On Equilibrium Analysis of Acyclic Multiclass Loss Networks under Admission Control

Ashraf Al Daoud[a], Murat Alanyali[b], David Starobinski[b]

[a] Department of Communications Engineering, German-Jordanian University, Jordan
[b] Department of Electrical and Computer Engineering, Boston University, USA

Abstract

We consider equilibrium analysis of several dynamic resource sharing policies for multiclass loss networks with acyclic topologies. The policies of interest are based on the principle of prioritizing classes via thresholding or reservation. We show that under each policy the equilibrium network state is a Markov random field and we obtain closed form expressions for the conditional probabilities therein. Such representations drastically reduce the computational complexity of blocking probability and revenue calculations. We provide revenue comparison of the considered policies and several extensions of the applied analytical technique.

Keywords: Dynamic Resource Allocation, Markov Random Fields.

1. Introduction

Revenue management via admission control arises in a variety of contexts ranging from logistic service provision to telephone networks. This paper is motivated by the renewed interest due to recent regulatory progress that paves the way to dynamic spectrum access by so-called primary and secondary users. In this context, resources are radio channels that cannot be utilized simultaneously in neighboring locations due to interference; and a spectrum provider may augment its revenue from a primary customer base by providing opportunistic access to secondary users subject to an admission policy.

We analyze four admission control policies on linear (and more generally acyclic) network topologies that may model highways and beaches. Considered policies are based on the concepts of reservation, which is optimal for an isolated resource [1], and thresholding [2]. Equilibrium analyses of each policy under standard Markovian models are computationally intense. We show that for linear network topologies, equilibrium system state under each policy can be explicitly represented as a Markov random field. The conditional probabilities in such representation are identified, thereby leading to substantial computational convenience in calculating blocking probabilities and network revenue. We provide a numerical comparison of revenue rates under the considered policies, and a discussion of generalizations of the applied techniques.

Connections between Markov random fields and loss networks have been previously studied. Exact analyses in the literature are limited to uncontrolled networks with acyclic topologies [3, 4, 5, 6]. Here the term topology specifies a structure for routes in the generic loss network formulation as specified by [5]. Analyses of these papers rely also on reversibility of network state in the uncontrolled case. This property holds for threshold policies, for which equilibrium distributions can be obtained in closed form but entail computational difficulty due to normalizing constants [2]. Reversibility generally does not hold under reservation policies. In turn these policies have been studied via fixed-point [5, 7, 8] and asymptotically exact approximations [9, 10]. The technical contribution of the present paper is exact analysis of controlled loss networks under two canonical admission philosophies in a particular setting. The analysis here also relies on reversibility, however the setting is different than those of [3, 4, 5, 6] due to both admission control and the composition of considered routes; in turn the analyses of these papers do not directly apply here. The present paper also provides a streamlined approach to identifying the field potentials in terms of probability transition matrices, which is a key to the alluded computational savings.

2. Resource Sharing Policies

Consider $2N + 1$ nodes indexed by the integers

$$\{-N, -N + 1, \cdots, N - 1, N\}$$

Email addresses: ashraf.aldaoud@gju.edu.jo (Ashraf Al Daoud), alanyali@bu.edu (Murat Alanyali), staro@bu.edu (David Starobinski)
[1] This work was funded by NSF grant CCF-0964652

and let an edge be drawn between nodes i and j if $|i - j| = 1$. We shall interpret each edge as a collection of C resources. We consider two types of dynamic requests for resources: At each node, type k ($k = 1, 2$) requests arrive according to a Poisson process with rate λ_k. If honored, such a request generates a revenue of r_k units and it holds one resource at each edge that is incident to the node of arrival. Arrival processes across nodes and types are independent, and holding times are exponentially distributed with unit mean, independently of the history prior to the request arrival.

