

Marginal Productivity Index Policies for Admission Control and Routing to Parallel Multi-server Loss Queues with Reneging

José Niño-Mora*

Universidad Carlos III de Madrid
Department of Statistics
C/ Madrid 126
E-28903 Getafe (Madrid), Spain
jnimora@alum.mit.edu
<http://alum.mit.edu/www/jnimora>

Abstract. This paper addresses the problem of designing tractable dynamic admission control and/or routing policies in a Markovian model of parallel multi-server loss queues with reneging, which seek to optimize performance objectives of concern. Such problems are relevant in a variety of systems that provide distributed telecommunication or computing services. The paper shows the direct applicability of the author's results in Niño-Mora (2002) [Dynamic allocation indices for restless projects and queueing admission control: a polyhedral approach. *Math. Program.* **93** 361–413], where index policies based on restless bandits were developed in a broader setting. A well-grounded and tractable index policy for admission control and/or routing is thus proposed for the model of concern. Results of preliminary computational experiments are reported, showing that the proposed index policy is nearly optimal, and substantially outperforms conventional benchmark policies in the instances investigated.

Keywords: dynamic routing; admission control; index policies; parallel queues; multi-server; loss queues.

1 Introduction

This paper addresses the problem of designing tractable dynamic admission control and/or routing policies in a Markovian model of parallel multi-server loss queues with reneging, which seek to optimize performance objectives of concern. A single customer class arrives to the system as a Poisson stream with rate λ .

* The author's work has been supported in part by the Spanish Ministry of Education & Science under grants BEC2000-1027, MTM2004-02334 and a Ramón y Cajal Investigator Award, by CAM-UC3M's grant UC3M-MTM-05-075, and by the European Union's Network of Excellence Euro-NGI. He has presented earlier versions of this work at the Schloss Dagstuhl Seminar on Scheduling in Computer and Manufacturing Systems (Wadern, Germany, 2002), and at the Workshop on Topics in Computer Communication and Networks (Cambridge Univ., UK, 2002).

Upon a customer's arrival, the controller decides (i) whether to admit the customer into the system or to reject it; and, if admitted, (ii) to which of K queues in parallel to route the customer for service. Queue $k \in \mathbb{K} \triangleq \{1, \dots, K\}$, which serves customers in FCFS order, is endowed with a pool of m_k identical parallel exponential servers, each working at rate μ_k , and has a finite buffer of size $n_k \geq m_k$. We will denote by $X_k(t)$ the state of queue k at time t , given by the number of customers it holds waiting or in service, and by $a_k(t) \in \{0, 1\}$ the action indicator taking value 1 when a customer arriving at time t would *not* be routed to queue k .

Once in a queue, customers that have not started service become impatient, being prone to renege from the system. As introduced in Palm (1957), we assume that the reneging-time distribution is exponentially distributed with rate θ_k at queue k , and that interarrival, service and reneging times are mutually independent.

The following types of cost and reward accrue: (1) *holding costs*, at the convex nondecreasing rate $h_k(j_k)$ per unit time that j_k customers are in queue k ; (2) *reneging costs*, at rate c_k per customer that reneges from queue k ; (3) *completion rewards*, at rate r_k per service completed at queue k ; and (4) *loss costs*, at rate ν per rejected customer.

We will find it convenient to write the equivalent total net cost rate per unit time when the joint system state is $\mathbf{j} = (j_k)_{k \in \mathbb{K}}$ and joint action $\mathbf{a} = (a_k)_{k \in \mathbb{K}}$ prevails as

$$-(K - 1)\lambda\nu + \sum_{k \in \mathbb{K}} \{C_k(j_k) + \nu Q_k^{a_k}(j_k)\}, \tag{1}$$

where

$$C_k(j_k) \triangleq h_j(j_k) + (j_k - m_k)^+ c_k \theta_k - \min(j_k, m_k) r_k \mu_k$$

and

$$Q_k^{a_k} \triangleq \lambda a_k.$$

We will consider two versions of the problem: (i) the case where the admission control capability is enabled; and (ii) the case where it is not, in which rejections occur only when all buffers are full. The system is operated by adopting an admission control and routing policy (for version (i)), or just a routing policy (for version (ii)), denoted by π , which is drawn from the corresponding class Π of history-dependent randomized policies.

