

X-RAY IMAGERY AS THE RECORD OF ALL DATA OF INTEREST IN HYPERVELOCITY IMPACT FRAGMENT STUDIES

John B. Bacon⁽¹⁾, Anthony R. Allen⁽²⁾, Jose M. Ferrer⁽²⁾, John N. Opiela⁽³⁾, Melissa A. Ward⁽³⁾

⁽¹⁾ NASA JSC 28101 NASA Parkway Mail Code XI5, Houston TX 77058 USA: John.Bacon-1@nasa.gov

⁽²⁾ University of Florida, Gainesville FL USA: allen.anthony@ufl.edu, joseferrer@ufl.edu

⁽³⁾ Jacobs/JETS Contract NASA JSC 28101 NASA Parkway Mail Code XI5, Houston TX 77058 USA
John.N.Opiela@nasa.gov, Melissa.A.Ward@nasa.gov

ABSTRACT

Laboratory study of hypervelocity spacecraft fragmentation has traditionally involved the collection and analysis of fragments that were caught in deceleration material surrounding the impact. This process has typically involved the disintegration of the catchment material either through chemical dissolution, or through physical excavation to recover the fragments. Due to the scale of the three impact tests—the Satellite Orbital Debris Characterization Impact Test (SOCIT), the DebrisSat satellite impact test, and the DebrisLV launch vehicle impact test—the latter two using more than 12 cubic meters of polyurethane foam to capture the fragments, these projects have used x-ray imagery to precisely locate and thus, to more efficiently extract fragments in the soft-catch material. Three years into the DebrisSat fragment extraction process, a side study was initiated to explore what additional information could be discerned from the x-rays, with significant results. This study was instrumental to a rapid replacement and retooling as the project was forced to replace the x-ray system around which the extraction process had been based.

The revised process continues to map the debris for extraction. The project has, in parallel, systematically addressed the limits/tolerances of what x-rays can reveal about size, shape, density, mass, velocity, energy, and deformation/damage of the fragment during the deceleration in the catchment material while replicating the original extraction mapping function. All of these features have been optimized or have sufficient understanding to characterize the basic factors that will define a complete data set extracted solely from x-ray imagery. It is an ideal time to develop such a process, with extracted fragments providing “ground truth” against image-only data, and abundant available imagery of the same fragments under both the prior and replacement x-ray technologies, which have several fundamentally different characteristics.

This paper addresses the types and quality of hypervelocity fragmentation data that can be and has been extracted from x-rays. It further addresses the question of whether and under what circumstances future

hypervelocity experiments can use x-ray methods to largely—or to completely—avoid the extraction process in recording all appropriate results. Lastly, this paper addresses lessons learned and how future efforts can be further optimized.

1 BACKGROUND

When the DebrisSat project set the requirements for its soft-catch materials, it built upon the earlier work of the SOCIT test [1]. The original foam material used in SOCIT was no longer available, so the remaining choices were narrowed to polyurethane, polyisocyanurate, perlite, and common acoustical ceiling tiles. Of these, only polyurethane provided minimal friability, the desired panel densities, and a uniform cross-section. It was understood from SOCIT experience that the benefits of this material for preparing and performing the impact test could also pose difficulties for post-impact extraction of the fragments. The current study has evolved techniques that could make these issues moot.

For both the precursor DebrisLV and for the main DebrisSat impact tests, the chamber’s “soft-catch arena” was constructed from bundled stacks of foam panels with varying density. Fig. 1 shows the DebrisSat target suspended within the test chamber before the impact. The upper (radially inward) panels of each stack had the lowest density (0.048 g/cm³, 3 lb/ft³), attempting to slow the fragments with minimal damage. The middle layers of each stack were of medium density (0.096 g/cm³, 6 lb/ft³) and the lower (radially outward) layers were of high-density foam (0.193 g/cm³, 12 lb/ft³). The lowest layer of each stack, facing the test chamber wall, was plywood. To accommodate existing materials, fixtures, and handling methods, each bundle measured 2 feet wide by 4 feet long. Bundle depth and layer composition varied by position within the test chamber, whether uprange, side wall, or downrange. The uprange partial wall and the side walls are visible in Fig. 1.

After the impact test, all foam bundles and loose debris were removed from the test chamber and shipped to the University of Florida at Gainesville for fragment extraction and characterization. SOCIT data analysis [2] showed that x-ray imagery could reveal the locations and relative sizes of fragments embedded within the foam

panels, so this method was selected to identify fragment locations. Soft-catch bundles were separated into their component panels.




Figure 1. Pre-collision DebrisSat and catch foam arrangement. Each foam region is a stack of varying density panels seven or more layers deep, illustrated at right. The project seeks to characterize all parameters of over 215,000 fragments > 2 mm in largest dimension embedded in the 400+ large foam panels



Figure 2. A scanned panel's negative x-ray image showing embedded debris of different (brighter) colors associated with their material and thickness as well as many (darker, larger) cavities resulting from the ballistic penetration of the embedded fragments. The methods discussed in this paper use the color and saturation of discovered fragments to derive their properties and full 3D shape (Fig. 4). We use cavity information (see Fig. 14) to assess the fragments' initial energy, deformation and breakage during capture. Over 1300 fragments > 2 mm are identified in this panel. Previous extraction and measurement techniques (Fig. 3) would take one person several weeks to record the same data that is available in minutes—along with a wealth of contextual data—under the alternative x-ray approach.

Using an x-ray camera normally used for civil engineering small-sample material measurements, a series of images was composited into a single image of a given panel (Fig. 2). There was some residual parallax caused by the multiple wide-field-of-view images and a significant series of registration wires that blocked part of the view. Automatic feature-detection software identified fragments in the image (Fig. 3), which were then projected as a series of “find” boxes onto the actual foam panel to locate the fragments. Because of

substantial residual parallax these boxes could only bound regions where the fragments were certain to lie. Then locations of embedded fragments were marked and carefully extracted from the foam. The location process was two-dimensional, leaving the fragment depth unknown. Details of this process can be found in [3]. In the original process, the x-ray images were stored but not put to any further use.




Figure 3. The original x-ray system created in a stitched assembly of 15 images with an obtrusive registration grid. Regions of suspected debris were marked in red boxes. This image was then projected on the panel itself to allow physical marking of the panel for later extraction. (From Shiotani, et. al, 2017 [4])




Figure 4. The raw image of a captured 7 mm nut in the foam, next to the autonomously generated 3D probability map of expected filled voxels under the processes described in this paper. Note the clear, regular hexagonal exterior, < 2 mm thick walls, and circular center emerging purely from analysis of sets of exceptionally noisy raw x-ray images (four each in two views). The actual extracted object is shown. The middle 3D voxel rendering is > 3x more refined than the simplified rendering that is used to generate and compare shape models for the next generation of orbital debris modeling. The methods discussed herein are thus likely capable of generating the global data set for fragment shape analyses, pending final statistical validation of the software.

