

DESIGN OF THE BANK OF AMERICA CALIFORNIA DATA NETWORK

Muayyad Al-Chalabi and William J. Liss

Muayyad Al-Chalabi is a member of technical staff in the Network Design and Optimization Department at AT&T Bell Laboratories in Holmdel, New Jersey. **William J. Liss** is head of the Network Management and Host Interfaces Department at AT&T Bell Laboratories in Liberty Corner, New Jersey. Mr. Al-Chalabi works on customer studies and design and optimization of integrated voice and data networks. He joined the company in 1984. He has a B.Sc. from the University of Manchester Institute of Science and Technology and an M.S. and Professional Degree from Columbia University, all in electrical engineering; and an Sc.M. in applied mathematics from Brown University. Mr. Liss is responsible for developing software for the StarKeeper® network-management-system product and for high-speed multi- (continued on page 106)

This paper describes Bank of America's California data network, a corporate utility network that uses AT&T's Datakit® virtual-circuit-switch technology to consolidate multiple networks into a single network. The network is managed by AT&T's StarKeeper® network-management system and Dataphone® II level-IV system controller. The network architecture consists of an access network that connects branches to hubs, and a backbone network that interconnects hubs and data centers. We address three fundamental areas: optimization of the network topology and node configuration, performance analysis, and network management. Network optimization deals with determining the optimal number of nodes, access facilities engineering, and backbone configuration and routing. Performance analysis shows the end-to-end delays for the applications. In the network management area, we discuss administration, disaster recovery, network monitoring, and data collection.

Introduction

The Bank of America (BoFA) has about 1000 branch offices in California. Each branch supports multiple applications such as teller and banking platforms, automated-teller machines (ATMs), private banking, and real-estate loan processing. (Panel 1 defines terms and acronyms used in this paper.) These applications were introduced into the branches over several years, depending on the needs of the time. Each application has its own multidrop network and its own protocol and devices. The presence of these discrete networks poses problems in managing resources, adding new applications, and adapting to traffic changes.

As a technology leader, Bank of America recognized that operating multiple, independent networks did not conform to the new

banking needs. Therefore, it decided to integrate the transport of these networks at the point of origin—i.e., the branch—over a single utility network.

The Bank of America is currently deploying its new California data network (CDN). To provide physical transport and management for this network, it has selected a solution based on AT&T's Datakit® virtual-circuit switch (VCS) and StarKeeper® network-management system (NMS). The open architecture and flexibility of Datakit VCS will allow the physical-transport part of CDN to be a true utility network that will serve a wide variety of current and future data applications.

CDN is designed to meet the following objectives:

- Reduce costs through consolidation of existing physical networks.
- Increase network availability through robust network design and end-to-end network management.
- Support a multivendor environment.
- Deploy new services and applications more rapidly.
- Protect the current investment in deployed application hardware.
- Provide the capability to implement a contingency operations plan (COP) for data-center backup to protect against disasters (e.g., fire, chemical leaks, and floods).

This paper presents BofA's current mode of operations, technology tradeoffs for design approaches, methods used to design the new network, and implementation of this large and complex data network. Because the network described here represents the solution as of a given time, some details cited may change as BofA's needs and plans evolve.

Analysis of the Current Network

This section describes the data-transport environment at BofA before implementation of the new utility network based on Datakit VCS. Most of the situations discussed are common to other large, corporate data networks that have multiple applications with a variety of protocols and are not peculiar to BofA. They arise largely

because increased intercorporate competition (with its effect on data applications) is juxtapositioned with rapidly changing data-communications technology. BofA has been proactive over the last few years in analyzing the shortcomings of this approach to networking and has recognized the need for a new network to support its business plans in the future.

Branch-Office Network Topology. Each BofA branch office in California normally supports several banking applications, and each application connects to data centers in Los Angeles or San Francisco through a dedicated, point-to-point or multidrop, analog private-line network. The Concord, California, facility is both BofA's main network-management center (NMC) and its backup data center. This deployment produces a large mesh of disjointed private-line networks, with each branch office served by "drops" from several of the networks.

Figure 1 shows how this arrangement serves three branch-office applications (e.g., ATM, banking platform or teller, and investment). When branch office needs are added to the larger service centers and internal BofA applications (shown at the bottom of Figure 1), BofA has over 60 logical and/or physical data networks. Each network consists of a variety of data-transport and network-management technologies.

Network Operations. Of particular interest are two aspects of the network operations: cost and quality of the communications service.

Network cost consists of two main components: facilities and operations staff. Multidrop networks do not use analog private lines efficiently. Studies done by BofA's telecommunications organization showed that most BofA applications use only 15 to 20 percent of the network capacity. In addition, the cost of analog private lines continues to increase at a faster rate than that of high-capacity digital facilities, such as AT&T's digital data service (DDS) or Accunet® T1.5 service. Because these private-line networks involve a wide variety of technologies and operations procedures, operations costs are also high as

Panel 1. Acronyms and Terms

3174	a 3270-compatible terminal/system	FEP	front-end processor
3270	IBM's synchronous terminal/cluster controller system	Hub	a module used to identify SACs and BACs
ATM	automated-teller machine	ITN	integrated-transport network
BAC	backbone-access center	LATA	local-access transport area
BAU	Branch Automation Utility	LSN	location serial number
BofA	Bank of America	NMC	network-management center
BSC	binary synchronous communications	NMS	network-management system
CC	cluster controller	PAD	packet assembler/disassembler
CCITT	International Telegraph and Telephone Consultative Committee	PDD	predefined destination
CDN	Bank of America's California data network	PI	programmer's interface
COOLS	Community Office On-Line System	RS-232	an EIA physical/electrical connector standard
COP	contingency operations planning	RTA	response-time analysis
CSU	channel service unit	SAC	service area concentrator
DAF	distributed application function	SAM	synchronous/asynchronous multiplexer
DCE	data communications equipment	SDLC	synchronous data link control
DDS	digital data service	TDM	time-division multiplexing
DSU	data service unit	TSM	transparent synchronous module
E-INOS	Enhanced Interactive Network Optimization System	TY	an asynchronous-connection module
EIA	Electronic Industries Association	UNMA	unified network-management architecture
		VCS	virtual-circuit switch
		X.25	a CCITT packet-switching protocol

reflected in the large numbers of operational staff and high training costs.

Several things influence a user's perception of *network service quality*: the quality of data transmission, number of network outages, and time required to repair these outages. The analog private-line facilities used in the existing BofA network are more prone to transmission errors than are DDS facilities. Facility or equipment outages will generally affect the data network user (e.g., ATM customer or BofA branch-office teller) because no alternate-path routing is available in most multidrop private-line networks. In addition, the time required to diagnose and repair problems in a network that involves a variety of different vendor equipment tends to be longer because:

- Multivendor equipment interface problems are often difficult to diagnose.
- The maintenance and operations staff must be familiar with the wide variety of network equipment and, therefore, is not as proficient with each component.

Contingency Operations. As with many corporations, BofA would like to be able to redirect large segments of data traffic to backup data centers when a disaster occurs. This capability is rapidly being mandated, especially for the banking industry. Most likely, a "common-disaster" transfer cannot be done with point-to-point or multidrop private-line networks.

Applications Support. In the past, the banking environment was stable, and bank-related data applications

