On Science and Technology, Teague carried the merit principle in staffing further than it customarily had been in the House. Staff members were hired according to qualifications for a particular job and were expected to serve all committee members regardless of party. Present and past staffers attest that as job applicants they were never asked their party affiliation. A recruiting system was developed under which criteria were set for open positions and then applicants were interviewed by senior staff members whose recommendations were taken seriously. Science and technology soon outdistanced other committees in the number of Ph.D. scientists and engineers employed.

The system has not worked frictionlessly. Republicans chafed for years because of tight limits on staff available to the minority. When Representative Charles A. Mosher of Ohio was ranking Republican he agreed generally with Teague on staffing principles. But Mosher retired 2 years ago and the current ranking Republican, John W. Wydler (R–N.Y.) has pressed the prerogatives of the minority to appoint some staff as allowed in recent reforms. Observers say, however, that quality control on staff has been maintained.

Teague's heir apparent as chairman is Representative Don Fuqua (D–Fla.). An eight-termer, he is regarded as knowledgeable and diplomatic and is expected to have no trouble in getting his chairmanship confirmed by the caucus or in dealing with members of the committee. Like Teague he is a conservative and, also like Teague, his experience as ranking majority member has been mainly on space matters, since he has chaired the subcommittee on space science and applications. But during the past Congress the chairman's health has been poor and Fuqua has worked closely with Teague, giving Fuqua a runnning start on the chairman's job. Fuqua, of course, faces changed circumstances in which a chairman is expected to reign but not to rule. Nor can he draw on the kind of congressional capital that Teague built up during his three decades in the House.

As for Teague, it is hard to sum up his career. He never felt impelled to run for the Senate or to seek the topmost rungs of the leadership ladder in the House. Presumably because of his own experience, he felt an obligation to a special constituency and left an indelible imprint on veterans' programs.

The move late in his career to the chairmanship of the Science and Technology Committee put him at the center of national science affairs, and his own standing in the House enabled him to strengthen decisively his committee's position.

His personal impact on science and technology issues is difficult to assess in detail. It is evident, for example, that when the National Science Foundation's education directorate was under attack because of a controversial school behavioral science course, Teague's actions as chairman helped keep the congressional reaction under control. And clearly, Teague played an important role in the come back of the science adviser. But it is probably not for specific actions that Teague has earned the high regard in which scientists who have dealt with him almost invariably seem to hold him. Teague has evinced an old-fashioned, perhaps somewhat unreflective but nonetheless unselfish, sense of public service and it is really through this quality that he has left his mark on science policy.

—JOHN WALSH

RESEARCH NEWS

Array Processors: Maxi Number Crunching for a Mini Price

Stimulated by ever cheaper and smaller components in the shape of microelectronic circuits, the minicomputer and, coming a decade later, the microcomputer have made it possible for researchers to own and operate their own computers, which they then dedicate to particular ends by way of the programs they write. A new trend, also driven by microelectronics, is now in view in which specialization is achieved by the organization of the computer circuits (hardware) themselves. At the leading edge of this changing scene is the array processor.

A device especially designed to carry out, at high speed, arithmetic operations on large arrays of numbers, such as vectors and matrices, the array processor can, with some limitations, transform a minicomputer into a computer for large-scale scientific computation—that is, into a number cruncher. Because of the modest financial investment required, about that of a minicomputer itself, it is becoming economically feasible for some researchers to do involved computations that once were not affordable, while others are able to process enormous volumes of experimental data, even if the computations themselves are not demanding. And, interestingly enough, far from hastening the demise of the central computing facility, the array processor may reverse the fortunes of those unable to purchase the newest maxicomputers by elevating a pedestrian large or mainframe computer into the near supercomputer class, although again there are some important qualifications—no one is getting something for nothing.

The largest group of users of array processors is in the field of signal processing, and the devices were originally designed to meet the needs of this group. One kind of signal processing is the analysis of radar reflections in "real time"—that is, as they are received—for enemy aircraft or missiles among a host of false signals and a large background "noise." A second type of signal processing is image analysis of microwave, infrared, or optical satellite photographs. Here, real-time response may not be necessary, but the volume of data generated is overwhelming—a Landsat-type satellite produces photographs at the rate of one each 20 seconds.

The champion users of array processors are those in the geophysical exploration business, primarily hunters of oil and natural gas. The now dominant means of accomplishing the exploration is, according to Donald Townsend of the Exploration Services Division of Geosource, Inc., in Houston, by seismic reflection surveys.

