## CASES IN NETWORK DESIGN

Edited by William E. Bracker, Jr. and Ray Sarch



## CASES IN NETWORK DESIGN

Edited by William E. Bracker, Jr., University of Arizona and Ray Sarch, Executive Technical Editor, Data Communications





McGraw-Hill Information Systems Company 1221 Avenue of the Americas New York, New York 10020



- Basic Guide to Data Communications. Edited by Ray Sarch. 1985, 360 pages, softcover.
- Cases in Network Design. Edited by William E. Bracker, Jr. and Ray Sarc. 1985, 280 pages, softcover.
- Computer Message Systems. By Jacques Vallee. 1984, 176 pages, clothbound.
- **Data Communications: A Comprehensive Approach.** By Gilbert Held and Ray Sarch. 1983, 441 pages, clothbound.
- Data Network Design Strategies. Edited by Ray Sarch, Executive Technical Editor, Data Communications. 1983, 273 pages, softcover.
- **Interface Proceedings.** Edited by Data Communications Magazine. Annual publication, softcover.
- The Local Network Handbook. Edited by George R. Davis, former Editor-in-Chief, Data Communications. 1982, 260 pages, softcover.
- Local Networking: Analysis, Implementation and Global Interconnection. By Michael L. Rothberg. 1985, 350 pages, clothbound.
- McGraw-Hill's Compilation of Data Communications Standards, Edition II. Edited by Harold C. Folts. 1982, 1,923 pages, clothbound.
- Teleconferencing and Beyond: Communications in the Office of the Future. By Robert Johansen, with others contributing. 1984, 206 pages, clothbound.
- Teletext and Videotex in the United States: Market Potential, Technology, and Public Policy Issues. By J. Tydeman, H. Lipinski, R. Adler, M. Nyhan, L. Zwimpfer. 1982, 312 pages, clothbound.

Copyright © 1985, McGraw-Hill, Inc. All rights reserved. Printed in the United States of America. Except as permitted under the United States Copyright Act of 1976, no part of this publication may be reproduced or distributed in any form or by any means, or stored in a data base or retrieval system, without the prior written permission of the publisher.

### Library of Congress Cataloging in Publication Data Main entry under title:

Cases in network design.

(Data communications book series) Includes index.

1. Data transmission systems—Case studies.
2. Computer networks—Case studies.
1. Bracker, William E., 1947—
11. Sarch, Ray. III. Series: McGraw-Hill Data Communications book series.
TK5105.C38 1985 384 85-10500
ISBN 0-07-606926-5

## CASES IN NETWORK DESIGN

Edited by William E. Bracker, Jr., University of Arizona and Ray Sarch, Executive Technical Editor, Data Communications



E8762347

## Preface

This book is divided into nine sections, each devoted to specific areas of data network cases. Most readers who are potential network implementers, network improvers, or those seeking examples of actual implementations to supplement academic data communications material should find the contents beneficial. All the cases originally appeared in DATA COMMUNICATIONS magazine as feature articles and should be even more valuable when presented in this format.

The case studies herein are by no means an exhaustive list. However, they do provide a representative sampling of network design cases. We have provided articles that describe major network implementations involving real-world applications.

We have also attempted to include as diverse a sample of network designs as possible. The reader should beware that, because of the dynamic and evolving nature of data communications, a network can be obsolete even as it is described in the literature. However, the process leading up to the design, together with the description of the network created, can often be directly applicable to network designs currently in progress.

The editors recommend that "Data Network Design Strategies" be used as a companion volume. Thereby, the technology and techniques in the network design process can be fully appreciated.



## Table of Contents

### VII Preface

### 1 SECTION 1 **BANKING AND BROKERAGE**

- Canadians exchange stocks in a mixed trading network, Donald R. Unruh (February 1981)
- Brokers invest in bisync to secure a reliable network, William A. McLoughlin (April 1981)
- London Stock Exchange bullish about videotex, Peter Bennett (August 1981)
- Bank data networks: Moving millions electronically, Edwin E. Mier (April 1981)
- 32 Bank validates transactions with home-grown data authentication, Esti Weissman, Louis H. Phillips, and Joe Wallace (July 1982)
- Reliability: The bottom line of Chase Manhattan's ATM network, Charles Tritto (November 1982)
- American Stock Exchange trades old network for streamlined setup, John Ambler (November 1983)

### 47 SECTION 2 **MILITARY APPLICATIONS**

The MX network: Designed to work when all else fails, Joseph St. Armand (September 1981)

### 61 SECTION 3 **AIRLINES**

- Concentration is the key to TWA's new network, Edward MacDonald (January 1981)
- Interconnect net navigates airline reservations, A Special Report (September 1982)

### 75 SECTION 4 LOCAL NETWORKS

- Local net user relates trials, tribulations, Ronald W. Gibson (March 1983)
- Considerations in interconnecting diverse local nets Richard H. Sherman and Melvin G. Gable (June 1983)

