

UCRL-JC-132806

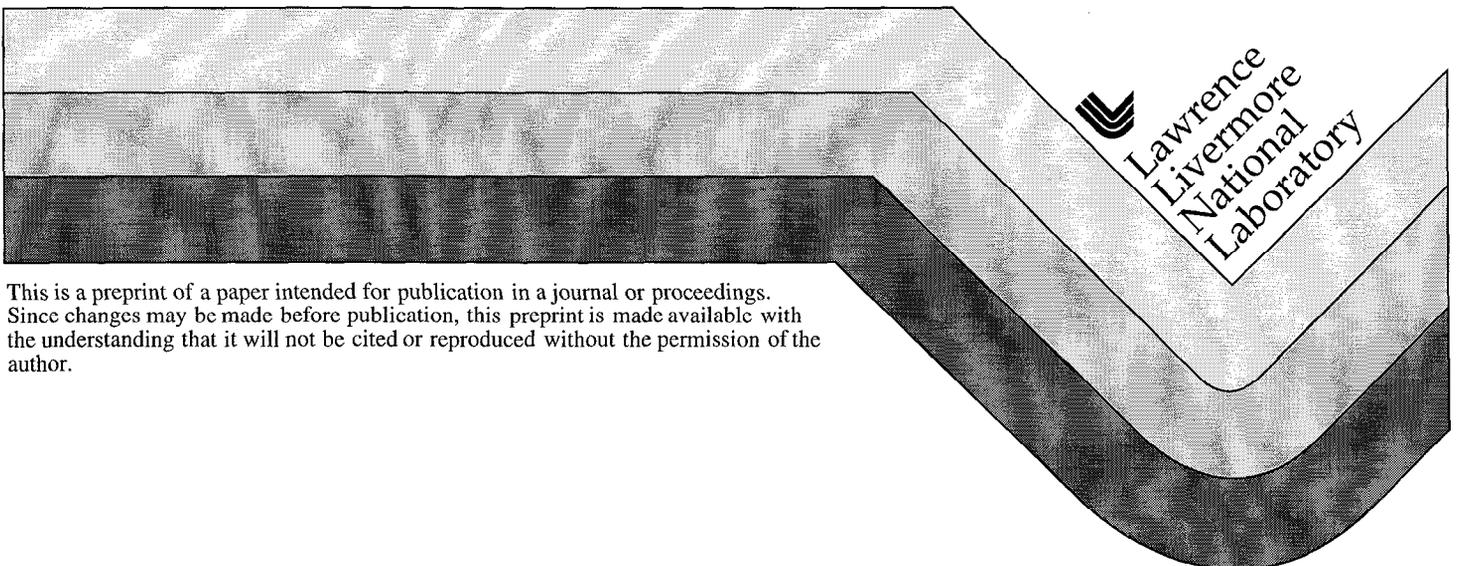
PREPRINT

# Radiographic Simulations and Analysis for ASCI

D.M. Slone  
M.B. Aufderheide  
A. Schach von Wittenau

This paper was prepared for submittal to the  
*Ninth Society for Industrial and Applied Mathematics Conference on  
Parallel Processing for Scientific Computing*  
San Antonio, TX  
March 22-24, 1999

December 18, 1999



This is a preprint of a paper intended for publication in a journal or proceedings.  
Since changes may be made before publication, this preprint is made available with  
the understanding that it will not be cited or reproduced without the permission of the  
author.