A large pad carried beneath a heavy truck is excited to vibrate at a known frequency, thus launching waves into the earth. Over a 10-second interval, the frequency is swept from, for example, 10 to 50 hertz. An array of up to 48 microphones (geophones) records, every few milliseconds over a period of 15 to 20 seconds, samples of the intensity and times of arrival of sound waves traveling through the earth and reflecting from boundaries between structures, such as rock layers and fluid-filled porous mate-

0036-8075/79/0112-0156$00.75/0

rials, with different sound velocities. The truck moves 50 meters or so along a line that may be several kilometers long, and the whole process is repeated. A similar procedure, but using ships in place of trucks and compressed air charges instead of a vibrating pad, enables marine seismic reflection data to be gathered. At a recent computer conference, Carl Savit of Western Geophysical Company, also in Houston, estimated that, last year, between 10^{14} and 10^{15} bits of data were recorded worldwide in this way.

To translate this jumble into a form meaningful to geophysicists—that is, a map showing the depth and location of the reflecting structures, because some of these are signals of oil- or gas-containing formations—much computation involving operations such as Fourier transforms and correlation and convolution integrals is necessary. Every sample requires, as a minimum, 100 arithmetic operations, and the total load is equivalent to the work output of a few dozen of IBM's largest general-purpose computer. At a few million dollars per survey, the prospective investment is large, even for the oil business. Moreover, the exploration companies prefer to do some of their data processing in the field, an environment possibly fit for minicomputers but certainly not for mainframes. The answer that has evolved over several generations of technology is the array processor.

Conceptually, the array processor is an outgrowth of what computer people call a peripheral processor or, less elegantly, a stunt box. The early minicomputers, for example, did not do floating point arithmetic in which a number is represented by the appropriate significant digits and an exponent. The advantage of floating point arithmetic is that operations on very large and very small numbers are possible without loss of the significant digits; but such operations cause the computer to slow down. Thus some companies, and minicomputer manufacturers quickly followed, sold special circuits that could be plugged into the mini to do floating point operations. Fancier stunt boxes since developed include devices to execute fast Fourier transform algorithms for signal processing and Fourier transform spectroscopies.

The advantage offered by stunt boxes is speed, which derives from the fact that they are nonprogrammable (hard-wired) devices that do only one operation, but therefore do it very well. An elaborate version of this type of device, developed for the oil exploration companies, was trade-marked the Array Transform Processor. Built by Raytheon Data Systems of Norwood, Massachusetts, it could execute 14 commonly called operations when attached to a host minicomputer. Many hundreds of Array Transform Processors were, and still are, used in geophysical exploration, according to Townsend.

The next step was to increase the versatility of stunt boxes by making them programmable and by adding to them their own internal memory, so that they could operate without being limited by the unimpressive speed at which the host minicomputer could pass data into and fetch it back from the box. The first company that achieved an array processor with these capabilities and that was able to market it successfully in both the geophysical exploration and general scientific research communities was Floating Point Systems of Portland, Oregon. (For others, in a now highly competitive business, see box.) The Floating Point Systems array processor, in turn, is built around concepts developed by Glen Culler of Culler/Harrison, Inc., in Santa Barbara, California.

Parallel Is Faster

Culler started with the idea of building a programmable, special-purpose computer. In the course of his work, Culler realized that a limitation on the processor speed was, as described by John Backus of IBM's San Jose research laboratory, the von Neumann bottleneck, after the mathematician John von Neumann, one of the fathers of the stored program computer. To execute an instruction, the computer must locate the instruction in memory, interpret the instruction, locate the data in memory to be acted upon, bring the data to the central processor, perform the operation, return the new data to memory, and so on. The bottleneck is that all communication is by way of a single pathway (bus) between a single memory and the logic and control units, so that only one of the functions can be taking place at any instant and each must occur in the proper sequence. For any given piece of data, this sequence cannot be violated; but, if many numbers, for example, are going to be treated in a similar way, operation is considerably speeded if all functions are going on in parallel, each with respect to a different number. In the final version of the array processor as implemented by Floating Point Systems, multiple communications pathways and spatially separated memories allow up to ten functions to be going on simultaneously, with each function taking place at a different location within the array processor.

