Fairness Through Linearly Increasing Holding Costs in Systems of Parallel Servers with Setup Delays

Esa Hyytiä
Department of Communications and Networking
Aalto University, Finland

Rhonda Righter
Department of Industrial Engineering and Operations Research
University of California Berkeley

Abstract—We consider a system of parallel servers, where arriving jobs are routed to one of the servers upon arrival. The standard objective of minimizing the mean sojourn time (delay) does not enforce any kind of fairness in the system and it is acceptable, e.g., to delay one job a lot if it reduces the sojourn time of some other jobs. We take fairness into account by defining a linearly increasing instantaneous holding cost rate by two job-specific non-negative random variables, (α_i, β_i), that can depend on the service time. We focus on first-come-first-served (FCFS) and preemptive last-come-first-served (LCFS) scheduling, and derive the so-called value functions for the corresponding M/G/1 queues. Then we apply these results and obtain cost-aware dispatching policies by means of policy improvement and lookahead. The policies are finally evaluated numerically.

I. INTRODUCTION

We consider a system of parallel servers, where arriving jobs are routed to one of the servers upon arrival. This system is known as a dispatching system and it arises in many contexts such us manufacturing and transportation systems, as well as modern ICT systems such as web server farms. Optionally, servers can be switched off when they become idle in order to save energy. This, induces a setup delay when a server is switched back on in response to an arriving job (see [1], [2]).

The standard objective is to minimize the mean sojourn time (delay) or its weighted version, where jobs have individual holding costs. However, this objective does not enforce any kind of fairness and it is acceptable, e.g., to delay one job a lot if it reduces the sojourn time of some other jobs. However, longer waiting time is likely to increase the "anxiety level" of a queueing customer. In this paper, we take fairness into account by defining a linearly increasing holding cost model, where the holding cost of job i is defined by two non-negative parameters (α_i, β_i), giving the instantaneous holding cost rate

$$\alpha_i + \beta_i t,$$

where t denotes the time the job has already spent in the system. With $\alpha_i = 1$ and $\beta_i = 0$ one obtains the standard objective of minimizing the mean delay. The holding cost structure is illustrated in Fig. 1, where upper lines correspond to case with $\alpha = 1$ and $\beta = 0.35$, and the lower to the standard sojourn time cost model. The left figure illustrates the holding cost rate (per unit time), and the right figure corresponds to the total cost incurred. Our cost structure is general in the sense that (α_i, β_i) are random variables that can depend on the service requirement.
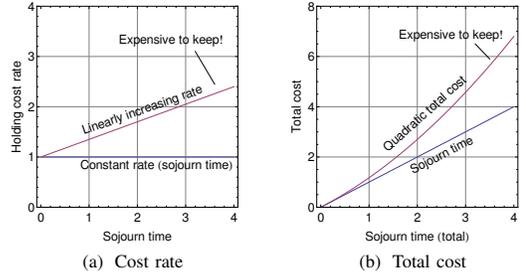

Fig. 1. A linear holding cost model promotes fairness by increasing the cost rate as a function of sojourn time.

We focus on the two standard scheduling disciplines, first-come-first-served (FCFS) and preemptive last-come-first-served (LCFS). FCFS is the best nonpreemptive discipline (for convex identical costs) [3], and also optimal when service times are unknown, but have increasing failure rate. Moreover, deviating from FCFS might cause social injustice [4].

The main contributions of this paper are: (i) closed-form expressions for the mean cost according to the general cost structure for M/G/1-FCFS and LCFS queues subject to random i.i.d. setup delays; (ii) closed-form expressions for the corresponding size-aware value functions yielding efficient cost- and state-aware dispatching policies, (iii) a comparison of M/G/1-FCFS with M/G/1-LCFS, including the relative impacts of variability on costs. As a special case one obtains the mean squared sojourn time in M/G/1-FCFS and LCFS queues. Additionally, several interesting theoretical observations are discussed separately. For example, we find that, although for a single queue LCFS can be better than FCFS when service times are highly variable, a good dispatching policy can mitigate this effect.

II. PRELIMINARIES

Each job i is characterized by its size x_i and the holding cost rate parameters (α_i, β_i). The constant component of the cost function corresponds to the weighted sojourn time T,

$$\int_0^T \alpha_i \, dt = \alpha_i T =: a_i T,$$

and the linear component to the weighted squared sojourn time

$$\int_0^T \beta_i t \, dt = \frac{\beta_i}{2} T^2 =: b_i T^2.$$

1

That is, our cost function is equivalent to assuming that the jobs pay a total cost which is a function of the sojourn time T and its square T^2,

$$C_i = a_i T_i + b_i T_i^2, \qquad \text{where } \begin{cases} a_i = \alpha_i, \\ b_i = \beta_i/2, \end{cases}$$

The parameters of jobs arriving in future are i.i.d. random variables, $(X_i, A_i, B_i) \sim (X, A, B)$, where A_i and B_i may depend on X_i. This enables performance metrics such as *slowdown*, where $A_i = 1/X_i$ and $B_i = 0$. We assume that (x_i, a_i, b_i) becomes known upon the arrival. For FCFS, the payment can occur upon arrival to the queue, whereas with LCFS the payment must be delayed until departure.

A. Setup Delay

Many results are easier to derive first for a fixed setup delay s, and then to generalize to the case where the setup delays are i.i.d. random variables $S_i \sim S$. Let Z denote an arbitrary performance measure (e.g., T^2). For M/G/1 with setup delays,

$$\mathbb{E}[Z] = \int \mathbb{E}[Z|s]\, \pi(s)\, dF(s), \tag{1}$$

where $\mathbb{E}[Z|s]$ is the expected value of Z given a fixed setup delay s, and (see [5])

$$\pi(s) = \frac{1 + \lambda s}{1 + \lambda \mathbb{E}[S]}. $$

B. Value functions

The central notion for developing good dispatching heuristics is the value function for each queue in isolation,

$$v(\mathbf{z}) := \lim_{t \to \infty} \mathbb{E}[V_z(t) - r \cdot t], \tag{2}$$

where $V_z(t)$ denotes the costs incurred during $(0, t)$ when initially in state z, and r is the mean cost rate. We have implicitly assumed a stable and ergodic system ($\lambda \mathbb{E}[X] < 1$). Moreover, we consider size-aware state information, where the (remaining) service times of all jobs, as well as the (remaining) setup delay, are available. The important quantity for dispatching is actually the difference, $v(\mathbf{z}') - v(\mathbf{z})$, which describes how much more or less it costs to start an ergodic system from state $\mathbf{z}'$ instead of state $\mathbf{z}$. Finally, it is easy to see that the value function for our cost structure decomposes,

$$v(\mathbf{z}) = v_A(\mathbf{z}) + v_B(\mathbf{z}), \tag{3}$$

where $v_A(\mathbf{z})$, corresponding to AT, depends on A, and $v_B(\mathbf{z})$, corresponding to BT^2, on B.

III. Analysis of M/G/1-FCFS

The cost structure with respect to waiting and sojourn time, including their squares, in an M/G/1-FCFS has been already discussed in [6]. Here we first restate some of the main results, and then derive an explicit expression for the value function with respect to T^2. Finally, we utilize (3) and obtain the value function with respect to our cost rate model.

