State Dependent Routing for Multirate Loss Networks

Ren-Hung Hwang, James F. Kurose and Don Towsley
Dept. of Computer Science
University of Massachusetts
Amherst, MA 01003

Abstract

Markov decision theory has been successfully used in adaptive routing in both traditional single-rate circuit-switched networks and multirate ISDN networks. While the computational complexity required in the case of single rate circuit-switched networks is low, it becomes extremely high in multirate ISDN networks. In this paper, we propose an approach towards adaptive routing in multirate networks using a Markov decision theoretic framework which maintains low computational complexity while still providing quite accurate routing information. In this approach, each link is modeled as a birth-death process to reduce the state space size and a policy iteration method from Markov decision theory is iteratively applied to achieve better network performance. Our results show that routing algorithms based on this approach yields better performance than least-load path routing (LLP) without incurring any significant increase in computational complexity.

1 Introduction

The problem of optimal adaptive routing of telephone traffic has been of considerable interest since the early 80's [1]. Due to advances in switching technologies, state-dependent routing schemes have become feasible (e.g., DNHR [2], DCR [3]). A number of these routing schemes have been developed based on Markov decision processes (MDP's) [1,4]. The MDP approach has been shown to produce a routing policy that is more general and efficient than DNHR [1]. The policies derived from the MDP approach require the evaluation of different cost functions. In one case, Krishnan and Ott in [1] have shown that the cost function can be estimated efficiently. More recently, Dziong *et. al.* [4] introduced the maximum revenue criterion and link shadow prices (first introduced by Kelly [5]) into this approach and demonstrated even further improvements.

Implementing adaptive routing schemes in Integrated Service Digital Networks (ISDN) is extremely difficult because of the complexity of the network. In [6], Dziong and Mason first extended their previous work on traditional circuit-switched networks to multirate circuit-switched networks with calls with heterogeneous bandwidth requirements. Because of the enormous complexity of the exact model, a statistical link independence assumption was made in order to achieve an implementable routing scheme. However, since the size of the link state space grows very quickly as the number of classes of traffic increases, the computational complexity involved in their routing scheme is still tremendous. Thus, the goal of this present study is to investigate a new state-dependent routing scheme for multirate networks based on a simplified link model. In [7], a rough approximation scheme was proposed to achieve this goal. In this paper, a more accurate approximation is proposed without additional computational costs. We also show that this scheme can yield competitive performance as compared to the schemes proposed in [6,8].

We adopt two important approaches in our scheme in order to avoid the problem of enormous computational complexity. We first model each link as a birth-death process by applying recent results from Chung and Ross [9]. We then form a Markov decision process based on this birth-death process, using a similar formulation as in the Adaptive State-Dependent Routing scheme (ASDR) in [10]. One difference, however, is that the concept of call loss in ASDR has been extended to revenue loss in order to consider the case of multiple classes of traffic.

Four routing policies are studied in this paper. The first three policies all apply Markov decision theory, using the simplified birth-death link model. The first policy is simply ASDR. The second policy is from the study in [6] in which the cost of rejecting a multi-link call is divided evenly among the links on the path. The third policy we study is similar to the modified MDPD policy proposed in [8], in which the path cost is taken into account in both the average link shadow price and the state-dependent link shadow price. The last policy we study is the least-loaded path routing scheme (LLP) which has been shown to be very effective in past studies (e.g., [2]). Our simulation results show that the ASDR-like policy requires the least computational complexity and on-line measurements and yields the best performance for fully connected, well-dimensioned topologies. We also compare routing policies based on our approximated link shadow prices with the results in [8] and find that they are very competitive.

The remainder of this paper is structured as follows. In section 2, we first summarize the formulation of the routing problem for multirate networks as a Markov decision process. We then show how the Markov decision process can be decomposed into a set of separable link processes using the link independence assumption. In order to further reduce the size of the state space, a simplified link model for links with multiple classes of traffic is discussed in section 3. At the end of section 3, we describe how to form a Markov decision process based on the simplified link model and derive a set of simple expressions for estimating the expected revenue loss for adding a call at each state. The four routing policies we study are described in section 4. Our simulation results are presented in section 5. Section 6 summarizes our results.

2 The MDP Approach Towards Routing

In this section, we briefly describe the formulation of the routing problem in multirate networks as a continuous-time Markov decision process [6]. We then show how the resulting Markov decision process can be decomposed into a set of link processes based on a link independence assumption.

