A MULTISCALE FRAMEWORK FOR PREDICTING THE PERFORMANCE OF FIBER/MATRIX COMPOSITES

A Dissertation

by

MICHAEL KEITH BALLARD II

Submitted to the Office of Graduate and Professional Studies of
Texas A&M University
in partial fulfillment of the requirements for the degree of

DOCTOR OF PHILOSOPHY

Chair of Committee,	John Whitcomb
Committee Members,	James Boyd
	Darren Hartl
	Anastasia Muliana
Head of Department,	Rodney Bowersox

December 2018

Major Subject: Aerospace Engineering

ABSTRACT

To enable higher fidelity studies of laminated and 3D textile composites, a scalable finite element framework was developed for predicting the performance of fiber/matrix composites across scales. Effective design paradigms and lessons learned are presented. Using the developed framework, new insights into the behavior of laminated and woven composites were discovered.

For a [0/90]s and [±45/0/90]s laminated composite, the classical free-edge problem was revisited with the heterogeneous microstructure directly modeled, which showed that the local heterogeneity greatly affects the predicted stresses along the ply interface. Accounting for the microscale heterogeneity removed the singularity at the ply interface and dramatically reduced the predicted interlaminar stresses near a free-edge. However, the heterogeneous microstructure was also shown to induce a complex stress distribution away from the free-edge due to the interaction of fibers near the ply interface, since close fibers were shown to induce compressive stress concentrations. The fiber arrangement had a significant effect on the local stresses, with a more uniform fiber arrangement resulting in lower peak stresses. Finally, the region needed to accurately predict the microscale stresses near the ply interface was shown to be much smaller then entire ply.

For two types of orthogonally woven textile composites, nonidealized textile models were created and subjected to a variety of loads, providing insight into how load is distributed throughout the complex tow architecture and the locations of critical stresses. By comparing the stresses of a textile model with and without binders, the binders were shown

to greatly affect the distributions of stress a tensile load but not in-plane shear. Variations in the local fiber volume fraction within the tows were shown to significantly affect the magnitude of critical stress concentrations but did not change where the critical stresses occurred. Finally, accounting for plasticity in the neat matrix pocket of the textile was shown to only affect the localized region near where binders traverse the thickness of the textile.

DEDICATION

I would like to first dedicate this dissertation to my wife. She has been an incredible supporter throughout the completion of this dissertation. Her love motivated me to work hard for us, and I cannot imagine going through life without her. Second, I would like to dedicate this dissertation to my family, who have loved and supported me through my long academic career. I would not be where I am or who I am without them.

ACKNOWLEDGEMENTS

First, I thank God for the opportunity to study a very small part of His creation at such a great university.

I am also deeply thankful for the inspiration and help from others, which has made this work possible. In particular, Dr. Whitcomb has been a great adviser and friend. He was integral to this work, offering many tools developed by himself and other graduate students before me. He has guided the research topics, given me inspiration, and imparted a great deal of knowledge through my undergraduate and graduate work. This work would not have been possible without him. In addition, my committee members have taught me in and out of the classroom. Their suggestions and guidance for this dissertation has been invaluable.

I am very grateful for the funding I have received from the NDSEG fellowship and the supercomputing facility at Texas A&M. Without their support and resources, none of this work would have been possible.

CONTRIBUTORS AND FUNDING SOURCES

Contributors

This work was supported by a dissertation committee consisting of Professor John Whitcomb (advisor), Professor James Boyd, and Professor Darren of the Department of Aerospace Engineering and Professor Anastasia Muliana of the Department of Mechanical Engineering.

All work for the dissertation was completed by the student, under the advisement of Professor John Whitcomb of the Department of Aerospace Engineering.

Funding Sources

This work was supported by the Department of Defense (DoD) through the National Defense Science & Engineering Graduate Fellowship (NDSEG) Program. In addition, the authors acknowledge the Texas A&M Supercomputing Facility (http://sc.tamu.edu/) for providing computing resources used in conducting the research reported in this paper.

TABLE OF CONTENTS

Page

LIST OF FIGURES

Page

LIST OF TABLES

Page

1 INTRODUCTION

1.1 Overview of Problem

The design of high-performance, composite materials is crucial to creating transformative capabilities in the aerospace industry. A computational framework that allows the design from the microscale, where fibers and matrix are modeled discretely, to achieve performance at larger scales would expedite the discovery, design and insertion of entirely new materials. Though the computational resources available to researchers continues to increase, most computational tools used today do not leverage the large, parallel architectures of modern supercomputers, except for tools that rely on explicit methods, which introduce many challenges and uncertainties. The design of multifunctional materials in particular has the chance to transform the industry but also hinges on the ability to predict the effect of microstructural changes on the response at higher scales. Composites with microvascular networks could allow self-healing by transporting reacting materials to cracks to form a structural resin or active cooling by pumping a coolant through the material. [1] [2] [3] Microvascular networks also have the potential to create embedded, tunable antennas within a structure [4], increasing the response of shape memory polymer structures [5]. Manipulating fluids inside microvascular channels via microscale valves, mixing chambers, and other components have many potential applications, such as chemical synthesis, and has been developed into a distinct field, microfluidics. [6] Embedded sensing fibers within a composite offers the ability to sense deformation and damage in real-time.

[1] [2] Shape memory materials have the potential of constituting morphing structures without the need of discrete actuators, saving cost and weight and creating a possibility for many new technologies. [2] [3] Each of these applications have been proposed and studied to some extent by researchers, but the technologies are still far from being integrated into commercial aerospace products, aside from a few research collaborations. Additionally, these technologies often rely on more complex composite designs that lack a clear length-scale separation. One such type of complex composite is three-dimensional (3D) textiles, which offer a unique potential for multifunctionality by allowing functional constituents to be woven directly into the fabric before the matrix is added. The ability to predict the response of complex composites with high fidelity is crucial to enabling and accelerating the development of transformative materials.

For both laminated and woven composites, a deeper understanding of where stress concentrations develop and the causes behind them is needed before functional constituents can be effectively integrated. Laminated composites exhibit many different types of failure modes, ranging from intra-ply cracking to delamination. The stresses near a free-edge are known to be critical and can induce free-edge delaminations, so it is important to understand the effect of the free-edge in order to accurately predict failure of laminated composites. The effect of the free-edge has been studied by numerous researchers for many years, However, researchers have almost always modeled the fiber/matrix plies as transversely isotropic, homogeneous materials because of the complexity of creating a suitable mesh for directly modeling the heterogeneous microstructure, the computational expense of analyzing

such a model, and the relatively few computational tools capable of leveraging large, parallel computing architectures. Treating each ply as a homogeneous, transversely isotropic material neglects the local effects of stiff fibers in a relatively soft matrix. Due to the much smaller scale of an individual fiber compared to that of a ply, the perturbation in the stress field due to the actual heterogeneity will remain localized to a small region, but if the aim of an analysis is to predict when a ply crack or delamination will initiate, then local stress concentrations can be very significant.

Two-dimensional textile composites have been routinely used within the aerospace industry, but 3D textiles have received a much slower adoption due to their challenging characterization and manufacturing. Three dimensional textiles use through thickness tows that tie the layers together to offer increased toughness, out-of-plane properties, and impact resistance. Three-dimensional woven composites have been used for some applications, such as blade containment systems and body armor, and they offer the potential for multifunctional applications, such as integrating active cooling, self-healing, or a tunable antenna into the material. [1] [2] [4] [7] However, the complex geometry makes it difficult to create realistic textile models and the resulting models typically require significant computational resources. Consequently, relatively few works in the literature have focused on predicting the stresses within 3D textiles. To study where stress concentrations develop under different loads and how load is transferred within the complex tow architectures of 3D textiles, a scalable finite element framework that leverages modern parallel computing resources is required.

For both laminated and textile composites, a new depth of understanding is needed to guide the future of composite design. For laminated composites, more insight into the stresses that develop at the microscale in critical situations should be pursued in order to realize how damage initiates at the microscale and grows across scales. On the other hand, 3D textile composites have been studied far less exhaustively as laminated composites due to the complex tow architecture. Significant research is still needed among the community to understand where damage will initiate at the microscale and what modeling details are important However, conventional finite element codes that do not leverage high-performance computing (HPC) resources will be insufficient to provide the desired insight.

1.2 Current State of the Art

For many years, researchers have used classical laminate theories, finite difference models, and finite element analysis (FEA) to characterize the effect of the free-edge for various types of layups with great success. [8] [9] [10] The effects of the stacking sequence and ply thickness have been studied thoroughly. However, researchers have almost always modeled the fiber/matrix plies as transversely isotropic, homogeneous materials. Some researchers have employed multiscale strategies to capture the effects of fibers in critical regions. [11] [12] In addition, some previous studies, have modeled discrete fibers near the free-edge, but relatively few fibers (approximately 10 fibers) were considered due to the prohibitive computational challenges. [13] [14]

The analysis of 3D textiles presents many challenges. Realistic models are difficult to create due to the complex tow architecture, meshes are often large and require significant

computational resources, and stresses within the textiles are complex, making visualization and communication difficult. Consequently, there are relatively few works in the literature on the topic compared to 2D textiles.

Textile models can be divided into two categories: idealized and nonidealized. For idealized models, the tow paths and cross-sections are specified without any need for analysis. For nonidealized models, tow paths and cross-sections are specified a priori without any need for simulating a change in the tow geometry. This results in very straight tows, such as the example shown in Figure 1.1.

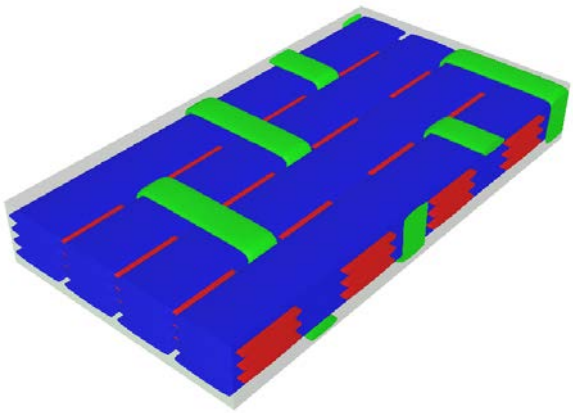

Figure 1.1. Example of an idealized textile model using TexGen

Over the last decade, researchers have developed numerous different techniques for creating 3D textile models. Among the tools for creating idealized 3D textile models, TexGen is a very popular tool created by Lin et. al. for creating idealized textile models. The tool uses spline interpolations through control points for the tow paths and assumes the tows have an idealized cross-sectional shape. [15] Lomov et. al. developed another

widely used tool called WiseTex. The tows have an idealized cross-section, but the tow paths are modified from an initial idealized path such that the deformational energy within the tows is minimized, which captures some of the realism of tow paths within 3D textiles. [16] Finally, Potter et. al. proposed a method for creating voxel meshes for idealized tow architectures. [17]

Among the methods for creating nonidealized 3D textile models, three categories of strategies exist among researchers. The most popular strategy is the digital element method. The strategy relies on having multiple digital chains per a tow, similar to fibers within a tow but with less than 100 digital chains. Using a combination of elastic and contact analyses, the processing a 3D textile preform undergoes is simulated using virtual molds and tows modeled using the digital chains. This strategy began with Wang and Sun in 2000 when they published Ref. [18]. Since then, a several researchers have developed the concept in tools for creating nonidealized 3D textile models. Two very similar tools were created: Virtual Textile Morphology Suite (VTMS) and Digital Fabric Mechanics Analyzer (DFMA). DFMA focuses on fabrics that do not have a matrix surrounding the tows, though it is more robust for creating fabric architectures. [19] [20] However, Drach et. al. developed an in-house method for removing interpenetrations between tows and creating a conforming mesh for the matrix. [21] The Air Force Research Lab (AFRL) developed VTMS for modeling composite materials. [19] [22] VTMS simulates the effects of processing on the tow geometry, creates surface representations of the tows, and utilizes a mesh independent method to predict the response of the textile model. The mesh independent method was

developed so that meshes of each constituent did not have to be compatible and could even overlap. Both VTMS and DFMA model the digital chains as rigid spheres connected by pins to elastic rod elements. The rod elements can carry load axially, while interpenetration is not allowed between rigid spheres. Aside from those, Green et. al. developed a similar digital element approach using beam elements instead of elastic rods with spheres. [23] Diverging from the digital chain methodology, Wintiba et. al. developed a method for creating nonidealized textile models with periodic geometry by modifying the cross-sections of idealized tows where interpenetrations occur. The geometry is "tensioned" to force tows to move during the analysis and create contact. [24] Finally, Stig et al. modeled the yarns as inflatable tubes. The yarns are initially shrunk and then re-inflated during a contact analysis, causing the tow cross-sections to deform. [25] It should be noted that DFMA and VTMS are the only tools in the open literature that attempt to simulate the processing a textile preform undergoes.

Using the methods described to create 3D textile models, researchers have employed finite element analysis to predict the response of 3D textiles for over a decade, beginning in the mid-1990's. [26] As researchers created more realistic models of the tow architecture, the ability to predict useful mesoscale properties from the models became a reality. In 2011, Prodromou et. al. developed a method of cells approach to predicting the mescoscale mechanical properties of 3D textiles. [27] In 2013, Hallal et. al. reviewed the available literature and showed that the numerical methods at the time provided promising tools for predicting the elastic moduli between 6% for most types of 3D textiles. [28]

More recently, researchers have started to use textile models to predict elastic properties and strengths for the textile composites. Dai and Cunningham used a relatively simple, idealized unit-cell with up to 120,000 elements. The model was used within a mosaic macro model to predict the strength of a textile panel with a hole. The simulates predicted the strength with 20% and predicted a similar progression of damage to those observed in experiments. [29] Saleh et. al. used an idealized textile model to predict the progression of damage under tensile loads along the x- and y-directions. The model was shown to match mesoscale quantities within 10% and qualitatively predicted a realistic progression of damage within each type of tow. For tension along the warps, they predicted that damage initiates in the binders where they bend begin to traverse the thickness, but critical damage quickly develops within the wefts. Within the experiments conducted, the progression of damage for an orthogonally woven textile under tension along the warps was: initial damage near binders where they begin or end to traverse through the textile thickness, matrix cracking in the tows perpendicular to the load, then finally delamination right before fiber breakage. [30] Green et. al. used a voxel continuum damage approach to predicting the damage with 3D woven composites. The results showed that waviness within the tows, even within the nominally straight warps and wefts, significantly affected the results. The models were created using their digital element method. [31]

In 2014, Drach et. al. created the most refined textile models, which remains the most refined models in the open literature until this work. The most refined mesh in the study had 3.2 million nodes. A series of meshes were created, and the elastic moduli of the textile

were predicted. It was shown that the predictions changed up to 10% when the mesh was refined from 500k nodes to over 3.2 million It was unclear if convergence for the elastic moduli was achieved with the 3.2 million node mesh, but it remains the highest standard for a high-fidelity model and highlighted the need for more refined meshes [21] In 2017, Shreyas Joglekar and Mark Pankow predicted the kinking of a orthogonally woven composite under compressive loads. An eigenvalue analysis was used to predict the buckling load for a textile model consisting of 300,000 linear tetrahedral elements. The predicted buckling load using an idealized and nonidealized textile model, created using TexGen and DFMA respectively, were compared to experimental results. A comparison of the tow architecture showed that the DFMA model had greater similarity to actual specimen, and the prediction of the buckling load was significantly more accurate when using the nonidealized model. [32] Finally, Drach et. al. predicted the progression of damage during curing within a orthogonally woven composite and compared four common failure criterion for the matrix. The study showed that a parabolic criterion was most suitable for the epoxy matrix at predicting realistic damage within the matrix. DFMA was used to create realistic textile models, and an in-house MATLAB code was used to remove interpenetrations between tows before a mesh for the matrix was created. The results were compared to an actual IM7/RTM6 and showed good agreement when an appropriate failure criterion was used. [33]

A very recent review of 3D textile by Saleh and Soutis showed that the community has focused on through thickness properties of 3D textiles but the in-plane properties remain

not well understood. The review considered three different categories of weaves: orthogonal, layer to layer, and angle interlock. Most the literature addressing in-plane properties focused on the tensile properties, leaving a gap of knowledge for textiles under in-plane shear loads. In general, it was shown that the orthogonal and angle interlock textiles exhibited improved delamination resistance compared to layer to layer textiles. [34] Also very recently, Huang reviewed the current literature for the prediction of mechanical properties of 3D textiles. The review highlighted that there is a lack of literature addressing textiles with z-pinning, such as orthogonally woven composites, under in-plane shear loads. [35]

From the mid-1990's the community has progressed to be able to predict mesoscale properties, such as elastic moduli and tensile strengths, within 10%. However, a study of 3D textiles under in-plane loads and a detailed investigation of the stress distributions with the 3D textiles is lacking in the body of literature. No published work has rigorously investigated the locations of stress concentrations and how load is distributed within the textile composite, but these insights will be necessary to effectively integrate multifunctional constituents and design better performing 3D textile composites.

1.3 Strategy

This work relies on the development of a finite element framework that leverages HPC resources to investigate fiber/matrix composites with greater detail than the current state of the art. A suite of preprocessing tools that leverage a combination of novel algorithms

and third-party tools was developed to create complex microscale models with random fiber arrangements, models of laminated composites with the heterogeneous microstructure directly modeled within the plies, and nonidealized 3D textile models. Since the meshes for these models can easily have more than 30 million nodes, a scalable finite element code was developed to predict the response of linear and nonlinear materials. The analysis of large models creates the challenge of storing large amounts of data from the distributed analyses. Consequently, a new file format was designed to store the data needed to describe the mesh, local coordinate systems, and a variety of field data types. Finally, a collection of Python scripts and a custom plugin for ParaView, allowing ParaView to efficiently read the newly developed file format in parallel, were developed to enable the visualization of the complex finite element data.

Leveraging this new finite element framework, both laminated and woven composites will be investigated with greater detail than past studies. The heterogeneous fiber/matrix microstructure within a laminated composite is directly modeled for two types of laminated composite layups: a crossply [0/90]s layup and quasi-isotropic [±45/0/90]s layup.. Simulating the laminate under uniaxial tension, the effect of discrete fibers on the stresses along the ply interface close to the free-edge is explored. The model with the heterogeneous microstructure directly modeled is also compared to the classical case where each ply is treated as a homogeneous material to highlight the differences that emerge when the heterogeneous microstructure is considered.

The finite element framework is used to explore the behavior of orthogonally woven composites under different in-plane loads using nonidealized textile models. Two types textiles are considered, a thin 1x1 orthogonal weave, which is a relatively simple 3D textile, and a thicker 2x2 orthogonal weave, which is more realistic to the 3D textiles used in industry. The aim is to provide insight into how load is transferred throughout the complex tow architecture, locations of critical stress, the role of the binders, and the effect of several modeling parameters.

2 ANALYSIS FRAMEWORK

2.1 Overview

Predicting the behavior of fiber/matrix composites presents difficult challenges at several lengths scales. At the microscale, fibers and matrix must be accounted for discretely, but since the microstructure is random, a valid model can become large enough to challenge a typical workstation. However, to better understand the distribution of possible responses within a composite, an ensemble of fiber/matrix models should be considered, which increases the computational demands of a microscale study significantly. At the mesoscale, laminated composites have historically been treated as homogeneous materials, which works well if the purpose of a simulation is to investigate the mesoscale response. However, if localized phenomena, such as damage, or small region on the scale of a fiber diameter, such as free-edge stresses, are of interest, then the microscale must be considered. When the laminated composite is modeled as a homogeneous material, a typical workstation might suffice for an analysis, but when the individual fibers are discretely modeled within a laminate composite, the simulation requires the use of high-performance computing. Finally, three-dimensional textile composites exhibit a complex geometry. The textile models require large meshes and produce large amounts of output for post-processing. All of these challenges create the need for a robust, scalable analysis framework. An effective analysis framework should leverage computational resources ranging from a single workstation to a large supercomputer, provide a toolset for efficiently creating models of fiber/matrix composites at the micro- and mesoscales, automate many tasks via scripts to

allow parametric studies or Monte-Carlo simulations, and include a toolset for visualizing the data resulting from large analyses.

This chapter proposes an analysis framework designed for the analysis of fiber/matrix composites, delineates the underlying algorithms, and highlights any third-party tools that are particularly useful. The proposed framework provides the tools necessary to study laminated and textile composites at the micro- and mesoscales. Unlike much of the work in the literature, most of the tools were developed in-house as part of this doctoral work. The decision to develop such a broad set of tools was based on the limitations of commercial software. The complexities of three-dimensional textiles and scale of computational resources required for many of the analyses required the development of a FEA framework capable of leveraging HPC systems. Figure 2.1 provides an overview of the analysis framework, indicating components that are third-party libraries and components developed to support this work. In the architecture diagram, a component with any part on top of a component in next lower level depends on the component from the lower level. For example, Beta2 depends on the BetaMesh, Sparse Math, and Dense Math libraries. Some of these components will be referenced throughout this chapter, but implementation details is outside the scope of this manuscript. This chapter discussing the analysis framework is divided into two sections. The first section discusses the preprocessing framework, consisting of new algorithms and third-party tools, used the microscale fiber/matrix models, laminate composite models with the heterogeneous microstructure directly modeled, and the 3D textile models. The next chapter provides a discussion of the development of the

scalable finite element framework used in this work, focusing on highlighting the challenges of distributed computing, lessons learned throughout the development, and effective design paradigms.

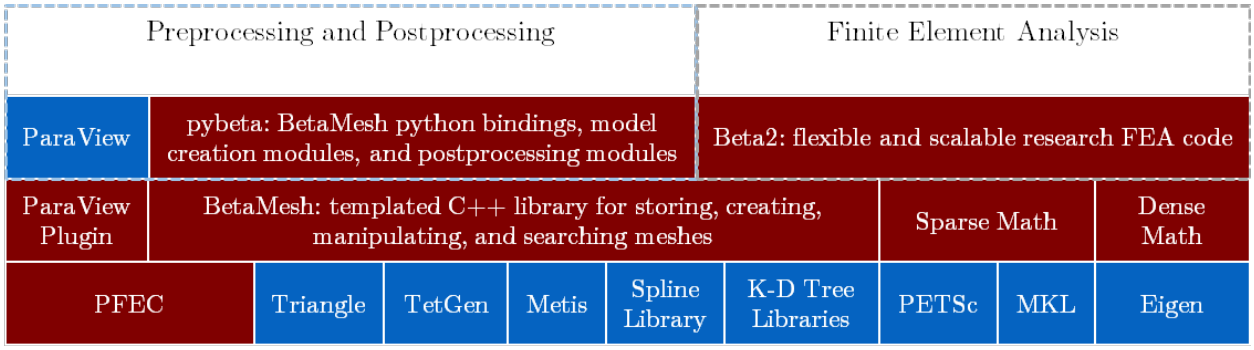

Note: MPI and OpenMP parallelization libraries are used across many components

Figure 2.1. Overview of analysis framework architecture (blue components indicate third-party libraries and maroon components indicate the ones developed)

2.2 Preprocessing

This section describes the algorithms and tools used to create the finite element models of unidirectional, laminated, and 3D woven textile composites. The first subsection describes the algorithms for creating and meshing random fiber arrangements needed for microscale analyses. Though these types of models have been created by other researchers for years, this work proposes a more general model generation technique, capable of creating complex RVE shapes. The second section presents the algorithms for creating the multiscale laminated composite models with discretely modeled fibers in the 0^0 and 90^0 plies. The third section outlines the suite of tools and algorithms used to create 3D textile models with

non-idealized geometry. Finally, the last section will share lessons learned, including proposed improvements of the preprocessing tools and strategies found to be particularly effective.

2.2.1 Fiber/Matrix Models

For years, many researchers have successfully created fiber/matrix models with random fiber arrangements. Some focus on statistically matching the arrangement observed in actual specimen, while other focus on efficiently creating random arrangements. However, all the existing methods have focused on square or rectangular RVE's. There are several reasons one might want to generalize the algorithms for a non-rectangular shape. For example, in Ref. [36], it was shown that square RVE's of random fiber/matrix composites tend towards transverse isotropy as the number of fibers in the periodic RVE increases, but as expected, in the limit of RVE sizes, the effective properties derived from the square RVEs were unable to result in transversely isotropic properties. Other RVE shapes will result in effective properties that approach exact transversely isotropy as the number of fibers increases, such as hexagonally shaped RVEs. Another possible reason, beyond the scope of the current work, is to allow the simulation of actual tow cross-sections at the microscale within a multiscale analysis, capturing the finite boundaries of the tow cross-sectional shape.

This work proposes a more general set of algorithms for generating 2D random fiber/matrix models for any closed shape. The discussion will be limited to 2D models, though it could be extended to 3D models in a straightforward manner. The strategy can be divided into the two algorithms: one for generating the fiber arrangements and one for

meshing the fiber arrangements. Each of the algorithms will be described in the following two sections respectively.

Algorithm for Creating Fiber Arrangements

The algorithm for creating random fiber arrangements can be divided into three steps:

1. Specification of RVE boundaries
2. Initial placement of fibers
3. Iterative removal of overlap between fibers

The boundaries of the RVE are defined by a list of points that are connected by straight boundary segments going counter-clockwise around the closed RVE shape. The direction is important because the outward and inward normal vectors for each boundary segment will be inferred from the connectivity. This work will assume faceted geometry because it simplifies the triangulation of the closed area in step 2 and the collision detection during step 3. Given this assumption, exactly two points will define each boundary segment.

Each boundary segment can be designated as one of three types: finite, symmetric, and periodic. A finite boundary does not allow fibers to cross during iterations. A symmetric boundary forces a fiber to either directly lie on the boundary, which cuts the fiber in half, or not cross the boundary at all, so symmetry can be preserved. A periodic boundary requires a pair of boundary segments and allows a fiber to cross one boundary segment with the remaining part of the fiber appearing at the other boundary segment in the pair. Figure 2.2 shows an example of fibers crossing periodic boundaries within a periodic square RVE. To ensure that periodicity is possible for a pair of boundary segments, it is useful to identify

matching boundary segments. Boundary segments m and n are a matching pair only if the lengths of the two segments are equal within numerical tolerance, which is covered by the next condition but provides an efficient initial check, and if $|(\boldsymbol{p}_1^m - \boldsymbol{p}_0^n) - (\boldsymbol{p}_0^m - \boldsymbol{p}_1^n)| < \varepsilon$, where $\boldsymbol{p}_i^j$ is the coordinate of point i of segment j and ε is the numerical tolerance used to test if two vectors are equivalent. Refer to Figure 2.3 for an illustration of identifying matching pairs. It should be noted that all boundary segments having a matching pair is not sufficient for periodicity of the RVE, consider an octagonal RVE. However, some models may only have periodicity in one direction, such a small periodic slice of a unidirectional lamina with the entire thickness modeled. Once the boundary type for each segment has been specified, step 2 begins.

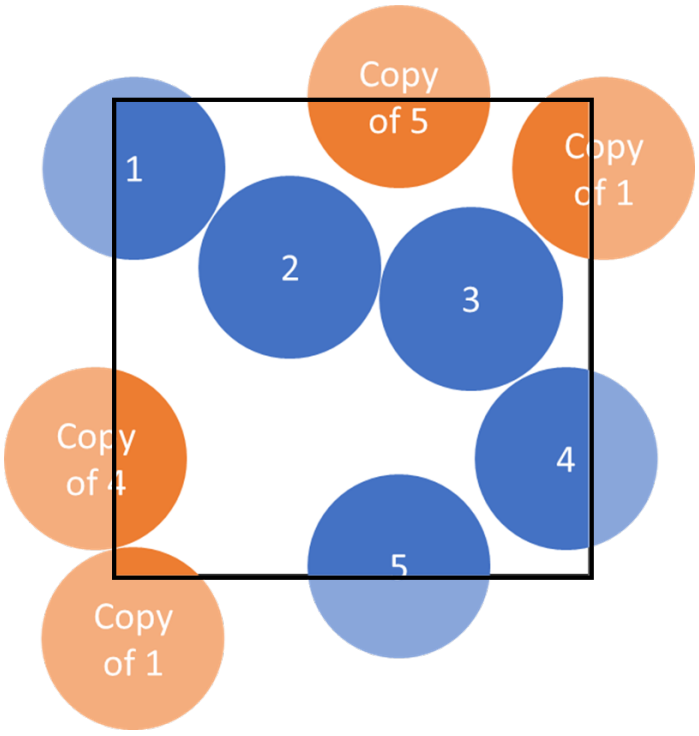

Figure 2.2. Example of 5 fibers within a periodic square RVE, showing copies of fibers that cross periodic boundaries

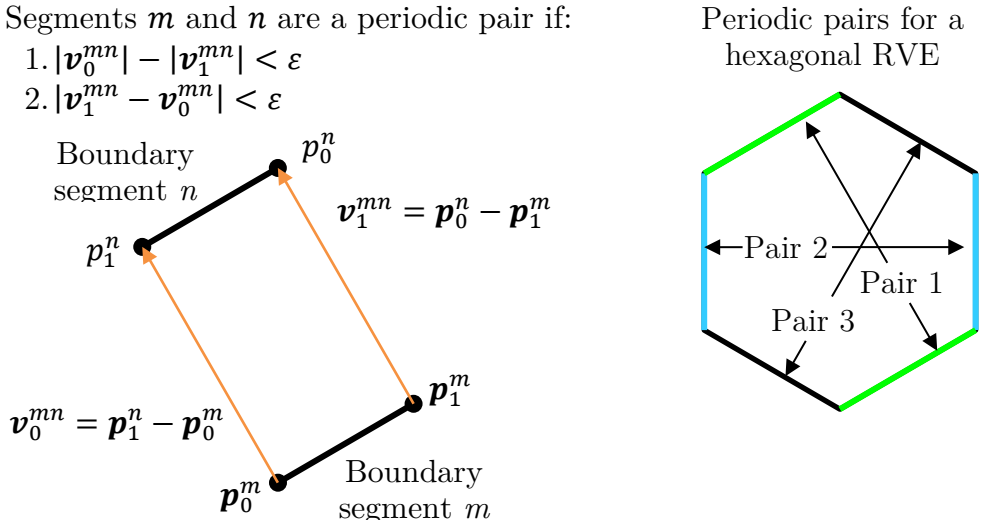

Figure 2.3. Identification of periodic pairs of boundary segments

Creating the initial placement of fibers, step 2, generally results in overlapping of fibers. Though the fibers can overlap at this point, the initial fiber arrangement should satisfy the boundaries of the RVE, which will depend on the boundary type of each segment. Figure 2.4 shows the flow chart of step 2, including illustrations of some of the substeps for a hexagonal RVE with 8 fibers.

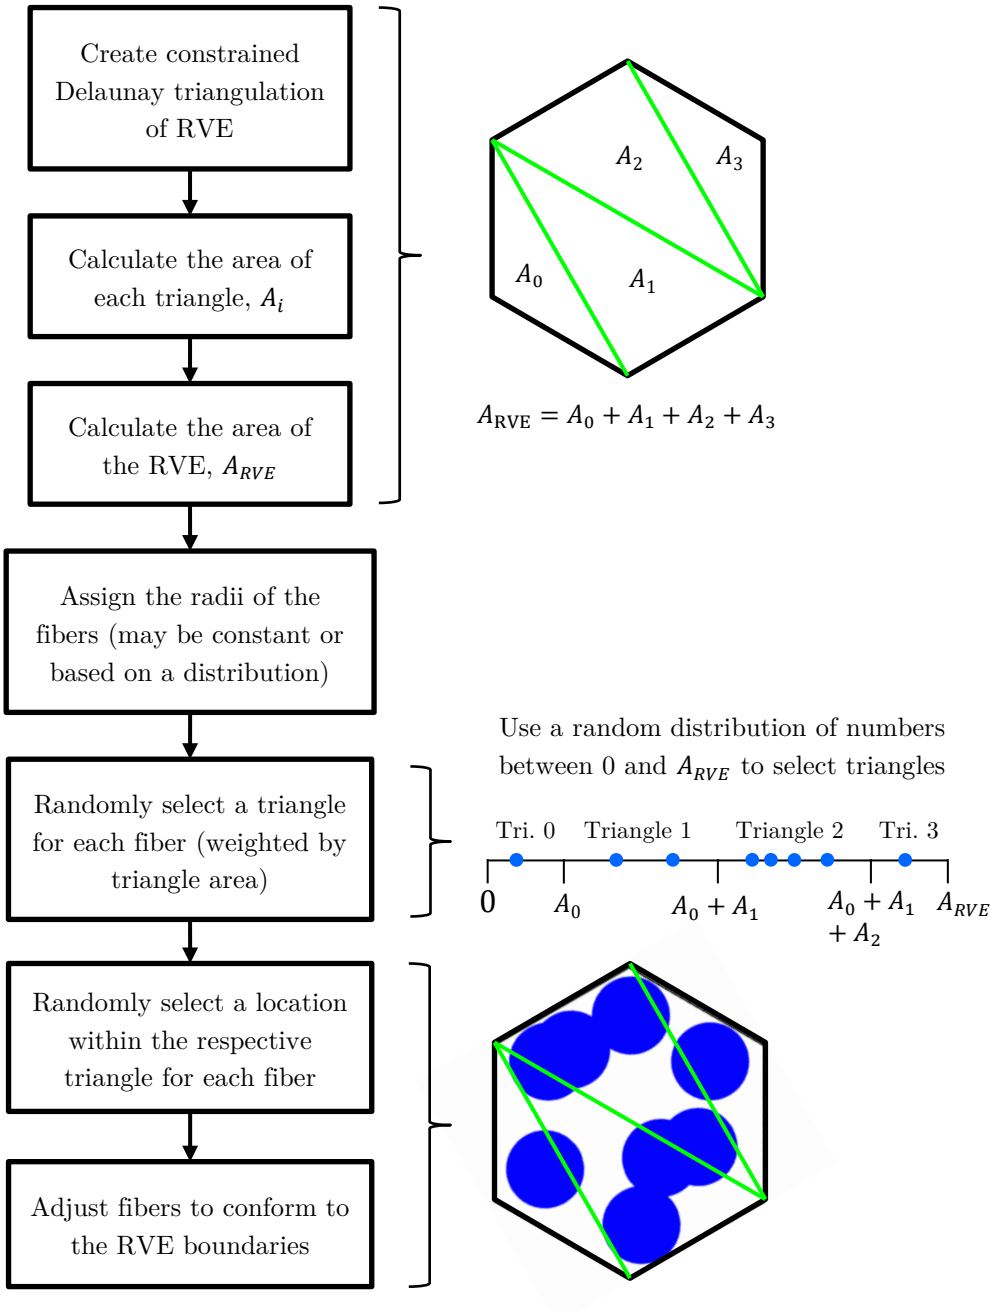

Figure 2.4. Flow chart of procedure for initially placing fibers (step 2) with illustrations for a hexagonal RVE with 8 fibers (shown on the right)

First, the boundary segments are used to create a constrained Delaunay triangulation of the RVE domain. While Delaunay triangulation is a method for creating triangles meeting the Delaunay condition given a set of points, constrained Delaunay triangulation is an extension of Delaunay triangulation that ensures the triangulation includes specified edges, which in this case is the RVE boundary segments. The triangulation of the RVE domain is then used when placing the fibers within the RVE. For each fiber, a random triangle is chosen, weighted by the area triangles so that large triangles have a higher likelihood of being chosen. This is done by creating a random value between 0 and the area of the RVE, A_{RVE}, for each fiber. To determine which triangle each fiber should be placed within, the range of values between 0 and A_{RVE} are discretized based on the area of each triangle, see Figure 2.4 for an example. With a triangle randomly chosen for each fiber, each fiber is randomly placed within the respective triangle. For fiber i in triangle j, the coordinate of the center of the fiber, $\boldsymbol{c}_i$, is given by Eqn. (2.1), where r and s are two random numbers between 0 and 1 and $\boldsymbol{q}_0^j$, $\boldsymbol{q}_1^j$, and $\boldsymbol{q}_2^j$ are the three vertex coordinates of triangle j.

$$\boldsymbol{c}_i = \left(1 - \sqrt{r}\right)\boldsymbol{q}_0^j + \sqrt{r}(1 - s)\boldsymbol{q}_1^j + \sqrt{r}s\boldsymbol{q}_2^j \tag{2.1}$$

With each fiber randomly placed within the RVE, fibers near the boundary of the RVE might not satisfy the requirement of a nearby boundary segment. So, each fiber is adjusted to conform to the RVE boundaries, which will be one of three cases:

1. If the fiber crosses a finite boundary, the fiber is moved along the inward normal of the boundary segment so that the fiber no longer crosses the boundary.
2. If the fiber crosses a symmetric boundary, the fiber is moved to lie perfectly on the boundary if it is closer than $\frac{r_i}{2}$; otherwise, the fiber is moved away from the boundary segment along the inward normal of the boundary segment such that any overlap is removed.
3. If the fiber is near a periodic boundary, a periodic copy is created on the other side of the periodic pair of boundary segments. If a fiber is near or crosses two periodic boundary segments, an extra periodic copy of the fiber should be created along the sum of the two periodicity vectors, refer to Figure 2.2 and Figure 2.5.

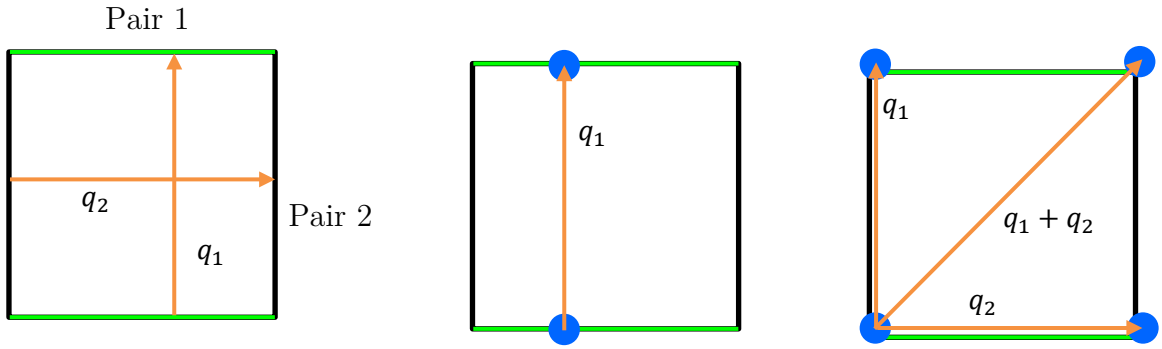

a) Periodic pairs of boundary segments and their periodicity vectors, q_i

b) Creating a periodic copy of a fiber crossing one periodic boundary

c) Creating periodic copies of a fiber crossing two periodic boundaries

Figure 2.5. Creating periodic copies of fibers near or crossing a periodic boundary segment

Once the fibers have been randomly placed within the RVE, any overlap between fibers must be removed. Step 3 involves an iterative procedure proposed by Ross McLendon in Ref. [37]. The iterative procedure used in this work is largely unchanged from the method proposed in Ref. [37]. Figure 2.6 shows the flow chart for the procedure used to remove the overlap between fibers.

To begin, given a cutoff distance for detecting neighboring fibers, d_n, the neighbors closer than d_n of each fiber are identified using a k-d tree search. This avoids having to check if each fiber overlaps every other fiber in the RVE during the iterations, which would be an n^2 problem.

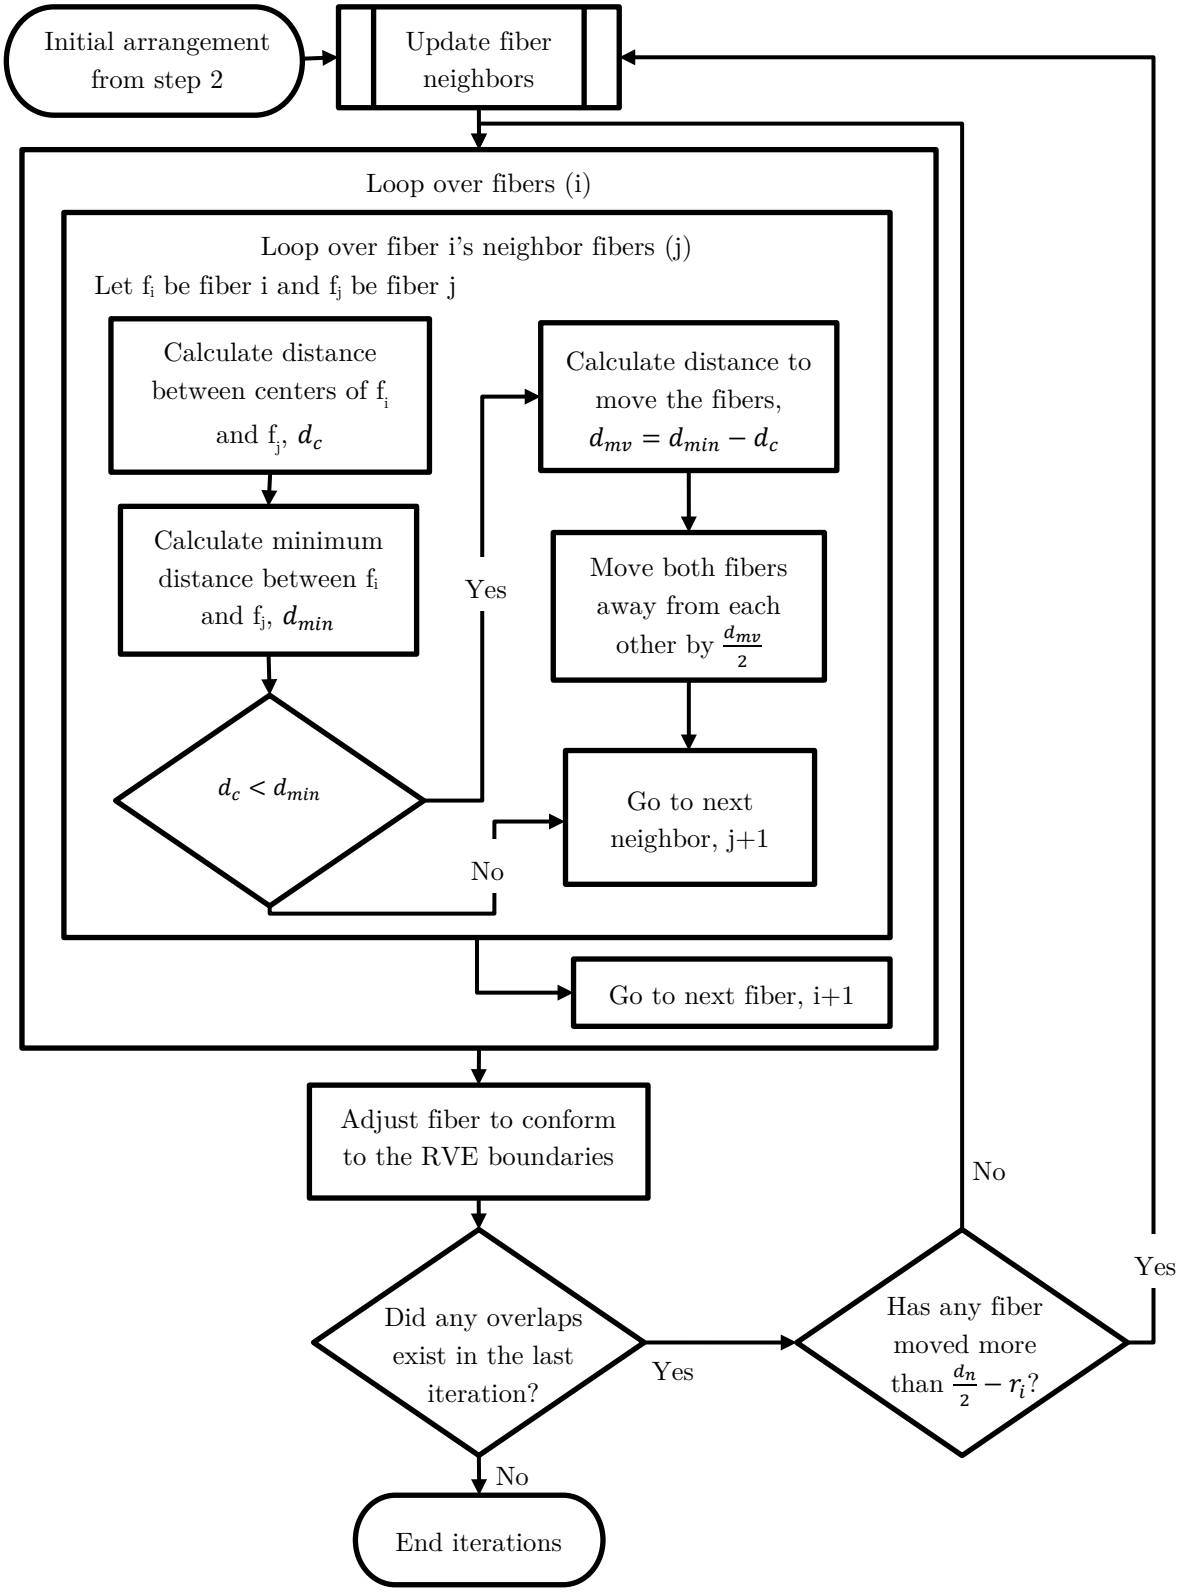

Figure 2.6. Flow chart of procedure for removing fiber overlap (step 3)

Next, for each fiber, the distance between the fiber and its neighbors is checked to ensure that the fibers are far enough apart. The minimum distance between two fibers, i and j, is given by $d_{min} = r_i + r_j + d_s$, where r_i is the radius of fiber i, r_j is the radius of fiber j, and d_s is the minimum space required between fibers. Specifying a minimum space that must exist between fibers effectively controls how random the realization will become. As the minimum space between fibers increases, the fiber arrangement tends toward hexagonal packing, though it is prohibitive to precisely create hexagonal packing from an initial random placement of fibers. It is convenient to propose a minimum spacing fraction, ξ, that is the distance required between fibers divided by the average of the radii of the two fibers. One value of ξ will be specified for an entire fiber arrangement, while the minimum distance between two fibers is given by $d_s = \frac{(r_i + r_j)}{2}\xi$. If two fibers are too close to each other, both fibers are moved apart from each along a line through the fiber centers.

Once each fiber is far enough away from its neighbors, the fibers are adjusted to conform to the RVE boundaries, using the same method in step 2. Fibers crossing finite or symmetric boundary segments are adjusted to lie within the RVE, and where needed, periodic copies of fibers that cross periodic boundary segments are created. If an iteration occurs were no overlap between fibers exists, then the final fiber arrangement is determined.

Algorithm for Meshing Fiber Arrangements

The previous section discussed the method for creating the arrangement of fibers, which includes the RVE boundaries, fiber radii, and fiber centers. To create a valid mesh for FEA, a mesh for the RVE is needed. There are many third-party meshing libraries that do

a great job at creating quality meshes given a boundary mesh, but a boundary mesh must be created for the basic geometry first. The algorithm for creating the boundary and area meshes can be divided into four steps:

1. Creation of boundary mesh for each fiber
2. Creation of boundary mesh for the RVE boundaries
3. Constrained Delaunay triangulation of fibers and matrix
4. Adjustment of mesh to conform to each fiber's circular boundary

For the first step of the algorithm, creating a boundary mesh for fibers that do not cross a boundary of the RVE is straightforward. Points placed around the circumference of the fiber and line elements connecting the points create a faceted closed circle. However, for a fiber crossing a periodic or symmetric boundary, the step becomes more complex. Recall that during initial placement and the iterative removal of overlap between the fibers, fiber can cross periodic boundaries, but periodic copies of the fibers are made. It is common for either the original fiber or the periodic copy of the fiber to get pushed outside the RVE during iterations. Boundary meshes must still be created for fibers that barely cross a periodic boundary segment.

Figure 2.7 provides illustrations of the procedure for creating a boundary mesh for fibers that cross boundary segments. First, the intersections of the fiber with any boundary segments are determined, requiring a line segment-circle intersection calculation, see Figure 2.7a. Next, the triangulation of the RVE domain from the previous algorithm is used to test if the fiber center lies within the RVE. If the fiber center does not lie within any

triangles and there are no intersections between the fiber and any boundaries, the fiber is thrown out. If the fiber lies within the RVE domain or intersects a boundary, then a surface mesh is needed for at least part of the fiber. Like before, points are initially evenly placed around the circumference of the fiber, refer to Figure 2.7b. This work used 20 points to start with. Points that are very close to the intersection between a fiber and RVE boundary are removed, refer to Figure 2.7c. At this point, a list of points around the fiber, including any intersections with RVE boundaries has been determined, which are then sorted going counter-clockwise around the fiber, refer to Figure 2.7d. If two intersection points lie next to each other without another point between them, one is added. An example of this is shown in Figure 2.7e with the addition of point 1 for the lower fiber. As illustrated in Figure 2.7g, any points lying outside the RVE domain are removed, see Figure 2.7f. If more than one point remains, line segments are created between each point, except when two intersection points are next to each other in the list of points, since this part of the boundary mesh will be taken care of in step 2 when the RVE boundary is meshed.

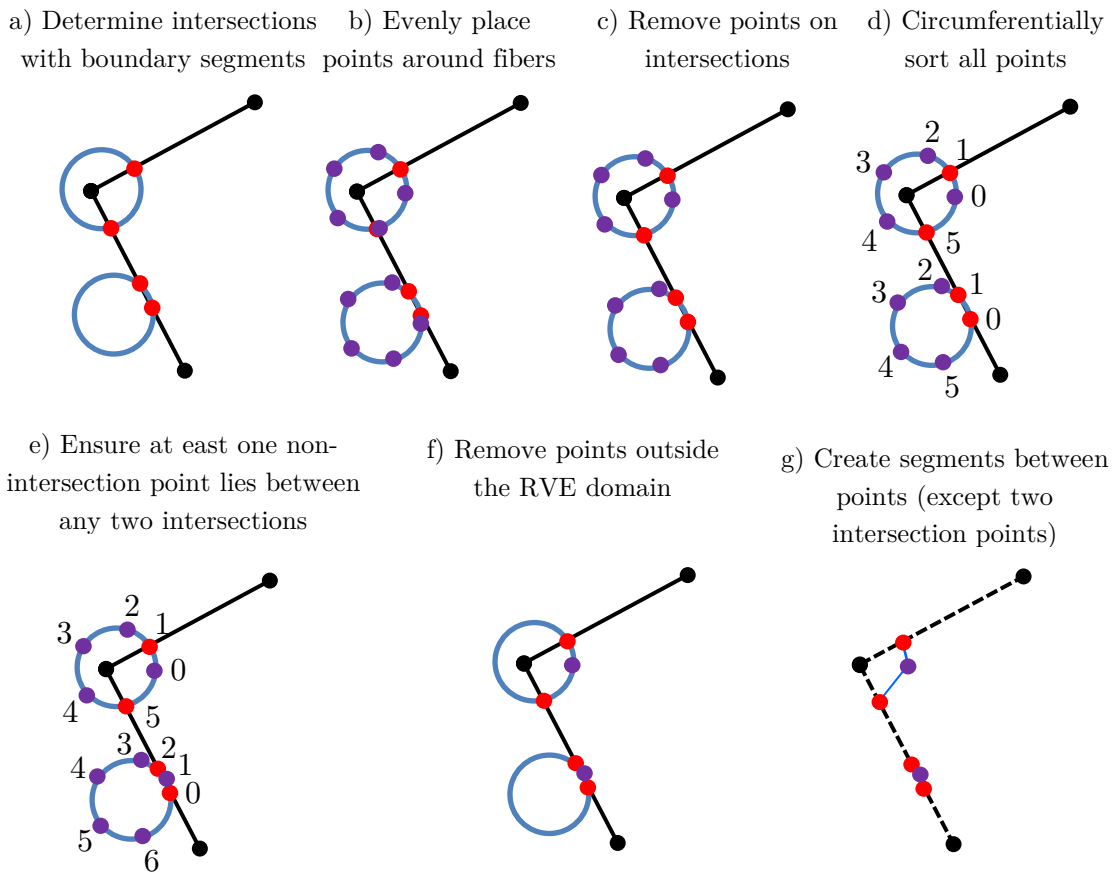

Figure 2.7. Illustration of procedure for creating boundary mesh for fibers that cross periodic boundary segments

Next is step 2, which is creating the boundary mesh for the RVE boundaries. Line elements are created through each point on the RVE boundary, including any intersection points created during step 1. To make this efficient, each intersection between a fiber and boundary segment in step 1 should be tracked. Once the boundary mesh of the fibers and RVE boundary have been created, a third-party meshing library is used to create the mesh of the domain. For this work, a library called Triangle was used. [38] Finally, any new

nodes in the resulting mesh that lie on line elements that defined the boundary of a fiber should be moved to lie on the fiber circumference. If quadratic triangles are used for the FEA analysis, then the mesh closely captures the circular shape of the fibers. Figure 2.8 shows a typical mesh of a periodic hexagonal RVE with 25 fibers.

