A Resource-Allocation Queueing Fairness Measure

David Raz
School of Computer Science
Tel-Aviv University, Tel-Aviv, Israel
davidraz@post.tau.ac.il

Hanoch Levy
School of Computer Science
Tel-Aviv University, Tel-Aviv, Israel
hanoch@cs.tau.ac.il

Benjamin Avi-Itzhak
RUTCOR, Rutgers University
New Brunswick, NJ, USA
aviitzha@rutcor.rutgers.edu

ABSTRACT

Fairness is a major issue in the operation of queues, perhaps it is the reason why queues were formed in the first place. Recent studies show that the fairness of a queueing system is important to customers not less than the actual delay they experience. Despite this observation little research has been conducted to study fairness in queues, and no commonly agreed upon measure of queue fairness exists. Two recent research exceptions are Avi-Itzhak and Levy [1], where a fairness measure is proposed, and Wierman and Harchol-Balter [18] (this conference, 2003), where a *criterion* is proposed for classifying service policies as fair or unfair; the criterion focuses on customer service requirement and deals with fairness with respect to service times.

In this work we recognize that the inherent behavior of a queueing system is governed by two major factors: Job *seniority* (arrival times) and job *service requirement* (service time). Thus, it is desired that a queueing fairness measure would account for both. To this end we propose a Resource Allocation Queueing Fairness Measure, (RAQFM), that accounts for both relative job seniority and relative service time. The measure allows accounting for individual job discrimination as well as system unfairness. The system measure forms a full scale that can be used to evaluate the level of unfairness under various queueing disciplines. We present several basic properties of the measure. We derive the individual measure as well as the system measure for an M/M/1 queue under five fundamental service policies: Processor Sharing (PS), First Come First Served (FCFS), Non-Preemptive Last Come First Served (NP-LCFS), Preemptive Last Come First Served (P-LCFS), and Random Order of Service (ROS). The results of RAQFM are then compared to those of Wierman and Harchol-Balter [18], and the quite intriguing observed differences are discussed.

Categories and Subject Descriptors

C.4 [**Performance of Systems**]: Performance Attributes—*Fairness*; F.2.2 [**Nonnumerical Algorithms and Problems**]: Sequencing and Scheduling; G.3 [**Probability and Statistics**]: Queuing Theory

General Terms

Performance, Measurement

Keywords

Fairness, FCFS, job scheduling, M/M/1, processor sharing, PS, queue disciplines, resource allocation, unfairness

1. INTRODUCTION

Queueing models have long served as key models in a wide variety of fields and applications, including computer systems, human services systems (be it on the web server or at a bank office) and telecommunications systems, to mention just a few. The issue of "fairness" is raised frequently in the context of evaluating queueing policies. In fact, one may argue that one of the main reasons, perhaps the utmost reason, for using a queue, be it in a public office or at a computer system, is to provide some type of fairness to the jobs being served. Evidence to the importance of this measure are common descriptions of queueing systems, such as "this system is fair" or "this queue is not fair". Scientific evidence of the importance of fairness of queues was recently provided in Rafaeli et. al. [13, 14] where psychological experiments demonstrated that for humans waiting in queues, the issue of fairness can be more important than the duration of the wait. In many applications, where an important question is how to operate the queueing system, the two issues most often considered are "efficiency" (which is mostly attributed to the mean waiting time in the system) and "fairness" While efficiency has been extensively studied and is well understood, there is still no commonly agreed upon theoretical yardsticks for measuring "fairness" in queueing systems.

The fairness factor associated with waiting in queues has been recognized in many articles and applications. Larson [8] in his discussion of the dis-utility of waiting, recognizes the central role played by 'Social Justice', and its perception by customers. Aspects of fairness in queues were discussed earlier by quite a number of authors: Palm [10] deals with judging the annoyance caused by congestion, Mann [9] discusses the queue as a social system and Whitt [17] addresses overtaking in queues, to mention just three. In the context of computer applications, the issue of fairness in web servers was discussed by Harchol-Balter et. al. [5]. Service scheduling issues in telecommunication systems at large, and

SIGMETRICS/Performance'04, June 12–16, 2004, New York, NY, USA.

with regards to web servers in particular, have become more crucial with the rapid and diverse increased usage of these systems. As a result, the interest in the fairness levels associated with the different schedules has greatly increased and the fairness factor is recently more often referred to in many applied and theoretic publications.

Despite its great importance, little study has been conducted on the subject of queue fairness and how to quantify it. As a result, the issue of fairness is not understood and agreed upon measures do not exist. Some research exceptions are Wang and Morris [16], Avi-Itzhak and Levy [1], and Wierman and Harchol-Balter [18]. In Wang and Morris [16] the Q-factor is proposed, that measures the performance, relative to global FCFS (namely FCFS among the customers of all sources), as observed by the customer stream treated worst under the worst possible combination of stream loads. Fairness is thus related to classes of customers defined via a measure of extreme treatment given by the service provider with respect to customer streams. In Avi-itzhak and Levy [1] measures based on order of service have been devised, starting with an axiomatic approach. In Wierman and Harchol-Balter [18], the expected normalized delay as function of service time, $E[T(x)/x]$, derived for all service times x, is used as a *criterion* (as opposed to a *measure*) for evaluating whether a system is fair or unfair. A large volume of literature exists on weighted fair queueing (e.g. Demers, Keshav and Shenker [3], Greenberg and Madras [4], Parekh and Gallager [11, 12]); however that work does not deal with *fairness to jobs* (rather it deals with fairness to streams, which fits communications systems), which is the subject of this work, and thus is out of the scope of this paper.

Our departure point for this work is the observation that the behavior of a queueing system is governed by two major physical factors, *job seniority*[1] and *job service requirements*. In all queueing analysis they serve (in the form of arrival times and service times), together with the server policy, to derive the system performance (e.g., expected delay). To this end, a complete fairness measure should account for both. To demonstrate this issue, consider the following case, which perhaps is best depicted in the supermarket queue setup: Mrs. Short arrives to a supermarket queue with a couple of items and finds in front of her Mr. Long with an overflowing cart. The question of whether it is fair to serve Short before Long, and the dilemma associated with this question, is rooted in the contradicting factors of *seniority difference* (working to the benefit of Long) and *service requirement difference* (working to the benefit of Short). The prior work mentioned above focused on one of these factors. Avi-Itzhak and Levy [1] focused on the issue of seniority with a modification of the measure to account for service times. Wierman and Harchol-Balter [18] has successfully captured the service time differences between jobs and provided interesting results regarding which policies are fair in this regard. However, that work does not account for seniority differences.

Very intriguing and drastically contradicting results may be obtained if only one of the factors is accounted for. Thus, a measure that accounts only for seniority differences, as shown in Avi-Itzhak and Levy [1], will rank FCFS as the most fair policy and LCFS as the most unfair policy, in the case of an equal service times system. In contrast, a measure that accounts only for service time differences, such as the criterion developed in Wierman and Harchol-Balter [18] implies that Preemptive LCFS is almost the only fair policy, while FCFS is always unfair. The objective of this work is to propose a measure that will account both for *seniority differences* and for *service time differences*, and be convenient for analysts to work with. To achieve this, our approach is to focus on the server resources and examine how fairly they are allocated to the jobs. This measure is based on the belief that at every epoch all jobs present in the system deserve an equal share of the server's attention, as is the case with Processor Sharing (see analysis, as early as [6, 7, 2] followed by many others) and that deviations from it creates discriminations (positive or negative). Accounting for these discriminations and summarizing them yields a measure of unfairness.

A unique property of RAQFM (in contrast to [16], [18]) is that it tracks the job inter-relations and the resulting unfairness throughout the queue progress process. As a result, it allows to understand fairness at the individual job level (and not only at the level of job classes). This allows us to evaluate individual discrimination and the unfairness of specific scenarios, at the one hand, and the overall unfairness of a system or a policy, on the other hand.

Our work starts (Section 3) with a description of RAQFM, the reasoning behind it, and a review of some results derived in a forthcoming paper by the authors. Similarly to other common queueing measures, such as mean delay, RAQFM is constructed from a measure for individual customer experience (discrimination) whose summary statistics yields a measure of system unfairness. Examination of several simple cases demonstrate that the measure is sensitive both to differences in customer seniority and to differences in customer service times. Having defined the measure we then (Section 4) conduct the analysis of four basic queueing disciplines (the analysis of the Processor Sharing is trivial, and is provided in Section 3): First Come First Served (FCFS), Non-Preemptive Last Come First Served (NP-LCFS) Preemptive Last Come First Served (P-LCFS), and Random Order of Service (ROS), under the M/M/1 paradigm. For all these disciplines we derive the expected discrimination experienced by an arriving customer as a function of the number of customers it finds in the system upon arrival. The results shed light on customer discrimination as a function of the queue situation it encounters. We then (Section 5) derive the system unfairness, measured via the variance of discrimination, for these four policies. Having derived these measures we then review them and discuss them (Section 6); we also compare them to the recent results of Wierman and Harchol-Balter [18]. Concluding remarks are given in Section 7.

