

Index Policies for a Multi-Class Queue with Convex Holding Cost and Abandonments*

M. Larrañaga^{1,2,5}, U. Ayesta^{2,3,4,5}, I.M. Verloop^{1,5}

¹CNRS, IRIT, 2 rue C. Carmichel, F-31071 Toulouse, France.

²CNRS, LAAS, 7 avenue du colonel Roche, F-31400 Toulouse, France

³IKERBASQUE, Basque Foundation for Science, 48011 Bilbao, Spain

⁴UPV/EHU, University of the Basque Country, 20018 Donostia, Spain

⁵Univ. de Toulouse, INP, LAAS, F-31400 Toulouse, France

ABSTRACT

We investigate a resource allocation problem in a multi-class server with convex holding costs and user impatience under the average cost criterion. In general, the optimal policy has a complex dependency on all the input parameters and state information. Our main contribution is to derive index policies that can serve as heuristics and are shown to give good performance. Our index policy attributes to each class an index, which depends on the number of customers currently present in that class. The index values are obtained by solving a relaxed version of the optimal stochastic control problem and combining results from restless multi-armed bandits and queueing theory. They can be expressed as a function of the steady-state distribution probabilities of a one-dimensional birth-and-death process. For linear holding cost, the index can be calculated in closed-form and turns out to be independent of the arrival rates and the number of customers present. In the case of no abandonments and linear holding cost, our index coincides with the $c\mu$ -rule, which is known to be optimal in this simple setting. For general convex holding cost we derive properties of the index value in limiting regimes: we consider the behavior of the index (i) as the number of customers in a class grows large, which allows us to derive the asymptotic structure of the index policies, and (ii) as the abandonment rate vanishes, which allows us to retrieve an index policy proposed for the multi-class M/M/1 queue with convex holding cost and no abandonments. In fact, in a multi-server environment it follows from recent advances that the index policy is asymptotically optimal for linear holding cost. To obtain further insights into the index policy, we consider the fluid version of the relaxed problem and derive a closed-form expression for the fluid index. The latter coincides with the stochastic model

*The PhD fellowship of Maialen Larrañaga is funded by a research grant of the Foundation Airbus Group (<http://fondation.airbus-group.com/>).

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGMETRICS'14, June 16–20, 2014, Austin, Texas, USA.
Copyright 2014 ACM 978-1-4503-2789-3/14/06 ...\$15.00.
<http://dx.doi.org/10.1145/2591971.2591983>.

in case of linear holding costs. For arbitrary convex holding cost the fluid index can be seen as the $Gc\mu/\theta$ -rule, that is, including abandonments into the generalized $c\mu$ -rule ($Gc\mu$ -rule). Numerical experiments show that our index policies become optimal as the load in the system increases.

1. INTRODUCTION

In this paper our objective is to develop a unifying framework to obtain well performing control policies in a multi-class single-server queue with convex holding costs and impatient customers. The single-server queue is the canonical model to study resource allocation problems and it can be considered as one of the most classical decision problems. It has been widely studied due to its applicability to any situation where a single-resource is shared by multiple concurrent customers. Abandonment or renegeing takes place when customers, unsatisfied of their long waiting time, decide to voluntarily leave the system. It has a huge impact in various real life applications such as the Internet or call centers, where customers may abandon while waiting in the queue, or even while being served. In the presence of abandonments and/or convex holding cost, a characterization of the optimal control is out of reach.

When the holding costs are linear and customers are not impatient, a classical result shows that the celebrated $c\mu$ -rule rule is optimal, that is, to serve the classes in decreasing order of priority according to the product $c_k\mu_k$, where c_k is the holding cost per class- k customer, and μ_k^{-1} is the mean service requirement of class- k customers, [13, 18]. The $c\mu$ -rule is a so-called *index policy*, that is, the solution to the stochastic control problem is characterized by an *index*, $c_k\mu_k$, which determines which customer is optimal to serve. This simple structure of the optimal policy vanishes however in the presence of convex costs and/or impatient customers. The optimal policy will in general be a complex function of all the input parameters function and the number of customers present in all the classes.

Optimality of index policies has enjoyed a great popularity. The solution to a complex control problem that, a priori, might depend on the entire state space, turns out to have a strikingly simple structure. For instance, in the case of the $c\mu$ -rule, the solution does not depend on the number of customers in the various classes. Another classical result that can be seen as an index policy is the optimality of Shortest-Remaining-Processing-Time (SRPT), where the index of each customer is given by its remaining service

time. Both examples fit the general context of Multi-Armed Bandit Problems (MABP). A MABP is a particular case of a Markov Decision Process: at every decision epoch the scheduler needs to select one *bandit*, and an associated reward is accrued. The state of this selected bandit evolves stochastically, while the state of all other bandits remain *frozen*. The scheduler knows the state of all bandits, the rewards in every state, and the transition probabilities, and aims at maximizing the total average reward. In a ground-breaking result Gittins showed that the optimal policy that solves a MABP is an index-rule, nowadays commonly referred to as Gittins' index [19]. Thus, for each bandit, one calculates an index that depends only on its own current state and stochastic evolution. The optimal policy activates in each decision epoch the bandit with highest current index.

Despite its generality, in multiple cases of practical interest the problem cannot be cast as a MABP. In a seminal work [35], Whittle introduced the so-called Restless Multi-Armed Bandit Problems (RMABP), a generalization of the standard MABP. In a RMABP all bandits in the system incur a cost. The scheduler selects a number of bandits to be made active. However, all bandits might evolve over time according to a stochastic kernel that depends on whether the bandit is selected for service or not. The objective is to determine a control policy that optimizes the average performance criterion. RMABP provides a more general modeling framework, but its solution has in general a complex structure that might depend on the entire state-space description. Whittle considered a relaxed version of the problem (where the restriction on the number of *active* bandits needs to be respected on average only, and not in every decision epoch), and showed that the solution to the relaxed problem is of index type, referred to as *Whittle's index*. Whittle then defined a heuristic for the original problem where in every decision epoch the bandit with highest Whittle index is selected. It has been shown that the Whittle index policy performs strikingly well, see [28] for a discussion, and can be shown to be asymptotically optimal, see [33, 31]. The latter explains the importance given in the literature to calculate Whittle's index. In order to calculate Whittle's index there are two main difficulties, first one needs to establish a technical property known as *indexability*, and second the calculation of the index might be involved or even infeasible.

In our main contribution of the paper, we verify indexability and calculate Whittle's index for the average cost criterion of the multi-class queue with abandonments and convex cost. In fact, our model can be written as a RMABP where each class is represented by a bandit and the state of a bandit describes the number of customers in that class. The evolution of the number of customers being birth-and-death, the bandit is of birth-and-death type. An important observation we make is that the Whittle index we obtain, which is expressed as a function of the steady-state probabilities, is in fact applicable for any birth-and-death bandit. This is a simple observation that has far reaching consequences since it allows to derive Whittle's index for a general class of control problems, as will be explained in the paper. Note that indexability would be needed to be established on a case-by-case basis. For the abandonment model with convex holding cost, we prove indexability by showing that threshold policies are optimal for the relaxed optimization problem and using properties of the steady-state distributions.

Having characterized Whittle's index in terms of steady-state distributions, we then apply it to various cases. In the case of linear holding cost, we show that the Whittle's index is a constant that is independent of the number of customers in the system and of the arrival rate. In fact, this index policy (with linear holding cost) coincides with the index policies as proposed in [7] and [5], for specific model assumptions, and is asymptotically optimal for a multi-server environment. For general convex holding cost we derive properties of the index value in limiting regimes: we consider the behavior of the index (i) as the number of customers in a class grows large, which allows us to derive the asymptotic structure of the index policies, and (ii) as the abandonment rate vanishes, which allows us to retrieve an index policy proposed for the multi-class M/M/1 queue with convex holding cost and no abandonments.

Our index is expressed as a function of the steady-state probabilities and it can thus efficiently be calculated, but it does not always allow to obtain qualitative insights. We therefore formulate a fluid version of the relaxed optimization problem, where the objective is *bias optimality*, i.e., to determine the policy that minimizes the cost of bringing the fluid to its equilibrium. We show how to derive an index for the fluid model, and we compare it with Whittle's index as obtained for the stochastic model. The advantage of the fluid approach lies in its relatively simple expressions compared to the stochastic one. It shows equivalence with the $Gc\mu/\theta$ -rule, that is, including abandonments into the generalized $c\mu$ -rule ($Gc\mu$ -rule) and provides useful insights on the dependence on the parameters. With linear holding cost the Whittle index and the fluid index are identical.

Numerical experiments show that our index policies work well and become optimal as the load in the system is large.

In summary the main contributions of this paper are:

- Unifying approach to obtain Whittle's index for scheduling problems under average cost criterion.
- For a multi-class queue with convex holding costs and abandonments we prove indexability and obtain Whittle's index as a function of steady-state probabilities.
- For linear holding costs Whittle's index is independent on the arrival rate and state.
- Development of a fluid-based approach to derive a closed-form formula for the indices for general holding cost.

The paper is organized as follows. In Section 2 we give an overview of related work and in Section 3 we describe the model. In Section 4 we present the relaxation of the original problem and show that threshold policies are optimal. We establish indexability and calculate Whittle's index under the average cost criterion. In Section 5 we explain a heuristic index policy, based on Whittle's index, for the original optimization problem. In Section 6 we calculate Whittle's index for linear holding cost and derive properties for general convex holding costs. In Section 7 we calculate the index for an M/M/1 queue without abandonments. Section 8 describes the asymptotic optimality result, and in Section 9 we present the fluid model and derive the fluid index. Finally, in Section 10 we numerically evaluate the performance of Whittle's index policy and the fluid index policy.

2. RELATED WORK

There are four main literature bodies that are relevant to our work: literature on (i) index policies for resource

allocation problems, (ii) scheduling with convex costs, (iii) scheduling in the presence of impatient customers, and (iv) fluid-based scheduling. We provide below a brief summary of some of the main contributions in each of the domains.

(i) The seminal work on the optimality of index policies for MABP is in the book by Gittins et. al. [19]. The optimality of the $c\mu$ -rule, i.e., strict priority is given according to the indices $c\mu$, in a multi-class single server queue for average reward and discounted cost criteria, in the preemptive and non-preemptive cases, is shown in [13, 18]. Index policies for RMABP were introduced in the seminal paper [35]. In [28] the author develops an algorithm that allows to establish whether a problem is indexable, and if yes, to numerically calculate, in an efficient way, Whittle's index. Under the assumption that an ODE has an equilibrium point and that all bandits are symmetric, in [33] it is shown that Whittle's index policy is asymptotically optimal as the number of bandits and the number of bandits that can be made active grow to infinity, while their ratio is kept constant. This result is generalized in [31] to the case in which there are various classes of bandits, and new bandits can arrive over time. In addition to resource allocation problems, Whittle's index has been applied in a wide variety of cases, including opportunistic spectrum access, website morphing, pharmaceutical trials and many others, see for example [19, Chapter 6]. The recent survey paper [21] is a good up-to-date reference on the application of index policies in scheduling.