A second kind of parallel organization is pipelining. The two basic arithmetic operations, addition and multiplication, are themselves divisible into suboperations. Conventionally, while a set of two numbers is being added, for example, a second set of numbers, also to be added, must wait until the first set is completely through the logic unit called the adder before being acted upon. In pipelined computers, the first suboperation can begin on the second set of numbers as soon as the second suboperation begins on the first set. Although the total time to add any two numbers is not decreased by pipelining, when many additions are to be performed together, the total time is reduced by a factor equal to the number of suboperations.

Both types of parallelism are therefore ideally suited to numerical calculations of large arrays of numbers, and, at last, there is the reason for the name array processors. There are, however, considerable differences in the design and capability of array processors; not all are pipelined, for example. According to George Hallnor of the TRW Systems Group, McLean, Virginia, these differences make each manufacturer's array processor best suited to a particular class of problems.

The actual benefits accrued in the use of an array processor depend on several factors. One published report, by Richard Bucy and K. D. Senne of the University of Southern California (USC), compares the performance of an array processor and some of the premier number crunchers in the execution of a typical signal processing algorithm using, as the measure of performance, the number of floating point operations per second divided by the cost of the processor. The USC researchers found that the array processor was more than 17 times as cost-effective (and only 2.5 times as slow) as today's recognized supercomputer, the CRAY-1 made by Cray Research, Inc., which in turn was four times as effective as a Control Data Corporation CDC 7600, the standard high-speed scientific computer.

A word to the wise, however, is that such comparisons can be misleading in that real computer programs consist of much more than floating point arithmetic operations. The actual benefit in using an array processor therefore depends sensitively on how much number crunching is to be done, expressed as a fraction of the total program running time. Frank Phelan of the Scripps Institution of Oceanography, an enthusiastic user of an array processor, realistically notes that the speed of a minicomputer augmented with an array processor can be comparable to that of a CDC 7600 only if the array processor can be kept busy.

A second essential is to do the number crunching in blocks. This point has been quantitatively analyzed by John Newkirk and Jon Claerbout of Stanford University, who compared the speed of their array processor with that of their host minicomputer alone. The Stanford investigators showed, for example, that an array processor could multiply two vectors faster than the host only when the vectors had 60 or more components. The problem is the time that it takes the host minicomputer to pass the information needed for the multiplication to the array processor. The solution, say some observers, is that, having the two vectors in the array processor, a programmer would do the utmost to carry out as many operations as possible on these quantities before returning the results to the host or, as is now possible, before directly passing data to an external mass storage device.

A related complication is associated with programming. Many manufacturers have now developed a library of standard mathematical subroutines that run in the array processor. A programmer can more or less write a normal program in Fortran and simply insert, at appropriate points in the program, statements that tell the minicomputer to execute the desired subroutine on the array processor. The problem is that, if the subroutine is as simple as a vector multiplication, then the inefficiency described by Newkirk and Claerbout comes into play. Even for more complex subroutines, time is still lost communicating program information between the host and array processor. In the end, to keep the array processor busy, the programmer must learn to write a lower level computer language than Fortran (assembly language) or, in some cases, even change the basic instructions of the array processor by a technique known as microcoding.

In assembly language, the programmer must learn to keep track of the various functions that together constitute the execution of an instruction. Probably nothing hindered the early acceptance of minicomputers more than the prospect, for the computer neophyte who thought learning Fortran was accomplishment enough, of learning assembly language. E. Randolph Cole of USC's Information Sciences Institute, Marina Del Ray, estimates that, because of the need to keep track of several things in parallel, assembly language programming is half again as difficult for an array processor as for a minicomputer.

Array processor compilers, programs that can translate Fortran into the array processor's assembly language, are just now becoming available, but they are not yet efficient. Peter Buhl, of the Lamont Doherty Geophysical Observatory, Palisades, New York, judges that for his application, which is similar to the marine seismic reflection studies of the geophysical exploration companies, his minicomputer-array processor system is about 100 times faster than the mini alone, provided that he programs in assembly language. Using the compiler to generate the assembly language code slows things down to the point that the array processor-equipped system is *only* ten times faster. However, programmers

Choosing an Array Processor

A prospective purchaser can count on finding, if he searches diligently, 40 or more companies that will sell him an "array processor." In short, there is as yet no uniformly accepted notion of what constitutes an array processor and what is really something else, but some guidelines can be given.

