

USING MINIS AND MICROS

Last month we discussed why some users would rather have their own computers than share the services provided on a large central computer. This month we look at the other side of the coin—why some information systems departments would rather use mini-computers and micro-computers to perform some applications. In most cases, existing work is not being off-loaded from mainframes; instead, it is new work that is being put onto the smaller machines. Even some companies that are strongly oriented toward mainframes are finding that small machines can be useful. Here is what we found happening in this fast-changing arena of mini- and micro-computers.

Metropolitan Life Insurance Company, with headquarters in New York City, is the second largest life insurance company in the U.S., according to *Fortune* magazine. It offers personal life and health, group life and health, property, and casualty insurance as well as annuity and pension plans. It has assets of over \$50 billion and employs about 48,000 people.

Metropolitan Life has some 20 large IBM mainframes, in five data centers, plus over 100 mini-computers—used mainly for data communication, remote printing, and remote summary reporting at their numerous smaller regional data centers.

For the past three years, Metropolitan Life has been exploring ways to provide more end-user computing. They are exper-

imenting with IBM's Information Center concept (which we discussed in June 1981) and they are testing a distributed system using IBM's VM (virtual machine) operating system in a network environment. The VM operating system is used on two mainframes at the New York City data center, an IBM 370/168-AP ('attached processor'—IBM's parlance for two processors hooked together) and a 3033-AP.

In addition, last year the company purchased a 'plug-compatible mini-computer' also to run VM. The computer is an F4000 from Formation, Inc., of Mt. Laurel, New Jersey. The machine has two megabytes of memory, three 3350-type disk drives, a tape drive, and some Formation and IBM CRT terminals. The cost of the system was \$160,000.

The F4000 is basically equivalent to an IBM 4331-1 in processing power, with an estimated 0.2 MIPS (million instructions per second) rating. In addition, it uses mini-computer peripherals rather than mainframe ones.

The information systems department at Metropolitan Life (the electronic installations department), had several reasons for choosing the Formation machine rather than a traditional mainframe.

First, they needed a machine that used the IBM 370 instruction set and could run VM. They wanted this feature because they are planning to create a network of 'operator-less' VM machines.

Second, they wanted to minimize equipment size and cost. The F4000 is being used to experiment with office automation and 'user-controlled computing'—where a user department manages its own computing resources. This lower-cost system, which could be placed in a user department, made such experimentation feasible.

And third, they chose Formation because this vendor is interested in working with them on this test. Formation has participated in Metropolitan's planning sessions and has even been willing to make hardware modifications, if necessary. This 'value added' co-operation from Formation has been important to Metropolitan in their research and development work.

Metropolitan Life's first F4000 was uncrated on a Thursday afternoon last October. The user disks were copied from the IBM 370/168-AP, and the system was and up and running the next day—without anyone working overtime Thursday night. The biggest installation problem was making sure that the floor under the disk drive was level. The company installed a few wall partitions around the computer, so the area can be locked at night; further, no special flooring or air conditioning were needed.

As mentioned, the Formation machine is intended to operate in an operator-less VM environment. Toward this end, Metropolitan Life is also working with Spartacus Computers, of Burlington, Massachusetts, a firm doing development work in this operator-less computer area.

The F4000 has been installed in the pension services department, where the users only need to change printer paper and mount tapes to op-

erate the machine. However, the system has not reached the fully operator-less stage, because: (1) the company is not yet able to start it up each morning from the 370/168 at the data center (they are waiting for the software from IBM on this), and (2) sufficient system data cannot yet be collected by the mainframe. Before the company can create a network of operator-less machines, these two problems will need to be solved.

Metropolitan Life foresees selected applications being developed centrally and then downloaded to the appropriate VM machines on the network. Such a network could provide cost-effective computing, because there would be no need for information systems staffs at the machines. Also, the network would provide back-up facilities.

The first major application that was run on the F4000 was a financial management system to assist actuaries and pension personnel in studying the cash flow of existing and proposed pension accounts. This decision support system was written on the 370/168 and then transferred to the F4000 over a 4800 bps line. Since the application runs under VM, there was no need to worry about smaller memory size, we were told. Their major concern in choosing a mini-computer was having less overall capacity than on a mainframe, but the other positive features of the system have diminished this concern. In fact, they feel that the F4000 is providing them with a very cost-effective system for their pioneering efforts.

The options in smaller machines

Last month we looked at the use of smaller machines from the users' point of view. This month we look at the other side of the issue—why some information systems departments are also putting in mini- and micro-computers.

We believe that most information systems managers would rather install mainframes, if at all possible, because they are easier to control and maintain. They are compatible with the company's existing machines. Mainframes reside in a computer center that can be more easily secured physically. Maintenance people can be located centrally. Remote access to the system can generally be controlled. And standards can be

created and more easily enforced with respect to programming methodologies and conventions, data, and communications.

However, a growing number of data processing departments are installing smaller machines. In this issue we describe several circumstances where smaller machines may be more appropriate than mainframes. But first we begin by looking at the various types of smaller machines available today; the choices are really quite varied. The following discussion of micros is based on the new edition of our book, *So You Are Thinking About a Small Business Computer* (Reference 1).

Single-user micros. At the 'smallest end' are the 8-bit single-user micro-computers. These 'business micros' first appeared on the marketplace in the late 1970s. They usually have from 32K to 64K bytes of main memory and floppy disk storage. The floppy disks are 8-inch, 5-1/4 inch, or (coming) 3-1/2 inch diameter sizes, and store from 92K to over one megabyte of data per disk. Examples of these systems are the Apple II, Radio Shack TRS-80 Models II and III, AVL Eagle II and III, and Xerox 820.

