MDP Routing in ATM Networks Using Virtual Path Concept

Ren-Hung Hwang, James F. Kurose, and Don Towsley

Department of Computer Science & Information Engineering
National Chung Cheng University
Chia-Yi, Taiwan, R.O.C.

Department of Computer Science
University of Massachusetts
Amherst, Massachusetts 01003
U.S.A.

Abstract

The Virtual Path (VP) concept has been proposed to simplify traffic control and resource management in future B-ISDN's. In particular, call setup processing can be significantly reduced when resources are reserved on VP's. However, this advantage is offset by a decrease in statistical multiplexing gains of the networks. The focus of this paper is on how to improve bandwidth efficiency through adaptive routing when capacity is reserved on all VP's. We first examine two VP capacity reservation strategies. We then design and evaluate computationally feasible Markov decision process-based routing algorithms and show that the network blocking probability can be significantly reduced by MDP routing.

1 Introduction

The CCITT specifications on future Broadband ISDN propose the Virtual Path (VP) concept as a component of traffic control and resource management in B-ISDN's [1,2]. The use of VP's for traffic control has several advantages. For example, VP's can be used to simplify call admission control, to segregate traffic types which requiring different QOS, to facilitate aggregation with similar performance requirements, to simplify network architecture, and to enhance network services [2]. An additional advantage is the reduction of processing complexity at the switches (nodes) within the network.

The use of VP's also plays an important role in network resource management. By reserving resources on VP's, such as buffer space and bandwidth, processing requirements for call establishment are substantially decreased. This is because a new virtual circuit (VC) [1] can be established by executing an admission control function at the source node of a VP without requiring any call processing at transit nodes along the VP. However, *a priori* reservation of resources on VP's also reduces the statistical multiplexing gains of the network, resulting in an increased probability that an arriving call will be blocked. Thus, there exists a trade-off between reduced control processing requirements and reduced resource utilizations. Throughout this paper, we will use the term "VP capacity" to refer to the resources reserved on VP's.

The purpose of this paper is to demonstrate how to reduce the call blocking probability through Markov Decision Process-based (MDP) routing algorithms when capacity is reserved on VP's. We first explore two VP capacity reservation strategies: a deterministic strategy and a statistical strategy. We then design and evaluate MDP routing algorithms for homogeneous VP networks. (In a homogeneous VP network, all applications are assumed to require identical end-to-end QOS guarantees.) In our numerical results, we show that MDP routing algorithms are able to reduce the network blocking probability significantly.

A related previous work is [3], in which Least Loaded Path-based (LLP) routing algorithms were developed to reduce the network call blocking probability for homogeneous VP networks. Our research differs from this work in the following respects. First, instead of using an LLP approach, we develop MDP-based routing algorithms for VP networks. Second, the study in [3] is limited to the case in which each VP only supports one QOS requirement. In our study, this constraint is relaxed. Finally, in addition to studying a deterministic VP capacity reservation strategy, we also propose a statistical strategy.

Most other earlier research in VP networks has focused on studying the trade-off between control cost, i.e., the amount of processing required at switches, and transmission cost, i.e., the amount of bandwidth required at physical links, for a given network topology and traffic pattern. For example, in [4], Burgin presents a cost-benefit analysis of reserving capacity on VP's in a homogeneous network. Hui *et al.*, in [5], also study the reduction of bandwidth efficiency resulting from reserving capacity on VP's. The study by Ohta *et al.* [6] is rather different from [4] and [5]. Instead of reserving capacity on VP's for fixed length periods, VP capacities are dynamically increased/decreased on demand in units of "blocks". By assuming that every call is homogeneous and that the traffic offered to each VP is identical, the authors study the effects of the block size on the trade-off between control cost and transmission cost.

The remainder of this paper is organized as follows: in section 2, a VP network is described in more detail. Two VP capacity reservation strategies are then described in section 3. In section 4, four MDP-based routing algorithms are proposed. The performance of the four routing algorithms are evaluated and compared to the performances of two baseline algorithms in section 5. Finally, section 6 summaries this paper.

2 Network Model

2.1 A Virtual Path Network

A VP is broadly defined as a labeled path (a set of adjacent links) which may carry more than one virtual circuit. Virtual circuits which belong to the same VP have the same

This research was supported in part by the Office of Naval Research under contract N00014-90-J-1293, the National Science Foundation under grant NCR-9116183.

Ren-Hung Hwang is supported in part by National Science Council of Taiwan, R.O.C., under contract NSC-83-0408-E-194-006-T.

[1] Throughout this paper, the terms "call", "VC", and "connection", are interchangeable.

endpoints. One advantage of introducing VP's into ATM networks is that all virtual circuits in a VP can be transported, processed, and managed in a similar manner.

A logical VP network can be defined by viewing the VP endpoints as nodes and VP's as links. Consider the physical network shown in Figure 1(a), which consists of 5 switches (nodes) and 12 unidirectional physical links. If a VP is assigned to each source-destination pair, as shown in Figure 1(b), an N-node fully connected VP network is formed consisting of N(N-1) unidirectional *logical links* (VP's).

One rationale for using VP's for traffic control is to separate different traffic types in order to prevent them from interacting or interfering with each other [2]. For example, different types of traffic (e.g., voice and data) may have very different traffic characteristics and QOS requirements; the question of how to multiplex two or more diverse traffic classes while providing different QOS requirements at a switch is a very complicated, open problem. However, the problem can be much simplified if different types of traffic are separated by assigning a VP with dedicated resources to each type of traffic. As a consequence of this separation, more than one VP may be established between the same source-destination pair, with each carrying different types of traffic.

Throughout this paper we will assume that *applications* in the VP network require identical *end-to-end* QOS guarantees. Such a VP network is referred to as a homogeneous VP network in [7]. However, unlike [7], we will assume that individual virtual circuits (VC's) in a VP connection may have different QOS requirements, as described in [2]. That is, a virtual path may provide more than one QOS class. The rationale for providing this flexibility is that when a VC is routed on an alternate path (i.e., a non-direct path consisting of two "concatenated" VP's) a more stringent QOS will be required on both of the two individual VP's constituting this alternate path, in order to meet the specified *end-to-end* QOS requirement.