2. MODEL AND NOTATIONS

For proposing the fairness measure we consider a very general queueing model, where no restrictive assumptions have to be made. To this end we consider a single-server queueing system[2], that is subject to the arrivals of a stream of customers, denoted $C_1, C_2, \ldots$, who arrive at the system at that order. We let a_i denote the arrival epoch of customer

[1] The terms "job" and "customer" are used interchangeably throughout the paper

[2] The Multi-server queueing system is discussed in a forthcoming work

C_i and s_i the total service requirement (measured in time units) of C_i. The departure epoch of C_i is denoted d_i. We assume that the server is work conserving and non-idling. The service order used by the server can be arbitrary. Further, the server does not have to serve the customers one by one; rather it can allocate pieces of service time, of arbitrary length, to the customers, or even serve customers concurrently. To account for this, we have $s_i(t)dt$ denote the infinitesimal amount of service C_i receives at epoch t (alternatively, define $s_i(t)$ to be the fraction of the capacity of the server devoted to C_i, at epoch t). Note that in any non-idling system, the total amount of service rate applied at every epoch t at which the system is not empty, is unity, i.e.,

$$\sum_{i|t\in(a_i,d_i)} s_i(t) = 1. \tag{1}$$

If C_i is served alone at t then obviously $s_i(t) = 1$. Also we have $\int_{a_i}^{d_i} s_i(t)dt = s_i$.

Once the fairness measure is defined, we will turn into analyzing several specific queueing disciplines. For the sake of that analysis we will deal with an M/M/1 type system. That is, the arrivals of customers to the system follow a Poisson process with rate λ and the service requirements of the customers are i.i.d. random variables whose duration is exponential with mean $1/\mu$.

Within the framework of the M/M/1 system, we will analyze the fairness of four fundamental service orders: 1) First Come First Served (FCFS). 2) Non-Preemptive Last Come First Served (NP-LCFS) in which when the server picks a job for service it always picks the youngest (least senior) job in the system; Once a job is in service, its service is not preempted by new arrivals. 3) Preemptive Last Come First Served (P-LCFS) which is identical to NP-LCFS in the order of service but in which a new arrival does preempt the job in service and starts get served upon arrival. 4) Random Order of Service (ROS). In this policy whenever the server selects a job for service, it does so by randomly selecting, with equal probability, a single job from all jobs present at the queue. Note that for ROS we consider a non-preemptive policy, i.e. Once a job is in service, its service is not preempted by new arrivals.

3. INTRODUCING RAQFM

3.1 Temporal discrimination and 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 [3] of the system resources. Denote by $N(t)$ the number of customers present in the system at epoch t Then that principle implies that the share of the server resources a customer deserves at t is simply given by $1/N(t)$. We call this quantity the *warranted service* of C_i at epoch t and denote it $R_i(t)$.

The *temporal discrimination* of C_i at epoch t, denoted $D_i(t)$, is defined as the amount of service granted to C_i at t, minus the warranted service of C_i at t, that is:

$$D_i(t) = s_i(t) - R_i(t) = s_i(t) - 1/N(t). \tag{2}$$

The overall *discrimination* of C_i is the discrimination it accumulates over its stay in the system; we denote it by D_i and its value is:

$$D_i = \int_{a_i}^{d_i} D_i(t)dt = \int_{a_i}^{d_i} (s_i(t) - 1/N(t))dt. \tag{3}$$

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.*

An important property of this measure is that it obeys, for every non-idling service conserving system, and for every t: $\sum_i D_i(t) = 0$, that is, that every positive discrimination is balanced with negative discrimination. This results from the fact that when the system is non-empty $\sum_i s_i(t) = 1$ (due to non-idling) and $\sum_i R_i(t) = N(t)(1/N(t)) = 1$. An important outcome of this property is that if D denotes the discrimination of an arbitrary customer where the system is in steady state, we get that the expected discrimination is zero, that is:

$$E[D] = 0. \tag{4}$$

A full proof is given in [15].

An interesting value, which we will analyze and discuss in the sequel, is the expected discrimination experienced by a tagged customer, and conditioned on the number of customers (k) the tagged customer finds in the system upon its arrival. This is denoted $E[D|k]$. A similar notation for the conditional second moment is $E[D^2|k]$.

3.2 System Measure of Unfairness

To measure the unfairness of a system and of a service policy across all customers, that is, to measure the *system unfairness*, one would choose some summary statistics measure over the variables D_i. We note that the most natural choice, the first moment of customer discrimination, cannot serve as a system unfairness measure, Since $E[D] = 0$. Thus, a very natural choice is that of the variance of customer discrimination, that is $Var[D]$, which we will adopt in this work[4]. This measure is also identical to the second moment of discrimination, $Var[D] = E[D^2]$, resulting from the fact that $E[D] = 0$.

It is interesting to note that the use of the variance is perhaps natural as fairness inherently deals with differences in treatment of customers.

3.3 Unfairness of a Scenario

A *scenario* is a sequence of arrivals, service actions and departures, occurring in a queue to a finite set of customers $C_i, i = 1, \ldots, L$. To evaluate the unfairness of a scenario one can run (emulate) the service as conducted in the system, and compute the set of individual discriminations $D_i, i = 1, \ldots, L$. The *scenario unfairness* can then be obtained via the statistical second moment (which equals the variance in our case) of these discriminations, that is $\frac{1}{L}\sum_{i=1}^{L}(D_i)^2$.

[3]In the model described here the underlying assumption is that all customers are “born equal”, and thus no weights are assigned to them. In an ongoing research we deal with a weighted version of the measure.

[4]Alternative measures can be $E[|D|]$ and $E[D|D > 0]$ which seem to be mathematically less convenient and which are beyond the scope of this paper.

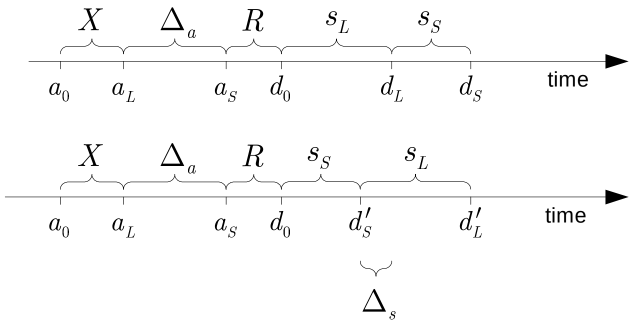

Figure 1: Mr. Long and Mrs. Short - Seniority vs. Service Requirement

3.4 Basic Properties of RAQFM

Below we review some basic properties of the fairness measure. The proof of these properties is provided in a forthcoming work.

1. Under the processor sharing scheduling (PS), the individual discrimination obeys: $D_i = 0$ for every i, regardless of the arrivals and service times. Thus the variance of discrimination, that is, the system unfairness, obeys: $Var[D] = E[D^2] = 0$. That means that PS is the most fair policy under this measure.

2. The value of individual discrimination is bounded from below and above by $-W_i/2 < D_i < s_i$, where W_i is the waiting time of C_i and s_i is its service time.

3. The value of system unfairness, for a system with *arbitrary service times and arrivals*, achieved via the measure of $E[D^2]$, is bounded by: $0 \le E[D^2] < \frac{N_{max}}{2} s_{max}$, where s_{max} is the maximal service requirement and N_{max} is the maximal number of customers served in a busy period (if the maximum exists).

3.4.1 *Sensitivity of the measure to seniority difference and to service requirement difference*

In this section we demonstrate the sensitivity of RAQFM to both seniority differences and service requirement differences. To do this, we recall the basic dilemma regarding the prioritization of Mr. Long vs. Mrs. Short presented in the introduction and consider a scenario that is a special case of that problem.

Consider the following scenario. C_0, serving as a "by-stander" for this analysis, arrives at the empty system at a_0 and leaves at d_0. At $a_0 < a_L < d_0$ arrives C_L (Long) with a service requirement of s_L. At $a_L < a_S < d_0$ arrives C_S (Short) with a service requirement of $s_S < s_L$. We now ask the question, whether it is more fair, under RAQFM, to serve C_S ahead of C_L.

Let X denote the length of the interval (a_0, a_L) $X = a_L - a_0$. Let R denote the remainder of the service requirement left for C_0 after a_S, i.e. $R = d_0 - a_S$. Let Δ_a denote the difference in seniority between C_S and C_L, i.e. $\Delta_a = a_S - a_L$. Let Δ_s denote the difference in service requirement between C_S and C_L, i.e. $\Delta_a = a_L - a_S$.

