LLR Routing in VP-based ATM Networks

Ren-Hung Hwang

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

Abstract

A Virtual Path (VP) concept has been proposed to simplify traffic control and resource management in future B-ISDN. Especially, significant call setup processing can be reduced when resources are reserved on VP's. In this paper, four adaptive routing algorithms based on the Least Loaded Routing *(LLR) concept are designed and evaluated for VP-based ATM networks in which resources are reserved on VP's. Our simulation results show that network blocking probability can be significantly reduced by LLR routing.*

1 Introduction

In ATM (Asynchronous Transfer Mode) networks, the Virtual Path (VP) concept was proposed to reduce the control costs by grouping connections sharing common paths through the network into a single unit [1, 2]. A VP can be viewed as a single logical direct link between a source node and a destination node which may consist of two or more physical links which together form a path. An ATM network with VP's is often referred to as VP-based ATM networks.

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.

In this paper, we study the routing problem in VP-based ATM networks when capacity is reserved on VP's. In particular, we are interested in homogeneous VP subnetworks in which all applications are assumed to require identical end-to-end QOS guarantees [3, 4]. Our approach is based on the *Least Loaded Routing* (LLR) concept, a well known, simple and efficient routing concept developed for circuit-switched networks. In our numerical results, we show that our LLR-based routing algorithms are able to reduce the network blocking probability significantly.

Routing algorithms based on the LLR approach for homogeneous VP networks are also studied in [3]. Our research differs from this work in the following respects. First, 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. Second, our routing algorithms are designed based on two more efficient VP capacity reservation strategies proposed in [4]. Finally, a dynamic VP capacity sharing concept is proposed to further reduce the network blocking probability. The same routing problem is also studied in our previous work [4] in which routing algorithms based on Markov decision theory were proposed.

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 proposed in [4] are then summarized in section 3. In section 4, four routing algorithms based on the LLR concept are proposed. The performance of the four routing algorithms are evaluated and compared to the performances of two base-line algorithms in section 5. Finally, section 6 summaries this paper.

2 Network Model

2.1 A Virtual Path Network

We assume that VP's are used for separating different traffic types [2]. That is, a VP with dedicated resources is assigned to each type of traffic between a source-destination pair. 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.

In this paper, we will only focus on homogeneous VP-based ATM networks [5]. A homogeneous VP-based ATM network is a subnetwork of a VP-based ATM network in which all connections require identical *end-to-end* QOS guarantees. We also assume that one VP is assigned to each source-destination pair. Therefore, a logical, fully connected, VP-based ATM network can be defined by viewing the VP endpoints as nodes and VP's as links, as shown in Figure 1.

2.2 Traffic Model, QOS Criterion, and Call Admission Scheme

The traffic model, QOS criterion, and call admission scheme used in [4] is also adopted in our numerical study. The reader is referred to [4] for the details.

Each traffic source is described by a two-state fluid-flow model [6] in which a source is either in an "on" state, transmitting at its peak rate, or in an "off" state, transmitting

This research was supported in part by the 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.

at zero bit rate. The durations of the "on" and the "off" states are assumed to be exponentially distributed random variables. Therefore, 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) [7]. 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.

The QOS requirement by a connection 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. 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.

The call admission scheme we adopted is the one proposed by Guérin *et al.* [7], which is based on the fluid-flow approximation model of [6]. In summary, the maximum number of sources, described by (r, ρ, b) with a QOS requirement $Q(D, \epsilon)$, that can be multiplexed onto a link of capacity C bps is given by

$$N = \lfloor \frac{C}{\hat{c}} \rfloor \tag{1}$$

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

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

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

3 Resource Management

In this paper, we assume certain amount of capacity is reserved on each VP. In this section, we first summarized two VP capacity reservation strategies proposed in [4]. We then address the issue of supporting multiple classes of QOS on a VP (or physical link).

3.1 VP Capacity Reservation Strategies

In [4], we proposed two VP capacity reservation strategies: a *deterministic* strategy and a *statistical* strategy. The basic distinction between these two strategies is whether separate capacity is reserved to the VP's. The deterministic strategy reserves dedicated resources (such as buffer space and bandwidth) to each VP while the statistical strategy allows statistical cell multiplexing among the cells from different VP's on a physical link.

