

A new approach to synchrotron energy-dispersive X-ray diffraction computed tomography

Olivier Lazzari,^{a,b} Christopher K. Egan,^{c*} Simon D. M. Jacques,^{a,b,c} Taha Sochi,^{a,b} Marco Di Michiel,^d Robert J. Cernik^c and Paul Barnes^{a,b}

^aSchool of Crystallography, Birkbeck College, Malet Street, London WC1E 7HX, UK, ^bDepartment of Chemistry, University College London, 20 Gordon Street, London WC1H 0AJ, UK, ^cSchool of Materials, The University of Manchester, Manchester M13 9PL, UK, and ^dESRF, 6 rue Jules Horowitz, 38000 Grenoble, France. E-mail: christopher.egan@manchester.ac.uk

A new data collection strategy for performing synchrotron energy-dispersive X-ray diffraction computed tomography has been devised. This method is analogous to angle-dispersive X-ray diffraction whose diffraction signal originates from a line formed by intersection of the incident X-ray beam and the sample. Energy resolution is preserved by using a collimator which defines a small sampling voxel. This voxel is translated in a series of parallel straight lines covering the whole sample and the operation is repeated at different rotation angles, thus generating one diffraction pattern per translation and rotation step. The method has been tested by imaging a specially designed phantom object, devised to be a demanding validator for X-ray diffraction imaging. The relative strengths and weaknesses of the method have been analysed with respect to the classic angle-dispersive technique. The reconstruction accuracy of the method is good, although an absorption correction is required for lower energy diffraction because of the large path lengths involved. The spatial resolution is only limited to the width of the scanning beam owing to the novel collection strategy. The current temporal resolution is poor, with a scan taking several hours. The method is best suited to studying large objects (*e.g.* for engineering and materials science applications) because it does not suffer from diffraction peak broadening effects irrespective of the sample size, in contrast to the angle-dispersive case.

Keywords: energy-dispersive diffraction; imaging; computed tomography.

1. Introduction

Tomography refers to the cross-sectional imaging of the interior of bulk objects using transmitted or scattered radiation (Kak & Slaney, 1988). The most common form of tomography is X-ray transmission computed tomography whereby an intense hard X-ray beam is used to illuminate an object at many different angles of rotation producing a large number of two-dimensional radiographs. These radiographs are then computed to produce a three-dimensional image of the bulk interior of the object (Brooks & Di Chiro, 1976). This approach to tomographic imaging has revolutionized radiological diagnostics for medical applications and has also been extensively used for materials science research (Stock, 2008; Brenner & Hall, 2007). Another approach to computed tomography imaging is to use the scattered radiation to form an image of the object. This offers some advantages over standard absorption tomography since the scattered radiation can hold extra information about the sample being analysed. For example, in X-ray fluorescence tomography the energy of

the emitted radiation is characteristic of the element or compound, thus element-specific imaging can be performed (de Jonge & Vogt, 2010). There are several disadvantages associated with such techniques, however, most notably (i) the generally longer acquisition times owing to the weak intensity of the scattered radiation and (ii) the difficulties with constructing an image of the object with the data available. The former can be overcome by using an intense and well defined incident beam, like those obtained at a synchrotron. This paper deals specifically with X-ray diffraction (XRD) scatter tomography as has been previously demonstrated (Harding *et al.*, 1987; Grant *et al.*, 1995; Kleuker, 1997). In this method the scattered radiation originates from diffracted X-rays from crystalline regions within the object. By measuring the scattering angle of the diffracted X-rays the inter-atomic spacing, d , of the crystalline object can be calculated using Bragg's law,

$$\lambda = 2d \sin \theta, \quad (1)$$

where θ is the Bragg angle for diffraction and λ is the wavelength of the diffracted light. Each crystalline material within the sample gives rise to a unique diffraction signature, leaving the fingerprint of that particular material. This method, when combined with tomography, therefore enables material-specific imaging. Additional information about the state of the sample is also available within the diffraction pattern which can also be exploited. For example, the relative positions of the peaks in the diffraction pattern can be mapped to give a measure of strain or crystallite size within the object.

There are two standard approaches to measuring an X-ray diffraction pattern: the first and most well known method uses a monochromatic X-ray beam (*i.e.* fixed λ) and the information about d is dispersed in angle (the so-called angle-dispersive case). The diffraction pattern can be recorded by scanning a point detector through an angular range, or by using one-dimensional or two-dimensional detectors that span this angular range (Harding *et al.*, 1987). The second method uses a polychromatic X-ray beam, with a fixed angle of diffraction (by virtue of a collimator assembly), and the energy of diffracted photons (the so-called energy-dispersive case) is measured using an energy-dispersive point detector (0D) (Hall *et al.*, 1998; Lazzari *et al.*, 2009). Both methods can be exploited to perform tomography, and are termed tomographic angle-dispersive diffraction imaging (TADDI) and tomographic energy-dispersive diffraction imaging (TEDDI), respectively. In this paper we describe a new approach to performing energy-dispersive computed tomography that is akin to the traditional angle-dispersive case, whilst still preserving energy resolution. The method is based upon a translation of the scattering voxel through the sample object whilst collecting a single diffraction pattern, akin to traditional angle-dispersive XRD tomography. This new technique overcomes many previous limitations of energy-dispersive diffraction, in particular spatial resolution.

