



This is an electronic reprint of the original article. This reprint may differ from the original in pagination and typographic detail.

Aalto, Samuli

Whittle index approach to multiserver scheduling with impatient customers and DHR service times

Published in: QUEUEING SYSTEMS

DOI: 10.1007/s11134-024-09902-5

Published: 01/06/2024

Document Version Publisher's PDF, also known as Version of record

Published under the following license: CC BY

Please cite the original version:

Aalto, S. (2024). Whittle index approach to multiserver scheduling with impatient customers and DHR service times. *QUEUEING SYSTEMS*, *107*(1-2), 1-30. https://doi.org/10.1007/s11134-024-09902-5

This material is protected by copyright and other intellectual property rights, and duplication or sale of all or part of any of the repository collections is not permitted, except that material may be duplicated by you for your research use or educational purposes in electronic or print form. You must obtain permission for any other use. Electronic or print copies may not be offered, whether for sale or otherwise to anyone who is not an authorised user.



# Whittle index approach to multiserver scheduling with impatient customers and DHR service times

Samuli Aalto<sup>1</sup>

Received: 31 March 2023 / Revised: 31 August 2023 / Accepted: 3 January 2024 / Published online: 19 February 2024 © The Author(s) 2024

# Abstract

We consider the optimal scheduling problem in a multiserver queue with impatient customers belonging to multiple classes. We assume that each customer has a random abandonment time, after which the customer leaves the system if its service has not been completed before that. In addition, we assume that the scheduler is not able to anticipate the expiration of the abandonment times but only knows their distributions and how long each customer has been in the system. Many papers consider this scheduling problem under Poisson arrivals and linear holding costs assuming further that both the service times and the abandonment times have exponential distributions. Even with these additional assumptions, the exact solution is known only in very few special cases. To tackle this tricky problem, we apply the Whittle index approach. Unlike the earlier papers, which were restricted to exponential service times, we allow the service time distributions for which the hazard rate is decreasing. The Whittle index approach is applied to the discrete-time multiserver queueing problem with discounted costs. As our main theoretical result, we prove that the related relaxed optimization problem is indexable and derive the corresponding Whittle index explicitly. Based on this discrete-time result, we develop a reasonable heuristic for the original continuous-time multiserver scheduling problem. The performance of the resulting policy is evaluated in the M/G/M setup by numerical simulations, which demonstrate that it, indeed, gives better performance than the other policies included in the comparison.

Keywords Optimal scheduling  $\cdot$  Multiserver queue  $\cdot$  Impatient customer  $\cdot$  Abandonment  $\cdot$  Whittle index  $\cdot$  DHR

Mathematics Subject Classification  $~60K25\cdot 90B22\cdot 90B36\cdot 68M20$ 

Samuli Aalto samuli.aalto@aalto.fi

<sup>&</sup>lt;sup>1</sup> Department of Information and Communications Engineering, Aalto University, Espoo, Finland

## **1** Introduction

We consider the optimal scheduling problem in a multiserver queue with impatient customers belonging to multiple classes. In some cases, customer impatience is best described by deterministic deadlines that are known to the scheduler. However, in this paper, we assume that each customer has a random abandonment time, after which the customer leaves the system if its service has not been completed before that. We also assume that the scheduler is not able to anticipate the expiration of the abandonment times but only knows their distributions and how long each customer has been in the system. We are interested in optimizing the scheduling policy among all nonanticipating policies (thus, allowing preemption but ignoring the remaining service and abandonment times) with an objective function that takes into account both the average holding costs of the customers in the system and the average abandonment penalties related to the customers that leave the system before their service has been completed.

Many papers consider this scheduling problem under Poisson arrivals and linear holding costs assuming further that both the service times and the abandonment times have exponential distributions. Even with these additional assumptions, the exact solution is known only in very few special cases, see [7, 8, 10].<sup>1</sup> Atar et al. [2, 3] approach this tricky problem by fluid-scaling methods. They introduce a modification of the well-known  $c\mu$ -rule, called the  $c\mu/\theta$ -rule, which takes into account the abandonments, and prove that it is asymptotically optimal in overload traffic conditions. Larrañaga et al. [14] approximate the stochastic model by a deterministic fluid model and derive a heuristic policy for the original problem, which is a combination of the  $c\mu/\theta$  rules.

Another approach applied to tackle the original problem under these exponential assumptions is based on the so-called Whittle index, originally developed in the context of restless bandit problems [26]. The idea is that the original constrained optimization problem with a finite upper limit for the scheduled customers in each time epoch is first relaxed by requiring that the constraint is satisfied just on average in time. This makes the problem much more tractable by decomposing it to separate subproblems per each customer in the system. However, when applying this approach, it has to be justified that the related relaxed problem is *indexable* [26]. Only for such problems, it is possible to derive a unique Whittle index. If the relaxed problem is indexable, the resulting Whittle index policy is known to be asymptotically optimal under certain technical conditions [23, 24]. Larrañaga et al. [15, 23] show that the relaxation of the scheduling problem is, indeed, indexable under these exponential assumptions, and derive the related Whittle index. For linear holding costs, the Whittle index policy proves to be the same  $c\mu/\theta$ -rule as derived in [2, 3] by fluid-scaling methods. The Whittle index approach is applied also in Ayesta et al. [5, 6] under exponential assumptions and linear holding costs. Due to some small differences in the model assumptions, they

<sup>&</sup>lt;sup>1</sup> This is due to both multiple parallel servers and customer impatience. For a single-server system without abandonments, the optimal policy is known to be the  $c\mu$ -rule under the exponential assumptions [9]. In addition, for an M/G/1 queue with general service time distributions (but still without abandonments), the Gittins index policy is optimal minimizing the average holding costs [1, 11, 21] and near-optimal in the corresponding multiserver case [20].

end up in a slightly modified version of the  $c\mu/\theta$ -rule when determining the Whittle index policy.

In some papers, the previous model assumptions are generalized by allowing more generally distributed abandonment times [12, 16]. However, in this paper, we stick to the exponential abandonment times but allow more general service times than just exponential. This setup is considered in [4, 13] but, in both of these papers, the study is limited to non-preemptive scheduling policies while our study will cover all preemptive policies, as well.

In this paper, we assume that the service times are DHR, i.e., the hazard rate of the service time distribution is decreasing<sup>2</sup> for each class. We apply the Whittle index approach to find out a reasonable heuristic among the nonanticipating policies for the optimal scheduling problem with DHR service times and exponential abandonment times. As far as we know, such results have not been presented earlier. By numerical simulations, we furthermore demonstrate that the resulting Whittle index policy outperforms the other scheduling policies included in the comparison.

As in [5], we apply the Whittle index approach in the closed version<sup>3</sup> of the corresponding discrete-time model with discounted costs (see Sects. 2, 3). The related relaxed problem is shown to be indexable and the corresponding Whittle index is derived explicitly in Sect. 4 (see Theorem 1), which is the main theoretical result of the paper. In Sect. 5, we move from discounted to undiscounted costs. In addition, we move from the discrete-time setup to the original scheduling problem in continuous time, and determine the Whittle index policy in this setting. The performance of the resulting Whittle index policy is evaluated by numerical simulations in Sect. 6, and Sect. 7 concludes the paper.

## 2 Scheduling problem in discrete time with discounted costs

We consider the following multiserver optimal scheduling problem in discrete time. Assume that there are *M* homogeneous parallel servers. Let the time slots be indexed by  $t \in \{1, 2, ...\}$ . In the beginning of the first time slot, there are *K* customers. Let  $S_k$  denote the random *service time* of customer  $k, k \in \{1, 2, ..., K\}$ , taking values in  $\{1, 2, ...\}$ . The service time distribution of customer *k* is general, and the service times are assumed to be independent. Let  $\mu_k(n), n \in \{0, 1, ...\}$ , denote the corresponding *discrete hazard rate* [22], i.e., the conditional probability that the service time of customer *k* is equal to n + 1, given that it is strictly greater than *n*,

$$\mu_k(n) = P\{S_k = n+1 \mid S_k > n\}.$$
(1)

In addition, let  $D_k$  denote the random *life time* of customer k, i.e., the time interval from the beginning until the customer leaves the system due to abandonment (unless

<sup>&</sup>lt;sup>2</sup> In this paper, we use the terms "decreasing" and "increasing" in their weak forms. Thus, e.g., a decreasing function is not required to be strictly decreasing.