2.2 Traffic Model, QOS Criterion, and Call Admission Scheme

In this section, we describe the traffic model, QOS criterion, and call admission scheme used in our numerical study. We emphasize, however, that the design of our routing policies is independent of the underlying traffic model, QOS criterion, and call admission policy.

The traffic model that we use throughout our numerical examples is a two-state fluid-flow model [8]. This model describes a source which is either in an "on" state, transmitting at its peak rate, or in an "off" state, transmitting at zero bit rate. The durations of the "on" and the "off" states are assumed to be exponentially distributed random variables. When this model is used to describe a packet voice source, the source will stay in the "on" state for an average of 352 msec, transmitting at 32 Kbps, or stay in the "off" state for an average of 650 msec, transmitting at zero bits per second. Based on this two-state fluid-flow model, a traffic source can be described by three parameters: its peak rate, r (bps); utilization, ρ, i.e., the fraction of time the source is at "on" state; and the mean "on" period, b (seconds) [9]. For example, the voice source can be described by the triple $(32000, 0.3513, 0.352)$.

Throughout this paper, we assume that VC's arrive at each source-destination pair according to a Poisson process and the call holding times are assumed to be exponentially distributed random variables with mean of unity.

Different types of traffic may require different QOS guarantees. Throughout our examples, we assume that the QOS criterion is specified by $Q(D, \epsilon)$, which means that no more than a fraction, ϵ, of cells of a virtual circuit should experience more than D seconds of queueing delay. Note that, unlike [3], the value of D does not include the propagation delay. We also assume that the buffer space allocated for each link (or VP) only holds up to D seconds of cells and arriving cells are discarded if the buffer is full.

Recently, a number of call admission schemes have been proposed for high-speed networks (e.g. [9, 10]). We use the admission control scheme proposed by Guérin *et al.* [9], which is based on the fluid-flow approximation model of [8]. In this model, a number of homogeneous sources described by (r, ρ, b) with a QOS requirement $Q(D, \epsilon)$ are multiplexed onto a link of capacity C bps. Based on fluid-flow approximation techniques, the maximum number of sources that can be multiplexed onto the link without violating the QOS constraint is given by

$$N = \lfloor \frac{C}{\hat{c}} \rfloor \qquad (1)$$

where $\hat{c}$ is called the equivalent capacity of a single source and is given by

$$\hat{c} = \frac{\alpha b(1-\rho)r - x + \sqrt{[\alpha b(1-\rho)r - x]^2 + 4x\alpha b\rho(1-\rho)r}}{2\alpha b(1-\rho)} \qquad (2)$$

where $\alpha = \ln(1/\epsilon)$ and $x = D \times C$.

3 Resource Management

In this section, we first describe two VP capacity reservation strategies. We will observe that implementing the second reservation strategy requires establishing multiple classes of QOS on a physical link. Thus, at the end of this section, we also discuss possible solutions for providing multiple classes of QOS on a physical link.

3.1 VP Capacity Reservation

Having noted that call setup processing can be eliminated from the transit nodes when resources are reserved on VP's, it is necessary to examine different strategies for reserving resources on VP's. In the following, we present two VP capacity reservation strategies: a *deterministic* strategy [4, 3] and a *statistical* strategy [11].

The deterministic strategy, advocated by [4, 3], reserves separate link capacity for each VP passing through the link. This approach treats a VP as if it were a physical link. A certain amount of buffer space at the source node and bandwidth at each physical link on this VP are *dedicated* to this VP. Since there is no statistical multiplexing performed among VP's at the link-level, the sum of the reserved VP bandwidths on a link is not permitted to exceed the total capacity of this link. However, statistical cell multiplexing is still performed among virtual circuits within a VP. Therefore, it is possible for the offered peak rate to a VP to exceed its allocated bandwidth over a short period of time. In such a case, cells that cannot be transmitted are buffered at the source node of the VP. Since the instantaneous transmission rate of a VP is limited to its reserved bandwidth, the buffer space required and the queueing delays incurred at transit nodes/links of the VP are very small, and are assumed to be negligible [4, 3].

The statistical strategy, inspired by [11], allows statistical multiplexing of cells from different VP's on a physical link. In this approach, instead of explicitly reserving buffer space and link bandwidth for VP's, each VP is allocated a *dedicated* number of virtual circuits, each with a guaranteed QOS, such that call setup processing can still be done locally at the source node so long as the number of existing virtual circuits within the VP is less than this number. The actual reservation is done by pretending as if each VP has accepted as many as the dedicated number of virtual circuits into the network. The advantage of the statistical approach is that it provides better statistical cell multiplexing gain than the deterministic strategy. However, this advantage is offset by the fact that when some connections are routed over multiple VP's, a more stringent QOS guarantee needs to be provided by each component VP.

Two problems can be identified when implementing the statistical strategy. First, the traffic characteristics of a virtual circuit may change as the traffic travels along the path; these changes must be characterized. Second, we must determine how to assign the end-to-end QOS requirement to each link such that the overall end-to-end QOS can be satisfied. Our approach to these two problems, is to assume that the traffic characteristics of a virtual circuit are unaffected as the traffic travels along the path and to assign an equal amount of the end-to-end QOS of a virtual circuit to each link on the path, the so called equal allocation (EQ) policy [11]. As a result, when a physical link is shared by more than one VP, the QOS requirement of different VP's may differ under this strategy even though we assume that all connections in a homogeneous VP network require the same *end-to-end* QOS guarantee. This issue of how to establish multiple classes of QOS service on a single link is deferred to the next section.

When capacity is reserved on VP's, it may be desirable to dynamically adjust the allocation to improve link bandwidth utilization as well as to adapt to dynamic changes in network traffic flows. In this paper, we assume that the reallocation of VP capacity should be done periodically on a much longer time scale than the interarrival time of successive calls [4]. Furthermore, we assume that the time interval between two VP capacity reallocations is significantly larger than the call setup time. Under this assumption, routing of virtual circuits can be performed as if the topology and the capacity of the VP network were fixed [3].

