A Geophysical Outlook–Part 4

New vector super computers promote seismic advancements

H. Roice Nelson, Jr., General Manager, Allied Geophysical Laboratories, University of Houston

20-second summary

An introduction to and a layman's summary of the needs for, architecture of and expected uses of the new vector computers in exploration geophysics is presented in this article, the fourth in a series on new exploration technologies. Array processors that perform specialized number crunching tasks have been tied to most geophysical processing systems over the last 10 years. However, the use of vector computers in seismic data processing is just beginning to be tested and put on-line in the largest of the major oil companies. The $10 to $20-million price tag has something to do with this. The use of these vector machines further requires a complete rethinking and reworking of seismic processing packages before the parallel processing techniques that they afford will provide cost effective increases in processing speed. Major areas of expected application of this technology in exploration geophysics are also summarized.

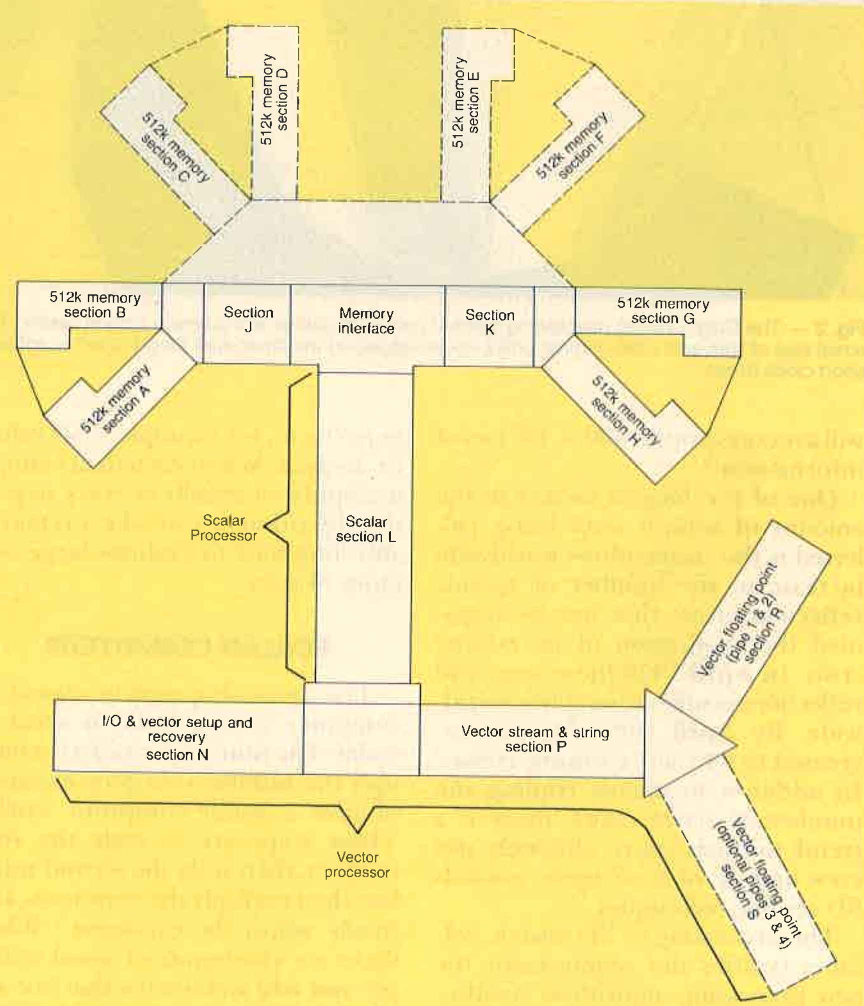

Fig. 1 — CDC's Cyber Model 205 vector computer system floor plan, as shown, takes up relatively little floor space, only about 250 square feet, in comparison to the work load it performs.

COMPUTATION REQUIREMENTS for geophysical processing historically have increased at least as fast as the development of new computer technologies. The need for increased processing power is the result of two major factors: the sheer volume of data that is being collected today, and the new processing algorithms or procedures that require instant access to larger portions of the data volume being evaluated.