#### DISCLAIMER

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

# Chapter 1

## Radiographic Simulations and Analysis for ASCI\*

Dale M. Slone<sup>†</sup>    Maurice B. Aufderheide<sup>‡</sup>    Alexis Schach von Wittenau<sup>‡</sup>

### Abstract

In this paper, we describe our work on developing quantitatively accurate radiographic simulation and analysis tools for ASCI hydro codes. We have extended the ability of HADES, our code which simulates radiography through a mesh, to treat the complex meshes used in ASCI calculations. Our ultimate goal is to allow direct comparison between experimental radiographs and full physics simulated radiographs of ASCI calculations. We describe the ray-tracing algorithm we have developed for fast, accurate simulation of dynamic radiographs with the meshes used in ALE3D, an LLNL ASCI code. Spectral effects and material compositions are included. In addition to the newness of the mesh types, the distributed nature of domain decomposed problems requires special treatment by the radiographic code. Because of the size of such problems, we have parallelized the radiographic simulation, in order to have quick turnaround time. Presently, this is done using the domain decomposition from the hydro code. We demonstrate good parallel scaling as the size of the problem is increased. We show a comparison between an experimental radiograph of a high explosive detonation and a simulated radiograph of an ALE3D calculation. We conclude with a discussion of future work.

### 1 Radiography and ASCI

The Accelerated Strategic Computing Initiative (ASCI) is a part of the Department of Energy's strategy to shift from nuclear weapons testing to an approach based on computational and experimental methods[2]. Specifically, ASCI will accelerate the development of simulation capabilities needed to ensure confidence in a safe and reliable stockpile. ASCI will provide computational capabilities that will help scientists understand aging weapons, predict when components will have to be replaced and evaluate the implications of changes in materials and fabrication processes. This science-based understanding is essential to ensure that changes brought about through aging or remanufacturing will not adversely affect the enduring stockpile. The Science-Based Stockpile Stewardship program will develop the new means to assess the performance of nuclear stockpile systems, predict their safety and reliability, and certify their functionality.

Transmission radiography is standard technique used in nondestructive evaluation. In this technique, a source such as a bremsstrahlung X-ray source is used to illuminate the object of interest. A detector array of some sort is placed behind the object and the attenuated radiation is recorded. The result is a 2D projection of the object onto the

---

\*This work was supported under the auspices of the U.S. Dept. of Energy by Lawrence Livermore National Laboratory under contract No. W-7405-Eng-48.

<sup>†</sup>Computer Scientist, Scientific Computing Applications Division, Lawrence Livermore National Laboratory, Livermore, CA.

<sup>‡</sup>Physicist, B Division, Lawrence Livermore National Laboratory, Livermore, CA.

imaging plane. If the object is viewed from many different angles, tomography can be used to reconstruct the object. All radiographs are corrupted with secondary radiation of some kind which can confuse one's interpretation of transmission radiographs. Also, if the object is fairly complex, interpreting its 2D projection can be problematic. For these reasons, it is often useful to simulate radiographs of the object using an appropriate model of the object of interest and the radiographic system. Often the model may be a finite element or finite difference mesh containing a discretized model of the object.

We have been developing HADES as a code to simulate radiography through hydrodynamic meshes. Applications include radiography of high explosive detonations [3], liner implosions [4], and industrial non-destructive evaluation. Until recently, HADES was only able to treat some 2D mesh types (meshes which assume a cylindrical symmetry in the object). In this paper, we describe how we have extended our capabilities to a 3D ASCII mesh. HADES uses ray-tracing techniques to compute the transmission radiographs through meshes, and objects of uniform density such as collimators, spheres, plates, cylinders and cones. This allows the user to obtain fast estimates for radiographs, including experimental geometry and layout, spectral effects, blurring and statistical fluctuations using dose to photon number conversions. HADES is able to simulate transmission radiography for spectral and monochromatic x-ray beams and high energy ( $> 800\text{MeV}$ ) protons. Eventually, neutron radiography will also be included. The code has been ported to SGI and DEC workstations, Crays, and ASCII Blue Pacific.

It should be noted that HADES only simulates first order, *direct* processes, not second order effects such as scattering and secondary bremsstrahlung. In many (most) radiographic measurements, these secondary processes are present in the radiograph, but not separable from the direct transmission image. It is not the goal of HADES to compute these processes. Monte Carlo codes such as COG and MCNP [5] can be used to study these secondary processes. The primary goal of HADES is to provide a fast, accurate estimate of direct transmission radiography, for quick parameter studies. Secondary processes, calculated elsewhere, can be added in consistently later.