A. Mean costs

In the general case, the coefficients (A, B) for T and T^2 in the cost function may depend on the service time of the job, X. Recalling that with FCFS, $T = W + X$, we first assume a general cost structure where the cost is a random polynomial function of kth degree,

$$C = \sum_{i=0}^{k} C_i T^i = \sum_{i=0}^{k} C_i (W + X)^i,$$

where C_i denotes the random coefficient that may depend on X, but not on the waiting time W. The above reduces to

$$C = \sum_{i=0}^{k} C_i \sum_{j=0}^{i} \binom{i}{j} W^j X^{i-j} = \sum_{i=0}^{k} D_i W^i,$$

where $D_i = \sum_{j=i}^{k} \binom{j}{i} C_j X^{j-i}$. The D_i are i.i.d. random variables that depend on X but not on W (or the scheduling discipline). Therefore, the cost of a job in the general case is

$$\mathbb{E}[C] = \sum_{i=0}^{k} \mathbb{E}[D_i]\, \mathbb{E}[W^i].$$

With our cost structure, $C_0 = 0$, $C_1 = A$, and $C_2 = B$, which gives $D_0 = AX + BX^2$, $D_1 = A + 2BX$, and $D_2 = B$. The cost of a job is

$$C = AX + BX^2 + (A + 2BX)W + BW^2. \tag{4}$$

Therefore, the mean cost of a job can be written as

$$\mathbb{E}[C] = \mathbb{E}[AT] + \mathbb{E}[BT^2], \tag{5}$$

where

$$\mathbb{E}[AT] = \mathbb{E}[AX] + \mathbb{E}[A]\mathbb{E}[W],$$
$$\mathbb{E}[BT^2] = \mathbb{E}[BX^2] + 2\mathbb{E}[BX]\mathbb{E}[W] + \mathbb{E}[B]\mathbb{E}[W^2].$$

The mean waiting time and the mean squared waiting time in an M/G/1-FCFS with a fixed setup time s are given, e.g., in [6]. Those results can be generalized to a system where setup times are arbitrary i.i.d. random variables $S_i \sim S$, using (1):

Lemma 1: The mean waiting time and the mean squared waiting time in an M/G/1-FCFS with setup time S are

$$\mathbb{E}[W] = \frac{\lambda \mathbb{E}[X^2]}{2(1 - \rho)} + \frac{2\mathbb{E}[S] + \lambda \mathbb{E}[S^2]}{2(1 + \lambda \mathbb{E}[S])}, \tag{6}$$

$$\mathbb{E}[W^2] = \frac{3\lambda^2 \mathbb{E}[X^2]^2 + 2\lambda(1 - \rho)\mathbb{E}[X^3]}{6(1 - \rho)^2} + \frac{\lambda(2\mathbb{E}[S] + \lambda \mathbb{E}[S^2])\mathbb{E}[X^2]}{2(1 - \rho)(1 + \lambda \mathbb{E}[S])} + \frac{3\mathbb{E}[S^2] + \lambda \mathbb{E}[S^3]}{3(1 + \lambda \mathbb{E}[S])}. \tag{7}$$

We note that (6) and (7) can be also obtained, e.g., from [7, Theorem 6]. Now, from (5), (6) and (7), we can give explicit expressions for the mean cost $\mathbb{E}[C]$:

Theorem 2: The mean costs of a job in an M/G/1-FCFS queue when (A, B) may depend on X are

$$\mathbb{E}[AT] = \mathbb{E}[AX] + \mathbb{E}[A]\left(\frac{\lambda \mathbb{E}[X^2]}{2(1 - \rho)} + \frac{2\mathbb{E}[S] + \lambda \mathbb{E}[S^2]}{2(1 + \lambda \mathbb{E}[S])}\right), \tag{8}$$

2

$$\boxed{\begin{aligned}
\mathbb{E}[BT^2] &= \mathbb{E}[BX^2]+ \\
&\mathbb{E}[BX]\left(\frac{\lambda\,\mathbb{E}[X^2]}{1-\rho} + \frac{2\,\mathbb{E}[S] + \lambda\,\mathbb{E}[S^2]}{1+\lambda\,\mathbb{E}[S]}\right) \\
&+\mathbb{E}[B]\Bigg[\frac{3\lambda^2\,\mathbb{E}[X^2]^2 + 2\lambda(1-\rho)\,\mathbb{E}[X^3]}{6(1-\rho)^2} + \\
&\frac{\lambda(2\,\mathbb{E}[S] + \lambda\,\mathbb{E}[S^2])\mathbb{E}[X^2]}{2(1-\rho)(1+\lambda\,\mathbb{E}[S])} + \frac{3\,\mathbb{E}[S^2] + \lambda\,\mathbb{E}[S^3]}{3(1+\lambda\,\mathbb{E}[S])}\Bigg].
\end{aligned}} \quad (9)$$

By setting $A = 1$ in (8) we get that $\mathbb{E}[T]$ depends on the first two moments of X and S, and by setting $B = 1$, we get that $\mathbb{E}[T^2]$ depends on the first three moments of X and S,

$$\begin{aligned}
\mathbb{E}[T^2] &= \mathbb{E}[X^2] + \mathbb{E}[X]\left(\frac{\lambda\,\mathbb{E}[X^2]}{1-\rho} + \frac{2\,\mathbb{E}[S] + \lambda\,\mathbb{E}[S^2]}{1+\lambda\,\mathbb{E}[S]}\right) \\
&+ \frac{3\lambda^2\,\mathbb{E}[X^2]^2 + 2\lambda(1-\rho)\,\mathbb{E}[X^3]}{6(1-\rho)^2} \\
&+ \frac{\lambda(2\,\mathbb{E}[S] + \lambda\,\mathbb{E}[S^2])\mathbb{E}[X^2]}{2(1-\rho)(1+\lambda\,\mathbb{E}[S])} + \frac{3\,\mathbb{E}[S^2] + \lambda\,\mathbb{E}[S^3]}{3(1+\lambda\,\mathbb{E}[S])}.
\end{aligned} \quad (10)$$

B. Value functions

Next we consider the effect of different starting states of the system. More specifically, we derive the value function for an M/G/1-FCFS queue with a setup delay with respect to the general cost structure. With FCFS, we adopt the convention that the costs are incurred upon arrival. This is feasible as the current state of the system defines the waiting and sojourn times. Consequently, the value functions characterize the state of the system for the jobs arriving in the future. This means that a sufficient state description is the virtual backlog u that includes the possible remaining setup delay δ.

The explicit expression for the value functions with respect to waiting time for systems with a fixed setup time s are already available from [6],

$$v_W(u) - v_W(0) = \frac{\lambda\,u}{2(1-\rho)}\left(u - \frac{s(2+\lambda s)}{1+\lambda s}\right), \quad (11)$$

$$\begin{aligned}
v_{W2}(u)-v_{W2}(0) &= \frac{\lambda u}{1-\rho}\Bigg[\frac{u^2}{3} - \frac{s^2(3+\lambda s)}{3(1+\lambda s)} \\
&+ \frac{\lambda\,\mathbb{E}[X^2]}{2(1-\rho)}\left(u - \frac{s(2+\lambda s)}{1+\lambda s}\right)\Bigg].
\end{aligned} \quad (12)$$

When setup times are arbitrary i.i.d. random variables we have

Lemma 3:

$$v_W(u) - v_W(0) = \frac{\lambda\,u}{2(1-\rho)}\left[u - \frac{2\,\mathbb{E}[S] + \lambda\,\mathbb{E}[S^2]}{1+\lambda\,\mathbb{E}[S]}\right], \quad (13)$$

and

$$\begin{aligned}
v_{W2}(u)-v_{W2}(0) &= \frac{\lambda u}{1-\rho}\Bigg[\frac{u^2}{3} - \frac{3\,\mathbb{E}[S^2] + \lambda\,\mathbb{E}[S^3]}{3(1+\lambda\,\mathbb{E}[S])} \\
&+ \frac{\lambda\,\mathbb{E}[X^2]}{2(1-\rho)}\left(u - \frac{2\,\mathbb{E}[S] + \lambda\,\mathbb{E}[S^2]}{1+\lambda\,\mathbb{E}[S]}\right)\Bigg].
\end{aligned} \quad (14)$$

Proof: In general, when the value function depends only on the virtual backlog u, we have

$$v(u) = \mathbb{E}[V(u)] - r\,\mathbb{E}[\tilde{B}(u)] + v(0),$$

where $V(u)$ denotes the mean cost the system incurs during the remaining busy period, $\mathbb{E}[\tilde{B}(u)]$ is the mean length of the remaining busy period,

$$\mathbb{E}[\tilde{B}(u)] = \frac{u}{1-\rho},$$

and r is the mean cost rate. We refer to the system with fixed setup delay s with superscript s, and superscript S corresponds to a system with i.i.d. random setup delays $S_i \sim S$. Clearly, $V^{(s)}(u) = V^{(S)}(u)$ as the setup delay does not affect the costs incurred during the *remaining* busy period. Therefore, $v^{(S)}(u) - v^{(S)}(0) = v^{(s)}(u) - v^{(s)}(0) + (r^{(S)} - r^{(s)})u/(1-\rho)$. Using (11) for $v_W^{(s)}(u) - v_W^{(s)}(0)$ ((12) for $v_{W2}^{(s)}(u) - v_{W2}^{(s)}(0)$) and (6) for $r_W^{(S)} - r_W^{(s)}$ ((7) for $r_{W2}^{(S)} - r_{W2}^{(s)}$) completes the proof. ∎

