

# Computer Tomography with Laser-generated X-rays at Gemini

Contact: [stephen.dann@stfc.ac.uk](mailto:stephen.dann@stfc.ac.uk)

**S. J. D. Dann, C. Thornton, C. D. Armstrong, C. D. Baird, C. M. Brenner, D. R. Symes**  
*Central Laser Facility,  
STFC Rutherford Appleton Laboratory,  
Didcot, OX11 0QX, UK.*

**M. J. V. Streeter, J.-N. Gruse, J. C. Wood, N. C. Lopes, S. P. D. Mangles, Z. Najmudin**  
*John Adams Institute for Accelerator Science,  
Imperial College London,  
London, SW7 2AZ, UK*

**C. D. Murphy, C. I. D. Underwood**  
*York Plasma Institute, Department of Physics,  
University of York, Heslington, YO10 5DQ, UK*

**L. R. Pickard**  
*National Composites Centre,  
Bristol and Bath Science Park, Feynman Way Central,  
Bristol, BS16 7FS, UK*

**K. D. Potter**  
*Advanced Composites Collaboration for Science and  
Innovation (ACCIS), University of Bristol,  
Bristol, BS8 1TR, UK*

## Abstract

We describe recent CT imaging work at Gemini. We focus on aspects of the scanning and reconstruction processes that are needed to deal with the poor repeatability of current laser-driven X-ray sources. In particular, we describe golden ratio scanning, a useful technique employed with sources of degrading quality, and a custom beam profile correction step, applied to correct for the pulse-to-pulse changes in the X-ray beam profile and direction.

## 1 Introduction

Computer tomography (CT) is a well-established technique for non-invasively imaging the interior structure of an object. While its best known applications are in medicine, CT is widely used in science and industry with scanners ranging from commercial bench-top devices to beamlines at Diamond and ISIS. The most common forms of CT use the absorption of X-rays within a sample to generate contrast. With spatially coherent X-ray sources phase differences can also contribute to image contrast, but we do not consider the effect of phase in this report.

X-ray radiography of industrial samples has already been demonstrated on the Gemini laser [1], as has CT of biological samples [2]. These use the so-called betatron X-rays generated in a laser wakefield accelerator. We applied the same techniques to a variety of industrial samples and refined the process further, to acquire and analyse the data more efficiently. In particular, we automated the process of scanning, reducing the chance of human error, and we applied the same analysis software used at Diamond Light Source.




Figure 1: Experimental schematic of the laser-plasma x-ray source. A -  $f/40$  parabola; B - Gas cell; C - Polyimide tape; D - Diverter magnet; E - Lanex screen; F - Target chamber wall; G - X-ray window; H - Rotation stage; I - Sample; J - Detector.

## 2 Experimental setup

The experiment was carried out using the Gemini South beam in an  $f/40$  configuration to deliver  $\sim 10$  J with 45 fs duration onto the interaction point (IP). The focal spot was measured to be  $40 \mu\text{m}$  giving a peak intensity  $\approx 7 \times 10^{18} \text{ W cm}^{-2}$ . Placed at the IP was a fixed-length gas cell which contained a mixed gas of helium with a 2% nitrogen doping maintained at a fixed pressure of 80 mbar. After the IP was a laser beam block, this was a 25  $\mu\text{m}$  thick polyimide tape that was refreshed on each laser shot using a tape drive. After the laser was separated from the electron and X-ray beam, the electrons were diverted by a 1 T magnet to an electron spectrometer and the X-rays continued to propagate to the end of the target chamber where they passed through a gold-coated polyimide sheet, used as an additional laser beam block in case the primary laser beam block failed. Finally, the X-rays exited the target chamber through a 250  $\mu\text{m}$  thick




Figure 2: The evolution of X-ray yield through a typical scan. The dashed line indicates the chosen threshold for usable data. Over the whole scan 54% of shots had a signal level above this threshold.

polyimide window. The X-rays pass through a sample mounted on a Standa rotation stage before being detected on an Andor iKon-L indirect X-ray detection camera. A diagram of the set-up is shown in figure 1.

Each X-ray image recorded in this way is a *projection*, containing information on the integrated density along the path of the X-ray beam. Projections are recorded as the sample is rotated through angles from  $0^\circ$  to  $180^\circ$ , and then they are converted to a 3-d volume using computer software. The Gemini laser was run in an automated mode, with the stage being rotated automatically after each shot. Data was also recorded automatically, making the whole process smooth and efficient.

