An Architecture for Broadband Virtual Networks Under Customer Control

Mun Choon Chan, Hisaya Hadama* and Rolf Stadler

Center for Telecommunications Research
Columbia University, New York, N.Y., USA
{mcchan, stadler}@ctr.columbia.edu

** NTT Optical Systems Laboratories*
1-2356 Take Yokosuka Kanagawa 238-03 Japan
hadama@exa.onlab.ntt.jp

Abstract

Emerging ATM-based Virtual Private Network (VPN) services offer customers a flexible way to interconnect Customer Premises Networks (CPNs) via high-speed links. Compared with traditional leased lines, these services allow for rapid provisioning of VPN bandwidth through cooperative control between customer and provider. Customers can dynamically renegotiate the VPN bandwidth according to their current needs, paying only for the resources they actually use. In order to meet the various requirements and demands of different classes of VPN customers, a VPN provider must provide customers with the flexibility to choose their own control schemes and objectives.

The focus of this paper is on enhancing the customer's capability of controlling a VPN. First, we propose a new scheme for a broadband VPN service, which is based on the *Virtual Path Group* (VPG) concept. In our scheme, the customer performs VP control operations without interacting with the VPN provider, thus enabling the following merits: (1) the customer can share bandwidth among VPs that traverse the same physical network link in the provider's domain, thus using the VPN bandwidth more efficiently; (2) customers can perform VP control operations according to their own requirements and control objectives. Second, we outline an architecture for a *customer-operated control system*, which utilizes a VPG-based VPN service. The system is structured into three layers of control, which execute on different time scales. The functionalities of these layers are call processing, VP control, and VPN control, respectively. Finally, we evaluate the effectiveness of the control system, with respect to VP control.

Keywords: Virtual Private Networks, Enterprise Networking, Network Services, Network Architectures, Broadband Networks, Network Control

This work was supported by the Department of the Air Force under contract F30602-94-C-0150. It was conducted using the resources of the Cornell Theory Center (SP2 parallel machine).

0-7803-2518-4/96 $5.00 © 1996 IEEE

Broadband Virtual Private Networks

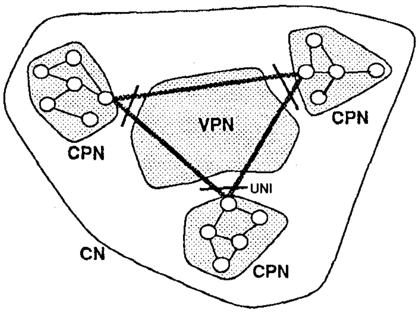

- Virtual Private Network (VPN) service offered by public network provider or third party provider
- Customer Premises Networks (CPNs) can be interconnected using a VPN service
- Customer Network (CN) combines all CPNs and the VPN into a virtual enterprise network
- Broadband VPN services
  - allow for rapid provisioning
  - provide customer-provider cooperative control of VPN bandwidth
  - are accessed via UNI interfaces

A broadband virtual private network (VPN) is a service that provides transparent broadband transmission capability between islands of customer premises networks (CPNs). Transparency refers to both call processing and end-to-end quality-of-service (QOS) guarantees. A VPN is a central building block for constructing a global customer network (CN) which interconnects CPNs.

Service providers are beginning to offer broadband VPN services using VP-based ATM transport networks [ATS93]. These services replace today's leased lines based on STM (SDH/SONET) and offer flexibility by allowing a customer to dynamically request adjustments in the VPN capacity from the VPN provider. Since networks typically exhibit a dynamic traffic pattern, such a technique of rapid provisioning will result in lower cost for the customer, because pricing is expected be based on the VPN capacity per time interval allocated to the customer network. A VPN is accessed via common user-network physical interfaces (UNIs).

VPN services may be offered by third party service providers, since VPN traffic may be carried over several different public networks [SPN93].

To utilize a VPN service, customers operate a control system that interfaces their traffic control and management systems with the VPN service. This system can execute control functions, such as call processing and call resource management, without interacting with the provider. Customer operated control benefits the customer in terms of fast execution of control operations and flexibility in choosing QOS objectives, as well as control schemes independent of the provider. In this paper, we suggest a new broadband VPN service which gives the customer more levels of control than other existing or proposed VPN services/schemes. We outline the architecture of a customer control system for this VPN service, without describing in detail the provider control system.

