RUTCOR RESEARCH REPORT

SERVICE TIME VARIABILITY AND JOB SCHEDULING FAIRNESS

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

RRR 12-2005, February 2005

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

[a]Department of Computer Science, Columbia University, New-York, NY, USA, elibrosh@cs.columbia.edu
[b]School of Computer Science, Tel-Aviv University, Tel-Aviv, Israel, hanoch@cs.tau.ac.il
[c]RUTCOR, Rutgers University, New Brunswick, NJ, USA, avi-itzha@rutcor.rutgers.edu

Rutcor Research Report
RRR 12-2005, February 2005

SERVICE TIME VARIABILITY AND JOB SCHEDULING FAIRNESS

Abstract. Fairness is an inherent and fundamental factor of queue service disciplines in a large variety of queueing applications. Service time variability across jobs is an important factor affecting both system performance and scheduling rules (for example, computer systems that prioritize short jobs over long jobs). Service time variability and its effects on mean response times have been studied extensively. However, its effect on queue fairness has not been researched. This work studies the effect of service time variability on queue fairness. We use the RAQFM queue fairness measure, whose analysis for the case of the M/M/1 queue was provided in [25], and study it under a wider variety of service time distributions (rather than exponential only) with a large range of service time variability. This serves two objectives: 1) Extend the understanding of queue fairness, and 2) Examine the capabilities and properties of RAQFM as a fairness measure. For the LCFS-PR scheduling we use a new approach and provide an analysis of the M/G/1 system; this is the first analysis of RAQFM for a non-Markovian system. We show that for this system the fairness depends on the first two moments of the service time and only on them. We also show that under LCFS-PR the expected discrimination of a job, conditioned on the service time, equals zero for every service time. For other service disciplines (FCFS, LCFS-NPR, ROS-NPR, ROS-PR) we approximate service time distributions by Coxian distributions and demonstrate a Markovian-type approach for deriving the RAQFM fairness level of M/Cox/1 systems. The analysis reveals that queue fairness is sensitive to service time variability and that the fairness ranking of common scheduling policies (e.g. FCFS, LCFS, ROS) depends on this parameter. The results demonstrate that the fairness values of RAQFM widely agree with common intuition, and thus provide further confidence in this metric.

1 Introduction

Queueing systems appear in a wide variety of applications such as computer systems, communication systems, web services and call centers, as well as in airports, public offices and many others. Queueing Theory has been used for nearly a century to study the performance of such systems and how to operate them efficiently. Service times and their distributions play an important role in affecting the performance of queueing systems, and the scheduling policies used. One can mention the Pollaczek-Khinchin formula (see queueing theory text books, e.g. [13], [7]) where for the M/G/1 system the average delay is proportional to the second moment of the service time. Accounting for service times in scheduling policies has been widely studied, mainly in the context of optimizing mean system delay or mean delay cost. A well known result in this context is the so-called μc rule, according to which minimization of the M/G/1 waiting cost is achieved by serving the customers with greatest μc value first (service completion rate times cost per minute); see [8] Chapter 3.3 and [14], [16].

Fairness has been recognized as a highly important performance aspect in queues. This recognition can be found in past studies such as [15] [26] [20] [17] and [28]. Recent experimental studies of the reaction of humans to various queue situations ([21] and [22]) have shown that fairness in the queue is very important to humans, perhaps some times even more than the wait itself. In practice, fairness aspects seem to affect scheduling policies, in some cases, not less than the wish to minimize mean waiting time (or weighted mean waiting time). However, fairness considerations have rarely been expressed quantitatively, simply since queue fairness quantification was not available until quite recently.

The interest in computer job scheduling and in their fairness has recently raised interest in *quantitatively evaluating queue scheduling fairness.* Work in this area has been done in [2], [5], [4], [29], [25], [23] [3] and [27][1]. To quantify queue fairness one must first select a measure (yardstick) of queue fairness. To this end, four different approaches have been proposed, recently: 1. In [2] measures based on order of service have been devised. 2. The slowdown (a.k.a. stretch, normalized response time) was proposed as a metric of unfairness: In [5] the max slowdown serves as indication of unfairness, in [4] the *max mean slowdown* is used to evaluate the unfairness of the SRPT scheduling policy and in [29] as a criterion for evaluating whether a system is fair or unfair. 3. In [25] an analysis of the *resources allocated* by the system to the various customers forms the base for a fairness measure named Resource Allocation Queueing Fairness Measure (RAQFM). 4. In [27] an approach based on counting the number of order violation and size violation events is proposed.

As discussed in [3] the first approach focuses on the relative arrival times of customers while the second approach focuses on their relative service times; as such both approaches have difficulties accounting for the tradeoff between relative seniority (the time spent in the system since arrival) and service requirement. The reader may recognize this tradeoff and

[1] The reader may question whether the fairness measures developed in the analysis of Weighted Fair Queueing, like Absolute Fairness Bound and Relative Fairness Bound (e.g., [10], [12], ch. 9 pp. 209-261, [9], [30]) should be considered. Those measures seem to fit well streams of packets and less so individual jobs, on which our focus is in this work.

its importance from his/her daily life experience where a very short job arrives to the queue just shortly after a very long job (e.g in a supermarket), raising the common dilemma of whom it is more fair to serve first. The third approach [25], [23] focuses on the resources of the system and their allocation in a manner that satisfies "social justice" perception, and thus allows to deal with the tradeoff between service requirement and seniority. The results derived in there, and demonstrated mainly by application to M/M/1 type systems, show that the measure is indeed sensitive to both factors and reacts properly (intuitively) in a variety of cases of interest. The fourth approach is very recent and is in a too early stage to be evaluated. We will therefore adopt RAQFM as our fairness evaluation metrics in this study.

For a quantitative measure to be widely accepted it must reflect, quantitatively, the qualitative (or intuitive, or "educated guess") assessment, by most people, of the relative levels of fairness of different queue disciplines, for the widest range of situations. Thus, one of our primary goals in this paper is to widen the demonstrated range where RAQFM does exactly this. Since fairness in a queue is mainly obtained by a desirable preference balance between a customer's seniority and length of service demanded by him, service times variance is intuitively a major factor. The larger the variance, the more prevalent are situations with acute conflict between seniority and service length. At the extreme, where all service times are the same there is no conflict at all, and FCFS is the fairest of all non-preemptive disciplines [2]. As the variance of service times increases we expect the unfairness of all non-preemptive disciplines to increase. On the other hand, LCFS, which is extremely unfair, may intuitively be more fair than FCFS once we allow preemption, i.e. FCFS may be less fair than LCFS-PR (Last-Come-First-Served-Preemptive- Resume) if the variance of service times is large. The intuitive reasoning behind this is that a customer with very long service time is likely to be preempted by one with shorter service time, thus achieving a better balance between seniority and service time. As stated above, the prior work on RAQFM [25], [23] focused on exponential service times (M/M/1 type models) and therefore could not address, directly, the service times variability factor.

This deficiency is overcome in this work by applying our analyses to the M/G/1 system in the LCFS-PR case and to the $M/G_{cox}/1$ system in the non-preemptive cases, where G_{cox} stands for Coxian distribution as approximation to more general distributions. The results derived in this work do add validity to the RAQFM measure by showing that it reflects well the intuitive impacts of the service variability factor on the level of fairness of a system. Thus, this measure may be a good choice for addressing performance questions as well as operational questions, such as: 1) To what degree (quantitatively) service time variability affects job scheduling fairness, 2) How fair are common scheduling disciplines, as a function of the job size variability, and 3) Which scheduling disciplines achieve higher job fairness (as function of the job variability). Since job fairness is one of the major concerns in choosing a scheduling disciplines, answers to these questions should be useful to system designers and operators. Further goals achieved in this work are (i) extension of fairness analyses conducted in prior studies to systems with various levels of service time variability, and (ii) further understanding of the properties of RAQFM, as a metric for fairness evaluation. Note that none of the prior studies dealt with the effect of service time variability on fairness level: i) [2], [27] did not address the problem, ii) [29] provided only a criterion and not a measure, and iii) [25] dealt only with exponential service times (M/M/1). Also, the results derived in this work further close a seemingly perceptual gap between the results derived in [2], concluding that FCFS is more fair than LCFS, to those derived in [29], concluding that LCFS-PR is more fair than FCFS. Our results show that depending on the service time variability either of the views may show up, and, further, that RAQFM properly reacts to these parameter changes and takes either of the views in a proper manner.

We start by presenting the model and reviewing the RAQFM measure (Section 2). We then (Section 3) turn to the analysis of the LCFS-PR and PS disciplines. In (Section 3.1), we provide a complete analysis of the queue unfairness (expressed as the second moment of discrimination) in the Last-Come-First-Served Preemptive-Resume (LCFS-PR) M/G/1 system. This is the first study where RAQFM is evaluated for a non-Markovian system. The method leads to a simple numerical recursion for evaluating the individual discriminations as well as the system's unfairness in this system. The results derived show that system unfairness under LCFS-PR directly depends on the first two moments of the service times and only on them. That is, service variability is a major factor affecting queue fairness.

A striking and seemingly 'paradoxical' result is derived for this system, stating that the expected value of discrimination, conditioned on the service time of a job, *equals zero* for all service times. This, misleadingly, seems to hint that LCFS-PR is very fair. However, the numerical evaluation of the fairness measure for this system (which accounts for temporal unfair situations, resulting from seniority violation by LCFS-PR) demonstrates that it is *very unfair*. The section is concluded in Section 3.2, where we recall that the unfairness of Processor Sharing (PS) is 0 in all single server systems, including the M/G/1 model (thus it is the most fair policy), and regardless of service variability.

We next (Section 4) turn to analyze the FCFS, LCFS non- preemptive (LCFS-NPR), Random-Order-of-Service Non-preemptive (ROS-NPR) and ROS-PR. We realize that the analysis of RAQFM for the M/G/1 model might be quite challenging. The reason for this is that the performance measure of fairness (at least as used in RAQFM) is inherently more involved (mathematically) than the performance measure of waiting times. This is so since the latter involves the measures of individual jobs while the former involves a comparative measuring between different jobs. To overcome this difficulty we turn to the commonly used approach of approximating a general service time distribution by a Coxian distribution, by matching the moments of the distributions, and analyzing the Markovian model with the Coxian service time distribution. In Section 4 we first discuss the approximating procedure and then analyze the corresponding Markovian models. The analysis is carried out via a set of recursive equations, which can be solved numerically to yield the individual job discrimination as well as system unfairness. To provide some insight into the behavior of the non-preemptive policies we provide in Section 5 an approximate analysis of discrimination in these systems, leading to some closed form approximate expressions. That analysis demonstrates that in non-preemptive systems, in the presence of highly variable service times, the positive discrimination experienced by the long jobs is the dominant factor in the system unfairness.