The value function for an M/G/1-FCFS queue with respect to sojourn time is the same as with respect to waiting time, $v_T(u) = v_W(u)$, given in (13). For the squared sojourn time, there is a small change to $v_{W2}(u)$ given in (14):

Lemma 4: The relative value function for an M/G/1-FCFS queue with respect to the squared sojourn time T^2 is

$$\boxed{\begin{aligned}
v_{T2}(u) - v_{T2}(0) &= \frac{\lambda u}{1-\rho}\Bigg[\frac{u^2}{3} - \frac{3\,\mathbb{E}[S^2] + \lambda\,\mathbb{E}[S^3]}{3(1+\lambda\,\mathbb{E}[S])} \\
&+ \left(\frac{\lambda\mathbb{E}[X^2]}{2(1-\rho)} + \mathbb{E}[X]\right)\left(u - \frac{2\,\mathbb{E}[S] + \lambda\,\mathbb{E}[S^2]}{1+\lambda\,\mathbb{E}[S]}\right)\Bigg].
\end{aligned}} \quad (15)$$

Proof: As in [6], we can consider the remaining busy period for a queue initially in state u, where u is the virtual backlog, that includes any remaining setup time. By definition,

$$v_{T2}(u) = \mathbb{E}[T_1^2 + \ldots + T_{N_u}^2] - r_{T2}\,\mathbb{E}[\tilde{B}(u)] + v_{T2}(0),$$

where N_u is the number of jobs arriving during the remaining busy period, the T_i are their sojourn times, and $\tilde{B}(u)$ is the length of the remaining busy period. As $T_i = W_i + X_i$,

$$\begin{aligned}
v_{T2}(u) - v_{T2}(0) &= \mathbb{E}[W_1^2 + \ldots + W_{N_u}^2] \\
&+ 2\,\mathbb{E}[W_1 X_1 + \ldots + W_{N_u} X_{N_u}] \\
&+ \mathbb{E}[X_1^2 + \ldots + X_{N_u}^2] - r_{T2}\,\mathbb{E}[\tilde{B}(u)].
\end{aligned} \quad (16)$$

The mean cost rate for the squared sojourn time is

$$r_{T2} = \lambda\,\mathbb{E}[T^2] = \lambda\,\mathbb{E}[W^2] + 2\rho\,\mathbb{E}[W] + \lambda\,\mathbb{E}[X^2].$$

Substituting this into (16) gives

$$\begin{aligned}
v_{T2}(u) - v_{T2}(0) &= \mathbb{E}[W_1^2 + \ldots + W_{N_u}^2] - \lambda\,\mathbb{E}[W^2]\,\mathbb{E}[\tilde{B}(u)] \\
&+ 2\,\mathbb{E}[W_1 X_1 + \ldots + W_{N_u} X_{N_u}] + \mathbb{E}[X_1^2 + \ldots + X_{N_u}^2] \\
&- (2\rho\,\mathbb{E}[W] + \lambda\,\mathbb{E}[X^2])\mathbb{E}[\tilde{B}(u)],
\end{aligned} \quad (17)$$

which, because of the FCFS discipline, reduces to

$$\begin{aligned}
v_{T2}(u) - v_{T2}(0) &= v_{W2}(u) - v_{W2}(0) \\
&+ 2\,\mathbb{E}[W_1 + \ldots + W_{N_u}]\mathbb{E}[X] + \mathbb{E}[N_u]\,\mathbb{E}[X^2] \\
&- (2\rho\,\mathbb{E}[W] + \lambda\,\mathbb{E}[X^2])\frac{u}{1-\rho}.
\end{aligned} \quad (18)$$

For $\mathbb{E}[W_1 + \ldots + W_{N_u}]$ we have

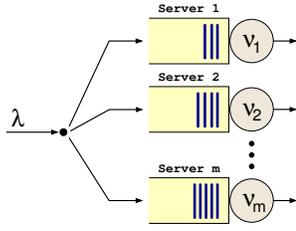$$v_W(u) - v_W(0) = \mathbb{E}[W_1 + \ldots + W_{N_u}] - \lambda\,\mathbb{E}[W]\frac{u}{1-\rho}.$$

Fig. 2. System of m parallel heterogeneous servers.

From (13), we obtain

$$\mathbb{E}[W_1 + \ldots + W_{N_u}] = \frac{\lambda}{2(1-\rho)}\left(u^2 + \frac{\lambda\,\mathbb{E}[X^2]}{1-\rho}u\right), \quad (19)$$

which is independent of the setup delay[1] S. Then,

$$\mathbb{E}[N_u] = \lambda\,\mathbb{E}[\tilde{B}(u)] = \frac{\lambda u}{1-\rho}. \quad (20)$$

Substituting (19) and (20) into (18) gives

$$v_{T2}(u) - v_{T2}(0) = v_{W2}(u) - v_{W2}(0) + \frac{\rho}{1-\rho}\left(u^2 + \frac{\lambda\,\mathbb{E}[X^2]}{1-\rho}u\right)$$
$$+ \frac{\lambda\mathbb{E}[X^2]}{1-\rho}u - (2\rho\,\mathbb{E}[W] + \lambda\,\mathbb{E}[X^2])\frac{u}{1-\rho},$$

which gives that $v_{T2}(u) - v_{T2}(0)$ is equal to

$$v_{W2}(u) - v_{W2}(0) + \frac{\rho}{1-\rho}\left(u^2 + \frac{\lambda\,\mathbb{E}[X^2]}{1-\rho}u - 2\,\mathbb{E}[W]\,u\right).$$

Using (6) for $\mathbb{E}[W]$ gives that $v_{T2}(u) - v_{T2}(0)$ is equal to

$$v_{W2}(u) - v_{W2}(0) + \frac{\rho}{1-\rho}\left(u^2 - \frac{2\,\mathbb{E}[S] + \lambda\,\mathbb{E}[S^2]}{1 + \lambda\,\mathbb{E}[S]}u\right). \quad (21)$$

Combining (14) and (21) gives (15). ∎

Now we suppose that we have random coefficients (A, B) that may depend on the service requirement X. Again, the value function can be derived:

Theorem 5: The value function of an M/G/1-FCFS queue with respect to total cost is given by

$$v(u) - v(0) = \frac{\lambda u}{2(1-\rho)}\left[h_1\left(u^2 - \frac{3\,\mathbb{E}[S^2] + \lambda\,\mathbb{E}[S^3]}{3(1 + \lambda\,\mathbb{E}[S])}\right) + h_2\left(u - \frac{2\,\mathbb{E}[S] + \lambda\,\mathbb{E}[S^2]}{1 + \lambda\,\mathbb{E}[S]}\right)\right], \quad (22)$$

where (A, B) may depend on X, and

$$h_1 = \frac{2\,\mathbb{E}[B]}{3}, \quad h_2 = \mathbb{E}[A] + 2\,\mathbb{E}[BX] + \frac{\lambda\mathbb{E}[X^2]\mathbb{E}[B]}{1-\rho}.$$

Proof: From (4) we deduce that

$$v(u) = \mathbb{E}[A + 2BX]\,v_W(u) + \mathbb{E}[B]\,v_{W2}(u).$$

Substituting (13) and (14) into above gives

$$\mathbb{E}[A + 2BX]\left[\frac{\lambda u}{2(1-\rho)}\left(u - \frac{2\,\mathbb{E}[S] + \lambda\,\mathbb{E}[S^2]}{1 + \lambda\,\mathbb{E}[S]}\right)\right] +$$
$$\frac{\lambda u\mathbb{E}[B]}{1-\rho}\left[\frac{u^2}{3} + \frac{\lambda\,\mathbb{E}[X^2]}{2(1-\rho)}\left(u - \frac{2\mathbb{E}[S] + \lambda\mathbb{E}[S^2]}{1 + \lambda\,\mathbb{E}[S]}\right)\frac{3\mathbb{E}[S^2] + \lambda\mathbb{E}[S^3]}{3(1 + \lambda\,\mathbb{E}[S])}\right]$$

which yields (22). ∎

[1] The virtual backlog u includes the possible remaining setup time. That is the only thing that matters to the jobs arriving during *the current busy period*.

