2
30/84

mfr (2)

DR-0641-6

I-18278

E&TR

Energy and Technology Review

Lawrence Livermore National Laboratory

October 1984

DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

Prepared for DOE under contract
No. W-7405-Eng-48

About the Journal

The Lawrence Livermore National Laboratory, operated by the University of California for the United States Department of Energy, was established in 1952 to do research on nuclear weapons and magnetic fusion energy. Since then, we have added other major programs, including laser fusion and laser isotope separation, biomedical and environmental sciences, and applied energy technology. These programs, in turn, require research in basic scientific disciplines, including chemistry and materials science, computer science and technology, engineering, and physics. The Laboratory also carries out a variety of projects for other Federal agencies. The *Energy and Technology Review* is published monthly to report on unclassified work in all our programs. A companion journal, the *Research Monthly*, reports on weapons research and other classified programs. Titles from recent issues of the *Energy and Technology Review* are listed opposite the inside back cover.

100000

E&TR

UCRL--52000-84-10

DE95 003245

Energy and Technology Review

Lawrence Livermore National Laboratory

October 1984

E. O. Lawrence Awards for 1984	1
Entering the World of Multiprocessing	2
One way to decrease computing time for a given calculation is to divide it into parts that can run simultaneously on several separate central processors, all of which share a common memory.	
Uranium Resources and the Development of Fission and Fusion Breeder Reactors	12
Our studies suggest strong reasons to support intensive research and development of breeding blankets for hybrid fusion reactors. According to present estimates, one fusion hybrid breeder could make enough fuel to support 5 to 20 conventional light-water reactors.	
Supracompression of High-Explosive Detonation Products	24
We have measured the shock-wave velocities in high-explosive detonation products to considerably higher precision than before. These new data have caused us to reexamine the way we describe detonations in our hydrodynamic codes.	
Abstracts	33

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

MASTER

UCRL-52000-84-10
Distribution Category UC-2
October 1984

About the Cover

The Laboratory has recently installed this Cray X-MP/2 computer, consisting of two identical central processing units that share a 4-million-word memory. To users linked to our network, the machine appears to be two independent units. We are developing software that soon will enable the processors to work simultaneously on the same problem, resulting in a significant reduction in computational time. The article beginning on p. 2 describes the Laboratory's multiprocessing system.

Scientific Editors

Lyle Cox
James Frank

General Editors

Robert A. Condouris
Richard B. Crawford
Lauren de Vore
Mary Ann Esser
Palmer T. Van Dyke

Graphics Designer

John Schuster

Artists

Frank Marquez
Kimberly Smith

Compositor

Louisa Cardoza

E. O. Lawrence Awards for 1984

Among the six recipients of the Ernest Orlando Lawrence Award for 1984 were two LLNL physicists, Peter Hagelstein and Robert Laughlin. Established in 1959 to honor the late Dr. Lawrence, a Nobel Laureate, the awards are given to U.S. citizens relatively early in their careers for "especially meritorious contributions to the development, use, or control of atomic energy." The awards, granted by the Department of Energy, now include the fields of chemistry, life sciences, materials science, national security, physics, and reactors.

Dr. Lawrence's Nobel prize was awarded in 1939 for his invention of the cyclotron and its subsequent application to basic research, which included the discovery of new elements and isotopes. Lawrence was also interested in medical applications of the cyclotron. Together with his brother John, he did much work in this field, including the irradiation of malignant tissues with neutrons. During World War II, he was a leader in developing the atomic bomb. After the war, Lawrence played a guiding role in establishing the Atomic Energy Commission and two major Department of Energy facilities, Lawrence Berkeley Laboratory and Lawrence Livermore National Laboratory, which bear his name and testify to his many contributions to science and to the nation.

This year's Lawrence awards bring to a total of 15 the number of Laboratory researchers who have been selected for this award. Each new winner will receive a medal, a citation, and \$10 000.

Dr. Hagelstein was cited for "exceptional contributions to national security through his innovation and creativity in x-ray laser physics, including the prolific conception of x-ray laser schemes, the analysis of x-ray lasing phenomena, and the creation of extraordinary computational modeling tools." Hagelstein started his career at the Laboratory in 1976. As a member of the

Physics Department, he has contributed to the fields of vacuum ultraviolet and soft-x-ray laser physics, atomic structure, electron collisional processes, plasma processes not in local thermodynamic equilibrium, radiation transport, and physical modeling with computers. He completed his undergraduate and graduate work at the Massachusetts Institute of Technology, where he received his Ph.D. from the Department of Electrical Engineering and Computer Science in 1981.

Dr. Laughlin obtained his Ph.D. in physics from M.I.T. in 1979. He came to the Laboratory from Bell Telephone Laboratories, where he held a postdoctoral appointment in the solid-state theory group. He was for two years an IBM Fellow. He is currently a research physicist in the Physics Department. His award recognizes his demonstration that the fractional quantum Hall effect, a new phenomenon recently discovered in a semiconductor, is produced by the condensation of electrons into a previously unknown form of quantum liquid—in effect, a new state of matter. This liquid is characterized by particle-like excitations that carry exactly one-third the charge on the electron. Laughlin's theory identifies the fractional quantum Hall effect with the first laboratory observation of a particle with a precisely quantized fractional charge.

Entering the World of Multiprocessing

As a direct way of decreasing the real time to calculate a problem's solution, we are designing software methods that enable many computers to work on the same problem at the same time. Much of our multiuser networking experience is directly applicable to this problem.

For further information contact
Harry L. Nelson (415) 422-4204 or
Pierre J. Du Bois (415) 422-4007.

Throughout the Laboratory's 30 years of work with the world's fastest computers, we have found that none was fast enough—for long. The process of devising ways to solve computational problems has always produced new concepts for which the fastest available machine was too slow. Fortunately, computer manufacturers have obliged by producing ever-faster machines. The increased speed is used in three main ways: to reduce a computation's run time, to expand its scope, or to increase its accuracy.

However, many of our workhorse programs of the 1950s are unused today, although they would run thousands of times faster on the new machines (a few seconds instead of hours). As it takes considerable effort and setup time to prepare a job for a computational run, users want the machine to work at least several minutes in return. As a result, computer programs (sometimes called codes) have a life cycle; as the hardware

becomes faster, we must either scrap the old codes, enhance them, or rewrite them to include more and better physics.

An upper limit on computing time (about ten hours) arises from the turnaround problem; the solution must be produced soon enough to affect ongoing activities. A program that takes a week to predict tomorrow's weather is of no more use than one that takes ten minutes—but gets it wrong. Programmers must decide, for each new generation of computers, the trade-off between what they *can* put in without making a program run too long and what they *must* include to obtain a reliable result.

An Historical View
Now, as in the past, we can see that adding more physics to some current codes would produce qualitatively more valuable results, but it is not clear that we can run enough such problems, within our time and budgetary

constraints, to achieve what we want with presently available hardware. We are discussing what can be done about this situation with computer designers throughout the world, and projects to stimulate research into methods of increasing computer power at tolerable cost are being funded.

For the past 30 years, large increases in overall computer speed have come mainly through advances in the speed of basic hardware components. More recently, most of the increase has resulted from improvements in architecture (that is, the arrangement of components within a computer). These advances might never have become widely available without the development of software, both programming languages and operating systems, capable of utilizing the raw speed of the new machines. Figure 1 and Table 1 present some historical data on computer technology and some projections.

Computers Now in Development

Computer designers are still trying to improve the speed and reliability of computer components, but they almost unanimously agree that the next order-of-magnitude increase in computer performance will have to come from computer architecture. Current efforts focus on some form of multiprocessing, in which a computational problem is divided into parts that can be processed simultaneously.

The major force behind this trend is economic. The cost of designing a major computer today far exceeds the cost of building it, and the only way to keep the price down is by spreading the design cost over many units. One way to do this is to bundle many copies of a machine together, all sharing the same memory and all (at least in theory) able to work on the same program at once.

The finite speed at which electrical signals propagate imposes physical constraints that prevent designers from making computers faster by simply making them bigger. Many of today's switching components operate in less

than 1 ns. Thus, the fact that a signal moves between two components at only about 20 cm/ns places a severe practical limit on the distance between

Fig. 1

The past and the projected future of computer improvement at LLNL. (The CPC, at the lower end of the overall improvement curve, is a card-program calculator.) Each decade for the past 30 years has seen a tenfold increase in the speed of the fastest machine at the Laboratory, a trend that shows no sign of slackening. Although component improvements have slowed, the contribution of computer architecture has maintained the rate of overall improvement.

Table 1 A capsule history of computer development, listing the main types of components and architecture employed during each decade, as well as the major software developments that enabled wide access to the hardware features. The shaded area indicates expected future developments.

Decade	Components	Architecture	Software
1950-59	Vacuum tubes Rotating storage	Index registers Parallel input/output	Assembly language
1960-69	Transistors Core memory	Independent functional units	Optimizing compilers
1970-79	Integrated circuits Kilobit chips	Pipelines (vector functional units)	Vectorizing compilers
1980-89	Large-scale integration.	Multiple processors	

components in a machine that attempts to apply all of its resources to a single task.

Another factor that makes the multiprocessor approach attractive is that system software initially developed to accommodate time sharing on networks of computers is directly applicable to managing multiple tasks within a multiprocessor. Several such software spinoff efforts are being carried out at LLNL, elsewhere in the United States, in Japan, and in several European Common Market countries.

Finally, a multiprocessor need not always be used on a single shared program. Within a time-sharing, multiuser network, each individual processor may also be used independently, together with a portion of its memory. Hence, existing applications codes run with little or no revision.

An early example of multiprocessing on a Cray X-MP computer is the champion computer chess program Cray-Blitz (see the box on p. 6).

The Laboratory's Computer Network

At present, the four Cray 1 computers and one Cray X-MP/2 of the Livermore Computer Center (LCC) make up the bulk of our computing power. Each Cray machine is connected to the users and to the other computers by a communication network known as Octopus (Fig. 2), making the LCC a multiuser, multicomputer facility. Each of the current computers is serviced by operating-system software called the Livermore Time-Sharing System (LTSS). (An operating system is a specialized collection of programs and data that governs the operation of a computer or a

network.) Figure 3 schematically depicts the computer system from the user's point of view.

As the Octopus network evolved into its present structure, we found ourselves maintaining software for a growing list of different kinds of equipment, each serving a particular function, each with its own unique communication protocol. This situation developed because:

- We had to bring new equipment on-line quickly.
- There were many people designing software in isolation.

Fig. 2

The Octopus network as of August 1984. The dark vertical lines in the center are a wide-band bus (the 50-Mbit/s HYPERchannel) that connects the present suite of LLNL computers (three CDC 7600s, four Cray 1s, and the new Cray X-MP/2) with the major facilities (mass storage and output devices) and the users (ranging from other computers to individuals at terminals).

• There was no standard to which software had to conform. We needed a way to maintain the existing network software efficiently while allowing new hardware and software to be developed and introduced.