Lastly (Section 6) we turn to conduct a numerical evaluation of the models, examining their fairness sensitivity to service time variability. The major findings are: 1. Service variability significantly affects the fairness of job scheduling policies, including their it relative (fairness) ranking. 2. At high service time variability: At most loads the non-preemptive policies are the most unfair. At high load LCFS-PR is the most unfair. ROS-PR seems to be the most fair almost at all loads. 3. At low service time variability: At most ranges the policies maintain order of fairness: FCFS > ROS-NPR > ROS-PR > LCFS-NPR > LCFS-PR. 3. Preemption is quite effective in achieving job fairness for highly variable service times. Nonetheless, if it is used in conjunction with the LCFS policy (which highly violates seniority) the system can become very unfair.

To summarize, the major contribution of this paper is in several aspects: 1) It significantly widens the set of systems for which the RAQFM measure agrees with common intuition, thus increases the confidence in this measure. 2) It provides, for the first time, an exact analysis of RAQFM for an M/G/1 system (LCFS-PR). 3) It contributes to the understanding of fairness in the context of variable service times. 4) It demonstrates that the analysis of fairness via RAQFM can be effectively extended to many systems with general service time via a mapping to Coxian distributions (further extension is carried out in [24] where RAQFM is applied to multi-queue multi-server systems).

2 Model, Notation and Review of RAQFM in a Single Server System

2.1 Model and Notation

Consider a single server queueing system. The system is subject to a stream of arriving jobs (customers), $C_1, C_2, \ldots$, arriving at this order. Let a_i and e_i denote the arrival and exit (departure) epochs of C_i respectively. Let S_i be a random variable denoting the service requirement (measured in time units) of C_i, where $S_1, S_2, ...$ are i.i.d as S. Let $s^{(1)} = E[S]$, $s^{(2)} = E[S^2]$ $\sigma_S^2 = E[S^2] - (E[S])^2$ and $\gamma_S = \sigma_S / E[S]$, where γ_S is called the coefficient of variation.

At each epoch t the server grants service at rate $x_i(t) \geq 0$ to C_i. Let $N(t)$ denote the number of customers in the system at epoch t. The system is work-conserving, i.e. $\int_{a_i}^{e_i} x_i(t)dt = s_i$. The server has a service rate of one unit and is non-idling, i.e. $\forall t, N(t) > 0 \Rightarrow \sum_i x_i(t) = 1$.

2.2 Individual Customer Discrimination

The fundamental principle underlying RAQFM is the belief that at every epoch t, all customers present in the system deserve an equal share of the system's limited resources (equal piece of the pie). This principle implies that the share of the server's resources a customer deserves at t is simply given by $1/N(t)$. This quantity is called the *momentary warranted service* of C_i at epoch t. Summing this for C_i yields $R_i \stackrel{def}{=} \int_{a_i}^{e_i} dt/N(t)$, the *warranted service* of C_i. The *(overall) discrimination* of C_i, denoted D_i is the difference between the warranted service and the granted service. Since the granted service is $S_i = \int_{a_i}^{e_i} x_i(t)dt$, then

$$D_i = S_i - R_i = S_i - \int_{a_i}^{e_i} dt/N(t). \tag{1}$$

A positive (negative) value of D_i means that a customer receives better (worse) treatment than it fairly deserves, and therefore it is *positively (negatively) discriminated.*

Since D_i consists of the difference between S_i and R_i, we may view S_i as the "positive discrimination" and denote it by $D_i^+ = S_i$, and R_i as the "negative discrimination" and denote it $D_i^- = -R_i$. Define D^+ and D^- to be the steady state limiting values of D_i^+ and D_i^- respectively.

An alternative way to define D_i is to define the *momentary discrimination* of C_i at epoch t as $\delta_i(t) \stackrel{def}{=} x_i(t) - 1/N(t)$, and then the overall discrimination of C_i is $D_i = \int_{a_i}^{e_i} \delta_i(t)dt$. An important property of this measure is that it obeys, for every non-idling work-conserving system, and for every t: $\sum_i \delta_i(t) = 0$, that is, every positive discrimination is balanced by negative discrimination. This results from the fact that when the system is non-empty $\sum_i x_i(t) = 1$ (due to non-idling) and the overall momentarily warranted service at such epoch is 1 as well. An important outcome of this property is that if D is a random variable denoting the discrimination of an arbitrary customer when the system is in steady state, then $E[D] = 0$, namely the expected discrimination is zero. The proof was derived in [23].

2.3 System Measure of Unfairness

To measure the unfairness of a system, using a particular policy, across all customers, that is, to measure the *system unfairness*, one would choose some summary statistics measure over the values D_i, or using the distribution of D, where D is a random variable denoting the discrimination of an arbitrary customer when the system is in steady state. Since $E[D] = 0$, a natural choice is $E[D^2]$ (which equals the variance in this case) and which we denote $F_{D^2} = E[D^2]$. Other optional measures are the mean of absolute deviations $E[|D|]$ and the mean negative discrimination, $-E[D|D < 0]$. Throughout this paper, “unfairness” refers to F_{D^2}.

3 Analysis of M/G/1: The Fairness of LCFS-PR and PS

3.1 Analysis of Fairness in the LCFS-PR System

In this section we analyze the fairness and discriminations experienced in the LCFS Preemptive (Resume) system. This is the first time in which an analysis of RAQFM is provided for non-Markovian service times. Consider a tagged customer C arriving at the LCFS-PR system. Let $k \geq 0$ be the number of customers it finds upon arrival. C enters service immediately, and these k customers will remain in the system until C leaves. Recall that S denotes a random variable representing the service time of C, with moments $s^{(1)}$ and $s^{(2)}$.

While C is served, customers arrive at the system at rate λ. Once the first such customer arrives, it preempts C, starting a sub-busy-period, at the end of which (after all customers arriving depart, including itself) the service of C resumes. Let N be a random variable denoting the number of arrivals during S; this is exactly the number of times C will be preempted and a sub-busy-period will start. Since the arrival process is Poisson, we have:

$$E[N] = \lambda E[S]; \quad E[N^2] = \lambda^2 E[S^2] + \lambda E[S]. \tag{2}$$

Let $D|k$ be a random variable denoting the discrimination experienced by C conditioned on the number of customers (k) it finds in the system upon arrival. Let $D^+|k$ and $D^-|k$ be the conditional positive discrimination (granted service) and negative discrimination (warranted service). Let $D^{SE}|k$ and $D^Q|k$ be the conditional discriminations experienced by C while in service and while in the queue, respectively. Assuming that the value of S is s, we have:

$$D^{SE}|k = (1 - 1/(k+1))s; \qquad D^Q|k = \tilde{D}_1|k + \tilde{D}_2|k + ... + \tilde{D}_N|k, \tag{3}$$

where $\tilde{D}_i|k$ is a random variable denoting the total discrimination experienced by C at the ith sub-busy period. Note that while N depends on S, the variables $\tilde{D}_i|k$ are i.i.d as $\tilde{D}|k$ (which denotes the discrimination experienced by C during an *arbitrary* sub-busy period). The following claim establishes a key relation between the variables $D|k$ and $\tilde{D}|k$:

Proposition 3.1. *For $k = 0, 1, ...$ the random variable $\tilde{D}|k$ is identical to the variable $D^-|k+1$.*

Proof. $\tilde{D}|k$ is the discrimination C experiences from the moment a new customer, say C', arrives (while C is in service) and until the sub-busy period of C' ends. Since during this sub-busy period both C and C' are in the system their negative discrimination during this period is identical. Further, since C' sees exactly $k+1$ customers upon arrival, its discrimination is distributed as $D|k+1$ and its negative discrimination is distributed as $D^-|k+1$. □

Thus, the first and the second moments of $D|k$ are given by:

$$E[D|k, S=s] = E\left[s(1 - \frac{1}{k+1}) + Y(s,k)\right]; \quad E[D^2|k, S=s] = E\left[s(1 - \frac{1}{k+1}) + Y(s,k)\right]^2, \tag{4}$$

where $Y(s,k) = \sum_{i=1}^{N(s)} \tilde{D}_i|k$ and $N(s)$ is the number of Poisson arrivals in time duartion s. Hence

$$E[Y(s,k)] = \lambda s E[\tilde{D}|k]; \quad E[(Y(s,k))^2] = \lambda s \sigma_{\tilde{D}|k} + (\lambda^2 s^2 + \lambda s) E[\tilde{D}|k]^2,$$

where $\sigma_{\tilde{D}|k}$ is the variance of $\tilde{D}|k$. Taking the expectations in Eqs. 4, unconditioning on $S = s$ and using Proposition 3.1 we get

$$E[D|k] = s^{(1)} \left[(1 - \frac{1}{k+1}) + \lambda E[D^-|k+1 \right] \quad (5)$$

$$E[D^2|k] = s^{(2)} \left(1 - \frac{1}{k+1} + \lambda E[D^-|k+1] \right)^2 + \lambda s^{(1)} E[(D^-)^2|k+1]. \quad (6)$$

Now, to derive the first two moments of $D^-|k$ one can repeat the analysis above only for the negative part of the discrimination, leading to equations similar to Eqs. 5 and 6:

$$E[D^-|k] = s^{(1)} \left[\frac{-1}{k+1} + \lambda E[D^-|k+1 \right], \quad (7)$$

$$E[(D^-)^2|k] = s^{(2)} \left(\frac{-1}{k+1} + \lambda E[D^-|k+1] \right)^2 + \lambda s^{(1)} E[(D^-)^2|k+1]. \quad (8)$$

Eq. 7 can be solved by successive substitution to yield:

$$E[D^-|0] = -s^{(1)} \sum_{i=1}^{\infty} \frac{\rho^{i-1}}{i} = \frac{s^{(1)}}{\rho} ln(1-\rho); \quad E[D^-|k] = -\frac{s^{(1)}}{\rho^{k+1}} \sum_{i=0}^{\infty} \frac{\rho^{k+i+1}}{k+i+1} \quad (9)$$

Finally, let g_k be the probability that C sees k customers upon arrival (equalling, due to PASTA, to the corresponding steady state probability). The LCFS-PR system is a symmetric queue as defined by [11, Section 3.3]. Therefore[2] $g_k = (1-\rho)\rho^k, \quad k = 0, 1, 2, ...$ where $\rho = \lambda E[S]$. Thus we have $E[D^2] = \sum_{k=0}^{\infty} (1-\rho)\rho^k E[D^2|k]$ which together with Eq. 6 yields:

$$E[D^2] = s^{(2)} \sum_{k=0}^{\infty} (1-\rho)\rho^k \left(1 - \frac{1}{k+1} + \lambda E[D^-|k+1] \right)^2 + \rho \sum_{k=0}^{\infty} (1-\rho)\rho^k E[(D^-)^2|k+1]. \quad (10)$$

Eq. 10 demonstrates a direct dependency of the second moment of discrimination on the first two moments of service time. Further, we may conclude the following important corollary:

Corollary 3.1. *The unfairness of the M/G/1 system with the LCFS-PR service regime, measured by the RAQFM measure (via the second moment of discrimination) depends on the first two moments of the service time S, and does not depend on higher moments of S.*

The following theorem establishes an important property of LCFS-PR by which all job sizes are treated equally:

[2] Note that this form results also directly from the following simple argument: In the LCFS-PR a customer leaves behind him k customers iff he encounters k customers upon arrival. A customer leaves behind $k+1$ customers iff he preempts a customer who encountered k upon arrival. Since $E[N] = \lambda E[S] = \rho$ we have $g_{k+1} = \rho g_k$ implying $g_k = (1-\rho)\rho^k$, where g_k is the probability of encountering k upon arrival.