Approaches to Broadband VPNs

VPN schemes level of bandwidth sharing	ATM leased line [ATS93]	VPN based on VC XCs [FGC95]	VPG-based VPN our scheme
- bandwidth of a physical link - shared by traffic from different customers	VP capacity control with customer-provider negotiation [SAY95]	——	VPG bandwidth control with customer-provider negotiation
- dedicated bandwidth to a customer - shared by different source-destination (s-d) traffic	——	call by call VC setup/release by customer	VP capacity control by customer
- dedicated bandwidth to a s-d pair - shared by the same s-d traffic	call by call VC setup/release by customer		call by call VC setup/release by customer

The above figure compares different schemes for broadband VPNs in terms of the possibility of sharing bandwidth on different levels. ATM leased line services allow VPs of different customers to share the bandwidth of a physical link through VP capacity control [ATS93]. To achieve this control, public network providers need mechanisms to keep the sum of the VP capacities below the capacity of the physical link. The capacity of a single VP can be shared by VCs with the same source-destination traffic through the customer's call-by-call VC setup/release procedures.

In order to allow a customer to share bandwidth among connections with traffic of different source-destinations, S. Fotedar et al. [FGC95] proposed a VPN service which adopts VCs as end-to-end logical links. Such a VC traverses one or more VPs in a VPN. Since a VP accommodates VCs with different source-destination traffic, the VP capacity can be statistically shared by traffic with different source-destination pairs, controlled by a VC admission control mechanism in the customer domain. This scheme, however, requires VC cross connects in the carrier's infrastructure, which increases service cost. Moreover, the customer performs VC admission control in a centralized manner, which does not allow for rapid call admission control, especially in a large-scale network.

In our scheme of a VPG-based VPN, the capacity of a VP is shared by VCs with the same source and destination through the customer's VC setup/release procedures. The VPG bandwidth that is allocated to a customer can be shared through VP capacity control which is performed by the customer. The VPG bandwidth can be changed through negotiation between the customer and the VPN provider. The most important feature of our scheme is that we can introduce a VP capacity control scheme that is independent of both the VPN provider's operations and VC setup/release procedures. This enables customers to optimize VP capacity control according to their own requirements and objectives.

Our Approach: A VPG-based VPN

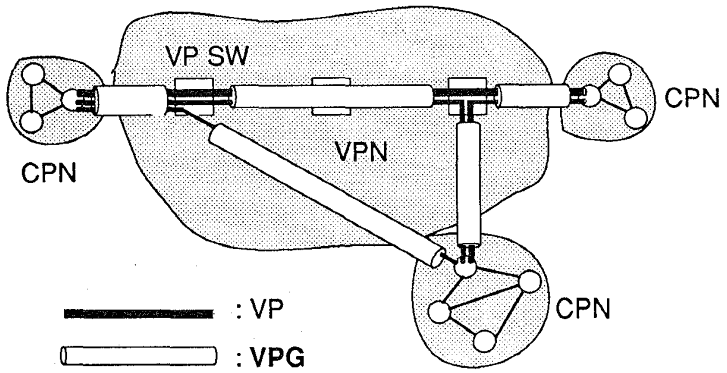

Definition of a Virtual Path Group (VPG):

- **logical link connecting VP switches in a public broadband network or its termination points**
- **has bandwidth allocated to accommodate a bundle of VPs**

Benefits: Enhanced controllability for a customer

- **customer's own VP capacity control for bandwidth sharing in a VPG**
- **customer's own VP restoration control utilizing a backup VPs**
- **customer-provider cooperative VPG capacity reallocation**

Virtual paths (VPs) provide direct logical links between Customer Premises Networks (CPNs). This ensures that VC setup procedures can be executed by the customer without interaction with the carriers. To enhance the customer's capability for VPN control, we introduce the *Virtual Path Group* (VPG) concept. A VPG is defined as a logical link within the public network provider's ATM network. The above figure shows a *VPG-based Virtual Private Network* connecting 3 CPNs. A VPG is permanently set up between two VP cross connect nodes or between a VP cross connect node and a CPN switch that acts as a customer access point for the VPN service. A VPG accommodates a bundle of VPs that interconnect customer access points. The VPN provider allocates bandwidth to a VPG, which defines the maximum total capacity for all VPs within the VPG. A VPG-based VPN consists of a set of interconnected VPGs.

