Energy and Technology Review

May 1988

Scientific Editor

Andrew J. Poggio

Editorial Staff

Lauren de Vore	Robert D. Kirvel
Kevin Gleason	Nona M. Sanford
Derek P. Hendry	Shirley O. Tatt

Art Staff

Lynn M. Costa	Raymond Marazzi
Paul M. Harding	John Zych
George Kittinos	

Compositor

Louisa Cardozo

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 or the University of California and shall not be used for advertising or product endorsement purposes.

Measuring the Performance of Supercomputers

1

We analyze the problems of evaluating supercomputer performance and present the results of measurements with our small benchmark: 7 samples of Fortran code.

Microwave Tokamak Experiment

15

We are using the free-electron laser as a microwave source for the tokamak to provide localized heating in the magnetically confined plasma and to evaluate heating and current-drive physics.

Gas Cooling of Laser Disks

25

We have explored the possibility of cooling laser disks by gas flow and conclude that this method is practical for cooling solid-state megawatt lasers.

Abstracts

35

Prepared for the United States of America
by
National Technical Information Service
U.S. Department of Commerce
5285 Port Royal Road
Springfield, Virginia 22161
Price: Unclassified printed copy, \$07.50; microfiche, \$05.00.

UCRL-52000-88-5
Distribution Category: UC-2
May 1988

MASTER

DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

Measuring the Performance of Supercomputers

We have developed efficient means of estimating supercomputer performance for typical workloads of numerically intensive physics computations.

For further information contact
Frank H. McMahon (415) 422-1647.

An important part of our mission at LLNL is the application of the most advanced computers to solve problems related to national defense. New computers appear or are proposed in such abundance that we are continually called upon to evaluate them. Over the last decade, we have developed an effective procedure for the preliminary evaluation process.

In evaluating a new or proposed computer, we cannot merely read the vendor's specifications, since they are generally not written in terms that make it possible to predict actual performance in a typical Laboratory computing environment. We must devise a test of the important computing demands by realistically representing computations that may be performed and by exercising all of the new computer hardware capabilities.

"Benchmarking" is the term given to the process of evaluating an entire computer system, or some aspect of it, by having it perform a typical set of realistic computations. The task or test itself is termed a benchmark. (This term and others are defined in the glossary on p. 13.) Table 1 lists some of the different methods that have been used to define or benchmark computer performance, with steeply increasing sample sizes and cost requirements.

Benchmarking the performance of new computers with upwardly compatible hardware and software systems presents few problems. However, benchmarking new types of high-performance computers, particularly supercomputers, has become increasingly complex. In these computers, the efficiency of both application coding and compiled coding depends greatly on the effective use of ever-larger combinations of parallel hardware resources. At present, the sensitivity of supercomputer performance to programming efficiency ranges up to two orders of magnitude. This makes it very difficult to demonstrate the future value of a new supercomputer before a maturely programmed workload and compiler are developed.

In the traditional method of benchmarking large numerical applications, the performance of old programs on new machines is measured. This approach produces a definitive evaluation only for those

users who are not seriously interested in exploiting new capabilities. Fortran compilers cannot be expected to transform old computer programs into the most efficient parallel forms using the most appropriate mathematical methods. The effective use of new parallel capabilities goes beyond direct conversion of the application programs. At the least, some rethinking and reprogramming of the applications (and benchmarks) is required.¹ Experience has shown that restructuring a large numerical application for vector processing may require a year-long effort and the expertise of the program's originators.^{2,3} The performance of these new programs or benchmarks must then be generalized to infer the future performance of the entire workload. The magnitude of these efforts exceeds the scope, expertise, and time available for traditional efforts.⁴

Finally, as new computational methods are developed, the character of the workload changes and old

Table 1. Methods of benchmarking computer performance.

Method	Lines of programming
Cycle time or add time	1
Single function (dot or matrix product)	10
Kernels of source code (Whetstone, LFK, etc.)	10 ³
New application prototypes	10 ⁴
Current numerical applications	10 ⁵ -10 ⁶
Entire workload	10 ⁶ -10 ⁷

benchmark projections are soon outdated. Thus, traditional benchmarking methods are not sufficient to demonstrate accurately the future performance of innovative applications on innovative computer systems. Since definitive evaluations of supercomputers usually are not possible, we have developed a small, preliminary computer test that measures the effective numerical performance range with very little effort. Inferences based on these test data have proven to be as accurate as inferences based on large-scale benchmark tests. We refer to our test as a benchmark, even though it is smaller than other benchmarking programs.

Benchmarking CPU Performance

Performance measurements of a central processing unit (CPU) can be a sufficient test of computer capacity when the workload is dominated by long-running numerical applications. These computations have been CPU-limited for decades (rather than, for example, limited by the input/output or I/O device). The enormous

increase in working memory sizes in this decade (two to three orders of magnitude) ensures that the performance of numerical applications will be more CPU-limited than in the past (Figure 1). Efficient time-sharing operating systems will also ensure that performance is more CPU-limited. This limitation is the driving force behind the development of new, faster architectures (supercomputers).

The advent of parallel CPU facilities has introduced a large variation in the speed of execution of different CPU instructions, especially in the CPUs of vector supercomputers. Effective parallel performance depends on an appropriate structuring of the Fortran program and on an effective optimizing compiler. Therefore, realistic analyses of computer performance must be based on a broad sample of compiled Fortran constructs covering the entire performance range.

The Livermore Fortran Kernels

The Livermore Fortran Kernels (LFK) benchmark is a set of Fortran

code samples designed to test representative computer performance on physics workloads. We developed it in 1970 to study the efficiency of CPU code generated by the LRLTRAN compilers⁵ for the CDC-6600 and CDC-7600 computers. The number of separate code samples was originally 12, later 24. The kernels make up a benchmark designed to characterize the usage of Fortran in computational physics programs. They include real samples of floating-point computations taken from diverse scientific applications, not the synthetic constructions of the Whetstone benchmark, for example.⁶ Many kernels represent widely used, generic computations, e.g., dot and matrix (SAXPY) products, polynomials, first sum and differences, first-order recurrences, matrix solvers, and array searches. The kernels incorporate generic Fortran computations and their basic code blocks. Some kernels typify often-used Fortran constructs and some contain Fortran statements that are difficult to compile into efficient machine code. The kernels were selected to represent the best and near-worst case Fortran practice so that a realistic, floating-point performance range could be measured. Also, each kernel characterizes a large class of loops that have approximately the same size, structure, and hence performance. The Fortran samples and all available CPU performance data are documented.⁷

The kernels are small enough to be quickly coded in assembly language and timed on new computers that are still in the design process. A CPU performance analysis can begin as soon as the CPU instruction set and instruction timings are defined in the design proposal. When a new machine

Figure 1. Evolution of memory size and speed. The enormous increase in working memory sizes in the past decade (two to three orders of magnitude) ensures that the performance of numerical applications will be CPU-limited.

comes into use, the same kernel tests can also measure the improving performance of the Fortran compiler as it matures.

How many lines of Fortran must be benchmarked to estimate the performance of most scientific applications? We think that, thanks to the generic character of many Fortran numerical expressions, the sample should be on the order of 1000 lines, not 10 000 lines. For example, in Kuck's study of concurrency in Fortran programs, his sample of 20 small programs contained less than 1000 lines of Fortran.⁸ Also, Knuth's empirical study of Fortran programs concluded that "a small number of basic patterns account for most of the programming constructions in use."⁹ The length of a satisfactory sample also depends on the complexity of the hardware architecture. Performance of most of the early, relatively simple, sequential computers was predicted by simple tests, such as the dot-product function test.

Compared to real applications and to typical benchmark programs, the LFK benchmark is very much smaller, more portable, and requires less effort to install and run. Table 1 lists some alternate methods for measuring computer performance that range in size from one machine instruction to millions of lines of code for the workload of a single computer center.

Metric for CPU Performance

Frequently, CPU performance is quoted as the maximum hardware instruction rate in millions of instructions per second (Mips) or as the measured rate for compiled, abstract instructions (Whetstones per second),⁶ also in Mips. However, the Mips rate is not a satisfactory metric for numerical computation because

a given computation can be accomplished in different ways by different instruction sets, of which there are many different types. In extreme cases, a computation that takes several thousand instructions in one computer may be accomplished by a single instruction in a vector computer.

Researchers in computational physics use systems of partial differential equations from mathematical physics to simulate the evolution of physical systems.² The mathematical methods depend on real-valued functions, and the algorithms are programmed, almost exclusively, in Fortran floating-point operations (flops). These floating-point operations are the necessary operations on any computer: the total number, for a particular simulation, is invariant. Thus, a meaningful computation rate can be measured by counting the number of floating-point operations and then dividing by the total execution time of a program. The first version of the LFK benchmark (1970) originated the CPU performance metric of millions of floating-point operations per second (Mflops).

Fortran Performance Variance

The execution rates of vector computers extends over two orders of magnitude,¹⁰ but CPU performance measurements are commonly reduced to a single Mips or Mflops value with no indication of variance. If the Mips or Mflops value is a peak rate, the vendor specification can be quite misleading. For example, the Cyber-205 literature rates this machine at 800 Mflops.¹¹ However, a simple run of the LFK benchmark reveals that some scalar Fortran computations execute at less than 1% of the quoted rate. Thus, thorough

computer evaluations require measurements of the variance in CPU performance. In our benchmarking, we estimate variance from the 24 kernel execution rates. (Variance should always be quoted along with the performance-range statistics.)

The CPU performance of the LFK samples depends on the efficiency of the Fortran programmer's technique, on an optimizing compiler, on an effective operating system, and on sufficient high-speed CPU facilities. Naturally, the performance range measured by the LFK benchmark will improve as the computer system evolves and will peak when all subsystems are mature.

One cannot expect to measure a static or definitive value for computer performance early in the computer's life cycle, when it is most needed for assessment. This uncertainty is greatest for new, highly innovative supercomputers. Given the large number of tradeoffs in the design of the subsystems noted above, an intrinsic or unique value for computer performance lacks meaning. In this context of uncertainty, the CPU performance range measured by the LFK benchmark has reasonable accuracy and the variance is realistic.

Early Results

One of the first performance results from executing the original 12 LFKs on the (then-new) CDC-6600 and CDC-7600 was that the average speedup and variance (from 6600 to 7600) for the kernels was consistent with the average speedup and variance of major numerical applications in the workload (4.5 ± 1.5). This result suggested that the LFK samples adequately represented the critical CPU behavior of numerical applications and that the

benchmark could be used to predict the CPU performance speedup of new computers.

Coverage of the performance range by the LFK benchmark can be demonstrated by plotting the CPU rates of the kernels as a function of the number of floating-point operations per loop (see Figure 2). The number of floating-point operations between loops is one measure of local parallelism within the bounds of a basic code block in a machine. For example, code blocks with many operations per loop can sustain concurrent operations for scalar or vector instructions (chaining).