1.1 Original Side-Study of X-rays as a Data Source

In the spring of 2019, a one-semester internship was created to explore expanded science opportunities for archived images during a long wait for repair of the previously mentioned x-ray system. The initial goal of the side study was the exploration of the data types and of the quality available from the x-rays. The side study's

minimum goal was to use fragment data processed only from x-ray images to create a 3D database of each fragment's size and location at its terminal location within the test chamber, with a goal of building an energy distribution and isotropy map.

Such data was not available in the extracted fragment database, because the original location had been so vaguely defined and recorded to an even lower precision into rough grids approximately 20×30 cm on a side. Even with the residual parallax of the original x-ray stitched image set, the photographic coordinates of the silhouetted shapes had a much better characterized X and Y location than the official database. The depth was known only to within the thickness of the panel. The shapes proved to be remarkably crisp, allowing the team to examine statistics of shape and size in addition to 3D location within the tunnel.

As the team spent more time examining the details of the images, it soon became evident that there was a strong potential to completely eradicate the need for manual extraction of foam panels, potentially relying on x-rays as the sole data source. This was especially intriguing because broad statistical information could immediately be gleaned on particles substantially smaller than the 2 mm minimum extracted size. Thus the “old” x-ray data immediately was useful in extending the statistical data set to lower sizes. Significant work beyond the scope of the internship would be required to bring the full development of x-ray data to fruition, however. The in-depth study of x-ray imagery was a fortuitous advantage when the original x-ray camera was declared to be irreparable.

The 3D mapping of the debris within the test chamber was straightforward. By mapping every foam panel's corner coordinates into the test chamber, the 3D plane of the image's thin-sectioned 2D debris field could be precisely mapped relative to the center of the impact (Fig. 5). This method could either provide a pixel-by-pixel mapping of each debris item in 3D-space, or (more rapidly) a map of color-coded spheres of equivalent diameter to match the fragment's size. The fragment silhouette size and locations could be directly entered to the 3D database based solely on image processing and a small software routine. An example of a rectilinear map is shown in Fig. 5, but the debris could similarly be mapped by penetration depth and radial coordinates. From this initial data set it became clear from 3D graphic rotations and alignment (Fig.6) that at least some of the fragments were breaking in the process of capture.

Only a partial data set was achieved, representing the extent of archived images at the time of the original camera's failure. This partial data set will in the future be combined with the data from the other 50% of the panels obtained from the replacement system.




Figure 5. Color-coded debris field map of the DebrisSat test, generated from archived panel x-rays (Allen and Bacon, [7].) Densely populated front panels (blue fragments) dominate the coloring, but other depths are visible. The data can similarly be represented in radial coordinates around the target center with the radius to each fragment correlated to foam penetration depth, a parameter correlated with its initial energy. The archive contained only a subset of debris-embedded panels.

The initial system was a scientific-grade x-ray imager with 0.09 mm resolution and 16-bit contrast depth, covering a field substantially smaller than any one panel. Fragment size, location, and 2D shape data could be extracted from each image scanned. Debris boundary information in this phase of the project was gathered via the use of the Canny Edge Detection algorithm described later in this paper [6]. By searching for local changes in intensity gradient throughout each x-ray image, fragment shadows occupying even less than 10 pixels (~ 1 mm) could be detected. It was a naive goal, perhaps, to correlate the data collected from many of the tens of thousands of already excavated fragments to calibrate the accuracy of the image-based fragment detection models. The biggest issue in this plan lay in the inadequate precision within the official database identifying a fragment's origin once it was extracted, making cross correlation exceptionally difficult.

Concurrently, the team was troubleshooting errors with the physical x-ray system that arose from widely varied exposure levels in each of the contributing frames. This contrast problem resulted from high parallax in wide-angle “macro” exposures provided by the camera. The dramatically varying background led to the selection of gradient-based edge detection in this early project, with a goal for more uniform background levels in the new system.

Although it was very difficult to do correlations of the x-ray images with their recorded lab-measured fragments, the 3D x-ray-only database offered preliminary insight into the “big picture” statistics of the breakup. Not only were signs of bending and breaking of fragments made apparent, but the extent of these fragment deformities as

they decelerated to full capture could be measured as well. In studying the full-panel debris fields and the previously ignored cavity areas, numerous geometric effects surfaced that were not immediately visible before (Fig. 6). The techniques studied during the internship indicated that near-instant size and energy distributions, spatial anisotropies, and detailed physics of capture could be derived from the same raw image used for simple extraction mapping. [7]. Several months were consumed in developing preliminary approaches to extracting such data from the archived images, before it became clear that an entirely new process would soon be necessary. When the opportunity arose to acquire a new system, many process improvements immediately suggested themselves from this initial work, and all efforts focused on developing the new system.




Figure 6. Context for the debris field is assessed across adjacent panels. Debris is color-coded in this example by the depth of its embedding panel, and the view is aligned along the nominal flight vector from the center of the impact. Note among other dense groupings the broken wire fragment in upper right, with three pieces found in adjacent panels. These were clearly broken apart during capture, not in the original impact.

2 NEW PROCESS

The immediate goal in the hardware replacement was to re-create the process for identifying the locations of embedded debris. However, the lessons learned from the internship led to additional goals and enabling requirements that could further refine the prospect of image-only data collection.

The recovery steps were as follows:

- 1) Acquire a new x-ray imager (A surplussed airport luggage scanner).

¹ The improvements included: 1) removal of all parallax for better 2D location, 2) triangulation for accurate depth, 3) painting of the actual sought fragment with color-coded material, 4) in-frame identification of and subsequent file naming relative to the particular panel being imaged, and 5) the creation of a virtual grid with alignment references that would allow correlation to a real (physical) grid overlay during the extraction process.

- 2) Characterize and calibrate the scanner for scientific use.
- 3) Build infrastructure to assure precise repeatable alignment for several engineering and science objectives.
- 4) Develop software to replicate and improve¹ the extraction mapping function
- 5) Develop enhanced software to extract all further available data from the x-ray image.

Steps 2-5 were concurrent and evolutionary with step 4 being the most urgent goal. Progress was significantly hindered by the 2020 COVID pandemic. The project focused its software development cycle to analyze re-imaged, previously extracted panels to confirm that no debris had been missed. (This independent confirmation of complete extraction allowed the project to clear many cubic meters of needed storage space.) Further, a single front-layer, low density, heavily embedded panel (inventory #DSF005: Fig. 1) had been imaged by the original camera but not yet extracted. This became a key item for the re-tooling of the system, and allowed a direct image-to-image comparison of hundreds of fragments. The combination of previously-extracted dense panels and the one yet-to-be-extracted front panel provided a rich data set that allowed the new system to be evaluated head-to-head with its predecessor.

The subsequent sections of this paper outline the details of the revised process, the quality and limits of the data that can be provided with this technology and prospects for future capabilities of this method.

3 NEW SCANNER

The acquired surplus luggage scanner, shown in Fig. 7, is a 2010 Smiths Detection HiScan 6046 standard airport hand luggage belt-driven scanner² with a tunnel cavity 60 cm wide by 46 cm high, although the latter dimension is heavily encroached by lead-lined flaps that provide significant lateral load to any object transported into them. The more these flaps deflect, the larger the resistive load.

