

VIRTUAL PRIVATE-LINE PERFORMANCE AND CUSTOMER COST IMPACTS

Diane D. Sheng

Diane D. Sheng is a supervisor in the Data Network Analysis Department at AT&T Bell Laboratories in Holmdel, New Jersey. Her group is involved in the competitive assessment and technical evaluation of potential AT&T data networking products and services. She joined the company in 1978 and has a B.S. in mathematical sciences, an M.S. in statistics, and a Ph.D. in operations research, all from Stanford University.

There has been considerable activity in the data networking area to develop architectures for virtual private-line networks as alternatives to private-line networks. This paper provides a comparative performance report of virtual private-line networks and private-line networks for SNA™/SDLC and 3270 BSC applications. (SNA is a trademark of IBM Corporation.) In particular, we discuss the effects on end-to-end performance of various virtual private-line protocol-handling capabilities in terms of the degree to which they produce better or worse network performance levels than private-line networks. Beyond performance optimization, it is also important to understand the customer's view of the economic impact of the various protocol-handling capabilities of virtual private-line networks. So, we also discuss the effects of these virtual private-line networking capabilities and suggest that they can lower a customer's facilities, modem, and data service unit costs.

Introduction

In this paper, we investigate the performance that can be achieved with virtual private-line (VPL) networks as compared to private-line (PL) networks. The virtual private-line networks addressed here are based on the integrated data architecture (IDA) model presented elsewhere in this issue,¹ and the measure of performance tracked is, primarily, the end-to-end response time that end users see. For comparison, we give the end-to-end response times experienced with today's voice-grade private-line (VGPL) networks and Dataphone® digital service PL networks, a digital data service (DDS) network. (Panel 1 defines acronyms and terms used in this paper.)

In looking at VPL performance, it is of interest to isolate the effects of the different protocol handling capabilities that are potentially

Panel 1. Acronyms and Terms in This Paper	
ACK	acknowledgment
BSC	binary synchronous communications
CC	cluster controller
CPE	customer-premises equipment
DDS	digital data service
dly_{inq}	queueing time at the CC for an inquiry
dly_{resp}	queueing time at the FEP for a response
DSU	data service unit
EOT	end of transmission
FEP	front-end processor
Globecom	IEEE Global Telecommunications Conference
ICC	IEEE International Conference on Communications
IDA	integrated data architecture
IEEE	Institute of Electrical and Electronics Engineers
INQ	inquiry message
LAPD	link-access procedures for D channel
LEC	local exchange carrier
PL	private line
POP	point of presence
RESP	response message
SDLC	synchronous data link control
SO	serving office
SW	frame-relay switch
TA	terminal adapter
TA_H	terminal adapter at the host end
TA_T	terminal adapter at the terminal end
$trns_{inq}$	transport time for the inquiry from CC to FEP
$trns_{resp}$	transport time for the response from FEP to CC
VGPL	voice-grade private line
VPL	virtual private line
WC	wire center

implemented within the terminal adapters (TAs), a key component in the IDA framework. In this paper, we look at the performance of VPL networks where the TAs are combined with multiplexing functionality, either on the customer's premises or within carrier points-of-presence (POPs) at the edges of the high-speed, frame-relay network. We focus on synchronous communications and—despite the trend toward SDLC—quantify the performance of both SDLC and BSC applications, because BSC applications still represent a large segment of the data communications embedded base. (SDLC is synchronous

data link control, and BSC is binary synchronous communications. They are data-link-layer protocols.^{2,3}) However, we confine our analysis to 3270 BSC as an example of standard BSC operation. (The 3270 is IBM's synchronous terminal/cluster controller system.) Of course, analysis of 3270 BSC may not necessarily cover the PL and VPL performance of numerous variants of BSC.

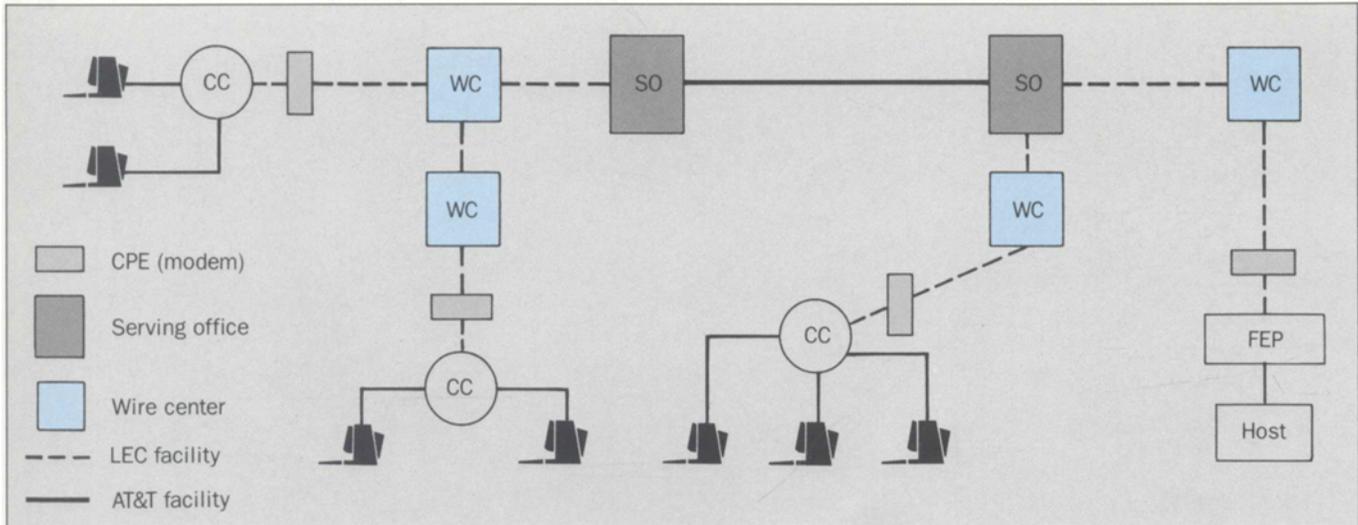
Beyond the end-to-end performance effects of various TA capabilities, it is also important to understand how providing those capabilities affects a customer's VPL costs. To illustrate some of that impact, we also look at the potential effects that VPL networking capabilities have on customer networking costs. As with performance, we view these cost effects in comparison to VGPL and DDS PL networking costs. Here, we focus only on transport costs and not on the comparative costs of providing feature functionality—such as network management or disaster recovery—with PL and VPL networks.

In laying out the performance comparisons between VGPL and PL networks and the related cost effects, we use a "slice" of a hypothetical wide-area networking customer as a reference point. Next, we present this customer model and the salient characteristics of the PL and VPL networks analyzed. Then, we discuss individual protocol-handling capabilities and quantify their effects on customer-perceived performance and costs. We do this for VPL networks that use VGPL and/or DDS access arrangements with a frame-relay backbone network.

Private-Line Customer Networking Model

Figure 1 is an abstract illustration of private networks today. As depicted, most private networks provide communications between a large-business customer's many remote sites and its handful of data centers. Today, these customers typically use VGPLs for transport.

The most cost-effective use of VGPLs involves multidrops of several collocated (i.e., in the same place) and/or noncollocated cluster controllers (CCs) onto one private line. Such multidrop private lines typically are physically bridged in AT&T serving offices; however, physical bridging in local-exchange-carrier (LEC) wire centers, or



even on a customer's premises, is also possible. For security and performance reasons (namely, response time enhancement), some customers choose not to use multi-drop PLs. Instead, they use point-to-point PLs, where each CC makes up a single tail circuit (a short private line on the CC side). In addition, for reliability, many customers use DDS PLs for transport. With both VGPL and DDS PLs, the predominant speeds used today are 4.8 and 9.6 kb/s (kilobits per second).

Throughout this paper, we use a hypothetical wide-area network customer for discussion purposes. Figure 2 pictures a portion of this customer's total private-line network. This portion involves sixteen PLs; eight transport SDLC transactions, and eight transport BSC transactions to and from the same FEP (front-end processor). Among the eight SDLC PLs are:

- One 9.6-kb/s and one 4.8-kb/s multipoint VGPL
- One 9.6-kb/s and one 4.8-kb/s point-to-point VGPL
- One 9.6-kb/s and one 4.8-kb/s multipoint DDS circuit
- One 9.6-kb/s and one 4.8-kb/s point-to-point DDS circuit.

The same composition applies for the eight BSC PLs.

We assume that each multipoint line has 5 drops

Figure 1. A typical private-line network. CC is a cluster controller, CPE is customer-premises equipment, FEP is the front-end processor, LEC is a local-exchange carrier, SO is a servicing office, and WC is the wire center.

and all 16 of the PL circuits are 1200 miles long. Each of the customer's remote sites has both SDLC and BSC data communications needs. The portion of the customer's network shown has 14 remote sites in all, and the 48 CCs are distributed among these remote sites. Individual sites have from one to eight CCs.

We selected the specific PL configurations in Figure 2 to illustrate the various levels of performance that can be achieved with PLs and VPLs, as well as their different aspects for customer networking costs. We can view Figure 2 as a representative slice of many large networks today. Each such network may typically have anywhere from 5 to 25 times as many remote sites and CCs.

Private-Line Performance. A key performance measure is the response time that a user sees. We define response time here as the time between arrival of the last byte of an inquiry (INQ) message at the CC and receipt of the last byte of the response (RESP) message back at the