Theorem 3.1. *Under LCFS-PR the following properties hold: (1) The expected discrimination of a job, conditioned on its service time, $E[D|S = s]$, equals cs for some constant c. (2) The expected discrimination of a job, conditioned on its service time, $E[D|S = s]$, is equal to zero.*

Proof. Taking expectation on Eq. 4 over k (after substitution of $Y(s, k)$) leads to the proof of (1). Property (1), together with the fact that the expected discrimination (unconditioned) obeys $E[D] = 0$ for any policy (see Section 2.2), lead to Property (2). □

Remark 3.1. *Claim 1 in Theorem 3.1 resembles the well known property of LCFS-PR (see text books, e.g. [14]) whereby the expected value of the conditional sojourn time (conditioned on the service time s) is proportional to s. This property was used in [29] (see Section 1) to conclude that LCFS-PR is always fair (via the slow-down fairness metrics).*

Remark 3.2 (Fairness of LCFS-PR)**.** *One may incorrectly interpret Theorem 3.1 as suggesting that LCFS-PR is very fair. In fact, LCFS-PR is extremely unfair, as demonstrated in Section 6, due to its violation of seniority (and which is captured by the $E[D^2]$ metric).*

The analysis can be generalized to yield the Laplace Stieltjes Transform (LST) and thus higher moments (and distribution) of customer discrimination. Let $B^*(\omega) = E[e^{-\omega S}]$. Let $D^*(\omega) = E[e^{-\omega D}]$ and $D^{-*}(\omega) = E[e^{-\omega D^-}]$ and let $D^*(\omega|k)$ and $D^{-*}(\omega|k)$ be these transforms conditioned on the number of customers seen on arrival (k). Then we have:

$$D^*(\omega) = \sum_{k=0}^{\infty}(1-\rho)\rho^k D^*(\omega|k); \qquad D^*(\omega|k) = B^*((1-1/k)\omega + \lambda - \lambda D^{-*}(\omega|k+1)), \tag{11}$$

$$D^{-*}(\omega|k) = B^*(-\omega/k + \lambda - \lambda D^{-*}(\omega|k+1)). \tag{12}$$

3.2 The Fairness of the PS System

The Processor Sharing discipline is utmost fair under assumption of general arrival and service times. This follows directly from the equitable resource allocation principle. Formal discussion is provided in [25].

4 M/G/1: Analysis of Various Scheduling Disciplines via Coxian Approximation

In this section we analyze the FCFS, LCFS-NPR, ROS-PR and ROS-NPR policies. The analysis approach used is to take the distribution of the service time, and approximate it by a Coxian distribution. This leads to a Markovian model for which the discrimination and fairness are then derived. For lack of space we provide only the analysis of FCFS and ROS-PR. The analysis of the other models can be found in the appendix.

4.1 Approximating General Service Time distributions by Coxian Distributions

For the purpose of approximating a general service time distribution we use a second order [3] moment matching technique [1, Chapter 2.5]. In particular, we fit a phase-type distribution, either Coxian or Erlangian, to the mean, $s^{(1)}$, and the coefficient of variation, γ_S, of the given service time random variable S. We distinguish between two cases: (a) When $0 < \gamma_S < 1$ we seek an integer k such that $\frac{1}{k} \approx {\gamma_S}^2$, and fit a k-stage Erlang distribution, E_k, (see Section 4.2) with mean $s^{(1)}$. To match an arbitrary $0 < \gamma_S < 1$ it is possible to use a more sophisticated distribution such as the mixed Erlang distribution which selects between E_k and E_{k-1} with some fixed probability. However, using the basic Erlang distribution leads to simpler recurrence equations and therefore we use it for the analysis. (b) When $\gamma_S > 1$ we use a Coxian-2 distribution [1, Chapter 2.4] which is composed of two exponential stages with mean lengths μ_i, $i = 1, 2$ where the move from the first stage to the second one is with probability p_1, and with probability $1 - p_1$ the service ends after the first stage. For the approximation we use the following parameters, suggested by [18]: $\mu_1 = 2s^{(1)}$, $\alpha = 0.5/\gamma_S$, $\mu_2 = \mu_1\alpha$.

4.2 Conditional Discrimination in M/E$_r$/1

Consider the M/E$_r$/1 where the service time distribution is Erlang with r exponential stages. For this distribution the service is assumed to consists of a sequence of r i.i.d phases (stages) $1, \ldots, r$, each exponentially distributed with parameter $r\mu$.

In a work conserving non-idling M/E$_r$/1 system the time between the arrival of a customer and its departure can be viewed as 'slotted' by arrivals and stage completions (where a 'slot' corresponds to the duration at which the system remains in a particular state). Let $T_i, i = 1, 2, \ldots$ be the duration of the i-th slot, then $T_i, i = 1, 2, \ldots$ are i.i.d. random variables exponentially distributed with parameter $\lambda + r\mu$ and first two moments $t^{(1)} = \frac{1}{\lambda + r\mu}$ and $t^{(2)} = \frac{2}{(\lambda+r\mu)^2} = 2(t^{(1)})^2$. The probabilities that a slot ends with an arrival or a stage completion are denoted by $\tilde{\lambda}$ and $\tilde{\mu}$ respectively (where $\rho = \lambda/\mu < 1$):

$$\tilde{\lambda} = \lambda/(\lambda + r\mu) = \rho/(r + \rho); \qquad \tilde{\mu} = r\mu/(\lambda + r\mu). \tag{13}$$

The system unfairness, given by $E[D^2]$, can be expressed as:

$$E[D^2] = P_0 E[D^2|0, 1] + \sum_{k=1}^{\infty} \sum_{j=1}^{r} P_{k,j} E[D^2|k, j], \tag{14}$$

where $P_{k,j}$, $k \geq 1$ is the probability of finding k customers upon arrival and the served customer in stage j, $P_0 = 1 - \rho$ is the probability of finding an empty system, and $E[D^2|k, j]$ is the second moment of D for a customer who arrives to find k customers in the system and the served one in stage j. Note that $P_{k,j}$ can be derived or computed via standard techniques for solving steady state balance equations (see, for example, [13]).

[3] A third order matching technique was proposed in [19] and can alternatively be taken.

4.2.1 FCFS

For a tagged customer C residing in the system let a denote the number of customers ahead of C, let b denote the number of customers behind C, and let $j = 1, \ldots, r$ denote the stage of service in which the served customer is currently found. C is said to be in state $\mathcal{S}_{a,b,j}$. Due to the memoryless properties of the system, the state $\mathcal{S}_{a,b,j}$ captures all that is needed for predicting the future of C. The momentary discrimination at state $\mathcal{S}_{a,b,j}$ is independent of the current service stage. We denote it by $c(a, b)$, where j is omitted (and where $\Delta()$ is the indicator function):

$$c(a, b) = \Delta(a = 0) - \frac{1}{a + b + \Delta(a > 0)}. \tag{15}$$

Let $D(a, b, j)$ denote the accumulated discrimination of C during a walk starting at state $\mathcal{S}_{a,b,j}$ and ending at the departure of C, and let $d(a, b, j) \stackrel{def}{=} d^{(1)}(a, b, j)$ and $d^{(2)}(a, b, j)$ be the first and second moments of $D(a, b, j)$. Then

$$E[D|k, j] = E[d(k, 0, j)]; \qquad E[D^2|k, j] = E[d^{(2)}(k, 0, j)]. \tag{16}$$

Assume C is in state $\mathcal{S}_{a,b,j}$. C will encounter one of the two following evens: (1) A new customer arrives into the system. The probability of this event is $\tilde{\lambda}$. Afterwards, C will move to state $\mathcal{S}_{a,b+1,j}$. (2) A customer completes its current stage. The probability of this event is $\tilde{\mu}$. If C is not being served ($a \neq 0$) it will move to $\mathcal{S}_{a,b,j+1}$ if $j \neq r$ or to $\mathcal{S}_{a-1,b,j}$ if $j = r$. If C is being served ($a = 0$) it will move to $\mathcal{S}_{0,b,j+1}$ if $j \neq r$ or will leave the system if $j = r$. Thus, we have:

$$D(a, b, j) = \begin{cases} Tc(a, b) + D(a, b + 1, j) & \text{w.p } \tilde{\lambda} \\ Tc(a, b) + D(a, b, j + 1) & \text{w.p } \tilde{\mu},\ j \neq r \\ Tc(a, b) + \Delta(a > 0)D(a - 1, b, j) & \text{w.p } \tilde{\mu},\ j = r, \end{cases} \tag{17}$$

where T is the duration of the current slot.Taking expectation leads to the recursive expression:

$$d(a, b, j) = \begin{cases} t^{(1)}c(a, b) + \tilde{\lambda}d(a, b + 1, j) + \tilde{\mu}d(a, b, j + 1) & j \neq r \\ t^{(1)}c(a, b) + \tilde{\lambda}d(a, b + 1, j) + \Delta(a > 0)\tilde{\mu}d(a - 1, b, 1) & j = r. \end{cases} \tag{18}$$

Squaring Eq. 17, expanding the resulting quadratic terms and taking expextation on them yields:

$$\begin{aligned} d^{(2)}(a, b, j) = {} & t^{(2)}(c(a, b))^2 + \tilde{\lambda}d^{(2)}(a, b + 1, j) + 2t^{(1)}c(a, b)\tilde{\lambda}d(a, b + 1, j) + \\ & \begin{cases} \tilde{\mu}d^{(2)}(a, b, j + 1) + 2t^{(1)}c(a, b)\tilde{\mu}d(a, b, j + 1)) & j \neq r \\ \Delta(a > 0)(\tilde{\mu}d^{(2)}(a - 1, b, j) + 2t^{(1)}c(a, b)\tilde{\mu}d(a - 1, b, j)) & j = r. \end{cases} \end{aligned} \tag{19}$$

These recursive relations (equations (19), (18)) combined with equations (16), (15), (14), (13) can be used, via numerical computation, to derive the system unfairness measure, $F_{D^2} = E[D^2]$.