These single-user micros are generally for individual employee use; however, some companies are installing them for a specialized application, such as business graphics, where they serve a group or a department.

More recently, single-user 16-bit micros have appeared on the marketplace. There are three families of 16-bit micro-processor chips: the Intel 8086/8088 (the 8088 is used in the IBM Personal Computer), the Motorola MC68000 (used in the TRS-80 Model 16), and the Zilog Z8000 (used in the Onyx System; with the UNIX operating system, this becomes a multi-user system). Some of these systems can have up to 256K bytes of memory, while others offer one megabyte or more. They often use a hard disk which can store 5 to 10 megabytes of data or more. Floppy disk storage may also be used, for main storage or for loading and unloading files.

New developments in this area seem to be 'happening daily.' A few examples are: (1) portable systems, about the size of a briefcase, (2) dual processor systems, with 8-bit and 16-bit processors in one system, to ease upward migra-

tion, (3) high resolution graphics systems, and (4) concurrent operating systems, allowing users to perform several tasks at one time.

Multi-user micros. Another recent option is 8-bit or 16-bit micro-computers with multi-user operating systems. These systems generally have from 128K to 1 megabyte of main memory and usually come with a hard disk system, with capacities in the 10 to 40 megabyte range. The 8-bit micros can support from 2 to (typically) 5 terminals, using multiple 64K byte blocks of memory, while the 16-bit machines can support from 3 to 24 terminals. These systems appear appropriate for use in a department or at a remote site, such as a sales office or warehouse. Examples of the 8-bit systems are the Cromemco Z2 and Vector Graphic 2600, while 16-bit systems are the Onyx C8002, Fortune Systems 32:16, NEC 200 series, and Zilog System 8000. Micros with 32-bit word lengths are just beginning to appear on the marketplace.

Network of micros. Another fairly recent option is the micro-computer network. In the current offerings, single-user micros are connected via a cable (which can range from a twisted pair to coaxial cable). Some of the micros serve as work stations while others are network servers, such as printer servers or file servers (for shared use of hard disk). These networks can grow to be quite large—supporting up to 64 work stations on one network; also, several networks can be inter-connected through 'gateway' servers. Examples of these systems are the TeleVideo and Intertec offerings, which use their own computers, and the Nestar Cluster/One and the Corvus Omninet, which provide networks for other brands of computers.

Another type of communication link is between micros and mainframes. Schindler (Reference 2) reports on four companies that offer such links. Micro-Integration of Friendsville, Maryland, and Winterhalter and Associates of Dexter, Michigan, offer software products for communicating between mainly CP/M-based micros and IBM mainframes. Local Data Inc. of Torrance, California, and Expander Inc. of Pittsburgh, Pennsylvania, offer hardware products for protocol conversion so that micros can communicate with IBM mainframes.

Traditional mini-computers. The next option is the more traditional 16-bit, multi-user mini-computer, such as the DEC PDP-11, Hewlett-Packard 3000, Data General Nova, and Four-Phase System IV. These systems generally have from 128K to 512K bytes of memory and use hard disks to store from 5M to 200M bytes of data. They are often used for interactive business applications, and typically can support up to 32 terminals. They have been used in business both as stand-alone systems—as the only computer in smaller companies and in departments and remote offices of larger companies—and in company networks. Sometimes they perform a specific application, such as being the front-end communications processor for one or more mainframes.

Second-hand mini-computers. An interesting option is 'second-hand' mini-computers and peripherals. As Kincaid (Reference 3) discusses, second-hand minis and their peripherals cost less and often can be acquired more quickly than new equipment—and some may be relatively new. Common reasons for buying second-hand equipment are: for experimentation in 'low budget' projects, to provide backup for new or installed equipment, for small remote operations, and to provide spare parts for installed systems.

Kincaid cautions that only knowledgeable computer buyers should look into second-hand computers, because these machines do not come with the traditional warranties from the manufacturer. *The Computer Economics Report* (Reference 4) recently surveyed computer manufacturers about their policies regarding their used computers. They report that manufacturers are generally willing to bring a used system up to its latest revision level and then put it under a maintenance contract; however, each case is negotiated separately, based on the condition of the used equipment.

Super-minis. A newer option in the mini-computer field is the 32-bit 'super-mini.' These are available from Prime, Perkin-Elmer, DEC, Data General, and others. These multi-user systems have from 500K to 16M bytes of main memory and can provide disk storage of over 2B bytes. They are actually much more powerful than the traditional business mainframes of 10 years ago,

and they are equivalent to the smaller mainframes on the market today. They are generally used for interactive processing and can support up to one hundred or more on-line terminals.

Network of minis. There are also networks of minis. For example, Wang, Datapoint, DEC, and others offer network software to link their machines together to form distributed systems. In these systems, some of the processors may serve as application processors, while others may be dedicated to specific functions, such as data communications or file management. These systems can grow to be quite large—possibly connecting scores of processors and communicating over long distances, such as between Europe and the U.S.

We have 'neatly categorized' these various smaller machine options to mainframes. But, from a capability standpoint, they are not always smaller than mainframes, so the categories are not really as neat as might first appear. Also, the new micro-computer systems are beginning to compete directly with the mini-computers, and even with some mainframes, so that the boundaries between the categories are becoming even more blurred. Some people say that such categories are now meaningless, but no newer classification that we have seen appears any more useful.

Technical issues. Moving work from one computer to another presents two major technical problems of interest here: code conversion and data conversion. Several of the mini-computer vendors, and also some independent companies, offer products to deal with these conversion problems.

One solution is the creation of an IBM plug-compatible mini-computer, such as the F4000 offered by Formation, Inc. Programs and data running under the IBM VM operating system can be moved directly onto the F4000 without the need to convert either.

