



INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

*Optimal Routing Problems
and
Multimodularity*

Eitan Altman — Sandjai Bhulai — Bruno Gaujal — Arie Hordijk

N° 3727

June, 1999

THÈME 1

*R*apport
de recherche



Optimal Routing Problems and Multimodularity*

Eitan Altman[†], Sandjai Bhulai[‡], Bruno Gaujal[§], Arie Hordijk[¶]

Thème 1 – Réseaux et systèmes
Projet Mistral, Sloop

Rapport de recherche n° 3727 – June, 1999 – 25 pages

Abstract: In this paper we study the static assignment of packets to M parallel heterogeneous servers with no buffers. Blocked packets are lost and the objective is to minimize the average number of lost packets. The paper is divided into three parts. In the first part we formulate the problem as a Markov Decision Problem with partial observation and we derive an equivalent full observation problem. From this full observation problem we derive some structural results. We establish, in particular, the existence of an optimal periodic policy. The second part of the paper deals with computational aspects of the problem. We use multimodularity to derive an algorithm to compute optimal policies and show that for two servers the cost function is multimodular. In the third part we consider the same problem with well known arrival processes, such as the Poisson Process, the Markov Modulated Poisson Process and the Markovian Arrival Process.

Key-words: Multimodularity, optimal control, partially observed Markov Decision Process, static routing.

* The research presented in this paper has been supported by the Van-Gogh grant N. 98001

[†] INRIA, B.P. 93, 06902 Sophia Antipolis Cedex, France. E-mail: altman@sophia.inria.fr

[‡] Vrije Universiteit, Faculty of Sciences, Amsterdam, The Netherlands. E-mail: sbhulai@cs.vu.nl

[§] LORIA, B.P. 102, 54602 Villers-les-Nancy, France. E-mail: gaujal@loria.fr

[¶] Leiden University, P.O. Box 9512, Leiden, The Netherlands. E-mail: hordijk@wi.leidenuniv.nl

Problèmes de routage optimal et multimodularité

Résumé : Dans cet article, nous étudions l'attribution statique des paquets à M serveurs parallèles hétérogènes sans files d'attente. Les paquets bloqués sont perdus et l'objectif est de minimiser le nombre moyen de paquets perdus. Cet article est divisé en trois parties. Dans la première partie nous formulons le problème comme une chaîne de Markov contrôlée avec observation partielle. Ensuite nous décrivons un problème équivalent avec observation complet. De ce problème d'observation complet nous obtenons quelques résultats structurels. Nous montrons en particulier l'existence d'une politique optimale périodique. La deuxième partie de cet article traite les aspects numériques du problème. Nous employons la multimodularité pour obtenir un algorithme qui calcule des stratégies optimales et nous montrons que la fonction de coût est multimodulaire pour le cas de deux serveurs. Dans la troisième partie nous considérons le même problème avec quelques processus d'arrivée bien connus, tels que le processus de Poisson, le processus Poisson modulé et le processus d'arrivée Markovien.

Mots-clés : Multimodularité, contrôle optimal, chaîne de Markov contrôlée avec observation partielle, routage statique.

1 Introduction

In this paper we study the problem of optimal routing of arriving packets into a system having a fixed number of servers and no queues. The objective is to maximize the throughput, or equivalently to minimize the losses, when the controller has no information on the state of the servers. This problem can be classified as an open-loop optimal control problem. Similar problems have already been studied by Combé and Boxma [5], Rosberg and Towsley [18] and Koole [11]. However, exact results derived using these models depend on the assumption that the interarrival times of the arrival process are independent identically distributed and on the assumption that the service times are exponentially distributed.

Hajek [8] introduced a novel approach to solve such problems in the context of queueing systems, based on multimodularity. Elementary properties of multimodular functions and their relations with open-loop optimal control problems are further developed in Altman, Gaujal and Hordijk [1]. In Altman, Gaujal, Hordijk and Koole [2] a problem related to the one in this paper has been solved. Some multimodular and Schur convexity properties of the costs were obtained, which enabled to give some characterizations of optimal policies.

In this paper we obtain new insight on the structure of optimal policies using techniques from Markov decision processes. We show that our problem has a remarkable and surprising Markovian structure, even though the arrival process is a general stationary process. This allows us to establish the existence of optimal periodic policies. This implies in fact that only a finite number of states are needed in the Markov Decision Process for solving the problem.

We further obtain in this paper new multimodular properties as well as analytical expressions for the costs under general types of arrival processes (Poisson Process, Markov Modulated Poisson Process, Markovian Arrival Processes). We also study approaches for efficient numerical solutions of the control problem.

2 Optimal Routing to Servers

In this section we consider the problem of optimal routing of arriving packets into a system with a fixed number of servers having no waiting room. The servers work independently of each other and have service times, with general service distributions, independent of the arrival process. We assume that the process of interarrival times is stationary and that the controller has no information on the state of the servers. The objective in this setting is to maximize the expected throughput or equivalently to minimize the loss rate.

2.1 Problem Formulation

Consider a system, to which packets arrive according to an arrival process $(T_n)_{n \in \mathbb{N}}$, where T_n represents the arrival epoch. We use the convention that $T_1 = 0$ and we assume that the process $(\tau_n)_{n \in \mathbb{N}}$ of interarrival times $\tau_n = T_{n+1} - T_n$ is stationary, thus $(\tau_m, \dots, \tau_{m+h})$ and $(\tau_n, \dots, \tau_{n+h})$ are equal in distribution for all m, n and h .

Assume that the system has M parallel servers with no waiting room and independent service times, independent of the arrival process. Upon arrival each packet must be routed to one of the M servers in the system and is served with a service time, which has a general service distribution G_m when routed to server $m \in \{1, \dots, M\}$. If there is still a packet present at the server, where the packet is routed to, then the packet in service is lost (we call this the preemptive discipline). In the special case of an exponential distribution, one can consider instead the case of the non-preemptive discipline in which the arriving packet is lost; the results below will still hold. The model is illustrated in Figure 1.




Figure 1: Optimal routing to servers

A routing policy is a sequence $\pi = (\pi_1, \pi_2, \dots)$, where $\pi_n \in \{1, \dots, M\}$ means that the n^{th} packet is sent to server π_n . The actions are taken without any knowledge of the state of the servers, which means that the routing is *static*. The routing problem is to find a routing policy that maximizes the expected throughput of the packets in the system. Since the loss rate is the initial given input rate minus the actual throughput, this is equivalent to minimizing the average number of lost packets in the system. Note that whenever an actual arrival finds a packet at the server there is a loss, thus the cost function can depend on the interarrival times exclusively.

Let the random variable $L_n(\pi)$ denote the number of losses just after the n^{th} action is taken under a policy π . Then the mathematical formulation of the previously stated

objective is to minimize the following quantity over the set of all routing policies.

$$\limsup_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n \mathbb{E}^\pi L_i(\pi). \quad (1)$$

2.2 Markov Decision Problem with Partial Information

Using an approach similar to Koole [11], we formulate the model under consideration as a Markov decision problem with partially observed states. This will allow us to show that there exists an optimal periodic policy in our general setting of stationary arrival process and general independent service times.