In the deterministic strategy, each VP can be viewed 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. When a physical link is shared by several VP's, the cells of different VP's are transmitted in a TDM-like manner such that the instantaneous bit rate transmitted on a VP is limited to its reserved bandwidth. Intuitively, this can be viewed as if no multiplexing were performed among VP's on the link-level. Thus, the maximum number of virtual circuits that can be supported by a VP without violating a given QOS guarantee can be computed locally at the source node, based on the knowledge about the the size of the reserved buffer space and the amount of reserved bandwidth.

The statistical strategy 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.

As in in [4], two assumptions were made in implementing the statistical strategy. First, we assume that the traffic characteristics of a virtual circuit are unaffected as the traffic travels along a VP. Second, we assume that the equal allocation (EQ) policy were adopted which assigns an equal amount of the end-to-end QOS of a virtual circuit to each link on a VP. As a consequence, 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.

3.2 Multiple Classes of QOS

Supporting multiple classes of QOS on a physical link or VP is required in the following two scenarios: (1) in the implementation of the statistical VP capacity reservation strategy, (2) when both direct calls (i.e., calls carried on a VP) and alternate calls (i.e., calls carried on an alternate path consists of two VP's) are carried on a VP simultaneously. In [4], a simple cell transfer control scheme, referred to as ADAPTIVE, is proposed for providing multiple classes of QOS on a physical link or VP. The ADAPTIVE scheme is a combination of two schemes: the UNIFORM scheme and the SEPARATE scheme.

In the UNIFORM scheme, all cells are provided with the most stringent QOS. 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.

On the other hand, the SEPARATE scheme separates cells according to their QOS requirements and allocates 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 cell multiplexing gains because of the partitioning of the resources. The reader is referred to [4] for the algorithm for determining the amount of buffer space and bandwidth required to dedicate to each class of QOS in the SEPARATE scheme.

The ADAPTIVE scheme is to take the advantage of both schemes. That is, the ADAPTIVE scheme will adopt the UNIFORM scheme if it supports more virtual circuits than the SEPARATE scheme. Otherwise, the ADAPTIVE scheme will adopt the SEPARATE scheme. In Figure 2, we compare and contrast these three schemes 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. From Figure 2, we can observe that the SEPARATE scheme is a complement to the UNIFORM scheme while the ADAPTIVE scheme takes advantages of both schemes.

4 VP Routing

Routing in homogeneous VP networks is very similar to routing in circuit-switched networks. Each virtual path can be viewed as a logical link with allocated capacity. Thus, as shown in Figure 1, a VP network can be thought of as a fully connected network even though the physical network itself is sparse. Existing adaptive routing algorithms in circuit-switched networks can be classified into two categories: *Least Loaded Routing*-based (LLR-based) [8,9] and *Markov Decision Process*-based (MDP-based) [10,11]. The LLR approach tries to route an incoming call to the direct link, the link between the source node and destination node, first. If the call is blocked (because of no free circuits), the "least busy" path, i.e., the path that has the maximum number of free circuits, is then tried. The MDP approach, on the other hand, formulates the routing problem as a Markov decision process. In the following, we propose routing algorithms for VP networks based on the LLR approach. MDP-based routing algorithms for VP networks were studied in our previous work [4].

4.1 LLR Routing in VP Networks

Since LLR-based routing algorithms route a call to alternate paths when it is blocked on the direct VP, let us first define the set of candidate alternate paths for each source-destination pair to be all possible paths, which consist of two VP's, that connecting the source node to the destination node. As in the previous section, we assume that the traffic characteristics of a virtual circuit are unaffected as the traffic travels through the first VP of an alternate path. However, 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 is 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. (Recall that a homogeneous VP network handles only one type of traffic with homogeneous end-to-end QOS requirements and alternate paths are limited to at most two VP's.) The ADAPTIVE scheme proposed in the previous section is applicable for providing two QOS classes on a VP with dedicated capacity.

In the case of the statistical VP capacity reservation strategy, providing two classes of QOS on a VP is much more complicated. Thus we assume that only one class of QOS is provided on a VP. Furthermore, as shown in Figure 3, VP's which consist of fewer physical links (e.g., VP 1 in Figure 3) is provided more stringent QOS than VP's which consist of more physical links (e.g., VP 3 in Figure 3). Therefore, besides limiting the alternate paths to consist of at most two VP's, a path is considered as a candidate alternate path only if the sum of the QOS of the two VP's constituting this path satisfies the end-to-end QOS requirement.