### 89 SECTION 5 **PACKET SWITCHING**

- An award-winning packet-switching selection program Stan Badowski and David Hanson (April 1983)
- Anatomy of a packet-switching overhaul, Joseph F. Haughney (June 1982)
- Packet network proven best for optimizing a library operation, William F. Adiletta (February 1984)
- Lessons learned in building an international packet network, Dan Muecke (May 1984)

## 119 SECTION 6 PBX/VOICE/ELECTRONIC MAIL/IMAGE

- Complex revamp of network steers user to a digital PBX solution, Michael R. Ham (July 1983)
- Modeling an electronic mail network: A primer, Robert Graff (August 1983)

- 133 Canada combines voice, data, and images on national satellite network, Pierre M. Hubert and Jacques Gaumond (January 1984)
- 137 Private network integrates data, voice, and video communications, W.G. Phile, D.L. Green, and L.J. Cole (May 1984)

### 145 SECTION 7

### LIGHTWAVE/FIBER OPTICS

- Lightwave is firm's link to greater productivity, Don Jacobs and Paul Dobson (June 1981)
- Network annals: From paper tape to lightwaves, Gerald H. Ely (September 1980)
- Breaking urban data bottlenecks with a private lightwave bypass, Don Jacobs (April 1984)

### 165 SECTION 8

### **DISTRIBUTED DBMS**


- Distributed DBMS no longer just a concept, James B. Rothnie, Jr. (January 1980)
- Piggybacking word processing on a DDP network, Mort Miller and Joyce Drumm (May 1982)
- Tightening up software on a distributed database, George B. Houston (April 1983)

### 187 SECTION 9

### **NEW NETWORKS AND UPGRADES**

- Network combines two approaches for peak performance, John Myers and Stuart Walker (January 1980)
- Forest-products firm dedicates minis to many locations. Thomas B. Walsh (May 1980)
- The month-to-month implementation of a private network, Henry J. Miller (June 1980)
- How and why one company upgraded its network, Ken Conley (October 1980)
- Necessity is mother of homegrown DDP network, Scott G. Abbey, David N. Bertollo, and Joseph R. Geller (February
- 219 ug enforcement officials nab a simple solution to network woes, Roy Bell and Dexter Thomas (March 1982)
- 安全的 Property of the Property Migration to SNA: Sometimes it's the only game in town,
  - College network passes test, despite limited computer resources, Mike Higgins (April 1983)
  - Homemade micro software for synchronous communications, Eduardo H. Martinez (September 1983)
  - Data communications network evolution has lessons for all, J.T. Sinclair, W.E. Lewis, and R.B. Wall (September
  - Reporting tools help new network to meet schedule as well as budget, Robert O'Brien, Hiram Ortiz, and Vince Coughlin (October 1983)
  - Building a network around an IBM 4300 mainframe, David P. Levin (December 1983)
  - How networking helped one industry cut paper mountains to bits, Helmut J. Tissler (March 1984)
  - Combining the best of SNA and X.25 architectures, Steven Holmes and Miles Fleming (June 1984)

### 273 INDEX



# Canadians exchange stocks in a mixed trading network

Donald R. Unruh, Toronto Stock Exchange, Ontario, Canada

When trading on the Toronto Stock Exchange reached record highs, so did communications costs. But Canadian ingenuity bred a hybrid network to balance costs and efficiency and expand services.

The Toronto Stock Exchange's (TSE) market information services are updated with data on more than 150,000 trading transactions a day, from all over North America. In addition, Candat (which refers collectively to the network and databases) responds daily to over 350,000 inquiries entered at more than 1,500 terminals across Canada and the United States. Many of the same access terminals are used in the Computer-Assisted Trading System (CATS), which handles between 15,000 and 18,000 trading transactions a week—approximately 25 to 30 percent of the Toronto Stock Exchange's total volume.

The number of Candat terminals has increased 50 percent annually over the last several years. This has been possible through the evolution of a hybrid computer network connecting both the many databases of trading data and the widely distributed users of the information services. The network is partially private, using local analog circuits and the TransCanada Telephone System's Dataroute digital communication service, while the public portion is based on the Datapac packet-switching service. The private portion of the network supports TSE's market information service, Autodat, and the CATS trading system. The public portion handles the exchange's two new information-retrieval services, Instadat and Quicknet.

The introduction of public packet-switching communications has been significant to the stock exchange in two ways. First, TSE has been able to offer additional market information services, while limiting its involve-

ment in the communications business. Secondly, the public network lets TSE offer these services competitively and profitably without geographical constraints. Consequently, TSE expects that in the near future all of its market information services, including those currently offered through the private network, will be accessed through packet-switching services.

Autodat, which has been in operation for about five years, is the predominant user of the Candat information network. It accounts for about 80 percent of the total inquiries, encompasses more than 1,350 terminals, and is used to some degree by virtually all of TSE's 79 member firms. Moreover, the Autodat computers, terminals, and communications facilities also handle the CATS trading system.