Further, we were beginning to feel the need for software that was modular and portable (that is, functionally "clean" and easily movable from one kind of computer to another), written in a high-level programming language, and easy to maintain over the software life cycle. These concepts were considered

important because of the increasing size of the network, the cost of support personnel, and the difficulty in acquiring people interested in the kind of work to be done.

The installation of the HYPERchannel, a multibus, high-speed data channel¹ (50 Mbit/bus), made it possible to completely restructure the LCC network. But the successor to LTSS, named NLTSS (Network LTSS), was more than a mere replacement. We wanted it to become a network operating system that is distributable and extensible (that is, capable of residing in more than one processor simultaneously and with a readily expandable basic instruction set) and, as appropriate, to support multiprocessing.

LINCS and NLTSS

To achieve such a new order, we had to provide the NLTSS with a framework within which to operate.² We designed a seven-layered network-

Fig. 3

The user's view of the Livermore Computer Center is a set of concentric circles of software and hardware, each of which communicates only with those it touches. The user, on the outside, can access only the applications systems. The user can, however, use certain applications codes to communicate with the programming system.

communications architecture, patterned after the model sponsored by the International Standards Organization (ISO). This framework is known as LINCOS (Livermore Interactive Network Communications System).³ However, the LINCOS architecture is more precise than the present ISO model in that each of its seven layers is completely defined (see the box on p. 8).

NLTS is a purely transaction-oriented, message-passing system whose functionality is determined by the various "server" processes available in the network address space common to it and other LINCOS processes around the network. (A process is usually the smallest unit of work schedulable by the operating system; a server process offers some service, such as management of

Multiprocessing and Computer Chess

Starting in 1974, a World Championship Tournament for Computer Chess has been played at three-year intervals. At the fourth tourney, in October 1983, the Cray-Blitz program established itself as the Computer World Champion. The programming team consisted of Robert Hyatt and Albert Gower of the University of Southern Mississippi and Harry Nelson of LLNL.

Four enhancements added during 1982 and 1983 were responsible for this victory:

- Moving the code to the Cray X-MP computer.
- Adding multiprocessing capability.
- Rewriting many subroutines in machine code.
- Incorporating new opening moves.

The Blitz chess program was begun in 1975 by Hyatt. After several years of development, it held a rating around 1700 (Class B) running on a VAX and was clearly in need of a faster computer. In 1980, Cray Research agreed to support Hyatt with Cray 1 machine time. The converted code, renamed Cray-Blitz (C-B), ran about ten times faster and played chess more than 300 points better (putting it in the Expert class).

One member of the team (Nelson) analyzed C-B's operation and discovered that it was spending a fair portion of its calculational time in one particular subroutine, which was then rewritten for higher speed. In the 1981 annual U.S. tourney, sponsored by the Association for Computing Machinery, C-B tied for second with Nuchess, the Northwestern University entry. The winner was Belle, from Bell Laboratories, which had won the world championship in 1980.

In any position, the collection of all legal plays by one of the players is called a "ply," and two plays make one full move. Belle uses a single-purpose computer designed specifically to play chess and relies mostly on speed rather than analysis, examining more than 100 000 positions per second and looking ahead in a typical situation to a depth of eight plies. C-B, to avoid wasting time on unprofitable plays, assigns a score to each play in the first few plies and explores only checks, captures, and promotions at deeper plies.

Studies with Belle seem to indicate that a program advances by about 200 rating points with each additional ply searched. With further improvement, C-B managed a tie for first with Belle and Nuchess in 1982. In 1983, Belle entered the human U.S. Open and finished with a rating of 2203, the first chess program to achieve the rating of Master. (The highest human performance rating is Bobby Fischer's 2750.)

In 1983, Cray Research released an experimental two-processor version of a multiprocessing operating system for their X-MP computer and offered its use for the upcoming World Championship Tourney. Hyatt and Nelson reworked C-B for use on the X-MP, beginning only two months before the tourney. Dual processing enables the two processors to share the ply-1 moves, after which they do the deepening on their own, finally sharing the backed-up score when they return to ply 1.

Because of the short time available, the program still had some bugs when the tournament began. When these appeared, C-B was ordered to drop back to previously proved routines, and further debugging was done overnight.

On the last night of the tournament, with its speed up to 20 000 positions per second, C-B met Belle. Since Gower had studied Belle's "book," C-B was ready with an innovation. As a result, C-B won a pawn at move 20, held on to this advantage while surviving a couple of program bugs, and finally won on move 54. Belle resigned when C-B's extra pawn advanced to the seventh rank, threatening either to promote or force Belle's king to a square where C-B's queen would check and mate in 3 moves.

Future plans for C-B include generalizing the coding so as to be ready for larger numbers of processors. Now that chess programs are moving into the Master class, there is a great deal of interest in having them play in human tournaments. This interest may fade if the programs ever begin to beat the best human players consistently, but that is unlikely to happen in this century.

disk files, directories, and time clocks.) This modular structure, and the use of a high-level language in its coding, makes the system readily extensible and portable.

The NLTSS architecture consists of a resident kernel (the essential portion of the operating system) and a number of server processes that manage the resources of the local computer on which they execute. NLTSS operates as a multi-tasked system in the sense that each server process spawns a co-server for each active customer interface. These are synchronized through a co-server kernel that manages the co-server semaphores. (A semaphore is a global variable through which co-servers communicate to effect synchronized access to shared data.) It will be relatively simple to adapt this semaphore mechanism to whatever hardware becomes available in the future to support multi-tasking.

Entering the World of Multiprocessing

As a first step into multiprocessor hardware, we have recently installed a Cray X-MP/2 multiprocessor (Fig. 4) that consists of two identical central processing units (CPUs), each compatible with a Cray 1 CPU. These two new CPUs share a 4-million-word primary memory, together with circuitry enabling the processors to communicate with each other, with secondary memory (disks), and with external devices. Testing suggests that each processor is about 25% faster than a Cray 1 on the average, with variations from equal to double speed for specific programs. A successor, the X-MP/4 (with four CPUs), is on order.

At present, the X-MP/2 is available to users on the Octopus network under the current LTSS system. In this environment, it appears to be two independent CPUs. We are working on enhancements to allow both processors to share the work of a single program within the NLTSS system.

Multiprocessing in Operation

Let us consider multiprocessing on an X-MP from two points of view, that of the operating system, whose job it is to

partition the machine's resources among competing users, and that of a user who would like to improve the performance of a program through multi-tasking.

The X-MP consists of a number of independent processors (units of hardware) running in parallel. Similarly, we may think of a program as consisting of a number of independent units of software, that is, sections capable of executing in parallel (possibly using the same reentrant logic thread, a program segment that can be used concurrently by several processes). Each such unit, called a task, becomes a process when packaged with information about the operating system and the hardware environment within which it is to run. Almost all of our present programs have only one process.

Fig. 4

The Cray X-MP/2 consists of two identical central processing units that share a 4-million-word primary memory, together with circuitry enabling the processors to communicate with each other, with secondary memory (16 disks, with 75 million words), and with external devices.

The LINC Architecture

Because the programming language will manage all the required interfaces, the layered structure of the LINC architecture need be of no concern to the applications user. The programmer must know only the words and forms that access the various available LINC resources. But to provide some understanding of the structure, it is useful to consider its two functions: message handling (lower four layers)

and message interpretation (upper three layers).

The message-handling section (unshaded area in the figure) is responsible for reliable delivery of a message (in a packet) to the correct destination exactly once. If the message is too long for one packet, these same four layers will split it up into segments and deliver the various segments in proper sequence exactly once.

The three higher layers of the structure (upper shaded area in the figure) convey thoughts. It is here that the association between sender and receiver (speaker and listener) is maintained (layer 5), the content of one or more packets is organized into a data structure usable by the application program (layer 6), and the transmitted thought is interpreted by the application (layer 7).

The layered structure is valuable because each layer might be represented on the same or various pieces of hardware. For instance, each column of layers may reside on a different Cray computer. The HYPERchannel (physical flow) connects the two columns.

In the stylized diagram, each layer of the structure communicates logically only with its equivalent layer in the same or some other system (dashed colored arrows). To do so, however, the message must work its way down through the lower layers in one column and then back up through the lower layers in the other column (black arrows).

Tokens mentioned in the table are machine-independent information parcels describing a piece of data. The description includes a data type, a usage, the length of the value field, and the actual data value.⁴ Data messages are considered to contain unstructured bit strings.

A monologue, as used in this context, is an association between a speaker and a listener, each having a unique network address. It covers everything involved in transmitting a message one way. Messages are considered to contain either control information or data.

7	Application	Operates on decoded tokens
6	Presentation	Decodes tokens
5	Session	Monologue management
4	Transport	Assurance, flow control
3	Network	Packet routing
2	Link	Bulk data movement, support of layer 1
1	Physical	Network device-to-device communication

Once it is executing as an NLTSS process on an X-MP, a program may create other processes, either as daughters or as clones of itself. Daughter processes are given a domain of access by the parent process, but each daughter process is assigned a unique memory space and unique communication-port addresses.

Cloned processes, in contrast, share the same domain of access, the same memory space, and the same communication-port addresses. A program able to arrange itself into multiple logic threads operating on the same data space in parallel is capable of multiprocessing. It does this by causing the various logic threads to execute in the clone processes.

Figure 5 diagrams such sequences. Although the NLTSS operating system places no restraints on the number of clones that may be created, the higher-level tasking support libraries that are accessible through programming languages will normally limit the number of processes in the program to the number of processors in the computer. For example, a program on the X-MP/4 typically will have one to four processes.

A programmer, in writing his applications program, can define units of computation that may proceed in parallel by declaring tasks. A task consists of a program starting point and some initialized registers. The newly declared task is entered on a list of tasks waiting for execution. As one of the cloned processes becomes available, the task list is scanned for an available task. If there is one, the process executes it; if not, it waits. A program on an X-MP may contain any number of tasks.

Creating clones, declaring tasks, and placing tasks into processes are all functions of the applications program. Scheduling the processes into processors is the function of the operating system. To create clone processes, it is necessary to invoke the NLTSS process server, but this need be done only once during the lifetime of the program. (A process does not die when it stops working. It just waits until it is needed.)

Time Sharing

In a typical time-sharing, interactive situation on our computers, the operating

system is transmitting and receiving messages from several hundred users and is keeping track of several thousand program files located in secondary (disk) memory, each of which is in one of several possible states. Principal among these states are:

1. Idle.
2. Waiting for a message.
3. Waiting for primary memory.
4. Copied to primary memory.
5. Executing.

In general, idling processes (state 1) will be started and advance to state 2; processes in state 2 will receive a message and advance to state 3; processes in state 3 will be moved to main memory (state 4). Processes in state 4 will execute (state 5) for a while and then drop back to state 1 if execution is complete, state 2 when they need input, state 3 when the system decides that it needs its primary memory for some other purpose, or back to state 4 if only the processor is needed (Fig. 6).