Another approach is to use conversion products. These provide 'automatic' conversion of all or part of programs or data. Usually these products convert software or data from one specific computer to another specific computer, so they are not generalized conversion aids. Also,

they often do not perform the full conversion task; some effort by programmers is necessary.

A third approach is to completely rewrite everything—programs and data files. This is the least desirable way to distribute work, but sometimes it is the only available option. For this approach, there are also products that can be used as 'conversion aids.' They consist of application generators and data management systems, which allow developers (and perhaps users) to more quickly create the desired programs and define the file formats, based on the design of the old system. These products are appearing for many makes of mini-computers and for some micro-computers. Special programs may still be needed to convert files from old formats to the new ones—although some data management systems provide routines for this also.

When to consider smaller machines

What are the circumstances that lead information systems management to consider using smaller computers? Here are the ones that we have come across in our studies: (1) when the current computer must be replaced, (2) when the applications are specialized, (3) when a remote site is too small to have its own mainframe computer, and (4) when the company wants to try out new systems ideas.

When the current computer must be replaced

There comes a time when every computer *must* be replaced—it just gets 'too old.' But what if there is no easy upwards migration to a new computer? This situation has occurred at some smaller computer sites where the workload has not grown rapidly, so their second-generation computer has been adequate for ten years or more and they have just stayed with it. But finally the old technology used in that computer begins to restrict what management wants to do, so the company is forced to upgrade to a computer that uses newer technology.

Often, floor space limitations, cost constraints, and user requirements make a mini-computer (rather than a mainframe) the most appropriate choice for the new computer. In such cases, the company must either buy application packages

to work on the new machine or completely rewrite their existing applications—and for the latter, they need an economical way to rewrite their software. Consider the case of the IIT Research Institute.

IIT Research Institute

The IIT Research Institute is a not-for-profit corporation involved in applied research projects for industry and government. The Institute, which is based in Chicago, Illinois, has some 1500 employees who are working on projects in air pollution, electronics, life sciences, and numerous other areas.

In 1967 the Institute acquired a Honeywell H-200—at that time, a medium-size mainframe computer—to perform its financial computing. As with most computers of that generation, by 1980, the computer had become relatively expensive to maintain. Even more important, it was difficult to find programmers who knew its programming language (Easycoder, an assembly language). Also, it was just too small (32K bytes of memory) to easily accommodate the Institute's normal size programs. And it used old technology—punched cards and magnetic tape.

So in 1980 the Institute began looking for a new computer. They wanted one that was larger and would support multi-tasking. Also, they needed facilities to help them rapidly convert their data and rewrite their assembly language programs in COBOL. Further, they wanted to provide their users with an end-user query language.

They settled on a Prime 550 mini-computer from Prime Computer of Framingham, Massachusetts. It has 1.25M bytes of memory, two 300 megabyte disk drives, a nine-track tape drive, a 600 line per minute printer, and 25 CRT terminals. This mini-computer is actually much more powerful than the mainframe it replaced. They essentially chose the Prime computer because of two software products available for it—a COBOL program generator from David R. Black and Associates of Northfield, Illinois, and an 'end user programming facility,' called INFO, from Henco, of Wellesley, Massachusetts. We have previously discussed both of these products, in our August and June 1981 issues, respectively.

Since the programmers at the Institute had little (if any) COBOL programming experience, and since they had some 452 programs and 200 data files to convert, management decided to simply rewrite (not redesign) the programs using the program generator. This, they felt, would speed up conversion considerably. Many of the programs were simple data entry and report generation programs, and they were easily re-written using the DRB Program Generator. A programmer would review the old program to find out which data files were involved. Then, using a CRT, he would lay out the data record, the report format, and the data entry formats on the screen. Using these, the generator would create the necessary programs. Many of the programs were re-written in from two hours to two days time (including testing) using this approach.

Also, to speed up conversion, the Institute designed a 'universal record format,' for passing data among the payroll, accounts payable, and general ledger systems. It was also used for passing data from the various project cost accounts to the Institute's general financial systems. Each of the 1000 on-going projects keeps track of accumulated costs using a project accounting system. Data from these files is accumulated for consolidation using this 'universal' record.

Fortunately, converting the data from the seven-track H-200 format to a nine-track format was very easy, and was performed by an outside service bureau. Then a programmer wrote a program to convert that data to the Prime format.

In all, the total conversion of the 452 programs, amounting to 350,000 lines of code, was completed in about one year (with 3 work-years of effort). The major delays involved the users; many of them wanted new and different reports, which often led to re-designing files, screen formats, etc., thus lengthening the conversion process. Also, educating the users to use terminals rather than fill out the familiar keypunch forms was more time-consuming than expected; users would call up data processing whenever they had the slightest problem. Most users still rely on paper reports rather than terminals, we were told.

The conversion required re-writing all of the programs, and could not have been accom-

plished nearly as rapidly without the program generator, they feel. Now that conversion is completed, the data processing department is re-examining the programs to see which need to be re-designed, extended, and so on.

So one reason why a company might choose a modern mini-computer is to take advantage of its program generator or data management system for converting old programs and data efficiently.

When the applications are specialized

When a mainframe system approaches being fully loaded with work, there are several options that data processing management can take. They may be able to add memory, software, and peripherals to the existing system, thereby upgrading it. Or they may replace the current machine with a more powerful one. Or they may order another general-purpose, plug-compatible mainframe for the new work, leaving the existing work on the current machine. Or they may decide to put some of their new, more specialized work on smaller (and possibly non-compatible) computers.