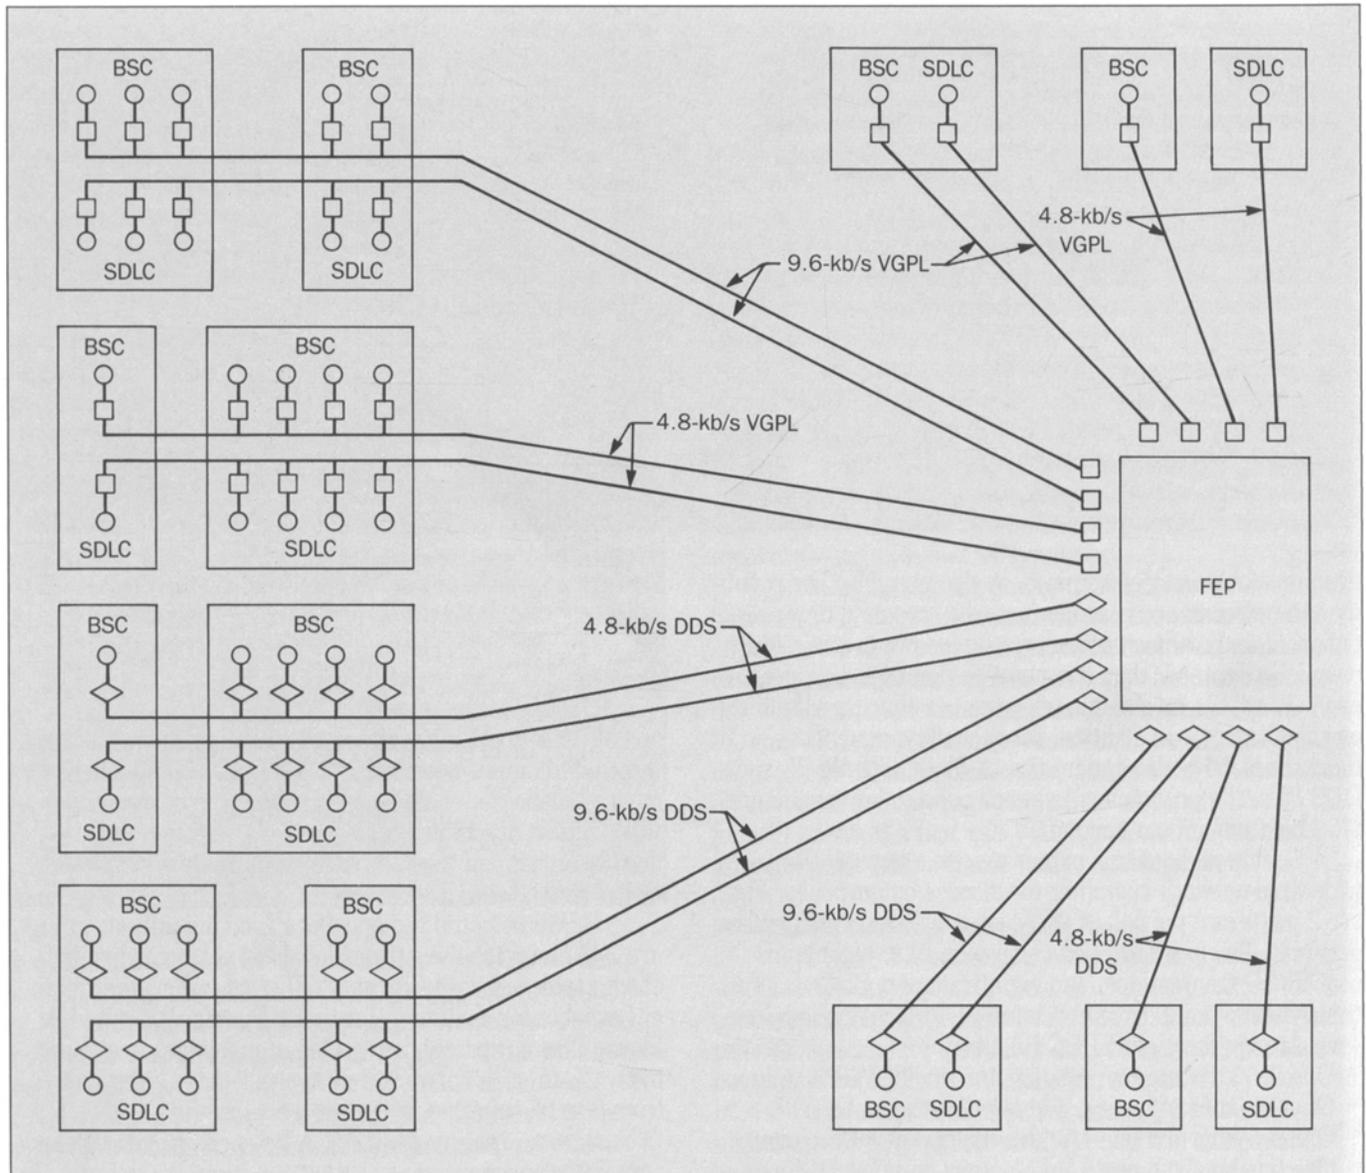


Figure 2. Portion of a customer's network. BSC is binary synchronous communications, FEP is a front-end processor,

SDLC is synchronous data link control, DDS is digital data service, and VGPL is a voice-grade private line.

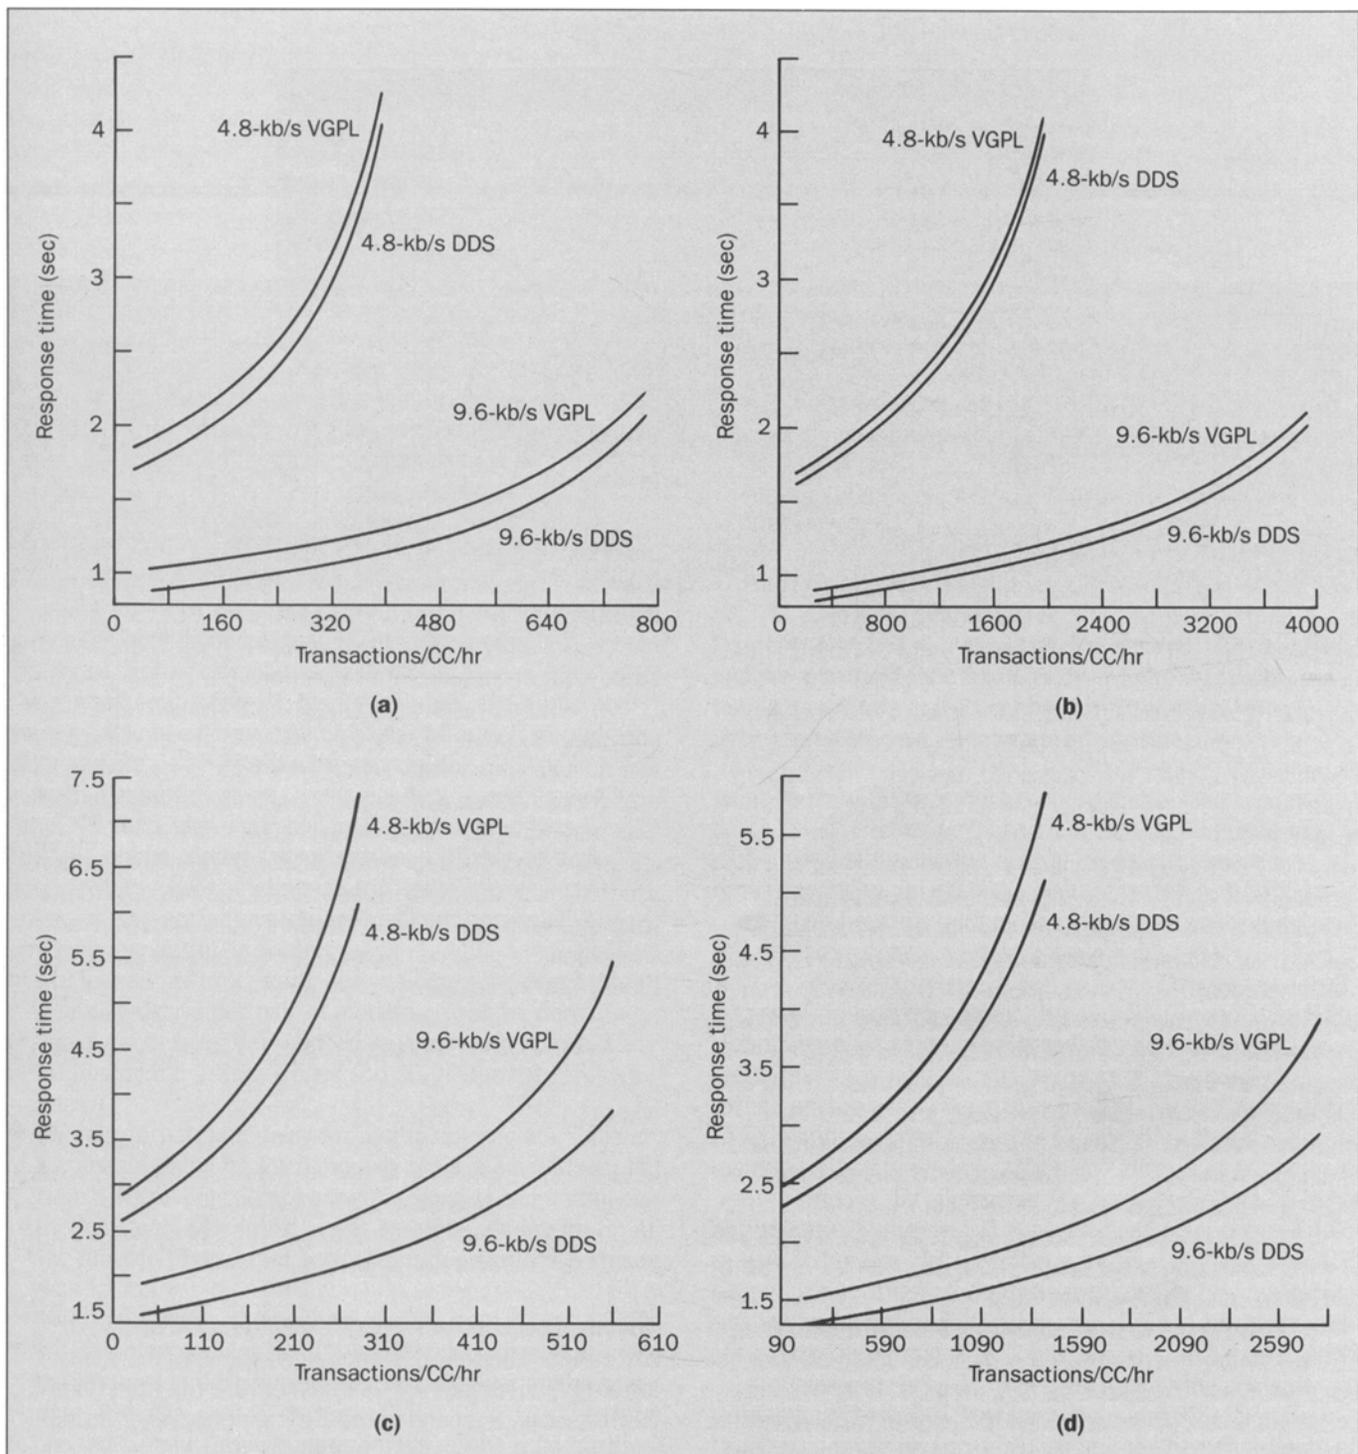


Figure 3. Private-line performance for (a) multipoint and (b) point-to-point SDLC, and for (c) multipoint and (d) point-to-

point BSC. CC is a cluster controller, DDS is digital data service, and VGPL is a voice-grade private line.

Table I. Customer PL and VPL Facilities and Equipment Costs

Facilities	Cost per month (\$)		Equipment	One-time cost (\$)	
	PL	VPL		PL	VPL
VGPL	14,657	16,614	Modems	80,000	120,000
DDS	14,697	16,575	DSUs	32,000	48,000
Total	29,354	33,189	Total	112,000	168,000

NOTE: DDS = digital data service; DSU = data service unit; PL = private line; VGPL = voice-grade private line; VPL = virtual private line.

CC. Excluded are host, FEP, and CC processing times and the communication delays between the end (terminal) devices and the CCs and between the FEP and the host. These processing times and delays depend highly on the specific devices used and are constant in their contribution to both PL and VPL performance. To determine the total transaction time a user sees, add fixed communication delays and host or terminal processing times to the response time estimates presented here. [CC, FEP, and host processing times vary greatly across applications and devices. Host processing times can be a major contributor to total response times and can typically vary from 1 to 7 seconds. FEP and CC processing times can typically vary from 50 to 250 ms (milliseconds) and from 1 to 50 ms, respectively.]

Figure 3 shows typical response times for the different SDLC and BSC circuits in the customer model described above. In Figures 3a and 3b, the particular SDLC transaction modeled consists of an 80-byte INQ and a 810-byte RESP, typical in interactive transactions seen in the banking industry. The INQ and RESP thus correspond to one and three individual SDLC frames, respectively. We assume that all SDLC CCs are half-duplex CCs (i.e., the CC cannot receive while it is sending, nor send while it receives). In Figures 3c and 3d, the BSC transaction modeled consists of a 100-byte INQ and a 1000-byte RESP, where each is transported as one block.

The transaction rates in Figure 3 cover a 5- to 75-percent utilization range of the PL in its most heavily used direction, e.g., from FEP to CC.

In Figure 3a, for example, a 5-percent utilization of the 4.8-kb/s multipoint PL corresponds to 26 transactions per hour for each of the five drops, and 75 percent corresponds to 391 transactions per hour per CC. The 9.6-kb/s multipoint PL can carry double the amount of traffic.

In Figure 3b, 5-percent use of the 4.8-kb/s point-to-point PL corresponds to 131 transactions per hour and

75 percent corresponds to 1964 transactions per hour. Similarly, 5-percent use of the 9.6-kb/s point-to-point PL occurs at 262 transactions per hour and 75 percent occurs at 3927 transactions per hour.

In Figure 3c, 5-percent utilization of the 4.8-kb/s multipoint PL corresponds to 19 transactions per hour per CC and 75 percent corresponds to 280 transactions per hour per CC. For the 9.6-kb/s multipoint PLs, 5-percent use occurs at 37 transactions per hour per CC and 75 percent occurs at 560 transactions per hour per CC.

In Figures 3b and 3d, for both the 4.8- and 9.6-kb/s point-to-point PLs, 5- and 75-percent use occurs at five times as many transactions per hour per CC as in Figures 3a and 3c because there are now one-fifth as many CCs per PL.