In order to guarantee cell level QOS in the carrier's network, policing functions (Usage Parameter Control) are required at the entrance of each VPG. Note that there is no need for a VPG identifier in the ATM cell header, since cells are transmitted by VP cross connect nodes based on their VP identifier. Only the network management systems must know about the routes of the VPGs, their assigned bandwidth, and the VPs associated with them. VPs and VPGs are set up by the network management system of the VPN provider during the VPN configuration phase.

The VPG concept enhances the customer's capability for VP capacity control. A customer can change the VP capacities, within the limits of the VPG capacities, without affecting other customers. The VPG bandwidth can be shared by VPs with different source-destination pairs, without negotiation between the customer and the VPN provider. Customers can independently achieve the optimum balance between the resources needed for VP control and the resources needed to handle the traffic load.

Control Layers of the Customer VPN Control System

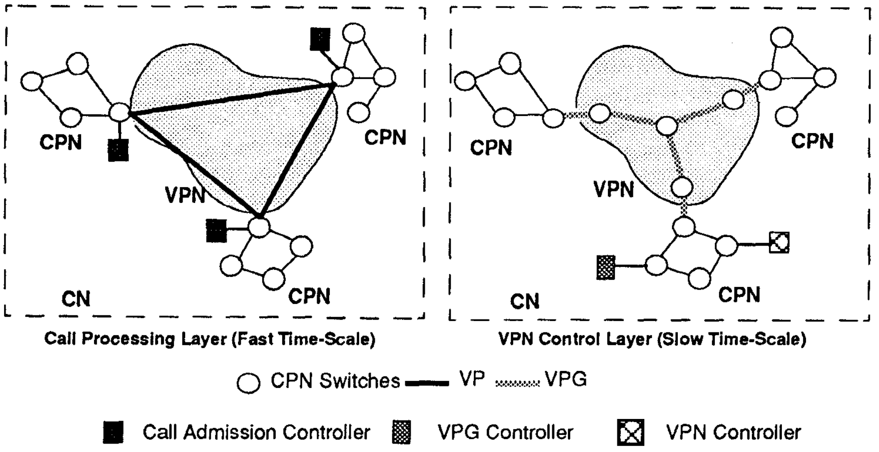
CPN
CPN
VPN
CN
CPN
CPN
CPN
VPN
CN
CPN
○ CPN Switches ▬ VP ▒ VPG
■ Call Admission Controller ▨ VPG Controller ☒ VPN Controller

Call Processing Layer (Fast Time-Scale)

VPN Control Layer (Slow Time-Scale)

Time Scales of Control		
Call Processing	Fast (< 1s)	VC Admission Control
VP Control	Medium (5s - 100s)	VP Bandwidth Allocation
VPN Control	Slow (> 100s)	VPG Bandwidth Allocation

We outline the architecture of a control system operated by the customer to utilize a VPG-based VPN service. The figure provides two views that correspond to two different layers of control in this system. The left side shows the view of the call processing system which sets up and releases calls in the customer network. This system knows about the VPs as logical network links, but has no knowledge about the VPGs. The right side shows the view of the VPN control layer, which performs adaptive VPN capacity allocation in cooperation with the VPN provider. In between these two layers is the VPG control layer, where VPG bandwidth is allocated to VPs -- in other words where VP capacity control is executed.

Controllers in different layers interact asynchronously with each other, which allows them to run on different time scales. Each of the 3 layers executes on a different time scale.

On the fastest time scale, a call admission controller, associated with a VP, decides whether a call can be admitted into the VPN, based on the VP capacity and its current utilization. The admission control policy ensures that enough capacity is available, such that cell-level QOS can be guaranteed for all VCs that are accepted. Admission controllers run on the time scale of the call arrival and departure rates. On a medium time scale, VPG controllers dynamically change the amount of VPG bandwidth allocated to associated VPs. This control scheme enables customers to exploit variations in utilization among VPs that traverse the same VPG, allowing them to share bandwidth between VPs of different source-destination pairs. In order to guarantee QOS, the sum of the VP capacities must be less than or equal to the capacity of the VPG link. On the slowest time scale, the bandwidth of the VPG links are renegotiated by the VPN controller based on usage patterns, blocking constraints per VP, etc. The objective is to minimize the VPN bandwidth cost, while observing the customer's QOS requirements.

