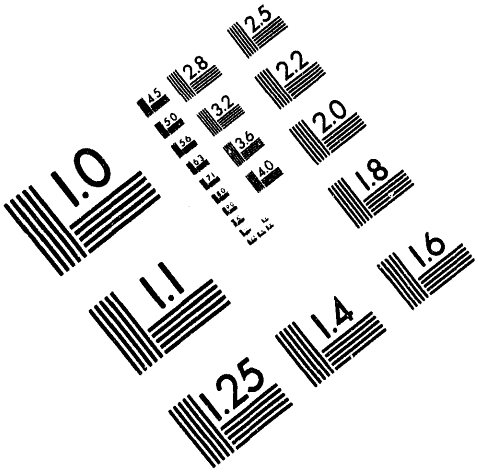
1.0
1.1
1.25
1.4
1.6
1.8
2.0
2.2
2.5
2.8
3.2
3.6
4.0
4.5
5.0
5.6

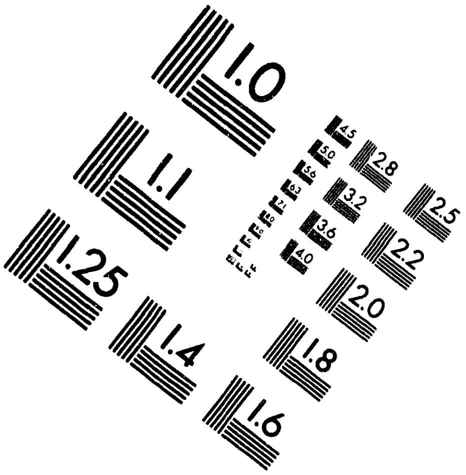
1.0
1.1
1.25
1.4
1.6
1.8
2.0
2.2
2.5
2.8
3.2
3.6
4.0
4.5
5.0
5.6


Centimeter
1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
mm
1
2
3
4
5
Inches

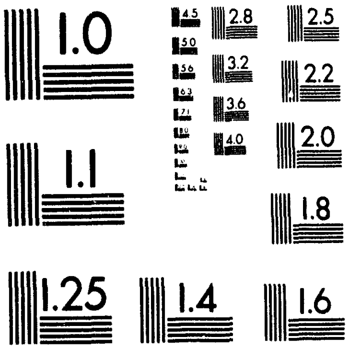
1.0
1.1
1.25
1.4
1.6
1.8
2.0
2.2
2.5
2.8
3.2
3.6
4.0
4.5
5.0
5.6
6.3
7.1

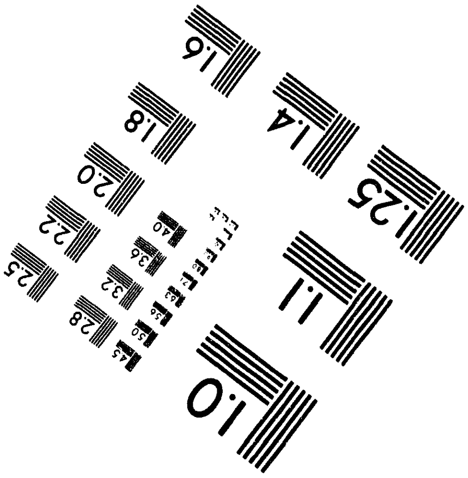
1.0
1.1
1.25
1.4
1.6
1.8
2.0
2.2
2.5
2.8
3.2
3.6
4.0
4.5
5.0
5.6

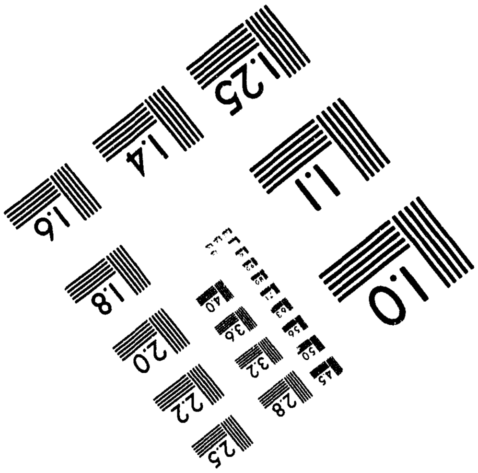
1.0
1.1
1.25
1.4
1.6
1.8
2.0
2.2
2.5
2.8
3.2
3.6
4.0
4.5
5.0
5.6

1 of 3

THRUST AREA REPORT • UCRL 53868-93

ENGINEERING RESEARCH, DEVELOPMENT AND TECHNOLOGY

FY93

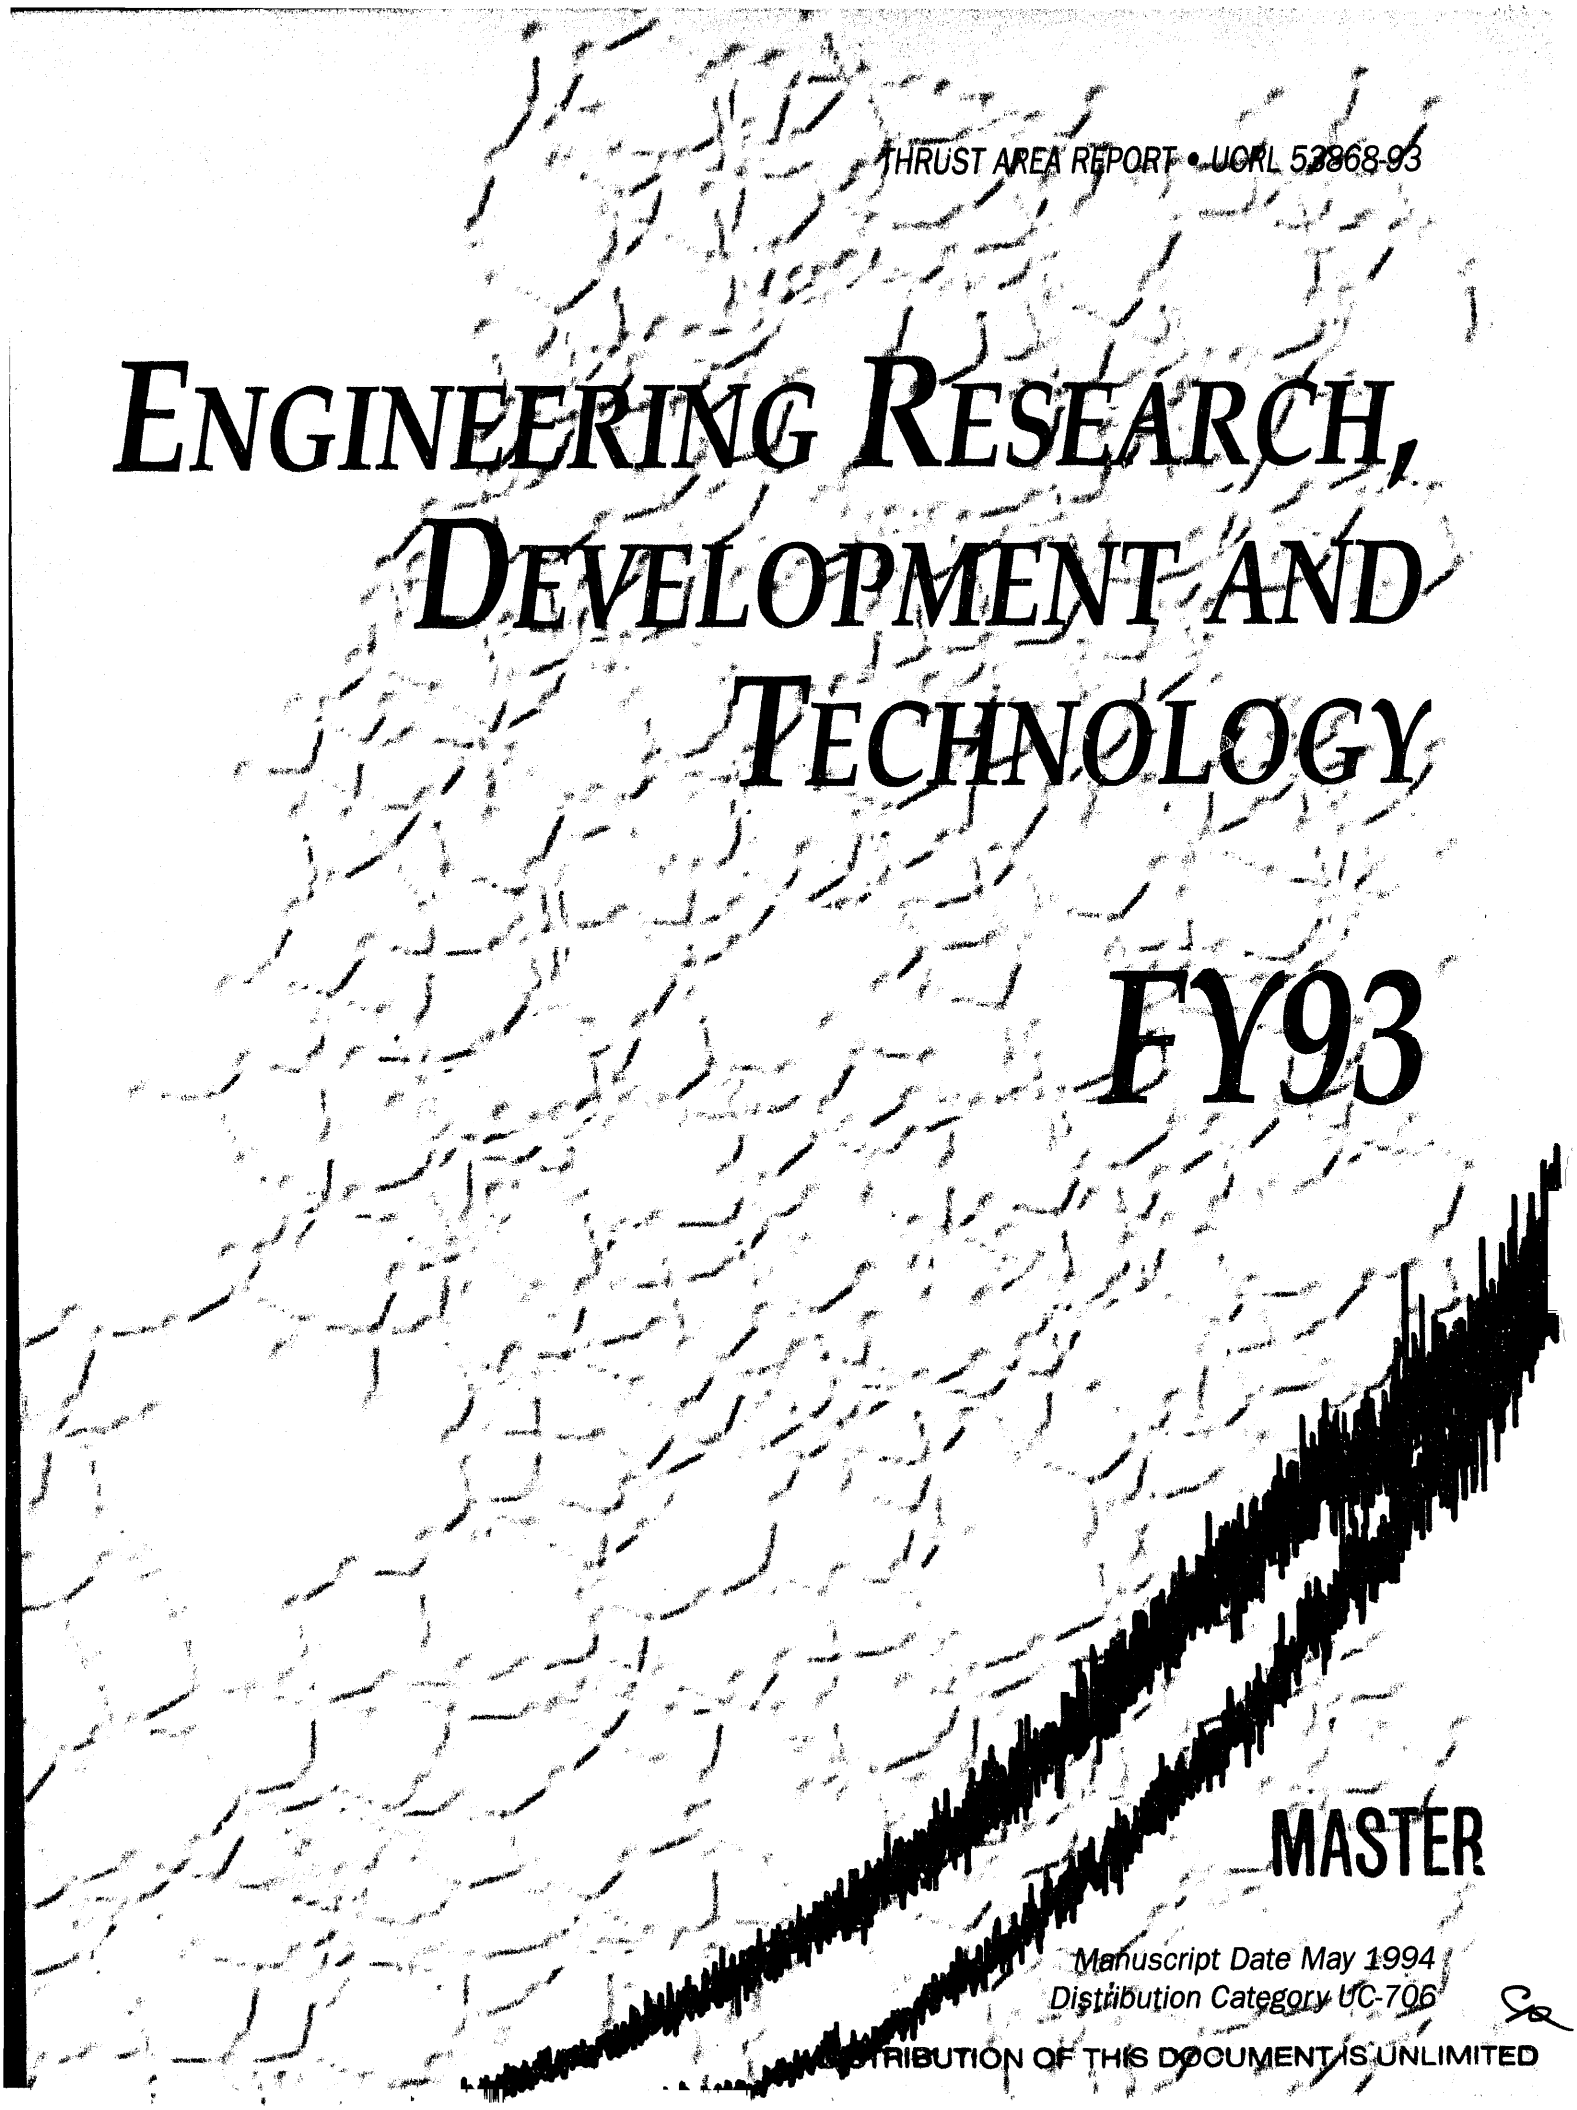

MASTER

Manuscript Date May 1994
Distribution Category UC-706

DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

1. Computational Electronics and Electromagnetics

2. Computational Mechanics

3. Diagnostics and Microelectronics

Overview
Joseph W. Balch, Thrust Area Leader

4. Emerging Technologies

Overview
Shin-yee Lu, Thrust Area Leader

5. Fabrication Technology

6. Materials Science and Engineering

7. Power Conversion Technologies

8. Nondestructive Evaluation

9. Remote Sensing, Imaging, and Signal Engineering

Author Index

Introduction

The mission of the Engineering Research, Development, and Technology Program at Lawrence Livermore National Laboratory (LLNL) is to develop the technical staff, tools, and facilities needed to support current and future LLNL programs. Our efforts are guided by a dual-benefit research and development strategy that supports Department of Energy missions, such as national security through nuclear deterrence and economic competitiveness through partnerships with U.S. industry.

To accomplish this mission, the Engineering Research, Development, and Technology Program has two important goals: (1) identify key technologies relevant to LLNL programs and (2) conduct high-quality research and development to enhance our capabilities in these technologies. To focus our efforts, we identify technology *thrust areas* and technical leaders for each area. The thrust areas are integrated engineering activities, staffed by personnel from Electronics Engineering, Mechanical Engineering, and other LLNL organizations.

The thrust area leaders are accountable to me for the quality and progress of their activities, but have sufficient latitude to manage the resources allocated to them. They are expected to establish strong links to LLNL programs and to industry; to outside and inside experts to review the quality and direction of their work; to use university contacts to supplement and complement their efforts; and to be knowledgeable in the state of the art outside of LLNL.

For the past several years, the Engineering Research, Development, and Technology Program has played a significant role in the technology transfer process, as the foundation for many of the Cooperative Research and Development Agreements now in place between LLNL programs, universities, and U.S. industries.

This annual report, organized by thrust area, describes our activities for the fiscal year 1993. The report provides timely summaries of objectives, methods, and results from nine thrust areas for this fiscal year: Computational Electronics and Electromagnetics; Computational Mechanics; Diagnostics and Microelectronics; Fabrication Technology; Materials Science and Engineering; Power Conversion Technologies; Nondestructive Evaluation; Remote Sensing, Imaging, and Signal Engineering; and Emerging Technologies.

Readers desiring more information are encouraged to contact the individual thrust area leaders or authors.

Roger W. Werne
Associate Director for Engineering
and Technology Transfer

Computational Electronics and Electromagnetics

The Computational Electronics and Electromagnetics thrust area is a focal point for computer modeling activities in electronics and electromagnetics in the Electronics Engineering Department of Lawrence Livermore National Laboratory (LLNL). Traditionally, we have focused our efforts in technical areas of importance to existing and developing LLNL programs, and this continues to form the basis for much of our research. A relatively new and increasingly important emphasis for the thrust area is the formation of partnerships with industry and the application of our simulation technology and expertise to the solution of problems faced by industry.

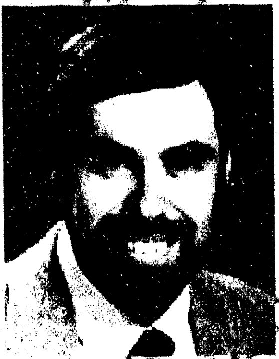

Our FY-93 work is described in the five reports that follow: (1) Parallel Computers and Three-Dimensional Computational Electromagnetics; (2) Computational Integrated Photonics; (3) Development of Millimeter-Wave Sources and Components for the Heating of Fusion Plasmas; (4) Electromagnetic Modeling for Target-Rich Embedded Environments; and (5) Acoustic, Finite-Difference, Time-Domain Technique Development.

John F. DeFord
Thrust Area Leader

Section 1

1. Computational Electronics and Electromagnetics

Overview
John F. DeFord, Thrust Area Leader

Parallel Computers and Three-Dimensional Computational Electromagnetics

Niel K. Madsen
Engineering Research Division
Electronics Engineering

We have continued to enhance our ability to use new massively parallel processing computers to solve time-domain electromagnetic problems. New vectorization techniques have improved the performance of our code DSI3D by factors of 5 to 15, depending on the computer used. New radiation boundary conditions and far-field transformations now allow us to compute radar cross-section values for complex objects. A new parallel-data extraction code has been developed that allows the extraction of data subsets from large problems, which have been run on parallel computers, for subsequent post-processing on workstations with enhanced graphics capabilities. A new charged-particle-pushing version of DSI3D is under development. Finally, DSI3D has become a focal point for several new Cooperative Research and Development Agreement activities with industrial companies such as Lockheed Advanced Development Company, Varian, Hughes Electron Dynamics Division, General Atomic, and Cray.

Introduction

The solution of physical problems whose behavior is governed by Maxwell's equations has been our focus for several years. Our activity in this area has resulted in the development of an electromagnetic simulation code DSI3D. For the past two years, we have concentrated on making the simulation code run efficiently on the new massively parallel processing (MPP) computers. These new computers will allow us to model much larger problems more accurately than in the past.

Most of the new MPP computers now have processors with vector processing units that can process certain types of vector data very quickly. Performance improvements of factors ranging from 5 to 15 have been observed on various types of processors.

New capabilities have also been added to DSI3D that now allow the solution of open region problems (problems without physical constraining boundaries). These new absorbing or radiation boundary conditions and far-field transformations allow for the evaluation of electromagnetic fields at distances that are not actually within the computational space. These are useful for computing radar cross-sections (RCS) of aircraft. Modern stealth aircraft are designed to have low RCS so that they are difficult to detect through the use of radar.

The use of MPP computers will now allow us to compute problems than are considerably larger than those previously possible to compute. The amounts of data produced are potentially so large that it will be increasingly difficult to move the entire data sets from the MPP computer to smaller graphics workstations typically used to visualize the data. Therefore, we have developed a new parallel program that allows us to extract relevant sub-parts of the data that are of specific interest to us.

To be able to address types of plasma physics problems, a new charged-particle-pushing capability has been added to DSI3D. This work presents many new challenges for the development of the code.

As a result of our increased simulation capabilities, DSI3D has become a common focal point as a simulation tool for the work specified in a number of Cooperative Research and Development Agreements (CRADA's). The companies involved include Lockheed Advanced Development Company (LADC), Varian, Hughes Electron Dynamics Division, General Atomic, and Cray.

Progress

Vectorization

In solving very large three-dimensional simulation problems, it is almost always essential to take advantage of special features of various computers so that they may be used more efficiently. Many modern computers have the ability to process vector data (arrays or strings of similar type data) much more efficiently than they can process scalar data (single variables or values). If one can organize the scalar data into vectors of considerable length, and the computational sequence so that all the data of a particular type can be processed at once, then great improvements in efficiency can usually be achieved. DSI3D data is complicated by the fact that it often is derived from unstructured grids and data structures, making vectorization difficult. Over the past year, we have discovered a way to vectorize the DSI3D data: reorder the data so that many partial dot products of vectors of similar size can be computed simultaneously. Previously, the DSI3D code ran at about a speed of 15 to 20 MFLOP (million floating points operations per second) on a Cray-C90 single processor. With the new structuring of the data, we now observe performance as high as 300 MFLOP on a single processor, improving the performance by a factor of about 20. On the new Meiko CS-2 MPP computer, we observe improvements of about a factor of 5 per processor. Improvements of these magnitudes make the relatively small amount of extra work in reorganizing the data well worthwhile. These improvements in efficiency will also allow us to solve larger problems much faster.

New Boundary Condition Capability

Since it is clearly impossible to discretize all of space for most problems, it is essential that one have an ability to truncate the computational space so that it includes only the essential regions of specific interest. One example of this type of problem is that of assessing the interaction of radar pulses with an aircraft flying at a reasonable altitude. The radar signal usually originates a long distance (in terms of the signal wavelength) from the aircraft. The incident signal interacts with the aircraft and is reflected in various directions. The radar receiver antenna is often located at the site of the signal origination (i.e., at a considerable distance from the aircraft). If one were to attempt to solve Maxwell's equations for the entire region including all the space between the radar and the aircraft, one would quickly decide that this would not be feasible with any computer imaginable. The problem would simply be too large and costly. However, one can solve this problem by solving Maxwell's equations in only a relatively small region surrounding the aircraft. To accomplish this in a practical sense, one needs the ability to incorporate absorbing-radiation boundary conditions into a simulation code. These boundary conditions are designed to allow outwardly radiating electromagnetic fields that are scattered from the aircraft to simply exit the problem region without being spuriously reflected back into the region, where they might erroneously re-interact with the aircraft. There are a variety of techniques with varying degrees of accuracy that have been developed to address this specific problem. However, none of these methods has ever been used for an unstructured-grid Maxwell solver such as DSI3D.

During FY-93, we implemented and tested two of these techniques, the second-order Mur condition and the stabilized Liao boundary condition. We chose to use the Liao condition since it can be implemented with fewer restrictions. In conjunction with the radiation condition, we have implemented far-field transformations of the Luebbers–Kunz type to use data obtained near the scattering object (aircraft) to compute the fields at large distances from the scattering object. These two enhancements give us the general capability to solve open-region problems, and in particular RCS problems.

Parallel Data Extraction for Visualization

Getting large-problem data sets onto and off MPP computers can be a considerable and sometimes impossible problem. The data is frequently much too large even to fit onto smaller workstations typically used for visualizing the data. We have designed and implemented a new parallel tool, DSI3DIO, to be used in conjunction with DSI3D, which runs on the MPP computer and simply extracts desired data subsets from the various different processors. The data to be extracted can be specified by various material-type parameters. The volumetric data is obtained from the restart files left behind by an executing simulation. The data is assembled from the many processors, and only its exterior surface information is retained, assembled, and transferred to the workstation for visualization. Ultimately, we will want to perform visualization and rendering of images directly on the MPP computer, but this process gives us the ability to display desired data with the currently available software on low-cost and performance workstations.

Charge-Particle Pushing and Tracking

To be able to solve certain low-density plasma physics problems, we have begun the development and implementation of a capability to include the motion of charged particles in a self-consistent manner with the Maxwell field-solution process on unstructured grids. This type of problem has been solved before, using only structured grids. Again, the additional geometric flexibility gained through the use of unstructured grids is needed. We currently have implemented the charged-particle capability for grids that consist of either hexahedral cells or tetrahedral cells, exclusively. Mixed-cell-type grids cannot be handled currently. The real challenge in the development of a parallel processing version of this new code is that of dealing with the 'load balancing' issue created because the particles can move freely. Since most of the computational costs involve computing the particle locations, trajectories, and charge distributions, a static partitioning of the problem domain will not suffice for efficient use of the MPP computer. We are currently addressing this issue.

Future Work

DSI3D has become of considerable interest to various industrial companies. We have established a CRADA with LADC to pursue further developments in the specific area of computing RCS's of aircraft and various objects. A CRADA with Varian has been established to assist in developing and designing windows used with their high-power gyrotrons, to allow the fields to escape from the gyrotron where they are generated and to be transmitted to where they will be used. For Hughes Electron Dynamics Division, DSI3D will be used to analyze the behavior and performance of their traveling wave tubes, which are used for space communication. The CRADA with General Atomic is concerned with the heating of plasmas in tokamak fusion reactors by using powerful electromagnetic fields transmitted into the reactor through waveguides. DSI3D will be used to study techniques for making this process more efficient. Finally, Cray is interested in developing application and simulation codes such as DSI3D that will perform very efficiently on their new designs of MPP computers.

Computational Integrated Photonics

Jeffrey S. Kallman and
Raymond J. Hawkins
Engineering Research Division
Electronics Engineering

Ewan M. Wright
Optical Sciences Center
University of Arizona
Tucson, Arizona

Teresa L. Swatloski
Applications Systems Division
Computation Directorate

Richard W. Ziolkowski
Department of Electrical
and Computer Engineering
University of Arizona
Tucson, Arizona

Jonathan P. Heritage
Department of Electrical
and Computer Engineering
University of California
Davis, California

In FY-93, we accomplished several tasks. In the area of finite difference time domain (FDTD) simulation, we developed two formulations for the straightforward inclusion of gain: a means for modeling gain by negative conductivity, and a way of modeling wavelength-dependent gain via anomalous dispersion. Using these techniques, we simulated higher-order mode lasing in optical traveling-wave amplifiers. We developed pseudo-spectral codes to model both flat and curved gratings, and introduced a second version of BEEMER, which we distributed to new sites. We continued investigations using short-time, nonlinear FDTD to model a number of devices. We completed the final report on the vertical impedance-matched photodiode. We simulated and compared to experiment annealed, proton-exchange waveguide formation and spectral characteristics. We also investigated dark solitons.

Introduction

Computational integrated photonics (CIP) is the area of computational physics that studies the propagation of light in optical fibers and in integrated optical circuits (the photonic equivalent of electronic circuits). The purpose of CIP simulation is to develop the computational tools that will support the design of photonic and optoelectronic integrated devices. These devices will form the basis of future high-speed and high-bandwidth information processing systems, and are key to the future industrial competitiveness of the U.S. CIP has, in general, two thrusts: (1) to develop predictive models of photonic device behavior that can be used reliably to enhance the speed with which designs are optimized for applications; and (2) to further the ability to describe the linear and nonlinear processes that occur and can be exploited in real photonic devices.

Our efforts in FY-93 focused on seven areas: (1) extension of the capabilities of linear finite difference time domain (FDTD) codes; (2) pseudo-spectral modeling of grating structures; (3) modeling the formation of optical structures by the process of annealed proton exchange; (4) development of new modeling techniques for the analysis of vertical impedance-matched photodiodes; (5) investigation of nonlinear devices via nonlinear FDTD methods; (6) investigations into the nature and uses of dark solitons; and (7) development of the user interface for our two-dimensional device modeling code BEEMER to meet outside industry standards so the code can be more widely used.

Progress

Linear FDTD Extensions

Semiconductor laser amplifiers are of particular use in many fiber-optics communications applications. The traveling wave amplifier (TWA) is especially attractive because of its wide bandwidth and high saturation-output power.[1] Oscillation in the cavity of a TWA is prevented by reducing the reflectivity of the end facets. This has been accomplished both through dielectric anti-reflection coatings on the facets and/or by tilting the facets. Tilting the facets (or, equivalently, the waveguide, with respect to the facet cleave) is the simpler way of reducing the facet reflectivity. The application of dielectric anti-reflection coatings requires extremely tight control over both the refractive indices and the thickness of the dielectric layers.[2] Tilted facets have been demonstrated in a TWA, with a peak gain of 19 dB at a wavelength of 1.49 μ.[3] This device did not, however, have a gain-independent output field. At an injection current level of 50 mA, the far-field intensity was single lobed, indicating that it was dominated by the fundamental mode. At an injection current level of 70 mA, the far field was dual lobed, indicating an increased contribution from the first-order mode. At an injection current of 100 mA, the far field was dominated by the first-order mode, and the device oscillated. We studied the tilted waveguide TWA experimentally, and observed the transition from amplifier to first-order-mode oscillator.

To perform these numerical experiments, we needed to find a way of incorporating gain into the material models used in our FDTD codes. We discovered two means by which this could be accomplished[4]: the first involved assigning a negative conductivity to a material; the second used a negative anomalous dispersion.

For many years, positive conductivity has been used by the microwave community to simulate dispersive and absorbing media. It occurred to us that if positive conductivity could be used to simulate the attenuation aspects of a dispersive medium, then a negative conductivity might be useful, and comparatively simple, for simulating a medium with gain. Testing this idea was straightforward, and we found that it worked quite well. Also straightforward was modeling analytically what it means for a material to have a negative conductivity. We found that, for a plane wave moving in the z direction in a medium with negative conductivity, the transverse electric field can be described by:

$$E(z,t) = E_0\, e^{-i\omega t}\, e^{i|k|z\cos(\theta/2)}\, e^{-|k|z\sin(\theta/2)},$$

where

$$|k| = \left(\frac{\omega^4 n^4}{c^4} + \mu^2\sigma^2\omega^2\right)^{1/4},$$

and

$$\theta = \arctan\left(\frac{\mu\sigma\omega}{\omega^4 n^4 / c^4}\right).$$

In the above equations, ω is the angular frequency of the wave, μ is the magnetic permeability of the material, σ is the conductivity of the material, n is the index of refraction of the material, and c is the speed of light in a vacuum. The gain comes from the term $-|k|\sin(\theta/2)$. If the conductivity is positive, then the angle θ is positive, and this term (in the exponential) acts as an attenuator. If the conductivity is negative, then the angle θ is negative, and the term acts as a gain.

Since the use of conductivity to model dispersion is what prompted us to use negative conductivity to model gain, it was a natural step to use negative dispersion as another method to model gain. We had previously developed and implemented a simple method of modeling dispersion[5]; we tested this idea and found it useful. The analytic model is a bit more complicated for this gain simulation, but there are advantages to using this formulation. The gain in this model is wavelength dependent, therefore any desired gain curve can be approximated using a number of resonances. These results are reported in a paper to be published in the near future.

Pseudo-Spectral Grating Models

It has been estimated that there is more than a terahertz of bandwidth available for use on modern optical fiber, but only a few percent of that bandwidth are currently used. The reason for this disparity is the difficulty of getting signals onto and off the fiber. Lasers can be directly modulated to frequencies in the low GHz, and very-high-speed modulators can put signals on a fiber with bandwidths in the tens of GHz, but at present, nothing uses the complete bandwidth of a fiber. There are a number of proposed ways of exploiting this opportunity. At the Lawrence Livermore National Laboratory (LLNL), we are exploring wavelength division multiplexing (WDM), using a device based on a Rowland Circle grating, which has the following property: if light is launched from a small spot on the Rowland Circle, the diffracted light that comes off the grating will be in focus on the Rowland Circle. This means that if we have a number of disparate wavelengths of light, all carrying signal information coming down a single optical waveguide that ends on the Rowland Circle and illuminates the Rowland Circle grating, the grating will separate the individual wavelengths and put them at different spots on the Rowland Circle, where they can be led away in individual fibers. The signals modulating them can then be detected. If the input light were to be used to illuminate a flat grating, separate focusing optics would be required.

We used a pseudo-spectral scalar field method to model the effects of a Rowland Circle grating. This program, GREAT_CIRCLE, has a fully integrated graphical user interface and allows the user to specify parameters such as materials, indices, wavelengths, input and output regions, and the configuration of the teeth. This program's interface and output are illustrated in Fig. 1.

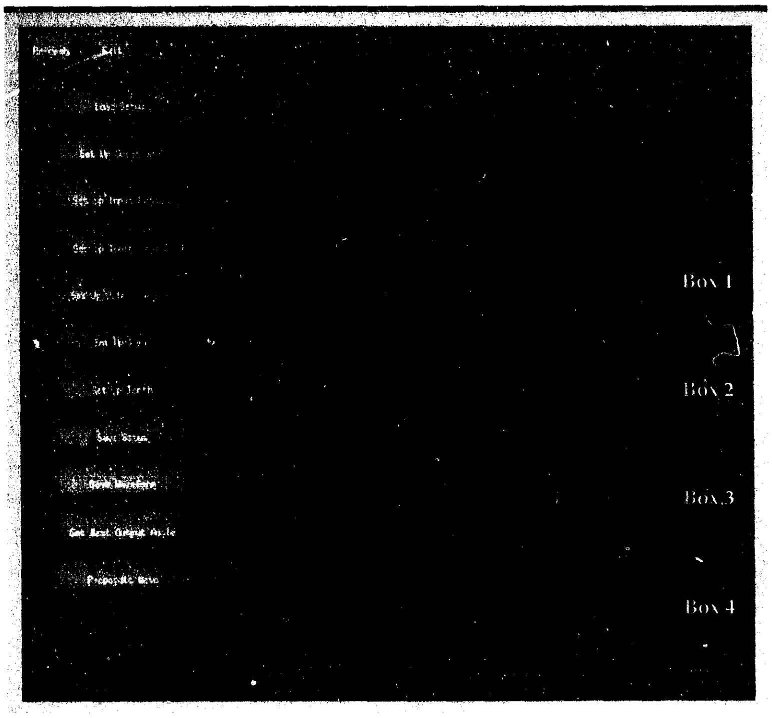

Figure 1. Graphical user interface to the program GREAT_CIRCLE. This interface allows the user to define the Rowland Circle, where the input and output waveguides should be placed, the form of the grating teeth, and the input and expected output mode shapes. The program computes the result of the input mode shape's hitting the grating teeth and propagating to the output waveguide, where it is compared to the expected output mode shape. The waveform in Box 1 is the input mode shape; Box 2 and Box 3 contain the grating shape and the intensity of the light hitting the teeth, respectively. In Box 4, the lower curve is the intensity of the output light; and the upper curve is the expected intensity, if all the input light had reached the output waveguide. The square box (upper left) contains the physical layout of the simulation.

Modeling Annealed Proton Exchange Waveguides

To have a multitude of optical devices on a single chip, it is necessary to be able to direct light from place to place. Light moves from place to place on a chip via waveguides, and one of the more recent methods of forming waveguides is through annealed proton exchange. This is a relatively new process and so is an art rather than established engineering practice. One of our endeavors this past year has been to model the formation of these waveguides and compare our simulations with samples produced at Stanford University. The diffusion process was modeled for a situation that was used to create an actual on-chip waveguide, and an estimate of the refractive index profile was generated. This was fed to a pseudo-spectral, beam-propagation-method simulator, and a mode spectrum was generated for each of a number of wavelengths. This was compared to experiment to see how well simulation matched experiment. Results are promising.

Vertical Impedance-Matched Photodetectors

One of the critical components in an optical communication system is the photodetector. This is the device that converts optical signals into electrical signals, and it must be as fast as possible. Large photodetectors have large capacitances, and are therefore slow. We have found a way to create smaller photodetectors in the vertical impedance-matched photodiode.

In conventional photodetector designs, optical energy is converted to electrical signals in an absorbing layer grown directly adjacent to the waveguide core, and energy is transferred from the waveguide to the absorber via evanescent field coupling. These conventional devices typically have detector lengths on the order of 100 μ.

In response to the need for shorter detector lengths, Deri and Wada[6] proposed the vertically integrated, impedance-matched waveguide/photodetector, in which a thin matching layer is inserted between the waveguide core and the absorbing layer. By judicious design of this matching layer, the length of the absorber can be reduced by a factor of 5, with a consequent reduction in capacitance and increase in speed.

This year, we developed a new coupled mode theory that provided a particularly convenient means by which the physics of these devices could be modeled.[7] This coupled mode theory allows a reduction of the device design parameter space, and extraction of the key parameters for device optimization—a formidable, if not impossible, task based on time-consuming, beam-propagation-method simulations alone. Using the coupled-mode theory, we discovered a new regime of operation of these devices in which the spatial transients could be eliminated, so that the diode absorption is exponential as in conventional devices.

Nonlinear FDTD Modeling

We have continued our development of a multidimensional, full-wave, vector Maxwell's equation simulator that models problems describing the interaction of ultra-short pulsed beams with a nonlinear material having a finite response time. This simulator is built around a FDTD engine that solves Maxwell's equations in two space dimensions and time. We have extended our previous simulator, which used a phenomenological time relaxation (Debye) model of a nonlinear Kerr material, to one that incorporates a Lorentz linear-dispersion model with both a Raman and an instantaneous, nonlinear Kerr model of the medium. This improved nonlinear FDTD (NL-FDTD) simulator has been used to study a variety of problems involving nonlinear, optical-guided waveguides and output couplers.

The NL-FDTD approach can readily handle very complex, realistic structures. The TE and TM nonlinear optics problems that were treated with the NL-FDTD simulator highlight differences between (1) the scalar and the vector approaches and (2) the effects of the finite response time of the medium. The NL-FDTD method is beginning to resolve several very basic physics and engineering issues concerning the behavior of the full electromagnetic field during its interaction with a non-linear, finite-response-time medium. In particular, during this past year we have:

(1) shown the viability of combining the linear Lorentz dispersion model, the instantaneous Kerr nonlinear model, and the Raman nonlinear model with the FDTD engine;

(2) studied one- and two-dimensional solitary waves with the improved NL-FDTD simulator;

(3) revisited the TM and TE models of an optical diode (linear-nonlinear interface switch) with the improved NL-FDTD simulator;

(4) characterized the performance of various linear–nonlinear slab waveguide structures; and

(5) characterized the performance of grating-assisted output couplers. We have shown that these coupler geometries can be designed to achieve beam-steering that is intensity-dependent; i.e., the output coupling angle is directly dependent on the intensity of the exciting wave.

In all of these analyses, we have identified the role of the longitudinal field component and the resulting transverse power flows in the associated scattering-coupling processes. Comparisons between the nonlinear structures and their linear counterparts were performed for all cases to quantify the effects of the nonlinear response of the medium on the basic device designs.

This work has been reported and well received in a variety of refereed journals and at professional meetings.[8–11] Our future efforts will include extending the NL-FDTD model to more complex propagation and scattering geometries.

Dark Soliton Investigations

We are also examining the phenomenon of dark solitons as a possible means of rapid switching of optical signals. A dark soliton is formed in a nonlinear material with a negative Kerr nonlinearity when an optical field with a 180° phase jump is introduced. The Kerr nonlinearity changes the index of refraction as a function of the intensity of the optical electric field, so a negative Kerr nonlinearity causes a decrease in the index of refraction of the material when a bright light is shown through it. In a linear material, a 180° phase jump spreads out and disperses, but in a material with a negative Kerr nonlinearity, the phase jump remains sharp. In addition, the index of refraction of the illuminated area is decreased everywhere except at the phase jump. A low power signal-carrying beam can be inserted into the phase jump region, and the higher index of refraction will keep it confined to the phase jump region. By varying the angle of the phase jump, the phase jump can be made to propagate at an angle, and thus the signal beam can be steered. A simulation of this effect can be seen in **Fig. 2**.

MOTIF for BEEMER

Our final effort in FY-93 involved bringing the user interface of BEEMER up to an industry-standard MOTIF interface. BEEMER was the first of our codes to be converted to MOTIF; others will follow during FY-94. Using the MOTIF interface, we were able to port BEEMER to the IBM RS/6000 Risc station and to HP 700 series computers, as well as install the program at the IBM Almaden Research Center and the Hewlett Packard Laboratories. **Figure 2** shows this new interface in use to simulate the propagation of a dark soliton.

Future Work

We will continue our efforts in the design of novel integrated optical devices, both for LLNL programs and for industry. It is our intention to transfer TSARLITE to a MOTIF interface and then to industry. Our development of linear FDTD applications to integrated optics will be extended to three-dimensional structures. Our studies of nonlinear FDTD will continue in the area of nonlinear waveguides and couplers.

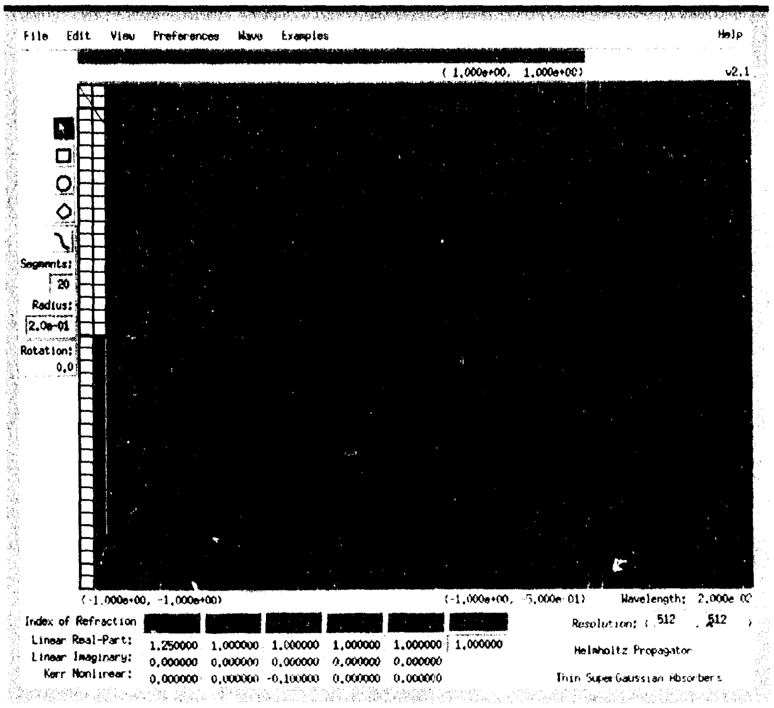

(1,000e+00, 1.000e+00)

Segments: 20

Radius: 2.0e-01

Rotation: 0.0

(-1.000e+00, -1,000e+00)

(-1,000e+00, -5,000e-01)

Wavelength: 2.000e 02

Index of Refraction						
Linear Real-Part:	1.250000	1,000000	1.000000	1,000000	1.000000	1,000000
Linear Imaginary:	0,000000	0,000000	0.000000	0,000000	0.000000	
Kerr Nonlinear:	0,000000	0,000000	-0,100000	0,000000	0.000000	

Resolution: (512 , 512)

Helmholtz Propagator

Thin SuperGaussian Absorbers

Figure 2. Simulation of a dark soliton propagating in a negative Kerr medium. The input to this simulation was a uniform bright field, half of which underwent a ~ 180° phase jump when it propagated through the gray material (a linear material with an index of refraction higher than the substrate). The wave with its phase jump entered the black material (a nonlinear material with a negative Kerr nonlinearity), and the phase jump, which takes the form of a sharp dip, was propagated without degradation. The edges of the simulation have rolled off because the simulation uses absorbing regions to keep light that leaves the simulation from the bottom of the display, from reentering at the top. This simulation was performed on the latest version of BEEMER, which has a MOTIF-based user interface. The user can now specify the objects in the simulation directly on the screen using the mouse, and all simulation parameters are available for manipulation, either directly on the screen or via selections from the menu at the top of the screen.

---

1. T. Mukai, K. Inoue, and T. Saitoh, *Electron. Lett.* **23**, 396 (1987).
2. T. Saitoh and T. Mukai, *J. Lightwave Technol.* **LT-3**, 288 (1985).
3. C.E. Zah *et al.*, *Electron. Lett.* **23**, 990 (1987).
4. R.J. Hawkins and J.S. Kallman, *Lasing in Tilted-Waveguide Semiconductor Laser Amplifiers*, Lawrence Livermore National Laboratory, Livermore, California, UCRL-JC-114272; accepted for publication in *Optics and Quantum Electronics* (1993).
5. R.J. Hawkins and J.S. Kallman, *J. Lightwave Technol.* **11**(11), 1872 (1993).
6. R. Deri and O. Wada, *Appl. Phys. Lett.* **55**, 2712 (1989).
7. E.M. Wright, R.J. Hawkins, and R.J. Deri, *Proc. Integrated Photonics Research Conf.* **10**, 238, (1993).
8. R.W. Ziolkowski and J.B. Judkins, "Applications of Discrete Methods to Pulse Propagation in Nonlinear Media: Self-Focusing and Linear–Nonlinear Interfaces," *Proc. Radio Science for the 1992 URSI EM Theory Symposium*, to appear in 1994.
9. D.B. Davidson and R.W. Ziolkowski, *Body of Revolution Finite Difference Time Domain Modelling of Space–Time Focusing by a Three-Dimensional Lens*, University of Arizona; accepted for publication in *JOSA A* (1994).
10. R.W. Ziolkowski and D.B. Davidson, *Designer Pulsed-Beams for Enhanced Space–Time Focusing*, University of Arizona; accepted for publication in *Opt. Lett.* (1994).
11. R.W. Ziolkowski and J.B. Judkins, *NL-FDTD Modeling of Linear and Nonlinear Corrugated Waveguides*, University of Arizona; accepted for publication in *JOSA B Opt. Phys.* (1994).

Development of Millimeter-Wave Sources and Components for the Heating of Fusion Plasmas

Clifford C. Shang
Engineering Research Division
Electronics Engineering

K. Felch and R.V. Schumacher
Varian Associates
Palo Alto, California

Malcolm Caplan
Magnetic Fusion Energy
Energy Programs

H.-U. Nickel and M. Thumm
ITP Kernforschungszentrum
Karlsruhe, Germany

George W. Kamin
Isotope Separation and Advanced
Manufacturing Technology Program
Laser Programs

W.A. Urbanus and A.V. Tulupov
FOM Institute
The Netherlands

Teresa L. Swatloski
Applications Systems Division
Computation Directorate

We are examining high-average-power, millimeter-wave (mm-wave) sources and components as part of an ongoing technology collaboration between Lawrence Livermore National Laboratory and U. S. industry. We have modeled a number of corrugated and distributed mm-wave windows as well as auxiliary gyrotron components for high-average-power application. We also examined a new, multi-megawatt, free electron maser concept using a DC accelerator system with depressed collector, based on the 1-MW device now being built in The Netherlands. Our work looks at an extrapolation of the Dutch design to output microwave power of over 4 MW.

Introduction

In FY-93, we examined two sets of microwave problems as part of our Cooperative Research and Development Agreement (CRADA) work with Varian and our international collaboration in studying millimeter-waves (mm-waves). The first set of problems involves design of high-average-power mm-wave structures, and the second involves design of a free-electron maser (FEM) for the heating of fusion plasmas. The principle computational features in modeling the mm-wave structures for high-average-power involve the launching of modes, modeling of lossy dielectrics, and development of

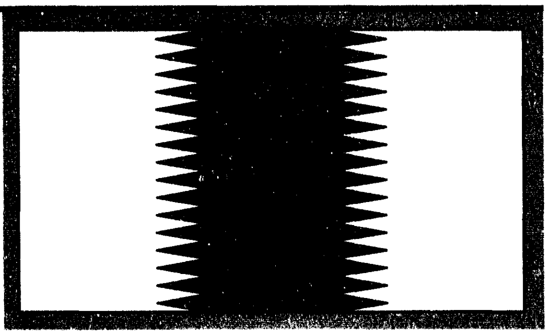

Figure 1. Configuration of concentric grooved window used in high-power broadband window study.

microwave diagnostics in detailed field codes. The modeling of the FEM required the development and application of a suite of fully three-dimensional (3-D), microwave particle-interaction codes.

The use of high-power microwaves to heat the plasma in a magnetic fusion energy (MFE) reactor at the electron-cyclotron resonance can yield a number of benefits, including bulk-heating and pre-ionization of the plasma, reaction startup, and instability suppression. The use of electron-cyclotron heating (ECH) in tokamak and stellerator reactors has been studied in many significant MFE experiments, including C-mod at the Massachusetts Institute of Technology and DIII-D at General Atomics in the U.S.; Compass at Culham, England; and T-10 at the Kurchatov Institute, Russia.

Operating parameters of interest for ECH applications include frequencies in the 140- to 250-GHz range and output power in the vicinity of 1 MW per bottle.[1] Currently, the fixed-frequency, mm-wave source available for use in the 1-MW range is the gyrotron. Understanding of the microwave properties of high-average-power radio frequency (RF) structures is crucial to the design of gyrotron tubes, amplifier devices, and future mm-wave sources. Dissipation of the RF (ohmic loss) and excessive mode conversion are often limiting

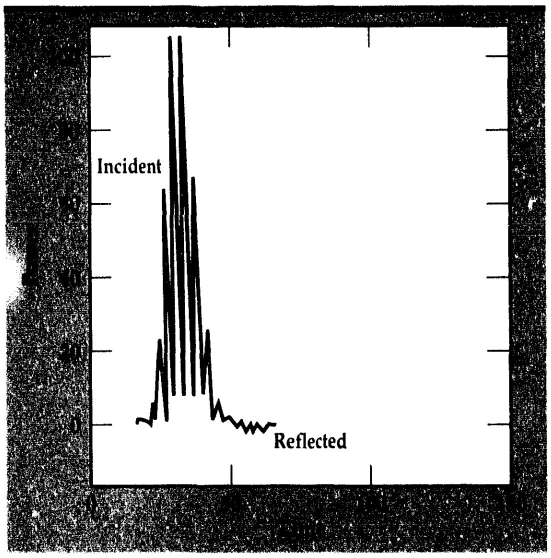

Figure 2. Experimental results of examination of field intensity distribution in corrugated window.

Table 1. Extrapolation of FOM Institute design to high power.

	FOM design	[illegible]
Net microwave output	1.3 MW	4.7 MW
Beam current	12 A	30 A
Beam voltage	1.75 MeV	1.75 MeV
Reflection coefficient	29%	20%
Beam emittance	50 π mm m rad	90 π mm m rad
Beam radius	.92 mm	1.28 mm
Overall total length	158 cm	138 cm
Number full periods Section 1	20	17
Number full periods Section 2	14	13
Peak wiggler field Section 1	2.0 kG	2.0 kG
Peak wiggler field Section 2	1.6 kG	1.5 kG
Inter-wiggler gap	6.0 cm	6.4 cm
Wiggler period	4 cm	4 cm
Waveguide mode	HE_{11} rectangular	HE_{11} rectangular
Waveguide height	2 cm	2 cm
Waveguide width	1.5 cm	1.5 cm
Frequency	200 GHz	200 GHz
High-voltage loss current	< 40 mA	< 100 mA
Depressed collector efficiency	90%	90%
Wall plug —>microwave efficiency	40%	50%

factors in the performance and robustness of the overall device. These issues and others pertaining to mm-wave devices may be investigated using a combination of particle interaction and field codes.

Progress

Electrical Analysis of Wideband and Distributed Windows, Using Time-Dependent Field Codes

Windows, which provide the barrier to maintain the vacuum envelope in a microwave tube, are critical components in high-average-power microwave sources, especially at millimeter wavelengths. As RF power levels approach the 100's of kW's to MW range, operating continuous wave (CW), the window assembly experiences severe thermal and mechanical stresses. Depending on the source, the bandwidth of the window may be less than 1 GHz for gyrotron oscillators or up to ~ 20 GHz for a tunable FEM. The bandwidth requirements give rise to a number of window configurations where the common goal is locally distributed heat dissipation. In order to better understand the transmission and RF properties of these microwave structures, we use detailed time-dependent field solvers.

A number of general time-dependent field solvers[2] have been developed to examine scattering and radiation problems. These algorithms work both on logically regular as well as completely unstructured meshes.[3] The basic algorithm has also been modified to handle both magnetically and electrically dispersive materials. Typically, the numerical solutions are second-order accurate and conditionally stable. The practical limitation to application of the codes involves the electrical size of the structure to be modeled. To obtain the proper spectral resolution requires approximately 10 mesh points per wavelength corresponding to the highest frequency of interest. In practice, large computational problems in 3 D (volumetric) are 10^3 λ^3 in size.

An ongoing window program at ITP Kernforschungszentrum Karlsruhe (KfK)[4] has been aimed at developing high-power, broadband windows. As a start in analyzing wideband windows, we examined the corrugations in the concentric grooved window (**Fig. 1**).

The groove period is $\lambda/3$, and the groove depth is λ, referenced to 140 GHz in free space. We launched a TE_{01} (parallel to grooves), 140-GHz (± 20% in frequency) pulse, and obtained 98% transmission in power (**Fig. 2**). Adiabatic theory was inadequate to account for the physics of the match. In fact, the corrugations were found to focus the RF into the bulk dielectric portion of the window. As can be seen in **Fig. 3**, the peak field intensity occurs in the corrugations.

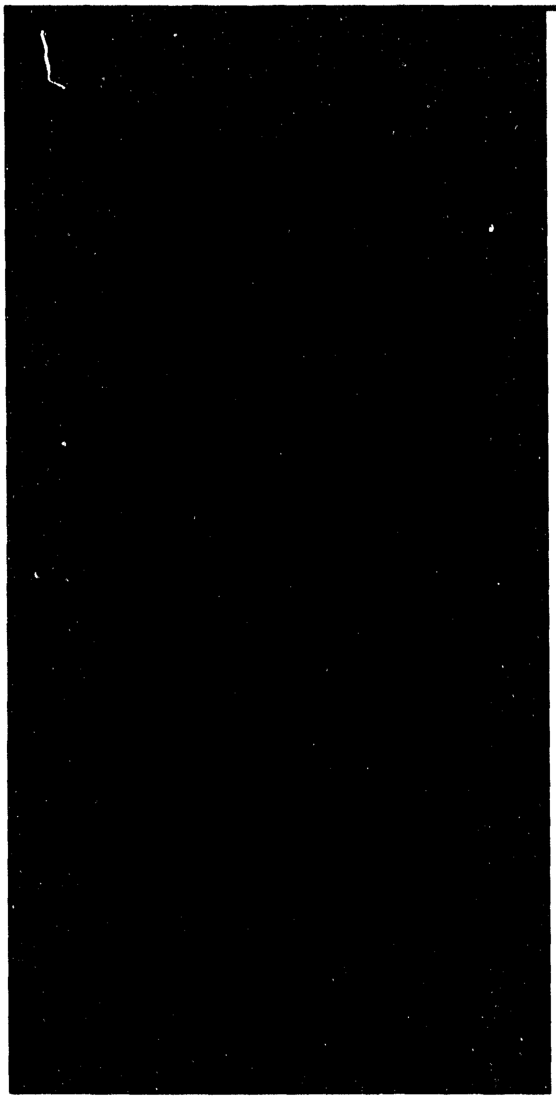

Figure 3. Photo showing peak field intensity in window corrugations.

A cropped, adiabatically tapered window with rectangular corrugations was fabricated and cold-tested at KfK. The calculational results were obtained by exciting single-frequency, x-directed electric currents, thereby launching waves with E perpendicular to the grooves. A key aspect for maintaining stability of the solution was to place the radiating boundary conditions several wavelengths from the scattering surface of the window. This prevented interaction of the evanescent fields with the numerical radiation boundary condition. For diagnostic purposes, at appropriate transverse slices on the vacuum side of the window, E × H is computed and summed over the surface directly from the field values. Good qualitative agreement between cold-test reflection data and calculational

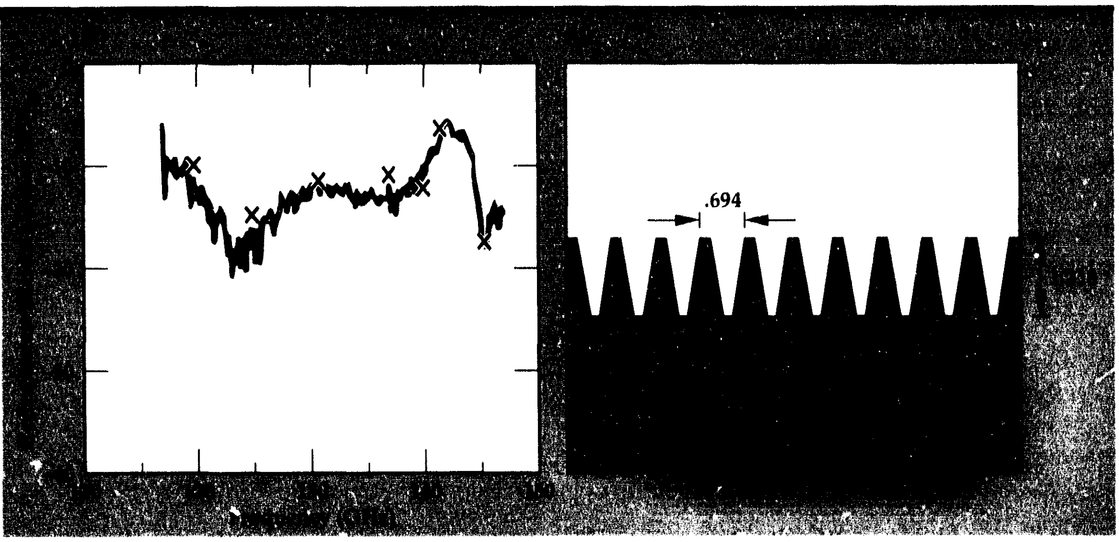

Figure 4. ***(a) Comparison of cold-test reflection data and calculations for cropped, adiabatically tapered window with rectangular corrugation, shown in (b).***

results can be found in **Fig. 4**. It should be possible to obtain closer correlation to measurements through better geometrical fidelity in the model (e.g., finer, unstructured mesh). However, examination of solution convergence awaits further analysis.

We are currently examining the electrical effects of introducing distributed cooling, which are discontinuous material junctions to window structures. The materials include metal cooling strips or dielectric cooling channels. By exploiting the focusing effect of the corrugations, we hope to obtain good transmission characteristics over a wide bandwidth.

Future work will involve coupling the RF calculations for the window assembly to thermal-mechanical codes.

Design of a Tunable, 4-MW FEM for Heating Fusion Plasmas

The FOM Institute for Plasma Physics, The Netherlands, is now constructing a FEM to be used as a high-frequency, tunable microwave source for heating fusion plasmas.[1] This source has been designed to ultimately operate CW at the 1-MW power level over an adjustable tuning range of 150 to 250 GHz. The design philosophy is to use a high voltage DC beam system with depressed collector in order to make the overall wall plug efficiency 40 to 50%. The high voltage, 1.75-MV power supply provides only loss current (~ 30 mA), while the 12-ampere-beam current is supplied by the 100- to 200-kV collector supplies.

A compatible microwave interaction circuit, coupling system, and wiggler magnet are shown in **Fig. 5**. The rectangular corrugated circuit operating in HE_{11} mode is very low loss, capable of handling multi-megawatts of power CW. The stepped waveguide system allows feedback and output coupling in highly overmoded guide while maintaining mode purity. The two-stage, stepped undulator allows for increased electronic efficiency while maintaining high-quality focusing.

There is an interest in extrapolating this design to higher powers in order to reduce cost/kW and

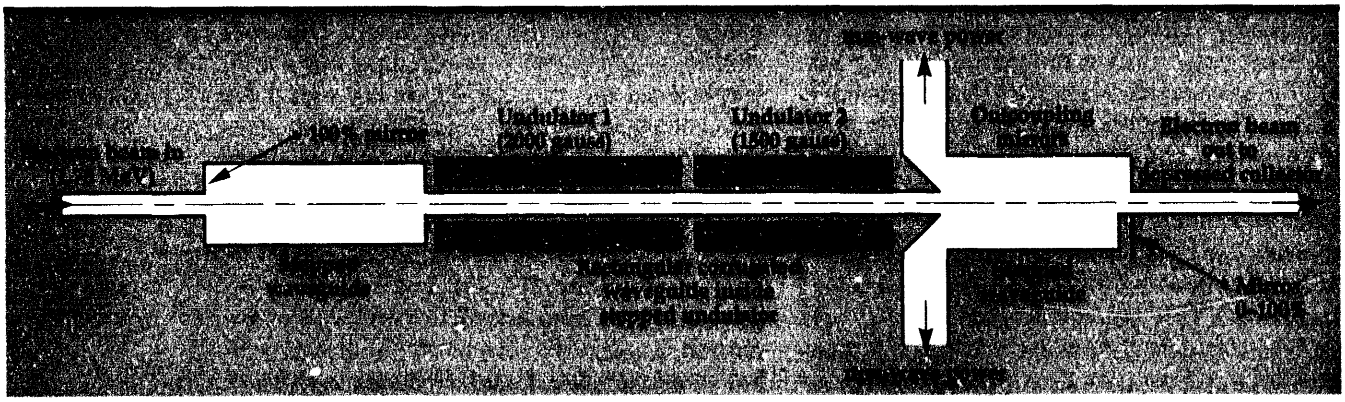

Figure 5. Schematic of mm-wave cavity of FEM showing step undulator, waveguide, and reflection/outcoupling system.

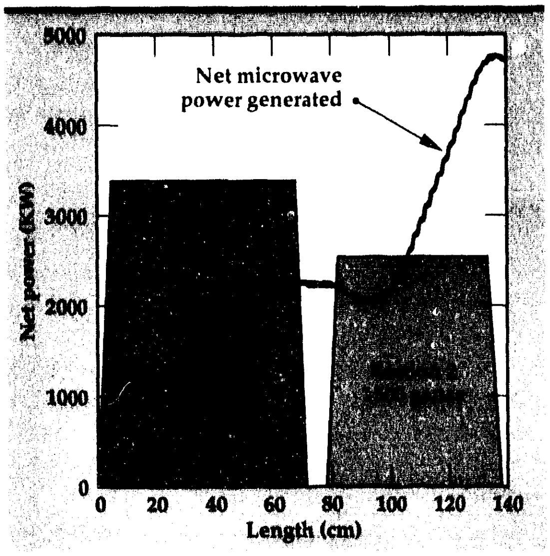

Figure 6. Net microwave power generated vs length for extrapolated high-power design of Table 1.

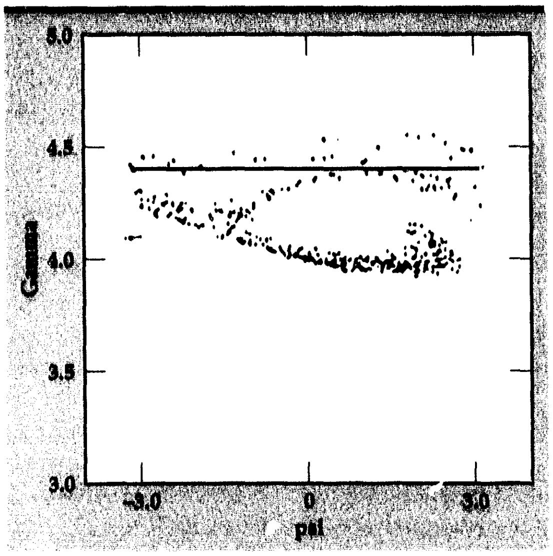

Figure 8. Energy of spent beam electrons vs phase angle at undulator exit.

develop a more compact microwave system. The most straightforward extrapolation is to increase the current beyond 12 A to as high as 30 A, keeping as much of the system the same while ensuring the integrity of the beam focusing system and CW power-handling capability. The cost savings occur because the high voltage supply costs scale slowly with current.

High Current Design

It is important to note that the 1.75-MeV beam radius is determined by emittance in the 10- to 30-A range. As the beam current increases, the emittance is also allowed to increase to keep the electron charge density from increasing excessively, causing problems due to ionization of background gas. The increase in emittance does not degrade microwave interaction efficiency.

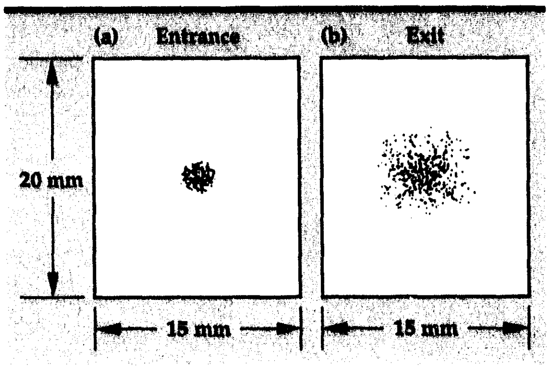

Figure 7. Transverse beam cross-section in rectangular waveguide at (a) entrance to undulator 1 and (b) exit side of undulator 2 after traversing step wiggler interaction region.

The original focusing system for the low-current design can be used even though the beam radius increases by 40%, the charge density increases by 20%, and the emittance increases by a factor of 2. The low-loss, corrugated waveguide system can handle up to 5 MW with acceptable wall loading (< 1 kW/cm^2). Further stepping of the waveguide would be required to separate the power into four output ports, each having corrugated, distributed cooled windows capable of handling ~ 1 MW CW power. **Table 1** gives the modified 30-A design from the original FOM Institute design. It is still assumed that the depressed collector can recover 90% of the power and dissipate up to 5 MW.

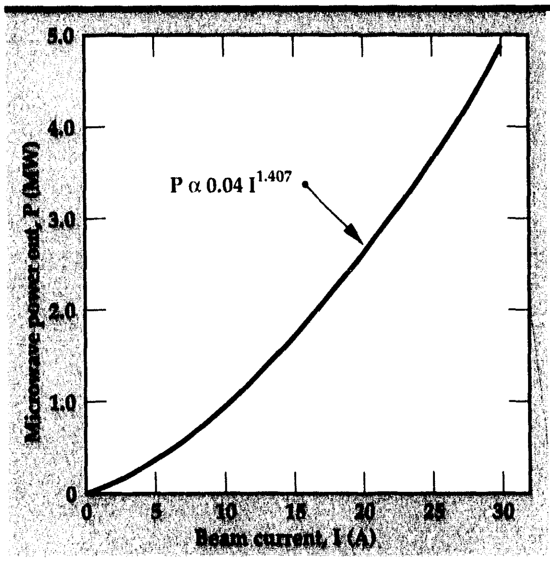

Figure 9. The dependence of microwave output power on beam current for the configuration.

Computer Simulations

The FEM performance was modeled using a fully self-consistent, 3-D, multimode, non-wiggle-averaged computer code including both AC and DC space charge effects.[5] The beam is modeled as several hundred particles, which represent a specified emittance and charge density profile.

Figure 6 shows the prediction for steady-state microwave power generated as the beam traverses the wiggler. The step-tapered wiggler enhances the power by almost 2.5 times to 4.7 MW. **Figure 7** shows the beam cross-section upon entering and exiting the wiggler. The beam increases by a factor of 2 in size, but is still well confined and centered in the original waveguide structure. **Figure 8** shows the electron energy spread in the beam after the interaction to be about 300 kV.

A number of simulations were performed at various currents between 12 and 30 A with each case being optimized. The results are shown in **Fig. 9**, which shows how the output power scales with beam current, assuming a roughly fixed charge density. The dependency goes faster than linear ($\sim I^{1.4}$) because the FEM efficiency increases with current.

The 1-MW FEM now being built at FOM Institute could be upgraded to as much as 4.7 MW by increasing the current from 12 to 30 A, still keeping many of the key design parameters constant (voltage, wiggler configuration, microwave waveguide system).

The upgraded design can greatly reduce the cost/kW since power supply costs at fixed voltage increase slowly with current (much less than linear), while output power increases faster than linear with current.

Conclusions

The electrical properties of mm-wave sources and structures were studied using a combination of particle-interaction and detailed electromagnetic field codes. Good agreement between calculation and experiment for component design has been obtained. Microwave simulations suggest that future RF heating systems, which are economically attractive, may be developed for the electron-cyclotron heating of fusion plasmas.

Future Work

In FY-94, newly approved CRADAs with General Atomics and Hughes will substantially broaden the scope of Lawrence Livermore National Laboratory's (LLNL) technology collaboration with industry in the microwave field, both in the mm-wave and radar frequency regimes.

Acknowledgements

We acknowledge several helpful discussions with Prof. J. Feinstein of Stanford University and Dr. C. Moeller of General Atomics. The kindly advice and support of Drs. J. DeFord and J. Sefcik of LLNL are gratefully acknowledged.

---

1. A.G. Verhoeven *et al.*, "The 1-MW, 200-GHz FOM-FUSION FEM," *Proc. 17th Int. Conf. Infrared and Millimeter Waves* (Pasadena, California), (December 1992).
2. K.S. Yee, *IEEE Trans. Antennas Propag.* **AP-14**, 302 (May 1966).
3. N.K. Madsen, *A Divergence Preserving Discrete Surface Integration Algorthm for Time-Domain Maxwell Equations*, Lawrence Livermore National Laboratory, Livermore, California, UCRL-JC-109787 (1993).
4. H.-U. Nickel, H. Massler, and M. Thumm, "Development of Broadband Vacuum Windows for High-Power Millimeter Wave Systems," *Proc. 18th Int. Conf. Infrared and Millimeter Waves* (Colchester, U.K.), (September 1993).
5. M. Caplan, "Predicted Performance of a DC Beam-Driven FEM Oscillator Designed for Fusion Applications at 200–250 GHz," *Proc. 14th Int. Free Electron Laser Conf.* (Kobe, Japan), (1992).

Electromagnetic Modeling for Target-Rich Embedded Environments

Scott D. Nelson
Defense Sciences Engineering Division
Electronics Engineering

We have investigated the feasibility of designing an electromagnetic (EM) radar system to examine the internal structure of concrete used in the highway industry. The central project involved the coordination of EM modeling, imaging, code design, and experimentation to design the hardware/software system. Modeling generated data for the imaging effort and allowed the precise control of individual parameters in the model. Modeling also supported the experimental effort to investigate the antennas and pulsers used in the experimental hardware setup.

The modeling effort consisted of three phases: (1) verification of embedded scattering models to uniform theoretical predictions; (2) analysis of the physical antenna models used in the project and modeling effort; and (3) three-dimensional modeling of the configuration with typical targets, antennas, and concrete.

Introduction

The Ground Penetrating Imaging Radar Project was funded by the Lawrence Livermore National Laboratory (LLNL) Power Conversion Technologies (formerly Microwave and Pulsed Power) Thrust Area, in conjunction with the Computational Electronics and Electromagnetics Thrust Area and other LLNL thrust areas. Our goal is to investigate the feasibility of designing an electromagnetic (EM) radar system to examine the internal structure of concrete structures typically found in the highway industry.[1] This work was a continuation of previous efforts.[2]

Progress

Two-dimensional (2-D) and three-dimensional (3-D) models were constructed, and the results were compared with experimental and uniform theoretical data. The 2-D data served as a preliminary test of the scattering models used in the AMOS 2.5D finite difference time domain (FDTD) EM modeling code.[3] The 3-D data generated via the TSAR 3-D FDTD code[4] served as a verification for the concrete model and was the method used to create time domain waveforms for the image-reconstruction algorithms.[5] The use of the 3-D code was necessary to simulate the complex target geometries used in the problem.

The 2-D comparison showed good agreement between the theoretical results for a uniform random distribution of targets and the FDTD model for the effective material parameters. The 3-D results compare the time domain waveform seen from a physical antenna modeled using TSAR, to an experiment performed with the antenna. Later models then introduced the concrete and targets into the problem. These results were then used to generate data for the imaging teams and to resolve hardware issues.

2-D Clutter Verification Effort

A region composed of scattering bodies was modeled in two dimensions (2D) to compare with the Maxwell–Garnet mixing formula,[6] which is the real part of **Eq. 1**, which shows the relationship between packing density and effective permittivity.

$$\varepsilon_{eff}=\varepsilon\left\{\left[\frac{1+\dfrac{2f_s(\varepsilon_s-\varepsilon)}{\varepsilon_s+2\varepsilon}}{1-\dfrac{f_s(\varepsilon_s-\varepsilon)}{\varepsilon_s+2\varepsilon}}\right]+i2f_s k^3 a^3\left|\frac{\dfrac{\varepsilon_s-\varepsilon}{\varepsilon_s+2\varepsilon}}{1-\dfrac{f_s(\varepsilon_s-\varepsilon)}{\varepsilon_s+2\varepsilon}}\right|^2\right\}$$

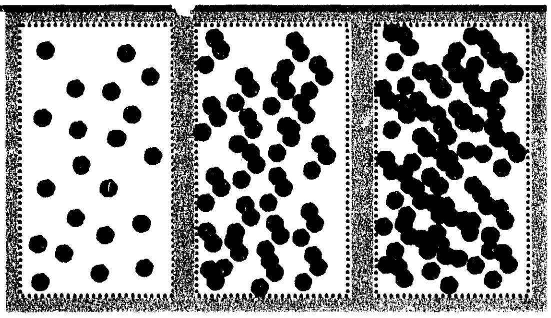

Figure 1. Packing geometries (10%, 30%, 50% density) used to check the model results against the mixing formula.

Figure 2. Antenna model used for the experimental comparison. This antenna is a dual-ridged broadband, conical horn antenna manufactured by EM Systems.

The form used in the comparison is shown in **Eq. 2**, which shows the relationship between packing density and effective real permittivity.

$$\varepsilon_{e_r}\varepsilon_0 = \frac{\varepsilon_s\varepsilon_0+\varepsilon_0^2+f\varepsilon_s\varepsilon_0-f\varepsilon_0^2}{\varepsilon_s+\varepsilon_0-f\varepsilon_s+f\varepsilon_0}$$

The region of interest is composed of uniform objects with 10%, 30%, and 50% fill ratios using dielectric aggregate. These values were chosen to span those aggregate concentrations seen in concrete. In this case, the targets in the model space were oriented in the same direction and are randomly separated, as shown in **Fig. 1**.

This simulation addressed the issue of target/clutter size vs pulse width. The general rule seen earlier[2] is that targets less than one third pulse width in diameter will scatter expanding spherical waves. In performing this series of 2-D verification measurements, this relationship was again confirmed. It was also necessary to justify the packing density checkerboard that was used in the model, since a finite difference square-cell grid was used. In this case, a random cellular offset was used to avoid resonances formed by rows of equally spaced clutter.

3-D Antenna Model Comparison

A broadband antenna model was used to compare the model with experiments. The antenna used in the model is shown in **Fig. 2** and is a model of the broadband, double-ridged, conical antenna. The time domain waveforms for the modeled case are shown in **Fig. 3** for the 3-D physical antenna.

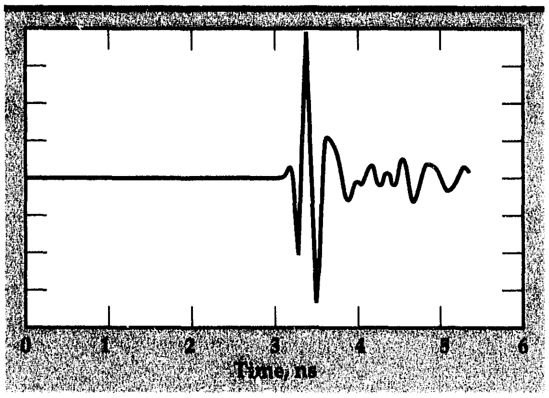

Figure 3. Time domain waveform for the 3-D, broadband antenna model, showing the fast leading pulse and the later time ringing.

3-D Concrete Simulation

Preliminary work for 3-D geometries used a complex target and clutter geometry but a simple, rectangular, horn antenna model. **Figure 4** shows a model using several rebar targets, three rebaric voids, a spherical air void, and a simple antenna. This waveguide-fed rectangular horn antenna exhibited the fundamental ringing seen when sending a pulse through a short section of waveguide and was expected. The pulse width was made short enough so that the ringing wave packet did not interfere with the imaging reconstruction.

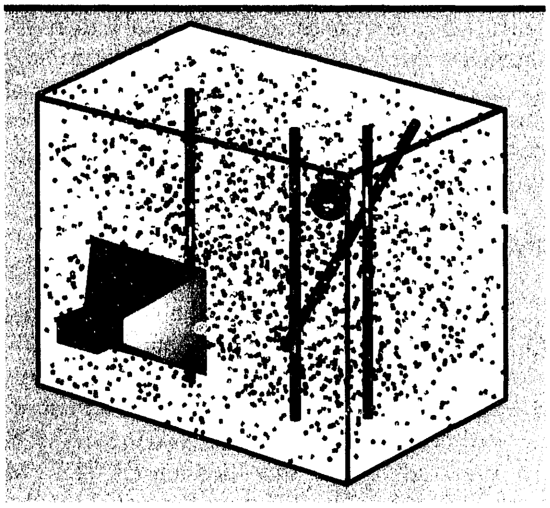

Figure 4. 3-D geometry used for the concrete simulation, containing randomly oriented cubical aggregate particles. The block has been outlined for illustration purposes, and the cement removed to reveal the aggregate, rebars, rebaric voids, and spherical void.

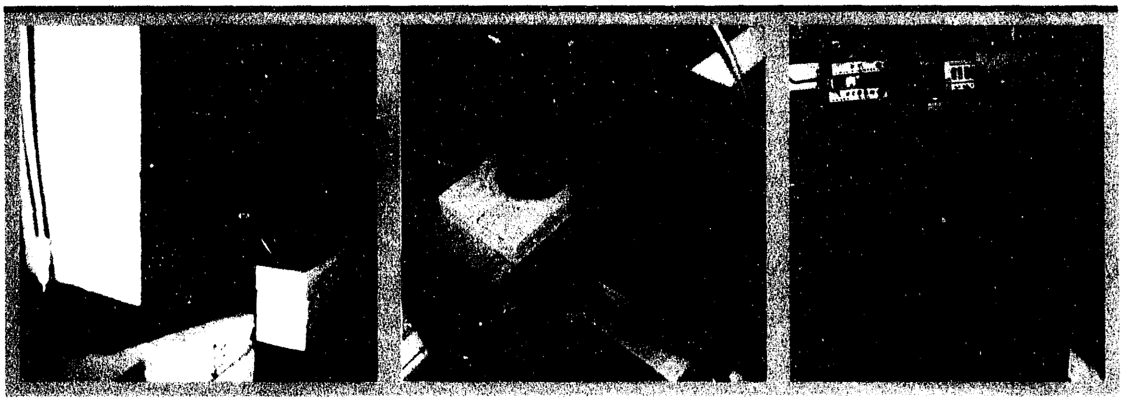

Figure 5. Experimental geometry showing concrete test slab, conical horn antennas, sampling scope, and absorbing material. The antennas were positioned so that the energy passed through the concrete; the antenna mutual coupling was included in the model.

The experimental geometry used for the 3-D concrete slab verification is shown in **Fig. 5**. In this case, broadband horn antennas were used to maintain the pulse fidelity. The energy was launched through the concrete test slab by the transmit antenna, which was connected to the pulse generator. The receiver antenna was located 90° from the transmitter. The pulser generated a 400-ps, full-width half-maximum (FWHM) Gaussian-like pulse, which the antenna then differentiated. The receiving antenna was connected to a Tektronix CSA803 sampling scope to digitize the waveform (see **Fig. 5**).

Conclusions

The aggregate clutter models used were not sensitive to orientation as long as they remained small with respect to the pulse width. For clutter targets larger than one third of a pulse width in diameter, the clutter orientation is important. The packing density of the clutter for dense aggregate masses is more critical in 2D than three dimensions (3D) since diffraction occurs more readily in 3D (since there are more diffraction surfaces). For dense clutter geometries, it is more important in 2D to ensure that there is a propagation path through the model when using high permittivity or high conductivity materials. These issues all trace back to the 2-D nature of the models, with circles being cylinders.

The 3-D models used contained targets and randomly oriented clutter; the 3-D mesher converted the randomly oriented cubes into multifaceted blobs. The imaging algorithms used[7] were able to image the aggregate clutter particles, and so a homogeneous effective material model was ***not suitable***. This was confirmed using the experimental data.[5] This sizing limit greatly increased the time required to do the 3-D modeling, but is very important because the aggregate clutter is diffractive, scattering, shadowing, and is imaged as discrete particles.

Future Work

A realistic ground model is being planned to include stratified earth, moisture gradations, rock effects, and possible bush root effects in 3D as part of other projects. Some of these effects have already been demonstrated in 2D. Also, the variability of the scanning antenna's flight path is being examined from a statistical standpoint to better simulate real world data-acquisition conditions. This will allow for positional errors in the locatability of the antennas. All previous experiments used precisely known antenna locations.

Acknowledgements

Thanks go to E. Johansson and J. Mast for their input on the imaging effort; to J. Warhus for his hardware and pulser design efforts; to J. DeFord for his advice and timely modifications to the AMOS code; and to S. Pennock for his assistance with TSAR.

---

1. J.P. Warhus, S.D. Nelson, J.M. Hernandez, E.M. Johansson, H. Lee, and B. Douglass, "Ground Penetrating Imaging Radar for Bridge Inspection," *Engineering Research, Development, and Technology*, Lawrence Livermore National Laboratory, Livermore, California, UCRL-53868-92, 7-5 (1992).
2. S.D. Nelson and C.A. Avalle, "Electromagnetic Modeling and Experiments for Dispersive Media," *Engineering Research, Development, and*

Technology, Lawrence Livermore National Laboratory, Livermore, California, UCRL-53868-92, 1-21 (1992).

3. J. DeFord, G. Kamin, and G.D. Craig, *Development and Applications of Dispersive Soft Ferrite Models for Time-Domain Simulation,* Lawrence Livermore National Laboratory, Livermore, California, UCRL-JC- 109495 (1992).
4. R.R. McLeod, *Temporal Scattering and Response Software,* Lawrence Livermore National Laboratory, Livermore, California, UCRL-MA-104861 (1992).
5. E.M. Johansson and J. Mast, "Time- and Frequency-Based 3-D Image Reconstruction," *Proc. 2nd Government Workshop on GPR, Advanced Ground Penetrating Radar Technologies and Applications,* (Ohio State University), (October 26–28, 1993).
6. J.A. Kong, *Electromagnetic Wave Theory,* John Wiley Sons (ISBN 0-471-52214-7), 1990.
7. J.E. Mast, *Microwave Pulse-Echo Radar Imaging for the Nondestructive Evaluation of Civil Structures,* Ph.D. Thesis, University of Illinois at Urbana–Champaign (1993).

Acoustic, Finite-Difference, Time-Domain Technique Development

Karl Kunz
Engineering Research Division
Electronics Engineering

A close analog exists between the behavior of sound waves in an ideal gas and the radiated waves of electromagnetics. This analog has been exploited to obtain an acoustic, finite-difference, time-domain (AFDTD) technique capable of treating small signal vibrations in elastic media, such as air, water, and metal, with the important feature of bending motion included in the behavior of the metal. This bending motion is particularly important when the metal is formed into sheets or plates. Bending motion does not have an analog in electromagnetics, but can be readily appended to the acoustic treatment since it appears as a single additional term in the force equation for plate motion, which is otherwise analogous to the electromagnetic wave equation.

The AFDTD technique has been implemented in a code architecture that duplicates the electromagnetic, finite-difference, time-domain technique code. The main difference in the implementation is the form of the first-order coupled differential equations obtained from the wave equation. The gradient of pressure and divergence of velocity appear in these equations in the place of curls of the electric and magnetic fields. Other small changes exist as well, but the codes are essentially interchangeable. The pre- and post-processing for model construction and response-data evaluation of the electromagnetic code, in the form of the TSAR code at Lawrence Livermore National Laboratory, can be used for the acoustic version.

Comparison of the AFDTD code with analytically derived results for compression waves has shown agreement to within 1 dB or less. The more difficult comparison, with experimentally obtained response data from a submerged plate with bending modes, has also been made.

The initial comparison of experiment and code with respect to bending modes is very encouraging. A variety of applications is possible, pending validation of the bending phenomenon. The applications include acoustic-radiation-pattern predictions for a submerged object; mine detection analysis; structural noise analysis for cars; acoustic barrier analysis; and symphonic hall/auditorium predictions and speaker enclosure modeling.

Computational Mechanics

The Computational Mechanics thrust area sponsors investigation into the underlying solid, structural, and fluid mechanics and heat transfer necessary for the development of state-of-the-art general purpose computational software. The scale of computational capability spans office workstations, departmental computer servers, and Cray-class supercomputers. Our codes have achieved world fame through our broad collaborators program, in addition to their strong support of ongoing Lawrence Livermore National Laboratory (LLNL) programs. Several technology transfer initiatives have been based on these established codes, teaming LLNL analysts and researchers with counterparts in industry, extending code capability to specific industrial interests of casting, metalforming, and automobile crash dynamics. The next-generation solid/structural mechanics code, ParaDyn, is targeted toward massively parallel computers, which will extend performance from gigaflop to teraflop power.

Our work for FY-93 is described in the following nine articles: (1) Seismic Response Computations for a Long Span Bridge; (2) ParaDyn: New Generation Solid/-Structural Mechanics Codes for Massively Parallel Processors; (3) Microchannel Boiling Mechanisms Leading to Burnout; (4) Rigid-Body Mechanics for Finite Element Analysis; (5) HYDRA: A Flow Solver for Transient Incompressible Viscous Flow; (6) Automatic Contact in DYNA3D for Vehicle Crashworthiness; (7) TRIM3D: A Three-Dimensional Radiation Heat Transfer Code for Participating Media; (8) GRIZ: Visualization of Finite Element Analysis Results on Unstructured Grids; and (9) PING: An Explicit Finite Element Code for Linear Structural Acoustics.

Gerald L. Goudreau
Thrust Area Leader

Section 2

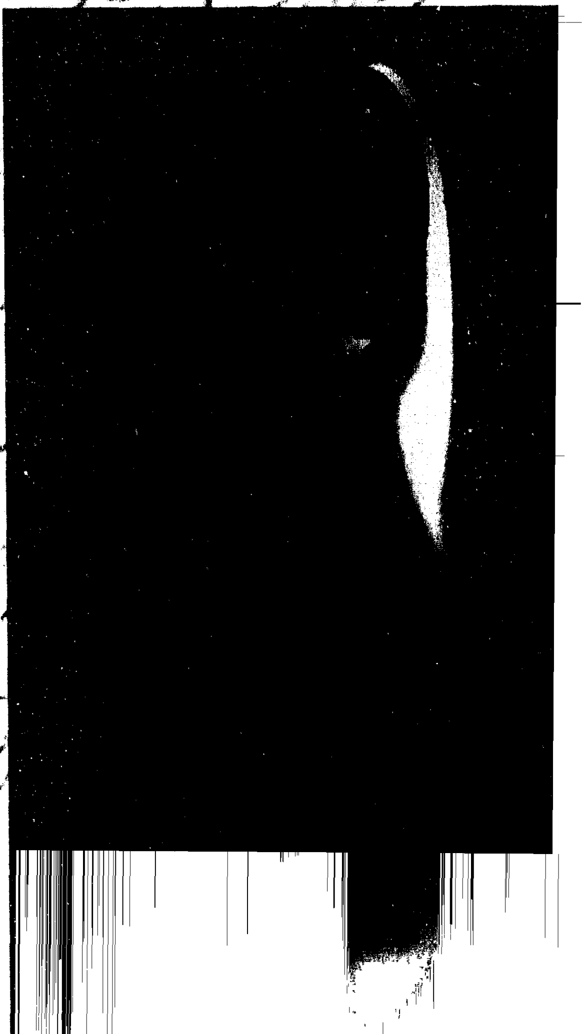

2. Computational Mechanics

Seismic Response Computations for a Long Span Bridge

David B. McCallen
Structural Mechanics Group
Mechanical Engineering

We are performing large-scale numerical computations to simulate the earthquake response of a major long-span bridge that crosses the San Francisco Bay. The overall objective of the study is to estimate the response of the bridge to potential large-magnitude earthquakes generated on the nearby San Andreas and Hayward earthquake faults. Generation of a realistic model of the bridge system is complicated by the existence of large pile group foundations that extend deep into soft, saturated clay soils, and by the numerous expansion joints that segment the overall bridge structure. In the current study, advanced, nonlinear, finite element technology is being applied to rigorously model the detailed behavior of the bridge system and to shed light on the influence of the foundations and joints of the bridge.

Introduction

The 1989 Loma Prieta earthquake in northern California served as a reminder of the vulnerability of many manmade structures to extreme natural events. The loss of life and the economic impact were remarkably acute given the fact that the earthquake epicenter was nearly 70 km away from the major metropolitan areas surrounding the San Francisco Bay. The strong ground motion generated by the Loma Prieta earthquake was only on the order of ten-s duration, and it has been estimated that a few more seconds of shaking would have resulted in the collapse or severe damage of many more structures.[1] The prestigious Governor's Board of Inquiry on the 1989 Loma Prieta earthquake concluded:

> "Future earthquakes in California are inevitable. Earthquakes larger than Loma Prieta with more intense ground shaking will occur in urban areas and have severe consequences—too large to continue 'business as usual'... The Loma Prieta earthquake should be considered a clear and powerful warning to the people of California... More aggressive efforts to mitigate the consequences of earthquakes are needed if their disastrous potential is to be minimized and one of the fundamental of responsibilities of government is fulfilled—to provide for public safety."[1]

With established programs in the earth sciences and engineering, with extensive capabilities in seismological instrumentation, numerical computations, and supercomputing, and with experience as the lead Department of Energy Laboratory for natural phenomenon hazards design-criteria development, Lawrence Livermore National Laboratory (LLNL) is well positioned to assist in the mitigation of seismic hazards for the State of California and for the U.S. in general.

Over the past two years, LLNL engineers and seismologists have been cooperating on a Laboratory Directed Research and Development project, the Computational Earthquake Initiative, aimed at developing a methodology chain for predicting the response of critical structures to large earthquakes.[2] This project is bringing together 'off-the-shelf' numerical analysis technology developed for defense programs, state-of-the-art numerical techniques from Princeton University, and recent seismological research developments from LLNL to develop a fault-to-structure procedure for estimating the response of a critical structure to a specific earthquake emanating from a given fault. The procedure under development (see **Fig. 1**) consists of three distinct parts: (1) seismologists generate a suite of synthetic earthquake time histories for bedrock below the site, using small level earthquake measurements (i.e., M = 2–3) as empirical Green's functions[3,4]; (2) geotechnical engineers analyze the dynamic site response of overlying soil, including the effects of pore water pressure variation and potential liquefaction, in order to estimate the motion of the ground surface at the site based on a methodology and computer program developed at Princeton University[5]; and (3) structural engineers use LLNL nonlinear finite element software to perform nonlinear analysis of the structure, or the structure and a surrounding soil island, to estimate the response of the structural system.[6,7]

Investigation and evaluation of the methodologies under development are being carried out with case studies of two transportation structures. One structure is the Painter Street Bridge in Rio Dell, California; the second structure is the Dumbarton Bridge at San Francisco Bay in California. The Dumbarton Bridge study is being performed at the suggestion of the California Department of Transportation (Caltrans). The structural analysis computations for the Dumbarton Bridge, the southern most crossing of the San Francisco Bay (see **Fig. 2**), are the focus of the current report.

This bridge is approximately one and one third miles long and consists of three distinct structures: two approaches constructed with reinforced concrete girders supporting the roadway, and a main span constructed with steel box girders supporting the roadway. The bridge structure contains eight thermal expansion joints located at various piers, and two expansion joints located at deck hinges in two of the bridge spans. The two approach structures are supported by large pile groups extending through approximately 50 ft of San Francisco Bay mud. The center span is also supported on pile group foundations, which extend free standing through approximately 50 ft of water into a more competent layer of soil beneath the bay.

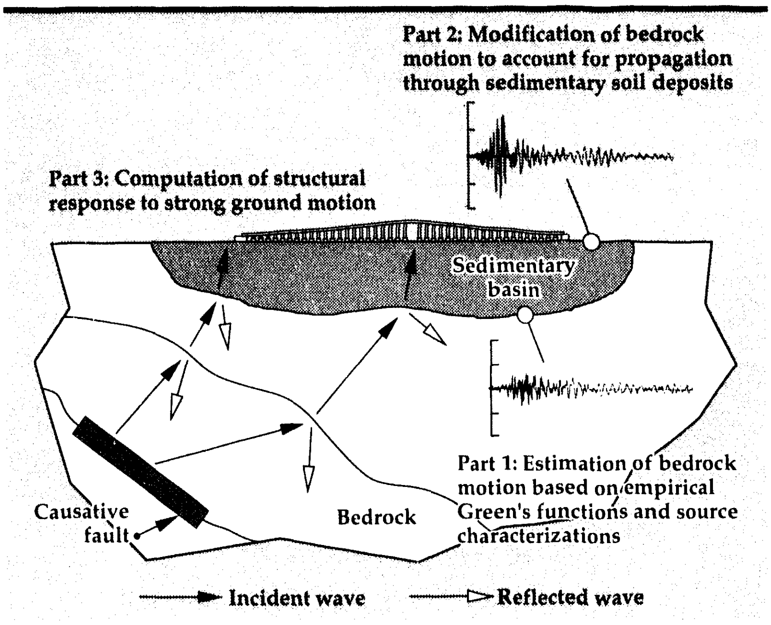

Figure 1. Steps in the methodology chain for prediction of the response of a structure to a large earthquake generated on a specified fault.

The Dumbarton Bridge was the only major San Francisco Bay crossing that was instrumented at the time of the 1989 Loma Prieta earthquake. The California Department of Mines and Geology's strong-motion instrumentation array measured the response of this bridge during the earthquake. The instrumentation layout on the bridge is shown in **Fig. 3**. The 'free-field' recording, which is intended to be representative of the surface ground motion at the bridge site, was measured at the ground surface approximately three quarters of a mile west of the bridge. A second free-field recording station east of the bridge failed to trigger properly during the Loma Prieta earthquake, thus the west free-field instrument is the only record of ground motion for the bridge. The instrumentation array on the bridge is relatively sparse, and the free-field instrument is distant enough from the bridge that it is impossible to ascertain if it is precisely representative of free-field motion for this bridge. Nonetheless, earthquake data records for a bridge this size are very few in number, and any data available for model validation and interpretation of structural response is somewhat of a luxury.

Global Bridge Model

Modeling the entire structure of a bridge the size of the Dumbarton is a time-intensive, difficult task. The definition of the bridge geometry alone can require weeks. The bridge geometry, defined by the design curvature and elevations of the roadway, and the bridge member sections and structural details, are all defined in a set of detailed drawings consisting of approximately 150 pages. For the current study, a solid model database was constructed for the entire bridge structure to allow visualization of the structural system from any vantage point. The solid model database was then used to generate a geometry file of critical wire-frame geometry points as definition of the structural elements for a finite element mesh-generation routine. The solid geometry model of the bridge is shown in **Fig 4**. The establishment of the solid model database allows easy generation of the global bridge model, and it also allows generation of a detailed-substructure model for any critical part of the bridge.

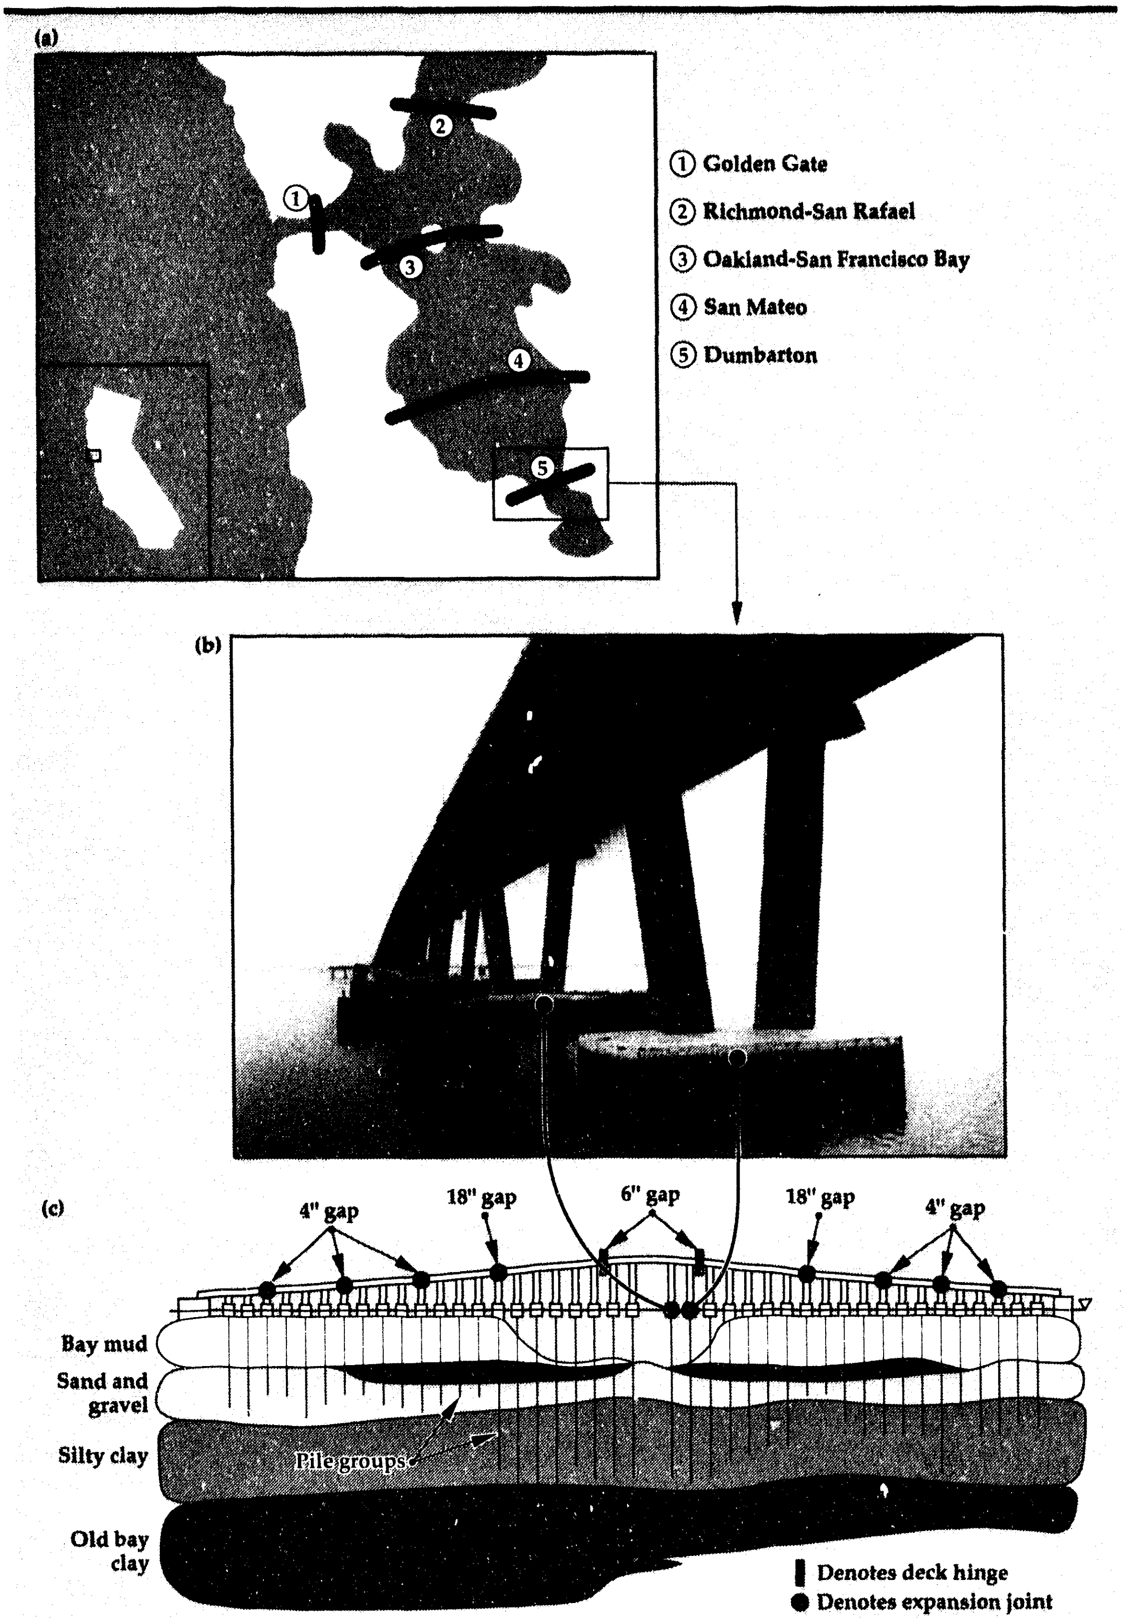

Figure 2. The Dumbarton Bridge: (a) location of major bay crossings, (b) Dumbarton Bridge, (c) profile of bridge and sedimentary deposits.

The global bridge model was constructed for the nonlinear finite element program NIKE3D.[8] The NIKE3D model of the bridge consists of 2272 shell elements for the deck, 3329 beam elements to represent the pier columns and deck girders, and 336 solid elements for the pier footings. The model contains 25,364 active degrees of freedom.

To represent the compliance of the individual pile foundations, foundation stiffness matrices were generated for selected representative piers. To develop the stiffness matrices of the large pile groups, unit loads were applied to a detailed three-dimensional substructure model of the foundation system to generate a six-by-six flexibility matrix of the pile/soil system (**Fig. 5**). Inversion of the flexibility matrix provides a six-by-six stiffness matrix characterization of the foundation, which is placed at the base of each pier in the global finite element model as indicated by **Fig. 5**.

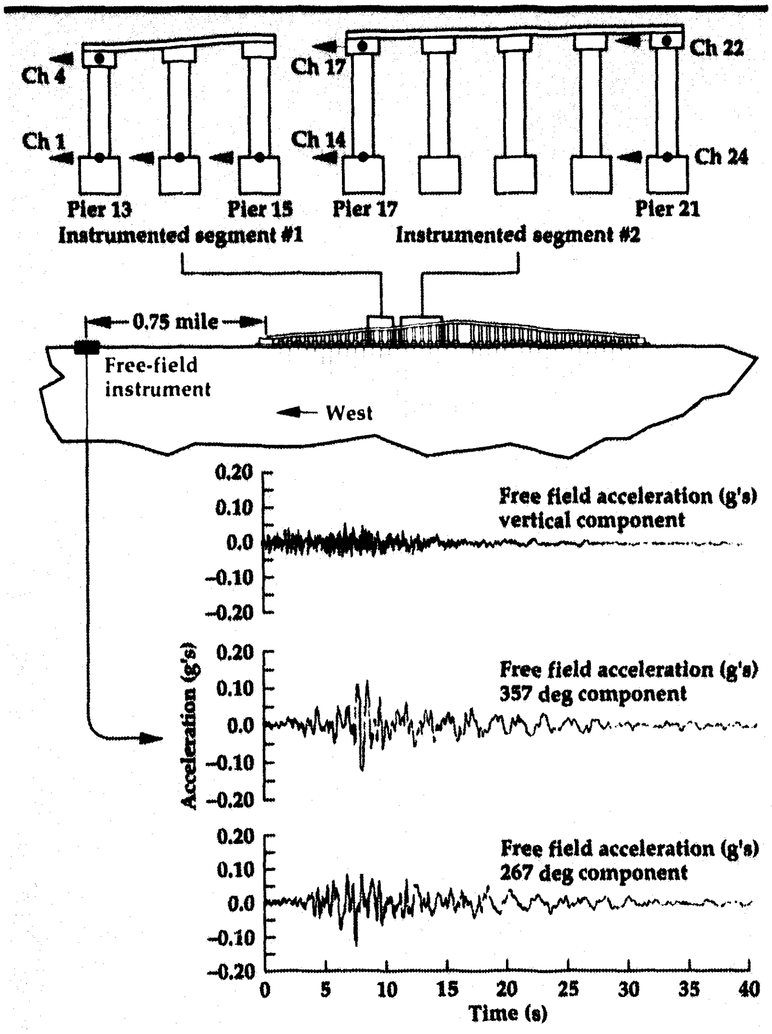

Figure 3. California Department of Mines and Geology instrumentation array on the Dumbarton Bridge (longitudinal channels shown) and Loma Prieta free-field records.

To represent the amplitude-dependent properties of the soil in the substructure models, equivalent linear soil properties were chosen, based on the soil strain levels observed in a University of California Berkeley (UCB) study[9] of the Dumbarton Bridge. In the UCB study, a one-dimensional wave propagation solution was performed for the Dumbarton site soil.