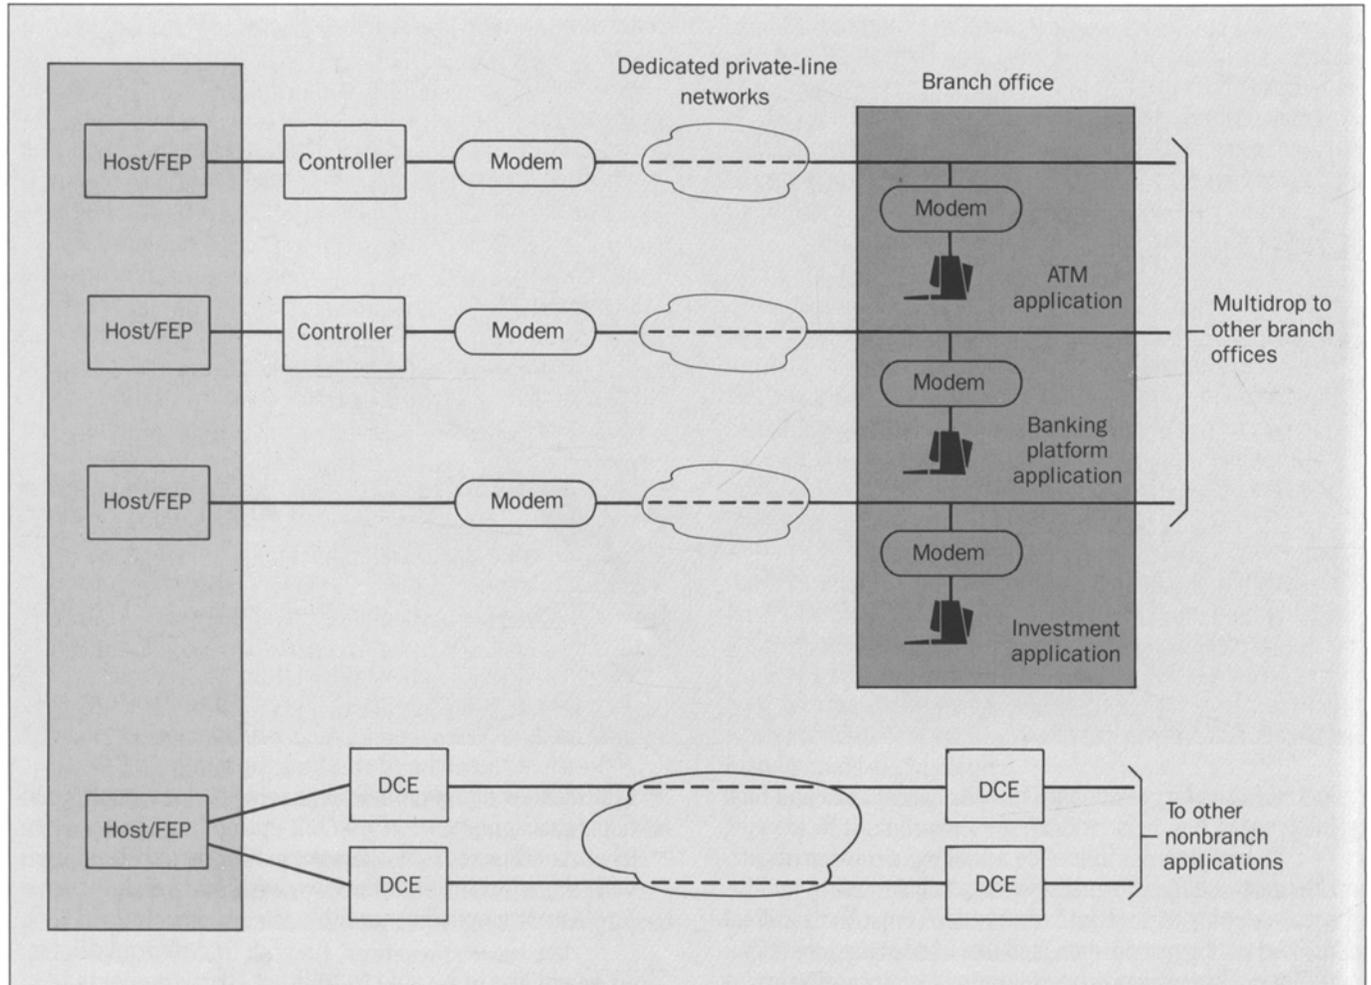


Figure 1. The Bank of America network in California before use of the Datakit virtual-circuit switch. ATM stands for automated-teller machine, DCE is data communications equipment, and FEP is a front-end processor.

remained relatively unchanged for years.

There is much more change in today's deregulated environment, and banks face competition from nonbanking concerns. In this dynamic environment, banks must be able to respond quickly with new applications. But such changes in a private-line network frequently require tearing down and rebuilding large pieces of the network. The

time and cost of doing this cannot be tolerated in today's business environment.

Technology Assessment for CDN

Bank of America extensively studied two networking alternatives for CDN: packet-switching technology and time-division multiplexing (TDM). Two implementations of packet switching—X.25 packet switching and Datakit virtual-circuit switching—were studied. (X.25 is a packet-switching technology that uses protocols defined by CCITT, the International Telephone and Telegraph Consultative Committee.) Tradeoffs were evaluated in the following areas:

- Ease of network operations, administration, and maintenance
- Ability to support multiple protocols
- Flexibility to accommodate future growth with minimal change to installed hardware or circuits
- Improved network-response time.

Time-Division Multiplexing. In a TDM implementation, a physical channel is established between the end devices (cluster controllers, or CCs) at the branch and front-end processors (FEPs) at the data center. A multi-port, split-stream modem is placed at the branch. The speed of the modem ranges from 9.6 to 19.2 kb/s (kilobits per second), and it has four to six input channels, each operating at 2.4 kb/s. In the backbone network, the low-speed channels are multiplexed into 1.544-Mb/s (megabits per second) channels through T1.5 multiplexers.

Although TDM technology consolidates current applications over a single circuit that emanates from each branch, it only replicates today's inflexible environment. Further, TDM has two disadvantages:

- The ports and bandwidth requirements at most branches exceed the capacity and speed of the modem technology.
- Low use (less than 20 percent) of the low-speed circuits results in low use of the T1.5 backbone network.

As a result, the high-speed circuits are used inefficiently.

Packet Switching. Both X.25 and Datakit VCS effi-

ciently use private-line facilities in a distributed, statistically derived way.

BofA's large embedded base of diverse equipment from several vendors prohibits changing any of the current network's equipment. In an X.25 network, a complete protocol specification—implemented in a packet assembler/disassembler (PAD)—is needed for each data-network user's device because protocols terminate at the edges of the network. The myriad protocols installed in BofA networks prohibit this option because of development and operations costs.

A typical X.25 network implementation consists of:

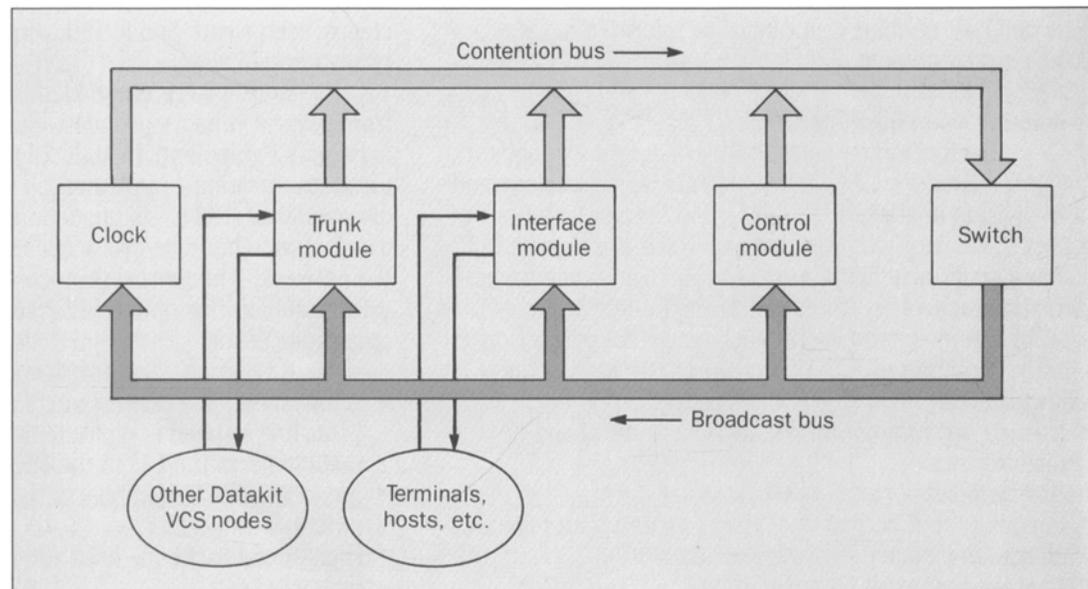
- A low-speed line that connects the CC to the PAD. (This line's speed is equivalent to the speed of the line that connects the CC to the FEP.)
- An access line that connects the PAD to the X.25 switch.
- High-speed backbone links (56 kb/s) that interconnect the switches.
- Either low-speed lines that connect the host PADs to the FEPs, or a direct connection between the switches and the FEPs.

Because of the protocol termination, store-and-forward techniques are used at both the access and egress of the access network and in the backbone networks. End-to-end delay is a function of frame size, packet overhead, access-line speed, backbone-line speed, and switch-processing delay. When the access-line speed is low, this component dominates the other delay components and greatly increases end-to-end delay.

Datakit VCS was selected for CDN because it offers the following features:

- Transparent support for multiple protocols through programmable interface modules. The software can be downloaded into each port-card module, which makes it extremely attractive to BofA for accommodating diverse protocols with minimal hardware change.
- High-speed trunking (1.544 Mb/s).
- "Pipelined frames" on both the access and backbone networks. Synchronous data-link control (SDLC),¹

Figure 2. Internal architecture of the Datakit virtual-circuit switch.



92

binary synchronous communications (BSC),² and other synchronous frames are broken down into smaller Datakit VCS frames and transported on the network. Pipelining reduces end-to-end delay.

- Dynamic configuration of virtual multipoint lines.
- End-to-end network management through StarKeeper NMS and the StarKeeper network troubleshooter.³ Because of the enhanced on-line diagnostics and trouble isolation capability, network-operations personnel become more productive and overall network availability increases.

The next two sections briefly describe the AT&T products—Datakit VCS and StarKeeper NMS.

Datakit VCS Overview

The Datakit virtual-circuit switch is a high-performance, state-of-the-art packet switch that has been designed to interconnect data-processing equipment. Based on a “kit” of compatible modules, the Datakit VCS can be configured to serve a wide range of network appli-

cations. High-speed, hardware-based packet switching—combined with the efficiency of virtual-circuit switching—gives Datakit VCS its high throughput, low delay characteristic. Figure 2 is an overview of the internal architecture of Datakit VCS.