Let $n_{i,k}$ denote the number of type k requests in service at node i. We shall denote the detailed occupancy of node i by

$$\mathbf{n}_i \doteq (n_{i,1}, n_{i,2})$$

and the total occupancy of the node by

$$s_i \doteq n_{i,1} + n_{i,2}.$$

Hence the vector $(\mathbf{n}_{-N}, \cdots, \mathbf{n}_N)$ of detailed occupancies necessarily satisfies

$$s_i + s_{i+1} \leq C, \quad \text{for all } \; i = -N, \cdots, N-1. \tag{1}$$

We consider the following four policies to share the available resources between the two types of requests:

1. **Complete Sharing Policy:** A request, of either type, is admitted whenever the network has resources available to accommodate the request. Namely, a request is admitted if and only if its inclusion preserves condition (1).
2. **Threshold Policy:** This policy prioritizes type 1 requests by enforcing a threshold T ($0 \leq T \leq C$) on the number of resources that can be unavailable due to ongoing type 2 requests. Namely, a type 1 request is admitted if that does not violate conditions (1), whereas a type 2 request is admitted if, in addition to (1), its inclusion preserves

$$n_{i,2} + n_{i+1,2} \leq T, \quad i = -N, \cdots, N-1. \tag{2}$$

3. **Complete Partitioning Policy:** Under a complete partitioning policy, the resources are statically divided between the two types so that K resources are allocated to type 1 and the remaining $C - K$ resources are available exclusively to type 2 requests. A type 1 request is then admitted if its inclusion preserves

$$n_{i,1} + n_{i+1,1} \leq K, \quad i = -N, \cdots, N-1.$$

A type 2 request is admitted if its inclusion preserves

$$n_{i,2} + n_{i+1,2} \leq C - K, \quad i = -N, \cdots, N-1.$$

4. **Reservation Policy:** This policy prioritizes type 1 requests by rejecting type 2 requests above a certain occupancy level. Specifically, each node has a reservation parameter R ($0 \leq R \leq C$). A type 1 request is admitted if its admission preserves conditions (1), whereas a type 2 request is admitted if, besides (1), its admission preserves the total number of requests in service at the node below R; i.e., *after* inclusion

$$s_i \leq R, \quad i = -N, \cdots, N.$$

Note that the reservation parameter acts on s_i rather than $s_i + s_{i+1}$ so the policy differs slightly from trunk reservation policies considered earlier for circuit switched networks [7].

We point out that complete sharing and complete partitioning can be analyzed via special cases of threshold and reservation policies; so it suffices to analyze only the latter two policies.

3. Equilibrium Node Occupancies

We define the network state $\mathbf{n}$ as

$$\mathbf{n} \doteq (\mathbf{n}_{-N}, \cdots, \mathbf{n}_N).$$

Under each of the above policies, $\mathbf{n}$ is an irreducible Markov process with a finite state space; therefore it possesses a unique equilibrium distribution π. The long-term rate of revenue generation is denoted by

$$W \doteq \sum_{i=-N}^{N} \sum_{k=1,2} r_k \lambda_k (1 - B_k^{(i)}),$$

where $B_k^{(i)}$ is the blocking probability of type k requests at node i. By the PASTA property, $B_k^{(i)}$ can be expressed in terms of the equilibrium distribution π. In this section, we provide closed form expressions for π that are particularly convenient for computing the blocking probabilities and in turn the revenue rate under each policy.

We define an occupancy state $\mathbf{n}$ to be *admitting* for type k at node i if a request of type k can be admitted at the node i at that network state. Let us denote such states by $\mathcal{A}^{(i,k)}$. For example, for a threshold policy, admitting states for type 2 requests at node i, $\mathcal{A}^{(i,2)}$, are the states $\mathbf{n}$ that satisfy both conditions (1) and (2) with strict inequality. Due to the PASTA property, blocking probability of type k requests at node i is given by

$$B_k^{(i)} = 1 - \sum_{\mathbf{n}:\mathbf{n}\in\mathcal{A}^{(i,k)}} \pi(\mathbf{n}). \tag{3}$$

For large values of parameters N and C, determining blocking probabilities via expression (3) is computationally intense. This is mainly due to two reasons: (i) The state space of the occupancy process $\mathbf{n}$ increases with the network size and the number of request types. Even in cases when the distribution π has a product form (as for the complete sharing policy), computation of π requires determining a normalizing constant that is NP-hard [11], and (ii) once π is obtained, computation of $B_k^{(i)}$ requires a large number of summands due to the statistical dependence among node occupancies.