4.2 LLR Routing Algorithms

In this section, we describe four LLR routing algorithms under the two different VP capacity reservation strategies. We first describe the LLR routing algorithm under the deterministic reservation strategy in detail. Other algorithms are described only in the manner in which they differ from the first algorithm.

4.2.1 Routing Policies under Deterministic Reservation Strategy

In the following description of LLR routing algorithms, the free capacity of a VP is measured by the maximum number of direct calls that can be added to the VP. As shown in Figure 2, 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. For example, consider the link studied in Figure 2. Let class 1 calls be alternate calls and class 2 calls be direct calls. (Recall that the QOS of class 1 is more stringent than that of class 2.) If there are no alternate calls, the link can support up to 65 direct VC's without violating the QOS requirement. However, when one alternate call is added to the link, only 63 direct VC's can be supported (under the ADAPTIVE scheme).

Algorithm 1: The deterministic reservation LLR algorithm (LLR_D):

When a call arrives at the source node, following steps are performed:

1. The call is first offered to the direct VP.
2. If the direct VP does not have enough capacity available to carry the call, the call is offered to the alternate path with the maximum free capacity where the free capacity of an alternate path is defined as the minimum free capacity over the path's VP's. The free capacity of a VP on an alternate path is given by $Cap(\ell) - t(\ell)$ where $t(\ell)$ is the trunk reservation level of the VP which will be defined later.
3. If the call still cannot be carried by the alternate path with the maximum free capacity, then the call is blocked.

To avoid network performance degradation under high traffic load, a trunk reservation concept is adopted in LLR-based routing algorithms [12]. The optimal trunk reservation problem in circuit-switched networks has been well studied in the literature [13,14]. In our study, near optimal trunk reservation levels are obtained based on the approximation technique proposed in [4]. The reader is referred to [15].

Algorithm 2: The deterministic reservation LLR algorithm with dynamic VP capacity sharing (LLR_DS):

In this algorithm, the *VP capacity borrowing* concept, proposed in [4], is adopted 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 VP 1 and VP 3 in Figure 3. When a VC offered to VP 1 is blocked, instead of offering the VC to an alternate path which might consist of VP 3 and some other VP, the VC could be carried on VP 1 if VP 1 could "borrow" some capacity from VP 3 to carry this VC. The "borrowed" capacity would be "returned" as soon as one of existing VC's on VP 1 terminates. Borrowing and returning capacity from one VP to another is simple to do and need to be performed only at the common source node; no processing is required at transit nodes. Thus, call setup processing can still be done at the source node.

The LLR_DS algorithm operates very similar to LLR_D, except that an additional step is added between Step 1 and Step 2 in the LLR_DS algorithm. This extra step is to allow dynamic VP capacity sharing between two VP's. That is, add the following step between Step 1 and Step 2 in the LLR_DS algorithm:

1(a): If the call is blocked on the direct VP and there exists another VP with free capacity that can be shared by the direct VP, the call is routed to that VP.

4.2.2 Routing Policies under Statistical Reservation Strategy

As stated at the end of section 4.1, only one class of QOS is established on each VP under the statistical reservation strategy. Therefore, the major difference between the routing policies under the statistical reservation strategy and the routing policies under the deterministic reservation strategy is the selection of the set of possible alternate paths for each source-destination pair. 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.

Algorithm 3: the statistical reservation LLR algorithm (LLR_S):

The LLR_S operates exactly the same as the LLR_D algorithm except that the selection of the set of possible alternate paths for each O-D pair and the optimal trunk reservation levels are different. Because only one class of QOS is provided on each VP under the statistical reservation strategy, the optimal trunk reservation problem is exactly the same as in the traditional circuit-switched networks. In [15], we formulate the optimal trunk reservation problem as a Markov decision process [13], and derive a set of simple expressions for computing the optimal trunk reservation levels.

Algorithm 4: the statistical reservation LLR algorithm with dynamic VP capacity sharing (LLR_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 LLR_SS algorithm has the same restriction on selecting the set of possible paths as LLR_S. The setting of trunk reservation levels is also the same as LLR_S.

5 Numerical Examples

In this section, we study the performance of the four routing algorithms via a simple 5-node network through simulation [4].