2. Method

X-ray diffraction computed tomography requires that the detector be positioned out of the direct transmitted beam, thus collecting data at some specified angle which defines the diffraction condition (2θ). In the TADDI set-up this is typically performed using an area detector positioned behind the sample using a central beam stop. The observed diffraction pattern can then be azimuthally integrated yielding a linear diffraction pattern with a high signal-to-noise ratio. In order to form an image of the object, a pencil-beam is used which limits the diffraction to specified regions within the sample. The sample is then translated and rotated, and a diffraction pattern is recorded at each position and angle. In the TEDDI configuration the data collection strategy is slightly different because of the requirement of a collimator. The collimator not only fixes the Bragg angle for diffraction but also defines a small region (a voxel) within the sample where the diffracted radiation has come from (see Fig. 1*a*). In principle, the sample could now be translated in three dimensions, scanning the voxel throughout the object, and a three-dimensional tomo-

gram could be simply pieced together. Indeed this has been performed previously (see, for example, Lazzari *et al.*, 2009); however, there are issues with deconvolving the data owing to the highly squashed shape of the voxel. To overcome this, we chose to use a computed tomography approach whereby the voxel is continuously scanned through the sample in a direction parallel to the incident beam whilst collecting the energy-dispersed data. This results in one diffraction pattern per sample translation (y) and rotation (ω) creating a data set of size $N_y \times N_\omega \times k$, where k is the number of energy channels.

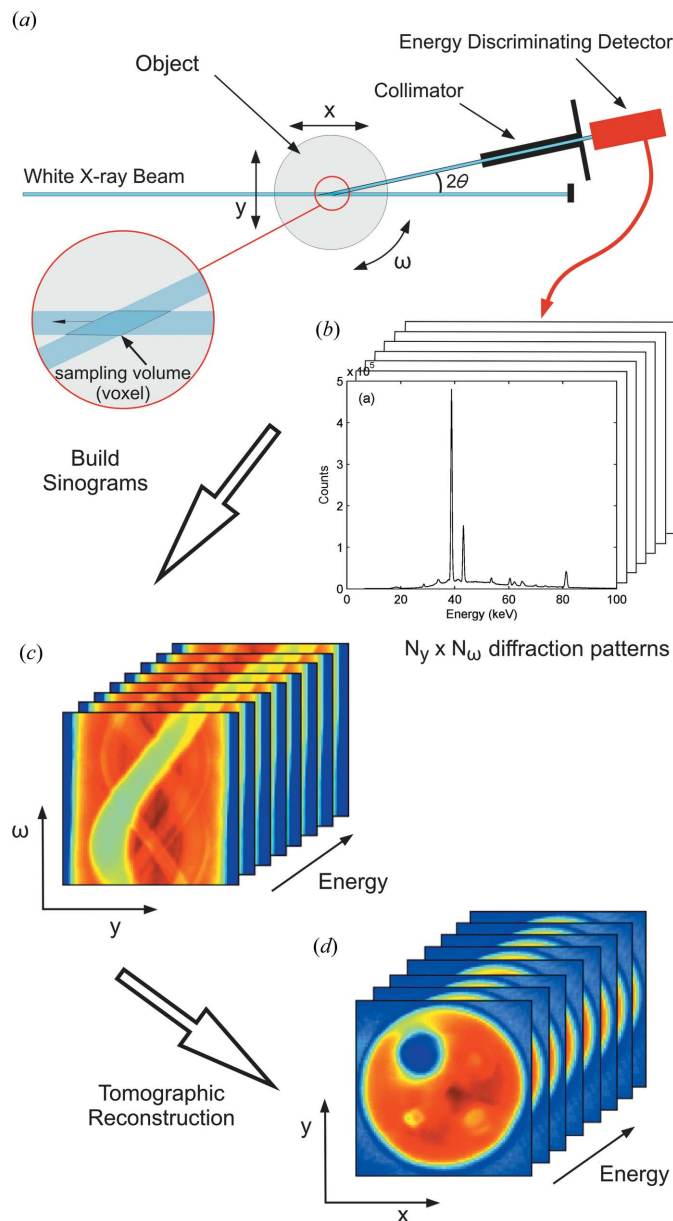


Figure 1
(*a, b*) Schematic of the basic experimental set-up for energy-dispersive X-ray diffraction computed tomography. The voxel is continuously scanned through the sample in the x direction whilst collecting data. This constitutes the collection of one diffraction pattern. This is repeated for a number of translations (N_y) and rotations (N_ω) as shown in (*b*) from which sinograms are built (*c*). Tomograms (*d*) are reconstructed from this data set.

