

Algorithms for automated meshing and unit cell analysis of periodic composites with hierarchical tri-quadratic tetrahedral elements

Hyung Joo Kim¹ and Colby C. Swan^{2,*},[†]

¹*Department of Mechanical Engineering, Center for Computer-Aided Design, The University of Iowa, Iowa City, IA 52242, U.S.A.*

²*Department of Civil and Environmental Engineering, Center for Computer-Aided Design, The University of Iowa, Iowa City, IA 52242, U.S.A.*

SUMMARY

Unit cell homogenization techniques together with the finite element method are very effective for computing equivalent mechanical properties of composites and heterogeneous materials systems. For systems with very complicated material arrangements, traditional, manual mesh generation can be a considerable obstacle to usage of these techniques. This problem is addressed here by developing automated meshing techniques that start from a hierarchical quad-tree (in 2D) or oc-tree (in 3D) mesh of pixel or voxel elements. From the pixel/voxel mesh, algorithms are presented for successive element splitting and nodal shifting to arrive at final meshes that accurately capture both material arrangements and constituent volume fractions, and the material-scale stress and strain fields within the composite under different modalities of loading. The performance and associated convergence behaviour of the proposed techniques are demonstrated on both densely packed fibre and particulate composites, and on 3D textile-reinforced composites. Copyright © 2003 John Wiley & Sons, Ltd.

KEY WORDS: homogenization; periodic media; mesh refinement; effective properties; textile composites; particulate composites

1. INTRODUCTION

1.1. The challenge

Unit cell computational homogenization methods for heterogeneous periodic media typically involve constructing a continuum model of a composite's heterogeneous material structure

*Correspondence to: Colby C. Swan, Department of Civil and Environmental Engineering, Center for Computer-Aided Design, The University of Iowa, Iowa City, IA 52242, U.S.A.

[†]E-mail: colby-swan@uiowa.edu

Contract/grant sponsor: National Science Foundation; contract/grant number: CMS-9713335, CMS-9896284

Contract/grant sponsor: Whitaker Foundation; contract/grant number: 96-0636

Received 29 August 2002

Revised 24 September 2002

Accepted 27 January 2003

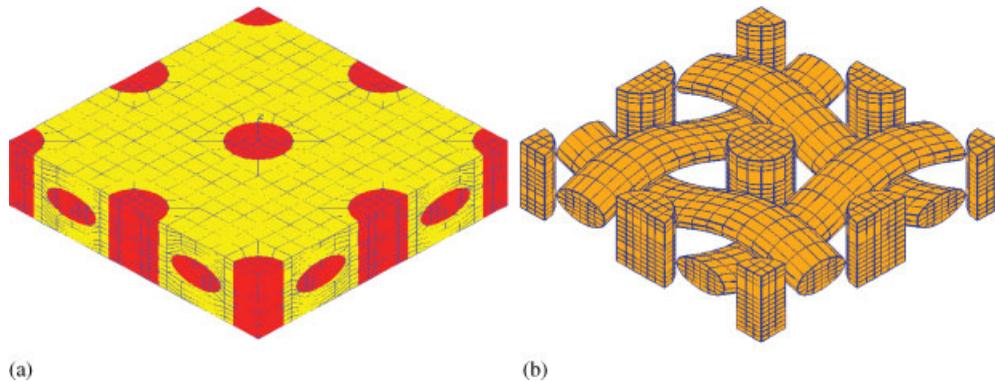


Figure 1. Traditional structured mesh finite element unit cell model of plain-weave textile with orthogonal inserts.

(Figure 1), and then applying loads to the model with imposed periodic boundary conditions [1–3]. The same basic methods apply equally well to aperiodic heterogeneous media, with the exception that one must generally work with stochastic realizations of the representative volume element (r.v.e.) rather than a well-defined unit cell (see, for example, Reference [4]). While these types of computational homogenization methods for composites are proving increasingly effective there remain many very practical challenges to be addressed. One is that the model of a composite's unit cell (or r.v.e.) must adequately capture material arrangements and constituent volume fractions as well as the local stress/strain fields in the unit cell or r.v.e. under a variety of loading conditions. Errors in both resolution of material volume fractions and material-scale stress and strain fields can lead to appreciable error in computing the effective properties of composites (see, for example, Reference [5]) which means that both types of errors require due consideration.

These requirements give rise to two inherent challenges associated with finite element-based computational homogenization:

- The computational cost (both required memory and the number of CPU operations) when dealing with three-dimensional composite material systems can be considerable, thus requiring usage of rapid and efficient computing techniques and solvers [6];
- The development of suitable material-scale finite element models that capture material arrangements, interface conditions, and stress/strain fields can be painstaking to create, and can require considerable amounts of the analyst's time whenever a new material arrangement or fibre volume fraction is to be considered.

The objective of this paper is to deal with second issue in a way that facilitates extending the realm of material systems to which computational homogenization can be readily and easily applied. In a preceding work [7], the authors extended the works of Kikuchi and Hollister [8–10] and Tareda *et al.* [11] developing and testing hierarchical pixel/voxel-based meshing techniques and applying them to textile composites. Here the hierarchical pixel/voxel meshing framework for periodic composites is further extended to achieve unit cell models of quadratic triangles (2D) or tetrahedra (3D) using combinations of element splitting and nodal shifting.

Since the proposed element splitting techniques are here presented in the context of unit cell analysis of periodic media, the specific meshing concerns associated with application of periodic boundary conditions to the unit cell model are addressed. The essential meshing and model development methods apply equally well, however, to aperiodic media as well.

Two of the most significant and challenging aspects of constructing meshes for general unit cell analysis of periodic composites are those associated with achieving full compatibility between the matrix and reinforcement phases (*the two-way meshing problem*), and that of applying periodic boundary conditions to the unit cell (*the external face-matching problem*). To highlight the general two-way meshing problem, it is helpful to consider the plane weave composite [12] with orthogonal inserts whose unit cell is shown in Figure 1. Meshing of the individual yarns in this unit cell model is actually quite simple, starting with a circular cylinder for each yarn and then applying a sequence of distortion, translation, and rotation transformations. Here this (Figure 1(b)) is called the *a priori meshing* of the individual yarns. On the other hand, meshing of the polymer matrix phase between the individual yarns in a way that achieves elements with favourable aspect ratios and yields full displacement compatibility between the yarn and matrix phases is much more challenging. While a number of ad-hoc and specific meshing algorithms can be developed to construct the polymer infill for specific textile patterns, such algorithms typically need to be re-developed (at considerable time cost) each time a new textile pattern is considered. Unstructured meshing techniques like Delaunay triangulation with advancing front techniques could be used to mesh the domain external to a given yarn, using the surface discretization of that yarn. However, as the front begins to approach other yarns or objects, the resulting fine mesh of nodes and elements will not match up with the *a priori* surface discretization of the yarns resulting in incompatibility of the displacement field at the yarn–matrix interface. The problem is that advancing front techniques only mesh one side of a given surface, whereas for composite materials, it is necessary that both sides of the surface be meshed in a way that achieves full compatibility. In this sense, two-sided meshing techniques are needed.

Within the context of a standard FEM framework, attention is devoted here to solving the two-way meshing problem for unit cell or r.v.e. models containing internal surfaces of discontinuity. It is worth noting briefly that recent and alternative approaches to the same problem have treated it less as a meshing issue by focusing primarily on enrichment of the FEM basis functions in the vicinity of the discontinuities. A few specific examples are those of Strouboulis *et al.* [13] who developed a generalized finite element method (GFEM) and the works of Belytschko *et al.* [14–16] who developed an extended finite element method (XFEM) for treating general discontinuities. In these approaches, the approximation functions in the FEM model are enriched using partition of unity concepts to incorporate local analytical solution characteristics in the vicinity of inclusions or discontinuities. While the meshing methods developed herein can be considered local ‘h’ refinement methods, those just cited correspond more closely to local ‘p’ refinement methods.

In unit cell analysis of periodic media, the shape of the unit cell domain is typically a parallelepiped with three pairs of external matching faces [3]. Since the unit cell repeats itself indefinitely to form the periodic medium, each unit cell will have six neighbouring unit cells with which it shares common faces. In the most general of cases, one of the most straightforward ways to achieve both periodicity and continuity of the displacement field between adjacent unit cells is to require that the mesh of the unit cell have identical surface discretizations for both surfaces of each surface pair. This insures a one-to-one correspondence

of all nodes lying on the external surfaces with the nodes on the matching pair surface. The issues associated with achieving mesh discretizations that satisfy these criteria are those related to the so-called *external face-matching problem*.

1.2. Prior works on automated meshing

A variety of what are now somewhat conventional approaches to capturing material boundaries with triangular (2D) and tetrahedral (3D) elements have been successfully developed and used over the past decade [17]. The classes of methods being reviewed here can be separated into two categories, those based on (1) structured meshes and (2) non-structured meshes. Creation of structured meshes based on transformations and mappings usually proceeds by creating a mesh with regular polygons (2D) or regular polyhedra (3D), and then mapping the meshed object into one of the same topology, but a different shape [18, 19]. Creation of the individual yarns of the textile composite unit cell in Figure 1(b) is an example of using mappings and transformations on a number of different objects.

In another structured approach, regular grids of nodes and cells are constructed in space, and the object to be meshed is inserted into the same space. Those elements lying completely outside of the space occupied by the object to be meshed are removed, while those cells containing boundaries of the object are truncated in a manner that captures the boundary. The recognized characteristics of the grid-base approach are that: (1) they are one-sided; (2) the finer the grid, the better the resulting mesh will be, since the proportion of internal elements that are well-shaped will be higher; and (3) object features small compared to the grid spacing size 'h' are lost in the meshing process. To facilitate meshing of objects with fine features, hierarchical grids using quadtree and octree multi-resolution approaches [20] have been developed as extensions of grid-based approaches. The hierarchical voxel-pixel-based meshing techniques proposed and developed in Reference [7] are two-sided hierarchical structured grid techniques.

In unstructured meshing, so-called Advancing Front Techniques [21, 22] based on Delaunay Triangulation [23, 24] are quite common. The first problem to be addressed with advancing front techniques is constructing a description of the meshing domain's boundary in terms of discrete data segments such as line segments (2D) or triangles (3D). In the data segments, the nodal sequence must be carefully chosen to distinguish between the interior and exterior of the meshing domain. Starting from the exterior boundaries and working inward, the advancing front techniques then fill the interior domain with triangles or tetrahedrons using Delaunay triangulation. In such approaches, the domain being meshed is typically defined by its boundaries, and only the domain on one-side of the boundary is actually meshed. Since both sides of material boundaries must typically be meshed for micro-mechanical analysis of composite materials, the one-sided methods cited above do not appear well suited for analysis of composite materials.

A major issue with triangulation techniques is ensuring that individual elements have quality shapes. Tests on Jacobian determinants have been essential to identify poorly performing elements, but when dealing with higher order displacement fields, the numerical performance of the mesh is also very sensitive to obtuse angles as shown in Reference [25]. Common quality indicators based on element geometry, and frequently used in pre- and post-processing mesh improvement schemes have thus been developed [24, 26]. Despite the good properties of triangulation techniques, they still require much user preprocessing both to determine material

boundaries and to adjust the mesh size. Attempts to extend triangulation techniques to 3D modelling are still ongoing with one of the major issues being the degeneration of element shapes and aspect ratios under repetitive local mesh refinements [27–29].