A minimum configuration for an Autodat subscriber is a CRT display and special keyboard with 165 keys for both Autodat and CATS services. A variety of trading information may be retrieved and displayed on the divided CRT screen. In a typical assortment of display data (Fig. 1), the upper screen area might contain either a conventional horizontal stock ticker or Autodat's unique "waterfall" vertical ticker. The central screen area would likely be devoted to a "stock-watch" display, which includes as many as 24 stocks or market indices requested by the user. The lower screen area is a work space for interactive inquiry at the keyboard, although all screen areas can be used for this purpose.

By entering the stock symbol and exchange code, a trader can request quotes on as many as six stocks,

18.0/ 1		T	ROC L.P	P C							
		İ	MCL		500	46.7	/+ 62	-07	4,175	2/	
29.5/ 2	29.7/	Ţ	STE VO	. А	100	30.2	/+ 15	-63	2,555	+.2/	
		T	KSR		200	11.6	/+ 22	-64	7,200	+.17	
T N.A			32.3/	32.2/+		5 T II		24.0/		24.1/-	28,57
T MF				28.2/-					30,3/	70 1 -	- :
M MF	28	3.2/	28.3/	28.3/	12,74			30.0/	30.2/	30.2/	30
V SIR		.05	3.10	3.05	93	a v vo	)LU			11 190	000
T \$I		7	695686	197.06	361,96	O T SE	}M		278936	86.68	27,10
T \$GL			215907	307.28	35,96	1 T \$1	10		779398	226.22	53,10
T MF			28.2/			.2/-	26,184		28.1/		
M MF			28.2/	28.3/		.3/	12,740		28.0/	28.3/	28.0/
V MF						.4/		NC			
N MF	2	3		23.2/ 0							
T B			47.0/	47.17	47	.0/-	11,128	+.17	47.0/	47.1/	47.0
										4.4	. 79.70
											:32:30.

1. Partitioned screen. Subscribers to this Toronto Stock Exchange service receive different outputs on a single, partitioned CRT screen. The top portion is a "waterfall"

display of broadcast stock data, while the center and bottom portions are stationary displays that are updated periodically through individual transactions.

options, bonds, or commodities in each screen area. The tickers in the upper screen section are constantly moving, and the market information in the two lower areas is stationary, but automatically updated after every quote change or trade.

In addition to offering trading data similar to that of other market information services. Autodat provides several unique types of market information. A "previous-trades" display, for example, lists all trades in a given stock in successive pages, starting with the 10 most recent trades. A "block-trades" display (based on a dollar value of \$50,000 or greater) shows the 10 most recent trades, along with the buyer and seller broker numbers. Information for the stock option market, which has been growing rapidly, includes the last price of a stock, its bid and asking price, and the last prices of up to 15 options simultaneously. A "multipage-options" display expands on the options information by showing the volume traded, the net change from previous close, the open, high, and low prices, and the total open interest.

The Autodat responses to the 250,000 daily terminal inquiries to the central Candat database are generated at the TSE's network center (Fig. 2). The database includes stocks and options traded on the Toronto, Montreal, Vancouver, and Alberta exchanges and, through the Securities Industries Automation Corporation (SIAC), the New York Stock Exchange (NYSE) and the American Stock Exchange (ASE). There is also a link to the Quotron financial information service in New York City, which covers all U.S. stock exchanges other than the NYSE and ASE. Two-thirds of the inquiries are for TSE-traded stocks and options, about 20 percent for U.S. stocks, and the rest for

shares on other Canadian exchanges.

Instadat is a low-cost market information service. It provides access to the Candat database for inquiry, but does not include any tickers or automatic updates of displays, such as Autodat's "stock-watch" feature. Instadat, which began in mid-1979, has already reached a daily volume of up to 80,000 inquiries through 180 terminals.

Quicknet offers the same market information as Instadat, but with a message-processing facility for multiple-office brokerage firms. It currently serves three national brokerage firms with over 100 inquiry terminals at 75 branch offices throughout Canada. Messages can be transmitted between two terminals or broadcast to all terminals. Most terminal-to-terminal messages are trading orders, while broadcast messages usually contain research and accounting information.


The CATS trading system originated in November 1977 as a pilot project and has since become the official trading system for inactive stocks. It was the first and still is the most advanced automated trading system, and its stocks account for about 700 of TSE's 1,150 securities. The CATS system was created to eliminate inactive posts from the trading floor to accommodate increased options trading. The changeover was planned so that no dealer would have to use both manual and automated trading procedures.

### Trading network's structure

These market information services are provided through the Candat network as shown in Figure 2. The basic network elements are the host facility (at the TSE network center in Toronto), lines to sources of market data, lines to users of market data, and the hardware

2. Overall network. The TSE network includes three distinct subnetworks, each providing substantially different information services. In addition, a number of periph-

eral communications links to other stock exchanges and services, local update terminals, and remote newspaper typesetters are also supported on-line.



and communications facilities associated with each of the three Candat services.