3.2 Multiple Classes of QOS

The implementation of the statistical VP capacity reservation strategy requires that a physical link accommodate multiple service classes with different QOS requirements. In order to achieve high bandwidth utilization, a sophisticated cell transfer control scheme can be developed. However, it is expensive to implement such a complicated cell control scheme at switches and the resulting QOS is difficult to predict. In the following, we focus on non-priority schemes which are relatively easy to implement.

A simple approach is to uniformly provide the most stringent QOS to all cells [3, 11]. Since all cells are treated equally, no special cell control scheme is required. However, link utilization will be decreased because an unnecessarily stringent QOS is provided to *all* cells, regardless of their actual QOS requirements. We will refer to this approach as the UNIFORM scheme.

A complementary approach to the UNIFORM scheme is to separate cells according to their QOS requirements and allocate dedicated buffer space and bandwidth to each class of QOS. Cells with the same QOS requirement are treated equally and are served by their dedicated bandwidth based on a first-come-first-serve discipline. The advantage of this approach is that cells will receive their actual QOS requirements. However, this advantage is offset by the reduction in the statistical multiplexing gain because of the partitioning of the resources. We will refer to this approach as the SEPARATE scheme.

In order to implement the SEPARATE scheme, we need to determine the amount of buffer space and bandwidth to dedicate to each class of QOS. Let us consider an example where two classes of QOS, say $Q(D_1, \epsilon_1)$ and $Q(D_2, \epsilon_2)$, are to be established on a link of capacity C. Let N_1 and N_2 be the maximum number of virtual circuits of these two classes that can be simultaneously supported without violating the QOS requirements. Assume without loss of generality that $Q(D_1, \epsilon_1)$ is more stringent than $Q(D_2, \epsilon_2)$. Our approach to determine the amount of buffer space and bandwidth to dedicate to each class of QOS is as follows. For a given value of N_1, we first determine the buffer size, B_1, and the amount of bandwidth, C_1, required by the class 1 traffic by solving the following equations:

$$B_1 = D_1 \cdot C_1 \tag{3}$$
$$C_1 = \hat{c}_1 N_1 \tag{4}$$
$$B_1 > 0 \tag{5}$$
$$C_1 > 0 \tag{6}$$

where $\hat{c}_1$ is given by equation (2) using $\alpha = \ln(1/\epsilon_2)$ and $x = B_1$. After a simple calculation, we have

$$C_1 = \frac{(\alpha b(1-\rho) + \rho N_1 D_1) r N_1}{\alpha b(1-\rho) + D_1 N_1}. \tag{7}$$

(Recall that we have homogeneous traffic sources described by (r, ρ, b).)

The maximum number of class 2 calls that can be supported, N_2, is then given by equation (1) using $\alpha = \ln(1/\epsilon_2)$ and $x = D_2 \cdot (C - C_1)$.

Extending this algorithm to a more general case with arbitrary classes of QOS is very simple, and, thus, is omitted.

We can compare and contrast the SEPARATE and UNIFORM scheme via a simple example in which two sets of homogeneous voice sources, characterized by the same two-state fluid-flow model with parameters $(32000, 0.3513, 0.352)$, and differing only in their QOS requirements are presented to a link. In Figure 2, we compare the maximum number of class 1 and class 2 virtual circuits that can be supported on the link under these two schemes. A point (N_1, N_2) on the curve of a cell control scheme in Figure 2 can be interpreted as follows: in order to support N_1 class 1 virtual circuits, N_2 is the maximum number of class 2 virtual circuits that can be supported on the link without violating the QOS requirement under this cell control scheme.

We observe that when the number of virtual circuits of the first class, N_1, is small, the SEPARATE scheme can support more class 2 VC's than the UNIFORM scheme.

On the other hand, when N_1 is large, we see the opposite result. The rationale for such an observation is that, as stated above, when N_1 is small, the capacity required to support N_1 class 1 VC's is also small which implies the reduction in the statistical multiplexing gains due to the separation of resources is also small. Thus, by providing a less stringent QOS to the class 2 VC's, the SEPARATE scheme can support more class 2 VC's. However, as N_1 increases, the reduction in the statistical multiplexing gains due to the separation of resources also increases. Hence, at some point, the UNIFORM scheme becomes better than the SEPARATE scheme.

We adopt a hybrid scheme, referred to as the ADAPTIVE scheme, which is a combination of the UNIFORM scheme and the SEPARATE scheme. Consider the previous case where two classes of QOS are to be established. For a given N_1, the ADAPTIVE scheme will adopt the UNIFORM scheme if it supports more class 2 virtual circuits than the SEPARATE scheme. Otherwise, the ADAPTIVE scheme will adopt the SEPARATE scheme.

4 VP Routing

When resources are reserved on VP's in homogeneous VP networks, routing in VP networks is very similar to routing in traditional single-rate circuit-switched networks. Therefore, in this section, we study how routing algorithms in circuit-switched networks can be applied to VP networks.

Existing adaptive routing algorithms in circuit-switched networks can be classified into two categories: *Least-Load Path*-based (LLP-based) [12] and *Markov Decision Process*-based (MDP-based) [13]. The LLP approach tries to route calls to the "least busy" path, i.e., the path that has the maximum number of free circuits. The MDP approach, on the other hand, formulates the routing problem as a Markov decision process. Since the MDP approach has been shown to be more efficient than the LLP approach in other contexts [13,14], we will focus on MDP-based routing in VP networks.

4.1 MDP Routing in Circuit-Switched Networks

The formulation of the routing problem in traditional circuit-switched networks as a Markov decision process has been studied as early as 1966 [15]. However, solving an exact Markov decision process to obtain an optimal routing policy is computationally infeasible for any realistic network due to the huge state space.