1.3. Proposed new methods

Owing to the concerns raised above with existing automated meshing techniques, voxel-based-meshing techniques were considered [7] due to their many attractive characteristics, including two-sidedness and the ability to achieve matching discretizations of external faces. Despite these positive attributes, voxel-based techniques, since they are non-conforming and two-sided, necessarily use elements that generally contain more than one material to resolve material interfaces. Owing to the uncertainty on how to prescribe material properties for these heterogeneous elements containing more than one material, the convergence characteristics of voxel-based meshes are somewhat slower than might be achieved with conforming meshes that employ only homogeneous elements, with each containing only a single material. In this work, we therefore attempt to combine the positive attributes of voxel-based meshing with the positive attributes of conforming meshes.

The proposed methods make usage of appropriate material domain identification functions (MIDFs). Once such functions are developed, a hierarchical pixel or voxel mesh of any desired resolution can be developed using the techniques described in Reference [7]. Since the pixels or voxels in the vicinity of material interfaces will contain a mixture of materials, it is desirable to subdivide them into homogeneous triangles or tetrahedra. This is done in two stages. In the first, all pixels or voxels are subdivided into triangles or tetrahedra and the resulting mesh is called the base triangle mesh or the base tetrahedral mesh. The prefix base- is used to indicate that base-triangles or tetrahedra cannot necessarily be used in the final mesh since each could contain more than one material. In the second stage, those base triangles or tetrahedra that are heterogeneous are modified using both nodal shifting and element splitting techniques to achieve a mesh of homogeneous triangles or tetrahedra.

The rationale behind using a background mesh of hierarchical pixels or voxels is that such geometric objects can be successively refined any number of times without any degradation of their shapes or aspect ratios. Though there are several well-known ways to subdivide triangles and tetrahedrons into finer ones while preserving the original favourable shape and aspect ratio, the reliance on and successive refinement of a background mesh of hierarchical pixels or voxels is used to accommodate data hierarchy. Consequently, instead of refining the triangular or tetrahedral mesh directly, the background pixel or voxel mesh is refined using hierarchical quad-tree or oc-tree data structures and then the refined pixels or voxels are subdivided again into triangles or tetrahedra. Note that this approach and the resulting data hierarchy facilitate the realization of unit cell meshes with external face-matching properties.

In the proposed meshing techniques, the concept of MDIFs play a central role in process automation, since they hold all the information describing the arrangement of materials in the composite. Without MDIFs, most unstructured meshing techniques require discrete data segments [such as *least-qualified elements* with Delaunay triangulation; *exterior meshes* with advancing front methods] to characterize the material boundaries. Such discrete data segments must be updated and augmented as mesh refinement proceeds. An additional benefit of MDIFs is that they enable resolution of position-dependent anisotropy. For example, in the case of textile composites, the yarns are essentially uni-directionally aligned fibre–matrix composites

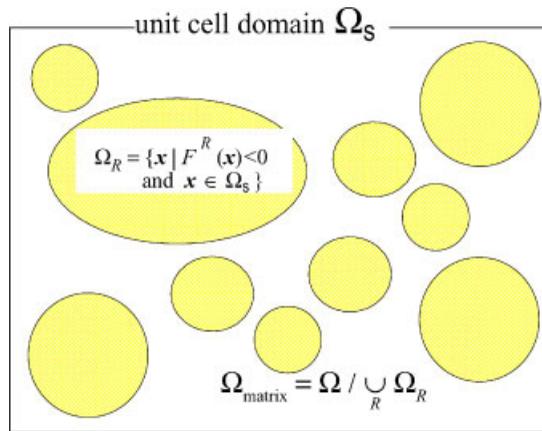


Figure 2. Definition of MDIFs.

that feature strong transverse isotropy with much higher stiffness and strength in the longitudinal direction than in transverse directions. Moreover, since individual yarns are warped to create an interwoven reinforcement structure, the axis of transverse isotropy changes smoothly with distance along a given yarn, which makes anisotropy position-dependent. Note that when the MDIF is of analytic form it can provide the axial-direction of transverse anisotropy at any point in yarn's domain.

In this work, final meshes are generated using bi-quadratic triangular elements and tri-quadratic tetrahedral elements. There are numerous advantages to using such elements over linear triangles and tetrahedra. First the numerical performance of such elements is superior to that of linear elements, since they are less prone to mesh-locking phenomena [30]. Secondly, quadratic elements can potentially facilitate the resolution of highly curved material interfaces, thus resulting in reduced meshing errors associated with volume fraction truncation.

2. DESCRIPTION OF NEW MESHING TECHNIQUES

2.1. The material domain identification function

Most composites are typically fabricated with continuous and/or discrete reinforcing fibres or particles embedded in a continuous matrix material. In a three-dimensional spatial domain, the material region occupied by the R -th reinforcing fibre or particle can be described mathematically either with a single function $F^R(\mathbf{X}) : \Omega_s \mapsto \mathfrak{R}$, or a set of different functions as follows:

$$\Omega^R = \{ \mathbf{X} \in \Omega^S \mid F^R(\mathbf{X}) < 0 \} \quad (1)$$

These mathematical representations permit one to test whether any given material point lies inside of a specific reinforcing object R (Figure 2) or outside of that object. Specifically, a material point \mathbf{X} lies inside of the R -th reinforcing object when $F^R(\mathbf{X}) < 0$ and either outside

of the yarn or on the boundary otherwise. Furthermore, at least one material interface always lies between two points that have function values of different sign.

In general, a composite material's unit cell will contain a finite number of N_R different reinforcing objects. Mathematically, the functions F^R describing each of these objects can be combined to yield a function that takes material points in the unit cell as arguments and returns the value of the reinforcing object in which that point lies. Such functions are here called MDIFs, and have the following typical form.

$$M(\mathbf{X}) = \begin{cases} J & \text{iff } F^J(\mathbf{X}) < 0 \text{ and } J \in \{1, 2, 3, \dots, N^R\} \\ 0 & \text{iff } F^J(\mathbf{X}) \geq 0 \quad \forall J \in \{1, 2, 3, \dots, N^R\} \end{cases} \quad (2)$$

In this way, whenever the MDIF returns a value of zero, the argument point is taken to lie in the matrix phase, and when the MDIF returns a positive value, it lies in one of the composite's reinforcing objects. An underlying assumption here is that the material description functions of Equation (1) do not intersect.

Although in the current work it is assumed that analytical expressions for the MDIFs are available, there are cases where analytical MDIFs are not available but rather must be derived from CT-scans of the heterogeneous domains to be meshed. In such cases, the MDIFs would need to be derived from material density values provided at discrete grid points [9–11] that are independent of the model to be developed. If the material density ρ throughout the domain were interpolated using suitable basis functions associated with the grid points, then material interfaces could be identified as those surfaces across which material density crosses defined threshold values. Accordingly, a given spatial point \mathbf{X} in the model domain could be identified with a specific material J if the density at \mathbf{X} corresponds to that of material J . That is,

$$M(\mathbf{X}) = \begin{cases} J & \text{if } \rho(\mathbf{X}) \in [\rho_J - \frac{1}{2}\delta_J, \rho_J + \frac{1}{2}\delta_J) \text{ and } J \in \{1, 2, \dots, N^R\} \\ 0 & \text{otherwise} \end{cases} \quad (3)$$

where ρ_J denotes the mean density of the J -th material in the domain, and δ_J denotes a range of density centred about ρ_J to be associated with material J .

In the mesh of a unit cell model, the fraction of each element that is occupied by a given reinforcing object can be determined to any desired level of precision by sampling material points from that element either randomly and evaluating Equation (1). Thus for each finite element comprising the composite's unit cell, one can readily determine a volume fraction associated with the reinforcing object of interest. In the limit of infinite mesh resolution of the unit cell ($N \rightarrow \infty$) the reinforcing object volume fraction associated with each element tends to discrete values of 0 and 1. However, for practical mesh resolutions, many of the elements will typically contain intermediate reinforcement volume fractions in the range $\phi_R \in (0, 1)$.

Another important issue when analysing certain classes of materials such as textile composites is that of material hierarchy. A composite material is made of one or more distinct materials, and one or more of these materials can themselves be composites whose effective properties are anisotropic. In the specific case of textile composites, the yarns are essentially uni-directionally aligned fibre–matrix composites that feature strong transverse isotropy with much higher stiffness and strength in the longitudinal direction than in transverse directions. Moreover, since individual yarns are warped to create an interwoven reinforcement structure, the axis of transverse isotropy changes smoothly with distance along a given yarn. When an

anisotropic material phase occupies a given element e , the orientation of the material phase m in that element is sampled at the centroid $\bar{\mathbf{X}}_e$ of the element. The director $d(\bar{\mathbf{X}}_e)$ for axisymmetry of a transversely isotropic material at centroid $\bar{\mathbf{X}}_e$ is a unit vector aligned with that material's local axial direction.

The textile composite unit cell model (Figure 1) is composed of four woven yarns and nine orthogonal inserts. It is here assumed that each yarn was initially a straight cylinder with an elliptical cross-section, but then deformed so that the centreline curve is sinusoidal and so that cross-sections remain orthogonal to the original axis of the cylinder (Figure 3). In addition, the centreline curve of the yarn lies on a user-defined, $\xi\eta$ -plane defined by three points A, B and C . Points A and B are the two end points of the yarn's centreline axis and C is an auxiliary point that determines the plane. The mathematical description of each woven yarn in local co-ordinates can be written as

$$\begin{aligned} \Omega^R &= \{\mathbf{X} \mid F^R(\mathbf{X}) < 0\} \\ &= \{|\xi| \hat{F}^R(\xi) < 0\} \end{aligned} \tag{4a}$$

where

$$\hat{F}^R(\xi) = \left(\frac{\eta - C_L(\xi)}{r_\eta} \right)^2 + \left(\frac{\zeta}{r_\zeta} \right)^2 \text{ and centreline } C_L(\xi) = l_\eta \sin \left(\frac{2\pi\xi}{l_\xi} \right) \tag{4b}$$

In the preceding, r_η, r_ζ are radii of the yarn's elliptical cross-section in local directions η, ζ , respectively, and l_ξ, l_η are the amplitude and wavelength of the sinusoidal curve $C_L(\xi)$ that defines the centreline of woven yarn object in local co-ordinates $\xi - \eta$ (Figure 3).