## 2 Ray-Tracing for ASCII Hydro Codes

The physics of transmission radiography can be simply stated. Radiation from a (nearly point-like) source shines through the object(s) of interest and the fraction which is not attenuated is measured at a detector plane behind the object. The detection plane is pixelized in the calculation in order to mimic actual detector properties or as a necessary discretization of the problem. For all of the radiographic probes described above, the amount of attenuation is a function of the pathlength ( $\int \rho \mu dl$ , where  $\rho$  is material density,  $\mu$  is the mass attenuation coefficient, and  $l$  is the pathlength) between the source and detector pixels.

There are two ways to think about ray-tracing through such a problem. In the first approach, one forms a ray from the source to each detector pixel. For each ray, one steps from the source to the detector pixel, incrementing any contributions to the pathlength from whatever objects are in the problem. For a mesh, one tracks through the mesh zone by zone, computing how the ray traverses the mesh. When the detector pixel is reached, the total is saved and the loop continues on to the next source-detector ray. This approach is similar to what is done in Monte-Carlo particle tracking. In the second approach, instead of looping over rays, one loops over objects in the problem. For each object, one computes the shadow of the object on the detector plane. As one loops over the objects, the composite

shadow of the whole system is built up. These approaches are equivalent as long as the intensity in each detector pixel is only a function of the pathlength along the ray between the source and the detector.

HADES uses the latter approach globally, to combine hydro meshes with various objects of uniform density. For each object, HADES uses the former approach, tracking through the object of interest. Tracking through meshes becomes more problematic as the geometry and connectivity of the mesh becomes more complex.

The Livermore ASCI code ALE3D [1] uses a mesh in which the zone can be approximated as a tetrakis hexahedron. This zone is a “hexahedron” whose eight vertices are allowed to be shifted from their normal positions. As a result of this shifting, each “face” of the hexahedron is broken into four triangular subfaces formed using the four vertices on this face and their spatial average. Because of the size of these problems, these meshes are typically decomposed into tens to hundreds of domains, making the connectivity even more complex.

Because of these factors, it was decided to apply the second ray-tracing approach to these meshes. HADES opens the hydro dump file, and loops over all the zones in all domains of the problem, and computes the shadow of each zone on the detector pixel array. For each zone, the maximal possible shadow on the detector plane is computed, indicating the bundle of rays which must be considered in more detail. Each of these rays are then tested for intersections with the twenty-four faces of the zone. From these intersections, the pathlength contribution to each ray is computed and incremented.

This algorithm is easily adapted to zones of other shapes. In fact, it was adapted from a version which approximated the zones as exact hexahedra. It would have been easy to build a version to deal with tetrahedra. We have been asked whether it would be more efficient to tetrahedralize the mesh first so that only tetrahedra would be ray-traced. Even if one discounts the computational cost of tetrahedralizing the mesh, this approach would still be more inefficient than our current implementation, because every face of every zone must be queried for intersections with the ray. Our algorithm requires 24 intersection calculations for each zone, whereas tetrahedralizing first would result in  $4 \times 24 = 96$  intersection calculations. In addition, our philosophy is to radiograph the hydro model with as great a fidelity as possible, so that there is no chance that remeshing artifacts compromise the simulation.

In this algorithm, there is no effort made to track rays through the mesh. The ray sums are assembled stochastically (as randomly as the zones were stored in the file). This randomness of the algorithm is its greatest weakness. Never during the calculation is it possible to assemble a record of how a ray tracks through all the elements of the simulated geometry. Thus it is not possible to do a line integral along the ray, which may be useful in future refinements which might try to estimate x-ray scattering backgrounds or proton energy loss effects. There are two main advantages to this approach. First, it is easy to implement. Second, it is easy to parallelize. In fact, one can use the domain decomposition used in the hydro calculation.

