Microstructure reconstruction of fibrous C/C composites from X-ray microtomography

Julio Martin-Herrero, Christian Germain

Microstructure reconstruction of fibrous C/C composites from X-ray microtomography

J. Martín-Herrero [a,*], Ch. Germain [b]

[a] *Dept. Signal Theory and Communications, Univ. of Vigo, 36310-Vigo SPAIN*

[b] *Equipe Signal et Image, LAPS-UMR 5131 CNRS, Univ. Bordeaux1, 33405-Talence FRANCE*

Abstract

The precise characterization of the microstructure of fibrous composites is essential for an accurate determination of their properties and behaviour. However, fabrication processes usually introduce serious deviations from simple ideal spatial arrangements of preform fibres and matrix. The characterization of composites has been performed using destructive and non destructive testing techniques, many of them based on image analysis. Carbon fibres in carbon matrix composites, however, pose a difficult challenge to image analysis systems due to the poor contrast between the different phases. We describe a procedure for the reconstruction of the true microstructure of carbon/carbon composites using phase contrast X-ray microtomography, show some results, and discuss its performance.

Key words: Carbon/carbon composites, Chemical vapor infiltration, Image analysis, Microstructure

* Corresponding author. FAX: +34 986 812 116
Email address: julio@uvigo.es (J. Martín-Herrero).

1 Introduction

The special properties of fibre reinforced composites responsible for their success are based on their inherent anisotropy, a characteristic dictated by the spatial arrangement of the constituent phases. Indeed, the microscopic arrangement of the complex fibrous structure and matrix plays a major role in the determination of the macroscopic properties of the material, such as their elasticity, strength, permeability, or thermal conductivity. As such, there is a whole range of numerical methods designed to estimate bulk properties assuming the characteristics of the phases and their spatial arrangement are known, such as those based on Finite Elements Models. Among the structural characteristics that influence the properties of the material are those regarding the fibres, such as the distribution of lengths and orientations, waviness, curvature, and volume fraction, and similar quantities for the matrix and void phases. Manufacturing techniques are fine tuned to control these characteristics, but the tuning process usually requires repeated cycles of fabrication and comparison with characterized samples. The recent advances in computer aided modelling and simulation have not eliminated the necessity of experimental work, but rather have increased the demand of accurate experimental data for training, testing, and validation, once the early models based on ideal, regular, or uniform structures have been superseded. Moreover, no matter how much fine tuning, there is no such a thing as the perfect manufacturing process and, therefore, different factors in the fabrication chain introduce in the real microstructure of the product deviations from the theoretical design which affect the predicted performance of the material.

Costly, limited, error-prone manual methods [1, 2] for the characterization

of samples of fibrous composites have been outdated by automatic methods, such as those in [3–20], using modern imaging systems [21] and digital image processing [22]. Most of them are based on optical microscopy, and therefore limited to the study of a 2D cross-section, serial sectioning, or, at most, limited-depth volumes of translucent materials using optical confocal microscopy [23]. However, only some aspects of the 3D microstructure can be reconstructed from a 2D section, such as fiber orientation [20], subject to a series of assumptions [24]. Producing and analyzing series of close, thin parallel slices of the samples [25], such as in [7, 9, 11, 19], is labour intensive and adds the problem of matching among slices.

X-ray microtomography [26, 27] is a non destructive technique that produces accurate images of 3D volumes by reconstruction from multiple X-ray projections [28], allowing the direct characterization of the 3D microstructure of samples with significant thickness. The huge amount of data produced by this kind of imaging system for large enough samples, even at moderate resolutions, precludes any manual processing such as fibre identification. However, 2D image processing techniques usually do not have a direct translation into the 3D domain, and processing a 3D microtomograph just as a series of 2D sections is clearly a waste of valuable information. Therefore, specific 3D techniques are needed [29, 30], but the combined effect of the increased amount of data to process and the added dimension (most processing algorithms scale geometrically with the dimension of the dataset) calls for extreme care in the selection of the methods and their optimization.

Carbon/carbon (C/C) fibre reinforced composites [31] are made of carbon fibres embedded in a continuous matrix of carbon. The fact of both matrix and fibres being fabricated from carbon produces a unique combination of

properties, including very low thermal expansion coefficients and high thermal conductivity, retaining their mechanical properties at high temperatures ($> 2000°C$ in nonoxidizing atmospheres), high specific strength, excellent resistance to abrasion and ablation, high resistance to thermal shock, very high elastic modulus, low density, high electrical conductivity, low hygroscopicity, nonbrittle failure, low biological rejection, resistance to chemical corrosion, and reasonable machinability [32].

However, C/C composites pose a significant challenge to X-ray imaging systems due to the close densities of inclusions and matrix, yielding very poor contrast among the constituent phases. Sensitivity can be increased for composites made up of materials with neighbouring densities such as C/C composites by using coherent X-rays (e.g. from a synchrotron source) and varying the distance from sample to detector [27]. Yet this technique, called holotomography, applies only to very small C/C samples for which all fibres show the same orientation. Indeed, in larger samples, or in samples where fibres perpendicular to the tomography occur, these fibres cause very high cumulated phase lag values which saturate the detector. The resulting image is blurred, rendering the separation of fibres and matrix impossible [33]. The alternative, also requiring a coherent X-ray source, is phase contrast tomography, also called edge-detection mode CMT, where interfaces between constituent phases are detected through the interference patterns arising from the out of phase waves propagated through materials with different refraction indices [34].

Previous attempts have been made to characterise the microstructure of a C/C composite using X-ray microtomography [18, 34]. However, the carbon matrix proved to be very difficult to separate from the carbon fibres, and only the separation of the porosity from the solid phases was achieved. Later, the

method was succesfully applied to the separation of the three phases in an Al-SiC_p composite [35, 36]. On the other hand, the X-ray microtomography segmentation procedures in [29, 30] deal with easily separable phases, such as glass fibres in a polypropylene matrix, therefore the constituents are clearly resolved in the microtomographs and the segmentation phase is straightforwardly addressed by simple thresholding, ineffective in the case of a C/C composite synchrotron X-ray microtomography.

In the following, we will show a procedure for the reconstruction of the 3D microstructure of a C/C fibrous composite from a phase contrast X-ray microtomography, based on advanced image processing techniques, allowing the separation of fibres, porosity, and matrix, and the subsequent analysis, characterization, and even visualization by means of enhanced reality techniques as an aid to design and analysis. Section 2 describes a sample we use as study case to illustrate the method, described in section 3. The description is accompanied by some results obtained with the study sample. Finally, section 4 is a brief discussion and section 5 closes with a short summary.

2 A C/C fibre reinforced composite microtomographic image

The image in Figure 1 was obtained at the European Synchrotron Radiation Facility (ESRF) ID19 High Resolution Diffraction Topography Beamline, dedicated to radiography, microtomography, and diffraction imaging experiments [27, 34]. The X-ray beams produced at third generation synchrotron radiation facilities such as the ESRF at Grenoble (France) have a high degree of coherence. This results from the small source size, σ, about 50 μm, and the large distance from source to sample, L, in the 100 m range, so that

the transverse coherence length, dc $= \lambda L/2\sigma$, is in the 100 μm range. Phase changes occur at the edges of a particle or porosity embedded in a matrix having a different index of refraction, and out of phase transmitted radiations produce an interference fringe marking the location of the interface. The image shown was obtained using phase contrast with fixed distance from sample to detector [34].

The sample comes from the French aeronautics industrial company Snecma Propulsion Solide (SPS). SPS has decades of experience in isothermal chemical vapour infiltration (CVI) [32, 37, 38] for the development of proprietary C/C, C/SiC, and SiC/SiC composites for industries as varied as heat treatment, silicon, electronics, glass, nuclear power, and space. Figure 1 shows a $200 \times 200 \times 200$ portion of a synchrotron microtomograph of a sample of one of SPS's C/C composites at an incomplete densification stage (sample "CC2" in [34]). Figure 2 shows a 200×200 slice of the microtomograph in Figure 1.