System parameters:		According to the basic policy:	
m	number of servers	λ_i	arrival rate to server i
ν_i	service rate of server i	X_i	job size distribution
s_i	setup delay of server i		to server i
λ	total job arrival rate		
X	job size distribution		

C. Dispatching and policy iteration

Let us next consider the complete system of parallel servers illustrated in Fig. 2. The system parameters are given in Table I. The primary application of the value functions is to develop a dispatching heuristic based on policy improvement, i.e., we can carry out the first policy iteration (FPI) step. We start with an arbitrary static policy α_0, referred to as the basic policy. A simple example is the Bernoulli split (RND), which chooses the server for each job at random with some probabilities $p_1, \ldots, p_m$. With a static policy, the arrival process to each server i is (independent) Poisson process with some rate λ_i and i.i.d. service times X_i, and the system *decomposes* to m independent M/G/1 queues. Hence, the value function is the sum of the server-specific value functions,

$$v(\mathbf{z}) = \sum_i v^{(i)}(u_i),$$

where $\mathbf{z} = (u_1, \ldots, u_m)$ denotes the state of system with u_i denoting the backlog in server i. Given the value function, the policy iteration step can be carried out. We compute the cost of assigning the new job to queue i, and then choose the queue with the smallest expected cost. Formally, let

$$u_i \oplus x = \begin{cases} u_i + x/\nu_i, & \text{if } u_i > 0, \\ s_i + x/\nu_i, & \text{if } u_i = 0, \end{cases}$$

i.e., $u_i \oplus x$ denotes the backlog in queue i once a job with size x has been added (including a possible setup period). Then,

$$\alpha_{\text{FPI}}(\mathbf{z}, x) = \arg\min_i a_i(u_i, x),$$

where $a_i(u_i, x)$ denotes the so-called *admission cost*,

$$a_i(u_i, x) = c_i(u_i, x) + v_i(u_i \oplus x) - v_i(u_i).$$

When minimizing the costs according to T^2,

$$a_i(u_i, x) = (u_i \oplus x)^2 + v_{T2}^{(i)}(u_i \oplus x) - v_{T2}^{(i)}(u_i).$$

Note that the server-specific value functions $v_{T2}^{(i)}(u)$ depend on the static basic policy α_0 through the corresponding arrival process (λ_i, X_i), see Table I.

Given the value function with respect to the total costs $AT + BT^2$, we can similarly obtain the corresponding FPI policy.

D. Lookahead

Lookahead is a recent idea to close the gap between FPI and the optimal policy determined by repeated policy iteration [8]. The basic idea is to consider jointly the assignment of the present and the next job. To this end, it is convenient to define the so-called Lookahead value:

Definition 1: $L(\mathbf{z}, \lambda^*, X^*)$ is the expected difference in cumulative cost between the queue in equilibrium and a queue initially in state $\mathbf{z}$ receiving the next job X^* at Poisson rate λ^* after which the arrival process based on the basic policy resumes. The convention is to denote the event that a job is assigned elsewhere by an empty job, $X^* = \emptyset$.

The lookahead policy (LH) considers actions (i_0, i_1), where i_0 denotes the server of the current job, and i_1 the server where the next job is (tentatively) assigned. As with FPI, LH chooses the server with the smallest cost,

$$\alpha_{\text{LH}}(\mathbf{z}, x) = \arg\min_{i_0} \left(\min_{i_1} a_{i_0, i_1} \right). \tag{23}$$

We can proceed similarly as with FPI, and choosing the action to drop this and the next job as the reference, the admission cost for action (i, i) is

$$
\begin{aligned}
a_{i,i}(\mathbf{z}, x) = c_i(u_i, x) + \\
L^{(i)}(u_i \oplus x, \lambda^*, X^*/\nu_i) - L^{(i)}(u_i, \lambda^*, \emptyset),
\end{aligned} \tag{24}
$$

and for $i \neq j$,

$$
\begin{aligned}
a_{i,j}(\mathbf{z}, x) = c_i(u_i, x) + \\
L^{(i)}(u_i \oplus x, \lambda^*, \emptyset) - L^{(i)}(u_i, \lambda^*, \emptyset) + \\
L^{(j)}(u_j, \lambda^*, X^*/\nu_j) - L^{(j)}(u_i, \lambda^*, \emptyset).
\end{aligned} \tag{25}
$$

Due to lack of space, we leave the explicit expressions to reader, but show both FPI and LH in action later in Section VI.

IV. ANALYSIS OF M/G/1-LCFS

Let us next consider preemptive LCFS scheduling, i.e., where the job that has arrived (if any) most recently receives service. This means that LCFS works as a stack. It has strong insensitivity properties [5]. For example, the mean sojourn time in an M/G/1 with preemptive LCFS and an i.i.d. random setup delay S is [5, Eq. (7)]

$$\mathbb{E}[T] = \frac{\mathbb{E}[X]}{1 - \rho} + \frac{\mathbb{E}[S] + (\lambda/2)\mathbb{E}[S^2]}{1 + \lambda \mathbb{E}[S]}, \tag{26}$$

i.e., it depends only on the service time through its mean.

A. Busy periods

In general, the analysis of LCFS reduces to the analysis of the busy periods as each arriving job can be seen to start a mini busy period at the end of which it then departs. In order to deduce the corresponding results with respect to the squared sojourn time, we let $\tilde{B}(u)$ denote the length of the remaining busy period in a work-conserving M/G/1 queue with the initial backlog of u including a possible (remaining) setup delay δ. The first two moments of the length of the busy period are [9]

$$\mathbb{E}[\tilde{B}] = \frac{\mathbb{E}[X]}{1 - \rho}, \quad \text{and} \quad \mathbb{E}[(\tilde{B})^2] = \frac{\mathbb{E}[X^2]}{(1 - \rho)^3}.$$

We need similar expressions for the remaining busy period:

Lemma 6:

$$\mathbb{E}[\tilde{B}(u)] = \frac{u}{1 - \rho}. \tag{27}$$

and

$$\mathbb{E}[\tilde{B}(u)^2] = \frac{\lambda \mathbb{E}[X^2] u}{(1 - \rho)^3} + \frac{u^2}{(1 - \rho)^2}. \tag{28}$$
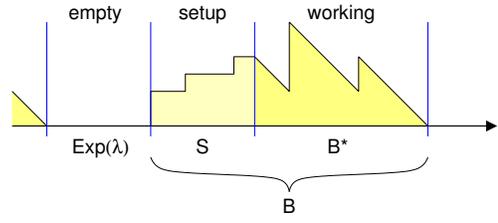

Fig. 3. *One busy cycle consists of three phases: an empty system, a setup phase of s, and a working phase during which jobs are actually processed.*

Proof: The remaining busy period is equal to a random sum,

$$\tilde{B}(u) = u + \tilde{B}_1 + \ldots + \tilde{B}_N,$$

where the $\tilde{B}_i$ are i.i.d. busy periods and $N \sim \text{Poisson}(\lambda u)$. Thus,

$$\mathbb{E}[\tilde{B}(u)] = u + \lambda u \, \mathbb{E}[\tilde{B}] = \frac{u}{1 - \rho}.$$

Similarly, for the second moment we get, by conditioning,

$$\mathbb{E}[\tilde{B}(u)^2 | N] = u^2 + N \, \mathbb{E}[\tilde{B}^2] + 2N \, \mathbb{E}[\tilde{B}] \, u + N(N-1) \, \mathbb{E}[\tilde{B}]^2.$$

Consequently, unconditioning gives (28). ∎

B. Mean costs

In this section, we will discuss the mean costs in an M/G/1-LCFS queue. To this end, let Z denote the sum of the service time of a job X and the remaining setup delay R upon arrival, $Z = R + X$. Thus, without new arrivals, Z would be the sojourn time of the given job. In what follows, we need the first two moments of Z.