<sup>&</sup>lt;sup>3</sup> Following [25, 26], we call the dynamic setup with Poisson arrivals and *K* classes of customers as the *open* version of the problem, while the *closed* version refers to the transient case where there are *K* (heterogeneous) customers in the system at time 0 and no new arrivals later on.

its service has already been completed prior to that moment). We assume here that the life time of customer k,  $D_k$ , is independently and geometrically distributed taking values in  $\{1, 2, ...\}$  with a fixed abandonment probability  $\theta_k > 0$ .

Customers are served according to a *nonanticipating*<sup>4</sup> scheduling policy  $\pi$  that allows preemptions. Let  $\Pi$  denote the family of such disciplines. In the beginning of any time slot  $t \in \{1, 2, ...\}$ , the scheduler chooses at most M of the customers for service (during that time slot). Denote  $A_k^{\pi}(t) = 1$  if customer k is chosen for service in the beginning of time slot t; otherwise, let  $A_k^{\pi}(t) = 0$ . Thus, for any policy  $\pi$  and time slot t, we have the following constraint:

$$\sum_{k=1}^{K} A_k^{\pi}(t) \le M.$$
<sup>(2)</sup>

When making the decision in the beginning of a time slot, say t, a nonanticipating policy  $\pi$  does not know the exact service times  $S_k$  nor the abandonment times  $D_k$  but only the attained services of all the customers prior to that time slot t. Let  $X_k^{\pi}(t)$  denote the state of customer k in the beginning of time slot t taking values in

$$\mathcal{S} = \{0, 1, \ldots\} \cup \{A, B\}.$$

If customer k is still in the system in the beginning of time slot t, its state  $X_k^{\pi}(t) \in \{0, 1, \ldots\}$  refers to the attained service prior to time slot t. However, as soon as customer k leaves the system, its state is marked by symbol A if the reason is abandonment and by symbol B if the reason is service completion. As for a possible abandonment, we assume that customer k leaves the system due to abandonment in time slot t if the service of the customer has not been completed prior to time slot t and  $D_k = t$ . Therefore, even if customer k was scheduled in the beginning of time slot t, it leaves the system due to abandonment if  $D_k = t$ .

For any customer k, holding costs are accumulated at rate  $h_k > 0$  until the whole customer is completed or the customer leaves due to abandonment. In addition, if customer k leaves the system due to abandonment, it incurs a penalty cost of  $d_k$  (as a lump sum). When the customer has left the system, no more costs accumulate any more. All costs are *discounted* in time by factor  $\beta \in (0, 1)$ . The objective function in our scheduling problem is, thus, given by

$$E\left[\sum_{t=1}^{\infty}\sum_{k=1}^{K}\beta^{t-1}\left(h_{k}+d_{k}\mathbf{1}_{\{D_{k}=t\}}\right)\mathbf{1}_{\{X_{k}^{\pi}(t)\in\{0,1,\ldots\}\}}\right].$$
(3)

The aim is to find the optimal scheduling policy  $\pi$  that minimizes the expected discounted costs (3) subject to the strict capacity constraint (2) for all time slots *t* and assuming that the scheduling decisions in each time slot *t* are based on the states  $X_k^{\pi}(t)$  of all the customers.

<sup>&</sup>lt;sup>4</sup> A nonanticipating scheduling policy is aware of the attained services and the times that the customers have already spent in the system, but it does not have any knowledge of the remaining service nor the remaining life times of the customers.

## 3 Whittle index approach to the scheduling problem

Due to abandonments, the original optimization problem described in the previous section belongs, even in the single-server case M = 1, to the class of *restless bandit* problems, which are known to be mathematically intractable. In [26], Whittle proposed to relax such problems by first replacing the strict capacity constraint of the scheduler (such as (2) in our problem) by a time-averaged one, and then considering the Lagrangian version of the relaxed problem, which makes the problem separable. If the relaxed problem can be proved to be *indexable*, this leads to so-called Whittle index rule, which solves the relaxed problem, and serves as a reasonable heuristic for the original problem.

In this paper, we apply Whittle's approach [26], which results in the following (separate) subproblems for each customer k: Find the optimal policy  $\pi$  that minimizes the objective function

$$f_{k,\beta}^{\pi} + \nu g_{k,\beta}^{\pi},\tag{4}$$

where  $\nu$  can be interpreted as the unit price of work,  $f_{k,\beta}^{\pi}$  as the expected discounted costs of customer k, and  $g_{k,\beta}^{\pi}$  as the expected discounted amount of work allocated to customer k,

$$f_{k,\beta}^{\pi} = E\left[\sum_{t=1}^{\infty} \beta^{t-1} (h_k + d_k \mathbf{1}_{\{D_k=t\}}) \mathbf{1}_{\{X_k^{\pi}(t) \in \{0,1,\ldots\}\}}\right],$$
$$g_{k,\beta}^{\pi} = E\left[\sum_{t=1}^{\infty} \beta^{t-1} A_k^{\pi}(t)\right].$$

The separable subproblems (4) are considered in the context of Markov decision processes. The possible actions  $a_k \in \mathcal{A} = \{0, 1\}$  are "to schedule" ( $a_k = 1$ ) and "not to schedule" ( $a_k = 0$ ).

Consider any customer k. Let  $q_k(y|x, a) \ge 0$  denote the transition probability from state  $x \in S$  to state  $y \in S$  after action  $a \in A$ . The nonzero transition probabilities in our model are as follows:

$$\begin{aligned}
q_k(A|n, a) &= \theta_k, & n \in \{0, 1, ...\}, a \in \mathcal{A}, \\
q_k(n|n, 0) &= 1 - \theta_k, & n \in \{0, 1, ...\}, \\
q_k(B|n, 1) &= (1 - \theta_k)\mu_k(n), & n \in \{0, 1, ...\}, \\
q_k(n + 1|n, 1) &= (1 - \theta_k)(1 - \mu_k(n)), n \in \{0, 1, ...\}, \\
q_k(x|x, a) &= 1, & x \in \{A, B\}, a \in \mathcal{A}.
\end{aligned}$$
(5)

Note that the states A and B are absorbing for any policy. Let then  $c_k(x, a)$  denote the expected immediate cost in state  $x \in S$  after action  $a \in A$ . In our model,

$$c_k(n, a) = h_k + \theta_k d_k + a\nu, \ n \in \{0, 1, \ldots\}, \ a \in \mathcal{A},$$
  

$$c_k(x, a) = a\nu, \qquad x \in \{A, B\}, \ a \in \mathcal{A}.$$
(6)

🖉 Springer

Since the state space S is discrete, the action space A is finite, and the expected immediate costs are bounded, the optimal policy belongs to the class of stationary policies [19, Thm. 6.3]. For each stationary policy  $\pi$ , the decisions  $A_k^{\pi}(t)$  related to customer k are deterministic depending just on the current state  $X_k^{\pi}(t)$ ,

$$A_k^{\pi}(t) = \begin{cases} 1, & \text{if } X_k^{\pi}(t) \in \mathcal{B}^{\pi}, \\ 0, & \text{otherwise,} \end{cases}$$

where  $\mathcal{B}^{\pi} \subset \mathcal{S}$  is called the *activity set* of policy  $\pi$ .

Finally, let  $V_{k,\beta}(x; \nu)$  denote the *value function* for state  $x \in S$  related to the minimization of the expected discounted costs (4) with Lagrangian parameter  $\nu$ . The corresponding *optimality equations* [19, Thm. 6.1] read as follows:

$$V_{k,\beta}(x;\nu) = \min\{0,\nu\} + \beta V_{k,\beta}(x;\nu), \quad x \in \{A, B\},$$

$$V_{k,\beta}(n;\nu) = h_k + \theta_k d_k + \min\left\{\beta\left(\theta_k V_{k,\beta}(A;\nu) + (1-\theta_k)V_{k,\beta}(n;\nu)\right), \\ \nu + \beta\left(\theta_k V_{k,\beta}(A;\nu) + (1-\theta_k)(\mu_k(n)V_{k,\beta}(B;\nu) + (1-\mu_k(n))V_{k,\beta}(n+1;\nu))\right)\right\},$$

$$n \in \{0, 1, \ldots\}.$$
(7)

In addition, let  $V_{k,\beta}^{\pi}(x; \nu)$  denote the corresponding value function for policy  $\pi$ . The policy  $\pi_k^*$  for which

$$V_{k,\beta}^{\pi_k^*}(x;\nu) = V_{k,\beta}(x;\nu)$$

for all  $x \in S$  is said to be  $(\nu, \beta)$ -optimal for customer k.

Let us conclude this section by defining the *indexability* property, which is not automatically guaranteed for genuine restless bandit problems [26].

