

Admission and routing of soft real-time jobs to multiclusters: Design and comparison of index policies

José Niño-Mora

Department of Statistics

Carlos III University of Madrid

28903 Getafe (Madrid), Spain

jnimora@alum.mit.edu, <http://orcid.org/0000-0002-2172-3983>

Published in *Computers & Operations Research*, vol. 39, 3431–3444, 2012

DOI: [10.1016/j.cor.2012.05.004](https://doi.org/10.1016/j.cor.2012.05.004)

Abstract

Motivated by time-sensitive e-service applications, we consider the design of effective policies in a Markovian model for the dynamic control of both admission and routing of a single class of real-time transactions to multiple heterogeneous clusters of web servers, each having its own queue and server pool. Transactions come with response-time deadlines, staying until completion if the latter are missed. Per job rejection and deadline-miss penalties are incurred. Since computing an optimal policy is intractable, we aim to design near optimal heuristic policies that are tractable for large-scale systems. Four policies are developed: the static optimal Bernoulli-splitting (BS) policy, and three index policies, based respectively on individually optimal (IO) actions, one-step policy improvement (PI), and restless bandit (RB) indexation. A computational study demonstrates that PI is the best of such policies, being consistently near optimal. In the pure-routing case, both the PI and RB policies are nearly optimal.

Keywords: e-services; admission control; routing; multiclusters; parallel multiserver queues; soft real-time; deadlines; index policies

MSC (2010): 90B22; 90C40

1 Introduction

Motivated by time-sensitive online data services (e-services) applications, this paper considers the design of effective policies for the dynamic control of both admission and routing of a single class of real-time transactions to multiple clusters of web servers, in the setting of a queueing model of Markov decision process (MDP) type. Service resources comprise n parallel clusters of servers, where each cluster $k \in \mathbb{K} \triangleq \{1, \dots, n\}$ has its own dedicated queue with unlimited buffer space for holding transaction requests waiting or being processed, and a pool of m_k identical servers, each working at a speed of μ_k (transactions per time unit). The system is heterogeneous in that both server speeds and server-pool sizes may differ across clusters. Transactions arrive as a Poisson process with rate λ , with their execution times being independent and exponentially distributed with unit rate. To ensure stability, it is assumed that the incoming load $\rho \triangleq \lambda / \sum_k m_k \mu_k$ is less than unity. Upon arrival of a transaction request, the system controller immediately and irrevocably decides whether to admit or reject it. If admitted, the controller further decides at once to which cluster to route it, basing decisions on the history of previous queue lengths and actions. Once in a cluster queue, pending transactions are scheduled on a first-come first-served (FCFS) basis.

Transactions are delay-sensitive of *soft real-time* type, coming with timeliness requirements on their individual response or waiting times given in the form of *soft deadlines* from their times of arrival, i.e., deadlines that can be missed without totally destroying the value of the transactions. Hence, transactions with missed deadlines remain in the system until completed.

The prime motivation for the study of such a system model comes from certain time-sensitive e-service applications, where it is desirable but not mandatory to meet given transaction deadlines, and which increasingly rely on heterogeneous multicluster web-server architectures. Examples of such soft real-time applications include e-commerce (see Bertini et al. (2010)), real-time databases, e.g., for online stock trading (see Kao and García-Molina (1996) and Kang et al. (2007)), and scientific computing (see Zhu (2001) and Plankensteiner et al. (2010)). See also He et al. (2006), which considers a similar model to the present one, yet focusing on the *pure-routing* case, i.e., without admission control.

It should be mentioned that in other important applications, notably in call centers, transactions are subject to *firm deadlines*, meaning that customers abandon if they are put on hold for too long. Use of index policies based on restless bandits for the resulting models with abandonment, or reneging, has been proposed in Niño-Mora (2007). A thorough study of such models is the subject of another paper by the author, as yet unpublished.

The incorporation of admission control in the present model is motivated by the observation that, at times when the system is heavily congested, admitting newly arriving transactions will almost guarantee that their deadlines will be missed. Hence, in such circumstances it may be beneficial to reject new transactions if doing so is not too costly.

The main performance metrics to evaluate the Quality-of-Service (QoS) provided to users under a given control policy π in such a system are the *rejection ratio* p^π , which is the long-run average fraction of transactions that are rejected upfront, and the *deadline miss ratio* q^π , which is the long-run average fraction of transactions that are admitted only to miss their deadlines. While one would like to have both p^π and q^π small, it is intuitively clear that there is trade-off between them. We can visualize such a trade-off as the lower boundary in the *region of achievable performance* of (p^π, q^π) pairs, spanned under all admissible policies π , which is displayed in Figure 1 for the instance with $n = 2$ clusters with server-pool sizes $\mathbf{m} = (4, 8)$, server speeds $\boldsymbol{\mu} = (5, 3)$, and $\rho = 0.95$.

A common approach to policy design in such systems is to specify in advance upper bounds \bar{p} and \bar{q} on the metrics p^π and q^π , and then try to construct a policy π that satisfies the resulting constraints $p^\pi \leq \bar{p}$ and $q^\pi \leq \bar{q}$. Another common approach is to aim to minimize one of the metrics, say p^π , subject to an upper bound constraint on the other, e.g., $q^\pi \leq \bar{q}$.

In contrast, in this work we aim to jointly optimize the pair (p^π, q^π) , by setting economic incentives to the service provider, considering a *performance objective* to minimize of the form $Rp^\pi + Cq^\pi$, where $R > 0$ represents a cost per rejected transaction, and $C > 0$ is a cost per missed deadline. Notice that the pure-routing case considered in He et al. (2006) can be seen as the special case of the present model where the rejection cost is R is sufficiently large.

Since minimizing the cost $Rp^\pi + Cq^\pi$ is equivalent to maximizing the net reward $R(1 - p^\pi) - Cq^\pi$, such an objective models a Service Level Agreement (SLA) contract, which sets the fee R to be collected by the service




Figure 1: The achievable performance region of (p^π, q^π) pairs.

provider from each admitted user, and incorporates an economic compensation scheme by which a refund C is issued to a user whose QoS target deadline is missed. While the above *performance optimization problem* can be formulated as a countable-state average-cost MDP, computing an optimal policy by solving numerically the corresponding *dynamic programming* (DP) equations is computationally intractable, even for a moderate number of clusters. Hence, the main aim of this paper is to design high-performance heuristic policies that can be computed with low complexity in large-scale systems.

In the sequel we will focus, with no loss of generality, on the case $C = 1$, which results by normalizing the above objective by dividing it by C . Further, we will use interchangeably the terms “queue” and “cluster,” since each cluster has its own queue, and we will use the term “job” instead of “transaction.”

1.1 Related prior work

Most related prior work on soft-real time systems has focused on pure-routing models, and has only considered *static* policies, which do not use system state information. Static policies decompose routing decisions into a *load balancing* or *workload allocation* rule, which determines the proportion of jobs to be assigned to each queue, and a *job dispatching* rule, which specifies where to send each job. An admission policy may be incorporated by fixing the proportion of jobs to be rejected. A simple static policy is a *Bernoulli splitting* (BS), which splits the arrival stream according to fixed probabilities. Kallmes and Cassandra (1995) addresses computation of the optimal BS in the special case of the present model, including admission and routing, where each queue has a single server and the rejection and deadline miss penalties are equal ($R = C = 1$). This calls for solving a convex separable *nonlinear program* (NLP). He et al. (2006) shows how to compute the optimal workload allocation for the special pure-routing case of the present model, considering BS as well as round-robin job dispatching.

As for dynamic policies, which use some degree of state information, prior work has addressed the *performance evaluation* of the deadline-miss ratio for some ad hoc routing policies, in homogeneous systems with parallel single-server queues. Kao and García-Molina (1996) considers routing policies that use deadline information, which are shown to outperform symmetric BS routing. Zhu (2001) reports on a comparative simulation study of four routing policies for the same model, one of which is the *shortest queue routing* (SQR) rule. It is shown that SQR is the best policy among those considered.

In different models from the one considered here, index policies have been proposed for routing to parallel queues based on one-step *policy improvement* (PI) over the optimal BS (see Krishnan (1990)), and also based on *restless bandit* (RB) indexation, introduced in general in Whittle (1988) and first applied in Niño-Mora (2002) to problems of admission control and routing to parallel queues (see also Niño-Mora (2007)).

Effective heuristic index policies have also been proposed and investigated for dynamic routing in certain call-center models that incorporate both multiple customer classes and multiple server pools with different skills.

See, e.g., [Bhulai \(2009\)](#), which derives a PI-based policy that is shown to perform extremely well, and [Gurvich and Whitt \(2010\)](#), where the proposed policies can achieved desired QoS levels. Note, however, that, even in the case of a single customer class the models considered in such papers are fundamentally different from the one addressed here, not only because they do not incorporate rejections and consider either no deadlines or firm deadlines (with abandonment) rather than soft deadlines, but more importantly because they assume that customers (from each class) join a *common queue* upon arrival, to be *later* routed to available servers. In contrast, in the present model customer transactions that are admitted are immediately routed upon arrival to one of *multiple parallel queues*, each corresponding to a multiserver cluster.

1.2 Extending prior work

This paper extends the aforementioned research in the following ways: (1) the model incorporates dynamic control of admission, rather than considering only static or only pure-routing policies; (2) instead of addressing the performance evaluation of given ad hoc policies, the aim is *performance optimization*, i.e., to *design* heuristic policies that can be computed with low complexity and come close to optimizing the performance objective of concern; (3) the focus is on systems with parallel *heterogeneous* multi-server queues, motivated by the prevalence of heterogeneity in real multicluster systems; (4) the computation of the optimal BS is extended from the special case $R = 1$ (recall that $C \equiv 1$) to the general case where R may differ from 1; (5) new dynamic index policies type are introduced, based on three different methods (individually optimal decisions, one-step policy improvement, and restless bandits); and, finally, all the heuristic policies considered are benchmarked against the optimal performance in a numerical study.

We will focus on the intuitively appealing class of *index policies*, which determine admission and routing actions on the basis of n numeric *routing indices*, where index $\nu_k(x_k)$ is attached to queue $k \in \mathbb{K}$ as a non-decreasing function of its current *state* x_k , the latter giving the number of jobs waiting or being served in the cluster's queue. The resulting routing policy sends an admitted job to a queue with currently lowest index value, breaking ties arbitrarily.

Such routing policies aim to approximate the optimal routing policy, which is shown below in Section 2.3 to have the following form: a newly arriving job that finds the system in the joint state $\mathbf{x} = (x_k)$, if admitted, is assigned to a queue k attaining the minimum value of a certain quantity $\nu_k^*(\mathbf{x})$, which is thus an *optimal routing index* which depends on the joint state \mathbf{x} , rather than on the state x_k of queue k alone. The optimal index $\nu_k^*(\mathbf{x})$ measures the long-run expected increment in total costs accrued that results by routing arrivals in state \mathbf{x} to queue k , given that all other actions are taken optimally.

It is further shown in Section 2.3 that it is optimal to reject an arrival finding the system in the joint state $\mathbf{x} = (x_k)$ iff $R \leq \gamma^*(\mathbf{x})$, where $\gamma^*(\mathbf{x})$ is the *optimal global admission index* defined by $\gamma^*(\mathbf{x}) \triangleq \min_k \nu_k^*(\mathbf{x})$.

Hence, for consistency with such optimal admission policies, given heuristic routing indices $\nu_k(x_k)$ we will base admission decisions on a corresponding global admission index $\gamma(\mathbf{x}) \triangleq \min_k \nu_k(x_k)$. The resulting policy admits an arrival finding the system in state \mathbf{x} if $\gamma(\mathbf{x}) < R$, and rejects it otherwise.

This raises the issue of how to define such routing indices so that they can be computed quickly and yield high-performance policies. An insight on which we will draw is that the routing index $\nu_k(x_k)$ should be a proxy measure for the optimal index $\nu_k^*(\mathbf{x})$, which, as pointed out above, is the incremental cost of routing an arrival to queue k in state \mathbf{x} .