Data packets are switched within Datakit VCS through the combined actions of its contention bus, broadcast bus, switch module, and clock module. A terminal-connected interface module, which wants to transmit data to an interface module connected to another host (or a trunk module connected to a distant Datakit VCS), will offer the data packets to the contention bus. The contention bus (and broadcast bus) operates at a rate of 8 Mb/s. If we account for packet overhead, this allows an effective throughput of 44,000 Datakit VCS packets per second, or about 5.6 Mb/s.

Packet switching is fast because it is implemented completely in hardware. The *contention bus* operates in a synchronous way and completely avoids collisions. At the beginning of a packet cycle, all interface modules that have

something to transmit contend for this bus. One, and only one, module wins the contention bus and transmits a packet. A module continues to contend for this bus until it wins and transmits a packet.

The *switch module* is responsible for switching a transmitted packet to the proper destination. Each packet on the contention bus contains the address of the packet's source. The switch module translates this source address into a destination address and places the packet on the *broadcast bus*. The *receiving module* will monitor the broadcast bus and remove packets addressed to it.

The high-speed Datakit packet switch is used to implement a virtual-circuit service between terminals and host computers. This technology has two characteristics that are important to wide-area networks: efficient use of facility bandwidth, and low nodal-switch delay.

In the Datakit VCS, network bandwidth is not preallocated to a circuit but is used only on demand, which permits highly efficient use of bandwidth and reduces trunk costs. Cross-switch delay is also small—typically 2 to 9 ms (milliseconds)—because of a combination of the previously described high-speed backplane, small packet sizes, and hardware-based packet switching. This is in contrast to other, traditional packet-switching techniques that may require software processing of the protocol up through level 3 (i.e., the network layer) to do packet routing.

Features. Three features of the Datakit VCS interface modules are particularly important in BofA's multi-drop, multiprotocol, synchronous environment: support for multiple protocols, emulation of multipoint lines, and pipelines for data transmission.

Multiple protocol support. The Datakit VCS, by design, simultaneously and efficiently handles different protocol interfaces (including asynchronous, bisynchronous, SDLC, and X.25). The port-interface function of its transparent synchronous module (TSM) was designed to be transparent to the data being transported. In addition, each port on the TSM can be configured for a different protocol. This allows the user to choose the best vendor's

equipment for each application (e.g., ATM, teller, etc.) and still be able to support the resulting mixture on a single Datakit VCS network.

Multipoint line emulation. In a Datakit VCS network, multiple appearances of CC ports on the branch-office side can be multiplexed to appear as a single multidrop line to an FEP connected to the host side. In general, Datakit VCS will broadcast outbound data to all devices defined as members of this logical multipoint circuit.

For standard 3270 BSC and SDLC protocols,^{1,2,4} Datakit VCS intelligently "fans out" traffic to cluster controllers. (The 3270 is IBM's synchronous terminal/cluster controller system.) For example, consider what happens as the Datakit VCS node receives a frame of SDLC data from the FEP. The node maps the frame's SNA™ network address to a specific virtual circuit and routes the data only to its destination rather than to all devices defined on the multipoint. (SNA is a trademark of IBM Corporation.) This is much more efficient than unintelligent broadcasting, and conserves backbone-network bandwidth.

Data pipelining. On an X.25 packet network, nodes must wait until they receive an entire data frame before forwarding it to the next network node. This is done for routing and error-correction purposes. Unfortunately, this successive store-and-forward technique increases end-to-end network delay.

The Datakit VCS pipelining feature breaks up the large data frames into smaller data packets. As each packet is filled, it is routed through the network to the receiving device. In this way, data at the beginning of the frame completely traverses the network before data at the end of the frame is received. This largely eliminates the store-and-forward delay of other packet networks and is particularly useful in SDLC data transmission, which is prevalent at BofA.

Switch Capacity. Datakit VCS hardware is available in a variety of sizes:

- Model 2000 can be configured as two 4-shelf cabinets for large sites that have up to about 1000 ports (e.g., terminals), depending on the specific mixture of port-

card types.

- Model 500 is a two-shelf cabinet for medium sites, with up to about 200 ports, depending on the specific mix of port cards.
- The Datakit VCS concentrator (a synchronous/asynchronous multiplexer, or SAM) is a family of single-shelf units that serve from 4 to 64 ports.

Network-Management Overview

The StarKeeper NMS, a multifunction operations system based on UNIX® System V, supports Datakit VCS networks. It runs on a variety of AT&T 3B computers and is one element of AT&T's unified network-management architecture (UNMA) offering. StarKeeper NMS provides switch surveillance and maintenance capabilities, and centralized collection, analysis, and reporting of traffic and billing data generated by the Datakit VCS.

While the list of StarKeeper NMS features is lengthy, a few have proved to be of particular importance to BofA: color displays of network status, a programmable user interface, and an alarm system.

Network Graphics Displays. The StarKeeper NMS presents color-graphics displays of the network. Different views of the network are possible, from a network overview to the view of a node and the modules it contains.

When network faults are displayed, the color of the failed network entity changes dynamically according to the severity of the failure. This allows BofA's network-operations-management personnel to monitor overall network status at a glance.

Programmer's Interface. A user-programmable interface allows a network administrator to write and incorporate UNIX system shell scripts or C programs into the StarKeeper NMS. (A *shell script* is a program written in the UNIX system's shell language.) Then, the NMS programmer's interface (PI) subsystem receives messages from the network and distributes them to these programs or shell scripts. Either the Datakit VCS nodes or StarKeeper NMS may generate the messages that may be distributed to the programs or scripts.

It is expected that BofA will use this feature in the future as its need increases for customized features for its network.

Enhanced Alarm Conditioning. The network administrator may set alarm thresholds and redefine alarm severity. Thresholds may be set to sound an alarm for every occurrence of an event after the n th occurrence, or to sound it after every n events. Separate thresholds can be configured for each type of alarm.

The administrator may also reassign an alarm's severity. Initially, each alarm is placed into one of three categories—minor, major, and critical—but any alarm may be moved to a different severity level.

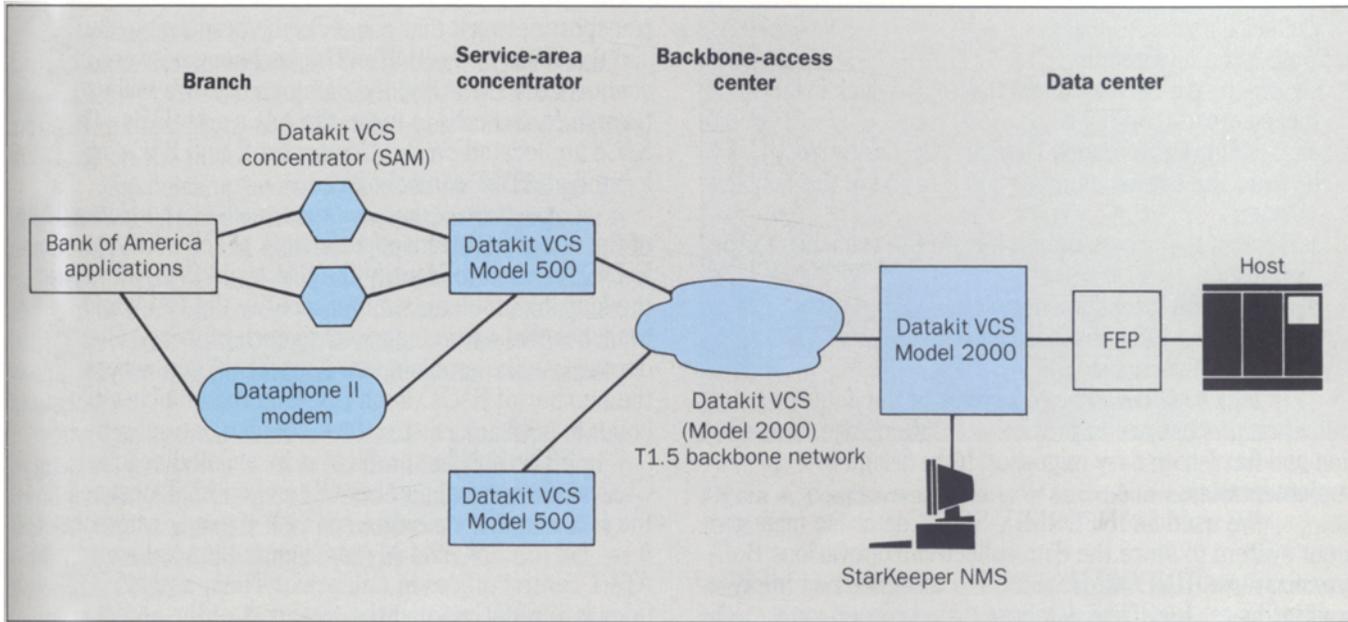
StarKeeper NMS also provides alarm import, export, and filtering functions. A StarKeeper NMS can send alarms to another StarKeeper NMS, where a single consolidated-alarm summary will be made available to the NMC operator. A filtering capability permits the administrator to specify which of the alarms are to be exported. Also, controller alarms from a Dataphone® II level-IV system (i.e., from modems) can be imported.

