

Voxel-based meshing and unit-cell analysis of textile composites

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

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

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

SUMMARY

Unit-cell homogenization techniques are frequently used together with the finite element method to compute effective mechanical properties for a wide range of different composites and heterogeneous materials systems. For systems with very complicated material arrangements, mesh generation can be a considerable obstacle to usage of these techniques. In this work, pixel-based (2D) and voxel-based (3D) meshing concepts borrowed from image processing are thus developed and employed to construct the finite element models used in computing the micro-scale stress and strain fields in the composite. The potential advantage of these techniques is that generation of unit-cell models can be automated, thus requiring far less human time than traditional finite element models. Essential ideas and algorithms for implementation of proposed techniques are presented. In addition, a new error estimator based on sensitivity of virtual strain energy to mesh refinement is presented and applied. The computational costs and rate of convergence for the proposed methods are presented for three different mesh-refinement algorithms: uniform refinement; selective refinement based on material boundary resolution; and adaptive refinement based on error estimation. Copyright © 2003 John Wiley & Sons, Ltd.

KEY WORDS: periodic media; unit-cell-analysis; mesh refinement; homogenization; textile composites; effective properties; automated meshing

1. INTRODUCTION

1.1. Overview

Computational unit-cell homogenization methods typically involve constructing a continuum model of a composite's heterogeneous material structure (Figure 1), and then applying loads

*Correspondence to: Colby C. Swan, Department of Civil and Environmental Engineering, Center for Computer-Aided Design, The University of Iowa, 4016 Seamans Center, 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 21 May 2001
Revised 19 December 2001
Accepted 22 April 2002*

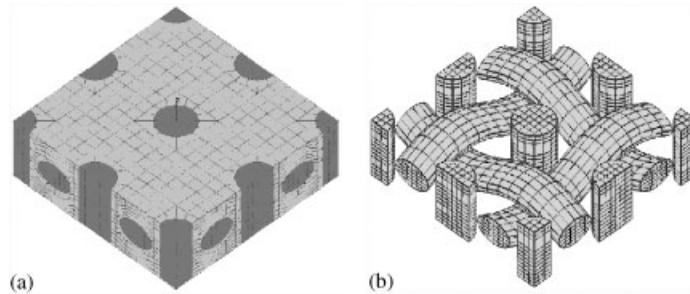


Figure 1. Traditional finite element unit-cell model of plane-weave textile with orthogonal inserts: (a) shows the model with both yarns and matrix; and (b) shows only the woven yarns. The mesh of this model is 'structured'.

to the model with imposed periodic boundary conditions [1–3] to study the property-structure behaviour. While both spectral [4] and finite element [5–9] analysis methods have been proposed and implemented in unit-cell modelling of textile composites or other periodic media, the latter appear more robust due to their ability to handle general, non-linear, material scale behaviours. Nevertheless, there are many very practical challenges to be addressed in computational unit-cell analysis and homogenization of composite materials. The model of a composite's unit-cell must adequately capture material arrangements and volume fractions as well as the local stress/strain fields in the unit-cell under a variety of loading conditions. Wentorf *et al.* [9] have noted that errors in both resolution of material interfaces and material-scale stress and strain fields can lead to appreciable error in computing the effective properties of composites. Therefore, both types of errors require consideration in unit-cell analysis.

There are alternatives to unit-cell analysis techniques that attempt to homogenize textile composites using hybrid micro-mechanics-based rules of mixtures in tandem with geometric decomposition of the textile into portions that are either aligned or orthogonal to the macroscopic strain in the composite. Such techniques have been applied both to textile composites with linear elastic constituents [10] and even inelastic constituents [11, 12]. The current unit-cell analysis framework being developed here is more computationally involved than the preceding techniques but also more general, and thus useful for exploring the property-structure characteristics of textile composites.

The requirements that unit-cell analysis techniques use models that are sufficiently refined to capture both material arrangements and also the material scale stress–strain fields within the composite 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 unit-cell analysis of three-dimensional composite material systems can be considerable, thus requiring usage of rapid and efficient high performance computing techniques [9];
- 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 unit-cell analysis and computational homogenization can be applied. Specifically, novel, voxel-based meshing and analysis techniques based on bilinear and trilinear elements are presented, tested, and discussed here, within the context of unit-cell analysis of periodic textile composites. In a subsequent work [13] the voxel-based unit-cell analysis framework developed here is extended to semi-structured, adaptive meshing techniques based on tri-quadratic tetrahedral elements.

1.2. The basic difficulty

Textile composites are one particular class of materials that serve to illustrate the potential difficulties in both meshing and unit-cell analysis of media with fairly complicated material structures. Considering the plain weave composite with orthogonal inserts whose unit-cell is shown in Figure 1, it can be noted that the meshing of individual yarns in the unit-cell is actually quite simple, whereas the meshing of the polymer matrix phase between the individual yarns is much more difficult. For each individual yarn in the textile's unit-cell, one can start with a mesh of a circular cylinder, and then apply a number of deformation, translation, and rotation transformations to each to obtain the basic textile shown in Figure 1(b). Though meshing of the textile itself is quite straightforward, meshing of the polymer matrix infill between the yarns in such a way that elements with favourable aspect ratios are achieved and such that full compatibility of the displacement field across material boundaries is achieved is somewhat more difficult. 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. For this reason, we seek to develop and test here, and in other works [13], robust, automated meshing techniques that can be easily applied to very broad classes of composite materials.

1.3. Review of mesh generation based on triangles and tetrahedra

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 [14]. The classes of methods being reviewed here can be separated into two categories: (1) those based on structured meshes and (2) those based on non-structured meshes.

1.3.1. Structured mesh generation approaches. 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 [15, 16]. 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.

Another structured approach, similar to the voxel-based techniques to be introduced in the main body of this article, is a grid-based technique. In such techniques, a regular grid of nodes and cells is constructed in space, and the object to be meshed is inserted into the same space. Those cells 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

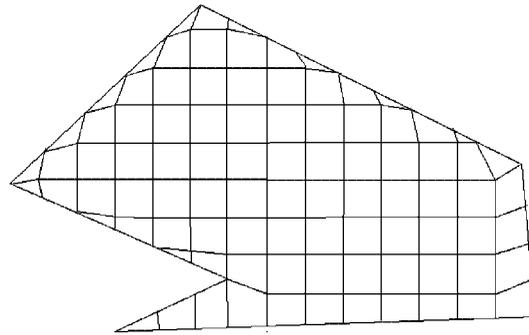


Figure 2. Grid-based approach of Hole [14] in which interior grid cells form good elements, and those along the boundary must be adjusted or trimmed.

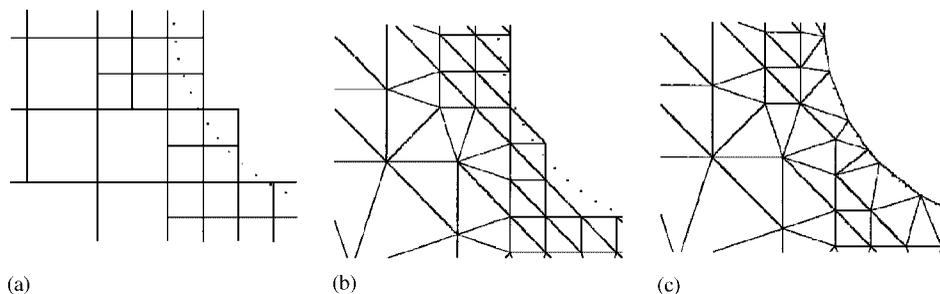


Figure 3. Details of hierarchical, structured grid-based method [17]: (a) quadtree grid around boundary; (b) setting-up mesh by connecting centroids of elements; and (c) translate nodes to boundary.

captures the boundary (Figure 2). The recognized characteristics of the grid-base approach are that: (1) 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 (2) 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 [17] have been developed as extensions of grid-based approaches (Figure 3).

1.3.2. Unstructured mesh generation approaches. Unstructured mesh generation techniques are typically based on triangular elements in two-dimensions and tetrahedral elements in three-dimensions. 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 by Babuška and Aziz [18]. Common quality indicators based on element geometry, and frequently used in pre- and post-processing mesh improvement schemes have thus been developed [19, 20]). 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. One of the major difficulties being addressed

is degeneration of element shapes and aspect ratios under repetitive local mesh refinements [21–23].

In unstructured meshing, advancing front techniques [24, 25], Delaunay triangulation [26, 19] and combinations of the two [27, 28] are quite common. The first problem to be addressed is constructing a description of the meshing domain's boundary in terms of discrete data segments such as line segments/triangles (advancing front method) or triangles/tetrahedrons with minimum aspect ratio (Delaunay triangulation). The data segments must be carefully chosen to distinguish between the interior and exterior of the meshing domain.

In unstructured approaches, the domain being meshed is defined by its boundaries, and typically only the domain on one-side of the boundary is actually meshed. For micro-mechanical analysis of composite materials, there exist multiple singly connected material domain entities (inclusions) and the boundaries between these different entities and the matrix medium that occupies the interstices are essentially embedded within. Since both the matrix and inclusion phases must generally be meshed, some modifications to one-sided meshing would be needed in unstructured approaches. For this reason and others, we consider below voxel-based meshing techniques that are inherently two-sided.

1.4. Precedents in voxel-based meshing

Voxel-based meshing is a three-dimensional analogue to two-dimensional pixel-based meshing. It has been previously used in automotive applications for meshing of complicated three-dimensional components for FEM analysis [29] and also in biomechanics to construct unit-cell models of trabecular bone using CT-scan data [30, 31]. Since trabecular bone is a quasi-random composite material in which the material structure cannot be described mathematically, CT-scan data is essential to generate the model. Accordingly, the resolution of material boundaries that can be achieved in unit-cell models of trabecular bone is limited by the resolution of the CT-scan data. Unlike quasi-random trabecular bone, many composites have a nearly periodic or nearly regular material structure that can be described mathematically. Therefore, when voxel methods are applied to such composites, infinite material resolution is theoretically possible, although realistically, the size of the FEM analysis problem imposes a limit on material resolution. In this study a variety of relatively simple voxel meshing algorithms are developed and applied to the analysis of fibrous textile composites.