The operation of such a system raises the following optimization problems: (i) find a policy minimizing the expected total discounted value of net costs accrued,

$$\min_{\pi \in \Pi} \mathbb{E}_i^\pi \left[\int_0^\infty e^{-\alpha t} \sum_{k \in \mathbb{K}} \{C_k(X_k(t)) + \nu Q_k^{a_k(t)}\} dt \right], \tag{2}$$

where $\alpha > 0$ is the discount rate and $\mathbf{i} = (i_k)_{k \in \mathbb{K}}$ is the initial joint state; and (ii) find a policy minimizing the long-run average net cost rate per unit time,

$$\min_{\pi \in \Pi} \limsup_{T \rightarrow \infty} \frac{1}{T} \mathbb{E}_i^\pi \left[\int_0^T \sum_{k \in \mathbb{K}} \{C_k(X_k(t)) + \nu Q_k^{a_k(t)}\} dt \right]. \tag{3}$$

Notice that in (2)–(3) we have disregarded the additive constant $-(K-1)\lambda\nu$ in (1).

Such problems are relevant in a variety of applications, most notably in the provision of geographically distributed telecommunication or computing services. Thus, e.g., in a distributed call center, calls may be initially accepted or rejected, depending on current congestion levels. If accepted, they are routed to one of multiple operator pools. The present model assumes that operators within a pool are homogeneous, whereas their skills might differ across pools. It further captures the fact that customers are prone to become impatient, possibly abandoning the system before receiving service. A similar situation arises in the operation of a distributed grid of computing nodes, consisting of workstation clusters, to which jobs are to be dynamically routed. For earlier related work see, e.g., Houck (1987), Bassamboo et al. (2005), and references therein.

While the above problems are readily formulated as discrete-time finite Markov decision processes (MDPs), their solution via the conventional dynamic programming approach is computationally intractable in large-scale models, due to the exponential growth of the state space's size on the number of queues. We will thus focus attention on the goals of designing, computing and testing well-grounded heuristic policies that are readily implementable.

Since in such problems the controller must dynamically assess the relative values of alternative rejection and routing actions, it is both intuitively appealing and practical to do so based on an *index policy*. In the present model, such policies are based on attaching to each queue k an *index* $\nu_k(j_k)$, which can be thought of as a *measure of undesirability for routing a customer to queue k* , given as a function of its current state j_k . Then, in the problem version with admission control capability, an arriving customer would be admitted if $\nu > \nu_k(j_k)$ for at least one queue k with $j_k < n_k$, i.e., if the cost of rejecting the customer exceeds the undesirability of routing it to some nonfull queue; otherwise, the customer would be rejected. If accepted, the customer would then be routed to a queue with *smallest* current index value, among nonfull queues k for which $\nu > \nu_k(j_k)$.

Indeed, for problem version (ii), the classic *Shortest Queue Routing* and *Shortest Expected Delay Routing* rules are examples of such index policies, with the former known to be optimal in special symmetric cases. See, e.g., Winston (1977), Johri (1989), and Hordijk and Koole (1990). We are thus led to address the issue of how to define appropriate indices $\nu_k(j_k)$ for the above problems.

Such an issue was actually resolved by the author in a broader setting in Niño-Mora (2002a). That paper, which drew on earlier work in Whittle (1988) and in Niño-Mora (2001), introduced the idea that problems of dynamic admission control and/or routing to parallel queues can be formulated as *restless bandit problems* (RBPs). It also established the existence of a corresponding *marginal productivity index* (MPI) for the constituent *bandits* (which correspond to a single queue subject to admission control), under rather general birth-death dynamics and nonlinear cost rate functions. Further, it furnished efficient index-computing algorithms, as well as closed index formulae for some special cases. See also the presentation in Niño-Mora (2002b).