Lemma 7:

$$\mathbb{E}[Z] = \mathbb{E}[X] + \frac{(1 - \rho)}{1 + \lambda \mathbb{E}[S]} \left[\mathbb{E}[S] + \frac{\lambda}{2} \mathbb{E}[S^2] \right], \tag{29}$$

$$
\begin{aligned}
\mathbb{E}[Z^2] = \mathbb{E}[X^2] \; + \\
\frac{1 - \rho}{1 + \lambda \mathbb{E}[S]} \left[2 \mathbb{E}[X]\mathbb{E}[S] + (1 + \rho) \mathbb{E}[S^2] + \frac{\lambda}{3} \mathbb{E}[S^3] \right].
\end{aligned} \tag{30}
$$

Proof: The first moment, $\mathbb{E}[Z]$, is already available from [5], while $\mathbb{E}[Z^2]$ can be obtained using PASTA. Conditioning on $X = x$ and $S = s$,

$$\mathbb{E}[Z^2 | x, s] = \frac{1/\lambda}{\mathbb{E}[\ell]} (x + s)^2 + \frac{s}{\mathbb{E}[\ell]} \int_0^s \frac{1}{s} (x + \delta)^2 \, d\delta + \frac{\mathbb{E}[\tilde{B}^*]}{\mathbb{E}[\ell]} x^2$$

where ℓ is the busy-idle-setup cycle length, and $\tilde{B}^*$ is the remaining busy period once the setup phase has completed (see Fig. 3). From [5], we have

$$\mathbb{E}[\tilde{B}^* | S = s] = \frac{(1 + \lambda s)\mathbb{E}[X]}{1 - \rho}, \tag{31}$$

$$\mathbb{E}[\ell | S = s] = \frac{1/\lambda + s}{1 - \rho}. \tag{32}$$

Consequently,

$$\mathbb{E}[Z^2 | x, s] = x^2 + \frac{1 - \rho}{3(1 + \lambda s)} s(6x + 3(1 + \lambda x)s + \lambda s^2),$$

and therefore

$$\mathbb{E}[Z^2|s] = \mathbb{E}[X^2] + \frac{(1-\rho)s}{1+\lambda s}\left(2\,\mathbb{E}[X] + (1+\lambda\,\mathbb{E}[X])s + \frac{\lambda}{3}s^2\right).$$

Then, using (1), yields (30) after some manipulation. ∎

Similarly, we will need the first two moments of the remaining setup delay R an arbitrary arriving job sees:

Corollary 8:

$$\mathbb{E}[R] = \frac{1-\rho}{1+\lambda\,\mathbb{E}[S]}\left(\mathbb{E}[S] + \frac{\lambda}{2}\,\mathbb{E}[S^2]\right), \qquad (33)$$

$$\mathbb{E}[R^2] = \frac{1-\rho}{1+\lambda\,\mathbb{E}[S]}\left(\mathbb{E}[S^2] + \frac{\lambda}{3}\,\mathbb{E}[S^3]\right). \qquad (34)$$

Proof: First we recall that $Z = X + R$ and (33) follows immediately from (29). Similarly,

$$\mathbb{E}[Z^2] = \mathbb{E}[(X+R)^2] = \mathbb{E}[X^2] + 2\,\mathbb{E}[X]\,\mathbb{E}[R] + \mathbb{E}[R^2],$$

and utilizing (30) and (33) yields (34). ∎

Next we assume that (A, B) may depend on the service requirement X.

Theorem 9: The mean costs of a job in an M/G/1-LCFS queue when (A, B) may depend on X are

$$\boxed{\mathbb{E}[AT] = \frac{\mathbb{E}[A]}{1+\lambda\,\mathbb{E}[S]}\left(\mathbb{E}[S] + \frac{\lambda\,\mathbb{E}[S^2]}{2}\right) + \frac{\mathbb{E}[AX]}{1-\rho}.} \qquad (35)$$

and

$$\boxed{\begin{aligned}\mathbb{E}[BT^2] &= \frac{\lambda\,\mathbb{E}[X^2]\mathbb{E}[BX]}{(1-\rho)^3} + \frac{\mathbb{E}[BX^2]}{(1-\rho)^2} \\ &+ \frac{1}{(1-\rho)(1+\lambda\,\mathbb{E}[S])}\Big[\mathbb{E}[B]\left(\mathbb{E}[S^2] + \frac{\lambda}{3}\,\mathbb{E}[S^3]\right) \\ &+ \left(\frac{\lambda\,\mathbb{E}[X^2]}{1-\rho}\mathbb{E}[B] + 2\,\mathbb{E}[BX]\right)\left(\mathbb{E}[S] + \frac{\lambda}{2}\,\mathbb{E}[S^2]\right)\Big].\end{aligned}} \qquad (36)$$

Proof: First,

$$\mathbb{E}[AT] = \mathbb{E}[A\tilde{B}(Z)],$$

where $Z = R + X$. Conditioning

$$\begin{aligned}\mathbb{E}[AT] &= \mathbb{E}[\mathbb{E}[A\tilde{B}(Z)|Z]] \\ &= \mathbb{E}[\frac{AZ}{1-\rho}] = \frac{\mathbb{E}[A]\mathbb{E}[R] + \mathbb{E}[AX]}{1-\rho}.\end{aligned}$$

Substituting $\mathbb{E}[R]$ from (33) gives (35).

For the quadratic cost term $\mathbb{E}[BT^2]$ we obtain similarly

$$\mathbb{E}[BT^2] = \mathbb{E}[\mathbb{E}[BT^2|X, R, B]],$$

which gives, from (26),

$$\begin{aligned}\mathbb{E}[BT^2|X, R, B] &= \frac{\lambda\,\mathbb{E}[X^2]BZ}{(1-\rho)^3} + \frac{BZ^2}{(1-\rho)^2} \\ &= \frac{\lambda\,\mathbb{E}[X^2]BX + BR}{(1-\rho)^3} + \frac{BR^2 + 2BRX + BX^2}{(1-\rho)^2}.\end{aligned}$$

Then substituting $\mathbb{E}[R]$ and $\mathbb{E}[R^2]$ from Corollary 8 and taking the expectation yields (36). ∎

We note that the extra term due to setup delay in $\mathbb{E}[AT]$ is the same as with FCFS. However, for $\mathbb{E}[BT^2]$ this is no longer the case. The result for the mean squared sojourn time in an M/G/1-LCFS follows as a special case:

Corollary 10: The mean squared sojourn time in M/G/1-LCFS with a setup delay S is

$$\boxed{\begin{aligned}\mathbb{E}[T^2] &= \frac{\mathbb{E}[X^2]}{(1-\rho)^3}\left(1 + \frac{\lambda(1-\rho)(\mathbb{E}[S] + \frac{\lambda}{2}\,\mathbb{E}[S^2])}{1+\lambda\,\mathbb{E}[S]}\right) \\ &+ \frac{2\,\mathbb{E}[X]\,\mathbb{E}[S] + (1+\rho)\,\mathbb{E}[S^2] + \frac{\lambda}{3}\,\mathbb{E}[S^3]}{(1-\rho)(1+\lambda\,\mathbb{E}[S])}.\end{aligned}} \qquad (37)$$

Proof: Set $B = b_i = 1$ in (36). ∎

Similarly, (26) can be re-derived from (35) with $A = a_i = 1$.

C. Value function

The state description for a size-aware M/G/1-LCFS queue with a setup delay s is $\mathbf{z} = (\delta; (\tau_1, \Delta_1), \ldots, (\tau_N, \Delta_n))$, where δ denotes the remaining setup time, and (τ_i, Δ_i) are the time job i has already spent in the system and its remaining service time, respectively. In the general case with job-specific cost factors the state is given by $\mathbf{z} = (\delta; (\tau_1, \Delta_1, a_1, b_1), \ldots, (\tau_N, \Delta_n, a_n, b_n))$. First we consider the general case and assume that the cost factors (A, B) may depend on the service requirement X. Let $u_i = \delta + \sum_{j=1}^{i}\Delta_j$, $u = u_n$, and job 1 (if any) is currently receiving service.