5.1 Performance Metric

As in [4], the performance metric we are studied includes three components: the overall VC blocking probability of the network, the call admission cost, and the resource management cost. We assume the network topology remains fixed during the evaluation of a routing algorithm. Thus, the network blocking probability is used to measure the transmission cost. On the other hand, the call admission cost and the resource management cost contribute to control costs. The admission cost associated with a VC is defined as the number of times the call admission function is invoked in setting up the VC. The resource management cost is defined as the average number of times the resource management function is invoked in setting up the VC.

5.2 Network Configuration

The topology of the 5-node network is illustrated in Figure 1(a). The bandwidth and buffer space for each link is configured as follows: 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) while the remaining links 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 in [4]. 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. For 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 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 Algorithms

The performance of the four routing algorithms is evaluated through simulations. As in [4], the performance of the four routing algorithms is compared to the performance of two base-line routing algorithms: the DIRECT algorithm and the NOVP algorithm. In the DIRECT algorithm, 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 algorithm 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 LLR 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 performance of the four LLR-based routing algorithms and two base-line algorithms are compared in Figure 4. 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 is discarded. The vertical lines about each point indicate the 95 percent confidence interval. The following observations, which is similar to the results of [4], can be made from Figure 4.

1. LLR routing algorithms are able to provide significantly lower network blocking probability with a slightly increased control cost as compared to the DIRECT algorithm.
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 algorithms that allow one VP to borrow capacity from another VP yield significantly lower network blocking probability than algorithms that do not.
4. Routing algorithms under the statistical VP capacity reservation strategy provide, in general, lower network blocking probability and resource management cost than algorithms with the deterministic strategy. (In fact, routing algorithms 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 algorithms under the deterministic strategy yield lower network blocking probability because the sizes of the possible path sets are larger under these algorithms.

In [15], the performance the four LLR-based routing algorithms is compared to the performance of the four MDP-based routing algorithms studied in [4]. We observe that MDP-based routing algorithms yield slightly better network blocking probabilities than the LLR-based routing algorithms.

6 Summary

In this paper, LLR routing in VP networks with reserved VP capacity was studied. We first summarized two VP capacity reservation strategies proposed in [4]. 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. Four LLR routing algorithms were then designed and evaluated under these two strategies. The performance metrics studied include network blocking probability, call admission cost, and resource management cost. The performance of the four LLR routing algorithms was first compared to the performance of two base-line routing algorithms. Our numerical results showed that LLR routing algorithms were able to provide significantly lower network blocking probability with a slightly increased control cost. We also observed that routing algorithms with a statistical reservation strategy yield, in general, a lower network blocking probability and resource management cost than algorithms with the deterministic strategy. Finally, the NOVP algorithm, which does not reserve capacity on VP's, provided the lowest network blocking probability. However, this was offset by the increase in control cost. In [15], we observe that the MDP routing algorithms only slightly outperforms the LLR routing algorithms.

References

[1] CCITT SG. XVIII, "Some Key Features on UNI and NNI Considering ATM (A framework for discussion)." Seoul Meeting, 25 Jan. – 5 Feb. 1988.

[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] R.-H. Hwang, J. F. Kurose, and D. Towsley, "MDP Routing in ATM Networks Using the Virtual Path Concept," *INFOCOM'94*, pp. 1509–1518, June 1994.

[5] 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.

[6] 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.

[7] 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.

[8] G. R. Ash, "Use of a Trunk Status Map for Real-Time DNHR," in *Eleventh International Teletraffic Congress*, (Kyoto, Japan), September 1985.

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

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

[11] K. R. Krishnan, "Markov Decision Algorithms for Dynamic Routing," *IEEE Communications Magazine*, vol. 28, pp. 66–69, October 1990.

[12] J. M. Akinpelu, "The Overload Performance of Engineered Networks with Nonhierarchical and Hierarchical Routing," in *Tenth International Teletraffic Congress*, June 1983.

[13] P. B. Key, "Optimal Control and Trunk Reservation in Loss Networks," *Probability in the Engineering and Informational Sciences*, vol. 4, pp. 203–242, 1990.

[14] M. I. Reiman, "Optimal Trunk Reservation for a Critically Loaded Link," in *13th International Teletraffic Congress*, June 1991.

[15] R.-H. Hwang, "Routing in VP-based ATM Networks," *Submitted to Journal of Information Science and Engineering*, July 1994.

