Los Alamos Technical Associates, Inc.
6501 Americas Parkway, NE, Suite 900, Albuquerque, NM 87110

LA-SUB--94-1166

DIGITAL SPALL RADIOGRAPH ANALYSIS SYSTEM

REPORT ON SIMULATED THREE-DIMENSIONAL DIGITAL SPALL IMAGE RECONSTRUCTION FIDELITY

By
Curtis L. Harris
January 1990

*This work is funded by the
Advanced Technology Assessment Center (ATAC)
at
Los Alamos National Laboratory
Contract Number 9-XF8-2109M-1*

MASTER

Los Alamos Technical Associates, Inc.

Telephone Number: (505) 884-3800
Telefax Number: (505) 884-9346

DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

A handwritten signature in black ink, appearing to be 'CH' or similar, located below the distribution statement.

DISCLAIMER

**Portions of this document may be illegible
in electronic image products. Images are
produced from the best available original
document.**

DISCLAIMER

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

DIGITAL SPALL RADIOGRAPH ANALYSIS SYSTEM

Digital Spall Image Reconstruction Fidelity

SUMMARY

This report describes progress on work to develop a cost effective, rapid response system for measuring momentum and kinetic energy of spall for the Advanced Technology Assessment Center (ATAC) Armor/Anti-Armor (A³) program at Los Alamos National Laboratory.

The system will exploit data contained in two sets of simultaneous co-planar flash radiographs taken along the center line of anticipated spall motion. Data contained in each set (which is proportional to the mass and z-number of the spall material intersected by the exposing x-ray at each point) is digitized and used to construct a three dimensional model (called the reconstructed spall image) that approximates the original spall cloud. From the model the mass of spall fragments is computed.

The two sets of radiographs, separated in time, represent the spall configuration at two instants of time. Spall fragments from the first instant are matched with those from the second instant to determine velocity.

Evaluation of the fidelity of candidate reconstruction algorithms is the highest priority task in this development program for the obvious reason that the efficacy of the projected spall analysis system depends upon the fidelity of the reconstruction techniques.

The purpose of this document is to report the results of analysis of the fidelity of best reconstruction procedure (for one radiograph set) investigated to date.

The reconstruction procedure uses data from four simultaneous radiographs representing two sides and two diagonals of a cube. The procedure makes use of an available space algorithm, two probabilistic devices (a mass placement probability heuristic, and a mass clumping heuristic), and a stochastic procedure for mass that

cannot be placed by the algorithm or either of the heuristics. The procedure is fully described in the body of the report.

Reconstruction is performed on image "slices", thin rectangular volumes perpendicular to the line of motion of the spall cloud that pass through the simultaneous radiographs at an interval corresponding to a narrow slice of the spall image). Reconstructed spall fragment sections are "projected" into the slice based on data contained in all four simultaneous radiographs.

Reconstruction fidelity is measured by comparing the reconstructed spall image to an original 3-dimensional spall image from whose radiographs the reconstruction derives. A detailed description of how the original images are generated and how the radiographs are made is contained in the body of the report.

Eight hundred seventy seven specimens (spall slice images) were generated, their radiographs were derived, and their reconstructions were computed from data contained in these radiographs. For slices containing few fragment sections (twelve or fewer) reconstruction fidelity was good; for five or fewer spall fragment sections per slice, reconstruction fidelity was excellent, the mean fidelity for this subsample exceeding .95 (where 1.00 is perfect reconstruction). This is significant because for spall radiographs that we have analyzed, in by far the majority of cases when slice width is chosen to be sufficiently small, the number of apparent spall sections is almost always sufficiently small to ensure high fidelity reconstruction.

Given this fact, we conclude that the performance of the subject procedure is acceptable as a basis of the planned spall analysis system and that four radiographic views provide sufficient data for accurate reconstruction. In addition, even for slices containing a larger number of fragment sections, the fidelity of reconstruction of the individual slice can be estimated based on analysis discussed below. That is, in all cases the fidelity of the reconstruction will be known to the spall analyst.

In addition simply to measuring the fidelity of each algorithm, knowledge has been gained relating algorithm performance to radiograph properties so that reconstruction fidelity can be estimated based on radiograph properties.

