
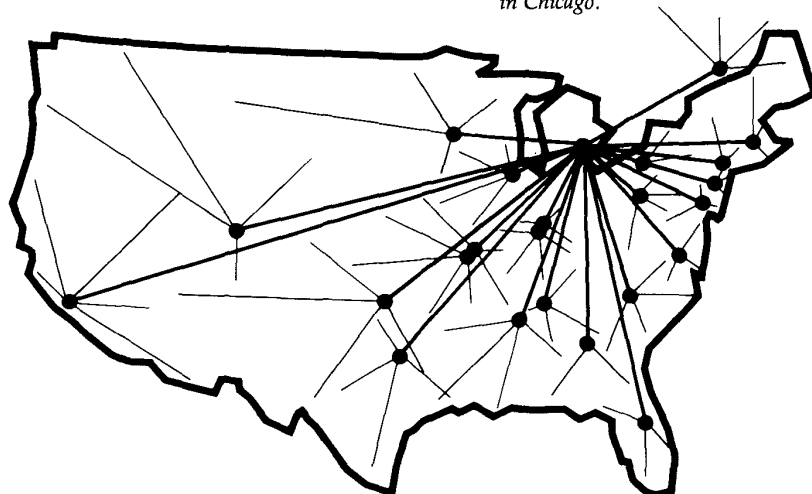


Case Study: The CIRRUS Banking Network



This is the third case study to appear in Communications. This time, case-study editors David Gifford and Alfred Spector interviewed the president and the director of systems and operations at CIRRUS Systems, Inc., a funds transfer network, to see how distributed systems are being employed in the banking industry. As the following transcript shows, the CIRRUS network provides security, lost message recovery, database consistency via atomic actions, recovery after network partition, and a framework for interorganization cooperation in a distributed system. The CIRRUS system is not the most complex computer system to have been described in these case studies; rather, it is an example of how a few well-chosen ideas can be the basis of a simple and useful system. The interviews were conducted on March 29, 1985, at the headquarters of CIRRUS Systems, Inc., in Chicago.



THE CIRRUS BANKING NETWORK

The CIRRUS banking network makes coast-to-coast automatic banking transactions possible. The system will soon be able to handle international currency transactions and point-of-sale transactions in stores.

DAVID GIFFORD and ALFRED SPECTOR

Interview 1. Bruce Burchfield, president of CIRRUS Systems, Inc., talks about the development of the CIRRUS network, its present scope, and plans for its future.

DG. Bruce, could you start by telling us something about the history of the CIRRUS network?

Burchfield. I'd be glad to. Bankers have traditionally provided customers with personalized service. That's the nature of the business. By the 1970s, most banks had automated their "back-room" operations, but were still using people for all their customer interaction. During the 1970s, though, both technology and customer demographics changed. On the one hand, automated teller machines (ATMs) were becoming reliable and cost effective, and on the other, customer demographics were changing and banks were under pressure to provide more convenient access to their services.

Bankers couldn't really afford to provide essential full-branch service for 16 or 18 hours a day, so they installed ATMs to keep certain essential services available around the clock. The next step was to situate these ATMs away from existing bank branches. The choice for bankers was either to build \$1 million branches that would cost a half a million a year to operate, or to install 24-hour ATMs for about a \$100 thousand each and then \$50 thousand a year for maintenance. From the banker's perspective, ATMs are a very cost-effective way of making their services available in more places and at more times.

Then, around 1976, bankers realized they could reduce their costs even further by sharing ATMs. Since there's usually a certain amount of excess capacity on an ATM that's only serving one bank's customers, some banks began to sell their unused capacity by sharing their resources with client banks and then charging a set fee for every transaction processed. Reciprocal shar-

ing arrangements also provide participating banks with better distribution channels, which can mean a competitive advantage.

These developments have allowed banks to reduce staffing and cut back on hours in their regular branches, and to deal with competitive threats from companies like Sears and American Express that are entering the financial-services business with a cost structure that would drive unmodernized banks out of business. It's a traditional paradigm—when the structure of an industry changes, the old leaders are displaced if they can't keep pace.

DG. So by the late 1970s you knew not only that ATMs were going to be important, but that ATM sharing was also going to be important. How did the idea of CIRRUS begin to take shape?

Burchfield. Banks had already formed local and regional sharing arrangements. For example, I had developed a shared automated-teller-machine network in metropolitan Chicago called Cash Station. The next logical step was to provide access for customers regardless of where in the country they happened to be.