To produce these PL response-time analyses, we used approximate analytical performance models^{4,5} derived from queueing theory and experience based on software simulations of multipoint circuits.⁶⁻⁹ Mean PL response time is characterized by:

$$dly_{inq} + trns_{inq} + dly_{resp} + trns_{resp} \quad (1)$$

where

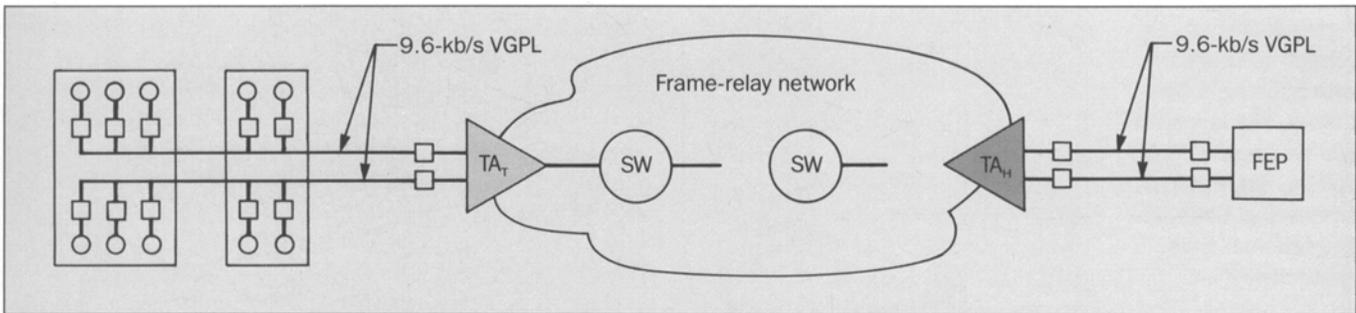
dly_{inq} = queueing time at the CC for an inquiry

$trns_{inq}$ = transport time for the inquiry CC-to-FEP

dly_{resp} = queueing time at the FEP for a response

$trns_{resp}$ = transport time for the response FEP-to-CC.

The main contributors in expression (1), are dly_{inq} , which covers the time that INQ messages spend in the CC waiting to be polled, and dly_{resp} , which covers the time that RESP messages spend in the FEP waiting for selection and start of transmission back to the CC. The transport time contributors in expression (1), $trns_{inq}$ and $trns_{resp}$, depend on the speed of the PL, size of the INQ and RESP



messages, and distance-dependent propagation delay of the PL connection.

In estimating PL response times, the performance models used here pay particular attention to capturing dly_{inq} and dly_{resp} . To do this, they model the FEP polling and selecting activities as quasi-independent queueing systems. In particular, FEP polling is represented by a multiple-queue, cyclic-server system with server walk-times, while FEP selecting is represented by a single-queue, single-server system with server vacations. These queueing models incorporate:

- Modem signal-processing delays. These are constant one-way delays per pair of modems. (The modem signal-processing delays used in generating Figure 3 are 20 ms for the 4.8-kb/s lines and 30 ms for the 9.6-kb/s lines.)
- Modem turnaround delay for multipoint circuits. This is a constant delay for each modem at the CC end from the time a secondary device raises request-to-send until the modem replies by raising clear-to-send. (In Figure 3, the modem turnaround delay for both 4.8 and 9.6 kb/s was 15 ms.)
- $trns_{inq}$ and $trns_{resp}$ themselves.
- Effects of important SDLC and 3270 BSC protocol parameters and FEP system generation parameters on dly_{inq} and dly_{resp} .

(Cole and Kumar provide more details on the relationships among these various factors.^{4,5})

In Figure 3, we see better response times for

Figure 4. A typical virtual private-line network. FEP is a front-end processor; SW is a frame-relay switch; TA_H and TA_T are terminal adapters at the host and terminal ends, respectively; and VGPL is a voice-grade private line.

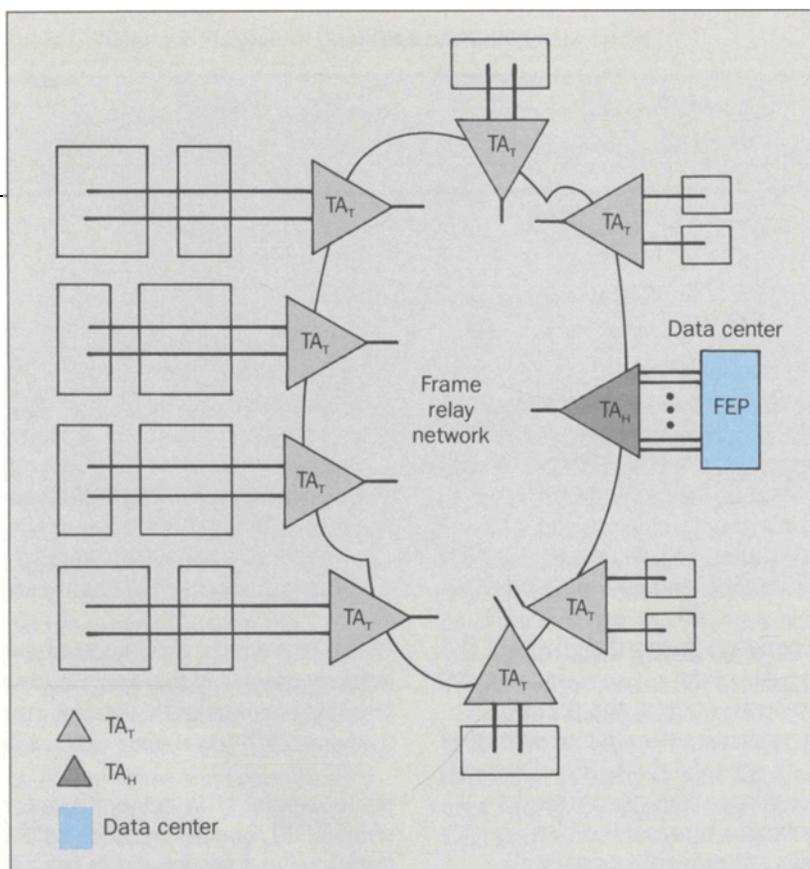
point-to-point DDS circuits than for same-speed, point-to-point VGPL circuits because DDS circuits eliminate the modem signal-processing delays. The better response times for multipoint DDS circuits compared to VGPL circuits result because both modem signal-processing delays and modem turnaround delays are eliminated.

When we compare Figure 3d to 3c, we see better response times for the point-to-point VGPL and DDS BSC circuits than for same-speed, multipoint, VGPL BSC circuits—similarly a result of elimination of the modem turnaround delays. This is also evident if we compare the SDLC point-to-point results of Figure 3b to the SDLC multipoint results of Figure 3a.

Finally, by comparing Figure 3c to 3a and 3d to 3b, we can see the overall superior performance of SDLC to BSC. The considerably worse response times for BSC circuits reflect the half-duplex nature of the BSC protocol and the much larger BSC block used here (1000 bytes) relative to the maximum SDLC frame size (256 bytes).

Private-Line Costs. When we look at the costs of private-line networks, key components are the monthly tariffs paid for transport facilities and the one-time costs of the customer's modems and data service units (DSUs). Table I shows typical facilities and equipment costs for the

Figure 5. Virtual private-line network with hubbing concentrators. FEP is a front-end processor; TA_H and TA_T are terminal adapters at the host and terminal ends, respectively.



54

customer model of Figure 2; costs of the host, FEP, CCs, and all other end devices are excluded.

In Table I, we have assumed costs of \$2000 and \$3000 for all 4.8- and 9.6-kb/s modems, respectively, and \$1000 for all 4.8- and 9.6-kb/s DSUs. We also assume that the local channel portion of the facilities are, on average, 20 miles long at both the CC and FEP ends of the PLs, and that the interoffice connection portion of the facilities is 1160 miles for each PL.

Virtual Private-Line Networks

Figure 4 shows simple private-line replacement by a virtual private-line network. Here, both a 9.6-kb/s SDLC multipoint tail circuit and a 9.6-kb/s BSC multipoint tail circuit use the same VPL network connection. The TAs at the edges of the frame-relay network (TA_T at the terminal end, and TA_H at the host end) segment and wrap the native protocol messages into LAPD frames for individual transport across the network. (LAPD stands for link-access procedures for D channel.) TA_T and TA_H

also concentrate the traffic from separate virtual circuits onto single physical links between the TAs and the frame-relay switches.

The economical design of a VPL network must reflect optimal placement of TA concentrators and frame-relay switches relative to the locations of the customer's remote sites and data centers.¹⁰⁻¹⁶ To illustrate the performance associated with VPL networking, we will discuss a progression of three possible networking configurations:

- Large-capacity TA concentrators at hubbing sites
- Host TA (TA_H) at the customer's data center
- Host TA at the customer's data center and small-capacity TA concentrators at the customer's remote sites.

TAs Used as Hubbing Sites. The first configuration uses large-capacity TA concentrators as hubbing sites to consolidate traffic from several remote sites onto the backbone network. Figure 5 shows this configuration for our customer model. Here, we assume the hubbing TA concentrators are placed 100 miles, on average, from the remote

