42

A Diameter Based Method for Virtual Path Layout in ATM Networks

Jimmy P. Wong, Rajesh K. Pankaj
University of Toronto
10 King's College Road, Toronto, Canada M5S 1A4. Phone: +1 416 978 0806.
Fax: +1 416 978 4425. email: wongja,pankaj@comm.toronto.edu

Abstract

ATM networks use virtual paths (VPs) to route information from source to destination. By using VPs, the call set-up and switching costs can be reduced. In this paper, we consider the problem of selecting where VPs should be placed. We propose an algorithm that is based on reducing the network diameter. The performance of this *diameter* method is compared with another heuristic that uses a *clustering* algorithm. Using simulation, it is shown that the Diameter method performs better than the Clustering method in reducing the average connection cost.

Keywords

ATM, Virtual Paths, Network Management

1 INTRODUCTION

The use of virtual paths is an important concept associated with Asynchronous Transfer Mode (ATM) (Ohta, 1988), (Burgin, 1991), (Sato, 1990). A virtual path (VP) is a logical connection between two nodes in a network. Each VP consists of a sequence of one or more physical links. A network of VPs forms a higher layer that is logically separate from the underlying physical network. From this VP network, individual connections, or virtual channels (VCs), can be routed. Each VP can be used by many virtual channels. That is, VCs are multiplexed together on to a VP and transported with a common identifier, called a Virtual Path Identifier (VPI). Individual VCs within a virtual path are distinguished by their Virtual Channel Identifier (VCI).

The advantages of using virtual paths include the following. At call set-up, the routing tables at the intermediate nodes of an existing VP do not need to be updated. This reduces the call set-up delays. By grouping many connections (VCs) into single units (VPs), the switching costs are reduced. In particular, at the intermediate nodes of a VP, only the VPI label of a VC needs to be processed. As shown in (Burgin, 1991), more than 90% of the processing time can be saved if virtual paths are used. However, this improvement can only be obtained if the VPs are assigned efficiently. That is, given a network topology and traffic distribution, establish a system of virtual paths so that the network performance is optimized. This involves finding, for each VP, the VP end-nodes (or terminators), the actual route between the terminator nodes, and the path capacity.

In what sense is a VP network optimal? In papers by Cheng and Lin (Cheng, 1994) and Ahn *et al.* (Ahn, 1994),

L. Mason et al. (eds.), *Broadband Communications*

VPs are assigned to minimize the call blocking probability. Chlamtac *et al.* (Chlamtac, 1993) use a *Clustering* algorithm to establish VPs. Another issue in assigning VPs is the effect of failures in the network. Murakami and Kim (Murakami, 1994) propose a VP routing scheme which minimizes the expected amount of lost flow due to a network failure.

In this paper, we propose a solution to a simplified form of the above VP assignment problem. Given a network topology and traffic demands, a VP network is established to minimize the *average connection cost* under the constraint that the number of VPs assigned in the network is limited. We then also consider this problem under an additional constraint that the number of VPs on each link is less than a prespecified bound. The bounds on the number of VPs in the network and the number of VPs over a link are important consideration because of the limited number of available VPI addresses. If separate VP layouts are needed for different classes of service supported by the network, the limit on the number of VPI addresses can become a significant constraint. In our algorithm, we consider a single set of VPs. If different VP layouts are required for different classes of service, the algorithm should be run separately to determine the VP layout for each class of service. A bound on the number of VPs traversing a given link also limits the effect of a single link failure on the VP layout.

We consider only the problem of finding the VP terminators and the actual path between the end-nodes of each VP. We do not determine the bandwidth allocated to each VP. The work on bandwidth assignment in conjunction with the algorithms proposed in this paper is under progress.

Note that the *cost* associated in routing a connection can be a general cost function. However, one advantage of the use of virtual paths is a reduction in call set-up and switching costs. Therefore, in this paper, we consider a cost function that is representative of the call set-up and switching costs.