2.1 Network Model

Let us consider a network consisting of a set of nodes, N, a set of links (trunk groups), L, and a set of all possible Origin-Destination (O-D) pairs, W. Associated with each O-D pair w, $w \in W$, is a set of possible paths. The network handles K classes of traffic labeled $k = 1, ..., K$. Assume a class k call requires b_k units of bandwidth and arrives to O-D pair w according to a Poisson process with rate λ_k^w. (Without loss of generality, we assume $1 = b_1 \leq b_2 \leq ... \leq b_{K-1} \leq b_K$.) Let $\lambda_k = \sum_{w \in W} \lambda_k^w$ be the total exogenous arrival rate of class k calls to the network. The

This work is supported in part by the Office of Naval Research under contract N00014-90-J-1293 and an equipment grant from the National Science Foundation CER-DCR-8500332.

call holding time for a class k call is assumed to be exponentially distributed with mean $1/\mu_k$.

When a call arrives to the network, it will be either carried on a path or lost, depending on the routing policy being used. For each class k call carried on the network, we assume that it brings $r_k = b_k/\mu_k$ units of revenue to the network. The network performance metric in which we are interested is measured by *reward loss* [6,8], which is defined by:

$$reward\ loss = \frac{\sum_{k=1}^{K} r_k \lambda_k B_k}{\sum_{k=1}^{K} r_k \lambda_k}$$

where B_k is the blocking probability of class k traffic.

The routing policies in which we are interested are state-dependent policies in that they use network state information to make routing decisions. Let π be an arbitrary routing policy. When a call arrives to the network, policy π can either route the call on the most "efficient" path (based on some computation using current state information) or reject it. The objective of a routing scheme is to minimize the reward loss of the network or, equivalently, to minimize the expected revenue loss of the network.

2.2 Problem Formulation

In this section, the routing problem in a multirate network is formulated as a Markov decision process. First, let us describe the state of the network by a $K \times J$ matrix $\mathbf{x} = [x_{k,j}]$ where $x_{k,j}$ denotes the number of class k calls carried on path j, where J is the total number of distinct paths in the network.

We then define the state transition probability matrix for a given policy, π. Let δ_k^j be a $K \times J$ matrix with 1 in the position of row k, column j and zeros in all other positions. Let $\Delta_k^w(\mathbf{x}, \pi)$ be either a matrix equivalent to δ_k^j if policy π will route an incoming class k call of O-D pair w on path j or a zero matrix if policy π will reject an incoming class k call. The state transitions for policy π in state $\mathbf{x}$ are:

- $\mathbf{x} \to \mathbf{x} + \Delta_k^w(\mathbf{x}, \pi)$ with rate λ_k^w: this transition corresponds to an arrival of class k call of O-D pair w.
- $\mathbf{x} \to \mathbf{x} - \delta_k^j$ with rate $x_{k,j}\mu_k$: this transition corresponds to a departure of class k call on path j.

Now let us define the cost of the Markov process defined by policy π. Let $q(\mathbf{x}, \pi)$ denote the revenue loss rate while the process is in state $\mathbf{x}$ under policy π. Here, $q(\mathbf{x}, \pi)$ is given by:

$$q(\mathbf{x}, \pi) = \sum_{k \in S(\mathbf{x}, \pi)} r_k \lambda_k$$

where $S(\mathbf{x}, \pi)$ is the set of traffic classes that will be blocked under policy π given that the network is in state $\mathbf{x}$.

Let $V_{\mathbf{X}}(t, \pi)$ be the expected revenue loss of the process in time t with $\mathbf{x}$ as its initial state. If the process is completely ergodic, $V_{\mathbf{X}}(t, \pi)$ is given by [1,11]:

$$V_{\mathbf{X}}(t, \pi) = g(\pi)t + v(\mathbf{x}, \pi)$$

where $g(\pi)$ is called the loss rate of policy π [1] (also, the "gain" of the policy in [11]) and $v(\mathbf{x}, \pi)$ is the relative value of state $\mathbf{x}$ under policy π. In this paper, we will assume that all processes defined by all possible policies are completely ergodic [6].

From [11], $g(\pi)$, for any arbitrary policy π, can be obtained by solving the following set of equations:

$$\begin{aligned} g(\pi) &= q(\mathbf{x}, \pi) + \sum_k \sum_w \lambda_k^w [v(\mathbf{x} + \Delta_k^w(\mathbf{x}, \pi), \pi) - v(\mathbf{x}, \pi)] \\ &\quad + \sum_k \sum_j x_{k,j}\mu_k [v(\mathbf{x} - \delta_k^j, \pi) - v(\mathbf{x}, \pi)] \quad \forall \mathbf{x}. \end{aligned} \tag{1}$$

An optimal policy, π^*, is the policy such that $V_{\mathbf{X}}(t, \pi^*) = min_\pi\{V_{\mathbf{X}}(t, \pi)\}$ for all $\mathbf{x}$. That is, an optimal policy is the policy that minimizes the expected loss revenue in any time t with any possible initial state. Based on Markov decision theory [11], an optimal policy can be found by a policy iteration procedure based on the following iteration cycle:

- **Value Determination Step:** for a given policy, solve the set of equations defined in (1) by setting one of the relative values to zero.
- **Policy Improvement Step:** find the policy that minimizes the right hand side of equation (1) for each state $\mathbf{x}$ using the relative values obtained from the previous step. This policy is called an improved policy [11].

