A CRAY RESEARCH, INC. PUBLICATION

CRAY CHANNELS

Volume 6, Number 3

ANNOUNCEMENT!
CRAY X-MP, SSD set new standards

FEATURE ARTICLES:

DEPARTMENTS:

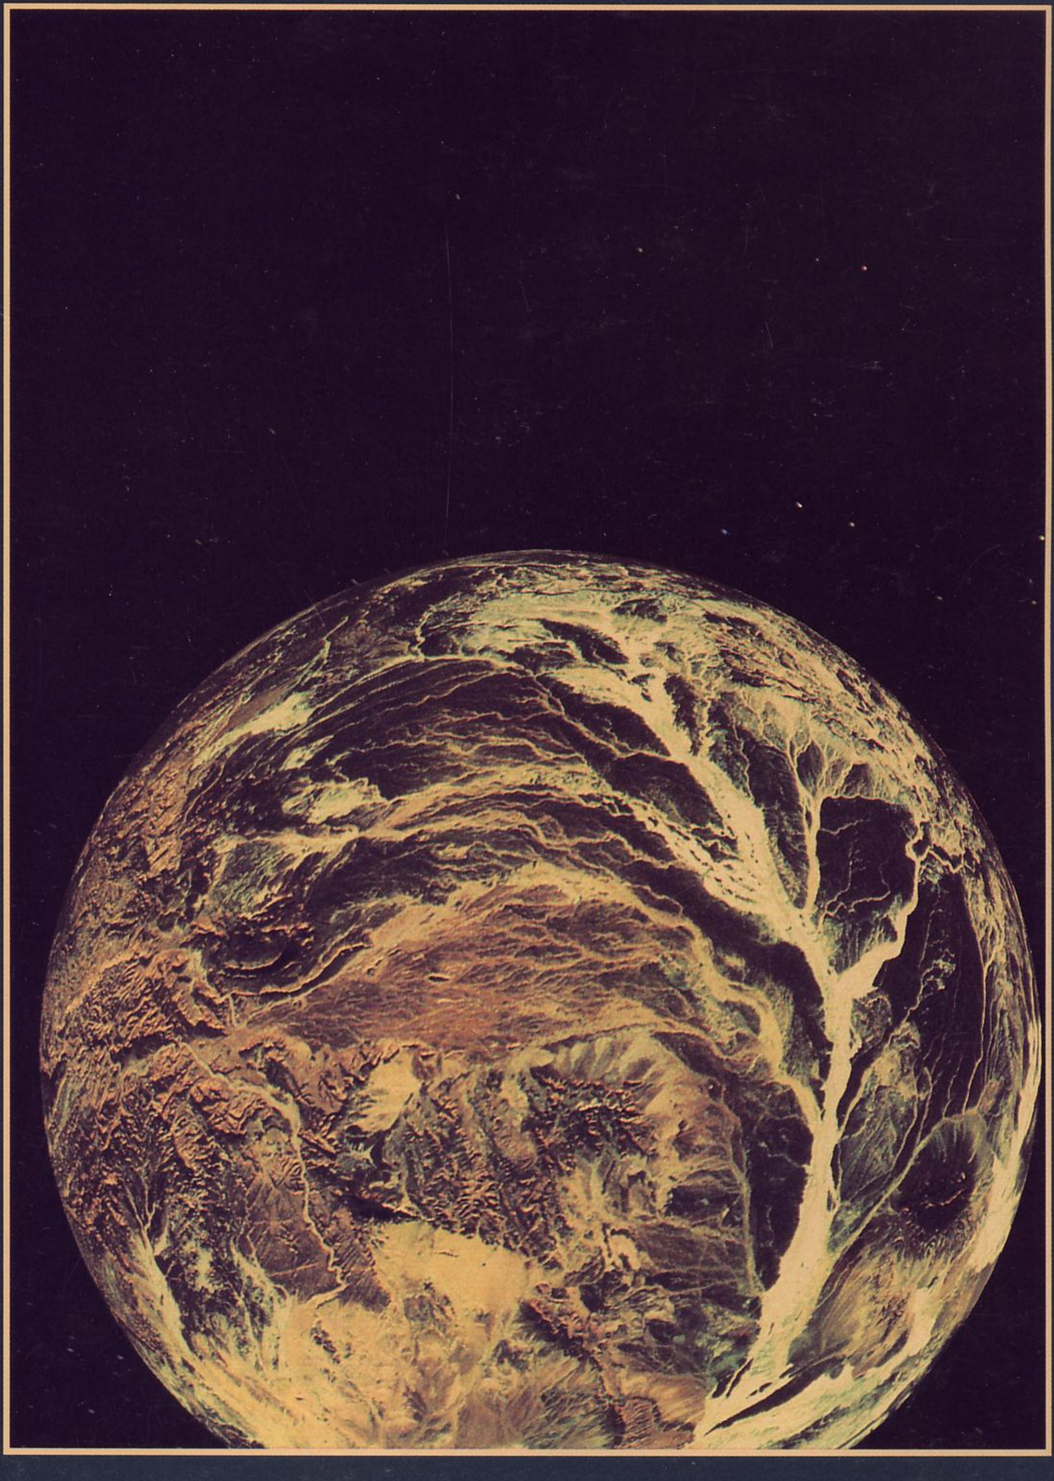

IN THIS ISSUE

Research is at least half of what Cray is all about. Yes, it's true that the company produces close to 25 new supercomputers in a year, and approximately half of the company's employees are involved in system manufacture. But when it gets right down to it, Cray's sterling research and development efforts are what breathe life into the rest of our endeavors.

In July of this year, the company advanced the state-of-the-art in general purpose supercomputing. The CRAY X-MP/48 and its associated peripheral equipment offer new heights in overall supercomputing performance. The four-processor, eight-million word system is innovative in its own right, incorporating advanced technology packaged in a new way. In addition, the Cray development group introduced new single-processor CRAY X-MP systems, providing higher performance levels for lower cost than ever before. New peripheral products announced concurrently ensure that the computational power of these new systems is not held captive by I/O limitations.

Cray Research remains committed to the development of supercomputers, now and in the future. The recent product introductions are evidence of that dedication to providing powerful computing.

Our article describing the CRAY X-MP Series of Computer Systems begins on page 2. Other articles in the issue include a summary of multitasking techniques used for a Monte Carlo algorithm, and an expose on research being done at the university level with a CRAY computer. Our regular departments are packed with interesting pieces about SIGGRAPH '84 and the CRAY as animal matchmaker. We hope you'll enjoy this fall 1984 issue of CRAY CHANNELS.

On the cover a desert planet is shown, which was created by processing and projecting onto a sphere information from the Landsat Earth-imaging satellite. Three separate light frequency bands were combined into this composite image from *Desert Planet Simulation*, a computer-generated film sequence. The sequence was generated on a CRAY X-MP/2 at Cray's Men Heights facility. Each frame took approximately 45 seconds to generate at a resolution of about 2000 by 1500 pixels per frame. (Photo © 1984, Geometric Productions, Berkeley, California)

CRAY CHANNELS

A Cray Research, Inc. publication

Volume 6, Number 3

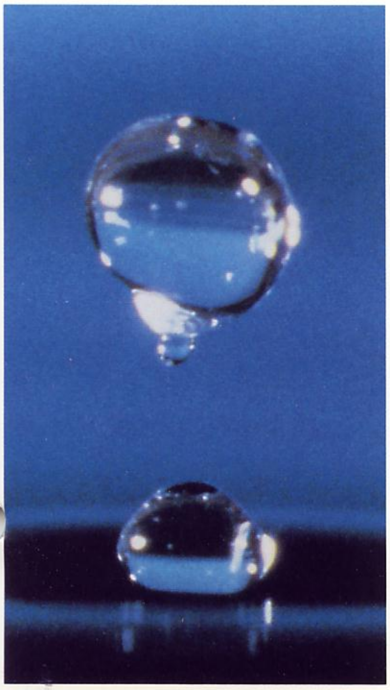

10

FEATURES

DEPARTMENTS

Credits

Editor
Carol Smith

Contributing Writer
Ken Jopp

Design
Cindy Erickson

CRAY CHANNELS is a quarterly publication of Cray Research, Inc., 608 Second Avenue South, Minneapolis, MN 55402. It is intended for users of Cray computer systems and others interested in the company and its products. Subscription requests, feature story ideas and news items submitted to the editor at the above address are welcomed.

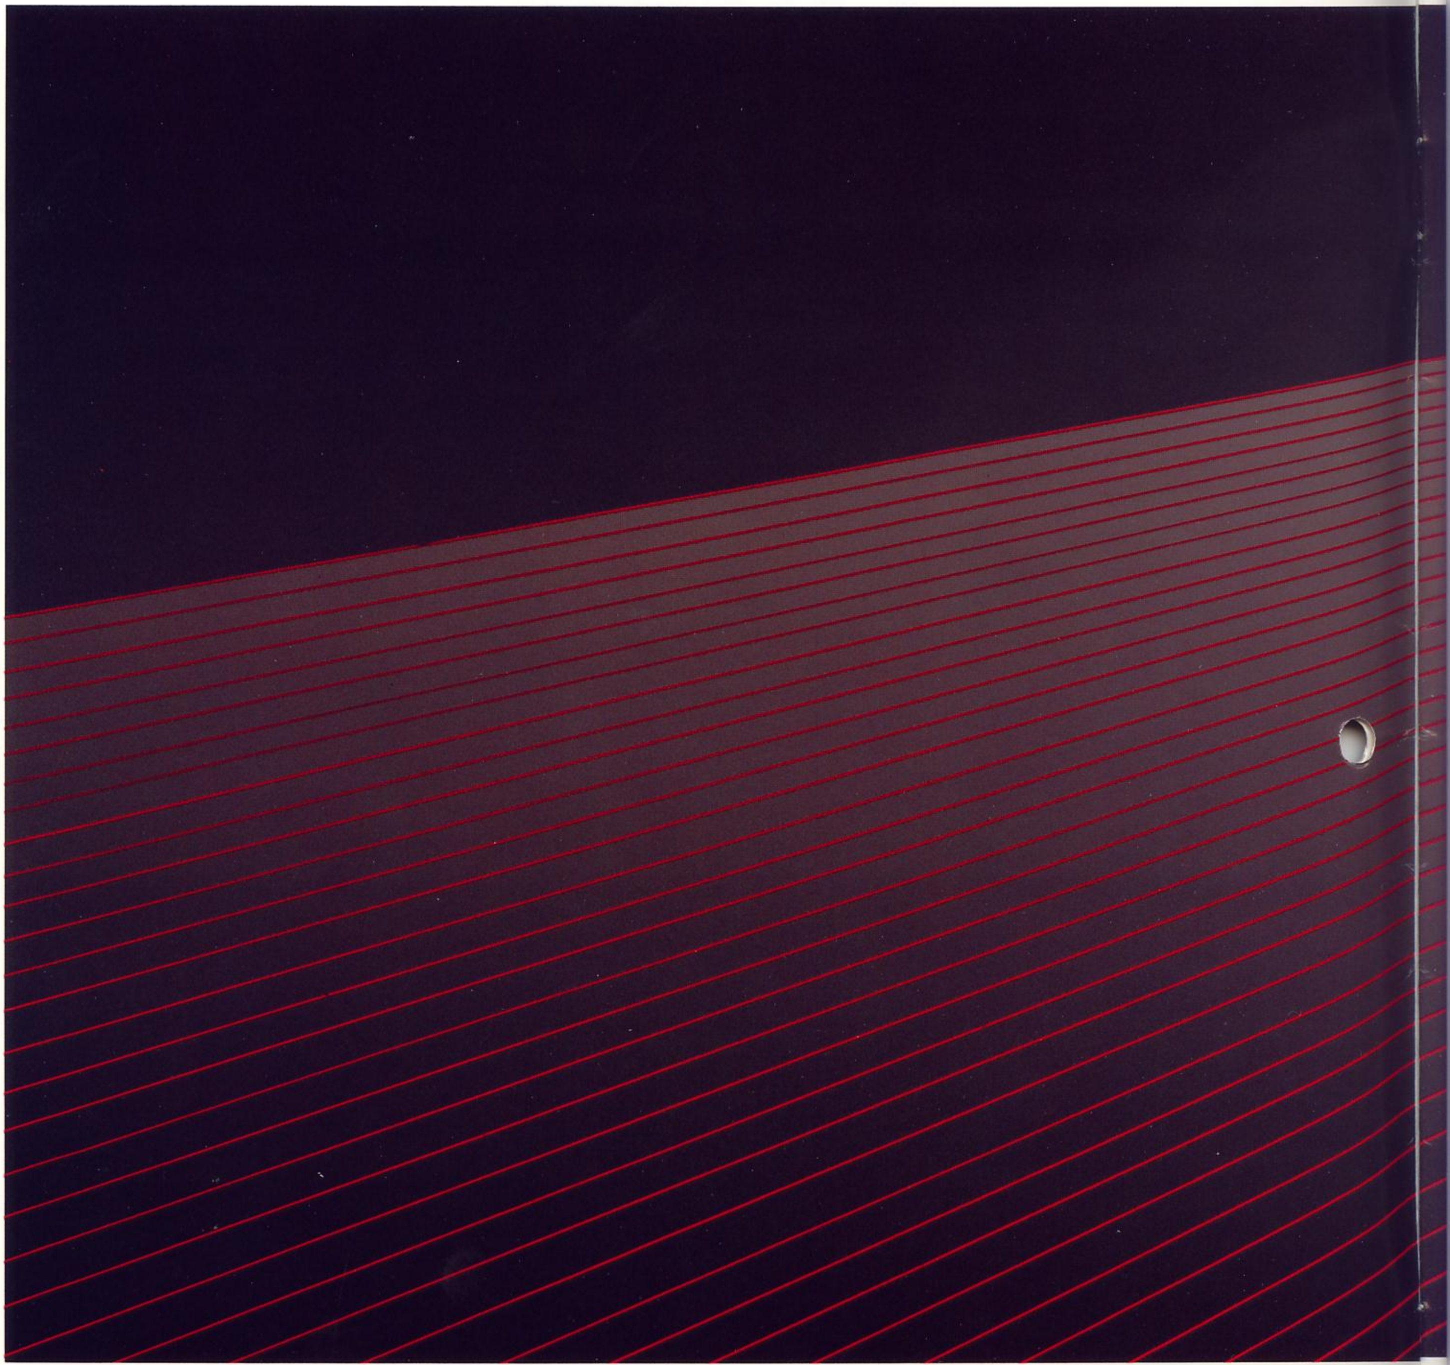

Often, less is more. But when it comes to computational power for very large-scale problems, more is often not enough. Taking that challenge to heart, the Cray Development Group headed by Steve Chen developed new versions of the CRAY X-MP supercomputer. The X-MP line now incorporates one, two, four or eight million words of central memory and one, two or four central processing units (CPUs) in a single mainframe.