Theorem 11: The relative value function with respect to the cost term AT in an M/G/1-LCFS is

$$\boxed{\begin{aligned}v_A(\mathbf{z}) - v_A(0) &= \sum_{i=1}^{n} a_i\left(\tau_i + \frac{u_i}{1-\rho}\right) + \\ &\frac{\lambda\,\mathbb{E}[A]}{1-\rho}\left(\frac{\delta^2}{2} - \frac{u}{1+\lambda\,\mathbb{E}[S]}\left(\mathbb{E}[S] + \frac{\lambda\,\mathbb{E}[S^2]}{2}\right)\right).\end{aligned}} \qquad (38)$$

Proof: Clearly,

$$v(\mathbf{z}) = \mathbb{E}[V_0(\mathbf{z}) + V_1(\mathbf{z}) - r\tilde{B}(\mathbf{z})] + v(0),$$

where $V_0(\mathbf{z})$ denotes the costs the present jobs incur, $V_1(\mathbf{z})$ the costs the jobs arriving during the remaining busy period incur, and r is the mean cost rate, $\lambda\,\mathbb{E}[AT]$. First,

$$\mathbb{E}[V_0(\mathbf{z})] = \sum_i a_i\left(\tau_i + \frac{u_i}{1-\rho}\right).$$

Then,

$$\mathbb{E}[V_1(\mathbf{z})] = \int_0^{\delta}\lambda\frac{\mathbb{E}[A(t+X)]}{1-\rho} + \lambda\left(\frac{u}{1-\rho} - \delta\right)\frac{\mathbb{E}[AX]}{1-\rho},$$

where the integral corresponds to the mean contribution by the jobs arriving during the remaining setup delay δ, and the second term to the jobs arriving after that. This gives

$$\mathbb{E}[V_1(\mathbf{z})] = \frac{\lambda}{1-\rho}\left(\frac{\delta^2}{2}\,\mathbb{E}[A] + \frac{\mathbb{E}[AX]u}{1-\rho}\right).$$

Finally, the mean cost rate is given by (35) and

$$\mathbb{E}[r\tilde{B}(\mathbf{z})] = \frac{\lambda u}{1-\rho}\left(\frac{\mathbb{E}[A]}{1+\lambda\,\mathbb{E}[S]}\left(\mathbb{E}[S] + \frac{\lambda\,\mathbb{E}[S^2]}{2}\right) + \frac{\mathbb{E}[AX]}{1-\rho}\right).$$

When combining these, the terms with $\mathbb{E}[AX]$ cancel each other, and (38) is obtained. ∎

Note that here we assume a general setup delay distribution S for future setups, but also assume the exact information regarding the remaining setup delay δ in the current busy period. Moreover, even though the mean cost $\mathbb{E}[AT]$ depends on $\mathbb{E}[AX]$ according to (35), the corresponding value function depends only on the means $\mathbb{E}[A]$ and $\mathbb{E}[X]$.

Theorem 12: The value function with respect to the cost term BT^2 in an M/G/1-LCFS is

$$
\begin{aligned}
&v_B(\mathbf{z}) - v_B(0) = \\
&\sum_i b_i \left[\tau_i^2 + \frac{u_i}{1-\rho}\left(2\tau_i + \frac{\lambda\,\mathbb{E}[X^2]}{(1-\rho)^2} \right) + \frac{u_i^2}{(1-\rho)^2} \right] \\
&+ \frac{\lambda\delta^2}{(1-\rho)^2}\left(\frac{\lambda\,\mathbb{E}[X^2]\,\mathbb{E}[B]}{2(1-\rho)} + \mathbb{E}[BX] + \frac{\delta\,\mathbb{E}[B]}{3} \right) \\
&- \frac{\lambda u}{(1-\rho)^2(1+\lambda\,\mathbb{E}[S])}\left[\mathbb{E}[B]\left(\mathbb{E}[S^2] + \frac{\lambda}{3}\,\mathbb{E}[S^3] \right) \right. \\
&\left. + \left(\frac{\lambda\,\mathbb{E}[X^2]\,\mathbb{E}[B]}{1-\rho} + 2\,\mathbb{E}[BX] \right)\left(\mathbb{E}[S] + \frac{\lambda}{2}\,\mathbb{E}[S^2] \right) \right].
\end{aligned} \tag{39}
$$

Proof: We prove this by considering the remaining busy period of length $\tilde{B}(u)$. Again we can write

$$
v(\mathbf{z}) = \mathbb{E}[V_0(\mathbf{z}) + V_1(\mathbf{z}) - r \cdot \tilde{B}(u)] + v(0),
$$

where $V_0(\mathbf{z})$ denotes the costs the present n jobs incur, $V_1(\mathbf{z})$ the costs the jobs arriving during the remaining busy period will incur, and r is the mean cost rate, where now $r = \lambda\,\mathbb{E}[T^2]$, and $\mathbb{E}[T^2]$ is given by (37). For $\mathbb{E}[V_0(\mathbf{z})]$ we obtain

$$
\begin{aligned}
\mathbb{E}[V_0(\mathbf{z})] &= \sum_i b_i \mathbb{E}[(\tau_i + \tilde{B}(u_i))^2] \\
&= \sum_i b_i \left[\tau_i^2 + \frac{u_i}{1-\rho}\left(2\tau_i + \frac{\lambda\,\mathbb{E}[X^2]}{(1-\rho)^2} \right) + \frac{u_i^2}{(1-\rho)^2} \right],
\end{aligned}
$$

as the remaining sojourn time of job i corresponds to the (conditional) remaining busy period, where Lemma 6 can be utilized for each job i. The mean cost of a job arriving during the remaining setup delay δ is

$$
\begin{aligned}
\mathbb{E}[B\tilde{B}(X+\delta)] &= \mathbb{E}[B\left(\frac{\lambda\,\mathbb{E}[X^2](X+\delta)}{(1-\rho)^3} + \frac{(X+\delta)^2}{(1-\rho)^2} \right)] \\
&= \frac{\lambda\,\mathbb{E}[X^2]}{(1-\rho)^3}(\mathbb{E}[BX] + \delta\,\mathbb{E}[B]) + \frac{\mathbb{E}[BX^2] + 2\delta\,\mathbb{E}[BX] + \delta^2\,\mathbb{E}[B]}{(1-\rho)^2}.
\end{aligned}
$$

For jobs arriving during the remaining setup and busy period, we take the remaining setup delay into account using the above, and denote $f(t) = \mathbb{E}[B\tilde{B}(X+t)]$. From PASTA,

$$
\mathbb{E}[V_1(\mathbf{z})] = \int_0^\delta \lambda f(t)\,dt + \lambda\left(\frac{u}{1-\rho} - \delta \right) f(0).
$$

The first integral corresponds to new jobs arriving during the remaining setup period, and the latter to new jobs that see the server already running. These reduce to

$$
\begin{aligned}
\mathbb{E}[V_1(\mathbf{z})] =& \frac{\lambda\delta}{6(1-\rho)^3}\left(3\,\mathbb{E}[X^2](2\lambda\,\mathbb{E}[BX] + \lambda\delta\,\mathbb{E}[B]) + \right. \\
& 2(1-\rho)(3\,\mathbb{E}[BX^2] + \delta(3\,\mathbb{E}[BX] + \delta\,\mathbb{E}[B]))\Big) \\
& + \lambda\left(\frac{u}{1-\rho} - \delta \right)\left(\frac{\lambda\,\mathbb{E}[X^2]\,\mathbb{E}[BX]}{(1-\rho)^3} + \frac{\mathbb{E}[BX^2]}{(1-\rho)^2} \right).
\end{aligned}
$$

The mean cost rate, $\lambda\,\mathbb{E}[BT^2]$, times the mean duration of the remaining busy period, $\mathbb{E}[\tilde{B}(u)]$, is

$$
\frac{\lambda u}{1-\rho}\,\mathbb{E}[BT^2],
$$

where $\mathbb{E}[BT^2]$ is given in (36). Then, for $\mathbb{E}[V1(\mathbf{z})] - (\lambda u/(1-\rho))\mathbb{E}[BT^2]$, after some manipulation, one obtains

$$
\begin{aligned}
&\frac{\lambda\delta^2}{(1-\rho)^3}\left(\frac{\lambda\,\mathbb{E}[X^2]\,\mathbb{E}[B]}{2} + (1-\rho)\left(\mathbb{E}[BX] + \frac{\delta\,\mathbb{E}[B]}{3} \right) \right) \\
&- \frac{\lambda u}{(1-\rho)^2(1+\lambda\,\mathbb{E}[S])}\left[\mathbb{E}[B]\left(\mathbb{E}[S^2] + \frac{\lambda}{3}\,\mathbb{E}[S^3] \right) \right. \\
&\left. + \left(\frac{\lambda\,\mathbb{E}[X^2]}{1-\rho}\,\mathbb{E}[B] + 2\,\mathbb{E}[BX] \right)\left(\mathbb{E}[S] + \frac{\lambda}{2}\,\mathbb{E}[S^2] \right) \right].
\end{aligned}
$$

Finally, adding $\mathbb{E}[V_0(\mathbf{z})]$ yields (39). ∎

We note that again the value function is more "robust" than the mean cost rate as (39) does not depend on $\mathbb{E}[BX^2]$, whereas the mean cost $\mathbb{E}[BT^2]$ does. Moreover, the value function with respect to the squared sojourn time with LCFS is a quadratic function of the workloads u_i, whereas the corresponding value function with FCFS depends also on u^3. That is, *adding very large jobs is more expensive than FCFS than under LCFS*. This holds both for the sojourn time and its square. The value function with respect to the sojourn time is given in [5, Eq. (10)].