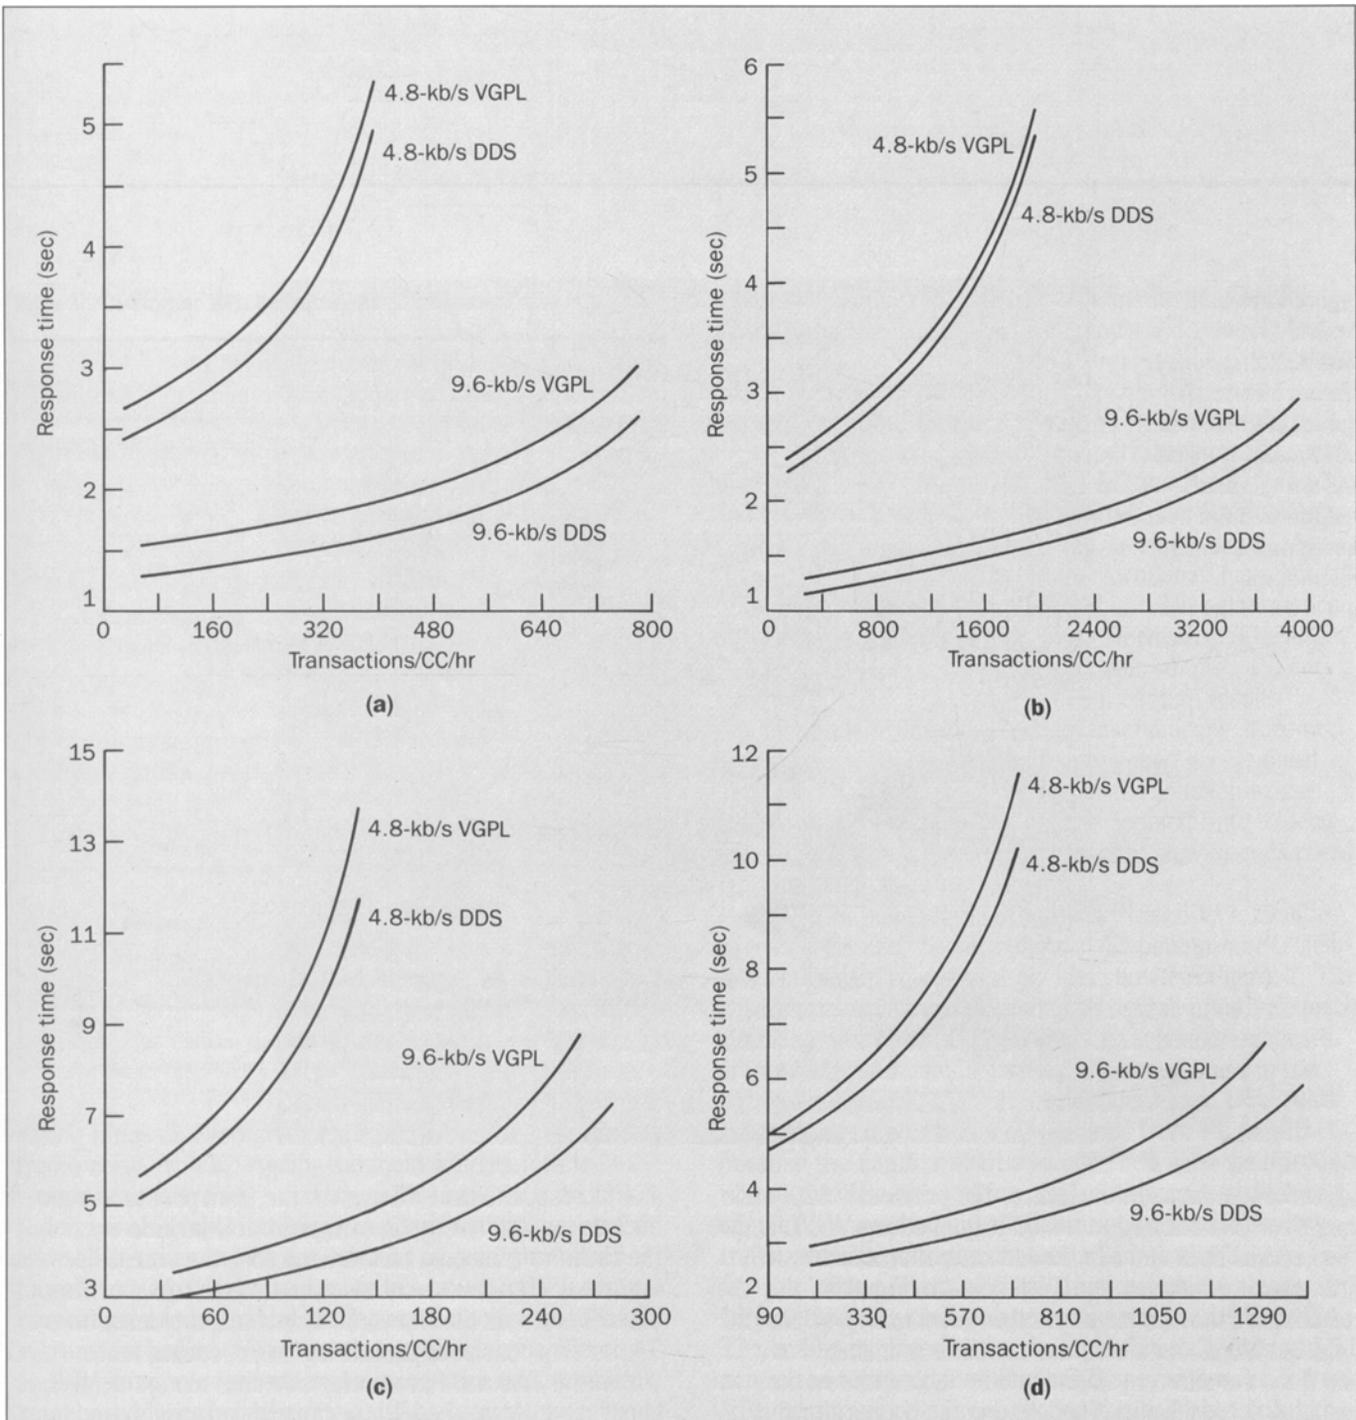
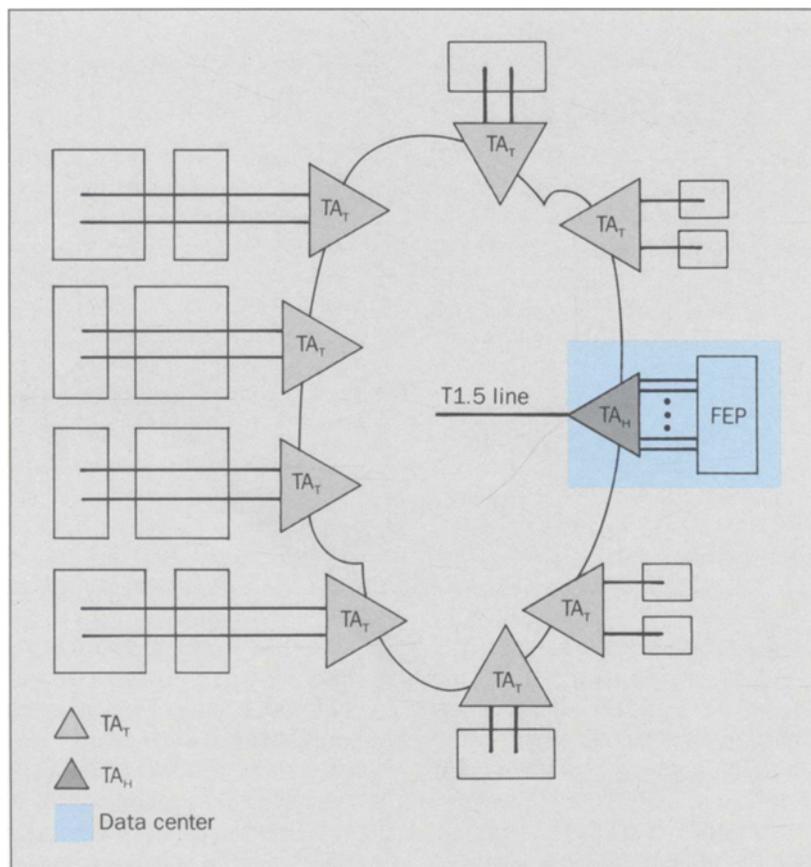


Figure 6. Virtual private-line performance for (a) multipoint and (b) point-to-point SDLC, and for (c) multipoint and (d)

point-to-point BSC. CC is a cluster controller, DDS is digital data service, and VGPL is a voice-grade private line.

Figure 7. Virtual private-line network with hubbing concentrators and data-center premises concentrators. TA_H and TA_T are terminal adapters at the host and terminal ends, respectively.



56

sites and the data center. Hence, the previously 1200-mile long PLs have been replaced by 100-mile long VGPL and DDS access lines and a backbone network consisting of TAs, packet switches, and T1.5 links. In Figure 6, the resulting VPL response times are shown in a and b for the different SDLC circuits in the customer model, and in c and d for the different BSC circuits. Table I shows the resulting access facilities and modem/DSU costs for this VPL networking configuration; as before, the costs of the TA concentrators and switches themselves have been excluded.

Expression (1) represents the mean VPL (as well

as PL) response time. However, the transport time variables, $trms_{inq}$ and $trms_{resp}$, in expression (1) include not only the times to transport messages across the access lines into the VPL network but also the time to transport from TA to TA across the network. The transport times from TA_T to TA_H consist of processing and queueing times within the TAs and frame-relay switches along the VPL connection. Also, the VPL performance models used here address the dynamic fluctuation in these delay components within the network for individual polls, and INQ and RESP messages. (To generate these performance curves, we used LAPD frame sizes of 40 bytes on average and proc-

Table II. Customer VPL Facilities and Equipment Costs

Facilities	Cost per month (\$)	Equipment	One-time cost (\$)
a. Data-center concentration			
VGPL	10,580	Modems	80,000
DDS	10,620	DSUs	32,000
T1.5	7,013	DSUs	6,000
Total	28,213	Total	118,000
b. Remote-site concentration			
VGPL	8,110	Modems	62,000
DDS	9,177	DSUs	22,000
T1.5	7,013	DSUs	6,000
Total	24,300	Total	90,000
c. Remote-site concentration and fanout			
VGPL	9,050	Modems	68,000
DDS	10,132	DSUs	24,000
T1.5	7,013	DSUs	6,000
Total	26,195	Total	98,000

NOTE: DDS = digital data service; DSU = data service unit; PL = private line; VGPL = voice-grade private line; VPL = virtual private line.

essing times on average of 4 ms per TA and 1 ms per frame-relay switch.) Again, the main contributors to response time are dy_{mq} and dy_{resp} . But now their estimation must also capture their dependency on polling cycles and select times, which are affected by transport times from TA to TA. Here, we apply the same modeling structure of quasi-independent polling and selecting queueing systems as before. However, the particular cyclic server and single-server-with-vacation queueing models^{4,5} are augmented to include the effects of transport times from TA to TA.

If we compare Figure 6a to Figure 3a, the performance that the customer sees is about 35 percent worse at 4.8 kb/s and 40 percent worse at 9.6 kb/s for

multipoint SDLC VPL circuits than for the corresponding PL circuits. Similarly, when we compare Figures 6b and 3b, the performance of the point-to-point SDLC circuits is about 44 percent worse for the 4.8-kb/s VPLs and 62 percent worse for the 9.6-kb/s VPLs than for the PLs.

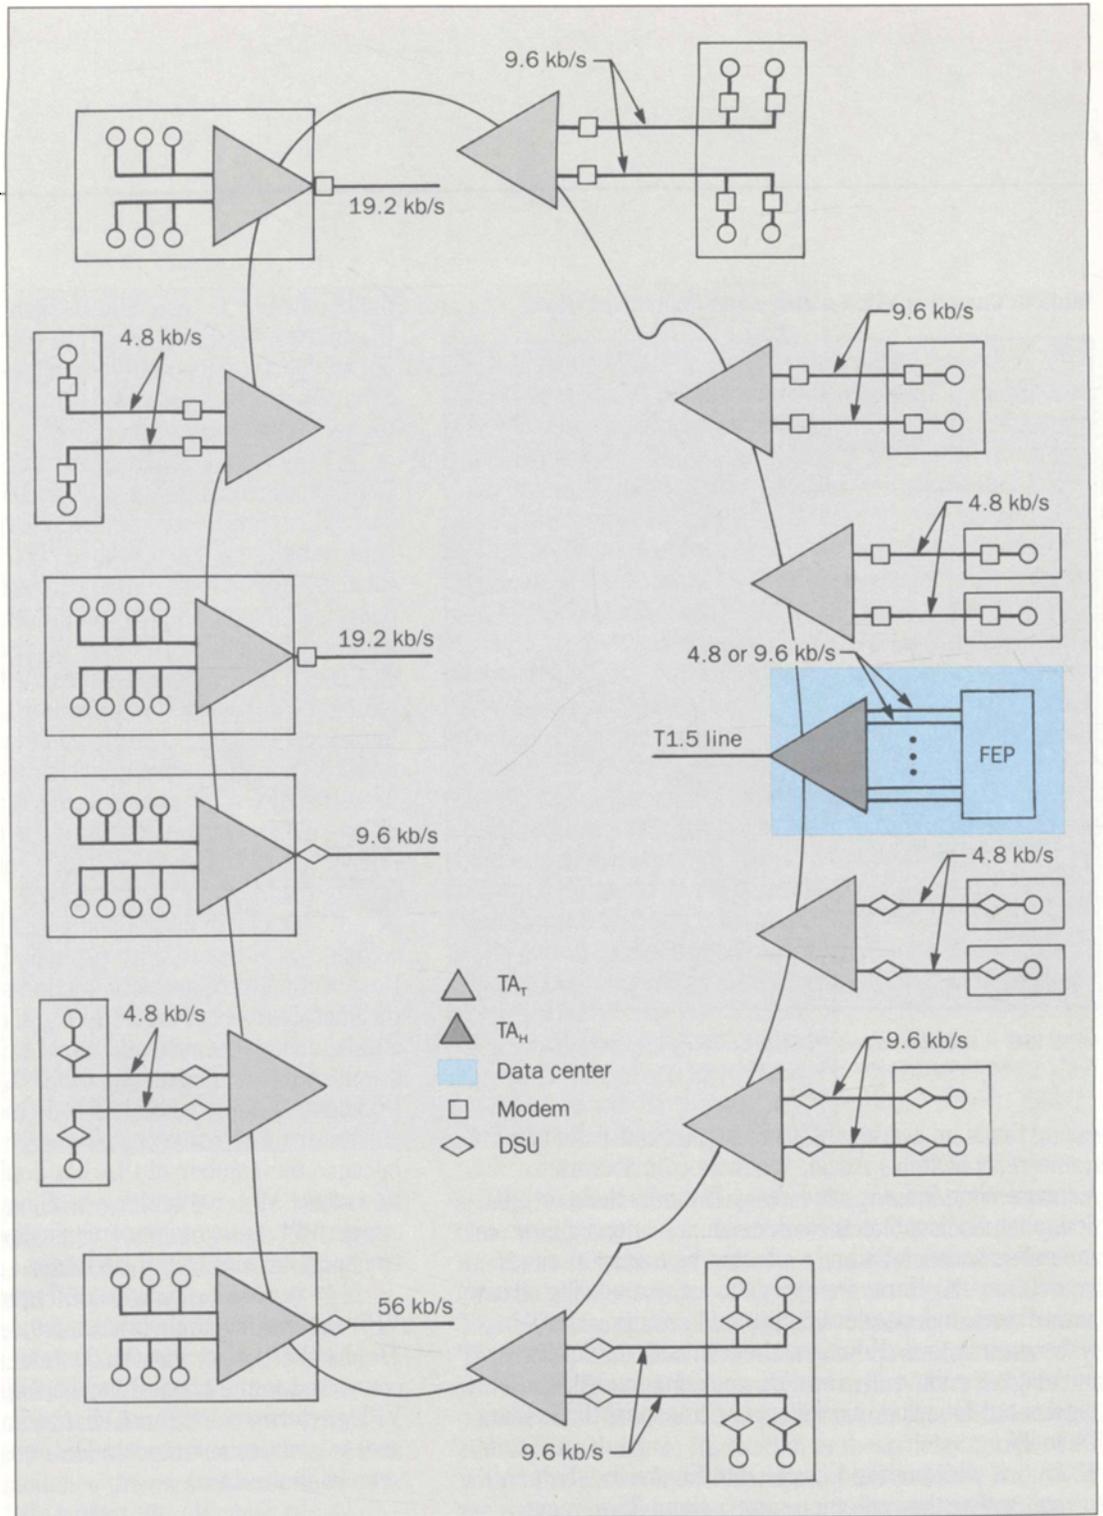
For the BSC circuits, the performance difference between the VPL (Figures 6c and 6d) and PL (Figures 3c and 3d) circuits is comparatively larger. For multipoint circuits at both 4.8 and 9.6 kb/s, VPL performance is more than 80 percent worse than PL performance; for point-to-point circuits at both 4.8 and 9.6 kb/s, VPL performance is about 90 percent worse than PL performance. The larger SDLC and BSC response times for this simple VPL networking configuration relative to PL circuits result fundamentally from the effects of inserting any store-and-forward network between two slow-speed access lines. That the BSC VPL response times are so much worse clearly illustrates the effects of the doubled time to accumulate or insert large BSC blocks onto slow-speed access lines into and out of the network.