Here we show that the alluded difficulties can be avoided in linear (more generally acyclic) network topologies for the considered sharing policies. Namely, we establish that under each policy the distribution π admits the representation

$$\pi(\mathbf{n}) \;=\; \hat{\pi}_N(\mathbf{n}_0)\prod_{i=0}^{N-1} P_{N-i}(\mathbf{n}_i,\mathbf{n}_{i+1})P_{N-i}(\mathbf{n}_{-i},\mathbf{n}_{-i-1}), \tag{4}$$

where $\{\hat{\pi}_N(\mathbf{n}_0) : \; n_{0,1}, n_{0,2} = 0, 1, \cdots, C\}$ is a probability distribution and P_j, $j = 1, \cdots, N$ are stochastic matrices. Here, enumeration of matrix rows and columns with two-dimensional vectors $\mathbf{n}_i$ is arbitrary but consistent for all matrices P_j. The structure of form (4) yields that π is a Markov random field with respect to the linear network graph. In the terminology of Markov random fields, the vector $\hat{\pi}_N$ is a node potential and the matrices P_j are edge potentials.

Under the representation (4), the equilibrium variables

$$\mathbf{n}_i, \mathbf{n}_{i+1}, \mathbf{n}_{i+2}, \cdots$$

form a non-homogeneous Markov process whose transition probabilities are determined by the stochastic matrices P_j. An immediate consequence of this observation is that the marginal distribution of the network load at the neighborhood of a given node admits a succinct characterization, and thereby leads to efficient computation of blocking probabilities at that node. For example, for a threshold policy with threshold T, blocking probability of type 2 requests at node 0 is given by

$$B_2^{(0)} = 1 - \sum_{(\mathbf{n}_{-1},\mathbf{n}_0,\mathbf{n}_1)\in\mathcal{A}} P_N(\mathbf{n}_0,\mathbf{n}_1)\hat{\pi}_N(\mathbf{n}_0)P_N(\mathbf{n}_0,\mathbf{n}_{-1}), \tag{5}$$

where $\mathcal{A}$ is the set of triplets $(\mathbf{n}_{-1}, \mathbf{n}_0, \mathbf{n}_1)$ such that

$$\begin{aligned} s_0 + s_1 &< C, & s_0 + s_{-1} &< C,\\ n_{0,2} + n_{1,2} &< T, & n_{0,2} + n_{-1,2} &< T. \end{aligned}$$

Blocking probabilities at other nodes can be obtained through similar computations. Note that the number of possible values of $\mathbf{n}_i$ is at most quadratic in C and membership in the set $\mathcal{A}$ is verified in triplets; so (5) entails $O(C^6)$ operations.

We next identify the matrices P_j and the distribution $\hat{\pi}_N$ that render the form (4) an equilibrium distribution for the node occupancies under a threshold policy with parameter T. This will be done by first finding a nonnegative matrix Q and a nonnegative vector ϕ such that the product

$$\phi(\mathbf{n}_0)\prod_{i=0}^{N-1} Q(\mathbf{n}_i,\mathbf{n}_{i+1})Q(\mathbf{n}_{-i},\mathbf{n}_{-i-1}) \tag{6}$$

satisfies the detailed balance equations for the process $\mathbf{n}$. This would imply that the process is reversible. Furthermore, it would imply that expression (6) would give the equilibrium distribution of π after being properly normalized to become a probability distribution on the state space associated with the threshold policy. Such a normalization can be imposed on the individual factors in (6) by defining $P_j(\mathbf{x},\mathbf{y})$ from (4) as

$$P_j(\mathbf{x},\mathbf{y}) \doteq \frac{Q(\mathbf{x},\mathbf{y})\sum_{\mathbf{z}} Q^{j-1}(\mathbf{y},\mathbf{z})}{\sum_{\mathbf{z}} Q^j(\mathbf{x},\mathbf{z})}, \quad j = 1, \cdots, N \tag{7}$$

where Q^j denotes the j^{th} matrix power of Q, and

$$\hat{\pi}_N(\mathbf{x}) \doteq G\phi(\mathbf{x})\left(\sum_{\mathbf{z}} Q^N(\mathbf{x},\mathbf{z})\right)^2 \tag{8}$$

where

$$G^{-1} = \sum_{\mathbf{x}} \phi(\mathbf{x})\left(\sum_{\mathbf{z}} Q^N(\mathbf{x},\mathbf{z})\right)^2. \tag{9}$$