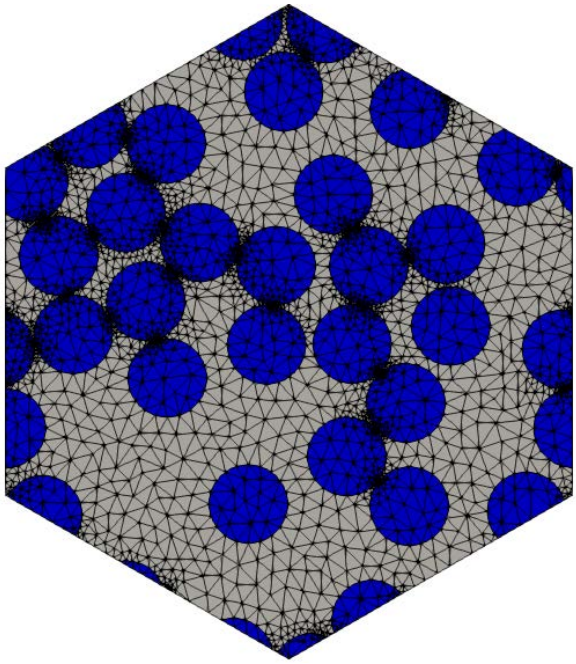

Figure 2.8. Typical mesh of a periodic hexagonal RVE with 25 fibers and a fiber volume fraction of 40%

2.2.2 Multiscale Laminated Composite Models

Researchers have used various approaches to predict the behavior of materials across scales. Some used a surrogate microscale model at each quadrature point, referred to as the FE^2 method, while others use complex shape functions that can better capture the response

of the underlying microstructure. As computational resources have increased, it is possible to consider much larger models, provided that this analysis framework can exploit HPC resources. Instead of modeling a lamina as a homogeneous material or employing reduced order models to capture the response of the microstructure, this work will revisit the free-edge problem using a multiscale model that directly models the random fiber/matrix microstructure of a laminated composite.

In this work, the 0-90 ply interface is the focus of the free-edge analyses. To study the effect of the heterogeneous microstructure on the stresses that develop along the ply interface near the free-edge, the fibers and matrix of the 0^0 and 90^0 plies will be modeled discretely. The method proposed herein is limited to discretely modeling the 0^0 and 90^0 plies. Studies requiring discretely modeled plies with different orientations will require a more complex meshing strategy than used herein.

Even with increasing HPC resources, the entire domain of a typical composite specimen is prohibitively large for directly modeling the fibers and matrix, but a much smaller unit cell can be considered that is representative of the entire specimen. For a free-edge analysis, it is important that the width of the model is large enough so that the free-edge effect on one side of the specimen does not affect the stress state on other side. For this work, models will be wider than 10 ply thicknesses. However, in the direction of the uniaxial load, the model can be much smaller.

To begin, fiber/matrix cross-sections are created for the 0^0 and 90^0 plies. The method for creating fiber/matrix models discussed in the last section is used to make the cross-

sections. The cross-sections are translated and rotated to be where the respective ply will exist, and then the cross-sections are extruded in such a way to keep the nodes along the ply interface compatible, refer to Figure 2.9.

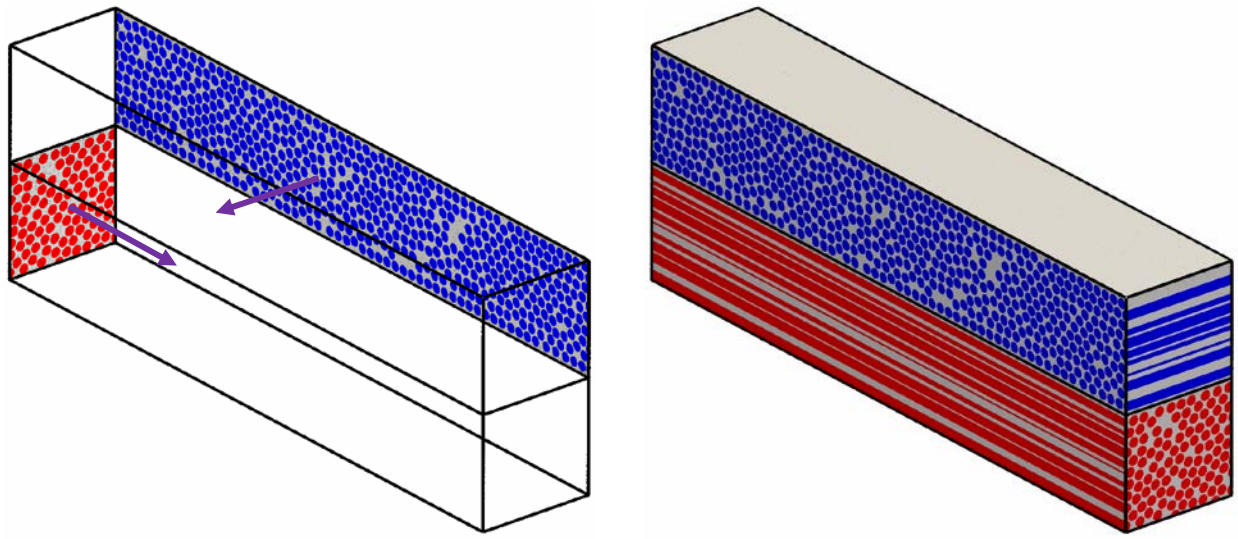

Figure 2.9. Extrusion of ply cross-sections to create volume mesh of 0^0 and 90^0 plies in a crossply laminate

Plies away from the interface of interest can be modeled as a homogeneous material, which does not require as refined of mesh as those used for the discretely modeled plies. However, a ply modeled as a homogeneous material that is connected to a discretely modeled ply requires a transition mesh, transitioning from the small element size of the discretely modeled ply to a larger element size. For the other homogeneous plies, a structured mesh can be used.

For discretely modeled plies, quadratic wedges elements are used, since the triangle fiber/matrix meshes are extruded to create the volume meshes for the laminae. In the structured meshes, quadratic hexahedral elements are used. Consequently, the transition mesh must connect to quadratic quadrilateral faces. Within the transition mesh, quadratic pyramid elements are used to change from quadrilateral faces to triangle faces. Then, tetrahedral elements are created, using a general tetrahedral mesh library, TetGen, between the pyramid elements to form the rest of the transition mesh. [39] Figure 2.10 shows an illustrative example of a transition mesh between two hexahedral meshes with different refinements..

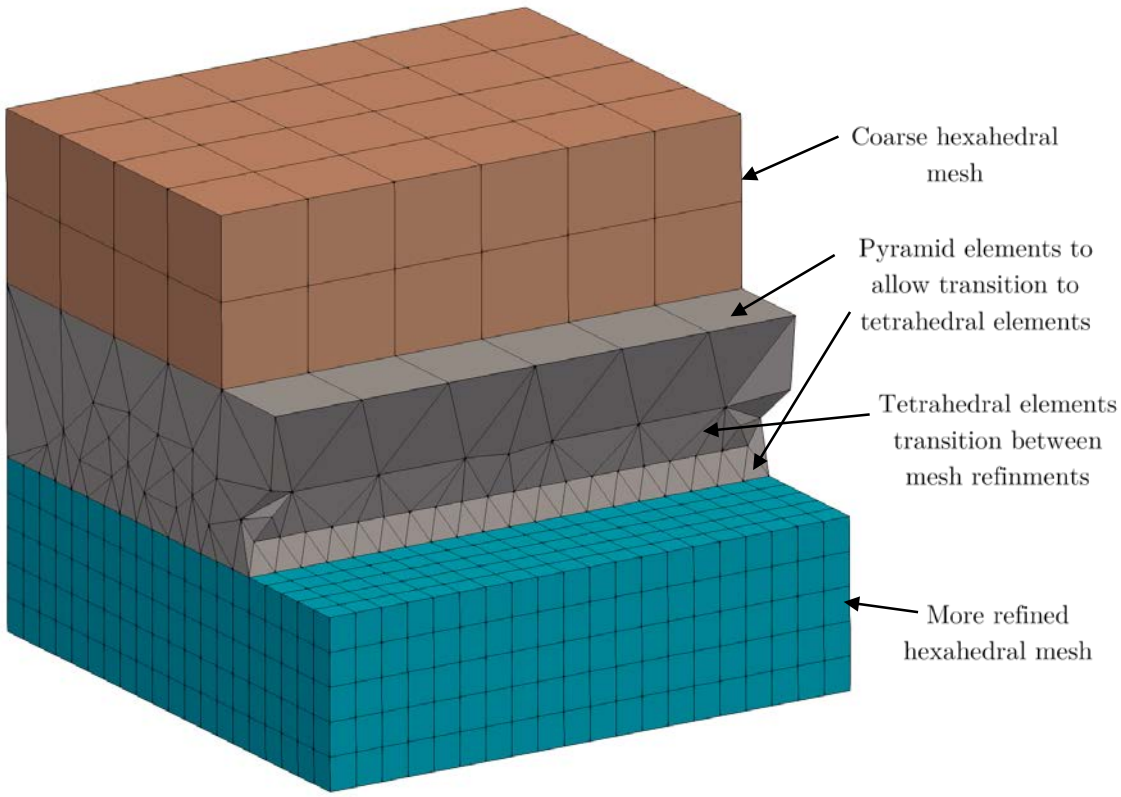

Figure 2.10. Illustrative model with refined mesh layer, coarser mesh layer, and a transition mesh layer, which uses pyramid and tetrahedral elements

For more realistic example, Figure 2.11 shows a model for a [±45/0/90]s laminated composite with both heterogeneous and homogeneous plies within the same multiscale model. The $+45^0$ ply, the top ply, is modeled using a structured mesh of hexahedral elements. The -45^0 ply, the second ply from the top, is modeled using a combination of pyramid and tetrahedral elements. Finally, the 0^0 and 90^0 plies are modeled as discrete fibers and matrix using wedge elements.

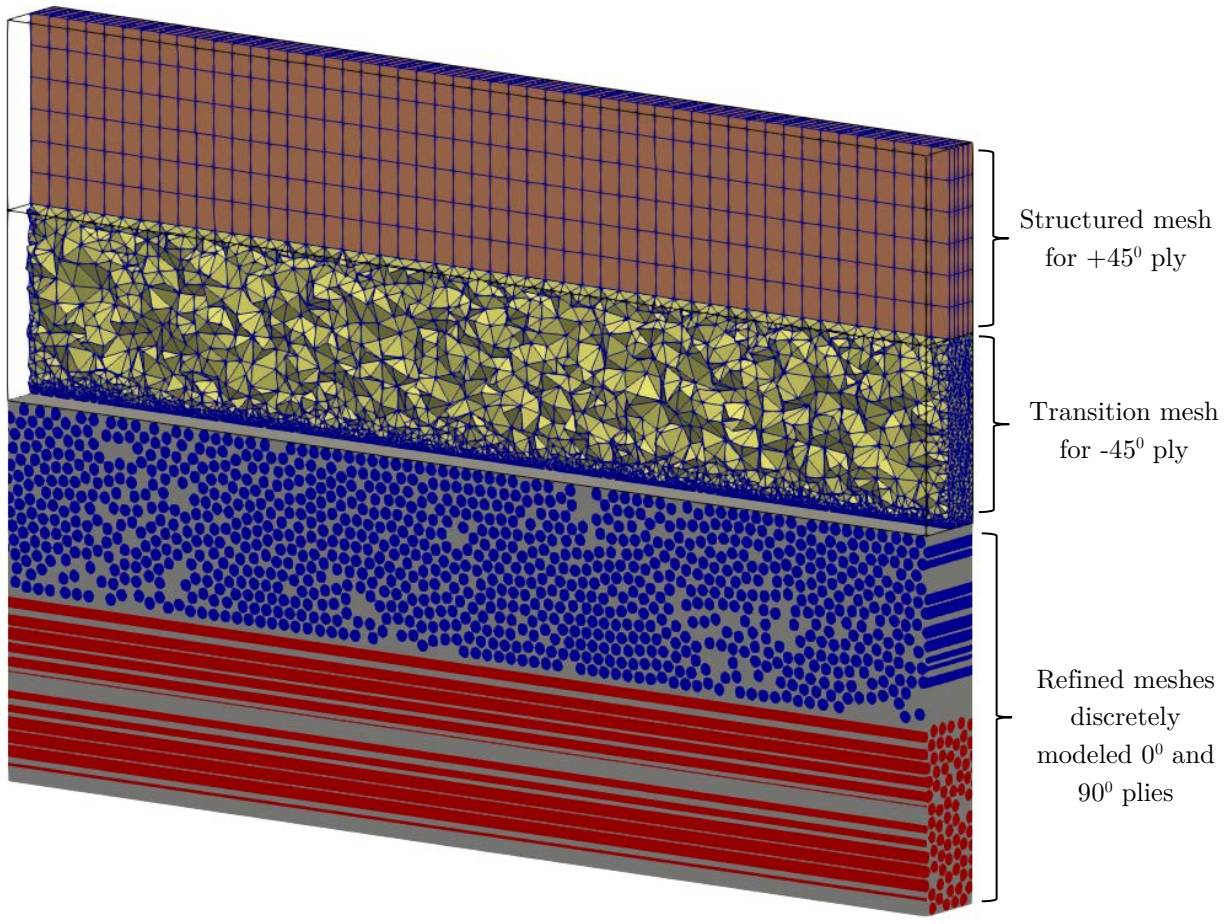

Figure 2.11. A multiscale model for a [±45/0/90]s laminated composite with mixed heterogeneous and homogeneous plies (top two plies are clipped to show interior mesh refinement)

2.2.3 Three-Dimensional Textile Models

Though researchers have extensively studied two-dimensional textiles, simulating the behavior of three-dimensional textiles has received less attention due to the complex geometry and need for large meshes. Several tools exist that can create 3D textile models, which can be separated into two categories: idealized and non-idealized. Two tools specialize in creating non-idealized textile geometries: Virtual Textile Morphology Suite (VTMS) and Digital Fabric Mechanics Analyzer (DFMA). DFMA focuses on fabrics that do not have a matrix surrounding the tows, though it is more robust for creating fabric architectures. [19] [20] The Air Force Research Lab (AFRL) developed VTMS for modeling composite materials. [19] [22] VTMS simulates the effects of processing on the tow geometry, creates surface representations of the tows, and utilizes a mesh independent method to predict the response of the textile model. The mesh independent method was developed so that meshes of each constituent did not have to be compatible and could even overlap. This introduces error near the interfaces tow-matrix interfaces, which are regions of great interest when predicting damage. Since this work utilizes a standard FEA formulation, conforming meshes of the matrix and tows without interpenetrations is necessary, but VTMS is not capable of creating the compatible meshes. Consequently, this work will use VTMS to create non-idealized geometry for orthogonally woven textiles, but an in-house tool will use the geometry to create a finite element mesh of the textile composite. Figure 2.12 shows an overview of the process used to create 3D textile models.

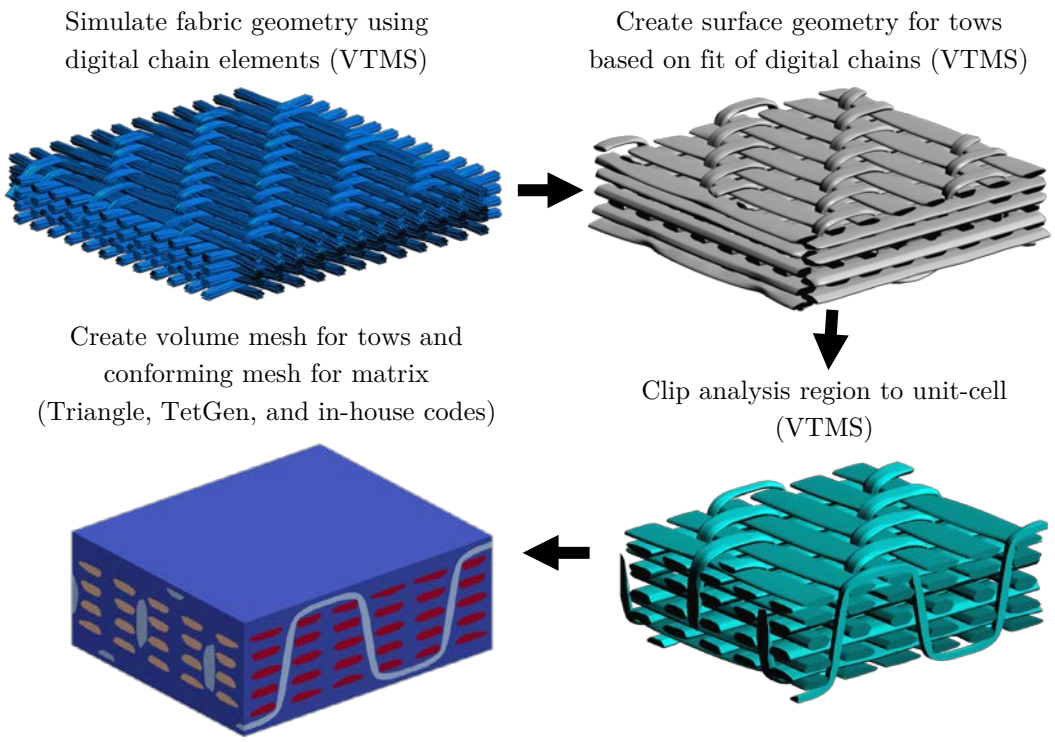

Figure 2.12. Overview of process used to create 3D textile models

Using VTMS to Create Non-Idealized Geometries for 3D Textiles

VTMS models the effects of processing through the use of "digital chains", which consist of rigid spheres connected by pins to elastic rod elements. The rod elements can carry load axially, while interpenetration is not allowed between rigid spheres. When any two rigid spheres very come close to each other, a very stiff contact element is created between them to prevent the interpenetration. If enough rigid spheres are used along the digital chain, a single digital chain resembles the behavior of a solid fiber, except that the digital chain does not have a bending stiffness like an actual fiber. Ideally, a digital chain should be used for

each fiber within the tow, but this is currently computationally prohibitive. Although VTMS can use multiple digital chains to represent each tow, it is more computationally efficient to start with a single digital chain for each tow. The process to create the textile model involves several steps, and each digital chain representing a two will be replaced with multiple digital chains at a later step.

To begin, a weave pattern is specified in VTMS, and the initial tow geometry results in a lot of space between digital chains, as shown in Figure 2.13. The space between digital chains is removed during the relaxation steps. Since contact is only enforced for the rigid spheres, it is important to ensure that the rod elements are short enough to prevent digital chains from passing through each other during relaxation steps. However, it is also important to make the elastic rod elements long enough to prevent the rigid spheres within a single digital chain from touching each other, which would result in numerical instability. For this work, the length of the rod elements divided by the diameter of the rigid spheres were specified to be 0.6. An example of a digital chain that represents one of the binders is shown in Figure 2.14. Unfortunately, the spheres are not plotted to scale within VTMS, so the spheres in Figure 2.14 are not shown to scale.

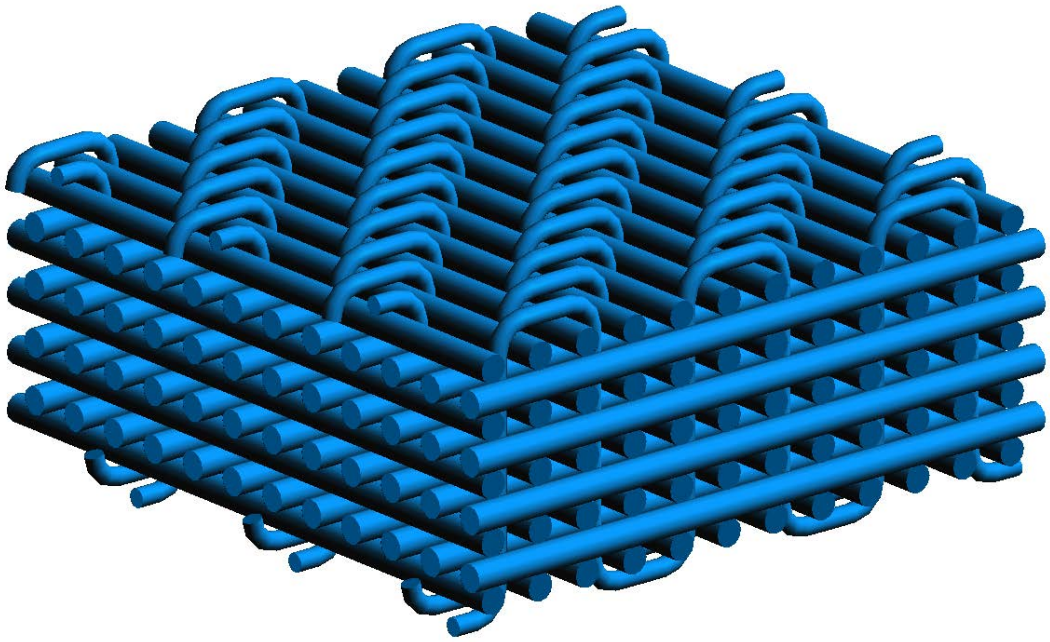

Figure 2.13. Initial configuration of orthogonally woven textile

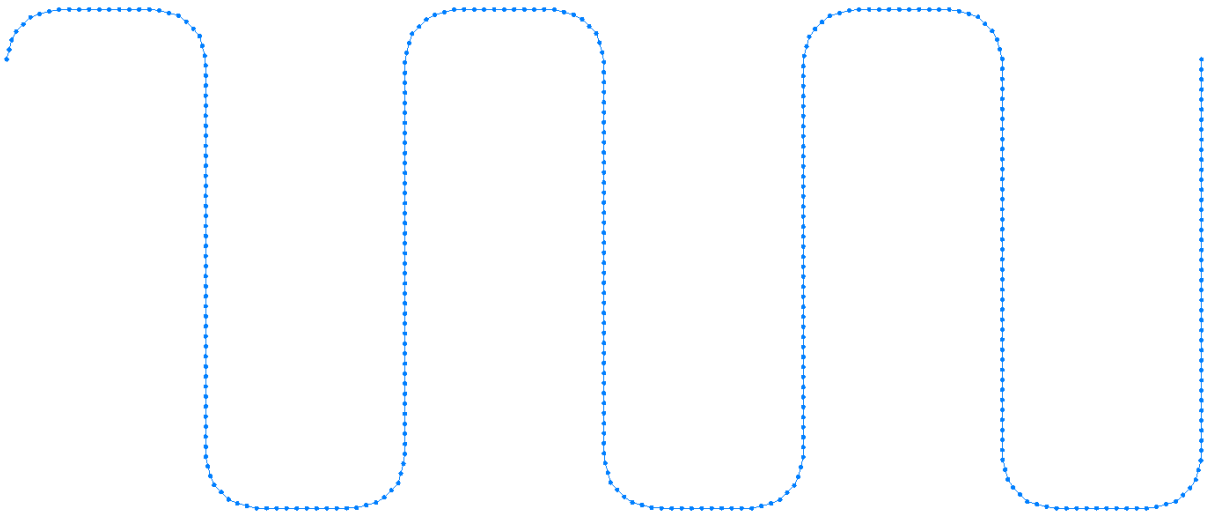

Figure 2.14. Mesh of a digital chain that represents a binder (spheres are not to scale)

A relaxation step is fundamentally an elasticity and contact analysis. The rod elements of the digital chains expand or contract, and contact is modeled between the rigid spheres. Before the relaxation, the ends of the digital chains are fixed, and a tensile prestress is specified within each digital chain, which will encourage the rod elements to contract. During a relaxation step, the prestress within the digital chains will cause the digital chain to move by itself, even without contact with another digital chain. The motion of the digital chains is reduced via a damping factor, which the user must control. For the first relaxation steps, a small damping factor is needed to prevent digital chains from moving too far. After much of the space has been removed between fibers, a larger damping factor can be used. Actual textiles do not exhibit nearly as much space between tows as the model shown in Figure 2.13. So, to encourage the digital chains to move towards the midplane of the model, the ends of the digital chains were modified outside of VTMS. VTMS can output the fabric geometry to a file between each step. So, a Python script was developed to move the ends of the digital chains to the midplane, and the modified fabric was imported back into VTMS. An example of model with the ends of the digital chains moved to the midplane is shown in Figure 2.15. Relaxation steps with just tension in the digital chains will consolidate the textile some in the out-of-plane direction, though not enough to create a realistic geometry.

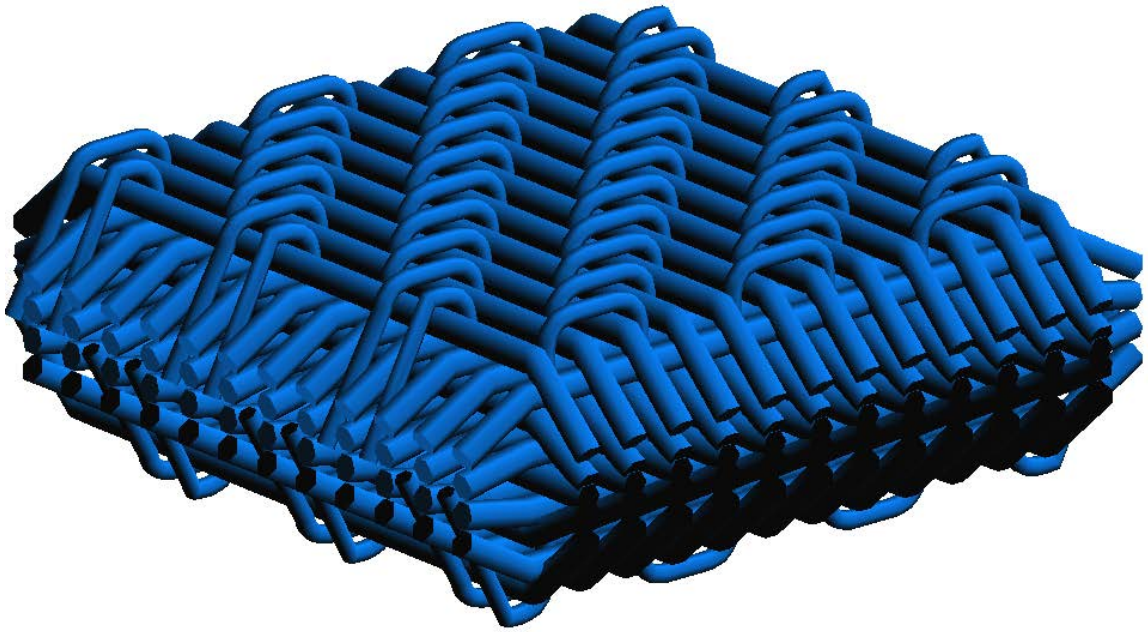

Figure 2.15. Configuration after ends of digital chains have been moved towards the midplane

In addition to the tension within the digital chains causing the digital chains to move, rigid or flexible molds can be used to compact the model in the out-of-plane direction during relaxation steps to yield a more realistic textile geometry. Figure 2.16 shows an example of a textile model with rigid planes used for compaction. Multiple relaxation steps are required to remove the unrealistic space between the digital chains. Typically, 10 to 20 relaxation steps are required for an orthogonally woven model with 5 weft layers. Thicker models require more relaxation steps to remove the space between digital chains. Figure 2.17 shows a textile model after several relaxation steps.

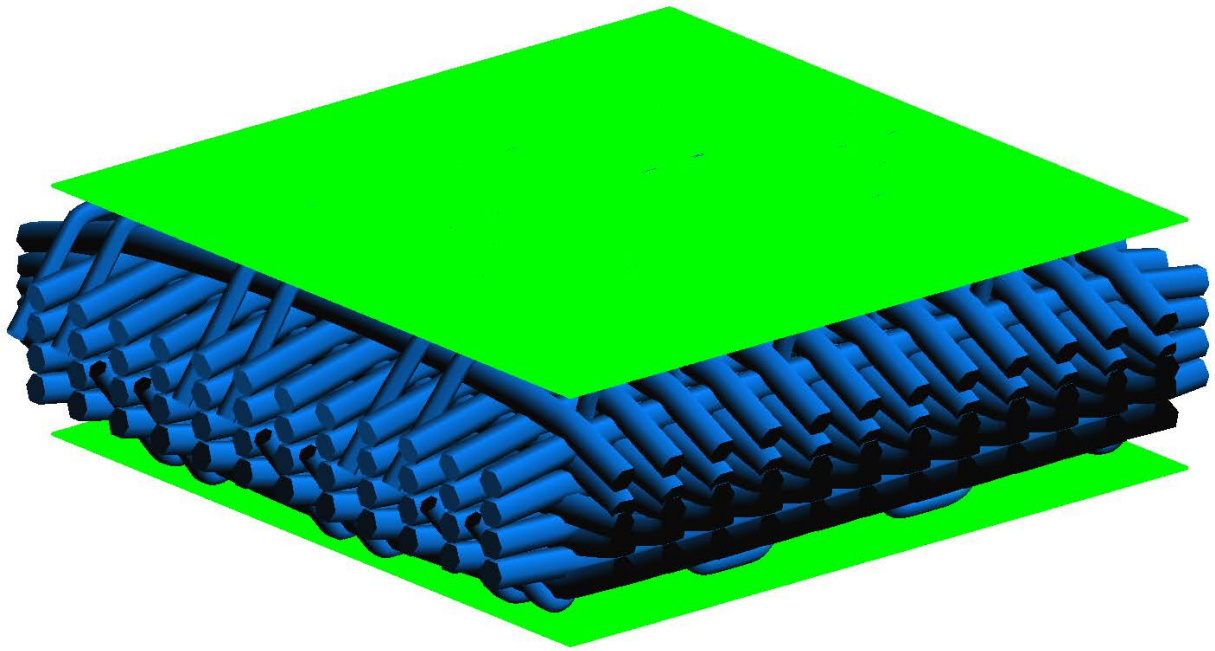

Figure 2.16. Textile model with rigid planes for compaction during relaxation steps

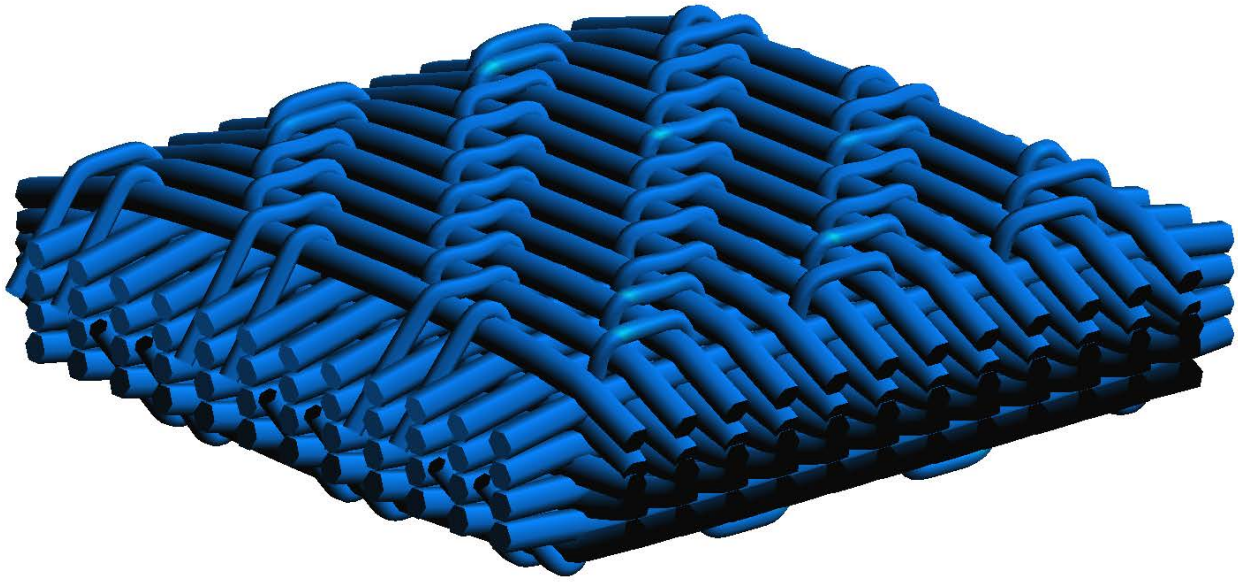

Figure 2.17. Textile model using one digital chain per a tow after several relaxation steps

After the textile model with one digital chain per a tow resembles a realistic textile more closely, each digital chain is replaced with multiple digital chains. When only one digital chain is used per a tow, the cross-section of a tow cannot deform and remains circular, but when multiple digital chains are used to represent a tow, the tow cross-section is defined by the collection digital chains and can therefore deform as the digital chains move relative to each other. As the model size or number of digital chains increases, it quickly becomes computationally expensive, so a compromise is needed between fidelity and computational cost. For the 2x2 orthogonal weave shown in the preceding figures, 10 digital chains per a tow were used, and each relaxation step required about an hour on an overclocked Intel Core i7-5820k processor with overclocked DDR4 RAM at a frequency of 2800 MHz.

Figure 2.18. Textile model with 10 digital chains per a tow

When multiple digital chains are used per a tow, the number of contact elements increases nonlinearly. As a result, the elasticity and contact problems become much more complicated and expensive. To ensure digital chains do not pass through each other during a relaxation step, the rigid planes must be moved in smaller increments and the damping factor, which reduces how much the geometry changes during a relaxation step, must be carefully chosen. As a result, when multiple digital chains are used per a tow, more relaxation steps are required than when one digital chain is used per a tow.

Removing the unrealistic space between tows near the midplane was one of the major challenges in generating relatively thick 3D textile models in VTMS. As the rigid planes compact the textile model, the cross-sections of the outer layers of tows deform significantly before the space between tows is removed. This issue exists because contact is the mechanism used for removing the space between tows. Compaction of the model is to some extent realistic, but an actual specimen is compacted starting with an initial state that does not have large amounts of space between tows. For future studies, a different approach is needed for removing the space between tows before compaction is simulated. However, with the current features of VTMS, a tedious work around was partially successful. Typically, just the ends of the digital chains are fixed during relaxation steps, but VTMS allows the user to specify the displacements of all the nodes within a digital chain. So, tows can be specified not to move, while a rigid plane compacts specific layers at a time. The digital chains with specified zero displacements are still considered for contact, so digital chains that can move do not pass through others. By specifying the digital chain within

selected tows to have zero displacements, a few layers of tows can be compacted at a time, which will remove the space between tows near the midplane without as much unrealistic deformation in the cross-sections of the outer tows. Figure 2.19 shows the textile model after the warps and wefts have been compacted, while every node in the binders were specified to have zero displacements.

Figure 2.19. Textile model after relaxation and compaction of warps and wefts (every node within the binders were specified to have zero displacement)

After all of the warps and wefts have been incrementally compacted, the displacements in the warps and wefts are specified to be zero, and the binders are allowed to move. A

large tension is specified for the binders to remove the excessive slack. During this step, the rigid planes are not used to compact the textile. Figure 2.20 shows the textile model after multiple relaxation steps. Recall that the elastic rod elements of a digital chain may shorten in length during the relaxation steps due to the tensile prestress. The rod elements can shrink significantly during this step due to large amounts of slack in the binders, see Region A in Figure 2.19. Because of this, each digital chain within the binders must be remeshed between relaxation steps to ensure the elastic rods are not so short that rigid spheres within the same digital chain come into contact with each other, which would cause numerical instability.

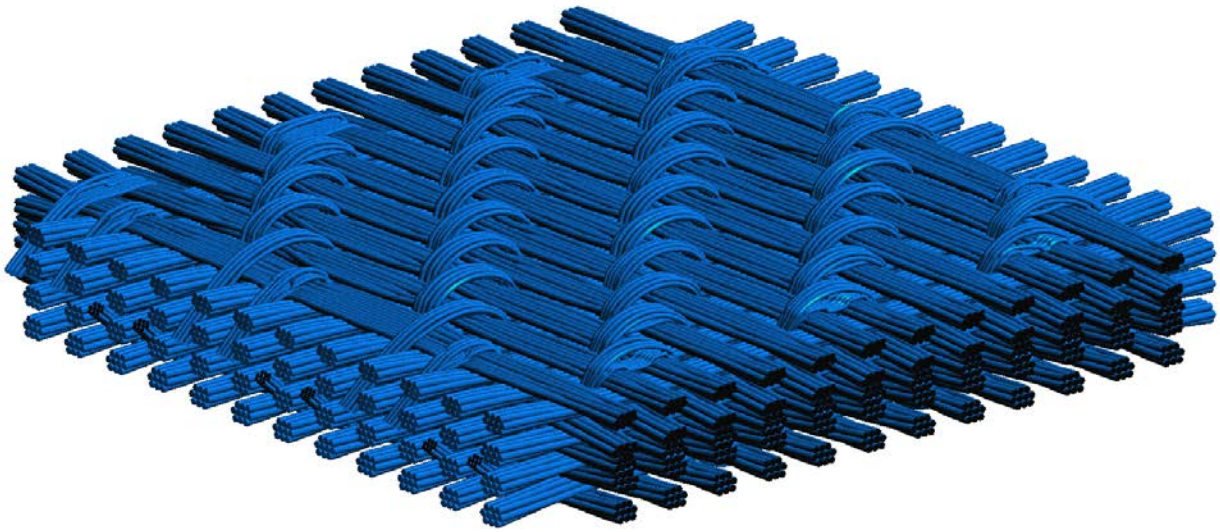

Figure 2.20. Textile model after relaxation of all digital chains (no compaction through rigid planes yet)

Finally, only the ends of each digital chain in the textile model are fixed, and rigid planar molds are used to compact the textile. This results in flatting of the binders on the outside of weave, where the rigid planes come into contact with the digital chains, as shown in Figure 2.21. For the textile model shown, more than three days of wall-time was required to perform all the relaxation steps.

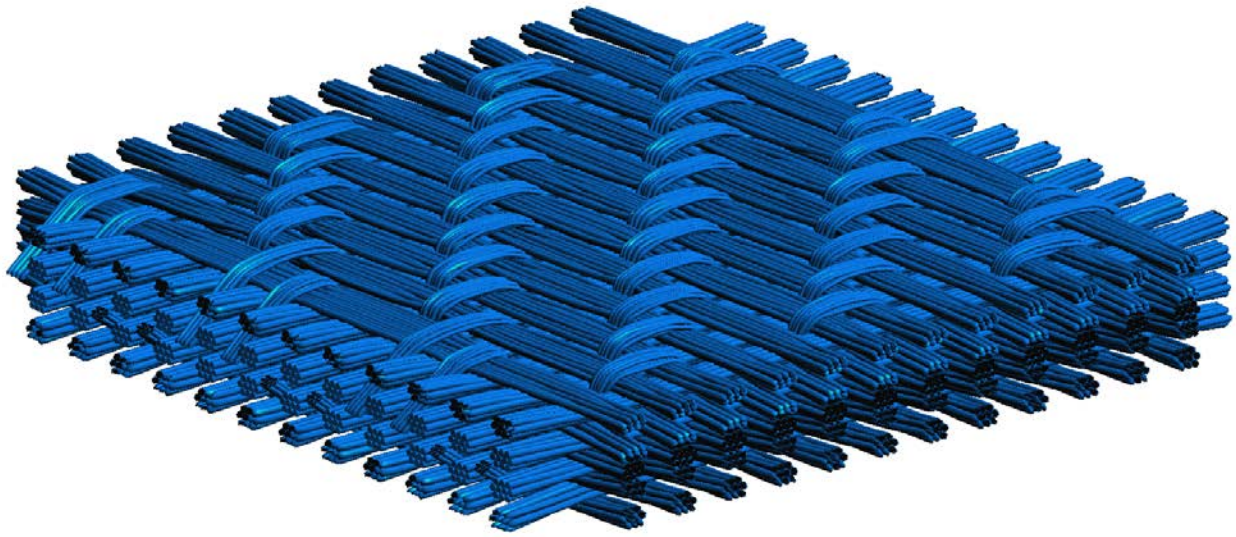

Figure 2.21. Textile model after compaction of all digital chains using rigid planes

The next step to create a finite model for a textile composite is to create a surface representation of each tow. To create a surface for a tow, cross-sections are determined along the path of the tow by fitting a spline around the bundle of digital chains that represent the tow. However, a few issues emerged when a creating the surfaces of the tows.

First, VTMS failed to create some cross-sections, resulting in gaps along the tow path and extreme distortion near the ends of the tows, as shown in Figure 2.22.

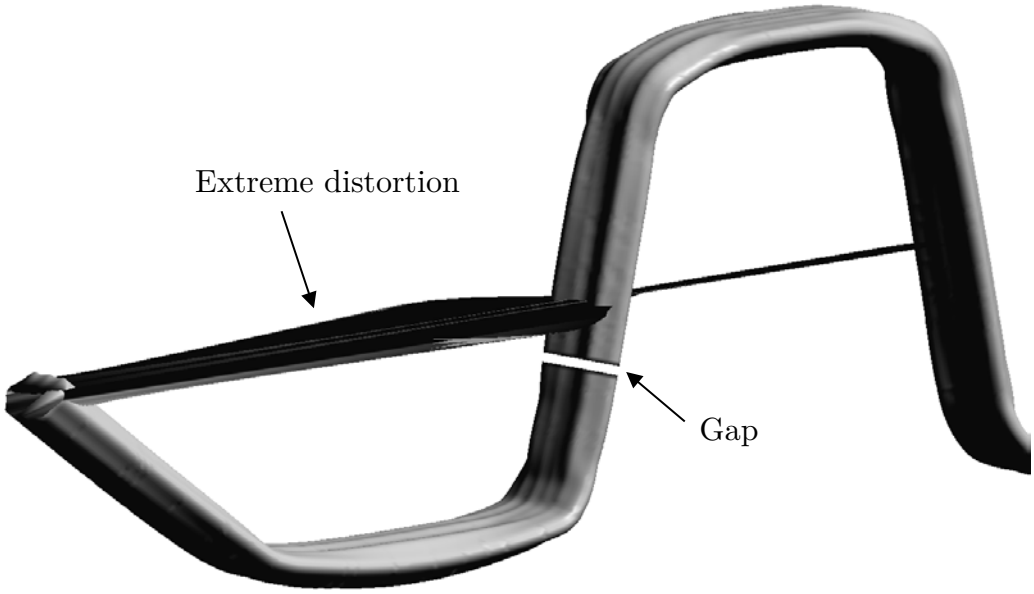

Figure 2.22. Surface for a binder tow with extreme distortion and gaps

To alleviate the distortion near the ends to the tows, a Python script was developed to clip the digital chains. The result is shown in Figure 2.23. However, VTMS still failed to produce some cross-sections, resulting in gaps in the tow surfaces, even after the digital chains were clipped, as shown in Figure 2.24. Refining the number of cross-sections created for tows with discontinuous surfaces fixed the issue.

Figure 2.23. Textile model with digital chains clipped to smaller region to avoid surface meshing issues

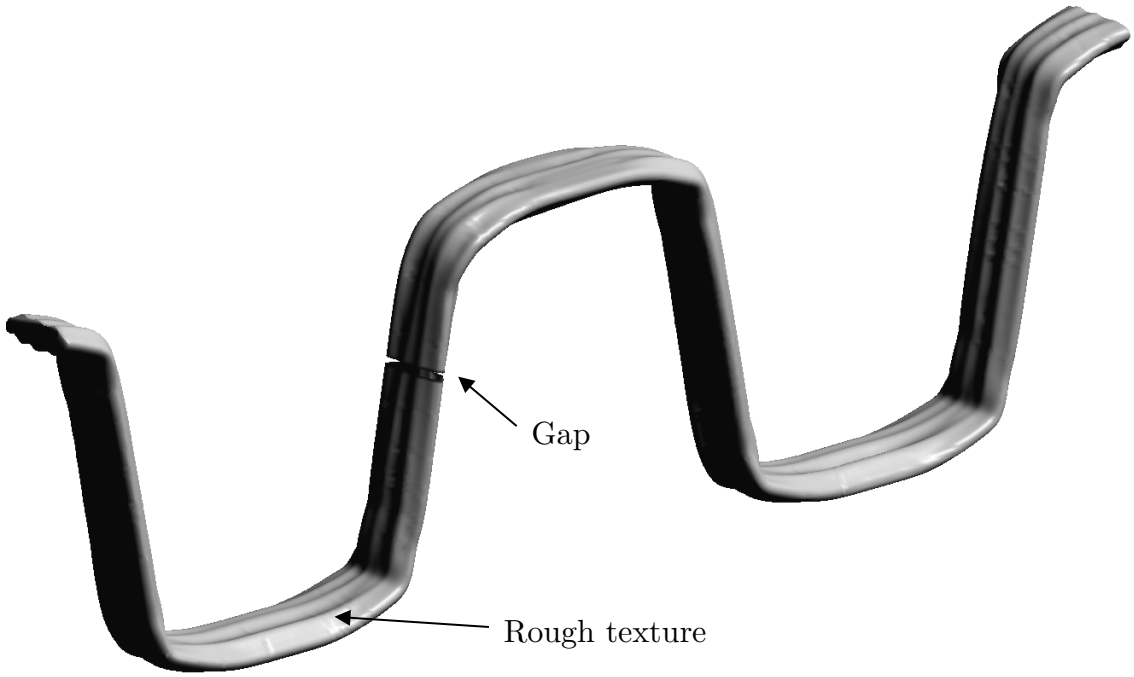

Figure 2.24. Surface for a binder tow after clipping the digital chains with a gap within the surface

The second issue encountered when creating the tow surfaces was unrealistic rough texture in the tow surface, which is also shown in Figure 2.24. When VTMS creates the cross-sections, a spline is fit around the bundle of digital chains that define the tow. The minimum radius of curvature for the spline is the radius of one of the digital chains by default, resulting in a greatly textured surface for the tow. By request, AFRL added an option to control the minimum radius of curve of the splines. By specifying a large minimum radius of curvature (called the magnification factor under "Dial" settings in VTMS), the tow surfaces were much smoother, though some perturbations remained, as shown Figure 2.25.

Figure 2.25. Tow surface using splines with a minimum radius of curvature 10x that of the digital chains

By clipping the digital chains, modifying the number of cross-sections created for each tow, and increasing the minimum radius of curvature used for the splines that defined the tow cross-sections, smooth surface representations of the tow geometry were created, as shown in Figure 2.26. The next steps would be clip the tow surfaces down to the analysis region, create a volume mesh for each tow, and create a compatible mesh for the matrix that surrounds the tows. However, generating the volume meshes that can be used for standard FEA requires that there are no interpenetrations between tows. Unfortunately, the contact problem with digital chains only approximates the tow surfaces. Consequently, many interpenetrations typically exist between the tow surfaces. VTMS includes a method for removing interpenetrations, but the algorithm is not precise enough to allow the creation of a conforming, non-intersecting volume mesh for each tow. Consequently, to remove the interpenetrations, the cross-sections for each tow are shrunk towards the centroid by a small amount, referred to as compensation within VTMS, resulting in some matrix material between each tow and its neighboring tows. The tow surfaces after compensation is shown in Figure 2.27. This removes many of the interpenetrations, but some very small interpenetrating areas typically still exist. The remaining interpenetrations were addressed by manually moving some tows after the next step of clipping the geometry down to the analysis region.

Figure 2.26. Tow surfaces based on the group of digital chains used for each tow

Figure 2.27. Tow surfaces after cross-sections are contracted

Next, the tows are clipped in-plane down to the analysis region, as shown in Figure 2.28. It should be noted that it is important to specify a tight tolerance for the clipping of the tows within VTMS. Otherwise, the nodes on clip planes will not lie precisely enough on the respective clip plane.

Figure 2.28. Tow surface clipped down to the analysis region

When VTMS creates the clipped tows, the surfaces are closed on all sides and described by a surface mesh of general polygons. Some polygons on the boundary of the clip region can have over a thousand sides and be quite complex, but most are faceted triangles and quadrilaterals. The in-house tool takes the surface geometry in this state, and creates the finite element model, but a few issues persist that require manual modification of the tow surfaces. Some interpenetrations typically occur where binders cross over wefts, and a significant amount of space between tows near the midplane is created when the tow cross-

sections are shrunk, as shown in Figure 2.29a. Tows are manually moved around to remove interpenetrations remaining interpenetrations and more evenly distribute the tows through the thickness of the model, as shown Figure 2.29. To manually rearrange the tows, translations were specified for single or even groups of tows.