**Definition 1** The relaxed optimization problem with objective function (4) and related to customer *k* is *indexable* if, for any state  $x \in S$ , there exists  $W_{k,\beta}(x) \in [-\infty, \infty]$  such that

- (i) decision a = 1 (to schedule customer k) is optimal in state x if and only if ν ≤ W<sub>k,β</sub>(x);
- (ii) decision a = 0 (not to schedule customer k) is optimal in state x if and only if  $v \ge W_{k,\beta}(x)$ .

If the problem is indexable, the corresponding index  $W_{k,\beta}(x)$  is called the *Whittle index*.

Note that, according to this definition, the two actions are equally good (and, thus, optimal) in state x if and only if  $v = W_{k,\beta}(x)$ .

In the following section, we prove that the relaxed optimization problem with objective function (4) is, indeed, indexable for any *k* under the assumption that the discrete hazard rate  $\mu_k(n)$  is a decreasing function of *n*.

## 4 Whittle index for DHR service times

In this section, we solve the relaxed optimization problem with objective function (4) for a single customer, say customer k. Since we are all the time considering a single customer, we leave out the related subscript k to lighten the notation. The relaxed problem is shown to be indexable and the corresponding Whittle index is derived explicitly (Theorem 1), which is the main theoretical result of the whole paper.

Throughout this section, we assume that the discrete hazard rate  $\mu(n)$  of the customer's service time distribution is a decreasing function of *n*, i.e.,

$$\mu(0) \ge \mu(1) \ge \mu(2) \ge \dots \ge \mu(\infty),\tag{8}$$

where we have defined

$$\mu(\infty) = \lim_{n \to \infty} \mu(n) \ge 0.$$
(9)

Under this assumption, we prove that the relaxed optimization problem with objective function (4) is indexable. In addition, we derive the corresponding Whittle index values  $W_{\beta}(x)$  for any state  $x \in S$  by solving the problem (4) for any  $\nu \in (-\infty, \infty)$ .

**Theorem 1** Consider a single customer under assumption (8).<sup>5</sup> The relaxed optimization problem with objective function (4) is indexable, and the corresponding Whittle indexes are given by the following formulas:

$$W_{\beta}(n) = (h + \theta d) \frac{\beta(1 - \theta)}{1 - \beta(1 - \theta)} \mu(n), \quad n \in \{0, 1, \dots, \},$$
  

$$W_{\beta}(A) = W_{\beta}(B) = 0.$$
(10)

Proof Note first that, by the DHR assumption, we clearly have

$$W_{\beta}(0) \ge W_{\beta}(1) \ge W_{\beta}(2) \ge \dots \ge W_{\beta}(\infty), \tag{11}$$

where we have defined

$$W_{\beta}(\infty) = \lim_{n \to \infty} W_{\beta}(n) = (h + \theta d) \frac{\beta(1 - \theta)}{1 - \beta(1 - \theta)} \mu(\infty) \ge 0.$$
(12)

The main proof is given below in five parts  $(1^{\circ}-5^{\circ})$ . For the proof, we partition the possible values of  $\nu$  into separate intervals and solve the relaxed optimization problem with objective function (4) in these intervals by utilizing the optimality equations (7). This is reflected by the five parts of the main proof.

In parts 1°–4°, we consider the nonnegative values of  $v, v \ge 0$ . From optimality equations (7), we see that, in this case, the minimum expected discounted cost  $V_{\beta}(x; v)$  for the absorbing states  $x \in \{A, B\}$  clearly equals 0 and is achieved by the policies  $\pi$ 

<sup>&</sup>lt;sup>5</sup> In the sequel, this is called the *DHR assumption*.

that choose action 0 in states A and B. It follows that, for any  $\nu \ge 0$ , the optimality equations (7) for the remaining states  $n \in \{0, 1, ...\}$  read as follows:

$$V_{\beta}(n; \nu) = h + \theta d + \min \left\{ \beta (1 - \theta) V_{\beta}(n; \nu), \\ \nu + \beta (1 - \theta) (1 - \mu(n)) V_{\beta}(n+1; \nu) \right\}.$$
(13)

In the last part 5°, we consider the nonpositive values of  $\nu$ ,  $\nu \leq 0$ . In this case, the minimum expected discounted cost  $V_{\beta}(x; \nu)$  for the absorbing states  $x \in \{A, B\}$  equals clearly  $\nu/(1 - \beta)$  and is achieved by those policies  $\pi$  that choose action 1 in states A and B. Moreover, for any  $\nu \leq 0$ , the optimality equations (7) for the remaining states  $n \in \{0, 1, \ldots\}$  read as follows:

$$V_{\beta}(n;\nu) = h + \theta d + \beta \theta \frac{\nu}{1-\beta} + \min\left\{\beta(1-\theta)V_{\beta}(n;\nu), \\ \nu + \beta(1-\theta)\left(\mu(n)\frac{\nu}{1-\beta} + (1-\mu(n))V_{\beta}(n+1;\nu)\right)\right\},$$
(14)  
$$n \in \{0, 1, \ldots\}.$$

1° Let  $\nu \in [W_{\beta}(0), \infty)$ . We prove that the policy  $\pi$  with activity set

$$\mathcal{B}^{\pi} = \emptyset,$$

according to which user k is not scheduled in any state  $x \in S$ , is  $(\nu, \beta)$ -optimal. It remains to prove that policy  $\pi$  is optimal in any state  $n \in \{0, 1, ...\}$ .

We start the proof by first deriving the value function  $V^{\pi}_{\beta}(n; \nu)$  for policy  $\pi$  from the following equations:

$$V_{\beta}^{\pi}(n;\nu) = h + \theta d + \beta (1-\theta) V_{\beta}^{\pi}(n;\nu), \quad n \in \{0, 1, \ldots\}.$$
(15)

The unique solution of these linear equations is clearly given by

$$V_{\beta}^{\pi}(n;\nu) = \frac{h+\theta d}{1-\beta(1-\theta)}, \quad n \in \{0, 1, \ldots\}.$$
 (16)

Let  $n \in \{0, 1, ...\}$ . By (16), the following condition for optimality of  $\pi$  in state n [based on the optimality equation (13)],

$$\beta(1-\theta)V_{\beta}^{\pi}(n;\nu) \leq \nu + \beta(1-\theta)(1-\mu(n))V_{\beta}^{\pi}(n+1;\nu),$$

is easily shown to be equivalent with

$$\nu \ge (h + \theta d) \,\frac{\beta(1 - \theta)}{1 - \beta(1 - \theta)} \,\mu(n). \tag{17}$$

D Springer

This condition is satisfied due to the DHR assumption, since we assumed that

$$\nu \ge W_{\beta}(0) = (h + \theta d) \frac{\beta(1-\theta)}{1-\beta(1-\theta)} \mu(0),$$

which completes the proof of claim 1°.

2° Let  $\nu \in [W_{\beta}(1), W_{\beta}(0)]$ . We prove that the policy  $\pi$  with activity set

$$\mathcal{B}^{\pi} = \{0\},\$$

according to which user k is scheduled just in state 0, is  $(\nu, \beta)$ -optimal. It remains to prove that policy  $\pi$  is optimal in any state  $n \in \{0, 1, ...\}$ .

We start the proof by first deriving the value function  $V^{\pi}_{\beta}(n; \nu)$  for policy  $\pi$  from the following equations:

$$V_{\beta}^{\pi}(0;\nu) = h + \theta d + \nu + \beta(1-\theta)(1-\mu(0))V_{\beta}^{\pi}(1;\nu),$$
  

$$V_{\beta}^{\pi}(n;\nu) = h + \theta d + \beta(1-\theta)V_{\beta}^{\pi}(n;\nu), \quad n \in \{1, 2, \ldots\}.$$
(18)

The unique solution of these linear equations is given by

$$V_{\beta}^{\pi}(0; \nu) = \frac{h + \theta d}{1 - \beta(1 - \theta)} \left( 1 - \beta(1 - \theta)\mu(0) \right) + \nu,$$
  

$$V_{\beta}^{\pi}(n; \nu) = \frac{h + \theta d}{1 - \beta(1 - \theta)}, \quad n \in \{1, 2, \ldots\}.$$
(19)

By (19), the following condition for optimality of  $\pi$  in state 0,

$$\beta(1-\theta)V_{\beta}^{\pi}(0;\nu) \ge \nu + \beta(1-\theta)(1-\mu(0))V_{\beta}^{\pi}(1;\nu),$$

is easily shown to be equivalent with

$$\nu \le (h + \theta d) \,\frac{\beta(1-\theta)}{1 - \beta(1-\theta)} \,\mu(0),\tag{20}$$

where the right hand side equals  $W_{\beta}(0)$  by (10).