To design the routing indices $\nu_k(x_k)$ we will deploy three alternative general methods that have been proposed in the literature for obtaining priority indices. Two such methods correspond to different assessments of the incremental cost of assigning an arrival to a queue. At one extreme, the *individually optimal* (IO) method *underestimates* such a cost by disregarding future arrivals. At the opposite extreme, the RB method *overestimates* such a cost by taking into account all future arrivals, as though there were no other queues to share the load. Such methods yield routing indices $\nu_k^{\text{IO}}(x_k)$ and $\nu_k^{\text{RB}}(x_k)$ that can be evaluated separately for each cluster k , as they depend only on its individual parameters. The third method we will consider, which is to carry out one PI step of the policy iteration algorithm for MDPs, starting from the optimal BS, gives routing indices $\nu_k^{\text{PI}}(x_k)$ that depend on the parameters of all clusters. Such a method has been used in [Krishnan \(1990\)](#) to obtain a heuristic index policy for job assignment to parallel multiserver queues with a minimum average response-time objective.

To the best of the author's knowledge, none of the above policy-design methods has been applied in prior work to the present model. Only the optimal BS method has been applied to special cases: in Kallmes and Cassandras (1995) for the case of single-server queues $m_k \equiv 1$, and in He et al. (2006) for the pure-routing case.

The main objective of the present work is to deploy the BS, PI, IO, and RB policy design methods in the present model, addressing the efficient computation of the resulting policies, and to identify the best among such policies through a numerical study.

The remainder of the paper is organized as follows. Section 2 describes the model and formulates the optimal control problem. Section 3 addresses the efficient evaluation of the IO routing index in the cases of constant, uniform, and exponentially distributed job deadlines. Section 4 discusses the PI routing index, as well as the prerequisite computation of the optimal BS. Section 5 considers the RB routing index. Section 6 reports on the results of a numerical study benchmarking the above policies, against each other and against the optimal performance. Section 7 concludes.

2 The model

2.1 Model description and problem formulation

Consider a system with n parallel clusters of servers, with cluster $k \in \mathbb{K}$ having its own queue with unlimited buffer space and m_k identical exponential servers, each working at rate μ_k . Jobs arrive as a Poisson stream with rate λ , with job interarrival and service times being independent. Upon arrival of a job, it is decided (i) whether to admit or reject it; and, if admitted, (ii) to which queue to route it.

Jobs come with a (relative) deadline τ , giving the time interval from arrival to deadline expiration, which is independent of arrival and service times, being drawn in i.i.d. fashion from a general cumulative distribution function (CDF) $F(t)$ with finite mean $\mathbb{E}[\tau] = 1/\theta$. We will consider the cases of deterministic, uniform and exponential deadlines, focusing on the case of *deadlines to the end of service*, where τ represents an individual *response-time deadline*, although the ensuing analyses readily extend to the case of deadlines to the beginning of service. Admitted jobs whose deadlines are missed stay nevertheless in the system until service completion. A cost of $R > 0$ is incurred per rejected job, whereas a cost of $C = 1$ is incurred per missed deadline.

Denote by $X_k(t)$ the state of queue k at time t , giving the number of jobs its queue holds waiting or being served, and by $a_k(t) \in \{0, 1\}$ the binary *action* that is set to 0 if a job arriving at time t would be routed to queue k . One may imagine that each queue k has its own *entry gate* and *gatekeeper*, who can open ($a_k(t) = 0$) or shut ($a_k(t) = 1$) the queue's entry gate. The requirement that a job can be routed to at most one queue is thus formulated by the constraint that at most one gate can be open at any time:

$$\sum_{k \in \mathbb{K}} (1 - a_k(t)) \leq 1, \quad t \geq 0. \quad (1)$$

The following condition will be assumed to hold.

Assumption 1. *At each queue k , the queue discipline is FCFS.*

Assumption 1, which excludes use of alternative queue disciplines such as *earliest deadline first*, is needed to ensure that the response-time distribution of a job assigned to a queue k is determined only by the queue state x_k found on arrival. Denoting by $T_k(x_k)$ a random variable with such a *conditional response-time distribution*, and letting τ be the job's deadline, the expected deadline-miss cost incurred by a job assigned to queue k in state x_k is

$$P_k^{\text{miss}}(x_k) \triangleq \mathbb{P}\{T_k(x_k) > \tau\} = \int_0^\infty \mathbb{P}\{T_k(x_k) > t\} dF(t). \quad (2)$$

We will find convenient to formulate the lump costs incurred due to job rejections and deadline misses in terms of continually accruing costs, by considering that the total cost rate incurred per unit time when the joint action $\mathbf{a} = (a_k)_{k \in \mathbb{K}}$ is taken in the joint state $\mathbf{x} = (x_k)_{k \in \mathbb{K}}$ is given by

$$\lambda R \left[1 - \sum_{k \in \mathbb{K}} (1 - a_k) \right] + \sum_{k \in \mathbb{K}} \lambda P_k^{\text{miss}}(x_k) (1 - a_k). \quad (3)$$

We will also find convenient to reformulate (3) to separable form as

$$-(n-1)\lambda R + \sum_{k \in \mathbb{K}} \{ \lambda R a_k + \lambda P_k^{\text{miss}}(x_k)(1-a_k) \}. \quad (4)$$

Consider the problem of finding an *average-cost optimal policy*, which minimizes the long-run expected average cost per unit time. Disregarding the constant $-(n-1)\lambda R$ in (4), we can formulate such a problem as

$$\min_{\pi \in \Pi} \limsup_{T \rightarrow \infty} \frac{1}{T} \mathbb{E}_{\mathbf{x}^0}^{\pi} \left[\int_0^T \sum_{k \in \mathbb{K}} \{ \lambda R a_k(t) + \lambda P_k^{\text{miss}}(X_k(t))(1-a_k(t)) \} dt \right], \quad (5)$$

where “lim sup” is the limit superior, $\mathbb{E}_{\mathbf{x}^0}^{\pi}[\cdot]$ denotes expectation under policy π conditioned on the initial joint state being equal to $\mathbf{x}^0 = (x_k^0)_{k \in \mathbb{K}}$, and Π is the class of *admissible policies*, which are allowed to depend only on the history of system states and actions encountered in the past, and may be randomized.

By standard results, problem (5) has an optimal policy that is *stationary deterministic*, i.e., in which the action to take at every state \mathbf{x} is a deterministic function of \mathbf{x} . Both optimal policies and the optimal cost rate are independent of the initial state, and they are characterized by the *DP average-cost optimality equation* (see (6) below). Yet, finding an optimal policy by numerical solution of such equations is, due to the curse of dimensionality, computationally intractable for instances with even a moderate number n of queues. Hence, we aim to design and compare heuristic index policies that can be computed with low complexity and perform well, as discussed in Section 1.

2.2 Charging deadline-miss costs via action-independent holding costs

Application of two of the index design methods considered in this paper, PI and RB (see Sections 4 and 5), would be considerably simplified if, rather than charging the deadline-miss costs incurred at each queue k in terms of action-dependent cost rates as in formulation (5), they were charged in terms of action-independent holding costs, which continually accrued at a certain rate $h_k(x_k)$ per unit time when the queue holds x_k jobs.

Hence, to prepare the ground for the analyses in Sections 4 and 5, we next address the reformulation of problem (5) to an equivalent problem where deadline-miss costs are accounted for via action-independent holding cost rates $h_k(x_k)$. For such a purpose, we start by formulating the *DP optimality equations* characterizing the optimal policies for average-cost problem (5), which are stated in terms of variables v and $b(\mathbf{x})$: for every system state \mathbf{x} ,

$$v = \min \begin{cases} n\lambda R - \sum_{l \in \mathbb{K}} \bar{\mu}_l(x_l) \Delta_l b(\mathbf{x}) \\ (n-1)\lambda R + \lambda P_k^{\text{miss}}(x_k) + \lambda \Delta_k b(\mathbf{x} + \mathbf{e}_k) - \sum_{l \in \mathbb{K}: x_l \geq 1} \bar{\mu}_l(x_l) \Delta_l b(\mathbf{x}), k \in \mathbb{K}, \end{cases} \quad (6)$$

where we write $\bar{\mu}_k(x_k) \triangleq \min(x_k, m_k) \mu_k$ and $\Delta_k b(\mathbf{x}) \triangleq b(\mathbf{x}) - b(\mathbf{x} - \mathbf{e}_k)$, \mathbf{e}_k being the k th unit coordinate vector in \mathbb{R}^n . By standard results, the DP equations (6) have a solution v^* and $b^*(\mathbf{x})$, which is unique up to a constant additive term for the $b^*(\mathbf{x})$, where v^* is the optimal average cost per unit time and $b^*(\mathbf{x})$ is an optimal relative cost function.

Let us now define the holding cost rate $h_k(\cdot)$ for each queue $k \in \mathbb{K}$ by

$$h_k(0) \triangleq 0, \quad h_k(x_k) \triangleq \bar{\mu}_k(x_k) P_k^{\text{miss}}(x_k - 1), \quad x_k \geq 1, \quad (7)$$

and consider the following modified admission control and routing problem, where two types of costs are incurred: rejection costs, as before, and holding costs, which accrue at each queue k at rate $h_k(x_k)$ while it holds x_k jobs:

$$\min_{\pi \in \Pi} \limsup_{T \rightarrow \infty} \frac{1}{T} \mathbb{E}_{\mathbf{x}^0}^{\pi} \left[\int_0^T \sum_{k \in \mathbb{K}} \{ R \lambda a_k(t) + h_k(X_k(t)) \} dt \right]. \quad (8)$$

The optimal policies for the modified problem (8) are characterized by its DP optimality equations: for every system state \mathbf{x} ,

$$\tilde{v} = \min \begin{cases} n\lambda R + \sum_{l \in \mathbb{K}} h_l(x_l) - \sum_{l \in \mathbb{K}} \bar{\mu}_l(x_l) \Delta_l \tilde{b}(\mathbf{x}) \\ (n-1)\lambda R + \sum_{l \in \mathbb{K}} h_l(x_l) + \lambda \Delta_k \tilde{b}(\mathbf{x} + \mathbf{e}_k) - \sum_{l \in \mathbb{K}: x_l \geq 1} \bar{\mu}_l(x_l) \Delta_l \tilde{b}(\mathbf{x}), k \in \mathbb{K}. \end{cases} \quad (9)$$

Equations (9) have a solution \tilde{v}^* and $\tilde{b}^*(\mathbf{x})$, which is unique up to a constant additive term for the $\tilde{b}^*(\mathbf{x})$, where \tilde{v}^* is the optimal average cost per unit time and $\tilde{b}^*(\mathbf{x})$ is an optimal relative cost function.

The following result ensures that, charging costs at each queue using the holding cost rates defined in (7), we obtain an equivalent problem.

Proposition 1. *Problems (5) and (8) are equivalent, having the same optimal policies and average cost per unit time.*

Proof. For each queue k , define the function $\tilde{C}_k(\cdot)$ by the recursion

$$\tilde{C}_k(0) \triangleq 0, \quad \tilde{C}_k(x_k + 1) \triangleq \tilde{C}_k(x_k) + P_k^{\text{miss}}(x_k), \quad x_k \geq 0. \quad (10)$$

The result follows by checking that $v^* \triangleq \tilde{v}^*$ and $b^*(\mathbf{x}) \triangleq \tilde{b}^*(\mathbf{x}) - \sum_k \tilde{C}_k(x_k)$ satisfy the DP equations (6) iff \tilde{v}^* and the $\tilde{b}^*(\mathbf{x})$ satisfy (9). \square

2.3 The structure of optimal admission and routing policies

In light of (9) and Proposition 1, an optimal policy for problem (5) can be formulated in terms of the relative cost increments $\Delta_k \tilde{b}^*(\mathbf{x})$ for problem (8): reject an arrival in state \mathbf{x} if $R \leq \min_k \Delta_k \tilde{b}^*(\mathbf{x} + \mathbf{e}_k)$; otherwise, admit the arrival and route it to a queue k with minimum $\Delta_k \tilde{b}^*(\mathbf{x} + \mathbf{e}_k)$. Hence, optimal admission actions are characterized by a global admission index, given by $\gamma^*(\mathbf{x}) \triangleq \min_k \Delta_k \tilde{b}^*(\mathbf{x} + \mathbf{e}_k)$, whereas optimal routing actions are characterized by indices $\nu_k^*(\mathbf{x}) \triangleq \Delta_k \tilde{b}^*(\mathbf{x} + \mathbf{e}_k)$ that depend on the system state \mathbf{x} .