2.3 Decomposition of the MDP

Clearly, the system state space is too large for a direct solution. Thus, approaches such as those used in [1,4] have proposed a link independence assumption and a consequent decomposition of the path reward (cost) into a set of separable link rewards (costs). Here we adopt the approach used in [1] and extend the definition of $q(\mathbf{x}, \pi)$ of a given policy to the expected revenue that will be lost under a given policy π.

The approach used in [1] is called separable routing. The idea is to restrict the policies to the subset of all policies operating under the assumption that the cost functions associated with adding a call to the network, $v(\mathbf{x} + \delta_k^j, \pi) - v(\mathbf{x}, \pi)$, are *separable*, i.e., $\forall \mathbf{x}$

$$v(\mathbf{x} + \delta_k^j, \pi) - v(\mathbf{x}, \pi) = \sum_{\ell \in j} [v_\ell(\mathbf{y}_\ell + \mathbf{e}_k, \pi) - v_\ell(\mathbf{y}_\ell, \pi)] \tag{2}$$

where $\mathbf{e}_k$ is a vector of size K with a 1 in position k and zeros in all other positions and $\mathbf{y}_\ell$ is the state of link ℓ on path j corresponding to the network state $\mathbf{x}$. That is, given the network state $\mathbf{x}$, each link $\ell \in L$ of the network can be described by a vector $\mathbf{y}_\ell = (y_{\ell,1}, ..., y_{\ell,K})$ where $y_{\ell,k}$ denotes the number of class k calls carried on link ℓ. To simplify equation (2), let

$$p_k^\ell(\mathbf{y}_\ell) = v_\ell(\mathbf{y}_\ell + \mathbf{e}_k, \pi) - v_\ell(\mathbf{y}_\ell, \pi).$$

Equation (2) can then be rewritten as

$$v(\mathbf{x} + \delta_k^j, \pi) - v(\mathbf{x}, \pi) = \sum_{\ell \in j} p_k^\ell(\mathbf{y}_\ell).$$

Here $p_k^\ell(\mathbf{y}_\ell)$ is referred to as the state-dependent link shadow price [6], and can be viewed as the cost of adding a class k call on link ℓ when the link is in state $\mathbf{y}_\ell$.

Given the separability assumption, the path cost required in the policy improvement step of the policy iteration procedure can be obtained by solving for the state-dependent link shadow prices at each individual link on the path. That is, for a given policy, we can form a Markov process based on a single link, rather than on the state. We can thus now focus on a single link, ℓ, in isolation and form a Markov process, for a given policy, to obtain the state-dependent link shadow prices at this link. We assume that [1,4,6,10]:

1. calls from various classes arrive to this link according to independent Poisson processes with rate λ_k^ℓ.
2. a call carried on an n-link path behaves like n independent calls, i.e., a call terminates independently on each of the n single links.
3. call holding times for class k calls are exponentially distributed with mean $1/\mu_k$.

For a given policy π, the rate of revenue loss for state $\mathbf{y}_\ell$, $q(\mathbf{y}_\ell, \pi)$, is given by:

$$q_\ell(\mathbf{y}_\ell, \pi) = \sum_{k \in S_\ell(\mathbf{y}_\ell, \pi)} r_k^\ell \lambda_k^\ell$$

where r_k^ℓ is the expected revenue loss for rejecting a class-k call on this link (which is referred to as *link loss*) and $S_\ell(\mathbf{y}, \pi)$ is the set of traffic classes that will be blocked under policy π given that the link is in state $\mathbf{y}_\ell$.

The relative values, $v_\ell(\mathbf{y}_\ell, \pi)$, of link ℓ are obtained by solving the set of equations in the value determination step:

$$g_\ell(\pi) = q_\ell(\mathbf{y}_\ell, \pi) + \sum_k \lambda_k^\ell [v_\ell(\mathbf{y}_\ell + \mathbf{e}_k(\mathbf{y}_\ell, \pi), \pi) - v_\ell(\mathbf{y}_\ell, \pi)] + \sum_k \sum_j y_{\ell,k} \mu_k [v_\ell(\mathbf{y}_\ell - \mathbf{e}_k, \pi) - v_\ell(\mathbf{y}_\ell, \pi)] \quad \forall \mathbf{y}_\ell \quad (3)$$

where $\mathbf{e}_k(\mathbf{y}_\ell, \pi) = \mathbf{e}_k$ if policy π accepts a class k arrival given that the link is in state $\mathbf{y}_\ell$. Otherwise, $\mathbf{e}_k(\mathbf{y}_\ell, \pi)$ is a zero vector.

3 Link Model Simplification

The goal of simplifying the link model is to be able to derive a set of simple expressions for computing the state-dependent *link shadow prices*. Our approach is to develop an approximation to the link state dynamics that yields a birth-death process. This will allow us, as in [1], to derive a set of simple expressions for computing link shadow prices.

