pi'$ .  $\square$

To address the service time preference test, we note that due to  $s_i \neq s_j$  one can interpret the test in two ways. One way focuses on the order by which  $J_i$  and  $J_j$  are scheduled to *enter service*; the completion times of the interchanged jobs, in this case, may change due to the interchange. The alternative way focuses on *interchanging the service completion times* of  $J_i$  and  $J_j$ . Such an interchange requires allowing interruptible policies; also it is required that both  $\pi$  and  $\pi'$  are feasible.

The question of whether the *service-entrance* version of the test holds for SQF remains open in this paper. Below we show that the strong *service completion* version of the test holds for SQF.

**Theorem 3.2.** *The SQF measure satisfies the strong service time preference test (service completion interchange version).*

*Proof.* Assume that jobs  $J_i$  and  $J_j$  obey  $a_i = a_j$  and  $s_i < s_j$ . Let  $\pi$  and  $\pi'$  be two feasible schedules defined as in the test definition (service completion version), thus the service completion times of  $J_i$  and  $J_j$ ,  $d_i$  and  $d_j$ , are interchanged between  $\pi$  and  $\pi'$ . Let  $f(\pi)$  be the *unfairness value* associated by SQF with  $\pi$ . Observe that the sum of completion times remains unchanged under both policies and thus the constant  $c$  and the discrimination inflicted on all jobs except of  $J_1$  and  $J_2$  are the same under both policies. We thus can express the change in the unfairness value as

$$f(\pi) - f(\pi') = \frac{1}{n} \left( (t_i - s_i c)^2 + (t_j - s_j c)^2 - (t_j - s_i c)^2 - (t_i - s_j c)^2 \right)$$

where  $t_i$  is the sojourn time of  $J_i$  under  $\pi$ . This expression can be simplified to  $\frac{-2c(s_i - s_j)(t_i - t_j)}{n} < 0$ , implying that  $f(\pi) < f(\pi')$  and that  $\pi$  is more 'fair' than  $\pi'$ .  $\square$

### 3.4 Optimal Scheduling and Zero Measure

We start this section by observing that in a system with infinite processing capacity a trivial resource allocation (scheduling) can achieve zero value of SQF, namely optimal fairness:

**Observation 3.1.** *In an  $M/G/\infty$  system the discrimination of every job is zero and the system's SQF value obeys  $\int E(T(x) - 1 \cdot x)^2 b(x) dx = 0$ .*

We next turn to single server systems and examine the feasibility of an optimal scheduler that achieves zero unfairness. For the sake of comparison note that two unfairness measures which were proposed in the past do have optimal scheduling policies. For the measure proposed in [2], which focuses on job seniority (disregarding job size), FCFS always achieves minimal unfairness (which can be pegged to zero). For the measure proposed in [8], PS always achieves zero unfairness (which is optimal).

We use sample path analysis (defined in Section 3.3) to study the system conditions that lead to the desired zero unfairness value. We consider a scheduling policy to be optimal for a certain sample path if the unfairness value associated by SQF for this path (and schedule) is zero.

**Lemma 3.1.** *There exist sample paths for which zero measure cannot be achieved by any non-preemptive policy.*

*Proof.* Let  $s_1$  be the size of the first job served in a busy period consisting of more than one job, and let  $S$  be the total size of all the other jobs in the sample path. The contribution of the discrimination of this job to the total unfairness is  $(s_1 - s_1(s_1 + S + D))/(s_1 + S)^2$ , where  $D$  is the accumulated waiting time of all other jobs. This quantity equals zero only when  $D$  is zero, which is not feasible.  $\square$

Observe that if all busy periods consist of a single job then we have zero unfairness. This implies that under light loads (i.e., with high probability large portion of the busy periods consist of a single job) any scheduling policy will result in a low unfairness value.

If preemption is allowed (like in PS for example) then in certain cases one can come up with optimal scheduling leading to an unfairness measure of zero.