Recent developments in multichannel recording systems[1] and in 3D (three-dimensional) seismic acquisition techniques[2] have greatly increased the amount of data that is being collected. For example, a 24-channel shothole crew will collect on the order of 24 × 10^6 bits of information per day.[a] While about 300 × 10^6 bits of information will be collected by a 48-channel VIBROSEIS crew.[b] Recording 32-bit words for higher dynamic range and using 3D recording procedures, a 120-channel VIBROSEIS crew will result in collection of about 1.250 × 10^9 bits of information per day.[c] The daily amount of data collected by a 1,024-channel sign-bit VIBROSEIS crew

[a] 50 shots/day × 24 channels/shot × 5 seconds/channel × 250 samples/second × 16 bits/sample = 24 × 10^6 bits/day
[b] 100 vibrator points/day × 48 channels/V.P. × 16 seconds × 250 samples/second × 16 bits/sample = 307 × 10^6 bits/day
[c] 80 vibrator points/day × 120 channels/V.P. × 16 seconds × 250 samples/second × 32 bits/sample = 1.229 × 10^9 bits/day

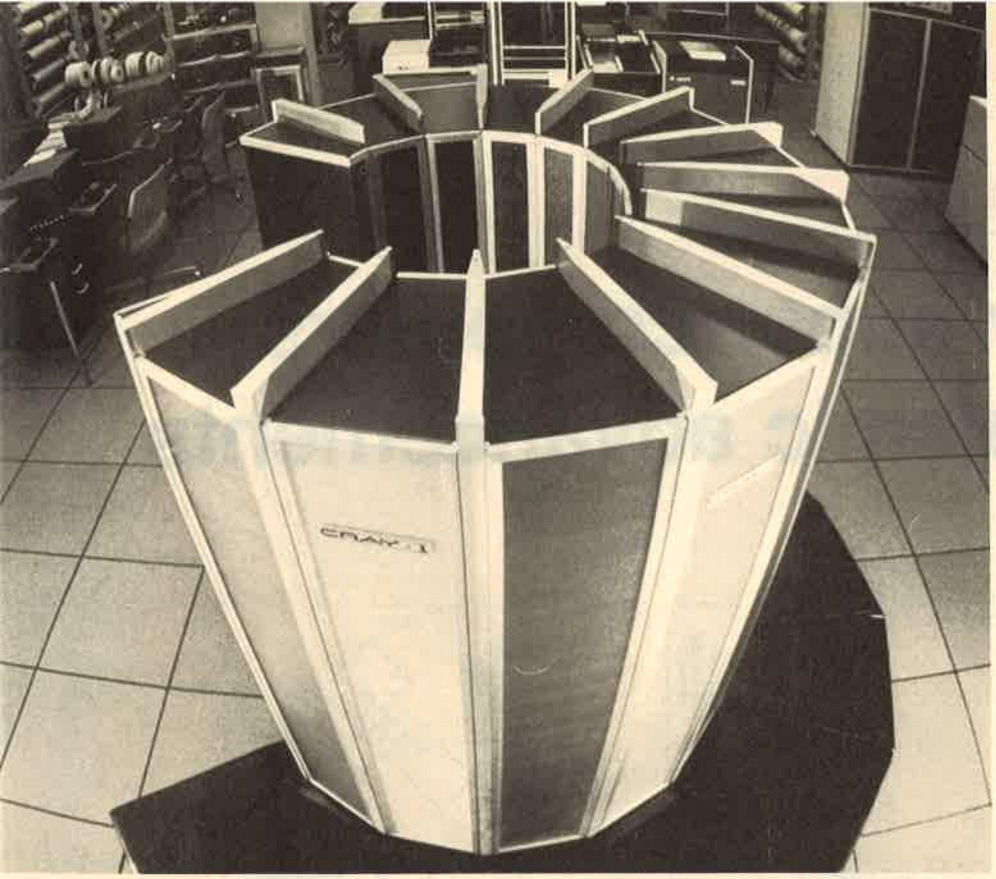

Fig. 2 — The Cray-1 vector processing system looks somewhat like a bench sofa in reality. The small size of this, and other similar units, is due largely to the short wire length used to achieve short clock times.