In actuality, the seismic input motion for a bridge like the Dumbarton will vary somewhat from pier to pier across the structure. The time history of pier motions differs across the structure as a result of wave propagation, differential wave velocities in near-site rocks and soils, and wave emission delays at the fault rupture.[10] For the bridge response predictions for the Loma Prieta earthquake, the free-field motion measured west of the bridge was assumed as input across the entire bridge. Lacking any more detailed information about the seismic input motion, this simple idealization of the seismic motion was assumed as a first approximation to the motion of the bridge resulting from the earthquake. As discussed in a subsequent section, in future work predicting the response of the bridge to an M = 8+ earthquake, the variation of input motion across the bridge will be explicitly modeled.

Despite the fact that the free-field accelerations for the Loma Prieta earthquake were only on the order of 0.15 g (see **Fig. 3**), there was clear evidence of impact occurring between adjacent segments of the bridge. The measured acceleration time histories at Pier 17 are shown in **Fig. 6**. The acceleration time histories measured longitudinal to the bridge (channel 17) show evidence of impact occurring between bridge segments. The influence of expansion joint impact on the dynamic response of a large bridge system is not currently well understood. Computational study of this phenomenon[11] indicates that it has a potential nontrivial influence on the global dynamics.

Since the purpose of the current study was to analyze the response of the Dumbarton Bridge to large magnitude earthquakes, in which both the amplitude and duration of shaking are significantly longer than that experienced in the Loma Prieta earthquake, it was deemed necessary to model the nonlinear behavior of the bridge expansion joints. Representation of impact phenomena in a finite element model significantly complicates the numerical computation task because it introduces strong nonlinearities into the system. In the current study, the penalty contact surface capability in NIKE3D was used to model the potential impact between adjacent bridge segments.

A typical Dumbarton Bridge expansion joint is shown in **Fig. 7**. Under very small amplitude motion, adjacent bridge segments "A" and "B" can move longitudinally, essentially independently of one another. The four inches of polystyrene is a joint filler that is intended to expand and contract with the adjacent bridge segments as the bridge heats up and cools down during the day. The adjacent bridge segments can move towards each other in a relatively unimpeded fashion until the existing gap closes and the adjacent bridge segments (i.e., segment A and segment B in **Fig. 7**) contact each other in a 'Mode 1' impact. Each bridge segment can also move relatively freely away from the pier until it individually impacts the pier in a 'Mode 2' or 'Mode 3' impact. The bridge segments can also move apart approximately six inches from their initial position unimpeded until 'tension only' cable and bar restraining elements go into tension and effectively couple the adjacent segments together, as indicated in 'Mode 4.' All of these potential modes of bridge interaction have been incorporated in the NIKE3D finite element model of the bridge, through the use of contact surfaces for contact/impact and discrete elements to represent the tension-only behavior of restrainer cables and bars. The NIKE3D joint model is shown schematically in **Fig. 7**.

Dumbarton Bridge Response in the Loma Prieta Earthquake

Direct time-history integration analyses have been performed to compute the response of the Dumbarton Bridge in the Loma Prieta earthquake. To assess the influence of expansion joints and foundations, parametric studies are being performed with the foundations and expansion joints included and neglected, respectively. Study of the bridge response measurements from the

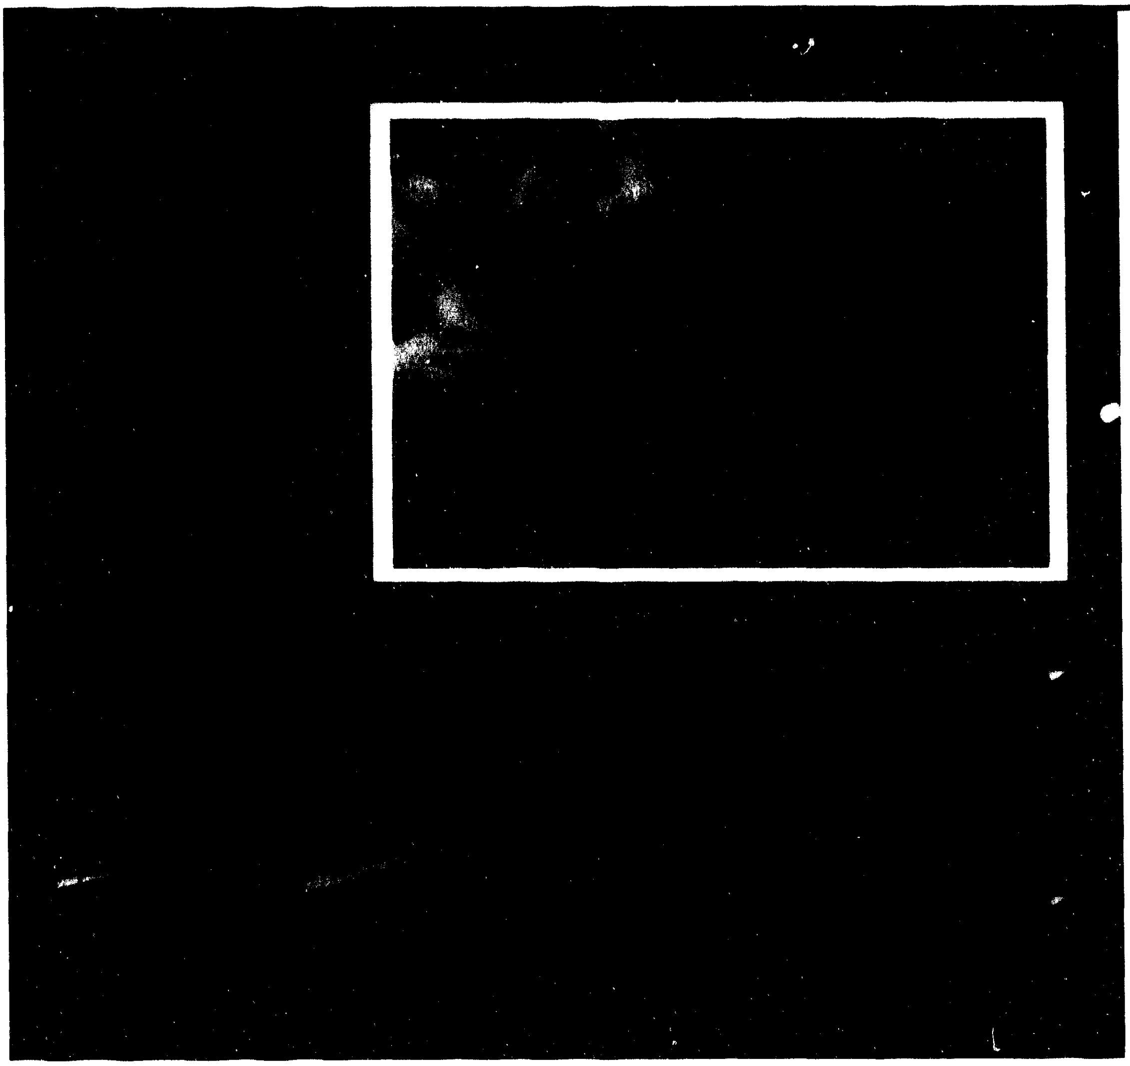

Figure 4. Three-dimensional solid model of the Dumbarton Bridge.

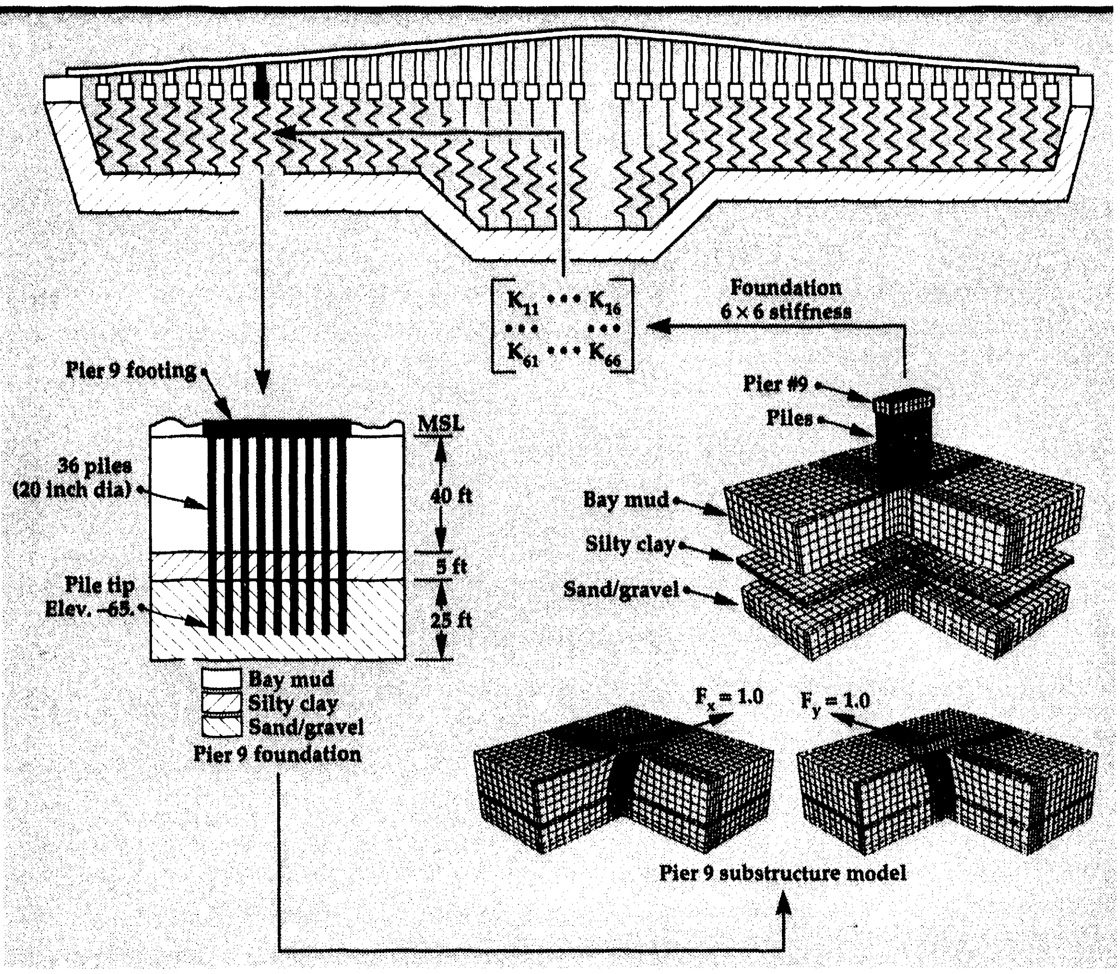

Figure 5. Representation of foundation stiffness matrix in global bridge model.

Loma Prieta earthquake indicated that impact appeared to be occurring only at the western deck hinge (see **Fig. 2**, **Fig. 6**, and **Fig. 8**). This was evidenced by the fact that acceleration spikes were observed in the longitudinal measurement stations in instrumented segment #2 (see **Fig. 3**), but were essentially absent in the longitudinal stations of segment #1 (see **Fig. 3**). Segment #2 and segment #1 are separated by a large, eighteen-inch expansion gap, and a check of displacements at channel 4 and channel 17 (**Fig. 3**) confirms that deflections were too small to allow collision to occur at that location.

Exaggerated displaced shapes computed with the finite element model of the Dumbarton are shown in **Fig. 8a** at selected time steps. At $t = 19.8$ s, the relative displacement is such that both deck hinges are opened farther than the nominal

TABLE 1. Finite element computational requirements for Loma Prieta earthquake time history computations.

Model	Number of time steps	Number of equilibrium iterations	Number of stiffness formations	Total Cray YMP solution time (h)
Linear model (no gaps)	1500	4578	3	2.17
Nonlinear model (6-in. gaps)	1500	4919	46	4.44
Nonlinear model (3-in. gaps)	1500	5199	52	4.86

six-inch gap at that particular instant in time. Computed and measured displacement time histories at the top of Pier 13 are shown in **Fig. 8b**. The computed displacements are in reasonable agreement with the observed displacements, and the accuracy of the model is commensurate with the accuracy obtained by the UCB team in their Loma Prieta response predictions.[9]

Computed acceleration time histories for the top of Pier #21 are shown in **Fig. 9** for two different cases. In the first case (**Fig. 9a**), contact surfaces and discrete elements were included in the model, and the initial gaps at the expansion joints were given the nominal design values indicated in **Fig. 2**. The true gap distances existing at the time of the Loma Prieta earthquake were a function of the temperature, and the earthquake occurred on a very warm afternoon in which the expansion joints were probably closed a certain amount as a result of thermal expansion. In addition, the existing gaps at the time of the earthquake were a function of previous loading history and any prior settlement or movement to which the bridge had been subjected.

When the bridge was analyzed assuming the nominal design expansion-gap distances, impact was observed at the west deck hinge (see **Fig. 2** and **Fig. 8**). The computations did not exhibit any impact in the piers associated with "instrumented segment #1." This concurred with the observed measured bridge response in that there was impact visible only in the recordings in the eastern instrumented segment of the bridge, i.e., "instrumented segment #2" in **Fig. 3**. However, comparison with the measured acceleration time history at Pier 21 indicated that the impact occurred much later in the time history than was observed in the bridge response data (see **Figs. 9a** and **9b**). Additional response calculations were performed with the deck hinge gap reduced from six inches to three inches, which is entirely possible given the high temperature on the day of the earthquake. With the smaller gap, impacts are clearly visible in the computed response earlier in the earthquake; however, the impact times do not precisely coincide with the observed impact times. Further investigation of the impact phenomenon, with consideration of various initial gap sizes, is ongoing.

Summary and Future Work

Performing nonlinear response calculations for such a large model is a very computationally intensive task. When this study was initiated, it was unclear if full-scale nonlinear response calculations were feasible for a bridge the size of the Dumbarton. Seismic analyses for a magnitude 8+ earthquake will require thousands of time steps, and it was unclear if a nonlinear analysis was feasible even on a Cray YMP. There are many potential pitfalls in a nonlinear analysis of this size, which can quickly make the calculations intractable. Simply determining appropriate error norms for evaluating convergence under base excitation loading was a question mark.

The analyses performed to date have indicated that the nonlinear computations, while time intensive, can be performed with appropriate convergence and within realistic supercomputer times. NIKE3D uses a quasi-Newton solution scheme in which the Broyden-Fletcher-Goldfarb-Shanno algorithm is used to perform stiffness updates. The full-system, instantaneous stiffness matrix is only reformed as required, when the program senses a failure to converge appropriately. The nonlinear solution parameters for the three different cases shown in **Fig. 9** are given in **Table 1**.

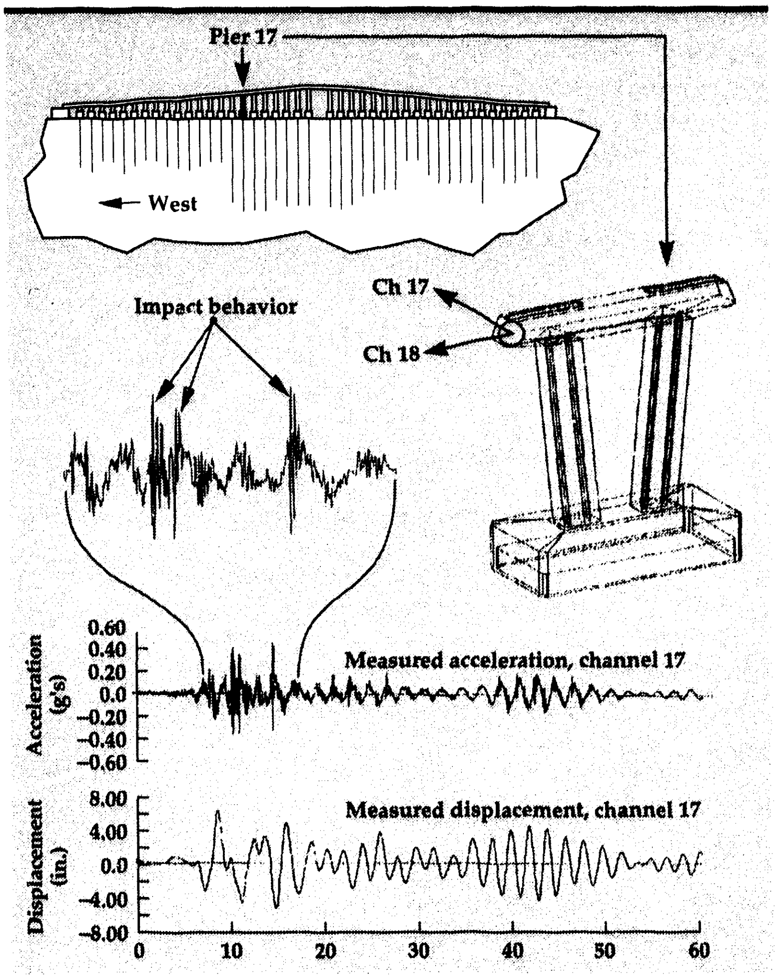

Figure 6. Structural response measurements at the top of Pier 17 for the Loma Prieta earthquake.

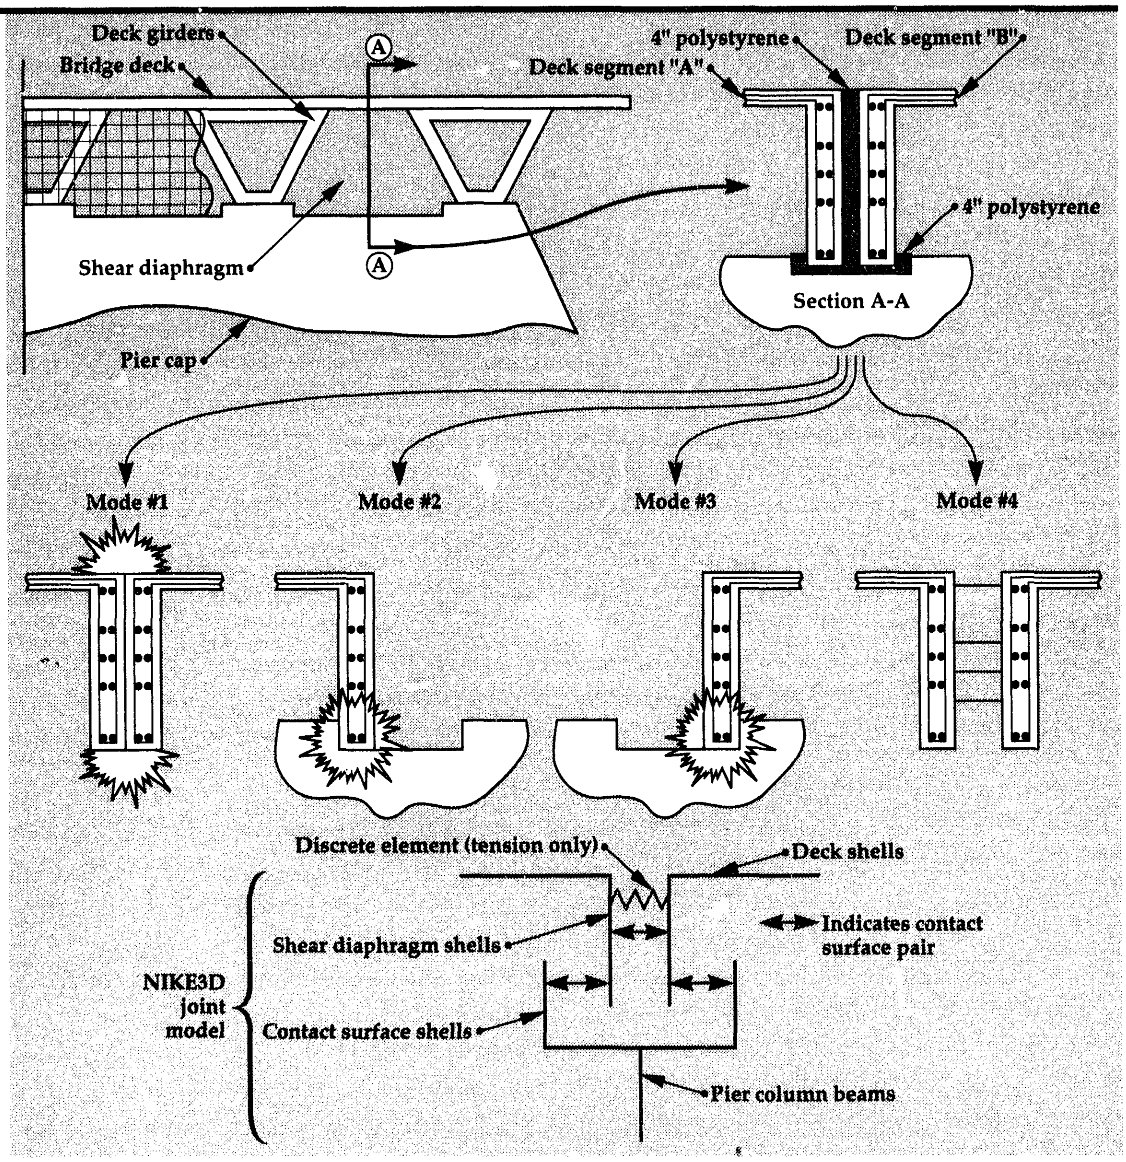

Figure 7. Typical 4-in. expansion joint detail and NIKE3D model.

The total Cray YMP solution time for the most costly computation (including joint banging with an initial gap of three inches) was approximately 4.5 h. Initial computations with an extreme magnitude 8+ synthetic earthquake record indicate that approximately 20 s of extremely strong ground motion, with extensive banging of many expansion joints, could be performed with about 8 h of Cray YMP computer time. It appears that a complete nonlinear solution for a magnitude 8+ earthquake is within our grasp.

Efforts to date at modeling the complex mechanical behavior of the bridge expansion joints are very encouraging. The model exhibits behavior that is commensurate with that observed in the field. However, additional work will be carried out in the coming year to refine our bridge joint models. For example, there are two thrust buffers, which are mechanical devices for softening impact, located at the deck hinges. These are not currently included in our model, and future effort will involve quantifying the behavior of these devices and including them in our mathematical bridge model. Parametric studies are also being performed to develop a clearer picture of the influence of the initial gap dimensions of the expansion joints. The effect of the bridge deck impacts on the bridge displacements and member forces will be evaluated. For the relatively mild Loma Prieta earthquake, the deck impacts had a significant effect on accelerations (**Fig. 9**), but only a minor influence on displacements. For a larger magnitude earthquake, we believe the influence of deck collisions will be more significant.

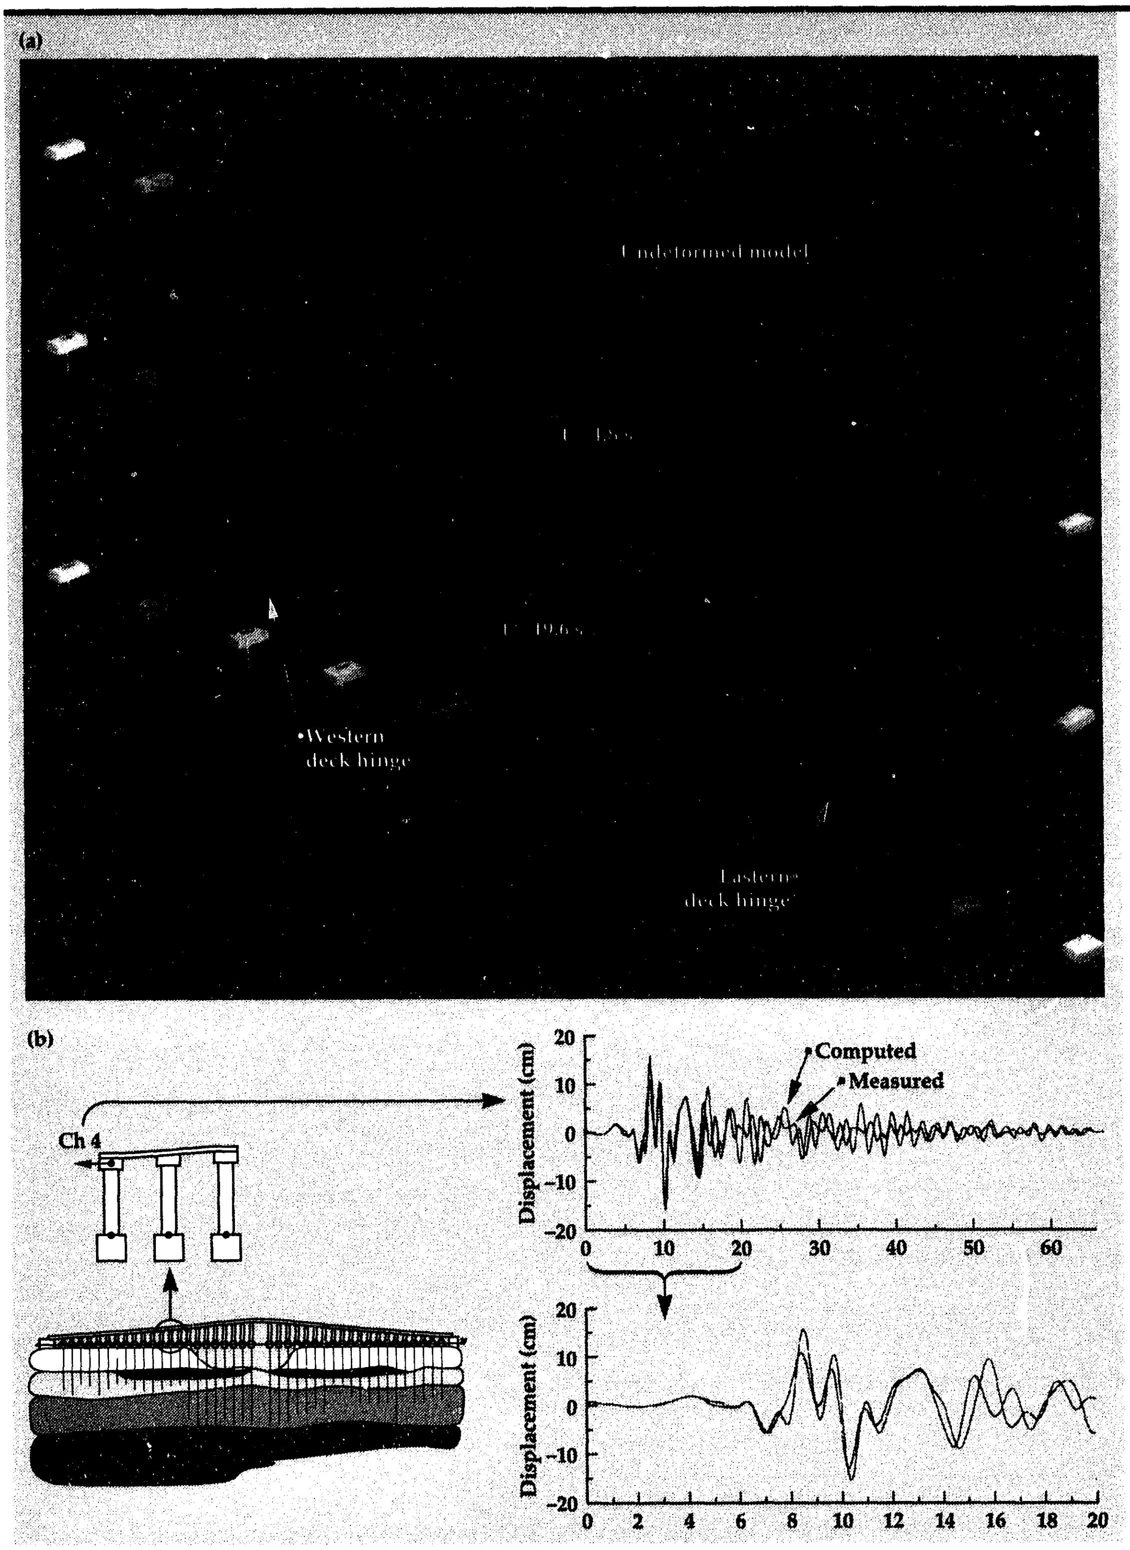

Figure 8. Computed Dumbarton Bridge response for the Loma Prieta earthquake: (a) displaced shape at selected time steps, and (b) computed displacement time history at Pier 13.

Seismologists working on the Computational Earthquake Initiative are in the process of developing synthetic ground-motion predictions for the Dumbarton site for large-magnitude earthquakes on the Hayward and San Andreas faults. The ground motion predictions will include the effects of differential motion across the length of the Dumbarton Bridge. Geotechnical engineers, in turn, will modify these earthquake motions to account for the effects of the site soil response (site drilling performed by Caltrans has provided material properties for the 700-ft-deep Dumbarton soil

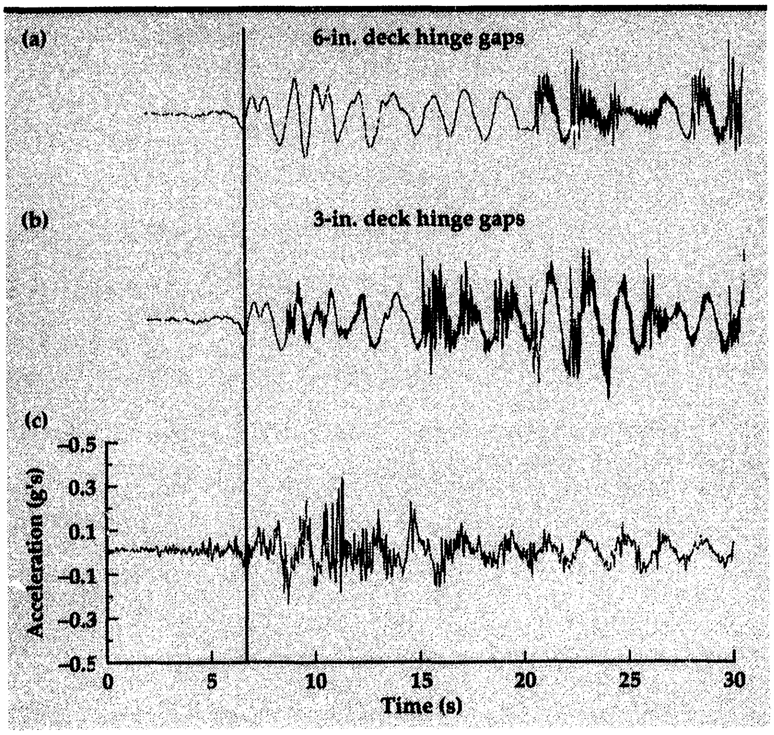

Figure 9. Acceleration time histories for Pier #21: (a) computed assuming a 6-in. expansion gap at deck hinges, (b) computed assuming a 3-in. expansion gap at deck hinges, and (c) measured.

profile) and of the nonlinear interaction between the pile foundations and the supporting San Francisco Bay mud. The large-earthquake ground motions will be used in the coming year to shake the Dumbarton structural model.

Acknowledgements

The author would like to acknowledge the continued support of Dr. G. Goudreau of the Mechanical Engineering Department and Dr. F. Heuze of the Earth Sciences Department at LLNL. Thanks also to Dr. B. Maker of the Methods Development Group for implementation of special features in the NIKE3D program and to M. Loomis for the generation of computer graphics representation of the bridge and computational results.

1. G.W. Housner (chairman), *Competing Against Time,* Report to Governor George Deukmejian from the Governor's Board of Inquiry on the 1989 Loma Prieta Earthquake (May 1990).
2. D.B. McCallen, F.E. Heuze, and P.W. Kasameyer, "A Methodology for Calculating the Response of Structures to Earthquakes on Specific Faults," *Soil Dynamics and Earthquake Engineering VI, Proc. 6th Int. Conf. Soil Dynamics and Earthquake Engineering,* Elsevier Applied Science (New York, New York), 1993.
3. L. Hutchings, *Bull. Seismological Soc. Amer.* **81**(5), (1991).
4. S.P. Jarpe and P.W. Kasameyer, *Validation of a Methodology for Predicting Broadband Strong Motion Time Histories from Empirical Green's Functions,* Lawrence Livermore National Laboratory, Livermore, California, UCRL-JC-113782 (1993).
5. R. Popescu and J.H. Prevost, "Numerical Class 'A' Predictions for Model Nos. 1, 2, 3, 4a, 6, 7, 11, and 12," *Verification of Numerical Procedures for the Analysis of Soil Liquefaction Problems Vol. 1, Proc. Int. VELACS Conf.,* A.A. Balkema (Brookfield, Vermont), 1993.
6. D.B. McCallen and K.M. Romstad, *Dynamic Response of a Short-Span Reinforced Concrete Overpass,* Lawrence Livermore National Laboratory, Livermore, California, UCRL-JC-113411 (1993).
7. D.B. McCallen and G.L. Goudreau, *Large-Scale Computations in Analysis of Structures,* Lawrence Livermore National Laboratory, Livermore, California, UCRL-JC-115223 (1993)
8. B.N. Maker, R.M. Ferencz, and J.O. Hallquist, *NIKE3D: A Nonlinear, Implicit, Three-Dimensional Finite Element Code for Solid and Structural Mechanics—User Manual,* Lawrence Livermore National Laboratory, Livermore, California, UCRL-MA-105268 (1991).
9. G.L. Fenves, F.C. Filipou, and D.T. Sze, *Response of the Dumbarton Bridge in the Loma Prieta Earthquake,* University of California Berkeley Earthquake Engineering Research Center, Berkeley, California, Report EERC-92/02 (1992).
10. B.A. Bolt, "Seismic Strong Motion Synthetics," *Proc. Seminar on Seismic Design and Retrofit of Bridges,* California Department of Transportation (Berkeley, California), (June 1992).
11. S.P. Singh and G.L. Fenves, "Earthquake Response of Multi-Span Freeway Structures," *Proc. 2nd Ann. Seismic Research Workshop,* California Department of Transportation (Sacramento, California), (March 1993).

ParaDyn: New Generation Solid/Structural Mechanics Codes for Massively Parallel Processors

Carol G. Hoover
National Energy Research
Supercomputer Center
Computation Directorate

James D. Maltby
Nuclear Test Engineering
Mechanical Engineering

Anthony J. De Groot
Engineering Research Division
Electronics Engineering

Robert G. Whirley
Nuclear Explosives Engineering
Mechanical Engineering

The goal of this project is algorithm research and development leading to a massively parallel version of explicit finite-element algorithms in the DYNA3D, nonlinear, solid and structural mechanics program. We discuss a design using message passing for the DYNA3D finite-element algorithm (the constitutive update and time integration) designed for the Lawrence Livermore National Laboratory Meiko CS-2 system. We summarize speed and scalability results for a data-parallel implementation of DYNA2D on a CM-5, and message-passing versions of DYNA3D on the Meiko CS-2.

Introduction

There are several challenging tasks that must be accomplished before massively parallel computing in solid and structural mechanics becomes a useful computational engineering tool for production tasks. The development of both parallel methods for the general unstructured meshes and multi-body interface contact algorithms exemplified in the DYNA3D computer program[1] is a fundamental step for progress in this field. An unstructured grid is identified as having a non-regular connectivity and spacing of the mesh indices and may contain several parts. An inherent property of finite-element algorithms is the indirect addressing (scatter operation) that occurs in the accumulation at the nodes of force contributions from connected elements. On a parallel computer, this results in irregular communication patterns between processors. Another characteristic of DYNA3D algorithms, advantageous for long time simulations, is the use of a low-order element-integration technique. On parallel systems, this means even greater efforts must be taken to reduce the interprocessor communication times. A goal in this research is to investigate parallel algorithms for DYNA3D that limit the communication times to a small fraction of the overall time. To achieve this goal, we are developing techniques for partitioning the problem onto the processors that minimize the nodal data communicated to other processors.

Contact algorithms involve searching for line and surface intersections on a global mesh. Because the searches are global over a set of processors, new algorithms must be designed for distributed-memory computers. Fast table-lookup techniques, sorts, and merges must be investigated and their performance analyzed for large deformation applications. Repartitioning of the contact surfaces and the elements will be necessary for applications in which contact surfaces change rapidly during the evolution of the problem.

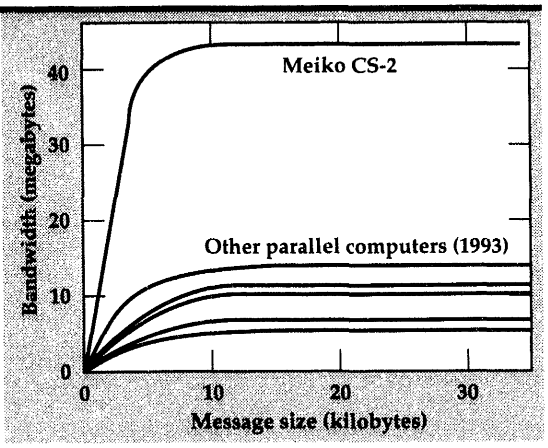

Figure 1. A comparison of network performance for massively parallel computers in 1993. These measurements represent maximum network performance resulting from using the lowest-level system software or from programming a communications co-processor.

The 256-processor Meiko CS-2 computer at Lawrence Livermore National Laboratory is a very promising architecture for designing efficient algorithms in the DYNA3D program. The most advantageous feature is the low message latency of the CS-2 network compared to other competitive computer systems. Message latency is a measure of the time required for preparing the data to be sent out through the internal communications network. For the best performance, the communications co-processor can be programmed to obtain network latencies as low as 9.4 μs. For more portable programming, Meiko has provided a message-passing subroutine library with an interface matching the Intel NX library. The measured latency for this library is roughly 70 ns. This low latency feature of the network compared to other systems is illustrated in **Fig. 1**. Each processing unit in the CS-2 consists of a 50-MHz Viking superscalar processor and two 45-MHz Fujitsu vector pipes. Peak speed for a 256-node system is 50 GFLOP.

Advanced applications in solid mechanics, such as multicar crash and airliner safety simulations, that would benefit from gigaflop speeds and gigaword memories are abundant. The algorithm research we are pursuing in this project will provide the opportunity for better turnaround in engineering design processes as well as the possibility for more complex applications.

Progress

Algorithm Research

Static load balancing implies a problem partitioning at initialization only, whereas dynamic load balancing repartitions the problem periodically during the course of the calculation. The partitioning must balance the computational load among the processors while minimizing the interprocessor communication time. Partitioning techniques are applicable to a wide selection of algorithms involving grid connectivities and matrix methods. Recursive spectral bisection techniques have evolved recently from this research activity. These techniques use the adjacency graph (the Laplacian matrix) of the connectivity matrix to find vertices in a minimum-size separator list for a partition. Simon[2] developed a heuristic method to compute the separator list for a general mesh connectivity. The algorithm uses components of the eigenvector associated with the smallest nonzero eigenvalue of the Laplacian matrix. We have tested this method, using Simon's routines, for both speed and partitioning quality. We have also investigated newer multilevel techniques,[3] and measured factors of 8 to 20 speed increases depending on the mesh size.

The design of a flexible data structure to link the partitioning output into the DYNA3D element models and contact algorithms is an essential component in this algorithm research. We have designed the data structures to support a single-program, multiple-data methodology on the Meiko CS-2. This method assumes each processor executes the full DYNA3D program on its portion of the grid and communicates nodal point data to other processors as needed during the calculation. Our element-partitioning technique generates a file of partitioning information based on the element edge and face lists calculated from recursive spectral bisection. The partition file contains (1) a unique assignment of elements to processors, (2) an assignment of nodes to any processor containing a connected element, and (3) a list of 'edge nodes' by global node number. The edge nodes are defined as those nodes that must communicate information between processors. From these data, read during the DYNA3D input phase, we generate local to global node- and element-number lists and the list of global node numbers for edge nodes. These are sufficient data structures to set up the communication of element forces to edge nodes during the force assembly step.

This method has several advantages. The single processor version of the DYNA3D input file may be used without changes. Partitioning for a different number of processors involves only regenerating the partition file without any need to regenerate the DYNA3D input file. The generation step for the partition file is flexible enough so that we can modify it for algorithm storage optimization as well as for overlapping calculations and communications. **Figure 2** illustrates the software developed to include the partitioning step for massively parallel computers.

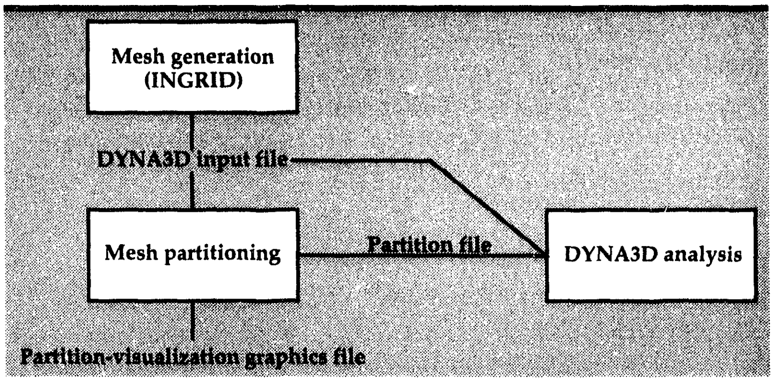

Figure 2. Mesh-partitioning software added for analysis performed on massively parallel computers. The mesh-partitioning software is added between the mesh generation and the DYNA3D analysis.

We have written a solid-element kernel program using algorithms extracted from DYNA3D and retaining all the characteristics of the DYNA3D programming for solid elements. The kernel includes the elastic and plastic material models and sufficient boundary conditions to test the standard bar impact data case. We have tested the parallel version of this kernel on up to 32 processors on the Meiko for a thousand-cycle calculation for the plastic deformation of a bar impacting a wall.

DYNA2D Results on the CM-5

Our progress on the DYNA2D project has included achieving speeds that are ten times Y-MP production run speeds, developing the storage strategies for a data-parallel model, and measuring performance. The key finding is that algorithm kernels are essential for projecting performance on massively parallel computers for DYNA3D.

The final part of this work has been the completion of the performance analysis and a projection of performance over the next two years on the CM-5 for unstructured grid problems using the Thinking Machines' data parallel software model. The version 3.0 (CMF and CMSSL) software is unable to distinguish between the indirect addressing (the scatter and gather steps) that occur on-processor and those that occur off-processor. The result of this is that *all* indirect addressing incurs the expense of the message-passing latency, whether or not there is any real communication needed. This severely limits the performance of a data-parallel model. We project that if this software problem was corrected, the performance for the bar impact problem would improve from .530 µs per element cycle to .133 µs per element cycle on 512 processors. A rough estimate of the gigaflop rates are 1 GFLOP and 4 GFLOP, respectively. The two-dimensional contact algorithms performed less efficiently, yielding speeds less than 1 GFLOP. Further large speed improvements for DYNA2D are limited by the bandwidth between the SPARC registers and memory (20 MB/s shared among four vector pipes) compared to the bandwidth from the vector processors to their own memory (128 MB/s).

Preliminary Timing Results for the Meiko CS-2

During the pre-acceptance period for the 64-processor Meiko CS-2, we verified the correctness of the physical results for the bar impact problem, using our message-passing algorithms. Very preliminary speedup and timing results for the solid-element kernel are shown in **Fig. 3**. These measurements were made using the superscalar processor and no vector processing. Problem sizes of 971, 7776, and 62,208 elements were run on one, two, four, eight, sixteen, and thirty-two processors. The 62,208-element problem uses less than one half of the processor memory for a single processor and thus, represents a very small problem size for the class of problems targeted for the Meiko. The speedup is reasonable and corresponds to an efficiency of 91% for thirty-two processors. The 971-element problem illustrates the lack of scalability for calculations in which most of the problem is communicated between processors at each timestep.

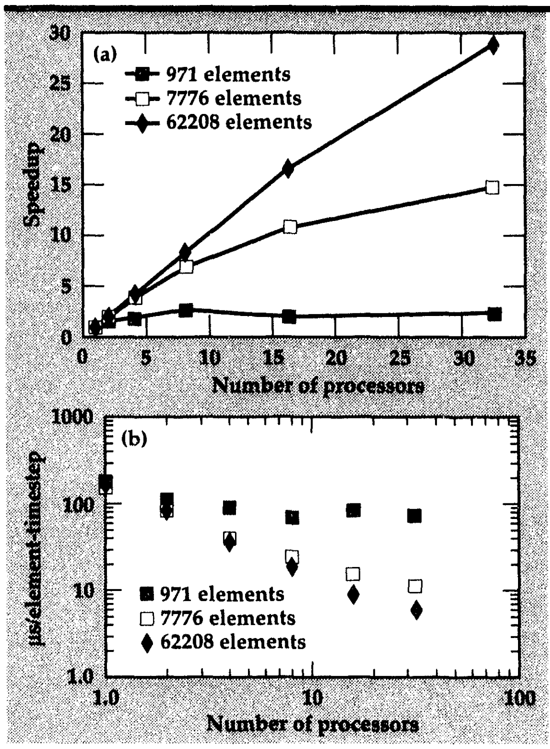

Figure 3.
(a) Speedup of the DYNA3D kernel on the Meiko CS-2;
(b) execution time for the DYNA3D kernel on the Meiko CS-2.

We expect significantly different and improved results in the future. We plan to tune our algorithms to improve communication speeds as well as to take advantage of vectorization. We expect better network performance as a result of better system-routing tables. Similarly, computational speeds will improve as the compiler software for vectorization evolves and matures.

Future Work

The principal effort in the coming year will be research in parallel-contact algorithms and dynamic domain-decomposition techniques. Once contact algorithms are implemented, we will demonstrate large-scale applications such as penetration and crash worthiness.

Future work will focus on further optimization and parallel-algorithm development for the DYNA3D program on the Meiko CS-2 computer. The investigations will include developing and testing message-passing strategies (such as overlapping communications and calculations), optimizing message-buffering techniques, extending the vector algorithms to address data-locality issues, and performing detailed parallel-speed and performance measurements. Message passing will be developed first with the Intel NX library interface, and later selected routines will be implemented with Meiko-specific software if optimization requires it. Further performance tests may lead to new algorithms that reduce storage as well as increase computational speed using the superscalar and vector processors. A goal in this work is to achieve speeds above 10 GFLOP with a 256-processor system.

Acknowledgements

We received very valuable work and advice from G. Montry of Southwest Software in implementing message passing in the solid-element kernel on an Intel. We especially wish to thank the staff at the Livermore Computer Center for the early use of the Meiko CS-2 and for their excellent support.

---

1. R.G. Whirley and B.E. Englemann, *DYNA3D: A Nonlinear, Explicit, Three-Dimensional Finite Element Code for Solid and Structural Mechanics—User Manual*, Lawrence Livermore National Laboratory, Livermore, California, UCRL-MA-107254, Rev. 1 (1993).
2. H.D. Simon, *Computing Systems in Engineering* 2(2/3), 135 (1991).
3. S.T. Bernard and H.D. Simon, "A Fast Multilevel Implementation of Recursive Spectral Bisection for Partitioning Unstructured Problems," *Proc. Sixth SIAM Conf. Parallel Processing for Scientific Computing*, (March 22–24, 1993), D.E. Keyes, M.R. Leuze, L.R. Petzold, and D.A. Reed (Eds.), 1993.

Microchannel Boiling Mechanisms Leading to Burnout

Charles S. Landram
Nuclear Test Engineering
Mechanical Engineering

Robert A. Riddle
Applied Research Engineering
Mechanical Engineering

We are analyzing the thermal performance of microchannel heat sinks to extend their applied heat loads beyond coolant single-phase limits. Ours is the first investigation of boiling in the narrow (50-μm) microchannels having typically high-aspect-ratio (of order 10/1) flow cross-sections. The prescription of local, wall-coolant, interfacial, two-phase correlations first required development of a validated, approximate, thermal-model accounting for conjugate heat transfer. The strongest mechanism for heat transfer in two-phase microchannel flow was found to be saturated boiling in a channel region near the heated base. When this region dried out, burnout occurred, both in the computations and in the experiment.

Introduction

Thermal analyses of microchannel heat sinks—μ-sized, laminar-flow cooling channels fabricated into semi-conductive substrates—have been limited to single-phase coolant flows.[1,2] Fluid entering the base-heated parallel channels, which are separated by highly conducting walls, exits as a thermally developed flow (**Fig. 1**). As the base heating flux is increased, the fluid at the channel exit, where the coolant and wall temperatures are greatest, begins subcooled and/or saturated boiling. Further increases in heat flux can be sustained, with the incipience of boiling moving axially inward into the channels, until a crisis exists at the exit. Use of these high-performance heat sinks with increasing heat flux mandates an understanding of this crisis, which is the purpose of this work.

Analysis of the boiling events is a three-dimensional (3-D) problem. In addition to the axial variations, the various boiling regions can occur within each cross-section in the longitudinal

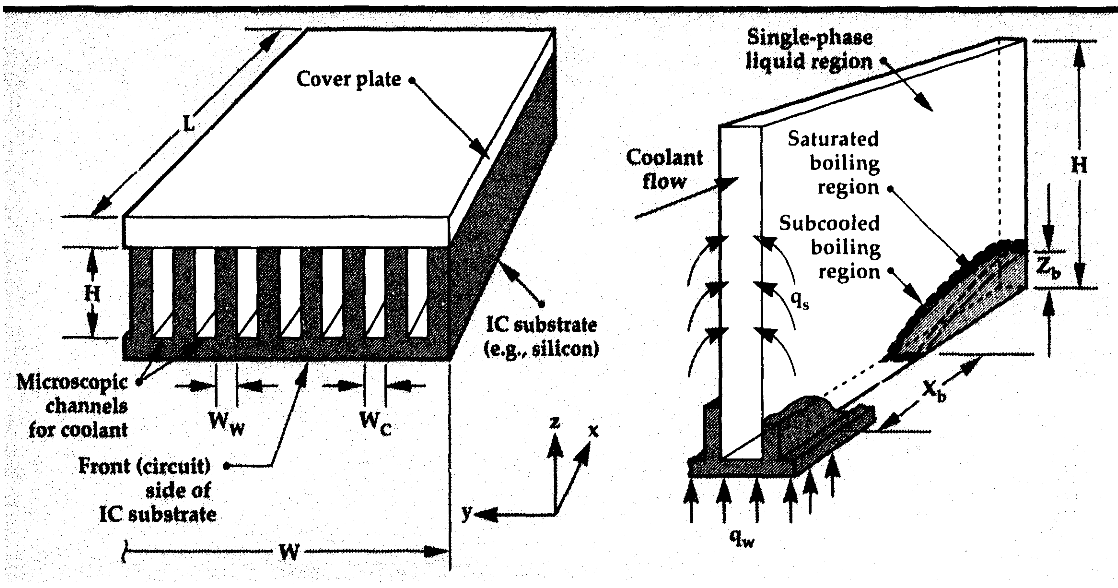

Figure 1. The 3-D boiling phenomena occurring within the channel.

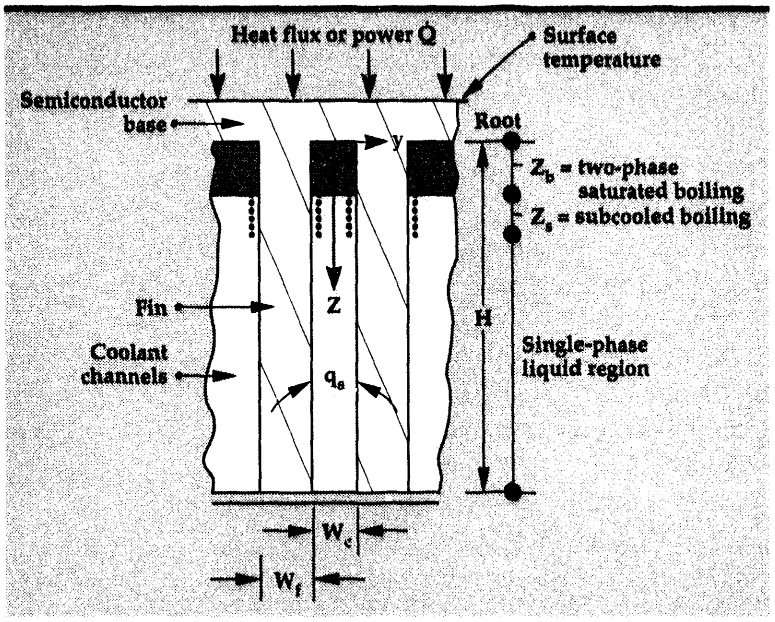

Figure 2. The channel cross-section at a fixed axial position, with several flow regimes.

and lateral directions within the channel (**Fig. 2**). Within any cross-section, single-phase flow, subcooled and saturated boiling can occur as separated, but thermally coupled flows. Computations of this complicated flow have not previously been attempted, not only owing to its 3-D nature, but also because of the necessity to treat the fluid-wall interface as a conjugate situation; i.e., a surface heat transfer coefficient is unknown and cannot be specified.

Progress

We first developed an approximate model to accommodate a prescribed, fluid-wall, thermal interfacial conductance for either single-phase or subcooled boiling, whichever qualified. This model is shown to represent the conjugate problem, yet allow interfacial specifications. The saturated boiling region is treated separately. All the regions are coupled by an overall enthalpy balance and the prevailing local regional requirements.

Approximate Single-Phase Model

The entire cross-section, $w_c/2 < y < w_c/2$, $0 < Z < H$ (**Fig. 2**), is first considered in single-phase flow. For fully-developed laminar flow with uniform base-heat flux, q_w, the energy equation for the coolant can be integrated over the lateral y range to give a local longitudinal coolant temperature, $T_0(Z)$, averaged in the lateral direction as

$$\frac{w_c}{2}k_c\frac{d^2T_0}{dZ^2}+q_s(Z)=f(Z)\cdot\frac{q_w\left(w_c+w_f\right)}{2H}, \qquad (1)$$

where k_c is the coolant thermal conductivity, $q_s(Z)$ is the local interfacial heat flux between the separating wall and the coolant, and f(Z) is defined by the known local velocity u normalized to its mean value u_m:

Table 1. Criteria for subcooled boiling interfacial heat flux.*

$q_s = A + B$

Criterion	A	B
$T_{w_{ONB}} < T_w < T_{w_{FDB}}$	$h^1(T_{sat}-T_0)$	$\dfrac{1}{\varphi^2(T_w-T_{sat})^2}$
$T_w > T_{w_{FDB}}$; $T_0 < T_{sat}$	0	$\dfrac{1}{\varphi^2(T_w-T_{sat})^2}$
$T_w < T_{w_{ONB}}$	$h^1(T_w-T_0)$	0

where

$$T_{w_{ONB}}=T_{sat}+\frac{h^1\varphi^2}{2}\left[1+\sqrt{1+\frac{4}{h^1\varphi^2}\left(T_{sat}-T_0\right)}\right]$$

$$T_{w_{FDB}}=T_{sat}+\sqrt{1.4}\left(T_{w_{ONB}}-T_{sat}\right)$$

*from Collier, Ref. 4. Values of the parameter φ are also given in Ref. 4.

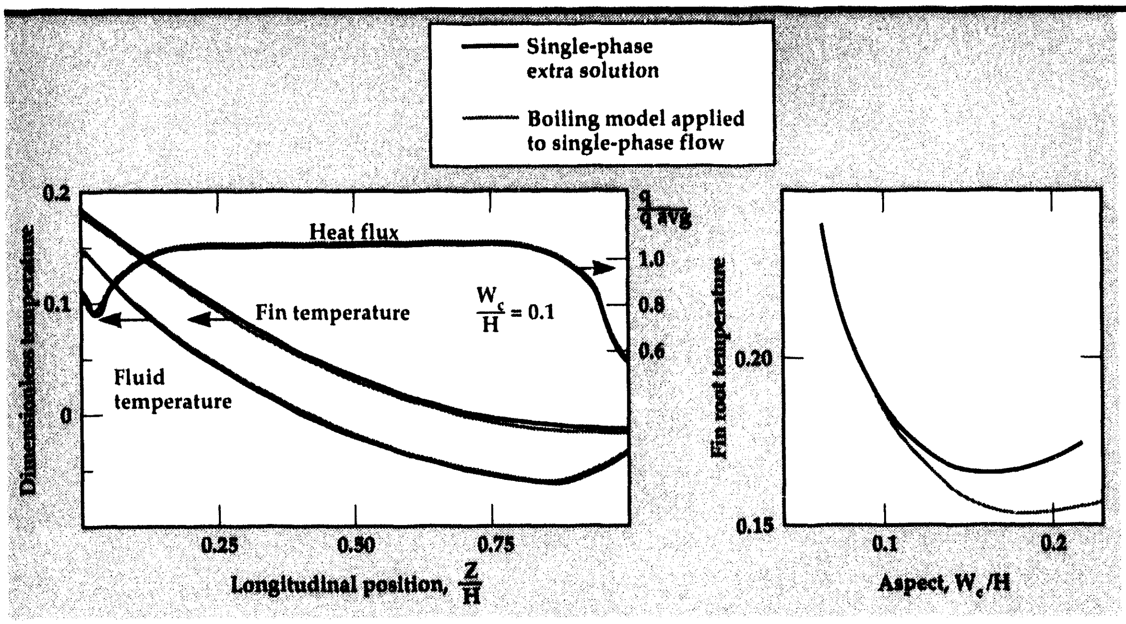

Figure 3. The prescriptive conjugate model used here for the single-phase and subcooled boiling regimes, validated for the case of the entire cross-section in single-phase flow. The model is best for large aspect ratios and predicts their optimum values.

$$f(Z) = \frac{2}{w_c}\int_0^{w_c/2} \frac{u(y,Z)}{u_m}\,dy\,.$$

Other symbols shown in **Eq. 1** are defined in **Fig. 2**.

Integrating in the lateral direction across the separating wall, called a fin, of thickness w_f, gives its local longitudinal temperature $T_w(Z)$ as

$$\frac{w_f}{2}k_f\,\frac{d^2T_w}{dZ^2} - q_s(Z) = 0\,, \qquad (2)$$

where k_f is the fin thermal conductivity, and q_s is the aforementioned interfacial heat flux. The temperature across the fin is taken to be uniform.[2]

The sum of **Eqs. 1** and **2** is independent of q_s, so that with either T_0 or T_w known, the other can be determined. Thus far, our formulation is exact within its framework; however, closure will require postulate of the interfacial heat flux q_s in terms of local variables T_0 and T_w.

We now focus on interfacial flux. For single-phase internal flows, the traditional heat transfer coefficient is defined in terms of the difference between local wall temperature and fluid bulk (or mixed mean fluid temperature). This definition is abandoned here and replaced by a definition based upon local wall and fluid temperatures. For single-phase flow, we assume

$$q_s = h^1\left[T_w(Z) - T_0(Z)\right], \qquad (3a)$$

where h^1 is determined as if heat is transferred in the lateral direction only:

$$h^1 \cong \frac{140}{17}\,\frac{k_c}{2w_c}\,. \qquad (3b)$$

It should be underscored that **Eqs. 3a** and **3b** are hypothetical and only have validity if they provide results that are in essential agreement with those given by the conjugate solution.[1,2] The value of the treatment here is in terms of an interfacial prescription, which in the boiling regime allows local, phenomenologically based correlations to be introduced. It is therefore crucially important that **Eqs. 3a** and **3b** be validated once the system of **Eqs. 1** through **3** is solved with the following boundary conditions:

$$\frac{dT_o}{dZ} = \frac{dT_w}{dZ} = 0 \quad \text{at} \quad Z = H$$

$$-\frac{dT_w}{dZ} = \frac{q_w\left(w_c + w_f\right)}{w_f k_f} \quad \text{at} \quad Z = 0$$

$$T_b = \frac{1}{H}\int_0^H f(Z)T_0(Z)dZ\,\cdot$$

The last condition is used to relate local single-phase results to the bulk temperature, T_b, which is obtained from an overall coolant energy balance. The solution to the above system of equations is analytical, and those details are omitted for brevity. The validity of the method is shown by the excellent agreement between local heat flux and fin and fluid temperatures, shown in **Fig. 3** for the case considered. Although **Eqs. 3a** and **3b** provided reasonable results for large aspect ratios, W_c/H, their validity lessened for small aspect ratios, as illustrated in the inset of **Fig. 3**.

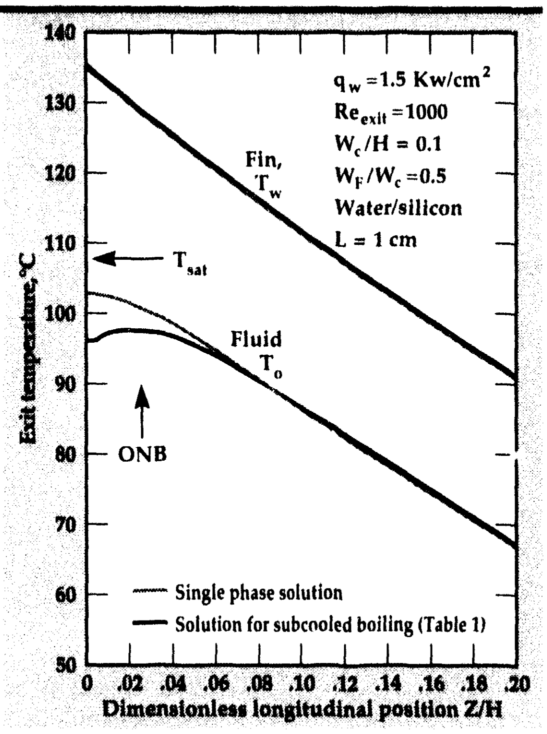

***Figure 4.* Longitudinal locations that qualify for local subcooled boiling, occurring near the root. The wall temperature is not perturbed by subcooled boiling.**

Subcooled Boiling

Subcooled boiling first occurs in longitudinal positions in the vicinity of the base (Z = 0). In the same cross-section, the more remote longitudinal positions will remain in single-phase flow. Because q_s can be considered a local prescription, regions undergoing subcooled boiling can be accommodated by substituting the appropriate interfacial flux along $0 < Z < H$ instead of **Eq. 3a.** These criteria are given in **Table 1** along with the qualifying wall-superheat conditions.

The solutions to the appropriate **Eqs. 1** and **2** and **Table 1** are now obtained numerically, accounting for the regions that are either in single-phase flow or in subcooled boiling. The heat flux here is not raised to levels causing saturated boiling, which is to be considered subsequently. The results are shown in **Fig. 4.** For the case shown, the onset of subcooled boiling occurs very near the root (Z = .007 H), and fully developed subcooled boiling does not occur, even as saturated boiling ($T_0 = T_{sat}$) is approached. Also, the fin temperature, T_w, is practically unperturbed by subcooled boiling. Several other cases demonstrated similar results, supporting our conclusion that subcooled boiling phenomena are not significant, at least for microchannel flows having substantial aspect ratios H/w_c of order 10.

Saturated Boiling

The fluid temperature $T_0(Z)$ cannot exceed the saturation temperature at the prevailing pressure. At sufficiently high heat fluxes and at particular axial positions, there will exist a longitudinal range $(0 < Z < Z_b)$ in which saturated boiling occurs, and beyond which $(Z > Z_b)$ the remainder of the coolant in the cross-section is in subcooled boiling and single-phase flow (primarily the latter based on the previous discussion).

The saturated region $(0 < Z < Z_b)$ heat flux can be shown to be governed by the annular flow criteria,[3] in which the local heat flux is given by

$$q_s = h_{tp}^1 \left(T_w - T_{sat} \right) . \qquad (4a)$$

The local conductance h_{tp}^1 is generally comprised of two components, a forced convection component and another contributed by nucleate boiling, the latter shown to be inoperative because (a) nucleation sites are not easily replenished by liquid in the μ-sized channels and, (b) separate calculations showed its credit to be overly optimistic relative to available experimental data (to be shown subsequently). The forced convection component is given in terms of the Martinelli parameter X and the single-phase conductance h^1 (**Eq. 3b**) by[4]:

$$\frac{h_{tp}^1}{h^1} = \frac{3.5}{\sqrt{X}} . \qquad (4b)$$

The Martinelli parameter X, applied to laminar flow, was related to the spatially averaged, mass quality ψ in the saturated region, assumed independent of longitudinal position Z in $0 < Z < Z_b$. In place of **Eq. 1**, the dependence of quality ψ on axial position was determined by an energy balance at any location within the saturated region. The result yielded a first-order, non-linear differential equation that was numerically integrated beyond the axial position x_b, where saturated boiling commenced; i.e., at $T_0(0) = T_{sat}$. **Equation 2** was used to determine T_w.

Model Closure

The bulk enthalpy at an axial position L is established by the overall heat balance between the inlet and that axial position:

$$e_{bulk} = c_p T_i + \frac{q_w \left(w_c + w_F \right) L}{\dot{m}} , \qquad (5a)$$

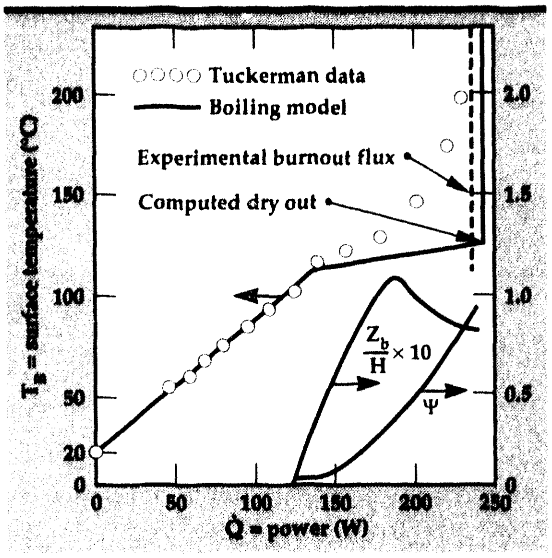

***Figure 5.* Saturated boiling in a region near the root, the key mechanism for two-phase microchannel cooling. Burnout occurs when the saturated region dries out.**

where $\dot{m}$ is the mass rate of flow through the channel, and T_i is the inlet temperature. The enthalpy at that same cross-section is given by

$$e_{bulk} = \left(h_f + \psi h_{fg}\right)\frac{Z_b}{H} + \frac{c_p}{H}\int_{Z_b}^{H} f(Z)T_0(Z)dZ\,, \qquad (5b)$$

where h_f and h_{fg} are the enthalpy of the saturated liquid and heat of vaporization, respectively.

Equations 5a and **5b** are equated, and the single-phase temperature $T_0(Z)$ in $Z_b < Z < H$ is found analytically by solving **Eqs. 1** through **3**, with the boundary conditions at $Z = Z_b$:

$$T_0(Z_b) = T_{sat} \qquad (6)$$

and

$$-\frac{dT_w}{dZ}(Z_b) = \frac{q_w(w_c + w_F)}{w_F k_F}(1 - F)\,, \qquad (7)$$

where F is the fraction of energy absorbed in the saturated regions. The solution is iterative, and unique values of quality, saturated region size Z_b and root temperature $T_w(0)$, were computed.

A comparison between model calculations and the experimental results of Tuckerman[5] is shown in **Fig. 5**. The surface temperature T_B was determined by solving the heat conduction equation across the base thickness.[2] The model demonstrates that the strongest mechanism for enhanced heat transfer in boiling microchannel flow is local saturated boiling in $0 < Z < Z_b$, where the large energy sink, given by the term $\psi h_{fg} Z_b$, is responsible for nearly doubling the single-phase heat flux before burnout is experienced. The model predicts the measured burnout heat flux at the channel exit for the case presented, and furthermore, burnout is observed to coincide with saturated-region dry out, when $\psi = 1$ (phase change energy is no longer available).

Future Work

The model optimistically predicts the measured base temperature shown in **Fig. 5**. Based on Woods' data[6] for saturated boiling in tubes, McAdams[7] ascribes the diminished experimental performance to a breakup of the liquid wall annulus film, especially at mass qualities exceeding 30%. This is unaccounted for in our saturated region model (i.e., **Eq. 4b**).

Subsequent work will modify the model by using a saturated-region, thermal conductivity based on mixture quality as opposed to the traditional practice of basing it on the liquid value. This will effectively attempt to account for the annular liquid-film breakdown. We will also compare the model to our current data obtained at higher flow rates and larger flow widths.

---

1. C.S. Landram, *An Exact Solution for Conjugate Longitudinal Fin-Fluid Heat Transfer in Internal Flow Including Optimization*, Lawrence Livermore National Laboratory, Livermore, California, UCRL-JC-103249 (1990).
2. C.S. Landram, *ASME Heat Transfer in Electronic Equipment* **HTD-171**, 127 (1991).
3. C.E. Dengler and J.N. Addoms, *Chem. Eng. Progress Symp. Series* **52**(18), 95 (1956).
4. J.G. Collier, *Convective Boiling and Condensation*, McGraw–Hill (New York, New York), 147 and 203, 1972.
5. D.B. Tuckerman, *Heat-Transfer Microstructures for Integrated Circuits*, Ph.D. Thesis, Stanford University, Palo Alto, California; also Lawrence Livermore National Laboratory, Livermore, California, UCRL-53515 (1984).
6. W.K. Woods, *Heat Transfer in Saturated Flow Boiling*, Sc.D. Thesis in Chem. Eng., Massachusetts Institute of Technology, Cambridge, Massachusetts (1940).
7. W.H. McAdams, *Heat Transmission*, McGraw–Hill (New York, New York), 398 (Figs. 14–26), 1954.

Rigid-Body Mechanics for Finite Element Analysis

Bradley N. Maker
Defense Technologies Engineering Division
Mechanical Engineering

Perhaps the most common approximation in engineering is that, relative to its neighbors, a system component is structurally rigid. Closed-form analytics frequently rely on the assumption of rigidity to simplify large mechanical systems into manageable sub-systems. Rigidity can often be justified in that it provides a worst-case stress or loading scenario. We have now developed the assumption for use in our solid- and structural-mechanics finite element code, NIKE3D.

Introduction

Finite element methods have traditionally been applied to problems where complexities in geometry or material behavior preclude closed-form analytic solutions. The power of the finite element approximation has allowed engineers to include considerable detail in their structural response calculations, potentially eliminating the need for simplifying assumptions. However, as engineers investigate processes with nonlinear material and geometric behavior, computational resources are quickly exhausted, and the assumption of rigidity once again becomes valuable. An excellent example is in sheet metal forming, where the assumption of rigid tooling is very reasonable.

Finite elements may be rendered rigid in two ways. The simplest method is to apply very large elastic moduli to an otherwise deformable element. This 'brute force' method requires no code modifications, yet has several shortcomings. Rigid elements still need to be processed during stress, strain, and stiffness evaluations. Nodal degrees of freedom associated with rigid elements remain independently active. And large disparities in material properties lead to numerical round-off errors.

Progress

Our rigid material implementation is based instead on rigid-body mechanics. In this method, a group of finite elements may be defined to be a rigid body. The center of mass for these elements is computed, and (at most) six degrees of freedom are assigned to the body—three for translation and three for rotation. The global degrees of freedom associated with all of the nodes in the rigid body are then replaced by the six rigid-body degrees of freedom, condensing the size of the model. Nodal loads and boundary conditions are resolved onto the center of mass coordinates. The motion of the rigid body is then determined by solving a coupled set of equilibrium equations for the rigid body and any other deformable elements in the model. Finally, the current coordinates of rigid-body nodes are updated using the center of mass coordinates and the rigid-body assumption.

This approach results in a method where stress, strain, and stiffness computations for rigid elements may be skipped. Numerical errors from disparate material properties are avoided. Most importantly, the total number of independent degrees of freedom in the model is reduced. This allows very large numbers of elements to be included in rigid bodies without adding significant computing cost to the model, and allows existing models to be executed more quickly by invoking the assumption of rigidity where appropriate.

Formulation

The position vector x of a finite element node point may be written as

$$\mathbf{x} = \mathbf{X} + \mathbf{u} \, ,$$

where **X** is the initial coordinate of the point, and **u** is its displacement vector. If the point is associated with a rigid body, then

$$\mathbf{x} = \overset{cm}{\mathbf{X}} + \mathbf{a} \, ,$$

where $\overset{cm}{\mathbf{X}}$ is the current position of the center of mass, and **a** is the current vector from the center of mass to the point. The vector **a** may be written in terms of $\mathbf{a}_0$, its value at $t = 0$, and a skew-symmetric rotation matrix Λ:

$$\mathbf{a} = \Lambda \mathbf{a}_0 \, .$$

The incremental displacement relationships required for NIKE3D[1] are obtained by linearization of the above expressions:

$$\Delta\mathbf{u} = \Delta \overset{cm}{\mathbf{X}} + \Delta\Lambda\mathbf{a}_0 \, .$$

Linearization of the rotation matrix Λ, presented by Simo,[2] results in expressions of a convenient form:

$$\Delta\Lambda\mathbf{a}_0 = \mathbf{R}\Delta\boldsymbol{\theta} = \begin{bmatrix} 0 & a_z & -a_y \\ -a_z & 0 & a_x \\ a_y & -a_x & 0 \end{bmatrix} \begin{Bmatrix} \Delta\theta_x \\ \Delta\theta_y \\ \Delta\theta_z \end{Bmatrix} ,$$

where $\mathbf{a} = \{a_x, a_y, a_z\}^T$, and $\{\Delta\theta_x, \Delta\theta_y, \Delta\theta_z\}$ are the rotation increments of the rigid body expressed in global coordinates.

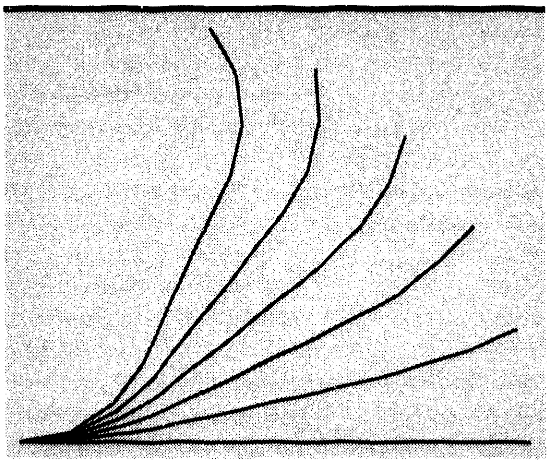

Figure 1. The centerline of an elastic cantilevered beam loaded with an end moment, shown at several stages of deformation. The beam is discretized with ten elements, the central four elements being rigid. Load is applied in ten increments. The rigid behavior of the central portion of the beam through finite rigid-body translation and rotation is demonstrated.

For a model containing both deformable and rigid components, the nodal degrees of freedom may be grouped, and the above expressions used to obtain a condensed set of unknowns:

$$\begin{Bmatrix} \Delta\mathbf{u}^D \\ \Delta\mathbf{u}^R \end{Bmatrix} = \begin{bmatrix} \mathbf{I} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{I} & \mathbf{R} \end{bmatrix} \begin{Bmatrix} \Delta\mathbf{u}^D \\ \Delta \overset{cm}{\mathbf{X}} \\ \Delta\boldsymbol{\theta} \end{Bmatrix} \quad \Rightarrow \quad \Delta\mathbf{u} = \mathbf{A}\Delta\hat{\mathbf{u}} \, ,$$

where the $\hat{(\cdot)}$ superscript denotes a condensed degree of freedom vector. Substituting this expression into the discrete form of the principle of virtual work, we obtain expressions for the condensed finite element stiffness matrix and residual vector for the coupled deformable/rigid system:

$$\hat{\mathbf{K}}\Delta\hat{\mathbf{u}} = \hat{\mathbf{F}} \, , \quad \hat{\mathbf{K}} = \mathbf{A}^T\mathbf{K}\mathbf{A} \, , \quad \text{and} \quad \hat{\mathbf{F}} = \mathbf{A}^T\mathbf{F} \, .$$

This condensed system is solved by NIKE3D's standard nonlinear equation solver, returning rigid-body translation and rotation increments $\Delta\overset{cm}{\mathbf{X}}$ and $\Delta\boldsymbol{\theta}$. These are used to update the center of mass, and finally the nodal coordinates. We employ quaternion parameters to characterize the rotational configuration of the rigid body, allowing finite rotation increments in a single step without inducing deformation.

Implementation

The rigid-body capability was implemented in NIKE3D as a material attribute. This proves convenient for users, since each physical part in a finite element model is typically assigned a unique material number. By selecting the material model #20, the part is identified as a rigid body. Several rigid parts may be included in a single model. If two or more rigid parts share common node points, they are automatically merged and treated as a single rigid body. Rigid bodies may be constructed from any of the element types available in NIKE3D—beams, shells, and/or solids.

All of the boundary conditions in NIKE3D are available for use with rigid bodies. These include prescribed displacement, force, pressure, and body forces. In the case of prescribed displacements, the user must prescribe the motion of the center of mass of the rigid body. Nodal displacement boundary conditions are ignored. Slide surface or contact boundary conditions are also available for use with rigid bodies. Penalty parameters for these interfaces are chosen using reasonable values of elastic moduli, which are entered as rigid-body material properties.

Applications

We anticipate that the rigid material will be heavily used in NIKE3D. Our preliminary analyses have demonstrated the validity of the implementation, and some of the potential cost savings from the rigid assumption. We present three examples using rigid-beam, solid, and shell elements:

(1) *Elastic/rigid beam.* A simple cantilevered beam loaded with a tip moment is shown in **Fig. 1.** The beam is discretized using ten elements. The four center elements are defined as rigid, while the remaining elements have elastic properties. Load is applied in ten increments. The beam centerline is shown for several stages during the loading. The rigidity of the center section of the beam through finite translation and rotation is apparent.

(2) *Hydroformed sheet metal cover.* A hydroforming simulation presented earlier[3] was repeated using the rigid material. **Figure 2** shows an aluminum sheet formed into a three-dimensional pan. In the hydroform process, the male tooling emerges from an initially flat blank holder surface, while subsequent external pressure drives the sheet to conform to the tooling contour. The tooling and blank holder were modeled using one layer of solid, eight-node elements.

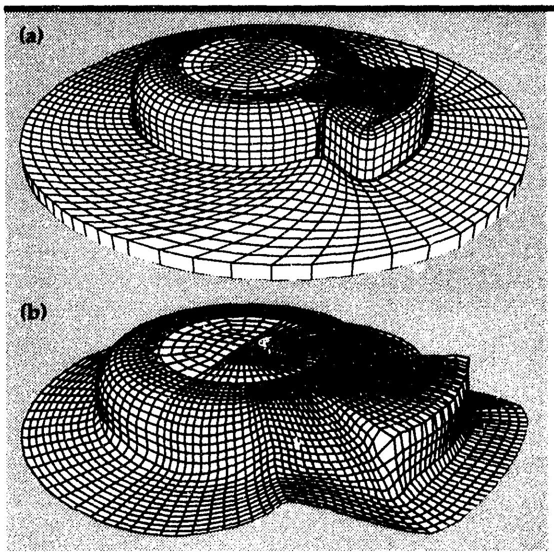

Figure 2. Results from a NIKE3D hydroforming simulation, showing the tooling (a) and finished part (b). The tooling model is constructed using a single layer of eight-node solid elements, making it an ideal candidate for the rigid material, which reduced runtime by over 90%.

Table 1. Cost breakdowns for hydroforming analysis with deformable and rigid tooling. The element-by-element (EBE) iterative linear equation solver[4] was required with the deformable tooling model to fit the linear system into available core memory. Notice that the tooling (brick elements) consumed nearly half of the computer time in the first model, and virtually none when the rigid material was used.

Model Characteristics		
sheet: 1650 shell elements		
tooling: 1550 brick elements		
Linear System	***deformable***	***rigid***
equations	16000	10000
storage (Mword) (direct)	17.0	3.2
(EBE)	6.6	
CPU Cost Breakdown	***deformable***	***rigid***
initialization	0.5%	2.0%
slide surfaces	1%	53%
brick elements	47%	0%
shell elements	12%	21%
linear solver (EBE)	40%	21%
Total cost (Cray Y/MP)	500 min	37 min

Table 1 shows the cost breakdown associated with two NIKE3D simulations: one with deformable tooling, the other with rigid tooling. Applying the rigid material decreased the total simulation time by over 90%, from 500 down to 37 CPU minutes on our Cray Y/MP machine. A large portion of this savings was realized by skipping the processing of the solid elements, which comprise the tooling, saving nearly 250 CPU minutes. The decreased size of the condensed linear system accounted for a savings of nearly 200 CPU minutes. In addition, the rigid tooling simplified enforcement of the contact constraint between the sheet and tooling, since only one side of this interface was deformable. This resulted in fewer iterations in the nonlinear equation solver while searching for equilibrium during each load step.

(3) *Superplastic-formed drape die.* Rigid shell elements were used to model the tooling for analysis of the drape die problem shown in **Fig. 3.** The sheet metal is initially flat, and fully supported around the perimeter of the die cavity. Pressure is applied to the sheet, blowing it into the die cavity. Near the end of the simulation, pressure is increased dramatically to form the sharp corners of the drape die.

Future Work

The rigid-body capability has been implemented in NIKE3D for static analysis only.

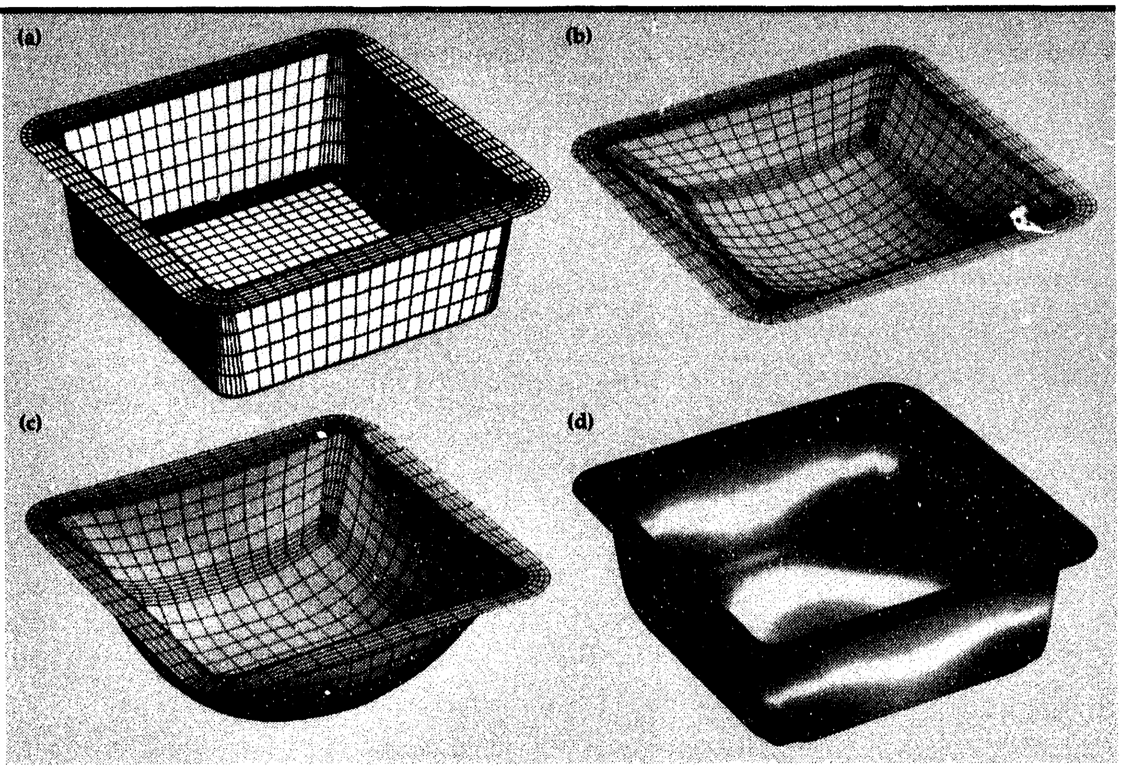

Figure 3. Results from a NIKE3D superplastic forming analysis. The tooling (a) is represented using rigid shell elements. Pressure is applied to the initially flat sheet metal blank, which deforms into the tooling (b), (c). The finished part is shown in (d), where the darkest shade indicates effective plastic strain in excess of 100%.

Additional work is required to extend the capability to dynamic problems. Rigid-body dynamics will require solution of an additional (rotational) equilibrium equation, and generate a non-symmetric set of linear equations. Additional cost savings may be realized by recognizing the presence of rigid bodies in NIKE3D's contact algorithms, which for some cases may reduce their cost by 50%. These enhancements will increase NIKE3D's attractiveness for metal forming and highway barrier post-crash simulations.

---

1. B.N. Maker, R.M. Ferencz, and J.O. Hallquist, *NIKE3D: A Nonlinear, Implicit, Three-Dimensional Finite Element Code for Solid and Structural Mechanics—User Manual*, Lawrence Livermore National Laboratory, Livermore, California, UCRL-MA-105268 (1991).
2. J.C. Simo and L. Vu-Quoc, *CMAME* **66**, 125 (1988).
3. B.N. Maker, *Finite Element Modeling of a Hydroformed Sheet Metal Cover*, Lawrence Livermore National Laboratory, Livermore, California, UCID-21614 (1988).
4. R.M. Ferencz, *Element-by-Element Preconditioning Techniques for Large-Scale, Vectorized Finite Element Analysis in Nonlinear Solid and Structural Mechanics*, Ph.D. Thesis, Stanford University, Palo Alto, California (1989).

HYDRA: A Flow Solver for Transient Incompressible Viscous Flow

Mark A. Christon
Nuclear Explosives Engineering
Mechanical Engineering

This article describes continuing research and code development efforts for a high performance flow solver that addresses the incompressible class of complex geometry transient flow problems that require very-high-resolution meshes. Research efforts for FY-93 have focused on the algorithm-to-architecture mapping issues involving vector, single instruction-multiple data and multiple instruction-multiple data parallel computers, and the investigation of efficient methods for solving the consistent pressure Poisson equation. Several example flow problems are presented to illustrate the class of problems being considered.

Introduction

The simulation of flow fields about vehicles and in turbomachinery remains one of the computational grand challenges for the 1990's. An example of this class of computational fluid dynamics problem is the transient simulation of flow around a full submarine. It is anticipated that more than one million elements will be required to resolve important flow-field features such as shed vortices and regions of separated flow. In addition to the high degree of spatial discretization, the temporal resolution is also demanding, ultimately requiring the effective mapping of the flow-solution algorithm to current vector and parallel supercomputer architectures.

HYDRA is a finite element code that solves the transient, incompressible Navier–Stokes equations, and is based primarily upon the work of Gresho *et al.*[1–4] HYDRA makes use of advanced solution algorithms for both implicit and explicit time integration. In the case of the second-order projection algorithm (implicit),[3,4] a consistent-mass predictor in conjunction with a lumped-mass corrector legitimately decouples the velocity and pressure fields, thereby reducing both memory and CPU requirements relative to traditional, fully coupled solution strategies. The consistent-mass predictor retains phase speed accuracy, while the lumped-mass corrector (projection) maintains a divergence-free velocity field. Both the predictor and the corrector steps are amenable to solution via direct or preconditioned iterative techniques, making it possible to tune the algorithm to the computing platform; i.e., parallel, vector, or super-scalar. The second-order projection algorithm can accurately track shed vortices, and is amenable to the incorporation of either simple or complex (multi-equation) turbulence submodels appropriate for the driving applications.

The explicit solution algorithm[1,2] sacrifices some of the phase speed accuracy of the fractional step algorithm for the sake of minimizing memory and CPU requirements. However, the momentum equations are still decoupled in the explicit algorithm. While both the diffusive and Courant stability limits must be respected in the explicit algorithm, balancing tensor diffusivity is used to ameliorate the restrictive diffusive stability limit in the explicit algorithm. This in combination with single-point integration and hourglass stabilization makes the explicit algorithm computationally efficient, and because of this, the explicit algorithm has been the focus of the initial parallelization efforts.

This work is part of a collaborative effort involving the Mechanical Engineering and Physics Departments, and Military Applications Program at Lawrence Livermore National Laboratory (LLNL). While our effort addresses one of the National Grand Challenges, namely simulating flow fields about vehicles and in turbomachinery, this computational fluid dynamics (CFD) capability is unique because it also finds

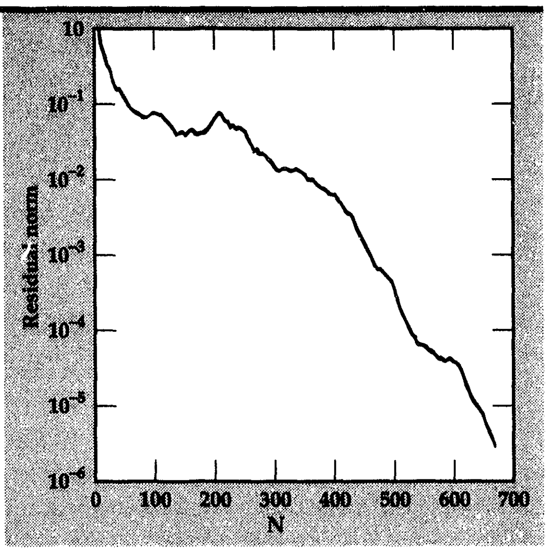

Figure 1. Convergence history of element-by-element, diagonally scaled conjugate gradient.

application within multiple divisions at LLNL, the U.S. Department of Energy, and in U.S. industry. Applications within LLNL include the study of casting processes, heavy-gas dispersion, and flow in the planetary boundary layer. There is also immediate application in industries critical to U.S. competitiveness in the world economy, such as the automotive industry, where CFD is being used to augment engineering design in the areas of vehicle aerodynamics vehicle heating, and air conditioning, and engine and underhood cooling.

Progress

Development efforts over the past year have included the implementation of solution options in HYDRA for the two-dimensional (2-D) transient Navier–Stokes equations, and the 2-D scalar advection-diffusion equation. The 2-D solution options complement HYDRA's existing 3-D capabilities. In addition, the applicability of algebraic multigrid for solution of the pressure Poisson equation (PPE) was investigated. This section summarizes recent developments, and the results of the pressure Poisson solver investigation.

Shown below are the semi-discrete form of the transient, incompressible Navier–Stokes equations with the consistent PPE, and a single scalar advection-diffusion equation for temperature transport.

$$M\dot{\mathbf{u}} + K(\mathbf{u})\mathbf{u} + CP = \mathbf{f} \qquad (1)$$

$$C^T M_l^{-1} CP = C^T M_l^{-1}[\mathbf{f} - K(\mathbf{u})\mathbf{u}] \qquad (2)$$

$$M\dot{T} + K_T(\mathbf{u})T = Q \qquad (3)$$

Efficient solution of the coupled system of equations for velocity and temperature as a function of time relies heavily upon the rapid solution of the consistent PPE (**Eq.2**). Convergence rates using the diagonally scaled conjugate gradient method[5] are typically poor, requiring a large number of iterations. **Figure 1** shows a typical convergence history for the PPE at a given time step for a problem requiring 20,550 elements. A reduction of five orders of magnitude in the residual norm for the PPE required nearly 700 iterations. The computational work to reduce the residual norm is proportional to the number of elements in the mesh, and a very small residual is required to maintain a divergence-free velocity field.

Multigrid methods are efficient schemes for solving partial differential equations on global grids and have been applied to elliptic problems with good success. Most multigrid methods achieve enhanced convergence rates by effectively damping low-amplitude error components on successively coarse grids, thus reducing the total computational work required to solve a problem. On the coarsest grid, the cost of solving the residual problem is negligible, but the finest grid is chosen to deliver a required degree of accuracy.[6] Typical multigrid methods are not readily applicable to the inherently unstructured grids used in the finite element method. For this reason, a 'black box' multigrid solver was used. Algebraic multigrid (AMG) solves a system of equations using the same ideas popularized in conventional multigrid.[7]

Table 1. PPE solvers.

PPE solver	Cycle time (μs/El.–Δt)	Relative cycle time	Relative memory usage
Band Gaussian	22.0	1.0	1.0
PVS Solver	8.8	0.4	0.75
EBE/DSCG	34.5	1.57	0.0
AMG V-Cycle	31.9	1.45	1.3

Table 1 summarizes the results of a comparison among a band Gaussian elimination solver, the PVS (parallel-vector) solver,[8] element-by-element diagonally scaled conjugate gradient (EBE/DSCG), and AMG using a full V-cycle for a 3-D problem requiring 20,550 elements. The memory usage and element cycle time for HYDRA's explicit time-integration scheme has been normalized relative to the band Gaussian solver for the sake of comparison. All computations were performed on a 128-MW CRAY Y-MP. The data in **Table 1** demonstrates no clear advantage when using AMG. The element cycle time delivered by PVS (single processor) essentially matches the solid element cycle time of DYNA3D.

HYDRA makes use of element grouping with both topological and spectral-domain decomposition methods for obtaining vectorization, fine-grained parallelism, or coarse-grained parallelism. Because the element data structures for the vectorized version of HYDRA are adjustable, they are also used for the single instruction-multiple data (SIMD) (CM-2/CM-5) implementation, where element-level operations are performed in a lock-step parallel fashion.

For the CRAY architecture, the vector block size is configured as 128 (twice the length of the vector pipeline). In the case of the CM-2/CM-200, the element block size is configured as a multiple of the minimum virtual processor ratio (four) and the number of available processing elements. For the CM-5, the block size is configured as an integral multiple of the processor vector pipeline length and of the number of available processors enabling processor pipelined operations in conjunction with SIMD parallelism.

The element cycle time on the CM-5 is approximately 72 μs per element per time step using the EBE/DSCG solver with partitioned scatter operations. The use of the partitioned scatter utility provides a factor of approximately three to four speedup over the standard FORTRAN-90 syntax for the scatter-assembly step of the matrix vector multiplication in the EBE/DSCG algorithm. The most current parallelization efforts are being directed towards the use of spectral domain decomposition methods[9] in a message-passing programming model for the Meiko CS-2.

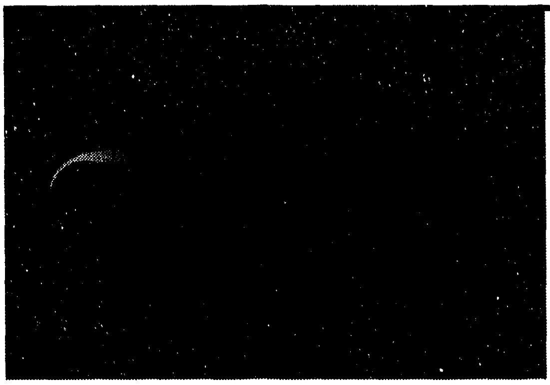

Figure 2. Snapshot of vorticity field during a vortex-shedding cycle.

Applications

In this section, three examples are presented to illustrate HYDRA's application to both 2-D and 3-D transient flow problems. The first example is the flow past a cylinder at Re = 250 based on diameter. This calculation illustrates the formation of the well-known Karman vortex sheet downstream of the cylinder. **Figure 2** shows a snapshot of the vorticity field during a vortex shedding cycle at a simulated time of 90 units. The primary vortex below the cylinder centerline is rotating counter-clockwise, while the upper vortex is rotating clockwise. Also present in this snapshot are attached secondary vortices in the near wake of the cylinder. The onset of shedding occurs after approximately 70 simulated time units, with a resulting Strouhal number of 0.21. A mesh with 16,116 nodes and 15,840 elements was used with a resulting cycle time of 8.77 μs per element time step on a CRAY Y-MP when using the PVS solver.

The second example also illustrates HYDRA's 2-D capability. **Figure 3** shows the vorticity field for the flow about an airfoil at a small angle of attack. The flow problem made use of the Massachusetts Institute of Technology's flapping foil experiment to define the airfoil configuration, angle of attack, and inlet flow conditions. For this flow problem, Re = 10,000 based upon the airfoil

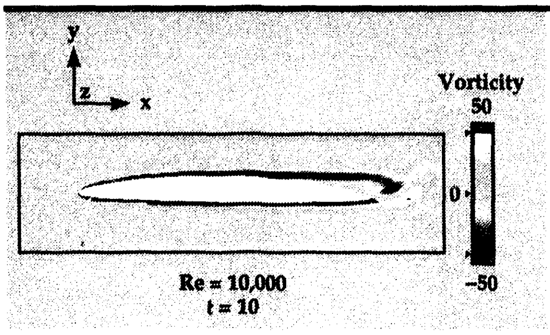

Figure 3. Vorticity field for flow about an airfoil at a small angle of attack.

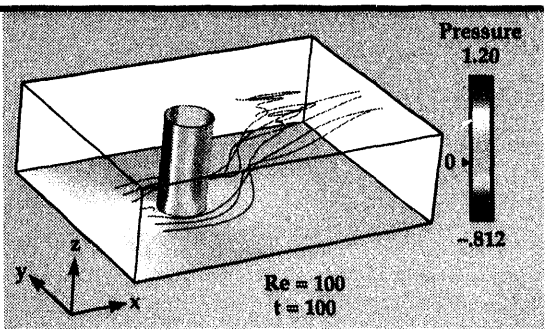

Figure 4. Pressure distribution for flow past a cylinder and a plate.

chord length. The snapshot in time shows a vortex being shed from the trailing edge of the airfoil, and clearly delineates the flow separation point on the upper side of the airfoil. The vortex sheet and separated vortex on the top of the airfoil rotate clockwise, and the vortex sheet and separated vortex on the bottom of the airfoil rotate counter-clockwise.

The final 3-D example consists of the flow past a cylinder and a plate. **Figure 4** shows the pressure distribution on the cylinder and the plate, and particle traces during the initial flow development. The darker regions on the cylinder indicate high pressure corresponding to stagnation points in the flow. After approximately 125 time units, periodic vortex shedding develops downstream of the cylinder, similar to what happens in the 2-D vortex shedding example. The complex particle paths downstream of the cylinder are due to the fully 3-D nature of the vorticity generated at the juncture between the cylinder and the plate and in the wake of the cylinder. Such flow situations are of interest in the onset of turbulent flow, and in turbulent flow situations for submarines.

Future Work

While work will continue on developing efficient methods for solving the PPE, code development efforts will focus on implementing, testing, and evaluating turbulence models. This work will build on the already robust and high-performance capability of HYDRA, to deliver not only a direct numerical simulation capability, but also a turbulence modeling capability.

---

1. P.M. Gresho, S.T. Chan, R.L. Lee, and C.D. Upson, *Int. J. Numer. Methods Fluids* **4**, 557 (1984).
2. P.M. Gresho, S.T. Chan, R.L. Lee, and C.D. Upson, *Int. J. Numer. Methods Fluids* **4**, 619 (1984).
3. P.M Gresho, [Part 1: Theory], *Int. J. Numer. Methods Fluids* **11**, 587 (1990).
4. P.M Gresho and S.T. Chan, [Part 2: Applications], *Int. J. Numer. Methods Fluids* **11**, 621 (1990).
5. F.S. Beckman, *Mathematical Methods for Digital Computers* **1**, 62 (1965).
6. W.L. Briggs, *Multigrid Tutorial*, Lancaster Press (Lancaster, Pennsylvania), 1987.
7. J.W. Ruge and K. Stuben, "Algebraic Multigrid," *Multigrid Methods*, S.F. McCormack (Ed.), Society for Industrial and Applied Mathematics Series: Frontiers in Applied Mathematics (Philadelphia, Pennsylvania), 1987.
8. O.O. Storaasli, D.T. Nguyen, and T.K. Agarwal, "A Parallel-Vector Algorithm for Rapid Structural Analysis on High-Performance Computers," NASA Technical Memorandum 102614 (April 1990).
9. H.D. Simon, *Computing Systems in Engineering* **2**(2), 135 (1991).

Automatic Contact in DYNA3D for Vehicle Crashworthiness

Robert G. Whirley and
Bruce E. Engelmann
Nuclear Explosives Engineering
Mechanical Engineering

This paper presents a new formulation for the automatic definition and treatment of mechanical contact in explicit, nonlinear, finite element analysis. Automatic contact offers the benefits of significantly reduced model construction time and fewer opportunities for user error, but faces significant challenges in reliability and computational costs. We have used a new four-step automatic contact algorithm. Key aspects of the proposed method include (1) automatic identification of adjacent and opposite surfaces in the global search phase, and (2) the use of a smoothly varying surface normal that allows a consistent treatment of shell intersection and corner contact conditions without *ad hoc* rules. Three examples are given to illustrate the performance of the newly proposed algorithm in the public DYNA3D code.

Introduction

Almost all finite element models of vehicle crashworthiness involve contact to some extent. This contact may arise from many sources, including impact of some part of the vehicle with a barrier, or internal contact between two objects within a crushing cavity. Self-contact may arise between two surfaces of a single body due to local buckling and folding deformation. The dynamic nature of vehicle crashworthiness contact problems, combined with their complex geometry and highly nonlinear material behavior, makes them well suited for analysis using an explicit, nonlinear, finite element code such as DYNA3D[1] from Lawrence Livermore National Laboratory.

Contact problems are traditionally modeled within nonlinear finite element codes by defining a list of nodes or segments (element faces) comprising a contact surface, and then defining which contact surfaces should be checked against each other for contact detection. This approach has the advantage that the analyst can use intuition and insight to minimize the size of the contact surface area. Since contact represents a significant fraction of the cost of a crashworthiness analysis, this is an important factor in minimizing the overall cost of running the analysis.

Therefore, the total cost of performing an analysis could be reduced if automatic contact could be used. In addition, automatic contact can lead to improved robustness and reliability arising from fewer errors in complex models constructed by less-experienced analysts. Also, automatic contact may be more robust than defined-surface contact in complex crashworthiness models because objects dynamically deform in ways not anticipated by even an experienced analyst, requiring that the model be augmented with new contact definitions and that the analysis be rerun. Considering these benefits, it now seems natural to examine the question of automating the treatment of contact within a complex crashworthiness model.

There are two major obstacles to widespread use and acceptance of automatic contact treatments. First, concerns exist regarding the robustness of automatic contact on complex crashworthiness problems. These concerns are well founded, since large vehicle models contain some of the most difficult contact situations. However, it is these same large vehicle models which would yield the most time and money savings in reduced model preparation time if automatic contact could be confidently used. The other major concern is that automatic contact will be sufficiently more expensive than defined-surface contact, so that any savings in manpower costs would be eaten away by increased computing costs. In estimating this impact, it should be recognized that once a large model is constructed, it is usually analyzed several times under differing conditions. Thus, increased computer costs accumulated over the life of the analysis model could potentially outweigh savings in reduced model-generation time. In this situation, automatic contact may still yield benefits in reduced turnaround time, even if the cost is marginally increased.

Progress

Overview of New Contact Algorithm

A successful automatic-contact algorithm is more than just a surface generation procedure used in conjunction with standard single-surface or two-surface contact treatments. Automatic contact generates much more complicated surface geometries, and mixtures of shell and solid-element contact segments in complex surface intersections are much more common than in defined-surface contact. This characteristic of automatic contact means that the contact search methodologies must be especially robust to correctly handle these cases, and this requirement has motivated the development of the approach described below.

The contact algorithm presented in this paper is based on the widely used node-on-segment concept. In this approach, the two surfaces to be considered are arbitrarily labeled as a slave surface and a master surface, and slave nodes are prevented from penetrating through the master surface. In the case of automatic contact, each node in the contact surface plays the role of a slave node, and is tested against a set of 'master surface regions,' defined as the subset of the entire contact surface that is spatially close to the slave node. Variations of this node-on-segment approach have been used with success in defined-surface contact implementations in many nonlinear, finite element codes for crashworthiness analysis.

In the following, it is useful to distinguish between the contact *segment*, which is the bilinear isoparametric surface containing the four nodes, and the contact *surface*, which is the actual surface of the material containing any thickness offsets from the segment. A contact surface is therefore coincident with the corresponding contact segment for a solid (continuum) element, but differs by a half-thickness for a shell element.

The complete, automatic contact algorithm proposed here consists of four main steps: (1) contact surface construction, (2) global search, (3) local search, and (4) constraint enforcement. The contact-surface construction step generates contact segments for the entire exterior boundary of each body in the problem. The global search step then locates, for each slave node, a set of master surface regions that are spatially close and likely to be penetrated. For each slave node, the local search then carefully searches the set of master surface regions to find the contact point. Finally, the constraint enforcement step performs the calculations to transmit interface pressure and prevent interpenetration. Each of these steps is further discussed below.

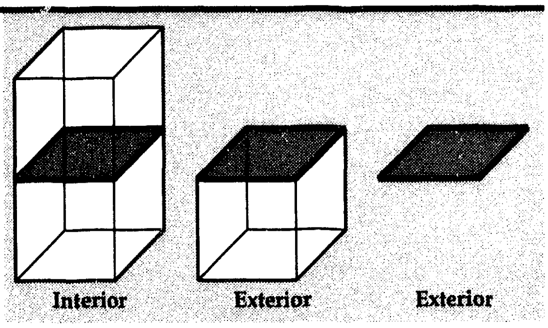

Figure 1. ***Illustration of the process in which the exterior boundary of a body is found as the set of segments with zero or one continuum elements attached.***

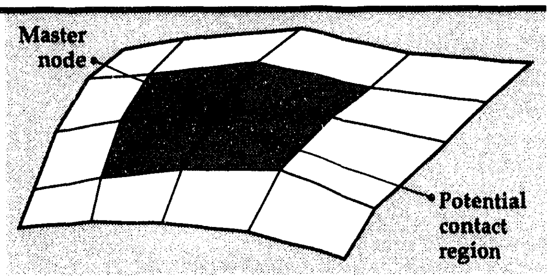

Figure 2. ***Master surface region in segments around a master node.***

Generation of Contact Surface

The first step in automatic contact, performed during initialization, is the generation of a contact surface for the entire outer boundary of the model. The outer boundary is determined using the procedure described for adaptive slide surfaces in Whirley and Engelmann.^{2} In summary, all segments (element faces) in the model that comprise possible contact segments are first given unique numbers (i.e., segments that differ only in the ordering of their bounding nodes are considered the same face). Next, the number of elements attached to each segment is determined; there is a maximum of two 8-node continuum elements and one 4-node shell element that may be attached to a given segment. The exterior segments are then quickly determined as the set of segments with zero or one continuum element attached (see **Fig. 1**). For each exterior segment, a thickness is computed based on the maximum nodal thickness of any node in the segment. An effective stiffness for use in penalty contact calculations is also computed and stored for each segment.