Given the unavailability of such optimal admission and routing indices, this paper compares, in terms of the performance of the resulting policies, several alternative ways of defining a global admission index $\gamma(\mathbf{x})$ and routing indices $\nu_k(x_k)$ that can be evaluated with low computational complexity.

3 IO index

The simplest dynamic routing policy we consider is to route an arrival to a queue where the probability of missing its deadline, conditioned on the number of jobs found on such a queue upon arrival, is minimal. The resulting IO routing index is given by such a conditional deadline-miss probability, being $\nu_k^{\text{IO}}(x_k) \triangleq P_k^{\text{miss}}(x_k)$, where $P_k^{\text{miss}}(x_k)$ is defined in (2). Note that $\nu_k^{\text{IO}}(0) = \dots = \nu_k^{\text{IO}}(m_k - 1)$, and that $\nu_k^{\text{IO}}(x_k)$ converges monotonically up to unity as $x_k \rightarrow \infty$:

$$\bar{\nu}_k^{\text{IO}}(\infty) = \lim_{x_k \rightarrow \infty} \nu_k^{\text{IO}}(x_k) = \lim_{x_k \rightarrow \infty} P_k^{\text{miss}}(x_k) = 1. \quad (11)$$

Hence, the admission policy based on the IO routing index, having admission index $\gamma^{\text{IO}}(\mathbf{x}) = \min_k \nu_k^{\text{IO}}(x_k)$, will not reject any job in the case $R \geq 1$.

We next address the efficient evaluation of $P_k^{\text{miss}}(x_k)$ in the cases of constant, uniform, and exponential response-time deadlines. Since the focus is on a single queue, the label k is dropped below from the notation, writing, e.g., $P^{\text{miss}}(x)$.

3.1 Constant response-time deadlines

In the case where jobs have a constant response-time deadline $\tau \equiv t$, $P^{\text{miss}}(x)$ equals $P_m(x; t) \triangleq \mathbb{P}\{T(x) > t\}$, where m is the number of servers in the queue and $T(x)$ denotes a random variable distributed as the response

time (under FCFS) of an arrival who finds x jobs present in the queue. To evaluate $P_m(x; t)$, we will need the following probabilities, where $A_m(j) \sim \text{Erlang}(j, m\mu)$:

$$Q_m(j; t) \triangleq \mathbb{P}\{A_m(j) > t\} = e^{-m\mu t} \sum_{l=0}^{j-1} \frac{(m\mu t)^l}{l!}, \quad j = 1, 2, \dots \quad (12)$$

For fixed m and t , the $Q_m(j; t)$ are efficiently evaluated by the following second-order linear recursion, where $\Delta Q_m(j; t) \triangleq Q_m(j; t) - Q_m(j-1; t)$:

$$\begin{aligned} Q_m(1; t) &= e^{-m\mu t}, \quad \Delta Q_m(2; t) = m\mu t Q_m(1; t) \\ \Delta Q_m(j+1; t) &= (m\mu t/j) \Delta Q_m(j; t), \quad j \geq 2. \end{aligned} \quad (13)$$

The following result evaluates $P_m(x; t)$ in terms of $Q_m(\cdot, t)$ and $Q_{m-1}(\cdot, t)$. Along with (13), it allows us to compute $P^{\text{miss}}(0), \dots, P^{\text{miss}}(x)$ in $O(x)$ time.

Proposition 2. (a) $P_1(x; t) = Q_1(x+1; t)$, for $x \geq 0$.

(b) If $m \geq 2$, then $P_m(x; t) = e^{-\mu t}$ for $0 \leq x < m$; and, for $x \geq m$,

$$P_m(x; t) = Q_m(x-m+1; t) + e^{-\mu t} \frac{1 - Q_{m-1}(x-m+1; t)}{(1 - 1/m)^{x-m+1}}.$$

Proof. (a) If $m = 1$, $T(x) \sim \text{Erlang}(x+1, \mu)$, which yields the result.

(b) If $m \geq 2$, then $P_m(x; t) = e^{-\mu t}$ for $0 \leq x < m$, as in this case the arriving job finds some free server. Consider now the case $x \geq m$. Then, the conditional response time is decomposed as $T(x) \sim A + \xi$, where $A \sim \text{Erlang}(x-m+1, m\mu)$ and $\xi \sim \text{Exp}(\mu)$ are independent random variables giving the job's conditional waiting time and its service time, respectively. Hence, denoting by $f_A(t)$ the density of A , and writing $j = x - m + 1$,

$$\begin{aligned} P_m(x; t) &= \mathbb{P}\{A + \xi > t\} \\ &= \mathbb{P}\{A + \xi > t \mid A > t\} \mathbb{P}\{A > t\} + \int_0^t \mathbb{P}\{A + \xi > t \mid A = s\} f_A(s) ds \\ &= \mathbb{P}\{A > t\} + \int_0^t \mathbb{P}\{\xi > t - s\} f_A(s) ds = Q_m(j; t) + \int_0^t e^{-\mu(t-s)} f_A(s) ds. \end{aligned}$$

On the other hand, letting $A' \sim \text{Erlang}(j, (m-1)\mu)$, we have

$$\begin{aligned} \int_0^t e^{-\mu(t-s)} f_A(s) ds &= \int_0^t e^{-\mu(t-s)} (m\mu)^j \frac{s^{j-1}}{(j-1)!} e^{-m\mu s} ds \\ &= \frac{e^{-\mu t}}{(1 - 1/m)^j} \int_0^t ((m-1)\mu)^j \frac{s^{j-1}}{(j-1)!} e^{-(m-1)\mu s} ds \\ &= \frac{e^{-\mu t}}{(1 - 1/m)^j} \int_0^t f_{A'}(s) ds = e^{-\mu t} \frac{1 - Q_{m-1}(j; t)}{(1 - 1/m)^j}. \end{aligned}$$

□

3.2 Uniform response-time deadlines

Assuming that jobs have response-time deadlines $\tau \sim \text{Unif}[t_1, t_2]$, $P^{\text{miss}}(x) = P_m^*(x)/(t_2 - t_1)$, where $P_m^*(x) \triangleq \int_{t_1}^{t_2} P_m(x; t) dt$. To evaluate $P_m^*(x)$, we will use quantities $Q_m^*(j) \triangleq \int_{t_1}^{t_2} Q_m(j; t) dt$ and $R_{m-1}^*(j) \triangleq \int_{t_1}^{t_2} e^{-\mu t} Q_{m-1}(j; t) dt$.

Both $Q_m^*(j)$ and $R_{m-1}^*(j)$ are readily evaluated in terms of the $Q_m(\cdot; \cdot)$ and $Q_{m-1}(\cdot; \cdot)$ using integration by parts, which yields

$$Q_m^*(j) = t_2 Q_m(j; t_2) - t_1 Q_m(j; t_1) + \frac{j}{m\mu} [Q_m(j+1; t_1) - Q_m(j+1; t_2)], \quad (14)$$

and

$$R_{m-1}^*(j) = \frac{1}{\mu} \left[e^{-\mu t_1} Q_{m-1}(j; t_1) - e^{-\mu t_2} Q_{m-1}(j; t_2) - \frac{Q_m(j; t_1) - Q_m(j; t_2)}{(1 + 1/(m-1))^j} \right]. \quad (15)$$

The next result, which follows from Proposition 2, evaluates $P_m^*(x)$ in terms of $Q_m^*(\cdot)$ and $R_{m-1}^*(\cdot)$. Along with (14), (15), and the results in Section 3.1, this allows us to evaluate $P^{\text{miss}}(x)$, for $i = 0, \dots, j$, in $O(j)$ time.

Proposition 3. (a) $P_1^*(x) = Q_1^*(x+1)$, for $x \geq 0$.

(b) If $m \geq 2$, then $P_m^*(x) = (e^{-\mu t_1} - e^{-\mu t_2})/\mu$ for $0 \leq x < m$; for $x \geq m$,

$$P_m^*(x) = Q_m^*(x-m+1) + \frac{P_m^*(0) - R_{m-1}^*(x-m+1)}{(1-1/m)^{x-m+1}}.$$

3.3 Exponential response-time deadlines

Assuming that jobs have response-time deadlines $\tau \sim \text{Exp}(\theta)$, $P^{\text{miss}}(x) = \theta P_m^*(x; \theta)$, where $P_m^*(x; s) \triangleq \int_0^\infty e^{-st} P_m(x; t) dt$. To evaluate $P_m^*(x; \theta)$ we will use quantities $Q_m^*(j; s) \triangleq \int_0^\infty e^{-st} Q_m(j; t) dt$, which, by standard results, are evaluated in closed form as

$$Q_m^*(j; s) = \frac{1}{s} \left[1 - \left(\frac{m\mu}{s + m\mu} \right)^j \right], \quad j = 1, 2, \dots$$

The next result follows from Proposition 2 by taking Laplace transforms.

Proposition 4. (a) $P_1^*(x; \theta) = Q_1^*(x+1; \theta)$, for $x \geq 0$.

(b) If $m \geq 2$, then $P_m^*(x; \theta) = 1/(\mu + \theta)$ for $0 \leq x < m$; and, for $x \geq m$,

$$P_m^*(x; \theta) = Q_m^*(x-m+1; \theta) + \frac{\frac{1}{\mu + \theta} - Q_{m-1}^*(x-m+1; \mu + \theta)}{(1-1/m)^{x-m+1}}.$$

4 PI index

We next consider the PI method (see Krishnan (1990)), which has not been applied before to the present model. This method consists of two stages: (1) computing the optimal BS of the arrival stream; and (2) carrying out one step of the PI algorithm for MDPs, starting from the optimal BS.

4.1 First stage of the PI method: computing the optimal BS

The first stage is to compute the optimal BS of the arrival stream. In a BS, the actions taken upon job arrivals are independently drawn according to fixed probabilities: a job is rejected with probability λ_0/λ , and is assigned to queue k with probability λ_k/λ , for $k \in \mathbb{K}$. Hence, the input to each queue k is a Poisson process with rate λ_k , with the queues evolving independently. Consider the resulting M/M/ m_k queue k , which has offered load $r_k(\lambda_k) \triangleq \lambda_k/\mu_k$ and utilization factor $\rho_k(\lambda_k) \triangleq \lambda_k/(m_k\mu_k)$. Assuming that $\rho_k(\lambda_k) < 1$, let $\tilde{T}_k(\lambda_k)$ be a random variable with the queue's steady-state sojourn-time distribution, whose CDF $S_k(t; \lambda_k) \triangleq \mathbb{P}\{\tilde{T}_k(\lambda_k) \leq t\}$ is given by (see Gross et al. (2008, p. 72))

$$S_k(t; \lambda_k) = 1 - e^{-\mu_k t} - G_k(\lambda_k)(e^{-(m_k\mu_k - \lambda_k)t} - e^{-\mu_k t}), \quad (16)$$

where, denoting by $E_{2, m_k}(r)$ the Erlang-C formula for the probability that a job has to wait in the M/M/ m_k queue with offered load r ,

$$G_k(\lambda_k) \triangleq \frac{E_{2, m_k}(r_k(\lambda_k))}{1 - m_k + r_k(\lambda_k)}. \quad (17)$$

Yet, note that (16) is only valid when $r_k(\lambda_k) \neq m_k - 1$, i.e., when $\lambda_k \neq (m_k - 1)\mu_k$. In the case $r_k(\lambda_k) = m_k - 1$, we have

$$S_k(t; \lambda_k) = 1 - (1 + \mu_k E_{2, m_k}(m_k - 1)t)e^{-\mu_k t}. \quad (18)$$

Hence, the deadline-miss probability $\bar{P}_k^{\text{miss}}(\lambda_k) \triangleq \mathbb{P}\{\tilde{T}_k(\lambda_k) > \tau\}$ for a randomly arriving job that is admitted and assigned to queue k is given by

$$\begin{aligned} \bar{P}_k^{\text{miss}}(\lambda_k) &= \int_0^\infty (1 - S_k(t; \lambda_k)) dF(t) \\ &= \begin{cases} \phi(\mu_k) + G_k(\lambda_k)(\phi(m_k\mu_k - \lambda_k) - \phi(\mu_k)), & \lambda_k \neq (m_k - 1)\mu_k \\ \phi(\mu_k) - \mu_k E_{2, m_k}(m_k - 1)\phi'(\mu_k), & \lambda_k = (m_k - 1)\mu_k, \end{cases} \end{aligned} \quad (19)$$

where $\phi(s) \triangleq \mathbb{E}[e^{-s\tau}] = \int_0^\infty e^{-st} dF(t)$ is the Laplace transform of the deadline τ . We will consider the cases $\tau \equiv t$, with $\phi(s) = e^{-st}$, $\tau \sim \text{Unif}[t_1, t_2]$, with $\phi(s) = (e^{-t_1 s} - e^{-t_2 s})/(s(t_2 - t_1))$, and $\tau \sim \text{Exp}(\theta)$, with $\phi(s) = \theta/(s + \theta)$.

