

Customer Management and Control of Broadband VPN Services

M.C. Chan, A. A. Lazar and R. Stadler
Center for Telecommunications Research
Columbia University, New York, NY 10027
{mcchan, aurel, stadler}@ctr.columbia.edu

Abstract

We present an architecture for customer management and control of a broadband VPN service. The architecture is aimed at giving the VPN customer a high level of control over the traffic on the VPN, such that end-to-end requirements for the customer's enterprise network can be met. We describe how different control and management objectives can be achieved with this architecture. Its design includes a generic resource controller, which can be specialized in order to realize a large class of control schemes, following a customer's specific requirements. We have implemented a prototype of this architecture on a high-performance emulation platform. The prototype allows us to validate the management and control functionality of the customer control system and to demonstrate the performance characteristics of different realizations of the architecture.

Keywords

Virtual Private Networks, Management of Broadband Services, Customer Control, Management Architectures, Prototyping

1. INTRODUCTION

Broadband technology has the potential to change corporate networking in major ways. Broadband networks are aimed at providing quality-of-service (QOS), thus making it possible to support real-time services like voice and video communication, in addition to best-effort data delivery. Due to their ability to integrate different services on the cell-level, they provide a promising platform for distributed multimedia applications that are emerging today. Furthermore, the advent of broadband technology will enable the integration of today's separate corporate networks (voice network, data network), which often rely on different public services

(e.g., leased lines for voice traffic and LAN interconnection, frame-relay service for low-volume data exchange) into a single enterprise network, using a single Virtual Private Network (VPN).

Corporations want to control and manage their enterprise networks according to their own control objectives and management strategies. This implies that a corporate customer, using a VPN service, needs the capability to control and manage its traffic on the VPN--possibly in cooperation with the provider. For the designer of an enterprise network, the question arises, which part of the control functionality is executed in the customer's domain and which part in the provider's domain. More precisely: which functions are performed by the customer alone, which by the provider alone, and which in the form of customer-provider cooperative control.

There are strong reasons for *customer control*, i.e., for running traffic management functions in the customer domain. First, different customers pursue different control and management objectives while running their enterprise networks. For example, customer requirements concerning the traffic carried on in a VPN are very diverse with respect to supporting multimedia traffic with different performance characteristics and performance requirements. Some customers may want to operate a multiclass network with several traffic classes for both real-time and non real-time traffic; others may want to support just one class of traffic with peak rate allocation. Some might want to implement a call priority scheme which enables calls of higher priority to pre-empt those of lower priority when the network is congested; others may want to apply other control schemes in case of congestion. Providers face difficulties in their efforts to accommodate such diverse requirements. Customers who know their requirements better than the providers may be in a better position to execute control according to their objectives. Also, operations under customer control can be executed faster than those performed in cooperation with the provider, since no negotiation is required. For example, setting up connections over a VPN can be done by the customer in a distributed way, based only on local information. This allows customers to engineer or configure their traffic control systems in such a way that short connection set-up times can be achieved, which is required by some applications.

Second, customers want provider-independent control in order to meet special requirements for the enterprise network [ZER92]. For example, usage collection that permits billing at a level of detail beyond the provider's capability, such as billing at an application level, may be needed. Furthermore, the partitioning of the VPN by the customer may be required to implement sophisticated access control mechanisms, which prevent unauthorized access to certain partitions of the network. Also, automatic fall-back mechanisms may be desirable for critical applications that need high network reliability.

Third, moving the responsibility for VPN traffic management from the provider to the customer accelerates the introduction of Broadband VPN services. Specifically, public VPN services based on Constant Bit Rate (CBR) Virtual Paths (VPs) can be provided efficiently today [ATS93, FOT95]. However, such a service requires resource control by the customers, since they will be billed based on allocated bandwidth--even if they do not use it.

In this paper, we present an architecture for customer-based management and control of a broadband VPN service. We outline how different control and management objectives can be achieved with this architecture. An element of this architecture is the design of a generic resource controller, which can be specialized in order to realize a large class of control schemes, following a customer's specific requirements. Further, we present a prototype imple-

mentation of this system in a high-performance emulation environment. The prototype allows us to demonstrate performance characteristics of the customer control system and to validate the management and control functionality.

The paper is organized as follows. Section 2 describes different broadband VPN services from a customer's perspective, specifically a VPG-based VPN service, which gives the customer a high level of control. Section 3 discusses customer control and management objectives for a VPN. Section 4 presents our architecture for a customer operated control and management system for a VPG-based VPN service. Finally, Section 5 describes our experience with a prototype implementation of the architecture and the emulation platform we use for prototyping.

2. BROADBAND VPN SERVICES

A broadband virtual private network (VPN) is a service that provides broadband transmission capability between islands of customer premises networks (CPNs) (Figure 1). It is a central building block for constructing a global enterprise network (EN) which interconnects geographically separate CPNs. A VPN service involves several administrative domains: the customer domain, the domain of the VPN service provider--also called "value added service provider" (VASP)--, and one or more carrier domains [SCH93]. As a result, it is necessary to address the aspects of multi-domain management in the context of VPN service management and provisioning ([LEW95], [TSC95]). The scope of this paper is limited to the customer domain and the interaction between the customer domain and the VPN provider domain.