(ii) A seminal paper on scheduling in the presence of convex costs is [30], where the author introduced the Generalized- $c\mu$ -rule ($Gc\mu$) and showed its optimality in heavy-traffic for convex delay cost. The $Gc\mu$ -rule associates to each class- i customer with experienced delay d_i the index $C'_i(d_i)\mu_i$, where $C_i(\cdot)$ denotes the class- i delay cost. The optimality of the $Gc\mu$ -rule in a heavy-traffic setting with multiple servers was established in [27]. In [1] the authors calculate Whittle's index policy for a multi-class queue with general holding cost functions. In [12], convex holding costs are considered as well and, taking a stochastic approach, the author obtains an index rule that consists on first-order differences of the cost function, rather than on its derivatives.

(iii) The impact of abandonments has attracted considerable interest from the research community, with a surge in recent years. To illustrate the latter, we can mention the recent Special Issue in Queueing Systems on queueing systems with abandonments [23] and the survey paper [15] on abandonments in a many-server setting. Related literature that is more close to our present work consists of papers that deal with optimal scheduling or control aspects of multi-class queueing systems in the presence of abandonments, see for instance [20, 2, 4, 16, 5, 7, 3, 24, 26, 10]. Note that, with the exception of [3], these papers consider linear holding cost. In the case of one server the authors of [16, 10] show that (for exponential distributed service requirements and impatience times) under an additional condition on the ordering of the abandonment rates an index policy is optimal for linear holding cost. In the case of no arrivals and non-preemptive service, the authors of [2] provide partial characterizations of the optimal policy and show that an optimal policy is typically state dependent. It is worthwhile to mention that [2] is inspired by a patient triage problem which illustrates that abandonments are as well an important issue in other areas than information technology. As far as the authors are aware, the above two settings are the only

ones for which structural optimality results have been obtained. State-dependent heuristics for the multi-class queue are proposed in [2] for two classes and no arrivals and in [20] for an arbitrary number of classes including new arrivals. In [7] the authors obtain Whittle's index for a multi-class abandonment queue without arrivals, that is, each customer is a bandit and the state of a bandit is either *present* or *departed*. In an overload setting the abandonment queue has been studied under a fluid scaling in [4, 5], where the authors scale the number of servers and the arrival rate and show that an index rule is asymptotically fluid optimal. In our analysis we will show how the indices of [7] and [4, 5] coincide with the Whittle's index rule in the case of linear holding costs and in the presence of arrivals. In [26] the optimal policy is obtained for two classes of customers for a fluid approximation of the stochastic model, which allows to propose a heuristic for the stochastic model for an arbitrary number of classes. We finally mention [3, 24] where the authors derive index policies by studying the Brownian control problem arising in heavy traffic. In [3] general delay costs are considered while in [24] the impatience of customers has a general distribution with increasing failure rate.

(iv) The approach of using the fluid control model to find an approximation for the stochastic optimization problem finds its roots in the pioneering works by Avram et al. [6] and Weiss [34]. It is remarkable that in some cases the optimal control for the fluid model coincides with the optimal solution for the stochastic problem. See for example [6] where this is shown for the $c\mu$ -rule in a multi-class single-server queue and [9] where this is shown for Klimov's rule in a multi-class queue with feedback. For other cases, researchers have aimed at establishing that the fluid control is asymptotically optimal, that is, the fluid-based control is optimal for the stochastic optimization problem after a suitable scaling, see for example [8, 17, 32]. We conclude by mentioning that the fluid approach owes its popularity to the groundbreaking result stating that if the fluid model drains in finite time, the stochastic process is stable, see [14].

3. MODEL DESCRIPTION

We consider a multi-class single-server queue with K classes of customers. Class- k customers arrive according to a Poisson process with rate λ_k and have an exponentially distributed service requirement with mean $1/\mu_k$, $k = 1, \dots, K$. We denote by $\rho_k := \lambda_k/\mu_k$ the traffic load of class k , and by $\rho := \sum_{k=1}^K \rho_k$ the total load to the system. We model abandonments of customers in the following way:

- Any class- k customer not served abandons after an exponentially distributed amount of time with mean $1/\theta_k$, $k = 1, \dots, K$, with $\theta_k > 0$.
- A class- k customer that is being served abandons after an exponentially distributed amount of time with mean $1/\theta'_k$, $k = 1, \dots, K$, with $\theta'_k \geq 0$.

The server has capacity 1 and can serve at most one customer at a time, where the service can be preemptive. We make the following natural assumption: $\mu_k + \theta'_k \geq \theta_k$, for all k . That is, for a class- k customer its departure rate is higher when being served than when not being served.

At each moment in time, a policy φ decides which class is served. Because of the Markov property, we can focus on policies that only base their decisions on the current number

of customers present in the various classes. For a given policy φ , $N_k^\varphi(t)$ denotes the number of class- k customers in the system at time t , (hence, including the one in service), and $\vec{N}^\varphi(t) = (N_1^\varphi(t), \dots, N_K^\varphi(t))$. Let $S_k^\varphi(\vec{N}^\varphi(t)) \in \{0, 1\}$ represent the service capacity devoted to class- k customers at time t under policy φ . The constraint on the service amount devoted to each class is $S_k^\varphi(\vec{n}) = 0$ if $n_k = 0$ and

$$\sum_{k=1}^K S_k^\varphi(\vec{n}) \leq 1. \quad (1)$$

The above describes a birth-and-death process with transition rates: $q_k^\varphi(\vec{n}, \vec{n} + \vec{e}_k) = \lambda_k$, and

$$q_k^\varphi(\vec{n}, \vec{n} - \vec{e}_k) = \mu_k S_k^\varphi(\vec{n}) + \theta_k(n_k - S_k^\varphi(\vec{n})) + \theta'_k S_k^\varphi(\vec{n}), \quad (2)$$

for $n_k > 0$, with \vec{e}_k a K -dimensional vector with all zeros except for the k -th component which is equal to 1.

Let $C_k(n, a)$ denote the cost per unit of time when there are n class- k customers in the system and when either class k is not served (if $a = 0$), or when class k is served (if $a = 1$). We assume $C_k(\cdot, 0)$ and $C_k(\cdot, 1)$ are convex and non-decreasing functions and satisfy for all $n \geq 0$,

$$C_k(n, 0) - C_k((n-1)^+, 0) \leq C_k(n+1, 1) - C_k(n, 1) \leq C_k(n+1, 0) - C_k(n, 0). \quad (3)$$

Observe that if $C_k(0, 0) \geq C_k(0, 1)$, then (3) implies that, for all n , $C_k(n, 0) \geq C_k(n, 1)$. We also note that (3) is always satisfied when (i) $C_k(n, a) = C_k(n)$, or when (ii) $C_k(n, a) = C_k((n-a)^+)$. Case (i) represents holding costs for customers in the *system*, while (ii) represents holding costs for customers in the *queue*.

We further introduce a cost d_k for every class- k customer that abandoned the system when not being served and a cost d'_k for a class- k customer that abandoned the system while being served.

The objective of the paper is to find the optimal scheduling policy φ under the average-cost criteria, that is,

$$\limsup_{T \rightarrow \infty} \sum_{k=1}^K \frac{1}{T} \mathbb{E} \left(\int_0^T C_k(N_k^\varphi(t), S_k^\varphi(\vec{N}^\varphi(t))) dt + d_k R_k^\varphi(T) + d'_k R'_k{}^\varphi(T) \right), \quad (4)$$

where $R_k^\varphi(T)$ and $R'_k{}^\varphi(T)$ denote the number of class- k customers that abandoned the queue while waiting and while being served respectively, in the interval $[0, T]$ under policy φ . We have $\mathbb{E}(R_k^\varphi(T)) = \theta_k \mathbb{E} \left(\int_0^T (N_k^\varphi(t) - S_k^\varphi(\vec{N}^\varphi(t))) dt \right)$ and $\mathbb{E}(R'_k{}^\varphi(T)) = \theta'_k \mathbb{E} \left(\int_0^T S_k^\varphi(\vec{N}^\varphi(t)) dt \right)$, by Dynkin's formula. We now introduce the following notation to denote the total cost in state n_k under action $a \in \{0, 1\}$:

$$\tilde{C}_k(n_k, a) := C_k(n_k, a) + d_k \theta_k (n_k - a)^+ + d'_k \theta'_k \min(a, n_k), \quad (5)$$

so that the objective (4) can be equivalently written as

$$\limsup_{T \rightarrow \infty} \sum_{k=1}^K \frac{1}{T} \mathbb{E} \left(\int_0^T \tilde{C}_k(N_k^\varphi(t), S_k^\varphi(\vec{N}^\varphi(t))) dt \right). \quad (6)$$

The above described stochastic control problems have proved to be very difficult to solve. Already for the special case of linear holding cost, deriving structural properties of optimal

policies is extremely challenging. For example, in [16] optimal dynamic scheduling is studied for two classes of customers ($K = 2$), $d_k = d'_k$, $\theta_k = \theta'_k$, $\mu_1 = \mu_2 = 1$, and linear holding cost. Define $\tilde{c}_k := c_k + d_k \mu_k$. For the special case where $\tilde{c}_1 \geq \tilde{c}_2$ and $\theta_1 \leq \theta_2$, the authors show that it is optimal to give strict priority to class 1, see [16, Theorem 3.5]. It is intuitively clear that giving priority to class 1 is the optimal thing to do, since serving class 1 myopically minimizes the (holding and abandonment) cost and in addition it is advantageous to keep the maximum number of class-2 customers in the system (without idling), since they have the highest abandonment rate. In [10] optimal dynamic scheduling is studied for $C_k(n, a) = c_k n$, $d_k = d'_k$, and either $\theta_k = \theta'_k$ or $\theta'_k = 0$. For the special case where the classes can be ordered such that $\tilde{c}_1 \geq \dots \geq \tilde{c}_K$, $\tilde{c}_1(\mu_1 + \theta'_1 - \theta_1) \geq \dots \geq \tilde{c}_K(\mu_K + \theta_K - \theta'_K)$, and $\tilde{c}_1(\mu_1 + \theta'_1 - \theta_1)/\theta_1 \geq \dots \geq \tilde{c}_K(\mu_K + \theta'_K - \theta_K)/\theta_K$, the authors show that it is optimal to give strict priority to the class having the highest index.

Outside these special parameter settings, or for convex holding cost, an optimal policy is expected to be state dependent, and as far as the authors are aware, no (structural) results exist for this stochastic optimal control problem.

In order to obtain insights into optimal control for *convex* holding cost, in this paper we will solve a relaxed version of the optimization problem that allows us to propose a heuristic for the original model. This relaxation technique is described in the next section.