Fig. 1. A synchrotron microtomograph of a 0.0033 mm^3 (side length < 0.15 mm) cubic sample of a C/C composite. Voxel size is 0.745 μm.

Given that we are dealing with a C/C composite, there is no doubt that the

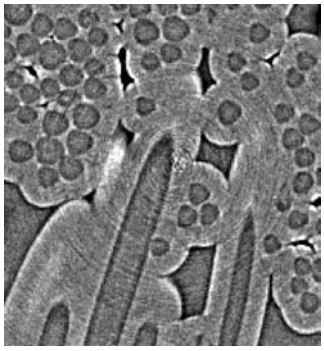

Fig. 2. A slice (0.0222 mm^2) of the microtomograph in Figure 1.

image in Figure 1 provides a very good discrimination between fibres and matrix to the human observer. The porosity typical of the CVI densification process [39–41] can also be clearly seen in the image. However, Figures 1 and 2 also show a high level of noise (greylevel variations not due to differences in the material but to added noise, such that the same greylevels appear mixed in different phases), poor use of the dynamic range (few different meaningful greylevels, see below), and moderate resolution. Resolution is expected to improve in the near future. X-ray detection is performed through film or visible light scintillators, and, therefore, the system is diffraction limited. X-ray lenses to magnify the image before the detector, such as Kirkpatrick-Baez focusing devices, may overcome this limitation. Nevertheless, resolution is also limited by data bandwidth and storage capabilities.

However, limited resolution is not the worst characteristic of the image. Noise levels and poor quantization are. In Figure 2, apart from random noise, rings of varying intensity can be clearly seen around the tomography axis, an artefact due to variations in the efficiency of neighbouring pixels in the detector. Figure 3 shows the histogram of the slice in Figure 2. Out of 256 possible levels in the 8 bit quantization, only 53 have a pixel count different from zero. Moreover, 51 of the 53 levels are grouped in 25 groups, thus reducing the effective dynamic

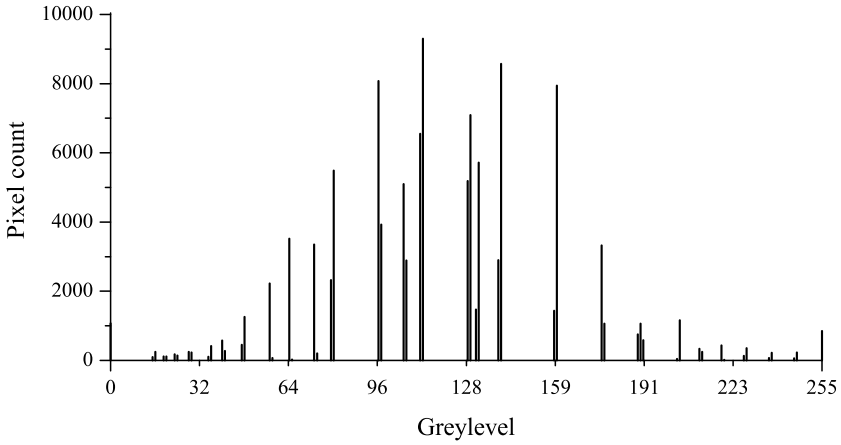

Fig. 3. Histogram of the image in Figure 2. There are 53 levels out of 256 possible levels in the dynamic range, grouped in 27 groups, thus in practice reducing the effective dynamic range to 27 levels.

range to 27 levels, i.e. the equivalent to less than 5 bit quantization.

The effect of the noise and the reduced dynamic range can be seen in the magnified details in Figure 4. Close scrutiny reveals that, in spite of what could be assumed from Figure 1, greyscale based per-pixel segmentation of fibres from matrix by simple thresholding is not possible. The same grey levels are found in all phases of the material. The human visual system is able to segment the image because it performs high level processing involving edge detection and pattern matching with circular/elliptical primitives [42–44] corresponding to the expected fibre cross-sections (see the small windows at the lower right corners). Note for instance the effect that just a sparse couple of brighter pixels have on the recognition of the fibre in the lower right corner of the upper part of Figure 4.

Within this context, our aim is to use image processing techniques for the automated extraction of fibres and porosity from the matrix, to allow the detailed characterization of the material.

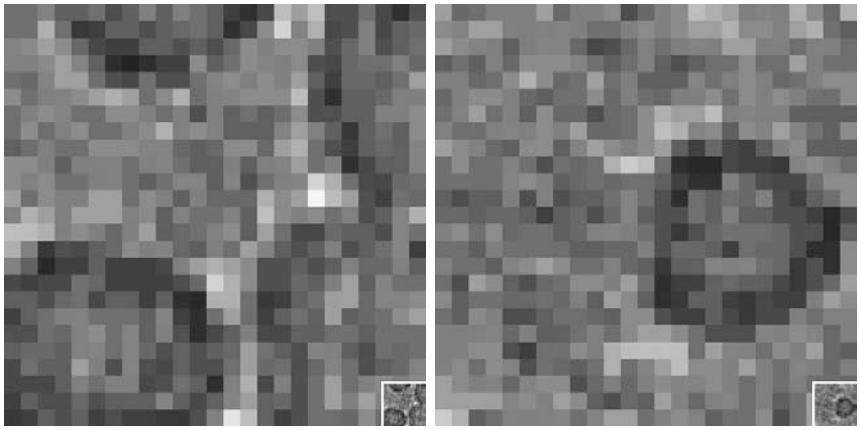

Fig. 4. Details of fibres/matrix in the image in Figure 2. The small windows at the lower right corners show the entire detail at normal scale.

3 Method

3.1 Preprocessing

As we have just seen, the grey levels of individual voxels are not enough to distinguish fibre voxels from matrix voxels. It is the edges —contrast among adjacent regions in the image— that determine the boundaries of the fibres. Fibres are restituted by our visual system from the edge information. Thus, we need the edges to define the fibres in the image. We have to reduce the noise level, but any noise filtering must respect the edges. Otherwise, we would be losing the information we are seeking, the basis of the segmentation procedure. The noise has a high spatial frequency, but low pass filtering would also blur the edges. Median filtering could preserve to some extent the edges, but the resolution is too poor: The edges are too narrow as to not be wiped away by any but the smallest of filters, which would be useless. A noise reduction method capable of efficiently removing high frequency noise while preserving edges is anisotropic diffusion [45, 46].

Perona and Malik [45] noted that the convolution of an image $I_0(x, y)$ with a Gaussian kernel,

$$K_\sigma(x, y) = \frac{1}{2\pi\sigma^2} \exp(-\frac{|x|^2 + |y|^2}{2\sigma^2}) \tag{1}$$

with standard deviation σ yields the same result as the solution of the isotropic diffusion (heat) partial differential equation,

$$\frac{\partial I(x, y, t)}{\partial t} = \mathrm{div}(\nabla I(x, y, t)), \tag{2}$$

where $I(x, y, t)$ is the image $I(x, y)$ at time $t = 0.5\sigma^2$, with initial conditions $I(x, y, 0) = I_0(x, y)$, and ∇I is the image gradient.