Fig. 5

Schematic diagram illustrating how two processes in a program select tasks from a task list and move them to one or another of two processors, as they become available, for execution.

Fig. 6

The possible states of a program on a time-sharing multiprocessor. States 1, 2, and 3 (idle, waiting for a message, waiting to be transferred to central memory) are for programs stored in disk memory. State 4 occurs when the program advances to central memory, and state 5, when the program is running in one of the central processor units. State 0 represents the user in the contemplation mode, thinking about which program to choose and how to use it.

There may be many user processes in state 4 in addition to the resident system process itself, which never drops to a lower state. On an X-MP/2, at most two processes (and on a Cray 1, one) are actually using a CPU to advance a calculation. Each of these gets exclusive use of a CPU for a short time (tens to hundreds of milliseconds), then waits while another process uses the same CPU. Various system processes also run periodically, sharing the available processors with the user processes. Charges are assessed to users for work done on a "per CPU" basis.

The User's View

Suppose now that a user has created a program for an X-MP/2 that consists of three tasks. Task A gets things going (obtains input files, generates starting values, and initializes the appropriate input/output devices); then it begins the first iteration by starting tasks B and C, which we suppose are capable of running simultaneously while sharing the same primary memory area without interfering with each other's data structures (Fig. 7). Task A remains idle while waiting for tasks B and C to complete and resumes after their completion to collect the results of their calculations, to decide whether another iteration is needed, to share any necessary information, to set up for the next iteration, and to restart tasks B and C. Let us assume further that task A works very quickly and that tasks B and C are well balanced, each requiring

Fig. 7

The logical flow of a program consisting of three tasks, two of which (B and C) can execute simultaneously in separate parts of a multiprocessor. Task A starts the computation and gets everything ready for tasks B and C. Colored blocks indicate whether B and C are executing simultaneously or alone. When B and C are both completed, task A integrates the results and restarts B and C. This alternation continues until the computation is finished. Task A functions only when neither B nor C is working.

about the same number of CPU cycles to do its job. The user wants this program, which previously took 75 minutes of CPU time on a Cray 1, to finish in 30 minutes on the X-MP/2.

The Operating System's View

When the user's program starts, the operating system sees it as one unit of computation, a single process that advances from state 1 to state 4. It then becomes two (or more) tasks, each of which must gain control of a CPU in order to progress. Since a second processor is available, a second (clone) process is created containing one of these additional tasks. (One can think of a process as a reusable container into which tasks are inserted for execution.) At the same time, there may be dozens of competing processes from various other users, including the system's own processes, which may also need a CPU.

These demands of other programs will interfere with the completion of our user's job and can cause its processes to run one at a time in the same or alternating CPUs. Finally, if the program must be returned to disk (to make way for a task of higher priority, for example), all of its processes must be stopped and their current state remembered, as well as the state of any of its tasks not yet completed.

Our user's program thus will still consume about 1 hour of CPU time, which may be split among both available CPUs. It will be able to run in half an hour of real time only if it has a high priority or if there is no competing work. In any case, the user will be charged for the full hour of CPU time whether or not his job runs alone or shares time with other tasks.

The user's multi-tasking program in this example has one advantage over the competition; whenever one of its processes resides in memory, all of them do. This improves the chances that more than one of its processes will acquire a CPU at the same time.

A program large enough to occupy all the primary memory of a multiprocessor such as the X-MP/4 would monopolize all its CPUs even though it was using only one of them. The accounting

procedure, in this case, will have no alternative but to charge such a program for all CPUs. It follows that designers of codes requiring a large fraction of the available memory will have a strong economic incentive to multiprocess.

Conclusions

Although LLNL users do not yet have to be concerned with multi-tasking of their applications codes, we believe that the benefits of faster turnaround and better use of their allocated time will soon outweigh the difficulties of restructuring long-running programs into a few logically independent tasks. This will be especially true as the number of processors assembled within a specific mainframe increases from the present two to the expected 16 (or more) by 1990.

Users should find that an X-MP, with NLTSS in place, will provide an adequate test bed for the logical development of a program's multi-tasking capabilities, leading to better utilization of tomorrow's hardware. As they take advantage of this capability, the simple world of Fig. 3 will gradually give way to the more complex one of Fig. 8.

Key Words: architecture—multiprocessing; computer—Cray 1, Cray X-MP; Livermore Computer Center; Livermore Time-Sharing System; multiprocessor; multi-tasking; Networked Livermore Time-Sharing System; Octopus Network; time sharing.

Fig. 8

LINC enables us to send commands in plain English text (which leave the terminal as strings of ASCII characters) to any program, computer, terminal, etc., in the entire Livermore Computer Center complex. Processes can be temporarily linked in a dynamic configuration (colored area) that communicates with another process.

Notes and References

1. The HYPERchannel is a high-speed data channel produced by Network Systems Corporation, Brooklyn Park, Minnesota.
2. P. J. Du Bois, *NLTSS Overview*, Lawrence Livermore National Laboratory, Rept. UCID-20035 (1983).
3. J. G. Fletcher and R. W. Watson, *An Overview of LINC Architecture*, Lawrence Livermore National Laboratory, Rept. UCID-19294 (1982).
4. J. A. Minton *et al.*, *The Syntax and Semantics of a LINC Control Monologue*, Lawrence Livermore National Laboratory, Rept. UCID-18852 (Rev. 2) (1984).

Uranium Resources and the Development of Fission and Fusion Breeder Reactors

Our studies reveal strong reasons to support intensive research and development on hybrid fusion breeder reactors that produce fuel for conventional fission reactors.

For further information contact
Claire E. Max (415) 422-5442.

One of the key factors affecting the potential contribution of light-water nuclear reactors to the future energy supplies of the U.S. and of the world is the magnitude of natural uranium resources. Although the supply of uranium can be extended in a number of ways (such as reprocessing the uranium and plutonium in spent-fuel rods, improving isotope separation, and deploying higher-efficiency light-water reactors or advanced converter reactors), the amount of economically recoverable uranium is, after all, finite. Unless we are willing and able to curtail energy use or to substitute alternative technologies, such as coal or pure fusion, sooner or later we must develop breeders of fissile fuel.

The supply of natural uranium is large, and neither we nor the rest of the non-Communist world is likely to run

out of it much before the year 2050. However, it takes a long time to develop a new energy technology to the point of commercial application and even longer for it to penetrate the market enough to make a difference. We have analyzed hypothetical development schedules for several types of breeder reactor from the perspective of future energy demands and uranium resources. According to our analysis, there is not much time to spare; there are definite advantages to starting experimental development of breeders, especially of the hybrid fusion variety, immediately.¹

United States Uranium Resources

The most reliable estimates of the domestic U.S. uranium supply, based on information from the uranium industry and government surveys, have

been published annually by the Grand Junction Area Office of the Department of Energy (DOE).² The total supply is divided into reserves and resources, with the latter subdivided into categories labeled "probable potential," "possible potential," and "speculative" to reflect a decreasing level of confidence in their existence and recoverability.

The uranium supply is further categorized according to the estimated "forward cost" of recovery. This is approximately the marginal cost of mining and milling, and is roughly half the total cost. Each forward-cost category represents all uranium available up to that forward cost. There are additional speculative resources at considerably higher recovery prices, but, for the present, let us limit ourselves to resources that cost up to \$100 per pound to recover, which translates into a fuel cost up to about five times the current rate.

In recent years, the DOE uranium data have been presented in terms of probability distributions within each forward-cost category. The probabilities are defined cumulatively, indicating confidence that there is at least the given amount of recoverable uranium in the category under consideration. Figure 1 depicts the DOE's 1982 estimate of U.S. total uranium resources in the 95th, 50th, and 5th percentiles of the cumulative probability distribution. At the 50% confidence level, the total U.S. resource for the \$100/lb forward-cost category amounts to 4.1 million tons of yellowcake (U_3O_8).

More than 4 million tons of uranium may seem like a generous supply, but those familiar with the uranium industry emphasize that there are serious obstacles to recovering it all. It takes time to find uranium, and about 15 years from the discovery of new ore to the first commercial production.³ Hence, a uranium producer must be convinced that there will be an adequate demand for uranium at least 15 years into the future. Otherwise, he will be tempted to invest his money elsewhere.

With the uncertainties in today's nuclear industry, there are no forecasts of steady demand. Reactor-construction

lead times are also long compared with characteristic periods of demand fluctuation (Fig. 2).⁴ The combination of long lead times and rapid fluctuations in projected demand makes the uranium market unusually unstable.

Figure 3 shows the enormous swings in uranium price in recent years.⁴ In view of the long lead times in bringing newly discovered uranium to market, such major cyclic variations make it all the more difficult to forecast the demand 15 years in the future. To the extent that mines close during the years when the

◀ Fig. 1

Probability distribution of total U.S. uranium supply (sum of reserves plus potential possible, and speculative resources) at various uranium costs.² Dollar amounts listed represent extraction or "forward" costs only, roughly one-half of the selling price.

▼ Fig. 2

Contrast between rapid swings in the annual change in electricity demand (black) and the long lead times required for reactor construction (color).

price is falling, thus raising the cost of future production, this cyclic behavior also shifts uranium resources into higher cost categories.

A further complicating feature is that most uranium is sold on long-term contracts, since fission-reactor owners like to be sure of their fuel supply for a substantial fraction of the 30- to 50-year planned life of their plants. Accordingly, there may be periods when projected uranium requirements and supplies are out of equilibrium with each other.

In view of these factors, the choice of figures on which to base an estimate of the total available uranium resource becomes a matter of judgment. Some experts argue that one should use a resource estimate corresponding to a low confidence level in Fig. 1, since old estimates of uranium supply have tended to be low. Others argue that, out of prudence, one should use a confidence level close to 100%. For the purposes of this discussion, we will use the numbers at the 50% confidence level (roughly 4.1 million tons of U_3O_8 at forward costs less than \$100/lb) as typical estimates of the total U.S. resource.

Simple Arguments Based on Resource Supply

Having obtained estimates of the U.S. uranium resource, we can use them to evaluate the relative need for fission and fusion hybrid breeders. First, however, we must review the salient features of fusion and fission breeders as they are presently envisioned⁵ and typical fuel-use parameters for light-water reactors (LWRs). The figures given, in all cases, are for reactors capable of generating 1 CW_e .

For a fission breeder, we can take the Liquid Metal Fast Breeder Reactor

(LMFBR) as a prototypical design. (Although the Clinch River breeder reactor was recently denied further U.S. Government funding, European and Japanese LMFBR programs are continuing. In the U.S., industrial studies are concentrating on LMFBR concepts having lower capital costs.) The LMFBR uses neutrons from plutonium fission to convert uranium-238 into more plutonium. Initial fuel loading for an LMFBR is 3000 kg of plutonium, and no further loading is needed (although a substantial amount of fuel is stored outside the reactor for rotation and reprocessing). Fuel production is about 210 kg of plutonium per year, or 6300 kg over a breeder's hypothetical 30-year life.