This paper is organized as follows. In Section 2, the VP assignment problem is defined. Section 3 describes the proposed VP layout algorithm and analyzes its computational complexity. In Section 4, some simulation results are presented to compare the proposed method with the Clustering algorithm proposed in (Chlamtac, 1993). Section 5 concludes the paper.

2 PROBLEM FORMULATION

The VP assignment problem is formulated as follows.

1. Let an undirected graph $G = (V, E)$ represent the physical network, where V denotes ATM switches and E represents physical links connecting the nodes. Let graph G have $|V| = N$ nodes and an arbitrary number of edges. We assume the network is connected.
2. The amount of traffic from node i to node j is given by ρ_{ij}. Let the total amount of traffic in the network be

$$\rho_T = \sum_{i,j \in V} \rho_{ij}. \tag{1}$$

3. A cost function $C(i, j)$ represents the cost of routing one unit of traffic from node i to node j. Assuming a connection transmits one unit of traffic, $C(i, j)$ represents the cost of routing the connection from i to j. As discussed previously, we assume $C(i, j)$ is associated with the call set-up and switching costs. Note that $C(i, j)$ is a function of the VPs assigned to the network.
4. Let M be the number of VPs assigned to the network. Each VP is defined by a terminator pair, (s, t), and the actual route from s to t. The cost of using VP (s, t) is given by

$$C_{VP}(s, t) = \beta + \alpha\, d(s, t) \tag{2}$$

where $d(s,t)$ is the number of physicals links traversed by the VP. The cost of using a physical link by itself is 1. Parameter β is a fixed cost in routing over the VP. This includes the cost associated with the selection of an available VCI label for the connection during call set-up and also the cost of processing of the VCI label at the end-nodes of the VP. $\alpha\, d(s,t)$ is a cost that is proportional to the length of the VP. It represents the cost of translating the VPI field at the intermediate nodes of the VP.
The variables α and β are chosen so that $C_{VP}(s,t) < d(s,t)$. That is, routing a connection over the VP is less costly than routing over the corresponding physical links. Finally, $C_{VP}(s,t)$ should not be less than $d(s,t)$ if node s is adjacent to node t. The cost of routing over a one-link VP should be greater than or equal to the cost of routing over the physical link itself.

5. Each physical link, e_k, in the network is restricted to having a maximum of η_k VPs routed over it. This constraint is motivated by two possible reasons. First, the number of available VPI addresses may be limited which puts a limit on the number of VPs that can traverse that link. Second, by limiting the number of VPs that can be on a physical link, we limit the sensitivity of the VP layout network in the event of a link failure.
The first of the two reasons can be important if we see the VP layout algorithm running separately for each class of service. If a larger number of classes of service are supported, each class will have only a limited number of VPI addresses. We define the *maximum link load* to be the largest number of VPs going through the same physical link.

Given the above constraints, assign a set of VPs to minimize

$$\bar{C} = \frac{\sum_{i,j\in V} \rho_{ij} C(i,j)}{\rho_T}. \tag{3}$$

$\bar{C}$ represents the average connection cost of the network given a VP layout. Note that if we let $D(i,j) = \rho_{ij}C(i,j)$, minimizing

$$\sum_{i,j\in V} D(i,j) \tag{4}$$

is equivalent to minimizing Eqn. 3.

In the next section, we propose an algorithm which attempts to minimize $\bar{C}$ while maintaining a small maximum link load.

3 VP ASSIGNMENT ALGORITHMS

In this section, we present a heuristic for assigning virtual paths that is based on reducing the network diameter. We will refer to it as the *Diameter method.* As a comparison, we also briefly discuss a method proposed in (Chlamtac, 1993). The computational complexity of both methods are then analyzed.

3.1 Diameter Method

In the Diameter method, a VP is assigned if it reduces the *diameter* of the network. The diameter of a network is the *distance* between the node-pair that is farthest apart. Distance can be defined as the minimum number of hops between the two nodes, the propagation delay, or some other measure. Therefore, depending on what *cost* the VP assignment is set to minimize, an appropriate distance measure is chosen.

The Diameter algorithm is iterative. In each iteration, a new VP is selected. Therefore, to assign M VPs, M iterations are performed. For each virtual path, the VP terminator-pair is first chosen, and then the actual route is selected.