Let $\bar{\mathcal{X}} = \mathbb{N}_0^M$ denote the state space. Let $\bar{x} = (\bar{x}_1, \dots, \bar{x}_M) \in \bar{\mathcal{X}}$. Then $\bar{x}_m = 0$ indicates that there is no packet at server m before the next decision. If $\bar{x}_m > 0$ then there is a packet present at server m and \bar{x}_m is the total number of arrivals that occurred at other servers since this packet arrived. We assume that the system is initially empty, thus the initial state is given by $(0, \dots, 0)$. Let $\bar{\mathcal{A}} = \{1, \dots, M\}$ denote the action set, where action $\bar{a} \in \bar{\mathcal{A}}$ corresponds to sending the packet to server \bar{a} . Let S_m be a random variable with distribution G_m , then define $f_m(n)$ by

$$f_m(n) = \mathbb{P}(S_m \geq \tau_1 + \dots + \tau_n). \quad (2)$$

Note that $f_m(n)$ represents the probability that a packet did not leave server m during n arrivals. Now the transition probabilities for the state components are given by

$$\bar{p}_m(\bar{x}_m, \bar{a}, \bar{x}'_m) = \begin{cases} 1, & \text{if } \bar{x}_m = 0, \bar{a} \neq m \text{ and } \bar{x}'_m = 0 \\ f_m(1), & \text{if } \bar{x}_m = 0, \bar{a} = m \text{ and } \bar{x}'_m = 1 \\ 1 - f_m(1), & \text{if } \bar{x}_m = 0, \bar{a} = m \text{ and } \bar{x}'_m = 0 \\ f_m(n+1), & \text{if } \bar{x}_m = n \text{ and } \bar{x}'_m = n+1 \text{ for } n \in \mathbb{N} \\ 1 - f_m(n+1), & \text{if } \bar{x}_m = n \text{ and } \bar{x}'_m = 0 \text{ for } n \in \mathbb{N} \\ 0, & \text{otherwise.} \end{cases}$$

The direct costs are given by

$$\bar{c}(\bar{x}, \bar{a}) = \begin{cases} 0, & \bar{x}_{\bar{a}} = 0 \\ 1, & \bar{x}_{\bar{a}} > 0. \end{cases}$$

The set of histories at epoch t of this Markov decision process is defined as the set $\bar{\mathcal{H}}_t = (\bar{\mathcal{X}} \times \bar{\mathcal{A}})^{t-1} \times \bar{\mathcal{X}}$. For a full state information Markov decision process, a policy

$\bar{\pi}$ is defined as a set of decision rules $(\bar{\pi}_1, \bar{\pi}_2, \dots)$ with $\bar{\pi}_t : \bar{\mathcal{H}}_t \rightarrow \bar{\mathcal{A}}$. Since we cannot observe the state of the servers in the above problem formulation, we restrict the policies to decision rules with $\bar{\pi}_t : \bar{\mathcal{A}}^{t-1} \rightarrow \bar{\mathcal{A}}$.

For each fixed policy $\bar{\pi}$ and each realization \bar{h}_t of a history, the random variable \bar{A}_t is given by $\bar{a}_t = \bar{\pi}_t(\bar{h}_t)$. The random variable \bar{X}_{t+1} takes values $\bar{x}_{t+1} \in \bar{\mathcal{X}}$ with probability $\bar{p}(\bar{x}_{t+1}, \bar{a}_t, \bar{x}_t)$. With these definitions the expected average cost criterion function $\bar{C}(\bar{\pi})$ is defined by

$$\bar{C}(\bar{\pi}) = \limsup_{T \rightarrow \infty} \mathbb{E} \frac{1}{T} \sum_{t=1}^T \bar{c}(\bar{X}_t, \bar{A}_t). \quad (3)$$

Let $\bar{\Pi}$ denote the set of all policies. The Markov decision problem is to find a policy $\bar{\pi}^*$, if it exists, such that $\bar{C}(\bar{\pi}^*) = \min\{\bar{C}(\bar{\pi}) \mid \bar{\pi} \in \bar{\Pi}\}$. This Markov decision model is characterized by the tuple $(\bar{\mathcal{X}}, \bar{\mathcal{A}}, \bar{p}, \bar{c})$.

Note that when the interarrival times are i.i.d. and the service times are exponential, the transition probabilities only depend on the presence of a packet at the server, whereas in our model we have to record the time a packet spends at the server. Therefore the model in Koole [11] uses the reduced state space $\{0, 1\}^M$.

2.3 Markov Decision Problem with Full Information

The standard approach to solve Markov decision problems with partial information is to transform these problems into equivalent Markov decision problems with full information. However, in general this transformation results in very big state spaces. Therefore algorithms which solve such Markov decision problems are often of limited use, since their complexity is very large and numerical results are difficult to obtain. In this section we formulate a Markov decision model $(\mathcal{X}, \mathcal{A}, p, c)$ with full state information, which has the property that it is equivalent to the partial information model and yet allows for numerical computation.

Let $\mathcal{X} = (\mathbb{N} \cup \{\infty\})^M$ be the state space. The m^{th} component x_m of $x \in \mathcal{X}$ denotes the number of arrivals that occurred since the last arrival to server m . We assume that the initial state in this case is (∞, \dots, ∞) . Let $\mathcal{A} = \bar{\mathcal{A}}$ be the corresponding action space. The transition probabilities are now given by

$$p(x, a, x') = \begin{cases} 1, & \text{if } x'_a = 1 \text{ and } x'_m = x_m + 1 \text{ for all } m \neq a \\ 0, & \text{otherwise.} \end{cases}$$

Let the direct cost $c(x, a) = f_a(x_a)$, where $f_a(\infty)$ equals zero. Note that the policies in Π are allowed to depend on X_t , since these random variables do not refer to the state

of the servers. The expected average cost criterion C defines the full state information Markov decision problem. The next theorem shows that the previously defined models are equivalent.

Theorem 2.1: The models $(\bar{\mathcal{X}}, \bar{\mathcal{A}}, \bar{p}, \bar{c})$ and $(\mathcal{X}, \mathcal{A}, p, c)$ are equivalent in the sense that policies in the two models can be transformed into each other leading to the same performances. Furthermore optimal policies exist for both models and the expected average cost of the optimal policies are equal.

Proof: Let $\bar{\pi} \in \bar{\Pi}$ be an arbitrary policy. Define the string of actions $(\bar{a}_1, \bar{a}_2, \dots)$ by $\bar{a}_1 = \bar{\pi}_1, \bar{a}_2 = \bar{\pi}_2(\bar{a}_1), \dots$. Now define the function $\Gamma : \bar{\Pi} \rightarrow \Pi$ such that $\pi = \Gamma(\bar{\pi})$ is the constant policy given by $\pi_t(h_t) = \bar{a}_t$ for all h_t .

From the definition of x_t it follows that the m^{th} component of x_t is equal to $(x_t)_m = t - \max_s \{\bar{a}_s = m\}$ if such an s exists and infinity otherwise. Now it follows that $\mathbb{P}((\bar{X}_t)_m = n)$ is non-zero only if $\bar{a}_{t-n} = m$. Therefore $\mathbb{P}((\bar{X}_t)_m = n) = f_m((x_t)_m)$ for $n = (x_t)_m$ and zero otherwise. Note that the expected cost only depends on the marginal probabilities. From the structure of the cost function and the transition probabilities it follows that $\mathbb{E} \bar{c}(\bar{X}_t, \bar{a}) = f_{\bar{a}}((x_t)_{\bar{a}}) = \mathbb{E} c(X_t, \bar{a})$. From this fact it follows that $\bar{C}(\bar{\pi}) = C(\Gamma(\bar{\pi}))$.

Let $\pi \in \Pi$ be an arbitrary policy. Note that there is only one sample path h_t , which occurs with non-zero probability, since the initial state is known and the transition probabilities are zero or one. Now define the equivalence relation \sim on Π such that $\pi \sim \pi'$ if $\pi_t(h_t) = \pi'_t(h_t)$. Now it directly follows that policies in the same class have the same cost. Furthermore it follows that the constant policy π'_t for each t is a representative element in the class, which is the image of a policy in $\bar{\Pi}$.