The expanded and enhanced CRAY X-MP Series includes six models based on the successful X-MP CPU introduced in 1982. Single-processor CRAY X-MP systems stand alongside the field-proven dual-processor X-MP computers. The CRAY X-MP/48 computer is the latest evidence of Cray Research's commitment to leadership in supercomputer performance. The system is configured with eight million words of central memory and four CPUs.

Introducing the enhanced CRAY X-MP Series of Computer Systems

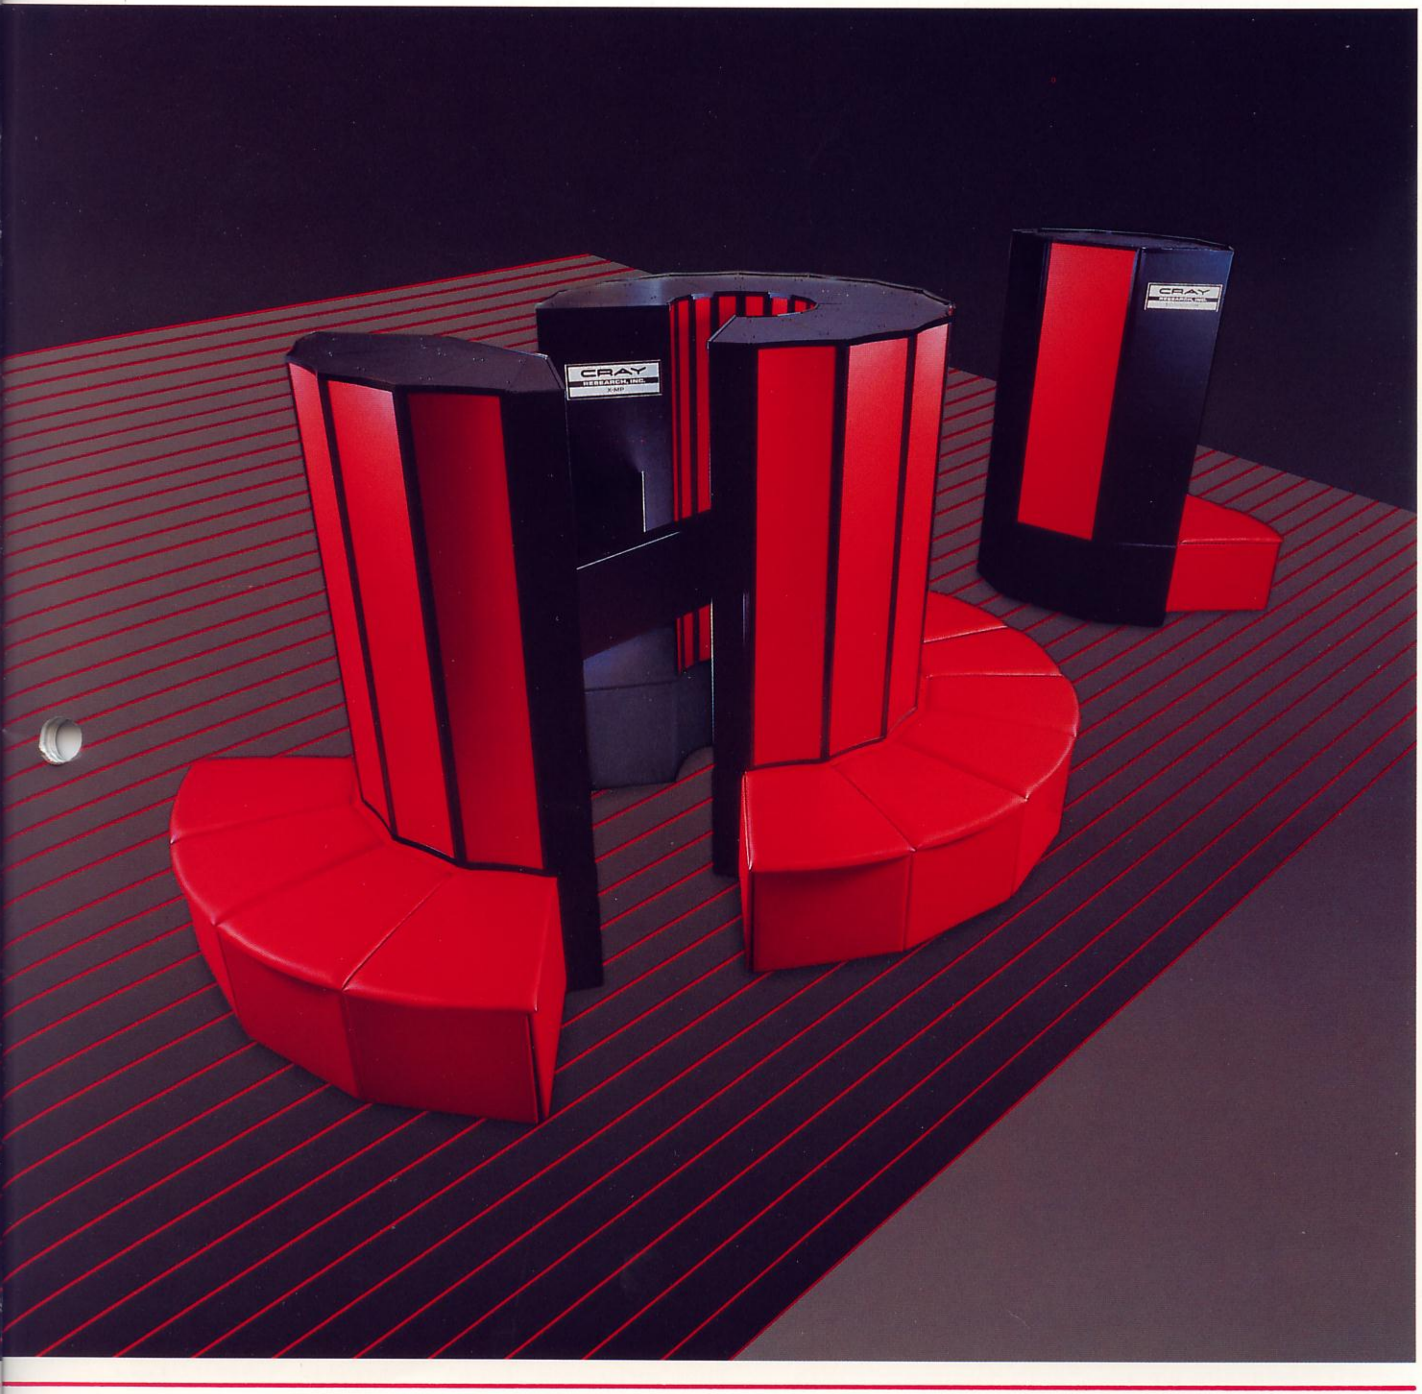

CRAY X-MP/48 overall performance typically can be up to ten times that of the CRAY-1. Besides the new mainframes, Cray Research also announced new input/output projects. The Solid-state Storage Device (SSD) now ranges in size from 256 to 1024 Megabytes (Mbytes) of very fast random-access secondary MOS memory. And the new DD-49 disk drive complements the increased computational power of the X-MP Series computers.

The CRAY X-MP multiprocessor configurations allow the user to employ multiprogramming, multiprocessing and multitasking techniques. As a result, job turnaround time can be reduced significantly. Programs that can be multitasked can also take advantage of vector and scalar processing capabilities, offering unusually fast compute speed and turnaround time. For instance, in certain applications, by combining vector processing and multitasking techniques, the CRAY X-MP/48 can offer a total speedup over conventional scalar processing ranging from 34 to 72 times. The speedup over scalar processing with CRAY X-MP/2 systems ranges from 18 to 38 times.

Steve Chen explained, "The full CRAY X-MP configuration — the CRAY X-MP/48 with the 1024-Mbyte SSD — is the first parallel scalar/vector machine with sufficient power for addressing large general 3-D simulation problems. We believe that CRAY X-MP computers will help scientists find the key to solving these problems."

Software developed for the CRAY-1 or any CRAY X-MP model can run on all other models of the X-MP Series, thus protecting user software investment. Over 200 applications programs can operate on X-MP computers. These programs are geared to solve problems in industries such as petroleum, aerospace, automotive and nuclear research.

CRAY X-MP computers can be integrated easily into a user's existing computer environment. Hardware and software interfaces for several manufacturers' equipment are available for the CRAY X-MP computers.

Chen commented, "I believe that with this announcement, we have a complete product line to offer for various user requirements in the coming years. Once again, we are able to demonstrate both our company's commitment to lead in the development of high performance machines, and our ability to deliver a real product, thus setting new performance/cost standards in this industry."

Chen acknowledged, "Others also are developing parallel processors. Multiprocessor architectures will become more prevalent as time goes on. Development of new algorithms and applications will be important as more parallel machines become available. The CRAY X-MP Series of Computers gives researchers the opportunity to explore future parallel algorithms to achieve even higher performance, while allowing the existing user's application programs to run faster without modification."

CRAY X-MP/48

The new top-of-the-line CRAY X-MP/48 is the most powerful computer system available today. The system is configured with four identical CPUs that share an eight-million-word ECL bipolar central memory arranged in 64 interleaved banks. Maximum memory bandwidth is 16 times that of the CRAY-1. The CPUs each have a 9.5 nanosecond (nsec) clock and a memory bank cycle time of 38 nsec. Each CPU on the X-MP/48 offers gather/scatter and compressed index vector instructions, which allow for the vectorized processing of randomly organized data. The mainframe is arranged in 12 columns in a 270° arc and requires the same electrical power as the CRAY-1.

The I/O Subsystem (IOS), which is an integral part of all CRAY X-MP computers, also contributes to the X-MP/48's outstanding performance. It acts as a data concentrator and data distribution point for the CRAY X-MP mainframe. The I/O Subsystem offers parallel disk drive capabilities, I/O buffering for disk-resident and buffer memory-resident datasets, on-line tape handling and efficient front-end system communication. Up to eight million words of buffer memory can be configured on the IOS, enabling faster and more efficient data access and processing by the CPUs. The IOS is standard on all CRAY X-MP models.

CRAY X-MP/2 models

The middle-range CRAY X-MP/22 and X-MP/24 computers, first introduced in 1982, are now arranged in six columns and are half the size of their predecessors. They feature two CPUs and two or four million words of ECL bipolar central memory arranged in 16 or 32 interleaved banks, respectively. Maximum memory bandwidth is eight times that of the CRAY-1. As with the X-MP/48, the CPUs can operate independently on different programs or can be harnessed together to operate on a single user program.

CRAY X-MP/1 models

The entry-level CRAY X-MP/11, X-MP/12 and X-MP/14 systems combine a single CRAY X-MP CPU with one, two or four million words of static MOS central memory, respectively. On the four-million-word system, memory is arranged in 32 interleaved banks, and on the one- and two-million word systems it is arranged in 16 banks. Maximum memory bandwidth is four times that of the CRAY-1. The CRAY X-MP/1 mainframe is arranged in six columns and is configured with the standard IOS.

Solid-state Storage Device

With 1024 Mbytes of MOS memory, the new model SSD is four times larger than its predecessor. The SSD offers transfer rates of 100 to 2000 Mbytes/sec (depending on channel configuration), access time of less than 50 microseconds and large storage capacity. It acts as a fast-access disk device for large datasets generated and manipulated repetitively by user programs, and also can be used by the system for temporary storage of system programs.

On some applications, a factor of three to six times speedup of program execution time has been observed with the SSD. The SSD has demonstrated its utility and reliability in a variety of applications such as weather forecasting, seismic processing, reservoir modeling, fluid flow computation and finite element analysis.

DD-49 disk drive

The new DD-49 disk drive has a 1200-Mbyte capacity and a 10 Mbyte/sec transfer rate. The unit offers twice the capacity and 2.5 times the effective transfer rate and access speed of Cray's DD-29 disk drive. The DD-49 is typically configured with the CRAY X-MP/48.

Product availability

The new CRAY X-MP/48 has been installed and operating since June at Cray's Mendota Heights computer center, where it is being used for software development, application demonstration and benchmarking. The first installation of an X-MP/48 at a customer site will take place in the fourth quarter of 1984. CRAY X-MP/2 and X-MP/1 systems are in production and being delivered now. □

Product highlights

CRAY X-MP mainframe highlights:

- Four processors sharing eight million words of ECL bipolar memory with the X-MP/48
- Two processors sharing two or four million words of ECL bipolar memory with X-MP/22 or X-MP/24 systems, respectively
- One processor with one, two or four million words of MOS memory, on the CRAY X-MP/11, 12 or 14 computers, respectively
- 9.5 nsec clock cycle
- 38 nsec (on X-MP/48 and X-MP/2) or 76 nsec (on X-MP/1) memory bank cycle time
- SECDED memory protection
- Four parallel memory ports per processor
- Flexible hardware chaining for vector operations
- Second vector logical unit
- Gather/scatter and compressed index vector support on CRAY X-MP/48
- Flexible processor clustering for multitasking applications
- Dedicated registers for efficient interprocessor communications and control

Input/output highlights:

- 6-Mbyte, 100-Mbyte and 1000-Mbyte channels
- I/O Subsystem with:
  - Parallel disk streaming capabilities, one controller per disk
  - I/O buffering for disk- and tape-resident datasets
  - Software support for parallel disk striping
  - Buffer memory-resident datasets
  - High-performance disk drives
  - On-line tape handling
  - Front-end system communication with IBM, CDC, Data General, DEC, Honeywell, Sperry and NSC network adapters

SSD highlights:

- Memory size ranging from 64 to 1024 Mbytes
- Support for:
  - One or two 1000-Mbyte channels for linkage to CRAY X-MP/48
  - One 1000-Mbyte channel for linkage to CRAY X-MP/2
  - One 100-Mbyte channel for linkage to CRAY X-MP/1
- SECDED memory protection
- Software support to allow existing programs to run without program modification
- Optional direct data path between SSD and IOS

Multitasking: a vectorized Monte Carlo algorithm on the CRAY X-MP/2

Y. Chauvet
Commissariat a l'Energie Atomique, France

This article provides a brief comparison between CRAY-1/S and CRAY X-MP hardware and describes the multitasking tools available with FORTRAN. CRAY-1/S and CRAY X-MP performances are compared on a vectorized Monte Carlo algorithm for neutron transport problems. The main characteristics of the algorithm and two possible methods for its parallelization are described. We conclude that tasks as short as 0.05 seconds can realize reasonable speed improvement when multitasked.