The differences between the prior x-ray camera and this replacement scanner are significant and many. The new system is a linear scanner, working more like a fax machine than a traditional camera, recording one column of pixels at a time while the imaged object moves through a vertical sensing plane defined by bounding sensors.

² When it was discovered that surplus air transportation luggage screening x-ray scanners could be acquired within the government for the cost of transportation, NASA acquired one immediately from the U.S. Transportation Security Administration (TSA), and it was installed within the panel storage and analysis lab.

This feature allows for an arbitrarily-long image with identical parallax in every column and essentially no parallax issue in any image row. Software easily removes the (identical) parallax in every column. Properly oriented, a full panel can be rendered in a single image in seconds, vs. the 15 images necessary in the prior system.

The new scanner has approximately 10x lower resolution in each dimension than its predecessor but still is just within the resolution limits necessary to achieve project goals of recording all fragments ≥ 2 mm in longest dimension.




Figure 7. The luggage scanner (left) and custom aluminum/steel/polycarbonate shoe (right). The shoe registers the debris-laden foam panels precisely within it, provides necessary beam energy filters, and shuttles like a piston with tight tolerance through the scanner. External roller tables and guides assure alignment at both ends. Blue weights at each end of the shoe provide needed traction against the lead-lined scanner curtains, and the ramped metal tabs at each end of the shoe help to mitigate accelerations, decelerations, and even belt wear under the resistive loads of those flaps.

Most important of all the distinctions is that the luggage scanner is a dual-energy system vs. its predecessor's single energy beam. The dual beam system is what lets a luggage detector identify the scanned material as a ratio of the absorption in two different wavelengths. This feature cannot discriminate one hydrocarbon from another, or titanium from steel, but it conveniently registers colors associated with the three material densities in NASA's orbital Debris Engineering Model (ORDEM): 1) low density (plastics), 2) light metals, and 3) dense metals.

4 CHARACTERIZE SCANNER FOR SCIENCE

The repurposed luggage scanner was not originally intended for the recording of detailed scientific data, especially at the extremely small scales that dominate fragmentation debris populations. Neither was it anticipated that multiple views would be needed to

³ This repeatability was necessary both to stack repeated exposures for signal-to-noise improvement and to precisely

correlate and triangulate to sub-millimeter precision.³ The registration issue is addressed in the next section with the discussion of the "shoe" and other infrastructure. Even with this reliable registration, however, the overall calibration of this high-parallax system represented a significant geometry challenge, akin to conducting a full full-volume survey with heavily distorted, wide-angle lenses in originally unknown locations relative to the imaged panel.

The scanner has a pair of co-planar sets of linear sensor rows along two adjacent edges of the tunnel, which receive a fan of x-rays from a common source deep below and to one side of the scan volume. Early attempts at calibration centred on the idea of building a geometric model of the panel within the scan tunnel and ray tracing from first principles to survey each recorded pixel. This ultimately proved futile, because the source is apparently not a pure point, amongst several other unknowns in the proprietary design of the machine. The simple geometric model would not converge, even with significant "fudge factors" to accommodate the observed nonlinearities.

Ultimately the system was calibrated empirically as a black box, using flat mesh grids to generate thousands of known correlated points at three known heights above well-registered tray locations. The empirical curves thus generated correlated points observed in the A and B views at three heights. These correlations on three parallel planes were used to interpolate the true height of each observed fragment inside a very small subset of the tunnel's volume.

A vertical location accuracy of approximately 1/8 of a millimeter is achieved by this method. Such accuracy is enabled by the tight tolerance in maintaining consistent and repeatable registration of the foam panels during the switch between the two views. This and other registration demands engendered a special focus on precision in developing the shoe.

The initial 3D shape boundary definition (known as "shape carving") of each sub-centimeter sized object relies upon the angle between the intercepting rays in the A and B views. While there are significant difficulties in calculating these vectors exactly, it has been found that the overall uncertainties in the 3D shape definition are driven by exposure noise and insignificantly, by small uncertainties in intercept angle. Thus, the approximate intercept angles from the geometric ray tracing exercise are used in this facet of the analysis.

5 BUILD ADDITIONAL INFRASTRUCTURE

To effectively utilize the scanner, a mounting rig called the "shoe" (Fig. 8) was developed to both protect the

triangulate and determine the shapes of the embedded debris objects.

delicate lightweight panels from the heavy lead flaps of the scanner, and to provide sub-millimeter precision in its task to repeatedly align and register each panel for multiple views. Thus the rig needed to be massive and snug within the scanner tunnel, similar to a piston. Rollers and alignment rails assure that this piston-like device smoothly enters and exits the tunnel in the same orientation on every pass. Precise interior mounting techniques suspend each foam catch panel in only two reposes that are used to both triangulate the location of the debris, and to create a rudimentary tomographic recreation of each item's 3D shape. The weights were adjusted to provide best possible traction on the belt (motive force against the retarding flaps) without causing the belt to slip on its drive rollers.




Figure 8. Shoe loaded with panel showing all key features

The primary function of the shoe is to provide two precise survey views of the panels that can be later correlated exactly to tolerances much smaller than the smallest dimension of the imaged fragments. The consistency and reliability of the scans have been ensured through multiple modifications and upgrades. The lid, shelf, supports, indexer, weights, and guard rail all contribute to achieving this purpose, and each went through numerous evolutions.

We had originally imagined that the longitudinal distortions could either be the result of gradual accelerations and decelerations of the belt drive, or alternatively, by instantaneous sticks and releases of the belt. An indexer was implemented such that either mechanism could be detected and corrected. As it turns out, only the former mechanism appears to be applicable. Thus, the indexing algorithm ultimately measures the relative stretching of 6 mm-wide regions defined by the centers of the interstitial spaces between a tightly spaced raft of 6 mm-square machine keys. Software then compresses each corresponding 6 mm slice of the image to the same reference width in every shot. These slices typically have adjustments of only one or two percent width, meaning very little distortion at the pixel level. (In the expanded images shown in Fig. 9, a 6 mm gap is represented by 90 pixels. This wedge is compressed

uniformly, meaning that the residual “real world” distortion per pixel for each 1% compression step is less than one micron).



Figure 9. Indexer Slip Example. The blue vertical reference lines reveal a varying drift of several millimeters along the travel direction of four successive images of a 1.22-meter panel. The indexer (above the alphanumeric panel ID) allows the images to be aligned to 70-micron tolerance in the lateral direction. Coupled with the cross-tunnel registration assured by precision mounting techniques, these steps enable image stacking of multiple exposures to achieve enormous signal-to-noise improvement, and to assure that the two perspective views are exactly aligned so that millimeter-scale fragments can be exactly surveyed within the entire panel expanse.