Recently, Krishnan and Ott [13] proposed a computationally feasible, MDP-based routing scheme, referred to as State-Dependent Routing (SDR) scheme. The SDR scheme uses a policy iteration procedure to obtain an improved policy of the initial policy, which is assumed to be a direct routing policy. In direct routing, calls are only allowed to be routed via a direct link. Under the assumption that the arrival process of each O-D pair is an independent Poisson process, the arrival process to each link is, therefore, also an independent Poisson process. To carry out one iteration of the policy iteration procedure, a set of linear equations is solved to obtain a set of *relative values* and the gain of this initial policy. However, this calculation can be made independently for each link because the arrival process to each link is assumed to be statistically independent and each call is routed to the direct link only. By solving the set of value determination operation equations for each link, the cost (measured in terms of the expected number of additional future calls which are blocked as a result of accepting a currently offered call) of adding a call to the link when it already has i calls in progress, which we will refer to as the state-dependent link shadow price at link state i, $p^\ell(i)$, is given by the following formula:

$$p^\ell(i) = \frac{E(\lambda^\ell, C_\ell)}{E(\lambda^\ell, i)}, \tag{8}$$

where $E()$ is the Erlang's B formula, C_ℓ is the link capacity and λ^ℓ is the arrival rate of the link offered traffic. In [13], the arrival rates of the link offered traffic correspond to the traffic assignments of the optimal non-alternate routing policy. Later studies [16] estimate the actual traffic load through on-line measurements during network operation rather than relying on such *a priori* traffic estimation. The state-dependent link shadow prices are then updated periodically.

In the SDR scheme, all calls have the same rewards and the objective of the MDP model is to minimize the number of lost calls per unit of time. In [16], the MDP model is extended to maximize revenues in the case where calls may have different reward parameters. Specifically, when a call is carried on a multilink path, the revenue of the call is divided into individual link revenues. By a similar decomposition of the MDP model into link processes, the state-dependent link shadow price at link state i, $p^\ell(i)$, is given by the following formula:

$$p^\ell(i) = \frac{r^\ell E(\lambda^\ell, C_\ell)}{E(\lambda^\ell, i)}, \tag{9}$$

where r^ℓ is the so-called link-level revenue loss. The value of r^ℓ will generally be different for each routing policy.

4.2 MDP Routing in VP Networks

To perform MDP routing in a VP network, we first define a set of possible paths for each source-destination pair. A candidate path may consist of one or more VP's. Since a VP network is very likely to be densely connected, we restrict all candidate paths to consist of at most two VP's. From a previous study on fully connected networks [14], we have found that priority MDP routing policies, which give priority to direct links, yield better performance than non-priority MDP routing policies. Thus, we consider MDP routing schemes in VP networks which always offer a call to the direct VP first. If a call is blocked on the direct VP, it is then offered to the alternate path with the minimum *path cost*. A call is rejected if the minimum path cost of all alternate paths is greater than the call reward. Before we develop MDP routing algorithms for VP networks, however, several issues need to be addressed.

When a call is to be routed on an alternate path consisting of two VP's, we again encounter the same problems found with the statistical strategy for reserving capacity on VP's, i.e., how to characterize the traffic characteristics of a VC as it travels along the alternate path and how to assign the end-to-end QOS requirement to each link of the alternate path. As in the previous section, we assume that the traffic characteristics of a virtual circuit are unaffected as the traffic travels through a VP. The problem of dividing the end-to-end QOS into local QOS is subtler because it depends on the underlying VP capacity reservation strategy. If a deterministic reservation strategy is used, the equal allocation policy, i.e., assigning an equal amount of the end-to-end QOS of the virtual circuit to each VP on the alternate path, is also adopted. As a consequence, two classes of QOS, one for direct calls and one for alternate calls, need to be established on a VP which contains alternate routed calls. The ADAPTIVE scheme proposed in the previous section is applicable for providing two QOS classes on a VP with dedicated capacity.

On the other hand, if a statistical VP capacity reservation strategy is adopted, the equal allocation policy becomes problematic because it is very difficult to provide two classes of QOS on a VP under this strategy. Therefore, under the statistical reservation strategy, our approach for providing alternate routing is to restrict the paths in the set of possible paths such that an alternate path is considered as "possible" only if the sum of the QOS of the two VP's on this path satisfies the end-to-end QOS requirement.

4.3 MDP Routing Policies

In this section, we describe four MDP routing policies under the two different VP capacity reservation strategies. In general, the following algorithm is shared by these four routing policies.

- **Initial step:**
  Use direct VP routing as the initial policy. Assign link-offered loads, λ^ℓ, accordingly. (λ^ℓ is the total offered link load on link ℓ.)
- **Loop**
  1. At regular time intervals of length Δ_t, use link-measurements to re-estimate the link-offered loads, λ^ℓ.
  2. Use the new estimated link-offered loads, λ^ℓ, from the previous step to compute new values of the link shadow prices. The computation of link shadow prices and the method by which link shadow prices are used to select the routing path, or reject the call, will be policy dependent, as discussed below.

The four MDP routing policies discussed below are quite similar. We first describe the priority MDPD policy under the deterministic reservation strategy in detail. Other policies are described only in the manner in which they differ from this policy. In the following description, we describe how link losses and link-offered loads are estimated, link shadow prices and path cost are computed, and the "best" path is selected under different routing policies.

4.3.1 Routing Policies under Deterministic Reservation Strategy

Policy 1: The deterministic reservation priority MDPD policy (MDPD_D):

This policy simply follows the priority MDPD policy specified in [14]. A detailed description follows.

Link loss: When a call arrives to the network, it will either be carried on a path or lost. We assume that each call carried on the network produces r units of revenue. From another point of view, the network loses r units of revenue for each call that is rejected. As with the MDPD policy in [14], the revenue loss for rejecting a call on an alternate path is divided into link-level revenue losses. In order to simplify the computation of link loss, i.e., r^ℓ in equation (9), we will only study a division rule which divides the revenue loss for rejecting an alternate call evenly among the links (VP's) on the alternate path [14]. The link-level revenue loss, r^ℓ, under this policy is defined as follows [17, 14]:

$$r^\ell = \frac{r\lambda_d^\ell + 0.5r\lambda_a^\ell}{\lambda_d^\ell + \lambda_a^\ell} \qquad (10)$$

where λ_d^ℓ is the estimated offered load on the direct path and λ_a^ℓ is the estimated offered load on all alternate paths. The reason why the cost of rejecting an alternate call is $0.5r$ in equation (10) is that all alternate paths consist of two links, and the revenue lost in rejecting an alternate call is split equally between the two links.