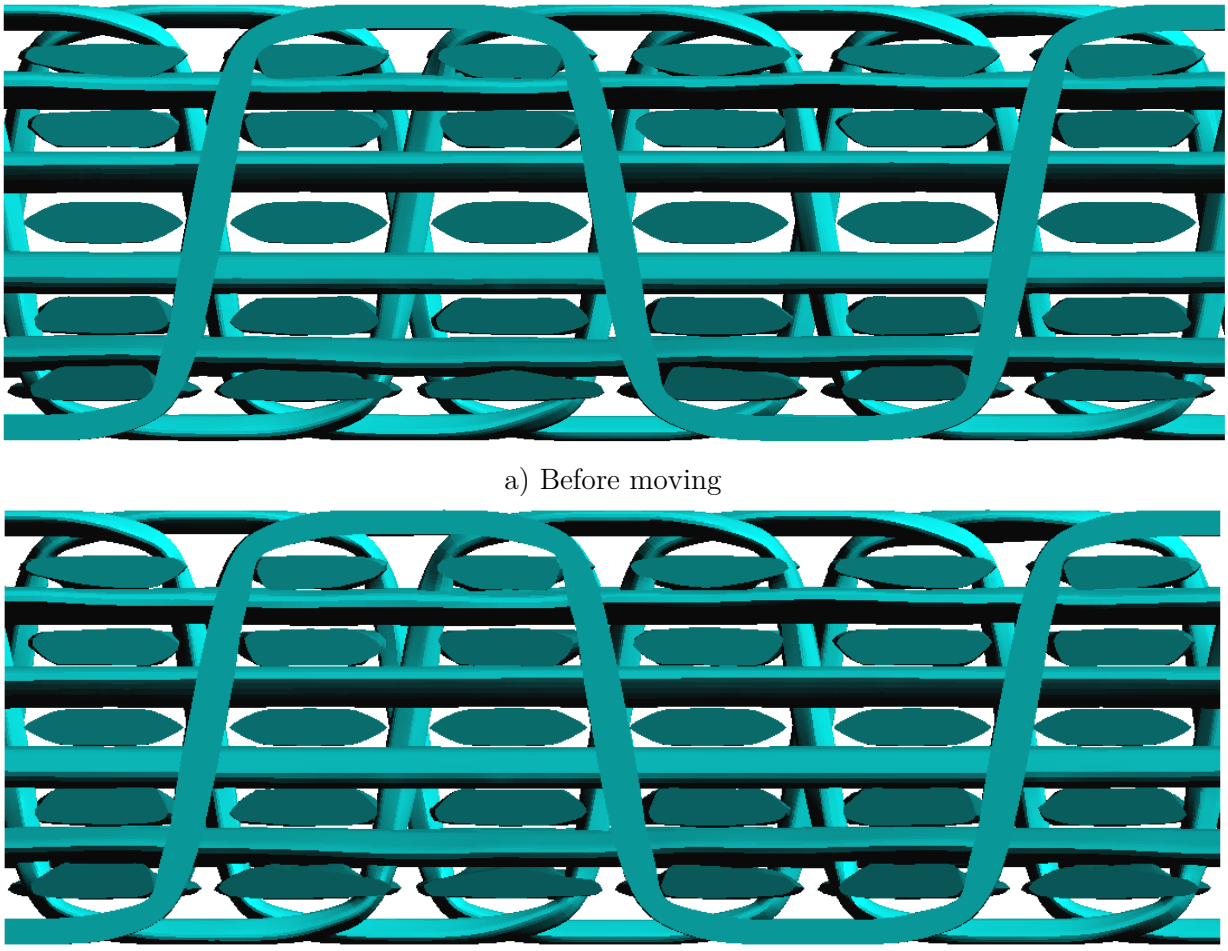

a) Before moving

b) After moving

Figure 2.29. Clipped tow surfaces before and after manual adjustment

Creating a Finite Element Model from VTMS Tow Geometry

Using the tow surfaces clipped down to the analysis region that VTMS creates, an in-house tool was developed to create the finite element model. The algorithm for creating the final finite element model can be divided into five steps:

1. Creation of surface meshes for the tows
2. Creation of a surface mesh for the matrix that is compatible with the tow surface meshes
3. Creation of a volume mesh for the matrix
4. Creation of volume meshes for the tows and determination of local data, such as the local coordinate systems and local fiber volume fraction
5. Consolidation of volume meshes and final modifications

Step 1: Creation of Surface Meshes for the Tows

The surface description of the tows that VTMS produces is a faceted, polygonal surface mesh. Some polygons have over a thousand edges, and some of the edges in the surface mesh can be very short. The polygonal surfaces could directly be given to third-party tetrahedral meshing library, like TetGen, but the volume mesh ends up having excessive refinement in some areas. [39] Third party meshing libraries do not typically coarsen the surface elements of the input, so TetGen creates very refined regions near surface elements with unusually small edges. Instead, a more uniform surface mesh for the tows was created from the polygonal surface mesh that VTMS creates.

VTMS writes clipped tow (.ctw) files that contain a description of the tows. The in-house code extracts the polygons and list of points that define the path of the tow, called medials, which are the centroids of the tow cross-sections, and stores the information within a data structure. The medials will be sued to define tow path, but VTMS sometimes fails to calculate the centroid of some cross-sections, resulting in "-1.#IND00000" in place of a number within the file. It is important to filter these medials out.

After all the .ctw files have been read, the polygons that lie on the in-plane boundaries of the textile model are identified, and the nodes are adjusted to lie exactly on the extent of the mesh, removing any numerical error in the coordinates. Steps later in this process depend on the boundary nodes exactly lying on the clipping plane. VTMS uses a coordinate system where the y-axis is the out-of-plane direction of the textile, but the out-of-plane direction is typically the z-axis in the literature, so the tow surface and medial coordinate are rotated. Next, the tow surfaces and medials are scaled so that the thickness of the textile model is a more realistic value. The textile model of the 2x2 orthogonal weave in this work is assumed to be 5.25 mm thick. Finally, all the surface meshes are translated so that the minimum coordinate in the textile mesh will lie at the origin.

At this point the surface mesh for each tow is modified so that the edge lengths are more uniform and every surface element is a triangle instead of a general polygon. The processes illustrated for a very simple mesh of rectangles in Figure 2.30. First, any edges that are less than ¼ of the average edge length are collapsed, resulting the two nodes being combined. This can result in polygons with less than three edges, which are removed.

When an edge is removed, the nodes are combined into one, but special care must be taken to ensure that if one or both of the two nodes lie on the boundary of the textile then the node lying on the fewest boundary planes is the one removed. For example, if one node lies on the $x = 0$ and $y = 0$ planes and another lies on the $x = 0$ plane, then the node lying on the $x = 0$ plane should be the one removed. See parts a and b of Figure 2.30 for an illustration of the approach. This strategy ensures that the boundary of the textile remains exactly planar.

After the short edges are removed, the unusually long edges are refined. Recall the goal is create a surface mesh with a more uniform mesh refinement. Edges larger than four standard deviations from the average edge length are subdivided into the edge length divided by the average edge length (rounded to the nearest integer) smaller edges. Refer to parts c and d of Figure 2.30.

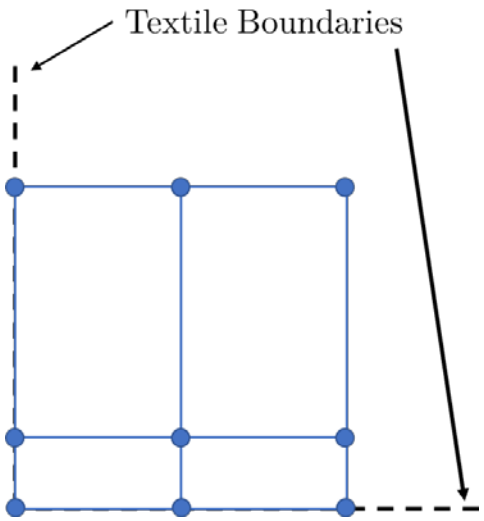

a) Initial polygon mesh with long and short edge lengths

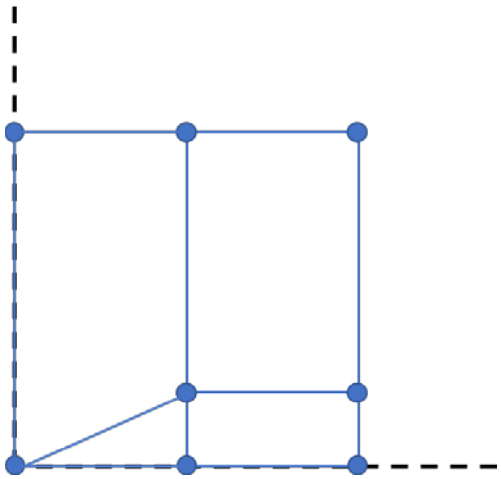

b) Mesh after first short edge is collapsed, bottom-left node is kept since it lies on more boundaries

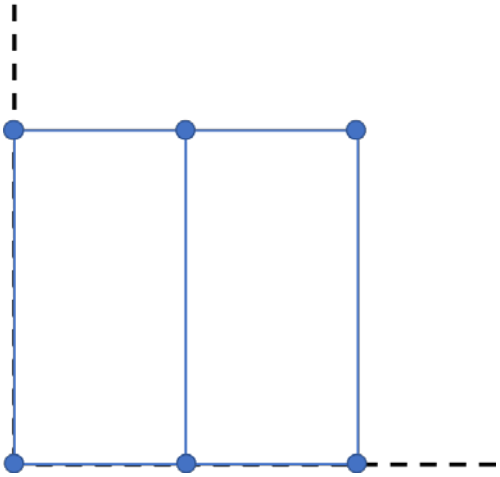

c) Mesh after all short edges have been collapsed

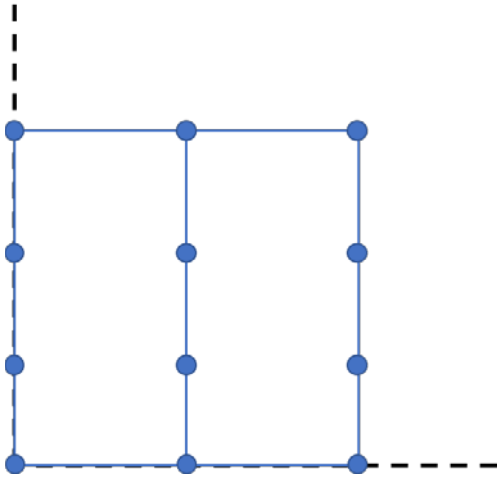

d) Mesh after long edges have been refined

Figure 2.30. Illustration of modification of polygonal meshes by collapsing short edges and refining long edges

The final step for creating the tow surface meshes is replacing the polygon elements with a submesh composed of triangles. The polygons created by VTMS are faceted, so all of the nodes in an element lie in the same plane, within numerical error. So, Triangle is used to create the triangle submesh that will replace the polygon. Since the polygon edges have been modified to be more uniform at this point, Triangle is only allowed to add nodes within the closed domain and not along any edges of the polygon. For the model in this work, the maximum area a single triangle can have was specified to be $1.4\bar{s}^2$, where $\bar{s}$ is the average edge length of the polygon. As mentioned before, some polygons can have over a thousand edges. The largest polygons in the surface meshes exist where a binder was cut by a clipping plane in VTMS. Figure 2.31 shows a triangle submesh for one of these polygons.

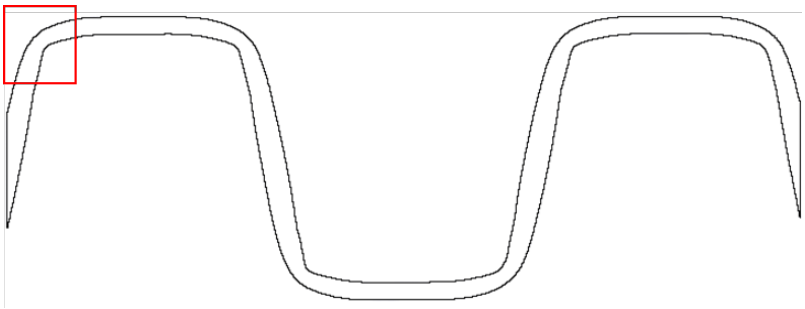

a) A large polygon from where a binder is cut by the textile boundary containing over a thousand edges

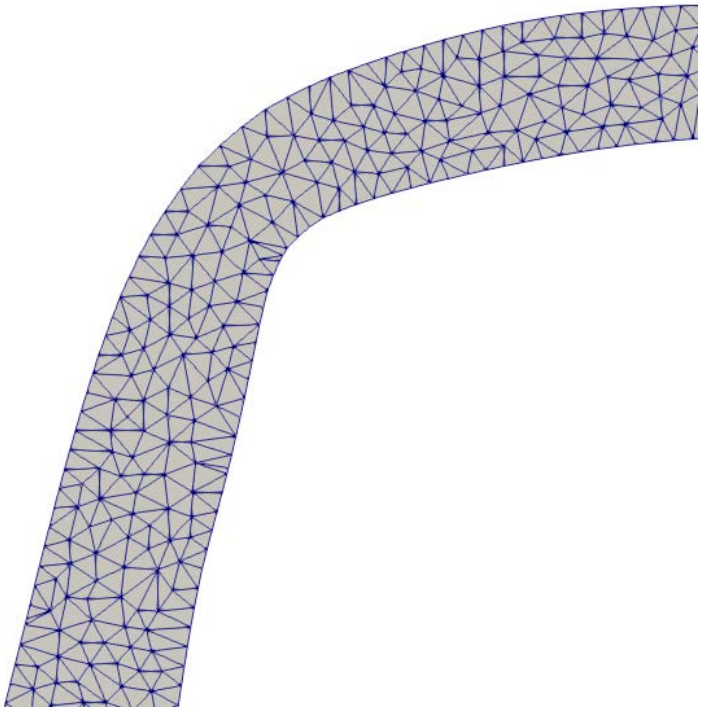

b) Closeup of a triangle submesh for highlighted red region of the polygon in part a

Figure 2.31. Submesh of large polygon from the surface of a binder tow that was cut by the clipping plane in VTMS

Once each polygon is replaced with the triangle submesh and duplicate nodes have been removed, the tow surface meshes are complete. Figure 2.32 shows the final surface mesh for one of the binders cut by a clipping plane within VTMS. This procedure results in a

significantly more uniform mesh refinement for the surface meshes, which allows quality volume meshes to be easily created.

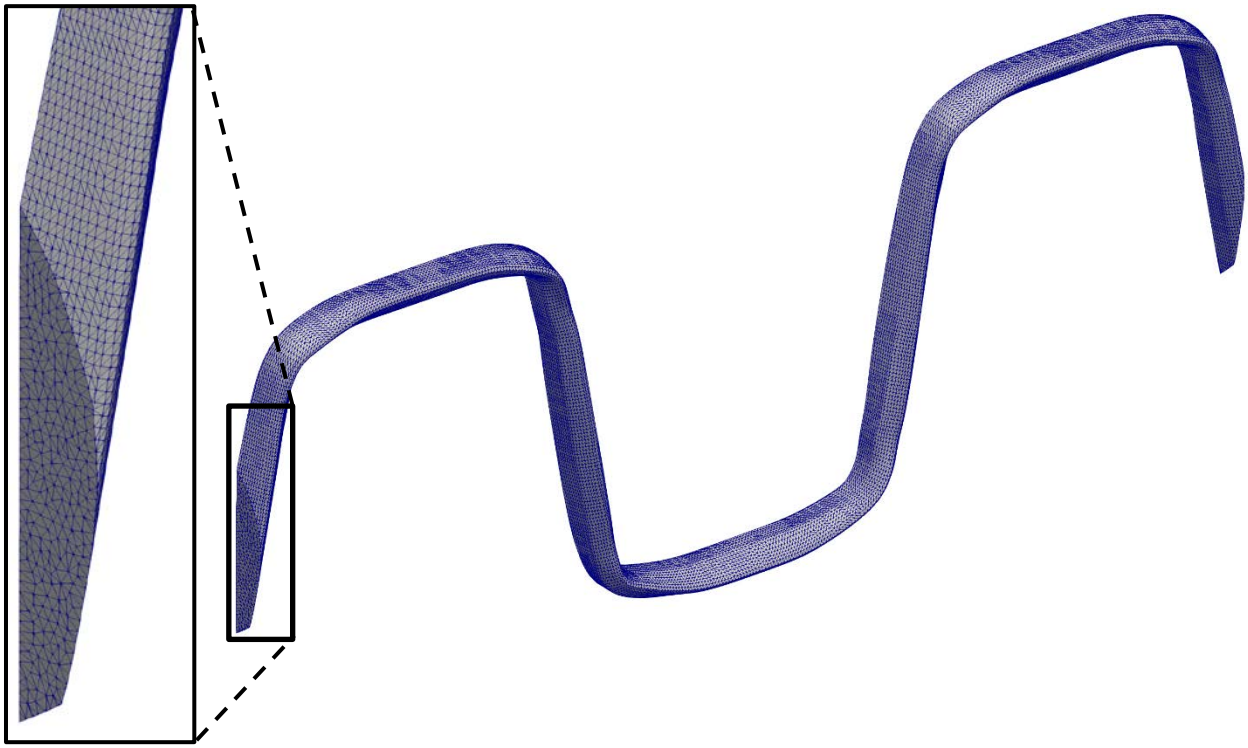

Figure 2.32. Final surface mesh for a binder tow that was cut by the clipping plane in VTMS

Step 2: Creation of a Surface Mesh for the Matrix

Standard FEA requires a compatible mesh for all constituents. To create a surface mesh that is compatible with the tow surface meshes, a surface mesh for each boundary face of the textile model is created, excluding regions where the tows intersect the textile boundary

since these regions are already meshed. To create a mesh for one of the boundaries of the textile model, the elements of the tow surface meshes that lie on the boundary are identified. Next, the all edges of those elements that are sheared by two elements are removed. which leaves the edges that define the intersection of the tow and the boundary of the textile model. Next, any isolated edges or nodes in the tow surface meshes (those not part of elements that lie on the boundary) that touch the textile model boundary are found. Finally, edges along the outside of the textile model boundary are created. Combining all the nodes and edges discussed in this paragraph and specifying where tows intersect the boundary to be holes, a triangle mesh is created for the boundary of the textile model. Figure 2.33 shows a mesh created for the boundary of the textile model that is normal to the y-axis with the most positive y-value.

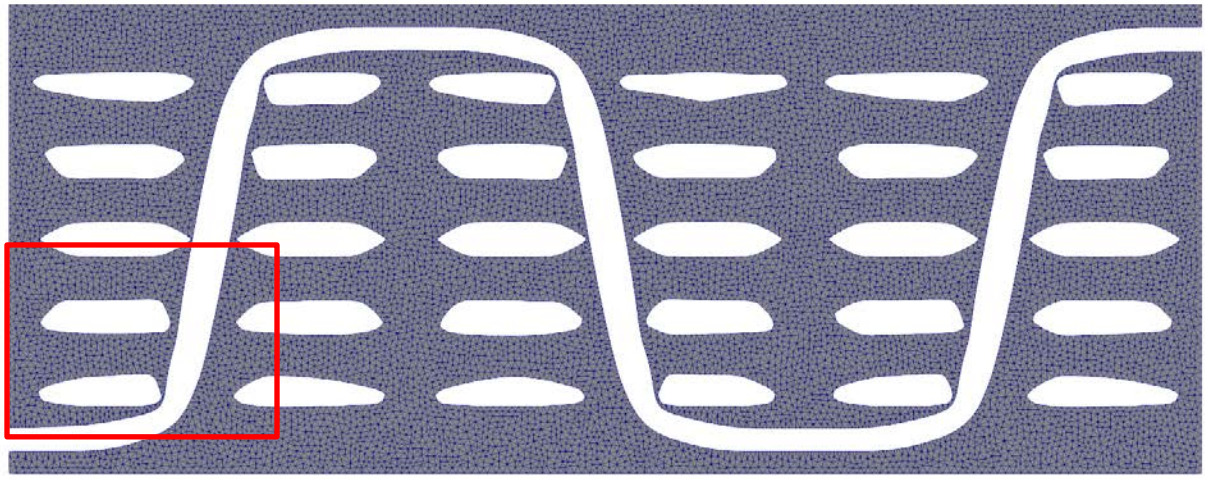

a) Mesh of entire textile model boundary

b) Mesh of region highlighted in part a

Figure 2.33. Mesh of textile model boundary that is normal to the y-axis

Next, a copy of the tow surface meshes is created and elements that lie on the textile boundary removed, as shown in Figure 2.34. The outward normal of all elements in the

copy of tow surface meshes is reversed, so that the outward normal of the surface points away from the matrix volume to be meshed. This is important since TetGen requires that the surface elements' normal vectors point outward from the volume to be meshed. The surface mesh of the neat matrix is created by combining the mesh of the textile boundaries and the modified copy of the tow surfaces. Figure 2.35 shows the surface of the neat matrix pocket of a 3D textile. In the figure, the surface of each tow is colored differently to clearly show that the matrix surface is compatible with the surfaces of the tows.

Figure 2.34. Surface meshes of tows without elements that lie on the textile boundaries

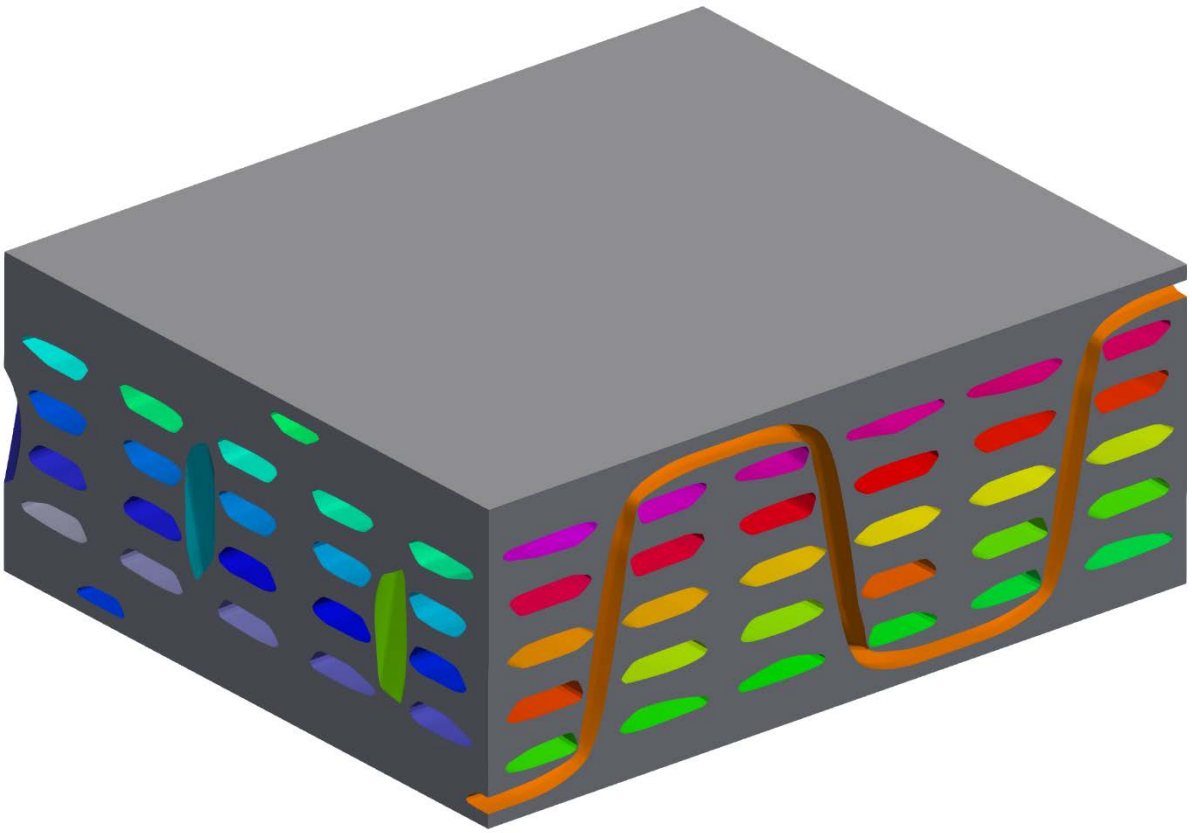

Figure 2.35. Surface mesh of the neat matrix pocket of a 3D textile (the surface of each tow is colored differently for clarity)

Step 3: Creating of a Volume Mesh for the Matrix

With a quality surface mesh for the matrix, creating a quality tetrahedral volume mesh becomes easy using TetGen. Due to the large textile files, it would be relatively expensive to use the command line TetGen program. In this work, the native C++ library for TetGen was used so that that file I/O could be avoided. TetGen includes a data structure to store the meshing flags for the native interface. Regarding a detail that is not well documented, it is important to ensure the "object" flag is set to tetgen::tetgenbehavior::POLY; otherwise,

the volume mesh will fail to stay inside the complex closed surface mesh. Additionally, the quality of the tetrahedral mesh heavily relies on the "minratio" and "mindihedral" flags. For the meshes used in this work, 1.1 and 14 were used respectively.

Step 4: Creation of Volume Meshes for the Tows and Determination of Local Data

This step begins with creating tetrahedral volume meshes for the tows. For this work, TetGen was used to create the volume mesh with the same parameters as described in step 3. The resulting volume meshes exhibited nearly uniform refinement within each tow, though the elements within binders were smaller than the elements within the warps and wefts, as shown in Figure 2.36. This was largely due to the higher number of cross-sections needed for the binders than the warps and wefts to avoid gaps within the surfaces, which was described earlier in the chapter. Because there were more cross-sections, a finer surface mesh was created, which led to a finer volume mesh. The binders are expected to experience complex stress states, so the increase mesh refinement may be justified, though no studies were conducted to determine an optimal refinement for each tow.

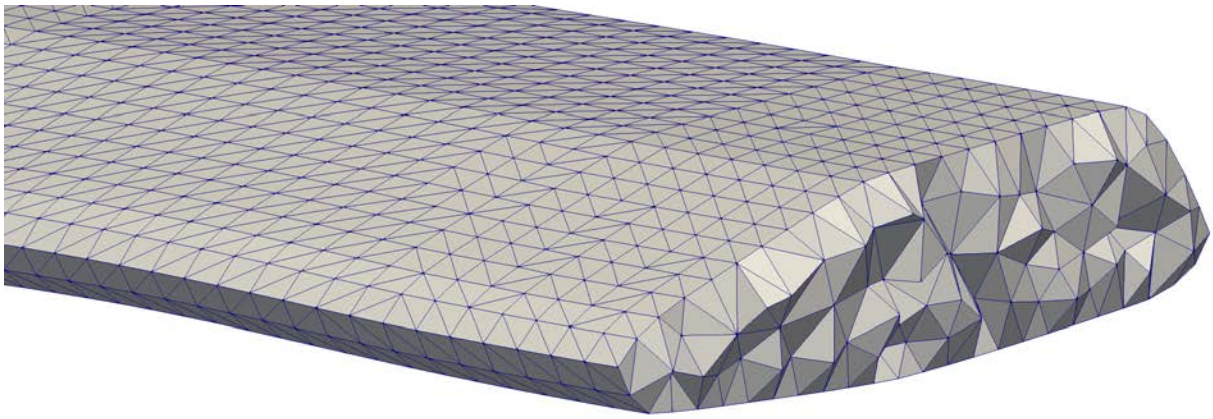

a) Warp

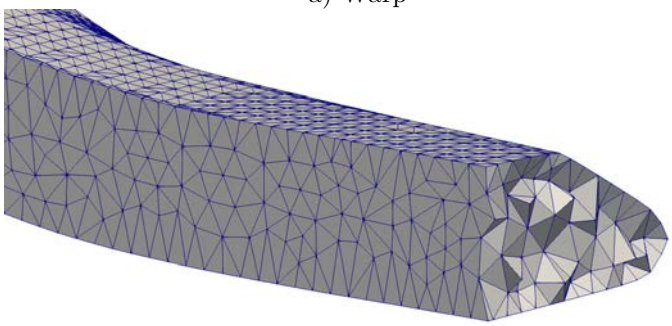

b) Binder

Figure 2.36. Clipped view of a typical warp and binder tow showing the interior mesh refinement of each (shown at same scale)

After the volume meshes are created for the tows, the local material coordinate system and local fiber volume fraction is determined. The local material coordinate system within the tows is calculated for each node in the volume meshes, while the local fiber volume fraction is calculated per element basis. To define the fiber direction within a tow, a piecewise cubic spline is fit to the medials (area centroids of each cross-section VTMS created), and the tangent of the spline is assumed to be the direction aligned with fibers in the tow. For each node within the tows, the local x-axis is determined by finding the closest

point on the spline of the tow path and using the unit tangent vector as the direction of the local x-axis. The local y-axis is assigned to be parallel to the cross product of a vector parallel to the global z-axis and the local x-axis. The local z-axis is defined by the cross product of the local x and y-axes, which define a right-handed coordinate system.

To assign the local fiber volume fraction, the cross-sectional area is calculated at 1000 sampled points along the tow's path and are compared to a nominal cross-sectional area. The nominal cross-sectional area of a tow is given by its volume divided by its length, $\bar{A}_i = L_i/V_i$. The volume of each tow is calculated by summing the volume of each element within the tow, which will be denoted by V_i for tow i. The length of each tow, L_i, is calculated using the sum of the distance between the medials.

Let t be a parametric coordinate along the tow path that goes from 0, at one end of a tow, to 1, at the other end of the tow. The medials are fit using a spline as a function of t. Next, the cross-sectional area is computed at 1000 sample points along the tow path, and the cross-sectional area, A, is fit using a spline as a function of t. Given a nominal fiber volume fraction for the tows, $\bar{V}_f$, the local fiber volume fraction for a position along the tow path is given by Eqn. (2.2), where $V_f(t)$ is the local fiber volume fraction as function of the parametric coordinate, t, and $A(t)$ is the cross-sectional area at t.

$$V_f(t) = \frac{A(t)}{\bar{A}_i}\bar{V}_f \tag{2.2}$$

The local volume fraction is computed at 1000 discrete samples along the tow path, and a spline fit of the local volume fraction as a function of t is also created. Given a (x, y, z) coordinate, the closest point on the tow path can be determined, along with the corresponding value of t. It should be noted that determining the value of t for a point on the tow path closest to another point requires a iterative solution, and the bisection method was used in this work to solve for t. For each element in a tow, the local fiber volume fraction is determined by computing the centroid of the element, finding t at the closest point on the spline of the tow path, and calculating the fiber volume fraction at that value of t using the spline fit. Figure 2.37 shows the local fiber volume fraction within the tows for an orthogonally woven textile with a nominal fiber volume fraction of 0.6. Since the cross-sections of the tows can deform significantly when VTMS simulates the relaxation and compaction of the tows, the local fiber volume fraction can vary significantly throughout the tows. The impact of the variation of fiber volume fraction is evaluated later in this work.

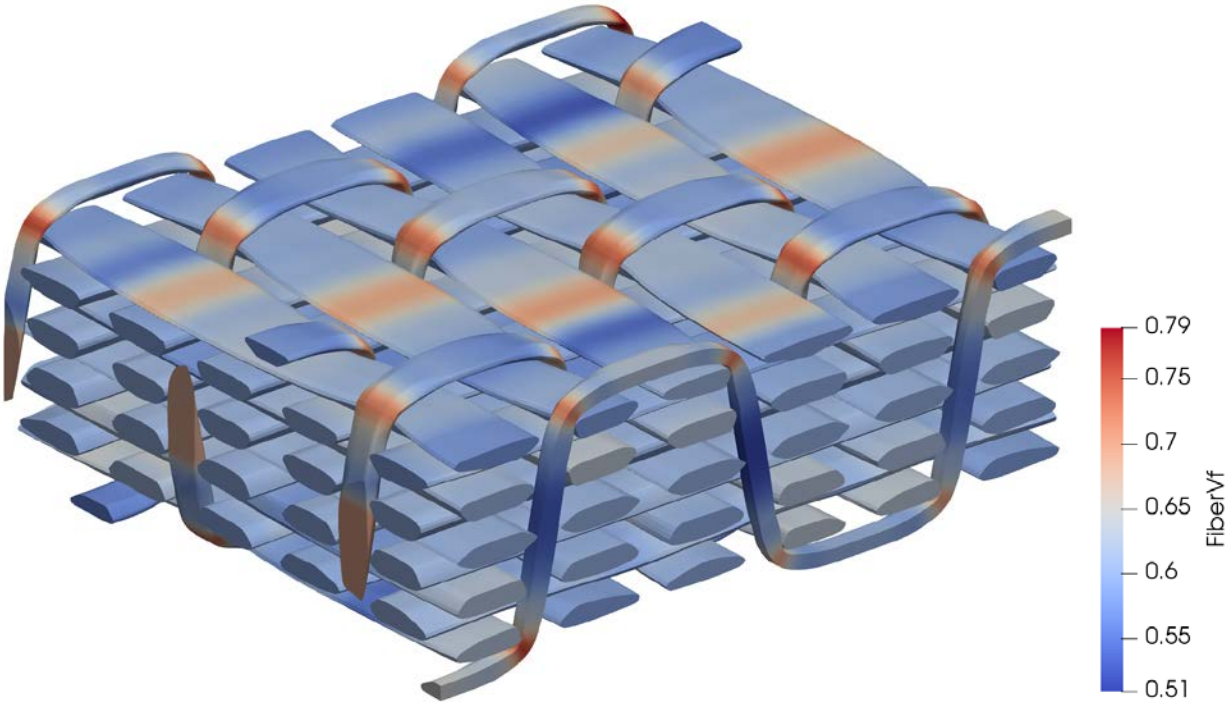

Figure 2.37. Contours of local fiber volume fraction in the tows with a nominal fiber volume fraction of 0.6

Step 5: Consolidation of volume meshes and final modifications

For the final step of creating the 3D textile models, the volume meshes of the tows and neat matrix are combined, duplicate nodes are removed, linear tetrahedrals are replaced with quadratic tetrahedrals, the mesh is divided into partitions for distributed computing, and the files are written in an in-house file format described later in this chapter.

2.2.4 Lessons Learned

A framework was developed to create fiber/matrix models with varying shapes and boundary types, multiscale laminate composite models, and non-idealized 3D textile models. Many lessons were learned from developing the framework, and this section describes some that may help other researchers developing similar tools.

First, regarding third-party libraries, several libraries were invaluable in this framework. Triangle and TetGen both provided great performance and integrated easily into the code. More general meshing libraries that create quality quadrilateral and hexahedral meshes would be a very nice addition, but few existing tools offer a C++ API and many require licensing. Two k-d tree libraries were used in the code to efficiently conduct spatial searches: libkdtree++ and nanoflann. [40] [41] They each take very different approaches to the data structures. Using a self-contained approach, libkdtree++ stores pointers to the objects (whether it be nodes, elements, etc.) within its own data structure and requires an accessor to get the coordinate of the object. It provides methods for finding other objects within a specified distance or the nearest objects to a given object. However, a limitation was discovered. The library threw exceptions when many objects were clustered in several different very small areas. If the objects were spaced out well, libkdtree++ required relatively little interface code and performed well. On the other hand, nanoflann builds a k-d tree on top of an existing vector of objects, and it does not modify the underlying vector. Building the k-d tree index information was expensive, and the templates used by nanoflann became quite complicated with already heavily templated classes within this framework. However, it was very memory efficient and fast once the index information was constructed. If objects are often inserted within an algorithm requiring k-d searches, then libkdtree++ is the way to go. If the objects are not inserted and the coordinates of the objects do not change when k-d searches are needed, then nanoflann seems the best choice.

Second, embracing generic coding for the mesh storage, geometry information, and algorithms that operate on meshes saved a lot of development time. Admittedly, templating within C++ can become quite verbose; so, the components of a framework that utilize numerous template should be selected carefully. Very few templates were used in other parts of the framework as whole, but the entire preprocessing code was based on C++ templates, allowing the finite element analysis and postprocessing codes to use the same code to store and manipulate the geometry, and perhaps most helpful, unified all the I/O algorithms for serial and distributed systems.

Third, creating Python bindings for all the preprocessing C++ code resulted in a convenient framework for creating very complicated models through scripting. Different Python scripts could create different types of models, whether it be fiber/matrix models, laminate composite models, or textile models. The Python scripts were organized into modules and then into a package. Providing the native libraries allowed other parts of the finite element and postprocessing C++ code to reuse a lot of the preprocessing code, while the Python package provided a more expressive way of creating models. The strategy worked quite well.

Finally, many issues were discovered with VTMS when creating large 3D textile models. The tools had not been tested for thick textile models before, and it was shown that relying on the build-in contact analyses to remove a lot of space between tows resulted in extreme deformation of the tow cross-sections before the space near the midplane was removed. Additionally, the user was required to guide VTMS through many of the steps and make

informed choices for many parameters. Because of this, the tool exhibited a steep learning curve to effectively utilize, and many work arounds were needed for existing issues. Another issue was the existence of interpenetrations between tows. Since contact is only modeled between the rigid spheres, interpenetrations between tows is inevitable when the textile is significantly compacted. Shrinking the tows worked, but unrealistic amounts of matrix resulted between tows. A more accurate approach to the contact problem or a method for creating compatible meshes between interpenetrating tows is needed. It should be noted that a third-party mesh boolean library, called Cork, was tested for removing the interpenetrating region between tows, but excessive refinement was created where tows interpenetrated near the textile boundaries. [42] Creating non-idealized textile models has come a long way in the last decade, but several issues make the tools very difficult to use for creating relatively thick textile composite models.

2.3 Development of a Scalable Finite Element Framework

The finite element software, called Beta2, was developed to be a scalable, extensible FEA framework. Predicting the performance of composite materials across scales requires significant computational resources. However, most of the commercial FEA programs used today does not scale to distributed systems for implicit analyses. Some leverage distributed computing for explicit analyses, but explicit analyses introduce other challenges, including smaller timesteps as the mesh refinement increases. For analyses of composites that include regions at the microscale, elements are often on the scale of hundreds of nanometers. If the aim is to predict the quasi-static response, then the number of timesteps becomes

prohibitive. Explicit methods are often adapted for distributed computing because they do not require solving a global system of equations, which is required for implicit methods. The FEA code developed for this work leverages third-party libraries for sparse linear algebra operations, including solving the system of linear equations, but distributed computing creates many other challenges than simply solving the equations. Supporting distributed computing affects:

- Creating large finite element models
- Processing input scripts
- Storing and reading the mesh (nodes, element connectivities, etc.)
- Storing global matrices and vectors
- Computing typical matrix-matrix, matrix-vector, and vector-vector operations between global matrices and vectors
- Assembling element matrices and vectors into the global system of equations
- Enforcing boundary conditions
- Solving the global system of equations
- Computing volume average and other model level quantities
- Writing the output data
- Visualizing the results

This list is not exhaustive, but it highlights most of the components of a finite element framework that requires modification due to distributed computing. Some processes are not directly affected but are instead affected by the increase in model size that comes with

scalability. For the types of analyses used in this work, FEA models typically contained less than a million degrees of freedom (DoFs) before distributed computing, but this work uses models with 50 to 100 million DoFs. Some analyses model nonlinear behavior, requiring many solves and easily result in over a 100 GB of data for a single nonlinear analysis. The larger models posed new challenges and required the development of several tools. This section aims to discuss the lessons learned, including successful and unsuccessful design patterns, from developing a scalable FEA framework.

2.3.1 History of the FEA Code

A FEA code, called Beta, had been developed over numerous years with the effort of various students and researchers within the research group of Dr. John Whitcomb at Texas A&M University. The FEA code for this work, Beta2, borrows many concepts and algorithms from its predecessor, but the entire code was rewritten to use modern C++ design, templates, exceptions, improve encapsulation, and support distributed computing. Before Beta2, an intermediate, proof of concept version of Beta was developed to leverage HPC. Due to the amount of code in Beta, some of the challenges required fundamental design changes that were not anticipated before the development of the intermediate version. After many fundamental issues persisted, the version was completely scrapped a year later. However, many lessons were learned from the development, and Beta2 was developed guided by a better understanding of what design patterns worked.

2.3.2 Mesh Partitioning

In a distributed analysis, it is important to identify which MPI rank is responsible for calculations and output of a specific node or element. In this work, a partition will refer to a fictitious subdomain of the analysis region, and the subdivision is purely for computational purposes. A single partition may or may not contain several constituents and can be arbitrary in shape, though it is computationally efficient to minimize size of the inter-partition boundaries. Continuous Galerkin fields are maintained across inter-partition boundaries, so partitioning an analysis region does not introduce any additional error. PETSc requires a continuous range of equations to be owned by an MPI rank. Since DoFs are typically nodal quantities in structural FEA, it logically follows that each node should be owned by exactly one MPI rank. Which MPI rank owns an element is less important, except for visualization or if elemental DoFs are introduced, as is common in computational fluid dynamics (CFD). None of the constitutive laws in this work require non-local calculations, so an MPI rank only requires the elements that include an owned node to assemble its respective rows of the global matrix. This fundamentally means that if three nodes within an element are owned by three different ranks then the element stiffness matrix is computed on all three ranks, since each rank would only assemble the terms for the DoFs it owns. This approach has the benefit of requiring no communication between ranks during assembly, but it also means that redundant computations are taking place.

However, it should be noted that there is another philosophy. Ownership could be on an element basis. This means that each element stiffness matrix is only assembled once,

but extra communication between ranks is required once assembly is complete. PETSc still requires each DoF to be owned by only rank, so several ranks may be computing terms that should appear in the row of equations owned by only one of the ranks. The ranks that computed terms that appear in a row they do not own must communicate the terms to rank that is responsible for the row in the system of equations. PETSc takes care of most of the communication, but it is unclear which method is more efficient. If it is cheap to compute the stiffness matrix for an element, then it may be more efficient for ownership to be on a nodal basis. If the constitutive model requires significant computation, then the communication cost of terms that are not owned by each rank to the respective owner is likely to become more efficient. Benchmarks were not performed in this work to suggest when one becomes more efficient.

For a distributed analysis, the decision was made that partitioning should be done before the analysis as part of the preprocessing. This makes load balancing more difficult, since the load on each rank must be estimated a priori, but it simplifies the code for the analysis, since the MPI rank that owns a node or element does not change throughout the analysis. New DoFs can be added during the analysis, but it does require reallocation of the global matrices and vectors. For a linear elastic analysis with 60 million DoFs, the time spent to solve the system of equations typically requires between 3 to 5 times longer than the time needed to allocate and assemble the system of equations. Though reallocation and assembly does have a cost, it is less significant than solving the system of equations. Consequently, strategies that improve the performance of the solver or several parts of the

process at once should be pursued first. Sometimes Metis/ParMetis was used to partition the mesh, but out of convenience, the mesh was often divided into rectangular, structured partitions. This caused some ranks to have about twice as many owned DoFs as other ranks, due to nonuniformity in the tetrahedral meshes. By more intelligently partitioning of the mesh could offer a significant speedup.

More intelligent static load balancing, which is done a priori, would be very useful, but it should be noted that dynamic load balancing, which occurs multiple times during the analysis, introduces a subtle issue. Dynamic load balancing can probably improve the performance of progressive damage analyses, since the calculation of stiffness matrices for elements exhibiting plasticity or damage may become more expensive and damage tends to localize. However, if visualization typically requires plotting quantities over the time of the analysis, then the mesh must be read and the data structures used to store the data must be reallocated every timestep. So dynamic load balancing may improve the time required for the analysis at the expense of visualization performance.

2.3.3 Sparse Storage and Solvers

Several different C++ libraries provide a suite of solvers and storage classes. The three most popular are Intel's Math Kernel Library (MKL), PETSc, and Trilinos. The main offering of MKL is an optimized implementation of Pardiso, a direct sparse solver. Direct solvers do not scale as well as iterative solvers, but they offer a robust method for solving symmetric, unsymmetric, and even indeterminant systems of equations. In this work,

MKL's Pardiso provided a reliable tool for determining appropriate tolerances for the iterative solvers, but it was not used for problems with more than 1 million DoFs.

Both PETSc and Trilinos offer a wide variety of solvers and preconditioners. Trilinos focuses on providing many components relevant to solving partial differential equations, including a lot more than just the storage and linear solvers. PETSc focuses more tightly on the storage and solvers. For this work, PETSc was chosen for its tighter scope. However, one issue to be aware of with PETSc is the changing interface. With each release, some function calls change, requiring some minor modifications to a code that uses PETSc with each upgrade. Visual Studio was used for the development of the framework, though the analyses were typically run on a Unix or Linux HPC system. It can be very difficult to compile PETSc on Windows with the dependencies and packages needed. Though success was found, researchers should be aware that the large sparse libraries can be difficult to compile on Windows. On the other hand, being able to use the Visual Studio debugger for small distributed analyses proved to be invaluable.

For linear elastic analyses, an additive Schwarz preconditioner was used with a conjugate gradient (CG) solver. The global stiffness matrix remains symmetric, which the CG solver leverages to offer good performance. For analyses with plasticity, the global stiffness matrix may not be symmetric, due to the non-associative flow rule discussed in the next chapter. Consequently, an additive Schwarz preconditioner was used with a bi-conjugate gradient stabilized (BiCGStab) solver. Other preconditioners, such as block factorization or

algebraic multigrid may offer improved performance, but they both require more parameters to tune and information about the problem.

Since several types of math libraries were used, it was important to maintain encapsulation of the sparse math routines from the rest of the FEA. To provide the most flexibility, a very thin interface was created, wrapping the sparse math libraries. Matrices and vectors had an integer ID, and calls to a sparse math routine from the FEA code required passing the ID's of the containers to use in the operation. This design pattern was found to be difficult to maintain. It was too easy for a developer to pass the ID for the wrong global container. To solve this issue, a class was developed to manage the global container ID's and required a strongly typed enumeration in the FEA code in place of an ID. For example, solving the class $Kq = F$ system of equations for linear elasticity looks like:

```
equations->Solve(glContainers.get(GlobalMatrixID::ElasticityStiffness),
                 glContainers.get(GlobalVectorID::RHS),
                 glContainers.get(GlobalVectorID::Displacement));
```

This approach made it very clear to a reader which containers were being used in a sparse matrix operation, but the class needed to translate the strongly typed enumerations into an integer is unnecessary. If the sparse math library was revamped, the vector and matrix enum types should be given as templates to the sparse math interface, and containers could be identified by the enumerations directly. Even with a verbose interface, this approach allowed several implementations of the sparse math libraries without a single change to the FEA code, and it opens the possibility of assembling and solving the global

systems using very different computational architectures, such as general-purpose computing on graphics cards (GPGPU) or coprocessor cards. Implementing either of these would require no changes outside the sparse math library implementation, though CPUs would still be computing the element matrices and vectors to assemble.

There are two types of calls that must be added in the FEA part of the code to make this design paradigm possible. First, the sparse library does not know when element-to-global assembly ends. It can detect when one begins (an assembly call is made since the end of a previous one), but a call is needed to indicate to the sparse math library that all the element matrices and vectors have been passed off for assembly. This is important since PETSc and other distributed libraries require post-assembly operations. The second type of call is updating the ghost values, which are the terms in the global vectors that are needed by an MPI rank but are not owned. Ghost values are often needed for the solution and incremental solution vectors to allow the elements to compute their stiffness matrix or stress state. Currently, the FEA code must let the sparse math library know to update the ghost values before it asks for parts of the solution vector that it may not own. However, with a more intelligent way of managing the global containers within the sparse math library, the FEA part of the code should not have to make this call. The sparse library could keep track of what ghost values it already has available, which containers have been modified, and communicate to other MPI ranks when a value is requested that it does not own or have an updated value available. It should be noted that the ghost values an MPI

rank could need is determined based on connectivity when the non-zeros are specified in the global matrix.

2.3.4 Boundary Conditions and Multipoint Constraints

In micromechanical analyses, multi-point constraints (MPCs), which are linear additional equations that relate a degree-of-freedom (DoF) to other DoFs, are very useful for imposing periodic boundary conditions. Many recent finite element codes use Lagrange multipliers or a constraint matrix to enforce multi-point constraints, which are both straightforward approaches. However, these approaches have a performance cost compared to directly modifying the global system of equations to account for the additional equations, since they make matrix-vector operations more expensive for the solver. This work employs a less popular methodology of enforcing MPCs. Additional equations are directly incorporated into the global system of equations during assembly. The method makes assembling element matrices and vectors into the global matrices and vectors more complicated but offers good performance and a single interface for all boundary conditions, including MPCs. A method was developed by Clint Chapman and improved by Ross McLendon for imposing multi-point constraints within a centralized finite element code. [37] [43] However, the problem becomes much more complex in a distributed computing framework, and significant development was needed to generalize the approach for distributed computing. The effort to generalize the method was successful, but it is unclear if the complexities and maintainability of the code outweigh the performance benefit. A description of the algorithms developed for a generalized incorporation of additional

equations within a distributed computing architecture is not provided within this manuscript.

2.3.5 Distributed Data Storage and Output

One of the issues encountered with distributed computing was the scale of data. Some plasticity/damage analyses required 200 GB of binary data for a single analysis. Generally, several different types of loads are considered, each requiring a separate analysis. If the effect of even one parameter is investigated, the amount of data can become quite large, and a method for efficiently storing, reading, and writing distributed data becomes a critical component in the framework. Within a distributed system, the data must be read in parallel to avoid every processor having to read the entire data set for a very small relevant subset of the data. Additionally, a method for storing several types of data was needed, including nodal, element, nodal by element (quantities extrapolated to the nodes within an element), and quadrature data. Several types of existing file formats were evaluated, but no file format was found that could both store all the types of data required and had a performant, open C++ library to interface with files using the format. Consequently, a new file format and I/O library was developed for the finite element data: the partitioned finite element compound (PFEC) file format.

The PFEC file format groups the data for each partition of the mesh and relies on a hierarchy of lookup tables to efficiently transverse the data. The data from all the partitions are stored within the same file. This decision was made because of typical file quotas on HPC systems. If a file was output for each partition at each time step, then around 10k

files would be typically be needed for a progressive damage analysis, but the HPC system used in this work limited users to 200k files total. During output, each process computes the space needed for the data that it is about to write, the root process allocates a blank file through the file system and writes the lookup tables, and then each process writes its data to the respective locations within the files. A meta file uses JSON to describe the collection of files created for a single analysis and what data exists within each. The output for a 60 million DoF analysis with flags to output all the relevant fields, except quadrature data, is typically between 30 and 40 GB per an output time step. Including the time to calculate derived quantities, the total time required to generated and write the output data requires just over two minutes with this approach. ASCII would have resulted in prohibitively large files for the precision needed and required much longer to read and write. HDF5 was considered as the backbone of the new file format, but it requires rectangular datasets and the C API is rather complex. With nodal by element and quadrature data, the number of values per an entry varies. So, a new, simple file format was developed instead. A capable file format and I/O interface is a necessity for a distributed FEA framework.

2.3.6 Custom Reader for ParaView

Visualizing datasets as large as 50 GB per a time step requires a sophisticated visualization tool. Especially for 3D textile models, it is very difficult to anticipate what types of plots will be useful to understand and communicate what is occurring within the material. Consequently, exploratory visualization was very important, and ParaView

provided everything that was needed. However, importing the data into ParaView was a challenge, since an the PFEC file format was used and ParaView does not natively support nodal by element data.

ParaView expects each data array to be either a nodal or elemental field. Nodal by element and quadrature data must be averaged at nodes before ParaView can display the data. However, the derived data, such as stress and strain, should not be averaged along constituent interfaces. A strategy tried early on was converting the data from PFEC to the binary VTU format. If each constituent in each partition had its own VTU file, then ParaView organized the data as a multiblock dataset and correctly displayed the appropriate discontinuities across constituent interfaces. However, this approach resulted in far too many files and was inconvenient due to the need for several file formats.

