In Proc. MILCOM '97, IEEE Military Communications Conference,
Monterey, CA (Nov. 1997).

VIRTUAL TREE PATH MANAGEMENT FOR BATTLEFIELD NETWORKS

Adarshpal S. Sethi and Prashant Ramarao
Department of Computer and Information Sciences
University of Delaware, Newark, DE 19716
{*sethi, prashant*}*@cis.udel.edu*

ABSTRACT

Virtual Trees are source-rooted trees in ATM networks in which VCs originating at the source but going to different destinations can share some of the bandwidth pre-allocated to the VT, thus providing an additional multiplexing advantage over Virtual Paths. Non-ATM protocols such as MIL-STD-188-220A can also take advantage of Virtual Trees because source concatenation provides a benefit very similar to ATM multiplexing. We describe the results of a simulation study which shows that Virtual Tree configurations have almost 40% lower cost than corresponding Virtual Path configurations. Applications of this work to dynamic routing schemes for battlefield wireless networks are outlined.

Keywords: Virtual Paths, Virtual Trees, MIL-STD-188-220A, Path Management.

INTRODUCTION

The technique of Virtual Paths (VPs) is used in ATM networks to perform bandwidth allocation for Virtual Channels (VCs) and to simplify setting up of VCs in response to connection requests The VP concept provides several advantages such as allowing simpler network architectures, eliminating the need for call-by-call routing resulting in very short connection setup times, and allowing easier implementation of dynamic bandwidth allocation schemes [1]. In our previous work, we have extended this concept of Virtual Paths to propose a new technique called Virtual Trees (VTs) that can be used for bandwidth allocation in ATM networks [2], [3]. A Virtual Tree corresponds to pre-allocated bandwidth along a set of links in the network that form a tree rooted at a source node and leading to various destinations. VTs retain all the advantages of VPs but have a better performance potential because of the increased multiplexing between traffic on VCs flowing from the same source to different destinations over common links. The effect is a reduced probability of blocking (rejecting connection requests), better utilization of allocated capacity, and flexibility in determining routes when the VT is set up.

Battlefield wireless networks can exploit the advantages of Virtual Trees because of source concatenation at the data link and physical layers which has an effect very similar to that of VC multiplexing in ATM networks. For instance, the MIL-STD-188-220A protocol [4] uses source concatenation at both the data link and physical layers. Paths to various destinations that have been set up in the form of a VT are better able to take advantage of source concatenation than paths that do not use a tree structure. For this reason, our work on Virtual Trees holds great promise for application to path and bandwidth management in battlefield networks.

We start this paper by briefly describing the concept of Virtual Trees and then present the results of a simulation study that compares bandwidth allocation for VTs and VPs. We conclude by outlining the potential application of these results to battlefield network path management.

VIRTUAL TREES

A Virtual Path (VP) is commonly defined as a bundle of Virtual Channels (VCs) delimited by two Virtual Path terminators. Multiple VCs may exist simultaneously within a VP between a given source-destination pair; these VCs are allowed to share the bandwidth pre-allocated to the VP they belong to.

Prepared through collaborative participation in the Advanced Telecommunications/Information Distribution Research Program (ATIRP) Consortium sponsored by the U.S. Army Research Laboratory under the Federated Laboratory Program Cooperative Agreement DAAL01-96-2-0002.

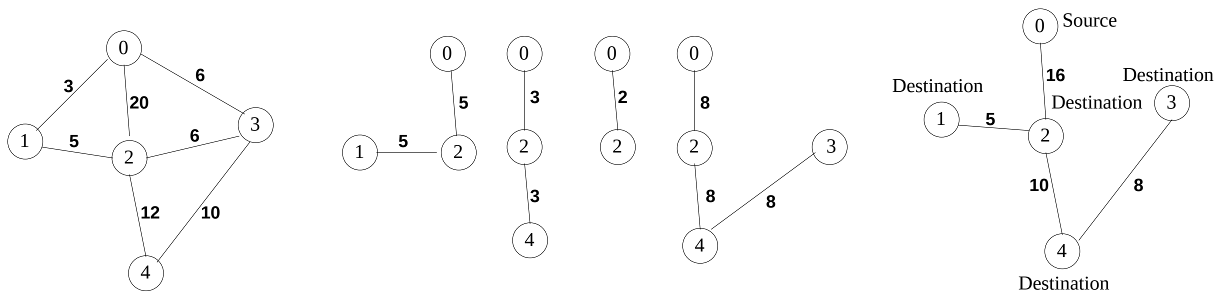

(a) (b) (c)