The selection of a VP proceeds as follows. Let $G = (V, E)$ represent the physical network. Let S_0 be the set of VPs selected in the previous iterations. We define $E' = E \cup S_0$ and a weighted digraph $G' = (V, E')$. G' represents the VP network overlayed on top of the physical network. Let g_k and g'_k be the weights associated with edges in G and G', respectively. The edges in G' associated with physical links are assigned a weight $g'_k = 1$. An edge representing virtual path (s, t) is assigned a weight $g'_k = C_{VP}(s, t)$ (refer to Eqn. 2).

Define another graph H and let the edges have weight h_k. H keeps track of the physical links on which new VPs can be routed. Note that VPs may not be routed over links that already have the maximum number of VPs routed over them. Let η_k be the maximum of VPs that can use physical link e_k. Let n_k be the number of VPs using link e_k. Let $Q = \{(i, j) \mid i, j \in V\}$ be the set of node-pairs in graph H where there exists no path from i to j. Initially, graph $H = G$, and hence is connected. However, when a link, e_k, cannot support any more VPs ($n_k = \eta_k$), it is removed from H. Therefore, H may become disconnected as VPs are added to the network. Q is the set of node-pairs where a path no longer exists due to the maximum allowable VP per link constraint.

In the Diameter algorithm, graph G' is used to find the VP terminator-pair and graph H is used to find the actual route. Define

$$D(i, j) = \rho_{ij} C(i, j) \tag{5}$$

as the *distance* from node i to node j in graph G'. Recall ρ_{ij} is the amount of traffic from i to j. Let $C(i, j)$, the cost of routing a unit of traffic from i to j, be the weight of the shortest path from i to j in graph G'.

The VP terminator-pair is selected as follows. A node-pair (s, t) that is maximally separated, according to the distance measure $D(s, t)$, is selected. By assigning a VP at (s, t), we can reduce the distance between the node-pair, and hence the diameter of the network can be reduced.

Once a VP terminator (s, t) is chosen, the actual route can be selected. Using graph H, a minimum weight path is found from s to t. Recall that each physical link e_k in the network is restricted to supporting only η_k VPs. Initially, edges in graph H have a weight $h_k = 1$. However, when a link is used by the maximum allowable number of VPs, the edge weight is changed to $h_k = \infty$. Therefore, these edges are removed from consideration when the paths for subsequent VPs are computed. Note that if a path in graph H cannot be found for the VP (s, t), then that node-pair is ignored and we proceed to find another terminator-pair.

The above process is repeated using the updated graphs G' and H to find the next VP terminator-pair and VP route, respectively. The Diameter method is summarized below.

Diameter Algorithm

Let graph $H = G(V, E)$. Edges in G and H have weight $g_k = 1$ and $h_k = 1$, respectively. Let $S_0 = \{\}$; initially no VPs are selected. Let $Q = \{\}$; initially graph H is connected. The number of VPs using link e_k is initially zero ($n_k = 0$ for all links). For each VP do:

1. Let $G' = (V, E \cup S_0)$. VP terminator-pairs that have been previously selected are added to G to form a new graph, G'. Note that physical links have weight $g'_k = 1$, while VP edges have weight $g'_k = C_{VP}(s, t)$.
2. Using graph G', find

$$D_{\max} = \max_{(i,j) \notin S_0 \text{ or } Q} D(i, j)$$

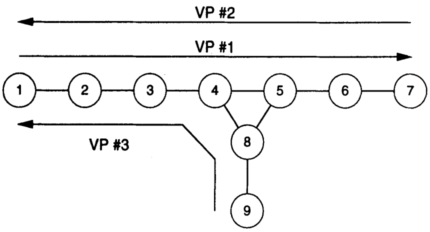

Figure 1 Applying Diameter Method to a 9-node Network

and the corresponding the node-pair, $x = (i, j)$, with $D(i, j) = D_{\max}$.