A brief system comparison

From a user's point of view, the CRAY X-MP offers at least three major improvements over the CRAY-1/S:

- A single job may access two CPUs sharing the same central memory (two or four million words phased in 16 or 32 banks).
- Each CPU has four independent paths to memory which allow a speed of 75 MOPS (million of operations per second) not counting memory references, on any diadic operation between vectors stored in memory. This is three times faster than the CRAY-1/S.
- The clock period is 9.5 nanoseconds (nsec) instead of 12.5 nsec for the CRAY-1/S, thus allowing scalar calculations to run 1.3 times faster.

Task and logical CPU concepts

Before describing the main routines supplied by Cray for implementing multitasking with FORTRAN, we must introduce the concept of tasks and logical CPUs.

Typically, a task is a part of a program that can be run in parallel with some other parts of the program. The number of tasks is not necessarily limited by the number of available physical CPUs. The attachment of a physical CPU to a task is done in two steps implementing the logical CPU concept. First the library scheduler assigns a task to a logical CPU and then COS assigns a logical CPU to a physical CPU. As a result, there is no guarantee that independent tasks will execute in any particular order.

FORTRAN multitasking tools

The primary multitasking tools can be classified into three categories. They are: task creation, critical sections monitoring and task synchronization mechanisms. A brief description of each follows.

Task creation

In practice, a task will be built by a special call to a subroutine. The actual beginning of its execution is defined by the system according to the global scheduling of the physical CPUs.

Task initiation is done by the following FORTRAN statement: CALL TSKSTART (*task array, name, param-list*). *Task array* is a control array built by the user job in order to identify the task, *name* is the name of the called subroutine and *param-list* is a list of parameters transmitted by their addresses to the called subroutine.

Synchronization can occur at the completion of a task by using the TSKWAIT routine. The identifier of a task given in the *task array* can be retrieved by calling the TSKVALUE routine.

Critical sections monitoring
A critical section is a segment of code that updates a memory area shared by several tasks. In such a section, in order to avoid conflicts, special LOCK VARIABLES are used through two routines: LOCKON and LOCKOFF.

First, a variable must be specified as being a lock variable by the statement CALL LOCKASGN (*name*). Afterward, the statement CALL LOCKON (*name*) is added at the beginning of the critical section in order to set the lock variable. The statement CALL LOCKOFF (*name*) is added after the last statement of the critical section to clear the lock variable when the shared memory is free again.

During execution, two different situations may occur, according to the state of the lock variable. If it is set, then the task encountering the CALL LOCKON (*name*) statement is suspended until another task clears the lock variable. Otherwise, the task sets the lock variable before advancing into the critical section and clears it at the very end of the section with the CALL LOCKOFF (*name*) statement.

Task synchronization mechanisms
It may be useful to synchronize task operation, as opposed to allowing tasks to wait for each other. To enable synchronization, special EVENT VARIABLES can be used in order to suspend or to allow the processing of a task according to the event variable state: Posted or Cleared.

First, a variable must be defined as being an event by the CALL EVASGN (*name*) statement. It is then initialized by using either the CALL EVPOST (*name*) statement or the CALL EVCLEAR (*name*) statement.

During execution, two different situations may occur, according to the state of the event variable. If it is posted when a task issues a CALL EVWAIT (*name*) statement, then the task continues. Otherwise, it will be suspended until another task posts the event.

Please realize that this short description of the primary multitasking tools is not exhaustive, but should be sufficient to understand the following examples.

Two examples of parallelization

The purpose of our algorithm is to study the behavior of a sample of particles moving into a two-dimensional axisymmetric Lagrangian geometry by using a Monte Carlo method.

The technique used to vectorize that type of algorithm will not be described here. Let us only mention that it is very similar to the "Stack-events" method introduced by F. W. Bobrowicz (1) and described in (2) for the particular case of neutron transport problems.

Two methods for parallelizing the algorithm will be presented here. The first one, which is most natural, allows us to measure the overhead inherent in multitasking when tasks become too small. The second, which is more CRAY X-MP/2 oriented, allows one to obtain the best speedup factor for this type of algorithm.

In the first method, the sample is split into n parts that can be handled in parallel by n tasks, as seen in Figure 1.

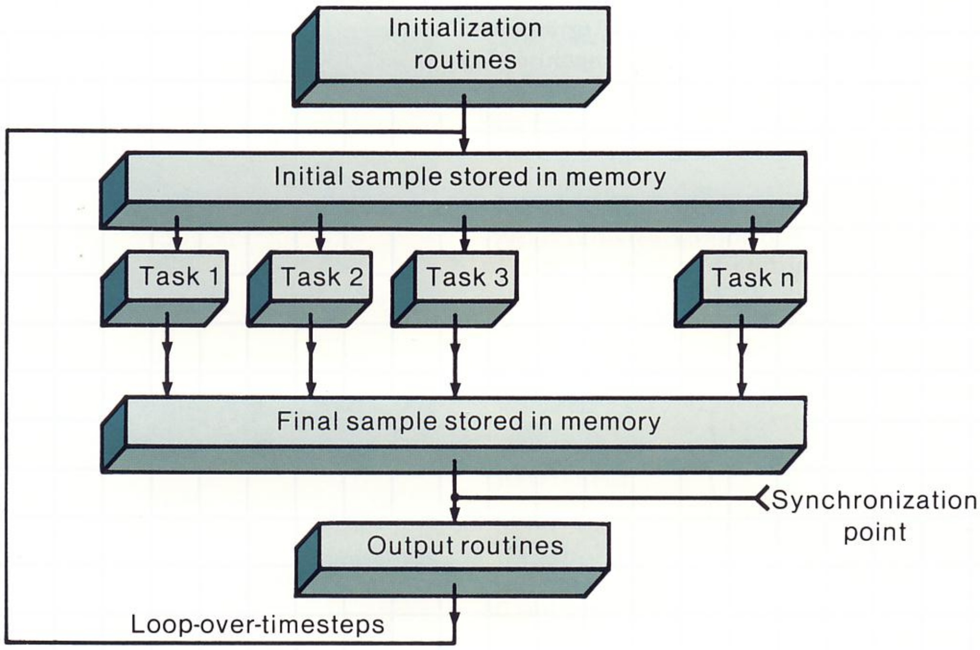

Figure 1.

Here the number of tasks could be greater than the number of physical CPUs and synchronization is done at completion of all tasks before calling the output routines and starting a new timestep. During successive timesteps, the sample size decreases because particles leak out of the system. Consequently, task lengths become shorter and overhead appears.

With the second example, we try to get the best speedup possible from the two CPUs. The following statement sequence illustrates the approach taken:

```
                    Initialization

      CALL TSKSTART (Tskarray, COMPUTES, param-list)

C                   Loop-over-timesteps
      DO 1 I = 1,NDT
      CALL EVPOST (START)
      CALL COMPUTES (param-list)
      CALL EVWAIT (COMPLETE)

                    Outputs
  1   CONTINUE
```

In this example, COMPUTES is the subroutine that studies a sample part that corresponds to the tasks of the first example. First, *before* the loop-over-timesteps, we define a task that will process the first half of the sample by calling COMPUTES. This task is synchronized by event variables which allow it to start and which give some information upon completion. Next we begin the loop-over-timesteps. This handles task synchronization and calls, by an ordinary statement, to the routine COMPUTES for the second half of the sample.

Calculation reproducibility

A characteristic of Monte Carlo method is to use random numbers generated by *recurrent* functions (RANF, for instance, on Cray or CDC machines). There is no guarantee that tasks will execute in any particular order. If one does not take care, *various runs will always give different results*!

To become independent of the scheduling, random number generators are considered as critical sections. They are used according to a lock variable in the following way:

```
      CALL TSKVALUE (NUTASK)              (Task identification)
      CALL LOCKON (LOCKRANF)              (Post lock variable)
      CALL RANSET (SEED (NUTASK))  }
      DO 1 I = 1,N                 }
  1   R(I) = RANF()                }      (Critical section)
      CALL RANGET (SEED (NUTASK))  }
      CALL LOCKOFF (LOCKRANF)             (Clear lock variable)
```

In this way, each task identified by a number has its own sequence of random numbers; two tasks running concurrently will never use the RANF function at the same time.

Code adaptations

Adaptations to the algorithm for multitasking were easy and required little time. Expression of parallelism through the multitasking library is very convenient even if it is not as perfect as theorists would like. Before going on, it is interesting to mention a mistake I made using the TSKSTART statement.

Question: Why would the two sequences that follow not produce the same results?

```
      DO 1 I = 1,NCPU
1     CALL COMPUTES (I)


      DO 1 I = 1,NCPU
1     CALL TSKSTART (Tskarray, COMPUTES, I)
```

Answer: With the first sequence, routine COMPUTES and incrementation of loop index I, are always done alternatively. In the second, it is very likely that the loop will be completed before the beginning of the first task and consequently, the value stored at the address of I will always be NCPU.

This example suggests that if we mistakenly assume that FORTRAN is always easy to use, it could lead us to new difficulties in multitasked jobs — and debugging such jobs surely will not be easy.

Measured results

The results compare calculation speeds between the CRAY-1/S using CFT 1.10, and the CRAY X-MP/2 with CFT 1.11 without multitasking. On the few scalar parts of the code we measured a speed ratio of 1.2 to 1.3. For vectorized routines, this ratio varies according to the vector lengths and the type of computation. Typical ratios are as follows:

Elements per vector	Ratio range
30	1.3 - 1.7
64	1.6 - 1.9
256	1.7 - 2.0

The ratio for the entire algorithm is 1.6 with CFT 1.11 and becomes 1.7 with CFT 1.13 on the CRAY X-MP/2.

The second set of results concerns speedup factors due to multitasking. The speedup factors measured on a dedicated machine are compared to a theoretical one calculated in the following manner:

If to is the time spent in non-multitasked parts of the algorithm and tp is the time spent by tasks running in parallel, then theoretical speedup is:

$$s = (to + tp)/(to + tp/2)$$

Over a timestep and for a realistic sample, we obtain the results listed in Figure 2.

	Number of tasks	Total CPU time (sec)	Total elapsed time (sec)	Measured speed-up	Theoretical speed-up	Difference
	2	2.53	1.303	1.942	1.953	0.6%
Example 1	8	2.60	1.342	1.937	1.962	1.3%
	16	2.65	1.373	1.930	1.955	1.3%
Example 2	2	2.67	1.364	1.957	1.961	0.2%

Figure 2.

Notes

For these four runs, theoretical speedup factors are somewhat different because the random sequences employed were not the same.

The four measured speedup factors are very close because, as one will see with the next results, overhead due to multitasking is negligible and speedup is only algorithm-dependent.

The last results are overhead evaluation. It is very easy to observe when overhead appears with this algorithm, as task lengths decrease at each timestep. To help determine overhead, we conducted two very similar runs on the CRAY X-MP/2. The second was modified to replace the task initiation statements with ordinary call statements. (Of course other statements concerning lock variables and synchronization mechanisms have been deleted.) Then, by evaluating the CPU time ratio between these two runs according to the task length, it was very easy to observe the reasonable threshold below which multitasking is not recommended.

Overhead with respect to task length

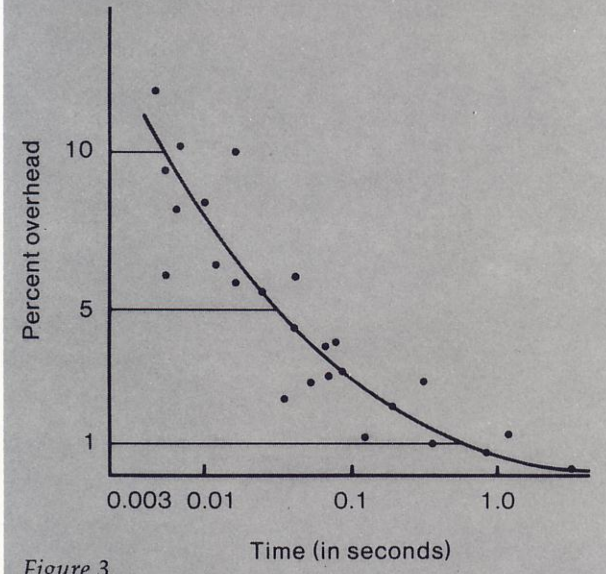

Figure 3.

As one can see by the curve in Figure 3, for this multitasking example, *if task length is greater than 0.05 second, overhead is less than 5%.* If task length is between 0.01 and 0.05 second, overhead is between 5 and 10%.

Conclusion

The multitasking tools available with FORTRAN under COS 1.13 are very convenient and are well suited to users' needs, though for some specific applications CRAY Assembly Language can be used to achieve higher performance for programs with very small task length. It is also important to keep in mind that FORTRAN must be used wisely in the multitasking environment.

For task length exceeding 0.05 seconds, measured speedup factors due to multitasking are very close to theoretical factors if jobs are running in a dedicated environment. Of course, in a multiprogramming environment with only two CPUs, the speedup factors often will be close to 1, unless all of memory is used by the job.