Corollary 13: The value function for M/G/1-LCFS with respect to the squared sojourn time with a setup delay s is

$$
\begin{aligned}
&v_{T2}(\mathbf{z}) - v_{T2}(0) = \\
&\sum_i \left[\tau_i^2 + \frac{u_i}{1-\rho}\left(2\tau_i + \frac{\lambda\,\mathbb{E}[X^2]}{(1-\rho)^2} \right) + \frac{u_i^2}{(1-\rho)^2} \right] \\
&+ \frac{\lambda}{(1-\rho)^2}\left[\delta^2\left(\frac{\lambda\,\mathbb{E}[X^2]}{2(1-\rho)} + \mathbb{E}[X] + \frac{\delta}{3} \right) - \right. \\
&\left. \frac{us}{1+\lambda s}\left(2\,\mathbb{E}[X] + s(1+\rho+\frac{\lambda s}{3}) + \frac{\lambda\,\mathbb{E}[X^2](2+\lambda s)}{2(1-\rho)} \right) \right],
\end{aligned} \tag{40}
$$

where $u_i = \delta + \Delta_1 + \ldots + \Delta_i$, and $u = u_n$.

Proof: Set $B = b_i = 1$ and $S = s$ (constant) in (39). ∎

D. Policy iteration and lookahead

Let us next return again to the complete system of m parallel servers. For LCFS, a sufficient state description for server i is

$$
\mathbf{z}_i = (\delta; (\tau_{i,1}, \Delta_{i,1}), \ldots, (\tau_{i,n_i}, \Delta_{i,n_i})),
$$

where δ denotes the remaining setup delay, and $(\tau_{i,j}, \Delta_{i,j})$ are the elapsed time and the remaining service time of job j of server i. The state of the system is $\mathbf{z} = (\mathbf{z}_1, \ldots, \mathbf{z}_m)$. We let $\mathbf{z}_i \oplus x$ denote the new state of server i with a job of size x added,

$$
\mathbf{z}_i \oplus x = \begin{cases} (\delta; (0,x), (\tau_{i,1}, \Delta_{i,1}), \ldots, (\tau_{i,n_i}, \Delta_{i,n_i})), & \text{if } n_i > 0, \\ (s; (0,x)), & \text{if } n_i = 0. \end{cases}
$$

To illustrate FPI, let us consider again the cost structure according to the squared sojourn time T^2. The value function is given in (40), and as the "squared sojourn time cost" is paid upon the departure, there is no immediate cost associated
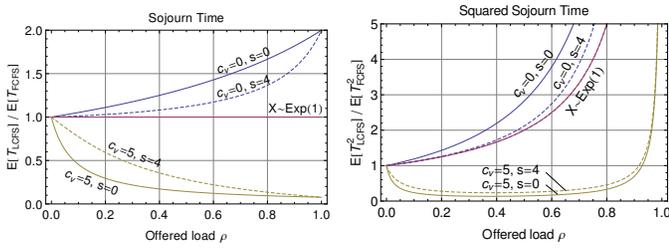
Fig. 4. Comparison of LCFS to FCFS in M/G/1 with respect to the mean sojourn time $\mathbb{E}[T]$ (left) and the mean squared sojourn time $\mathbb{E}[T^2]$ with (dashed line) and without (solid line) setup time.

Fig. 5. Simulation results for $\mathbb{E}[T^2]$ with FCFS (left) and LCFS (right).

with adding a job to a queue. Hence, FPI (see Section III-C) reduces to

$$\alpha_{\text{FPI}}(\mathbf{z}, x) = \arg\min_i \left(v_{T2}^{(i)}(\mathbf{z}_i \oplus x) - v_{T2}^{(i)}(\mathbf{z}_i) \right).$$

The lookahead policy can be derived similarly as for FCFS in Section III-D. We omit the details due to lack of space.

V. COMPARISON OF LCFS AND FCFS

With exponentially distributed service times (M/M/1-queue), FCFS and LCFS achieve the same mean sojourn time for any i.i.d. random setup time S, $\mathbb{E}[T]^{\text{(LCFS)}} = \mathbb{E}[T]^{\text{(FCFS)}}$ [5]. However, the well-known deficiency of LCFS (without setup delay) is the higher variance in the sojourn time for M/M/1 queues. Whether this is still the case with a setup delay is an interesting question.

Corollary 14: For an M/M/1-queue with a constant setup time s it holds that

$$\boxed{\frac{\mathbb{E}[T^2]^{\text{(LCFS)}}}{\mathbb{E}[T^2]^{\text{(FCFS)}}} = \frac{1}{1-\rho} \qquad \forall\, s \geq 0.}$$

Proof: This follows from (10) and (37), and the fact that $\mathbb{E}[X^2] = 2/\mu^2$ and $\mathbb{E}[X^3] = 6/\mu^3$. ∎

For general service times, with $\mathbb{E}[X]$, $\mathbb{E}[X^2]$, $\mathbb{E}[X^3]$ finite,

$$\lim_{\rho \to 1} \frac{\mathbb{E}[T^2]^{\text{(LCFS)}}}{\mathbb{E}[T^2]^{\text{(FCFS)}}} = \infty,$$

i.e., in the *heavy traffic* limit, LCFS is worse than FCFS with respect to the mean squared sojourn time with or without a setup delay. We note that this is not the case with the mean sojourn time metric, for which LCFS is superior when the coefficient of variation is sufficiently high.

Moreover, as $\mathbb{E}[T^2]$ depends on $\mathbb{E}[X^3]$ with FCFS, but only on $\mathbb{E}[X^2]$ with LCFS, it follows that for any $\rho < 1$, and a sequence of service times $\{X_i\}$ such that $\lim_{i \to \infty} \mathbb{E}[X_i^k] = \infty$ for $k = 2$ or $k = 3$, while $\mathbb{E}[X_i] = 1$, it holds that

$$\frac{\mathbb{E}[T^2]^{\text{(LCFS)}}}{\mathbb{E}[T^2]^{\text{(FCFS)}}} \to 0,$$

i.e., LCFS is asymptotically better than FCFS when the offered load is finite and the service times become heavy-tailed. Also this holds for queues with and without a setup delay.

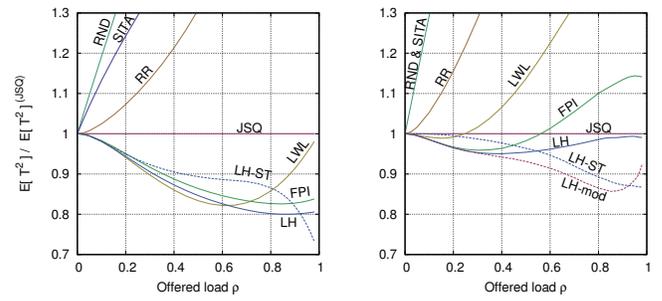Our observations are illustrated in Fig. 4, where the service times obey Weibull distribution with parameters such that $\mathbb{E}[X] = 1$ and the coefficient of variation, $c_v = \sigma/\mathbb{E}[X]$, is $c_v \in \{0, 1, 5\}$, corresponding to service times being constant, exponentially distributed and "heavy-tailed", respectively. We consider a queue without a setup delay (solid lines), and with a setup delay of $s = 4$ (dashed lines). The left figure illustrates the mean sojourn time ratio of LCFS to FCFS, and the right shows the corresponding ratio for the mean squared sojourn time. We can see that with the setup time the difference in performance (either way) becomes smaller.