Introducing in (2) as *diffusion conductance* or *diffusivity*, $g(s)$, a rapidly decreasing function of a greyscale edge detector such as the gradient magnitude, $s = |\nabla I|$,

$$\frac{\partial I(x, y, t)}{\partial t} = \mathrm{div}(g(|\nabla I(x, y, t)|)\nabla I(x, y, t)), \tag{3}$$

smoothing on both sides of edges becomes much stronger than across them. A diffusivity constant with time but varying with location (x, y) would make (2) a linear nonhomogeneous diffusion equation. However, if g is made a function of time, as in (3), the diffusion equation becomes nonlinear and nonhomogeneous, referred to as *anisotropic* in the image processing literature, even when conventional PDE terminology reserves the term for the case where the diffusivity is a tensor, varying both with location and direction. The consequence of this "pseudoanisotropy" is that only the magnitude but not the direction of the diffusion flux can be controlled. Noise close to edges remains unchanged due to the small flux in the vicinity of edges. To enable smoothing

parallel to edges, (3) must be generalized with a diffusivity matrix G with nonzero off diagonal elements (see, for instance, [46–50]), thus rendering it truly anisotropic.

Perona and Malik suggested

$$g(s) = \frac{1}{1 + s^2/\kappa^2} \tag{4}$$

and

$$g(s) = \exp(-s^2/\kappa^2) \tag{5}$$

as diffusivity functions. Since then, many other diffusivity functions have been proposed [46, 51]. In fact, Black *et al.* [51] demonstrated that anisotropic diffusion in the sense of (3) is the gradient descent of an estimation problem with a robust error norm induced by $g(s)$, thus providing a sound theoretical foundation to choose adequate diffusivity functions.

As several authors have revealed [52–55], (3) is an ill-posed problem, in the sense that images close to each other are likely to diverge during the process [54], but it can be stabilized by regularization. One common approach [53] is to smooth the variable of the diffusivity, i.e. to use a smoothed version of the image for the gradient in each step, as in

$$\frac{\partial I(x, y, t)}{\partial t} = \operatorname{div}(g(|\nabla I_\sigma(x, y, t)|)\nabla I(x, y, t)), \tag{6}$$

where $I_\sigma = K_\sigma * I$, with K_σ a suitable local convolution kernel of width σ, for instance a Gaussian kernel. However, Weickert and Benhamouda [56] proved that a standard spatial finite difference discretization is sufficient to turn (3) into a well posed system of nonlinear ordinary differential equations. Therefore,

direct implementations of the Perona-Malik filter tend to work reasonably well because of the regularizing effect of the discretization.

The extension to D dimensions is straightforward. The diffusion process is described by the equation

$$\frac{\partial I(\mathbf{x},t)}{\partial t} = \mathrm{div}(G \cdot \nabla I(\mathbf{x},t)), \tag{7}$$

where $\mathbf{x} \in \mathbb{R}^D$ and G is a square $D \times D$ diffusivity matrix. Equation (3) is (7) with G a diagonal 2×2 matrix with equal diagonal elements $g(|\nabla I|)$. Diffusion filtering is usually implemented using simple finite differences to approximate the image derivatives within an explicit or Euler-forward scheme.

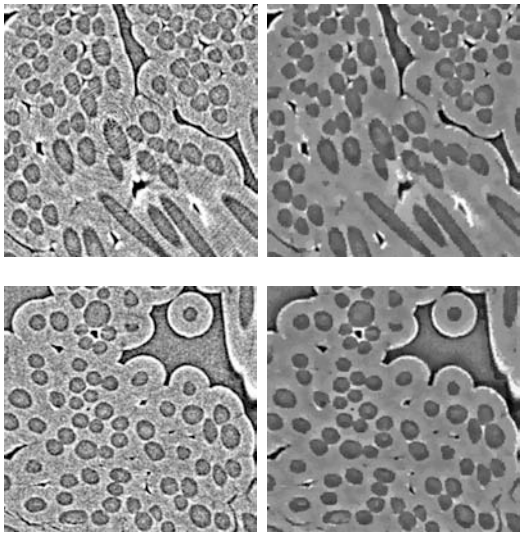

Fig. 5. Two orthogonal slices of the tomograph in Figure 1 before (left) and after (right) filtering with 3D anisotropic diffusion.

Figure 5 shows the results of applying (7) with (4) and $D = 3$ to the tomograph in Figure 1, and they are quite satisfactory. The relevant edges have been preserved while most of the noise has been wiped away. Note that the filter is equally efficient on the two types of noise present in the image, the random uniform noise and the circular patterns around the tomography axis.

3.2 Segmentation

However good the results in Figure 5 regarding noise reduction and mesoscale feature preservation, fibres cannot yet be separated on the basis of individual voxel grey levels. Many fibres show brighter nuclei inside darker boundaries, whose grey levels can also be encountered in the matrix. Our approach is masking all pixels within edges. For that purpose we use a differential profiler [57] along voxel rows. A differential profile is obtained by constructing an array of integers where every element represents the lowpass filtered gradient of the grey levels along a scan line. This is easily implemented using exclusively integer arithmetic. Let $I_{x,y,z}$ be the greyscale image. Let P_l be an integer vector to hold the profile of a scan line, and assume without loss of generality that the scan direction is along $+x$. Combining forward differencing

$$P_l = I_{l,y,z} - I_{l+1,y,z} \tag{8}$$

with lowpass filtering

$$P_l = \frac{1}{3}(P_{l-1} + P_l + P_{l+1}), \tag{9}$$

and forgetting about the dividing constant, we get

$$\begin{aligned} P_l &= I_{l-1,y,z} - I_{l,y,z} + I_{l,y,z} \\ &\quad - I_{l+1,y,z} + I_{l+1,y,z} - I_{l+2,y,z} \\ &= I_{l-1,y,z} - I_{l+2,y,z} \end{aligned} \tag{10}$$

and thus the whole procedure to obtain the profile is reduced to a sequence of integer subtractions of grey levels. Near-zero profile segments reflect parts of the image where no significant variations of grey level occur. Positive peaks reflect decreasing grey levels, such as when entering a fibre, and negative peaks

increasing bright, such as when leaving a fibre. The narrower the peak, the faster the variation. The higher the peak, the greater the variation. Thus, peak shape and magnitude tell us everything we need to know about intensity transitions along a given direction.

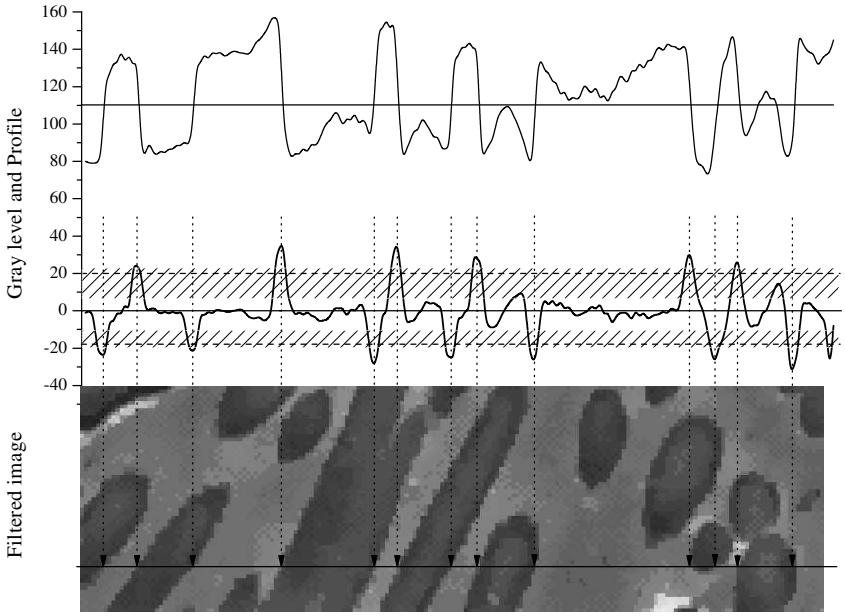

Fig. 6. Edge detection by profiling scan lines in slices of the tomograph in Figure 1 after filtering with 3D anisotropic diffusion. The patterned areas in the profile graph show the range of possible thresholds for successful detection of all fibres in the given line.