The CRAY X-MP/2 is a high-speed computer very well suited to the kind of production code with which we have experimented. In fact, with both vectorization and multitasking, the X-MP/2 happened to be 3.4 times faster than the CRAY-1/S. □

Acknowledgement

I would like to thank Mr. Azar and Mr. Guidici, both of Cray Research France, and Mr. Meurant and Mr. David, both of CEA, for many fruitful discussions and for their support at Mendota Heights, Minnesota, USA, where the tests were done.

References

Bobrowicz, F. W., *Vectorized Monte Carlo photon transport*, LA 9752, Los Alamos National Laboratory, May, 1983.

Chauvet, Y., Verwaerde, *Performance comparisons between CRAY-1/S and CYBER 205 computers*, Rapport CEA n. 2383, December, 1983.

Bubbles, drops and a whole lot more

Chemical engineering and materials science at the University of Minnesota

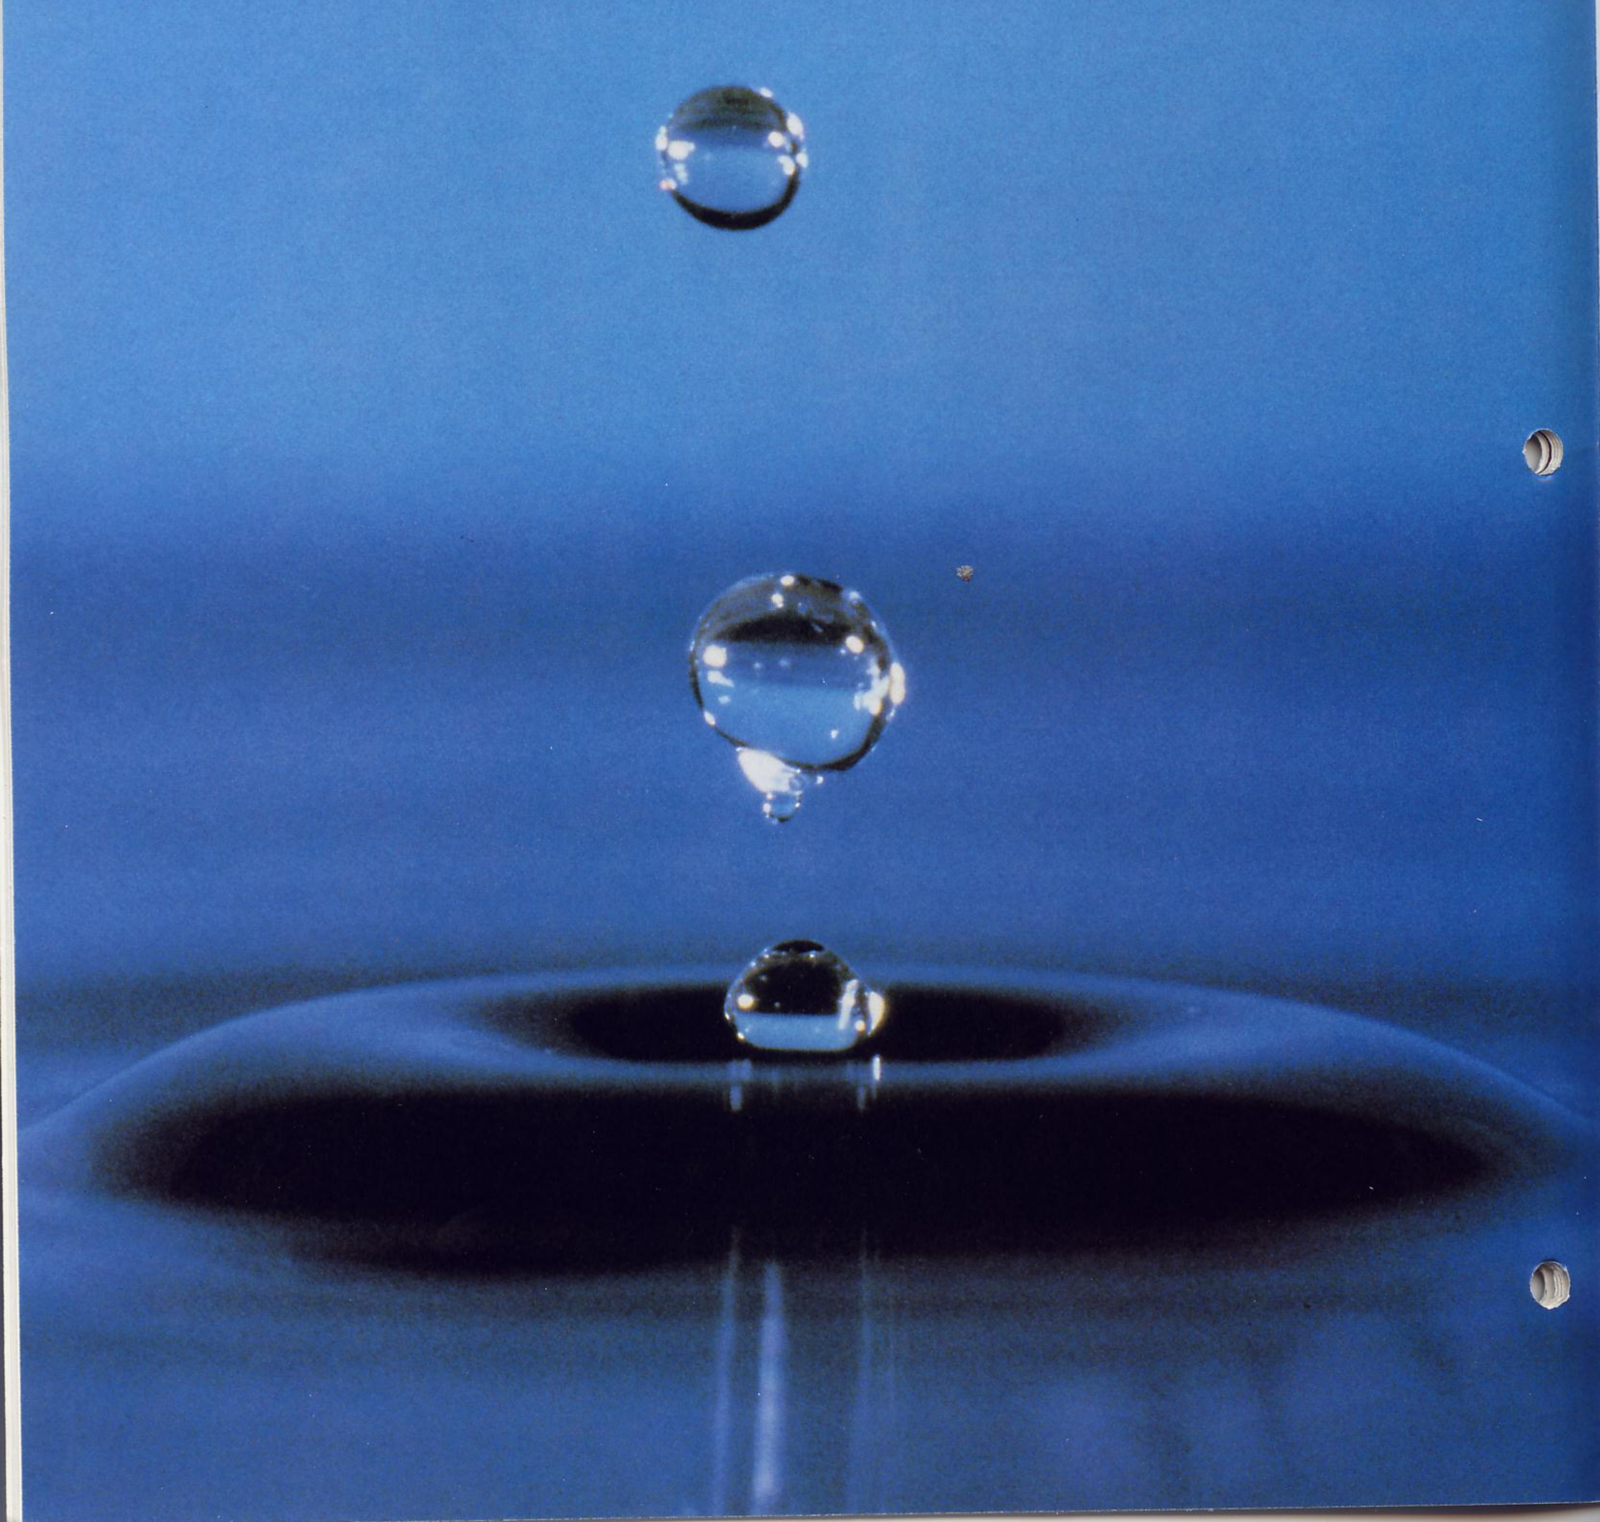

Chemical engineering and materials science are disciplines frequently called upon to help design or improve commercial and industrial operations. Scientists working in these disciplines study chemical and physical phenomena at both the micro- and macroscopic levels. The knowledge they gain can then be used in the creation of new commercial products and the innovation of new manufacturing processes.

Scientists working in the Chemical Engineering and Materials Science Department at the University of Minnesota (U of M) use the CRAY computer installed at the University extensively in their research. The heaviest users are a group of researchers studying the properties and behavior of fluids in such fields as fluid dynamics, fluid microstructure and microfluid mechanics. Their work is aimed at understanding the principles governing fluids at solid surfaces, fluid interfaces and fluid flows and diffusions. Simulating fluid behavior via computer has come to play a role equal in importance to the researchers' theories and laboratory experiments. Practical applications of this work include improving the accuracy of petroleum reservoir simulations and increasing the efficiency of coating flow processes. Coating flows are used in such areas as the manufacture of adhesive tapes, photographic film and magnetic recording media.

Dr. H. Ted Davis, head of the U of M's Chemical Engineering and Materials Science Department, offered some insight into the relationship between the department's two disciplines: "The chemical engineer tends to be interested in chemical processes, such as the catalytic process, the reaction process, the coating flow process and so forth. He is interested in how to control those processes and how to apply them to the creation of finished products." Davis explained that the materials scientist, on the other hand, is interested in creating novel materials and tends to study the substance itself, things like its composition and its electronic and mechanical properties.

Davis concluded, "The person studying the process of, say, coating a substrate with a plastic film, is better served if he understands the science of the material being deposited. The person attempting to invent novel materials is well served if he understands the processes by which materials are created. As you can see, there's a lot of overlap between the two disciplines."

Specialized analytical approach

Dr. Davis and Dr. L.E. "Skip" Scriven collaborate in directing the department's fluids research. Davis, Scriven and their collaborators have developed their own approach to solving problems in fluid behavior on the CRAY. Their method is an adaptation of finite element analysis and can be divided into several steps. For many purposes it involves subdomaining, which divides a complex system into manageable parts, including adaptive subdomaining, which appropriately sizes the parts for the given problem. "You don't know what size you want your subdomains to be until you have the problem solved, so the subdomains themselves become unknowns," explained Scriven. "In the jargon most people use, they say they need to calculate a mesh or a grid to solve the problem. We're incorporating the calculation of the grid needed to solve the problem into the problem itself, so we get it all done at once."

Subdomaining is followed by finite element basis functions and isoparametric mapping, which accurately maps irregular curved shapes onto a grid of squares through the use of variable scaling. "This is what the Flemish cartographer Gerhardus Mercator was trying to do," said Scriven. "He tried to put the Earth on a flat surface. But it got horribly distorted because he was trying to do it all with the same scale. We use different scalings for different parts, so the overall map varies in scale from place to place. Before the advent of computer-aided mathematics this wasn't possible for real systems."

Mapping is followed by Newton iteration and vectorized calculations for the coefficients of the basis functions. "This is the 'Minnesota brand' of analysis," said Scriven. "I don't know of anyone else who's doing all of these things in fluids research."

Bubbles and drops

Fluids research at the U of M grew out of initial studies in bubble and drop dynamics. Examples of common bubbles and drops include soap bubbles and raindrops. The chemical and physical laws governing their behavior apply to a wide range of disciplines — literally from nuclear physics to astrophysics. Discovering the principles governing bubble and drop dynamics, therefore, aids researchers in a variety of fields.

"A bubble or a drop is a relatively compact blob," explained Scriven. "It is a sphere, or its topological equivalent, of one fluid phase surrounded by another." In the case of a raindrop, the fluid surrounding the drop is the atmosphere itself. A drop is a relatively dense fluid surrounded by a less dense fluid. On the other hand, a bubble is a relatively less dense fluid surrounded by a denser one.

Principles of bubble and drop dynamics are applied to nuclear physics in a model of nuclear fission which treats very heavy atomic nuclei as rotating liquid drops. According to this theory, the properties determining how drops of water fuse and oscillate when brought into contact are responsible for the same behavior in very heavy atomic nuclei. Understanding this behavior in common drops can lead to predictions of the conditions under which very heavy atomic nuclei will fission or be created by fusion of smaller nuclei.

One of the most significant forces governing drop dynamics is surface tension. It determines the shapes of fluid drops and is a major factor determining how readily two fluids will mix. Surface tension stems from the expression of intermolecular forces within the zone separating the inside and outside of a drop. It arises because the net attractive forces among molecules parallel to the plane of a drop's surface are always greater than the net attractive forces perpendicular to that plane. Surface tension arising from the intermolecular attraction among water molecules, for example, prevents raindrops from dispersing into mist as they fall through the atmosphere. It also explains water's ability to penetrate some fabrics but not others. Surface tension magnitude varies proportionately with the density gradient between the interior of a drop and the fluid outside it.

The molecular theories enabling scientists to understand the mechanisms of surface tensions and interfacial phenomena were too complicated to solve before the advent of large computers like the CRAY. Now, with computer-aided finite element analysis, the nonlinear equations of the molecular theory of surfaces routinely yield the molecular insights necessary for understanding the principles governing surface tension behavior.