VI. NUMERICAL EXAMPLES

In this section, we will give some numerical examples that illustrate the use of the analytical results derived earlier. We consider a system of 2 parallel servers that process the arriving jobs according to FCFS or LCFS scheduling. The reference dispatching policies are: [10]

- RND is a random Bernoulli split that balances the load.
- SITA, size-interval-task-assignment, sends jobs shorter than ξ to server 1, and the rest to server two. Here we have chosen such ξ that balances the load (SITA-e).
- RR (round-robin) uses a sequential order $1, \ldots, m, 1, \ldots$
- JSQ, join-the-shortest-queue, chooses the server with less jobs. Ties are resolved in favor of server 1.
- LWL (least-work-left) chooses the queue with a shorter backlog. Ties are resolved in favor of server 1.
- FPI is the first policy iteration applied to SITA or RND.
- LH is the lookahead policy based on SITA or RND. We include also LH-ST based on minimizing the sojourn time.

A. FCFS

Consider a small system of two identical servers with service rates $\nu_1 = \nu_2 = 1$ so that the total capacity is $c_{\text{tot}} = 2$. The jobs arrive according to a Poisson process with rate λ and have exponentially distributed sizes $X \sim \text{Exp}(1)$. The objective is to minimize the squared sojourn time T^2.

Fig. 5 (left) depicts the numerical results with different dispatching policies. FPI and LH are based on SITA, which is a significantly better basic dispatch rule than RND for FCFS scheduling. On the x-axis is the offered load, $\rho = \lambda\,\mathbb{E}[X]/c_{\text{tot}}$, and the y-axis corresponds to the performance relative to JSQ. We observe that at small loads LWL, FPI and LH do a pretty good job. At higher loads, the performance of LWL degrades, and in general, LH achieves the best performance.
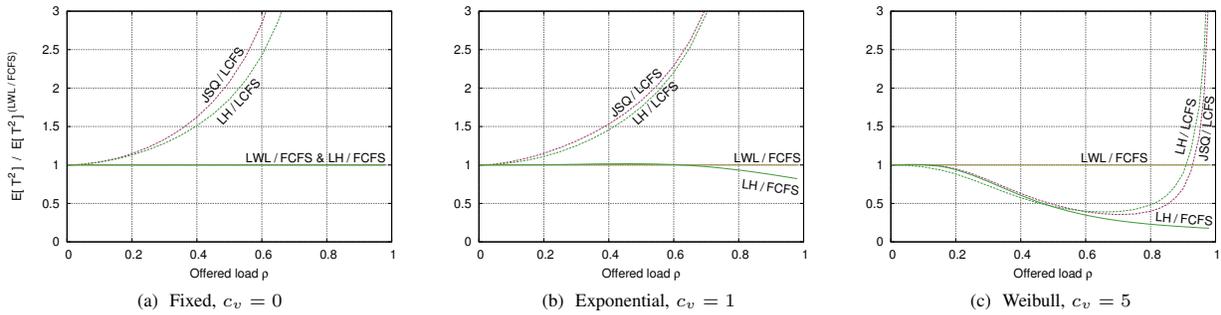
Fig. 6. Comparison of FCFS and LCFS with different job size distributions with respect to $\mathbb{E}[T^2]$.

B. LCFS

Next we consider the same two server system with LCFS scheduling, i.e., the example system comprises two identical servers, service times are exponentially distributed, and the objective is to minimize the squared sojourn time. Fig. 5 (right) depicts the numerical results. The basic policy (starting point) for FPI and LH is RND. The mean (squared) sojourn time is the same with RND and SITA. Under low load, FPI and LH achieve the best performance. However, as the load increases, the lookahead policies become superior. This is not surprising as the lookahead approach in general has a better "understanding" about the complete state of the system, whereas FPI assumes that the queues are isolated (via the static basic policy). Under a very high load, when $\rho \to 1$, LH-ST, focusing on the sojourn time instead of its square, seems to outperform LH. This applies to both FCFS and LCFS.

C. Comparison of FCFS and LCFS

In Section V, we compared FCFS to LCFS for $\mathbb{E}[T^2]$ in a single server queue with and without setup delays. Here we carry out a similar comparison for two parallel servers. Specifically, we assume that job sizes obey Weibull distribution with $c_v = 0, 1, 5$, so that $c_v = 0$ corresponds to fixed size, and with $c_v = 1$ one obtains the exponential distribution. For each case, we have chosen a good reference policy, LWL for FCFS and JSQ for LCFS, and the corresponding LH policy (based on SITA for FCFS, and RND for LCFS). Our expectation, based on the single server results of Section V, is that LCFS is better than FCFS with heavy-tailed job sizes, and vice versa.

Figure 6 depicts the numerical results. We use $\mathbb{E}[T^2]$ with FCFS and LWL as the reference and show the relative performance with each set of parameters, $\mathbb{E}[T^2]^{(x)}/\mathbb{E}[T^2]^{(\text{LWL,FCFS})}$. We see that a smart dispatching policy, LH based on SITA, with FCFS, achieves the lowest $\mathbb{E}[T^2]$ in overall. This applies also to the case with "heavy-tailed" job sizes, i.e., the dispatching policy manages to mitigate the deficiency in the scheduling.

Finally, we note that the results with a setup delay are similar; it does not change characteristics of FCFS or LCFS.

VII. CONCLUSIONS

We have considered general linearly increasing cost rates in the context of parallel FCFS and LCFS servers. Dispatching jobs under such cost structure turns out to be non-trivial, which is not surprising as the problem is generally intractable even for the mean sojourn time. Hence, we resort to FPI and LH approaches. To this end, we first derived the value functions for M/G/1-FCFS and -LCFS queues with respect to the general cost structure, which includes both the mean sojourn time $\mathbb{E}[T]$ and the mean squared sojourn time $\mathbb{E}[T^2]$ as special cases. The value functions enable the efficient cost-aware FPI and LH dispatching policies, which we evaluated under different job size distributions. We made several interesting observations. First, under a very high load and LCFS, a better performance was obtained if LH policy sought to minimize the mean sojourn time instead of its square. Second, parallel servers with LCFS seem to suffer more from unfairness than in a single server system, and FCFS combined with LH dispatching policy achieves constantly very high performance in terms of T^2 even when the job sizes varied a lot (a regime considered favorable for LCFS). In contrast, for the mean sojourn time, LCFS is manageable and can be significantly better than FCFS system [5]. This suggests that FCFS enables fairness and the task of the dispatching policy is to deal with the varying job sizes.

REFERENCES

[1] A. Gandhi, M. Harchol-Balter, and I. Adan, "Server farms with setup costs," *Performance Evaluation*, vol. 67, no. 11, pp. 1123–1138, 2010.
[2] V. Maccio and D. Down, "On optimal policies for energy-aware servers," in *In Proc. of MASCOTS*, San Francisco, US, Aug. 2013.
[3] R. Righter, J. G. Shanthikumar, and G. Yamazaki, "Extremal properties of the FIFO discipline in queueing networks," *Journal of Applied Probability*, vol. 29, no. 4, pp. 967–978, Dec. 1992.
[4] R. C. Larson, "OR forum – perspectives on queues: Social justice and the psychology of queueing," *Oper. Res.*, vol. 35, no. 6, 1987.
[5] E. Hyytiä, R. Righter, and S. Aalto, "Energy-aware job assignment in server farms with setup delays under LCFS and PS," in *26th International Teletraffic Congress (ITC'26)*, Sweden, Sep. 2014.
[6] ——, "Task assignment in a heterogeneous server farm with switching delays and general energy-aware cost structure," *Performance Evaluation*, vol. 75–76, no. 0, pp. 17–35, May-June 2014.
[7] P. D. Welch, "On a generalized M/G/1 queuing process in which the first customer of each busy period receives exceptional service," *Operations Research*, vol. 12, no. 5, pp. 736–752, 1964.
[8] E. Hyytiä, "Lookahead actions in dispatching to parallel queues," *Performance Evaluation*, vol. 70, no. 10, pp. 859–872, 2013.
[9] L. Kleinrock, *Queueing Systems, Volume I: Theory*. Wiley, 1975.
[10] M. Harchol-Balter, *Performance Modeling and Design of Computer Systems: Queueing Theory in Action*. Cambridge University Press, 2013.