CIRRUS was started by 10 banks that wanted to expand their regional networks into a national network. These included BayBank of Boston, Manufacturers Hanover of New York, Mellon Bank of Pittsburgh, and First Chicago. The first discussions were held in December of 1981, and by July of 1982 CIRRUS was incorporated. The initial service that CIRRUS set out to provide was cash access for people away from home.

We knew we needed a certain critical mass to make the network successful, and so we worked to build the network up quickly. As of March 1985, we have 46 states covered. In building this network, we've created a technical infrastructure that's unprecedented in the banking industry. The number of member banks has

© 1985 ACM 0001-0782/85/0800-0798 75¢

now grown to 1425, although only 16 of them are directly connected to the CIRRUS switch. For example, there are over 300 banks in Texas on CIRRUS, but they're all connected to a single node. The Texas regional network handles all intrastate traffic and forwards the remainder to CIRRUS via the Texas CIRRUS node. There are probably 150 million to 200 million transactions a year that are processed by CIRRUS member banks, and yet only a fraction of that goes through the CIRRUS switch. Thus, in many cases, CIRRUS is a brand name as opposed to a processing function. Our node-based architecture has worked out well, and we don't think the politics or the structure of the ATM market will ever be conducive to connecting individual banks directly.

DG. What other services do you now support besides cash access?

Burchfield. We support withdrawals and balance inquiry from checking, savings, and line-of-credit accounts, and we are also developing direct debit point-of-sale services that will allow us to put machines in retail outlets. This summer we're connecting to a Canadian ATM network, and we're planning to do automatic currency conversion for international transactions.

We're also considering the possibility of taking deposits over CIRRUS. Interbank deposit taking already exists on a regional level—on the Chicago network I was involved with, the bank that received a deposit would prove and verify the deposit envelope, process the checks, and return bounced checks to the depositor's bank. The actual transfer of funds between banks is handled electronically.

DG. When CIRRUS first started, did you analyze all the different exposures of the system and try to put together a study that could convince all of your participating banks that CIRRUS was going to be secure?

Burchfield. To resolve the security issue, we assigned liability for all of the things that could possibly go wrong. For example, all of our connection nodes have financial and data security responsibility for their traffic. Extensive logs are kept for audit purposes to help assign liability. This gives each CIRRUS node an active interest in running a secure system. The employees of each bank are bonded, and the switch is operated by a bank with extremely strong audit functions. We also have provisions in our contracts that prohibit banks from using transaction data for competitive analysis.

AS. With all this delegation, it sounds as though the CIRRUS organization itself doesn't have to worry about liability.

Burchfield. That's correct. For example, the recent failure of certain savings and loan organizations in Ohio didn't cause us any concern because their CIRRUS attachment node, Central Trust in Cincinnati, is liable for their activity. If the savings and loans went defunct, CIRRUS would not be out of any money because Central Trust is responsible for their transactions.



BRUCE BURCHFIELD

Bruce Burchfield is president and chief operating officer of CIRRUS Systems, Inc. He was formerly vice-president and manager of the electronic-banking service division at the First National Bank of Chicago, where he was responsible for the development and management of new electronic-banking services, including electronic funds transfer equipment and First Chicago's shared ATM network, Cash Station. Before joining First Chicago, Burchfield was an engineer for Reynolds Aluminum.

Central Trust actually shut their CIRRUS access off as soon as there was a problem.

DG. How are consumers protected against the misuse of their cards?


Burchfield. They're protected by Federal Regulation E, which says that a customer is not liable for the use of lost or stolen cards as long as the loss is reported within 2 days. When a lost card is reported within 60 days, but after 2 days—that is, within the amount of time it takes to get a bank statement—liability is limited to \$50. The issuing bank is responsible for any ATM-related problems that a customer may have. That's what we tell customers when they call us—to contact their issuing bank. That's really all we can do, since we have no authority relative to accounts.

DG. Do you provide any services directly to consumers?

Burchfield. Only one—an 800 number that consumers can call to locate a nearby CIRRUS machine. The database of ATM locations is a geographical matrix that's accessed by area code and telephone exchange. The 800 operator asks for this information and then provides the locations of some nearby terminals. This service is subcontracted to an outside firm.

DG. If I ran a bank and felt that I was at a competitive disadvantage because I didn't belong to an ATM network, and I wanted to join CIRRUS, how would you deal with my request?