1.5. Voxel-based meshing and adaptive refinement

To achieve convergence in both material resolution and computed effective material properties, adaptive mesh refinement methods are employed with the proposed voxel/pixel techniques. In recent years, much work has been done with error estimators and their usage in adaptive mesh refinement. Zienkiewicz and Zhu [32] developed a practical *a-posteriori* error estimator based on stress averaging or the so-called L_2 projection method. Since the error estimator uses a smoothed stress field from the FE approximation rather than the true stress field, its efficacy depends on the chosen manner of smoothing. It is now recognized [33] that the L_2 projection method is asymptotically correct only when the mesh discretization error is modest. Recently, new smoothing techniques like Super-convergent Patch Recovery [34] and Recovery by Equilibrium Patches [35] have been developed to enhance the precision in estimating the discretization error. In this work, we introduce a new adaptive mesh-refinement technique that is natural for usage in voxel/pixel meshing techniques. It is based on sensitivity analysis of

virtual strain energy in each element of the model to mesh refinement, which can also be interpreted as an indicator of weak-form dis-equilibrium when the stress field is projected onto smooth basis functions within the interior of individual unrefined elements.

1.6. Procedures in voxel-based meshing and overview

The basic procedures of voxel-based meshing and unit-cell analysis are as follows:

1. Develop suitable mathematical descriptions of the individual inclusion objects to be used in the composite.
2. Generate a uniform initial mesh of quadrilaterals (pixels, in two-dimensions) or hexahedra (voxels in three-dimensions) for the unit-cell model.
3. Using sampling techniques, determine the fraction of each voxel that lies within each of the inclusion objects.
4. Apply rules of mixtures to obtain upper and lower bounds on the mechanical properties of 'mixed' voxels that contain more than one material.
5. For a given macroscopic stress or strain loading on the unit-cell, solve for the heterogeneous displacement fields that lead to satisfaction of the weak form of the local equilibrium conditions.
6. Compute the corresponding macroscopic stresses and strains in the composite.
7. Note the discrepancy between 'bounds' on homogenized properties from unit-cell analysis associated with alternative treatments of mixed voxels.
8. If discrepancy between bounds is unacceptable, or if results have not yet converged, apply successive quadtree/octree mesh refinements and repeat above procedures until the discrepancies are sufficiently small.

These voxel-based meshing techniques for composite materials are investigated here because: (1) voxels do not suffer from shape distortion under repetitive refinement; (2) the methods are inherently two-sided; and (3) unit-cell models can be developed with negligible volume-fraction errors. In the sequel of this paper, the basic methodology of unit-cell homogenization of periodic composites is briefly reviewed. Then the concepts, issues and techniques employed in voxel-based meshing of textile composites are presented. The convergence characteristics of the voxel-based methods are then presented for two-dimensional aligned fibre composite examples, and then for three-dimensional fibrous textile composites.

2. COMPUTATIONAL HOMOGENIZATION

2.1. Notation for periodic composites

The general classes of composite materials under consideration (Figure 4) are assumed to occupy a domain Ω_b in \mathbb{R}^3 and to exhibit periodic microstructures in that a single wavelength sub-domain forms a unit-cell $\Omega_s = \prod_{i=1}^3]0, \lambda_i[$, where λ_i are the dimensions of a specific composite's unit-cell. Regions of the structural domain Ω_b are formed by repetition of a specific composite's unit-cell Ω_s . Material points in an undeformed unit-cell Ω_s are referenced by local material scale co-ordinates \mathbf{X} , while the spatial material scale co-ordinates \mathbf{x} locate material points in a deformed cell. Similarly, the Lagrangian and spatial macro-scale co-ordinates of

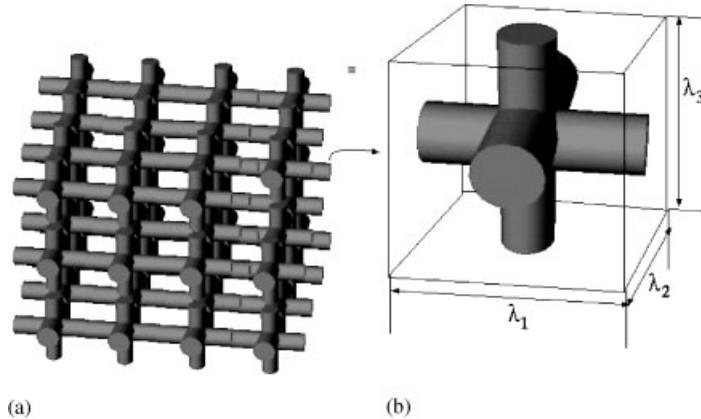


Figure 4. Schematic of a medium with periodic micro-structure: (a) structural domain Ω_b ; and (b) unit-cell Ω_s .

a point in the larger domain Ω_b are denoted by \mathbf{Y} and \mathbf{y} . The undeformed and deformed locations of a material point are related by displacement vectors \mathbf{u} such that

$$\mathbf{x} = \mathbf{X} + \mathbf{u}(\mathbf{X}) \quad \text{material scale}; \quad \mathbf{y} = \mathbf{Y} + \mathbf{u}(\mathbf{Y}) \quad \text{macroscale} \quad (1)$$

Material scale co-ordinates \mathbf{X} and \mathbf{x} are used in solving unit-cell homogenization problems, whereas macroscale co-ordinates \mathbf{Y} and \mathbf{y} are used in solving structural analysis problems with the effective medium approximation. (In this work attention is confined solely to perform unit-cell analysis at the material scale.)

At the material scale, and at the macroscale, the respective deformation gradients are defined as follows:

$$\mathbf{F}(\mathbf{X}) = \frac{\partial \mathbf{x}}{\partial \mathbf{X}}, \quad \Phi(\mathbf{Y}) = \frac{\partial \mathbf{y}}{\partial \mathbf{Y}} \quad (2)$$

If the length scale of variation l associated with deformation Φ in a composite structural element Ω_b is much larger than the wavelengths of the material microstructure ($l/\|\lambda\| \gg 1$) then the deformation and stress fields will be λ -periodic on the material scale such that

$$\mathbf{F}(X_I) = \mathbf{F}(X_I + n\lambda_I), \quad \sigma(X_I) = \sigma(X_I + n\lambda_I), \quad (I = 1, 2, 3) \quad (3)$$

where n is a small, but otherwise arbitrary integer, $\sigma(\mathbf{X})$ is the local Piola–Kirchoff stress tensor, and $\mathbf{F}(\mathbf{X})$ is the local deformation gradient. Periodicity of one field (deformation or stress) typically assures periodicity of the other. The periodicity of the stress and deformation field on the material-scale suggests that they admit additive decompositions into macroscopic contributions (\mathbf{S}, Φ) , which are constant on the micro-scale, and purely oscillatory contributions $[\sigma^*(\mathbf{X}), \mathbf{F}^*(\mathbf{X})]$, which can vary quite significantly on the unit-cell length scale:

$$\sigma(\mathbf{X}) = \mathbf{S} + \sigma^*(\mathbf{X}), \quad \mathbf{F}(\mathbf{X}) = \Phi + \mathbf{F}^*(\mathbf{X}) \quad (4)$$

The homogenized macroscopic tensors $(\mathbf{S}, \mathbf{\Phi})$ are defined as volume averages of $(\boldsymbol{\sigma}, \mathbf{F})$ over the undeformed unit-cell, and have the mathematical definitions:

$$\mathbf{S} = \langle \boldsymbol{\sigma} \rangle = \frac{1}{V} \int_{\Omega_s} \boldsymbol{\sigma} \, d\Omega_s, \quad \mathbf{\Phi} = \langle \mathbf{F} \rangle = \frac{1}{V} \int_{\Omega_s} \mathbf{F} \, d\Omega_x \quad (5)$$

where V is the initial volume of the unit-cell $\bar{\Omega}_s$. Combining (4) and (5) it can be shown that

$$\langle \boldsymbol{\sigma}^* \rangle = \mathbf{0}, \quad \langle \mathbf{F}^* \rangle = \mathbf{0} \quad (6)$$

With regard to the constant-periodic decomposition of the deformation field, it can also be shown by integrating (4)₂ that the displacement field over an individual unit-cell must admit a linear-periodic decomposition:

$$\mathbf{u}(\mathbf{X}) = [\mathbf{\Phi} - \mathbf{1}] \cdot \mathbf{X} + \mathbf{u}_{\text{per}}^*(\mathbf{X}) \quad (7)$$

where any constant or rigid body contributions to the displacement field \mathbf{u} are eliminated by proper restraints on the unit-cell domain, and imposing $\mathbf{\Phi} = \mathbf{R} \cdot \mathbf{U}$, where \mathbf{R} is a rotation operator, and \mathbf{U} is the right stretch tensor. The periodic displacement field contribution $\mathbf{u}_{\text{per}}^*(\mathbf{X})$ permits a heterogeneous (non-uniform) deformation field within the unit-cell domain $\bar{\Omega}_s$, satisfying local stress equilibrium $\sigma_{ll,l} = 0$.

2.2. The effective elastic stiffness tensor \mathbf{C}

The deformation-controlled homogenization problem involves imposing macroscopic deformation $\mathbf{\Phi}$ (or imposing $\mathbf{u} = [\mathbf{\Phi} - \mathbf{1}] \cdot \mathbf{X}$) on the cell of a composite and computing the resulting periodic minimal energy displacement field $\mathbf{u}^* : \bar{\Omega}_s \Rightarrow \mathfrak{R}^3$ satisfying the equilibrium equation ($\nabla \cdot \boldsymbol{\sigma} = 0$ on $\bar{\Omega}_s$). From the equilibrium local displacement field $\mathbf{u}(\mathbf{X})$, strain–displacement relations, and constitutive behaviours of the materials in the composites, the corresponding stress field $\boldsymbol{\sigma}(\mathbf{X})$ and its macroscopic counterpart \mathbf{S} can be computed. For the special case of a composite material being comprised of linearly elastic and perfectly bonded constituents, the functional form of the effective constitutive behaviour for the composite will necessarily be that of linear, anisotropic elasticity. Using deformation-controlled homogenization, the components of a composite's linearized anisotropic stiffness tensor \mathbf{C} can be calculated straightforwardly as follows: given a state of macroscopic deformation $\mathbf{\Phi}$ imposed on the unit-cell. The associated macroscopic strain (Green–Lagrange) is computed as

$$E_{KL} = \frac{1}{2} [\Phi_{kK} \Phi_{kL} - \delta_{KL}] \quad (8)$$

A weak form of the local equilibrium equation $\nabla \cdot \boldsymbol{\sigma}^{(KL)} = \mathbf{0}$, $\forall \mathbf{X} \in \bar{\Omega}_s$ is then solved (subject to periodic boundary conditions) for the unknown part of the material-scale displacement field $[\mathbf{u}^*(\mathbf{X})]^{(KL)}$ on the domain of the unit-cell $\bar{\Omega}_s$. In solving a weak form of the local equilibrium state equation over $\bar{\Omega}_s$ via the finite element method, the local strain and stress fields are explicitly computed. By the *average stress theorem* for composites subjected to macroscopically homogeneous deformation loading Φ_{kK} , the macroscopic stress tensor $S_{IJ}^{(kK)}$ corresponding to the applied macroscopic strain E_{KL} is simply the volume average of the

local stress field.

$$S_{IJ}^{(KL)} = \frac{1}{V} \int_{\Omega_s} \sigma_{IJ}^{(KL)} d\Omega_s \quad (9)$$

Accordingly, by solving the unit-cell problem for all appropriate and necessary permutations of the K and L indices of the macro strain tensor \mathbf{E} . The coefficients of the effective medium elastic stiffness tensor for the composite are simply

$$C_{IJKL} = \frac{S_{IJ}^{(KL)}}{E_{KL}} \quad (10)$$

where δ_{KL} is the Kronecker delta function.