The host facility consists of three mainframes and three front-end processors. In addition, several smaller computers serve as cluster controllers, communications processors, concentrators, and message processors within and outside the network center. Although the message traffic is very high and constantly flows in all directions, the system software is designed so that each computer handles only its specific functions.

The Candat mainframes are an IBM 3031, an IBM 4331, and a Control Data Corporation Omega 480-3. The 3031 maintains the Candat database and runs all application programs. It also operates the 4331 via a channel-to-channel adapter. The Omega 480-3 is used

for software development and production work for various Candat administrative reporting applications. A pair of Digital Equipment Corporation (DEC) PDP-11/60 minicomputers and a PDP-11/40 operate as front-end processors, primarily handling polling and data traffic between the mainframes and the computers and terminals in the network. The mainframes, front-end processors, tape drives, and the three sets of five Control Data 33301 disk storage units for the Candat database are linked through a T-bar Inc. 3915/3916 computer switch system.

Because of the periodically intense network activity, the terminals linked to the network demand high performance levels from all computers, particularly from those in the host facility and the DEC LSI-11 microcom-

puters that function as cluster controllers. In the peak trading hour between 10:00 and 11:00 A.M., for example, 40 messages per second flow from the front-end processors to the 3031 mainframe. As network activity grows, additional PDP-11/60s will be installed to handle the increased network-management load.

The application programs on the 3031 are run under the control of a TSE-written multitask monitor. In the 11/60 front end, these multitask monitors operate in conjunction with DEC's 2848 emulator software. To the DX-11 interface, the mainframe operates as if it were supporting relatively few IBM 2260 user terminals located at the network center. Logically, it is not aware of the communications environment beyond the DH-11 interfaces. On the output side of the front ends, polling is done through half-duplex asynchronous transmission, using a TSE-written subset of IBM's binary synchronous communications (BSC) protocol.

### **Database input and update**

The most active portions of the Candat database are the stocks and options traded on the Toronto Stock Exchange (80 percent of the value of all Canadian stock trading is done on the TSE). Dual PDP-11/20 minicomputers concentrate data for TSE stock transactions that are entered at 12 on-line keyboards (located below the main trading floor) and for transactions from other Canadian stock exchanges received from a 5-bit Baudot ticker. Data on options, quotes, and trades are entered directly into the front-end processors via local terminals.

A PDP-11/34 concentrator for the NYSE and ASE tickers is located at the SIAC offices in New York City. This enhances the quality of the transmission to Toronto and reduces the cost (the minicomputer's price was significantly lower in New York than in Toronto). The concentrator receives "last-sale" and "bid/ask" tickers from SIAC on 4.8-kbit/s receive-only lines and buffers them onto a disk in case the network center unexpectedly shuts down. The two inputs are then merged in alternating 64-character blocks and sent to Toronto via 9.6-kbit/s asynchronous transmission.

Users also have direct access to the Quotron service for U.S. stocks and options not maintained in the Candat database (most of these are Chicago-based options and commodities). In addition, the TSE ticker is reciprocally transmitted to Quotron for distribution to U.S. subscribers and to other market inquiry services in Canada, as well as to some Canadian and U.S. newspapers for remotely generating typesetting tapes.

### The private subnetwork


The Autodat subnetwork (Fig. 3) is based on a battery of cluster controllers, each managing data communications for as many as 20 inquiry terminals. The TSE currently has 78 active cluster controllers—45 at the network center, and the rest at customer sites.

The ports on each controller are distributed as much as possible among Autodat subscribers, so that no subscriber's terminals are linked to only one controller. This way, if a controller goes down and its terminals are temporarily inoperative, at least one other terminal

on another controller will maintain service. All but a few controllers are DEC LSI-11 microcomputers. The others (in the network center) are DEC PDP-11/05 units. Each microcomputer controls the Autodat-user interactions for the terminals in its cluster and maintains a local list of the stock-watch selections for all its terminals, automatically updating each item. Probably the controllers' most difficult task is handling the almost constant stream of data feeding the waterfall ticker. The greatest load for this occurs during the first hour of the 6-hour day, when approximately a quarter of the total terminal activity takes place.

The controllers at the network center are hardwired to the front-end processors and linked through standard phone lines to Autodat subscribers in metropolitan Toronto. Other controllers are located in major subscriber offices throughout Canada. These microcomputers may be hardwired to terminals at the same site or connected through local or long-distance leased lines to others. The remote cluster controllers are themselves linked to the network center through Dataroute

**3. Autodat subnetwork.** LSI-11 concentrators, at both the network center and user sites, handle Autodat traffic, which includes interactive trading transactions via CATS.




digital leased lines, which are substantially less-expensive than standard voice-grade lines. About half of the 1,350 Autodat terminals are linked to the network center over these facilities.