4.2.2 Preemptive ROS

Consider a preemptive ROS policy in which preemption occurs at all arrival instants, and any selection for service (either at preemption or at service completion) is done at random among all customers in the system (including the preempted one). Let $\overline{a} =< a_1, \ldots, a_r >$ be a vector of length r, where a_i is the number of customers other than C that need to complete $r-i+1$ stages of service. Let $a = \sum_i a_i$ denote the total number of customers in the system other than C. Let $c = 1, \ldots, r$ be an integer variable such that $r-c+1$ is the number of stages that C needs to complete. Let s be a boolean variable equalling 1 if C is in service and 0 if it is waiting. The state of C is denoted by $\mathcal{S}\overline{a}, c, s, j$. In this state a_i,$i = 1, \ldots, r$ customers need to complete $r-i+1$ stages, the tagged customer C needs to complete $r-c+1$ stages, the one in service is in its j-th stage, and it is C if $s = 1$ and not C if $s = 0$. To simplify the presentation of the recursive equations we define $\overline{1_j} =< 0, \ldots, 0, 1, 0, \ldots, 0 >$ to be a vector that its j-th element is 1 and the rest are zero.

When C is in state $\mathcal{S}_{\overline{a},c,s,j}$ it will encounter one of the following possible events:

1. If $s = 0$ the possible events are:

   (a) A customer arrives into the system and C is chosen to receive service next. The probability of this event is $\frac{\tilde{\lambda}}{a+2}$ and C will move to $\mathcal{S}_{\overline{a}+\overline{1_1},c,1,c}$.

   (b) A customer arrives into the system and a waiting customer (other than C) which is left with $r-k+1$ stages is chosen to receive service next. The probability of this event is $\tilde{\lambda}\frac{\tilde{a_k}}{a+2}$ where $\tilde{a_k} = a_k$ for $k = 2, \ldots, r$ and $\tilde{a_k} = a_1 + 1$ for $k = 1$. Then C will move to $\mathcal{S}_{\overline{a}+\overline{1_1},c,0,k}$.

   (c) A customer completes its current stage j, where $j \neq r$. The probability of this event is $\tilde{\mu}$ and C moves to $S_{\overline{a}+\overline{1_{j+1}}-\overline{1_{j-1}},c,0,j+1}$.

   (d) A customer completes service, leaves the system and C is chosen to receive service next. The probability of this event is $\tilde{\mu}/a$ and C will move to $\mathcal{S}_{\overline{a}-\overline{1_r},c,1,c}$.

   (e) A customer completes service, leaves the system and a waiting customer left with $r-k+1$ stages is chosen to receive service next. The probability of this event is $\tilde{\mu}\frac{a_k}{a}$ and C will move to $\mathcal{S}_{\overline{a}-\overline{1_r},c,0,k}$.

2. If $s = 1$ the possible events are:

   (a) Same as (1a).

   (b) Same as (1b) but C moves to $\mathcal{S}_{\overline{a}+\overline{1_1},j,0,k}$.

   (c) Same as (1c) but C moves to $S_{\overline{a},c+1,1,j+1}$.

   (d) The customer in service completes its service. The probability of this event is $\tilde{\mu}$ and C leaves the system.

For $s = 0$, $d(\overline{a}, c, s, j)$ can be expressed as

$$d(\overline{a}, c, 0, j) = t^{(1)}c(\overline{a}, s) + \frac{\tilde{\lambda}}{a+2}d(\overline{a} + \overline{1_1}, c, 1, c) + \sum_{i:a_i>0} \tilde{\lambda}\frac{\tilde{a}_i}{a+2}d(\overline{a} + \overline{1_1}, c, 0, i) + \begin{cases} \tilde{\mu}d(\overline{a} + \overline{1_{j+1}} - \overline{1_{j-1}}, c, 0, j+1) & j \neq r \\ \frac{\tilde{\mu}}{a}d(\overline{a} - \overline{1_r}, c, 1, c) + \sum_{i:a_i>0} \tilde{\mu}\frac{a_k}{a}d(\overline{a} - \overline{1_r}, c, 0, i) & j = r \end{cases} \quad (20)$$

For $s = 1$ it is given by

$$d(\overline{a}, c, 1, j) = t^{(1)}c(\overline{a}, s) + \frac{\tilde{\lambda}}{a+2}d(\overline{a} + \overline{1_1}, c, 1, c) + \sum_{i:a_i>0} \tilde{\lambda}\frac{\tilde{a}_i}{a+2}d(\overline{a} + \overline{1_1}, j, 0, i) + \begin{cases} \tilde{\mu}d(\overline{a}, c+1, 1, j+1) & j \neq r \\ 0 & j = r \end{cases} \quad (21)$$

The recursive equations for $d^{(2)}(\overline{a}, c, s, j)$ can be derived in a similar manner. They are given in the appendix (Section 8.1.3) .

A customer arrives to the system either at state $\mathcal{S}_{\overline{0},1,1,1}$ when it is empty, where $\overline{0}$ is a zero vector of length r, or at state $\mathcal{S}_{\overline{a},1,0,j}$ when it is serving a customer at stage j and the number of service stages remaining to the queued up customers is represented by $\overline{a}$. Then, for preemptive ROS

$$E[D^2] = P_0 d^{(2)}(\overline{0}, 1, 1, 1) + \sum_{\overline{a}} \sum_{j=1}^{r} P_{\overline{a},j} d^{(2)}(\overline{a}, 1, 0, j], \quad (22)$$

where $P_{\overline{a},j}$, is the probability of finding a customers upon arrival such that the number of customers left with $r - i + 1$ stages of service is a_i, and the served customer is in stage j.

4.2.3 Non-Preemptive LCFS

Here we give a short description of the state variable, $\mathcal{S}_{a,b,j}$, used to construct the recursive discrimination equations for this model. The full analysis is presented in the appendix (Section 8.1.1). Our approach is to preserve the notations of the FCFS model. At every slot let a denote the number of customers arrived earlier than C and thus to be served after C, and let b denote the number of customers arrived later than C and thus to be served before C. The state $S_{a,b,j}$ (where j is the stage of the customer in service) captures all that is needed for predicting the future of C.

4.2.4 Non Preemptive ROS

Similarly to the previous section we provide only a brief description of the state variable. Further details can be found in the appendix (Section 8.1.2). For a tagged customer C, denote by a the number of customers in the system other than C, and by s a boolean variable equaling 1 if C is in service and 0 if it is waiting. Then, in state $\mathcal{S}a, s, j$ there are a customers in addition to C, the one in service is in its j-th stage, and it is C if $s = 1$ and not C if $s = 0$.

4.3 Conditional Discrimination in M/Cox$_2$/1

For the analysis of M/Cox$_2$/1 model we preserve the notations of the the M/E$_r$/1 model and use the same state variables; this leads to simpler recursive equations than the M/E$_r$/1 equations (since there are only two stages). The full description of the M/Cox$_2$/1 model and its analysis is given in the appendix (Section 8.2).

4.4 Computation Complexity Perspectives

Observe that the complexity of the discrimination computation (i.e., the recursive equations complexity) is dependent upon the number of stages used to approximate the service time distribution. This implies that the solution of M/Cox$_2$/1 equations is of relatively lower concern since the distribution is represented by only 2 stages. Higher concern might be attributed to the M/E$_r$/1 model (r states), with high values of r. Nonetheless, as demonstrated in the numerical results section (Section 6) below, the discrimination and fairness values for this model converge as a function of r and thus, models involving large values of r are not needed.

Thus, in practice, for both models the computational complexity is based only on a small number of stages, and thus reasonable number of states and variables in the numerical computation.

5 Properties of Discrimination under Long Service Times

In this section we are interested in understanding the behavior of the discrimination function in the presence of high variability service times. We do this by focusing on non-preemptive systems and studying the discrimination during the service of customers with long service times. To this end, it will be convenient to break the discrimination of C_i into several components. Let D_i^{SE} and D_i^{Q} be the discriminations experienced by C_i while in service and while (waiting) in queue, respectively. We may further break D_i^{SE} into the positive discrimination observed in service, denoted and obeying $D_i^{SE+} = S_i$, and the corresponding negative discrimination D_i^{SE-}, obeying $D_i^{SE-} + D_i^{Q} = -R_i$ (which is the overall warranted service). Recall also the notations $D_i^{+} = S_i$ and $D_i^{-} = -R_i$ (see Section 2.2). For the corresponding steady state variables we use the same notation where the index i is omitted.

5.1 Expected Positive and Negative Discriminations

¿From $E[D] = 0$, it follows that:

Observation 5.1. *Under RAQFM, for any single server system and any work conserving policy, the expected values of the positive discrimination and of the negative discrimination are equal to each other:* $E[D^+] = -E[D^-]$.

5.2 Non Preemptive systems: The effect of a Long Service customer on the Discrimination of Other Customers

Consider a tagged customer, C, who resides in the system, and who encounters, during her waiting time, the service of a very long job (denote the customer with the long job by C_L).

To achieve some insight, we use a simplistic model and assume that customers are of two types whose service time are exponentially distributed with means $1/\mu_1$ and $1/\mu_2$ where the second type corresponds to the very large jobs and thus $\mu_2 << \mu_1$ and the mean service time of C_L is $1/\mu_2$. The arrival rate (Poisson) into the system is assumed to be λ.

Note that λ/μ_2 is not necessarily smaller than 1 (for stability). In fact, we are interested in cases where $\lambda/\mu_2 >> 1$. We also assume that service is non-preemptive, thus, once the service of the long job started it will be carried out to completion.