**Lemma 3.2.** *Consider a busy period consisting of  $N$  jobs,  $J_1, \dots, J_N$  with service times  $x_1, \dots, x_N$ , and arrival times  $a_1, \dots, a_N$ , respectively. Assume the busy period starts at  $a_1$  and let  $\hat{x} = \sum_{i=1}^N x_i$ . Let  $\pi$  be a scheduling policy that achieves a zero measure. Let  $J_k$ ,  $1 \leq k \leq N$  be a job for which the minimum  $\min_i \frac{a_1 + \hat{x} - a_i}{x_i}$  over all the  $N$  jobs is achieved. Let*

$\pi$  be a policy that achieves a zero measure of SQF. Then: (i)  $\pi$  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  $J_j$ ,  $j \neq k$ , should be  $\left(\frac{a_1 + \hat{x} - a_k}{x_k}\right) \frac{x_k}{x_j}$ .

*Proof.* The proof of (i) is by a way of contradiction. 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 jobs 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  $T_k/x_k < T_j/x_j$ . Thus, 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  $T_j - cx_j = T_k - cx_k$ .  $\square$

**Corollary 3.1.** *If all Jobs arrive at the same time then a necessary condition for obtaining optimal scheduling would be to schedule a job with the largest service time as the last job to complete service.*

We use a simple example to demonstrate how to obtain a zero measure for a given setting. Suppose we have 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. We assume that no other arrival occurs and a non-idling server (with processing speed 1). Using the previous lemma it is clear that  $J_3$  will complete service at time 20. Setting the completion time of any single job determines the constant  $c$  and thus for this particular setting  $c = \frac{20}{10} = 2$ . We could reach a zero measure if we serve  $x_1$  at time interval  $(0, 3)$  and don't let it leave the system before time 6, serve  $x_2$  at time interval  $(3, 10)$  and keep it in the system till time 14, and serve  $x_3$  during time interval  $(10, 20)$ .

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

**Theorem 3.3.** *For preemptive policies, there exists a sample path for which there exist no zero unfairness schedule.*

*Proof.* Assume a busy period consisting of three jobs with sizes 1, 1, 10 that arrive simultaneously at time zero. Due to lemma 3.2 the largest job must be scheduled to complete last

and thus  $c = 1.2$ . The value of  $c$  implies that each one of the other two jobs should complete service at time 1.2, which is not feasible.  $\square$

Observe that in all the zero unfairness computations above the constant  $c$  relies on full 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 sample paths for which there exists a schedule of zero unfairness. Within this class there exist paths for which a zero measure schedule must be anticipative.*

*Proof.* Consider a non-anticipative scheduler denoted by  $\pi$ , i.e.,  $\pi$  uses only past history to make the 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  $t = 0$ , and examine the schedule generated by  $\pi$  up 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  $t = 2$  provides that  $\pi$  fails to obtain zero measure. This follows since Lemma 3.2 implies that  $J_1$  should complete service at time 2 in the sample path consisting of  $J_1, J_2$  and  $J_3$ . For clarity we remark that a zero measure can be reached if  $J_1$  is served at time interval  $(1, 2)$ ,  $J_2$  served at time intervals  $(0, 1)$  and  $(2, 4)$ , and  $J_3$  served in parallel to  $J_2$  at time interval  $(2, 4)$  (b) Alternatively, if  $J_1$  was scheduled to complete service at time 2 or afterwards,  $\pi$  cannot obtain zero measure for the original sample path (i.e., the one consisting of  $J_1$  and  $J_2$ ). This is so since Lemma 3.2 implies that  $J_1$  should complete service at time 1.5.  $\square$

## 4 Slowdown Fairness Analysis

In the following analysis we consider single server systems with general arrivals where the service times are i.i.d random variables which are independent of the arrival process and having finite mean and finite variance (G/GI/1), and where the server is non-idling. For some of the results we address more specific systems (M/GI/1, M/M/1, M/D/1).