For each yarn, the transformation matrix \mathbf{Q} between XYZ -co-ordinates and $\xi\eta\zeta$ -co-ordinates is

$$\xi = \mathbf{Q} \cdot (\mathbf{X} - \mathbf{X}_A) \tag{5}$$

Since yarns are assumed to have transversely isotropic material properties, only the material director in the fibre direction, specified with respect to the local $\xi\eta\zeta$ co-ordinate system, is needed. The fibre director is assumed parallel to the yarn's centreline curve and is constant over each elliptic cross-section perpendicular to the ξ -axis. Accordingly, the material director of a sample point (ξ, η, ζ) in $\xi\eta\zeta$ -co-ordinates can be expressed as

$$\hat{\mathbf{d}}(\xi) = \frac{1}{\sqrt{1 + \left(\left(\frac{2\pi l_\eta}{l_\xi} \right) \cos \left(\frac{2\pi\xi}{l_\xi} \right) \right)^2}} \cdot \left(1, \left(\frac{2\pi l_\eta}{l_\xi} \right) \cos \left(\frac{2\pi\xi}{l_\xi} \right), 0 \right) \tag{6}$$

The material director is then transformed into the XYZ -co-ordinate system as follows:

$$\mathbf{d}(\mathbf{X}) = \mathbf{Q}^T \cdot \hat{\mathbf{d}}(\xi) \tag{7}$$

2.2. Construction of the background pixel/voxel mesh

In this work, it is assumed that a given unit cell model can initially be meshed with hierarchical pixel (in 2D) or voxel (in 3D) meshing techniques such as those described in Reference [7]. The hierarchy of the mesh is assumed to be described in an appropriate tree structure

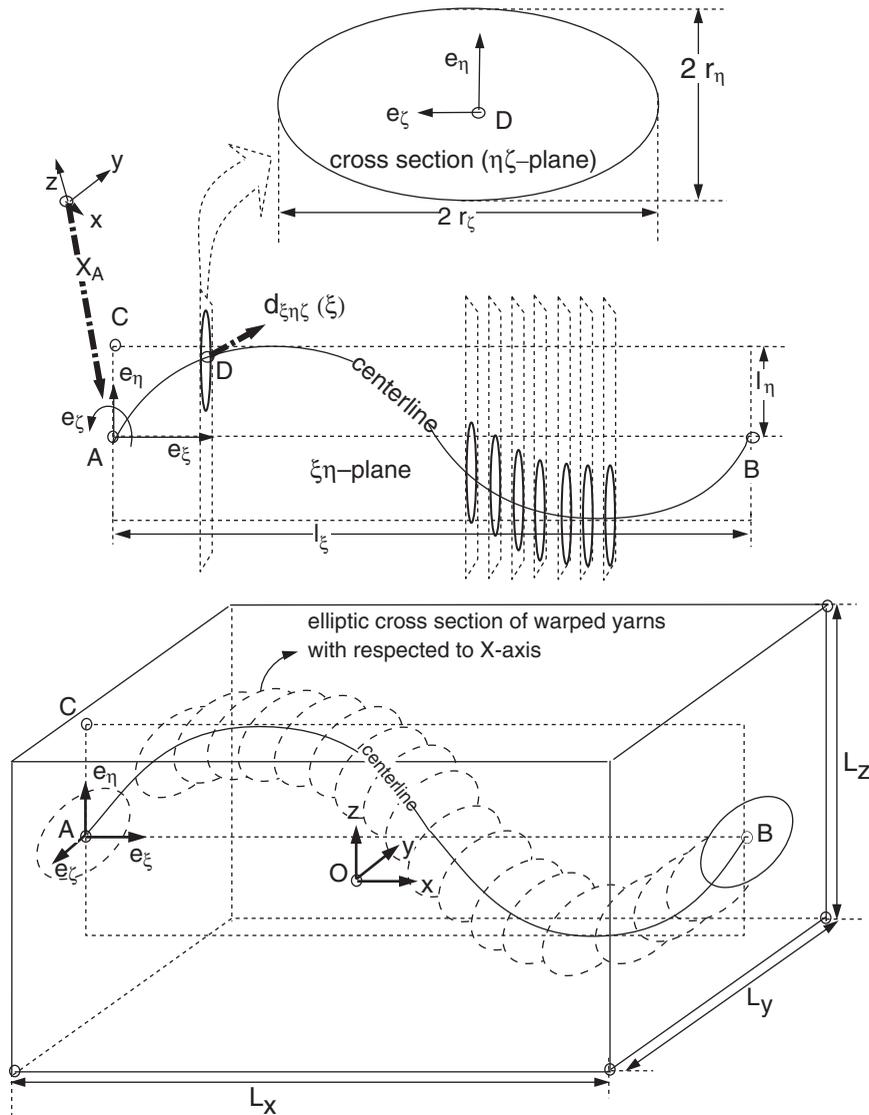


Figure 3. Description of a specific MDIF for warped textile yarns.

where each element has either children or parents, or both. If it is assumed that N^R distinct reinforcing objects with appropriate MIDFs comprise the unit cell model, then each pixel/voxel in the model will have N^R independent volume fraction values. If, for a given element, all of these values are either zero or unity, and their cumulative sum is also zero or unity, then that element will be said to be homogeneous in the sense that it is completely occupied by a single reinforcing object. Elements not satisfying this criterion are said to be heterogeneous. The pixel/voxel mesh can be successively refined uniformly using bisection techniques, until

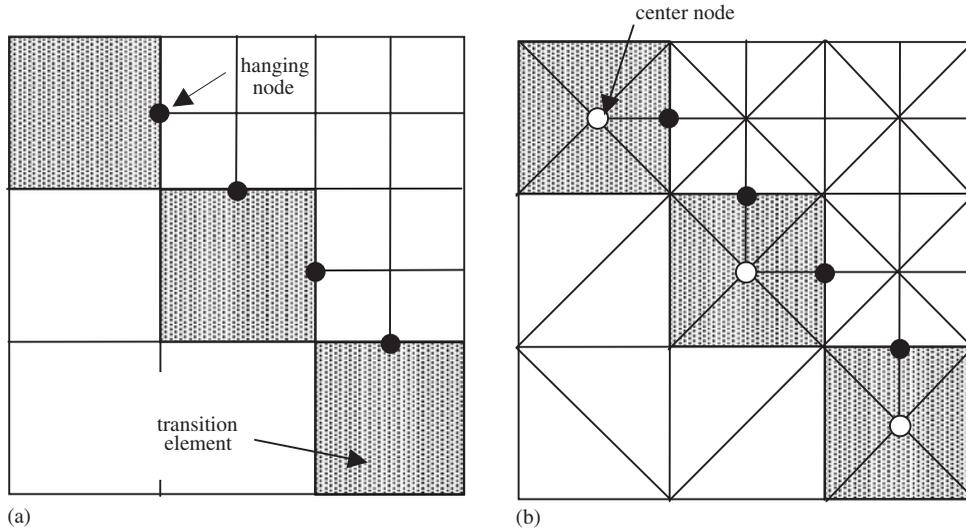


Figure 4. Splitting patterns for a pixel mesh into a base-triangle mesh: (a) pixel mesh; and (b) base-triangle mesh. Pixels with hanging nodes are called transition-pixels (shaded) and are given additional centre nodes to facilitate splitting.

most all pixels/voxels are sufficiently homogeneous. Alternatively, the mesh can be selectively refined, where only the heterogeneous elements and perhaps their neighbours, are successively bisected until a sufficient degree of interface resolution (and element homogeneity) is achieved.

With non-conforming pixel/voxel meshing techniques, the number of refinement cycles (in which some or all pixels/voxels are bisected) needed to clearly defined material interfaces can be very significant, resulting in FEM models that are very large. Alternatively, after a few cycles of refinement, it may be possible to split each pixel/voxel into triangles and tetrahedra that with a bit of additional manipulation can be made materially homogeneous.

2.3. Construction of base-triangular/tetrahedral elements

The fundamental idea at this stage of meshing is to split all pixels into triangles and all voxels into tetrahedra without regard to the location of material interfaces. In two dimensions, individual four-noded pixels without hanging nodes can simply be diagonally bisected with no additional concerns about displacement field continuity between adjoining triangles, and if the original pixels are square, then the triangles resulting from splitting will have aspect ratios of $\sqrt{2}$, where here aspect ratio is defined as the ratio of maximum to minimum edge length. When the pixels have hanging nodes (Figure 4), however, then it necessary to add a new central node. Triangular elements are then constructed out of the new centre node and all edge segments comprising the pixel's boundary. The essential techniques used in subdivision of pixels into triangles generalize directly to three-dimensions where voxels and their quadrilateral faces are split into tetrahedra.

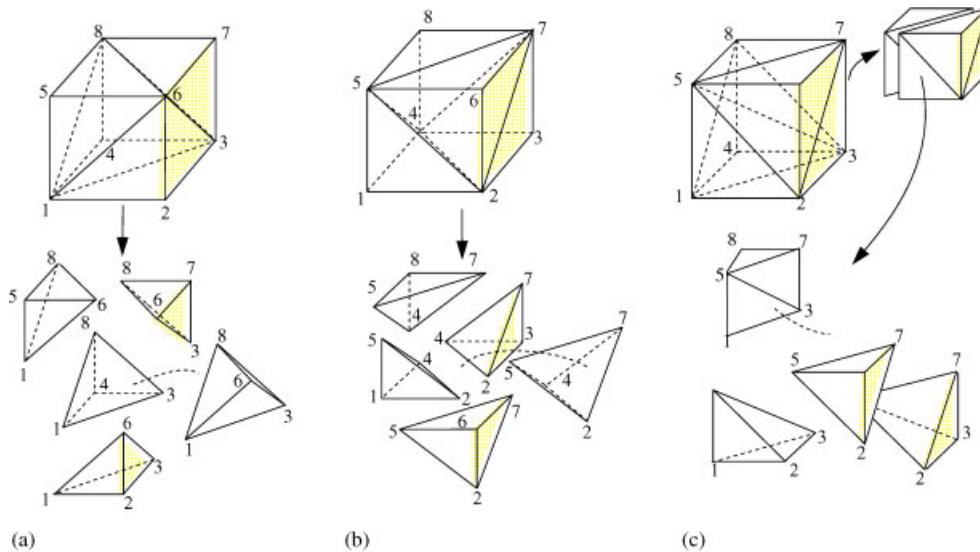


Figure 5. Patterns for splitting a voxel into base-tetrahedra without introducing any new nodes: (a) Hex5-even; (b) Hex5-odd; and (c) Hex6.

In three-dimensions, achieving compatibility of the displacement field or matching of face nodes between elements requires some caution. As shown in Figure 5 there are at least three techniques for splitting a voxel into tetrahedra without introducing any new nodes. The splitting pattern called HEX6 has the good property of nodal face-matching between neighbouring elements in that cutting edges of the faces on opposing faces of the voxel are always aligned. The HEX5-even and HEX5-odd splitting patterns are alternatively applied such that the diagonal edges cutting opposing faces of the voxel are never aligned, producing possible mis-matching of nodes and displacement fields between two elements descending from two adjacent voxels. Alternatively, if an alternating HEX5 even-odd splitting scheme is used, then the problem of face-matching can be solved. In the present work, HEX5 even-odd is employed as the basic splitting pattern for creating base-tetrahedral from voxels.

When cube-shaped voxels are split into tetrahedra using HEX5, the maximum aspect ratio of the tetrahedra is $\sqrt{2}$ whereas the maximum aspect ratio with HEX6 is $\sqrt{3}$. Furthermore, the alternative diagonal cutting directions with HEX5 accommodate face-matching within the mesh when voxels with hanging nodes are subdivided into tetrahedra.

When a voxel element with a hanging node on any of its faces or edges is to be subdivided into tetrahedra, then a body-centred node is first added to that voxel (Figure 6). The new central node and the six faces of the voxel then form six pyramidal polyhedra. The subdivision of the faces of the voxel (or the bases of each pyramid) then proceeds just as the subdivision of pixels above. Those faces with hanging nodes along any of their edges have a face-centred node added, and are then subdivided into triangles. The triangles of each voxel face, and the body-centred node of the voxel are then grouped to divide the voxel into a number of tetrahedra.