At this point we know that if there exists an optimal policy for the full observation problem, then there is also an optimal policy for the partially observed problem with the same value. In order to prove the existence of an optimal policy for our expected average cost problem, we first introduce the discounted cost:

$$V^\alpha(\pi, x) = \mathbb{E}^\pi \left[\sum_{t=1}^{\infty} \alpha^{t-1} C(X_t, A_t) \right].$$

The optimal discounted cost is defined as

$$V^\alpha(x) = \min\{V^\alpha(\pi, x) \mid \pi \in \Pi\}.$$

A sufficient condition for the existence of optimal stationary policy for the expected average cost is that $|V^\alpha(x) - V^\alpha(z)|$ is uniformly bounded in α and x for some z , where $V^\alpha(x)$ represents the minimal α -discounted costs (Theorem 2.2, Chapter 5 of Ross [19]).

Note that using the same policy for different initial states can only lead to a difference in cost when a packet is sent to a server for the first time. Since $0 \leq \alpha^t c(x, a) \leq 1$ it follows that the difference in cost cannot be larger than M . Now let π_z denote the α -optimal policy for initial state z and $V^\alpha(\pi_z, x)$ the value for the assignment done for initial state x using π_z . Then

$$V^\alpha(x) - V^\alpha(z) \leq V^\alpha(\pi_z, x) - V^\alpha(z) = V^\alpha(\pi_z, x) - V^\alpha(\pi_z, z) \leq M.$$

The proof is completed by repeating the argument for $V^\alpha(z) - V^\alpha(x)$ with π_x . \square

Instead of describing a policy using a sequence π , it will sometimes be more helpful to consider an equivalent description using time distances between packets routed to each server. More precisely, define for server m

$$\eta^m(0) = 0 \quad \text{and} \quad \eta^m(j) = \min\{i > \eta^m(j-1) \mid \pi_i = m\},$$

for $j \in \mathbb{N}$. Define the distance sequence δ^m by

$$\delta^m = (\delta_1^m, \delta_2^m, \dots), \quad \text{with} \quad \delta_j^m = \eta^m(j) - \eta^m(j-1),$$

for $j \in \mathbb{N}$. In order to show that these quantities are well defined, we show that each optimal policy uses every server infinitely many times. In the sequel we drop the bar in the notation, since we already established equivalence between the partial and full state information model.

Theorem 2.2: Every optimal policy $\pi = (\pi_1, \pi_2, \dots)$ has the property that every server is used infinitely many times, thus $\sup\{j \mid \pi_j = m\} = \infty$ for $m \in \{1, \dots, M\}$.

Proof: Assume that there is a server m , which is not used infinitely many times. Then there is an N such that $\pi_j \neq m$ for all $j \geq N$. Let $\pi' = (\pi_N, \pi_{N+1}, \dots)$. It follows that the difference of the total expected costs until any time n with $n > N$ between the policies π and π' is at most M . Since the criterion function looks at the expected average cost both policies have the same cost $C(\pi) = C(\pi')$.

Now choose n such that $f_m(n) < C(\pi')$. Note that this is possible since the cost incurred using an arbitrary policy is positive and since for all m , $\lim_{n \rightarrow \infty} f_m(n) = 0$. Define the policy π'' by

$$\pi''_j = \begin{cases} m, & \text{if } j \text{ is a multiple of } n \\ \pi'_{j - \lfloor j/n \rfloor}, & \text{otherwise.} \end{cases}$$

Note that the assignments for all j , which are not a multiple of n , are equal under both π' and π'' , but occur later under π'' and therefore $c(X_{j+\lfloor j/n \rfloor}(\pi''), A_{j+\lfloor j/n \rfloor}(\pi'')) \leq c(X_j(\pi'), A_j(\pi'))$ for all j which are not a multiple of n . It follows that $C(\pi'') \leq \frac{1}{n}f_m(n) + \frac{n-1}{n}C(\pi')$. By definition of n it follows that $C(\pi'') < C(\pi')$, thus π' is not optimal. This leads to a contradiction, because it follows that π is not optimal either. \square

Theorem 2.3: Let $f_m(n)$ be strictly positive and decreasing to zero in n for all m . Then there exists an optimal periodic policy.

Proof: Let π be an arbitrary policy and m be an arbitrary server. Suppose that the distance sequence $\delta = \delta^m$ for this server has the property that $\sup\{\delta_j \mid j \in \mathbb{N}\} = \infty$. We shall construct a policy π' with distance sequence $\delta^{\pi'}$ such that $\sup\{\delta_j^{\pi'} \mid j \in \mathbb{N}\}$ is finite and $C(\pi') \leq C(\pi)$.

Let $p = \min\{f_m(k) - f_m(k + 1) \mid m = 1, \dots, M, k = 1, \dots, M\}$. Note that p is positive, since f is positive and decreasing. Now choose n such that $2q = 2f_m(n) < p$. Note that this is possible, since $f_m(n)$ goes to zero as n goes to infinity. Since the supremum of the distance sequence is infinity, there is a $\delta_j > 2n + 2M + 1$. Consider the $2M + 1$ consecutive assignments starting n units after a packet is routed to server m for the j^{th} time. Since there are M servers, it follows that there is at least one server, to which a packet is assigned three times, say m' . Replace the second assignment to m' in this sequence by an assignment to server m . The situation is depicted in Figure 2.




Figure 2: Cost reduction

We consider the decrease in the total expected cost until the l th instant, where l is an arbitrary integer larger than the time at which a packet is routed to server m for the $j + 1$ st time. This decrease is given by

$$\begin{aligned} & [f_m(\delta_j) + f_{m'}(\delta_1) + f_{m'}(\delta_2)] - [f_m(\delta'_j) + f_m(\delta''_j) + f_{m'}(\delta_1 + \delta_2)] > \\ & [f_m(\delta_j) + f_{m'}(\delta_1) + f_{m'}(\delta_2)] - [2f_m(n) + f_{m'}(\delta_1 + 1)] = \\ & f_m(\delta_j) + f_{m'}(\delta_2) + [f_{m'}(\delta_1) - f_{m'}(\delta_1 + 1)] - [2f_m(n)] > \\ & f_m(\delta_j) + f_{m'}(\delta_2) + p - 2q > 0. \end{aligned}$$

The first inequality follows from the fact that $n < \delta'_j$, $n < \delta''_j$, $\delta_1 + \delta_2 \geq \delta_1 + 1$ and f is decreasing. The second inequality follows from the definition of p . Since f is positive it follows by construction of n that the last inequality holds.

Repeating this procedure for every $\delta_j > 2n + 2M + 1$ provides us a policy π' such that $C(\pi') \leq C(\pi)$ and $\sup\{\delta_j^{\pi'} \mid j \in \mathbb{N}\} < 2n + 2M + 1$. By repeating this procedure for every server, we get an optimal policy that visits a finite number of states. By Chapter 8 of Puterman [17] we know that the optimal policy can be chosen stationary. It follows that $\pi_n(h_n) = \pi_0(x_n)$. Since the state transitions are deterministic it follows that the optimal policy is periodic. \square

From (2) it follows that $f_m(n)$ satisfies the conditions in the previous theorem, since the random variables τ_n are positive. Therefore we know that the optimal policies are periodic. In the next section we will show how we can use this result in order to use multimodularity to find optimal routing policies.

2.4 Optimal Routing to Two Servers

In the previous section we showed that the optimal policy for all stationary arrival processes is periodic. In the case that there are only two servers in the system, we can say more about the structure of the optimal policy, using the fact that there is a full state information Markov decision process for this model.

In Liu and Towsley [13] it is pointed out that when the system has identical servers, thus the service distributions of the servers are equal, then the *round robin* policy is optimal, which means that sending a packet to each server alternatively is optimal. The next theorem shows that for two servers with exponential service times the optimal periodic policy is sending the first packet to the slowest server and then several packets to the fastest server.