Global Search

The objective of the global search is to identify and store a region of the body where a slave node may potentially come into contact with a master node. This is accomplished by defining a set of 'master nodes' that are close to a given slave node, using the procedure detailed below. The master surface region is then defined as the total area of all contact segments connected to master nodes as depicted in **Fig. 2**.

It is normally not necessary to perform the global search at every time step of an explicit transient analysis. Time steps are usually small due to stability limits, and since nodes undergo only small relative motions within a time step, data from a global search remains valid for several steps thereafter. Experience to date indicates that a fixed global search interval of ten steps works reasonably well, but clearly this interval should be chosen based on problem parameters such as peak velocities and element characteristic lengths. General algorithms for selecting the global search interval will be presented in future publications.

The first step in finding the set of 'close nodes' is to compute a characteristic length for each node on the contact surface. This nodal characteristic length is found as the maximum diagonal of any contact segment connected to a node. Slave nodes that are farther away than this characteristic length from a master node cannot possibly be in contact with any segment connected to the master node.

The next step is to perform a set of three nested, one-dimensional bucket sorts, as proposed in Benson and Hallquist.[3] The nodes from the bucket containing the slave node, along with neighboring buckets, are then combined to form the 'node pool' for that slave node. Next, a 'reduced node pool' is formed by eliminating from the node pool any nodes that either share a contact segment with the slave node or that are farther than their characteristic length away from the slave node, as shown in **Fig. 3**. This second test was motivated by the pinball contact algorithm proposed in Belytschko and Neal.[4]

This reduced node pool still contains more nodes than are needed from the global search to define regions of potential contact. Several approaches may be taken to construct the final 'master node list' from the reduced node pool.

The most robust approach that is still cost-competitive is to select for the master node list a minimal set of nodes in the reduced node pool, which identifies all potential contact regions. To construct this set, it is useful to introduce the concept of *adjacent segment* and *opposite segment* with respect to the slave node. An adjacent segment is a segment from which the path along the surface of the body to reach the slave node results in a monotonically decreasing distance to the slave node. An opposite segment is a segment which either (a) lies on a completely separate body from the slave node, or (b) lies on the same body as the slave node, but from which the path along the surface to reach the slave node does not result in a monotonically decreasing distance function. It is also useful to introduce the concept of nodal tracking: if contact is not found within a given master segment, then all surrounding segments are evaluated to see if one of them is potentially better. If one appears promising, then that segment is checked, and if nothing is found, then its most promising neighbor segment is checked, and so on. This tracking algorithm has the advantage of letting the master segment 'track' the slave node motion, but will fail if it can reach the segment containing the slave node (i.e., if applied to an adjacent region). Tracking is particularly useful in extending the interval between global searches, without sacrificing robustness of the overall contact algorithm. Thus, it is appropriate to store only the one closest node on each opposite contact region, and use the tracking algorithm to locate the contact point if it lies elsewhere in that

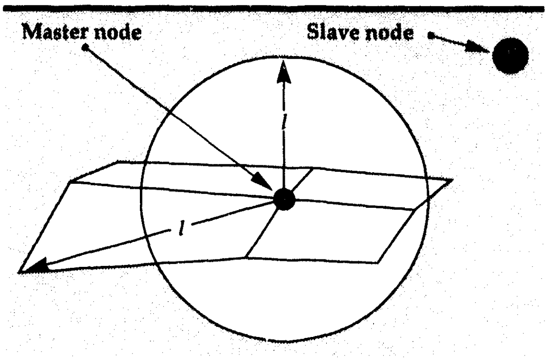

Figure 3. Formation of reduced node pool. Nodal characteristic lengths are used to limit the number of segments considered for potential contact with a slave node. This figure shows how the master node excluded by this test cannot be connected to a segment in contact with the slave node.

region. All reduced pool nodes on an adjacent contact region, which are within their characteristic length of the slave node, are included in the master node list since tracking cannot be used on these nodes.

For each node in the contact surface (i.e., each slave node), the algorithm used to build the master node list from the reduced pool node list is summarized below:

A. Loop over all nodes that share a segment with the slave node; for each of these nodes, use the list of segments connected to this node and flag all nodes of each segment as nontrackable. This prevents the tracking algorithm from finding the slave node in contact with itself.

B. Initialize number of master nodes (for current slave node) to zero.

C. Loop over all nodes in the reduced master node pool.

   1. Choose the node from the reduced master pool that is closest to the slave node, and mark it as taken.

   2. If this node is marked as nontrackable for this slave node and its distance from the slave node is greater than its characteristic length, then discard it and get the next node from the reduced master node pool.

   3. If this node is marked as nontrackable for this slave node but within its characteristic length, then store it as a valid nontrackable master node. Then, mark all nodes of all segments connected to this reduced pool node as nontrackable.

   4. If this node is marked as trackable for this slave node, then just mark all nodes of all segments connected to this reduced pool node as trackable. This means that a closer node has already been stored for this opposite surface so it is not necessary to store this node.

   5. If this node has not yet been marked for this slave node, then this must be the closest node on an opposite surface, so add it to the master node list as a trackable master node. Then, mark as trackable all nodes of all segments connected to this reduced pool node.

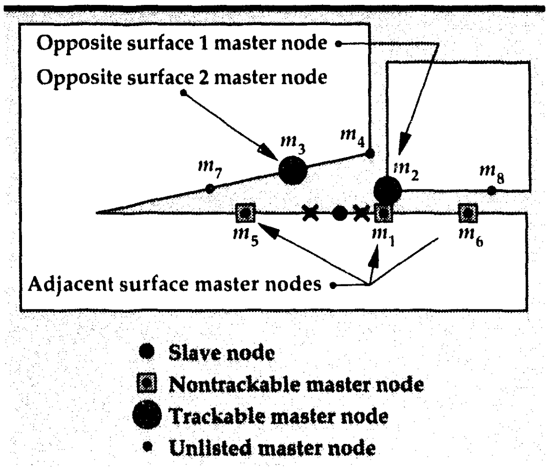

Figure 4. Action of the global search algorithm to construct the master node list from element surfaces with uneven spacing in both self-contact and two-surface contact (shown in two dimensions for clarity). This global search procedure also identifies opposite and adjacent contact regions, and trackable and nontrackable master nodes. Nodes included in the final master node list are enclosed by shaded circles for trackable nodes and shaded squares for nontrackable nodes.

The action of this algorithm is illustrated (in two dimensions for clarity) in **Fig. 4**. The slave node is represented by the large black circle, and reduced pool nodes that are assumed to have passed the characteristic length check are shown as small black circles. First, the algorithm eliminates from consideration reduced pool nodes that share a segment with the slave node, as indicated by the "X" over the two nodes in the figure. Next, reduced pool nodes are considered in order of increasing distance from the slave node, and in **Fig. 4** are sequentially labeled m_1 through m_8. Node m_1 is considered first, and has been marked as a nontrackable node in step A of the algorithm because of its proximity to the slave node. Node m_1 is therefore stored as a nontrackable master node, and all nodes that share a segment with m_1 (only m_6 in this case) are marked as nontrackable. Node m_2 is considered next, and has not yet been marked. Thus, node m_2 must be the closest node on an opposite surface, so it is stored as a trackable master node, and its neighbor nodes (m_8) are marked as trackable. Node m_3 is now considered, and is also found unmarked. Thus, m_3 represents the closest node on another opposite surface, so it is stored as a trackable master node, and its neighbor nodes (m_4 and m_7) are marked as trackable. The global search now considers node m_4, and it is found to have been marked as trackable so it is not stored and its neighbors are marked as trackable (no neighbors of node m_4 are within the characteristic length in this example, so no action is taken). Node m_5 is considered next, and is found marked as a nontrackable node from step A, so it is stored as a nontrackable master node. Node m_6 is treated similarly. Node m_7 is found marked as trackable (by operations on node m_3), and has no neighbors so no action is taken. Node m_8 is also already marked as trackable and has no neighbors, so no action is taken. Thus, the final master node list consists of nontrackable nodes m_1, m_5, and m_6, and trackable nodes m_2 and m_3.

Local Search

The objective of the local search is to determine if a slave node is in contact. If contact is found, then the local search algorithm also determines the precise location of the contact point within a contact segment.

If a slave node was in contact the previous step, then the local search begins by checking the master segment in which contact was previously found. If no contact is found, then all segments connected to the best node of the old contact segment are also checked. It may be desirable to apply an inexpensive screening criterion to these nodes to determine that they are really potential contact regions, before performing the detailed penetration check in each of the surrounding segments. If a slave node was not in contact the previous step, or was in contact but the above procedure did not detect contact this step, then the master node list from the most recent global search is used as the starting point. Nodes are selected from the master node list in order of increasing distance from the slave node. Once a master node is chosen, two inexpensive tests are first applied to verify that the slave node could potentially be in contact with one of the neighboring segments. If these tests indicate contact is possible, then all segments connected to the master node are checked for contact with the slave node. If no contact is found anywhere in this process, then the slave node is determined to not be in contact at all during this step. If contact is found, then one of the constraint enforcement algorithms, described in a subsequent section, is used to satisfy the contact conditions.

The remainder of this section first presents the local segment search procedure. The detailed contact check and contact-point calculation procedure used for candidate master segments is then described.

Local Segment Search. For a slave node that was in contact during the previous time step, the local search begins with a segment contact check in the master segment where contact was found last step. If a contact point is found on the interior of this segment, then the search is terminated unless the segment lies on a shell element and the contact point is within a half-thickness of the segment boundary. For shell segments with contact points within a half-thickness of the boundary, the search procedure must be continued to assure that the correct contact point is located for shell intersections and corners. If a contact point is found and the search terminated, then the selected constraint enforcement procedure is applied.

If contact is not found, the local search is expanded by performing a segment contact check on all segments connected to the closest node of the old contact segment; segments containing the slave node are not considered. At this point, it is possible in general to detect multiple contacts, especially if this region of the structure contains corners or shell-solid element intersections, so a hierarchy must be developed to determine the 'best' segment to use. This hierarchy is:

(1) any shell segment interior contact,

(2) a shell segment tube/pinball contact or solid-segment interior contact,

(3) a solid segment tube/pinball contact.

Segment tube/pinball contact indicates that a contact point was found on a segment boundary, as described in the "Segment Contact Check" section below. If more than one contact segment is found at the same level of the hierarchy, then the segment is chosen containing the contact point closest to the slave node. This corresponds to choosing the least penetrated solid segment or the most penetrated shell segment, and yields the correct behavior in all shell-solid intersections examined thus far. If the best segment found is only a solid segment tube/pinball contact, and other solid contact segments were found that were not in contact, then this situation is treated as no contact. If an acceptable contact segment has been found, then the selected constraint-enforcement procedure is applied, otherwise the local search is continued as described below.

If the slave node was not in contact last step, or if it was in contact but the above local search procedure did not locate a contact segment for this step, then the general local search is used. The master nodes found in the global search are examined in order of increasing distance from the slave node.

Before applying the detailed segment contact check described below, a master node from the

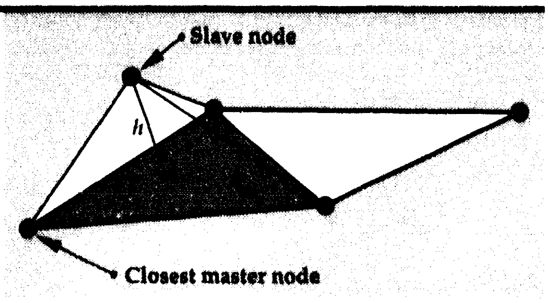

Figure 5. Tetra-hedron test used to screen candidate master nodes taken from the global search before applying the segment contact check.

global search is first subjected to two inexpensive screening tests to verify that it does locate a potential contact region. First, a 'tetrahedron test' is applied to each segment containing the master node, by evaluating the volume of a tetrahedron containing the slave node, the master node, and two adjacent nodes from the master segment, as shown in **Fig. 5**. In the tetrahedron test, a 'characteristic height' h is computed from the tetrahedron using

$$h = 3\frac{V_t}{A_{ts}}, \qquad (1)$$

where V_t is the volume of the tetrahedron and A_{ts} is the area of the tetrahedron base that is part of the master segment (note that h can be negative). The tetrahedron test is passed (i.e., contact is possible) if, for a master segment on a shell element,

$$t_m + t_s \geq h \geq -t_m - t_s\,, \qquad (2)$$

or for a master segment on a solid element,

$$t_s \geq h \geq t_s - t_m\,. \qquad (3)$$

Thus, this test is passed if any tetrahedron yields a sufficiently small positive volume or negative volume to indicate that penetration of the slave node through the master surface is possible.

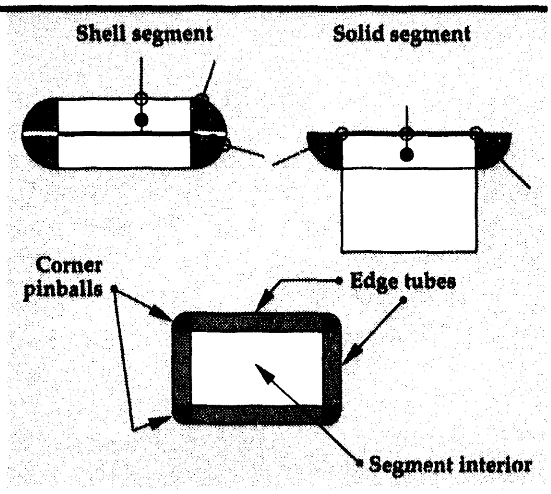

Figure 6. Contact segment showing the use of edge tubes and corner pinballs to compute the return direction vector $\hat{\mathbf{n}}$ if line or point contact is detected.

Next, a 'bounding box' test is performed by finding the minimum and maximum coordinates in each direction of the master node and its two adjacent nodes in the master segment, and adjusting outward using the slave thicknesses. This test is passed if the slave node is contained within any bounding box constructed using the master node.

If a master node passes both of the screening tests, then all segments connected to that master node are evaluated using the segment contact check described below. Once again, it is possible (and likely in complex structures) that multiple contacts will be found, and the hierarchy described above is used to choose the best contact. If more than one contact is found at the same level of the hierarchy, a slightly different resolution procedure is required. First, a comparison vector **v** is constructed: if the slave node was in contact the previous step, **v** is defined as the total assembled contact force on the slave node; if the slave node was not in contact the previous step, then **v** is defined as the relative velocity between the slave node and the master node,

$$\mathbf{v} = \mathbf{v}_s - \mathbf{v}_m\,. \qquad (4)$$

The appropriate contact segment is then chosen that produces the most negative result of $\min(\mathbf{v}\cdot\hat{\mathbf{n}})$.

If contact is found, then the master segment number containing the contact point is stored for the current slave node for use in future steps. The selected constraint enforcement algorithm may then be applied.

Segment Contact Check. The first step in the segment contact check is to compute the local isoparametric coordinates (ξ,η) of the closest point projection of the slave node onto the master segment. Let **L**(ξ,η) represent the parametric equation of the master surface, and let $\mathbf{x}_s$ be a position vector to the slave node; then for a closest-point projection, the contact point (ξ_c,η_c) must satisfy

$$\frac{\partial \mathbf{L}}{\partial \xi}(\xi_c,\eta_c)\cdot\left[\mathbf{x}_s - \mathbf{L}(\xi_c,\eta_c)\right] = 0\,, \quad \text{and} \qquad (5)$$

$$\frac{\partial \mathbf{L}}{\partial \eta}(\xi_c,\eta_c)\cdot\left[\mathbf{x}_s - \mathbf{L}(\xi_c,\eta_c)\right] = 0\,. \qquad (6)$$

These simultaneous equations are solved numerically using Newton iteration.

If the slave node is directly above or penetrated through the master segment, then the above algorithm will yield a contact point (ξ_c,η_c) that is inside the master segment. If the slave node is not directly above or below the master segment, then the above procedure may return a contact point outside the segment. If this happens, then the contact point is returned to the nearest edge of the segment, using a closest-point projection, and a flag is set to indicate potential 'line contact' if only one coordinate was out of the segment, or potential 'point contact' if both coordinates were out of the segment.

Once a contact point has been located, then the return direction is found and a contact check made. First, the thickness of the master segment at the contact point is evaluated based on the isoparametric shape functions and the nodal segment thicknesses. Experience has indicated that accurate calculation of this thickness is essential to the performance of this method. Next, the return direction vector $\hat{\mathbf{n}}$ is calculated. If the contact point was found to be in the interior of the master segment, then the return direction vector is simply the normal to the segment at the contact point. If the contact point was found outside the segment, then the return direction vector is taken from the appropriate 'tube' or 'pinball' placed on the segment edges and corners as shown in **Fig. 6**, and a flag is set. The use of the tube and pinball for evaluation of the return direction vector is an essential component of this algorithm. It assures that the normal varies continuously across shell edges, and greatly improves the basis for treatment of complex corner-contact geometries without *ad hoc* procedures.

The next step in the penetration check is to define a vector **t** from the contact point on the master segment to the slave node,

$$\mathbf{t} = \mathbf{x}_S - \mathbf{x}_C \,, \qquad (7)$$

where $\mathbf{x}_C$ is the position vector to the contact point. Next, compute the distance from the contact point to the slave node, with the sign convention taken from the normal **n** to the master segment (note that this normal may be arbitrarily oriented and is used only as a reference direction),

$$d = \|\mathbf{t}\| \frac{\mathbf{n} \cdot \mathbf{t}}{\|\mathbf{n} \cdot \mathbf{t}\|} \,. \qquad (8)$$

The return direction vector may now be computed by adjusting the orientation of **t** to be the same as that of **n**,

$$\hat{\mathbf{n}} = \mathbf{t} \frac{\mathbf{n} \cdot \mathbf{t}}{\|\mathbf{n} \cdot \mathbf{t}\|} \,. \qquad (9)$$

Note that if $\mathbf{n} \cdot \mathbf{t} = 0$, then the slave node lies on the master segment, $d = 0$, and **Eq. 9** cannot be used to determine the return direction vector. In this (rare) case, set $\hat{\mathbf{n}} = \mathbf{n}$.

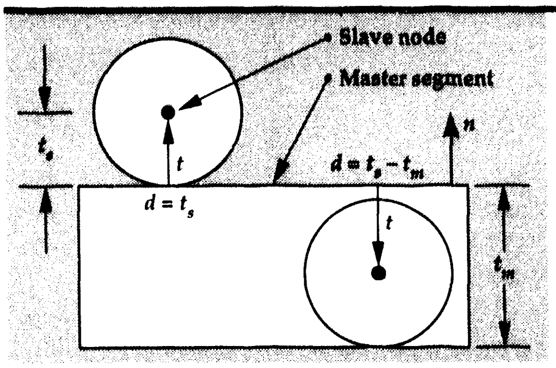

Figure 7. Contact detection test for master surface that lies over a continuum element. Limiting cases of $d = t_s$ (first contact) and $d = t_s - t_m$ (too penetrated to be in contact) are shown.

The contact surface penetration check may now be performed. Let t_m denote the half-thickness of the master surface at the contact point, and let t_s denote the half-thickness of the slave node. If a contact surface lies on shell elements, then the physical thickness is used. For solids, a representative thickness is chosen as the minimum of 60% of the smallest segment diagonal and the ratio of the element volume to the contact segment area.

If the master segment lies on a continuum element, then contact is detected if $d \le t_s$ and $d \ge t_s - t_m$. In this case, the penetration depth, a, considering all surface thicknesses may be computed as

$$a = d - t_s \,, \qquad (10)$$

as shown in **Fig. 7**. If the master segment lies on a shell element, then contact is detected if $d < t_s + t_m$ and $d > -(t_s + t_m)$. The penetration depth, a, may now be found from

$$a = d - \frac{d}{|d|}(t_s + t_m) \,, \qquad (11)$$

as illustrated in **Fig. 8**. Finally, an integer flag is computed that indicates which of the four nodes of the contact segment is closest to the contact point, and whether the contact point is within t_m of the segment

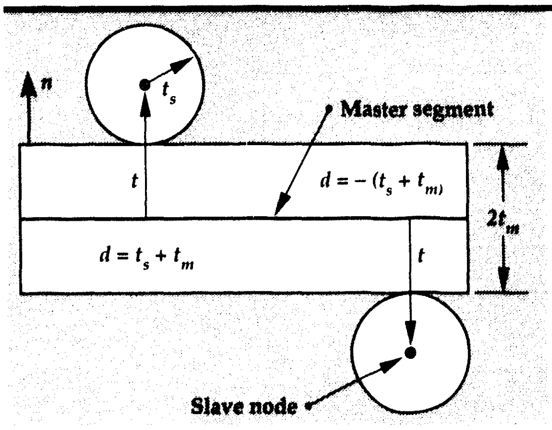

Figure 8. Contact detection and penetration test for master surface that lies over a shell element. Limiting cases of $d = t_s + t_m$ (just in contact with top surface) and $d = -(t_s + t_m)$ (just in contact with bottom surface) are shown.

border. This flag is used in the local segment search when checking old segments, as previously described.

Constraint Enforcement

Once the contact point has been located, a constraint enforcement algorithm must be used to enforce the contact conditions. The proposed formulation has been evaluated in DYNA3D with two constraint enforcement approaches: the penalty method, and a Lagrangian projection method. Each of these is briefly described below.

The penalty method enforces contact conditions by effectively placing a spring between a penetrating slave node and the corresponding master segment. A restoring force f_S is placed on the penetrating slave node to return it to the surface, and an equal and opposite force is placed on the master segment at the contact point. Nodal forces on the master segment are then computed by interpolation using the bilinear shape functions of the contact segment.

The slave node restoring force is found from

$$\mathbf{f}_s = K a \hat{\mathbf{n}} , \qquad (12)$$

where K is a penalty stiffness, a is the penetration depth, and $\hat{\mathbf{n}}$ is the return direction vector discussed above. The penalty stiffness is computed from the bulk modulus of the master segment material and the effective element thickness.

A Lagrangian projection method has been recently implemented for comparison with the penalty method. In contrast to the penalty method, the Lagrangian projection method gives exact results (zero interpenetration) for node-on-node contact, and has been found to generally yield more accurate pressure distributions than the penalty method. This approach is adapted from the forward-Lagrange method of Carpenter, Taylor, and Katona,[5] and is similar to the method proposed in Belytschko and Neal.[4] More experience on production DYNA3D analysis problems will be required to fully evaluate the strengths and weaknesses of each method.

Numerical Examples

Crush of a Wide Flange Beam. This example illustrates the modeling of complex self-contact conditions that occur during the buckling and collapse of a wide flange beam. Although generated somewhat artificially in this example by a pressure load on the upper surface, this type of deformation involving local buckling and sequential contact with several different surfaces separated by corners in a T-intersection, is quite typical of the behavior found in large-scale

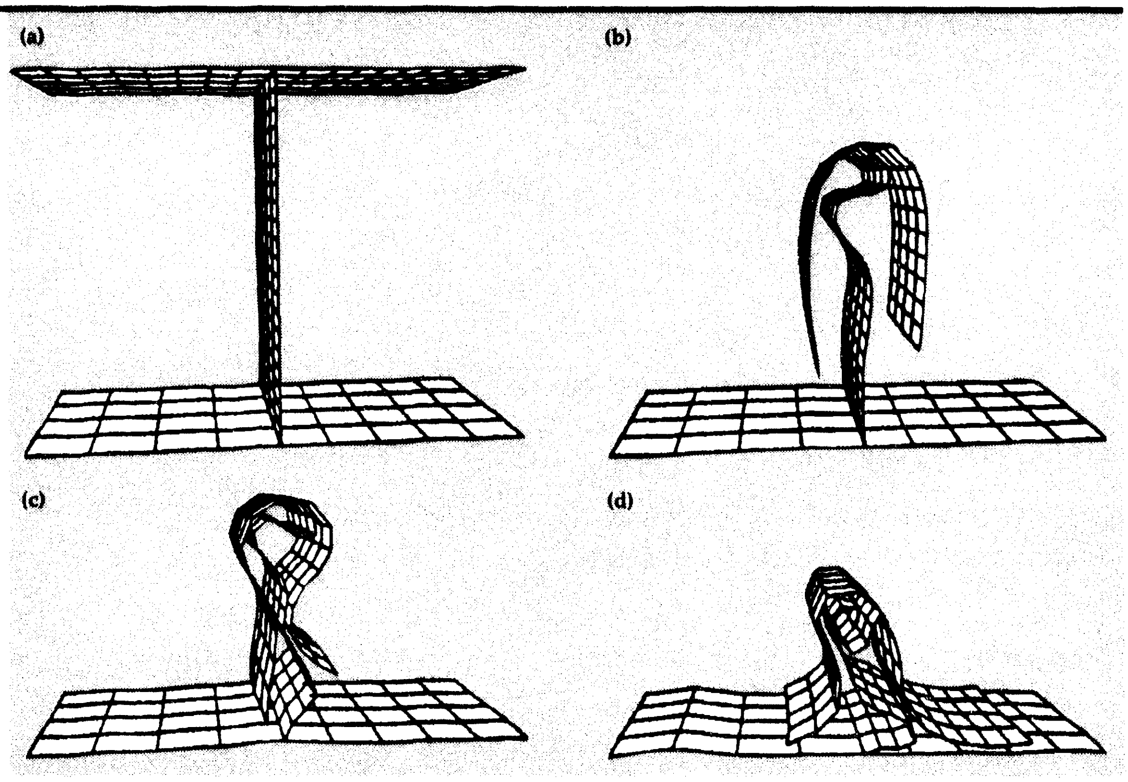

Figure 9. Time sequence of deformed meshes, showing collapse of a wide flange beam including buckling, self-contact, edge contact, and sliding surface contact.

vehicle crashworthiness models. Clearly, the successful solution of this problem requires a robust, self-contact capability along with proper treatment of the corner condition at the intersection of the web and lower flange.

A time sequence of deformed geometry plots is shown in **Fig. 9**. The top left image shows the undeformed geometry, and the top right image shows initial folding of the upper flanges due to a plastic hinge near the web, and initial buckling of the web itself. The lower left image shows a more developed buckle with self-contact in the web, and multiple points of contact between the web and the folded flanges. The upper left flange is about to strike the top surface of the bottom left flange near the T-intersection with the web. The lower right image shows complete collapse of the web and resulting contact with both the lower flange surface and the folded upper flanges in addition to complex self-contact. The folded upper flanges have impacted the lower flange and are sliding outward along its upper surface.

Wall of Blocks. To illustrate the advantages of automatic contact and the robustness of the proposed algorithm, the impact of a stiff elastic block onto the end of a wall of rigid blocks was simulated in DYNA3D. This problem would be very difficult to solve using defined-surface contact because of the large number of disjoint surfaces and potential contact pairs. In addition, this problem illustrates the importance of keeping an arbitrary number of nodes in the global search, since many slave nodes can have several potential master surface regions on opposite surfaces for contact. This problem also exercises most branches of the decision tree for contact at intersections.

The wall of rigid blocks was constructed with a small initial space between blocks to prevent contact from being detected at initialization, and thus forcing the run-time search algorithm to handle the complete contact detection problem. Although the very coarse mesh also causes some inaccuracy in the contact enforcement due to edge-to-edge penetration, this could be removed in a more realistic problem by simply refining the mesh. The coarse mesh is retained in this example, however, because it forces a larger number of master surfaces to be considered for each slave node and increases the number of corner and edge conditions, and is therefore a more demanding test on the search algorithm.

A time sequence of deformed geometry plots is shown in **Fig. 10**. The upper left image shows the initial geometry, and the upper right image shows that the first large motions occur in the loner blocks on the free end, as expected, with small motions evident in the impact area. The lower left and lower right images show the propagation of these disturbances to the first and last columns of blocks, and then to the entire wall. Note that contact conditions are correctly enforced in each of the figures.

Vehicle Frontal Impact. As a final example, the frontal impact of a vehicle front-end frame model was simulated using the proposed automatic contact algorithm in DYNA3D. The impact barrier was modeled using a rigid-wall unilateral-contact option, and the entire model was selected for

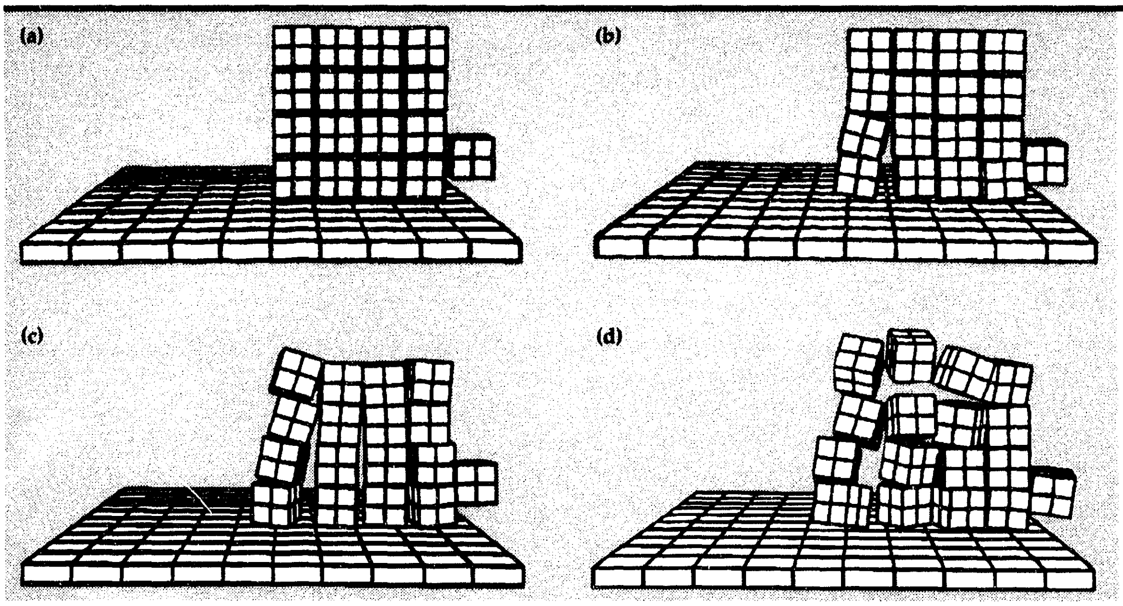

Figure 10. Time sequence showing the impact of a stiff elastic block on a wall of rigid blocks. This problem has a large number of potential contact surface pairs and complex contact conditions on edges and corners, and would be much more difficult to solve using defined-surface contact.

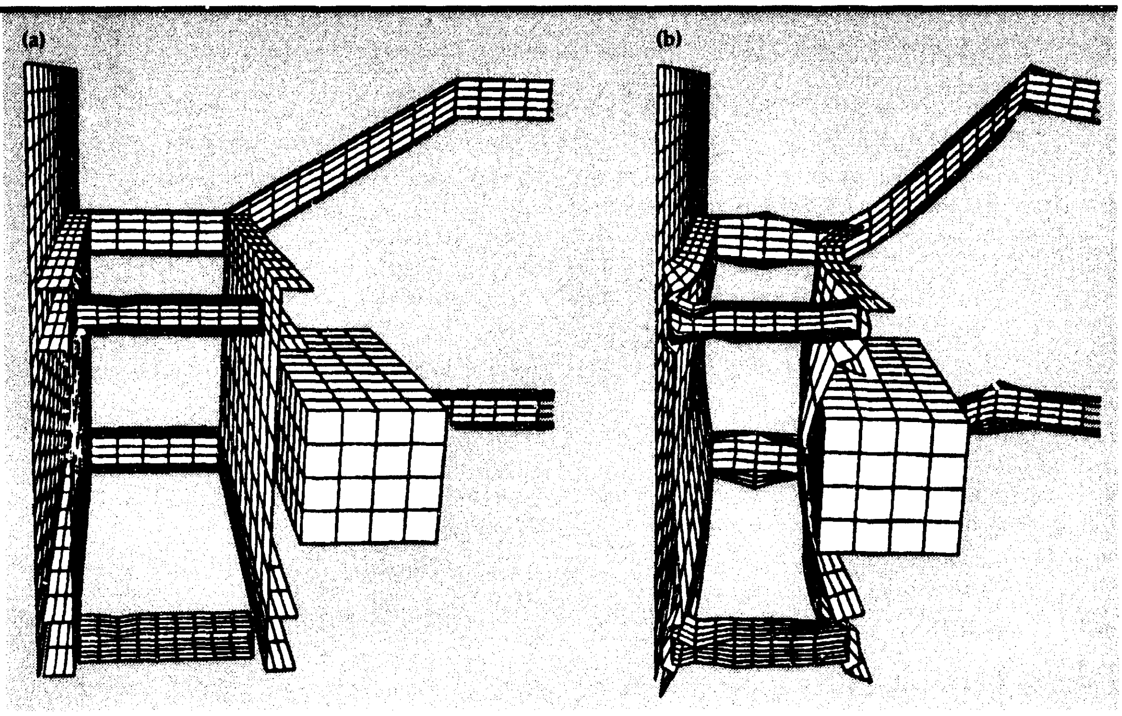

Figure 11. Comparison of original and deformed geometry for a vehicle front-end frame in a frontal impact using automatic contact in DYNA3D. Note the combination of self-contact and two-surface contact in various regions of the model.

automatic contact treatment. The original and deformed meshes of the half-symmetry model are compared in **Fig. 11**. The two front and middle crossmembers are constructed of steel channel sections, with a plate attached between the upper and lower members. The front channel sections collapse upon impact with the barrier, and the lower flange of the upper channel folds under and comes into self-contact. A similar self-contact occurs in the upper flange of the lower channel. Numerous two-surface contacts occur during the impact, including those between the upper and lower energy absorbers and the crossmembers, and between the rigid engine block and the plate between the aft crossmembers.

Although this model contains only 2000 shell elements and is quite coarse, it contains a number of complex shell intersections. The deformed mesh shows that the proposed automatic contact algorithm detected and correctly handled both the self-contact and two-body contact conditions. This analysis required 916 s of CPU time on an IBM RS/6000 model 550 workstation, of which approximately one third was consumed by the automatic contact calculations.

Future Work

Ongoing research is focused on refinement of several aspects of the algorithm. The minimization of the number of nodes stored by the global search is one area of continuing study. Other efforts concern the proper treatment of hypervelocity impact in the segment contact check. Developments are also underway to permit automatic contact to be activated by material groups in DYNA3D to improve efficiency and reduce costs when only a small part of a model is involved in contact. These developments will soon be available in the production version of DYNA3D. Experience and feedback from DYNA3D outside collaborators will no doubt highlight other areas where improvements can be made and are an important part of this research.

Implementation efforts are well underway to complete the vectorization of the algorithm to fully exploit hardware capabilities on CRAY and other vector machines. Strategies are also being developed to implement the automatic contact algorithm on massively parallel computers using a message-passing programming model that accommodates vector hardware on each processor, if available. These vector and parallel implementations in DYNA3D should allow very-high-performance computing resources to be successfully applied to complex vehicle crashworthiness problems.

---

1. R.G. Whirley and B.E. Engelman, *DYNA3D: A Nonlinear, Explicit, Three-Dimensional Finite Element Code for Solid and Structural Mechanics—User Manual,*

Lawrence Livermore National Laboratory, Livermore, California, UCRL-MA-107254 Rev. 1(1993).

2. R.G.Whirley and B.E. Engelmann, "Slidesurfaces with Adaptive New Definitions (SAND) for Transient Analysis," *Proc. Symp. New Methods in Transient Analysis*, ASME AMD **143**, 65 (1992).
3. D.J. Benson and J.O. Hallquist, *Computer Methods in Applied Mechanics and Engineering* **78**, 141 (1990).
4. T. Belytschko and M.O. Neal, "Contact-Impact by the Pinball Algorithm with Penalty, Projection, and Lagrangian Methods," *Proc. Symp. Computational Techniques for Impact, Penetration, and Perforation of Solids*, ASME AMD **103**, 97 (1989).
5. N.J. Carpenter, R.L. Taylor, and M.G. Katona, *Int. J. Numer. Methods Eng.* **32**, 103 (1991).

TRIM3D: A Three-Dimensional Radiation Heat Transfer Code for Participating Media

James D. Maltby
Nuclear Test Engineering
Mechanical Engineering

The TRIM3D radiative heat transfer code for participating media has been successfully linked to the TOPAZ3D conduction code and used to solve state-of-the-art coupled problems. A massively parallel version has been developed for the Meiko CS-2 massively parallel processor with significantly increased performance, and material property input has been automated. Future plans include an interface to the HYDRA fluids code.

Introduction

TRIM3D is a Monte Carlo code for solving thermal radiation heat transfer problems involving a participating medium. The medium may be absorbing, emitting, and scattering, with mixed diffuse and specular boundary conditions. The scattering model is anisotropic, with user input scattering phase functions. The code formulation allows non-homogeneous and non-isothermal problems to be solved. Output of the code is a matrix of Total-Exchange Areas (TEA's), which describe the radiant interchange among all surfaces and gas volumes in the geometry. A Monte Carlo approach to radiation heat transfer problems without a participating medium has proved very successful, resulting in the computer code MONT3D.[1] TRIM3D was developed during FY-92 and given preliminary testing and verification.

Progress

Coupling with TOPAZ3D

The TEA matrix produced by TRIM3D is dimensionless and does not contain any temperature information, so it must be passed to another code for flux calculations. A special version of the TOPAZ3D heat transfer code was developed during FY-93 that uses the TEA matrix to solve coupled conduction/convection problems. The basic elements used in TRIM3D are four node shells and eight node bricks. TOPAZ3D uses the TEA matrix to calculate surface radiative fluxes and volumetric heat generation rates. Since this matrix is temperature independent, boundary conditions and temperatures in an analysis may be changed without re-running TRIM3D. A similar approach has been very successful with MONT3D and TOPAZ3D, resulting in a large savings in computer time.

Material Property Generation

Also during FY-93, automated generation of banded gas properties was developed, using an Elasser narrow-band model.[2] These routines generate banded spectral data in TRIM3D format for

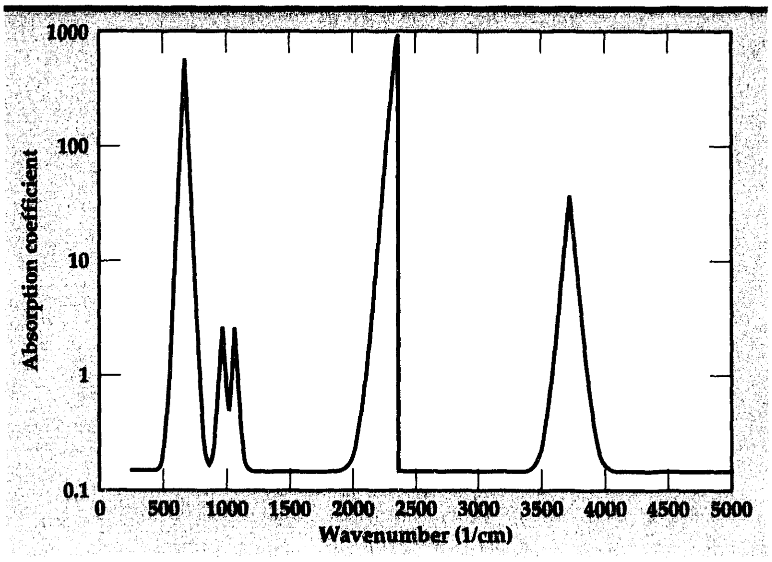

Figure 1. *Absorption coefficient vs wavenumber for a carbon dioxide, nitrogen, and carbon particle mixture.*

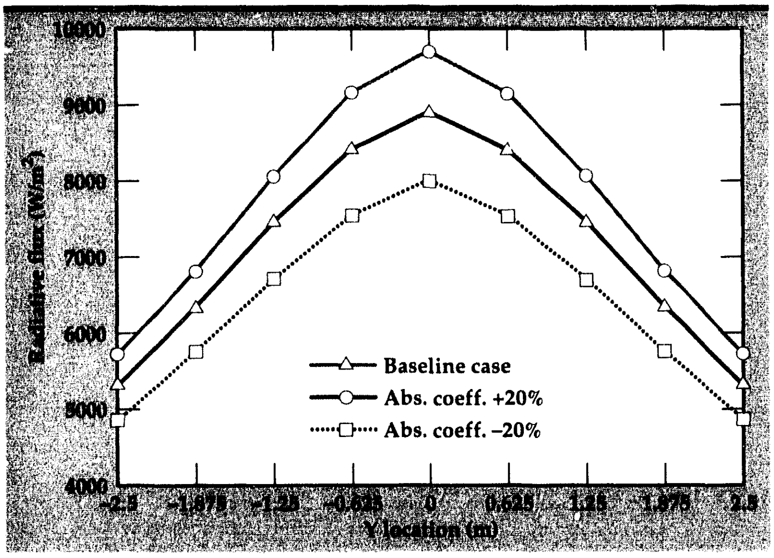

Figure 2. Changes in surface radiative flux distribution with absorption coefficient.

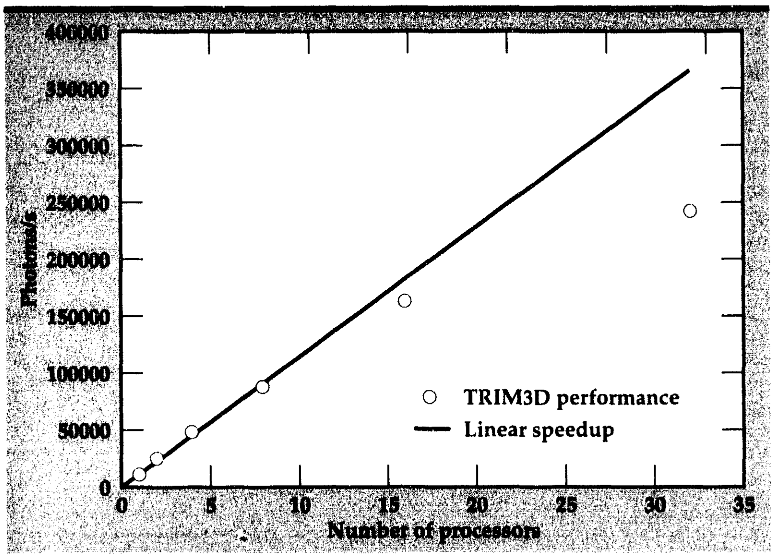

Figure 3. Parallel speedup on 32-node Meiko CS-2 parallel processor.

mixtures of absorbing gases, non-absorbing gases, and scattering particles. For inhomogenous problems, multiple material property data sets can be generated automatically. The spectral variations for gases can be quite complex, as shown in **Fig. 1**. **Figure 1** shows absorption coefficient vs wavenumber for a mixture of carbon dioxide, nitrogen, and carbon particles, generated using the gas property routines. Using user-input band wavelength breakpoints, the curve is integrated to produce the required coefficients. Additional routines in the same package generate curves of scattering phase function as a function of angle, so an entire material property data set may be produced in an automated manner.

Code Verification

One of the difficulties in verifying a participating medium code is the small number of problems with analytic solutions that exercise all the code features. To address this problem, a symposium was held at the 1992 ASME Heat Transfer Conference to assess the current capability for solving non-gray, anisotropically scattering, multi-dimensional radiation problems.[3] Thirty-four benchmark problems ranging from one to three dimensions, at optical depths from 0.1 to 10, were specified. TRIM3D was used to solve the entire benchmark suite, including a non-isothermal, inhomogenous, three-dimensional problem with non-black walls. Agreement with other techniques presented at the 1993 symposium was good, though only one other code (also Monte Carlo) was able to complete the entire range of problems.[4] TRIM3D was the only code demonstrated that coupled directly to a general purpose heat transfer code (TOPAZ3D) for solving multimode problems.

The 1994 Symposium on Radiative Heat Transfer in Participating Media will focus on material property effects. Radiative material properties are difficult to measure, particularly for gases and semi-transparent materials. TRIM3D was used to test the sensitivity of the most detailed benchmark problem to errors in material property inputs. Absorption coefficient, scattering coefficient, surface reflectivity, scattering phase function, and band wavelength breakpoints were all varied, and ranked according to magnitude of effect. **Figure 2** shows the effect of varying the absorption coefficient on the predicted wall radiative flux. These results will be presented at the 1994 symposium.

Parallel Processing on the Meiko CS-2

One problem with Monte Carlo methods is the large amount of computer time required to achieve good statistical accuracy. The benchmark problems above required 40 to 50 hours of computation time on a Sun SPARC-2 workstation. The random branching required by the simulation makes vectorization difficult and inefficient, so the CRAY's do not provide significant additional speedup. However, the statistical independence of the photon trials makes TRIM3D an ideal candidate for parallelization using a multiple instruction-multiple data or single program-multiple data (SPMD) model.

A version of TRIM3D for the Meiko CS-2 massively parallel processor was developed using an SPMD model and the NX-2 message passing library. The parallel speedup achieved was good, as seen in **Fig. 3**. A parallel efficiency of 67% was achieved on a 32-node partition, reducing the time required for calculating one of the ASME benchmark problems from 46 h on the SPARC-2 to 1 h 20 min on the Meiko.

Future Work

In FY-94, we will concentrate on developing applications for TRIM3D and a manual. Additional features are planned for the production version of the code to simplify the solution of large problems and make the code more 'user friendly.'

Common industrial furnace applications often require coupled calculation of participating radiative heat transfer and fluid flow, so a link to the HYDRA fluid dynamics code will be developed. Since the typical mesh densities required for radiative heat transfer analyses and fluid flow analyses differ greatly, the primary task will be to develop a formulation for mapping between the two grids. The mapping algorithm will have to be both accurate and efficient enough to be executed every fluid time step. Once developed, the approach will have to be verified against known solutions. If this is successful, the resulting coupled capability will be unique.

---

1. J.D. Maltby and P.J. Burns, *Numer. Heat Transfer* 9(2), (1991).
2. R. Siegel and J.R. Howell, *Thermal Radiation Heat Transfer,* 4th ed., Hemisphere Publishing Corp. (New York, New York), 1992.
3. T.W. Tong and R.D. Skyocypec, "Summary on Comparison of Radiative Heat Transfer Solutions for a Specified Problem," *Developments in Radiative Heat Transfer,* HTD **203** (New York, New York), 1992.
4. J.T. Farmer and J.R. Howell, "Monte Carlo Solution of Radiative Heat Transfer in a Three-Dimensional Enclosure with an Anisotropically Scattering, Spectrally Dependent, Inhomogenous Medium," *Developments in Radiative Heat Transfer,* HTD **203** (New York, New York), 1992.

GRIZ: Visualization of Finite Element Analysis Results on Unstructured Grids

Donald Dovey and
Michael D. Loomis
Applications Systems Division
Computation Directorate

GRIZ is a general-purpose post-processing application that supports interactive visualization of finite element analysis results on three-dimensional unstructured grids. GRIZ includes direct-to-videodisc animation capabilities and is being used as a production tool for creating engineering animations.

Introduction

Several developments have contributed to a need for a 'next generation' finite-element post-processing application. One is the availability of workstations with hardware acceleration for rendering three-dimensional (3-D) models, and a concurrent requirement by analysts for fast rendering of large problems to allow interactive interrogation of analysis results. A second factor is the recent development of new algorithms for computer visualization of volume data, including cut planes, isosurfaces (surfaces of constant result value in the mesh volume), vector grids, particle-following techniques for vector fields, and 'volume rendering' techniques. Finally, the ability to output animations to video makes the results of large time-dependent analyses more readily accessible to peers of the analyst and to non-analysts.

GRIZ is a new post-processing application that has been developed to meet the needs mentioned above. It provides a variety of visualization techniques, supports interactive analysis of results, includes libraries for computing derived variables, and provides animation capabilities and video output.

Progress

A variety of finite element analysis codes is supported by the Methods Development Group in the Mechanical Engineering Department of Lawrence Livermore National Laboratory (LLNL). The codes include non-linear solid mechanics codes (DYNA3D, NIKE3D), a heat-transfer code (TOPAZ3D), an acoustics code (PING), and a fluid analysis code (HYDRA). Currently, GRIZ provides post-processing capability for all these codes, as well as for electromagnetics codes (DSI3D, AMOS). The analysis results for all of these programs are provided on unstructured grids. The variety of applications supported has required that GRIZ have a general and flexible design.

Program Design

The GRIZ software can be loosely broken down into components: a file reader module, a user interface module, a derived variable module, a visualization algorithm module, a display module, and a video output module. These modules are described in the paragraphs below.

The file reader module reads in data in the 'plotfile' format used at LLNL. Finite element mesh connectivity and result data at successive states are included in this data.

The user interface is implemented using the Motif interface widget toolkit. A menu system, a command dialog for typing in commands, and view control widgets are included in the interface, along with an area for displaying the mesh. Moving the mouse in the display area with a mouse button held down can be used to interactively change the view. The mouse is also used to graphically pick on individual elements or nodes, which then display their result values.

Analysis codes like DYNA and NIKE output a limited number of variables at each state, from which numerous derived variables can be

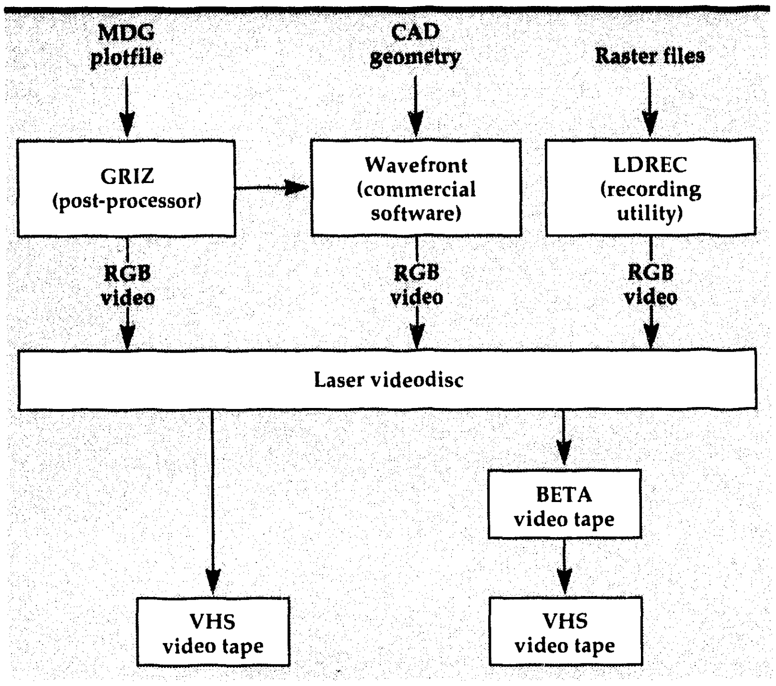

Figure 1. Data flow in animation production.

computed. GRIZ includes a module for calculating derived variables for the DYNA3D and NIKE3D codes. This module can be extended to include derived variable computations for other analysis codes.

The visualization module implements a number of visualization techniques, including color-mapped fringe plots of result values on the surface of a mesh, contour plots, cut planes, isosurfaces, vector grids, and particle traces. Capabilities for computing result minimums and maximums over all the states, setting cutoff values for the results display, and translating individual materials and turning on and off the display of individual materials in the mesh are also provided.

The visualization module uses a sorting algorithm to determine which faces are on the interior and exterior of the mesh.[1] Once the exterior faces are determined, a face-normal smoothing algorithm is applied to give the faceted surface of the mesh the appearance of being smooth. This smoothing helps the user understand the shape of the model and avoids distracting lighting artifacts.

The visualization module also provides the ability to interpolate result data between states over time. This ability is necessary in order to support smooth animation. In DYNA and NIKE analyses, material failure is simulated by deleting elements over the course of the simulation. GRIZ correctly handles these element deletions during interpolation and recomputes the visible surface of the mesh.

The display module provides rendering of the mesh and uses the SGI library for its display routines. Results are displayed on 3-D (volume), two-dimensional (2-D) (shell), and one-dimensional (beam) elements. On SGI workstations, the rendering is performed in hardware, which means that the image can be updated rapidly in comparison with software implementations. For small models, the data can be animated interactively.

The video output module permits the user to save individual frames of an animation directly to a write-once read-many (WORM) videodisc via a hardware video interface installed on the workstation. A single videodisc can hold 24 minutes of animation per side.

Animation Production

The GRIZ command set allows animations of substantial complexity to be generated when commands are scripted into a text file for batch processing. To facilitate this process, we have developed a script generator called FRAME, which uses key frame techniques to produce an explicit, frame-by-frame description of an animation in terms of GRIZ commands. FRAME allows beginning and ending values to be specified for the arguments of virtually any GRIZ command, and then smoothly interpolates those values over a given range of frames. The interpolation may be linear or cosine-based, or it may use ease-in and ease-out accelerations.

With the speed of hardware rendering, GRIZ may run FRAME-generated scripts and save several thousand frames directly to videodisc in one overnight run. As such, it is the preferred method of animation when production schedules are tight or when magnetic disk space is limited. On the other hand, one of the disadvantages of hardware rendering is that the rendering resolution is limited, which can sometimes lead to annoying aliasing artifacts during video playback.

To avoid this problem, an alternative animation path may be taken by using Wavefront rendering and animation software. GRIZ provides an interface for outputting data to Wavefront, where it can be rendered at high-resolution with shadows and reflections, transparency, texture maps, and full anti-aliasing. The tradeoff lies in the additional time needed to create the animation. Wavefront provides an additional benefit in that it enables translated computer-aided design data to be imported for rendering and animation. This provides another source of engineering data to be used along with analysis results and widens the potential client base for future animation production.

A third method of animation recording involves the use of a utility called LDREC (LaserDisc RECord). LDREC simply takes a series of existing raster images (in Silicon Graphics' native .rgb format), loads them one by one into the frame buffer, and records them in sequence onto the videodisc. LDREC was written primarily to provide an alternate way to record images that derive from applications other than those supported by GRIZ or Wavefront.

Once the individual frames of an animation have been recorded to videodisc, the various sequences and title frames comprising the final video may be edited together. While the final product will almost always be VHS video cassettes because of the widespread availability of VCR's in that format, there are generally two different ways to approach this. If the number of final tape copies is few and complex video editing is not required, it is quite possible to obtain excellent results by editing directly from the videodisc to VHS using the hardware available in the video rack. Otherwise, the best approach is to use a more fully equipped video editing facility and create a master tape in Beta format, from which an unlimited number of VHS dubs can be made.

Future Work

Continuing work on GRIZ will focus on extending and enhancing its current capabilities. An immediate task is to define and support a more flexible, extendible, machine-portable data format for analysis results. Support for time-history data (result data that is output at finer intervals for a selected set of elements or nodes) and more sophisticated X–Y plotting capabilities should be incorporated. Code-specific capabilities will continue to be implemented. Examples include flow ribbons for fluids codes, interface force display for solid mechanics codes, particle 'flight paths' for Monte Carlo codes, and derived variable libraries for the 2-D codes, DYNA2D and NIKE2D.

With the current trend of sending analysis codes to massively parallel processors, we will need to develop algorithms for performing efficient visualization of very large datasets; i.e., greater than one million elements. The current software should provide a good platform for developing algorithms to handle this type of data.

---

1. M.A. Christon and Thomas Spelce, "Visualization of High Resolution, Three-Dimensional, Nonlinear Finite Element Analyses," *Proc., Visualization 1992*, Boston, Massachusetts, 324 (October 1992).

PING: An Explicit Finite Element Code for Linear Structural Acoustics

Mark A. Christon
Nuclear Explosives Engineering
Mechanical Engineering

PING is an explicit finite element code for analyzing coupled acoustic fluid-structure interaction, where the propagation of sound waves is of primary interest in the fluid. This report consists of an overview of the coupled fluid-structure formulation, a discussion of the solution strategy, and an example of PING's application to the simulation of an acoustically excited elastic structure.

Introduction

The identification and reduction of the acoustic signature of submerged vehicles and structures are of the utmost importance in naval underwater vehicle design and in underwater exploration and recovery. A clear understanding of the interrelationship between hull geometry and acoustic signature is crucial to the design of 'quiet' hulls, and to the related goal of detecting and accurately identifying underwater vehicles and structures. The loading of a submerged body by either an external-incident acoustic pressure wave, or by internal mechanisms such as rotating machinery, excites both the elastic body and the fluid that surrounds it. The surrounding fluid, in effect, contributes to the overall inertia and effective elasticity of the submerged body, thereby changing the way it responds to either internal or external forces. Conversely, the submerged body influences the behavior of propagating disturbances in the surrounding fluid, which provides a feedback mechanism between the fluid and the body. The structural acoustics problem remains one of the computational 'grand challenges' facing the U.S., as defined under the Federal High Performance Computing Program.

To attack this problem, our structural acoustics effort uses an interdisciplinary team seeking to identify the fundamentally critical fluid-structure acoustic phenomena required to accurately simulate the behavior of submerged vehicles and structures and their associated acoustic characteristics. The three fundamental issues being addressed are (1) surface curvature modeling discrepancies and their ramifications for the transfer of energy between the submerged body and the fluid,[1] (2) the proliferation of propagating structural shell frequencies,[2] and (3) the unpredictable acoustic radiation from submerged stiffened panels.

Our research to date has relied on both analytical and numerical models to investigate the time-dependent behavior of complex submerged bodies with and without internal structure, and has yielded a time-domain analysis tool for performing both two-dimensional (2-D) and three-dimensional (3-D) structural acoustics simulations. The analysis tool, PING,[3] is an explicit, highly vectorized, acoustic fluid solver intimately coupled with the structural analysis code NIKE3D.[4] While the underwater applications have driven the PING development effort, the code is general purpose, and can treat virtually any acoustic working fluid and any combination of linear elastic structural materials.

The ultimate goal for this work is to apply the concepts and simulation tools developed by the working group to the evaluation of newly proposed naval design concepts for their acoustic performance. The following sections describe the progress to date on the structural acoustics problem, and the application of PING to a fluid-loaded, 1/24th-scale, streamline elastic hull.

Progress

PING is an explicit finite element code for analyzing 2-D and 3-D, transient, linear, fluid-structure interaction problems where the fluid may be represented by a perturbation pressure field. PING addresses the class of problems usually referred to as structural acoustics problems in the time domain. PING represents the union of NIKE3D and a stand-alone, explicit finite element code for solving the second order, scalar wave equation. NIKE3D, in effect, provides a base library of structural mechanics code. NIKE3D is a 3-D, implicit, nonlinear finite element code used for analyzing both static and moderate-rate dynamic problems in structural mechanics.[4]

Structural acoustics problems often consider the coupled fluid-structure interaction in the mid-range frequencies. PING has been specifically designed for such problems where spatial scales vary over a broad range, internal structural ribs and stiffeners are significant, and complex external structural geometry requires high-resolution meshes. While many researchers in the field of structural acoustics rely on frequency domain methods and boundary integral or boundary element techniques, PING is a time domain code that makes use of a 'fluid island' coupled to the structure. It has been estimated that the cost of computing a solution containing a moderate number of frequencies for a 3-D structure in the frequency domain can be reduced by about three orders of magnitude by using time domain methods. Furthermore, the time domain approach promises great flexibility because of its ability to treat non-periodic driving functions and broad-band pulsed input signals, allowing many frequencies to be captured in a single analysis. With PING, time history data such as fluid pressure, structural displacements, and accelerations can be measured directly.

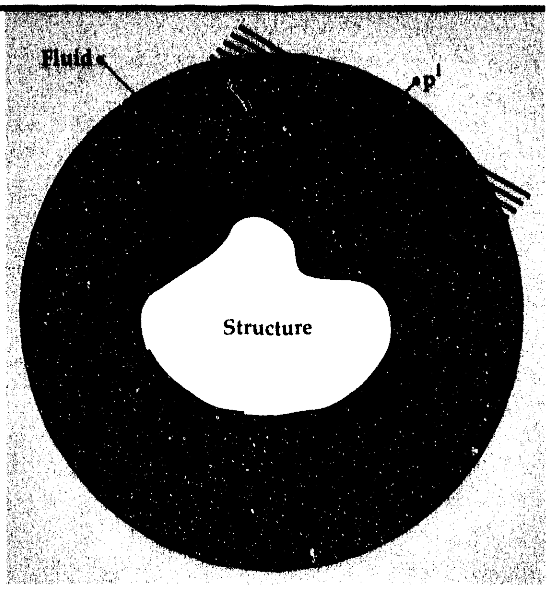

Figure 1. Fluid-structure domains with a single-incident plane wave.

Although NIKE3D treats both geometric and material nonlinearity, PING provides the capability of treating only linear structures, linear elastic materials, and a linear, fluid, perturbation pressure field. Although PING relies heavily on NIKE3D's rich suite of structural mechanics capabilities, NIKE3D's Newmark time integration scheme has been reformulated for PING in terms of total displacements. This fundamental modification follows Goudreau[5] and enables a single factorization of the structural equations to be used with many resolves in the coupled fluid-structure time marching scheme.

The initial software design trade-offs made in PING were based on the philosophy of maximizing performance and the scalability of problem size. For this reason, the acoustic fluid module has been fully vectorized with the hope of providing a balanced computational tool requiring nearly equal resources for the structural and the fluid components of the model.

Governing Equations

The equations for conservation of linear momentum for the structure are presented in indicial notation in **Eq. 1**, with the associated specification of either traction on Γ_1 as in **Eq. 2**, or displacement on Γ_2 as in **Eq. 3** and shown in **Fig. 1**. Also shown in **Fig. 1** is the fluid domain with its outer radiating boundary Γ_3 and a single incident plane wave, p_i.

$$\rho_s \ddot{\mathbf{u}}_i = \sigma_{ij,i} + \mathbf{f}_i \quad (1)$$

$$\sigma_{ij}\mathbf{n}_j = t_i \quad (2)$$

$$\mathbf{u}_i = \hat{\mathbf{u}}_i \ , \quad (3)$$

where σ is the Cauchy stress tensor, ρ_s is the density of the structure, **u** is the displacement, **f** is the body force per unit volume, $\mathbf{n}_i$ are the components of the unit outward surface normal vector, and $\ddot{}$ denotes differentiation with respect to time.

For the fluid, it is assumed that a single scalar field representing a perturbation pressure field may be used, i.e., the fluid is inviscid and isentropic, and only small disturbances are considered. While several forms of the linear wave equation for an acoustic fluid could be used,[6] the formulation in terms of the perturbation pressure is the most convenient for the coupled fluid-structure problem. By comparison, the acoustic propagation problem in the fluid could be formulated in terms of displacements,[7] but the use of a scalar field affects a 3:1 reduction in computational work.

In order to account for multiple incident plane waves, **Eq. 4** is used for the superposition of the perturbation pressure, p, with known incident pressure waves in order to represent the total perturbation pressure, Δp. In **Eq. 4**, N_{pw} is the number of incident plane waves. Using **Eq. 4**, it is possible to write the wave equation in terms of the perturbation pressure as shown in **Eq. 5**:

$$\delta p = \sum_{k=1}^{N_{pw}} p_k^i + p \tag{4}$$

$$\frac{1}{c^2}\ddot{p} = \nabla^2 p \tag{5}$$

where c is the fluid sonic velocity.

A Sommerfeld condition that relates the rate of change in the pressure to its normal derivative on the outer boundary of the fluid domain, as shown in **Eq. 6**, is used for the outer boundary radiation condition. This boundary condition allows the far-field to be approximated using a truncated computational domain with a non-reflecting outer boundary. **Equation 6** is exact in a one-dimensional sense for the wave equation, but it is only approximate in a multi-dimensional sense. Currently, more sophisticated (higher-order) treatments of this boundary are under consideration,[8,9] but are not yet implemented.

$$\frac{\partial p}{\partial t} = -c\frac{\partial p}{\partial \mathbf{n}} \tag{6}$$

With the proper specification of initial conditions (velocity and acceleration for the structure, and first- and second-time derivatives of the pressure for the fluid), the following interface conditions complete the definition of the structural acoustics problem. **Figure 2** illustrates the geometrical relationship between the fluid and structural domains via the outward normal, as presented in **Eq. 7**. **Equation 8** is a kinematic constraint relating the velocity of the structure to the fluid velocity. Using **Eqs. 7** and **8**, and performing a momentum balance in the fluid at the fluid-structure interface, **Eq. 9** may be obtained in terms of the total perturbation pressure. This equation relates the structural acceleration to the normal gradient of the pressure in the fluid, and is a key component in the fluid-structure coupling.

$$\mathbf{n}_s = -\mathbf{n}_f \tag{7}$$

$$\dot{\mathbf{u}}_s = \mathbf{v} \tag{8}$$

$$\mathbf{n}_s \cdot \ddot{\mathbf{u}}_s = -\frac{1}{\rho_f}\frac{\partial \delta p}{\partial \mathbf{n}_f} \tag{9}$$

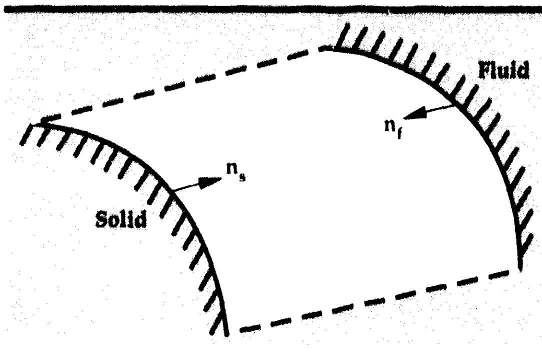

Figure 2. Interface normals.

The semi-discrete equations for the problem are obtained by proceeding with the standard Galerkin discretization.[10] The coupled formulation yields several additional terms. The first arises from the application of the Sommerfeld boundary condition on Γ_3 in **Fig. 3**. The forcing term for the fluid must include the gradient of the incoming pressure waves on the structure for a single-incident plane wave. Further, the formulation must account for the projection of the pressure forces from the fluid onto the structure, and the conjugate projection of acceleration onto the fluid at the interface. The additional pressure forces arise from the use of superposition in the simplification of the treatment of incident plane waves, and augment the force on the structure due to the perturbation pressure.

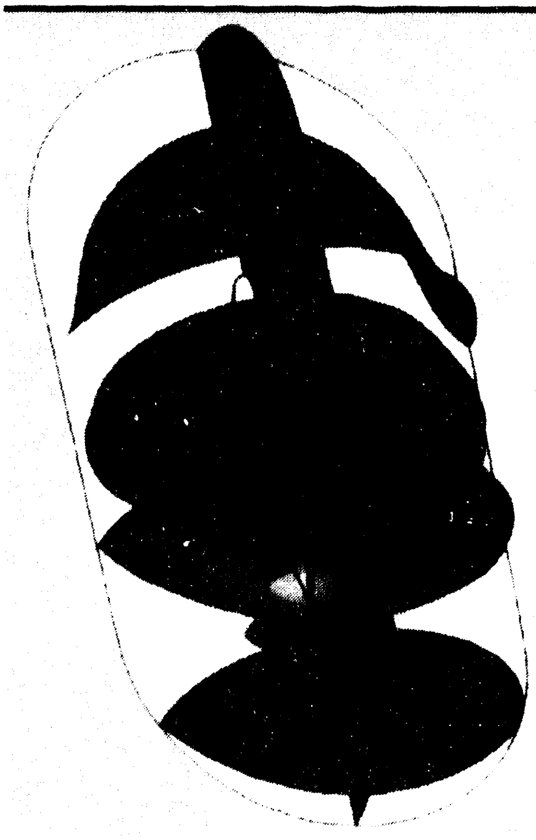

Figure 3. Elastic-scattered pressure field for streamline hull model.

Solution Strategy

The system of semi-discrete equations for the structural acoustics problem is non-symmetric and poses severe computational restrictions if a time-stepping scheme is applied directly. Because the goal for PING was to provide a high performance, robust code capable of treating large-scale structural acoustic problems, several trade-offs were required. To ameliorate the problems arising from the non-symmetric coupled system, the equations are treated in a segregated fashion, in which the structure is integrated implicitly in time and the fluid is integrated explicitly in time. Newmark integration (trapezoidal rule for $\beta = 1/4$, $\gamma = 1/2$) is used for the structure, and a second-order-accurate central difference scheme (with the radiating boundary terms centered) is used for the fluid. Because the time integration of the fluid is explicit, the overall stability of the coupled system is governed by the Courant limit in the fluid, while the time integration in the structure is unconditionally stable.[11-16] The integration scheme for the coupled system of equations is as follows:

1. Compute the fluid pressure right-hand-side using the current structural acceleration and data for the incident plane waves, then update the pressure.
2. Compute new forces on the structure using the current fluid pressure field, then update the structural response.
3. Update the velocity and the acceleration for the structure, and repeat steps 1–3.

Because the step-by-step algorithm is only conditionally stable, the Courant limit (based on the smallest element in the fluid mesh) is automatically computed, and the time-step size is adjusted to maintain stability. The explicit time integration in the fluid minimizes memory demands and is highly vectorized. The unconditionally stable integration in the structure allows the times step of the structure to be identical to that in the fluid, avoiding issues relating to subcycling.

Applications

To illustrate PING's capability to perform simulations of fluid-loaded elastic bodies, a Gaussian pressure pulse was used to excite a 1/24th-scale streamline hull model. **Figure 3** shows the resulting elastic-scattered pressure field at a point in the simulation where the incident wave is just traversing the stern of the vehicle. The elastic-scattered pressure field is represented in terms of isosurfaces. The acoustic radiation effects of a longitudinal traveling wave in the structure can be clearly identified in the fluid by the isosurfaces. This simulation required approximately 50,000 fluid elements and several thousand shell elements.

Future Work

Continuing research in linear structural acoustics will be directed to two primary areas: the first will be towards the development of higher-order radiation boundary conditions that are computationally efficient; the second will concentrate on the investigation of geometric scaling laws for acoustically excited structures.

---

1. C. Rosenkilde, *Studies on the Implementation of Normals and Curvatures I. The First of Mean Curvature*, Lawrence Livermore National Laboratory, Livermore, California, UCRL-ID-115168, Part I (1993).
2. J.B. Jones-Oliveira and L.P. Harten, *Transient Fluid-Solid Interaction of Submerged Spherical Shells Revisited with Consideration of Radiation Effects*, Lawrence Livermore National Laboratory, Livermore, California,UCRL-JC-114883 (1993).
3. M.A. Christon, *PING: An Explicit Finite Element Code for Linear Structural Acoustics—User Manual*, Lawrence Livermore National Laboratory, Livermore, California, UCRL-MA-114536, Draft (1993).
4. B.N. Maker, *NIKE3D: A Nonlinear, Implicit, Three-Dimensional Finite Element Code for Solid and Structural Mechanics—User Manual*, Lawrence Livermore National Laboratory, Livermore, California, UCRL-MA-105268 (1991).
5. G.L. Goudreau, *Evaluation of Discrete Methods for the Linear Dynamic Response of Elastic and Viscoelastic Solids*, Ph.D. Thesis, Dept. Civil Engineering, University of California, Berkeley, California, Report No. 69-15, 68 (1970).
6. P.R. Wallace, *Mathematical Analysis of Physical Problems*, Holt, Rinehart and Winston (New York, New York), 231 (1972).
7. H.C. Chen and R.L. Taylor, *Int. J. Numer. Methods Eng.* **29**, 683 (1990).
8. M. Johnsen and D.R. Lynch, "A Second-Order Radiation Boundary Condition for the Shallow Water Wave Equations on Two-Dimensional Unstructured Finite Element Grid," accepted for publication in *Int. J. Numer. Methods Fluids* (December 1993).

9. P.M Pinsky and L.L. Thompson, *J. Acoust. Soc. Am.* **91**, 1320 (1992).

10. O.C. Zienkiewicz and P. Bettes, *Int. J. Numer. Methods Eng.* **13**, 1 (1978).

11. T. Belytschko and R. Mullen, *Int. J. Numer. Methods Eng.* **12**, 1575 (1978)

12. T.J.R. Hughes and W.K. Liu, *J. Applied Mechanics* **45**, 375 (1978a).

13. T.J.R. Hughes and W.K. Liu, *J. Applied Mechanics* **45**, 371 (1978b).

14. K.C. Park, C.A. Felippa, and J.A. DeRuntz, "Stabilization of Staggered Solution Procedures for Fluid-Structure Interaction Analysis," *Proc. Computational Methods for Fluid-Structure Interaction Problems—ASME Winter Annual Meeting*, T. Belytschko (Ed.), AMD Vol. 26, 95 (1977).

15. K.C. Park, "Partitioned Transient Analysis Procedures for Coupled-Field Problems: Stability Analysis," *Transactions of the ASME*, Vol. 47, 370 (1980).

16. K.C. Park and C.A. Felippa, *J. Applied Mechanics* **47**, 919 (1980).

Diagnostics and Microelectronics

The Diagnostics and Microelectronics thrust area conducts activities in semiconductor devices and semiconductor fabrication technology for programs at Lawrence Livermore National Laboratory. Our multidisciplinary engineering and scientific staff use modern computational tools and semiconductor microfabrication equipment to develop high-performance devices. Our work concentrates on three broad technologies of semiconductor microdevices: (1) silicon and III-V semiconductor microelectronics, (2) lithium niobate-based and III-V semiconductor-based photonics, and (3) silicon-based micromachining for application to microstructures and microinstruments.

In FY-93, we worked on projects in three areas, described in the reports that follow: (1) photocathodes; (2) a fiber optic transceiver for interfacing digital superconducting electronics; and (3) a novel flow cytometer.

While our new Micro Technology Center, completed in FY-92, is primarily a solid-state device research facility, the emphasis of the thrust area is to solve problems relating to diagnostic and monitoring instrumentation in a variety of scientific investigations.

Joseph W. Balch
Thrust Area Leader

Section 3

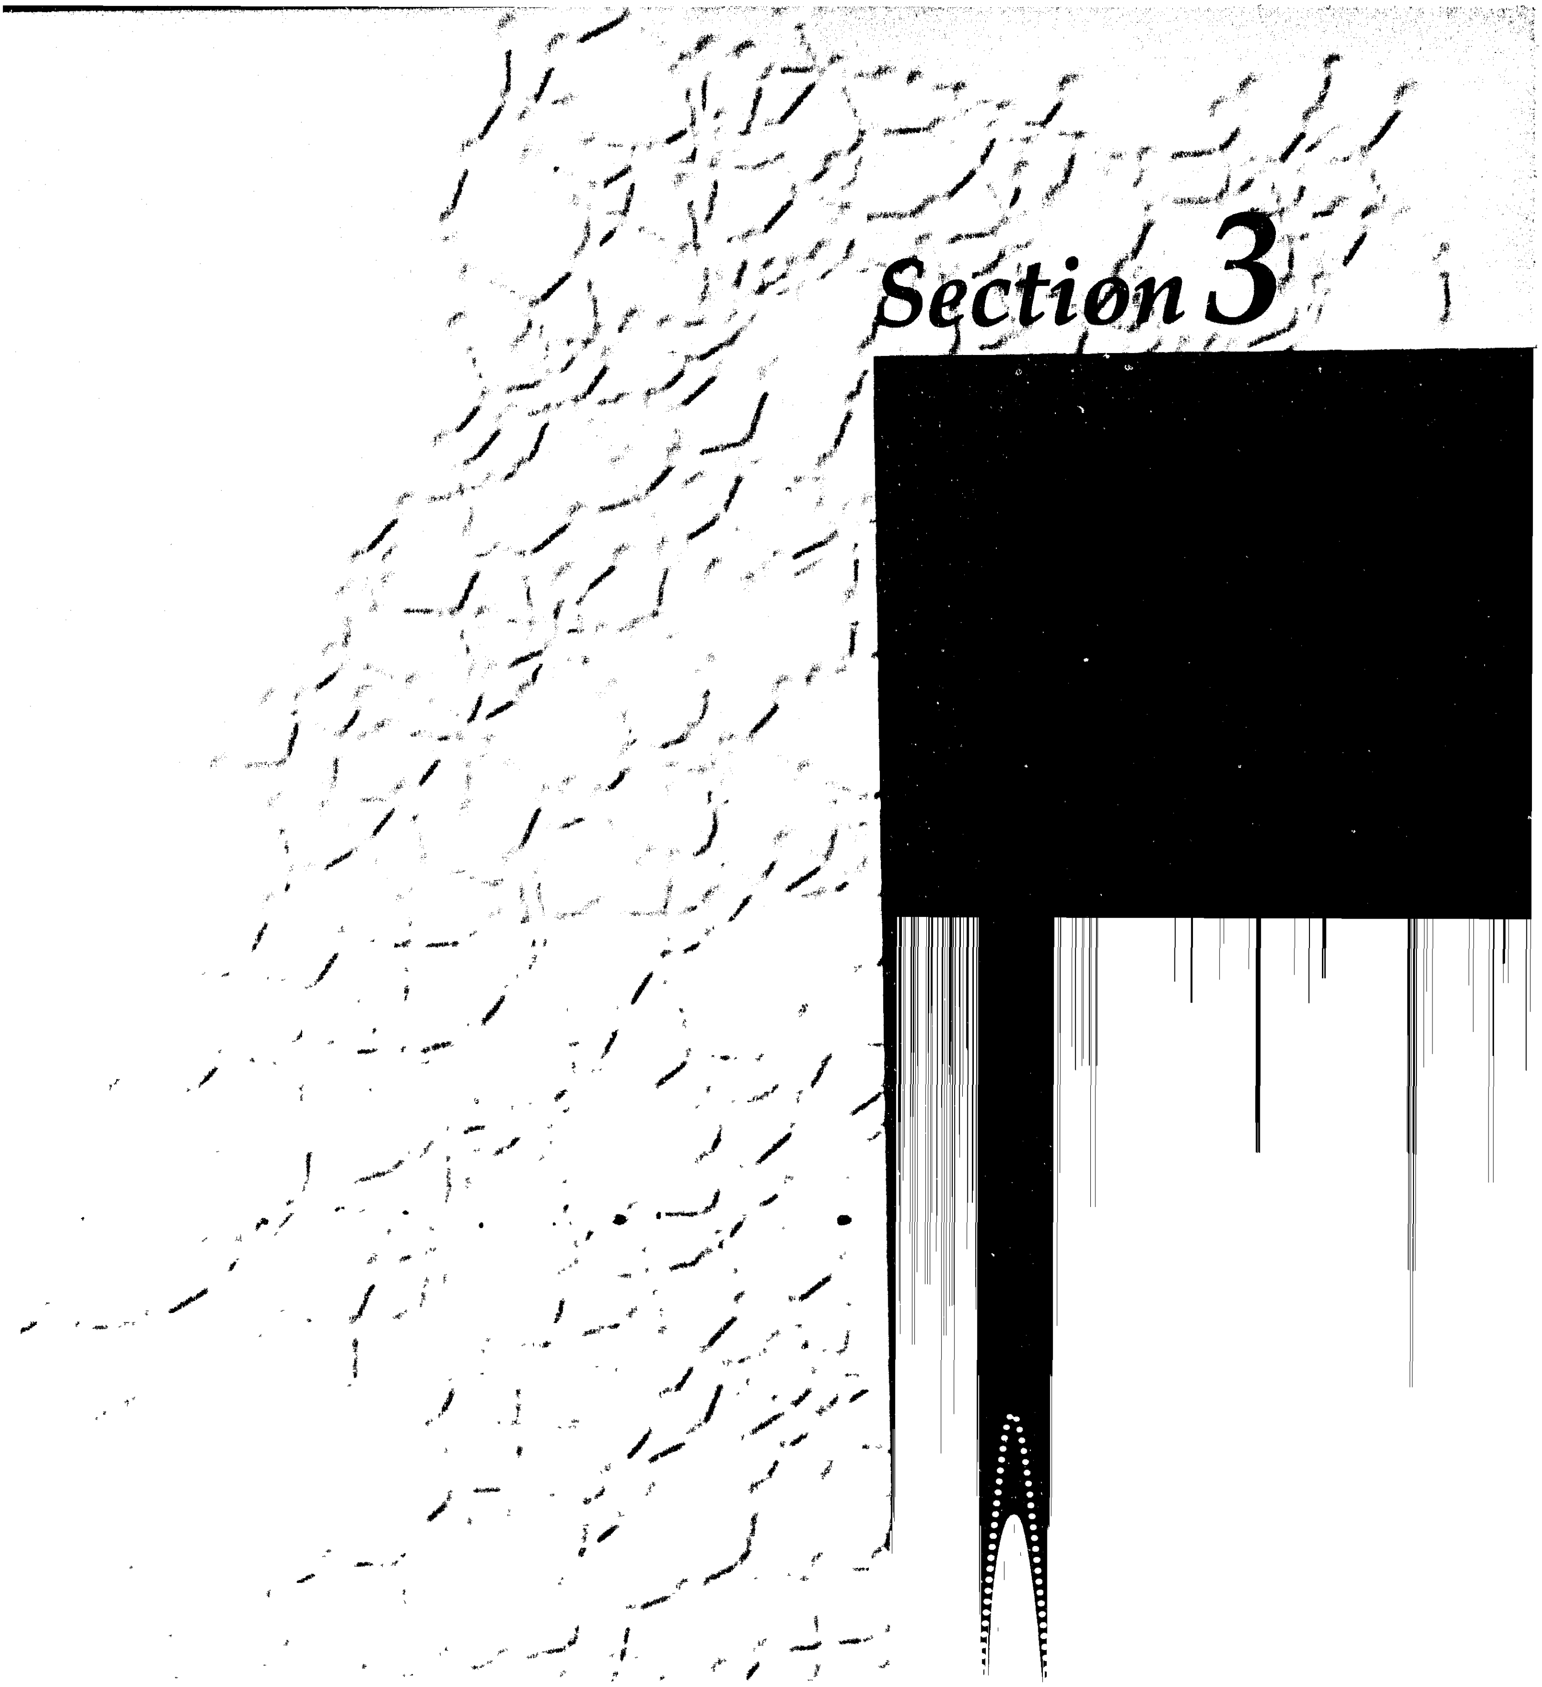

3. Diagnostics and Microelectronics

Overview

Joseph W. Balch, Thrust Area Leader

Infrared Photocathodes

Gregory A. Cooper and
Raymond P. Mariella, Jr.
Engineering Research Division
Electronics Engineering

Robert Chow
Materials Fabrication Division
Mechanical Engineering

Photocathodes with sensitivity to infrared (IR) wavelengths longer than 1 µm are not commercially available. There are several applications for such photocathodes at both Lawrence Livermore National Laboratory and the Department of Defense. We are using the precision and high quality of molecular beam epitaxy to grow structures that can absorb infrared photons, convert them into mobile electrons, and then transport those electrons to the surface for emission into the vacuum. Building on our previous work of growing IR absorbing materials and emitting surfaces, this year we have concentrated on the problem of electron transport.

Introduction

Of all materials tested, p-type GaAs has shown the highest quantum yield (η, defined as the number of emitted electrons per incident photon) for the detection of visible and near-infrared (IR) light. State-of-the-art commercial GaAs photocathodes can have quantum yields as high as 10% for wavelengths as long as 0.9 µm. There are other solid-state materials with higher efficiencies at converting photons into mobile electrons, but the ability to emit the electrons into a vacuum, as the photocathode does, offers greater radiation hardness. Photocathodes are therefore well suited to photon counting and to two-dimensional imaging when used in conjunction with electron multipliers, such as microchannel plates.

For longer wavelengths, GaAs is essentially a transparent material, and direct detection of IR light is not possible. Smaller bandgap semiconductors can absorb IR photons and generate mobile electrons, but such materials generally have a large energy barrier at the surface that prevents the efficient emission of the electrons into the vacuum. Other researchers, such as those at Varian/Intevac and Rockwell International, have proposed and demonstrated smaller bandgap photocathodes, where the surface barrier is overcome by transferring the electrons to a higher energy level, using a strong electric field during the transport to the surface. Although this concept works, it provides only a marginal increase in wavelength response (out to 1.4 µm) and doesn't appear to be extendible to longer wavelengths.

Our approach, molecular beam epitaxy (MBE), is to grow additional layers on a small-bandgap IR absorber. The energy bandgap of the additional layers is graded from the value of the absorber to the value of a good electron emitter, such as GaAs (**Fig. 1a**). MBE allows us to grade the energy gap without significantly changing the crystal structure, by substituting one elemental component with another. Maintaining the crystal structure is important to the overall efficiency of this concept. Crystal defects can interfere with the transport of electrons from the IR absorber to the emitter. In our design, we are taking advantage of our discovery (patent applied for) that p-type AlGaSb is as good an electron emitter as GaAs, and is lattice-matched to several IR absorbing materials that cover the wavelength range from 1 to 10 µm. An important difficulty with this approach, however, is that as the energy gap is increased, the energy of the electron conduction band also increases. This means that an electric field must be applied to push the electrons up the conduction band hill to the emitter (**Fig. 1b**). Supplying this electric field has been the goal of our work this year.

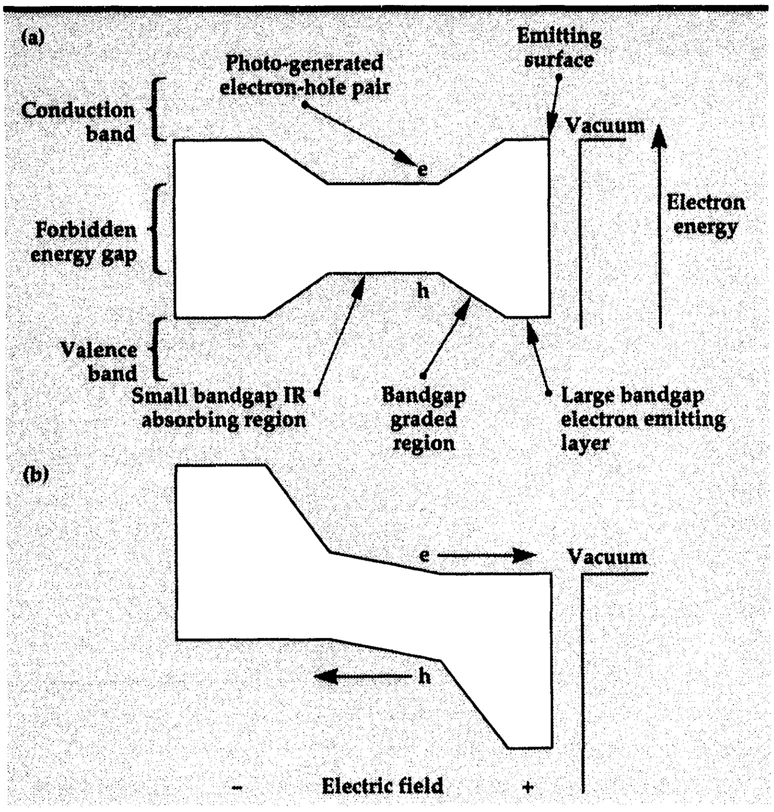

Figure 1. Energy band diagram of proposed IR photocathode. Photons with energies greater than the bandgap of the IR absorbing region are absorbed in this region, and the energy is used to promote an electron from the valence band, across the forbidden energy gap, into the conduction band, leaving a hole in the valence band. In the as-grown, unbiased structure (a), the photo-generated electron does not have enough energy to overcome the conduction band barrier to get into the emitting layer. With bias applied to the structure (b), the barrier is removed, and the electron may move freely to the surface to be emitted into the vacuum.

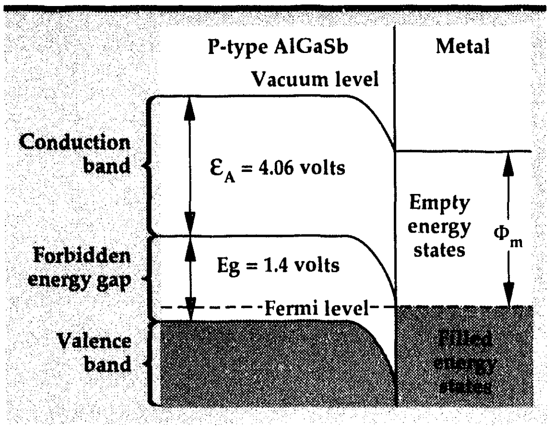

Figure 2. Energy band diagram of interface between p-type AlGaSb and a metal with a small work function, Φ_m. Mobile electron vacancies (holes) near the top of the valence band are unable to overcome the energy barrier to travel into the metal. Similarly, electrons near the top of the Fermi level in the metal are unable to move up the energy hill to get into the conduction band of the semiconductor. An ideal situation like this would limit electrical conduction across this interface to very small levels.

Progress

The electric field required to push the electrons up the conduction band hill to the emitter has to be produced by applying an electrical bias voltage across the device. The polarity of this bias is positive on the emitter side and negative on the IR absorber side. Our problem has been to find a way to apply the positive bias to the p-type emitter without injecting a large number of holes into the device. Based on published values of electron work functions, we have identified several metals that should form Schottky barriers to hole transport when applied to p-type AlGaSb (Fig. 2). We have seen no experimental evidence of significant barriers to current flow for any of the metals we have tried. We believe that this could be due to unknown effects on the AlGaSb surface or to unusually high hole concentrations in the AlGaSb. Other researchers have noted that unintentionally doped GaSb and AlGaSb tend to be strongly p-type. AlGaSb is a relatively new semiconductor material, and little published data exists regarding free surface characteristics and metal interfaces. One report, which we did not verify, suggests that GaSb grown on the <311> orientation should have lower intrinsic hole concentrations. We have been unable to determine the hole concentration using either Hall measurements or Polaron electrochemical techniques.

In an attempt to lower the hole concentration, we tried lowering the GaSb fraction in AlGaSb to zero, essentially using pure AlSb as the emitter material. Although AlSb was able to emit electrons, its performance was only about 1% that of the AlGaSb surface. We tried to make a Schottky barrier contact to AlSb using silver, but no barrier was observed in the electrical measurements.

We announced the availability of this technology and received several responses from U.S. companies. One company signed a non-disclosure agreement and visited Lawrence Livermore National Laboratory (LLNL) to discuss collaboration. No further technology transfer activity has occurred.

Future Work

There are no funds to continue this work at LLNL. We plan to publish what we have learned, with the hope that other researchers will continue this work.

Fiber Optic Transceiver for Interfacing Digital Superconducting Electronics

Jeffrey D. Morse
Engineering Research Division
Electronics Engineering

Kent G. McCammon
Defense Sciences Engineering Division
Electronics Engineering

We have conducted low temperature experiments to characterize the performance of high speed photodetectors and $LiNbO_3$ optical modulators in cryogenic environments down to 4.2K. Specific parameters quantified from this study include bandwidth, responsivity, V_π (half wave voltage), optical loss, and package stability. Metal-semiconductor-metal photodiodes fabricated from GaAs and InGaAs have been characterized. Results demonstrate that both the responsivity and bandwidth depend on temperature. Several Mach-Zehnder modulator designs were tested over the dc to > 20-GHz range. These results demonstrate the performance of $LiNbO_3$ modulators at cryogenic temperatures, and enable the design of high-sensitivity data links. The experimental results are coupled with optical data link analysis to estimate the performance of transmit and receive modules in cryogenic systems. This provides further insight into the critical issues associated with the integration of optoelectronic networks with superconducting signal-processing circuits.

Introduction

The continual need for increased performance in advanced signal-processing systems and applications has made the implementation of photonic data links within cryogenic environments very attractive. This arises, in part, from the reduction in system noise achievable, as well as from the possibility of integrating optoelectronic circuits with very-high-performance superconducting electronics. Advances in superconducting materials, devices, and circuits, along with improvements in cryogenic cooling systems, have made this circuit technology competitive in terms of power requirements and cost for high-performance, signal-processing systems. Superconducting circuits provide processing speeds that far exceed the present state of conventional electronics. Supercomputing, high-speed switching, and signal processing are some examples of applications that are driving the requirements for higher-speed electronics. Superconducting circuits can provide a solution to higher-speed processing in limited scale. A critical limitation for realizing large-scale systems using superconductors is the thermal interface between a liquid helium environment and room temperature. For example, the thermal interface limits the use of superconductors as a central processor in a computer. The fast processor in the dewar must interface with memory at room temperatures. The physical location of the superconducting processor and memory must be kept in close proximity for bit-rates > 5 Gbits/s due to transmission line dispersion effects. Also, the heat load due to the electrical cable interface limits the number of connections and reduces the lifetime of a remote system such as a satellite platform.[1] Replacing the electrical lines with optical fibers reduces the overall thermal load, eliminates electromagnetic interference, and allows increased, aggregate data capacity through the cryogenic interface.

Resolving the issues of packaging, operation, and integration of high-performance optoelectronic devices within cryogenic environments is the remaining step for realization of the systems described above. The fiber optic transmitter and receiver are the critical components of the overall fiber-optic system. Bi-directional communication requires a transceiver in the cryogenic environment interfacing with low-level signal levels.

Ideally, the transceiver components must operate at low temperature, and the fiber system must operate at high bit-rates with low noise.

Operation of LiNbO$_3$ modulators and semiconductor photodetectors was evaluated at cryogenic temperatures. For photodetectors, the responsivity and bandwidth were compared for room temperatures and cryogenic temperatures. LiNbO$_3$ modulators were tested for the effects of low temperatures on bandwidth, loss, and V_π, or half-wave voltage. This investigation has focused on the use of LiNbO$_3$ external modulators over laser diodes for integration with superconducting electronics. External modulators are passive optical modulators that dissipate heat only through optical scattering losses, and require only one bias voltage signal for operation. Also, the external modulator approach allows the laser source to be physically separated from the modulator. The laser is thus a CW source and can be chosen for wavelength, type, and power level required. The requirement for high sensitivity dictates the use of low-noise, solid-state laser sources. Examples of these lasers are diode-pumped Nd:YAG lasers in the 1.3-μm window and a diode-pumped erbium laser in the 1.5-μm window. The lasers not only give lower relative intensity noise over semiconductor lasers, but also can provide orders of magnitude increases in optical power available.

Figure 1. Cross-section of GaAs MSM photodiode structure.

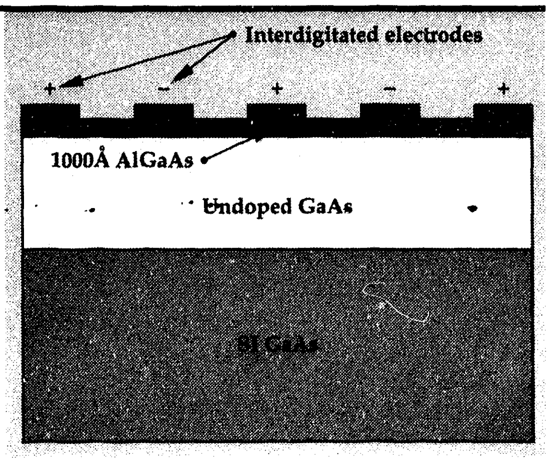

Figure 2. GaAs MSM photodetector response at various temperatures.

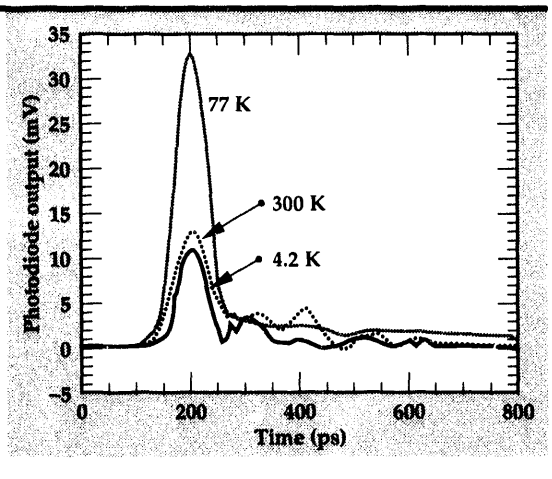

Photodetector Characterization

The performance of GaAs metal-semiconductor-metal (MSM) photodiodes has been characterized at cryogenic temperatures. MSM photodetectors were chosen for their simplicity and ease of fabrication, as well as for subsequent integration. Furthermore, MSM performance is comparable to that of PIN photodiodes, which are more complicated for monolithic integration. The cross-section of the MSM photodiode structure is illustrated in **Fig. 1**. The active region of the photodiode is a 1-μm, undoped GaAs layer grown by molecular beam epitaxy, and therefore of high quality. This layer is capped by a 1000-Å Al$_{0.5}$Ga$_{0.5}$As layer to improve the Schottky barrier, and to reduce the impact of surface defects on the GaAs active region. The fabrication sequence consists of patterning a Schottky barrier metal (Ti/Pt/Au for GaAs) in an interdigitated electrode configuration. The electrode spacing and width typically range from 0.1 to 5.0 μm. For this work, 1-μm electrode spacing designs were used. The length of the electrodes is 50 μm, and the overall active area of the photodetector is 50 μm × 50 μm, thus the design is well suited for coupling to an optical fiber having a core diameter of 50 μm or less, with minimal losses. The photodetectors were mounted on 50-Ω microstrip transmission lines coupled using SMA connectors. The photodetector is fiber pigtailed by aligning a 50-μm-core, multimode fiber, over the active region, then attached with UV-cured epoxy. The complete package design includes a stress relief fixture in order to minimize the impact of strain on the fiber pigtail, which may be catastrophic as the photodetector is brought down in temperature.

The GaAs MSM photodiodes were tested using a pulsed laser-diode source. The laser diode delivers 1.2-pJ, 20-ps full-width half-maximum, 810-nm optical pulses at a 10-KHz rate. For this experiment, the photodetector is cooled by direct immersion in liquid nitrogen or liquid helium. The resulting transient responses of the MSM photodiode are shown in **Fig. 2**. The bias voltage applied to the MSM is 1 V. The transient response tracks the optical signal to first order, and is limited by the mounting fixture to a rise time of ~ 25 ps, and a pulse width of ~ 35 ps. At room temperature, the falling edge of the transient response exhibits a tail, which is likely due to the associated time for complete sweepout of the charge by the electric field between the electrodes. At 77K, the photodetector response exhibits dramatic changes. The amplitude of the transient response increases by a factor of three. Furthermore, the tail on the falling edge is eliminated, with only a small reflection remaining, which is due to the package and bondwires. These effects are intuitive since the carrier mobility, hence drift velocity, increases as temperature decreases due to reduced phonon scattering. This is observed in the photodetector response, since the photocurrent and carrier transit time are directly related to the carrier velocity. At 4.2K, the amplitude of the photodiode response is roughly 10% less than at room temperature. The falling edge does not exhibit a tail, but the reflections remain. Interpretation of this result is difficult since competing processes exist that are not precisely quantified. The relevant physical processes include increased impurity scattering, which will dominate the phonon scattering at this temperature, thus some reduction in mobility is expected. The manner in which this effect impacts each type of free charge carrier will determine the effect on the photodetector transient response. Another physical effect that cannot be overlooked, is the increase of the semiconductor bandgap at low temperatures. This is particularly important for this case because the laser-diode optical wavelength is near the band edge. This results in an increase of the absorption depth, therefore charge carriers are photogenerated further from the surface, where the electric field strength is much less. The drift velocity of these carriers will be less than at high electric fields, and the contribution to photocurrent will be reduced.

Electro-Optic Modulator Testing

The electro-optic (EO) modulators tested at low temperatures were a combination of modulators designed and fabricated by Lawrence Livermore National Laboratory (LLNL), and commercially available devices from United Technology Photonics (UTP). All the modulators were made from x-cut LiNbO$_3$ and had fiber-pigtailed packaging. All the devices were Mach–Zehnders (MZ) with traveling wave electrodes except one lumped-electrode, balanced-bridge device design. In addition, two experimental LLNL devices were tested that were designed for 50-Ω impedance, with velocity matching for high bandwidth between the optical and microwave fields.

The modulator cooling process was implemented in two stages, using liquid nitrogen and liquid helium. The packaged modulators were placed on a cold finger inside a two-stage, insulated refrigeration system. The cold finger was gradually cooled by liquid nitrogen to slowly lower the temperature of the modulator. Temperature was monitored by a thermocouple gauge attached to the outside of the modulator package. Eventually, the modulator package was partially immersed in liquid nitrogen, where it was maintained to conduct measurements at 77K. To lower the temperature further, a jet stream of liquid helium was sprayed directly on the modulator package, using a helium transfer line. This approach allowed cooling to temperatures as low as 10K. This cooling technique gave minimal control of the temperature during cooling, but eliminated the need for a dewar with fiber-optic modifications. The next phase of testing will use a shielded helium dewar to contain the superconducting circuits and modulators.

The fiber pigtail-to-LiNbO$_3$ interface is the critical packaging concern. Thermal stresses at the interface fail if the rate of temperature change is excessive. We have encountered insertion loss increases for cooling rates that exceeded 20K/min. One device suffered a physical shift and permanent misalignment of the fiber-to-niobate interface. Another device shifted and caused a 3-dB loss while cooling at a rate of 26K/min. However, LLNL and commercial devices cooled at lower rates gave no additional insertion loss during cooling and cycling back to room temperature.

The experimentally measured temperature dependence of V_π for three different EO modulators is shown in **Fig. 3**. An 8% ± 1% increase in V_π, with temperature decreasing from 300K to 10K, was found for the three devices. The specific

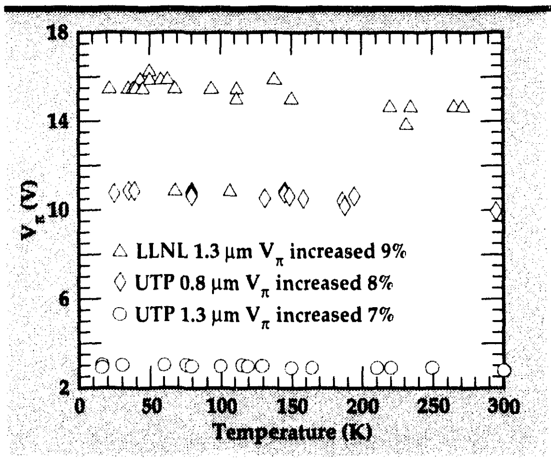

***Figure 3.* Temperature dependence of V_π for various LiNbO$_3$ EO modulators.**

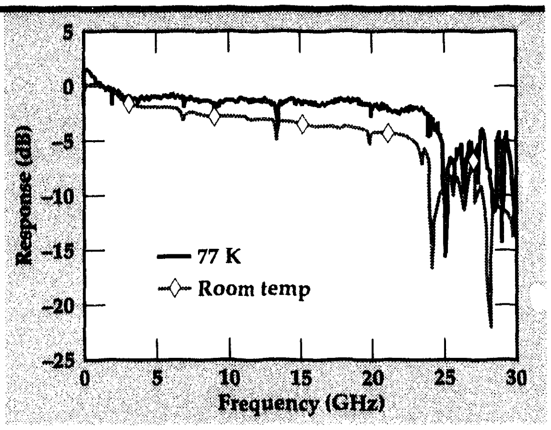

Figure 4. Temperature dependence of LLNL EO modulator frequency response.

devices tested consisted of an LLNL 1.3-µm, traveling-wave MZ; a UTP 0.8-µm, traveling-wave MZ; and a UTP 1.3-µm, lumped balanced-bridge modulator (MZ with a directional coupler for complementary outputs). All three devices exhibited the same linear rate of increase in V_π, within experimental error. All the devices tested were made with x-cut $LiNbO_3$.

Modulator frequency response was measured at cryogenic temperatures for a lumped-element device, a traveling-wave device, and a velocity-matched device. The swept frequency technique was used to determine the frequency response.[2] All three types of modulators tested showed no degradation in frequency response when cooled. A representative swept-frequency response is shown in **Fig. 4** for a LLNL fabricated modulator. The device is a traveling-wave MZ operating at 0.8 µm. This device has been designed to be velocity matched with 50-Ω characteristic impedance. The frequency response at liquid nitrogen temperature is improved over the room temperature response. Testing a larger sampling of velocity matched devices is underway to determine the temperature dependence of the frequency response. Frequency response improvements and V_π degradations at lower temperatures would allow tradeoffs in the modulator electrode design for cryogenic applications.

Conclusion

In summary, the performance of high-speed photodetectors and EO modulators has been characterized at cryogenic temperatures to determine the viability of interfacing superconducting electronics with fiber-optic transmitters and receivers. Experimental results have demonstrated that the performance of GaAs photodetectors and $LiNbO_3$ EO modulators is marginally degraded at most, and can even be improved. Further understanding will enable the design of EO components specifically for cryogenic environment applications.

Furthermore, the packaging of optoelectronic devices that include optical fiber pigtails has proven to be robust enough for cryogenic temperature transitions, given that appropriate care is taken in introducing the packaged components to the cryogenic environment.

---

1. A.R. Johnston, *SPIE Proc.* **1946**(38), 19.
2. O. Uehara, *Appl. Opt.* **17**(1), 68 (January 1978).

Novel Flow Cytometry

Raymond P. Mariella, Jr.
Engineering Research Division
Electronics Engineering

We report a new physical configuration for flow cytometry that greatly increases the signal-to-noise ratio for right-angle-scattered light and also greatly eases the alignment requirements. The new technique views the scattered light that is trapped within the optical waveguide of the flow stream in air.