## 4.1 Non Size-Based Non-Preemptive Policies

We consider non-preemptive non-size-based scheduling policies. From Eq. (6) we get that the unfairness of a non-preemptive non-sized-based policy under G/GI/1 is given by

$$\begin{aligned} SQF &= \int_x E \left[ \left( W(x) - x \frac{E[W]}{b_1} \right)^2 \right] b(x) dx = E[W^2] - \frac{2E[W]E[E[W(x)x]]}{b_1} + \frac{b_2 E^2[W]}{b_1^2} = \\ &= E[W^2] + E^2[W] \left( \frac{b_2}{b_1^2} - 2 \right), \end{aligned} \quad (9)$$

where the last equality follows since the waiting time distribution in the queue is the same for any service time  $x$  and thus  $E[E[W(x)x]] = E[xE[W]] = E[W]E[x]$ .

Now, all these service orders have the same expected sojourn time, namely  $E[T] = b_1 + E[W]$ , which in the case of M/GI/1 translates to:

$$E[T] = b_1 + \frac{\lambda b_2}{2(1 - \rho)}. \quad (10)$$

Since all non-preemptive service orders have the same expected waiting time we can conclude that:

**Theorem 4.1.** *The SQF fairness relative ranking of non-preemptive and non-service-time-based policies is determined only by the second moment of the waiting time. The unfairness value is positively linear in the second moment of the waiting time.*

Next we are interested in the relative fairness of various scheduling policies in this family of policies. A policy in this family can be described as a schedule that only depends on the relative seniority of the customers. We will show that under a general queueing framework (G/GI/1) preferential service to more senior customers always increases fairness.

**Lemma 4.1.** *Consider a G/GI/1 system. Let  $\pi$  be a non-preemptive non-service-time-based scheduling policy. Let  $J_i$  and  $J_j$  be two jobs which concurrently reside under  $\pi$  in the system, and for which  $a_i < a_j$ . Let  $X_i$  and  $X_j$  be random variables denoting their service times, respectively. Assume that  $\pi$  schedules  $J_i$  and  $J_j$  consecutively,  $J_i$  ahead of  $J_j$ . Let  $\pi'$  be a*

scheduling policy that is identical to  $\pi$  except that the scheduling of  $J_i$  and  $J_j$  is done on reverse order. Then the unfairness measure obeys  $SQF^\pi < SQF^{\pi'}$ .

*Proof.* Let  $P$  be a finite sample path (realization) that consists of the arrivals and service times,  $a_1, \dots, a_n$  and  $x_1, \dots, x_n$  where  $1 \leq i < j \leq n$ . 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  $w_P^i(\pi)$  denote the waiting time experiences by  $J_i$  under  $P$  and  $\pi$ . Similarly define the waiting under all paths  $(P, P')$  and scheduling  $(\pi, \pi')$  and both jobs  $(J_i, J_j)$ .

Let  $d = w_P^i(\pi)$  and let  $a_j - a_i = \Delta a > 0$ . Then we have  $w_P^j(\pi) = d + x_1 - \Delta a$ ,  $w_{P'}^i(\pi) = d, w_{P'}^j(\pi) = d + x_2 - \Delta a$ . Similarly for  $\pi'$  we have:  $w_P^i(\pi') = d - \Delta a$ ,  $w_P^j(\pi') = d + x_2$ ,  $w_{P'}^i(\pi') = d - \Delta a, w_{P'}^j(\pi') = d + x_1$ . It is easy to see that the sum of these variables is the same for  $\pi$  and  $\pi'$ , thus  $E[W]$  is not affected by the scheduling policy. Thus, from Eq. (9) we have only to examine  $E[W^2]$ . The contribution of these terms to  $E[W^2]$  under  $\pi$  is given by:

$$d^2 + (d + x_1 - \Delta a)^2 + d^2 + (d + x_2 - \Delta a)^2, \quad (11)$$

and their contribution under  $\pi'$  is given by:

$$(d - \Delta a)^2 + (d + x_2)^2 + (d - \Delta a)^2 + (d + x_1)^2, \quad (12)$$

and the difference between these contributions is  $-2\Delta a(x_1 + x_2) < 0$ . Since this holds for every sample path,  $E[W^2]$  under  $\pi$  is smaller than that under  $\pi'$ .  $\square$

**Corollary 4.1.** *The unfairness values under SQF for FCFS, NP-ROS and NP-LCFS in a G/GI/1 system obey  $SQF^{FCFS} < SQF^{NP-ROS} < SQF^{NP-LCFS}$ .*

Note 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.