will average around 400×10^6 bits of information.[d]

One of the biggest factors in the amount of seismic data being collected is the tremendous worldwide increase in the number of seismic reflection crews that has accompanied the realization of an energy crisis. In April 1976 there were 240 reflection seismic crews active worldwide. By April 1981 this had increased to 640 active seismic crews.[3] In addition to almost tripling the number of active crews, there is a trend towards more channels per crew and an obvious move towards 3D seismic techniques.

The processing of 3D seismic volumes typifies the requirement for new processing algorithms (mathematical procedures) that need instant access to larger portions of available data sets. When a volume of seismic data has a full 3D migration algorithm applied to it, all of the data within the radius of the migration aperture needs to be available for computing each output trace.[4] However, on a smaller scale, similar needs occur when CMP (common midpoint) gathers from two or three crossing lines are concurrently used to perform, for example, a 3D velocity analysis. When each data value is manipulated serially at every step of the algorithm, it can take a remarkably long time to evaluate large volumes of data.

SCALAR COMPUTERS

The processing used by almost all computer systems today is serial or scalar. The four steps needed to multiply two numbers is a good example of how a scalar computer works. These steps are to scale the first number, then scale the second number, then multiply the mantissas, and finally adjust the exponent.[5] When there are thousands of serial multiply and add statements that are applied to each trace, it can add up to a lot of computer time.

One way that the bottleneck inherent in scalar processing is overcome is to build electronic black boxes that sidetrack the problem. These systems are the various forms of array processors. They are designed to do repetitive operations on large arrays, like a seismic trace, quickly, and where possible to do more than one operation at the same time. When an array processor multiplies an array of numbers by one number, the array is fed into a pipe line. The first number of the array goes through the standard serial steps. However, as this number moves to the next step, another number follows it into the pipe line. Once the pipe line is full, there is one multiplication result coming out for each processing step or clock cycle. This can result in significant time savings when doing repetitive operations on long numerical arrays.

Array processors also have the capability to do several operations in parallel with one instruction. A typical example is to have a floating point addition, a memory read, a floating point multiplication, and an integer addition (bumping a counter) programmed with one statement to occur simultaneously.[5] In order to use the array processors efficiently the large arrays have to be programmed so they will be recognized as vectors. It is an important step for the processor to recognize these arrays as vectors. The application of arithmetic and logical operations to vectors is known as vector processing.

VECTOR COMPUTERS

The new super vector computers can be grossly described as giant array processors. A more accurate portrayal is to describe these systems as a combination of a host and a task processor. They have the same capabilities for pipe line and parallel processing. However, the machine controls automatically which operations are done in scalar mode and which statements are done in vector mode. The same logic that says that a computer-controlled multiprocessing environment is more efficient than a single-job processing mode applies to this aspect of vector computers. Many operations appear to the user to be performed in a serial fashion, with other operations being performed in a parallel mode, but all operations are issued in strict sequence from a single instruction stream. The parallel or vector processing mode allows the manipulation of many operands* by a single instruction as discussed above. The operating system permits functional concurrency wherever possible while retaining the logical soundness of the user's program.[6]

[d] 100 vibrator points/day × 1024 channels/V.P. × 16 seconds × 250 samples/second × 1 bit/sample = 410×10^6 bits/day

*Something (as a quantity or data) that is operated on (as in a mathematical operation).