The Customer VPN Control System

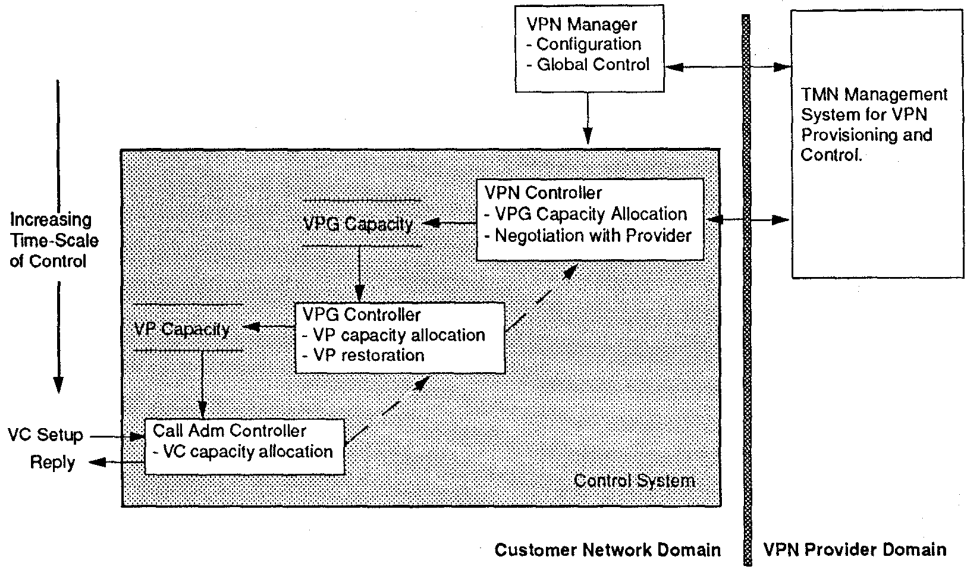

The control system is structured according to three control layers, one subsystem per layer. Interactions among the controllers can be realized in various ways. For instance, call admission controllers can send bandwidth requests to VPG controllers, triggered by a pressure function, or a VPG controller can periodically recompute the VP capacities and distribute them to call admission controllers. Also, VPG control can be executed either in a centralized or a distributed fashion. In our current implementation on a network emulation platform, there is one call admission controller for each VP and a single VPG controller and VPN controller for the whole customer network. Periodically, the VPG controller reallocates the capacity of VPG links among the VPs, using a weight function that takes into account the utilization, the offered load and the blocking constraints per VP. The control frequency of the VPG and the VPN controllers are among the parameters that can be changed by the customer management system [PAC95].

The above figure also shows the systems involved in the process of VPN provisioning. The information concerning VPG topology, VP topology and the mapping between them are exchanged during the VPN initialization phase and stored in the management systems of both the customer and provider. Knowledge about VPG-VP mapping is also required in the provider's control system which performs UPC per VPG. The use of VPGs has no influence on cell switching and transmission; this information is not needed in the transport network.

A customer control system was implemented on a network emulator which we built on a SP2 parallel machine [CHA96]. On this platform, the functional components of the control system are executed by a parallel simulation kernel, which allows us to specify the estimated communication delays among controllers and the processing delays of operations. The platform thus allows us to approximate the behavior and performance of a system built according to our architecture. We can visualize the state of the system in real-time and conduct performance measurements. Experimental results from this platform are presented later.