It is clear from Figure 2 that on these "pipelined" CPUs, the net performance of loop computation generally correlates with the volume of computation in the loop. The larger loops approach an upper limit in execution speed, both in the scalar

and vector cases, as the CPU resources become saturated with concurrent work.

The LFK loops in Figure 2 cover a range of code block sizes, and enough to measure the extremes of scalar and vector computation rates, which are the effective performance range for Fortran programs. Note that the scalar and vector measurements span performance range.

From these results, the LFK benchmark appears to be broad enough to cover the effective CPU performance range. Moreover, these measurements show that the performance of a numerical workload on the present generation of computers is characterized by the basic code block sizes used, as well as by vectorizability. These characteristics of Fortran source programs are more incidents of programming style than requirements of physics models. Thus, the

performance of one physics application is not a sure guide to the performance of other physics applications, because the applications may be programmed differently. The key to characterizing Fortran performance is to represent a diverse set of common Fortran practices; this is the intention of the LFK sampling.⁷

Harmonic Mflops Model of a Numerical Program

We have seen, from Figure 2, that loop structures generally execute at stable, characteristic rates that are proportional to the number of floating-point operations per branch (that is, the number of floating-point operations in a basic machine code block). Since a program is a collection of elementary code blocks, the CPU performance of a program may be characterized by the performance of its basic code blocks.

Therefore, the net computation rate of a program is simply the total number of floating-point operations N divided by the total CPU time T (remember that N is constant for a given problem). The CPU rate of an individual basic code block BCB_k is

$$\rho_k = n_k / t_k,$$

where n_k is the total number of floating-point operations computed by BCB_k and t_k is the total time required by BCB_k . The ρ_k values are generally stable rates but may be subject to some fluctuation due to competition for memory. The net rate of a program containing M basic code blocks is

$$\rho_{\text{net}} = \sum n_k / \sum t_k, \quad 1 \leq k \leq M, \text{ or}$$

$$\rho_{\text{net}} = \sum n_k / \sum (n_k / \rho_k), \quad 1 \leq k \leq M.$$

Thus, the net CPU rate of a program is the weighted harmonic mean of the

Figure 2. Early results (1973) for CPU performance as evaluated by 12 LFK samples. The trend lines indicate that performance increases with the number of floating-point operations per loop and then reaches a plateau.

rates of the basic code blocks, not the arithmetic mean of the rates. Note that the weights are the total computation work of each code block. The weighted harmonic mean is the harmonic Mflops model of a numerical program. The net CPU rate of a workload of numerical applications can be computed, in turn, as the harmonic mean rate of the individual programs.

In the LFK workload sensitivity analysis,⁷ we identify computation work fractions,

$$f_k = n_k/N,$$

where $\sum f_k = 1$. The weighted harmonic mean is

$$p_{\text{net}} = 1 / \sum (f_k / p_k), \quad 1 \leq k \leq M.$$

A two-component form of this harmonic Mflops model is used in the LFK benchmark to show the sensitivity of the net or harmonic mean rate of a program to code optimization. This reduced model is called a SISD/SIMD performance model:

$$\rho(f_k) = 1 / [(1 - f_k)/p_s + f_k/p_o].$$

The subscripts denote scalar and vector (or optimized) terms: p_s is approximated by the harmonic mean rate of the slower half of the LFK ranked by performance rate, and p_o is approximated by the harmonic mean rate of the fastest quartile of the LFK rates.

The LFK model synthesizes the disjoint scalar and vector CPU performance ranges shown in Figure 2. Figure 3 shows the harmonic variation of performance for all degrees of vectorization and for two vector lengths (471 and 19 elements). The CPU performance of vector computers is seen to be a hyperbolic function of vectorization

or optimization. The performance of most Fortran programs will be in the region bounded by the long-vector and short-vector performance curves.

Validating the LFK Benchmark

To assess the validity of the small LFK benchmark, we checked CPU performances on the LFK test against large benchmarks or application programs. Six extracts (over 7000 lines of Fortran) from the LLNL workload were chosen in 1978 to benchmark the Cray-1 computer. In

a separate analysis, we counted the total number of floating-point operations computed by each of the six benchmarks executing on the CDC-7600. Table 2 shows the net computation rate (Mflops) for each benchmark and the equi-weighted harmonic mean rates of the LFK kernels (then 14) run on the Cray-1. The harmonic mean rates for the six benchmarks are remarkably consistent with the harmonic mean rate of the LFK tests and with the net rate of benchmark B2, the most complete and significant benchmark.

Figure 3. The variation of performance of various computers as a function of the degree of vectorization. Two DO loop lengths were tested, 471 and 19 elements; the upper (solid) curve for each band is for the long loop, the lower (dashed) curve is for the short loop. The CPU performance of vector computers is clearly a function of optimization and vector length.

The FPS-264 is a 64-bit mini-supercomputer; it was first benchmarked with the LFK test in 1985. The harmonic mean performance of the FPS-264 is 0.74 of Cray-1 performance and 0.58 of Cray-XMP/1 performance. The harmonic mean usually correlates with and predicts the performance of poorly vectorized old programs. (These programs are sometimes called "dusty decks"—the image is one of taking an old deck of punched cards off the shelf and running it unchanged.) The performance of the first LLNL applications converted to the FPS-264 have confirmed the LFK-predicted performance level.

The speed-up ratios of Cray and FPS-264 performance, plotted in Figure 4, are based on the LFK SISD/SIMD models for these computers. The Cray computers are three to five times faster than

the FPS-264 for computations that are 95% vectorized. When the computations are less than 50% vectorized, the Cray computers are only 1.0 to 1.5 times faster than the FPS-264. In this "dusty deck" domain, it is reasonable to describe the performance of the FPS-264 as "Cray-class."

From Figure 4, we can infer the level of vectorization required for the CPU performance ratios obtained for several LLNL applications programs (e.g., DYNA2D). When the performance ratios are overlaid on the basic relation of Figure 4, the vectorization values are in agreement with the known levels of the codes.

LFK Measures on the LLNL Workload

The performance of an entire computer workload is a very

complex subject encompassing all of the user's demands on the system. In a scientific or engineering environment, however, most time is spent on numerical applications, which collectively are called the production workload. Under an efficient multiprogramming system, a production workload will be generally CPU-limited and its performance should be predictable with the LFK benchmark.

The Cray-XMP floating-point performance range for nine of the largest production applications in the LLNL workload is shown in Figure 5. The overall production-time rate of the Cray-XMP/1 (one CPU) for 1985 was approximately 24 Mflops. This sustained rate, derived from the rates of individual applications, has been verified by hardware monitors measuring all computation for extended periods of production time.¹² These sustained rate measurements vary from 22 to 25 Mflops/processor. Thus, the total throughput for an Cray-XMP/4 (four CPUs) approaches 100 Mflops.

The LFK performance range on the Cray-XMP/1 shown in Figure 5 is the result of an "as is" compilation; that is, there was no vector "tuning" of the Fortran source. However, vectorization of the LLNL workload has been in progress for more than a decade. Using the LFK SISD/SIMD model, we see in Figure 3 that a 24-Mflop net rate implies that the LLNL workload executes 80 to 90% of the computation with vector operations. We confirmed this LFK model prediction using the Cray-XMP hardware monitors, which indicate that close to 90% of the computation is performed by vector operations.¹²

Our timing analysis has shown that, for all practical purposes, the nine large production applications do

Table 2. Test performance (in Mflops) of the CDC-7600 and Cray-1 on the 1978 and LFK benchmarks.

Benchmarks	CDC-7600	Cray-1	
	Fortran as received	Fortran as received	Extensive rewriting
LFK (14 kernels)	3.4	6.7	13.7
B1	3.3	7.7	15.0
B2	3.2	6.8	12.7
B3	4.0	8.6	45.5
B4	2.5	5.1	21.6
B5	3.0	5.6	6.1
B6	2.8	8.9	11.3
Harmonic mean of benchmarks B1-B6	3.11	6.80	12.98
Performance ratios	1.0	2.19	4.17
Proportion of flops computed in vector mode	0 ^a	0.3	0.7

^aThe CDC-7600 is not a vector machine.

indeed characterize the performance of the entire LLNL workload. It remains to be proven whether the LFK benchmark performance can characterize the observed performance range of the nine applications.

The LFK performance statistics, based on the LLNL-written 1985 CIVIC 130i compiler, are plotted in Figure 5. The average performance rates of nine LLNL applications fall near the middle of the LFK performance distribution. A few of the applications are in the top quartile; this is due, in part, to some use of assembly language libraries. Clearly, the LFK benchmark characterizes the effective Fortran performance range of many mature vector applications with a dispersion in performance greater than the workload itself (Figure 5). In this case, the sustained geometric mean rate of LLNL applications of 24 Mflops is near the center of the LFK performance range, midway between the equi-weighted harmonic mean and arithmetic mean rates.

Figure 5 suggests two other conclusions:

(1) The equi-weighted arithmetic mean of the LFK rates correlates with the highly vectorized applications in the workload, because the arithmetic mean is dominated by the high vector rates. The LFK average rate is greater than 40 Mflops, implying over 90% vectorization of the computation shown in Figure 3. Optimists prefer this statistic.

(2) The equi-weighted harmonic mean of the LFK rates correlates with the poorly vectorized applications in the workload, because the harmonic mean is dominated by the low scalar rates. The harmonic mean is about 8.5. This implies only 40% vectorization in Figure 3, which

is consistent with the vectorization of 10 of the 24 LFK kernels in this case. Pessimists prefer this statistic.

These LFK benchmark statistics measure a realistic range within the entire range (1-1000 Mflops) of hardware performance of current vector processors. This effective performance range lies between the harmonic mean and the arithmetic mean rate. It corresponds to the effective performance range of

efficient Fortran programs, e.g., 8-40 Mflops on the Cray-XMP/1 (Figure 5). Many applications and workloads have been and can be predicted to be in this range. Measuring this range with the LFK makes the hardware performance range of vector processors more predictable. The uncertainty is reduced from two or three orders of magnitude to a factor of five, which is about the same as the Fortran

Figure 4. The performance of the Cray-1 and Cray-XMP relative to the FPS-264. The relative performance bands were calculated using LFK test data. The measured speedup of two LLNL application programs, DYN2D and DYN3D, agrees with the LFK test results. The relative performance of the Cray computers is less than twice that of the FPS-264 when vectorization is below 70%.

Figure 5. The Cray-XMP/1 floating-point performance range, as evaluated by the LFK benchmark, and nine of the largest production applications in the LLNL workload. (A, G, and H denote the arithmetic, geometric, and harmonic means. P denotes the peak CPU rate and O denotes a single measurement.)

performance range uncertainty of conventional, scalar computers (see Figure 6). The range of performance exhibited by the LLNL workload cannot be reduced to a single number, and the performance range should be used instead.