Thus, we can formulate the problem of finding the optimal BS as the following NLP, where $f_k(\lambda_k) \triangleq \lambda_k \bar{P}_k^{\text{miss}}(\lambda_k)$ is the *deadline-miss rate* for queue k :

$$\min \left\{ R\lambda_0 + \sum_{k \in \mathbb{K}} f_k(\lambda_k) : \sum_{k \in \mathbb{K}} \lambda_k = \lambda, \lambda_0 \geq 0, 0 \leq \lambda_k \leq m_k \mu_k, k \in \mathbb{K} \right\}. \quad (20)$$

Notice that, in problem (20), we allow a BS to saturate the queue at queue k by allocating to it an arrival rate equal to its service capacity ($\lambda_k = m_k \mu_k$). In such a case—which as we will see turns out to be relevant when λ and R are both large enough, and also in the pure-routing version of the problem, where $\lambda_0 \equiv 0$, when λ is large enough—the probability that a random job assigned to a saturated queue k misses its deadline is $\bar{P}_k^{\text{miss}}(m_k \mu_k) = 1$, and hence we define $f_k(m_k \mu_k) \triangleq m_k \mu_k$.

Kallmes and Cassandras (1995) presents a Lagrangian algorithm for solving the version of (20) with strict inequalities $\lambda_k < m_k \mu_k$ in the special case $R = 1$ with single-server queues $m_k = 1$, showing that taking $R = 1$ ensures that the optimal solution to (20) has $\lambda_k^* < m_k \mu_k$ for each k . As for the multiserver case $m_k \geq 2$, a Lagrangian algorithm is outlined in He et al. (2006), albeit only for the pure-routing case where $\lambda_0 \equiv 0$, and without discussing the possibility that some queue(s) may be saturated in an optimal solution.

We next discuss how to extend the results in such works to obtain the optimal BS $\boldsymbol{\lambda}^* = (\lambda_k^*)_{k=0}^n$ solving (20) for the present model, assuming that each function $\bar{P}_k^{\text{miss}}(\lambda_k)$ satisfies the following condition:

- (C) $\bar{P}_k^{\text{miss}}(\lambda_k)$ is continuous, increasing and strictly convex on $[0, m_k \mu_k]$, having a continuous and increasing derivative $(d/d\lambda_k)\bar{P}_k^{\text{miss}}(\lambda_k)$ on $(0, m_k \mu_k)$, with finite one-sided derivatives $(d/d\lambda_k)\bar{P}_k^{\text{miss}}(0^+)$ and $(d/d\lambda_k)\bar{P}_k^{\text{miss}}((m_k \mu_k)^-)$. $\alpha_k \triangleq f'_k(0^+)$ and $\beta_k \triangleq f'_k((m_k \mu_k)^-)$.

Under (C), which is readily verified in the cases of constant, uniform and exponential deadlines considered here (e.g., in the case $\tau \equiv t$ for a single-server queue, $\bar{P}_k^{\text{miss}}(\lambda_k) = e^{-(\mu_k - \lambda_k)t}$), the function $f_k(\lambda_k)$ inherits the same properties as $\bar{P}_k^{\text{miss}}(\lambda_k)$, since $f'_k(\lambda_k) = \bar{P}_k^{\text{miss}}(\lambda_k) + \lambda_k (d/d\lambda_k)\bar{P}_k^{\text{miss}}(\lambda_k)$ on $(0, m_k \mu_k)$, having one-sided derivatives $\alpha_k \triangleq f'_k(0^+)$ and $\beta_k \triangleq f'_k((m_k \mu_k)^-)$ that satisfy

$$0 < \alpha_k = \bar{P}_k^{\text{miss}}(0) = \phi(\mu_k) < 1 < 1 + m_k \mu_k (d/d\lambda_k)\bar{P}_k^{\text{miss}}((m_k \mu_k)^-) = \beta_k. \quad (21)$$

Hence, under (C), problem (20) has a unique optimal solution $\boldsymbol{\lambda}^*$.

Assume, by reordering if necessary, that $\alpha_1 \leq \alpha_2 \leq \dots \leq \alpha_n$. Note that, since $\alpha_k = \phi(\mu_k)$ and $\phi(s)$ is nonincreasing, such an ordering is equivalent to $\mu_1 \geq \mu_2 \geq \dots \geq \mu_n$. Further, let x_1, \dots, x_n be an ordering of the n queues for which $\beta_{x_1} \leq \beta_{x_2} \leq \dots \leq \beta_{x_n}$.

It follows from the above (cf. (21)) that, for any queue k and $\alpha \in [\alpha_k, \beta_k]$, the equation $f'_k(\lambda_k) = \alpha$ has a unique root $\Lambda_k^*(\alpha)$ in $[0, m_k \mu_k]$, where $\Lambda_k^*(\alpha)$ is continuous and increasing in α , and satisfies that $\Lambda_k^*(\alpha_k) = 0$ and $\Lambda_k^*(\beta_k) = m_k \mu_k$. For convenience of notation, we also define

$$\Lambda_k^*(\alpha) \triangleq 0, \quad \alpha < \alpha_k, \quad \text{and} \quad \Lambda_k^*(\alpha) \triangleq m_k \mu_k, \quad \alpha > \beta_k. \quad (22)$$

Now, by standard results in convex optimization, the optimal BS solving (20) is characterized by the following *Karush–Kuhn–Tucker (KKT) first-order optimality conditions*: a feasible BS λ^* for (20) is optimal iff there exists a Lagrange multiplier α^* for the equality constraint such that

$$\begin{aligned} R &\geq \alpha^*, & \text{with "=" if } \lambda_0^* > 0 \\ f'_k(\lambda_k^*) &= \alpha^*, & \text{for any queue } k \text{ with } 0 < \lambda_k^* < m_k \mu_k \\ \alpha_k &\geq \alpha^*, & \text{for any queue } k \text{ with } \lambda_k^* = 0 \\ \beta_k &\leq \alpha^*, & \text{for any queue } k \text{ with } \lambda_k^* = m_k \mu_k. \end{aligned} \tag{23}$$

Note that it follows from (23) that, since $\beta_k > 1$, in the case $R = 1$ the optimal BS will have $\lambda_k^* < m_k \mu_k$ for every queue k .

To compute the optimal BS λ^* for (20), we next present a ranking algorithm that extend to the current setting the algorithm given in Kallmes and Cassandras (1995, p. 321) for the special case $R = 1$ and $m_k \equiv 1$. The following algorithm determines an α^* such that the optimal BS satisfying (23) is given by $\lambda_k^* \triangleq \Lambda_k^*(\alpha^*)$ for $k \in \mathbb{K}$ —taking into account (22)—and $\lambda_0^* \triangleq \lambda - \sum_{k \in \mathbb{K}} \lambda_k^*$. Three cases need to be distinguished (note that case 3 does not arise when $R = 1$):

Case 1 (light load): for some $1 \leq l < n$, $\sum_{k=1}^l \Lambda_k^*(\alpha_l) < \lambda \leq \sum_{k=1}^l \Lambda_k^*(\alpha_{l+1})$. Letting $\tilde{\alpha}$ be the unique root in $(\alpha_l, \alpha_{l+1}]$ of $\sum_{k=1}^l \Lambda_k^*(\alpha) = \lambda$, take $\alpha^* \triangleq \min(\tilde{\alpha}, R)$. Note that $\lambda_k^* > 0$ for any $1 \leq k \leq l$ with $\alpha_k < R$, $\lambda_k^* = 0$ for any $1 \leq k \leq l$ with $\alpha_k \geq R$ and for $l < k \leq n$, and $\lambda_0^* = 0$ iff $\tilde{\alpha} \leq R$.

Case 2 (medium load): $\sum_{k \in \mathbb{K}} \Lambda_k^*(\alpha_n) < \lambda < \sum_{k \in \mathbb{K}} \Lambda_k^*(\beta_{x_1})$. Letting $\tilde{\alpha}$ be the unique root in (α_n, β_{x_1}) of $\sum_{k \in \mathbb{K}} \Lambda_k^*(\alpha) = \lambda$, take $\alpha^* \triangleq \min(\tilde{\alpha}, R)$. Note that $0 < \lambda_k^* < m_k \mu_k$ for any $k \in \mathbb{K}$ with $\alpha_k < R$, $\lambda_k^* = 0$ for any $k \in \mathbb{K}$ with $\alpha_k \geq R$, and $\lambda_0^* = 0$ iff $\tilde{\alpha} \leq R$.

Case 3 (heavy load): for some $1 \leq l < n$, $\sum_{k \in \mathbb{K}} \Lambda_k^*(\beta_{x_l}) \leq \lambda < \sum_{k \in \mathbb{K}} \Lambda_k^*(\beta_{x_{l+1}})$. Letting $\tilde{\alpha}$ be the unique root in $[\beta_{x_l}, \beta_{x_{l+1}})$ of $\sum_{k \in \mathbb{K}} \Lambda_k^*(\alpha) = \lambda$, take $\alpha^* \triangleq \min(\tilde{\alpha}, R)$. Note that $\lambda_k^* = m_k \mu_k$ for any $k \in \{x_1, \dots, x_l\}$ with $\beta_k \leq R$, $\lambda_k^* < m_k \mu_k$ for any $k \in \{x_1, \dots, x_l\}$ with $\beta_k > R$ and for $k \in \{x_{l+1}, \dots, x_n\}$, $\lambda_k^* > 0$ for any $k \in \mathbb{K}$ with $\alpha_k < R$, and $\lambda_0^* = 0$ iff $\tilde{\alpha} \leq R$.

Note that, in the light-load case, some queue(s) k will not be used at all under the optimal BS. In contrast, in the heavy-load case some queue(s) k may be saturated ($\lambda_k^* = m_k \mu_k$). Note also that, since we assumed $\rho \triangleq \lambda / \sum_k m_k \mu_k < 1$, no higher arrival rates λ need to be considered than those in Case 3. Another observation is that, if the rejection cost is $R \geq \beta_{x_n}$ then no job is rejected under the optimal BS. At the opposite extreme, if $R \leq \alpha_1$ then all jobs are rejected.