Evaluation of VP Control for a Single VPG

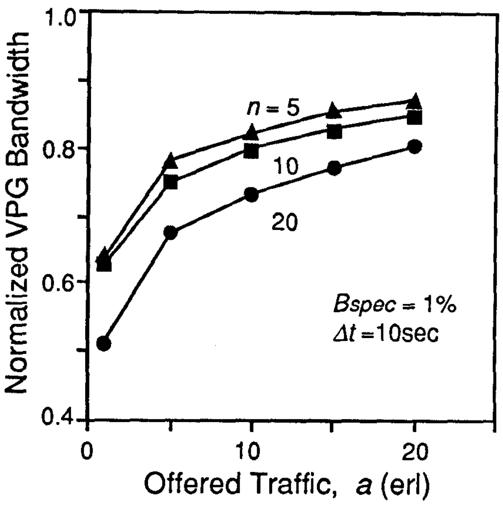

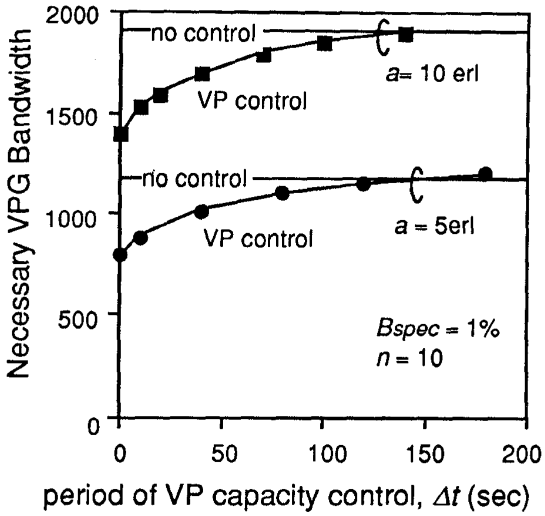

The customer's VP capacity control scheme was evaluated for the simple case which involves a single VPG. We determined the necessary VPG bandwidth to satisfy a specific blocking probability, using a periodic VP capacity control scheme. In each period, the VP capacities were recomputed following a simple heuristic, which takes into account the current VP utilization and the offered load.

For the evaluation, the network traffic was composed of two classes with different bandwidth requirements. A class 1 call needs one unit of bandwidth, while a class 2 call requires 10 units of bandwidth. The holding time of the calls of both classes was exponentially distributed with a mean of 100 seconds, and call arrivals were modeled as Poisson processes. We varied three parameters in the experiments: the number of VPs in the VPG link (*n*), the offered load (*a*), and the control period for changing the VP capacities (Δt). All VPs in the VPG link experienced the same offered load.

To compare the effectiveness of the control scheme, we define the normalized VPG capacity as the ratio of the VPG capacity needed to attain a specific call blocking probability (*Bspec*) with VP capacity control over that without control (fixed VP capacities). The figure on the left side shows normalized VPG capacities for 5, 10 and 20 VPs. It indicates that the control effect is especially large when the offered traffic per VP is small and the number of VPs multiplexed in the VPG is large. For example, when the offered traffic is 1 (erl) and the number of VPs is 10, we can reduce the necessary bandwidth by 37.3% for a control period of 10 seconds. The figure on the right side shows the necessary VPG bandwidth relative to the control period Δt. The control effect rapidly reduces as Δt increases. This means that the VP capacities can not effectively follow the traffic fluctuations when Δt is large. The figure indicates that VP capacity control for constant Poisson traffic is effective up to a control period of 100 seconds.

Evaluation for a Network Scenario (1)