### 3 Parallel Performance

We investigate the parallel performance of HADES by carrying out a scaling study on two different platforms. The test problem we use is simple, but exercises all of the HADES functionality necessary for more complex problems.

We used MPI (the mpich implementation) to achieve parallelism in HADES. The code uses a master-slave approach, where the slave nodes are used only when instructed to by the

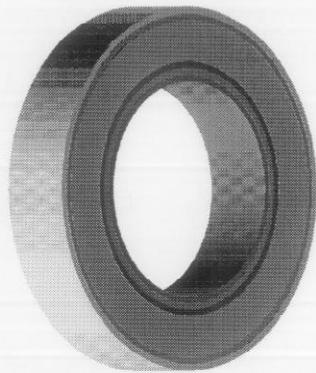


FIG. 1. *The 8 domain test cylinder.*

master node. The master reads the input and performs all the setup necessary to compute the ray-tracing through the object. The master wakes the slave and needs to communicate only twenty-four scalars, one 1D array (number of materials), and one 2D array (number of energy bins \* number of materials). Each domain of the input mesh is assigned to a logical processor. Each processor computes the pathlengths for all voxels within each each domain. Once all of the domains have been computed, a collective communication is used to sum all the pathlengths. After this reduction step, the master node has the full pathlength information from all of the domains in the problem.

### 3.1 Parallel Platforms

We used two parallel platforms available through the Livermore Computing Center. The first is one node (aspen) of the DEC Alpha cluster. Aspen is an AlphaServer 8400 with ten CPUs. Installed on aspen is a gang-scheduling system that puts each requested processor on its own CPU and ensures that the operating system schedules the job as a single process. Gang-scheduling treats the parallel job like any other job, which means that all of the processors get the same time slices at the same time. With the system needing one CPU for scheduling, we ran the scaling study using eight processors.

The second platform is the ASCI IBM Blue-Pacific machine. The partition available for large jobs consists of 168 nodes, with each node having four PowerPC 604e CPUs. Running jobs own the processors, so there is no competition for cycles with other jobs.

### 3.2 The Test Problem

We performed the scaling study using the same cylindrical test problem that is part of the ALE3D regression test suite. The problem is a right cylinder that is comprised of layers. Each layer is made from four domains, where each domain is one quadrant, and are added whole. Each layer contains six materials which are laid out in concentric shells. The two outer and one inner shells are metal, while the the middle layers are non-metal. There is also a central void region. Figure 1 shows the material layout for the 8 domain test cylinder rotated 45 deg around the  $y$  axis.

The test series is composed of a cylinder made from 1 – 128 layers (4 – 512 domains) plus a quadrant and a half of one layer. The problem is designed so that each domain does a constant amount of hydrodynamic work,  $\approx 10,000$  zones.

We used HADES to simulate radiographs along the  $y$  axis of the cylinder; the positions

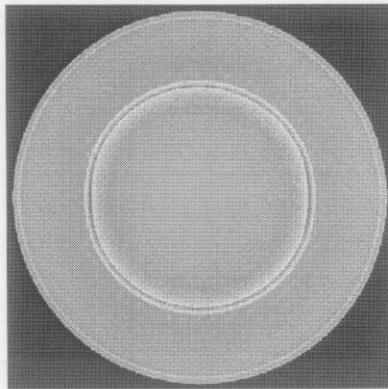


FIG. 2. Radiograph of the 8 domain full cylinder.

of the X-ray source and detector plane ensures that the entire object is in the field of view. Two X-ray sources were used: a  $4\text{MeV}$  monochromatic source and a  $17.5\text{MeV}$  bremsstrahlung spectral source using 44 energy bins. The resulting radiograph from the 8 domain full cylinder with the monochrome source is shown in Figure 2.