In the fusion breeder concept, neutrons from fusion reactions in the center of a reactor vessel convert uranium-238 or thorium-232 in a surrounding blanket into fissile fuel (plutonium-239 or uranium-233). The resulting fissile fuel can support several LWRs. There are two generic types of fusion breeder under consideration, distinguished by differences in the blanket in which the conversion takes place.

The breeding blanket of a fast-fission hybrid reactor contains mostly fertile material and generates considerable fission energy in the blanket itself. Such a blanket requires a fusion core with a Q of 2 to 4 (that is, one that produces 2 to 4 times as much thermal power as it takes to keep it running) and would generate enough fissile material to fuel about five LWRs of equivalent power.

The blanket of a suppressed-fission hybrid reactor consists mostly of a nonfissioning neutron multiplier, such as beryllium, plus a small fraction (less than 5%) of fertile material.⁶ This type of blanket requires a fusion core with a Q of 4 to 8 and would support 10 to 20 LWRs. In concept, suppressed-fission hybrids are safer and have less potential adverse environmental impact than either fast-fission hybrids or LMFBRs but require more advanced fusion technology, and their blanket designs are further from current experience.

Light-water reactors are currently fueled with uranium enriched in

Fig. 3

Historical variations in the price of yellowcake (U_3O_8) in constant 1981 dollars.

uranium-235. If we assume operation at 70% of capacity, a 30-year life, recycling of uranium-235 from the spent fuel, and a tails assay of 0.2%, the lifetime fuel required for one LWR is about 5500 tons of yellowcake equivalent. For straight-through operation without recycling, the fuel requirement rises to about 6300 tons.

Given this amount of uranium to fuel an LWR over its lifetime, the total U.S. uranium supply would suffice for about 750 LWRs. Each of these reactors, during its lifetime, would produce about 6.3 tons of recoverable plutonium. If all the reactors are allowed to finish their lives, and if we add to the fuel supply the plutonium currently in the U.S. spent-fuel stockpile (roughly 60 tons), we find that we could eventually fuel 1600 1-GW_e LMFBRs in the U.S.

In 1980, the total installed electric capacity for the U.S. was 570 GW_e, of which 57 GW_e were nuclear. It is clear that 1600 GW_e of LMFBRs should suffice to replace the entire complement of LWRs well into the next century. Even if we include in this simplest scenario the LMFBR's requirements for additional out-of-core fuel handling (reprocessing and rotation), there seems to be enough fuel to start up a replacement group of LMFBRs once the LWRs have depleted the supply of mined uranium. A comparable replacement group of hybrid fusion reactors could also be started up, since these need no fissionable fuel.

To decide when we must start building breeders of either category, we must reason backward from the date when LWRs will have depleted the economic supplies of uranium and ask how many breeders or hybrid reactors per year should have been built previously to bring a replacement fueling system on line at the time when the fuel runs out. A quantitative answer to this question requires us to choose a specific hypothetical growth scenario for U.S. electric generating capacity. Given the present state of the nuclear power industry, any such scenario is highly uncertain.

For total U.S. electricity demand, we use the DOE's Energy Information Agency projections to 1995 (Ref. 7) and a relatively low 1% annual growth rate

thereafter. Nuclear demand is assumed to rise from 16% of the total in 1995 to 20% in 2020 to 50% in 2050.

With 1% annual demand growth, the U.S. supply of economic uranium (that obtainable at less than \$100/lb forward cost) will be depleted about the year 2045, when the U.S. nuclear capacity is 670 GW_e. Similar but more uncertain calculations indicate that the uranium supplies of the non-Communist world will be depleted about 2050. Thus, for both the world and the U.S., the low-growth-rate scenarios predict that uranium resources should be adequate for the next 60 to 70 years.

Let us assume that the LMFBR will be introduced commercially for large-scale deployment in 2010 and that the introduction date for hybrid fusion reactors is 15 years later, in 2025. We take a hybrid-reactor support ratio of 10, intermediate between the fast-fission and suppressed-fission values. For this crude argument, we neglect the effect of the fuel formed in the new breeders themselves in delaying the date of uranium-supply depletion.

With these assumptions, we find that we need a construction rate of about 19 LMFBRs or 3.5 hybrid fusion reactors per year in the U.S. This is in addition to the new capacity required to replace existing aging power plants. A recent study estimated that after the year 2010, the U.S. will need 10 to 20 GW_e of new plants each year, even without overall growth in electric demand, just to meet replacement needs.⁸

The required construction rate for LMFBRs is high. Even neglecting replacement of aging capacity, the rate is at least 50% more than the total capacity of LWRs that would have been built in the same time period, had uranium depletion not been an issue. The required LMFBR construction rate is also more than twice our highest historical rate for LWRs, which is about 9 GW_e/y. This argument seems to imply that, to replace the LWR population with LMFBRs by 2045 without a crash program, we should not wait until 2010 but should start building them in quantity within the next 10 to 15 years.

By contrast, building 3.5 hybrid fusion reactors per year seems quite feasible. Such a construction rate would not compete with the building of new

▲ Fig. 4

Two predictions of U.S. nuclear generating capacity from 2000 to 2050 (from Ref. 9).

Fig. 5 ►

Cumulative U.S. uranium consumption from 2000 to 2075, assuming rapid market penetration using the (a) moderate and (b) high nuclear demand curves from Fig. 4. The LWR curve represents the consumption without breeders. The portion of the LWR curve lying in the colored band exceeds estimated U.S. uranium resources. Numbers in parentheses give introduction dates for breeders.

LWRs to replace aging reactors. The hypothesized date of commercial introduction is still more than 40 years away, allowing ample time for the necessary proof-of-principle experiments. The construction schedule for hybrid reactors would be eased still more if each could support more than ten LWRs.

As stated above, this argument is too simple in that even one breeder makes some fuel, stretching out the existing resource and delaying uranium depletion. However, detailed analyses that include this effect⁹⁻¹¹ agree qualitatively with the conclusions we reached using simple arguments.

Scenario 1: Rapid Deployment

To illustrate the technical potential of the two different breeder approaches, let us begin with a very optimistic assumption: that 10 years after their respective commercial introduction dates, both the LMFBR and the hybrid fusion technologies can be built as fast as needed to replace existing 30-year-old LWRs and to satisfy new electricity demand. (Historically, new energy technologies have taken 30 to 40 years to gain a 10% share of the market. We will examine below how this market-penetration rate changes the picture.)

We reproduce in this section some of the results of Ref. 9, which analyzed the introduction of hybrid fusion breeders and LMFBRs under a wide variety of assumptions. Figure 4 shows two scenarios for the growth of U.S. nuclear generating capacity. On the moderate-growth curve, total U.S. electrical capacity grows a bit less than 1% per year after 1995, with nuclear electricity reaching about 50% of the total in 2050. On the high-growth curve, there is a higher nuclear capacity in the year 2000, but the slopes of the high- and moderate-growth curves are similar.

Figure 5 shows predicted cumulative U.S. uranium consumption, calculated using a support ratio of 4 for the hybrid fusion breeder.⁹ Assuming rapid market penetration for the years 2000 to 2075, the trends for both high and moderate growth are similar. Uranium consumption levels off at the lowest

value for hybrid fusion reactors introduced in the year 2000 and at the next lowest for LMFBRs introduced in 2000. Even the highest level, for hybrid fusion reactors introduced in the year 2020, is still well below the estimated resource of 4.1 million tons.

Scenario 2: Conservative Deployment

The picture changes drastically when we conservatively assume slower market penetration. In reality, the market-penetration rate of a new technology will depend on economic criteria such as relative cost per kilowatt-hour, relative capital cost, and perceived risk of the technology in question. Some analysts have discussed these issues briefly,¹¹ but the usual method is to assume an *a priori* market-penetration algorithm for new technologies, independent of their relative economic attractiveness.^{9,10}

At this point, the choice of algorithm is quite arbitrary, since we are applying historical trends to economies that are much larger than any existing today. Figure 6 contrasts the rapid deployment assumed in the previous section with a slower, more conservative market-penetration rate.⁹ The substantially slower rate does not reach the 10-GW_e/y deployment level until almost 2050. If we scale France's current rate of constructing LWRs in proportion to its gross national product or to the size of its total electrical capacity, the U.S. should be able to exceed the 10-GW_e/y level well before 2050. The curves in Fig. 6 thus represent two extreme deployment schedules, rapid and conservative, with a most likely schedule lying somewhere between.

Figure 7 shows what happens to the projected U.S. consumption of uranium for moderate and high growth rates when we use a conservative market-penetration estimate as in Fig. 6. In both LMFBR cases, consumption shows no sign of leveling off before 2075. The hybrid fusion reactor, because of its higher support ratio, can level off the uranium consumption much earlier.

This effect would be all the more dramatic if we had considered support

ratios for the hybrid fusion reactor in the range of 10 to 20, appropriate for a suppressed-fission concept. In the latter

Fig. 6

Different rates of market penetration of breeder technologies assuming high growth in electricity demand. Rapid penetration (black line) assumes that 10 years after introduction, hybrid breeder reactors or LMFBRs can be deployed as fast as necessary to meet LWR replacement requirements and increased demand. Conservative penetration (colored curve) reflects an extrapolation of slower, historical rates of market penetration (from Ref. 9.)

Fig. 7

Cumulative U.S. uranium consumption from 2000 to 2075, assuming conservative market penetration of breeder technologies. (a) high and (b) moderate nuclear demand curves from Fig. 4 (Ref. 9). The portions of the curves lying in the colored band exceed estimated U.S. uranium resources.

case, the uranium consumption never rises above 4.1 million tons, whereas with a support ratio of about 4 it exceeds 4.1 million tons by a small amount. The LMFBR fails definitively (that is, fails to prevent uranium-supply exhaustion within the time scale considered) in the high-growth case.

The situation in the non-Communist world is qualitatively analogous to that of the U.S. for similar deployment assumptions, except that the time scale for uranium shortages is somewhat longer. The U.S. thus could experience a uranium shortage well before the world at large. This prospect raises the politically sensitive issue of importing uranium. Ore from Australia and Canada is already cheaper than U.S. ore, and this trend would be accentuated if the U.S. uranium supply were depleted before that in other countries. In view of the unhappy precedent in the oil market, the question of uranium imports is certainly worth further analysis.

Economic Considerations

Under present conditions in the U.S. utility industry, the introduction date and market-penetration rate of the hybrid fusion reactor and the LMFBR will ultimately be determined by economics. It is convenient to compare LMFBR and hybrid fusion reactor costs with those for LWRs, since neither of the two breeders will be relevant unless LWRs begin to achieve large market penetration. The price of uranium provides a convenient parameter for this comparison inasmuch as fuel costs contribute to the price of LWR electricity but not significantly to LMFBR and hybrid fusion reactor electricity, for which capital costs are always dominant.