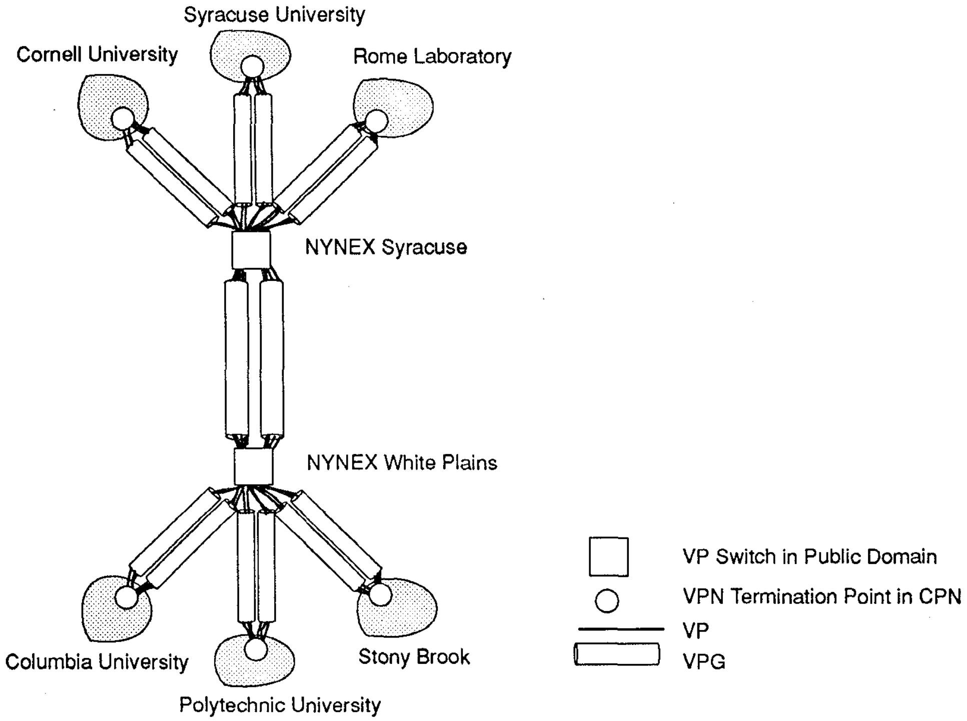

The design of the customer control architecture was evaluated in a scenario based on the topology of the NYNET testbed, an ATM WAN that connects various research laboratories in New York State. In this scenario, a VPN service interconnects 6 CPNs. The VPN contains 14 unidirectional VPGs which support 30 unidirectional VPs, connecting the 6 CPNs in a full mesh topology. The two VPGs in the middle carry 9 VPs, the remaining VPGs carry 5 VPs each.

In this evaluation, we extended the network topology from the previous one which includes a single VPG to a scenario which includes a network of VPGs. In this case the customer control system runs call admission controllers in all CPNs, one per VP. A centralized VPG controller estimates the utilization of the VPs, periodically recomputes the VP capacities, and distributes them, using a simple two phase protocol. The protocol ensures that no VP capacity is set to a value smaller than the VP utilization at the time the new capacity is received by the call admission controller.

The evaluation focused on VP capacity control, and VPN control was not performed, i.e., the VPG link capacities remained constant during the course of the experiments.

Evaluation for a Network Scenario (2)

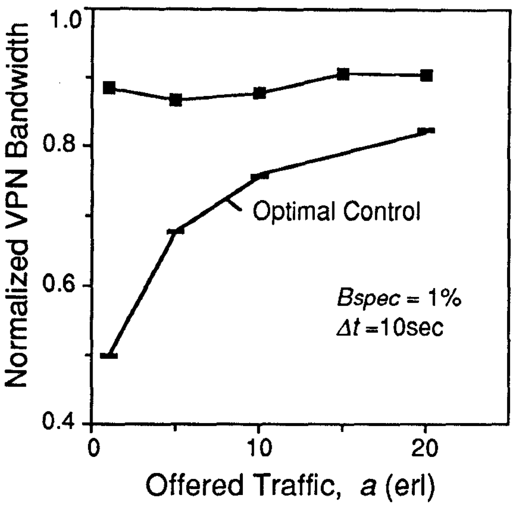

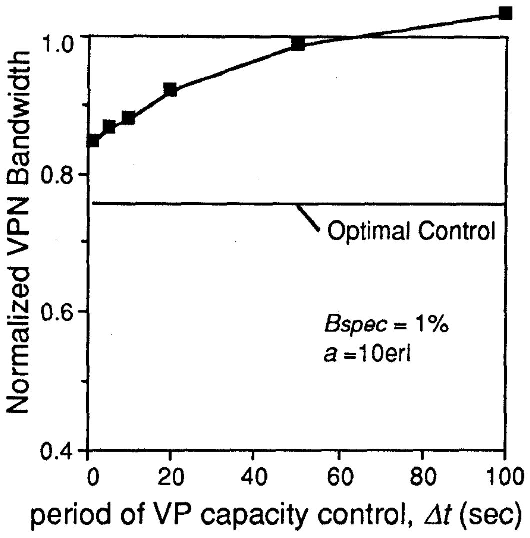

As in the previous evaluation, the network traffic is composed of two classes with different bandwidth requirements. A class 1 call needs one unit of bandwidth, while a class 2 call requires 10 units of bandwidth. The holding time of the calls of both classes is exponentially distributed with a mean of 100 seconds, and call arrivals are modeled as Poisson processes. All VPs are loaded with uniform traffic intensities.