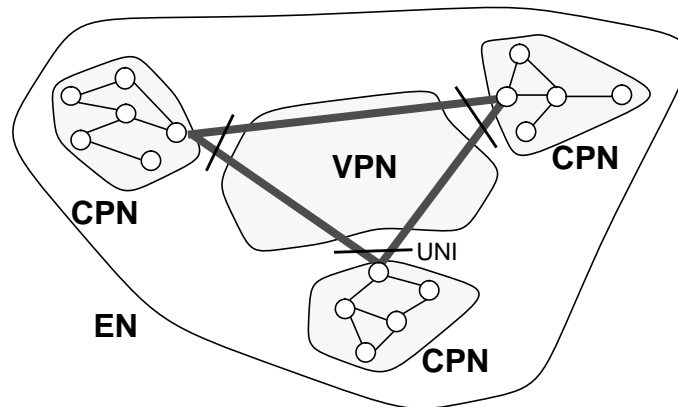


Figure 1 Customer's view of a virtual private network.

Traditionally, leased line circuits based on STM (SDH/SONET) technology have been used for providing VPN services [YAM91]. The speed of the circuit can be changed by customer-provider cooperative control. However, dynamic bandwidth adjustment for leased line circuits is inefficient and costly compared to ATM-based services, which place no restriction on the line speeds the customer can choose from [HAD94].

Service providers are beginning to offer broadband VPN services using ATM transport networks. Two common approaches are Virtual Circuit (VC)-based VPN services ([SAY95]) and VP-based VPN services [ATS93]. These services provide ATM logical links between separate CPNs. In the case of a VC-based VPN service, the customer requests a new VC from the provider for every call to be set up over the VPN. Bandwidth control and management between customer and provider is performed per VC. In the case of a VP-based VPN service, customers can perform their own call and resource control for a given VP, without negotiating with the

VPN provider. Bandwidth control and management between customer and provider is performed per VP. VC-based and VP-based VPN services replace today's leased line services. They offer customers more flexibility in dynamically requesting adjustments in the VPN capacity. 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 to be based on the VPN capacity per time interval allocated to the enterprise network. A VPN is accessed via common user-network physical interfaces (UNIs).

A Virtual Path Group (VPG)-based VPN service has been proposed to enhance customer control over the VPN [CHA96a]. The Virtual Path Group (VPG) concept has been introduced to simplify virtual path dynamic routing for rapid restoration in a carrier network [HAD89]. In a VPG-based VPN service, a VPG is defined as a logical link within the public network provider's ATM network. Figure 2 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 end-to-end 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.




Figure 2 A VPG-based virtual private network.

VPs and VPGs are set up by the network management system of the VPN provider during the VPN configuration phase. Only the network management systems must know about the routes of the VPGs, their assigned bandwidth, and the VPs associated with them. The use of VPGs has no impact on cell switching, as cells are transmitted by VP cross connect nodes based on their VP identifier. 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.

The VPG concept enhances the customer's capability for VP capacity control. It allows transparent signalling and dynamic VP bandwidth management within the customer domain. A customer can change the VP capacities, within the limits of the VPG capacities, without interacting with the provider. As a result, the VPG bandwidth can be shared by VPs with different source-destination pairs. Furthermore, customers can independently achieve the optimum balance between the resources needed for VP control and the resources needed to handle the traffic load.

3. CUSTOMER CONTROL AND MANAGEMENT OBJECTIVES

From the perspective of traffic control, the customer wants to achieve two sets of objectives. The first set relates to end-to-end QOS requirements for the traffic on the enterprise network, which translates into QOS objectives for the traffic that traverses the VPN. QOS objectives on the cell level are usually expressed in terms of bounds on end-to-end delays and error rates; on the call level QOS objectives include call blocking constraints and bounds on call set-up times. The second set relates to efficient use of VPN resources, primarily trunk bandwidth.

Efficient use of the VPN bandwidth can be achieved by exploiting statistical multiplexing at the cell- and/or the call-level. On the cell-level, multiplexing gains among calls with the same source-destination pair (with respect to. the VPN) can be achieved using the schemes described in [HYM91, ELW93]. Cell multiplexing among calls with different source-destination pairs can be performed based on the contract region concept [HYM94]. On the call-level, schemes for VP control (e.g. [OHT92]) can be used to exploit multiplexing among calls with the same source-destination pairs. Finally, the techniques described in [FOT95] and [CHA96a] can be used to multiplex calls with different source-destination pairs. Depending on the type of VPN service the provider offers, the customer can choose to implement one or more of the above multiplexing schemes in the customer control system.

In terms of managing the enterprise network, customers want the capabilities to control the bandwidth cost of the VPN service, define QOS objectives and set preferences and priorities for resource allocation to deal with congestion situations. These management objectives apply to the customer domain only and are different from customer to customer. They define the policies according to which the customer control system operates. Management capabilities can be realized by tuning controllers in the customer control system (Section 4.3). For illustration purposes, we describe below some of the management capabilities we have implemented in our prototype system.