Terminal communications within Autodat consists of either a broadcast line (for ticker transmission and automatic updating) or a line for database inquiry and response. The broadcast lines operate at 2.4 kbit/s, while most inquiry lines transmit at either 2.4 or 4.8 kbit/s. About 80 percent of the total Autodat message traffic is the broadcast of trading data.

CATS trading (also done on the Autodat subnetwork) begins with an inquiry on a particular stock. The Autodat terminal then displays a ranked list of the best bids and offers for that stock, along with the number of shares for each. The trader inputs his security number, the price and volume of the bid or offer, and whether it is a client or nonclient trade. The data, entered in response to prompting questions, is then edited by CATS and returned to the terminal display in a fixed format, together with the latest quote, for the

**4. Instadat subnetwork.** Datapac's packet-switching network serves as the backbone for this service. Updates are made on a transaction, instead of broadcast, basis.



trader's acknowledgment. When the data has been verified, the trader presses the enter key to submit the order. A response informs him immediately that either the trade has taken place (that is, his bid or order has been matched) or his order has been placed in a trading queue. Later, successful trades are typed out on an adjacent printer, and a note is displayed in a memo area at the bottom of the terminal screen.

The charge for Autodat is a single fee per month per terminal, which varies depending on the type of terminal equipment (14- or 23-inch display screen with customized keyboard and optional printer). Subscribers also pay a monthly access fee for each exchange. There is no additional charge for the CATS service, since TSE members are already billed a standard transaction fee for exchange trades.

## **Economics of packet-switching services**

Packet-switching communications was first considered as a way to eliminate the high communications cost in an earlier version of Candat. Until recently, the inquiry service used 75-bit/s telegraph lines to serve subscribers in remote areas not covered by the newly available Dataroute network. However, because of the high circuit charges, TSE's cost of service to these remote subscribers had become intolerably high.

TSE therefore did a study to compare the quality and cost of the existing telegraph network with the Datapac packet-switching network. It found that for about the same monthly cost of the telegraph facilities, packet-switching facilities could provide substantially more data throughput (three times the transmission speed). Packet-switching was discovered also to cost subscribers less. Dataroute (digital) data transmission between Toronto and Vancouver, for example, costs about 5 cents for the amount of data contained in one closely typed page. With Datapac, excluding any associated hardware, the cost of transmitting the same information between Toronto and Vancouver is less than 1 cent (based on a charge of 92 cents per thousand 256-character packets between the two cities).

In addition to solving the immediate cost problems for remote service, packet-switching communications gave TSE the means to offer broader Candat services without having to manage a private leased-line network. Since TransCanada could readily count the number of packets associated with a given subscriber, communications cost could be taken out of the Candat monthly fee, and the subscriber billed directly by TransCanada for Datapac usage.


The Instadat subnetwork (Fig. 4), the first Candat service based on packet-switching technology, is supported by a PDP-11/34, two PDP-11/05s, and a PDP-11/05 backup machine located at the network center. The PDP-11/34 communications processor provides service to more than 100 terminals of various types and models. All have been selected by subscribers from a list of Candat-approved terminals, and all operate at 300 or 1.2K bit/s, determined by the user.

The communications software running on the PDP-11/34 was developed jointly by TSE and TIL Systems Ltd. of Toronto. It is compatible with the CCITT X.25 communications standard and supports TransCanada's X3000 standard synchronous protocol. All Instadat terminals operate with the X3101 asynchronous protocol when connected through TransCanada's network interface module (NIM).

The tickers and automatic-update features are not yet available on Instadat terminals, because, unlike Autodat, Instadat has no broadcast capability. The user must therefore enter a new inquiry to access the database for every information request. The reasons for this are entirely economic: packet transmission of the almost continuous stream of broadcast data would result in extremely high network charges. Moreover, because packet-switching is a point-to-point process, the front ends must transmit separately to each terminal, which requires much more front-end processing power than for individual terminal-inquiry transactions. In contrast, a cluster controller in the Autodat subnetwork transmits once, simultaneously, to all multidropped terminals on a given leased line.

TSE recognizes the desirability of adding automatic

**5. Quicknet subnetwork.** PDP-11/34 minicomputers handle the message-switching operations of the Quicknet subnetwork. Subscribers consist of multibranch firms.



updating and of eventually providing Autodat-level services to all locations. A more economical automatic-updating method would require terminal users to specify the particular securities they would like updated. This would require far fewer packets, especially if update messages were consolidated in packets.

There are several ways to do this. As with the Autodat cluster controllers, software could be added to the communications processor so that it maintains tables of securities for each Instadat terminal. The tables would then be set up remotely and changed by the user, as they are for the Autodat stock watch. Alternately, the user hardware could be upgraded to intelligent terminals that could automatically request updates. Intelligent terminals might also permit a more economical leased-line network for broadcast data, reserving the packet-switching network for inquiry.

