M.I.T. Laboratory for Computer Science

Computer Systems Research Division

January 23, 1978

Request for Comments No. 155

DISTRIBUTED COMPUTER SYSTEM IN A BANK:

Notes on the First National City Bank

by Liba Svobodova

As part of a study to understand the nature of computing and the role of computers in a bank, I visited the First National City Bank in New York City. The First National City Bank, or shortly Citibank, is quite unique in its organization and management of computer facilities. Computer services are completely decentralized. The decentralization was adopted as a consequence of the growing cost of maintaining the original computer center that had 18 computers (IBM S/360, Burroughs, and also some quite ancient IBM machines) and large (and growing) programming staff. As more and more services were put on this centralized computer facility, the cost per transaction started growing, primarily as a result of growing cost of human labor. Software maintained at this large computer center was so complex (and apparently, poorly designed) that it took a sizeable group of programmers to check and debug various improvements and patches made by the previous group of programmers. In 1974, a proposal was made that the computing services be decentralized. A study of the bank's operations revealed sectors that were more or less self contained

and separate computers were placed into these individual sectors. Unfortunately, no direct communication was provided between computers in different sectors. All the necessary communication has to be carried out manually, either by means of various paper forms, or by shipping around magnetic tapes. Currently, Citibank is working on a development of a network to connect their individual computers within New York City. This network is called Citinet and is based on ARPANET technology.

The first part of this report discusses the structure of Citibank, the bank's present use of computers within this structure and the planned communications network, Citinet. The second part presents a rough model of a distributed computer system needed in this type of environment.

STRUCTURE AND OPERATIONS OF CITIBANK

The First National City Bank has a world-wide network of branches. It covers over 100 countries and has physical offices in over 300 cities. Services provided by the bank can be divided at the first level into retail services and corporate services. Corporate services (Citicorp) definitely present a more interesting case for study of distributed computing. At the Citicorp the business is viewed as a movement of information. The information on funds availability, information on foreign currency, information on account balances, information on loan status, money movements, money market instruments, buys and sells, delay orders and other items. Citicorp is divided into a number of compartments. The most important seem to be the National Banking Group (NBG), International Banking Group (IBG), and World Corporation Group (WCG).

Until 1974, most of the information processing needs were channeled to a central computer system managed by the operating group. As the volume of

transactions to be processed daily increased, more and more people were brought in to handle the load. The new people were not usually trained properly, and consequently created more errors. Thus, more people were needed to solve the added problems. Staffing levels appeared to be growing uncontrollably and correspondingly, the cost was rising approximately 15% per year. Finally, it became clear that the situation demanded a radical decision. The decision was to reorganize the information processing services, to decentralize them. The plan was to consolidate all information processing related to one product under one organizational unit. Each such unit is called a channel. Each channel has its own minicomputer-based system. Each channel reports to the division head, but now divisions are more manageable, approximately 500 to 600 persons for each division. This structure enhances the autonomy of individual divisions and in addition to reducing the overall cost of processing services in the bank, it has also a positive impact on how people feel about computing services.* This approach seems to have been beneficial not just from the point of management of computer resources. It helped consolidate support of customer services in such a way that many elaborate transactions can be now accomplished much faster and with just a small number of people. The bank management's view of the bank's goals and achievements is summarized in the statement from the CITICORP REPORTS/1976: "With our innovative approach to computer and communications technology, we

* I spoke to the manager of one of these small minicomputer based computer centers. His feelings about the whole issue of decentralization were very positive. It seems that this attitude is shared by most of the people in similar positions.

are rolling the clock backward and forward at the same time: backward to the era of personalized service and forward to a new era of speed, efficiency and quality in transaction processing."

COMMUNICATION NETWORK

Given the geographical spread of Citicorp, communication represents a real problem. The communications system used by the bank is very complex and has been growing in a somewhat disorganized way, similar to the growth of the centralized computer facility. This communication system supports primarily voice communication; computer communication is basically nonexistent. In 1975, the communication costs were 23 million dollars. The bank projects that use of satellites and packet switching could reduce the communication costs by over 50%. Previously, the communications facilities were separate from the centralized computing services. However, now with the redistribution of the computing power, the communication system should be integrated with the existing computer installations.

In 1975, the communications management became centralized and started on a program, the goal of which was to build a global network with modern message switching capabilities and security features suitable to operate in a bank environment. The Citibank's international communications network is called Globecom and it services 80 cities around the world. It has a voice grade spine that spans the world, servicing 22 cities directly, additional cities through switching centers: London serving Europe, Manama serving Middle East, Hong Kong serving Far East, and New York connecting East and West and serving 20 cities in South America and Africa. The heart of the network is Citiswitch in New York City that links the bank's internal offices with the outside world. The Citiswitch consists of two pairs of minicomputers that are linked

together. One pair (two PDP 11/45) is the actual switch that handles services such as message routing, message storage, and message retrieval; it can currently handle over 125 lines. The other pair of computers (PDP 11/30) is called Medic; this system performs message repair, message annotation, and maintains security.