Cost management allows the customer to define the maximum average cost of the VPN communication resources over a specific period of time. This capability is realized by setting constraints on the negotiation of VPN bandwidth between the customer and the provider. VP management allows the customer to manipulate VP bandwidth directly. Operationally, the control of the VP bandwidth can be executed either automatically by the customer control system or under direct control of the operator of the enterprise network. The operator can allocate a fixed amount of bandwidth to a VP, which must be respected by the control system. QOS and priority management operations define how calls are handled in the enterprise network. In our implementation, every call is characterized by a performance class and a priority class. The performance class of a call determines its QOS requirements. QOS management deals with managing the level of service provided to different performance classes. In particular, the customer can modify the blocking objectives of calls belonging to a performance class. The level of priority determines the relative importance of a call. In our scheme, a high priority call can pre-empt a call of lower priority in case of congestion. The customer can enable and disable priority control and can set blocking objectives for priority classes. The concepts of QOS and priority class are independent in the sense that a call that demands stringent QOS requirements can have low priority and vice versa. Finally, the above described management capabilities are orthogonal in the sense that they can be applied independently of one another.

4. A CUSTOMER CONTROL AND MANAGEMENT ARCHITECTURE

Figure 3 shows the systems involved in the provisioning and operation of a VPG-based VPN. In the provisioning phase, information concerning the VPG topology, the VP topology and the mapping between them is exchanged and stored in the management systems of the customer and the provider. Knowledge about the VPGs is also required in the provider's control system, which performs Usage Parameter Control (UPC) per VPG. The use of VPGs has no influence on cell switching and transmission, since cells are switched according to the VP identifiers in their headers. Figure 3 also shows the organization of the control system according to time-scales. The customer control system contains three classes of controllers: VP admission controller, VPG controller, and VPN controller. These controllers operate on different time-scales and run asynchronously.

We illustrate the interaction among these controllers with an example. Assume that one of the VPs experiences a sudden increase in traffic load. The VP admission controller associated with this VP admits calls as long as there is sufficient capacity. If there is not sufficient capacity available, calls are blocked. On a slower time scale, the VPG controller detects the congestion in this particular VP and attempts to allocate additional bandwidth to it. If the increase in traffic load is transient and, therefore, the demand for bandwidth drops after some time, the interaction stops here. Otherwise, if the congestion persists, the VPN controller, which runs on a slower time-scale, will request additional VPN capacity from the provider.

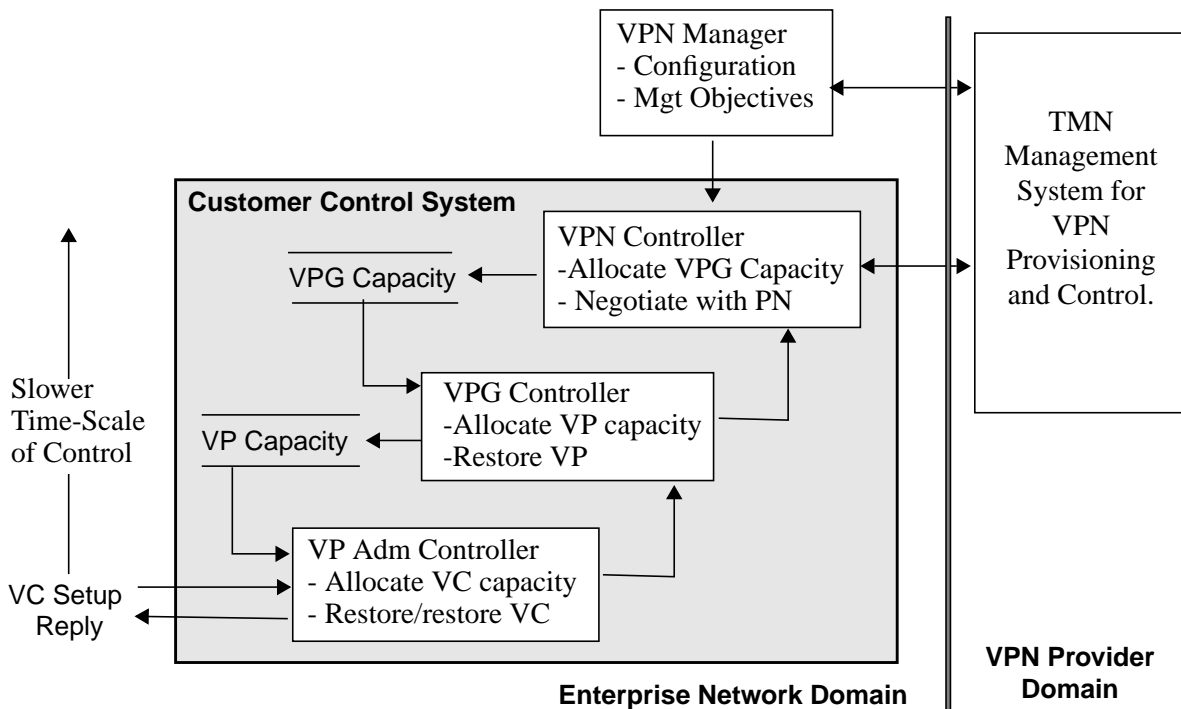


Figure 3 A functional model of the customer control system.