### 3 Source degradation and golden ratio scanning

For now, laser-driven X-ray sources are less reliable and less consistent than their traditional counterparts. One consequence of this is that the X-ray yield tends to decrease throughout the scan, as shown in figure 2. If the scan were performed linearly from  $0^\circ$  to  $180^\circ$ , the later angles would have fewer X-rays and hence worse signal-to-noise ratio. Furthermore, if the X-ray yield fell enough that the scan had to be terminated early, there would be a range of angles that were entirely missing, preventing the reconstruction process from proceeding properly.

In fact, rather than setting out to perform a scan with a given angular spacing, we would prefer a process of constant refinement, so that we could terminate the scan at any point and still have a reasonably complete dataset. This should have the additional benefit of spreading the high and low quality data throughout the complete range of angles. To do this, we take a number  $a$ , and at each




Figure 3: The spacing of samples in a golden ratio scan of 65 samples. The samples cover the range  $0^\circ$  to  $180^\circ$ . Their opposites (obtained by adding  $180^\circ$ ) are also shown to emphasise the complete coverage offered.

iteration  $i$  we rotate the sample to:

$$\theta_i = 180^\circ \text{fract}(ai) \quad (1)$$

$$\text{where } \text{fract } x = x - \lfloor x \rfloor \quad (2)$$

That is, we multiply  $a$  by the iteration count, take the fractional part and multiply it by the size of our desired interval ( $180^\circ$ ). Although any irrational number might work, the usual choice for  $a$  is the golden ratio [3],  $\varphi = (\sqrt{5} + 1)/2$ . Figure 3 shows that this covers the half-circle roughly uniformly, without obvious gaps. In fact this is true of any continuous sequence of samples, so that both high-quality and low-quality samples are distributed uniformly, and whenever you choose to stop the scan the dataset will be complete. Note however that the reconstruction process must be able to deal with non-uniform sample spacing.

### 4 Reconstruction process

Once the projections have been obtained, they must be *reconstructed* to form a 3-d volume. During the experiment, we performed this using Savu [4, 5], a comprehensive software suite developed at Diamond, which provides a uniform interface not only to the reconstruction process itself, but also to input, output, and preprocessing steps. We used Savu to read partly-preprocessed data from NeXuS files [6], then perform important steps such as filtering, alignment, and downsampling, before reconstructing the 3-d volume and writing the results to an HDF5 file. This was then visualised with the help of ParaView [7]; this pipeline allowed the results to be viewed in the control room within hours of the data being acquired.

Several important steps were performed with custom scripts, since there was no equivalent functionality in




Figure 4: An example of the beam profile correction process: (a) The original image; (b) the image after removing the sample and alignment wires; (c) the reconstructed X-ray beam profile; (d) the transmission map, obtained by dividing (a) by (c). The second row shows how the regularisation parameter  $\alpha$  affects the fitting process in this case: (e)  $\alpha = 10^{-3}$  is too large failing to fit to the beam profile correctly; (f)  $\alpha = 10^{-4}$  is better, and is used for (c-d); (g)  $\alpha = 10^{-5}$  is beginning to overfit slightly, but still giving reasonable results; (h)  $\alpha = 10^{-6}$  is too small, and the beam profile is badly overfitting to the data. Removing the regularisation entirely ( $\alpha = 0$ ; not shown here) has a similar effect.

Savu. Briefly, the processing pipeline consisted of the following steps:

1. Review the data, particularly the signal levels; eliminate all images with poor X-ray signal (using an ad-hoc threshold; 30 counts per pixel in this case) and identify suitable images to use as dark fields (no X-rays) and flat field (X-rays, but no sample);
2. Combine all remaining images with their corresponding sample angles to form a single NeXuS file containing raw data;
3. Perform beam profile correction (see below) to produce a file containing transmission percentages instead of raw X-ray counts;
4. Feed the data into Savu, which typically performed the following steps:
  - (a) Apply a median filter
  - (b) Translate each projection to correct for offset
  - (c) Downsample data to reduce computation time
  - (d) Compute the centre of rotation
  - (e) Reconstruct the volume using filtered back-projection (FBP)
5. Use a script to create an XDMF file, allowing the HDF5 volume data to be visualised in ParaView.

The main custom step here, which we describe in more detail, is the beam profile correction. The raw data is a value of X-ray signal per pixel, whereas the tomographic reconstruction needs a transmission fraction. The standard method to do this is to divide the image by a “flat field”, acquired using the same setup but with the sample removed. This assumes that the X-ray beam has the same shape and intensity on every shot, which is not typically the case with laser-driven sources. Instead, we attempt to infer the beam profile, and use that instead of a separately-measured flat field. This requires a border region in the image that the shadow of the sample will never touch. We then fit a suitable function to the data in this border, before extrapolating that fit across the central region to form the beam profile. While the natural choice of a function to fit might be a Gaussian, we instead use a polynomial in two variables. The fitting process is then linear, and we can apply regularisation.