However, the direct applicability of such work to models such as those addressed in this paper appears to have been overlooked. This paper thus sets out to clarify how such results bear on the present model. While the results are presented here in abridged form, a full version of this paper with complete analyses and extensive experimental results is currently under preparation. For related work on the theory and applications of restless bandit indexation, see also, e.g., Niño-Mora (2006a, 2006b).

The remainder of the paper is organized as follows. Section 2 discusses the reformulation of the above problems in the framework of the RBP, which allows us to deploy the corresponding MPI policy. Section 3 reviews the indexability analysis for appropriate single-queue admission control subproblems, along with index-computing algorithms, which are exploited to obtain closed-form index formulae in some special cases. Finally, Section 4 reports on the results of some preliminary computational experiments on the performance of the proposed MPI policy. These show that the proposed index policy is nearly optimal, and substantially outperforms conventional benchmark policies in the instances investigated.

2 RBP Reformulation and MPI Policy

The formulations given in (2)–(3) have been chosen to make it apparent that such problems fit into the framework of the RBP introduced by Whittle (1988). The RBP concerns the optimal dynamic allocation of effort to a collection of stochastic *projects*, modelled as binary-action (1/work/active; 0/rest/passive) MDPs. Whittle considered the problem version where a fixed number M of projects are to be engaged at each time. He showed that, for a limited class of restless projects, which he termed *indexable*, there exists a state-dependent index that characterizes their optimal policies, and proposed to use the resulting index policy for the multi-project RBP: engage at each time M projects with currently largest index values.

In the present model, as introduced in Niño-Mora (2002a), we identify a “project” with a single queue fed with a Poisson arrival stream, subject to admission control. It is convenient to imagine that such a control action is exercised by a *gatekeeper* who, when active, blocks access to the queue by shutting an entry gate, and, when passive, allows customers to enter the queue by letting the gate open. See Fig. 1. Notice that, when the queue’s buffer is full, both actions yield identical dynamics. We will thus term such a full state *uncontrollable*, and the remaining states, where there is an effective choice of action, *controllable*. Further, for consistency with the interpretation of action a_k used in (1), where $a_k = 1$ means *not* routing to queue k , we will adopt the convention that the active action is taken at uncontrollable states, i.e., $a_k = 1$ if queue k is full.

Now, to view the present model as a multi-project RBP we represent it as a collection of queues as above, *each fed by its own arrival stream with rate λ* and having its own gatekeeper, where the actions of gatekeepers are coordinated in such a way as to yield the required dynamics. Thus, in the problem version with

admission control capability, *at least* $M = K - 1$ projects are to be engaged at each time, so that at least $K - 1$ queues' entry gates are to be shut: if $K - 1$ gates are shut, an arriving customer would be routed to the queue whose gate is open, which must be nonfull; if the K gates are shut, the customer would be rejected. In the problem version without admission control capability, exactly exactly $K - 1$ queues' entry gates are shut at any time, and an arriving customer would be routed as in the previous case, provided there is at least one nonfull queue; if all queues are full, the K entry gates must be shut.

Now, to deploy restless bandit indexation we *decouple* problems (2)–(3) into the *single-project subproblems*

$$\min_{\pi_k \in \Pi_k} \mathbb{E}_{i_k}^{\pi_k} \left[\int_0^\infty e^{-\alpha t} \left\{ C_k(X_k(t)) + \nu Q_k^{a_k}(t) \right\} dt \right] \tag{4}$$

and

$$\min_{\pi_k \in \Pi_k} \limsup_{T \rightarrow \infty} \frac{1}{T} \mathbb{E}_{i_k}^{\pi_k} \left[\int_0^T \left\{ C_k(X_k(t)) + \nu Q_k^{a_k}(t) \right\} dt \right], \tag{5}$$

where i_k is the initial state of queue k , and Π_k is the class of admissible admission control policies for operating the queue *in isolation*.