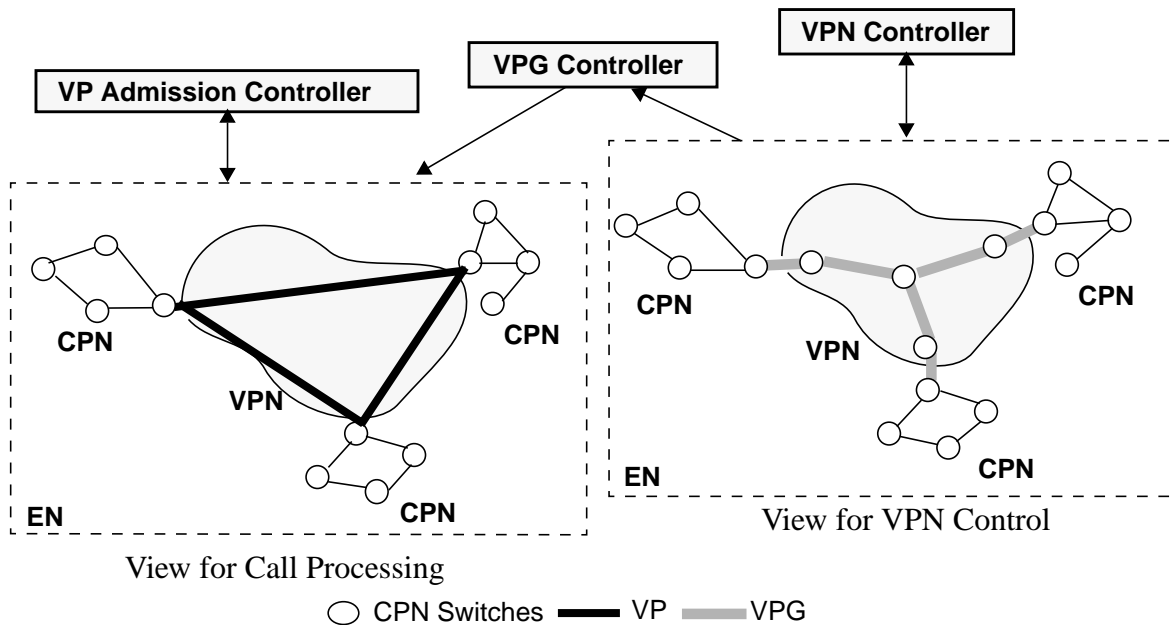


Figure 4 Network views the controllers operate on.

For the purpose of dynamic bandwidth control, a VPG-based VPN can be compared to an ATM network in which the link size can be varied. Therefore, controllers in the customer domain operate on two views of the network (Figure 4). The view on the left side of Figure 4 shows a network of end-to-end VPs which connect a set of CPNs. The view on the right shows a VPG network, which connects the same set of CPNs. The relationship between VPs and VPGs defines the mapping between both views.

The VP admission controller, which participates in call setup and release in the enterprise network, operates on the left view. The controller decides whether a call can be admitted into the VPN, based on the VP capacity, its current utilization and the admission control policy. The VP admission controller always ensures that enough capacity is available, such that cell-level QOS can be guaranteed for all calls that are accepted. The controller runs on the time scale of the call arrival and departure rates (seconds or below). There can be one VP admission controller per VP, or one for a set of VPs. The VPG controller operates on both views. Depending on the state of the VPs (in particular, traffic statistics and VP size) and the control objectives, it dynamically changes the amount of VPG bandwidth allocated to associated VPs. This controller enables customers to exploit variations in utilization among VPs that traverse the same VPG, allowing bandwidth between VPs of different source-destination pairs to be shared without interacting with the provider. In order to guarantee QOS, the sum of the VP capacities must be less than or equal to the capacity of the VPG link. The controller runs on a time-scale of seconds to minutes. The VPN controller operates on the right view. It is the only controller which interacts with the provider, and it runs on the slowest time scale of all the controllers (minutes or above). The VPN controller dynamically negotiates the bandwidth of the VPG links with the provider, based on traffic statistics and control objectives (e.g., minimizing the VPN cost), while observing the customer's QOS requirements.

4.1 Controller Design

Figure 5 shows the functional design of a VP admission controller and a VPG controller according to our implementation. In this design, the VP admission controller includes two objects: a VC capacity allocator and a coordinator. The allocator receives requests from a VC connection manager in the customer domain. The coordinator changes the capacity of the VP upon request from the VPG controller. It changes the capacity of the VP only when the bandwidth requirements of the active calls in the VP do not exceed the new capacity. The VPG controller includes four objects. The trigger object periodically initiates the VP capacity allocator to run the VP allocation algorithm. The coordinator sends the new VP capacities to the coordinators of the associated VP admission controllers, using a synchronization protocol. Finally, an estimator object collects statistics from the VP admission controllers. This data is used by the capacity allocator.

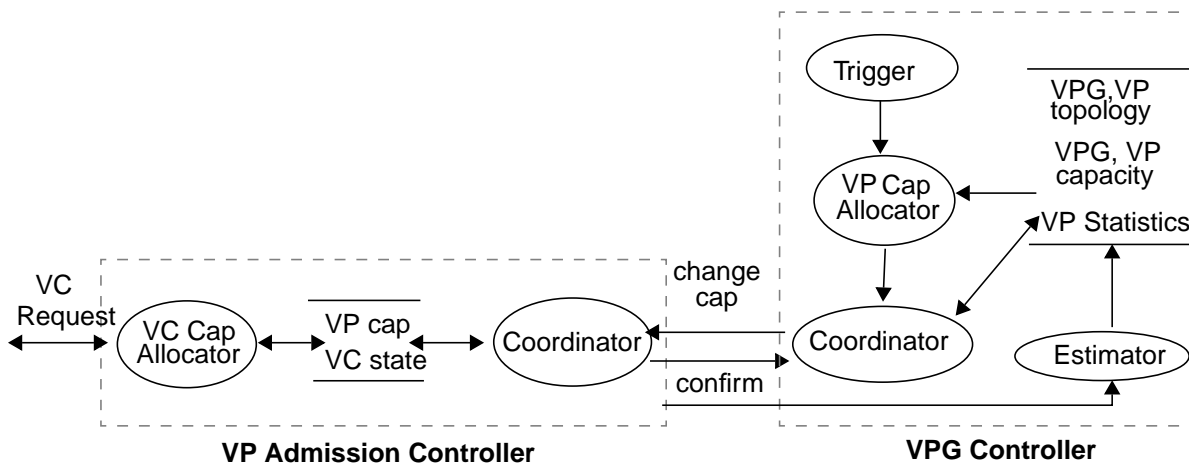


Figure 5 Functional model of a VP admission controller interacting with a VPG controller