Figure 6 shows a grey level profile (above) and the corresponding differential profile (center) along a line of the tomograph in an arbitrary direction (below). A tentative grey level threshold is plotted on the grey level profile (upper graph), to show the lack of robustness of a per-pixel approach due to the brighter nuclei of some fibres. The vertical arrows show the edges in the image corresponding to the peaks above or below tuneable thresholds (dashed lines) in the differential profile (lower graph). The range of possible thresholds that would achieve successful detection of all fibre edges in the line is highlighted

in the graph. The upper (in absolute value sense) bounds for the thresholds are dictated by two criteria: 1) The positive threshold should be low enough as to not miss any positive peak at the entrance of a fibre, and 2) the negative threshold should not miss any negative peak at the exit of a fibre. The lower bound for the positive threshold depends only on the peaks outside fibres (positive peaks inside fibres are irrelevant): the threshold should not allow any peak not lying at the entrance of a fibre. Conversely, the lower bound for the negative threshold depends only on the peaks inside fibres (negative peaks outside fibres are irrelevant): the negative threshold should be high enough as to exclude any negative peak not at the exit of a fibre.

To obtain a mask of fibre voxels, therefore, the voxel rows along the three possible orientations in the 3D image are profiled, such that all voxels along a differential profile from a positive peak to a negative peak are marked as foreground voxels, and the rest as background voxels. In the resulting mask, some fibres are "broken", i.e. the dark boundary does not surround the entire fibre, and therefore some profiles may result in spurious rectilinear spikes protruding from the fibre mask, due to the lack of a negative peak, or well similar rectilinear structures may be missing from the interior of a fibre due to the lack of a positive peak, or a spurious negative peak with sufficient magnitude inside a fibre.

Here is where spatial correlation among the different orientations in the 3D image enters, as interior fibre voxels not marked in a profile in a given direction are likely to be marked in the profiles along any of the transversal directions, and thus the *a priori* knowledge of the morphology of the fibres is enough, in a simple postprocessing stage, to get rid of the majority of the thin, rectilinear, spike-like 2D artifacts protruding from or entering "broken" fibres.

Porosity in the matrix is segmented by the same procedure, only the thresholds for the peak detection in the differential profile are tuned to the characteristic edges of the porosity regions, where transitions are more pronounced than in fibres, see the darker lacunar areas with a brighter rim in Figure 5. This is due to the contrast saturation effect in the matrix/air interface induced by the increased sensitivity of the tomograph in order to resolve carbon/carbon interfaces.

Both the fibre and the porosity masks can be "cleaned" with a labelling algorithm if their inspection reveals it to be necessary. Cleaning the masks consists in getting rid of remnants of matrix material that may be adhered to the outside face of fibre or porosity boundaries due to imprecisions in the exact localization of edges. This is achieved by 1) labelling all connected foreground voxels in the mask (fibres or porosity) with a unique label subject to the condition of their original grey levels being below a given threshold, and 2) labelling as background voxels (i.e. delete from the mask) all connected groups of remaining unlabelled voxels in the mask if any voxel in the group touches a background voxel (matrix). This ensures that the "bright voxel hunt" takes effect only in groups of voxels attached to the exterior boundary of fibres or porosity, without affecting the bright cores of fibres. The result is a purified mask of fibre or porosity voxels where bright voxels are only permitted in the interior of fibres.

The masks thus obtained, see Figure 7 and Figure 8 for some partial renderings of fibre bundles and porosity, can be combined into a unique material description mask, detailing the microstructure of the material. All voxels in the tomograph are classified as belonging to one of the three phases in the material: matrix, fibre, or porosity. This mask can then be directly fed into

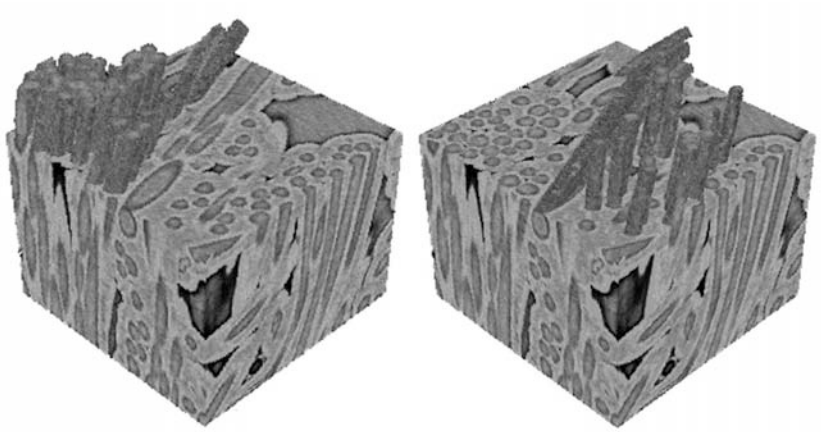

Fig. 7. Partial renderings of some fibre bundles, using the original grey levels in the tomograph.

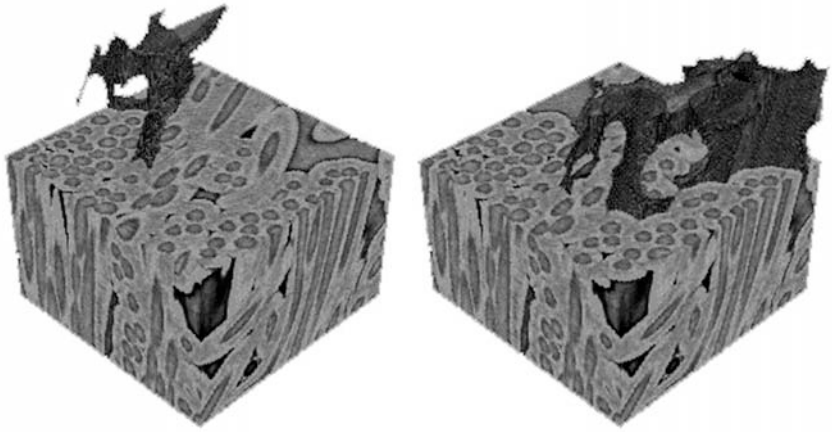

Fig. 8. Partial renderings of the porosity formed during the CVI process, using the original grey levels in the tomograph.

models requiring a precise description of the microstructure of the material. Thus, for instance, a Finite Element Model can be constructed just by assimilating the finite elements to the classified voxels.

3.3 Fibre individuation: The heavy ball

The microstructure of the sample is contained in the mask, indeed, but still some further processing is necessary in order to be able to perform a com-

plete characterization of the microstructure. Some usual characteristics, such as void content or volume fractions, can be directly computed from the phase masks. However, others not less important such as the distribution of fibre lengths, orientations, curvature, or waviness, cannot. Fibres have to be identified individually in order to be able to measure individual characteristics such as those just mentioned.

A connected component labelling algorithm could do it, if we could warrant that individual fibres do not touch each other in the mask. However, this is not the case. Fibres touch each other both because of the limited resolution of the image and because some of them actually touch each other in the material. Therefore a connected component algorithm is able only to separate fibre bunches, groups of fibres either touching each other or close enough as to not be resolved by the spatial resolution of the imaging system, see Figure 7.

We use a novel method, the *heavy ball*, to successfully separate individual fibres. It treats the fibre mask as a solid block drilled by wormholes, the fibres. The hole mouths are located in each side of the block, and a zero drag, zero gravity, high inertial ball is pushed into each hole in turn. The ball size is adapted to the diameter of each hole. The ball is forced to move continuously until it either finds the end of the fibre or leaves the block. The ball run along each fibre can be used to label the fibre with a unique identifier for subsequent characterization, or directly to compute the required parameters.