Let K be the number of customers present at the system when the service of C_L starts, or when C arrives, whichever is later; obviously $K \geq 2$ since both C_L and C reside in the system. Let $\bar{t}$ be the expected duration until the next event (service completion of C_L or arrival of a new customer), then $\bar{t} = 1/(\lambda + \mu_2)$. Let $D^-_{(k)}$ be the negative discrimination experienced by C during the service of C_L given that $K = k$ and let $d^{(1)-}_{(k)} = E[D^-_{(k)}]$ and $d^{(2)-}_{(k)} = E[(D^-_{(k)})^2]$.

Proposition 5.1. $d^{(1)-}_{(k)}$ *,* $k = 2, , 3, ...$ *is monotonically decreasing, i.e.* $d^{(1)-}_{(k+1)} < d^{(1)-}_{(k)}$.

Proof. The proof is carried out by examining two systems, one that starts with k customers and one that starts with $k + 1$ customers. If the systems are subject to exactly the same arrival and departure processes (until the departure of C), then at every epoch of arrival or departure the first system will have one less customer. Since the temporal negative discrimination at t is given by $1/N(t)$ it follows that the discrimination in the first system is larger than that in the second system for *every sample path.* This directly implies the monotonicity of the expected values as stated in the proposition. □

To bound the value of $d^{(1)-}_{(k)}$, it is now sufficient to bound the value of $d^{(1)-}_{(2)}$, which we do next. We look at the events occurring while C_L is served; these can be either an arrival or the service completion of C_L. Since both occur at exponential rates the expected duration until the next event is given by $\bar{t} = 1/(\lambda + \mu_2)$, the probability that the next event is an arrival is given by $p = \lambda/(\lambda + \mu_2)$ and the probability that the next event is a service completion is given by $1 - p$. Also, the negative momentary discrimination at the first interval (time until first event) is -1/2, at the second interval is given by -1/3 and so on. Thus, the over all expected negative discrimination accumulated by C during the service of C_L is given by:

$$d_{(2)}^{(1)-} = -\bar{t}\sum_{i=0}^{\infty} \frac{1}{2+i} p^i. \tag{23}$$

yielding:

$$d_{(2)}^{(1)-} = -\frac{\bar{t}}{p^2}[-p - ln(1-p)]. \tag{24}$$

For relatively large values of $1/\mu_2$ where p is close to 1, we have:

$$d_{(2)}^{(1)-} \approx -\frac{1}{\lambda+\mu_2} ln(\frac{1}{1-p}) \approx -\frac{1}{\lambda} ln(\frac{1}{1-p}) \approx \frac{-1}{\lambda} ln(\frac{\lambda}{\mu_2}). \tag{25}$$

5.3 Non Preemptive systems: Effect of a Long Service on the served Customer Discrimination

We now repeat the analysis performed in the previous section, but now we focus on the discrimination experienced by C_L while being served. Let K be the number of customers present at the system when the service of C_L starts. Assume that $K \geq 2$ (that is C_L is not alone). Let $D_{(k)}$ be the negative discrimination experienced by C_L during its service given that $K = k$ and let $d_{(k)}^{(1)} = E[D_{(k)}]$. Similarly to the previous section one can show that $d_{(k)}^{(1)}$ is monotonically non-decreasing in k. Also, similarly to Eq. 23 we get:

$$d_{(2)}^{(1)} = \bar{t}\sum_{i=0}^{\infty}\left(1 - \frac{1}{2+i}\right) p^i = \bar{t}\left(\frac{1}{1-p} - \sum_{i=0}^{\infty}\frac{1}{2+i}p^i\right). \tag{26}$$

Following the analysis of the previous section we get, when p approaches 1:

$$d_{(2)}^{(1)} \approx \frac{1}{\lambda+\mu_2}\left(\frac{1}{1-p} - ln(\frac{1}{1-p})\right) \approx \frac{1}{\lambda}\left(\frac{\lambda}{\mu_2} - ln(\frac{\lambda}{\mu_2})\right). \tag{27}$$

Observation 5.2. *Eqs. 27 and 25 reveal that while the discrimination of the served customer (C_L) is proportional to $\frac{1}{\mu_2}$ (mean service time), the negative discrimination of a waiting customer (C) is proportional only to its logarithm. Thus, the former is a much more significant factor in the overall system unfairness. This will explain some of the results reported in Section 6.*

6 Numerical results and Observations

6.1 Numerical Results

In this section we numerically evaluate the systems studied, aiming at examining their unfairness as a function of the system load, service time variability and scheduling policy. We

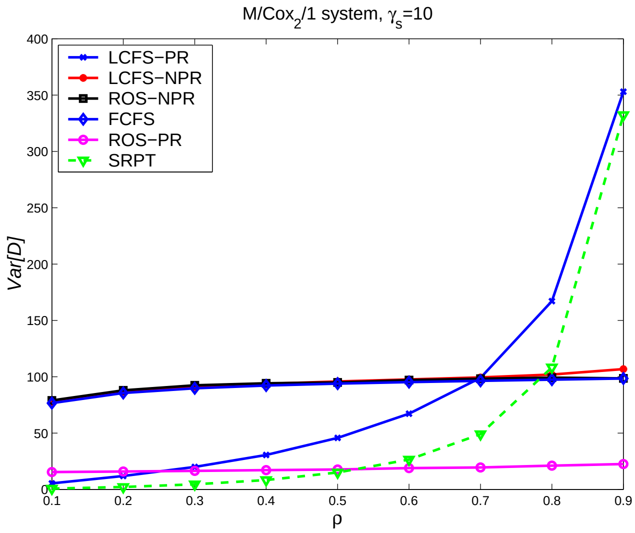

Figure 1: Unfairness for high variability system, M/COX_2/1 (Coefficient of variation $\gamma_s = 10$)

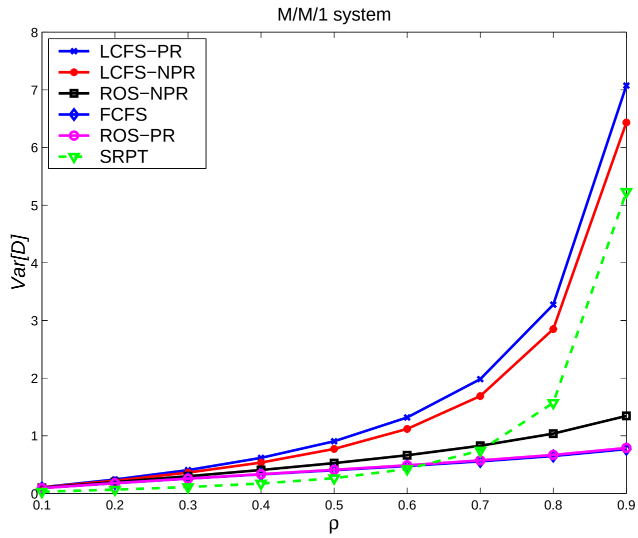

Figure 2: Unfairness for medium variability system, M/E_1/1 (Coefficient of variation $\gamma_s = 1$)

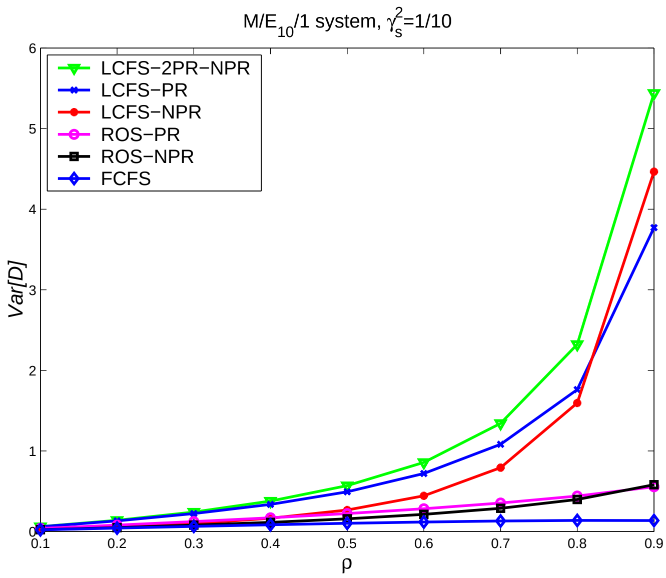

Figure 3: Unfairness for low variability system, M/E_{10}/1 (Coefficient of variation $\gamma_s = 1/\sqrt{10}$)

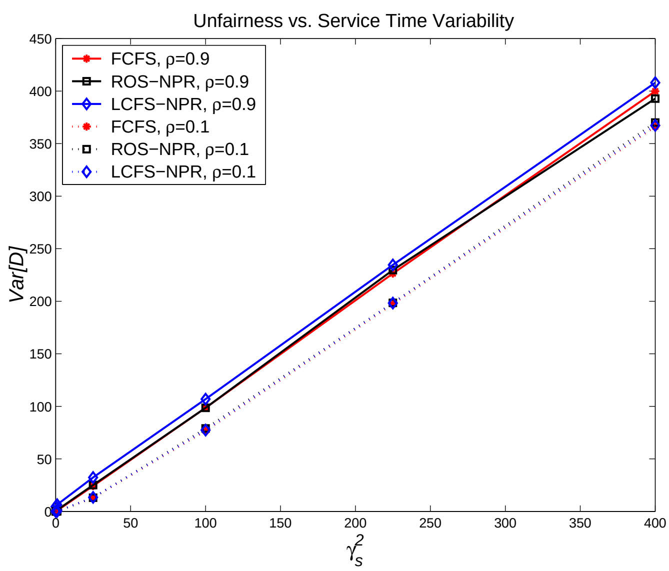

Figure 4: Unfairness of Non-Preemptive Policies as a Function of Service Time Variability

conduct an evaluation for all the scheduling policies studied under a wide range of service time variabilities. For low service time variability we use the M/E_r/1 model and for the high variability the M/COX_2/1 model.

In all cases examined we set the mean service time to $E[S] = 1$ and use the arrival rate λ to control the system load. We also vary the service time coefficient of variation γ_S. The examination is carried out for loads ranging between 0.1 and 0.9 and γ_S varies through the values $1/(\sqrt{2}0), 1/(\sqrt{1}5), 1/(\sqrt{1}0), 1/(\sqrt{5}), 1, 5, 10, 15, 20$. For compactness we do not present here all the results.

Figure 1 depicts the unfairness in the system under high service time variability $\gamma_S = 10$. The behavior for $\gamma_S = 5, 15, 20$ is quite similar and thus is not presented. Figure 2 depicts the unfairness in the system under medium service time variability, $\gamma_S = 1$ (the M/M/1 case). Figure 3 depicts the unfairness in the system under low service time variability, $\gamma_S = 1/(10^{0.5})$ (the M/E_{10}/1 system). Lastly, in Figure 4 we consider the non-preemptive policies (FCFS, ROS and LCFS) and depict their unfairness as a function of the service time variability, for several values of ρ.