A large portion of the communications within the bank is in the form of typed documents. Citibank is now designing a prototype of an electronic mail system that would further reduce the cost of communications. Citibank has installed its own computer controlled private telephone branch exchange for its New York offices; the electronic mail system is to be built on the top of it. The electronic mail system will provide facilities for editing and formatting documents, including facilities for storage, retrieval and automatic forwarding of the messages. In addition to messages reflecting day-to-day operations of individual sectors of the bank, the mail system will handle delivery of monthly management reports.

COMPUTERIZED SERVICES IN INDIVIDUAL BANKING GROUPS

Let us get back to the idea of a channel as an independent organizational unit. International Banking Group designed a model of a channel with the following characteristics. The channel has an on-line minicomputer with its own local data base. The input to the channel (i.e., the organizational unit) is, if possible, in the form acceptable directly by a computer. In addition to providing services to the IBG customers,the computer system in the channel supports automated data capture for corporate accounting and the management information system. The IBG apparently tries to use the same hardware-software system in all of its channels. They use PDP 11/70, with some specialized software that supports their particular needs such that the

managers of individual channels have to worry only about applications programs. One of their systems called MARTI is responsible for international money transfer and it currently handles 5.9 billion dollars annually. Another system called DARTS is used for foreign currency money transfers with 300 million dollars per day volume. This latter system incorporates "error-repellant processing"; it would be interesting to find out what it involves. So far, the decentralization trend seems to be confined to New York operations only. Large centralized data centers still exist in major overseas locations; England, Germany, Brazil, Japan. These centers are primarily based on an IBM system 3's, 360's and 370's. Smaller locations have NCR Ledger card systems. The plan is to extend the concept used in New York into overseas branches, that is, have services in branches and customer locations built around minicomputer work station systems. The electronic output from these work stations will be tied to country and regional computers and switches. using the Globecom facilities.

The World Corporation Group serves multi-national companies. This market is rather small. It involves approximately 450 companies, but each of these companies typically has a number of subsidiaries. The companies served by WCG have significant operations in five or more countries and annual sales volume in excess of 500 million dollars. Because of the size of these multi-national companies and the complexity of their transactions, the bank has to provide personalized services. Within this group the number of transactions may be small, but these transactions typically involve high dollar volume. While modern technology is needed, significant amount of human judgement will remain necessary in this business. The business in the World Corporation Group seem to require that a single clerk be able to handle several operations rather

than having different operations handled by different clerks. For large customers it may be desirable to dedicate a service center to each customer. Customer's access to the banking services should be possible directly via computer terminals. For multi-national customer, the bank would actually provide communicaiton for the customer's own subsidiaries using the global communication network as described before.

The National Banking Group provides a wide range of financial services for U.S.-based corporations, except the multinational accounts handled by the World Corporation Group and some specialized accounts. NBG is trying to improve customer services through various computer based products, such as Auto Lockbox and a computerized Letter of Credit system. NBG is also developing an electronic funds transfer service called Paytronics. Paytronics is an enhancement of existing Automated Clearing House (ACH) services. Currently there are about 15 automated clearing houses in operation. However, these clearing houses operate on a regional basis only.* Paytronics is intended to cover multiple ACH regions, thus freeing the participating companies from having to make multiple banking arrangements on their own.

This summarizes the major computer oriented developments in the three major groups of the bank, the International Banking Group, the World Corporation Group, and the National Banking Group. The Citibank also has to connect to a variety of the external agencies, in addition to the clearing house. An example is the federal reserve, and a private organization of major banks called Bankwire. The bank is also working on a system that would

---

* It is interesting to note how extensively the automated clearing houses are currently used. It is estimated that by the end of 1977, 85% of all U.S. households with checking accounts will be tied (through their banks) to an automated clearing house.

provide direct communication between the bank and a special class of customers, the brokerage houses.

PLANS FOR A DISTRIBUTED SYSTEM

As said earlier, the first phase of getting away from a large centralized system was decentralization of computing services. The next phase is to build a distributed system, that is, a system where individual computer-based stations can communicate directly. Decentralization was a major effort requiring redesigning most of the software. However, building a distributed system does not seem to be a lesser effort than the first phase, primarily because individual stations use different computer hardware and on the top of this heterogeneous hardware software has evolved in uncontrolled ways.