Let then  $n \in \{1, 2, ...\}$ . Again, by (19), the following condition for optimality of  $\pi$  in state n,

$$\beta(1-\theta)V_{\beta}^{\pi}(n;\nu) \le \nu + \beta(1-\theta)(1-\mu(n))V_{\beta}^{\pi}(n+1;\nu),$$

is easily be shown to be equivalent with condition

$$\nu \ge (h + \theta d) \,\frac{\beta(1 - \theta)}{1 - \beta(1 - \theta)} \,\mu(n). \tag{21}$$

Deringer

This condition is satisfied due to the DHR assumption, since we assumed that

$$\nu \ge W_{\beta}(1) = (h + \theta d) \frac{\beta(1 - \theta)}{1 - \beta(1 - \theta)} \mu(1),$$

which completes the proof of claim 2°.

3° Let  $m \in \{1, 2, ...\}$  and  $\nu \in [W_{\beta}(m + 1), W_{\beta}(m)]$ . We prove that the policy  $\pi$  with activity set

$$\mathcal{B}^{\pi} = \{0, 1, \ldots, m\},\$$

according to which user k is scheduled in states  $\{0, 1, ..., m\}$ , is  $(\nu, \beta)$ -optimal. It remains to prove that policy  $\pi$  is optimal in any state  $n \in \{0, 1, ...\}$ .

We start the proof by first deriving the value function  $V^{\pi}_{\beta}(n; \nu)$  for policy  $\pi$  from the following equations:

$$V_{\beta}^{\pi}(n;\nu) = h + \theta d + \nu + \beta(1-\theta)(1-\mu(n))V_{\beta}^{\pi}(n+1;\nu), \quad n \in \{0, 1, \dots, m\}, V_{\beta}^{\pi}(n;\nu) = h + \theta d + \beta(1-\theta)V_{\beta}^{\pi}(n;\nu), \quad n \in \{m+1, m+2, \dots\}.$$
(22)

The unique solution of these linear equations is given by

$$V_{\beta}^{\pi}(n;\nu) = \frac{h+\theta d}{1-\beta(1-\theta)} \left(1 - \sum_{i=0}^{m-n} \beta^{i+1}(1-\theta)^{i+1} \prod_{j=0}^{i-1} (1-\mu(n+j)) \mu(n+i)\right) + \nu \left(\sum_{i=0}^{m-n} \beta^{i}(1-\theta)^{i} \prod_{j=0}^{i-1} (1-\mu(n+j))\right), \quad n \in \{0, 1, \dots, m\},$$

$$V_{\beta}^{\pi}(n;\nu) = \frac{h+\theta d}{1-\beta(1-\theta)}, \quad n \in \{m+1, m+2, \dots\}.$$
(23)

By (23), the following condition for optimality of  $\pi$  in state m,

$$\beta(1-\theta)V_{\beta}^{\pi}(m;\nu) \ge \nu + \beta(1-\theta)(1-\mu(m))V_{\beta}^{\pi}(m+1;\nu),$$

is easily shown to be equivalent with

$$\nu \le (h + \theta d) \frac{\beta(1 - \theta)}{1 - \beta(1 - \theta)} \,\mu(m),\tag{24}$$

where the right hand side equals  $W_{\beta}(m)$  by (10).

Let then  $n \in \{m + 1, m + 2, ...\}$ . Again, by (23), the following condition for optimality of  $\pi$  in state n,

$$\beta(1-\theta)V_{\beta}^{\pi}(n;\nu) \leq \nu + \beta(1-\theta)(1-\mu(n))V_{\beta}^{\pi}(n+1;\nu),$$

D Springer

is easily be shown to be equivalent with condition

$$\nu \ge (h + \theta d) \frac{\beta(1 - \theta)}{1 - \beta(1 - \theta)} \mu(n).$$
<sup>(25)</sup>

This condition is satisfied due to the DHR assumption, since we assumed that

$$\nu \ge W_{\beta}(m+1) = (h+\theta d) \frac{\beta(1-\theta)}{1-\beta(1-\theta)} \mu(m+1).$$

It remains to prove that policy  $\pi$  is optimal in any state  $n \in \{0, 1, ..., m-1\}$  given that  $\nu \in [W_{\beta}(m+1), W_{\beta}(m)]$ . Let  $n \in \{0, 1, ..., m-1\}$ . By (23), the following condition for optimality of  $\pi$  in state n,

$$\beta(1-\theta)V_{\beta}^{\pi}(n;\nu) \ge \nu + \beta(1-\theta)(1-\mu(n))V_{\beta}^{\pi}(n+1;\nu),$$

can, by straightforward but tedious manipulations, be shown to be equivalent with condition

$$\nu \le (h+\theta d) \, \frac{\beta(1-\theta)}{1-\beta(1-\theta)} \, \frac{\sum_{i=0}^{m-n} \beta^i (1-\theta)^i \prod_{j=0}^{i-1} (1-\mu(n+j)) \, \mu(n+i)}{\sum_{i=0}^{m-n} \beta^i (1-\theta)^i \prod_{j=0}^{i-1} (1-\mu(n+j))}.$$

This condition clearly follows from (24) due to the DHR assumption, which completes the proof of claim  $3^{\circ}$ .

4° Let  $\nu \in [0, W_{\beta}(\infty)]$ . We prove that the policy  $\pi$  with activity set

$$\mathcal{B}^{\pi} = \{0, 1, \ldots\},\$$

according to which user k is scheduled in states  $\{0, 1, ...\}$ , is  $(\nu, \beta)$ -optimal. It remains to prove that policy  $\pi$  is optimal in any state  $n \in \{0, 1, ...\}$ .

We start the proof by first deriving the value function  $V^{\pi}_{\beta}(n; \nu)$  for policy  $\pi$  from the following equations:

$$V_{\beta}^{\pi}(n;\nu) = h + \theta d + \nu + \beta(1-\theta)(1-\mu(n))V_{\beta}^{\pi}(n+1;\nu), \quad n \in \{0, 1, \ldots\}.$$
(26)

The unique solution of these linear equations is given by

$$V_{\beta}^{\pi}(n;\nu) = \frac{h+\theta d}{1-\beta(1-\theta)} \left(1 - \sum_{i=0}^{\infty} \beta^{i+1}(1-\theta)^{i+1} \prod_{j=0}^{i-1} (1-\mu(n+j)) \mu(n+i)\right) + \nu \left(\sum_{i=0}^{\infty} \beta^{i}(1-\theta)^{i} \prod_{j=0}^{i-1} (1-\mu(n+j))\right), \quad n \in \{0,1,\ldots\}.$$
(27)

🖄 Springer

Let  $n \in \{0, 1, ...\}$ . By (27), the following condition for optimality of  $\pi$  in state n,

$$\beta(1-\theta)V_{\beta}^{\pi}(n;\nu) \ge \nu + \beta(1-\theta)(1-\mu(n))V_{\beta}^{\pi}(n+1;\nu),$$

can, by straightforward but tedious manipulations, be shown to be equivalent with condition

$$\nu \le (h+\theta d) \, \frac{\beta(1-\theta)}{1-\beta(1-\theta)} \, \frac{\sum_{i=0}^{\infty} \beta^i (1-\theta)^i \prod_{j=0}^{i-1} (1-\mu(n+j)) \, \mu(n+i)}{\sum_{i=0}^{\infty} \beta^i (1-\theta)^i \prod_{j=0}^{i-1} (1-\mu(n+j))}.$$

This condition is satisfied due to the DHR assumption, since we assumed that

$$\nu \leq W_{\beta}(\infty) = (h + \theta d) \frac{\beta(1 - \theta)}{1 - \beta(1 - \theta)} \mu(\infty),$$

which completes the proof of claim 4°.

5° Finally, let  $\nu \in (-\infty, 0]$ . We prove that the policy  $\pi$  with activity set

$$\mathcal{B}^{\pi}=\mathcal{S},$$

according to which user k is scheduled in all states, is  $(v, \beta)$ -optimal. It remains to prove that policy  $\pi$  is optimal in any state  $n \in \{0, 1, ...\}$ .