Both the TEDDI and TADDI methods produce a three-dimensional matrix data set. Each $N_y \times N_\omega$ ‘image’ is known as a sinogram, so we effectively have k sinograms corresponding to each energy or angle channel. As a simple starting point we can view each sinogram throughout the energy or angle range and observe how individual features within the object appear and disappear as we move through each k channel. From here we can select a region of interest in the diffraction pattern that specifies a particular diffraction peak and compute the corresponding tomogram for that region. We can repeat this for every material within our object and stitch these together to produce a reconstruction of the sample interior. Alternatively we can compute the tomogram for every channel (*i.e.* for every k value) producing a two-dimensional slice for each channel, in a method previously reported (Bleuet *et al.*, 2008). Here, tomograms are reconstructed using filtered back-projection, a standard for computed tomography (Kak & Slaney, 1988).

3. Experimental details

We have devised a phantom object to be a demanding validator for the TEDDI method. The object has been designed to contain regions of varying X-ray absorption (relatively low absorption materials with wax and bone ash to higher absorption elements with aluminium and iron) and different degrees of crystallinity (amorphous glass, semi-crystalline PEEK polymer, highly crystalline bone ash and iron oxide). The phantom object has an outer diameter of 36 mm and is schematically shown in Fig. 2.

TEDDI measurements were performed at station 16.4 of the Daresbury SRS synchrotron source before its closure in 2008 (Clark *et al.*, 1996). This station generated a high-flux white X-ray beam, from a superconducting three-pole wiggler magnet (peak field of 6 T), with a critical energy of 33 keV. The X-ray beam had a pencil-like geometry of diameter 0.5 mm. Energy-dispersive data were recorded using a liquid-nitrogen-cooled germanium detector using planar lapped molybdenum collimator plates. The collimator plates were 30 cm in length and were separated by 0.1 mm giving an aspect ratio of 3000:1. The sample–collimator distance was 100 mm. The collimator and detector were set to define an angle of diffraction of 4.5° which approximately covers d -spacings in




Figure 2 Diagram of the phantom object containing cylindrical features made from six materials with varying crystalline and X-ray absorption properties.

the range $0.8\text{--}7.9 \text{ \AA}$ for the given energy range of the source and detector. The detector count time was 10 s per measurement.

Data collection was performed by scanning the beam horizontally through the sample in a total of 80 steps with a step size of 0.5 mm. This was coupled with rotation around 180° in a total of 80 steps. This therefore yields 80×80 diffraction patterns. Crucially, owing to the long count times associated with the technique, the data acquisition took approximately 18 h to collect.

4. Results and discussion

Fig. 3 shows the summed total energy-dispersed diffraction (EDD) pattern obtained from the phantom object. The EDD pattern covers an energy range of between 10 keV and 100 keV and we observe 15 individual peaks within this range. These peaks correspond to well defined diffraction conditions from the highly crystalline (iron oxide, aluminium, *etc.*) and semi-crystalline (wax and PEEK) materials. Typically the peak full width at half-maximum (FWHM) is less than 1 keV. We also observe a broad spectrum covering an approximate energy range of between 25 and 70 keV. This broad spectrum

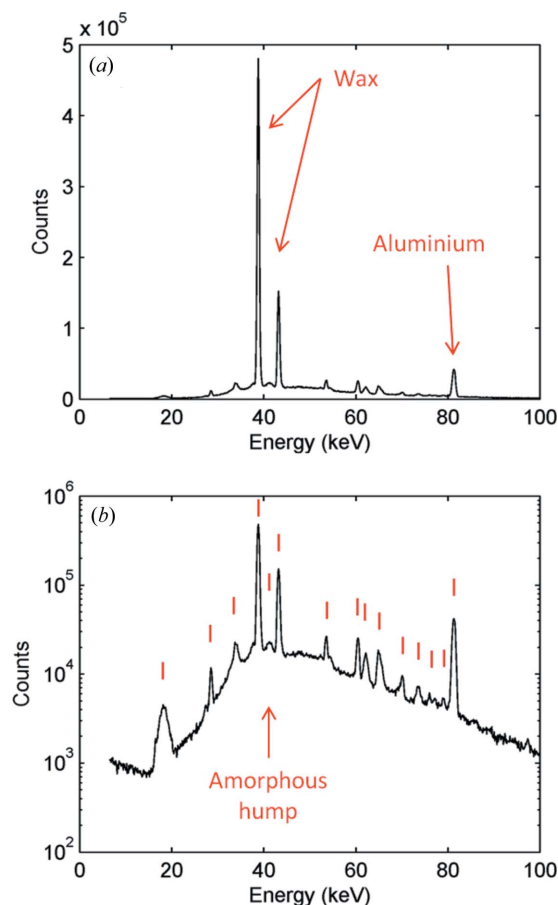


Figure 3 Cumulative sum of all energy-dispersive diffraction patterns obtained from the phantom object. (a) Linear scale, (b) logarithmic scale. Tick marks and labels identify the positions of major diffraction contributions.