The first step towards building a distributed system is to provide computer communication facilities. First National City Bank made a decision to use existing technology, specifically the ARPANET technology to develop a communication network that will connect their computers. There are some differences however. First, the Citibank decided to use Honeywell Level 6 minicomputers as IMP's, and contracted BBN to modify ARPANET protocols to run of these minicomputers. Initially it was planned to have a translation unit between a host and an IMP in an attempt to minimize the level of modification necessary in each of the existing computers within the bank. However, the designers felt that this would introduce initial complexity and therefore the translation unit approach was abandoned.

The first version of the network will use leased voice grade lines and eventually these lines are planned to be upgraded to 50 kilobit lines. So far, the bank has a charter for New York City only, but it is planned that the

network will extend to the branches abroad, hopefully through a satellite. How this ARPANET-like scheme fits together with the previously described communication facilities, the Globecom and the electronic mail system, is not clear, but the plan is to arrive at an integrated voice and data communications network.

The New York operations of the First National City bank are distributed throughout New York City. However, most of the corporate services are concentrated in just a few buildings, so there will be clusters of minicomputers in the same building. As a reliability consideration, it is planned to have maximum of two hosts per IMP. That means that with 200 minicomputers there will be about 100 IMPs in the system. This approach seems to be especially inefficient within individual buildings. A suggestion that the bank should look into the possibility of building a local network to connect minicomputers within each building was rejected on the basis that none of the local network technologies tried so far has proved itself to be sufficiently reliable and ready to be transferred to an environment such as theirs.

The distributed system has to reflect the structure and the processing needs of the bank. From the point of management and accounting the bank represents a hierarchical system: at the top is the corporate management, the level below being then managements of individual banking groups (sectors), and each banking group then has several channels (service management divisions) that perform different types of services. It must be possible for the managerial and accounting data to flow between the levels of the hierarchy, for example, the corporate management and the management of individual banking groups, but not directly between groups within the same level. However,

transactions needed to support corporate services may require assistance of several computer stations in the same service management division or even several computer stations in different divisions. Thus, on the level of service management divisions, there is a need for distributed processing both within each division and that spans the boundaries of individual divisions, possibly the boundaries of individual banking groups.

Finally, the network must provide an interface to the external environment. Specifically, it must be possible to connect to the federal reserve, to New York City clearing house and to Bankwire, which is a private institution connecting major banks in the U.S.A. A current proposal is to provide a single institutional interface that concentrates communications with any of these external systems; each of the banking groups then connects to this institutional interface.

The Citinet must support two classes of operations: 1) file transfer between computer stations to eliminate manual routing of magnetic tapes and disks, 2) distributed transactions to achieve real time updating of customer and bank files. Thus, two higher level protocols are planned for the network. Specifically, file transmit-receive protocol, which I believe is an adoption of the file transfer protocol used in the ARPANET, and transaction transmit receive protocol. Unfortunately the information available on the second protocol, which seems to be the more interesting of the two, is very skimpy. Basically, it is assumed that a transaction may need to go through several processors. It will be necessary to provide a directory or a set of directories to aid in the routing of the transaction in the network. A Citinet directory will take customer identification and transaction type and provide the physical address of the appropriate computer station.

The two network protocols promise to support exchange of information in two different modes, however, there is one very important problem that has not been addressed. The software in individual computers developed in individual ways, and there exist, for example, five different check processing systems and seven different letter of credit systems. It will be very very difficult to make these different application programs cooperate in a smooth way.

OPERATIONAL STATISTICS

Let us look at some numbers now. The bank processes one to two million transactions per day. Out of this, 40% are corporate transactions and 60% retail transactions. Though the corporate transactions represent only 40% of the total volume, money-wise they represent 70% of the dollar volume.

As said earlier, decentralization involved separation of services so there is a minicomputer dedicated to a service. However, there are special minicomputers dedicated to maintaining user accounts. About 50,000 transactions per day involve two computers: a service computer and the computer that has the account for the customer requesting the service. Crediting of customer accounts is done only once a day, at night. Movement of information between different computers is done mostly via magnetic tapes.

Accounting and management information is gathered on a daily basis. Presently, 19 magnetic tapes are processed daily by IAMIS, Integrated Accounting and Management Information System implemented on PDP 11. IAMIS will be described in more detail later.

Finally, the individual divisions communicate with the institutional interface as described earlier. The transaction rate at this interface is as follows. Several hundred transactions per day go to New York City clearing

house, approximately 25 transactions a day to federal reserve, and the same amount to Bankwire.

