The future of remote information processing systems

by M. J. TOBIAS and GRAYCE M. BOOTH

Honeywell Information Systems
Phoenix, Arizona

INTRODUCTION

Remote information processing is the computer community's wave of the future. For both computer users and for computer vendors it represents an enormous opportunity. To the user, it is a powerful tool with great potential to increase his business efficiency and profit. To the vendor, it opens a broad new market. Properly used, it can benefit both.

This paper discusses remote information processing networks in the rapidly growing area of on-line business applications. We have chosen to survey trends in this area, pointing out some of the important considerations which the potential designer or user of such a system cannot afford to overlook.

One aspect of remote processing is the use of multiple computer information networks. Large complex networks, such as ARPA, represent the leading edge of remote system technology. However, the authors feel that the majority of on-line business applications in the 70's will be implemented as individual systems, not as information utilities. Large scale use of shared or public information networks for on-line business applications will not be seen in the decade of the 70's, mostly because of problems of control, security, and privacy.

TRENDS

Much has been written recently about "the changing computer markets of the 1970's" and the industry's "coming of age," and still more in regard to the increasing sophistication of the computer user. Today's users are:

- Demanding system solutions to their business problems.
- Using an information systems approach to the control of their geographically dispersed organizations.
- Insisting that the emphasis be on INFORMATION as well as on computer power.

Basically, then, we can conclude that there is an increasing trend toward business oriented information systems whose characteristics are:

- Use of a data base—often large and possibly distributed among several branch offices.
- Communications orientation for information collection and distribution.
- Use of information processors (computers) for transformation and manipulation of information.

The needs which these users are expressing have of course been present all along. The historical objections to the computer have been:

- It is located far away from the field "where the action is."
- Only computer professionals can talk to it.
- It is very difficult to organize and store data within it.
- Somehow it always seems more suited for scientific than business problems.

Today these problems are gradually being solved, allowing the emergence of effective remote information processing systems for business use.

There are strong indications that the trend toward remote information processing is a significant one, and will continue to be so. Figure 1 shows projected shipments of new build computer equipment (for the industry as a whole) for batch vs on-line use. While the market estimates are relatively firm only through 1975, the trend can be projected with reasonable assurance farther into the future.

Though various market surveys may differ in detail,

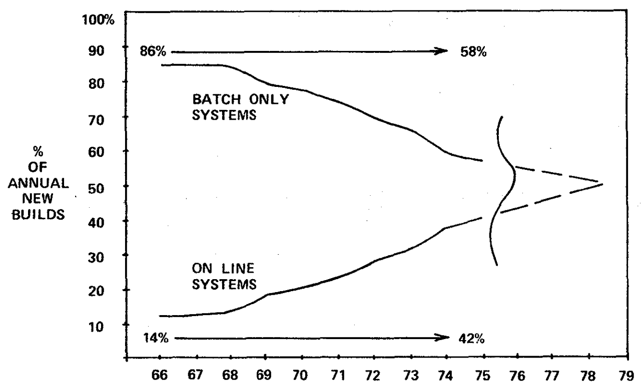

Figure 1—New build shipments forecast

all seem to agree that on-line information systems will be of increasing importance in the decade of the 70's.

APPLICATIONS

Before looking into the future of remote information processing, it may be interesting to briefly survey applications of the past and present. This can indeed be done briefly, since remote information processing is still in an early stage of development. This of course is one of the reasons that its future fascinates us.

Yesterday's applications

Remote processing began in the 1950's. The earliest examples were special purpose systems, and were largely concentrated in a few types of applications.

Airline reservations

In 1952, American Airlines installed the first computerized reservation system, which employed a special purpose computer, Teleregister's Magnetronic Reservisor. Several other airlines installed identical or similar systems. The first break in this pattern occurred in 1958, when Eastern Airlines installed a Univac File Computer, Model 1, marking the first use of a general purpose digital computer for airline reservations.[1]

Command and control

One cannot speak of early remote systems without mentioning SAGE. This on-line command and control system for the U.S. Air Force was planned early in the 1950's, and became operational in 1958.[2]

There were many other pioneering remote information systems. However, in general the cost and complexity of on-line systems deterred their use except when there was no alternative. The airlines, for example, simply could not handle their increasing reservation loads except through automation. Similarly, SAGE represented the only possible means of detecting and reacting to an enemy attack employing missiles and jet aircraft. The absolute need for these systems therefore overrode cost considerations.