There are several types of specialized work that can be put on smaller machines, such as some types of applications work, data communications processing, program development, and database management. Let us first consider the case of putting new applications on small machines.

New applications. Companies that are distributing *new* applications to smaller machines often have special requirements that cannot be met easily using their mainframes. Also, many times these are stand-alone applications that can be put on a less-costly smaller machine, leaving the mainframe for the larger, more complex applications.

To illustrate, one company we visited recently needed a meeting registration system because they periodically host meetings involving several hundred people. Since this application could not be justified on the company's mainframe computer, a programmer created it on a micro—a Radio Shack TRS-80 Model II with 64K bytes of memory. And he used a Radio Shack report generator package, called Profile. The system keeps

track of attendees and spouses, where they are staying during the conference, which airline flight they will arrive on, and which events or sessions they have signed up for. Also, the system contains a billing module, which creates bulk invoices for the participating companies. When a conference takes place, the company takes the micro to the conference and uses it on-site to obtain up-to-the-minute information. This is just one of the numerous uses the company has for this particular micro, but it has proved to be a valuable one.

Other instances of stand-alone applications that we have come across, where small machines have been appropriate, are a computer graphics system for the trust department of a bank, an inventory control system for a warehouse, and an editorial system for newspaper editors. These are applications that generally are not tightly coupled to the companies' production systems, so they can be put onto their own machines.

Along these same lines, a company may want to limit access by most employees to certain applications—applications that perhaps use confidential information, such as data on possible mergers, acquisitions, etc. Such sensitive applications can be put onto small computers where physical access is under tight control.

Another reason why a company might want to put an application on a smaller computer is so that they can buy a second computer for back-up. Providing back-up for a single crucial application may only be economically feasible with a smaller machine.

Also, one or more desired software packages may be available on smaller machines at very low prices. This may well be a main incentive for managers and executives getting their own personal business computers.

Application system development often is easier and faster on small machines, because their operating systems are not so complex. Therefore, the data processing department may choose to use a smaller machine for a project that could not be economically developed on a mainframe.

Existing applications. Companies may also choose to distribute *existing* general applications to smaller machines. Apparently they choose this alternative only as a last resort, based on com-

ments we have received from a number of sources. However, with the appearance of plug-compatible mini-computers, as well as program generators, data management systems, and relatively inexpensive application packages for smaller machines, off-loading existing applications may become more economical, and thus more common.

One application where off-loading the work is generally economical is data entry—where the remote terminals are tied into a small local computer rather than to the central site via long distance data communications. The reduction in data communication costs may justify the re-programming effort. When the data entry only involves entering and validating the information, some validation checks can be performed by the local computer, before the data is passed to the mainframe.

Data communications. Mier (Reference 5) describes the numerous roles that small machines, generally mini-computers, can play in communication networks. He sees them being used as: front-end communication processors, stand-alone network switching nodes, terminal cluster controllers, gateway processors to other networks, and distributed processors.

Although similar in nature, these different uses of computers in networks require somewhat different machine configurations. Many will need large amounts of main memory for protocol conversion, processing, routing, and/or buffering. Some may require supplemental mass storage or other peripherals. So each network node must be tailored to fit the specific communication functions it will perform. Mier believes that smaller machines will play important roles as specialized network processors in the coming years, because network protocols and architectures are becoming more complex, requiring more intelligence at the nodes.

While Mier is talking about using mini-computers at communication nodes in networks linking mainframes, within the past few years micro-based products specifically designed to serve as nodes in mini-computer networks have become available. For example, a micro-based statistical multiplexor handles 'bundling' of traffic from numerous asynchronous terminals to a host mini-

computer, which reduces the number of communication lines needed. Such products have been available for about five years; they provide error checking, asynchronous-to-synchronous conversion for more efficient data communications (if desired), and some can support multi-drop operations. Also, micro-based port selectors allow asynchronous terminals to communicate with several mini-computer hosts using a single communication link. Two companies offering such products are Micom of Chatsworth, California, and Racal-Milgo of Miami, Florida.

Program development. As we discussed in the October and November 1979 issues, programming work-station systems are available for off-loading program development work from mainframes. Since development work is often given 'second class status,' as compared to production work, and since it does not need to interact with currently running production work, application development can be treated as a specialized application and given its own machine.

The newest development that we are aware of in this area is IBM's Application Development Center concept. IBM is encouraging companies to off-load program development work onto a system that has a number of integrated development tools. IBM generally favors the use of a *mainframe* for a function such as this, but the concept is equally applicable for off-loading development work to a smaller machine.

Another approach involves off-loading development work to end users. One way to do this is via the IBM Information Center concept, which provides them with resources on a large machine, as we discussed in our June 1981 issue. Another way is to give the end users their own smaller machines that have powerful data management systems, as we discussed in March, May, and July 1981.

Database management. Back-end database machines are beginning to appear. Britton Lee, for example, offers two models, the IDM 500 and the IDM 200. Also, Tandem's mini-computers can be configured to provide database machine service. In March 1981 we described how one company uses a Tandem system in this manner. This company's database—a catalog of books in print—is growing at such a rapid rate that they

need to increase the database often and incrementally, while maintaining acceptable response times for users. Software AG and Storage Technology also offer database machines. The use of back-end database machines is just beginning; it remains to be seen how well user companies take to this approach.

So there are numerous 'special situations' where using a smaller machine is a practical alternative.

When the remote site is too small for a mainframe

Until recently, the most practical, yet sometimes still-too-expensive, way to provide computing to remote sites was to provide them with a computer terminal, a printer, and a modem—and let them communicate with a mainframe.