Link loads: The offered link loads, λ_d^ℓ and λ_a^ℓ, are estimated by

$$\lambda_d^\ell = \frac{\tilde{\lambda}_d^\ell}{1 - B^\ell},$$

$$\lambda_a^\ell = \frac{\tilde{\lambda}_a^\ell}{1 - B^\ell},$$

where $\tilde{\lambda}_d^\ell$ is the average number of direct calls carried per unit of mean call holding time, $\tilde{\lambda}_a^\ell$ is the average number of alternate calls carried per unit of time, and B^ℓ is the measured link blocking probability.

The total offered link load on link ℓ, λ^ℓ, is the sum of offered loads on the direct and alternate paths, i.e.,

$$\lambda^\ell = \lambda_d^\ell + \lambda_a^\ell.$$

Link shadow prices: Computing the link shadow prices in VP networks under the deterministic reservation strategy is much more complicated than for traditional circuit-switched networks. Thus approximation techniques are required. The complication arises from the fact that when a VC is routed on an alternate path, it requires a more stringent QOS at each VP on this alternate path. Thus, unlike in traditional circuit-switched networks, the link shadow prices in VP networks for adding a direct call and an alternate call are not equivalent for a given link state.

The main idea of our approximation technique is to measure the capacity required by an alternate call in units of the capacity required by a direct call. By doing so, a VP can then be viewed as a link with a fixed number of "circuits". Given that a VP has a fixed number of circuits, equation (9) can be applied to compute the link shadow prices for both direct calls and alternate calls. In the following, we first show how to measure the capacity required by an alternate call in units of the capacity required by a direct call. We then show how link shadow prices for both direct calls and alternate calls are estimated.

When an alternate call is added to a VP, the maximum number of direct calls that can be added to this VP may decrease by more than one. Thus, our approach is to define the capacity required by an alternate call as the difference between the number of the maximum number of direct VC's that can be supported on the VP before and after adding the alternate call. We refer to this capacity as the "virtual capacity" required by the alternate call.

In order to compute the virtual capacity required by an alternate call, we introduce some additional notation and functions. Let the link state be described by (i,j) where i is the number of existing direct VC's and j is the number of existing alternate VC's. Assume link ℓ can support up to N^ℓ direct VC's without violating the QOS requirement. On the other hand, when only alternate VC's are carried on the link, assume that N_a^ℓ is the maximum number of alternate VC's that can be supported. Since alternate calls require a more stringent QOS requirement, N_a^ℓ is no greater than N^ℓ, i.e., $N_a^\ell \leq N^\ell$. Let $\mathcal{N}^\ell(j)$ be the maximum number of direct VC's that can be supported on link ℓ when there exist j alternate virtual circuits on the link. Therefore, the quantity $\mathcal{N}^\ell(j) - \mathcal{N}^\ell(j+1)$ is the virtual capacity required to carry an alternate call when there already exists j alternate VC's on the link.

Since the capacity required to carry an alternate call is now measured in units of the capacity required by a direct call, we now can use equation (9) to compute link shadow prices as follows. Define $\mathcal{V}^\ell : \{0,...,N^\ell\} \times \{0,...,N_a^\ell\} \rightarrow \{0,...,N^\ell\}$ where $\mathcal{V}^\ell(i,j)$ is called the transformed link state of link ℓ when there are i direct VC's and j alternate VC's on link ℓ. The function $\mathcal{V}^\ell(i,j)$ is defined by

$$\mathcal{V}^\ell(i,j) = \begin{cases} i + \mathcal{N}^\ell(0) - \mathcal{N}^\ell(j), & \text{if } i \leq \mathcal{N}^\ell(j) \\ \text{undefined}, & \text{otherwise,} \end{cases}$$

where $\mathcal{N}^\ell(0) - \mathcal{N}^\ell(j)$ is the capacity, in units of the capacity required by a direct call, required to support j alternate VC's. Thus $\mathcal{V}^\ell(i,j)$ gives us the total capacity required to support i direct VC's and j alternate VC's.

The state dependent link shadow price, $p_d^\ell(i,j)$, for adding a direct call on link ℓ when there exists i direct VC's and j alternate VC's on the link is given by

$$p_d^\ell(i,j) = \frac{r^\ell E(\lambda^\ell, N^\ell)}{E(\lambda^\ell, \mathcal{V}^\ell(i,j))}. \tag{11}$$

Note that this exactly parallels equation (9).

The computation of the approximate link shadow prices for alternate calls needs to take into account the fact that adding an alternate call may decrease the maximum supportable number of direct calls. By the use of virtual capacity, the state-dependent link shadow prices for adding an alternate call, $p_a^\ell(i,j)$, are estimated by

$$p_a^\ell(i,j) = \sum_{v=\mathcal{V}^\ell(i,j)}^{\mathcal{V}^\ell(i,j+1)-1} \frac{r^\ell E(\lambda^\ell, N^\ell)}{E(\lambda^\ell, v)}. \tag{12}$$

The summation in equation (12) corresponds to the virtual capacity required to add an alternate call at link state (i,j).

Path cost: The cost of adding a VC on an alternate path with links ℓ_1, ℓ_2, in respective link states $(i_1,j_1), (i_2,j_2)$, is given by:

$$path\_cost = p_a^{\ell_1}(i_1,j_1) + p_a^{\ell_2}(i_2,j_2)$$

Path selection: Under the MDPD_D policy, a call is always offered to the direct path first. If it is blocked, i.e., $\mathcal{V}^\ell(i,j) = N^\ell$, the call is then offered to the alternate path with the minimum path cost. A call is rejected if the minimum path cost equals or exceeds the call reward, r.