LFK Comparisons of Conventional Processors

The performance of the LFK on conventional serial computers, such as the DEC VAX-785 or the IBM 3081K, generally shows a relatively small dispersion or range of computing rates (Figure 6). In these cases, the mean and median rates are nearly equal, and parametric statistics may be used to evaluate the significance of performance differences between different computers. Note that a statistical error in the mean rates of about 10% can be expected as a result of the LFK sample size (24 kernels, 72 samples). One of the central measures (the geometric mean) may be reasonably interpreted as a characteristic rate of computation for a particular computer/compiler system.

The uncertainty in CPU performance measured by the LFK benchmark is indicated by the standard deviation. Conventional computers with mature compilers generally have an LFK performance range in which the minimum rate is not much more than six times slower than the maximum rate. Lower minimum performance has often been observed in new systems under development because of an immature compiler and, more rarely, because of a bottleneck in the hardware architecture. In any case, the machine code produced for the slowest

kernels should be analyzed by the compiler staff for improvements.

LFK Comparisons of Vector Processors

The accuracy of performance ratings of supercomputer CPUs is limited by the variability of large diverse workloads and by the enormous variance in vector CPU performance. We use descriptive statistics and nonparametric statistical methods.¹³

Vector processors have multiple, tightly coupled, unbalanced computer elements and a serial section 10 to 100 times slower than the fast, vector pipelines. The performance of the LFK on vector computers, such as the Cray-XMP/1, generally shows a huge range of CPU performance—more than two orders of magnitude from the slowest to fastest kernel (Figure 7).

Figure 8 compares the performance of the Cray-XMP/1 with that of the NEC SX-2 on the LFK test. The LFK test executes the 24 kernels three times, with average DO loops of different lengths (18, 89, and 468) so we can compare performance on short and long vectors. Nearly three orders of magnitude of supercomputer performance is consistently correlated with very small scatter using the LFK test. Vector performance is at least forty times larger than scalar performance on the NEC SX-2. The relative performance advantage of the NEC SX-2 vs the Cray-XMP/1 increases sharply in the vector domain because the NEC SX-2 has four vector add pipelines and four vector multiply pipelines, whereas the Cray-XMP/1 has only one of each. Clearly, this

Figure 6. The range of CPU performance, in terms of numerical computation rate, as evaluated by the 24 kernels for various computers. In all cases, we used a 64-bit word size, 72 sample timings, and an average DO loop length of 166. (A, G, and H denote the arithmetic, geometric, and harmonic means.)

Figure 7. CPU performance ranges of the 24 kernels on several different vector computers. S denotes an all-scalar test. The kernels were then vectorized and executed three times, with short, medium, and long DO loops. The performance of the LFK on vector computers, such as the Cray-XMP/1, generally shows a wide range of CPU performance — more than two orders of magnitude from the slowest to fastest kernel. (A, G, and H denote the arithmetic, geometric, and harmonic means.)

regression curve exhibits the general character of the relative CPU performance of the Cray-XMP/1 and the NEC SX-2 over the very large ranges measured by the LFK benchmark.

The 1985 LLNL production workload on the Cray-XMP/1 (Figure 7) was measured to have a net, or average, rate of 22–25 Mflops. The geometric mean rate of the LFK executed on the Cray-XMP/1 CFT 1.15 system is 18.7–24.9 Mflops, depending on the average vector length. The close agreement of these mean rates is both useful and fortuitous; other Cray workloads are known to have a different composition and average rates.¹⁴ We may, however, infer that the average performance rate of the specific 1985 LLNL workload can be correlated with the geometric mean rate of the LFK on vector computers, such as the NEC SX-2, at least for the time being. The geometric mean rates of the LFK are 43.9–63.5 Mflops on the NEC SX-2, which implies a speedup of 2.35 to 2.55 times the Cray-XMP/1 on a per-processor basis for this specific 1985 workload.

We should also note that the average throughput rate of the LLNL workload on the four-processor Cray-XMP/4 averaged between 88 and 100 Mflops. Therefore, the Cray-XMP/4 system has a significant throughput advantage relative to the NEC SX-2 for the current LLNL workload.

LFK Comparisons of Multiprocessor Performance

Alliant Computer Corporation was the first developer of a multiprocessor system to demonstrate automatic parallel processing of all the LFK kernels

by a technique called microtasking. Asynchronous parallel execution is initiated by the Alliant compiler, which generates all of the necessary control code within the scope of a DO loop nest. Up to eight processors can be tightly coupled, and each processor has vector operations as well. The effective Fortran use of all these hardware resources was measured by several LFK tests,⁷ summarized in Figure 9.

Executing the LFK benchmark using a single processor of a multiprocessor computer provides a measure of the lower performance bound of the system. An upper performance bound can be projected to be N processors greater. The performance shown in Figure 9 reflects the fact that six of the 24 LFKs are recurrence equations, which require serial coding. Test performance would be better on

Figure 8. Comparison of the LFK-evaluated performance of the NEC SX-2 and Cray-XMP/1 computers. The LFK results indicate that the NEC SX-2 considerably outperforms the Cray-XMP/1. (A, G, and H denote the arithmetic, geometric, and harmonic means.)

Figure 9. CPU performance ranges of the 24 LFK on the Alliant multiprocessor. The arithmetic (A), geometric (G), and harmonic (H) means are shown for scalar processing (S) and for multiprocessing with one, four, and eight processors (FX1, FX4, and FX8, respectively). The curves show the increasing trends of the three means.

applications that permitted more parallel coding. The Alliant FX8 demonstrates that Fortran programming expresses sufficient parallelism for automatic and efficient multiprocessing on up to eight processors without explicit synchronization programming.

Tracking the Evolution of Supercomputers

The evolution of high-performance computers can be tracked using the performance ranges measured with the LFK benchmark. The evolution of computer hardware

capability, in terms of the performance line, the performance of efficiently programmed Fortran applications has been increasing by a factor of ten every eight years. The performance of inappropriately programmed applications has improved much less. In all cases, performance has grown exponentially for several decades. Figure 10 clearly shows that the problem of benchmarking supercomputers is increasingly more difficult because the entire performance range of individual supercomputers is also growing exponentially from less than one

order of magnitude in 1969 to about three orders of magnitude in 1985. Thus, attempts to reduce the performance of a supercomputer to a single number are unrealistic.

Conclusions

Traditional benchmark tests of computer performance are not suitable for computers using vector or parallel processing, where the range of CPU performance is more than ten times the performance range of conventional, serial processors. The effective CPU performance of supercomputers now ranges from a few to several hundred Mflops (million floating-point operations per second). Attempts to reduce this two-order-of-magnitude range of CPU performance to a single number are not valid and have produced much confusion. No single quotation of performance is sufficient.

The Livermore Fortran Kernels (LFK) test contains a broad sample of Fortran computations that can be used to measure an effective numerical performance range, thus avoiding the perils of a single performance "rating." A complete report of LFK test results gives six performance rates: the minimum, the harmonic, geometric, and arithmetic means, the maximum, and the standard deviation. These measurements show a realistic variance in Fortran CPU performance that has stood the test of time as supercomputers, their compilers, and the coding mature. Ultimately, the level of performance of a workload on a supercomputer relates most closely to the level of development of efficient, parallel computer codes

Figure 10. Supercomputer performance has increased dramatically (more than three orders of magnitude) in the past 25 years as computers, and particularly compilers, have evolved. (A, G, and H denote the arithmetic, geometric, and harmonic means.) The line shows the upper bound of the performance ranges.

Key Words: Livermore Fortran Kernels (LFK) — computer benchmark test; processor — conventional, multiple, parallel, vector; supercomputer — benchmarks, performance evaluation.

Notes and References

1. V. J. Kransky, E. D. Giroux, and G. A. Long, *Parallel Implementation of a Two-Dimensional Model*, Lawrence Livermore National Laboratory, Rept. UCRL-74894 (1973).
2. G. H. Rodrigue, E. D. Giroux, and M. W. Pratt, "Perspectives on Large Scale Scientific Computation," *IEEE Computer* 13 (10), 65-80 (1980).
3. W. R. Martin, P. F. Novak, and J. A. Rathkopf, "Monte Carlo Photon Transport on a Vector Supercomputer," *J. Res. Develop.* 30 (2), 193-202 (1986).
4. T. W. Keller, *Cray-1 Evaluation Final Report*, Los Alamos National Laboratory, Rept. LA-6456-MS (1976), p. 31.
5. S. F. Mendicino et al., "The LRL-TRAN Compiler," *Comm. ACM* 11, 747-759 (1968).
6. H. J. Currow and B. A. Wichman, "A Synthetic Benchmark," *The Computer J.* 19 (1), 43-49 (1976).
7. F. H. McMahon, *The Livermore Fortran Kernels. A Computer Test of the Numerical Performance Range*, Lawrence Livermore National Laboratory, Rept. UCRL-53745 (1986).
8. D. J. Kuck, Y. Muraoka, and S. C. Chen, "On the Number of Operations Simultaneously Executable in Fortran-Like Programs and Their Resulting Speedup," *IEEE Trans. Comput.* C-21 (12), 1293-1310 (1972).
9. C. S. Arnold, "Performance Evaluation of Three Automatic Vectorizer Packages," in *Proc. Int'l. Conf. on Parallel Processing* (IEEE Computer Society Press, Silver Spring, MD, 1982).
10. J. J. Hack, "Peak vs Sustained Performance in Highly Concurrent Vector Machines," *IEEE Computer* 19 (9), 11-19 (1986).
11. R. Schow, "CDC 205 Runs 800 Million Operations per Second," *Computers and Automation* 14 (19), 1 (1980).
12. H. J. Nelson, "Using the Performance Monitors on the Cray XMP 48," *Technical Report* 5 (9), 15-23 (1985).
13. P. J. Fleming and J. J. Wallace, "How Not to Lie with Statistics: The Correct Way to Summarize Benchmark Results," *Comm. ACM* 29 (3), 218-221 (1986).
14. R. H. Ewald, "An Overview of Computing at Los Alamos," in *Supercomputers in Theoretical and Experimental Science*, J. J. DeVries and P. Van Camp, Eds. (Plenum Press, New York, NY, 1985).

Glossary

Application program. A computer program used to solve specific, user-defined problems.

Basic code block. A block of procedural statements beginning with a statement label, called a branch target, that contains no other branch targets. Important optimizations of the machine operations within a basic code block are not constrained by the behavior of the rest of the program.

Central processing unit (CPU). The section of the computer hardware that performs computational operations; the CPU contains most of the logic circuits and operating registers in the computer.

Code. A computer program that contains the operating instructions for the computer.

Compiler. A program that translates high-level procedural language programs into the machine-language instructions that can then be executed by the computer.

DO loop. A block of Fortran statements that can be executed repeatedly by invoking the verb DO.

Floating-point operations. Operations on scaled-integer numbers that approximate real numbers and are represented by an integer part and an exponent part, as in scientific notation (e.g., 2×10^{-3}).

Fortran. The principal computer language used to program numerical applications for more than 30 years.