3.1 A Simplified Link Model

By defining the link state as the link occupancy, i.e., the number of busy circuits on the link, Wilkinson [12] showed that the cumulative distribution function (CDF) of the link occupancy of a link with non-Poisson arrivals can be very accurately approximated by a negative binomial distribution when the link blocking probability is higher than 0.001. Recently, Chung and Ross [9] have shown that such an approximation can also be applied to multirate loss networks, and that for the purposes of computing blocking probabilities, the approximated link steady state distribution (a negative binomial distribution) can be approximated well by the steady state distribution of a birth-death process. Our approach then is to form a Markov decision process based on this birth-death process. In the following, we summarize how a link with multiple classes of traffic can be approximated as a birth-death process using the technique in [9].

Let us describe the state of a link by the link occupancy, n. Given the link independence assumption of the previous section, the CDF of the link occupancy of each network link will be approximated by the following birth-death process. Let $\rho_k^\ell = \lambda_k^\ell / \mu_k$ be the offered load for class-k calls to this link and μ_1 be normalized to 1. Then the birth-death process has state space 0,...,C where C is the capacity of the link. When in state n, the process has a death rate of n and a birth rate of $\xi^2/\sigma^2 + n(1 - \xi/\sigma^2)$, where $\xi = \sum_{k=1}^K b_k \rho_k$ and $\sigma^2 = \sum_{k=1}^K b_k^2 \rho_k$. The numerical results in [9] show that the blocking probability of the link can be accurately approximated by the steady state distributions of this birth-death process under various traffic conditions, especially under heavy traffic conditions.

3.2 Applying Markov Decision Theory to the Simplified Link Model

In this section, we derive a set of simple expressions for computing the "approximate" state-dependent link shadow prices by forming a Markov process with cost using the simplified (birth-death) link model. Let us focus on a single link, ℓ. Let $\eta_k = r_k^\ell \lambda_k^\ell$. For convenience, we will omit the subscript (or superscript) ℓ in our notation. The set of equations in the value determination step (i.e., equation (3)) is given by:

$$\begin{aligned}
g &= 0 + \bar{\lambda}_0 (v(1) - v(0)) \\
g &= 0 + \bar{\lambda}_i (v(i+1) - v(i)) - i(v(i) - v(i-1)) \\
&\qquad 0 < i \le C - b_K \\
g &= \sum_{j=k}^K \alpha_j \eta_j + \bar{\lambda}_i (v(i+1) - v(i)) - i(v(i) - v(i-1)) \\
&\qquad C - b_k < i \le C - b_{k-1}, \ 2 \le k \le K \\
g &= \sum_{j=1}^K (1 - \alpha j) \eta_j - C(v(C) - v(C-1)).
\end{aligned}$$

The state-dependent link shadow prices, $p_k(i)$, are then computed as

$$p_k(i) = (v(i + b_k) - v(i)) / \mu_k$$

where $v(i) - v(i-1)$ is obtained by solving the above $C + 1$ equations.

$$\begin{aligned}
v(C) - v(C-1) &= \frac{\sum_{j=1}^K \eta_j}{\bar{\lambda}_{C-1}} \frac{E(\bar{\lambda}, C)}{E(\bar{\lambda}, C-1)} - \\
&\quad \frac{\sum_{j=2}^K \alpha_j \eta_j}{\bar{\lambda}_{C-1}} \frac{E(\bar{\lambda}, C)}{E(\bar{\lambda}, C-1)} (1 + E(\bar{\lambda}, C-1)) \\
v(i) - v(i-1) &= \frac{\mathbf{g}}{\bar{\lambda}_{i-1} E(\bar{\lambda}, i-1)} - \frac{\sum_{j=k}^K \alpha_j \eta j}{\bar{\lambda}_{i-1}}, \\
&\quad C - b_k + 2 \le i < C - b_{k-1} + 2, \ 2 \le k \le K \\
v(i) - v(i-1) &= \frac{\mathbf{g}}{\bar{\lambda}_{i-1} E(\bar{\lambda}, i-1)}, \\
&\quad 1 \le i < C - b_K + 2
\end{aligned}$$

where

$$E(\bar{\lambda}, i) = \frac{\frac{1}{i!} \prod_{j=0}^{i-1} \bar{\lambda}_j}{\sum_{n=0}^{i} \frac{1}{n!} \prod_{j=0}^{n-1} \bar{\lambda}_j}$$

and

$$\mathbf{g} = \sum_{j=1}^K (1 - \alpha_j) \eta_j - C(v(C) - v(C-1)).$$

In the preceding equations, we have introduced a class-dependent parameter, α_k. Ideally, we would like to set $\alpha_k = 0, \ \forall k$, since the approximation has already implicitly modelled all calls as requiring a single circuit. This can be seen by inspecting the above formulation for the case of a single class of traffic in which each call requires more than one circuit at a time. However, there are several reasons for setting $\alpha_k > 0$ for $k > 1$ (α_1 is always 0). First, when b/C is not small enough (where $b = max_k(b_k)$), the negative binomial distribution does not match the CDF of the link occupancy very well. Second, when mixing two or more classes of traffic, the approximate model cannot capture the fact that in the exact model there will be no class k arrivals when the link occupancy is less than b_k. We have observed that when b/C is sufficiently small, α_k can be set to 0 for all k. A complete discussion of how values of α_k's are chosen is deferred to Section 5.