Today, micro-computers provide other alternatives. They can permit very small remote sites to have quite powerful computing facilities on-site. Our discussion earlier in this report of 8-bit and 16-bit micros, and networks of micros, indicates some of the alternatives available.

Last month we described the case of school districts that need only part-time use of a micro-computer to prepare their monthly financial reports—so they use micros, acquired for student instruction, in non-instruction hours. Before this time, the small school districts had no economical automated alternative available to them.

The 'operator' problem. In the case of small remote sites, their 'computer operation' usually is too small to support a regular computer operator. So who handles the computer problems that occur? Generally, someone at the remote site makes it his or her duty to become knowledgeable in such matters.

But the problem is not limited to small sites. The problem exists even for remote sites with (say) one or two dozen employees. In such cases, a new possible solution is a mini-computer and operating system combination which permits 'operator-less' remote operation. Operator-less operation is often a requirement for communication processors, which may serve as nodes in a network, for example. But what about operator-less application processors?

Possible problems with operator-less networks. What are the problems associated with operator-less networks? One company with experience helping users design and implement distributed systems is Honeywell Information Systems. We talked with Grayce Booth, who has been involved in distributed systems work for a number of years. Based on her experiences with users, she told us that she has seen the following six major problem areas in implementing an operator-less network. Generally, the solutions to the problems are new types of software or changes to existing software, she says.

Need to down-load software. First, networks with operator-less remote computers need the facility to down-load software from the central site. This requires special software in both systems. Down-loading software to an up-and-running remote site usually is straight-forward. However, when there has been a system failure, software may need to be re-transmitted to the remote site, but only after the remote system has been restarted. For this, each operator-less system needs start-up routines in a read-only memory that allows it to be re-started from the central site. The central site also needs software to perform its part of this operation.

Using value-added networks. Second, even with the facilities in place for down-loading software, it may still not be possible to achieve an operator-less operation across a public or value-added network, Booth points out. This is because these networks use packet switching, which breaks up communications into separate packets, with each packet containing addressing information that must be stripped off at the remote site *before* the packet information can be used. If the remote system has gone down, it must be restarted via a packetized message. But how can that non-working system first strip off the addressing information and re-assemble the packets into their proper sequence, in order to receive the re-start message?

Automatic restart and recovery. Third, software at the remote site must contain automatic restart and recovery procedures. This seems evident, yet how many of the mini and micro operating systems on the market today have such features?

Systems software at the remote site. Fourth, along the same lines as the prior point, Booth notes that both the application and the system software at the remote site must assume that there is no human operator who can intercede when there is a problem. For most of the mini-computer systems on the market, the operating system and the application software assume that there *is* a human operator, and they ask for help from that operator when certain problems occur.

Database backup and recovery. Fifth, while mainframe database systems have sophisticated and efficient facilities for database backup and recovery, such is not the case in database systems for smaller machines. There are generally no backup tape units at remote sites, only floppy disks, which make for tedious backup. Possible solutions are the new streamer tape units coming to the market or using the mainframe as the backup.

Remote diagnosis. Sixth, another seemingly obvious point is that there needs to be some way to diagnose problems in the remote system from the central site. This includes diagnosing the remote software as well as the remote hardware. Programs to remotely diagnose software problems are not common today. And when you add to that the problem of using products from different vendors, remote diagnosis can be a severe problem.

Booth has seen these six areas as the major problems encountered by Honeywell users who have planned and implemented operator-less networks. While there are few such networks in operation today, she is optimistic and feels that all of these problems can be solved.

At a recent seminar, a computer operations manager discussed some of the problems that her organization is experiencing in just trying to *reduce* the number of operators, on night shifts and on weekends. If there are no operators on duty on a night shift, for instance, excellent hardware and software reliability are essential. Sophisticated monitors are needed for hazard detection. Operators can hear unusual noises in mechanical units, or can smell that something is overheating, and can turn off the equipment before severe damage is done. And if the site is

subject to repeated power failures, operators may be needed for getting things running again.

Operator-less computers surely are coming. But there are still a number of problems to be solved.

When you want to try out new systems ideas

The advent of micro-computers appears to have re-ignited the pioneering spirit in many information systems professionals. We are constantly hearing of exciting and innovative new projects that companies are undertaking—and practically all of them are being performed on smaller machines. In the office automation area especially, pilot projects are using minis and micros.

There appear to be several reasons for this phenomenon. One, of course, is cost. Experimenting on a small machine is less expensive, and therefore, it can be cost-justified more easily. Another reason is ease of use—the smaller machines are often easier to program, because they do not have big, multi-tasking operating systems to complicate matters. Yet another reason is that new, exciting packages are appearing on these smaller machines, for a wide variety of uses. And these packages are usually low cost.

One organization that is doing pioneering work in the area of computer-to-computer networks is the U.S. Federal Communications Commission. They have chosen smaller machines, one reason being that these machines allow their experimentation to evolve in more controllable and economical steps.

The FCC's project

The U.S. Federal Communications Commission (FCC) is the government agency responsible for regulating interstate tele-communications activities within the United States. The FCC authorizes use of the radio frequency spectrum, regulates interstate communications via common carriers, and oversees radio and television broadcast companies. They also establish tele-communication policy and co-ordinate with foreign governments on tele-communications matters. And the FCC licenses land, aviation, and marine communication users.

This last-named function—licensing—is an important one; they issue over one million licenses a year. Moreover, this licensing workload has grown enormously in recent years. From the mid-1930s through the end of 1970, for example, they had issued a cumulative total of some 5.5 million authorizations to providers of tele-communication services. By August of 1980, that figure was over 23 million—a 340% increase in just ten years. And the growth rate is expected to increase.