We now use the rejection cost rate ν as a parameter, and consider how the optimal policies for such subproblems vary as ν ranges over \mathbb{R} . We will say that any of the above queue k 's subproblems is *indexable* if there exists an *index* $\nu_k^*(j_k)$, defined for controllable states $0 \leq j_k < n_k$, such that, for any initial state i_k , it is optimal to take the active action, i.e., reject an arrival when the queue occupies state j_k , iff $\nu_k^*(j_k) \geq \nu$. In such case, we term $\nu_k^*(j_k)$ the queue's *marginal productivity index* (MPI). Such a term was introduced in Niño-Mora 2006a, motivated by the interpretation of the index given there as the state-dependent marginal cost vs. work trade-off rate.

Hence, if it were established that such an MPI exists for the above single-queue subproblems, we could use it to implement an index policy as discussed above.

3 Single-Queue Subproblems: Indexability and MPI Computation

In fact, the required indexability analysis was carried out in Niño-Mora (2002a), in a broader setting of admission control problems with birth-death dynamics and nonlinear cost rates, which includes the above subproblems as special cases. We review here such results, which further include an efficient index-computing algorithm, and highlight their application to the present model.

3.1 Review of Relevant Results in Niño-Mora (2002a)

Consider the system portrayed in Fig. 1, which represents a single-server facility catering to an incoming customer stream, which is endowed with a finite buffer capable of holding up to and including n customers, waiting or in service. Customer flow is regulated by a *gatekeeper*, who dynamically opens or shuts an *entry*




Fig. 1. Control of admission to a single queue

gate which customers must cross to enter the buffer; those finding a shut gate, or a full buffer, on arrival are rejected and lost.

The state $X(t)$, recording the number in system at times $t \geq 0$, evolves as a controlled *birth-death process* with state space $N = \{0, \dots, n\}$. While the system occupies state j , customers arrive at rate $\lambda(j)$ (being then admitted or rejected), and the server works at rate $\mu(j)$.

The system is governed by an *admission control policy* π , prescribing the action $a(t) \in \{0, 1\}$ to take at each time t , where action 1 means shutting the entry gate. As before, we assume that such an action must be taken when the queue is full.

The system continuously accrues holding costs, at rate $C(j)$ while in state j , along with rejection charges, at rate ν per customer rejected. The equivalent rejection cost rate under action a is therefore $\nu Q^a(j)$, where $Q^a(j) \triangleq \lambda(j)a$ is the rejection cost rate.

In such a setting, we consider the following optimization problems: (i) find a discount-optimal admission control policy,

$$\min_{\pi \in \Pi} \mathbb{E}_i^\pi \left[\int_0^\infty e^{-\alpha t} \{C(X(t)) + \nu Q^{a(t)}(X(t))\} dt \right], \tag{6}$$

where $\alpha > 0$ is the discount rate and i is the initial state; and (ii) find an average-optimal policy,

$$\min_{\pi \in \Pi} \limsup_{T \rightarrow \infty} \frac{1}{T} \mathbb{E}_i^\pi \left[\int_0^T \{C(X(t)) + \nu Q^{a(t)}(X(t))\} dt \right]. \tag{7}$$

We will assume that the model parameters satisfy the following regularity conditions, where we use the notation $\Delta x(j) \triangleq x(j) - x(j - 1)$, and write

$$d(j) = \mu(j) - \lambda(j),$$

and

$$\rho(j) = \frac{\lambda(j)}{\mu(j + 1)}.$$

Assumption 1. *The following conditions hold:*

- (i) *Concave nondecreasing $d(j)$: $0 \leq \Delta d(j + 1) \leq \Delta d(j)$, $1 \leq j \leq n - 1$, with $\Delta d(1) > 0$.*
- (ii) *Convex nondecreasing $C(j)$: $\Delta C(j + 1) \geq \Delta C(j) \geq 0$, $1 \leq j \leq n - 1$.*