Theorem 2.4: Suppose that the two servers have service times, which are exponentially distributed with rates μ_1 and μ_2 respectively such that $\mu_1 < \mu_2$. Then the optimal policy has the form $(1, 2, \dots, 2)^\infty$.

Proof: By Theorem 2.2 we know that all the servers are used infinitely many times. Thus the assignment to server 1 and consecutively to server 2 occurs infinitely many times. By Theorem 2.1 we know that the state after these occurrences are equal. Therefore the assignments between these occurrences must equal too. Now it immediately follows that there is an optimal policy of the form $\pi = (1, \dots, 1, 2, \dots, 2)$.

Now suppose that the number of 1's and 2's in such a policy are given by a and b respectively. Let $n = a + b$, then $C(\pi)$ is given by $nC(\pi) = (a - 1)f_1(1) + f_1(b + 1) + (b - 1)f_2(1) + f_2(a + 1)$. Let the policy π' be the policy where the first 1 is interchanged with the last 2. Then $nC(\pi') = (a - 2)f_1(1) + f_1(b) + f_1(2) + (b - 2)f_2(1) + f_2(a) + f_2(2)$. Now $n[C(\pi) - C(\pi')] = \{ [f_1(b + 1) - f_1(b)] - [f_1(2) - f_1(1)] \} + \{ [f_2(a + 1) - f_2(a)] - [f_2(2) - f_2(1)] \}$. From the convexity of $f_m(n)$ in n it follows that the previous quantity is positive. Therefore the form of the optimal policy is either $(1, 2, \dots, 2)$ or $(2, 1, \dots, 1)$. Since $\mu_1 < \mu_2$ and therefore $f_1(n) > f_2(n)$ for all $n \in \mathbb{N}$, it follows that $(1, 2, \dots, 2)$ is the optimal policy. \square

3 Multimodularity

Multimodularity was introduced by Hajek [8] for the study of admission control in queueing systems. The study focused on the assignment of a sequence of arrivals at a controller, who has to assign a certain fraction to an infinite queue. Hajek proved, by using multimodularity, that the optimal assignment has some regular structure, where the objective is to minimize the expected number of customers in the queue.

In this section we define multimodularity and study some properties of this concept. Furthermore we show that multimodularity provides us with a method to compute optimal routing policies. In particular we shall prove that in case we have two servers in the system, the cost function is multimodular.

3.1 General Properties

Let \mathbb{Z} be the set of integers. Let $e_i \in \mathbb{Z}^n$ for $i = 1, \dots, n$ denote the vector having all entries zero except for a 1 in the i^{th} entry. Let d_i be given by $d_i = e_{i-1} - e_i$ for $i = 2, \dots, n$. Then the base of vectors for multimodularity is defined as the collection

$\mathcal{B} = \{b_0 = -e_1, b_1 = d_2, \dots, b_{n-1} = d_n, b_n = e_n\}$. Thus we have

$$\begin{aligned} b_0 &= (-1, 0, 0, \dots, 0, 0, 0) \\ b_1 &= (1, -1, 0, \dots, 0, 0, 0) \\ b_2 &= (0, 1, -1, \dots, 0, 0, 0) \\ &\vdots \\ b_{n-2} &= (0, 0, 0, \dots, 1, -1, 0) \\ b_{n-1} &= (0, 0, 0, \dots, 0, 1, -1) \\ b_n &= (0, 0, 0, \dots, 0, 0, 1). \end{aligned}$$

Following Hajek, a function g defined on \mathbb{Z}^n is called *multimodular* if for all $x \in \mathbb{Z}^n$ and $b_i, b_j \in \mathcal{B}$ with $i \neq j$

$$g(x + b_i) + g(x + b_j) \geq g(x) + g(x + b_i + b_j). \quad (4)$$

Let g be a function on \mathbb{Z}^n . Define $\Delta_{e_i}g(x) = g(x + e_i) - g(x)$ for all $x \in \mathbb{Z}^n$. Since the unit vectors form a basis of \mathbb{Z}^n , every vector v is a linear combination of the unit vectors and can be written like $v = \sum_{i=1}^n \lambda_i e_i$. Then Δ_v is defined as $\Delta_v g(x) = \sum_{i=1}^n \lambda_i \Delta_{e_i}g(x)$. These definitions provide a useful characterization for multimodularity.

Lemma 3.1: ([1], Lemma 2.2) A function g is multimodular if and only if

$$\Delta_v \Delta_w g \leq 0,$$

for all $v, w \in \mathcal{B}$ with $v \neq w$.

In future discussions we need more than multimodularity with respect to a base \mathcal{B} only. It is often needed that the same function is also multimodular with respect to the base $-\mathcal{B}$. From the linearity of the operator Δ and Lemma 3.1 it immediately follows that g is also multimodular with respect to $-\mathcal{B}$.

3.2 Convexity

We define the notion of an atom in order to study some properties of multimodularity. In \mathbb{R}^n the convex hull of $n + 1$ affine independent points in \mathbb{Z}^n forms a *simplex*. This simplex, defined on x_0, \dots, x_n is called an *atom* if for some permutation σ of $(0, \dots, n)$

$$\begin{aligned} x_1 &= x_0 + b_{\sigma(0)}, \\ x_2 &= x_1 + b_{\sigma(1)}, \\ &\vdots \\ x_n &= x_{n-1} + b_{\sigma(n-1)}, \\ x_0 &= x_n + b_{\sigma(n)}. \end{aligned}$$

This atom is referred to as $A(x_0, \sigma)$ and the points x_0, \dots, x_n are called the *extreme points* of this atom. Each unit cube is partitioned into $n!$ atoms and all atoms together tile \mathbb{R}^n . We can use this to extend the function $g : \mathbb{Z}^n \rightarrow \mathbb{R}$ to the function $\bar{g} : \mathbb{R}^n \rightarrow \mathbb{R}$ as follows. If $x \in \mathbb{Z}^n$ then define the function by $\bar{g}(x) = g(x)$. Thus \bar{g} agrees with g on \mathbb{Z}^n . If x is not an extreme point of an atom, then x is contained in some atom A . The value of \bar{g} on such a point $x \in \mathbb{R}^n$ is obtained as the corresponding linear interpolation of the values of g on the extreme points of A . The following theorem shows the relation between multimodularity and convexity.

Theorem 3.2: ([8],[1], Theorem 2.1, [4]) A function g is multimodular if and only if the function \bar{g} is convex.

We defined multimodularity on the set of integers. Remark 2.1 of Altman, Gaujal and Hordijk [1] points out that for any convex subset S , which is a union of a set of atoms, multimodularity of g and convexity of \bar{g} still holds when restricted to S . Multimodularity on this convex subset means that (4) must hold for all $x \in S$ and $b_i, b_j \in \mathcal{B}$ with $i \neq j$, such that $x+b_i, x+b_j$ and $x+b_i+b_j \in S$. We will see that this is essential for the application in the next sections, where we use the set $S = \mathbb{N}_0^n$, with $\mathbb{N}_0 = \mathbb{N} \cup \{0\} = \{0, 1, 2, \dots\}$.

From the previous discussion it follows that multimodularity has a relation with convexity. In the next section we show how we can exploit this relation in order to minimize convex functions on the lattice.

3.3 Local Search

Let $g : \mathbb{Z}^n \rightarrow \mathbb{R}$ be an objective function in a mathematical model which has to be minimized. In general this is a hard problem to solve. In this section we show that if the function g is multimodular, then we can use a *local search* algorithm, which converges to the globally optimal solution. In Koole and Van der Sluis [12] the local search algorithm has been proven for the search space \mathbb{Z}^n . The following theorem shows that it also holds for any convex subset, which is a union of a set of atoms.