4. RELAXATION AND INDEXABILITY

The solution to (6) under constraint (1) cannot be solved in general. Following Whittle [35], we study the relaxed problem in which the constraint on the service devoted to each class must be satisfied on *average*, and not in every decision epoch. The control policy must thus satisfy

$$\limsup_{T \rightarrow \infty} \frac{1}{T} \int_0^T \sum_{k=1}^K S_k^\varphi(\vec{N}^\varphi(t)) dt \leq 1. \quad (7)$$

The objective of the relaxed problem is hence to determine the policy that solves (6) under constraint (7). This can be solved by considering the following unconstrained control problem: find a policy φ that minimizes

$$\limsup_{T \rightarrow \infty} \frac{1}{T} \mathbb{E} \left(\int_0^T \left(\sum_{k=1}^K \tilde{C}_k(N_k^\varphi(t), S_k^\varphi(\vec{N}^\varphi(t))) - W \left(1 - \sum_{k=1}^K S_k^\varphi(\vec{N}^\varphi(t)) \right) \right) dt \right), \quad (8)$$

where W is the Lagrange multiplier. The key observation made by Whittle is that problem (8) can be decomposed into K subproblems, each corresponding to a different class (or bandit when using terminology from the RMABP literature). Thus, the solution to (8) is obtained by combining the solution to the separate optimization problems.

For the remainder of this section we drop the dependency on the class, and for a given W we consider the individual optimization problem for a given class, that is, minimize

$$\limsup_{T \rightarrow \infty} \frac{1}{T} \mathbb{E} \left(\int_0^T \left(\tilde{C}(N^\varphi(t), S^\varphi(N^\varphi(t))) - W \left(1 - S^\varphi(N^\varphi(t)) \right) \right) dt \right), \quad (9)$$

where now $N^\varphi(t)$ is the state of a given class at time t . Under a stationarity assumption, we can invoke ergodicity to show that (9) is equivalent to minimizing

$$\mathbb{E}(\tilde{C}(N^\varphi, S^\varphi(N^\varphi)) - W\mathbb{E}(\mathbb{1}_{S^\varphi(N^\varphi)=0})), \quad (10)$$

where N^φ denotes the steady-state number of customers in a class under policy φ . We observe that the multiplier W can be interpreted as *subsidy for passivity*.

In summary, the relaxed optimization problem can be written as K independent one-dimensional Markov Decision Problems (9). In the next section we will determine the structure of the optimal control of the relaxed problem (9).

4.1 Threshold policies

In the following proposition we show that an optimal solution of the relaxed problem (9) is of threshold type, i.e., when the number of customers is above a certain threshold n , the class is served, and is not served otherwise. We denote by $\varphi = n$ the threshold policy with threshold n , that is, $S^n(m) = 1$ if $m > n$, and $S^n(m) = 0$ otherwise.

PROPOSITION 1. *There is an n such that the policy $\varphi = n$ is an optimal solution of the relaxed problem (9).*

PROOF. The value function $V(n)$ satisfies

$$\begin{aligned} &(\mu + \theta' + m\theta + \lambda)V(m) + g \\ &= \lambda V(m+1) + \theta(m-1)V((m-1)^+) \\ &+ \min\{\tilde{C}(m, 0) - W + (\mu + \theta')V(m) + \theta V((m-1)^+), \\ &\tilde{C}(m, 1) + (\mu + \theta')V((m-1)^+) + \theta V(m)\}, \end{aligned} \quad (11)$$

where g is the average cost incurred under an optimal policy. Proving optimality of a threshold policy is hence equivalent to showing that if it is optimal in (11) for state $m+1$, $m \geq 0$ to be passive, then it is also optimal in (11) for state m to be passive, i.e., $\tilde{C}(m+1, 0) - W + (\mu + \theta' - \theta)V(m+1) \leq \tilde{C}(m+1, 1) + (\mu + \theta' - \theta)V(m)$, implies $\tilde{C}(m, 0) - W + (\mu + \theta' - \theta)V(m) \leq \tilde{C}(m, 1) + (\mu + \theta' - \theta)V((m-1)^+)$. A sufficient condition for the above to be true is (3) together with the inequality $V(m+1) + V((m-1)^+) \geq 2V(m)$, for $m \geq 0$. The latter condition, convexity of the value function, will be proved below, which concludes the proof.

In case of bounded transition rates, one can uniformize the system and use Value-Iteration in order to prove convexity. However, our transition rates are unbounded. We therefore look to the truncated space, truncated by $L > 1$, and smooth the arrival transition rates as follows: $q_k(m, m+1) = \lambda(1 - m/L)$, $m = 0, \dots, L$. Denote by $V^L(m)$ the value function of the L -truncated system. By [11, Theorem 3.1] we have that $V^L(m) \rightarrow V(m)$ as $L \rightarrow \infty$. Hence, convexity of the function V is implied by convexity of V^L for all L , and we are left with proving the latter. The latter is uniformizable, hence we can use the Value-Iteration technique. The proof for convexity of V^L is available in [25]. \square

Below we write the steady-state distribution of threshold policy $\varphi = n$. We denote the steady-state probability of being in state i under policy $\varphi = n$ by $\pi^n(i)$, and have

$$\pi^n(i) = \prod_{m=1}^i \frac{q^n(m-1, m)}{q^n(m, m-1)} \pi^n(0), \quad i = 1, 2, \dots, \quad (12)$$

where $\pi^n(0) = \left(1 + \sum_{i=1}^{\infty} \prod_{m=1}^i \frac{q^n(m-1, m)}{q^n(m, m-1)}\right)^{-1}$.

4.2 Indexability and Whittle's index

Indexability is the property that allows to develop a heuristics for the original problem. This property requires to establish that as the Lagrange multiplier, or equivalently the subsidy for passivity, W , increases, the collection of states in which the optimal action is *passive* increases, i.e., the optimal threshold n increases. It was first introduced by Whittle [35] and we formalize it in the following definition.

DEFINITION 1. *A class is indexable if the set of states in which passive is an optimal action (denoted by $D(w)$) increases in W , that is, $W' < W \Rightarrow D(W') \subseteq D(W)$.*

An optimal solution of problem (9) is a threshold policy, or more specifically, if it is optimal to be passive in state m , $m \geq 1$, then it is also optimal to be passive in state $m-1$, see the proof of Proposition 1. We can therefore equivalently write the following definition for indexability.

DEFINITION 2. *Let $n(W)$ denote the largest value such that the threshold policy $n(W)$ minimizes (9). A class is indexable if $n(W)$ is non-decreasing in W , that is, $W' < W \Rightarrow n(W') \leq n(W)$.*

Provided we can establish indexability, the Whittle index in a state m is defined as the largest value for the subsidy such that the optimal policy is indifferent of the action in state m . Formally:

DEFINITION 3. *When a class is indexable, the Whittle's index in state m is defined by $W(m) := \inf\{W : m \leq n(W)\}$.*

The solution to the relaxed control problem (8) will then be to activate all classes k that are in a state n_k such that their Whittle's index exceeds the subsidy for passivity, i.e., $W_k(n_k) > W$. A standard Lagrangian argument shows that there exists a value of W for which the constraint (7) is binding, i.e., the optimal policy φ that solves Problem (8) will on average activate 1 class.

Obviously, the solution to the relaxed optimization problem is not feasible for the original problem. Following Whittle, we use Whittle's index to construct the following heuristic for the original problem (6) under the constraint (1): select in every decision epoch the class with *largest* Whittle index. We will formally describe this in Section 5.

To conclude this subsection we show that for the model under consideration, the classes are indexable.

PROPOSITION 2. *All classes are indexable.*

PROOF. Since an optimal policy for (9) is of threshold type, for a given subsidy W the optimal average cost under threshold n will be $g(W) := \min_n \{g^{(n)}(W)\}$, where

$$g^{(n)}(W) := \sum_{m=0}^{\infty} \tilde{C}(m, S^n(m)) \pi^n(m) - W \sum_{m=0}^n \pi^n(m). \quad (13)$$

The function $g(W)$ is a lower envelope of linear non-increasing functions in W (see Figure 1, where we depict the lower-envelope for the case of quadratic cost). It thus follows that $g(W)$ is a concave non-increasing function.

It follows directly that the right-derivative of $g(W)$ in W is given by $-\sum_{m=0}^{n(W)} \pi^{n(W)}(m)$. Moreover, we will prove below that $\sum_{m=0}^n \pi^n(m)$ is strictly increasing in n . Since $g(W)$ is concave in W , its second derivative is non-increasing in W . It hence follows that $n(W)$ is non-decreasing in W .

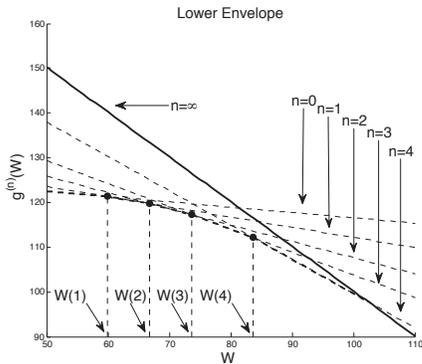


Figure 1: Lower envelop $g = \min_n \{g^{(n)}\}$ **when** $\tilde{C}(n, a) = (1 + 2\theta)n + 3n^2$, **for** $a = 0, 1$, **and** $\theta = 6, \lambda = 33, \mu = 10$.

We now prove that $\sum_{i=0}^n \pi^n(i)$ is strictly increasing, or equivalently, that $1 - \sum_{i=n+1}^{\infty} \pi^n(i)$ is strictly decreasing. Using (12), the latter is equivalent to verifying that

$$\frac{\sum_{m=n+1}^{\infty} \prod_{i=1}^m \frac{q^n(i-1, i)}{q^n(i, i-1)}}{\sum_{m=n}^{\infty} \prod_{i=1}^m \frac{q^{n-1}(i-1, i)}{q^{n-1}(i, i-1)}} < \frac{1 + \sum_{m=1}^{\infty} \prod_{i=1}^m \frac{q^n(i-1, i)}{q^n(i, i-1)}}{1 + \sum_{m=1}^{\infty} \prod_{i=1}^m \frac{q^{n-1}(i-1, i)}{q^{n-1}(i, i-1)}}. \quad (14)$$

holds for all n . Note that $q^n(m-1, m) = q^{n-1}(m-1, m)$ for all m and $q^n(m, m-1) = q^{n-1}(m, m-1)$ for all $m \neq n$. From the assumption $\mu + \theta' \geq \theta$ we have $q^n(n, n-1) \leq q^{n-1}(n, n-1)$. Hence, the left-hand-side of (14) is strictly less than 1, while the right-hand-side is larger than or equal to 1. This proves (14). \square

4.3 Derivation of Whittle's index

We are now in position of deriving Whittle's index. An optimal policy is fully characterized by a threshold n such that the passive action is prescribed for states $m \leq n$, and the active action for states $m > n$. Our key observation to derive Whittle's index is that it is not necessary to solve the optimality equation (11), but that it suffices to determine the average cost for threshold policies. In turn, the average reward g can be expressed as a function of the steady-state probabilities, which in the case of birth-and-death processes have a well-known solution. This explains the novelty of our approach with respect to the literature as it gives a unifying framework and determines the Whittle's index for any indexable birth-and-death process. We will come back to this at the end of this section.

We can now state the main result of the paper.

PROPOSITION 3. *If*

$$\frac{\mathbb{E}(\tilde{C}(N^n, S^n(N^n))) - \mathbb{E}(\tilde{C}(N^{n-1}, S^{n-1}(N^{n-1})))}{\sum_{m=0}^n \pi^n(m) - \sum_{m=0}^{n-1} \pi^{n-1}(m)}, \quad (15)$$

with $\pi^n(m)$ the steady-state probability of being in state m under policy n , is non-decreasing in n , then Whittle's index $W(n)$ is given by (15).