3. A VP terminator-pair is selected at $y = (s, t) = (i, j)$.
4. Find the shortest path and the corresponding shortest distance, $h(s, t)$, from node s to node t in graph H. If $h(s, t) < \infty$, then a feasible path was found. For each edge e_k used by the path do:

   - $n_k = n_k + 1$
   - If $n_k = \eta_k$, then set $h_k = \infty$

   If $h(s, t) = \infty$, then no feasible path was found. Therefore, a VP cannot be placed at (s, t). Let $Q = Q \cup \{(s, t)\}$ and go to Step 2.
5. Let $S_0 = S_0 \cup \{(s, t)\}$. VP (s, t) is added to the set of selected VPs, S_0.

As an example, we apply the Diameter algorithm to a 9-node 9-link network (Figure 1). We assume the traffic distribution is uniform ($\rho_{ij} = 1$ for all node-pairs) and the VP cost parameters are $\beta = 0.9$ and $\alpha = 0.1$. There is no restriction on the maximum number of VPs per link ($\eta_k = \infty$ for all links). Since node-pairs $(1, 7)$ and $(7, 1)$ are equally far apart, we arbitrarily place the first VP at $(1, 7)$. The second and third VPs are placed at $(7, 1)$ and $(9, 1)$, respectively.

The Diameter method reduces the diameter of the network in an attempt to minimize the average connection cost. However, assigning a VP between a node-pair that is farthest apart in the network is, in general, not optimal. For example, consider the 7-node 1-D network* in Figure 2. Assume the traffic distribution is uniform, $\beta = 0.9$ and $\alpha = 0.1$. For simplicity, we assign a single bidirectional VP. Using the Diameter method, a VP is placed at $(1, 7)$. However, the optimal VP location to minimize the average connection cost is at $(2, 6)$. Therefore, assigning a VP end-to-end in the 1-D network is not optimal.

By placing the VP terminators away from the end-nodes of the 1-D network, we can reduce the average connection cost further. But where is the optimal location?

Given an N node 1-D network with an end-to-end length equal to $\mathcal{L}$, the optimal location of a single bidirectional VP can be found. If we assume the traffic distribution is uniform, as $N \to \infty$, the optimal VP placement occurs

*In a 1-D network with nodes $\{1, 2, \ldots, N\}$, node i is connected to nodes $(i - 1)$ and $(i + 1)$.

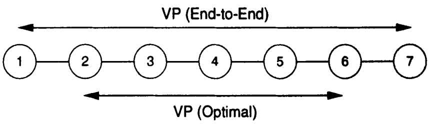

Figure 2 7-node 1-D Network

at (s, t), where s and t are the two nodes which are at a distance $a\mathcal{L}$ away from the end-nodes of the 1-D network. Parameter a is found using the following equation (see (Wong, 1995) for details).

$$a = \frac{(\alpha' + \beta' - 1)\left(-1 + 2\alpha' + \alpha'^2 + \sqrt{2\left(1 + \alpha'^2\right)}\right)}{2\left(\alpha' - 1\right)\left(1 + 4\alpha' + \alpha'^2\right)} \tag{6}$$

where $\alpha' = \alpha/\mathcal{L}$ and $\beta' = \beta/\mathcal{L}$. By using Eqn. 6, the optimal (or near-optimal) location of a VP can also be found for a 1-D network with $N < \infty$. Note that parameter a was computed under the assumption that $N = \infty$ and therefore it may not be optimal for finite N.

3.2 Modified Diameter Method

In Section 3.1, it was shown that assigning a VP end-to-end in a 1-D network is not optimal. In an attempt to improve the performance of the algorithm, we extend the results from Eqn. 6 to the Diameter method for general networks with uniformly distributed traffic.

Instead of assigning a VP between node-pairs that are furthest apart in the network, a VP is established between node-pairs that are *slightly* less separated than the maximum amount. The algorithm is implemented as follows. A node-pair (i, j) that is maximally separated, according to the distance measure $D(i, j)$, is chosen. Let $D_{max} = D(i, j)$ and let $\mathcal{P}$ represent the shortest path from i to j with length D_{max}. Consider path $\mathcal{P}$ as a 1-D network with length D_{max}. A VP is placed at (s, t) according to the results of Eqn. 6. That is, s and t are at a distance aD_{max} from the end-nodes of $\mathcal{P}$.