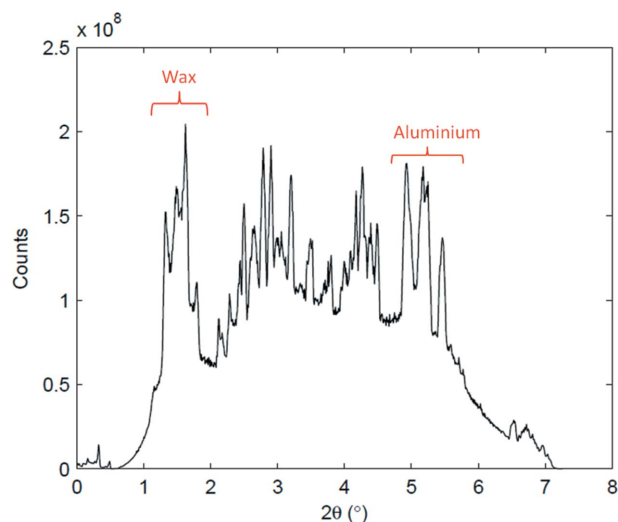


Figure 4 Angle-dispersive diffraction pattern collected from the same phantom object, showing significant peak broadening. Phase identification is almost impossible without prior knowledge of the sample composition. Data collected using an incident monochromatic X-ray beam of energy 90 keV.

corresponds to diffraction from the amorphous and semi-crystalline materials within the sample (wax, glass and PEEK).

Because the voxel acts like a point-source scatterer with a strictly defined angle 2θ , diffraction peaks are found to be well defined. If the angular acceptance of the collimator is increased such that scattering occurs from a much larger section of the sample, we would observe broadening of the diffraction peaks, meaning physical quantities like phases, particle size and strain could not be identified. By analogy, this is often observed in angle-dispersive diffraction from large objects, *i.e.* those with a large angular acceptance. As an example, Fig. 4 shows an angle-dispersive diffraction pattern collected from the same phantom object. We can see that there is significant peak broadening with many of the diffraction peaks overlapping. This means that it is almost impossible to determine the phases within the sample without *a priori* knowledge of the sample composition and indicates TADDI is completely unsuitable for studying large objects. Having said that, tomographic reconstruction can still be performed as shown in the supplementary information.¹ As a rule of thumb, we suggest a working range for TADDI as

$$\text{sample diameter} \times \tan(2\theta) < \text{effective detector pixel size.}$$

TEDDI, on the other hand, is well suited for studying large objects because of the small angular acceptance meaning it does not suffer from broadening effects. Unfortunately, it is less well suited to imaging smaller features within objects, *e.g.* thin coatings or particulates. Owing to the finite size of the diffraction lozenge, diffraction signals may be weak or drowned out when the voxel is stepwise moved through the sample. The length of the lozenge is described by