A much cleaner approach was to implement a custom reader for ParaView, but some unanticipated challenges arose. First, a competent knowledge of the VTK framework and ParaView pipeline is prerequisite for writing a plugin for ParaView. Initially, minimal documentation and examples existed, though the documentation has improved some since. Second, readers can be compiled in different ways depending on the operating system, but both methods require compiling the ParaView source. On Windows, the ParaView code base can be compiled into libraries, and as long as the CMake build is kept intact, the reader can be compiled into a DLL on other machines without needing to recompile the ParaView source. The DLL can be distributed to other Windows machines to use with the official ParaView binaries, but the DLL of the reader can only be loaded into the version

ParaView that matches the libraries that the reader was built against. If there is a version mismatch, ParaView will typically give a non-sensical error when loading the DLL. When a new version of ParaView is released, the ParaView source must be rebuilt before the reader can be updated for the new ParaView version.

An alternative method works for both Windows and Linux systems. The ParaView source can be built along with the reader source to create a binary of ParaView with the reader built in. However, it can be inconvenient to compile ParaView with all of the relevant packages, such as QT, MPI, and Python, each time a new version is released. This is unfortunately, the only viable strategy on Linux.

Since ParaView does not handle nodal by element or quadrature data, the reader had to compute the data at the nodes and perform the averaging internally. A multiblock data set was created within the reader with the structure shown in Figure 2.38. The strategy worked well within ParaView, and most filters that are prepackaged into ParaView worked with the multiblock data set when ghost cells were used. The averaging can be expensive for large models, so a postprocessing utility was developed to break up PFEC mesh and data files, so that one constituent exists per a partition and data is pre-averaged before being read into ParaView. The resulting files were quicker to read into ParaView, since most of the processing is already completed outside of the reader. In hindsight, it would be helpful to design the PFEC file format to mimic the structure required by ParaView. Averaging the data would still need to be done, but the change would have simplified reader significantly.

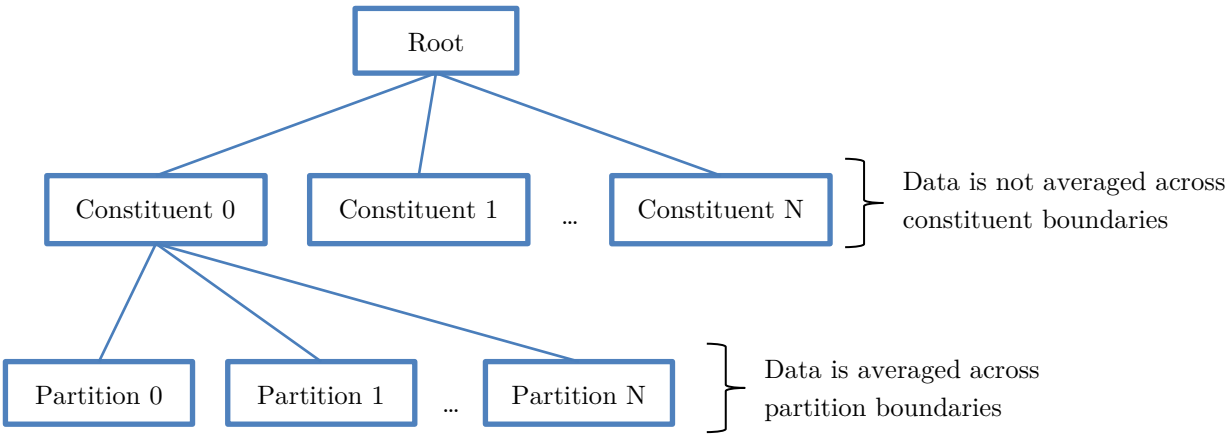

Figure 2.38. Structure of multiblock data set within ParaView

2.3.7 Python Script Framework for Postprocessing

To create the model, run the analyses, and post-process the data, an expressive high-level scripting interface was needed. This work used over 1000 analyses, which would be impractical to work with without automation. Additionally, many more figures than those used in this work needed to be created to provide insight into the locations of stress concentrations for a variety of models and configurations. By generating standard figures with an automated framework, the researchers focus could be on gaining the insight and creating figures to communicate the insights effectively.

A Python package, named pybeta, was developed to provide most of the preprocessing and postprocessing functionality. This Python package allowed many versatile libraries to

be used, such as NumPy, SciPy, and Matplotlib. The package was separated into five modules:

- mesh: Python wrappers for BetaMesh
- fiber_matrix: scripts for creating finite element models for fiber/matrix RVEs with general shapes, including a GUI for creating models
- analysis: scripts for managing ensembles of finite element analyses (very useful for ensembles of fiber/matrix analyses)
- utility: general utilities
- visualization: scripts for creating line plots, setting up ParaView visualizations, or creating complex animations in ParaView

3 THEORY

This chapter is divided in to three sections. The first section provides an overview of the classical finite element method, which is used in this work. Next, the theory for a plasticity model used to study the effect of accounting for plasticity within the neat matrix pocket is presented. Finally, the theory for using the finite element method predict the response of nonlinear materials is discussed.

3.1 Overview of Finite Element Method

This section provides a high-level derivation of the finite element method used in this work. First, the kinematics for the infinitesimal strain tensor are derived, followed by a brief description of Cauchy's law. Next, the conservation of linear and angular momentum are used to derive the governing equation and its weak form. An overview of linear elastic constitutive laws are provided. Finally, the classical finite element method is derived for linear elastic materials.

3.1.1 Kinematics

The Eulerian and Lagrangian descriptions are two ways of expressing the coordinates of a material. In the Eulerian description, properties are measured at fixed spatial coordinates, though the material may move. In the Lagrangian description, properties are measured at material points, which requires tracking how the material has moved from a reference configuration. Typically, an Eulerian description is adopted in the field of fluid mechanics, since tracking the movement of every fluid particle quickly becomes computational prohibitive. A Lagrangian description is usually used in the solid mechanics field, which

results in a mesh that deforms with the material. In the Lagrangian description, it is useful distinguish between the undeformed and deformed configurations, as illustrated in Figure 3.1. Let X_i be the i^{th} component of the coordinate for a material point before deformation occurs, and let x_i be the i^{th} component of the coordinate for a material point after experiencing deformation. There exists a function, which will be denoted by χ, that relations the undeformed configuration to the deformed configuration, as expressed in Equation (3.1).

$$x_i = \chi(X_i, t) \tag{3.1}$$

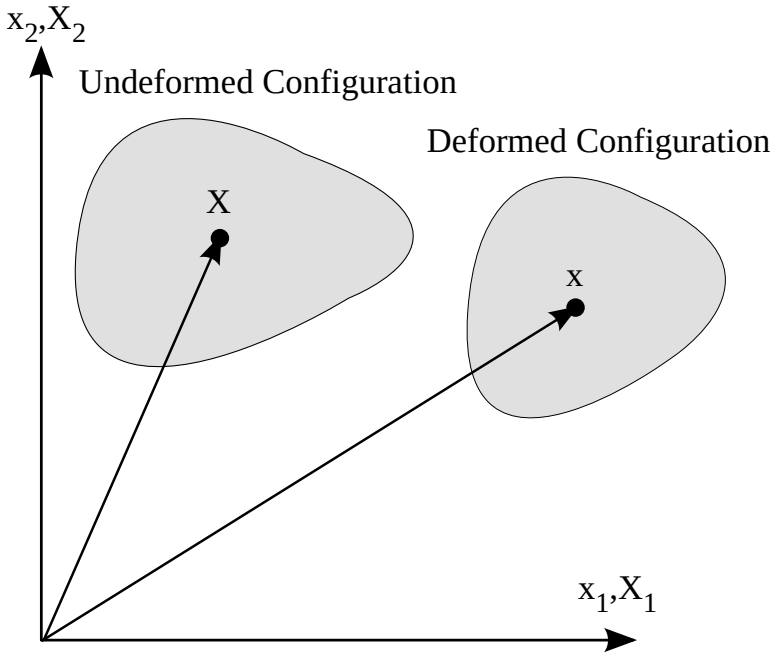

Figure 3.1. Illustration of the undeformed and deformed configurations

The displacement is defined as the difference between the deformed and undeformed configurations. In the Lagrangian description, Equation (3.2) expresses the displacement in terms of the undeformed coordinate and mapping function, χ.

$$u_i(X_i, t) = \chi(X_i, t) - X_i \tag{3.2}$$

Displacement has units of length, but it is more useful to develop a unitless measure of deformation. The deformation gradient as defined by Equation (3.3) offers a unitless measure of deformation. It noted that δ_{if} is the Kronecker delta. However, the tensor is generally nonsymmetric. More conveniently, the tensor can be made symmetric by multiplying with the transpose, forming the right Cauchy-Green deformation tensor, C_{ij}, which mathematically written in Equation (3.4). Finally, the Green-Lagrange strain tensor is defined in Equation (3.5).

$$F_{ij} = \frac{\partial x_i}{\partial X_j} = \frac{\partial u_i}{\partial X_j} + \delta_{ij} \tag{3.3}$$

$$C_{ij} = F_{ki}F_{kj} = \frac{\partial u_i}{\partial X_j} + \frac{\partial u_j}{\partial X_i} + \frac{\partial u_i}{\partial X_j}\frac{\partial u_j}{\partial X_i} + \delta_{ij} \tag{3.4}$$

$$E_{ij} = \frac{1}{2}\left(C_{ij} - \delta_{ij}\right) = \frac{1}{2}\left(\frac{\partial u_i}{\partial X_j} + \frac{\partial u_j}{\partial X_i} + \frac{\partial u_i}{\partial X_j}\frac{\partial u_j}{\partial X_i}\right) \tag{3.5}$$

Removing the second order term yields the linearized Green-Lagrange strain tensor, and assuming that $\frac{\partial u_i}{\partial X_j} \approx \frac{\partial u_i}{\partial x_j}$ finally results in the linearized infinitesimal strain tensor, $\boldsymbol{\varepsilon}$, as shown in Equation (3.6). With these assumptions, the infinitesimal strain tensor becomes equivalent in the Eulerian and Lagrangian descriptions, and at this point, x_i will be used as the coordinate system.

$$\varepsilon_{ij} = \frac{1}{2}\left(\frac{\partial u_i}{\partial X_j} + \frac{\partial u_j}{\partial X_i}\right) \tag{3.6}$$

3.1.2 Cauchy's Law

Traction, also known as the stress vector, is defined on a plane by the force, $\boldsymbol{f}$, divided by area, A. Cauchy's law relates the traction to a quantity known as the Cauchy stress tensor, $\boldsymbol{\sigma}$, as shown in Equation (3.7), where $\boldsymbol{n}$ is the unit vector normal to the plane. The Cauchy stress tensor is defined in the deformed configuration, and it is a measure of the force in the deformed configuration divided by the deformed area at a point. However, when infinitesimal strain is valid, the force and area remain approximately unchanged when switching between the deformed and undeformed configurations. This work only considers infinitesimal strains, so the Cauchy stress tensor is equivalent to other stress measures and will be referred to as simply the stress tensor.

$$t_i = \sigma_{ij} n_j \tag{3.7}$$

3.1.3 Conservation Equations

The balance of linear momentum requires that the rate of change of linear momentum equals the forces acting on a material of volume V. The balance of linear moment is given in Equation (3.8), where $\boldsymbol{t}$ is the traction, $\boldsymbol{b}$ is the body force, $\boldsymbol{\rho}$ is the density, and $\boldsymbol{v}$ is the velocity. Using Cauchy's law, shown in Equation (3.7), and the divergence theorem, the equation can be expressed as shown in Equation (3.9). Finally, using Reynold's transport theorem and the fact that the equation must hold for any arbitrary volume, the balance of linear momentum can be written in the useful form shown in Equation (3.10).

$$\int t_i dA + \int \rho b_i dV = \frac{d}{dt} \int \rho v_i dV \tag{3.8}$$

$$\int \left(\frac{\partial \sigma_{ji}}{\partial x_j} + \rho b_i \right) dV = \frac{d}{dt} \int \rho v_i dV \tag{3.9}$$

$$\frac{\partial \sigma_{ji}}{\partial x_j} + \rho b_i = \rho a_i \tag{3.10}$$

Similarly, the balance of angular momentum requires that the rate of change of angular momentum equals the moments acting on the material. This can written in tensor notation as shown in Equation (3.11), but it is more convenient to express the equation in indicial notation as shown in Equation (3.12). Using Cauchy's law, the product rule, and the

Reynold's transport theorem, the equation becomes Equation (3.13). Rearranging the equation as shown in Equation (3.14), it becomes clear that the right hand side is zero by the balance of linear of momentum, refer to Equation (3.10).

$$\oint \mathbf{x} \times \mathbf{t} dS + \int \mathbf{x} \times \rho \mathbf{b} dV = \frac{d}{dt} \int \mathbf{x} \times \rho \mathbf{v} dV \tag{3.11}$$

$$\oint \epsilon_{ijk} x_j t_k dS + \int \rho \epsilon_{ijk} x_j b_k dV = \frac{d}{dt} \int \rho \epsilon_{ijk} x_j v_k dV \tag{3.12}$$

$$\int \left(\epsilon_{ijk} \left(\delta_{jm} \sigma_{mk} + x_j \frac{\partial \sigma_{mk}}{\partial x_m} \right) + \rho \epsilon_{ijk} x_j b_k \right) dV = \int \rho \epsilon_{ijk} x_j a_k dV \tag{3.13}$$

$$\int \epsilon_{ijk} \delta_{jm} \sigma_{mk} dV = \int \epsilon_{ijk} x_j \left(\rho a_k - \frac{\partial \sigma_{mk}}{\partial x_m} - \rho b_k \right) dV \tag{3.14}$$

Using the conservation of linear momentum, the conservation of angular momentum is shown in Equation (3.11). Since the equation must hold for any volume, the integrand must be equal to 0. Using this and the identity $\varepsilon_{imn}\varepsilon_{ijk} = \delta_{mj}\delta_{nk} - \delta_{nj}\delta_{mk}$, it can be shown that the stress tensor is symmetric, as shown in Equation (3.12).

$$\int \epsilon_{ijk} \delta_{jm} \sigma_{mk} dV = \int \epsilon_{ijk} \sigma_{jk} dV = 0 \tag{3.15}$$

$$\epsilon_{ijk} \sigma_{jk} = \sigma_{jk} - \sigma_{kj} = 0 \tag{3.16}$$

3.1.4 Weak Form Derivation

The governing equation shown in Equation (3.10) was derived from the conservation of linear momentum, but a different form is needed to numerically solve the system of equations. This work focuses on the static equilibrium solution, so the governing equation becomes Equation (3.17). Using the principle of virtual work, the equation is multiplied by a variation of the displacement and integrated over the volume, as shown in Equation (3.18).

$$\frac{\partial \sigma_{ji}}{\partial x_i} + \rho b_i = 0 \tag{3.17}$$

$$\int \delta u_i \left(\frac{\partial \sigma_{ji}}{\partial x_i} + \rho b_i \right) dV = 0 \tag{3.18}$$

To avoid including a derivative of the stress tensor, the first term is first integrated by parts to yield Equation (3.19). Next, the divergence theorem is used to manipulate the first term, which results in Equation (3.20).

$$\int \left(\frac{\partial \delta u_i \sigma_{ji}}{\partial x_i} - \sigma_{ji} \delta \frac{\partial u_i}{\partial x_i} + \delta u_i \rho b_i \right) dV = 0 \tag{3.19}$$

$$\int u_i \sigma_{ji} n_j dA + \int \left(\delta u_i \rho b_i - \sigma_{ji} \delta \frac{\partial u_i}{\partial x_i} \right) dV = 0 \tag{3.20}$$

Finally, the weak form given in can be obtained by leveraging the definition of the infinitesimal strain tensor and the symmetry of the stress tensor. The weak form of the governing equation is shown in Equation (3.21).

$$\int t_i \delta u_i dA + \int \left(\rho b_i \delta u_i - \sigma_{ij} \delta \varepsilon_{ij}\right) dV = 0 \tag{3.21}$$

3.1.5 Constitutive Laws

Hooke proposed that in a continuous media stress is linearly related to strain. The relationship can be described by a 4th order tensor, C_{ijkl}, known as the stiffness matrix, as shown in Equation (3.22). I should be noted and emphasized here that the stress is linearly related to the elastic part of the strain tensor. If no eigenstrains exist, such as plastic strain, thermal expansion, moisture expansion, etc., then the strain tensor is equivalent to the elastic strain. For the this section, the strain tensor will be in place of elastic strain, but in the last section of this chapter, this restriction will be removed. The constitutive relationship can be inverted, and the matrix relating strain to stress, S_{ijkl}, is known as the compliance matrix, which is shown in Equation (3.23).

$$\sigma_{ij} = C_{ijkl} \varepsilon_{kl} \tag{3.22}$$

$$\varepsilon_{ij} = S_{ijkl} \sigma_{kl} \tag{3.23}$$

The strain and stress tensors have already been shown to be symmetric. Consequently, the constitutive matrix exhibits some symmetries. It follows that $C_{jikl} = C_{ijkl}$ and $C_{ijlk} = C_{ijkl}$, due to the symmetry of the stress and strain tensors. These relations are often referred to as minor symmetries, and they result in a reduction from 81 unique constants to only 36 unique constants. Finally, the stress can be expressed as a partial derivative of the strain energy density, U, with respect to the strain, as shown in Equation (3.24). This leads to a symmetry of the stiffness matrix as shown in Equation (3.25), known as major symmetry, further reducing the number of unique constants from 36 to 21.

$$\sigma_{ij} = \frac{\partial U}{\partial \varepsilon_{ij}} \tag{3.24}$$

$$C_{ijkl} = \frac{\partial^2 U}{\partial \varepsilon_{ij} \partial \varepsilon_{kl}} = C_{klij} \tag{3.25}$$

Because of the symmetries, Hooke's law can be written in a contracted notation, as shown in Equation (3.26). In contracted notation, the engineering strain is used by convention, which means the shear terms, ε_4, ε_5, and ε_6 are twice the tensorial shear strains. In indicial notation, the equation simply becomes $\sigma_i = C_{ij}\varepsilon_j$. Note that the stiffness matrix is still symmetric even in the contracted notation, $C_{ij} = C_{ji}$. The order of stresses and strains in the contracted notation is arbitrary, but the order used must be consistent. The finite element code used for this work uses the order of (11, 22, 33, 12, 23, 13), which is not as common as Voigt's notation of (11, 22, 33, 23, 13, 12). The relations between the

contracted and tensorial stresses and strains are shown in Equations (3.27) and (3.28), respectively.

$$\begin{bmatrix}\sigma_1\\ \sigma_2\\ \sigma_3\\ \sigma_4\\ \sigma_5\\ \sigma_6\end{bmatrix} = \begin{bmatrix}C_{11} & C_{12} & C_{13} & C_{14} & C_{15} & C_{16}\\ C_{12} & C_{22} & C_{23} & C_{24} & c_{25} & C_{26}\\ C_{13} & C_{23} & C_{33} & C_{34} & C_{35} & C_{36}\\ C_{14} & C_{24} & C_{34} & C_{44} & C_{45} & C_{46}\\ C_{15} & C_{25} & C_{35} & C_{45} & C_{55} & C_{56}\\ C_{16} & C_{26} & C_{36} & C_{46} & C_{56} & C_{66}\end{bmatrix} \cdot \begin{bmatrix}\varepsilon_1\\ \varepsilon_2\\ \varepsilon_3\\ \varepsilon_4\\ \varepsilon_5\\ \varepsilon_6\end{bmatrix} \tag{3.26}$$

$$\begin{bmatrix}\sigma_1\\ \sigma_2\\ \sigma_3\\ \sigma_4\\ \sigma_5\\ \sigma_6\end{bmatrix} = \begin{bmatrix}\sigma_{11}\\ \sigma_{22}\\ \sigma_{33}\\ \sigma_{12}\\ \sigma_{23}\\ \sigma_{13}\end{bmatrix} \tag{3.27}$$

$$\begin{bmatrix}\varepsilon_1\\ \varepsilon_2\\ \varepsilon_3\\ \varepsilon_4\\ \varepsilon_5\\ \varepsilon_6\end{bmatrix} = \begin{bmatrix}\varepsilon_{11}\\ \varepsilon_{22}\\ \varepsilon_{33}\\ 2\varepsilon_{12}\\ 2\varepsilon_{23}\\ 2\varepsilon_{13}\end{bmatrix} \tag{3.28}$$

Many materials exhibit a mirror or rotational symmetry, which reduce the number of unique constants needed to characterize the constitutive relationship. An isotropic exhibits mirror symmetry about every plane and rotational symmetry about every axis, reducing the number of unique parameters to two. Though there are only two parameters needed to express the constitutive relationship of an isotropic material, the compliance matrix is often written in terms of three engineering constants that have a physical interpretation. The Young's modulus relates the extensional deformation to the normal stress and is denoted by E. The shear modulus relates the shear deformation to the shear stresses and is denoted

by G. Finally, the Poisson's ratio relates the transverse and normal deformation and is denoted by ν. For anisotropic media, more constants are required and the subscripts for E, G, and ν indicate the direction and plane in which the respective property is measured.

An orthotropic material has no rotational symmetries but has three orthogonal planes of symmetry, namely the x-y, y-z, and x-z planes. The compliance matrix of an orthotropic material can be expressed in terms of the engineering constants as shown in Equation (3.29).

$$\begin{bmatrix} \varepsilon_{11} \\ \varepsilon_{22} \\ \varepsilon_{33} \\ 2\varepsilon_{12} \\ 2\varepsilon_{23} \\ 2\varepsilon_{13} \end{bmatrix} = \begin{bmatrix} \frac{1}{E_{11}} & -\frac{\nu_{21}}{E_{22}} & -\frac{\nu_{31}}{E_{33}} & 0 & 0 & 0 \\ -\frac{\nu_{12}}{E_{11}} & \frac{1}{E_{22}} & -\frac{\nu_{32}}{E_{33}} & 0 & 0 & 0 \\ -\frac{\nu_{13}}{E_{11}} & -\frac{\nu_{23}}{E_{22}} & \frac{1}{E_{33}} & 0 & 0 & 0 \\ 0 & 0 & 0 & \frac{1}{G_{12}} & 0 & 0 \\ 0 & 0 & 0 & 0 & \frac{1}{G_{23}} & 0 \\ 0 & 0 & 0 & 0 & 0 & \frac{1}{G_{13}} \end{bmatrix} \cdot \begin{bmatrix} \sigma_{11} \\ \sigma_{22} \\ \sigma_{33} \\ \sigma_{12} \\ \sigma_{23} \\ \sigma_{13} \end{bmatrix} \tag{3.29}$$

A transversely isotropic material has the symmetries exhibited by an orthotropic material plus a rotational symmetry about the x-axis. This type of material is of particular interest of this work since the carbon fibers are will be assumed to be transversely isotropic. The direction along the fiber axis is referred to as the longitudinal direction, and the directions perpendicular to the fiber axis are referred to as the transverse direction. Conveniently, the compliance matrix of a transversely isotropic material can be expressed

in terms of the engineering constants as shown in Equation (3.30), since $\nu_{23} = \nu_{32}$, $E_{22} = E_{33}$, and $G_{12} = G_{13}$. However, it should be noted that only five unique constants are needed to characterize the material.

$$\begin{bmatrix} \varepsilon_{11} \\ \varepsilon_{22} \\ \varepsilon_{33} \\ 2\varepsilon_{12} \\ 2\varepsilon_{23} \\ 2\varepsilon_{13} \end{bmatrix} = \begin{bmatrix} \frac{1}{E_{11}} & -\frac{\nu_{21}}{E_{22}} & -\frac{\nu_{21}}{E_{22}} & 0 & 0 & 0 \\ -\frac{\nu_{12}}{E_{11}} & \frac{1}{E_{22}} & -\frac{\nu_{23}}{E_{22}} & 0 & 0 & 0 \\ -\frac{\nu_{12}}{E_{11}} & -\frac{\nu_{23}}{E_{22}} & \frac{1}{E_{22}} & 0 & 0 & 0 \\ 0 & 0 & 0 & \frac{1}{G_{12}} & 0 & 0 \\ 0 & 0 & 0 & 0 & \frac{2(1+\nu_{23})}{E_{22}} & 0 \\ 0 & 0 & 0 & 0 & 0 & \frac{1}{G_{12}} \end{bmatrix} \cdot \begin{bmatrix} \sigma_{11} \\ \sigma_{22} \\ \sigma_{33} \\ \sigma_{12} \\ \sigma_{23} \\ \sigma_{13} \end{bmatrix} \tag{3.30}$$

3.1.6 Finite Element Method

The finite element method (FEM) was been developed by a series of many researchers to numerically solve partial differential equations. The method has been extensively used to solve the variational form of the governing equation shown in Equation (3.21). FEM relies on discretizing the domain into a collection of smaller domains, known as finite elements, which consist of a connected set of nodes. The displacements within a finite element are approximated using a set of interpolation functions, ψ. Within conventional FEM, one interpolation function exists per a node within a finite element. Each function has a value of 1 at the location of the respective node within the element and has a value

of 0 at the other nodal locations. Mathematically, this can be written as shown in Equation (3.31), where ψ^i is the interpolation function associated with node i and $\mathbf{x}^j$ is the coordinate of node j. Furthermore, the value of all of the interpolation functions sum to 1 at every value within the element, as expressed in Equation (3.32), where n is the number of nodes in the element. This property is known as the partition of unity. It should be noted that variations of the finite element method has been developed, and some of these methods rely on interpolation functions that do not satisfy one or both of the properties shown in Equations (3.31) and (3.32). These methods are known as extended finite element methods (X-FEM). The work herein uses a conventional FEA formulation.

$$\psi^i(\mathbf{x}^j) = \delta_{ij} \tag{3.31}$$

$$\sum_i^n \psi^i(x) = 1 \tag{3.32}$$

The interpolation functions are used to interpolate the value of quantifies that are known at the nodes to any location within the element. The displacement at a position within the element, **x,** is approximated by Equation (3.33), where subscripts indicate the component of the displacement vector and superscripts indicate the corresponding node. Similarly, the variation of the displacement, δu, at a position within the element is approximated by Equation (3.34).

$$u_i(\mathbf{x}) \approx \sum_k^n \psi^k(\mathbf{x}) u_i^k \tag{3.33}$$

$$\delta u_i(\mathbf{x}) \approx \sum_k^n \psi^k(\mathbf{x}) \delta u_i^k \tag{3.34}$$

For conciseness, it is convenient to organize the nodal displacements within the element as a vector, $\boldsymbol{q}$. Terms are organized first by component then by node. Equation (3.35) shows the terms in vector $\boldsymbol{q}$ for a 3D problem, which requires three components of displacement. This allows the variation of the displacement to be written in terms of the vector $\boldsymbol{q}$ more compactly as shown in Equation (3.36). Similarly, Equation (3.37) shows the variation of the strain in terms of the nodal displacement vector, $\boldsymbol{q}$, using contracted notation.

$$\boldsymbol{q} = \begin{bmatrix} u_0^0 \\ u_1^0 \\ u_2^0 \\ \vdots \\ u_0^n \\ u_1^n \\ u_2^n \end{bmatrix} \tag{3.35}$$

$$\delta u_i(\mathbf{x}) = \frac{\partial u_i(\mathbf{x})}{\partial q_k} \delta q_k \tag{3.36}$$

$$\delta \varepsilon_{\mathrm{i}}(\mathbf{x}) = \frac{\partial \varepsilon_i(\mathbf{x})}{\partial q_k} \delta q_k \tag{3.37}$$

It is convenient to express the derivative of the strain tensor with respect to the nodal displacements as a matrix, B_{ij}. Equation (3.38) shows the terms within the matrix. With the formation of the B matrix, the strain and the variation of the strain can conveniently be written as shown in Equations (3.39) and (3.40), respectively.

$$B_{ij}(\mathbf{x}) = \frac{\partial \varepsilon_i(\mathbf{x})}{\partial q_k} = \begin{bmatrix} \frac{\partial \psi^1(\mathbf{x})}{\partial x_1} & 0 & 0 & \dots & \frac{\partial \psi^n(\mathbf{x})}{\partial x_1} & 0 & 0 \\ 0 & \frac{\partial \psi^1(\mathbf{x})}{\partial x_2} & 0 & \dots & 0 & \frac{\partial \psi^n(\mathbf{x})}{\partial x_2} & 0 \\ 0 & 0 & \frac{\partial \psi^1(\mathbf{x})}{\partial x_3} & \dots & 0 & 0 & \frac{\partial \psi^n(\mathbf{x})}{\partial x_3} \\ \frac{\partial \psi^1(\mathbf{x})}{\partial x_1} & \frac{\partial \psi^1(\mathbf{x})}{\partial x_2} & 0 & \dots & \frac{\partial \psi^n(\mathbf{x})}{\partial x_1} & \frac{\partial \psi^n(\mathbf{x})}{\partial x_2} & 0 \\ 0 & \frac{\partial \psi^1(\mathbf{x})}{\partial x_2} & \frac{\partial \psi^1(\mathbf{x})}{\partial x_3} & \dots & 0 & \frac{\partial \psi^n(\mathbf{x})}{\partial x_2} & \frac{\partial \psi^n(\mathbf{x})}{\partial x_3} \\ \frac{\partial \psi^1(\mathbf{x})}{\partial x_1} & 0 & \frac{\partial \psi^1(\mathbf{x})}{\partial x_3} & \dots & \frac{\partial \psi^n(\mathbf{x})}{\partial x_1} & 0 & \frac{\partial \psi^n(\mathbf{x})}{\partial x_3} \end{bmatrix} \tag{3.38}$$

$$\varepsilon_{\mathrm{i}}(\mathbf{x}) = B_{ik}(\mathbf{x}) q_k \tag{3.39}$$

$$\delta\varepsilon_{\mathrm{i}}(\mathbf{x}) = B_{ik}(\mathbf{x}) \delta q_k \tag{3.40}$$

Using these expressions for strains and the constitutive law, the weak form of the governing equation can be written as shown in Equation (3.41). The variation of the nodal displacements, δq_k, can be factored out of the integrals, since it is the nodal displacements and not a function of position. Since δq_k is arbitrary, the integrals must equal zero. Equation (3.42) shows the equation after terms have been rearranged and δq_k has been factored out of the integrals.

$$\int t_i \frac{\partial u_i(\mathbf{x})}{\partial q_k} \delta q_k dA + \int \left(\rho b_i \frac{\partial u_i(\mathbf{x})}{\partial q_k} \delta q_k - C_{ij} B_{jm}(\mathbf{x}) q_m B_{ik}(\mathbf{x}) \delta q_k \right) dV = 0 \tag{3.41}$$

$$\int B_{ik}(\mathbf{x}) C_{ij} B_{jm}(\mathbf{x}) dV \, q_m = \int t_i \frac{\partial u_i(\mathbf{x})}{\partial q_k} dA + \int \rho b_i \frac{\partial u_i(\mathbf{x})}{\partial q_k} dV \tag{3.42}$$

Recall that this equation must hold for each finite element, so the integrals are over the domain of the finite element. The left-hand side of the equation forms a matrix, which is conventionally referred to as the stiffness matrix. The right-hand side of the equation forms a vector, known as the force vector. The stiffness matrix and force vector are mathematically written in Equations (3.43) and (3.44), respectively.

$$K_{ij} = \int B_{mi}(\mathbf{x}) C_{mn} B_{nj}(\mathbf{x}) dV \tag{3.43}$$

$$F_i = \int t_m \frac{\partial u_m(\mathbf{x})}{\partial q_i} dA + \int \rho b_m \frac{\partial u_m(\mathbf{x})}{\partial q_i} dV \tag{3.44}$$

This arrangement of terms results in a classical system of linear equations. Finally, Equation (3.45) expresses the system of equations for solving for the mechanical response of a solid material.

$$K_{ij} q_j = F_i \tag{3.45}$$

The displacements across element boundaries are held as continuous. Consequently, a global system of equations can be formed by adding all the element contributions. By solving the global system of equations, the displacements at every node in the mesh determined. With the displacements, the strain and stress within the element can be calculated using Equation (3.39) and the constitutive law.

3.2 Plasticity Model

Plastic deformation generally refers to permanent deformation within a material. Permanent deformation will cause a material to not recover its original shape after an applied load is removed. Some materials have a greater propensity to experience plastic deformation than others. Within the context of this work, carbon fibers and the epoxy matrix behave very differently outside the linear regime. Under longitudinal tension, carbon fibers behave linearly just before the point of fracture. On the other hand, an epoxy matrix exhibits significant permanent deformation before fracture occurs.

Figure 3.2 shows an illustration of a nonlinear stress-strain curve. Until point A, the stress-strain relationship is linear. The stress where the elastic regime ends and plastic deformation begins is referred to as the yield stress, σ_y. From A to B, plastic deformation is accumulated as yield occurs, but unloading still occurs elastically, as illustrated from point B. Additionally, the strain can be separated into two components: the plastic strain, ε_p, and the elastic strain ε_e. From point A to B, the change in stress decreases as the strain increases, which is referred to as strain hardening. From point B to C, the stress decreases

as the strain increases, which is an instability and referred to as strain softening. From point C to D, strain hardening continues until fracture, denoted by "x" at point D.

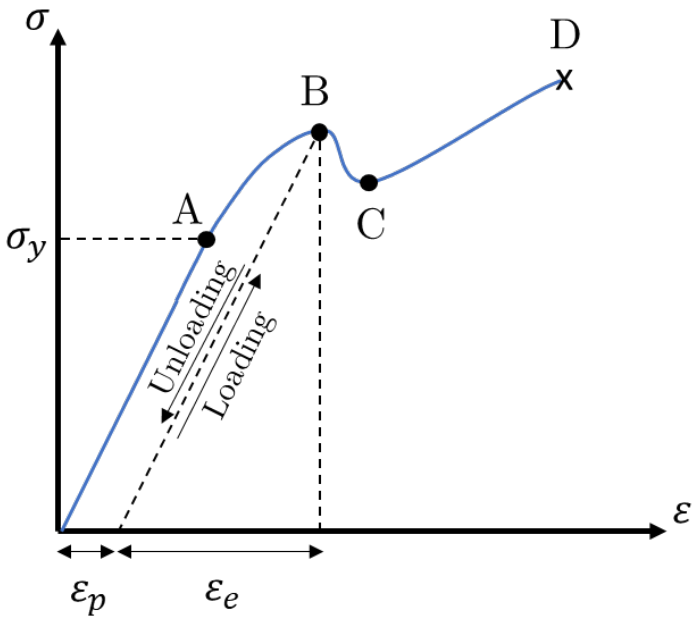

Figure 3.2. General stress-strain curve exhibiting nonlinearities

A plasticity model is meant to capture the nonlinear behavior before fracture occurs. Generally, the evolution of plastic deformation in epoxy resins depends on the rate of load and temperature. This work is limited to quasi-static loads at a single temperature, so a rate-independent plasticity model will be used. Additionally, many types of polymers, including thermoset epoxy resins, exhibit strain softening under a variety of loads, but the phenomenon is most pronounced under compressive loads. [44] [45] However, capturing the strain softening introduces several additional complexities, so this work will assume that only strain hardening occurs.

A rate-independent plasticity model can be divided into three components:

1. Yield criterion, which indicates when yielding should occur
2. Hardening laws, which define how the parameters of the yield criterion evolve as plastic deformation accumulates
3. Flow rule, which defines how plastic strain evolves during yielding

Each of these components must be defined for a plasticity model. The functions used for each component can vary significantly depending on the material. This work will consider a model developed by Melro et. al.. [46] [47] It is a relatively simple rate-independent, pressure-dependent plasticity model developed specifically for epoxy resin systems.

3.2.1 Yield Criterion

Many different yield criteria have been proposed by researchers over the years. One of the most famous being the Von Mises criterion, as shown in Equation (3.46), where $\boldsymbol{f}$ is the yield function and J_2 is the second invariant of the deviatoric stress tensor. [48] Equation (3.47) shows J_2 in terms of the components of the stress tensor. Yielding occurs when the yield function reaches 0., which occurs when J_2 reaches one third of the square of the yield stress. It makes no distinction between tensile and compressive strengths and does not include a pressure dependent term. However, the yield criterion has been widely used successfully for ductile metals.

$$f(\boldsymbol{\sigma}, \sigma_y) = 3J_2 - \sigma_y^2 = 0 \tag{3.46}$$

$$J_2 = \frac{1}{6}\left((\sigma_{11} - \sigma_{22})^2 + (\sigma_{22} - \sigma_{33})^2 + (\sigma_{33} - \sigma_{11})^2 + 6(\sigma_{12}^2 + \sigma_{23}^2 + \sigma_{13}^2)\right) \quad (3.47)$$

Due to the pressure sensitivity exhibited by epoxy resins, Melro et. al. proposed a paraboloidal yield criterion that included a pressure dependent term, based on the paraboloidal yield criterion proposed by Tschoegl. [46] [49] The yield criterion is written in Equation (3.46), where I_1 is the first invariant of the stress tensor, σ_t is the tensile yield strength, and σ_c is the compressive yield strength. The first invariant of the stress tensor is simply given by $I_1 = \sigma_{11} + \sigma_{22} + \sigma_{33}$. It should be noted that yielding will occur when $f = 0$, but the yield strengths can evolve as plastic strain accumulates.

$$f(\boldsymbol{\sigma}, \sigma_c, \sigma_t) = 6J_2 + 2I_i(\sigma_c - \sigma_t) - 2\sigma_c\sigma_t = 0 \quad (3.48)$$

Figure 3.3 shows a typical yield surface using the criterion given by Equation (3.46). The surface is plotted in principle stress space, and the hydrostatic axis lies along the line $\sigma_1 = \sigma_2 = \sigma_3$. It is important to note that the surface is symmetric about the hydrostatic axis and smooth, which are two properties that will simplify the computational implementation.

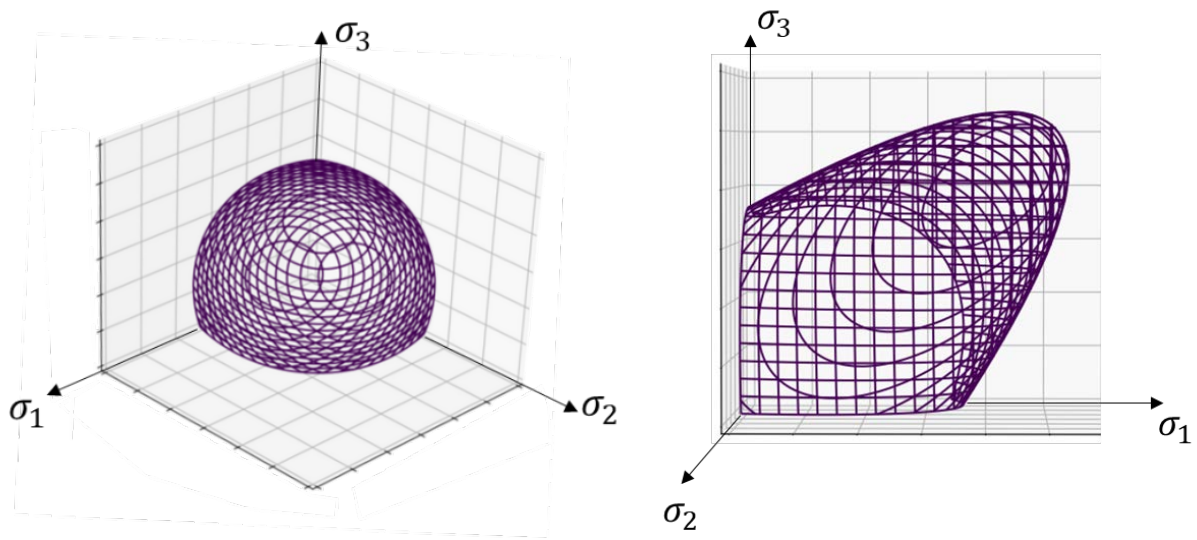

Figure 3.3. Two views of a typical paraboloidal yield surface given by Equation (3.46)

3.2.2 Hardening Laws

The yield criterion indicates when yielding will occur, but the parameters in the yield function are often not constant. As plastic strain accumulates, the stress required to cause further plastic deformation in the material can change, refer to the region between points A and B in Figure 3.2. There are two common categories of hardening: isotropic and kinematic. Considering surface in stress space defined by $f = 0$, isotropic hardening grows the surface, while kinematic hardening translates the surface. If the material experiences proportional loading, then isotropic and kinematic hardening give the same stress-strain response. However, when the proportions of stress components change, then the two types of hardening behave differently. In reality, a material exhibits both types of hardening. However, the analyses in this work focus on proportional load paths for the models, and though the local material may experience non-proportional loading, the degree of non-

proportionality should be relatively small. Consequently, only isotropic hardening will be considered in this work.

Ideally, careful experimental tests should inform the selection of hardening parameters. In general, the hardening law can simply be a spline fit of discrete experimental data. However, experimental data was not readily available for the epoxy resin systems considered in this work. Consequently, a simple power law will be used for isotropic hardening of the tensile and compressive strengths as a function of the effective plastic strain, ε_e^p. The effective plastic strain is a measure of the accumulated plastic strain. The effective plastic strain cannot decrease throughout any loading history, and it is defined incrementally, taking the history of plastic strain into account. Consider the case where a material is yielded in tension. If the loading is reversed and the material is yielded in compression, then the plastic strain tensor may be reduced, but the effective plastic strain continues to increase. The increment of effective plastic strain is defined in terms of the increment of plastic strain as shown in Equation (3.49), where k is the constant that results in $\varepsilon_e^p = \varepsilon_{11}^p$ under uniaxial load along the x-axis. Equation (3.50) shows the value of k for the yield function shown in Equation (3.46), where ν_p is the plastic Poisson ratio.

$$\Delta\varepsilon_e^p = \sqrt{k\Delta\varepsilon_{ij}^p \Delta\varepsilon_{ij}^p} \tag{3.49}$$

$$k = \frac{1}{1 + 2\nu_p^2} \tag{3.50}$$

The hardening laws used in this work take the form that is shown in Equations (3.51) and (3.52), where σ_{c0} and σ_{t0} are the compressive and tensile yield strengths before any yielding occurs, H_c and H_t are the compressive and tensile hardening moduli, and n_c and n_t are the compressive and tensile hardening exponents.

$$\sigma_c = \sigma_{c0} + H_c\left(\varepsilon_e^p\right)^{\eta_c} \tag{3.51}$$

$$\sigma_t = \sigma_{t0} + H_t\left(\varepsilon_e^p\right)^{\eta_t} \tag{3.52}$$

3.2.3 Flow Rule

The flow rule determines how the plastic strain tensor will evolve during yielding. It is useful to propose that the instantaneous rate of change in plastic strain is equal to the rate of change of a plastic multiplier, $\dot{\gamma}$, times a flow tensor, $\boldsymbol{N}$. The flow tensor indicates the proportionality of how each component of the plastic strain will change, while the rate of change of the plastic multiplier indicates how fast the plastic strain will change.

$$\dot{\boldsymbol{\varepsilon}}^p = \dot{\gamma}\boldsymbol{N} \tag{3.53}$$

For associate flow rules, the flow tensor is simply the derivative of the yield function with respect to the stress tensor, $\boldsymbol{\sigma}$. However, this would result in an unrealistic volumetric change during yielding when the yield function contains a pressure-dependent term. [46] Consequently, the flow rule will be based on a second potential, $\boldsymbol{g}$, called the flow potential. When the flow potential does not equal the yield function, the flow rule is called non-

associative. Equation (3.54) shows the flow potential, g, proposed by Melro et. al., where α is a material constant.

$$g(\boldsymbol{\sigma}) = 3J_2 + \frac{\alpha}{9} I_1^2 \tag{3.54}$$

The material parameter α is not independent and can be related to ν_p. It should the value that results in the proper volumetric deformation during yielding due to a uniaxial tensile load. The derivation can be found in Ref. [46], but the result in given in Equation (3.55).

$$\alpha = \frac{9}{2} \frac{1 - 2\nu_p}{1 + \nu_p} \tag{3.55}$$

Using the flow potential in Equation (3.54), the incremental form of the flow rule can be written as shown in Equation (3.56), where $\boldsymbol{S}$ is the deviatoric stress tensor and $\boldsymbol{I}$ is the identity tensor.

$$\Delta \boldsymbol{\varepsilon}^p = \Delta\gamma \frac{\partial g}{\partial \boldsymbol{\sigma}} = \Delta\gamma \left(3\boldsymbol{S} + \frac{2}{9} \alpha I_1 \boldsymbol{I} \right) \tag{3.56}$$

3.2.4 Return Mapping

With the yield criterion, hardening laws, and flow rule, the elastoplastic response of the material can be fully defined. Inside the elastic regime, the response of the material is

linear, but in the plastic regime, the set of equations become nonlinear. The equations would be prohibitive to solve analytically except for very simple cases. Consequently, a numerical method must be used to solve the nonlinear equation.

Solving the nonlinear equation can be framed as a nonlinear optimization problem subject to three constraints, also referred to as Karush-Kuhn-Tucker conditions. The first constraint is that the rate of change of the plastic multiplier is nonnegative. This results from the choice to have the flow tensor, $\boldsymbol{N}$, account for the direction of plastic flow and have the plastic multiplier determine the magnitude of the plastic strain to be created. The second condition is that the yield function must always remain non-positive. Yielding occurs when $f = 0$, and the material behaves elastically inside the yield surface, $f < 0$. Finally, the third condition is that the plastic strain can only be created when the yield surfaces is reached. These three conditions can mathematically be written as:

$$\dot{\gamma} \geq 0 \qquad f \leq 0 \qquad \dot{\gamma} f = 0 \tag{3.57}$$

Many researchers have proposed methods for solve the nonlinear equation, but the most popular category of methods is return mapping algorithms. [50] A return mapping algorithm relies on predicting the next stress state given a strain increment and, if the predicted stress state is outside the yield surface, correcting the stress state until it returns to the yield surface, hence the term return mapping.

Most methods use an elastic predictor, but how the stress state is corrected depends on the method. The two most widely used return mapping algorithms are the closest point projection method (CPPM) and the convex cutting plane algorithm (CPA). The difference between the two method lies in how the flow tensor is updated. During the correction steps, the CPPM uses the normal to the surface that the stress is being corrected *to* as the direction, whereas the CPA uses the normal to the surface the stress is being corrected *from* as the direction. The CCPM has been shown to outperform the CPA method, but the CPA method is easier to implement for complex plasticity models. [51] These methods work well for a wide range of plasticity models.

However, when a non-associative flow rule is used, then the normal of the yield surface, f, is not used. Instead, the normal of the surface formed by the flow potential, g, is used. In general, this can introduce many complications, but fortunately, the flow potential and yield surface are both fully differentiable, symmetric about the hydrostatic axis, and paraboloidal in shape, refer to Figure 3.3. Due to these properties, the return mapping path occurs in a plane that includes the elastic predictor stress state and the hydrostatic axis.

Figure 3.4 shows an illustrative example for the CCPM and CPA when an associative flow rule is used and for a return mapping when a non-associative flow rule is used. The increment in strain is used to determine the elastic predictor, and the corrections occur along different paths depending on the return mapping algorithm used. It is important to note that the three cases may not result in the same σ_{n+1}.

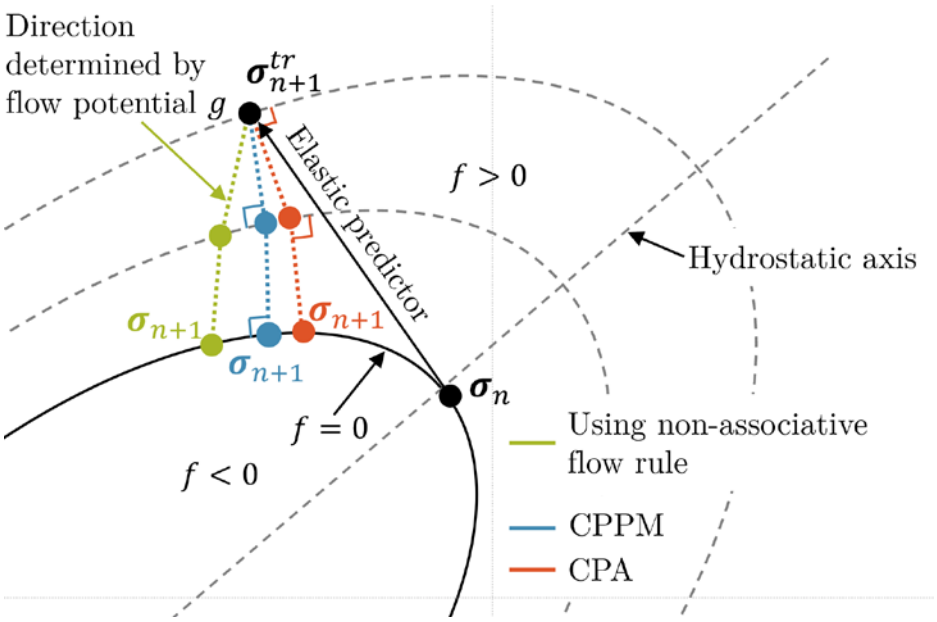

Figure 3.4. Illustrative comparison of return mapping algorithms

Using the discrete form of the plasticity equations, this relatively simple model becomes a single nonlinear equation as a function of $\Delta\gamma$. Many researchers, including Melro et. al., have used the Newton-Raphson method to solve the return mapping problem. Fundamentally, the aim is to find the root of the yield function, $f(\boldsymbol{\sigma}, \sigma_c, \sigma_t)$, given an elastic predictor. However, recall that σ_c and σ_t are a function of the effective plastic strain, ε_e^p, as shown in the hardening laws. The procedure is as follows:

1. Begin with a guess for $\Delta\gamma$
2. Calculate $\Delta\varepsilon_e^p$
3. Update the hardened yield strengths, σ_c and σ_t
4. Calculate f
5. Calculate $\frac{\partial f}{\partial \Delta\gamma}$

6. Calculate the next guess for $\Delta\gamma_{i+1} = \Delta\gamma_i - f_i \left(\frac{\partial f}{\partial \Delta\gamma}\Big|_i \right)^{-1}$

7. Repeat steps 2-6 until $abs(f) \leq tol$

The derivation for the relevant equations are given concisely in Ref. [46], but the final equations are provided here in Equations (3.58)-(3.64).