From our current vantage point in time, we can see that other applications had a similar serious need for on-line processing, to provide improved efficiency, cost reductions, and competitive advantages. However, these potential advantages were not recognized at the time, so these (to us) very real needs were ignored.

Today's applications

We shall look at today's remote information processing applications more generally than we looked at yesterday's. The accompanying papers in this session describe some interesting examples of today's uses. We will therefore speak generally of the types of applications which are installed, not specifically of the companies or agencies which use them.

On-line banking

Many of the larger banks have already moved into the use of on-line teller terminals to provide instant balance checking and updating.

Order entry

These systems, with the necessary supporting inventory management features, are in fairly wide use today. They are usually applied in the area of wholesale, rather than retail, sales.

Retail sales

Credit checks are being increasingly automated. A number of retail establishments have installed terminals at each cash register, so that credit can be checked prior to every sale.

Point-of-sale automation is being carried farther, in some cases, to encompass not only credit checks but sales slip preparation and even simultaneous inventory updating. One of the pioneering efforts of this type, the TRADAR system jointly developed by General Electric and the J.C. Penney Co., was discontinued, for reasons apparently not directly related to its technical feasibility.[3]

Computer aided instruction

CAI is being used in some schools, but is not yet in widespread use. Student teaching methods using computer-controlled consoles range from quite straightforward drilling to sophisticated interactive tutorials.

Law enforcement

On a small scale, law enforcement systems have many of the same requirements as the armed forces' command and control systems. Government funding is being heavily applied to these systems, especially the FBI's National Crime Information Center network, and the local/regional Law Enforcement Assistance Administration-funded systems.[4]

Network utility

It would be unrealistic to complete any discussion of today's applications without mentioning the ARPA network. While beyond the scope of the more common applications, it has amply proved its feasibility. It is now operational, tying together over 25 computers (as of early 1972) in a research-oriented network. The ARPA design goals are built around providing ½ second access from a terminal anywhere on the net to a host anywhere on the net, regardless of the number of nodes separating them. The great variety of host computers involved—Digital Equipment, Honeywell, IBM, Burroughs, etc.—and the consequent enormous potential for resource sharing make ARPA unique.[5]

Tomorrow's applications

It is now time to gaze into our crystal ball at the applications of the future. In general, we see these as logical outgrowths of today's systems, expanded in scale and scope, but not dramatically different in basic concepts. However, this growth and broadening of scope will enormously increase the computer's influence on society, as on-line systems move more and more into the mainstream of the average person's life.

Retail sales

There will be nearly universal computerization of point-of-sale recording, including instantaneous credit checks. The average retail store or supermarket of the late 1970's will have its cash registers on-line. They may be locally on-line, to the store's private computer, or linked to a regional or national network of many such stores.

This trend will also extend into a more hostile environment, the automobile service station. On-line credit checking, and possibly on-line account updating, will be standard.

Medical systems

On-line hospital management will be widely used. The computer will also be used for the storage of medical histories, and to aid in diagnosis.

Increased population mobility makes centralized medical records almost a necessity, and this need will be met via remote information processing. The ever-increasing pace of developments in medicine makes it impossible for any doctor to remain abreast of current trends, even in his area of specialization. An information processing system, however, can be kept up-to-date, so that all current and/or obscure diagnostic information is available to every doctor.

Law enforcement

Increasing crime rates will provide the incentive for further use of computers. The criminal, like the population in general, is increasingly mobile, forcing improved intercommunication between law enforcement agencies. The local/regional networks now being developed will gradually merge, creating a national network tying together all levels of law enforcement agencies.

Financial utility

Banks will continue to expand their use of remote information processing. On-line banking may merge (although perhaps not within this decade) with other systems to form the often discussed "financial utility."

The financial utility will be the means of creating the "checkless society," or even the "moneyless society." Each citizen will have a "money card," which he will use for all purchases. When the card is inserted in an appropriate terminal, the buyer's credit will be checked and, if good, his bank account will be reduced and the store's account increased.

This application will require a network of computers similar to that in ARPA. Certain technical problems (mostly of scale) are involved; however, these may well not be the limiting factor. Human reaction to this form of money management may be sufficiently negative to delay or limit its acceptance. In some form, and at some time, however, it is bound to come.

Computer utility

For our final example of tomorrow's remote information processing applications, the computer utility seems the obvious choice. Much has been written and spoken about the computer utility, and it is difficult or impossible to evaluate many of the predictions.