According to some observers, there really are two classes of device under the general heading of peripheral processor: the array processor and the signal processor. Others make the distinction with less enthusiasm. One type of array processor is that, described in the accompanying story, developed to meet the needs of the geophysical exploration companies. It is designed primarily to do floating point arithmetic operations on arrays (vectors and matrices) of numbers, and its internal structure reflects this. "Signal processors" were developed for government (often military) agencies interested in such applications as extracting target signals from radar reflections. Often these devices were constructed originally under government contract and only marketed later. Signal processors often do not do floating point operations, and hence are faster. They often have shorter computer word lengths than "array processors," and their internal structure is not oriented toward handling arrays of numbers. Several hundred of each type have, by now, been sold.

Array and signal processors further break up into three groups. The first comprises large expensive array processors, costing one to several million dollars, such as those made by IBM and the Control Data Corporation exclusively for their own mainframe computers. On the other end of the price scale, the second group consists of the logic circuits needed to make up an array processor mounted on circuit boards, ready for mounting in a minicomputer chassis. In this form, array processors can cost from $5000 to $20,000. Finally, in the center, lie the true, general-purpose, fully programmable array processors. These devices come housed in their own chassis, contain their own internal memory, and can be attached to a variety of minicomputers and, for some manufacturers, mainframe computers, although some vendors make array processors specifically for their own brand of minicomputer. These machines come priced from $20,000 to $150,000.

A recent article in a computer magazine* listed, in a table, 15 manufacturers, the models they offered, and the price range and performance of each. The article was written by the president of one array processor manufacturer. With this warning in mind, listed below are the names of the 15 (plus one additional) array processor makers: Control Data Corp., Minneapolis; Computer Design and Applications, Inc., Newton, Massachusetts; CSP, Inc., Burlington, Massachusetts; Data General Corp., Westboro, Massachusetts; Data West Corp., Scottsdale, Arizona; ESL, Inc., Sunnyvale, California; Floating Point Systems, Inc., Portland, Oregon; Honeywell, Inc., Waltham, Massachusetts; IBM, Manassas, Virginia; Magnavox (a militarized version of the CSPI array processor), Fort Wayne, Indiana; PCS, Saline, Missouri; Plessy, Inc., Gaithersburg, Maryland; Raytheon Data Systems, Norwood, Massachusetts; Signal Processing Systems, Inc., Waltham, Massachusetts; Stein Associates, Waltham, Massachusetts; and Westinghouse Electric Co., Baltimore, Maryland. Let the buyer beware!

—A.L.R.

**Mini-Micro Systems*, July 1978, p. 54.

point out that once the key subroutines of a problem are coded in assembly language, a compiler can be very useful for building up programs that can call many subroutines before returning control to the host.

A final point, for those doing high-precision computation, is that the array processors offer a shorter computer word length than available on maxi computers. The number of significant digits in a number is limited by the number of binary bits in the computer word. Some scientists have shied away from the array processor because of this limitation. Henry Schaefer and William Miller, of the University of California at Berkeley, who pioneered the used of high-performance minicomputers in quantum chemistry, are examples of those waiting for an array processor that can do high-precision arithmetic.

With some loss in processing speed, it is, however, possible to circumvent the short word length limitation in some cases. At Cornell University, John Wilkins needs to diagonalize matrices in his cal-

Surgeon General Seeks Physicians' Help in DES Alert

The surgeon general, Julius Richmond, issued a physician advisory last October to every M.D. and osteopath in the country, warning of serious hazards associated with the synthetic estrogen diethylstilbestrol (DES). It is a drug that perhaps 2 million American mothers were given during pregnancy when it was freely prescribed between 1940 and 1970. DES has been linked with cancer in the daughters of the women who took it, and it is suspected of having done some harm to the mothers and—recent studies suggest—to the sons as well.

A warning from the nation's chief medical officer presumably signals something akin to a national health emergency. But the DES advisory stirred no commotion, largely because the dangers have been publicized already, and because the use of DES reportedly dropped to a trickle once the cancer link was documented. A spokesman for the surgeon general, Michael White, said that he did not know how many physician advisories had been issued in the past, but he judged them to be rare and significant because he knew of only two: one on the carcinogenic effects of asbestos, sent out early in 1978, and the DES advisory, sent on 4 October.