Figure 1 illustrates the two possible orders of service. In the top half of the figure C_L is served before C_S, and in the bottom half the opposite.

If C_L is served before C_S, then C_L leaves the system at $d_L = d_0 + s_L$ and C_S leaves the system at $d_S = d_0 + s_S + s_L$. The discriminations of the three customers are therefore

$$\begin{aligned} D_0 &= (X + \Delta_a + R) - (X + \Delta_a/2 + R/3) \\ D_L &= s_L - (\Delta_a/2 + R/3 + s_L/2) \\ D_S &= s_S - (R/3 + s_L/2 + s_S), \end{aligned}$$

and the total unfairness is $(D_0^2 + D_L^2 + D_S^2)/3$.

If C_S is served before C_L, then C_L leaves the system at $d'_L = d_0 + s_S + s_L$ and C_S leaves the system at $d'_S = d_0 + s_S$. The discriminations of the three customers are therefore

$$\begin{aligned} D'_0 &= (X + \Delta_a + R) - (X + \Delta_a/2 + R/3) \\ D'_L &= s_L - (\Delta_a/2 + R/3 + s_S/2 + s_L) \\ D'_S &= s_S - (R/3 + s_S/2), \end{aligned}$$

and the total unfairness is $((D'_0)^2 + (D'_L)^2 + (D'_S)^2)/3$.

The difference between the fairness values of the two possibilities is therefore

$$\frac{(D'_0)^2 + (D'_L)^2 + (D'_S)^2 - D_0^2 + D_L^2 + D_S^2}{3} = \frac{(s_S + s_L)(\Delta_a - \Delta_s)}{6}. \quad (5)$$

Expression (5) now reveals that the unfairness difference in this case is monotonic in the difference between the seniority difference and the service requirement difference, that is in $\Delta_a - \Delta_s$. The higher this difference is the more unfair it is to serve Short first; the lower it is, the more unfair it is to serve Long first. The point of indifference, is exactly when these differences equal each other, that is $\Delta_a = \Delta_s$.

It is clear from this scenario that RAQFM accounts for both seniority difference and service requirement difference, in the proper direction.

4. CONDITIONAL DISCRIMINATION IN M/M/1

In this section we consider the M/M/1 system and derive $E[D|k]$, the expected customer discrimination, conditioned on the number of customers it finds in the system. While $E[D|k]$ is not used to derive the system unfairness, it is still an interesting measure that will serve to shed light on the situations at which systems are subject to high discrimination.

Let D be a random variable denoting the discrimination experienced by an arbitrary (tagged) customer when the system is in steady state. Let $E[D|k]$ denote the expected value of D conditioned on the number of customers, k, found in the system upon the tagged customer arrival. Similarly let $E[D^2|k]$ be the conditional second moment of D.

(queue and server) in

In a work conserving non-idling M/M/1 system the time between the arrival of a customer and its departure is slotted by arrivals and departures of customers. 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 $\lambda + \mu$ (and first two moments $t^{(1)} = \frac{1}{\lambda+\mu}$, $t^{(2)} = \frac{2}{(\lambda+\mu)^2} = 2(t^{(1)})^2$). The number of customers present in the system during slot i is denoted by N_i. We note that N_i changes by 1, up or down, at the end of the slot. The probabilities that a slot ends with an arrival or with a departure are denoted by $\tilde{\lambda}$ and $\tilde{\mu}$ respectively.

$$\tilde{\lambda} = \frac{\lambda}{\lambda+\mu} = \frac{\rho}{1+\rho} \qquad \tilde{\mu} = \frac{\mu}{\lambda+\mu} = \frac{1}{1+\rho}, \tag{6}$$

where $\rho = \lambda/\mu < 1$.

In the disciplines considered in this work, the customer in service receives the full service rate of the server for the duration of a whole slot. The discrimination, D, of an arbitrarily selected customer in steady-state can be written as

$$D = \sum_{i\in S} T_i - \sum_{i\in S\cup W} T_i/N_i, \tag{7}$$

where S is the set of slots in which the customer is being served and W is the set of slots in which it is waiting. Note that while $E[D]=0$, the conditional expected discrimination $E[D|k]$, is not necessarily zero and is a function of the state of the system, k, encountered by the customer. Below we derive the value of $E[D|k]$ for several key service policies.

4.1 FCFS

Let us consider a tagged customer C. At every slot let a denote the number of customers ahead of C and b the number of customers behind C. Due to the memoryless properties of the system, the state $\mathcal{S}_{a,b}$ captures all that is needed to predict the future of C. The momentary discrimination at state $\mathcal{S}_{a,b}$, is denoted $c(a,b)$ and is the rate of discrimination at which customer discrimination accumulates when the customer is at that state. From Equation (2) this is given by:

$$c(a,b) = \begin{cases} -\frac{1}{a+b+1} & a>0 \\ 1-\frac{1}{b+1} & a=0 \end{cases}. \tag{8}$$

The discrimination accumulated for the customer over a slot of length T at state $\mathcal{S}_{a,b}$ is $c(a,b)T$.

Let $D(a,b)$ denote the discrimination of a walk starting at state $\mathcal{S}_{a,b}$. Using this notation we get:

$$E[D|k] = E[D(k,0)]. \tag{9}$$

Assume C is in state $\mathcal{S}_{a,b}$ in slot i. C will encounter one of the two following evens at the next slot $i+1$.

1. A customer arrives into the system. The probability of this event is $\tilde{\lambda}$. Afterward C will move to state $\mathcal{S}_{a,b+1}$.
2. A customer leaves the system. The probability of this event is $\tilde{\mu}$. If C is not being served ($a\neq 0$) C will afterward move to state $\mathcal{S}_{a-1,b}$, else C leaves the system.

Thus for customers not in service

$$D(a,b) = \begin{cases} T_i c(a,b) + D(a,b+1) & \text{with probability } \tilde{\lambda} \\ T_i c(a,b) + D(a-1,b) & \text{with probability } \tilde{\mu} \end{cases}, \tag{10}$$

and for customers in service

$$D(a,b) = \begin{cases} T_i c(a,b) + D(a,b+1) & \text{with probability } \tilde{\lambda} \\ T_i c(a,b) & \text{with probability } \tilde{\mu} \end{cases}. \tag{11}$$

REMARK 4.1. *In both equations above the sums represent sums of independent random variables. This is correct since the future discrimination (expressed by the $D()$ variable on the right-hand-side) is independent of the length of the current slot (expressed by $T_i c(a,b)$); also the future state is independent of the length of the current slot (since regardless of the length, the transition probabilities are $\tilde{\lambda}$ and $\tilde{\mu}$).*

Let $d(a,b) \stackrel{def}{=} d^{(1)}(a,b)$ and $d^{(2)}(a,b)$ be the first and second moment of $D(a,b)$. Combining (10) and (11) and taking expectations leads to the following recursive expression:

$$d(a,b) = \begin{cases} t^{(1)}c(a,b) + \tilde{\lambda}d(a,b+1) + \tilde{\mu}d(a-1,b) & a>0 \\ t^{(1)}c(a,b) + \tilde{\lambda}d(a,b+1) & a=0 \end{cases}. \tag{12}$$

REMARK 4.2. *Note that the quantity $\frac{E[D|k]}{t^{(1)}}$ depends on λ and μ only through their ratio $\rho=\lambda/\mu$, and not on their individual values. This is true for all the policies studied in this paper.*

Equation (12) can be used, via numerical recursion, to compute the values of $d(a,b)$. From these, and using Equation (9), one can derive $E[D|k]$.

Alternatively, one can compute $E[D|k]$ using the result of the following theorem.

THEOREM 4.1. *The following equality holds:*

$$E[D|k] = t^{(1)} \sum_{b=0}^{\infty}\sum_{a=0}^{k} E[G(a,b)|k]c(a,b), \tag{13}$$

where $E[G(a,b)|k]$, the probability that a walk reaches $\mathcal{S}_{a,b}$ given that the customer sees k other customers on arrival is

$$E[G(a,b)|k] = \tilde{\lambda}^b\tilde{\mu}^{k-a}\binom{k-a+b}{b} = \frac{\rho^b}{(1+\rho)^{k+b-a}}\binom{k-a+b}{b}. \tag{14}$$

The proof is given in Appendix A.