4.2 Second stage of the PI method: PI step

The second stage of the PI method is to carry out one PI step of the policy iteration algorithm for MDPs, starting from the optimal BS policy, which yields a better policy. Recall from Section 2.3 that an optimal policy, given by a solution to the DP optimality equation (9) for problem (8), takes the following form: upon arrival of a job finding the system in state \mathbf{x} , (i) reject it if $R \leq \min_k \Delta_k \tilde{b}^*(\mathbf{x} + \mathbf{e}_k)$; otherwise, (ii) send it to a queue with lowest $\Delta_k \tilde{b}^*(\mathbf{x} + \mathbf{e}_k)$.

The PI method replaces the relative cost increments $\Delta_k \tilde{b}^*(\mathbf{x})$ by those under the optimal BS λ^* , which we denote by $\Delta_k \tilde{b}(\mathbf{x}; \lambda^*)$. Now, since under a BS policy the n queues evolve independently, separately accruing costs in the modified problem (8) at rate $\lambda R a_k(t) + h_k(X_k(t))$ for queue k , we have $\tilde{b}(\mathbf{x}; \lambda^*) \triangleq \sum_k \tilde{b}_k(x_k; \lambda_k^*)$, where $\tilde{b}_k^*(x_k; \lambda_k^*)$ is the relative cost function for queue k , considered in isolation. Hence, $\Delta_k \tilde{b}(\mathbf{x}; \lambda^*) = \Delta b_k(x_k; \lambda_k^*) \triangleq \tilde{b}_k(x_k; \lambda_k^*) - \tilde{b}_k(x_k - 1; \lambda_k^*)$. The resulting PI policy is: upon arrival of a job finding the system in state \mathbf{x} , (i) reject it if $R \leq \min_k \Delta \tilde{b}_k(x_k + 1; \lambda_k^*)$; otherwise, (ii) route it to a queue k with lowest $\Delta \tilde{b}_k(x_k + 1; \lambda_k^*)$.

Hence, the PI policy is an index policy, having routing index

$$\nu_k^{\text{PI}}(x_k) \triangleq \Delta \tilde{b}_k(x_k + 1; \lambda_k^*), \tag{24}$$

for queue k , and global admission index $\gamma^{\text{PI}}(\mathbf{x}) \triangleq \min_k \nu_k^{\text{PI}}(x_k)$.

To evaluate the PI index $\nu_k^{\text{PI}}(x_k)$ for queue k , we must solve the corresponding *Poisson equations*, which are readily reformulated in terms of the index as

$$\begin{aligned} f_k(\lambda_k^*) - \lambda_k^* \nu_k^{\text{PI}}(0) &= h_k(0) \\ f_k(\lambda_k^*) - \lambda_k^* \nu_k^{\text{PI}}(x_k) + \bar{\mu}_k(x_k) \nu_k^{\text{PI}}(x_k - 1) &= h_k(x_k), \quad x_k \geq 1, \end{aligned} \quad (25)$$

where $\bar{\mu}_k(x_k) \triangleq \mu_k \min(x_k, m_k)$, and $h_k(x_k)$ is the holding cost rate defined in (7). Thus, (25) gives a first-order linear recursion for computing the index values $\nu_k^{\text{PI}}(0), \dots, \nu_k^{\text{PI}}(x_k)$ in linear, $O(x_k)$ time, given λ_k^* and $f_k(\lambda_k^*)$.

We write below $\rho_k^* \triangleq \rho_k(\lambda_k^*)$, $\bar{P}_k^{\text{miss},*} \triangleq \bar{P}_k^{\text{miss}}(\lambda_k^*)$ and $f_k^* \triangleq f_k(\lambda_k^*)$. Note that $f_k^* = \lambda_k^* \bar{P}_k^{\text{miss},*}$.

Remark 1. The following results are readily obtained.

1. In the case of a queue k that is not used under the optimal BS, i.e., having $\lambda_k^* = 0$, (25) and (7) yield that $\nu_k^{\text{PI}}(x_k) = P_k^{\text{miss}}(x_k)$, and hence the PI index reduces to the IO index $\nu_k^{\text{IO}}(x_k)$.
2. In the case $0 < \lambda_k^* < m_k \mu_k$, i.e., $0 < \rho_k^* < 1$, $\nu_k^{\text{PI}}(0) = \bar{P}_k^{\text{miss},*}$ and

$$\nu_k^{\text{PI}}(x_k) = \frac{1}{(\rho_k^*)^{x_k}} \left[\nu_k^{\text{PI}}(m_k - 1) + \sum_{j_k=m_k-1}^{x_k-1} (\rho_k^*)^{j_k} (\rho_k^* \bar{P}_k^{\text{miss},*} - P_k^{\text{miss}}(j_k)) \right], \quad x_k \geq m_k. \quad (26)$$

3. In the case of a saturated queue ($\rho_k^* = 1$), $\bar{P}_k^{\text{miss},*} = 1$, $\nu_k^{\text{PI}}(0) = 1$, and

$$\nu_k^{\text{PI}}(x_k) = \nu_k^{\text{PI}}(m_k - 1) + \sum_{j_k=m_k-1}^{x_k-1} (1 - P_k^{\text{miss}}(j_k)), \quad x_k \geq m_k. \quad (27)$$

4. If $\nu_k^{\text{PI}}(x_k)$ converges as $x_k \rightarrow \infty$ to a finite limit $\nu_k^{\text{PI}}(\infty)$, then

$$\nu_k^{\text{PI}}(\infty) = \begin{cases} \frac{m_k \mu_k - f_k^*}{m_k \mu_k - \lambda_k^*} = \frac{1 - \rho_k^* \bar{P}_k^{\text{miss},*}}{1 - \rho_k^*} > 1, & \text{if } 0 < \rho_k^* < 1 \\ \nu_k^{\text{PI}}(m_k - 1) + \sum_{j_k=m_k-1}^{\infty} (1 - P_k^{\text{miss}}(j_k)), & \text{if } \rho_k^* = 1. \end{cases} \quad (28)$$

Based on extensive numerical experience, we conjecture that the PI index $\nu_k^{\text{PI}}(x_k)$ for the present model is monotone increasing and converges to a finite limit. See, e.g., Figure 2. Further, L'Hôpital's rule yields that $\lim_{\lambda_k^* \nearrow m_k \mu_k} (m_k \mu_k - f_k(\lambda_k^*)) / (m_k \mu_k - \lambda_k^*) = \beta_k \triangleq f'_k((m_k \mu_k)^-)$, suggesting that the limiting index in the case $\rho_k^* = 1$ should be $\nu_k^{\text{PI}}(\infty) = \beta_k$, which the author has also verified numerically. Yet, such experiments also reveal that the computed index $\hat{\nu}_k^{\text{PI}}(x_k)$ may occasionally diverge in practice, due to small inaccuracies in the computed value of f_k^* . We propose to overcome such numerical instability by setting the computed index value equal to the appropriate limit in (28) as soon as $\hat{\nu}_k^{\text{PI}}(x_k)$ shows any signs of diverging, e.g., by exceeding the limit $\nu_k^{\text{PI}}(\infty)$ in (28).

5 RB index

The RB routing index for a queue is obtained by adapting to the present model the approach and results introduced in Niño-Mora (2002) for the design of policies for control of admission and routing to parallel queues via Whittle's RB indexation (see Whittle (1988)). See also Niño-Mora (2007). To define Whittle's RB index, we consider the single-queue case ($n = 1$) of modified problem (8), concerning the optimal dynamic control of admission to an M/M/ m queue with arrival rate λ and service rate μ at each server. Note that we drop the queue label k . Jobs come with a response-time deadline with CDF $F(t)$. At the arrival instant of a job, the




Figure 2: The PI index for a queue with $m = 6$, $\mu = 0.5$, $\tau \sim \text{Unif}[2.5, 3.5]$, and $\lambda^* = 2.1$.

controller can choose to reject it, incurring a rejection penalty of R , or accept it. In the latter case, a unit penalty is incurred if the job's deadline is missed. As discussed in Section 2.2, we will account for deadline-miss costs via continuously accruing action-independent holding costs, which are incurred at rate $h(x)$ (given by (7)) per unit time when there are x jobs in the queue.

Consider the corresponding *optimal admission control problem*, which is to find an admission control policy π^* , drawn from the class Π of policies that base decisions on the history of queue states and actions, that minimizes the average rate of rejection and deadline-miss costs per unit time. Adapting the notation in (8) to the present setting, we formulate such a problem as

$$\min_{\pi \in \Pi} \limsup_{T \rightarrow \infty} \frac{1}{T} \mathbb{E}_{i^0}^{\pi} \left[\int_0^T \left\{ R\lambda a(t) + h(X(t)) \right\} dt \right]. \quad (29)$$

By standard results, problem (29) has an optimal policy that is stationary deterministic and independent of the initial state. Now, instead of considering a fixed rejection charge R , we will view it as a *parameter* that is allowed to take any real value $R \in \mathbb{R}$. Rather than solving each individual problem (29) separately for different values of R , we will characterize optimal policies for the entire parametric family of such *R-charge problems*, as $R \in \mathbb{R}$, in terms of a numeric index attached to queue states.

Let us say that problem (29) is *indexable* if there exists a numeric index $\nu^{\text{RB}}(x)$ attached to queue states x , which does not depend on R , such that, for any rejection charge $R \in \mathbb{R}$ and state x , it is optimal to reject an arrival when the queue holds x jobs iff $\nu^{\text{RB}}(x) \geq R$.

Note that, in particular, if problem (29) is indexable and the index $\nu^{\text{RB}}(x)$ is nondecreasing in the state x , then it follows that problem (29) is solved optimally by a *threshold policy* (reject an arrival iff the queue length is larger than a certain threshold). While the conventional approach to establish the optimality of threshold policies for admission control problems such as (29) relies on assuming convexity of the holding cost rate function $h(x)$ (see Stidham (1985)), we will be able to establish the stronger result of *indexability* with a nondecreasing index $\nu^{\text{RB}}(x)$, even though the holding cost rate $h(x)$ defined in (7) is *not convex* (note that $h(x) = m\mu P^{\text{miss}}(x-1)$ for $x \geq m$, and hence $h(x)$ grows to $m\mu$ as $x \rightarrow \infty$).

To establish that problem (29) is indexable with a monotone nondecreasing RB index $\nu^{\text{RB}}(x)$, and to evaluate the index, we will deploy the general sufficient conditions for indexability and the *adaptive-greedy index algorithm* introduced in Niño-Mora (2001, 2002), which do not require such convexity assumptions. Such conditions are formulated in Niño-Mora (2002) in terms of certain quantities $w^{S_y}(x)$, for $x \geq 0, y \geq 1$, which represent *marginal rejection measures*, and indices $\nu^{\text{RB}}(x)$, which are produced by the algorithm. Here, we will write $w(x) \triangleq w^{S_{x+2}}(x)$.