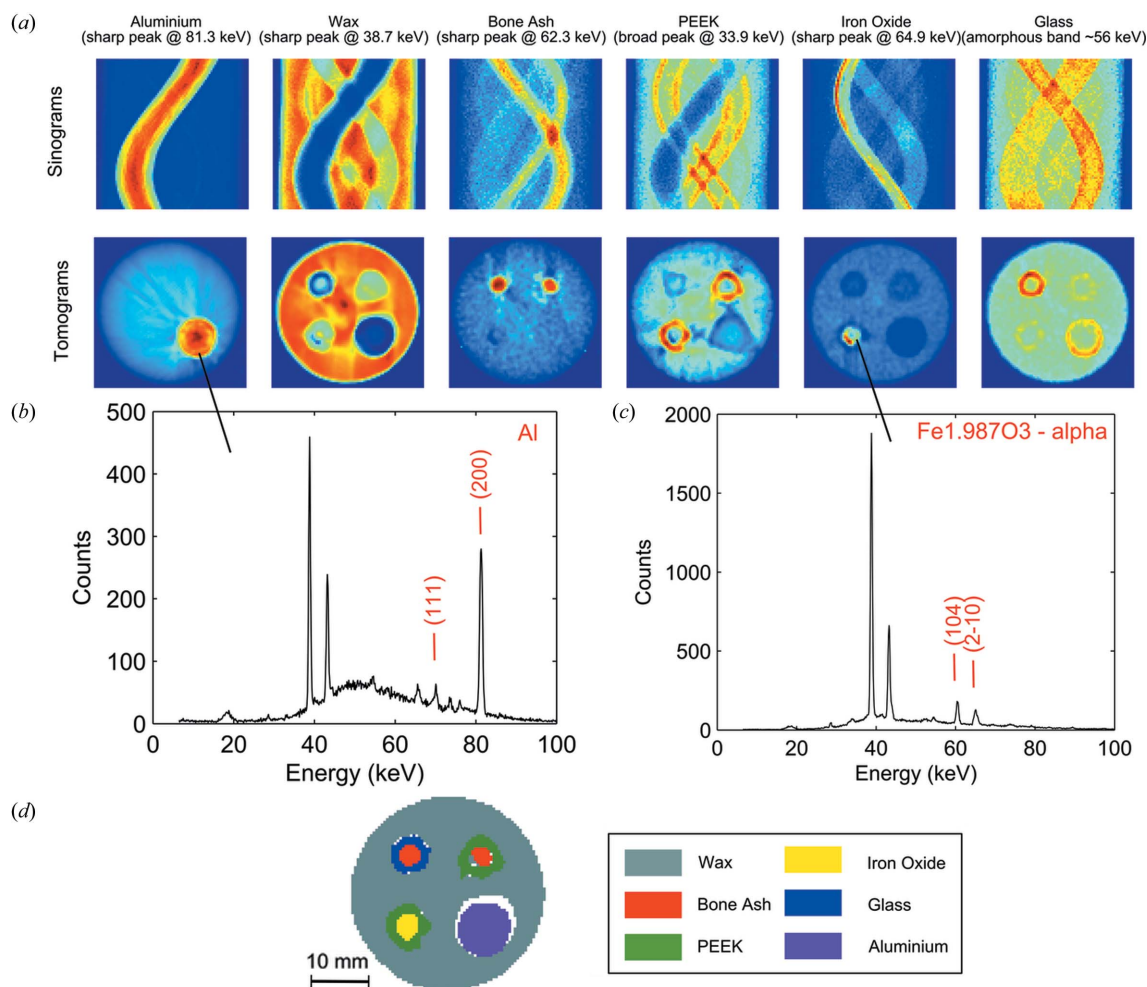
¹ Supplementary data for this paper are available from the IUCr electronic archives (Reference: VV5037). Services for accessing these data are described at the back of the journal.

$$L = (h/2)[\cot(2\theta - \alpha) + \cot(2\theta + \alpha)] + (e + d) \sin \alpha [\csc(2\theta - \alpha) + \csc(2\theta + \alpha)], \quad (2)$$

where h is the incident beam size (0.5 mm), α is half the angular acceptance of the collimator, e is the sample–collimator distance and d is half the length of the collimator [corrected form after Rowles (2011)]. For the data presented here the voxel length is calculated to be 8.47 mm, and is therefore extremely squashed in one dimension, effectively limiting the technique to a spatial resolution of about 8 mm. However, because of the novel data collection strategy employed (using a computed tomography approach as described above), the effect of having a squashed-shaped voxel can be negated. This is because the lozenge is continuously scanned through the sample in a direction parallel to the incident X-ray beam, *i.e.* along its longest dimension. This constitutes the collection of one diffraction pattern and therefore limits the spatial resolution to the size of the incident beam (0.5 mm), *i.e.* the same as is found in standard X-ray absorption computed tomography.

In order to correlate particular diffraction peaks to specific materials it is necessary to calculate powder diffraction patterns for suspected materials and compare these with the experimental data. An alternative method would involve using a Rietveld refinement technique to identify specific peaks and therefore materials within the object (Cernik *et al.*, 2011). Phase extraction proceeds by selecting an energy or angle channel and spatially mapping the intensity of that channel. An alternative method involves fitting the identified diffraction peak and spatially mapping fitted parameters such as peak width or integrated intensity for those phases. In this paper we choose the latter method since it is more sensitive to subtle variations in the data, for example in differentiating two overlapping peaks. Furthermore, it gives us the ability to map peak position (*i.e.* strain) or peak area (*i.e.* quantitative amount) for materials within the sample.

Phase extraction can either be performed before or after tomographic reconstruction. If phase extraction is to be performed before reconstruction, only a small number of reconstructions need to be performed corresponding to each phase that has been identified. An alternative method is to reconstruct the object for all energy channels and then perform phase extraction, which is what we do here. The phase-extracted sinograms and reconstructed images (tomograms) of the phantom object using the TEDDI technique are shown in Fig. 5(a). We show each of the six phases within the sample along with the energy of the diffraction peak from which the phase was extracted. Each phase was mapped by peak fitting using a Gaussian profile and presented as the integrated intensity of the diffraction peak. We also show two diffraction patterns taken from localized regions within the reconstructed data using the reverse analysis method of Bleuet *et al.* (2008), one from a central region on the aluminium phase (Fig. 5b) and the other from the iron oxide phase (Fig. 5c). Principal diffraction peaks from these phases are identified. Finally we present the colour-coded composition map of the phantom object (Fig. 5d).