4.1.1 Special Cases

$E[D|0]$, the expected discrimination for a customer arriving at an empty system can be easily derived from (13):

$$E[D|0] = t^{(1)}\sum_{b=0}^{\infty}\tilde{\lambda}^b\left(1-\frac{1}{b+1}\right) = t^{(1)}\left(\frac{1}{1-\tilde{\lambda}} - \sum_{b=0}^{\infty}\frac{\tilde{\lambda}^b}{b+1}\right) \tag{15}$$

The second part of this sum yields:

$$\sum_{b=0}^{\infty}\frac{\tilde{\lambda}^b}{b+1} = \frac{1}{\tilde{\lambda}}\int\sum_{b=0}^{\infty}\tilde{\lambda}^b d\tilde{\lambda} = \frac{1}{\tilde{\lambda}}\int\frac{1}{1-\tilde{\lambda}}d\tilde{\lambda} = \frac{1}{\tilde{\lambda}}\ln(1-\tilde{\lambda}), \tag{16}$$

and so

$$E[D|0] = t^{(1)}\left(\frac{1}{1-\tilde{\lambda}} + \frac{1}{\tilde{\lambda}}\ln(1-\tilde{\lambda})\right) = t^{(1)}\left(1+\rho - \frac{1+\rho}{\rho}\ln(1+\rho)\right). \tag{17}$$

Observe that $E[D|0]$ is a monotone-increasing function in ρ and thus its maximum value is reached when $\rho\to 1$ and its minimum value is reached when $\rho \to 0$, the high and low traffic bounds. For the high traffic bound

$$E[D|0] \xrightarrow{\rho \to 1} t^{(1)}(2 - 2\ln 2) \approx 0.613706 t^{(1)}, \qquad (18)$$

which is the maximum expected discrimination for a user in the FCFS service policy. For the low traffic bound

$$E[D|0] \xrightarrow{\rho \to 0} t^{(1)}(1 + 0 - 1) = 0. \qquad (19)$$

This is expected since when $\rho \to 0$ there are almost no arriving customers. Thus an arriving tagged customer finding the system empty will be served and leave without experiencing arrivals during its stay, and thus its discrimination approaches $D = 0$.

Another value easily derived is the limit of $E[D|k]$ for $\rho \to 0$. If $\rho \to 0$ it follows immediately that $\tilde{\lambda} \to 0$ and $\tilde{\mu} \to 1$ and thus $E[G(a,b)|k] = 0$ for $b > 0$ and so (13) becomes

$$E[D|k] \xrightarrow{\rho \to 0} t^{(1)} \sum_{a=0}^{k} c(0,a) = 0 + t^{(1)} \sum_{a=1}^{k} -\frac{1}{a+1} = -t^{(1)} H_k, \qquad (20)$$

where H_n is the n-th harmonic number.

For the numerical computation of other values using (13), the summation must be stopped at some large number, justified by the low probability of having very large values of b, i.e. $b((a,b)|k) \xrightarrow{b \to \infty} 0$. For the computation using (12) the same applies as $d(a,b) \xrightarrow{b \to \infty} 0$. This can be simply done by setting $d(a,b)$ to be zero for some large value of a and b. A second way to justify these is to look at a system with a limited (but large) queue.

4.1.2 Numerical results and properties

Figure 2 depicts the value of the conditional (normalized by $t^{(1)}$) discrimination $\frac{E[D|k]}{t^{(1)}}$ as a function of k, for some values of ρ. The normalized discrimination due to the fact mentioned in Remark 4.2, namely that this enables us to create a single plot that covers every value of λ and μ, by creating a plot of ρ. We will follow this convention throughout this paper.

All the special cases presented above can be verified in the figure. In addition we note the following properties:

1. The worst (most negative) discrimination a customer may experience is when the load approaches zero and the customer finds a very long queue; in this case the expected discrimination monotonically decreases in the queue length (k) and is unbounded. This seems to agree with common feelings where perhaps the most disappointing queue state one can encounter is a long queue when the load is very small.

2. The best (most positive) discrimination a customer may experience is to find an empty queue when the load is very high. This, again, seems to agree with common customer feelings.

3. Negative discrimination appears to monotonically increase with the queue size encounters. This seems to fit our intuition as well.

4. Normalized discrimination seems to monotonically decrease with the load, ρ.

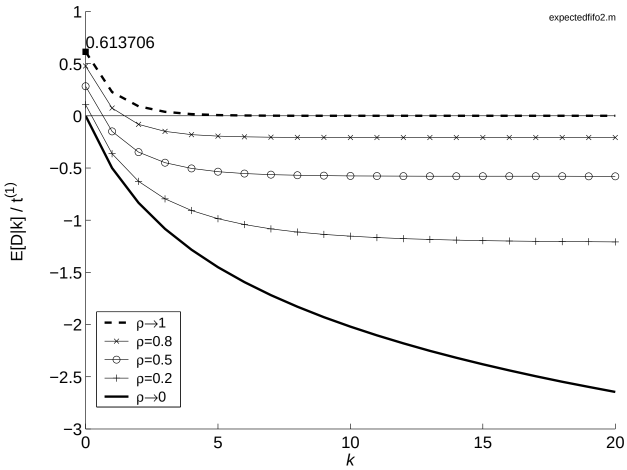

Figure 2: Conditional (normalized) Discrimination for M/M/1 under FCFS

4.2 Non-Preemptive LCFS

Let us consider a tagged customer C. For the sake of description we assume that the queue is ordered in order of arrival and customers are taken for service from the tail of the queue. At every slot let a denote the number of customers ahead of C at the queue (and thus to be served after C). Let b denote the number of customers behind C at the queue (and thus to be served before C).

Defining $c(a,b)$, the momentary discrimination of state $\mathcal{S}_{a,b}$, as in Section 4.1, and noting that the customer is served when $b = 0$ we have:

$$c(a,b) = \begin{cases} -\frac{1}{a+b+1} & b > 0 \\ 1 - \frac{1}{b+1} & b = 0 \end{cases}. \qquad (21)$$

Here too, when C is in state $\mathcal{S}_{a,b}$ in slot i it will encounter one of two possible evens at the next slot $i+1$:

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 $\mathcal{S}_{a+1,b}$, otherwise to $\mathcal{S}_{a,b+1}$

2. A customer leaves the system. The probability of this event is $\tilde{\mu}$. If C was in service, C leaves the system, otherwise C moves to $\mathcal{S}_{a,b-1}$.

Using the same notation, definition and method as in Section 4.1, this leads to the following recursive expression for expressing $D(a,b)$. For customers not under service

$$D(a,b) = \begin{cases} T_i c(a,b) + D(a,b+1) & \text{with probability } \tilde{\lambda} \\ T_i c(a,b) + D(a,b-1) & \text{with probability } \tilde{\mu} \end{cases}, \qquad (22)$$

and for customers under service

$$D(a,b) = \begin{cases} T_i c(a,b) + D(a+1,b) & \text{with probability } \tilde{\lambda} \\ T_i c(a,b) & \text{with probability } \tilde{\mu} \end{cases}. \qquad (23)$$

From these, and like in Section 4.1, we derive the recursive equations for computing $d(a,b)$:

$$d(a,b) = \begin{cases} t^{(1)}c(a,b) + \tilde{\lambda}d(a,b+1) + \tilde{\mu}d(a,b-1) & b > 0 \\ t^{(1)}c(a,b) + \tilde{\lambda}d(a+1,b) & b = 0 \end{cases}. \tag{24}$$

Observe that an arriving customer starts either with state $\mathcal{S}_{0,0}$ when the the system is empty, or with $\mathcal{S}_{1,k-1}$ when it is not empty, thus

$$E[D|k] = \begin{cases} d(k-1,1) & k > 0 \\ d(0,0) & k = 0 \end{cases}. \tag{25}$$

From (24) and (25) one can compute $E[D|k]$ recursively. A more efficient approach is presented in Theorem 4.2 and Lemma 4.1 as follows.

THEOREM 4.2. *Consider the infinite series a_n, b_n, $n = 1, 2, \ldots$ defined recursively as follows:*

$$a_1 = 0 \qquad b_1 = 1 \tag{26}$$

$$a_2 = t^{(1)}c_1 + \tilde{\mu}d^{(s)} \qquad b_2 = \tilde{\lambda} \tag{27}$$

$$\begin{aligned} a_n &= \frac{a_{n-1} + t^{(1)}b_{n-1}c_{n-1} - b_{n-1}\tilde{\mu}\frac{a_{n-2}}{b_{n-2}}}{1 - \frac{b_{n-1}\tilde{\mu}}{b_{n-2}}} \\ &= \frac{b_{n-2}(a_{n-1} + t^{(1)}b_{n-1}c_{n-1}) - \tilde{\mu}b_{n-1}a_{n-2}}{b_{n-2} - b_{n-1}\tilde{\mu}} \qquad ,n > 2 \end{aligned} \tag{28}$$

$$b_n = \frac{\tilde{\lambda}b_{n-1}}{1 - \frac{b_{n-1}\tilde{\mu}}{b_{n-2}}} = \frac{\tilde{\lambda}b_{n-1}b_{n-2}}{b_{n-2} - b_{n-1}\tilde{\mu}} \qquad ,n > 2 \tag{29}$$

where $c_n = c(k-1,n)$ and

$$d^{(s)} = t^{(1)}\left(1 + \rho - \sum_{a=0}^{\infty}(\frac{\rho}{1+\rho})^a\frac{1}{a+k}\right). \tag{30}$$

Then for every $n \geq 1$ and $k > 0$ the following equality (for expressing $E[D|k]$) holds:

$$E[D|k] = a_n + b_n d(k-1,n). \tag{31}$$

The proof is given in Appendix B.