Step 2:
$$\Delta\varepsilon_{\mathrm{e}}^{\mathrm{p}} = \Delta\gamma\sqrt{kA} \tag{3.58}$$

Step 2:
$$A = \frac{18J_2^{tr}}{(1+6G\Delta\gamma)^2} + \frac{4\alpha^2 (I_1^{tr})^2}{27(1+2K\alpha\Delta\gamma)^2} \tag{3.59}$$

Step 4:
$$f = \frac{6J_2^{tr}}{(1+6G\Delta\gamma)^2} + \frac{2I_1^{tr}}{1+2K\alpha\Delta\gamma}(\sigma_c - \sigma_t) - 2\sigma_c\sigma_t \tag{3.60}$$

Step 5:
$$\frac{\partial f}{\partial \Delta\gamma} = -\frac{72GJ_2^{tr}}{(1+6G\Delta\gamma)^3} - \frac{4K\alpha I_1^{tr}}{(1+2K\alpha\Delta\gamma)^2}(\sigma_c - \sigma_t) + \frac{2I_1^{tr}}{1+2K\alpha\Delta\gamma}\left(\frac{\partial\sigma_c}{\partial\Delta\gamma} - \frac{\partial\sigma_t}{\partial\Delta\gamma}\right) - 2\left(\sigma_t\frac{\partial\sigma_c}{\partial\Delta\gamma} + \sigma_c\frac{\partial\sigma_t}{\partial\Delta\gamma}\right) \tag{3.61}$$

Step 5:
$$\frac{\partial\sigma_{\mathrm{c}}}{\partial\Delta\gamma} = \frac{\partial\sigma_{\mathrm{c}}}{\partial\Delta\varepsilon_{\mathrm{e}}^{\mathrm{p}}}\frac{\partial\Delta\varepsilon_{\mathrm{e}}^{\mathrm{p}}}{\partial\Delta\gamma} = \eta_c H_c\left(\varepsilon_{\mathrm{e}}^{\mathrm{p}}\right)^{\eta_c - 1}\frac{\partial\Delta\varepsilon_{\mathrm{e}}^{\mathrm{p}}}{\partial\Delta\gamma} \tag{3.62}$$

Step 5:
$$\frac{\partial\sigma_t}{\partial\Delta\gamma} = \frac{\partial\sigma_t}{\partial\Delta\varepsilon_{\mathrm{e}}^{\mathrm{p}}}\frac{\partial\Delta\varepsilon_{\mathrm{e}}^{\mathrm{p}}}{\partial\Delta\gamma} = \eta_t H_t\left(\varepsilon_{\mathrm{e}}^{\mathrm{p}}\right)^{\eta_t - 1}\frac{\partial\Delta\varepsilon_{\mathrm{e}}^{\mathrm{p}}}{\partial\Delta\gamma} \tag{3.63}$$

Step 5:
$$\frac{\partial\Delta\varepsilon_{\mathrm{e}}^{\mathrm{p}}}{\partial\Delta\gamma} = \sqrt{kA} - \frac{\Delta\gamma\sqrt{k}}{2\sqrt{A}}\left(\frac{216GJ_2^{tr}}{(1+6G\Delta\gamma)^3} + \frac{16\alpha^3 KI_1^{tr}}{27(1+2K\alpha\Delta\gamma)^3}\right) \tag{3.64}$$

To clarify the relation between the solution of the nonlinear equation and the return mapping sequence, Figure 3.5 shows a typical f vs $\Delta\gamma$ curve along with an illustration of

the return mapping sequence in the plane that contains the elastic predictor and the hydrostatic axis. Recall that Newton-Raphson iteration is used to solve for $\Delta\gamma$ such that $f = 0$, which is point B in Figure 3.5. Because the function changes so rapidly near $\Delta\gamma = 0$, the number of iterations required to reach convergence greatly depends on the initial guess. Additionally, an initial guess for $\Delta\gamma$ beyond the solution will result in the next guess for $\Delta\gamma$ being negative without modification during the Newton-Raphson iterations, refer to points C and D in Figure 3.5. Since a negative plastic multiplier is not permissible, the previous guess for $\Delta\gamma$ is reduced by a factor of 4 if the next guess would be negative.

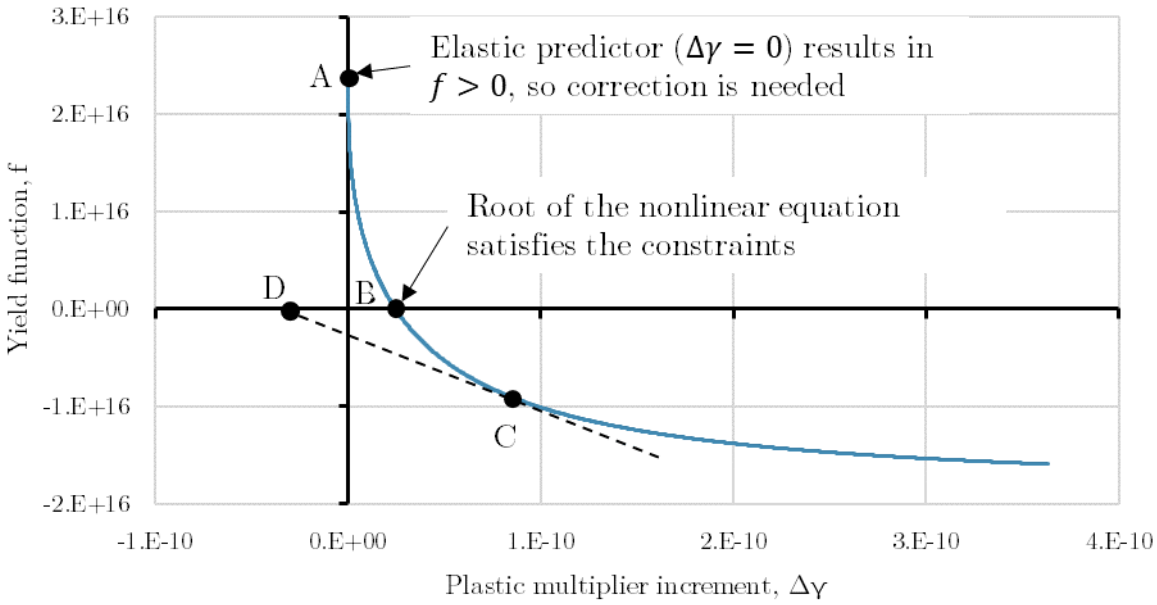

a) f vs $\Delta\gamma$ plot

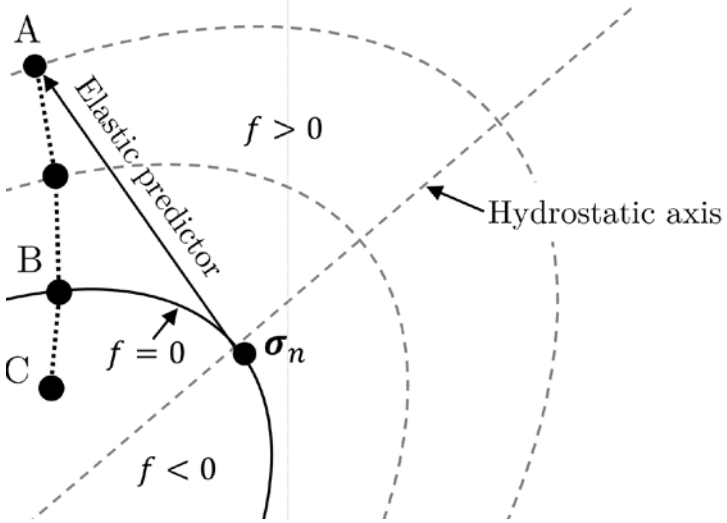

b) Illustration relating points from f vs $\Delta\gamma$ plot to the return mapping with the yield surfaces

Figure 3.5. Typical f vs. $\Delta\gamma$ curve for a load step the results in the elastic predictor lying outside the yield surface

It should be noted that some complications can occur for some hardening laws. For example, consider the two cases in Table 3.1 where the compressive and tensile hardening

parameters differ. In case 1, the tensile and compressive hardening moduli are the same, but the hardening exponents differ. In case 2, the hardening exponents are the same, but the hardening moduli differ. For both cases, a local maximum appears in the f vs $\Delta\gamma$ curve, as shown in Figure 3.6. The presence of a local maximum can result in divergence for the Newton-Raphson method. During the sequence of corrections, if a guess places $\Delta\gamma$ on the left side of the local maximum, the algorithm will drive $\Delta\gamma$ towards 0, which does not satisfy $f = 0$. In this work, the hardening parameters are chosen arbitrarily due to a lack of available experimental data. So, the same parameters will be used for tensile and compressive hardening. However, future work employing more complex hardening laws should expect to require a more powerful nonlinear solver than the Newton-Raphson method.

Table 3.1. Stress invariants and material properties that illustrate local minima

	Case 1	**Case 2**
I_1	28.7e7	-3.54e8
J_2	3.19e13	1.50e16
E	3.72 GPa	
ν	0.35	
σ_{c0}	60 MPa	
σ_{t0}	20 MPa	
H_c	160 MPa	70 MPa
η_c	0.1	0.2
H_t	160 MPa	
η_t	0.2	
ν_p	0.3	

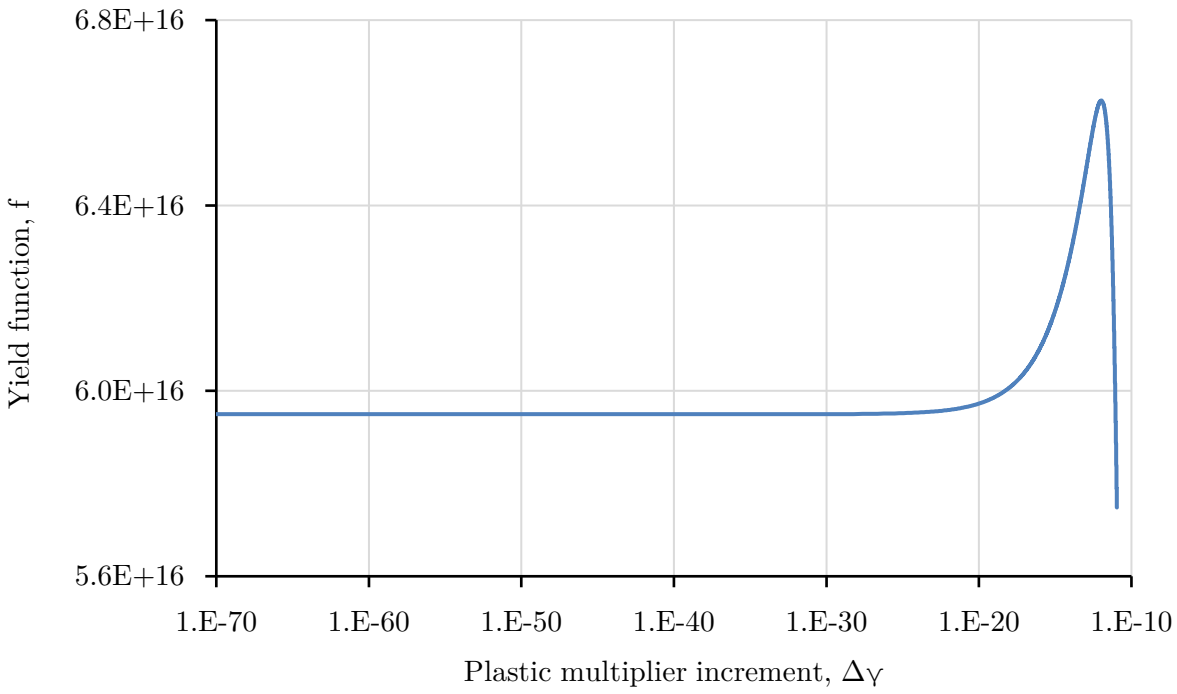

a) f vs $\Delta\gamma$ for case 1 ($\eta_c \neq \eta_t$)

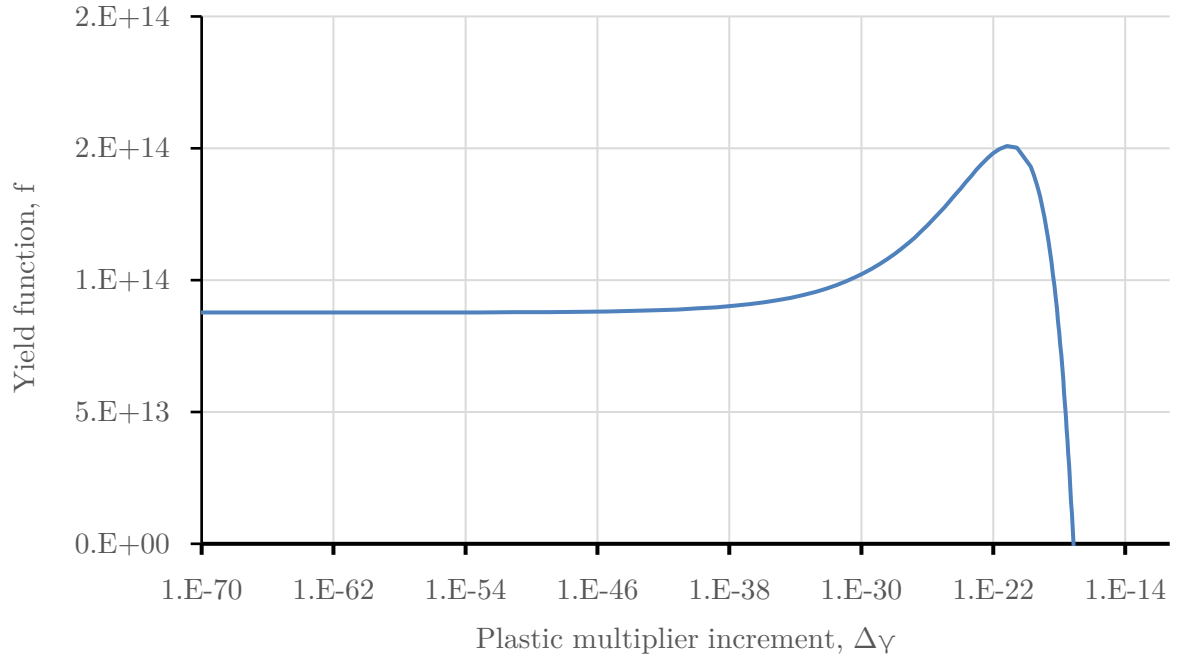

b) f vs $\Delta\gamma$ for case 2 ($H_c \neq H_t$)

Figure 3.6. f vs $\Delta\gamma$ for two cases that exhibit local maxima

3.3 Nonlinear Solution Method

In a simulation with a plastic material, the global system of equations becomes nonlinear. For each load increment, the system of equations becomes implicitly dependent on the displacement vector, since the stress is a nonlinear function of strain in a plastic material. Consequently, a global iteration strategy is needed to determine the displacement vector that satisfies equilibrium.

The governing equation in Equation (3.21) still holds, but the derivation of the finite element method must be modified to allow for a nonlinear constitutive law. The only term affected by introducing plasticity is the stress-strain relation. First consider a decomposition of the strain tensor into its elastic and plastic parts. Using contracted notation, Equation (3.65) relates the strain tensor, sometimes called the total strain tensor, to its elastic and plastic parts, ε_{ij}^{e} and ε_{ij}^{p} respectively. In light of this relation, it is important to recall that the stress is linearly related only to elastic part of the strain via a stiffness tensor, as shown in Equation (3.66). The plastic strain evolves as yielding occurs via the plasticity model discussed in the last section.

$$\varepsilon_i = \varepsilon_i^e - \varepsilon_i^p \tag{3.65}$$

$$\sigma_i = C_{ij}\varepsilon_j^e = C_{ij}\left(\varepsilon_j - \varepsilon_j^p\right) \tag{3.66}$$

When the strain tensor is equivalent to the elastic strain, the finite element method led to Equation (3.42). However, the presence of plastic strain introduces an extra term into the volume integral, resulting in Equation (3.67).

$$\int \left(B_{mi}(\mathbf{x})C_{mn}B_{nj}(\mathbf{x})q_j - B_{mi}(\mathbf{x})C_{mn}\varepsilon_n^p\right)dV = \int t_m \frac{\partial u_m(\mathbf{x})}{\partial q_i} dA + \int \rho b_m \frac{\partial u_m(\mathbf{x})}{\partial q_i}\, dV \tag{3.67}$$

With this new nonlinear term, the classical $Kq = F$ form becomes less useful. Instead, a Newton-Raphson iterative solve will be used to determine an increment of the displacements for each load step. First, it is useful to introduce a residual vector, R_i, which will be the difference between the right-hand side and left-hand side of Equation (3.67). However, the left-hand side can be manipulated to give a simpler form. Equation (3.68) shows the simplified form for the residual vector, which is fundamentally the difference between the externally applied loads and the internal forces. For equilibrium to be maintained, the residual must be zero.

$$R_i = \int t_m \frac{\partial u_m(\mathbf{x})}{\partial q_i} dA + \int \rho b_m \frac{\partial u_m(\mathbf{x})}{\partial q_i}\, dV - \int B_{mi}(\mathbf{x})\sigma_m\, dV \tag{3.68}$$

The residual can be expressed in terms of a Taylor series, as shown in Equation (3.69), where $R_I(\boldsymbol{q})$ indicates the residual vector evaluated with the nodal displacement vector $\mathbf{q}$. The Newton-Raphson method relies on the truncation of the high-order terms, which linearizes the residual function. Since the aim is to determine displacements that result in a residual vector that is zero, $R_i(\boldsymbol{q}+\boldsymbol{\Delta q})=0$, the equation becomes as shown in Equation (3.70).

$$R_i(\boldsymbol{q}+\boldsymbol{\Delta q}) = R_i(\boldsymbol{q}) + \frac{\partial R_i}{\partial q_j}\Delta q_j + \frac{\partial^2 R_i}{\partial q_j^2}\Delta q_j^2 + \cdots \tag{3.69}$$

$$-\frac{\partial R_i}{\partial q_j}\Delta q_j = R_i(\boldsymbol{q}) \tag{3.70}$$

In this form, Newton-Raphson iterations can be used to iteratively correct the displacement vector until the residual becomes close enough to zero. However, the derivative of the residual in terms of the displacements must first be defined. Using the form of the residual vector given Equation (3.68), it is important to recognize that the first two terms are independent of the nodal displacement vector. Recall that $\frac{\partial u_m(\mathbf{x})}{\partial q_i}$ is a matrix of shape functions evaluated at $\mathbf{x}$, and B_{mi} is independent of the nodal displacement vector, as it is a matrix of derivatives of the shape function with respect to the coordinates. Consequently, the derivative of the residual vector becomes Equation (3.71). Plugging Equation (3.71) back into Equation (3.70) yields a system of linear equations and reveals the tangent stiffness matrix, K_{ij}^{tan} .

$$\frac{\partial R_i}{\partial q_j} = -\int B_{mi}(\mathbf{x}) \frac{\partial \sigma_m}{\partial q_j}\, dV = -\int B_{mi}(\mathbf{x}) \frac{\partial \sigma_m}{\partial \varepsilon_n} \frac{\partial \varepsilon_n}{\partial q_j}\, dV$$
$$= -\int B_{mi}(\mathbf{x}) \frac{\partial \sigma_m}{\partial \varepsilon_n} B_{nj}(\mathbf{x})\, dV \tag{3.71}$$

$$\int B_{mi}(\mathbf{x}) \frac{\partial \sigma_m}{\partial \varepsilon_n} B_{nj}(\mathbf{x})\, dV\, \Delta q_j = K_{ij}^{tan} \Delta q_j = R_i(\boldsymbol{q}) \tag{3.72}$$

Now the algorithm becomes:

1. Provide an initial estimate for q_i
2. Calculate $R_i(\boldsymbol{q})$
3. Check convergence (if the residual is small enough then the algorithm is done)
4. Form K_{ij}^{tan}
5. Solve for Δq_i
6. Update the estimate for the displacement vector $q_i^{n+1} = q_i^n + \Delta q_i$
7. Repeat steps 2-6 until convergence is reached

In step 1, the initial estimate can be calculated in many ways. The Newton-Method is known to have trouble when the residual vector has local minima, which was explored in Ref. [52] for cohesive zone models. So, it is important to have a good initial estimate for q_i. This work will begin with an estimate based on solving Equation (3.72) with the residual calculated using the displacement vector from the previous load step, which is zero for the first load step. After a few initial load increments, the initial estimate for q_i will be based

on a linear extrapolation. Let $t-1$ be the pseudo time of the previous load step, and t be current pseudo time. The increment of the displacement for the last time step is given by $\Delta q_i^{t-1} = q_i^{t-1} - q_i^{t-2}$. Suppose the applied loads are scaled according to a load factor, L, and the change in the load factor is similarly given by $\Delta L^{t-1} = L^{t-1} - L^{t-2}$. Then as long as the proportionality of the loads remains constant between the last load increment and the current load increment, the estimate for q_i can be given by the extrapolation shown in Equation (3.73).

$$q_i^t = q_i^{t-1} + \frac{\Delta q_i^{t-1}}{\Delta L^{t-1}} \Delta L^t \tag{3.73}$$

It should be noted that if the proportionality of the load changes from the previous load increment to the current load increment, then the estimate for q_i should be based on a tangent solution, which requires a linear solve. However, this work only considers proportional loading, and this strategy removed the need to perform a linear solve for most of the time steps in the analysis, greatly speeding up the analyses.

In step 2, convergence is assessed based on a relative L2-norm, which is given by the L2-norm of the residual divided by the L2-norm of the internal forces. When the relative L2-norm became smaller than the tolerance for the analysis, 1e-6 in this work, then the solution was considered converged.

Finally, to form the tangent matrix, the derivative of the stress with respect to the total strain is required. For a linear elastic material, the derivative is simply the linear stiffness

matrix, C_{ij}. However, for a plastic material, the derivation for the term can be quite involved. Melro et. al. provides an expression for the term in Refs. [46] and [47]. However, the tangent stiffness did not result in acceptable convergence when plastic strain began to accumulate. Recently, van der Meer discovered some important terms were missing and proposed an improved formulation in Ref. [53]. The formulation was implemented in Beta2, but the iterations sometimes diverged. An independent derivation was developed for the model described in this chapter. However, it differed from the other two derivations mentioned and sometimes resulted in divergence. Unfortunately, this may be a problem with the implementation or the derivation of the tangent stiffness. However, the tangent stiffness could be calculated using a finite difference approach by perturbing each component of the strain tensor by a very small amount, 1e-7 in this work, and calculating the change in stress for each perturbation. After significant optimizations in the code, this method performed efficiently and resulted in stable convergence for all cases considered. The approach may be slightly slower than directly computing the tangent stiffness, but it does not require anything other than the calculating the stress given the strain.

It should be noted that each Newton-Raphson solution typically requires several solutions to a linear set of equations when the materials begin to yield. Additionally, multiple load increments will be needed to keep the solution stable. Consequently, the time required for a plasticity analysis can increase significantly.

4 MICROSCALE ANALYSES OF UNIDIRECTIONAL FIBER/MATRIX COMPOSITES

4.1 Overview

The overall aim of this work is to predict the response of composite materials across scales. In the case of laminated composites, the heterogeneous microstructure of the plies near the interface of interest will be directly modeled, which directly captures the microscale response within a mesoscale analysis. Plies that are away from the interface of interest will be modeled as a homogeneous material, but these ply-level properties can be experimentally measured. However, in the case of three-dimensional textiles, the nonidealized textile models result in a varying cross-sectional area along the path of the tow. In reality, the area of fiber in every cross-section of a tow remains approximately constant, but the local fiber volume fraction can change significantly, refer to Figure 2.37. To investigate the effect of the variation of the local fiber volume fraction on the stress distributions within 3D textiles, a relationship between the fiber volume fraction and material properties is needed. Supporting the study of 3D textiles that is presented in a later chapter, this chapter investigates the dependence of mesoscale linear elastic properties for unidirectional fiber-matrix composites on the fiber volume fraction.

There are a variety of analytical homogenization methods and FEM models with simple fiber arrangements that could be used to predict the effective properties, but accounting of the randomness of the heterogeneous microstructure has been shown to significantly affect the predicted effective properties, especially the predicted shear moduli and Poisson's ratios.

[36] [47] [37] However, accounting for the randomness of the fiber arrangements requires an ensemble of models to understand the distribution of effective properties. This work investigates the effect of fiber volume fraction on the effective properties used at the mesoscale by using an ensemble of random fiber/matrix RVEs for each fiber volume fraction considered. The mean effective properties from each ensemble are fit using a cubic spline. In a later chapter, the spline will be used at the mesoscale to interpolate the effective properties for each element in the tows of a 3D textile model.

The first section of this chapter provides the relevant background information. This is followed by a parametric study of the effect of the fiber volume fraction on the effective stiffness matrix for different fiber arrangements. Next, the results from the study are used to develop a mesoscale constitutive law. Finally, a summary of the microscale analyses is provided.

4.2 Background

This section begins with a description of the strategy for creating the random RVEs, the number of fiber arrangements in each ensemble, and the range of fiber volume fractions considered. Next, the meshes used for the fiber/matrix RVEs are briefly presented. Then, a detailed explanation of the periodic boundary conditions are provided, followed by the assumed microscale constituent properties. Finally, the method used to calculate the effective properties is discussed.

4.2.1 Random Fiber/Matrix RVEs

Historically, random fiber/matrix RVE have been generated with square shapes. [36] [37] [47] Researchers have developed several computational methods for generating the random fiber arrangements, and the square shape makes the implementation of an algorithms relatively simple. One popular method is to perturb a periodic arrangement of fibers, which has the advantage of successfully creating very high fiber volume fractions if a hexagonal packing is perturbed. [54] [55] [56] However, highly random fiber arrangements are not possible at low fiber volume fractions. Another method that is common in the recent literature is to randomly place fibers within the RVE and remove the overlap between them through an iterative contact algorithm. [57] [36] [58] [37] [47] It has been noted in Ref. [47] and [37] that the method performs well for fiber volume fractions up to 65%. The method was initially attempted for the fiber volume fractions considered in this work, which ranges from 40% to 80%, but the method failed to remove the overlap between fibers for some initial arrangements of fibers above 70% fiber volume fraction. The percentage of arrangements that failed in an ensemble increased as the fiber volume fraction increased. At 80% fiber volume fraction, it was very difficult to successfully create a nonoverlapping fiber arrangement, much less an ensemble of fiber arrangements. The underlying issue is that the maximum packing within a square RVE shape is lower than the 90.7% theoretical limit for hexagonal packing. Consequently, arrangements of fibers within a square RVE are more likely to lock up during the iterative removal of overlap between fibers.

To allow the reliable generation of random RVEs with high fiber volume fractions, the RVE shape, RVE size, and number of fibers should be selected such that a hexagonal close packing is possible. This maximizes the theoretical packing limit for the RVE and decreases the probability of locking to occur. It should also be noted that a periodic square RVE shape cannot result in transversely isotropic properties, since it is incapable of allowing a rotational symmetry that exists in transversely isotropic materials. Ref. [36] showed that as the number of fiber arrangements in an ensemble increases, the mean effective properties approaches tends towards orthotropy but not transverse isotropy. For the limit to correctly be transverse isotropy, an RVE shape that allows transverse isotropy must be used.

Figure 4.1 shows the first three hexagonally shaped RVEs in the series that allow hexagonal close packing. Note that the RVE boundaries cut through the fibers. The number of fibers within each RVE, N_{fibers}, is given by Equation (4.1), and the side length of the hexagonal RVE, a, is given by Equation (4.2), where L is the level shown in Figure 4.1, r is the average radius of the fibers, and V_f is the fiber volume fraction. The process for creating the hexagonal RVEs was discussed in chapter 2 of this manuscript. For this study, hexagonal RVEs with 25 fibers, which is level 3 in Figure 4.1, and a constant fiber radius of $5e\text{-}6\,m$ are used. It should be noted that Gherissi et. al. studied the variation of the fiber volume fraction within the cross-sections of tows in order to predict the size of RVE necessary to predict statistically significant properties for the tows. The study showed that an RVE with 15 fibers was appropriate, which this work exceeds. [59]

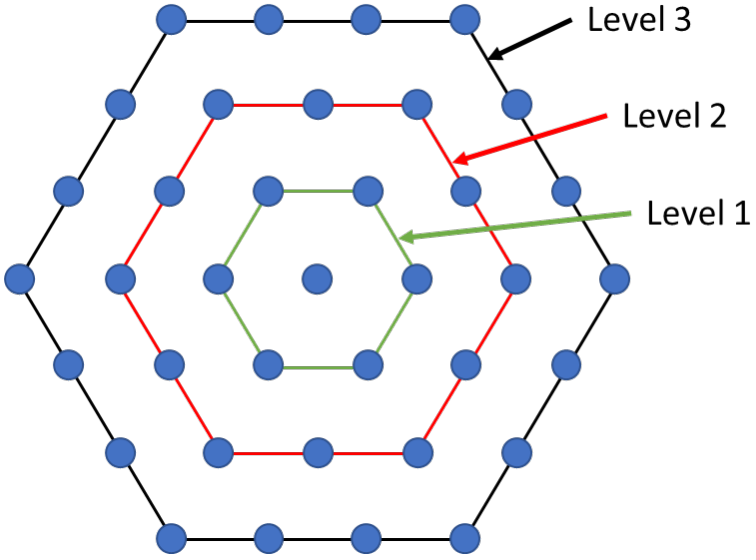

Figure 4.1. RVE sizes that allow hexagonal packing

$$N_{fibers} = 3L + 6\sum_{i=0}^{L-1} i \tag{4.1}$$

$$a = L\sqrt{\frac{2\pi r^2}{\sqrt{3}V_f}} \tag{4.2}$$

Since the fiber arrangements are random, an ensemble of models is used for each fiber volume fraction considered to capture the distribution of properties. In Ref. [36], the average engineering properties were plotted as a function of the number of fiber arrangements in the ensemble for a carbon fiber/epoxy matrix composite system. It was shown that the average properties of an ensemble with 50 random fiber arrangements

differed from a reference ensemble, which contained 500 fiber arrangements, by less than 2%. It should be noted that guidelines given in Ref. [36] for the number of fiber arraignments were in the context of inverse analyses. In an inverse analysis, the lamina level properties must be predicted quite accurately to accurately determine the constituent properties. Since this work does not conduct an inverse analysis, 50 random fiber arrangements were considered for each fiber volume fraction. Based on the range of predicted fiber volume fractions in the simulated 3D textile model, shown in Figure 2.37, 18 ensembles were created for fiber volume fractions ranging from 40% to 80%. Figure 4.2 provides two example RVEs for 40%, 60%, and 80% fiber volume fraction. As shown in the figure, the amount of randomness that can occur in the fiber arrangements decreases as the fiber volume fraction increases towards the hexagonal close packing limit.

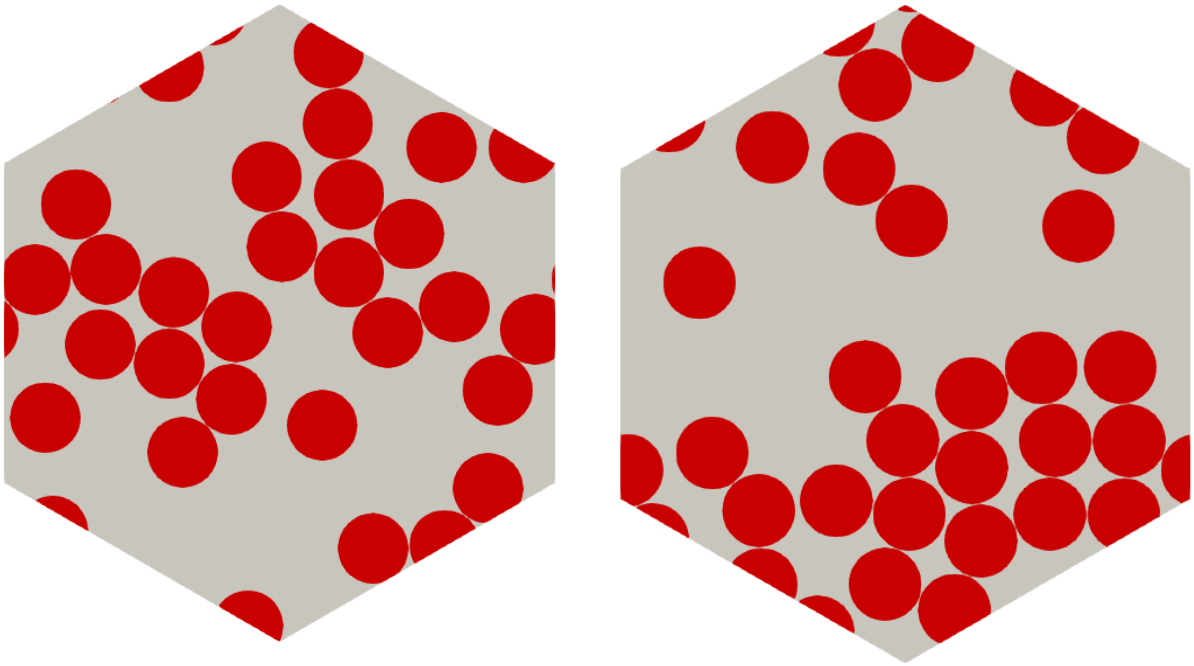

a) 40% V_f

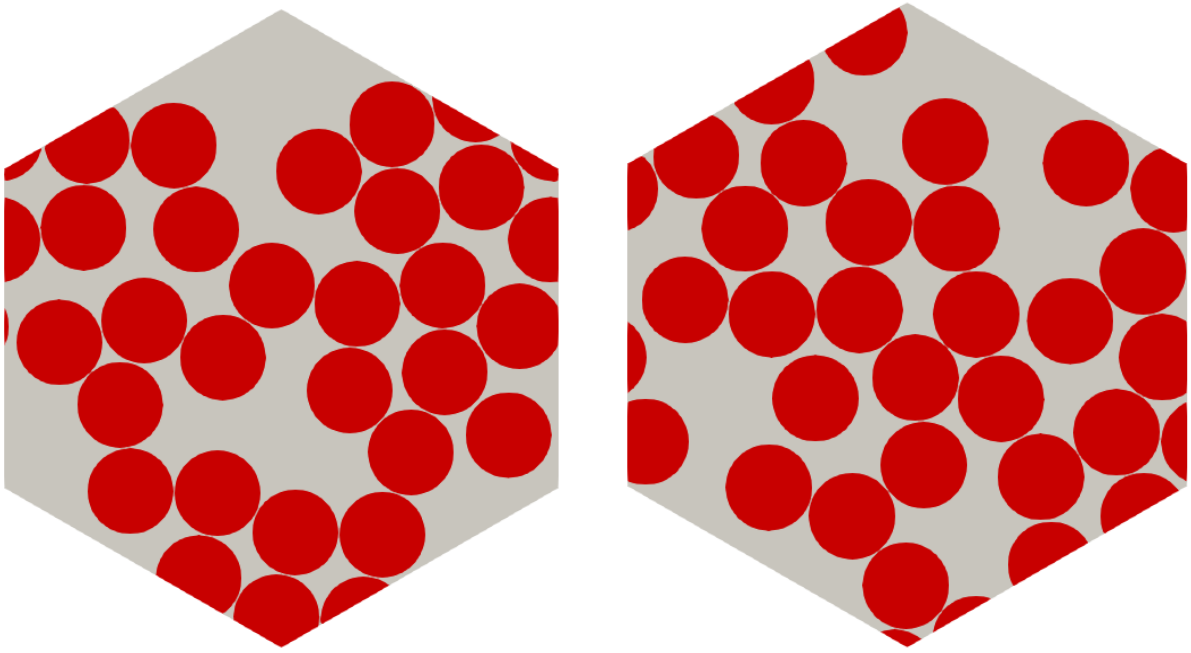

b) 60% V_f

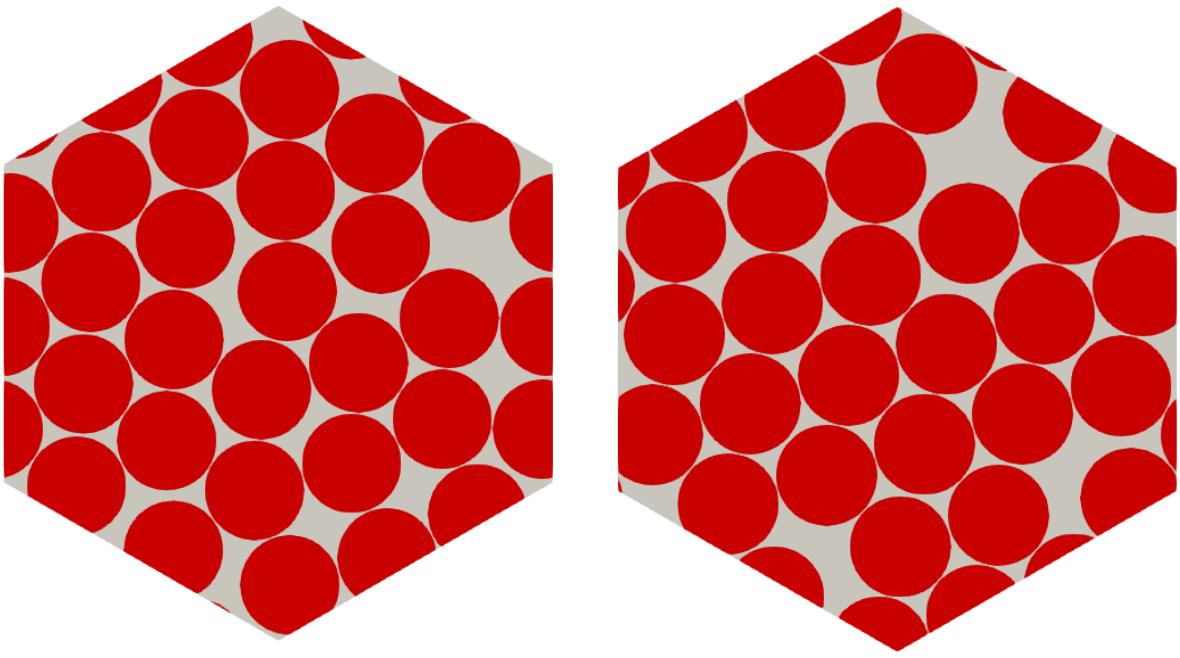

c) 80% V_f

Figure 4.2. Two random fiber arrangements for three selected fiber volume fractions, V_f

4.2.2 RVE Meshes

The 2D random fiber arrangements were meshed using quadratic triangle elements. Nodes were moved to conform to the circular fiber boundaries, and a higher mesh refinement was used in regions where fibers came close. Figure 4.3 shows a typical triangle mesh for a random fiber/matrix RVE with 60% fiber volume fraction. Though the 2D meshes could be directly used in a quasi-3D analysis to predict the 3D effective stiffness matrix, it was more convenient to use 3D meshes with the finite element framework used in this work. Consequently, the 2D meshes were extruded to create 3D meshes consisting of quadratic wedge elements. Conventionally, lamina properties are given where x_1 indicates the fiber direction, so the meshes were rotated so that the fibers were aligned x_1-axis.

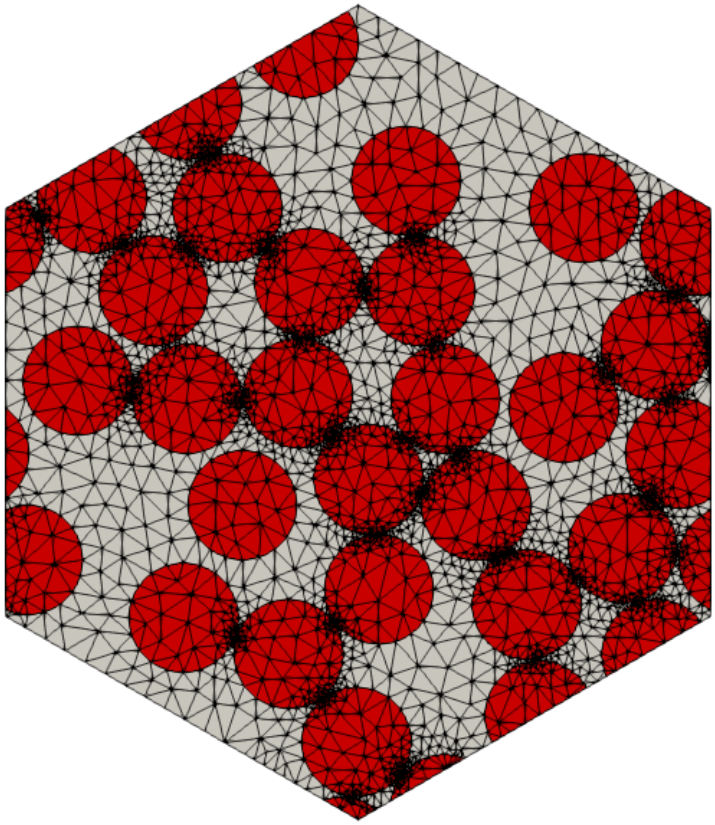

Figure 4.3. Mesh used for left fiber arrangement with 60% V_f in Figure 4.2

4.2.3 Periodic Boundary Conditions

The fiber/matrix RVEs were generated to be geometrically periodic, and the mesh was generated to place nodes along the boundaries periodically, so it is possible to exactly impose periodic boundary conditions. Periodic boundary conditions for the unit cell are concisely expressed in terms of the displacements along the unit cell boundaries in Equation (4.3) using indicial notation. Repeated indices imply summation, and d_k is a vector of periodicity that starts at some boundary coordinate x_j in one unit-cell and ends at the equivalent point the next unit cell.

$$u_i(x_j + d_j) = u_i(x_j) + \langle \frac{\partial u_i}{\partial x_k} \rangle d_k \tag{4.3}$$

For hexagonally shaped RVEs, the boundary conditions are explicitly given in Equations (4.4)-(4.6). The coordinate system and the parameters are defined in Figure 4.4. Equation (4.4) gives the conditions for periodicity, and Equations (4.5) and (4.6) are used to prevent rigid body motion.

$$\begin{gathered} u_i^{\bar{a}}(x_1, x_2, x_3) = u_i^a(x_1, x_2 - m, x_3 - n) + \left\langle \frac{\partial u_i}{\partial x_2} \right\rangle m + \left\langle \frac{\partial u_i}{\partial x_3} \right\rangle n \\ u_i^{\bar{b}}(x_1, x_2, x_3) = u_i^a(x_1, x_2 - m, x_3 + n) + \left\langle \frac{\partial u_i}{\partial x_2} \right\rangle m - \left\langle \frac{\partial u_i}{\partial x_3} \right\rangle n \\ u_i\left(x_1, \frac{d}{2}, x_3\right) = u_i\left(x_1, -\frac{d}{2}, x_3\right) + \left\langle \frac{\partial u_i}{\partial x_2} \right\rangle d \\ u_i\left(\frac{e}{2}, x_2, x_3\right) = u_i\left(-\frac{e}{2}, x_2, x_3\right) + \left\langle \frac{\partial u_i}{\partial x_1} \right\rangle e \end{gathered} \tag{4.4}$$

$$\left\langle \frac{\partial u_i}{\partial x_j} \right\rangle = \left\langle \frac{\partial u_j}{\partial x_i} \right\rangle \tag{4.5}$$

$$u_i(0,0,0) = 0 \tag{4.6}$$

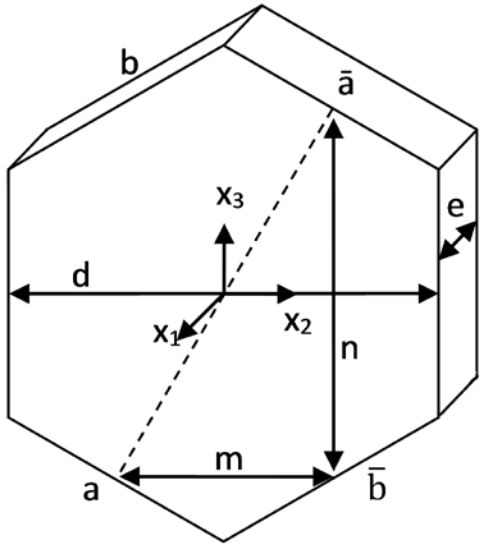

Figure 4.4. Dimension of hexagonal RVE relevant to periodic boundary conditions

4.2.4 Microscale Constituent Properties

Fibers and matrix were modeled discretely and assumed to be IM7 graphite fibers and 5220-4 epoxy matrix respectively. The properties for the epoxy matrix was taken from Ref. [60]. However, it was shown in Ref. [36] that the fiber properties need to be inversely determined if the effective properties are to accurately represent the homogenized microstructure. In Ref. [36], the apparent fiber properties were inversely determined using the effective ply properties and matrix properties in Ref. [60]. Those same apparent fiber properties determined in Ref. [36] were used in this paper, so that homogenizing random fiber/matrix models of unidirectional composites results in the effective properties in Ref. [60]. For this study, the constituent materials were assumed to be linearly elastic. Table 5.2 summaries the properties used for the fibers and matrix.

Table 4.1. Material properties for microscale study

IM7 Graphite Fibers [36]	
E_1^f (GPa)	276
E_2^f, E_3^f (GPa)	26.0
G_{12}^f, G_{13}^f (GPa)	20.7
G_{23}^f (GPa)	7.55
ν_{12}^f, ν_{13}^f	0.292

5220-4 Epoxy Matrix [60]	
E^m (GPa)	3.45
G^m (GPa)	1.278
ν^m	0.35

4.2.5 Calculation of Homogenized Stiffness Matrix

The heterogenous microscale RVEs were assumed to be infinitely repeating using periodic boundary conditions. The mean-field properties, also referred to as the effective or homogenized properties, of a material consisting of the repeated RVE are determined using the volume average stress and volume average strain under a variety of loads. In an analysis with a length scale larger than the RVE size, the unidirectional composite can be treated as a homogeneous material with the effective properties determined from the RVE analyses with little error introduced.

The effective stiffness matrix, which relates the volume average stress to the volume average strain, is determined by specifying only one volume average strain component at a time to be non-zero and calculating the volume average stresses. Consequently, the full effective stiffness matrix requires six load configurations. Each term of the effective stiffness matrix, $\tilde{C}_{ij}$, is calculated using Equation (4.7), where $\Delta\langle\sigma_i\rangle$ is the change in volume average stress from a zero volume average strain state and $\langle\varepsilon_j\rangle$ is the volume average strain. Since

no thermal or moisture effects were considered, zero volume average strain results in zero volume average stress for these analyses. A component of volume average strain can be imposed by specifying the respective volume average displacement gradient in the periodic boundary conditions.

$$\tilde{C}_{ij} = \frac{\Delta\langle\sigma_i\rangle}{\langle\varepsilon_j\rangle} \tag{4.7}$$

4.3 Effect of Fiber Volume Fraction on the Effective Stiffness Matrix

Since each different random fiber arrangement will result in a different effective stiffness matrix, an ensemble of fiber arrangements have a distribution of effective stiffness matrices. A Python script was written to run the various cases, read the output data, and store the effective stiffness matrix of each RVE for each fiber volume fraction considered. No assumption about isotropy was made and the effective properties were generally anisotropic

In general, the effective stiffness matrix has 21 unique terms. The terms of a linear elastic stiffness matrix can be divided into three groups: normal terms, normal-shear coupling terms, which capture normal stresses that occur for an applied shear strain and vice versa, and shear terms. For an orthotropic material, which exhibits three planes of symmetry, all normal-shear coupling terms and off-diagonal shear terms are zero. A uniform arrangement of fibers, such as square or hexagonal packing, exhibits orthotropic effective properties, due to mirror symmetries about three planes. However, with random fiber arrangements, the material is expected to be anisotropic. Because the 2D fiber arrangement

is extruded, the only symmetry will be a mirror symmetry about the x-y plane, resulting in what is classified as a monoclinic material. The stiffness matrix of a monoclinic material with a mirror symmetry about the y-z plane will have the form shown in Equation (4.8). Note that there are 13 unique terms in the stiffness matrix of a monoclinic material.

$$\begin{bmatrix}\sigma_{11}\\ \sigma_{22}\\ \sigma_{33}\\ \sigma_{12}\\ \sigma_{23}\\ \sigma_{13}\end{bmatrix} = \begin{bmatrix} C_{11} & C_{12} & C_{13} & 0 & C_{15} & 0 \\ & C_{22} & C_{23} & 0 & C_{25} & 0 \\ & & C_{33} & 0 & C_{35} & 0 \\ & & & C_{44} & 0 & C_{46} \\ & \text{sym} & & & C_{55} & 0 \\ & & & & & C_{66} \end{bmatrix} \cdot \begin{bmatrix}\varepsilon_{11}\\ \varepsilon_{22}\\ \varepsilon_{33}\\ 2\varepsilon_{12}\\ 2\varepsilon_{23}\\ 2\varepsilon_{13}\end{bmatrix} \tag{4.8}$$

Figure 4.5 shows the distributions of normal terms in the effective stiffness matrices as a function of fiber volume fraction. Error bars indicate ± 1 standard deviation. It is widely known that C_{11}, which is the diagonal stiffness term in the direction of the fibers, follows a rule of mixtures quite well. So, as expected, C_{11} exhibits a linear dependence on the fiber volume fraction and very little variation between fiber arrangements, as shown in Figure 4.5a. The coefficient of variation, which is the standard deviation divided by the mean, remained below 0.03% for all fiber volume fractions considered.

Unidirectional composites are typically treated as transversely isotropic materials. The assumptions made for the fiber/matrix models only assume a monoclinic material, and it is expected that many of the fiber arrangements do not exhibit transverse isotropy. However, if a sufficient number of different fiber arrangements is considered, the mean properties should tend towards transverse isotropy. Figure 4.5b and Figure 4.5c show the pairs of

normal terms that should be equal for a transversely isotropic material, namely C_{22}-C_{33} and C_{12}-C_{13}. As expected, the stiffness terms are not linear with respect to the volume fraction. The coefficient of variation for these terms was significantly higher than the coefficient of variation for C_{11}. C_{22} and C_{33} had a maximum coefficient of variation of 2.8%, while C_{12} and C_{13} had a maximum coefficient of variation of 1.7%. For the volume fractions considered, the mean of the distributions for C_{22} and C_{33} differed at most by 1.2%, while C_{12} and C_{13} differed at most by 0.73%. The mean effective properties did not exhibit transverse isotropy, but it is expected that the mean properties will become more transversely isotropic with more fiber arrangements in the ensemble or larger RVE sizes.