An alarm text redefinition capability exists that permits alarm messages to be tailored for a specific location. For instance, the alarm message could reference a manual page or include a telephone number to call for help. This is important to BofA because network operations personnel will receive alarm messages that contain useful information, which will increase their responsiveness and effectiveness when real problems arise.

Network Design

For the analysis and design of BofA's California data network, we used a proprietary AT&T network-design system, the Enhanced Interactive Network Optimization System (E-INOS).⁵ AT&T Bell Laboratories developed this system and continually enhances it to accommodate new services and new technology.

The design problem can be posed as follows: Suppose we have a set of branches $B = \{b_i\}$, where $i = 1, 2, \dots, 1000$; a set of data centers $D = \{d_j\}$,



where $j = 1, 2, 3$ (i.e., two primary data centers and one for backup); and a set of traffic requirements $R = \{r_{i,j}\}$, where $i \in B$ and $j \in D$. Then, minimize the objective function:

$$\sum_{i=1}^{i=1000} \sum_{j=1}^{j=2} c(r_{i,j})$$

where $c(r_{i,j})$ is the cost of the facilities generated to satisfy the requirements of connecting branch i to its primary data center j .

The problem is subject to the following constraints:

- Network availability should be better than 99.95.
- End-to-end network delay (excluding host and end-point delay) for a typical transaction (i.e., 300 characters from the terminal to the host, and a 750-character response from the host) should not exceed 3 seconds, on average,

Figure 3. New network architecture. FEP is a front-end processor, NMS means network-management system, and SAM is a synchronous/asynchronous multiplexer.

and 5 seconds for the 95th percentile. Of course, these constraints apply to a specific network configuration.

Because so many variables must be examined to design a network of BofA's size, an exhaustive search for an optimal solution is not feasible. In our approach to solve BofA's data networking problem, we decompose the problem into a set of subproblems that are solved in sequence and perform local optimization on subsets of the network variables. The process generates solutions in stages and progressively reduces the number of variables to be examined in successive stages.

The design process involves an iterative application of the following steps (which are discussed in the subsections that follow):

1. Collect data.
2. Define the architecture.
3. Configure the hardware and topology, which involves these steps:
 - a. Configure the branch Datakit VCS hardware.
 - b. Find the optimal number and location of the hubbing nodes.
 - c. Design the access network from the branches to the hubs.
 - d. Design the backbone network.
 - e. Configure the hubs' hardware.
4. Evaluate performance.

Data Collection. The objectives of the data-collection process are to provide accurate design information and furnish an easy migration from design to implementation.

We used an INFORMIX-SQL™ database management system to store the data collected from various BofA organizations. (INFORMIX-SQL is a trademark of Informix Software, Inc.) The database was designed with multiple tables that are cross-referenced by a location serial number (LSN). For example, branch-related information, such as street address and telephone number, are stored in the branch table. Transaction rates and sizes of messages for each application are saved in the traffic-profile table. Each application's cluster controllers, terminals, protocols, and speed are stored in the device table. Datakit VCS equipment and modems, channel service units (CSUs) or data service units (DSUs), channel banks, etc. are kept in the network equipment table.

Architecture Definition. BofA's network has a four-tier architecture (Figure 3):

1. Branch level
2. Service-area concentrator (SAC) level
3. Backbone-access center (BAC) level
4. Data-center level.

Tiers 2 and 3 are interrelated; they prove-in higher-speed facilities through concentration of traffic from the lower layer.

We will use the term *access network* to mean the

transport network that connects the branches to the SACs and the SACs to the BACs. The backbone network interconnects the BACs and the data centers. We will use the term *hubbing nodes* to mean the SACs and BACs. The SACs are located on BofA's premises, and BACs are located at AT&T central offices.

Topology and Hardware Configuration. The Hub module of the integrated-transport-network (ITN) subsystem in E-INOS is used to identify the SACs and BACs. To solve the hubbing problem, we first identify the BACs and the branches that home (connect) to each of them. Then, we decompose the problem into K subproblems, where K is the number of BACs. Each BAC and the branches that home to it form a cluster, thus defining one subproblem.

The hubbing problem is formulated as follows: Choose the optimal number (K) and optimal location of the BACs from a candidate set $CO = \{co_k\}$, where $k = 1, \dots, N_{co}$ and N_{co} represents the total number of AT&T central offices in California. Then, assign each branch to a BAC so that the total cost of the access network is minimized. Similarly, identify the SACs in each BAC cluster by selecting the optimal number (N_{sac}^q) and location of SACs from a candidate list $SAC^q = \{sac_{i,j}^q\}$, where $q = 1, \dots, K$ and $v = 1, \dots, M^q$, where M^q is the number of branches in a BAC cluster q . Two constraints are imposed on the solution for each cluster:

1. Each branch is assigned to one SAC:

$$\sum_{j=1}^{N_{sac}^q} x_{i,j}^q = 1$$

for $i = 1, \dots, M^q$.

2. A SAC can handle a maximum of 15 branches:

$$\sum_{i=1}^{M^q} x_{i,j}^q \leq 15$$

for $j = 1, \dots, N_{sac}^q$, where $x_{i,j}^q$ is the link that con-

nects branch i to SAC j in cluster q . Constraint 2 was included as a first-order performance constraint. A more detailed performance constraint is included in the design, and a detailed performance analysis was done on the final configuration.

The hubbing algorithm is a one-opt exchange (it does one interchange at a time) that requires the generation of an $N \times M$ cost matrix, where N is the total number of branches and M is the number of hubs. Each element of the matrix represents the cost of homing the branch to a hub. Each branch is homed to the nearest hub without exceeding the limit on the maximum number of branches.

The optimal number of BACs is a function of the branch concentration in an area, the traffic distribution, the line costs, and the number and location of AT&T central offices. As a constraint, no more than 150 branches (indirectly hubbed through a SAC) or 15 SACs can be homed to a BAC.

We performed a parametric analysis on the number of BACs, where we varied the number of nodes between 2 and 20 as presented in Figure 4. The results show that, from a total-cost perspective (access network, backbone network, and switch costs), all configurations with 5 to 12 BACs seem nearly equal. From both a reliability and growth perspective, we selected a nine-node design as the most cost-effective solution. Because of the high concentration of branches in the San Francisco and Los Angeles metropolitan areas, we paid special attention when determining these nodes to provide geographic diversity to minimize the probability of common facility failure.

To reduce the complexity of the SAC-placement problem, we chose the branches designated as "large" and "extra large" as the SAC candidate set. We then selected an optimum number of SACs per BAC cluster from the candidate set, and did a parametric analysis for each BAC cluster to find the optimal number and location of SACs. We found that 150 SACs satisfied BofA's requirement.

Branch configuration (end point). The branches repre-

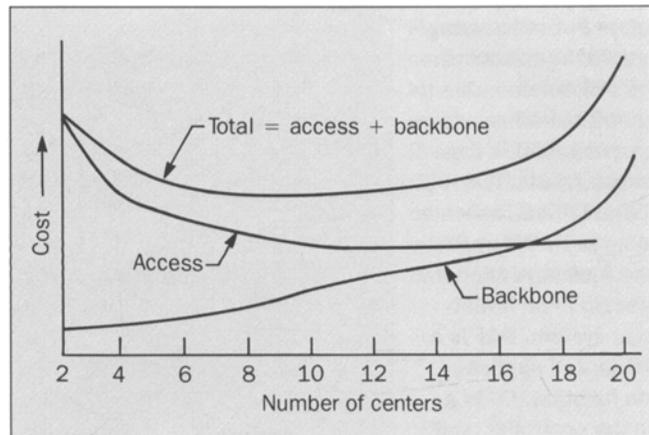
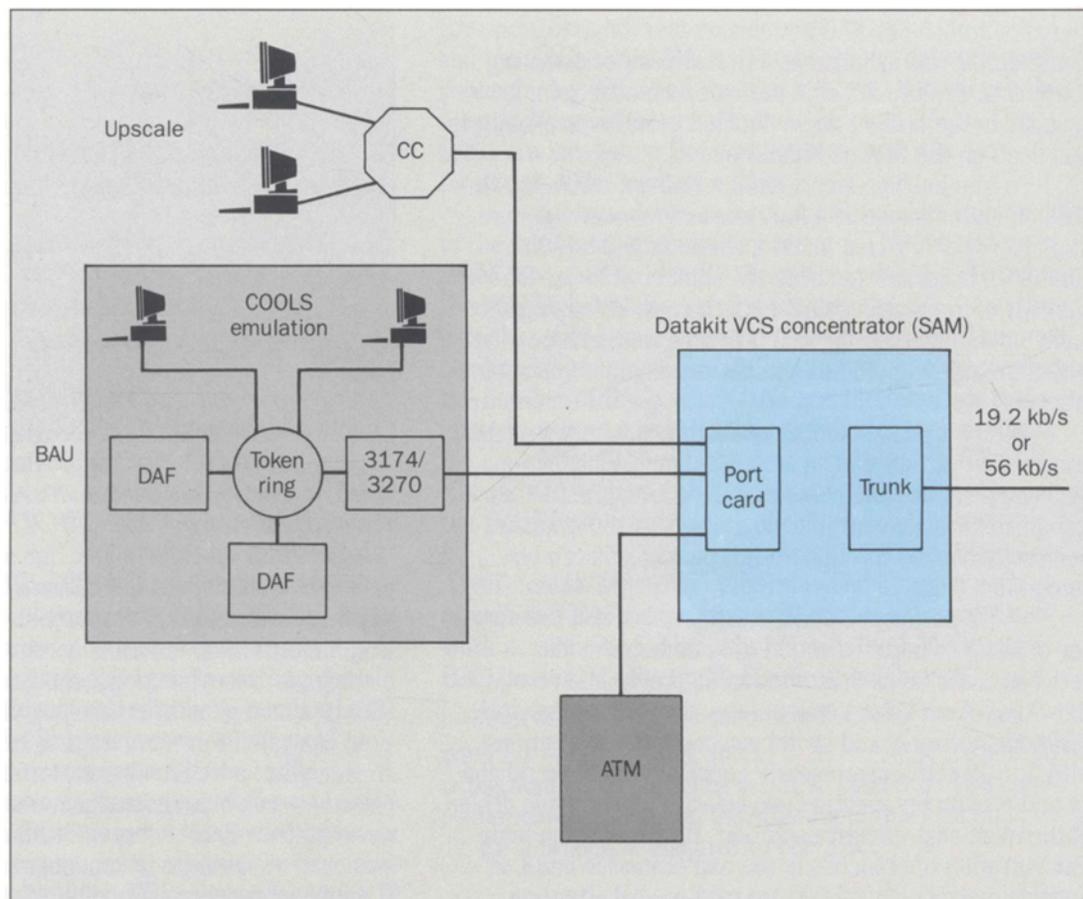