The break-even price of U_3O_8 is defined as that uranium price at which LWRs burning mined uranium would produce electricity at the same cost per kilowatt hour as either LMFBRs burning plutonium or a system of LWRs burning fuel from hybrid fusion breeder reactors.

Detailed analyses that include the effects of the LWR fuel cycle, plutonium pricing, and the LMFBR learning curve have shown that the minimum break-even uranium prices for the LMFBR are

in the range of \$120 to \$170/lb with a mean of \$145/lb for second- or third-generation LMFBRs.¹²⁻¹⁴ (For first-generation LMFBRs, the break-even prices are much higher, clustering around \$250 to \$300/lb uranium.)

Comparable analyses for the hybrid fusion reactors are much more uncertain, since none has ever been built. The hybrid reactor has a dramatic advantage over the LMFBR because of its larger support ratio, however, since this makes it possible to spread increased capital cost over all the client LWRs. Estimates of break-even costs for the hybrid fusion reactor range from \$50 to \$100/lb of uranium,^{15,16} or conservatively \$80 to \$100/lb. If these cost estimates prove out in practice, the hybrid reactor's economic advantage over the LMFBR will be substantial.

To estimate when the LMFBR and the hybrid fusion reactor might be introduced commercially without major government intervention, we must consult estimates of future uranium prices (Fig. 8).^{12,17} Let us use the break-even price ranges given above: \$90/lb for a hybrid fusion reactor and \$145/lb for a mature LMFBR. Then the middle-price curve of Fig. 8 indicates that these break-even prices will be reached when the cumulative uranium consumption is 0.7 and 2.2 million tons, respectively. On the LWR curves of Figs. 5a and 7a, these break-even levels of consumption are reached about the years 2005 and 2030, respectively.

By way of corroboration, a much more detailed analysis indicates that the probability that the LMFBR will become economically competitive with the LWR is zero before 2020, 40% by 2030, and 80% by 2040, in substantial agreement with our crude estimate.¹⁴

There are large uncertainties inherent in this calculation, but it is clear that the hybrid fusion reactor could become economically competitive as much as 25 years before the LMFBR. The effect of this difference on an analysis such as that depicted in Fig. 7 would be twofold: it would delay the date of LMFBR introduction until 2020 or even later, and it probably would make market penetration by the hybrid fusion reactor

faster than for the LMFBF. Both factors would improve the advantages of the hybrid fusion reactor over the LMFBF.

Analogous economic considerations make hybrid fusion reactors seem at least potentially interesting as production reactors for weapon-grade nuclear materials. A hybrid fusion reactor might produce up to six times as much fissile material as a standard production reactor of the same power. Even if its capital cost were three times as much, the net cost advantage per gram of output would still be attractive.

Future Strategies for Breeder Development

Assuming that we would be compelled to deploy fully some form of breeder reactor by the middle of the next century, the scenarios and the simple arguments developed here indicate that, without making unrealistic demands on the reactor construction industry, large-scale commercial application of the hybrid fusion reactor could begin as late as 2025. Because of the LMFBF's low support ratio, its deployment would have to begin by 2000 or 2010 at the latest (well before it could be economically competitive) to avoid inconveniently large construction rates later.

It is also possible to foresee a "window of opportunity" for hybrid fusion reactors. The window would open if:

- Fission power revives in the U.S., reaching perhaps 30 to 50% of total electric capacity in the period from 2020 to 2050.

- Total electricity demand grows at a moderate to fast pace, or if a rapid increase in U.S. nuclear capacity becomes necessary (to reduce the damaging environmental impact of carbon dioxide emissions, for example).¹⁸

- There is at least the promise of adequate reprocessing capacity. The window presumably would close when pure-fusion reactors become economically competitive with fission breeders or with LWRs fueled by hybrid fusion reactors.

A different aspect of the timing issue arises from the higher capital costs of advanced breeder systems relative to

those of LWRs, which translate into higher break-even uranium prices. These higher break-even uranium prices will, in turn, be reached at later dates, as illustrated in Table 1.

Fig. 8

Price of uranium (in constant 1982 dollars) as a function of cumulative U.S. uranium consumption (from Ref. 17).

Table 1 Effective capital costs of various reactor designs and the year of breakeven with LWR electricity costs.

Advanced system	Effective capital cost times LWR cost	Break-even uranium price, constant 1982 \$/lb	Probable year of break-even with LWRs	Range of break-even dates
Hybrid fusion reactor	1.1-1.2	80-100	2005	2000-2010
LMFBF (mature)	1.2-1.4	120-170	2030	2015-2040
Pure-fusion reactor	1.5-2.0?	175-275?	2050?	2040-2060?

^aBased on reactor design as presently conceived.

This type of relation emphasizes the importance of reducing projected capital costs in advanced systems. Efforts are currently underway to design LMFBRs with costs approaching those of LWRs.¹⁹ Capital-cost issues are also being reemphasized in the fusion community.²⁰

Of course, if the growth of commercial U.S. nuclear-fission power never recovers from its present slump, there would be little need for the hybrid fusion reactor and the LMFBR. Presumably, coal use would have to expand significantly and would then play a role somewhat analogous to that of LWRs in our scenarios. As coal is consumed, the rise in its price will be analogous to that of uranium (see Fig. 8). There would thus be some break-even price for coal above which pure-fusion power would be economical. It will take further study to determine whether the break-even date for pure fusion relative to coal would be earlier or later than that given in Table 1 relative to LWRs.

The Electric Utilities' Perspective

If nuclear fission does recover in the U.S., and if a large constellation of light-water reactors is finally built, the electric utilities may, understandably, be reluctant to replace their hard-won LWRs with LMFBRs. The hybrid fusion reactor, with its higher support ratio, would enable them to keep their large investment in LWR technology and expertise while fueling the reactors from a much smaller number of hybrid fusion reactors. Their investment in risky new technology would be further minimized if the hybrid fusion reactors were government-operated, perhaps in conjunction with adjacent secure facilities for reprocessing and waste treatment.

The electric utilities could also expect a good deal of downtime in the first generation of hybrid fusion reactors. Hence, at first, it would be undesirable to connect a 1-GW_e hybrid directly to the power grid; an unexpected outage would create too large a power swing. A preferable operating mode would be to run the hybrid fusion reactor off line as a fuel factory, using its power locally to run ancillary equipment and co-located

facilities. The hybrid fusion reactor is better suited for this fuel-factory role than the LMFBR because it produces much more fissionable fuel per unit of power.

Finally, the utility industry has considerable interest in having proven reprocessing technology available. Reprocessing, as an integral part of the hybrid fusion reactor and LMFBR economy, must be developed and brought to maturity at least simultaneously with the breeders. In the ideal case, reprocessing and waste-treatment facilities should be brought on line first to reduce the risk of investing in breeders or hybrid fusion reactors. Such considerations are particularly important for hybrid fusion reactors, which would require a much larger throughput for reprocessing than would LMFBRs.

Development Strategy for the Hybrid Fusion Reactor

Current DOE funding for hybrid fusion reactor research and development (\$150 000 per year) originates in the national magnetic fusion program. Since the magnetic fusion program faces the prospect of, at most, level funding for the foreseeable future and since fusion is its first priority, current plans are to cut hybrid fusion reactor research to zero within a few years. The Electric Power Research Institute provides a small level of funding for planning and technology studies for the hybrid fusion reactor but probably would not be able to mount a major hardware-development effort.

An unspoken reason for the fusion community's previous lack of vigorous support for the hybrid fusion reactor is a reluctance to "taint" fusion by association with fission power, currently low in public esteem. The formal argument is that there can be no hybrid fusion reactor without the rapid development of fusion technology itself, so the proper course is to use present scarce resources to push fusion development along. Achieving adequate fusion performance will make the hybrid fusion reactor much more credible. It will then be easier to obtain funding for research and development on specific issues such as blanket design and new reprocessing technologies.

This argument has merit, but it assumes that there will be plenty of time for research and design for the hybrid fusion reactor after the pure-fusion issues have been resolved. The validity of this assumption depends on the desired deployment date for the hybrid fusion reactor and on the anticipated rate of progress in the pure-fusion area.

What is the desired deployment date? The scenarios described above indicate that we probably will not run out of uranium even if we delay deployment of the hybrid fusion reactor until about 2025. However, economic projections suggest that it could be economically competitive much earlier, perhaps as early as 2005 to 2015, when the uranium price reaches \$80 to \$100/lb.

What is the anticipated rate of progress in the pure-fusion technology? In tokamaks, the proposed Tokamak Fusion Core Experiment (TFCX), the successor to the Tokamak Fusion Test Reactor (TFTR), is planned to be fully operating in the early to mid-1990s. If successful, it would have fusion and plasma parameters suitable for the core of a first-generation hybrid fusion reactor. In magnetic mirrors, the currently proposed Fusion Power Demonstration (FPD), a successor to the Magnetic Fusion Test Facility (MFTF-B) Upgrade, would likewise have fusion and plasma parameters acceptable for a hybrid fusion reactor if it operates as planned; the FPD would be fully operational in the mid- to late 1990s. Analogous inertial-confinement fusion options for the mid-1990s have yet to crystallize, but we can already estimate from the tokamak and mirror proposals that the physics of the hybrid fusion reactor's core will be resolved no sooner than 1995 to 2000.

If deployment of the hybrid fusion reactor were postponed to the latest possible moment (2025), 25 to 30 years would elapse between proof of the fusion-core physics and the starting date for large-scale commercial deployment. This period would be needed to develop and test detailed blanket designs, to build and run a hybrid fusion test reactor, and to build and run a full-scale commercial demonstration reactor. It is

certainly conceivable, but not guaranteed, that these tasks could be accomplished in the desired 25 to 30 years.

On the other hand, one can argue that if the hybrid fusion reactor really were to prove economical in about 2005, as currently predicted, the country could ill afford to miss out on the tremendous potential savings to the domestic energy economy that it represents (\$20 to \$50 billion between 2010 and 2020, for a range of reasonable assumptions). Nor would we want to miss the opportunity to establish a strong position in the large potential export market for these reactors or the fuel they manufacture. Waiting until the mid- to late 1990s before beginning any meaningful blanket research and development would almost certainly delay commercial deployment of a full-scale hybrid fusion reactor past 2005, partly nullifying its predicted favorable economics.

The estimated economic break-even date of 2005 is far from certain, being based on paper studies of blanket designs for the hybrid fusion reactor rather than on experiments. (Some paper studies have predicted break-even prices as low as \$50 to \$60/lb uranium, leading to even earlier dates.) Only after experience in designing, building, and operating prototype blanket modules can we decide whether our estimated break-even uranium price is realistic, too conservative, or too optimistic, and whether there is, in fact, a window of opportunity for exploitation of the hybrid fusion reactor beginning in 2005.