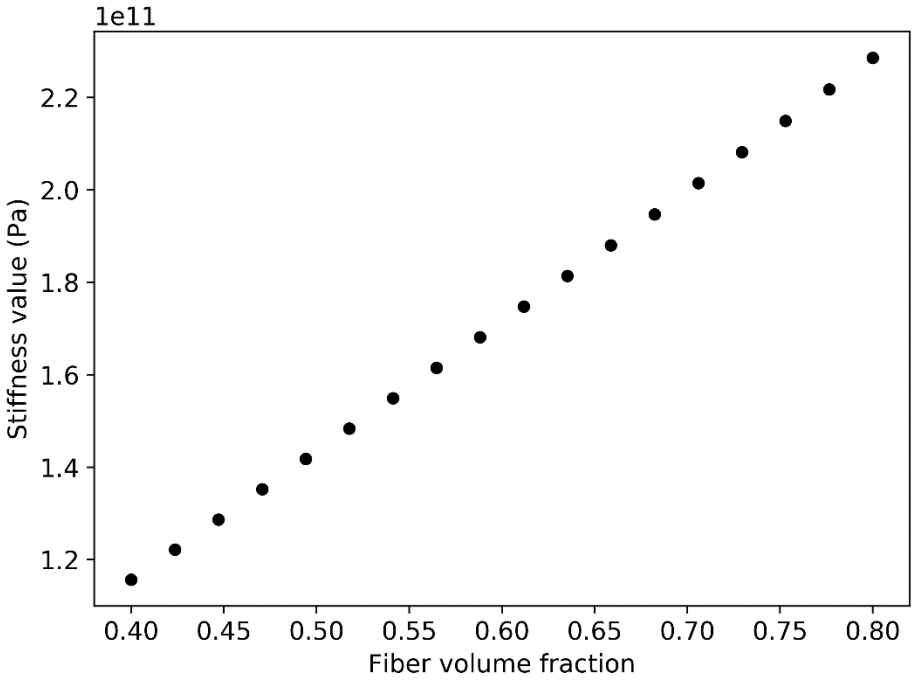

a) C_{11}

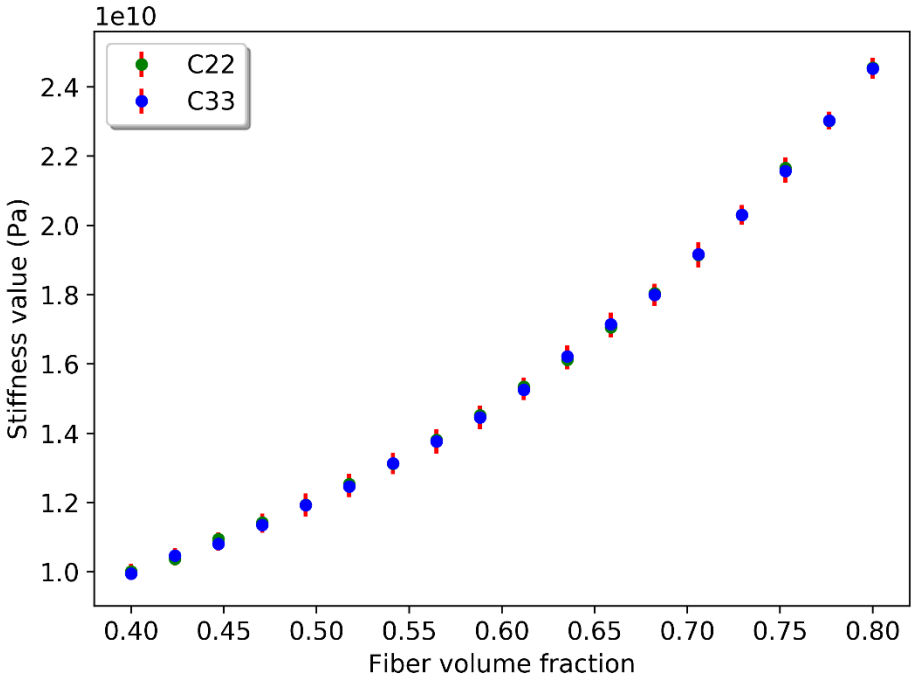

b) C_{22} and C_{33}

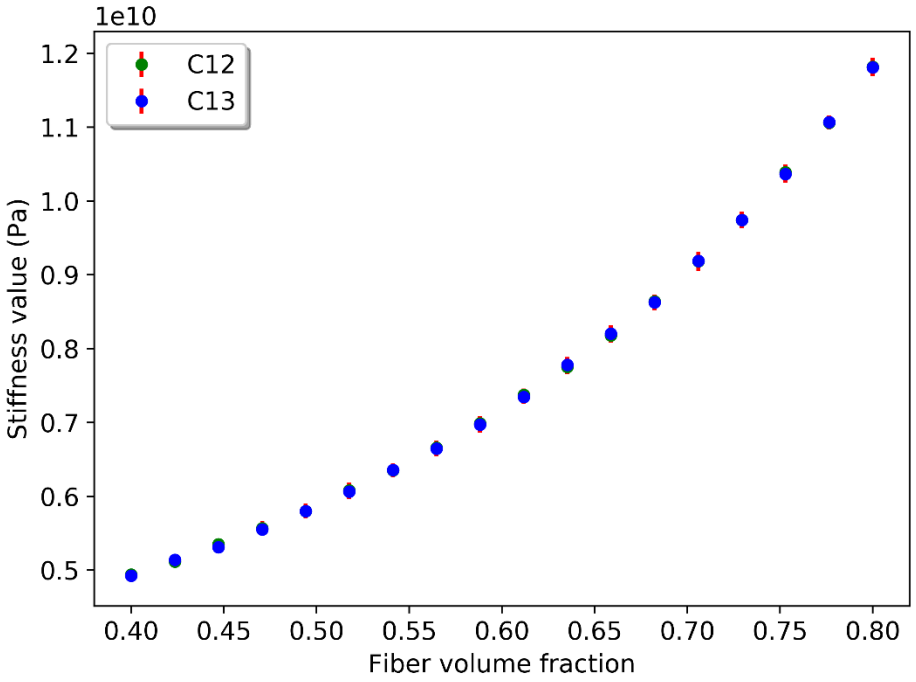

c) C_{12} and C_{13}

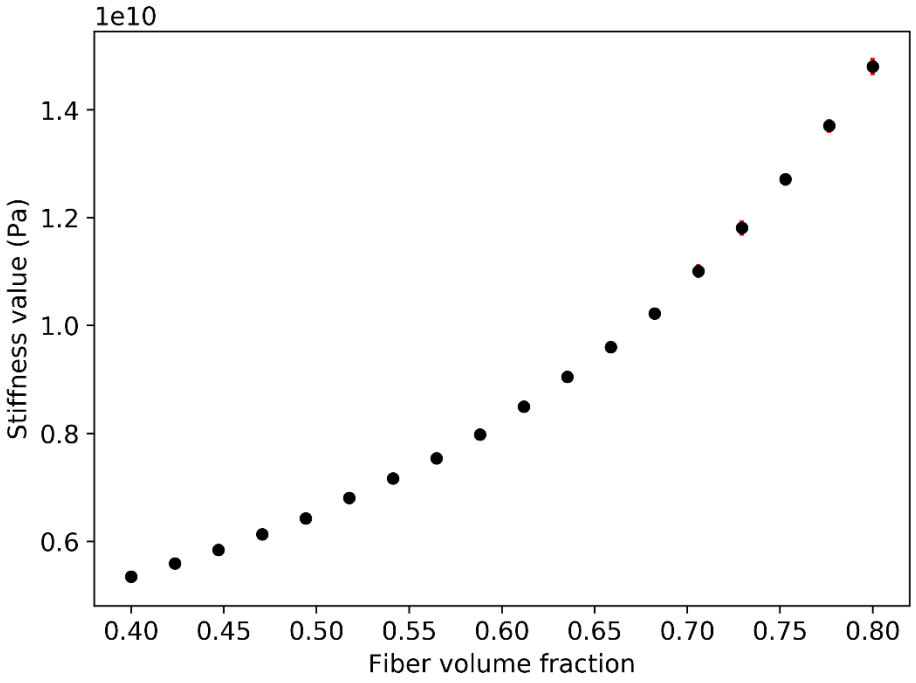

d) C_{23}

Figure 4.5. Mean effective normal stiffness terms as a function of fiber volume fraction with error bars indicating ±1 standard deviation

For a material that exhibits mirror symmetries about three planes, the normal-shear coupling terms become zero. Due to the general lack of symmetry about the x-y and x-z planes for random fiber arrangements, C_{15}, C_{25}, and C_{35} are expected to be non-zero for a random fiber arrangement. However, if a sufficient number of fiber arrangements are considered, then the mean value of the normal-shear coupling terms are expected to tend towards zero, though the standard deviation may be quite large. Figure 4.6 shows the mean effective shear-normal coupling stiffness terms as a function of fiber volume fraction with error bars indicating ± 1 standard deviation. The mean of all three non-zero normal-shear coupling terms remained between -3.8e7 and 4e7 Pa, which is about two orders of magnitude smaller than the smallest term on the diagonal of the effective stiffness matrix. However, the terms exhibited a significant standard deviation for every fiber volume fraction considered. The fiber volume fraction did not seem to affect the standard deviation significantly. The range of standard deviations for C_{15} remained between 4.8e7 and 9.2e7. The range of standard deviations for C_{25} remained between 9.1e7 and 1.9e8. Finally, the range of standard deviations for C_{35} remained between 7.7e7 and 1.6e8. Note that the magnitude for the standard deviation is much larger than the mean. However, the standard deviation of these coupling terms remained less than 5% of any of the mean of the normal stiffness terms.

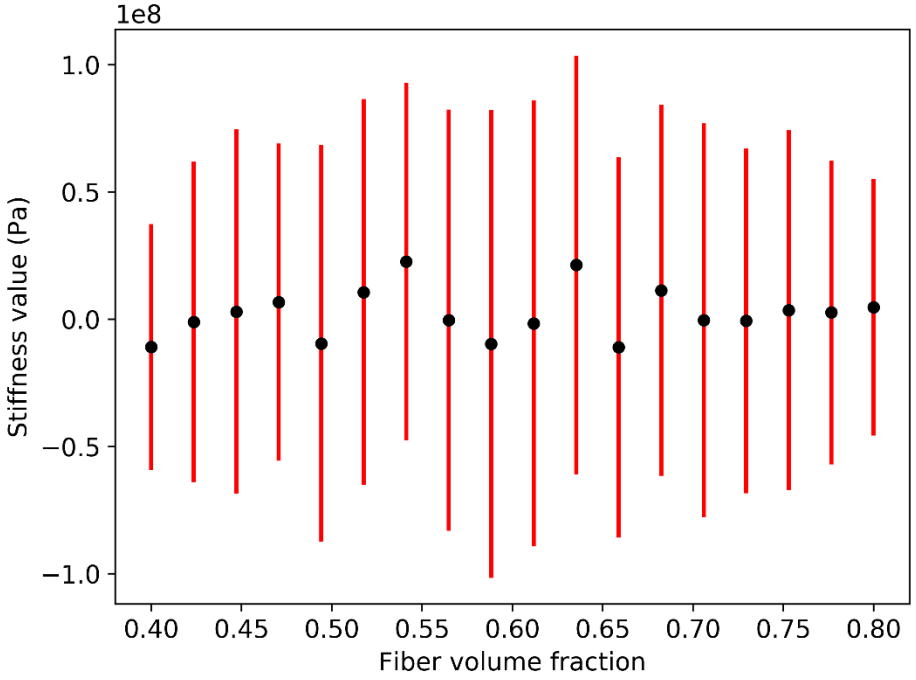
1e8
1.0
0.5
0.0
−0.5
−1.0
Stiffness value (Pa)
0.40
0.45
0.50
0.55
0.60
0.65
0.70
0.75
0.80
Fiber volume fraction

a) C_{15}

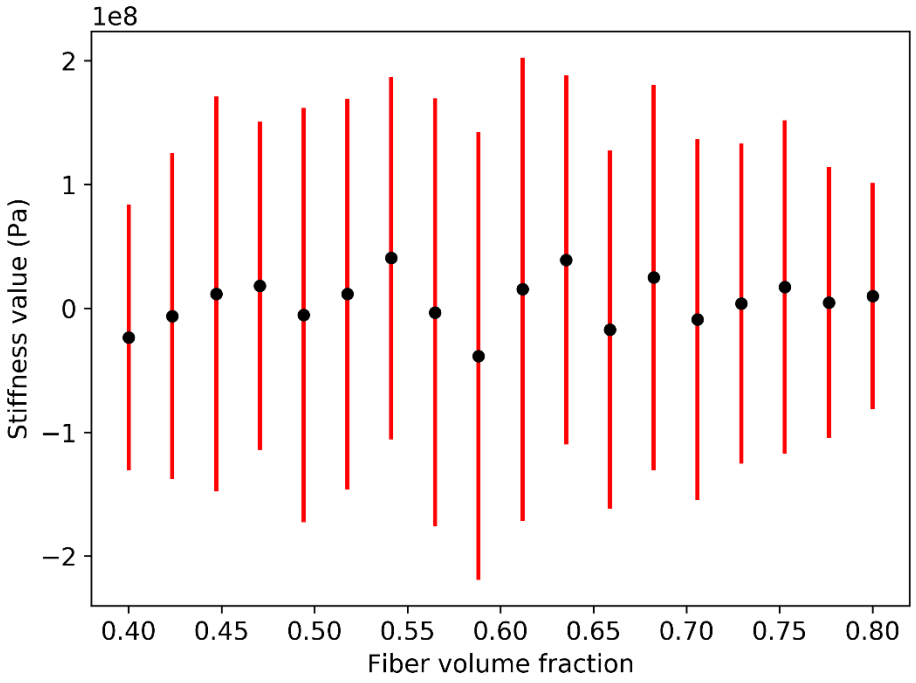
1e8
2
1
0
−1
−2
Stiffness value (Pa)
0.40
0.45
0.50
0.55
0.60
0.65
0.70
0.75
0.80
Fiber volume fraction

b) C_{25}

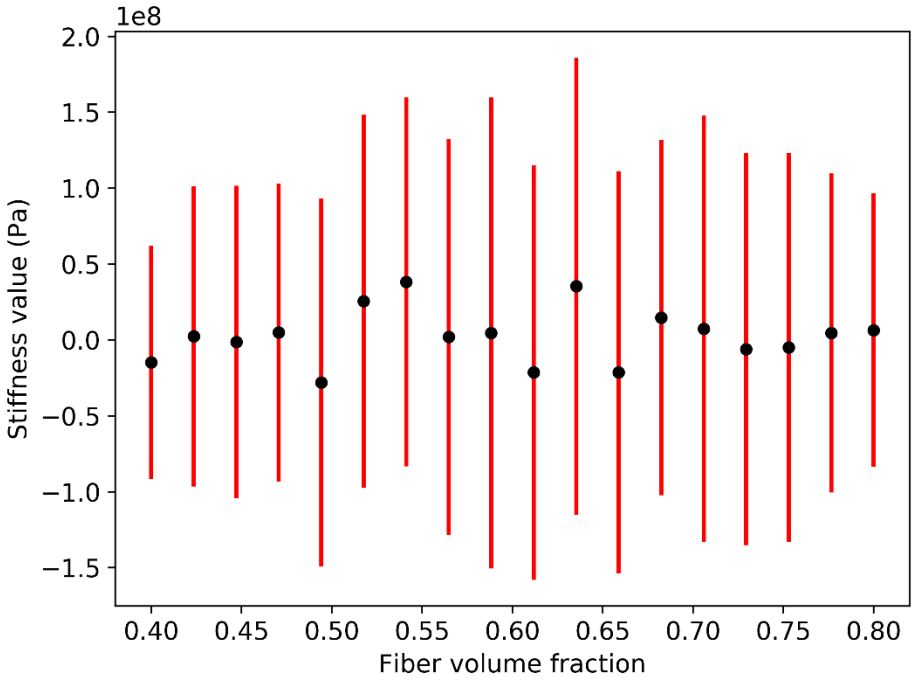

c) C_{35}

Figure 4.6. Mean effective shear-normal coupling stiffness terms as a function of fiber volume fraction with error bars indicating ±1 standard deviation

Figure 4.7 shows the mean effective shear stiffness terms as a function of fiber volume fraction with error bars indicating ±1 standard deviation. As shown in Figure 4.7a, the longitudinal shear terms, C_{44} and C_{66}, exhibited a nonlinear relationship with the fiber volume fraction and larger standard deviations than the normal stiffness terms shown in Figure 4.5. The coefficient of variance for both terms ranged from 2.1% to 7.3%. For a transversely isotropic material, C_{44} and C_{66} should be equal, but for the ensembles considered, the mean value of C_{44} and C_{66} differed by at most 3%, which is less than coefficients of variance for the terms. The in-plane shear stiffness term exhibited smaller standard deviations than the longitudinal terms, as shown in Figure 4.7b, with a coefficient

of variance of at most 2.2%. C_{46} is related to a coupling between the two directions of longitudinal shear, and the term is exactly zero for an orthotopic material. Similar to the shear-normal coupling terms, the mean value of C_{46} was quite small, but the standard deviation of the terms was significantly larger than the mean value. The standard deviation of C_{46} ranged from 1.3e8 to 2.7e8, making C_{46} the term with the largest standard deviations in the effective stiffness matrix. Since the terms along the diagonal of effective stiffness matrix increase with increasing fiber volume fraction, due to the stiff properties of the fiber relative to the epoxy matrix, the longitudinal shear coupling is less significant for larger fiber volume fractions. For the low fiber volume fractions considered, the standard deviation of C_{46} was about 5% of the longitudinal stiffness terms. With such a low magnitude relative to the diagonal shear stiffness terms, the different fiber arrangements did not exhibit a significant longitudinal shear coupling.

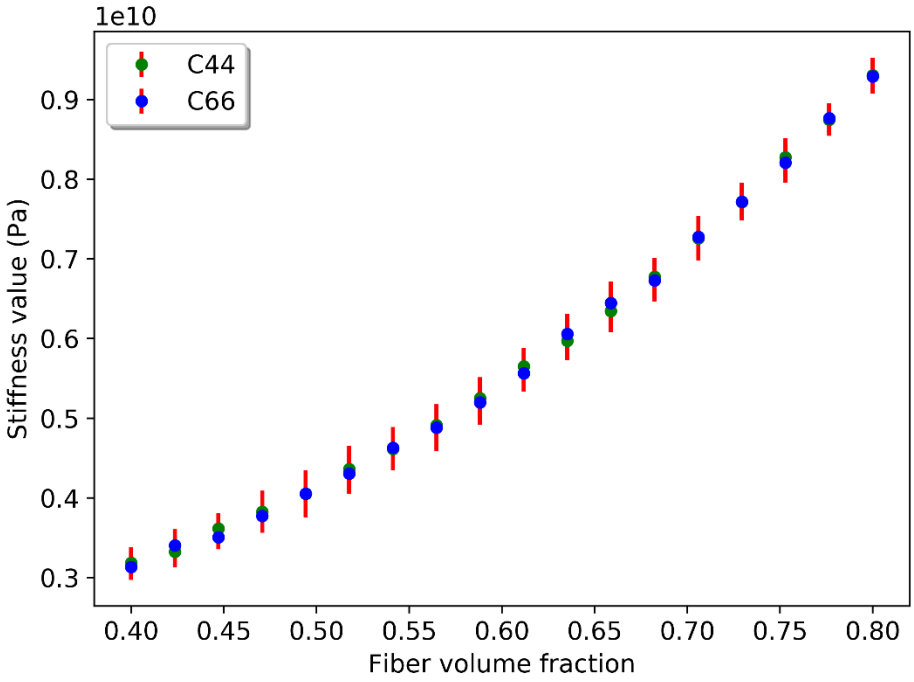

a) C_{44} and C_{66}

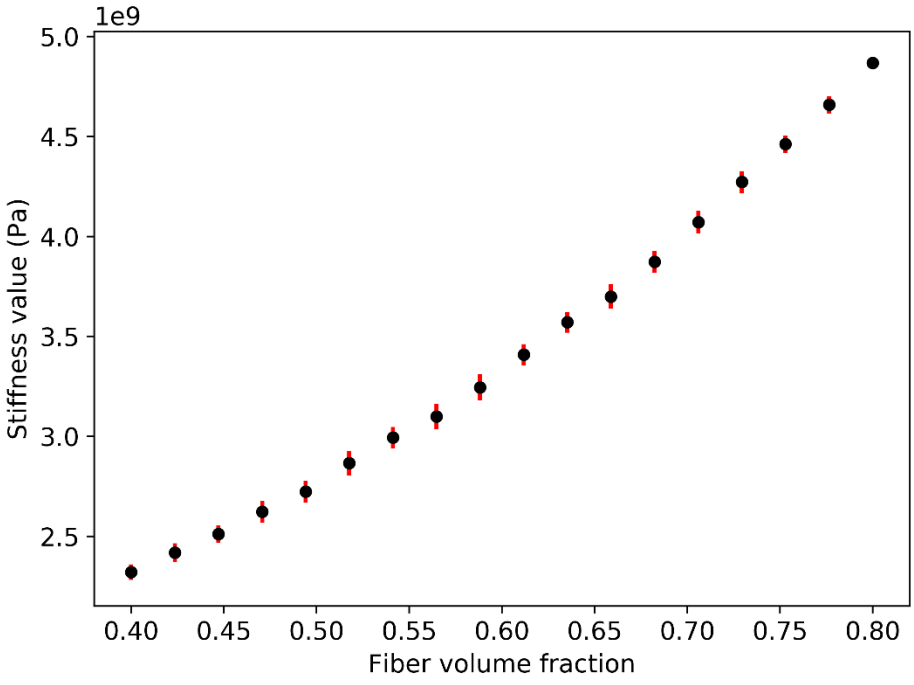

b) C_{55}

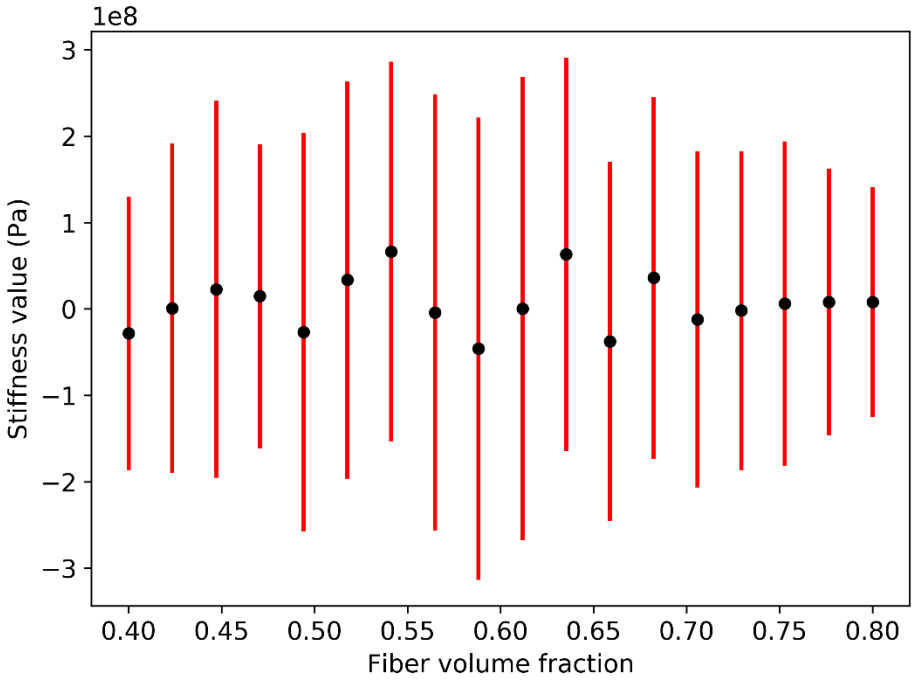

c) C_{46}

Figure 4.7. Mean effective shear stiffness terms as a function of fiber volume fraction with error bars indicating ±1 standard deviation

4.4 Development of a Mesoscale Constitutive Law

To use the effective properties of the ensembles of fiber arrangements at the mesoscale, a model must be developed that can provide the stiffness matrix for each quadrature point in a mesoscale analysis. In this work, the local fiber volume fraction is predicted for each element of the 3D textile models based on the change in cross-sectional area. A constitutive law was developed in the FEA framework that takes a list of effective stiffness matrices at discrete fiber volume fractions, fits each term using a cubic spline, and provides the interpolated stiffness matrix for any given volume fraction within the range of discrete fiber

volume fractions. The mean effective stiffness matrix for each volume fraction considered could be directly used at the mesoscale with this strategy. The properties were relatively close to transversely isotropic. However, the effect of locally anisotropic properties has not be explored. So, to avoid introducing more parameters in the textile analyses, the mean effective stiffness matrix for each ensemble will be modified to be exactly transversely isotropic, and each term as a function of fiber volume fraction will be smoothed.

4.4.1 Forcing Transverse Isotropy

To force the mean effective stiffness matrix for each fiber volume fraction to exhibit transverse isotropy, several terms were modified. First, the non-zero coupling terms, which are C_{15}, C_{25}, C_{35}, and C_{46}, were assumed to be zero. Second, C_{22} and C_{33} were set to be average of the two terms, so they had the same value. Likewise, the same was done to the C_{12} and C_{13} pair of terms and the C_{44} and C_{66} pair of terms. Lastly, the isotropic relationship between E_{22}, ν_{23}, and G_{23} was enforced, which is given in Equation (4.9). To enforce Equation (4.9), each mean effective stiffness matrix was inverted after the terms discussed in the previous two steps were modified, then the shear term of the compliance matrix, S_{46}, was modified. The modified compliance matrix was then inverted back to result in the transversely isotropic stiffness matrix.

$$G_{23} = \frac{E_{22}}{2(1+\nu_{23})} \tag{4.9}$$

Figure 4.8 shows the mean Young's and shear moduli as a function of fiber volume fraction after transverse isotropy was enforced with errors bars indicating ± 1 standard deviation. As expected, the longitudinal Young's modulus varied linearly with respect to the fiber volume fraction with little variation. The transverse Young's modulus and longitudinal shear modulus both varied nonlinearly with respect to the fiber volume fraction. The shear modulus exhibited the largest standard deviations of the three moduli. Considering the distributions of the effective stiffness matrices, the distributions of moduli were as expected.

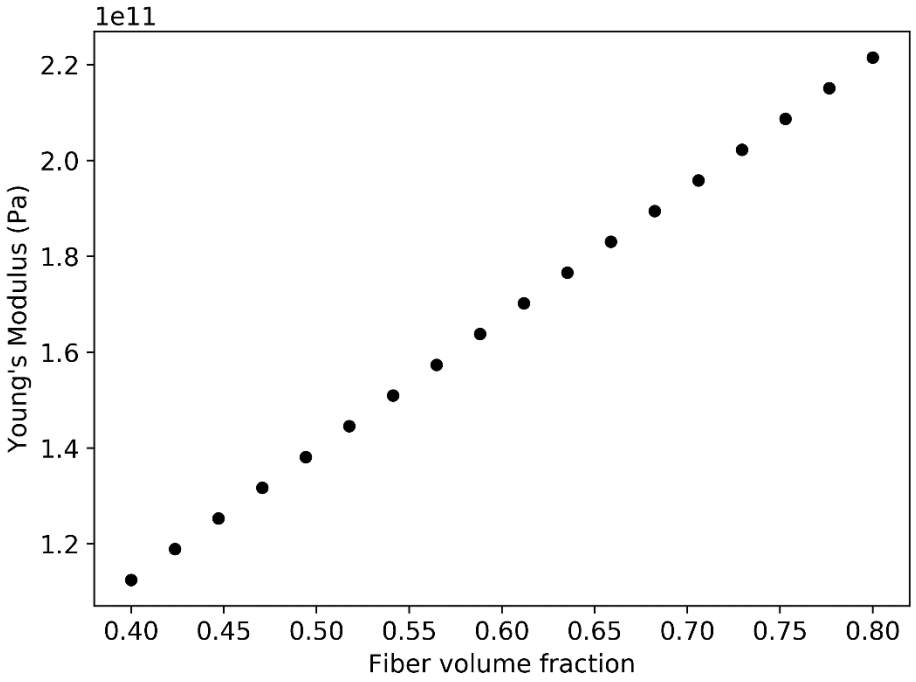

a) Longitudinal Young's modulus, E_{11}

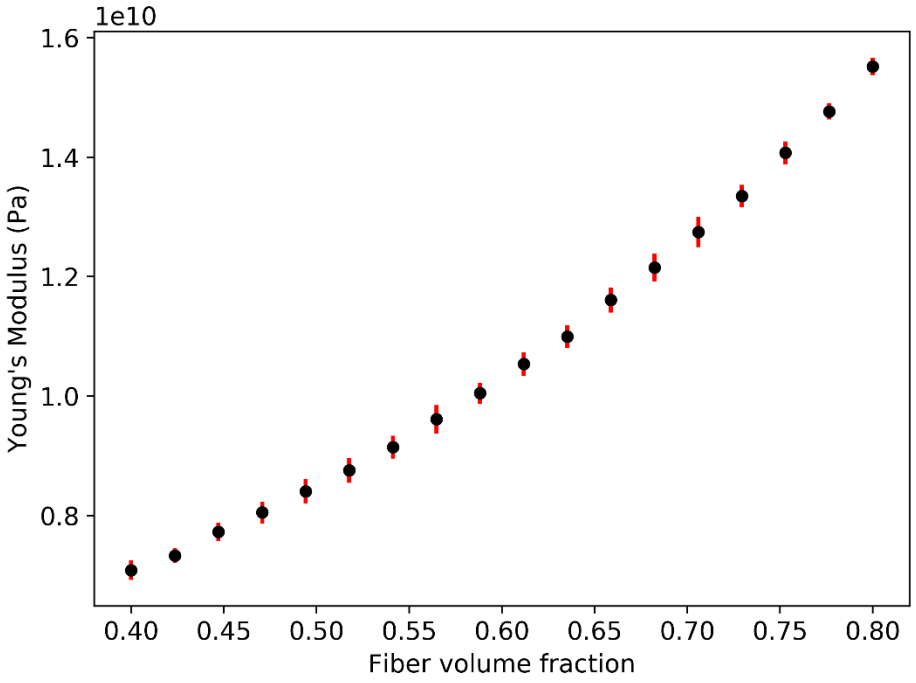

b) Transverse Young's Modulus, $E_{22} = E_{33}$

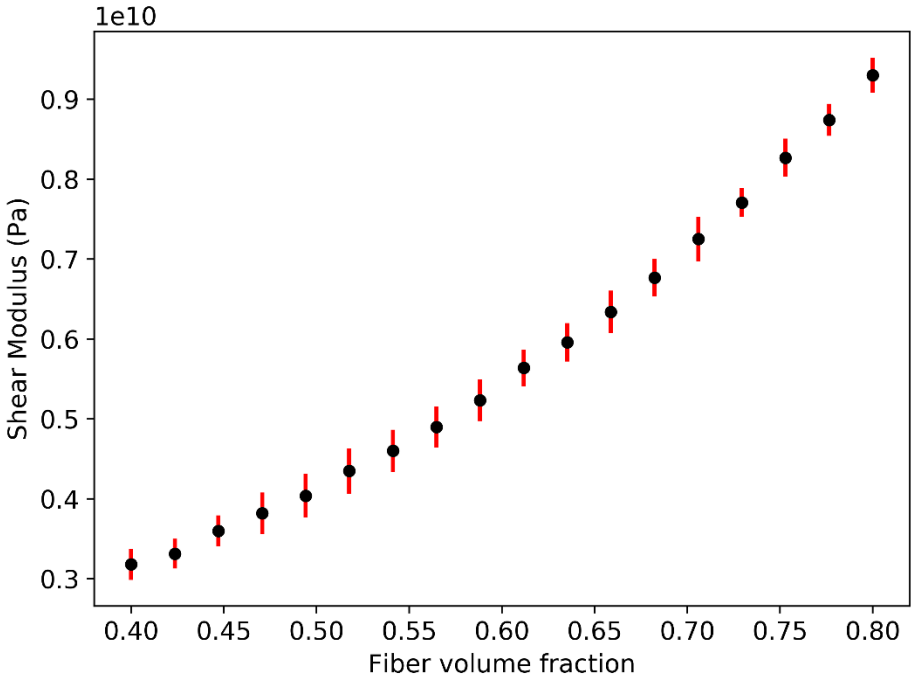

c) Longitudinal shear modulus, $G_{12} = G_{13}$

Figure 4.8. Mean Young's and shear moduli as a function of fiber volume fraction after transverse isotropy has been enforced with errors bars indicating ±1 standard deviation

Figure 4.9 shows the mean Poisson's ratios as a function of fiber volume fraction after transverse isotropy has been enforced with error bars indicating ±1 standard deviation. The major longitudinal Poisson's ratio, ν_{12}, decreased almost linearly with an increasing fiber volume fraction, and the standard deviation decreased with an increasing fiber volume fraction. Though the minor longitudinal Poisson's ratio, ν_{21}, can be computed from ν_{12}, E_{22}, and E_{11}, Figure 4.9b shows the nonlinear dependence of ν_{21} on the fiber volume fraction and large standard deviations. As shown in Figure 4.9c, the in-plane Poisson's ratio ν_{23}, varied nonlinearly as a function of the fiber volume fraction and exhibited large standard deviations within the fiber arrangements.

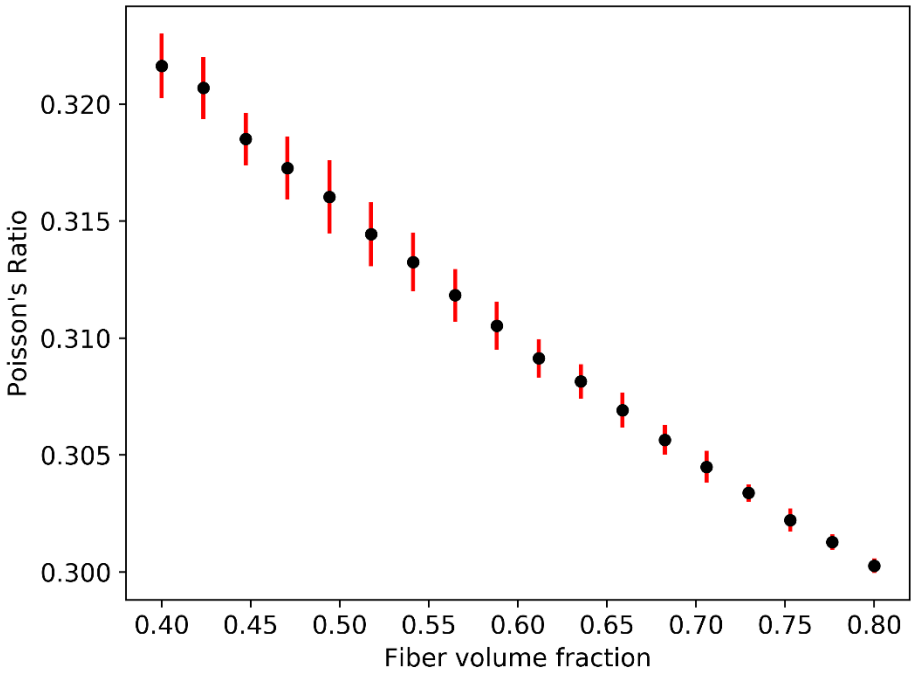

a) Major longitudinal Poisson's ratio $\nu_{12} = \nu_{13}$

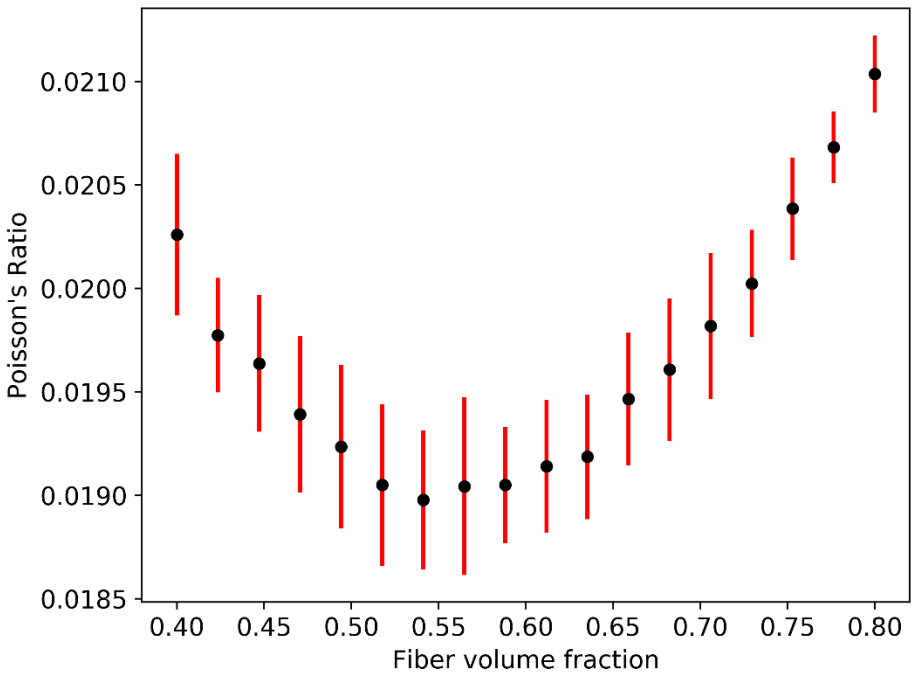

b) Minor longitudinal Poisson's ratio $\nu_{21} = \nu_{31}$

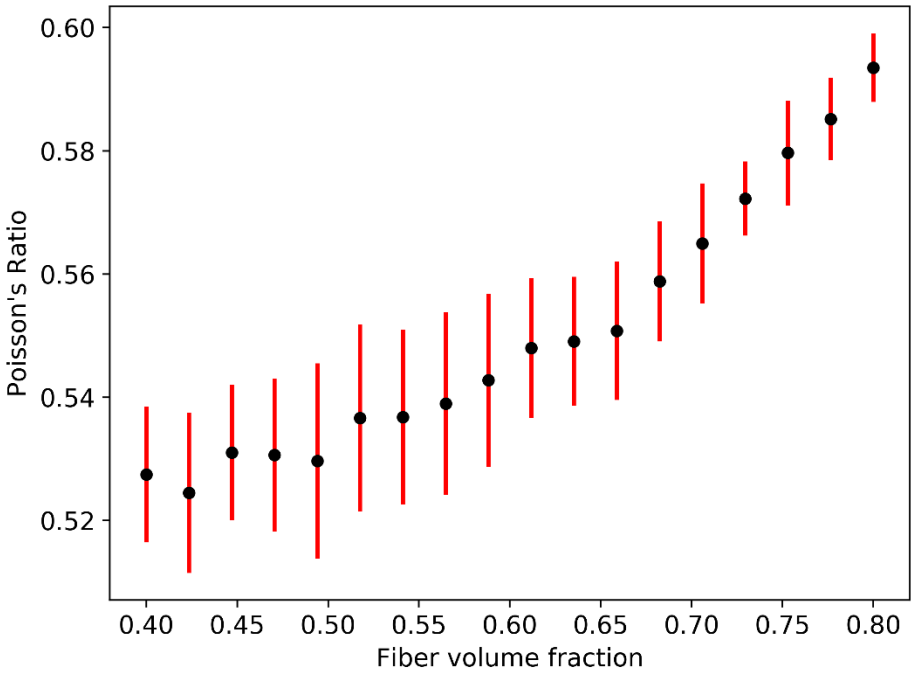

c) In-plane Poisson's ratio $\nu_{23} = \nu_{32}$

Figure 4.9. Mean Poisson's ratios as a function of fiber volume fraction after transverse isotropy has been enforced with error bars indicating ±1 standard deviation

4.4.2 Smoothing

Most of the engineering properties seemed to vary relatively smoothly as a function of the fiber volume fraction, but ν_{23} exhibited an irregular variation. To avoid any irregularities in the properties evaluated at the mesoscale, the nonzero mean effective stiffness terms were fit using a cubic polynomial. At the mesoscale, the effective stiffness matrix for each element in the tows was created using the cubic polynomial fits. Figure 4.10 shows the cubic fit for each nonzero term of the stiffness matrix.

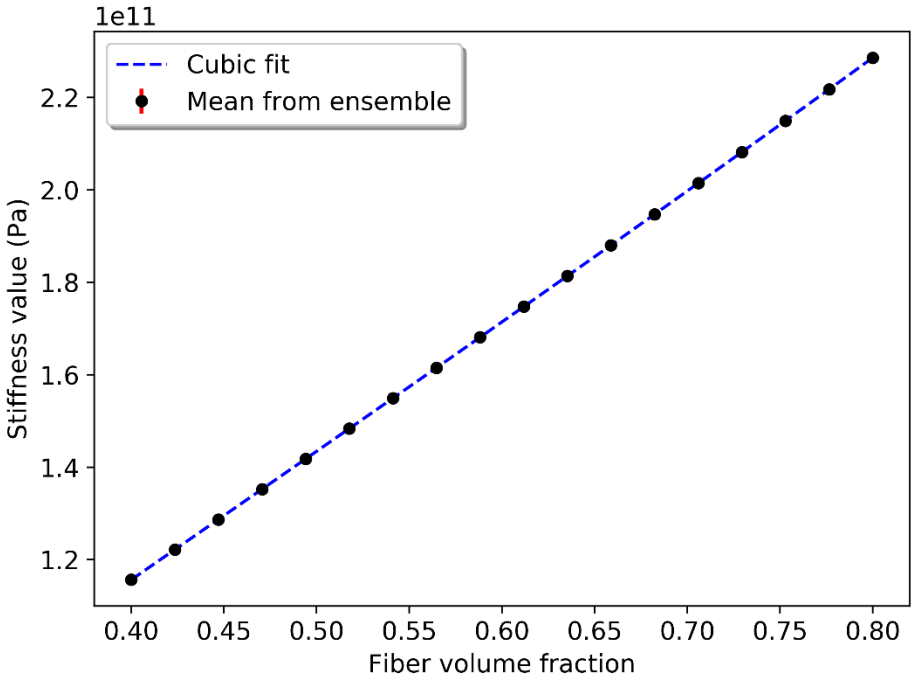

a) C_{11}

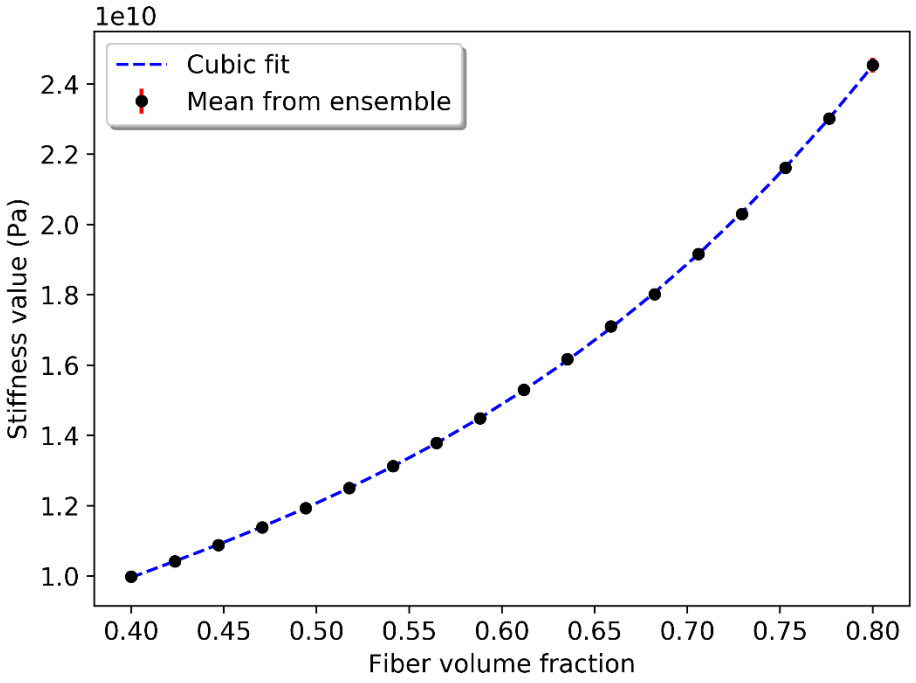

b) $C_{22} = C_{33}$

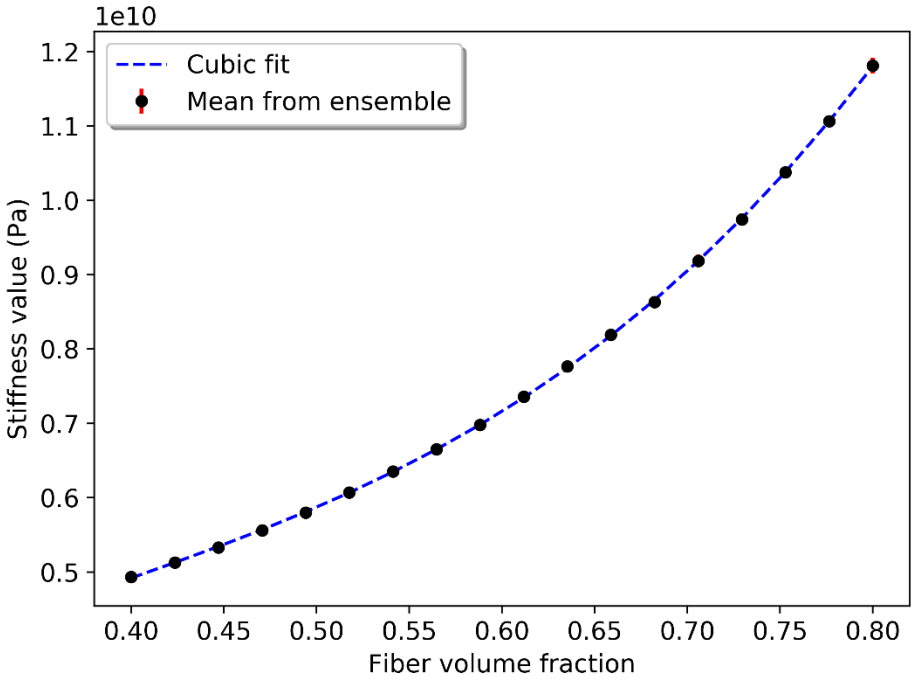

c) $C_{12} = C_{13}$

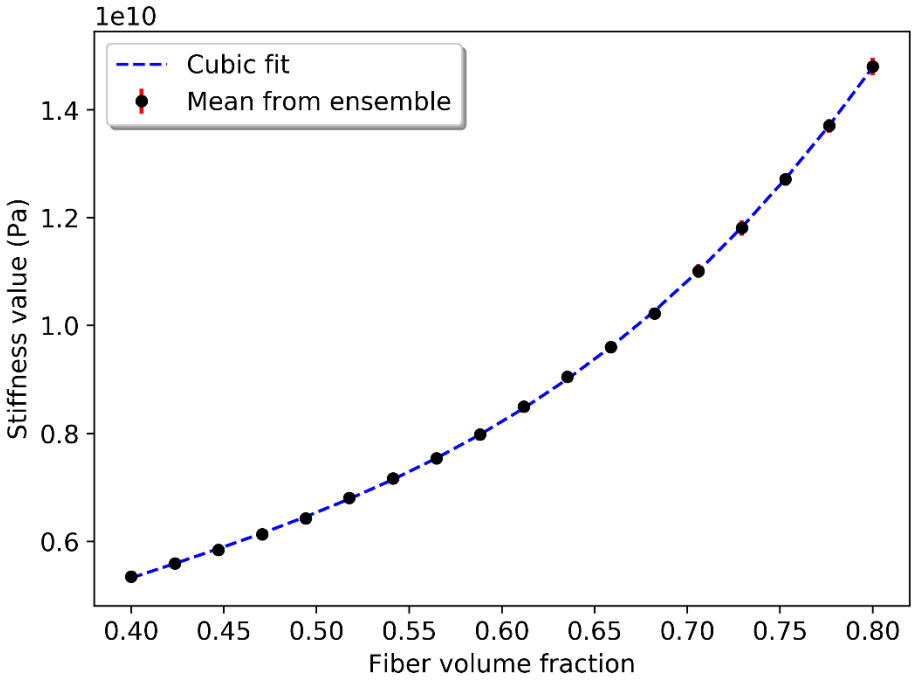

d) C_{23}

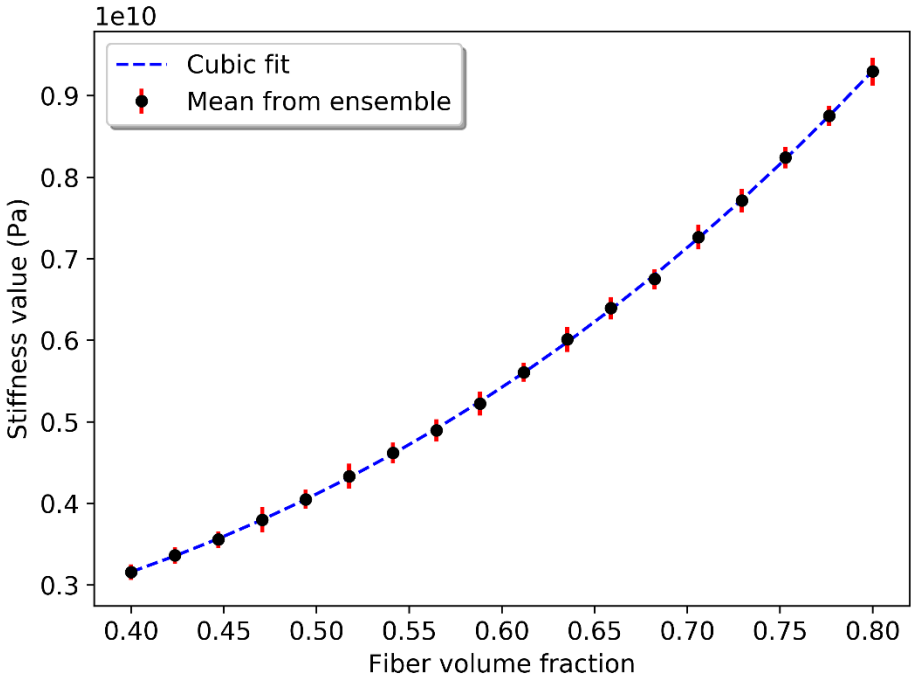

e) $C_{44} = C_{66}$

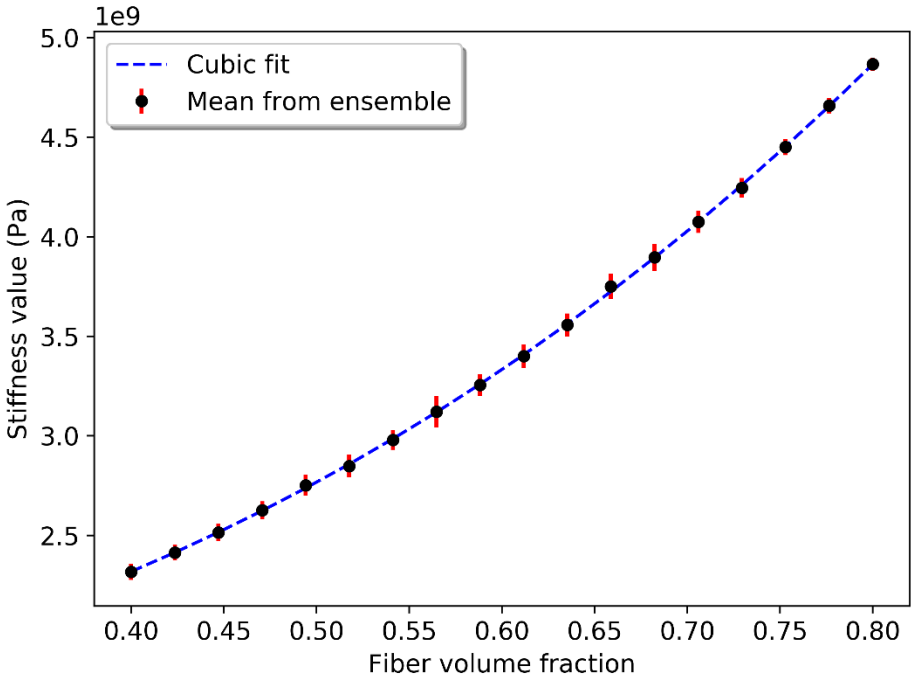

f) C_{55}

Figure 4.10. Cubic fit of nonzero effective stiffness terms

4.5 Summary

Within multiscale analyses, the microscale can be accounted for directly by modeling the heterogeneous microstructure within a mesoscale analysis or by using reduced order models. If local phenomena on the length scale smaller than a microscale RVE are of interest, then mean field homogenization can may introduce significant error, but mean field homogenization techniques have been proven to be very useful. Both strategies are used within this manuscript, but the 3D textiles used exhibited a significant variation of the cross-sectional area along tow paths, which would cause a variation in the local fiber volume fraction. It is computationally prohibitive to model each fiber within the tow directly, and a single ensemble of fiber/matrix RVEs is insufficient to prefect the mesoscale properties if the variation of fiber volume fraction is to be captured. A strategy for capturing the effect of changing local fiber volume fractions was proposed based on spline fits of the mean properties from multiple ensembles of fiber arrangements, each representing a particular fiber volume fraction.