The lid is a key feature of the shoe, as it is needed to protect the panels from the lead flaps at the start and end of the tunnel. The initial lid design consisted of a sheet of 3/16” paper-backed “foam core” that would rest across the top of the shoe. Foam core was selected in the belief that its low mass would create the lowest possible noise signature over and above the absorption of the panel itself. After further testing, it was determined that the whitening clay in the foam core’s paper face sheets was a dominant source of noise in the images, resulting from its heavy silicon and aluminium oxide content. The initial panel support system also had been made of foam core, more than doubling this unexpectedly bad source of noise. It was subsequently determined that stiff unfaced 1/2” polyurethane foam insulation (roughly 2.7 times the thickness of the originally-selected foam core, and half the thickness of the thinnest catch panel) provided little to no noise to our images and negligible net absorption. The 1/2” foam was thereafter used as the shelf material to support the scanned panels within the shoe.

Despite its solving of our shelf noise problem, the foam insulation was inadequate for the lid because the foam could easily bend and break after deflecting the heavy lead safety screens for even a few cycles. We quickly resolved that removal of the manufacturer-certified and warrantee-coupled lead flaps would generate untenable legal and potentially safety issues no matter how much work we might put into an engineering alternative shield system.

Another problem that the 1/2” foam insulation presented was its lack of x-ray absorption. The system was designed to scan luggage, and there was no such luggage casing material around our fragments, leading to

overexposure/saturation in both channels.

Both the structural integrity and some oversaturation issues were resolved by developing a polycarbonate lid. This lid fit over the top of the shoe and form-fit the leading and trailing edge metal deflection ramps as well as providing a vertical wall backing. This vertical section of the nominally horizontal lid was useful for both added structural moment (less deflection) and to continue its low-energy x-ray filtering function across all possible rows of the two foam views. Note that while absorptive, the polycarbonate is exceptionally uniform in its x-ray absorption, dramatically improving the random noise encountered with the foam core. It is a near-ideal filter.

Later, the team found that saturation and noise similarly were reduced for both the high and low energy beams by adding a thin sheet of aluminium under the shoe. This aluminium sheet acts a bit like a neutral density filter to bring the entire foam panel and its contents into the unsaturated absorption range of the scanner. In retrospect, a thicker sheet of aluminium may have been an optimum choice for a lid without an additional dedicated screen, but the final shoe configuration is fully calibrated and meets all requirements, so we have not elected to explore this option.

A method was needed to securely mount panels into the shoe and provide two positions in which the panels could be scanned. To accomplish this, necessary supports for the previously mentioned foam shelf were developed that would image outside of the largest panel area. As recounted above, the shelf was ultimately made of unfaced 1/2" foam insulation. The edges of the shelf outside of the panel image are reinforced with metal U-channel, and all handling and mounting interfaces have this aluminium reinforcement to mitigate the risk of damage and/or alignment creep to the delicate shelf itself.

The shelf supports are a modular set of three triangular polycarbonate plates, with the edges of the latter two aligning with the first plate to make different smooth, thin angled planes (*i.e.*, support #1 is a fixed triangular prism whose apex's sides define both planes. Supports 2 and 3 register against the lower triangle to smoothly continue the established two planes, one at a time, shown in Fig. 10.) These thin planar supports provide the two tightly registered views labelled "A" and "B" for our scans. The planes are 55 degrees from each other. The planes were chosen to get the lowest-parallax view possible for a 2'-wide panel (View A) and then the maximum angular separation possible in the same rig (View B), which is the demonstrably noisier view. The shelf supports were initially prototyped using 3/8" foam core but were ultimately upgraded to polycarbonate supports. These supports added increased rigidity and will not sag or otherwise degrade over time.

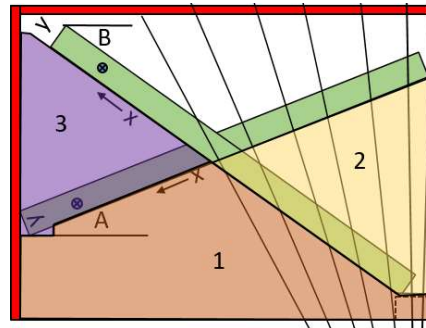


Figure 10. A cross section of the tunnel showing the numbered three support blocks and the A and B orientations of the same panel. A fan of x-rays from a source located deep below the bottom, right corner intercepts the red sensor arrays. A candidate embedded object is shown in each orientation at the left end of the panel. Precise sub-pixel registration allows us to survey the details of the fragment itself, in addition to merely triangulating its position in the panel.

The registration of the shelf to its supports is perhaps the most critical step in obtaining the required end-to-end data quality. This is akin to assuring that a survey theodolite is precisely centered over its reference stakes. The team iterated numerous variants before settling on the current system. Aluminum U channels at both ends of the foam shelf stiffen the lightweight material to take repeated loads, while 1/8" metal dowel pins rigidly register the shelf in an exactly repeatable way to the supports. The modular supports are similarly registered to the base with metal dowels. Shelf angle and position are thus assured in repeatable, standard longitudinal planes that are each normal to the detection plane.

A lesson learned from the prior x-ray work was that the image filenames of each panel were not well correlated with the foam panel imaged within. A separate database was required to keep these files correlated. Thus the shoe was designed to house large lead-sheet letters and numbers in an otherwise unused section of the frame to display the ID of the panel being imaged. Optical Character Recognition is used to name the adjusted images, to rename and sort the raw images, and to embed the panel name in all file names associated with the subsequent analyses.

The team iterated several hardware improvements to ruggedize and to simplify the end-to-end process. As of this writing, a panel's entire journey from deep storage through the process and back to storage takes well under 20 minutes using a single operator. Parallel preparation and imaging processes cut this time roughly in half when two operators are available. This compares to a physically remote, 1- to 2-hour process per panel under the prior system that required numerous alignments and considerable physical logistics.

6 IMAGE ANALYSIS: FINDING & EXTRACTION

The crucial function of the first part of the x-ray image analysis software is to identify and then to triangulate all identifiable fragments embedded in foam panels. With the background image noise reduced as far as possible by the optimization of the shoe filters, a variety of software techniques must then be used to best find real debris against that residual background noise. Large-scale and pixel-scale de-noising algorithms, edge detection algorithms, and precise thresholding are used.

A first software step is called image stacking: a practice common in astrophotography. In the stacking technique, precisely registered duplicate images are averaged at the pixel level. This has the effect of suppressing the (random) noise while leaving intact the (static) physical image. All of the shoe features are built to create these lowest-noise, precisely registered source images for stacking. A total of eight panel images (four in each of the A and B views) must be obtained to suppress the intrinsic noise in each view.

Additional precision and accuracy are gained by working with 10x expanded versions of the images. Prior to the stacking, each image is magnified by a factor of ten in each dimension, using a cubic spline interpolation between each original pixel and its neighbors. This spline smoothing has advantages in noise suppression, in precise edge-finding, and in sub-pixel shape determination. The expanded images have a pixel resolution of 70 microns, comparable to the raw images available from the prior scientific-grade x-ray imaging camera.

The residual noise level, in theory, performs like the inverse square root of the number of stacked images. Thus four images theoretically drop the noise by a factor of two. In practice, we see slightly better results.