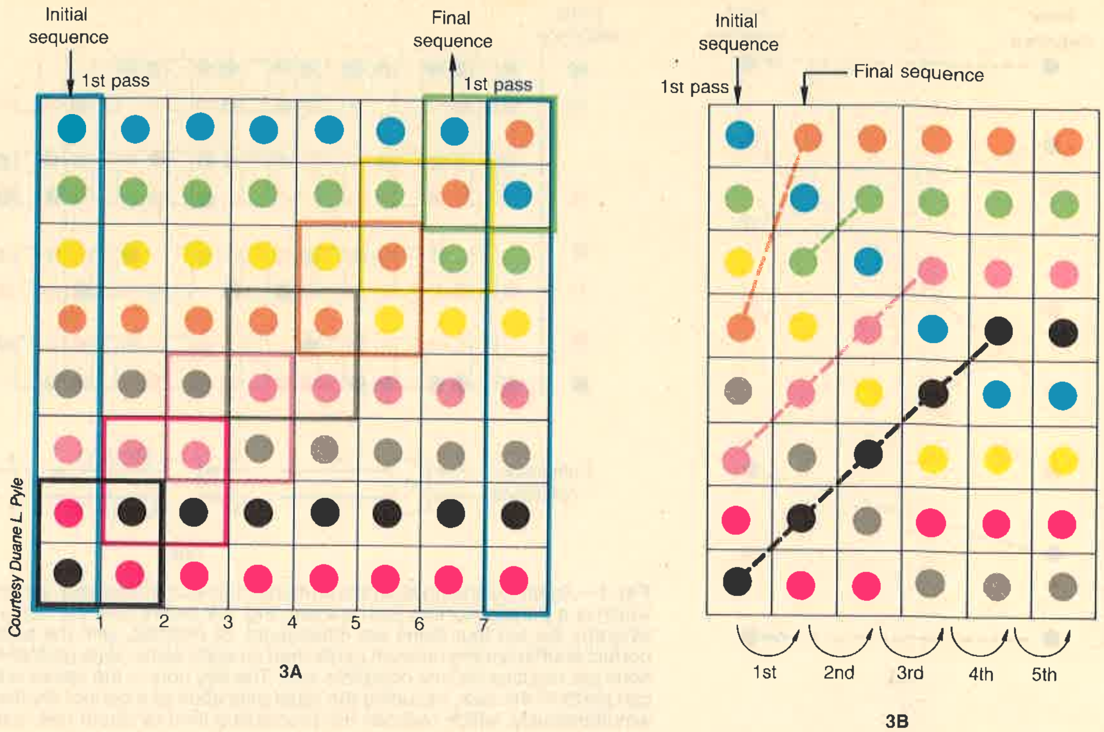

Fig. 3—In order for explorationists to take advantage of the amazingly rapid processing capabilities of the new vector super computers, the software used must be completely rethought. This example of a "bubble sort" illustrates one of the algorithms designed for vector software. Fig. 3A shows how seven comparisons are required for the first pass on an eight-item bubble sort, while 3B shows the five sequential passes needed for one complete bubble sort of the same eight items. With the volumes of data used in modern geophysical processing these serial procedures become prohibitive in processing time and cost effectiveness.

A good example of parallel processing capabilities is the program loops that occur so often in scientific and geophysical processing. The pervasive nature of program loops has resulted in hardware and software designs that make it convenient for assembly language/machine language programmers to code program loops. Higher language loop structures, like the Fortran 'DO' statement, play a central role in geophysical processing. These loops can often be represented in terms of algebraic operations performed on linear lists, or vectors, of operands. The vector computer takes advantage of the repetitious, similar computations involved in such loops by overlapping various stages of the computation.[7] This is accomplished in a similar manner to the pipe line processing discussed above.

One major advantage of vector processing over scalar processing is the elimination of overhead associated with maintenance of the loop control variable. Often these loops reduce to a simple sequence of instructions without backward branching. However, not all aspects of a problem lend themselves to vector processing. Scalar techniques still need to be applied to effectively do branching, as in a computed 'GO TO' statement. Consequently, the large vector machines provide both a scalar and a vector processor, with instructions and registers for both applications.[6,8]