Policy 2: The deterministic reservation priority MDPD policy with dynamic VP capacity sharing (MDPD_DS):

For this policy, we introduce a new concept, *VP capacity borrowing*, to further reduce the blocking probability of the VP network. The idea is to allow one VP to borrow capacity from another VP if they share the same source node and the set of links in the former VP is a subset of that of the latter. For example, consider the case where VP 1 and VP 2 share the same source node and the set of links of VP 1 is a subset of VP 2. When a VC offered to VP 1 is blocked, instead of offering the VC to an alternate path which might consist of VP 2 and some other VP, the VC could be carried on VP 1 if VP 1 could "borrow" some capacity from VP 2 to carry this VC. The "borrowed" capacity would be "returned" as soon as an existing VC on VP 1 terminates. Borrowing and returning capacity from one VP to another can be performed at the common source node without involving any processing at transit nodes. Therefore, call setup processing can also be done at the source node alone.

The MDPD_DS policy operates very similar to MDPD_D, except for some minor changes regarding path selection and link-load measurements. These changes are described below.

Link loads: The advantage of temporarily borrowing capacity from another VP to carry a VC is that this VC is still carried on a single VP and at the original QOS requirement. This differs from the MDPD_D policy, in which an alternate path with two VP's and a resulting more stringent per-link QOS would be required. Therefore, the VC can be viewed as a direct VC which is partially carried on the VP loaning the capacity. The methods used in MDPD_D for measuring the link carried loads for direct VC's and alternate VC's are modified slightly to reflect this property.

Path selection: As in MDPD_D, the MDPD_DS policy always offers a VC to the direct VP first. If the VC is blocked, MDPD_DS then tries to "borrow" some capacity from other VP's in order to temporarily assign this capacity to the direct VP. If the borrowing is successful, the VC is carried on the direct VP using this borrowed capacity. Otherwise, the VC is offered to the set of alternate paths in the same manner as MDPD_D.

4.3.2 Routing Policies under Statistical Reservation Strategy

A major difference between the routing policies under the statistical reservation strategy and the routing policies under the deterministic reservation strategy is the computation of the link shadow prices. As stated at the end of section 4.2, only one class of QOS is established on each VP under the statistical reservation strategy. Therefore, the computation of link shadow prices for the routing policies under the statistical reservation strategy is exactly the same as in traditional circuit-switched networks. That is, equation (9) can be applied directly to compute the link shadow prices. Furthermore, the costs of adding a direct call and an alternate call are the same. With the exception of this difference, routing policies under the statistical reservation strategy are very similar to the routing policies under the deterministic reservation strategy.

Policy 3: The statistical reservation priority MDPD policy (MDPD_S):

This policy operates slightly differently from MDPD_D in the selection of the set of possible paths for each source-destination pair. As discussed in section 4.2, an alternate path is considered as "possible" under the statistical reservation scheme only if the sum of the QOS of the two VP's on this path satisfies the end-to-end QOS requirement.

Policy 4: the statistical reservation priority MDPD policy with dynamic VP capacity sharing (MDPD_SS):

Borrowing and returning capacity from one VP to another VP is much easier to implement under the statistical reservation scheme because all cells are statistically multiplexed on physical links and the traffic characteristics of a VC are assumed to be unaffected as the VC travels through the network. The QOS received by a VC which is routed on the borrowed capacity is always guaranteed to be more stringent than the original QOS requirement because the VC only travels over a subset of the set of links of the borrowed VP. Thus, the only operation required at the source node is to adjust the maximum number of VC's that can be accepted on these two VP's.

The MDPD_SS policy has the same restriction on selecting the set of possible paths as MDPD_S.

5 Numerical Examples

In this section, we study the performance of the four routing algorithms via a simple 5-node network through simulation. We first describe the performance metric in which we are interested, followed by a detailed description of the network configuration. Finally, the performance of the four routing policies is compared with those of two base-line policies.

5.1 Performance Metric

The performance metric in which we are interested includes four components: the overall VC blocking probability of the network, the largest VC blocking probability of a source-destination pair among all source-destination pairs, the call admission cost, and the resource management cost. The first two metrics are considered as transmission costs while the last two metrics contribute to control costs. We assume the network topology remains fixed during the evaluation of a routing algorithm. Thus, the transmission cost of the network, which is measured by traffic-carrying capacity in [4], bandwidth cost in [5], and transmission efficiency in [6], is measured by the network blocking probability in our study.

The control cost is decomposed into a call admission cost and a resource management cost. The admission cost associated with a call is defined as the number of times the call admission function is invoked in setting up a VC. The resource management cost is defined as the average number of times the resource management function is invoked in setting up a VC.

5.2 Network Configuration

We study routing policies in a 5-node network illustrated in Figure 1(a). We assume links a, b, e and f all have a bandwidth of 4.632 Mbps, each with an input buffer space of 148.224 Kbits (i.e., 32 msec worth of data). The remaining links are assumed to have bandwidth 3.088 Mbps, each with input buffer space 98.816 Kbits. The VP configurations are shown in Figure 1(b). We assume that homogeneous voice sources, described by a two-state fluid flow model with parameters $(32000, 0.3513, 0.352)$ and QOS $Q(0.032, 10^{-6})$ are offered to the network.

5.3 VP Capacity Allocation

VP capacity is assigned as follows. Under the deterministic VP capacity reservation strategy, each VP is assigned a dedicated bandwidth of 1.544 Mbps and a buffer space of 49.408 Kbits. The ADAPTIVE scheme is used to establish two classes of QOS on a VP, one for direct VC's and the other one for alternate VC's. The resulting admissible region is shown in Figure 2.

Under the statistical VP capacity reservation strategy, the bandwidth and buffer space of each physical link are shared among all VP's. The number of VC's reserved for each VP is set such that it is as close as possible to the number of direct VC's that can be supported on the VP when the deterministic VP capacity reservation strategy is used. The ADAPTIVE scheme is also used for establishing multiple QOS classes on a physical link. However, the large number of VC's requiring a more stringent QOS have resulted in uniformly providing the most stringent QOS to all VC's on each physical link. The final assignment of the number of VC's reserved on each VP is to reserve 88 VC's on VP 0, 4, 10, and 14 and 64 VC's on each of the remaining VP's.