3. BASIC CONCEPTS OF VOXEL-BASED MESHING

3.1. Mathematical description of reinforcement

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 Ω_R occupied by each 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:

$$\Omega_R = \{\mathbf{X} | F^R(\mathbf{X}) < 0\} \quad (11)$$

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

In voxel-based meshing, the unit-cell domain Ω_s of the composite is discretized into a finite number of N non-intersecting volume elements Ω_e or voxels. The fraction of each voxel that is occupied by a given reinforcing object can be determined to any desired level of precision by sampling material points from the voxel either randomly or systematically (Figure 5). Thus for each voxel (or 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 voxel resolution of the unit-cell ($N \rightarrow \infty$) the reinforcing object volume fraction associated with each voxel tends to discrete values of 0 and 1. However, for practical mesh resolutions, many of the voxels will typically contain intermediate reinforcement volume fractions in the range $\phi_n \in (0, 1)$.

In order to compute the volume fraction of each voxel that is occupied by a given material object m , the voxel is subdivided into L non-intersecting volume sub-elements. At the centroid of each sub-element, the material indicator function (11) is then evaluated for the m th material component. If the centroid lies within material phase m , it is assumed that the entire volume of the sub-element lies within material phase m . Accordingly, the volume fraction of material m in the e th voxel comprising the unit-cell is computed as follows:

$$\phi_e^m = \frac{\sum_{l=1}^L \{V_e^l | F^m(\mathbf{\kappa}_e^l) < 0\}}{\sum_{l=1}^L V_e^l} \quad (12)$$

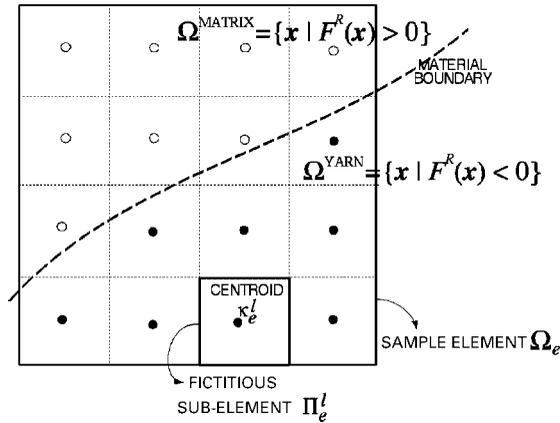


Figure 5. Schematic showing subdivision of a voxel into $L=16$ sub-elements to facilitate computation of the material volume fractions and material directors.

where V_e^l is the volume of the e th voxel's l th sub-element Π_e^l , and κ_e^l is the same sub-element's centroid. In the limit as $L \rightarrow \infty$ the computation of material volume fractions in each voxel tend to their exact values. In all of the computations of this effort, a value of $L = 100$ is used in two-dimensions, and $L = 1000$ in three-dimensions.

3.2. Material hierarchy and averaged material directors

Another important issue when analysing 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 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 m partially or even fully occupies a given voxel e , the average orientation of the material phase m in the voxel is needed.

The director for a transversely isotropic material at a point \mathbf{X} is a unit vector aligned with that material's local axial direction. To compute the averaged material director \mathbf{d}_e^m for transversely isotropic phase m occupying voxel e , the following algorithm is used:

$$\mathbf{d}_e^m = \frac{\sum_{l=1}^L \{\mathbf{d}^m(\kappa_e^l) * V_e^l | F^m(\kappa_e^l) < 0\}}{\varphi_e^m \sum_{l=1}^L V_e^l} \tag{13}$$

Above, $\mathbf{d}^m(\kappa_e^l)$ is the director of material m evaluated at points κ_e^l known to be lying within material phase m .

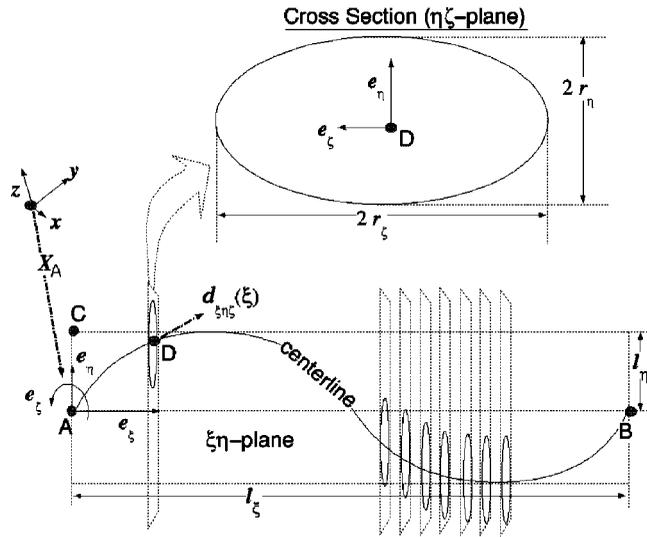


Figure 6. Geometric description of a woven yarn and its cross-section.

3.3. Geometry description

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 centerline curve is sinusoidal and so that cross-sections remain orthogonal to the original axis of the cylinder (Figure 6). 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, 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

$$\begin{aligned} \Omega^{\text{YARN}} &= \{\mathbf{X} \mid F^{\text{YARN}}(\mathbf{X}) < 0\} \\ &= \{\xi \mid \hat{F}^{\text{YARN}}(\xi) < 0\} \end{aligned} \tag{14a}$$

where

$$\hat{F}^{\text{YARN}}(\xi) = \left(\frac{\eta - C_L(\xi)}{r_\eta}\right)^2 + \left(\frac{\xi}{r_\zeta}\right)^2 \text{ and centreline } C_L(\xi) = l_\eta \sin\left(\frac{2\pi\xi}{l_\zeta}\right) \tag{14b}$$

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

$$\xi = \mathbf{Q} \cdot (\mathbf{X} - \mathbf{X}_A) \tag{15}$$

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

$$\hat{\mathbf{d}}(\xi) = \frac{1}{\sqrt{1 + \left(\frac{2\pi l_\eta}{l_\xi}\right) \cos\left(\frac{2\pi \zeta}{l_\xi}\right)^2}} \left(1, \left(\frac{2\pi l_\eta}{l_\xi}\right) \cos\left(\frac{2\pi \zeta}{l_\xi}\right), 0 \right) \quad (16)$$

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

$$\mathbf{d}(\mathbf{X}) = \mathbf{Q}^T \cdot \hat{\mathbf{d}}(\xi) \quad (17)$$

3.4. Effective properties of mixed voxels

In cases, where individual voxels contain a mixture of both reinforcement and matrix material, the associated mechanical properties of the voxel can be determined either from: (a) usage of combinations of the classical Voigt–Reuss-type mixing rules; or (b) constructing unit-cell models of impure voxels and employing homogenization to compute effective properties. Experimentation by the authors with both options has shown that option (a) is more efficient and effective than option (b). Thus, in the subsequent development, the effective elastic properties \mathbf{C}_e^* in a given voxel e are determined as follows:

$$\mathbf{C} = \alpha \mathbf{C}^{\text{Voigt}} + (1 - \alpha) \mathbf{C}^{\text{Reuss}} \quad (18)$$

in which

$$\mathbf{C}^{\text{Voigt}} = \phi_{\text{yarn}} \mathbf{C}^{\text{yarn}} + \phi_{\text{matrix}} \mathbf{C}^{\text{Matrix}} \quad (19a)$$

$$\mathbf{C}^{\text{Reuss}} = [\phi_{\text{yarn}} \mathbf{S}^{\text{yarn}} + \phi_{\text{matrix}} \mathbf{S}^{\text{Matrix}}]^{-1} \quad (19b)$$

Above, \mathbf{C} and \mathbf{S} denote, respectively, elastic stiffness and compliance tensors for the yarn and matrix phases. In Equation (18), $\alpha = 1$ gives the Voigt rule that provides an upper bound on the possible stiffness of any mixed voxel, and $\alpha = 0$ provides the Reuss upper bound on the compliance of any mixed voxel. While applied here for strictly linear elastic material behaviours, the Voigt and Reuss mixing rules can easily be used with a wide variety of material behaviours, as was demonstrated by Swan and Kosaka [37] in the context of structural topology optimization.

In this work, the polymer matrix phase is assumed to be isotropic, and the yarn phase transversely isotropic. Thus when computing the Voigt ($\mathbf{C}_e^{\text{Voigt}}$) or Reuss ($\mathbf{C}_e^{\text{Reuss}}$) stiffness of a given voxel, the yarn stiffness \mathbf{C}^{yarn} or compliance \mathbf{S}^{yarn} in (19) are computed based on the averaged material director $\bar{\mathbf{d}}_e^{\text{yarn}}$ which is obtained as described in Section 3.3.

3.5. Mesh refinement options

3.5.1. Uniform and selective refinement. In voxel-based meshing, increasing mesh refinement leads to both: (1) enhanced resolution of the material scale stress and strain fields; and (2) enhanced resolution of material arrangements. Uniform refinement of all voxels throughout a unit-cell model Ω_s increases both (1) and (2), but at a higher computational expense. More feasible is the *selective refinement* of those voxels that contain a mixture of materials,

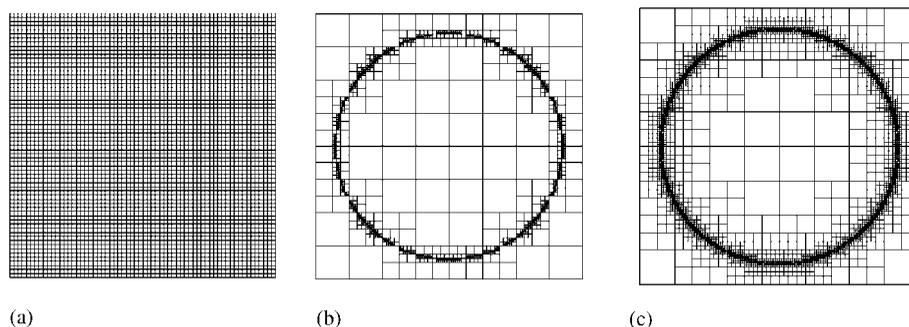


Figure 7. Meshes for: (a) uniform refinement; (b) selective refinement; (c) constrained selective refinement.

leading to a mesh with higher resolution in the vicinity of material boundaries, relatively coarse elements in the free-field, and a large spectrum of element sizes. A two-dimensional representation of an aligned fibre composite with both uniform and selective mesh refinement is shown in Figure 7. With *selective refinement*, voxels are refined only to better resolve material boundaries, and not to minimize error estimates associated with the displacement solution fields as is commonly done in adaptive h -refinement techniques. Thus, the criterion for determining whether a given voxel should be refined is $[\phi(1 - \phi) > C]$, where ϕ can be taken as the volume fraction of the matrix phase, and where C is a small positive parameter chosen to reflect the minimal acceptable level of voxel purity. In two-dimensions, planar elements undergoing *selective refinement* are simply subdivided into four planar sub-elements, while in three-dimension, hexahedral brick elements are subdivided into eight hexahedral sub-elements upon refinement.