Kernel. The portions of a computer program that consume much of the execution time; usually a nest of loops.

Microtasking. The ability of a system to efficiently multiprocess very small tasks.

continued on page 14

continued from page 13

Multiprogramming. The allocation of CPU and input/output (I/O) resources among several independent programs; the CPU and I/O devices run in parallel to maximize the utilization of these resources and throughput.

Multiprocessing. The use of several CPUs to execute, in parallel, several independent tasks in a single program or several independent programs.

Parallel processing. Simultaneous execution of two or more independent tasks to minimize the total processing time.

Pipelined CPU. See vector computer.

SISD/SIMD computer. A single-instruction, single-data processor (scalar) and a single-instruction, multiple-data processor (vector).

Supercomputer. A computer system with the maximum computational speed, memory size, and cost that is feasible in each epoch. Primarily used to simulate the evolution of complex physical systems by numerical modeling. (The "1985 epoch" supercomputers were called Class-VI machines, and typically had eight million words of memory, ran application programs at 10-80 Mflops, with burst speeds up to 200 Mflops, and cost roughly \$20 million. The next-generation, 1990 or Class-VII, supercomputers are expected to have 64 times more memory and eight times the speed.)

Upward compatibility. A new computer is "upwardly compatible" with respect to an older model if all old programs can be installed on the new computer without reprogramming. In strict terms, this includes machine-language programs.

Vector computer. A computer with hardware instructions that can each operate on a set of data elements, usually a one-dimensional subarray (a row or a column of an array). High speed is realized by streaming a set of elements through the hardware segments, all of which execute in parallel. This CPU design is termed a vector "pipeline."

Workload. All of the users' demands on a computer system; usually represented by the customers' current application programs.

Microwave Tokamak Experiment

The Microwave Tokamak Experiment, now under construction at the Laboratory, will use microwave heating from a free-electron laser. The intense microwave pulses will be injected into the tokamak to realize several goals, including a demonstration of the effects of localized heat deposition within magnetically confined plasma, a better understanding of energy confinement in tokamaks, and use of the new free-electron laser technology for plasma heating.

For further information contact
E. Rick Hooper (415) 422-6792,
William M. Nevins (415) 422-7032,
or Ronald H. Cohen (415) 422-9831.

A problem of increasing concern to the U.S. and the world is finding new sources of energy as fossil fuels become depleted. At LLNL, we are studying the prospects for generating energy from the controlled fusion of heavy isotopes of hydrogen (deuterium and tritium). The ultimate fuel for fusion — deuterium separated from the hydrogen in water — represents a fuel reserve that would last for billions of years.

Today, fusion energy research based on confinement by magnetic fields is carried out predominantly with the tokamak. This device contains a toroidal (doughnut-shaped) vacuum vessel; its toroidal magnetic field and current combine to hold the fusion fuel for long confinement times. Researchers around the world are actively exploring ways to heat the fuel to ignition temperatures and to drive the toroidal current continuously. At present, ohmic heating currents produce limited fuel temperatures, and transformer-induced ohmic current is limited to pulses of many seconds.

Our research at LLNL is centered on new methods for heating fuel and on novel schemes for driving current. We will pursue both of these tasks with microwave power generated by a free-electron laser (FEL) and a new high-repetition-rate accelerator. This research will be conducted in a new facility, the Microwave Tokamak Experiment (MTX), that combines the Alcator-C tokamak and free-electron lasers.

The Experimental Test Accelerator II (ETA II) was recently built at LLNL and is in its early stages of operation. The ETA II can produce short pulses of high-current electron beams at a peak power level of 15–25 GW. In an FEL, the electron beam is passed through a wiggler with a microwave source, and the resulting microwave signal is amplified to 6–10 GW at millimeter wavelengths. In essence, the FEL is a broadband, tunable amplifier with demonstrated high gain and efficiency. The FEL is undergoing rapid development at LLNL.¹

The Alcator-C tokamak was recently brought from the Massachusetts Institute of Technology (MIT) and installed at LLNL (Figure 1). We plan to carry out new experiments on the high-density plasma in the tokamak using pulsed, high-average-power, high-frequency microwaves. The FEL/Alcator-C tokamak combination is unmatched in terms of its ability to provide plasma fueling by frozen

deuterium pellets, localized heat deposition in the plasma, and magnetic fields and densities of interest for future reactor-grade experiments.

In this article, we first describe plasma and microwave heating. Next, we discuss the use of FELs for magnetic fusion. We then discuss the Microwave Tokamak Experiment (MTX), scheduled to begin in the fall of 1988.

Plasma Heating

In tokamak confinement devices, an electric current must be maintained within the plasma to keep it confined. Plasma is similar to a wire in that it has a finite resistance to the flow of current. The plasma current can be maintained by providing an electric field parallel to the plasma current. The electric field is created by treating the ring of confined plasma as the secondary circuit in a transformer; so long as current in the primary circuit is changing, an electric field can be maintained in the plasma.

Power is dissipated in maintaining a plasma current. Thus, the plasma current not only helps to confine the plasma but it also heats the plasma to temperatures of 1–2 keV. However, plasma resistance falls rapidly with increasing plasma temperature, so

Figure 1. Facility for the Microwave Tokamak Experiment (MTX). The Alcator-C tokamak, shown in detail at the top, will be used to carry out plasma-heating experiments. In the diagram, at left, is the Experimental Test Accelerator (ETA II) containing the electron-beam accelerator and FEL wiggler. The tokamak is located within the main building and has its own shielding. Power from the FEL will be transmitted about 19 m to the tokamak through a quasi-optical, microwave-transmission system that consists of an evacuated 0.6-m-diameter ducting network enclosing four focusing mirrors. The microwave beam will then be focused to enter the tokamak through a 4-cm-wide diameter port.

that it becomes difficult to heat plasma further using resistive effects alone.

Although a temperature of 1 keV may seem quite high, it is not sufficient to maintain fusion reactions. For fusion, plasma temperatures in the range of 10 to 20 keV are needed. Hence, auxiliary heating is needed in addition to the ohmic heating provided by the plasma current. We can think of the auxiliary heating as the match or pilot light used to light a stove; it is required to get the fire (or fusion burn in our case) started. Once the fusion burn begins, energy from the fusion reactions themselves will be sufficient to maintain the plasma temperature above 10 keV.

Many schemes are available for heating fusion plasmas. One of the most successful techniques has been the injection of energetic neutral-particle beams. LLNL has been involved in the research and development of such beams for many years. More recently, researchers have become interested in using another technique to heat ions within plasma. Because this scheme involves radio waves with frequencies on the order of the ion-cyclotron frequency (~ 100 MHz), we refer to this type of heating as ion-cyclotron heating. However, as promising as these two plasma-heating methods are, each one has drawbacks.

The energy deposited by energetic neutral beams tends to spread out over the plasma cross section rather than to remain concentrated near the center, where it would be most effective in heating. In addition, fusion reactors would have to be larger than those available in current experiments, and they would have to operate at a higher density. Substantial increases in the energy

of neutral-particle beams would be needed to enable them to penetrate the large, denser plasmas.

Ion-cyclotron heating, on the other hand, is able to deposit energy at the center of a plasma. Recent experiments at the Joint European Torus (JET) have demonstrated such deposition. Unfortunately, the low-frequency waves needed for ion-cyclotron heating are strongly affected by the plasma and have a different polarization than waves outside the plasma. An antenna or some other launching structure must therefore be placed near the plasma surface if these waves are to be coupled to the plasma. Such coupling remains somewhat problematic and depends on the properties of the edge plasma. In particular, ion-cyclotron heating does not seem to be compatible with the recently discovered, enhanced-confinement regime known as the H mode. As a result, our interest continues in developing new and more effective methods for heating tokamak plasmas.

One alternative is to use waves with frequencies near the electron-cyclotron resonance frequency to heat fusion plasmas. The use of these high-frequency (~1-mm-wavelength) waves has many advantages, including the fact that their propagation inside the plasma is similar to that outside. As a result, electron-cyclotron waves can be launched at some distance from a plasma; in fact, they can be bounced off mirrors much like a beam of visible light. The waves then propagate through the vacuum surrounding a plasma and into the plasma itself. Electron-cyclotron frequency waves are expected to propagate into plasma over a wide range of plasma edge conditions. The location of energy deposition within

the plasma can be controlled by the relation between the frequency of the wave and the magnetic field within the plasma. As a consequence, many of the difficulties we experience with other heating methods are overcome through the use of electron-cyclotron heating.

Successful experiments incorporating electron-cyclotron heating have been carried out on many devices, including the T-10 tokamak at the Kurchatov Institute in Moscow and the D-III tokamak at General Atomics in San Diego. The present work on D-III includes a collaboration with several Laboratory scientists who are participating in the effort. Experiments on D-III have demonstrated that electron-cyclotron heating does indeed deposit energy at the plasma center and that this technique of plasma heating is not sensitive to plasma edge conditions. Indeed, the experiments have demonstrated that electron-cyclotron heating is more effective than either neutral beams or ion-cyclotron heating in obtaining enhanced energy confinement.

In the past, many researchers were concerned that efficient, high-power radiofrequency (rf) sources at the high frequencies (200–500 GHz) needed for electron-cyclotron heating in tokamak reactors would not be available. The recent development at LLNL of efficient free-electron lasers in this frequency range has renewed interest in electron-cyclotron heating of fusion plasmas. Today, electron-cyclotron heating is a leading candidate for igniting the plasma in both the Compact Ignition Tokamak (CIT) that will be built at Princeton and for the proposed International Thermonuclear Engineering Reactor (ITER) that is being designed through an international

collaboration involving the Soviet Union, the European community, Japan, and the United States. LLNL is the lead laboratory for the U.S. in the ITER collaboration.

Recent advances in the efficient generation of high-frequency rf power have also motivated the MTX at the Laboratory. We hope to deliver up to 2 MW of average power from the FEL to the MTX. Such power should lead to substantial increases in plasma temperature and will demonstrate the use of intense pulses of microwave power to heat tokamak plasmas.

Free-Electron Lasers

Before discussing our plans for the MTX in more detail, it is important to understand the basic principles of the free-electron laser (FEL). An FEL directly converts the energy of a beam of electrons into electromagnetic radiation. The radiation is emitted as the electron beam bunches into many thin structures resembling pancakes or disks that radiate coherently. The bunching occurs when the electrons wiggle in a periodic array of magnets (the wiggler) and interact with the electromagnetic radiation. Dramatic increases in the efficiency of conversion between electron-beam power and radiation are possible when the wiggler magnetic-field strength or period decreases with distance along the wiggler (i.e., in a tapered wiggler).

It is also important to remember that the radiation from an FEL can be at any wavelength, from microwave to visible light and beyond. The radiation we are concerned with in the MTX is in the microwave range. The frequency at which an FEL operates is determined by the wiggler period, the electron energy, and the wiggler magnetic field.