The principles governing drop dynamics also find application in astrophysics, since stars are recognized as exhibiting some of the properties of fluid drops. Stars do not exhibit surface tension because their matter trails off gradually into space. But they can be accurately modeled as self-gravitating drops. "A star does not have an appreciable surface tension analog, but it has gravity," Scriven said. "Electrostatic attraction and gravity are governed by the same inverse square law, so computing predicted behavior for the two phenomena is virtually the same." Treating stars as rotating, self-gravitating fluid drops has led to a theory to explain the formation of binary stars. A binary star is a system of two stars that revolve around each other due to their mutual gravity. Under certain conditions, a rotating drop will distort into a two-lobed "dumbbell" shape, then split into two separate drops. It is now thought that the principles governing this behavior in common drops might also govern the formation of binary stars from single stars.

The study of drop dynamics offers some insight into an aspect of meteorology that fascinates most people — lightning. When exposed to an electric field, a fluid drop sometimes forms conical ends and "spits" small highly-charged daughter drops. In other cases, a charged drop elongates and fissions into two drops. Under other conditions, a charged spherical drop at rest will break up spontaneously.

These phenomena were observed early in this century but no method was available to accurately quantify these behaviors of charged drops. Recently, the phenomena were quantified with the help of the CRAY, by relating a drop's charge to the amplitude of its oscillation.

This knowledge has potential applications in meteorology since it is known that raindrops in thunderclouds carry electrical charges. Understanding how charged drops break-up may someday help to explain how electrical charges become separated in thunderclouds. This, in turn, will help explain the mechanism responsible for lightning.

Knowledge of the physics and chemistry of surface tension, along with other properties of fluids, can be applied to improving a variety of industrial processes. Two areas the U of M's researchers are currently exploring are improving the efficiency of petroleum reservoir simulations and fluid coating techniques. The remainder of this article focuses on these applications, which have become increasingly practicable with progress made possible by the University's CRAY.

Petroleum recovery

Oil companies stand to gain considerably from a better understanding of drop dynamics. A better understanding of surface tension, for example, can be used to predict the behavior of fluids trapped in underground rock, a situation oil companies deal with every day. Surface tension dictates the miscibility and flow rates of such fluids. These factors must be taken into account in computer simulations of petroleum reservoirs. Simulations are commonly used to assess the probable effectiveness of petroleum recovery operations.

A common misconception is that petroleum reservoirs are free-standing pools of underground oil. But nearly all underground oil is trapped inside the chaotic capillary-like porespaces of porous rock, typically sandstone, or carbonate rock such as limestone or dolomite. Recovering petroleum from these underground rock "sponges" generally involves three kinds of methods. Primary recovery methods rely on geologic pressures to push the petroleum up a drilled well to the surface. But underground pressures usually squeeze up only a small fraction of a given reservoir's content. Forcing water through a reservoir, a secondary recovery technique, can push to the surface significantly more petroleum. These two recovery methods combined will, on average, allow recovery of only about 30% of a given reservoir's content.

A third recovery method, called enhanced petroleum recovery, involves adding carbon dioxide or a surfactant (a detergent) to the injected water. These additives act to relieve surface tension at oil-water interfaces, allowing the two liquids to mix. This, in turn, allows the injected water to penetrate further into the small passages in the rock and to mobilize more of the oil trapped there. A successful enhanced recovery operation will typically bring to the surface an additional 10-15% of a reservoir's content, an improvement that could translate into billions of dollars in revenue from a single enhanced recovery operation.

Several major oil companies are using CRAY computers in different ways to simulate the gross flow of oil and water underground. The chemical engineering researchers at the University of Minnesota are using their CRAY to simulate oil-water interfaces and flow and dispersion rates inside porous microstructures. Their basic interest is to develop generalized theories to explain the dispersion and flow rates of fluids through media of varying permeability. Such theories may then be incorporated into reservoir simulation programs used by oil companies.

Using finite element analysis, the U of M researchers are able to make predictions about density gradients within and between fluids interfacing in rectangular porespaces. Similarly shaped porespaces are common in the type of carbonate rock found in many petroleum reservoirs. "We had a theory to predict these density gradients, but a theory is just a set of formulas. We didn't know what the formulas said until we were actually able to make predictions from them," said Scriven. "Getting a prediction was beyond conventional mathematics. But where conventional mathematics fails, either in principle or in cost — where it's just not feasible, we can step in today with computer-aided mathematics and compute the predictions. There's not a direct experimental comparison, but our prediction squares with the whole circumstantial case about conditions in small pores."

Being able to predict and map out the density gradients within and between fluids interfacing in microscopic pores allows scientists to determine which fluids in a porous rock adhere to the walls of its porespaces and how strongly. It also allows them to determine the shape of the menisci, or interfacial surfaces, between the fluids. These factors, in turn, directly influence how difficult it is to mobilize the fluids inside the rock, and once calculated, can be incorporated into reservoir simulations to improve their accuracy.

Another approach the chemical engineering researchers have taken to improving the accuracy of reservoir simulations is to construct three-dimensional simulations of porous rock. The flow of fluids through the rock is modeled as the electrical conductivity of electric circuitry. The flow of fluid through a permeable medium, such as a porous rock, is reliably comparable to the flow of electricity through circuits. A medium of varying permeability is accurately modeled by simulating a network of electrical conductors with varying proportions of randomly placed nonconducting insulators. The relative conductivity of the network is analogous to reservoir rock's permeability to the injected water solution used in secondary and enhanced recovery operations. The injected solutions are typically saline, and therefore electrically conducting. The relative nonconductivity of the simulated network is analogous to reservoir rock's permeability to oil, which does not conduct.

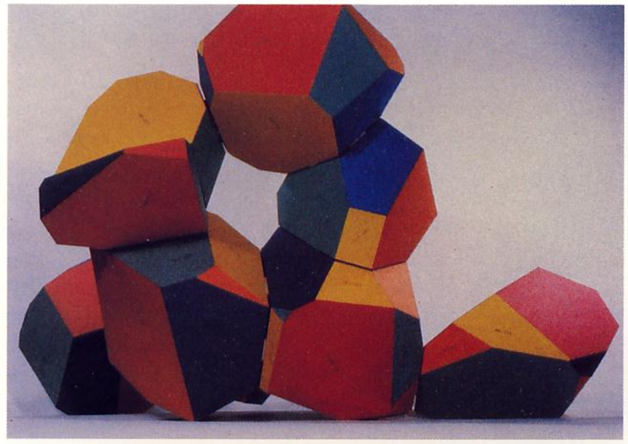

Figure 1. A three-dimensional model of irregular polyhedra based on a computer-generated Voronoi tessellation.

By a mathematical process called Voronoi tessellation, a computer simulates a three-dimensional space filled with irregularly shaped polyhedra. These polyhedra simulate irregularly shaped porespaces and rock granules as shown in the model in Figure 1. The computer randomly designates polyhedra as electrical conductors while monitoring the system's overall conductivity. When a critical number of designated conductors, or "the percolation threshold," is reached, the system begins conducting. Its conductivity increases continuously as the number of designated conductors is increased. Figure 2 illustrates the similarity of experimental results with those that were computed. The applicability of Voronoi tessellation to determination of rock permeability has been demonstrated experimentally by the chemical engineering researchers.

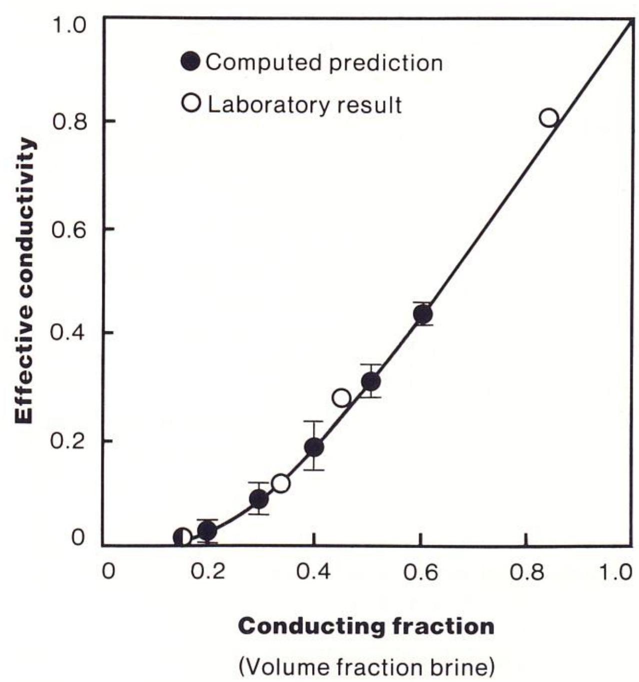

Figure 2. Finite element predictions of Voronoi conductivity are compared with the conductivity of a microemulsion determined by laboratory experiment.

Even by checking such simulations against numerous core samples taken from reservoirs, interpolating diffusion and flow rates between sample locations does not always yield an accurate map of a given reservoir. Gross features of the reservoir, such as underground temperature and pressure, must also be considered. In addition, features of the reservoir's structural geology, such as fissures and fractures, and the relative proportions of the hydrocarbons making up the petroleum itself must be taken into account. These factors will vary from place to place within a reservoir. Therefore, it is necessary to determine the rates at which the surfactant will diffuse and the mixture of fluids will flow through each part of the reservoir. Still, reservoir simulations can be very cost effective despite these difficulties. With potentially billions of dollars at stake in an enhanced recovery operation, a single correct decision based on computer simulation can pay for a CRAY many times over.

Coating flows

Another area of fluids research at the U of M's Chemical Engineering and Materials Science Department which has benefited from CRAY computing power is the prediction of coating flow configurations. Photographers, musicians, and anyone who has ever wrapped a birthday gift, have all benefited from the study of coating flows. Photographic films, audio and video tapes and the many varieties of tape adhesives all consist of a solid substrate which, during the manufacturing process, is coated with a thin liquid film. Constructing and operating equipment used for such coating processes requires a working knowledge of fluid dynamics, including interfacial phenomena and surface tension. The U of M researchers have successfully applied their brand of finite element analysis to fluid dynamics problems in thin film coating.

Typically, the goal of a thin film coating process is to create a liquid layer uniform in thickness and composition, at the fastest possible rate with the least wastage. Progress in coating technology generally means being able to coat thinner, more uniform films of higher solids content at higher speeds. But until recently, accurate theoretical models have lagged behind trial and error and intuition as tools for advancing coating technology.

Now some accurate theoretical models have been developed. Scriven and various collaborators have developed theories to precisely describe and predict the complicated shapes of flow fields found in real situations. "Most theories up to this point have been heavily idealized," said Scriven. "For example, the biophysicist studying the horse may say 'As a first approximation I will take the horse to be a sphere.' That's typical, because all the previously available mathematics was for spheres, bricks, cylinders and relatively simple shapes. In these coating flows, however, the configurations are complicated free films of peculiar shapes turning sharp corners."

A particular coating configuration that has been described successfully using the new theory is curtain coating, a process which involves pouring a liquid coat onto a moving substrate. Curtain coating involves free surface flow fields. These are characterized by two air-liquid interfaces and are affected by the interaction of viscous, pressure, capillary and often inertial and gravity forces. Surface tension at the air-liquid interfaces complicates flow field analysis. Additional complications are introduced by the poorly understood physics of the three-phase contact line where the solid substrate, the coating liquid and the displaced gas meet. The shapes of the free surfaces involved in coating flows are governed by nonlinear equations. Nonlinearity is significantly increased if the liquid being used to coat is non-Newtonian, which is the case with many industrial applications. Non-Newtonian liquids, such as polymer solutions, exhibit elasticity, and in other ways behave qualitatively differently from more common Newtonian liquids.

"Our method allows us to describe what's happening around the impact zone: how high the stresses are there, if anything is lodging there, or if there are any recirculation patterns there." explained Scriven. (See Figure 3.) "It also tells you what's happening at the edge of the pouring lip: if the liquid is falling

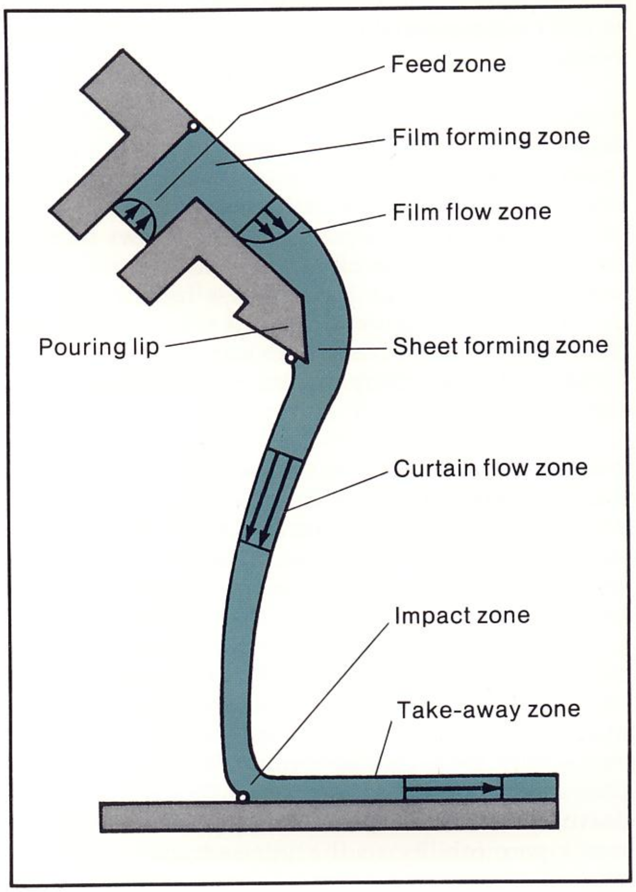

Figure 3. The diagram depicts distinct flow zones into which a curtain coating flow field is is divided for finite element analysis. Researchers predict the conditions under which a given liquid must be poured to achieve desired coating thickness.