3.5.2. Constrained selective refinement. Pure *selective refinement* of pixels/voxels to achieve high resolution of material boundaries has the pitfall that it can result in highly inaccurate numerical solutions. This is most likely to happen when highly refined elements are adjacent to much larger elements that have undergone far fewer generations of refinement (Figure 7(b)). Since they are adjacent to a larger element, the sequence of highly refined elements cannot capture local fluctuations in stresses and strains, since many of their associated degrees of freedom are constrained in order to maintain displacement continuity. To avoid this potential problem, when employing selective refinement of pixels/voxels, it is best to impose the constraint that adjoining pixels/voxels be no more than one generation apart in terms of their level of refinement. In this work, mesh refinement to achieve high resolution of material interfaces but with the constraint that adjacent elements must be no more than one generation apart is termed *constrained-selective refinement*. It should be observed (Figure 7(c)) that with *constrained selective refinement*, as elements in the vicinity of material interfaces are refined, elements in the free-field away from material boundaries are also refined. Practical details on the implementation of constrained selective mesh refinement are provided in Section 4.

3.5.3. Adaptive mesh refinement. As described above, *selective refinement* is based strictly on refinement of those voxels/pixels lying along material interfaces. With *constrained*

selective refinement, even those voxels removed from material interfaces will be refined as those voxels in the vicinity of material interfaces are increasingly refined. A question that needs to be addressed is whether *constrained selective* methods of refinement have performance characteristics similar to that of the so-called adaptive mesh refinement techniques that refine those regions having the highest energy errors.

The global equilibrium equation for the unit-cell can be written in variational form as

$$\delta\Pi = \delta W^{\text{int}} - \delta W^{\text{ext}} \quad (20a)$$

$$= \int_{\Omega} \delta\boldsymbol{\varepsilon}^{\text{T}} : \boldsymbol{\sigma} \, d\Omega - \int_{\Omega} \delta\mathbf{u}^{\text{T}} \mathbf{b} \, d\Omega - \int_{\Omega} \delta\mathbf{u}^{\text{T}} t \, d\Gamma \quad (20b)$$

$$= 0 \quad (20c)$$

where the first integral term in (20b) denotes the internal virtual work in the domain of analysis. A potentially useful criterion for mesh refinement is the sensitivity of virtual strain energy in each element or voxel to potential refinement. This sensitivity can be evaluated on an element-by-element basis and then ranked, giving a priority criterion for each voxel or element to be refined.

In two-dimensions, the isolated refinement of a given four-node bilinear quadrilateral finite element results in the introduction of five new nodes. Of these five nodes, only the center node will be free, while the remaining four on the boundaries will be enslaved to their pre-existing master nodes as described in Section 4. Consequently, the additional virtual strain in the element stemming from refinement can be attributed entirely to the additional degree of freedom at the centre of the refined element, or $\delta\boldsymbol{\varepsilon}(\mathbf{X}) = \mathbf{B}_c(\mathbf{X}) \cdot \delta\mathbf{u}_c$ where $\mathbf{B}_c(\mathbf{X})$ is the strain–displacement matrix associated with the introduced centre node, and $\delta\mathbf{u}_c$ is the virtual displacement associated with the introduced centre node. The shape function associated with the free centre node will typically be a bubble-function that vanishes on the boundaries of the element. (As an example, for a two-dimensional four-node bilinear element, the centre node function in the parent domain would be of the form $N^C(\mathbf{r}) = (1 - r^2)(1 - s^2)$.) The virtual strain energy in the e th element associated with refinement is simply:

$$(\delta W^{\text{int}})_e = \int_{\Omega_e} \delta\mathbf{u}_c \mathbf{B}_c^{\text{T}} \boldsymbol{\sigma} \, d\Omega_e \quad (21)$$

Taking the sensitivity of this virtual strain energy with respect to the virtual nodal displacement $\delta\mathbf{u}_c$, yields an unbalanced force on the hypothetically introduced centre node c associated with the current stress field $\boldsymbol{\sigma}$ within the e th element:

$$(\mathbf{f}_c^{\text{int}})_e = \frac{d(\delta W^{\text{int}})_e}{d(\delta\mathbf{u}_c)} = \int_{\Omega_e} \mathbf{B}_c^{\text{T}} \boldsymbol{\sigma} \, d\Omega_e \quad (22)$$

The norm of such an unbalanced force for each element in the mesh can then be used as a prioritization criterion for refinement of all elements in the mesh. It should be realized that the unbalanced force magnitude as expressed in (22) is not an error estimator for the stress within an element, but is instead an error indicator for the strain energy within an element. By virtue of its integral form, this criterion will give higher refinement priority to large elements,

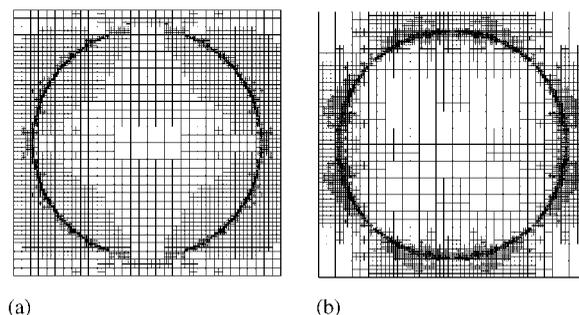


Figure 8. Adaptively refined, undeformed, voxel-based meshes. In both computations, mixed elements were modelled using the Voigt (iso-strain) mixing rule: (a) extensional test; and (b) shear test.

as opposed to smaller elements. Examples of two-dimensional unit-cell meshes after many generations of adaptive refinement are shown in Figure 8.

It is worth noting here that the proposed expression [Equation (22)] for sensitivity of virtual strain energy to mesh refinement is not exact. When a mixed voxel or element is refined, the modelled distribution of material within the space covered by that voxel will typically change. Such spatial re-distribution of materials within a voxel undergoing refinement has not been accounted for in Equation (22).

4. DISPLACEMENT CONTINUITY AND HANGING NODES

4.1. Description of hanging nodes

An advantage associated with pixel/voxel refinement using edge bisection is that the element aspect ratio is kept unchanged through countless generations of refinement. Thus if voxels in the original mesh have favourable shapes and aspect ratios, the favourable shape and aspect ratio are preserved throughout all successive refinements. In contrast, however, conventional unstructured meshing techniques using triangular/tetrahedral elements can suffer considerable difficulties as element shape and aspect ratios tend to degenerate with repeated generations of refinement. Although pixel/voxel techniques do not have problems pertaining to poor element aspect ratios under repeated refinement, the issue of the so-called hanging nodes is a potential problem requiring attention. In 2D, hanging nodes are those introduced along edges separating elements that have undergone different levels of successive refinement. In 3D, hanging nodes are those introduced along edges and faces separating elements having undergone different levels of refinement. Unless their degrees of freedom are constrained, such nodes induce displacement field discontinuity. Furthermore, in computational homogenization of periodic media where periodic boundary conditions must be imposed, additional care must be devoted to treatment of hanging nodes on the unit-cell's boundaries.

Edges are simple data structures associated with contiguous nodes of a given pixel or voxel element. All pixel elements have four edges and all voxel elements will have 12 edges. Whenever a pixel or voxel element is subdivided during mesh refinement, all of its associated edge data is erased, and edges of the four/eight descendant elements are generated instead.

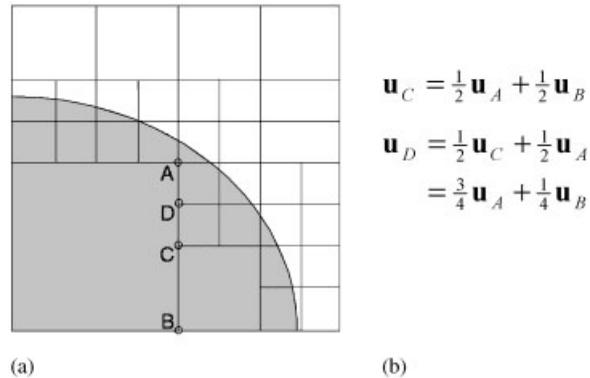


Figure 9. Schematic showing the need for nodal enslavement to enforce displacement continuity in selectively refined voxel or pixel meshes: (a) Hierarchical mesh with enslaved nodes; and (b) dependencies among master and slave nodes.

In the unit-cell model, any interior edge (i.e. an edge that is not on an exterior boundary of the model) is shared among two to four adjoining elements. When not every one of these adjoining elements is refined, the new node introduced on the interior edge is said to be a hanging node. To avoid discontinuity of the displacement field along this edge, the displacements of the hanging node should be enslaved to those of the end nodes on the edge. In this sense, the end nodes on the edge are said to be master nodes, while the hanging node is said to be a slave to the masters (Figure 9). When the other adjoining elements are subsequently refined, the hanging node no longer needs to be enslaved, and it becomes a regular free node. This is accomplished by using a few relatively simple data structures (see Box 1): *edge slave sets*

Box 1. Definition of *edge slave sets* and *potential master sets*.

(1) *edge slave sets* $\Omega(\bullet, \bullet)$;

$$\Omega(M_1, M_2) = \begin{cases} S, & \text{when a new node } S \text{ is introduced at the midpoint of an edge} \\ & \text{or diagonal with end-nodes } M_1 \text{ and } M_2. \\ 0, & \text{when the edge defined by end nodes } M_1 \text{ and } M_2 \text{ has no} \\ & \text{midpoint node.} \end{cases}$$

(2) *potential master sets* $\Psi(\bullet)$;

Each newly introduced node S , has a set of potential masters

$$\Psi(S) = \{M_1, \dots, M_N\}$$

where the members are pre-existing nodes at the end-points or corner-points of the edge or face along which S was introduced. Note that the set has two members when S is introduced along an edge, and four members when introduced on a face (see Figure 10).