Introduction

In previous flow cytometers (FCM's), the right-angle-scattered (RAS) light has been viewed perpendicular to the liquid flow, typically using a high numerical aperture (NA) microscope objective lens or fiber optics. Some of the difficulties associated with this approach are the very limited depth of field of high NA lenses, typically only a few µm, and the necessity to align precisely the exact focal point of the lens with that of the excitation light source. As is shown in **Figs. 1** and **2**, a NA 0.6 lens captures less than 14% of all solid angles. Our invention is the use of the unconfined aqueous flow stream itself as an optical waveguide, which captures approximately 17% of the solid angles. This very modest increase is augmented by the fact that all of the scattered light is trapped: there is no 'focal point' for this configuration. Alignment simply requires aligning the light source onto the flow stream; the liquid optical waveguide is then automatically 'aligned.' This approach provides robust, stable light collection. For the collection of elastically scattered light, another immense advantage occurs: the background of scattered light is extremely low when using the flow-stream waveguide (FSW), because the same physical properties that confine the desired light within the stream also keep random scattered light out.

Progress

In our studies, we have found that this new configuration gives the elastically scattered RAS signal a much higher signal-to-noise ratio than that of forward scattered light.

In building a system that uses the flow stream as an optical waveguide for collecting RAS, one must take into consideration that the effective NA of a water/air waveguide is $1/n_{\text{water}} = 1/1.33 = 0.75$. One could place the light detector directly into the flow, possibly with a conical lens to keep the flow from becoming too unstable. In this case, the area of the detector must be large enough to trap all of the rapidly diverging light from the FSW. For our research purposes, we found it

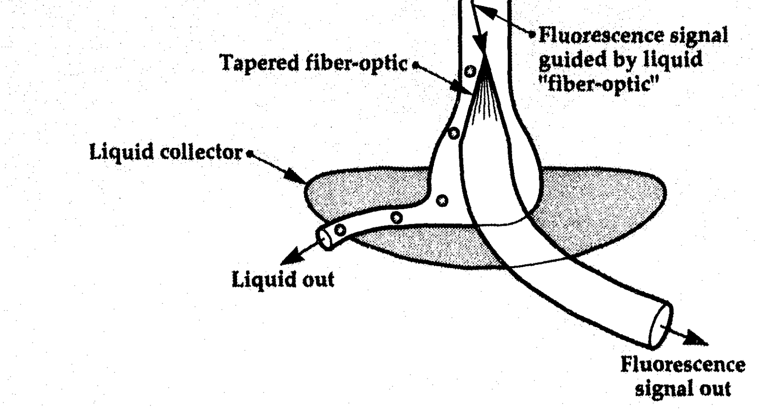

Figure 1. Schematic drawing of diode-laser-based flow cytometer. Cells in an aqueous carrier flow down through an orifice and form a free stream. When any cell passes through the laser beam, it will scatter (reflect) the laser light into the 'forward-light-scatter detector,' thus revealing the presence of a cell. This event is added to the total count. If there is a simultaneous RAS signal, then the cell can be specifically identified.

Figure 2. Plot of light collected vs subtended angle for s and p polarizations.

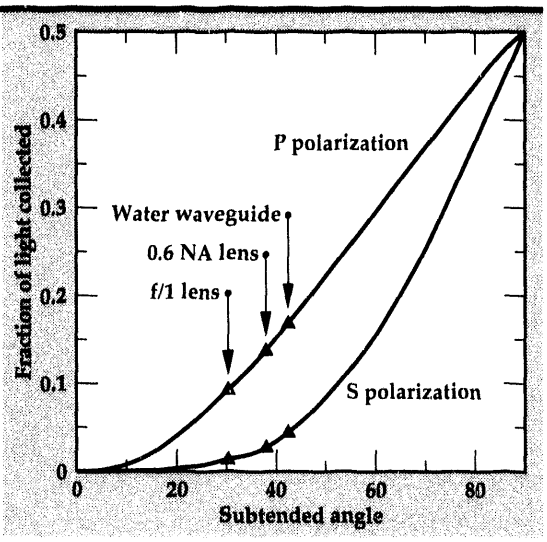

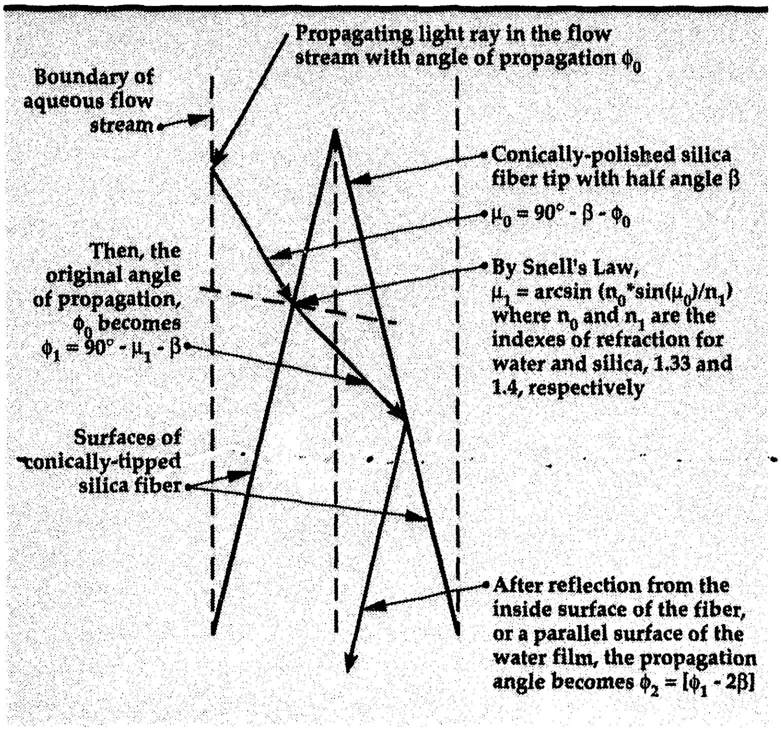

Figure 3. Schematic of the progress of a light ray as it passes from the aqueous flow stream, into the conical fiber tip, and reflects off sloping side of the conical shape. For example, for a 0.48-NA core fiber, any final angle of propagation, f_2, that is less than 20° will be successfully trapped in the FO.

convenient to place a conically polished fiber-optic waveguide (FO) in the flow stream to serve as an intermediate optical device to conduct the light from the FSW to the light detector. In this embodiment, it is easier to keep the flow stream from wetting the electrical contacts of the light detector. Because no commonly available FO has 0.75 NA, some of the large-angle-of-propagation light in the FSW would not be confined in the FO. Therefore, it is necessary to expand the beam and reduce the NA of the propagating light as it enters the FO; thus, one must use an FO whose core is larger than the diameter of the flow stream. This is illustrated in **Fig. 3**. In terms of ray tracing, by providing a conical taper at the transition from the FSW to the FO, the rays with ϕ larger than θ, the maximum angle of propagation for the NA of the FO, reflect off the angled surfaces of the water film over the tip and are redirected down the FO with reduced propagation angles, ϕ_2. There is a refraction of any propagating light as it passes from the water, with index of refraction n = 1.33 to the silica with n = 1.4. So, using Snell's Law, if half the tip angle is β, then as shown in **Fig. 3**,

$$n_1 \sin(\mu_0) = n_2 \sin(\mu_1)$$

$$\text{where} \qquad \mu_0 = 90 - \phi - \beta$$

$$\text{and} \qquad \phi_1 = 90 - \mu_1 - \beta .$$

The minimum length of the tapered transition must be sufficiently large so that all propagating rays with $\phi > \theta$ strike this slanted surface. The minimum length and width of the FO core can be estimated from:

$$\tan\phi = w_2 / x = w_1 / y$$

$$\text{and} \qquad \tan\beta = w_1 / l_1 = w_2 / l_2 ,$$

$$\text{with} \qquad l_2 = x + y + l_1 .$$

Thus,

$$l_2 = w_1 \left[(\cot\beta + \cot\phi) / (1 - \tan\beta * \cot\phi) \right], \qquad (1)$$

and, from above,

$$w_2 = l_2 * \tan\beta . \qquad (2)$$

For example, for 0.48-NA fiber with a 25° included tip angle (2β), l_2 and w_2 are 18.6*w_1 and 4.1*w_1, respectively, using the simple model shown in **Fig. 3**. Thus, for a 0.05-mm-diameter flow stream, the 0.48-NA fiber core needs to be approximately 0.2 mm or larger. The propagating rays, which cause **Eqs. 1** and **2** to attain their maximum value for l_2, and w_2 as well, are those with ϕ_1 only slightly larger than θ. The value of θ is normally quoted for an infinitely long fiber; for a short length of fiber, practically speaking, the effective NA for a few-cm-long fiber is larger, because the optical loss for ϕ slightly larger than θ is very small per cm of fiber. Therefore, one can use a short fiber with a core slightly smaller than that given by **Eq. 2** with no measurable decrease in performance, compared to a larger-core fiber.

The other important consideration in designing a tapered fiber to conduct RAS light from the FSW to a detector is the included conical angle of the tip. As is shown in **Fig. 3**, the final angle of propagation, φ_2, assuming a thin film of water covering the silica core's tip, is given by $\varphi_2 = |\varphi_1 - 2\beta|$. In the preceding considerations, we have modeled φ_1 of the light rays as either being less than θ, and thus, already being confined in the FO, or reflecting off the conical taper once and being redirected down the FO with $\varphi_2 < \theta$. Because of the small difference in the indices of refraction between water and silica, there is very little reflection from the tip's surface. Although it may be counterintuitive, making the included tip angle, 2β, smaller makes the required FO core diameter smaller, until $|\varphi_1 - 2\beta|$ becomes less than θ. For the case of the 0.48-NA FO, the smallest 2β that still provides $\varphi_2 < \theta$ is 11.5°. We have successfully polished several 0.48-NA fiber tips to $\beta = 12.5°$; that is, close to the 11.5° limit. Using continuous-wave diode lasers with power less than 10 mW, we have easily seen 1-V RAS pulses from 2.9-μm-diameter latex spheres, using inexpensive photodiode/pre-amp hybrids with 1-MΩ feedback resistor.

Emerging Technologies

The mission of the Emerging Technologies thrust area at Lawrence Livermore National Laboratory is to help individuals establish technology areas that have national and commercial impact, and are outside the scope of the existing thrust areas.

We continue to encourage innovative ideas that bring quality results to existing programs. We also take as our mission the encouragement of investment in new technology areas that are important to the economic competitiveness of this nation.

In FY-93, we have focused on four projects, summarized in this report: (1) Tire, Accident, Handling, and Roadway Safety; (2) Load Calculation and System Evaluation for Electric Vehicle Climate Control; (3) Paste Extrudable Explosive Aft Charge for Multi-Stage Munitions; and (4) Flexible Manufacturing for Photonics Device Assembly.

Shin-yee Lu
Thrust Area Leader

Section 4

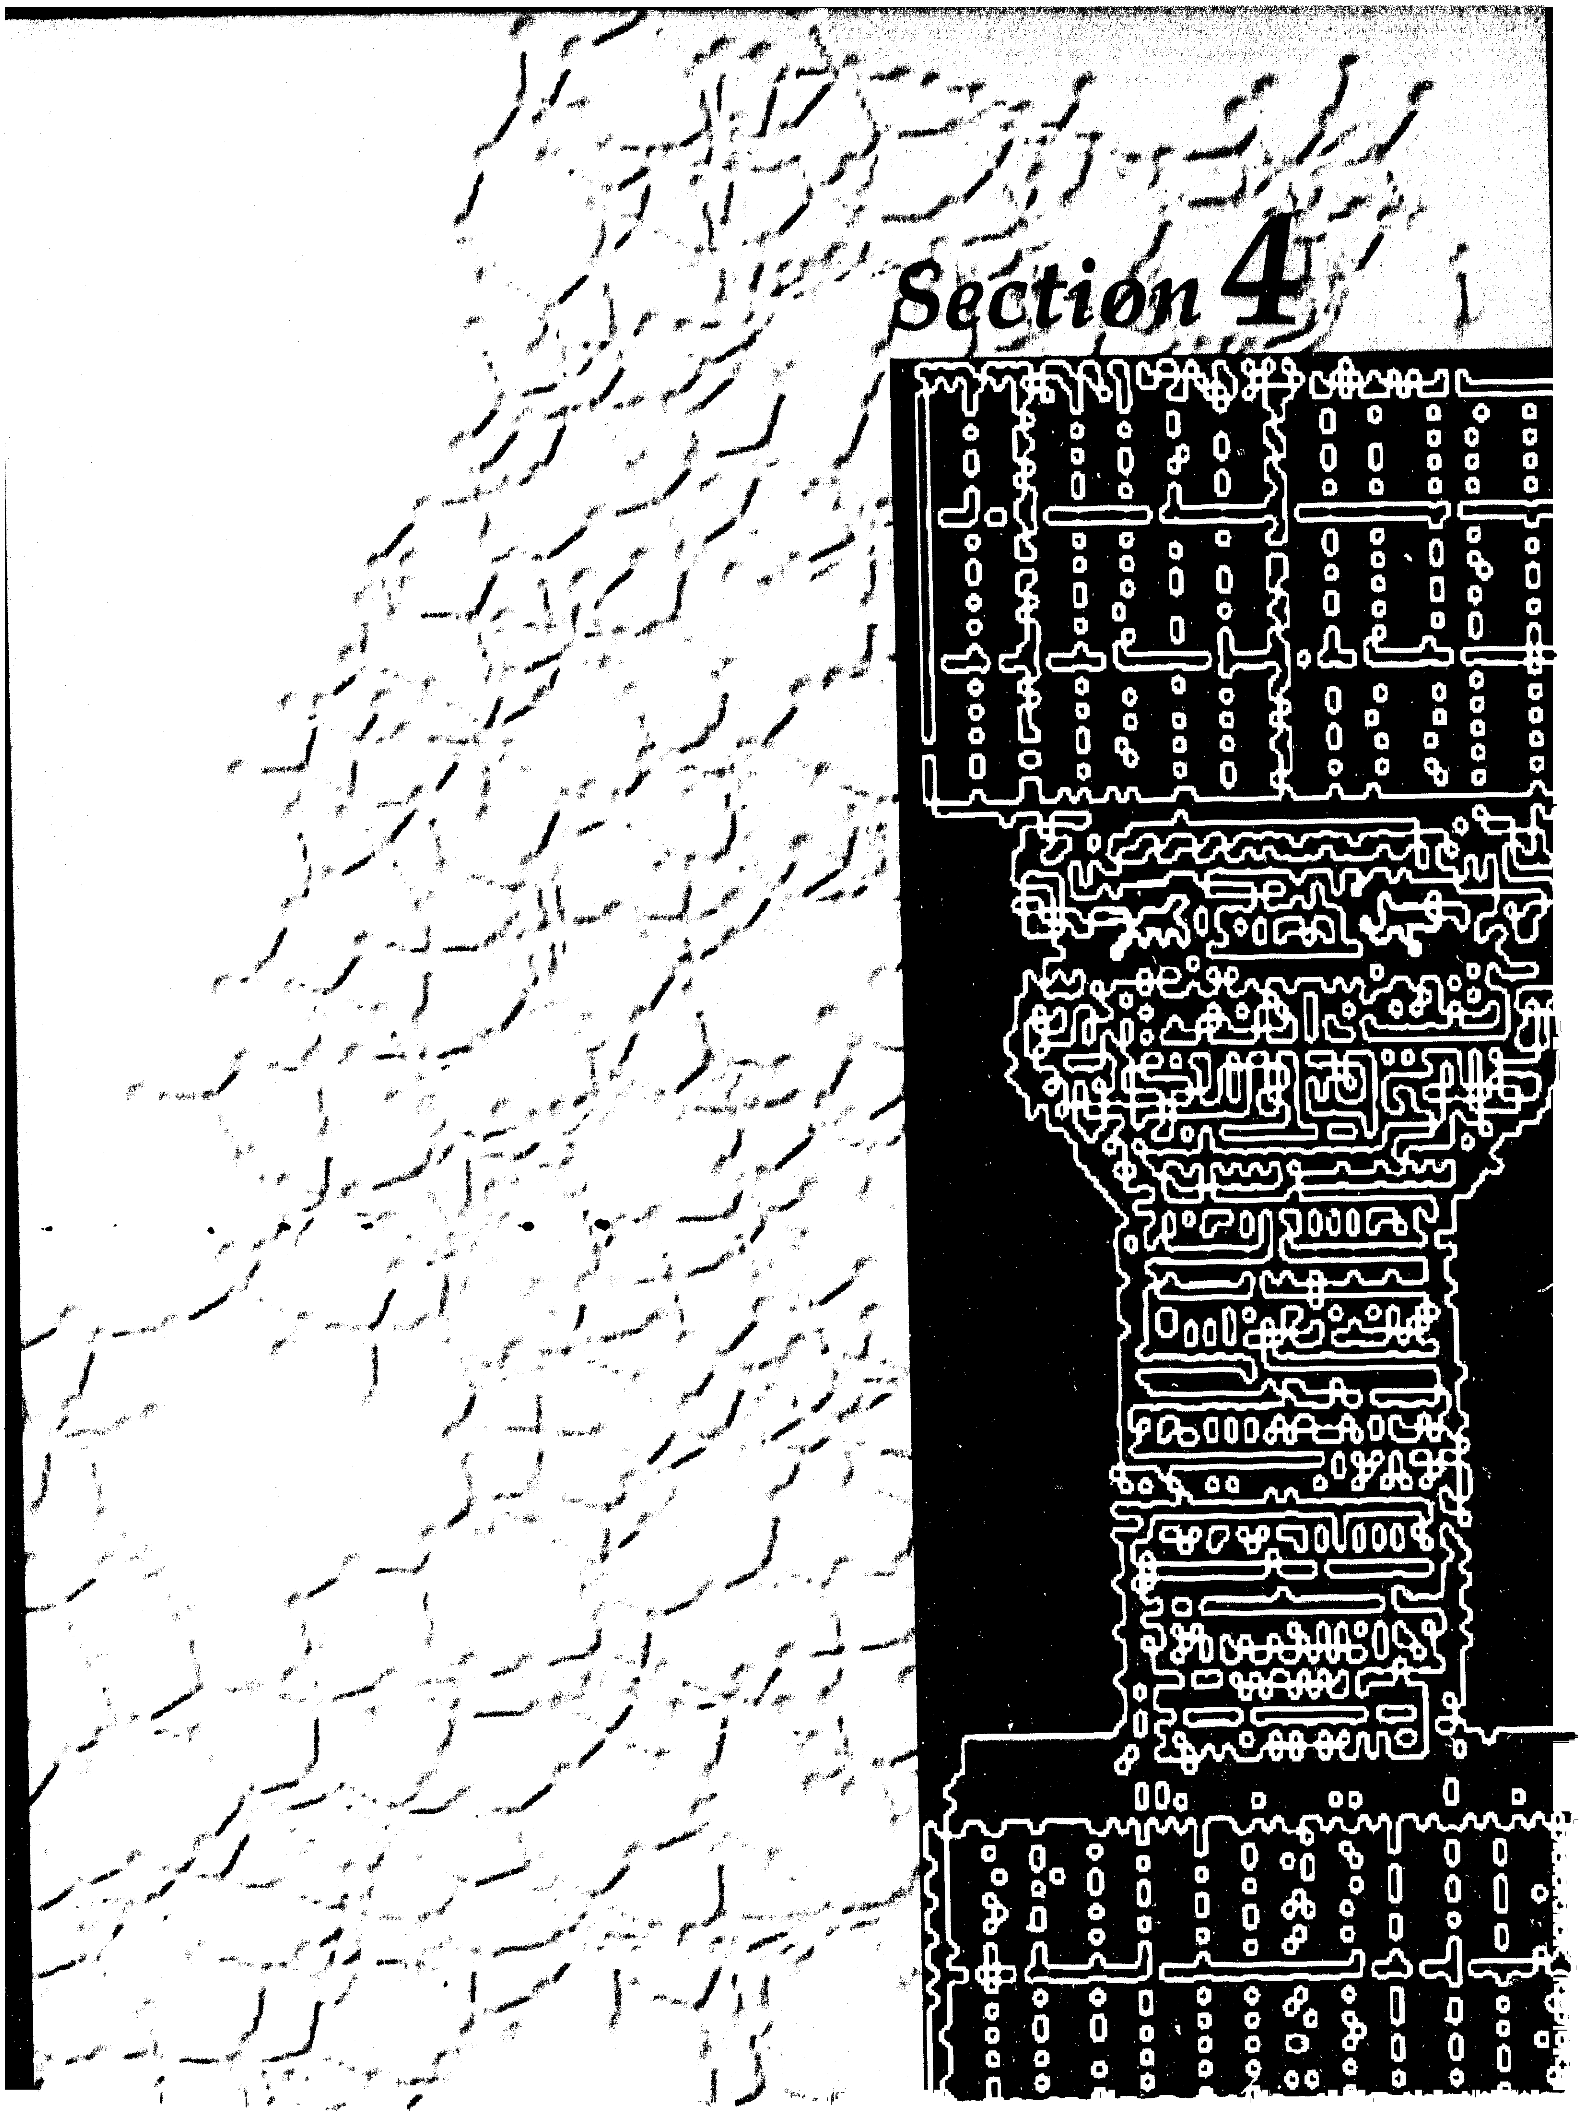

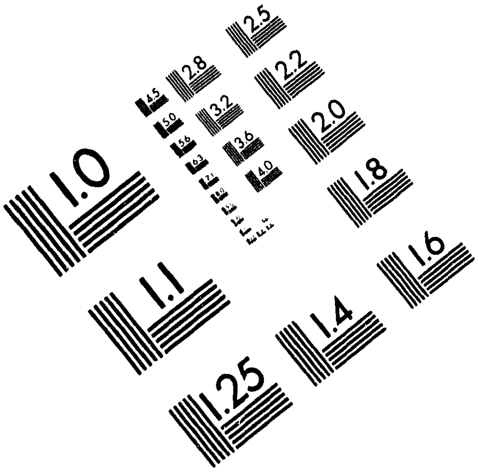

AIIM

Association for Information and Image Management

1100 Wayne Avenue, Suite 1100
Silver Spring, Maryland 20910

301/587-8202

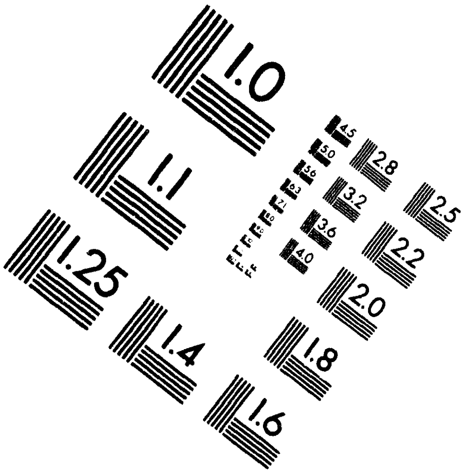

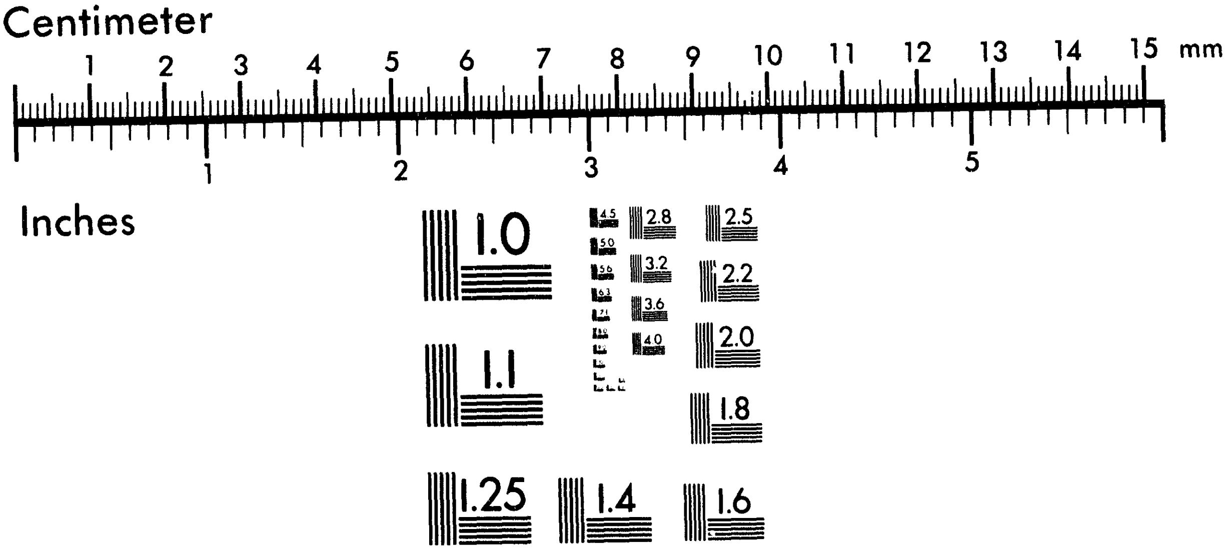

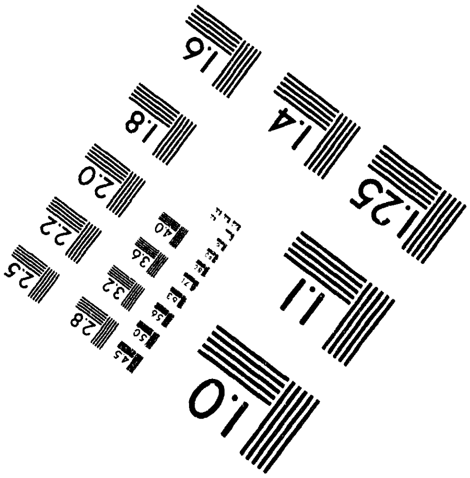

MANUFACTURED TO AIIM STANDARDS
BY APPLIED IMAGE, INC.

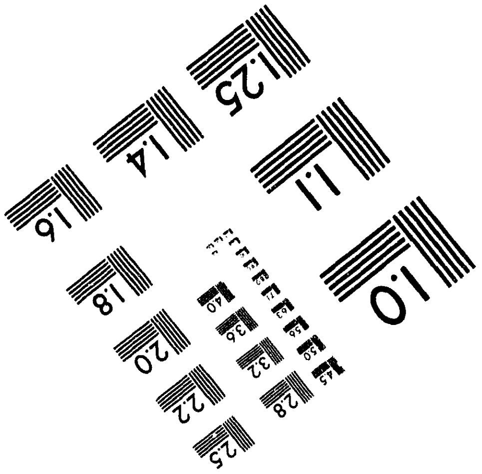

2 of 3

4. Emerging Technologies

Tire, Accident, Handling, and Roadway Safety

Roger W. Logan
Defense Technologies Engineering
Mechanical Engineering

We are developing technology for an integrated package for the analysis of vehicle handling and impact into roadside features and into other vehicles. The program involves the development and use of rigid-body algorithms and Lawrence Livermore National Laboratory's DYNA and NIKE finite-element codes. The goal is a tool for use by engineers in industry and at federal and state Departments of Transportation, allowing good quantitative results at the workstation level. At the same time, the work enhances our competency in furthering the development of DYNA and NIKE, and our expertise in spaceframe design for impact and crashworthiness.

Introduction

During its first year (FY-92), the Tire, Accident, Handling, and Roadway Safety project (TAHRS) made progress in the areas of vehicle handling and interfacing, roadside feature modeling, vehicle models, and plans for test and validation.[1] In FY-93, major areas of activity involved test and analysis comparison of aluminum tube crush, remote-controlled impact test of a vehicle, and design analysis for an aluminum-chassis concept vehicle. These activities made use of and enhanced our competencies in DYNA[2] analysis and testing, and in the ALFAMAX code developed in FY-92.

Progress

Design of Spaceframe Components

As part of an electric vehicle consortium, Lawrence Livermore National Laboratory (LLNL) and Kaiser Aluminum are conducting experimental and numerical studies on crashworthy aluminum spaceframe designs. We have jointly explored the effect of heat treat on crush behavior and duplicated the experimental behavior with finite-element simulations.[3] In the first stage of preparation for the design of a spaceframe aluminum chassis for a lightweight electric vehicle, we conducted a series of crush tests on extrusion sections, complemented by finite-element calculations of the crush behavior. The goal of this experimental/numerical study is to improve both the material behavior (strength with energy absorption) and the numerical modeling technology to allow scaling up to a full-vehicle finite-element model. Since one goal of finite-element technology advancements is to understand the limits of coarse meshing for workstation applications, the vehicle models might well be of the 2000–3000 node level. Thus, ability to capture the experimental crush is a critical issue.

The quasi-static axial crush tests were carried out on a Baldwin Universal Testing Machine with a 2300-kN capacity. 51-mm-square tubing was extruded in a 6XXX series Kaiser aluminum alloy.

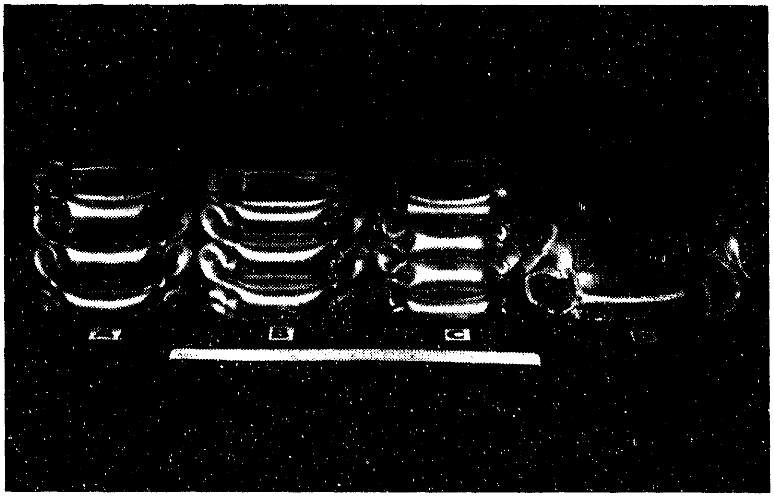

Figure 1. Experimental crush behavior of 6XXX aluminum extrusions with different heat treats. Note increased number of buckles in (c), and splitting in (d).

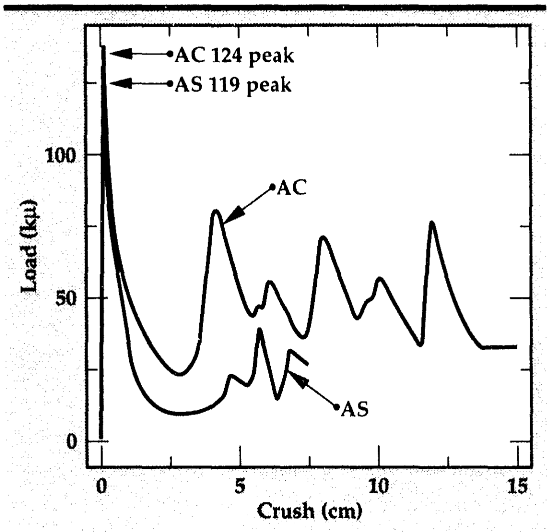

Figure 2. Load vs crush for aluminum extrusions, showing splitting in AS and high-energy absorption in AC.

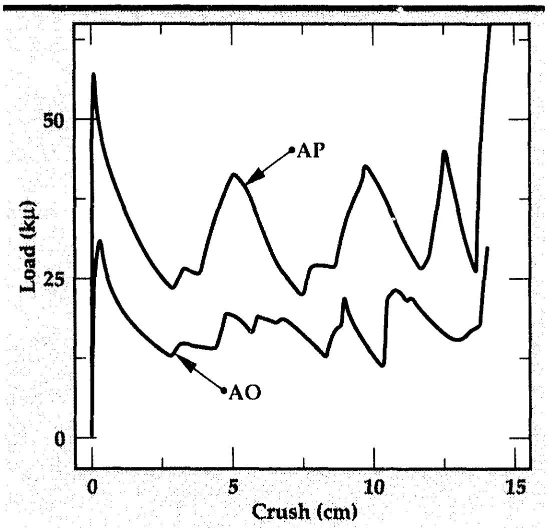

Figure 3. Load vs crush for axial compression of 6XXXA0 and 6XXXAP heat-treat aluminum extrusions.

By varying process parameters and heat treatment, two distinct failure modes can be obtained. In each case, the samples were 200 mm in length prior to the test. Four different processing and heat treatment variations were performed on the extruded material, and the results of the axial crush tests are shown in **Fig. 1** and discussed below in the order a, b, d, and c.

Figure 1a is the crushed extrusion that has been completely annealed, designated 6XXXA0. The material in this case demonstrates a buckle failure mode. The crush load (Pmax) for the sample was 30 kN and would clearly not be suitable for use in automotive structural applications. **Figure 1b** shows the extrusion in a partially heat-treated condition, designated 6XXXAP, where the buckling failure mode is retained, yet the crush load (Max) of 58 kN is still not suitable for structural application.

Figure 1d shows the deformation of the tube, designated 6XXXAS, that has a reasonable crush load (Max) of about 120 kN, but the crush failure mode is to bulge and eventually split. **Figure 2** shows the load vs crush curve (AS) for this material, where the load reaches the peak very quickly and drops off rapidly. The energy absorbed in this case is 15.6 kJ/kg. In the application of automotive spaceframes, the structure is required to withstand operational loads and have high impact-strength and energy-absorbing characteristics. High

Figure 4. DYNA3D calculation of crush of the 6XXXAP extrusion.

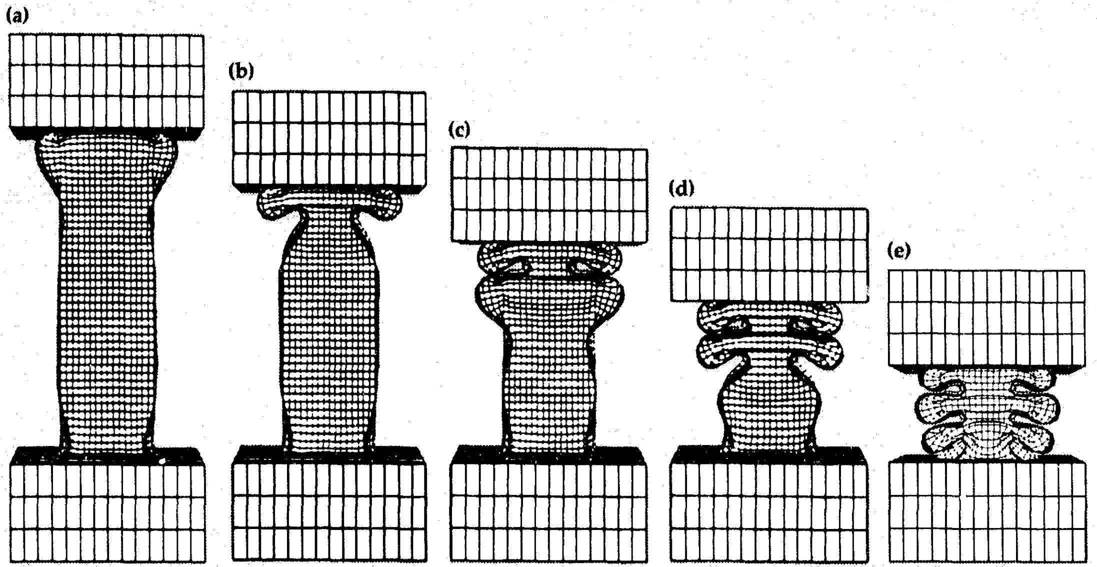

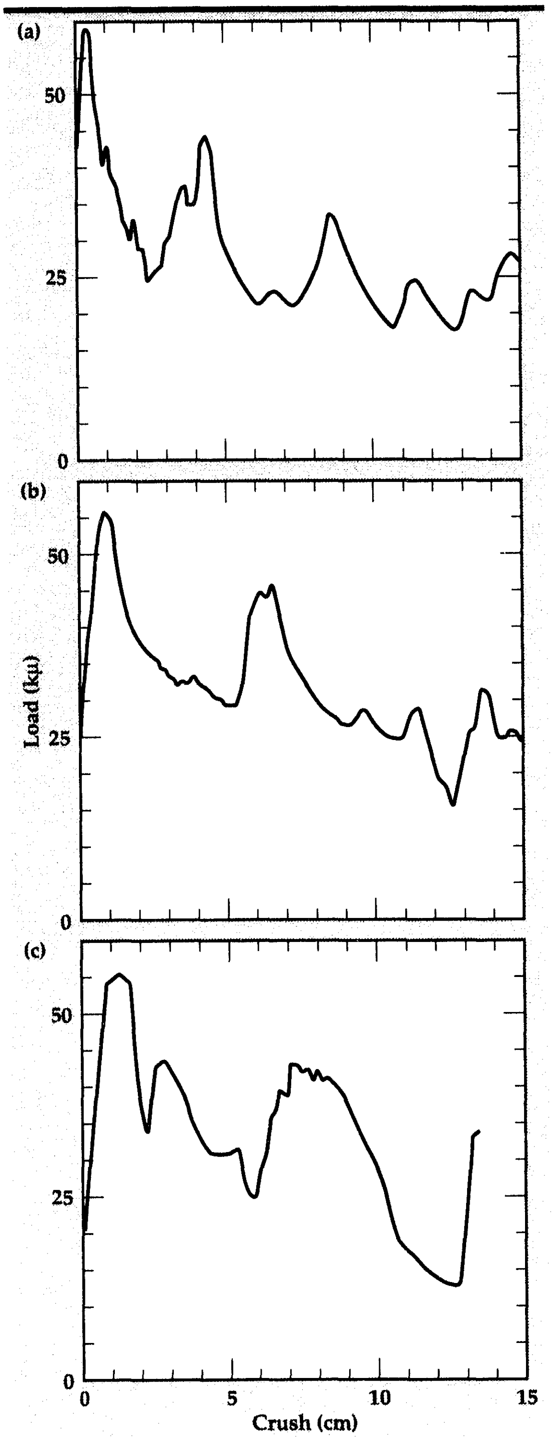

Figure 5. DYNA3D calculation of load vs crush of the 6XXXAP extrusion. Mesh size is (a) 3-mm, (b) 8-mm, and (c) 12-mm squares.

strength combined with high energy absorption is essential for reducing the acceleration of the passenger in a forward impact accident. The material in this case, although of high strength, should be further processed to be more suitable for impact-absorbing parts of the structure.

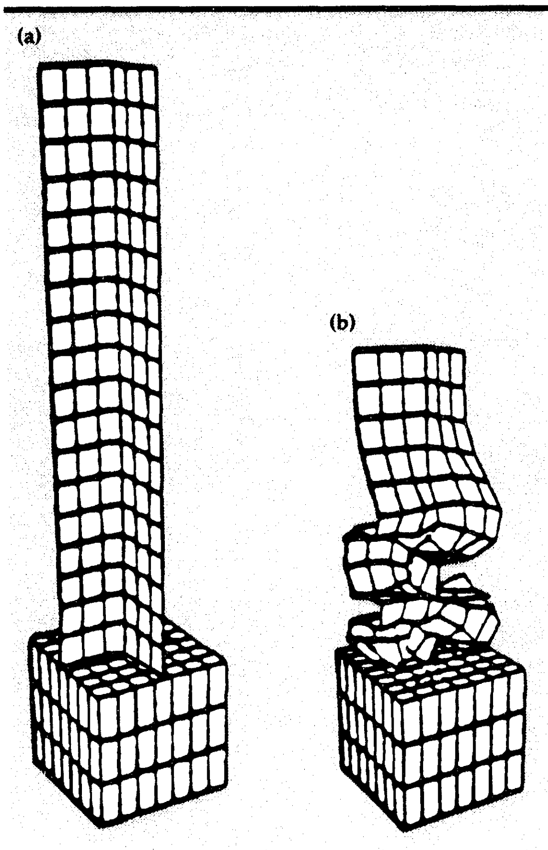

Figure 6. Calculated crush of 6XXXA0, tube using mesh size of 8 mm, near upper limit.

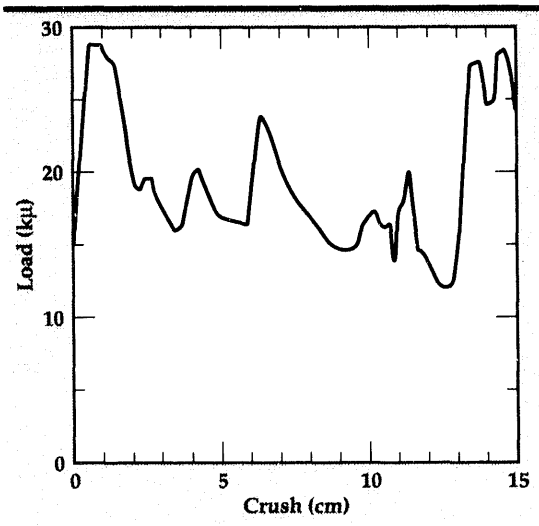

Figure 7. DYNA3D calculation of load vs crush of the 6XXXA0 extrusion. Mesh size is 8 mm.

Figure 1c shows the performance of the tube material, designated 6XXXAC, that is more suitable for the energy-absorbing sections of an automotive spaceframe. In this case, the material processing parameters and heat treatment were changed. The crush load (Max) is similar to the first sample (120 kN), but the failure mode is changed, and a succession of buckles is seen. **Figure 2** shows the load vs crush curve (AC) in this case, where the total energy absorption is

Figure 8. Remote steering system developed for LLNL test.

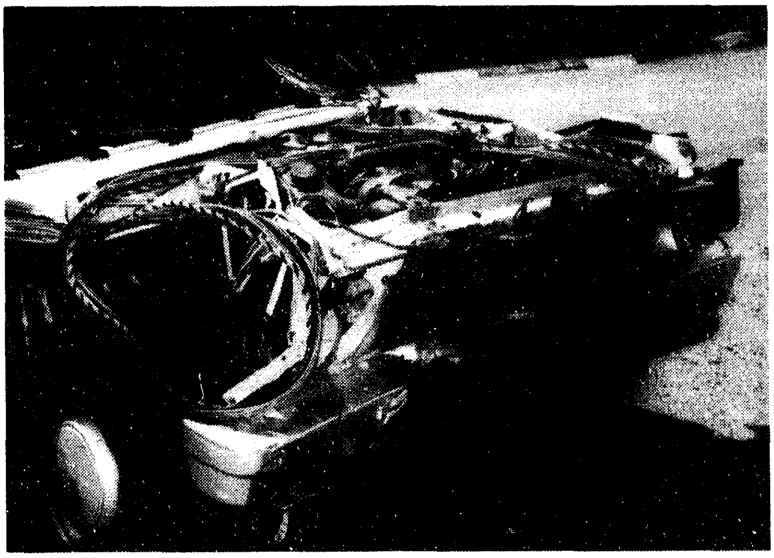

Figure 9. Results of frontal impact test at LLNL.

approximately 35 kJ/kg, more than twice the value seen in the first sample. After the initial peak, the curve tends to flatten except for the series of smaller peaks that correspond to each individual buckle. This material is more suited to the crush areas of the structure.

It is preferred that one optimum alloy be used in the structure, rather than several for different parts of the vehicle. In this way, manufacturing costs can be kept to a minimum. One single failure mode will also improve the accuracy of the modeling. The failure mode of buckling is desired, to improve energy absorption, and the geometric perturbations can be designed in, using the DYNA analysis, seen below. In this way, the average crush force can be retained, but the curve will be even flatter than shown in **Fig. 2**.

Axial Crush Modeling

The aluminum extrusions were modeled in axial crush with DYNA3D. The goals were to determine the finite-element mesh size necessary to match the experimentally observed behavior in the quasi-static tests described above, and to match the load vs crush behavior with stress-strain behavior as in static tension tests conducted as part of the study.

Figure 3 shows the crush force vs distance curves for the 6XXXA0 and 6XXXAP samples. Except for the high ratio of early peak to average crush level (which may be designed out if desired), the A0 curve is smooth and flat except for the perturbations where each of three successive buckles forms during crush. However, the average crush force of about 16 kN is too low for an optimized design. There is an improved situation in the force vs distance for the 6XXXAP heat treat, with average crush force now up to about 31 kN, but again with a good, flat curve except for buckle formation.

Figure 4 shows a sequence of the DYNA3D analysis of the crush test of the AP material. The sequence of formation of the buckles (location is random unless seeded) and the number and nature of the buckles agree well with the test.

In **Figs. 5a** through **5c**, we explore the effect of mesh size (3-mm, 8-mm, and 12-mm squares) on the nature of the force-distance in crush. For the 3-mm and 8-mm meshes, we achieve a good estimate of crush force, with some differences in the perturbations due to different patterns in predicted buckle formation.

Although its peak force is good, the 12-mm mesh is of marginal use for design of the section. It may be usable in a coarse-mesh vehicle model even without tuning of load vs crush, but it is near the limit of realistic modeling results to capture the onset and progression of buckles. However, since there is a factor of approximately 64 in CPU time from 3-mm to 12-mm squares, it is important to find this optimum mesh for extension to full-vehicle studies. Mesh sizes of 12 mm (or even 24 mm) are about the minimum size that can be used throughout a full-vehicle model and still retain a workstation-level capability. This is because of the increase in degrees of freedom with the simultaneous decrease in time-step size due to the smaller elements.

Figure 6 provides a good illustration of the difficulties in capturing buckling behavior as the mesh is coarsened. On the left is the starting mesh (with 4-fold symmetry) for crush of a 6XXXAP tube, and on the right is the crushed geometry with 8-mm elements. The buckles are forming correctly, but it is clear that we are approaching the coarsest mesh within reason to capture the full buckling during crush. Beyond this level, pseudo-methods may be more effective in a full-vehicle model. However, even at this level, **Fig. 7** shows that once again a good approximation of the force-displacement curve in crush is achieved. For the AP material, a full piecewise plasticity model is needed to capture the high post-yield peak as the first buckle forms.

Vehicle Test and Design Analysis

We conducted a full-scale barrier impact test at LLNL's Site 300, using a 1978 Ford Fairmont instrumented with strain gages and accelerometers. One goal of the test was to develop a capability for testing under remote-steered conditions of impact structures. The remote-piloted steering system developed (**Fig. 8**) permitted this. Impact speed was just under 30 mph into a rigid concrete barrier. The results (**Fig. 9**) were captured on high-speed video. Comparison to DYNA calculations is in progress using a simple vehicle model.

The ALFAMAX steady-state, vehicle-handling algorithm was improved by the addition of a bending correction to the anti-sway-bar equations and subsequent performance calculations for a prototype of the Running Chassis design. These calculations demonstrate the unusual behavior of the near-constant-power, electric powertrain in acceleration, as shown in the time history of speed (mph) and forward acceleration (G's) in **Fig. 10**.

In terms of lateral acceleration, ALFAMAX demonstrated the potential for oversteer given the initial target of near 50/50 weight distribution if road camber at the front is held near zero (using negative roll camber). This scenario develops as we proceed from condition (A) to (B) in **Fig. 11**, which shows the progression of lateral G's with oversteer angle. Condition (C) shows the effect of adding a large, front anti-sway bar, resulting in high cornering G's but with development of understeer in the limit to avoid unstable handling behavior as in (B).

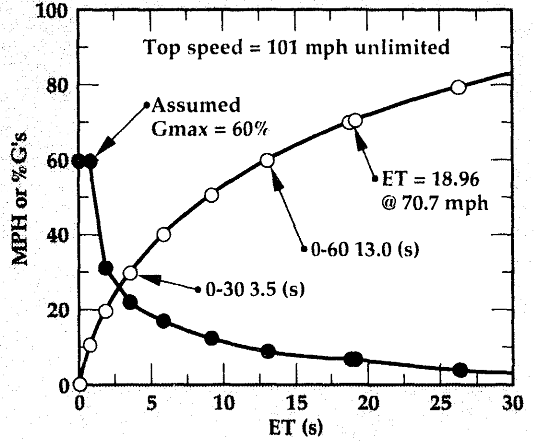

Figure 10. Forward acceleration performance of a prototype electric vehicle design.

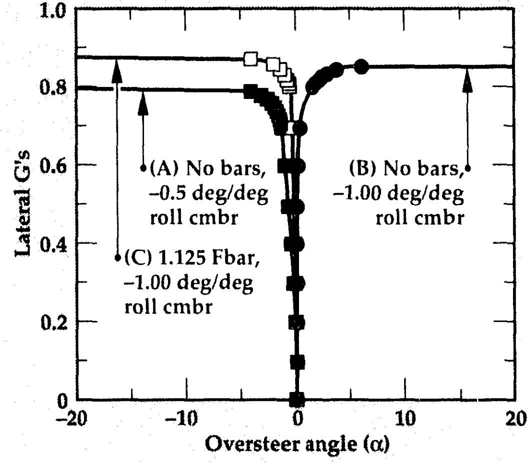

Figure 11. Lateral acceleration (steady-state skid-pad) for the prototype electric vehicle scenarios.

Acknowledgements

The author wishes to acknowledge the contributions of M.J. Burger and L.D. McMichael of LLNL, and of R.D. Parkinson of Kaiser Aluminum for the aluminum crush studies. K.A. Broughton and the TAHRS test teams of L.E. Sedlacek and R.B. Burdick developed the plan and systems for the LLNL Site 300 test.

---

1. R.W. Logan, "Tire, Accident, Handling, and Roadway Safety," *Engineering Research, Development, and Technology*, Lawrence Livermore National Laboratory, Livermore, California, UCRL-53868-92, 4-1 (1992).
2. R.G. Whirley, *DYNA3D: A Nonlinear, Explicit, Three-Dimensional Finite Element Code for Solid and Structural Mechanics—User Manual*, Lawrence Livermore National Laboratory, Livermore, California, UCRL-MA-107254 (1991).
3. R.W. Logan, M.J. Burger, L.D. McMichael, and R.D. Parkinson, "Crashworthiness Analysis Using Advanced Material Models in DYNA3D," *Crashworthiness and Occupant Protection in Transportation Systems*, AMD **169**, J.D. Reid and K.H. Yang (Eds.), ASME (1993).

Load Calculation and System Evaluation for Electric Vehicle Climate Control

Salvador Aceves-Saborio and
William J. Comfort
Nuclear Test Engineering
Mechanical Engineering

Providing air conditioning for electric vehicles (EV's) represents an important challenge, because vapor-compression air conditioners, which are common in gasoline-powered vehicles, may consume a substantial part of the total energy stored in the EV battery. Our work has two major parts: a cooling and heating load calculation for EV's, and an evaluation of several systems that can be used to provide the desired cooling and heating in EV's. Four cases are studied: short-range and full-range EV's are each analyzed twice, first with the regular vehicle equipment, and then with a fan and heat-reflecting windows, to reduce hot soak.

Our results indicate that for the batteries currently available for EV propulsion, an ice storage system has the minimum weight of all the systems considered. Vapor-compression air conditioners have the minimum weight for battery storage capacities above 270 kJ/kg.

Introduction

The state of California has mandated zero-emission vehicles by 1998. This opens the market for the introduction of electric vehicles (EV's), which currently represent the best developed technology for zero-emission vehicles. However, before EV's can be successfully marketed, five major issues must be addressed: cost, performance, convenience, passenger safety, and passenger comfort. Our work focuses on the last issue, evaluating alternative systems to accomplish appropriate climate control for EV's.

Heating and air conditioning (HVAC) are by far the car accessories that consume the most energy. HVAC does not represent a challenge for amply powered gasoline vehicles. On the other hand, EV's have a low energy-storage capacity, and HVAC may consume a substantial amount of the total energy stored, considerably reducing the vehicle range, which is one of the most important parameters for EV acceptability.

Recent legislation has allowed the use of combustion heating for EV's at ambient temperatures below 5°C. This legislation simplifies the problem of EV heating, which can be accomplished by using waste heat or electric resistance heating [or using the air conditioning (AC) system as a heater, if appropriate], at temperatures above 5°C, and turning to combustion heating as ambient temperature falls below 5°C. Air conditioning of EV's has, therefore, become the most important problem associated with EV climate control.

Progress

This report analyzes four combinations of vehicle equipment and driving scenarios. Short-range and full-range EV's are each analyzed twice, to evaluate the effect that load reduction systems may have on the thermal load: first with the regular vehicle equipment, and then with heat-reflecting (hot mirror) windows and a fan that circulates air when the vehicle is stopped, to reduce hot soak. Both vehicles are analyzed for a daily schedule of two trips, one starting from comfort conditions (assuming that the vehicle is cooled and protected from the sun during the charge process in the cooling season, and that the vehicle is heated during charge in the heating season), and the other starting from full hot or cold soak.

The trips have a duration of 30 min each for the short-range EV, and 60 min each for the full-range EV. A full description of the EV characteristics for the four cases is shown in **Table 1.**

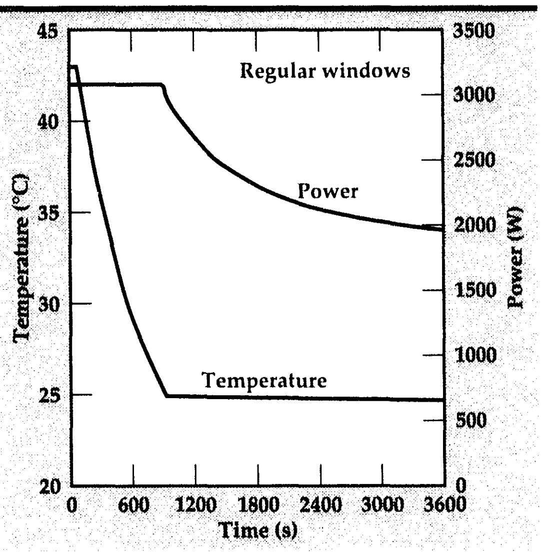

Figure 1. ***Air conditioner cooling power and air temperature inside the vehicle, as a function of time, for the regular glass vehicle.***

Our aim is to find the air conditioner that represents the best choice under a given set of requirements, and to determine how this best system changes as the requirements change. The figure of merit used is the overall weight of the system, which includes the weight of the battery required to provide the energy for operating the system. The best air conditioner has the minimum overall weight. This is equivalent to using the reduction in range as the figure of merit, considering that added system weight results in a decreased possibility of installing more batteries in the EV.

We have analyzed alternative systems for AC and have classified them into two categories. Systems in the first category, which are competitive with respect to vapor compression systems in overall system weight, have been analyzed with respect to their capacity to provide the heating required at temperatures above 5°C. Systems within this category are an ice-storage system and a zeolite-water adsorption system. A vapor-compression system is also analyzed, to compare its performance and total weight with those of the alternative systems. Systems in the second category are capable of providing the required AC. However, they either have an excessive system weight or a low efficiency, resulting in a high overall system weight, and are therefore not competitive with systems in the first category, at the current state of technology. A brief description of these systems is given, along with the possibilities for future system applicability.

Cooling- and Heating-Load Calculations

We have performed cooling and heating load calculations for the vehicles and driving conditions described in **Table 1**. **Figures 1** and **2** illustrate the results of our calculations. These figures show air conditioner cooling power and air temperature inside the vehicle, for the regular glass and for the hot mirror vehicle, respectively. The regular glass vehicle requires 3080 W of initial power to reach comfort conditions in 15 min. After this point, the air conditioner cooling power is controlled to keep a fixed (comfort) temperature inside the vehicle, and therefore the required power drops with time, approaching the steady-state value. The same behavior is observed in **Fig. 2** for the hot mirror vehicle, for which the required initial power is equal to 2070 W.

The contributions to the heating load in steady state include metabolic (negative contribution), ventilation, conduction, and defrosting.

Table 1. ***Electric vehicle characteristics.***

All cases considered have the dimensions of subcompact vehicles, and include the following thermal load reduction equipment:

(1) a roof insulated with 0.02 m of high-performance thermal insulation; and
(2) operation by minimizing the amount of external ventilation.

Vehicles are used for two daily trips. The first trip starts from comfort conditions; the second trip starts from full soak conditions. Load calculation considers two occupants during summer operation, and a single occupant during winter operation. The four cases considered are as follows:

Case 1. Short-range electric vehicle, used for two trips daily, each lasting 30 min. All the vehicle windows are made of regular glass, with heat-absorbing coating on the top 0.10 m of the windshield.

Case 2. Short-range electric vehicle, used for two trips daily, each lasting 30 min. All the vehicle windows are made of heat-reflecting (hot mirror) glass. A solar-powered fan operates whenever the vehicle is left in the sun, to reduce hot soak temperature.

Case 3. Full-range electric vehicle, used for two trips daily, each lasting 60 min. Same vehicle characteristics as for Case 1.

Case 4. Full-range electric vehicle, used for two trips daily, each lasting 60 min. Same vehicle characteristics as for Case 2.

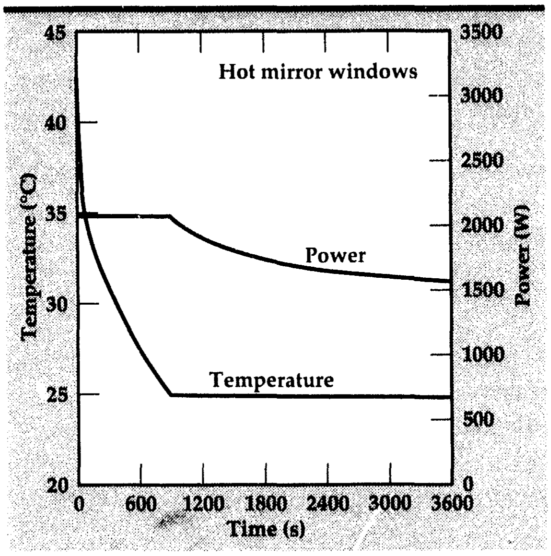

Figure 2. Air conditioner cooling power and air temperature inside the vehicle, as a function of time, for the hot mirror glass vehicle.

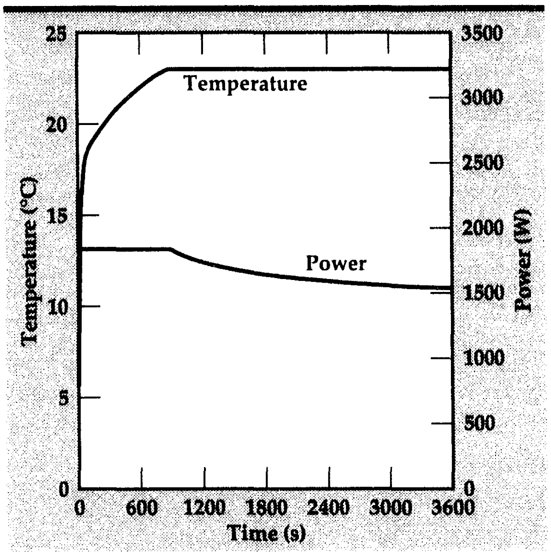

Figure 3. Required heating power and air temperature inside the vehicle, as a function of time.

The heating load from cold soak is calculated by assuming that all the vehicle surfaces have a temperature equal to the ambient temperature (5°C), at the moment in which vehicle operation starts. The required heating power and energy are calculated with the same computer code used for the transient cooling load analysis, assuming a heating period of 15 min, starting from hot soak.

Figure 3 illustrates the results, showing the temperature inside the vehicle, as well as the heating power as a function of time. The initial power requirement to reach comfort conditions in 15 min is equal to 1830 W.

Table 2 lists the energy required for cooling and heating the vehicle. The total energy requirement is obtained by adding the energy for the initial trip (starting from comfort conditions) to the energy of the second trip (starting from hot or cold soak conditions). The energy for the first trip is calculated by multiplying the steady-state load by the time length of the trip. The energy for the second trip is obtained by integrating the power requirement curves (**Figs.** 1, 2, and 3). The results for the cooling load are obtained for the ambient conditions existing in the desert southwest, where the cooling load is close to the maximum in the country. The cooling loads are used to calculate the system weight and energy consumption for the AC systems.

The results for the heating load indicate the maximum possible amount of heating that has to be provided to the vehicle before the combustion heater can be used. Although small compared to the heating requirements at low ambient temperatures if no combustion heating is allowed, the total energy requirement for heating is substantial, and could reduce the vehicle range considerably if only electric resistance heating is used to provide it. The following two alternatives are considered:

(1) Use the air conditioner to provide heating. The applicability of this alternative depends on the type of air conditioner being used.

(2) Use heat recovery from the power controller and motor, to heat the vehicle. Assuming a

Table 2. Cooling and heating energy requirements for all the cases considered. First trip starts from comfort conditions. Second trip starts from a full hot or cold soak.

		Energy required, J	
Cooling	1st trip	2nd trip	Total
1. Regular glass, 30 min	3.27×10^6	5.11×10^6	8.38×10^6
2. Hot mirror, 30 min	2.69×10^6	3.55×10^6	6.24×10^6
3. Regular glass, 1 h	6.54×10^6	8.90×10^6	1.54×10^7
4. Hot mirror, 1 h	5.38×10^6	6.50×10^6	1.19×10^7
Heating			
1. Ambient 5°C, 30 min	2.70×10^6	3.19×10^6	5.89×10^6
2. Ambient 5°C, 1 h	5.40×10^6	6.04×10^6	1.14×10^7

6-kW steady power for the vehicle, and power controller and motor efficiencies of 0.9 and 0.94 respectively, results in ~ 0.9 kW of waste heat, which is approximately equal to 50% of the required heating load, and could be used to complement the heating obtained from either resistance heating or air conditioner heating.

Air Conditioner System Evaluation

This section gives a description and evaluation of three systems used to satisfy the cooling loads specified in **Table 2**. These systems, vapor compression, ice storage, and adsorption air conditioners, have been classified as having the best available characteristics for EV applications. Descriptions in this report include the overall weight (system plus battery weight to provide energy for operation), which is used as the figure of merit in this evaluation.

The storage capacity for lead-acid batteries, used in this analysis, is approximately 122 kJ/kg at no load; 94 kJ/kg at the vehicle driving load; and 64 kJ/kg, considering a discharge efficiency of 86% and a depth of discharge of 80%. These figures, which are representative of high-performance, current technology lead-acid batteries,[1] are used for the base designs. However, battery storage capacity is expected to increase in the future, as better storage technologies are developed, or existing technologies become less expensive. We have taken this possibility into account by studying the effect of battery storage capacity on overall system weight for the three systems analyzed.

System descriptions also include a calculation of the heating performance for the three systems. Applicability of the AC systems for heating at temperatures above 5°C, for which combustion heating is not allowed, is of great importance to avoid significant range reductions during the winter months. Additional systems not considered to be competitive in terms of overall system weight are also described briefly.

Table 3. System weight by component and total volume for a vapor-compression air conditioner system, for the energy requirements given in Table 2 for the four cases described in Table 1.

	Case number			
System component	**1**	**2**	**3**	**4**
Air conditioner weight (kg)	24	21	24	21
Battery weight (kg)	60	44	109	85
Overall weight (kg)	84	65	133	106
Total volume (l)	44	34	69	55

Vapor Compression Air Conditioner

This system has been in use in most vehicles for a long time, and is therefore the system of choice, unless an alternative system proves to be substantially better in terms of range reduction and/or cost. Variable-speed scroll compressors driven by electric motors have been developed recently, and represent an improvement that suits the applicability of vapor compression air conditioners to EV's extremely well.

The system and battery weight for the four combinations of vehicle and trip length are listed in **Table 3**. The maximum AC power requirements for the vehicles considered in the present analysis are 3.08 kW and 2.07 kW for the regular glass and for the hot mirror glass vehicles, respectively. While no manufacturer of variable-speed compressors of the capacity required in the present work is known to the authors, a correlation has been derived from data published in manufacturers' catalogs for scroll compressor systems that predicts a 3-kg drop in system weight per kW of cooling power reduction. This correlation is applied to the calculation of the system weights given in **Table 3**. The calculations assume a system COP equal to 2.2, including parasitic losses.

The vapor-compression air conditioner can easily be adapted to operate as a heat pump, which can provide all of the heating required at temperatures above 5°C. Considering a heating COP equal to 2.5, the system provides the total heating requirement while consuming the energy content of 37 kg of batteries for two 30-min trips, and 71.5 kg of batteries for two 60-min trips.

Ice Storage System

In this system (**Fig. 4**), ice is stored in the shell side of a shell and tube heat exchanger. Ice is made while the batteries are recharged. Two system configurations are possible. The first configuration uses an external refrigerator to cool alcohol or other coolant with a low freezing point and appropriate characteristics. The external refrigerator is connected to the vehicle during charge, and the coolant circulates through the ice storage until all the water inside the storage freezes. Once the freezing process is completed, the external refrigerator is disconnected from the vehicle, remaining in the charging location. When AC is required during vehicle operation, the coolant is circulated through the ice storage and through a liquid-to-gas heat exchanger, cooling the air that circulates into the vehicle. This system has the double advantage of not requiring a refrigeration system onboard, and requiring very little battery energy during operation. Battery energy is only required for fan and pump operation. Both of these factors help in reducing overall system weight. Two disadvantages of this system are (1) the need to connect and disconnect the external refrigerator every day; and (2) the requirement for an external refrigeration system at the charging location.

The second possible configuration includes a small, low-power, vapor-compression refrigerator, which increases the system weight, as part of the vehicle equipment to freeze the water during charging. The vapor-compression freezer is designed to operate only when the batteries are being charged, thereby limiting the amount of battery energy required for system operation.

The system weights for the four combinations of vehicle characteristics and trip lengths are shown in **Table 4**. Weights for the two system configurations described above are identical, with the only exception that the second configuration includes the additional weight of the refrigeration system, estimated at 20 kg. The system is designed by using a computer model that calculates the required coolant tube spacing inside the ice storage to provide appropriate cooling power in all cases. The ice storage (shell and tube heat exchanger) and the water-to-air heat exchanger are made of aluminum for reduced weight. Considering 10% total losses, and an average freezer COP equal to 1.1, the system COP is estimated as 1.0. Both configurations of the system have a lower overall weight than the vapor-compression air conditioner, making this system an attractive choice for EV AC. An additional advantage is the possibility of being able to easily remove the ice storage (or at least the water) from the vehicle whenever cooling is not required. System volume is, however, higher than the volume for the vapor-compression system, due to the relatively low density of ice. This may be a limitation for high energy requirements.

Table 4 shows an estimated weight of 10 kg for a heat recovery system, which includes a heat exchanger, ducts, and a fan. This extra weight is necessary because ice-storage-system performance for heating is limited. The only way this system can be used for heating is by heating the water in the storage during the charge process, and using the hot water during vehicle operation. If water can be heated to 90℃, and heat losses are limited to 10%, the system designs for the regular glass vehicles (Cases 1 and 3) can provide approximately 80% of the heating load at 5℃ (**Table 2**). The system designs for the vehicle equipped with hot mirror windows (Cases 2 and 4) can provide about 60% of the heating load at 5℃. A heat-recovery system can provide approximately 50% of the heating load, and provides an adequate complement to the ice-storage system, for heating.

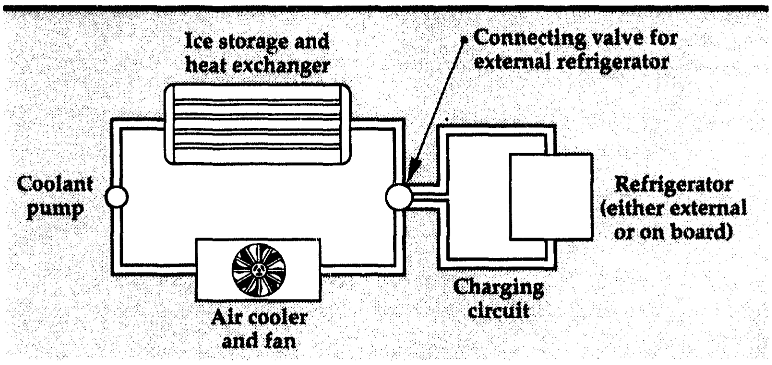

Figure 4. Schematic diagram of an ice-storage AC system.

Table 4. System weight by component and total volume for the two configurations of an ice-storage air conditioner system, for the energy requirements given in Table 2 for the four cases described in Table 1.

Configuration 1. No freezer in the vehicle.

	Case number			
System component	**1**	**2**	**3**	**4**
Ice stored (kg)	25	19	46	35
Ice storage/ Heat exchanger (kg)	7	6	11	9
Water-to-air heat exchanger, pump and fan (kg)	7	9	8	10
Heat recovery system (kg)	10	10	10	10
Battery weight (kg)	3	3	6	6
Overall weight (kg)	52	47	81	70
Total volume (l)	75	58	109	88

Configuration 2. Freezer in the vehicle

	Case number			
System component	1	2	3	4
Ice stored (kg)	25	19	46	35
Ice storage/ Heat exchanger (kg)	7	6	11	9
Water-to-air heat exchanger, pump and fan (kg)	7	9	8	10
Freezer (kg)	20	20	20	20
Heat recovery system (kg)	10	10	10	10
Battery weight (kg)	3	3	6	6
Overall weight (kg)	72	67	101	90
Total volume (l)	85	68	119	98

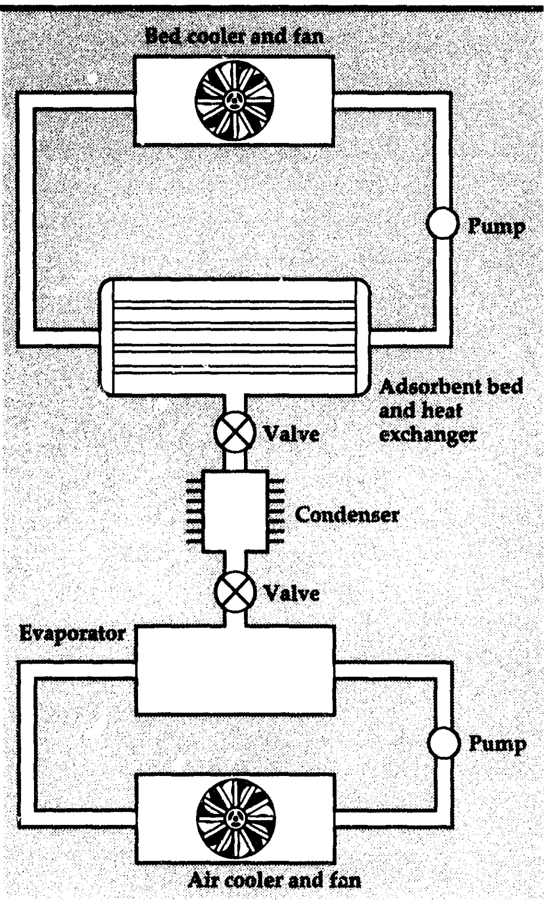

Figure 5. Schematic diagram of an adsorption AC system.

Adsorption Air Conditioner

Adsorption systems operate by taking advantage of the affinity that some substances such as silica gel, zeolites, and activated carbon (adsorbents), have for certain liquids or gases (adsorbates). This affinity tends to reduce the vapor pressure of the adsorbates, which, for liquid adsorbates, results in partial evaporation and cooling of the remaining liquid. The cold liquid can then be used for AC.

Figure 5 is a schematic diagram of an adsorption air conditioner. The operation cycle is divided into two parts, adsorption and regeneration. For adsorption, the valves between the evaporator and the adsorbent are opened, resulting in evaporation and a cold adsorbate that can then be used for cooling. However, adsorbents can only adsorb a limited amount of adsorbate before it is necessary to desorb (regenerate the adsorbent) to continue operating the system. The regeneration cycle is usually accomplished by heating the adsorbent to temperatures above 100 °C. Once the regeneration cycle is completed, the cooling cycle can be started again. This cyclic operation is very well suited for EV's, since the regeneration cycle can be done during the charge period, and the cooling effect is obtained during vehicle operation without consuming battery power, except for the power required for fan and pump operation. A more detailed description of adsorption systems is given by Suzuki.[2]

Attraction between the adsorbent and the adsorbate is due to either surface forces (physical sorption), or chemical bonds (chemisorption). Most systems studied to date, including the zeolite-water system analyzed in this section, are physical sorption systems, for which property data is more readily available. Zeolite is used because it can adsorb a great amount of vapor, and water is used because it has the highest latent heat of vaporization among the common liquids, providing a high cooling effect for a given adsorbed mass. This high cooling effect for water makes it the best choice, even though water adsorption systems necessarily operate under vacuum conditions, requiring the use of pressure vessels within the system, adding to the system weight.

The resulting weights and volumes for the adsorption systems designed for the four cases being considered, are given in **Table 5**. System COP's are on the order of 0.4, inferior to those obtained with the other systems described in this section.

Adsorption cycles can provide heating without difficulty, by using the energy released by the adsorbent when adsorption takes place. The heat of adsorption for the zeolite-water system

Table 5. System weight by component and total volume for a zeolite-water adsorption air conditioner, for the energy requirements given in Table 2 for the four cases described in Table 1.

System component	Case number 1	2	3	4
Adsorbent (kg)	28	19	44	32
Adsorbent storage (kg)	27	19	38	29
Water (kg)	8	8	11	10
Evaporator (kg)	4	4	5	4
Condenser (kg)	4	4	5	4
Hot heat exchanger, pump and fan (kg)	6	5	6	6
Cold heat exchanger, pump and fan (kg)	6	5	6	6
Battery weight (kg)	6	6	12	12
Overall weight (kg)	89	70	127	103
Total volume (l)	107	77	152	119

is approximately 30% higher than the latent heat of vaporization of water. Therefore, the zeolite-water system has the (theoretical) capability to provide 30% more heating than it can provide cooling. Consulting the heating requirements given in **Table 2**, it can be seen that the four adsorption systems are capable of providing enough heating at a 5°C ambient temperature.

Battery Capacity

The weight calculations for the four systems considered as appropriate air conditioners for EV's were carried out for lead-acid batteries which are currently the best developed technology for EV propulsion. However, alternative batteries with a higher energy-storage capacity are being developed, or are available at a high cost. As new technologies become available, or the cost of existing technologies drops, the results of the calculations shown in the previous sections are expected to change. **Figure 6** shows the effect of the battery capacity on system weight for the four systems. As expected, vapor-compression air conditioners benefit the most from improved batteries, and surpass the ice-storage system as the lightest system for large battery storage capacities. The battery storage capacity at which the vapor-compression air conditioner becomes the lightest system is a function of the energy requirement for the system, and is approximately equal to 270 kJ/kg for Case 1 (regular windows, two 30-min trips), dropping to values in the range between 200 and 250 kJ/kg for the other cases.

The figures also indicate that the ice-storage system that includes the freezer is also lighter than the vapor-compression air conditioner, for energy capacities under 150 kJ/kg, for all cases except Case 2 (the lowest energy requirement), for which the vapor-compression system is lighter, even for lead-acid batteries. Adsorption systems have a lower system weight than vapor-compression systems only for low battery capacities and high thermal loads (Cases 3 and 4).

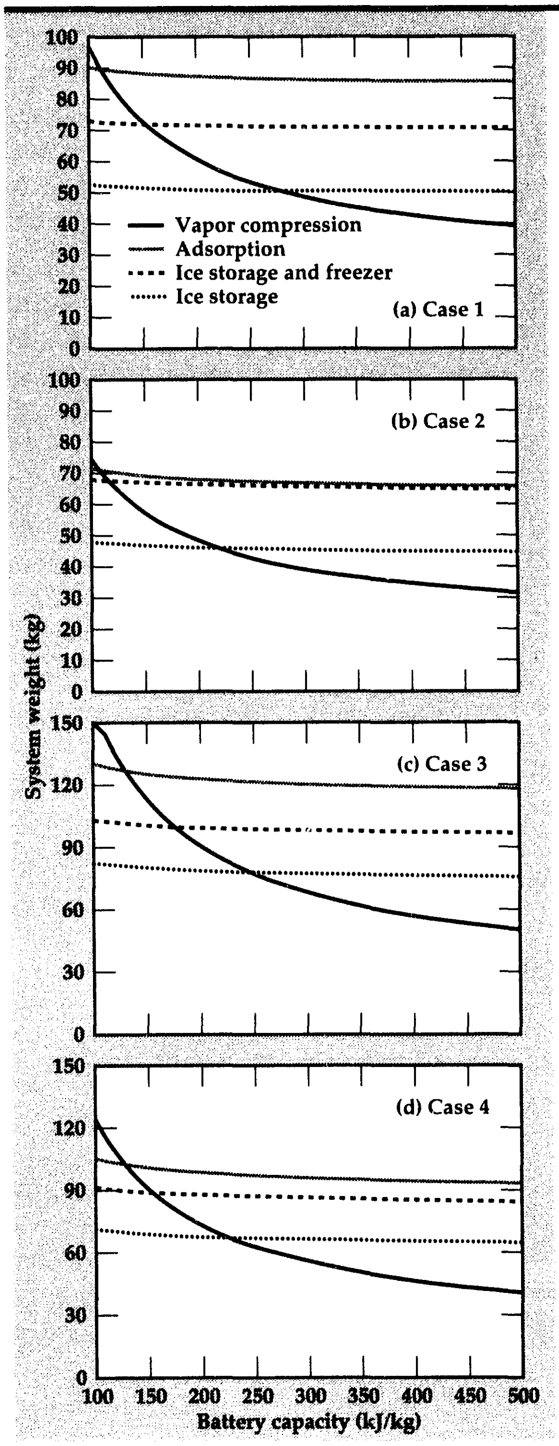

Figure 6. System weight as a function of battery capacity for the four cases described in Table 1. Battery capacities indicate energy storage at zero load (equal to 122 kJ/kg for lead-acid batteries), without taking into account the discharge efficiency and the maximum depth of discharge.

Other AC Systems

The following additional AC systems have been analyzed and considered inappropriate for EV cooling, due to a high system weight or high energy consumption, or both:

(1) *Thermoelectrics.* Current thermoelectric coolers consume a significant part of the total battery energy. Research projects currently under way are aimed at substantially improving thermoelectric cooler performance. However, even obtaining thermoelectric coolers with a figure of merit equal to 0.006, which is twice as large as currently available for thermoelectrics, yields a COP of only 1.35, calculated for a 40°C temperature difference between hot and cold sides. Improved EV

batteries can also help in making thermoelectrics a better choice for EV cooling.

(2) *Stirling cycles.* Stirling cycles have a theoretical efficiency equal to the Carnot efficiency, and therefore have the potential to reach higher efficiencies than those possible with a vapor-compression air conditioner. Stirling cycles are commonly used for cryogenic applications, and have the additional advantage of not requiring a CFC refrigerant. However, at this time, Stirling cycles are not competitive in cost, performance, or reliability, compared to vapor-compression air conditioners for EV's.

(3) *Brayton cycles.* Brayton air conditioners have as their major advantage a low system weight, which has resulted in a broad applicability to aircraft. However, their COP is usually less than 1, which results in an excessive energy consumption for EV's. It is doubtful if turbomachinery with the high efficiencies required will be developed in the near future.

(4) *Vacuum boiling of water.* In this system, a vacuum pump is used to reduce the pressure of a water container, to the point at which the water boils. This causes a drop in temperature for the liquid water that remains in the container, which can then be used for AC. The system has a high maximum-theoretical efficiency, but the vacuum pump required to boil enough water to cool the vehicle is excessively heavy, resulting in a large overall weight.

(5) *Heat-actuated air conditioners.* These systems, which include ejector, absorption, adsorption with regeneration, and metal hydride systems, cannot be applied with advantage to EV's due to the small amount and low temperature of the waste heat available, and also due to the low COP that these systems normally reach.

Conclusions

Our most important conclusions are as follows:

(1) Reducing the thermal loads is very important for providing appropriate AC without an excessive system weight. Heat transfer into the EV can be decreased considerably by using special window glasses, extra insulation, and reduced external ventilation, and by operating a fan when the vehicle is left in the sun, to reduce hot soak.

(2) Ice-storage systems represent the best alternative for EV AC. Ice-storage systems may be installed with or without a vapor-compression refrigerator (freezer) in the vehicle. While using an external freezer is more desirable in terms of system weight, it may be necessary to install a freezer in the vehicle if using the external freezer is considered impractical. In either case, ice-storage systems consume very little energy during EV operation, and have the lowest overall weight of all the systems analyzed.

(3) Ice-storage systems present two disadvantages: the large volume that they require, especially for large cooling requirements, and their inability to provide enough heating for an ambient temperature of 5°C. Therefore, this system needs to be complemented with a heat-recovery system to provide all the heating required.

(4) The zeolite-water adsorption systems analyzed have approximately the same overall weight as vapor-compression air conditioners. While they can provide all the required energy for heating down to 5°C, their COP is low, and they probably require further improvement before they can be extensively applied to EV AC.

(5) Increasing the battery storage capacity from its value for lead-acid batteries (122 kJ/kg) results in a great weight reduction for vapor-compression air conditioners. These systems eventually reach and surpass ice-storage systems as the system with the least overall weight. The battery storage capacity at which vapor-compression air conditioners become the lightest system varies as a function of the cooling energy requirement, from approximately 220 to 270 kJ/kg. Vapor-compression air conditioners will become the best method for EV cooling once batteries with this capacity become available at a reasonable price.

---

1. J. Hardin, "Battery Testing," *Proc. Automotive Technology Development Contractors Coordination Meeting* (Dearborn, Michigan), (1992).
2. M. Suzuki, *Adsorption Engineering*, Kodansha LTD, Tokyo, Japan (1990).

Paste Extrudable Explosive Aft Charge for Multi-Stage Munitions

Scott A. Ruddell,
Douglas R. Faux, and
Robert M. Kuklo
Nuclear Explosives Engineering
Mechanical Engineering

The paste extrudable explosive (PEX) aft charge project is a multi-year effort to determine the feasibility of using a paste explosive aft charge in a multi-stage munition and, ultimately, to demonstrate this in a fully functional weapon. Our FY-93 accomplishments were quite satisfying: predictions on the high probability of detonation with detonators embedded in the paste stream were correct; sympathetic detonation was found to be of no concern; all timing information extracted from the analysis and testing proved valid. With the exception of the detonator axial position vs time, the analytical models provided sound results. A series of successful tests culminating in an actual detonation shows that PEX holds great promise as a viable follow-through charge option in a multi-stage conventional munition.

Introduction

Multi-stage conventional munitions are used widely for demolition applications. These weapons typically have a two-stage warhead: a forward charge that produces a borehole in the target, and an aft charge that enters the borehole and then detonates, destroying the target. The aft charge is usually a steel-encased explosive that either enters the borehole by its own kinetic energy or is 'driven' into the borehole by a rocket or velocity augmenter.

To increase the versatility and reduce the weight of a portable, multi-stage munition, a paste extrudable explosive (PEX) aft charge that injects PEX into the borehole formed by the forward charge replaces the steel-encased aft charge. The PEX aft charge can be used with a smaller, less perfect borehole and provides greater coupling of the explosive with the target.

Progress

The PEX aft charge project is a multi-year effort to determine the feasibility of using a paste explosiveaft charge in a multi-stage munition and, ultimately, to demonstrate this in a fully functional weapon.

The PEX aft charge was conceived as an improvement for the man-portable penetration augmented munition (PAM) currently being developed for U.S. Special Operations Forces. PAM is a typical multi-stage munition with an additional forward shaped charge designed to remove hard metallic obstructions, such as rebar reinforcement in concrete found near the surface of a target, that may hinder proper emplacement of the steel-encased aft charge. By substituting PEX for the current PAM aft charge, the rebar cutting charge can be eliminated as well as the steel casing for the aft charge, resulting in a considerable weight savings. This not only improves the portability of the PAM but also enhances its successor, the standoff destruct munition (SODM, or 'flying PAM') illustrated in **Fig. 1**.

Our FY-93 development plan consisted of (1) multiple tests to evaluate the borehole fill characteristics and potential for sympathetic detonation of the PEX; (2) additional testing to study effectiveness of different initiation schemes; and (3) further testing and analysis to help validate the computer code simulations (CALE[1] and DYNA2D[2]).

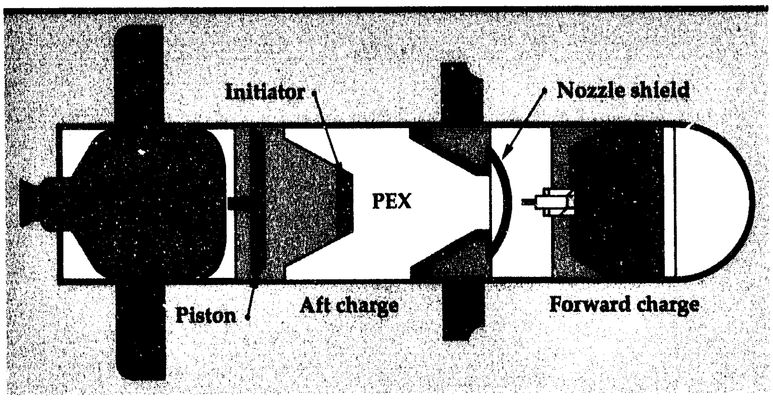

Figure 1. Conceptual sketch of a SODM with a PEX aft charge.

Mock PEX Testing and Analysis

All tests involved the use of a gas-driven injector system ('PEX gun') to eject five pounds of mock PEX at approximately 150 m/s out of a 3-in.-dia nozzle into a simulated borehole target. Through the use of high-speed-camera footage and flash-x ray imaging, the PEX flow characteristics both prior to and during the fill process were evaluated and compared to the analytical results.

Once again, the DYNA2D predictions for the PEX stream formation compared favorably with the experimental results. Validation of the CALE calculations was not possible with the particular test configuration and diagnostics chosen. More testing will be necessary to further explore the accuracy of the CALE results.

Pressure transducers were placed at strategic locations at the base of the target bore. These pressure records provided valuable information on the likelihood of a sympathetic detonation caused by impact at the bottom of the bore. All maximum pressures recorded were well below the predicted detonation pressure of the PEX.

Only one detonation scheme was pursued this fiscal year. Three, small, chemically delayed detonators were embedded in mock PEX, and their behavior monitored with flash x rays during a fill test. All three detonators remained with the paste throughout the 20-ms fill process. Unfortunately, the first-cut finite element model of the paste stream with embedded detonators did not adequately predict the detonator behavior observed experimentally. The errors were undoubtedly a result of the way the detonators were tied to the PEX in the model and not of a problem with the code itself. More work is needed to enhance the accuracy of this model.

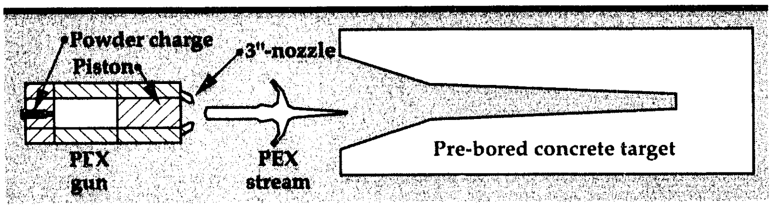

Figure 2. PEX stream formation shown in mid-flight.

Live PEX Test

To verify predictions, assumptions, and experimental findings, an actual detonation was necessary. The gun was loaded with real PEX and live delay detonators, and fired into a pre-bored concrete target (see **Fig. 2**). The PEX detonated exactly as predicted, resulting in significant damage to the concrete target. High-speed cameras were again used to capture the paste stream formation. No significant differences between the flow behaviors of the mock and the real PEX were observed.

Future Work

If this project continues, FY-94 objectives will be: (1) to design and test a blast shield to protect the aft charge from the explosive impact of the boring charge; (2) to devise a test series to validate the CALE calculations; (3) to modify the DYNA2D model to achieve a more accurate representation of the detonator behavior; (4) to evaluate an alternate initiation scheme involving a detonation cord pulled along with the PEX; and (5) to conduct a system prototype test combining the PEX with preferred detonation scheme, blast shield, and forward boring charge.

Acknowledgements

The authors wish to acknowledge those individuals who supported the bunker 851 PEX activities at the Lawrence Livermore National Laboratory Site 300 test facility. Without their cooperation, flexibility, and versatility, none of this could have been possible.

---

1. R. Tipton, *CALE—User Manual*, Lawrence Livermore National Laboratory, Livermore, California, UCRL-MA-920701 (1992).
2. J.O. Hallquist, *User's Manual for DYNA2D—An Explicit Two-Dimensional Hydrodynamic Finite Element Code with Interactive Rezoning and Graphical Display*, Lawrence Livermore National Laboratory, Livermore, California, UCID-18756, Rev. 3 (1988).

Flexible Manufacturing for Photonics Device Assembly

Shin-yee Lu
K. David Young
Engineering Research Division
Electronics Engineering

The assembly of photonics devices such as laser diodes, optical modulators, and opto-electronics (OE) multi-chip modules usually requires the placement of micron-size devices, and sub-micron precision attachment between optical fibers and diodes or waveguide modulators (pigtailing). This is a labor-intensive process. Studies done by the OE industry have shown that 95% of the cost of a pigtailed photonic device is attributed to the current practice of manual alignment and bonding techniques. At Lawrence Livermore National Laboratory, we are working to reduce the cost of packaging OE devices, through the use of automation.

Introduction

At present, the cost of opto-electronic (OE) devices is dominated by the effort required to package those devices into an integrated system. Components such as laser diodes and modulators, designed for high-performance applications, are single-mode devices. They must be connected using optical fibers or other types of waveguides with sub-micron alignment accuracies. Presently, OE packaging is usually performed by highly skilled technicians looking through microscopes and manually adjusting sub-micron positioning stages. For single-mode fibers, six degrees of freedom for positioning are sometimes required. Once alignment is achieved, the components must be held in place using epoxy, solder, or other attachment techniques, and realigned before the gluing process settles. This labor-intensive process results in only a few packages per day by each technician. The packaging costs are by far the highest fraction of the total cost of an assembled OE package. The consequences of this low-volume, labor-intensive process are readily apparent. The costs are too high to allow the advantages of fiber optics to penetrate such markets as on-chip interconnects, interboard connections in computers, and local area networks.

We believe that the packaging process must be automated to significantly reduce the costs of OE devices. The electronics industry has successfully reduced the costs of its products through the massive use of automation, including alignment, parts handling and feeding, and *in situ* quality control. Unfortunately, the sub-micron precisions and six-axis alignment required for OE packaging greatly exceed the requirements of electronics device automation. The automated systems developed to

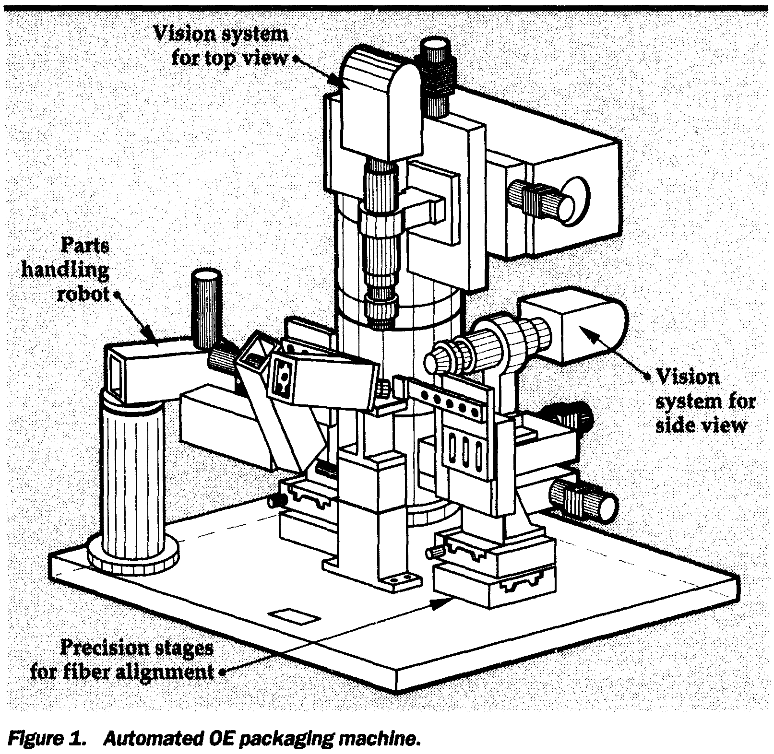

Figure 1. Automated OE packaging machine.

assemble integrated circuits cannot be applied to the problem of packaging OE circuits. We are currently designing and building a flexible workcell for the automation of the packaging of OE devices. The initial task is to align a single-mode fiber to each end of a Mach–Zehnder waveguide modulator (fiber pigtailing). We are also developing silicon micro-bench-based techniques to provide a packaging technology that is compatible with automated alignment and pigtailing.

Progress

The design of the workcell (**Fig. 1**) uses machine vision to provide position information for low resolutions (microns), long travel movement, and optic power throughput optimization to provide final high-resolution, small-travel alignment. The combined vision/optic power feedback approach provides a low-cost output metrology for high-resolution positioners. It also eliminates the need for high-precision fixturing, calibration, and temperature control to hold photonics devices in fixed locations at micron precision. The workcell should be able to handle a wide variety of different sizes and attachment geometries of photonic devices.

The current workcell has eleven axes: a five-axis dual camera system, and two sets of three-axis, high-precision stages for fiber manipulation. The dual camera system generates position information from the top and the side view of the waveguide modulator chip and fiber. The top camera subsystem has three degrees of freedom for moving along a plane above the chip and for focusing; the side camera subsystem has two degrees of freedom for moving along a side of the chip, as well as for focusing. The pixel resolution is 1.4 μ. The motion control hardware and software for the vision system are chosen such that the position accuracy is compatible with the image resolution. The reason for the high position repeatability requirement of the vision system is to register the two coordinate systems defined by two images due to: (1) the move from one field of view to the next along a certain trajectory; and (2) the move between two focal planes. The misalignment of the mechanical system is calibrated by using the vision system itself for output metrology.[1]

We have developed computer vision algorithms for fiber and waveguide detection, and for kinematic calibration of the camera system. We are currently developing control algorithms for fiber manipulation.

Future Work

We plan to demonstrate high-precision fiber alignment, and to characterize the alignment performance in the near future. We plan to start the assembly of a miniature robot arm for chip placement, and for fiber attachment using UV curing adhesive. The workcell will be ready to test pigtailing a lithium niobate waveguide modulator by the end of FY-94.

---

1. S.-y. Lu, M.D. Pocha, O.T. Strand, and K.D. Young, *Flexible Manufacturing for Photonics Device Assembly*, Lawrence Livermore National Laboratory, Livermore, California, UCRL-JC-114204 (1993).

Fabrication Technology

The mission of the Fabrication Technology thrust area at Lawrence Livermore National Laboratory (LLNL) is to have an adequate base of manufacturing technology, not necessarily in-house, to conduct the future business of LLNL. Our specific goals continue to be to (1) develop an understanding of fundamental fabrication processes; (2) construct general-purpose process models that will have wide applicability; (3) document our findings and models in journals; (4) transfer technology to LLNL programs, industry, and colleagues; and (5) develop continuing relationships with the industrial and academic communities to advance our collective understanding of fabrication processes. Our niche has been and will continue to be in bringing our strengths in precision engineering to bear on problems in manufacturing.

In pursuit of the above goals for FY-93, the thrust area sponsored three projects. First, we continued to develop applications for molecular modeling, a technology on the verge of becoming an engineering tool. Chronicled here is its application to the understanding of the plasticity observed in very small indentations of silica. In the second project, we began characterizing the limit to which a commercial, double-sided lap can generate flat optical components. This is done both to understand the engineering limits of a manufacturing process and to prepare ourselves for the purchase of many flat components required by the National Ignition Facility. Our third project is an exercise to capitalize on a diamond-metal composite, originally developed for its thermal properties, by using its concomitant unique structural properties. The chemical bond between the diamond and metal, which gives high thermal conductivity, also gives high mechanical strength. This makes it suitable as a grinding wheel.

Over the next year, we will complete our characterization of the double-sided lap and conclude the development of a diamond composite grinding wheel. We will embark on developing and characterizing a new series of beryllium-based multilayers for their high reflectivity at wavelengths under 13 nm. A better understanding of beryllium-based multilayers has direct benefit to soft x ray projection lithography.

Kenneth L Blaedel
Thrust Area Leader

Section 5

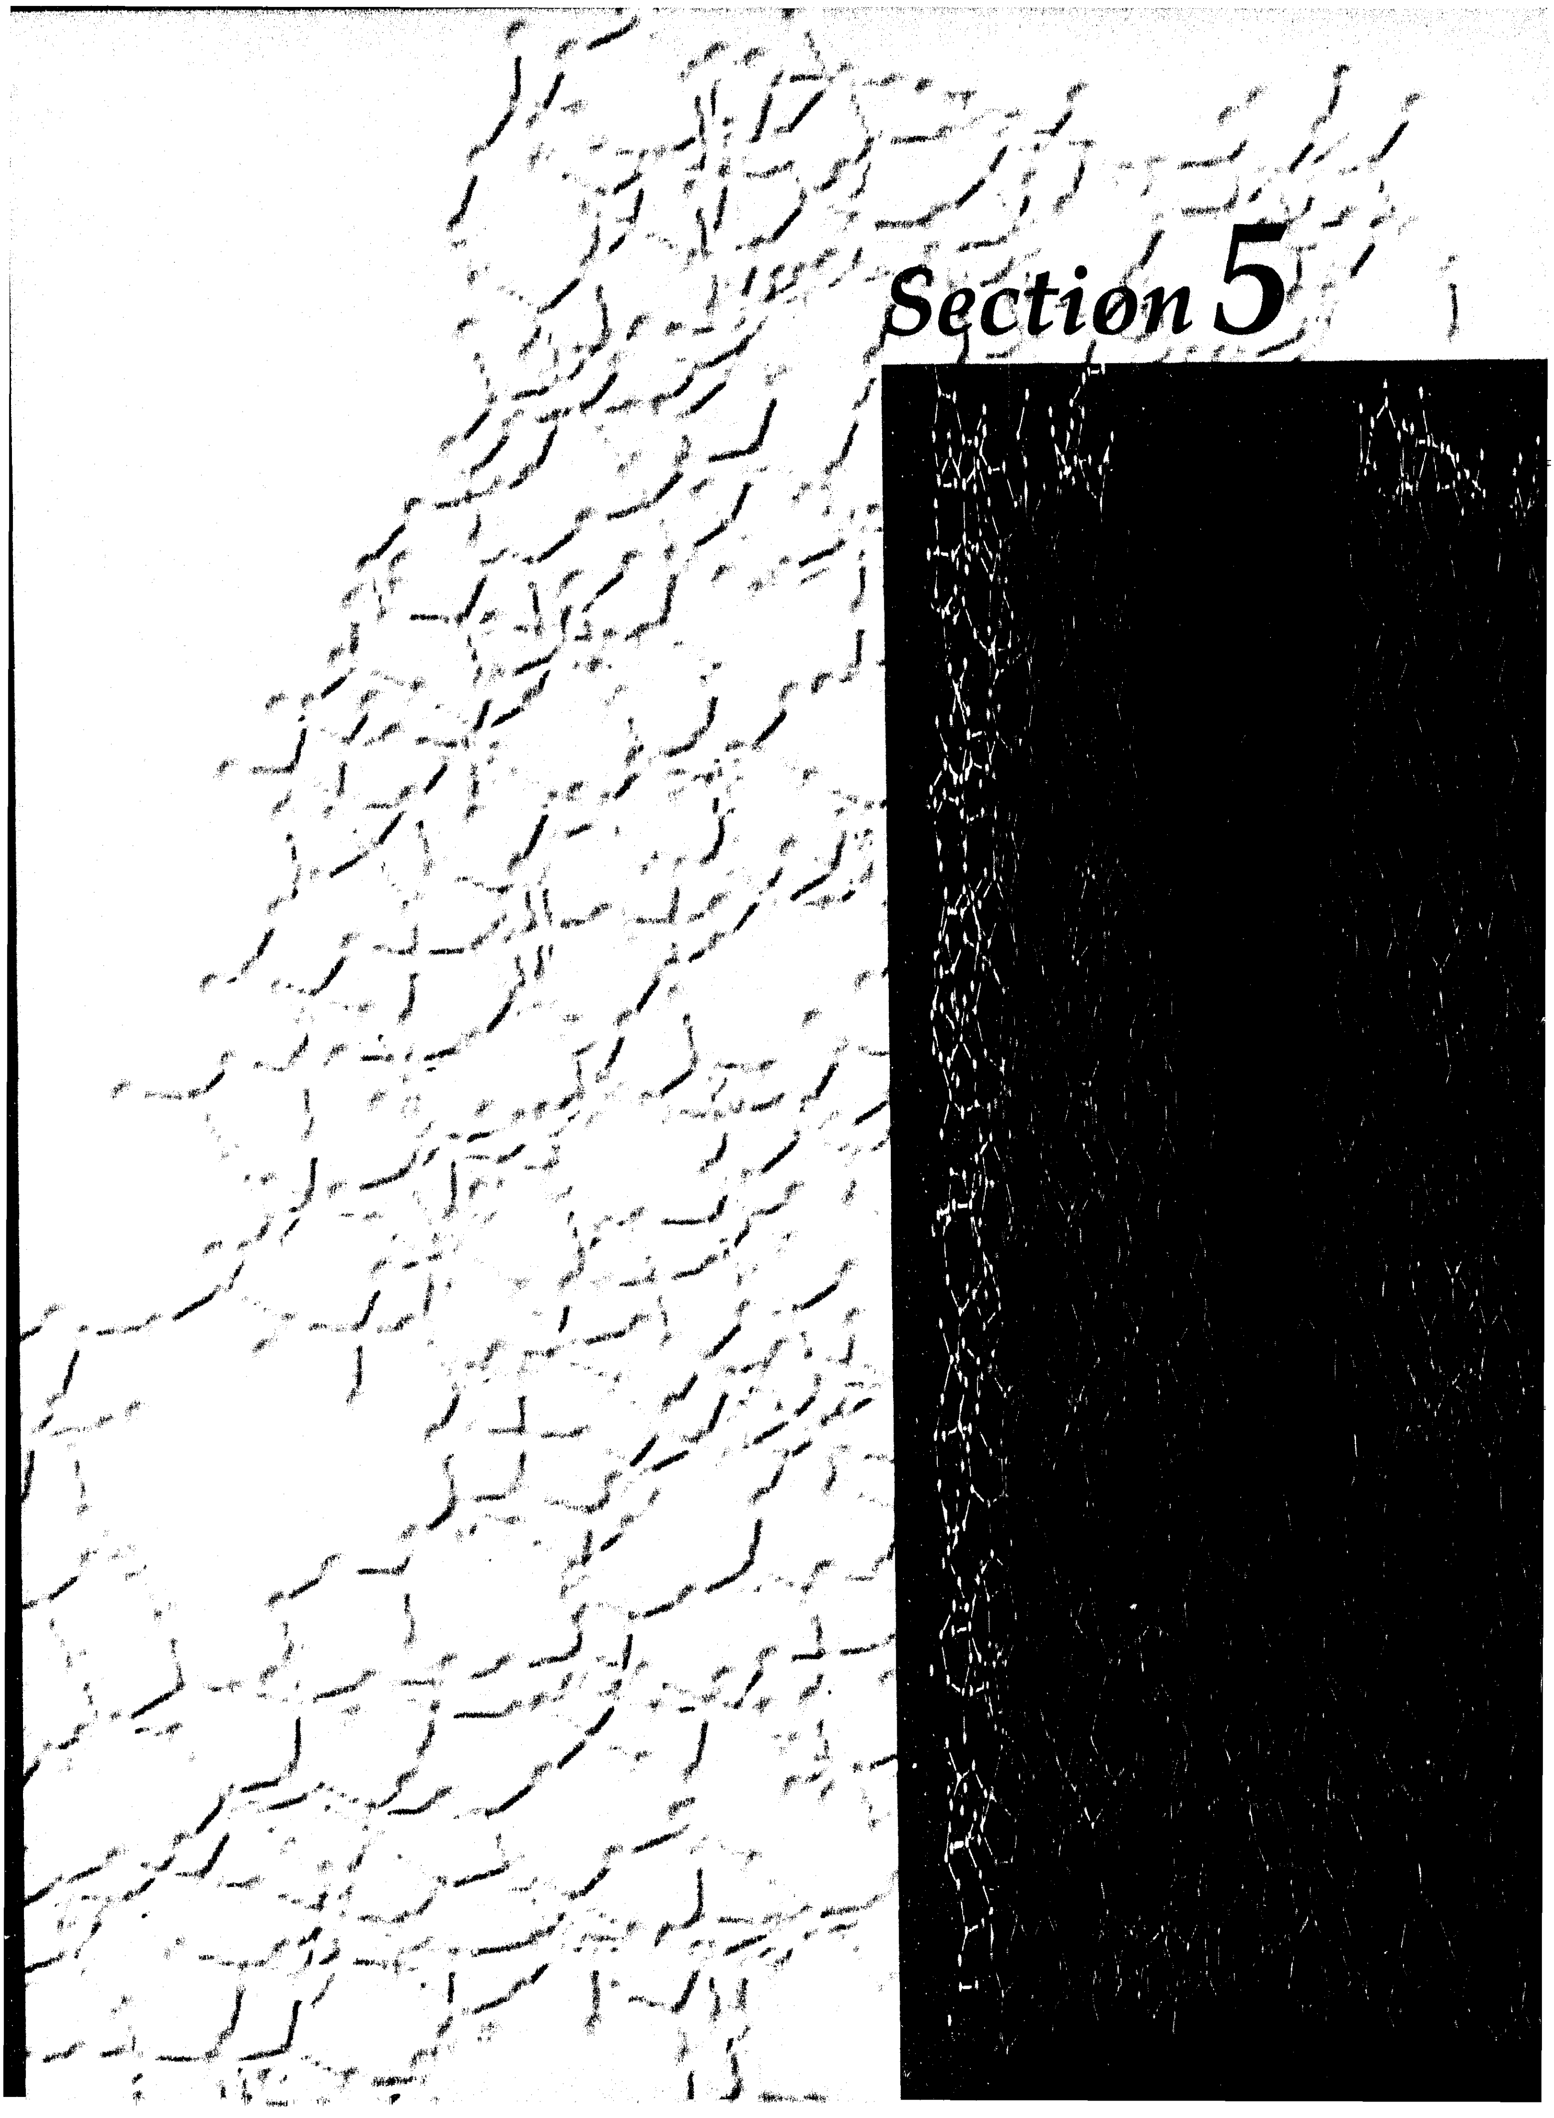

5. Fabrication Technology

Overview

Kenneth L. Blaedel, Thrust Area Leader

Simulation of Nanometer-Scale Plasticity in Silica Glass

James N. Glosli,
David B. Boercker, and
James Belak
Condensed Matter Physics Division
Physical Sciences Directorate

Irving F. Stowers
Energy, Manufacturing,
and Transportation Technologies

We applied our molecular dynamics model of fabrication processes to the examination of mechanisms for nanometer-scale plasticity and damage initiation in silica glass. Two types of computer experiments were carried out: (1) the indentation of a smooth silica surface by a sharp diamond-like tool, and (2) the propagation of a crack into a notched silica surface. The indentation simulations indicate that the strength of the silica surface is dominated by the elastic properties, and several nanometer-scale mechanisms of plasticity were identified. The critical stress for crack propagation was found to be in reasonable agreement with experiment, and a novel stick-slip phenomenon was observed.