Now, in Niño-Mora (2002a[Th. 7.2 and Cor. 7.1]) it is proven that, under Ass. 1, such problems are indexable, so that their optimal policies are characterized by an MPI $\nu^*(j)$. Further, the latter is consistent with the intuitive class of *threshold policies*, which prescribe to reject customers if the queue is long enough. Namely, the MPI is monotone nondecreasing in the queue length j :

$$\nu^*(0) \leq \nu^*(1) \leq \dots \leq \nu^*(n - 1).$$

That paper further gives a recursive index algorithm, which we include here for ease of reference:

$$\begin{aligned} \nu^*(0) &= \frac{\Delta C(1)}{\alpha + \Delta d(1)} \\ \nu^*(j) &= \nu^*(j - 1) + \frac{\Delta C(j + 1) - \nu^*(j - 1) (\alpha + \Delta d(j + 1))}{\alpha + \Delta d(j + 1) + \frac{w^{S(j+1)}(j - 1)}{\rho(j - 1)}}, \quad 1 \leq j \leq n - 1, \end{aligned} \tag{8}$$

where $S(j + 1) \triangleq \{j, \dots, n - 1\}$ is an *active-state set* corresponding to a threshold policy, and quantities $w^{S(j+1)}(j - 1)$ are *marginal workloads*, which are recursively computed by

$$\begin{aligned} w^{S(2)}(0) &= \lambda(0) \frac{\alpha + \Delta d(1)}{\alpha + \lambda(0) + \mu(1)} \\ w^{S(j+1)}(j - 1) &= \lambda(j - 1) \frac{\alpha + \Delta d(j) + \frac{w^{S(j)}(j - 2)}{\rho(j - 2)}}{a(j) (\alpha + \lambda(j - 1) + \mu(j))}, \quad 2 \leq j \leq n - 1. \end{aligned} \tag{9}$$

In (9), quantities $a(j)$ are also recursively computed, by letting $a(1) = 1$, and

$$a(j) = 1 - \frac{\lambda(j - 1)\mu(j - 1)}{(\alpha + \lambda(j - 2) + \mu(j - 1))(\alpha + \lambda(j - 1) + \mu(j))} \frac{1}{a(j - 1)}, \quad 2 \leq j \leq n. \tag{10}$$

For the average criterion, one need simply set the value $\alpha = 0$ in the above recursions.

3.2 Application to the Present Model

To apply the above results to the present model, we must identify the corresponding parameters for each single-queue subproblem, and verify that they satisfy Ass. 1. Clearly, the parameters of concern should be defined as follows, where we have dropped the queue label k in the notation:

$$\begin{aligned} \lambda(j) &\triangleq \lambda, \quad 0 \leq j \leq n, \\ \mu(j) &\triangleq (j - m)^+ \theta + \min(j, m) \mu, \quad 1 \leq j \leq n \\ C(j) &\triangleq h(j) + (j - m)^+ c \theta - \min(j, m) r \mu, \quad 0 \leq j \leq n. \end{aligned} \tag{11}$$

With such definitions, it is readily verified that Ass. 1 holds, and, therefore, the results reviewed above apply to the model of concern in this paper.

3.3 Some Closed-Form First- and Second-Order MPI Formulae

In certain special cases, it is possible to solve explicitly the above index recursions to obtain closed-form expressions for the MPI.

Thus, in the classic case of multiple $M/M/1$ queues (no renegeing) in parallel, with $h(j) = hj$ and $r = 0$, under the average criterion, it is shown in Niño-Mora (2002a) that the MPI is given by

$$\nu^*(j) = \frac{c}{\mu} \sum_{i=1}^{j+1} (1 + \dots + \rho^{i-1}) = \begin{cases} \frac{c}{\mu} \left[\frac{\rho^{j+2} - 1}{(\rho - 1)^2} - \frac{j + 2}{\rho - 1} \right] & \text{if } \rho \neq 1 \\ \frac{c}{\mu} \frac{(j + 1)(j + 2)}{2} & \text{if } \rho = 1, \end{cases} \tag{12}$$

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

Another interesting case is that where the objective of concern is to *maximize the system's throughput*. In the case of multiple $M/M/1$ queues in parallel, with $h(j) \equiv 0$ and $r = 1$, under the average criterion, the index recursion above yields the *constant* MPI $\nu^*(j) \equiv -1$. Since such an index is noninformative, we proceed as in Niño-Mora (2006b) by introducing a tie-breaking, *second-order MPI* $\gamma^*(j)$, based on the MacLaurin expansion on the discounted MPI,

$$\nu^*(j) = -1 + \alpha \gamma^*(j) + O(\alpha), \quad \text{as } \alpha \searrow 0.$$

Thus, the corresponding index policy routes customers to a queue with smallest second-order MPI.