Figure 5

(a) Sinograms and tomograms for each of the six phases present in the sample. Gaussian functions are fitted to the diffraction peaks specified and mapped as the peak area. (b) and (c) Localized diffraction patterns from aluminium and iron oxide regions. Principal diffraction peaks from these phases are identified. (d) Final colour-coded composition map obtained from the sample.

The first thing to note is that the wax phase has not accurately been reconstructed as the ‘holes’ in the wax where other materials lie are not circular. This is also true for the PEEK phase where the ‘rings’ are also found not to be exactly circular. These two effects are due to X-ray absorption from other phases within the sample, most notably the large aluminium cylinder. This can more clearly be seen in the sinograms. For example, if we look at the PEEK sinogram we can see that whenever the aluminium cylinder crosses the path of the PEEK there is significant absorption resulting in reduced intensity. This reduced intensity causes the PEEK phase to have an irregular shape when reconstructed. However, some of the other phases are very accurately reproduced, e.g. the aluminium and bone ash. This is largely due to the higher energy of these diffraction peaks and thus reduced X-ray absorption within the sample. The iron oxide phase has also been quite accurately mapped, although there does appear to be quite a large amount of self-absorption from this phase resulting in an apparent inhomogeneous distribution. One interesting point to note was that of the iron oxide crystal structure. It became apparent that only an α -phase

Fe₂O₃ or hematite diffraction pattern would match the iron oxide phase peaks, thus confidently identifying this particular phase. Interestingly, it was not possible to conclusively identify this phase from ADD data, as will be discussed later. It was more difficult to extract the amorphous glass phase, because there are no distinct diffraction peaks. Instead we performed a peak fit to the amorphous ‘hump’ [most clearly seen in the logarithmic plot in Fig. 3(b)] using three separate Gaussian profiles. The phase was extracted by plotting the peak height at about 56 keV. This method was quite successful since the reconstructed image is largely true to the original phantom object; however, there is a certain amount of disturbance from other phases within the object.

From the images of the extracted phases it is obvious that the TEDDI technique suffers from self-absorption for lower-energy diffracted photons. The extracted wax and PEEK phases, which have the lowest diffraction energies, are particularly affected. It is therefore necessary to apply an absorption correction to the data in order to obtain a true reconstruction of the object. This absorption correction should be applied without any prior knowledge of the object

morphology or material content. The absorption correction method requires the acquisition of an absorption map, which, for example, could be obtained from a standard X-ray absorption tomogram. Strictly, the absorption map is required to be energy-resolved in order to account for absorption throughout the energy range of interest. This, for example, could be achieved using the same energy-dispersive detector in transmission mode. However, there may be issues with high levels of flux owing to the limited count rates of such detectors, which could be mitigated, for example, by placing a known calibrated attenuator in the incident beam. Another option might be to obtain absorption maps using monochromatic X-rays, at a number of different energies that broadly span the energy range of interest (*e.g.* 30 keV, 50 keV, 70 keV) and interpolate between them, creating a pseudo-energy-resolved absorption map. This has the advantage of being significantly faster since area detectors can be used. This technique could also be applied at regions of interest in the diffraction pattern, for example at 38 keV if one wanted to correct the wax diffraction peak.

The absorption map can then be ‘forward projected’ creating an absorption sinogram, taking into account the trajectory of the diffracted photons by integrating over their path length, *i.e.* through an angle 2θ . With our collection geometry diffracted X-rays are only collected within one plane (defined by the collimator), which eliminates absorption artefacts from out-of-plane diffraction. This offers some advantages over other methods which collect diffracted X-rays through conic sections that can be significantly harder to correct for absorption. The absorption correction can now be applied to the original data using the Beer–Lambert law,

$$I_{\text{corr}}(E) = I_{\text{m}}(E) / \exp\left[\int \mu(E) dt\right], \quad (3)$$

where $I_{\text{corr}}(E)$ is the corrected intensity (as a function of energy), $I_{\text{m}}(E)$ is the measured intensity and $\mu(E)$ is the effective linear absorption coefficient for each pixel in the absorption map. Integration is performed over the pixelated path of the diffracted X-rays. Application of the absorption-correction algorithm, in this case correcting for X-ray absorption from the aluminium cylinder, is demonstrated in Fig. 6. As the absorption map is rotated, the absorption-correction sinogram is built up by integrating over the pixelated path of the diffracted photons and applying the Beer–Lambert law. The absorption-correction sinogram therefore has numerical values greater than 1, with 1 corresponding to no correction to be applied. In this case, note that the ‘intensity’ of the absorption-correction sinogram is not uniform. This is because as the aluminium cylinder is rotated it ‘shadows’ different portions of the sample; sometimes a large proportion of the sample is shadowed, sometimes only a small amount. Once the absorption-correction sinogram is built, it is multiplied with the original data sinogram and then back-projected to obtain the final (absorption-corrected) tomogram. The result of applying this correction (accounting for absorption from the aluminium cylinder) when reconstructing the wax phase is shown in Fig. 6(c). We can see that the correction has been successful in that the ‘hole’ in the wax