Theorem 3.3: Let g be a multimodular function on S , a convex subset which is a union of atoms. A point $x \in S$ is a global minimum if $g(x) \leq g(y)$ for all $y \neq x$ such that x and y are both extreme points of an atom in S .

Proof: Let $x \in S$ be a fixed point. Now suppose that there is a $z \in \mathbb{R}^n$ such that $x+z \in S$ and $\bar{g}(x+z) < \bar{g}(x) = g(x)$. We show that there is an atom $A(x, \sigma)$ in S with an extreme point y such that $g(y) < g(x)$.

Since any n vectors of \mathcal{B} form a basis of \mathbb{R}^n , we can write z as $z = \sum_{i=0}^n \beta_i b_i$. Furthermore, this can be done such that $\beta_i \geq 0$, since any $\beta_i b_i$ with $\beta_i < 0$ can be replaced by $-\beta_i \sum_{j=0, j \neq i}^n b_j$.

Now reorder the elements of \mathcal{B} as (b'_0, \dots, b'_n) and the elements $(\beta_0, \dots, \beta_n)$ as $(\beta'_0, \dots, \beta'_n)$ such that $\beta'_0 \geq \dots \geq \beta'_n \geq 0$ and $z = \sum_{i=0}^n \beta'_i b'_i$. Note that this notation is equivalent to the notation $z = \beta''_0 b'_0 + \beta''_1 (b'_0 + b'_1) + \dots + \beta''_n (b'_0 + \dots + b'_n)$ with all $\beta''_i \geq 0$ and with $b'_0 + \dots + b'_n = 0$. Now fix a non-zero point $z' = \alpha z$ with $\alpha < 1$ such that $\alpha (\beta''_0 + \dots + \beta''_{n-1}) \leq 1$. The set S is convex with $x, x + z \in S$, hence $x + z' \in S$. Since by Theorem 3.2 \bar{g} is convex and $\bar{g}(x + z) < g(x)$, it follows that $\bar{g}(x + z') < g(x)$.

Let σ be the permutation induced by $(\beta''_0, \dots, \beta''_n)$. Now consider the atom $A(x, \sigma)$, then by construction $x + z' \in A(x, \sigma)$. Since x is an extreme point and \bar{g} is linear, there must be another extreme point, say y , such that $g(y) < g(x)$. \square

This theorem shows that to check whether x is a globally optimal solution, it is sufficient to consider all extreme points of all atoms with x as an extreme point. When a particular extreme point of such an atom has a lower value than x , then we repeat the same algorithm with that point. Repeating this procedure is guaranteed to lead to the globally optimal solution.

Every multimodular function is also integer convex (see Theorem 2.2 in [1]). One could wonder if local search also works with integer convexity instead of multimodularity, which is a stronger property. The next counter-example in $(\{0, 1, 2\})^2$ shows that this is not true and that multimodularity is indeed the property needed for using local search. Define the function g such that

$$\begin{aligned} g(0, 2) &= -1, & g(1, 2) &= 2, & g(2, 2) &= 5, \\ g(0, 1) &= 2, & g(1, 1) &= 1, & g(2, 1) &= 0, \\ g(0, 0) &= 5, & g(1, 0) &= 4, & g(2, 0) &= 3. \end{aligned}$$

One can easily check that g is an integer convex function, but not multimodular since $g((1, 2) + b_0) + g((1, 2) + b_2) = 0 < 4 = g((1, 2)) + g((1, 2) + b_0 + b_2)$. Starting the local search algorithm at coordinate $(2, 1)$ shows that all neighbors have values, which are greater than 0. However, the global minimum is $g((0, 2)) = -1$.

Since all the extreme points of all atoms with x as an extreme point can be written as $x + \sum_{i=0}^n \alpha_i b_i$ with $\alpha_i \in \{0, 1\}$, $b_i \in \mathcal{B}$, the neighborhood of a point x consists of $2^{n+1} - 2$ points (we subtract 2 because when $\alpha_i = 0$ or when $\alpha_i = 1$ for all i , then the points coincide).

Although the complexity of the local search algorithm is big for large n , the algorithm is worthwhile studying. First of all, comparing to minimizing a convex function on the

lattice, this algorithm gives a reasonable improvement. Secondly, the algorithm can serve as a basis for heuristic methods.

In the next section we show that the cost function in the model with only two servers is multimodular.

3.4 Optimal Routing to Two Servers

In the previous section we showed that the optimal policy for all stationary arrival processes is periodic. In this case the notation of the distance sequence can be beneficially used to approach the decision problem. After the first assignment to server m , the distance sequence δ^m for server m is periodic, say with period $d(m)$. Since the servers are idle initially, the first assignment does not lead to a loss of a packet. Therefore in future discussions we will write $\pi = (\pi_1, \dots, \pi_n)$ for the periodic assignment sequence with period n and with a slight abuse of notation we denote the periodic distance sequence for server m by $\delta^m = (\delta_1^m, \dots, \delta_{d(m)}^m)$.

The periodicity reduces the cost function in complexity. Since we use the expected average number of lost packets as the cost function, we only have to consider the costs incurred during one period. The cost function for server m is given by

$$g_m(\pi) = \frac{1}{n} \sum_{j=1}^{d(m)} f_m(\delta_j^m). \quad (5)$$

The expected average number of lost packets in the total system is then given by

$$g(\pi) = \sum_{m=1}^M g_m(\pi) = \frac{1}{n} \sum_{m=1}^M \sum_{j=1}^{d(m)} f_m(\delta_j^m). \quad (6)$$

Note that equation (1) or equivalently (3) are exactly given by equation (6). It is tempting to formulate that $g_m(\pi)$ is multimodular in π for all m . Note that this is not necessarily true, since an operation $v \in \mathcal{B}$ applied to π leads to different changes in the distance sequences for different servers. In the case where we have two servers, we can prove multimodularity however, since an operation applied to the first server, leads to an opposite operation applied to the second server.

In the next discussion we assume that we have two servers and that π is a periodic sequence with period n . We will show that (6) is multimodular in π . After establishing multimodularity of the cost function for two servers, we can use the local search algorithm to find optimal policies for different server rates. In order to prove the multimodularity

we do not use (5), but instead we use $g'_m(\pi)$ defined by

$$g'_m(\pi) = \frac{1}{n} \sum_{j=1}^{d(1)} f_m(\delta_j^1). \quad (7)$$

This function only looks at the distance sequence assigned to the first queue. Note that $g(\pi)$ can now be expressed as $g(\pi) = g'_1(\pi) + g'_2(\vec{3} - \pi)$. We first prove that $g'_m(\pi)$ is multimodular in π .

Lemma 3.4: Let π be a fixed periodic policy with period n . Let $g'_m(\pi)$ be defined as in (7). Then $g'_m(\pi)$ is a multimodular function in π .

Proof: Since π is a periodic sequence, the distance sequence $\delta = \delta^1$ is also a periodic function, say with period $p = d(1)$. Now, define the function h_j for $j = 1, \dots, p$ by $h_j(\pi) = f_m(\delta_j)$. The function h_j represents the cost of the $(j+1)^{\text{st}}$ assignment by looking at the j^{th} interarrival time. We will first prove that this function is multimodular for $\mathcal{V} = \{b_1, \dots, b_{n-1}\}$.

Let $v, w \in \mathcal{V}$ with $v \neq w$. If none of these elements changes the length of the j^{th} interarrival time then $h_j(\pi) = h_j(\pi + v) = h_j(\pi + w) = h_j(\pi + v + w)$. Suppose that only one of the elements changes the length of the interarrival time, say v , then $h_j(\pi + v) = h_j(\pi + v + w)$ and $h_j(\pi) = h_j(\pi + w)$. In both cases the function $h_j(\pi)$ satisfies the conditions for multimodularity.