6.2 Observations and Properties

1. **Effect of variability:** Service time variability significantly affects the fairness experienced in the various disciplines. In fact, service time variability affects also the relative (fairness) ranking of the scheduling policies. For example, at $\rho = 0.6$, the relative fairness ranking for $\gamma_S = 10$ (Figure 1) is ROS-PR > LCFS-PR > FCFS ≈ ROS-NPR ≈ LCFS-NPR (where > should read as "more fair" and ≈ as "approximately identically fair"). In contrast, for $\gamma_S = 1$ (Figure 2) it is FCFS ≈ ROS-PR > ROS-NPR > LCFS-NPR > LCFS-PR.

2. **High variability service times:** We observe the following properties (demonstrated in Figure 1, and observed in all the high variability cases we examined):

   (a) The unfairness of all non-preemptive policies is about the same. The reason for this is that the dominant discrimination, in this case, is the positive discrimination of the long jobs (see the results derived in Section 5 Observation 5.2), which becomes dominant since unfairness is taken as the second moment of discrimination. Thus, the particular order of service has negligible effect on the overall discrimination, in these policies.

   (b) For low to medium loads the non-preemptive policies are the most unfair while the ROS-Preemptive is the most fair.

   (c) For high loads the LCFS-PR becomes the most unfair (while ROS-PR maintains its highest fairness rank).

3. **Low variability service times:** we observe:

(a) The unfairness values for the M/E_r/1 model are affected by the values of r but seem to converge once r reaches the values between 5 and 10.

(b) At most ranges of load the fairness relative ranking obeys FCFS > ROS > LCFS and $NPR > PR$. Both relations agree with common intuition and with prior results indicating that in the case of deterministic service (zero service time variability) serving jobs in the order they arrive is the most fair order among non-preemptive policies.

(c) At very high load, LCFS-NPR becomes more unfair than LCFS-PR. This *surprising result* is explained in detail in [6], and is due to the following: 1) The highest unfair situations occur when there are 2 customers in the system, and 2) At high load LCFS-PR reduces the number of these situations since it increases the number of customers in the system (compared to LCFS-NPR). This suggests that a hybrid policy, LCFS-2PR-NPR (preemption occurring only when there are 2 customers in the system), is more unfair than either LCFS-NPR or LCFS-PR, which indeed is verified in the figure.

4. **Linearity:** The unfairness of the Non-preemptive policies is roughly proportional (linear) to the square coefficient of variation of the service time (See Figure 4). This occurs at all loads; further, at medium to heavy loads these values are quite insensitive to the load. The explanation is similar to that of 2.a above. Such linearity is also observed (not depicted) for preemptive policies (LCFS-PR, ROS-PR).

6.2.1 Preemption and Ordering

Observe the ROS-PR and the LCFS-PR policies. First, it is striking to see that ROS-PR is very fair at most cases (the most fair in many cases). In contrast, the LCFS-PR is very unfair (most unfair) in most cases. This suggests that the *preemption* factor drives LCFS-PR towards fairness (discriminating against long jobs) while the LCFS factor (discrimination against senior jobs) drives it towards unfairness. The latter factor does not exist in ROS-PR, which is the reason it is very fair. This observation should be put in the context of previous research, e.g., [29] that found LCFS-PR to meet their criterion for being always fair (using the slow-down approach).

Thus, we may conclude that preemption serves as an instrumental tool for increasing fairness when job size variability is large. Such mechanism is however counteracted if it is accompanied by an order violating scheme (LCFS). To this end, we conjecture that a policy that will preempt jobs as in ROS-PR while maintaining a Round-Robin order is expected to be even more fair than ROS-PR. This model is more complicated to analyze and is currently studied.

These observations seem to fit with intuition and thus provide support to the validity of RAQFM.

6.3 Service-time Based policies

Service-time based policies typically provide some form of prioritization to short jobs. Common such policies are Shortest Job First (SJF) and Shortest Remaining Processing Time first (SRPT); While these policies have known advantages in reducing mean delay, one may ask then what is the fairness level of such policies as a function of the service time variance. While these are not the subject of this work, rather a subject for further work (their analysis under RAQFM is more complicated), we will only briefly comment on them.

First, some fairness properties of SJF can be observed from [23], where two classes with priorities are studied. It is shown that prioritization of short jobs (jobs belonging to one class) over long jobs (second class) increases system fairness in most cases. However – this is not always the case: When variability is relatively small, full priority given to short jobs may reduce fairness due to the long jobs being blocked for long time.

Second, consider SRPT for which we ran a simulation in the context of this work. The unfairness of SRPT is plotted for the high variability and medium variability cases (Figure 1 and Figure 2). One may observe that the unfairness of SRPT is low (relatively to the job-size independent policies studied here) at low loads. Nonetheless, its unfairness becomes relatively high at medium to high loads. It is interesting to note that such behavior was observed also via the slow-down fairness criterion [29] where it was shown that SPRT obeys that criterion for $\rho < 0.5$ and does not for high values of ρ. This results from the fact that SRPT provides strong priority to short jobs, on the account of the seniority of other jobs.

7 Concluding Remarks

This work aimed at understanding how service time variability affects fairness in queueing time as well as validating RAQFM as a fairness measure. We analyzed basic common service disciplines via models (an exact M/G/1 system for LCFS-PR and Coxian approximation for other disciplines) that accounted for the service time variability. The Coxian approach proves to be quite general and can be extended (to some limited extent) to approximate some cases of more complex disciplines, such as SRPT ; this is the subject of a current study. We showed that the system unfairness is significantly affected by the first two moments of the service time; for the LCFS-PR we showed that higher moments of service time do not affect the unfairness. For LCFS-PR we also showed that the expected discrimination conditioned on the job's service time is identical (zero) to all service times. We demonstrated that fairness is sensitive to service time variability including affecting the relative (fairness) ranking of the various scheduling policies. Our results shed more light on the subject of fairness in queueing systems, demonstrate that RAQFM can be calculated for a wide range of systems, and provide intuitive support for the RAQFM measure.

References

[1] I. Adan and J. Resing. *Queueing Theory.* http://www.cs.duke.edu/ fishhai/misc/queue.pdf, 2001.

[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 and applications. Technical Report RRR-26-2004, RUTCOR, Rutgers University, July 2004. `http://rutcor.rutgers.edu/pub/rrr/reports2004/26_2004.pdf`.

[4] N. Bansal and M. Harchol-Balter. Analysis of SRPT scheduling: investigating unfairness. In *Proceedings of ACM Sigmetrics 2001 Conference on Measurement and Modeling of Computer Systems*, pages 279–290, 2001.

[5] M. Bender, S. Chakrabarti, and S. Muthukrishnan. Flow and stretch metrics for scheduling continuous job streams. In *Proceedings of the 9th Annual ACMSIAM Symposium on Discrete Algorithms*, pages 270–279, San Francisco, CA, 1998.

[6] E. Brosh, B. Avi-Itzhak, and H. Levy. The Effect of Service Time variability on Job Scheduling Fairness. Technical Report. February 2005. `http://www.cs.columbia.edu/~elibrosh/Brosh_Levy_AviItzhak_2005.pdf`.

[7] R. B. Cooper. *Introduction to Queueing Theory.* North -Holland (Elsevier), 2nd edition, 1981.

[8] D. R. Cox and W. L. Smith. *Queues.* Methuen/Wiley, London, 1961.

[9] S. J. Golestani. A self-clocked fair queueing scheme for broadband application. In *Proc. IEEE INFOCOM*, pages 636–646, Toronto, Canada, June 1994.

[10] A. G. Greenberg and N. Madras. How fair is fair queueing? *Journal of the ACM*, 3 (39):568–598, 1992.

[11] F. P. Kelly. *Reversibility and Stochastic Networks.* Wiley and Sons, New York, 1979.

[12] S. Keshav. *An Engineering Approach to Computer Networking: ATM Networks, the Internet, and the Telephone Network.* Addison Wesley Professional, Reading, MA, 1997.

[13] L. Kleinrock. *Queueing Systems, Volume 1: Theory.* Wiley, 1975.

[14] L. Kleinrock. *Queueing Systems, Volume 2: Computer Applications.* Wiley, 1976.

[15] R. C. Larson. Perspective on queues: Social justice and the psychology of queueing. *Operations Research*, 35:895–905, Nov-Dec 1987.

[16] S. A. Lippman. On dynamic programming with unbounded rewards. *Management Science*, 21(11):1225–1233, 1975.

[17] I. Mann. Queue culture: The waiting line as a social system. *Am. J. Sociol.*, 75:340–354, 1969.

[18] R.A. Marie. Calculating equilibrium probabilities for $\lambda(n)/c_k/1/n$ queue. In *Proceedings Performance*, pages 117–125, Toronto, May 1980.

[19] T. Osogami and M. Harchol-Balter A Closed-form Solution for Mapping General Distributions to Minimal PH Distributions In *12th International Conference on Modelling Tools and Techniques for Computer and Communication System Performance Evalua-*

tion, pages 200–217, Urbana, Illinois, USA, September 2003.

[20] C. Palm. Methods of judging the annoyance caused by congestion. *Tele. (English Ed.)*, 2:1–20, 1953.

[21] A. Rafaeli, G. Barron, and K. Haber. The effects of queue structure on attitudes. *Journal of Service Research*, 5(2):125–139, 2002.

[22] A. Rafaeli, E. Kedmi, D. Vashdi, and G. Barron. Queues and fairness: A multiple study investigation. Faculty of Industrial Engineering and Management, Technion. Haifa, Israel. Under review, 2003. `http://iew3.technion.ac.il/Home/Users/anatr/JAP-Fairness-Submission.pdf`.

[23] D. Raz, B. Avi-Itzhak, and H. Levy. Classes, priorities and fairness in queueing systems. Technical Report RRR-21-2004, RUTCOR, Rutgers University, June 2004a. `http://rutcor.rutgers.edu/pub/rrr/reports2004/21_2004.pdf`.

[24] D. Raz, B. Avi-Itzhak, and H. Levy. Fair Operation of Multi-Server and Multi-Queue Systems Submitted.

[25] 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. Also in *Performance Evaluation Review*, 32(1):130-141.

[26] M. H. Rothkopf and P. Rech. Perspectives on queues: Combining queues is not always beneficial. *Operations Research*, 35:906–909, 1987.