It is readily verified that such a second-order MPI is computed recursively by

$$\begin{aligned} \gamma^*(0) &= 1/\mu \\ \gamma^*(j) &= \gamma^*(j - 1) + \frac{1}{\mu} (1 + \dots + \rho^j), \quad 1 \leq j < n. \end{aligned} \tag{13}$$

Further, the solution to such a recursion is as follows. In the case $\rho = \lambda/\mu \neq 1$,

$$\gamma^*(j) = \frac{j + 1 - (j + 3)\rho + \rho^2 + \rho^{j+1}}{\mu(\rho - 1)^2},$$

and, in the case $\rho = 1$,

$$\gamma^*(j) = \frac{2 + j + j^2}{2\mu}.$$

4 Some Computational Experiments

This section reports on some preliminary experimental results on the relative performance of the proposed MPI policy. More thorough results, resulting from a large-scale computational study, will be reported in the full version of this paper.

In the following experiments, the performance of the MPI policy is compared against the optimal performance, and against the performance of two benchmark policies: the classic *Shortest Queue* (SQ) routing policy, and the *Smallest Expected Cost* (SEC) routing policy, which is the individually optimal policy. The experiments were performed with MATLAB 2006b, using implementations developed by the author. The optimal performance was computed by solving the linear programming formulation of the corresponding DP equations, using the CPLEX solver interfaced with MATLAB via TOMLAB. The alternative policies considered were evaluated by solving the appropriate linear evaluation equations.

In the first experiment, we investigate how relative performance varies with the arrival rate λ , using the the base instance having $K = 2$ queues in parallel, and the following parameters. Buffer sizes are $(n_1, n_2) = (25, 25)$, the numbers of servers are $(m_1, m_2) = (4, 5)$, the service rates are $(\mu_1, \mu_2) = (0.60, 0.40)$, the renegeing rates are $(\theta_1, \theta_2) = (0.2, 0.2)$, the holding cost rates are $(h_1, h_2) = (1, 1)$, the renegeing cost rates are $(c_1, c_2) = (0.3, 0.3)$, the completion rewards are $(r_1, r_2) = (1, 1)$, and the cost per rejected customer is $\nu = 0$. We investigate the model under the long-run average criterion, and in the version where only routing decisions are allowed.




Fig. 2. Exp. 1: Relative suboptimality gaps of MPI, SEH and SQ policies as λ varies




Fig. 3. Exp. 2: Relative suboptimality gaps of MPI, SEH and SQ policies as θ varies




Fig. 4. Exp. 3: Performance of MPI and optimal (OPT) routing policies as ν varies

Fig. 2 plots the relative suboptimality gaps of the MPI, SEC and SQ policies as the arrival rate λ varies over the interval $[1, 7]$. The plot shows that the MPI policy is nearly optimal throughout such a range, consistently outperforming the SEC and SQ policies. The SEC policy is better than the SQ policy in relatively light traffic, whereas the opposite holds in heavier traffic. The plot further shows that the three policies are asymptotically optimal in heavy traffic.

In the second experiment, we investigated how relative performance varies with a constant reneging rate θ , using a base instance as in the previous experiment, fixing the arrival rate to $\lambda = 2$. Fig. 3 plots the relative suboptimality gaps




Fig. 5. Exp. 4: Relative suboptimality gaps of MPI, SEH and SQ policies as λ varies

of the MPI, SEC and SQ policies as the reneging rate θ varies over the interval $[0.1, 0.35]$. Again, the plot shows that the MPI policy is nearly optimal throughout such a range, consistently outperforming the SEC and SQ policies. Among the latter, the SEC policy is the better one, with its relative suboptimality gap ranging from between 8% to 10%.