The beam profile could easily be overfit to data in the border, causing problems when it is extrapolated across the centre. We solve this with *regularisation*, which penalises fits that have larger coefficients. This is based on the principle that overfitting usually results in larger coefficients. We take the  $L^1$  norm, so that the fitting problem becomes:

$$\text{Minimise: } \sum_i \left( y_i - \sum_j x_{ij} w_j \right)^2 + \alpha \sum_i |w_i| \quad (3)$$

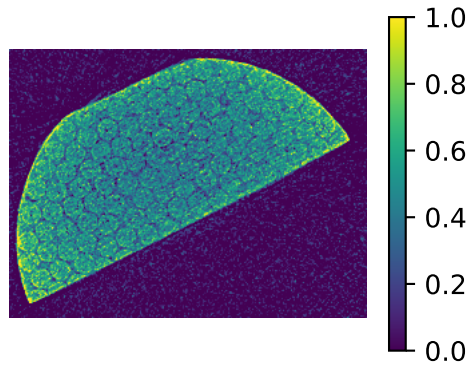


Figure 5: An average over 200 slices, reconstructed using gridrec. The colour scale indicates reconstructed density.

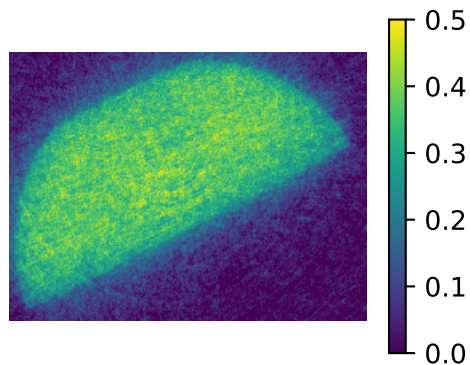


Figure 6: A single slice, reconstructed using SIRT. The colour scale indicates reconstructed density.

where  $\alpha$  determines the strength of the regularisation.

An example of this beam profile correction is shown in figure 4. This shows that the calculated beam profile matches the data closely, giving a clean transmission measurement. It also shows how important the selection of the regularisation coefficient  $\alpha$  is to getting reasonable results.

## 5 Reconstruction results

Data from one particular sample is shown in figures 5 and 6. This is a composite cylinder containing a kink band failure that has been used by the National Composites Centre to assess the effect of impact and compression, see also Ref. [1]. The entire dataset was reconstructed using the gridrec algorithm [8] available at the TomoPy github [9]. Although the shape of the sample is visible, the signal-to-noise ratio is very poor. However, the sample is very nearly uniform in the vertical direction

(the rotation axis), and averaging 200 slices results in figure 5. This shows the object’s internal structure, with cylinders packed together in a hexagonal pattern. The space between the cylinders appears to be filled with a slightly less dense material. Taking a single slice (sinogram) and reconstructing it with the SIRT (Simultaneous Iterative Reconstruction Technique) algorithm, set to 10 iterations, results in figure 6. Although the internal structure is not visible, the overall shape of the sample is clear.

## 6 Conclusion

There is great potential to use X-rays generated with laser-plasma accelerators for industrial tomographic scanning. When fully optimised, these sources should be able to provide rapid inspection with sub-micron resolution — a capability that is not possible with conventional CT machines. Translating this technology from proof-of-principle experiments to industrial application requires improvements in source characterisation and stability. This is a core objective of the EPAC facility under construction at the CLF. Because the properties of these X-ray beams are different to those produced in synchrotrons or benchtop X-ray scanners, new image processing techniques need to be developed. Here, we have started that process showing a background subtraction method specifically written for the Gemini betatron X-ray beams. In the longer term, these steps can be integrated into a CT pipeline using software that is standardised across the RAL user facilities.

## References

- [1] J.-N. Gruse *et al*, Nucl. Instrum. Meth. A **983** 164369 (2020) DOI: 10.1016/j.nima.2020.164369
- [2] J. M. Cole *et al*, PNAS **115** 6335–6340 (2018) DOI: 10.1073/pnas.1802314115
- [3] T. Köhler, IEEE Nuclear Science Symp. Conf. Record, **6** 3961–3965 (2004) DOI: 10.1109/NSSMIC.2004.1466745
- [4] R. C. Atwood *et al*, Phil. Trans. R. Soc. A **373** 20140398 (2015) DOI: 10.1098/rsta.2014.0398
- [5] <https://github.com/DiamondLightSource/Savu/>
- [6] <https://www.nexusformat.org/>
- [7] <https://paraview.org/>
- [8] B. A. Dowd *et al*, Proc. SPIE **3772** 224–236 (1999) DOI: 10.1117/12.363725
- [9] <https://github.com/tomopy/tomopy/>