The electrons in an FEL electron beam have a velocity very close to the speed of light and thus are called relativistic. As shown in Figure 2, these relativistic electrons are accelerated by the wiggler in the sense of changing direction rather than gaining energy. If nothing else were to occur in an FEL, the device would be a standard synchrotron radiation source, such as the one at the Stanford Synchrotron Radiation Laboratory. Under these

circumstances, the radiation from individual electrons in the beam would add incoherently.

In an FEL, however, the electrons physically bunch on the scale of one wavelength of the radiation, as viewed by an observer moving with the beam, to form a series of thin disks. The single-particle radiation from each electron in a disk has the same phase; the radiation amplitudes, rather than the powers, add together. Furthermore, because the disks are

separated by almost one wavelength, the radiation amplitudes from different disks also add in phase. The radiation power, which is proportional to the square of the radiation amplitude, then depends on the square of the number of electrons radiating in phase rather than on only the number of electrons.

At the single frequency designated for an FEL, which is 250 GHz in our MTX application, the coherence of the radiation (the linear adding of amplitudes) produces an enormous enhancement of radiation power. Thus, the FEL becomes a powerful tool: a highly efficient amplifier of microwaves.

Advantages of the FEL for the MTX

Previous experiments with high-power electron-cyclotron resonance heating have used gyrotrons. In a gyrotron, the power in a high-energy electron beam within a strong magnetic field is used to generate microwave power. The process, first demonstrated in a practical device in the Soviet Union, relies on the energy dependence of the cyclotron frequency of the beam electrons. The capability of these devices is impressive; indeed, Varian Associates is now working on a 1-MW, continuous-wave gyrotron operating at 140 GHz. Nevertheless, the FEL incorporates features that are particularly advantageous for fusion research. The advantages of the FEL for our MTX project include:

- The capability for higher frequencies. FELs are now being operated at infrared and visible wavelengths.
- Higher power capability per unit size. FELs with power up to 10 MW

(a) Side view

(b) Top view

Figure 2. Diagram of a wiggler in which magnets are arranged linearly every 5 cm. (a) Side view of a linear wiggler in which magnets provide a magnetic field that periodically reverses direction. On the midplane, the magnetic field lines point either up or down. As the electron beam moves across the magnetic field, each electron wiggles in response to a force at right angles to both the magnetic field and the direction of motion. (b) Top view of the wiggler, with the magnetic field directed into or out of the illustration, as denoted by the tails and tips of arrows, respectively. Only the top set of magnets is shown. The horizontal wiggles of electrons is represented by the undulating blue line.

are now being considered and would be a natural size for fusion devices of the future.

- **Excellent tunability.** The tuning curve of an FEL has a width of 5% (to the 80% gain points). This means that we can quickly change the frequency to control plasma heating. On a slow time scale (i.e., over seconds), the frequency can be changed by an almost arbitrarily large amount by changing the beam energy and wiggler magnetic field.

The FEL also provides power in short, extremely intense pulses. Because the pulse length is shorter than the electron exponentiation times (the buildup time for electrons from ionization or surface emission), transmission lines can carry very high average power to the fusion device with less concern for breakdown.

The absorption physics is strongly nonlinear and thus has some unique features (see the box on nonlinear absorption of electron cyclotron waves, p. 20). Therefore, we will need to conduct experiments to demonstrate that the power will be well absorbed and that the plasma energy losses will remain low in the presence of a train of intense pulses.

During our initial experiments, we will use the MTX free-electron laser to inject a single pulse of microwaves into the tokamak. The single pulse will allow us to evaluate the absorption of 140-GHz power at 2–4 GW for durations of 20–35 ns. Our initial tests will be followed by experiments in which we will inject a burst of up to 50 pulses at a repetition rate of 5 kHz. Such bursts will permit a more accurate measurement of absorption, will yield a first evaluation of heating efficiency, and will allow us to determine the

confinement of energy on a time scale comparable to the energy-confinement time. This time is 5–10 ms for an ohmic-heated plasma in the 5-T magnetic field needed for heating at 140 GHz. Later, we will undertake experiments at 250 GHz, culminating in up to 0.5 s of heating at an average power of 2 MW.

The Microwave Tokamak Experiment

Figure 1 shows our experimental facility, including the Alcator-C tokamak, the FEL, and the quasi-optical transmission line that will be used to bring the microwave beam to an entrance port of the vacuum vessel. Figure 3 shows two photographs of the actual tokamak during installation at LLNL. The vacuum vessel and magnetic field coils are placed inside a cryostat and are cooled by a spray of liquid nitrogen. Reducing the temperature with liquid nitrogen lowers the resistance of the magnetic field coils; however, 140 MW of power is still needed to obtain the 9-T field required for the absorption of microwaves at the electron-cyclotron frequency.

Our short-term goals for the MTX are to:

- Inject microwaves into the tokamak at a peak-power density greater than 50 MW/cm² using a windowless, quasi-optical system (including mirrors) to transmit power.
- Measure the absorption of intense microwave pulses and compare the results with predictions from nonlinear theory.
- Evaluate backscattering from parametric instabilities, which may be generated in the plasma. If necessary, we will determine mechanisms to keep the backscattering to low levels.

Figure 3. Photographs of the Alcator-C tokamak during installation at LLNL. (a) This view shows several of the 18 access ports, the poloidal field coils, and the toroidal field coils. (b) The white cryostat contains the vacuum chamber and magnets. Current for the toroidal field coils is supplied by the bus above the machine; other coils are fed from the side. The microwave transmission line is not yet installed.

Nonlinear Absorption of Electron Cyclotron Waves

Electrons move along field lines in spiral paths, which can be viewed as the vector sum of streaming along magnetic field lines plus rotation or gyration around field lines. Electron cyclotron waves are microwaves that interact with electrons by resonating with their gyro motion. In its simplest form, this means that the electric field of the cyclotron wave rotates with the electrons as they gyrate. When this condition is achieved, the gyro motion is steadily accelerated or decelerated by the waves. Whether the electron is accelerated or decelerated depends on the phase of the wave relative to that of the electron, i.e., the orientation of the rotating electric field relative to the electron. Most commonly, this acceleration continues until some event not intrinsic to the basic wave-particle interaction intervenes — for example, the electron exits the microwave beam. Typically, the energy gained by an electron during a single pass through the beam is small, but substantial energy can be transferred from the wave to the electrons after many passes.

When the wave electric field becomes very strong, the acceleration resulting from the wave can be self-

limiting. One mechanism for this results from the fact that as an electron gains energy, its frequency of gyration becomes slightly smaller (resulting from its relativistic increase in mass). Hence, a previously resonant electron can no longer quite keep up with the rotating electric field. After a while, the relative phase of the particle and the wave shifts enough that the electron begins to decelerate.

Ultimately, the entire process repeats. The electron, which is said to be nonlinearly trapped about the cyclotron resonance, goes through alternating accelerations and decelerations as long as it is in the microwave beam. These swings in energy during oscillations can be large compared to the initial energy of a typical electron in the plasma. The net result of electrons passing through a beam of resonant cyclotron waves is that about half of the electrons emerge with their energies increased by an amount comparable to the energy swings. The resulting transfer of energy from the microwave beam is less than would be the case in the absence of the self-limiting process, but it is still adequate to completely absorb the beam in a plasma such as that in the MTX.

• Determine the energy confinement of the heated plasma. Our results will be compared with high-power heating done with gyrotrons elsewhere in the world, notably on the T-10 and D-III tokamaks.

We also plan to study the electron-cyclotron heating of peaked-density-profile plasmas, which can be generated by injecting small pellets of frozen deuterium into the tokamak. When such plasmas were created in Alcator-C at MIT, they showed enhanced confinement; however, auxiliary heating was not available. In the Tokamak Fusion Test Reactor (TFTR) at Princeton, neutral-beam heating of pellet-injected plasmas has led to very hot plasmas. It is possible that when such "super shots" are combined with tritium plasmas, they may lead to the first demonstration of enough

deuterium-tritium neutron production in a magnetically confined plasma to produce more power than is absorbed. Our MTX tests will include localized heating of such profiles to evaluate their possible application to igniting a reactor.

Our long-range objectives are to use this unique FEL heating system to advance our understanding of tokamak physics and to develop ideas for improved tokamaks of the future. Especially important is the ability to deposit large quantities of heat in the plasma on a fast time scale. Specific long-range objectives include:

- Energy-confinement scaling.
- Studies of the transport mechanisms in plasma.
- Control of magnetohydrodynamic (MHD) instabilities, which perturb the confining magnetic field in a manner that decreases confinement.

• Shaping the plasma profile to optimize confinement.

• Studies of current drive (see the box on radiofrequency current drive, p. 21).

Our goal in current-drive research is to learn how to make plasmas in a steady-state tokamak rather than in the pulsed, transformer-driven tokamaks used today. Current-drive experiments would use the robust, nonlinear effects associated with the intense interaction with microwaves to generate a net flow of electrons along the magnetic field.

Operation of the tokamak at LLNL will begin in the fall of 1988, when ohmic-driven plasmas will be formed. Our first experiments using the FEL are planned for the winter of 1988, when we will inject single pulses of microwaves. In 1989, we plan to use multiple pulses (up to 50

Radiofrequency Current Drive

A tokamak requires current flowing around the torus to maintain the magnetic fields that confine the plasma. The most common technique for generating such a current is ohmically. The plasma is, in effect, made the secondary winding of a transformer, and an electric field is induced in it by varying the voltage across the primary winding. The induced electric field then drives the current. This approach is effective but precludes steady-state operation of the tokamak, since the primary voltage cannot be increased indefinitely. In recent years, it has been demonstrated theoretically and experimentally that it is possible to drive currents with waves; such currents could, in principle, be maintained continuously.

There are two basic mechanisms by which waves can drive currents. First, waves with electric-field components along the (mainly toroidal) magnetic field can directly accelerate electrons in the direction of the magnetic field, producing a current. The second mechanism is called anisotropic resistivity, and occurs even if the wave acts only to increase the energy of

gyration of the electrons (as is the case for cyclotron waves, see the box on p. 20). This second mechanism results from the fact that energetic electrons collide with other electrons and ions less frequently than do slower-moving electrons. Suppose that we took a group of electrons traveling in one direction along a magnetic field line, say to the right. Suddenly, without changing their motion along the field line, we completely eliminated their ability to collide (this could be achieved, for example, by giving them a very large gyration energy). Initially there would be no current. But, after some time, the remaining electrons would collide with the much heavier ions and re-randomize themselves, with about as many electrons going to the right as to the left. But now, our selected, collision-free group of electrons is still traveling to the right, producing a net current. This mechanism is actually more efficient than the first, with the consequence that cyclotron waves, which drive current mainly by anisotropic resistivity, are almost as efficient as waves with parallel electric field.

in a 10-ms burst) to increase the plasma stored energy, to permit energy accounting, and to begin studying the confinement of heat. In the longer term, our goal is to inject 2 MW of average power at a wavelength of 1 mm for 0.5 s.