Obviously, there exist many ways of realizing the above design, with respect to control algorithms, mechanisms for trigger realization, synchronization protocols, and centralized or distributed implementation of the controllers. For example, the control system may include one VP admission controller per VP or one centralized controller for the whole VPN. The same applies for VPG control. Also, VP 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 VP admission controllers. Similarly, the synchronization protocols between the VP admission controller and the VPG controller can be realized in different ways. One possibility is that the VP admission controller, upon receiving a request to change the VP size, checks whether the current utilization is above or below the new size. If the utilization is below, the VP size is changed and a confirmation is sent to the coordinator of the VPG controller. If it is not below, the VP size remains the same and a failure reply is sent instead. In another possible implementation, when the attempt for changing the VP size is not successful, the VP admission controller waits and blocks further calls from being admitted. Then, the utilization of the VP can only be decreased, as calls can leave but no new calls are admitted. When the utilization drops below the new size, the VP size is updated and the reply sent to the VPG controller. A customer's choice for a specific design of the control system is based upon its control objectives and requirements for the control system, which relate to system size, expected traffic and signalling load, efficiency of resource control and robustness of

the control system. In order to enable the realization of a large class of control objectives and control schemes, we have designed a generic controller as one of the building blocks of a customer control system. This generic controller enables many interaction patterns among controllers and is constructed in a modular way.

Figure 6 shows a functional model of the generic controller, which includes two sets of subcontrollers in a symmetrical design. One set of subcontrollers regulates the access to the resource, and the other set controls the size of the resource. The two sets of subcontrollers cooperate by accessing a shared data object, the resource graph. Each set of subcontrollers is made up of three functional components: trigger, allocator, and coordinator. The trigger decides when a computation should be done. The allocator performs the computation, which can be initiated by an external controller or by the trigger. The allocator that controls the access to the resource computes the amount of the resource that should be given to a particular request. The allocator that controls the size of the resource determines the resource capacity. A change of the resource capacity is coordinated by the coordinator object, which facilitates the interaction with other controllers. In particular, it implements the synchronization protocol needed to ensure that state changes among distributed controllers do not violate a set of resource constraints. The resource graph is modeled as two sets of weighted graphs, one representing the resource allocation and statistics, the other the resource capacity. Interfaces are provided to access and modify the relationship among these graphs.

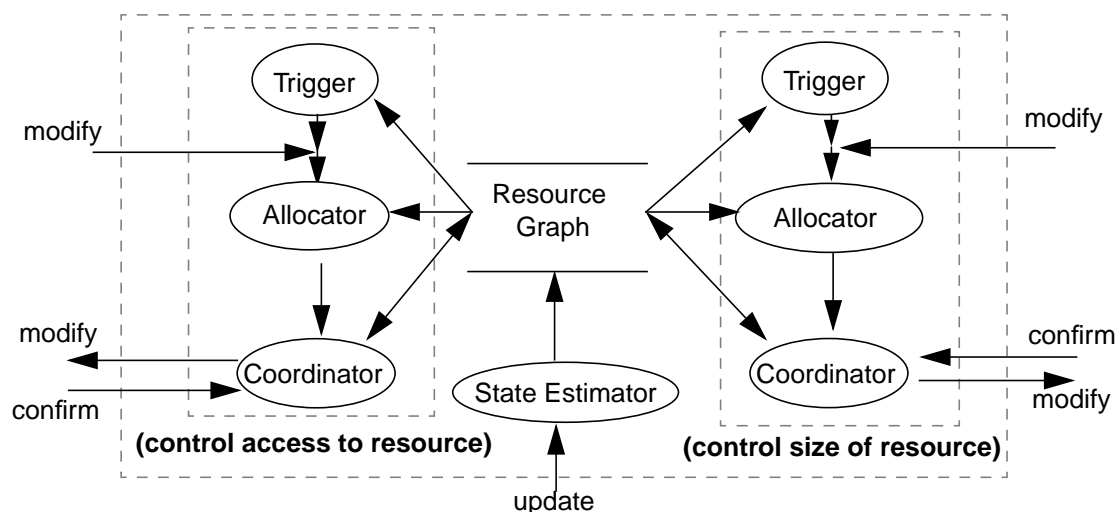


Figure 6 Functional model of a generic controller.

In our implementation, a generic controller is realized as a container class in C++, which includes as base classes the subcontrollers, trigger, allocator, coordinator, etc. Interfaces offered by these subcontrollers are implemented as virtual functions that are overloaded for a specific realization of the controllers.