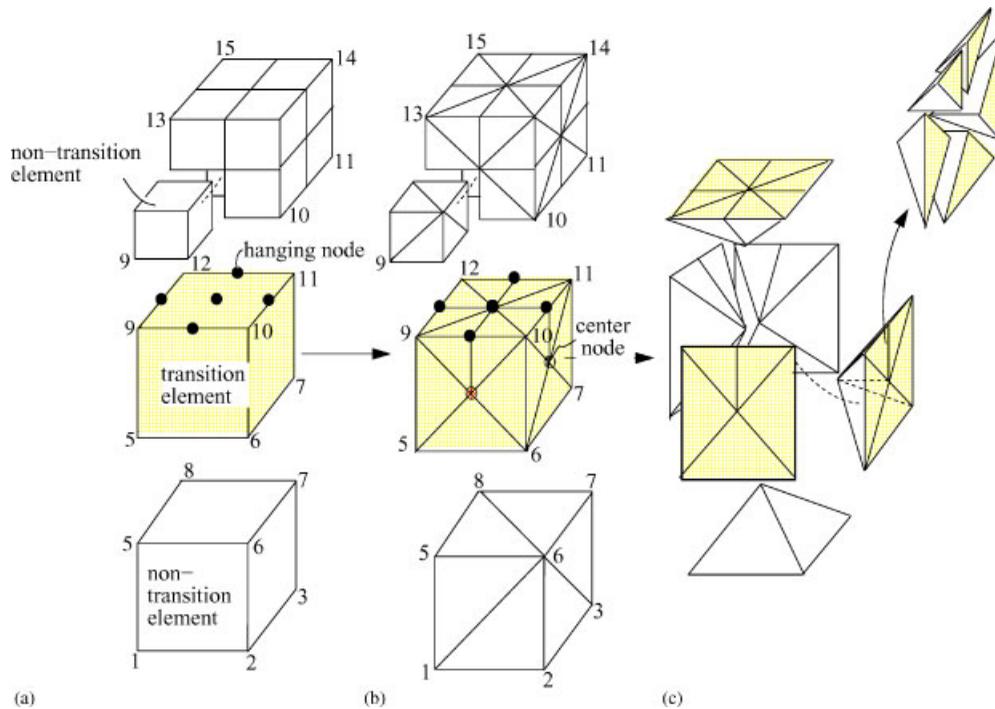


Figure 6. Splitting patterns for transition voxels with hanging nodes: (a) voxel mesh with a transition element having hanging nodes; (b) addition of a body-centred node; and (c) splitting of transition voxel into tetrahedral.

It is worth mentioning that once the base triangles and base tetrahedra are formed, the meshing data structure employed records the parent pixel or voxel from which each has descended. Consequently, if it is subsequently determined that the mesh of base triangles or tetrahedra is still not sufficiently refined to capture the finer details of material arrangements, the meshing algorithm resorts back to further refinement of the background pixel/voxel mesh. In this case, information from specific triangles and/or tetrahedra can be sent back to their parent pixels or voxels for selective mesh refinement of the background mesh.

2.4. Node-shifting and splitting of base triangles and tetrahedra

2.4.1. Nodal proximity regions and shifting criteria. For those base triangles or tetrahedra that are heterogeneous, it would be most favourable if their existing nodal points could simply be translated or shifted to coincide with their closest point projection onto the nearest material interface. In this way, a final homogeneous, conforming mesh could be achieved without adding any new nodes or elements, but simply by modifying nodal co-ordinates (conformal mapping). Note that nodal shifting without considering the shape of element domain joining the node could easily result in severe element distortion and bad element aspect ratios. Therefore, nodal shifting must be coupled together with the further splitting of base triangles and tetrahedra to achieve materially homogeneous elements.

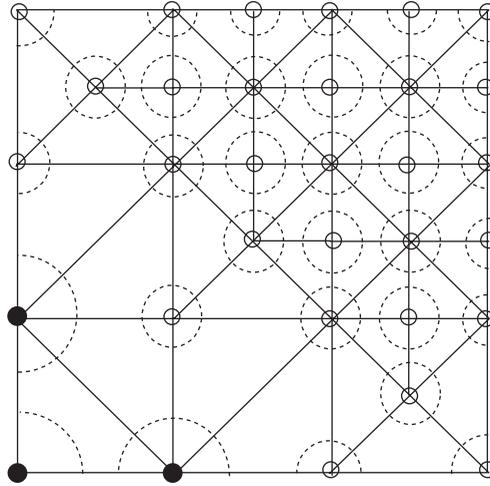


Figure 7. Schematic of proximity radii and proximity regions. Nodes belonging only to smaller elements have smaller proximity radii and regions.

The criteria on when to simply shift nodes, when to split the elements, and when to do both makes usage of a *proximity radius*, which is simply the maximum distance that a given node can be shifted to lie along a material interface (Figure 7). The area around each node swept out by its proximity radius is called its *proximity region*. The proximity radius for each node is specified heuristically as some fraction of the minimum edge length of the elements to which it belongs. In this way, nodes that serve as vertices of only large elements will have large proximity radii and regions, while nodes serving as vertices for small elements will have small proximity radii and regions. There is some flexibility in choosing the size of the proximity radius. Selection of the radii too large can result in degeneration of the base triangles or tetrahedra into those having undesirably large aspect ratios. On the other hand, selection of the proximity radii too small can result in unnecessary splitting of the base triangles or tetrahedral into those with very large aspect ratios as is shown in Figure 8. In this work, the proximity radius for each node has been successfully taken as one-third the minimum edge length of all the base elements to which the node belongs.

Once the base-mesh and nodal proximity regions have been established, a multi-dimensional search process is undertaken for each node in the mesh to determine whether or not a material interface passes within that node's proximity region. To facilitate the search of each node's proximity region, eight discrete search directions are used in two dimensions, and 27 are used in three dimensions (Figure 9). These numbers of search directions are selected only for implementational convenience and the need to have them evenly distributed over a plane (in 2D) or space (in 3D). Larger values that provide uniform directional coverage could also be considered, although the ability of these algorithms to resolve object boundaries is less dependent on the number of search directions than the size of base-elements around that object's material interface. Generally, as the curvature of a material interface increases, the size of elements around it will decrease considerably.

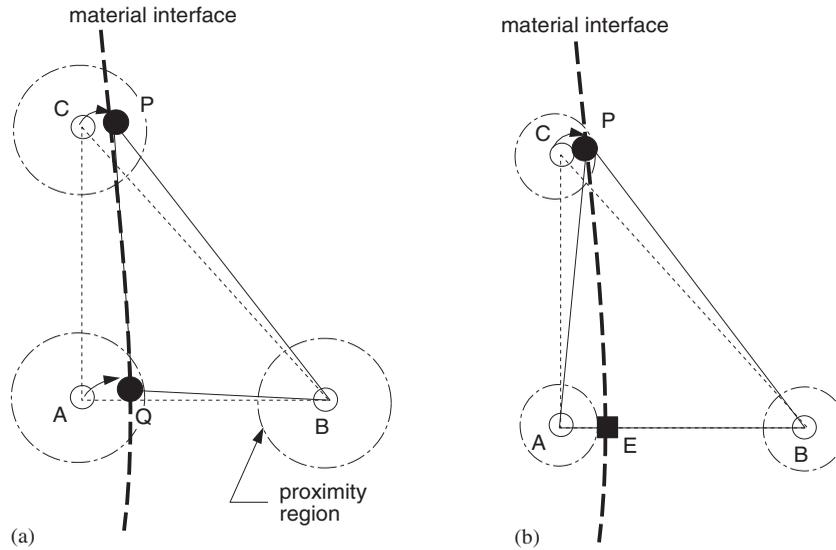


Figure 8. Example showing how proximity radius too small leads to high aspect ratio element after element splitting: (a) with interface node **A** having an adequate proximity radius no splitting of **BAC** is needed but with node shifting $\mathbf{BAC} \Rightarrow \mathbf{BQP}$ with small aspect ratio; and (b) with interface node **A** having a proximity radius too small, edge-cut **E** is introduced close to node **A** to capture material interface. The original element, triangle $\mathbf{BAC} \Rightarrow \mathbf{BEP} \cup \mathbf{EAP}$ with **EAP** having a large aspect ratio.

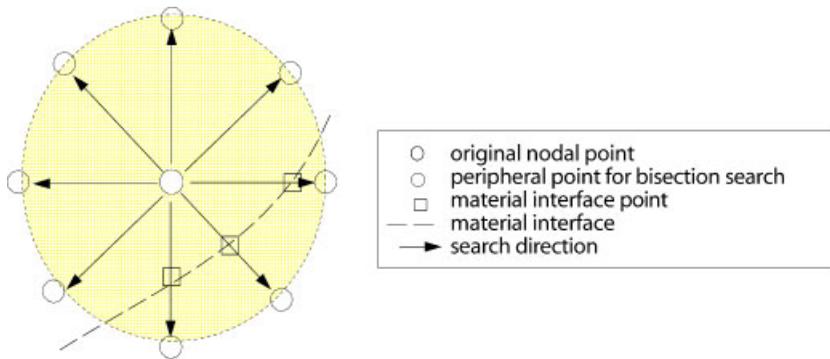


Figure 9. Nodal point and its proximity region in two-dimensions. Eight search directions are defined, and bisection searches are used to find interface points within the proximity region.

For each search direction, the point at the periphery of the proximity region is tested to see if it lies in the same material as the node under consideration. If the periphery point lies in a different material, a conventional bisection search is performed using the discrete-valued MDIFs that return the value of the reinforcing material object in which a test spatial

point resides. Points lying in the search direction and on a material interface are found by determination of the points at which the MDIFs change value. Since the search is performed in a number of different directions, a multitude of interface points can potentially be found for each node. In this case, the search direction with the minimum distance to a material interface is considered and the others are discarded. When a node has a single material interface passing through its proximity region, that node is then designated as an *interface node*. The co-ordinates of *interface nodes* are then changed to those of the nearest interface point found in the search, and the two materials associated with the interface are recorded in a nodal data structure. For *non-interface nodes*, the co-ordinates remain unchanged, but the material in which the node resides is recorded in the nodal data structure.

When the search of a node's proximity region finds more than one material interface, it is clear that the vicinity around that node needs a higher degree of mesh resolution than exists. Accordingly, the node is marked as a *multi-interface node*. If after the searching of all nodes' proximity regions there exist any *multi-interface nodes*, the meshing process then reverts back to the preceding voxel-based mesh, and those voxels to which the *multi-interface nodes* belong are further subdivided. The meshing process then returns once again to the voxel-splitting process to form the base tetrahedra. This cyclic process can continue as long as is necessary to achieve a base tetrahedral mesh with no *multi-interface nodes*.

2.4.2. Element splitting techniques. A *homogeneous* base-triangle or tetrahedron is one in which all vertex nodes lie either in the same material or on the material interface, whereas a *heterogeneous* element is one in which not all vertex nodes lie in the same material. When a base triangle or tetrahedron is *heterogeneous*, it is necessary to determine where the material interface(s) passes through the element. For *heterogeneous* elements, a bisection search, again using the MDIF, is thus performed along all edges connecting vertex nodes residing in different materials.