For locating the ball entrance points and estimating the required ball size, the ellipse fitting procedure depicted in [20, 58], see [59], is applied to each outer boundary (six 2D binary images) of the binary fibre mask. The method requires the labelling of the connected components in each boundary. During

the labelling, the centroid and the covariance matrix of the coordinates of the pixels in each connected component is computed. The centroid (average location) gives the center of the fitted ellipse, and the root of the lesser eigenvalue of the covariance matrix Σ gives its minor axis, which coincides with the fibre radius, r:

$$r^2 = 2(\Sigma_{11} + \Sigma_{22} - \sqrt{\Delta}), \tag{11}$$

where

$$\Delta = (\Sigma_{11} + \Sigma_{22})^2 - 4(\Sigma_{11}\Sigma_{22} - \Sigma_{12}\Sigma_{21}). \tag{12}$$

Touching fibres in the 2D image have to be split before the ellipse fitting procedure. Otherwise a single enclosing ellipse will be fitted to an entire group of touching fibres in the 2D image. The splitting of touching fibres can be performed by means of mathematical morphology operators, using successive erosions and dilations, such that fibres are shrunk until disconnection and then expanded again constrained to the original image. Disconnecting touching groups of pixels is a well known application of mathematical morphology [22].

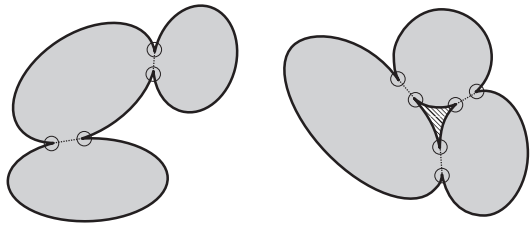

Fig. 9. Ellipse splitting procedure: left) Simple case; right) Complex case not considered in [58]. Small circles mark convex perimeter points. Dotted lines mark splitting lines. The greyed area within the group on the right is enclosed by an internal perimeter.

An alternative approach is described in [58]. The perimeters —foreground

pixels with at least one background neighbour— of the connected components in the 2D images are sequentially traversed and their curvature computed at each perimeter pixel. Pixels where the curvature becomes convex are recorded. If two fibres touch each other, there will be two convex points. A straight line joining the points can be drawn to split the fibres, so they can be labelled as two different components. If three or more fibres touch, it could be just a trivial extension of the two-fibre case, and then a line is drawn for each consecutive pair of convex points, see Figure 9 (left). However, it could also be a special case, not considered in [58], where three or more fibres touch each other, thus totally enclosing a background region inside the group, see Figure 9 (right). In this case, an odd number of convex points is found when traversing the external perimeter of the group. If this is the case, the internal perimeter (the perimeter of the enclosed background) has to be traced also, and the splitting lines should be drawn between pairs formed by a convex point from the external perimeter and the closest convex point in the internal perimeter, see Figure 9 (right). As many convex points have to be found in the internal perimeter as in the external perimeter. If they are not, another background area is enclosed within the group, and so on.

Therefore, if the aggregation pattern of the fibres in a given sample is not too complex, the convex point search is probably faster than mathematical morphology. Otherwise, mathematical morphology is probably a safer choice. No matter what method is used, splitting the touching fibres is a fast operation. Note that it is only performed on the external sides of the 3D domain of the fibre mask, $\mathcal{M}_{\mathrm{f}}$, six 2D binary images. Moreover, the purpose of ellipse fitting is locating entrance points for the heavy ball, and this is performed sequentially. Fibres traversed from the first side are likely to emerge on any

of the sides. Fibres are marked as they are traversed, to avoid repeated runs. Thus, as the process progresses through the sides, each side has less and less ellipses left to be fit.

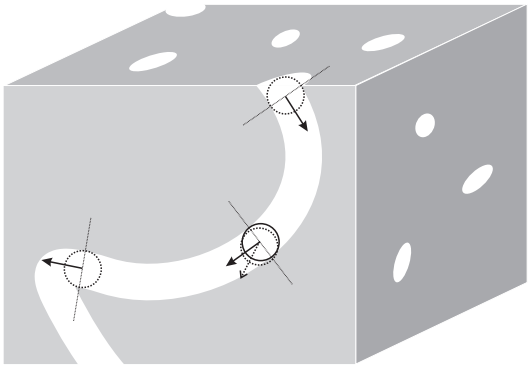

Fig. 10. Identification of individual fibres with the heavy ball method. Solid arrows show the direction of the previous move. The dotted arrow and circle represent a candidate movement. The thinner dotted line marks the limit for candidate moves depending on the previous move. Ellipses in the mask surface are best fits to the fibre sections and determine the radius and initial location of the ball.

Once the center and radius of the "wormhole" have been estimated, a heavy ball of the appropriate size —same radius as the fibre, i.e. minor axis of the fitted ellipse, r, equation 11— is placed centred in the hole mouth. Figure 10 illustrates the method. The ball has inertia, i.e. at step k it remembers the previous move, vector $\mathbf{m}^k$. This memory is initialized opposite to the normal of the current mask boundary, $\mathbf{m}^0 = -\hat{\mathbf{n}}$. The ball moves looking for the widest passage, one voxel at a time. The set of candidate locations for the centre of the ball, $\mathcal{C}^k$, is made of all the voxels C_i adjacent to the current centre C^k such that the direction of advance is at most at right angles with the previous movement. This is verified by means of the scalar product of the vector of the

previous move and the vector linking the current centre and the candidates:

$$\mathcal{C}^k = \left\{ \mathrm{C}_i \in \mathcal{M}_\mathrm{f} \mid \|\overline{\mathrm{C}^k\mathrm{C}_i}\|_0 = 1 \ \wedge \ \overline{\mathrm{C}^k\mathrm{C}_i} \cdot \mathbf{m}^k \geq 0 \right\} \tag{13}$$

For each step, the number $s(\mathrm{C}_i)$ of fibre voxels overlapped by the ball in each candidate location C_i is computed:

$$\mathcal{O}_i^k = \left\{ \mathrm{P} \in \mathcal{M}_\mathrm{f} \mid f(\mathrm{P}) = 1 \ \wedge \|\overline{\mathrm{C}_i\mathrm{P}}\|_2^2 \leq r^2 \ \right\} \tag{14}$$

$$s(\mathrm{C}_i) = \mathrm{I}\left(\mathcal{O}_i^k\right) \tag{15}$$

where $f : \mathcal{M}_\mathrm{f} \subset \mathbb{Z}^3 \to \{0, 1\}$ is the binary fibre mask.

If the maximum overlap is verified for more than one candidate location, the one closest to the previous direction is chosen, i.e. the one maximizing the scalar product of the intended and the previous movement:

$$\mathrm{C}^{k+1} = \arg\max_{\mathrm{C}_i \in \mathcal{C}^k} \left\{ s(\mathrm{C}_i) \wedge \overline{\mathrm{C}^k}\mathrm{C}_i \cdot \mathbf{m}^k \right\} \tag{16}$$

This step by step advance is repeated until the centre of the ball leaves the fibre, either because the fibre ended inside the mask, or because the ball is leaving the mask, $\mathrm{C}^k \notin \mathcal{M}_\mathrm{f}$.

During the advance of the ball the centre locations $\{\mathrm{C}^k\}$ can be recorded —for instance for a curvature analysis—, the length of the moves accumulated, $\sum_k \|\mathbf{m}^k\|_2$, —to estimate the length of the fibre—, the fibre voxels $\bigcup \mathcal{O}^k$ labelled with a unique identifier —for further characterization in a later occasion or involving several fibres, such as nearest neighbour and average distance computations—, and appropriate statistics can be computed, such as extreme locations of the fibre centres $\{\mathrm{C}^k\}$, or variations in the fiber width or size

by inspection of the cross sections at right angles with the current movement vector, $\mathbf{m}^k$.