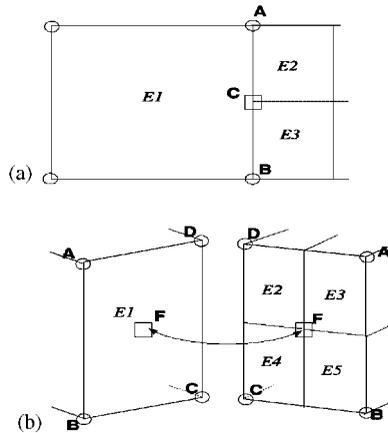


Figure 10. (a) A node 'C' hanging along an edge; and (b) a node 'F' hanging on an element face.

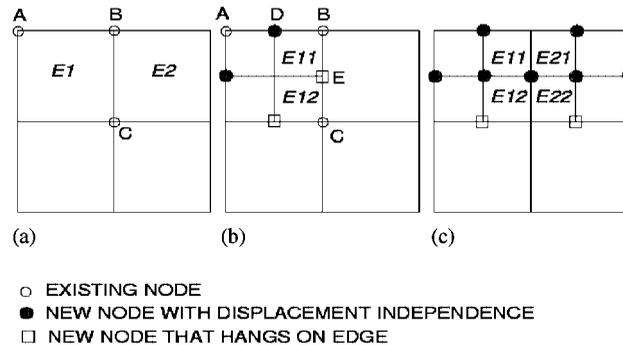


Figure 11. Illustration of hanging nodes with local mesh refinement: (a) edge BC is shared by two adjoining elements E1 and E2; (b) after refinement of E1, the introduced node E hangs along edge BC of E2, while node D is free since edge AB is on periphery of model and does not belong to any elements beyond E1; and (c) upon subsequent refinement of element E2, node E is set free.

and *potential master sets*. Each existing edge in the mesh has an *edge slave set* and each newly introduced node *S* has a set of *potential masters*.

During hierarchical quadtree or octree mesh refinement processes, when an existing pixel or voxel element is refined, new nodes can be introduced at mid-points on existing edges and faces, or at the centre of the element interior (Figure 10). When a node is considered for introduction along a given edge (Figure 11), the *edge slave set* is first checked. If that edge slave set is already occupied, then no node need be introduced along the edge considered. However, if that edge slave set is empty, then the new node is indeed introduced. The identifying numbers of the nodes defining the edge are then placed into the new node's *potential masters* set. Similarly, when a new node is considered for introduction on an element face, the edges for the two diagonals defining that face are first checked. If either of these diagonals' *edge slave sets* have a member, then the node need not actually be introduced. However, when both *edge slave sets* are empty, then the new node is introduced. In addition

the four nodes defining the two diagonal edges are then added to the new node's potential masters set.

Once refinement of pixels/voxels in the mesh is complete, it is then necessary to determine which nodes are still *hanging* and should thus have their degrees of freedom enslaved to their masters. In the course of subdividing pixels and voxels, many elements and edges are necessarily phased out and replaced by their children. If a given edge still exists in the refined mesh and if it has a node in its *edge slave set*, then that node should be enslaved to its masters.

4.2. Hierarchy of nodal enslavements

After multiple local mesh refinements, the nodal slave–master relations can form a hierarchical structure in which one node is enslaved to masters that may in turn be enslaved to other masters themselves. To ensure full compatibility of the displacement field in unit-cell FEM analysis, the displacement dependencies between masters and slaves must be resolved. Though the hierarchical relations can make the displacement dependencies complicated, the descending order of slave–master node numbers can facilitate resolution of the difficulty. One practice that helps in this regard is that any new node introduced during mesh refinement is given a number higher than that of all pre-existing nodes. Consequently, a newly introduced node can be potentially enslaved only to those nodes with smaller identifying numbers. The specific algorithm used to resolve hierarchical displacement dependencies in this work is presented in Box 2.

While both element count and node count increase quite drastically even with selective, constrained selective, and adaptive mesh refinement, the analysis problem size as determined by the degrees of freedom count actually increases more modestly than with uniform refinement due to constraints upon newly introduced nodes. This leads to comparative economy in unit-cell computations employing selective, constrained selective, or adaptive refinement compared with uniform mesh refinement.

4.3. Periodic boundary conditions and mesh refinement

4.3.1. Virtual nodal pair groups. In unit-cell computations, periodic boundary conditions are applied to those nodes on the exterior faces of the model. This creates a need for displacement compatibility on opposing exterior boundaries of the unit-cell (Figure 12) and this gives rise to additional concerns when adaptive or selective mesh refinement is utilized. For example in three dimensions, since all eight nodes at the vertices of the unit-cell occupy the same relative location within the unit-cell (Figure 13), they must all experience the same periodic displacement u^* . These eight nodes can be associated with each other using data structures here denoted as *virtual pairs*. Each vertex node of the unit-cell is a member of seven *virtual pairs*, and thus each has seven *virtual twins*. Nodes along interior points of exterior unit-cell edges are members of three *virtual pairs*, and thus have three *virtual twins*. 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*. 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.

4.3.2. Virtual pairs and mesh refinement. After local mesh refinement, it is not uncommon for nodes on exterior boundaries of the unit-cell to be members of incomplete virtual pairs. This

Box 2. Algorithm for resolving displacement dependency

The final form of nodal displacement dependency due to hierarchical refinement will be:

$$\hat{\mathbf{u}}_S \equiv \sum_k^{\tilde{n}_S} w_{\tilde{\Psi}(S,k)}^S \cdot \hat{\mathbf{u}}_{\tilde{\Psi}(S,k)},$$

where $\hat{\mathbf{u}}_S$ is the nodal displacement of slave S, $\tilde{\Psi}(S,k)$ is the k th master among \tilde{n}_S free masters enslaving node S and $w_{\tilde{\Psi}(S,k)}^S$ is a weight for displacement dependency of node S on that master node $\tilde{\Psi}(S,k)$.

For each node $S = 1, \dots, N_{node}$ in ascending order (where N_{node} is total number of nodes)
If S is a hanging node (a slave node),

1. Identify masters from the set of potential masters

$$\Psi(S) = \{\Psi(S,i)\} \quad i = 1, \dots, n_S,$$

(where, n_S is number of potential masters for slave S, see **Box 1**.)

2. Preliminary nodal dependency can be expressed as;

$$\hat{\mathbf{u}}_S \equiv \sum_k^{n_S} \frac{1}{n_S} \cdot \hat{\mathbf{u}}_{\Psi(S,k)}$$

Note that $\Psi(S,i)$ should be distinguished from free masters $\tilde{\Psi}(S,i)$.

For $k = 1, \dots, n_S$

If $\Psi(S,k)$ is a hanging node, to say, a slave node,

Retrieve displacement dependency $\hat{\mathbf{u}}_{\Psi(S,k)}$ of slave $\Psi(S,k)$ from previous step.

Note that $\Psi(S,k) < S$ by ascending order of numbering in master-slave relation.

Then, using the retrieved displacement dependency:

$$\hat{\mathbf{u}}_{\Psi(S,k)} \equiv \sum_j^{\tilde{n}_{\Psi(S,k)}} w_{\tilde{\Psi}(\Psi(S,k),j)}^S \cdot \hat{\mathbf{u}}_{\tilde{\Psi}(\Psi(S,k),j)}$$

End If

End Loop over k

3. Then the displacement dependency $\hat{\mathbf{u}}_S$ becomes;

$$\begin{aligned} \hat{\mathbf{u}}_S &\equiv \sum_k^{n_S} \frac{1}{n_S} \cdot \hat{\mathbf{u}}_{\Psi(S,k)} = \sum_k^{n_S} \frac{1}{n_S} \cdot \left\{ \sum_j^{\tilde{n}_{\Psi(S,k)}} w_{\tilde{\Psi}(\Psi(S,k),j)}^S \cdot \hat{\mathbf{u}}_{\tilde{\Psi}(\Psi(S,k),j)} \right\} \\ &= \sum_k^{n_S} w_{\tilde{\Psi}(S,k)}^S \cdot \hat{\mathbf{u}}_{\tilde{\Psi}(S,k)}, \text{ after contraction} \end{aligned}$$

where $w_{\tilde{\Psi}(S,k)}^S$ is a sum of weights $w_{\tilde{\Psi}(\Psi(S,k),j)}^S$ related to $\hat{\mathbf{u}}_{\tilde{\Psi}(S,k)}$.

End If

End Loop over S

occurs because opposing boundaries of the unit-cell can have different mesh discretizations and there may be no virtual twin for some nodes on the opposing boundaries. When a node on the boundary of the unit-cell is a member of an incomplete virtual pair, it should be enslaved to the nodes in its set of potential masters in order to maintain displacement continuity. The

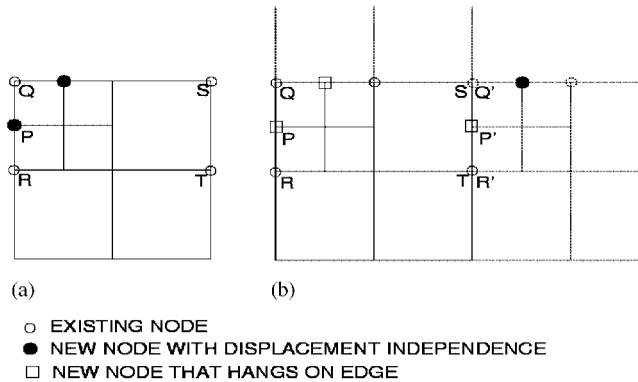


Figure 12. Illustration of hanging node issues associated with periodic boundary conditions: (a) without periodic boundary condition, node P on periphery of the model is free; and (b) with periodic boundary condition, edge QR has a virtual twin, edge ST, and thus Node P should hang on edge QR to ensure displacement compatibility.

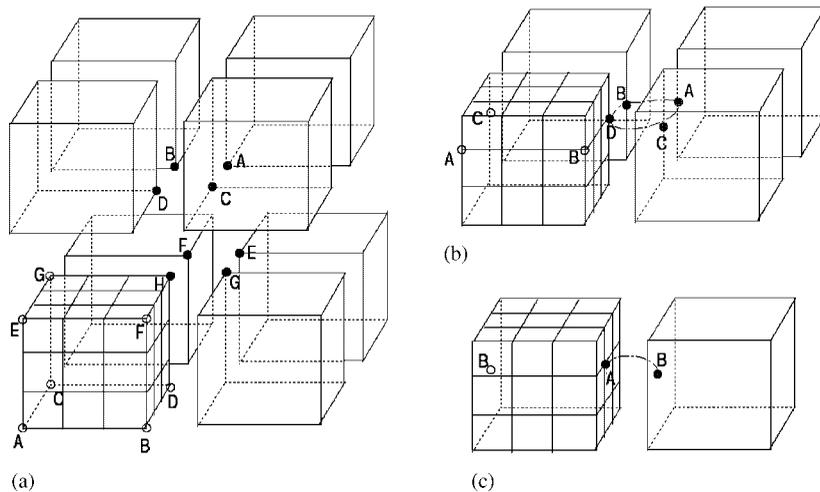


Figure 13. Nodes on the periphery of a unit-cell forming virtual pair groups when: (a) nodes are at the vertices of the unit-cell; (b) nodes lie along edges of the unit-cell; and (c) nodes lie within exterior faces.

master nodes will themselves be members of virtual pairs and thus enslaved to achieve periodic boundary conditions. If the master nodes are themselves members of incomplete virtual pairs, then the process continues in a recursive fashion.