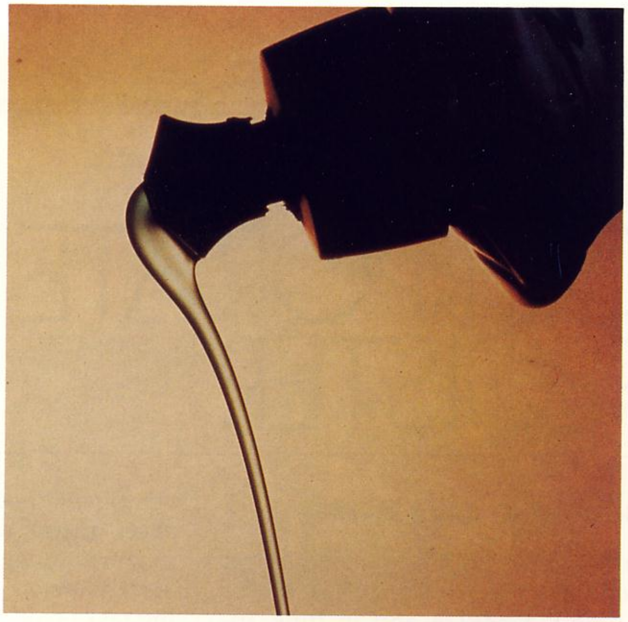

Figure 4. Pouring liquid from teapot demonstrates the "teapot effect." Detergent bottle shows lip designed to avoid "teapot effect"; design matches finite element predictions.

smoothly or if it is wetting back in the 'teapot' mode (as shown in Figure 4 photos.) But more importantly, these analyses tell us the operating range of the process. This is important because, for example, if you increase the flow rate too much, the liquid starts building up in the impact zone, which you don't want. If you slow the substrate too much, the same thing happens." Scriven added that analyses so far have been two-dimensional. His group is beginning to develop three-dimensional simulations which will take into account the edges of flow fields.

Particularly important to generating predictions from this theory has been the ability to employ the full Newton-iteration process to solve the entire set of linearized equations. In a particular example, predicting the shapes of curtain and roll coating flow fields (roll coating is a coating process involving application by rollers), the number of elements in a relatively coarse discretization ranged from 30 to 100 and the number of unknowns from about 400 to 1500. Some of the predictions, mainly with the smaller numbers of unknowns, took from 5 to 12 CPU seconds per iteration using a CDC Cyber 74. With the larger numbers of unknowns the calculations were made with a CRAY-1 and required from 0.5 to 1 second per iteration. A complete parameter study takes hundreds or thousands of iterations and finer discretizations are often needed.

Conclusion

Academic researchers are unique in their ability to investigate phenomena without having to keep an eye on potential profits. Their bottom line is advancing our understanding of the physical universe. Nonetheless, further understanding of the processes described in this article will undoubtedly pave the way for improving many commercial and industrial processes.

Better understanding of fluid behavior in permeable media has direct applications for improving the accuracy of petroleum reservoir simulations, which are an increasingly important tool for making multibillion dollar decisions. Improving simulation accuracy will increase the cost effectiveness of petroleum recovery operations and, in turn, benefit everyone who uses gasoline, oil, plastics and other petroleum products. Similarly, improving the efficiency of thin film liquid-coating operations promises to benefit both industries using these processes and their customers.

The University of Minnesota is one of only three universities in the United States to own a supercomputer. Its supercomputing ability allows researchers to test theories faster and more economically than they could solely with laboratory experiments. CRAY computing power has filled a particularly useful niche, in this case, by making feasible the complex computations needed to accurately simulate the varied behavior and properties of fluids in chemical engineering and materials science research. □

References

Benner, Robert Eugene Jr., *Equilibria, Stability and Bifurcations in the Physics of Fluid Interfaces,* Ph.D. thesis, University of Minnesota, 1983.

Coyle, Dennis Joseph, *The Fluid Mechanics of Roll Coating: Steady Flows, Stability and Rheology,* Ph.D. thesis, University of Minnesota, 1984.

Kistler, S.F., Scriven, L.E., *Coating Flow Theory by Finite Element and Asymptotic Analysis of the Navier-Stokes System,* from *International Journal for Numerical Methods in Fluids,* Vol. 4, pp. 207-229 (1984).

CORPORATE REGISTER

Boeing orders X-MP

In May, Boeing Computer Services Company (BCS) placed an order for a CRAY X-MP/24 and in June, installed a second CRAY-1 S/1000. The new X-MP will be installed during the first quarter of 1985. The CRAY-1/S was accepted in early July. The CRAY computers at BCS, in Bellevue, WA, serve the needs of the Boeing Company and customers located in the United States, Canada and Great Britain.

ARCO Oil and Gas orders X-MP

ARCO Oil and Gas Company, located in Dallas, TX, recently ordered a CRAY X-MP/24 computer system. ARCO Oil and Gas, a division of the Atlantic Richfield Company, explores for, produces and sells oil, natural gas and natural gas liquids in North America. The system will be installed during the fourth quarter of 1984.

Ford to get a CRAY

In August, Cray announced that Ford Motor Company ordered a CRAY X-MP/11 computer system. The system will be installed in the first quarter of 1985 at Ford's Engineering Computer Center in Dearborn, Michigan. Ford will use the system primarily for structural analysis. Stuart M. Frey, Vice President, Car Product Development, said, "The massive number-crunching capacity of this supercomputer will triple Ford Motor Company's ability to perform finite element analysis." He noted that wider application of computer modeling techniques to sheet metal design, vehicle crash performance, weight reduction, aerodynamics and other engineering tasks will improve product quality and performance, reduce the cost of prototyping parts and vehicles, and increase engineering productivity.

CTSS Consulting Group

With the availability of the first CRAY-1 in 1976, the U.S. national laboratories quickly recognized the need for CRAY-level computing power. The CRAY had to be integrated into existing computer networks, and communication between the other computers and the CRAY was a must. Consequently, the Livermore Time Sharing System (LTSS), which was used on CDC 7600, 6600 and Star 100 computers, was converted for use with CRAY systems.

The CRAY Time Sharing System (CTSS) provides an interactive interface between CRAY computers and all other systems in several laboratory networks. Today, CTSS, a public domain code, is used on 30-40% of all CRAY computers in the United States.

Recently, the Cray Western Region established the Cray CTSS Consulting Group to provide consulting services to customers using CTSS. Bing Young, manager for the new organization, explained, "Our group is truly a consulting organization, rather than a field support organization. We offer CTSS users assistance in developing and implementing software for new Cray hardware. In addition, we advise Cray Software Development on design that could make interfacing to CTSS easier, thereby increasing the probability that Cray-developed software can be used. The CTSS Consulting Group works with site analysts in developing interfaces between Cray software products and CTSS. When unusual system problems occur, the group is available to provide backup support for customers and on-site analysts."

Young explained that the consulting group is also an integral part of the CTSS Baseline System Working Group, a CTSS user group. That organization is a cooperative effort among all CTSS sites to create a basic system program library as a foundation for local systems. Said Young, "Each user tailors CTSS for their own particular needs. The baseline system integrates the best modifications and makes them available to all CTSS users."

The CTSS Group will assist the regions in assessing the new user's CTSS needs and in identifying the manpower needed to configure, run and maintain CTSS. CRAY CTSS customers will have primary responsibility for installing and maintaining CTSS on their systems. However, the consulting group will assist them with technical expertise and attention.

The CRAY CTSS Consulting Group is located in Livermore, California. Those interested in additional information about CTSS should contact their Cray representative.

APPLICATIONS IN DEPTH

Migration modeling on the CRAY

In seismic data processing, migration is a step that removes distortion from the seismic image that is imposed by the nature of the seismic method. Depth migration, in addition, removes distortion caused by refraction according to Snell's Law. MODMIG™ is a depth migration program that correctly migrates zero-offset data when lateral variations of velocity exist in the section.

The code is a 2-D post-stack migration based on the wave equation and operates in the frequency-space (F-X) domain. In the F-X domain the individual seismic traces have been transformed to the temporal frequency domain, but no transformation is made in the spatial direction. The space domain retains the flexibility required to account for lateral variations of velocity.

The traces are transformed to the frequency domain for the following reasons:

- Computational efficiency — downward continuation can be reduced to a one-dimensional convolution in the space direction for each frequency. When the velocity varies laterally, this becomes a space-variant convolution.
- Parallelism — each frequency component can be processed independent of the others. This means that the algorithm is naturally organized to exploit multiprocessors such as the CRAY X-MP.
- The downward continuation step size can be optimized separately for each frequency to control dispersion.
- Coherent noise, spatial aliasing and evanescent waves can be controlled separately for each frequency by controlling the spatial bandwidth. This is akin to a frequency-dependent aperture.

The downward continuation steps are depth steps; as a result, the output of the migration is in depth. MODMIG has the facility to rescale the depth output to time using the same velocity function used for the migration.

Because the migration is performed in depth and is recursive, the velocities used must be interval velocities. MODMIG accepts velocity input in the form of a detailed interval velocity model described in depth or in the form of stacking velocities at various locations along the line to be migrated.

Conversion of MODMIG for operation on CRAY computers was recently completed in a cooperative effort between GeoQuest International, Inc. and Cray Research at the request of a CRAY user who is also a MODMIG user.

On the CRAY, solution time is fast enough so that it is practical to use MODMIG iteratively as an inverse modeling program. The ability to make rapid iterations means that more iterations can be made to refine the solution.

In converting the program, careful attention was given to obtaining the best performance in the most frequently performed calculations. Overall, on a CRAY-1 M/1200, performance is 240 times that which was experienced on a VAX/11-780. In production the program is running 25 times faster on the CRAY-1/M than on a large-scale mainframe. On the CRAY X-MP, this performance can be multiplied by the number of processors.

Persons interested in additional information on MODMIG operation on CRAY systems should contact: GeoQuest International, Inc., 4605 Post Oak Place, Suite 130, Houston, TX 77027; telephone: (713) 627-7180.

Vehicle crash analysis on the CRAY

Crash/impact simulation of vehicle structures has been greatly enhanced by vector processing on CRAY computers using the DYCAST/GAC nonlinear structural dynamic finite element computer code. DYCAST/GAC is a general purpose code developed by Grumman Aerospace Company specifically directed toward solving vehicle crash problems. It is used to evaluate the crash behavior of helicopter sections and complete automobile structures involving plasticity, large deformations, collapse and material failure. The finite element method permits sufficient detail to model the behavior of individual components. Design studies can be performed by changing the geometry and material of individual components to explore the effect of changes.

Typically, such analyses require many restarts to complete the simulation, solving a portion of the problem each day. However, the CRAY allows weeks of such analysis work to be compressed into one day.

For more information about DYCAST/GAC contact: Robert Winter, Research and Development Center, Grumman Aerospace Corporation, Bethpage, NY 11714; telephone: (516) 575-2237.

Gaussian 82 on the CRAY

Gaussian 82 is a quantum chemistry software package with a broad range of uses in the chemical and pharmaceutical sciences. It allows the chemist to calculate the structure of both stable molecules and molecular structures too unstable to isolate and study by conventional methods.

Gaussian 82 is now available for operation on CRAY computers. It is more stable numerically than Gaussian 80 and has a broader range of capabilities. The CRAY implementation allows chemists to treat molecules previously considered too large for quantum mechanical calculations, as well as allowing calculations to be carried out at a higher level of theory, producing more correct results.

The package contains programs for calculation of one- and two-electron integrals over s, p, d and f contracted Gaussian functions. Self Consistent Field calculations may be performed for restricted or unrestricted Hartree-Fock wavefunctions. Automated geometry optimization to either energy minima or transition states may now be performed. Various one-electron properties of the wavefunction may be calculated, including Mulliken populations and electrostatic fields. Correlation energy calculations may be carried out using Moller-Plesset perturbation theory, configuration interaction, and coupled cluster theory. The package also contains facilities for analytic or numerical calculation of energy gradients and force constants.

Those interested in installing Gaussian 82 on CRAY systems should contact: Dr. John Pople, Department of Chemistry, Carnegie-Mellon University, Pittsburgh, PA 15213.

CRAY helps organize TIMETABLES

One doesn't often think of using a supercomputer in developing project plans and implementation schedules. But CRAY owners may just be interested in what a project management system operating on a CRAY computer can do for them.

What is project management? It is a comprehensive discipline covering the administration of projects from conceptualization to completion. It encompasses a set of principles and methods for sound project planning and control. Project management grew from a need to better plan and control complex projects outside the normal routines of a company. Its primary goals are to minimize the negative effects of major projects on the rest of the organization and to maximize the resource use to meet time, cost and quality objectives.

Sounds great. But what types of projects can benefit from project management? Major undertakings like design and development of a new product line, company mergers and acquisitions, plant expansions and construction, installation of computer hardware and software, contract proposal budgeting and planning are all endeavors that can benefit from project management. As the scale and scope of the project grows, so does the requirement for computer power to schedule it.

Consequently, TIMETABLE™ a computerized project management information system, has been converted for operation on CRAY computers. The system plans, schedules and monitors complex projects by network management methods. The system has been designed to be flexible and easy to use; it is a mature program, having been used in industry and government for over ten years. Large government agencies as well as Fortune 100 companies have implemented TIMETABLE effectively.

The system interfaces with a graphics program that can produce network diagrams, Gantt barcharts and cost or resource graphs. TIMETABLE's database capability allows simple user programs or system utility programs to access calculated data to produce special reports. Through interfaces with database management systems, TIMETABLE has a virtually unlimited ad hoc report generation capability. It is able to combine schedule data with other project data to produce information to guide the project manager's decision making.

The chart below illustrates the type of project information provided by the system.

Those interested in additional information about TIMETABLE should contact: Ray N. Sauer, AccuraTech, Inc., 5422 Chevy Chase Drive, Houston, TX 77056; telephone: (713) 960-9385.

SIGGRAPH '84

A universe of computer-generated images filled the Minneapolis Auditorium and Convention Center for three days this summer. The event was SIGGRAPH '84, the eleventh annual conference and exhibition on computer graphics and interactive techniques, sponsored by the Association for Computing Machinery.

SIGGRAPH '84 was an opportunity for novice and expert alike to view the latest in computer graphics technology. More than 200 companies exhibited their products, including high-speed printers and plotters, video digitizers and displays, electronic palettes, graphics software packages and, of course, computers. Two evenings were highlighted by screenings of computer-generated films and videos.

The conference ran from July 22nd to the 27th. It included 27 technical course offerings on such topics as the medical and biological applications of computer graphics and advanced topics in freeform curve and surface simulation. A technical program of papers and panel discussions addressed areas including international technology transfer, algorithms for painting and matting, microcomputer graphics and graphics standards.