We start the proof by first deriving the value function  $V^{\pi}_{\beta}(n; \nu)$  for policy  $\pi$  from the following equations:

$$V_{\beta}^{\pi}(n;\nu) = h + \theta d + \beta \theta \frac{\nu}{1-\beta} + \nu + \beta (1-\theta) \Big( \mu(n) \frac{\nu}{1-\beta} + (1-\theta)(1-\mu(n)) V_{\beta}^{\pi}(n+1;\nu) \Big), \quad (28)$$
  
$$n \in \{0, 1, \ldots\}.$$

The unique solution of these linear equations is given by

$$V_{\beta}^{\pi}(n;\nu) = \frac{h+\theta d}{1-\beta(1-\theta)} \left(1 - \sum_{i=0}^{\infty} \beta^{i+1}(1-\theta)^{i+1} \prod_{j=0}^{i-1} (1-\mu(n+j)) \mu(n+i)\right) + \frac{\nu}{1-\beta}, \quad n \in \{0, 1, \ldots\}.$$
(29)

Let  $n \in \{0, 1, ...\}$ . By (29), the following condition for optimality of  $\pi$  in state n [based on (14)],

$$\beta(1-\theta)V_{\beta}^{\pi}(n;\nu) \ge \nu + \beta(1-\theta)\Big(\mu(n)\frac{\nu}{1-\beta} + (1-\mu(n))V_{\beta}^{\pi}(n+1;\nu)\Big),$$

can, by straightforward but tedious manipulations, be shown to be equivalent with condition

$$\nu \le (h+\theta d)\beta(1-\theta)\sum_{i=0}^{\infty}\beta^i(1-\theta)^i\prod_{j=0}^{i-1}(1-\mu(n+j))\,\mu(n+i),$$

which follows from our assumption that  $\nu \le 0$  since the right hand side of the inequality above is clearly nonnegative. This completes the proof of claim 5°.

From 1° to 5° together, we deduce that the relaxed optimization problem with objective function (4) is indexable in any state  $x \in S$  and the corresponding Whittle index is given by (10).

### 5 Whittle index policy

In this section, we move from discounted to undiscounted costs. In addition, we move from the discrete-time setup to the original open version of the scheduling problem in continuous time, and determine the Whittle index policy in this setting.

Let us first consider *undiscounted* costs in the discrete-time model. Let  $W_{k,1}(n)$  denote the Whittle index for customer k in state n related to the undiscounted costs, which is derived from the discounted-cost Whittle index  $W_{k,\beta}(n)$  as follows:

$$W_{k,1}(n) = \lim_{\beta \to 1} W_{k,\beta}(n), \quad n \in \{0, 1, \ldots\}.$$

By (10), we have

$$W_{k,1}(n) = \left(\frac{h_k}{\theta_k} + d_k\right) (1 - \theta_k) \,\mu_k(n), \quad n \in \{0, 1, \dots, \},\tag{30}$$

assuming that the discrete hazard rate  $\mu_k(n)$  of customer k satisfies the DHR assumption (8).

When considering the continuous-time model related to the open version of the scheduling problem, the DHR assumption (8) reads as follows:

$$\mu_k(x) \ge \mu_k(y) \quad \text{for all } x < y, \tag{31}$$

i.e., the (continuous-time) hazard rate function  $\mu_k(x)$  is decreasing for all x > 0 and all customer classes *k*.

The continuous-time counterpart of the Whittle index is now determined by letting the time slot shrink down to 0 while, at the same time, scaling the service completion and abandonment probabilities accordingly. As the result, we get the following expression for the continuous-time Whittle index  $W_k(a)$  for customer class-k, where a > 0 refers to the attained service of the customer:

$$W_k(a) = \left(\frac{h_k}{\theta_k} + d_k\right) \mu_k(a), \quad a > 0.$$
(32)

🖄 Springer

Note that, in this continuous-time model, the holding cost rate  $h_k$  is given per time unit while the abandonment penalty  $d_k$  is still a lump sum. Geometric abandonment times are replaced by exponentially distributed abandonment times so that  $\theta_k$  now refers to the abandonment intensity per time unit (instead of the abandonment probability per time slot). In addition, the discrete hazard rate is replaced by the continuous-time hazard rate function  $\mu_k(a)$ . Note also the similarity with the corresponding Gittins index for the DHR service times [1]:

$$G_k(a) = h_k \mu_k(a), \quad a > 0.$$
 (33)

**Definition 2** Consider the original continuous-time scheduling problem with *M* homogeneous parallel servers and exponentially distributed abandonments. Assume that the hazard rate functions  $\mu_k(x)$  of all customer classes *k* satisfy the DHR assumption (31). At any time *t*, the *Whittle index policy* (WHI) chooses to serve

- (i) all the customers, if there are at most *M* customers in the system;
- (ii) those *M* customers that have the highest Whittle indexes  $W_k(a_k)$  given in (32), if there are more than *M* customers in the system.

As proposed in [20], the corresponding *Gittins index policy* (GIT) for the multiserver case is defined similarly, just replacing the comparison of Whittle indexes  $W_k(a)$  by the comparison of Gittins indexes  $G_k(a_k)$  given in (33).

Note also that, in the special case, where customer class k has an exponential service time distribution with mean  $1/\mu_k$ , the Whittle index  $W_k(a)$  is just constant,

$$W_k(a) = \left(\frac{h_k}{\theta_k} + d_k\right) \mu_k, \quad a > 0.$$
(34)

Thus, if all customer classes have exponential distributions, the Whittle index policy (WHI) is equivalent to the  $c\mu/\theta$  rule (CMTH) derived in [2, 3, 15, 23].

On the other hand, if there is just one customer class (K = 1) with a DHR service time distribution, both WHI and GIT reduce to the Forward–Backward policy (FB),<sup>6</sup> which is known to be the optimal scheduling policy in the ordinary single-server M/G/1 queue with DHR service times but without any abandonments [1, 18].

In the following section, we compare, by numerical simulations, the performance of the Whittle index policy (WHI) to the GIT, CMTH, and FB schedulers mentioned above. In addition, we include ordinary (nonanticipating) policies First-Come-First-Served (FCFS) and Processor-Sharing (PS) in our comparison, as well as the anticipating policy Earliest-Deadline-First (EDF). Note, however, that the EDF policy, which is known to be the optimal (anticipating) policy when service times are exponential and only the abandonment penalties are taken into account [17], utilizes the remaining life time information, which is not available for the nonanticipating policy WHI.

<sup>&</sup>lt;sup>6</sup> Scheduling policy FB chooses always the customer with the least attained service. It is also known as the Least-Attained-Service policy (LAS).

## **6 Numerical results**

In this section, we evaluate, by numerical simulations, the performance of the Whittle index policy (WHI) in the M/G/M queue with M servers, DHR service times, and exponentially distributed abandonments. We demonstrate that WHI gives systematically better performance than the nonanticipating policies CMTH, FB, PS, and FCFS mentioned in Sect. 5. As for the Gittins index policy (GIT), the difference in the performance is smaller, but still WHI appears to be better more often than GIT. In addition, we find that WHI even outperforms the anticipating policy EDF, when both holding costs and abandonment penalties are taken into account.

In our numerical simulations, the customers arrive according to a Poisson process at rate  $\lambda > 0$ , and they have independent service times. There may be multiple customer classes, but the service time distribution in each class *k* is of type DHR satisfying the continuous-time version (31) of the DHR assumption. In addition, the abandonment times of the customers in class *k* are independent and exponentially distributed with intensity  $\theta_k > 0$ . Note also that, due to abandonments, our queueing system is stable for any  $\lambda > 0$ .

#### 6.1 Single customer class

First we consider the single-server case (M = 1) where all the customers belong to the same class. For the service time, we apply the *Weibull distribution* with shape parameter  $\alpha > 0$  and the scale parameter  $\gamma > 0$ , for which the cumulative distribution function reads as

$$F(x) = 1 - e^{-(\gamma x)^{\alpha}}, \quad x > 0,$$

and the hazard rate function is given by

$$\mu(x) = \alpha \gamma (\gamma x)^{\alpha - 1}, \quad x > 0,$$

which is a decreasing function of x whenever  $0 < \alpha \le 1$ . In addition, the mean service time depends on the parameters as follows:

$$E[S] = \frac{1}{\gamma} \Gamma\left(1 + \frac{1}{\alpha}\right),$$

where  $\Gamma(\cdot)$  refers to the Euler gamma function. Note that the Weibull distribution with shape parameter  $\alpha = 1$  is, in fact, the same as the exponential distribution with mean  $1/\gamma$ .