4 Routing Policies

Four routing policies are studied in this paper. The first three policies are adaptive state-dependent routing algorithms based on Markov decision theory and use the approximate state-dependent link shadow prices we obtained in section 3. They differ in how link rewards and path costs are assigned. These three routing policies share the following algorithm:

1. Use direct, non-alternate routing as the initial policy. Assign link-loads $\{\lambda_k^\ell\}$ (where λ_k^ℓ is the offered load of class-k call on link ℓ) accordingly.
2. At regular intervals, e.g., ten maximum mean holding times [8], use link-measurements (which may depend on the policies) to re-estimate the link-offered loads.
3. Use the new estimated link-offered loads λ_k^ℓ from step (2) to compute new values of the link shadow prices. The routing policy then routes a new call according to the new set of link shadow prices. The method by which link shadow prices are used to select the routing path, or reject the call, is policy dependent.

4.1 The priority ASDR policy

Our ASDR scheme is a modified adaptive version of the SDR scheme proposed in [10]. The following is a detailed description of priority ASDR.

Call classification: In priority ASDR, calls are classified into K classes with a class k call requiring b_k circuits at a time and $1 = b_1 < b_2 < ... < b_{K-1} < b_K$. Note that as in the original ASDR, we do not distinguish between direct calls and alternate calls.

Link loss: The link loss is set to the expected revenue loss of rejecting a call. i.e., $r_k^\ell = r_k$.

Link-loads $\{\lambda_k^\ell\}$: For each link, ℓ, the following statistics are measured for each class of calls, $1 \le k \le K$.

- carried load $\tilde{\lambda}_k^\ell$: average number of calls carried per unit of time.
- blocking probability B_k^ℓ: the fraction of time that the link is in "blocking" states. Blocking states for class k are those state with less than b_k free circuits.

The offered load of class k call on link ℓ, λ_k^ℓ, is then computed as

$$\lambda_k^\ell = \frac{\tilde{\lambda}_k^\ell}{1 - B_k^\ell}$$

Path cost: The cost of adding a class k call on a path with links $j_1, ..., j_m$, in the respective states $i_1, ..., i_m$, is given by [1]:

$$path\_cost \quad = \quad \sum_{n=1}^{m} p_k^{j_n}(i_n) \qquad (4)$$

Path selection: The priority ASDR policy always routes a call to the direct link whenever the path cost of the direct link is less than the expected reward of the call. If a class k call is blocked on the direct link, the priority ASDR scheme then finds the path, among all possible paths, that has the minimum path cost. If the minimum path cost equals or exceeds the cost of rejecting the call, r_k, the call is rejected; otherwise, the call is routed on the minimum-cost path.

4.2 Priority MDPD policy and MDPD' policy

In [8], the priority MDPD and MDPD' routing policies are introduced and studied. In this paper, we study simplified versions of these two policies, in which (1) the state-dependent link shadow prices are the approximated values obtained in the previous section, and (2) certain on-line link-measurements are simplified. A detailed description follows.

Call classification: In these two policies, calls are first classified according to their link losses. Calls that have the same bandwidth requirement and call holding times are then aggregated into one class [8].

Link loss: When a call is carried on a multi-link path, it has been suggested in [4] that the reward of the call be divided into link rewards. This idea can also be applied to the case of dividing the revenue loss of rejecting a multi-link call into link losses. In our implementation, the call reward (loss) is divided evenly among the links on the path. When calls with the same bandwidth requirement and holding time but different losses are aggregated into one class, a modified link loss for this class of call is computed as follows.

Let us focus on a single link ℓ. Assume Γ is the set of traffic classes that are to be aggregated. For a given policy, let r_k^ℓ be the link loss of a class k call and λ_k^ℓ be the measured offered load which will be specified later. Then the modified link loss for the aggregated class of traffic, r_Γ^ℓ, is given by [6]:

$$r_\Gamma^\ell = \frac{\sum_{k\in\Gamma} r_k^\ell \lambda_k^\ell}{\sum_{i\in\Gamma} \lambda_k^\ell}$$

Link loads: Link loads are measured similarly as in the priority ASDR. However, in order to be able to decompose the call reward into link rewards, we need to measure, for each class of traffic on each link, the carried load of direct calls and the carried load of all alternate calls.