Perhaps the most unique feature of SIGGRAPH '84 was the premiere of the first totally computer-generated Omnimax film, which was shown at the Minnesota Science Museum. Omnimax film is 70mm wide and has frames three times the height of film shown in commercial theaters. The film is projected onto a domed screen virtually filling the audience's entire field of vision. The result is a convincing illusion of being "in" the projected scenes.

The computer-generated Omnimax film that premiered in conjunction with SIGGRAPH was a collage of short sequences demonstrating the varieties of computer graphics applications and styles. Ranging from swoops and dives over imaginary landscapes to trips through outer space and then into the world of living cells, the film provided a chair-gripping roller coaster ride for the eyes.

The film included four CRAY-generated sequences. "Desert Planet Simulation", prepared by Geometric Productions, Berkeley California, included the creation of a planet by "wrapping" a Landsat image around a computer-generated sphere. "The Final Frontier," prepared by Alan Barr, California Institute of Technology (Caltech), and Gray Lorig, Rensselaer Polytechnic Institute, illustrated advanced molecular modeling technology. "Space Station," prepared by Caltech scientists and artists, simulated a flight to a space colony housing such recognizable buildings as the United Nations headquarters and the Washington Monument. This sequence demonstrated advanced modeling techniques for architectural design. "Saturn Fly-by," from

TIMETABLE PROGRESS REPORT — BOOSTER PUMP STATION DEMONSTRATION PROJECT ------- FINAL NETWORK DATA AS OF 28JAN74 PAGE 1 — TIMETABLE BAR CHART REPORT PAGE 1

OVERALL SCHEDULE — ACCURATECH, INC. 5422 CHEVY CHASE DRIVE HOUSTON, TEXAS 77056 — RUN APR-27-1984 07:37:34 — OVERALL SCHEDULE RUN APR-27-1984 07:37:34

SEQUENCE: ES I SELECTION: — SPAN: 28JAN74 TO 08MAR74 — PLOTTED 1 DAY/COL. 1

NODES (I) *	(J)	DURATIONS ORG	REM C	FREE FLOAT	ACTIVITY DESCRIPTION AND CODING				START DATES -EARLY *	LATE-	FINISH DATES -EARLY *	LATE-	TOTAL FLOAT	28JAN74 11FEB74 25FEB74 MTWTFMTWTFMTWTFMTWTFMTWTFMTWTF:	NODES (I) *	(J)
1AA	1AD	5.0	5.0		SUBMIT PAINT CERTIFICATE/SAMPLE	PNTG	OF	ENGR	28JAN74	13FEB74	01FEB74	19FEB74	11.0	XXXXX------- ----- : . :	1AA	1AD
1DA	1DD	5.0	5.0		SUBMIT ELEC MATL/LIGHTS/PANEL	ELEC	OF	ENGR	28JAN74	01FEB74	01FEB74	07FEB74	4.0	XXXXX---- : . : . :	1DA	1DD
1GA	1GD	5.0	5.0		SUBMIT PLMB MATL/PUMP/CONTROLS	PLMB	OF	ENGR	28JAN74	05FEB74	01FEB74	11FEB74	6.0	XXXXX------- . : . :	1GA	1GD
1JA*	1JD	4.0	4.0		SUBMIT PREFAB MATL/BLDG DETAILS	GENL	OF	ENGR	28JAN74	28JAN74	31JAN74	31JAN74	0.0	CCCC . : . : . :	1JA*	1JD
1MA	1MD	5.0	5.0		SUBMIT CONC MATL AND MIX DESIGN	GENL	OF	ENGR	28JAN74	29JAN74	01FEB74	04FEB74	1.0	XXXXX- : . : . :	1MA	1MD
2AA	2AC	2.0	2.0		MOBILIZE PROJECT	GENL	OF	MISC	28JAN74	01FEB74	28JAN74	04FEB74	4.0	XX---- : . : . :	2AA	2AC
2EC	2CE	3.0	3.0		NOTIFY USERS OF OUTAGE THREE DAYS BEFORE THE INTERRUPTION OF SERVICE	GENL	IN	MISC	28JAN74	01FEB74	30JAN74	05FEB74	4.0	XXX---- : . : . :	2EC	2CE
2AC	2AE	2.0	2.0		GRADE SITE	GENL	EX	SITE	30JAN74	06FEB74	31JAN74	07FEB74	5.0	: XX----- : . : . :	2AC	2AE
2CC	2CE	1.0	1.0		SURVEY SITE	GENL	EX	SITE	30JAN74	05FEB74	30JAN74	05FEB74	4.0	: X---- : . : . :	2CC	2CE
1JD*	1JG	4.0	4.0		APPROVE PREFAB MATL/BLDG DETAILS	ARMY	OF	PURC	01FEB74	01FEB74	06FEB74	06FEB74	0.0	: CCCC : . : . :	1JD*	1JG
2AE	2AG	1.0	1.0	4.0	TRENCH FOR FOOTINGS	GENL	EX	CONS	01FEB74	08FEB74	01FEB74	08FEB74	5.0	: X-----: . : . :	2AE	2AG
1AD	1AG	5.0	5.0		APPROVE PAINT CERTIFICATE/SAMPLE	ARMY	OF	PURC	04FEB74	20FEB74	08FEB74	26FEB74	11.0	: XXXXX- ------------ . :	1AD	1AG
1DD	1DG	5.0	5.0		APPROVE ELEC MATL/LIGHTS/PANEL	ARMY	OF	PURC	04FEB74	08FEB74	08FEB74	15FEB74	4.0	: XXXXX- ---. : . :	1DD	1DG
1GD	1GG	5.0	5.0		APPROVE PLMB MATL/PUMP/CONTROLS	ARMY	OF	PURC	04FEB74	13FEB74	08FEB74	19FEB74	6.0	: XXXXX- ----- : . :	1GD	1GG
1MD	1MG	3.0	3.0		APPROVE CONC MATL AND MIX DESIGN	ARMY	OF	PURC	04FEB74	05FEB74	06FEB74	07FEB74	1.0	: XXX- : . : . :	1MD	1MG
2CE	2CG	3.0	3.0		CUT AND STUB-UP PIPE MAIN	PLMB	EX	CONS	05FEB74 E	06FEB74	07FEB74	08FEB74	1.0	: .XXX-: . : . :	2CE	2CG
[illegible]	1JJ	10.0	10.0		DELIVER PREFAB MATL/BLDG DETAILS	GENL	OF	PROC	07FEB74	07FEB74	21FEB74	21FEB74	0.0	: . CCC CCCCCCC . :	1JG*	1JJ
		[illegible]	1.0		DELIVER CONC MATL AND MIX DESIGN	GENL	OF	PROC	[illegible]	[illegible]	[illegible]7FEB74	08FEB74	1.0	:	1MG	1MJ
					FORM AND POUR SLAB/FOUNDATION						[illegible]	14FEB74	[illegible]		[illegible]	2AI

Caltech and the JPL Computer Graphics Laboratory, used data from the Voyager space probe to simulate a flight to the planet Saturn, including a plunge through its rings.

Tim Kay, a Caltech student who worked on the latter two sequences, explained, "They were both done using basically the same program, a modified ray-tracing algorithm. Because of the endless detail that can be incorporated into a realistic graphics scene, and because we need 24 frames per second of viewing time, making movies is one of the few applications I know of that can challenge even a CRAY. As an intern at Cray, I'm working to develop more efficient algorithms for advanced graphics. Having a CRAY handy lets me check my work frequently. Otherwise, I'd have to wait all night to see a picture!"

Back at the Minneapolis Auditorium exhibition hall, one of the most popular exhibits was the computer-generated hologram display at the Cray booth. A joint project of Ford Motor Company., Advanced Dimensional Displays, Inc., and Cray Research, the hologram depicted a futuristic automobile designed by Ford engineers. Mathematical descriptions of the body surfaces were converted to a fine polygonal mesh of over 35,000 polygons using Ford's Product Design Graphics System.

Cray Research contributed computer time and technical assistance to generate the images used to create the hologram. The polygonal mesh was used as input to a special version of MOVIE.BYU for image generation on a CRAY-1/S. Two hundred seventy frames at one-third degree intervals, for a total viewing arc of 90 degrees, were then generated and displayed on a RAMTEK display for photography.

Advanced Dimensional Displays, Inc. combined the computer-generated images into the hologram.

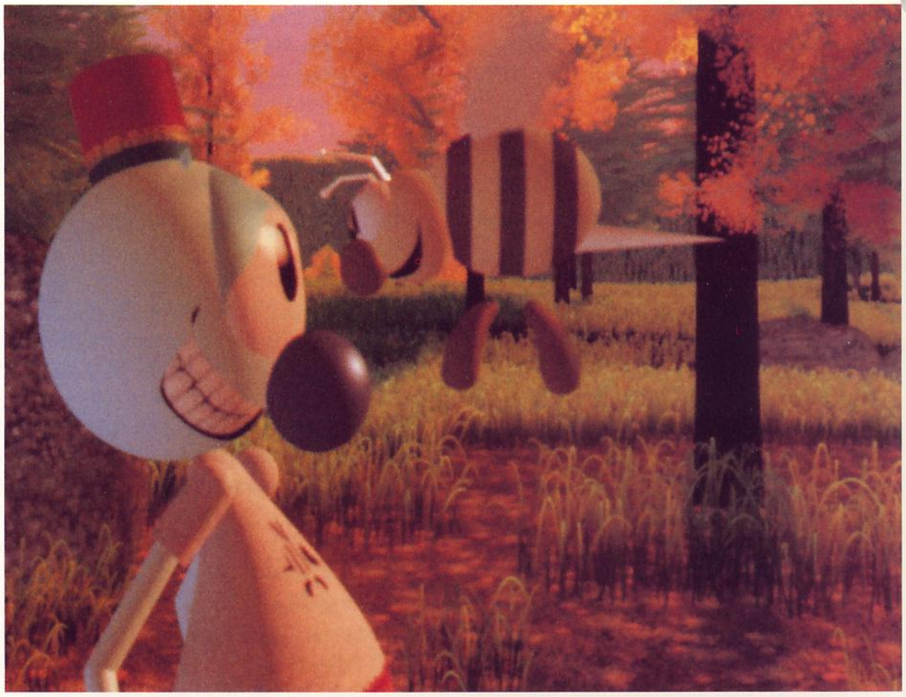

CRAY-generated characters from "The Adventures of Andre and Wally B." © 1984, Lucasfilm.

Each image was used to expose a narrow vertical strip of the nearly seven-foot-long holographic film. The viewer could choose any point of view from the total 270 frames. Each eye viewed only one strip of the hologram at a time, so the viewer saw two views simultaneously, creating stereo depth perception.

SIGGRAPH's two evening film and video screenings, the "Electronic Theater," included submissions from across the United States, as well as from Brazil, France, Italy and Japan. During the second evening's show, Lucasfilm premiered a preliminary version of an animated short entitled "The Adventures of Andre and Wally B.," the first computer-generated motion graphic to include motion blurring consistently on an entire character for an entire sequence. (The evening's emcee noted that nearly one fourth of the technical papers presented at SIGGRAPH were from Lucasfilm.) The characters in the "Adventures..." sequence were generated at Cray's Mendota Heights facility using two CRAY X-MP computers. Over 500 control parameters were used to define the main character. This was the first major program written in the C language to run on a CRAY.

In 11 years, SIGGRAPH has grown from three days of refereed paper sessions and informal idea exchanges to an annual redefinition of the state of the art in computer graphics. From about 30 exhibitors at SIGGRAPH '77 to well over 200 this year, the event has expanded to include virtually every element of computer graphics research, design and presentation. This year's presentations made clear the usefulness of supercomputers like the CRAY for generating realistic three-dimensional motion-graphics, as well as simulating three-dimensional images via computer-generated holograms. The best graphics presented at SIGGRAPH '84 were both dazzling and convincing. What we can expect to see in another eleven years, only a clairvoyant futurist would have any business guessing.

USER NEWS

Animals find unlikely benefactor in the CRAY

While no one has yet marketed computer dating as a way to ensure human survival, computerized matchmaking for endangered animal species is an idea whose time has apparently come. A CRAY-1 has helped chart potential mates for endangered and other captively held animals in an effort to monitor their populations for loss of genetic diversity and potential inbreeding. The effort was organized by the International Species Inventory System (ISIS), an animal data bank providing zoos and animal research centers with genealogical and medical information on captively held animals. ISIS uses the CRAY installed at the University of Minnesota.

ISIS is supported by various organizations, including the American Association of Zoological Parks and Aquariums, the National Institutes of Health, and the American Association of Zoo Veterinarians. Headquartered at the Minnesota Zoological Garden (MZG), ISIS houses lineages and locations of 60,000 captively held birds and mammals, including more than 2200 species from over 180 zoos and animal research centers. The information is made available to participating institutions to help them in several ways, one of which is to locate animals of particular blood lines for captive breeding.

"ISIS collects information such as age, sex, parentage, place of birth and circumstance of death for each registered animal," explained Nathan Flesness, ISIS director. "It includes a subsystem that contains pooled laboratory data relating to blood chemistry for normal individuals. This is especially helpful to zoo veterinarians who may not have access to normal blood properties for each species they examine." ISIS also compiles a microfiche Species Distribution Report showing all the data for all known individuals of every bird and mammal species in captivity. (The version printed on paper runs some 3000 pages and weighs about 45 pounds.) The service began in 1974 and has grown at a rate of about one new member institution per month. Participating institutions send in data forms on their collections and ISIS, in turn, distributes the information to the other participants.

Genealogical data is important to zoos since loss of genetic diversity and inbreeding are perhaps the greatest problems in captivity breeding. Genes that can cause deformities, sterility and early death are common in all animals. However, they are usually recessive and can be masked by the contribution of a corresponding dominant gene from either parent. But with inbreeding, the chances increase that both parents will share some of the same harmful recessive genes, making it less likely they will be masked.

Genealogical record keeping is growing in importance since captive animal populations are primarily made up of zoo-bred individuals. "Of all new acquisitions made by U.S. zoos last year, 90 percent were zoo-bred offspring," explained Flesness. "Breeders are able to locate individuals of varied bloodlines to ensure a thorough mix of genes among their populations by referring to ISIS."