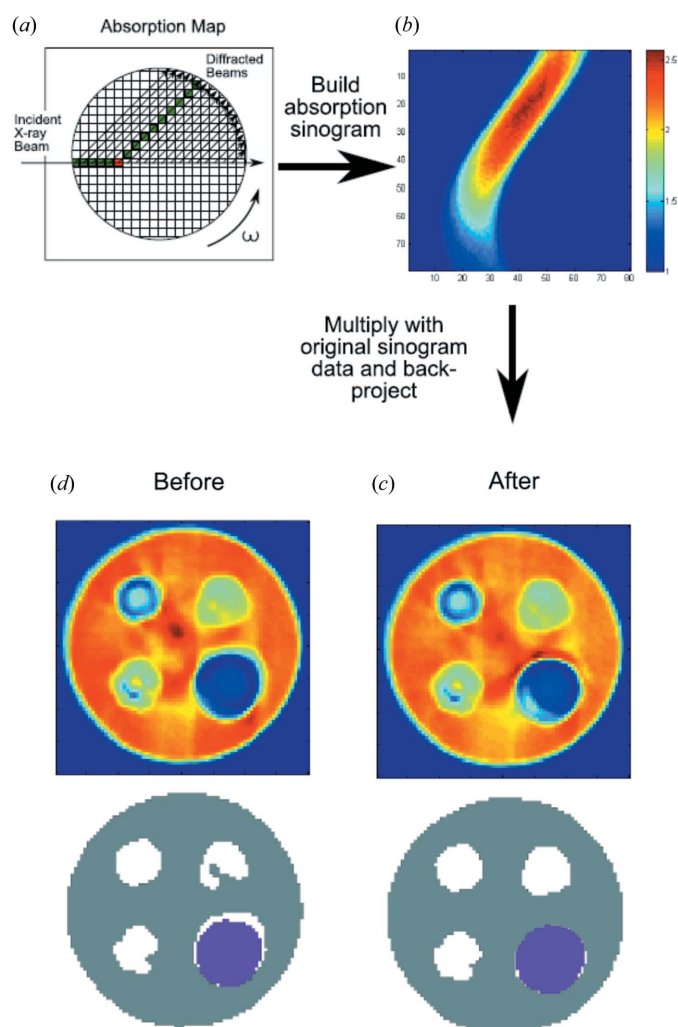


Figure 6 Correcting the wax phase reconstruction for X-ray absorption by the aluminium cylinder. (a) For each projection the absorption map is integrated over the pixelated route of the diffracted photons. This is then summed along the length of the incident beam path, thus simulating the TEDDI data collection strategy. The Beer–Lambert law was then applied to each projection to build the absorption-correction sinogram. This is then multiplied with the original data sinogram and then back-projected to finally obtain the corrected image (c). This compares with the uncorrected original (d).

(grey) now has a tight fit around where the aluminium cylinder (purple) sits, as compared with the original reconstruction (Fig. 6d). Thus, with the application of the absorption correction we can obtain a better representation of the original sample. Whilst this correction has been demonstrated for TEDDI, it is not difficult to adapt it for TADDI; all that is required is to adjust how the integration over the pixelated route is performed, *i.e.* taking into account diffraction into a cone.

One of the major limitations of the TEDDI method, discussed earlier in the paper, was the significant collection time required to image an entire sample. This was found to be partly due to the fact that the voxel had to be continuously scanned through the sample in a direction parallel to the X-ray beam for every translation and rotation, a process that took

approximately 10 s for this sample. However, the main temporal constraint is found in the count-rate limitations of energy-dispersive detectors. For TADDI, on the other hand, data can be recorded in a single static measurement, thus taking only about 0.5 s or faster (Jacques *et al.*, 2011). This makes TADDI approximately 40× faster than TEDDI, and therefore is a severe limitation in the case of time-resolved experiments. However, this time disadvantage is for a *single* collimator and detector system (the so-called zero-dimensional case), and the TEDDI technique offers significant scalability using array collimators and detectors, drastically increasing the speed of this method. Situations can be envisioned whereby several single point systems could be combined into a linear one-dimensional array thus reducing the time for each measurement. If 20 zero-dimensional systems are combined into a one-dimensional array, we can match the time resolution of the TADDI technique. Further, if we could build array collimators coupled with pixelated detectors forming two-dimensional arrays, we would drastically reduce the time to make a measurement. For example, if we had a 20 × 20 collimator and detector array (Cernik *et al.*, 2008; Tunna *et al.*, 2006) we could almost achieve a 10× speed advantage over the angle-dispersive case. TADDI, on the other hand, has very limited scalability in this respect since it is only possible to use one detector for each measurement. Situations could be envisioned where diffraction patterns could be recorded from an entire sample in a single static measurement using suitably large detector and collimator array systems.