Whenever a heterogeneous element has a material interface intersecting one of its edges at a point not residing in either of the end-node's proximity regions, then it may be necessary to split or subdivide that element into sub-triangles or sub-tetrahedra. In the current framework, two very simple rules are used to guide the splitting of triangles and tetrahedra:

1. Cuts are never introduced along edges connecting end-nodes, where one or both of the nodes is an interface node; and
2. Triangles, in both two and three dimensions, are only split using one or three edge-cuts, but never two cuts [28].

The first rule precludes the creation of extremely slender elements with high aspect ratios, and the second rule precludes tetrahedral face-mismatching in three-dimensions (Figure 10). Although there are generally ten ways to split tetrahedra using edge-cuts (Figure 11), only four of these splitting patterns are actually permitted under the second rule cited above. Whenever a triangular element or face has two edges with interface points, an additional point must be added to the third edge to facilitate the *auxiliary edge-cut*. In three-dimensions, the addition of an auxiliary edge-cut to one face of a tetrahedron can create a total of two edge-cuts on another of its faces. In that case, an additional edge-cut would be required on that triangular face. This iterative process should be continued until no faces of the tetrahedron in question have two edge-cuts.

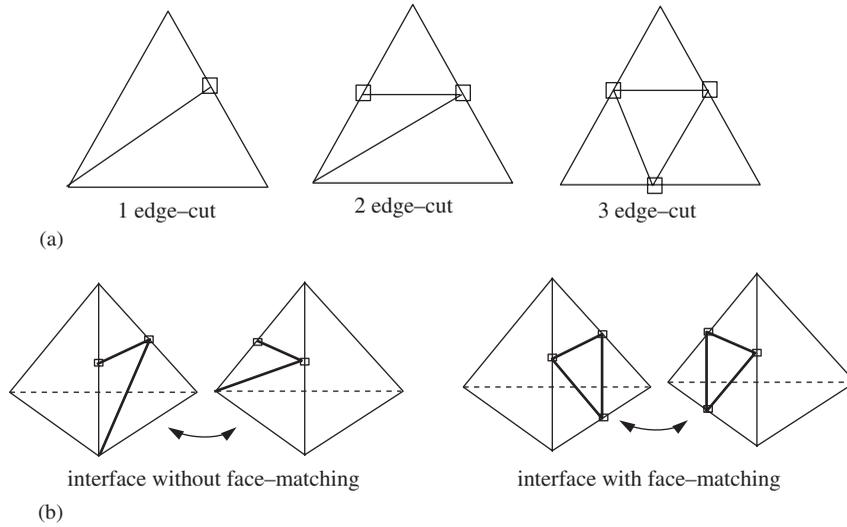


Figure 10. General splitting patterns for triangles: (a) one, two, and three edge-cuts; and (b) illustration of how two-edge cuts can lead to face mis-matching between adjacent tetrahedra in three-dimensions.

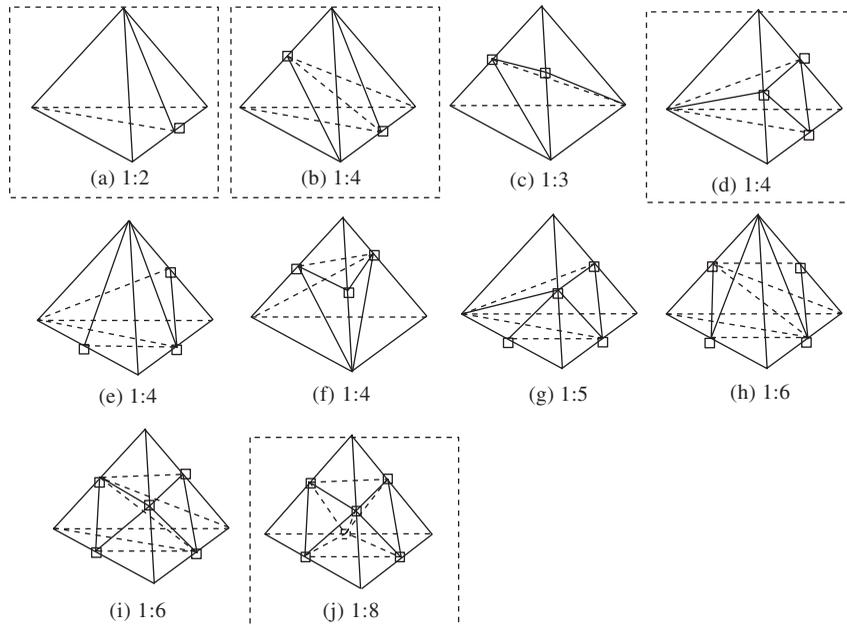


Figure 11. Ten general patterns for splitting a tetrahedron into sub-tetrahedra using edge bisection. Of the ten general patterns, only four (a, b, d, and j) are permissible due to the requirement on triangle splitting with either one or three edge-cuts.

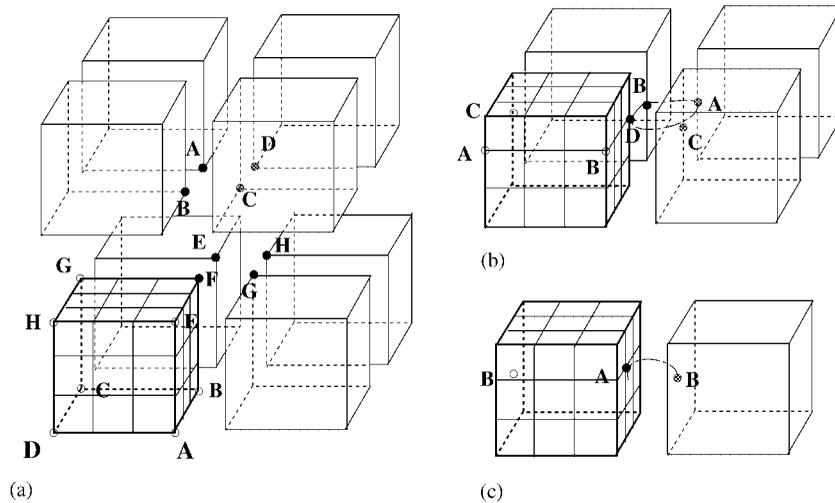


Figure 12. Schematic of adjacent unit cells for a periodic medium, with: (a) virtual pairs of nodes at the unit cell's eight vertices; (b) virtual pairs of nodes along edges; and (c) virtual pairs on exterior faces.

2.5. Periodic boundary condition considerations

In unit cell computations, periodic boundary conditions are generally applied to those nodes on the exterior faces of the model. This creates a need for displacement compatibility on opposing exterior faces of the unit cell (Figure 12).¹ For example in three dimensions, since all eight nodes at the vertices of the unit cell model occupy the same relative location within the unit cell (Figure 12(a)), they must all experience the same periodic displacement \mathbf{u}^* [3]. These eight nodes can be associated with each other using data structures here denoted as virtual pairs [7]. Each vertex node of the unit cell is a member of seven virtual pairs, and thus each has seven virtual twins. Nodes along unit cell edges are members of three virtual pairs, and thus have three virtual twins (Figure 12(b)) while nodes lying on the interior of exterior surfaces of the unit cell are members of only one virtual pair and thus have only a single virtual twin (Figure 12(c)). When nodes are members of virtual pairs, they are enslaved to each other and their degrees of freedom \mathbf{u}^* are identical. Effectively this is achieved simply by giving the twin nodes' degrees of freedom identical equation numbers in the system of discrete finite element equations.

Beyond using virtual pairs, it is also necessary that the mesh discretizations of corresponding external faces be identical. This is facilitated from the outset by choosing the initial number of voxels in the principal co-ordinate directions of the initial mesh to be even. When the voxels are then split into tetrahedra using the alternating HEX5-even and HEX5-odd methods, the base meshing of the unit cell exterior faces will be compatible. Figure 13 shows a

¹In the event that aperiodic media are being considered and periodic boundary conditions will not be imposed, model development can proceed straightforwardly neglecting all of the issues pertaining to periodicity raised in Section 2.5.

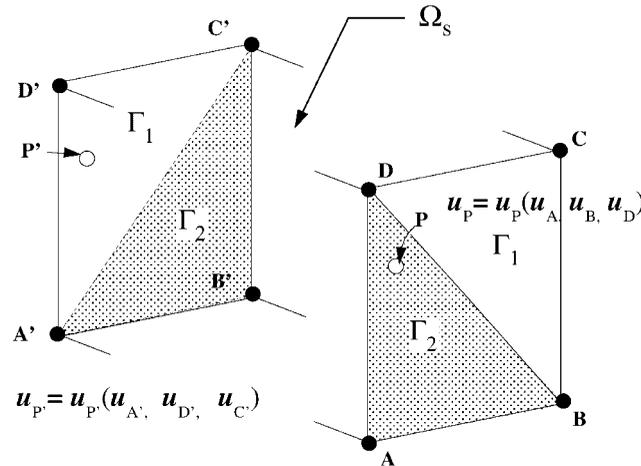


Figure 13. Schematic demonstrating inability to achieve displacement field periodicity on opposing exterior faces of the unit cell domain Ω_s when the faces have different mesh discretizations.

simple example where corresponding external faces of a unit cell are meshed incompatibly. In this example, even if the virtual node pairs (A–A', B–B', C–C' and D–D') are all properly enslaved, the displacements \mathbf{u}^* of interior virtual pair points P and P' will not necessarily be equal since they are each functionally dependent on different sets of nodes that are not corresponding virtual pairs.

In the current framework, the mesh of the unit cell is constructed in three stages: (1) creation of the background pixel/voxel mesh; (2) splitting of pixels/voxels into triangles/tetrahedra to form the base-triangular or tetrahedral mesh; and (3) additional node-shifting and element splitting to achieve the final triangular/tetrahedral mesh. It is implicitly assumed that the background pixel/voxel mesh has already been established using the techniques described at length in Reference [7] and that this background mesh has produced identical discretizations of exterior faces of the unit cell, and that all nodes on the exterior faces are members of between one and seven virtual pair groups. From this point, the creation of the base-triangular/tetrahedral mesh, and the final mesh proceed as described above, with the caveat that the following considerations must be heeded to facilitate enforcement of periodicity on the unit cell response:

1. Whenever a new node is introduced to an exterior face of the model the corresponding virtual pair nodes must also be added on opposing faces.
2. The shifting (if any) for two nodes that are members of the same virtual pair should be identical.

The first consideration above assures that each node on the exterior faces of the unit cell will have at least one virtual twin. The second permits exterior faces of the unit cell to undergo compatible distortion during the node-shifting process. Figure 14 illustrates the identical node-shifting of virtual twins \mathbf{N} and \mathbf{N}' . Although this could conceivably be achieved by considering virtual replicas of all reinforcing objects in the unit cell, it is much simpler to make virtual

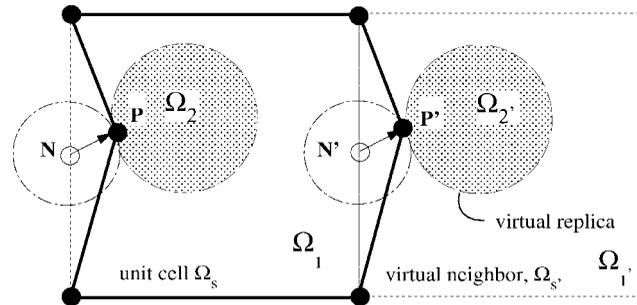


Figure 14. Identical node-shifting of virtual twins on exterior boundary of unit-cell.