Flesness said access to the University's CRAY has been useful in performing necessary genetic matrix calculations. The mathematics used to generate a matrix involves only simple algebraic equations, but a typical problem might plot 1000 animals against 1000 potential mates, yielding a matrix of one million cells. "We've really been using the CRAY as a big number cruncher," Flesness said. "We've thought about putting the program through a microcomputer and just letting it run for 48 hours, but the problem is getting stable voltage for that long. Things like lightning can cause voltage fluctuations that make it hard to get reliable performance over that long a time." In contrast to the 48 hours a microcomputer might take, the CRAY generates a one-million-cell matrix "in a second or two," Flesness said.

To generate a genetic matrix, the founders of a population are placed in the upper left corner of the matrix grid. The founders of a population are animals originally taken from the wild so their pedigrees are unknown. The matrix then builds itself from the founders by recursion. "The CRAY computer's large memory is useful," Flesness said, "since we can store half the matrix, 500,000 cells, in its core." (Because the matrix is symmetrical about its diagonal, only one half needs to be stored.)

The genetic matrix program calculates the inbreeding coefficient and the relationship to the founders for all possible offspring of all possible matings. The inbreeding coefficient is the chance that at any given gene locus the offspring would receive identical genes from each parent. Identical pairs of genes indicate the possible relationship of both parents to a common ancestor, and increase the possibility of potentially harmful genes being expressed in their offspring.

A nilgiri tahr, one of several endangered species being monitored with the help of a CRAY computer. (Photo credit: John Perrone, Minnesota Zoo)

Animals for which matrices have been generated include the cheetah, the Siberian tiger and the nilgiri tahr, an endangered mountain dweller distantly related to the antelope. The nilgiri tahr matrix was recently run for the MZG and the Memphis Zoo, which are the only U.S. zoos housing the tahr. The matrix was used to select specific pairs of animals for future mating. The MZG also referred to ISIS when assembling its musk-oxen herd and troop of dusky leaf langurs (a type of Asian monkey).

Flesness explained that ISIS's primary role is to disseminate information to its members. "It's hard to know specifically how all of that information is used," he added. "But generally members want to find out who has a certain animal of a certain age and sex, or they want to know who else has a species they've just acquired so they can ask them questions about breeding and raising offspring." Flesness said there are plans to expand ISIS's data bank to include information on captively held reptiles and amphibians. Expansion plans depend on grant funding; membership fees currently account for only one third of ISIS's budget.

Flesness acknowledged that, "Comparisons to computer dating or Noah's Ark are sometimes hard to avoid when explaining ISIS," but he added, "Such flippant comparisons should not obscure ISIS's important function." With wildlife habitats being chain-sawed, paved and cultivated at an increasing rate, zoos are stepping in as refuges for displaced and threatened animal species. As more species come to depend on zoos for sanctuary, the importance of recording lineages and medical data for the animals increases significantly. Perhaps it is ironic that one of the world's most sophisticated technologies is being used to help preserve animals driven from the wild by civilization. But at this point some species' survival depends on ensuring a continued monitoring and conscious mixing of their gene pools, which, in turn, requires sophisticated information storage and processing equipment like the CRAY.

NASA Research Institute plans for future

The long-term goals of the U.S. space program critically depend on advances in computer science and technology. To help meet this need, NASA last year created the Research Institute for Advanced Computer Science (RIACS) at NASA Ames Research Center in California. RIACS' mission is to research and design an advanced computer simulation laboratory and participate in joint projects with other computer-related research groups at the Ames Center.

RIACS will concentrate on researching two computer subsystems to be interposed between the scientist running a simulation and his computing system. The two subsystems being researched are a computational system that would include a library of software parts, and an expert system to assist scientists in using the software library. "Both will be high-performance systems," noted RIACS research engineer George Adams. "I suspect both will require the development of a new computer architecture, as well as new software."

As envisioned by RIACS, the software library will house a collection of subroutines, or "solution blocks," to provide scientists with easy access to a large number of prepackaged algorithms. In the case of computational fluid dynamics, for example, each solution block would designate a specific algorithm and mesh density. The solution blocks will have standard interfaces, allowing them to be pieced together to form larger programs. The library will also contain a programming scheme allowing a programmer to create new parts for the library, either by direct expression of an algorithm or by combining parts already inventoried.

RIACS forsees a situation where a scientist could build a simulation program by mapping solution blocks directly over a graphic image of the object being designed or modified. For example, an aerodynamicist working on an airfoil simulation would scan the on-line library of pre-packaged algorithms and meshes. He would then select the one that describes a given portion of the airfoil and indicate where on the airfoil that algorithm and mesh density is to be applied. An instance of that solution block would then appear as a translucent cube on the appropriate part of the airfoil image. By piecing together all of the required solution blocks, the scientist would eventually cover the entire airfoil, completing the simulation program.

The second subsystem RIACS envisions is an expert system to assist scientists in constructing simulation programs from the subroutines in the software library. The expert system would check to see if the assembled parts were compatible and if they were pieced together logically. "We envision an expert system for each discipline," explained Adams. "Scientists will be able to interact with the system in the language of their scientific specialty. We think we can significantly improve efficiency by minimizing the amount of effort the scientist has to spend accommodating the computer." RIACS will work to develop expert systems for NASA in aerodynamics, computational chemistry and mission planning. But "RIACS also expects this will facilitate the development of expert systems for other sciences," emphasized Adams.

RIACS is also providing technical advice to developers of NASA's Numerical Aerodynamic Simulation (NAS) facility. The NAS facility will operate the supercomputer "engine" that will eventually run the simulation laboratory just described. "Our initial configuration is just getting set up now," commented NAS Deputy Manager, Don Ciffone. "But the system will always be evolving," he added. "Our philosophy is to assemble the most advanced machines available. We'll always be pushing the state of the art."

Ciffone said NAS has ordered a CRAY-2 to be installed in 1985. A CRAY X-MP/12 is being installed as an interim computer until the CRAY-2 becomes available. "The CRAY is going to be the main high-speed processor for our initial configuration," Ciffone said.

Once completed, the NAS facility will be used by NASA and other aerospace manufacturers to conduct wind tunnel simulations. "The cost of wind tunnel time is pricing U.S. airplane manufacturers out of the world market," commented Ciffone. With a sufficiently powerful complex of supercomputers, NAS hopes to allow scientists to simulate airflows over aircraft and significantly reduce the number of wind tunnel hours required for aircraft development, thus lowering development costs. Ciffone emphasized that the NAS facility is not intended to be used only by NASA Ames researchers, but will be available on a time-sharing basis to industry and academia, as well as government researchers.

RIACS will assist in the design of the NAS facility by studying ways to extend the UNIX operating system, which NAS will use as its standard user interface, between multiple UNIX-controlled machines. RIACS is also currently sponsoring studies to produce an analytical model of NAS's initial configuration and to answer performance questions for the benchmark workload.

NASA created RIACS because it recognized the importance of advances in computer science and technology to its future goals. However, once the RIACS blueprint for an advanced simulation lab is formulated, it will be applicable to fields as varied as molecular engineering and astronomy. Taken together, RIACS and NAS promise to open the doors to research that is not feasible at present.

El Niño under close scrutiny

If you lived on the planet Earth during the winter of 1982-83 you may remember the unusually severe storms that pummeled the Pacific coast of the United States, the blizzards that blew across the western U.S. and the widespread drought that disabled parts of the western Pacific, Africa and Mexico. All of this havoc now appears to have been the work of a child. El Niño, (Spanish for "The Child") is a seasonal influx of warm water off the coast of South America which usually occurs around Christmas time. Every three to six years El Niño's intensity increases to the point where weather patterns are disrupted globally. Scientists working with a CRAY-1 at the National Center for Atmospheric Research (NCAR) have developed an accurate model of El Niño with which they hope to analyze its past and present effects and predict future ones.

Flooding from a rainfall of 16 inches. Monitoring El Niño may help to predict disasters such as this. (Photo credit: National Oceanic and Atmospheric Administration)

During ordinary years, storms that form or intensify near the east coast of Asia move across the Pacific Ocean to the North American coast, causing the rainy climates of Washington and Oregon. Eventually they cross the Rocky Mountains and continue eastward. But during years of intense El Niños, weather patterns change significantly. By a physical mechanism not yet understood, the Pacific storm track is diverted northward to Alaska, rearranging storm paths across the United States. For example, in the Winter of 1976-77, an intense El Niño caused this rearrangement. The western United States experienced an unusually warm, dry winter, while the East had an extremely cold and snowy winter as far south as Florida.

Maurice Blackmon of NCAR's Global Climate Modeling Group along with John Geisler (University of Utah) and Eric Pitcher (University of Miami) conducted four long-term simulations of El Niño using NCAR's Community Climate Model (CCM). CCM is a three-dimensional spectral general circulation model that simulates atmospheric behavior in terms of interacting fields, such as wind and pressure fields. These fields, in turn, are analyzed in terms of spherical harmonics. Previous general circulation models were based on a longitude-latitude grid, which was much less precise than the approach CCM uses.

Combining the new approach with high-speed computing has considerably increased the practicality of conducting long-term atmospheric simulations. Modeling one day's atmospheric activity with CCM took 110 seconds on the CRAY-1. This was carried out in thirty-minute time steps, each of which required solving partial differential equations for nearly a half-million variables. "It's now feasible to do multi-year simulations," commented Blackmon. "Five to eight years ago when we ran general circulation simulations, a simulation of one month was considered long."

To plot the effects of El Niño, Blackmon and his collaborators established a basic warm anomaly for the equatorial Pacific sea surface, based on an average of six different extreme El Niños. Then they ran a control experiment followed by three simulations using three different sea-surface temperatures: the norm plus the anomaly, the norm plus twice the anomaly and the norm minus the anomaly. Each simulation represented 1200 real days from which the investigators extracted 90-day-averaged samples.

The model successfully simulated numerous meteorological anomalies that corresponded to observations, including a change in the standing wave (the normal wave pattern that influences storm direction), a change in the position and strength of the jet stream in the Northern Hemisphere and a change in the storm routes.

Blackmon and his collaborators will continue to investigate the effects of El Niño and will focus on determining why the effects of one differ so much from those of another. Preliminary results suggest this variation may be due to the position of the warm-water anomaly. In 1976-77 the warm water was near the international dateline, but in 1982-83 it stayed near the South American coast.

Blackmon noted that the mechanism responsible for El Niño's periodic changes is not yet understood. He added that extreme El Niños do not correspond to other periodic phenomena, such at sunspots or solar flares, or to geological phenomena such as volcanic activity. "It appears to be entirely an oceanographic phenomenon," Blackmon said.

If his current ideas about El Niño are right, Blackmon added, it should be possible to give three month advance warnings of future extreme El Niños. Advance warnings would be valuable to utility companies, which could anticipate sharp rises or falls in consumption rates, and to farmers who could better plan their seasonal timetables and irrigation needs.

In its 1983 annual report, NCAR concluded: "The high-speed CRAY-1 computer enabled the researchers to conduct this experiment, which was nearly five times larger than any similar previous ones. One of the 1200 day simulations took the CRAY-1 35 hours to run; previous computers would have required 140 hours for a model this size."

The Last Starfighter

In case you haven't read your newspaper's movie guide lately, you may be interested in knowing that the movie "The Last Starfighter" is appearing in theaters around the United States. Yes, Digital Productions' efforts on behalf of Lorimar Productions have come to fruition — and the result is quite amazing. Some 25 minutes of the movie were produced with Digital Productions' CRAY computers.

Computer-generated scenes are cut with live action in such a way that most viewers cannot tell the difference between the real and the imaginary. In one instance, the hero takes off in a space-like automobile, constructed in the traditional fashion by Lorimar, while a spaceship, created by Digital Productions, flies off into outer space. It's virtually impossible to tell that they were created in different media when they are on the screen.

With the CRAY X-MP, Digital Productions is able to generate 400,000 to 1.5 million polygons per image. The hero-ship in Starfighter, the Gunstar, is composed of 750,000 polygons.

On a CRAY-1/S and then a CRAY X-MP, the computer-generated portions of Starfighter took 12 months of production with three months in preproduction for designing and planning. The CRAY X-MP gives Digital Productions the capability of producing two 90-minute feature films a year. Of course, the X-MP takes it all in stride.

WORLDWIDE BUSINESS CENTERS

Domestic

Corporate Headquarters
608 Second Avenue South
Minneapolis, MN 55402
(612)333-5889

Government Relations Office
1828 L Street NW
Suite 201
Washington, DC 20036
(202)775-0155

Central Region Office
5330 Manhattan Circle, Suite F
Boulder, CO 80303
(303)499-3055
Albuquerque
Chicago
Minneapolis

Eastern Region Office
11710 Beltsville Drive
Suite 500
Beltsville, MD 20705
(301) 595-5100
Atlanta
Boston
Laurel
Pittsburgh

Petroleum Region Office
5858 Westheimer, Suite 500
Houston, TX 77057
(713) 975-8998
Dallas
Tulsa

Western Region Office
5776 Stoneridge Mall Road
The Atrium, Suite 350
Pleasanton, CA 94566
(415) 463-2800
Los Angeles
Seattle

International

Cray Canada Inc.
4141 Yonge Street, Suite 302
Toronto, Ontario
Canada M2P 2A8
(416) 299-2729

Cray Research France, S.A.
Rue de Tilsitt
75017 Paris, France
01-766-01-55

Cray Research GmbH
Perhamerstrasse 31
8000 Munich 21, West Germany
89-56014-0

Cray Research Japan, Limited
W 1154, Shin Aoyama Building
1-1 Minami-Aoyama 1-chome
Minato-Ku, Tokoyo 107, Japan
03-403-3471

Cray Research (UK) Ltd.
Cray House
London Road, Bracknell
Berkshire RG12 25Y
United Kingdom
0344-485971

CRAY RESEARCH, INC.

Corporate Headquarters
Cray Research, Inc.
608 Second Avenue South
Minneapolis, MN 55402
Tel: (612) 333-5889