Finally, with regard to the sensitivity of the TEDDI method, we estimate that the detection limit is around 1 wt%, although this will be highly dependent upon the sample in question. This figure is about 10× less sensitive than that quoted for the ADD technique (Bleuet *et al.*, 2008), but we feel this is fair considering the limited azimuthal range the collimator/detector samples. Having said that, the EDD method does take advantage of the broadband nature of the white beam and it does not require flux-limiting optics, which should increase the signal-to-noise ratio.

5. Summary and conclusions

In summary, we have devised a new data collection strategy for performing energy-dispersive X-ray diffraction computed tomography. This method overcomes the problems associated with the highly elongated shape of the diffraction lozenge, meaning EDD tomography can now be used to study systems with small features (*e.g.* thin coatings or particulates). The method has also been shown not to suffer from sample broadening effects because the diffracting volume is essentially a point source, meaning the TEDDI method is well suited to studying large objects (>3 mm). This is in contrast to

angle-dispersive diffraction which is severely affected by sample broadening effects and therefore not appropriate for studying larger objects. For sufficiently absorbing materials an absorption correction is required in order to produce an accurate representation of the sample. Details of an absorption-correction algorithm have been outlined in this paper. Unfortunately the current single-point TEDDI method is very slow meaning the method is limited to studying static systems. However, there are significant opportunities for scalability through the manufacture of array collimator and detector systems. In theory this can reduce acquisition times by a factor of 20 using modestly sized array systems, allowing for *in situ* dynamic studies.

The authors wish to thank: EPSRC and CCLRC for financial support, ESRF and SRS for allocated beam time, and Mike Miller, formerly of SRS, for his efforts in modifying the station software to facilitate the energy-dispersive diffraction computed tomography.

References

- Bleuet, P., Welcomme, E., Dooryhée, E., Susini, J., Hodeau, J. L. & Walter, P. (2008). *Nat. Mater.* **7**, 468–472.
- Brenner, D. J. & Hall, E. J. (2007). *N. Engl. J. Med.* **357**, 2277–2284.
- Brooks, R. A. & Di Chiro, G. (1976). *Phys. Med. Biol.* **21**, 689–732.
- Cernik, R. J., Hansson, C. C. T., Martin, C. M., Preuss, M., Attallah, M., Korsunsky, A. M., Belnoue, J. P., Jun, T. S., Barnes, P., Jacques, S., Sochi, T. & Lazzari, O. (2011). *J. Appl. Cryst.* **44**, 150–157.
- Cernik, R. J., Khor, K. H. & Hansson, C. (2008). *J. R. Soc. Interface*, **5**, 477–481.
- Clark, S. M., Cernik, R. J., Grant, A., York, S., Atkinson, P. A., Gallagher, A., Stokes, D. G., Gregory, S. R., Harris, N., Smith, W., Hancock, M., Miller, M. C., Ackroyd, K., Farrow, R., Frances, R. & Ohare, D. (1996). *Mater. Sci. Forum*, **228**, 213–217.
- Grant, J. A., Morgan, M. J., Davis, J. R. & Wells, P. (1995). *J. Opt. Soc. Am. A*, **12**, 291–300.
- Hall, C., Barnes, P., Cockcroft, J. K., Colston, S. L., Hausermann, D., Jacques, S. D. M., Jupe, A. C. & Kunz, M. (1998). *Nucl. Instrum. Methods Phys. Res. B*, **140**, 253–257.
- Harding, G., Kosanetzky, J. & Neitzel, U. (1987). *Med. Phys.* **14**, 515–525.
- Jacques, S. D. M., Di Michiel, M., Beale, A. M., Sochi, T., O'Brien, M. G., Espinosa-Alonso, L., Weckhuysen, B. M. & Barnes, P. (2011). *Angew. Chem. Int. Ed.* **50**, 10148–10152.
- Jonge, M. D. de & Vogt, S. (2010). *Curr. Opin. Struct. Biol.* **20**, 606–614.
- Kak, A. C. & Slaney, M. (1988). *Principles of Computerized Tomographic Imaging*. Philadelphia: SIAM.
- Kleuker, U. (1997). *X-ray Diffraction Computed Tomography: A Survey and Description*. Bellingham: International Society Optical Engineering.
- Lazzari, O., Jacques, S., Sochi, T. & Barnes, P. (2009). *Analyst*, **134**, 1802–1807.
- Rowles, M. R. (2011). *J. Synchrotron Rad.* **18**, 938–941.
- Stock, S. R. (2008). *Int. Mater. Rev.* **53**, 129–181.
- Tunna, L., Barclay, P., Cernik, R. J., Khor, K. H., O'Neill, W. & Seller, P. (2006). *Meas. Sci. Technol.* **17**, 1767–1775.