Seismic processing has had a major impact on the development of vector processing. For example, by the early 1960s the need for a rapid convolution method, or procedure to locate a given pattern on a seismic trace, was defined. The petroleum industry therefore sponsored the development of "convolver boxes." The main idea was to access the data as little as possible and maximize the arithmetic performed during each access. By the late 1960s these boxes had developed into the first floating point convolvers. The similarity of the convolution product algorithm to most other algorithms of linear algebra generated research and development of more generally applicable array processors.[9] By the beginning of the 1970s various seismic contractors and major oil companies had better defined their needs and had worked with hardware manufacturers to develop 32-bit systems for seismic and 64-bit systems for reservoir modeling, where rounding error requires more precision.[10] These systems developed and featured convolution, matrix multiplication, FFTs (fast Fourier transforms), recursive filtering and a variety of simpler algorithms.

By 1974, when CDC (Control Data Corp.) introduced the STAR-100, array processors in seismic data processing were well accepted. This was one of the first examples of an integrated scalar and vector processor. Around this same time, Seymour Cray, one of CDC's top computer designers, left CDC and the competing super vector computers of today began to be developed.

Architecture is quite different for the Cray and CYBER vector processing systems. However, it is not the purpose of this article to directly compare the systems or to make recommendations as to which is the best to buy for any specific reason. Therefore, in writing about the capabilities of similar types of equipment of which the author is aware, the descriptions will not be tied to one vendor or the other. The different vector processing systems and their associated models cover a wide range of specifications.

The manner in which vectors are

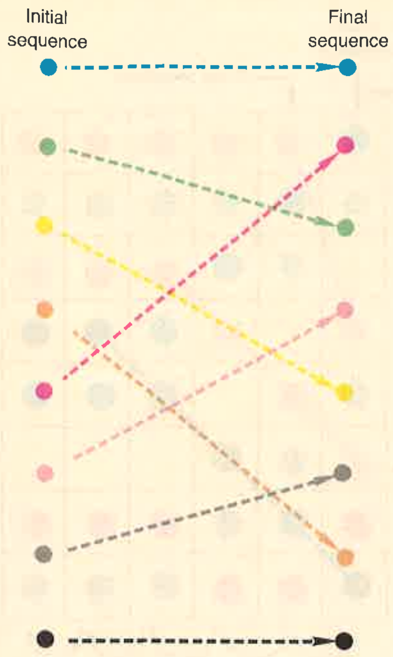

Courtesy Duane L. Pyle

4A

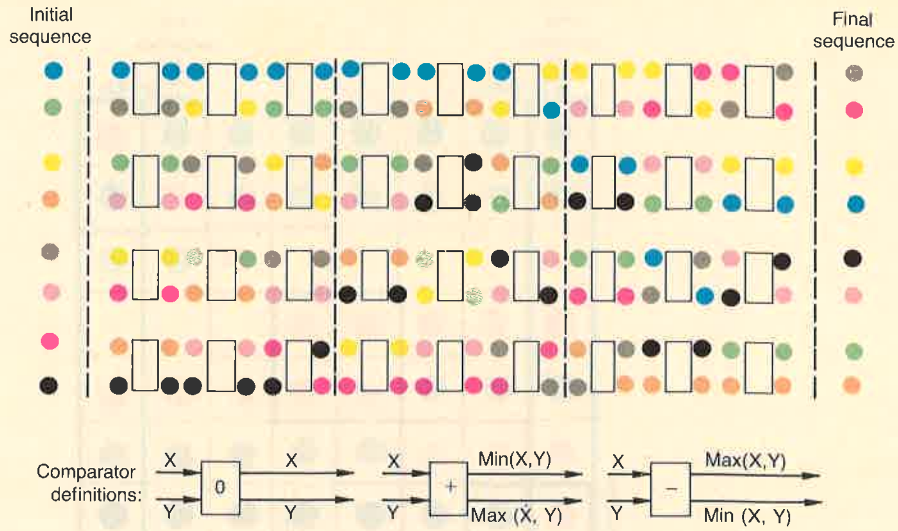

4B

Fig. 4—Again illustrating algorithms utilized in vector processing, a method called the perfect shuffle, which is a parallel sorting tool, is used. Fig. 4A shows how the name perfect shuffle was derived, whereby the top four items are interleaved, or shuffled, with the bottom four. Fig. 4B illustrates a perfect shuffle sorting network performed on eight items. Nine perfect shuffles involving 36 comparisons are required for one complete sort. The key here is the speed with which the vector processor can perform the task, including the rapid execution of a perfect shuffle and doing four comparisons simultaneously, which reduces the processing time by about one quarter.