We estimate that the scanner’s high- and low-energy beams appear to share a common generator and single power input. This means that the sum of the energies emitted in the two beams is generally stable, while the momentary noise in each individual beam is prone to larger noise. (A small surge in the low beam energy is generally congruent with a near-equal fall in high energy, at a pixel-sized scale.). The scanner’s internal software works to produce a compensating third color (green in the red-green-blue image, with red and blue holding the high- and low-energy information, respectively). This green saturation is created in such a mathematical way to result in the lowest-noise grayscale image, while red and blue appear to be tweaked to maintain a color well categorized by material, no matter the thickness of the imaged fragment (thin steel is still distinguished from thick aluminum or thicker plastic, e.g., see Fig. 11). Because the noise is noticeably lower there, we generally work in the grayscale within our debris-finding algorithms. We

do employ the color information to assess material type of the debris once we have confidence that the boundary is properly located. The high-energy beam is also referenced in some cases as a discriminator of the boundary between thin metals and the oftentimes highly compressed foam that surrounds them (discussed later in this paper).

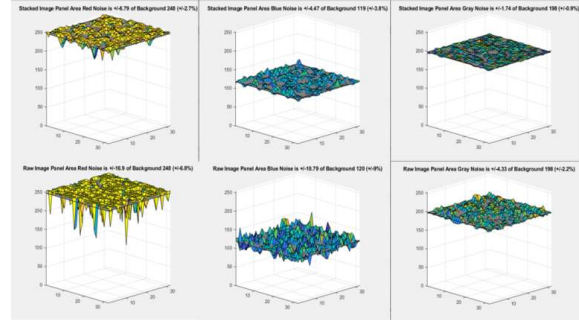


Figure 11. Background noise surfaces in (left-right) red, blue, and gray. The noise levels of a 4-stacked set are shown above, and of a single contributing exposure, below. The addition of a thin aluminium sheet below the test area was necessary to avoid saturation in the high-energy channel. This step reduced the relative noise in the image by 27% at the expense of saturating a limited number of pixels in the densest found fragments.

Other than these few cases, we have found that simple grayscale thresholding against a uniform background is the lowest-noise way to find the bulk of the discernible fragments in a panel. In optimizing this thresholding, we have an algorithm to remove the very gradual intensity changes in all three colors and grayscale over the width of the panel caused by several geometric and hardware effects. The result is a creamy smooth background of uniform bit level with all cavities and densifications beyond the residual noise level easily thresholded above or below it. We draw the actual color from the original x-ray image, however; since we believe that our manipulations can give rise to errors in depth estimation.

After image stacking and background levelling, one last trick is to look at the residual intrinsic “white” noise of the image. Here we run the danger of missing pixel-sized real items, but in general, such techniques are useful to suppress the randomized noise with little loss of information once we get to extraction of real features. During development, over 20 image de-noising methods were explored to improve image detection. A comparison of the explored algorithms is listed in Tab. 1. These de-noising algorithms were compared using the following metrics: average grayscale pixel value, the standard deviation of the pixel intensity (i.e., the average magnitude of the noise), and percent background noise. Ultimately, an edge-preserving blur with a threshold (the *weiner2* algorithm in MATLAB) was chosen as the optimal denoising algorithm. It uses a bilateral threshold

to make edges more or less pronounced. It works by enhancing dark fragments but decreases the intensity gradient around holes.

Finally, the team focuses special techniques on panel edges and cavities where the background gray level varies (sometimes dramatically) from the average panel background (Fig. 14). We employ software to recognize that the otherwise uniform background has changed, and then switch from grayscale thresholding to a local gradient-based, edge-finding technique called Canny edge detection in these limited regions. (The same edge detection method used in the original study of the old x-rays, which had substantial broad background gradients that made thresholding difficult).




Figure 12. Edge finding within partially filled pixels at the edge of a found item. The saturation and color are both in transition mid-pixel in the unexpanded left image from the found object that occupies one part of the original pixel to the background foam that occupies the rest. In the 10x enlargement (middle), thresholding identifies that the original pixel has some non-foam content, and suppresses all pure foam to a black background, showing the interpolated gradient of color and saturation. It is then possible to locate the edge within the interpolated 10-pixel field, not by grayscale but by color. The illustrated found object (right) is 10 mm long and 1 mm wide: likely a wire shred.

The Canny algorithm represents the best fit edge of an object to a single pixel tolerance. This enables data to be collected on fragments without any significant size or shape distortion, especially in the 10x enlargement. The function was parameterized via an exhaustive analysis of the various input parameter combinations.

Every “finding” algorithm was evaluated at each combination of parameters by comparing it to a “ground truth” image. This ground truth image was created by manually selecting visible fragments in the grayscale x-ray image under several contrast stretching steps. For each parameter combination, the ground-truth image and auto-detected image were compared via the use of a Precision-Recall curve. The precision of any such “finding” algorithm is defined as the ratio of true positives (all the real ones found) divided by the sum of true positives and the false positives: *i.e.*, all the

fragments the algorithm thinks it found. A precision of 0.85 says that 15% of the identified items are not real. Recall is the other key metric, defined as the ratio of the true positives found by the algorithm divided by the sum of all the real items the algorithm found plus all the real ones it missed. A recall of 0.9 says that the algorithm missed 10% of the real items that are actually there.

A precision-recall curve is particularly useful when one data class is significantly greater than the other (in this case foam pixels greatly overpopulate fragment pixels). Precision-Recall summarizes the trade-off between the true positive rate and the positive predictive value for a predictive model using different probability thresholds, and is our basis of self-checking the usefulness of all of our algorithms. While we optimized this under the new x-ray scanner and software, we do not have comparable performance assessments of the prior method. Generally, we achieve 90% precision and recall in our new method, with the false positives and false negatives dominated by the smaller hydrocarbon items. We are working to use context to remove compressed foam as it is the dominant source of false positive “fragments”.

6.1 Fragment Database Built from the Image

The border-finding techniques described above result in a data structure of the regions/pixels in the expanded image that are identified to be fragments. All pixel values and numerous region properties are maintained for each such fragment in a data structure for each view. Since the A and B views are precisely registered, it is a simple matter to correlate the data sets to identify and triangulate the same fragment in each view. These confirmed fragments are painted in their recorded material-dependent color in a 2D map of the panel with sub-millimeter precision in parallax-free X-Y coordinates. This map replaces the original “Find Box” map of the prior system (shown previously in Fig. 3). Similarly, it is projected onto the panel where felt markers identify the “dig here” locations. However, there are numerous improvements highlighted at the end of this section.

Of course, at some point we reach the limits of signal to noise when seeking fragments. In the limiting case of thin hydrocarbon fragments, many such debris items often are cleanly resolved only in the A view, which has a narrow range of ray angles near the normal in the panel, and thus a minimum of foam-related noise. These A-view-only fragments are also represented in the find map, with the color replaced by a bright cyan. These fragments are known only in their “A” row coordinate and exact column, which leaves ambiguous their actual depth, and thus their position along the slightly diagonal ray that connects them with top and bottom penetration of the panel. In the case of the cyan fragments, the team must core along the known ray angle to find the fragment. In triangulated (full color) fragments, the team knows to core straight down. Although the true depth is known, the

current process does not use this feature, and instead finds the marked debris purely by digging along the suggested vector. These techniques all represent a substantial improvement over the prior “search box” practice in guiding students to buried fragments for extraction.