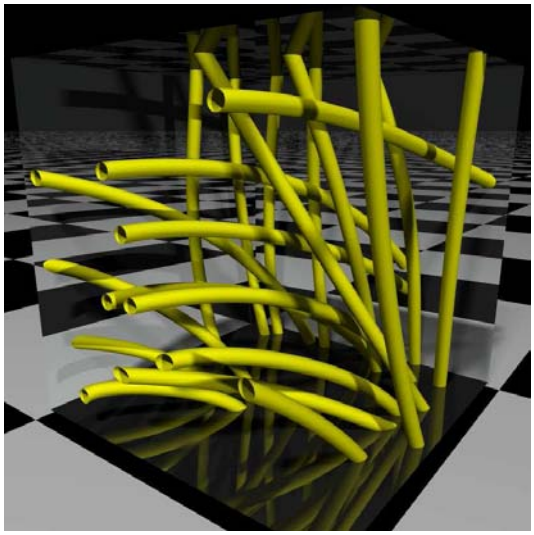

Fig. 11. Rendering of some of the fibres in the sample in Figure 1 after separation with the heavy ball. Only some fibres are shown to avoid visual clutter.

Figure 11 shows some isolated fibres in the sample in Figure 1, where the output of the heavy ball algorithm has been used to produce a ray tracing rendering.

3.4 Labelling algorithms

Several steps in the methodology depicted above use connected components labelling algorithms, both in 2D and 3D.

2D labelling is required for the ellipse fitting procedure. The components targeted by the algorithm are, if the previous processing was successful, small to medium compact rounded shapes. Here the simple classical recursive labelling algorithm [60, 61], also known as region growing or region burning, is likely to work fast and well. Moreover, the labelling aims at computing moments of

each connected component, and this can be easily done during the labelling with the recursive method. The limited size of the components is likely to prevent any stack overflow problem arising from its recursive nature, the cause why it is usually discarded in benefit of the two-pass iterative algorithm.

In the classical two-pass iterative method [62, 63], first the image is traversed in raster scan order assigning temporary labels to the components, and then a second pass resolves conflicts in components with multiple labels, thus requiring an equivalence table to record label conflicts, usually handled with a Union-Find data structure [64]. Therefore, in exchange of the robustness implicit in its iterative nature, it is far more complex than the recursive algorithm, and characterizing components during the labelling is not straightforward, as components are not labelled sequentially and labels are not resolved until the second pass.

Recently a hybrid algorithm has been proposed [65, 66], that has proven to be faster and more robust for compact objects than the classical recursive algorithm, whereas maintaining its simplicity and versatility. The 2D labelling in the ellipse fitting procedure is therefore a good instance for adoption of the hybrid technique.

3D labelling is used to purify the phase mask. Labelling in 3D, however, can be trickier. The components are not moderately sized anymore. We are dealing now with three dimensions, and the recursion depth and the required stack grow accordingly. Thus serious doubts arise about the robustness of the algorithms relying on recursion. The classical recursive algorithm can be rewritten in iterative form, by using a custom stack. However, this comes at the cost of increased complexity.

	Rec. depth		Time (ms)	
Recursive	281 607		146.9	+62%
Iterative (Stack)	252 341	89.6%	167.3	+84%
Hybrid along x	27 617	9.8%	108.9	+20%
Hybrid along y	26 246	9.3%	99.7	+10%
Hybrid along z	15 481	5.5%	90.9	
Two-pass	N/A	N/A	497.8	×5.5

Table 1

Performance data in the labelling of the fibre mask with different labelling algorithms.

We analyzed the performance and robustness in 3D of the different alternatives in the labelling of the phase mask. Table 1 summarizes the results, including different scan directions for the hybrid algorithm, which is anisotropic, in the sense that it favours a direction of choice. Note the effect of the inherent anisotropy of the fibrous components, even when many fibres in the sample deviate significantly from the main orientation, z, specially towards y, see Figure 11 and axes in Figure 1. All data were obtained on a iPentium M 2.13 GHz, 1 GByte RAM, OS MS WindowsXP. Both time and stack figures were clearly favourable to hybrid labelling. A detailed analysis can be found in [66].

4 Discussion

The study case accompanying the description of the method just described shows its effectivity. It deals with a sample showing all the problems typical of

a C/C composite synchrotron tomograph, and demonstrates how the method is able to tackle them and results in a detailed description of the microstructure of the sample, either in the form of a phase mask suitable for numerical models working on explicit spatial domains, or up to the point of a detailed individual characterization of each fibre in the sample. Specifically, the heavy ball is not only able to separate the fibres, but it also produces a set of data permitting a whole range of measurements, from simple statistics, such as fibre length and orientation, to the parameterization of fibres as 3D curves, surfaces, or solids, see Figure 11.

With respect to the computational overhead, clearly the heaviest stage is the anisotropic diffusion filtering. Completing the 30 to 40 iterations needed to get results such as those in Figure 5 on the 200^3 sample takes several minutes. However, techniques do exist capable of accelerating significantly anisotropic diffusion [67].

The differential profiler is very fast, the whole procedure in all sections along the three directions, including the correlation stage, takes about 2 seconds.

With respect to the stages involving labelling algorithms, time figures are provided in the previous section, and a detailed analysis of these techniques can be found in [65,66]. If the right choice is made in the selection of the labelling algorithm, labelling times for moderately sized samples are very small. However, it should be noted that, in spite of the high reduction in stack overhead of the hybrid technique, it still uses recursivity, and this could be a source of trouble as sample sizes grow. For the 3D labelling of big samples, it may be necessary to resort to block techniques, where the sample is labelled by blocks [68], or to use an iterative technique.

Regarding the heavy ball, computation time depends on the length and radius of the fibre, but the average fibre, 200 voxel long and 14 voxel in diameter, takes about 60 ms, including dumping the center locations to a text buffer for export. Future implementations of the algorithm can yet be optimized for speed, for instance by avoiding repeated computations in the overlap for candidate moves at each step.

With respect to parametric complexity, tuneable parameters are only found in two of the stages, the preprocessing filtering and the differential profiler. Anisotropic filtering uses three parameters, the noise scale σ in (6), the feature scale κ in (4), and the number of iterations. These usually require careful fine tuning when dealing with applications related to enhancement of images with multiple levels of detail and where the visual appearance of the final result is a major factor. However, 1) in our case the noise scale and the feature scale are well defined and well apart, and 2) we use anisotropic diffusion to provide suitable input for the differential profiler, which has a high degree of robustness and is itself tunable.

The result is that, as long as the noise is sufficiently eliminated and the mesoscale features of the microstructure are well preserved, the sensitivity to the parameters of the diffusion filter is very low. Differences among parameter configurations can be handled by the tuneable thresholds in the differential profiler. We recommend plotting a couple of profiles anywhere through the sample every, say, ten iterations of the anisotropic diffusion filter, so that the progress in the edge enhancement and noise reduction can be tracked, a suitable stopping point for the diffusion can be chosen, and adequate thresholds for the profiler can be set.

Regarding precision, anisotropic diffusion has the interesting property of preserving not only the existence of edges at the feature scale, but also their location. With the differential profiler, on the other hand, it is possible to determine with high accuracy the position of edges, and also to choose the outer side of the features (first threshold crossing point for 'entry" peaks and second for "exit" peaks), the inner side (vice versa), or exactly the edge (maximum between threshold crossing points), see Figure 6.