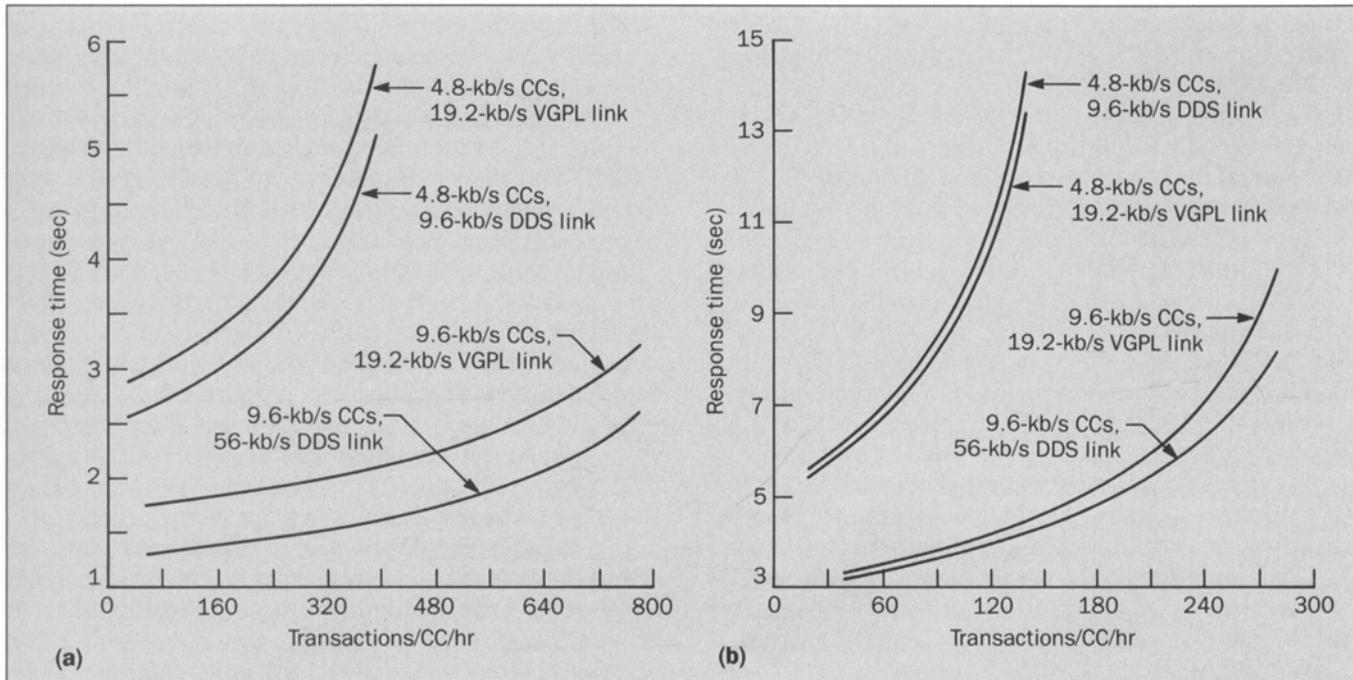
Because Table I compares PL and VPL facilities, we can see a cost penalty of about 13 percent in PL facilities for the entire network portion shown in Figure 2. The facilities costs increase despite the overall greatly reduced PL distances, because additional LEC channel expenses are introduced in connecting the PL access lines to the backbone network. The one-time modem and DSU costs substantially increase (by 50 percent) from PL to VPL, because the number of PLs has doubled with the introduction of the VPL networking configuration. Overall, this simple VPL networking configuration may not offer any transport savings to the customer.

TA Placed at Data Center. Figure 7 shows the second VPL networking configuration we considered. Here, the TA_H has been placed at the customer's data center and connected to the frame-relay backbone via one T1.5 trunk. VPL performance has not changed appreciably, but the access facilities and modem/DSU costs have dropped as seen in Table IIa.

In Table IIa, we used a cost of \$3000 for the T1.5 DSUs. The facilities costs here have dropped 15 percent

Figure 8. Virtual private-line network with hubbing, data-center, and remote-site concentrators. FEP is a front-end processor.





over the facilities costs for VPL in Table I, while the modem/DSU costs have dropped 30 percent over the modem/DSU costs for VPL in Table I. These cost savings reflect the replacement of 16 low-speed access lines to the data center in our customer model with one T1.5 access line. The data center—for most large-business customers—typically handles many more circuits than the 16 in our hypothetical model, which suggests even greater savings in terms of data-center access costs for this VPL networking configuration.

TAs at Data Center and Remote Sites. The last VPL networking configuration considered here involves placing small-capacity TA concentrators at the customer's remote sites as well (Figure 8). For our customer model, customer premises TA₁s have been placed at remote sites that have five or more CCs. These TA₁s have been homed (connected) to the frame-relay switches by 9.6-kb/s VGPL,

Figure 9. Multipoint virtual private-line performance for (a) SDLC and (b) BSC with remote-site concentrators. CC is a cluster controller, DDS is digital data service, and VGPL is a voice-grade private line.

19.2-kb/s DDS, or 56-kb/s DDS links.

What is the effect of introducing a slow-speed 9.6- or 19.2-kb/s link between the customer premises TA₁s and the switches? As we can see by comparing a and b of Figure 9 to a and c of Figure 6, the transaction times increase. Only the multipoint access lines have changed in their VPL configuration, hence Figure 9 still represents the performance of the point-to-point VPL circuits.

As with consolidation of data-center access lines, consolidation of remote-site access further reduces the access facilities and modem/DSU costs (Table IIb). Comparing a and b in Table II, we find an additional 14-percent

savings in facilities and 24 percent in modems/DSUs with remote-site-access consolidation. Now, compare the costs in Table IIb for this configuration to the original PL costs in Table I. We see in Table IIb an overall 17-percent savings in facilities and 20-percent savings in modems/DSUs. Therefore, if the purchase price of the customer premises TA concentrators is low enough (and depending on the customer's amortization period for equipment costs), then a VPL networking configuration that involves data-center and remote-site concentrators may offer significant economical savings across the customer's entire network.

Other Possibilities. Beyond simple private-line replacement, the TAs in a VPL network may implement various protocol-handling capabilities. For example, the TAs may allow the network to break up physical multipoint-access lines and, instead, logically bridge drops on a VPL. Or they may permit the network to mismatch speeds between a higher-speed access line to the FEP and low-speed tail circuits to the CCs. In the sections that follow, we present these potential TA capabilities one-by-one and discuss their effect on performance and costs for the VPL networking configuration shown in Figure 8.

Pipelining effects. *Pipelining* refers to a destination TA's ability to begin creating an outgoing SDLC or BSC frame before it receives all the LAPD frames for that SDLC (or BSC) frame.

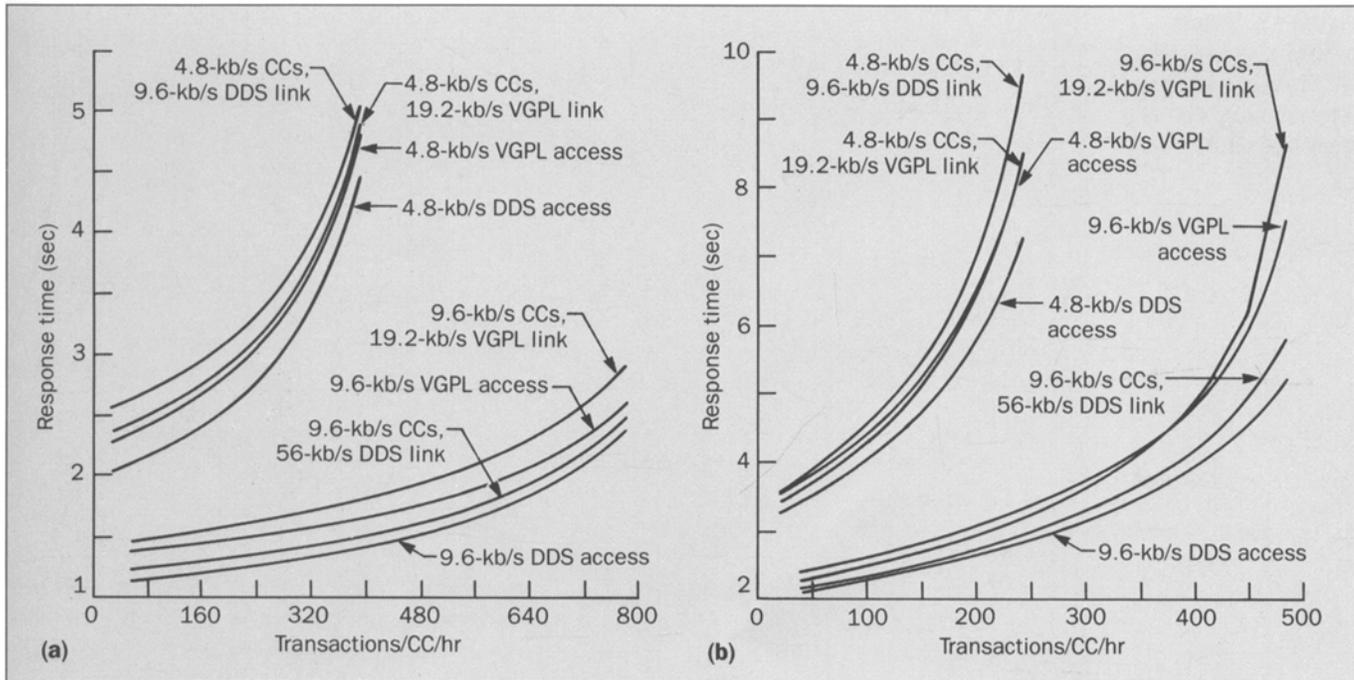
If the network's originating and destination access lines are of equal speeds, and if the delays do not vary for individual LAPD frames traversing the network, then the destination TA may begin transmitting an outgoing SDLC (BSC) frame as its first LAPD frame is being received. Each succeeding LAPD frame should arrive before the precise time required to maintain synchronous transmission of the SDLC (BSC) frame as the frame is transmitted out of the network to the host or CC. Without pipelining, the destination TA would have to accumulate the entire SDLC (BSC) frame, which is delayed by the time for the originating TA to receive the SDLC (BSC) frame.