Figure 4. Cost versus number of backbone-access centers (BACs).

sent the end points of the network; that is, BofA's points of presence in an area. Application hardware—such as the Community Office On-Line System (COOLS) for banking platform or teller functions; the Branch Automation Utility (BAU), a new-generation banking platform or teller function; Upscale for investment functions; and ATMs—are connected to a VCS concentrator. The type of equipment placed at each branch is a function of the number of devices at the branch and the traffic generated. Small branches or satellites of branches will receive a Dataphone II multiport modem—Model 2096 (9.6 kb/s), Model 2144 (14.4 kb/s), or Model 2192 (19.2 kb/s)—depending on the ports and bandwidth required. In BofA's network, 400 of the branches are classified as modem branches. In 700 branches, a Datakit VCS concentrator (a SAM) will be installed. The number of ports used and the amount of traffic generated or received by the branch determines the trunk speed from the branch. Either 19.2 kb/s or 56 kb/s may be selected, based on a cost-versus-performance tradeoff.

Figure 5 shows a typical branch configuration using SAMs.

Figure 5. Branch configuration (end point). The ATM is an automated-teller machine, BAU is the Branch Automation Utility, COOLS is the Community Office On-Line System, and Upscale is an investment system. DAF is a distributed application function. CC is a cluster controller, and SAM is a synchronous/asynchronous multiplexer.



98

Service-area concentrator. We determined that BofA's network required 150 SACs, which typically are large branches or service centers. A SAC provides an intermediate hubbing point between the end points (branches) and the BACs, and aggregates traffic from branches to a BAC. In addition, SACs generate their own traffic, i.e., applications that reside at the branch.

The SACs allow BofA to tie new end points directly into the network with minimal carrier involvement. A SAC may consist of a one- or two-shelf Datakit VCS Model 500, or a Model 2000. Most SACs have a Datakit

VCS Model 500, while a few large sites have a Model 2000.

Figure 6 shows the basic SAC configuration. The nodes contain:

- TSMs—These modules interface to the end devices.
- T1.5 modules—These modules are connected to the remote branches that have a SAM.
- DDS trunk modules—This module connects a SAC to the nearest BAC.

Access network. The access network consists of two components: access links from branch to SAC, and access

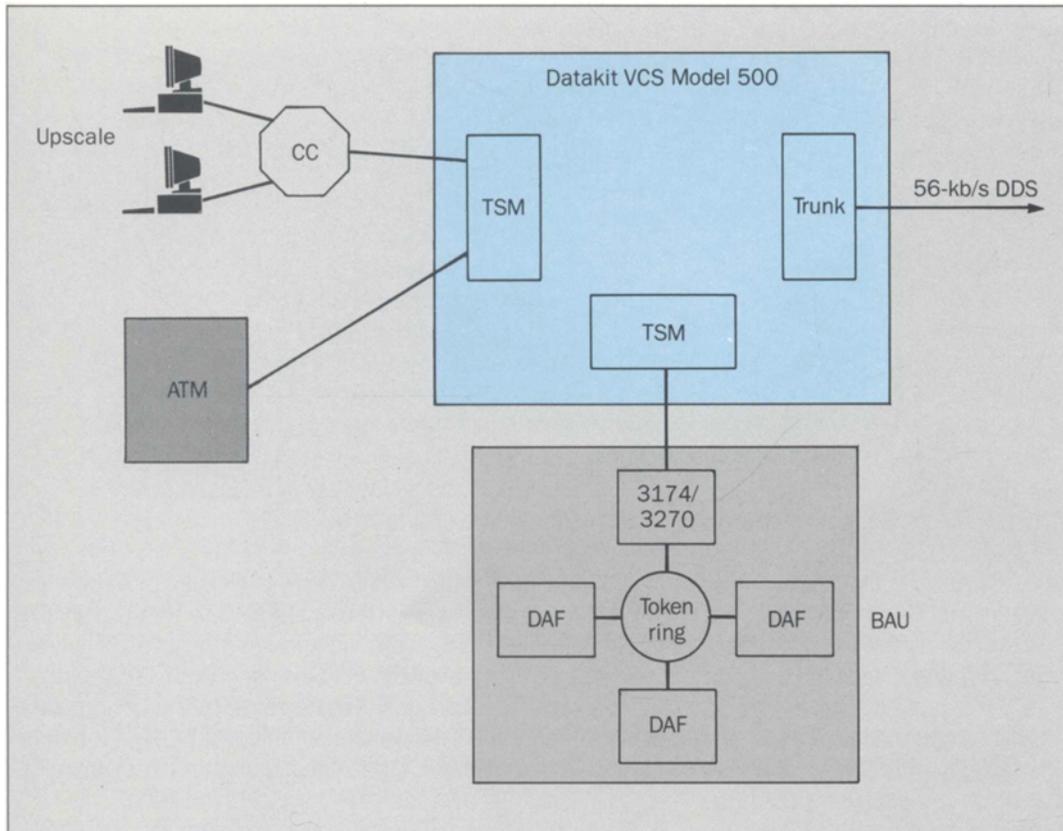


Figure 6. Service-area concentrator. ATM is an automated-teller machine, BAU is the Branch Automation Utility. CC is a cluster controller, and TSM is a transparent synchronous module.

99

links from SAC to BAC.

We considered three services for connecting the branches to SACs: analog, 9.6-kb/s DDS, and 56-kb/s DDS. For this portion of the network, 9.6-kb/s DDS (type-2 local channels) is about 41 percent more expensive than an analog network, while 56-kb/s DDS (type-2 local channel) is 95 percent more expensive than analog.

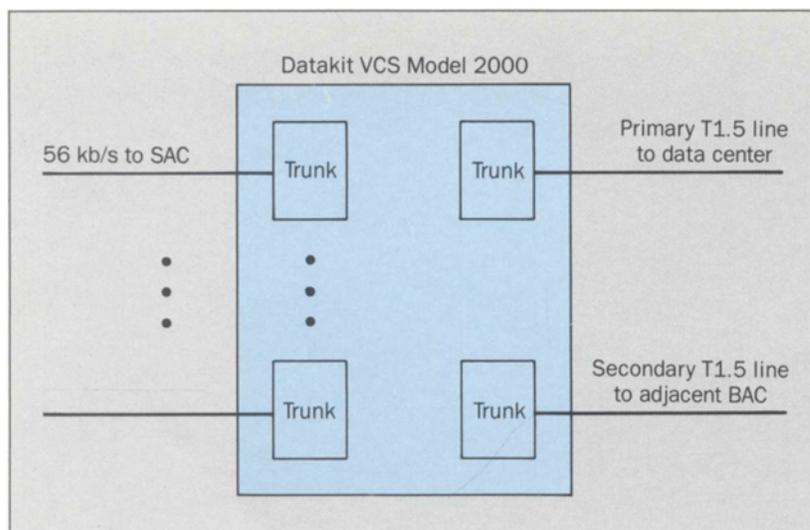
We selected analog facilities with modems that operate at 19.2 kb/s where the traffic allowed, and selected 56-kb/s DDS for the remainder. These are short-distance access facilities and mostly intra-LATA (i.e., within the local-access transport area). The bandwidth requirement from the SAC to the BAC depends on the

number of branches that home into the SAC and the traffic the SAC generates. The SACs are homed into the BACs by 56-kb/s DDS facilities.