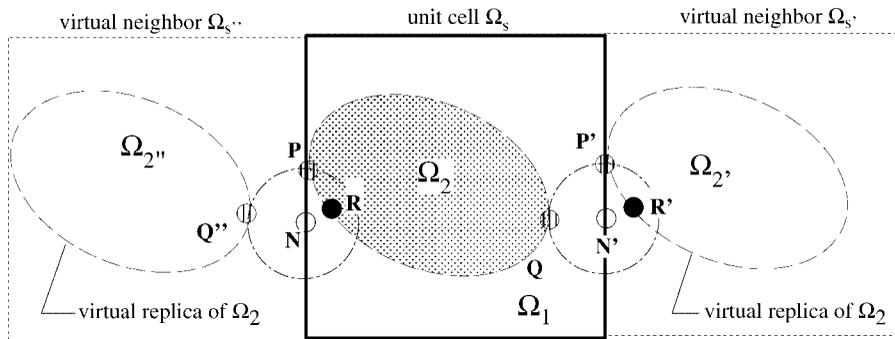


Figure 15. Nodes N and N' are designated as *multi-interface nodes*, since two material interfaces pass through their proximity region.

twins share the same node attributes (i.e. the designation as *interface nodes* with the same position shift).

When a node has more than one material interface passing through its proximity region, it is designated a *multi-interface node*, and the base voxels to which it belongs must undergo further refinement (Section 2.4.1). For nodes lying on the external faces of the unit cell the detection of multiple interfaces passing through the nodal proximity region can be a challenge (Figure 15). Strictly speaking when searching the proximity region of a node on the unit cell boundary, one should consider not only all of the reinforcing objects in the unit cell, but also their *virtual replicas* in adjoining unit cells. In the example shown in Figure 15, node N and its virtual twin N' would be designated *multi-interface nodes* and the mesh would be refined until this was no longer the case.

2.6. Symmetry reductions of unit cell models

Unit cells of periodic media will often feature one or more planes of symmetry that can be exploited to achieve considerable efficiency in unit cell analysis, since only a fractional portion of the complete unit cell Ω_s needs to be meshed and modeled. When developing a reduced unit cell model that exploits symmetry, the fact that heterogeneous periodic displacements \mathbf{u}^*

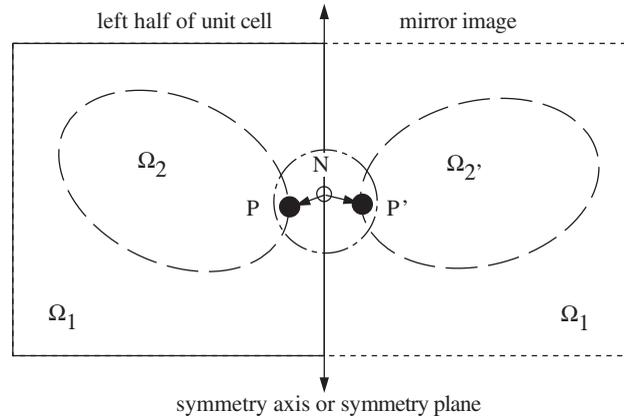


Figure 16. Symmetric unit cell of which only half needs to be modelled. During meshing, nodes along the symmetry plane can be shifted only within the plane.

vanish on the symmetry planes means that the opposing exterior surfaces of the model need not have identical mesh discretizations. In such cases, many of the periodicity considerations in the preceding subsection can be neglected. Nevertheless, in meshing the reduced unit cell model, when nodes lying on any of the symmetry plane(s) are found to be *interface nodes*, they can only be shifted to a new position within the symmetry plane. As the example of Figure 16 shows, when the position shift of an interface node on a symmetry boundary would move it off of that boundary, then the mesh in the vicinity of that node requires further refinement to resolve the relevant material interfaces.

2.7. Positioning of intermediate nodes in quadratic elements

In all of the preceding discussion, the positioning of nodes and the description of elements has been entirely in terms of the nodes at the vertices of triangular and tetrahedral elements. As stated at the outset, six-noded quadratic triangular elements (in 2D) and 10-noded quadratic tetrahedral elements (in 3D) have been used in this effort since they have superior numerical performance characteristics with respect to mesh-locking phenomena. The positioning of the intermediate nodes along all edges of triangular or tetrahedral elements remains to be specified. Clearly, the simplest way to position the intermediate nodes on element edges is to place them at the mid-point of the line segment formed by the end nodes. In the case of elements with edges whose ends are interface nodes, however, a better approach might be to position the intermediate nodes so as to lie along the same material interface as the end nodes. In doing so, the mesh of the unit cell model would achieve piecewise quadratic, as opposed to piecewise linear, resolution of material interfaces. In principle this would reduce the error associated with truncation of reinforcement phase volume fractions when their interfaces are represented in piecewise linear or piecewise planar approximations. This issue is not as simple as it appears to be on its face, however, since direct pursuit of such an approach often leads to quadratic elements with negative Jacobian determinants [31]. Accordingly, in all of the example problems solved below, the intermediate nodes were simply placed at the midpoints of the line segments connecting the two vertex nodes.

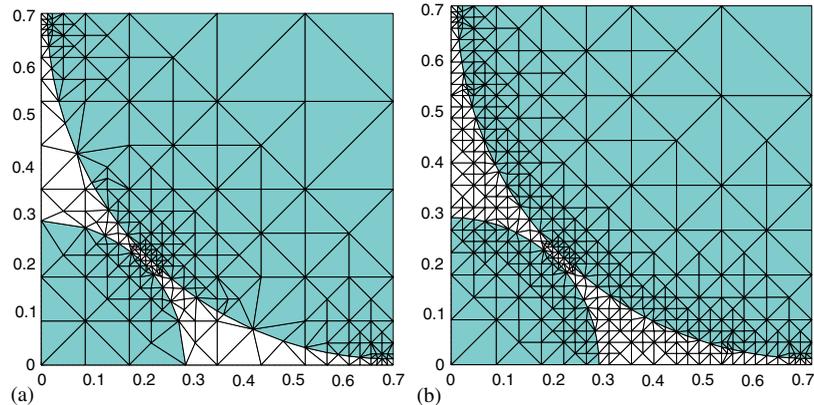


Figure 17. Quarter unit cell aligned fibre-composite mesh with two fibre sizes and two mesh-refinement options in 2D: (a) the minimum refinement needed to resolve the gap between adjacent fibres is used; and (b) refinement exceeds the minimum needed to resolve spacing between fibres.

3. DEMONSTRATIVE RESULTS

To demonstrate the robustness and efficacy of proposed meshing techniques, three different types of composite unit-cells are modelled here: (1) an aligned fibre composite at very high fibre volume fraction, achieved using multiple fibre diameters; (2) a particulate composite at moderately high volume fraction of reinforcement; and (3) a three-dimensional continuous textile-reinforced composite material. The adequacy of the proposed meshing techniques is explored by considering the convergence behaviour of the computed effective elastic constants from the unit cell models. For the composites considered here, the geometrical complexity of the reinforcing material objects is not a major difficulty since spheres, cylinders, and warped yarns with elliptical cross-sections are fairly simple and straight-forward to describe with MDIFs. In the examples presented, a more immediate concern is that caused when the reinforcing objects are packed very tightly with minimal spacing between them, as this typically requires high degrees of mesh resolution simply to capture the material interfaces.

3.1. *Densely packed aligned fibre composite*

In this example, the material system is a composite of aligned silicon-carbide fibres (Si-C) and an indium-tin matrix (In-Sn). Such composites are of interest because of the high stiffness of the fibres and the high viscoelastic damping characteristics of the matrix [32]. To achieve both high stiffness and high mechanical damping behaviours, composites with high fibre volume fractions are considered. One way to achieve this is to employ aligned fibres with different diameters to enhance their packing in the transverse plane. Using two different diameters, a fibre volume fraction of 92% is obtained with the unit cell model whose meshes are shown in Figure 17. Since the unit cell features two planes of symmetry, unit cell analysis can be performed on 1/4 unit cell models. Since the fibres are extremely close together in this model, a moderately high degree of mesh refinement is required to resolve the spacing

Table I. Computed effective moduli of aligned-fibre composite with two fibre diameters and a 92% fibre volume fraction.

Properties of SiC/InSn fibre composite		
	E	ν (Poisson's ratio)
Fiber (silicon-carbide)	400 GPa	0.14
Matrix (indium-tin)	20.3 GPa	0.35
Computed effective material properties (values are identical to three figures with both meshing options)		
C_{11} (GPa)	C_{12} (GPa)	C_{66} (GPa)
253	65.7	94.9
Mesh specifications with different mesh refinement options		
	Number of elements	Number of nodes
Minimum refinement	2921	1424
Additional refinement	12329	6124

between fibres. The computed effective elastic moduli of Table I indicate that the mesh that is just sufficiently refined to resolve the gap between fibres is also adequate to provide elastic moduli to three digits of precision.

3.2. Particulate composites

In this example, the methods are applied in three-dimensions to create the unit cell of a particulate composite of silicon-carbide (Si-C) and indium-tin (In-Sn). The inclusions are assumed to be spherical, of uniform size, and arranged in a face-centred cubic (FCC) pattern, which gives up to about 70% volume fraction when uniform particles are closely packed. Although the effective material constants of such a composite feature cubic symmetry, they are not isotropic. Owing to cubic symmetry, the effective elastic constants are identical in each of the composite's principal planes of material symmetry. The three effective elastic constants computed for the medium are the principal Young's modulus $E(=E_1=E_2=E_3)$, the Poisson's ratio in the principal planes $\nu(=\nu_{12}=\nu_{23}=\nu_{31})$ and the shear modulus in the principal planes $G(=G_{23}=G_{13}=G_{12})$. These first two constants can be determined using 1/8 unit cell models, whereas the shear modulus can be computed using 1/4 of the unit cell as shown in Figure 18. These meshes for these partial unit cell models are shown in Figure 19.

The convergence behaviour of the computed elastic constants in Figure 20 shows that the effective stiffness moduli increase with increasing mesh refinement. This potentially surprising behaviour is due to the fact that intermediate nodes of the quadratic tetrahedral elements were positioned midway along element edges, leading to truncation of the SiC volume fraction at low mesh resolutions.