Burchfield. To answer that question, let me start by telling you something about our structure (see Figure 1). We have three types of members—Principal, Associate, and Corresponding members. Principal and Associate members have direct links to the CIRRUS switch and



The CIRRUS network consists of 16 4.8 kbit/second lines that radiate from the CIRRUS switch at the National Bank of Detroit (NBD) to the various nodes. Principal and Associate members are directly connected to the CIRRUS switch at NBD. A Principal member has exclusive marketing rights for

a particular area, and an Associate member has direct access, but not exclusive rights. Corresponding members are connected to the switch through Principal or Associate members.

FIGURE 1. Topology of the CIRRUS Network

the right to franchise CIRRUS to other financial institutions. Principals and Associates are primarily large financial institutions that operate regional networks.

These Principal and Associate members can bring other banks into the network as Corresponding members. These are generally smaller institutions for which a direct connection would not make economic sense. Therefore, whenever a small bank approaches CIRRUS directly, I refer it to the Principal or Associate member that offers CIRRUS in the bank's state.

DG. What would happen if Bank X joined a regional ATM network like Cash Stream, without actually joining CIRRUS, and tried to send transactions through the CIRRUS switch via the Cash Stream CIRRUS node?

Burchfield. Well, first of all, Cash Stream could sign up Bank X as a member of their regional network, and therefore as a member of CIRRUS, but if Cash Stream actually did send the CIRRUS switch a transaction that originated at Bank X, without working through a regional network, that transaction would be rejected because Bank X would not be in the financial institution table at the CIRRUS switch. It wouldn't be a valid transaction.

AS. Who are your direct competitors?

Burchfield. There's the Plus network, which is run by a group of banks, and Master Teller, which is operated by Master Card. As of the end of 1984, Plus had 4700 ATMs, Master Teller had 3250, and CIRRUS had 6500. Plus has a competing networks rule that prevents their members from joining another network. Master Teller doesn't have such a rule and neither do we, so some banks belong to both CIRRUS and Master Teller. Master Teller is very big on credit functions, of course, but their card base is limited as far as debit transactions or checking accounts, so there's a certain advantage for customers in being connected to the two networks.

DG. So if there were two banks that were connected to both CIRRUS and Master Teller, they would have two different ways of routing transactions between themselves. Does that make transaction switching a competitive business?

Burchfield. Transaction switching is not competitive right now because the bank that decides which network to use is not the bank that pays for the transaction. The destination bank has to pay for the transaction and would thus like to receive it over the more

economical network; the acquiring bank, on the other hand, wants to issue the transaction in the most convenient way, which is probably the least favorable to the user.

AS. How do CIRRUS member banks recover the cost of processing transactions for other member banks?

Burchfield. The issuing bank pays the deploying bank 50 cents for a withdrawal and 25 cents for a balance inquiry. The switch gets an additional 25 cents per transaction. Right now, most CIRRUS member banks offer CIRRUS services for free, but my assumption is that members may eventually pass their costs on to the customer. Customers will then have to compare the cost and convenience of CIRRUS with traveler's checks and other alternatives.

AS. Do you foresee the establishment of gateways for connecting CIRRUS, Master Teller, and Plus?

Burchfield. We don't think there's any need for connecting ATM services. We see our independence as a competitive advantage. For point-of-sale services, though, it makes sense to allow merchants to accept different kinds of cards. When point-of-sale machines are connected to more than one network, it becomes logical to standardize fees and message formats. These issues are already being addressed by the American Bankers Association and the Electronic Funds Transfer Association (a multi-industry trade association for electronic funds transfer). CIRRUS, Plus, and Master Teller are already meeting to talk about technical standards, to try to build a common gauge railroad, so to speak, so that point-of-sale transactions won't happen in a fragmented way. We're all very competitive, but we realize that it's not going to work unless there are standards.

Interview 2. Jay Levy, director of systems and operations for CIRRUS Systems, Inc., talks about the design and implementation of the network.

AS. Jay, we'd like you to take us over some of the technical details of the CIRRUS system. Could you start by giving us an idea of how it's organized?

Levy. Sure. CIRRUS is a star network, with a message switch at the National Bank of Detroit that currently connects to 16 nodes via 4.8 kbit/second dedicated lines. If a line fails, it's replaced by a dial-up circuit that runs at the same speed. The communications protocols are bisync point to point. The CIRRUS switch is the secondary point on each line, so each node processor has primary access for the bid on its line.

The switch itself is a Tandem NONSTOP II system with four CPUs (See Figure 2). Each CPU has 2.4 megabytes of main memory. The system has a pair of 256-Mbyte drives, a pair of 128-Mbyte drives, and a pair of 64-Mbyte drives. These are shared by all of the CPUs. Most of the switching functions are handled by one Tandem CPU with another acting as backup. When we run tests, we take two of the four CPUs and configure a test system.