MTX Diagnostics

We require a wide variety of diagnostic devices to measure the basic properties of the plasma and to evaluate heating. We are constructing the basic set of instruments at the MTX is being assembled and commissioned. In general, the diagnostics will provide measurements of plasma-density and temperature profiles, plasma impurities, radiated power, and other primary parameters. Table 1 is a list of the basic diagnostics and the quantities they will measure. Many of the diagnostics are major

Table 1. Basic plasma diagnostics for the MTX and the quantities measured.

Diagnostic	Quantity measured
Far-infrared interferometer	Profile of electron density
Electron-cyclotron emission	Profile of electron temperature
Extreme ultraviolet	
SPREAD spectrometer	Impurity spectra
Magnet loops	Plasma field and position
Soft x ray	X ray emission
Limiter x ray	X rays from run-away electrons
Visible bremsstrahlung	Average impurity level
Bolometers	Radiated power
Plasma camera	Plasma position and H-alpha emission ^a
FIG/RGA ^b	Gas pressure and impurities
Thomson scattering	Electron temperature and density

^aH-alpha emission is light emitted by neutral hydrogen or deuterium atoms, in this case due to excitation by plasma electrons. Its intensity is closely related to the source of plasma particles.

^bFast ion gauge/residual gas analyzer.

Table 2. Measurement capabilities for single-pulse MTX experiments.^a

Issue	Instrument	Quantity measured
Opacity	Calorimeter	Power transmission through plasma
	Microwave detectors	Time-resolved transmission
	Visible bremsstrahlung	Heating location
	Fast magnetic loop	Total energy absorbed
Backscatter	Microwave detectors	Backscatter near 180 deg
	Heterodyne receiver	Frequency shift of backscatter
Absorption physics	Fast, soft x ray	Hot electron energy
	Thomson scattering	Hot electron distribution
	Fast electron-cyclotron emission (single channel)	Hot electron energy

^aDiagnostics for evaluating backscattering and absorption require both basic plasma diagnostic devices and new, fast diagnostic instruments.

undertakings in themselves. One example is the far-infrared (FIR) interferometer, shown in Figure 4. Debugging of these instruments will take place during initial tokamak operation, which is scheduled for the fall of 1988.

We will use additional diagnostics for measuring the absorption of microwave power, monitoring any backscattering, and determining the effects on electrons. Table 2 and Figure 5 show some of our current plans for such measurements.

Calorimetric measurement of the energy transmitted through the plasma is central to our initial experiments. Reducing the energy

Figure 4. Far-infrared (FIR) interferometer. This diagnostic device is being designed and constructed in collaboration with researchers at UCLA to measure the plasma-density profile on the MTX. One infrared beam with a wavelength of 185 m is passed through the plasma, while the second beam with a frequency difference of 1 MHz is used as a local oscillator. Schottky detectors measure the phase shift of the transmitted radiation along several chords. These data can be analyzed to yield the radial density profile. A planned upgrade includes a second set of detectors to determine the Faraday rotation of the radiation. Analysis of this rotation yields the discharge-current profile in the plasma. The massive stand is designed to maintain the optical path length to within 1/50th of a wavelength during the 0.5-s plasma discharge so that noise arising from vibrations can be minimized.

transmitted will provide a first demonstration of absorption, and the results can be compared with those from theoretical calculations. However, this measurement will not be conclusive until we can account for the total injected energy. Such an accounting will probably require the injection of enough microwave pulses to raise the average plasma temperature significantly. This is the goal of the second series of experiments, in which the length of the burst of pulses (10 ms)

will be comparable to the energy confinement time.

We are interacting with many other laboratories around the world to obtain the difficult measurements we require, to evaluate the absorption physics, and to understand the plasma response to heating. Table 3 lists our interactions on MTX diagnostics, which include collaborative efforts with researchers at General Atomics, the University of Maryland, the University of California at Los Angeles, and the

Japanese Atomic Energy Research Institute.

Summary

The Microwave Tokamak Experiment (MTX) is a prime example of applying state-of-the-art technology to advance the quest for future energy supplies. In our work, the free-electron laser is combined with the magnetic confinement approach to fusion. This experiment, which will soon be operational, will

Figure 5. MTX microwave diagnostics. We will monitor power transmitted through the plasma by a calorimeter, shown here at the far right, and microwave receivers. A wire grid in the beamline will reflect a small fraction of power to determine the incident and reflected power. Small holes in the transmission mirror will also monitor incident and reflected signals.

Table 3. Collaborative efforts on MTX diagnostics.

Instrument	Collaborator	Measurement
Electron-cyclotron emission (basic)	General Atomics	Electron temperature profile
Magnetics (fast)	General Atomics	Plasma energy change during electron-cyclotron resonance heating
Electron-cyclotron emission (fast)	University of Maryland	Fast-electron heating during FEL electron-cyclotron resonance heating
Far-infrared interferometer and polarimeter	University of California, Los Angeles	Density profile; current-density profile
Soft x ray (fast)	Japanese Atomic Energy Research Institute	Fast-electron heating by FEL electron-cyclotron resonance heating; ion temperature (fast response)
Neutrons	Japanese Atomic Energy Research Institute	
Microwaves	Japanese Atomic Energy Research Institute and Toshiba	Total average power; spatial power distribution; polarization; incident polarization; and reflected power

provide us with an opportunity to study dense plasmas heated by powers unprecedented in the electron-cyclotron frequency range required by the especially high magnetic fields used with the MTX and needed for reactors. The possibilities of addressing new physics, of studying energy-confinement mechanisms in the tokamak, and of advancing the state of fusion research make this an exciting challenge for physicists and engineers both at LLNL and at the laboratories around the world that are collaborating on the MTX.

Key Words: Alcator-C tokamak; free-electron laser (FEL); magnetic fusion energy; Microwave Tokamak Experiment (MTX).

Notes and References

1. See the December 1986 issue of *Energy and Technology Review* (UCRL-52000-86-12), pp. 8-13 and 33-41, for more details on the new free-electron laser technology and its use in the Microwave Tokamak Experiment.

Gas Cooling of Laser Disks

Solid-state disk amplifiers can produce near-diffraction-limited, repetitively pulsed output beams if they are cooled on their surfaces by a heat-transfer medium of excellent optical quality. Our experiments and models show that effective cooling can be provided by flowing gas in adjacent channels across laser disk surfaces.

For further information contact
Georg F. Albrecht (415) 423-7562
or Steven B. Sutton (415) 422-0322.

Since the earliest days of lasers, the scaling of solid-state lasers to significant average powers (i.e., relatively high pulse-repetition rates) while maintaining high beam quality has been considered a difficult problem. It was thought that waste heat, generated in the optical pumping process, could not be removed in such a manner as to reduce thermally induced optical distortions to an acceptable level. However, recent analysis^{1,2} showed that this problem could be overcome by using a geometry in which thin disks are optically pumped and heat is removed from their large faces and across the thinnest dimension. This amplifier architecture is very similar to the Nova laser disk amplifiers developed for single-shot fusion experiments,³ but with a cooling subsystem added.

The thrust of our present work is to assess quantitatively the performance potential of the flow-cooling disk concept by taking experimental measurements of optical beam quality over a wide range of cooling parameters and by developing and validating quantitative design codes. Our work to date supports the conclusion that we can design adequately cooled

solid-state laser amplifiers approaching a megawatt of high-quality output power per square meter of aperture.

The active medium in a solid-state laser disk amplifier is a thin, rectangular plate or slab with side lengths in a ratio of about 2:1. This plate is oriented with respect to the output beam direction at an angle that minimizes reflection losses (i.e., Brewster's angle). To effectively cool such a plate, heat must be removed from the largest surface across the thinnest dimension. Since the laser beam traverses the slab through its large faces, the cooling

medium that draws heat from the slab must have excellent optical quality and must be unaffected by the heat-removal process. These considerations lead to the choice of a gas as the coolant.

Accordingly, we have investigated the flow of a gas over a hot surface in a confined flow channel (Figure 1). We must understand not only the bulk flow properties and their effect on eventual optical distortions of the slab and the separation windows but also the optical properties of the gas flow itself. We have therefore undertaken an experimental and modeling effort to examine in detail

Figure 1. A cross section of the geometry of an amplifier disk, the cooling gas, and the laser beam in a gas-cooled disk amplifier.

the fluid dynamics, heat transfer, and optical properties of a section of channel flow.

Cooling Channel Geometry

Typical channel geometry is shown in Figure 2. Cooling effectiveness in the channels can be characterized by parameters that affect laser system performance (e.g., slab temperature and temperature gradient along the slab). Each of these parameters is dependent on a number of flow and fluid parameters, such as Mach number, pressure, and the flow geometry. For example, Figure 3 depicts the effect of flow Mach number and pressure on the temperature drop in the thermal boundary layer. At first glance, it appears that two ways to reduce the slab temperature are to increase the Mach number or increase the gas pressure. However, frictional drag (which consumes compressor power) increases with pressure and with the cube of the Mach number; increased pressure also increases the distortion of optical components in the system. To maintain a high overall system

efficiency, we wish to use the minimum amount of compressor power needed to achieve the desired degree of cooling that is consistent with acceptable optical distortions. Thus, proper cooling of the laser slab requires the correct choice of flow geometry, gas type, and flow conditions to extract the heat at tolerable levels of compressor power and optical distortion.

Flow Design

Flow design begins with the nozzle. A properly designed nozzle efficiently exchanges internal energy for kinetic energy (pressure for flow) while introducing little disturbance into the flow stream. The nozzle takes low-velocity gas provided by the driving compressor and accelerates it with a significant reduction in the static pressure and temperature of the gas but with negligible reduction in the stagnation pressure and temperature. Flow exiting the nozzle has a high velocity, which is required for efficient cooling, and a nearly uniform velocity profile.

As the flow travels down the narrow channel, the presence of the walls causes the formation of viscous boundary layers. A boundary layer is a zone of accommodation between regions of different flow characteristics. In a viscous boundary layer, the transition is between a high-velocity core flow and the zero-velocity flow at the channel walls. At the channel entrance, where the viscous influence zone is thinnest, frictional forces are at their largest. As the flow develops (see Figure 4), the viscous boundary layer grows and the velocity gradient and shear stresses at the wall decrease. The viscous development continues until the wall effects and the velocity profile have stabilized. The retardation of flow is accompanied by a decrease in both the static and stagnation pressures.

Gas Temperature, Velocity, and Density

Now let us consider gas temperature. In the configuration shown in Figure 4, the walls upstream of the slab are adiabatic (no heat transfer). In this region, shear-generated energy causes a slight warming of the gas wherever there

Figure 2. A schematic representation of a cooling channel, showing typical channel geometry. An inlet nozzle takes high-pressure, low-velocity gas from the system compressor and injects it into the flow channel. The cooling gas is conditioned both hydrodynamically and thermally in the development portion of the flow channel. The intermediate region of the channel encompasses the laser slab, where heat is removed by the fluid. Beyond the laser slab (downstream) is a relaxation region, followed by a diffuser element where kinetic energy is recovered from the flow.