The modified Diameter method is summarized below. Note that only Step 3 in the algorithm is changed.

Diameter Algorithm (Modified)

Let $S_0 = \{\}$ and let $Q = \{\}$. Let $n_k = 0$ for all links. For each VP do:

1. Let $G' = (V, E \cup S_0)$.
2. Using graph G', find

   $$D_{\max} = \max_{(i,j) \notin S_0 \text{ or } Q} D(i, j)$$

   and the corresponding the node-pair, $x = (i, j)$, with $D(i, j) = D_{\max}$.
3. Let $\mathcal{P}$ be the path from i to j in G' with length $D_{\max}$. Model path $\mathcal{P}$ as a 1-D network with uniform traffic and assign the VP terminator-pair (s, t) according to Eqn. 6.
4. Find the shortest path from node s to node t in graph H. If no feasible path is found, a VP cannot be placed at (s, t). Let $Q = Q \cup \{(s, t)\}$ and go to Step 2.

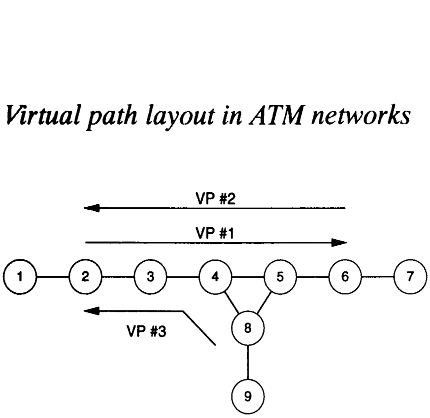

Figure 3 Applying Modified Diameter Method to a 9-node Network

5. Let $S_0 = S_0 \cup \{(s,t)\}$. VP (s,t) is added to the set of selected VPs, S_0.

As an example, consider the same network from Figure 1. Using the Modified Diameter algorithm, VPs are placed at $(2,6)$, $(6,2)$, and $(8,2)$ (Figure 3). Note that in most cases, there is not a node which is exactly a distance aD_{max} from the end-node of $\mathcal{P}$. Therefore, we choose the node that is nearest to that position.

In summary, for a general network with a uniform traffic distribution, the above algorithm can be used to assign VPs. However, when the traffic distribution is non-uniform, it is difficult to determine the optimal placement of a VP in a 1-D network. For simplicity, under non-uniform traffic, the VP is placed at the end-nodes of $\mathcal{P}$, i.e., the Diameter algorithm in Section 3.1 is used.

Both the Diameter and Modified Diameter methods were simulated over a variety of networks. Simulation results are presented in Section 4. We also simulated the Clustering method (Chlamtac, 1993) to compare its performance with the Diameter method.

The Clustering algorithm was proposed by Chlamtac, Faragó, and Zhang (Chlamtac, 1993) (Chlamtac, 1994). In the Clustering method, VPs are chosen which are far away from each other. In each iteration, a new VP is selected which is farthest away from existing VPs. In what sense is a VP *far away* from another VP? Define $D(x,y)$ as the *distance* between node-pairs $x = (i,j)$ and $y = (s,t)$. Let

$$D(x,y) = \rho_{ij}\left[d(i,s) + C_{VP}(s,t) + d(t,j)\right] \tag{7}$$

where $d(a,b)$ is the minimum physical hop distance from node a to node b and $C_{VP}(s,t)$ is the cost of a connection using VP (s,t). VPs can be selected that are far apart according to Eqn. 7 (see (Chlamtac, 1993) for details).

3.3 Computational Complexity Analysis

In this section, we analyze the computational complexity of the Diameter method and the Clustering algorithm.