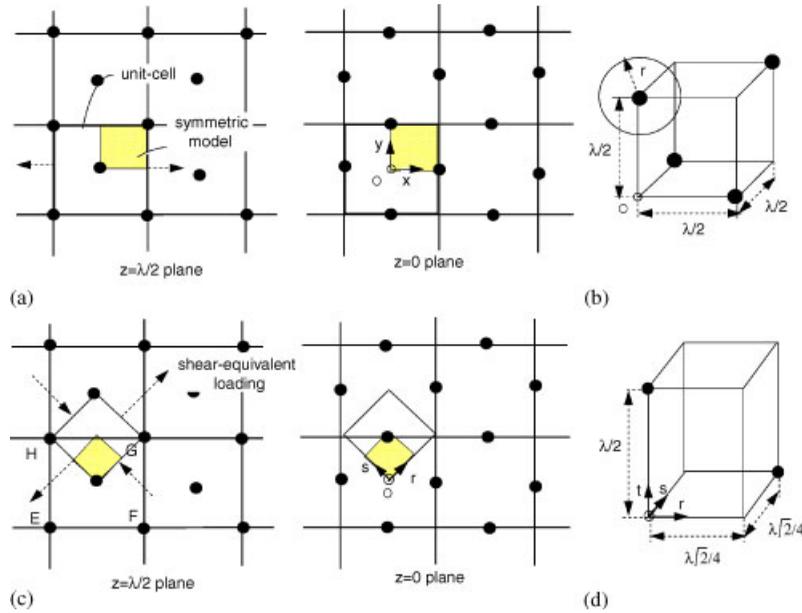


Figure 18. (a) Face-centred cubic arrangement of spherical SiC particles in the composite; (b) 1/8 unit cell model and boundary conditions used to compute extensional moduli; (c) extensional loading on rotated unit cell to obtain shear moduli; and (d) dimensions of the reduced unit cell model.

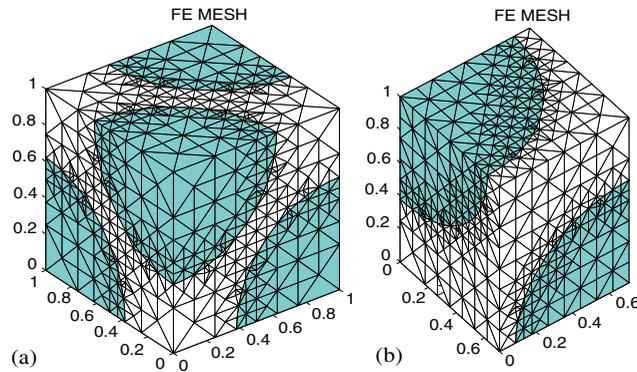


Figure 19. FE meshes for particulate SiC-InSn composite (50% particle volume fraction): (a) isometric view of model used in unit cell analysis to compute extensional moduli of particulate composite; and (b) isometric view of model used to compute shear moduli of the same particulate composite.

3.3. Textile composite

As a final example, the proposed meshing techniques are applied to unit cell analysis of the textile composite whose unit cell was originally shown in Figure 1. Here the linear

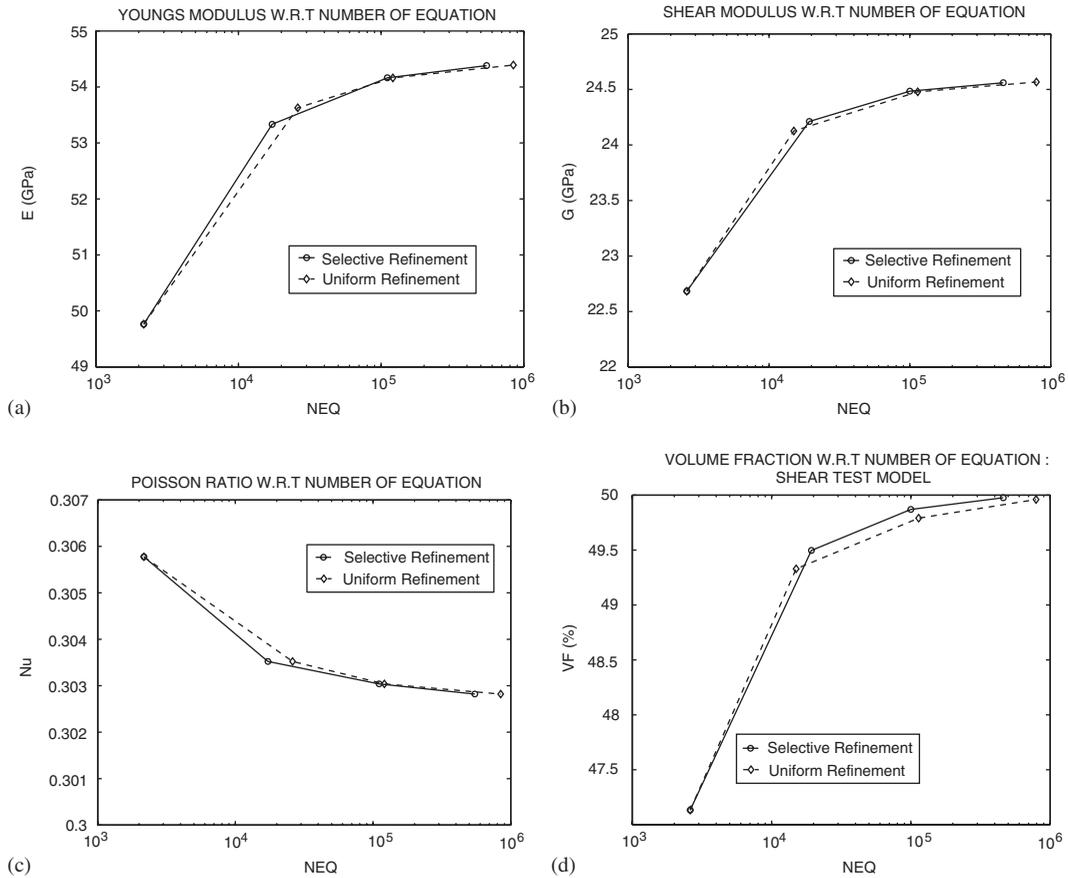


Figure 20. Convergence of computed characteristics of particulate SiC-InSn composite with mesh refinement (50% particle volume fraction): (a) principal Young's modulus E ; (b) principal shear modulus G ; (c) Poisson's ratio ν ; and (d) actual particulate reinforcing volume fraction.

elastic moduli of the textile are computed, and in other forthcoming works, the geometrically non-linear response characteristics of the composite are also computed with the same unit cell model. The dimensions and geometrical descriptions of yarns and the unit cell model accompany Figure 3 and are presented in Table II. The yarns are modelled as transversely isotropic aligned-fibre composites with 70 and 30% volume fractions of E-glass and cast epoxy matrix, respectively. In the unit cell model, the local fibre orientation in the yarns is taken into account as described in Section 2.2. Since the material arrangement within the unit cell shows orthotropy with identical properties in the 1- and 2-directions, six constants, $E_1 = E_2$, E_3 , $G_{13} = G_{23}$, G_{12} , $\nu_{13} = \nu_{23}$, and ν_{12} are necessary to describe the effective elastic constants of this composite. The unit cell mesh used in this analysis is shown in Figure 21, and the computed effective elastic constants are presented in Table III.

Table II. Dimensions of textile composite with 41% yarn volume fraction (cf. Figure 3).

Overall dimensions							
L_x			L_y			L_z	
40			40			15	
Yarn reinforcements							
Yarn index	Auxiliary points defining local $\xi\eta\zeta$ -co-ordinates			Characteristic measures for woven yarns			
	A	B	C	l_η	l_ζ	r_η	r_ζ
1	(-20, 10, 0)	(20, 10, 0)	(-20, -10, 1)	4	40	2.5	4.7
2	(-20, -10, 0)	(20, -10, 0)	(-20, -10, -1)	4	40	2.5	4.7
3	(10, -20, 0)	(10, 20, 0)	(10, -20, 1)	4	40	2.5	4.7
4	(-10, -20, 0)	(-10, 20, 0)	(-10, -20, -1)	4	40	2.5	4.7
5	(-20, -20, -7.5)	(-20, -20, 7.5)	(-20, 0, -7.5)	0	15	4.7	4.7
6	(-20, 20, -7.5)	(-20, 20, 7.5)	(-20, 0, -7.5)	0	15	4.7	4.7
7	(20, -20, -7.5)	(20, -20, 7.5)	(20, 0, -7.5)	0	15	4.7	4.7
8	(20, 20, -7.5)	(20, 20, 7.5)	(20, 0, -7.5)	0	15	4.7	4.7
9	(-20, 0, -7.5)	(-20, 0, 7.5)	(0, 0, -7.5)	0	15	4.7	4.7
10	(20, 0, -7.5)	(20, 0, 7.5)	(0, 0, -7.5)	0	15	4.7	4.7
11	(0, -20, -7.5)	(0, -20, 7.5)	(0, 0, -7.5)	0	15	4.7	4.7
12	(0, 20, -7.5)	(0, 20, 7.5)	(0, 0, -7.5)	0	15	4.7	4.7
13	(0, 0, -7.5)	(0, 0, 7.5)	(-20, 0, -7.5)	0	15	4.7	4.7

4. DISCUSSION

The methods that have been described and developed here are a direct extension of the hierarchical voxel/pixel techniques described in Reference [7]. In the preceding work, discrete resolution of material interfaces within composites could generally be achieved only in the limit of infinite mesh refinement. Here, the background mesh of pixels and voxels can undergo fewer cycles of refinement, following which background hierarchical mesh of pixels/voxels is transformed into a mesh of quadratic triangles or tetrahedra using element splitting and nodal shifting to fall on material interfaces. Once a description of the material arrangements in a composite is generated through the so-called MDIFs, the complete mesh generation process can be fully automated, and need not require graphical interaction on the part of the user.

The two-sided meshing algorithm presented here has been developed and tested in the context of periodic, multi-material composite material systems. While it is certainly appropriate for this application, it applies to a variety of other less specialized applications where multiple regions of a domain must be meshed in a two-sided way that preserves the conformity between the regions. For example, meshes for fluid–structure interaction applications could be achieved with these techniques.

In this work, attention has been confined primarily to meshing with the objective of achieving high geometric fidelity of material arrangements as described by the so-called MDIFs. Though not developed and applied here, the underlying meshing data structures and rules