Backbone-access center. A BAC acts as the major hubbing point for the SACs and routes traffic to the data centers. BACs are located at an AT&T central office that offers continuous, 24-hour operations. Many central offices are unattended during off-shift hours.

The BACs have the switching fabric (a Datakit VCS Model 2000) needed to route data to any data center on the network during normal operations. In disaster mode, the BACs can rehome data flow from the failed data center to the backup data center. They also route traffic

Figure 7. Backbone-access center (BAC) configuration. SAC is a service-area concentrator.



100

around a backbone link or node failure. The Datakit VCSs at BACs contain only trunk modules, DDS trunk modules to connect to the SACs, and T1.5 trunk modules to connect to the backbone network. Figure 7 shows a typical BAC configuration.

Backbone network. The ITN route module was used to configure the optimal backbone network. The module finds the most cost-effective set of services and routes for all the requirements between the BACs. When evaluating the backbone network, we considered a mix of DDS and Accunet T1.5 services. To transport traffic from the branches to the data centers and still meet BofA's end-to-end response-time requirements, we determined that Accunet T1.5 service—which provides 1.544 Mb/s—was the only feasible choice for the backbone links.

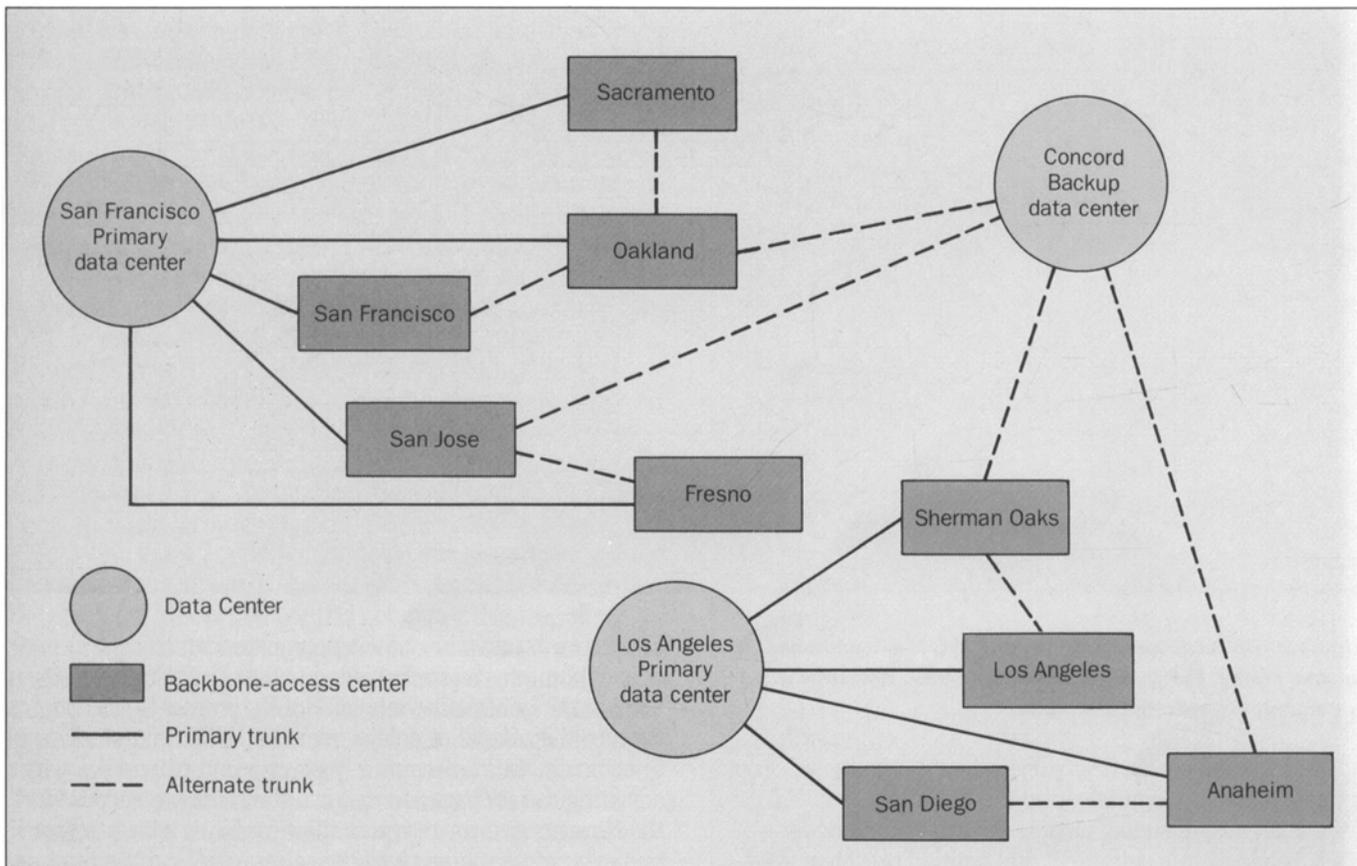
Each virtual circuit is assigned both a primary route and a secondary route. These routes are provided over link-disjoint paths; i.e., they do not have a common node pair. Therefore, the network will transport traffic from any BAC to the data center, even if a link or node failure occurs along the path.

Each BAC's traffic is routed via a direct 1.544-

Mb/s link to its primary data center because of the high volume of traffic carried from each BAC to its primary data center. An alternate route via another BAC is provided to increase network availability. To satisfy the COP requirement, each BAC has an alternate route to the backup data center in Concord. However, the number of hops for the alternate routes and COP is constrained to a maximum of two. Figure 8 shows the backbone configuration.

Data center. The bank has two primary data centers, San Francisco and Los Angeles, which are backed up by a third data center in Concord. The data centers house mainframe computers that serve the applications. A data center represents an end point in that it generates traffic to the applications' controllers at the branches and receives data from these devices. All data-center nodes are interconnected in a mesh by lightwave fiber trunks.

As Figure 9 shows, there are two possible arrangements for the data center. In one arrangement, data is transported at RS-232 port levels through a matrix switch to an FEP. This may involve long runs of bulky EIA (Electronic Industries Association) cabling that may exceed RS-232 distance limits. The other arrangement



provides data transport through a fiber trunk to a small Datakit VCS node or SAM located next to the FEP. This will provide the switching and access functions of a matrix switch, while eliminating the long lengths of EIA cabling. Another advantage of this arrangement is that it uses Datakit VCS technology common to CDN and is under common StarKeeper NMS control, thus increasing manageability and reliability.

Each data center is configured to support its region. For example, Los Angeles supports the southern region, while San Francisco supports the northern one.

Figure 8. Backbone-access centers.

Concord is configured to support either data center in a disaster. StarKeeper NMS is located at Concord.

Performance Evaluation. The final step in the network design process is to evaluate performance and fine-tune the configuration. Delay considerations strongly affect the sizing of access and backbone links and the routing along the backbone network. The response-time-analysis (RTA) module of the E-INOS Data subsystem is used to predict the end-to-end delays and perform parametric analysis on

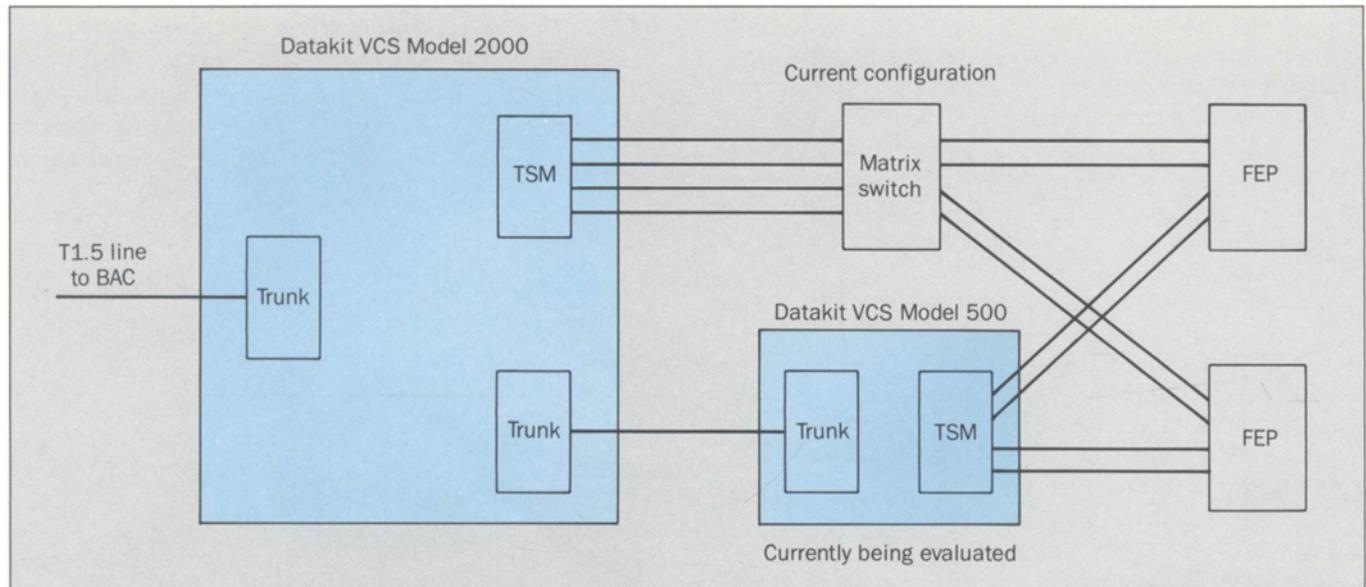


Figure 9. Data center configuration. BAC is a backbone-access center, FEP is a front-end processor, and TSM is a transparent synchronous module.

several variables—such as a line's speed, number of drops on a virtual multipoint line, and the line's protocol. RTA uses analytical queueing models for several protocols—such as SDLC (point-to-point and multipoint), binary synchronous, and polled-mode asynchronous. *Response time* is defined as the time elapsed between when a user depresses the ENTER key on a terminal and when the terminal receives the last character of the host's response.