JAY LEVY

Jay Levy is director of systems and operations for CIRRUS Systems, Inc. Prior to joining CIRRUS, Levy was an assistant vice-president at the National Bank of Detroit (NBD), in charge of design and development pertaining to NBD's operation of the CIRRUS ATM interchange switch. He was also responsible for developing the ATM program from a proprietary system to a regional network, and assisted in the development of telephone-bill payment services, videotex banking applications, and international retail banking operations.

The 256-Mbyte disks are used for daily transactions—one disk drive for primary and one for backup. The 128-Mbyte disks are used for control files, processor status files, terminal files, institution-level files, and holiday records. They also contain the routing tables and the system definition file, although these are kept in main memory when the system is in operation. The 64-Mbyte disks contain system files.

The switch sits next to another Tandem system at the National Bank of Detroit that could be reconfigured to take over the CIRRUS load if the primary system failed. About the only thing that would cause a real disaster is a power failure, since there is no backup



The CIRRUS switch uses redundant CPUs and disk drives to maintain a high degree of availability and reliability.

FIGURE 2. CIRRUS Switch Configuration

power supply at the switch site. However, NBD does have a backup site that wouldn't be affected by a power failure in Detroit.

DG. How much log data do you accumulate on a day-to-day basis?

Levy. A log record is 411 bytes. There are two log records written for each transaction. One is written into the log file that corresponds to the transaction's anticipated date of settlement. The log file for a Monday would contain between 24,000 and 40,000 records. The switch stops accepting activity for a current date at 8 P.M. Detroit time, which means cutting over to the new log file. Just before cutover, approximately 99 percent of our traffic is already going to the new date of settlement, as specified by the terminal owner.

The other log record is to a backup transaction file that's in time-stamp order. We've been fortunate in that we've never had to use the backup log file.

AS. What do you do with the log that you've accumulated each day?

Levy. The log is the source of the batch reports that we send to each node bank's data-processing organization each day. A node's report lists all of its activity for the day in time-stamp order. The report has separate sections for inbound and outbound traffic. The batch run is also used to verify the on-line settlement totals that indicate the balance of each node with the switch. If the batch run produces totals that are consistent with the on-line totals, settlement is effected by transferring funds between "due from" and "due to" accounts that are maintained by each node bank at NBD. The actual transfers are accomplished by the money-management departments of member banks. We only do settlement on a node-to-node basis—each node must deal separately with its member banks.

AS. What happens when the batch node balances don't agree with the on-line node balances?

Levy. We look for the reason. Typically we'll be out 25 or 50 cents for a missed transaction fee somewhere. We've never had a serious discrepancy in a consistency check.

AS. What is the peak transaction load on the switch?

Levy. The switch was designed to handle 2 transactions per second, which is 16 messages a second. I don't think we've ever achieved that peak—we have never exceeded 500,000 transactions a month. The total elapsed time for a transaction averages between 12 and 15 seconds from request to cash dispense. This includes the processing time at a node bank for approval and debit.

AS. Suppose I had an account at BayBank and that I went to an ATM run by Landmark Savings and Loan in Pittsburgh, which happens to be a member of the Cash Stream network. Mellon runs Cash Stream and

is a node on the CIRRUS network. What happens when I put my BayBank card in the Landmark ATM and request \$100?

Levy. You would get a *customer lead through* on the screen that would look exactly like the lead through for regular Landmark customers. The ATM wouldn't know who you were yet, and so its processor would ask you for your personal identification number (which it would temporarily buffer), what type of account you wanted to access (i.e., savings or checking), and the dollar amount of your transaction.

By this time the Landmark ATM would have already read Track II of the magnetic strip on the back of your BayBank card, and determined your account number, expiration date, and some other variable data that are dependent on BayBank. The first 3 to 11 digits of an account number specify the issuing bank. Track II holds 40 bytes, although we're only using 19 of them right now. Tracks I and III we don't use at all—Track I is for the airline industry, and Track III is used in off-line banking transactions, but not in the interchange environment.

To get back to your transaction, though, the host processor at Landmark has not been involved at all up to this point. Once the ATM has collected all of the data it needs for your \$100 withdrawal, it sends a message to its host at Landmark, which determines that the account number on your BayBank card is not a Landmark account, and creates a message to Cash Stream. The Landmark host also writes a suspense record to its log so that it can time out the transaction and free the ATM if Cash Stream does not reply after a certain amount of time. The Landmark host would probably give Cash Stream about 34 seconds before doing this. Most ATMs also have internal timers—we suggest that ATMs abort transactions and return consumers' cards if they do not hear from their hosts in 45 seconds.