The sums in the above definitions run through all possible values of associated entries, which will be specified below. Note that under definitions (7)-(9) the product on the right hand side of (4) is equal to

$$\frac{\hat{\pi}_N(\mathbf{n}_0)}{\left(\sum_{\mathbf{z}} Q^N(\mathbf{n}_0,\mathbf{z})\right)^2}\prod_{i=0}^{N-1} Q(\mathbf{n}_i,\mathbf{n}_{i+1})Q(\mathbf{n}_{-i},\mathbf{n}_{-i-1}). \tag{10}$$

It is easy to verify that if (6) satisfies the detailed balance equations for $\mathbf{n}$, then so does (10); therefore (4) holds.

We next carry out the program outlined above. Let

$$M \doteq \{\mathbf{x} = (x_1, x_2) \in \mathbb{Z}_+^2 : x_1 + x_2 \leq C,\ x_2 \leq T\}.$$

Theorem 3.1. *(Threshold policy): Consider a threshold policy with parameter T. Let Q be the $|M| \times |M|$ matrix defined by setting for each $\mathbf{x}, \mathbf{y} \in M$*

$$Q(\mathbf{x}, \mathbf{y}) = \begin{cases} \frac{\lambda_1^{y_1}}{y_1!} \frac{\lambda_2^{y_2}}{y_2!} & \text{if } \begin{cases} \sum_{k=1}^2 x_k + y_k \leq C \text{ and} \\ x_2 + y_2 \leq T \end{cases} \\ 0 & \text{otherwise.} \end{cases}$$

Let ϕ be the M-dimensional vector defined by setting

$$\phi(\mathbf{x}) = \frac{\lambda_1^{x_1}}{x_1!} \frac{\lambda_2^{x_2}}{x_2!}, \qquad \mathbf{x} \in M.$$

Then (6) satisfies the detailed balance equations for $\mathbf{n}$; in turn the equilibrium distribution π of $\mathbf{n}$ is given by (4) subject to definitions (7)-(9).

Proof. Let $I\{.\}$ be the indicator function that is 1 if its argument is true and 0 otherwise. Also define

$$\mathbf{e}_1 \doteq (1, 0) \quad \text{and} \quad \mathbf{e}_2 \doteq (0, 1),$$

and for each $i = -N, \cdots, N$ the following vectors of tuples that have the same dimensionality as $\mathbf{n}$:

$$\mathbf{e}_{i,1} \doteq \left((0,0), \cdots, \underbrace{(1,0)}_{i^{th} entry}, \cdots, (0,0) \right),$$

$$\mathbf{e}_{i,2} \doteq \left((0,0), \cdots, \underbrace{(0,1)}_{i^{th} entry}, \cdots, (0,0) \right).$$

Note that jumps of the process $\mathbf{n}$ are of the form $\pm\mathbf{e}_{i,k}$ for some node i and some type k.

Let us refer to elements of the state space of the network occupancy process as feasible states. For a threshold policy with parameter T, state $\mathbf{n}$ is feasible if and only if it satisfies conditions (1) and (2). For each feasible state $\mathbf{n}$ let

$$\mu(\mathbf{n}) \doteq \phi(\mathbf{n}_0) \prod_{i=0}^{N-1} Q(\mathbf{n}_i, \mathbf{n}_{i+1}) Q(\mathbf{n}_{-i}, \mathbf{n}_{-i-1}).$$

For $i \neq 0$ and $k = 1, 2$

$$\begin{aligned} \frac{\mu(\mathbf{n} + \mathbf{e}_{i,k})}{\mu(\mathbf{n})} &= \frac{Q(\mathbf{n}_{i-1}, \mathbf{n}_i + \mathbf{e}_k) Q(\mathbf{n}_i + \mathbf{e}_k, \mathbf{n}_{i+1})}{Q(\mathbf{n}_{i-1}, \mathbf{n}_i) Q(\mathbf{n}_i, \mathbf{n}_{i+1})} \\ &= \frac{\lambda_k}{n_{i,k} + 1} I\{\mathbf{n} + \mathbf{e}_{i,k} \text{ is feasible}\}. \end{aligned}$$