Path cost: The method of computing path costs in priority MDPD is the same as in ASDR.

Consider now the MDPD' policy. It was demonstrated in [8] that priority should not be always given to direct links. Instead, [8] proposes a method of modifying the path net-gain to improve the performance of the routing scheme without giving priority to direct calls. The idea is to increase the probability of choosing paths with a higher derivative of the average network reward with respect to the rate of calls carried on the path. It was shown in [8] that the first derivative of the average network reward with respect to the rate of the class k calls carried on a path with links $j_1, ..., j_m$ is given by $(r_k - \sum_{n=1}^{m} \bar{p}_k^{j_n})$, where $\bar{p}_k^{j_n}$ is the average link shadow price measured on-line from the previous policy iteration cycle. To increase the probability of choosing paths with a higher derivative, the path cost, equation (4), is modified as follows [8]:

$$path\_cost' = \sum_{n=1}^{m} [(1-\alpha) p_k^{j_n}(i_n) + \alpha \bar{p}_k^{j_n}]$$

Path selection: Priority MDPD uses the same method of path selection as priority ASDR. MDPD' uses a similar algorithm except no priority is given to direct links.

4.3 The LLP policy

The last policy we studied is least-loaded path (LLP) routing. LLP has been shown to be very efficient in the past and is most efficient with an appropriate trunk reservation policy. Under LLP with a trunk reservation policy, a call is first offered to the direct link. If it is blocked, then the path with the maximum number of path circuits is tried, where the path circuit is defined as the minimum number of free circuits over the path's links. The number of free circuits of a link is defined as $C_s - x_s - t_s$ where C_s is link capacity, x_s is the number of busy circuits, and t_s is the trunk reservation level. The trunk reservation level on each link has been tuned to yield the best performance in our simulation results.

5 Numerical Results

5.1 Comparing the approximate link shadow prices with exact results

We first examine the accuracy of our approximate shadow cost model for a link with more than one class of traffic. We first examine the case where all classes of traffic have the same mean call holding time. By setting $\alpha_2 = 0.1$, as we can see in Figure 1, the approximate state-dependent link shadow prices still match the exact values very well. Notice that a state in our simplified model corresponds to a set of states in the exact model. In Figure 1, the minimum and the maximum of the state-dependent link shadow prices of the corresponding set of states in the exact model are plotted.

The need for class-dependent parameters, α_k's, in computing link shadow prices can be seen in Figure 1(a). As we can see, the cost of adding a class 1 call is more expensive when the link state is 95 (i.e., $C - b_2$) than when the link state is 96. Intuitively, this is because when the link state is $C - b_2$, the cost for accepting a class 1 call includes both blocking future class 1 calls as well as class 2 calls. However, the cost for accepting a class 1 call when the link state is $C - b_2 + 1$ includes only blocking future class 1 calls because a class 2 call will be blocked regardless of whether a class 1 call is accepted at this state. The simplified link model cannot capture such a phenomenon because all classes of calls have been modeled as requiring a single circuit. The introduction of α_k's is to compensate for this deficiency. Note that this phenomenon occurs only in link state i, $C - b_2 \le i \le C$. Equivalently, the setting of α_k's also only affects the shadow prices for those states. From our experiments, we have found that the choices of α_i's are very sensitive to the number of classes present in the network and are less sensitive to the mean or variance of the traffic and the approximated link shadow prices match the exact values very well under various configurations by setting $\alpha_2 = 0.1$ in the case of $K = 2$ and $\alpha_2 = 0.1, \alpha_3 = 0.15$ in the case of $K = 3$.

In applying Markov decision theory to the approximate model, we have implicitly made the assumption that the exact state-dependent link shadow prices of states in the exact model that result in the same overall link occupancy should be roughly the same. This is validated in our numerical results for the case where all classes of calls have the same mean call holding time and the ratio of b/C is small enough.

When we extend the approximation model to the general case in which each class of traffic has a different mean call holding time, we observe that the implicit assumption we have made becomes questionable. That is, for a given state in the simplified model, the link shadow prices of the corresponding states in the exact model can be quite different. This is because when the mean holding time of a wide-band call is much longer than that of a narrow-band call, the expected revenue loss (i.e., link shadow price) of a state with a wide band call will be much higher than the loss of a state with equivalent number of narrow band calls. However, we are still interested in investigating how a routing policy will perform even when the shadow price information is inaccurate.

5.2 Comparison of four routing policies

We now study the performance of the four routing policies on a nonsymmetric 7-node fully connected network under various traffic conditions (Table 1). The performance of the four previously discussed routing policies is shown in Table 2. (95% confidence intervals for simulation results are also indicated in the table.)

Let us first consider the case in which the link shadow prices obtained from our approximate model are quite accurate (configurations 1 and 2). We observe that all adaptive state-dependent routing algorithms perform better than LLP under various traffic loads. Priority ASDR and priority MDPD have very competitive performance and perform marginally better than MDPD'. All of these results confirm the results presented in [6,8].