processed is a major difference between the Cray and CYBER systems. One system loads each vector into a vector register and then adds this to another vector register. This is theoretically a better way to work with short vectors, but also gives good results for long vectors. The other vector processor uses memory-to-memory streaming. This requires additional startup time to fill the pipes, but works extremely well on longer vectors.

Real memory in vector processors normally ranges up to 1 million 64-bit words. However, this can go as high as 4 million 64-bit words. One system has this memory arranged in 16 banks with a bank cycle time of 4 clock periods. This makes very efficient random access memory. The other system has 2 × 10^{12} words of virtual memory and 16 input/output ports with a total I/O rate of 3.2 × 10^9 bits per second. The arithmetic, logical and shift operations are handled in the pipes or functional units. There are up to 4 of these in one system and 12 in the other.

One of the most fascinating aspects of these computers is their relatively small size. Fig. 1 shows a floor plan of a CDC-205. The unit is about 6 feet in height and does not require more than about 250 square feet of floor space. The Cray-1 (Fig. 2) looks somewhat like a bench sofa, and likewise takes up little floor space, as shown. Despite the small size, these systems require several of the largest available scalar computers (front end machines) to feed them enough data to keep the processing stream going. The small size has largely to do with the short wire lengths required to get the phenomenal clock cycle times.

Computation speeds. Computer capabilities are measured by the speed that instructions and mathematical operations can occur. A large scalar computer will operate at about 8 MIPS (million instructions per second). Vector computers will run at rates between 50 and 800 Mflps (megaflops or millions of floating point operations per second). A machine that does 800 Mflps has been designed, but has not yet been built. It will require 4 pipe lines, linked triads, and short precision (32-bit words) to accomplish this speed. (A linked triad is a special kind of vector operation of the form a scalar times a vector plus a vector.)

A single Mflp is equivalent to between 3 and 5 MIPS because floating point arithmetic has more complex instructions. This means that a vector processor manipulates between 25 and 400 times as much data as a large scalar computer. However, experience shows that the vector systems typically are running at 16 to 20 Mflps or 8 to 10 times the speed of a large serial system.[10]

A major determinant in specifying the processing rate is the clock period. The clock period is the time that passes between each processing step, once the pipes are full. Typical array processors attached to a VAX 11/780 minicomputer have a clock rate of 124 ns (nanoseconds). One ns is 10^{-9} seconds. A beam of light travels 30 cm in 1 ns. The available vector processors have clock rates of 12.5 ns and 20 ns.

Software has to be completely rethought in order to take advantage of the extra speed of a vector processor. To illustrate this rethinking procedure, a programming procedure for sorting will be evaluated for a serial algorithm and a parallel algorithm. Probably everyone is acquainted with some form of sorting, whether it be in preparing to balance the check book or in alphabetizing a list of names.

Serial sorting is illustrated using a bubble sort, a procedure whereby specified items appear to "bubble" to the top of a column or list. To start with there are 8 items in random order that are to be sorted, as shown in the left hand column of Fig. 3A. The bottom two items in the left hand column are first compared. The larger number is placed on top of the smaller as illustrated in the second column. There are a total of

7 comparisons required to make the first pass.