4.4. Implementation of constrained selective refinement

Practically, constrained-selective refinement can be implemented using the data structures mentioned previously. Figure 14 shows an example in which refinement of an element leads to

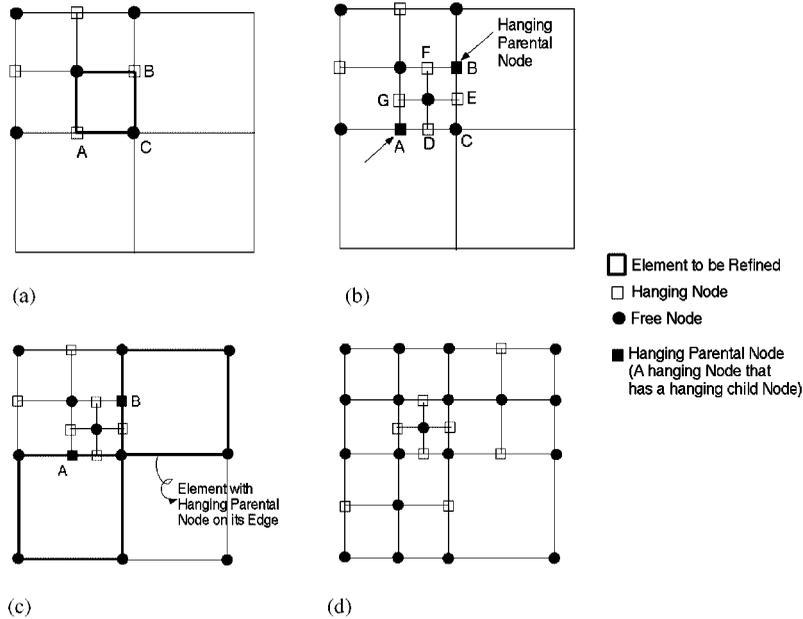


Figure 14. Details on implementation of constrained mesh refinement where adjoining elements can be no more than one generation apart: (a) an element is marked for refinement; (b) a newly introduced node E in the refined element hangs on edge BC and similarly node D hangs on edge AC; (c) elements in which parents of the new nodes are themselves hanging nodes are marked for refinement; and (d) such elements are refined, and parents (e.g. node B) of the new nodes (e.g. node E) are set free.

introduction of a new node (a *child*) that hangs along an edge in which one or more of the edge's end nodes (*parents*) are themselves hanging nodes. In such cases, the nodes just introduced are said to be *children*, and their potential masters are their *parents*. When either of the parents is themselves a hanging node along an edge, the element to which the edge on which the parent hangs is then marked for refinement. In a similar fashion, Figure 15 shows an example where newly introduced *child* nodes lying along the external boundaries of the unit-cell are enslaved to *parent* nodes that are members of *incomplete virtual pairs*. In such cases, the *grandparents* of the new child nodes (that is, the parents of the parent that is a member of an incomplete virtual pair) are identified. These *grandparents* are necessarily members of complete virtual pairs. The edge defined by the virtual twins of the *grandparents* is identified, and the element to which it belongs is then marked for refinement.

5. DEMONSTRATIVE EXAMPLES

5.1. Two-dimensional applications

The efficiency and convergence characteristics of the proposed voxel-based meshing and modelling techniques are first tested in two-dimensions on a relatively simple example. The results of unit-cell homogenization with voxel-based techniques for an aligned fibre composite are

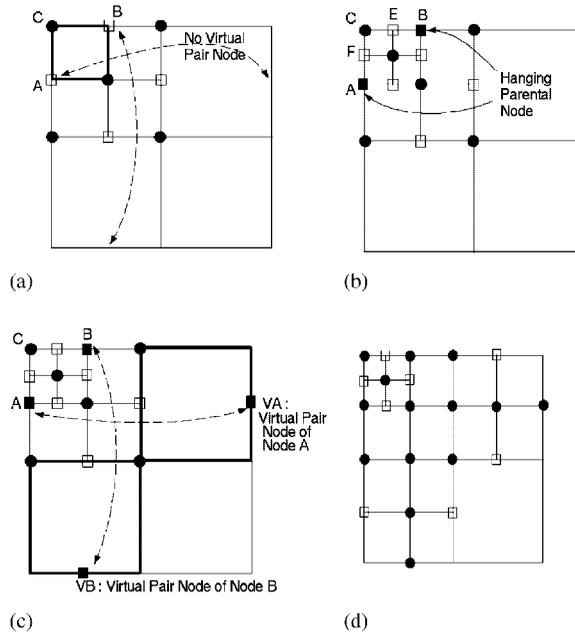


Figure 15. Constrained selective refinement with *external* hanging parents: (a) element with nodes A and B is marked for refinement. Due to periodic boundary conditions, nodes A and B are children whose parents are members of virtual pair groups; (b) the new child nodes E and F have parents (C,B) and (C,A); (c) nodes A and B are set free by adding virtual siblings nodes VA and VB; and (d) final mesh.

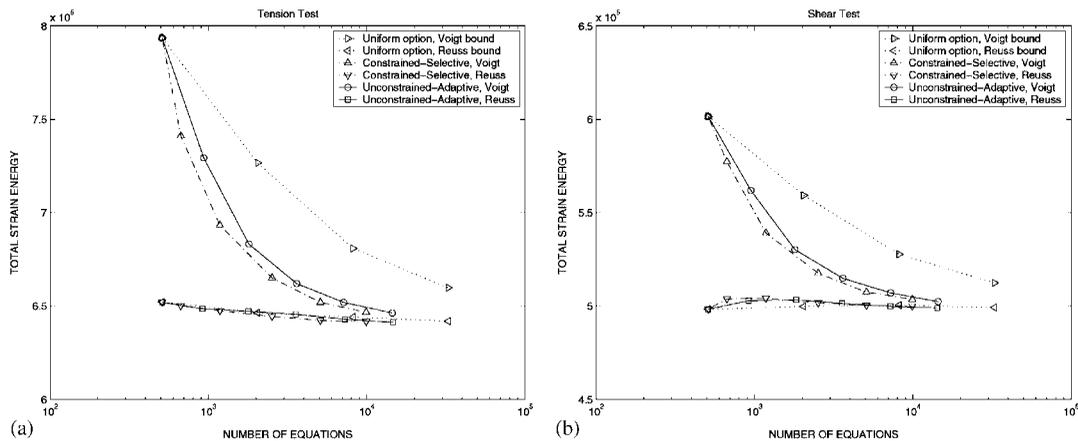


Figure 16. Convergence of computed total strain energies for: (a) transverse extension; and (b) transverse shear of the unit-cell shown in Figures 7 and 8. For each test, three sets of curves associated with the different mesh refinement options are shown (uniform, selective and adaptive refinement).

presented in Figure 16 with three different mesh refinement options: (i) uniform refinement; (ii) unconstrained selective refinement; and (iii) adaptive refinement. For each mesh discretization, the pixels/voxels containing a mixture of fibre and matrix material were modelled first using the Voigt rule of mixtures (19a), and again using the Reuss rule of mixtures (19b). The Voigt rule provided an upper bound on the stiffness of the mixed voxels, whereas the Reuss rule provides an upper bound on the compliance of mixed voxels. Hence for a given mesh discretization, the spread between the unit-cell analysis results with these two treatments reflects the remaining uncertainty in results associated with mixed voxels. Hence, the respective curves with Voigt and Reuss treatment of mixed voxels do not constitute upper and lower bounds on the true solution. This is because additional error or uncertainty in the model exists due to mesh coarseness.

The results of Figure 16 demonstrate the relative efficiency of the three mesh refinement alternatives. Uniform refinement clearly demonstrates a very slow rate of convergence. Unconstrained selective refinement, on the other hand, shows good convergence in the early stages of refinement, but with increasing mesh refinement, the rate of convergence is comparatively slow. The reason for this behaviour is clear: unconstrained selective refinement deals only with the error due to mixed voxels and not the error due to mesh coarseness. For adaptive mesh refinement based on the error indicator of (21), convergence of unit-cell analysis results is achieved without the aid of mesh size control (or the constraint that adjoining elements can be no farther than one generation of refinement apart) since the error estimator can detect overly coarse mesh regions automatically. The initial mesh used to generate the analysis results of Figure 16 had an element resolution of 16×16 or 256 bilinear continuum elements. The final meshes with uniform, unconstrained selective, constrained selective refinement were shown in Figure 7, while the final meshes based on adaptive refinement were shown in Figure 8.

It is worth noting here that while adaptive refinement worked quite well in the preceding examples, there are cases where it will not work so well. Specifically, as noted recently by Fuenmayor *et al.* [33], the success of adaptive refinement is highly sensitive to the initial mesh resolution. If the initial mesh is too coarse, the solution based on adaptive mesh refinement may never converge to the true solution. This is demonstrated in the results of Figure 17(b) where the solution based on a coarse initial mesh of 4×4 bilinear continuum elements, together with usage of the Reuss rule of mixtures, never converges to the true solution. The results of Figure 17(a) do show convergence but are somewhat counterintuitive. With adaptive mesh refinement, the Voigt treatment of mixed voxels and the Reuss treatment of mixed voxels can lead to different distributions of stress/strain in the unit-cell, and thus different error estimators. Consequently, the successive generations of adaptively refined meshes using the Voigt rule of mixtures will differ from the corresponding generations of adaptively refined meshes using the Reuss rule of mixtures. In this sense, voxel-based unit-cell analysis with purely adaptive mesh refinement techniques is more involved than the techniques based on constrained-selective refinement. This fact coupled with the observation that analysis starting with overly coarse initial meshes may not converge indicates that mesh refinement based strictly on error estimates may not be the best way to proceed.