LEMMA 4.1. *The series b_n given above satisfies*

$$0 < b_n/b_{n-1} < \rho. \tag{32}$$

In other words, $b_n \xrightarrow{n\to\infty} 0$.

PROOF. The proof is by induction.
Induction base: for $n = 2$

$$b_2/b_1 = \tilde{\lambda}/1 = \frac{\rho}{1+\rho} < \rho, \tag{33}$$

and of course $\tilde{\lambda} > 0$.
Induction assumption: $0 < b_{n-1}/b_{n-2} < \rho$
Induction step: From (29), b_n can be written as

$$b_n = \frac{\tilde{\lambda}b_{n-1}}{1 - \frac{b_{n-1}\tilde{\mu}}{b_{n-2}}} = \frac{\frac{\rho}{1+\rho}b_{n-1}}{1 - \frac{b_{n-1}\frac{1}{1+\rho}}{b_{n-2}}} = \rho b_{n-1}\frac{1}{1 + \rho - \frac{b_{n-1}}{b_{n-2}}}. \tag{34}$$

From the assumption

$$\frac{b_{n-1}}{b_{n-2}} < \rho \Rightarrow 1 + \rho + \frac{b_{n-1}}{b_{n-2}} > 1 \Rightarrow 0 < \frac{1}{1 + \rho - \frac{b_{n-1}}{b_{n-2}}} < 1. \tag{35}$$

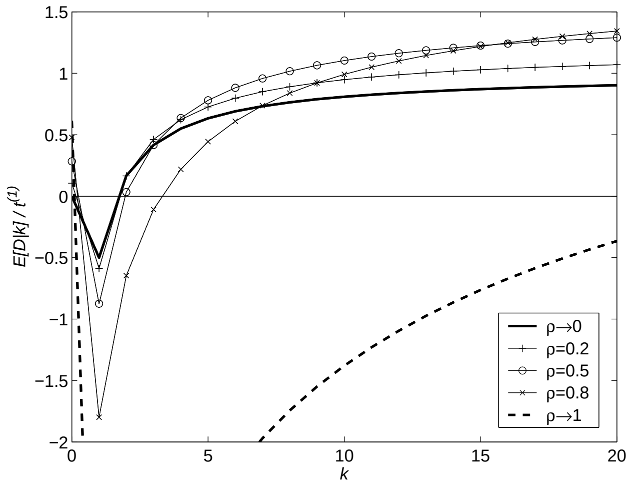

Figure 3: Conditional (normalized) Discrimination for M/M/1 under NP-LCFS

Substituting this in (34) yields

$$0 < b_n < \rho b_{n-1} \tag{36}$$

proving the lemma. □

Theorem 4.2 and Lemma 4.1 provide an iterative way (as opposed to a recursive one) for evaluation $E[D|k]$. The idea is that one can take an arbitrary value of n, compute a_n and b_n (iteratively) and $d(k-1,n)$ (recursively) and from them get $E[D|k]$. Now, from Lemma 4.1 $b_n \xrightarrow{n\to\infty} 0$ while $d(k-1,n)$ is bounded, we get that $a_n \xrightarrow{n\to\infty} E[D|k]$. That is $E[D|k]$ can be computed from the limit of a_n when $n \to \infty$.

4.2.1 Special Cases

First note that for $k = 0$ all non-preemptive service policies yield the same results thus

$$E[D|0] = t^{(1)}\left(1 + \rho - \frac{1+\rho}{\rho}\ln(1+\rho)\right), \tag{37}$$

$$E[D|0] \xrightarrow{\rho\to 1} t^{(1)}(2 - 2\ln 2) \approx 0.613706, \tag{38}$$

$$E[D|0] \xrightarrow{\rho\to 0} 0. \tag{39}$$

It is also simple to evaluate $E[D|k]$ when $\rho \to 0$. In this case an arriving customer will almost surely wait for the current customer to be served, then be served herself, i.e.

$$E[D|k] \xrightarrow{\rho\to 0} t^{(1)}\left(\frac{1}{k+1} + 1 - \frac{1}{k}\right) = \frac{t^{(1)}(k^2 - k - 1)}{k(k+1)}. \tag{40}$$

4.2.2 Numerical results and properties

Figure 3 depicts the conditional (normalized by $t^{(1)}$) discrimination $\frac{E[D|k]}{t^{(1)}}$ as a function of k, for some values of ρ. We note the following properties:

1. The highest value of negative discrimination is achieved at very high load and when $k = 1$. In this case the customer arrives and finds only the customer in service ahead of him, but then he ends up waiting for long time due to customers arriving after him. Thus a very high negative discrimination fits the natural intuition.

2. For all load values the highest value of negative discrimination is achieved when $k = 1$, which is, again, intuitive.

3. Relatively large values of positive discrimination are achieved when k is large. In these cases the arriving customer gets positively discriminated since it is served ahead of $k - 1$ of those customers.

4.3 Preemptive LCFS (P-LCFS)

The analysis done in Section 4.2 for NP-LCFS holds also to P-LCFS when it comes to state definition and to the momentary discrimination. However, when C is in state $\mathcal{S}_{a,b}$ in slot i it will encounter a slightly different set of possible evens at the next slot $i + 1$:

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

2. A customer leaves the system. The probability of this event is $\tilde{\mu}$. If C was in service, C leaves the system, otherwise C moves to $\mathcal{S}_{a,b-1}$.

Using the same method and notation as in Section 4.1, this leads to the following recursive expression for computing $d(a, s)$:

$$d(a,b) = \begin{cases} t^{(1)}c(a,b) + \tilde{\lambda}d(a,b+1) + \tilde{\mu}d(a,b-1) & b > 0 \\ t^{(1)}c(a,b) + \tilde{\lambda}d(a,b+1) & b = 0 \end{cases}. \quad (41)$$

Observe that an arriving customer starts service immediately and thus

$$E[D|k] = d(k, 0). \quad (42)$$

4.3.1 Special Cases

It is simple to evaluate $E[D|k]$ when $\rho \to 0$. In this case an arriving customer will almost surely be served in full and leave the system, thus

$$E[D|k] \xrightarrow{\rho\to 0} t^{(1)}\left(1 - \frac{1}{k+1}\right) = \frac{t^{(1)}k}{k+1} \xrightarrow{k\to\infty} t^{(1)}. \quad (43)$$

4.3.2 Numerical results and properties

Figure 4 depicts the conditional (normalized by $t^{(1)}$) discrimination $\frac{E[D|k]}{t^{(1)}}$ as a function of k, for some values of ρ. We note the following properties:

1. The highest value of negative discrimination is achieved at very high load and when $k = 0$. In this case the customer arrives and finds no customer in service ahead of him, but then he has a high chance of ending up waiting for long time due to customers arriving after him. Thus a high negative discrimination fits the natural intuition.

2. For all load values the highest value of negative discrimination is achieved when $k = 0$, which is, again, intuitive.

3. Relatively large values of positive discrimination are achieved when k is large. In these cases the arriving customer gets positively discriminated since it is served ahead of those k customers.

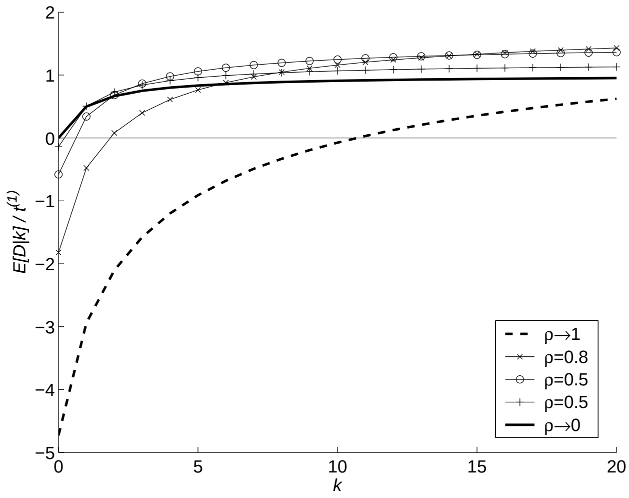

Figure 4: Conditional (normalized) Discrimination for M/M/1 under P-LCFS

4.4 Random Order of Service (ROS)

Let us consider a tagged customer C. Let a denote the number of customers in the system other than c. Consider a boolean state s which is 1 if the customer is in service and 0 if it is waiting. Due to the memoryless properties of the system, the state $\mathcal{S}_{a,s}$ captures all that is needed to predict the future of C. The momentary discrimination at state $\mathcal{S}_{a,s}$, is denoted $c(a, s)$ and is the rate of discrimination at which customer discrimination accumulates when the customer is at that state. From Equation (2) this is given by:

$$c(a,s) = \begin{cases} -\frac{1}{a+1} & s = 0 \\ 1 - \frac{1}{a+1} & s = 1 \end{cases}. \quad (44)$$

The discrimination accumulated for the customer over a slot of length T at state $\mathcal{S}_{a,s}$ is $c(a, s)T$.