The design of the generic controller shown in Figure 6 has brought us the following benefits. First, it was possible for us to design and implement all three classes of controllers --VP admission controller, VPG controllers, and VPN controller-- as a refinement of the generic controller class. For example, the VP admission controller in Figure 5 has two “non-trivial” controller objects --the VC resource allocator and the coordinator-- and five “trivial” controller objects. (Trivial controller objects can be thought of as objects which perform no action except that of forwarding data to another object. They are not shown in Figure 5). The VPG controller

contains four non-trivial controller objects and three trivial objects. Second, based on the generic controller design, we were able to realize different control schemes that attempt to achieve different control objectives for the customer control system. Realizing different control schemes is often possible by exchanging a set of subcontrollers in the system. For example, we implemented two classes of VC capacity allocators, realizing different VP admission schemes. One scheme aims at achieving call blocking objectives related to performance classes. The other scheme realizes call pre-emption in case of congestion, taking into account the priority of a call.

4.2 Enabling Management Objectives

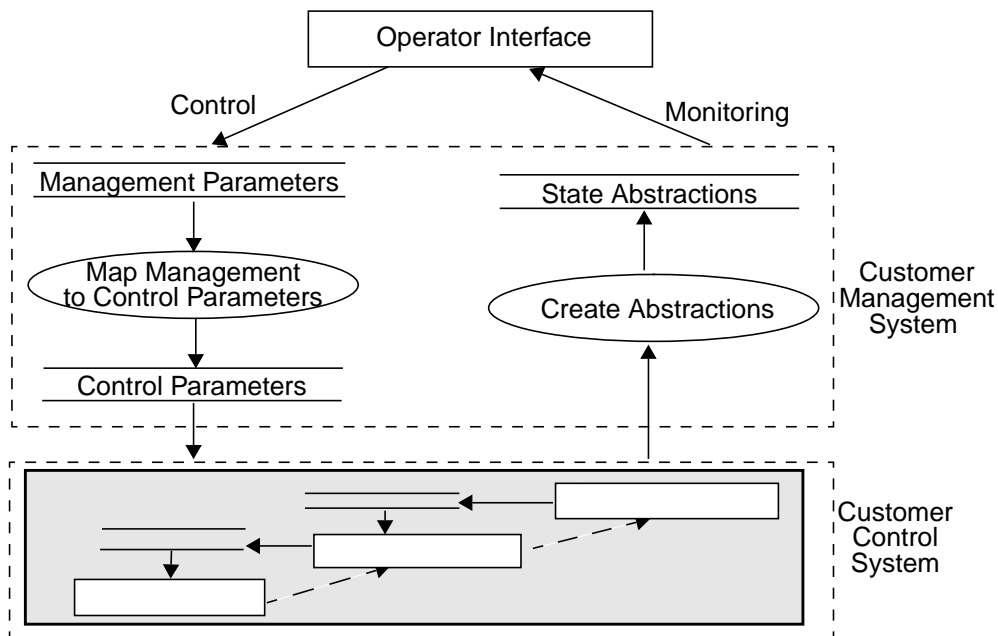


Figure 7 Framework for customer management.

The customer operates a management system to control and monitor the traffic on the enterprise network. A part of this system manages the traffic over the VPN, which is the focus of this paper. Examples of management capabilities that are related to the VPN service include controlling the bandwidth cost of the VPN service, VP bandwidth management, and QOS management (see Section 3). In the following, we describe how the management objectives outlined in Section 3 can be realized. Figure 7 shows our framework for implementing management capabilities. In this framework, management parameters, which directly relate to management objectives, are mapped onto control parameters, which influence the behavior of the controllers, and are subsequently distributed to the controllers in the customer control system [PAC95]. In our implementation, management parameters are made available to the operator of the enterprise network through the management console (Figure 8). A management parameter can be mapped onto control parameters for one or more classes of controllers. For example, cost management operations affect only the VPN controller. Allocating a specific capacity to a VP through a VP management operation affects both the VPG and the VPN controllers. QOS management operations, such as setting call blocking objectives, generally affect all classes of controllers. In response to a change in blocking objectives, the VP admission controller adjusts its admission policy, the VPG controller changes the VP allocation strategy, and the VPN controller negotiates the VPG sizes according to the new bandwidth requirements.