With an increasing workload such as this in sight, in 1979 the automated data processing steering committee was formed to study the commission's data processing operations and create a long range plan for all of the FCC's data processing for the 1980s. To date, the main data processing workload for the commission has been performed on two Honeywell 66/60 mainframe computers.

Einar Stefferud of Network Management Associates, of Huntington Beach, California, has acted as a consultant to the FCC on this long range plan. The plan, described in *The Future of Electronic Information Handling at the FCC: Blueprint for the 80s* (Reference 6), has been approved by the commission and is being implemented. We will briefly discuss this plan and the FCC's progress on it to date.

The information management network plan. The FCC sees a continuing need for large computers, to maintain central files, provide central database management facilities, and perform large-scale computations. They also see the desirability of distributed data processing, to handle the numerous different requirements of users. They view their users as FCC staff, other federal agency employees, and the public.

Their plan envisions three different types of computer systems. *Single-user systems*, employing micro-computers, will support individuals with local computation, text editing, and access to other computers. *Multiple-user systems*, somewhat larger in size, will be used to support 10 to 30 users within a group—say, an office, a laboratory, and so on. These systems will provide the single-user systems with software program libraries, larger files, and communication facilities to other computer systems. *Remote-*

utility systems, using mainframe computers, will provide large central databases and large-scale computation capabilities, just as they do today.

All of these kinds of systems are in existence today. What is not now available is the network and software to inter-connect them, provide data transport and mail facilities, and mediate the device and computer incompatibilities.

The FCC's network plan has two components—local networks and a remote network—each performing different functions. A *local network* will provide 'close coupling' between the single-user work stations and the multiple-user systems, within a local work area. These networks thus must provide very fast data and program transfer capabilities. *The remote network* will provide users with access and file transfer capabilities among the multiple-user systems, remote utility systems, and public networks.

Both types of networks will need to provide similar functions, including (1) work station access to any computer, (2) file transfer between computers (including work stations) with appropriate file access controls, and (3) electronic mail, which differs from file transfer in some respects.

The FCC sees database management systems being a major tool for creating both specialized applications and commission-wide applications for this network system.

The plan calls for an evolutionary seven-year implementation, 1981 through 1987. Where possible, existing systems will be incorporated into the overall network plan.

Implementation to date. The first phase is nearing completion under direction of the associate managing director for information management, who is responsible for the entire data and office automation program for the FCC. The FCC has begun with the multi-user systems and the remote network, which they call FCCNET. And they have established a permanent network management staff to manage this new network. They have acquired six 16-bit multi-user micro-computers from Onyx Computers of San Jose, California. Each system has either 500K or 1 megabyte of main memory, one or more 18 megabyte hard disk units, and a 12.5 megabyte tape system

for initial program loading and file backup. The price of a basic system was about \$31,000.

One of these systems is used by the FCC laboratory in Columbia, Maryland, for tracking the testing and certification of communication devices. Two more are located in Gettysburg, Pennsylvania, where they are used to track the processing of FCC license applications. The remaining three are located at FCC headquarters in Washington, D.C.

One of these headquarters micros serves solely as a mail relay switch. It also connects the Honeywell mainframes to FCCNET. In the next phase, users of the existing mail system on the Honeywell mainframes will be able to exchange messages over FCCNET as well. The Honeywell mail system is used extensively by staff at headquarters and by the outlying field offices for internal FCC communications.

The other two micros at FCC headquarters are currently used for preparing and then routing urgent correspondence among FCC officials. Examples of such uses are responses to letters from members of congress which need to be drafted and then approved by several FCC officials, or drafting replies to consumer inquiries.

Each Onyx has five to seven terminals connected to it and each runs an identical copy of the 'network mail transfer' software. Each communicates with the other micros via a loosely coupled network, involving conventional data communications. The principal use of network mail is to transmit messages, but they assign a broad meaning to the term 'message.' It includes files, inter-office memos, software program updates, data records for processing on other computers, and even documents for backup filing on alternate computers.

The communication protocol and multi-channel memo distribution facility (MMDF) that are used for FCCNET mail were developed for the Computer Science Network (CSNET). This is a new network established to link computer science centers in educational institutions, industry, and government to share knowledge and resources. It has been funded primarily by the National Science Foundation. The MMDF software was developed mainly at the University of Delaware by the Department of Electrical Engineer-

ing. So participants on CSNET and FCCNET can communicate with each other because they use the same communication protocols.

Currently, then, the FCC has installed six multi-user micro-computers and a loosely-coupled network for sending messages among them. During the next phase, they plan to add additional micro- or mini-computers to FCCNET for more applications, and connect the Honeywell mainframes and some field office micro-computers to the FCCNET mail system. They also plan to install higher bandwidth local network connections among the headquarters computers. Work will also begin on finding suitable single-user work-station computers as the micro evolution continues to provide more capability for less money.

Conclusion. As we have seen, smaller machines are being used by information systems departments to replace obsolete equipment, off-load existing specialized applications, distribute new work, provide remote sites with computing facilities, and experiment with new ideas. As small systems become more powerful, we expect information systems executives to take advantage of them in still other ways, so this list will grow.