Fig. 1. (a) A computer network G. Numbers along lines show link capacity. (b) A set of Virtual Paths in G originating at source 0 with numbers indicating bandwidth allocated to VP. (c) A Virtual Tree VT_0 from G rooted at node 0.

We have proposed a new technique called Virtual Trees (VTs) [3], [2] in which pre-defined bandwidth is available for use by a given source node, just as in VPs. Whereas in VPs, the path and the corresponding bandwidth is defined separately for each destination, in VTs some of the links and their bandwidths could be shared by VCs going to different destinations. This takes the form of a tree rooted at the source node whence the name Virtual Tree. In a VP, different VCs may share the bandwidth of the VP resulting in better multiplexing of the traffic than if the network allocated bandwidth individually for each VC [5]. However, multiplexing in VPs is limited to traffic flowing between the same source-destination pair. In Virtual Trees, the multiplexing effect of VPs is enhanced by permitting sharing of bandwidth among traffic flowing to different destinations although it is still limited to traffic originating at a single source. The effect is a reduced probability of blocking (rejecting connection requests), better utilization of allocated capacity, and flexibility in determining routes when the VT is set up. These effects are obtained by preserving all the major advantages of VPs with only the slight additional cost of a somewhat more complex call admission algorithm at the source node.

The concept of a VT and its relationship with the paths for source-destination pairs is depicted in Figure 1. Figure 1(a) shows a network with five nodes connected by links with specified bandwidths. The traffic from each source-destination pair is transmitted along the route selected for each one of them. Part (b) shows a set of VPs for the source-destination pairs (0-1), (0-2), (0-3), and (0-4) with a bandwidth allocated to each VP based on predicted traffic patterns. Part (c) shows a VT rooted at the source node 0 for the same set of source-destination pairs. Bandwidth is now allocated to each link of the VT as shown. The total bandwidth allocated to a VT's link is the equivalent bandwidth of all the traffic multiplexed in it which is likely to be smaller than the sum of the individual bandwidths used by the VPs.

BANDWIDTH ALLOCATION FOR VIRTUAL TREES

Given a network topology with various physical link capacities, and a predicted traffic pattern between various source-destination pairs, we need algorithms to construct VT configurations (sets of VTs that can carry the given traffic). Finding such configurations that are optimal (minimize total network cost) is an NP-Hard problem [6]. We have developed two approximate optimization algorithms for finding VT and VP configurations that, while not yielding an optimal solution, provide good solutions in a reasonable amount of time [7]. The two algorithms named *Graph Reduction* and *Underloaded Neighbors* both start with shortest path configurations and modify them in different ways to try and satisfy link capacity constraints. Both algorithms start by constructing *Unconstrained* paths in which link capacities are not taken into account during the process of bandwidth allocation, but then differ in the way in which link capacity constraints are used to modify the paths. A description of the algorithms is available in [7] and is omitted here for reasons of space; instead, we present the results of a simulation study comparing the performance of the algorithms on VTs and VPs with that of the unconstrained algorithm.