Introduction

High-precision machining of glass and ceramic surfaces impacts a large number of fabrication technologies. The principal difficulty is the brittle nature of the material, which results in cracking and other forms of damage. However, there is a growing body of experimental evidence that under certain conditions, usually associated with small depths of cut, even these materials will behave in a ductile manner, and material removal can occur with minimal surface damage. The ability to understand and manipulate this 'ductile-brittle' transition is, therefore, critical to improving the economic viability of state-of-the-art machining capabilities.

Progress

In FY-93, we extended our molecular dynamics model of fabrication processes to examine mechanisms of nanometer-scale plasticity and damage initiation in silica glass. During these simulations, we evaluated the response of a material subjected to an external force, by following the response of every atom in the material. Two types of computer experiments were carried out: (1) the indentation of a smooth silica surface by a sharp diamond-like tool, and (2) the propagation of a crack into a notched silica surface.

One of the goals of our work is to describe the microscopic processes that give rise to macroscopic phenomena in terms of the forces between atoms. The nature of these interatomic forces in silica is beginning to be understood, and both ionic and covalent potential models have been proposed. The energetics of these models lead to tetrahedral coordination of oxygen atoms around each silicon atom, and simulations give comparable agreement with the experimental structural properties of glass. We used the covalent model of Stixrude and Bukowinski,[1] which reproduces the experimental structure and elastic properties of quartz and silica glass.

Amorphous silica glass[2] forms a random network of interconnected tetrahedra. Each tetrahedron is composed of a silicon atom at the center and four oxygen atoms at the corners. These oxygen atoms, shared by neighboring tetrahedra, are known as bridging oxygens. To create the initial slab, we reproduced a configuration of 192 atoms from a bulk simulation of amorphous silica in a 4-×-4-×-4 array. The resulting large cell, containing 12,288 atoms, was 'cooked' at a very high temperature and then quenched to a moderate temperature, with periodic boundary conditions in all three directions. At about 1000K,

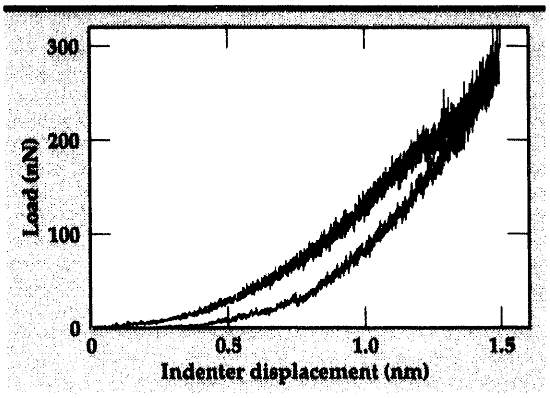

***Figure 1.* The load as a function of indentation while indenting a smooth silica surface with a sharp diamond-like tool. The upper curve is the load during indentation, and the lower curve is the load during removal. The kink at an indentation of 1.25 nm corresponds to the onset of plasticity.**

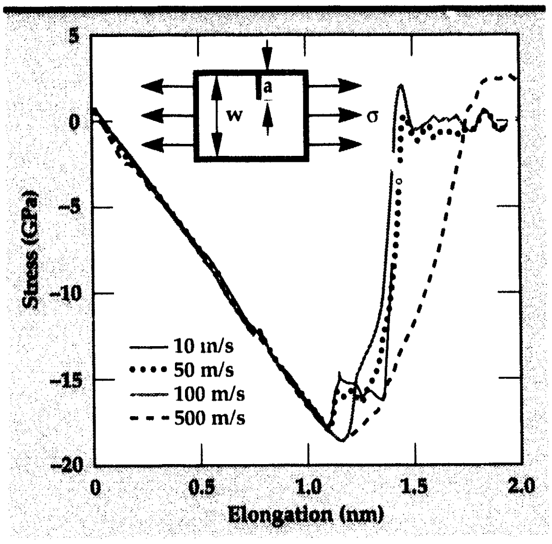

***Figure 2.* Stress as a function of elongation of the simulation cell (strain) while stretching a notched silica surface at 10, 50, 100, and 500 m/s. The geometry of the computer experiment is shown in the inset.**

the periodicity in the vertical direction was removed, and the surface was further quenched to room temperature (300K). When the periodicity in the z-direction was removed, the slab expanded somewhat in that direction, and the top of the slab was observed to oscillate up and down. Care was taken to assure that the expansion was complete and that the oscillations were damped away before any computer experiments were performed.

Indentation

Our indentation computer experiment was carried out using the same geometry as in our study of the indentation of metal surfaces. A sharp diamond tool, in the shape of the corner of a cube, was created by cleaving along the three {100} planes in the diamond structure. Several layers of atoms were removed from the tip, leaving an effective tip radius of about 1 nm. The carbon atoms in the tool, fixed at their lattice positions, interact with the silicon and oxygen surface atoms through the repulsive part of a Lennard–Jones potential. The bottom-most layer of surface atoms was also rigid, forming a hard boundary, and periodic boundaries were applied within the plane of the surface. The computer experiment consists of moving the tool into the surface at a constant velocity and observing the response of the surface atoms.

The load (total force) on the tool as a function of indentation is shown in **Fig. 1** for an indentation rate of 50 m/s at room temperature. The curve is reminiscent of recent nano-indentation experiments,[3] though we do not observe any crack formation in these simulations. For indentations beneath 1.25 nm, the surface responds elastically, and the loading curve is reversible upon removing the tool. At an indentation of 1.25 nm, we observe a kink in the loading curve. At this point in the simulation, a significant rearrangement of the tetrahedral network occurs directly beneath the tool, accompanied by an increase in the number of silicon-oxygen bonds and a rise in the local temperature. Reversing the tool from beyond this point leads to the hysteresis loop as shown. The relatively small hysteresis indicates that most of the deformation is accommodated elastically. The elastic modulus may be estimated from the slope of the unloading curve. We find Y = 50 ± 10 GPa, which compares favorably with the experimental value of 70 GPa. The contact pressure (hardness) estimated from the load and the observed area of contact for an indentation of 1.5 nm is about 15 GPa, which is in reasonable agreement with recent micro-indentation studies.

Crack Propagation

On the macroscopic scale, surface deformation and material removal in glass and ceramic materials occur primarily through the creation and propagation of cracks. Cracks propagate (elongate) to relieve elastic stress at the expense of increasing the free surface area and breaking bonds at the crack tip. To see the extent to which an atomistic model can reproduce the observed macroscopic material response and provide details not otherwise obtainable, we have simulated the propagation of a crack into a silica surface. The geometry is shown in the inset to **Fig. 2**. We use the same simulation cell as in the indentation, with a notch of length **a** in the surface and a free lower boundary condition. The computer experiment consists of pulling apart the sample at a constant velocity (constant strain rate).

In **Fig. 2**, we show the stress (force per unit area) on the left and right hand edges as a function of the elongation of the simulation cell for four different velocities. The critical stress (σ_c) for crack propagation depends on the geometry and the length of the pre-existing crack, with longer cracks yielding at lower applied stress. In fracture mechanics,[4] a material-dependent stress intensity factor (K) is defined. For our geometry, $K = Y(a/w)\sigma_c a^{1/2}$, with the geometrical factors included in the function $Y(a/w)$. From the observed critical stress, we estimate $K = 2.8 \pm 0.3$ MPa·m$^{1/2}$. This value is somewhat larger than the experimental value for silica, $K_{exp} = 1.67$ MPa·m$^{1/2}$, but considering the ambiguity in the measurement of K, the order of magnitude agreement is quite satisfying.

The stress curve in **Fig. 2** displays a plateau after the initial yield, especially at the slower strain rates. In the simulation, we observe the crack length to elongate by about 30%, stop, and then yield again at a lower stress. The fracture mechanics model predicts that the critical stress for the longer crack length should decease by a factor of $1.3^{1/2}$, or about 15%, which is approximately the value observed in the simulations. Stick-slip phenomena during crack propagation have been observed in fiber-reinforced materials. The sticking is determined by the microstructure. In our simulation, the crack stops after entering a large void, as shown in **Fig. 3**. The void allows additional stress relaxation, which halts the crack tip until sufficient stress builds up again.

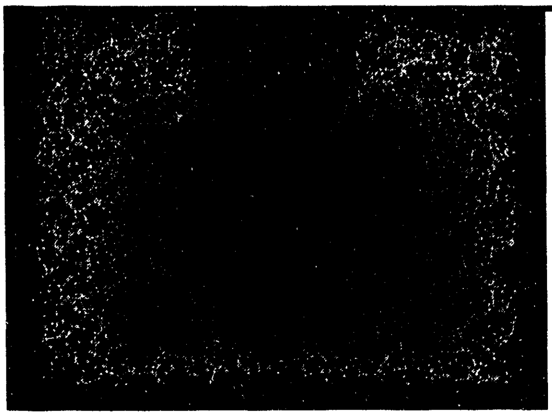

***Figure 3.* The propagation of a crack into a notched silica surface. The crack tip halts as it enters the large void to the right. This corresponds to the plateau in the stress curve shown in Fig. 2.**

Future Work

For the past few years, we have used atomic-scale simulation techniques to study fabrication processes involving metal and ceramic surfaces. We are beginning to understand the microscopic processes that give rise to the macroscopic phenomenologies that guide the practicing engineer. In the molecular dynamics model, we routinely simulate processes for which no macroscopic understanding exists. For this new information to be of practical engineering use, we will need to develop new phenomenologies. This will be the focus of our future work.

---

1. L. Stixrude and M.S.T. Bukowinski, *Phys. Rev. B* 3 **44**, 2523 (1991).
2. R.B. Sosman, *The Phases of Silica*, Rutgers University Press (New Brunswick, New Jersey), 1965.
3. R.F. Cook and G.M. Pharr, *J. Am. Ceram. Soc.* **73**, 787 (1990).
4. D. Broak, *Elementary Engineering Fracture Mechanics*, Martinus Nijhoff Publishers (Dordrecht, The Netherlands), 1986.

Mechanical Characterization of Copper–Diamond Composites

John A. Kerns
Applied Research Engineering
Mechanical Engineering

Daniel M. Makowiecki
Materials Division
Chemistry and Materials Science Department

The main goals of this project were to measure the tensile properties of a copper–diamond composite (CDC) material and to demonstrate that a grinding wheel could be manufactured using the CDC material as the abrasive. Tensile properties have been measured with limited success because of the high failure rate in manufacturing the dog bone test specimens. The basic conclusion of the tensile test is that this material has low ductility and, therefore, the failure mechanism is not brittle. The second conclusion is that a grinding wheel made using the CDC material is possible. Finally, this project has led to the development of a new concept in making grinding wheels that is the subject of current research and possible technology transfer initiatives.

Introduction

The goals of this project were to measure the stress-strain relation of a copper–diamond composite (CDC) developed at Lawrence Livermore National Laboratory (LLNL), and to make and test a grinding wheel made of this material. A test matrix was developed that satisfied the proposed tensile test requirements. A total of 21 test samples were defined from the test matrix. Ten tensile test specimens failed during the manufacturing process, which limited the amount of data available for analysis. The data that was taken shows that the CDC material fails at relatively low strain, but the failure mechanism has a ductile component. The unexpected high failure rate has led to improved manufacturing processes. (Replacement samples for this study were not made because of budgetary constraints.)

Several grinding hubs were manufactured and are available for attaching the CDC material to the grinding surfaces. We successfully demonstrated the capability of attaching the CDC material to a steel substrate for use as a grinding wheel. We are also developing a single-layer diamond grinding wheel that is more versatile than a wheel made using the composite material. A demonstration single-layer wheel was made and successfully tested.

Progress

Mechanical Tests

The first objective of this research was to determine the tensile properties of the CDC material by measuring the load vs deflection on specifically designed test specimens, and converting this to a stress-strain diagram. The test specimens, which were made into a dog bone shape as shown in **Fig. 1**, were considerably smaller than those required by standard ASTM test procedures. The sample test sizes were intentionally kept small because of the time and effort required to make

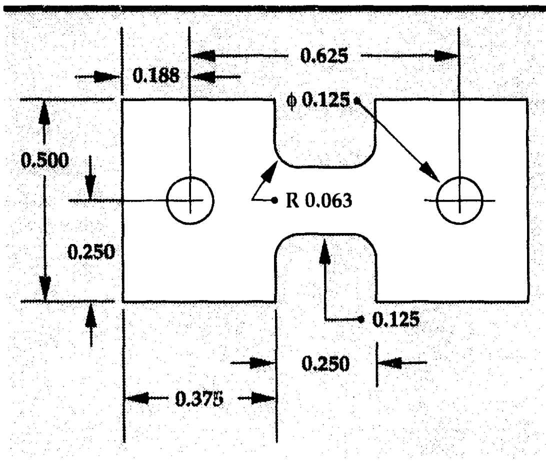

Figure 1. Dog bone test sample used in the tensile tests. Thickness of the test pieces ranged from 0.0268 in. to 0.0589 in.

CDC material in the present facilities at LLNL. A finite element analysis of the test specimen geometry verified that the test section in the dog bone was subject to a uniform stress.

To characterize the CDC material, we developed a test matrix (**Table 1**) that included composites made from commercially available copper-coated diamonds (DeBeers) and LLNL samples made using diamond powders from Beta Diamond and Engis. The Beta and Engis diamonds were coated, using processes developed at LLNL, and made into CDC test specimens. The test matrix of two LLNL CDC's and three commercial test specimens for each diamond size was deemed adequate for statistical and manufacturing purposes.

Manufacturing problems were not anticipated because thermal test specimens using the same diamond suppliers and diamond sizes had been made for each of these same diamond powders for another project. Several problems, however, did arise during the manufacture of the test specimens. The first problem was excessive shrinking and warping that occurred in the manufacture of the 6- to 12-μ Beta and Engis diamond test specimens. The shrinking and warping occurred because of the manufacturing process in which the diamond powder is first coated by physical vapor deposition with tungsten, then with copper. Additional copper is then electrolessly plated onto the coated powder. The amount of plating material was estimated using information obtained from studying the commercially coated copper–diamond powder. The next processing step is infiltrating the coated diamond that has been pressed into a rectangular form, with a copper–silver alloy. The amount of infiltrated alloy was estimated based on previous experience with composites made with the commercially coated copper–diamond powder. The shrinking occurred because of the electroless copper plating onto the small diamond. Scanning electron microscope (SEM) photos showed the plated copper on the Beta and Engis 6- to 12-μ diamond powder to be porous and the coated diamond to be agglomerated. The porosity and agglomeration are functions of both the amount of copper that is plated onto the part and the plating rate. The porosity around the diamond particles allowed the pressed compacts to shrink during the infiltration process. Warping occurs because an excessive amount of alloy was infiltrated into the pressed composite. The excess alloy accumulates on one surface of the compact, and the different expansion coefficients between the compact and the alloy cause the part to warp. The distortion caused by the plating and the excess alloy in the Beta and Engis 6- to 12-μ samples was so severe that it was impossible to make the small tensile test specimens shown in **Fig. 1**.

The commercially coated diamond did not shrink because the copper coatings were not as highly agglomerated. Since earlier experiments had determined the correct amount of alloy to be used in the infiltration process, there was very little warping in the commercial composites, and there was very little excess alloy on the finished part. The shrinking and warping were insignificant in the DeBeers commercial composites, and we were able to make tensile test specimens.

Unlike the 6- to 12-μ rectangular CDC, samples made with the Beta and Engis 24- to 36-μ diamond powder did not shrink. SEM analysis of this powder showed less agglomeration than for the 6- to 12-μ powder. Less agglomeration in the plated powder produces a higher density compact that will shrink less during infiltration. Problems associated with fabricating rectangular CDC samples with the LLNL coated powder can be eliminated by adjusting the parameters of the electroless plating process. Although there was no noticeable shrinkage, the parts still warped because of excess alloy. The warping was a significant fraction of the test sample thickness.

An electric discharge machining process cut the rectangular copper–diamond composite pieces into the flat and parallel dog bone shapes needed for the tensile test. During the machining process, the CDC parts had to be assembled onto the machine several times. This handling process could have work-hardened the test samples and made them

Table 1. Sample test matrix for mechanical property characterization of the CDC material. Listed are the number of samples that were planned, the number made into dog bone test specimens, and the number actually tested.

	6- to 12-μ			24- to 36-μ			40- to 60-μ		
	Plan	Made	Test	Plan	Made	Test	Plan	Made	Test
Beta	2	0	0	2	2	0	2	2	1
Engis	2	0	0	2	2	0	2	2	2
DeBrs	3	3	3	3	3	3	3	3	2

even more brittle than they are in the fully annealed state. Because the diamond composite material is brittle, small bending loads will easily crack the thin dog bone section. This accounted for the loss of all four of the 24- to 36-µ samples made from Beta and Engis diamond powder, and two of the 40- to 60-µ samples.

Tensile tests were performed using the remaining eleven test samples. The results of the tests are given in **Table 2**, and a typical stress-strain plot is shown in **Fig. 2**. Consistent test data was obtained for the 6- to 12-µ and the 24- to 36-µ test samples. The 40- to 60-µ data is inconsistent because the stresses and strains at failure vary within the test group, and this data is not comparable to the test data for the compacts made with the smaller size diamond. The dog bone samples were made with a cross-section of 20 or more diamonds in the shortest dimension. The 40- to 60-µ dog bones have an average of 21 to 30 diamonds in the thickness dimension, compared to 48 to 52 diamonds for the 24- to 36-µ and 97 to 126 diamonds in the 6- to 12-µ dog bones. The smaller the number of diamonds across the shortest dimension, the more anisotropic the material will become. This is the cause of the inconsistency in the data for the compacts made with the larger diamond sizes.

The tensile test data taken on the DeBeers 6- to 12-µ and 24- to 36-µ samples show that these diamond composites fail at relatively low strain, approximately 3.5% to 5.5%. The stress at failure varies from 49 to 63.5 ksi, and working stresses given by the 0.1% and 0.2% offset are 29 to 41 ksi. Although the material fails at a low strain, it exhibits low ductility as seen from a typical stress-strain diagram shown in **Fig. 2**. The LLNL-developed CDC material was not tested in this size range. We can only speculate that the LLNL material would perform as well or better than the samples made from the commercial powder obtained from DeBeers. Some justification for this statement can be made by examining the LLNL and DeBeers 40- to 60-µ data and noting that there are no large discrepancies between this data.

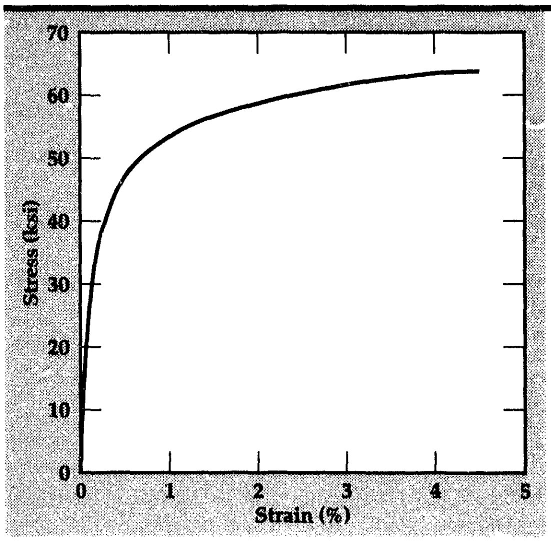

Figure 2. ***Stress-strain diagram for a 6- to 12-µ dog bone test sample made using commercially available, copper-coated, DeBeers diamond powder.***

Grinding Wheel Development

The next phase of this project was the development of a grinding wheel that incorporated the CDC material as the abrasive surface. The tasks were to demonstrate that the CDC material could be bonded to a substrate material that could be used as a grinding wheel hub; to make a hub; to bond the CDC material to the hub; and to test the wheel.

Table 2. ***Summary of the tensile test data.***

Size range (µ)	Diamond type	Stress 0.1% off (ksi)	Stress 0.2% off (ksi)	Failure strain (%)	Stress (ksi)	Young's modulus (Mpsi)
6 to 12	DeBeers	36	41	3.5	65	28
6 to 12	DeBeers	37	41	3.8	60.5	28
6 to 12	DeBeers	36	43	4.5	63.5	28
24 to 36	DeBeers	29	33	5.5	50	30
24 to 36	DeBeers	30	33	4.5	49	30
24 to 36	DeBeers	32	34	5.5	50	30
40 to 60	Beta	28	n/a	0.28	17	17
40 to 60	Engis	23.5	28.5	0.285	16.7	16.7
40 to 60	Engis	23	27	0.78	34.4	34.4
40 to 60	DeBeers	22	23.5	0.58	24	24
40 to 60	DeBeers	19	22	0.35	20.8	20.8

The first phase of this project, i.e., bonding the CDC material to a substrate, was demonstrated by attaching the CDC material to a steel substrate. Metallography analysis showed that the copper–silver alloy infiltrated the compact, and wet the steel surface to provide a good bond. In this trial, there did not appear to be gross shrinkage or warping of the composite material, and the surface was smooth. Because this was a demonstration of bonding capability, we did not attempt to measure the surface characteristics quantitatively. Quantitative measurements were planned for the grinding wheel.

The next steps were to build a grinding hub and bond the CDC material to it. One hub of titanium and several hubs of steel were made. At this time a single-layer diamond grinding wheel was determined to be more attractive than a multi-layer diamond wheel using the CDC material. Several initial attempts were made to demonstrate that a single layer of diamonds could be attached to a base material. The success of these initial attempts led to a concept of attaching a single layer of diamonds to a ceramic substrate, and then to the manufacture of a demonstration ceramic hub, single-layer, diamond grinding wheel. The ceramic surface was carefully measured during all phases of manufacture. The results of this measurement show that the surface does have only a single layer of diamonds and that the manufacturing process did not cause the surface to distort. Grinding tests show that the diamonds on this wheel do not pull out and therefore can only wear out, and that the ground surfaces were of very high quality and smooth. This concept is unique because it will enable a manufacturer to build a replenishable complex grinding surface with a single layer of diamonds. The complex surface is machined onto the substrate of the grinding wheel, and this surface remains fixed during any subsequent replenishing of the grinding wheel. The replenishment is accomplished by chemically stripping the worn diamonds from the substrate and then recoating the substrate with a new layer of diamonds. This manufacturing process is the subject of a patent application and ongoing research and technology transfer efforts.

Acknowledgements

We would like to thank the following people at LLNL for their help in this project: R. Rosen for his help in measuring the tensile properties of the CDC material; C. Alford, M. McKernan, and P. Ramsey for their efforts in coating diamond powder and maintaining the coating equipment; L. Wagner for his brazing expertise and insights into the infiltration process; R. Juntz for his help and expertise in developing the single-layer diamond grinding wheel; J. Yoshiyama for his high-quality SEM photos and insights; P. Davis for consultation and evaluation, and for providing the hubs for the single-layer diamond grinding wheel; K. Blaedel for his suggestions, consultation, and support for this project; and N. Colella for the opportunity to start this project during the Brilliant Pebbles development and for the encouragement to continue this work.

Materials Science and Engineering

The objective of the Materials Science and Engineering thrust area is to enhance our understanding of the physical and mechanical behavior and the correlations among processing, structure, and properties for structural materials that are of interest to Lawrence Livermore National Laboratory (LLNL) programs and U.S. industry. We work in (1) composite materials, including metal matrix and polymer matrix composites; (2) superplastic materials, which are crystalline solids capable of extremely high elongations; and (3) process modeling, in which we enhance our ability to model material processing, using LLNL's finite element codes.

Our activities for FY-93 are focused on all three technical areas. Details are provided in the following six reports: (1) Diffusion Bonding of Superplastic Aluminum Alloys; (2) Processing and Characterization of Laminated Metal Composites; (3) Cure Monitoring of Polymer Matrix Composites; (4) Casting Process Modeling; (5) Strength of Al and Al–Mg/Alumina Bonds Prepared Using Ultrahigh Vacuum Diffusion Bonding; and (6) Multi-Axial Superplastic Forming Apparatus.

Donald R. Lesuer
Thrust Area Leader

Section 6

6. Materials Science and Engineering

Diffusion Bonding of Superplastic Aluminum Alloys

Anne J. Sunwoo
Engineering Sciences
Mechanical Engineering

We have successfully demonstrated diffusion bonding of aluminum alloy 8090, using the production-feasible diffusion bonding (DB) condition. We also evaluated the effect of surface chemistry on the DB properties of 8090 for both interlayered and bare surfaces. Two interlayer elements, copper and zinc, affected the alloy differently. With the same interlayer thickness and bonding parameters, the tensile properties of the copper interlayer specimens of 8090 were found to be substantially better than those of the bare surface and zinc interlayer specimens. For the copper interlayer specimens of 8090, the tensile strength of 200 MPa was obtained. For the zinc interlayer specimens, the advantages of using a low-temperature transient liquid phase bonding were not observed.

Introduction

Strong interest exists in diffusion bonding (DB) and superplastic forming (SPF) technology because of the need to reduce manufacturing cost and component weight.[1] Concurrent DB–SPF processing is also considered to be an energy-saving manufacturing technology since it simplifies the production of complex components. DB is an attractive manufacturing option for applications where the preservation of the base metal microstructure and, in turn, mechanical properties, is imperative in the bond area. The process uses either solid-state or transient liquid phase bonding to produce a bond with microstructure continuity in the joint. In addition, there is no localized thermal gradient present to induce distortion or to create residual stresses in the component.

Despite the strong interest by both the aerospace[2] and automotive[3] industries to extend this technology to aluminum (Al) alloys, technical progress in this specific area has been relatively slow. So far, only certain titanium alloys have been used in DB–SPF processing.[1] The extreme difficulty associated with the DB of Al alloys is a readily formed, tenacious Al oxide layer. The solubility of oxygen is less than 0.01 at -% in aluminum, and the stability of oxide is such that it will not decompose at desired bonding temperatures, even in vacuum. The oxide layer acts as an effective barrier preventing atoms from interdiffusing across the bond interface and, consequently, yields a poor metallurgical bond. However, the lithium (Li)-Al alloy 8090 has exhibited successful solid-state diffusion bonding by modifying the Al oxide to less stable, discontinuous, Li-rich Li-Al spinels.[4] They appear as discrete, brittle particles on the bond interface. Joints that were solid-state diffusion bonded in vacuum showed a planer microstructure at the interface and shear strengths similar to those of the base material. Moreover, the 8090 alloy has lower density, and higher specific strengths and modulus than the other high-strength Al alloys, which are attractive attributes for weight-restricted aerospace applications. It is also quench-insensitive and can precipitate strengthening phases during cooling.

Progress

As the technology moves from the laboratory to production, the DB process has to be production-feasible and cost-effective. A comparative study of DB of SPF Al alloys has been initiated at the Lawrence Livermore National Laboratory (LLNL) to determine the effect of surface chemistry on the DB properties of the Al alloy, 8090 (in wt.-% 2.4Li-1.18Cu-0.57Mg-0.14Zr-Al). The integrity of the diffusion bonds was evaluated for both

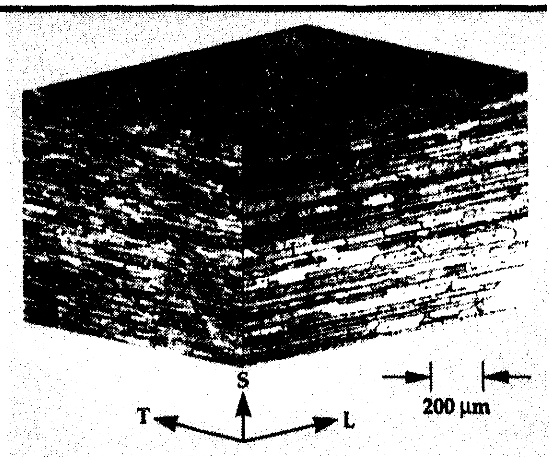

Figure 1. 8090 Al alloy shown in three orientations.

interlayered and bare surfaces. Two inter-layer elements, copper (Cu) and zinc (Zn), were compared. Although the eutectic temperature of Al-Cu is 548°C, a thin Cu layer in contact with 8090 has been shown to lower its eutectic temperature to ~521°C.[4] In 8090, Cu is one of the primary alloying elements but has a limited solubility in Al at the bonding temperature. Zn, on the other hand, forms a considerably lower eutectic temperature (380°C) with Al and is highly soluble in Al. The diffusivity of Zn in Al is much faster than that of Cu, but Zn forms a more thermodynamically stable oxide. These subtle metallurgical differences will affect the transient liquid phase (TLP) formation at the interface, which in turn will influence the bond quality.

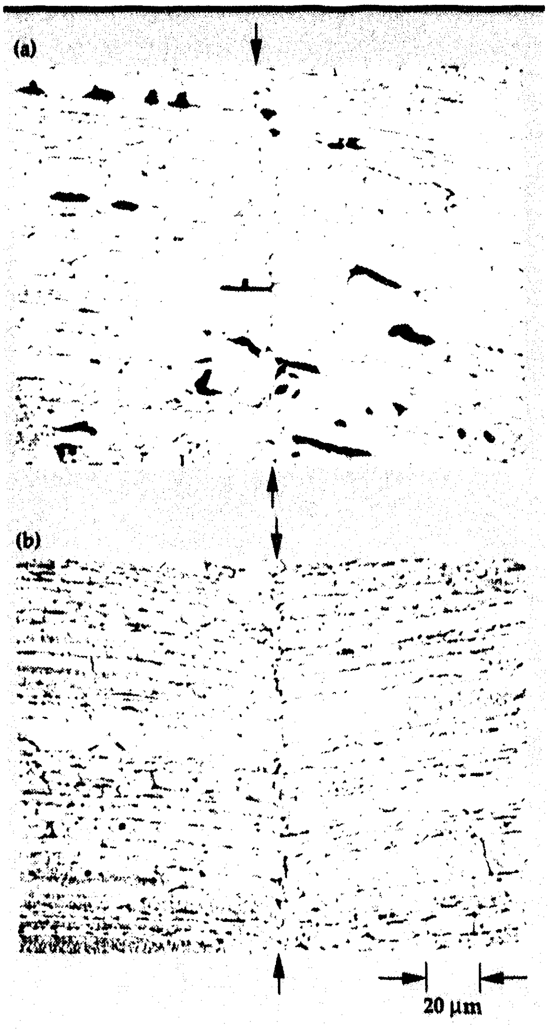

Figure 2. Cross-sections of the diffusion-bonded 8090 joints showing the effect of transient liquid phase bonding (a) with a Cu interlayer and (b) with a Zn interlayer. The interface is indicated by the arrows.

Experimental

In this work, a 9.5-mm-thick plate of 8090 was used to determine the bond tensile strength. To enhance grain boundary diffusion at the bonding interface, the surfaces to be bonded were matched in the short-transverse orientation to take advantage of elongated grains. **Figure 1** shows the microstructure of 8090 in three orientations. Electroplating was used to coat Cu and Zn interlayers on Al surfaces. The interlayer thickness was 10 µm.

The DB apparatus consisted of two ceramic platens each 50 mm in diameter and 150 mm thick, a quartz tube 63.5 mm in diameter, weights, and a tubular furnace. The specimens with matching Cu-Cu and Zn-Zn surfaces were sandwiched between ceramic platens prior to placement inside the quartz tube. The tube was sealed and backfilled with argon; a predetermined weight, which was held in place by the movable crosshead of a test machine, was loaded on the ceramic platens. The quartz tube was placed inside a water-cooled, tubular quartz furnace. Two chromel-alumel thermocouples were used to monitor the temperature, one embedded in the sample and other near the sample for furnace temperature control.

Under a static pressure of 1.2 MPa (175 psi), the samples were heated to 550°C at 18°C/min and held for 15 min to ensure melting of the Cu interlayer, lowered to 500°C for 60 min, and then slowly furnace-cooled. The bonded interfaces and fracture surfaces were examined using optical metallography and scanning electron microscopy fractography, respectively. The bond quality was determined by tensile tests at room temperature. The specimens were pulled at a constant crosshead speed of 1.27 mm/min with the bond located at the center of the 25.4-mm gage length.

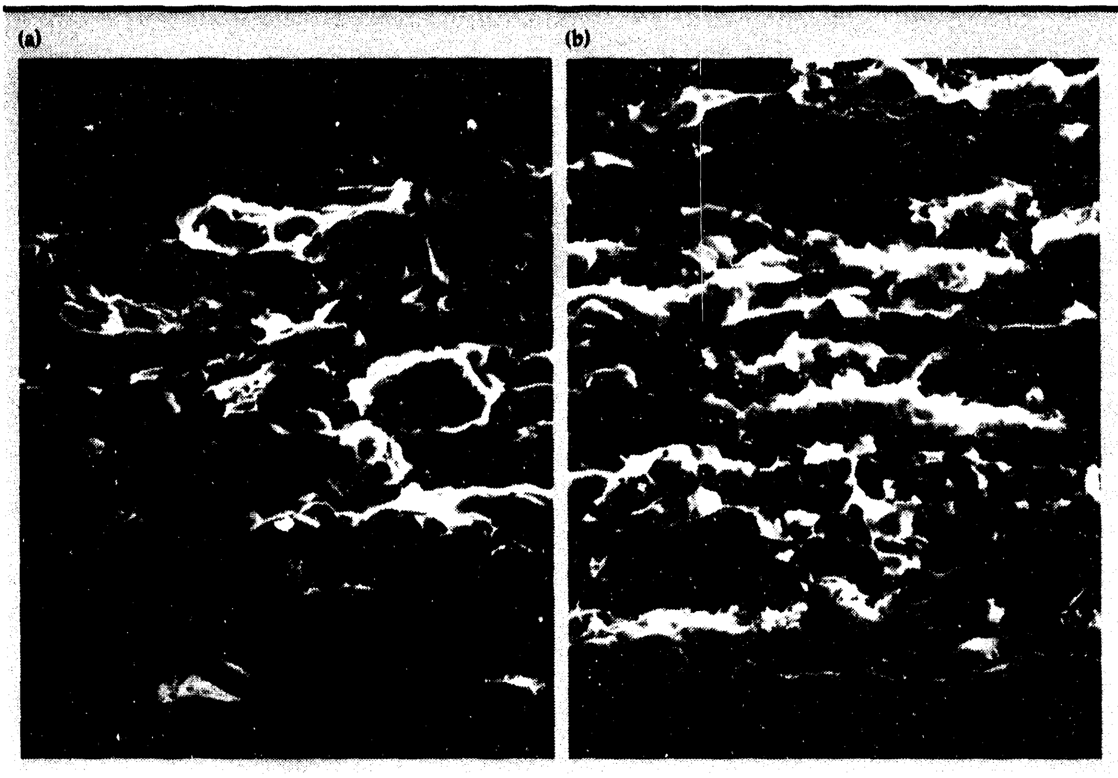

Figure 3. Scanning electron microscopy fracture surfaces of the diffusion-bonded 8090 joints (a) with a Cu interlayer and (b) with a Zn interlayer.

Results and Discussion

For the bare surface DB of 8090, the benefit of Li-Al spinel formation at the bonding surfaces was not found. In a non-vacuum environment, the specimen simply did not bond. Although there was sufficient pressure available to produce interfacial contact of the bonding surfaces, in an argon atmosphere with a partial pressure of oxygen, even the freshly exposed metal surfaces would reoxide and prevent the metal/metal contact required to form the bond.[5]

The primary advantage of using an interlayer is that at the bonding temperature, melting at the interlayer interface forms a TLP, which leaves a clean surface to be bonded and also helps to flatten the surfaces. However, the desorption, transportation, and reprecipitation phenomena that occur at the melt region[6] will vary depending on the reaction between substrate and interlayer. The difference in the Cu and Zn TLP formation and their effects on DB is shown in the cross-sections of polished DB joints of 8090 (**Fig. 2**). For the Cu interlayer specimens, the bond interface was difficult to detect (**Fig. 2a**). A discrete dispersion of Cu-rich phases near the bond area provided the only observable indication of the bond location in the joint. Three ternary intermetallic compounds that are in equilibrium with the Al-rich solid solution are T_B ($Al_{7.5}Cu_4Li$), T_1 (Al_2CuLi), and T_2 (Al_2CuLi_3).[7] These particles were only present within 125 µm from the interface. A similar observation was made by Bushby and Scott[8] where they found that Cu had diffused up to 150 µm into the Al when held at 550°C for 100 min. Also note that these Cu-rich phases did not form on all the grain boundaries. Diffusion of Cu or TLP into the base metal through the grain boundary region is strongly influenced by the grain misorientation, whether at random (high-angle) or ordered (low-angle) boundaries.[9] In the case of the random high-angle boundaries, higher grain boundary energy promotes deeper liquid penetration and the formation of a smaller dihedral angle with the liquid phase. On the other hand, a much smaller effect of liquid penetration and change in the solid–liquid interface profile is seen with low-angle boundaries, based on their lower boundary energy.

For the Zn interlayer specimens, the reaction at the interface was significantly different. The cross-section of the DB specimen showed more distinct outline of the bond interface (**Fig. 2b**). Unlike Cu, Zn does not form a common ternary intermetallic compound with any of the alloying elements in 8090, but in contact with Al undergoes a phase change at ~385℃,[10] which is close to the eutectic temperature of the Zn-Al system. At the bonding temperatures of 500 and 550℃, the Zn-rich TLP diffuses rapidly into the matrix, but the areas with eutectic composition will be the last to solidify. The noticeable difference in composition was only seen at the bond interface of an unetched specimen. Both Cu and Zn interlayer specimens contained nonplanar interface profiles, and the presence of extensive plastic deformation was evident in the grains near the interface where they were severely misaligned.

To effectively compare the bond quality, the 8090 base metal was reprocessed using the DB parameters. The tensile test results are listed in **Table 1**. With the same interlayer thickness and bonding parameters, the tensile properties of the Cu interlayer specimens of both alloys were found to be substantially better than those of the Zn interlayer specimens. The integrity of the bond was determined using a joint efficiency, which is a ratio of the joint strength and base metal strength. For the Cu interlayer 8090 specimens, the joint efficiency of 50% (200 MPa) was obtained. In comparison, only a 25 % joint efficiency was obtained for the Zn interlayer 8090 specimens.

Figure 3 shows the fracture surfaces of the Cu and Zn interlayer specimens. The typical fracture behavior of 8090 alloy was seen in the fracture surface of a Cu interlayer specimen. The surface revealed ductile dimples with delamination along the grain boundaries (**Fig. 3a**). Because the delamination occurred along the boundaries that contained Cu-rich phases, the grains shown in the fracture surface appeared much larger than those in the Zn interlayer fracture surface (**Fig. 3b**). Since Zn is highly susceptible to oxidation, for Zn interlayer specimens, the areas that have not established an interfacial contact are exposed to an argon atmosphere with a partial pressure of oxygen. Thus, the crack propagated along the bond interface, resulting in predominantly intergranular fracture with some ductile dimples and contributing to a poor bond strength. The fracture surface showed the pancake-like microstructure of the base metal.

The shear strength of the bond was qualitatively estimated using a theoretical relation,[11] $\tau = 0.6\sigma_L$, where σ_L is longitudinal tensile strength (see **Table 1**). It is difficult to make a direct comparison between the present results and others due to the fact that many reported data are in the post-DB T6 condition. With the post-DB heat treatment, the Cu-rich phases seen in the grain boundaries of the Cu interlayer specimens of 8090 would dissolve and reprecipitate in the matrix, thereby improving the joint strength. In addition, the joints are produced in vacuum with much longer time than the ones used in this work. Despite these differences, the as-bonded shear strengths of 120 MPa for the Cu interlayer specimens of 8090 are comparable to the values reported by Dunford and Patridge.[4] A commercially pure Al with the Cu interlayer exhibited a maximum shear strength value of ~50 MPa after 60 min at 550℃ under 3 MPa pressure in vacuum; a low strength was attributed to Al_7Cu_2Fe particles found on the bond interface.[7] The estimated shear strengths of the Zn electroplated specimens of 8090 (58 MPa) are much lower than those of the Zn physical-vapor-deposited specimens of 8090 by Ridley and Livesey,[10] where they reported an average shear strength value of 144 MPa for the joints produced using pressure of 3 MPa at 515℃ for 2.5 h. However, the actual shear strength requirements of the bonds for aircraft structures are generally on the order of 10–20 MPa,[12] and the present results satisfied the requirements. The greater importance for the bond property is reported to be its resistance to peeling, particularly during the SPF process.

Table 1. Tensile and estimated shear strengths of diffusion-bonded Al alloy, 8090.

	Fracture strength MPa	Estimated shear strength MPa
8090	395±20	223
Cu Interlayer 8090	200±7	120
Zn Interlayer 8090	96.5±10	58

Future Work

A further study is needed to better understand the effect of Zn on the diffusion bonding of Al alloys.

1. E.D. Weisert and G.W. Stacher, "Concurrent Superplastic Forming and Diffusion Bonding of Titanium," *Superplastic Forming of Structural Alloys*, N.E. Paton and C.H. Hamilton (Eds.), TMS-AIME (Warrendale, Pennsylvania), 1982.

2. P.J. Winkler, T. Heinrich, R. Keyte, and R.A. Ricks, "Bonding and Superplastic Forming of Al-Li Alloy AA8090 for Commercial Applications," *Aluminum Lithium*, M. Peters and P.J. Winkler (Eds.), DGM Informationsgesellschaft, Verlag, 1992.
3. J.E. Allison and G.S. Cole, *JOM*, **19** (1993).
4. D.V. Dunford and P.G. Partridge, *Mater. Sci. Tech.* **8**, 385 (1992).
5 E.R. Maddrell and E.R. Wallach, *Recent Trends in Welding Science and Technology*, ASM International (Metals Park, Ohio), 540, 1990.
6. R.M. German, *Liquid Phase Sintering*, Plenum Press (New York), 1985.
7. H.K. Hardy and J.M. Silcock, *J. Inst. Metals* **84**, 423 (1955–56).
8. R.S. Bushby and V.D. Scott, *Mater. Sci. Tech.* **9**, 417 (1993).
9. H. Kokawa, C.H. Lee, and T.H. North, *Metall. Trans. A* **22A**, 1627 (1991).
10. N. Ridley and D.W. Livesey, "Diffusion Bonding of Superplastic Al-Li (8090) Using a Zinc Interlayer," *Aluminum Lithium*, M. Peters and P.J. Winkler (Eds.), DGM Informationsgesellschaft, Verlag, 1992.
11. ASM, *Source Book on Selection and Fabrication of Al Alloys*, ASM International (Metals Park, Ohio), 1, 1978.
12. J. Pilling and N. Ridley, *Mater. Sci. Tech.* **3**, 353 (1987).

Processing and Characterization of Laminated Metal Composites

Chol K. Syn,
Susan L. Stoner, and
Donald R. Lesuer
Engineering Sciences
Mechanical Engineering

O.D. Sherby
Department of Materials Science
and Engineering
Stanford University
Palo Alto, California

Multi-layer laminated metal composites of Al 6090-25 vol% SiC_P and Al 5182 were prepared by hot pressing alternating layers of the component materials in an argon gas atmosphere. Tensile and fracture toughness were measured as a function of volume percent of the component materials under T6 heat-treated and untreated conditions. T6 heat treatment on the aluminum (Al) laminates induced a substantial increase in the yield and tensile strength but a decrease in tensile ductility and interlayer bond strength. Fracture toughness measured both in the crack arrester and crack divider orientations showed a definite increasing trend with the volume percent of the Al 6090-25 vol% SiC_P component.

Introduction

The idea of laminating different metals and alloys to form a composite material that exploits the good properties of the constituent materials has been known from antiquity: *The Illiad* of Homer, e.g., describes Achilles' shield, made of two outer layers each of bronze and tin, and one middle layer of gold.[1] The idea has been used in many industrial applications.[2] However, most of the current, industrial, metal-based laminates contain only two or three layers, and are used to save material cost and improve mostly wear or corrosion resistance. Recent studies[3–5] show that *multilayer* laminated metal composites (LMC's) can have superior damage-critical properties, such as fracture toughness or fatigue resistance, superior to that of the component materials. Damage crack propagation in a laminate of dissimilar materials is inherently difficult since an interface can act as a barrier to the crack propagation, especially when such an interface delaminates at the crack tip and blunts the crack.[6] A good example is the LMC formed by press-bonding alternate layers of ultrahigh carbon steel (UHCS) and mild steel.[4] In this LMC, the dynamic fracture toughness was far higher than that of either component material. LMC's can also have damping capacity superior to that of the component materials,[7] which can be very useful in structures requiring high acoustic damping.

Progress

We initiated the present research in FY-91 to investigate the influence of processing and structural variables on the mechanical properties of multilayer LMC's made of two constituents, one ductile but tough and the other brittle but strong. We chose in FY-91 to study two LMC systems, UHCS/brass and Aluminum (Al)/Al-SiC_P, and continued our study in FY-92. Al-SiC_P represents an aluminum-based SiC particulate, reinforced metal-matrix-composite (MMC) material. These two systems were chosen to show that the toughness at ambient temperature of hard and brittle UHCS and Al-SiC_P MMC can be enhanced substantially by lamination with ductile but tough counterparts.[3] In FY-92, we studied the influence of the surface preparation of the component materials and the layer thickness of the laminates on their interfacial microstructure and mechanical properties.[8–10] The main thrust for FY-93 was to study the fracture toughness of the Al/Al-SiC_P laminates as a function of the volume fraction of Al-SiC_P MMC component.[11]

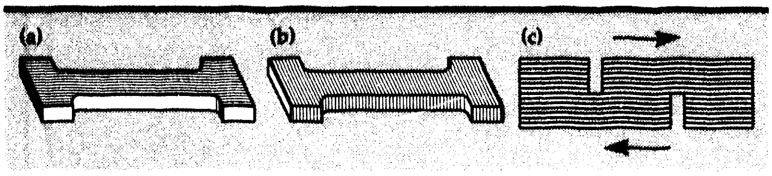

Figure 1. Geometries of tensile samples in the (a) parallel and (b) normal orientations and (c) the lap shear test sample.

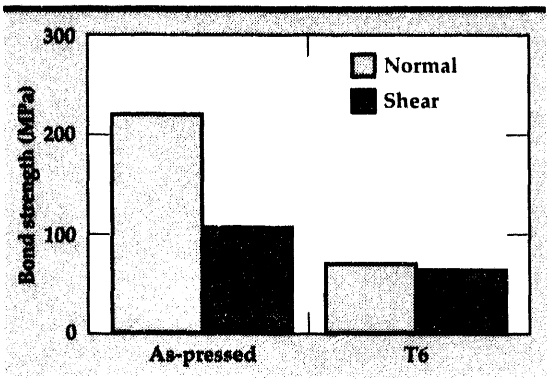

Figure 2. 'Normal' and 'shear' interface bond strengths before and after the T6 heat treatment.

Experimental Procedure

Processing. Al 5182 (Al-4.5%Mg-0.35%Mn) and Al 6090-SiCp MMC (Al-1.0%Mg-0.6%Si-0.28%Cu-0.2%Cr matrix with 25 vol% SiC particulate reinforcement) were obtained in the form of plates and sheets of different thicknesses from commercial sources. The materials were sliced into 50 mm-×-50 mm squares. Al 5182 and Al 6061-SiCp squares were descaled using acid solution. Laminates of different volume percents were prepared by hot pressing a stack of alternate layers of the component materials of appropriate thicknesses. Each stack of ~ 50-mm height was pressed at 450°C in an argon gas atmosphere to a third to fourth of its initial height. Such a large reduction ensured good bonding at interfaces. Average layer thickness of a laminate of 50-50 vol% was about 750 μm after initial pressing. Some laminates were given the T6 heat treatment, which was to soak at 530°C for about 75 min followed by water quenching to room temperature, and temper at 160°C for 16 hours followed by air cooling to room temperature.

Testing. Tensile tests were performed with flat specimens cut with the tensile axis parallel to the layers, as shown in **Fig. 1a**. Interlayer interface bond-strength tests were conducted with two types of samples: (1) with flat tensile specimens cut with the tensile axis perpendicular to the layers for bond strength in the direction normal to the layers (**Fig. 1b**), and (2) with 'lap shear' test specimens for the shear bond strength (see **Fig. 1c**). Fracture toughness was measured with chevron-notch three-point bend bar specimens in which the notch was cut either in the *crack arrester* or *crack divider* orientation. In the crack arrester orientation, the crack front propagates in the thickness direction, cutting the layers sequentially. In the crack divider orientation, the crack front propagates through the laminate cutting all the layers simultaneously.

Experimental Results

Tensile Properties. Tensile properties for the laminates containing 50-50 vol% of the component materials are summarized in **Table 1** in the as-pressed, and T6 heat-treated conditions. **Table 1** also includes the tensile properties of the component materials for comparison.

Bond Strengths. Bond strengths were measured for the 50-50 vol% laminates in the as-pressed and T6 treated conditions, and the results are plotted in **Fig. 2. Figure 2** shows substantial reductions (~ 70% for the 'normal' and ~ 40% for the 'shear' directions) in the bond strengths after the T6 treatment, indicating that the T6 treatment can be an effective way to change the interfacial bond strengths.

Fracture Toughness. Toughness values measured for the 50-50 vol% laminates with and without the T6 treatment are shown in **Fig. 3** for the crack divider and crack arrester orientations.

Table 1. Tensile properties of Al 6090/SiC/25p-Al 5182 laminates containing 50-50 volume percent of the component materials.

Processing	Average Layer Thickness, μm	Yield Strength MPa (ksi)	Ultimate Tensile Strength, MPa(ksi)	Elongation %
As-Pressed	750	162(23.5)	266(38.5)	16.9
Pressed, T6	750	232(33.6)	333(48.2)	7.2
Al 5182	—	130(18.9)	275(39.9)	25
Al 6090/SiC/25p	—	265(38.4)	310(45.0)	4
Al 6090/SiC/25p, T6	—	422(61.2)	510(72.7)	3.2

Fracture toughness of the component materials, Al 6090/SiC/25p MMC with T6 treatment, and Al 5182, was not measured in the present study, mainly due to the lack of sufficiently thick materials. For the Al 6090/SiC/25p MMC component, two toughness values have been reported,[12,13] and the higher value of the two is included in **Fig. 3**. Since the toughness value for the Al 5182 component is not available in literature, the value for Al 5083 H321,[14] whose composition is rather close to that of the Al 5182, is used in **Fig. 3**.

The toughness of the laminates does not change noticeably from the average value of ~ 30 MPa·m$^{1/2}$ regardless of the T6 treatment for both crack divider and arrester orientations. These results contrast with the substantial change in the tensile properties and interfacial bond strengths before and after the same T6 treatment. It seems certain that the fracture toughness of laminated materials is determined by several simultaneously or sequentially operating mechanisms. The relative dominance of these mechanisms can be changed by the relative strengths of the component materials and the interface bond strength. An example of the interdependence of interface and component material properties is available from the data in **Figs. 2** and **3** and **Table 1**. Heat treatment decreases the interfacial bond strength (which should increase toughness through delamination effects, as will be shown later) and increases the flow stress of the MMC layers (which should decrease toughness). The net result of these two competing effects is shown in **Fig. 3**. The toughness has changed very little with heat treatment. Nonetheless, the results shown in **Fig. 3** indicate that the laminated composites are very forgiving about the influence of heat treatment and other processing variables as long as the layers are reasonably well bonded to ensure the structural integrity, as shown elsewhere.[3,9,10]

The fracture toughness variation for both crack divider and arrester orientations as a function of the volume percent of the MMC component is shown in **Fig. 4** for the T6 treated laminates. It is significant to note that the toughness shows a slight but definitely increasing trend with the vol% of the MMC component for both orientations. The horizontal axis in **Fig. 4** also shows the global vol% of the SiC reinforcement. Results in **Fig. 4** clearly show that with lamination, the fracture toughness of 6090/SiC/25p MMC can increase with increasing global vol% of reinforcement. An increased vol% of the MMC or global vol% of the SiCp reinforcement will also lead to increased yield and tensile strengths of the laminate. Tensile properties of a composite material are known to follow the rule-of-average of those of the component materials.[15] These results indicate that with lamination technology, both fracture and tensile strengths can be enhanced by increasing the vol% of MMC

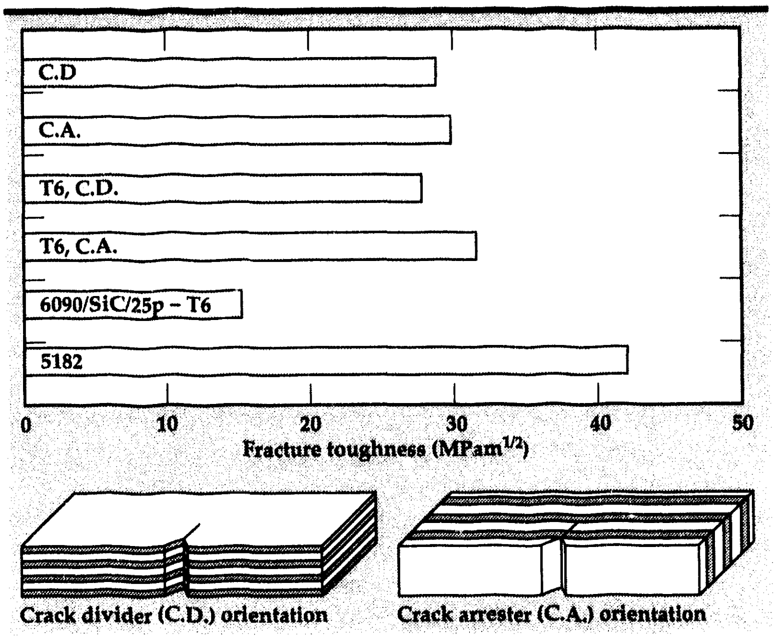

Figure 3. Fracture toughness of Al 6090/SiC/25p-Al 5182 laminates (containing 50-50 vol% of the component materials) with and without T6 treatment for the crack divider (C.D.) and crack arrester (C.A.) orientations.

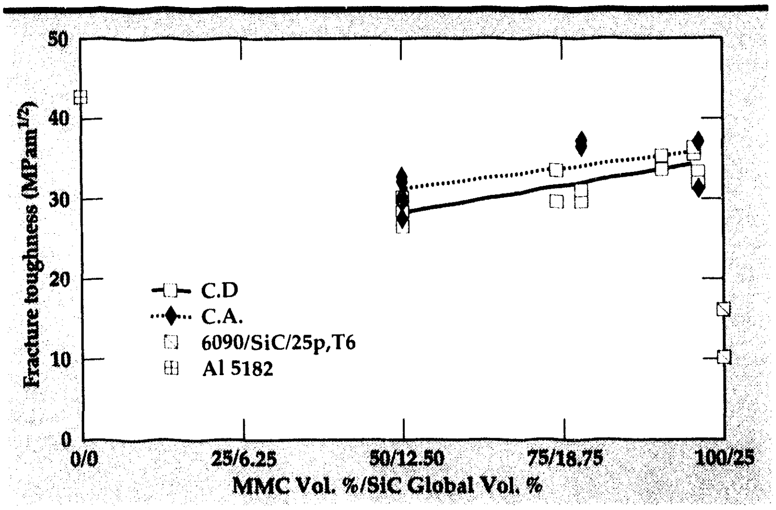

Figure 4. Fracture toughness of Al 6090/SiC/25p-Al 5182 laminates vs the vol% of the MMC component and the global vol% of the SiC$_p$ reinforcement.

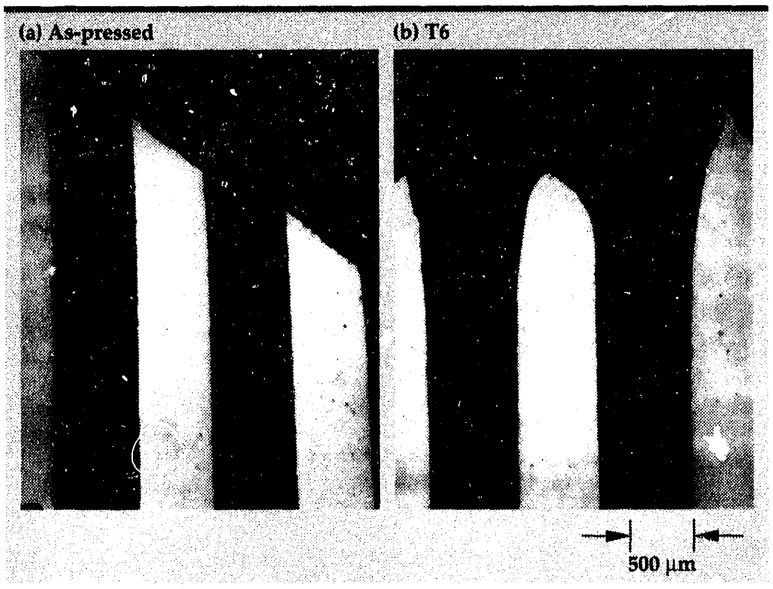

***Figure 5.** Cross-sectional views of tensile tested laminates with and without T6 treatment.*

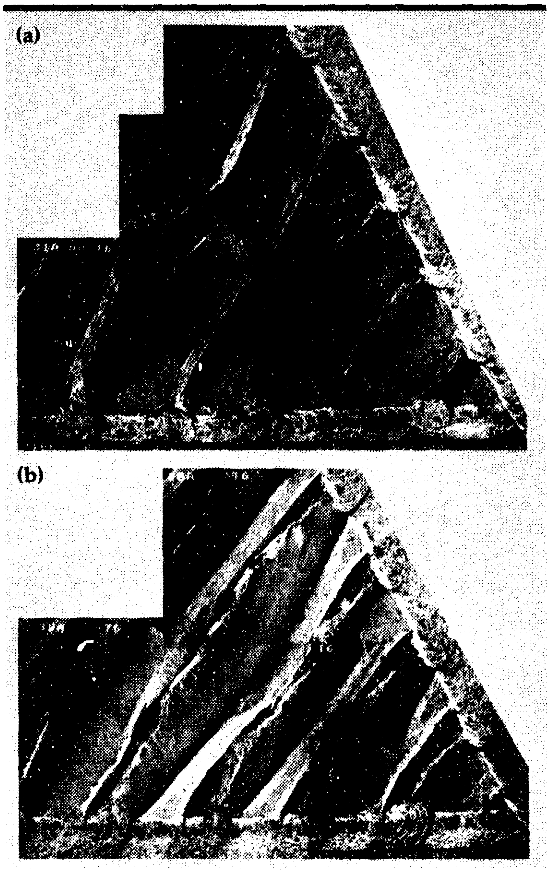

***Figure 6.** Fracture surfaces of the 50-50 vol% laminates in the (a) as-pressed and (b) T6-treated conditions for the crack arrester orientation. Direction for crack propagation is from the lower right-hand to the upper left-hand corners in both (a) and (b).*

or reinforcing agent. The global vol% of the reinforcing agent in a laminate may be increased by using an MMC that contains a higher vol% reinforcement.

Fracture Behavior. Fracture surfaces of tensile and fracture toughness tested samples and cross-sectional microstructures of tensile tested samples were examined to determine fracture morphology and the interfacial delamination behavior of the laminates. **Figure 5** represents the cross-sectional views of the tensile tested samples in the as-pressed and T6-treated conditions. It is clearly seen in **Fig. 5** that the T6 treatment leads to large-scale interfacial delamination as a result of the reduction in interfacial bond strength (**Fig. 2**). Despite the large effect of the T6 treatment on the delamination behavior in the tensile tested samples, the influence of such treatment on the fracture surface seems to be rather subtle. **Figure 6** shows the fracture surfaces of the laminates with 50-50 vol% in the as-pressed (**Fig. 6a**) and T6-treated (**Fig. 6b**) conditions when tested in the crack arrester orientation. In **Fig. 6**, the fracture crack was initiated in the lower right-hand corner and propagated to the upper left-hand corner. In the as-pressed condition, every other interface of the layers was delaminated while every interface was separated in the T6-treated condition. **Figure 7** shows the fracture surface morphology of the laminates in the T6-treated condition with 50, 75, and 97 vol% of the MMC component. Results are given for the crack divider orientation. Every interlayer interface is seen to have delaminated. Every Al 5182 and 6090 MMC layer shows a shear fracture mode. For the 97%-3% laminate, the rather thick MMC layers exhibited tensile fracture in the early stage of crack growth but immediately transitioned to the shear fracture mode.

Future Work

We are continuing to characterize the Al/Al-SiC laminates regarding their (1) damping capacity in the audible frequency range; (2) fatigue behavior; (3) response to ballistic impact; and (4) deformation behavior at elevated temperature. We are planning to make LMC's of other light materials such as Ti and Mg alloys (1) containing a high damping capacity material as a component, and (2) containing an intermetallic or superalloy as a component. These new laminates will be tested for their strength, ductility, toughness, and other characteristics, as for the Al laminates reported here.

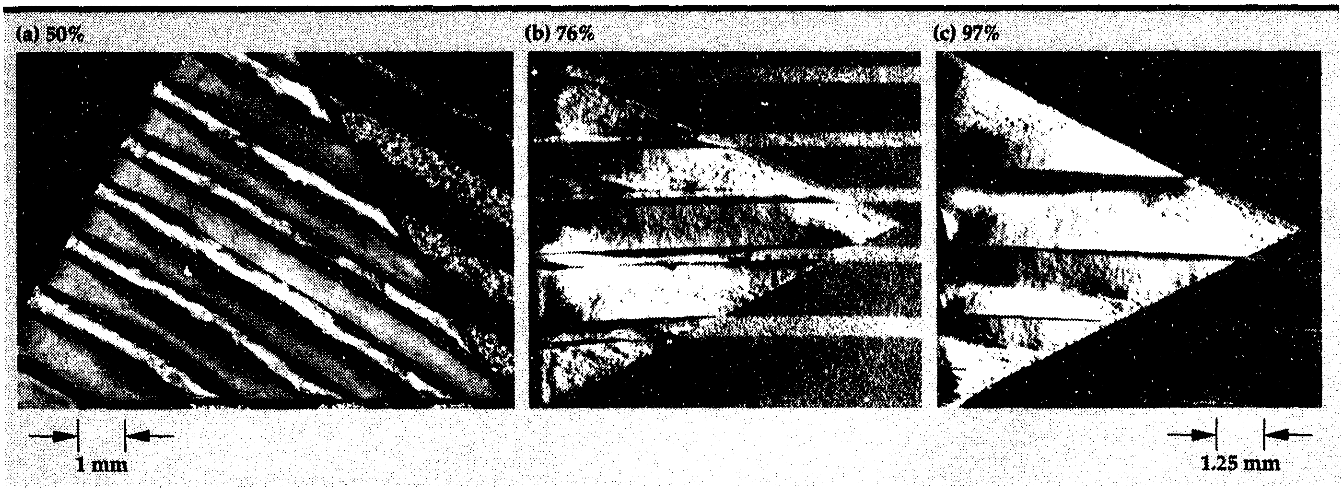

Figure 7. Fracture surface of the T6-treated laminates with (a) 50, (b) 76, and (c) 97 vol% of the MMC component tested in the crack divider orientation.

Acknowledgements

The authors appreciate assistance provided by C. Steffani for descaling Al alloys; R. Otto for pressing laminates; W. Andrade and G. Stebbins for preparing test samples; A. Shields for conducting mechanical property tests; and J. Ferreira for metallography.

---

1. *The Iliad* of Homer, translated by R. Lattimore, University of Chicago Press (Chicago, Illinois), 411 (lines 270–272), 1951.
2. E.S. Wright and A.P. Levitt, "Laminated Metal Composites," *Metallic Matrix Composites*, K.G. Kreider (Ed.), Academic Press (New York, New York), 37, 1974.
3. C.K. Syn, D.R. Lesuer, K.L. Cadwell, O.D. Sherby, and K. Brown, "Laminated Metal Composites of Ultrahigh Carbon Steel/Brass and Al/Al-SiC: Processing and Properties," *Proc. Conf. Developments in Ceramic and Metal-Matrix Composites*, K. Upadhya (Ed.), TMS, 311, 1991.
4. D.W. Kum, T. Oyama, J. Wadsworth, and O.D. Sherby, *J. Mech. Phys. Solids* **31**, 173 (1983).
5. O.D. Sherby, S. Lee, R. Koch, T. Sumi, and J. Wolfenstine, *Materials and Manufacturing Processes* **5**, 363 (1990).
6. J. Cook and J.E. Gordon, "A Mechanism for the Control of Crack Propagation in All-Brittle Systems," *Proc. Royal Soc. London* **282**, 508, 1964.
7. B.P. Bonner, D.R. Lesuer, C.K. Syn, and O.D. Sherby, "Damping Measurements for Ultrahigh Carbon Steel/Brass Laminates," *Proc. Symp. Damping of Multiphase Inorganic Materials, ASM Materials Week, November 1–5, 1992, Chicago, Illinois*, R. Bhagat (Ed.), ASM International, 179, 1992.
8. C.K. Syn, D.R. Lesuer, J. Wolfenstine, and O.D. Sherby, *Metall. Trans. A* **24A**, 1847 (1993).
9. C.K. Syn, D.R. Lesuer, and O.D. Sherby, "Processing and Mechanical Properties of Laminated Metal Composites of Al/Al-25 vol.% SiC and Ultrahigh Carbon Steel/Brass," *Adv. Synthesis of Engineered Structural Materials*, J.J. Moore, E.J. Lavernia, and F.H. Froes (Eds.), ASM Int., 149 (1992).
10. C.K. Syn, D.R. Lesuer, and O.D. Sherby, "Processing and Mechanical Properties of Laminated Metal Composites of Al 6090-25 vol.% SiCp and Al 5182," *Proc. Int. Symp. Light Materials for Transportation Systems, June 20–23, 1993, Kyongju, Korea*, N.J. Kim (Ed.), Center for Advanced Aerospace Materials, Pohang University of Science and Technology, Pohang, Korea, 763, 1993.
11. C.K. Syn, S. Stoner, D.R. Lesuer, and O.D. Sherby, "Influence of Volume Fraction of Component Materials and Interlayer Bond Strength on Fracture Toughness of Multi-Layer Al 6090-25 vol.% SiCp and Al5182 Laminates," to be published in *Proc. TMS Conf. High Performance Ceramic and Metal Matrix Composites*, (San Francisco, California, February 27–March 3, 1994).
12. Y. Flom and R.J. Arsenault, *J. Metals*, **38**(7), 31 (1986).
13. M.G. McKimpson, E.L. Pohlenz, and S.R. Thompson, *J. Metals* **45**(1), 26 (1993).
14. Wrought Aluminum Selector Chart, in *Structural Alloys Handbook 1992*, MCIC, J.H. Holt, H. Mindlin, and C.Y. Ho (Eds.), CINDAS/Purdue University (West Lafayette, Indiana).
15. S. Lee, J. Wolfenstine, and O.D. Sherby, *J. Compos. Mater.* **25**, 842 (1991).

Cure Monitoring of Polymer Matrix Composites

Keith C. Hong
Chemical Sciences Division
Chemistry and Materials Science Department

Thomas M. Vess
Environmental Sciences Division
Physics Department

In FY-93, we continued our work on remote cure monitoring. We demonstrated our on-line Raman scattering technique for cyanate ester, a one-component resin. We also applied our methods to polymer-matrix fiber-reinforced composite materials.

Introduction

As the civilian applications for advanced materials, such as polymer matrix composites, become increasingly important, the emphasis has shifted from performance to manufacturing processes. Poor quality control and high rejection rate of the finished parts have resulted in higher costs, which render these materials less competitive in the market place. Therefore, 'smart' and careful monitoring of the processing procedure, before and during the cure, are desperately needed.

Currently, there are several techniques that are used for following the cure reactions of polymer matrix composites. With recent advances in diode laser and other optical instruments, Raman spectroscopy has become a major analytical tool for characterizing polymers and chemical reactions. Advantages over existing methods are numerous, including direct chemical information, freedom from background corrections, and minimum sample fluorescence. A new project was initiated in FY-92 to take advantage of all these unique features. The fiber optics technique was combined with Raman spectroscopy to monitor cure reactions of resins and composites remotely and *in situ*. During FY-92, we demonstrated the feasibility of remote cure monitoring of a two-part epoxy resin, using 200-μm-diameter quartz optical fibers. We validated the Raman scattering technique by demonstrating a direct proportionality between Raman peak ratios and epoxide group concentration.

Progress

In FY-93, we extended our remote, on-line Raman technique to other systems. It was demonstrated for monitoring cure of cyanate ester, a one-component resin with higher temperature capability than epoxy. The resin used was AroCy L-10 from Rhone-Poulenc. **Figure 1** shows the Raman spectra of uncured and fully cured cyanate ester resin. The extent of chemical reaction was followed by plotting scattering intensity as a function of Raman shift. The most noticeable change was the disappearance, in the cured resin on the bottom curve, of the strong vibrational doublet bands of the cyanate C≡N groups at 2234 cm^{-1} and 2268 cm^{-1}. This indicated the formation of the triazine ring from three separate cyanate groups. Moreover, the temperature of the resin during cure was also determined. The intensity ratio of the Stokes to the anti-Stokes of a vibrational band (634 cm^{-1} in this case) was shown to be similar to the temperature of the thermocouple embedded in the resin. This validated the Boltzmann equation, which states that the logarithm of this ratio is inversely proportional to the system temperature. This was the first time that the chemical reaction and temperature of a resin had been measured simultaneously with a single probe. These results were subsequently published and presented at national and international conferences.[1–3]

In addition to neat resins, fiber-optic Raman spectroscopy was also applicable to polymer-matrix fiber-reinforced composite materials. Raman vibrational bands of epoxy resin within the composites containing continuous Kevlar fibers were readily identified with a simple

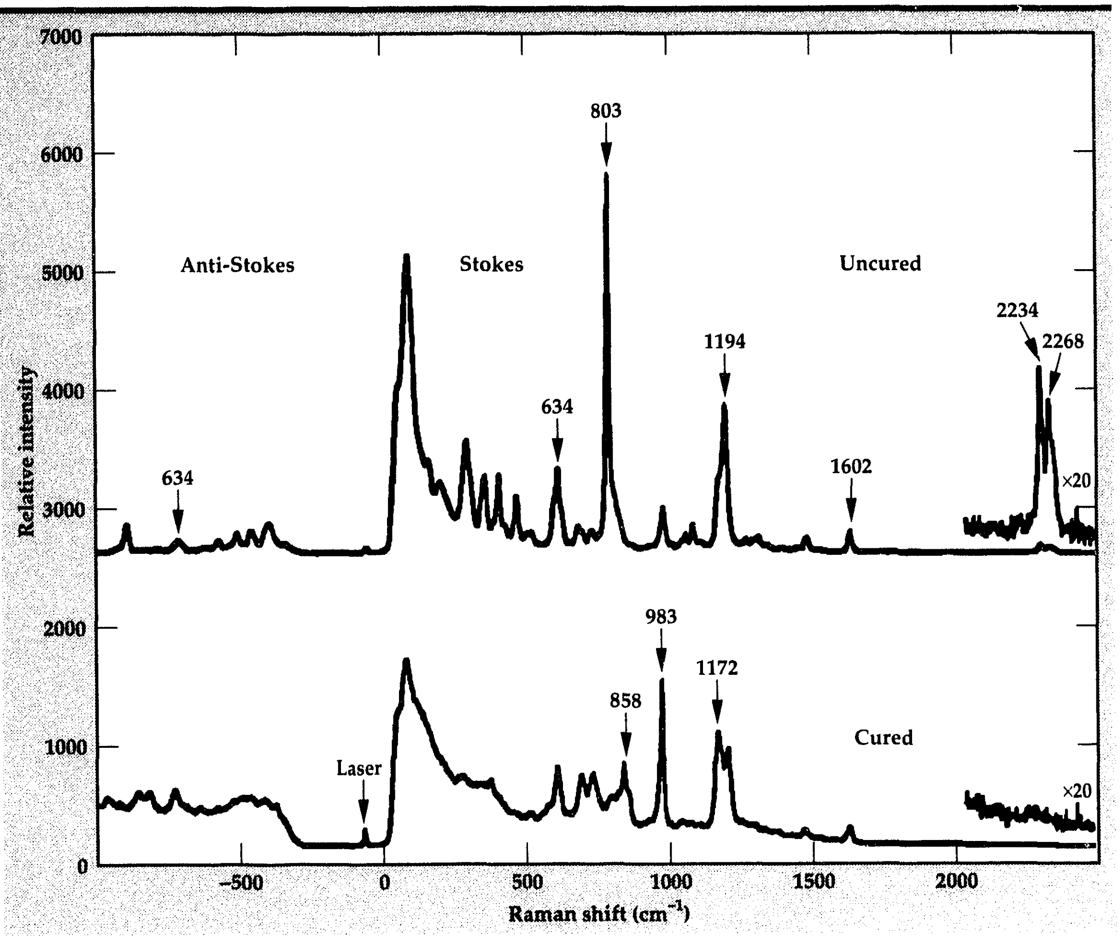

***Figure 1.* Raman spectra of uncured (top) and cured (bottom) cyanate ester resin.**

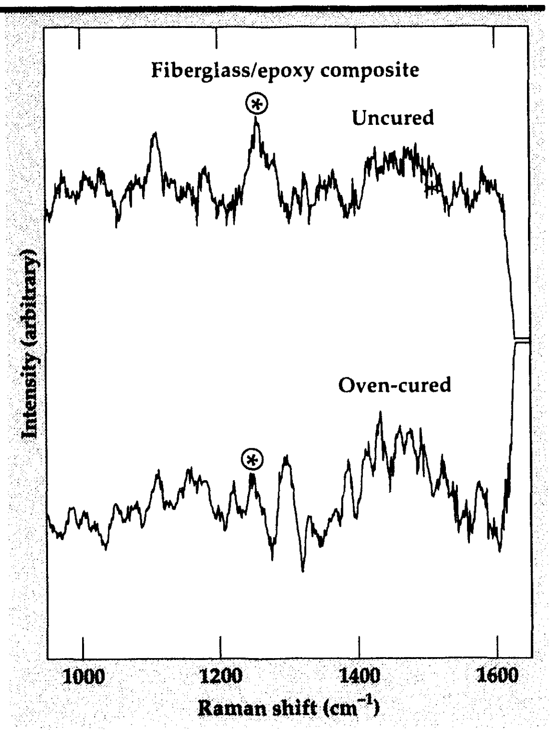

***Figure 2.* Raman spectra of composites of glass fiber in epoxy, uncured (top) and oven-cured (bottom).**

background subtraction procedure. While this was very encouraging, the fiber content of Kevlar was relatively low, ~ 30% by volume. A much higher fiber content, 60%±2% by volume, was needed for most applications. By using a laser light source with a shorter wavelength output (647 nm, instead of 773 nm), good Raman spectra of composites of glass fiber (56% by volume) in epoxy were obtained, and are displayed in **Fig. 2**. The glass fiber was flat woven mat rather than unidirectional bundles. The oven-cured sample corresponded to eight hours of cure at 100°C. The reduction of the peak around 1250 cm^{-1} (bottom curve) agrees very well with our early study of the loss in peak intensity at 1240 cm^{-1} after cure.[4]

Results from FY-92 and FY-93 have demonstrated that fiber optic Raman spectroscopy is a powerful technique for monitoring cure reactions and temperatures simultaneously for epoxy and cyanate ester. This can be done remotely and on a real-time basis. This same technology can be used equally well for composites of glass fiber and

Kevlar in epoxy. It is clear that this optical tool is applicable to automated manufacturing processes and as 'smart' sensors for polymer matrix composites.

---

1. K.C. Hong, T.M. Vess, R.E. Lyon, K.E. Chike, J.F. Aust, and M.L. Myrick, "Remote Cure Monitoring of Polymeric Resins by Laser Raman Spectroscopy," *Proc. SAMPE Int. Symp.* **38**(1), 427 (1993).
2. T.M. Vess and K.C. Hong, "Characterization of Curing Processes of High Temperature Polymers and Composites by Fiber Optic Laser Raman Spectroscopy," *Proc. 1993 Spring Meeting Materials Research Society* (San Francisco, California), (April 12–16, 1993).
3. K.C. Hong and T.M. Vess, "Remote, On-Line Cure Monitoring of Polymer Resins and Composites by Fiber Optic Raman Spectroscopy," *Proc. Gordon Research Conf.* (Plymouth, New Hampshire), (June 27–July 2, 1993).
4. M.L. Myrick, S.M. Angel, R.E. Lyon, and T.M. Vess, *SAMPE J.* **28**(4), 37 (1992).

Casting Process Modeling

Arthur B. Shapiro
Nuclear Test Engineering
Mechanical Engineering

For three years, we have worked on a computer code to predict the final shape and strength of cast-metal parts and to solve process problems such as porosity and deformation. Our work has led to a Cooperative Research and Development Agreement, through which we will continue development of the code, validate the code with casting experiments, and commercialize the code for industrial use.

Introduction

The age of high-precision industrial casting is upon us as we watch American companies, interested in cost savings, buy less expensive cast-metal parts from Europe and Japan in preference to more expensive American manufactured products. To stay competitive, we have been working on a computer code that can predict, with accuracy, the final shape and strength of cast-metal parts, as well as problems we might encounter in the casting process, such as porosity and deformation.

By simulating the casting on a computer, rather than constructing prototypes, we can cut costs and reduce the development cycle from the design inception to the introduction of the cast part into the marketplace. This computer code can also be used for process modeling, leading to intelligent process control that can greatly reduce the amount of material waste and energy consumption during manufacture. Our emphasis is on the production of components for the aerospace industry. The work also supports objectives for the precision die casting of weapon components at Lawrence Livermore National Laboratory (LLNL).

Progress

FY-93 was the third and final year in which Casting Process Modeling received funds through the Material Science and Engineering thrust area. Our work led to the execution of a Cooperative Research and Development Agreement (CRADA) in Casting Process Modeling between LLNL, Howmet Corporation, a manufacturer of cast aerospace components, and UES, Incorporated, a commercial engineering software company.

During FY-93, we added solidification kinetics to the computer code. Casting parts involves the transfer of heat, accompanied by phase change in a finite thickness interface between the liquid phase and the solid phase. The interface region, containing a mixture of solidified crystals (i.e., grains) and liquid, is also called the mushy zone. Upon solidification, the latent heat contained in the liquid is released locally to the mixture of crystals and liquid. The surrounding material conducts the heat away, permitting further solidification. If the local temperature gradient near the mushy zone is so small as to inhibit local heat loss, the latent heat release will raise the local temperature enough to remelt portions of the newly solidified crystals. This temporary reheating is called recalescence. Solidification kinetic models address the above phenomenon by modeling the physics of crystal nucleation and growth, and crystal-to-crystal interactions.

Future Work

Through the CRADA, we will continue development of the computer code, validate the code with casting experiments, and commercialize the code for industrial use.

Strength of Al and Al-Mg/Alumina Bonds Prepared Using Ultrahigh Vacuum Diffusion Bonding

Wayne E. King,
Geoffrey H. Campbell, and
Walter L. Wien
Condensed Matter
Sciences Division
Chemistry and Materials Science Department

Susan L. Stoner
Engineering Sciences
Mechanical Engineering

We have measured the cross-breaking strength of Al and Al-Mg alloys bonded with alumina. Diffusion bonding of Al and Al-Mg alloys requires significantly more bonding time than previously thought to obtain complete bonding. In contrast to previous diffusion bonding studies, fracture morphologies are similar to those obtained in bonds formed by liquid phase reaction; i.e., bonds are as strong or stronger than the ceramic, and fracture tends to propagate in the metal for pure Al and near the interface or in the ceramic for the alloys. There are indications that the fracture morphology depends on Mg content and therefore on plasticity in the metal.

Introduction

Interfacial toughness, Γ_i, is the energy required to propagate a crack stably along the interface for a unit increase in crack area. This energy is dissipated through a number of mechanisms, e.g., creation of new surface area (either advancing the crack, incidental crack branching, or microcracking); plastic deformation; and phonon emission. The magnitude of the dissipated energy will depend on a complex interplay of the operative mechanisms. Understanding and controlling interfacial toughness depends on relating this interplay to the atomic, micro-, and macrostructure of the interface, which in turn determines its adhesion and strength. A relation describing this complex interaction was first proposed by Evans and Rühle[1]:

$$\Gamma_i = F(W_{AD} - W_S) * G(\textbf{plasticity, roughness, ...}), \qquad (1)$$

where W_{AD} is the work of adhesion, and W_S is the work of segregation. It is possible to measure both Γ_i and the peak in the work of adhesion curve experimentally. The coupling of adhesion with the dissipation of energy through plasticity and roughness precludes analytical unfolding of this relation. Determination of the functions F and G requires a series of experiments in which adhesion, plasticity, and roughness are systematically varied. In the first of this series of experiments, we present results from measurement of the diffusion bonding kinetics and cross-breaking strength of Al and Al-Mg alloys bonded with alumina. Mg content was varied from 0 to 5% to systematically change the yield strength of the metal.

Progress

Bonding Kinetics of Al and Al-Mg with Sapphire

To assess the appropriate bonding conditions for the fabrication of mechanical test specimens, we have determined the conditions under which completely dense interfaces form during ultrahigh vacuum diffusion bonding, by studying the evolution of bonded interfacial areas of Al and Al-Mg/sapphire interfaces as a function of bonding pressure and time. The as-received surfaces of the metal foil specimens and the polished surfaces of the alumina were sputter-cleaned using 1-keV Xe^+ ions. Auger electron spectra were acquired to

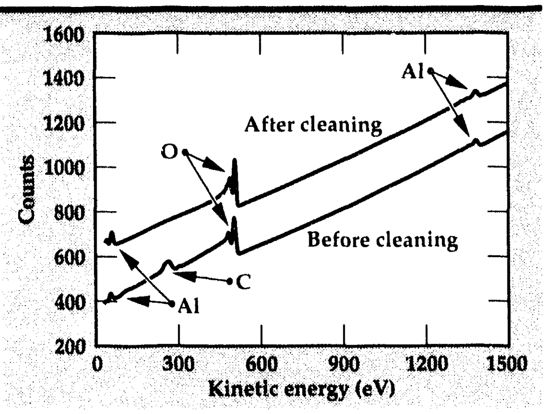

Figure 1. Effect of sputter cleaning on auger electron spectra from surface of alumina ceramic.

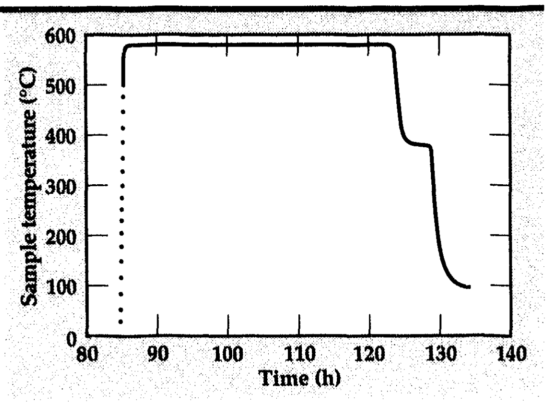

Figure 2. Typical time/temperature history for bonding of Al and Al-Mg alloys with alumina.

verify that surface contaminants had been removed. A typical example is shown in **Fig. 1**. The cleaned surfaces were placed together in the bonding chamber. The applied forces resulted in pressures from 1 to 20 MPa at the bonding interfaces.

The bonding pressure and time were varied to study their effects. The temperature was 550°C; pressures were 5 and 20 MPa; and the times were 4, 18, and 38 h. Bonds formed for shorter times at lower pressures fell apart after exposure to ambient atmosphere. During the diffusion-bonding heat cycle (**Fig. 2**), the vacuum in the bonding chamber never exceeded 1×10^{-8} Torr.

The characterization of the bonded area at the interface was done by optical microscopy, observing the interface through the sapphire. Under normal-incidence light conditions, the bonded areas are light in contrast, because the Al is conforming to the polished surface of the sapphire and reflecting light back into the objective. Results are shown in **Fig. 3**. The results indicate that achieving a fully dense interface in the Al/alumina system at 20 MPa pressure requires bonding times of 20 h at 600°C. Al-Mg alloys require 38 h at 550 to 580°C, and 20 MPa to attain a fully dense interface. This observation is remarkable in light of previous work where diffusion bonds of Al were carried out at times < 1 h.[2–4]

Cross-breaking Strength

Al and Al-Mg/alumina samples were bonded for fracture testing. These samples were diffusion bonds of polycrystalline, high purity Al foils (~ 100 µm thick), and Al-Mg alloy foils (100 and 400 µm thick), to 20-cm, high-purity, polycrystalline alumina rods. The 'sandwich' geometry was selected to minimize the contribution of bulk plasticity to the fracture-resistance measurements.[5] Surfaces of the metal foils and alumina rods were sputter-cleaned *in situ* prior to bonding, which was carried out as above. After removal from the bonding machine, each sample was cut into seven 3 x 3 x 40-mm bars and loaded on a 4-point bend fixture for testing.

Load and mid-point displacement were monitored as a function of time for a constant crosshead displacement rate. Cross-breaking strengths were determined from the maximum in the load vs displacement curves. **Table 1** shows the measured cross-breaking strengths and associated error bars (standard deviation) for the alloys investigated. All bonds are of approximately equal strength and are near that of the ceramic.

Figure 4a shows a typical load vs time curve for the pure Al samples; **Fig. 4b** shows a micrograph from a scanning electron microscope (SEM), of the corresponding fracture surface. The gradual decrease of the load after the maximum load was reached

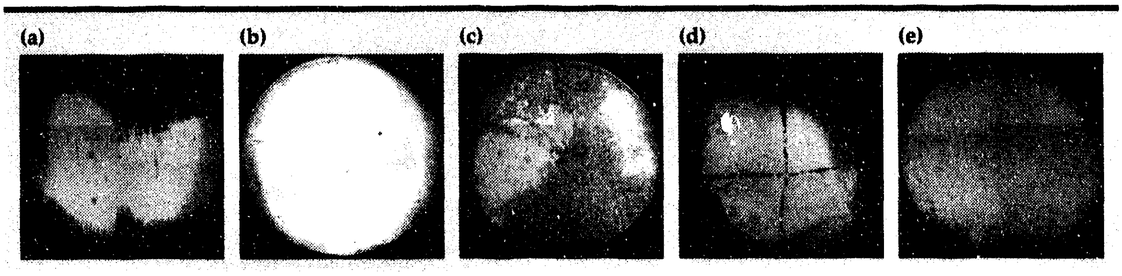

Figure 3. Optical micrographs depicting the bonding of Al-Mg alloys (a and b) with pure Al, and with a sapphire (c-e). (a) 20 MPa, 550°C, 18 h; (b) 20 MPa, 550°C, 38 h; (c) 5 MPA, 550°C, 4 h; (d) 20 MPa, 550°C, 4 h; (e) 20 MPa, 550°C, 18 h.

Table 1. Cross-breaking strengths of samples studied in this investigation.

Material	Cross-breaking strength (MPa)	Standard deviation (MPa)
Alumina	346	65
Pure Al	289	19
Pure Al	309	43
Al-2.5%Mg	288	13
Al-3%Mg	240	61
Al-3%Mg	306	48
Al-4%Mg	162	127
Al-4.6%Mg	278	7

indicates that the bond underwent a ductile rupture. Evidence of this mode of fracture is observed in the SEM image where the fracture path was through the Al, leaving ductile ligaments on both fracture surfaces. The bonded interface was sufficiently strong that the fracture propagated in the metal. This is in agreement with previous work on liquid-phase, reaction-bonded Al/alumina.[6]

Figure 5a shows a typical load vs time curve for the Al-3%Mg samples; **Fig. 5b** shows an SEM image of the corresponding fracture surface. In this case, we observed a markedly sharper decrease of the load after the maximum load was reached, compared with that for pure Al. This indicates that the bond ruptured with less ductility. This observation was confirmed by the SEM image that shows fracture near or on the interface with appreciable ductile ligaments remaining, which may have bridged the crack face.

Figure 6a shows a typical load vs time curve for the Al-4.6%Mg samples; **Fig. 6b** shows an SEM image of the corresponding fracture surface. In this case, the bond was observed to fracture brittlely when the maximum load was reached. This observation is confirmed by the SEM image, which shows fracture had propagated entirely through the ceramic, in agreement with previous observations for liquid phase-bonded Al-4%Mg bonded with alumina.[6]

Future Work

To further investigate the interface properties, we must use a stronger ceramic, sapphire. It is expected that with the reduced flaw population of the sapphire compared with polycrystalline alumina, there will be reduced tendency for fracture to initiate on the tensile face of the ceramic or, for near interface fractures, to divert into the ceramic.

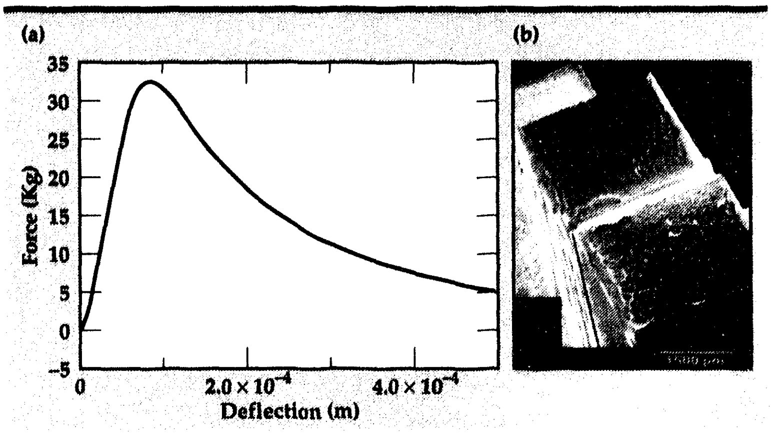

Figure 4. Load vs mid-point displacement at left (a) and SEM image from fracture surface (b) of pure Al bond.

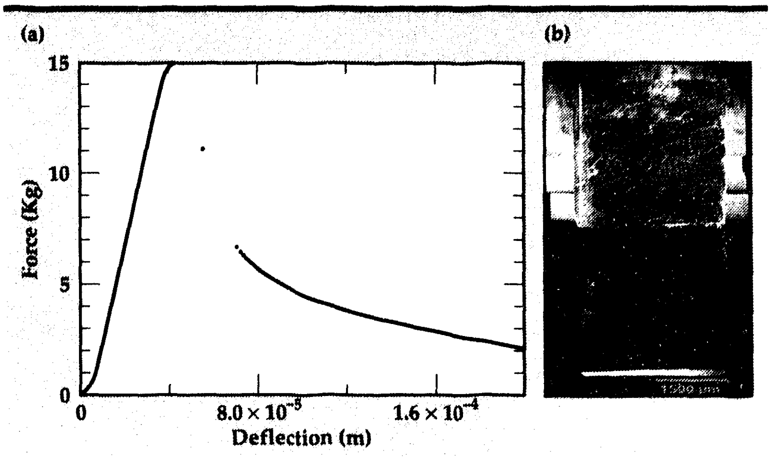

Figure 5. Load vs mid-point displacement at left (a) and SEM image from fracture surface (b) of Al-3% Mg bond.

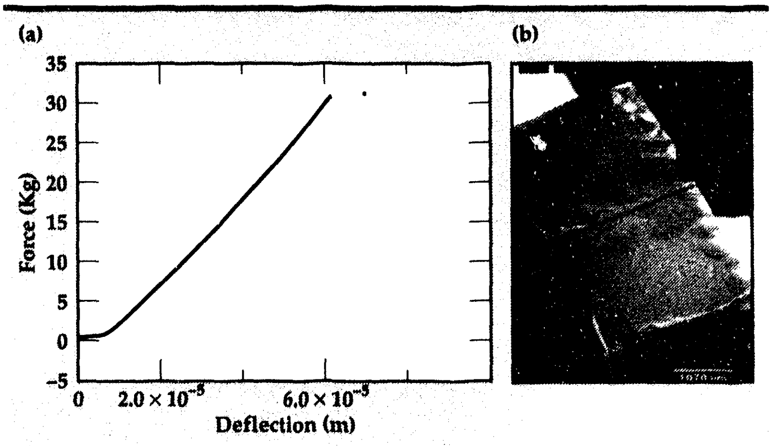

Figure 6. Load vs mid-point displacement at left (a) and SEM image from fracture surface (b) of Al-4.6% Mg bond.

---

1. A.G. Evans and M. Rühle, *Mater. Res. Bull.* XV(10), 46 (1990).
2. G. Heidt and G. Heimke, *Ber. Dt. Keram. Ges.* 50(9), 303 (1973).

3. M.G. Nicholas, "Ceramic-Metal Interfaces," *Surfaces and Interfaces of Ceramic Materials*, L.-C. Dufour, C. Monty, and G. Petot-Ervas (Eds.), NATO ASI Series E: Applied Sciences **173**, Kluwer Academic Publishers, 393 (1989).
4. B. Derby, "The Influence of Surface Roughness on Interface Formation in Metal/Ceramic Diffusion Bonds," *Ceramic Microstructures '86: Role of Interfaces*, J.A. Pask and A.G. Evans (Eds.), Materials Science Research **21**, Plenum Press (New York, New York), 319 (1986).
5. A.G. Varias, Z. Suo, and C.F. Shih, *J. Mech. Phys. Solids* **39**(7), 963 (1991).
6. B.J. Dalgleish, K.P. Trumble, and A.G. Evans, *Acta Metall.* **37**(7), 1923 (1989).

Multi-Axial Superplastic Forming Apparatus

Anne J. Sunwoo and
Susan L. Stoner
Engineering Sciences
Mechanical Engineering

We have proposed improvements to our multi-axial superplastic forming apparatus. The apparatus provides many new capabilities for studying the anisotropic behavior of superplastic alloys.

Introduction

The use of superplastic forming (SPF) technology is rapidly growing for both aerospace and automotive applications. The characterization of SPF materials has always been limited to uniaxial tensile testing. Because the materials behave anisotropically, knowledge of the mechanical properties in the multi-axial condition becomes important as technology moves from research to production.

To meet this need, we designed and built a multi-axial SPF apparatus. The major advantage of this system is that it enables us to study the anisotropic behavior of superplastic alloys. The forming behavior of superplastic materials can only be realistically simulated from multi-axial test conditions. From formability studies, we can develop forming limit diagrams, which are needed for industrial applications. The post-formed properties of SPF components can also be determined.

At Lawrence Livermore National Laboratory (LLNL), finite element analysis is currently being used to predict SPF behavior. Our apparatus provides a method of experimental validation as well as experimental data for this analysis.

Cavitation, which occurs in some superplastic Al alloys, has an adverse effect on the mechanical properties of post-formed components. It is known that the imposition of hydrostatic pressure during forming can significantly reduce or eliminate cavitation. Since our apparatus provides for the use of back pressure during forming operations, its effect on cavitation and, hence, on forming limit diagrams, can be studied.

Using this apparatus, we have successfully produced hemispherical domes of superplastic aluminum alloys.

Progress

We have proposed several improvements to the apparatus to effectively characterize a wide range of superplastic materials. To use the SPF apparatus at 1000°C, we will use new die assembly material and redesign the assembly. A finite element analysis of the dies and the furnace will be used to optimize the design of the die assembly such that heat-up time is minimum and thermal gradients in the forming part are reduced.

To study different SPF characteristics, a conical Upper Die will be made. With this geometry, much higher strains can be achieved, and post-form properties can be attained from the completed part. The current hemispherical shape can easily be modeled and allows straightforward correlation of multi-axial to uniaxial data.

A final improvement to the system includes developing software to control the pressure/time history used for forming. This software will use feedback from a linear variable differential transformer (LVDT) to adjust the forming and back pressures appropriately to maintain a constant strain rate during the forming operation. In the current system, only a 'pseudo' constant strain rate can be maintained, manually, by adjusting pressures based on the LVDT readout. The software will also be used to collect pressure, displacement, and temperature data.

This apparatus is an asset to current and future programs at LLNL involving superplasticity, providing many capabilities for weapons programs and technology transfer projects.

Power Conversion Technologies

The goal of the Power Conversion Technologies thrust area is to develop and apply advanced power electronic, pulsed power and microwave technologies to support existing and emerging programmatic requirements at Lawrence Livermore National Laboratory (LLNL). Our main objective is to work on nationally important problems while enhancing our basic understanding of enabling technologies.

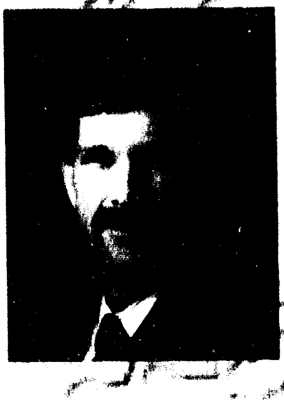

During FY-93, we concentrated our research efforts on the eight project areas described in this report: (1) Ferroelectric Cathode Measurements; (2) Pulsed Plasma Processing of Effluent Pollutants and Toxic Chemicals; (3) Plasma Afterburner for Treatment of Effluents from Solid-Waste Processing; (4) Environmentally Safe Testing of Refrigerants; (5) Compact Pulsed Power Supplies; (6) Advanced Ground-Penetrating, Imaging Radar for Bridge Inspection; (7) High-Power, Electron Beam-Induced Switching in Diamond; and (8) Applying Statistical Electromagnetic Theory to Mode-Stirred Chamber Measurements.

Mark A. Newton
Thrust Area Leader

Section 7

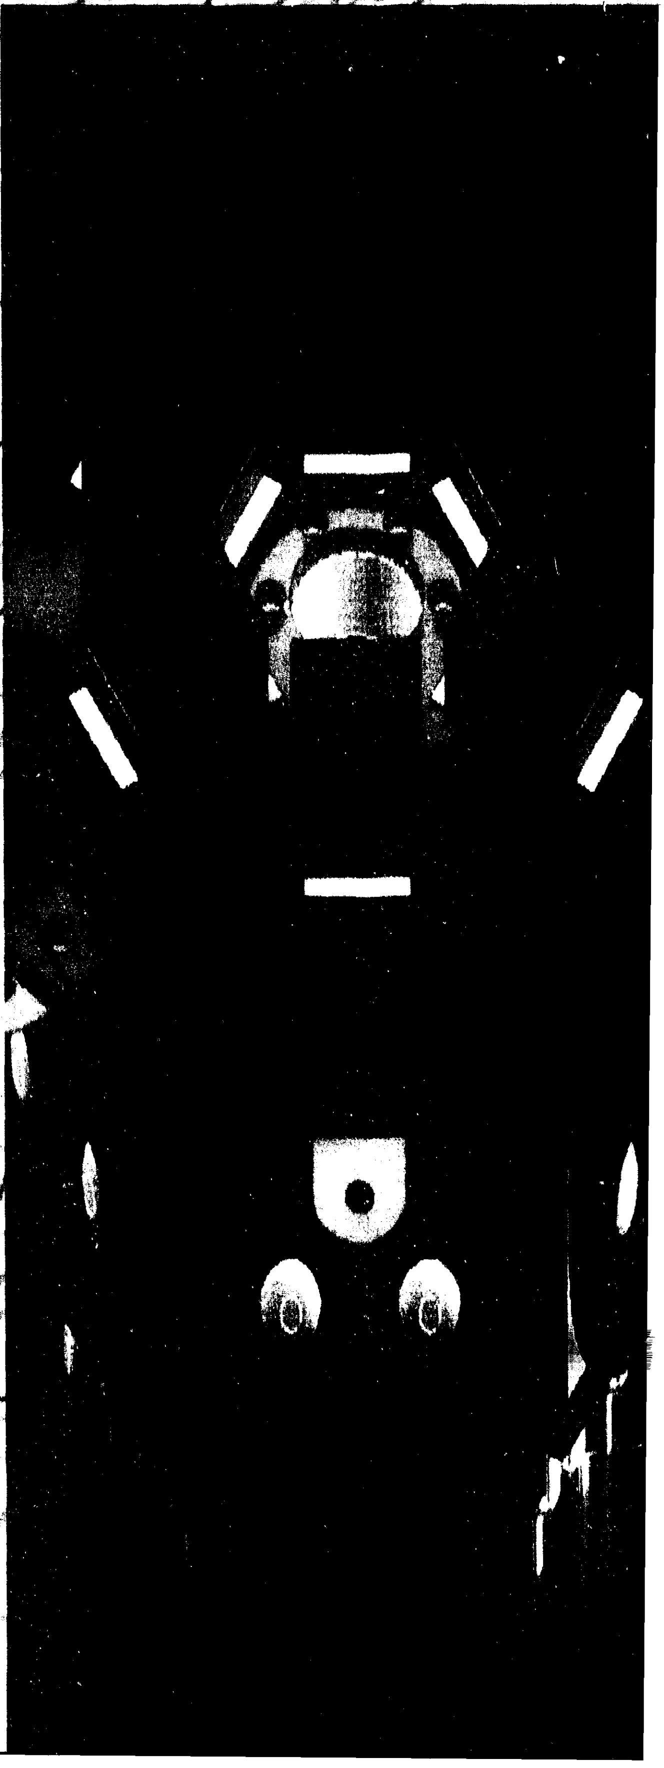

7. Power Conversion Technologies

Ferroelectric Cathode Measurements

Stephen E. Sampayan and
David O. Trimble
Laser Engineering Division
Electronics Engineering

George J. Caporaso and
Glen A. Westenskow
Laser Programs

Clifford L. Holmes
Applied Research Engineering
Mechanical Engineering

We have initiated an investigation of electron emission from ferroelectric cathodes. Our experimental apparatus consisted of an electron diode and a 100-kV, 12-ohm, 70-ns, pulsed high-voltage power source. A planar triode modulator that can be driven by a synthesized waveform generator initiates the polarization inversion and allows inversion pulse tailoring. Our initial measurements indicate that emission current densities above the Child–Langmuir space charge limit, J_{CL}, are possible. We believe that this effect is based on a non-zero initial energy of the emitted electrons and is strongly coupled to relative timing between the inversion pulse and application of the main anode-cathode pulse. We have also initiated brightness measurements of the emitted beam and estimate a preliminary lower bound on the order of 5×10^9 A-m^{-2}-rad^{-2}.

Introduction

New research into fast-switching, ferroelectric emitters promises to dramatically influence the state of the art in pulsed electron device technology. This new class of electron emitter technology has shown extremely high current densities (of order 1 kA-cm^{-2}) and high brightness (10^{10} A-m^{-2}-rad^{-2}).[1,2] The fact that these emitters can be operated at non-UHV pressures and do not require elevated temperatures or a pulsed laser system, makes them a very attractive alternative to conventional cathode technology.

The emission from a ferroelectric is believed to result from the expulsion of charge from the surface and sub-surface layers.[1,3] This effect occurs upon switching the internal polarization by a fast voltage pulse. Typical durations of the emitted beam, in a Zirconium-based material, can be of order 200 ns.[4] Since the bound charge is quite high (0.01–0.1 mC-cm^{-2}), the resultant electric fields impart significant energy to the emitted electrons. This latter phenomenon has been verified by direct measurement[5] and explains measured current densities in excess of those predicted by the Child–Langmuir space charge limit for electrons with zero initial kinetic energy.[6]

One of our motivations for pursuing ferroelectric emitter technology is upgrading the cathode brightness of the high-power, pulsed accelerators at the Lawrence Livermore National Laboratory. The Advanced Test Accelerator or ATA, a 50-MeV, 10-kA electron accelerator,[7] and the Flash X-Ray Radiography Accelerator or FXR, a 20-MeV, 4-kA electron accelerator,[8] use velvet cathodes for electron production. Brightness for a velvet cathode is typically an order of magnitude below that of a high brightness dispenser cathode used on the Experimental Test Accelerator II, a 7.5-MeV, 2.5-kA electron accelerator.[9] To implement a dispenser cathode on either accelerator, however, would require major upgrades to the injector. In addition,

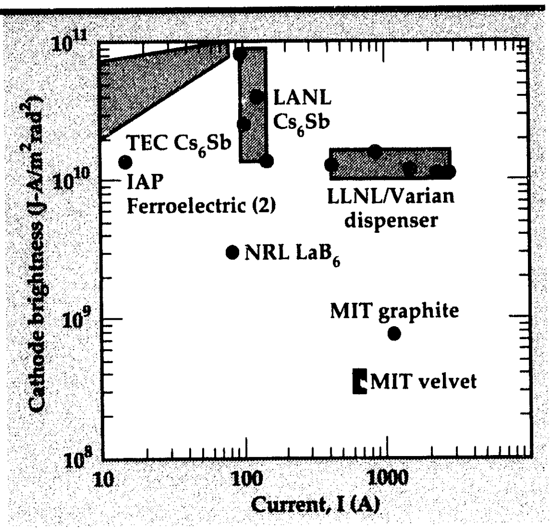

Figure 1. Cathode brightness vs current for various cathodes.

there are applications where gating the electron beam is required. Again, significant modifications to the injector would be required in these applications.