### 3.3 Results

We ran the test problem on both machines with 1–512 domains at powers of two (note that we ran a 96 domain problem instead of 128) and both X-ray sources. Since each domain of the mesh is assigned to its own processor and the number of zones in each domain is held constant, the problem is inherently load balanced and the serial runs exhibit linear scaling with problem size. As seen in Figure 3, both platforms show good scaling for the cylindrical scaling problem. The timing numbers shown are all for the  $4\text{MeV}$  monochromatic X-ray source. The spectral X-ray source shows similar behavior, with runtimes  $\approx 1.75$  longer since there is no parallelization over energy bins for the source.

Figure 3 shows wall clock time for the DEC machine aspen. Because the DEC is not a dedicated parallel machine however, the effect of running part of the test series during the day, when the machine becomes heavily loaded manifests itself clearly. The wall clock times for test runs on a heavily loaded machine exhibited a decrease in speedup from 8 to as low as 4. The parallel timings for the medium and fine mesh refinements are not flat due to the 8 cpu limit imposed by the platform. The curves do show a linear slope however that tracks the serial timing slopes quite well.

On the IBM, as seen in Figure 3, the wall clock time increases faster with respect to increasing number of domains due to the large amount of communication required by the master node sending startup data to all of the slave nodes and resulting contention for the communication network. For the 512 domain problem, the user time exhibits a speedup of 365, while the wall clock time speedup is 143. This decrease in speedup of  $\approx 2.5$  is purely due to the large increase in communication cost.

We also took the sixty-four domain problem and ran it on differing numbers of processors to test the overloading of multiple domains on individual processors with the resulting loss of constant work per processor. The code shows good scaling as expected, since the problem is still well load balanced with respect to the amount of work per domain (see Figure 4).

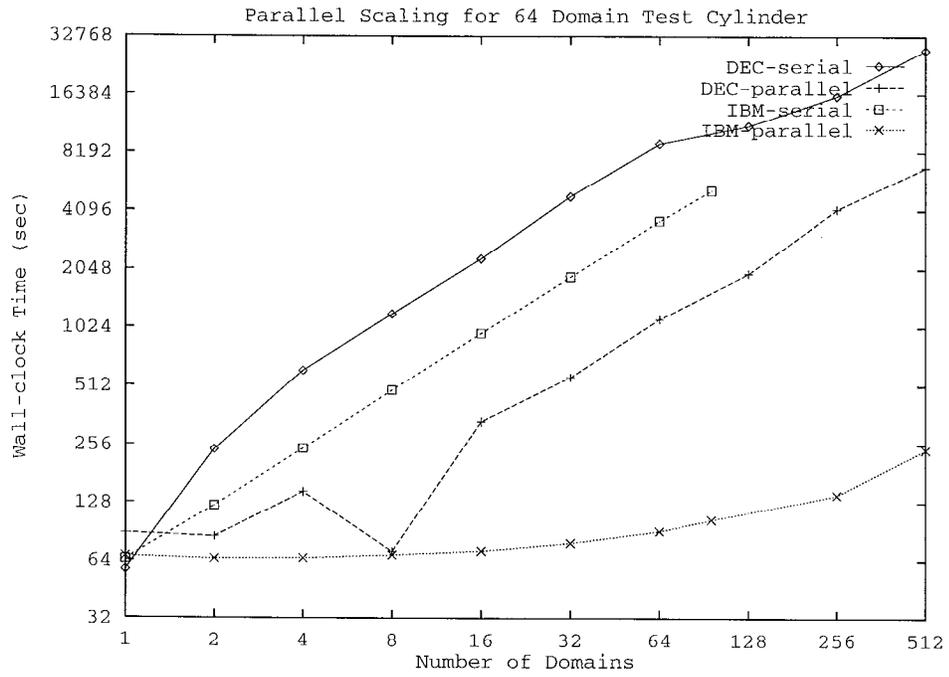


FIG. 3. Scaling of run time with number of domains on the DEC and IBM machines. The parallel curve for the DEC isn't linear due to having only 8 available processors in the box.