The only potential source of trouble are brighter cores within the fibres which sometimes may percolate to the outside of the fibre through a broken edge, creating a connection with the matrix material and leaving a gap for the differential profiler. However, correlations among profiles in the three spatial directions decrease their probability, and, moreover, the path of the heavy ball is not likely to be affected by these spots. Nevertheless, in significant numbers they could affect some of the statistics. An inspection of the phase mask will reveal the extent of the problem, if any, and advice about the opportunity of a finer tuning of the profiler thresholds or the anisotropic diffusion before characterization.

Note also that the procedure described for the application of the heavy ball would miss any fibre entirely embedded *inside* the sample. If a low but significant probability of short fibres entirely embedded in the sample is expected, the interior of the fibre mask may be scanned after completing the procedure, looking for fibre voxels that were not tracked by the heavy ball, in order to apply it there. Alternatively, if the material is expected to have a large proportion of short (with respect to sample size) fibres, the heavy ball can be applied as if peeling an onion: first the heavy ball is applied to all six outer boundaries of the sample, then the procedure is repeated a little deeper into

the sample for any fibre section not yet tracked, and so on, until the volume being inspected is too small as to contain any fibres.

Finally, note that the diffusion process in the preprocessing stage involves mixing contributions from neighbouring voxels. This implies the existence of border effects near the image boundaries. Edge detection in the profiler, in turn, uses finite differences involving separated voxels, and it is not able to detect edges if they are just at the boundary of the image. Moreover, in the correlation among profiles, the extreme ends of profiles count with much less information than the rest. All this added up results in the outer boundaries of the resulting mask having a high degree of unaccuracy in the determination of phase type. Therefore, it is a good practice to discard a few layers of voxels around the mask before any further processing.

5 Conclusion

We have detailed and illustrated a multistage methodology able to extract and characterize the true microstructure of a fibrous C/C composite using phase constrast X-ray microtomography. Since C/C composites are probably a worst-case benchmark, the method can also be applied to other cylindrical fibre composites. The processing chain is made of 1) a preprocessing stage where 3D anisotropic diffusion filtering is used to reduce the noise while preserving the features in the sample; 2) a differential profiling method to detect areas between intensity edges on every 2D cross section parallel to the reference system, which are then refined by correlating the outputs in 3D; 3) an optional "cleaning phase" using a 3D labelling algorithm to purify the resulting mask; and 4) a novel fibre separation process which requires a ellipse fitting procedure

applied to the outer boundaries of the 3D mask followed by the use of what we call a heavy ball to individually identify each fibre and characterize it.

We hope to have contributed to demonstrate the ability of phase contrast X-ray tomography to produce sufficient information as to extract the true microstructure of composites with constituents with close densities, such as C/C composites, when appropriate image processing techniques are used.

Acknowledgements

The authors wish to thank Prof. G. L. Vignoles and O. Coindreau (LCTS), Snecma Propulsion Solide, and the ESRF (European Synchrotron Radiation Facility) ID 19 team for providing us with the C/C composite 3D image.

References

[1] Fakirov S, Fakirova C. Direct determination of the orientation of short glass fibers in an injection-molded poly(ethylene terephthalate) system. Polym Compos. 1985;6:41–46.

[2] Bay RS, Tucker III CL. Stereological measurement and error estimates for three-dimensional fibre orientation. Polym Eng Sci. 1992;32(4):240–253.

[3] Fischer G, Eyerer P. Measuring the spatial orientation of short fibre reinforced thermoplastics by image analysis. Polym Compos. 1988;9(4):297–304.

[4] Wetherhold RC, Scott PD. Prediction of thermoelastic properties in short-fiber composites using image analysis techniques. Compos Sci Technol. 1990;37(4):393–410.

[5] Clarke AR, Davidson NC, Archenhold G. Measurements of fibre directions in reinforced polymer composites. J Microsc (Oxford). 1993;171:69–79.

[6] Guild FJ, Summerscales J. Microstructural image analysis applied to fibre composite materials: a review. Composites. 1993;24:383–393.

[7] Paluch B. Analysis of geometric imperfections in unidirectionally-reinforced composites. In: Proceedings of the European Conference on Composite Materials (ECCM6); 1993. p. 305–310.

[8] Mattfeldt T, Clarke AR, Archenhold G. Estimation of the directional distribution of spatial fibre processes using stereology and confocal scanning laser microscopy. J Microsc (Oxford). 1994;173(2):87–101.

[9] Clarke AR, Archenhold G, Davidson NC. A novel technique for determining the 3D spatial distribution of glass fibres in polymer composites. Compos Sci Technol. 1995;55:75–91.

[10] Paluch B. Analysis of geometric imperfections in unidirectional reinforced composites. J Compos Mater. 1996;30(4):454–485.

[11] Davidson NC, Clarke AR, Archenhold G. Large-area, high-resolution image analysis of composite materials. J Microsc (Oxford). 1997;185(2):233–242.

[12] Joyce PJ, Kugler D, Moon TJ. A technique for characterizing process induced fiber waviness in unidirectional composite laminates using optical microscopy. J Compos Mater. 1997;31(17):1694–1734.

[13] Zhu YT, Blumenthal WR, Lowe TC. Determination of non-symmetric 3-D fiber-orientation distribution and average fiber length in short-fiber composites. J Compos Mater. 1997;31(13):1287–1301.

[14] Clarke AR, Archenhold G, Davidson NC. 3D Confocal microscopy of glass fibre reinforced composites. In: Summerscales J, editor. Microstructural

Characterisation of Fibre-Reinforced Composites. Cambridge: Woodhead Publishing; 1998. p. 55–137.

[15] Mlekusch B. Fibre orientation in short-fibre-reinforced thermoplastics II: Quantitative measurements by image analysis. Compos Sci Technol. 1999;59:547–560.

[16] Zak G, Park MHCB, Benhabib B. Estimation of average fibre length in short-fibre composites by a two-section method. Compos Sci Technol. 2000;60:1763–1772.

[17] Eberhardt CN, Clarke AR. Fibre orientation measurements in short glass fibre composites. Part I: automated, high-angular resolution measurement by confocal microscopy. Compos Sci Technol. 2001;61:1389–1400.

[18] Vignoles GL. Image segmentation for hard X-ray phase contrast images of C/C composites. Carbon. 2001;39:167–173.

[19] Lee YH, Lee SW, Youn JR, Chung K, Kang TJ. Characterization of fiber orientation in short fiber reinforced composites with an image processing technique. Mater Res Innov. 2002;6:65–72.

[20] Blanc R, Germain C, Da Costa JP, Baylou P, Cataldi M. Fiber orientation measurements in composite materials. Compos Pt A-Appl Sci Manuf. 2006;37:197–206.

[21] Clarke AR, Eberhardt CN. Microscopy Techniques for Materials Science. Cambridge: Woodhead Publishing; 2002.

[22] Gonzalez RC, Woods RE. Digital Image Processing. 2nd ed. Upper Saddle River, NJ: Prentice Hall; 2002.

[23] Wilson T. Confocal Microscopy. London: Academic Press; 1993.

[24] Russ JC, Dehoff RT. Practical Stereology. 2nd ed. New York: Plenum Press; 2000.

[25] Sterio DC. The unbiased estimation of number and sizes of arbitrary particles using the dissector. J Microsc. 1984;134(2):127–136.

[26] Maisl M, Scherer T, Reiter H, Hirsekorn S. Nondestructive investigation of fibre reinforced composites by X-ray computed tomography. In: Holler P, Hauk V, Dobmann G, Ruud C, Green R, editors. Nondestructive Characterisation of Materials. Berlin: Springer-Verlag; 1988. p. 147–154.

[27] Coindreau O, Vignoles G, Cloetens P. Direct 3D microscale imaging of carbon-carbon composites with computed holotomography. Nucl Instrum Methods Phys Res Sect B-Beam Interact Mater Atoms. 2003;200:308–314.