INTEGRATED ACCOUNTING AND MANAGEMENT INFORMATION SYSTEM

Though the bank is talking about processing transactions representing customer services in a distributed way, the first application using Citinet is expected to be the Integrated Accounting and Management Information System (IAMIS). IAMIS generates management reports and also provides interactive access to various data bases maintained by the bank such as data bases on the bank customers, bank products and organizational data bases.

The processing of information for IAMIS is done in two stages. The first stage, called Citiproof, is a data capture system that captures and edits information from over 150 channels. Presently the input to Citiproof comes on magnetic tapes. The second stage involves data base creation and maintenance. Information from individual channels processed on nth day appears in (incorporated into) the IAMIS data bases at the beginning of (n+1)st day. The data base creation and maintenance subsystem of IAMIS also generates monthly management reports and supports online inquiries to the data bases. IAMIS supplies management reports both to the top level management, that is, the corporate management, and to the management of individual banking groups. IAMIS processes approximately 10 million bytes per day.

IAMIS is now in the first phase of its evolution, where the data capture, data base creation, and generation of management information reports are centralized.* The second phase involves decentralization of the MIS portion,

---

* IAMIS is implemented presently on a system built around two PDP 11/70 processors running under RSX-11D operating system. One processor is used for data capture, the second for data base creation and maintenance and data retrieval.

that is, use of the information in the data base. And finally, the third phase will decentralize also the data capture and the data base creation. Actually, IAMIS has not completed its first phase yet. So far data capture and processing is done on a monthly basis. Daily data capture and processing will occur later. Also, as Citinet becomes operational, individual channels (service management divisions) will be able to feed IAMIS directly rather than via magnetic tapes. The third phase will eliminate the necessity to ship information into the central utility because the data capture will be located within individual banking groups and possibly individual channels and correspondingly, will eliminate the need to ship reports back to individual groups. However, the reconcilement of transactions that involve several groups will be more difficult in this fully distributed system.

Finally, some information about the IAMIS data bases. The system has files for the 55,000 of banks corporate customers. In this group of customers there exist 20,000 different relationships with a bank. On the total, the customers have 200,000 accounts using 2,000 of the different services provided by the bank. Because the way the software for individual systems and individual channels evolved, there is a great discrepancy in what is used as customer identification, for example. Thus, the central data base of the IAMIS has to maintain up to 19 "alias" names per customer which only adds to the confusion and complexity of the system. IAMIS must protect the data maintained in the data bases from unauthorized access. Accessability is oriented around the customer (the customer or the parent company of the customer must be in the responsibility domicile of the requestor) or the account (the account must be in the earnings domicile of the requestor). While there is a need to share information across the boundaries of individual

banking groups, some information must never cross these lines. These restrictions stem primarily from federal regulations.

To support the fully distributed version of IAMIS, it is of course necessary to have the computer communication network operational. However, a distribution of IAMIS is a change that may necessitate a major rethinking of the data base organization and algorithms for query processing and report generation.

SUMMARY

Some preliminary conclusions can be made on the basis of the previous discussion of the First National City Bank. First, from the point of the data bases maintained by the bank, the system is very static. That is, the structure of the data bases does not change, though the content of individual records in individual files may change quite frequently. The main problem in this type of system seems to be the partitioning of data bases and correspondingly, the implementation of the services such that the number of transactions that have to cross the boundaries of individual channels or even individual banking groups would be minimized. The first approach was to dedicate a computer system to a service or group of services. The current trend seems to be to recombine services such that the set of services used by a group of customers or a single large customer can be provided at one computer site and possibly even give the customer a direct access to the bank services from the customer's premises.

The research being done in our group definitely reflects the distributed processing problems of the bank. One problem is naming of services. In some of the bank's channels there are several systems that provide the same kind of service, yet it is necessary to establish communication explicitely with one

of these systems rather than being able to ask simply for the service and have the system decide which of the processors will provide this service. The generic naming of services would be a very useful feature in this type of environment. Another problem is protection, of course. The bank apparently does provide some encryption in international money transfers but I don't think that the problem of protection has been properly considered in the Citinet. Local network technology and internetworking (especially interconnecting geographically spread local networks) are also important issues.

With the Citinet it sould be possible to process a large amount of transactions within the bank in much shorter time. However, there is still the problem of handling the actual physical paper checks because these paper checks have to be returned to the customers. There is definitely a bottleneck in how much paper can be handled in a day. Currently, the bank handles 1,000,000 checks per day out of which 250,000 are corporate checks. Due to this bottleneck the only natural next step seems to be a complete electronic funds transfer. Finally, Citinet seems to be viewed mainly as a tool needed to support (further automate) the current model of operations in the bank. Unfortunately, only a banking expert could decide whether the advent of distributed processing should lead to redefinition of the model of banking.