### **Quicknet subnetwork**

The packet-switching software that links Quicknet with the Datapac network was also developed by TIL Systems, which provides its TIL Model 258 video unit as the standard terminal. The user can incorporate any of several compatible hard-copy printers with the video terminal. Each subscribing firm has its own PDP-11/34 message processor, which can handle more than 100 terminals (Fig. 5). The message processor can be located at the subscribing firm's main offices, at the network center, or at any other advantageous location. Depending on their distance from the network center, message processors can communicate with the frontend machines either through hardwired connections, local leased lines, or the Dataroute network.

Quicknet provides multi-office subscribers with a message-switching facility that is faster and less-expensive than competitive Canadian wire message services. Messages print out on a Quicknet terminal at 30 characters or more per second, while other services operate at rates of 10 characters per second. Quicknet messages are switched by the message processor using a store-and-forward procedure.

Instadat users are charged a monthly fee for access to Candat (which varies depending on whether or not they are TSE members) and an inquiry fee (currently \$5 per 1,000 inquiries). In addition, they lease the terminals directly from a terminal vendor and pay exchange access fees and Datapac charges. The total monthly cost to a TSE member (excluding data fees payable to other exchanges) typically ranges from \$215 to \$290 per month per terminal. Similarly, Quicknet customers pay a terminal leasing fee to TIL Systems, exchange access fees, and Datapac charges.

A significant aspect of communications at the TSE is that the Candat services, because of packet-switching, have been given international scope. As mentioned earlier, the Toronto Stock Exchange plans to convert all services to packet switching within the next two years because of the technology's international reach and cost-effectiveness. Terminals for Quicknet have already been installed in Western Europe. Ultimately, exchange members will be able to trade on the TSE directly, and in real time, from anywhere on earth.

# Brokers invest in bisync to secure a reliable network

William A. McLoughlin, Drexel Burnham Lambert Inc., New York, N.Y.

Requiring more speed, capacity, and backup than its original network allowed, one brokerage firm scrapped a teletypewriter service, preferring to design its own medium-speed bisynchronous network.

In the spring of 1979, Drexel Burnham Lambert Inc. (DBL), a brokerage house for stocks, commodities, and options, routed buy-and-sell orders received at its lower Manhattan headquarters from its branches around the country to the nearby New York and American Stock Exchanges. It returned order verifications to the branch offices that had originated them. For the 46-year-old firm, this was business as usual, except that the order traffic was transmitted more than 30 times faster than it had been previously—up to 2.4 kbit/s instead of 75 bit/s.

A new store-and-forward message-switching computer network caused the leap in speed over the former teletypewriter network. Before the implementation, DBL had relied on a securities-industry timesharing service called Brokerage Transactions Services Inc. and 75-bit/s teletypewriter lines to link its 56 domestic and international offices. But in the fall of 1976, DBL decided to develop in-house a bisynchronous network.

The heart of the new network consists of two central processing units, an IBM System 370/158 Model 3 and an IBM 3031 mainframe computer. The 370 operates on-line, and the 3031 performs off-line batch processing, communications software testing, and program development. The on-line computer drives both the local and the remote communications networks under IBM's TCS/TCAM (telecommunications control system/telecommunications access method) software. Developed especially for the brokerage industry, this software handles message switching, routing, queuing,

retrieving, and error detection. It supports a variety of terminals, including the IBM 3270 and teletypewriter devices, which facilitated the changeover from the original network's teletypewriter equipment.

The on-line setup also supports a real-time order-processing function via IBM's securities order match (SOM) software. SOM accepts orders and messages from registered brokers at remote branch offices in accordance with the New York Stock Exchange Floor Communications Standards.

These orders are validated and stored in an on-line disk file of active orders. At the same time, they are switched to the appropriate exchange floor and terminal. Execution reports coming back from the floor are matched against the original orders in the order file. An execution notice is then transmitted to the originating branch office, and a trade report is recorded on an activity file for later processing.

DBL's network is made up of three local loops—one on-line, two off-line—and a remote network. The on-line local network consists of a mixture of approximately 60 IBM 3278 CRTs and IBM 3287 printers connected to two IBM 3274 control units that, for communications with host computers, can support up to 32 terminals or printers connected by coaxial cable. It connects many of the firm's major departments.

The other two local networks, which are not on-line, each connect to an IBM 3274 control unit with IBM 3278 CRTs. One network transmits inquiries about accounts, securities, and stock records. The other is

used by programmers for software development.

The remote network, which consists of two IBM 3705 communications controllers acting as front ends, connects domestic and international offices. One 3705 is on-line, while the other is for backup and testing. Either 3705 can attach to either CPU through a channel-interface enable/disable switch on the 3705's front panel.

Each 3705 operates under an emulation program. The on-line 3705 presently supports 164 remote devices connected to 18 bisynchronous circuits and 24 teletypewriter circuits.