Next, we consider the special case of M/GI/1 systems under steady state, where we explicitly derive closed form expressions for the SQF value. We then use them to compare

the fairness of three policies, FCFS, non-preemptive LCFS (NP-LCFS), non-preemptive ROS (NP-ROS). The second moment of the waiting time for FCFS ordering is given by  $E[W^2]^{FCFS} = 2E^2[W] + \frac{\lambda b_3}{3(1-\rho)}$ . The second moment for NP-LCFS and NP-ROS orderings is  $E[W^2]^{NP-LCFS} = \frac{1}{1-\rho}E[W^2]^{FCFS}$  and  $E[W^2]^{NP-ROS} = \frac{1}{1-\rho/2}E[W^2]^{FCFS}$ , respectively.

Using the above we can derive the slowdown fairness measure,  $SQF$ , for these policies

$$\begin{aligned} SQF^{FCFS} &= \frac{\lambda b_3}{3(1-\rho)} + \frac{\lambda^2 b_2^2}{4(1-\rho)^2} \frac{b_2}{b_1^2} \\ SQF^{NP-ROS} &= \frac{\lambda b_3}{3(1-\rho)(1-\rho/2)} + \frac{\lambda^2 b_2^2}{4(1-\rho)^2} \left( \frac{b_2}{b_1^2} + \frac{2\rho}{2-\rho} \right) \\ SQF^{NP-LCFS} &= \frac{\lambda b_3}{3(1-\rho)^2} + \frac{\lambda^2 b_2^2}{4(1-\rho)^2} \left( \frac{b_2}{b_1^2} + \frac{2\rho}{1-\rho} \right). \end{aligned} \quad (13)$$

Thus, in agreement with intuition, the fairness ranking is  $FCFS < NP-ROS < NP-LCFS$ .

## 4.2 Preemptive policies

We start with the analysis of policies that can be set up as symmetric queues (see [4]). These policies include processor sharing (PS), preemptive last-come-first-served (P-LCFS), and a variety of policies that lie between  $PS$  and  $P-LCFS$ . An example of such a policy is *finite-processor-sharing* where the processor is shared (not necessary equally) by all jobs when the number of jobs in the system does not exceed  $r$ , and shared by only  $r$  of them and the remaining ones are kept waiting when the number of jobs present exceeds  $r$ . Further details and variance formulas for such hybrid policies can be found [1].

It can be easily verified that all members of the symmetric queue policy class have the same expected sojourn time [14], which is given by  $E[T(x)] = \frac{x}{1-\rho}$  for a job requiring  $x$  units of service. By substituting this value into (7) we have that the SQF value for symmetric queues is

$$SQF^{SQ} = \int (Var[T(x)] + (E[T(x)] - cx)^2) b(x) dx = \int Var[T(x)] b(x) dx, \quad (14)$$

where the last equality follows since the constant  $c$  in a symmetric queue is  $\frac{1}{1-\rho}$ . Thus, we can state that

**Theorem 4.2.** *The SQF fairness relative ranking of preemptive and non-service-time-based policies that can be setup as symmetric queues is determined only by the variance of the waiting time.*

Note the similarity of this result to that of Theorem 4.1 where in both cases the relative ranking of non-service-time-based policies (either preemptive or non-preemptive) follows the variance of the sojourn time.

The variance of the sojourn time under P-LCFS is (taken from [13])  $Var[T(x)]^{P-LCFS} = \frac{\lambda x b_2}{(1-\rho)^3}$ . Using equation (14) we can derive a closed form expression for the SQF value under P-LCFS

$$SQF^{P-LCFS} = \int \frac{\lambda x b_2}{(1-\rho)^3} b(x) dx = \frac{\rho b_2}{(1-\rho)^3}. \quad (15)$$