We simulated the unconstrained and the two con-

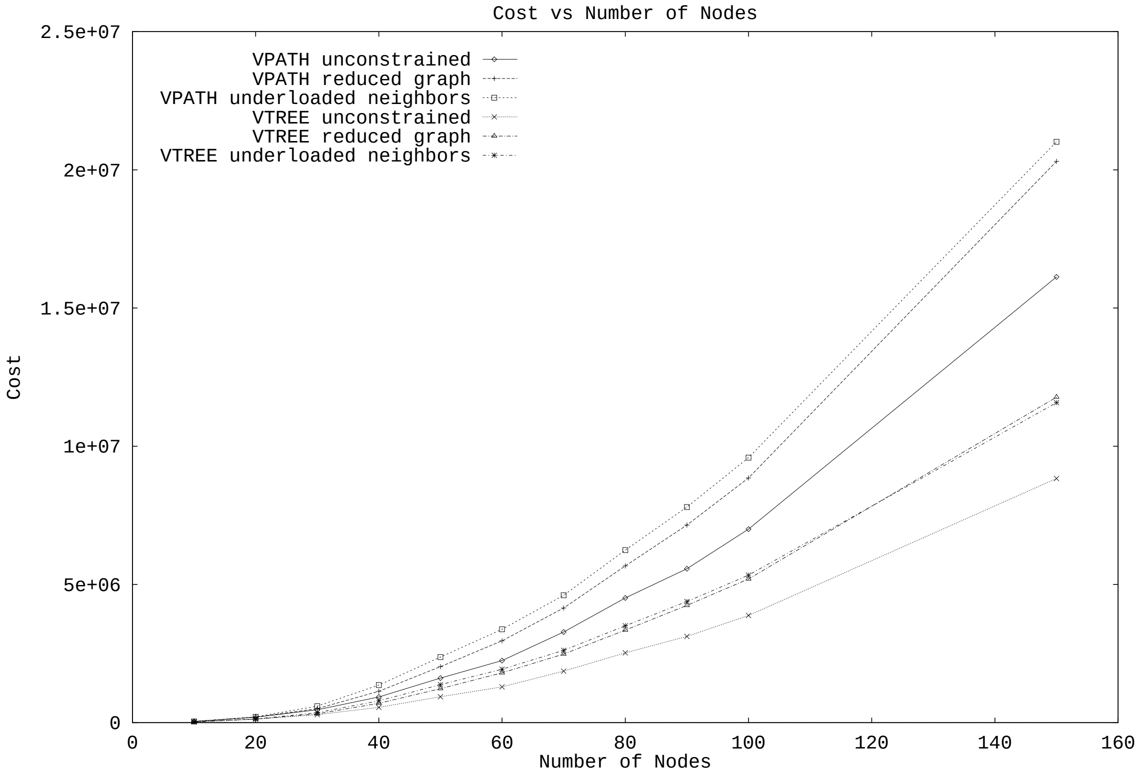

Fig. 2. Simulation results for Approximate Algorithms for VTs and VPs.