In speaking of the computer utility, we mean the very general, highly reliable, information network which provides data, message, and computing services to an extremely wide spectrum of society. We also mean systems utilizing large and flexible information processors to provide low cost computing to the average person—the home owner, the housewife, the student, and so on—as well as to business and industry.

In many respects the computer utility represents the "goal" of remote information processing. When the day arrives on which the computer utility is both feasible and practical, all of the problems discussed during this session will have been solved. The most important of these problems will relate to the acceptance of the utility by a society confident of the reliability and data integrity and security of the total system. When (or possibly if) this day arrives, the computer will indeed have had an ineradicable effect on society.[6]

BUILDING AN INFORMATION NETWORK

The implementor of any remote information processing system must realize that he is faced with a networking problem. Every remotely oriented system requires a network, where a network is defined as a set of hardware and software elements which provides for the collection, processing, and distribution of information.

The problem of creating an information network with optimum characteristics is logically the same regardless of scale. A small information processor surrounded by a few terminals forms a network. So does a nationwide hookup of half a dozen geographically separated information processors, tied together via communications lines, interfacing with hundreds of terminals, as

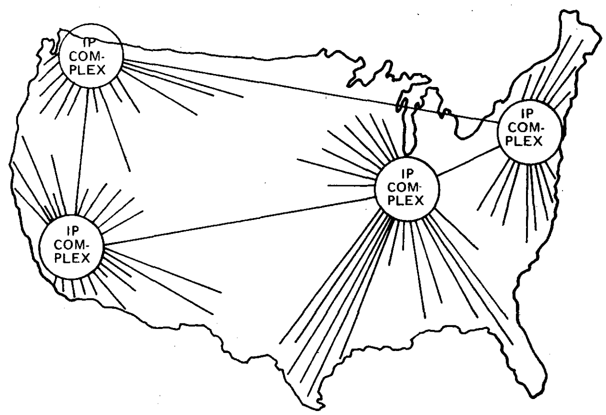

Figure 2—Large information network

shown in Figure 2. The major difference is in scale, not in the techniques of networking.

Before discussing, in the following sections, the hardware and software elements needed to form an information network, it is important to point out the characteristics shared by all remote information processing systems.

Flexibility

As a company becomes increasingly involved with remote information processing, the business becomes closely intertwined with the system. Thus, the information system must be as flexible and capable of growth as the business itself.

Availability and integrity

A company which implements a remote information processing system (for purposes other than simple time sharing) is committing an important part of its communication and control structure to that system. When a company commits its business (or part of it) to the information processing system, its dependence on that system becomes very great. High availability, or continued operation of the system, is therefore required. The user must carefully evaluate the degree of availability needed; this is an expensive commodity, but an absolutely essential one.

Provisions for protecting data integrity are also essential. The heart of any remote information processing system is its data base. These large data bases are often the only repository of vital business data. It is not only impractical but impossible to maintain this information in any manual form. It is therefore absolutely essential that this data be protected against loss or damage.

Optimum cost/performance ratio

Remote information processing systems are usually more expensive than batch systems, but many are more costly than necessary. There are many techniques for reducing costs, especially communications facilities costs, while retaining desired performance. Some of the commonly used cost reduction techniques, such as multiplexing, concentration, and line switching, are described in the following section on Network Hardware Elements. These must be thoroughly investigated by the user who wishes to install an effective remote information processing system.

NETWORK HARDWARE ELEMENTS

An information network is formed of a number of hardware elements, which can be classified as follows:

- Terminals
- Distribution facilities—lines, trunks
- Network (communications) processors
- Information processors
- Multiplexors and/or concentrators
- Switching devices

Some of the more important features of each of these elements are discussed in the following sections.

Terminals

The proper choice of terminals is important to the success of any information network. Great flexibility in terminal design is important, since the terminal is the part of the system closest to the user and with the most need to adjust to his individual requirements.

Many categories of terminals will be used in the 1970's, including teletypewriters, CRT displays, remote batch, data preparation, and graphics. The two most important types, however, will be displays and application specialized terminals.

Displays will normally operate on voice grade lines, at line speeds up to 9600 or 10800 bps (the higher speeds will be needed to satisfactorily perform screen filling/emptying operations). Displays will most often be used on multi-drop lines, in poll and select mode, with occasional use on single station private or dial lines when heavy utilization makes multi-drop lines impractical.