Again, assume C is in state $\mathcal{S}_{a,s}$ in slot i. C will encounter one of the following events at the next slot $i + 1$, assuming C is not being served in slot i ($s = 0$)

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

2. A customer leaves the system, and C is chosen to receive service next. The probability of this event is $\tilde{\mu}/a$. C will move to $\mathcal{S}_{a-1,0}$.

3. A customer 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 $\mathcal{S}_{a-1,1}$.

If C is being served in slot i it will encounter one of the following evens at the next slot $i + 1$

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

2. C leaves the system. The probability of this event is $\tilde{\mu}$.

Using the same method and notation as in Section 4.1, this leads to the following recursive expression for computing

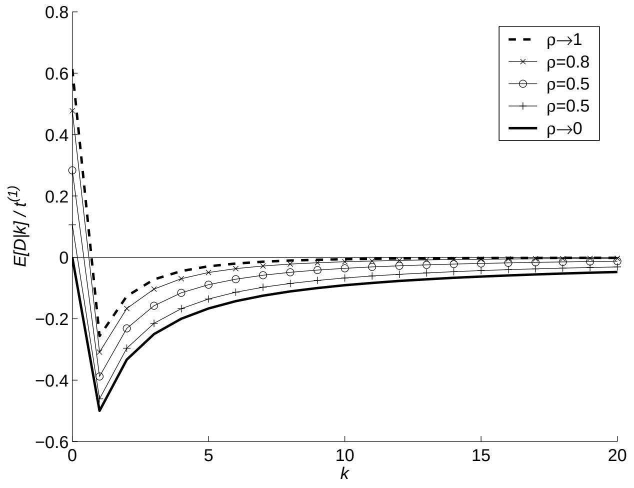

Figure 5: Conditional Discrimination for M/M/1 under ROS

$d(a,s)$:

$$d(a,s) = \begin{cases} t^{(1)}c(a,s) + \tilde{\lambda} d(a+1,0) \\ \quad + \dfrac{\tilde{\mu}}{a} d(a-1,1) + \dfrac{\tilde{\mu}(a-1)}{a} d(a-1,0) & s=0 \\ t^{(1)}c(a,b) + \tilde{\lambda} d(a+1,1) & s=1 \end{cases}. \tag{45}$$

Observe that a customer arrives at the system either at $\mathcal{S}_{0,1}$ when the the system is empty, or at $\mathcal{S}_{k,0}$ when it is not, thus

$$E[D|k] = \begin{cases} d(0,1) & k=0 \\ d(k,0) & k>0 \end{cases}. \tag{46}$$

4.4.1 Special Cases

Again, for $k=0$ all non-preemptive service policies yield the same results, the results derived in Section 4.2.1 regarding $E[D|0]$ hold here as well.

4.4.2 Numerical results and properties

Figure 5 depicts the conditional (normalized by $t^{(1)}$) discrimination $\frac{E[D|k]}{t^{(1)}}$ as a function of k, for some values of ρ. We note the following properties:

1. The discrimination values are significantly lower (both positive and negative) compared to Non-Preemptive LCFS. This demonstrates the fact that ROS violates the natural service order (FCFS) significantly less than Non-Preemptive LCFS.

2. For all load values, the discrimination reaches zero when k is large. Recall that for the Processor Sharing policy (PS) the discrimination value is constantly zero. This means that for large values of k ROS behaves like PS. We conjecture that this is even more so for preemptive variants of ROS.

5. DERIVATION OF SYSTEM UNFAIRNESS IN M/M/1

In this section we derive the variance of the discrimination for the service disciplines under study in an M/M/1 system; this will serve as a measure of system unfairness. As shown above, we have $E[D]=0$ and therefore $Var[D]=E[D^2]$. Conditioning this value on the number of customers seen on arrival, k, yields

$$Var[D] = \sum_{k=0}^{\infty} P_k E[D^2|k], \tag{47}$$

where $P_k = (1-\rho)\rho^k$ is the probability of finding k customers on arrival. The analysis below derives $E[D^2|k]$, thus yielding $Var[D]$.

5.1 FCFS

Similarly to (9)

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

From (10) for customers not in service

$$D(a,b)^2 = \begin{cases} (T_i c(a,b) + D(a,b+1))^2 & \text{with probability } \tilde{\lambda} \\ (T_i c(a,b) + D(a-1,b))^2 & \text{with probability } \tilde{\mu} \end{cases}. \tag{49}$$

Similarly, from (11) for customers in service

$$D(a,b)^2 = \begin{cases} (T_i c(a,b) + D(a,b+1))^2 & \text{with probability } \tilde{\lambda} \\ (T_i c(a,b))^2 & \text{with probability } \tilde{\mu} \end{cases}. \tag{50}$$

We now Combine (49) and (50) and take their expectations, by first expanding the quadratic form. In doing so we take advantage of the fact that $c(a,b)$ is a constant, and that the sum represents a sum of *independent* random variables (see Remark 4.1) and thus the expected value of the product equals to the product of the expected values; for example, $E[2T_i c(a,b) D(a,b{+}1)] = 2c(a,b)E[T_i]E[D(a,b{+}1)]$. All this leads to the following recursive expressions for $d^{(2)}(a,b)$. For $a>0$

$$\begin{aligned} d^{(2)}(a,b) &= t^{(2)}(c(a,b))^2 + \tilde{\lambda} d^{(2)}(a,b+1) + \tilde{\mu} d^{(2)}(a-1,b) \\ &\quad + 2t^{(1)}c(a,b)(\tilde{\lambda} d(a,b+1) + \tilde{\mu} d(a-1,b)), \end{aligned} \tag{51}$$

and for $a=0$

$$\begin{aligned} d^{(2)}(a,b) &= t^{(2)}(c(a,b))^2 + \tilde{\lambda} d^{(2)}(a,b+1) \\ &\quad + 2t^{(1)}c(a,b)\tilde{\lambda} d(a,b+1). \end{aligned} \tag{52}$$

Using the recursive expression for computing $d(a,b)$ given in (12), and substituting $t^{(2)} = 2(t^{(1)})^2$, we can write for $a>0$

$$\begin{aligned} d^{(2)}(a,b) &= 2(t^{(1)})^2(c(a,b))^2 + \tilde{\lambda} d^{(2)}(a,b{+}1) + \tilde{\mu} d^{(2)}(a{-}1,b) \\ &\quad + 2t^{(1)}c(a,b)(d(a,b) - t^{(1)}c(a,b)) \\ &= \tilde{\lambda} d^{(2)}(a,b+1) + \tilde{\mu} d^{(2)}(a-1,b) - 2t^{(1)}c(a,b)d(a,b) \end{aligned} \tag{53}$$

and for $a=0$

$$\begin{aligned} d^{(2)}(a,b) &= 2(t^{(1)})^2(c(a,b))^2 + \tilde{\lambda} d^{(2)}(a,b+1) \\ &\quad + 2t^{(1)}c(a,b)(d(a,b) - t^{(1)}c(a,b)) \\ &= \tilde{\lambda} d^{(2)}(a,b+1) - 2t^{(1)}c(a,b)d(a,b) \end{aligned} \tag{54}$$

REMARK 5.1. *Note that from Remark 4.2, $\frac{E[D^2|k]}{t^{(2)}}$ also depends on λ and μ only through their ratio ρ and thus can be computed from the knowledge of ρ alone. This is true for all the policies studied in this paper.*

5.2 Non-Preemptive LCFS

Using the same arguments as in Section 5.1 leads to the following recursive expression for $d^{(2)}(a,b)$. For $b > 0$

$$d^{(2)}(a,b) = t^{(2)}(c(a,b))^2 + \tilde{\lambda} d^{(2)}(a,b+1) + \tilde{\mu} d^{(2)}(a,b-1) + 2t^{(1)}c(a,b)(\tilde{\lambda} d(a,b+1) + \tilde{\mu} d(a,b-1)) \quad (55)$$

and for $b = 0$

$$d^{(2)}(a,b) = t^{(2)}(c(a,b))^2 + \tilde{\lambda} d^{(2)}(a+1,b) + 2t^{(1)}c(a,b)\tilde{\lambda} d(a+1,b). \quad (56)$$