PROOF. Let $\tilde{W}(n)$ be the value for the subsidy such that the average cost under threshold policy n is equal to that under policy $n-1$. Hence, using (10), we have that for all $n \geq 1$, $\mathbb{E}(\tilde{C}(N^n, S^n(N^n)) - \tilde{W}(n)\mathbb{E}(\mathbb{1}_{S^n(N^n)=0}))$ is equal

to $\mathbb{E}(\tilde{C}(N^{n-1}, S^{n-1}(N^{n-1})) - \tilde{W}(n)\mathbb{E}(\mathbb{1}_{S^{n-1}(N^{n-1})=0}))$. For threshold policy n we have $\mathbb{E}(\mathbb{1}_{S^n(N^n)=0}) = \sum_{m=0}^n \pi^n(m)$, hence $\tilde{W}(n)$ is given by (15).

It can be verified that $\tilde{W}(n)$ being non-decreasing, implies that $g(\tilde{W}(n)) = g^{(n)}(\tilde{W}(n)) = g^{(n-1)}(\tilde{W}(n))$. In addition, since $\frac{dg^{(n)}(W)}{dW} = -\sum_{m=0}^n \pi^n(m)$ is decreasing in n (see proof of Proposition 2), we have $g(W) = g^{(n-1)}(W)$ for $\tilde{W}(n-1) \leq W \leq \tilde{W}(n)$. This implies that Whittle's index is given by $W(n) = \tilde{W}(n)$. \square

Equation (15) can be numerically computed, since the cost function and the steady-state probabilities are known. In Section 6 closed-form expressions and limiting properties for Whittle's index will be derived for special cases.

A few comments are in order. The first concerns the form of (15). The numerator in (15) can be interpreted as the increase in cost by deciding to become passive in state n and keeping all other actions unchanged, and similarly the denominator can be understood as the corresponding increase of passivity rate for the process, measured by the additional probability in which a subsidy is received. Thus $W(n)$ can be interpreted as a measure of increased cost per unit of increased passivity, a term coined as Marginal Productivity Index by Niño-Mora [28].

The second comment regards the applicability of Whittle's index (15) in other contexts. Indeed we can outline a general recipe to develop Whittle's indices for bandits whose evolution can be described by general birth-and-death processes:

- (i) Establish optimality of monotone policies (as in Proposition 1).
- (ii) Establish indexability (as in Proposition 2).
- (iii) If (i) and (ii) can be established, then Whittle's index is given by Proposition 3, where the steady-state probabilities are as in (12).

Steps (i) and (ii) are model dependent, but (iii) is immediate and the index will always be given by Proposition 3.

To the best of our knowledge we are the first to observe that for bandits whose evolution can be described by a birth-and-death process, one can get an explicit closed-form expression for Whittle's index. Perhaps a reason for this lies in the difficulty to solve the optimality equation (11), which has two unknowns g and $V(m)$, and this has led researcher to circumvent this difficulty by considering the discounted cost first, equating the total discounted costs as done in Proposition 3 for average cost and then taking the limit in order to retrieve an index for the average cost case. This is for instance the approach taken in [1] to derive an index for convex costs without abandonments or in [19, Section 6.5] for bi-directional bandits in which the *active* and *passive* actions push the process in opposite directions. In [22] the authors develop an algorithm to calculate an index in a multi-class queue with admission control. All these models have in common that after the relaxation, the bandits are birth-and-death, and the obtained Whittle's index is thus equal to (15). We will explain in Section 7 how to derive the index of [1] using this approach. Regarding the bi-directional bandit it can be directly checked that Index (15) is equivalent to the index of [19, Theorem 6.4]. Finally, we note that by adapting the cost structure we obtain that index (15) is equivalent to that of [22, Theorem 2].

Having made this remark on the applicability of (15) in a more wider context, in the remainder of the paper we will discuss the properties of Whittle's index (15) in the context of a queue with convex costs and abandonments.

5. WHITTLE'S INDEX POLICY

In this section we describe how the solution to the relaxed optimization problem is used to obtain a heuristic for the original stochastic model. The optimal control for the relaxed problem is not feasible for the original stochastic model, since in the latter at most one class can be served at a time. Whittle [35] therefore proposed the following heuristic, which is nowadays known as Whittle's index policy:

DEFINITION 4 (WHITTLE'S INDEX POLICY). *Assume at time t we are in state $\vec{N}(t) = \vec{n}$. Whittle's index policy prescribes to serve the class k having currently the highest non-negative Whittle's index $W_k(n_k)$,*

$$W_k(n_k) = \frac{\mathbb{E}(\tilde{C}_k(N_k^{n_k}, S^{n_k}(N_k^{n_k}))) - \mathbb{E}(\tilde{C}_k(N_k^{n_k-1}, S^{n_k-1}(N_k^{n_k-1})))}{\sum_{m=0}^{n_k} \pi_k^{n_k}(m) - \sum_{m=0}^{n_k-1} \pi_k^{n_k-1}(m)}, \quad (16)$$

where $\pi_k^{n_k}(m)$ is the steady-state probability in state m under policy n_k for class k .

Note that in case all classes have a negative index, we define that Whittle's index policy will keep the server idle (until there is a class having a positive value for its index). This follows, since, when the Whittle's index is negative, in the relaxed problem you will keep the class passive even though a negative subsidy is given.

When $\tilde{C}_k(m_k, 0) \geq \tilde{C}_k(m_k, 1)$ for all m_k , the Whittle index $W_k(n_k)$ will always be positive. This can be seen as follows. Recall that $W_k(n_k)$ refers to the value of W such that a threshold policy n_k is an optimal solution of the relaxed problem. Hence, for all $m_k \leq n_k$, it is optimal to keep the class passive, that is, $\tilde{C}_k(m_k, 0) - W_k(n_k) + (\mu_k + \theta'_k - \theta_k)V(m_k) \leq \tilde{C}_k(m_k, 1) + (\mu_k + \theta'_k - \theta_k)V(m_k - 1)$, as we saw in the proof of Proposition 1. Since $\tilde{C}_k(m_k, 0) \geq \tilde{C}_k(m_k, 1)$, $\mu_k + \theta'_k \geq \theta_k$, and $V(\cdot)$ is non-decreasing (see proof of Proposition 1), it follows that $W_k(n_k) \geq 0$.

Instead, when $\tilde{C}_k(m_k, 0) < \tilde{C}_k(m_k, 1)$ for an m_k , $W_k(n_k)$ can be negative for certain states n_k . For example, when $\theta'_k = \theta_k$ and $d'_k \gg d_k$. Then, even though the total departure rate of class- k customers is highest when serving class k , for certain states n_k it might be better not to serve class k . The latter follows since having a class- k customer abandon while being served, will incur a much higher cost than when it abandons while waiting. The latter is taken care of by the possibility of the Whittle index $W_k(n_k)$ to be negative.

From the practical point of view, the interest of index (16) lies in the fact that the index of class k does not depend on the number of customers present in the other classes j , $j \neq k$. Hence, it provides a systematic way to derive simple implementable policies which we will show perform very well, see Section 10, and which in fact are asymptotically optimal in certain settings, see Section 8.

6. CASE STUDIES

In this section we further investigate properties of the obtained Whittle's index (16). In Section 6.1 we obtain that the index is state-independent for linear holding cost. In

Section 6.2 we derive asymptotic properties of the index for general convex holding cost functions.

6.1 Linear holding cost

In this section we consider linear holding cost, that is, $C_k(n_k, a) = c_k(n_k - a)^+ + c'_k \min(n_k, a)$. Hence, under this function, any class- k customer in the queue contributes with c_k to the cost, and a class- k customer in service contributes with c'_k to the cost. In particular, if $c'_k = c_k$, then C_k represents the linear holding cost of customers in the *system* and if $c'_k = 0$ then C_k represents the linear holding cost of customers in the *queue*. These two holding cost functions have been considered in the literature in the context of abandonments, for example [7] considers the former, while [4] takes the latter. From our formula (16) we will be able to obtain a full characterization of Whittle's index. Interestingly, we show that the Whittle's index becomes state-independent and does not depend on the arrival rate λ_k .

It will be convenient to define $\tilde{c}_k := c_k + d_k \theta_k$, $k = 1, \dots, K$, which can be interpreted as the total cost per unit of time incurred by a customer who can abandon the system.

We now state the main result of this section. The proof can be found in the technical report [25].

PROPOSITION 4. *Assume linear holding cost $C_k(n_k, a) = c_k(n_k - a)^+ + c'_k \min(n_k, a)$. Then, the Whittle's index for class k is*

$$W_k(n_k) = \frac{\tilde{c}_k(\mu_k + \theta'_k)}{\theta_k} - \tilde{c}'_k, \text{ for all } n_k. \quad (17)$$

An interesting feature of (17) is that it is independent of the arrival rate λ_k and independent on the number of class- k customers present, n_k . In Section 6.2 we will show that this observation only holds for linear holding costs.

The index (17) allows for the following interpretation. Consider there is only one class- k customer in the system and no future arrivals, we then have $\tilde{C}_k(1, 1) = \tilde{c}'_k$, $\tilde{C}_k(1, 0) = \tilde{c}_k$, $q_k^1(1, 0) = \theta_k$, $q_k^0(1, 0) = \mu_k + \theta'_k$. Index (17) can equivalently be written as $(\mu_k + \theta'_k) \left(\frac{\tilde{c}_k}{\theta_k} - \frac{\tilde{c}'_k}{\mu_k + \theta'_k} \right)$, which is equal to $q_k^0(1, 0) \left(\frac{\tilde{C}_k(1, 0)}{q_k^1(1, 0)} - \frac{\tilde{C}_k(1, 1)}{q_k^0(1, 0)} \right)$. Hence, the index can be interpreted as the reduction in cost when making a class- k bandit active instead of keeping him passive (the term within the brackets) during a time lag equal to the departure time in the active phase.

We now consider some particular cases that have been studied in the literature. For example, let us consider first the case in which all customers can abandon the system, i.e., $\theta'_k = \theta_k$, for $k = 1, \dots, K$, and that the cost for abandonment is the same for both *active* and *passive*, so $d_k = d'_k$. Let us consider two cases: In the first case all customers in the system incur a holding cost. This implies that $c_k = c'_k$, and thus $\tilde{c}_k = \tilde{c}'_k$. Substituting into (17) we get $W_k(n_k) = \frac{\tilde{c}_k \mu_k}{\theta_k}$. In the second case we consider that only customers in the queue incur a holding cost, so we take $c'_k = 0$, which gives $\tilde{c}_k - \tilde{c}'_k = c_k$, and upon substitution in (17) we get the index $W_k(n_k) = \frac{\tilde{c}_k \mu_k}{\theta_k} + c_k$.