5.4 Performance Evaluation of Routing Policies

The performance of the four routing policies is evaluated through simulation. To evaluate the trade-off between transmission and control costs as a result of MDP routing and reserving VP capacity, we also compare the performance of the four routing policies to two base-line routing policies: the DIRECT policy and the NOVP policy. In the DIRECT policy, a deterministic reservation strategy is adopted and all VC's are offered to direct VP's only. A VC blocked on the direct path is rejected. The NOVP policy does not reserve capacity on VP's. Instead, calls are set up in the traditional hop-by-hop fashion. In NOVP, alternate routing is also based on the MDP approach. For a fair comparison, paths in the set of possible paths for each O-D pair are limited to those paths which consist of at most two links. The ADAPTIVE scheme is also adopted for establishing multiple classes of QOS on a physical link.

The performance of the four MDP routing algorithms and two base-line algorithms are compared in Figure 3, Table 1 and Table 2. Each simulation point is observed over 10 independent runs. Each run corresponds to 5000 units of call holding time. For each run, statistics gathered in the initial 10% of the simulation time are discarded. The vertical lines in Figure 3 about each point indicate the 95 percent confidence interval. The following observations can be made from the simulation results.

1. MDP routing algorithms are able to provide significantly lower network blocking probability with a slightly increased control cost as compared to the DIRECT policy.
2. The lowest network blocking probability is achieved by not reserving capacity on VP's. However, this advantage is offset by the significant increase in control cost as shown in Table 1 and Table 2.
3. Routing policies that allow one VP to borrow capacity from another VP yield significantly lower network blocking probability than policies that do not.
4. Routing policies coupled with the statistical VP capacity reservation strategy provide, in general, lower network blocking probabilities and resource management costs than policies coupled with the deterministic strategy. (In fact, routing policies with the statistical strategy do not invoke any resource management function at all.) One exception occurs when the traffic load is very light (e.g., 50 in this example). In this case, routing policies under the deterministic strategy yield lower network blocking probability because the sizes of the possible path sets are larger under these policies.

6 Summary and Discussion

In this paper, MDP routing in VP networks with reserved VP capacity was studied. We first examined two VP capacity reservation strategies. The deterministic strategy reserves dedicated resources on VP's and the statistical strategy, instead of reserving dedicated resources, reserves a dedicated number of VC's, each with a guaranteed QOS, for each VP. Although the second strategy was able to achieve better statistical multiplexing gains, physical links were required to provide a more stringent QOS. We also proposed a ADAPTIVE scheme for more efficiently establishing multiple classes of QOS on a physical link or VP.

Four MDP routing policies were designed and evaluated in this study. The performance metrics studied include network blocking probability, call admission cost, and resource management cost. The performance of the four MDP routing policies was compared to the performance of two baseline routing policies. Our numerical results showed that MDP routing algorithms were able to provide significantly lower network blocking probability with a slightly increased control cost. We also observed that routing policies with a statistical reservation strategy yield, in general, a lower network blocking probability and resource management cost than policies with the deterministic strategy. Finally, the NOVP policy, which does not reserve capacity on VP's, provided the lowest network blocking probability. However, this was offset by the increase in control cost.

References

[1] CCITT SG. XVIII, "Draft Recommendation I.311: B-ISDN General Network Aspects." Geneva, June 1990.

[2] CCITT SG. XVIII, "Draft Recommendation I.371: Traffic Control and Congestion Control in B-ISDN." Geneva, June 1992.

[3] S. Gupta, K. Ross, and M. E. Zarki, "Routing in Virtual Path Based ATM Networks," *GLOBECOM'92*, pp. 571–575, December 1992.

[4] J. L. Burgin, "Management of Capacity and Control in Broadband ISDN," *Intl. J. Digital and Analog Cabled Systems*, vol. 2, pp. 155–165, 1989.

[5] J. Y. Hui, M. B. Gursoy, N. Moayeri, and R. D. Yates, "A Layered Broadband Switching Architecture with Physical or Virtual Path Configurations," *IEEE Journal on Selected Areas in Communications*, vol. 9, pp. 1416–1426, December 1991.

[6] S. Ohta and K. Sato, "Dynamic Bandwidth Control of the Virtual Path in an Asynchronous Transfer Mode Network," *IEEE Transactions on Communications*, vol. 40, pp. 1239–1247, July 1992.

[7] S. Gupta, K. Ross, and M. E. Zarki, "On Routing in ATM Networks," *Modelling and Performance Evaluation of ATM Technology, IFIP Workshop TC6*, p. 10.2, January 1993.

[8] D. Anick, D. Mitra, and M. M. Sondhi, "Stochastic Theory of a Data-Handling System with Multiple Sources," *Bell Systems Technical Journal*, vol. 61, pp. 1871–1894, October 1982.

[9] R. Guérin, H. Ahmadi, and M. Naghshineh, "Equivalent Capacity and Its Application to Bandwidth Allocation in High-Speed Networks," *IEEE Journal on Selected Areas in Communications*, vol. 9, pp. 968–981, September 1991.

[10] A. I. Elwalid and D. Mitra, "Effective Bandwidth of General Markovian Traffic Sources and Admission Control of High Speed Networks," *IEEE/ACM Transactions on Networking*, vol. 1, pp. 329–343, June 1993.

[11] Y. Sato and K. Sato, "Virtual Path and Link Capacity Design for ATM Networks," *IEEE Journal on Selected Areas in Communications*, vol. 9, pp. 104–111, January 1991.

[12] J. Regnier, P. Blondeau, and W. H. Cameron, "Grade of Service of a Dynamic Call Routing System," in *Tenth International Teletraffic Congress*, June 1983.

[13] T. J. Ott and K. R. Krishnan, "State Dependent Routing of Telephone Traffic and the Use of Separable Routing Schemes," in *Eleventh International Teletraffic Congress*, (Kyoto, Japan), September 1985.