The predicted favorable economics of the hybrid fusion reactor suggest that blanket development should be pursued simultaneously with fusion research and development. This would enable us, by the mid-1990s, to assess both the practicality and economics of alternative concepts and the operating characteristics of the fusion plasma core. If, at that time, the hybrid fusion reactor's economics continue to look favorable, and if, by the mid-1990s, nuclear fission power has begun a significant revival, the country would then be well poised to have a hybrid demonstration reactor operating about 2005.

Preliminary experimental data on many of these blanket issues can be obtained with nonfusion facilities such as research LWRs, the Rotating Target Neutron Source (RTNS), liquid-lithium test loops, and pulsed electron-beam machines. Table 2 describes the experimental parameters and facilities appropriate for studying some critical blanket technologies.²¹ It might take 3 to 5 years to complete these experiments. Conceptual design work on hybrid fusion reactors could be supported concurrently, with increased attention to fission-suppressed blanket concepts. These experiments and design studies would lay the foundation for design, testing, and deployment of prototype blanket modules on one of the next-generation fusion facilities in the early to mid-1990s.

Table 2 Test facilities for near-term experiments on blanket technology for the hybrid fusion breeder reactor.

Critical parameter or system	Required information or performance	Possible test facility
Lifetime of solid-fuel form under cyclic loads	Exposure to: ~10 ⁵ pulses, tokamak, ~10 ³ pulses, ICF	Research reactors, pulsed electron-beam facilities
Cladding/coolant compatibility	Corrosion of ferritic and stainless steel alloys by liquid metals and molten salts	Research reactors and closed-loop facilities, including effects of magnetic field
Atom production rates, atoms per 14-MeV neutron: tritium fissile	~1.1 ~0.5-1.5	TFTR and accelerator neutron sources
High-energy neutron cross sections	Secondary-neutron emission spectra from 8- to 14-MeV neutrons	Accelerator neutron sources
Fission-suppressed blanket designs	Design data on fluid, particulate, and small-sphere fuel forms	Research and development on fission-suppressed designs
Tritium retention and extraction from lithium materials	Tritium transport and solubility in liquid and solid lithium compounds	Tritium experiments at Mound Laboratory and Argonne National Laboratory

This specific goal of designing and building one or more hybrid fusion reactor blanket modules for deployment on an early-1990s fusion machine has several attractions. It would:

- Provide crucial practical experience that cannot be obtained from theoretical design studies.
- Give thrust and focus to what has been a somewhat diffuse U.S. development effort for the hybrid fusion reactor.
- Be an essential step on the way to a hybrid fusion test reactor.

On the basis of similar arguments, the USSR is planning to deploy a fusion-hybrid blanket around the OTR tokamak, its next proposed major fusion device.

The Laboratory's MFTF-B Upgrade, as presently proposed, would have adequate wall loadings in its central regions to test the major blanket developments needed for a hybrid fusion reactor. The first U.S. test bed for blankets using tokamak geometry might be the Engineering Test Reactor (ETR) rather than the TFCX, because, as presently conceived, the latter would not accommodate these blanket modules. Blanket tests using inertially confined fusion would have to await the development of a fusion driver capable of a high repetition rate; no single-shot test facility would provide adequate time-averaged wall loadings. Wall loadings of about 1 MW/m² might be obtained by the year 2000 using high-repetition-rate, solid-state or krypton fluoride lasers, or a heavy-ion ICF driver with a 10% duty cycle.

Conclusions

Successful completion of hybrid blanket-module tests on an early or mid-1990s fusion facility would preserve important options for this country. If fission power were again growing well in the U.S. and if the economics of the hybrid fusion reactor proved favorable, we would be in a good position to build an experimental facility specifically focused on hybrid applications. This would be a test or prototype hybrid fusion reactor, to be operational by about 2005.

Some among the fusion community might regard such a proposed hybrid fusion reactor as a threat, since it would of necessity divert funds from research on pure-fusion power. However, this viewpoint overlooks an essential fact. As shown in Table 1, because of their projected high capital costs, pure-fusion reactors as presently conceived probably would not be economically competitive with LWRs before about 2050. The hybrid might provide a way to deploy fusion technology commercially 30 to 50 years earlier, providing valuable practical experience and insight into possible fusion cost-saving measures. In this view, early deployment of the hybrid fusion reactor could help rather than hinder the development of pure-fusion power by illuminating the path to lower-cost reactor designs. □

Key Words: reactor—breeder, fusion, fission, hybrid, light water, liquid-metal fast-breeder; uranium resource.

Notes and References

1. C. E. Max, *Uranium Resources and Their Implications for Fission Breeder and Fusion Hybrid Development*, Lawrence Livermore National Laboratory, Rept. UCRL-90820 (1984).
2. *Statistical Data of the U.S. Uranium Industry*, U.S. Department of Energy, Grand Junction, Colorado, Rept. GJO-100(82) (1982).
3. *Uranium: Resources, Production, and Demand*, OECD Nuclear Energy Agency and International Atomic Energy Agency (1982).
4. "Uranium and Nuclear Energy: 1982," in *Proc. 7th Intl. Symp. Uranium Institute* (Buitenvorthe, London, 1983), p. 58.
5. A. V. Nero, Jr., *A Guidebook to Nuclear Reactors* (University of California Press, Berkeley, 1979); figures quoted are typical values from Tables 10-1 and 13-2.
6. J. D. Lee and R. W. Moir, "Fission-Suppressed Blanket for Fissile-Fuel-Breeding Fusion Reactors," *J. Fusion Energy* 1, 299 (1981).
7. *U. S. Commercial Nuclear Power*, Energy Information Agency, U.S. Department of Energy, Rept. DOE/EIA-0315 (1982).
8. E. E. Hill et al., *Nuclear Fuel Cycles for Post-1990*, Lawrence Livermore National Laboratory, Rept. UCID-20014 (1984).
9. S. I. Abdel-Khalik, "On the Role of Fusion-Fission Hybrids in the Nuclear Future of the United States and the Western World," *Nuclear Technology/Fusion* 3, 53 (1983).
10. J. A. Maniscalco et al., *The Fusion Fuel Factory—An Early Application of Nuclear Fusion that Makes Sense*, Lawrence Livermore National Laboratory, Rept. UCRL-87801, Rev. 1 (1982).
11. W. Haefele, Ed., *Energy in a Finite World: A Global Systems Analysis*, (Ballinger, Cambridge, Massachusetts, 1981).
12. R. L. Civiak, *The Economic Competitiveness of Breeder Reactors Compared to Light Water Reactors*, Congressional Research Service, U.S. Library of Congress (1982).
13. B. G. Chow, *Economic Comparison of Breeders and Light Water Reactors*, Pan Heuristics, Marina del Rey, California, Rept. PH78-11-220-34 (1979); also see B. G. Chow, "Comparative Economics of the Breeder and Light Water Reactor," *Energy Policy*, December 1980, p. 293.
14. M. K. White and E. T. Merrill, *Evaluation of the Commercial Fast Breeder Introduction Date*, Battelle Pacific Northwest Laboratory, Rept. PNL-3597 (1981).
15. D. H. Berwald et al., *Fission-Suppressed Hybrid Reactor—The Fusion Breeder*, Lawrence Livermore National Laboratory, Rept. UCID-19638 (1982).
16. *Preliminary Feasibility Assessment of Fusion-Fission Hybrids. Final Report*, Electric Power Research Institute, Palo Alto, California, EPRI Project RP-1463 (December 1981).
17. *Nuclear Proliferation and Civilian Nuclear Power. Report of the Nonproliferation Alternative Systems Assessment Program*, U.S. Department of Energy, Rept. DOE/NE-0001 (1980).
18. Possible climatic consequences of a global warming trend caused by rising levels of carbon dioxide in the atmosphere are discussed in the September 1984 *Energy and Technology Review* (UCRL-52000-84-9), p. 1.
19. W. H. Arnold et al., "Design Approaches to Achieve Competitive LMFBR Capital Costs," *Power Engineering*, September 1982.
20. J. H. Nuckolls, "The Feasibility of Inertial Confinement Fusion" *Physics Today*, 39, 24 (September 1982); *Economic Potential of Inertial Fusion*, Lawrence Livermore National Laboratory, Rept. UCRL-90758 (1984).
21. *Preliminary Feasibility Assessment of Fusion-Fission Hybrids, Summary Report*, Electric Power Research Institute, Palo Alto, California, EPRI Project RP-1436 (December 1980).

Supracompression of High-Explosive Detonation Products

Reexamination of the equation of state for high-explosive detonation products has revealed inadequacies in our treatment of the equation of state and has posed a troubling dilemma concerning the Chapman-Jouguet model for the detonation state.

For further information contact
Edward L. Lee (415) 422-1316 or
Leroy G. Green (415) 422-6890.

In a number of weapons applications, the detonation products of high explosives are supracompressed, that is, they are driven to compression states significantly greater than those reached in planar detonations. Supracompressed states (also known as overdriven states) are produced by shock reflection and by a convergent geometry. We have reproduced such states in the laboratory by using a two-stage, light-gas gun to shock-compress explosives with flat, high-velocity projectiles (impactors). To measure these supracompression states, we have developed high-precision experimental methods. Our results have forced us to reexamine the accepted equations of state for detonation products.

A review of the work by others^{1,2} on the supracompressed detonation of plastic-bonded explosive (PBX 9404) renewed our interest in supracompressed detonation and the equation of state (EOS) of detonation products. Detonation products generally are complex mixtures,

the molecules of which can interact in many different ways and with different intermolecular potentials. It has been noted that the Jones-Wilkins-Lee (JWL) EOS,³ derived from cylinder-test data⁴ using currently accepted values of the Chapman-Jouguet (C-J) pressure, does not predict the behavior of pressure vs relative volume for the supracompressed states. (The C-J pressure is the pressure at the plane behind the detonation front where the explosive has reacted completely to its equilibrium products. See the box on p. 28.)

We reexamined the reported results by performing supracompressed-state detonation measurements with new high-precision experiments.⁵ In doing so, we made the most extensive set of measurements on PBX 9404 (94% HMX, 3% NC, 3% CEP) and a less extensive set on the high explosives LX-07-2 (90% HMX, 10% Viton) and RX-26-AF (45% TATB, 48.5% HMX, 6.5% Estane). Here, we focus primarily on the PBX 9404 results.

Experimental Methods

Overdriven detonation experiments performed in the United Kingdom in the early 1960s employed explosively accelerated metal flyer plates, or impactors, that impinged on a similar buffer plate to generate a powerful shock in the buffer plate.⁶ The transit of the shock transmitted to explosive specimens mounted on the buffer plate was monitored by means of pins. Some years later, investigators at the Ballistics Research Laboratory in Aberdeen, Maryland, used a similar flyer-plate technique but monitored shock histories with a streak camera.¹

Our experiments, using a two-stage, light-gas gun and an ultrafast pin technique, have yielded more precise results. The methods we used to generate shocks at supracompressions in the test explosive had been used

previously to study the high-pressure properties of inert solids and liquids. Our impactors were carried by projectiles fired from the two-stage, light-gas gun (Fig. 1a) at velocities as high as 7.5 km/s. Supracompressed detonation was induced in the target explosive either by direct impact of the explosive by the high-velocity impactors or by the transfer of the shock to the explosive through an intermediate buffer plate.