The second equality above follows since both $Q(\mathbf{n}_{i-1}, \mathbf{n}_i + \mathbf{e}_k)$ and $Q(\mathbf{n}_i + \mathbf{e}_k, \mathbf{n}_{i+1})$ are nonzero if only if

1) $s_i + s_{i\pm1} + 1 \leq C$, and

2) $n_{i,2} + n_{i\pm1,2} + 1 \leq T$ when $k = 2$.

In that case $\mathbf{n} + \mathbf{e}_{i,k}$ is feasible because $\mathbf{n}$ is feasible and it differs from $\mathbf{n} + \mathbf{e}_{i,k}$ only in the ith entry; furthermore

$$\begin{aligned} Q(\mathbf{n}_{i-1}, \mathbf{n}_i + \mathbf{e}_k)/Q(\mathbf{n}_{i-1}, \mathbf{n}_i) &= \lambda_k/(n_{i,k} + 1) \\ Q(\mathbf{n}_i + \mathbf{e}_k, \mathbf{n}_{i+1})/Q(\mathbf{n}_i, \mathbf{n}_{i+1}) &= 1 \end{aligned}$$

due to the definition of Q.

For $i = 0$,

$$\begin{aligned} \frac{\mu(\mathbf{n} + \mathbf{e}_{0,k})}{\mu(\mathbf{n})} &= \frac{Q(\mathbf{n}_0 + \mathbf{e}_k, \mathbf{n}_{-1}) \phi(\mathbf{n}_0 + \mathbf{e}_k) Q(\mathbf{n}_0 + \mathbf{e}_k, \mathbf{n}_1)}{Q(\mathbf{n}_0, \mathbf{n}_{-1}) \phi(\mathbf{n}_0) Q(\mathbf{n}_0, \mathbf{n}_1)} \\ &= \frac{\lambda_k}{n_{0,k} + 1} I\{\mathbf{n} + \mathbf{e}_{0,k} \text{ is feasible}\}. \end{aligned}$$

The second equality above follows from the definitions of the matrix Q and the vector ϕ. Therefore μ solves the detailed balance equations for $\mathbf{n}$; and in turn π is given by (4). ☐

We next turn to the reservation policy. The network occupancy $\mathbf{n}$ is generally not reversible under reservation policies. Perhaps the easiest way to verify this is to consider a single node (i.e. $N = 0$) and inspect the asymmetries in the state transition diagram of $\mathbf{n}$. While the techniques of Theorem 3.1 cannot be applied to obtain the equilibrium distribution π due to irreversibility, we shall show that a coarser state descriptor

$$\mathbf{s} \doteq (s_{-N}, \cdots, s_N)$$

is reversible and it can be analyzed by analogous techniques. (Recall that $s_i = n_{i,1} + n_{i,2}$ is the total occupancy of node i.) Under reservation policies, $\mathbf{s}$ is a Markov process since request admission decisions are based on total node occupancies and request holding times have identical statistics for the two types.

Let π_S denote the unique equilibrium distribution of $\mathbf{s}$ under a reservation policy with parameter R. We establish reversibility of $\mathbf{s}$ by showing that π_S admits the following form

$$\pi_S(\mathbf{s}) = \hat{\pi}_N(s_0) \prod_{i=0}^{N-1} P_{N-i}(s_i, s_{i+1}) P_{N-i}(s_{-i}, s_{-i-1}), \tag{11}$$

where $\hat{\pi}_N(s_0) : s_0 = 0, 1, 2, \cdots, C$ is a probability vector and $P_j,\ j = 1, \cdots, N$ are $(C + 1) \times (C + 1)$ stochastic matrices. Towards that end, it is enough to show that detailed balance equations for **s** are satisfied by

$$\phi(s_0)\prod_{i=0}^{N-1} Q(s_i, s_{i+1})Q(s_{-i}, s_{-i-1})$$

for some nonnegative $(C + 1)$-dimensional vector ϕ and $(C + 1) \times (C + 1)$ matrix Q, since then one can set

$$P_j(x, y) = \frac{Q(x, y)\sum_{z=0}^{C} Q^{j-1}(y, z)}{\sum_{z=0}^{C} Q^j(x, z)}, \quad j = 1, \cdots, N \qquad (12)$$

and

$$\hat{\pi}_N(x) = G\phi(x)(\sum_{z=0}^{C} Q^N(x, z))^2 \qquad (13)$$

where $x, y = 0, 1, \cdots, C$ and G is a normalizing constant.