Now suppose that v adds and w decreases the length of the j^{th} interarrival time by one. Then $h_j(\pi + v) \leq h_j(\pi) = h_j(\pi + v + w) \leq h_j(\pi + w)$. Since h_j is a convex decreasing function, it follows that $h_j(\pi + w) - h_j(\pi + v + w) \geq h_j(\pi) - h_j(\pi + v)$. Now the multimodularity condition in equation (4) directly follows by rearranging the terms. Since $g'_m(\pi)$ is a sum of $h_j(\pi)$ it follows that $g'_m(\pi)$ is multimodular for \mathcal{V} .

Now consider the elements b_0 and b_n and note that the application of b_0 and b_n to π splits an interarrival period and merges two interarrival periods respectively. Therefore

$$\begin{aligned} n g'_m(\pi + b_0) &= n g'_m(\pi) - f_m(\delta_1) - f_m(\delta_p) + f_m(\delta_1 + \delta_p), \\ n g'_m(\pi + b_n) &= n g'_m(\pi) - f_m(\delta_p) + f_m(\delta_p - 1) + f_m(1), \\ n g'_m(\pi + b_0 + b_n) &= n g'_m(\pi) - f_m(\delta_1) - f_m(\delta_p) + f_m(\delta_1 + 1) + f_m(\delta_p - 1). \end{aligned}$$

Now $n[g'_m(\pi + b_0) + g'_m(\pi + b_n) - g'_m(\pi) - g'_m(\pi + b_0 + b_n)] = [f_m(\delta_1 + \delta_p) + f_m(1)] - [f_m(\delta_1 + 1) + f_m(\delta_p)]$. Let $k = \delta_1 + \delta_p + 1$. Since the function $f_m(x) + f_m(y)$ with $x + y = k$ is a symmetric and convex function, it follows from Proposition C2 of Chapter 3 of Marshall

and Olkin [15], that $f_m(x) + f_m(y)$ is also Schur-convex. Since $(\delta_1 + 1, \delta_p) \prec (\delta_1 + \delta_p, 1)$, the quantity above is non-negative.

In the case that we use $w = b_0$ and $v \in \mathcal{V}$ such that v does not alter δ_1 , then it follows that $g'_m(\pi + v + w) = g'_m(\pi + v) + g'_m(\pi + w) - g'_m(\pi)$. The same holds for $w = b_n$ and $v \in \mathcal{V}$ such that v does not alter δ_p . Suppose that v does alter δ_1 , then we have $n[g'_m(\pi + b_0) + g'_m(\pi + v) - g'_m(\pi) - g'_m(\pi + b_0 + v)] = [f_m(\delta_1 + \delta_p) + f_m(\delta_1 - 1)] - [f_m(\delta_1 + \delta_p - 1) + f_m(\delta_1)]$. When v alters δ_p we have $n[g'_m(\pi + b_n) + g'_m(\pi + v) - g'_m(\pi) - g'_m(\pi + b_n + v)] = [f_m(\delta_p + 1) + f_m(l)] - [f_m(\delta_p) + f_m(l + 1)]$ for some $l < \delta_p$. Now by applying the same argument as in the case of b_0 and b_n we derive multimodularity of $g'_m(\pi)$ for the base \mathcal{B} . \square

Now we will prove that $g(\pi)$, which is given by $g(\pi) = g'_1(\pi) + g'_2(\vec{3} - \pi)$ is multimodular. The proof is based on the fact that if a function is multimodular with respect to a base \mathcal{B} , then it is also multimodular with respect to $-\mathcal{B}$.

Theorem 3.5: Let g'_1 and g'_2 be multimodular functions. Then the function $g(\pi)$ given by $g(\pi) = c_1 g'_1(\pi) + c_2 g'_2(\vec{3} - \pi)$ for positive constants c_1 and c_2 is multimodular in π .

Proof: Let $v, w \in \mathcal{B}$, such that $v \neq w$. Then

$$\begin{aligned}
 & g(\pi+v) + g(\pi+w) \\
 &= c_1 g'_1(\pi+v) + c_2 g'_2(\vec{3} - \pi - v) + c_1 g'_1(\pi+w) + c_2 g'_2(\vec{3} - \pi - w) \\
 &= c_1 [g'_1(\pi+v) + g'_1(\pi+w)] + c_2 [g'_2(\vec{3} - \pi - v) + g'_2(\vec{3} - \pi - w)] \\
 &\geq c_1 [g'_1(\pi) + g'_1(\pi+v+w)] + c_2 [g'_2(\vec{3} - \pi) + g'_2(\vec{3} - \pi - v - w)] \\
 &= c_1 g'_1(\pi) + c_2 g'_2(\vec{3} - \pi) + c_1 g'_1(\pi+v+w) + c_2 g'_2(\vec{3} - \pi - v - w) \\
 &= g(\pi) + g(\pi+v+w).
 \end{aligned}$$

The inequality in the fourth line holds, since g'_1 is multimodular with respect to \mathcal{B} and g'_2 is multimodular with respect to $-\mathcal{B}$. \square

4 Arrival Processes

In the previous sections we derived an expression for the structure of the optimal policy, provided that the arrival process has stationary interarrival times. However, we were not able to explicitly formulate an expression for the period of this policy. In this section we assume that the service times are exponentially distributed with service rate μ_m and we look at some special cases of stationary arrival processes. We derive an explicit expression for the cost function, which will enable us to compute the optimal policy analytically when all parameters are specified.

4.1 Poisson Process

Consider a Poisson process with rate λ . The process of interarrival times of this Poisson process is independent identically distributed with probability density $f(t) = \lambda \exp(-\lambda t)$. Therefore it follows that (2) is given by

$$f_m(n) = \mathbb{E} \exp \left(-\mu_m \sum_{k=1}^n \tau_k \right) = \left[\mathbb{E} \exp(-\mu_m \tau_1) \right]^n.$$

Note that the fact that the process of interarrival times is i.i.d. simplifies the cost function. Using the probability density of the interarrival times, the quantity between the brackets in the right-hand side is given by

$$\mathbb{E} \exp(-\mu_m \tau_1) = \int_0^{\infty} \exp(-\mu_m t) \lambda \exp(-\lambda t) dt = \frac{\lambda}{\lambda + \mu_m}.$$

Now it follows that the cost function is given by

$$g(\pi) = \frac{1}{n} \sum_{m=1}^M \sum_{j=1}^{d(m)} \left[\frac{\lambda}{\lambda + \mu_m} \right]^{\delta_j^m}.$$

By choosing the Poisson process as the arrival process for our model, it becomes relatively simple to calculate the optimal policies analytically. For example: suppose that we have a Poisson process with rate $\lambda = 1$ and suppose that the rate of the first server is $\mu_1 = 1$. We can study the behavior of the period of the optimal policy as we increase the rate of the second server $\mu_2 \geq \mu_1$ as follows. Since the optimal policy is of the form $(1, 2, \dots, 2)$ we can parameterize the cost function by the period n and the server speed μ_2 given by

$$g(n, \mu_2) = \frac{1}{n} \left(\frac{1}{2} \right)^n + \frac{1}{n} \left(\frac{1}{1 + \mu_2} \right)^2 + \frac{n-2}{n} \left(\frac{1}{1 + \mu_2} \right).$$

By solving the equations $g(n, \mu_2) = g(n+1, \mu_2)$ for $n \geq 2$ we can compute the server rates μ_2 where the optimal policy changes. For example: the optimal policy changes from $(1, 2)$ to $(1, 2, 2)$ when $\mu_2 \geq 1 + \sqrt{2}$. The results of the computation are depicted in Figure 3.