### **Probing network requirements**

Before setting up the domestic branch-office bisynchronous network, DBL studied a week's worth of traffic statistics on message size and type, frequency, and peak periods at each branch. From this information, DBL developed a traffic profile and priority profile for each branch office. These statistics and growth projections were then submitted to AT&T Long Lines, IBM, and General Telephone and Electronics Information Services for estimates on network development.

Each vendor returned a branch-office network-circuit configuration that met DBL's requirements of a 5-second message response time for 90 percent of all transactions. The proposals indicated that three to four multipoint circuits at 2.4 kbit/s would deliver that level of service. They also showed that, although beyond the specification, a six-circuit network cost only 5 percent more than a four-circuit configuration.

After reviewing the vendors' recommendations, DBL adopted AT&T's six-circuit configuration, but with a split of three digital and three analog circuits because of availability. [The performance of AT&T's digital circuits is superior to that of its analog circuits. The only drawbacks of digital circuits are the long lead times for ordering and the limited geographic availability.] Since then, DBL has installed mainly multipoint 4.8-kbit/s circuits.

At most domestic branch offices, the firm knew it needed at least one CRT terminal and printer. But several key locations had separate securities and commodities operations requiring additional equipment.

It was agreed that using IBM software with non-IBM terminals could lead to finger-pointing by the non-IBM-ers when problems arose. The headquarters local network therefore would use IBM 3274 controllers, 3278 CRTs, and 3287 printers, and the remote network would use IBM 3276 CRT/controllers, instead of 3274s, in addition to 3278 CRTs and 3287 printers.

This covered all DBL local and domestic locations except the New York Stock Exchange and the American Stock Exchange. For these sites, DBL sought seven CRTs and printers that had the multifunction keyboard capabilities of the aging Sycor Inc. terminals it was using, but that could operate under the IBM 3270 protocol. The terminals were to include new functions relating to options trading, as well as format-error-detection features to catch errors at the operator level, rather than at the CPU level. DBL's search ended at Megadata Corporation, in Bohemia, N.Y. Megadata's SIR 1000 CRT/controller could act as an IBM 3271

cluster controller and could also drive a Teletype Model 40 printer for receiving buy-and-sell orders. In addition, it conformed to the stock exchanges' type-font requirements. Megadata agreed to custom-design the terminal's keyboard.

To promote a gradual changeover, the balance of the network, including several domestic correspondent offices and a few offices at the headquarters, initially kept their teletypewriting devices. Those at headquarters were retained as backup units. During a network outage, the teletypewriter paper tape would store traffic until the system went up again. Today, because DBL has made other provisions for redundancy, this backup arrangement is no longer necessary, and all but two teletypewriters have been eliminated.

### The international connection

For several reasons, most of them economic, DBL decided to retain its international point-to-point teletype-writer terminals for the new network. In the international branches, the point-to-point arrangement is part of a "speech-plus" setup known as SVD, or simultaneous voice data, provided by ITT. It allows a user with a voice requirement to derive as many as five 75-bit/s data channels in addition to the voice channel. In the original network, all international teletypewriter channels came to DBL headquarters from ITT's Technical Operations Center (TOC), located one block away.


Each circuit was routed into the building through teletypewriter hard-copy monitors. These hard-copy monitors had two functions: to provide a headquarters reconciliation copy of all international transactions and to communicate directly with international locations in the event of a Brokerage Transaction Services Inc. computer failure. In that case, to keep a portion of the network operating, a DBL supervisor would flip a switch eliminating the monitors from the circuit. Then, DBL could communicate point-to-point with remote offices.

### Influencing factors

Since the network represents the lifeline by which DBL conducts its business, the firm planned its network around three reliability criteria: (1) redundancy of key components, (2) simplicity of design and operation, and (3) diagnostic capabilities.

In the local network, DBL had two IBM 3274 control units, each supporting close to its maximum of 32 devices. These control units were critical to network operation, because the loss of one would put half of DBL's headquarters off-line. To ensure redundancy of key components, DBL installed a third control unit configured to back up either of the two on-line controllers. For simplicity of operations, this third control unit was switchable. On the channel side, it can be switched through an IBM 2914 matrix channel switch. On the device side, all three controllers are cabled into T-Bar Inc.'s Model 5182 two-for-one (two on-line, one backup controller) switch assemblies and its Model 5106 switching control unit. If a control-unit fails, operators can begin restart procedures by entering an off-line command for the failed unit from one of the two network-control CRTs or replacing the failed unit on the

**1. Local control.** When the switch is in the N, or normal, position, controllers 1 and 2 are on-line; in position A, 2 and 3 are on-line; and in position B, 1 and 3 are on-line.



channel with the standby 3274 from the 2914 panel. As shown in Figure 1, the T-Bar cabinet switching control unit, reacting to a failed IBM 3274 control unit, would switch all 32 devices from the faulty 3274 to the standby unit. An on-line command restores the 32 devices. Total downtime lasts less than 1 minute.