Because of the high-speed backplanes of the Datakit VCS nodes and the use of pipelining, which permits data transmission to begin before an entire frame is received, the delays that the Datakit™ network introduces are an insignificant portion of the overall end-to-end delay. The access speed and the protocols between the end devices (CC and FEP) are the primary factors in delay.

Tests of multipoint lines show that transporting messages across a Datakit network adds only 2 to 8 per-

cent more transaction time as compared to a direct connection (physical wire). If we adjust the parameters used in the generic models for BofA's protocols, we can capture the effects of delays over the Datakit network. Specifically, we can accurately predict end-to-end delay by adjusting the propagation delay in the model to account for the time needed to transport the first bit of a frame from one edge of the network to the other edge and the time needed to process and recognize a frame at the CC or FEP.

Figure 10 shows the results on delay for varying the link speed between the CC and the FEP, and the effects of adding CCs to a virtual multipoint circuit. These results are for an SDLC line, and typical message sizes of 300 characters from the CC to the host and a 750-character response from the host. The end-to-end delay calculations do not include host delays.

Elsewhere in this issue, Sheng presents a more comprehensive performance analysis of generic reference connections for virtual private-line networks.⁴

Network-Management-System Design. As we have dis-

cussed, various components of the Datakit VCS and Dataphone II product lines are used throughout the network to provide an optimum data transport solution. This involves modems (Models 2096, 2144, and 2192), DSUs (Model 2656), and Datakit VCS hardware (Models 2000, 500, and SAM). Designing a network-management system for a network of this size and diversity, while also addressing special BofA needs, represents a challenge. The AT&T UNMA provides a framework for such network management.

Figures 11 and 12 illustrate the solution selected for the Bank of America. Figure 11 shows the SAC-level connectivity of diagnostic or alarm channels associated with the Datakit VCS and Dataphone II equipment. Diagnostic or alarm information from branch-office modems or DSUs is carried back to the SAC-resident modem or DSU over the same facilities that carry user data (i.e., "in-band"). Similarly, branch-office SAM diagnostic or alarm information is carried back in-band to the SAC-resident, Datakit VCS Model 500 control computer. Hence, all appropriate equipment-diagnostic information (from remote branch offices) that is being homed to this SAC site can be accessed through the Datakit VCS Model 500 or a local modem or DSU. This consolidated diagnostic and alarm information reenters the data network via a standard Datakit VCS asynchronous connection (a TY module). From here, the information flows out over either the primary trunk to a BAC or a secondary trunk to an adjacent SAC and into the backbone-data network.

BofA's network-management center is located at the backup data center in Concord (Figure 12). All (Datakit VCS and Dataphone II) network diagnostic and alarm information is delivered to the two Datakit VCS Model 2000 nodes at Concord. Four network-management systems (two StarKeeper NMSs and two Dataphone II level-IV system controllers) are required to monitor this large network. Each system is connected to both Datakit VCSs to avoid a major network-management-system disruption if either Datakit VCS or its associated T1.5 trunks fails. Other StarKeeper NMS and Dataphone II level-IV system controllers, normally used for testing and special develop-

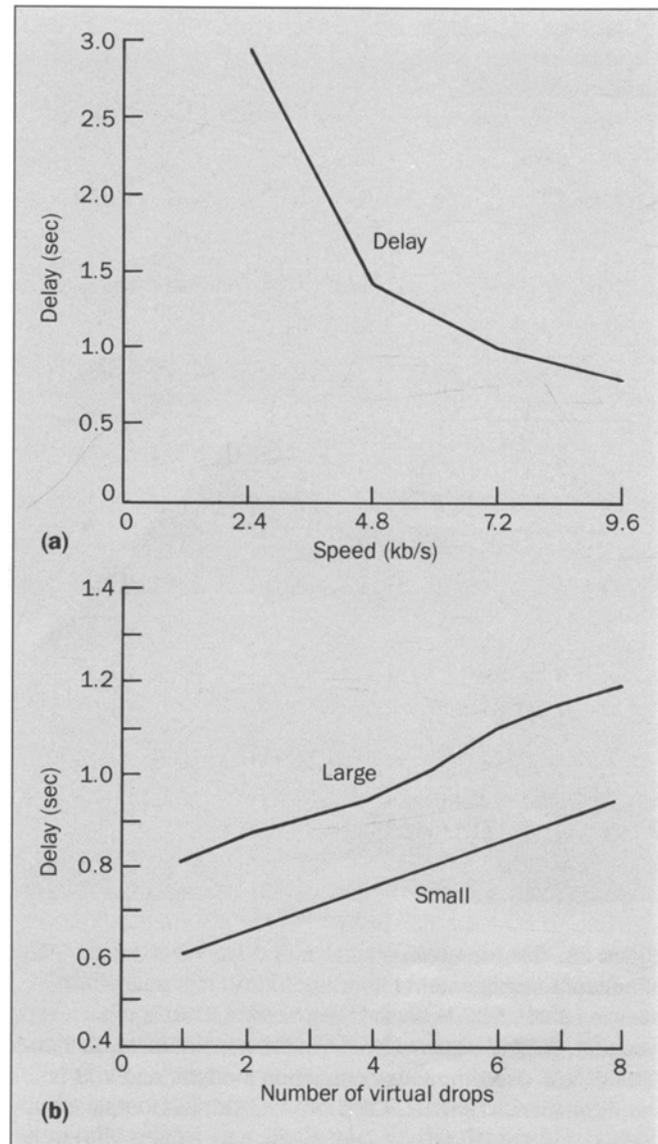


Figure 10. Performance results. (a) Delay versus speed on a synchronous data-link control line. (b) Delay versus number of virtual drops (synchronous data-link control line at 9.6 kb/s).