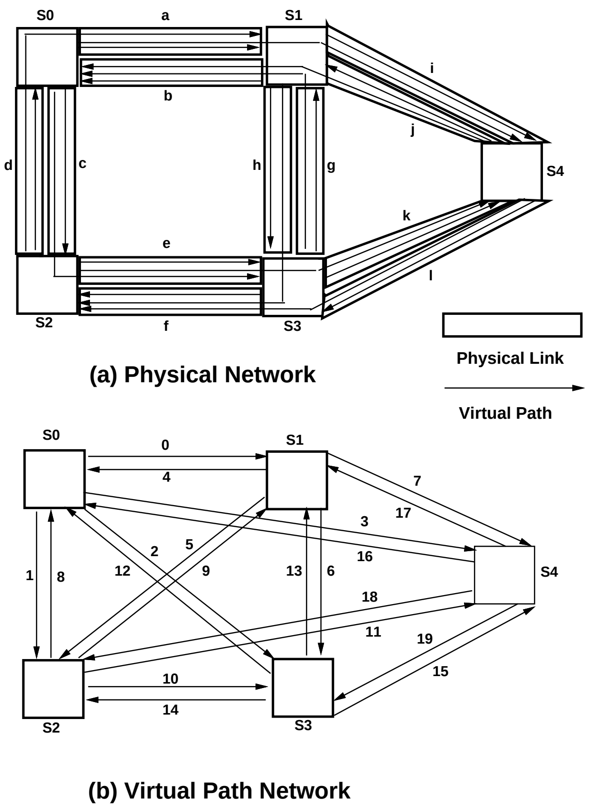

Figure 1: Example: A 5-node network and its corresponding VP network.

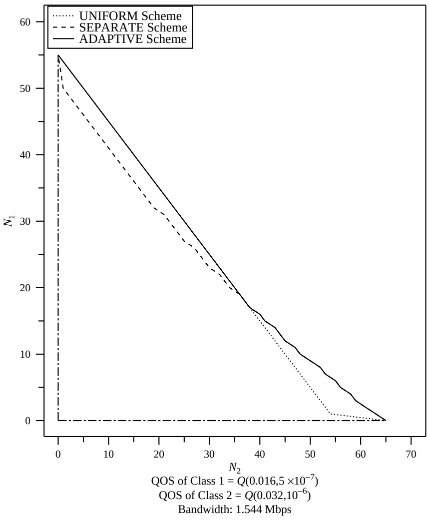

Figure 2: Comparison of the maximum number of VC's that can be supported under different cell control schemes.

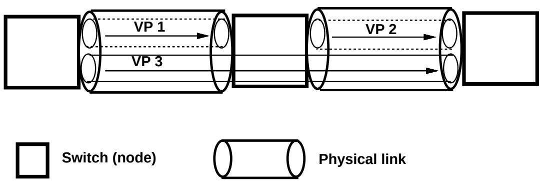

Figure 3: Example: A simple VP network

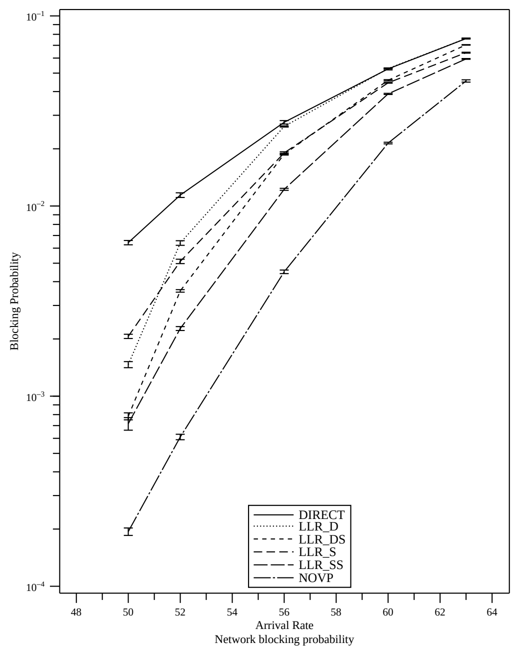

Figure 4: Comparison of network blocking probability of different routing policies

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
LLR_D	1.008	1.009	1.005	1.001	1.000
LLR_DS	1.004	1.005	1.002	1.000	1.000
LLR_S	1.006	1.009	1.011	1.004	1.004
LLR_SS	1.004	1.006	1.008	1.004	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
LLR_D	0.015	0.019	0.010	0.001	0.000
LLR_DS	0.019	0.031	0.050	0.044	0.042
LLR_S	0.000	0.000	0.000	0.000	0.000
LLR_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.					