There appear to be three options as to how to proceed from here. The most cautious would be to conduct further reconstruction algorithm development and analysis. This option would develop slice-to-slice reconstruction enhancements; develop a second set of enhancements that more fully exploit data from the diagonal radiographs, only a portion of which is currently being used; and systematically analyze a large number of existing spall radiographs to verify that the number of fragment sections per slice can be limited sufficiently to ensure high fidelity results; and systematically analyze performance relative to mass contained in each slice.

The second option would be to accept the current results as sufficient and to proceed with development of a four-view quarter-scale system as an intermediate step to a full-scale system capable of detailed measurement of momentum of individual spall fragments and detailed analysis of mass and momentum distribution within spall clouds.

The third option (given the current state of funding, probably the most reasonable) would be to develop a two-view system to measure aggregate momentum of the spall cloud that could be upgraded to a four-view detailed analysis system at a later date.

DEVELOPMENT FRAMEWORK

In order to ensure that reconstruction fidelity is measured in a meaningful way, we have designed and implemented a development framework capable of generating a statistically significant number of spall image specimens.

Figure 1 diagrams the framework in which reconstruction algorithms are developed.

Synthetic spall images, produced by the spall image simulator, comprising some deterministic elements and some stochastic elements, are saved in a spall image store from which they can be drawn as many times as needed for analysis.

The synthetic radiograph generator computes radiographs simulating x-ray heads placed radially in a plane perpendicular to the spall specimen axis. These sets of radiographs are saved to be used as they are needed.

Figure 1. Spall Reconstruction Algorithm Development Framework

The image reconstruction algorithm module then uses synthetic radiograph sets to generate reconstructed images for the various algorithms in the reconstruction algorithm library.

The image correlation analyzer compares each reconstructed image with its corresponding original spall image and stores the results, along with properties of the related radiographs, in a reconstruction fidelity data base.

The analysis module then provides statistical treatments of the entire body of reconstruction fidelity data and for various subsets of that data defined by radiograph properties and reconstruction procedures.

Major elements of the spall image simulator are diagramed in Figure 2. The simulator allows the researcher to define position and fragment data that are relevant to specific reconstruction issues. Its analytical component (fragment position computation)

determines spall placement coordinates stochastically within constraints in the envelope and concentration specifications. Its discrete realization component (fragment selection) makes stochastic choices of actual fragments to be used constrained by a total mass specification and other selection criteria provided by the user. The spall image assembly module then produces a synthetic spall image which is held in the spall image store for analysis.

Figure 2. Spall Image Simulator

This spall image simulator is not intended to be nor is it claimed to be a simulator of the dynamics of spall generation. It is only intended as a tool for generating spall images with properties that are important to the problem of three-dimensional reconstruction. It is an implement for placing spall fragments in three-space for the purpose of algorithm testing. It incorporates no knowledge of how those fragments came to be where they are.

Synthetic spall images are used as baselines for algorithm development for three reasons. The first is that facilities for making n simultaneous coplaner radiographs of actual spall do not exist where n is larger than two or three. Second, the spall pattern "frozen" by the radiographic process exists only for a few tens of nanoseconds; thus, it

provides no baseline against which to measure the performance of reconstruction algorithms. Third, due to the statistical requirements of reconstruction algorithm validation, a very large number of specimens will be required; it is much more feasible to create these synthetically than on the ballistic testing range. Therefore, synthetic spall data are used as the starting point for reconstruction validation.

SYNTHETIC SPALL IMAGES

In order to understand what is represented by the material presented below, it is necessary to understand the conceptual framework from which that material derives.

Figure 3 shows the arrangement of radiographic elements presumed by the simulator. The direction of spall motion is perpendicular to the surface of the figure. A number of x-ray sources are arranged radially about the anticipated center line of spall motion. Radiographic film cassettes are placed opposite the x-ray sources. When the spall cloud enters the cylinder in the field of view of the radiographic elements, a timing mechanism causes the x-ray sources to activate simultaneously, exposing the film for a few tens of nanoseconds. This procedure creates radiographs which contain the information needed to measure spall mass.

In the algorithm development testing system, a three-dimensional synthetic spall image is generated in a memory array within the computer. For the purpose of reconstruction algorithm testing, this synthetic spall image represents the spall cloud. Synthetic radiographs are generated by computing ray intersections with volumetric pixels (voxels) comprising the synthetic spall image.

For ease of representation both for viewing and analysis the three-dimensional synthetic spall images are divided into slices one voxel thick.