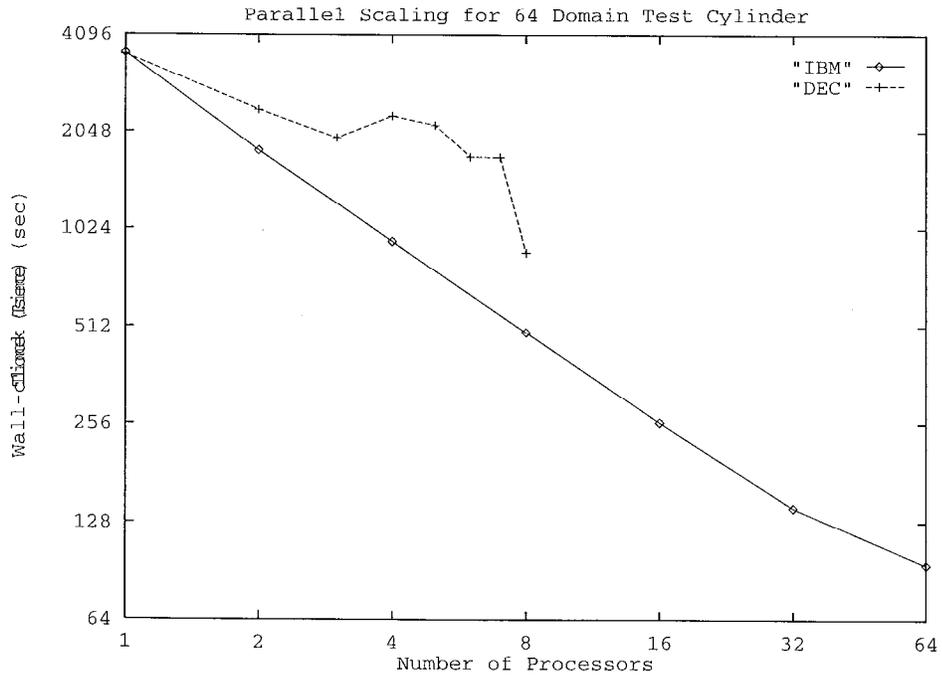


FIG. 4. Scaling of run time for constant domain sized problem.

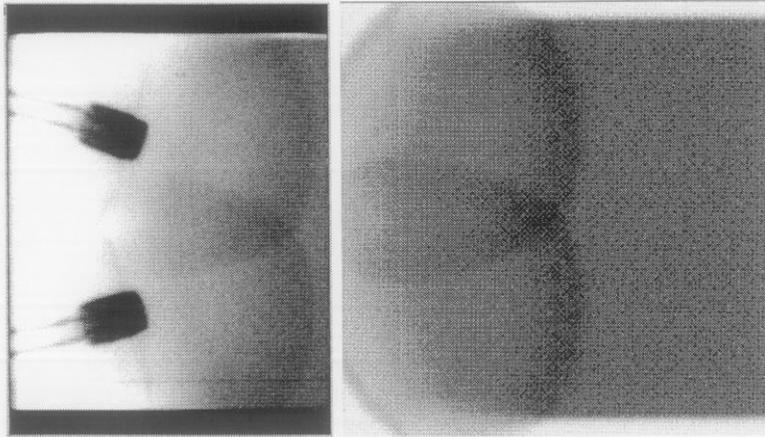


FIG. 5. *Experimental radiograph and HADES simulation of Shot 1561B.*

#### 4 Application to a High Explosive Detonation Experiment

Shot 1561B was performed on October 28, 1997 at the Department of Energy's Big Explosives Experimental Facility (BEEF) at the Nevada Test Site. A rectangular block of LX-14 high explosive (depth of 2.5 cm, width of 4.6 cm, and height 6.5 cm) was detonated simultaneously at two locations on the top face. These detonators generated two outward moving detonation fronts which begin passing through one another after roughly  $1.15 \mu\text{s}$ .