Figure 3. The effect of Mach number and pressure on the convective heat transfer temperature drop in a 4.5-mm channel with a surface heat flux of 2 W/cm^2 .

exists a strong velocity gradient. In the slab region, in order to remove heat, the surface temperature must be greater than the gas temperature. As with the velocity, this temperature difference results in the formation of a thermal boundary layer with steep temperature gradients in the gas near the surface. The temperature gradients allow energy to "diffuse," again primarily through turbulent action, through the boundary layer and out into the core flow, thus increasing the thickness of the thermal boundary layer. Like the viscous boundary layer, this thermal boundary layer grows across the channel until something resembling a fully developed profile is achieved. Note that the viscous profile is symmetric whereas the thermal profile is asymmetric, because heat is being removed from only one side of the channel.

Since we use a gas for cooling, changes in temperature and pressure strongly affect the gas density. Figure 4 includes the density profile through the channel, which is seen to vary inversely with the temperature. Any variation down the channel of either the cross-stream density gradients or the bulk density can introduce aberrations into the extraction beam.

The high-velocity gas decelerates in the diffuser before it is returned to the compressor, converting some kinetic energy to internal energy. This causes an increase in both the static temperature and pressure. Diffusers are not as efficient as nozzles. It is difficult to transfer more than about 70% of the available kinetic energy back into internal energy.

The frictional losses in the cooling channel and the recovery losses in the diffuser contribute to the pressure difference necessary to

Figure 4. Diagram of flow development in the cooling channel.

Figure 5. Calculated Mach numbers and a comparison of calculated (curves) and measured (data points) static pressures in a straight and a tapered (0.05-deg) channel.

drive the flow through the cooling loop. Since the compressor power is directly related to the overall stagnation pressure drop, we must design the least lossy, most thermally efficient cooling arrangement possible. The compressor power and thermal conditioning power represent sacrifices in the single-shot efficiency that must be made in order to operate the laser at average power.

To optimize the cooling system, we must understand gas flow and heat transfer from a system viewpoint. We must also understand the effect of heat transfer and flow behavior on optical distortions in the laser materials. To do this, we used three types of models:

- (1) Analytical solutions to the one-dimensional flow equations for the special flow conditions of constant area with friction (Fanno flow) and variable area with constant Mach number, friction, and heat addition (a tapered channel).
- (2) Numerical integration of the general one-dimensional flow equations that include the effects of area change, friction, and heating, using experimental correlations for boundary effects.
- (3) Numerical integration of the two-dimensional Navier-Stokes equations using high-order turbulence modeling. For this, we employed the TEXSTAN computer program⁴ that integrates the steady, compressible, boundary-layer form of the Navier-Stokes equations for mass, momentum, and energy.

Experiments

We performed many extensively instrumented experiments to collect information on the bulk character of the channel flow and to gather data for validating the various numerical approaches. In all cases, we used nitrogen as the cooling fluid.

Pressure

We measured the static pressures in the channel, since they are associated with flow-separation window pressure distortion. We also concentrated on the surface temperature and the gas temperature, since they influence thermal distortion of the slab and flow separation window, which in turn creates optical distortions.

We measured static pressures upstream, downstream, and along the channel with 16 static pressure taps and a pressure gauge. Figure 5 compares the measured and predicted static pressure for a constant-area channel and a tapered channel (0.05-deg taper). The associated flow Mach numbers are also shown. Good agreement between measured and calculated static pressures is achieved through the nozzle and the channel. However, the models are not sophisticated enough to model the diffuser flow. Figure 5 clearly shows both the effect of the nozzle in increasing the velocity and reducing the static pressure, and the effect of the diffuser in recovering some of the nozzle pressure drop. Figure 5 also shows the effect of a very slight channel taper on reducing the static pressure drop at moderate Mach numbers, thus offering a geometry variation that can be used to reduce pressure imbalance distortions in the optical components.

Temperature

There are two issues that must be considered with respect to temperature: (1) the magnitude of the surface temperature, which is a measure of the slab temperature and influences lasing gain and optical distortions, and (2) the thermal profiles of the boundary layer, which pertain to scattering losses within the cooling gas. We have simulated

thermal load with an electrically heated indium/tin-oxide resistive coating at the slab location. We measured gas temperature within the thermal boundary layer with a small thermistor mounted to a hot-wire probe support. The surface temperature was measured using thermal imagery.

Figure 6 compares probe temperatures as calculated by the TEXTAN code and as measured by an intrusive thermal probe for flow in a 3-mm-thick straight channel. In interpreting the measured temperature, it is important to note that flow stagnates at the leading edge of any blunt object. Thus the thermal probe does not necessarily see the static temperature of the gas but rather a temperature more closely associated with stagnation

conditions. Thermal conduction within the gas reduces the measured temperature somewhat below the stagnation temperature. The additional lines on the ordinates of Figure 6 give values for the calculated and measured surface temperatures. The calculations capture very well both the shape and magnitude of the gas temperature. Figure 6 also shows the very steep temperature gradients that occur in the vicinity of the wall.

The TEXTAN-calculated slab-surface temperature profile reveals the importance of properly considering thermal boundary layer development. It is only after a relatively long development length that the TEXTAN-calculated surface-temperature gradient approaches the value associated with simple caloric heating. This has caused us to consider thermally preconditioning the flow upstream of the slab, using resistive heating, to minimize the cooling variations along the active medium.

So far, we have explored bulk flow properties of compressible channel flow as they relate to pressure and temperature profiles. These profiles determine how much surface heat flux can reasonably be removed and indicate the resulting deformations in the slab and the channel separation windows, which in turn determine some of the beam aberrations in the laser system. In general, our results show the predictive accuracy of our numerical approaches and their usefulness as tools for designing future experiments and for exploring optimized configurations that we cannot investigate experimentally.

Figure 6. Comparison of measured and calculated probe temperatures in a Mach-0.4 flow with a surface heat flux of 1.67 W/cm^2 . Upstream temperatures are shown in black; downstream temperatures in color; data points denote measured temperatures; curves denote calculated temperatures. The additional black and colored lines and data point near 330 K on the ordinate give values for the calculated and measured surface temperatures.

Optical Properties of the Gas Flow

In a generic disk amplifier architecture, the beam must traverse two types of gas flow. The first is channel flow, whose bulk properties we discussed above. The second type of flow occurs in the regions between

the disks, which we call the drift region (see Figure 7). It is not desirable to simply trap gas in the drift region, since thermal convection within these long-optical-path-length regions would give rise to severe beam distortions. Therefore, the flow in the drift region must be

conditioned to moderate thermal convection currents and temperature gradients in the gas.

Turbulent Scattering. Turbulent scattering, like all scattering, arises from variations in the index of refraction in the medium traversed by the beam.⁴ Since the index of refraction of a gas depends on its density (which in turn depends on temperature and pressure), fluctuations in the pressure and temperature cause scattering.

Pressure fluctuations arise mainly from local turbulent velocity fluctuations and acoustic waves propagating through the gas. The acoustic waves are related to boundary layer instabilities, like the trend toward transitory velocity stall at the diffuser walls.

Temperature fluctuations in the flow channel arise from the boundary layer at the heated slab. Since turbulence generation consists, so to speak, of the removal of a small parcel of sublayer fluid through a shear-driven instability, a piece of very hot gas from near the wall is carried out into midstream, where it is registered by a probe beam as a temperature fluctuation.

In the drift region the same basic process occurs. Possible large-scale temperature gradients across the reflector plate must be averaged out quickly or we will have much the same large-scale temperature problem as we encountered with convective recirculation. To accomplish this, we have the gas leave the reflector plate through an array of small holes or jets. The many jets generate a background of high turbulent intensity that rapidly mixes the drift region gas and breaks up the large-scale temperature gradients into small-scale local temperature fluctuations.

Schlieren Experiments. To investigate experimentally the problem of light scattering by a

Figure 7. The drift region is the space between amplifier disks. The flow there must be conditioned to minimize beam distortions.

Figure 8. A Schlieren setup visualizes refractive index gradients by passing only deflected rays past the knife edge (a). If the scattered light is collected in a detector (b), the amount can be measured.

turbulent flow, we used the Schlieren method (Figure 8), commonly used to look at refractive index variations. The collimated beam from a laser is directed through the flow channel or the drift region. Most of the light rays pass undisturbed through the region in question, and are imaged by a lens to form a far-field diffraction pattern that corresponds to the aperture that the test object presents to the beam. Some of the rays, however, are deflected.

If we place an obstacle, like a knife edge, at the focus, we can

stop the undeflected rays from propagating further and only the deflected rays travel past the knife edge. The deflected light can then be used to form an image of the flow channel or drift region. This image maps the locations in the aperture where the ray deflection took place, thus indicating how uniformly scatterers are distributed across the aperture. Since the pulse from the laser is much shorter than the time it takes the gas to flow across the aperture, the laser pulse freezes the motion of the gas in the channel. Finally, as indicated in Figure 8, we can collect all the scattered light in a photodetector and determine the ratio of the deflected light to the intensity incident from the laser. In this way, quantitative measurements of the scattered intensity can be made. Below, we give a few examples of such measurements as they were made with the flow channel and the drift region.

Figure 9 shows the focal plane of the lens that collects the scattered light. We used a cross mask (instead of a knife edge) to block out the cross-shaped diffraction pattern arising from the square aperture that the flow channel window presents to the probe beam. Around the mask we clearly see the elliptical halo of light that was scattered by the turbulent flow features in the channel. Several quantities of interest can be deduced from this picture. First, the light is typically scattered into an elliptical cone of 1.3×0.75 mrad, which corresponds to an average size of the scatterers in the flow channel of 1.5×0.8 mm. Hence we can say that the average eddy that causes the scattering is elliptical in shape, with the long direction aligned with the flow and, at 0.8 mm, has a thickness of about one-quarter of that of the flow channel. If we form an image of the flow channel with the scattered light, we obtain the picture in

Figure 10a. This shows that the scatterers are uniformly distributed across the aperture and scattering levels are so low that optical imperfections become clearly visible. Figure 10b is a picture obtained with a different nozzle, where the distribution of scatterers, and thus the

Figure 10. Schlieren pictures of two flow channels. (a) A 3-mm-wide flow channel at Mach 0.64. The flow is well conditioned, the distribution of scatterers across the aperture is uniform, and overall scattering levels are low. (b) A 5-mm-wide flow channel at Mach 0.3. The long features extending from top to bottom are most likely vortices generated by faulty nozzle design and not controlled by a sufficient flow development length. The distribution of scatterers is nonuniform and overall scattering levels are high.