Figure 4 illustrates the correspondence between a coordinate space slice and a one pixel wide strip of one of its radiographs. The simulator generates synthetic spall images

Figure 3. Schematic Arrangement for Exposing Multiple Radiographs Taken at Same Position in Line of Fire

Figure 4. Reconstruction Geometry

as a collection of slices containing distributions of spall fragment sections as viewed from the target (not from the point of view of the radiograph).

RECONSTRUCTION APPROACH

Perfect reconstruction requires the number of independent radiographs to be equal to the order of the reconstruction space. The resolution needed for this system may drive the reconstruction space order to be 4096 or greater. It is clearly impractical to configure a radiographic system with this number of x-ray sources. The price that is paid for reconstructing from a lesser number of radiographs is that the resulting reconstructions will contain uncertainty.

Development and verification of reconstruction algorithms that approximate the original spall fragment mass and position with acceptable fidelity from n radiographs where n is economically feasible is the specific challenge of the current work.

A number of reconstruction algorithms and heuristics have been developed using two and three views. As anticipated, the fidelity of these was found to be insufficient to support meaningful spall cloud reconstruction. The results for three views, however, were encouraging enough to motivate examination of procedures based on larger numbers of views. And some ideas developed for two and three views were incorporated in a four view approach.

Reconstruction is performed on one slice at a time. The reconstruction space is a rectangle having a number of cells equal to the number of cells in the slice. Data from the related radiograph strips are used to determine the placement of mass in the reconstruction space according to the procedures described below. These procedures depend upon two important assumptions. The first is that the spall consists of a single x-ray absorbing species (although work is planned to generalize the procedures to accommodate multiple species via multi-spectral x-ray methods).

The second is that pixel sizes and light source intensity have been calibrated in such a way that the digital image of a volume of material one pixel in cross-section and the thickness of one half that of the dimension of the smallest object of interest is the digital threshold for digitizing the image. This assumption allows for reconstruction by unit masses.

Reconstruction by unit masses means that in the reconstructed image each cell (voxel) either contains a unit mass or nothing. This may seem to be a severe limitation, but it must be remembered that reconstruction takes place in a digital space and that the amount of error resulting from this treatment can be suppressed to a level that is tolerable by "increasing the fineness of the mesh", that is, by making the pixel/voxel size sufficiently small.

Throughout the following discussion reference will be made to the example shown in Figure 5 Original Spall Slice Example and to other figures constructed to help explain some aspect of the reconstruction process or its results. This is a spall slice drawn from the specimen set the analysis of which is contained in the next section of this report. It was chosen because it illustrates all of the features needed in the section below and because it was well but not perfectly reconstructed. The view in the radiograph is as seen end-on from target, not parallel to the line of fire as most radiographs are made. The objects in the figure are spall fragment sections (i. e., the portion of spall fragments lying in the planar volume of the slice. Since the slice is one pixel thick, the sections are one pixel thick (thus, they are all represented in the same color). The reader should be careful not to confuse this figure with a radiograph.

The example's reconstruction was based on the data shown in Figure 6 Slice Radiographs for Example in Figure 5; in this figure different colors represent different quantities of mass penetrated by the x-rays, the blue end of the spectrum representing greater mass, the red end less mass.

The first step in reconstruction is to develop an initial available space matrix for the slice to be reconstructed. The available space matrix is the same size as the reconstruction space. Its cells are initialized by multiplying the values of the pixels from each of the radiograph strips corresponding to that cell together and setting the cell to zero if the product is zero and to one if the product is non-zero. Every cell that contains a one, then represents a voxel in which it is possible to place a unit mass. Cells with zero value represent voxels where mass is prohibited by the data contained in the radiograph set. This is the first probability cut. It is a rough cut, to be sure. For, in general, its non-zero cells contain many more spaces than there are unit masses to place, but the mass placement prohibition it imposes is of great importance. The initial

Figure 5. Original Spall Slice Example

Figure 6. Slice Radiographs for the Example in Figure 5 (in order from top to bottom a) side view, b) upper left diagonal view, c) top view, d) upper right diagonal view, and e) mass color key)

available space matrix for the example is congruent with Figure 7 Reconstruction Probability Space, except the available space matrix is monochromatic, representing zeros and ones.

The available space matrix is updated each time a unit mass is placed by one of the procedures below.