Theorem 3.2. *(Reservation policy): Given a reservation policy with parameter R, let the matrix $(Q(x, y) : x, y = 0, 1, \cdots, C)$ be defined by setting*

$$Q(x, y) = \begin{cases} \frac{(\lambda_1+\lambda_2)^y}{y!} & \text{if } x + y \le C, \text{ and } y \le R \\ \frac{(\lambda_1+\lambda_2)^R \lambda_1^{y-R}}{y!} & \text{if } x + y \le C, \text{ and } y > R \\ 0 & \text{otherwise,} \end{cases}$$

and for $x = 0, 1, \cdots, C$ let

$$\phi(x) = \begin{cases} \frac{(\lambda_1+\lambda_2)^x}{x!} & \text{if } x < R \\ \frac{(\lambda_1+\lambda_2)^R \lambda_1^{(x-R)}}{x!} & \text{if } x \ge R. \end{cases}$$

Then π_S is given by (11) subject to definitions (12)–(13).

Proof. The theorem is proved by verifying the detailed balance equations as in the proof of Theorem 3.1 and thus the proof is omitted. □

4. Approximations for Large Topologies

In this section we consider an approximate method that entails further computational reductions, and that is asymptotically exact as N tends to ∞. The approximation is based on the observation that in a large linear topology most of the nodes are far away from the two network boundaries; if the teletraffic statistics are homogeneous in the network then marginal distributions of the network occupancy **n** become almost uniform and the concept of a typical node emerges as N tends ∞. A good candidate for a typical node is the center node whose limiting distribution is determined by the asymptotic characterization of P_N and $\hat{\pi}_N$.

Consider first the threshold policy. Let P_j and $\hat{\pi}_N$ be defined by equalities (7) and (8) respectively, and define

$$\begin{aligned} P(\mathbf{x}, \mathbf{y}) &= \lim_{N\to\infty} P_N(\mathbf{x}, \mathbf{y}) \\ \hat{\pi}(\mathbf{x}) &= \lim_{N\to\infty} \hat{\pi}_N(\mathbf{x}) \end{aligned}$$

for $\mathbf{x}, \mathbf{y} \in M$. The above limits exist and can be characterized. Namely, Q is a non-negative primitive matrix, i.e., there exists an integer $j > 0$ such that all entries of Q^j are strictly positive. Therefore by [12, Theorem 1.2]

$$\lim_{j\to\infty} \frac{Q^j(\mathbf{x}, \mathbf{y})}{\sigma^j} = \mathbf{w}(\mathbf{x})\mathbf{v}(\mathbf{y}), \qquad (14)$$

where σ is the largest eigenvalue of Q and $\mathbf{w} = (\mathbf{w}(\mathbf{x}) : \mathbf{x} \in M)$ and $\mathbf{v} = (\mathbf{v}(\mathbf{x}) : \mathbf{x} \in M)$ are respectively the corresponding right and left eigenvectors, normalized such that $\mathbf{v}'\mathbf{w} = 1$. It can be verified that Q is irreducible so **w** and **v** are unique up to a multiplicative constant and in turn the right hand side of equality (14) is well-defined. By substituting (14) in (7) and (8) we obtain respectively

$$\begin{aligned} P(\mathbf{x}, \mathbf{y}) &= \frac{Q(\mathbf{x}, \mathbf{y})\mathbf{w}(\mathbf{y})}{\sigma \mathbf{w}(\mathbf{x})} \\ \hat{\pi}(\mathbf{x}) &= G_o \frac{\lambda_1^{x_1}}{x_1!}\frac{\lambda_2^{x_2}}{x_2!}\mathbf{w}(\mathbf{x})^2 \end{aligned}$$

for some normalizing constant G_o.