[27] 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, Italy.

[28] W. Whitt. The amount of overtaking in a network of queues. *Networks*, 14(3):411–426, 1984.

[29] 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.

[30] Y. Zhou and H. Sethu. On the relationship between absolute and relative fairness bounds. *IEEE Communication Letters*, 6(1):37–39, January 2002.

8 Appendix

8.1 Conditional Discrimination in M/E$_r$/1

8.1.1 Non-Preemptive LCFS

Let C denote the tagged customer. For consistency we assume that the queue is ordered in order of arrival and customers are admitted into service from the tail of the queue. At every slot let a denote the number of customers arrived earlier than C and thus to be served after C. Let b denote the number of customers arrived later than C and thus to be served before C. The state $S_{a,b,j}$ (where j is the stage of the customer in service) captures all that is needed for predicting the future of C. Note that using this description a customer is served in $\mathcal{S}_{a,0,j}$, i.e., when $b = 0$.

An arriving customer starts either in state $\mathcal{S}_{0,0,1}$ when the system is empty, or in state $\mathcal{S}_{k-1,1,j}$ when the system is serving a customer at stage j. Thus

$$E[D|k,j] = \begin{cases} E[d(k-1,1,j)] & k > 0 \\ E[d(0,0,0)] & k = j \end{cases} \tag{28}$$

$$E[D^2|k,j] = \begin{cases} E[d^{(2)}(k-1,1,j)] & k > 0 \\ E[d^{(2)}(0,0,1)] & k = j. \end{cases} \tag{29}$$

Using the same notations as in Section 4.2.1 we have:

$$c(a,b) = \begin{cases} -\frac{1}{a+b+1} & b > 0 \\ 1 - \frac{1}{a+1} & b = 0. \end{cases} \tag{30}$$

Here too, when C is in state $\mathcal{S}_{a,b,j}$ it will encounter one of two possible events:

1. A customer arrives into the system. The probability of this event is $\tilde{\lambda}$. If C was in service ($b = 0$) it will move to $\mathcal{S}_{a+1,0,j}$, otherwise to $\mathcal{S}_{a,b+1,j}$
2. A customer completes its current stage. The probability of this event is $\tilde{\mu}$. If C was in service and $j = r$ it leaves the system. If C wasn't in service it moves to $\mathcal{S}_{a,b,j+1}$ if $j \neq r$ or to $\mathcal{S}_{a,b-1,1}$ if $j = r$.

This leads to the following recursive expressions for $D(a,b,j)$. For $b > 0$

$$D(a,b,j) = \begin{cases} Tc(a,b) + D(a,b+1,j) & \text{w.p } \tilde{\lambda} \\ Tc(a,b) + D(a,b,j+1) & \text{w.p } \tilde{\mu},\ j \neq r \\ Tc(a,b) + D(a,b-1,1) & \text{w.p } \tilde{\mu},\ j = r, \end{cases} \tag{31}$$

and for $b = 0$

$$D(a,0,j) = \begin{cases} Tc(a,0) + D(a+1,0,j) & \text{w.p } \tilde{\lambda} \\ Tc(a,0) + D(a,0,j+1) & \text{w.p } \tilde{\mu},\ j \neq r \\ Tc(a,0) & \text{w.p } \tilde{\mu},\ j = r. \end{cases} \tag{32}$$

From these, as in Section 4.2.1, we derive the recursive equations for expressing $d(a,b,j)$ and $d^{(2)}(a,b,j)$.

For $b > 0$

$$d(a,b,j) = \begin{cases} t^{(1)}c(a,b) + \tilde{\lambda}d(a,b+1,j) + \tilde{\mu}d(a,b,j+1) & j \neq r \\ t^{(1)}c(a,b) + \tilde{\lambda}d(a,b+1,j) + \tilde{\mu}d(a,b-1,1) & j = r \end{cases} \tag{33}$$

$$d^{(2)}(a,b,j) = t^{(2)}(c(a,b))^2 + \tilde{\lambda}d^{(2)}(a,b+1,j) + 2t^{(1)}c(a,b)\tilde{\lambda}d(a,b+1,j))+ \\ \begin{cases} \tilde{\mu}d^{(2)}(a,b,j+1) + 2t^{(1)}c(a,b)\tilde{\mu}d(a,b,j+1) & j \neq r \\ \tilde{\mu}d^{(2)}(a,b-1,j) + 2t^{(1)}c(a,b)\tilde{\mu}d(a,b-1,1) & j = r, \end{cases} \tag{34}$$

and for $b = 0$

$$d(a,0,j) = \begin{cases} t^{(1)}c(a,0) + \tilde{\lambda}d(a+1,0,j)+ & \\ \tilde{\mu}d(a,0,j+1) & j \neq r \\ t^{(1)}c(a,0) + \tilde{\lambda}d(a+1,0,j) & j = r \end{cases} \tag{35}$$

$$d^{(2)}(a,0,j) = t^{(2)}(c(a,0))^2 + \tilde{\lambda}d^{(2)}(a+1,0,j) + 2t^{(1)}c(a,0)\tilde{\lambda}d(a+1,0,j))+ \\ \begin{cases} \tilde{\mu}d^{(2)}(a,0,j+1)+ & \\ 2t^{(1)}c(a,0)\tilde{\mu}d(a,0,j+1) & j \neq r \\ 0 & j = r. \end{cases} \tag{36}$$

8.1.2 Non-Preemptive ROS

For a tagged customer C, let a denote the number of customers in the system other than C. Consider a boolean variable s which is 1 if C is in service and 0 if it is waiting. The notation used in this section remains unchanged, except that b is replaced by s. In state $\mathcal{S}a, s, j$ there are a customers in addition to C, the one in service is in its j-th stage, and it is C if $s = 1$ and not C if $s = 0$.

The momentary discrimination at this state is $c(a,s)$,

$$c(a,s) = \begin{cases} -\frac{1}{a+1} & s = 0 \\ 1 - \frac{1}{a+1} & s = 1. \end{cases} \tag{37}$$

A customer arrives to the system either at state $\mathcal{S}_{0,1,1}$ when it is empty, or at state $\mathcal{S}_{k,0,j}$ when it is serving a customer at stage j. Then

$$E[D|k,j] = \begin{cases} E[d(k,0,j)] & k > 0 \\ E[d(0,1,1)] & k = 0 \end{cases} \tag{38}$$

$$E[D^2|k,j] = \begin{cases} E[d^{(2)}(k,0,j)] & k > 0 \\ E[d^{(2)}(0,1,1)] & k = 0. \end{cases} \tag{39}$$

When C is in state $\mathcal{S}_{a,s,j}$ it will encounter one of the following possible events:

1. If $s = 0$ the possible events are:

(a) A customer arrives into the system. The probability of this event is $\tilde{\lambda}$ and C will move to $\mathcal{S}_{a+1,0,j}$.

(b) A customer completes its current stage j, where $j \neq r$. The probability of this event is $\tilde{\mu}$ and C moves to $S_(a, 0, j+1)$.

(c) A customer completes service, leaves the system and C is chosen to receive service next. The probability of this event is $\tilde{\mu}/a$. C will move to $\mathcal{S}_{a-1,1,1}$.

(d) A customer completes service, leaves the system and C is not chosen to receive service next. The probability of this event is $\tilde{\mu}(a-1)/a$. C will move to $\mathcal{S}_{a-1,0,1}$.

2. If $s = 1$ the possible events are:

(a) Same as (1a) but C will move to $\mathcal{S}_{a+1,1,j}$.

(b) Same as (1b) but C will move to $S_{a,1,j+1}$.

(c) The customer in service completes its service. The probability of this event is $\tilde{\mu}$ and C leaves the system.

Using the same method as in the previous section this leads to the following recursive expressions:

$$d(a,s,j) = t^{(1)}c(a,s) + \tilde{\lambda}d(a+1,s,j) + \begin{cases} \tilde{\mu}d(a,s,j+1) & j \neq r \\ \frac{\tilde{\mu}}{a}d(a-1,1,1) + \tilde{\mu}\frac{a-1}{a}d(a-1,0,1) & j = r, s = 0 \\ 0 & j = r, s = 1 \end{cases} \tag{40}$$

$$d^{(2)}(a,s,j) = t^{(2)}(c(a,s))^2 + \tilde{\lambda}d^{(2)}(a+1,s,j) + 2t^{(1)}c(a,s)\tilde{\lambda}d(a+1,s,j) + \begin{cases} \tilde{\mu}d^{(2)}(a,s,j+1) + 2t^{(1)}c(a,b)\tilde{\mu}d(a,s,j+1)) & j \neq r \\ \frac{\tilde{\mu}}{a}d^{(2)}(a-1,1,1) + \tilde{\mu}\frac{a-1}{a}d^{(2)}(a-1,0,1) + & \\ 2t^{(1)}c(a,s)(\frac{\tilde{\mu}}{a}d(a-1,1,1) + \tilde{\mu}\frac{a-1}{a}d(a-1,0,1)) & j = r, s = 0 \\ 0 & j = r, s = 1. \end{cases} \tag{41}$$

8.1.3 Preemptive ROS

Here we give the recursive equation for the second moment of the discrimination, which is based on the notations and the analysis method described in Section 4.2.2. Following Section