In Step 2 of the Diameter algorithm, finding the node-pair that is farthest apart consists of first finding the distance between every node-pair. The time taken by Warshall-Floyd's shortest path algorithm to compute these distances is $O(N^3)$,† where N is the number of nodes in the network. After finding the VP terminator-pair, computing the actual route using Dijkstra's method requires $O(N^2)$. Therefore, each VP requires $O(N^3)$ computations. Since we are assigning M VPs, the overall complexity of the Diameter method is $O(N^3 M)$.

† A function $f(n)$ is $O(g(n))$ if $\exists c > 0$ such that $f(n) < cg(n)$ for all n sufficiently large.

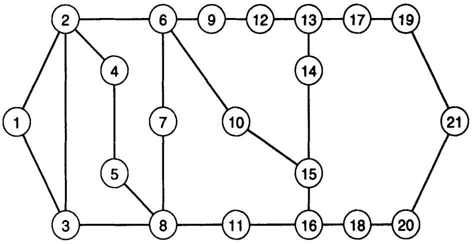

Figure 4 21-node 26-link ARPA2 Network

The computation complexity of the Clustering method is as follows. For the distance measure given in Eqn. 7, it can be shown that the complexity of the Clustering algorithm is $O(N^3)$.

From the above analysis, the Clustering algorithm has a lower computational complexity. However, in the following simulations, we show that the Diameter method performs better than the Clustering algorithm in reducing the average connection cost. If we assume that the time a VP network remains fixed is in the order of hours, the better performance the Diameter method outweighs the lower computational cost of the Clustering algorithm.

4 SIMULATION RESULTS

In this section, simulation results are presented to compare the performance of the Diameter method with the Clustering algorithm. The VP assignment algorithms are applied to some general networks under uniform and non-uniform traffic. The average connection cost, $\bar{C}$, is found as a function of the number of VPs, M, assigned to the network. In addition, we compare the algorithms' sensitivity to link failures.

4.1 Uniform Traffic

A uniform traffic distribution is defined as follows. The amount of traffic from node i to j is $\rho_{ij} = 1$ for all node-pairs. In the following simulations, the VP cost parameters are $\beta = 0.9$ and $\alpha = 0.1$, and the number of VPs per link is unrestricted.

Results for a 21-node 26-link ARPA2 network (Figure 4) is given in Figure 5. We see that the Diameter method performs as well or better than the Clustering method for all M. $\bar{C}$ becomes smaller as more VPs are assigned to the network. However, successive VPs decrease $\bar{C}$ by a smaller and smaller amount.

Figure 5 also shows the performance when VPs are assigned using the Modified Diameter method as opposed to the Diameter method. For small M, the Modified Diameter method gives better performance over the Diameter method. However, for large M, there is very little difference between the two methods. In some cases (e.g.,. $M = 10$), the Diameter method actually performs better than the modified method. This occurs because the Modified Diameter method is optimal only in the sense of establishing a single VP in a 1-D network.

From these results the Diameter method performs better than the Clustering algorithm under a uniform traffic distribution (refer to (Wong, 1995) for more results). Note that we do not guarantee that the Diameter method will

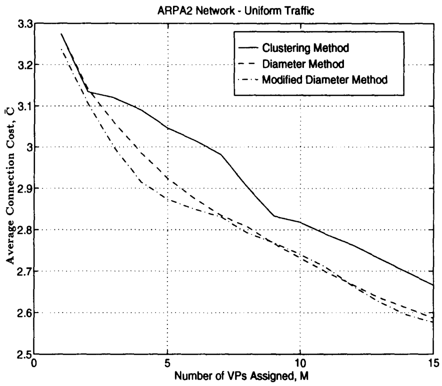

Figure 5 ARPA2 Network under Uniform Traffic.

always perform better than the Clustering method. For some networks, the Clustering algorithm gave a lower average connection cost.

The benefits of using the Modified Diameter method over the Diameter method is more pronounced in large networks. In a small 1-D network, there is little difference between assigning a VP optimally as opposed to end-to-end. For example, the benefit of an optimal VP assignment is much larger in a 20-node network than in a 3-node network. Therefore, in the ARPA2 network, where the average distance between two nodes is small, the performance is approximately the same for the two methods. However, for a large network, where the typical distance between two nodes is larger, the benefits of the Modified Diameter algorithm is emphasized (see (Wong, 1995)).