In practice, however, the network introduces random delays that cause variations in delay; thus, succeeding LAPD frames may not arrive in time to maintain synchron-

ous transmission. "Gaps" in the transmission of an SDLC frame introduce frame errors that result in eventual frame retransmission. This probability of synchronization loss can be reduced to acceptably low levels if the beginning of transmission of each SDLC frame is delayed by an amount known as the build-out delay. The build-out delay increases the probability that succeeding LAPD frames will arrive before a preceding forwarding unit has finished transmission.

Figure 10a demonstrates the VPL performance advantage for the customer's SDLC circuits obtained by pipelining. If we compare the VGPL and DDS access curves here to those of Figure 6a, we see that pipelining reduced the VPL response times by about 33 percent for the multipoint circuits. (Figure 10a does not explicitly show the response times for the point-to-point SDLC circuits. But reductions in their response times from pipelining are of similar magnitude.) If we compare the remote-site-concentrator curves of Figure 10a to Figure 9a, we see that pipelining has reduced the VPL response times by a comparatively smaller amount (less than 10 percent) for SDLC transactions. In this comparison, the presence of slow-speed links between the customer premises TAs and the frame-relay switches counteracts the performance advantages of pipelining, further illustrating that the degree to which pipelining optimizes VPL performance is proportional to the ratio of networking link speeds to access line speeds. Overall, when we compare all the 9.6-kb/s access curves (with and without remote-site concentrators) of Figure 10a to Figure 3a, we see that the performance difference between the VPL and PL transaction times has now been reduced to about 20 percent for multipoint circuits.

Pipelining also offers significant advantages for BSC performance, as seen for the customer's BSC circuits in Figure 10b. VPL performance is now about 26 percent worse than PL performance for the 9.6-kb/s VGPL multipoint access circuits and 34 percent worse for the 9.6-kb/s DDS multipoint access circuits. Again, pipelining produces similar reductions for the point-to-point BSC response times. Moreover, with BSC transport, pipelining is handled



more simply than with SDLC transport. Build-out delays are not needed, and the destination TA may begin transmission onto the outbound line as soon as the first LAPD frame of the BSC block arrives. If the TA finished transmission of a LAPD frame before the next LAPD frame of the BSC block arrives, the TA can insert SYNC (synchronization) characters into the BSC block to avoid synchronization loss. (Note that the use of pipelining does not affect the customer's costs for facilities and modems/DSUs.)

Fanout and speed mismatching effects. *Fanout* or *logical multipointing* refers to the TA's ability to bridge several TA_T concentrator ports onto a single TA_H concentrator port. As an example, consider the single TA_T concentrator multipoint line that supports two CCs (TA_T in the upper right and bottom right corners of Figure 8).

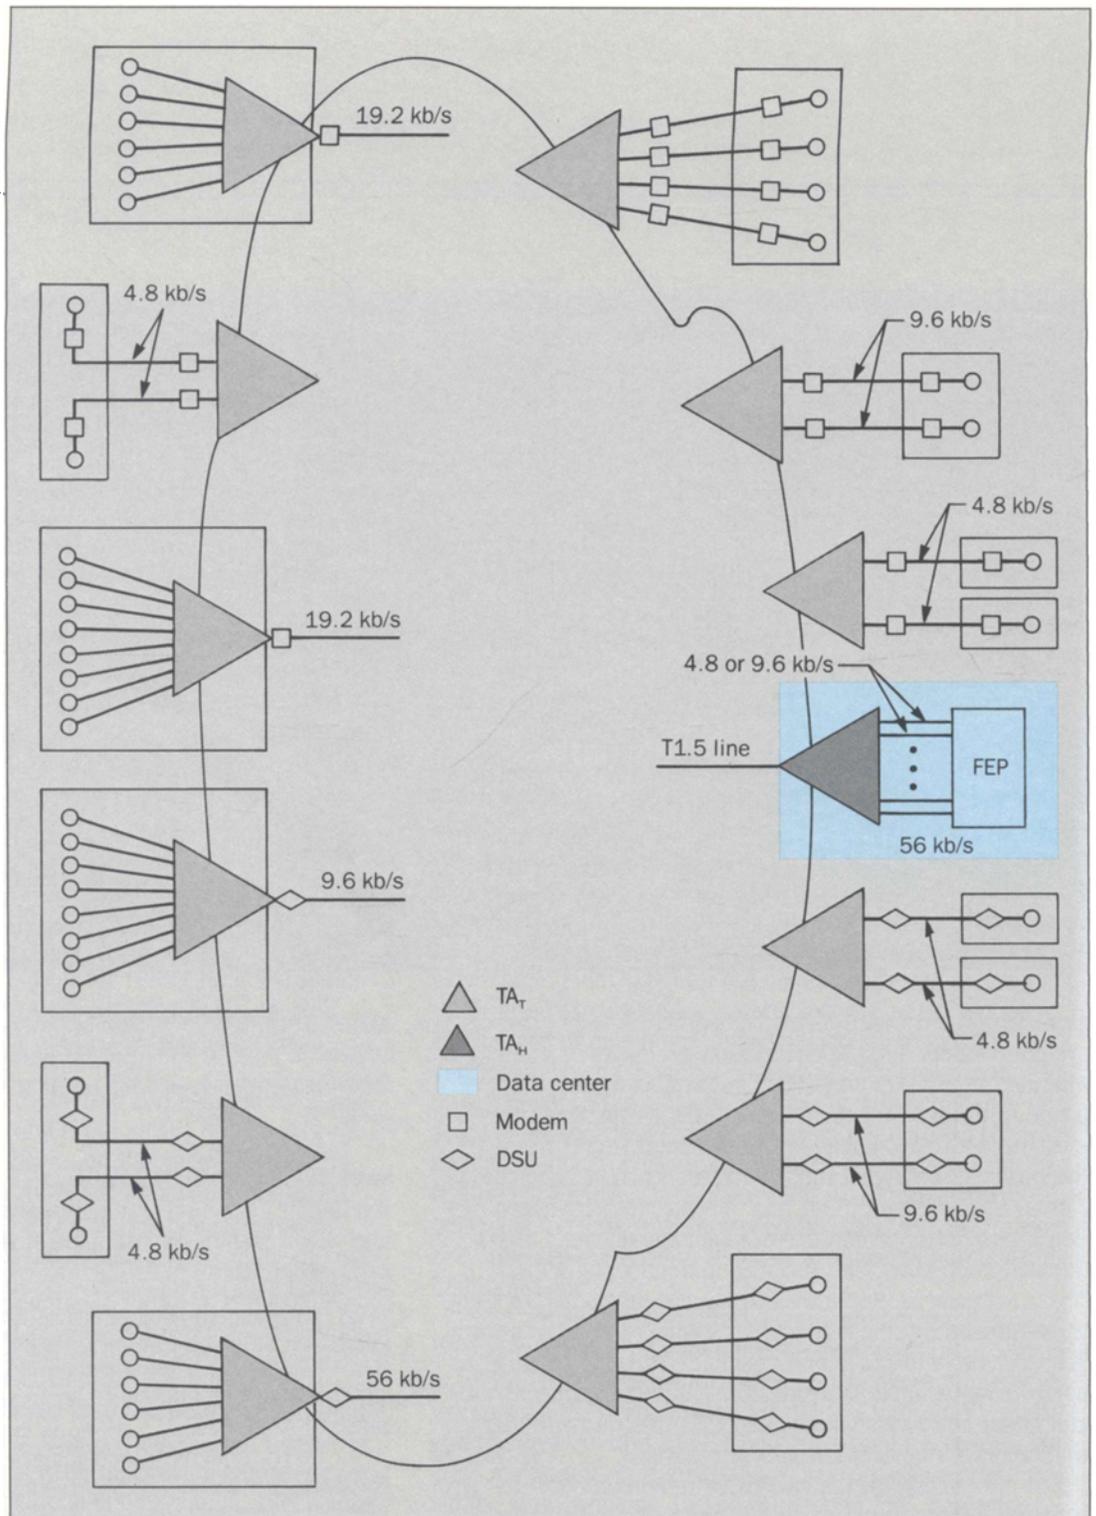
With fanout, the two CCs are spread out over two separate TA_T concentrator ports; each CC connects over a point-to-point line into the TA. The FEP's view of the net-

Figure 10. Multipoint virtual private-line performance for (a) SDLC and (b) BSC with pipelining. CC is a cluster controller, DDS is digital data service, and VGPL is a voice-grade private line.

work is unchanged; it still "sees" two CCs on a single multipoint line. As a result, the queuing delays for polls and RESP frames experienced at the FEP remain the same when fanout is used. However, the queuing delays for INQ frames at each CC can be reduced because of less polling delay, originally caused by polls queuing behind RESP frames on leaving the TA_T s to the particular CC. In addition, with fanout, there are no modem turnaround times at the CCs.

To reduce the queuing delays at the FEP, we could use a high-speed, 56-kb/s access line to the FEP. Here, the VPL network would be doing speed mismatching (e.g., enabling two external end points to communicate

Figure 11. Virtual private-line network with fanout and speed mismatching. FEP is a front-end processor.