Recent ferroelectric emitter data taken at 4 A-cm^{-2}, 15 A total, showed a brightness of order 10^{10} A-m^{-2}-rad^{-2} (RMS). Comparison with other cathode brightness data vs current for various cathodes[10] is shown in **Fig. 1.** The experiment used a 1–2-kV inversion pulse with a variable DC power supply across the anode-cathode (A-K) gap[2]; i.e., emission was controlled by the inversion pulse. This data was taken at an A-K potential of 10 kV.

Our ultimate goal is to determine if a similar brightness is attainable for kiloampere electron beams. Further, since the emission appears to be dependent on the inversion pulse, we are also trying to determine to what extent the emission can be controlled by the inversion pulse. We report here on our apparatus and preliminary measurements.

Progress

Apparatus

The cathode test stand is shown schematically in **Fig. 2.** This apparatus was previously used to measure brightness of osmium-coated dispenser cathodes, but was modified to allow exploration of ferroelectric emitters.[10]

A pre-poled, 1-mm-thick × 5.1-cm-dia lead titanate-zirconate piezo-ceramic (LTZ-2)[11] was placed behind a highly transparent grid. The emission side was polished; the grid was a square mesh composed of 0.025-mm-dia tungsten wires spaced approximately 0.75 mm on centers. Sample orientation was such that the bound positive charge (negative screening charge) was positioned toward the A-K gap. All but a 3.8-cm-dia central area of the sample was apertured. The grid-emitter assembly was recessed into the cathode shroud approximately 0.23 cm; the distance from the emitter surface to the anode was approximately 0.6 cm. A copper plate mask, 0.16-cm thick, served as the anode. The entire vacuum system and housing were fabricated from stainless steel; ceramics were used as insulators. Turbo-molecular pumps were used throughout. Base pressures in the 10^{-7} T range were easily achieved and maintained.

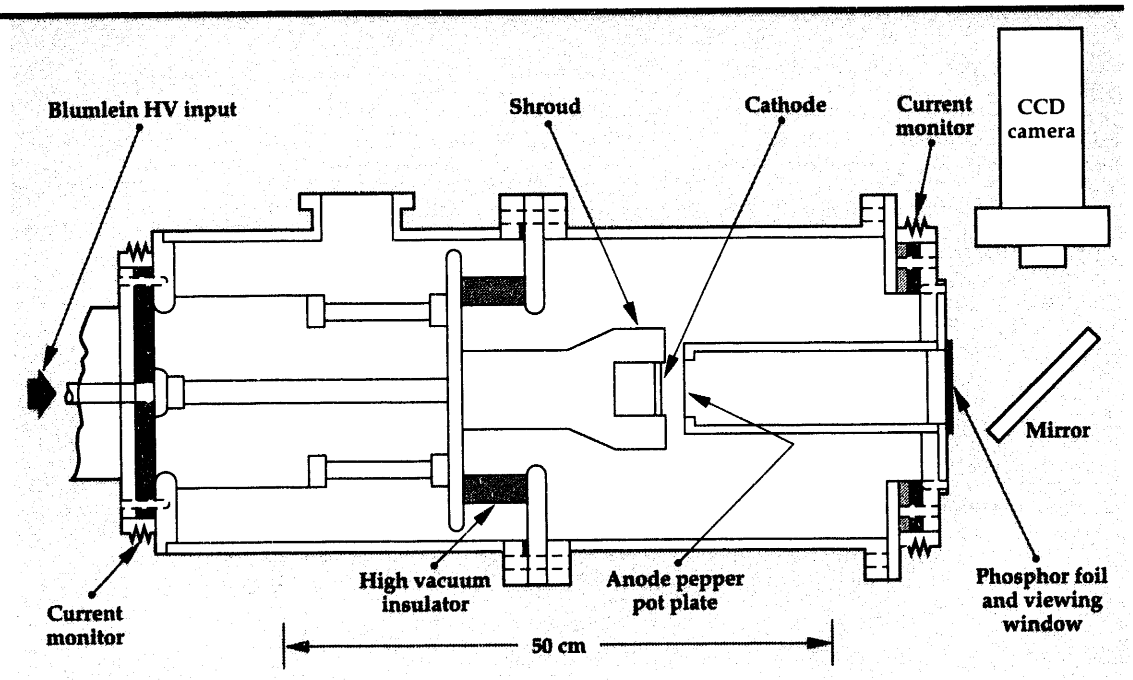

Figure 2. Cathode test stand schematic.

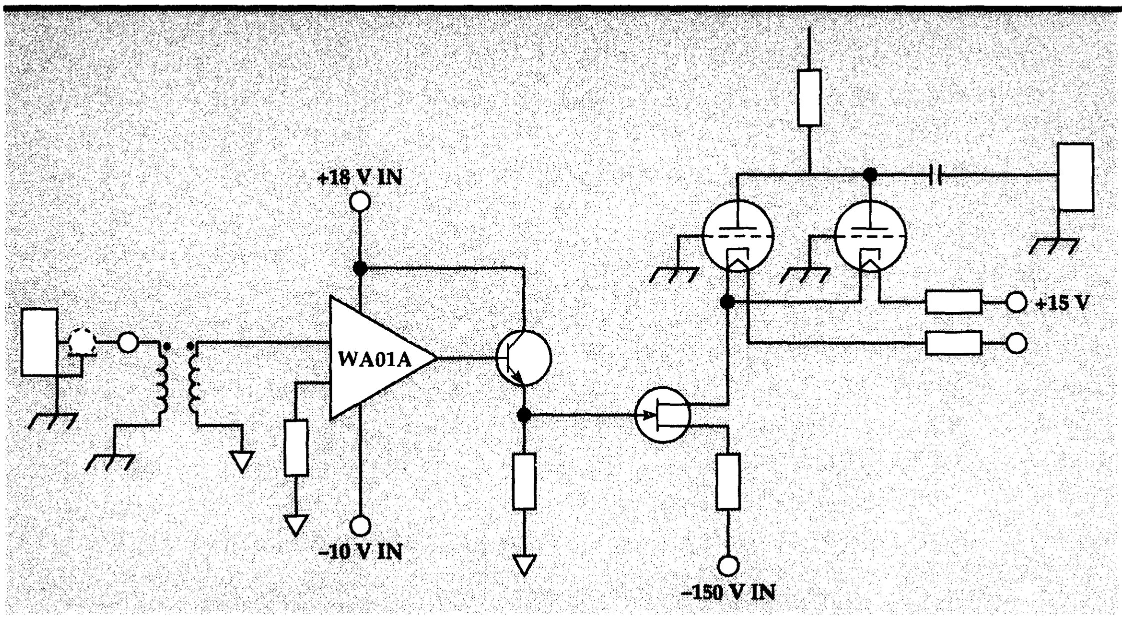

Figure 3. Inversion pulse modulator. Sixteen parallel units provide the polarization pulse.

Main current was measured by resistive current monitors; A-K potential was measured by a capacitive probe calibrated by a standard resistive probe. Data was recorded with fast analog oscilloscopes.

Brightness measurements were performed with an anode mask consisting of nine 50-μ holes placed in an 'X' pattern and a fast phosphor, ZnO(Ga), to image the pattern after a 24.5-cm drift distance. Charging was eliminated with a conductive film placed over the phosphor substrate. An image-intensified gated camera allowed observation of the phosphor image. Gate width used was typically 40 ns.

The polarization inversion pulse was coupled to the back plane of the ferroelectric emitter via a 10-ohm coaxial cable. The modulator system was at ground potential while the ferroelectric assembly and cathode stalk were elevated to the accelerating potential. The anode was maintained essentially at ground potential through the resistive current monitor. The 10-ohm coaxial cable wound into an inductor provided the isolation impedance from the cathode stalk to ground.

A polarization inversion pulse was supplied by a set of 32 planar triodes in parallel (**Fig. 3**).[12] A total current of approximately 0.5 kA at a maximum voltage of 20 kV could be obtained by discharging a 0.05-μF capacitor connected to the ferroelectric backplane via the 10-ohm coaxial cable. The bandwidth of the system was approximately 50 MHz.

The high-voltage pulsed power was supplied by a water-filled Blumlein. For double pulse capability, we used a magnetically delayed low pressure switch (MDLPS).[13] Double pulses as closely spaced as 50-μs were possible. The system was capable of 5 Hz continuous repetition rate. Although the present MDLPS limited output to below 100 kV, the Blumlein and power supply system were capable of output pulses on the order of 250 kV into 12 ohm.

Experimental Results

Emission from the ferroelectric emitter was achieved with a negative 1–2-kV pulse. Approximately 20,000 shots were placed on the ferroelectric emitter at 1 pps, with no discernible degradation in the emission characteristics. A-K potential was varied from approximately 10 to 60 kV. Typical waveforms are shown in **Fig. 4**. Peak currents, J, of approximately 400 A were achieved during these initial measurements. As was generally observed

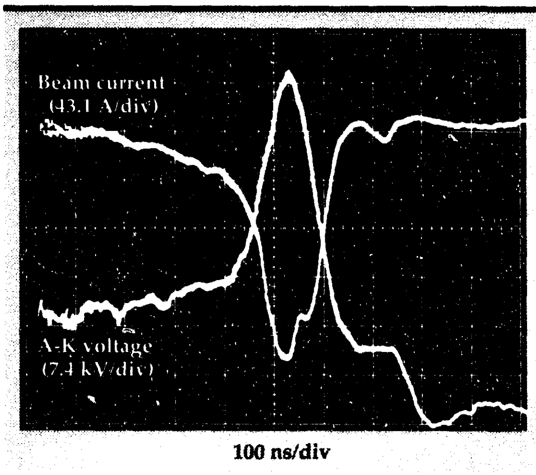

Figure 4. Typical A-K voltage and current pulse.

at lower A-K potentials, the measured currents were unipolar, consistent with electron flow, and above the Child–Langmuir limit, J_{CL}, for electrons with zero initial kinetic energy.

During initiation of the inversion processes, we characteristically observed multiple discharges between the wire mesh and the ferroelectric emitter. Nevertheless, unipolar current flow is generally observed in the main A-K gap. We also correlated the timing of these discharges with the onset of electron current. In general, we observed discharge initiation upon application of the inversion pulse. Strong electron emission did not occur until several hundred nanoseconds later, however.

To optimize the emission, timing between the inversion pulse and the main A-K pulse was varied (**Fig. 5**). For a fixed A-K potential of 37 kV, the emitted current varied from zero to approximately 250 A over this timing variation. The optimum timing varied slowly about the peak, showing a 25% decrease in the emission current for a timing variation of ± 100 ns from the optimum.

The emitted current from the ferroelectric was measured as a function of the applied A-K potential (**Fig. 6**). In this data, the A-K potential was varied from 5 kV to 50 kV; the emitted current varied from 95 A to 411 A for the 11.3 cm^2 cathode. Some scatter, on the order of 10%, was present in the data and consistent with that previously measured by others.[4] The data also indicated that the emitted current exceeded J_{CL} at lower A-K potentials. The largest ratio of J/J_{CL} occurred at low energies, on average, with much decreased ratios at increased energies.

A plot of the ratio of J/J_{CL} as a function of electric field for two different A-K gaps reveals consistent behavior and decreased emission at higher electric fields (**Fig. 7**). From this data, the emitter appears to be emission-limited above an electric field of 20 kV/cm. We speculate that this effect may result from the interaction of the A-K pulse and inversion pulse.

Initial brightness measurements were attempted with the multi-aperture mask. The timing between the A-K potential and polarization inversion pulse to the applied A-K potential timing was varied to ensure an optimum. Data was taken at A-K potentials of 16 kV and 33 kV. Peak gap current varied from 86 A at 16 kV to 207 A at 33 kV.

During the measurement, we observed distortion of the mask pattern on the phosphor. Further, the mask image and the intensity of the image of the beamlets at the phosphor were inconsistent from shot to shot at the lower A-K potential.

Figure 5.
Electron current as a function of the timing between polarization inversion pulse and beginning of emission.

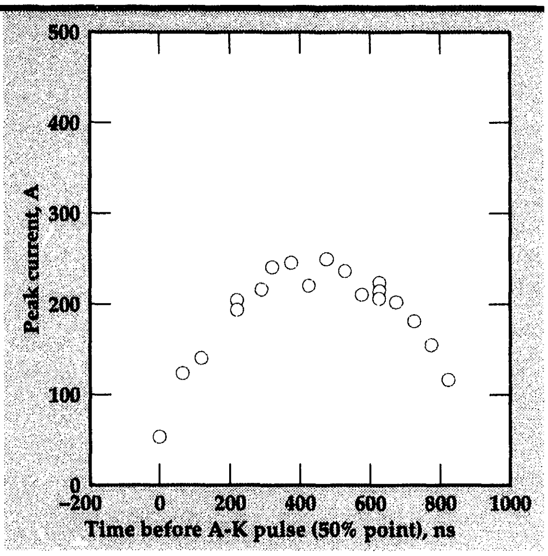

Figure 6.
Voltage-current characteristics of A-K gap.

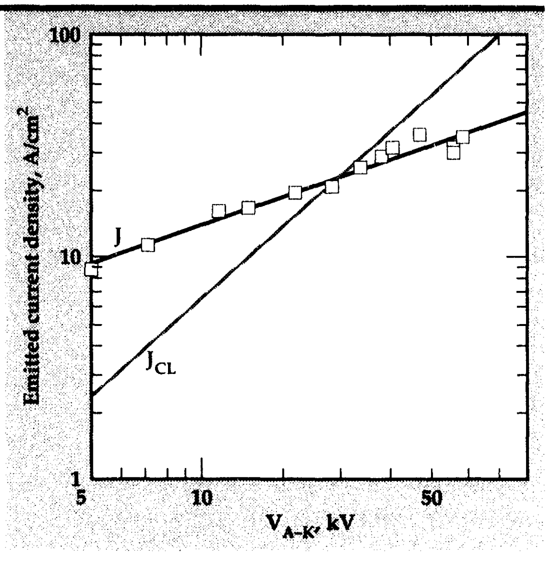

Figure 7. Space charge limit ratio as a function of electric field.

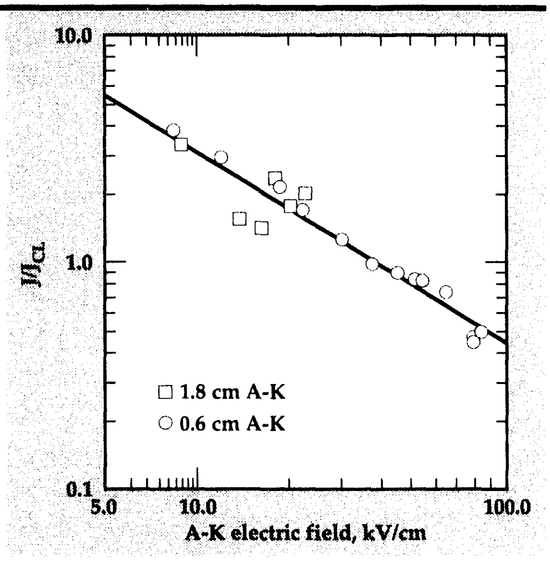

From the expansion of the beamlets, and inferring the space-charge expansion by aperture hole size and current density at the cathode surface, we estimated the lower bound of the beam brightness (based on the 2 sigma or edge of the distribution). As in similar measurements,[10] we estimated this lower bound at 5×10^9 A/m^2-rad^2 at 206 A and 10^{10} A/m^2-rad^2 at 86 A.

Future Work

In FY-94, we plan to modify the grid structure on the ferroelectric surface, investigate alternate emitter materials, and make additional measurements.

---

1. H. Gundel, H. Riege, J. Handerek, and K. Zioutas, *Pulsed Electron Emission from Ferroelectrics*, CERN Report No. CERN/PS/88-66 (AR), (1988).
2. G. Kirkman, B. Jiang, and N. Reinhardt, "A High Brightness Electron Beam Produced by a Ferroelectric Cathode," presented at the IEEE International Electron Devices Meeting (San Francisco, California), (1992).
3. L. Schachter, J.D. Ivers, J.A. Nation, and G.S. Kerslick, *J. Appl. Phys.* **73**(12), 8097 (1993).
4. J.D. Ivers, R. Advani, J.A. Nation, and L. Schachter, *Proc. 1991 IEEE Particle Accelerator Conf.* (San Francisco, California), 2104 (1991).
5. H. Gundel, H. Riege, E.J.N. Wilson, J. Handerek, and K. Zioutas, *Nucl. Instrum. Methods* **A280**, 1 (1989).
6. S.E. Sampayan, G.J. Caporaso, C.L. Holmes, E. Lauer, D. Prosnitz, D.O. Trimble, and G.A. Westenskow, "Emission from Ferroelectric Cathodes," presented at the Workshop on High Intensity Electron Sources (Legnaro, Italy), (1993).
7. L.L. Reginato, *The Advanced Test Accelerator (ATA): A 50-MeV, 10-kA Induction Linac*, Lawrence Livermore National Laboratory, Livermore, California, UCRL-88877 (1983).
8. B. Kulke, D.S. Ravenscroft, and G.E. Vogtlin, *Design and Performance of the 1.5-MV Injector for FXR*, Lawrence Livermore National Laboratory, Livermore, California, UCRL-85060 (1981).
9. J.C. Clark *et al.*, *Design and Initial Operation of ETA-II Induction Accelerator*, U.S. Department of Energy Report No. DOE ER-40150-126 (1988).
10. W.C. Turner, Y.-J. Chen, W.E. Nexsen, M.C. Green, G. Miran, and A.V. Nordquist, "High-Brightness, High-Current Density Cathode for Induction Linas FEL's," *Proc. 1988 Linear Accelerator Conf.* (Williamsberg, Virginia), (1988).
11. Transducer Products, Route 63 North, Goshen, Connecticut.
12. E.E. Bowles and W.C. Turner, *Digest of Technical Papers, 7th IEEE Pulsed Power Conf.* (Monterey, California), 43 (1989).
13. S.E. Sampayan, H.C. Kirbie, E.J. Lauer, A.N. Payne, D. Prosnitz, and D.O. Trimble, "Ferroelectric Cathode Measurements," *Digest of Technical Papers, 9th IEEE Pulsed Power Conf.* (Albuquerque, New Mexico), (1993).

Pulsed Plasma Processing of Effluent Pollutants and Toxic Chemicals

George E. Vogtlin
Defense Sciences Engineering Division
Electronics Engineering

We are exploring the efficiency of pulsed plasma processing in the removal of nitrogen dioxide, nitrogen oxide, and other pollutants. This process uses an electrical discharge to create chemical radicals from air molecules. These radicals can react with pollutants and form harmless compounds. We also use additives such as hydrocarbons to improve the efficiency of the removal. The efficient removal of nitrogen dioxide has required the presence of dilute aqueous solution of ammonia.

Introduction

Pulsed plasma devices produce a discharge in which a majority of the electrical energy goes into the production of energetic electrons, rather than into gas heating. Even though the electrons are short-lived under atmospheric pressure conditions and rarely collide with a pollutant molecule, they undergo many collisions with the dominant background gas molecules, thus producing radicals that, in turn, decompose the toxic compounds. The efficiency of the approach arises from the fact that the radicals have long lifetimes and react selectively with the contaminant molecules.

The device we use for this discharge is a wire in a tube. High voltage is applied in pulses shorter than 200 ns. These high-voltage pulses achieve non-thermal conditions through the production of short-lived microdischarges. Non-thermal means the electrons are heated, not the gas molecules. These microdischarges yield a large improvement in the power efficiency because, within the short lifetime of each microdischarge, the ions do not experience significant movement and therefore do not contribute to the power consumption. The short lifetime of these microdischarges is accomplished with the use of very short high-voltage pulses (pulsed corona discharge) and/or with the use of dielectric coatings on the electrodes (dielectric barrier discharge).

It is known that non-thermal techniques are particularly efficient only when the pollutant is present in very small concentrations. At higher initial pollutant concentrations, the electrical energy consumption becomes comparable, or may even exceed, that of thermal techniques. For most applications, the energy consumption can only be reduced by using reagents that effectively recycle the active radicals. The removal of NO_x using non-thermal plasmas can be implemented without using catalysts or additives, but only at a high electrical energy cost. The use of additives, such as ammonia or urea-based substances, becomes necessary in most cases just as in the best available 'conventional' NO_x control technologies. In our most recent work, we investigated the use of hydrocarbon additives to reduce the non-thermal $DeNO_x$ energy cost.

Progress

Experimental Setup

The gas system, plasma reactor structure, electrode geometries, experimental parameters, and diagnostics are as follows.

The principal experimental apparatus used in this effort is a closed-loop gas system with mixtures of bottled gas to simulate the flue gas. A drawing of the system is shown in **Fig. 1**. This system permits the introduction of various gas combinations prior to testing. Analysis can be conducted during or after these tests. The recirculated gas can be heated to determine its effect on the electrical discharge, chemical reactions, and water content.

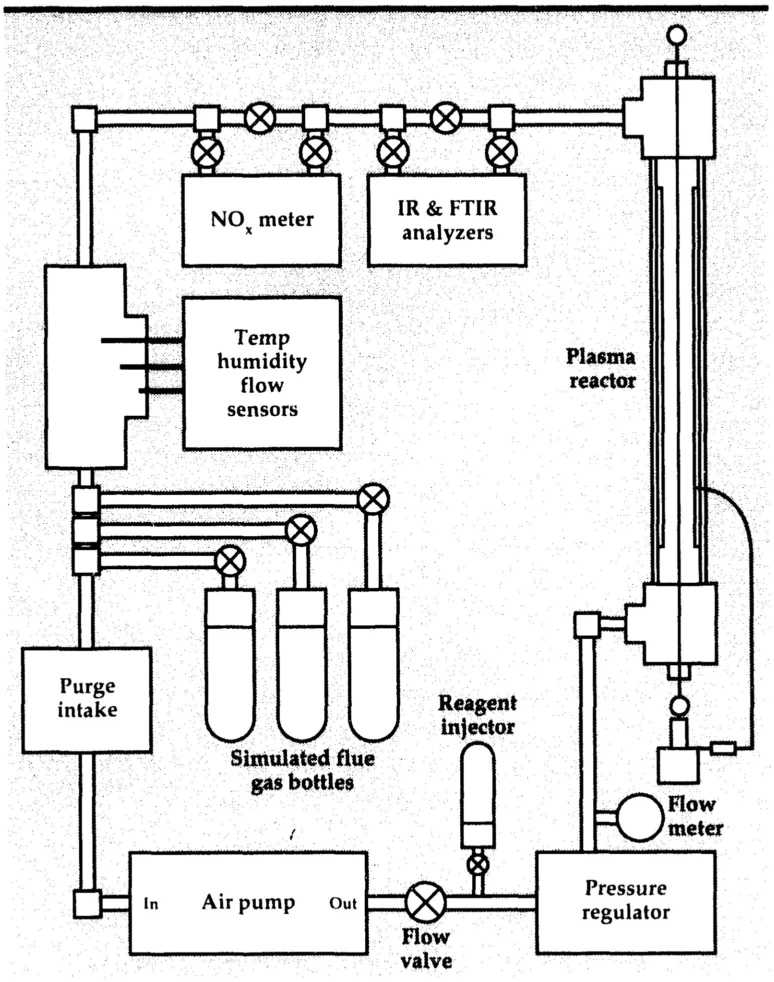

Figure 1. Schematic of the gas system.

Figure 2. Structure of the pulsed plasma reactor.

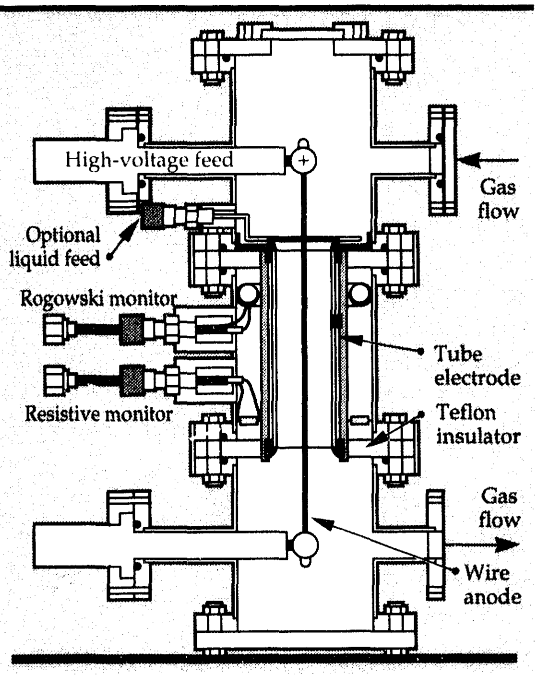

The closed-loop system permits testing of processes without a requirement for the development of complex and expensive power supplies for the one pass treatment. It also permits high flow rates without an extensive gas mixing and heating system. One pass treatment can be simulated by adjusting the gas flow rate and residence time.

The processing chamber has been designed with a 2-in.-dia outer pipe. This tube can be used as an electrode; other geometries of smaller dimensions can be placed inside. A cross-section is shown in **Fig. 2**. This processor can accommodate electrode geometries up to 2 in. in diameter. The current reaction chamber is 15.24 cm in length, but can be increased as needed to match the impedance of the cable to that of the processor, for maximum energy transfer.

The corona discharge that we are investigating is generated from electrodes, one of which has a stress on its surface that is greater than five times the average stress. This ratio is commonly called f^* (peak E field/average E field). Below an f^* of five, breakdown occurs before corona occurs.[1] We have used f^*'s of 7.7 or greater. Breakdown after the onset of corona discharge can occur but is a function of the stress and pulse length. We have constructed a processor that can accommodate many electrode shapes. **Figure 3** shows a discharge of a 1.5-mm-dia wire 2 cm from a ground plane. This discharge was a 100-ns pulse at 40 kV.

The structure of the corona is a streamer originating at the positive wire. The wire is bent up at the ends, and the streamers follow the electric field. A cathode spot is visible at the lower end on the ground plane. One or more of the streamers is much brighter and will eventually result in a high temperature arc. This arc formation results in a low-impedance short circuit and creates nitric oxides.

The voltage and pulse length applied to electrodes determine the time to breakdown if the f^* is above five. As the voltage increases, breakdown occurs at an earlier time. Open-shutter photographic imaging of this process using different pulse lengths shows first faint streamers that have crossed the gap except for a small space at the cathode (low stress electrode). These streamers then grow in intensity, and a spot appears on the cathode surface opposite the streamer. One of the streamers and its spot grow brighter than the others, then proceed to breakdown.

One scenario for this process is as follows: the streamers have an impedance greater than tens of kilo-ohms. As the streamer approaches the cathode, it must charge an increasing capacitance.

The limited charging current through the streamer reduces the potential across the streamer to cathode gap. This slows the streamer progress. The approaching streamer increases the stress on the cathode, which develops a cathode spot. Breakdown is then due to the heating of the cathode spot.

It has been observed that if the voltage is raised close to breakdown, then NO is produced, not removed. In other cases, even though the NO removal rate increases with voltage, the efficiency decreases. This implies that the optimum voltage for efficient NO removal be much less than the breakdown threshold.

The pulse generator used in our tests is capable of 50 kV into a matched load or 100 kV into a high impedance load. Pulse lengths can be varied from 10 ns to greater than 1 μs depending on the charging cable length. The risetime from this cable pulser is 5 ns. The risetime of the pulse can be varied by the addition of an RC network, and the pulse rate can be increased up to 25 pps. Most of the data presented here were taken using a pulse rate of around 20 pps and a pulse length of around 80 ns. Typical voltage and current waveforms are shown in **Fig. 4**.

Diagnostic systems have been used to analyze the light emissions of the discharge. These include a monochromator and an open-shutter camera. Devices to measure the results of chemical reactions include a chemical NO_x meter, a chemiluminescence NO_x meter, and an infrared (IR) and Fourier transform infrared red (FTIR) analyzer.

The chemical NO_x meter was calibrated on a daily basis using a calibration bottle of 100 ppmv of nitric oxide in dry nitrogen. The chemiluminescence NO_x meter was checked for linearity against a bottle of 100 and 510 ppmv of NO in nitrogen and was linear to one percent.

We are using the IR and FTIR analyzers to monitor the levels of secondary pollutants and harmful by-products, for example N_2O, that may be produced by the plasma processing. The FTIR analyzer provides quantitative determination of several components because it measures the IR absorption spectrum over a wide range of wavelengths.

The measurement of energy is essential to determine the efficiency. We measure the energy by measuring the voltage and current as a function of time and then integrate the product as a function of time. The pulse rate is measured by a counter. To prevent reflections, a load resistor is included at the end of a short transmission line. This is where the voltage is measured with a Physics International voltage divider. The current is measured with a 0.1-ohm resistor in the return path.

***Figure 3.* Coronal discharge showing early arc formation.**

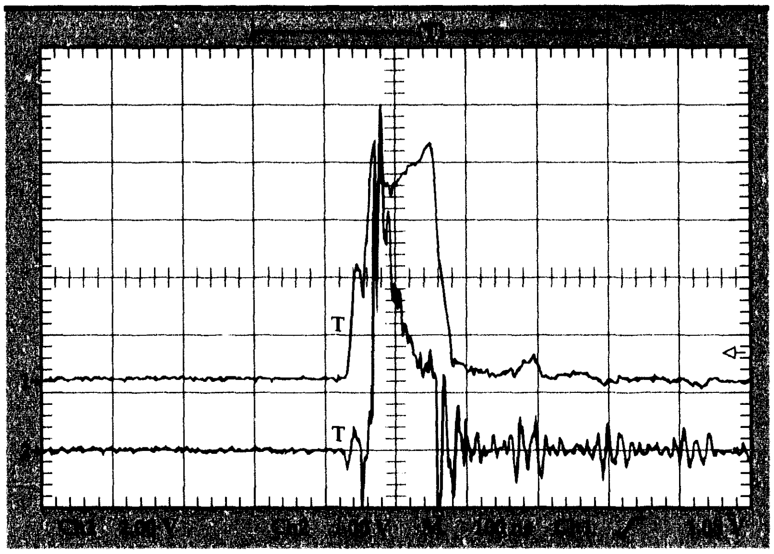

***Figure 4.* Typical voltage and current waveforms in our plasma reactor. The voltage (upper curve) is 10 kV per division. The current (lower curve) is 98 A per division. The time is 100 ns per division.**

We have imaged the discharge using an open-shutter camera. This, combined with the ability to generate pulse lengths of from 20 ns to 1000 ns, permits time resolution studies of the streamer evolution. We have also imaged the discharge with a monochrometer to measure the electron temperature.

Experimental Results

The following results were obtained using the wire-cylinder pulsed corona system. The gas flow was recirculated but was in turbulent flow through

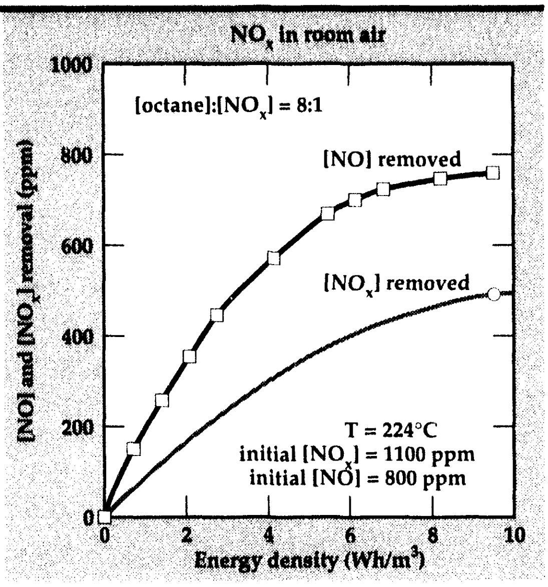

Figure 5. Measured NO and NO_x removal as a function of energy density input into the discharge. The gas temperature was 224°C. The gas was room air containing initial concentrations of 1100 ppm and 800 ppm of NO_x and NO, respectively. The n-octane additive was added at a molar ratio to initial [NO_x] of 8:1.

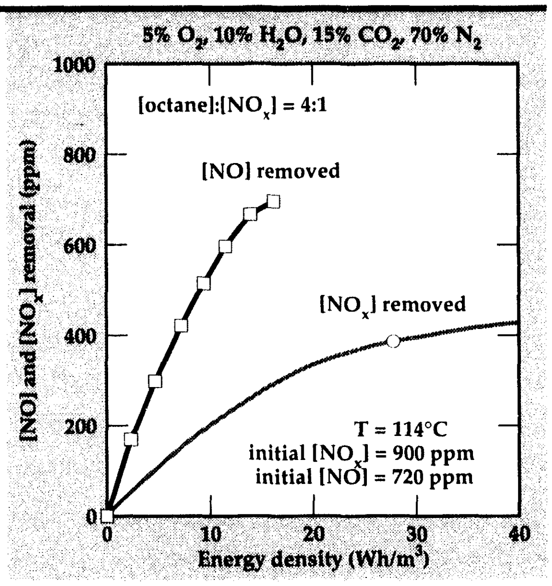

Figure 6. Measured NO and NO_x removal as a function of energy density input into the discharge. The gas temperature was 114°C. The gas consisted of 5% O_2, 10% H_2O, 15% CO_2, and 70% N_2, with initial NO_x and NO concentrations of 900 ppm and 720 ppm, respectively. The n-octane additive was added at a molar ratio to initial [NO_x] of 4:1.

the plasma reactor. **Figure 5** shows the removal of NO and NO_x in air at a gas temperature of 224°C, using a relatively large amount of n-octane additive. The molar ratio of n-octane to initial NO_x used in this run was 8:1. About 750 ppm of the initial 800 ppm NO was removed with an energy consumption of 9.5 Wh/m³. About 500 ppm of the initial 1100 ppm NO_x was removed with the same energy consumption. This energy consumption corresponds to 18 eV per removed NO molecule, and 29 eV per removed NO_x molecule.

Figure 6 shows the removal of NO and NO_x in a gas mixture consisting of 5% O_2, 10% H_2O, 15% CO_2, and 70% N_2. The gas temperature was 114°C. The molar ratio of n-octane to initial NO_x used in this run was 4:1. Almost 100% of the initial 720 ppm NO was removed at an energy consumption of 16 Wh/m³, corresponding to 28 eV per removed NO molecule. About 45% of the initial 900 ppm NO_x was removed at an energy consumption of 28 Wh/m³, corresponding to 80 eV per removed NO_x molecule.

The above results show that hydrocarbon additives could be very effective in reducing the energy cost for non-thermal DeNO_x. For the same energy input, higher NO_x reductions are observed with higher hydrocarbon additive concentrations. However, care has to be taken with regards to secondary pollution production due to unreacted hydrocarbons ('slip'). It is therefore desirable to achieve the NO_x reduction using the least amount of additive at an acceptable energy consumption. These results show that it is rather energy intensive to remove large amounts (> 500 ppm) of NO_x at acceptable injection levels of additives (molar ratios of 1:1 or less).

We believe the hydrocarbon additive serves by recycling the hydroxyl radicals during the oxidation and reduction of NO. The efficiency of a particular hydrocarbon is therefore limited by the reaction rate of this hydrocarbon with hydroxyl radicals. **Table 1**[2] shows the reaction rates for a variety of hydrocarbons. Note that n-octane and C_2H_4 have almost the same reaction rate with

Table 1. Interaction of hydroxyl radicals with hydrocarbons.

Species	Rate (cm³/molecule-s)	Rate (298 K) (cm³/molecule-s)	Rate (400 K) (cm³/molecule-s)
CH_4	$6.95\times10^{-18}T^2\exp(-1280/T)$	8.41×10^{-15}	4.53×10^{-14}
C_2H_6	$1.37\times10^{-17}T^2\exp(-444/T)$	2.74×10^{-13}	7.22×10^{-13}
C_3H_8	$1.27\times10^{-17}T^2\exp(+14/T)$	1.18×10^{-12}	2.10×10^{-12}
C_4H_{10}	$1.55\times10^{-11}\exp(-540/T)$	2.55×10^{-12}	4.02×10^{-12}
$CH(CH_3)_3$	$9.58\times10^{-18}T^2\exp(+305/T)$	2.37×10^{-12}	3.29×10^{-12}
Cyclohexane	$2.73\times10^{-11}\exp(-390/T)$	7.38×10^{-12}	1.03×10^{-11}
n-Octane	$3.12\times10^{-11}\exp(-380/T)$	8.72×10^{-12}	1.21×10^{-11}
C_2H_4	$2.15\times10^{-12}\exp(+411/T)$	8.53×10^{-12}	6.01×10^{-12}

OH at room temperature. However, at higher gas temperatures, C_2H_4 becomes less effective because of its negative activation energy.

Another option is to convert the nitric oxide to nitrogen dioxide and then remove the NO_2 by a chemical reaction. In our experiments, NO_2 would not react efficiently with ammonia in gaseous form up to temperatures of 200°C. Efficient chemical removal of NO_2 was achieved by passing the gas over a very dilute aqueous solution of ammonia. This removal produces ammonium nitrate, which can be used as a fertilizer.

Conclusion

We have investigated the effect of hydrocarbon additives in lowering the energy consumption of non-thermal $DeNO_x$ using a pulsed corona reactor. Based on our measurements of the energy consumption for the removal of various amounts of NO_x, we have made estimates of the power consumption for implementing this technique. Non-thermal $DeNO_x$ techniques are particularly efficient when the initial concentration of NO_x is low. At high concentrations of NO_x, the electrical power consumption of this technique is rather high, and it will be necessary to use or combine this technique with thermal and/or catalytic methods. The implementation will largely depend on how many additives and how much electrical power consumption and final NO_x concentration are tolerable for a particular application. The removal of nitrogen dioxide efficiency can be improved by the presence of a dilute aqueous solution of ammonia.

---

1. L.L. Alston, *High Voltage Technology*, Oxford University Press (London, England), 1968.
2. R. Atkinson, *Chem. Rev.* **86**, 69 (1986).

Plasma Afterburner for Treatment of Effluents from Solid-Waste Processing

Robert A. Anderson,
David A. Goerz, and
Roger A. Richardson
Defense Sciences Engineering Division
Electronics Engineering

John A. Kerns
Applied Research Engineering
Mechanical Engineering

Our objective is to develop technology and expertise for the destruction and stabilization of hazardous mixed wastes, using high-temperature plasma systems. Plasma technologies for environmental cleanup and waste management are rapidly emerging. In particular, plasma-arc vitrification has proven effective in encapsulating inorganic compounds in a non-leachable, glass-like slag while also destroying organic compounds by the extreme heat of the process. At Lawrence Livermore National Laboratory (LLNL), we are building a science-based program to develop thermal plasma processing technologies for the treatment of hazardous, radioactive, and mixed wastes. We have formed a multi-disciplinary project team to draw upon LLNL's advanced computational capability and experimental facilities. To help provide focus, we are collaborating with an industrial partner (Retech, Inc.) to develop the apparatus needed to improve their method for high-temperature disposal of hazardous waste materials. We have assembled an RF plasma torch and afterburner reaction chamber in the laboratory, and have made preliminary measurements to determine (1) electrical parameters, (2) plasma temperature, (3) gas flow and mass balance, and (4) chemical composition of the high-temperature gas. From this work, we developed a diagnostic module that will be used to measure the off-gas constituents of a 1/4-scale plasma-arc furnace. From these results and the results of our modeling effort, the hardware can be designed.

Introduction

Plasma torches have been used for decades in many industrial applications. The technology is well established in metallurgical processing, materials synthesis, and surface coating. In recent years, plasma technologies have appeared in innovative hazardous-waste processing systems.[1] The rationale for using plasma torches is their ability to destroy organic and biological components while decreasing the volume of waste and producing a non-leachable, glass-like slag suitable for landfill. Retech, Inc. (Retech) has been a leader in developing plasma-arc furnaces for waste destruction and disposal.[2] The method of using heat generated from a plasma torch to treat metal and organic contaminated soil was first demonstrated on a large scale at the Superfund site at Butte, Montana.[3] This work was jointly sponsored by the Department of Energy and the Environmental Protection Agency.

Figure 1 shows a conceptual hazardous-waste-treatment system. The primary reactor, a plasma-arc centrifugal treatment (PACT) system, uses heat from a plasma torch to melt and vitrify solid feed material. Organic components are vaporized and decomposed by the intense heat and become an effluent gas that must be further processed. The key to making such a system practical and environmentally acceptable is the successful development of a cost-effective method to treat the hazardous-gas effluent. We are working in collaboration with Retech to develop and demonstrate a plasma-fired, secondary reactor (afterburner) that will satisfy the need.

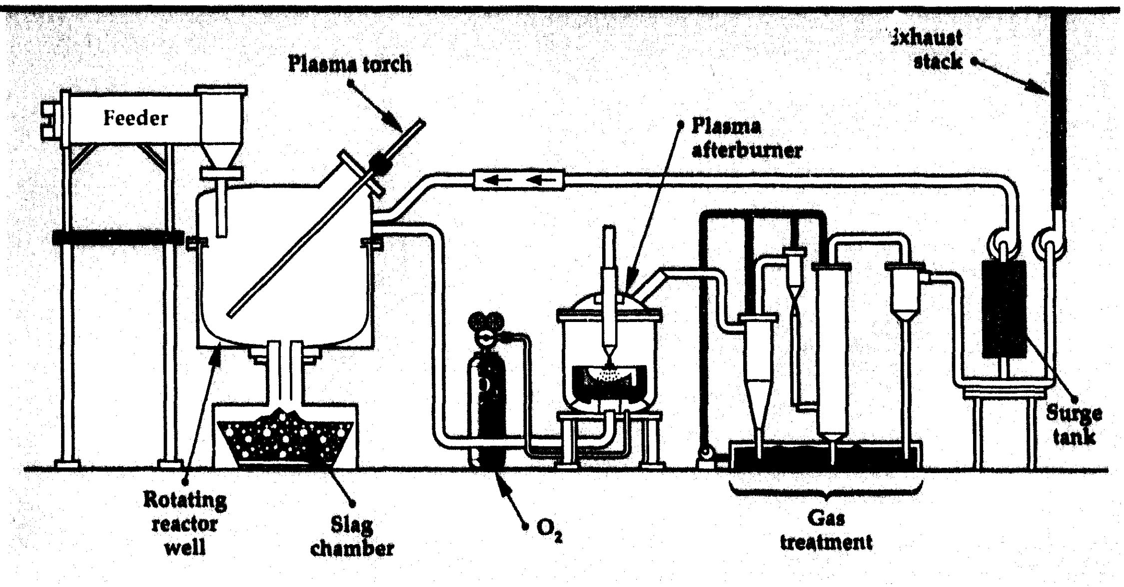

Figure 1. Conceptual hazardous-waste-treatment system.

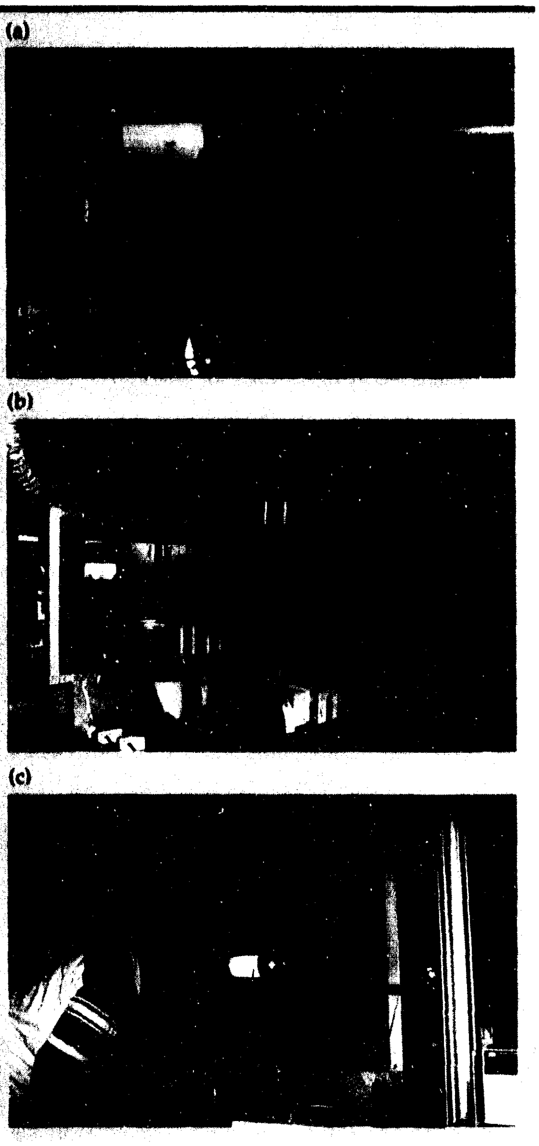

Figure 2. Photo showing (a) a close-up of the plasma torch; (b) the laboratory setting; and (c) the torch in operation.

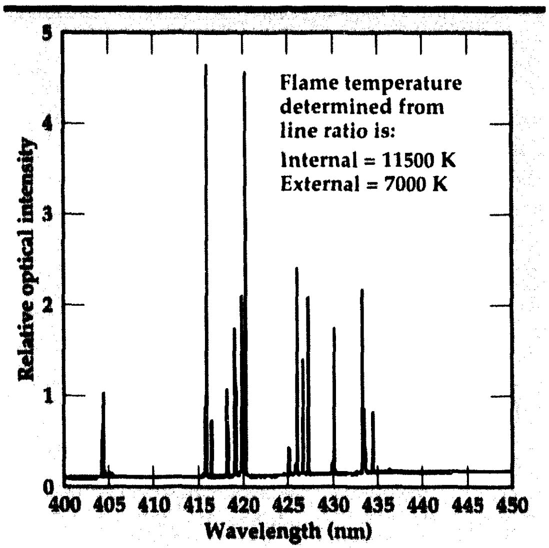

Figure 3. Plasma-torch-flame optical spectrum measured with a 2-m monochrometer.

Progress

With support from the Power Conversion Technologies (formerly Microwave and Pulsed Power) Thrust Area, a facility was established to conduct laboratory experiments, develop plasma diagnostics, and perform measurements needed to validate computer simulation results.

We made use of an existing 20-kW RF plasma torch, shown in **Fig. 2**, and assembled the ancillary equipment needed to perform routine experiments. The temperature of the torch flame was determined from optical emissions measured with a 2-m monochrometer. **Figure 3** shows the optical

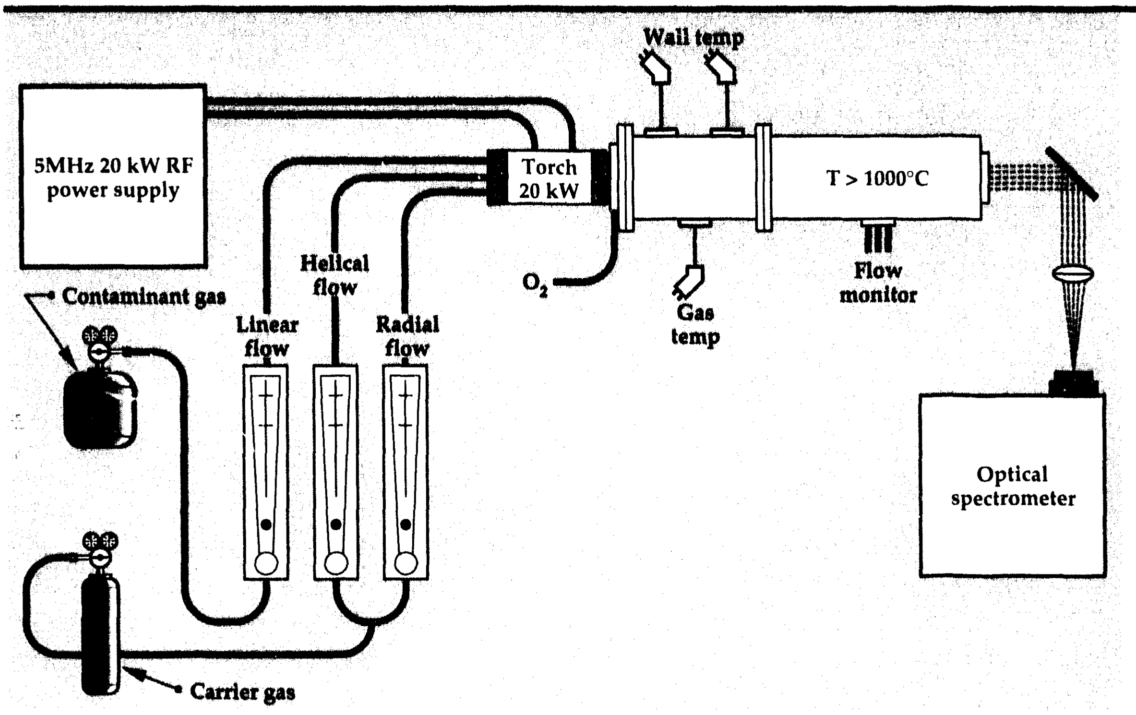

Figure 4. Schematic of laboratory plasma torch and reaction chamber instrumented for temperature measurements.

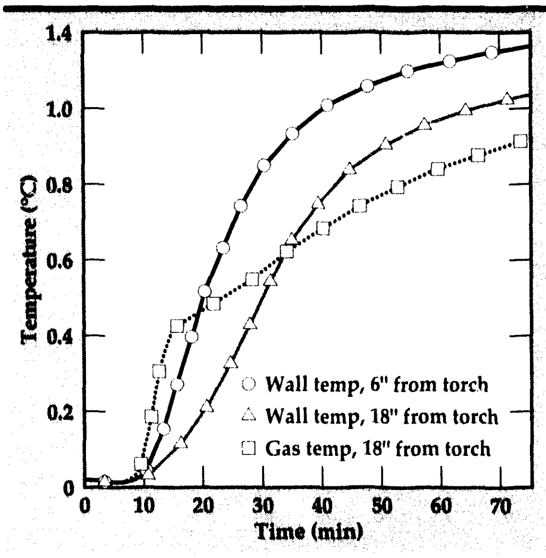

Figure 5. Measured plasma-torch, reaction-chamber temperature vs time.

spectrum from a view across the output of the torch operating with argon gas. Flame temperatures typically ranged from 7000 to 12,000K.

We constructed a high-temperature reaction chamber and used it to investigate important issues addressing how to maximize heat transfer to the effluent gas. **Figure 4** shows a schematic of the laboratory setup and reaction chamber. **Figure 5** shows typical temperature measurements at different positions during operation. We tested various schemes for directly or indirectly heating an effluent gas with the extremely hot plasma torch. We then constructed a two-section, water-cooled calorimeter to study the balance between radiative and convective heat transport. We also made high-temperature gas measurements using thermocouples with radiation shields to isolate and decouple from the colder walls. Our work so far has been without benefit of gas chemical analysis. Future experiments will make use of a gas chromatograph/mass spectrometer to do on-line characterization of process gas constituents.

Our experimental work has provided valuable insights and experience, which we are folding into the new design of an afterburner for a 50-kW RF plasma torch.We will use the afterburner to perform controlled experiments on chemical decomposition at high temperatures. We have also developed a diagnostic package to measure and characterize effluent gas from Retech's PACT-1.5 system. The module, shown in **Fig. 6**, is designed to maintain the gas at high temperature while recording gas flow and temperature and permitting gas sampling for off-line chemical analysis.

Accurate computer modeling of the hydrodynamic flow and the chemical kinetics in the afterburner is extremely important in designing the hardware geometry, and determining the oxidizer additives and the chemical composition of the effluent gas. The chemical reactions expected in the afterburner are very complex and require powerful calculational capability to model the process. Our collaborators in Lawrence Livermore National Laboratory's (LLNL's) Defense Sciences

Figure 6. Photo of off-gas sampling apparatus to be used for experiments on Retech's PACT-1.5 testbed.

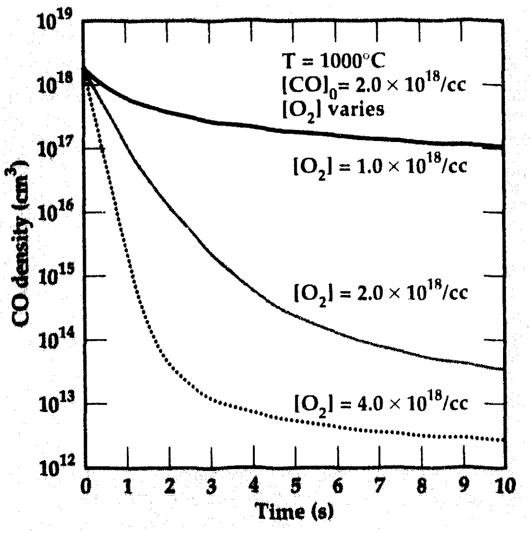

Figure 7. Preliminary calculation of oxidation rates vs O_2 concentration.

(A-Division) have modified two existing chemistry kinetics codes and have begun to apply the codes to our problems.

Figure 7 shows preliminary results from simple calculations showing oxidation rates for varying oxygen concentrations. Ongoing modeling is making use of a recently developed one-dimensional reaction code capable of handling a large number of species. Efforts are also underway to adapt a three-dimensional hydrodynamics and chemical kinetics code to operate on a massively parallel computer so the performance of a complete plasma afterburner system can be modeled. The model would include the container geometry and material, the plasma physics with turbulent flow, and the chemical reactions with varying conditions. The computer modeling allows parameter studies to be done quickly and easily at a fraction of the cost of laboratory and field experiments. Benchmarking of the codes has been done for cases where a reliable theoretical basis exists, such as in running the code until a steady-state value is calculated and then comparing that with the theoretical value.

Future Work

A Cooperative Research and Development Agreement between LLNL and Retech has been signed. LLNL's contribution to the overall effort will include (1) engineering design and systems analysis, (2) code development and computer simulations, (3) laboratory experiments and scientific validation, and (4) process diagnostics and chemical analysis. Retech will design and construct the prototype plasma-fired afterburner.

We will work with Retech to refine the conceptual system design, based on technology tradeoffs and economic analysis. We will use our diagnostic package to measure and characterize the off-gas from Retech's 1/4-scale PACT system. The data from these tests will be used to establish performance requirements for our laboratory-scale plasma afterburner. We will perform controlled laboratory experiments to validate computer simulation results and investigate the important plasma-chemistry reactions. Based on the study of the laboratory system, we will collaborate with Retech engineers on the design of a prototype plasma afterburner sized to handle the effluent gas from one of their smaller PACT systems.

---

1. "World Progress in Plasma Applications," *Proc. 2nd Int. Plasma Symposium*, EPRI Center for Materials Production, Palo Alto, California, CMP Report No. 93-4 (May 1993).
2. Retech, Inc., "Apparatus and Method for High Temperature Disposal of Hazardous Waste Materials," U.S. Patent No. 5,005,494.
3. Retech, Inc., *Plasma Centrifugal Furnace—Applications Analysis Report*, EPA/540/A5-91/007 (June 1992).

Testing of Potential Environmentally Safe Refrigerants

W. Ray Cravey and
Gary R. Mease
Defense Sciences Engineering Division
Electronics Engineering

Ruth A. Hawley-Fedder and
Jane Cupps
Condensed Matter and
Analytical Sciences Division
Chemistry and Materials Science Department

We have investigated the electrical breakdown parameters that contribute to arc-induced perfluoroisobutylene (PFIB) formation in alternative refrigerants and dielectrics. We performed an extensive experimental matrix to determine the important electrical parameters that contribute to PFIB formation. Measurements were made on the electrical parameters associated with the breakdown and on the amount of PFIB that was produced.

Introduction

With the imminent phase-out of chlorofluorocarbons (CFC's) during the next few years, both government laboratories and industry have been working to develop environmentally safe refrigerants. Although there are many proposed short-term alternatives, no refrigerant has proven to be an ideal replacement for current CFC's. One stumbling block associated with the replacements is the potential to generate toxic gases, such as perfluoroisobutylene (PFIB), when the fluids are subjected to high electrical stresses and/or breakdown.

During the first year of our project, we developed a test stand for evaluating alternative refrigerants.[1] This year, we used the test stand to determine the key electrical parameters that contribute to PFIB production. A comprehensive test plan was designed to provide the necessary data to characterize the toxic by-product production as a function of energy, pulse repetition frequency (prf), and electrode geometry.

Progress

We have completed an extensive experimental matrix using our CFC-replacement-fluid test stand. We have conducted several experiments that were designed to evaluate the formation of toxic by-products as a function of pulsed breakdown parameters. Our experiments studied the formation rate as a function of prf, energy, gap spacing, and electrode geometry.

Experiment

We conducted all our experiments using two fluids. The first fluid, which was used for most of the testing, has been evaluated in the past and was used as the reference. The second fluid, which is a new experimental fluid, showed extremely low amounts of PFIB formation in initial testing. For each experiment, we measured the gap voltage, current, and arc-voltage drop.

A diode probe was designed to measure the arc voltage as a function of time. With breakdown voltages ranging from 25 to 30 kV in many instances, it was difficult to measure arc voltages of only 30 to 60 V. We designed a simple diode probe that allowed us to use a high-bandwidth, low-voltage probe during the conduction phase of the breakdown. The probe is made up of a string of fast-recover diodes in series with a bias supply, a current-limiting resistor, and the test spark gap. The diodes are biased to conduct when the gap voltage is below a set reference level of 250 V. Once the gap voltage reaches the reference level, the diodes are forward biased, and current begins to flow through series circuit. The current was limited to a few milliamps so as not to sustain the arc or add any significant energy to the discharge. A diagram showing the diode probe and the system is shown in **Fig. 1.**

Figure 1. Schematic diagram of experimental system.

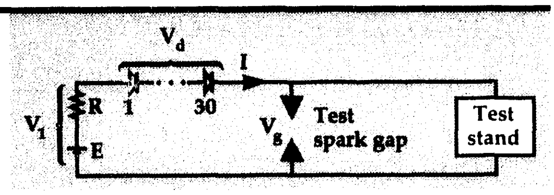

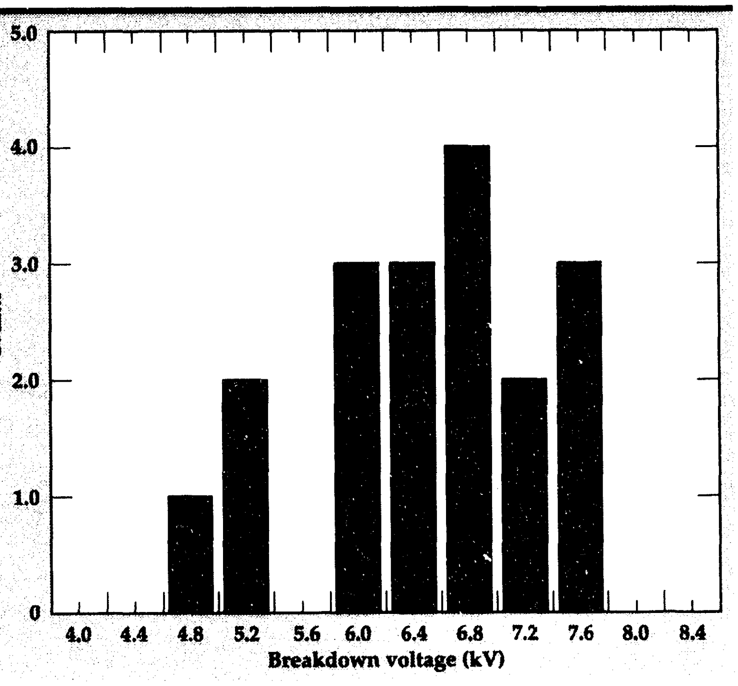

Figure 2. Histogram of breakdown voltage for 2- and 8-mil gap spacing.

Figure 3. Typical probe data set.

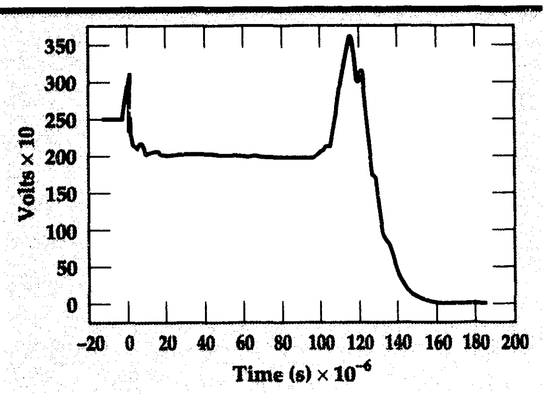

Analysis of the fluid was performed using gas chromatography with electron capture detection. A 100-μl sample of the fluid was withdrawn from the test stand and placed in a 10-ml reacti-vial containing hexane (UV grade). Prior to analysis, 1 to 2 μl of the hexane solution was withdrawn from the vial and injected into the gas chromatograph. The instrument was standardized using a gas standard prepared at Lawrence Livermore National Laboratory (LLNL) with an approximate composition of 1 ppmv PFIB in nitrogen. The exact composition of the LLNL gas standard was determined by comparison with standards that had been prepared by 3M Corporation in St. Paul, Minnesota. The LLNL gas standard was used to develop a calibration curve that related area counts to pg of PFIB injected. Two different column configurations were used for the analyses. For the majority of the analyses, a 105-m, DB-1 capillary column was used. The alternative configuration was a dual-column method using a Porapak-Q column in series after a Diethylhexyl Sebacate column, with column switching after about 1.5 min.

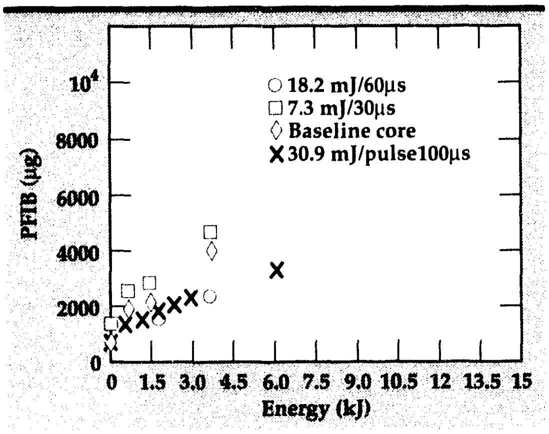

Figure 4. Plot of PFIB formation as a function of energy for three pulse widths.

Data Collection and Reduction

Breakdown voltage, conduction voltage, and conduction current were randomly measured during each experimental run. A minimum of three sets were recorded for each series of the experiment. In some cases, several sets were recorded to provide a statistical representation of the breakdown voltage and pulse energy.

The breakdown voltage was taken directly for the peak digitized-gap-voltage measurement, using a computer. A mean breakdown voltage was determined from several shots recorded for one series. A histogram of the breakdown voltage for a 2- and 8-mil gap spacing is illustrated in **Fig. 2.** Each histogram bar represents twenty samples taken at random during a 50-Hz continuous cycle.

The arc voltage was calculated from the diode probe voltage and current. A typical probe data set is shown in **Fig. 3.** Initially, the diode string is reverse-biased, and the voltage measured is simply the power supply voltage, E, which in this case is 250 V. Once the voltage across the gap collapses, the voltage drops below E, and the diode string begins to conduct. Next, the voltage decreases to the supply voltage minus the voltage drop across the resistor. At this point, the voltage across the gap is simply the measured voltage minus the voltage drop across the diode string. The voltage across the diode string was determined from the probe current and the measured current-voltage characteristics.

The energy was calculated using both digitized-current and diode-probe voltage measurements on a computer. The arc voltage and current were multiplied to determine the instantaneous power, then integrated with respect to time to yield the arc energy.

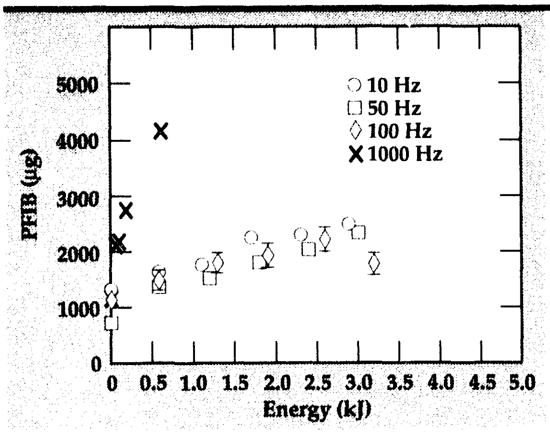

Figure 5. Plot of PFIB formation as a function of energy for four prf's.

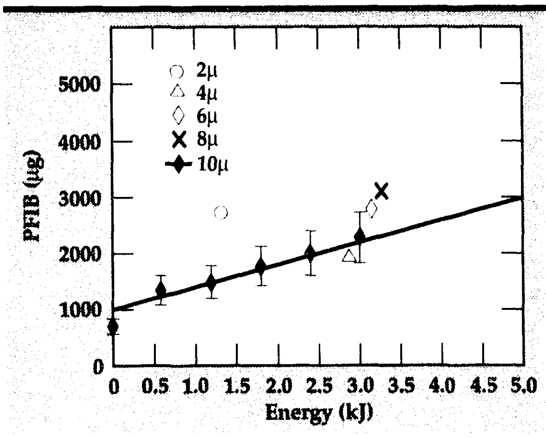

Figure 6. Plot of PFIB formation as a function of energy for varied gap spacing, using disk electrodes.

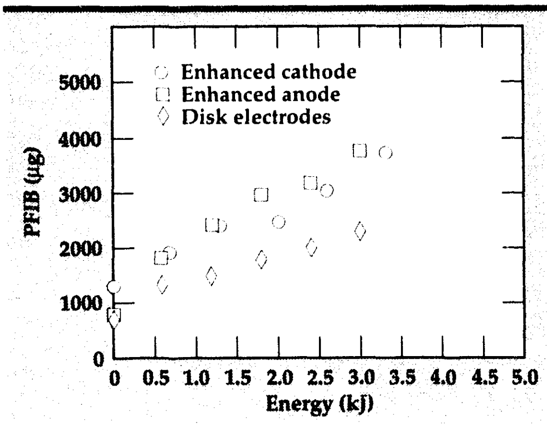

Figure 7. Plot of PFIB formation as a function of energy for three electrode geometries.

Experimental Results

The first experiment was designed to show if the PFIB formation is a thermal or non-thermal process when it is subjected to arc discharges. We varied the energy going into the discharge by varying the pulse width. A plot showing the PFIB generation as a function of energy is shown if **Fig. 4**. All runs in this experiment were done with a prf of 50 Hz. Three pulse widths were used, 30, 60, and 100 μs, with corresponding energy/pulse of 7.3, 18.2, and 30.9 mJ. A high initial baseline was recorded for the 30-μs case. The baseline was corrected and was also plotted. In general, there is clearly a linear trend with respect to energy for all three cases. The 60- and 100-μs pulses follow the same trend within a few percent, with the 30-μs series out to 20%.

We also measured the dependence of PFIB formation on prf. Four data sets were taken, for 10, 50, 100, and 1000 Hz. A 100-μs pulse width was used for all sets excluding the 1000-Hz test. As seen in **Fig. 5**, except for the 1000-Hz prf experiment, there is no dependency on prf.

A comparison of PFIB formation as a function of energy and gap spacing is illustrated in **Fig. 6**. Disk electrodes were used for these experiments; 100,000 pulses were discharged through the gap at a prf of 50 Hz and a pulse width of 100 μs. A sample was taken at the end of each run. The gap was increased from 2 to 10 mils, in increments of two mils. As seen in **Fig. 6**, there is little dependence on the gap spacing when the data is corrected for arc energy. This supports the theory that the formation of PFIB is energy-dependent and is a thermal dissociation process rather than a nonthermal one.

To determine the dependence on electrode shape, we tested three electrode geometries: disk, enhanced anode, and enhanced cathode. As illustrated in **Fig. 7**, both the enhanced anode and enhanced cathode geometries appear to follow linear trends with respect to energy. It is difficult to draw any conclusion based on electrode shape. The enhanced geometries were tested at similar times, and this may account for their close relationship; the disk electrode data was analyzed earlier. Throughout the experiments, we noticed a variation in PFIB count from month to month.

Future Work

Last year, we identified companies in the refrigeration industry as possible partners for the facility and capabilities we have built up over the last year. In the majority of industrial refrigeration systems, the refrigerant is circulated through the compressor motor, for cooling and lubrication. In large systems, the voltage level can exceed 1000 V, leading to the occasional and often catastrophic electrical breakdown of the fluid. The unique capabilities for testing for arc-induced toxic by-products were presented to representatives of the Air Conditioning and Refrigeration Technology Institute (ARTI) technical committee. It was agreed that LLNL possesses unique capabilities to test for such arc-induced toxic by-products. We are now working on a contract with ARTI to test three proposed refrigerant/oil blends to be used as short-term replacements in industrial chillers.

---

1. W.R. Cravey, W.R. Luedtka, R.A. Hawley-Fedder, and L. Foiles, "Testing of CFC Replacement Fluids for Arc-Induced Toxic By-Products," *Engineering Research, Development, and Technology*, Lawrence Livermore National Laboratory, Livermore, California, UCRL-53868-92, 7-19 (1993).

Compact Pulsed Power Supplies

W. Ray Cravey and
Gary R. Mease
Defense Sciences Engineering Division
Electronics Engineering

Hugh C. Kirbie,
Steven A. Hawkins, and
Craig W. Ollis
Laser Engineering Division
Electronics Engineering

We have completed the first stage of development of a compact, modulator module for driving pulsed plasma loads. The module is capable of generating 800-V, 150-ns pulses at a pulse repetition frequency (prf) of 10-kHz continuous operation. Voltages as high as 50 kV, with pulse widths on the order of 200 ns and prf's of > 50 kHz, can be achieved by stacking additional modules. At Lawrence Livermore National Laboratory, we have developed a non-linear magnetic core model for circuit design and analysis of the Heavy Ion Fusion Project's MOSFET- (metal oxide semiconductor field effect transistor) switched modulator. The PSpice model was developed to simulate the performance of the modulators and help predict the behavior of the Metglas induction core. We have assessed the usefulness of MOSFET-switched modulators to drive pulsed plasma loads, and found a natural match between the pulse agility of the modulators and the dynamic load characteristics of the reactor cell. We modified the modulator to produce 12-kV, 10-A, 200-ns pulses, with the addition of a 12-turn output transformer. Initial tests demonstrate the generation of corona discharges at prf's as high as 200 kHz.

Introduction

There have been dramatic improvements in the power handling capabilities of high-power MOSFET[1] switches over the past year, with no signs of slowing down.[2] The Internal Confinement Fusion (ICF) program's Heavy Ion Fusion Project (HIF) at Lawrence Livermore National Laboratory (LLNL) has been developing MOSFET-switched modulator technology for applications in recirculating accelerators.[3] Over the last year, we have improved the current handling capabilities of the switches by a factor of ten.

In November 1993, a partnership between the Microwave and Pulsed Power Thrust Area and the HIF Project was formed. We have been working to adapt FET switching technology to pulsed plasma processing. In this application, a 20-kV pulse is used to drive a reactor cell that comprises a 6-mil wire in a 1-in. tube. The applied voltage on the wire initiates a non-thermal plasma in an exhaust gas. The non-thermal plasma initiates a chemical reaction in the gas and forms OH radicals, which in turn remove NO_x from the exhaust gas.

While the HIF modulators have broken new ground in accelerator cell technology, the switch arrangement that is used is not the best for a compact, inexpensive, pulsed-plasma power supply. We designed a new switch array that mitigates the need for the costly switch isolation used in the HIF machines. The new arrangement is based on the voltage adder concept used for accelerator cells for many years. The basic

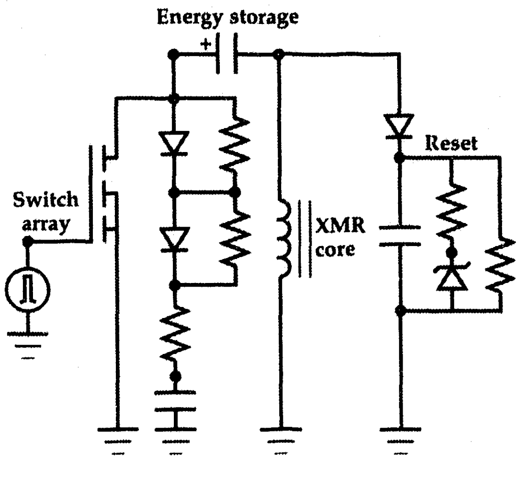

***Figure 1.* Schematic diagram of basic MOSFET-switched modulator.**

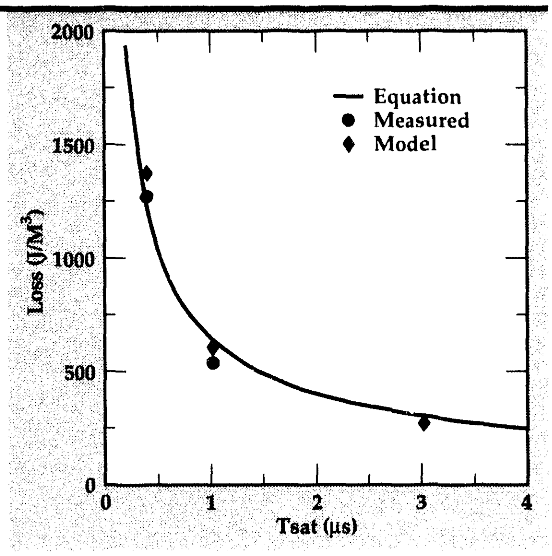

Figure 2. Plot of small-core loss data showing good agreement between measured data, an empirical equation, and PSpice simulation for 2605S-3A Metglas core.

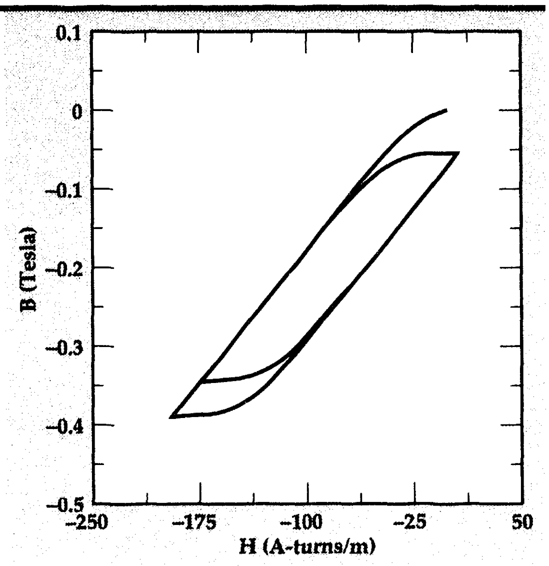

Figure 3. PSpice simulation of the B-H loop for a four-pulse burst.

circuit architecture is identical to the original HIF modulator concept. The key to maintaining a compact size in this arrangement is to integrate the energy storage, switching, and induction cells into a single package.

Progress

Heavy-ion beams are under study as a future drive source for inertial fusion. LLNL and Lawrence Berkeley Laboratory have proposed a circular accelerator known as a recirculator. A recirculator requires the pulsed power supplies to operate from 50 to 100 kHz prf, and continuously variable pulse widths ranging from 400 ns to 1 μs. The HIF group has selected the circuit architecture shown in **Fig. 1**, which generates voltage pulses for acceleration and reset with a single on/off command.

Modeling

One major concern with the modulator's performance was the ability for core rest between pulses. We developed a non-linear, Metglas core model to help understand the behavior of the core trajectory. The core model in PSpice, which uses the Jiles–Atherton model for analyzing non-linear cores,[4,5] was adjusted to simulate the Metglas core by adjusting several parameters. The model was verified by comparing the calculated core losses from the hysteresis with the empirical core loss equation from Kuenning[6] and small core test data taken earlier (see **Fig. 2**). A PSpice simulation of the B-H trajectory for a four-pulse burst and the measured B-H trajectory from the modulator are shown in **Fig. 3**.

Pulsed Plasma Processing

We have constructed a small, 5-kV induction accelerator cell powered by a MOSFET switch array. The Metglas induction cell is surrounded by 24 FET's and several capacitor banks to yield a unique integration of source and cell. The cell prf, pulse width, pulse amplitude, and core reset interval were all controlled directly. The modulator was adapted to generate a non-thermal plasma discharge at low energies and high prf's. The output voltage of the modulator was increased with the use of an output step-up transformer. Twelve turns were wrapped around the modulator core as shown in **Fig. 4**. The output from the modulator was connected to a wire loop-plane

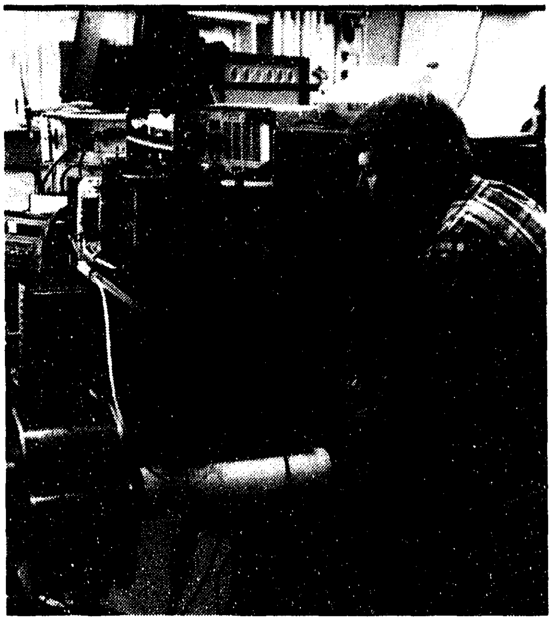

Figure 4. Modulator core in HIF Project experiment.

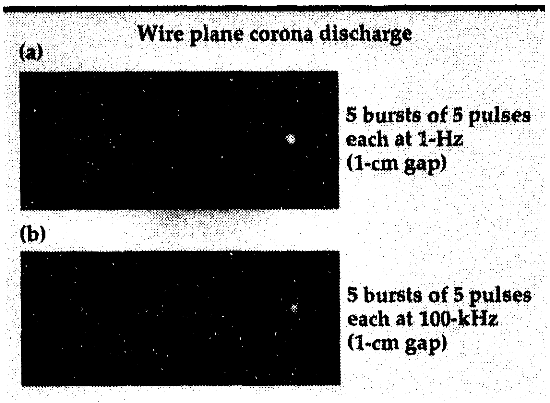

Figure 5. Photos of wire-plane corona discharges (1-cm gap) of five bursts of five pulses each, at (a) 1 Hz and (b) 100 kHz.

electrode geometry. Peak voltages of 12 kV with rise times of 50 ns and pulse widths of 200 ns were generated. **Figure 5** shows two open-shutter photographs of the corona discharge taken for a 1-Hz burst and a 100-kHz burst of 25 pulses each.

Significant progress has been made towards the next generation of compact, pulse agile, high reliability modulators. We have designed and are currently testing an 800-V MOSFET module. The basic circuit architecture is the same as that shown in **Fig. 1**. The 1-×-6 module consists of 12 APT1001RBN switches in a parallel configuration. The switches are arranged around a small Metglas core and connected in series to eight Roderstien capacitors (**Fig. 6**). Each module was designed to handle a minimum of 800 V and 500 A. Each module will be able to switch 500 W average power. We plan to stack 25 modules in series to obtain a 20-kV, 500-A pulse with an average power of 12.5 kW.

We have tested a single module to 800 V and 200 A. We have operated the module at 10 kHz continuous operation, and as high as 300 kHz at reduced voltages. A 5-pulse burst at 300 kHz is shown in **Fig. 7**. The final design will be arranged around a small Metglas core designed for minimum losses and size. We will charge the system with a floating power supply arrangement that will maintain an average charge across each capacitor, but will not require a charging resistor or other power-dissipating elements. Energy that is recovered from the reset circuit will be reintroduced into the charging circuit, further increasing efficiency.

Figure 6. Photo of 1-×-6 module, with switches arranged around a small Metglas core.

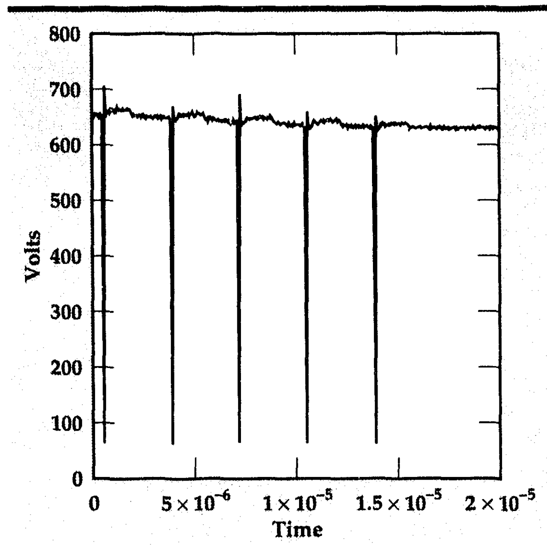

Figure 7. Five-pulse burst from 1-×-6 module.

Future Work

Our next goal is to build several MOSFET modules and demonstrate their performance when stacked in a voltage-adder configuration. We are working to develop a 25-kV, 10-kW modulator to be used in support of the pulsed-plasma processing laboratory at LLNL. A computer-aided-design solid model of the second 1-×-6 prototype is illustrated in **Fig. 8**. Its overall size has been greatly reduced by replacing the larger Roderstien capacitors with new, high-density, surface-mount technology.

Beginning in FY-94, HIF will build a small-scale recirculator for evaluating the feasibility of this type of accelerator as a future driver for inertial fusion. The design concept behind the 1-×-6 modulator will have direct applicability to the induction-accelerating cell modulators

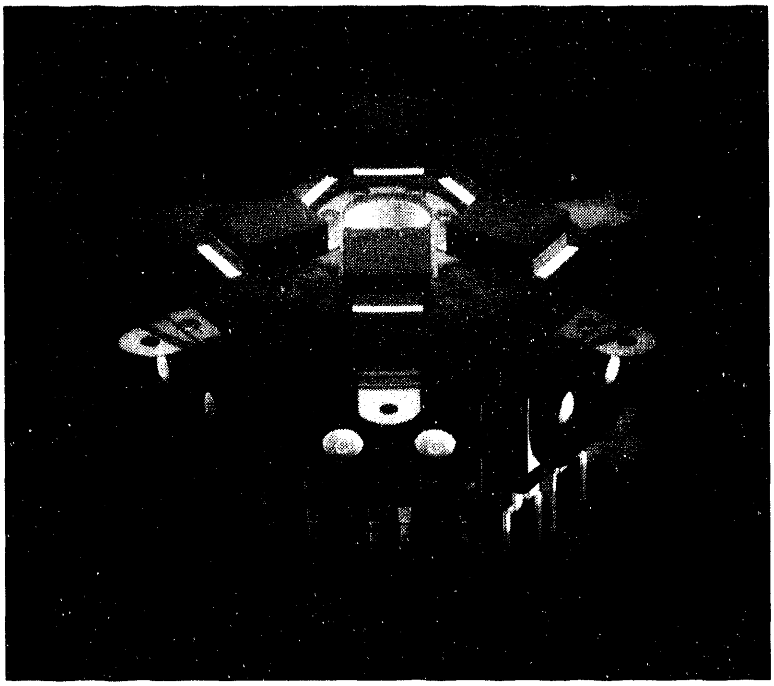

Figure 8. CAD solid model of our latest 1-×-6 prototype module.

needed for this small ion-recirculator experiment. The induction accelerating cells require a 600-V drive with variable pulse width capability and a maximum prf capability in excess of 100 kHz. Implementation of this technology into a recirculating induction-accelerator cell will require repackaging and resizing the components, but the basic concept has been proven. This project is the beginning of the development of a compact source technology that may find use in a wide variety of applications.

---

1. C.W. Ollis, W.R. Cravey, S.A. Hawkins, H.C. Kirbie, and M.A. Newton, "SPICE Modeling of a FET-Switched Induction Accelerator Cell," *Proc. 19th IEEE Int. Pulsed Power Conf.* (Albuquerque, New Mexico), (June 1993).
2. B.J. Baliga, *IEEE Trans. Electron Devices* **38**(7), (July 1991).
3. H.C. Kirbie, W.R. Cravey, S.A. Hawkins, M.A. Newton, and C.W. Ollis, "A FET-Switched Induction Accelerator Cell," *Proc. 19th IEEE Int. Pulsed Power Conf.* (Albuquerque, New Mexico), (June 1993).
4. PSpice is a registered trademark of MicroSim Corp., 23175 La Cadena Drive, Laguna Hills, California 92653.
5. D.C. Jiles and D.L. Atherton, *J. Magnetic Materials* **61**, (1968).
6. R. Kuenning, R. Stone, M. Welborn, G. Caporoso, T. McCarville, K. Avery, and K. Whitham, *Long Pulse Accelerator*, Lawrence Livermore National Laboratory, Livermore, California, Beam Research Memo RM-89-43 (October 1989).

Advanced Ground-Penetrating, Imaging Radar for Bridge Inspection

John P. Warhus,
Scott D. Nelson, and
Jeffrey E. Mast
Defense Sciences Engineering Division
Electronics Engineering

Erik M. Johansson
Laser Engineering Division
Electronics Engineering

During FY-93, we continued with development and experimental evaluation of components and system concepts aimed at improving ground-penetrating imaging radar (GPIR) for nondestructive evaluation of bridge decks and other high-value concrete structures. We developed and implemented a laboratory test bed, including features to facilitate component testing, antenna system configuration evaluation, and collection of experimental data from realistic test objects. In addition, we developed pulse generators and antennas for evaluation and use in antenna configuration studies. This project was part of a cooperative effort with the Computational Electronics and Electromagnetics and Remote Imaging and Signal Engineering Thrust Areas, which contributed signal- and image-processing algorithm and software development and modeling support.

Introduction

Inspection of high-value structures, like bridges, is an application of Ground Penetrating Radar (GPR) technology, which is growing in importance. In a typical inspection application, inspectors use GPR to locate structural components, like reinforcing bars embedded in concrete, to avoid weakening the structure while collecting core samples for detailed inspection. Advanced GPR, integrated with imaging technologies for use as a nondestructive evaluation (NDE) tool, can

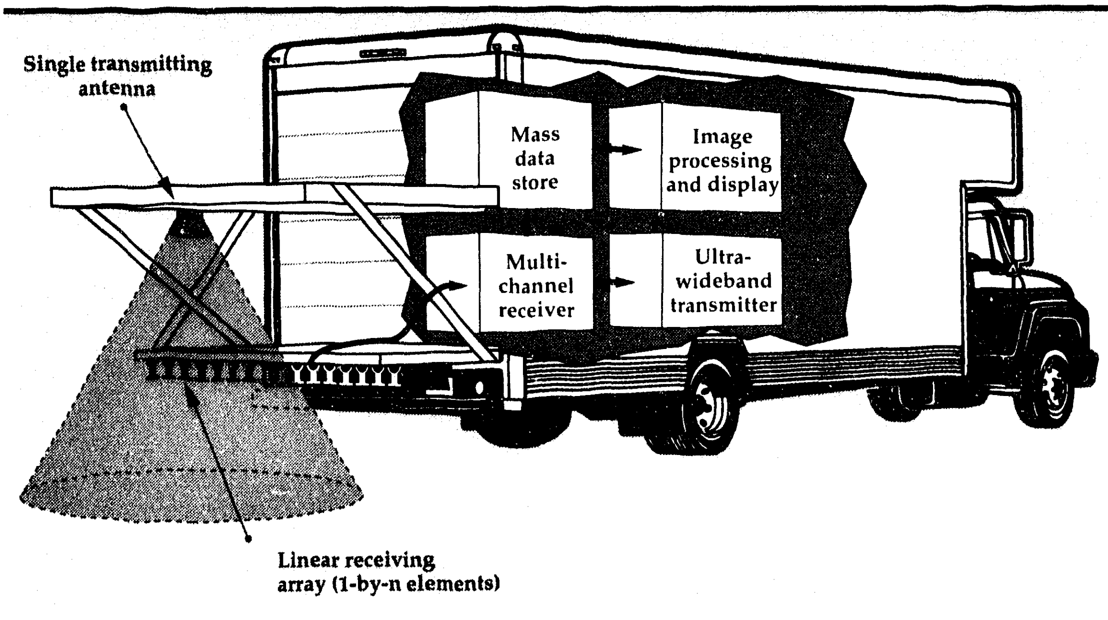

Figure 1. GPIR bridge-deck-inspection concept.

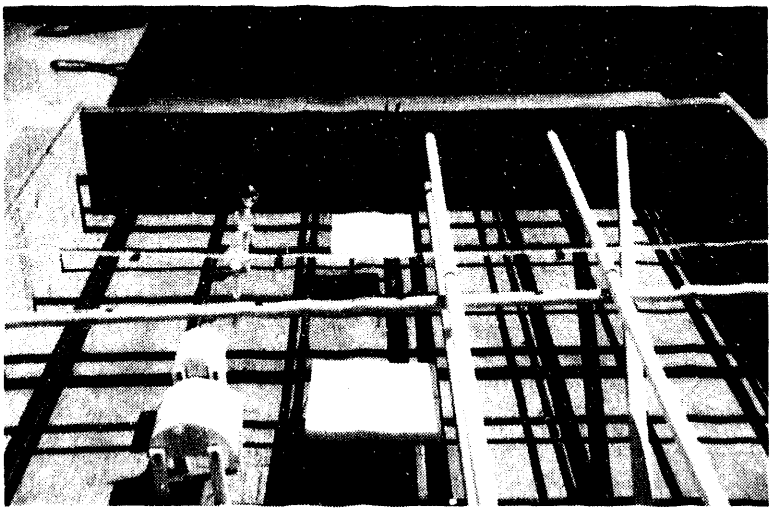

Figure 2. Test bed prior to concrete pour.

provide the capability to quickly locate and characterize construction flaws and wear- or age-induced damage in these structures without the need for destructive techniques like coring. The bridge deck and its wearing surface are the most vulnerable parts of a bridge to damage from routine service, and they are particularly well suited for inspection using a vehicle-mounted inspection system.

More than 40% of the 578,000 highway bridges in the U.S. are either structurally deficient or functionally obsolete.[1] These conditions can limit bridge utility and, if they are not properly monitored and maintained, pose a safety threat to bridge users. An advanced GPR-based inspection system has the potential of addressing critical national and international needs for reliable, cost-effective NDE of bridges.

In an advanced bridge deck inspection system, like the one shown in **Fig. 1**, a mobile Ground Penetrating Imaging Radar (GPIR) gathers data for high-resolution image reconstruction of embedded defects and features. High-quality images, processed in centrally located computing centers, allow visualization of internal structure, permitting evaluation of deck conditions from data that are usually obtainable only by destructive means.

In this project, our goal is to demonstrate enhanced GPIR performance by the application of advanced hardware and software. Among the desired enhancements are: traffic-lane-wide coverage enabled by new transmitters, antennas and large-aperture receiving arrays; and high-speed data acquisition and accompanying inspection-vehicle speed made feasible by multiple receiver channels using state-of-the-art data transmission and storage equipment. In addition, downrange and cross-range resolutions can be improved by increasing transmitted pulse bandwidth and applying synthetic-aperture radar data-processing techniques.

Progress

During FY-93, our efforts were aimed at experimental validation of system-level concepts developed in FY-92.[2] We established a laboratory enabling test and evaluation of GPIR components and system configurations, and development of experimental data to use for validation of computer simulations (performed through the Computational Electronics and Electromagnetics Thrust Area) and evaluation of imaging- and signal-processing algorithms and techniques (developed by the Remote Imaging and Signal Engineering Thrust Area). We designed, developed, and characterized ultra-wide-bandwidth (UWB) antennas and pulse generators and used them in the experimental test bed to generate data for image processing. We also tested available

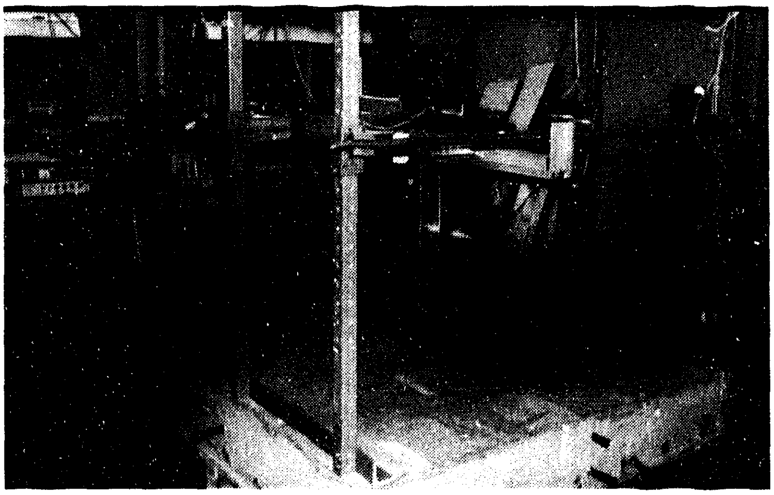

Figure 3. The GPIR laboratory showing a typical experimental setup.

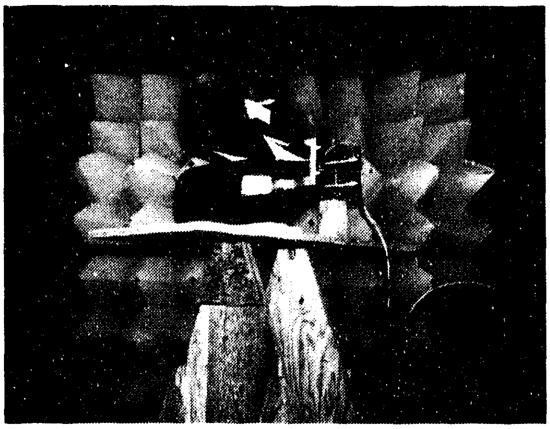

Figure 4. Slotline-bowtie antenna setup in anechoic chamber.

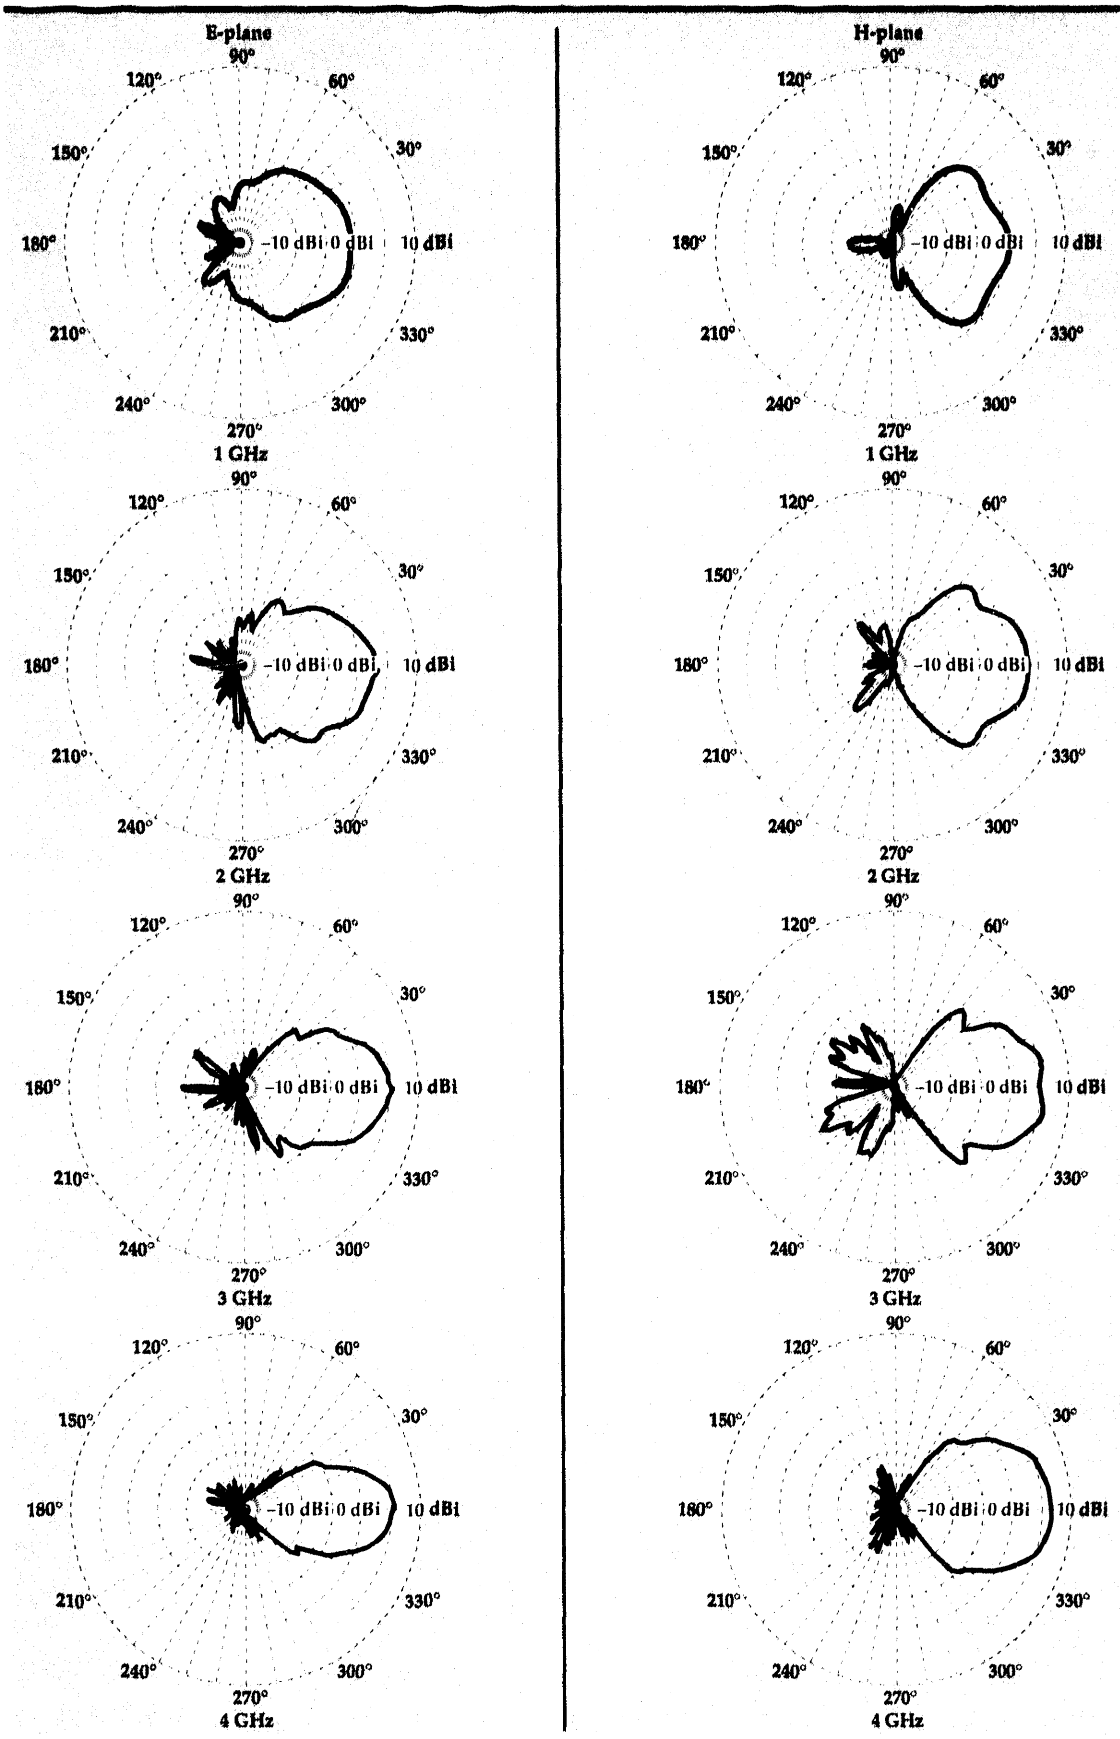

Figure 5. Antenna patterns for slotline-bowtie antenna.

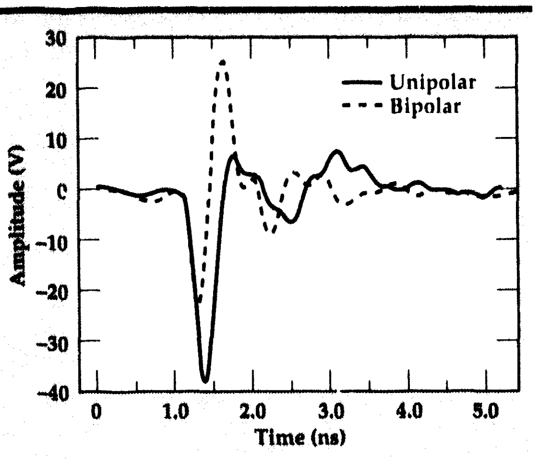

Figure 6. Output pulse from step-recovery diode impulse generator.

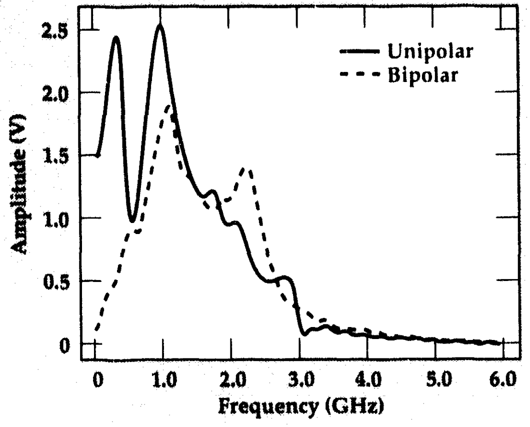

Figure 7. Spectrum of step-recovery diode impulse generator.

off-the-shelf antennas and pulse generators to provide benchmarks against which we could evaluate our own designs, and to determine their viability as possible components in an advanced GPIR.

GPIR Laboratory and Test Bed

The GPIR laboratory is equipped with both standard and special test equipment needed to support UWB radar experimentation and data acquisition. Standard test equipment includes oscilloscopes, impulse generators, UWB microwave amplifiers, power supplies, motor controllers, and meters. Specialized equipment, developed under the ongoing program, includes impulse generators and transmitting and receiving antennas designed specifically for the inspection application. Data-acquisition and control hardware used in experiments are under the control of a dedicated instrument-control computer.

A test bed, designed to provide a realistic test object during early development, is set up in the laboratory. Shown in **Fig. 2** before concrete was poured, the test bed is a concrete slab (2 m-x-2 m-x-0.3 m thick) containing reinforcing bars (rebars), both fixed and removable, and other objects designed to provide a means for evaluating data acquisition and imaging performance. Concrete poured in the test bed was prepared and handled in accordance with specifications used by the California Department of Transportation for concrete bridge decks. The test bed is equipped with motor-driven slide mechanisms designed to move antennas accurately and repeatably over the slab, simulating both the motion of an inspection vehicle over a roadbed or bridge deck and a fixed linear array of receiving antennas mounted on the vehicle. A typical experimental setup in the GPIR laboratory is shown in **Fig. 3**.

Antennas and Pulse Generators

We designed and characterized UWB antennas for transmitting and for use in receiving arrays. A slotline-bowtie antenna^3 designed for use as a transmitting antenna is shown in **Fig. 4** set up in an anechoic chamber for antenna-pattern measurements. This antenna was designed to provide good antenna-pattern control and impulse response over a relatively wide band. Characterization measurements were made over a frequency range from 500 MHz to 8 GHz. The measurements showed that the antenna provided usable gain and a well-behaved radiation pattern over the full range. Antenna patterns measured for this antenna from 1 GHz to 4 GHz are shown in **Fig. 5**.

We also assembled and tested UWB pulse generators using step-recovery diodes, to evaluate their potential as impulse sources for a GPIR inspection system. Key requirements for the generator include peak power of 100 W to a few kilowatts, pulse widths of 100 to 300 ps, and pulse repetition frequencies (PRF) up to 5 MHz. In the impulse generators we tested, avalanche transistor pulse generators were used to drive series stacks of step-recovery diodes.

Output pulses and spectra for both unipolar and bipolar impulse generators are shown in **Figs. 6** and **7**. The unipolar pulse shown in the figures has peak power of ~ 30 W, usable bandwidth of greater than 3 GHz, and was operated at a PRF > 1 MHz. The PRF was limited by the ability of the avalanche transistor to operate at high repetition rates. Further work is needed to complete the evaluation of this technique for GPIR impulse generation.