4.2 Non-Uniform Traffic

In this section, we compare the average connection cost of the two VP assignment algorithms given a non-uniform traffic distribution. We define a non-uniform traffic distribution as follows. The amount of traffic from node i to node j is:

$$\rho_{ij} = U_{ij} \text{ for all } i, j \in V$$

where U_{ij} is a random variable uniformly distributed over $[0, 1]$. Note that the amount of traffic of each node-pair is chosen independently of the other node-pairs. Also, there is no restriction on the number of VPs that can use a particular link.

A number of non-uniform traffic distributions were generated. For a given network, the average connection cost was obtained for each of these traffic distributions. The mean and 90% confidence interval of $\bar{C}$ are shown in the following figures.

Figure 6 show results for the 21-node 26-link ARPA2 network. As in the uniform traffic case, the Diameter method performs better than the Clustering method.

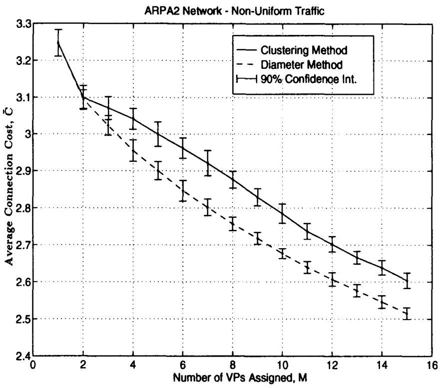

Figure 6 ARPA2 Network under Non-Uniform Traffic.

4.3 Limit on Number of VPs per Link

In the previous simulations, we have not limited the number of virtual paths that can use a physical link. In this section, we assume that the number of virtual paths that can use a physical link must be less than a prespecified limit. This limit is likely to be imposed by the limited number of VP addresses available. The limit may also be imposed to reduce the sensitivity of the VP layout to a link failure. We compare the performance only the Diameter algorithm has this restriction. The Clustering method can assign any number of VPs on a link.

In the following, we apply both algorithms to a number of 40-node 80-link networks. The maximum link load, L, is found from the resulting VP networks. These simulations are performed under uniform traffic and without restriction on the number of VPs per link. In Figure 7, the mean and 90% confidence interval of L is found as a function of M, the number of VPs assigned to the network.

As the number of VPs (M) increases, L also increases. In general, using the Clustering method results in a VP layout that has a lower maximum link load compared to the unrestricted Diameter method, and hence is less sensitive to link failures. This occurs because the Clustering algorithm assigns VPs which are far apart from each other. Therefore, the VP routes are less likely to have common links which contribute to a higher link load.

In the following simulations, we limit the maximum number of VPs that can use a link. Let $\eta_k = n$ for all links. That is, each physical link can support a maximum of n VPs. Note that this restriction applies only to the Diameter algorithm. We can extend this restriction to the Clustering method by modifying its algorithm. However, in the simulations, this was not done.

Figure 8 is a graph of $\bar{C}$ vs. n for a 21-node 26-link ARPA2 network under a non-uniform traffic distribution using the Diameter method. The number of VPs assigned is $M = 5, 10, 15$. As n decreases, the average connection cost increases.

When we restrict the number of VPs per link, VPs that previously used a certain path must be routed to another path that is possibly longer. This increases the VP cost and hence the average connection cost may also increase. More importantly, the VP restriction may actually prevent a VP from being established if no feasible route is found. The Diameter algorithm is forced to choose a less beneficial VP, resulting in a higher average cost.

For comparison, the graph also shows $\bar{C}$ and L using the Clustering method (shown by the 'o' on the graph). When

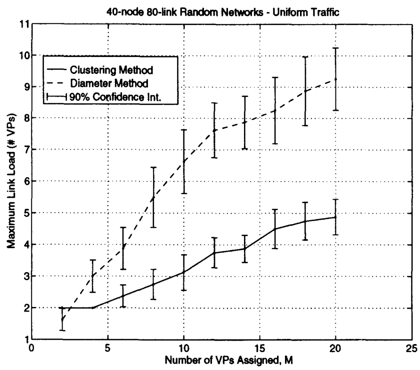

Figure 7 40-node 80-link Networks with Limited # VPs per link.