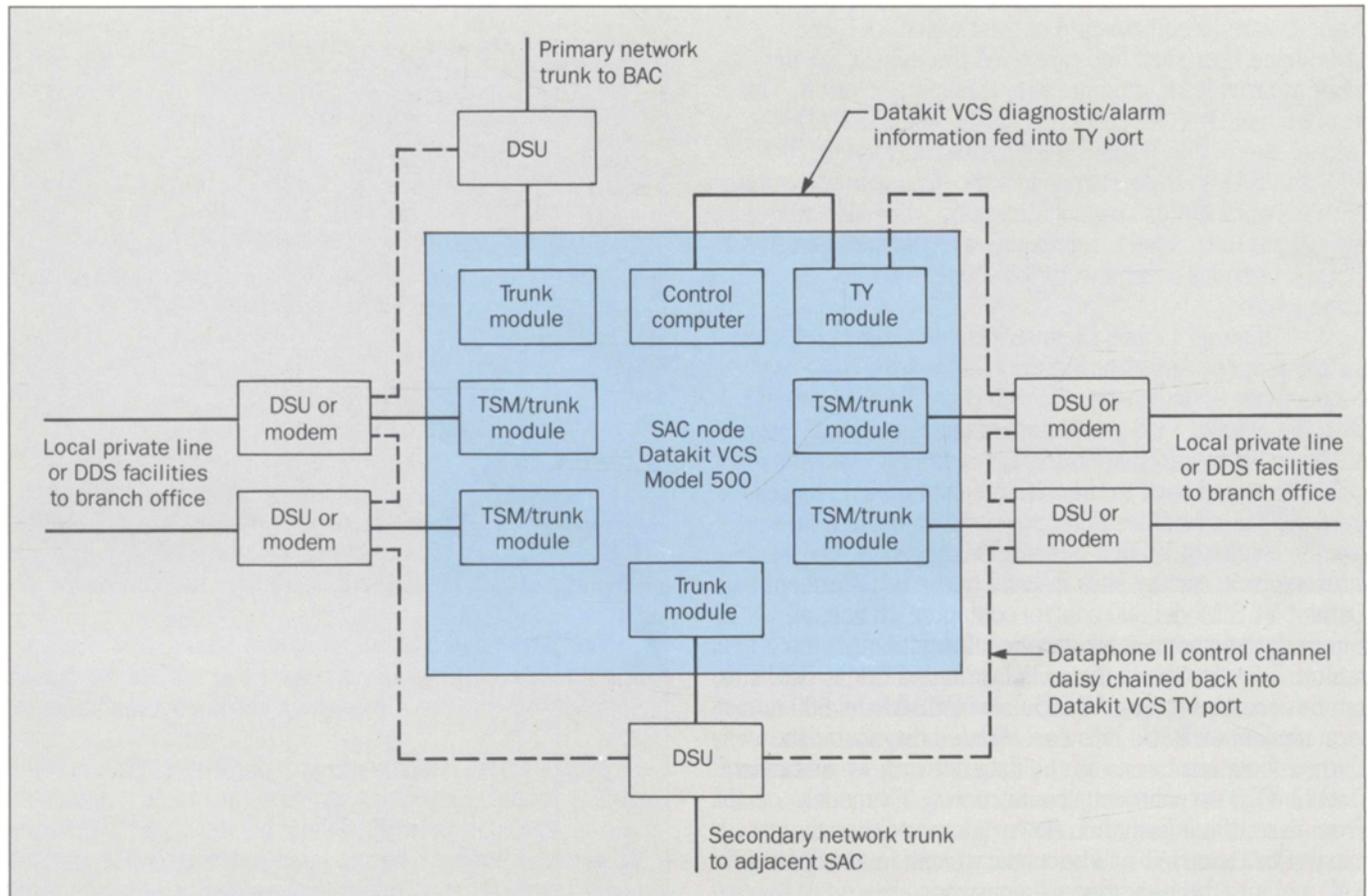


Figure 11. Service-area-concentrator (SAC) level connectivity of network-management information. BAC is a backbone-access center, DDS is digital data service, DSU is a data service unit, TSM is a transparent synchronous module, TY is a Datakit VCS asynchronous-connection module, and VCS is the virtual-circuit switch.

ment, are available to act as backup if a main unit fails.

The resulting network-management system is both comprehensive in its coverage and robust in various failure scenarios.

Contingency Operations Design. Under disaster conditions, Bank of America needs the ability to redirect to the backup data center at Concord all data traffic that normally homes on a main data center. This may be necessary to respond to an earthquake, a fire, or massive equipment failure. This function is well-suited to a Datakit VCS/

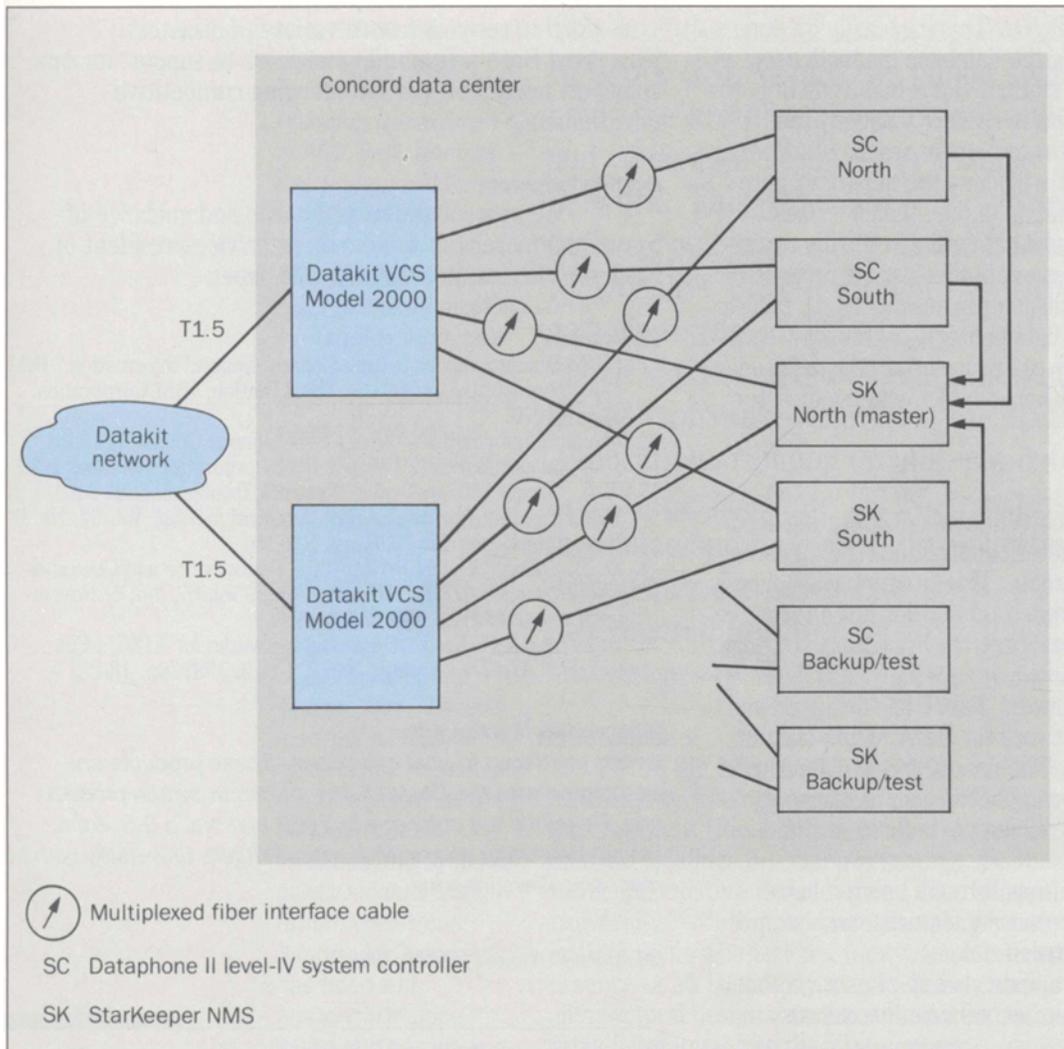


Figure 12. Bank of America's network-management center at Concord, California.

105

StarKeeper NMS network.

Normally, data that enters the network at any point (e.g., a branch office) receives the address of the predefined destination (PDD) selected for that circuit. Because point-to-point or multipoint private lines previ-

ously supported most of the traffic, this PDD address generally relates to a single port on the data-center Datakit VCS and, therefore, a single FEP port. To reroute this port's traffic to another FEP in a backup data center requires only that the PDD assigned at the entering

Datakit VCS module be changed. This can easily be done for a single port by changing the database manually from a StarKeeper NMS operator station. But if hundreds of ports are involved, as might happen in a major disaster, this process could become lengthy and error prone. StarKeeper NMS provides a solution: It will allow the network operator to load a data file in a "batch" mode, which it uses to drive a routine that changes all databases in the appropriate Datakit VCS nodes. These data files can be prepared in advance to match possible disaster scenarios (e.g., failure of the entire San Francisco data center). In this way, each file can be checked carefully to ensure accuracy and then quickly invoked, using StarKeeper NMS, when a disaster occurs.

Conclusion

Bank of America's current mode of operating separate analog networks for different applications is expensive, difficult to manage, and inflexible. This situation will become even more complicated as branch automation and alternative banking systems are introduced into the branches to off-load activities from branch personnel.

In response to this need, AT&T has designed and implemented a new data network for BofA, using Datakit VCS and StarKeeper NMS. The new network will provide:

- Cost containment by allowing better use of transport bandwidth and lowering training costs for operations staff
- Enhanced network operations through using reliable switch hardware, trunk rerouting features, and comprehensive network management
- Flexibility by supporting a wide variety of user protocols and providing the ability to reconfigure the network on

demand to recover from a variety of disasters. This gives BofA a true utility network to support its data transport needs in a quickly changing competitive environment.

Acknowledgment

We appreciate the assistance and guidance of Steve Westmoreland, system director/vice president of Bank of America, in preparing this paper.

References

1. "IBM Synchronous Data Link Control General Information," IBM Document No. GA27-3093-2, Third Edition, IBM Corporation, March 1979.
2. "General Information—Binary Synchronous Communications," *IBM Systems Reference Library*, IBM Corporation, October 1970.
3. T. E. Marques, "StarKeeper® Network Troubleshooter: An Expert System Product," *AT&T Technical Journal*, Vol. 67, No. 6, November/December 1988, pp. 137-154.
4. D. D. Sheng, "Virtual Private-Line Performance and Customer Cost Impacts," *AT&T Technical Journal*, Vol. 67, No. 6, November/December 1988, pp. 47-68.
5. S. Lin and J. H. Rath, "Designing Networks for AT&T's Customers," *AT&T Technology*, Vol. 2, No. 2, 1987, pp. 18-25.

Biographies (continued)

plexed interfaces to host computers. These products are associated with the Datakit® virtual-circuit-switch product line. He joined the company in 1969 and has a B.S. from Lehigh University and an M.S. from Rutgers University, both in mechanical engineering.

(Manuscript received July 29, 1988)