In our numerical experiments, we let the shape parameter  $\alpha$  vary but choose the scale parameter  $\gamma = \Gamma(1 + \frac{1}{\alpha})$  so that E[S] = 1 for any  $\alpha$ . In addition, we have fixed the abandonment rate to be  $\theta = 1/8$  so that the mean abandonment time is eight times longer than the mean service time. Holding costs are accrued at rate h = 1, and the abandonment penalty cost takes value d = 10.

Since there is just a single class of customers with DHR service times, the Whittle index policy (WHI) (as well as the Gittins index policy (GIT)) is the same as the FB policy. Based on simulations, we have estimated its performance against policies FCFS, PS, and EDF with varying inverse shape parameter values  $(1/\alpha \in [1, 3])$  and two different loads ( $\lambda \in \{1, 2\}$ ). In each simulation run with fixed model parameters and scheduling policy, we have gathered the system statistics until there are  $10^6$  customer arrivals. The results for the lower load ( $\lambda = 1$ ) are shown in Fig. 1, and for the higher load ( $\lambda = 2$ ) in Fig. 2. In these figures, the mean total/holding/abandonment costs refer to the estimated mean total/holding/abandonment costs per time unit (and not per customer).

As seen from these figures, when  $1/\alpha = 1$  (i.e., the service times are exponential), the mean total costs, holding costs, and abandonment costs are essentially the same for the nonanticipating policies FCFS, PS, and WHI, which is in line with theory. For any  $1/\alpha > 1$  and for both load levels ( $\lambda \in \{1, 2\}$ ), WHI is systematically better than FCFS and PS both in the mean total cost sense as well as for the two components (holding and abandonment costs) separately. In addition, the performance of PS is much closer to that of WHI than the performance of FCFS.

When comparing WHI to EDF, we see that EDF typically produces smaller mean abandonment costs than WHI, which is also in line with theory. However, for the higher load ( $\lambda = 2$ ), the difference is very small and almost within the random variations related to the numerical simulations. On the other hand, the mean holding costs of EDF are systematically much higher than those of WHI. As a result, WHI outperforms EDF when both holding costs and abandonment costs are taken into account.

#### 6.2 Multiple customer classes

Next we consider the single-server case (M = 1) where there are two customer classes. Class-1 customers have exponential services times with mean  $E[S_1] = 1$ , and class-2 customers have Weibull services times with a varying inverse shape parameter  $1/\alpha_2 \in [1, 3]$  but with a fixed mean  $E[S_2] = 1$ . The system is studied under two different total loads,  $\lambda \in \{1, 2\}$ , assuming that both classes generate an equal load:  $\lambda_1 = \lambda_2 = \lambda/2$ . The other parameters are chosen as follows:

$$1/\theta_1 = 1/\theta_2 = 8$$
,  $h_1 = h_2 = 1$ ,  $d_1 = 10$ ,  $d_2 = 20$ .

In this case, the Whittle indexes for the two classes are

$$W_1(x) = \left(\frac{h_1}{\theta_1} + d_1\right) \frac{1}{E[S_1]} = 18,$$
  
$$W_2(x) = \left(\frac{h_2}{\theta_2} + d_2\right) \mu_2(x) = 28 \,\mu_2(x),$$

where the decreasing hazard rate of class-2 customers is given by

$$\mu_2(x) = \alpha_2 \gamma_2 (\gamma_2 x)^{\alpha_2 - 1}$$



**Fig. 1** Single-class case, lower load ( $\lambda = 1$ ): simulated average costs for policies FCFS, PS, WHI, and EDF as a function of the inverse Weibull shape parameter  $1/\alpha$ . **a** Mean total costs (top), **b** mean holding costs (middle), **c** mean abandonment costs (bottom)



**Fig. 2** Single-class case, higher load ( $\lambda = 2$ ): simulated average costs for policies FCFS, PS, WHI, and EDF as a function of the inverse Weibull shape parameter  $1/\alpha$ . **a** Mean total costs (top), **b** mean holding costs (middle), **c** mean abandonment costs (bottom)



**Fig.3** Two-class case: Whittle index as a function of the attained service for exponential distribution (green) and Weibull distributions with the inverse shape parameters  $1/\alpha_2 = 2$  (blue) and  $1/\alpha_2 = 3$  (red) (Color figure online)

with  $\gamma_2 = \Gamma(1 + \frac{1}{\alpha_2})$ . Since, for any  $0 < \alpha_2 < 1$ ,

$$\lim_{x \to 0} \mu_2(x) = \infty, \quad \lim_{x \to \infty} \mu_2(x) = 0,$$

there is a unique borderline attained service value  $a^*(\alpha_2) > 0$  such that

$$W_1(a^*(\alpha_2)) = W_2(a^*(\alpha_2)).$$

This is illustrated in Fig. 3, where we see the constant Whittle index  $W_1(x)$  and the decreasing Whittle index  $W_2(x)$  for the inverse shape parameter values  $1/\alpha_2 \in \{2, 3\}$ . The borderline attained services are  $a^*(1/2) = 1.210$  and  $a^*(1/3) = 0.915$ , respectively.

Thus, the Whittle index policy (WHI) is as follows for any  $0 < \alpha_2 < 1$ . Whenever there are class-2 customers with attained service less than the borderline value  $a^*(\alpha_2)$ , they are served in the FB manner. But as soon as the attained services of all class-2 customers are at least the same as the borderline value  $a^*(\alpha_2)$ , the server starts to serve class-1 customers until there are no longer any class-1 customers or a new class-2 customer arrives. At this point, the server returns to serve class-2 customers in the FB manner. Since the Whittle index for any class-1 customer is the same, the way how the server serves these customers can be any work-conserving discipline. In our experiment, we apply the robust PS policy for class-1 customers whenever they are served.

The Gittins index policy (GIT) is very similar. In this case, the unique borderline attained service value  $\tilde{a}(\alpha_2) > 0$  is determined from equation

$$G_1(\tilde{a}(\alpha_2)) = G_2(\tilde{a}(\alpha_2)),$$



**Fig. 4** Two-class case: Gittins index as a function of the attained service for exponential distribution (green) and Weibull distributions with the inverse shape parameters  $1/\alpha_2 = 2$  (blue) and  $1/\alpha_2 = 3$  (red)(Color figure online)

where the Gittins indexes for the two classes are given by

$$G_1(x) = \frac{h_1}{E[S_1]} = 1, \quad G_2(x) = h_2 \mu_2(x) = \mu_2(x).$$

This is also illustrated in Fig. 4, where we see the constant Gittins index  $G_1(x)$  and the decreasing Gittins index  $G_2(x)$  for the inverse shape parameter values  $1/\alpha_2 \in \{2, 3\}$ . The borderline attained services are now  $\tilde{a}(1/2) = 0.500$  and  $\tilde{a}(1/3) = 0.471$ , respectively.

In addition, since

$$\left(\frac{h_1}{\theta_1} + d_1\right)\frac{1}{E[S_1]} = 18, \quad \left(\frac{h_2}{\theta_2} + d_2\right)\frac{1}{E[S_2]} = 28,$$

the  $c\mu/\theta$  rule (CMTH) is in this case a priority policy that gives always full priority to class-2 customers. As mentioned in the previous section, in the special case  $\alpha_2 = 1$ (i.e., with exponential service times in both classes), the Whittle index policy (WHI) is the same as the  $c\mu/\theta$  rule (CMTH). On the other hand, it is shown in [4] that CMTH is asymptotically optimal among non-preemptive scheduling policies for general service times.

Based on simulations, we have estimated the performance of the Whittle index policy (WHI) against policies FCFS, PS, FB, GIT, and CMTH with the varying inverse shape parameter values  $(1/\alpha_2 \in [1, 3])$  and two different loads ( $\lambda \in \{1, 2\}$ ). In each simulation run with fixed model parameters and scheduling policy, we have gathered the system statistics until there are  $10^6$  customer arrivals. The results for the lower load ( $\lambda = 1$ ) are shown in Fig. 5, and for the higher load ( $\lambda = 2$ ) in Fig. 6. In these figures, the mean total costs refer to the estimated mean total costs per time unit (and not per customer). Therefore, the classwise mean total costs sum up to the mean total costs induced by all customers.