The first procedure is an algorithm called determined mass placement. This procedure is exercised when the sum of the non-zero cells in an available space matrix vector (row, column, or diagonal) is exactly equal to the value of the pixel of the radiograph strip corresponding to that vector. When this is the case, there is one and only one way in which the number of unit masses represented by the pixel value can be distributed along the vector. This procedure places a unit mass in each reconstruction space cell corresponding to a non-zero cell in the determined vector. This procedure takes precedence over all of the others.

The reconstructed image of the example slice is shown in Figure 8 Reconstructed Slice. The blue areas show sections reconstructed by this algorithmic method. Pink areas indicate portions reconstructed by one of the procedures discussed below.

To deal with situations in which the determined mass placement algorithm cannot be used, heuristic procedures are needed. The first of these is a quasi-probabilistic procedure that depends upon a weighted version of the available space matrix, which is called the reconstruction probability space. (See Figure 7 Reconstruction Probability Space.) In the figure areas shown in colors toward the blue end of the spectrum are more probable locations for mass than those shown in colors toward the red end. Mass placement is forbidden in the areas shown in black. This space, like the available space matrix, is congruent with the reconstruction space.

Each of its cells is initialized by first computing the sum of slice radiograph pixels corresponding to it, and then multiplying that sum by its counterpart in the available space matrix. The theory is that mass is more likely to occur in the slice along vectors whose radiograph pixel values are large than where they are small and most likely to occur at points along those vectors where they intersect other vectors with high pixel value. This heuristic provides a method for identifying those intersections (i.e., cells).

Figure 7. Reconstruction Probability Space

Figure 8. Reconstructed Slice

When no further mass can be placed by the determined mass algorithm, this heuristic procedure is attempted. It is actually exercised when there is a maximum value (i.e., a single value greater than all others) in the reconstruction probability space. In this case, a unit mass is placed in the counterpart reconstruction matrix cell. After that unit mass is placed, the corresponding available space matrix cell is set to zero to indicate that cell is no longer available, and all of the corresponding radiographs are also updated to show one less unit mass to be placed. The reconstruction probability space cell is set to zero as a result of the zeroing of the counterpart available space matrix cell. This heuristic takes precedence over the remaining two.

When no maximum reconstruction probability value exists, a second heuristic, the clumping heuristic, is invoked. This heuristic maintains another weighting of the reconstruction space in which each available cell is given a weight to indicate its proximity to mass that has already been placed. This procedure operates on all of the cells whose reconstruction probability is maximal (i.e., one of the set of cells whose reconstruction probability is greater than the rest). Within this "most probable set" it evaluates the clumping weight. If a single cell has a clumping weight greater than others in the set, it places a unit mass in the corresponding reconstruction space cell, after which available space and mass bookkeeping is as above.

When no maximum clumping weight exists within the maximal reconstruction probability cell set, the last heuristic is invoked. This procedure makes a random choice of a cell within the maximal clumping weight subset of the maximal reconstruction probability set and places a unit mass at the appropriate location in the reconstruction space.

It is from this procedure that reconstruction uncertainty arises. It is not, however, the case that all erroneous mass placements are actually committed by this procedure. When this procedure misplaces a mass unit, that misplacement may cause the other procedures, even the determined mass algorithm to make errors as well. To investigate the magnitude of the uncertainty resulting from this process, the work described in the following section was carried out.

EVALUATION

In order to evaluate the performance of candidate reconstruction algorithms it was necessary to devise a metric that is directly meaningful with respect to error in momentum measurement. This meant that the metric needed a spatial error component and a mass error component.

The most obvious and simplest metric is Point Wise Fidelity, that is, the sum of the products of corresponding cells of the original and reconstructed slices divided by the number of non-zero cells in the original slice. This metric was rejected, however, for three reasons. First, it has no clear meaning with respect to momentum because it does not distinguish spatial and mass errors. Second, it is sensitive to shape and orientation differences between original and reconstructed fragments, elements that are not relevant to momentum. And, third, a double penalty is incurred for a misplaced unit mass, once for the cell it should occupy and again for the cell in which it was actually placed.

The metric chosen for evaluation is called the Fragment Weighted Sum Fidelity (FWSF). The FWSF is computed in three steps. First, each reconstructed fragment is matched with its presumed antecedent; this correlation is made on the basis of best fit of mass and center of mass for the fragments. This matching is illustrated for the example slice in Figures 9a Color Coded Original and 9b Color Coded Reconstruction. The colors indicate which reconstructed fragment is matched with which original fragment. The two small dark blue dots barely visible in Figure 9b are misplaced point masses that have no antecedents.