A complete bubble sort of these 8 items takes 5 sequential passes as is illustrated in Figure 3B. The larger numbers appear to bubble to the top. There are 28 serial comparisons required to do this sort.[11]

Parallel sorting is described using an algorithm called the perfect shuffle. Fig. 4a shows how the name was derived. The 8 items in the left hand corner are to be sorted. They are first divided in half, and then interleaved in the same way a card deck is shuffled. Each of the 9 times the 8 items are shuffled they are passed through one of 3 comparators. These comparators pass the numbers, put the smaller number on top, or put the larger number on top as illustrated in Fig. 4B. A perfect shuffle sorting of 8 items requires 36 comparisons of which 8 are redundant. However, if these comparisons can be performed simultaneously, then only 9 comparison clock periods are required for the perfect shuffle sort example. This further reduces to 7 clock periods when the 2 redundant comparisons are accounted for. The time to actually perform the perfect shuffle must also be considered. If this algorithm is to work, it is important to be able to do the perfect shuffle fast, as is possible on a vector processor.[11]

These examples point out that it is not enough for explorationists and programmers to understand how to use vector processing capabilities. The algorithms must also be designed or redesigned for the vector computer environment. The data must be structured correctly so that data motion will take advantage of the vector processing capabilities.[12] At the computation rates being discussed, data input and output can be critical to achieving cost effective processing.

Data motion, the rate at which data can be input into the computer, can be as important in using vector computers as the computation speed. An example of this is shown by some research occurring at the University of Houston's Research Computation Laboratory. The object is to extend a Fourier method wave equation modeling program[13] to three dimensions and vectorize it for use on the CDC-205 vector processor. The program utilizes five arrays of which three are calculated for each time step. Each of these arrays involves 16 million components requiring a total storage of 80 million words. However, only two million words can occupy the physical memory at any one time, and so the data is divided and stored on disc. Three of the arrays are accessed twice for each time step. Preliminary calculations indicate that the first pass is compute bound and the second pass is I/O bound.[14]

GEOPHYSICAL APPLICATIONS

Numerical reservoir modeling has been the single most important use of vector super computers in the petroleum industry to date. This application involves simulating the physical behavior of pressures, saturations, densities and similar quantities within the porous rocks and sands from which oil and gas are extracted. After the initial discovery well and the subsequent development wells delineate how large a field is, the reservoir engineer enters the picture. The very large numerical models are used to help decide how many more wells need to be drilled, where to drill them and how fast they can be produced. The numerical simulators solve simultaneously for three dependent variables at each of several hundred time steps for up to 12,000 nodes.[15] This scale of number crunching can be approached in several areas of geophysical processing.

Three-dimensional seismic techniques provide the geophysical basis for using vector processors in seismic processing. Many of the 3D processing algorithms should be able to be vectorized and used cost effectively on a vector processor. Especially processes like a full 3D migration that cannot realistically be accomplished on a standard scalar processor. There are several types of 3D migration for regularized (stacked) data: Kirchhoff summation, finite difference, Fourier transform, and direct inversion. Each procedure gives the same basic answer. However, vectorization of the different procedures will probably show more efficient processing for one of the methods than for another. It is possible that Kirchhoff summation techniques can be used to do a 3D migration of irregular data (prestack) using vector processing.[16]

Many of the more basic procedures can be vectorized for considerable processing time savings. The most basic processing procedures include filtering, deconvolution and correlation. VIBROSEIS correlation is an example of a processing procedure that probably can be done as a matrix multiply very rapidly on a vector processor. Again, velocity analysis is an ideal candidate for vectorization when working with 3D velocity analysis algorithms.

It is not unreasonable to project that the parallel logical operations associated with a vector processor will have great value in data base management, and possibly even interactive image processing and interpretation. The new display technologies are rapidly being tied to computers. As complex, large number crunching processing and numerical modeling techniques are vectorized, these display devices will provide an operator window to interactively work with and modify data sets. Assuming some basic input/output problems can be solved, it is estimated that a 2D Fourier synthetic modeling program that takes ten hours to run on a VAX will take five minutes on a Cyber-205.[17] This same program took between 6 and 10 hours in a VAX 11/780 multiprocessing environment.[10] Once the data transfer problems are solved, it is reasonable to project that this will become an interactive interpretation aid.