[14] R.-H. Hwang, J. F. Kurose, and D. Towsley, "State Dependent Routing for Multirate Loss Networks," *GLOBECOM'92*, pp. 565–570, December 1992.

[15] V. E. Beneš, "Programming and Control Problems Arising from Optimal Routing in Telephone Networks," *Bell Systems Technical Journal*, vol. 45, pp. 1373–1439, 1966.

[16] Z. Dziong *et al.*, "On Adaptive Call Routing Strategy for Circuit Switched Networks - Maximum Reward Approach," in *Twelfth International Teletraffic Congress*, (Torino), June 1988.

[17] Z. Dziong and L. Mason, "Control of Multi-Service Loss Networks," in *Proc. of the 28th Conference on Decision and Control*, (Tampa, Florida), December 1989.

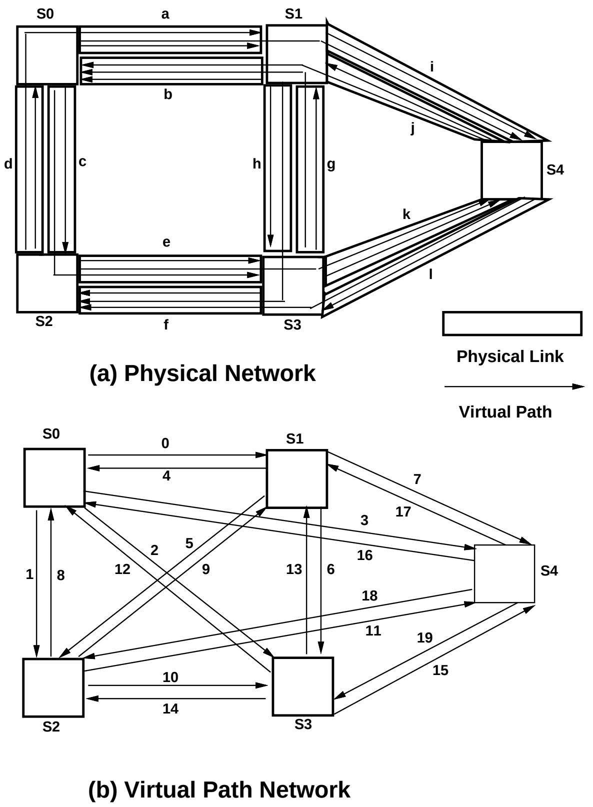

Figure 1: Example: A 5-node network and its corresponding VP network.

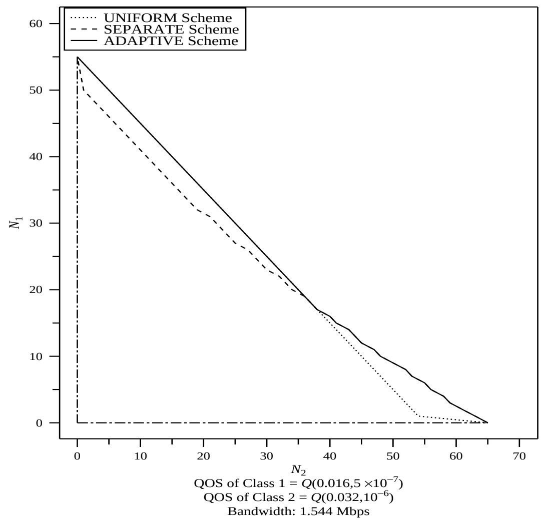

Figure 2: Comparison of the maximum number of VC's that can be supported under different cell control schemes.

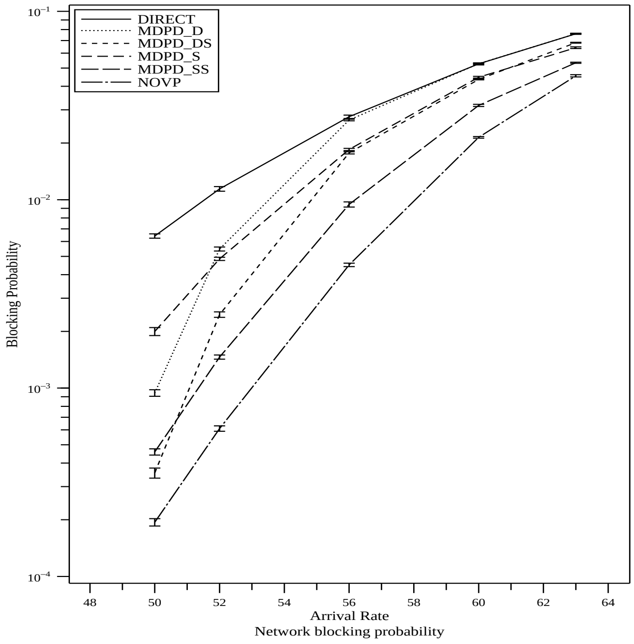

Figure 3: Comparison of network blocking probability of different routing policies

Table 1: Comparison of call admission cost of different routing policies.

Policy	Arrival rates				
	50	52	56	60	63
DIRECT	1.000	1.000	1.000	1.000	1.000
MDPD_D	1.010	1.016	1.012	1.003	1.000
MDPD_DS	1.005	1.008	1.005	1.001	1.000
MDPD_S	1.007	1.011	1.014	1.004	1.004
MDPD_SS	1.005	1.008	1.005	1.001	1.000
NOVP	1.400	1.400	1.402	1.406	1.410
*All 95% confidence intervals are less than 0.001.					

Table 2: Comparison of resource management cost of different routing policies.

Policy	Arrival rates				
	50	52	56	60	63
DIRECT	0.000	0.000	0.000	0.000	0.000
MDPD_D	0.019	0.032	0.025	0.005	0.000
MDPD_DS	0.022	0.039	0.061	0.060	0.060
MDPD_S	0.000	0.000	0.000	0.000	0.000
MDPD_SS	0.000	0.000	0.000	0.000	0.000
NOVP	1.400	1.400	1.402	1.406	1.410
*All 95% confidence intervals are less than 0.001.					