5.2. Full three-dimensional textile composite analysis

The proposed voxel techniques are now applied to unit-cell analysis of the textile composite whose unit-cell was shown in Figure 1. The dimensions and geometrical descriptions of yarns,

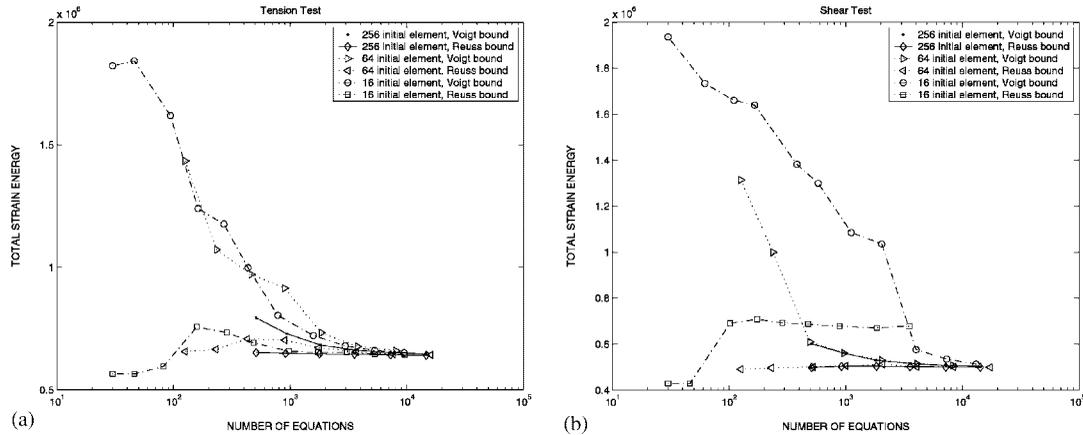


Figure 17. Convergence with adaptive mesh refinement of computed total strain energies for the unit-cell of Figures 7 and 8 under: (a) transverse extension; and (b) transverse shear. If the initial mesh is sufficiently fine, the computed strain energies converge, whereas if the initial mesh is too coarse, the strain energy does not necessarily converge.

along with material properties used are presented in Tables I and II and Figures 6 and 20. The yarns are assumed to be transversely isotropic aligned-fibre composites with 70 and 30% volume fractions of E-glass fibres and epoxy matrix, respectively. Local orientation of yarns was taken into account as described in Section 3.3. In Figure 18, the convergence of computed elastic moduli is presented based on unit-cell analysis with two mesh refinement options: (i) uniform and (ii) constrained selective approach. (The adaptive refinement option was excluded from consideration here since, as shown in two-dimensional case study, its convergence characteristics are strongly dependent on the initial coarseness of FE mesh and there is no estimating the initial coarseness in advance when dealing with complex material structure.) Of the two options considered, constrained-selective mesh refinement leads to a more rapid rate of convergence. However, even after four cycles of constrained-selective mesh refinement, and even with over 460 000 DOF in the unit-cell model, there is still approximately 1–4% error/uncertainty in the computed moduli of the composite. The final voxel mesh of the textile after four generations of constrained selective refinement is shown in Figures 19(a) and 20.

To compare the results of voxel-based unit-cell analysis with those representative of unstructured meshing [13], computed moduli are shown in Figure 18 for the mesh based on 10-noded tri-quadratic tetrahedra mesh shown in Figure 19(b). The effective composite moduli computed with the tetrahedral mesh are actually less than those computed by the Reuss lower bound and the refined voxel mesh. The reason for this slight discrepancy is the volume fraction errors that occur when going from a voxel mesh to a discrete tetrahedra-based mesh. Specifically, the yarn volume-fraction in the voxel mesh of Figure 19(a) is 54.57% with 171 968 elements, whereas that in the tetrahedral mesh of Figure 19(b) is 54.11% with 221 024 elements. The discretization error in constituent volume fractions can be virtually eliminated with voxel-based methods, whereas it can never be entirely eliminated in tetrahedra-based methods.

Table I. Dimensions for unit-cell models and relative dimensions characterizing yarn reinforcements in textile model.

Dimensions for unit-cell models								
2D model dimensions (Figures 7, 8 and 20)			3D model dimensions (Figures 6 and 20)					
L_x	L_y	R	L_x	L_y	L_z			
20 μm	20 μm	9.5 μm	4 mm	4 mm	1.5 mm			
Geometrical characterization of yarns in textile composite unit-cell								
Yarn Index	Type*	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	E	(-20,-10,0)	(20,-10,0)	(-20,-10,-1)	3	40	2.5	4.5
2	E	(-20,10,0)	(20,10,0)	(-20,10,1)	3	40	2.5	4.5
3	E	(-10,-20,0)	(-10,20,0)	(-10,-20,1)	3	40	2.5	4.5
4	E	(10,-20,0)	(10,20,0)	(10,-20,-1)	3	40	2.5	4.5
5	R	(-20,-20,-6)	(-20,-20,6)	(-20,-20,0)	0	12	5	5
6	R	(20,20,-6)	(20,20,6)	(20,-20,0)	0	12	5	5
7	R	(20,-20,-6)	(20,-20,6)	(20,20,0)	0	12	5	5
8	R	(-20,20,-6)	(-20,20,6)	(-20,-20,0)	0	12	5	5
9	R	(-20,0,-6)	(-20,0,6)	(20,0,0)	0	12	5	5
10	R	(0,-20,-6)	(0,-20,6)	(0,20,0)	0	12	5	5
11	R	(20,0,-6)	(20,0,6)	(-20,0,0)	0	12	5	5
12	R	(0,20,-6)	(0,20,6)	(0,-20,0)	0	12	5	5
13	R	(0,0,-6)	(0,0,6)	(0,-20,0)	0	12	5	5

*Type E indicates the cross-section of yarn in ζ -direction is elliptical and R indicates that rectangular.

Table II. Material properties for individual fibres and fibre-matrix yarns. All moduli are provided in units of GPa.

Young's modulus E		Poisson's ratio ν			
E-Glass Fibre	72.4	0.20			
Epoxy Matrix	2.76	0.35			
Yarn (70% fibre volume fraction, transversely-isotropic, fibre alignment with x_3 -axis)					
$E_1(E_2)$	E_3	$G_{13}(G_{23})$	G_{12}	ν_{12}	$\nu_{13}(\nu_{23})$
17.3	51.7	6.48	6.48	0.332	0.0768

6. DISCUSSION AND CONCLUSIONS

In this work, novel meshing concepts based on voxel/pixel image analysis techniques have been developed, implemented, and tested in pursuit of robust and automated techniques that can be applied to composite materials with very complicated material structures. As noted

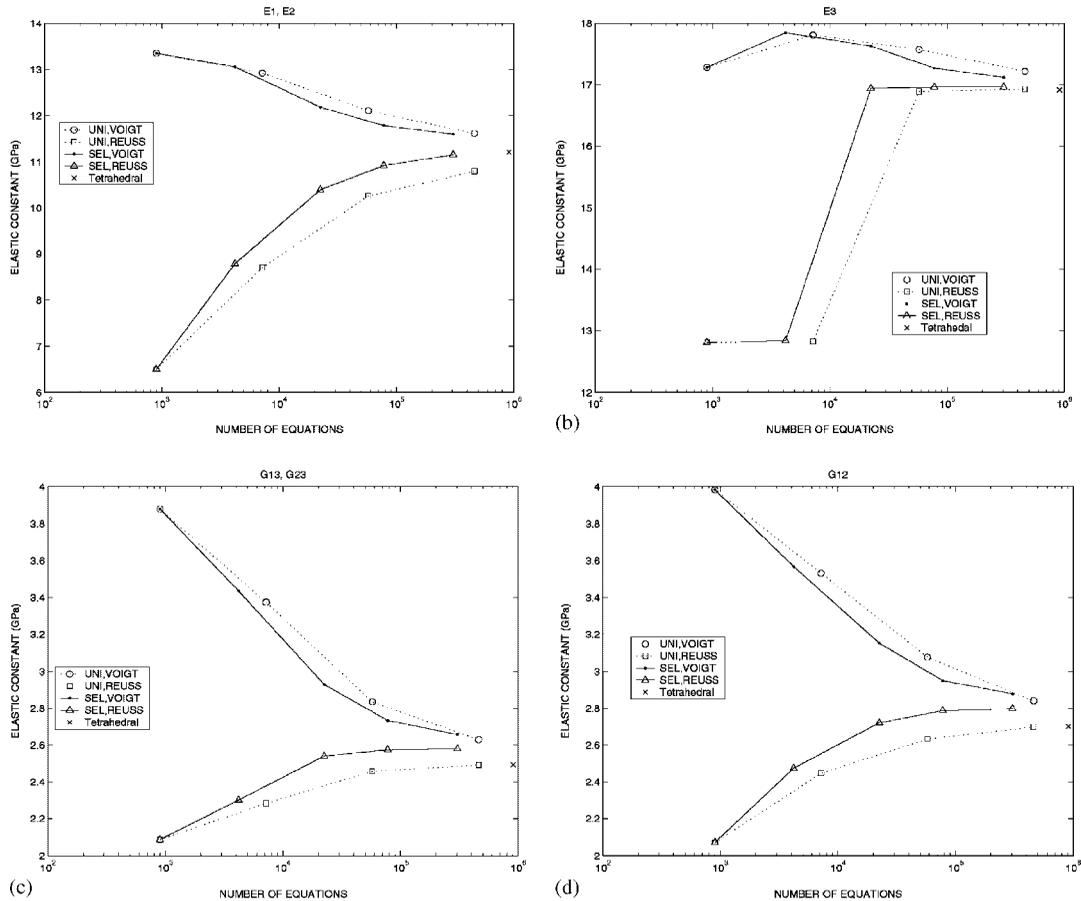


Figure 18. Convergence of computed full textile elastic moduli for various mesh refinement options: (a) shows the in-plane Young's moduli $E_{11} = E_{22}$; (b) shows the transverse Young's moduli E_{33} ; (c) shows the transverse shear moduli $E_{13} = E_{23}$; and (d) shows the in-plane shear modulus E_{12} .

in the Section 1, many of the fundamental concepts for voxel-based meshing were originally presented by Hollister *et al.* [30]. Here, however, the original techniques have been extended considerably with many practical details addressed, including: (i) mathematical descriptions of reinforcing objects such as undulating fibrous yarns; (ii) alternative local mesh refinement options and the special considerations needed in application of periodic boundary conditions; and (iii) the relative performance of the alternative mesh refinement options.