**Fig. 5** Two-class case, lower load ( $\lambda = 1$ ): simulated average total costs for policies FCFS, PS, FB, CMTH, WHI, and GIT as a function of the inverse Weibull shape parameter  $1/\alpha_2$ . **a** All customers (top), **b** class-1 customers (middle), **c** class-2 customers (bottom)



**Fig. 6** Two-class case, higer load ( $\lambda = 2$ ): simulated average total costs for policies FCFS, PS, FB, CMTH, WHI, and GIT as a function of the inverse Weibull shape parameter  $1/\alpha_2$ . **a** All customers (top), **b** class-1 customers (middle), **c** class-2 customers (bottom)

As seen from these figures, when  $1/\alpha_2 = 1$ , the mean total costs (for all customers or classwisely) are essentially the same for the policies FCFS, PS, and FB. On the other hand, when  $1/\alpha_2 = 1$ , the mean total costs (for all customers or classwisely) are also essentially the same for the policies WHI and CMTH. Both of these results are in line with theory.

When comparing WHI to FCFS and PS, we make similar observations as earlier in the single-class example: WHI is systematically better than FCFS and PS. In addition, the performance of PS is much closer to that of WHI than the performance of FCFS. When comparing WHI to FB and GIT, we see that WHI is clearly better than FB and GIT when  $1/\alpha_2$  is sufficiently small. For greater values of  $1/\alpha_2$ , WHI still outperforms FB and GIT but the performance difference between the two policies becomes much smaller. For any  $1/\alpha_2 > 1$ , WHI is also systematically better than CMTH. The priority rule CMTH clearly favors class-2 customers giving them much better performance than for the customers in class 1. The performance of WHI mimics that of CMTH when  $1/\alpha_2$  is sufficiently small, but for greater values of  $1/\alpha_2$ , its performance is more comparable to that of FB and GIT.

#### 6.3 Random model parameters

Now we consider the single-server case (M = 1) where we still have two customer classes but we generate certain model parameters randomly and separately for each simulation run to study how systematically the performance of different scheduling policies behaves when compared to each other.

As before, class-1 customers have exponential services times with mean  $E[S_1] = 1$ , while class-2 customers have Weibull services times with mean  $E[S_2] = 1$  and two different inverse shape parameter values,  $1/\alpha_2 \in \{2, 3\}$ . The system is again studied under two different total loads,  $\lambda \in \{1, 2\}$ , assuming that both classes generate an equal load:  $\lambda_1 = \lambda_2 = \lambda/2$ . In addition, we fix the holding costs for both classes  $(h_1 = h_2 = 1)$  but generate the abandonment penalties  $d_1$  and  $d_2$  separately and independently from the uniform distribution over the interval [5, 25] and the mean abandonment times  $1/\theta_1$  and  $1/\theta_2$  separately and independently from the uniform distribution over the interval [4, 12]. Note that, if such a random parameter combination satisfies

$$\frac{1}{\theta_1} + d_1 < \frac{1}{\theta_2} + d_2,$$

the  $c\mu/\theta$  rule (CMTH) is the priority policy that gives always full priority to class-2 customers. Otherwise it gives full priority to class-1 customers.

Based on simulations, we have estimated the performance of the Whittle index policy (WHI) against policies FCFS, PS, FB, GIT, and CMTH with two different loads  $\lambda$ , two different inverse shape parameter values  $1/\alpha_2$ , and N = 50 different random parameter combinations  $(1/\theta_1, d_1, 1/\theta_2, d_2)$ . The same parameter combinations are applied to all competing scheduling policies. In each simulation run with fixed model parameters and scheduling policy, we have gathered the system statistics until there are  $10^6$  customer arrivals. The performance comparison between the six policies is based on the estimated mean total costs for each simulation run separately.

	WHI	GIT	FB	CMTH	PS	FCFS			
$(a) \lambda = 1$									
Rank 1	31	18	1	0	0	0			
Rank 2	17	19	11	3	0	0			
Rank 3	2	9	33	6	0	0			
Rank 4	0	4	5	7	34	0			
Rank 5	0	0	0	34	16	0			
Rank 6	0	0	0	0	0	50			
Avg rank	1.42	1.98	2.84	4.44	4.32	6.00			
Avg cost	4.395	4.425	4.502	4.770	4.660	5.681			
Cost diff	-	0.72%	2.43%	8.49%	6.04%	29.56%			
$(b) \lambda = 2$									
Rank 1	43	7	0	0	0	0			
Rank 2	7	27	11	5	0	0			
Rank 3	0	10	32	8	0	0			
Rank 4	0	6	7	4	33	0			
Rank 5	0	0	0	33	17	0			
Rank 6	0	0	0	0	0	50			
Avg rank	1.14	2.30	2.92	4.30	4.34	6.00			
Avg cost	15.300	15.680	15.974	17.087	16.639	20.991			
Cost diff	—	2.71%	4.53%	11.45%	8.82%	37.40%			

**Table 1** Two-class case, inverse shape parameter  $1/\alpha_2 = 2$ : Comparison of policies WHI, GIT, FB, CMTH, PS, and FCFS with N = 50 random parameter combinations. (a) Lower load ( $\lambda = 1$ ), (b) higher load ( $\lambda = 2$ )

## *Inverse shape parameter* $1/\alpha_2 = 2$ .

Let us first consider the results related to the inverse shape parameter value  $1/\alpha_2 = 2$ . The results for the two loads ( $\lambda \in \{1, 2\}$ ) are presented in Table 1. After the title row, the next six rows give the number of different ranks in the policy comparison covering the 50 parameter combinations. Thus, the second row (Rank 1) indicates how many times (out of 50) the policy (in a fixed column) has been the best one among the six policies, while the seventh row (Rank 6) expresses how many times it has had the worst estimated performance. The following row (Avg rank) gives the average rank, and the next one (Avg cost) the average estimated mean total costs over the 50 different parameter combinations. Finally, in the last row (Cost diff), we have calculated the relative difference in the average estimated mean total costs between WHI and the other six policies.

From Table 1, we see that, overall, the Whittle index policy (WHI) outperforms the other policies having 31 first positions, 17 second positions, and 2 third positions in the lower load. In the higher load, the difference is even more clear with 43 first positions and 7 second positions for WHI. The second best policy is clearly GIT taking even 18 first positions in the lower load and 7 of them in the higher load. In the lower load, the difference between WHI and GIT seems to be quite small. FB, which is the third

best policy, takes mostly the third position. The  $c\mu/\theta$  rule (CMTH) has some second positions in both loads but usually it takes just the fifth position, while the PS policy is usually the fourth one. The worst performance in all the trials is generated by the FCFS policy. In addition, we observe that the relative difference between WHI and the other policies is systematically larger in the higher load (when compared to the lower load).

#### Inverse shape parameter $1/\alpha_2 = 3$ .

Now we consider the results related to the inverse shape parameter value  $1/\alpha_2 = 3$ . The results for the two loads ( $\lambda \in \{1, 2\}$ ) are presented in Table 2. The structure of Table 2 is the same that of Table 1.

In general, the results look very similar to the previous case  $(1/\alpha_2 = 2)$ . From Table 2, we see that the Whittle index policy (WHI) again outperforms the other policies having now 30 first positions, 19 second positions, and only 1 third position in the lower load. In the higher load, the difference is again more apparent with 42 first positions and 8 second positions for WHI. The performance order of the other policies remains the same: the second best is GIT, then come FB, PS, and CMTH, and the worst one is again FCFS. In the lower load, the difference between WHI and GIT seems to be even smaller when compared to the previous case. However, in the higher load, the relative difference between WHI and all the other policies is again systematically larger.

#### 6.4 Multiple servers

Finally, we consider the case with multiple parallel servers  $(M \ge 1)$ . All the customers belong to the same class, and they have Weibull services times with mean E[S] = 1 and inverse shape parameter  $1/\alpha = 2$ . The system is studied under a varying number of parallel servers,  $M \in \{1, 2, ..., 8\}$ , and two different relative loads,  $\lambda/M \in \{1, 2\}$ . The other parameters are chosen as follows:

$$1/\theta = 8, \quad h = 1, \quad d = 10.$$

Since there is just a single class of customers, the Whittle index policy (WHI) (as well as the Gittins index policy (GIT)) is the same as the FB policy. Based on simulations, we have estimated its performance against policies FCFS, PS, and EDF. As before, the system statistics have been gathered until there are  $10^6$  customer arrivals. The results for the lower relative load ( $\lambda/M = 1$ ) are shown in Fig. 7, and for the higher relative load ( $\lambda/M = 2$ ) in Fig.8. In these figures, the mean total/holding/abandonment costs refer to the estimated mean total/holding/abandonment costs per time unit.