We then study the case in which the approximate link shadow prices are not very accurate (configurations 3 and 4). We observe that even when inaccurate link shadow prices are used, the adaptive routing schemes still outperform the LLP scheme. As in the previous case, the priority ASDR and the priority MDPD schemes still have very competitive performance. The only exception is in the overload condition (configuration 4), where MDPD' performs equivalently to ASDR and MDPD. Thus we conclude that our simplified link model provides a very simple way of estimating state-dependent link shadow prices and that priority ASDR is a simple, but effective, routing policy for fully connected, well-dimensioned networks.

5.3 Routing policies in ill-dimensioned networks

In [8] Dziong and Mason also studied an ill-dimensioned network, W8N (Table 3), in order to demonstrate that priority MDPD and LLP can perform poorly while the modified MDPD scheme, MDPD', still yields excellent performance. In this section, we study the performance of our priority MDPD, LLP and modified MDPD schemes on the same network with the same traffic loads. The purpose of this study is to demonstrate how performance degrades when less accurate link shadow prices are used. Table 4 compares the performance of our routing policies with Dziong and Mason's results. We observe that the poor performance of priority MDPD and LLP are also confirmed in our study. It is very interesting to observe that our priority MDPD yields even better performance than the one presented in [8]. On the other hand, we do see some performance degradation when we compare the two modified MDPD schemes. Fortunately, such degradation is not very significant. Considering the decrease in complexity, we thus believe that this makes our approach towards computing approximate link shadow prices quite useful. Furthermore, as the link capacity increases, the approximated link shadow prices will becomes more accurate and we would expect the degradation to be reduced.

6 Summary

In this paper, we have described a simplified link model which can be used to derive a set of simple expressions to estimate state-dependent link shadow prices. We also have shown that the approximate link shadow prices are very accurate in several nontrivial regimes. Even in the case where the approximate link shadow prices are not very accurate, we have shown that the degradation of the performance of routing policies due to these inaccuracies is not very significant.

ACKNOWLEDGEMENT

The authors wish to thank Dr. Z. Dziong for providing the detailed traffic pattern for the W8N network and his valuable comments on an earlier draft of this paper.

References

[1] T. J. Ott and K. R. Krishnan, "State Dependent Routing of Telephone Traffic and the Use of Separable Routing Schemes," in *ITC'11*, September 1985.

[2] G. R. Ash, "Use of a Trunk Status Map for Real-Time DNHR," in *ITC'11*, September 1985.

[3] J. Regnier, P. Blondeau, and W. H. Cameron, "Grade of Service of a Dynamic Call Routing System," in *ITC'10*, June 1983.

[4] Z. Dziong *et al.*, "On Adaptive Call Routing Strategy for Circuit Switched Networks - Maximum Reward Approach," in *ITC'12*, June 1988.

Link / O-D pair	Link Cap.	Configuration 1			Configuration 2		
		λ_1	λ_2	μ_2	λ_1	λ_2	μ_2
1-7	100	75	0.5	1	82.5	0.5	1
8-11	200	38	28	1	40	30	1
12-15	50	40	0	1	43	0	1
16-18	85	30	6	1	34	7	1
19-20	150	30	20	1	35	21	1
21	125	60	9	1	60	10	1
Link / O-D pair	Link Cap.	Configuration 3			Configuration 4		
		λ_1	λ_2	μ_2	λ_1	λ_2	μ_2
1-7	100	75	0.05	0.1	82.5	0.05	0.1
8-11	200	38	2.8	0.1	40	3.0	0.1
12-15	50	40	0	0.1	43	0	0.1
16-18	85	30	0.6	0.1	34	0.7	0.1
19-20	150	30	2.0	0.1	35	2.1	0.1
21	125	60	0.9	0.1	60	1.0	0.1

Table 1: Topology and traffic configurations, $b_2 = 5$.

Policy	Confg. 1	Confg. 2	Confg. 3	Confg. 4
LLP	1.50 ± 0.02	4.93 ± 0.02	1.49 ± 0.05	4.86 ± 0.08
ASDR	1.00 ± 0.01	4.57 ± 0.04	0.90 ± 0.03	4.47 ± 0.09
MDPD	1.02 ± 0.01	4.74 ± 0.04	0.87 ± 0.03	4.40 ± 0.04
MDPD'	1.18 ± 0.03	4.84 ± 0.03	1.02 ± 0.08	4.48 ± 0.05

Table 2: Reward losses [%] for different policies and configurations.