When your transaction arrives at the Cash Stream switch, Cash Stream validates it to make sure that it has a valid business date, dollar amount, and so forth. Cash Stream would then look up your bank number in a routing table, determine that BayBank is not a Cash Stream member bank, and then create a message to CIRRUS in a variation of ANSI X9.2 format. The message would contain a lot of information, including a 40-character description of the ATM, to comply with government regulations. Once the message is created, Cash Stream logs it and sends it to the CIRRUS switch via a bisynch point-to-point protocol. Cash Stream waits up to 24 seconds for a response before aborting the transaction.

CIRRUS receives the message, looks up the issuing bank number in its routing table, and then forwards the message to the CIRRUS node that the issuing bank is connected to, which in this case is BayBank. The CIRRUS switch sets a timer for 20 seconds and then waits for the response from BayBank. BayBank receives the message, verifies that the personal identification number (PIN) you supplied is correct, and then checks your balance. If you check out, BayBank writes a memo

post record on your account indicating that a \$100 withdrawal is in progress, and then sends an approval message back through CIRRUS, Mellon, and Landmark to the ATM, which dispenses your \$100.

The memo post record that BayBank has written is simply noting that a withdrawal is in progress. It will *not* cause \$100 to be deducted from your account. That happens with the second round of messages, which I haven't yet discussed.

When the ATM has actually dispensed cash into your hands, it sends a message to Landmark saying that it has given you \$100. The completion message goes from the ATM through Landmark and Mellon to CIRRUS, and then CIRRUS sends a message back to Landmark acknowledging the completion message. CIRRUS also sends the completion message on to BayBank, which uses it to create a "postable" record that is used to deduct the money from your account. A completion acknowledgment then goes back from BayBank to CIRRUS. Altogether, your transaction has required four CIRRUS messages—two for the initial approval and two for the completion.

If CIRRUS does not receive a completion acknowledgment from BayBank after 120 seconds, the switch continues sending until it does receive one. If for some reason BayBank cannot deal with the completion message, they call us and we purge it from our file.

Approximately half of our member banks use this protocol. The other half use a shorter protocol that eliminates the second two messages. In our example, this would mean that the CIRRUS switch would still get the completion message from the ATM, but that it would not forward it to BayBank. When this shorter protocol is used, the issuing bank must simply assume that a withdrawal has actually taken place. Obviously, if there's a problem then a message must be sent to the issuing bank instructing it to compensate for the debit it posted.

AS. What would happen if BayBank had sent an approval to the Landmark ATM and I had received the \$100, but the completion message had not returned from the ATM to BayBank because the Mellon Cash Stream system had failed?

Levy. If the CIRRUS switch does not get a completion message from Mellon within 120 seconds, then Mellon would be marked down by the switch and would be placed off line for 10 seconds. CIRRUS would send a message to BayBank saying that the cash had not been dispensed. After the 10 seconds had elapsed, we would bring Mellon on line by issuing a series of network-management messages. If the completion message from Mellon came in after the 120-second time had expired, but within the same settlement day, the CIRRUS switch would interrogate the log file to see if there was a record for that message. If there was, the switch would automatically batch adjusting entries to both sides of the transaction to ensure that the settlement was handled properly. If the late completion message came in after the settlement cutover for the day, or never came

in at all, Landmark, as the ATM operator, would have to present an off-line paper adjustment to the switch, along with some kind of physical evidence of the transaction, like an audit tape or an internally produced report indicating that the transaction had occurred, and a CIRRUS report indicating that they hadn't received value for the transaction. Without that documentation they wouldn't get their money. This kind of thing is not uncommon—some ATMs don't automatically send completion messages to the host. The only way to discover what's happened is to get a status of prior transactions after a new transaction starts.

DG. Could you review the way the time-out intervals on your system interact?

Levy. We have a 20-second timer at the switch for messages to nodes; we advise sending nodes to set a 24-second timer to allow some transmission time. A bank that's not directly connected to the switch, like Landmark in your previous example, may wait as long as 34 seconds to allow the intermediary bank some processing time. We typically tell the ATM issuers to set their timers for about 45 seconds; if the terminal doesn't get a response within that time, it shuts down and returns the customer's card.

DG. The two message protocols you described—the four-message protocol and the two-message protocol—make radically different assumptions about when to deduct money from an account, don't they?