4.2 Markov Modulated Poisson Process

The MMPP [9] is constructed by varying the arrival rate of a Poisson process according to an underlying continuous time Markov process, which is independent of the arrival




Figure 3: Relationship between n and μ_2

process. Therefore let $\{X_n \mid n \geq 0\}$ be a continuous time irreducible Markov process with m states. When the Markov process is in state i , arrivals occur according to a Poisson process with rate λ_i . Let p_{ij} denote the transition probability to go from state i to state j . Then the infinitesimal generator Q is given by

$$Q = \begin{pmatrix} -p_1 & p_{12} & \cdots & p_{1m} \\ p_{21} & -p_2 & \cdots & p_{2m} \\ \vdots & \vdots & \ddots & \vdots \\ p_{m1} & p_{m2} & \cdots & -p_m \end{pmatrix}, \text{ with } p_i = \sum_{\substack{j=1 \\ j \neq i}}^m p_{ij}.$$

Let $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_m)$ be the matrix with the arrival rates on the diagonal and $\lambda = (\lambda_1, \dots, \lambda_m)$ the vector of arrival rates. With this notation, we can use the matrix analytic approach to derive a formula for (2).

Define the matrix $F(t)$ is as follows. The elements $F_{ij}(t)$ are given by the conditional probabilities $\mathbb{P}(X_{n+1} = j, \tau_{n+1} \leq t \mid X_n = i)$ for $n \geq 1$. It follows that $F(\infty)$ given by $(\Lambda - Q)^{-1}\Lambda$ is a stochastic matrix.

Theorem 4.1: (Section 5.3 in [16]) The sequence $\{(X_n, \tau_n) \mid n \geq 0\}$ is a Markov renewal sequence with transition probability matrix $F(t)$ given by

$$F(t) = \int_0^t e^{(Q-\Lambda)u} du \Lambda = [I - e^{(Q-\Lambda)t}] (\Lambda - Q)^{-1} \Lambda.$$

The MMPP is fully parameterized by specifying the initial probability vector q , the infinitesimal generator Q of the Markov process and the vector λ of arrival rates. Let the row vector s be the steady state vector of the Markov process. Then s satisfies the equations $sQ = 0$ and $se = 1$, where $e = (1, \dots, 1)$. Define the row vector q by

$$q = \frac{1}{s\lambda} \cdot s\Lambda.$$

Then by Fisher and Meier-Hellstern [7] we know that q is the stationary vector of $F(\infty)$ and makes the MMPP interarrival process stationary. We can then use (2) to express the cost function for the MMPP. By Theorem 2.3 we know that the optimal policy is periodic. Therefore the cost is completely characterized by equations (5) and (6). If we assume that the underlying Markov chain of the MMPP is ergodic, it converges to the stationary regime as $t \rightarrow \infty$, and then the average cost is the same for any initial distribution.

In order to find an explicit expression for the cost function, we compute the Laplace-Stieltjes transform $f^*(\mu)$ of the matrix F . Since F is a matrix, we use matrix operations in order to derive $f^*(\mu)$, which will also be a matrix. The interpretation of the elements $f_{ij}^*(\mu)$ are given by $\mathbb{E}[e^{-\mu\tau_{n+1}} \mathbb{1}_{\{X_{n+1}=j\}} \mid X_n = i]$ for $n \geq 1$, where $\mathbb{1}$ is the indicator function. Let I denote the identity matrix, then $f^*(\mu)$ is given by

$$\begin{aligned} f^*(\mu) &= \int_0^\infty e^{-\mu t} F(dt) = \int_0^\infty e^{-\mu t} e^{(Q-\Lambda)t} (\Lambda - Q)(\Lambda - Q)^{-1} \Lambda dt \\ &= \int_0^\infty e^{-(\mu I - Q + \Lambda)t} dt \Lambda = (\mu I - Q + \Lambda)^{-1} \Lambda. \end{aligned}$$

In order to find an explicit expression for the cost function, we have to compute (2) first. This is done in the next lemma. Note that we do not need the assumption of independence of the interarrival times to derive this result.

Lemma 4.2: Let $f^*(\mu)$ be the Laplace-Stieltjes transform of F , where F is a transition probability matrix of a stationary arrival process. Then

$$f_m(n) = \mathbb{E} \exp\left(-\mu \sum_{k=1}^n \tau_k\right) = q [f^*(\mu)]^n e.$$

Proof: Define a matrix Q_n with entries $Q_n(i, j)$ given by

$$Q_n(i, j) = \mathbb{E} \left[\exp \left(-\mu \sum_{k=1}^n \tau_k \right) \mathbb{1}_{\{X_n=j\}} \mid X_0 = i \right].$$

Note that Q_1 is given by $f^*(\mu)$. By using the stationarity of the arrival process it follows that $Q_n(i, j)$ is recursively defined by

$$\begin{aligned} Q_n(i, j) &= \sum_{l=1}^m Q_{n-1}(i, l) \mathbb{E} \left[\exp(-\mu\tau_n) \mathbb{1}_{\{X_n=j\}} \mid X_{n-1} = l \right] \\ &= \sum_{l=1}^m Q_{n-1}(i, l) \mathbb{E} \left[\exp(-\mu\tau_1) \mathbb{1}_{\{X_1=j\}} \mid X_0 = l \right] \\ &= \sum_{l=1}^m Q_{n-1}(i, l) \cdot Q_1(l, j). \end{aligned}$$

Note that the last line exactly denotes the matrix product, thus $Q_n = Q_{n-1} \cdot Q_1$. By induction it follows that $Q_n = Q_1^n$. Then it follows that

$$f_m(n) = \mathbb{E} \exp \left(-\mu \sum_{k=1}^n \tau_k \right) = \sum_{i=1}^m \sum_{j=1}^m \mathbb{P}(X_0 = i) Q_n(i, j) = q [f^*(\mu)]^n e.$$

The last equation holds since the row vector q is the initial state of the Markov process and summing over all j is the same as right multiplying by e . \square

The previous lemma enables us to express the cost function (6) explicitly for the MMPP. The cost is now given by

$$g(\pi) = \frac{1}{n} \sum_{m=1}^M \sum_{j=1}^{d(m)} q [(\mu_m I - Q + \Lambda)^{-1} \Lambda]^{\delta_j^n} e.$$

Note that although we know the structure of the optimal policy, it is not intuitively clear that it is optimal in the case of the MMPP. The following example will clarify this statement. Suppose that one has an MMPP with two states. Choose the rates λ_1 and λ_2 of the Poisson processes such that the policies would have period 2 and 3 respectively if the MMPP is not allowed to change state. One could expect that if the transition probabilities to go to another state are very small, the optimal policy should be a mixture of both policies. But this is not the case as Theorem 2.4 shows.

4.3 Markovian Arrival Process

The Markovian arrival process model (MAP) is a broad subclass of models for arrival processes. It has the special property that every marked point process is the weak limit of a sequence of Markovian arrival processes (see Asmussen and Koole [3]). In practice this means that very general point processes can be approximated by appropriate MAP's. The utility of the MAP follows from the fact that it is a versatile, yet tractable, family of models, which captures both bursty inputs and regular inputs.

Markovian arrival processes can be modeled in a variety of ways. The model for the MAP described in Koole [10] has the advantage that the model is parameterized with few parameters. In Chapter 10 of Dshalalov [6] it is shown that the MAP can be modeled such that it is a natural generalization of the Poisson process. Furthermore it is shown that there is an analogous description for discrete PH-distributions. In this section we use the model formulation described in Lucantoni, Meier-Hellstern and Neuts [14], which reduces the computational complexity by using a matrix analytic approach. From this model it also easily follows that the MMPP is a special case of the MAP.