Based on these results, it can be stated that the voxel-based meshing techniques proposed here can indeed be used to rapidly and automatically generate 3D unit-cell models of composites with fairly complicated material arrangements such as textiles. It is further noted that the associated models typically involve bilinear quadrilaterals in two-dimensions and trilinear hexahedra in three-dimensions, with both types of elements having excellent shape and aspect

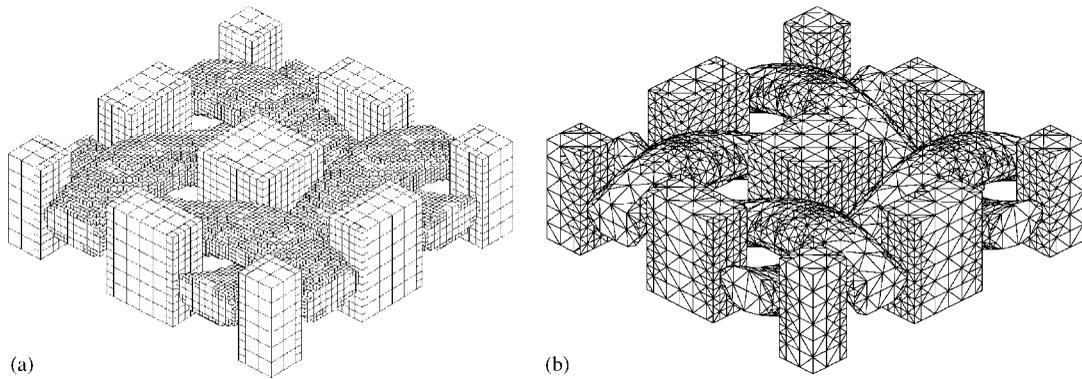


Figure 19. (a) Voxel-based mesh after four cycles of constrained selective refinement, with only elements of yarn volume fraction exceeding 50% shown; and (b) mesh of the textile composite based on 10-noded tetrahedral elements.

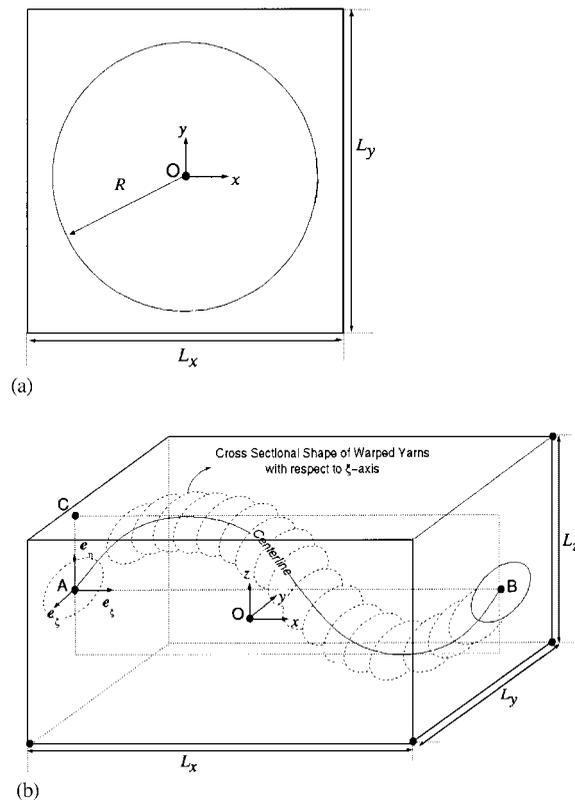


Figure 20. Dimensions of unit-cell model: (a) 2D model; and (b) 3D model: also shows three auxiliary points A, B, and C that define local $\xi\eta\zeta$ -co-ordinates.

ratio characteristics under successive generations of refinement. In addition, whereas meshing of textile composites with linear tetrahedral continuum elements has been shown by Wentdorf *et al.* [9] to result in potentially large material volume fraction errors that may outweigh typical finite element approximation errors, the proposed voxel-based techniques can capture the volume fractions of the different composite constituents to any desired level of precision even with coarse meshing. Despite these positive attributes, voxel/pixel based meshing techniques are somewhat more expensive computationally than models based on more conventional meshing techniques. Hence, there is a tradeoff with usage of voxel-based meshing and unit-cell analysis of composite materials. The human time required to generate suitable unit-cell models is greatly reduced with the proposed techniques, but the computational cost required to perform the unit-cell analysis (including convergence analysis of the computed effective material properties) is somewhat more expensive.

The added computational expense associated with voxel-based methods can potentially be mitigated using methods that take advantage of the hierarchy of the meshes employed. For example, the hierarchical parent-child data structures used in creating successive generations of increasingly refined voxel-based meshes can conceivably be used in scalable multi-grid solving techniques. Specifically, the hierarchical element and node information can be used to automatically generate the prolongation and restriction operators needed with multi-grid solvers. Furthermore, the existence and natural progression of coarse to fine voxel-based meshes alleviates the difficulty of generating coarse grid model approximations (c.f. Reference [36]).

While the proposed voxel-based meshing and unit-cell analysis techniques have both strengths and drawbacks, the hierarchical structure of the techniques makes them quite general and easy to implement. In an extension of the techniques presented here, hierarchical voxel-based meshing is used [13] for a few generations of refinement, followed by a transition to meshing with materially pure quadratic tetrahedral continuum elements. The hierarchical voxel-based meshes with their excellent aspect ratios and high refinement at material interfaces provide an excellent point of departure for meshing with pure, quadratic tetrahedral elements.

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 and 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. Dumont S, Lebon F. Wavelet-Galerkin method for periodic heterogeneous media. *Computers and Structures* 1996; **61**(1):66–65.
5. Chapman CD, Whitcomb JD. Thermally induced damage initiation and growth in plane and stain weave carbon-carbon composites. *Mechanics of Composite Materials and Structures* 2000; **7**(2):177–194.

6. Choi J, Tamma KK. Woven fabric composites—Part I: Prediction of homogenized elastic properties and micromechanical damage analysis. *International Journal for Numerical Methods in Engineering* 2001; **50**: 2285–2298.
7. Swan CC. Unit cell analysis of 3-D graphite textile-reinforced polymer matrix composites. In *Advances in Composite Materials and Mechanics*, Maji A (ed.). ASCE: New York, 1999.
8. Swan CC, Kosaka I. Homogenization-based analysis and design of composites. *Computers and Structures* 1997; **64**:603–621.
9. Wentorf RR, Collar R, Shephard MS, Fish J. Automated modelling for complex woven mesostructures. *Computer Methods in Applied Mechanics and Engineering* 1999; **172**:273–291.
10. Vandeurzen P, Ivens J, Verpoest I. A three-dimensional micromechanical analysis of woven fabric composites: Parts I & II. *Composites Science and Technology* 1996; **56**:1303–1327.
11. Tabiei A, Jiang Y. Woven fabric composite material model with material nonlinearity for the finite element simulation. *International Journal of Solids and Structures* 1999; **36**(18):2757–2771.
12. Jiang Y, Tabiei A, Simites, J. A novel micro-mechanics approach to the derivation of constitutive equations for global/local analysis of plane weave fabric composites. *Composites Science and Technology* 2000; **60**: 1825–1833.
13. Kim HJ, Swan CC. Automated meshing and unit-cell analysis of periodic composites with hierarchical quadratic tetrahedra, in review. accepted for publication in *International Journal for Numerical Methods in Engineering* (2003).
14. Hole K. Finite element mesh generation methods: a review and classification. *Computer-Aided Design* 1988; **20**(1):27–38.
15. Thompson JF, Thomas FC, Mastin CW. Automated numerical generation of body fitted curvilinear coordinate system for field containing any number of arbitrary two-dimensional bodies. *Journal of Computational Physics* 1974; **15**:299–319.
16. Tasy TK, Hsu FS. Numerical grid generation of an irregular region. *International Journal for Numerical Methods in Engineering* 1997; **40**:343–356.
17. Greaves DM, Borthwick AGL. Hierarchical tree-based finite element mesh generation. *International Journal for Numerical Methods in Engineering* 1992; **29**(4):513–520.
18. Babuška I, Aziz AK. On the angle condition in the finite element method. *SIAM Journal on Numerical Analysis* 1976; **13**(2):214–226.
19. Borouchaki H, George PL. Aspects of 2-D Delaunay mesh generation. *International Journal for Numerical Methods in Engineering* 1997; **40**:1957–1975.
20. Field DA. Qualitative measures for initial meshes. *International Journal for Numerical Methods in Engineering* 2000; **47**:887–906.
21. Arnold DN, Mukherjee A, Pouly L. Locally adapted tetrahedral meshes using bisection. *SIAM Journal on Scientific Computing* 2000; **32**(2):431–448.
22. Cougny HL, Shephard MS. Parallel refinement and coarsening of tetrahedral meshes. *International Journal for Numerical Methods in Engineering* 1999; **46**:1101–1125.
23. Krížek M, Pradlová J. Nonobtuse tetrahedral partitions. *Numerical Methods for Partial Differential Equations* 2000; **16**(3):327–334.
24. 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.
25. Löhner R, Oñate E. An advancing front point generation technique. *Communications in Numerical Methods in Engineering* 1998; **14**:1097–1108.
26. 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.
27. Borouchaki H, Laug P, George PL. Parametric surface meshing using a combined advancing-front generalized Delaunay approach. *International Journal for Numerical Methods in Engineering* 2000; **49**:233–259.
28. Radovitzky R, Ortiz M. Tetrahedral mesh generation based on node insertion in crystal lattice arrangements and advancing-front-Delaunay triangulation. *Computer Methods in Applied Mechanics and Engineering* 2000; **187**:543–569.
29. Kikuchi N. Keynote lecture on digital-image based modelling. *Proceedings of the International Symposium on Optimization and Innovative Design*, Tokyo, 1997, July.
30. Hollister SJ, Riemer B. Digital image based finite element analysis for bone microstructure using conjugate gradient and Gaussian filter techniques in SPIE. *Mathematical Methods in Medical Imaging II* 1993; **2035**: 95–106.
31. 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.
32. Zienkiewicz OC, Zhu JZ. A simple error estimator and adaptive procedure for practical engineering analysis. *Computer Methods in Applied Mechanics and Engineering* 1992; **24**:337–357.
33. Fuenmayor FJ, Oliver JL, Rodenas JJ. Extension of the Zienkiewicz–Zhu error estimator to shape sensitivity analysis. *International Journal for Numerical Methods in Engineering* 1997; **40**:1413–1433.

34. Zienkiewicz OC, Zhu JZ. The super-convergence patch recovery and a-posteriori error estimates; Part I. The recovery technique Part II. error estimates and adaptivity. *International Journal for Numerical Methods in Engineering* 1992; **33**:(Part I) 1331–1365 (Part II) 1365–1382.
35. Boroomand B, Zienkiewicz OC. An improved REP recovery and effectivity robustness test. *International Journal for Numerical Methods in Engineering* 1997; **40**:3247–3247.
36. Adams M, Taylor RL. Parallel multigrid solvers for 3D-unstructured large deformation elasticity and plasticity finite element problems. *Finite Elements in Analysis and Design* 2000; **36**:197–214.
37. Swan CC, Kosaka I. Voigt–Reuss topology optimization for structures with non-linear material behaviours. *International Journal for Numerical Methods in Engineering* 1997; **40**:3785–3814.