We now assume that only customers in the queue can abandon, that is, the customer in service will not abandon, hence $\theta'_k = 0$, for $k = 1, \dots, K$. This is the model assumption of [7] and [4]. We first assume that all customers in the system incur a holding cost, that is, $c_k = c'_k$, and we thus get $\tilde{c}'_k = \tilde{c}_k$. From (17) we get $W_k(n) = \frac{\tilde{c}_k \mu_k}{\theta_k} - c_k$. We can similarly calculate the index in the case in which only customers

in the queue incur a holding cost, i.e., $c'_k = 0$, to obtain the index $W_k(n_k) = \frac{\tilde{c}_k \mu_k}{\theta_k}$. These two last indices have been derived in [7] and [4], respectively. More specifically, [7] derives the index $\frac{\tilde{c}_k \mu_k}{\theta_k} - c_k$ when studying one customer and no future arrivals. Interestingly, we observe that the index remains the same in the presence of random arrivals as considered in this paper. When the customer in service does not contribute to the holding cost our model coincides with that analyzed in [4], where it is shown that the index rule $\frac{\tilde{c}_k \mu_k}{\theta_k}$ is asymptotically fluid optimal in a multi-server queue in overload ($\rho > 1$). We therefore conclude that the Whittle's index, $W_k(n_k)$, we have derived, retrieves index policies that had been proposed in the literature when studying the system in special parameter regimes.

To finish this subsection we now provide an intuition to understand the result of Proposition 4 in the case $\theta'_k = \theta_k$ and $c_k = c'_k$. In this setting, at any moment in time, all customers in the system incur a holding cost c_k and can abandon at rate θ_k . Substituting $\mathbb{E}(\tilde{C}_k(N_k^n, S_k^n(N_k^n))) = \tilde{c}_k \mathbb{E}(N_k^n)$ and $W_k(n) = \frac{\tilde{c}_k \mu_k}{\theta_k}$ in (16), we get the relation

$$\theta_k (\mathbb{E}(N_k^{n-1}) - \mathbb{E}(N_k^n)) = \mu_k \left(\sum_{m=n}^{\infty} \pi_k^{n-1}(m) - \sum_{m=n+1}^{\infty} \pi_k^n(m) \right),$$

which can be seen as a rate conservation. Indeed, the term on the left-hand-side represents the difference in the average number of customers that abandons the system per time unit when comparing both policies n and $n-1$. The right-hand-side represents the difference in the average number of customers that is served per time unit when comparing both policies n and $n-1$. The left-hand-side being equal to the right-hand-side is exactly the rate conservation.

6.2 Convex increasing holding cost

In this section we characterize Whittle's index for general convex non-decreasing holding cost functions. We first note that the cost associated to abandonments of customers are linear functions. We can thus use the result of Proposition 4 to rewrite Whittle's index as

$$W_k(n_k) = d_k(\mu_k + \theta'_k) - d'_k \theta'_k + \frac{\mathbb{E}(C_k(N_k^{n_k}, S_k^{n_k}(N_k^{n_k}))) - \mathbb{E}(C_k(N_k^{n_k-1}, S_k^{n_k-1}(N_k^{n_k-1})))}{\sum_{m=0}^{n_k} \pi_k^{n_k}(m) - \sum_{m=0}^{n_k-1} \pi_k^{n_k-1}(m)}. \quad (18)$$

In the remainder of this section, we will focus on the second term in (18), that is, the term corresponding to the holding cost. We will characterize Whittle's index for large state values and we will observe that for non-linear holding cost the index $W_k(n_k)$ is dependent on n_k , that is, is state-dependent.

We assume that the holding costs $C_k(n_k, 1)$ and $C_k(n_k, 0)$ are upper bounded by polynomials of finite degrees $P_k < \infty$ and $Q_k < \infty$, respectively. Hence, we can write $C_k(n_k, a) = E_k(n_k, a) + o(n_k)$, for large values of n_k , where $E_k(n_k, 1) = \sum_{i=0}^{P_k} C_k^{(P_k, i)} n_k^i$, with

$$C_k^{(P_k, i)} := \lim_{n_k \rightarrow \infty} \frac{C_k(n_k, 1) - \sum_{j=i+1}^{P_k} C_k^{(P_k, j)} n_k^j}{n_k^i},$$

and $E_k(n_k, 0) = \sum_{i=0}^{Q_k} E_k^{(Q_k, i)} n_k^i$, with

$$E_k^{(Q_k, i)} := \lim_{n_k \rightarrow \infty} \frac{C_k(n_k, 0) - \sum_{j=i+1}^{Q_k} E_k^{(Q_k, j)} n_k^j}{n_k^i}.$$

We assume w.l.o.g. that P_k is such that $C_k^{(P_k, P_k)} > 0$ and Q_k such that $E_k^{(Q_k, Q_k)} > 0$.

In the following proposition we give the expression for Whittle's index for large states. The proof can be found in the technical report [25]:

PROPOSITION 5. *Assume that $C_k(n_k, 1)$ and $C_k(n_k, 0)$ are upper bounded by a polynomial of degree P_k and Q_k respectively. Then, we have $W_k(n_k) = W_k^\infty(n_k) + o(1)$, as $n \rightarrow \infty$, where $W_k^\infty(n_k) := d_k(\mu_k + \theta'_k) - d'_k \theta'_k + W_k^c(n_k)$ and*

$$W_k^c(n_k) := (E_k(n_k, 0) - E_k(n_k, 1)) + (\mu_k + \theta'_k - \theta_k)/\theta_k \cdot \left(\frac{E_k(n_k, 0)}{n_k} + \sum_{i=2}^{P_k} C_k^{(P_k, i)} \sum_{j=0}^{i-2} n_k^{i-2-j} \left(\frac{\lambda_k}{\theta_k} \right)^{j+1} \right). \quad (19)$$

Assume for example that $P_k = Q_k \geq 2$ and $C_k^{(P_k, P_k)} = E_k^{(Q_k, Q_k)}$, then for states that are large enough, in the case when $C_k(n_k, a) = C_k(n_k)$ or $C_k(n_k, a) = C_k((n_k - a)^+)$, the value of $W_k^\infty(n_k)$ is determined by the highest polynomial. For this example, this is given by

$$\left(E_k^{(P_k, P_k-1)} - C_k^{(P_k, P_k-1)} + \frac{\mu_k + \theta'_k - \theta_k}{\theta_k} E_k^{(P_k, P_k)} \right) n_k^{P_k-1}. \quad (20)$$

The latter is independent of the arrival rate λ_k , and hence, so is W_k^∞ for large enough states. This robust index (20) can serve as an approximation for Whittle's index policy when there are a large number of customers in the system. In particular, it can be of interest in overload settings.

7. M/M/1 MULTI-CLASS QUEUE

The multi-class M/M/1 queue without abandonments has received lot of attention from the research community. In the case of linear holding cost, the $c\mu$ -index rule has been proved to be optimal in two main settings: (i) with exponential distributed service times and preemptive scheduling [13], and (ii) general service time distributions and non-preemptive scheduling [18]. A brief explanation of the optimality of an index rule is that having a linear holding cost c_k for a class- k customer per unit of time is equivalent to a problem where a reward c_k is received upon service completion (and no holding cost) [19, Section 4.9]. The latter can be seen as a MABP, for which an index rule (in this case $c\mu$) is optimal¹. However, this equivalence holds only for linear holding costs, which explains why for general cost functions the structure of the optimal scheduling policy is no longer of index type. In that context, a fruitful approach has been to derive scheduling policies with near-optimal performance or asymptotically optimal performance in a limiting regime, see the references as stated in Section 2.

In this section, we derive an index policy for the multi-class M/M/1 system by considering the limit of our Whittle index as the abandonment rate tends to 0. Note that the Whittle's index $W_k(n_k)$ goes to ∞ as $\theta_k \rightarrow 0$, and it turns out that when scaling the index by θ_k we get a non-trivial limit. The proof of the next proposition may be found in the technical report [25].

¹This is known as the *tax* formulation of a MABP, see [19, Section 4.9].

PROPOSITION 6. Assume $C_k(n_k, a) = C_k(n_k)$, $a = 0, 1$, $\theta'_k = \theta_k$, and $d_k = d'_k = 0$. Then,

$$\lim_{\theta_k \rightarrow 0} \theta_k W_k(n_k) = \frac{\mu_k(1 - \rho_k)}{\rho_k} \cdot \left[\sum_{m=0}^{\infty} \rho_k^m (1 - \rho_k) C_k(n_k - 1 + m) - C_k(n_k - 1) \right]. \quad (21)$$

A heuristic for the M/M/1 queue with as objective to minimize the holding cost can now be derived as follows. Set $\theta_k = \theta'_k$ for all k and consider the index multiplied by θ_k as $\theta_k \rightarrow 0$. A heuristic is then to give priority according to the index as given in (21).

In case of linear holding costs $C_k(n_k) = c_k n_k$, the index (21) coincides with the $c_k \mu_k$ -rule. For general holding cost the index in (21) was also obtained in Glazebrook *et.al* [1] (see also Section [19, Section 6.5]) by carrying out a model-dependent analysis, which consists in considering first the total discounted holding cost criterion, calculating the corresponding Whittle's index, and afterwards taking the limit in the discounting factor. In that case too, indexability needs to be established. Proposition 6 illustrates the versatility of our approach where after having established indexability and obtained the index for the general multi-class model with abandonments, one can apply it to more specific settings such as in the case of no abandonments ($\theta_k = 0$).

As pointed out in [19, Section 6.5] applying directly the average cost criteria to the M/M/1 queue without abandonments gives no meaningful index. Consider an M/M/1 queue with threshold policy n , where the taken action is *passive* for all states below and equal to n , and *active* for all states above n . A classical queueing theory result shows that in the absence of abandonments, under policy n the probability of being in state n will be $1 - \rho$, i.e., independently of where the threshold is set. This result that the subsidy obtained is $W(1 - \rho)$, that is, independent of the policy n , and therefore, the subsidy does not allow to "calibrate" the states. In our approach this is circumvented by obtaining an index for the, well-defined, case with abandonments and then letting $\theta \rightarrow 0$, while in [19, Section 6.5] this is circumvented by looking at the discounted problem and scaling the immediate cost.

For large values of n_k , the index (21) is approximately equal to $C'_k(n_k) \mu_k$, which we refer to as the $Gc\mu$ -rule. This rule was introduced in [30] for convex delay cost. The equivalence with the $Gc\mu$ rule can be seen as follows. We have for n_k large, $\sum_{m=0}^{\infty} \rho_k^m (1 - \rho_k) C_k(n_k - 1 + m) - C_k(n_k - 1) \sum_{m=0}^{\infty} \rho_k^m (1 - \rho_k) = (1 - \rho_k) \sum_{m=0}^{\infty} \rho_k^m (C(n_k - 1 + m) - C(n_k - 1)) \approx (1 - \rho_k) \sum_{m=0}^{\infty} m \rho_k^m C'(n_k - 1) = C'(n_k) \frac{\rho_k}{(1 - \rho_k)}$, where we used that for n_k large with respect to m , we have $\frac{C(n_k - 1 + m) - C(n_k - 1)}{m} \approx C'(n_k)$ and that large values of m have a negligible weight on the summation. Hence, we can write $\lim_{\theta_k \rightarrow 0} \theta_k W_k(n_k) \approx C'_k(n_k) \mu_k$.