DES was first linked in 1971 with a rare form of vaginal cancer (clear cell carcinoma) in daughters of women who used the drug during pregnancy, and in the same year the Food and Drug Administration (FDA) ordered that DES be labeled as contraindicated for what was then a common use—the prevention of miscarriage. Oddly, DES was rather widely used for this purpose through 1971, even though a study done as early as 1953 at the University of Chicago found that it was ineffective as an aid to pregnant women. A number of studies since then have strengthened the suspicions about DES, and in 1976 the Department of Health, Education, and Welfare (HEW) began a major campaign to alert the public to the hazards. In March 1978, the secretary of HEW created a task force to look into the question and make policy recommendations. The report came out in September and the advisory in October.

In addition to recommending that the surgeon general himself alert physicians to the problem, the task force reached the following conclusions:

- Between 4 million and 6 million people were exposed to significant doses of DES and must be watched for symptoms of cancer or infertility. This includes the mothers who used DES, their sons, and particularly their daughters, for whom the evidence of risk is strongest.
- The risk of clear cell carcinoma in the daughters is "well established," but small; the incidence is expected to be between 1.4 per 1,000 and 1.4 per 10,000 in the exposed population through age 24.
- The risk of cancer in the mothers is "unproved," although there is "cause for serious concern" that exposure to DES may be linked with a higher risk of developing cancer of the breast, cervix, uterus, or ovary. More studies are needed.
- There is "no firm evidence" of higher rates of infertility or cancer among DES sons, but "recent studies clearly show an excess of genital abnormalities" in this group, the sort of abnormalities that can signal cancer. Again, studies are needed.
- Physicians who prescribed DES should warn their patients of the risks and urge them to come in for checkups. Daughters of DES mothers should be examined at least once a year, beginning at age 14. This counseling, HEW says, should be provided free of charge.
- DES should be prescribed only for FDA-approved uses: estrogen replacement in young women with serious deficiencies, therapy for menopausal symptoms, treatment of cancer, and treatment of osteoporosis (for which it is expected to be approved). The drug has been approved for use in suppressing lactation in new mothers, but the FDA is in the process of canceling this one from the list. Although DES is occasionally given as a postcoital contraceptive (known as the morning-after pill), and although one company has proposed to market such a pill, the FDA has not approved the drug for this use. Bernard St. Raymond of the FDA said the drug is available for other purposes, and there is "no way you can force the physicians not to prescribe it as a contraceptive." Apparently it is used as an alternative to abortion, most frequently by college women. The task force recommended: "Postcoital contraception with estrogens in any woman should be restricted to situations where no alternative is judged acceptable by a fully informed patient and her physician."

The medical profession is less than exuberant about the advisory, for it asks physicians to do a considerable amount of research in old files, it gives them the onerous task of informing patients that their doctor has given them a bad drug, and it suggests that all this be done gratis. "Maybe in five years we'll be out of this quagmire," said Ervin Nichols, a spokesman for the Washington, D.C., office of the American College of Obstetricians and Gynecologists. He felt that the decision to publish the advisory was "way out of proportion" to the size of the problem. Many physicians, he said, already have made an effort to tell the patients who received DES about the risks, and he doubted that there would be any major campaign to search the files for new names. Doctors will probably restrict themselves to answering queries about DES when they come in.

—E.M.

0036-8075/79/0112-0159$00.50/0

culations of the electronic structure of solids. In this application, the required electron wave functions are initially calculated with insufficient precision because of the short word length. Wilkins finds than an iterative procedure well known to applied mathematicians provide the extra precision he is seeking.

What these caveats add up to is that there is a definite place where array processors are of extreme benefit and some where they are not, and that anyone contemplating acquiring one should, at this point, be willing to spend a little time getting to know the machine at a pretty intimate level, if the maximum possible performance is the goal. Those who have made this commitment have been more than satisfied.

Kent Wilson, of the University of California at San Diego, had one of the very first array processors manufactured and has a newer model now. Wilson is interested in molecular dynamics calculations in liquids. In essence, given 100 atoms, their positions and velocities, and the forces between them, what happens? The problem is a straightforward classical mechanics exercise of integrating Newton's force law for each atom. For each set of initial conditions, some 300 coupled differential equations must be intergrated for several thousand time steps, depending on how long the atoms and molecules are to be tracked. However, most of the possible initial coordinate and velocity sets are not of interest because, if the equations of motion were integrated, nothing much (formation of molecules or reactions between molecules, for example) would happen.