Similarly to (25)

$$E[D^2|k] = \begin{cases} d^{(2)}(k-1,1) & k > 0 \\ d^{(2)}(0,0) & k = 0 \end{cases}. \quad (57)$$

5.3 Preemptive LCFS

Using the same arguments as in Section 5.1 leads to the following recursive expression for $d^{(2)}(a,b)$. For $b > 0$

$$d^{(2)}(a,b) = t^{(2)}(c(a,b))^2 + \tilde{\lambda} d^{(2)}(a,b+1) + \tilde{\mu} d^{(2)}(a,b-1) + 2t^{(1)}c(a,b)(\tilde{\lambda} d(a,b+1) + \tilde{\mu} d(a,b-1)) \quad (58)$$

and for $b = 0$

$$d^{(2)}(a,b) = t^{(2)}(c(a,b))^2 + \tilde{\lambda} d^{(2)}(a,b+1) + 2t^{(1)}c(a,b)\tilde{\lambda} d(a,b+1). \quad (59)$$

Similarly to (42)

$$E[D^2|k] = d^{(2)}(k,0). \quad (60)$$

5.4 Random Order of Service (ROS)

Using the same arguments as in Section 5.1 leads to the following recursive expression for $d(a,s)^{(2)}$. For $s = 0$

$$d^{(2)}(a,s) = t^{(2)}(c(a,0))^2 + \tilde{\lambda} d^{(2)}(a+1,0) + \frac{\tilde{\mu}}{a} d^{(2)}(a-1,1) + \frac{\tilde{\mu}(a-1)}{a} d^{(2)}(a-1,0) + 2t^{(1)}c(a,b)\left(\tilde{\lambda} d(a+1,0) + \frac{\tilde{\mu}}{a} d(a-1,1) + \frac{\tilde{\mu}(a-1)}{a} d(a-1,0)\right) \quad (61)$$

and for $s = 1$

$$d^{(2)}(a,s) = t^{(2)}(c(a,0))^2 + \tilde{\lambda} d^{(2)}(a+1,0) + 2t^{(1)}c(a,b)\tilde{\lambda} d(a+1,0) \quad (62)$$

Similarly to (46)

$$E[D^2|k] = \begin{cases} d^{(2)}(k,0) & k > 0 \\ d^{(2)}(0,1) & k = 0 \end{cases}. \quad (63)$$

6. NUMERICAL RESULTS AND DISCUSSION

Figure 6 depicts $\frac{Var[D]}{t^{(2)}}$ as a function of ρ for the policies studied in this paper. The figure demonstrates the following properties:

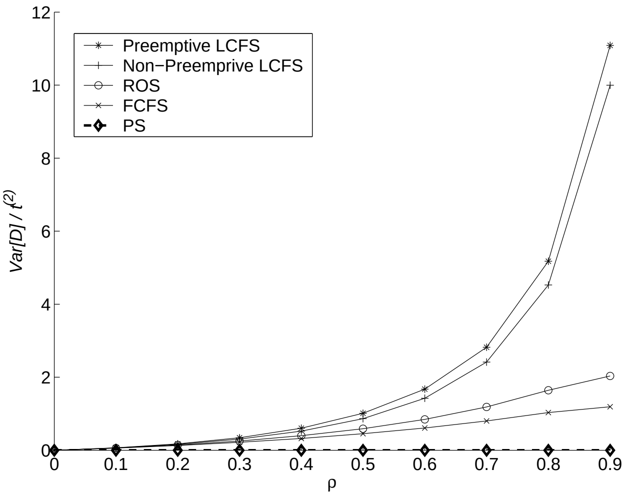

Figure 6: System Unfairness (Variance of discrimination) for M/M/1

1. In terms of unfairness the policies may be ranked as $P - LCFS > NP - LCFS > ROS > FCFS > PS$. That is, Preemptive LCFS is the most unfair and Processor sharing is the most fair. This dominance is observed for all values of ρ, except for $\rho = 0$.

2. At $\rho = 0$ the unfairness measures of all policies converge, as expected, and all policies experience full fairness, ($Var[D] = 0$). It is important to note that this results from the fact that most customers find an empty system upon arrival and thus are subject to no discrimination. Nonetheless if one conditions the discrimination on the system not being empty the discrimination is not zero.

3. Processor sharing is the only policy (of the policies studied) whose fairness is absolute (unfairness = 0).

4. The system unfairness (as measured by $Var[D]$) seems to be monotonically increasing in the system utilization ρ for all policies. For FCFS and ROS the increase is modest, and the values at $\rho = 0.9$ are relatively modest. In fact the slopes of their curves suggest that when $\rho \to 1$ their values may converge to a finite number (in contrast, for example, to the expected delay, which blows up at $\rho = 1$). For the LCFS policies the increase is very drastic at high loads

It is interesting to compare these results with other queueing fairness criteria, in particular the results derived in Wierman and Harchol-Balter [18]. That work considered a large set of scheduling policies and classified them to 1) Always fair, 2) Sometimes fair, and 3) Always unfair. Fairness, in the context of that work is defined by considering the expected normalized system time $E[T(x)/x]$ of jobs whose service time is x and examining whether $E[T(x)/x] \leq 1/(1-\rho)$ for all x. If for certain x $E[T(x)/x] > 1/(1-\rho)$ then jobs of size x are treated unfairly and the policy is unfair.

The policies studied both in that work and in this study are $FCFS$, PS and P-LCFS. For these policies under the above classification both PS and P-LCFS are shown in [18] to be always fair. FCFS, in contrast, is shown to be always unfair.

Our results are in agreement regarding PS (which serves as the departure point for both measures), but differ significantly on FCFS and P-LCFS. The reason seems to be that the $E[T(x)/x]$ classifier deals with the expected performance of a size-x job, averaged over all these jobs over all situations. The classifier emphasis is therefore on service times and it does not account directly for relative seniority. Thus P-LCFS whose seniority treatment is extremely unfair (as it discriminates against old jobs) is classified as fair, while FCFS whose seniority treatment is extremely fair is classified as unfair.

7. CONCLUDING REMARKS

This paper dealt with the issue of fairness in queueing systems, and aimed at deriving a measure that will express unfairness in various systems and that can be used to compare various service policies to each other. Recognizing that both *seniority* and *service requirements* must play significant role in scheduling decisions, we focused on proposing a measure, RAQFM, that accounts for both quantities. RAQFM is constructed by accounting for the individual discrimination attributed to each job in the system and summarizing them. This approach allows to use RAQFM for measuring individual job discrimination, unfairness of scenarios, and unfairness of systems and service policies. Further, the measure allows the use of common queueing theory techniques for evaluating the measure.

We demonstrated the sensitivity of RAQFM to seniority and service requirement by considering a specific case-study scenario. We studied the measure and demonstrated its properties by deriving both the conditional discrimination and the system unfairness for the M/M/1 system under four fundamental service disciplines. Evaluation of special cases of the results suggest that its predictions seem to agree with common intuition about fair and unfair queueing situations.

Our analysis leads to a simple ranking of the service policies examined for the M/M/1 system, independent of the load situation, according to their relative fairness (as evaluated by RAQFM). The ranking, stated from most-fair to least-fair, is PS, FCFS, ROS, NP-LCFS and P-LCFS.

RAQFM can of course be applied to any single server system (e.g. $M/G/1$, $G/G/1$); The question of how to analyze it in a compact form remains open.

Studying of queueing fairness has just began and more research is needed to further understand it. In particular, our ongoing research deals with expanding the analysis of RAQFM to a wider variety of queueing systems and service policies. Good understanding of fairness and proper quantification of it will allow researchers and practitioners to quantitatively account for fairness, in addition to the traditional measure of efficiency, in designing queueing systems and scheduling policies. A comparison of the various measures of fairness in queues is important in order to better understand the subject as well as understand the situations at which each of the measures should be applied.

8. REFERENCES

[1] B. Avi-Itzhak and H. Levy. On measuring fairness in queues. *Advances of Applied Probability*, to appear, 2004.

[2] E. G. Coffman, Jr., R. R. Muntz, and H. Trotter. Waiting time distribution for processor-sharing systems. *J. ACM*, 17:123–130, 1970.

[3] A. Demers, S. Keshav, and S. Shenker. Analysis and simulation of a fair queueing algorithm. *Internetworking Research and Experience*, 1:3–26, 1990.