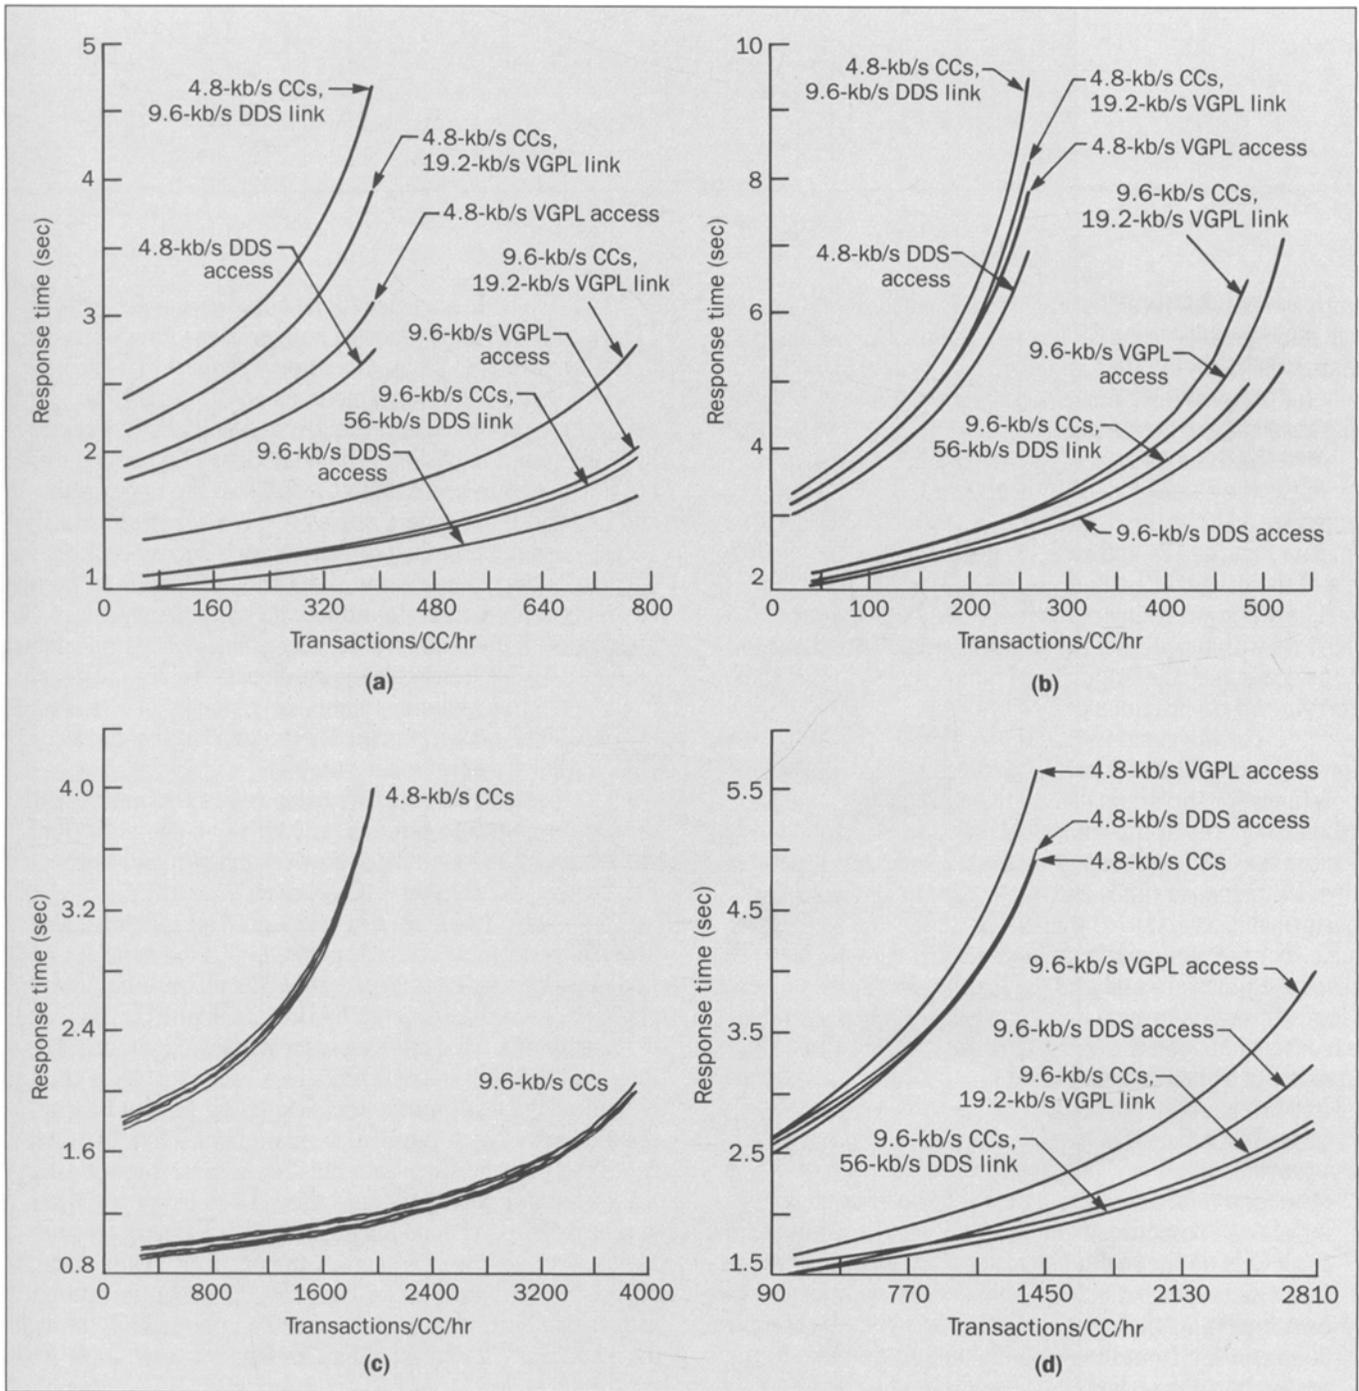


Figure 12. Virtual private-line performance for (a) SDLC and (b) BSC with pipelining, fanout, and speed mismatching; and for (c) SDLC and (d) BSC with pipelining, fanout, speed mis-

matching, and local polling. CC is a cluster controller, DDS is digital data service, and VGPL is a voice-grade private line.

with the network at differing speeds). Thus, together, fanout and a high-speed FEP line can reduce the polling cycle times for VPL networks. Figure 11 shows this architecturally for our customer model, and Figures 12a and 12b show the resulting performance.

With fanout, each VPL access line carries less traffic at the same level of transactions per CC. Thus, when we address the 5- to 75-percent utilization of access lines in Figures 12a and 12b, we imply far greater use levels of the original PLs in Figure 3. (We imply similar results when we compare a and b of Figures 12 and 10.) That is, with fanout at x transactions per CC per hour for one access line, the original multipoint access line would be carrying $5x$ transactions per hour.

For the customer's SDLC circuits, VPL response times (Figure 12a) now cross over relative to PL transaction times for the originally multipoint circuits (Figure 3) when compared at the same VPL and PL utilization levels. At low use, VPL response times are 13 percent greater than PL response times and, at the higher use levels, VPL performance is about 20 percent better than PL performance. For the customer's multipoint BSC circuits, the VPL response times are still greater than PL response times. However, the difference in response times at comparable use levels between the VPLs (Figure 10b) and PLs (Figures 3c) has been reduced to less than 20 percent across all use levels.

When fanout is implemented, the VPL network can provide logical bridging in two ways:

- *Broadcast multipointing*—The $TA_{H,S}$ replicate all commands and responses from the FEP, and transmit them to all CCs on the multipoint. Each CC, then, "hears" all these messages but acts on only those meant for it.
- *Selective multipointing*—The $TA_{H,S}$ "read" each command and response from the FEP, and only route them to the destination access line for the addressed CC.

Because selective multipointing can greatly reduce the traffic load on the network (compared to broadcast multipointing), it clearly is preferred from a performance viewpoint.

Table IIc also shows the costs associated with Figure 11's VPL networking configuration. As we can see in the diagram and table, fanout may require that the customer purchase additional modems and DSUs for the separate access lines into the $TA_{T,S}$. Also, facilities costs have increased because of the additional PLs needed to provide separate access into the $TA_{T,S}$ in the upper right and bottom right corners of Figure 11. Higher-speed FEP access does not increase the customer's access costs because "access lines" from the TA_H to the FEP are on the customer's premises and, hence, do not generate any facilities costs. If high-speed FEP access lines had been applied in Figure 5's VPL networking configuration, however, we would need to include the additional expense of one 56-kb/s DDS line per tail circuit for the FEP to the carrier POP 20 miles away.

Polling effects. *Local polling* refers to the ability of the $TA_{T,S}$ to emulate the FEP and do CC-polling functions and of the $TA_{H,S}$ to emulate the CCs and respond to the FEP polls.

Here, when the CC sends an INQ message, the TA_T still segments the INQ message into the smaller LAPD frames for transport across the network but intercepts EOTs and ACKs sent by the CC. (An EOT is an end of transmission, and an ACK is an acknowledgment.) At the other side of the network, the TA_H reconstructs the INQ message and queues it while waiting for the FEP to poll that CC. When the FEP sends the poll, the TA_H intercepts the poll and transmits the INQ message to the FEP.

Local polling reduces the FEP polling cycle time in that polling cycles no longer include the delays for polls to pass from TA_T to TA_H across the network. Thus, local polling permits a higher polling rate. However, mitigating factors can detract from performance optimization through local polling. With local polling, an INQ message must first wait to be polled at the CC and then, after transmission from TA_T to TA_H , the INQ message must wait again to be polled by the FEP. Without local polling, once the INQ message has been passed from TA_T to TA_H , the message is passed to the FEP without further waiting.

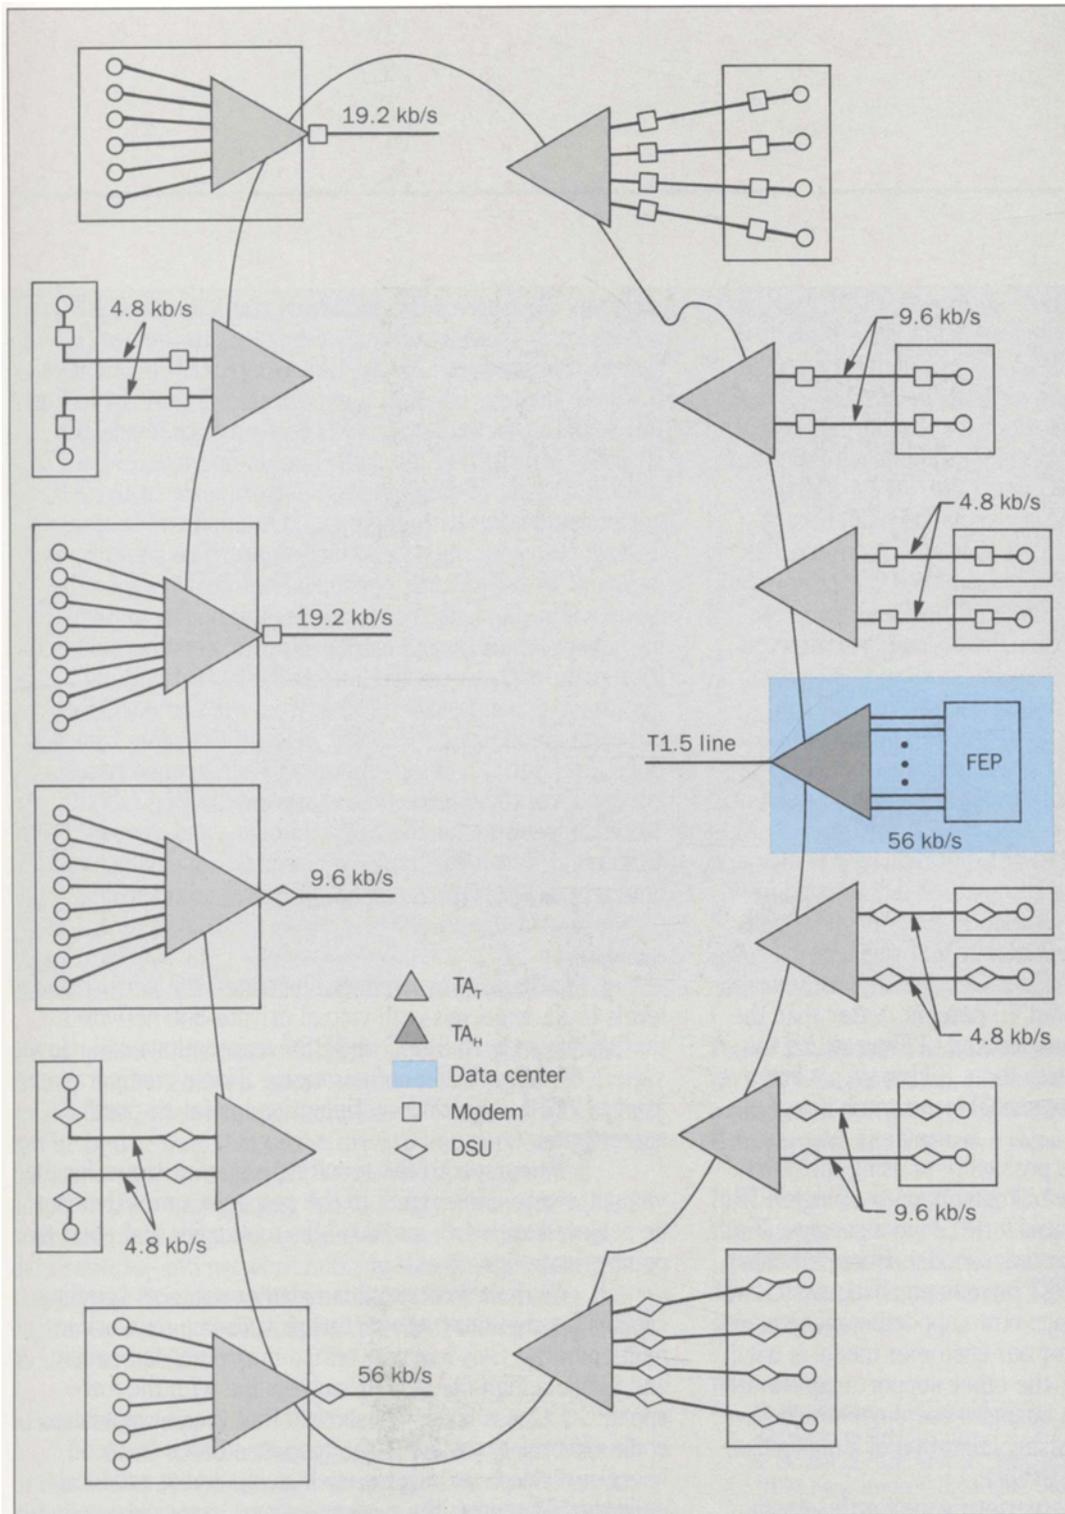


Figure 13. Virtual private-line network with fanout and a high-speed, multiplexed front-end processor (FEP) port.