As seen from Fig. 7, the performance order of the four policies remains the same for any  $M \in \{1, 2, ..., 8\}$  in the lower relative load: WHI being the best and FCFS the worst. The absolute difference in mean total costs is not changing much for the PS and FCFS policies (when compared to WHI), while, for the EDF policy, the difference narrows as M increases. As for the cost components, EDF minimizes the mean abandonment costs for any M, but its mean holding costs are so high that it just takes the third position in the overall comparison.

					-	
	WHI	GIT	FB	CMTH	PS	FCFS
$(a) \lambda = 1$						
Rank 1	30	19	1	0	0	0
Rank 2	19	24	7	0	0	0
Rank 3	1	7	41	1	0	0
Rank 4	0	0	1	3	46	0
Rank 5	0	0	0	46	4	0
Rank 6	0	0	0	0	0	50
Avg rank	1.42	1.76	2.84	4.90	4.08	6.00
Avg cost	3.551	3.564	3.628	4.082	3.778	5.157
Cost diff	-	0.40%	2.11%	15.04%	6.34%	46.03%
$(b) \lambda = 2$						
Rank 1	42	8	0	0	0	0
Rank 2	8	33	9	0	0	0
Rank 3	0	9	41	0	0	0
Rank 4	0	0	0	4	46	0
Rank 5	0	0	0	46	4	0
Rank 6	0	0	0	0	0	50
Avg rank	1.16	2.02	2.82	4.92	4.08	6.00
Avg cost	11.742	11.892	12.207	14.530	12.816	18.895
Cost diff	-	1.34%	3.87%	23.35%	9.02%	61.56%

**Table 2** Two-class case, inverse shape parameter  $1/\alpha_2 = 3$ : Comparison of policies WHI, GIT, FB, CMTH, PS, and FCFS with N = 50 random parameter combinations. (a) Lower load ( $\lambda = 1$ ), (b) higher load ( $\lambda = 2$ )

In the higher relative load, the results are somehow different, as seen from Fig. 8. The Whittle index policy (WHI) is still consistently the best when the mean total costs are compared, but now EDF proves to be the worst. While it is still minimizing the mean abandonment costs for any M, the mean holding costs are now much worse than for any other policy in this comparison. Another difference to the lower relative load is that the absolute difference in the mean total costs is clearly increasing for all policies (when compared to WHI) as M increases.

# 7 Conclusions

We considered the optimal scheduling problem in a multiserver queue with impatient customers belonging to multiple classes with an objective function that takes into account both holding costs and abandonment penalties. Many papers consider this scheduling problem under Poisson arrivals and linear holding costs assuming further that both the service times and the abandonment times have exponential distributions. Even with these additional assumptions, the exact solution is known only in very few special cases.

To find a reasonable heuristic solution for this tricky problem, we applied the Whittle index approach. Unlike in the earlier papers, which were restricted to exponential



**Fig. 7** Multiserver case, lower relative load  $(\lambda/M = 1)$ : simulated average costs for policies FCFS, PS, WHI, and EDF as a function of the number of servers *M*. **a** Mean total costs (top), **b** mean holding costs (middle), **c** mean abandonment costs (bottom)



**Fig. 8** Multiserver case, higher relative load ( $\lambda/M = 2$ ): simulated average costs for policies FCFS, PS, WHI, and EDF as a function of the number of servers *M*. **a** Mean total costs (top), **b** mean holding costs (middle), **c** mean abandonment costs (bottom)

service times, we allowed the service time distributions for which the hazard rate is decreasing. We considered the discrete-time multiserver queueing problem with discounted costs. As our main theoretical result, we proved that the related relaxed optimization problem indexable and derived the corresponding Whittle index explicitly.

Based on this discrete-time result, we developed the Whittle index policy for the original continuous-time multiserver scheduling problem. The performance of the resulting policy was first evaluated in the M/G/1 setup by numerical simulations, which demonstrated that it gives better performance than the other policies included in the comparison. The nonanticipating Whittle index policy even outperformed the anticipating EDF policy when both holding costs and abandonment penalties were taken into account. In addition, we evaluated the performance of the Whittle index policy in the multiserver setting. Again it outperformed consistently the other policies included in the comparison.

In future, we would like to study whether it is possible to derive, by the Whittle index approach, a reasonable scheduling policy for a multiserver queue with impatient customers where the service time distributions are even more general. Another line of research is to allow more general holding costs.

Funding Open Access funding provided by Aalto University.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit http://creativecommons.org/licenses/by/4.0/.

## References

- Aalto, S., Ayesta, U., Righter, R.: On the Gittins index in the M/G/1 queue. Queueing Syst. 63, 437–458 (2009)
- 2. Atar, R., Giat, C., Shimkin, N.: The  $c\mu/\theta$  rule for many-server queues with abandonment. Oper. Res. **58**, 1427–1439 (2010)
- 3. Atar, R., Giat, C., Shimkin, N.: On the asymptotic optimality of the  $c\mu/\theta$  rule under ergodic costs. Queueing Syst. **67**, 127–144 (2011)
- Atar, R., Kaspi, H., Shimkin, N.: Fluid limits for many-server systems with reneging under a priority policy. Math. Oper. Res. 39, 672–696 (2014)
- Ayesta, U., Jacko, P., Novak, V.: A nearly-optimal index rule for scheduling of users with abandonment. In: Proc. of IEEE Infocom, pp. 2849–2857 (2011)
- Ayesta, U., Jacko, P., Novak, V.: Scheduling of multi-class multi-server queueing systems with abandonments. J. Sched. 20, 129–145 (2017)
- Bhulai, S., Blok, H., Spieksma, F.: K competing queues with customer abandonment: optimality of a generalised cμ-rule by the smoothed rate truncation method. Ann. Oper. Res. 317, 387–416 (2022)
- Chen, G., Gayon, J., Lemairec, P.: Stochastic scheduling with abandonment: necessary and sufficient conditions for the optimality of a strict priority policy. Oper. Res. 71(5), 1789–1793 (2023)
- 9. Cox, D., Smith, W.: Queues. Methuen (1961)
- Down, D., Koole, G., Lewis, M.: Dynamic control of a single-server system with abandonments. Queueing Syst. 67, 63–90 (2011)

- 11. Gittins, J.: Multi-armed Bandit Allocation Indices. Wiley, Hoboken (1989)
- Kim, J., Randhawa, R., Ward, A.: Dynamic scheduling in a many-server, multiclass system: the role of customer impatience in large systems. Manuf. Serv. Oper. Manag. 20, 285–301 (2017)
- Kim, J., Ward, A.: Dynamic scheduling of a GI/GI/1+GI queue with multiple customer classes. Queueing Syst. 75, 339–384 (2013)
- Larrañaga, M., Ayesta, U., Verloop, I.: Dynamic fluid-based scheduling in a multi-class abandonment queue. Perform. Eval. 70, 841–858 (2013)
- Larrañaga, M., Ayesta, U., Verloop, I.: Asymptotically optimal index policies for an abandonment queue with convex holding cost. Queueing Syst. 81, 99–169 (2015)
- 16. Long, Z., Shimkin, N., Zhang, H., Zhang, J.: Dynamic scheduling of multiclass many-server queues with abandonment: the generalized  $c\mu/h$  rule. Oper. Res. **68**, 1218–1230 (2020)
- Moyal, P.: On queues with impatience: stability, and the optimality of earliest deadline first. Queueing Syst. 75, 211–242 (2013)
- Nuyens, M., Wierman, A.: The foreground-background queue: a survey. Perform. Eval. 65, 286–307 (2004)
- 19. Ross, S.: Applied Probability Models with Optimization Applications. Holden-Day, Toronto (1970)
- Scully, Z., Grosof, I., Harchol-Balter, M.: The Gittins policy is nearly optimal in the M/G/k under extremely general conditions. Proc. ACM Meas. Anal. Comput. Syst. 4, 1–29 (2020)
- 21. Scully, Z., Harchol-Balter, M.: The Gittins policy in the M/G/1 queue. In: Proc. of WiOpt (2021)
- Shaked, M., Shanthikumar, J., Valdez-Torres, J.: Discrete hazard rate functions. Comput. Oper. Res. 22, 391–402 (1995)
- Verloop, I.: Asymptotically optimal priority policies for indexable and nonindexable restless bandits. Ann. Appl. Probab. 26, 1947–1995 (2016)
- 24. Weber, R., Weiss, G.: On an index policy for restless bandits. J. Appl. Probab. 27, 637–648 (1990)
- 25. Whittle, P.: Arm-acquiring bandits. Ann. Probab. 9, 284–292 (1981)
- Whittle, P.: Restless bandits: activity allocation in a changing world. J. Appl. Probab. 25A, 287–298 (1988)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.