Application specialized terminals will include teller sets for banking, factory data collection devices, stock broker sets, airline reservation sets, and other similar devices.

These specialized terminals will be heavily used in transaction processing systems. The terminal operator will, in many cases, use the terminal only part-time, and will therefore not be an expert operator. His lack of specialization as a terminal operator will often be compensated for by terminal specialization.

Distribution facilities

The distribution network is the means by which information flows between nodes in the network. Distribution facilities are provided by the common carriers—AT&T, Western Union, etc.—or potentially by the newer carriers such as MCI and Datran.

Digital networks will represent a significant advancement in the time frame beyond 1975. Both AT&T and the specialized carriers are beginning to install such facilities. The advantages of such networks include:

- Emphasis on digital traffic
- Improved switched line facility
- Low error rate
- High availability
- Significantly reduced cost

The reduction of user cost is a primary advantage of such networks. In the Datran network, for example, regional transmission (i.e., transmission beyond the local trunk) will be independent of distance. The charges proposed are much lower than current rates, especially at the higher transmission speeds of 9600 or 14400 bps.

It is concluded, then, that distribution facilities are becoming more flexible, more reliable, and more economical. This trend should continue as competition strengthens in this field. The resulting improvements will be manyfold for remote information processing.

Network processor

A network processor is a computer specifically designed for the control of data communications. Network processors are vital to the success of remote information processing systems. A network processor can be used either as a front-end to an information processor, or as a free-standing system performing functions such as complex concentration and/or message switching.

The important reasons for the use of a front-end network processor are:

- Off-loading of network oriented functions from the information processor.
- Greater flexibility and versatility.
- Increased system availability.

It has been found in practice that the network processor design should emphasize flexibility and memory capacity, rather than the maximum in processing speed. Its primary tasks are those of data manipulation and transliteration, not "number crunching." However, the instruction set of the network processor must necessarily lend itself to efficient character and byte data handling. If this is not the case, processing speed may indeed become a limiting factor. It has also been found advantageous for the front-end network processor's instruction set to be closely related to that of the associated information processor.

Some of the important design requirements of a network processor are these:

- Memory size up to 256K bytes.
- Fast access disk interface (the disk space in front-end applications should ideally be shared with the information processor in the "delta" configura- shown in Figure 3).
- Optional console, card reader, and printer interfaces.
- Optional non-interruptable power source.
- Ability to be configured with full redundancy.
- Program controlled reconfiguration and communications line switching.
- Interfaces to narrow band, voice band, and broadband lines, both dedicated and dial-up.
- Ability to support several hundred concurrent user messages—possibly as high as 1000 in 1975 and beyond.
- Floating dial-out channels.
- Integrated modems.
- Ability to interface with other network elements such as multiplexors, line switches, etc.
- Ability to interface with more than one information processor.

The above listed requirements are probably not fully met by any network processor design in existence today. Also, the rather large load requirement listed is representative of only a very few present day applications. However, the trend toward large loads and very high availability requirements is unmistakable. Thus,

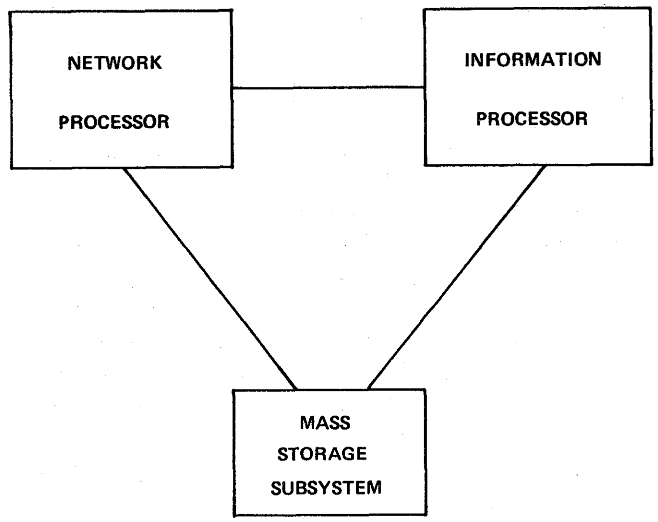

Figure 3—"Delta" configuration

these requirements are not blue-sky—they are for tomorrow's applications.

Two aspects of the network processor requirements are sufficiently interesting to warrant further discussion.