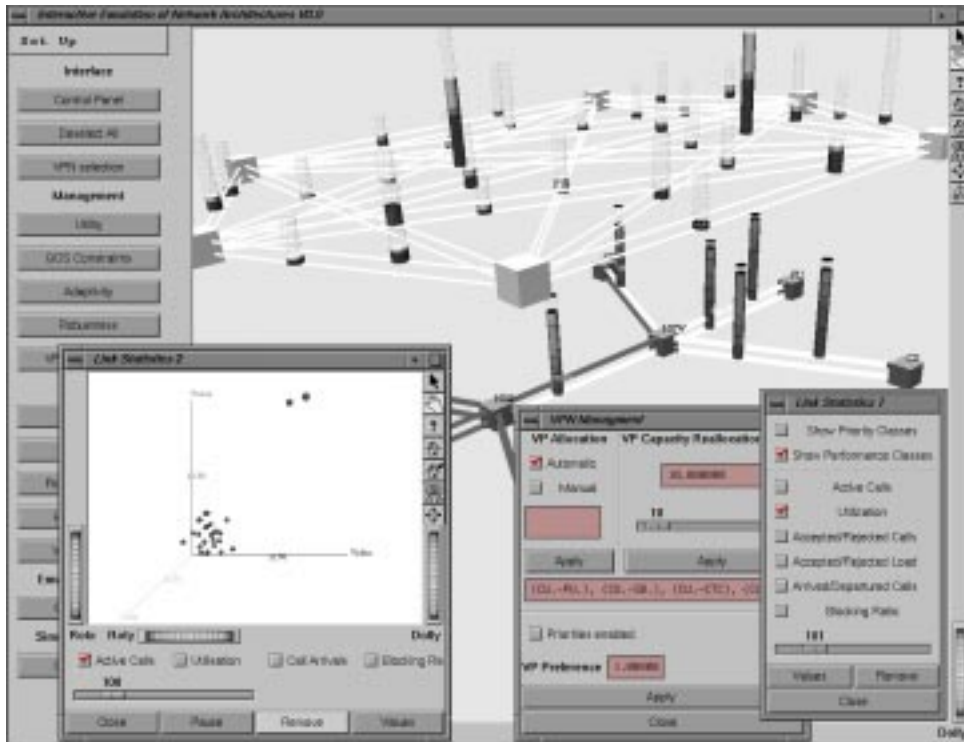


Figure 8 Management console for a VPG-based VPN service. The upper layer represents the VP network, the lower layer the VPG network. The vertical bars on the VP network indicate the utilization, the vertical bars on the VPG network the allocation of VPG bandwidth to VPs.

Figure 8 shows the screen of the management console that we have implemented for customer management of a VPG-based VPN service. Both layers of the VPN are visible. The upper layer represents the VP network, the lower layer the VPG network. The vertical bars on the VP network show the current utilization of the VPs. The three segments of a particular bar correspond to the three traffic classes supported in our particular system. The outline of the cylinders indicate the currently allocated VP capacities. The vertical bars on the VPG network give the allocation of the VPG bandwidth to the VPs. A “cloud view” on the lower left corner shows the number of active calls in the VPs. Each axis corresponds to a traffic class. In this specific snapshot, one can see that two of the VPs experience a much higher load than the others. The interface in Figure 8 allows an operator to perform management operations and observe the reaction of these operations on the global state of the system.

5. PROTOTYPING AND EVALUATING THE ARCHITECTURE

Two main tasks are involved in the development of a network architecture: the development of a software system and the design and analysis of algorithms. The first task focuses on software engineering aspects to satisfy the system requirements. The second concentrates on developing control functions that meet performance objectives. A thorough evaluation of the performance characteristics of a network control system has to take into account both of these aspects. Our approach to evaluating a target architecture is to build a software prototype, designed according to this architecture, which runs the intended algorithms [CHA96b].

The emulation platform consists of four building blocks: parallel simulation kernel, emulation support, real-time visualization and interactive control, and emulated system (Figure 9).




Figure 9 Building blocks of the interactive emulation platform.

The module for real-time visualization and interactive control contains an interface which provides 3-D visual abstractions of the system state. The emulation support module coordinates the exchange of control and monitoring messages between the graphical interface and the emulated system. It reads the states of the emulated system, and performs filtering and abstraction operations before making the information available for visualization. Control information from the user is mapped onto a set of control parameters that are interpreted by the emulated system.

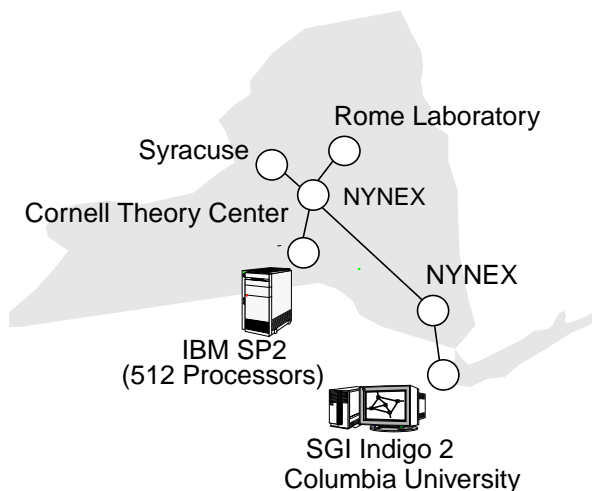


Figure 10 (a) Hardware configuration of the interactive emulation platform.

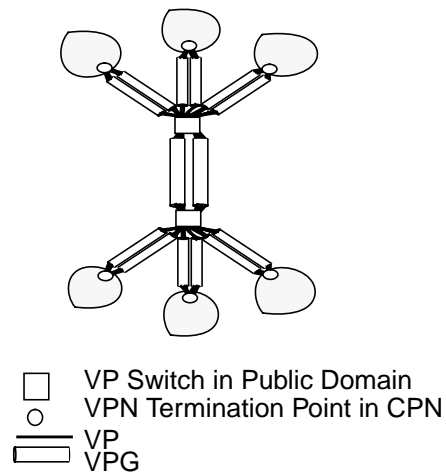


Figure 10 (b) Network topology used in the evaluation.

In our implementation, both the emulated system and the simulation kernel (coded in C++ and MPI) run on a SP2 supercomputer located at the Cornell Theory Center (CTC) in Ithaca, New York. The real-time visualization and interactive control module resides on an SGI Indigo2 workstation at Columbia University (Figure 10(a)). It is written using Open Inventor, a 3D graphics tool kit based on Open GL. The emulation support module is distributed on the two machines. These machines communicate through NYNET, an ATM network that connects several research laboratories in New York State.