A final improvement is that the software generates a reference grid of exact 1/2” squares anchored at a series of four registration pins inserted into the foam during the shelf mounting process. (The pins are located along the panel at a line that represents the position of the ray normal to the shelf; this leaves the registration pins as pure dots in the A view). Using the registration pins and a 1/2” wire mesh screen, the students can always cross-reference the extraction map to any chunk of the real panel that contains one of the reference pins. This method was particularly useful in debugging and calibrating the end-to-end process but remains a useful tool in some cases

Table 1: results of available image noise reduction software algorithms on imaged catch panels.

TABLE I LIST OF PREPROCESSING TECHNIQUES				
Denoise approach	std	avg	+/- background noise %	
Min-Max [0,1]	0	1	0	
Median Filter	2.67	218.1768	1.2	
Edge Preserve Blur w/ thresh	2.89	217.416	1.3	
Edge Preserve Blur	3.1	217.5898	1.4	
Adaptive Gaussian	4.51	217.5742	2.1	
Adaptive Median + Adaptive Gaussian	6.68	210.291	3.2	
Color Regularization	0.01	0.3809	3.7	
Grayscale	8.18	217.5781	3.8	
Regularization [0,2]	0.06	1.6117	3.8	
Adaptive Median	8.44	210.3311	4	
2-Normalization	0	0.0371	4.5	
Gaussian Blur	10.78	238.7197	4.5	
Scale	1.85	24.3096	7.6	
Averaging		~12	~10	
Original Image	25.58	238.7598	10.7	
Min-Max [-1,1]	0.41	1.7913	22.7	
Histogram Equalization	47.03	100.7058	46.7	
Standard Scaler	1.01	0.0426	2.37E+03	
Center	25.5	1.0705	2.38E+03	

6.2 Thresholding vs. Gradient Edge Detection

Thresholding against a smoothed background is the preliminary bulk image processing method. However, this method does not detect debris inside cavities well because the background average gray level shifts significantly relative to the faintest object we are seeking to compare it with. Ideally, such cavities and edges should help discriminate the faint objects. A single value of thresholding works reasonably well for the A view, but in the highly inclined B view, panel edges are particularly pronounced and a large part of the image. Edge detection using the Canny method [6] is employed in some cases where we know that thresholding may be breaking down, as in all edges of the panel in the B view and cavity edges in either view (Fig. 13).




Figure 13. “Finds” in the same debris-laden panel under Thresholding (top) and under Canny Edge Detection (bottom). Although edge detection methods generally result in more false positives than thresholding, we employ them whenever we know that we are in a zone of differing net absorption (edges and cavities) or whenever we need to seek and confirm a missing object found in one view but not the other. (Images from Allen and Bacon [7])

The Canny method uses two thresholds applied to pixel edges in an image instead of a binary partition on either side of a gray level. Automatic software gradient-bounding values and two standard deviations of a Gaussian filter are used.

Using the observed boundaries of the panel (oftentimes only a broken chunk) and the cavities within it, we look for the areas where standard thresholding method is suspect and run the Canny edge detection on those areas. After getting an automatic threshold value in the first Canny run, the threshold is manually adjusted higher or lower to achieve better results.

7 IMAGE ANALYSIS: FRAGMENT DATA

As discussed above, the larger goal of the current work is not just to replicate or improve the task of locating debris for extraction, but to augment or even to replace the current practice of detailed fragment analysis that follows laborious extraction. Two enhanced software avenues have been developed:

- 1) The contextual data that can be obtained by examining the orientation, deformation, breakage, and energy profile of fragments as recorded with respect to their trajectory through the foam and
- 2) Material, density, shape, and volume of the fragments themselves, determined solely from imagery.

The former of these avenues is largely complete, while the latter avenue still is under intense development, calibration, and validation. Both avenues ultimately are limited by the same signal-to-noise problems, which limit the size and density of the fragments that can be identified for extraction. However, our experience is that

enormously useful information can be obtained from this technology in both areas. We explain some of the techniques and limits here.

7.1 Contextual results

The DebrisSat project originally had hoped to characterize the velocities of the fragments from high-speed impact footage. Although the impact films were completely saturated (leading to no useful speed data), the x-ray study offers a serendipitous alternative.

From the Tate equations of terminal ballistics [8], the energy and mass of a projectile are correlated to the cavity dimensions and to the compressive yield strength Y_c of the decelerating medium. Using the orientation of the cavity and its volume from such forensic ballistics techniques, we can determine direction and initial energy, and from that (and a mass estimate of the end fragment), we can derive total velocity for most of the larger fragments. Ultimately, these various inputs may lead to a better model of the immediate post-collision state of hypervelocity impact fragments, including their speed, isotropy, and innate strength under the collision's massive heating and shear. These calculations have some limiting noise, as outlined below.

To do their job effectively, the foam panels necessarily have very low density, resulting in a small absorption of even the low-energy x-rays and negligible absorption of the high-energy rays. A nominal, 1" thick, first-surface, low-density panel (the layer and type of panel where most of the debris embedded) has a mean gray-level bit depth of 6 units of the available 256. With a limiting gray noise level of about ± 4 units, there are fundamental limits to the size of cavity features that can reliably be extracted. Such cavities hold the information about trajectory and energy of the embedded fragments at their tips and sides. To find the cavities, our software searches for large connected regions that statistically are less exposed than full-thickness foam, and seeks to interpret the smallest crevices at the edges of such regions using context to discriminate marginal features amidst the noise. This search is done in the low-noise (gray) image layer due to limited bit depth (Fig. 14).




Figure 14. Enhanced cavities in the same panel as shown previously in Fig. 2. Note the bright, densified compressed foam at the end of most cavities.

One of the properties recorded with each fragment's image is its equivalent ellipse, including that ellipse's major axis azimuth/orientation. When we gathered and sorted the images of the debris itself and color-coded the equivalent ellipse orientation as being either aligned within 45 degrees of the flight vector or 45 degrees from the normal to the flight vector, we were able to confirm that the fragments orient in a streamlined manner with strong statistical bias.

We further noted that deformations (particularly evident as bends in wire-like objects) were predominantly in the ram direction as the fragment captured, leading us to conclude that most of the observed deformations of slender fragments resulted from the (compressive) capture event, and not from the (expansive) fragmentation event (Fig. 15). Further, the existence of a streamlining effect leads us to conclude that shear is a major factor in the deceleration vs pure compression. This shear lends credence to the hypothesis that sub-fragments along a cavity wall were shed during capture and are not closely spaced fragments traveling in close association with each other.




Figure 15. Fragments sorted by size in one front-surface panel, in their proper orientation and size as imaged in the panel. The velocity vector is superimposed in blue. All items whose best-fit ellipse is within 45 degrees of ("aligned with") the flight vector are in red, with greater than 45 degrees ("broadside to the vector") painted in white. The majority of fragments exhibit strong alignment with the flight vector, and deformations appear to be correlated with the compression in the ram direction.