The "small" communications oriented user requires the same degree of sophistication and functionality as the large user. Thus, the growth of the front-end network processor function on site is largely related to load handling capability. This makes the network processor design even more difficult, in order to provide for the "small" load economically yet grow to the very large loads as well.

Though the communications front-end processor has, in the past, been thought of as merely an instrument of the information processor, it is now clear that it must function quite independently of the information processor. For example, the front-end processor must remain operational even when its information processor fails. In this way contact with the rest of the network is not lost, input messages are still received, output messages can still be sent, and—most important—full service can be brought back with the minimum of difficulty and time delay after information processor repair is effected. Front-end processor independence is also needed as this processor takes on more and more tasks related not to information processor interface, but to the network (i.e., message switching).

Information processor

The information processor is of course where the actual application processing takes place. Perhaps the most important hardware feature of an information processor to be used for remote processing is modularity.

In a remote system, the information processor should be capable of at least fail soft, and for some applications fail safe, operation. Fail soft is aided by hardware modularity, with the possibility of configuring multiple processors, multiple memory modules, multiple I/O connections, and so on. This modularity must be such that failure of any hardware element, such as a memory module, does not prevent continued operation of the remainder of the system.

Multiplexors/concentrators

Both multiplexors and concentrators can be used to greatly decrease communications line costs. Multiplexors are passive, non-programmable devices, while concentrators are programmable.

A multiplexor splits up a single communications line so that it can carry a number of multiplexed messages concurrently. There are two types of multiplexors, Time Division (TDM) and Frequency Division (FDM).

A wide range of manufacturers offer TDM's and FDM's. General Electric, Timeplex, Computer Transmission Corp. (Multitran TDM), and GTE Information Systems are only a few of these. AT&T also offers multiplexors.

A concentrator is a form of network processor, whose function is to compress data on many lower speed lines onto fewer higher speed lines. Because it is programmable, a concentrator is more flexible and potentially more powerful than a multiplexor. As with other network processors, it is also capable of assuming other network oriented tasks.

The Honeywell Model 730/50 is a typical concentrator built around a minicomputer. It concentrates up to 128 low speed lines (up to 300 baud) or up to 64 medium speed (up to 2400 baud) onto 1 to 4 medium speed lines (up to 10,800 baud) leading to a remote network processor or information processor.

Switching devices

Reconfiguration line switches are used in many networks. Their most common use is to switch communications lines between two or more network processors.

Figure 4 shows a typical fail safe system. The reconfiguration switch shown might be the Honeywell Line Transfer Device 355, which can switch up to 96 lines under either manual or program control. The diagram also shows a "Deadman Timer," such as the Honeywell Computer Monitor Adapter, typically used for automatic failure detection in configurations of this type.

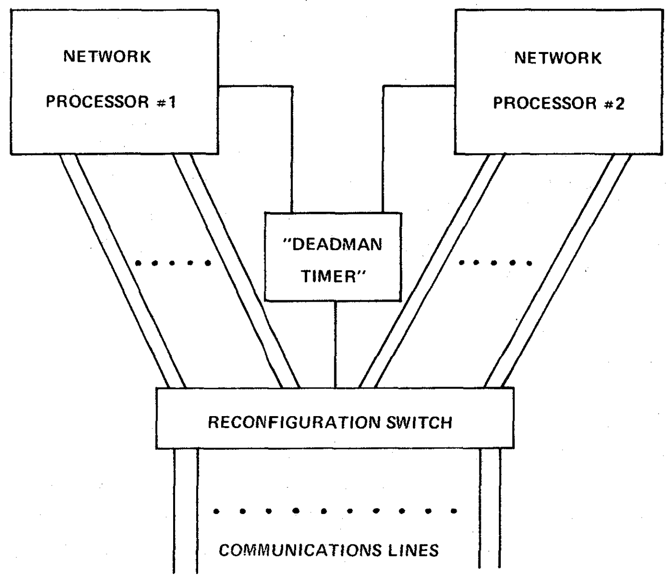

Figure 4—Fail safe configuration

SUPPORTING NETWORK SOFTWARE

The implementor of an information network must also look closely at software requirements. It is not possible, in a paper of this length, to discuss—or even list—all of the functions needed. Therefore, only a few of the more important are included here.

Terminal interface

The system interface seen by the terminal operator must be appropriate to the application, and to the operator's level of training. Particular care must be taken to support the type of operator who uses the terminal, not as a full time task, but as a subsidiary part of his job.