Levy. Traditionally, the U.S. banking industry won't ever deduct anything immediately. We have very few real-time posting, checking, or savings applications in the U.S. That convention doesn't always hold true for other countries, though. In Canada, for instance, there are systems that use real-time posting, which means that a single customer transaction can post several items on a cardholder's account. In the U.S., banks use on-line systems for maintaining account balances. Posted records from a system are merged with other instrument activity (such as checks) during batch processing. The daily batch run is what actually deducts the money from accounts. The result of a daily batch run is a file containing the starting account balances for the system for the next day, as well as records for pending transactions.

DG. From what you've been saying, there seems to be little concern that any of the participating banks might be doing anything dishonest. What would happen if a member bank immediately acknowledged that an ATM had dispensed cash when in fact it hadn't. Wouldn't that bank be credited for \$100 while the customer was left standing in the rain without any money in his hand?

Levy. Customers would let us know if there was anything of that nature going on. Also, there's a significant amount of federal regulation that protects customers in

those situations, so it's hardly in the member bank's best interest to do such a thing.

AS. But the safety and security of your network depend very much on the safety and security of its constituent parts.

Levy. Absolutely. CIRRUS was developed as a pipeline between banks that had been doing this kind of thing for 10 or 12 years before we came along. We're another link of potential authorization for the banks. All of the policies and procedures that banks had evolved to protect themselves from error and fraud were carried over directly to the CIRRUS environment.

DG. You mentioned before that you partition any part of your network that doesn't respond within a certain amount of time. Is that an important part of your system strategy?

Levy. Yes. I'm probably one of the few proponents of this kind of partitioning. The more popular strategy is to keep resending messages until a response is received. I believe that the financial implications of transaction processing in the ATM environment require this kind of scrupulousness. It has proved easier for us to balance the network by guaranteeing that transactions flow out of the store-and-forward process in a logical sequence. We've also found that if you give a node time to recover you can often prevent more serious problems.


DG. How often do banks go off line?

Levy. We're working to get that rate down to a minimum. We have a standard that no bank should be off line more than five percent of the time, averaged over a month. Some of our members are well above that mark, and some are well below it, but most are off line no more than five percent of the time. Nearly all of our member banks use highly available systems, but as I mentioned earlier, if a node does not get a completion message in time from an ATM, CIRRUS may take the entire node off line. Of course in the case of a delayed completion, a message goes out 10 seconds after a node goes off line to bring it back up again. These short ups and downs have a minimal effect on transaction processing.

DG. CIRRUS itself could also go down, couldn't it?

Levy. It's not likely. The average availability of the CIRRUS switch exceeds 98 percent on any given day (see Figure 3). We do not allow scheduled outages of the CIRRUS switch. Actually, we allow one a year, but it still goes against that 98 percent. Tandem equipment does not require outages for hardware maintenance.


DG. Considering communication-line failures, central-site failures, and node failures, what's the probability that one will be able to access another at any given time?



The average uptime on the CIRRUS switch is between 98 and 100 percent over a 24-hour period. Scheduled downtime

is counted as unavailable time.

FIGURE 3. Average Availability on the CIRRUS Switch



Network availability is computed as the probability that one CIRRUS node can access another at any given time. Network availability can be affected by failures in any part of the

network, from the switch to the communication lines to the nodes.

FIGURE 4. Average Availability on the CIRRUS Network

Levy. About 97 percent (see Figure 4). In February, for example, the average network uptime was 96.3 percent, which covered a range for individual nodes from 90.13 to 99.6. The 90.13 node was either having a line problem or a processor problem. The mean daily average availability for the switch during February was 99.62 percent.

AS. What kept you off the other 0.38 percent of the time during February?

Levy. The backup file just got incredibly large over the three-day Washington's Birthday weekend, and we had to shut the system down for a minute or two (or fifteen) to reallocate some file space. There was also a date problem because some of our members wanted to process on that Monday.

DG. The switch was closed?

Levy. The switch observes the seven holidays that all 12 Federal Reserve offices observe as federal holidays. Because the National Bank of Detroit isn't open, there's no way to move the money, and if the Fed's not open, there's no way to fund an account at NBD. Certain banks that close for Lincoln's Birthday stay open for Washington's, though, and they were trying to present transactions when there wasn't a valid log file open,

and that probably led to some downtime.

AS. What fraction of the transactions that reach CIRRUS have to be rejected because they're not destined for a member bank?

Levy. About 15 percent.

AS. Could you describe the message protocol you use?