REFERENCES

1. Canning, R. G. and N. C. Leeper, *So You Are Thinking About a Small Business Computer*, Prentice-Hall, Inc. (Englewood Cliffs, New Jersey 07632), July 1982; 216 pages; price: \$11.95.
2. Schindler, P.E. Jr., "Third-party vendors offer micro-to-mainframe solutions," *Information Systems News*, CMP Publications, Inc. (333 East Shore Road, Manhasset, New York 11030), June 15, 1981, pp. 28; price \$1.25.
3. Kincaid, M. "Large companies using small computers—a trend to 'used,'" *Data Management*, DPMA (505 Busse Highway, Park Ridge, Illinois 60068), January 1981, pp. 18-22; price \$3.00.
4. "Monthly newsletter examines maintenance policies on used gear," *Computerworld* (Box 880, 375 Cochituate Rd., Framingham, Massachusetts 01701), February 22, 1982, pp. 11; price \$1.25. This article contains a brief description of some of the survey results. For more information about the survey on vendors' maintenance policies for their used equipment, contact: *Computer Economics Report*, Computer Economics Inc. (150 Fifth Ave., New York, New York 10001). This monthly newsletter costs \$295 per year.
5. Mier, E., "Minicomputers' new network roles," *Data Communications*, McGraw-Hill (1221 Avenue of the Americas, New York, New York 10020) October 1980, pp. 77-93; price \$2.00.
6. *The future of electronic information handling at the FCC: Blueprint for the 80s*, October 31, 1980, will be available later this year from the U.S. government's National Technical Information Service (5285 Port Royal Road, Springfield, Virginia 22161). Order number and price not yet available.

It is becoming apparent that training office employees to use, and feel comfortable with, terminals and personal work stations is a challenging subject. For one thing, there will be so many employees to train. The problem is especially acute for training executives and managers, because it is unreasonable to expect them to spend many hours in classes. Fortunately, the computer itself can help provide this needed training, via computer assisted learning. Next month we discuss how computers are being used in such training.

The 'end user interface' is taking on increasing importance in office systems, because end users may refuse to use a system that is not 'comfortable' to them or is one they cannot remember how to use. Month after next we delve into the question of the end user interface and discuss two types of more 'friendly' query systems—those using natural languages and those having several levels of interfaces so to accomodate both novice users and experienced users.

EDP ANALYZER is published monthly and copyright© 1982 by Canning Publications, Inc. 925 Anza Avenue, Vista, California 92083. All rights reserved. Photocopying this report for personal use is permitted under the conditions stated at the bottom of the first page. Also, see Declaration of Principles on page 15.

COMMENTARY

MANAGING REMOTE APPLICATION DEVELOPMENT WORK

by Burton Grad, Heights Information Technology Service Inc., Tarrytown, N.Y.

(Because of the growing interest in tele-commuting and the management of 'remote' employees, we asked Burton Grad, President of Heights, to discuss some of the main considerations in managing off-premises data processing people. Heights is successfully using flexible-time, off-premises professionals, as described in this Commentary—Ed.)

Heights Information has been in the business of providing software services since April 1979. Luanne James and I established the company in co-operation with F International, a firm based in the U.K. which has been in the software services business since 1962.

One of our fundamental premises in starting Heights was that, in spite of the 'programmer shortage,' there were a substantial number of highly qualified people available to perform application development work, but that they wanted more flexible work schedules than most employers provide. This flexibility might involve working at home, working less than normal hours, or not working for periods of time because of family commitments or travel opportunities.

A second premise for our business was that clients would be more interested in getting their jobs done well at competitive prices than in requiring that the person assigned be at their site while doing the work.

Regarding the first premise, we have indeed found data processing professionals who want flexible working arrangements. We recruit as 'panel members' people with a minimum of three to five years of actual computer programming and analysis experience. While locating specific skills is difficult, it has not been hard to find qualified people with 'standard' IBM programming backgrounds. Our panel members sign an agreement form which states that they are independent contractors and are free to obtain work on their own, as well as through Heights.

Currently, our activities are centered around New York City and the San Francisco bay area. As of April 1982, Heights had over 80 panel members in the New York, New Jersey, and Connecticut area, and over 150 in the San Francisco, Oakland, and San Jose, California area. We do perform work for clients in other locations, when the work can be done by panel members in the current branches.

Regarding our second premise—that we could find employers willing to use people on a remote or flexible time basis—this has turned out to be a much more difficult factor than we anticipated. There is a great deal of educational work which must be done with clients before they will even try this different way of getting the work done—particularly in larger companies that have previously used outside professional software services firms; they usu-

ally insist that the work be done on their site, on a full-time basis. Smaller companies are more willing to work in this new mode, as long as the project plan is well spelled out and they receive frequent progress updates.

Here is how we answer two concerns—about using remote people—that are often expressed to us by prospective clients.

How will we know how much actual work effort is going into the job? Every significant application development project deserves to have specifications carefully spelled out, a specific work plan put in place, deliverables and milestones established, and frequent progress reports against the plan. There must be substantial on-going communication during the project to maintain effective liaison and insure that the client is satisfied on how the work is proceeding. When this type of professional project management is followed, we have found that the success ratio of work done on an at-home or part-time basis has been as high as with the projects where we have put the person on-site on a full-time basis.

How can we (the client) provide for frequent communications with the remote person, to discuss changes in the specifications or work procedures? At Heights, we have found that we must provide an additional level of interface and support, to insure that communications are clear and continuing, and that expectations are realistic on both sides. One way that we do this is by using a project manager, who does little (if any) programming but who is the main communication link between the client and the project team. It is up to the project manager to see that the needed communication flows both ways. In addition, we have found that on almost every project there has to be significant on-site time, in order to collect information, work out mutual understandings, and test and install the product.

We have such confidence in this approach, based on our experience since 1979, that we are usually prepared to make fixed-price bids on work where we have developed the specifications or where they are clearly spelled out by the client.

Clients have often told us how pleased they were with the high level of dedication of our 'remote' panel members. They felt they paid only for actual work time, and projects were completed on (or close to) schedule.