In contrast to P-LCFS, the sojourn time variance under PS does not have such a simple representation. [1] presents the following formulation for the variance

$$Var[T^2(x)]^{PS} = \frac{\lambda x b_2}{(1-\rho)^3} + R(x) - R(0),$$

where

$$R(x) = \frac{2}{(1-\rho)} \int_x^\infty (y-x)(1-F(y)) dy$$

$$R(0) = \frac{E[W^2]^{FCFS}}{(1-\rho)^2},$$

and  $F(\cdot)$  is the equilibrium CDF of the delay in an M/G/1-FCFS queue. Nonetheless, [1] shows that  $Var[T^2(x)]^{PS} \leq Var[T^2(x)]^{P-LCFS}$  and thus from (14) it immediately follows that

$$SQF^{PS} \leq SQF^{P-LCFS}. \quad (16)$$

The analysis of size-based policies, such as SRPT and SJF, is out of the scope of this paper and is a subject of ongoing work.

### 4.3 Fairness in M/D/1 and M/M/1 models

Next we compare the fairness of the following policies, P-LCFS, PS, FCFS, non-preemptive LCFS (NP-LCFS) and non-preemptive ROS (NP-ROS) under M/D/1 and M/M/1 models. The second moment of the waiting time for FCFS ordering is given by  $E[W^2]^{FCFS} = 2E^2[W] + \frac{\lambda b_3}{3(1-\rho)}$ . The second moment for NP-LCFS and NP-ROS orderings is  $E[W^2]^{NP-LCFS} = \frac{1}{1-\rho}E[W^2]^{FCFS}$  and  $E[W^2]^{NP-ROS} = \frac{1}{1-\rho/2}E[W^2]^{FCFS}$  respectively.

We compare these policies under deterministic service time (namely,  $b_2 = (b_1)^2$  and  $b_3 = (b_1)^3$ ) as well as under exponential service time (namely,  $b_2 = 2(b_1)^2$  and  $b_3 = 6(b_1)^3$ ).

To compute the unfairness under M/M/1/PS we plug the sojourn time variance from [5],  $Var[T(x)]_{M/M/1}^{PS} = \frac{2\rho x}{(1-\rho)^3} - \frac{2\rho}{\mu(1-\rho)^4} (1 - e^{-\mu x(1-\rho)})$ , into (14) and compute the integral.

Table 1 provides the fairness values of the various policies under the two models. Thus

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

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

for an M/D/1 the fairness ranking is  $SQF > SQF^{NP-LCFS} > SQF^{NP-ROS} > SQF^{FCFS}$  whereas for M/M/1 the ranking is  $SQF^{P-LCFS} = SQF^{NP-LCFS} > SQF^{NP-ROS} > SQF^{FCFS}$  (recall that for both systems we have  $SQF^{P-LCFS} \geq SQF^{PS}$ ).

## 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, SJF and SRPT). Policies whose SQF value is not available in closed form from the previous section (SJF and 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 the 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  $b_2 = 99$ ). Figures 1 – 3 show the corresponding unfairness plots.