For these cases, the language used by the terminal operator must be simple and self-explanatory. Clear diagnostics must be provided in case of error, and software assistance must be available when the operator is presented with unusual conditions.

Load leveling

Some networks of multiple information processors may require a load leveling function. This allows jobs to be distributed between the processors to even out work loads. If the processors have unequal capacity, or are of different types, load leveling can also detect jobs too large to be run on a smaller system, or jobs which require a specific type of processor, and direct them to the appropriate location for execution. As large distributed information systems become more widely used, load leveling in various forms will be increasingly necessary. However, vendor-supplied software for this function is rare today (if not non-existent).

Distributed data base

As networks with distributed computing capabilities become prevalent, distributed data bases will naturally develop. A distributed data base exists when two or more information processors which communicate via the same network allow creation and sharing of permanent files.

Without going into the complexities of distributed data bases—which are many—it can be pointed out that a single job may require access to more than one part of a distributed data base. This can be achieved either by transmitting file data or by transmitting jobs or tasks. Both of these methods cause increased communications load, and may require very high speed transmission.

Message management

Applications which process remote input and output have, in general, a more complex task than those which handle local I/O. Ideally, an application program should be able to treat a message from a CRT identically with one from a card reader, and be able to send output identically whether to a CRT or to an on-line printer. In practice, it is seldom this simple.

One approach to this problem has been provided in the COBOL communications-oriented verbs RECEIVE and SEND. The approach is an unfortunate one, in the authors' opinions, since it would be more consistent to use READ and WRITE. However, CODASYL has chosen to use different verbs, so the COBOL community must accept this.[7]

One pitfall which the user must watch for in a vendor's implementation of these verbs is whether or not device independence is provided. CODASYL has not specified whether terminal characteristics are to be handled by system software or by the user program. The former, except in very rare cases, is much preferable.

For example, the verb RECEIVE should cause a line (or other logical increment) of input to be presented to the program, stripped of terminal control characters and blank filled if necessary (for short lines or imbedded horizontal tabs). If this is not done, a great deal of tedious—and sometimes complex—processing is required in the application program. Output should similarly be prepared as a line or page by the program, with all necessary terminal control provided by the system.

Although it makes no difference to the application where (outside of its boundaries) terminal-specific processing takes place, this is important to the system as a whole. Terminal characteristics should be handled entirely in the network processor, allowing the information processor to ignore these characteristics. This allows much greater flexibility to change/add terminal types, and/or to redirect messages to alternate destinations.

In the context of remote input/output, a software generator approach to creating the coding to handle terminal specifics is an absolute necessity. It is unrealistic to expect any vendor to provide terminal handling software for every terminal on the market. Instead, the use of a generator package can make it relatively simple to add types of terminals to the network. This is the only realistic way to approach the variability inherent in remote terminals.

Data base management

The designer of a remote information processing system must look very closely at data base processing, since this will be the heart of the system. A data base management software system of considerable sophistication is normally required.

The most important point to be considered, which is sometimes overlooked, is the need for concurrent access to the data base. If any of the on-line applications update the data base, there is a potential problem of access conflicts.

Even with update access, these problems can be minimized if each application can be structured to read only one record and hold it the minimum possible time before updating and releasing it. In this comparatively simple case a software system for record locking can prevent any other access to a record being updated.

More complex situations occur if one or more applications must update multiple records. Locking records in this case can cause the "deadly embrace" situation shown in Figure 5. The data base management software should detect and resolve these conflicts.

To improve throughput, locking may be used only for updating programs, with inquiries allowed to run without check. This approach can cause the inquiry mode application to see inconsistent data base status, as shown in Figure 6. Programs must be written to take this into account; otherwise they may malfunction when a "shifting data base" condition arises.

There is a considerable difference in complexity between the simplest case of concurrent update (one record read and updated by each program) and the most complex (many concurrent applications each updating many records).

Whenever possible, on-line applications should be structured to achieve the simpler case. One way to do this is to perform the minimum updating necessary on-line. For example, an order entry system might be set up to update credit limits, inventory balance, and production and shipment schedules. Analysis might show that update of the production and shipment schedules could be postponed for batch handling, either once daily (usually at night) or in slack on-line periods. This would significantly decrease the amount of on-line updating, therefore decreasing the overhead involved in locking records and monitoring for interference.