In Figures 12c and 12d, we show the VPL response times for the networking configuration in Figure 11 with local polling assumed at the TAs. As is evident for both the customer's SDLC and BSC circuits, the benefits of polling cycle reduction have offset the double queueing delays introduced with local polling. For all configurations with the same CC speed (4.8 or 9.6 kb/s) in Figure 12c, we can see that local polling has also removed any significant differences between the response times from customer premises TAs connected by 9.6-, 19.2-, or 56-kb/s links to the frame-relay switches. For the SDLC circuits, VPL response times are now up to 10 percent better than the PL response times at the equivalent line-utilization levels. As mentioned above, with the VPL fanout configuration, each access line can carry more traffic. If we now compare VPL performance to PL performance at the same transactions rate per CC, the VPL SDLC multipoint response times are about 30 percent better than PL.

For the BSC circuits, VPL performance is now about 10 percent better than PL for the DDS multipoint access circuits and about 17 percent better for the VGPL multipoint access circuits at the same line-utilization levels. At the same transactions rate per CC, the VPL BSC multipoint response times are about 40 percent better than the PL response times. Again, as denoted in Figure 12d, the response times associated with the 4.8-kb/s CCs—*independent of the customer premises TA-to-switch link speed*—are close.

This introduces the possibility of using the VPL network to multiplex several tail circuits onto a single FEP port, thus offering the potential for FEP cost savings. Figure 13 shows this for our customer model. Here, the many 56-kb/s access lines at the FEP have been collapsed to two 56-kb/s FEP access lines, each now supporting many more CCs. One FEP access line for our customer model is used to support the 24 BSC CCs; the other supports the 24 SDLC CCs. (Any real, large-business customer would have to observe BSC's line-addressing limitation of a maximum of 32 CCs per multiplexed FEP line.)

Figure 14 shows the performance for this confi-

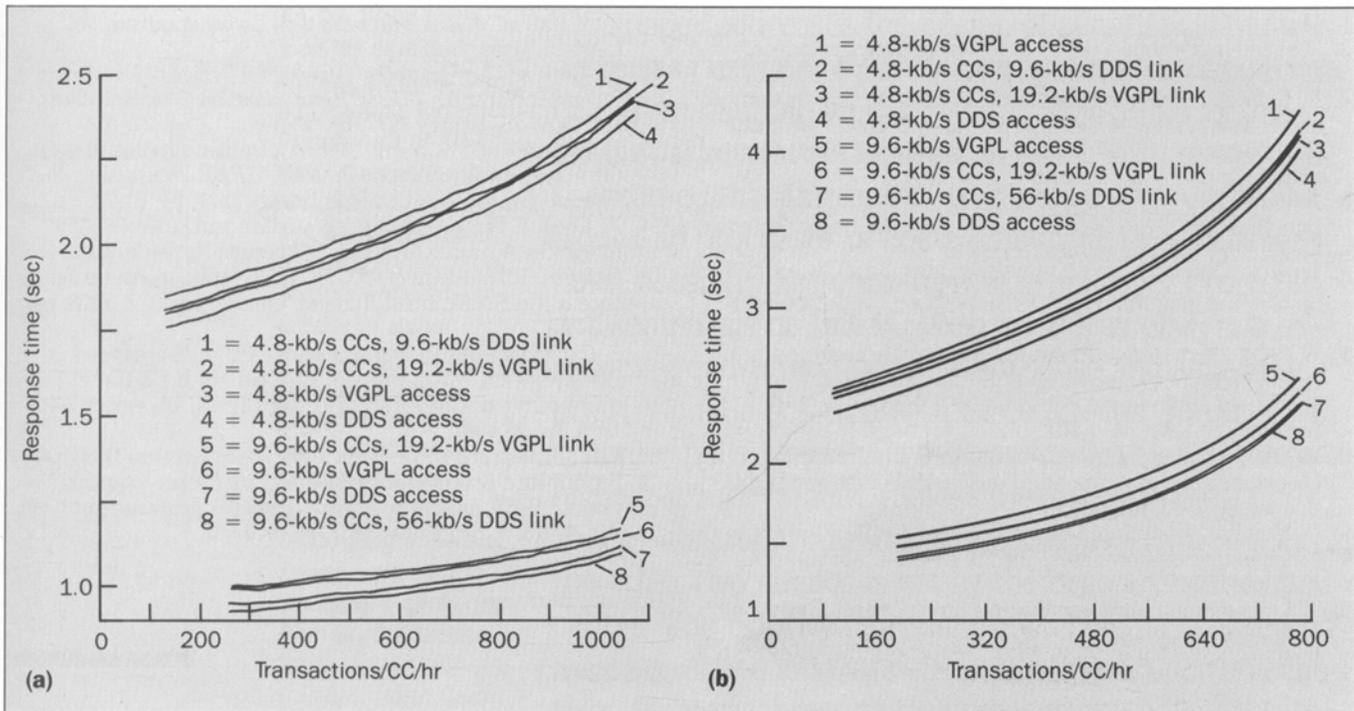
guration. Because of the additional traffic on the FEP access lines, the network may carry less traffic per CC than that in Figures 12c and 12d, where FEP ports were not consolidated. Overall, if we compare Figure 14 to Figure 3, the SDLC and BSC VPL performance levels still surpass that of PL at the same access-line-use levels. (To take advantage of the improved performance from FEP port consolidation onto high-speed multiplexed access lines, a customer must be willing to sacrifice some transparency. With FEP port consolidation, changes must be made within the FEP. For example, the polling and selecting service-order tables within the FEP must be modified.) Moreover, each access line into this VPL networking configuration can handle more traffic from the same CC than the comparable speed PL line. Also, at the same transactions level per CC, the multipoint VPL response times are significantly smaller than those for the PLs (about 30 percent better for the SDLC circuits, and over 40 percent better for the BSC circuits), especially as the number of transactions per CC per hour increases.

Summary

In this paper, we have illustrated the performance levels to be expected with virtual private-line networks that transport both SDLC and BSC transactions and use a variety of networking configurations. These configurations involve VGPL and DDS multipoint and point-to-point access lines.

We compared the resulting response times for the virtual private-line circuits to the response times that can be achieved with 4.8- and 9.6-kb/s multipoint and point-to-point private-line circuits.

We described terminal-adaptor protocol-handling capabilities that allow the virtual private-line network to multipoint the CCs logically onto a multipoint tail circuit and to use a high-speed FEP access line with the low-speed CC access lines. We showed that these capabilities could effectively minimize the response times that end users see. Also, techniques such as pipelining and local polling could reduce the response times associated with a



virtual private-line network. If these possibilities are judiciously combined (dependent on the SDLC or BSC application), the virtual private-line network's performance can be better than that of the private-line network carrying the same amount of traffic. Moreover, if the terminal adapters have multiplexing functionality, placing small and large terminal adapters at the customer's remote sites, data centers, and network hubbing sites can produce virtual private-line networking configurations that offer the customer overall savings in private-line facilities and modem and DSU costs.

Acknowledgments

Many people throughout AT&T have contributed to this work, either through discussions or modeling computations. In particular, we offer special acknowledgments

Figure 14. Virtual private-line performance for (a) SDLC and (b) BSC with pipelining, fanout, local polling, and high-speed multiplexed FEP port. CC is a cluster controller, DDS is digital data service, and VGPL is a voice-grade private line.

to R. Cole, Y.-F. Lam, and R. Moats for helping to prepare the figures used here. M. Chiosi, S. Chitariance, B. Doshi, T. Kerrigan, S. Kashdan, A. Kumar, A. Kuthyar, J. Medamana, H. Nguyen, and J. Whitehead also merit acknowledgments.

References

1. K. A. Boakye, J. C. Kaufeld, Jr., and J. W. Palmer, "AT&T Data Networking Architecture," *AT&T Technical Journal*, Vol. 67, No. 6, November/December 1988, pp. 23-34.
2. "IBM Synchronous Data Link Control General Information," IBM Doc. No. GA27-3093-2, Third Edition, IBM Corporation,

- March 1979.
3. "General Information—Binary Synchronous Communications," *IBM Systems Reference Library*, IBM Corporation, October 1970.
 4. R. G. Cole, "Performance Models of Binary Synchronous Communications Multipoint Circuits over Virtual Private-Line Frame Relay Networks," *AT&T Technical Journal*, Vol. 67, No. 5, September/October 1988, pp. 41-56.
 5. A. Kumar, "Performance of SNA/SDLC Over Virtual Circuits in a Data Network," *AT&T Technical Journal*, Vol. 67, No. 5, September/October 1988, pp. 27-40.
 6. J. P. Buzen, "An Overview of Performance Prediction in MVS Systems and SNA Networks," *IEEE Fall Joint Computer Conference*, Dallas, Texas, November 2-6, 1986, pp. 751-759.
 7. W. Chou, "Performance Evaluation of Polled Multipoint Teleprocessing Networks," *Proceedings of the Computer Networking Symposium*, Institute of Electrical and Electronics Engineers, December 8, 1981, pp. 1-16.
 8. W. Chou, "Analysis of Data/Computer Networks," *Computer Communications*, W. Chou, ed., Prentice-Hall, Englewood Cliffs, New Jersey, 1983, Chapter 10.
 9. O. Hashida, "Analysis of Multiqueue," *Review of the Electrical Communications Laboratories*, Vol. 20, Nos. 3-4, March-April 1972, pp. 189-199.
 10. T. S. Tanenbaum, *Computer Networks*, Prentice-Hall, Englewood Cliffs, New Jersey, 1981.
 11. R. R. Boorstyn and H. Frank, "Large-Scale Network Topological Optimization," *IEEE Transactions on Communications*, Vol. COM-25, January 1977, pp. 29-47.
 12. M. Gerla and L. Kleinrock, "Topological Design of Distributed Computer Networks," *IEEE Transactions on Communications*, Vol. COM-25, January 1977, pp. 48-60.
 13. P. McGregor and D. Shen, "Network Design: An Algorithm for the Access Facility Location Problem," *IEEE Transactions on Communications*, Vol. COM-25, January 1977, pp. 61-73.
 14. W. Chou, F. Ferrante, M. Balagangadhar, and L. Gerke, "An Integrated Approach to Optimally Locating Network Access Facilities," *Proceedings of ICC '78*, IEEE International Conference on Communications, Toronto, Canada, June 4-7, 1978, pp. 335-345.
 15. S. Chitariance, "A Packet-Switched Network Analysis and Design System," *Proceedings of Globecom '86*, IEEE Global Telecommunications Conference, Houston, Texas, December 1-4, 1986, pp. 983-987.
 16. C. L. Monma and D. D. Sheng, "Backbone Network Design and Performance Analysis: A Methodology for Packet Switching Networks," *IEEE Journal on Selected Areas in Communications*, Vol. SAC-4, No. 6, September 1986, pp. 946-965.

(Manuscript received July 29, 1988)