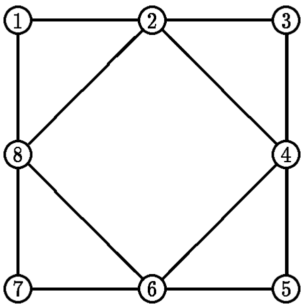

set of O-D pairs	nominal		overload	
	λ_1	λ_2	λ_1	λ_2
A	4.2	0.021	5.1	0.0255
B	56.0	0.224	68.0	0.2720

set	O-D pairs
A	(1,2), (2,3), (3,4), (4,5), (5,6), (6,7), (7,8), (1,8)
B	(2,4), (4,6), (6,8), (2,8), (2,6), (4,8)

Table 3: W8N topology and traffic configurations: Link cap. = 120, $b_2 = 12$, $\mu_2 = 0.1$.

Traffic load	Policies studied		
	pri. MDPD	MDPD'	LLP
nominal	12.17 ± 0.07	0.95 ± 0.07	10.13 ± 0.08
overload	17.74 ± 0.22	5.60 ± 0.12	12.98 ± 0.15
Traffic load	Policies in [8]		
	pri. MDPD	MDPD'	LLP
nominal	25.81 ± 0.37	0.70 ± 0.14	12.81 ± 0.20
overload	33.87 ± 0.25	4.92 ± 0.29	18.43 ± 0.23

Table 4: Reward losses [%] for policies studied in this paper and in [8]

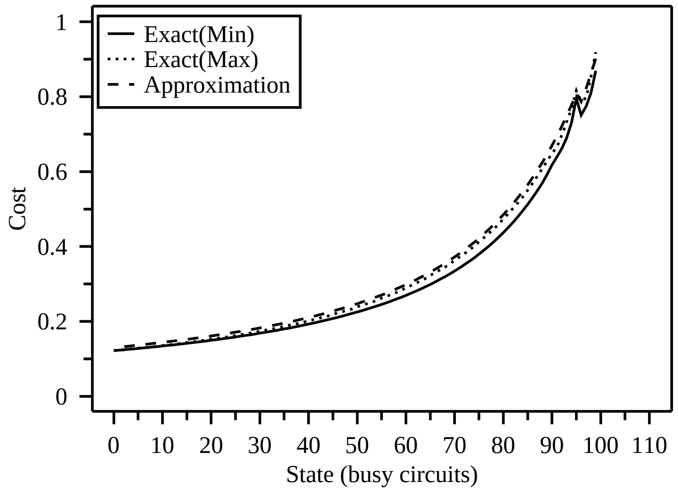

(a) Cost of adding one class 1 call
C=100, λ_1=50, λ_2=10
b_1=1, b_2=5

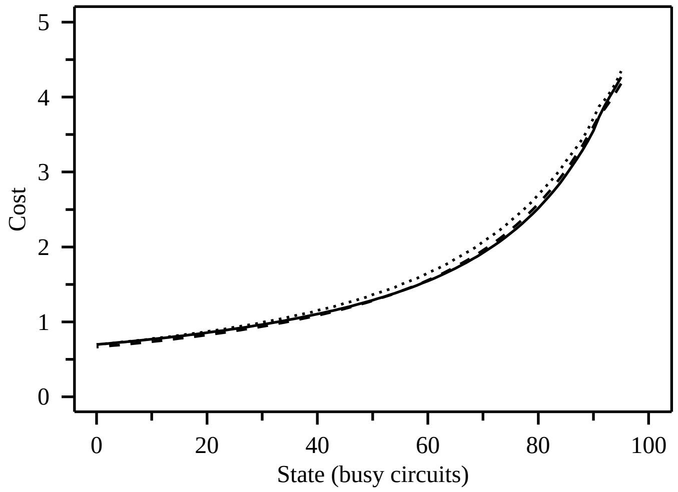

(b) Cost of adding one class 2 call

Figure 1: Applying MDP to approximated model

[5] F. P. Kelly, "Routing in Circuit Switched Networks: Optimization, Shadow Prices and Decentralization," *Advanced Applied Probability*, vol. 20, 1988.

[6] Z. Dziong and L. Mason, "Control of Multi-Service Loss Networks," in *CDC'28*, December 1989.

[7] Z. Dziong, K.-Q. Liao, L. Mason, and N. Tetreault, "Bandwidth Management in ATM Networks," in *ITC'13*, June 1991.

[8] Z. Dziong and L. Mason, "An Analysis of Near Optimal Call Admission and Routing Model for Multi-Service Loss Networks," in *INFOCOM'92*, May 1992.

[9] S.-P. Chung and K. W. Ross, "Reduced Load Approximations for Multirate Loss Networks," *To appear in IEEE Transactions on Communications.*

[10] K. R. Krishnan, "Adaptive State-Dependent Traffic Routing Using On-Line Trunk-Group Measurements," in *ITC'13*, June 1991.

[11] R. A. Howard, *Dynamic Programming and Markov Processes.* John Wiley & Sons, Inc., 1960.

[12] R. I. Wilkinson, "Theories of Toll Traffic Engineering in the USA," *Bell Systems Technical Journal*, vol. 40, pp. 421–514, 1956.