In the original remote network, two 3705s could access either CPU through front-panel enable/disable switches. But it was necessary to find a simple way to attach the 40-plus remote circuits to either CPU. The firm selected a Spectron digital switch patch (DSP) unit. Each DSP handles one bisync circuit or 12 two-

2. Local-loop backup. At DBL headquarters, the modem or DSU interface is linked to an RS-232-C switch. This switch is linked to the Spectron switch. If a local loop or

wire teletypewriter circuits. In addition to switching, the DSP allows the 3705 bisync port to be patched to any modem and vice versa. It also enables technicians to monitor the individual bisync circuits.


But the DSP units, although providing 3705 switching for the teletypewriter circuits, could not provide the necessary teletypewriter-circuit patching and monitoring diagnostic capabilities. So, DBL ran each teletypewriter circuit through an ADC Products voice-frequency patch field before attaching it to the DSP.

To ensure simple operation, DBL installed a Spectron remote master switch (RMS), which allows gang-transferring of all DSPs and all circuits from one 3705 to another. One reason DBL chose Spectron's DSP was that the device could be manually, as well as electrically, operated. This would prove particularly important if the digital power module (DPM), which electrically drives the DSPs, failed. In that case, an operator could manually transfer remote circuits from one 3705 to the other. This override feature was not available on some vendors' front-end switches.

Something else DBL liked about the Spectron RMS was that it lent itself to the firm's future plans. DBL knew that one day it would eliminate teletypewriter circuits. The eight DSPs, which are presently associated with the teletype circuits, could then be retrofitted for bisync circuits.


To implement the new remote network, DBL installed Bell System modems and digital service units (DSUs) and Bell System lines. However, it was known from past experience both that the weakest link in the carrier's provisions was the local loop and that losing the central site modem or DSU on a multipoint circuit would disable many stations. So, for backup, in the event of a modem or loop failure, two modems and DSUs and two local loops were installed at DBL head-quarters for each multipoint circuit. Each modem or

modem fails, a network supervisor can switch over to the alternate device to link the circuit to an on-line 3705, while AT&T repairs the defective component.



3. Dallas to New York. DBL can switch the 3276 control units of Schneider Bernet & Hickman, the firm's domestic correspondent in Dallas, to another circuit, despite one

circuit's total failure. This keeps the four SBH offices involved functioning through the use of modem-sharing devices and RS-232-C switches.



DSU was wired to its own local loop. The two local loops were bridged at the telephone company's central office. At DBL headquarters, the customer interface connection of each modem and DSU was linked to the RS-232-C switch also provided by Bell and housed in the same cabinet with the modems and DSUs.

This switch was linked to the Spectron DSP switch (Fig. 2). In the event of a modem or local-loop failure, an operator can flip one switch connecting the alternate modem or DSU to the on-line 3705. This arrangement enables DBL to get back in service within minutes, while allowing AT&T to repair the defective component. When this arrangement was first installed, DBL had only three analog circuits. The firm has since installed four more for reasons of space and economy.

The third reliability consideration was diagnostics. From experience, DBL knew it needed:

- A pair of d.c. milliammeters to measure a line current of 62.5 milliamps required by the 3705 on each side of the circuit.
- A bias and distortion meter to ensure the receipt and transmission of "pear-shaped" signals (unbiased and undistorted). This and the milliammeters were installed in the Spectron cabinet that houses the diagnostic equipment.
- Two teletypewriter keyboard send-receive units for hard-copy monitoring of low-speed data and control characters. These devices were provided by Western Union and wired into the ADC patch fields so that they could be patched into a circuit whenever necessary.

To monitor bisync devices, DBL learned that the handiest unit was a data-line monitor, which could dis-

play on a CRT device the actual data and control characters generated from and to the 3705.

At the time, approximately eight companies provided such equipment, but DBL chose the Spectron Datascope Model 502-KB, based on its established track record of simple operation. This interactive unit allows passive monitoring of data; trapping and counting various events such as polls, NAKs (negative acknowledgments), and RVIs (reverse interrupts); as well as emulating the CPU and testing new devices. In addition to these functions, it emulates a terminal and enables testing of a new 3705 port. These capabilities required Datascope programming.

Recently, DBL added to the diagnostic console an RTA-270 response time analyzer for 3270 devices, manufactured by Teleprocessing Products Inc. of Simi Valley, Calif. This unit can passively capture various measurements, such as poll-to-poll and bid-to-poll times, CPU and terminal delays, the number of messages, and the number of characters per message. These measurements can be viewed at any time from the front-panel LEDs.

This analyzer has proven extremely useful in determining the impact of broadcast messages on response time and bid-to-poll delays. DBL hopes to put IBM's CICS (customer information control system) on-line to its branch offices and believes that the RTA-270 will again prove invaluable in monitoring circuit loading and response times.

As in any business, certain remote locations are more important to the overall operation than others. DBL identified eight such offices. Again, the firm was