We used two kinds of buffer plates, 1100 aluminum and an explosive plate matched to the target explosive. We have found that an explosive plate significantly reduces systematic error in calculating the final particle velocity. In addition to minimizing the effects of uncertainties in the Hugoniot of the aluminum buffer plate, the explosive buffer plate eliminates uncertainties in the effect of material strength on the

Fig. 1

(a) The two-stage, light-gas gun uses a charge of gunpowder to accelerate a heavy projectile (piston) into a long, smooth-bore barrel filled with hydrogen gas. The compressed hydrogen ruptures a diaphragm at the end of the barrel and accelerates a lightweight projectile (impactor) down a smaller-diameter, smooth-bore barrel into the target explosive. (b) Impactor and buffer-target assembly. The 1100 aluminum impactor strikes the buffer plate, creating a shock wave that detonates the target explosive. As the 13 pins are crushed by the shock wave, they short and send signals through the cables to oscilloscopes, which accurately measure the arrival times of the shock wave at different target levels and locations.

release adiabat.² Thus, even though its dimensions are somewhat less certain than those of the aluminum plate, the explosive plate reduces net error. (However, we have thus far observed no systematic difference between the data obtained with either plate.)

Shock arrivals were monitored by an ultrafast pin technique. Figure 1b illustrates the projectile-impactor system we used for many of the experiments and a typical target configuration.

Assembly of Target Explosives

To obtain the full potential of the experiment, we must precisely control the thickness, flatness, and surface finish of the target explosive. Any error in dimension exceeding a few micrometres contributes to the total error of the experiment. The target explosives used in our tests were compacted in high-pressure dies and then machined to the required dimensions for the test specimens. The densities were measured before machining. We then lapped and polished the test specimens to obtain a mirror-like surface finish, necessary for attaining the maximum possible accuracy in the experiments. Careful lapping creates surfaces that are flat to at least $1\text{ }\mu\text{m}$.

Assembly of the target explosive is a meticulous process. We measure the thickness of the buffer plate to the nearest micrometre at specific index points to permit calculation of surface flatness. The thickness of the target explosive specimen is similarly measured. We join the two with a very thin (less

than $3\text{ }\mu\text{m}$ thick) adhesive bond and then repeat all measurements.

The mating of the buffer-target assembly to the 13 measurement pins must be done with great care. The anodized aluminum pin insulation must be very thin and fragile to achieve the desired 0.2-ns uncertainty in pin closure time when the shock wave hits the pin. As time differences are referenced to a common fiducial timing signal, we must measure the transmission times of all pin cables and delay lines to about 0.01 ns. Both the pin and fiducial signals are recorded on Tektronix 519 oscilloscopes, calibrated for each experiment with a 200-MHz sine wave. Each oscilloscope receives signals from two pins, located at two different levels, plus the fiducial signal. The time resolution for each signal, as measured from the fiducial, is about 0.5 ns. Figure 2 shows a typical oscilloscope trace of the fiducial signal and two pin signals.

The time after the fiducial signal to the first signal received from each pin, the cable-length corrections, the thickness of the target explosive at each pin position, and the coordinates of each pin referenced to the center pin are entered into a computer, which computes the tilt angle, tilt axis, shock transit time, average shock velocity, and standard error. The shock transit time includes a correction for the symmetric parabolic distortion of the impactor, which is determined by the difference between the shock arrival time for the center pin and for the average of the other pins on the same level.

For an 1100 aluminum impactor 3 mm thick, the time difference between the center pin and the average of the pins on a 9-mm radius is usually 2 ns or less but can be greater. Uncertainty in impactor distortion contributes significantly to the overall uncertainty in shock velocity. To compute the particle velocity and pressure in the target explosive, we use the measured shock velocity, the known Hugoniot equations of state of the impactor and buffer plate (when not the same as the target explosive), the densities of the impactor, buffer plate, and target explosive, and the projectile velocity, which is measured to better than 0.1% by a flash x-ray technique.

Fig. 2

Typical signal output from the pins on a target explosive. The first signal shown is a fiducial common to all seven oscilloscope traces; the second is from the center pin of the target, and the remaining five are from pins on the second level. A sine wave is used to determine the sweep rate of each oscilloscope. The time resolution for each signal is about 0.5 ns.

The Region of Supracompression

At supracompression (pressures much greater than the C-J detonation pressure), the experimental results for all the explosives measured thus far show a systematic difference when compared with calculations based on the usual JWL EOS. Figure 3 shows typical results for the PBX 9404 and PETN high explosives.¹

For both explosives, it is a relatively simple matter to adjust EOS parameters to match the data at very high pressures. At supracompression, the dependence of shock velocity on particle velocity becomes very sensitive to the Grueneisen parameter, γ . This parameter is a measure of the thermal component of the pressure in the EOS. If, instead of being constant, γ is a function of volume V , then $\gamma(V)$ at very high pressure can

Fig. 3

Results of supracompressed detonation experiments on PBX 9404 and PETN compared with JWL EOS calculations. (a) For PBX 9404, there is a systematic difference between the experimental data and calculations from the accepted C-J detonation pressure (37 GPa). The black data bars, with an uncertainty of one standard deviation, are from Ref. 1; the more accurate colored bars, which span two standard deviations, are from our experiments. Decreasing the C-J detonation pressure to 34 GPa enables us to match the data at large supracompression. (b) For PETN, in contrast, the measured C-J state, the expansion states, and the supracompressed states are reasonably consistent with the calculated curve (normalized to the PETN density used in the data, which are from Ref. 1.)

Deriving an Equation of State

In designing high explosives for weapon applications, we must have an accurate EOS to describe the detonating explosive's energy-release history. The explosive is first compressed to about 75% of its original volume by the shock wave of the detonation and is simultaneously converted to its highly compressed gaseous detonation products. The rapid expansion of these compressed gases, far beyond the original volume of the explosive, is what enables the explosive to do work, such as accelerating metal.

The EOS is most easily formalized as the relationship between the pressure generated in the explosion products as a function of the volume occupied and the consequent kinetic and internal energy of the explosion products. If this functional relationship accurately represents the behavior of an explosive, we can predict the acceleration it will impart to a metal shell in a particular geometric configuration.

A wide variety of EOS functional descriptions has been used to characterize explosive performance. Some, derivable from first principles, are quite simple in form; others include empirical elements and are more complex. The Jones-Wilkins-Lee (JWL) EOS, the one most generally accepted and widely used for weapons hydrodynamic calculations, has the form:

$$P = A(1 - \omega/R_1 V') \exp(-R_1 V') \\ + B(1 - \omega/R_2 V') \exp(-R_2 V') \\ + \omega E/V',$$

where P is pressure, V' is the ratio of the detonation-product volume to the original volume of the undetonated explosive ($V' = V/V_0$), and E is the internal energy per unit volume. To obtain the EOS of a specific explosive, it is necessary to determine the values of the five coefficients (A , B , R_1 , R_2 , and ω).

The coefficients of the JWL EOS have been determined for a large number of explosives, some wholly from experimental data and some from a

combination of experiments and complementary calculations.

For ease in discussing the EOS of explosive detonation products, we break the Hugoniot curve (see figure) into three distinct regions. The pivotal region, known as the Chapman-Jouguet (C-J) point, is the maximum steady-state detonation pressure (P_{C-J}). The region above this (colored portion of the curve) we call the supracompression region (or overdriven detonation) because only detonations driven by a large overshock or convergent geometry can reach states represented there. Below the C-J point is the expansion adiabat.

The profile of a steady-state detonation consists of a leading shock-wave front followed by a chemical reaction zone in which the shocked explosive reacts and turns into hot compressed gases. The C-J state lies between the reaction zone and the nonsteady rarefaction or Taylor wave, in which the reaction products expand and cool. Briefly, the C-J state is the place in the detonation front where chemical reaction is complete, the products are in equilibrium, and the flow is sonic.

The measurements that enable one to determine the coefficients of the JWL equation are:

- The explosive's original density, ρ_0 .
- The total chemical energy available, E_0 , obtained from detonation calorimetry.
- The Chapman-Jouguet conditions, which include the detonation velocity, D , and the maximum steady-state detonation pressure, P_{C-J} .
- The metal-acceleration capability, derived from the radius-time history of a rapidly expanding copper cylinder driven by the detonating explosive.

Precise data in the expansion region currently come from cylinder tests. In such tests, explosive packed into a copper tube is initiated at one end by an explosive lens that generates a planar detonation wave. We measure the radial expansion of the cylinder as it is moved by the expansion of the detonation products. This provides a measurement,

although not a direct one, of the expansion state.

Precise experimental measurement of the C-J pressure in a detonating explosive is very difficult. *In situ* pressure gauges can measure the C-J pressure to about 5%. Particle velocity can be similarly measured to about 2 or 3%, from which the C-J pressure may be calculated with essentially the same accuracy. The most common method is to measure the explosion-induced shock or particle velocity in a reference material with known properties, from which the C-J pressure can be inferred. This method incurs an error of several percent from the approximations required to solve the equations and additional errors from differentiating position-time data and extrapolating the results. New methods, using laser velocity interferometers and hydrodynamic-code calculations, should reduce the error in a C-J pressure measurement to one percent or less.

Measurements above the C-J region require precise experiments using a variety of guns in which projectiles are accelerated to high velocities and impact high-explosive samples to achieve a supracompressed condition.

Complementary to these measurements is the limiting thermodynamic energy condition (determined by the explosive's elemental formula, heat of formation, and density), which the EOS must obey. Some of the performance parameters can be estimated from a thermodynamic-hydrodynamic code. These, in turn, can help to set the value of the JWL coefficients.

The EOS of a high explosive is derived by combining the best available data of the kinds described above. The detonation pressure and detonation velocity fix a particular point at the start of the energy release, and the thermodynamic energy fixes the total energy available. The metal-acceleration history of an exploding cylinder further defines the energy release path. Through a series of iterations and supporting two-dimensional calculations, we then determine the exact parameters of the JWL EOS that best predict the cylindrical expansion observed.

The cylinder test itself, however, is insensitive to the C-J pressure and to the early expansion of the detonation products. A range of EOS values will fit the same cylinder-test results, depending on the C-J pressure used.

be adjusted to match the data. The values of γ remain within the range of theoretically acceptable values.

The Region Near C-J Compressions

Only an unsatisfactory choice of alternatives in our JWL EOS makes it possible to reproduce the behavior of both compression and expansion states for PBX 9404 in hydrodynamic calculations. One choice is to decrease the generally accepted value of the detonation pressure P_{CJ} by about 10% (for example, from 37 to 34 GPa). This change makes it possible to calculate both expansion and supracompression states. It fails, of course, to calculate the particle velocity (U_p) measured in experiments designed to yield P_{CJ} . (The detonation pressure equals the product of the particle velocity, the detonation velocity, and the initial explosive density. Since the detonation velocity and the initial explosive density are well known and fixed, any assumed change in P_{CJ} shifts the value of the particle velocity at the P_{CJ} point away from the experimental values.)