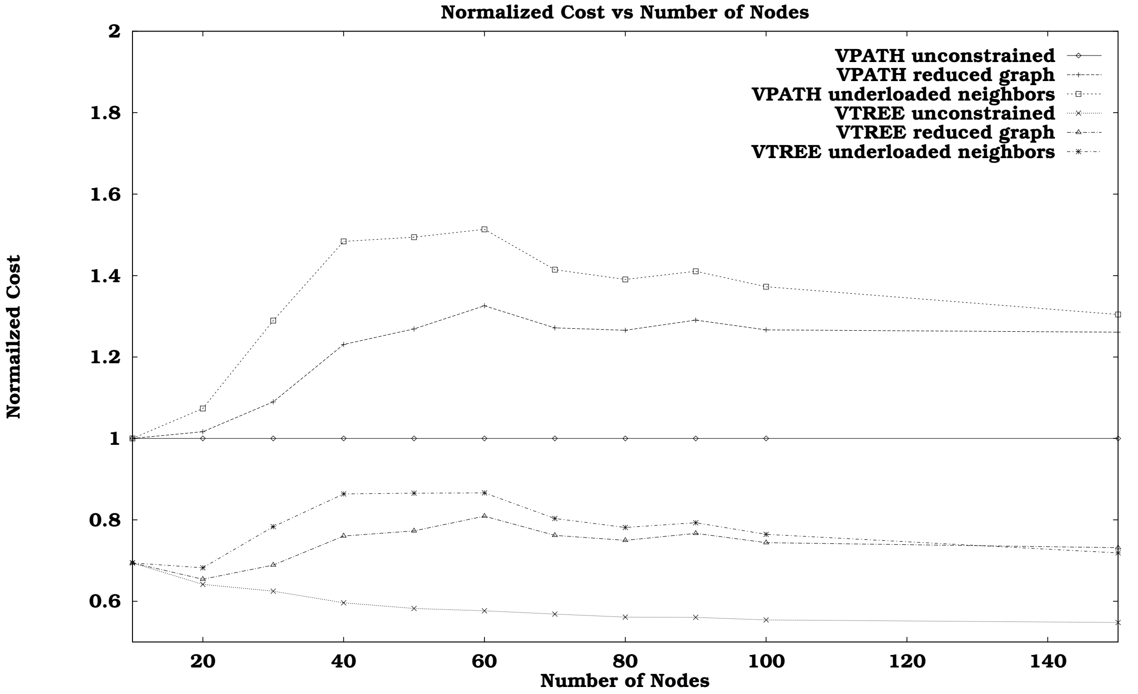

Fig. 3. Normalized Costs vs Number of Nodes for VTs and VPs.

strained algorithms for both VPs and VTs on a set of randomly generated network topologies satisfying certain specified constraints. Specifically, the number of nodes, N, in the network was varied from 10 to 200. For each size, the links were generated randomly satisfying minimum and maximum degree constraints, such that the number of links was of the order $NlogN$ and the network was connected. Each link was assigned the same cost and the same bandwidth. Also, equal traffic was assumed between all source-destination pairs. The total cost for both VP and VT configurations was computed with each of the three algorithms (the unconstrained and the two constrained). These costs were then averaged over multiple runs with independently generated networks for each network size to factor out statistical variations. We also computed the average normalized cost for both VPs and VTs, where the normalized cost for each network is the cost of the VP or VT configuration for that network obtained with the appropriate constraining algorithm divided by the Unconstrained VP cost for that network.

Figure 2 shows the results of this experiment. This figure shows the variation of the average cost with number of nodes in the network for both VP and VT configurations in the unconstrained case as well as in the two constrained cases. As expected, the average cost increases as the network size increases. This is because, although the traffic generated between each source-destination pair is constant, the number of such pairs increases at a rate that is order N^2 where N is the number of nodes. Also, the constrained costs are higher than the corresponding unconstrained cost for each of the VP and VT cases. However, VTs give a lower cost than VPs in all cases.

Figure 3 shows the variation of normalized cost with number of nodes. Since all costs are normalized by the unconstrained VP cost, the cost for the Unconstrained VP configurations is the constant 1.0 for all numbers of nodes. This figure allows us to quantify the improvement obtained by each algorithm over the Unconstrained VP configuration. The normalized cost for the Unconstrained VT configurations initially decreases as the number of nodes is increased because small network sizes provide less opportunity to multiplex multiple paths over common links. But with more than 70 nodes, this factor does not play an important role and the normalized cost of VTs settles down to slightly less than 0.6. This represents a more than 40% improvement over VPs. When capacity constraints are introduced, the cost of the resulting configurations is obviously more than the cost of the corresponding unconstrained configurations for both VPs and VTs. It appears that in the long run the normalized constrained costs stabilize at around 0.75 for VTs and around 1.25 for VPs. This again represents a 40% improvement for VTs over VPs.

APPLICATION TO WIRELESS BATTLEFIELD NETWORKS

Recently, Virtual Trees have been proposed to be used in two different contexts in wireless networks for managing handoff between mobile terminals and base stations. Acampora and Naghshineh [8] use a tree called a Virtual Connection Tree to modify existing VCs when a mobile node is handed off to a new base station. However, their tree is not rooted at the source like our tree is and also uses distinct paths from the root to the base stations, thus not taking advantage of multiplexing over overlapping links. Veeraraghavan et al. [9] use a Virtual Sink Tree to find paths to a given destination. A Sink Tree is rooted at the destination instead of the source, and takes advantage of multiplexing over common links, but VCI setup is slower.