Numerical example. In Table 1 we compare the suboptimality of the $C'(n)\mu$ -rule and index-rule (21) in an M/M/1 queue without abandonments. Note that when $\theta_k = 0$, for all k , we need to assume $\sum_{k=1}^K \rho_k < 1$ in order to assure stability of the system. Consider 4 classes of customers with the following parameters: $\mu_1 = 16, \mu_2 = 27, \mu_3 = 12$ and $\mu_4 = 21$, $\rho_1 = 3\rho/9, \rho_2 = \rho/9, \rho_3 = 5\rho/9$ and $\rho_4 = \rho/9$. The holding cost of each class are cubic, $C_k(n_k) := \alpha_k + \beta_k n_k + \gamma_k n_k^2 + \delta_k n_k^3$, for which (21) simplifies to: $\beta_k \mu_k + \gamma_k \mu_k \left(\frac{3\rho_k - 1}{1 - \rho_k} + \right.$

ρ	0.11	0.21	0.31	0.41
(21)	4.25e-06	1.51e-05	6.07e-06	5.02e-07
$C'(n)\mu$	0.0072	0.0636	0.1002	0.1320
ρ	0.51	0.61	0.71	0.81
(21)	0.008	0.0291	0.0919	1.7129
$C'(n)\mu$	0.1689	0.3616	1.8280	4.9539

Table 1: Suboptimality gap

$2n_k) + \delta_k \mu_k \left(3n_k^2 + 3 \left(\frac{2\rho_k - 1}{1 - \rho_k} \right) n_k + \frac{4\rho_k^2 + \rho_k + 1}{(1 - \rho_k)^2} \right)$. We take the particular example: $C_1(n_1) = 6n_1 + 2n_1^2 + 2n_1^3$, $C_2(n_2) = 2n_2 + 2n_2^2 + 2n_2^3$, $C_3(n_3) = n_3 + n_3^2 + 3n_3^3$ and $C_4(n_4) = 8n_4 + 2n_4^3$. We observe that for this example the $C'(n)\mu$ -rule is outperformed by the index-rule (21), but both policies give nearly optimal performance.

8. ASYMPTOTIC OPTIMALITY

In this section we will discuss asymptotic optimality of Whittle's index policy.

For linear holding cost, asymptotic optimality can be derived directly from [31]. Assume there are M servers and the arrival rate of class- k customers is $M\lambda_k$. Let W_k be the state-independent index as given in (17). In [31, Proposition 6.2] it is shown that the index policy (W), where at each moment in time a server serves a customers having highest non-negative index W_k , is asymptotically optimal in the following sense: for any policy φ ,

$$\lim_{M \rightarrow \infty} \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{k=1}^K \mathbb{E} \left(\int_0^T C(N_k(t), S_k^{(W)}(t)) dt \right) \leq \liminf_{M \rightarrow \infty} \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{k=1}^K \mathbb{E} \left(\int_0^T C(N_k(t), S_k^{(\varphi)}(t)) dt \right).$$

For general holding cost, we can not derive asymptotic optimality. We do expect however that under certain conditions one would have the following. Assume there are M servers and $x_k M$ queues where class- k customers arrive with rate λ_k , $k = 1, \dots, K^2$. A queue can be served by at most one server. In bandit terminology this represents having $x_k M$ class- k bandits whose state (that is, the number of customers in the queue) has values in $\mathcal{S} := \{0, 1, \dots\}$, and the scheduler needs to decide which M bandits to make active (so which M queues to serve). In case the state space \mathcal{S} would have been finite, the result in [33, 31] implies (under certain conditions) asymptotic optimality of Whittle's index policy as $M \rightarrow \infty$. However, for infinite state space, as is the case for our model, no result is known so far.

9. FLUID INDEX

In Section 4 we derived the optimal policy of the relaxed optimization problem (9), which was described by the index value (3). Unfortunately, for non-linear holding cost the index could not be written in closed-form. In this section we will therefore solve the fluid version of the relaxed optimization problem (9), that is, we only take into account the average behavior of the system. This will allow to obtain a

²This can represent for example a setting where there are $x_k M$ class- k flows having newly arriving packets (represented by customers).

closed-form expression for the fluid index. In Section 9.1 we describe the fluid control problem we need to solve and in Section 9.2 we obtain the solution and the fluid index. In addition, we compare the fluid index with the index for the stochastic model.

9.1 Fluid model description

We approximate the stochastic model as presented in Section 3 by a deterministic fluid model, where only the mean dynamics are taken into account. Let $m_k(t)$ be the amount of class- k fluid and $s_k(t)$ the control parameter. The dynamics is then given by

$$\begin{aligned} \frac{dm_k(t)}{dt} = & \lambda_k - \mu_k s_k(t) - \theta_k(m_k(t) - s_k(t))^+ \\ & - \theta'_k \min(s_k(t), m_k(t)). \end{aligned}$$

and the control $s_k(\cdot)$ needs to satisfy $0 \leq s_k(t) \leq 1$, $m_k(t) \geq 0$. At time t , the cost for the fluid model is defined as $\tilde{C}(m(t), s(t)) = (1-s(t))\tilde{C}(m(t), 0) + s(t)\tilde{C}(m(t), 1) - W(1-s(t))$, since $s(t)$ can be viewed as the fraction of time the passive action is taken at time t . The cost functions are continuous in both m and s . Note that we have used the same notation as in the stochastic model where the cost functions were discrete in m and s (slight abuse of notation). Assume

$$\frac{dC(m, 1)}{dm} \leq \frac{dC(m, 0)}{dm}, \quad (22)$$

which is the continuous equivalence of the RHS of (3).

An equilibrium point (\bar{m}_k, \bar{s}_k) of $\frac{dm_k(t)}{dt}$ is such that $0 = \lambda_k - \mu_k \bar{s}_k - \theta_k(\bar{m}_k - \bar{s}_k)^+ - \theta'_k \min(\bar{m}_k, \bar{s}_k)$, $\bar{s}_k \in [0, 1]$, and $\bar{m}_k \geq 0$. For the remainder of this section we drop the dependence on classes from the notation.

In the stochastic model the aim is to minimize (9), that is, to minimize the time-average cost minus the subsidy obtained. In equilibrium, \bar{s} is the average amount of time the system is passive, hence, the fluid version of (9) will be to find the equilibrium point that minimizes $TC(\bar{s}) := (1-\bar{s})\tilde{C}(\bar{m}, 0) + \bar{s}\tilde{C}(\bar{m}, 1) - W(1-\bar{s})$. We denote by (m^*, s^*) an optimal equilibrium point and define

$$TC^* := (1-s^*)\tilde{C}(m^*, 0) + s^*\tilde{C}(m^*, 1) - W(1-s^*). \quad (23)$$

Since the time-average criteria will be attained by several controls, in the next section we will study controls that are *bias-optimal*. That is, among all controls that reach the optimal equilibrium point, a bias-optimal control is the one that minimizes the cost to get to this point.

9.2 Fluid index for bias optimality

Having characterized the optimal equilibrium point in the previous section, the question is which control minimizes the cost to get to this equilibrium, referred to as *bias-optimality*. Hence, our aim is to find the control that minimizes

$$\int_0^\infty (\tilde{C}(m(t), s(t)) - W(1-s(t)) - TC^*) dt. \quad (24)$$

That is, minimize the total cost over time minus the optimal cost in equilibrium. The optimal solution to the fluid bias optimal problem is stated below. For the sake of clarity, we will focus on the case $\theta = \theta'$ and we refer to the technical report [25] for the general setting. Full details on the proof can be found in [25].

PROPOSITION 7. *Assume $\theta = \theta'$ and $d = d'$. An optimal control for the relaxed fluid problem (24) is $s^*(t) = 1$ if $W < w(m(t))$ and $s^*(t) = 0$ otherwise, with*

$$\begin{aligned} w(m) := & C(m, 0) - C(m, 1) + d\mu \\ & + \begin{cases} w^{(1)}(m) & \text{if } m < ((\lambda - \mu)/\theta)^+, \\ w^{(2)}(m) & \text{if } ((\lambda - \mu)/\theta)^+ \leq m \leq \lambda/\theta, \\ w^{(3)}(m) & \text{if } m > \lambda/\theta, \end{cases} \end{aligned}$$

with

$$\begin{aligned} w^{(1)}(m) &= \frac{\mu}{\theta} \frac{C(\frac{\lambda-\mu}{\theta}, 1) - C(m, 1)}{(\lambda - \mu)/\theta - m}, \\ w^{(2)}(m) &= \frac{(\lambda - \theta m) \frac{d}{dm} C(m, 1) + (\theta m + \mu - \lambda) \frac{d}{dm} C(m, 0)}{\theta}, \\ w^{(3)}(m) &= \frac{\mu}{\theta} \frac{C(m, 0) - C(\lambda/\theta, 0)}{m - \lambda/\theta}. \end{aligned}$$

The fluid index $w(m)$ is non-decreasing and continuous.

PROOF. The holding cost functions $C(m, a)$ are convex and non-decreasing, hence $(C(m, a) - C(m', a))/(m - m')$, $m \geq m'$ is non-decreasing in m and m' . Together with (22) this implies that the functions $w^{(1)}$, $w^{(2)}$, and $w^{(3)}$ are non-decreasing in m . Continuity of $w^{(1)}$, $w^{(2)}$, and $w^{(3)}$ follows from the fact that $\lim_{m \uparrow (\lambda-\mu)/\theta} \frac{C(\frac{\lambda-\mu}{\theta}, 1) - C(m, 1)}{(\lambda-\mu)/\theta - m} = \frac{dC(m, 1)}{dm}$ and $\lim_{m \downarrow \lambda/\theta} \frac{C(m, 0) - C(\lambda/\theta, 0)}{m - \lambda/\theta} = \frac{dC(m, 0)}{dm}$.

An optimal control that minimizes (24) is of threshold type, that is, there is an m' such that when $m \leq m'$ it is optimal to set $s(t) = 0$ and when $m > m'$ it is optimal to set $s(t) = 1$, see [25]. Recall that the optimal equilibrium point is denoted by (m^*, s^*) and minimizes $TC(\bar{s})$.

First assume $s^* \in (0, \min(\lambda/\mu, 1))$. In that case, the optimal equilibrium point satisfies $((\lambda - \mu)/\theta)^+ < m^* < \lambda/\theta$ and, due to the convexity of $TC(\bar{s})$ in \bar{s} , we have $\frac{dTC(s^*)}{ds} = 0$. Using that $0 = \lambda - \mu s^* - \theta m^*$, $\frac{dTC(s^*)}{ds} = 0$ can equivalently be written as $W = C(m^*, 0) - C(m^*, 1) + d\mu + w^{(2)}(m^*)$. Hence, for this choice of W , the optimal equilibrium point is given by m^* . Note that $dm(t)/dt$ is strictly negative when $s(t) = 1$ and $m(t) \in (((\lambda - \mu)/\theta)^+, \infty]$, and strictly positive when $s(t) = 0$ and $m(t) \in [0, \lambda/\theta)$. It hence follows directly that the only possible threshold policy that has m^* as equilibrium point, is the threshold policy m^* . Hence, we conclude that for $W = C(m^*, 0) - C(m^*, 1) + d\mu + w^{(2)}(m^*)$, the bias-optimal policy is m^* .