Fig. 4

Adiabatic Γ (a function that is proportional to the rate of change of pressure with respect to volume at constant entropy) for shock states and expansion states of PBX 9404 detonation products. By changing the detonation pressure from 37 to 34 GPa in the JWL EOS, it is possible to calculate the supracompression states.

Another—unsatisfactory—choice is to describe the EOS of PBX 9404 with arbitrary terms that yield pressure derivatives exhibiting extreme variations near the C-J compression point (Fig. 4). It is theoretically possible to match the resulting EOS with all the observations in this manner. However, such an EOS would be completely arbitrary, with properties that could not be explained on the basis of an equilibrium C-J state.

The PETN data provide a striking contrast to these results. In PETN, the C-J state measurements, the expansion states, and the supracompressed states are all consistent with a single EOS containing a variable γ (Figs. 3b and 5).

The Chapman-Jouguet Dilemma

Figure 3a compares the results of the supracompressed PBX 9404 experiments with the data from Ref. 1 and with the JWL EOS calculations. As can be seen, our results are consistent with those of Ref. 1. In fact, the new measurements in the region near C-J compression bring the difficulties of interpretation into even sharper focus.

The generally accepted description of the C-J condition requires a triple tangency between the Hugoniot, the isentrope, and the Rayleigh line in the pressure-volume plane, as illustrated in Fig. 6a. (The Hugoniot function specifies the locus of states that are possible immediately after the passage of the shock front. The Rayleigh line graphically represents the conservation conditions along the shock front. The isentrope is the path that detonation products follow as they expand and work is extracted from them.) To satisfy this triple-tangency condition, the locus of supracompression states in the U_s - U_p (shock velocity vs particle velocity) plane (Fig. 6b) must describe a curve whose slope dU_s/dU_p approaches zero as U_s approaches its lower limit, the detonation velocity. This is true regardless of the EOS, so long as it and its derivatives are continuous.

The results of supracompression experiments, with the exception of those for PETN, are at variance with the above criteria. We find the slope dU_s/dU_p is

greater than zero at points very close to the C-J compression values. If, as shown in Fig. 3a, the C-J pressure (nominally 37 GPa) is arbitrarily lowered to 34 GPa, the calculated locus can be brought into agreement with the data.

Such an EOS and an associated C-J pressure are inconsistent with accurate plane-wave experiments designed to measure the C-J pressure. There appears to be a systematic difference between impactor-supported and self-supported detonation waves. We do not understand the exact nature of this difference. However, two observations are worthy of note.

First, the discrepancy for the explosive PETN reported in Ref. 1, if it exists at all, appears to be much smaller than that for PBX 9404, LX-07, or RX-26. (For PETN, the generally accepted value of C-J detonation pressure is 31.7 GPa at a density ρ_0 of 1.765 g/cm³.) Moreover, PETN is a relatively oxygen-rich explosive whose detonation products contain considerably less carbon than the other explosives.

Second, in an impactor-supported shock wave, changes in the state of the detonation products are much more readily communicated to the wave front than in the case of a self-supported wave. Thus, if there are reactions in the detonation products that reduce the pressure, the pressure change will reach the front in the supracompression measurements more rapidly than in the C-J measurements.

These two observations suggest that irreversible, but relatively slow, carbon-condensation reactions may occur in the shock-front reaction zone, releasing heat but decreasing the pressure. The effect of such a phenomenon on C-J measurements would be difficult to determine, as the pressure changes are not large and the pressure drop would act as a rarefaction, thereby communicating slowly, if at all, with the detonation front. In the impactor-supported tests, in contrast, the pressure drop produced by the reaction should travel toward the shock front at the speed of sound (5 to 6 mm/ μ s) in the detonation products. To describe this behavior, we will have to incorporate

Fig. 5

In contrast to the PBX 9404 data, the PETN data can be described by a single EOS with a variable γ . The initial density of PETN is 1.725 g/cm³, and the detonation pressure P_{C-J} is taken as 30 GPa.

Fig. 6

Illustration of C-J and supracompressed shock states in (a) the pressure-volume plane and (b) the shock-velocity/particle-velocity plane. The tangency of the Rayleigh line in (a) with the Hugoniot of the reaction products defines the C-J condition.

rate-dependent condensation processes in the detonation models.⁷ We are currently investigating such model modifications.

Summary

We have measured the high-compression behavior of detonation products to considerably higher precision than before. We have incorporated this information in our hydrodynamic EOS models, albeit with reservations about the C-J model. In this regard, we now have strong evidence of a conflict between impactor-driven and self-supporting detonative shock waves. Almost certainly, we will need to develop more general models for the detonation state. We are currently conducting further experimental and theoretical investigations of the behavior of shock waves and rarefactions in the supracompression region. □

Key Words: Chapman-Jouguet pressure; detonation product; equation of state; high explosive—LX-07-2, PBX 9404, PETN, RX-26-AF; overdriven detonation; shock compression; supracompression.

Notes and References

1. J. H. Kineke and C. E. West, presented at *Fifth Symp. (Intl.) on Detonation*, U.S. Naval Ordnance Laboratory, Pasadena, California, Repts. ACR-184 and ONR-0851-0057 (1970).
2. M. Van Thiel, E. L. Lee, and S. Cochran, "Effects of Overdriven Shock States on the Equation of State of 9404 Explosive," *J. Appl. Phys.* **54**, 11 (1983).
3. E. L. Lee, H. C. Hornig, and J. W. Kury, *Adiabatic Expansion of High-Explosive Detonation Products*, Lawrence Livermore National Laboratory, Rept. UCRL-50422 (1981).
4. H. C. Hornig, E. L. Lee, M. Finger, and J. E. Kurie, presented at *Fifth Symp. (Intl.) on Detonation*, U.S. Naval Ordnance Laboratory, Pasadena, California, Repts. ACR-184 and ONR 0851-0057 (1970).
5. E. L. Lee, M. Van Thiel, L. G. Green, and A. Mitchell, "Detonation Product EOS: The Region Above Chapman-Jouguet Pressure," in *Shock Waves in Condensed Matter*, J. R. Asay, R. A. Graham, and G. K. Straub, Eds. (Elsevier Science Publishers, The Netherlands, 1984).
6. I. C. Skidmore and S. Hart, "The Equation of State of Detonation Products Behind Overdriven Detonation Waves in Composition B," presented at *Fourth Symp. (Intl.) on Detonation*, U.S. Naval Ordnance Laboratory, White Oak, Maryland, October 12-15, 1965.
7. E. L. Lee and C. M. Tarver, "Phenomenological Model of Shock Initiation in Heterogeneous Explosives," *Phys. Fluids* **23**, 12 (1980).

Entering the World of Multiprocessing

As a direct way to increase the speed of computation, to expand the size of problems it is practical to attempt, and to decrease the real time required to calculate a problem's solution, we are designing software methods that enable many computers to share the work on a single problem. Our new Cray X-MP contains two independent CPUs within the same mainframe, and this number is expected to increase to 16 or more by 1990. Much of our multiuser networking experience is directly applicable to this problem.

Contact: Harry L. Nelson (415) 422-4204 or Pierre J. Du Bois (415) 422-4007.

Uranium Resources and the Development of Fission and Fusion Breeder Reactors

According to present estimates, economically useful supplies of uranium will be exhausted in about 60 to 70 years if the U.S. nuclear power industry recovers from its present slump. To allow time to build enough replacement plants, it will be necessary to develop some other source of energy long before this happens. Our studies indicate that it will be difficult to develop an economically competitive fission breeder reactor in time to ease the shortage. Hybrid fusion reactors, however, each of which might make enough fuel for 5 to 20 light-water reactors, appear to offer a safe and cost-effective alternative. Research to develop the breeding blankets for such hybrid reactors should start as soon as possible.

Contact: Claire E. Max (415) 422-5442.

Supracompression of High-Explosive Detonation Products

The results of new, high-precision experiments with our light-gas gun, in which the detonation products of high explosives are driven to compression states considerably higher than those reached in planar detonations, have forced us to reexamine our treatment of the equation of state for high-explosive detonation products and have revealed ways in which it was inadequate. This finding poses a troubling dilemma concerning the Chapman-Jouguet model for the detonation state. There appears to be a systematic difference between impactor-supported and self-supported detonation waves, the nature of which is still under investigation.

Contact: Edward L. Lee (415) 422-1316 or Leroy G. Green (415) 422-6890.

Recent Titles

Articles published in recent issues of the *Energy and Technology Review* are grouped below according to their chief sponsors, the Assistant Secretaries of the U.S. Department of Energy. Research funded by other Federal agencies is listed under Work for Others.

ENERGY RESEARCH

Uranium Resources and the Development of Fission and Fusion Breeder Reactors (October 1984)
Climatic Warming and Carbon Dioxide (September 1984)
Response of Plants to Increased Carbon Dioxide (September 1984)
Energy-Flow Diagrams: An LLNL Contribution to Energy Analysis (May 1984)
Magma Evolution (May 1984)
Single-Chromosome Gene Libraries of Human DNA (Brief—May 1984)

FOSSIL ENERGY

Computer Simulation of Particulate Flow (May 1984)
Dispersion Tests of Liquefied Gaseous Fuels (Brief—May 1984)

DEFENSE PROGRAMS

Military Application

Supracompression of High-Explosive Detonation Products (October 1984)
Raman Spectroscopy of Shocked Water (Brief—April 1984)
Silicone Cushions for Engineering Applications (April 1984)
Synchrotron-Radiation Research Facility (Brief—April 1984)

LABORATORY REVIEWS

Entering the World of Multiprocessing (October 1984)
Earthquake Safety at LLNL (August 1984)
Geology of the Livermore Valley (August 1984)
Monitoring Groundwater Quality at Site 306 (August 1984)
State of the Laboratory Issue (July 1984)
LLNL in the Developing Livermore Valley (April 1984)

WORK FOR OTHERS

Defense Nuclear Agency

Simulating Nuclear Cratering Events at the Pacific Proving Ground (September 1984)

Nuclear Regulatory Commission

Earthquake-Hazard Analysis for Nuclear Power Plants (June 1984)
Human Factors in Reactor Facility Design (June 1984)
Reactor-Coolant Piping: Realistic Failure Assessment (June 1984)

United States Geological Survey

The Strength of Ice (April 1984)

Disclaimer

This document was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor the University of California nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial products, process, or service by trade name, trademark, *manufacturer*, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government thereof, and shall not be used for advertising or product endorsement purposes.

Printed in the United States of America
Available from
National Technical Information Service
U.S. Department of Commerce
5285 Port Royal Road
Springfield, Virginia 22161
Price: Printed Copy \$8.50, Microfiche \$4.50