Model Formulation

Let $\{X_n | n \geq 0\}$ be a continuous time irreducible Markov process with m states. Assume that the Markov process is in state i . The sojourn time in this state is exponentially distributed with parameter γ_i . After this time has elapsed, there are two transition possibilities. Either the Markov process moves to state j with probability p_{ij} with generating an arrival or the process moves to state $j \neq i$ with probability q_{ij} without generating an arrival. With this definition it follows that

$$\sum_{\substack{j=1 \\ j \neq i}}^m q_{ij} + \sum_{j=1}^m p_{ij} = 1,$$

for all $1 \leq i \leq m$. This definition also gives rise to a natural description of the model in terms of matrix algebra. Define the matrix D with elements $D_{ij} = \gamma_i q_{ij}$ for $i \neq j$. Set the elements D_{ii} equal to $-\gamma_i$. Define the matrix E with elements $E_{ij} = \gamma_i p_{ij}$. The interpretation of these matrices is given as follows. The elementary probability that there is no arrival in an infinitesimal interval of length dt when the Markov process moves from state i to state j is given by $D_{ij} dt$. A similar interpretation holds for E , but in this case it represents the elementary probability that an arrival occurs. The infinitesimal generator of the Markov process is then given by $D + E$. Note that the MMPP can be derived by choosing $D = Q - \Lambda$ and $E = \Lambda$.

In order to derive an explicit expression for the cost function, we use the same approach as in the case of the MMPP. The transition probability matrix $F(t)$ of the Markov renewal

process $\{(X_n, \tau_n) \mid n \geq 0\}$, given in Lucantoni, Meier-Hellstern and Neuts [14], is of the form

$$F(t) = \int_0^t e^{Du} \, du = [I - e^{Dt}] (-D^{-1}E),$$

where again the elements $F_{ij}(t)$ of the matrix $F(t)$ are given by $\mathbb{P}(X_{n+1} = j, \tau_{n+1} \leq t \mid X_n = i)$ for $n \geq 1$. It also follows that $F(\infty)$ defined by $-D^{-1}E$ is a stochastic matrix. Let the row vector s be the steady state vector of the Markov process. Then s satisfies the equations $s(D + E) = 0$ and $se = 1$. Define the vector row vector q by

$$q = \frac{1}{sEe} sE.$$

Then q is the stationary vector of $F(\infty)$. This fact can be easily seen upon noting that $sE = s(D + E - D) = s(D + E) - sD = -sD$. With this observation it follows that $qF(\infty) = (sEe)^{-1} sDD^{-1}E = q$. The MAP defined by q , D and E has stationary interarrival times. In this case we know that the optimal policy is periodic and the cost is thus completely determined by equations (5) and (6).

The Laplace-Stieltjes transform $f^*(\mu)$ of the matrix F is given by

$$\begin{aligned} f^*(\mu) &= \int_0^\infty e^{-\mu t} F(dt) = \int_0^\infty e^{-\mu t} e^{Dt} (-D)(-D^{-1})E \, dt \\ &= \int_0^\infty e^{-(\mu I - D)t} \, dt E = (\mu I - D)^{-1}E. \end{aligned}$$

The interpretation of f^* is given by the elements $f_{ij}^*(\mu)$, which represent the expectation $\mathbb{E}[e^{\mu\tau_{n+1}} \mathbb{1}_{\{X_{n+1}=j\}} \mid X_n = i]$. By Lemma 4.2 we know that (2) is given by the product of f^* . Therefore the cost function, when using the MAP as arrival process, is given by

$$g(\pi) = \frac{1}{n} \sum_{m=1}^M \sum_{j=1}^{d(m)} q[(\mu_m I - D)^{-1}E]^{\delta_j^m} e.$$

References

- [1] E. Altman, B. Gaujal and A. Hordijk. Multimodularity, convexity and optimization properties. Technical Report TW-97-07, Leiden University, 1997.
- [2] E. Altman, B. Gaujal, A. Hordijk and G. Koole. Optimal admission, routing and service assignment control: the case of single buffer queues. Technical Report, INRIA, 1998.

-
- [3] S. Asmussen and G. Koole. Marked point processes as limits of Markovian arrival streams. *Journal of Applied Probability*, 30:365–372, 1993.
 - [4] M. Bartroli and S. Stidham, Jr., Towards a unified theory of structure of optimal policies for control of network of queues. Technical report, Department of Operations Research, University of North Carolina, Chapel Hill, 1987.
 - [5] M.B. Combé and O.J. Boxma. Optimization of static traffic allocation policies. *Theoretical Computer Science*, 125:17–43, 1994.
 - [6] J.H. Dshalalow. *Advances in Queueing: Theory, Methods and Open Problems*. CRC Press, 1995
 - [7] W. Fisher and K.S. Meier-Hellstern. The Markov-modulated Poisson process (MMPP) cookbook. *Performance Evaluation*, 18:149–171, 1992.
 - [8] B. Hajek. Extremal splitting of point processes. *Mathematics of Operations Research*, 10:543–556, 1985.
 - [9] H. Heffes and D.M. Lucantoni. A Markov modulated characterization of packetized voice and data traffic and related statistical multiplexer performance. *IEEE Journal on Selected Areas in Communications*, 4:856–868, 1986.
 - [10] G. Koole. *Stochastic Scheduling and Dynamic Programming*. PhD thesis, University of Leiden, 1992.
 - [11] G. Koole. On the static assignment to parallel servers. *IEEE Transactions on Automatic Control*, 1999. (to be published).
 - [12] G. Koole and E. van der Sluis. An optimal local search procedure for manpower scheduling in call centers. Vrije Universiteit Amsterdam.
 - [13] Z. Liu and D. Towsley. Optimality of the round-robin routing policy. *Journal of Applied Probability*, 31:466–475, 1994.
 - [14] D.M. Lucantoni, K.S. Meier-Hellstern and M.F. Neuts. A single-server queue with server vacations and a class of non-renewal arrival processes. *Advances in Applied Probability*, 22:675–705, 1990.
 - [15] A.W. Marshall and I. Olkin. *Inequalities: Theory of Majorization and its Applications*. Academic Press, 1979.

- [16] M.F. Neuts. *Structured Stochastic Matrices of M/G/1 Type and their Applications*. Marcel Dekker, 1989.
- [17] M.L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, 1994.
- [18] Z. Rosberg and D. Towsley. Customer routing to parallel servers with different rates. *IEEE Transactions on Automatic Control*, 30:1140–1143, 1985.
- [19] S.M. Ross. *Introduction to Stochastic Dynamic Programming*. Academic Press, 1993.



Unité de recherche INRIA Sophia Antipolis
2004, route des Lucioles - B.P. 93 - 06902 Sophia Antipolis Cedex (France)

Unité de recherche INRIA Lorraine : Technopôle de Nancy-Brabois - Campus scientifique
615, rue du Jardin Botanique - B.P. 101 - 54602 Villers lès Nancy Cedex (France)

Unité de recherche INRIA Rennes : IRISA, Campus universitaire de Beaulieu - 35042 Rennes Cedex (France)

Unité de recherche INRIA Rhône-Alpes : 655, avenue de l'Europe - 38330 Montbonnot St Martin (France)

Unité de recherche INRIA Rocquencourt : Domaine de Voluceau - Rocquencourt - B.P. 105 - 78153 Le Chesnay Cedex (France)

Éditeur
INRIA - Domaine de Voluceau - Rocquencourt, B.P. 105 - 78153 Le Chesnay Cedex (France)
<http://www.inria.fr>
ISSN 0249-6399