Conversely, we have been able to observe that many threaded fasteners and their associated nuts and washers entered the foam along different cavities, and thus may be modeled as being dissociated in the fragmentation phase, and not during capture. We are exploring refinements to our software to help us better discriminate between these competing mechanisms.

8 DETECTION LIMITS

Similar to the problem of finding voids against this faint foam background, small hydrocarbon fragments are

difficult to identify and to characterize when their total absorption is a fraction of the panel that surrounds them. With only six bins of gray bit depth in even the lightest panel, any hydrocarbon fragment less than 1/6th the areal mass of the sparse foam is virtually impossible to spot or characterize, and slightly thicker ones are difficult unless each item covers a statistically large number of pixels.

The confidence that an absorptive region of the x-ray is a real hydrocarbon fragment grows with both its size and mass/area ratio. A carbon needle several millimeters long may be identified with high confidence when observed end-on in a single pixel (absorption well beyond statistical noise level) but lead to low confidence when seen broadside, where its absorption is in family with the background noise. Unless we get two views of it, however, it is hard to develop the full mass and material properties of a found item. We do attempt to identify probable hydrocarbon fragments for extraction (the cyan ones, discussed above), but light hydrocarbon fragments reach noise limits much sooner than other materials.

Compounding the problem is the fact that decelerating fragments regularly are observed to crush the surrounding foam into locally dense areas that can easily be confused with an embedded hydrocarbon fragment. We still are developing the software to recognize such compressed foam regions by their shape, density, and context on the edge or especially, the end of a large cavity. In general, though, much of the derived small hydrocarbon debris data still is of lower quality than the information that we get on all metals and dense/larger hydrocarbons.

9 X-RAY ADVANTAGES

With those limits acknowledged, nevertheless there are remarkable things that can be discerned by the bit depths recorded in both energy levels and in the gray channel.

First, the relative balance of the high- to low-energy absorption is the key datum of the luggage scanner and is the means by which security agencies identify the materials of every object in a luggage scan. We similarly use these colors to identify the material of each fragment into one of three major groups: plastics, low density metals, and high density metals. Although it would be interesting to know whether a material was in fact iron, nickel, or tungsten, such refined information is not part of the damage calculations from the overall debris environment model, so the 3-bucket categorization is sufficient for the foreseeable future (Fig. 16).

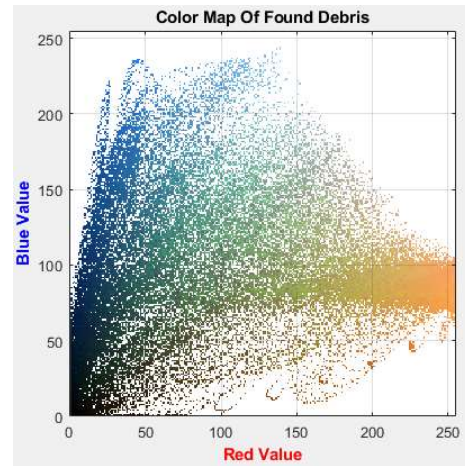


Figure 16. Color pixels within the images of debris only. All three primary colors—especially the green layer—are manipulated within the machine to result in a gray scale that corresponds to the net energy absorbed, while also enhancing colors. Red and blue layers correspond roughly but not linearly with transmittance in the high- and low-energy beams, respectively. Note that the steel band (blue) low-energy transmittance exceeds that of the pure foam (orange) in which it is embedded. Thus, the foam appears to be more absorptive than steel. This is an artifact of the rendering function that is designed to indicate both material and depth to the human eye. The color family along any of these fan “rays” of color (each ray is a different debris item of constant material) is visually in the same family and corresponds to the material density. The grayscale value corresponds to the cube root of the thickness. Dense metal is blue, light metal is green, and plastics render in orange, in general. The horizontal orange region is dominated by pixels on the splined boundaries of other (metallic) debris.

Second, since the inception of x-ray technology, the overall bit depth (particularly in the gray scale) has been useful in estimating the physical depth of an object. Mathematically, the x-ray absorption is linear with cube root of areal density. Thus, it is possible to estimate the total depth of an object from a single view, and in fact, we have calibrated many materials to give accurate depth estimates based solely on exposure and material type. An example of the theoretical exposure profile of straight pin vs. its measured profile in our system is shown in Fig. 17, showing exceptional agreement. It is crucial to note, however, that the actual shape is still slightly ambiguous: a cylinder is not distinguishable from a half-cylinder of twice its density using this technique. We call this the “Cylinder vs. Quonset Hut Ambiguity.” For this reason, we need a minimum of two views to create a reasonable understanding of the true shape.

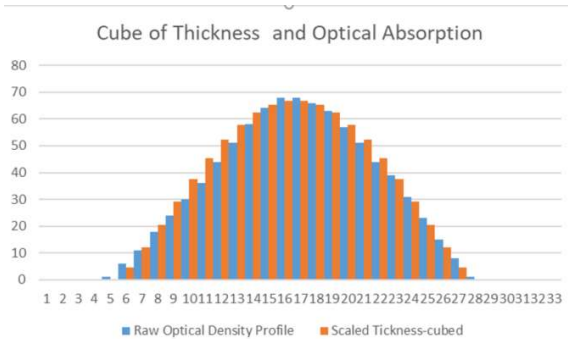


Figure 17. X-ray exposure of a cylindrical steel pin. The raw optical density of the x-ray is in blue. The scaled expected-response of the cylinder is in orange as the cube of the material depth along a ray, and exactly matches the splined curve within bit quantization. Through such calibrations (per material) it is possible to compute material thickness along the ray trace, thus adding additional insight to the fragment's 3D shape beyond what is learned from simplistic shape carving.

With two views, we achieve a rudimentary shape carve of the silhouettes. This is not as good as laboratory-based shape carving for several reasons, including dramatically poorer resolution, non-orthogonality of the views, and randomized alignment of the views relative to the object's principal axes. However, with moderately accurate depth estimates from the pixel-by-pixel exposure, we can do mathematics similar to that used in medical computer-aided tomography ("CAT" scans) to shuffle the depth bars around to reach a self-consistent probabilistic estimate of a true 3D shape. With the rudimentary shape carve defining the limiting envelope, we mathematically dither the locations of the voxel columns along the rays until the object above the A view is most consistent with the object as estimated by the B view. The "cylinder vs. Quonset hut" ambiguity can to some extent be resolved in this manner, although the exposure data is noisy, and the "answer" results in a set of probabilities that a particular 3D voxel is filled or empty (refer to Fig. 4).

10 Results

The methods described here are capturing boundaries, material, and depth information in two views of embedded debris 2 mm or greater with 0.07 mm resolution. Precision and Recall of the finding method are approximately 90% each with "ground truth" provided by thorough disintegration of the panel and a check of found objects vs. the software-generated map of predicted finds. The false positives are dominated by compressed foam confused as a hydrocarbon fragment, and false negatives are almost universally small, thin hydrocarbon fragments. Metal fragments are very reliably identified, except for short wires < 0.7 mm diameter. (Fig 18)

We are currently tuning the precision of the automated software that determines the actual volume, mass, and true shape of each fragment. We plan to publish these findings later this year, but the results shown in Fig. 4 are encouraging, with dimensions observed to be within small fractions of a millimeter of lab values.