5.1 Evaluating the Customer Control and Management System

In our evaluation, we implemented several versions of the VPN control system on the emulation platform. We have also built implementations of the management capabilities discussed in Section 3, namely, VP management, QOS management and priority management. Management operations are performed through the graphical interface (Figure 8). In the following, we summarize one of the results of the evaluation, in which we studied the effectiveness of resource control under constant traffic load. A more complete description of the experimental

results can be found in [CHA96a]. The performance of the customer control system was evaluated in a scenario based on the topology of the NYNET testbed. 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 (Figure 10(b)).

There is one VP admission controller per VP, executing a complete sharing policy. A centralized VPG controller periodically recomputes the capacities that are allocated to the VPs, by estimating the expected utilization of the VPs for the next control cycle. The VPG controller distributes the bandwidth to the VP admission controllers using a two phase protocol. This protocol ensures that the sum of the capacities of the VPs within a VPG does not exceed the capacity of the VPG. No VPN control is performed, i.e., the VPG link capacities remain constant during the course of the experiments.




Figure 11 Performance Evaluation of the customer control system.

We model the network traffic load, the processing time of the controllers and the time delay to send a message from one controller to another. 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. We vary 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 experience the same offered load.

We define the normalized VPG capacity as the ratio of the VPG capacity needed to attain a specific call blocking probability (B_{spec}) with VP capacity control over that without control (fixed VP capacities). The plot on the left side of Figure 11 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. The figure also shows the necessary VPN bandwidth for different traffic loads; the figure on the right side of Figure 11 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 distance

between the curves for the optimum control and our scheme in the left figure suggest that there is room for improving the algorithms and protocols we are running, specifically in the case where the offered traffic is low.

6. REFERENCES

- [ATS93] T. Aoyama, I. Tokizawa, and K. Sato, "ATM VP-Based Broadband Networks for Multimedia Services," *IEEE Communications Magazine*, April 1993, pp. 30-39.
- [CHA96a] M.C. Chan, H. Hadama and R. Stadler, "An Architecture for Broadband Virtual Networks under Customer Control," *IEEE NOMS*, April 1996.
- [CHA96b] M.C. Chan, G. Pacifici and R. Stadler, "Prototyping Network Architectures on a Supercomputer," *HPDC-5*, August 1996.
- [FOT95] S. Fotedar, M. Gerla, P. Crocetti, and L. Fratta, "ATM Virtual Private Networks," *Communications of the ACM*, vol. 38, no. 2, Feb. 1995.
- [ELW93] A.I. Elwalid, D. Mitra, "Effective Bandwidth of General Markovian Traffic Sources and Admission Control of High Speed Networks," *IEEE/ACM Transactions on Networking*, Vol. 1, No. 3, pp. 329-343.
- [HAD89] H. Hadama, 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).
- [HAD94] H. Hadama, T. Izaki, and I. Tokizawa, "Cost Comparison of STM and ATM Transport Networks," *NETWORKS'94*.
- [HYM91] J. M. Hyman, A. A. Lazar, G. Pacifici, "Real-Time Scheduling with Quality of Service Constraints," *IEEE Journal on Selected Areas in Communications*, September 1991.
- [HYM94] Hyman, J.M., Lazar, A.A. and Pacifici, G., "VC, VP and VN Resource Assignment Strategies for Broadband Networks", *Proceedings of the 4th International Workshop on Network and Operating System Support for Digital Audio and Video*, Vol. 846, Springer-Verlag, 1994.
- [LEW95] D. Lewis, S O'Connell, W. Donnelly, L. Bjerring, "Experiences in Multi-domain Management System Development," in *IFIP/IEEE ISINM*, Santa Barbara, 1995, pp. 494-505.
- [OHT92] S. Ohta, K. Sato, "Dynamic Bandwidth Control of the Virtual Path in an Asynchronous Transfer Mode Network," *IEEE Trans. Comm. Technol.*, 40, 7, pp. 1239-1247.
- [PAC95] G. Pacifici and R. Stadler, "Integrating Resource Control and Performance Management in Multimedia Networks," in *Proceeding of the IEEE ICC*, 1995.
- [SAY95] T. Saydam and J.P. Gaspoz, "Object-Oriented Design of a VPN Bandwidth Management System," in *IFIP/IEEE ISINM*, Santa Barbara, 1995.
- [SCH93] J.M. Schneider, T. Preuss, and P.S. Nielsen, "Management of Virtual Private Networks for Integrated Broadband Communication," in *Proc. of ACM SIGCOMM '93*, pp. 224-237.
- [TSC95] M. Tschichholz, J. Hall, S. Abeck, R. Wies, "Information Aspects and Future Directions in an Integrated Telecommunications and Enterprise Management Environment," *Journal of Network and Systems Management*, Vol. 3, No. 1, 1995, pp.111-138.
- [YAM91] T. Yamamura, T. Yasushi, N. Fujii, "A Study on an End Customer Controlled Circuit Reconfiguration System for Leased Line Network," *ISINM*, 1991, pp. 383-394.
- [ZER92] T.G.Zerbiec, "Considering the Past and Anticipating the Future for Private Data Networks", *IEEE Communication*, March 1992, pp.36-46.

Acknowledgments

This research was supported by the Department of the Air Force, Rome Laboratory, under contract F30602-94-C-0150. It was conducted using the resources of the Cornell Theory Center.