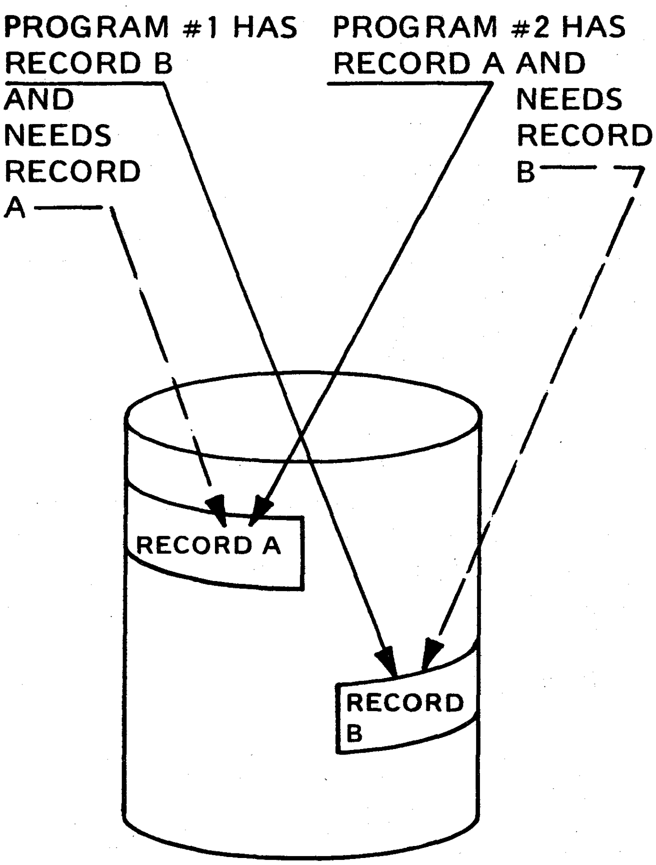

Figure 5—"Deadly embrace"

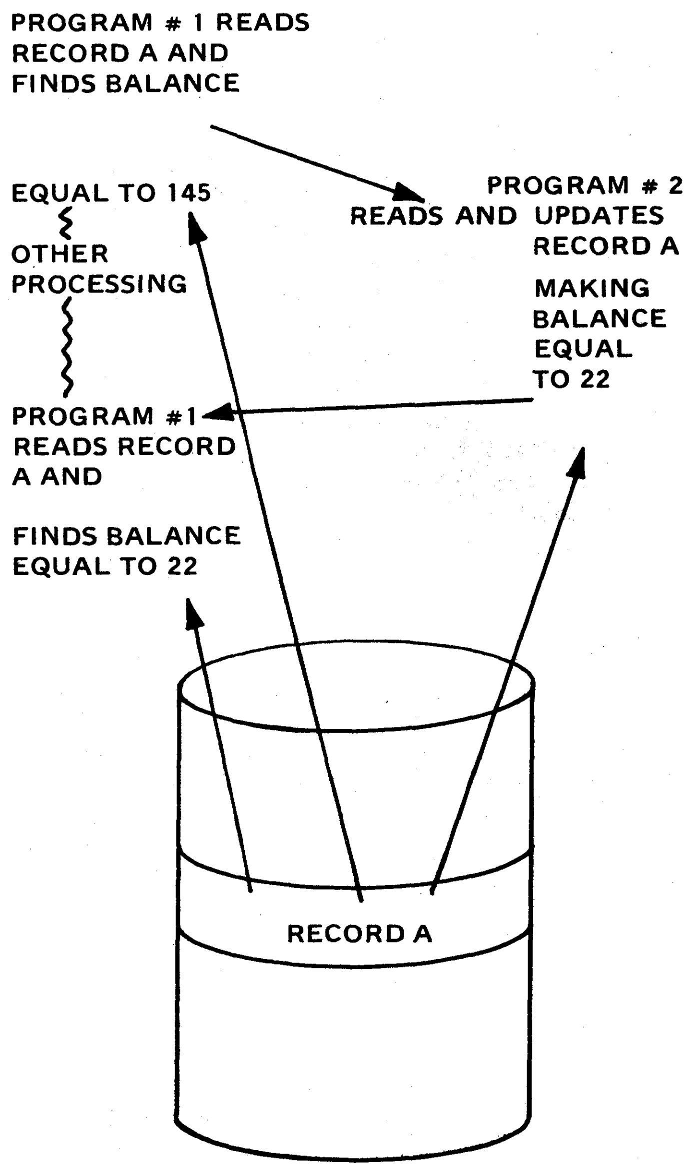

Figure 6—Shifting data base

This area should be extensively studied during system design. Unnecessarily complicated on-line update applications can be very costly in terms of software complexity, debugging difficulties, overhead, and in operational delays associated with locking and resolving access conflicts.