In the present model, such an index algorithm reduces to a coupled set of first-order linear recursions (see

Niño-Mora (2002, pp. 396–397)). First, compute

$$\begin{aligned} z(1) &= 1; \quad z(x) = 1 - \frac{\lambda \bar{\mu}(x-1)}{[\lambda + \bar{\mu}(x-1)][\lambda + \bar{\mu}(x)]z(x-1)}, \quad x \geq 2, \\ w(0) &= \frac{\lambda \mu}{\lambda + \mu}; \quad w(x-1) = \lambda \frac{\Delta \bar{\mu}(x) + \frac{w(x-2)}{\bar{\rho}(x-2)}}{[\lambda + \bar{\mu}(x)]z(x)}, \quad x \geq 2, \end{aligned} \quad (30)$$

where $\bar{\rho}(x) \triangleq \lambda/\bar{\mu}(x+1)$, $\bar{\mu}(x) \triangleq \min(x, m)\mu$, and $\Delta \bar{\mu}(x) \triangleq \bar{\mu}(x) - \bar{\mu}(x-1)$. Then, recursively compute the index $\nu^{\text{RB}}(x)$ by

$$\nu^{\text{RB}}(0) = \frac{h(1)}{\mu}; \quad \nu^{\text{RB}}(x) = \nu^{\text{RB}}(x-1) + \frac{\Delta h(x+1) - \nu^{\text{RB}}(x-1)\Delta \bar{\mu}(x+1)}{\Delta \bar{\mu}(x+1) + \frac{w(x-1)}{\bar{\rho}(x-1)}}. \quad (31)$$

We can simplify (30)–(31) by noting that $\Delta \bar{\mu}(x) = \mu$ for $1 \leq x \leq m$, and $\Delta \bar{\mu}(x) = 0$ for $x \geq m+1$. Further, from (7) and $P^{\text{miss}}(x) = P^{\text{miss}}(0)$ for $1 \leq x < m$ (see Section 3), we have $\Delta h(x) \triangleq h(x) - h(x-1) = \mu P^{\text{miss}}(0)$ for $i = 1, \dots, m$, whereas $\Delta h(x) = m\mu \Delta P^{\text{miss}}(x-1)$ for $x \geq m+1$. Hence,

$$\nu^{\text{RB}}(x) = \begin{cases} P^{\text{miss}}(0), & 0 \leq x \leq m-1 \\ \nu^{\text{RB}}(x-1) + \lambda \frac{\Delta P^{\text{miss}}(x)}{w(x-1)}, & x \geq m. \end{cases} \quad (32)$$

Such a set of first-order linear recursions, along with those in Section 3 for $P^{\text{miss}}(x)$ (which equals $\nu^{\text{IO}}(x)$), efficiently computes the index values $\nu^{\text{RB}}(0), \dots, \nu^{\text{RB}}(x)$ in $O(x)$ time.

In the single-server case $m = 1$ such recursions are readily solved:

$$z(x) = \begin{cases} \frac{\rho^{x+1} - 1}{(\rho + 1)(\rho^x - 1)} & \text{if } \rho \neq 1 \\ \frac{x+1}{2x} & \text{if } \rho = 1 \end{cases} \quad \text{and} \quad w(x) = \begin{cases} \lambda \frac{\rho - 1}{\rho^{x+2} - 1} & \text{if } \rho \neq 1 \\ \frac{\lambda}{x+2} & \text{if } \rho = 1. \end{cases}$$

Further, assuming deterministic deadlines $\tau \equiv t$, we have $\Delta P^{\text{miss}}(x) = e^{-\mu t}(\mu t)^x/x!$. Hence, in such a case the RB index is given by $\nu^{\text{RB}}(0) = e^{-\mu t}$, and, for $x \geq 1$,

$$\nu^{\text{RB}}(x) = \begin{cases} \frac{e^{-\mu t}}{\rho - 1} \sum_{j=0}^x \frac{\rho(\lambda t)^j - (\mu t)^j}{j!} & \text{if } \rho \neq 1 \\ e^{-\mu t} \sum_{j=0}^x \frac{(j+1)(\mu t)^j}{j!} & \text{if } \rho = 1. \end{cases} \quad (33)$$

The sufficient conditions for indexability referred to above, as they apply to the present model, are:

- (i) The marginal rejection measures $w^{S_y}(x)$ are positive; and
- (ii) the index $\nu^{\text{RB}}(x)$ is monotone nondecreasing in x .

The following result is established in Niño-Mora (2001, Cor. 2) and Niño-Mora (2002, Th. 6.3), at increasing levels of generality.

Theorem 1. *Under conditions (i)–(ii) the model is indexable, with index $\nu^{\text{RB}}(x)$.*

We are now ready to establish that the present model is indexable.

Proposition 5. *Problem (29) is indexable, and $\nu^{\text{RB}}(x)$ is its RB index.*

Proof. Regarding condition (i), its satisfaction is ensured by Niño-Mora (2002, Prop. 7.2) when service rates $\bar{\mu}(x)$ are concave nondecreasing in x , as is the case here. As for condition (ii), it follows from (32), noting that $P^{\text{miss}}(x)$ is nondecreasing and $w(x-1) > 0$. Invoking Theorem 1 completes the proof. \square

It must be emphasized that, in the system model of concern in this paper, with n parallel queues, the above RB index is to be computed separately for each of the queues, using for queue k in the above formulae the overall arrival rate λ and the parameters m_k and μ_k . Hence, unless the parallel-queue system is lightly loaded, one will typically have $\rho_k \triangleq \lambda/(m_k\mu_k) > 1$ for some or even all queues k , even while the system load $\lambda/\sum_k m_k\mu_k$ remains below unity. In such a typical case, the computation and use of the RB index raises the following issues:

1. It is readily verified that, for a queue with $\lambda > m\mu$, i.e., $\rho > 1$, $\lim_{x \rightarrow \infty} w(x) = 0$. Since $\lim_{x \rightarrow \infty} \Delta P^{\text{miss}}(x) = 0$, computation of the RB index via (32) can be numerically unstable.
2. Numerical experiments reveal that the RB index $\nu^{\text{RB}}(x)$ grows extremely fast with λ , converging to a finite limit $\nu^{\text{RB}}(\infty)$ as $x \rightarrow \infty$, which can be very large for a queue with utilization $\rho > 1$. The cause of this phenomenon is apparent in the single-server case $m = 1$ with a deterministic deadline $\tau \equiv t$, in which (33) yields that the index $\nu^{\text{RB}}(x)$ converges as $x \rightarrow \infty$ to the limit

$$\nu^{\text{RB}}(\infty) = \begin{cases} \frac{\rho e^{(\lambda-\mu)t} - 1}{\rho - 1} & \text{if } \rho \neq 1 \\ 1 + \mu t & \text{if } \rho = 1. \end{cases} \quad (34)$$

Hence, $\nu^{\text{RB}}(\infty)$ grows exponentially with λ . Thus, e.g., for a queue with $m = 1$, $\mu = 1$, $\tau \equiv 2$, and $\lambda = 5$, $\nu^{\text{RB}}(x)$ converges as $x \rightarrow \infty$ to $\nu^{\text{RB}}(\infty) = 3725.95$, as shown in Figure 3. As a consequence, when used to design an admission policy (reject an arrival if the rejection charge R is lower than the current index value at every queue), the resulting admission policy may grossly overreject jobs relative to what is optimal.

3. The convergence of the index $\nu^{\text{RB}}(x)$ to the finite limit $\nu^{\text{RB}}(\infty)$ implies that, in the single-queue admission control problem (29), it is optimal to admit *all* arrivals iff the rejection cost per job R is at or above $\nu^{\text{RB}}(\infty)$.




Figure 3: The RB index for a queue with $m = 1$, $\mu = 1$, $\tau \equiv 2$, and $\lambda = 5$.

6 Numerical experiments

This section reports on the results of a numerical study, whose aim is to benchmark the cost performance of the policies considered above (the static BS policy, and the dynamic IO, PI, and RB index policies), both against each other and against the optimal cost performance, in a variety of scenarios.

To accurately evaluate the optimal cost performance, attention was restricted to instances with $n = 2$ queues (clusters). The optimal average cost per job was evaluated by solving with CPLEX the linear programming

formulation of the average-cost DP equations (9), truncating the buffer size of each queue to 60 jobs (it was verified that increasing buffer sizes gave the same results for the instances considered). As for the performance of each of the dynamic index policies, it was evaluated by solving with MATLAB the corresponding Poisson equations.

The study encompasses both the model of optimal joint control of admission and routing, and the pure-routing model where the admission control capability is disabled.

The experiments were designed to assess the effect, on the performance of each policy, of varying the system load, the rejection cost, the degree of heterogeneity in either cluster server speeds or server-pool sizes, and the magnitude and distribution of the response-time deadline. Recall from Section 1 that the deadline-miss cost is normalized to the value $C = 1$. Further, in light of the alternative interpretation given in Section 1 of R as the admission fee per job and C as the deadline-miss refund, we only consider instances where $R \geq C$, i.e., where the refund per missed deadline is a fraction of the admission fee, as this appears more reasonable in practice than the case $R < C$. Note that, in such cases, the IO index policy does not reject any jobs, since its value is below 1.

6.1 Effect of system load

We consider a base instance having $n = 2$ queues, with server-pool sizes $\mathbf{m} = (4, 8)$ and service rates $\boldsymbol{\mu} = (5, 3)$. The response-time deadline is deterministic, being $\tau \equiv 1$. Figure 4 plots, for each of four values of the rejection cost $R \in \{1, 5, 20, \infty\}$, the average cost per job under each policy, and under an optimal policy, against the system load $\rho \triangleq \lambda / \sum_k m_k \mu_k$, which ranges over $\rho \in [0.1, 0.95]$ by changing the arrival rate λ .

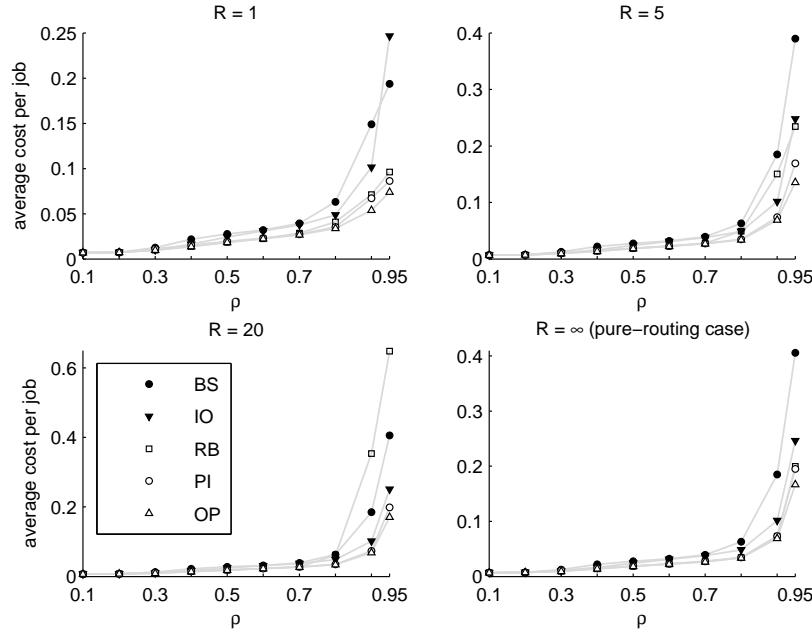


Figure 4: Effect of system load with $\tau \equiv 1$.

The results are displayed in Figure 4, which shows that, while all the policies considered are nearly optimal under light loads, they exhibit significant differences in their cost performances under heavier loads. For the model with admission control, the first three panes show that the PI policy consistently achieves a near-optimal performance. The performance of the RB policy steeply deteriorates as the load gets heavier in cases with high rejection costs, becoming the worst policy for $\rho = 0.9$ and $R = 20$. Recall that, as noted above, the RB index can reach extremely high values under heavy loads, and hence yields admission policies that overreject. As for the IO policy, its optimality loss worsens under as the load gets heavier in cases with low rejection costs (in

contrast to the RB policy). This is to be expected, since in such cases it will be optimal to reject some jobs, whereas the IO policy does not reject any job. The optimal BS policy steadily worsens as the load gets higher.

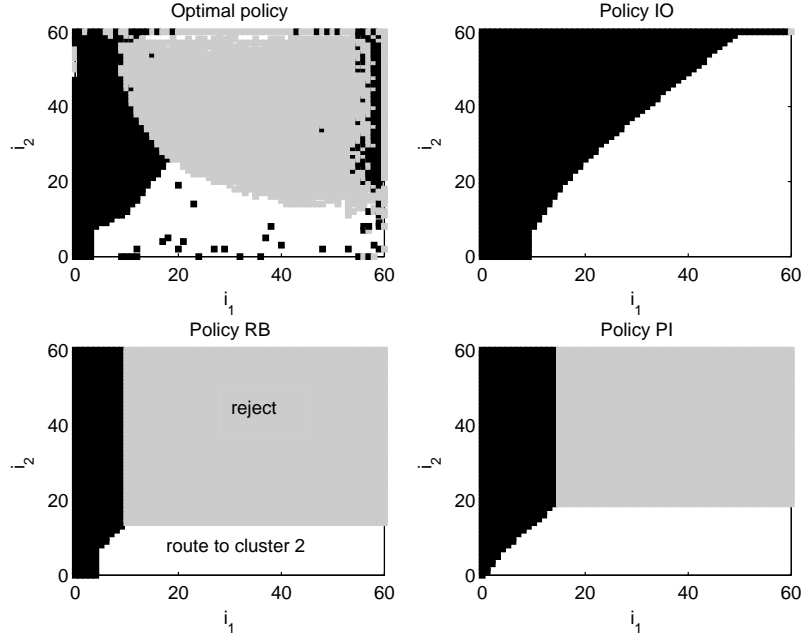


Figure 5: Structure of computed policies for the case $\tau \equiv 1$, $\rho = 0.9$, $R = 4$.