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

[5] M. Harchol-Balter, B. Schroeder, N. Bansal, and M. Agrawal. Size-based scheduling to improve web performance. *ACM Transactions on Computer Systems*, 21(2):207–233, May 2003.

[6] L. Kleinrock. Analysis of a time-shared processor. *Nav. Res. Log. Quarterly*, 11:59–73, 1964.

[7] L. Kleinrock. Time-shared systems: A theoretical treatment. *J. ACM*, 14:242–261, 1967.

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

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

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

[11] A. Parekh and R. G. Gallager. A generalized processor sharing approach to flow control in integrated services networks: The single node case. *IEEE/ACM Trans. Networking*, 1:344–357, June 1993.

[12] A. Parekh and R. G. Gallager. A generalized processor sharing approach to flow control in integrated services networks: The multiple node case. *IEEE/ACM Trans. Networking*, 2:137–150, 1994.

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

[14] 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.

[15] D. Raz, B. Avi-Itzhak, and H. Levy. Classes, priorities and fairness in queueing systems. Forthcoming, 2004.

[16] Y. T. Wang and R. J. T. Morris. Load sharing in distributed systems. *IEEE Trans. on computers*, C-34(3):204–217, 1985.

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

[18] 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*, San Diego, CA, June 2003.

APPENDIX

A. PROOF OF THEOREM 4.1

PROOF. Let L be a random variable denoting the walk a tagged customer C walks in the Markovian chain defined in Section 4.1. Let $\mathbb{L}$ be the field of all possible walks. Let l be an instance of L. Let $G(a, b)$ be a binary random variable denoting whether C encounters $\mathcal{S}_{a,b}$. The expected value of the discrimination, given l is:

$$E[D|L=l] = t^{(1)} \sum_{a=0}^{\infty}\sum_{b=0}^{\infty} E[G(a,b)|L=l]c(a,b). \quad (64)$$

Let $P[L=l|k]$ be the probability that $L=l$, given that k customers were seen on arrival. Using conditional distribution

$$E[D|k] = \sum_{l\in\mathbb{L}} E[D|k, L=l]P[L=l|k], \quad (65)$$

where $E[D|k, L=l]$ is the expected value of the discrimination, given l and k, and is defined only for path l that starts with k customers. Note that $L=l$ determines k and thus $E[D|k, L=l] = E[D|L=l]$, and (65) becomes

$$\begin{aligned}E[D|k] &= \sum_{l\in\mathbb{L}} E[D|L=l]P[L=l|k]\\ &= t^{(1)} \sum_{L\in\mathbb{L}}\sum_{a=0}^{\infty}\sum_{b=0}^{\infty} E[G(a,b)|L=l]c(a,b)P[L=l|k]\\ &= t^{(1)} \sum_{a=0}^{\infty}\sum_{b=0}^{\infty} c(a,b)\sum_{l\in\mathbb{L}} E[G(a,b)|L=l]P[L=l|k], \quad (66)\end{aligned}$$

where the second equality is due to (64) and the third one is due to changing the order of summation.

Using conditional distribution

$$\begin{aligned}E[G(a,b)|k] &= \sum_{l\in\mathbb{L}} E[G(a,b)|k, L=l]P[L=l|k]\\ &= \sum_{l\in\mathbb{L}} E[G(a,b)|L=l]P[L=l|k], \quad (67)\end{aligned}$$

where the second equality is due to the fact that l determines k. Substituting this into (66) yields

$$E[D|k] = t^{(1)} \sum_{a=0}^{\infty}\sum_{b=0}^{\infty} E[G(a,b)|k]c(a,b). \quad (68)$$

Observe that $E[G(a,b)|k] = 0$ for $a > k$ and thus

$$E[D|k] = t^{(1)} \sum_{b=0}^{\infty}\sum_{a=0}^{k} E[G((a,b)|k]c(a,b). \quad (69)$$

We now evaluate $E[G(a,b)|k]$, the probability that a walk starting at $\mathcal{S}_{k,0}$ will encounter $\mathcal{S}_{a,b}$. One may note that a walk will do this if and only if the first $k-a+b$ events in the walk consist of exactly b arrival events and $k-a$ departure events. The probability of such a walk is given by the binomial distribution:

$$\begin{aligned}E[G(a,b)|k] &= \tilde{\lambda}^b \tilde{\mu}^{k-a}\binom{k-a+b}{b}\\ &= \frac{\rho^b}{(1+\rho)^{k+b-a}}\binom{k-a+b}{b}. \quad (70)\end{aligned}$$

The second equality is due to (6). □

B. PROOF OF THEOREM 4.2

PROOF. The correctness of (31) for $n=1$ is immediate from (25) for $k>0$.

Next we prove (31) for $n=2$. Define $d^{(s)} \overset{def}{=} d(k-1,0)$. From (24)

$$E[D|k] = d(k-1,1) = t^{(1)}c(k-1,1) + \tilde{\lambda}d(k-1,2) + \tilde{\mu}d^{(s)}, \quad (71)$$

To prove the correctness of the expression given in (30) note that $d(k-1,0)$ can be evaluated according to (24) as follows

$$\begin{aligned}d(k-1,0) &= t^{(1)}c(k-1,0) + \tilde{\lambda}d(k,0)\\ &= t^{(1)}c(k-1,0) + \tilde{\lambda}(t^{(1)}c(k,0) + \tilde{\lambda}d(k+1,0)) = \ldots\\ &= t^{(1)}\sum_{a=k}^{\infty}\tilde{\lambda}^{a-k}c(a-1,0) = t^{(1)}\sum_{a=k}^{\infty}(\frac{\rho}{1+\rho})^{a-k}(1-\frac{1}{a})\\ &= 1+\rho-\sum_{a=0}^{\infty}(\frac{\rho}{1+\rho})^a\frac{1}{a+k}, \quad (72)\end{aligned}$$

leading immediately to (30)[5]. Thus the theorem is proved for $n=1,2$.

Next we prove the theorem for $n\geq 3$ by induction.

Induction base: correctness for $n=1,2$ was prover above.

Induction assumption: $E[D|k] = a_m + b_m d(k-1,m)$ for $m=1,2,\ldots,n-2,n-1$. For brevity denote $c_n = c(n,k-1)$ and $d_n = d(n,k-1)$. Note that using this notation the assumption can be written as follows

$$d_1 = a_m + b_m d_m \quad (73)$$

Induction step: we now prove the correctness for $m=n$. From the assumption for $m=n-2$

$$d_1 = a_{n-2} + b_{n-2}d_{n-2} \Rightarrow d_{n-2} = \frac{d_1 - a_{n-2}}{b_{n-2}}. \quad (74)$$

Note that using the abbreviated notation (24) for $a=k-1$ can be written as follows

$$d_n = t^{(1)}c_n + \tilde{\lambda}d_{n+1} + \tilde{\mu}d_{n-1}. \quad (75)$$

Evaluating this at $n-1$ yields

$$d_{n-1} = t^{(1)}c_{n-1} + \tilde{\lambda}d_n + \tilde{\mu}d_{n-2}. \quad (76)$$

From the assumption for $m=n-1$

$$d_1 = a_{n-1} + b_{n-1}d_{n-1}. \quad (77)$$

Substituting d_{n-1} from (76) yields

$$d_1 = a_{n-1} + b_{n-1}(t^{(1)}c_{n-1} + \tilde{\lambda}d_n + \tilde{\mu}d_{n-2}). \quad (78)$$

Substituting d_{n-2} from (74) yields

$$d_1 = a_{n-1} + b_{n-1}(t^{(1)}c_{n-1} + \tilde{\lambda}d_n + \tilde{\mu}\frac{d_1 - a_{n-2}}{b_{n-2}}). \quad (79)$$

This leads to

$$d_1(1-\frac{\tilde{\mu}b_{n-1}}{b_{n-2}}) = a_{n-1} + t^{(1)}b_{n-1}c_{n-1} - \tilde{\mu}b_{n-1}\frac{a_{n-2}}{b_{n-2}} + \tilde{\lambda}b_{n-1}d_n, \quad (80)$$

and then

$$d_1 = \frac{a_{n-1} + t^{(1)}b_{n-1}c_{n-1} - \tilde{\mu}b_{n-1}\frac{a_{n-2}}{b_{n-2}}}{1-\frac{\tilde{\mu}b_{n-1}}{b_{n-2}}} + \frac{\tilde{\lambda}b_{n-1}}{1-\frac{\tilde{\mu}b_{n-1}}{b_{n-2}}}d_n, \quad (81)$$

and by comparison to (31) the theorem is proved. □

[5] $d^{(s)}$ can also be written as $1+\rho-\frac{{}_2F_1(k,1;k+1,\frac{r}{1+r})}{k}$ where ${}_2F_1(a,b;c;x)$ is Gauss's Hyper-geometric Function which might yield to even further reduction