The analysis for the cases $s^* = 0$ or $s^* = \min(\lambda/\mu, 1)$ is more involved. We define the cost-to-go function $K^n(m, W)$ as the cost (24) under the threshold policy n when starting in state m and assuming a subsidy W for passivity. First assume $s^* = 0$, and let $m > \lambda/\theta = m^*$. After some calculations we obtain that for any such m , $K^n(m, W)$ is minimized in $n \geq \lambda/\theta$ when $W = C(n, 0) - C(n, 1) + d\mu + w^{(3)}(n)$. Hence, we conclude that the threshold policy n , with $n \geq \lambda/\theta$, is bias-optimal for this specific choice of subsidy. Similarly, setting $s^* = \min(\lambda/\mu, 1)$ we obtain that a threshold policy n , with $n \leq ((\lambda - \mu)/\theta)^+$, is bias-optimal when $W = C(n, 0) - C(n, 1) + d\mu + w^{(1)}(n)$.

Hence, we conclude that n is an optimal threshold policy when $W = w(n)$. This, together with the fact that $w(n)$ is non-decreasing in n , proves that if $W < w(m(t))$, then a threshold policy m' , with $m' < m(t)$ is bias-optimal, hence $s^*(t) = 1$. Similarly, if $W \geq w(m(t))$, then $s^*(t) = 0$. \square

We make the following observations:

- For linear holding cost, the fluid indices are state-independent and coincide with those of the stochastic model as stated in Proposition 4.
- Assume $C(m, a) = C(m)$. In that case,

$$w^{(2)}(m) = \frac{\mu}{\theta} \frac{d}{dm} \tilde{C}(m),$$

which corresponds to the $Gc\mu/\theta$ -rule. The terms $w^{(1)}(m)$ and $w^{(3)}(m)$ reduce to

$$w^{(1)}(m) = \frac{\mu}{\theta} \frac{\tilde{C}((\lambda - \mu)/\theta) - \tilde{C}(m)}{(\lambda - \mu)/\theta - m},$$

$$w^{(3)}(m) = \frac{\mu}{\theta} \frac{\tilde{C}(m) - \tilde{C}(\lambda/\theta)}{m - \lambda/\theta}.$$

We refer to [12] where index policies based on first-order difference have also been proposed and are shown to empty the system with the lowest cost possible in a single server multi-class queue without abandonments and no future arrivals.

- For large enough values of m , the fluid index $w(m)$ is determined by the value of $C(m, 0) - C(m, 1) + d\mu + \frac{\mu}{\theta} \frac{C(m, 0)}{m}$ (see $w^{(3)}(m)$). In case of polynomial bounded cost functions, as considered in Proposition 5, this coincides with the highest polynomial as determined in (20) for the stochastic model with $P = Q$ and $C^{(P, P)} = E^{(Q, Q)}$.

Having solved the fluid version of the relaxed problem, we can now propose the following heuristic for the stochastic model based on the fluid index $w(m)$.

DEFINITION 5 (FLUID INDEX POLICY). *Assume at time t we are in state $\vec{N}(t) = \vec{n}$. The fluid index policy prescribes to serve the class k having currently the highest non-negative fluid index $w_k(n_k)$.*

10. NUMERICAL ANALYSIS

The objective of the present section is to show in which regimes the Whittle index policy $W(n)$ (Equation (15)) performs well. We will focus on holding cost functions of the shape $C_k(n_k, a) = C_k(n_k)$, that is, the holding cost is a function of the number of class- k customers in the *system*. Throughout this section we assume $\theta_k = \theta'_k$ and $d'_k = d_k$, so that all class- k customers abandon at the same rate and same cost. Hence, $\tilde{C}_k(n_k, a)$ reduces to $C_k(n_k) + d_k\theta_k n_k$.

In Section 10.1 we compare the structure of Whittle's index policy with the structure of the optimal policy, which is numerically obtained through value iteration [29]. In Section 10.2 we then numerically compare the performance of the index policies with that of the optimal policy.

10.1 Structure of different policies

We compare the structure of the different index policies and the optimal policy for linear and convex holding cost.

10.1.1 Linear holding cost

We performed numerical analysis for a wide range of parameters and observed that for linear holding cost $C_k(n, a) = c_k n$ the optimal policy is of the following structure: when (N_1, \dots, N_K) is close enough to the origin (and N_i denotes

the number of class- i customers in the system), it is optimal to prioritize classes according to the $\tilde{c}\mu$ -rule, otherwise to prioritize classes according to the $\tilde{c}\mu/\theta$ -rule, where $\tilde{c}_k := c_k + d_k\theta_k$, see Figure 2 (left) with $\epsilon = 0$ as described in the next section. Hence, the Whittle's index (which corresponds to the $\tilde{c}\mu/\theta$ -rule in the linear case) captures the optimal action for states that are not too close to the origin.

10.1.2 General holding cost

To discuss the structure of index policies for general holding cost, we focus on two classes of customers ($K = 2$). In a state (N_1, N_2) , the action taken by Whittle's index rule is to serve the class having highest value $W_k(N_k)$. Since $W_k(N_k)$ is a non-decreasing function, this implies that there is an increasing switching curve (SC) such that when (N_1, N_2) is below the SC, Whittle's index policy serves class 1 and for any state (N_1, N_2) above the curve the policy serves class 2. Note that for linear holding cost this switching curve collapses either to the vertical or horizontal axis.

By simulations we observed that an optimal policy is as well of switching curve type. For example, in Figure 2 (left) we plot the switching curve of the optimal policy with the following holding cost: $C_1(n) = n + \epsilon n^2$ and $C_2(n) = n$ (parameters $\lambda = [9, 10]$, $\mu = [14, 16]$, $\theta = [2, 0.05]$, $d = [4, 0.3]$). When $\epsilon = 0$, we obtain a decreasing switching curve, which describes the behavior of the optimal policy for linear cost as explained in Section 10.1.1. As ϵ becomes positive, the switching curve becomes increasing. In addition, ϵ becomes larger, and hence the quadratic cost of class 1 increases, and therefore, class 1 gets priority in a larger region.

We now compare the actions taken under Whittle's index policy and the optimal policy. We consider an example with quadratic costs $C_1(n) = (c_{11} + d_1\theta_1)n + c_{21}n^2$ and $C_2(n) = (c_{12} + d_2\theta_2)n + c_{22}n^2$, and set the following parameters $\mu = [15, 18]$; $\theta = [4, 7]$; $c_1 = [1, 4]$; $c_2 = [2, 1]$; $d = [8, 6.5]$. In Figures 2 (middle and right) we plot the optimal actions (obtained by value iteration) for load $\rho = 0.8$ and $\rho = 2.5$, respectively, and compare it to the actions taken under Whittle's index policy. We observe that the optimal policy can be described by a switching curve. In addition the optimal policy coincides with that of Whittle's index $W(n)$ in almost all the state space as the workload increases. We also plotted the switching curve corresponding to the fluid index $w(n)$ and observe a very good fit.

10.2 Performance evaluation

In this section we evaluate numerically the performance of the index policies. This is carried out by computing the relative sub optimality gap between the average cost of the optimal solution and an index policy. In order to compute this we use the Value Iteration algorithm [29].

We saw in Section 7 that the index policy with index (21) performs very well in an M/M/1 multi-class systems (when there are no abandonments). We considered cubic costs and 4 classes of customers and compared the Generalized index rule ($Gc\mu$) and the index-rule of (21) and we observed there that the latter performs slightly better than the $Gc\mu$ -rule.

In this section we will consider scenarios allowing abandonments. We will evaluate the following indices: (i) the Whittle index $W(n)$ (Equation (15)), (ii) the Whittle index for large states $W^\infty(n)$ and (iii) the fluid index $w(n)$. We compare these to the two index policies proposed for a multi-class queue without abandonments: the $Gc\mu$ -rule, and the

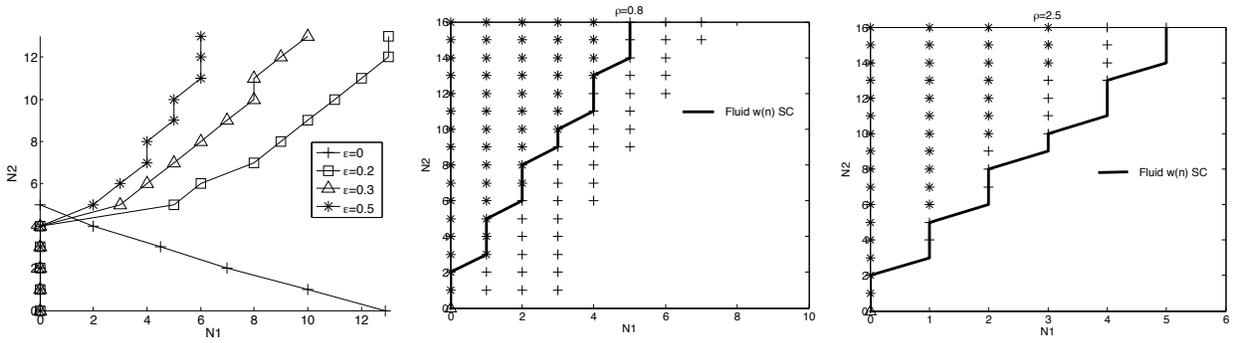


Figure 2: (Left:) Switching curves of the optimal policy for varying holding cost (from linear to quadratic). (Middle and right:) Actions under the optimal policy, the index policy $W(n)$, and the fluid index policy for quadratic holding cost. Area with “+”: $W(n)$ serves class 1 while it is optimal to serve class 2, Area with “*”: $W(n)$ serves class 2, which is also optimal, and in the white area $W(n)$ serve class 1, which is also optimal.

index-rule corresponding to (21) which is an approximation of $W(n)$ for θ close to zero. We will analyze two different scenarios: (1) varying the workload ρ , and (2) varying the abandonment rates θ_k .

10.2.1 Varying Workload

In this section we aim at observing the behavior of index policies for varying workload.

Example with linear holding cost: We set $C_k(n, a) = c_k n$, $\mu = [15, 25]$, $\theta = [4, 2]$, $c = [1, 1]$, $d = [5, 3.2]$, and let $\rho = \sum_{k=1}^2 \lambda_k / \mu_k$ vary in the interval $[0, 3.4]$, with $\lambda_1 / \mu_1 = \lambda_2 / \mu_2$. For linear holding costs, the indices $W(n)$, $W^\infty(n)$ and $w(n)$ reduce to the $\tilde{c}\mu/\theta$ -rule and the indices $Gc\mu$ and (21) reduce to the $\tilde{c}\mu$ -rule, with $\tilde{c}_k = c_k + d_k \theta_k$.

In Figure 3 (left) we observe that the $\tilde{c}\mu$ -rule is optimal in underload, while the $\tilde{c}\mu/\theta$ -rule performs optimal in overload. The latter can be explained by the structure of the optimal policy as discussed in Section 10.1.1. In a state far from the origin, the optimal action is to serve according to $\tilde{c}\mu/\theta$, which is the region in which the process will live in overload, explaining why the $\tilde{c}\mu/\theta$ -rule performs well in this case. In underload, apparently the effect of abandonments is not that important and the $\tilde{c}\mu$ -rule performs very well.