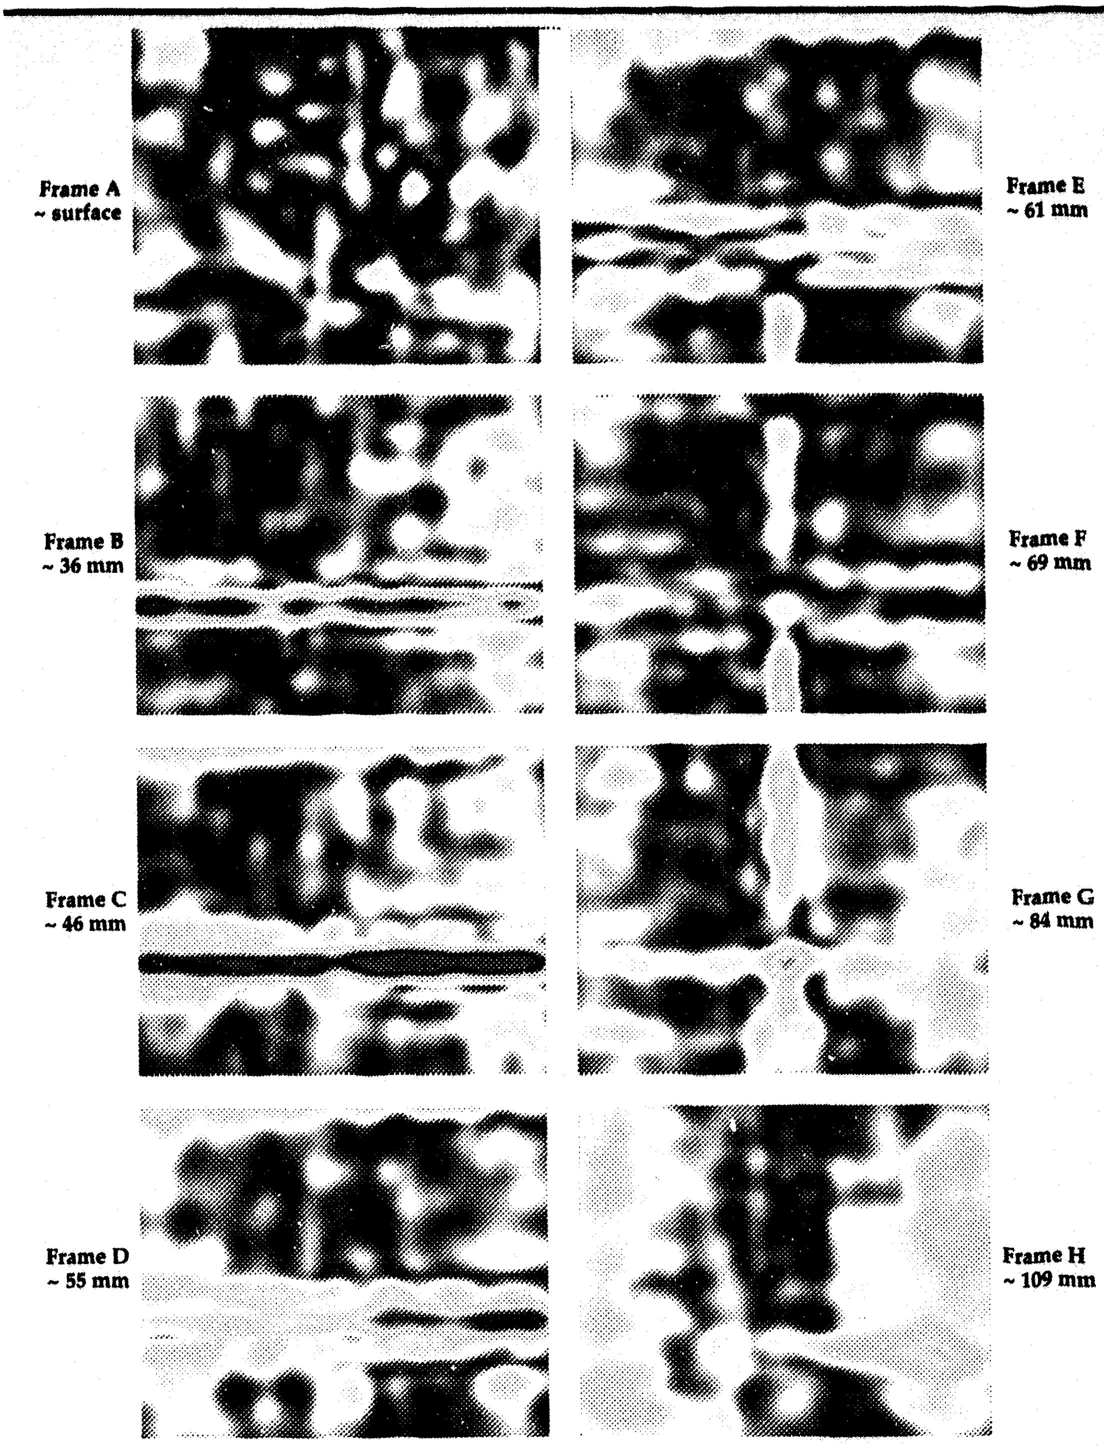

Figure 8. Images reconstructed from experimental data.

Experimental Results

In experiments conducted in FY-93, we collected data and reconstructed images of features embedded within a concrete test bed. In those experiments, we collected data using monostatic antennas driven by a low-power impulse generator (~ 1 W, peak). Data were acquired with the receiving antenna configured to receive the cross-polarized returns scattered from targets and clutter in the illuminated volume. Data recording used a sampling oscilloscope with a UWB microwave front-end preamplifier. Antennas were mounted 50 mm above the concrete surface, and moved over the test bed while data were collected over a rectangular sampling grid at 12.7-mm spatial intervals.

Figure 8 shows a sequence of two-dimensional images that are planar slices from a three-dimensional image reconstruction; the slices are parallel to the surface of the concrete. The depth calculated in the reconstruction is indicated for each frame. The sequence starts near the concrete surface (Frame A) and progresses through the volume.

Within the imaged volume are four cylindrical voids; these voids are in locations where removable rebars have been removed. The first two voids are parallel with each other, with top surfaces about 63 mm below and nearly parallel with the concrete surface. One of the voids, at the top of the image, is only partially visible in Frames C through E; the other is visible in Frames B through E.

The third void is beneath (at a depth ~ 94 mm below the surface) and perpendicular to the first two, and it slopes slightly deeper into the concrete at the end near the top of the image. In Frame D, near the bottom of the frame, the third void first appears, and it is visible in Frames D through G. The difference in depth, within the image, from one end of the third void to the other, is approximately 18 mm; its sloping orientation is evident in the image sequence.

A portion of the fourth void is visible in the lower right corner of Frames G and H of the sequence. This void starts at a depth ~ 117 mm and slopes deeper into the concrete and slightly toward the top of the frame. Because of the low output power of the impulse generator used in this experiment, scattered signal power was too low to permit detection of objects deeper than ~ 125 mm in the concrete.

Future Work

Our continuing technical efforts are aimed toward advancing radar data-acquisition hardware to improve our ability to recover data from deeper in concrete, while maintaining UWB. Improvements here will enable high-resolution imaging of objects that are buried more deeply in the material. Our efforts will focus on development and testing of higher power impulse generators and receiver front ends with improved noise performance and temporal gain control. We will continue with evaluation of antenna configurations that are better suited to the bridge-deck-inspection application.

Another important goal for the next fiscal year is development of a scaled, demonstration, prototype bridge-deck-inspection system for an external funding sponsor. That effort will require collaboration with industry to be successful.

Acknowledgements

We wish to thank D. Goodman, Remote Imaging and Sensor Engineering Thrust Area Leader, and J. DeFord, Computational Electronics and Electromagnetics Thrust Area Leader, for their support in supplying resources needed to develop and evaluate imaging techniques and perform electromagnetic modeling for this project.

---

1. *Our Nation's Highways: Selected Facts and Figures*, U.S. Department of Transportation, Federal Highway Administration, Publ. No. FHWA-PL-90-024.
2. J.P. Warhus, J.M. Hernandez, S.D. Nelson, E.M. Johansson, and H. Lee, "Ground Penetrating, Imaging Radar for Bridge Inspection," *Engineering Research, Development, and Technology*, Lawrence Livermore National Laboratory, Livermore, California, UCRL-53868-92 (March 1993).
3. A.K.Y. Lai, A.L. Sinopoli, and W.D. Burnside, *IEEE Trans on Antennas and Propagation* **AP-40**, 755 (July 1992).
4. S.D. Nelson, "Electromagnetic Modeling for Target-Rich Embedded Environments," *Engineering Research, Development, and Technology*, Lawrence Livermore National Laboratory, Livermore, California, UCRL-53868-93 (March 1994).
5. E.M. Johansson and J.E. Mast, "Imaging Algorithms for Synthetic Aperture Ultra-Wideband Radar," *Engineering Research, Development, and Technology*, Lawrence Livermore National Laboratory, Livermore, California, UCRL-53868-93 (March 1994).

High-Power, Electron Beam-Induced Switching In Diamond

Raymond D. Scarpetti, Jr.,
W. Wayne Hofer, and
Don R. Kania
Defense Sciences Engineering Division
Electronics Engineering

K.H. Schoenbach,
R.P. Joshi, C. Molina, and
R.P. Brinkmann
Old Dominion University
Norfolk, Virginia

We are developing a high-voltage, high-average-power, electron beam-controlled diamond switch that could significantly impact high power solid-state electronics in industrial and defense applications. An electron beam-controlled, thin film diamond could switch, with high efficiency, well over 100 kW average power at MHz frequencies greater than 5 kV. This performance is due to the excellent thermal and electronic properties of diamond, the high efficiency achieved with electron beam control, and the demonstrated effectiveness of microchannel cooling.

Our electron beam penetration-depth measurements agree with our Monte Carlo calculations. We have not observed electron beam damage in diamond for beam energies up to 150 keV. This report describes our experimental and calculational results and research objectives.

Introduction

Diamond has many properties that make it ideal for high-average-power, high-voltage, high-frequency solid-state switching. Diamond has a dielectric strength of 1–10 MV/cm[1], which is more than an order of magnitude greater than that of Si or GaAs, and four times greater than that of SiC. Diamond also has the highest thermal conductivity of any solid at room temperature, up to 20 W/cm-K, which is over ten times that of Si and GaAs and five times that of SiC. Diamond's high electron and hole mobility (> 1000 cm^2/Vs) allow for operating at MHz frequencies. The 5.47eV band gap results in lower dark current, higher temperature operation, and hardness to radiation as compared to Si and GaAs semiconductor devices. However, this wide band gap makes it difficult to ionize using lasers. UV lasers are inefficient and expensive, while lower frequency lasers couple poorly to the high absorption length of diamond. One very efficient method of generating charge carriers in wide band gap materials such as diamond is an electron beam.

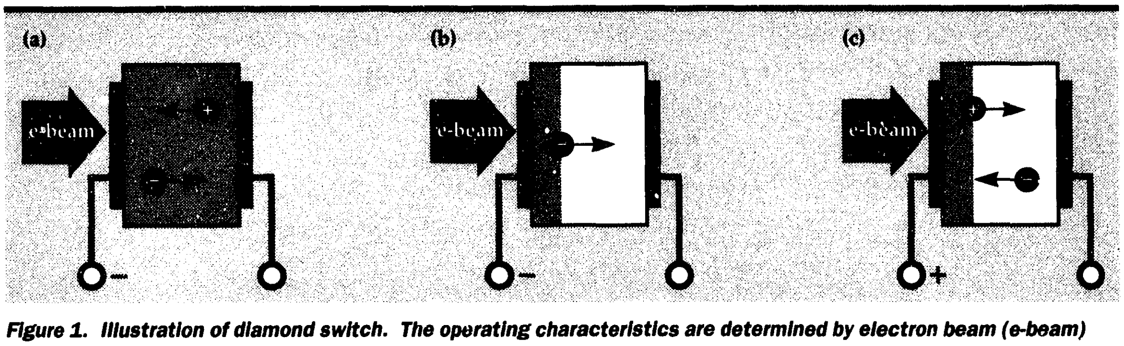

Figure 1. Illustration of diamond switch. The operating characteristics are determined by electron beam (e-beam) penetration depth and type of contact: (a) blocking contacts, diamond fully ionized; (b) electron-injecting contacts, diamond partially ionized; and (c) electron-injecting contacts, different bias.

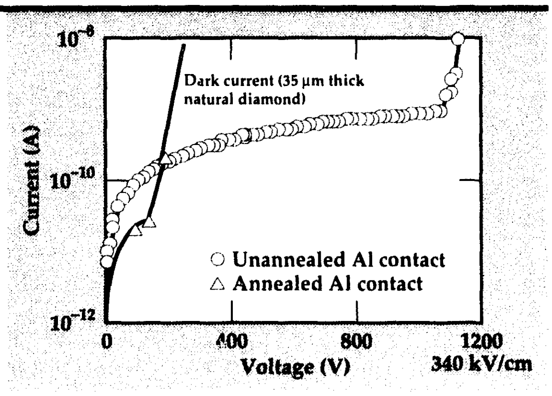

Figure 2. Experimental results, showing that un-annealed aluminum acts as a blocking contact, significantly increasing the switch holdoff voltage.

Advantages of electron beam activation are broad energy ranges and easily tailored energy waveforms, high wall plug efficiencies, lower costs, and higher sustained intensities than lasers. Until now, interest in diamond devices has been limited because of their high cost and the limited supply of high purity material. However, recent advances in processing technology and the successful growth of high quality artificial diamond using vapor phase deposition, are resulting in purer materials with less defects and, potentially, at greatly reduced cost. The result is a viable material for high-power, solid-state switching in support of the power electronics industry as well as for military applications. Lawrence Livermore National Laboratory (LLNL) and Old Dominion University have begun to address electron penetration range in diamond, bulk breakdown voltage, the role of crystal defects, the effects of contacts on charge injection and dark current,[2] and the development of a theoretical and modeling base. All of these are necessary to realize the potential of electron beam-controlled diamond devices as high-average-power solid-state switches.

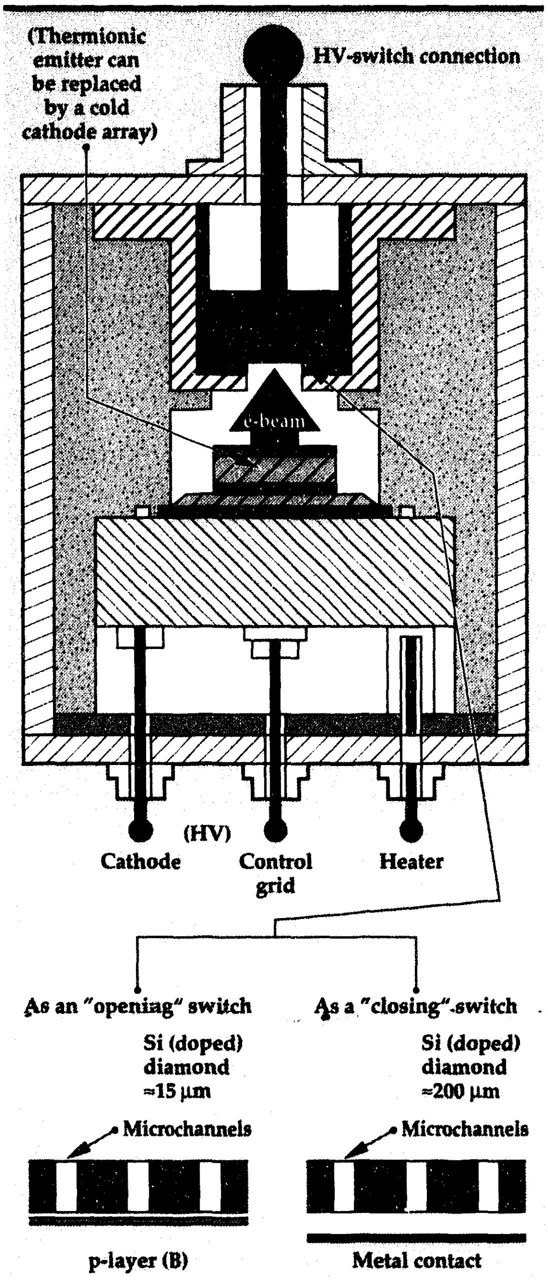

Figure 3. Illustration of e-beam-controlled diamond used as an 'opening' or 'closing' switch.

Progress

Switch Operation

There are three modes of operation based on incident electron beam energy and type of contacts used, as illustrated in **Figs. 1a, 1b,** and **1c**. In the first mode, blocking contacts are used, and the electron range equals the switch thickness. In this configuration (see **Fig. 1a**), the electron beam provides the only source of charge carriers through electron beam ionization. Since the electron-hole pairs are generated simultaneously throughout the entire region, the turn-on time is determined by the trap filling time in the diamond, which is related to the electron beam current and can be extremely fast. Blocking contacts are used to prevent charge-carrier injection through the contacts and to allow for the full use of the dielectric strength of the diamond. Experimental results (**Fig. 2**) show unannealed aluminum acts as a good blocking contact with applied electric fields of 340 kV/cm across the diamond before the onset of dark current, while the use of annealed aluminum contacts results in high dark currents at electric fields of 50 kV/cm.

The current gain of such a device is determined by the ratio of primary electron energy and the mean ionization energy for diamond, which is approximately 15 eV. Therefore, a single 150-keV electron will generate 10,000 electron-hole pairs. Upon termination of the electron beam, the switch opens in a time determined by either the carrier lifetime or carrier sweep time, whichever is shorter. A practical limit for switch thickness and therefore holdoff voltage is determined by the maximum electron energy that can impinge on the diamond and not damage the crystal structure. For diamond, this is between 150 and 200 keV, which relates to the electron range in diamond of 40 to 50 µm. Diamond film thickness on this order can support switch voltages of approximately 5 kV. A concept for a compact electron beam-controlled diamond switch projected to operate at average powers of 100's of kW, voltages of 5 kV, and switching frequencies of > 10 MHz, is shown in **Fig. 3**. Switching speed is estimated in nanoseconds with an overall expected efficiency of > 95%. This switch concept utilizes thin film diamond with its high thermal conductivity to be able to handle the high average powers. Microchannel cooling,[3] which is being developed at LLNL, allows for a power dissipation of up to 70 W/cm^2 °C.

In the second mode, a contact allows for electron injection at the cathode. In this configuration, the electron beam energy is lowered such that the electrons penetrate and ionize only a fraction of the switch volume in the immediate vicinity of the cathode (see **Fig. 1b**). The switch operates space charge limited with a current, J_s, which is dependent upon the switch voltage given by:

$$J_s = \varepsilon_r \, \varepsilon_0 \, \mu_n \left(V^2 / d^3 \right), \qquad (1)$$

where ε_r and ε_0 are the relative and absolute dielectric constants, μ_n is the electron mobility, V is the applied voltage, and d is the thickness of the diamond. Typical current gains for a switch of this type is 10^4. The advantage of a switch operating in a single injection mode is a lower beam energy and therefore a simpler system. However, the tradeoff is possible switch jitter due to establishment of the drift current in the base region of the switch.

In the third mode, the switch is operated in the double injection mode, that is, both of the contacts are charge injecting. Double injection allows for higher current and power gains over the single injection mode. Here, the switch is forward biased rather than reversed biased (see **Fig. 1c**). The contact opposite the electron beam-irradiated face acts as an electron injector, while the electron beam generates electron-hole pairs close to the

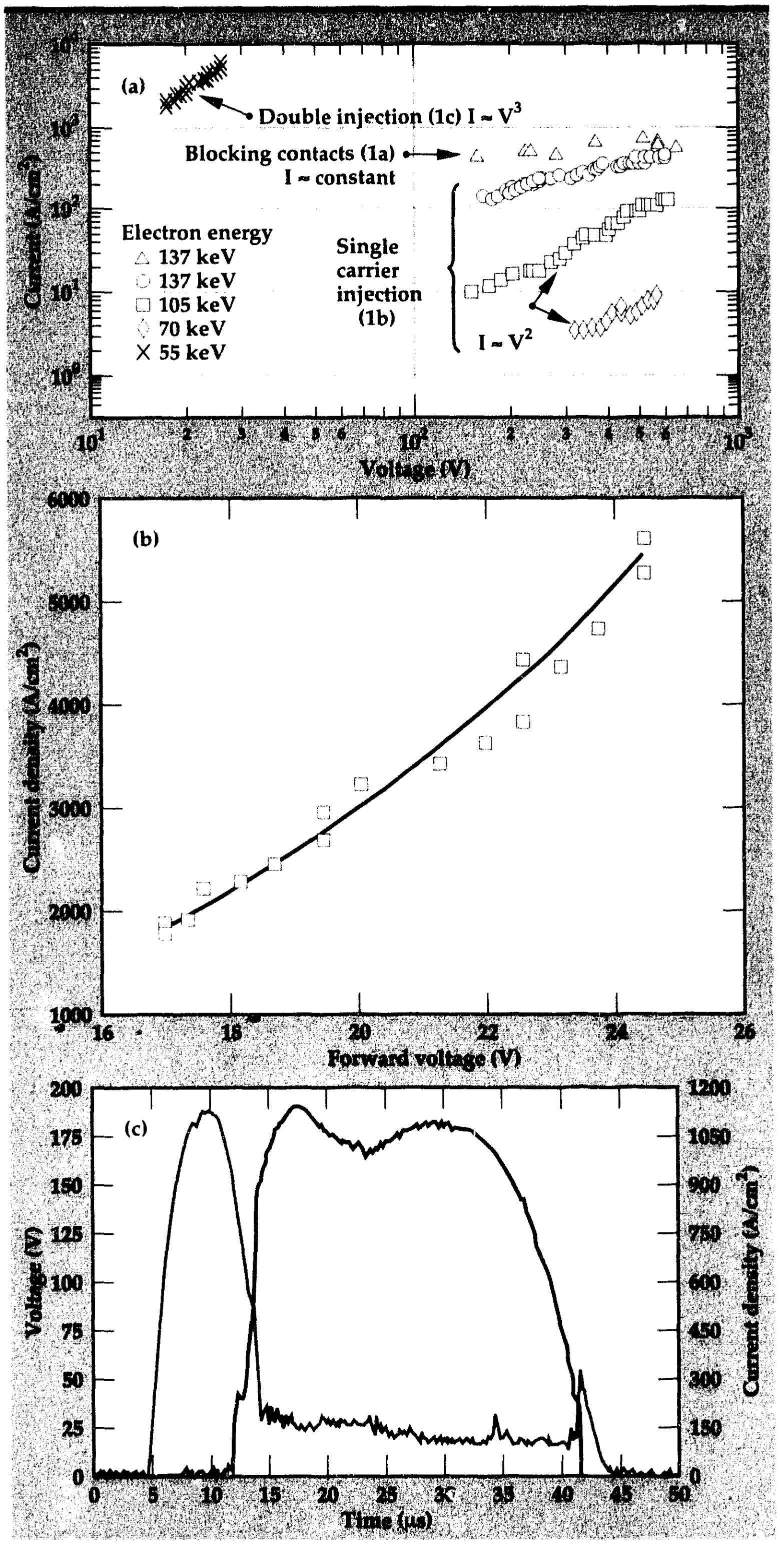

***Figure 4.** (a) Experimental results using blocking, single-, and double-charge injecting contacts; expected voltage dependencies (50-µm-thick natural diamond) are exhibited. (b) Results from experiments using double-charge injection, producing current gains of 10^6. (c) Voltage and current waveforms using double-charge injection.*

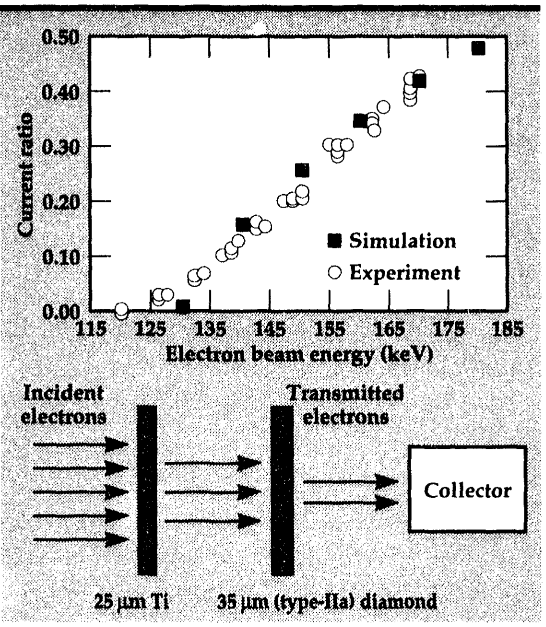

Figure 5. Electron fraction transmitted to collector.

irradiated contact, and serves as a hole emitter. The resulting current is given by:

$$J_d = \varepsilon_r \varepsilon_0 \mu_n \mu_p \tau \left(V^3 / d^5 \right), \tag{2}$$

where μ_n and μ_p are the electron and hole mobilities, and τ is the charge carrier decay time.

Assuming 10-µm-thick diamond, a forward voltage of 25 V, electron and hole mobilities of 1000 cm^2/Vs, a carrier decay time of 1 ns, and a relative dielectric constant of 5.7, the calculated current density, from **Eq. 2**, is 7.88 kA/cm^2. As will be discussed shortly, current densities of this order were obtained using natural diamond having a thickness of five times the value used above. This is possibly due to the fact that after trap filling occurs, the carrier decay time vastly increases (due to the removal of one loss mechanism, i.e., trapping).

Experimental results using blocking contacts, single carrier injection, and double injection, are shown in **Figs. 4a, 4b**, and **4c**. The results to note are:

(1) Switching using blocking contacts results in a current that is approximately independent on voltage. This is expected since the carrier generation is controlled by the incident electron beam and not the voltage across the switch.

(2) For single charge injection, a quadratic dependence on voltage, as described in **Eq. 1**, is seen. The increased current with increasing electron beam energy is likely due to the deeper penetration of the electrons into the bulk diamond, and consequently, the shortening of the effective gap length, d.

(3) For double charge injection, a cubic dependence upon voltage is seen and agrees nicely with **Eq. 2**.

These results are better seen in the expanded plot in **Fig. 4b**. Current gains of 10^6 have been obtained using 50-µm-thick diamond and a 55-keV electron beam having a current density of 6 mA/cm^2. The switch passed a current density of 2–5 kA/cm^2 and had a forward voltage drop of 18–25 V. Voltage and current waveforms for a case with similar parameters are shown in **Fig. 4c**.

Penetration of Electrons and the Generation of Free Carriers in Diamond

To determine the electron distribution and the spatially dependent source functions that result in the induced conductivity in the diamond films, electron beam penetration as a function of electron incident energy was experimentally studied and compared to numerical Monte Carlo simulations. The two are in good agreement, and the results, along with calculated carrier-generation profiles within the bulk diamond as a function of incident electron energies, are discussed in detail elsewhere.[4] The experimental setup consisted of a diode having a thermionic cathode that was powered by a discrete element 40-kV pulse-forming network,which output into a 1:11 step-up transformer. The output voltage pulse at the diode was variable, from 100 to 190 keV and from 1 to 15 µs in pulse length at current densities of up to 35 mA/cm^2. The emitted electron beam then passed through a 25-µm titanium foil and into a 35-µm-thick (type II-a) diamond sample. The electrons emerging from the diamond were collected in a graphite Faraday cup. The total collected current is a direct measure of the fraction of

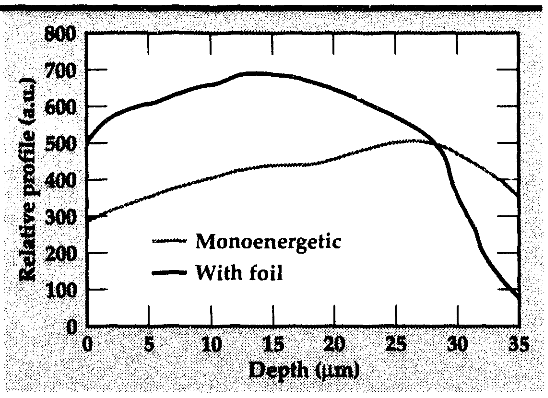

Figure 6. Profiles of carrier generation in 35-µm-thick diamond.

electrons transmitted, which relates to the electron energy loss processes and the stopping power in diamond. The resulting energy loss and diffusion of the initially monochromatic electron beam after passing through the anode have been taken into account. The experimental and theoretical results of the transmitted fraction vs incident electron energy are shown in **Fig. 5** and are in very good agreement with each other. With the confidence that the Monte Carlo simulation is accurate, free-charge carrier profiles throughout the diamond can be determined.

This information is important for efficient switch design since it relates to the drift and diffusive flow of charge within the diamond. Results from the Monte Carlo simulations relating to profiles of excess carrier generation in diamond are shown in **Fig. 6** for an electron energy of 135 keV.

Conclusions

Diamond has many properties that make it ideal for high-average-power, high-voltage, high-frequency solid-state switching. The type of contact used, the biasing, and the electron beam energy all determine the operating parameters of the switch. Results of experiments using blocking and charge-injecting contacts agree well with theory. Electron beam penetration and the generation of charge carriers in diamond were studied computationally, and the results agree very well with experiments.

Future Work

In our further experimental and computational efforts, we will investigate the influence of contact behavior and charge injection on dark current, switch performance, and recovery. Continued steady-state and transient-analysis modeling of charge transport in diamond will provide a better understanding of the physics leading to optimization of switch performance.

---

1. L.S. Pan, S. Han, and D.R. Kania, *Diamond Films and Technology* **2**, 99 (1992).
2. W.W. Hofer, K.H. Schoenbach, and R.P. Joshi, "High Average Power Switching in Diamond," *Proc. Twentieth Power Modular Symposium*, (Myrtle Beach, South Carolina), 241 (1992).
3. R. Beach, B. Benett, D. Mundinger, R. Solarz, B. Freitas, B. Comaskey, and M. Emanuel, *Modular Microchannel-Cooled Heatsinks for High-Average-Power Laser Diode Arrays*, Lawrence Livermore National Laboratory, Livermore, California, UCRL-JC-106942 (1991).
4. R.P. Joshi, K.H. Schoenbach, C. Molina, and W.W. Hofer, *J. Appl. Phys.* **74**, 1568 (1993).

Applying Statistical Electromagnetic Theory to Mode-Stirred Chamber Measurements

Richard A. Zacharias and
Wang C. Ng
Defense Systems Engineering Division
Electronics Engineering

T.H. Lehman
Consultant

We used statistical electromagnetics to characterize the behavior of microwave fields in mode stirred chambers (MSC's) and in aircraft. We did this to assess the feasibility of using an MSC to determine microwave upset thresholds of electronic subsystems that operate in large metal cavities, such as aircraft avionics. By demonstrating that an MSC can generate an electromagnetic environment with the same statistics that an installed electronic subsystem would experience, it follows that the subsystem can be tested uninstalled. Subsystem testing would provide tremendous cost savings over full system tests.

Introduction

Modern transportation systems and military systems are increasingly dependent on sophisticated electronic controls. At the same time, the potential for electromagnetic (EM) susceptibility of electronic systems is increasing as they use more and faster digital circuits in an EM environment being shared by more and more users. EM susceptibility assessments are required to ensure safety and performance of these systems. Full-system, full-threat susceptibility testing is often too expensive to be practical. Computer models alone have been unsuccessful in accurately predicting susceptibility, especially at high frequencies. Also, as the system size becomes many wavelengths at microwave frequencies, the coupling (pick-up of EM energy onto interior components) becomes extremely complex, to the point that deterministic descriptions cannot be made. We believe that the mode stirred chamber (MSC) and statistical EM will provide a new method to make economical and viable microwave-effects assessments of subsystems.

The MSC is a metal-walled chamber, large relative to a wavelength, into which EM energy is injected.[1] It usually contains a stirrer paddle or frequency modulator to randomly change resonant field (mode) patterns. Our analysis tools are based on statistical electromagnetics.[2,3] If an MSC can be used to generate the same EM environment that an electronic subsystem would experience when installed in a system cavity, the subsystem could be tested without installation. MSC testing could provide tremendous cost savings over full system tests.

Four issues must be resolved before certified, quantitative subsystem assessments can be made in MSC's:

(1) MSC tests must be shown to be repeatable, predictable, and rigorous. This is necessary for the technique to be accepted by government regulators and industry.

(2) The power distribution in the MSC must match that found in the system cavity. This ensures that the subsystem sees the same EM power environment as if it were installed.

(3) Coherence (phase-correlation) length must be understood, and the effects of structures near the subsystem taken into account. A structure installed in the cavity will interact with a subsystem if their separation is within a coherence length. Therefore, structures within this length will need to be simulated in the MSC for accurate results.

(4) The transfer function describing coupling to devices in the subsystem must be shown to be separable into a product of coupling from

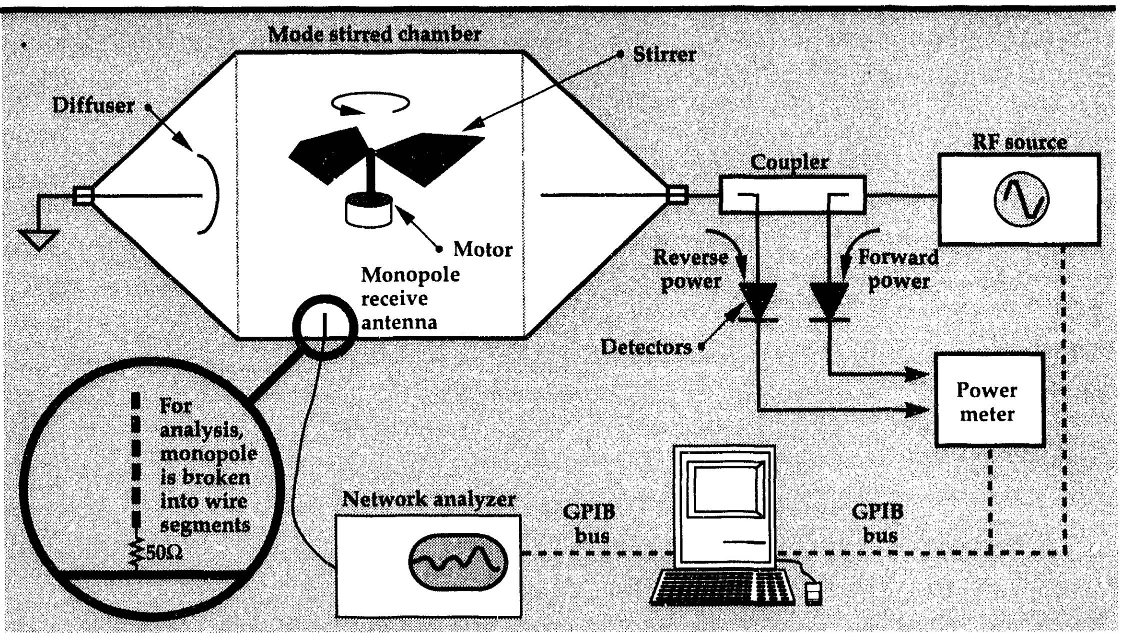

Figure 1. Test configuration for coherence-length experiments.

outside the system cavity to its interior, and coupling from the cavity interior to the devices within the subsystem. An example of a non-separable case would be when the major coupling into the subsystem is through a waveguide that exits the cavity. In this case, the random field environment in the cavity is immaterial, and MSC tests would not provide accurate results.

Progress

In FY-92, we developed a small MSC facility and instrumentation. We made measurements in the MSC and in a commercial aircraft that showed that the power distributions matched each other and the theory.[4] This year, we validated the predicted coherence and evaluated cases where a lack of separability might limit MSC usefulness. We also showed that frequency stirring can be used as a means to increase measurement speed. In this report, we will describe only the coherence length experiments.

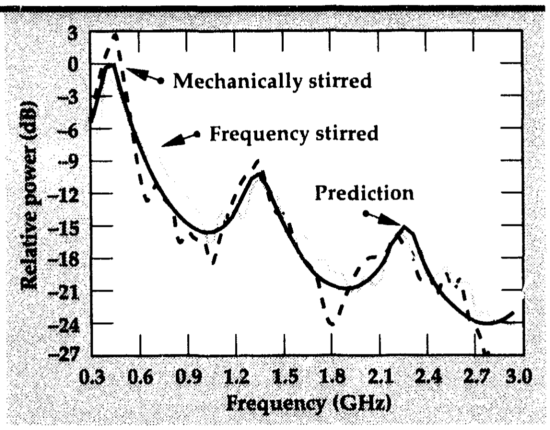

Figure 2. Measured and predicted power received by a 16-cm monopole in a complex field.

The theoretical prediction of the correlation function of the field at two points r_1 and r_2 within a complex cavity is:

$$\rho(r_1, r_2) = \frac{\sin k|r_1 - r_2|}{k|r_1 - r_2|}, \qquad (1)$$

where k is the wave number, which is a function of frequency.[3] Rather than trying to measure fields at two points in space and having to deal with measurement error accumulation, we chose an indirect means to validate this relation by using it to predict the EM power picked up by a simple monopole wire antenna as a function of frequency.

The power to such a monopole was measured in our MSC. The set-up for this experiment is shown in **Fig. 1**. Power was injected into the chamber, which was stirred to randomize the interior fields. The power to the monopole was measured using a network analyzer. Measurements were made for both frequency stirring (stepped frequencies over a 1% bandwidth about each center frequency) and mechanical stirring. The average measured responses for 1600 sample points are shown in **Fig. 2**. The mechanical-stirring data has been smoothed over three points (50-MHz spacing).

A method-of-moments calculation was used to predict the power received by the monopole. The wire was modeled as a series of wire segments, as depicted in **Fig. 1**. Each segment was assumed to be exposed to fields having the previously validated statistical distribution.[4]

The phase correlation of the fields was assumed to be governed by **Eq. 1**. A method-of-moments calculation was made to sum the contribution to the total load current made by each wire segment. The predicted average power to the load is shown in **Fig. 2**.

The measured and predicted monopole frequency response agree quite well, indicating that the derived correlation function is correct. The measured and predicted standard deviation of the monopole pickup also compared well. To our knowledge, these relationships have never been validated before. A (single) measurement of the monopole power measured in our EMPEROR facility with plane wave illumination is shown in **Fig. 3**, along with its predicted received-average power in the random field. With plane wave illumination, the fields at all points illuminating the monopole are phase coherent, so the nulls become deeper than those for the random field environment.

Future Work

The four issues listed in the introduction have been addressed over the last two years. The end product of this study is a set of theoretically and experimentally validated, EM-susceptibility-assessment tools to allow EM effects testing of subsystems at high frequency, while avoiding expensive, high-power, full-system tests. National Aeronautics and Space Administration (NASA) funding for this research will continue, with emphasis on developing methods of assessing the EM survivability of future generations of aircraft, while minimizing or eliminating full aircraft testing.

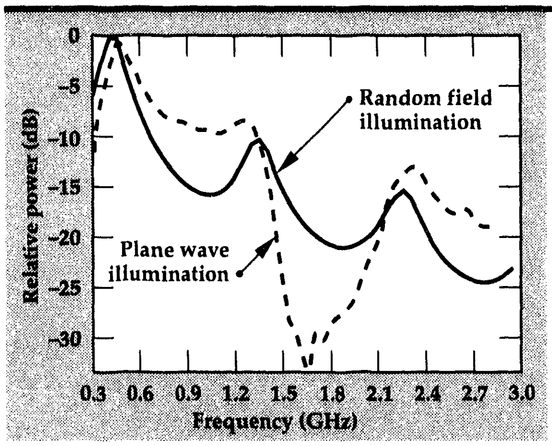

Figure 3. ***Measured power received by a 16-cm monopole under plane wave illumination compared to the predicted power received in a complex field.***

Acknowledgement

This work was funded in part by the NASA Langley Research Center.

---

1. M. Crawford and G. Koepke, *Design, Evaluation, and Use of a Reverberation Chamber for Performing Electromagnetic Susceptibility/Vulnerability Measurements*, National Bureau of Standards, Boulder, Colorado, Technical Note 1092 (1986).
2. R. Price *et al.*, *Determination of the Statistical Distribution of Electromagnetic Field Amplitudes in Complex Cavities*, Jaycor, Albuquerque, New Mexico, Report No. 88JAL129 (1988).
3. T.H. Lehman, *The Statistics of Electromagnetic Fields in Cavities with Complex Shapes*, Phillips Laboratory, Albuquerque, New Mexico, Interaction Note 494 (1993).
4. R.A. Zacharias and C.A. Avalle, "Applying Statistical Electromagnetic Theory to Mode Stirred Chamber Measurements," *Engineering Research, Development, and Technology*, Lawrence Livermore National Laboratory, Livermore, California, UCRL-53868-92, 7-23 (1993).

Nondestructive Evaluation

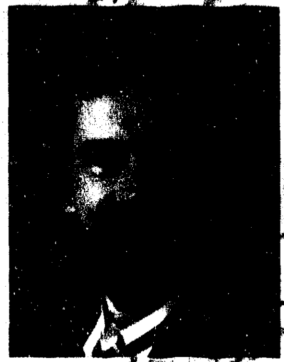

The Nondestructive Evaluation (NDE) thrust area supports initiatives that advance inspection science and technology. The goal of the NDE thrust area is to provide cutting-edge technologies that have promise of inspection tools three to five years in the future.

NDE provides materials characterization inspections, finished parts, and complex objects to find flaws and fabrication defects and to determine their physical and chemical characteristics. NDE also encompasses process monitoring and control sensors and the monitoring of in-service damage. For concurrent engineering, NDE becomes a frontline technology and strongly impacts issues of certification and of life prediction and extension.

In FY-93, the primary contributions of the NDE thrust area, described in the reports that follow, were in computed tomography, digital radiography, and laser ultrasonics.

Harry E. Martz
Thrust Area Leader

Section 8

8. Nondestructive Evaluation

Overview
Harry E. Martz, Thrust Area Leader

Computed Tomography

Stephen G. Azevedo
Engineering Research Division
Electronics Engineering

Daniel J. Schneberk
Applications Systems Division
Computations

Harry E. Martz
Manufacturing and
Materials Engineering Division
Mechanical Engineering

G. Patrick Roberson
Defense Sciences Engineering Division
Electronics Engineering

Our continuing work in scientific and industrial computed tomography (CT) has spanned all aspects of theoretical and experimental practice. We are applying the state-of-the-art methods of data acquisition, processing, analysis, and understanding associated with penetrating-radiation imaging. This report describes recent progress in our CT applications for FY-93. We have sought to advance the state of the art in CT technology, while at the same time actively support programs at Lawrence Livermore National Laboratory (LLNL) and new business initiatives. Our goal is to provide reliable and efficient nondestructive evaluation techniques for use in probing the internal structure of fabricated objects and materials associated with a broad spectrum of LLNL programs, industrial partnerships, and academic research.

Introduction

Computed imaging, where raw measurements must be substantially processed before an image can be displayed or analyzed, continues to be an active area of research and development for medical diagnoses, material science, and industrial nondestructive evaluation (NDE). Lawrence Livermore National Laboratory (LLNL) has recently had significant and broad impact in computed tomography (CT) of x-ray and gamma-ray measurements. First used in the 1970's as a medical diagnostic tool, CT was adapted to industrial and other nonmedical purposes in the mid-1980's, where it showed superior imaging when used with standard radiographic techniques.

Single-projection radiography hides crucial information: the overlapping of internal features obscures parts of these features, and their depth remains unknown. CT technology was developed to retrieve complete three-dimensional (3-D) information. It involves acquiring multiple radiographic images of the sample object at different angles, then collecting intensity information from one or many detectors for processing in a computer. The final 3-D image, generated by mathematically combining these images, gives the relative locations and dimensions of features within the object, as well as external details.

While medical CT technology is well developed due to the relative invariability of the human anatomy, CT for scientific and industrial purposes continues to be an active field of study and experimentation. The enormous applicability of this technology to objects of extreme differences in size, shape, and material make-up makes CT an ongoing area of scientific interest. LLNL's unmatched facility of experimental CT scanners and computers has made possible an array of industrial, scientific, and military applications. While continuing to serve the NDE needs of LLNL programs, the publications and presentations of CT projects at technical conferences form the bases for new collaborations with U.S. industry.

For the past seven years, we have maintained a balance of support to LLNL programs, research with academia, development of hardware and software, and technology transfer to industry. Our CT research and development efforts have been concentrated in three main areas: (1) scanners, (2) software tools, and (3) applications.[1–7]

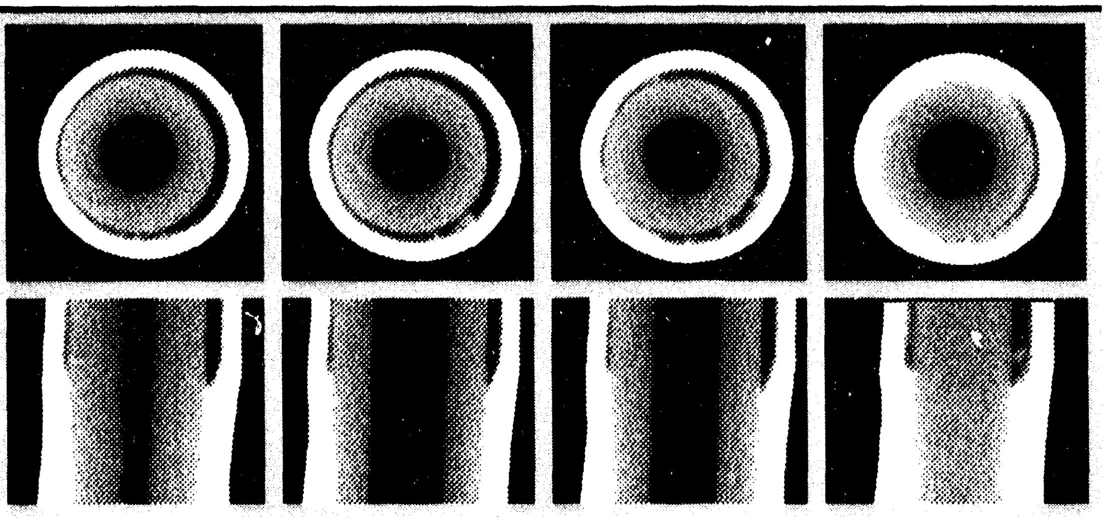

Figure 1. Horizontal (top) and vertical (bottom) cross-sectional CT slices through representative braze tube welds.

This report describes our progress in the three areas in FY-93. Applications have influenced the direction of our research and development strategies. We have sought to advance the state of the art in CT technology, while at the same time actively supporting LLNL programs and new business initiatives. Lastly, we outline our future plans in the field of CT.

Progress

In FY-93, we worked on various applications, each of which has driven the direction of our research and development into new NDE areas. We present our progress in scanners and software by describing these CT areas of application.

Braze Tube Imaging

The high-density scintillator glass researched last year[1,8] has provided a capability for low- to medium-energy scanning at substantial spatial resolution (~ 50 µm). The simplest configuration of this detector consists of a camera focused through a turning mirror into a monolithic piece of glass, thereby protecting the camera from the x ray fluence. MDCAT is an instance of this kind of scanner, equipped with an inexpensive COHU 4910 camera interfaced to a MATROX IMAGER-AT frame-grabber on a DELL PC. Each rotational view from this type of scanner is a two-dimensional (2-D) digital radiograph, which can serve as input to a 3-D, cone-beam, tomographic reconstruction algorithm. The output of the scanner is a 3-D cube of data that can be viewed from any direction.

As part of a collaborative effort between engineers at the Pantex Plant in Amarillo, Texas and at LLNL, we applied our MDCAT scanner to the 3-D inspection of braze tube welds (3 to 5 mm in diameter). After some initial radiographic evaluation, a suspect braze tube was scanned at 125 kVp and 5.0 mA, using the Phillips 320-kVp, x ray-machine source. **Figure 1** contains a number of CT cross-sectional images from the 3-D cube of data

Table 1. Relationships among basic geologic parameters, CT parameters, and inferred parameters of special interest.

Basic parameters	CT parameters	Inferred parameters
Density	Energy and attenuation	Sediment type
Percent moisture	Water contrast	Moisture content
Elemental makeup (e.g., mineral content)	Energy resolution	Chemical analysis
Roundness	Spatial resolution	Porosity
Grain-size distribution	Histogram analysis	Porosity
Pore throat size	3-D histogram	Permeability
Connectivity	3-D void volume	Flow direction (permeability tensor)
Fracture size and direction	Volumetric analysis	Fracture volume

that resulted from the scan. The top row of images contains horizontal slices at different vertical positions down the length of the braze joint. The bottom row of images shows vertical slices through the cube from top to bottom. The middle left slice in the lower row contains a small crack discovered with this scanner, and later verified with metalographic sectioning. This and other applications of this scanner have been useful to inspection problems in the LLNL weapons program.

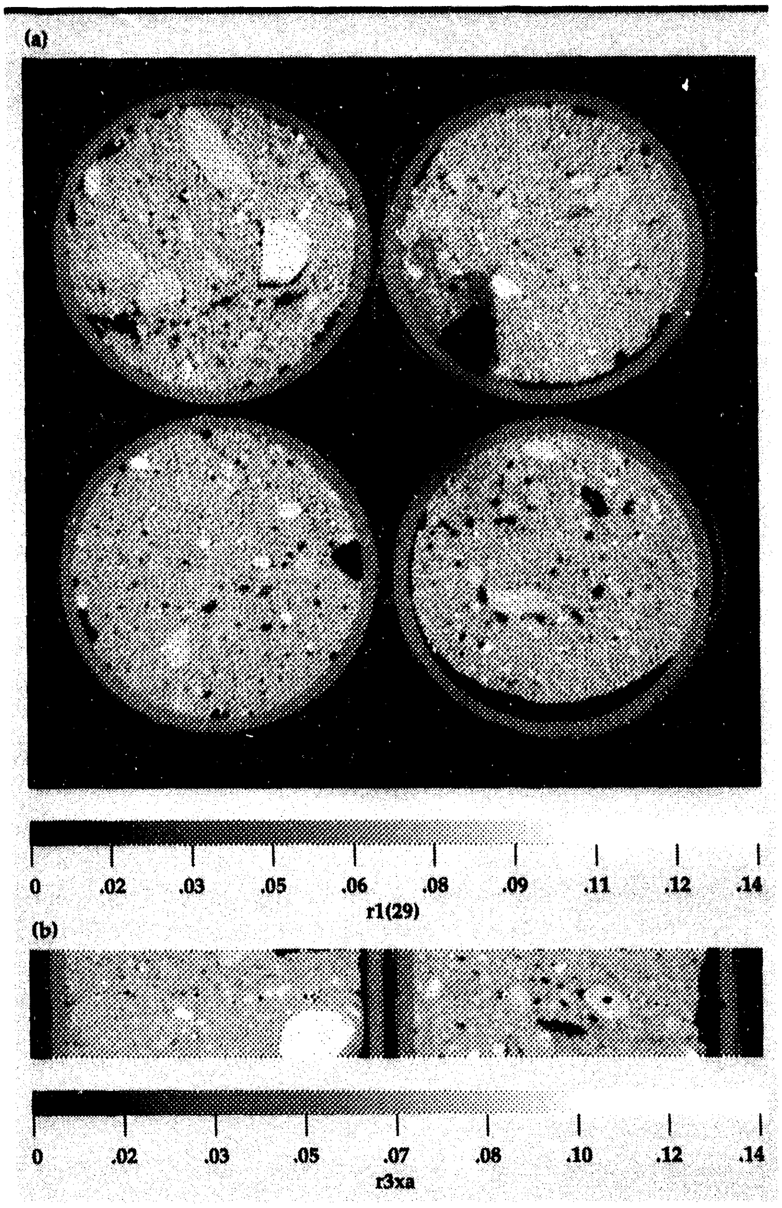

Figure 2. Horizontal (top) and vertical (bottom) cross-sectional CT slices through representative 1-in. soil samples using the LCAT scanner.

Soil Samples

For the past two years, we have performed 3-D CT scanning of soil samples for LLNL's Environmental Protection Department (EPD). More than 100 3-in., soil-drill samples have been scanned to select the locations for further analysis. This has included higher-resolution scanning and/or thin-section examination. **Figure 2** contains a representative CT slice acquired on the fan-beam LCAT scanner. LCAT consists of a 160-kVp x ray source and a collimated, Thomson linear-detector array of 1024 photodiodes. Only recently have we used the MDCAT scanner with the smaller pixel sizes on small samples. **Figure 3** contains a representative image from this type of scanning. The wide spatial resolution and energy range of our scanners provide a fully developed capability for a wide range of objects.

More recently, we have begun an in-depth study of soil parameters measured by CT. Environmental regulations and scientific research into the study of ground core samples have led to methods for understanding soil parameters in the field. The distribution and movement of toxic materials in the earth are functions of the permeability and porosity of the strata. Potential clean-up of these contaminants depends on the physical parameters of the strata. The parameters that are needed have been determined by recent research results[9] and are shown in the left column of **Table 1**. Current techniques for determining these parameters are labor-intensive, inconsistent, destructive to the sample, and expensive. One method involves sending the core samples to laboratories that measure porosity and permeability by 'flow-through' methods. The flow is only measured in one direction or, at best, orthogonal directions. Since this method is destructive and not closely associated with the sample region, a high degree of variability is exhibited between different laboratories. Another method uses thin-section preparation and examination through high-powered microscopes. For example, the connectivity of pore spaces has been shown in preliminary studies[10,11] to correlate with permeability and porosity of the samples. Quantitative image analysis of the microscope images is a powerful technique for estimating rock properties such as porosity, permeability, and pore-size distribution. Again, while the necessary properties are 3-D in nature, thin-section work is inherently 2-D.

CT, on the other hand, provides 3-D volumetric information that can be quantitatively accurate and dimensionally correct. In addition, it is completely nondestructive, allowing more consistent results and effective correlation between tests. The per-voxel information obtained from CT is similar to data from flow-through and thin-section tests.

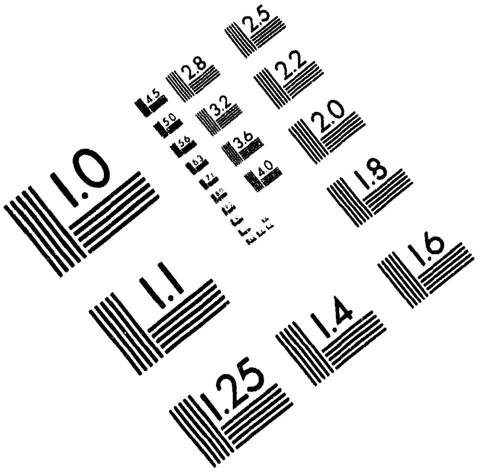

AIIM

Association for Information and Image Management

1100 Wayne Avenue, Suite 1100
Silver Spring, Maryland 20910

301/587-8202

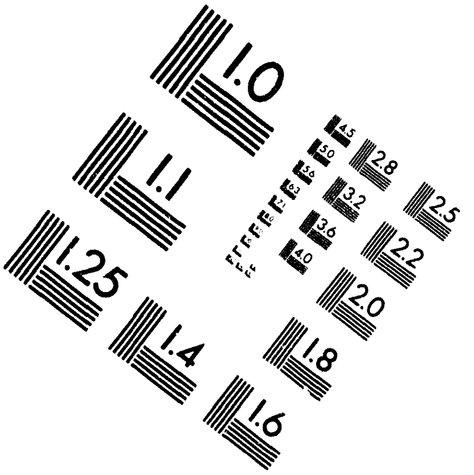

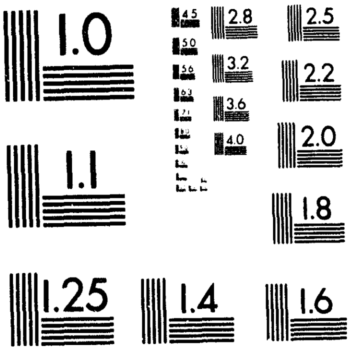

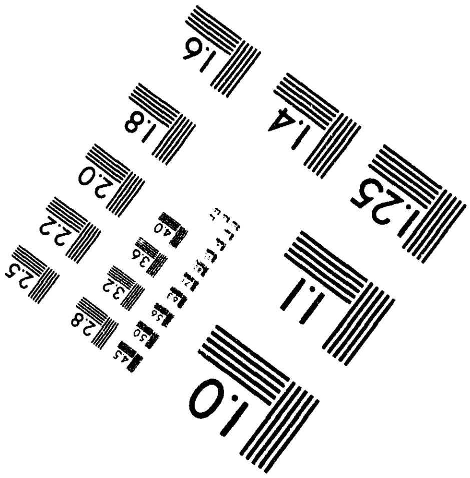

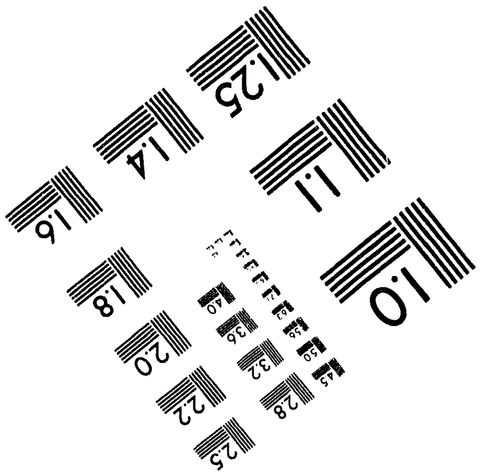

MANUFACTURED TO AIIM STANDARDS
BY APPLIED IMAGE, INC.

3 of 3

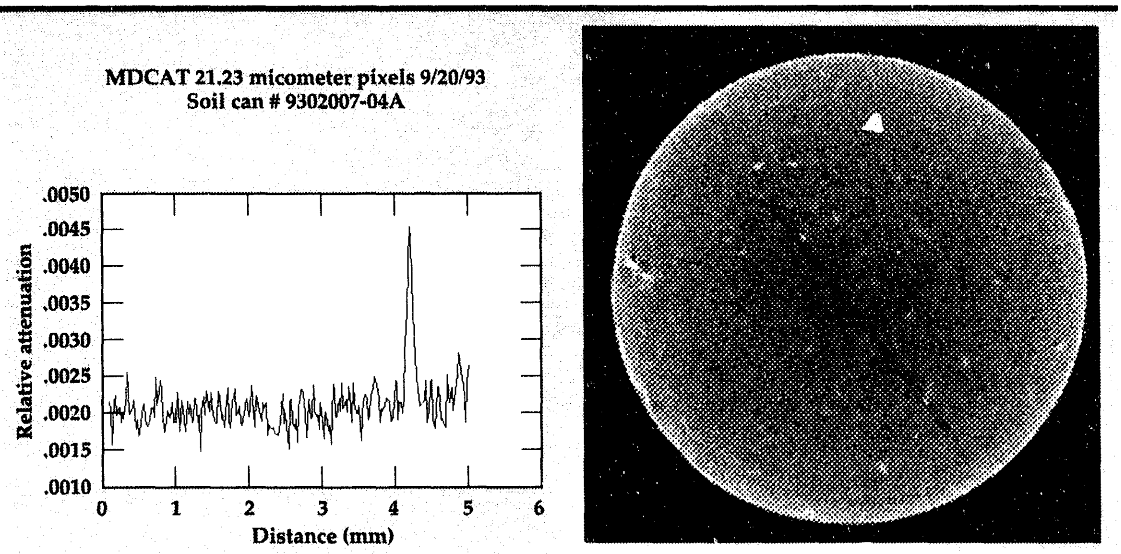

Figure 3. High-resolution CT slices through small sandstone soil samples.

The data is also digital and stored in computers, thus allowing a wide range of powerful image-processing techniques and automated analysis.

Recent work by Drs. Ehrlich and Prince of the University of South Carolina has proven that thin-section analyses can duplicate data measured in the laboratory, using the flow-through methods. We are working with them and with researchers in the Environmental Restoration Division of EPD to understand and quantify the correlations among thin-section, flow-through, and CT analyses. **Table 1** shows the parameters of interest and the CT parameters that can be brought into the analysis.

The overall objective of this project is to develop x-ray CT into an analytical tool that can be used to rapidly scan sediment cores and soil samples, deriving the physical measurements necessary for geological and hydrological assessment of contaminated Department of Defense and Department of Energy (DOE) sites. This testing is being done in such a manner that the core material is preserved for additional testing, including chemical analyses to determine the types and concentrations of contaminants. A field-transportable CT system for core sample studies is being proposed. Not only will CT provide a fast, accurate method for the characterization of porosity in sediments, but it will also provide a series of valuable measures of the pore network, including pore volume density, pore diameter, throat diameter, and the 3-D interconnectivity of the pore network. Using this data, it should be possible to derive both permeability and a permeability tensor, which is a measure of the directional permeability heretofore unavailable for hydrologic modeling.

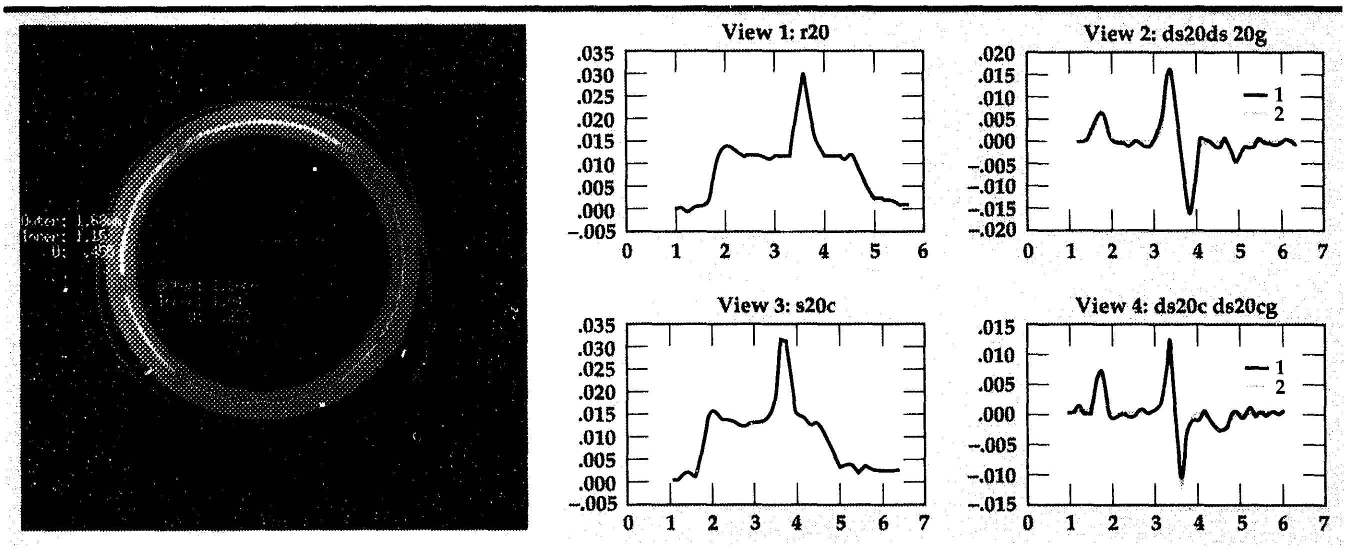

Figure 4. Wall-thickness measurements of a nuclear fuel rod. For all plots, distances are in mm.

Dimensional Measurements

Most work in the CT project until FY-93 has concentrated on improving our imaging capabilities; i.e., constantly enhancing features such as resolution, contrast, quantitativeness, and energy discrimination. This last year, however, we have been systematically applying further analysis tools to the images to gain other crucial information about the object under inspection. This includes such parameters as physical dimensions (e.g., wall thicknesses, radii of curvature, angular separation), defect volume or count, and moments (e.g., centroid, axes of symmetry). All of these measurements have associated error bounds that are determined from the noise levels inherent in the CT scan.

We are developing tools to measure and report both the dimensional measurements and the bounds on those measurements, for a variety of applications. Applications with industry involve non-contact metrology (2-D/3-D, interior/exterior dimensional measurements) that complement or replace coordinate measurement machines (used for art-to-part, physical-model-geometry acquisition, CAD/CIM model generation, and rapid prototyping), reverse engineering (automatic mesh generation and new part design), and competitive assessment.

An example of the use of CT for internal measurement is shown in **Fig. 4**. A cross-sectional CT slice of a nuclear fuel rod for the Savannah River New Production Reactor is at the left. The white region is a uranium core, surrounded on both sides by aluminum cladding. A plastic cover surrounds the entire fuel tube to prevent contamination. The wall thicknesses of both cladding and core at various locations around the tube were of interest to the researchers. A method of measuring these quantities was incorporated into the VIEW software. This involves the use of line traces (middle two plots) through the tube wall, and derivatives of those traces (right two plots). Several estimation procedures were developed[12] to find the edge centers, then widths were calculated to an accuracy of better than one pixel. This same technique has been applied to internal and external measurements of turbine blades, Caterpillar components, weapons assemblies, and automotive parts. Pitch-angle measurements of Caterpillar fuel injectors, taken from the object centroid, are shown in **Fig. 5**.

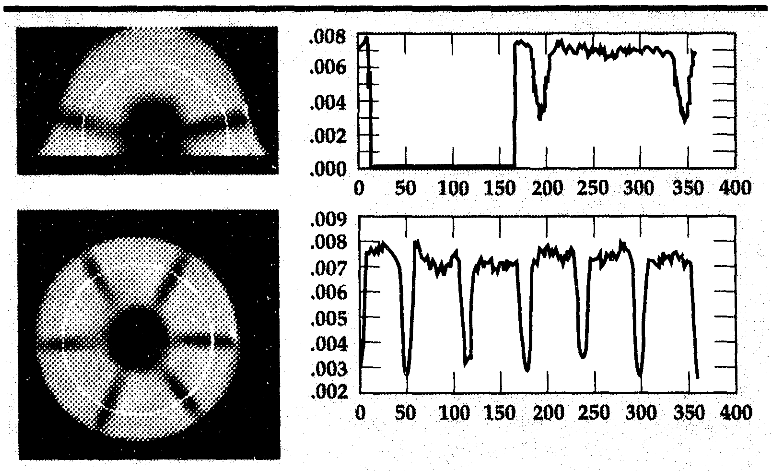

Figure 5. Pitch-angle measurements for a Caterpillar fuel injector.

Another measurement application is shown in **Fig. 6**. The CT cross-section of a cast piston with shrinkage defects is shown at left. We used the VISION code[13] to segment the void areas from the surrounding material, and to determine volume percent of voids (only the voids are shown at right). This particular application can be done in three dimensions with volumetric CT data, as is shown in **Fig. 6** for two dimensions.

Most of our work to this point has been with 2-D images; our future work will focus on 3-D measurements. Making dimensional measurements with 2-D tomographic scanners involves an error, because only the features that are perpendicular to the 2-D slice plane can be dimensioned with the full accuracy CT can provide. A number of researchers have extended the accuracy of CT dimensional measurements to sub-pixel intervals. However, this is only valid if the feature is in a certain orientation relative to the scan geometry.

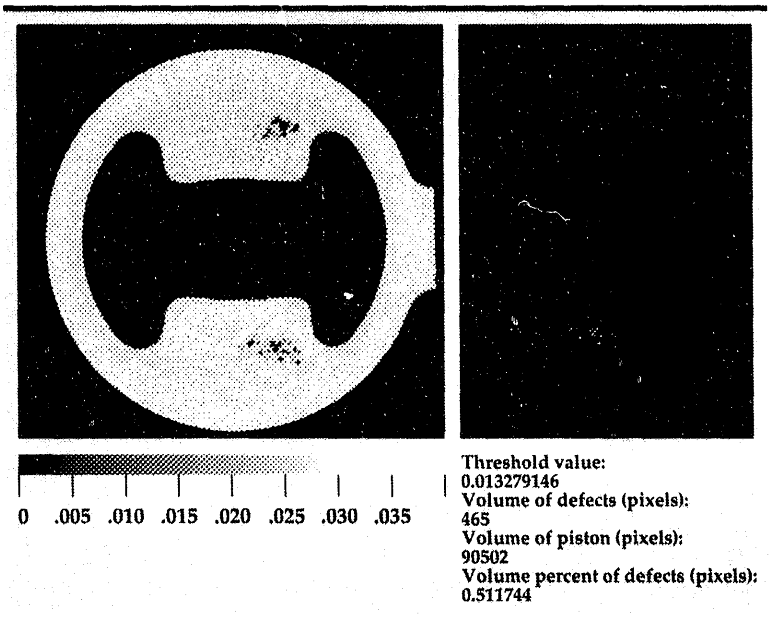

Threshold value:
0.013279146
Volume of defects (pixels):
465
Volume of piston (pixels):
90502
Volume percent of defects (pixels):
0.511744

Figure 6. Shrinkage measurements in a cast piston.

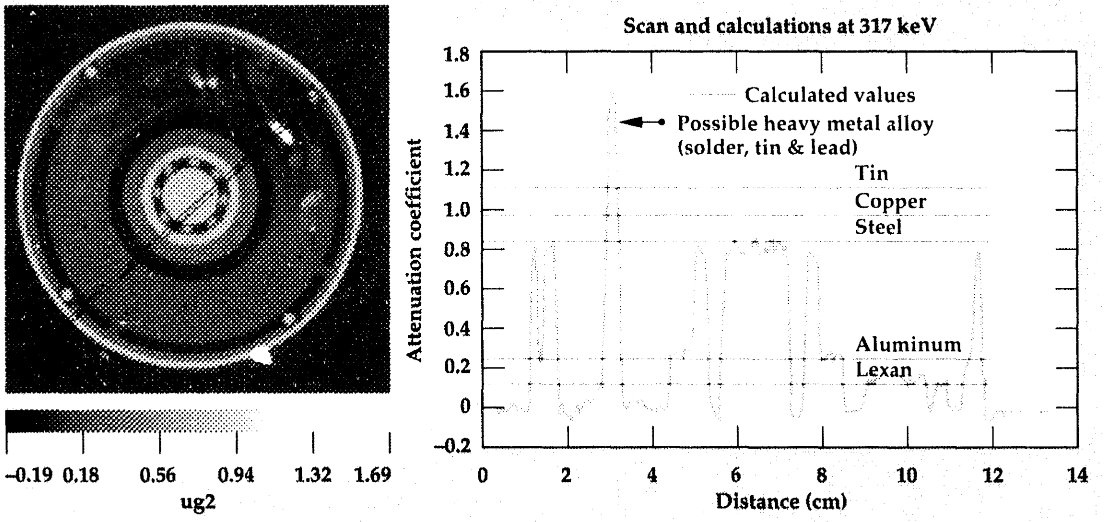

Figure 7. (a) Representative CT image of a DC motor. The colorbar below the image provides a scale relating the different colors (shades of gray here) in the CT image to the linear attenuation coefficient in cm^{-1}. The attenuation variations are related to changes in density (r) and/or effective atomic number (Z_{eff}); (b) a one-dimensional profile (black line) through the CT image showing the variation of the linear attenuation coefficient within the motor. The measured attenuation coefficient is compared to theoretically calculated values of various materials [labeled and shown by colored (shades of gray here) flat lines]. This enables us to determine whether the imaged item contains homogeneous or heterogeneous metals, or radioactive, combustible, or unknown material(s).

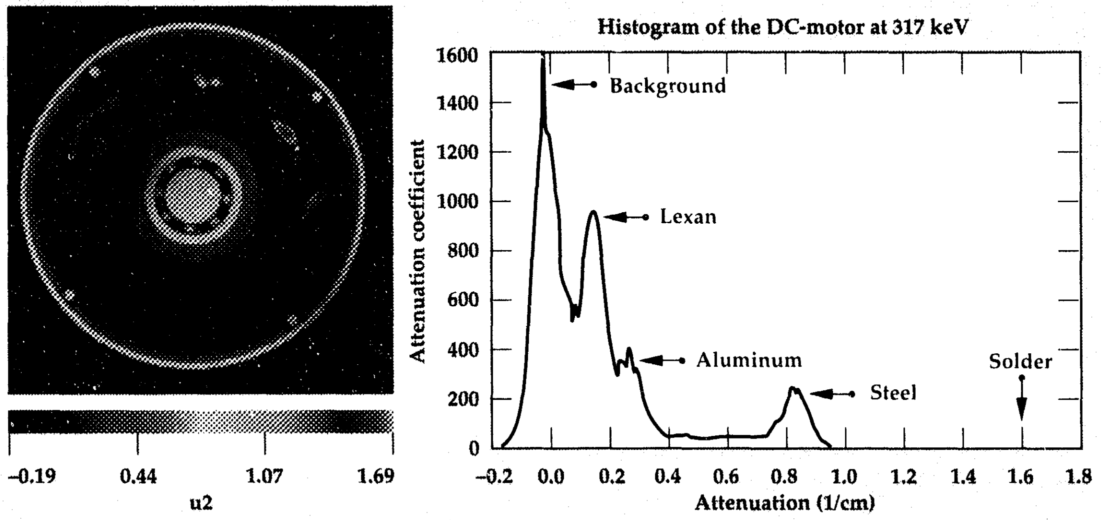

Figure 8. Histogram of the CT image, used to identify the materials within the individual waste item and to indicate their relative amounts. If the histogram contains only two peaks, one for background (air) and one for steel, the item would be classified as a homogeneous metal. The histogram of this CT image contains many peaks and, therefore, indicates the presence of heterogeneous metals.

Our ongoing work with industrial partners has emphasized the complicated 3-D structure of some of the simplest components (e.g., turbine blades). We see the possibility for CT to perform high-quality, non-contact metrology, but a key requirement is the ability to measure in all three dimensions with equal error tolerance. We have an aggressive program to develop 3-D metrology tools with CT-scanned data.

Mixed Waste Characterization

The Robotics Technology Development Program's Waste Facilities Operations (WFO) has been established (1) to develop, demonstrate, test, and evaluate robotics technology for existing and planned DOE waste facilities, where sorting, characterizing, and/or repackaging radioactive, hazardous, and/or mixed waste will occur; and

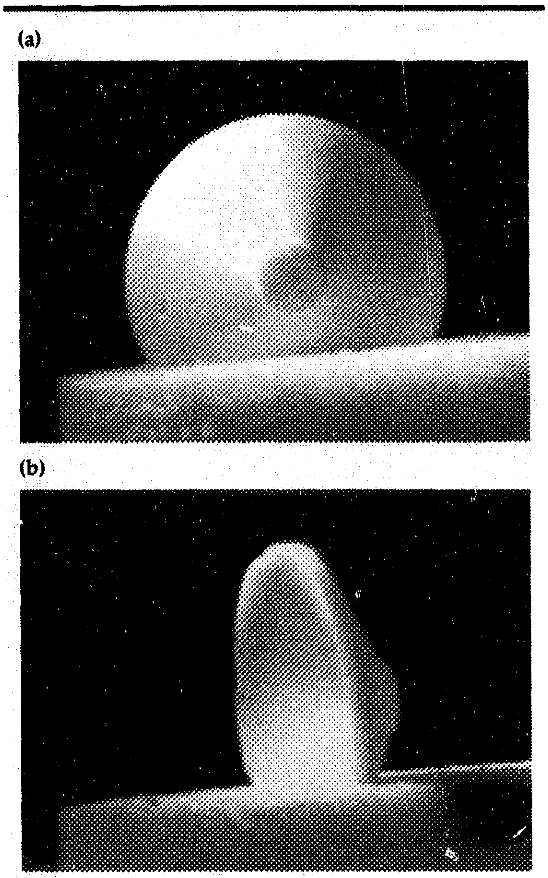

Figure 9. Photographs of a high-explosive hollow shell.

(2) to transfer this technology to industry. The goals of the WFO program are to provide technologies that remove humans from both radiological and physically dangerous environments in waste facilities; to improve the quality of operations; to increase the speed of operations; to reduce the manpower requirements for operation; to reduce life-cycle costs; and to meet federal, state, and local regulations. The Savannah River Technology Center in Savannah, Georgia is designated as the WFO team coordinator, with support from Fernald Environmental Management Project, Idaho National Engineering Laboratory, LLNL, Oak Ridge National Laboratory, and Sandia National Laboratories.

The specific LLNL tasks in the WFO program include the development and implementation of integrated waste-stream-characterization technologies such as non-invasive waste-drum and item characterization using CT, structured-lighting-based object definition, metal/nonmetal object discrimination, radiation detection, and metal-composition determination. This year, we have applied advanced, nondestructive screening and characterization techniques (e.g., fluorescence, gauging, radiography, and CT) to better sort individual items removed from within 208-l (55-gal) waste drums. We will apply these techniques to optimize the plant-scale throughput for the screening, characterization, and sorting of these individual items. For example, radiography would be used for initial screening of the waste item. If these results reveal that the item is homogeneous, a simple transmission gauge could aid the characterization. However, if the item is heterogeneous, the item should be characterized using a CT scanner to help identify its composition. Lastly, fluorescence would be used to confirm these results. These characterization techniques will aid the sorting of items into the following categories: (1) homogenous metals; (2) heterogeneous metals; (3) radioactive materials; (4) combustibles; or (5) unknown material(s).

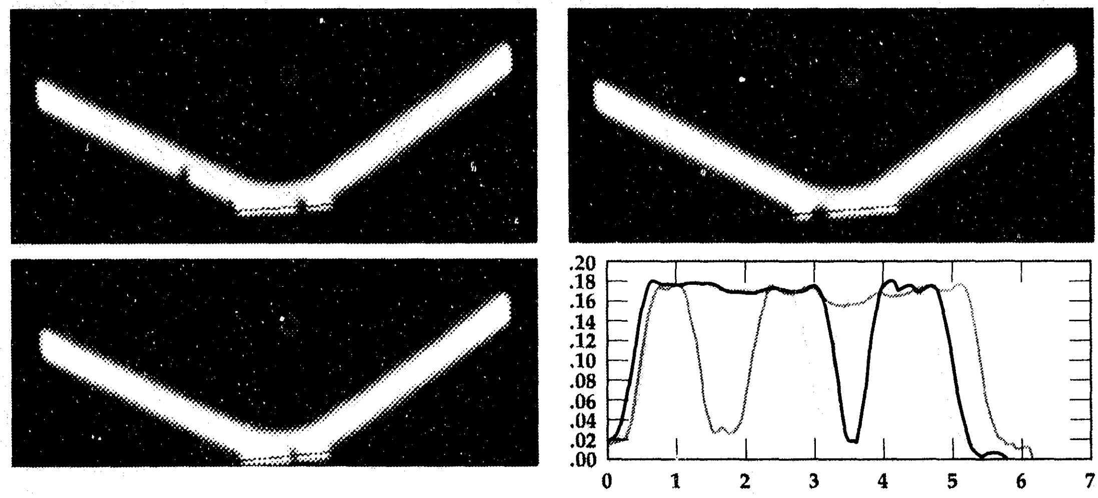

Figure 10. Reconstructed tomograms of the high-explosive hollow shell.

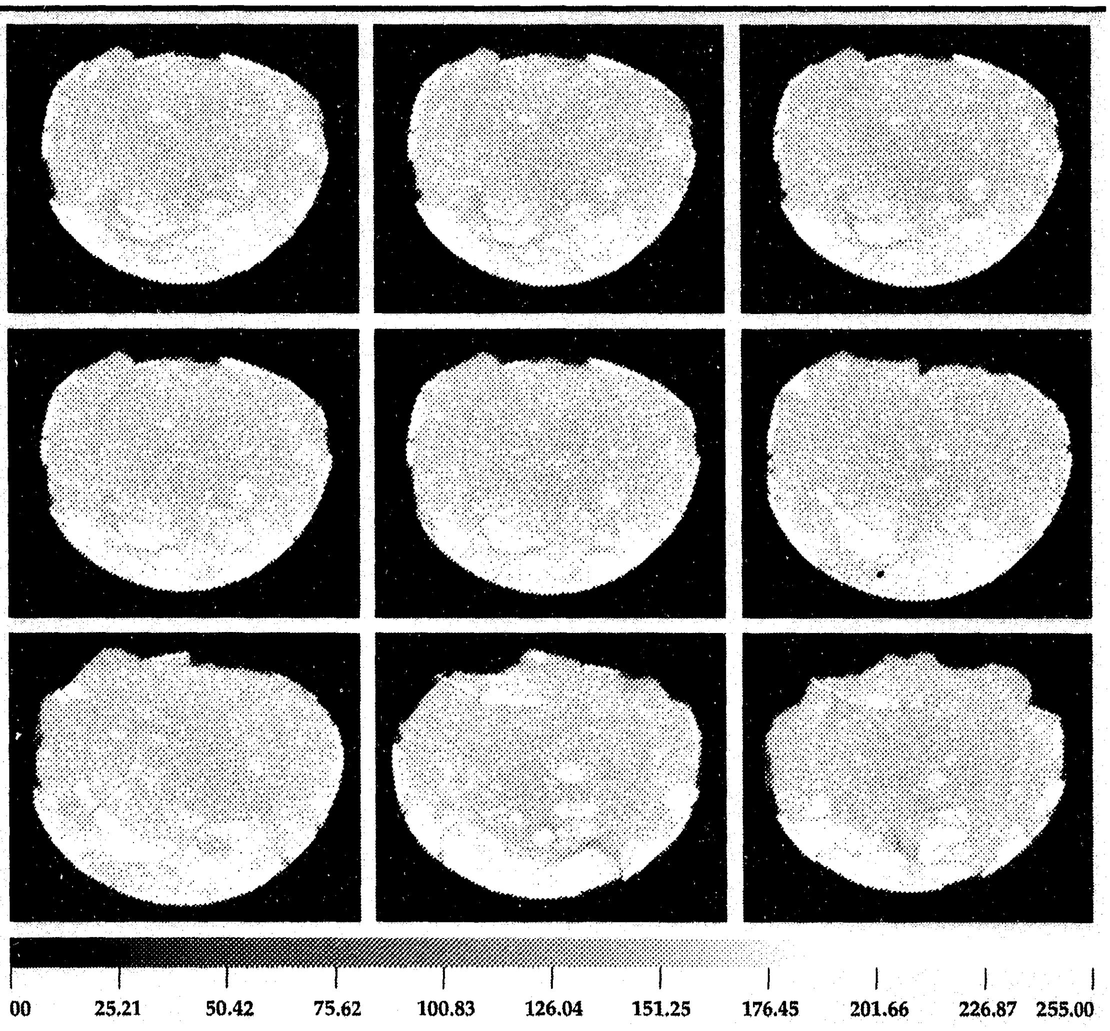

Figure 11. CT of dinosaur eggs.

We have used available LLNL CT scanners[3,14] to develop methods to optimize the 3-D characterization of mock-waste drums and individual waste items.[15,16] Examples are shown in **Figs. 7** and **8**. One scanner used for this work incorporates a well-collimated single source (isotopic or x ray tube) and a single, nuclear-spectroscopy-based detector. Each data point or pixel in the reconstructed images produced from this scanner represents the linear attenuation coefficient for a particular volume element within the object being inspected. This is the same value provided in tables of cross-sections. The experimental CT results can be compared to theoretically predicted values, to identify materials; no additional measurements are required. Furthermore, ratios and differences of CT reconstructed images taken at different energies reduce to functions that can be used to calculate the material properties for each voxel.[17] We are investigating methods to optimize the characterization of these items and minimize scan time.

Additional Applications

For the last few years, we have used our various CT systems to solve diverse experimental problems. In each case, CT has been demonstrated to be an effective tool for nondestructive inspection and / or characterization. In this section, we describe some additional NDE research problems investigated during FY-93.

Small, High-Explosive Shaped Charges. In cooperation with the Pantex Plant in Amarillo, Texas, we are investigating small, high-explosive shaped charges, using CT.[18] The fan-beam LCAT system is used to identify inhomogeneities, cracks, voids, and inclusions. This system provides 3-D images at 150-mm spatial resolution. Due to the energy integrating nature of this scanner, however, it cannot provide quantitative attenuation data. For quantitative absorption coefficients and densities, a collimated single-beam system with a spectroscopy detector is used.[19] The effective atomic

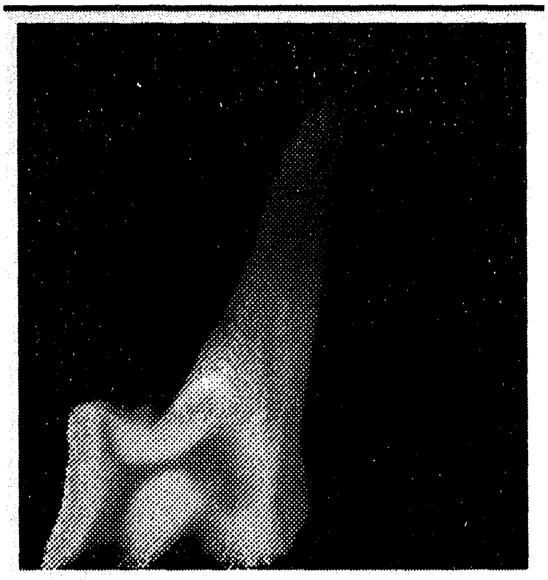

Figure 12. CT of prehistoric mammal teeth.

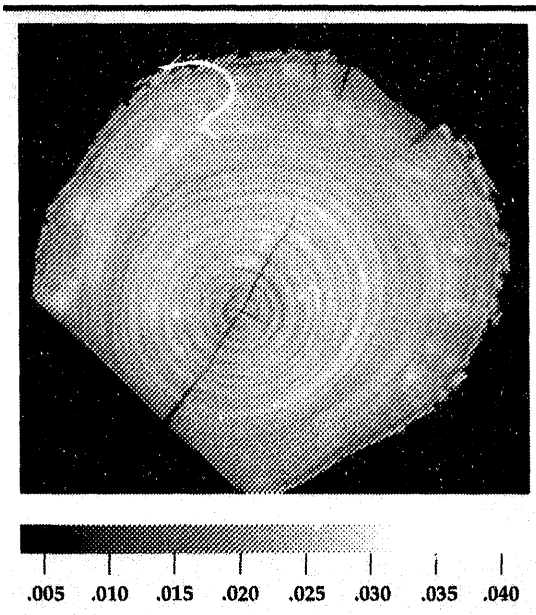

Figure 13. Large scan (1500 angles) of black-cherry log sample.

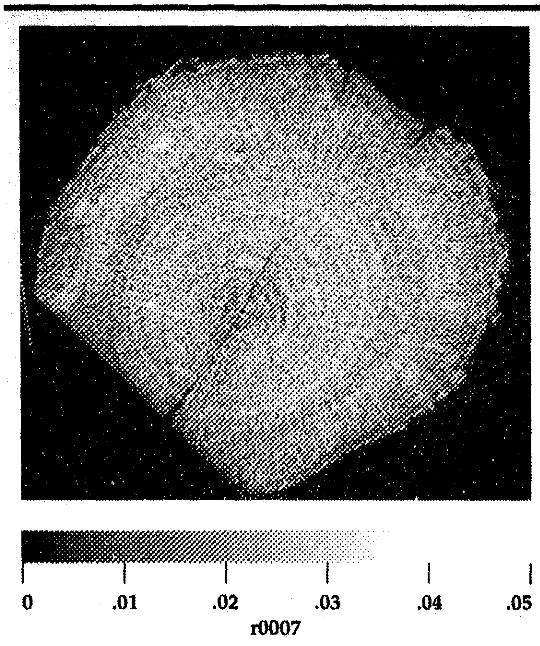

Figure 14. Lower-resolution scan (75 angles) of black-cherry log sample.

number of the materials is determined from ratios of measurements at different energies. We thereby obtain cross-sectional images of the chemistry, which can reveal segregation of explosive and binder.

Figure 9 shows photographs of an explosive hollow shell, which has six small holes drilled in it to test our ability to find small flaws. The shell diameter is 23 mm (0.9 in.); the holes range from ~ 0.2 mm to 0.77 mm (8 to 30 mils). **Figure 10** shows tomographic slices (using LCAT) at the positions of the holes; all are clearly visible. The geometric magnification for these images was 3.5x. The specimen is seen to be uniform in density in these images; this is confirmed by the spectroscopy-based CT results.

Superconducting Computer Research. A consortium consisting of industrial (Conductus, TRW), academic (University of California Berkeley, Stanford University), and government (National Institute of Standards and Technology, LLNL) researchers has undertaken to design and build the fastest processor available (100-ps clock rate, or 10 GHz) using superconducting Josephson–Junction (JJ) logic running at liquid-helium temperatures. Recently, Fujitsu in Japan produced an 8-bit, superconducting, digital signal processing (DSP) chip, but due to a corporate business decision, it was not commercialized. However, more recent progress by Conductus and TRW make the JJ design more feasible. This consortium has been funded for a five-year research program to develop a superconducting processor prototype. A DSP design was chosen for the prototype[20] because of its relative simplicity, its pipeline structure for fast throughput, and its lack of on-chip memory (which is difficult to fabricate in JJ's).

At LLNL, we direct the JJDSP design for massive image-processing problems such as 3-D, medical-image reconstruction (either x-ray CT or magnetic-resonance imaging) and volumetric display. An overall design has been specified[21] for fast image- and signal-processing using the JJDSP, currently under development for the Advanced Technology Program in the U.S. Department of Commerce. This processor is expected to be completed in the next three years and to have a 100-fold speed improvement over current DSP technology. The processor design may be influenced by the results of the algorithms developed and the software tests conducted, to optimize performance.

The algorithms we develop have far-reaching potential beyond the JJDSP hardware. Indeed, the same algorithms could be implemented on fast, silicon, DSP chips that are currently available for improved volume-reconstruction and rendering of medical images. It is now impossible to perform these tasks in real time, but with the JJDSP, this becomes a possibility. Other applications could be fast industrial-imaging for process control in manufacturing, real-time image processing for high-definition television applications, synthetic aperture radar processing, and other signal- and image-processing applications.

Archaeological Objects. We have used CT to investigate two fossilized dinosaur eggs (**Fig. 11**) and several pre-historic mammalian teeth (**Fig. 12**) that are of interest to archaeologists, under the direction of Prof. J. Keith Rigby at the University of Notre Dame. In all cases, the objects are most interesting and valuable intact, and the scientists do not want to section them for their studies. Medical x-ray CT scanners do not provide adequate imaging in most cases to visualize the tiny internal features in these objects, nor sometimes to even penetrate the dense rock.

The 65-million-year-old dinosaur egg from southern China (**Fig. 11**) has some clear internal features that are of interest to the scientists, but the 3-D results are inconclusive as to the type of animal, due to significant crushing of the skeleton during the fossilization process. The radiographic contrast between the fossilized bone of the embryo and the surrounding rock is very small. Even detectors with the highest dynamic range fail to make out the entire skeleton. We are continuing to improve both contrast and spatial resolution of our scanners (see Future Work below).

The mammalian teeth are found in the northwestern U.S. and provide archaeologists with a look at the evolutionary process of these rodent-sized creatures over millions of years, starting before the dinosaur extinction event. Internal features, such as the relative thickness of the dentine and enamel layers, provide clues about the lives of these creatures, such as what they ate and their age at death. We have used our scanners to image some of these teeth (**Fig. 12**) and are seeking funding for continued research.

Timber Analysis. In conjunction with Dr. Dan Schmoldt at the U.S. Forest Service (USFS), we have investigated the possibility of using x-ray CT to inspect felled hardwood logs for optimal milling strategies. Each log entering a mill may potentially be cut in a variety of ways, whether it be for veneer, plywood, or lumber. Because of the slow growth of hardwoods and the demand for their beauty, mills attempt to minimize any waste due to internal cracks, knots, or other imperfections in the tree. This is done now by experts in the mill, using cursory external inspections.

Clearly, CT can provide improved imaging of each cross-section of the log (**Fig. 13**), but the cost of the scanner and the time of inspection may prohibit the use of this technique on the production line. The scan in **Fig. 13** involved 1500 projection angles and took about an hour to produce on our LCAT scanner. We have applied some limited-data reconstruction methods[22] to investigate what could be done in less time. The scan in **Fig. 14** involved 75 projection angles and took less than a minute, and most of the major features are still visible. Further improvements in speed can be done in the future. We are preparing papers and proposals for future funding of this work.

Future Work

The CT project has benefited from strong elements of research, development, technology transfer, and an interesting variety of applications, each of which is essential to our growth. We plan to continue this successful combination of elements in the coming fiscal year. Two major thrusts will be (1) continued upgrades in 3-D measurement and analysis tools, including interfacing to CAD/CAM databases; and (2) incorporation of two, large-format, high-dynamic-range cameras to improve resolution and contrast of our medium-energy scanners. We will continue to improve the CT scanners built by our group and the software tools that extend the range of applications for CT technology. We expect to continue to participate in the transfer of our technology within LLNL, around the DOE nuclear weapons complex, and to commercial areas of investigation, as well as assist in their scaling up to fieldable industrial uses.

Acknowledgements

The authors would like to acknowledge the valuable contributions to this project by J. Haskins, L. Hester, D. Rikard, R. Ryon, and E. Updike, all from the NDE Section at LLNL. We would also like to thank D. Camp (Nuclear Chemistry), J. Stout and B. Qualheim (Environmental Restoration Division), B. Bonner (Earth Sciences), and J.E. Hernandez and J. Kaschmitter (Engineering Research Division) for their support and encouragement. Especially helpful on particular projects are D. Schmoldt (USFS) on the log samples, J.K. Rigby and K. Sauer (University

of Notre Dame) on the archaeological studies, R. Ehrlich and C. Prince (University of South Carolina) on the soil studies, and J. Holscher (Purdue University) on the dimensional measurements.

---

1. H.E. Martz, S.G. Azevedo, D.J. Schneberk, and G.P. Roberson, *Computed Tomography*, Lawrence Livermore National Laboratory, Livermore, California, UCRL-ID-114541 (1993).
2. H.E. Martz, D.J. Schneberk, G.P. Roberson, and S.G. Azevedo, "Computed Tomography," *Engineering Research, Development, and Technology*, Lawrence Livermore National Laboratory, Livermore, California, UCRL-53868-91, 7-6 (1992).
3. H.E. Martz, D.W. Perkins, S.G. Azevedo, D.J. Schneberk, M.F. Skeate, and G.P. Roberson, "Computerized Tomography," *Engineering Research and Development Thrust Area Report FY90*, Lawrence Livermore National Laboratory, Livermore, California, UCRL-53868-90, 7-14 (1991).
4. H.E. Martz, S.G. Azevedo, G.P. Roberson, D.J. Schneberk, and M.F. Skeate, "Computerized Tomography," *Engineering Research and Development Thrust Area Report FY89*, Lawrence Livermore National Laboratory, Livermore, California, UCRL-53868-89, 7-15 (1990).
5. S.G. Azevedo, H.E. Martz, and G.P. Roberson, "Computerized Tomography Reconstruction Technologies," *Energy and Technology Review*, Lawrence Livermore National Laboratory, Livermore, California, UCRL-52000-90-11·12 (November/December 1990).
6. H.E. Martz, G.P. Roberson, and D.J. Schneberk, "Progess in Computerized Tomography Scanners," *Energy and Technology Review*, Lawrence Livermore National Laboratory, Livermore, California, UCRL-52000-90-11·12 (November/December 1990).
7. H.E. Martz, D.J. Schneberk, and G.P. Roberson, "Applications of Computerized Tomography," *Energy and Technology Review*, Lawrence Livermore National Laboratory, Livermore, California, UCRL-52000-90-11·12 (November/December 1990).
8. R.C. Placious, D. Polansky, H. Berger, C. Bueno, C.L. Vosberg, R.A. Betz, and D.J. Rogerson, *Mater. Eval.* **49** (11), 1419 (November 1991).
9. C.M. Prince and R. Ehrlich, *Mathematical Geology* **22**(3), 333 (1990).
10. C.A. McCreesh, R. Ehrlich, and S.J. Crabtree, *AAPG Bull.* **75**(10), 1563 (1991).
11. R. Ehrlich, S.J. Crabtree, K.O. Horkowitz, and J.P. Horkowitz, *AAPG Bull.* **75**(10), 1547 (1993).
12. (a) J. Holscher and S.G. Azevedo, *Dimensional Measurements from Images*, Lawrence Livermore National Laboratory, Livermore, California, UCID in publication(1993).

    (b) G.P. Roberson and H.E. Martz, *Feasibility Study for the Application of Computed Tomography to Savannah River Reactor Fuel Tubes, Phase 2*, Lawrence Livermore National Laboratory, Livermore, California, UCRL in publication (1994).
13. J.E. Hernandez, *Using VISION*, Lawrence Livermore National Laboratory, Livermore, California, UCRL-MA-112377 (1993).
14. K.W. Dolan, R.W. Ryon, D.J. Schneberk, H.E. Martz, and R.D. Rikard, "Explosives Detection Limitations Using Dual-Energy Radiography and Computed Tomography," *Proc. First Int. Symp. Explosives Detection Technology* (Atlantic City, New Jersey), (November 13–15, 1991), 252 (1992); also Lawrence Livermore National Laboratory, Livermore, California, UCRL-JC-108254 (1991).
15. D.C. Camp and H.E. Martz, "Nondestructive Characterization of TRU and LLW Mixed-Wastes Using Active and Passive Gamma-Ray Spectrometry/Computed Tomography," *Proc. Information Exchange Meeting on Characterization, Sensors, and Monitoring Technologies* (Dallas, Texas), (July 15–16, 1992).
16. H.E. Martz, G.P. Roberson, C. Robert-Coutant, D.J. Schneberk, and D.C. Camp, "Quantitative Waste-Form Assay Using Gamma-Ray Computed Tomography," *Proc. Int. Symp. Spectrosc. and Struct. Molec. Nuclei* (Tallahassee, Florida), (March 27–29, 1992); also Lawrence Livermore National Laboratory, Livermore, California, UCRL-JC-110648 (1992).
17. D.J. Schneberk, H.E. Martz, and S.G. Azevedo, "Multiple-Energy Techniques in Industrial Computerized Tomography," *Review of Progress in Quantitative Nondestructive Evaluation*, D.O. Thompson and D.W. Chimenti (Eds.), Plenum Press (New York), 451 (1991).
18. H.E. Martz, D.J. Schneberk, S.G. Azevedo, and S.K. Lynch, "Computerized Tomography of High Explosives," *Nondestructive Characterization of Materials*, Vol. IV, C.O. Ruud *et al.* (Eds.), Plenum Press (New York), 187 (1991).
19. H.E. Martz, G.P. Roberson, D.J. Schneberk, and S.G. Azevedo, "Nuclear-Spectroscopy-Based, First-Generation, Computerized Tomography Scanners," *IEEE Trans. Nucl. Sci.* **38**, 623 (1991).
20. J. Spargo, S. Azevedo, J. Kaschmitter, M. Lewis, A. Silver, and S. Whiteley, *A Pipelined Superconducting Digital Processor*, Lawrence Livermore National Laboratory, Livermore, California, UCRL-JC-114594 (1993).
21. J.L. Kaschmitter, S.G. Azevedo, G.P. Roberson, and D.J. Schneberk, *JJDSP Functional Description*, Lawrence Livermore National Laboratory, Livermore, California, UCRL-CR-116943 (1994).
22. S.G. Azevedo, *Model-Based Computed Tomography for Nondestructive Evaluation*, Ph.D. Dissertation, University of California Davis, Davis, California (1991).

X ray Imaging: Digital Radiography

Kenneth W. Dolan,
Jerry J. Haskins,
Dwight E. Perkins, and
R. Derrill Rikard
Engineering Sciences
Mechanical Engineering

We are developing several x ray imaging systems for digital radiographic data acquisition, digital display, real-time and near real-time monitoring, and digital post-processing that offer alternatives to and advantages over traditional film imaging methods. The digital imaging systems are based on linear diode arrays and area diode arrays that are optically coupled to x ray-sensitive scintillator materials. Large objects can be imaged in near real time with extended linear arrays, and in real time with lens-focused area arrays. Relatively high spatial resolution and high contrast that approach capabilities of film have been obtained. Large-area imaging with somewhat reduced spatial resolution has also been demonstrated.

Introduction

Traditional x-ray radiographic inspection is based on interaction of a time-integrated, spatially modulated, transmitted x ray beam with a film recording medium. Since integration time is of the order of minutes, the traditional film-based method precludes viewing dynamic events, interactive multiple or continuous rotational viewing for near 100% inspection, and high-volume production line and surveillance monitoring. The conventional film-based method is also limited by dynamic range, i.e., the ability to image large variations in thickness or density, typically 100:1 variations with film. Digital imaging systems based on discrete detectors have achieved dynamic ranges of 10000:1. While film has provided unequaled spatial resolution, typically 50 to 100 line pairs/mm or better, new advances in phosphor screens, scintillation glass (monolithic and fiber optic plates), charge-coupled device (CCD) and intensified-CCD cameras, and discrete diode arrays are providing improved spatial resolution for digital radiography that approach film quality. The digital nature, high dynamic range, and improved spatial resolution of digital imaging systems make them strong candidates for automated inspection. Digital image enhancement, interactive workstation evaluation, and archival digital data storage are also strong motivations for developing digital imaging. Evaluation, acquisition, and application of digital imaging hardware and systems are key steps to improving our capabilities in this area.

Progress

Linear Diode Array System

We are investigating system parameters that affect performance of a linear diode array (LDA) digital radiography system to determine potential for programmatic inspection applications. The LDA digital radiography system, constructed from commercial components, consists of a 1024-element, solid-state diode array; 160-kVp, constant-potential x ray tube; fan beam collimator; automated vertical and rotational translation stages; and structural fixturing. Each cell of the linear diode array consists of a photodiode element optically coupled to a scintillator. X ray energy converted to light by the scintillator is detected by the photodiode, and an electronic signal proportional to the scintillator light output is obtained from the photodiode by readout electronics. The LDA contains 16 chip modules with 64 photodiodes on each module. Electronic readout is 12-bit, providing 4000:1 dynamic range. Readout time for one scan line (1024 elements) can be varied from 6 ms to 100 ms. Charge integration time for each element is identical to readout time. The 1024 photodiode elements on 0.45-mm pitch provide a 460-mm (18-in.) long by 0.6-mm (0.024-in.) sensitive area.

Different geometric magnifications of the object under inspection can be obtained by varying the position of the object relative to the source and detector array. With the object positioned midway between source and detector, a 2:1 geometric magnification and nominal spatial resolution of 0.2 mm (0.008 in.) are obtained. Objects of up to 230-mm (9-in.) diameter can be imaged at 2:1 magnification, and up to 460-mm (18-in.) at 1:1 magnification. The vertical stage provides 12 in. of linear translation for digital radiography. The rotational stage can be used to obtain different radiographic views, but is used primarily for obtaining a complete set of rotational views for single-slice computed tomography reconstruction. Readout times of the LDA are fast enough so that single-slice tomography views can be obtained as the object rotates without producing motion blur for slow rotational speeds, e.g., 6 rpm. At 50-ms readout time and 6 rpm, 400 single-slice rotational views are obtained in 20 s.

The system has been used to investigate near-real-time digital imaging and computed tomography applications such as for metal-matrix-composite uniformity and defect studies, for small-explosive sample uniformity and defect studies, for soil sample leach-rate studies, for demonstration of detection of mock explosives in packages and baggage, and for detection of altered electronic components using digital subtraction of before and after images.

Phosphor Screen and Intensified-CCD Camera System

A real-time radiography and digital imaging system constructed for large-area imaging consists of a mosaic array of phosphor screens, an intensified-CCD camera with wide angle lens; a 45-deg, front-surface-silvered mirror; a light-tight enclosure; and lead shielding for the camera. The size of the screen is dictated by the area desired to be imaged. Screen areas of up to 42 × 51 in., composed of nine phosphor screens, and as small as a single, 14-x-17-in. screen have been used.

Screen. The screen is a rare earth phosphor screen of the type used in 14-x-17-in. film packs for film radiography. The phosphor screen was selected for its high brightness at 9-MeV x ray energy, peak light output in the green for compatibility with the cameras, and short persistence for compatibility with dynamic events. Comparison studies previously done at Lawrence Livermore National Laboratory showed that thick gadolinium oxysulfide (GOS) screens met these criteria. Commercially available GOS screens were selected. The screens were epoxy bonded to a sheet of 0.063-in.-thick copper to provide a stiff backing, to convert some of the incident x rays to electrons to improve screen efficiency, and to attenuate some of the lower-energy scattered x rays. The epoxy bond is compatible with low temperature operation (0°C), although no provisions are made to prevent moisture condensation on the screens. A dry air flush would be needed for low-temperature conditions. The screens were mounted in a 3-x-3 mosaic array for large-area imaging. The corners of the screens come from the manufacturer already trimmed with rounded corners, and no attempt was made to square the corners, which would have reduced overall screen dimensions. Instead, the diamond-shaped blank areas at the screen interstices are used as fiducials to obtain exact measurements in the recorded image, and light-emitting diodes are mounted in the blank areas to check camera operation, focusing, and alignment. The copper sheet with screens attached was mounted on a 0.5-in.-thick aluminum plate for attachment to the imager structure.

Cameras. Criteria for camera selection were high sensitivity, high gain, high resolution, anti-blooming, automatic gain, and rugged design. The cameras chosen were intensified-CCD video cameras with a sensitivity of better than 10^{-6} foot-candles, overall gain of 10^6, operational through 5-g shock (5 to 200 Hz), and operational temperature range of 0 to 50°C with storage at −20 to +50°C. Scan rate was 59.94 fields/s interlaced. The intensifier element of this camera is a microchannel plate (MCP) electron multiplier. The automatic gain feature of the camera is accomplished by automatic continuous adjustment of the voltage to the MCP element. The voltage is adjusted to limit the current on the MCP element. The voltage can change within the scanning time of a single frame, but the gain is the same for all lines in the frame. This means that frame-to-frame comparisons require normalization, but that image features within a single frame can be used for calibration within that frame. The CCD has 756 horizontal by 485 vertical usable elements, providing the standard

3/2 video aspect ratio. The non-blooming feature in the readout electronics of the CCD prevents charge from bleeding from next-neighbor elements on readout. Output from the camera is standard RS-170. The cameras are capable of providing a spatial resolution of 756 picture elements (pixels) in the long dimension of the image. However, image digitization is done at 512 × 512 so that the spatial resolution for the large-area imager, determined by the long screen dimension (51 in.) divided by 512, is approximately 0.1 in./pixel.