In the third experiment, we investigated how relative performance varies with the rejection cost rate ν , in the problem version with admission control capability, using a two-queue base instance with parameters $(n_1, n_2) = (25, 25)$, $(m_1, m_2) = (4, 5)$, $(\mu_1, \mu_2) = (0.60, 0.40)$, $(\theta_1, \theta_2) = (0.2, 0.2)$, $(h_1, h_2) = (1, 1)$, $(c_1, c_2) = (0.3, 0.3)$, $(r_1, r_2) = (3, 3)$, and $\lambda = 4$. Fig. 4 plots the average cost incurred by the MPI policy and the optimal average cost, as ν ranges over the interval $[1, 4]$. Again, the plot shows that the MPI policy is nearly optimal throughout such a range.

Finally, in the fourth experiment, we investigated how relative performance varies with the arrival rate λ in a problem where the objective is to maximize expected total discounted throughput, using the the base instance having $K = 2$ queues in parallel, and the following parameters: $(n_1, n_2) = (25, 25)$, $(m_1, m_2) = (4, 5)$, $(\mu_1, \mu_2) = (0.60, 0.40)$, $(\theta_1, \theta_2) = (0.2, 0.2)$, $(h_1, h_2) = (0, 0)$, $(c_1, c_2) = (0, 0)$, $(r_1, r_2) = (1, 1)$, and the cost per rejected customer is $\nu = 0$. We investigate the model under the discounted criterion, with discount rate $\alpha = 0.01$, letting λ range over $[1, 7]$.

The results are shown in Fig. 5, where it is seen that, again, the MPI policy is nearly optimal throughout the parameter range, while the SEC policy exhibits a moderately poor performance, and the SQ policy shows a substantially worse performance.

References

- Bassamboo, A., Harrison, J. M., Zeevi, A.: Dynamic routing and admission control in high-volume service systems: asymptotic analysis via multi-scale fluid limits. *Queueing Syst.* **51** (2005) 249–285
- Hordijk, A., Koole, G.: On the optimality of the generalized shortest queue policy. *Probab. Eng. Inf. Sci.* **4** (1990) 477–487
- Houck, D. J.: Comparison of policies for routing customers to parallel queueing systems. *Oper. Res.* **35** (1987) 306–310
- Johri, P. K.: Optimality of the shortest line discipline with state-dependent service rates. *European J. Oper. Res.* **41** (1989) 157–161
- Niño-Mora, J.: Restless bandits, partial conservation laws and indexability. *Adv. in Appl. Probab.* **33** (2001) 76–98
- Niño-Mora, J.: Dynamic allocation indices for restless projects and queueing admission control: a polyhedral approach. *Math. Program.* **93** (2002a) 361–413
- Niño-Mora, J.: Admission control and routing to parallel queues via Whittle’s restless bandit index policy. Presentation at the Workshop on Topics in Computer Communication and Networks, Isaac Newton Institute for Mathematical Sciences, Cambridge Univ., UK (available at <http://www.newton.cam.ac.uk/webseminars/pg+ws/2002/CMPW03/1217/nino-mora/>) (2002b)
- Niño-Mora, J.: Restless bandit marginal productivity indices, diminishing returns and optimal control of make-to-order/make-to-stock $M/G/1$ queues. *Math. Oper. Res.* **31** (2006a) 50–84
- Niño-Mora, J.: Marginal productivity index policies for scheduling a multiclass delay/loss-sensitive queue. *Queueing Syst.* **54** (2006b) 281–312
- Palm, C.: Contributions to the theory of delay systems. *Tele* **1** (1957) 37–67
- Whittle, P.: Restless bandits: Activity allocation in a changing world. A Celebration of Applied Probability, J. Gani (ed.), *J. Appl. Probab.* **25A** (1988) 287–298
- Winston, W.: Optimality of the shortest line discipline. *J. Appl. Probab.* **14** (1977) 181–189