The blocking probability of type 1 requests arriving at a typical node can now be approximated by

$$B_1^{(0)} = 1 - \sum_{\mathbf{x}} \hat{\pi}(\mathbf{x}) \left(1 - \sum_{\mathbf{y}: x_1+x_2+y_1+y_2=C} P(\mathbf{x}, \mathbf{y}) \right)^2,$$

which is the probability that a random vector $(\mathbf{x}_-, \mathbf{x}_o, \mathbf{x}_+) \in M^3$ with distribution $P(\mathbf{x}_o, \mathbf{x}_-)\hat{\pi}(\mathbf{x}_o)P(\mathbf{x}_o, \mathbf{x}_+)$ satisfies $x_{-,1} + x_{-,2} + x_{o,1} + x_{o,2} = C$ or $x_{+,1} + x_{+,2} + x_{o,1} + x_{o,2} = C$. Analogously, for type 2 requests

$$B_2^{(0)} = 1 - \sum_{\mathbf{x}} \hat{\pi}(\mathbf{x}) \left(1 - \sum_{\mathbf{y}: x_2+y_2=T} P(\mathbf{x}, \mathbf{y}) \right)^2.$$

The revenue rate *per node*, $\frac{W}{2N+1}$, is then approximately

$$r_1\lambda_1(1 - B_1^{(0)}) + r_2\lambda_2(1 - B_2^{(0)}). \qquad (15)$$

This approximation is asymptotically exact as $N \to \infty$.

A similar procedure can be applied to the reservation policy. In this case, the limits $\hat{\pi}(x) = \lim_{N\to\infty} \hat{\pi}_N(x)$ and $P(x, y) =$

$\lim_{N\to\infty} P_N(x, y)$ are given by

$$\hat{\pi}(x) = \begin{cases} G_* \frac{(\lambda_1+\lambda_2)^x}{x!}\mathbf{w}(x)^2 & 0 \le x < R \\ G_* \frac{(\lambda_1+\lambda_2)^T \lambda_1^{(x-T)}}{x!}\mathbf{w}(x)^2 & R \le x \le C, \end{cases}$$

where G_* is a constant, and

$$P(x, y) = \frac{Q(x, y)\mathbf{w}(y)}{\sigma \mathbf{w}(x)}.$$

Here, as before, **w** and **v** are properly normalized right and left eigenvectors corresponding to the largest eigenvalue σ of the matrix Q defined in Theorem 3.2. In turn,

$$B_1^{(0)} = 1 - \sum_{x=0}^{C} \hat{\pi}(x)\left(1 - \sum_{y:x+y=C} P(x, y) \right)^2,$$

$$B_2^{(0)} = 1 - \sum_{x=0}^{R-1} \hat{\pi}(x)\left(1 - \sum_{y:x+y=C} P(x, y) \right)^2$$

and rate of revenue per node is approximated as in (15).

5. Revenue Comparison

In this section, we compute the highest revenue per node under each policy (over possible policy parameters) for the case $N = \infty$. For example, for the threshold policy, the revenue values are computed for all possible threshold parameters and the maximum value is picked. We refer to such value as W_{opt}. Figure 1 depicts W_{opt} under each policy for different values of r_2 with values on both axes normalized by r_1. The results are obtained for $C = 10$ channels and $\lambda_1 = \lambda_2 = 5.0$ requests per unit time. Note here that the displayed revenue values for the complete sharing policy are restricted to the range $0.6 \le \frac{r_2}{r_1} \le 1.0$ as those achieved below this range are fairly low. The figure shows that for $r_2 << r_1$, admitting type 2 requests leads to a loss in revenue and hence to negative profitability. In this case all the policies except the complete sharing policy have mechanisms to block type 2 traffic and thus perform better. For $r_2 > 0.4r_1$ the reservation policy starts performing strictly better than the other policies. For $r_2 > 0.6r_1$ the threshold policy outperforms the complete partitioning policy, and the revenue gap widens as r_2 increases since type 2 traffic becomes more rewarding and statistical multiplexing yields strictly better results than the separation of the two traffic streams. Finally, if $r_2 = r_1$, the maximum revenue under the threshold policy is achieved by letting $T = C$, in which case the policy becomes complete sharing and hence the two policies yield the same revenue.