4.2.2 $d^{(2)}(\overline{a}, c, s, j)$ can be expressed as follows. For $s = 0$:

$$d^{(2)}(\overline{a}, c, 0, j) = t^{(2)}c(\overline{a}, s)^2 + \frac{\tilde{\lambda}}{a+2}d^{(2)}(\overline{a} + \overline{1_1}, c, 1, c) + 2t^{(1)}c(\overline{a}, s)\frac{\tilde{\lambda}}{a+2}d(\overline{a} + \overline{1_1}, c, 1, c) +$$
$$\sum_{i:a_i>0} \tilde{\lambda}\frac{\tilde{a}_i}{a+2}d^{(2)}(\overline{a} + \overline{1_1}, c, 0, i) + 2t^{(1)}c(\overline{a}, s) \sum_{i:a_i>0} \tilde{\lambda}\frac{\tilde{a}_i}{a+2}d(\overline{a} + \overline{1_1}, c, 0, i) +$$
$$\begin{cases} \tilde{\mu}d^{(2)}(\overline{a} + \overline{1_{j+1}} - \overline{1_{j-1}}, c, 0, j+1) + 2t^{(1)}c(\overline{a}, s)\tilde{\mu}d(\overline{a} + \overline{1_{j+1}} - \overline{1_{j-1}}, c, 0, j+1) & j \neq r \\ \frac{\tilde{\mu}}{a}d^{(2)}(\overline{a} - \overline{1_r}, c, 1, c) + \sum_{i:a_i>0} \tilde{\mu}\frac{a_k}{a}d^{(2)}(\overline{a} - \overline{1_r}, c, 0, i) + 2t^{(1)}c(\overline{a}, s)(\frac{\tilde{\mu}}{a}d(\overline{a} - \overline{1_r}, c, 1, c) + & \\ \sum_{i:a_i>0} \tilde{\mu}\frac{a_k}{a}d(\overline{a} - \overline{1_r}, c, 0, i)) & j = r. \end{cases} \tag{42}$$

and for $s = 1$:

$$d^{(2)}(\overline{a}, c, 1, j) = t^{(2)}c(\overline{a}, s)^2 + \frac{\tilde{\lambda}}{a+2}d^{(2)}(\overline{a} + \overline{1_1}, c, 1, c) + 2t^{(1)}c(\overline{a}, s)\frac{\tilde{\lambda}}{a+2}d(\overline{a} + \overline{1_1}, c, 1, c) +$$
$$\sum_{i:a_i>0} \tilde{\lambda}\frac{\tilde{a}_i}{a+2}d^{(2)}(\overline{a} + \overline{1_1}, j, 0, i) + 2t^{(1)}c(\overline{a}, s) \sum_{i:a_i>0} \tilde{\lambda}\frac{\tilde{a}_i}{a+2}d(\overline{a} + \overline{1_1}, j, 0, i) +$$
$$\begin{cases} \tilde{\mu}d^{(2)}(\overline{a}, c+1, 1, j+1) + & \\ 2t^{(1)}c(\overline{a}, s)\tilde{\mu}d(\overline{a}, c+1, 1, j+1) & j \neq r \\ 0 & j = r. \end{cases} \tag{43}$$

8.2 Conditional Discrimination in M/Cox$_2$/1

Consider the M/Cox$_2$/1 where the service time distribution is a two-stage Coxian. For this distribution the service is assumed to be composed of two serially arranged stages, where a new customer that enters service starts with stage 1 and after its completion enters stage 2 with probability p_1. The mean length of stage i is μ_i, $i = 1, 2$.

Similarly to M/E$_r$/1 the time between the arrival of a customer and its departure is slotted by arrivals and stage completions, however, for Coxian service the slot duration is dependant upon the stage of service in which the served customer is currently found. Let $T_{i,j}$, $j = 1, 2$, $i = 1, 2, \ldots$ be the duration of the i-th slot, where j is the stage of the served customer. Then, $T_{i,j}$ is a random variable exponentially distributed with parameter $\lambda + \mu_j$; the first two moments of $T_{i,j}$ are $t_j^{(1)} = \frac{1}{\lambda+\mu_j}$ and $t_j^{(2)} = \frac{2}{(\lambda+\mu_j)^2} = 2(t_j^{(1)})^2$. The probabilities that a slot ends with an arrival or with a stage completion are denoted by $\tilde{\lambda}_j$ and $\tilde{\mu}_j$ respectively.

$$\tilde{\lambda}_j = \frac{\lambda}{\lambda + \mu_j} \qquad \tilde{\mu}_j = \frac{\mu_j}{\lambda + \mu_j}. \tag{44}$$

Similarly to Eq. 14 the system unfairness is expressed as:

$$E[D^2] = P_0 E[D^2|0, 1] + \sum_{k=1}^{\infty}\sum_{j=1}^{2} P_{k,j}E[D^2|k, j]. \tag{45}$$

8.2.1 FCFS

We preserve the notations of Section 4.2.1 and denote by $\mathcal{S}_{a,b,j}$ the state of a tagged customer C, where a is the number of customers ahead of C, b is the number of customers behind C, and j is the stage of customer in service.

The momentary discrimination at state $\mathcal{S}_{a,b,j}$ is given by

$$c(a,b) = \begin{cases} -\frac{1}{a+b+1} & b > 0 \\ 1 - \frac{1}{b+1} & b = 0. \end{cases} \tag{46}$$

Similarly to Eq. (16)

$$E[D|k,j] = E[d(k,0,j)] \tag{47}$$

$$E[D^2|k,j] = E[d^{(2)}(k,0,j)]. \tag{48}$$

Using a similar method of analysis as in Section 4.2.1, we have that

$$d(a,b,j) = t_j^{(1)} c(a,b) + \tilde{\lambda}_j d(a,b+1,j) + \begin{cases} p_1 \tilde{\mu}_j d(a,b,j+1) + \Delta(a>0)(1-p_1)\tilde{\mu}_j d(a-1,b,1) & j=1 \\ \Delta(a>0)\tilde{\mu}_j d(a-1,b,1) & j=2, \end{cases} \tag{49}$$

where T_j is the duration of the current slot and the served customer is located at stage j during this slot. The second moment is given by

$$\begin{aligned} d^{(2)}(a,b,j) = {} & t_j^{(2)}(c(a,b))^2 + \tilde{\lambda} d^{(2)}(a,b+1,j) + 2t_j^{(1)} c(a,b)\tilde{\lambda} d(a,b+1,j) + \\ & \begin{cases} p_1\tilde{\mu}_j d^{(2)}(a,b,j+1) + \Delta(a>0)(1-p_1)\tilde{\mu}_j d^{(2)}(a-1,b,1) + \\ 2t_j^{(1)} c(a,b)(p_1\tilde{\mu}_j d^{(2)}(a,b,j+1) + \Delta(a>0)(1-p_1)\tilde{\mu}_j d^{(2)}(a-1,b,1)) & j=1 \\ \Delta(a>0)(\tilde{\mu}_j d^{(2)}(a-1,b,j) + 2t_j^{(1)} c(a,b)\tilde{\mu}_j d(a-1,b,j))) & j=2. \end{cases} \end{aligned} \tag{50}$$

8.2.2 Non-Preemptive LCFS

We preserve the notations of Section 8.1.1 and denote by $\mathcal{S}_{a,b,j}$ the state of C, where a is the number of customers arrived earlier than C and thus to be served after C, b is the number of customers arrived later than C and thus to be served before C, and j is the stage of the served customer. The conditional discrimination (Eqs. (28), (29)) and the momentary discrimination (Eq. (30)) remains the same as in Section 8.1.1.

By examining the possible events that C encounter we get that for $b > 0$

$$d(a,b,j) = t_j c(a,b) + \tilde{\lambda}_j d(a,b+1,j) + \begin{cases} p_1\tilde{\mu}_j d(a,b,j+1) + (1-p_1)\tilde{\mu}_j d(a,b-1,1) & j=1 \\ \tilde{\mu}_j d(a,b-1,1) & j=2 \end{cases} \tag{51}$$

$$d^{(2)}(a,b,j) = t_j^{(2)}(c(a,b))^2 + \tilde{\lambda}_j d^{(2)}(a,b+1,j) + 2t_j^{(1)}c(a,b)d(a,b+1,j) + \begin{cases} p_1\tilde{\mu}_j d^{(2)}(a,b,j+1) + (1-p_1)\tilde{\mu}_j d^{(2)}(a,b-1,1) + 2t_j^{(1)}c(a,b)(p_1\tilde{\mu}_j d(a,b,j+1) + \\ (1-p_1)\tilde{\mu}_j d(a,b-1,1)) & j=1 \\ \tilde{\mu}_j d^{(2)}(a,b-1,1) + 2t_j^{(1)}c(a,b)\tilde{\mu}_j d(a,b-1,1) & j=2. \end{cases} \tag{52}$$

For $b = 0$

$$d(a,0,j) = t_j c(a,0) + \tilde{\lambda}_j d(a+1,0,j) + \begin{cases} p_1\tilde{\mu}_j d(a,0,j+1) & j=1 \\ 0 & j=2 \end{cases} \tag{53}$$

$$d^{(2)}(a,0,j) = t_j^{(2)}(c(a,0))^2 + \tilde{\lambda}_j d^{(2)}(a+1,0,j) + 2t_j^{(1)}c(a,0)d(a+1,0,j) + \begin{cases} p_1\tilde{\mu}_j d^{(2)}(a,0,j+1) + 2t_j^{(1)}c(a,0)p_1\tilde{\mu}_j d(a,0,j+1) & j=1 \\ 0 & j=2. \end{cases} \tag{54}$$

8.2.3 Non Preemptive ROS

We preserve the notations of Section 8.1.2 and denote by $\mathcal{S}a,s,j$ the state of C, where a is the number of customers in the system other than C, s is 1 if C is in service and 0 if it is waiting, and j is the state of the served customer. The conditional discrimination (Eqs. (38), (39)) and the momentary discrimination (Eq. (37)) remains the same as in Section 8.1.2.

By examining the possible events that C encounter we get that

$$d(a,s,j) = t_j^{(1)}c(a,s) + \tilde{\lambda}_j d(a+1,s,j) + \begin{cases} p_1\tilde{\mu}_j d(a,s,j+1) + \Delta(a>0)(1-p_1)\tilde{\mu}_j(\frac{1}{a}d(a-1,1,1) + \frac{a-1}{a}d(a-1,0,1)) & j=1 \\ \Delta(a>0)\tilde{\mu}_j(\frac{1}{a}d(a-1,1,1) + \frac{a-1}{a}d(a-1,0,1)) & j=2 \end{cases} \tag{55}$$

$$d^{(2)}(a,b,j) = t_j^{(2)}(c(a,s))^2 + \tilde{\lambda}_j d^{(2)}(a+1,s,j) + 2t_j^{(1)}c(a,s)\tilde{\lambda}_j d(a+1,s,j) + \begin{cases} p_1\tilde{\mu}_j(d^{(2)}(a,s,j+1) + 2t_j^{(1)}c(a,s)d(a,s,j+1)) + \Delta(a>0)(1-p_1)\tilde{\mu}_j(\frac{1}{a}d^{(2)}(a-1,1,1) + \\ \frac{a-1}{a}d^{(2)}(a-1,0,1) + 2t_j^{(1)}c(a,s)(\frac{1}{a}d(a-1,1,1) + \frac{a-1}{a}d(a-1,0,1)) & j=1 \\ \Delta(a>0)\tilde{\mu}_j(\frac{1}{a}d^{(2)}(a-1,1,1) + \frac{a-1}{a}d^{(2)}(a-1,0,1) + 2t_j^{(1)}c(a,s)(\frac{1}{a}d^{(2)}(a-1,1,1) + \\ \frac{a-1}{a}d^{(2)}(a-1,0,1)) & j=2. \end{cases} \tag{56}$$