Figure 5 plots the structure of the index policies and the optimal policy that was computed in one of the instances. It is apparent that the RB policy rejects too many jobs. Although the shape of the rejection region differs under the PI policy and under an optimal policy, the former's performance is near optimal. Note that, in the plot of the optimal policy, the actions at some joint states (x_1, x_2) differ from the overall pattern (e.g., there are some black dots inside the white region). This is due to numerical effects, as the linear programs from which such a solution results are rather large. In fact, we have found that in most instances the structure of the optimal policy is obscured due to such effects.

As for the pure-routing model, the lower right pane in Figure 4 shows that both the PI and the RB policies are nearly optimal for the range of loads considered. The IO policy deviates more from the optimal performance in heavy traffic. The performance of the BS policy worsens substantially as traffic gets heavier.

Figures 6 and 7 show the results of carrying out this experiment for two nondeterministic cases of the response-time deadline distribution: $\tau \sim \text{Exp}(1)$ and $\tau \sim \text{Uniform}[0.3, 1.7]$. Note that, in each case, the mean response-time deadline is 1. The results are qualitatively similar to those for a deterministic response-time deadline $\tau \equiv 1$. Yet, notice that the level of average costs increase with the variability of the response-time deadline distribution, being highest in the exponential case, and lowest in the deterministic case.

6.2 Effect of rejection cost

The second experiment aims to benchmark the performance of the policies considered as the rejection cost R varies. We consider the same base instance as in the first experiment. Figure 8 plots, for each of four values of the system load $\rho \in \{0.7, 0.8, 0.9, 0.95\}$, the average cost per job obtained under each policy, and under an optimal policy, against the rejection cost $R \in \{1, 2, \dots, 20\}$.

The results show that the performance of all policies considered, except RB, is not substantially affected by the rejection cost R . Such is also the case for the RB policy when the system load is low, but, as the load gets higher, its performance severely degrades as R grows. This is due to the fact that, for high loads, the RB index takes very large values, and hence the resulting admission policy rejects too many jobs.

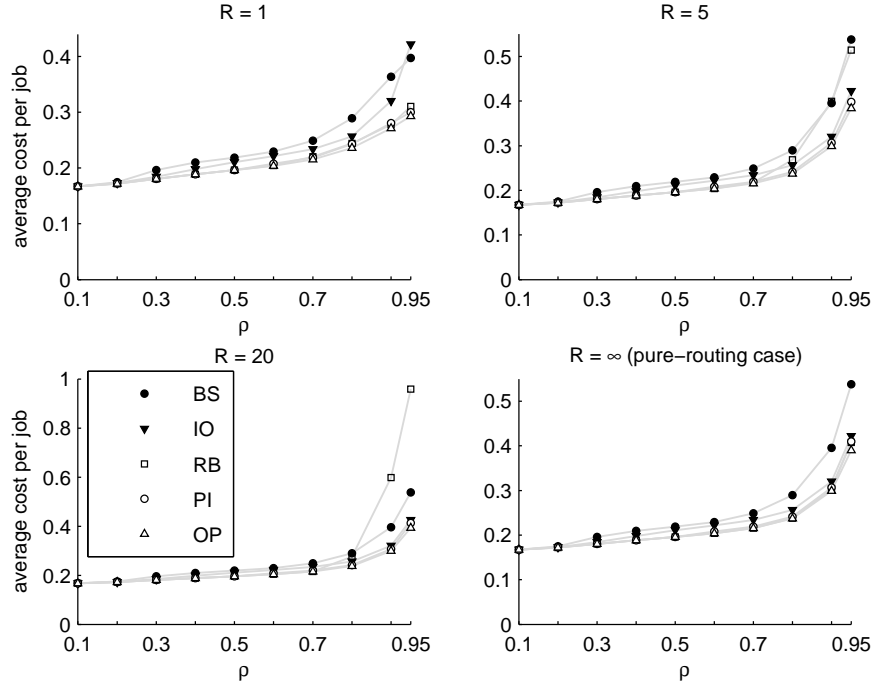


Figure 6: Effect of system load with $\tau \sim \text{Exp}(1)$.

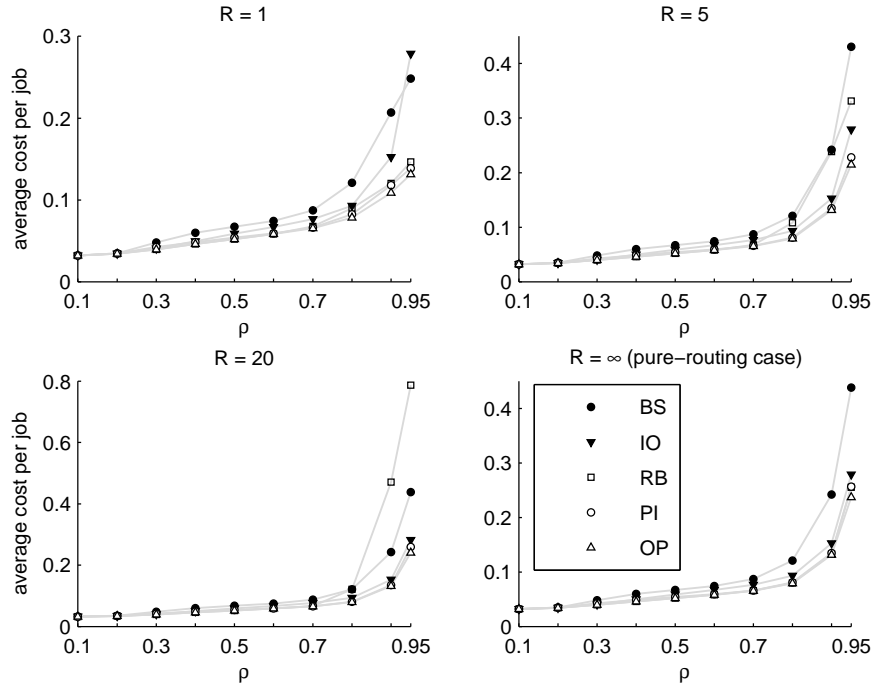


Figure 7: Effect of system load with $\tau \sim \text{Uniform}[0.3, 1.7]$.

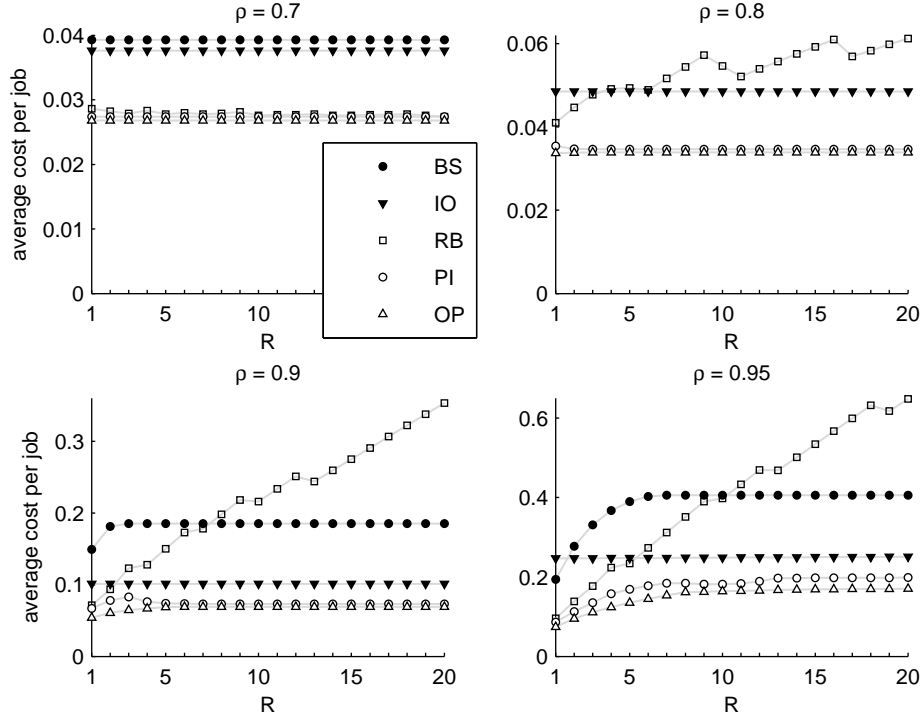


Figure 8: Effect of rejection cost.

6.3 Effect of heterogeneity in cluster server speeds

The third experiment aims to compare the policies as server speeds at the two queues range from a highly heterogeneous system to a homogeneous system. We consider a base instance with homogeneous server-pool sizes $\mathbf{m} = (4, 4)$, system load $\rho = 0.9$, and response-time deadline $\tau \equiv 1$. Figure 9 plots, for each of four values of the rejection cost $R \in \{1, 5, 20, \infty\}$, the average cost per job under each policy, and under an optimal policy, against five combinations of service rates: $\boldsymbol{\mu} \in \{(9, 1), (8, 2), (7, 3), (6, 4), (5, 5)\}$. Note that the average speed per server is the same in each combination.

For the model with admission control, the first three panes in Figure 9 show that the cost performances of both the BS and the RB policies deteriorate as cluster server speeds get more homogeneous, with BS (resp. RB) being the worse of the two for lower (resp. higher) rejection costs. As for the IO policy, its optimality loss displays the opposite behavior, becoming closer to optimal as cluster server speeds get more homogeneous. Note that IO is the worst policy for $R = 1$. The PI policy is, again, consistently near optimal.

As for the pure-routing model, the lower right pane in Figure 9 shows that the optimality loss of the BS policy deteriorates substantially as the queues become more homogeneous. Policies RB and IO display the opposite behavior, becoming closer to optimal as the clusters get more homogeneous. The best of the index policies is PI, followed by RB and then IO.

6.4 Effect of heterogeneity in cluster server-pool sizes

The aim of the fourth experiment is to benchmark the performance of the policies as server-pool sizes at the two queues range from a highly heterogeneous system to a homogeneous system. We consider a base instance with homogeneous server speeds $\boldsymbol{\mu} = (5, 5)$, system load $\rho = 0.9$, and response-time deadline $\tau \equiv 1$. Figure 10 plots, for each of four values of the rejection cost $R \in \{1, 5, 20, \infty\}$, the average cost per job under each policy, and under an optimal policy, against five combinations of server-pool sizes: $\mathbf{m} \in \{(9, 1), (8, 2), (7, 3), (6, 4), (5, 5)\}$.