Wilson is working on a solution that involves helping the researcher interact with the program by way of a touch interface and thereby use the chemist's intuition to full advantage. The touch interface allows the researcher to feel, with his hand, the forces on particles selected by way of a cathode-ray tube display. Moreover, with the interface, he can guide particles into "interesting" trajectories since the device applies, through the computer, the translational and rotational forces necessary to move them.

Wilson uses an array processor coupled to a minicomputer to carry out the integration of the equations of motion, a conceptually simple task but one which involves repeating the operations vast numbers of times. Without an array processor, Wilson says, a computer with the capability of a CDC 6600 or 7600 would be needed, but there is no way a single researcher could afford one.

Suppose one wanted to simulate a system of 10^6 particles, not just 100. This is the problem faced by those engaging in plasma simulations for use, among other reasons, in modeling fusion energy reactors. The particles in the plasma—a gas of ionized atoms and electrons—generate, by their electrical charge and motion, electric and magnetic fields. These fields in turn provide the forces that move the particles around. The sheer number of particles, however, imposes a constraint: how is one to get the data describing the positions, velocities, and forces on a million particles into and out of the array processor fast enough to take advantage of its speed. A second kind of plasma simulation that treats the system of charged particles as a continuous fluid meets the same limitation.

Low Cost as Important as Speed

Researchers at the plasma physics center of the University of California at Los Angeles (UCLA), headed by John Dawson, attacked this problem with a special computer, built around a commercial array processor, from Culler/Harrison. The computer has several unusual features. One is a set of four processors, called input-output processors, each of which controls the flow of information between one high-speed disk memory and the array processor, thereby bypassing the slow process of communicating data between the host processor and the array processor. As Culler points out, however, such an arrangement is only useful if the actual numerical calculation is so simple that it can be done quickly. If it is long and slow, then there is no need for a special system to provide data at such a rapid rate.

According to some test runs of the plasma simulation system by Dawson, Robert Huff, and Cheng-chin Wu of UCLA, the Culler/Harrison computer is about 200 times more cost-effective than UCLA's IBM 360/91, as measured by the cost to purchase the former and the rate charged by the university's computer center for use of the latter.

Although almost all array processors have been attached to minicomputers, they can just as well be made to augment the capabilities of the maxicomputers. A physicist at Cornell, Kenneth Wilson, has followed this tack. Wilson has been interested in applying a numerical random sampling (Monte Carlo) method, already highly developed in statistical mechanics, to an altogether different realm, that of calculating the strong nuclear force between quarks, the currently popular entities presumed to be the constituents of many elementary particles, such as the proton and neutron. A particular problem, as yet unsolved by theorists, is why quarks are not found as free particles. Any Monte Carlo calculation requires considerable numerical calculation, and Wilson's application is exceptionally demanding. Cornell's central computer, an IBM 370/168, is a little slow but mainly too expensive for him to use. An alternative, the CDC 7600, accessible at the Lawrence Berkeley Laboratory by way of telephone lines, is fast enough, less expensive to use than Cornell's own machine, but still too costly for the extensive, and untested, computation envisaged by Wilson.

As soon as Wilson became aware of the array processor, he set about exploring the best way to acquire one. In part because of a lack of a suitable minicomputer host at Cornell and in part because Wilson wanted access to the services provided by the computer center, a new arrangement was forged. A consortium of a half-dozen Cornell scientists pooled funds and bought a two-thirds interest in a new model array processor designed to be attached to a mainframe rather than to a mini. The remaining one-third was purchased by the computer center, and the array processor was hooked up to the IBM computer there.

What is already highly evident, Wilson claims, is that it is much cheaper to run a program on a array processor, even if there is no increase in speed, because the array processor is so inexpensive. Wilson, who programs in assembly language extensively, finds his programs run about three times faster and 100 times more cheaply than they would on the IBM machine alone. Others, writing Fortran programs that are compiled with a compiler developed at Cornell, find their programs running more slowly than on the IBM computer, but they are still saving considerable money.

Alec Grimison, head of the Cornell array processor project, says that the computer center makes its share of the array processor open to all users who have access to the Cornell computer system, although this has not been widely advertised. Nonetheless, word-of-mouth transmission has been rapid and the rate of appearance of new users is just about as high as can be handled.

All in all, as the word spreads and researchers become more familiar with the array processor concept, and if more efficient compilers are developed, these devices are sure to become more popular. And facilities such as that at Cornell may provide a way for many who would not risk the effort otherwise to make friends with the remarkable array processor.—ARTHUR L. ROBINSON