Three Scandiflash 450 KeV Flash X-Ray heads were set along a line which was 2.1 meters from the block. Each head was set to fire at a different time. Film was placed in each field of view so that 3 non-overlapping images could be obtained. In Figure 5, on the left, we show the image captured at  $2.19 \mu\text{s}$ . This problem is inherently 3D and ALE3D was used to simulate it. In Figure 5, on the right, we show a simulated radiograph using the ALE3D model, which had roughly 2.4 million zones. In the simulation we have included the spot and film blurs, the finite dose provided by the machine, and the bremsstrahlung spectrum appropriate for the machine. The qualitative agreement between the simulation and the radiograph is fairly good, particularly as regards the curvature of the detonation fronts. Future work will involve quantitative comparisons between the images, in addition to using the simulation to assess features seen in the experimental image.

#### 5 Conclusions

In this paper we have described our latest work on HADES, which supports ASCI development at LLNL. We have successfully extended our ray-tracing approach to use the complex mesh in ALE3D. Our approach has been to treat all zones stochastically, adding their contributions at the detector plane. The algorithm is easily extensible to other 3D mesh types. This approach has been parallelized and show excellent parallel scaling behavior. We have applied these techniques to some of the newest ASCI calculations and are obtaining promising agreement between simulation and experimental images.

There is a great deal of work to do. We plan to extend our coverage of 3D meshes to other codes as time and interests dictate. We will refine our 3D algorithm in several ways. First, we will seek better load balancing for the calculations. Second, we will examine the use of threads to fully exploit the multiprocessor nodes on ASCI Blue-Pacific. Third, we will seek better ray-tracing algorithms to supersede our current approach. A number of generalized Monte Carlo libraries are under development which might enable us to build

efficient ray-tracing algorithms over a domain-decomposed problem. Last, we will eagerly await and support the more detailed ASCI calculations which are one the horizon.

Another application which we will undertake in the near future is the parallelization of the 2D algorithm. Here we will decompose over ray bundles rather than over the spatial domain. This approach is possible because 2D problems are typically small enough to fit on each processor. This development will allow us to treat finite spot size in a rigorously correct way. It will also allow the accurate and efficient simulation of pinholes and coded aperature masks.

## 6 Acknowledgements

This work is the culmination of a long collaborative effort with many people at LLNL. The ALE3D team has been extremely generous with their time and resources, in particular, Richard Sharp, Barb Kornblum, Juliana Hsu, Scott Futral, and Rob Neely. Dave Morgan and Pat Egan were very generous with their time and the 1561B data. Norm Back has been a great help in tutoring us in the intricacies of radiography. We thank all of you for your help. Research at LLNL is supported under Contract W-7405-ENG-48 from the US Dept. of Energy.

## References

- [1] *The ALE3D Website*, <http://www.llnl.gov/bdiv/ale3d>.
- [2] *Accelerated Strategic Computing Initiative (ASCI)*, <http://www.llnl.gov/asci/>.
- [3] M. B. Aufderheide, H. Vantine, P. O. Egan, and D. Morgan, *X-ray Imaging and Reconstruction of a Detonation Front*, NIM A, (1998), in press.
- [4] E. A. Chandler, "Use of the Pegasus Z-Pinch Machine to Study Inertial Instabilities in Aluminum: A Preliminary Report", in Proceedings of the Sixth International Workshop on the Physics of Compressible Turbulent Mixing, G. Jourdan and L Houas, 1997.
- [5] *The COG Code System*, [http://www-phys.llnl.gov/N\\_Div/COG/](http://www-phys.llnl.gov/N_Div/COG/).