(a) No flow

(b) Flow (Mach 0.6)

Figure 9. Intensity distribution in the focal plane with the flow channel as the test object. The halo of scattered light around the mask covering the diffraction pattern is clearly visible. Its shape and size allow us to draw conclusions about the shape and size of the turbulent features in the flow channel.

heat transfer properties of the gas, is no longer uniform across the aperture.

Proportion of Light Scattered.

We also determined the proportion of light scattered as a function of Mach number and channel width. The results (Figure 11) show that the scattered intensity scales as the Mach number to the fourth power. This strong dependence derives from a dependence of the scattered intensity on the square of density fluctuation, which in turn depends on the square of the Mach number. If we draw a vertical line through the plot at fixed Mach number, we find the scaling of scattered intensity with channel width. A good fit to the data is achieved by scaling the scattered intensity with the channel width to the 1.4th power. If the scatterers were

all much smaller than the channel width, then, for small losses, one would expect a scaling that is linear with channel width. Deviations from this linear relation arise from the largest turbulence features in the flow, which have a dimension comparable to the channel width. (We should mention here that if one measures the scattering loss with an entirely different technique, one obtains a scaling law that is more closely proportional to the channel width. We are therefore inclined to look at the 1.4th power scaling as the upper limit.)

Effect of Thermal Boundary Layer. Refractive index fluctuations associated with pressure and temperature fluctuations within the thermal boundary layer produce turbulent scattering. Hence, one

would expect the amount of light scattered to depend on the amount of surface heat flux pulled from the slab. Figure 12 quantifies this effect. It is clear that a slower flow (i.e., Mach 0.3 vs Mach 0.5) permits a hotter boundary layer, which means larger temperature fluctuations, which in turn means a larger thermal contribution to turbulent scattering of the probe beam.

The same methods described above for diagnosing turbulent scattering in the flow channel can be used on the drift region. The average flow velocity in the drift region is at most a few meters per second, corresponding to Mach numbers below 0.01; thus, scattering fluctuations from this source should be negligible. Turbulent scattering in the drift region depends essentially on temperature fluctuations only. These will arise, as mentioned above, from temperature gradients across the reflector plates through which the jets enter into the drift region.

In our experiment, we generated such a temperature difference by heating the center of the reflector plate and then measuring the temperature difference between

Figure 11. The scattered intensity as a function of the Mach number of the flow, measured for three different channel widths (2, 3, and 5 mm) with the modified Schlieren setup of Figure 8b. The scattering scales as the fourth power of the Mach number and the 1.4th power of the channel width.

Figure 12. Enhancement factor of turbulent scattering due to passage of the gas over a hot surface. In a slower flow, the gas is heated to higher temperatures, which increases the magnitude of the temperature fluctuations and in turn increases the turbulent scattering.

center and edge. Since one would expect larger turbulent features at the lower flow rates prevalent in the drift region, we used a version of the setup in Figure 8 in which the probe beams makes a double pass through the drift region. This doubles the scattering angles and makes the halo of scattered light in the focus easier to observe. Figure 13 shows the focal plane for the Schlieren setup when the drift region is examined. The halo of scattered light is now circular and covers a much narrower cone. This means that the average turbulent feature size in the drift region is around 1–2 cm, which compares well with the spacing of the jets as they leave the reflector plate.

Figure 14 shows the approximate scattered intensity in the drift region as a function of center-to-edge temperature difference of the reflector plate through which the jets enter the drift region. As expected, the scattering is essentially independent of flow speed and depends on the temperature gradient only. It is equally clear that large-scale temperature gradients across the reflector plate must be avoided. For a given illumination nonuniformity, this can be achieved by water cooling the reflector from the back, which, for operation at significant average power, has to be done anyway.

Our turbulent scattering measurements so far have given us a good idea of the magnitude of the scattering, some of its most important scaling, and the origin and suppression of some undesirable flow instabilities. We have also found that under the conditions at which we eventually intend to operate an actual amplifier, the loss due to turbulent scattering can be held within very acceptable limits, particularly when helium is used instead of nitrogen. Our present efforts in this area are

(a) No flow

(b) Fast flow; no temperature gradient

(c) Slow flow; temperature gradient

Figure 13. Intensity distribution in the focal plane with the drift region as the test object for (a) no flow, (b) flow at 1700 liters/min and no temperature gradient, and (c) flow at 850 liters/min and a temperature gradient of 17 K. The scattering occurs into much smaller angles and the average turbulent feature is not distorted into an ellipse. (Compare this figure with Figure 9 for the flow channel.)

Figure 14. Scattered intensity as a function of center-to-edge temperature difference for different flow rates. The scattering is independent of flow rate. At the very low Mach numbers involved (around 0.003), this is to be expected.

directed at an increased understanding of the turbulence-generating mechanism in the drift region so that we can eventually achieve the required turbulent intensity with a well-understood minimum in compressor power requirements.

Conclusions

From our work to date, we are confident that we can design adequately cooled solid-state laser amplifiers approaching a megawatt of high-quality output power per square meter of aperture. With this ability to scale to such output powers while preserving beam quality plus the encouraging developments of diode laser pumping, the solid-state laser is truly technically competitive with respect to average power for very high-power applications, such as a laser driver for a fusion reactor.

Key Words: computer code — TEXSTAN; gas cooling; gas flow; turbulent; laser; disk amplifier; solid state

Notes and References

1. J. I. Emmett, W. F. Krupke, and J. B. Trenholme, "The Future Development of High-Power Solid-State Lasers," *Soviet J. Quantum Electron.* **13**, 1 (1983).
2. J. I. Emmett, W. F. Krupke, and W. R. Sooy, *The Potential of High-Average-Power Solid-State Lasers*, Lawrence Livermore National Laboratory, Rept. LCR-53571 (1984).
3. Disk amplifier development is described in the LLNL *Laser Program Annual Reports* from 1974 on.
4. TEXSTAN is a variation, produced by M. E. Crawford, of the STAN5 computer program; M. E. Crawford and W. M. Kays, STAN5—A Program for Numerical Computation of Two-Dimensional Internal and External Boundary Layer Flows, National Aeronautics and Space Administration, Rept. NASA CR-2742 (November 1976).
5. See for example, V. I. Tatarski, "The Effects of the Turbulent Atmosphere Wave Propagation," available from DoC, NTIS, ISBN-0-7065-0680-4. Translated from Russian in 1971.

Abstracts

Measuring the Performance of Supercomputers

We have developed an efficient way of estimating or benchmarking supercomputer performance for typical workloads of numerically intensive physics computations. The effective CPU (central processor unit) performance of supercomputers now ranges from a few to several hundred Mflops (million floating-point operations per second). Attempts to reduce this two-order-of-magnitude range of CPU performance to a single number are not valid and have produced much confusion. No single quotation of performance is sufficient. The Livermore Fortran Kernels (LFK) test constitutes a compact, efficient benchmark for Fortran computations that can be used to measure a realistic numerical performance range. Compared to application programs and to typical benchmark programs, the LFK benchmark is much smaller, much faster, and less expensive to run. The LFK may be used to simulate or measure the performance and efficiency of the CPU and to determine the accuracy of compiled Fortran code. The relative performance measurements based on LFK benchmarks generally correlate with the speedups measured with CPU-limited benchmarks, application programs, and computing system workloads. Ultimately, the level of performance of a workload on a supercomputer relates most closely to the level of development of efficient, parallel computer codes.

Contact: Frank H. McMahon (415) 422-1647.

Microwave Tokamak Experiment

In the Microwave Tokamak Experiment (MTX), now under construction at the Laboratory, the free-electron laser is combined with the magnetic confinement approach to fusion. Microwave heating from a free-electron laser will be used to heat the magnetically confined plasma. New diagnostics are being developed for the MTX, an international collaborative effort that will be operational in late 1988. The MTX will allow us to study dense plasmas heated by powers unprecedented in the electron-cyclotron frequency range, to address new physics, to investigate energy-confinement mechanisms in the tokamak, and to advance the state of fusion research.

Contact: E. Rick Hooper (415) 422-6792, William M. Nevins (415) 422-7032, or Ronald H. Cohen (415) 422-9831.

Gas Cooling of Laser Disks

Our experiments and calculations show that effective cooling of solid-state laser disk amplifiers can be provided by flowing gas in adjacent channels across laser disk surfaces. Turbulent scattering measurements give us a good idea of the magnitude of the scattering, some of its most important scaling, and the origin and suppression of some undesirable flow instabilities. The measurements have also shown that under the conditions at which we eventually intend to operate an actual amplifier, the loss due to turbulent scattering can be held within very acceptable limits. Our data and calculations indicate that we can design laser amplifier beamlines that approach a megawatt of average output power per square meter of aperture in a high-quality beam. The ability to scale to such output powers while preserving beam quality makes the solid-state laser a technically competitive candidate for very high-power applications, such as a laser driver for a fusion reactor.

Contact: Georg F. Albrecht (415) 423-7562 or Steven B. Sutton (415) 422-0322.

Recent Titles

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

DEFENSE PROGRAMS

Inertial Fusion

Gas Cooling of Laser Disks (May 1988)

Military Application

Measuring the Performance of Supercomputers (May 1988)

Composite Materials: Characterizing Their Mechanical Properties (April 1988)

Discovery of Two New Long-Lived Isotopes: ^{261}Lr and ^{262}Lr (March 1988)

Search for $^{272}[109]$ in a New Region of Stability (March 1988)

The Lightning-Invulnerable Device System (March 1988)

Detonation of Explosives in Accidents (January-February 1988)

The Detonation of High Explosives (January-February 1988)

Formulating High-Explosive Materials (January-February 1988)

High-Explosive Diagnostics (January-February 1988)

Synthesis of Energetic Materials (January-February 1988)

Synchrotron Radiation and Nuclear Test Diagnostics (November-December 1987)

Safeguards and Security Upgrades at LLNL (October 1987)

Safeguards and Security

Measuring Special Nuclear Material (October 1987)

Risk Management and Safeguards Evaluation (October 1987)

Technical Security (October 1987)

Validating Safeguards and Security's Effectiveness (October 1987)

ENERGY RESEARCH

Microwave Tokamak Experiment (May 1988)

Development of an Aluminum-Air Battery for an Electric Vehicle (April 1988)

The Rotating Target Neutron Source (March 1988)

LLNL

Director's Initiative Research

Recombinant DNA Biotechnology and the Bioconversion of Metals (April 1988)

Overviews

Elemental and Chemical-State Imaging (November-December 1987)

Multilayer Optics for X Rays, Soft X Rays, and Extreme Ultraviolet Light (November-December 1987)

Synchrotron Radiation in X-Ray Science and Technology (November-December 1987)

The X-Ray Calibration and Standards Facility (November-December 1987)

X-Ray Absorption Spectroscopy (November-December 1987)

X-Ray Diffraction Applications (November-December 1987)

WORK FOR OTHERS

Defense Advanced Research Projects Agency

Gas Cooling of Laser Disks (May 1988)

Department of State

Measuring Special Nuclear Material (October 1987)