The hexagonally shaped RVEs successfully allowed the reliable generation of fiber arrangements with relatively high fiber volume fractions. The FEA analyses did show that random fiber arrangements can exhibit couplings that do not exist for orthotropic isotropic materials, but the effect is expected to be on the order of 5% based on the standard deviation of the coupling terms compared to the diagonal stiffness terms. Though the mean properties did not exhibit transverse isotropy exactly, terms that would be equal were within 3% of

each other, and terms that would be zero were less than 5% of the diagonal terms. Larger RVEs and more fiber arrangements are expected to tend toward transverse isotropy.

Though this work uses the mean effective properties as a function of fiber volume fraction at the mesoscale, the standard deviation of the effective properties could be used when computing the local stiffness matrix by randomly selecting each stiffness term from a statistical distribution. This strategy would account for the microscale anisotropy at the next scale. Since this work focuses on the elastic response within the tows of the 3D models, the local anisotropy would have relatively little effect on the predicted stress distributions within the textile. However, the local anisotropy might become important within progressive analyses, since damage is such a local phenomenon, and the subject warrants future study.

The use of Monte Carlo simulations of random fiber/matrix RVEs to predict statistically significant properties to be used at the mesoscale is not new. However, this work proposes a method for generating random fiber/matrix RVEs that is valid for a much wider range of volume fractions and uses many ensembles of fiber arrangements to develop a constitutive law that will be used to account for the variation of the local fiber volume fraction in 3D textile models.

5 BRIDGING THE MICROSCALE AND MESOSCALE: REVISITING THE CLASSICAL FREE-EDGE PROBLEM*

5.1 Overview

For many years, researchers have used classical laminate theories, finite difference models, and finite element analysis (FEA) to characterize the effect of the free-edge for various types of layups with great success. [8] [9] [10] However, researchers have almost always modeled the fiber/matrix plies as transversely isotropic, homogeneous materials which neglects the local effects of stiff fibers in a relatively soft matrix. Due to the much smaller scale of an individual fiber compared to that of a ply, the perturbation in the stress field due to the actual heterogeneity will remain localized to a small region, but if the aim of an analysis is to predict when a ply crack or delamination will initiate, then local stress concentrations can be very significant. Some previous studies, have modeled discrete fibers near the free-edge, but relatively few fibers (approximately 10 fibers) were considered due to the prohibitive computational challenges. [13] [14]

With the finite element framework developed or this work and the computational resource available through the Texas A&M University's Supercomputing Center, the accuracy of approximating a ply as a homogenous material will be evaluated by directly modeling the heterogeneous microstructure within a section of a composite laminate. A

---

* Parts of this material originally appeared in the Journal of Composite Materials published by Sage Publications and the Proceedings of the 21st International Conference on Composite Materials. It is reprinted here with permission of the publisher.

much larger section of the microstructure is modeled at the microscale than has been done previously. Even with the increase in computational power, a thin-ply laminated composite will be considered due to a limitation of computational resources.

This chapter aims to investigate the effect of discrete fibers on the stresses along the ply interface close to the free-edge. The case where the heterogenous fiber/matrix microstructure is directly modeled will be compared to the case where each ply is treated as a homogeneous material. Two types of laminated composite layups will be considered: a crossply [0/90]s layup and quasi-isotropic [±45/0/90]s layup.

This chapter begins with a section that describes the composite laminates and the smaller analysis regions considered, the finite element model used for the analysis of the laminated composite, and the material properties used. Next, the results are divided into two section. The first results section investigates the effect of heterogeneity on the microscale stresses near a free-edge, the effect of fiber arrangement, and the interaction between fibers near the ply interface in a crossply [0/90]s laminate. The second results section studies a quasi-isotropic [±45/0/90]s laminate, exploring the effect of the heterogeneous microstructure and evaluating the effect of only considering a smaller microscale region. Finally, a summary of the results is provided, including suggestions for future studies.

5.2 Configurations

The following sections describe the composite laminate and the analysis region, the finite element model (including simplified models for investing the interaction of fibers), and the fiber and matrix properties and the homogenized properties for a lamina.

5.2.1 Composite Laminates and Analysis Region

Two types of laminated composites, [0/90]s and [±45/0/90]s, under uniaxial tension were modeled to understand the interlaminar stresses that develop near the free-edge. The [0/90]s laminate was selected because it is the simplest laminate that exhibits free-edge stresses, while the [±45/0/90]s laminate is often used in the industry. The domain of a typical composite specimen is prohibitively large for directly modeling the fibers and matrix, but a much smaller unit cell can be considered that is representative of the entire specimen. For the coordinate system shown in Figure 5.1, the laminate is assumed to be symmetric about the $y = 0$ and $z = 0$ planes. It should be noted that the [±45/0/90]s laminate is not actually symmetric about the $y = 0$ plane due to the ±45 plies, but the error is believed to be insignificant. Due to these assumed symmetries, one quarter of the specimen can be considered. Even with high-performance computing resources, there are limitations to the size of model. Along the x-axis, a relatively small slice of the laminate was modeled, but this is believed to have an insignificant effect a fiber diameter away from the boundaries normal to the x-axis. Because of these approximations, the analysis region for each type of laminate is reduced to that shown in Figure 5.1.

The boundary conditions applied to the analysis region for each of the two configurations are shown in Figure 5.2. Any boundary conditions not explicitly shown in Figure 5.2 are zero-traction conditions. The applied displacement, U_0, is chosen such that the model experiences 1% volume average strain in the x-direction. The analysis region of both laminates had a full width of about 11.5 ply thicknesses. The laminate [0/90]s had a full height of 4 plies, while the [±45/0/90]s laminate had a full height of 8 plies, but as mentioned earlier, only half of the width and height will be modeled herein. The dimensions [0/90]s analysis region as labelled in Figure 5.1a are as follows: $L = 2.88e\text{-}4\ m$, $W = 5.24e\text{-}5\ m$, and $H = 1e\text{-}4\ m$. The dimensions for the [±45/0/90]s analysis region as labelled in Figure 5.1b are as follows: $L = 4.36e\text{-}4\ m$, $W = 3.05e\text{-}5\ m$, and $H = 3e\text{-}4\ m$. Carbon fibers were assumed to have a diameter of $5\ \mu m$ and volume fraction of 60%. The analysis region chosen results in 550 fibers in the 0^0 ply for both types of laminates.

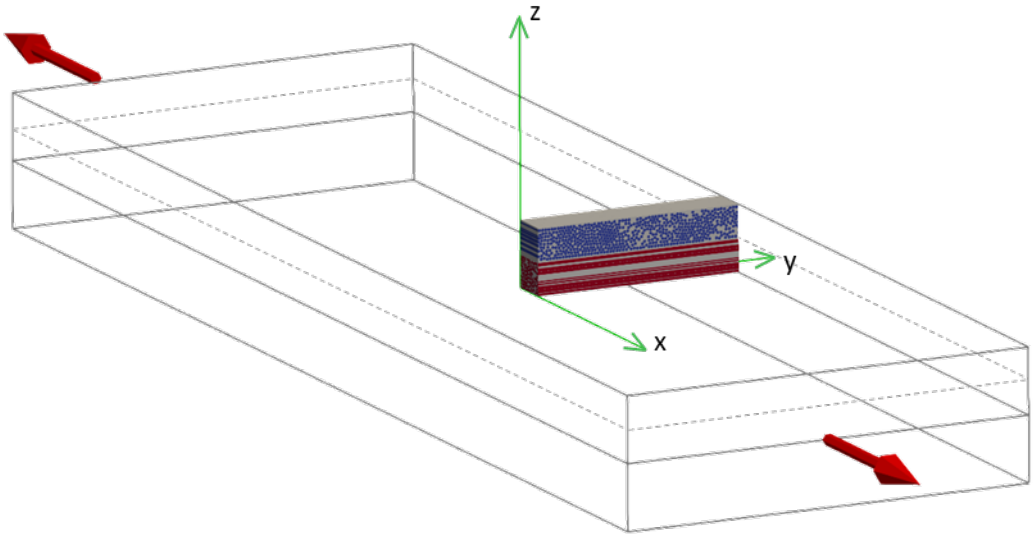

a) [0/90]s

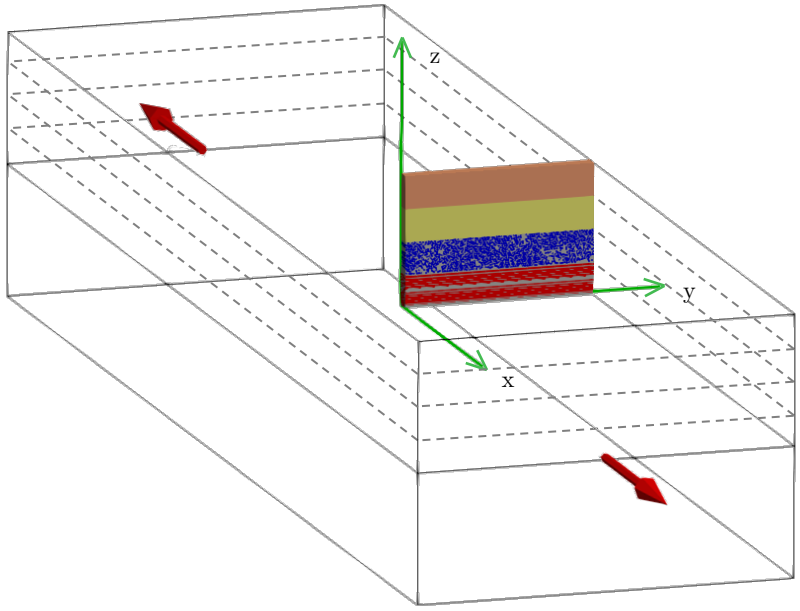

b) [±45/0/90]s

Figure 5.1. Laminated composite under uniaxial tension in direction of red arrows and the analysis region modeled

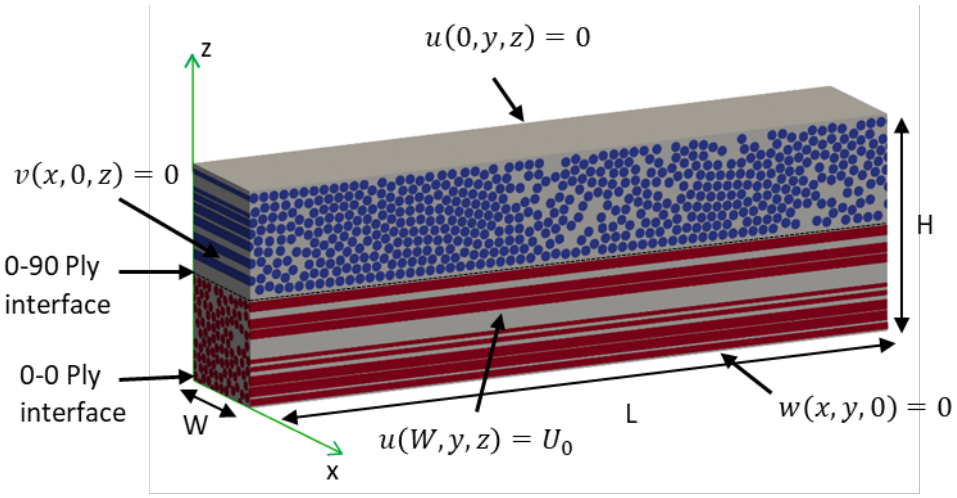

a) [0/90]s

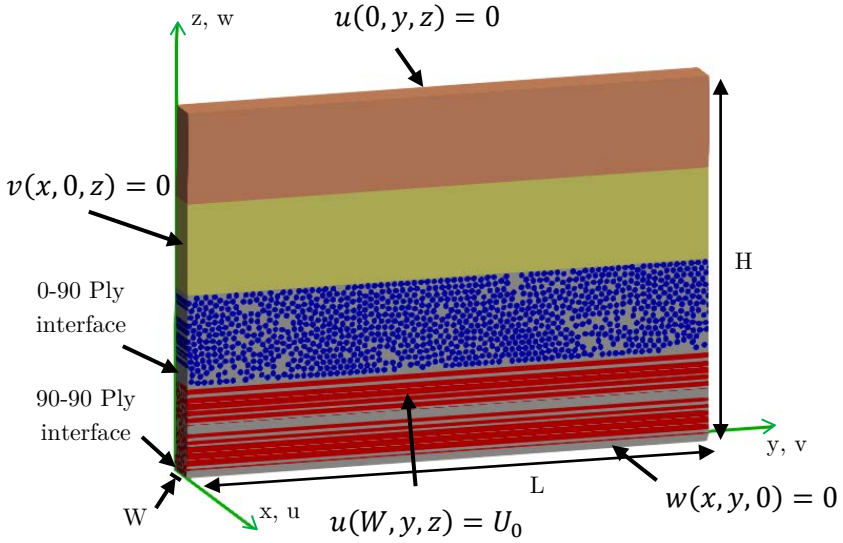

b) [±45/0/90]s

Figure 5.2. Boundary conditions for analysis region shown in Figure 5.1

5.2.2 Finite Element Model

This section starts with a description of the finite element model used for the unidirectional 0^0 and 90^0 plies that are modeled at the microscale. Next, the finite element models used for the [0/90]s and [±45/0/90]s laminates are presented. To allow the investigation of the effect of the microscale region size, a series of finite element models for

the [±45/0/90]s laminate, each with a different microscale region, is described. Finally, a simple two fiber model is given for investigating the interaction of fibers near the interface.

Unidirectional Ply Model with Heterogeneous Microstructure

This work used randomly generated fiber arrangements to represent the cross-sections of each ply. The arrangements of fibers were created by randomly positioning fibers within the cross-sections and iteratively adjusting fiber positions to remove any spatial interference. The algorithm for generating the fiber arrangements is described in an earlier chapter of this manuscript. For this study, a fiber volume fraction, V_f, of 60% was assumed for each ply. For the 0^0 ply, which has fibers aligned with the x-axis, the random arrangement was generated for the y-z cross section, which is shown in Figure 5.1b. For this cross-section, the ply boundaries were assumed to be planar and fibers were not allowed to touch the top and bottom boundaries normal to the z-axis during the generation of the random arrangement. For the 90^0 ply, which has fibers parallel to the y-axis, the random arrangement was generated for the x-z cross-section, while not allowing fibers to touch the top edge normal to the z-axis. However, the ply was assumed to be symmetric about $z = 0$, so fibers could either touch or cross exactly halfway to maintain a symmetry of the geometry. This is not completely realistic, since fibers can physically cross the mid-plane, resulting in no perfect symmetry about the mid-plane, but assuming that the laminate was symmetric about the mid-plane allowed half of the entire laminate thickness to be modeled.

For the [0/90]s laminate, two random fiber arrangements were considered to investigate the effect of fiber arrangements, which are shown in Figure 5.3. Realization A was generated

with fibers being at least 3% of a fiber radius away from another fiber, while the fibers in Realization B were specified to be at least 13% of a fiber radius away from another fiber. Once the geometry was created, a conforming quadratic triangle mesh was generated using a third-party library called Triangle. [38]

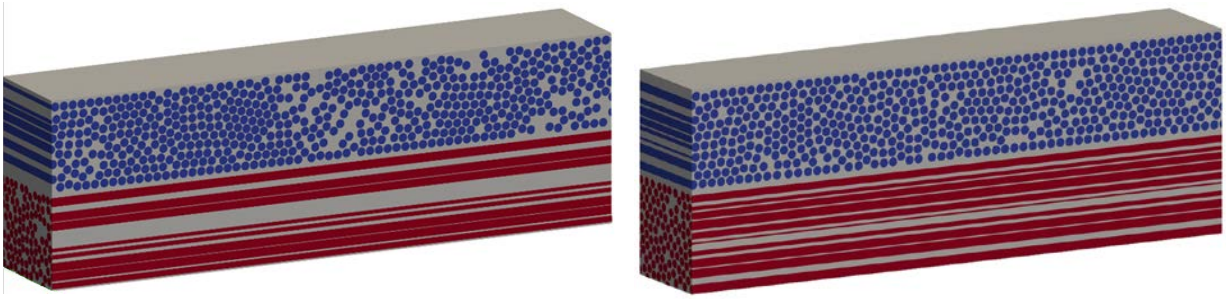

a) Fiber arrangement A

b) Fiber arrangement B

Figure 5.3. Two randomly generated fiber arrangements with different spacing between fibers used to investigate the effect of fiber arrangement

Laminated Composite Model

Although a [±45/0/90]s laminate is considered, only the 0^0 and 90^0 plies are modeled at the microscale since the 0-90 and 90-90 interfaces are the focus of this paper. To create a mesh for the 0^0 and 90^0 plies, the 2D cross-sectional meshes composed of quadratic triangles were extruded in such a way as to create a compatible mesh at the 0-90 ply interface, resulting in a 3D mesh of quadratic wedge elements for the 0^0 and 90^0 plies.

For the [0/90]s laminate, meshes for the two fiber arrangements shown Figure 5.4 was created by extruding the individual ply meshes in such a way to maintain compatibility at the ply interface. It was assumed that the fibers are straight. The mesh for arrangements

A and B contained about 14 million and 10 million nodes respectively and about 5.5 million and 3.8 million elements respectively. To compare the heterogeneous model to a homogeneous model, the free-edge behavior using homogenized laminae was also calculated. The mesh used for the homogeneous case was made of 20-node hexahedral elements and contained about 240,000 nodes. This mesh was selected for convenience in comparing the heterogenous and homogenized model predictions. If computational efficiency had been an issue, a much smaller quasi-3D model would have been used. [61]

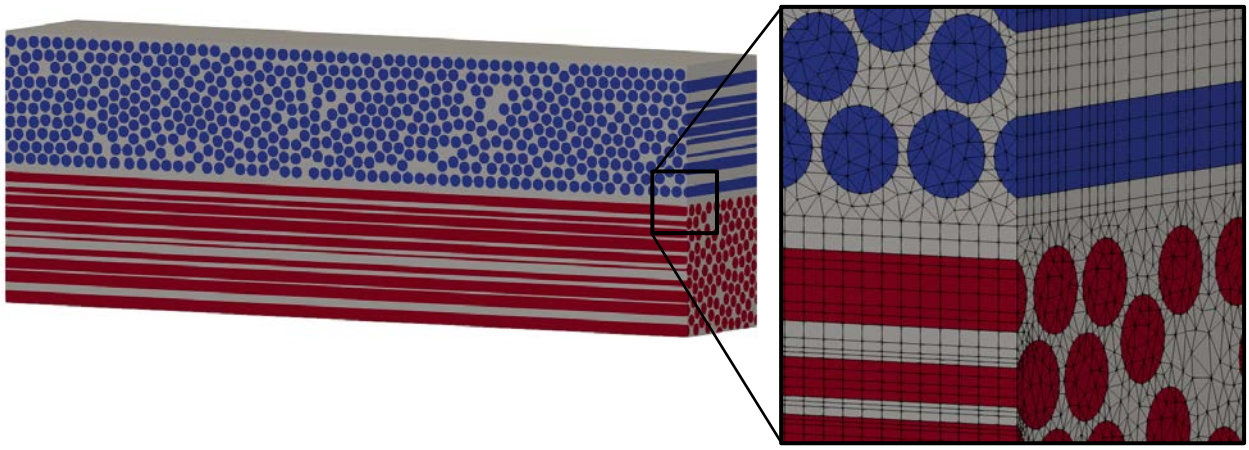

Figure 5.4. Mesh used for a [0/90]s laminate, highlighting the mesh refinement

For the [±45/0/90]s laminate, a 3D mesh was created first for the 0^0 and 90^0 plies, followed by creating meshes for the $+45^0$ and -45^0 plies. The same technique was used to to create the meshes for the 0^0 and 90^0 plies. For the $+45^0$ ply, which is the outermost ply of the laminate, a structured mesh of quadratic hexahedron elements was used. The number

of elements in the x, y, and z directions were 10, 50, and 6 respectively. The mesh can be coarse since the $+45^0$ ply is far away from the ply interfaces of interest and is included primarily to create a high Poisson's ratio. A mesh for the -45^0 ply must transition from the very fine mesh of the 90^0 ply to the coarse refinement of the $+45^0$ ply. Consequently, a combination of quadratic pyramid and tetrahedral elements were used to create a mesh to transition between the disparate mesh refinements of the neighboring plies. The mesh was generated using an in-house mesh generation software and TetGen. A more detailed description of how the models were created is provided earlier in the manuscript in chapter 2. The total number of elements within each ply is shown in Table 5.1. The mesh for the laminate is shown in Figure 5.5. To illustrate the interior refinement of the tetrahedral mesh, half of the $\pm 45^0$ plies are cut away in the figure. Additionally, the mesh refinement for the 0^0 and 90^0 plies are shown in a zoomed in view for a small region for clarity.

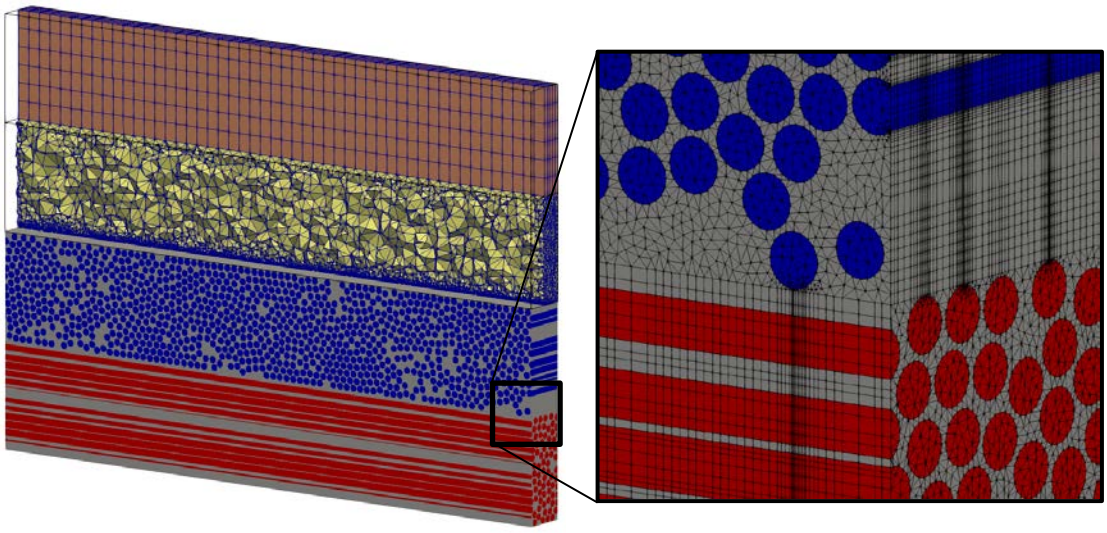

Figure 5.5. Mesh used for a [±45/0/90]s laminate, highlighting the mesh refinement

Table 5.1. Number of elements within each ply of the [±45/0/90]s model

Ply	# of Elements
$+45^0$	3,000
-45^0	1,076,784
0^0	15,141,140
90^0	6,539,604

Models Used for the Investigation of the Effect of the Microscale Region Size

For the [±45/0/90]s laminate, several cases will be considered with varying sizes of microscale regions for the 0^0 and 90^0 plies to provide insight into what size of a region should be modeled at the microscale to accurately predict interlaminar stresses near a free-edge. In future studies, much coarser meshes should be used for areas outside the microscale region. However, for convenience, the same mesh will be used in this study, and the material

properties of the elements will be changed from the fibers and matrix properties to the set of homogeneous properties used to model the plies at the mesoscale. To compare a model where fibers are modeled discretely to the classical homogeneous case, the same mesh described in the previous section was also used for the classical homogeneous case. The mesh is far more refined than necessary to accurately predict stresses for the classical homogeneous case, but using the same mesh is convenient for comparing models. For this study, the microscale region of 0^0 and 90^0 plies was reduced in the z- and y-directions, leaving the region near the 0-90 interface and free edge as discrete fibers and matrix. First, a microscale region that extends 25% and 10% of the ply thickness from the 0-90 ply interface will be considered. The two cases are shown in Figure 5.6. Next, the model with a microscale region of 10% of the ply-thickness in the z-direction, shown in Figure 5.6, was used as the starting point, and the microscale region is reduced in the y-direction, leaving one- and two-ply thicknesses from the free edge modeled as discrete fibers and matrix. Both models are shown in Figure 5.7.

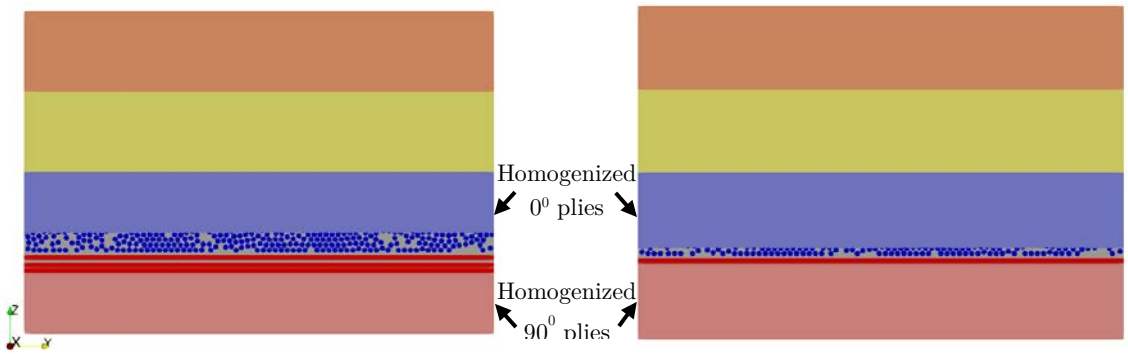

Figure 5.6. Model of a [±45/0/90]s laminate with 25% (left) and 10% (right) of the 0^0 and 90^0 plies modeled at the microscale

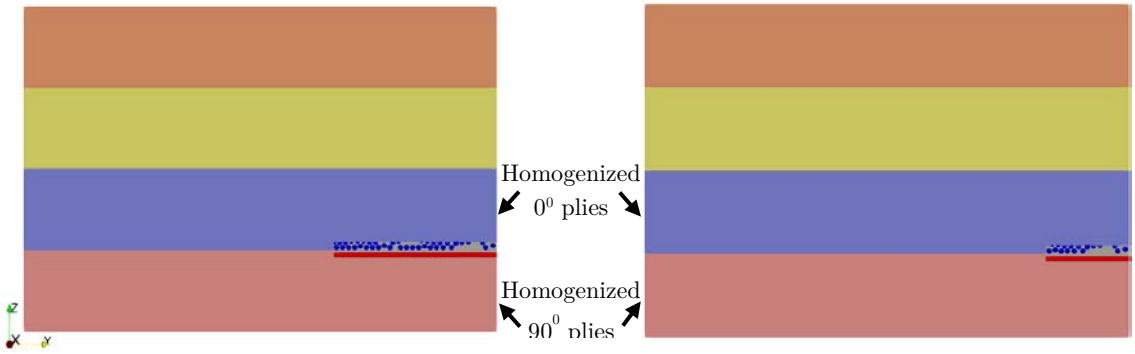

Figure 5.7. Model of a [±45/0/90]s laminate with only two (left) and one (right) ply thicknesses from the free-edge of the 0^0 and 90^0 plies modeled at the microscale

Simplified Models for Investigation of Fiber Interaction

It is known from 2D analyses of fiber/matrix composites that fibers in close proximity can interact to create stress concentrations. [58] It is expected that along the ply interfaces of a laminated composite, a similar phenomenon can occur, but it is difficult to investigate the interaction of fibers with a very large finite element model like the one described in the previous section. Consequently, two simplified models will be used. First, a periodic RVE of a single fiber embedded in matrix, which is shown in Figure 5.8a and results in a square fiber arrangement, will be used study the interaction of fibers in a unidirectional laminate. Second, a periodic RVE of two crossing fibers within matrix, which is shown in Figure 5.8b, will be used to study the interaction of fibers in a situation closer to a cross-ply laminate.

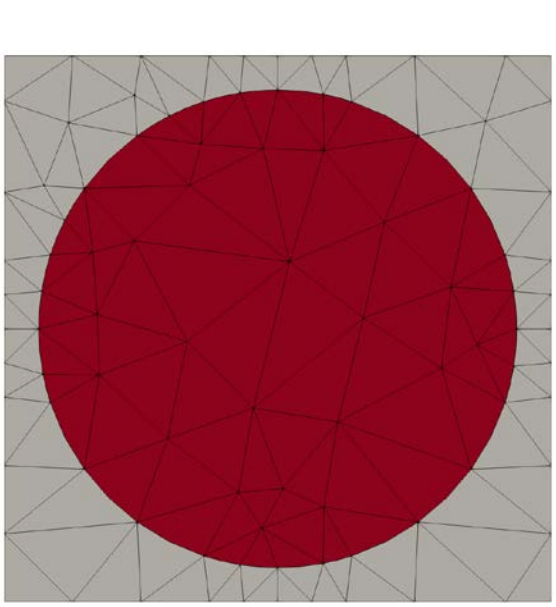

a) Periodic 2D RVE for unidirectional lamina

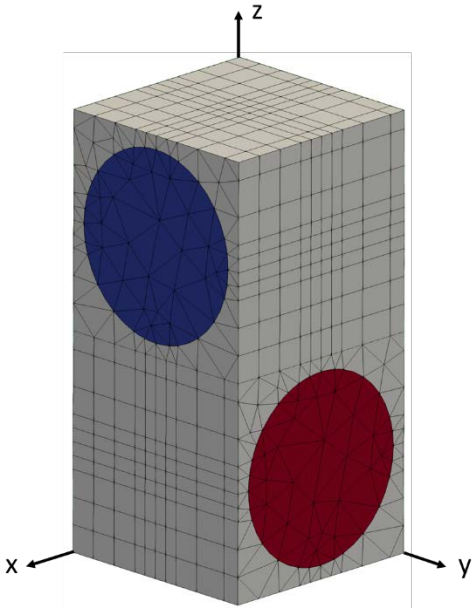

a) Periodic 3D RVE with crossing fibers

Figure 5.8. Simplified models for studying interaction between fibers

The simplified periodic fiber/matrix RVEs were generated to be geometrically periodic, and the mesh was generated to place nodes along the boundaries periodically, so it is possible to exactly impose periodic boundary conditions. Periodic boundary conditions for the unit cell are concisely expressed in terms of the displacements along the unit cell boundaries in Equation (5.1) using indicial notation. Repeated indices imply summation, and d_k is a vector of periodicity that starts at some boundary coordinate x_j in one unit-cell and ends at the equivalent point the next unit cell. A more detailed description of periodic boundary conditions can be found in Refs. [58] and [36].

$$u_i(x_j + d_j) = u_i(x_j) + \langle \frac{\partial u_i}{\partial x_k} \rangle d_k \tag{5.1}$$

5.2.3 Material Properties

Two approximations were used for the unidirectional laminae in these studies. For the model using homogeneous orthotropic properties for the plies, properties were taken from Ref. [60]. In the microscale region, fibers and matrix were modeled discretely and assumed to be IM7 graphite fibers and 5220-4 epoxy matrix respectively. The properties for the epoxy matrix was taken from Ref. [60]. However, it was shown in Ref. [36] that the fiber properties need to be inversely determined if the effective properties are to accurately represent the homogenized microstructure. In Ref. [36], the apparent fiber properties were inversely determined using the effective ply properties and matrix properties in Ref. [60]. Those same apparent fiber properties determined in Ref. [36] were used in this paper, so that homogenizing unidirectional laminae modeled as discrete fibers and matrix results in the effective properties in Ref. [60]. For all cases, the materials were assumed to remain linearly elastic. Table 5.2 summaries the properties used for the fibers, matrix, and homogenized plies.

Table 5.2. Material properties for free-edge study

IM7 Graphite Fibers [36]	
E_1^f (GPa)	276
E_2^f, E_3^f (GPa)	26.0
G_{12}^f, G_{13}^f (GPa)	20.7
G_{23}^f (GPa)	7.55
ν_{12}^f, ν_{13}^f	0.292

5220-4 Epoxy Matrix [60]	
E^m (GPa)	3.45
G^m (GPa)	1.278
ν^m	0.35

Effective IM7/5220-4 Lamina [60]	
E_1 (GPa)	165.5
E_2, E_3 (GPa)	10.34
G_{12}, G_{12} (GPa)	5.792
G_{23} (GPa)	3.315
ν_{12}, ν_{13}	0.31
ν_{23}	0.56

5.3 Free-Edge Analysis of a [0/90]s Crossply Laminate

This section investigates the effect of modeling the heterogeneous microstructure for a cross-ply laminate subjected to uniaxial tension in the 0^0 direction. This section begins with a comparison of the deformed cross-section shapes for the heterogeneous and homogeneous models. Next, the interlaminar normal stress distributions are compared, followed by a comparison of the interlaminar shear stress distributions. The effects of the discrete fibers and the free-edge are separated, providing an understanding of each and their contribution to the overall behavior. Next, the effect of fiber arrangement is investigated by comparing the interlaminar normal stresses for two different fiber arrangements. Finally, the section

ends with simplified studies on the effect of the matrix Poisson's ratio and the interaction between neighboring fibers with the aim to help explain the behavior observed in the other sections.

5.3.1 Comparison of Deformed Cross-Sections

Figure 5.9a shows the deformed y-z cross-section at $x = 0$ for the heterogeneous model, with an outline showing the deformed boundaries of the homogeneous model. The deformed cross-sections in Figure 5.9 were exaggerated by 43x to clearly illustrate the differences, since the applied 1% strain results in very small displacements. The deformed shape was quite similar for the two models. The largest displacement in the y-direction for the homogeneous model was 11% less than that of the heterogeneous model and occurred at the top left corner of Figure 5.9a. This means the effective properties used for the homogeneous model are not exactly representative of the heterogeneous microstructure. This is not surprising, since the fiber and matrix properties were chosen such that, for a periodic quasi-3D RVE, the effective properties based on the volume average stress-strain curves matched the lamina properties that were experimentally measured in Ref. [60]. However, each ply considered in this paper has a finite thickness. As a result, the 0^0 ply is only periodic along the x-axis, and the effective properties may not be as representative of this finite case as they are representative of the infinite periodic array considered in Ref. [36]. This discrepancy concerning using effective properties based on an infinite domain for an analysis of a finite domain contributes to the difference between the deformed cross-sections in Figure 5.9a. However, there is also the issue of the fiber volume fraction not being uniform

within the 0^0 ply, which can be observed in Figure 5.9a. The non-uniform fiber volume fraction also contributes to difference between the two displacement fields, but it is difficult to separate these two effects.

At the microscale, there were some notable differences between the two deformed shapes, but these differences were due to the perturbations caused by the discrete modeling of the heterogeneous microstructure. Along the top edge of Figure 5.9a, some of these perturbations can be observed, which are highlighted in Figure 5.9b.

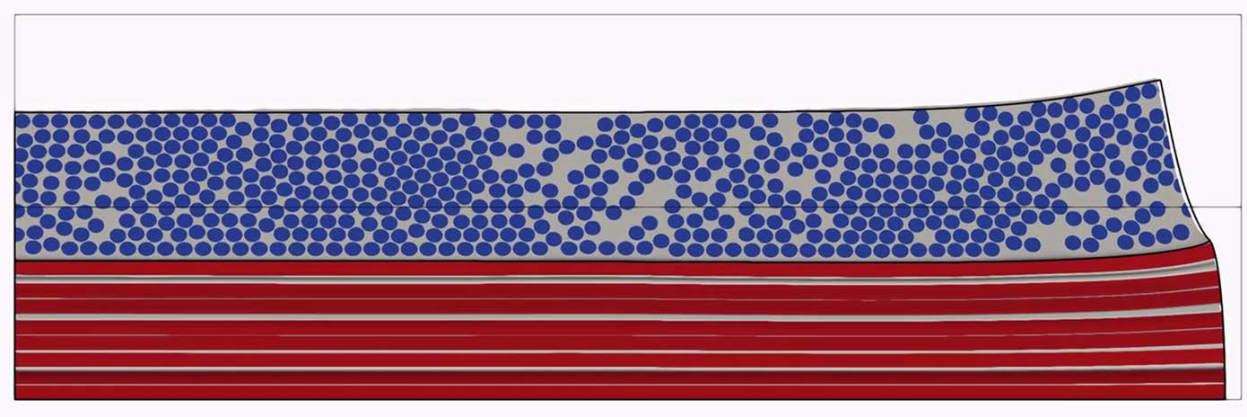

a) Deformed discrete fiber/matrix case with outline showing deformed homogenized case (43x displacement magnification)

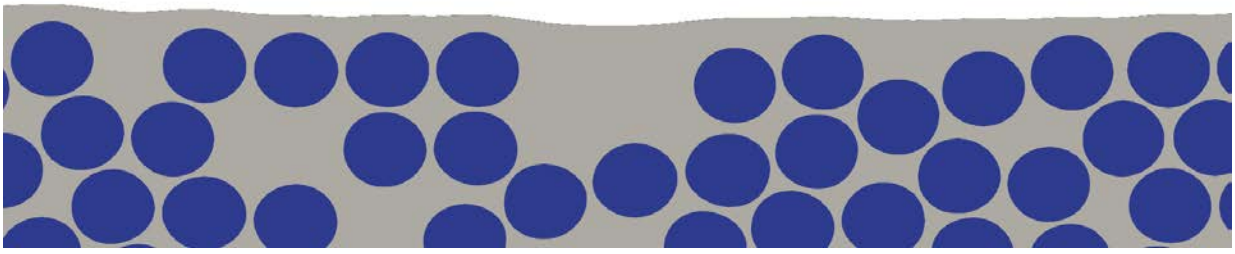

b) Deformed section along top edge of figure above showing local perturbations

Figure 5.9. Deformed y-z cross-section at x = 0

5.3.2 Interlaminar Normal Stress

The effect of the free-edge on the interlaminar stresses has been extensively studied using homogenized plies, but this paper directly models the fiber/matrix microstructure models within the plies. For the homogeneous model under uniaxial extension, σ_{zz} is the only significant interlaminar stress along the 90-90 ply interface, while σ_{zz} and σ_{yz} are both significant along the 0-90 ply interface. However, where fibers and matrix are modeled, σ_{xz} is not necessarily small.

Figure 5.10 shows contours for σ_{zz} along the 0-90 and 90-90 ply interfaces. To showcase the differences, a small region of the interface near the free-edge is shaded as a solid color, displaced in the z-direction in proportion to the magnitude of σ_{zz}, and compared to the case where plies are modeled using a homogeneous material, which is shown as a line. The clearest difference is along the 0-90 interface near the free-edge, where the homogeneous model predicts a very large normal stress due to a singularity, while the heterogeneous model predicts a much lower stress. In the homogeneous model, a combination of the free-edge and material discontinuity at the ply interface creates a singularity, where $\sigma_{zz}\left(x, y, \frac{H}{2}\right) \rightarrow \infty$ as $y \rightarrow L$. However, σ_{zz} cannot reach infinity in the finite element model and reaches a finite maximum value that is dependent on the mesh refinement. In the heterogeneous model, there is no such singularity at the ply interface, since there is only matrix material at the intersection of the free-edge and ply interface. Furthermore, the maximum stress along the ply interface in the heterogeneous model does not occur at the free-edge, but rather slightly away from it, where there is a matrix pocket.

There are two mechanisms contributing to the large difference of stress distributions for the two models. First, there is a local effect of the heterogeneous microstructure. The fibers near the ply interface and the large mismatch between fiber and matrix properties create a non-zero σ_{zz} even far away from the free-edge. The effect of fibers on the interlaminar normal stress can be seen along the entire ply interface in Figure 5.10. By contrast, the homogeneous model predicts that $\sigma_{zz}\left(x, y, \frac{H}{2}\right)$ goes to zero far away from the free-edge. Second, there is the effect of the free-edge that occurs at a larger scale than individual fibers. The free-edge still affects the interlaminar stresses, even if it no longer introduces a singularity at the ply interface for the heterogeneous model. It is beneficial to separate these two factors to understand the importance of each, which is done in a later section.

Comparing Figure 5.10a and b shows that when modeling the heterogeneous microstructure, σ_{zz} varies more sharply and reaches greater extrema for the 90-90 interface than the 0-90 interface. However, this is the opposite relationship that is predicted by the homogenous model, which predicts a much higher stress along the 0-90 interface due to the singularity. The reason for the heterogeneous model predicting a higher stress at along the 90-90 interface can be attributed to how fibers and matrix interact. Along the 0-90 interface, fibers are not allowed to cross the ply boundary in the method used to generate the random fiber arrangement. The fibers from each ply still interact with each other to create a complex stress distribution, but due to the distance between fibers, the interaction is weaker. Along the 90-90 interface, fibers can cross the boundary in the method used in

this paper, resulting in fibers that are very close to each other on the interface and a greater range of stresses. Additionally, there is a difference at the free-edge. When fibers and matrix are modeled discretely, there will exist a singularity at the intersection of the free-edge and fiber/matrix interfaces. Along the 90-90 ply interface, fiber/matrix interfaces intersect the free-edge, while fiber/matrix interfaces are further away from the 0-90 ply interface.

Qualitatively, Figure 5.10 shows that the normal interlaminar stresses are very different for the two models, but it is difficult to compare the plots quantitatively. To look at the differences more carefully, the interlaminar normal stress is plotted from the free-edge to the mid-plane along four paths for the heterogenous model. For the homogeneous model, there is no variation along the x-direction. For this discussion, the comparison will be limited to the 0-90 ply interface.

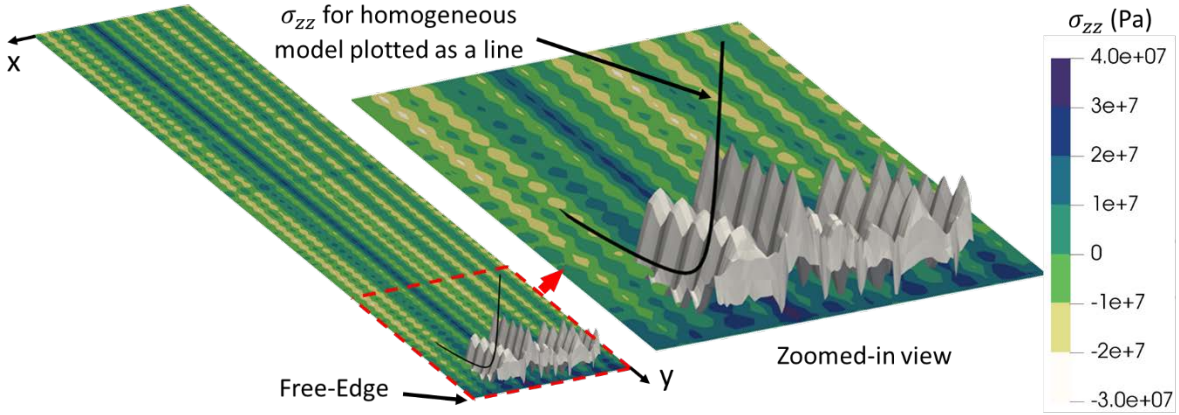

a) 0-90 ply interface

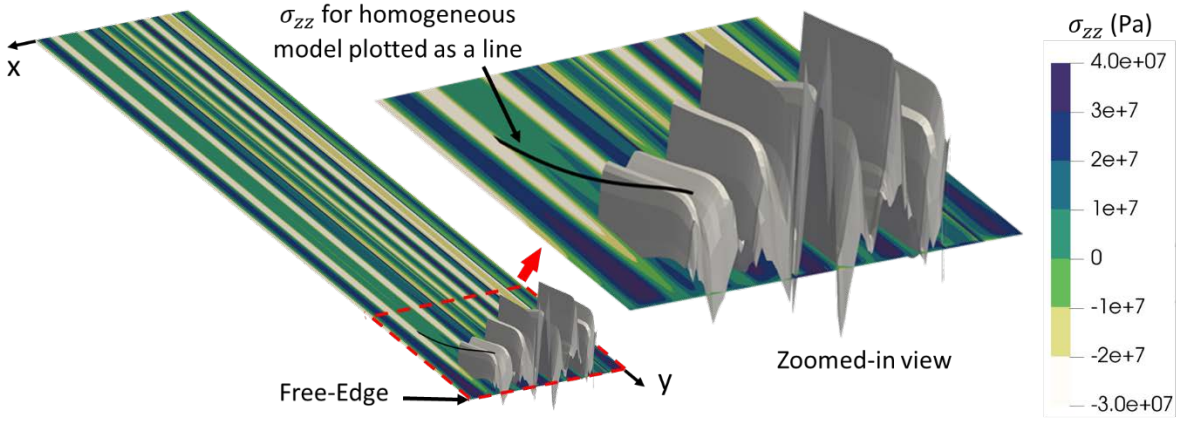

b) 90-90 ply interface

Figure 5.10. σ_{zz} contours for the ply interfaces (a region that is half a ply thickness from the free-edge and two fiber diameters from the edge of the model in the x-direction is deformed in proportion to σ_{zz})

Figure 5.11 shows the interlaminar normal stress as a function of the distance from the free-edge normalized by the ply thickness, which will be denoted d. Note that d is a unitless quantity. Four evenly spaced paths along the interface were considered parallel to the y-axis, which are illustrated in Figure 5.11. The normalized distance from the free-edge is

calculated by the equation below, where L and H are dimensions of the model shown in Figure 5.1b.

$$d = \frac{2(L - y)}{H} \tag{1}$$

The homogeneous model exhibits the singularity at the free-edge as mentioned earlier. As the distance from the free-edge increases, there is a sharp decrease in σ_{zz} for about 1 ply thickness from the free-edge and then it increases to approach 0. For the heterogeneous model, the stress along all four paths considered generally tends to decrease until about 1 ply thickness from the free-edge, but there are also significant oscillations and σ_{zz} drops significantly at the free-edge. Based on the results in Figure 5.11, the paths can be grouped into two categories. The two paths nearest to a 90^0 fiber (1 and 2) had similar stresses, and the two paths closer to matrix pockets (3 and 4) had similar stress distributions. However, there is a large difference in σ_{zz} between these two categories. An important difference between the heterogeneous and homogenized models occurs right at the free-edge. The homogeneous model exhibits a singularity, while the discrete case decreases locally to a finite value, since there is no material discontinuity at the ply interface.

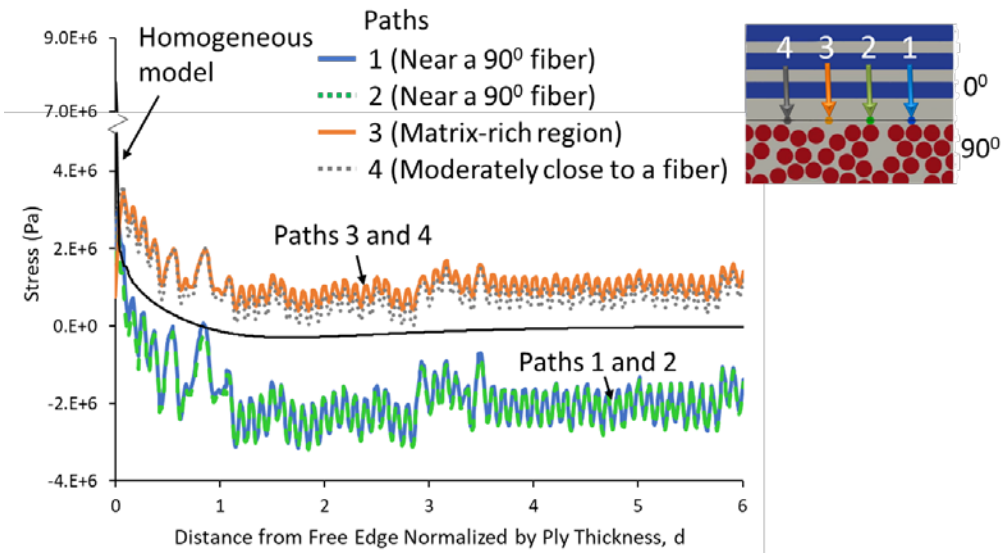

Figure 5.11. 0-90 interlaminar normal stress plotted along lines parallel to the y-axis versus the distance from the free-edge normalized by the ply thickness

5.3.3 Interlaminar Shear Stress

A comparison of the interlaminar shear stress, σ_{yz}, along the 0-90 ply interface for the heterogeneous and homogeneous models is shown in Figure 5.12. Compared to the predicted interlaminar normal stress shown in the previous section, both models predicted a more similar distribution of interlaminar shear stress. The shear stress becomes more severe within a fiber diameter of the free edge and then trends towards zero away from the free-edge. Like σ_{zz}, the heterogeneous model predicts an oscillation of σ_{yz} due to the interaction of fibers near the ply interface, but the smaller oscillations indicate that the 0^0 fibers had a smaller effect on σ_{yz} than σ_{zz}. Additionally, the distributions of σ_{yz} along all four of the paths considered were more similar than for σ_{zz}, indicating that the 90^0 fibers have less of an effect on σ_{yz} than σ_{zz}. For σ_{zz}, the stress became more compressive when fibers came

in proximity to each other. However, Figure 5.12b shows that σ_{yz} becomes less severe directly above a 0^0 fiber, but shear stress concentrations form near the sides of 0^0 fibers, shown by the peaks before and after each 0^0 fiber near the interface. Though the heterogenous microstructure does affect the predicted interlaminar shear stress, the effect is much less than observed for the interlaminar normal stress.