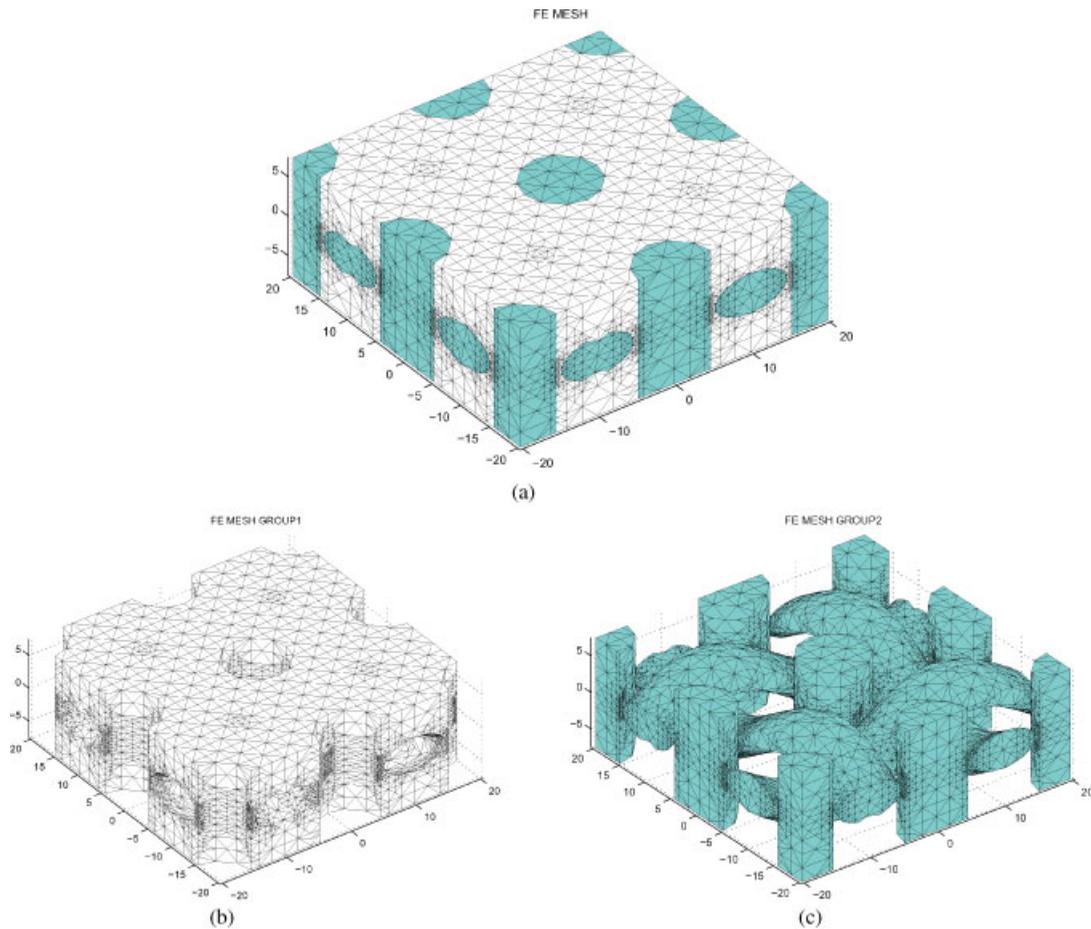


Figure 21. FE mesh for unit-cell model of textile composite, isometric view; (a) complete unit-cell; (b) matrix only; and (c) yarns only.

can be straightforwardly applied to adaptive analysis in which regions of the mesh with high local error estimates are refined either progressively, or *a posteriori*. The attractive aspect of the current meshing framework, however, is that as the mesh is adaptively refined to reduce solution errors, it would also be refined in a manner that leads to enhanced resolution of the underlying material arrangements.

If the proposed framework were to be used in adaptive mesh refinement, specific triangular/tetrahedral elements in the mesh would be tagged for refinement in *a posteriori* error analysis. The tagging for refinement would be passed back to the parent pixels or voxels in the hierarchical background mesh, and the background mesh would then undergo an additional cycle of selective refinement in which only the tagged pixels or voxels and their neighbours would be refined. The refined background pixel/voxel mesh would then undergo transformation to a mesh of discrete triangles and tetrahedra for further analysis.

Table III. Computed effective material properties of textile composite with 41% yarn volume fraction.

Properties of E-glass fibre and epoxy matrix						
	E			ν (Poisson's ratio)		
Fiber (glass fibre)	72.4 GPa			0.20		
Matrix (epoxy)	2.76 GPa			0.35		
Properties of yarn (70% Fiber volume fraction)						
	$E_1(=E_2)$	E_3	$G_{13}(=G_{23})$	G_{12}	$\nu_{13}(=\nu_{23})$	ν_{12}
Yarn	17.3 GPa	51.8 GPa	6.48 GPa	6.48 GPa	0.330	0.076
Computed moduli of textile composite						
Composite	7.99 GPa	12.4 GPa	2.11 GPa	2.11 GPa	0.173	0.326
Mesh specification for textile composite						
Volume fraction	Number of nodes			Number of elements		
40.7%	199 201			140 016		

Some of the demonstrative examples presented above show that coarse meshes tend to underestimate the volume fractions of convex reinforcing objects, since their boundaries are approximated here as piecewise linear or piecewise planar. If the intermediate nodes of edges could be positioned to fall on material interfaces, as opposed to merely the midpoint of the segment connecting end nodes, then the quadratic elements used in this study would achieve piecewise quadratic approximation of material interfaces. Such an approach requires some caution, though, since if not carefully applied, it can result in triangular elements and tetrahedra featuring negative jacobian determinants. As a future extension of this work, algorithms for positioning the mid-nodes of 6-noded triangles and 10-noded tetrahedra so that they fall on the same material interface as their end nodes without causing degeneration of the element characteristics could be addressed.

One of the novel aspects of the proposed meshing algorithms presented here is the usage of analytical material domain identification functions to describe material arrangements. For many types of composite materials, it could conceivably be quite challenging to obtain the analytical functions to describe material arrangements. Consequently, this is a subject that may require further investigation or development.

5. CONCLUSIONS

The proposed meshing and analysis techniques presented here are based on a novel approach in which the material arrangements in composites are described through collections of analytical functions. Since the analytical material descriptions are independent of the mesh

resolution, the meshes can be iteratively refined to achieve the described material arrangements with progressively increasing fidelity. Good element quality is achieved in this approach by using a background mesh of pixels/voxels that can be selectively refined any number of times using edge bisection while maintaining the initial element shape and aspect ratios of the original mesh. The excellent performance of the proposed meshing techniques has been demonstrated on examples involving development of unit cell models for composite materials with non-trivial material arrangements, and in computing the effective elastic constants of these composites using finite element analysis.

ACKNOWLEDGEMENTS

This work was funded by National Science Foundation research grants CMS-9713335, CMS-9896284, and a research grant from the Whitaker Foundation 96-0636. The support from both of these foundations is gratefully acknowledged.

REFERENCES

1. Suquet PM. Elements of homogenization for inelastic solid mechanics. In *Homogenization Techniques for Composite Media*, Sanchez-Palencia E, Zaoui A (eds). Springer: Berlin, 1987; 193–278.
2. Guedes J, Kikuchi N. Preprocessing postprocessing for materials based on homogenization method and adaptive finite element methods. *Computer Methods in Applied Mechanics and Engineering* 1991; **83**:143–198.
3. Swan CC. Techniques for stress and strain controlled homogenization of inelastic periodic composites. *Computer Methods in Applied Mechanics and Engineering* 1994; **117**:249–267.
4. Zohdi TI, Wriggers P. Aspects of the computational testing of the mechanical properties of microheterogeneous material samples. *International Journal for Numerical Methods in Engineering* 2001; **50**:2573–2599.
5. Wentorf RR, Collar R, Shephard MS, Fish J. Automated modeling for complex woven mesostructures. *Computer Methods in Applied Mechanics and Engineering* 1999; **172**:273–291.
6. Swan CC, Kosaka I. Homogenization-based analysis and design of composites. *Computers and Structures* 1997; **64**:603–621.
7. Kim HJ, Swan CC. Voxel-based meshing and unit cell analysis of textile composites. *International Journal for Numerical Methods in Engineering* 2003; **56**(7):977–1006.
8. Kikuchi N. Keynote lecture on digital-image based modeling. *Proceedings of the International Symposium on Optimization and Innovative Design*, Tokyo, July 1997.
9. Hollister SJ, Riemer B. Digital image based finite element analysis for bone microstructure using conjugate gradient and Gaussian filter techniques. *SPIE Proceedings: Mathematical Methods in Medical Imaging II* 1993; 95–106.
10. Hollister SJ, Kikuchi N. Homogenization theory and digital imaging: a basis for studying the mechanics and design principles of bone tissue. *Biotechnology and Bioengineering* 1994; **43**(7):586–596.
11. Tareda K, Miura T, Kikuchi N. Digital image-based modeling applied to the homogenization analysis of composite materials. *Computational Mechanics* 1997; **20**:331–346.
12. Swan CC. Unit cell analysis of 3-D graphite textile-reinforced polymer matrix composites. In *Advances in Composite Materials and Mechanics*, Arup Maji (ed). ASCE: New York, 1999.
13. Strouboulis N, Coppers K, Babuska I. The generalized finite element method: an example of its implementation and illustration of its performance. *International Journal for Numerical Methods in Engineering* 2000; **47**: 1401–1417.
14. Moes N, Dolbow J, Belytschko T. A finite element method for crack growth without remeshing. *International Journal for Numerical Methods in Engineering* 1999; **46**:131–150.
15. Belytschko T, Moes N, Usui S, Parimi C. Arbitrary discontinuities in finite elements. *International Journal for Numerical Methods in Engineering* 2001; **50**:993–1013.
16. Sukumar N, Chopp DL, Moes N, Belytschko T. Modeling holes and inclusions by level sets in the extended finite-element method. *Computer Methods in Applied Mechanics and Engineering* 2001; **190**:6183–6200.
17. Hole K. Finite element mesh generation methods: a review and classification. *Computer Aided Design* 1988; **20**(1):27–38.
18. Thompson JF, Thomas FC, Mastin CW. Automated numerical generation of body fitted curvilinear co-ordinate system for field containing any number of arbitrary two-dimensional bodies. *Journal of Computational Physics* 1974; **15**:299–319.

19. Tasy TK, Hsu FS. Numerical grid generation of an irregular region. *International Journal for Numerical Methods in Engineering* 1997; **40**:343–356.
20. Greaves DM, Borthwick AGL. Hierarchical tree-based finite element mesh generation. *International Journal for Numerical Methods in Engineering* 1992; **29**(4):513–520.
21. Chan CT, Anastasiou K. An automatic tetrahedral mesh generation scheme by the advancing front method. *Communications in Numerical Methods in Engineering* 1997; **13**:33–46.
22. Löhner R, Oñate E. An advancing front point generation technique. *Communications in Numerical Methods in Engineering* 1998; **14**:1097–1108.
23. Krysl P, Ortiz M. Variational Delaunay approach to the generation of tetrahedral finite element meshes. *International Journal for Numerical Methods in Engineering* 2001; **50**(7):1681–1700.
24. Houman B, Paul LG. Aspects of 2-D Delaunay mesh generation. *International Journal for Numerical Methods in Engineering* 1997; **40**:1957–1975.
25. Babuška I, Aziz AK. On the angle condition in the finite element method. *SIAM Journal on Numerical Analysis* 1976; **13**(2):214–226.
26. Field DA. Qualitative measures for initial meshes. *International Journal for Numerical Methods in Engineering* 2000; **47**:887–906.
27. Arnold DN, Mukherjee A, Pouly L. Locally adapted tetrahedral meshes using bisection. *SIAM Journal on Scientific Computing* 2000; **32**(2):431–448.
28. Cougny HL, Shephard MS. Parallel refinement and coarsening of tetrahedral meshes. *International Journal for Numerical Methods in Engineering* 1999; **46**:1101–1125.
29. Křížek M, Pradlová J. Non-obtuse tetrahedral partitions. *Numerical Methods in Partial Differential Equations* 2000; **16**(3):327–334.
30. Naylor DJ. Filling space with tetrahedra. *International Journal for Numerical Methods in Engineering* 1999; **44**:1383–1395.
31. Ahmed ZI, Canaan SA, Saigal S. Mid-node admissible spaces for quadratic triangular arbitrarily curved 2D finite elements. *International Journal for Numerical Methods in Engineering* 2001; **50**:253–272.
32. Kim HJ, Swan CC, Lakes RS. Computational studies on high-stiffness high-damping SiC–InSn particulate reinforced composites. *International Journal of Solids and Structures* 2002; **39**(23):5799–5812.