[28] Herman GT. Image Reconstruction from Projections: The Fundamentals of Computed Tomography. New York: Academic Press; 1980.

[29] Jia X, Gopinathan N, Williams RA, Eberhardt CN, Clarke AR. X-ray microtomography facilitated modelling of microstructures. In: Proceedings 2nd World Congress on Industrial Tomography. Hannover, Germany; 2001. p. 451–460.

[30] Eberhardt CN, Clarke AR. Automated reconstruction of curvilinear fibres from 3D datasets acquired by X-ray microtomography. J Microsc (Oxford). 2002;206(1):41–53.

[31] Buckley JD, Edie DD. Carbon-carbon materials and composites. Parkridge, NJ: William Andrew Publishing/Noyes; 1993.

[32] Christin F. Design, fabrication and application of thermostructural composites like C/C, C/SiC and SiC/SiC composites. Adv Eng Mater. 2002;4(12):903–912.

[33] Coindreau O. Etude 3D des preformes fibreuses: Interaction entre phenomenes

physico-chimiques et geometrie [PhD Thesis]. Université Bordeaux1. France; 2003.

[34] Coindreau O, Vignoles GL. Assessment of structural and transport properties in fibrous C/C composite preforms as digitized by X-ray CMT. Part I: Image acquisition and geometrical properties. J Mater Res. 2005;20:2328–2339.

[35] Velhinho A, Sequeira PD, Martins R, Vignoles G, Fernandes FB, Botas JD, et al. X-ray tomographic imaging of Al/SiC_p functionally graded composites fabricated by centrifugal casting. Nucl Instrum Meth B. 2003;200:295–302.

[36] Velhinho A, Sequeira PD, Martins R, Vignoles G, Fernandes FB, Botas JD, et al. Evaluation of Al-SiC wetting characteristics in functionally graded metal-matrix composites by synchrotron radiation microtomography. Mater Sci Forum. 2003;423–425:263–268.

[37] Bickerdike R, Brown A, Huges G, Ranson H. The deposition of pyrolytic carbon in the pores of bonded and unbonded carbon powders. In: Proceedings of the 5th AMCS Biennial Conference on Carbon; 1962. p. 575–582.

[38] Besmann TM. Processing science for chemical vapor infiltration. Oak Ridge: Oak Ridge National Laboratory; 1990.

[39] Jin S, Wang X, Starr TL, Chen X. Robust numerical simulation of porosity evolution in Chemical Vapor Infiltration I: Two space dimension. J Comp Phys. 2000;162:467–482.

[40] Jin S, Wang X. Robust numerical simulation of porosity evolution in Chemical Vapor Infiltration II. Two-dimensional anisotropic fronts. J Comp Phys. 2002;179:557–577.

[41] Jin S, Wang X. Robust numerical simulation of porosity evolution in Chemical Vapor Infiltration III: Three space dimension. J Comp Phys. 2003;186:582–595.

[42] Wertheimer M. Untersuchungen zur Lehre von der Gestalt II. Psycol Forsch. 1923;4:301–350. Translation published as Wertheimer M. Laws of Organization in Perceptual Forms. In: Ellis W, editor. A source book of Gestalt psychology. London: Routledge & Kegan Paul; 1938. p. 71–88.

[43] Marr D. Vision. San Francisco, CA: Freeman Publishers; 1982.

[44] DeValois RL, DeValois KK. Spatial Vision. Oxford: Oxford University Press; 1990.

[45] Perona P, Malik J. Scale-space and edge detection using anisotropic diffusion. IEEE Trans Pattern Anal Mach Intell. 1990;12:629–639.

[46] Weickert J. Anisotropic Diffusion in Image Processing. Stuttgart, Germany: Teubner; 1998.

[47] Weickert J. Theoretical foundations of anisotropic diffusion in image processing. In: Kropatsch W, Klette R, Solina F, editors. Theoretical Foundations of Computer Vision. vol. 11 of Computing Supplement. Wien: Springer-Verlag; 1996. p. 221–236.

[48] Weickert J. Coherence-enhancing diffusion filtering. Int J Comput Vis. 1999;31(2/3):111–127.

[49] Frangakis AS, Hegerl R. Nonlinear anisotropic diffusion in three-dimensional electron microscopy. In: Scale-Space Theories in Computer Vision. vol. 1682 of Lecture Notes in Computer Science. Springer-Verlag; 1999. p. 386–397.

[50] Tschumperlé D, Deriche R. Vector-Valued Image Regularization with PDEs: A Common Framework for Different Applications. IEEE Trans Pattern Anal Mach Intell. 2005 April;27(4):1–12.

[51] Black MJ, Sapiro G, Marimont DH, Heeger D. Robust anisotropic diffusion. IEEE Trans Image Process. 1998;7:421–432.

[52] Nitzberg M, Shiota T. Nonlinear image filtering with edge and corner enhancement. IEEE Trans Pattern Anal Mach Intell. 1992;14:826–833.

[53] Catté F, Lions PL, Morel JM, Coll T. Image selective smoothing and edge detection by nonlinear diffusion. SIAM J Numer Anal. 1992;29:182–193.

[54] You YL, Xu W, Tannenbaum A, Kaveh M. Behavioral analysis of anisotropic diffusion in image processing. IEEE Trans Image Process. 1996;5:1539–1553.

[55] Kichenassamy M. The Perona-Malik paradox. SIAM J Appl Math. 1997;57:1328–1342.

[56] Weickert J, Benhamouda B. A semidiscrete nonlinear scale-space theory and its relation to the Perona-Malik paradox. In: Solina F, Kropatsch W, Klette R, Bajcsy R, editors. Advances in Computer Vision. Wien: Springer-Verlag; 1997. p. 1–10.

[57] Martín-Herrero J, Alba-Castro JL. High speed machine vision: The canned tuna case. In: Billingsley J, editor. Mechatronics and Machine Vision in Practice: Future Trends. Hertfordshire, UK: Research Studies Press; 2003. p. 177–186.

[58] Clarke AR, Eberhardt CN. An automated fibre orientation measurement design. In: Microscopy Techniques for Materials Science. Cambridge: Woodhead Publishing; 2002. p. 202–207.

[59] Ahn SJ, Rauh W, Warnecke HJ. Least-square orthogonal distances fitting of circle, sphere, ellipse, hyperbola, and parabola. Patt Recogn. 2001;34:2283–2303.

[60] Brice CR, Fennema CL. Scene analysis using regions. Artif Intell. 1970;1(3-4):205–226.

[61] Luger GF, Stubblefield WA. Artificial Intelligence: Structures and Strategies for Complex Problem Solving. 2nd ed. New York: Chapman Hall; 1993.

[62] Rosenfeld A, Pfaltz J. Sequential operations in digital picture processing. J ACM. 1966;13:471–494.

[63] Dillencourt MB, Samet H, Tamminen M. A general approach to connected-component labeling for arbitrary image representations. J ACM. 1992;39:253–280.

[64] Tarjan RE. Efficiency of a good but not linear set union algorithm. J ACM. 1975;22(2):215–225.

[65] Martín-Herrero J. Hybrid cluster identification. J Phys A: Math Gen. 2004;37:9377–9386.

[66] Martín-Herrero J. Hybrid object labelling in digital images. Mach Vis Appl. 2007;(in press).

[67] Weickert J, ter Haar Romeny BM, Viergever MA. Efficient and reliable schemes for nonlinear diffusion filtering. IEEE Trans Image Process. 1998 March;7(3):398–410.

[68] Martín-Herrero J, Peón-Fernández J. Alternative techniques for cluster labelling on percolation theory. J Phys A: Math Gen. 2000;33:1827–1840.