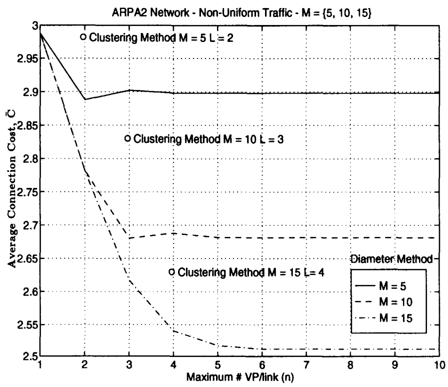

Figure 8 ARPA2 Network under Non-Uniform Traffic.

the VP assignment is unconstrained (n large), the Diameter method performs better than the Clustering method. However, even in some instances when the Diameter algorithm is restricted (n small), it still performs better than the Clustering method. For example, consider the $M = 15$ results. The VP assignment using the Clustering algorithm gives an average connection cost of $\bar{C} = 2.63$ and a maximum link load of $L = 4$. With $n = 3$, the average cost, using the Diameter method, is $\bar{C} = 2.62$. In this case, the Diameter method provides better performance in both average connection cost and link failure sensitivity.

In selecting the location of a VP, the Diameter algorithm uses the information from the currently assigned VPs. It does not consider the possibility of subsequent VPs. For example, given that we wish to establish two virtual paths in a network, the first VP is assigned under the assumption that only one VP is to be selected. A more efficient method is to consider assigning both VPs simultaneously in an attempt to minimize the average cost. However, this becomes more complex as M increases. The Diameter method trades off optimality for simplicity.

5 CONCLUSION

In this paper, we presented an algorithm for finding a system of VPs. Given a network with traffic distribution ρ_{ij}, the location of the VP terminators and the actual path between the end-nodes for each VP was found. Conceptually, the algorithm is simple and easy to implement. It has good performance in reducing the average connection cost while also limiting the number of VPs that can be routed on a physical link. The algorithm was compared with the Clustering algorithm proposed in (Chlamtac, 1993). It was shown that the our method has a slightly higher computational complexity than the Clustering method but performed better in terms of reducing the connection cost. Therefore, it is a good candidate for practical implementation.

REFERENCES

Ahn, S., Tsang, R. P. and Du, D. H. C. (1994) Virtual Path Layout Design on ATM Networks. *INFOCOM'94*. 192-200.

Burgin, J. and Dorman D. (1991) Broadband ISDN Resource Management: The Role of Virtual Paths. *IEEE Communications Magazine*, 29(1991).

Cheng, K.-T. and Lin, Y.-S. (1994) On the Joint Virtual Path Assignment and Virtual Circuit Routing Problem in ATM Networks. *GLOBECOM'94*. 777-782.

Chlamtac, I., Faragó, A. and Zhang, T. (1993) How to Establish and Utilize Virtual Paths in ATM Networks. *Proceedings of IEEE ICC'93*. 1368-1372.

Chlamtac, I., Faragó, A. and Zhang, T. (1994) Optimizing the System of Virtual Paths. *IEEE/ACM Transactions on Networking*. 581-586.

Murakami, K. and Kim, H. S. (1994) Near-Optimal Virtual Path Routing for Survivable ATM Networks. *INFOCOM'94*. 208-215.

Ohta, S., Sato, K. J. and Tokizawa, I. (1988) A Dynamically Controllable ATM Transport Network Based on Virtual Path Concept. *GLOBECOM '88*. 1272–1276.

Sato, K. I., Ohta, S. and Tokisawa, I. (1990) Broad-Band ATM Network Architecture Based on Virtual Paths. *IEEE Transactions on Communications*, Vol. 38, No. 8.

Wong, J. P. (1995) *A Diameter Based Method to Establish Virtual Paths in ATM Networks*. Master's Thesis, University of Toronto.