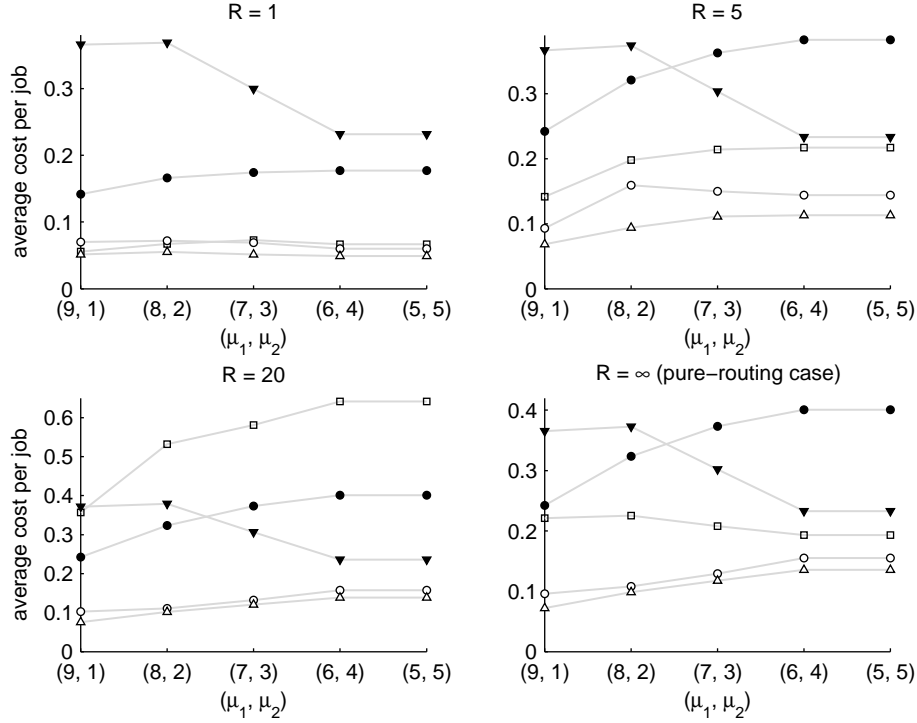


Figure 9: Effect of heterogeneity in cluster server speeds.

Note that the total number of servers is the same in each combination. The results of this experiment are qualitatively similar to those of the previous experiment in terms of ranking of the policies. Again, PI is the only policy that is consistently near optimal.

6.5 Effect of magnitude of response-time

The fifth experiment benchmarks the performance of the policies as the response-time deadline, which is taken to be deterministic with $\tau \equiv t$, ranges from a low to a high value. We consider a base instance with server-pool sizes $\mathbf{m} = (4, 8)$, service rates $\boldsymbol{\mu} = (5, 3)$, and system load $\rho = 0.9$. Figure 11 plots, for each of four values of the rejection cost $R \in \{1, 5, 20, \infty\}$, the average cost per job incurred under each policy, and under an optimal policy, against the response-time deadline $t = 1/\theta$, with $\theta \in \{0.1, 0.5, 1, 1.5, 2, 2.5, 3, 3.5, 4, 4.5\}$.

For the model with admission control, the first three panes in Figure 11 show that, for the model with admission control, the PI policy is consistently near optimal, with the other policies showing substantial optimality losses in some cases. Thus, the worst policy is IO for low R , BS for medium R , and RB for large R . For the pure-routing case, the fourth pane in the figure shows that PI and RB are both nearly optimal.

7 Discussion and conclusions

This paper has investigated the use of four policy design methods to obtain policies for control of admission and routing in a distributed soft real-time system with a parallel multicluster architecture: the static BS policy, and the dynamic IO, PI, and RB index policies. The above analyses and numerical results reveal the following facts and insights:

1. All the heuristic index policies considered can be computed efficiently in linear time via first-order linear recursions.

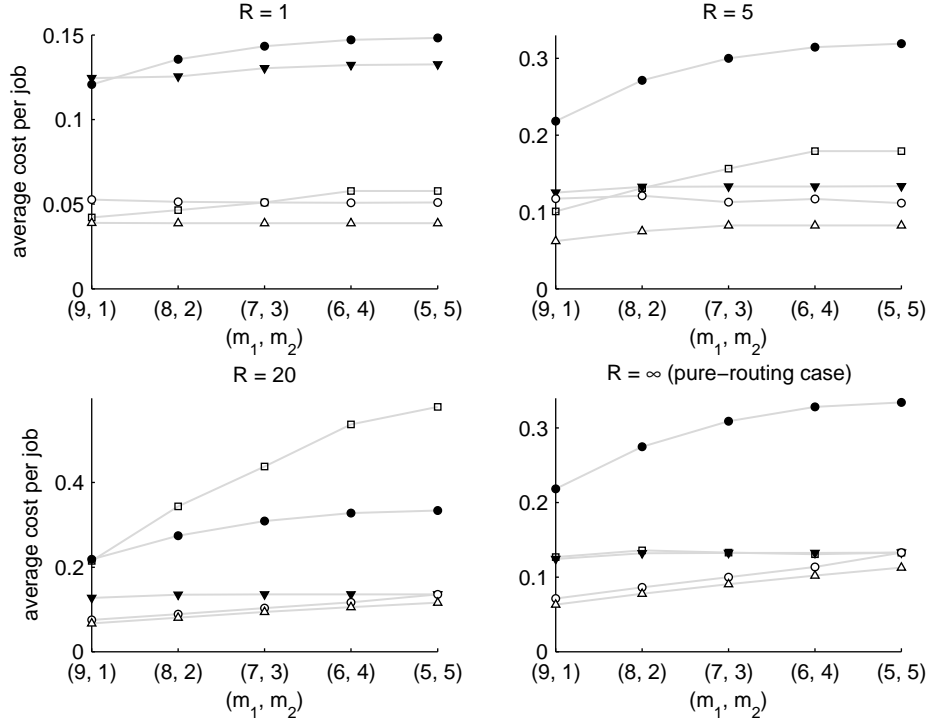


Figure 10: Effect of heterogeneity in cluster server-pool sizes.

2. The admission policy based on the IO index does not reject any job when $R \geq C$, even though doing so can substantially improve performance.
3. The admission policy based on the RB index tends to grossly overreject jobs when both the system load and the rejection cost are high, since under heavy loads the RB index can take extremely high values. This causes a severe performance degradation of the RB policy in such cases. Yet, in the pure-routing case, the RB routing index yields nearly optimal policies.
4. The optimal BS policy can be efficiently computed in the present model. Extending the range of rejection cost values R beyond the case $R = 1$ considered in prior work to $R \geq 1$ leads to policies that may saturate one or more queues when the load is high.
5. The optimal BS policy can be substantially outperformed by good dynamic policies.
6. Overall, the best policy among those considered for admission control and routing is the PI policy, which consistently yields nearly optimal policies. As for the pure-routing model, both the PI and RB index policies are consistently near optimal.

An interesting issue for future research would be to analyze the performance of the proposed PI index policy, or of another heuristic policy that performed as well, elucidating its asymptotic performance in instances with many servers, as the analysis of many-server models has received substantial attention in related literature (see [Gurvich and Whitt \(2010\)](#) and references therein).

Acknowledgments

The author acknowledges partial support for this work by the Spanish Ministry of Science and Innovation's projects MTM2007-63140 and MTM2010-20808. He presented a preliminary outline of this work at the Twelfth

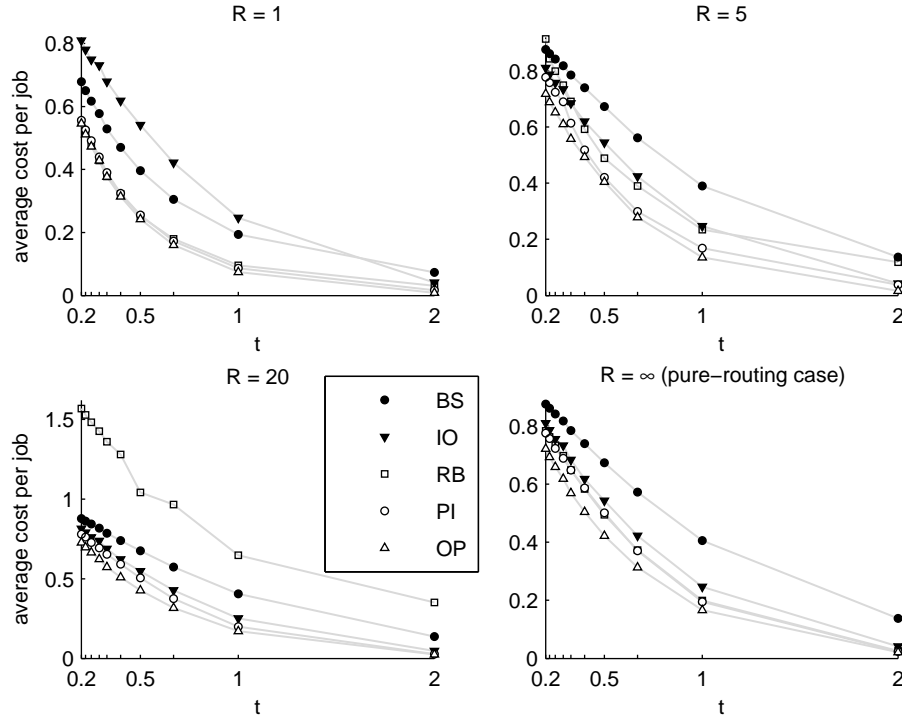


Figure 11: Effect of magnitude of response-time deadline (deterministic case with $\tau \equiv t$).

Workshop on MATHematical performance Modeling and Analysis (MAMA 2010), in New York, which appeared in the corresponding proceedings (Niño Mora (2010)).

References

- Bertini, L., Leite, J. C. B., Mossé, D., 2010. Power and performance control of soft real-time web server clusters. *Information Processing Letters* 110, 767–773.
- Bhulai, S., 2009. Dynamic routing policies for multiskill call centers. *Probability in the Engineering and Information Sciences* 23, 101–119.
- Gross, D., Shortle, J. F., Thompson, J. M., Harris, C. M., 2008. *Fundamentals of Queueing Theory*, 4th Edition. Wiley, Hoboken, NJ.
- Gurvich, I., Whitt, W., 2010. Service-level differentiation in many-server service systems via queue-ratio routing. *Operations Research* 58, 316–328.
- He, L., Jarvis, S. A., Spooner, D. P., Jiang, H., Dillenberger, D. N., Nudd, G. R., 2006. Allocating non-real-time and soft real-time jobs in multiclusters. *IEEE Transactions on Parallel and Distributed Systems* 17, 99–112.
- Kallmes, M. H., Cassandras, C. G., 1995. Two approaches to optimal routing and admission control in systems with real-time traffic. *Journal of Optimization Theory and Applications* 84, 311–338.
- Kang, K. D., Sin, P. H., Oh, J., H., S., Son, S. H., 2007. A real-time database testbed and performance evaluation. In: *Proceedings of the 13th IEEE International Conference on Embedded and Real-Time Computing Systems and Applications*, pp. 319–326.

- Kao, B., García-Molina, H., 1996. Scheduling soft real-time jobs over dual non-real-time servers. *IEEE Transactions on Parallel and Distributed Systems* 7, 56–68.
- Krishnan, K. R., 1990. Joining the right queue: A state-dependent decision rule. *IEEE Transactions on Automatic Control* 35, 104–108.
- Niño-Mora, J., 2001. Restless bandits, partial conservation laws and indexability. *Advances in Applied Probability* 33, 76–98.
- Niño-Mora, J., 2002. Dynamic allocation indices for restless projects and queueing admission control: A polyhedral approach. *Mathematical Programming* 93, 361–413.
- Niño-Mora, J., 2007. Marginal productivity index policies for admission control and routing to parallel multi-server loss queues with reneging. In: *Proceedings of the first EuroFGI International Conference on Network Control and Optimization (NETCOOP '07)*. Lecture Notes in Computer Science, vol. 4465. Springer, Berlin, pp. 138–149.
- Niño Mora, J., 2010. Index policies for admission and routing of soft real-time traffic to parallel queues. *ACM SIGMETRICS Performance Evaluation Review* 38, 21–23.
- Plankensteiner, K., Prodan, R., Fahringer, T., 2010. Scheduling scientific workflows to meet soft deadlines in the absence of failure models. In: D’Ambra, P., Guarracino, M., Talia, D., editors. *Euro-Par 2010, Part I*. Lecture Notes in Computer Science, vol. 6271. Springer, Berlin, pp. 367–378.
- Stidham, Jr., S., 1985. Optimal control of admission to a queueing system. *IEEE Transactions on Automatic Control* 30, 705–713.
- Whittle, P., 1988. Restless bandits: Activity allocation in a changing world. *Journal of Applied Probability* 25A, 287–298.
- Zhu, W., 2001. Allocating soft real-time tasks on cluster. *Simulation* 77, 219–229.