Integrity protection

Integrity is a broad heading, covering a number of hardware and software features. Both hardware and software must be properly oriented toward achieving adequate integrity.

Fail soft or fail safe configurations must provide graceful degradation when a non-essential component fails. Restart/recovery features must be provided, in case of total system failure.

Data protection is extremely important, so that data base information and incoming/outgoing messages will not be lost, duplicated, or garbled. Features such as file change journaling and automatic recovery are especially important, to protect against file damage in case of application program or system failure.

Test and diagnostic routines must be available, to test communications components, network and information processor and their components while the remainder of the system is processing normally. Off-line T&D is also necessary; however use of off-line T&D will generally be minimal, as most error diagnosis will be required while the system is on-line.

Security

Security (including privacy protection) features must include, as a minimum, terminal and user validation procedures independent of the information processor (because of terminal/terminal exchanges in some systems). Detailed user validation, including password and permission checks, may also be performed in the information processor. More complex security features, such as data encoding, may be required in some networks.

Debug aids

It is important that adequate debugging aids be available and specifically applicable to the on-line data base and business processing environment. Too often, on-line software systems are designed for production use only, forgetting that both initially and on a continuing basis system/program changes will be required. The prudent user will avoid this trap by insisting on adequate test features.

In testing on-line applications, it is very helpful to be able to easily load and test an on-line application from any convenient location, perhaps a teletypewriter, perhaps a card reader. To allow this, the software system must provide device independent I/O, and also the ability to obtain snapshots and traces interactively during test execution. Interactive debugging, when properly supported, can enormously improve programmer productivity.

Testing must often take place concurrently with production work. This is particularly true in systems which process actively for all or most of the 24 hour day. In these systems, all debug aids must be structured to allow concurrent test and production operations.

Finally, aids are needed to allow testing of data base update programs. The best aids allow test programs to be run, simulating update, without actually altering any file data. If test aids are not available, the user will find that he must manually set up test data bases, and quite likely also alter his programs so that they can use these data bases.

CONCLUSION

Remote information processing systems are technically feasible today. The most complex forms of such systems, information networks, are being implemented in small—but growing—numbers.[7] The ARPA network referred to earlier is an outstanding example of these.

These networks represent the leading edge of technology in at least some respects. They have proven the feasibility of distributed information systems, but not necessarily the general applicability of this approach. Most of the remote information processing of the 70's will still be in the form of individual on-line business systems. In many cases these will be networks which include multiple information processors, but single user owned. The mass use of public networks to solve business problems on-line will not be seen for some time, largely for reasons of security, privacy, and control.

Much work remains to be done to improve software technology, particularly in the areas of message management, data base management, and integrity protection. Innovative users, perhaps more than the hardware/software vendors, are leading the attack on these problems. The papers which follow describe some of the systems which embody today's state of the art in remote information processing systems.

ACKNOWLEDGMENT

The authors would like to acknowledge the contribution of Hal B. Becker, from whose work on structuring and organizing Information Networking we have borrowed freely.

REFERENCES

1 R A MC AVOY
Reservations communications utilizing a general purpose digital computer
Proceedings of the Eastern Joint Computer Conference
The Institute of Radio Engineers Inc NY NY 1957

2 R R EVERETT C A ZRAKET
H D BENINGTON

SAGE—A data-processing system for air defense
Proceedings of the Eastern Joint Computer Conference
The Institute of Radio Engineers Inc NY NY 1957

3 R M PETERSEN
TRADAR: Death of a retailer's dream
Datamation Vol 17 No 11 June 1 1971

4 P HIRSCH
LEAA: Who guards the guardians?
Datamation Vol 17 No 12 June 15 1971

5 L G ROBERTS B D WESSLER
Computer network development to achieve resource sharing
AFIPS Conference Proceedings SJCC May 1970

6 D F PARKHILL
The challenge of the computer utility
Addison-Wesley Publishing Company Reading Mass 1966

7 *First report of the Communications Task Group to the CODASYL programming language committee on the COBOL extensions to handle communications processing*
COBOL PLC Item No 68114 May 1969

8 D J FARBER
Networks: An introduction
Datamation Vol 18 No 4 April 1972