Should you, then, try to use people on a remote or part-time basis, either as employees or from a service such as Heights? With proper personnel qualifications and good project management, there is no reason why you cannot get good productivity from people who need flexible work locations or time schedules. However, managing remote employees does require that extra level of support that I discussed above. Also, you must analyze the effect that this approach may have on your other data processing employees.

It is our belief that the use of people who need flexible schedules and work locations will substantially increase the available programming resources. Also, we think that it will, in the long run, prove to be a common way of doing business in the data processing field. Don't close your eyes to the modern 'electronic cottage.'

SUBJECTS COVERED BY EDP ANALYZER IN PRIOR YEARS

1979 (Volume 17)

<i>Number</i>	<i>Coverage</i>
1. The Analysis of User Needs	H
2. The Production of Better Software	H
3. Program Design Techniques	H
4. How to Prepare for the Coming Changes	K
5. Computer Support for Managers	C,A,D
6. What Information Do Managers Need?	C,H
7. The Security of Managers' Information	L,C,A
8. Tools for Building an EIS	C
9. How to Use Advanced Technology	K,B,D
10. Programming Work-Stations	H,B
11. Stand-alone Programming Work-Stations	H,B
12. Progress Toward System Integrity	L,H

1980 (Volume 18)

<i>Number</i>	<i>Coverage</i>
1. Managing the Computer Workload	I
2. How Companies are Preparing for Change	K
3. Introducing Advanced Technology	K
4. Risk Assessment for Distributed Systems	L,E,A
5. An Update on Corporate EFT	M
6. In Your Future:Local Computer Networks	F,B
7. Quantitative Methods for Capacity Planning	I
8. Finding Qualified EDP Personnel	J
9. Various Paths to Electronic Mail	D,M
10. Tools for Building Distributed Systems	E,B,F
11. Educating Executives on New Technology	K
12. Get Ready for Managerial Work-Stations	C,A,B

1981 (Volume 19)

<i>Number</i>	<i>Coverage</i>
1. The Coming Impact of New Technology	K,A,B
2. Energy Management Systems	M
3. DBMS for Mini-Computers	G,B
4. The Challenge of "Increased Productivity"	J,K,A
5. "Programming" by End Users	C,H,B,G
6. Supporting End User Programming	C,H,B,K
7. A New View of Data Dictionaries	G,B
8. Easing the Software Maintenance Burden	H,B,G
9. Developing Systems by Prototyping	H,B,G
10. Application System Design Aids	H
11. A New Approach to Local Networks	F,K
12. Portable Software for Small Machines	B,H

1982 (Volume 20)

<i>Number</i>	<i>Coverage</i>
1. Practical Office Automation	A,B,C,K
2. Computer Graphics for Business	K,C,B
3. Interesting Decision Support Systems	C,B,H,A
4. Can Tele-communications Replace Travel?	A,F,J,L
5. The Human Side of Office Automation	A,J,K
6. Some Users Want Their Own Computers	B,C,K
7. Using Minis and Micros	B,E

Coverage code:

A Office automation	E Distributed systems	I Computer operations
B Using minis & micros	F Data communications	J Personnel
C Managerial uses of computers	G Data management and database	K Introducing new technology
D Computer message systems	H Analysis, design, programming	L Security, privacy, integrity
		M New application areas

(List of subjects prior to 1978 sent upon request)

Prices: For a one-year subscription, the U.S. price is \$60. For Canada and Mexico, the price is \$60 in U.S. dollars, for surface delivery, and \$67 for air mail delivery. For all other countries, the price is \$72, including AIR MAIL delivery.

Back issue prices: \$7 per copy for the U.S., Canada, and Mexico; \$8 per copy for all other countries, sent via AIR MAIL.

Reduced prices are in effect for multiple copy subscriptions, multiple year subscriptions, and for larger quantities of a back issue. Write for details. Agency orders are limited to single copy subscriptions for one-, two-, and three-years only.

Please include payment with order. For U.S. subscribers, you can use your Visa or MasterCard charge card; include your card name, number, and card expiration date on your order.

For payments from outside the U.S., in order to obtain the above prices, take your choice of three options: (1) use an international money order, (2) pay in U.S. dollars with a check drawn on a bank in the U.S., or (3) use any of the following charge cards: Visa, MasterCard, Eurocard, Access Card, Standard Bank/Kaart, Union Card International, or Diamond Card International. Please be sure to include your card name, number, and card expiration date on your order.

Editorial: Richard G. Canning, Editor and Publisher; Barbara McNurlin, Associate Editor. While the contents of this report are based on the best information available to us, we cannot guarantee them.

Missing Issues: Please report the non-receipt of an issue within one month of normal receiving date; missing issues requested after this time will be supplied at the regular back-issue price.

Copying: Photocopying this report for personal use is permitted under the conditions stated at the bottom of the first page. Other than that, no part of this report may be reprinted, or reproduced or utilized in any form or by any electronic, mechanical, or other means, now known or hereafter invented, including photocopying and recording, or in any information storage and retrieval system, without permission in writing from the Publisher.

Address: Canning Publications, Inc., 925 Anza Avenue, Vista, California 92083. Phone: (714) 724-3233, 724-5900.

Microfilm: EDP Analyzer is available in microform, from University Microfilms International, Dept. P.R., (1) 300 North Zeeb Road, Ann Arbor, Mich. 48106, or (2) 30-32 Mortimer Street, London WIN 7RA, U.K.

Declaration of Principles: This publication is designed to provide accurate and authoritative information in regard to the subject matter covered. It is sold with the understanding that the publisher is not engaged in rendering legal, accounting, or other professional service. If legal advice or other expert assistance is required, the services of a competent professional person should be sought. — From a Declaration of Principles jointly adopted by a Committee of the American Bar Association and a Committee of Publishers.