Levy. It's called the ANSI Interchange Message Specification for Debit and Credit Card Messages on Financial Institutions—ANSI X9.2-1980. It's a guideline, as all standards are guidelines. It defines 64 fields that are usable in a financial message, and it indicates which fields are mandatory for the type of message being sent. At CIRRUS we've enhanced this standard to meet our particular needs. We use all of the mandatory fields, along with some additional ones. There's a bitmap at the head of every message that shows which fields are being used. We have a fixed set of about 40 message formats that we use. The smallest uses about 170 bytes; the largest, a reversal message, over 400. Reversals are long because they contain a great deal of data from the original transaction. For each type of message, we have both a long and a short format. The long format has all of the fields filled in, and the short one requires that missing fields be filled in by the switch.

AS. Could you briefly summarize the types of messages that you use?

Levy. There are request, reply, completion, confirmation, adjustment, reconciliation, and network-management messages. For each type of financial transaction, there's an entire set of request, reply, completion, and confirmation messages.

DG. Are point-of-sale transactions handled separately?

Levy. Point-of-sale transactions, once they're developed, will use the same messages with various fields indicating a POS transaction. For POS we've created a new type of transaction that is a preauthorized request. This would be used at places like gas stations where intended purchases don't always match up to actual purchases (because who knows exactly how much gas their tank can hold). The preauthorization might be for \$20; the customer would pump \$18.50 worth of gas, and the POS terminal would then send subsequent messages out to build the posting record for \$18.50. The memo post, though, would be held for \$20. We want to use a different set of messages for preauthorization because we don't want anyone to confuse them with real purchases.

AS. Is preauthorization similar to the failure condition that occurs when an ATM can only distribute part of a requested sum?

Levy. Yes, except here the discrepancy is the norm instead of the exception.

DG. What are your message latency goals at the switch, and how can you maintain them as your volume increases?

Levy. We don't allow the switch any more than 2 seconds for processing a message. The actual time is usually somewhere between 0.7 and 1 second. The switch is modular enough that if we ever needed to we could open up a mirror of the application so that there could be multiple processors running the application simultaneously, although there would only be one transaction running in any processor at any time. I think we've actually done this at peak processing time.

AS. Do you spool log data onto magnetic tape for permanent safekeeping?

Levy. All the important files, including the log files, are backed up every night. Once a week we back up the entire system, including the application program, and transport it to an off-site storage location. We also put a digested form of the log on microfiche and keep it for five to seven years, to meet regulatory requirements.

DG. How would you compare your electronic funds transfer network to the Fed-wire?

Levy. Well, the Fed-wire is a wholesale-level EFT network, and CIRRUS is a retail-level EFT network. The average dollar value per transaction on the Fed-wire is probably a million times what we handle on CIRRUS.

The aggregate for a whole day of CIRRUS activity would barely warrant a wire on the Fed-wire system. That's part of the reason that we aren't nearly as security conscious as the Fed-wire is—our exposures are not nearly as great. Also, the Fed-wire has much more significant and stringent requirements with regard to who can initiate what type of transaction and how transactions can be effected.

AS. Who writes the programs for CIRRUS?

Levy. We contract the system services from the National Bank of Detroit, which maintains a programming staff to support the switch. NBD subcontracted the writing of the switch to a software house, but did all the design and testing on its own, and NBD owns the code. The software is relatively stable now. I would be surprised if NBD had more than the equivalent of one full-time person budgeted for maintaining the CIRRUS software. When a new release of applications is ready to be released into the system, it has to be certified by the system software group at NBD.

DG. Are new applications like point of sale contracted out to a software house?

Levy. As a matter of fact, we're moving our processing to Mellon Bank because they were the lowest bidder for the POS product. Mellon is going to have to redevelop the switch software.

DG. How many lines of code are in your on-line switch application?

Levy. The on-line switch software, which includes message routing and editing, is approximately 9000 lines of TAL, which is the Tandem application language.

DG. How big is your routing table?

Levy. Currently it consists of about 3500 records, with about 20 bytes per record. A routing-table entry includes a bank's prefix, which is from 3 to 11 digits, and an identification of the processor that services the bank. The routing table is called the FIT, for *financial-institution table*. We send an up-to-date copy of the FIT to our member banks every day. It's our intention to build an on-line message that keeps FIT copies up-to-date. The problem is that bringing a new member on can mean bringing as many as 100 new banks on with it. A hundred new FIT entries at one time is a data-management problem that we haven't solved yet.

AS. Do you test a member bank's system before letting it on the CIRRUS network?