Optics. The cameras are lens coupled to the imaging screen with a wide angle lens and a 45° turning mirror. The lens for each camera is 25-mm dia, with 12.5-mm focal length and f-stop of f1.4, and provides a large depth of field at the camera-to-screen distance of 66.8 in. The mirror is a front-surface-silvered, 9-x-9-x-1-in. glass. A single camera, or two redundant cameras, are mounted on their sides to use the 3/2 camera aspect ratio to obtain maximum resolution in the long (vertical) screen dimension. Because the cameras are mounted on their sides, all of the video images when displayed on video monitors appear rotated 90°. This is corrected in post-processing digitization of the video images where the 512-x-512 digitized images are rotated 90°, so that the image is displayed normally.

Light Enclosure. The customized light-tight box is fabricated from 3/4-in. plywood. The inside of the light-tight box housing cameras and mirror is painted black and lined with black cloth. The screen assembly is mounted on a rectangular tapered horn. The covering of the horn was made of perforated sheet metal. Black cloth is fitted to the outside of the horn, and provides light-tight coupling of the horn/screen assembly to the camera/mirror enclosure box.

Shielding. The CCD element and other solid-state components in the camera are susceptible to electronic noise and intrinsic damage from x ray exposure. Lead shielding consisting of 8 lead plates, each 1 × 12 × 12 in., is positioned next to the camera to eliminate direct exposure to x rays. Lead plate, 1-in. thick, protected the cameras on all other sides except for the lens side. Lead glass can be used to protect the lens side of the camera enclosure, and although this has not been done to date, no undue camera degradation has been observed.

The system has been used to image 55-gal waste drums, large castings, and a Titan motor section with spatial resolution 0.1 in. in large-area format (42 × 51 in.), and 0.03 in. in small-area format (14 × 17 in.).

High-Density Scintillating Glass and CCD Camera System

New developments with high-density glass scintillators have been used to construct real-time imaging systems from video components for both digital radiography and computed tomography. Commercially developed, terbium-activated glass with higher density than previously available provides improved absorption of x rays and improved brightness. Results are higher image definition as evidenced by improved image contrast and spatial resolution. The characteristic decay time of the new glass is longer (3 ms) as compared to typical phosphor screens, e.g., GOS (1 ms), but is well within the requirements for standard video imaging (30 ms/frame). Effects of variable integration times, focusing methods, reflective coatings, light output as a function of x ray energy, and recovery from radiation damage have been studied.

The high density glass has been used with a variety of cameras and lenses in light-tight enclosures with 45° bending mirror. The cameras are shielded from direct and scattered x ray exposure with lead sheet up to 4 in. thick to reduce x ray-induced noise and component degradation. Lead glass is also used in the optical path to reduce x ray-induced noise.

Output brightness of the glass has been studied in the x ray energy range of 70 keV to 300 keV, with additional data points at 4 MeV and 9 MeV. X ray brightness is nearly linear with output, and is increased by almost a factor of 2 by using a reflective coating on one side of the glass. We have obtained improved contrast and resolution by focusing into the glass rather than just at the back surface (near surface to the camera). Luminescence loss due to radiation exposure has been measured as less than 5% after 3-Mrad exposure, and we have found that 100% recovery of luminescence is obtained by annealing. Resolution with the glass has been measured at 20 line pair/mm at 70 keV and between 2 and 4 line pair/mm at 9 MeV. We have also obtained indications that we may not have obtained optimum resolution and contrast, and that by improving the camera and optics system, additional gains may be possible.

Systems based on the high-density glass scintillator have been used for imaging high-density and large/thick components at 4 and 9 MeV, for high definition of small iron- and nickel-based components (e.g., turbine blades, fuel injectors, pinch welds) at medium (160 to 320 kV) to high energies (4 MeV), and for imaging lower-density or smaller components at lower energies (below 160 kV).

Future Work

We plan to implement improvements in our digital radiography systems and analysis methods for dimensioning from digitized images that satisfy more stringent inspection requirements. Improvements in spatial resolution, contrast, and image analysis will be core aspects of this work. Our goal is to demonstrate practical aspects and usefulness of digital radiography for part dimensioning, detection of changes in internal components or overall part configuration between inspections, flaw detection and quantification, and part certification.

Acknowledgements

We would like to acknowledge the lead technical contributions of D. Schneberk, S. Azevedo, H. Martz, and P. Roberson, and the technical support and contributions of G. Ford and E. Updike.

Signal-to-Noise Improvements for Laser Ultrasonics

Graham H. Thomas and
Albert E. Brown
Engineering Sciences
Mechanical Engineering

Sherman Min and
Larry Peng
Sandia National Laboratories
Livermore, California

We are developing an ultrasonic evaluation technique for contamination-sensitive components in hostile environments. Laser-generated ultrasonics is an attractive alternative to traditional ultrasonic nondestructive evaluation because it allows remote, noncontacting, ultrasonic nondestructive evaluation to be performed. Laser ultrasonics has several other advantages such as broad band excitation, multimode acoustic energy generation, and adaptability to scanning complex shapes. A weakness of laser acoustics is the instrumentation sensitivity, or the signal-to-noise levels produced. We have developed an advanced filtering technique to minimize the noise and improve system sensitivity.

Introduction

Ultrasonic nondestructive evaluation is a valuable technology for material characterization and defect classification. Laser-based ultrasonics allows us to explore many new applications. For example, laser ultrasonics can be performed in hostile environments such as in a furnace or a glove box.

A major drawback for laser acoustics in many applications is its low signal-to-noise ratio compared to the common piezoelectric (PZT) transducer. Our approach to this problem is from a signal processor standpoint with the goal of identifying and applying an algorithm capable of extracting meaningful information from a noisy data set. In this cursory examination, we have designed a model-based estimation, or Kalman, filter in an attempt to improve the signal-to-noise ratio of a single-shot data set.

Progress

Laser acoustics relies upon a laser-induced thermal expansion, or a shock front that weakens into an elastic wave, for the generation of ultrasound.[1] Optical methods such as interferometry are then used to detect this energy.[2] While laser acoustics is considered a fairly mature laboratory technique in research environments, a robust turnkey system has yet to be developed for widespread use in industrial applications,[3,4] due in part to the inherent complexity of optics, but also to the poor sensitivity of optical technology relative to its PZT counterpart.[5] This sensitivity problem is being attacked from several different angles, from signal generation optimization[6–8] to interferometer design modifications. Our approach is to investigate the possible improvements that can be made after the data has been collected, using advanced techniques developed by the signal processing community. In this initial effort, we have designed a filter that incorporates a model of the system into the processing algorithm, thus preserving basic system behaviors while removing white noise.

In signal processing, we are interested in extracting useful information from an input signal while filtering out the unnecessary information. When the noise and/or the signal is contained in a narrow frequency band, a deterministic filter (i.e., a bandpass Butterworth filter) can be used quite effectively. Unlike electromagnetic acoustic transducers, PZT's, and other resonant devices, however, the ultrasonic signature of a laser source is broad band in frequency content. When combined with an optical interferometer that captures events exceeding the rf regime, decoupling noise from signal without significant information loss can become an onerous task. Under proper vibration isolation, rf shielding, and thermal control, the dominant noise component is white interferometer noise. Thus, our system is an ideal candidate for a Kalman filter, a state-space processor based on statistical principles that uses *a priori* knowledge of the system to remove random disturbances.

In contrast to deterministic filters that deal with repeatable signals, model-based estimation filters handle measurements that have been corrupted by random disturbances (non-repeatable signals). This class of filters is often referred to as optimal model-based processors in the sense that mathematical models have been used to characterize the process generating the signal, and the processor has been designed to produce a signal estimate that results in a minimum error variance. From an energy standpoint, a Kalman filter views a noise spike as an impulsive system response and thus rejects such fluctuations, producing an uncorrelated sequence called the innovations. The whiteness of the innovations sequence is also a measure of the filter's optimality.

Future Work

We are improving our laser acoustic facility and demonstrating its applicability to ultrasonic inspection problems. We are investigating applications of laser acoustics to nondestructive evaluation problems at Lawrence Livermore National Laboratory and for industries in the United States. Specifically, we will continue to explore applications for laser acoustics to control selected manufacturing processes, such as welder alignment, composite curing, and plutonium processing. Since each application of laser-generated ultrasonics entails a customized system, we need to have a thorough understanding of the fundamental capabilities and limitations of the technology to design the optimal inspection facility.

---

1. C.B. Scruby, R.J. Dewhurst, D.A. Hutchins, and S.B. Palmer, "Laser Generation of Ultrasound in Metals," *Research Techniques in Nondestructive Testing*, Vol. V, Academic Press (London, England), 281-327, 1982.
2. J.P. Monchalin, *IEEE Trans. Ultrason. Ferroelec. Freq. Contrl.* **33**(5), (September 1986).
3. General Dynamics, in conjunction with UltraOptec, a Canadian manufacturer of laser equipment, has reportedly developed a scanning laser-ultrasonic system for the inspection of composites.
4. Surface Combustion, Inc. and Textron Defense Systems presented their Workpiece Analyzer (WPA), a laser-based instrument for the measurement of temperature in industrial furnaces, at the Sixth International Symposium on Nondestructive Characterization of Materials, Oahu, Hawaii (June 1993).
5. J.W. Wagner, "Breaking the Sensitivity Barrier: The Challenge for Laser-Ultrasonics," IEEE Ultrasonics Symposium (October 1992).
6. T.W. Murray, J.S. Steckenrider, and J.W. Wagner, "Practical Systems for Pulsed Laser Array Generation of Ultrasound," Sixth International Symposium on Nondestructive Characterization of Materials, Oahu, Hawaii (June 1993).
7. H. Nishino and Y. Tsukahara, *Appl. Phys. Lett.* **62** (17), (April 1993).
8. C.E. Duffer, P.L. Peterman, and C.P. Burger, "Enhancement of Laser-Induced Surface Acoustic Waves for NDE," SEM Spring Conference on Experimental Mechanics (1993).

Infrared Computed Tomography

Nancy A. Del Grande and
Philip F. Durbin
Engineering Sciences
Mechanical Engineering

Michael R. Gorvad
Applications Systems Division
Computation Directorate

To find hidden corrosion, we are developing infrared computed tomography (IRCT) for slowing down costly damage to airframes, bridge decks, and pipelines. We use IRCT, with two thermal infrared bands, to provide time sequences of high-contrast temperature maps for flash-heated structures. These maps unfold the location and amount of hidden corrosion damage. During FY-93, we used IRCT to locate corrosion pockets within a Boeing 737 airframe. During FY-94, we plan to apply IRCT for corrosion detection within concrete bridge decks and exposed hazardous product pipelines.

Introduction

During the well-publicized Aloha Airlines accident, damage from corrosion caused the aircraft skin to rip apart from the Boeing 737 airframe. To maintain aircraft reliability, the Air Force spends over $700 M annually for corrosion-related maintenance actions. To mitigate airframe corrosion damage, the Federal Aviation Administration (FAA) Technical Center sponsors research and development within the Aging Aircraft Nondestructive Inspection (NDI) Technology Research and Development Program.

In 1993, the FAA Aging Aircraft NDI Technology Program supported thermal infrared imaging efforts at Wayne State University (WSU) and Lawrence Livermore National Laboratory (LLNL). WSU uses a single-band infrared (SBIR) technique for thermal-wave imaging of corrosion and disbonds within flash-heated aircraft structures.[1] At LLNL, we use a dual-band infrared (DBIR) technique for infrared computed tomography (IRCT) to provide quantitative corrosion damage detection within flash-heated airframes.[2–5] We use IRCT, with DBIR image ratios, to improve temperature contrast, remove clutter, and quantify corrosion damage within aircraft structures.

Our advanced development approach[6] would provide an early warning of low-level (5% or 6%) material-loss effects from corrosion thinning. The IRCT technique, which uses DBIR image ratios, is expected to establish LLNL's lead in NDI corrosion detection technology by meeting these stringent requirements. If successful, IRCT technology will provide a rapid, reliable, wide-area imaging, corrosion-detection tool for *in situ*, non-intrusive inspections of large structures. This tool, which does not presently exist, would address the need for early warning signals of corrosion damage and minimize the cost for corrosion repair.

Progress

During FY-93, we inspected the Boeing 737 Jetliner at the FAA/AANC (Aging Aircraft NDI Center) Test Bed Facility at Sandia National Laboratories in Albuquerque, New Mexico. We detected and quantified corrosion-related thickness losses that occurred within the lap splice of the flash-heated Boeing 737 aircraft skin. Our dynamic temperature and thermal-inertia maps depicted corrosion pockets inside a lap splice that was situated beneath the galley and the latrine. We photographed 'pillowing' at this site, which implied the expansion of corrosion by-products within the riveted lap splice. This report discusses our mapping procedures and the results of this demonstration.[7]

Emissivity-Corrected Temperature Maps

Using DBIR image ratios (from DBIR cameras that scan infrared at 3–5 μm and 8–12 μm), we enhance surface temperature contrast and remove

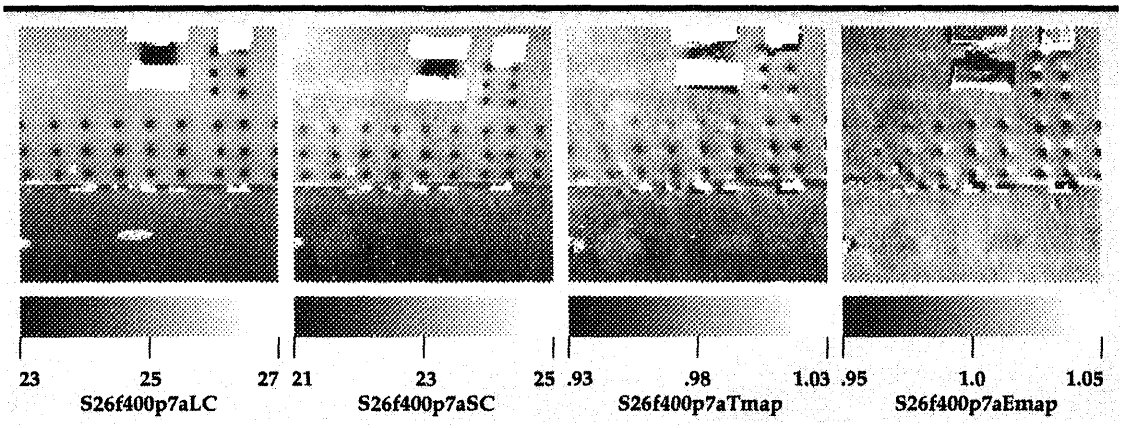

S26f400p7aLC S26f400p7aSC S26f400p7aTmap S26f400p7aEmap

***Figure 1.** Maps (left to right) of the corroded Boeing 737 lap splice structure, showing 8–12 μm and 3–5 μm apparent temperatures, dual-band infrared enhanced contrast temperatures (relative scale), and E-ratios (relative scale). Note apparent temperature and E-ratio patterns for black cloth tape (white, hot) and metal tape (black, cold) markers at top center, and for black paper tape marker (white, hot) at top right corner. Also, note E-ratio patterns for non-uniformly applied black paint at bottom.*

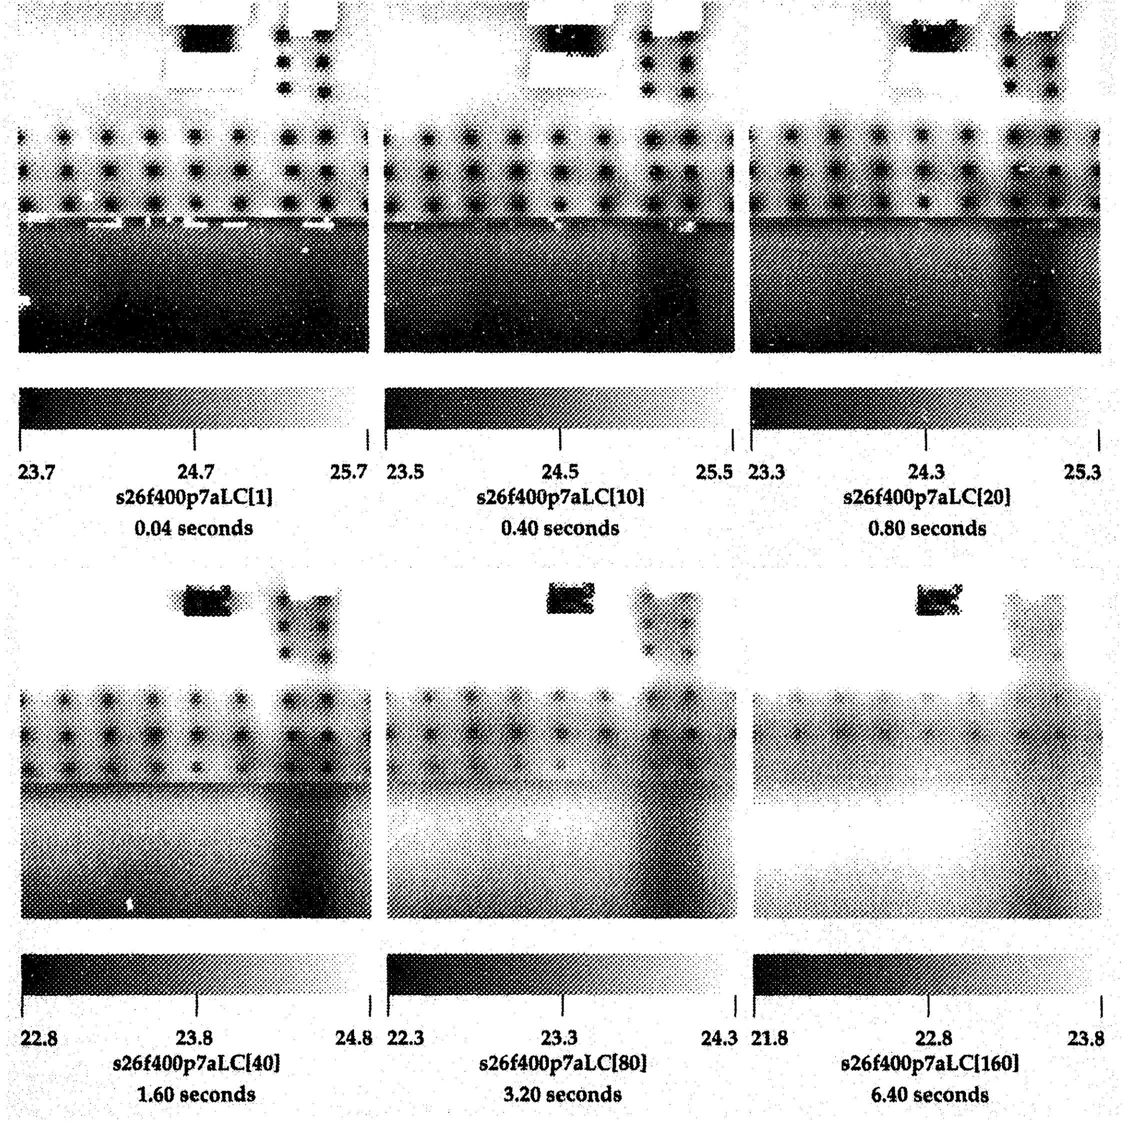

s26f400p7aLC[1]
0.04 seconds

s26f400p7aLC[10]
0.40 seconds

s26f400p7aLC[20]
0.80 seconds

s26f400p7aLC[40]
1.60 seconds

s26f400p7aLC[80]
3.20 seconds

s26f400p7aLC[160]
6.40 seconds

***Figure 2.** Dynamic temperature variations in °C for the Boeing 737 lap splice at different times after the heat flash, based on the 8- to 12-μm apparent temperature maps. Note warmer (lighter) corroded area within lap splice, near center of image.*

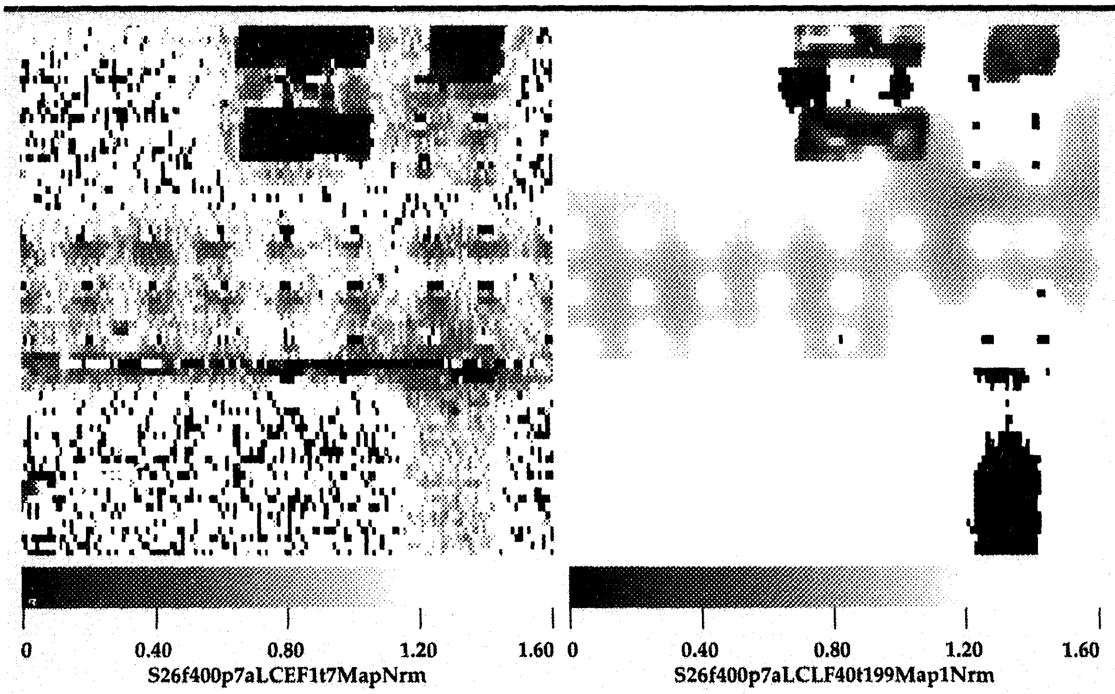

Figure 3. ***Composite thermal inertia maps of the Boeing 737 airframe lap splice. Note the dark 'butterfly-like' defect site (at upper right) on the late-time thermal inertia map where corrosion entered the lap splice from inside the aircraft, beneath the galley and the latrine. Also, note the low thermal inertia values for front-surface cloth and paper tape markers (at top center and right) and back surface tear straps (at bottom right).***

the mask of surface emissivity clutter (from dirt, dents, tape, markings, ink, sealants, uneven paint, paint stripper, exposed metal, and roughness variations). This clarifies interpretation of subtle heat flow anomalies associated with hidden defects and corrosion. We compute DBIR image ratios of enhanced temperature contrast (T^5) and emissivity-ratio (E-ratio) maps, based on an expansion of Planck's radiation law,[8] which has been used successfully for other applications[9–20]:

$$(T/T_{av})^5 = (S/S_{av})/(L/L_{av}) \qquad (1)$$

$$\text{E - ratio} = (L/L_{av})^2/(S/S_{av}), \qquad (2)$$

where **S** is the short-wavelength intensity (e.g., I_5), S_{av} is the average value of the pixels in **S**, **L** is the long wavelength intensity (e.g., I_{10}), and L_{av} is the average value of the pixels in **L**. See **Fig. 1** for the Boeing 737 images of the above temperature and E-ratio maps that allowed us to distinguish corrosion-related thickness loss effects and surface emissivity clutter.

Dynamic Temperature Changes for Flash-Heated Lap Splice

Corrosion within the Boeing 737 (epoxy-bonded) lap splice causes disbonding. Trapped by-products act as an insulator, delaying heat transfer by conduction from the front to the back surface. This effect is shown in **Fig. 2** by the near-constant temperature contrast from 0.4 to 3.2 s, based on our measurements of the FAA/AANC-owned Boeing 737 aircraft that we inspected.

Corrosion-related material loss effects are best measured at 0.4 s after the heat flash. Temperatures at 0.4 s are sensitive to material loss effects within a lap splice, and insensitive to timing uncertainties. The timing is early enough to provide a good temperature contrast for sites with and without material loss from corrosion. At later times, trapped materials mask the temperature-time history that characterizes material-loss effects for corrosion sites.

We established the correlation between percent thickness loss and above-ambient surface temperature rise, at 0.4 s after the heat flash, based on our measurements for five specimens, which averaged a 24±5% thickness loss per °C temperature rise. These specimens included an F-18 partially corroded wing box (with a 2.9-mm uncorroded thickness) and four 1.0-mm- to 3.9-mm-thickness aluminum panels with milled flat-bottom holes that had thickness losses ranging from 6% to 62%.

Thermal Inertia Mapping

We developed thermal inertia (or effusivity) maps, which had been used previously for other applications,[21,22] by solving the heat transfer equation for a thick panel with an instantaneous surface heat flux[23,24]:

$$T(x,t)=\frac{q}{\sqrt{4\pi k\rho c t}}\exp\left(-\frac{x^2}{4\alpha t}\right), \qquad (3)$$

where T is temperature, x is the distance from the surface, k is thermal conductivity, ρ is density, c is heat capacity, α is thermal diffusivity, t is time, and q is the surface heat flux. For a semi-infinite solid approximation, the surface temperature is proportional to the inverse square root of time. In practice, we map the fuselage composite thermal inertia, $(k\pi c)^{1/2}$, based on the inverse slope of the surface temperature vs inverse square root of time. Composite thermal inertia maps characterize shallow skin defects within the lap splice at early times (< 0.3 s), and deeper skin defects within the lap splice at late times (> 0.4 s). Late-time-composite thermal inertia maps depict where corrosion-related thickness losses occur.

We note in **Fig. 3** (right side) a butterfly-like pattern at the upper right side. This is where the corrosive activity invaded the inside (upper lap edge) of the Boeing 737 lap splice on Stringer 26, near station F400.7, beneath the galley and the latrine. Typical visible signs of corrosion were also evident (e.g., pillowing), resulting from the increased volume of corrosion by-products within the lap splice, between the rivet heads.

Future Work

The use of E-ratio maps to identify (then remove) clutter sites on co-registered temperature maps is unique at LLNL. It significantly improves the signal-to-noise ratio (S/N) for siting corrosion damage, compared to the S/N achievable with SBIR methods. It clarifies interpretation of corrosion damage by removing the mask from clutter.

We have also inspected panels removed from the airframe of a Boeing KC-135 at Tinker Air Force Base in Oklahoma. The Boeing KC-135 lap splices had ripples that produced 1.2°C surface temperature rises, 0.4 s after the heat flash. To distinguish corrosion from ripple defects, we will use temperature timegrams, taken every 550 µs, from 0 to 120 ms after the heat flash. We expect temperature timegrams to distinguish between defect types, reducing the cost associated with false corrosion calls.

Our future plans are to apply IRCT corrosion inspection to address concrete bridge deck issues, which concern the Federal Highway Administration, and exposed pipeline issues, which concern Dow Chemical.

Acknowledgements

We acknowledge support from the FAA Aging Aircraft NDI R&D Program and from the Nondestructive Evaluation (NDE) Thrust Area at LLNL. We thank C. Seher, P.K. Bhagat, and D. Galella at the FAA Technical Center for their helpful suggestions. We are grateful for the support of G. Phipps, C. Jones, D. Harmon, and P. Walter at Sandia National Laboratories, Albuquerque, New Mexico. We acknowledge A. Shapiro (LLNL) for adapting TOPAZ3D to model heat transfer for aircraft structures. We appreciate the support of S. Kulkarni, NDE Section Leader, H. Martz, NDE Thrust Area Leader, and K. Dolan, FAA Project Leader.

---

1. L.D. Favro, P.K. Kuo, R.L. Thomas, T. Ahmed, and Y.X. Wang, "Thermal Wave Imaging of Corrosion and Disbonds in Aircraft Structures," *Proc. SPIE Conf. 2001: Nondestructive Inspection of Aging Aircraft*, (San Diego, California), M.T. Valley, N.K. Del Grande, and A.S. Kobayashi (Eds.), (1993).
2. P.F. Durbin, N.K. Del Grande, K.W. Dolan, D.E. Perkins, and A.B. Shapiro, "Dual-Band Infrared Thermography for Quantitative Nondestructive Evaluation," *Joint Army, Navy, NASA, Air Force (JANNAF) NDE Subcommittee Meeting Proc.*, (April 1993).
3. N.K. Del Grande, K.W. Dolan, P.F. Durbin, M.R. Gorvad, B.T. Kornblum, D.E. Perkins, D.J. Schneberk, and A.B. Shapiro, "Three-Dimensional Dynamic Thermal Imaging of Structural Flaws by Dual-band Infrared Computed Tomography," *Proc. SPIE Conf. 1942: Underground and Obscured Object Imaging and Detection*, (Orlando, Florida), N.K. Del Grande, I. Cindrich, and P. Johnson (Eds.), (1993).
4. N.K. Del Grande, K.W. Dolan, P.F. Durbin, M.R. Gorvad, and A.B. Shapiro, "Dynamic Thermal Tomography for Nondestructive Inspection of Aging Aircraft," *Proc. SPIE Conf. 2001: Nondestructive Inspection of Aging Aircraft*, (San Diego, California), M.T. Valley, N.K. Del Grande, and A.S. Kobayashi (Eds.), (1993).
5. N.K. Del Grande, "Dual-Band Infrared Imaging for Quantitative Corrosion Detection in Aging Aircraft," *Proc. American Society for Nondestructive Testing (ASNT) Meeting*, (November 1993).
6. N.K. Del Grande, K.W. Dolan, P.F. Durbin, and D.E. Perkins, "Emissivity-Corrected Infrared Method for Imaging Anomalous Structural Heat Flows" (patent pending), 1993.
7. N.K. Del Grande, K.W. Dolan, P.F. Durbin, M.R. Gorvad, and A.B. Shapiro, *Dual-Band Infrared (DBIR) Imaging Inspections of Boeing 737 and KC-135*

Aircraft Panels, Lawrence Livermore National Laboratory, Livermore, California, UCRL-CR-115237 (1993).

8. L.A. LeSchack and N.K. Del Grande, [see Appendix for derivation of Eq. 1 and Eq. 2], *Geophysics* **41**, 1318 (1976).
9. P.F. Durbin and N.K. Del Grande, "Dual-Band Infrared Imaging For Concrete Bridge Deck Inspection," *Proc. Structural Materials Technology NDT Conf.* (to be published in 1994).
10. N.K. Del Grande, P.F. Durbin, and D.E. Perkins, "Dual-Band Infrared Imaging Applications: Locating Buried Minefields, Mapping Sea Ice, and Inspecting Aging Aircraft," *Review of Progress in Quantitative Nondestructive Evaluation*, D.O. Thompson and D.E. Chimenti (Eds.), Plenum Press (New York, New York), **12A**, 465 (1993).
11. J.E. Lewis, N.K. Del Grande, I. McKendry, P.F. Durbin, and M. Lepparanta, "Thermal Mapping," *ERS-1 Baltic Sea Ice Calibration/Validation Post-Experiment Report/Pipor/Finland*, Finnish Institute of Marine Research, Helsinki, Finland, Report 1992 (9), M. Lepparanta and M. Lensu (Eds.), (1992).
12. N.K. Del Grande, P.F. Durbin, M.R. Gorvad, D.E. Perkins, G.A. Clark, J. E. Hernandez, and R.J. Sherwood, "Dual-Band Infrared Capabilities for Imaging Buried Object Sites," *Proc. SPIE Conf. 1942: Underground and Obscured Object Imaging and Detection*, (Orlando, Florida), N.K. Del Grande, I. Cindrich, and P. Johnson (Eds.), (1993).
13. G.A. Clark, J.E. Hernandez, S.K. Sengupta, R.J. Sherwood, P.C. Schaich, M.R. Buhl, R.J. Kane, M.J. Barth, and N.K. Del Grande, "Sensor Feature Fusion for Detecting Buried Objects," *Proc. SPIE Conf. 1942: Underground and Obscured Object Imaging and Detection*, (Orlando, Florida), N.K. Del Grande, I. Cindrich, and P. Johnson (Eds.), (1993).
14. N.K. Del Grande, "Airborne Detection Of Buried Minefields," *Energy and Technology Review*, Lawrence Livermore National Laboratory, Livermore, California, UCRL-52000-91-12 (1991).
15. N.K. Del Grande, G.A. Clark, P.F. Durbin, D.J. Fields, J.E. Hernandez, and R.J. Sherwood, *Surveillance Technologies* **1479**, 335 (1991).
16. N.K. Del Grande, "Sensor Fusion Methodology for Remote Detection of Buried Land Mines," *Proc. 3rd. National Symposium on Sensor Fusion*, Vol. 1, (Orlando, Florida), IIAC/ERIM, 407 (August 1990).
17. N.K. Del Grande, "Temperature Evaluated Mine Position Survey (TEMPS) Application of Dual-Band Infrared Methodology," *Proc. 1990 Meeting IRIS Specialty Group on Passive Sensors*, IRIA/ERIM-sponsored symposium (March 1990).
18. N.K. Del Grande, *Geothermal Res. Council Bull.* **14**, 3 (March 1985).
19. N.K. Del Grande, *Geothermal Res. Council Trans.* **5**, 71 (1978).
20. N.K. Del Grande, "Method for Identifying Anomalous Terrestrial Heat Flows," Patent No. 4,005,289 (1977).
21. E.G. Grinzato, C. Bressan, P.G. Bison, A. Mazzoldi, P. Baggio, and C. Bonacina, "Evaluation of Moisture Content in Porous Material by Dynamic Energy Balance," *Proc. SPIE Conf. 1682: Thermosense XIV*, (Orlando, Florida), J.K. Eklund (Ed.), 213 (1992).
22. P.M. Delpech, D.M. Boscher, F. Lepoutre, A.A. Deom, and D.L. Balageas, "Quantitative Nondestructive Evaluation of Carbon-Carbon Composites by Pulsed Infrared Thermography," *Review of Progress in Quantitative Nondestructive Evaluation*, Vol. 12B, D.O. Thompson and D.E. Chimenti (Eds.), Plenum Press (New York and London), 1297 (1993).
23. H.S. Carslaw and J.C. Jaeger, *Conduction of Heat in Solids*, 2nd Ed., Oxford University Press (London, England), 101, 112, and 259 (1980).
24. J.W.M. Spicer, "Thermographic NDT," *SPIE Short Course Notes, SC37, Thermosense '93*, Section 5, 13 (1993).

Remote Sensing, Imaging, and Signal Engineering

Signal and image processing, which have always been important support for existing programs at Lawrence Livermore National Laboratory (LLNL), are becoming central to the formation of new programs.

The Remote Sensing, Imaging, and Signal Engineering (RISE) thrust area has been very active in working to define new directions.

We also maintain and continue to build our technical base in signal and image processing in support of existing programs, through such applications as diagnostic image processing, multi-aperture impulse radar, image sequence analysis, and adaptive optics technology.

Over the past several years, RISE has developed a series of computer software systems for signal and image processing. These systems provide portability among the many computer systems used at LLNL and give us a platform for transferring the results of specific research and development projects to application areas. Our major signal- and image-processing systems, VIEW and VISION, are used by several major LLNL programs and have been distributed to many university, industry, and government sites.

Work in RISE involves a diverse set of sciences and technologies ranging from optical physics to microbiology to advanced computer architectures. Collaboration with other thrust areas, such as Nondestructive Evaluation and Computational Electronics and Electromagnetics, and with other LLNL organizations, such as the Physics Department and the Biomedical Sciences Division, is central to our continuing work in innovative imaging and signal-processing applications.

Dennis M. Goodman
Thrust Area Leader

Section 9

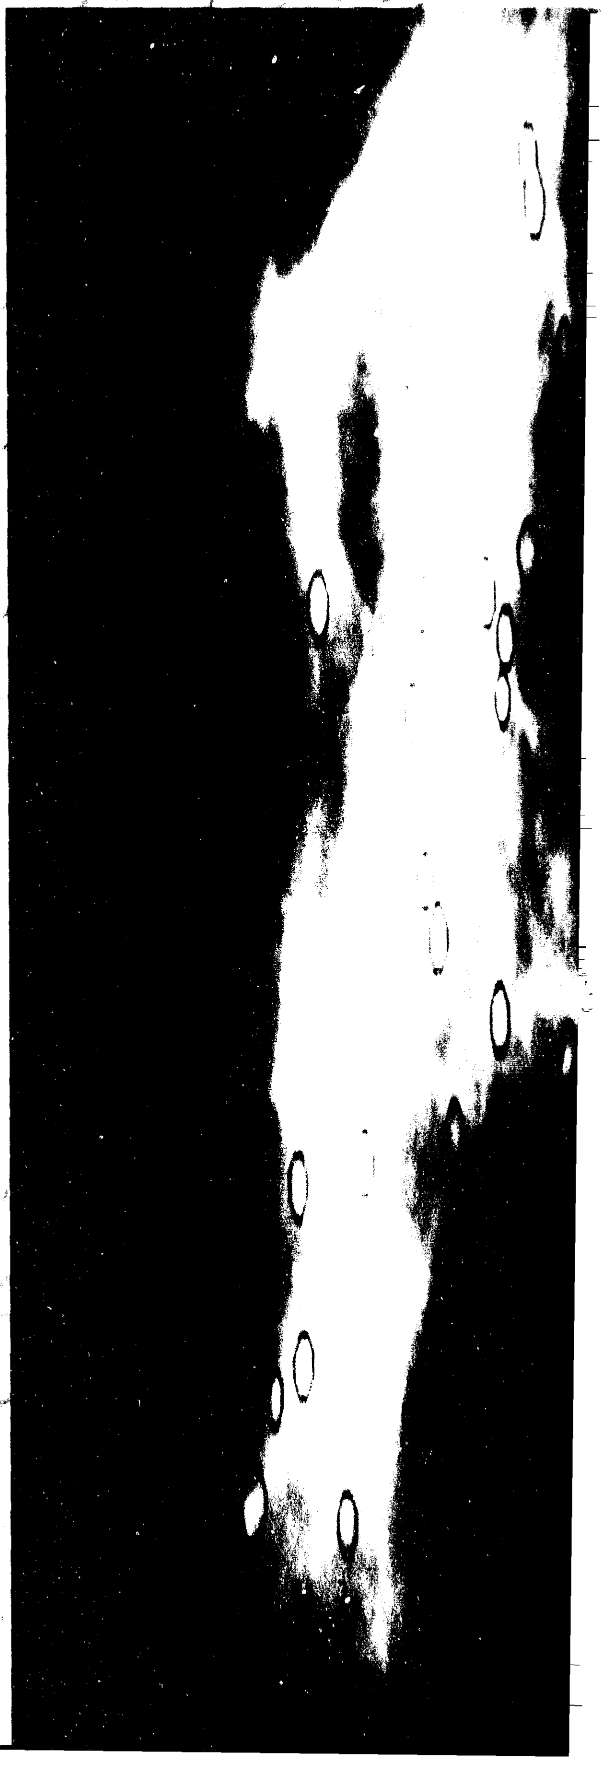

9. Remote Sensing, Imaging, and Signal Engineering

VISION Development

Jose E. Hernandez,
Robert J. Sherwood, and
Scott R. Whitman
Engineering Research Division
Electronics Engineering

VISION is a flexible and extensible object-oriented programming environment for prototyping computer-vision and pattern-recognition algorithms. This year's effort focused on three major areas: documentation, graphics, and support for new applications.

Introduction

For the past three years, several members of the Engineering Research Division at Lawrence Livermore National Laboratory (LLNL) have been developing an object-oriented programming environment for computer vision and pattern recognition known as VISION. VISION is a hybrid system consisting of more than 100,000 lines of code written in C, CLOS,[1–3] and FORTRAN. (CLOS, the Common Lisp Object System, provides object-oriented programming capabilities within the Common Lisp language.) The system was developed primarily to provide a technology base at LLNL in computer vision and pattern recognition. VISION provides a rich and flexible programming environment consisting of more than 20 data structures and over 100 functions advertised to the user for providing a high-level embedded language useful for data analysis, computer vision, and pattern recognition. (An embedded language is a language built on top of another language, in this case Lisp, and typically specialized for a particular domain, in this case, computer vision and pattern recognition.) Users can program algorithms directly in C or FORTRAN and link them at run-time with VISION.

VISION has been used as a prototyping environment for several important projects at LLNL. Some examples include: the Airborne Standoff Mine Detection Project for detecting buried mines; several projects in stereo vision and robotics grasp planning for sorting hazardous waste; the INSENS Project for developing pattern-recognition algorithms for detecting illegal border-crossing activities; and the Deadeye Project for detecting and tracking bullets in flight.

VISION has also been used as a prototyping environment by several students participating in LLNL student summer programs. For example, a student from California State University at Fresno developed an iterative algorithm for inserting a peg into a hole, using an approach that combines vision with force sensing; a student from the University of California at Berkeley developed an algorithm for stereo vision, using simulated annealing techniques; a student from Stanford University developed a prototype of a web-inspection system for fabric inspection (see **Fig. 1**).

Finally, VISION has been very successful in generating interest in collaborative ventures. Last year, we used VISION to collaborate with EG&G Astrophysics in developing pattern-recognition algorithms for detecting explosives in luggage for airport security. VISION also led to our participation in the Computer Aided Fabric Evaluation (CAFE) Project under the AMTEX Cooperative Research and Development Agreement. In collaboration with other national laboratories and the American textile industry, LLNL will be developing several web-inspection systems over the next five years for automating the inspection of textile fabrics.

Progress

Documentation

About 75% of the capabilities in VISION have been documented in a users manual.[4] The manual is available as a Postscript file and can be printed directly by users, using most standard laser printers. Also, a hypertext version of the manual is available on-line. This hypertext document allows users to navigate quickly through the manual and find documentation and examples for the many capabilities in VISION. Source

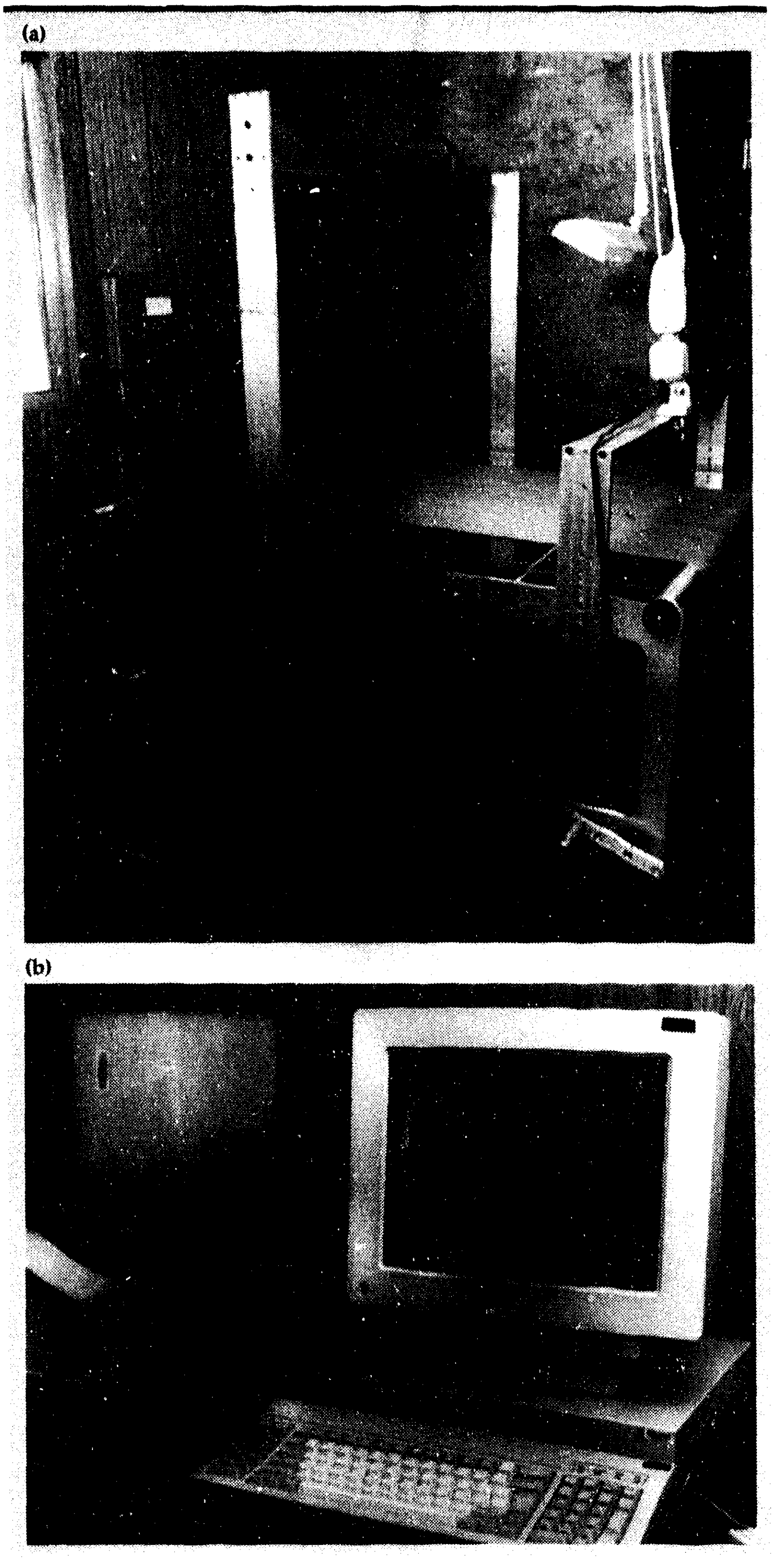

Figure l. A vision-based, web-inspection system for fabric inspection. This prototype was developed under the VISION environment.

code can also be found easily using GNU Emacs'[5] etags capability (see chapter 3 of the VISION manual for more details).

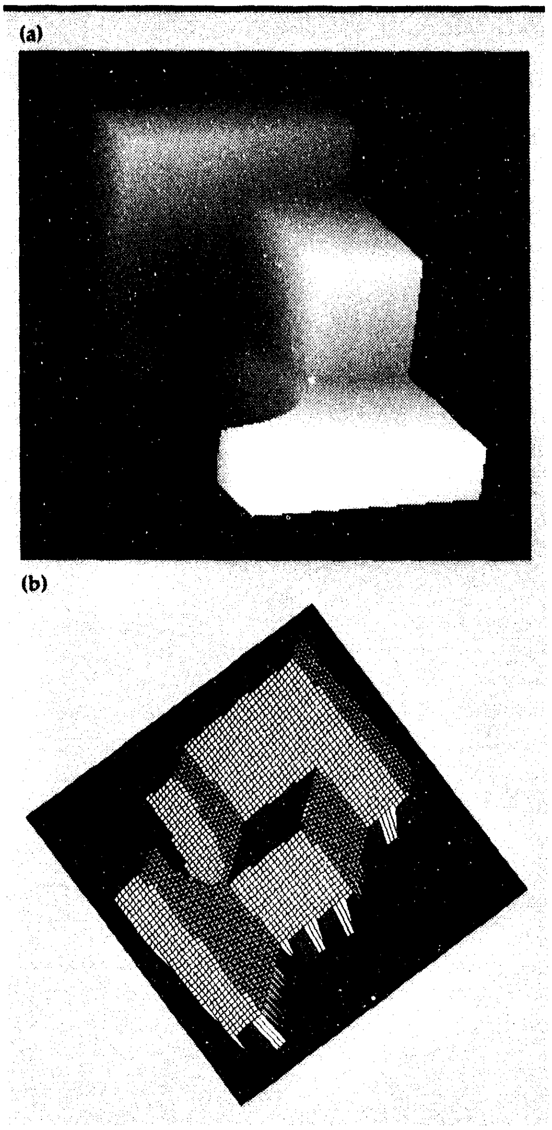

Figure 2. Illustration of 2.5-D graphics capabilities developed to aid in the analysis of range data for machine-vision applications: (a) range data of a 3-D object; (b) segmented range data. Different gray levels are used to visualize the segmented surface patches; the wire mesh preserves the range information.

Graphics

The graphics capabilities in VISION are currently based primarily on two different software libraries. DIGLIB,[6] a device-independent graphics library, forms the foundation for all the one-dimensional and two-dimensional (2-D) scatter plots in VISION, which are typically used in signal-processing and pattern-recognition applications. DispX11, a set of C functions written on top of Xlib, provides the foundation for many of the graphics capabilities in VISION for image processing and image analysis. This year, we were able to improve the integration of these two libraries and provide a more uniform interface to VISION users. This was achieved by developing a new driver for DIGLIB capable of using multiple windows created and managed by the DispX11 software.

Finally, several new graphics capabilities were added to VISION to visualize segmented 2.5-D data (i.e., data representing the surface of a three-dimensional object) and three-dimensional (3-D) data. To minimize the development effort, we took advantage of several software modules from the image-processing package VIEW[7] for providing the low-level functionality of the new graphics functions.

The 2.5-D graphics capabilities were developed to aid in the analysis of range data for machine vision applications. These new capabilities were used on a project for developing grasp-planning algorithms for picking unknown piled objects with a robot.[8] An example of the new 2.5-D graphics capabilities is shown in **Fig. 2**. The new 3-D graphics capabilities were developed to aid in the analysis of volumetric data. These new capabilities were used to aid in the development of pattern-recognition algorithms for detecting explosives in luggage for aviation security. An example of the new 3-D graphics capabilities is shown in **Fig. 3**.

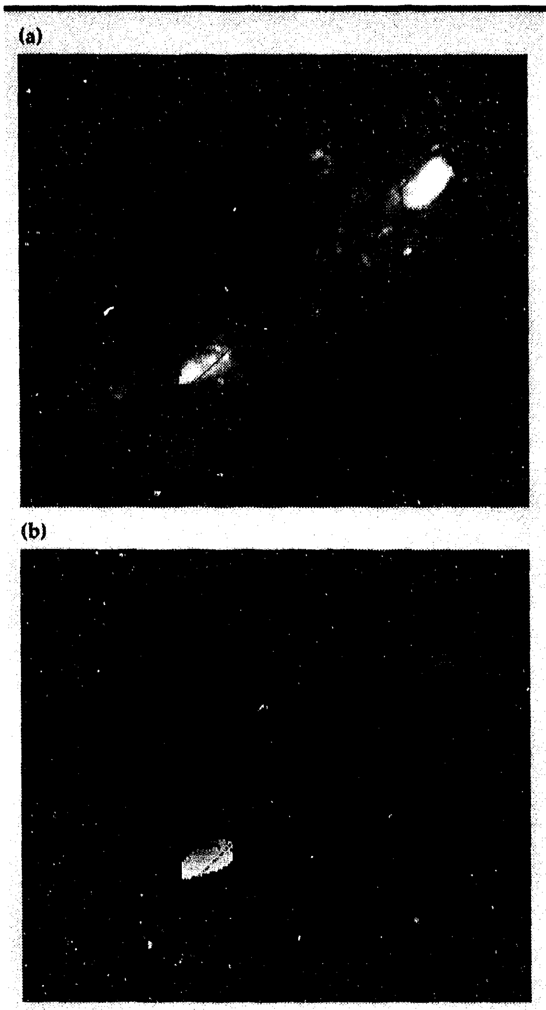

Figure 3. Illustration of 3-D segmentation techniques used in an algorithm for detecting explosives in luggage for aviation security: (a) 3-D reconstruction of a suitcase from an x ray system; (b) potential bomb threat objects are segmented in the 3-D data. Different gray levels are used to visualize the voxels for each object.

Future Work

We expect that new capabilities will continue to be added to VISION as we support new projects. Our current emphasis is in the design and development of real-time image-processing and pattern-recognition systems. VISION will still continue to play an important role in this new direction by providing an environment for quickly prototyping and testing algorithms before committing to a real-time implementation. The development of the Deadeye system has been a good example of this development process.

Acknowledgements

The authors want to acknowledge P. Fitch who helped fund the development of the web-inspection demonstration. Thanks are also extended to S. Lu who helped form the new Machine Vision Laboratory, which houses the web-inspection system for the CAFE project; B. Lawver who developed the data-acquisition software for the web-inspection system; and J. Brase and T. Yorkey, the developers of the software modules in VIEW.

---

1. G.L. Steele, Jr., *Common Lisp: The Language, 2nd Edition*, Digital Press (Burlington, Massachusetts), 1990.
2. J.A. Lawless and N.M. Miller, *Understanding CLOS: The Common Lisp Object System*, Digital Press (Burlington, Massachusetts), 1991.
3. S.A. Keene, *Object-Oriented Programming in Common Lisp*, Addison-Wesley (Reading, Massachusetts), 1989.
4. J.E. Hernandez, *Using VISION—User Manual*, Lawrence Livermore National Laboratory, Livermore, California, UCRL-MA-112337 DRAFT (1993).
5. R. Stallman, *GNU Emacs Manual*, Free Software Foundation.
6. H. Brand, *DIGLIB, Device Independent Graphics Library*, Lawrence Livermore National Laboratory, Livermore, California, (unpublished, 1985).
7. H. Szoke, J. Brase, V. Miller, and M. Wieting, *The VIEW Signal and Image Processing System*, Lawrence Livermore National Laboratory, Livermore, California, UCID-21368 (1988).
8. R.K. Johnson, "Vision-Based Autonomous Grasping of Unknown Piled Objects," *Engineering Research, Development, and Technology*, Lawrence Livermore National Laboratory, Livermore, California, UCRL-53868-93 (1994).

Imaging Algorithms for Synthetic Aperture Ultra-Wideband Radar

Erik M. Johansson
Laser Engineering Division
Electronics Engineering

Jeffrey E. Mast
Defense Sciences Engineering Division
Electronics Engineering

Synthetic aperture ultra-wideband radar imaging is an important emerging technology with the potential to produce high-resolution images for a variety of applications, including the nondestructive evaluation of civil structures, minefield clearing, unexploded ordnance detection and removal, hazardous waste identification, and detection of miscellaneous buried objects. In FY-93, we developed imaging software for two- and three-dimensional problems for several geometries, and produced high-resolution images of a laboratory concrete test-bed and of a simulated mine field located at Lawrence Livermore National Laboratory's Nevada Test Site.

Introduction

Ultra-wideband (UWB) radar is unique in that it does not use a carrier frequency to transmit signals. Instead, extremely short electromagnetic pulses are transmitted. Assuming the target or area being illuminated is composed of point-like scatterers, the received signal consists of a sum of scaled and delayed versions of the transmitted signal. This is known as pulse-echo radar.[1] The received signals give range information about the location of targets, but no azimuth information other than that deduced from knowledge of the antenna beam width and area of illumination. However, when multiple UWB antennas are used (or synthesized by moving a single antenna system), it is possible to determine azimuth information about a target as well.

As an antenna operating in monostatic mode (where the transmitter and receiver are collocated) is moved past a single point-target in a homogeneous transmission medium, the received pulse-echoes form a hyperbolic arc in space and time (see **Fig. 1**). For a single target, it is easy to determine the location of the target from the peak of the arc and knowledge of the characteristics of the transmission medium. However, for multiple targets in more complex media, the space–time image is quite difficult to interpret, and gives little insight as to the nature of the target area. By focusing each of the arcs resulting from a point scatterer into a single point, it is possible to form an accurate spatial image of the target area (see **Fig. 2**). We have developed two-dimensional (2-D) and three-dimensional (3-D) algorithms for several imaging geometries and implemented software that performs this focusing, i.e., translation from raw spatiotemporal data to high-resolution spatial images.

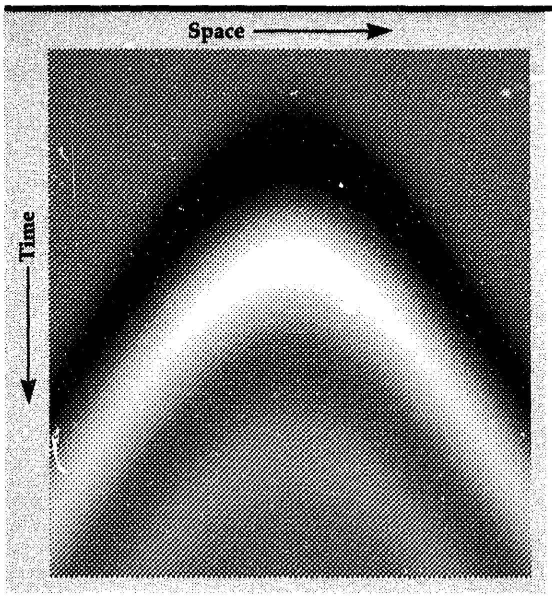

Figure 1.
Raw pulse-echo radar data from a monostatic system as an aperture is synthesized by moving an antenna past a single point-target. The pulse-echoes form a hyperbolic arc in space and time.

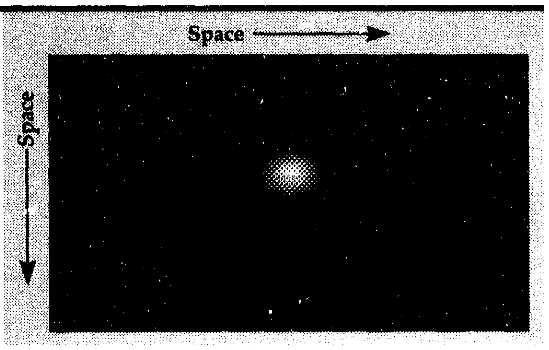

***Figure 2.* The 2-D image formed by 'focusing' the pulse-echoes shown in Fig. 1.**

Progress

In our algorithm development, we have taken two approaches to the image formation problem, a time-domain solution and a frequency-domain solution. Both methods produce similar results, giving us the flexibility of choosing the solution that is best suited to the problem at hand. The time-domain solution uses a delay-and-sum beamforming technique to focus the 'beam' of the synthetic aperture at each pixel in the image, and then estimate the pixel's value from the raw data. The frequency-domain method uses diffraction tomography techniques to backward propagate the received wavefield to estimate the target area.[1] The received radar data can be viewed as samples of this received wavefield.

In addition to the image formation problem, we have had to develop methods to account for different imaging geometries (monostatic, bistatic, and multistatic); layered variations in the dielectric constant of the transmission media; removal of the coupling pulse between the transmitter and the receiver; removal of the ground reflection; signal propagation loss; deconvolution of the transmitted pulse; and the dispersive and attenuative effects of conductive media on the propagating pulse.

During FY-93, we supported several programs at Lawrence Livermore National Laboratory (LLNL): the ground penetrating imaging radar (GPIR) project,[2,3] the mine detection project, and the UWB radar imaging effort. For the GPIR project, we processed data from electromagnetic simulations[4] as well as real data from our laboratory concrete test-bed, which is similar in composition to a typical bridge deck. We formed high-resolution images of this data, and were

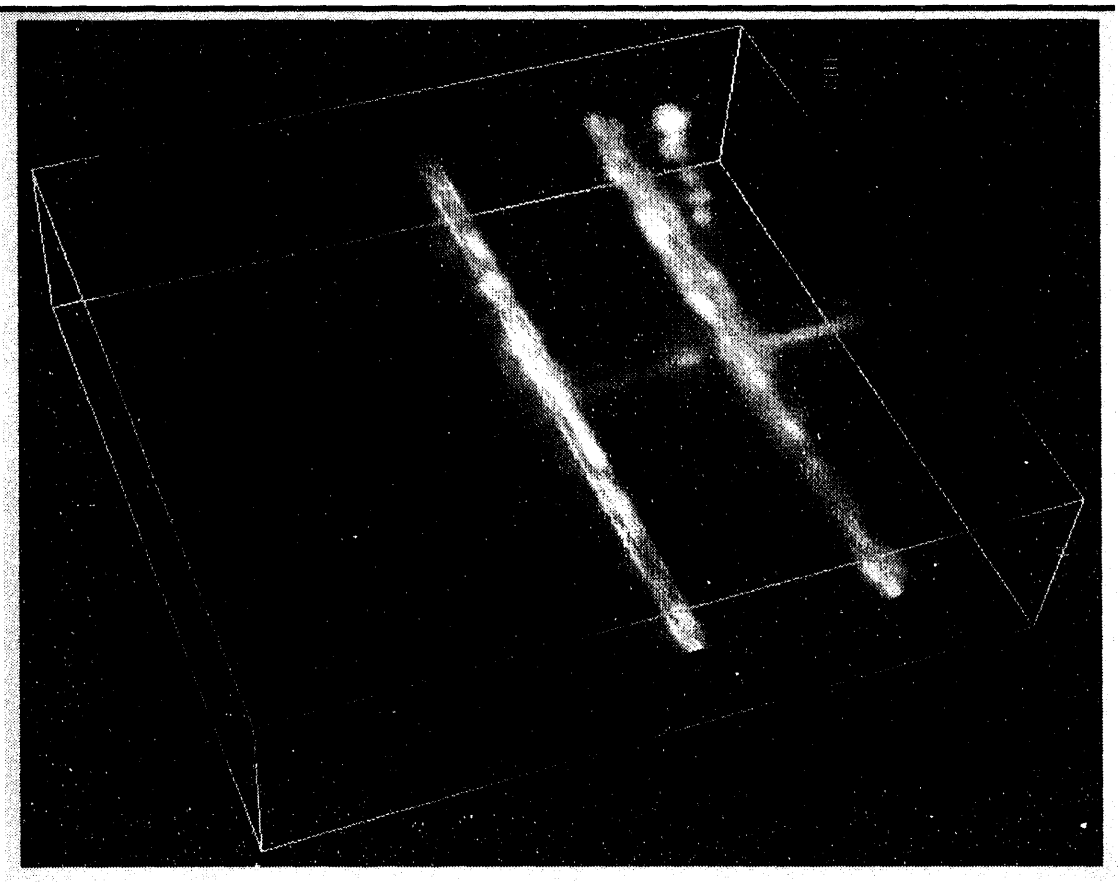

***Figure 3.* A high-resolution, 3-D rendering of the laboratory concrete test slab. The steel rebars are clearly visible.**

Figure 4.
A 2-D map of the ground plane of a minefield. The image is overlaid with a truth map showing the actual positions of the mines and fiducial markers.

able to identify steel reinforcing bars (rebars), voids, and various defects. A high-resolution, 3-D reconstruction of the concrete test slab, showing several rebars, is shown in **Fig. 3**. For the mine detection project, we successfully imaged radar data from a simulated mine field at LLNL's Nevada Test Site, producing accurate, 2-D, high-resolution maps of the ground plane showing the locations of surrogate metal and plastic mines and fiducial markers. A typical minefield image, overlaid with location markers indicating the positions of the mines and fiducials, is shown in **Fig. 4**. We also developed algorithms suitable for multistatic and airborne applications for the laser program's UWB radar imaging project.

Future Work

In FY-94, we plan to continue our research in UWB radar imaging, with the primary emphasis on supporting the GPIR project. We will continue research in the areas of attenuation compensation; coupling pulse removal (the coupling pulse is not uniform throughout all antennas in bistatic and multistatic systems); improved pulse deconvolution techniques; compensation for clutter due to aggregate; and beamwidth compensation. We also plan to investigate the effects of spatial decimation on image quality, to process new simulation and test-bed data, and to perform a parallel-processor-system evaluation and design.

Acknowledgements

We would like to thank our collaborators J. Warhus and S. Nelson; D. Goodman, Remote Sensing, Imaging, and Signal Engineering Thrust Area Leader; J. DeFord, Computational Electronics and Electromagnetics Thrust Area Leader; and Mark Newton, Power Conversion Technologies (formerly Microwave and Pulsed Power) Thrust Area Leader, for their support of this project.

---

1. J.E. Mast, *Microwave Pulse-Echo Radar Imaging for the Non-Destructive Evaluation of Civil Structures*, Ph.D. Thesis, University of Illinois at Urbana-Champaign (1993).
2. J.P. Warhus, S.D. Nelson, J.M. Hernandez, E.M. Johansson, H. Lee, and B. Douglass, "Ground Penetrating Imaging Radar for Bridge Inspection," *Engineering Research, Development, and Technology*, Lawrence Livermore National Laboratory, Livermore, California, UCRL-53868-92, 7-5 (1993).
3. J.P. Warhus, J.E. Mast, S.D. Nelson, and E.M. Johansson, *Ground Penetrating Imaging Radar Development for Bridge Deck and Road Bed inspection*, Lawrence Livermore National Laboratory, Livermore, California, UCRL-ID-113954 (1993).
4. S.D. Nelson and C.A. Avalle, "Electromagnetic Modeling and Experiments for Dispersive Media," *Engineering Research, Development, and Technology*, Lawrence Livermore National Laboratory, Livermore, California, UCRL-53868-92, 1-21 (1993).

Object Classification for Breast Cancer Diagnosis

Laura N. Mascio
Engineering Research Division
Electronics Engineering

Clinton M. Logan
Engineering Sciences
Mechanical Engineering

Jose M. Hernandez
Defense Sciences Engineering Division
Electronics Engineering

At Lawrence Livermore National Laboratory, we are developing a systems approach to better mammography. We have made significant technical advances in producing high quality images.

Introduction

In the past decade alone, breast cancer has killed a number of women larger than the number of people who have died in all American wars combined. There is no known cure for breast cancer and no good indicators for its cause. In fact, we know so little about the mechanisms of breast cancer, that about 75% of women who develop breast cancer have no known risk factors.

Until some cure is formulated, the only weapon against breast cancer is early detection. Currently, x ray mammography is the only means for detecting one of the earliest indicators for treatable breast cancer: microcalcifications, which are small calcium-rich mineral deposits that are often coincident with cancer cells.

When microcalcifications are imaged onto x ray film, using standard mammography equipment, the quality of the image is one important factor determining the detectability of the small, subtle objects. At Lawrence Livermore National Laboratory, one part of our systems approach to better mammography involves the design of a better x-ray screening unit that will yield better images while giving a lower radiation dose to the patient.

Other factors determining the detectability of elusive microcalcifications are human factors. In one day, radiologists may manually scan hundreds of films with a magnifying glass, which is tedious work and prone to error. Time of day, distractions, fatigue, and skill level can also affect a radiologist's diagnosis. Currently, there is an estimated 20% false-negative rate in detecting cancers that were indicated by microcalcifications. One third of these were retroactively rated as being clearly present in the mammogram, but were simply overlooked. Subjectivity and other human factors can be minimized with computer-aided diagnosis, which we are developing.

The software development for computer-aided diagnosis of digital mammograms is a key element of our systems engineering approach to improving early detection of breast cancer. This year, we were given funding from the Department of Energy as part of a Cooperative Research and Development Agreement (CRADA) with Fischer Imaging Corporation, to develop a filmless, full-field, digital-mammography screening unit. This filmless unit will produce higher-quality mammographic images than conventional film screen systems, and will give us direct digital data for our computer-aided diagnosis software project.

Progress

We have built the tools necessary to develop a first-generation object classifier. Our technical advances are as follows:

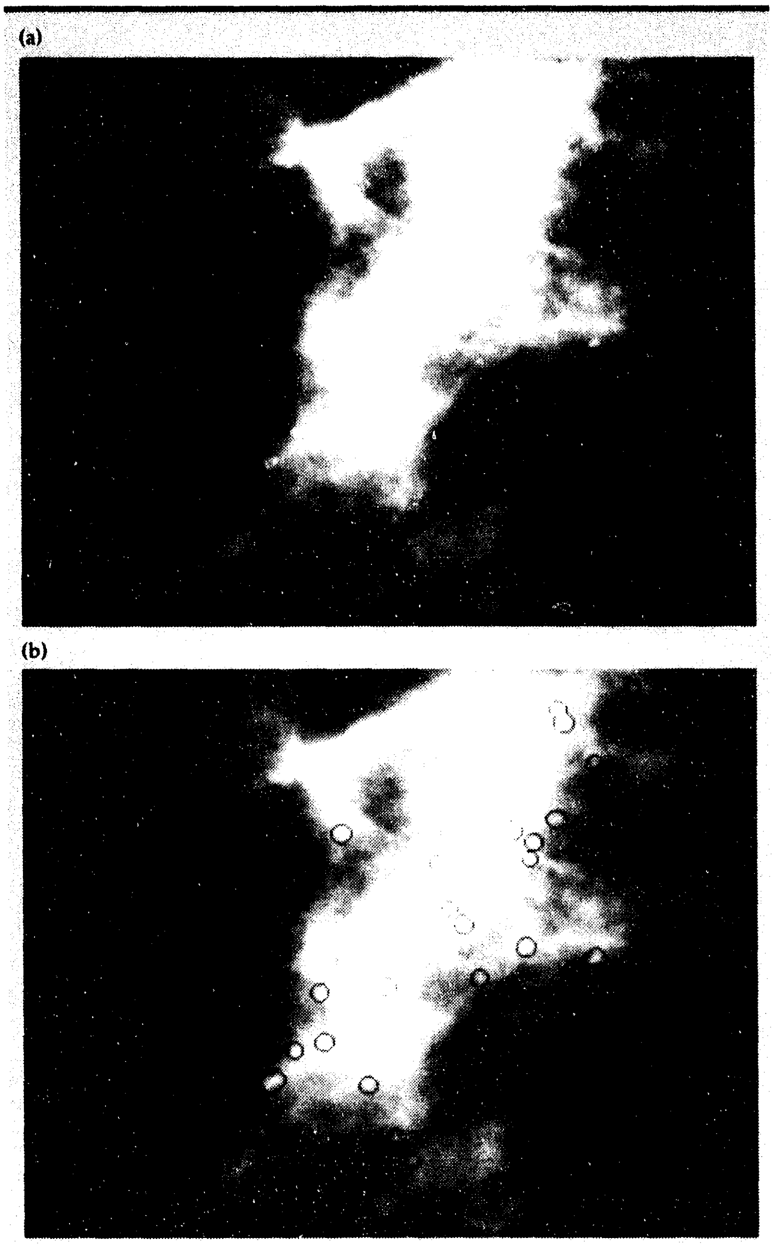

***Figure 1.** (a) Part of a digitized mammogram; (b) the results of computer analysis to assist in detecting microcalcifications.*

(1) we developed a scheme for displaying the approximately 50-Mb images, with magnification capabilities for areas of interest;

(2) we devised a built-in, motif-based (point and click) interface so radiologists can interactively give us a 'gold standard' truth for each detected object;

(3) we implemented algorithms to compute more than 60 feature measurements for detected objects; and

(4) we completed the initial stages of a microcalcification classifier to reduce false alarms, by designing a discriminant analysis scheme that can separate microcalcifications from artifacts, with a 0.085 error rate (on a scale of 0 to 1).

We are currently finalizing the design for a scheme that will discriminate between marker and non-marker microcalcifications.

As a result of the high level of human interest in the subject of breast cancer, and of the match with the current theme of combined use for weapons technology, the details of this project have been highly publicized through various media. We have given invited presentations at universities and technology transfer forums such as the University of California at Davis and the NASA/NCI Digital Mammography Workshop in Bethesda, Maryland.

Our preliminary efforts in microcalcification detection (see **Fig. 1**) have been presented to SPIE's Medical Imaging 1993 Conference. We have also published a manuscript in the proceedings.[1]

Future Work

We are working on proposals for CRADA's for the computerized inspection of mammograms. We will also continue to be active in technical collaborations both in the U.S. and abroad.

---

[1] L.N. Mascio, J.M. Hernandez, and C.M. Logan, "Automated Analysis for Microcalcifications in High-Resolution Digital Mammograms," *Proc. Medical Imaging 1993: Image Processing*, M.H. Loew (Ed.), *Proc. SPIE* **1898**, 472 (1993).

Image Sequence Analysis Using Spatio-Temporal Texture

Sailes K. Sengupta,
Gregory A. Clark, and
Freddie L. Barnes
Engineering Research Division
Electronics Engineering

Paul C. Schaich
Defense Sciences Engineering Division
Electronics Engineering

We have developed and coded an algorithm for motion pattern classification based on spatio-temporal texture. The algorithm has been implemented and tested for the detection of wakes in simulated data with a relatively low signal-to-noise ratio (0.7 dB). Using a 'hold one out' method, a detection probability of 100% with a 0% false alarm rate has been achieved on the limited number of samples (47 in each category) tested. The actual detection can be displayed in the form of a movie that can effectively show the submarine tracks based on the detected wake locations.

Introduction

Analysis and classification of motion types of objects in the three-dimensional (3-D) world based on image sequences (movies) of these objects is an important task in computer vision. Among various potential applications are (1) detection of wakes created by different types of vessels on the sea; (2) automated surveillance; (3) intruder detection; (4) industrial monitoring to detect malfunctioning moving parts; (5) monitoring satellite imagery for developing storm systems; and (6) monitoring crowds for early signs of disturbances. These motions share a common characteristic; namely, spatial and temporal uniformity characterized by structural or statistical self-similarity in space and time. This uniformity can be exploited to generate certain invariant features called the 'spatio-temporal textures' associated with each type of motion and used for the correct identification of the motion types. In the past, analyses have focused on the 'reconstruction' of the 3-D motion from the image sequence in essentially two steps: (1) determine the motion as a projection on the image plane, based on the raw image sequence; and (2) find the 3-D information from the two-dimensional projection. In the present work, we bypass this geometric reconstruction approach and instead use the first- and second-order spatio-temporal textural features to identify the motion types.[1] This report describes the spatio-temporal features used in our classification scheme and provides the experimental results obtained in a study of simulated submarine 'wake' data.

Progress

Two sets of features have been considered. The first set[1] has been based on the statistical distribution of the magnitude and gradient of the normal component of optical flow. These features are invariant under rotation, translation, and scale change. The second set consists of temporally smoothed(averaged over three temporally consecutive frames) spatial texture features. The features in the second set are basically first- and second-order spatial features.[2]

Experimental Results

Simulated wake data consisted of 147 consecutive frames at simulated intervals of approximately 2 s. In **Fig. 1**, the resulting wakes are shown at a signal-to-noise ratio of one. The features described in the previous section were computed for a 16-x-16 window over three consecutive instants. Note that we needed windows over three temporally consecutive frames because

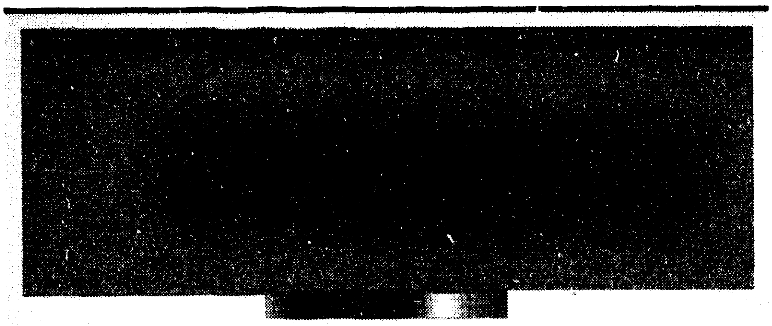

Figure 1. Simulated wake data.

the second-order temporal textural features needed the computation of second-order partial derivatives in space and time. Two types of water motion considered were wakes and ocean (background). The sample windows were chosen from areas in the imagery covering these two motion types. The computed feature vectors for the selected samples were subjected to a feature selection process based on sequential forward selection[3] (SFS), and the selected feature vectors were input to a probabilistic neural network[3] (PNN) for classification in a supervised mode. The number of training samples in each category was 47. Based on the five selected features using SFS, we trained the classifier using the PNN. We tested the efficiency of the classifier by the 'hold one out' method, where each sample is held out in turn and tested for its membership class by the PNN classifier trained with the remaining samples.

The SFS was required to select five 'best' features. Of the selected features, four were from the second set, consisting of first- and second-order spatial texture, and one was from the first set, consisting of optical flow features. The wakes were identified correctly 100% of the time with no false alarm. Features from the first or second set alone were unable to produce 100% accuracy in classification.

---

1. R.C. Nelson and R. Polana, *Comp. Vis. Graph. Img. Proc.: Image Understanding* **56**(1), 78 (1992).
2. R. Welch, S.K. Sengupta, A.K. Goroch, P. Rabindra, N. Rangaraj, and M.S. Navar, *J. Appl. Meteor.* **31**(5), 405 (1992).
3. P. Rabindra, S.K. Sengupta, and R.M. Welch, *Environmetrics* **3**(2), 121 (1992).

Adaptive Optics Technologies for Laser Guide Stars

James M. Brase,
Donald T. Gavel, and
Kenneth Avicola
Laser Engineering Division
Electronics Engineering

The goal of the adaptive optics (AO) technologies project is to develop an engineering technical base in AO to support new program activities in the areas of remote sensing and laser beam propagation for scientific, defense, and energy applications. We provide advanced AO technology to the Laser Guide Star Director's Initiative at Lawrence Livermore National Laboratory, which, in turn, provides a test bed for technology development.

Introduction

Adaptive optics (AO) allows imaging and beam control systems to actively adjust to changing environments, such as the atmosphere or the heating of optical elements. AO couples advances in sensors, computers, and optical components to foster potentially revolutionary performance increases in optical systems. For example, the resolution of the Keck 10-m telescope can be increased by a factor of 100 using AO systems.

For the past two years, the engineering departments, the Laser Guide Star Director's Initiative, and the Atomic Vapor Laser Isotope Separation (AVLIS) Program have collaborated to develop AO capabilities at Lawrence Livermore National Laboratory. We have developed a complete system-analysis and error-budgeting system. An end-to-end simulator with multiple, moving-phase-screen, atmospheric models and extended reference and object models provides more complete performance predictions. These capabilities have recently been used to predict atmospheric effects of airborne lasers for weapons and illuminator systems, and for establishing atmospheric correction requirements for power beaming to low-earth-orbit satellites.

Progress

A typical AO system has a sensor for measuring the wavefront perturbation to be corrected, a high-speed computer for calculating the required correction, and an optical correction element, usually a deformable mirror as shown in **Fig. 1**. The work within the project has concentrated on the

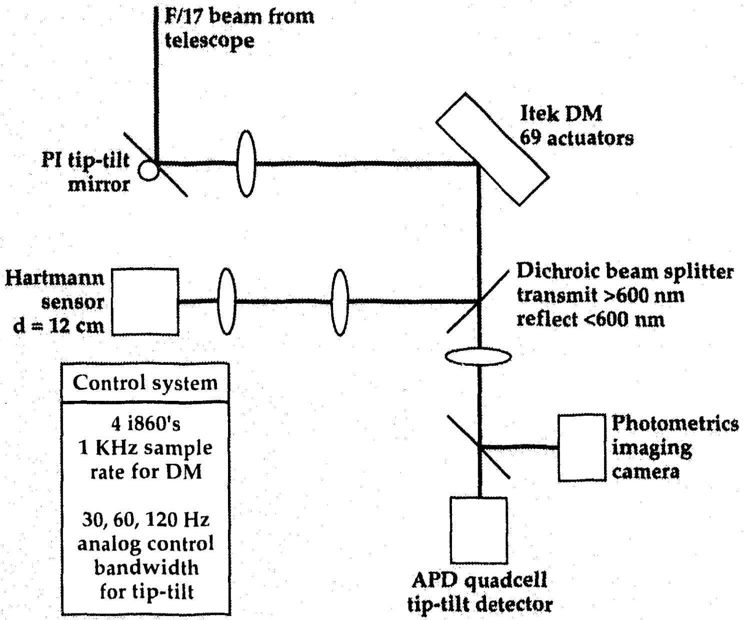

Figure 1. A block diagram of the AO system for the 40-in. telescope at Lick Observatory.

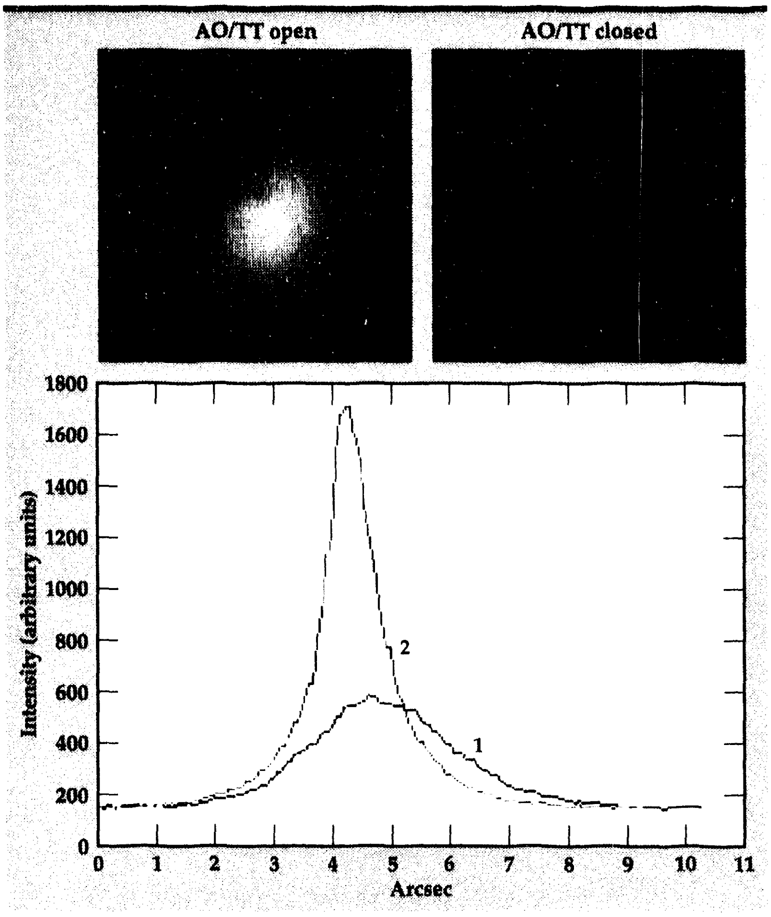

Figure 2. Initial results from the Lick Observatory AO system. The images are of the star α Aurigae at a wavelength of 1 μm. The uncorrected image is on the left, and the corrected is on the right. At the bottom are lineouts from the images, which show the FWHM and peak height improvements.

first two elements, the wavefront sensor and the wavefront control computer. In collaboration with Kodak, we have fielded a 1- to 2-kHz intensified Hartmann sensor and successfully tested it with both natural stars and a sodium laser guide star—the first ever measurement of this type. This high-speed wavefront sensor technology has been transferred to the AVLIS program for beam quality measurements.

Only a few years ago, wavefront control computers required multiple racks of dedicated electronics to keep up with the atmosphere. We have recently demonstrated a scalable, parallel, wave front control architecture, based entirely on commercially available board-level products, at speeds adequate for our first-generation AO systems (69 subapertures, ~ 100 MFLOP). It readily scales to large apertures for near-IR application (241 subapertures, ~ 500 MFLOP).

Our final objective in this project has been to demonstrate a fully operational system for real-time correction of the atmospheric turbulence, incorporating our wavefront sensor and control technology. We accomplished this goal in a series of experiments at Lick Observatory.

Our system, shown in **Fig. 1**, incorporates a 69-actuator Itek DM with the Kodak-based sensor and our scalable, wavefront control computer described above. A separate control loop, based on a silicon avalanche photodiode detector and a 3-axis piezoelectric mirror, controls the overall image position, or tip-tilt (TT). The corrected images are formed on a Photometrics camera with a 1024-x-1280 charge-coupled-device detector.

The results of initial system tests are shown in **Fig. 2**. These are images of the star α Aurigae at a wavelength of 1 μm. The first is uncorrected with a full-width half-maximum (FWHM) of about 3 arcsec. In the second image, the tip-tilt loop and DM loops have been closed. The FWHM is reduced to less than 1 arcsec, and the peak intensity is increased by nearly a factor of four. These results are consistent with the expected performance in conditions of poor visibility.

Future Work

AO technology has been demonstrated for the correction of the atmospheric turbulence that degrades astronomical images. We are continuing work on the development of Laser Guide Star systems, for the Lick and Keck Observatories, that will tremendously expand the sky coverage of these systems. The modest image improvements shown here are the first steps on the path to much more dramatic improvements for larger telescopes. AO technology represents a fundamental change in systems that use light. Their potential applications are exciting.

Vision-Based Autonomous Grasping of Unknown Piled Objects

Robert K. Johnson
Engineering Research Division
Electronics Engineering

Computer vision techniques have been used to develop a vision-based grasping capability for autonomously picking and placing unknown piled objects. This work is currently being applied to the problem of hazardous waste sorting in support of the Department of Energy's Mixed Waste Operations Program.

Introduction

The Department of Energy has an urgent need for automated systems to perform hazardous waste characterization, sorting, and processing. The Environmental Restoration and Waste Management–Applied Technology (ER&WM–AT) Robotics Program at Lawrence Livermore National Laboratory recently has been developing robotics technology to address this problem. To achieve autonomous processing of the waste objects, computer vision is required to determine grasp position and orientation.

In FY-92, we developed a vision-based grasping capability, which was limited to singulated unknown objects.[1] Stereo image processing was used to obtain the required three-dimensional (3-D) range data, and a simplified grasp analysis was performed for each reconstructed object. In FY-93, we have modified and extended this work to the more complex and realistic case of unknown piled objects.

Progress

A structured light system developed at Oak Ridge National Laboratory was used to obtain the 3-D range information.[2] This system provided an increase in both efficiency and accuracy over the previous stereo camera system. New algorithms were developed to process the structured light data and to segment the resulting range image into graspable regions. In addition, a robust grasp-planning strategy was developed to achieve the desired goal of picking the top unknown object from a pile. These capabilities have been implemented within the VISION programming environment[3] and demonstrated through simulations with real, structured light data of piled objects. The total CPU time required to process the structured light data and determine an optimal grasp is approximately 10 s on a Sun SPARCstation-2 for a typical scene. This work is currently being integrated into the ER&WM–AT Robotics Program's Interactive Controls Laboratory to produce a live demonstration of autonomous waste processing.

Range Image Processing

A range image is formed by interpolating the initial range data obtained from the structured light system. The raw structured light data consists of a random list of 3-D surface samples of

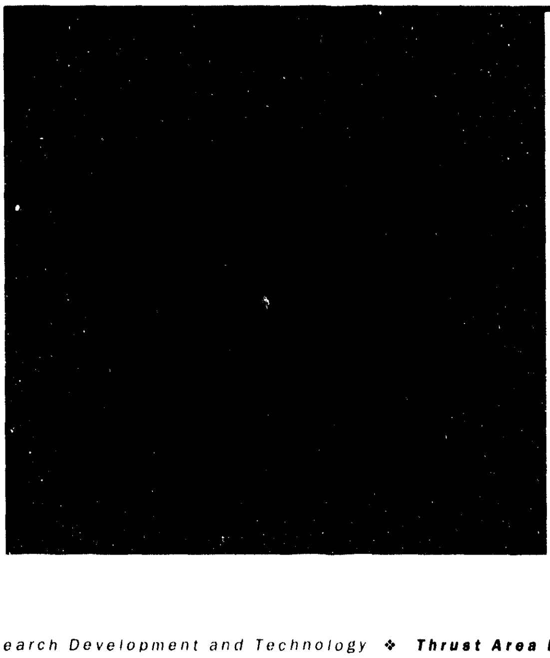

Figure 1. Initial range data obtained with structured light system for scene of piled objects.

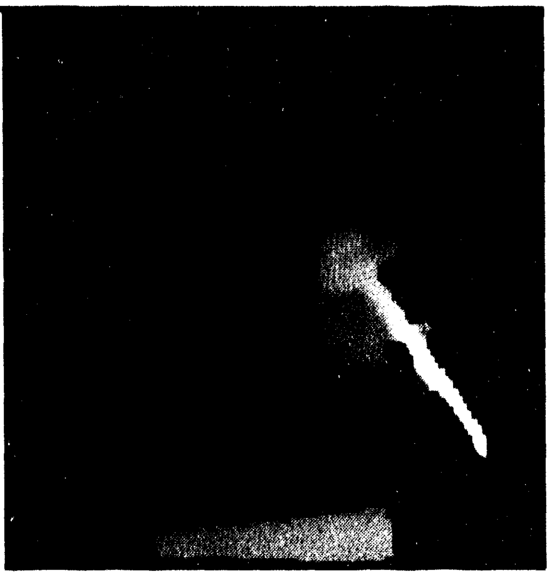

Figure 2. Range image formed from interpolation of initial range data.

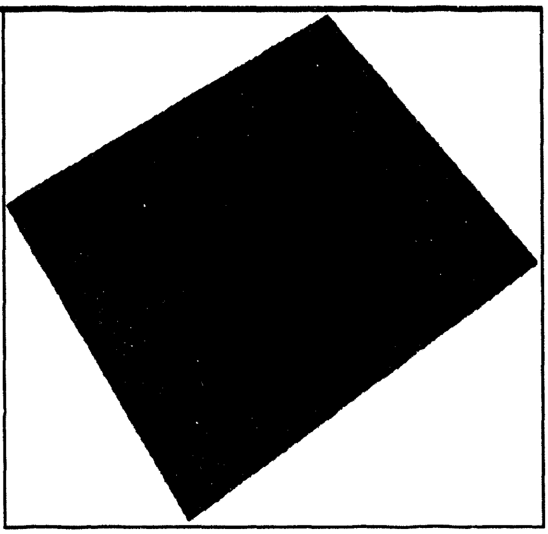

Figure 3. Surface plot of range image.

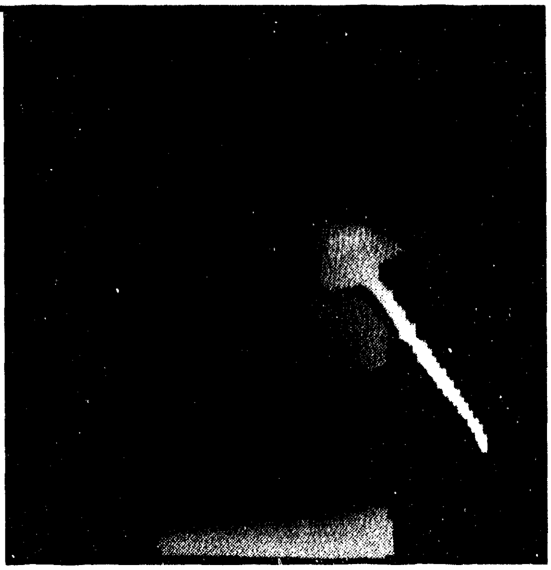

Figure 4. Edge-filtered range image.

the area in the field of view. We first grid the z-coordinate data into a uniform two-dimensional (2-D) array by using the corresponding x,y coordinates to determine the appropriate array element position. We typically use a sampling density on the order of 1 mm in x and *y*. Averaging is used if more than one sample falls into a particular array element. The initial gridded range data for a typical scene of piled objects is shown in **Fig. 1**. We then use a 2-D filtering technique to produce a continuous range image with strong edges preserved, as shown in **Fig. 2**. We first apply a dilation followed by an erosion, with a data-dependent kernel size, to make the surface sampling continuous. We then apply a Gaussian-weighted smoothing filter while preserving the background and thus the edges. A surface plot of the range image in **Fig. 2** is shown in **Fig. 3**.

An edge-based segmentation algorithm was developed to isolate the regions of interest to consider for grasping. The current grasp analysis is limited to parallel grippers oriented with their tool axis perpendicular to the range image plane. Therefore, object edges appropriate for grasping will also be fairly perpendicular to the range image plane. We locate these edges by computing the surface normals, using 3-x-3 Sobel kernels to compute the necessary derivatives.[4] Next we cut along the edges to form an edge-filtered range image, as shown in **Fig. 4**. A region segmentation with size constraints is then performed to locate the main regions of interest, which have edges with the desired orientation, as shown in **Fig. 5**.

Grasp Planning

Our current grasp analysis is constrained to four degrees of freedom, including three parameters for position and one parameter for orientation of a parallel gripper. The segmented regions of interest are first ordered, based on their maximum heights, and then processed from the top down until either a valid grasp is found or the number of regions is exhausted. Attainable grasps are identified for a particular region by analyzing its shape with respect to the gripper dimensions. A 2-D, symmetric, Euclidean skeleton[1] is first computed, and each skeletal point serves as a potential grasp site to be analyzed. The attainable grasps are ranked from best to worst, based on a weighted least-squares combination of seven grasp features: three measurements of parallel line boundary deviation, two clearance measurements between the grip-

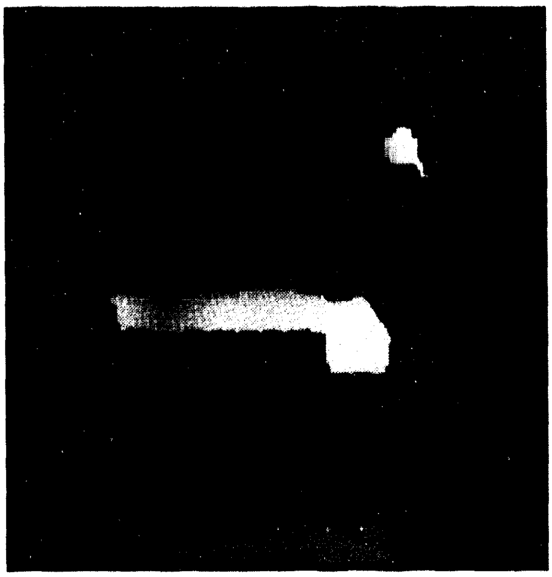

Figure 5. Segmented regions of interest.

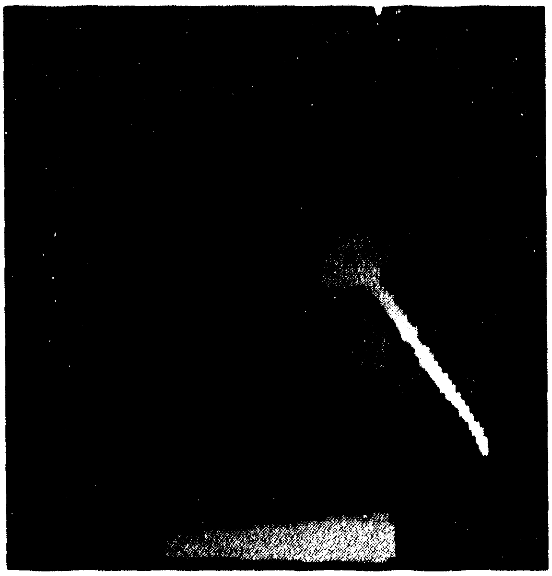

Figure 6. Optimal grasp for top region segment.

per and its surroundings, the maximum height under the gripper, and the distance from the centroid of the region. The optimality criterion can be modified as desired by simply adjusting the relative weightings of these seven features. An example of an optimal grasp for the top region of **Fig. 5** is shown in **Fig. 6**.

Acknowledgements

The author would like to thank the ER&WM–AT Robotics Program and, in particular, the Automation and Robotics Group for their support and for the use of their facilities.

---

1. S.Y. Lu, R.K. Johnson, and J.E. Hernandez, "Vision-based Grasping for Autonomous Sorting of Unknown Objects," *Engineering Research, Development, and Technology*, Lawrence Livermore National Laboratory, Livermore, California, UCRL-53868-92, 9-1 (1993).
2. F.W. DePiero and R.L. Kress, *Camera Calibration for SCOPE—Surface Characterization and Object Pose Equipment*, Oak Ridge National Laboratory, Oak Ridge, Tennessee, draft (1993).
3. J.E. Hernandez and M.R. Buhl, "VISION: An Object-Oriented Environment for Computer Vision and Pattern Recognition," *Engineering Research, Development, and Technology*, Lawrence Livermore National Laboratory, Livermore, California, UCRL-53868-92, 9-15 (1993).
4. D.H. Ballard and C.M. Brown, *Computer Vision*, Prentice-Hall (Englewood Cliffs, New Jersey), 1982.

Author Index

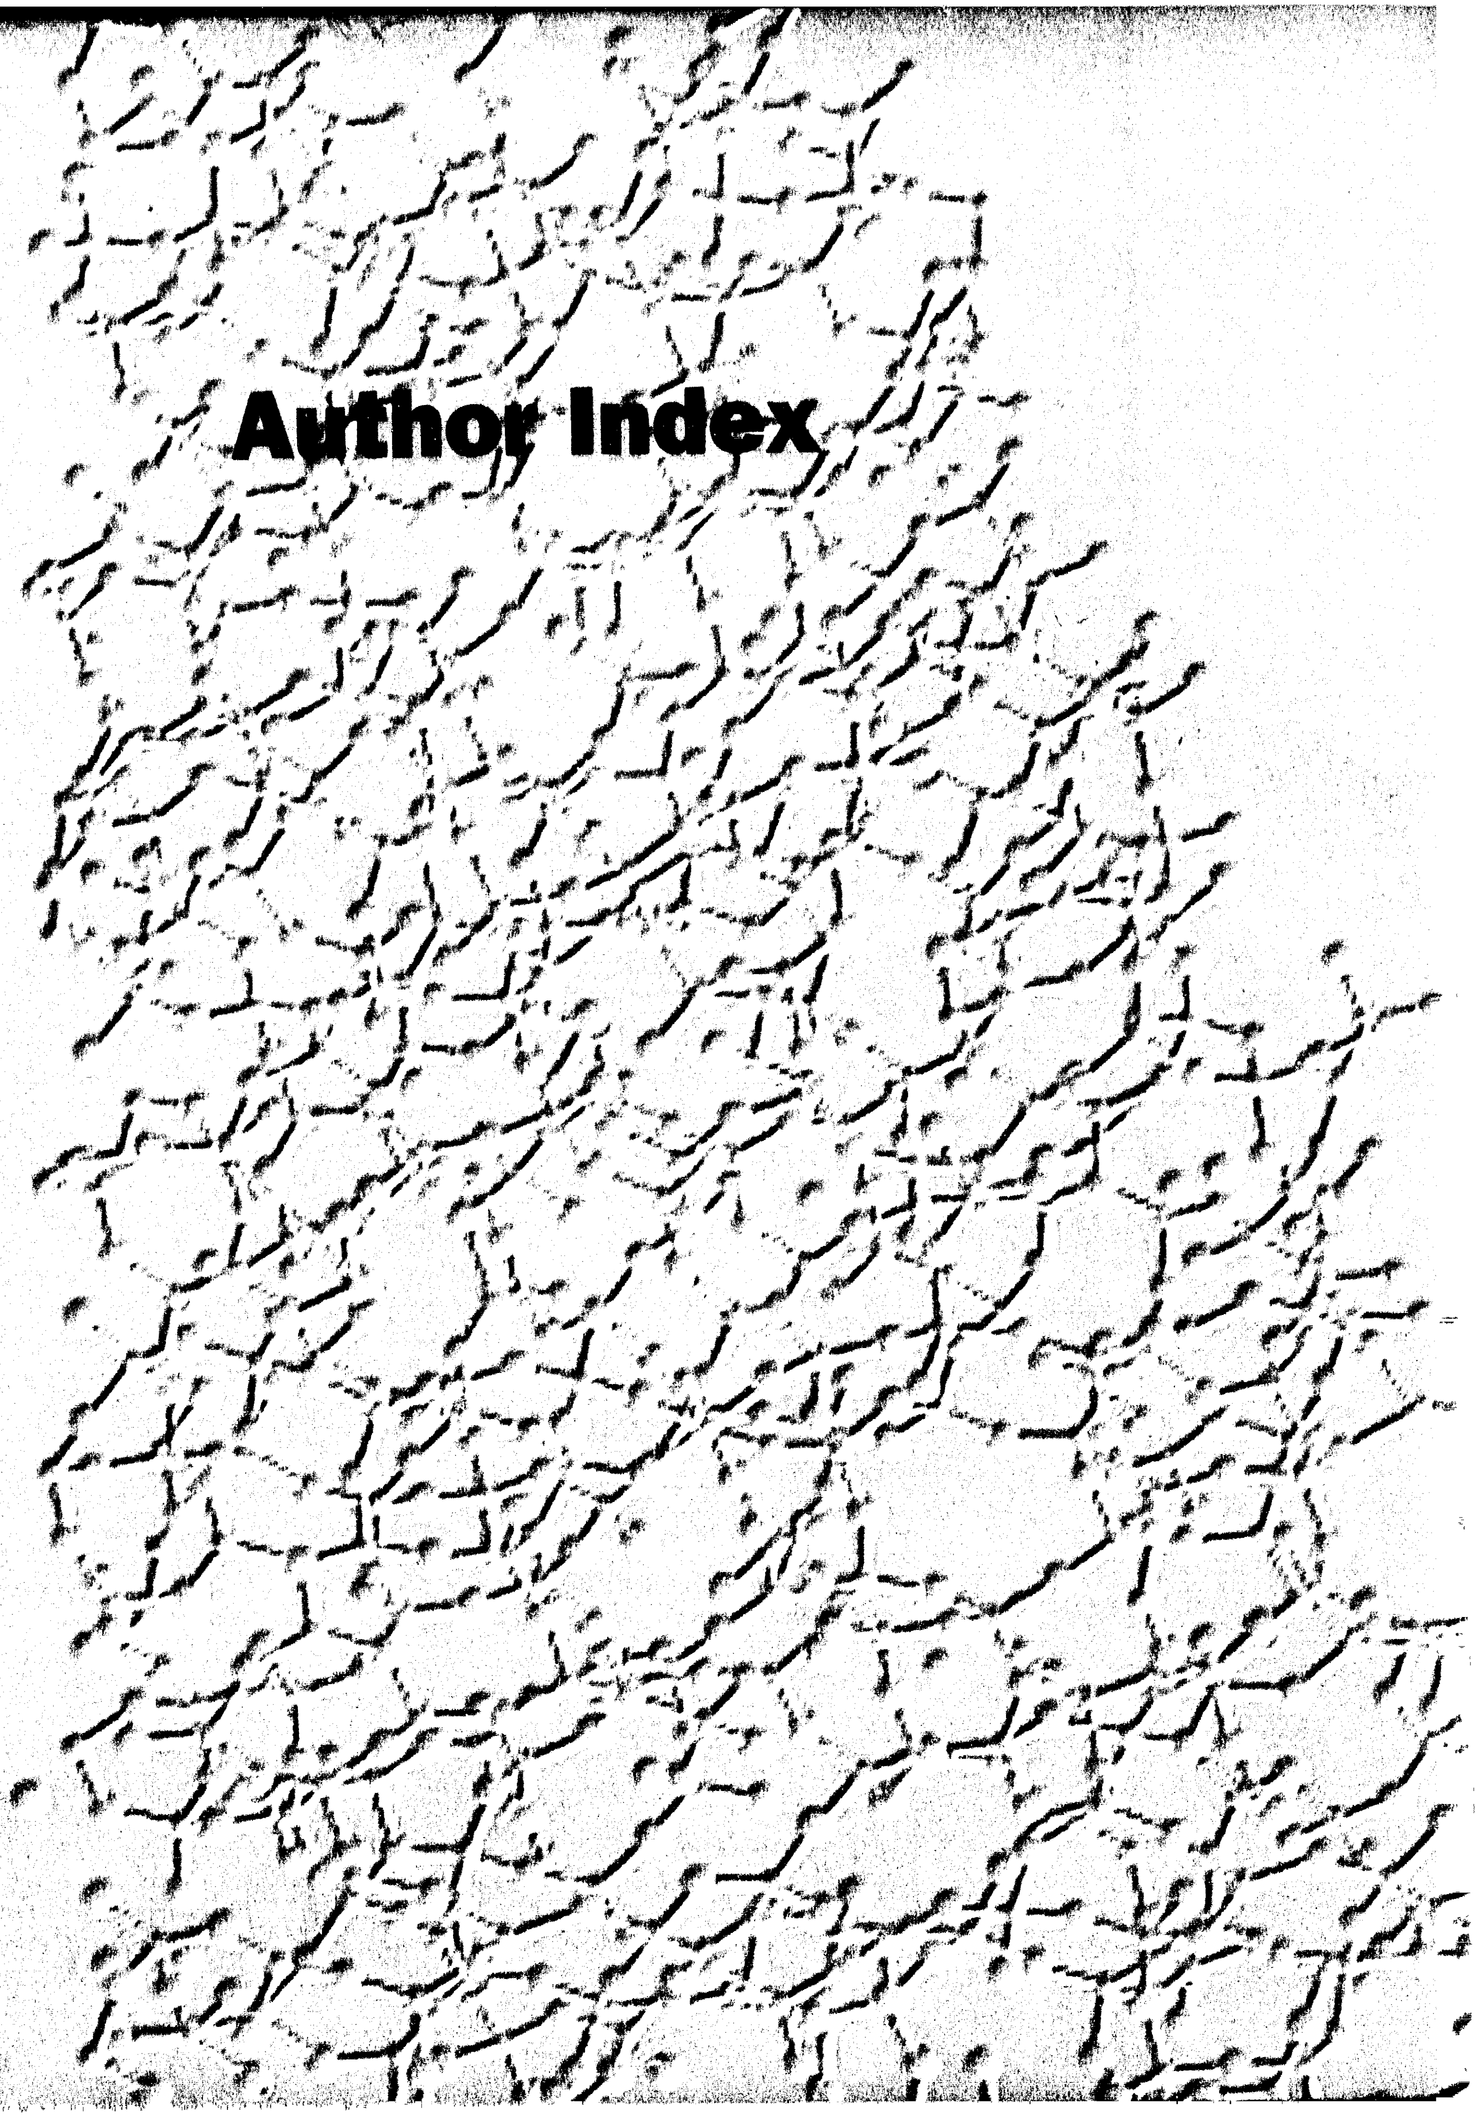

DATE
FILMED
7 /27/94

END