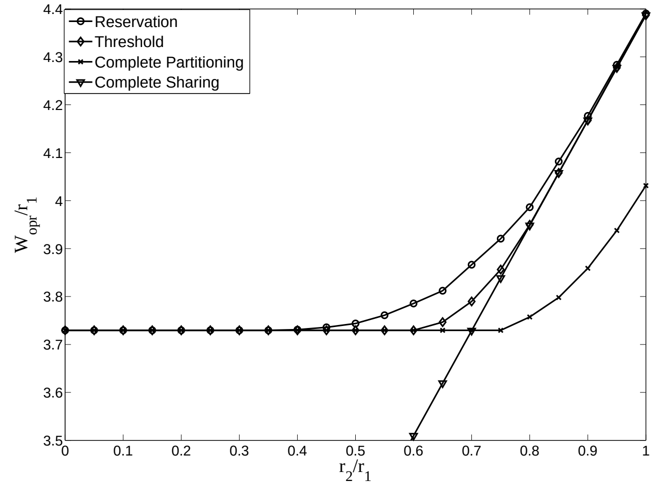

Figure 1: Optimal revenue rates per node for an infinitely extended topology with $C = 10$ and $\lambda_1 = \lambda_2 = 5.0$.

6. Generalizations

The analytical technique of Section 3 can be generalized in a number of directions at the expense of some notational burden. Specifically, analysis of (i) networks with node-dependent arrival rates, (ii) threshold and reservation policies with node-dependent parameters, and (iii) three or more request types each subject to different reservation parameters can be carried out by redefining matrices ϕ and Q in a way to mimic the general procedure followed here. The analysis can also be generalized to arbitrary acyclic (tree) topologies. This extension differs from the program presented here not in the guessing of matrices Q and π, but in the mechanism of passage to P_j and $\hat{\pi}$. Although it is rather rudimentary given the present analysis, the extension requires proper notation and it is omitted for space limitations.

References

[1] B. Miller, ”A queueing reward system with several customer classes,” *Management Sci.*, vol. 16, pp. 234-245, 1967.

[2] S. Jordan and P. Varaiya, “Control of multiple service, multiple resource communication networks,” *IEEE Transactions on Communications*, vol. 42, no. 11, pp. 2979–2988, Nov. 1994.

[3] F. Spitzer, “Random fields and interacting particle systems,” *Mathematical Association of America*, 1971.

[4] F. P. Kelly, ”Stochastic models of computer communication systems,” *Journal of the Royal Statistical Society*, vol. 47, pp. 379–395, 1985.

[5] S. Zachary and I. Zeidins, ”Loss networks and Markov random fields,” *Journal of Applied Probability*, vol. 36, pp. 403–414, 1999.

[6] M. Alanyali, ”Blocking bounds for random channel selection on tree topologies,” *IEEE Transactions on Information Theory*, vol. 53, no. 12, pp. 4750–4755, 2007.

[7] F. P. Kelly, Routing and Capacity Allocation in Networks with Trunk Reservation, *Math. of Oper. Res.*, vol. 15, no. 4, pp. 771–793, 1990.

[8] M. Bebbington, P. Pollett and I. Zeidins, ”Two-link approximation schemes for loss networks with linear structure and trunk reservation,” *Telecommunication Systems* vol. 19, no. 2, pp. 187–207, 2002.

[9] P.J. Hunt and C.N. Laws, ”Asymptotically optimal loss network control,” *Math. of Oper. Res.*, vol. 18, pp. 880–900, 1991.

[10] P.J. Hunt and C.N. Laws, ”Optimization via trunk reservation in single resource loss systems under heavy traffic,” *The Annals of Applied Probability*, vol. 7, pp. 1058–1079, 1997.

[11] M. Luby, and E. Vigoda, “Approximately Counting Up to Four,” *Proc. of 29th ACM Symp. on Theory of Computing*, pp. 682-687, 1997.

[12] E. Senata *Non-negative matrices and Markov chains*, Springer-Verlag, 1981.