Next, the errors in fragment mass and position are computed. The mass error is equal to the mass difference divided by the mass of the original fragment. The position error is equal to the difference in centers of mass divided by the diagonal of the slice. The fragment reconstruction error, then, is equal to the product of the mass error and the position error.

Last, the the weighted sum error for the slice is computed as the sum of the products of fragment reconstruction errors with their corresponding masses divided by the total mass in the original spall slice. The FWSF is equal to one minus this weighted sum.

Figure 9. a) Color Coded Original, and b) Color Coded Reconstruction

Eight hundred seventy seven specimen slices were generated for the analysis. In Figure 10 Reconstruction Performance, FWSF for all of these specimens is plotted against the number of fragment sections per slice and the trend of these data is shown in Figure 11 which plots the mean FWSF as a function of number of fragment sections per slice.

These data are particularly significant in view of an analysis we performed of actual spall radiographs which found that when slices are made sufficiently narrow, the number of fragment sections per slice rarely exceeds five. This was true even for radiographs with many hundreds of spall fragments. Slices can be made as narrow as

Figure 10. Reconstruction Performance by Fragment Section Count

Figure 11. Reconstruction Performance Trend

.02 mm (as measured on the film; this is even smaller object width, due to magnification). This implies that the mean FWSF can be expected to be approximately .95 most of the time.

When the number of fragment sections in a slice exceeds five, the reconstruction fidelity can be improved by using data from neighboring slices with fewer fragment sections. This procedure takes advantage of a clumping assumption similar to that used by the in-slice reconstruction procedure. It assumes that fragments extend across slices and uses data from neighboring high fidelity reconstructed slices to modify values in the reconstruction probability space to improve the likelihood that mass will be correctly placed by the procedures described above. Although this hypothesis has yet to be tested quantitatively, that it will have some beneficial effect is intuitively obvious.

Figure 12 Reconstruction Performance by Mass plots FWSF by total aggregate mass in the original slice. This graph tends to reinforce the conclusion that the reconstruction procedure described above works well in relatively sparse mass conditions. Since relatively sparse mass conditions tend to obtain for spall and since that tendency can be reinforced by narrow slice selection when necessary, these are encouraging results.

On the basis of all of these results, we conclude that high fidelity reconstructions can be achieved from four linearly independent simultaneous radiographs using the procedure described above and that the fidelity can be estimated from the number of fragments and the amount of mass in each slice. Thus, in most cases acceptable reconstructions can be achieved, therefore accurate momentum measurements can be made most of the time, and experimenters will know when the reconstruction fidelity is acceptable and when it is not.

We return briefly to the question of how best to measure reconstruction fidelity. Figure 13 Comparison of Fragment Weighted Sum Fidelity with Raw Point Wise Fidelity plots these two measures for each of the specimens. As is apparent from the plot, there are many cases where one measure is high, the other is low. This disagreement is due to sensitivity of the later to reconstructed fragment section shape and rotation difference from original. This comparison reinforces the decision to reject a point wise measure for one that has a direct physical interpretation.

Figure 12. Reconstruction Performance by Mass

Figure 13. Comparison of Fragment Weighted Sum Fidelity with Raw Point Wise Fidelity

Figure 14 Comparison of Fragment Weighted Sum Fidelity with Fragment Sum Fidelity plots the FWSF against a non-weighted sum alternative measure. The two are well correlated, but the simple sum appears to give an overly optimistic picture. So, although it makes small difference, we prefer the weighted sum.

Figure 14. Comparison of Fragment Weighted Sum Fidelity with Fragment Sum Fidelity

The FWSF has two shortcomings. When separate original sections are merged in the reconstruction, its fragment matching step can become confused and make anti-intuitive section matches. And, the denominator of the positional error component (the diagonal of the reconstruction space) tends to force that component to be smaller than it should be, thereby inflating the overall fragment reconstruction fidelity.

In the former case, for all instances we have examined, the effect has been to yield a falsely low FWSF, which slightly depresses performance as viewed statistically. This does no harm, and probably beneficially moderates the natural enthusiasm of the developers for the performance of their own reconstruction techniques.

In the later case, we are planning to revise the positional error component by using the effective diameter (i. e. the square root of the area) of the initial available space matrix in place of the diagonal of the reconstruction space as the denominator of this error component. The effect of this change is expected to be a change of approximately 2% in the FWSF.

All in all, however, we think that the FWSF is a good measure of performance as it stands.