Figure 1: The SQF unfairness of common scheduling policies in an M/D/1 system

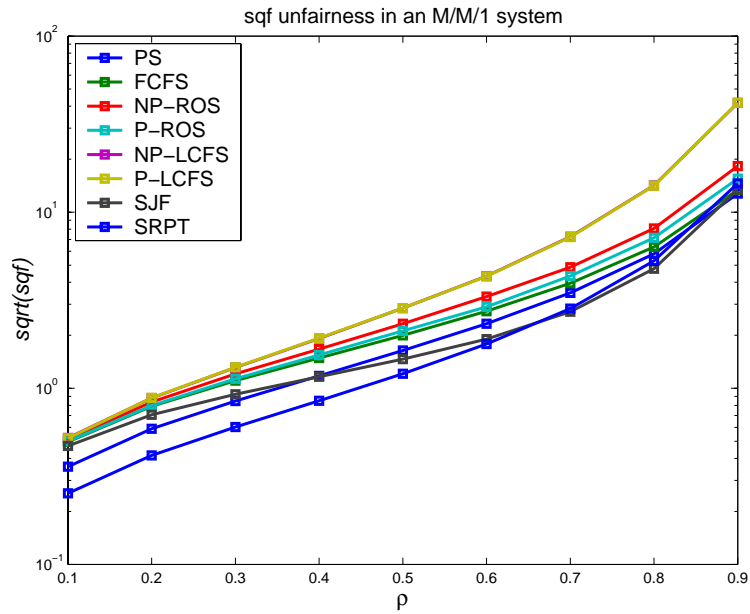


Figure 2: The SQF unfairness of common scheduling policies in an M/M/1 system

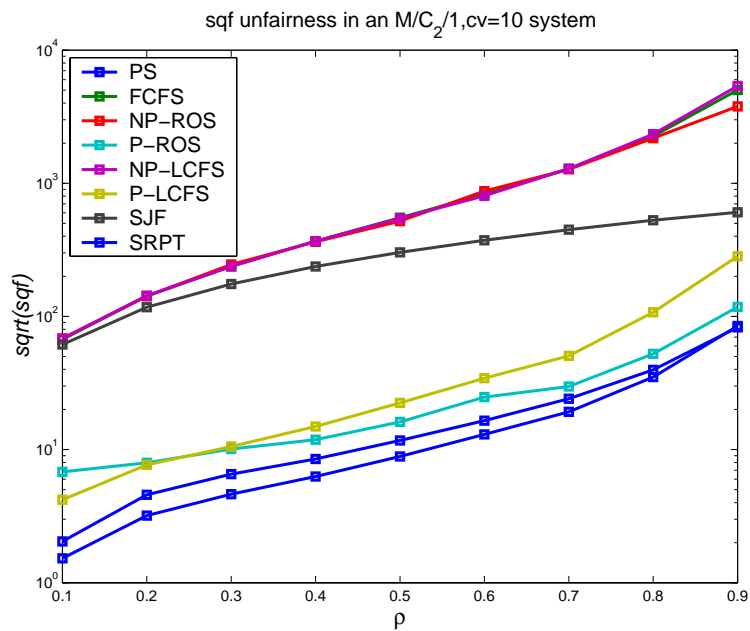


Figure 3: The SQF unfairness of common scheduling policies in an M/GI/1 system with coefficient of variation 10