It is also reasonable to speculate that as inversion and forward modeling techniques are improved there will be the development of computer interpretation procedures. The computer would interpret the raw data using pattern recognition procedures, and a forward modeling algorithm would be applied to the interpretation to make a seismic section. The result would be compared to the raw data, a new interpretation made based on the difference in the two sections and an iterative procedure followed until a satisfactory comparison is achieved. This is only a beginning in describing the potential of vector processors in geophysical processing and analysis.

SUMMARY

Geophysical processing has had a major impact on the development of array processors, and will continue to have an effect as computer technology moves into the world of the vector super computers. The computation requirements necessary to correctly process a 3D seismic survey will be a major motivating force in these developments. These techniques bring reflection seismology into better conformance with reservoir engineering and similar computer fitting procedures can be developed. There is still a lot of development that needs to occur, specifically in the process of software vectorization and data motion. However, there are many petroleum applications that appear to be ideal for this environment.

ACKNOWLEDGMENTS

Many of the ideas presented in this article are from conversations with professors at the Allied Geophysical Laboratories. Special thanks go to Olin Johnson, Duane Pyle, and especially my family.

LITERATURE CITED

1 Nelson, H.R., Jr., "Trends in Multichannel Seismic Recording Systems," *World Oil*, Nov. 1981.
2 Nelson, H.R., Jr., "3D Seismic Techniques," *World Oil*, Dec. 1981.
3 SEG News Bulletin, *GEOPHYSICS*, V. 46, No. 11, p. A-67, 1981.
4 French, W.S., "Two Dimensional and Three Dimensional Migration of Model Experiment Reflection Profiles," *GEOPHYSICS*, V. 39, No. 3, p. 265-277, 1974.
5 Verm, R., University of Houston, Image Processing Laboratory, personal communication, 1981.
6 Technical Description CDC Cyber 200/Model 205, 36 pp., 1980.
7 Pyle, L.D., "A 3D Migration Program for the CDC CYBER 205," Seismic Acoustics Laboratory, fourth year Semi-Annual Progress Review, V. 7, p. 01-05, 1981.
8 Johnson, P.M., "An Introduction to Vector Processing," reprint *Computer Design*, 9 pp., February 1978.
9 Johnson, O.G., Belzer, J. (et al), "Convolution Array Processing," *Encyclopedia of Computer Science and Technology*, Petroleum Industry section, V. 12, p. 52-57, 1979.
10 Johnson, O.G., University of Houston, Computer Science Department, personal communication, 1981.
11 Pyle, L.D., University of Houston, Computer Science Department, personal communication, 1981.
12 Edwards, C.A.M., Control Data Corporation, Petroleum Technology Center, personal communication, 1981.
13 Kosloff, D.D., and E. Baysal, "Forward Modeling By a Fourier Method," Seismic Acoustics Laboratory, fourth year Semi-Annual Progress Review, V. 7, p. H1-H22, 1981.
14 Johnson, O.G., Progress Report on the 3D Wave Equation Program for CDC CYBER 205, Seismic Acoustics Laboratory, fourth year Semi-Annual Progress Review, V. 7, p. I1-I5.
15 Peaceman, D., Belzer, J. (et al), "Numerical Reservoir Modeling," *Encyclopedia of Computer Science and Technology*, Petroleum Industry Section, V. 12, p. 57-63, 1979.
16 Gardner, G.H.F., University of Houston, Allied Geophysical Laboratories, personal communication, 1981.
17 Kosloff, D.D., visiting professor of geophysics, University of Houston, personal communication, 1981. ■

About the author

H. ROICE NELSON, *Jr., received a B.S. degree in geophysics from the University of Utah in 1974 and an M.B.A. degree from Southern Methodist University in 1981. He was employed by Amoco Production Co. in the summers of 1970 and 1973. From 1974 to 1980 he held various geophysical positions with Mobil Exploration and Production Services, Inc., in Dallas, Texas. He became a senior researcher with the Seismic Acoustics Laboratory at the University of Houston in January 1980 and held that position until recently being appointed general manager of the Allied Geophysical Laboratories, associated with the same university. He is currently serving as a member of the SEG Research Committee.*