Since a tactical battlefield network will not have fixed base stations to manage handoff, but is instead likely to have peer-to-peer communications, the main use of Virtual Trees in battlefield networks will be to manage the setup of end-to-end paths between nodes that wish to communicate with each other. These networks are not likely to use ATM technology, so the multiplexing of cells between VCs cannot be exploited in the way that VPs and VTs have done in ATM networks. However, there are similar advantages in wireless networks of using source concatenation of packets flowing to different destinations. The MIL-STD-188-220A protocol suite [4] being developed for Combat Net Radio (CNR) for digitized battlefields uses source concatenation at both the data link and physical layers. Because of the large net access delays, low bandwidths, and random access protocols used for media access, it makes sense to transmit multiple data link PDUs within a single transmission when a node gains access to the channel. For this reason, it is advantageous to set up the network layer paths in the form of trees to derive maximum benefit from shared links in the paths leading from the same source to different destinations. The algorithms developed by us to set up Virtual Tree configurations can be used for this purpose.

Another aspect of tactical battlefield networks is the dynamically changing configuration due to both mobility of the nodes and enemy action. While an initial VT configuration can be used by the nodes at start-up, this will rapidly become sub-optimal as the underlying network topology changes. We propose to use a partially decentralized approach to handle this problem. Whenever a link in the network goes down because it is no longer available, or when a new link is detected, an indication is provided by the data link layer to the network layer that results in a topology update. With our approach, each node keeps track of the network topology in its immediate vicinity and makes local adjustments to its own VTs as the topology changes. Each subnet or a small number of subnets collectively have a designated path manager node that keeps track of subnet topology. Periodically, the path manager runs an incremental algorithm to generate new VT configurations and updates the configurations being used by all nodes within its domain. The use of a path manager prevents congestion hot spots from arising, a possibility if the nodes act in a purely distributed manner without any coordination. On the other hand, the partially decentralized nature of the algorithm that allows the nodes to make changes to their own VTs will ensure that a failure of the path manager does not result in a total breakdown of communication.

CONCLUSIONS

In summary, we have shown that the use of Virtual Trees can give significant cost and performance advantages over the use of Virtual Paths for bandwidth allocation and management. In non-ATM wireless networks, similar advantages can be obtained by using source concatenation of data link and physical layer frames. Because of these advantages, the use of Virtual Trees can result in better algorithms for dynamic route allocation and path management using tree structures in networks with changing topologies such as tactical battlefield networks.

The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory or the U.S. Government.

REFERENCES

[1] Y. Sato and K. Sato. Virtual Path and link capacity design for ATM networks. *IEEE Journal on Selected Areas in Communications*, 9(1):104–111, January 1991.

[2] A.S. Sethi. A model for Virtual Tree bandwidth allocation in ATM networks. In *Proc. Infocom '95, 14th Annual Joint Conference of the IEEE Computer and Communications Societies*, pages 1222–1229, Boston, MA, 1995.

[3] A.S. Sethi and A. Mock. Virtual Trees - a new technique for bandwidth allocation in ATM networks. In *Proc. Second Intl. Conf. on Telecommunication Systems, Modeling and Analysis*, pages 113–119, Nashville,TN, March 1994.

[4] Military Standard - Interoperability Standard for Digital Message Device Subsystems (MIL-STD 188-220A). July 1996.

[5] J.A.S. Monteiro, M. Gerla, and L. Fratta. Statistical multiplexing in ATM networks. *Performance Evaluation*, 12(3):157–167, June 1991.

[6] I. Chlamtac, A. Farago, and T. Zhang. Optimizing the system of Virtual Paths. *IEEE/ACM Transactions on Networking*, 2(6):581–587, December 1994.

[7] A.S. Sethi and P. Ramarao. Bandwidth allocation algorithms for Virtual Trees. Technical Report 97-09, Dept. of Computer & Information Sciences, University of Delawar e, Newark, DE, December 1996.

[8] A.S. Acampora and M. Naghshineh. An architecture and methodology for mobile-executed handoff in cellular ATM networks. *IEEE Journal on Selected Areas in Communications*, 12(8):1365–1375, October 1994.

[9] M. Veeraraghavan, M. Karol, and K.Y. Eng. Implementation and analysis of connection-management procedures in a wireless ATM LAN. In *Proc. IEEE International Conference on Universal Personal Communications*, 1996.