Figure 18. The virtual 1/2" grid and the location map for a small chunk of foam. Cyan objects are those found only in the A view. These (if real) are located along a ray angle illustrated at the far right of the grid. All other colors represent real debris identified in two views, in the colors associated with their material density. (Blue/black is steel, green is aluminum, and orange is plastic). One expects to find these directly under their image in this parallax-free rendering. The red grid axes align on one of the registration pins (blue object at origin). A second registration pin is visible far right, exactly on a grid intersection as expected. The pins are cyan because, being longer than the panel is deep, they cannot be exactly correlated in the B view to lie within the panel, so are rejected as unconfirmed "finds" in the A view.

The orientation of the fragments relative to the flight vector indicates that they typically align in a streamlined way during the deceleration process. It is evident that deformation of wires is common at the ram end during such deceleration. It remains TBD and the subject of current work to confirm a hypothesis that other shapes suffer similar deformation. It is clear, however, that at least some of the fragments (both plate and wire shapes) break during the capture process, and it is a subject of current study to characterize what such breakage means to the statistical models that we generate from the extracted fragments.

11 Future Work

We are comfortable that we have optimized the image quality and registration, such that the remaining developments of this method will be exclusively in post-processing software. A newer-generation scanner would likely improve the signal-to-noise and would greatly

improve the 3D definition of each fragment shape and density. However, we have concluded that the current process already meets the resolution and accuracy requirements sufficient to conclude our statistical study of fragment shape, size, and material, and to compare this data against simplified geometric families of shapes that ultimately will define our debris model.

Our long-baselined process continues to analyze all extracted fragments in the laboratory, while the x-ray data process is perfected in parallel. In selected cases we can identify the precise coordinates of any analyzed object with its x-ray image pair, providing best possible calibration of the new process to the old. (Remember that the existing archive did not record location well enough to correlate back to the x-ray image). We have elected to calibrate the system by analyzing fragments in the x-ray first, and then reserving certain challenging fragments for head-to-head comparison, during this iterative improvement process. The new scanner-based extraction process is baselined, and all future software effort is centered on calibrating the fragment's derived data and in expanding our context characterizations for energy, deformation, and shedding. This latter work is aimed at correcting the statistical data on shape and size to account for breakage and shedding that occurred in the capture process, while the former work may lead to orders-of-magnitude speed improvements in DebrisLV and future fragmentation studies.

Thus, the team is in the process of recording and archiving all necessary images of the hundreds of remaining panels of both the DebrisSat and DebrisLV impact tests, using the described imaging techniques. While the DebrisSat project will have a precise, laboratory-generated data set to support the development of the new Orbital Debris Engineering Model (ORDEM) version 4.0, we believe that the x-ray techniques we have defined herein will be essential in defining the immediate post-impact debris field prior to its deceleration. Even more crucially, given their multiple order-of-magnitude time improvement in generating fragment data, our x-ray techniques may be the sole source of data that can be obtained on the equivalent DebrisLV impact test in time to support ORDEM 4.0.

The project is just finishing the imaging of every panel of both experiments. Once complete, the hardware can be retired. Since NASA owns the scanner and the shoe, it is hoped that we may find other ballistic tests within the U.S. Government or academia to which this hardware and these techniques can be applied.

12 Conclusions

A well-registered, two-energy x-ray luggage scanner has been demonstrated to provide useful data in all parameters of interest in fragmentation events, for fragments ≥ 2 mm in longest dimension. The intrinsic

imaging limits of resolution and of saturation (high and low) for an industrial-grade scanner have proven to be acceptable for scientific work, although improvements can be obtained in several parameters with higher resolution and higher dynamic range scientific equipment. Because such equipment can readily be exchanged at no cost between government agencies, such equipment has proven to be more easily accessible and cheaper than many more exotic options.

Newer luggage scanning equipment may prove to be substantially more powerful than this first implementation, since the state-of-the-art has evolved to provide more of a medical CAT-scan capability than the utilized *circa* 2010 surplus scanner provided (9). The ultimate limits of this technology are found where the fragments are of comparable size to three pixels, and/or are of a chemical composition and mass/area ratio that is less than or equal to that of the medium that embeds them. In such cases, extraction and benchtop measurement may be required to complete the statistical assessment of the fragment population.

Although the details of fragment shape, size, and mass are valuable and comparable in accuracy and precision to those obtained by competing methods for all denser, metallic fragments, the most significant advance provided by x-ray analysis is the contextual information associated with the cavitation, breakage, and deformation along the impact flight path. Such context significantly aids in understanding the original state of the fragments before capture. Combined with the two orders-of-magnitude speed increase in data gathering and analysis over prior manual methods, x-ray-based fragment analysis is a formidable new tool for characterizing fragmentation events.

13 References

1. Hogg, D.M., Cunningham, T.M., and Isbell, W.M. (1993). Final Report on the SOCIT Series of Hypervelocity Impact Tests, Report No. WL-TR-93-7025, Wright Laboratory, Armament Directorate, Wright-Patterson Air Force Base, Ohio.
2. McKnight, D.S., Johnson, N.L., Fudge, M.L., and Maclay, T.D. (1995). Satellite Orbital Debris Characterization Impact Test (SOCIT) Series Data Collection Report, Kaman Sciences Corporation, Alexandria, Virginia.
3. Rivero, M., Kleespies, J., Patankar, K., *et. al.* (2015). "Characterization of Debris from the DebrisSat Hypervelocity Test." In *Proceedings 66th International Astronautical Congress*, Jerusalem, Israel, IAC-15- A6.29x30343.
4. Shiotani, B., Rivero, M., Carrasquilla, M., Allen, S.(2017). "Characterizing DebrisSat Fragments: So Many Fragments, So Much Data, and So Little

- Time.” In *Proceedings 68th International Astronautical Congress*, Adelaide, Australia, IAC-17.A6.3.6x41656.
5. Allen, S. (2018). “Characterizing DebrisSat Fragments – Preliminary Results.” In *Proceedings 69th International Astronautical Congress*, Bremen, Germany, IAC-18.A6.3.4x45117.
 6. Canny, J. (1986). “A Computational Approach to Edge Detection.” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **8**(6), 679–698.
 7. Allen, A.R. and Bacon, J.B. (2019). “Macro-Scale Findings of the DebrisSat Debris Field Obtained from X-Rays of the Catch Panels,” In *Proceedings 1st International Orbital Debris Conference*, Sugarland, TX, USA.
 8. Tate, A. (1969). “Further results in the theory of long rod penetration.” *Journal of the Mechanics and Physics of Solids* **17**(3), 141-150.
 9. Transportation Security Agency (2021) “TSA Air Cargo Screening Technology List (ACSTL)” Version 11.5 February 10, 2021. url: https://www.tsa.gov/sites/default/files/non-ssi_acstl.pdf