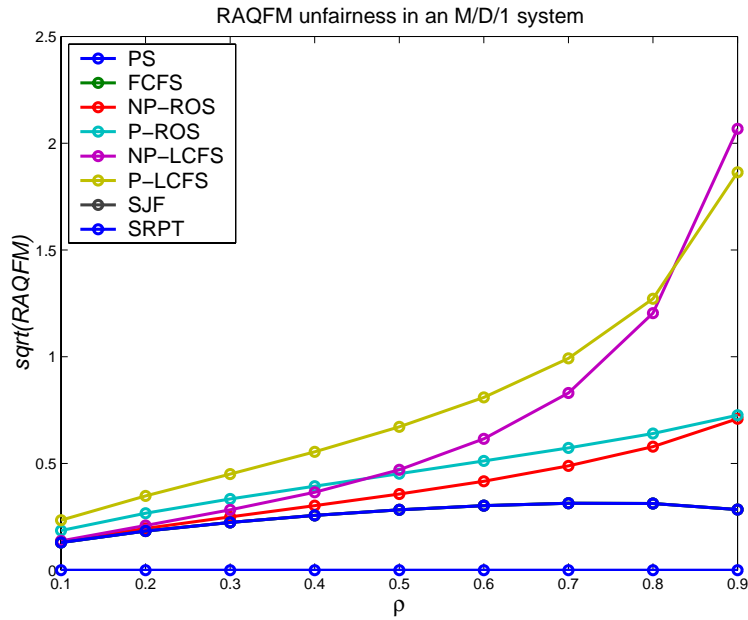


Figure 4: The RAQFM unfairness of common scheduling policies in an M/D/1 system

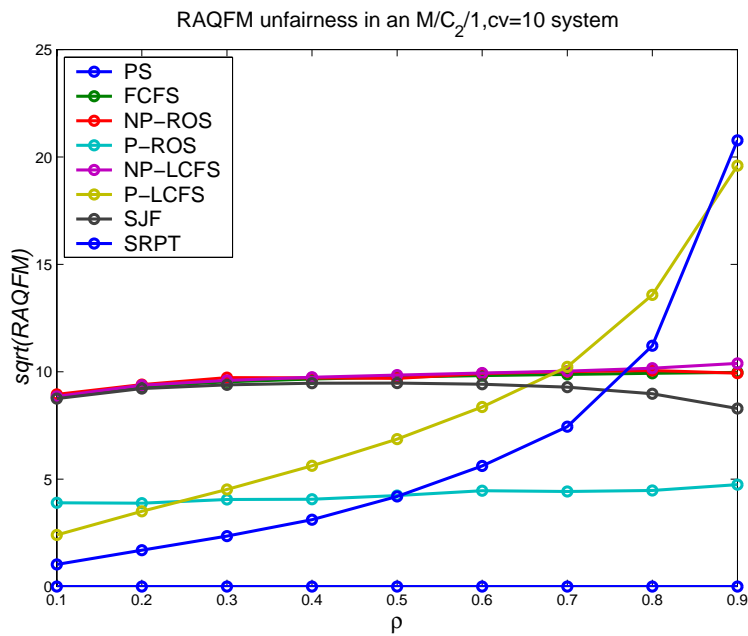


Figure 5: The RAQFM unfairness of common scheduling policies in an M/GI/1 system with coefficient of variation 10

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 Eq. (13) and (15) reveals that

$$\frac{SQF^{FCFS}}{SQF^{P-LCFS}} \approx \frac{\lambda b_2^2(1-\rho)}{4b_1^3}. \quad (17)$$

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. (17) 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  $\lim_{\rho \rightarrow 1} SQF^{FCFS} / DSF^{P-LCFS} = 0$ , 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 [12] 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 [8], whose unfairness values are depicted in Figure 4 and Figure 5, 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 <sup>2</sup>.

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

1. For all policies examined the unfairness measure increases with service time variability.

One can observe it either from the expressions derived in Section 4 or from the figures.

---

<sup>2</sup>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  $\rho \rightarrow 1$  both measures consider P-LCFS to be more unfair

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. 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 [12] 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.

## References

- [1] B. Avi-Itzhak and S. Halfin. Server sharing with a limited number of service positions and symmetric queues. *Journal of Applied Probability*, (24):990–1000, 1987.
- [2] B. Avi-Itzhak and H. Levy. On measuring fairness in queues. *Advances in Applied Probability*, 36(3):919–936, September 2004.
- [3] B. Avi-Itzhak, H. Levy, and D. Raz. Quantifying fairness in queueing systems: Principles, approaches and applicability. Technical Report RRR-25-2005, RUTCOR, Rutgers University, 2005.
- [4] F. P. Kelly. *Reversibility and Stochastic Networks*. Wiley, New York, 1979.
- [5] L. Kleinrock. *Queueing Systems, Volume 2: Computer Applications*. Wiley, 1976.
- [6] R. C. Larson. Perspective on queues: Social justice and the psychology of queueing. *Operations Research*, 35:895–905, Nov-Dec 1987.
- [7] A. Rafaeli, G. Barron, and K. Haber. The effects of queue structure on attitudes. *Journal of Service Research*, 5(2):125–139, 2002.
- [8] D. Raz, H. Levy, and 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*, pages 130–141, New York, NY, June 2004. (*Performance Evaluation Review*, 32(1):130-141).
- [9] M. H. Rothkopf and P. Rech. Perspectives on queues: Combining queues is not always beneficial. *Operations Research*, 35:906–909, 1987.
- [10] W. Sandmann. A discrimination frequency based queueing fairness measure with regard to job seniority and service requirement. Accepted for the 1st Euro NGI Conference on Next Generation Internet Networks Traffic Engineering, April 2005.
- [11] E. Sherry-Gordon. *New Problems in Queues: Social Injustice and Server Production Management*. PhD thesis, MIT, May 1987.

- [12] A. Wierman and 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*, pages 238 – 249, San Diego, CA, June 2003.
- [13] A. Wierman and 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*, pages 229–239, Banff, Alberta, Canada, June 2005.
- [14] S. Yashkov. Properties of invariance of probabilistic models of adaptive scheduling in shared-use systems. *Autom. Control Comput. Sci.*, 12:56–62, 1980.