Levy. Yes, we certify nodes with a battery of test cases. Every situation we can think of is attempted prior to putting a node on the network. The book that we keep the test cases in is over three inches thick—I believe it contains over 2500 test transactions. When we certify a node, we ask the candidates to have an ATM available in a test environment. We send them plastic cards that run against a simulated card proces-

sor at the switch in Detroit. The switch can also send simulated ATM activity out to them. We find that using an actual ATM is the best way to test.

AS. How is a typical test case described?

Levy. In great detail. This is a *very* critical part of our operation. I think part of the success of the network, in fact, can be attributed to the stringency of our testing. We send candidate banks a list of cards and a set of conditions that the cards have to meet. Some of these cards specify accounts that don't exist or that have incorrect balances. A test case instructs the candidate bank to use a particular card for a particular transaction, and then indicates the expected result. We test weekend and holiday processing and all the exception conditions. We also keep our set of test cases up-to-date as we add new functions to the switch.

AS. What percent of your budget eventually goes for testing?

Levy. I would say that as a percentage of overall development, testing is easily 25 percent. In an environment like ours, we have to be very careful.

AS. Have you had any situations where a member bank did not perform properly?

Levy. Yes, we've had member banks that stopped sending completion messages for one reason or another. On one occasion we had 1500 adjustments for one day's activity. These had to be presented on paper. This swamped their operation center as well as ours. We have also had requests for adjustments that came in as much as 50 business days after a transaction. It's difficult to accept activity that late.

DG. Could you discuss your basic strategy for security?

Levy. Within the banking environment, there's a great deal of physical security in addition to application-level security. Most any data that are maintained in a bank's data-processing center are considered confidential and secure. The only part of the system that we don't normally have physical control over is our communication lines, which are protected with link encryption provided by the Racal-Milgo DataCrypter II with a public key option. This is a commercial product that uses DES for encrypting messages. It includes a public key algorithm that allows us to distribute new master keys from time to time.

AS. Have your communication lines ever been actually attacked?

Levy. Not that I know. The biggest area for fraud in ATM banking is card and PIN compromise: Either a family member takes a card and cleans out an account, or someone obtains a card and its PIN by force. It's as easy to defraud an ATM that way as to tap the line.

AS. Are you interested in "smart cards"?

Levy. They would allow us to put a great deal more data on plastic and would let us do away with PIN

verification at the issuing bank. They would also allow us to maintain value on cards, so that customers would be able to put, say, a thousand dollars worth of value on a card that could then be spent without on-line verification. The thing that concerns me is how could we get the postable information to the cardholder's bank and actually deplete his account balance. I see the biggest application of chip cards in this country for things like food stamps where the value is depleted but the transactions don't have to be posted. It's not clear what application smart-card technology will have in banking, which may be why it's taking so long to be developed here. There's also the problem of refitting hundreds of thousands of terminals, and there would have to be a standard for where to put the contacts on the cards.

DG. Have you considered digital signature verification for point-of-sale transactions?

Levy. I have my doubts. Handwriting can be compromised. And it's hard to say what kinds of controls we're going to need until we really find out how easy it's going to be to deploy POS. I'm not convinced that signature technology is that much more secure than PINs.

DG. What's the largest transaction that can be issued over the CIRRUS network for a cash withdrawal?

Levy. Theoretically, the field size would allow for a \$100,000 transaction, but I don't know of a machine in the country that would let you do that. There are machines in Las Vegas that would probably let you take \$1000 out in one shot.

AS. Will future ATMs have change-making capabilities?

Levy. Yes, there's a machine that makes change, and five-bill dispensers that can give ones, fives, tens, twenties, and fifties or hundreds. The more these machines can do, the more truly they'll be able to replace human tellers, which is really valuable in some applications. There may not be a need for change-making capability at O'Hare Airport, but if someone's going to an ATM regularly for their everyday banking business, it should be able to cash their paycheck. Not everyone's paid in increments of \$10.

DG. I think that's all the questions we have for you. Bruce, Jay, thanks for telling us about the CIRRUS system.

CR Categories and Subject Descriptors: C.4 [Performance of Systems]: reliability, availability, and serviceability; D.2.9 [Software Engineering]: Management—software quality assurance; D.4.0 [Operating Systems]: General; J.1 [Administrative Data Processing]: financial; K.6.3 [Management of Computing and Information Systems]: Software Management

General Terms: Management, Reliability, Verification

Additional Key Words and Phrases: ATM, CIRRUS, EFT systems

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Association for Computing Machinery. To copy otherwise, or to republish, requires a fee and/or specific permission.