The figure on the left side shows the necessary VPN bandwidth for different traffic loads; the figure on the right gives the necessary VPN bandwidth for different control periods. The VPN capacity is computed as the sum of the VPG capacities. The figures also contain the lower limits for VPN bandwidth, which are calculated assuming complete VPG bandwidth sharing by all calls in the VPN. They approximate the performance of an optimal control scheme.

The figures clearly indicate that the VP capacity control scheme used in this network scenario performs less effectively than the scheme applied in the one link scenario. A decrease in effectiveness can be expected, since, in this scenario, each VP traverses several VPGs, and spare bandwidth in all affected VPGs must be available in order to increase the bandwidth of a single VP.

The distance between the curves for the optimum control and our VP scheme in the left figure suggest that there is room for improving our scheme, specifically in the case where the offered traffic is low.

Note that effective control is probably not the most important benefit for customers of a VPG-based VPN service. This might be the capability to reallocate VP bandwidth according to the customers' current needs--independent of the provider, who can deny such a request or execute it much more slowly than can be done by the customer control system.

Discussion

Customer Control of VPN

- **allows for execution of operations according to customer's requirements and control objectives**
- **increases efficiency and reliability of a VPN service for a customer**

VPN service based on Virtual Path Groups provides three levels of controls

- **VC: call admission and set up by customer**
- **VP: VP bandwidth allocation by customer**
- **VPN: VPN bandwidth allocation by customer-provider cooperative control**

The customer operated VPN control system

- **is structured into three layers, according to different levels of controls**
- **layers operate on different time scales, interact asynchronously**

A promising way to realize an efficient and reliable VPN service on a multi-carrier infrastructure is to enhance the capabilities of customers to control their VPN. In order to meet the various requirements and demands of different classes of VPN customers, the VPN provider has to support customers with the flexibility of choosing their own control schemes and objectives.

We propose a new VPN service, based on the VPG concept, which allows a customer to perform three levels of control: call admission and processing, VP bandwidth allocation, and VPN bandwidth control. While call processing and VP control are performed without interaction with the VPN provider, adaptive management of the VPN bandwidth is executed by cooperative control between customer and provider.

The VPN control system, operated by the customer, is structured into three layers, reflecting the three levels of controls. The layers operate on different time scales and interact asynchronously with each other. Specifically, the VPG concept allows for medium time-scale control between the call processing layer and the VPN control layer.

[ATS93] T. Aoyama, I. Tokizawa, K. Sato, "ATM VP-Based Broadband Networks for Multimedia Services," IEEE Communications Magazine, April 1993, pp. 30-39.

[CHA96] M.C. Chan, G. Pacifici and R. Stadler, "A Platform for real-time visualization and interactive simulation of large multimedia networks", 4th Int'l Workshop on Parallel and Distributed Real-Time Systems (WPDRTS), April 15-16 1996, Honolulu, Hawaii.

[FGC95] S. Fotedar, M. Gerla, P. Crocetti, and L. Fratta, "ATM Virtual Private Networks," Communications of the ACM, vol. 38, no. 2, Feb. 1995.

[HIS89] H. Hisaya and S. Ohta, "Routing control of virtual paths in large-scale ATM-based transport networks," Trans. of IEICE, vol. J72-B-1, no.11, pp 970-978, 1989 (in Japanese).

[PAC95] G. Pacifici and R. Stadler, "Integrating resource control and performance management in multimedia networks," in Proceeding of the IEEE International Conference on Communications, Seattle, WA, June 1995.

[SAY95] T. Saydam and J.P. Gaspoz, "Object-oriented design of a VPN bandwidth management system," in IFIP/IEEE International Symposium on Integrated Network Management, Santa Barbara, California, 1995.

[SPN93] J.M. Schneider, T. Preuss, and P.S.Nielsen, "Management of Virtual Private Networks for Integrated Broadband Communication," in Proceeding of the ACM SIGCOMM '93, pp. 224-237.