Example with quadratic holding cost: Consider the following parameters: $\mu = [15, 18]$, $\theta = [4, 7]$, $c_1 = [1, 4]$, $c_2 = [2, 1]$, $d = [8, 6.5]$, and we let λ vary, but keeping $\lambda_1 / \mu_1 = \lambda_2 / \mu_2$. We assume quadratic costs $C_1(n) = (c_{11} + d_1 \theta_1)n + c_{21} n^2$ and $C_2(n) = (c_{12} + d_2 \theta_2)n + c_{22} n^2$. See Figure 3 (middle) for the sub-optimality gap.

We observe that for low load the $Gc\mu$ -rule and the index-rule (21) behave very well. However, as the load grows larger, the sub-optimality gap of these θ -independent policies grows large, while our Whittle index policy $W(n)$, the Whittle index policy for large states $W^\infty(n)$ and the fluid index policy $w(n)$ become near optimal as $\rho > 2$ in this example. Hence, our index policies are very suitable for the overload setting, which are in practice of main importance. Note that the jump around $\rho = 2$ for the index-rule (21) is a result of undefined values around $\lambda_k = \mu_k$.

10.2.2 Varying abandonment rates

In this section we evaluate the performance of the index policies for varying abandonment rates.

Linear holding cost: In this case, the 5 index policies mentioned above reduce to the $\tilde{c}\mu/\theta$ -rule and the $\tilde{c}\mu$ -rule, as explained in Section 10.2.1. As $\theta_k \rightarrow 0$, we observed in simulations that the $\tilde{c}\mu$ -rule performs optimal, while the $\tilde{c}\mu/\theta$ -rule might perform very bad when the abandonment rates are negligibly small. It is known that for the non-reneging case, the $\tilde{c}\mu$ -rule is optimal in underload (the celebrated $c\mu$ -rule for a multi-class M/M/1 queue). The $\tilde{c}\mu/\theta = (c + d\theta)\mu/\theta$ index might however give an opposite priority rule when θ 's are very small, which explains the non-optimality of the $\tilde{c}\mu/\theta$ -rule when θ_k 's are very small.

Quadratic holding cost: Consider a system with 2 classes of customers. We assume quadratic holding costs $C_1(n) = \tilde{c}_{11}n + c_{21}n^2$ where, $\tilde{c}_{11} = (c_{11} + d_1 \theta_1)$, and $C_2(n) = \tilde{c}_{21}n + c_{22}n^2$, where $\tilde{c}_{21} = (c_{21} + d_2 \theta_2)$ and fix the following parameters: $\lambda = [4, 5]$, $\mu = [15, 17]$, $c_1 = [1, 4]$, $c_2 = [5, 1]$, $d = [2, 3]$, $\theta_1 = \epsilon_1 p$ and $\theta_2 = \epsilon_2 p$, where $\epsilon_1 = 0.05$ and $\epsilon_2 = 0.01$, and let p vary. Hence, $\rho = \sum_k \rho_k < 1$ so that the stability of the system is assured as $\theta_k \rightarrow 0$.

In Figures 3 (right) we plot the sub-optimality gap as p varies from 0 to 200, hence θ_1 and θ_2 range from $[0, 10]$ and $[0, 2]$, respectively. We observe for the θ -dependent indices a sub-optimality gap of 25% around $p = 0$. As θ grows large, this gap disappears however very fast. Note that the θ -independent indices work well, as we are in an underload scenario.

References

- [1] P. S. Ansell, K. D. Glazebrook, J. Niño-Mora, and M. O’Keeffe. Whittle’s index policy for a multi-class queueing system with convex holding costs. *Mathematical Methods of Operations Research*, 57(1):21–39, 2003.
- [2] N.T. Argon, S. Ziya, and R. Righter. Scheduling impatient jobs in a clearing system with insights on patient triage in mass-casualty incidents. *Prob. Eng. Inf. Sci.*, 22(3):301–332, 2010.
- [3] B. Ata and M.H. Tongarlak. On scheduling a multiclass queue with abandonments under general delay costs. *Queueing Systems*, 74:65–104, 2013.
- [4] R. Atar, C. Giat, and N. Shimkin. The $c\mu/\theta$ rule for many-server queues with abandonment. *Operation Research*, 58(5):1427–1439, 2010.

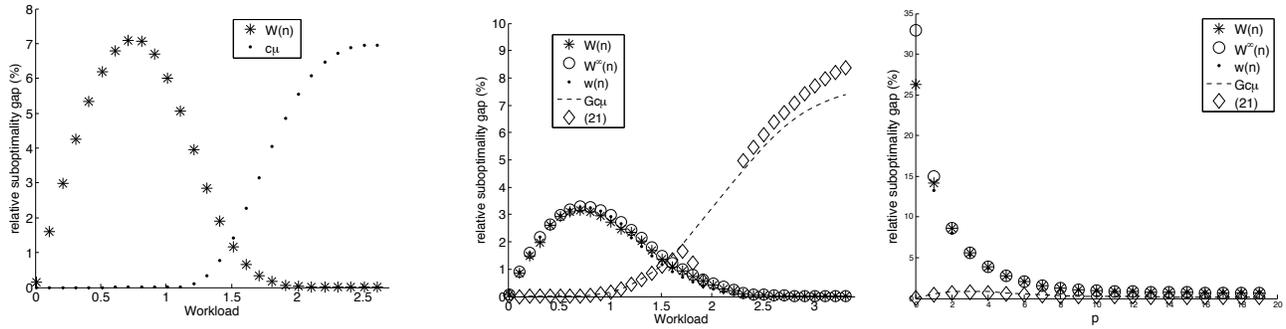


Figure 3: Left: sub optimality for linear holding cost, as ρ increases. Middle: sub optimality for quadratic holding cost as ρ increases. Right: sub optimality for quadratic holding cost as p ($\theta = p\epsilon$) varies.

- [5] R. Atar, C. Giat, and N. Shimkin. On the asymptotic optimality of the $c\mu/\theta$ rule under ergodic cost. *Queueing Systems*, 67:127–144, 2011.
- [6] F. Avram, D. Bertsimas, and M. Richard. Optimization of multiclass queuing networks: a linear control approach. *Stochastic Networks*, eds. F.P. Kelly and R.J. Williams, pages 199–234, 1995.
- [7] U. Ayesta, P. Jacko, and V. Novak. A nearly-optimal index rule for scheduling of users with abandonment. In *IEEE Infocom 2011*, 2011.
- [8] N. Bäuerle. Asymptotic optimality of tracking policies in stochastic networks. *Ann. Appl. Probab.*, 10:1065–1083, 2000.
- [9] N. Bäuerle and U. Rieder. Optimal control of single-server fluid networks. *Queueing Syst.*, 35:185–200, 2000.
- [10] S. Bhulai, H. Blok, and F.M. Spieksma. k computing queues with customer abandonment: optimality of a generalized $c\mu$ -rule by the smoothed rate truncation method. *Work in progress*, 2014.
- [11] S. Bhulai, A.C. Brooms, and F.M. Spieksma. On structural properties of the value function for an unbounded jump markov process with an application to a processor sharing retrial queue. *QUESTA*, 76:425–446, 2014.
- [12] C.F. Bispo. The single-server scheduling problem with convex costs. *Queueing Systems*, 73:261–294, 2013.
- [13] C. Buyukkoc, P. Varaya, and J. Walrand. The $c\mu$ rule revisited. *Adv. Appl. Prob.*, 17:237–238, 1985.
- [14] J.G. Dai. On positive harris recurrence of multiclass queueing networks: a unified approach via fluid limit models. *Annals of Applied Probability*, 5:49–77, 1995.
- [15] J.G. Dai and S. He. Many-server queues with customer abandonment: A survey of diffusion and fluid approximations. *J. Syst. Sci. Syst. Eng.*, 21:1–36, 2012.
- [16] D.G. Down, G. Koole, and M.E. Lewis. Dynamic control of a single server system with abandonments. *Queueing Systems*, 67:63–90, 2011.
- [17] A. Gajrat and A. Hordijk. Fluid approximation of a controlled multiclass tandem network. *Queueing Systems*, 35:349–380, 2000.
- [18] E. Gelenbe and I. Mitrani. *Analysis and Synthesis of Computer Systems*. London: Academic Press, 1980.
- [19] J.C. Gittins, K. Glazebrook, and R. Weber. *Multiclass Bandit Allocation Indices*. Wiley, 2011.
- [20] K.D. Glazebrook, P.S. Ansell, R.T. Dunn, and R.R. Lumley. On the optimal allocation of service to impatient tasks. *J. Appl. Prob.*, 41:51–72, 2004.
- [21] K.D. Glazebrook, D.J. Hodge, C. Kirkbride, and R.J. Minty. Stochastic scheduling: A short history of index policies and new approaches to index generation for dynamic resource allocation. *J. of Sched.*, 2013.
- [22] K.D. Glazebrook, C. Kirkbride, and J. Ouenniche. Index policies for the admission control and routing of impatient customers to heterogeneous service stations. *Operations Research*, 57:975–989, 2009.
- [23] J. Hasenbein and D. Perry (Eds). Special issue on queueing systems with abandonments. *Queueing Systems*, 75(2–4):111–384, 2013.
- [24] J. Kim and A. Ward. Dynamic scheduling of a GI/GI/1+GI queue with multiple customer classes. *Queueing Systems*, 75(2-4):339–384, 2013.
- [25] M. Larrañaga, U. Ayesta, and I.M. Verloop. Index policies for a multi-class queue with convex holding cost and abandonments. *LAAS technical report 14098*.
- [26] M. Larrañaga, U. Ayesta, and I.M. Verloop. Dynamic fluid-based scheduling in a multi-class abandonment queue. *Performance Evaluation*, 70:841–858, 2013.
- [27] A. Mandelbaum and S. Stolyar. Scheduling flexible servers with convex delay costs: Heavy-traffic optimality of the generalized $c\mu$ -rule. *Operations Research*, 52(6):836–855, 2004.
- [28] J. Niño-Mora. Dynamic priority allocation via restless bandit marginal productivity indices. *TOP*, 15(2):161–198, 2007.
- [29] M. L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. Wiley, 2005.
- [30] J. A. van Mieghem. Dynamic scheduling with convex delay costs: The generalized $c\mu$ rule. *The Annals of Applied Probability*, 5(3):808–833, 1995.
- [31] I.M. Verloop. Asymptotic optimal control of multiclass restless bandits. CNRS Technical Report, hal-00743781, September 2013.
- [32] I.M. Verloop and R. Núñez-Queija. Asymptotically optimal parallel resource assignment with interference. *Queueing Systems*, 65(1):43–92, 2010.
- [33] R.R. Weber and G. Weiss. On an index policy for restless bandits. *J. Appl. Prob.*, 27:637–648, 1990.
- [34] G. Weiss. On optimal draining of reentrant fluid lines. *Stochastic Networks*, eds. F.P. Kelly and R.J. Williams, pages 91–103, 1995.
- [35] P. Whittle. Restless bandits: Activity allocation in a changing world. *J. Appl. Prob.*, 25:287–298, 1988.