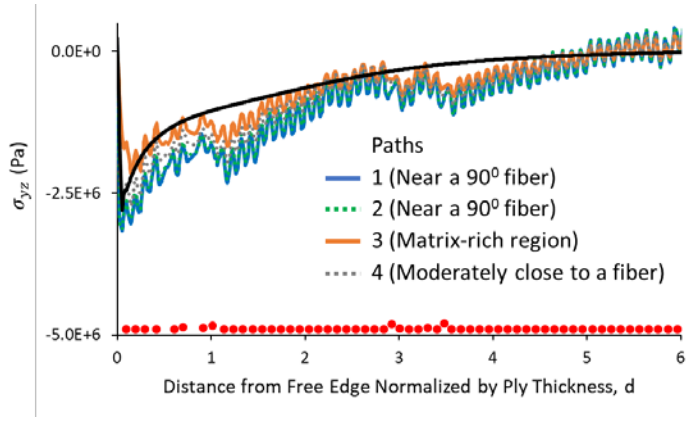

a) Predicted σ_{yz} for the entire half-width of the laminate

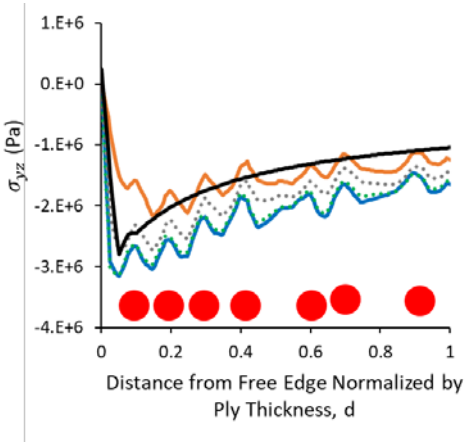

b) Predicted σ_{yz} within one ply-thickness of the free-edge

Figure 5.12. 0-90 interlaminar shear stress plotted along lines parallel to the y-axis versus the distance from the free-edge normalized by the ply thickness (positions of 0^0 fibers near the ply interface are shown using circles)

5.3.4 Separation of the Effect of the Free-Edge and Heterogeneity

As mentioned before, there are two effects contributing to the complex stress distributions for the heterogeneous model. First, the presence of free-edge will perturb the stresses in the material close to it. Second, the interactions of the 0^0 and 90^0 fibers near the

ply interface will cause perturbations of the stresses, even if the material is away from the free-edge. To isolate the effect of the free-edge, superposition will be used. The effect of the free-edge will be the response shown in Figures 5 and 6 minus the response of the material if it was away from the free edge, which will be referred to as the interior response. To predict the interior response, a different set of boundary conditions are applied to the model to make the material near $y = L$ be far from a free-edge. Instead of the $y = L$ plane being the free-edge, the $y = 0$ will be free, and the $y = L$ face will become a plane of symmetry. The stresses within three ply thicknesses of the $y = L$ will be assumed to be minimally affected by the free-edge, leaving the effect of the discrete fibers.

Figure 7 shows the normal stresses for the interior response. Like before, paths near a 90^0 fiber, namely 1 and 2, exhibited very similar stress distributions that were compressive in contrast to the tensile values near matrix-rich regions, namely 3 and 4. Furthermore, the troughs of the stress oscillations closely align with 0^0 fiber positions, which indicates that when fibers in both plies lie close to each other, a compressive stress state is formed. This is a similar trend to one observed in 2D fiber/matrix analyses of random RVEs. McLendon et. al. observed that compressive stresses form when fibers come close together in a uniaxial lamina under tension. [58] Similarly, a compressive effect is caused by fibers being close on each side of the ply interface, when the laminate is subjected to a tensile load. This behavior is important to consider if the strength of laminates near critical areas, such as a free-edge, is to be optimized.

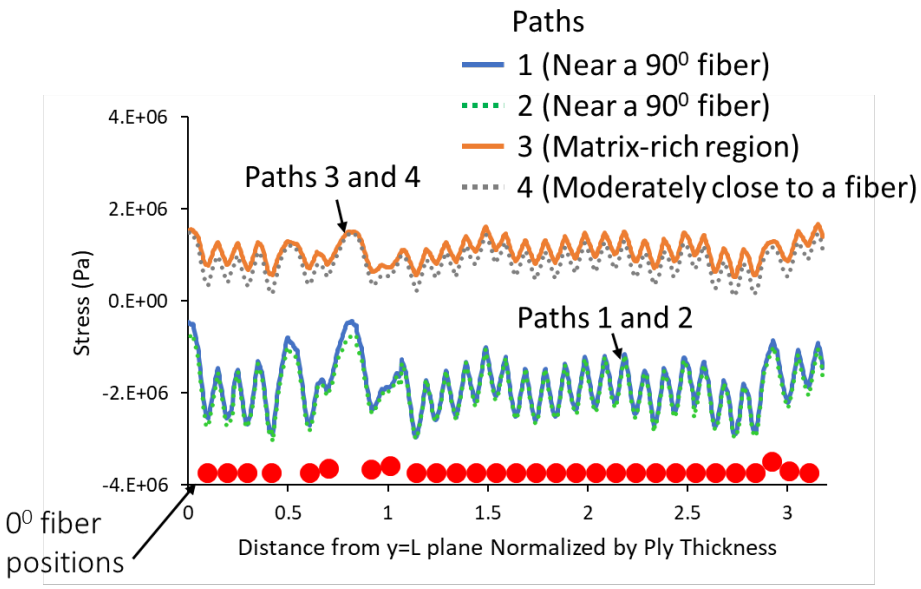

Figure 5.7. σ_{zz} for interior effect using heterogeneous model

Since linear elasticity is assumed in this paper, superposition can be used to isolate the effect of the free-edge by subtracting the response in interior response in Figure 7 from the response in Figure 6. Figure 8 shows the isolated free-edge effect along the 0-90 interface compared to the response of the homogeneous model, as well as 0^0 fiber positions. Figure 8 may be viewed as a fairer comparison of the predicted free-edge effects. With the free-edge effect isolated, the stress distribution for the two models match much better. There is still a significant difference very close to the free-edge due to the stress singularity in the homogenous model, as shown in Figure 8b, but the stress distributions are quite similar about 10% of a ply thickness away from the free-edge and further. There are still oscillations in σ_{zz} for the heterogeneous model indicating that the 0^0 fibers near the interface affect the

stresses induced by the free-edge. However, there is very little difference between the four paths considered, except very close to the free-edge.

Figure 5.12 shows that the range of stresses due to the free-edge effect is almost the same as the range of stresses in the interior due to interacting fibers. However, it should be remembered that the analysis region, shown in Figure 5.3, corresponds to a lamina thickness of $0.05\,mm$, which would be used for a thin-ply laminate. Plies that are more often used in industry are twice as thick. The severity and extent of free-edge stresses is dependent on the ply thickness and layup. Also, a [0/90]s laminate is not expected to exhibit a severe free-edge effect compared to some thicker laminates, such as a [±45/0/90]s laminate. Such laminates should be examined in futures studies.

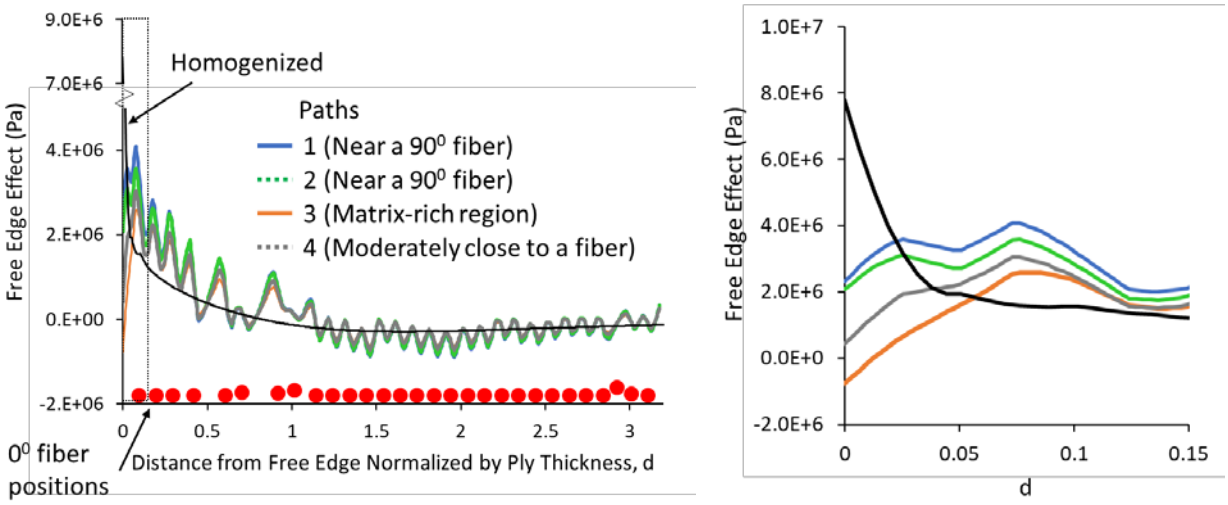

a) Free-edge effect within 3 ply thicknesses of the free-edge in addition to 0^0 fiber positions

b) Within 15% of a ply-thickness from the free-edge (zoom-in of dotted box in part a)

Figure 5.8. σ_{zz} for free-edge effect using the homogeneous and heterogeneous models

5.3.5 Effect of Fiber Arrangement

The fiber arrangement used for the previous results in this paper contains a lot of variation ((see realization A in Figure 5.3a). To study the effect of fiber arrangement, a more uniform fiber arrangement, as shown in Figure 5.3b, was considered and the interlaminar stresses for the two models were compared. Figure 5.13 shows the interlaminar normal stress, σ_{zz}, for both realizations. The stress distributions were quite similar, and the largest difference occurred where a large matrix pocket existed in realization A. Overall, the variation of σ_{zz} was more severe for realization A, due to the sometimes larger and smaller spacing between fibers than realization B. Similarly, Figure 5.14 shows the interlaminar shear stress, σ_{yz}, for both realizations, which also showed that realization A exhibits more variation of σ_{yz} due to the more varied microstructure. Though the trends are the very similar, randomness in the microstructure can causes relatively large differences in the interlaminar stresses.

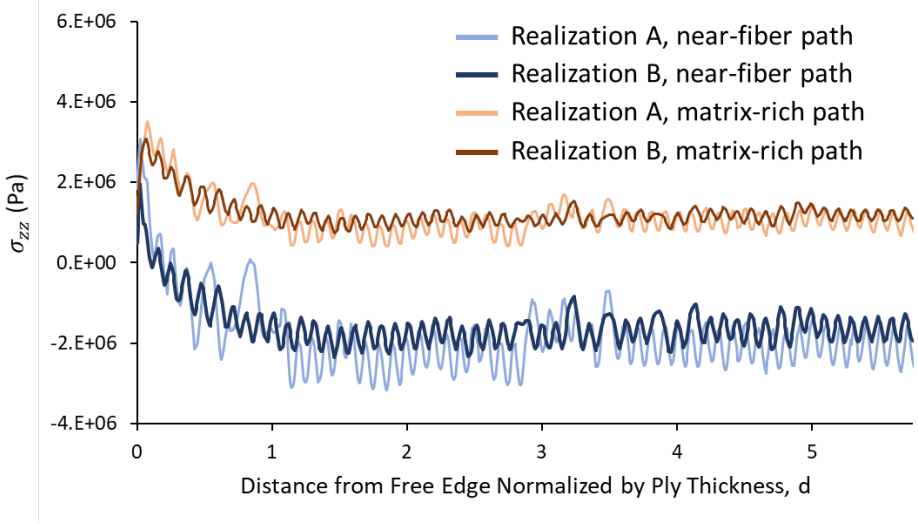

Figure 5.13. σ_{zz} for both realizations in Figure 5.3 for a near-fiber and matrix-rich path

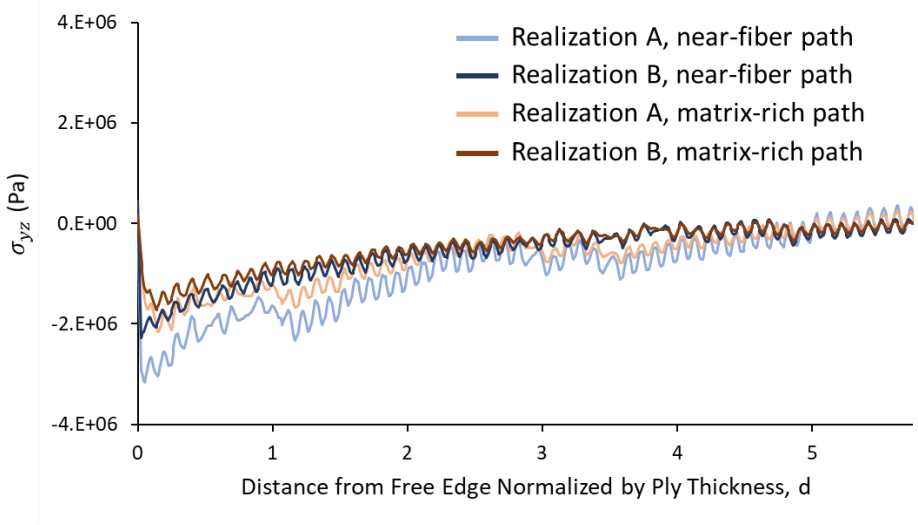

Figure 5.14. σ_{yz} for both realizations in Figure 5.3 for a near-fiber and matrix-rich path

5.3.6 Interaction Between Fibers and the Effect of $\boldsymbol{\nu^m}$

In the previous sections, it was shown that the heterogeneous microstructure led to complex stress distributions. It was observed in this paper and Refs. [58] and [62] that compressive stress concentrations form when fibers get close to each other. This section aims to explain the cause of the observed behavior. First, consider a 0^0 unidirectional composite modeled as an infinite array of fibers in a square arrangement, see Figure 5.8a, under uniaxial tension. For an applied strain of 1% along the fiber direction, Figure 5.15 shows the normal and shear stress distributions using the properties shown in Table 5.2, which were based on values in the literature. The non-zero stresses in Figure 5.15 are the result of the mismatch of the Poisson's ratios between the fibers and matrix. The important Poisson's ratios to compare for this case are ν_{12}^f, ν_{13}^f, since the tensile load is aligned with the fiber direction, and ν^m. However, the fibers are assumed to be transversely isotropic, so $\nu_{12}^f = \nu_{13}^f$, so this discussion will only refer to ν_{12}^f. When $\nu^m > \nu_{12}^f$, the effective ν_{xz} of the lamina is greater than ν_{12}^f, which means that the fibers will not contract transversely as much as the volume average of the lamina. This causes a compressive stress in the radial direction of the fiber and a tensile stress in the tangential direction, which can be observed in Figure 5.15a. As load is transferred between the tensile and compressive regions, shear stresses form as shown in Figure 5.15b. However, if ν^m was changed to be 0.292, the same value as ν_{12}^f, then σ_{xx} would be the only non-zero stress component. Additionally, it should be noted that if $\nu_{12}^f > \nu^m$ then the sign of the stresses in Figure 5.15 will be reversed. When there is a Poisson's ratio mismatch, stresses will develop even if the fiber is far away from

another fiber, but as fibers get close, more severe stresses develop, which explains the observation in Ref. [58].

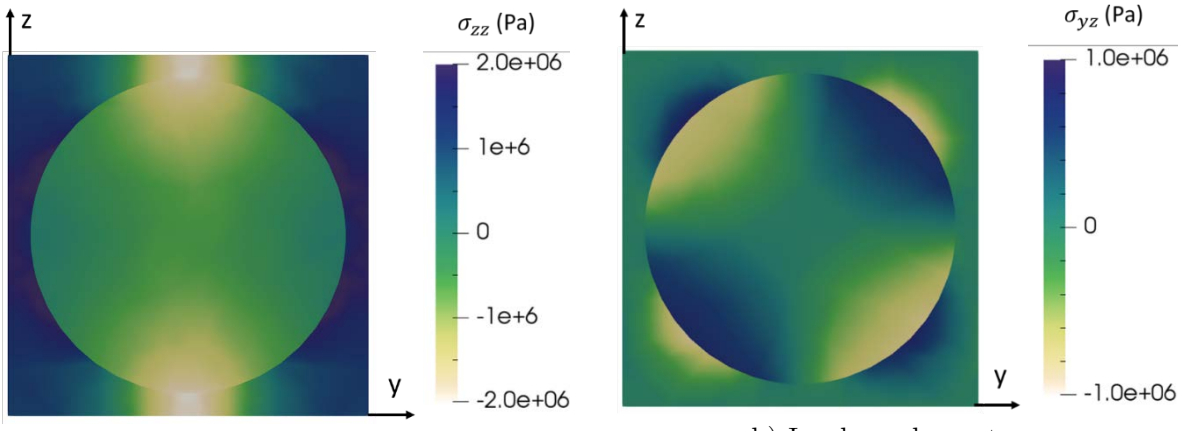

a) Normal stress, σ_{zz} b) In-plane shear stress, σ_{yz}

Figure 5.15. Stress distribution for a single fiber RVE subjected uniaxial 1% strain in x-direction using properties in Table 5.2

In a [0/90]s cross-ply under tension in the 0^0 direction, the 90^0 ply will primarily experience tension in the transverse direction, and similar to a ply under longitudinal tension, the Poisson's ratio significantly affects the stresses that will develop. For the properties in Table 5.2 and the case of transverse tension, the matrix will take most of the deformation in the loading direction, u_y, since E^m is much lower than E_2^f. If the positive-z plane of the RVE was left unconstrained, deformation along the y-axis will induce a displacement in the z-direction due to the Poisson effect, and since most of the deformation along the y-axis will occur in the matrix region, a much larger displacement in the z-direction will develop in the matrix region, as shown in Figure 5.16a. However, when

periodic boundary conditions are applied, the neighboring RVE's constrain the z-displacement, which develops tensile stresses in the matrix regions between fibers along the y-axis and compressive stresses between fibers along the z-axis, as shown in Figure 5.16b. Though it is not shown here, as the Poisson's ratios of the fibers and matrix, ν_{23}^{f} and ν^{m}, are reduced, σ_{zz} is also reduced.

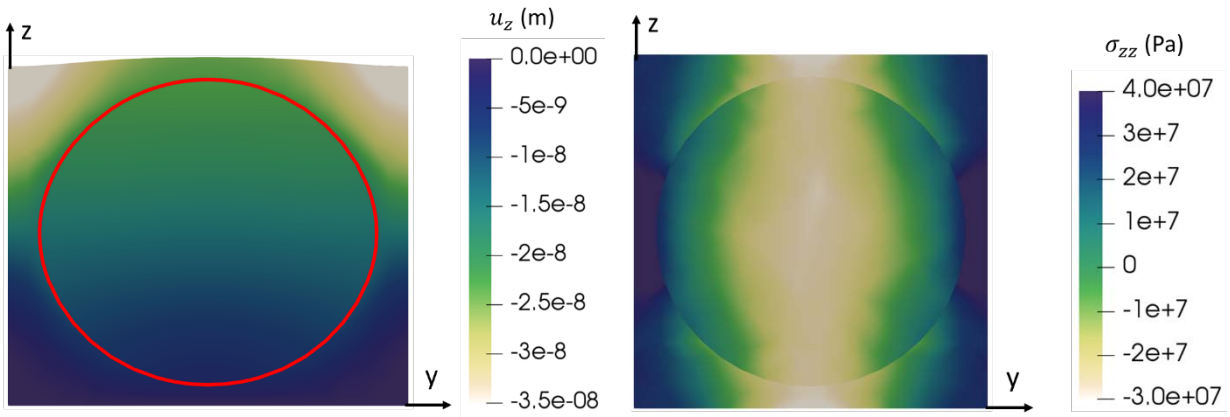

a) Z-displacement for case where top of RVE is unconstrained (deformation magnified by 10x)

b) Distribution of σ_{zz} for case where RVE is periodic

Figure 5.16. A single fiber RVE subjected to 1% strain in transverse y-direction using the properties in Table 5.2 for two different cases

Understanding the effect of Poisson's ratio and interaction of fibers for a unidirectional lamina under longitudinal and transverse tension help provide insight into the interlaminar stresses that develop in a cross-ply laminate under uniaxial tension. For a cross-ply laminate under tension in the 0^0 direction, there are three mechanisms that develop stresses along the ply interface. First, significant σ_{zz} and σ_{yz} stresses will develop in the 0^0 plies due to

a mismatch between ν^m and ν_{12}^f. Second, significant σ_{zz} stresses will develop in the 90^0 ply due to the Poisson effect, and the magnitude of the stress is dependent on primarily ν^m and secondarily ν_{23}^f. The stress decreases as the Poisson's ratios decrease. Third, interlaminar shear stresses will develop as load is being transferred between the 0^0 and 90^0 plies, due to the large mismatch of Poisson's ratios for the two plies.

For the properties in Table 5.2, $\nu^m > \nu_{12}^f$, so along the ply interface, σ_{zz} will be compressive near 0^0 fibers. Additionally, ν^m and ν_{23}^f are relatively large, so strong compressive stresses form along the ply interface near 90^0 fibers. This behavior is illustrated in Figure 5.17. However, the compressive contribution from the 0^0 ply would be reduced significantly if $\nu^m = \nu_{12}^f$, and the compressive contribution from the 90^0 ply would be reduced if ν^m was reduced. Furthermore, the contribution from the 0^0 ply to the interlaminar normal stress would be tensile if $\nu^m < \nu_{12}^f$. Without modifying the fiber properties, there is an optimal ν^m between ν_{12}^f and 0.0 that minimizes the peak stresses. Figure 5.18 shows the effect of reducing of the Poisson's ratio of the matrix for the case where $\nu^m = 0.16$ and $G^m = 1.49\ GPa$, to maintain isotropy. The interlaminar normal stress is reduced significantly by reducing ν^m, but the interlaminar shear stress remains significant due to load being transferred between the 0^0 and 90^0 plies. However, it should be noted that it is unlikely that an epoxy matrix can be manufactured with such a small value of ν^m as 0.16. Typically, fillers, such as nanoinclusions or polymer nanofibers, are added to epoxy resins to modify the properties, but it has not been shown in the literature that a Poisson ratio of 0.16 for the matrix is possible with an epoxy system.

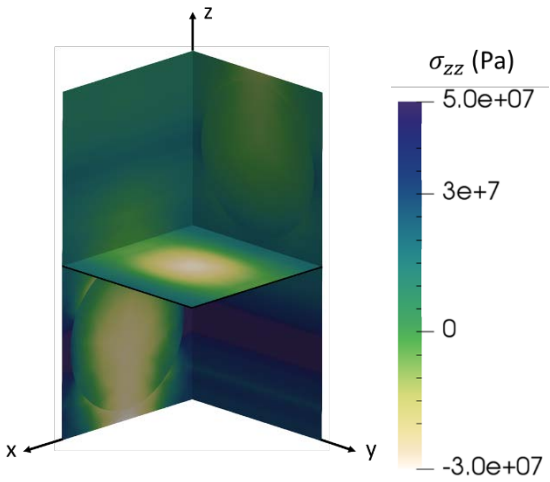

a) Distribution of σ_{zz}

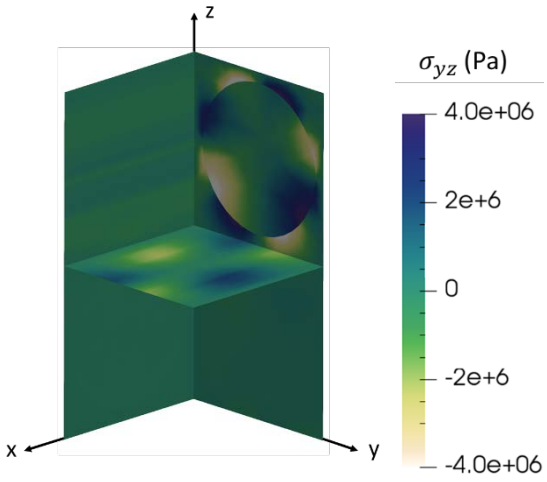

b) Distribution of σ_{yz}

Figure 5.17. Stress distribution for RVE of two crossing fibers subjected uniaxial 1% strain in x-direction using properties in Table 5.2

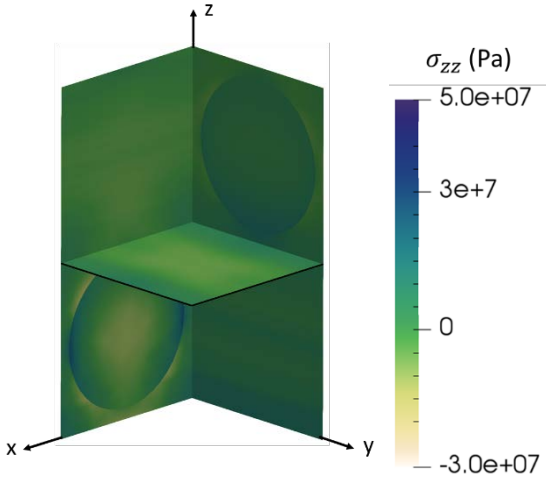

a) Distribution of σ_{zz}

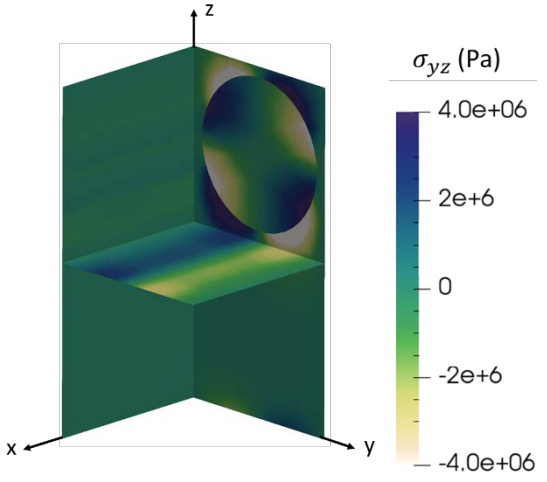

b) Distribution of σ_{yz}

Figure 5.18. Stress distribution for RVE of two crossing fibers subjected uniaxial 1% strain in x-direction setting $\boldsymbol{\nu^m = 0.16}$ ($\boldsymbol{G^m}$ changes to maintain isotropy)

The cross-ply laminate considered in this paper has a more complex fiber arrangement, a finite thickness, and a free-edge effect, so it expected that changing ν^m will not have as significant an effect as for the periodic square arrangement shown in Figure 5.18. To clearly illustrate the effect of reducing ν^m for the cross-ply laminate, Figure 5.19 shows the interlaminar normal stress with $\nu^m = 0.16$. As expected, reducing ν^m significantly reduces the effect of heterogeneity, which can be more clearly observed in regions away from the free-edge. It should be noted that the optimal ν^m for the periodic square fiber arrangement is probably not the optimal ν^m for the cross-ply laminate, due to the differences mentioned earlier. Optimization studies are needed to understand what set of microscale properties would result in the maximum reduction in the extreme stresses in both plies.

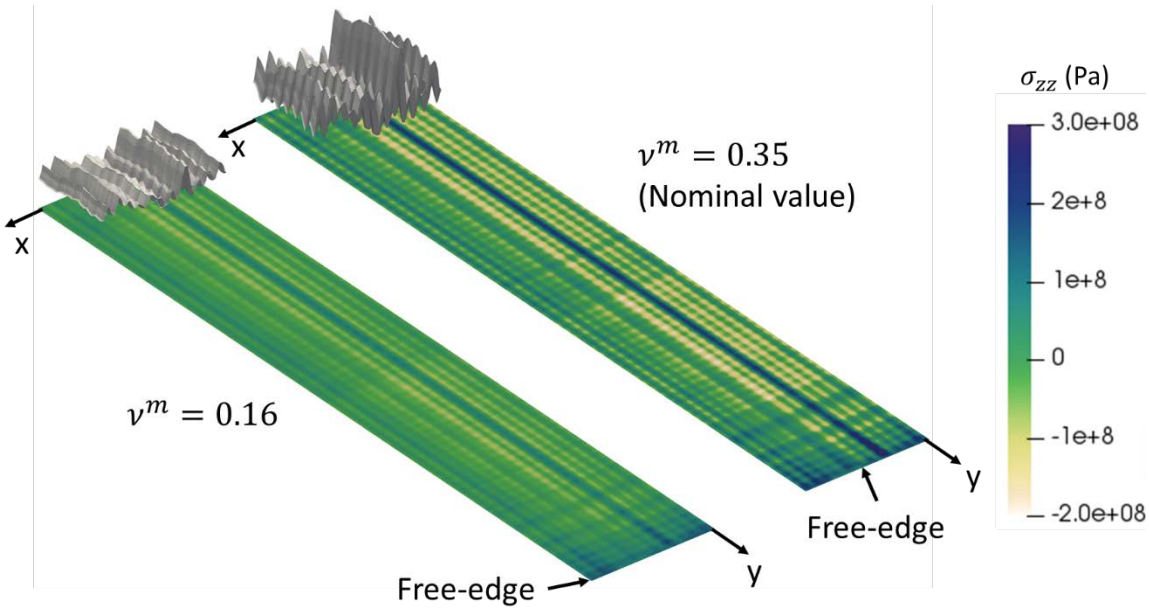

Figure 5.19. Comparison of σ_{zz} along 0-90 interface for two different values of ν^m (region far from free edge is deformed proportional to σ_{zz} to illustrate difference)

5.3.7 Summary

The classical free-edge problem was revisited for a crossply laminate with more detail than has been done before in the literature, and the predictions of a heterogeneous model, which models the fibers and matrix discretely, was compared to the predictions of a homogeneous model, which models each ply as a homogenous material. A few fiber diameters from the free-edge, the trend of the interlaminar normal and shear stresses along the ply interface were similar between the two models. However, the interaction of fibers near the ply interface caused significant local differences, due to compressive stress concentrations that developed when fibers came close to each other, and very different stresses close to the free-edge, as expected, due the singularity at the intersection of the free-edge and 0/90 ply interface for the homogenous model. The effect of the free-edge and heterogeneity were separated. With the free-edge effect isolated, predictions between the models matched well, expect for within a few fiber diameters of the free-edge. Isolating the effect of heterogeneity showed that fibers and matrix pockets significantly affect the stress state at the microscale. A comparison of the predicted interlaminar stresses for highly random and more uniform fiber arrangement showed that the fiber arrangement can significantly affect the local stresses and the peak stresses are reduced when the microstructure is more ordered. By investigating a set of more simple problems, it was shown that the effect of heterogeneity for a cross-ply laminate under uniaxial tension primarily depends on the Poisson's ratios. For a graphite fiber and epoxy matrix composite,

there is a value of ν^m between 0 and ν_{12}^f that reduces the effect of heterogeneity and minimizes the peak local interlaminar stresses.

This study focused on the simplest layup possible, a crossply laminate, but the laminate is known to exhibit a small free-edge effect. Insights into the effect of microscale heterogeneity on the local stresses near a free-edge were discovered, and but the next section considers a more realistic quasi-isotropic layup, which is known to exhibit a more pronounced free-edge effect.

5.4 Free-Edge Analysis of a [±45/0/90]s Crossply Laminate

This section explores the effect of microscale heterogeneity on stresses near a free-edge in a quasi-isotropic [±45/0/90]s laminate and assesses the microscale region size needed to accurately predict the stresses near the free-edge. The results can be separated into three primary sections. First, the effect of modeling discrete fibers and matrix is investigated by comparing a model where the entire 0^0 and 90^0 plies are modeled as discrete fibers and matrix to a classical homogeneous model, which treats all plies as homogeneous materials. This is followed by two sections that explore the effect of reducing the size of the microscale region in the 0^0 and 90^0 plies, first in the direction normal to the ply interface and followed by the direction normal to the free-edge.

5.4.1 Effect of Modeling Discrete Fibers and Matrix

Several models with varying sizes of microscale regions are considered in this paper. This section focuses on comparing a classical homogeneous model, where all plies are treated as a homogeneous, orthotropic material, to a model where the microstructure of the 0^0 and

90^0 plies are discretely modeled with a random fiber arrangement. In both cases, the $\pm 45^0$ plies are treated as a homogeneous material, since the stresses along the 0-90 and 90-90 interfaces are the focus of this paper.

Figure 5.20 shows the deformed y-z cross-section at $x = W/2$ for the case where the 0^0 and 90^0 plies are modeled as discrete fibers and matrix, along with an outline showing the deformed boundary of the classical homogeneous model. The deformation was exaggerated by 20x to clearly illustrate the differences, since the applied 1% strain results in very small displacements. The y and z components of displacements along the boundary of the cross-sections matched well and differed at most by 6%. The largest differences along the boundary were along the free-edge in the top three plies.

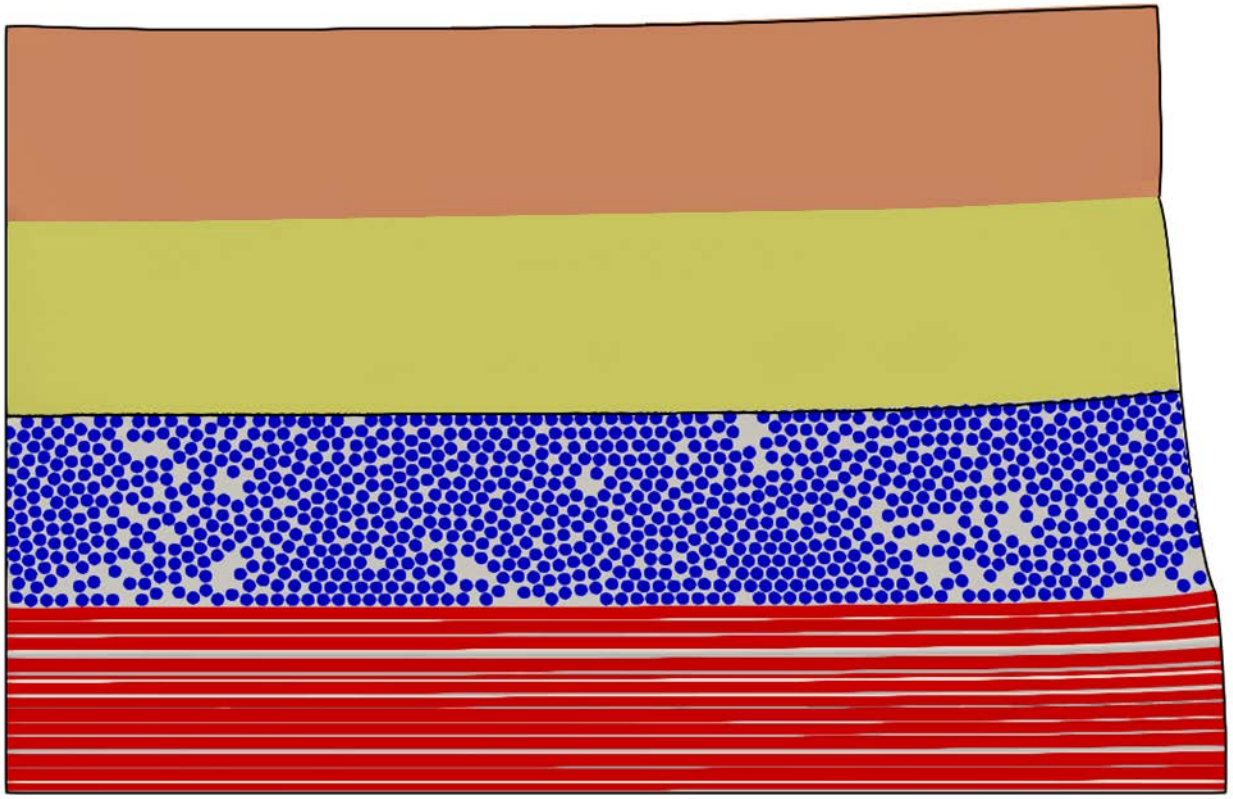

Figure 5.20. Deformed discrete fiber/matrix case with outline showing deformed classical homogeneous model (20x displacement magnification)

For the classical homogeneous model under uniaxial extension, σ_{zz} is the only significant interlaminar stress along the 90-90 ply interface, while σ_{zz} and σ_{yz} are both significant along the 0-90 ply interface. However, where fibers and matrix are modeled discretely, σ_{xz} is not necessarily small. This study will limit the scope of the discussion to the interlaminar normal stress, σ_{zz}, along the ply interfaces since it is known that this layup experiences predominantly mode I failure. Fig. 6a shows contours for σ_{zz} for the entire 0-90 and 90-90 ply interfaces. The complex stress pattern was due to the interactions of fibers near the

interface. As noted in Refs. [58] and [62], compressive stresses form when fibers come into close proximity. Ref. [62] also observed that tensile stress concentrations form near matrix pockets. Consequently, along the 0-90 ply interface, the interaction of 0^0 and 90^0 fibers created a variation of σ_{zz} in two directions, while σ_{zz} along the 90-90 ply interface exhibited variation in only one direction due to the microstructure. In addition to the effect of the heterogeneous microstructure, the mismatch of effective Poisson's ratios between plies caused an elevated σ_{zz} near the free-edge. Qualitatively, Figure 5.21a shows that the free-edge has a noticeable effect further into the laminate along the 90-90 ply interface. Figure 5.21b focuses on the stresses near the free-edge. In the figure, the surface of the ply interface is warped proportional to σ_{zz}, and the stress for the classical homogeneous model is shown as a solid line. The classical homogeneous model captures the trend that the stresses increase near the free-edge, but does not capture the significant variation due to the microstructure. Figure 5.21b highlights the more complex stress pattern along the 0-90 ply interface due to the interaction of fibers and shows that the stress at the free-edge is higher for the 0-90 interface than for the 90-90 interface.

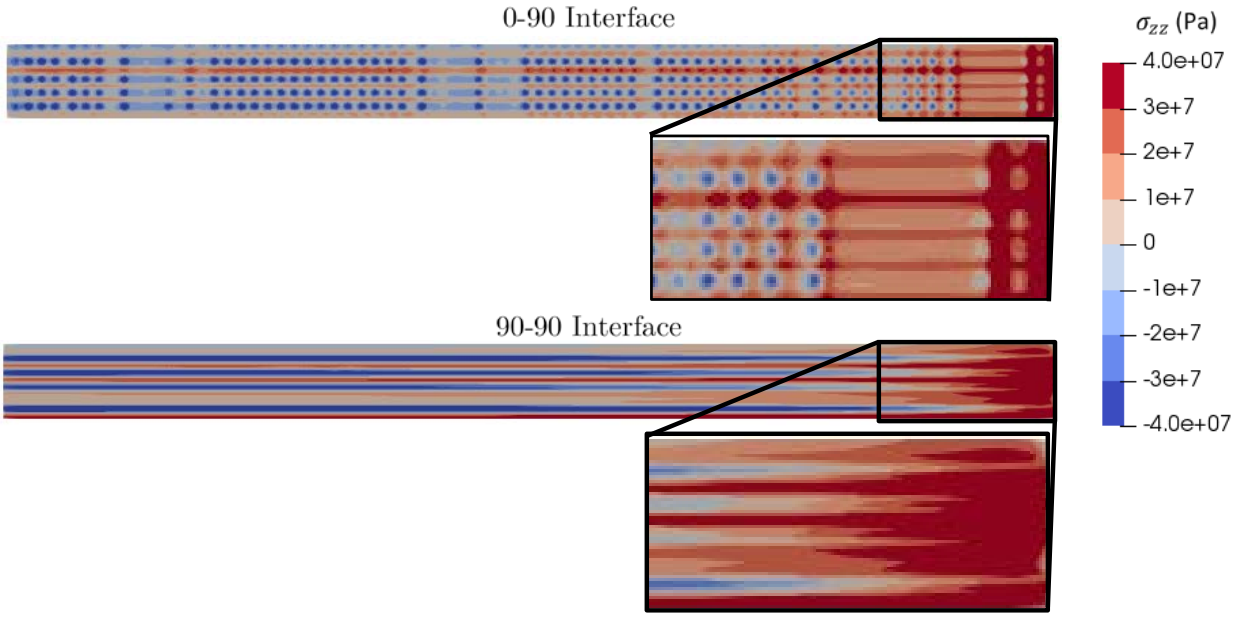

a) Comparison of $\boldsymbol{\sigma_{zz}}$ contours for the 0-90 and 90-90 ply interfaces for the model where fibers and matrix are discretely modeled

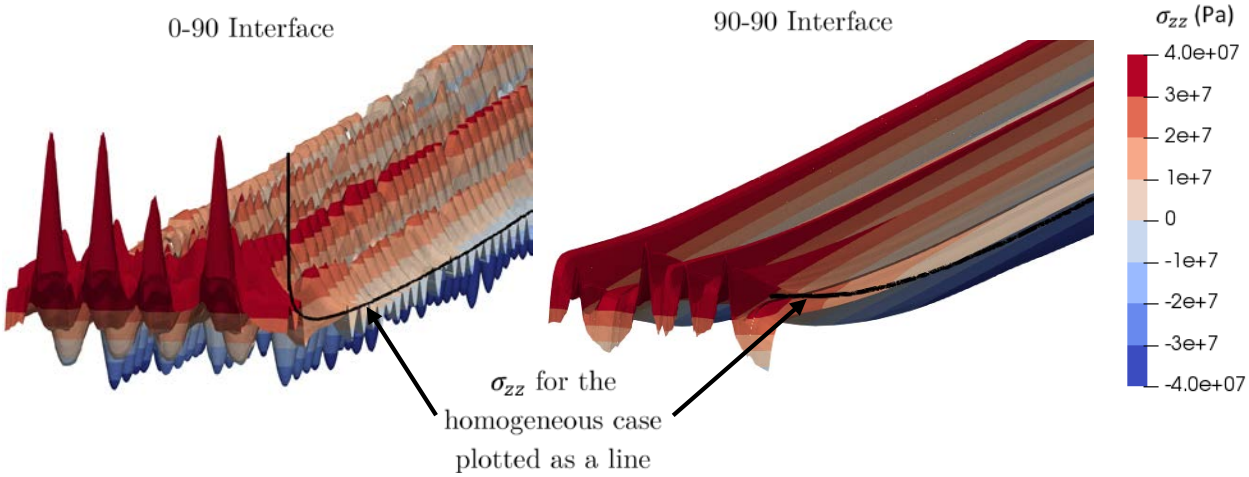

b) Comparison of $\boldsymbol{\sigma_{zz}}$ near the free-edge along the 0-90 and 90-90 ply interfaces (surfaces are deformed proportional to $\boldsymbol{\sigma_{zz}}$) along with the response of the classical homogeneous model

Figure 5.21. Comparison of $\boldsymbol{\sigma_{zz}}$ at the 0-90 and 90-90 ply interfaces

Because the peak stress is higher and the stress field more complex for the 0-90 ply interface, this study limits further investigation of the effect of modeling the discrete microstructure to the 0-90 ply interface. For a more quantitative comparison, σ_{zz} along the

0-90 ply interface is compared for several paths chosen that are parallel to the y-axis. Ref. [62] showed that σ_{zz} varies very differently paths chosen near a fiber versus a matrix pocket. Consequently, two paths were considered: one near a 90^0 fiber (at $x = 2e\text{-}5\ m$) and one near a matrix rich region in the 90^0 ply (at $x = 1.1e\text{-}5\ m$). The locations of these two paths are illustrated in Figure 5.22.

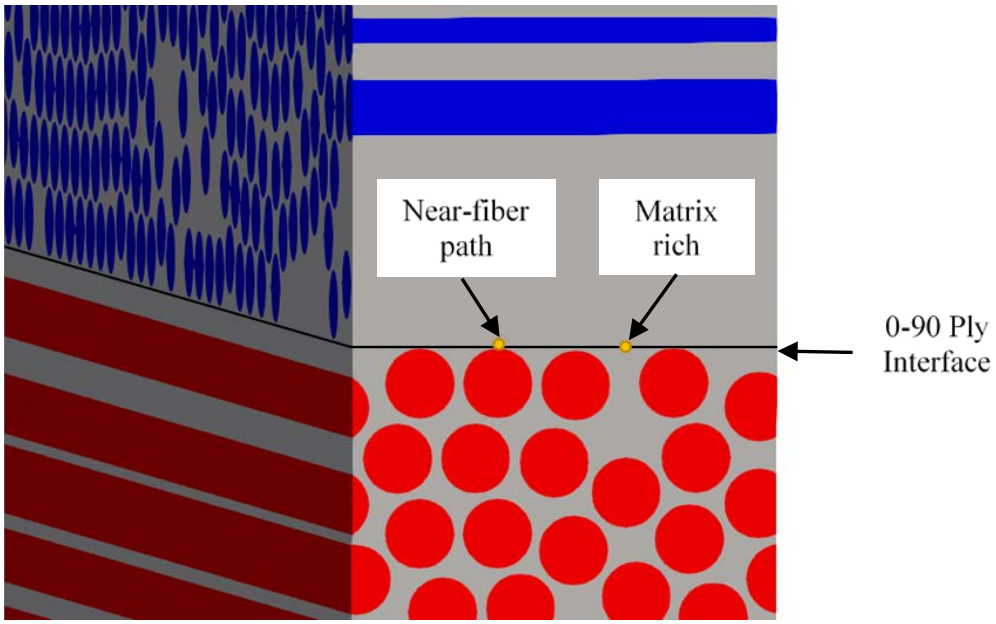

Figure 5.22. Illustration of two paths used for line plots of stress along the 0-90 interface

Figure 5.23 shows σ_{zz} along the two paths in Figure 5.22 for the case where the microstructure of the 0^0 and 90^0 plies is modeled discretely in addition to σ_{zz} for the classical homogeneous model. For the path near a 90^0 fiber, the peak σ_{zz} occurred at the free-edge, while the maximum stress for the matrix-rich path occurred about two fiber diameters from

the free-edge due to a matrix pocket in the 0^0 ply. The peak σ_{zz} was much higher for the near-fiber path, but the average σ_{zz} the matrix-rich path was higher over the width of the specimen. For both paths, σ_{zz} exhibited significant oscillations due to interaction between 0^0 and 90^0 fibers. The valleys of the oscillations corresponded to centers of 0^0 fibers, which is related to the tendency of fibers in close proximity to cause a compressive stress as noted earlier.

Figure 5.24 shows the same data as Figure 5.23 but for a region near the free-edge. The figure shows that σ_{zz} increased near the free-edge for the near-fiber path, while the stress decreased near the free-edge for the matrix-rich path. This difference was likely due to a nearby singularity at the intersection of the fiber/matrix interface and free-edge. For the near-fiber path, a singularity was close to the ply 0-90 interface, which caused an elevated stress. On the other hand, the matrix-rich path experienced a peak stress a few fiber diameters away from the free edge, where there was a large matrix pocket in the 0^0 ply for the RVE considered. For some values of y, σ_{zz} was a different sign and more than 200% different between the two paths considered. Though displacements along the boundary of the model matched well, the stresses along the ply interfaces differed significantly when fibers and matrix were modeled discretely.

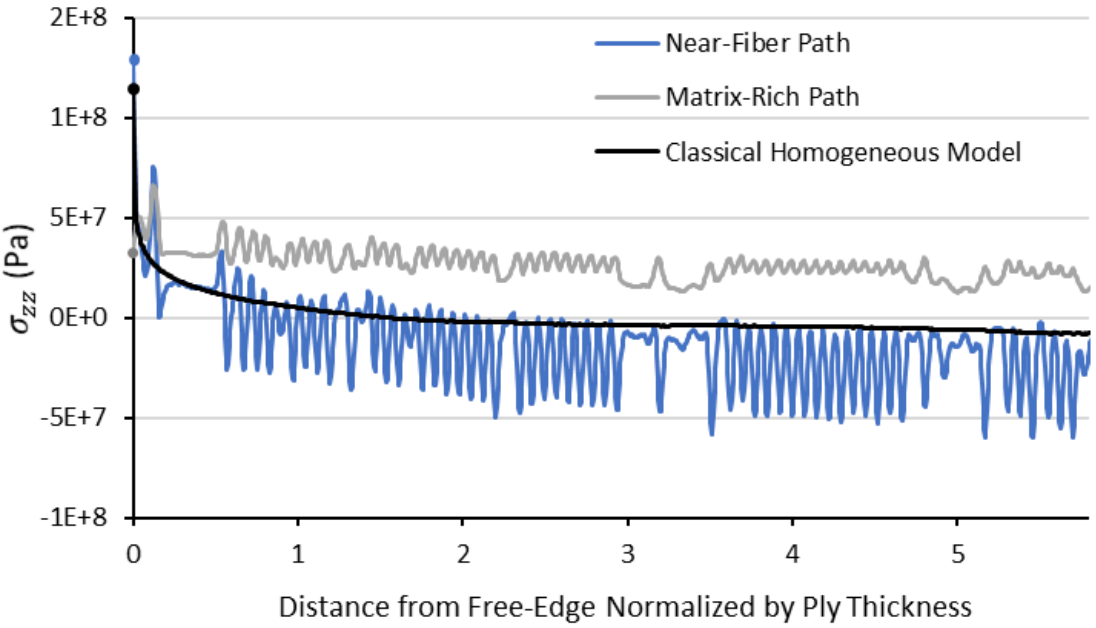

Figure 5.23. σ_{zz} along selected paths vs. the distance from the free-edge normalized by the ply thickness for the case where the microstructure of the 0^0 and 90^0 plies is modeled discretely

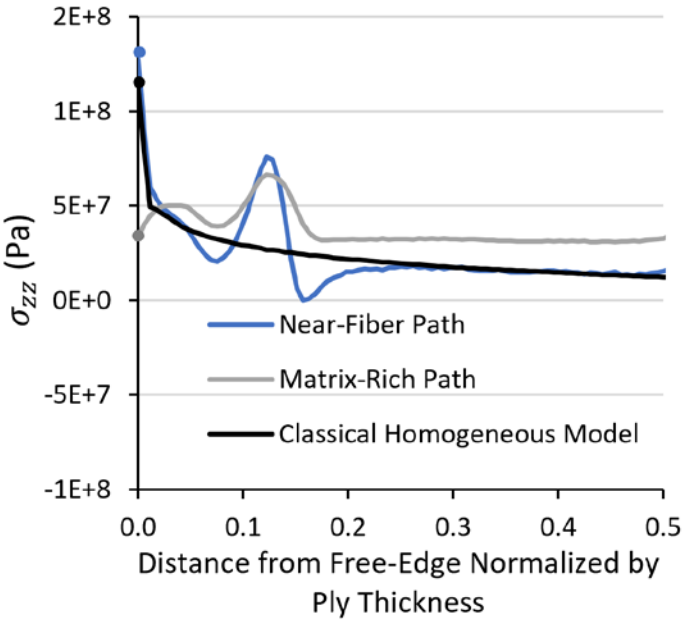

Figure 5.24. σ_{zz} along selected paths from Figure 5.23 shown for half a ply-thickness from the free-edge

5.4.2 Reducing the Microscale Region Normal to the Ply Interface

In the previous section, the entire 0^0 and 90^0 plies were modeled at the microscale, but this section investigates the effect of reducing the microscale region in the 0^0 and 90^0 plies in the direction normal to the ply interface, which is along the z-axis. For the region of the 0^0 and 90^0 plies modeled at the mesoscale, homogeneous, orthotropic properties are used. The reference model contains several types of error, but the difference between models with a reduced microscale region and the reference model represents an additional error due to reducing the microscale region. Two cases with reduced microscale regions are compared to the case where the entire 0^0 and 90^0 plies are modeled using discrete fibers and matrix. In one case, 25% of the plies are modeled at the microscale, and in the other, only 10% of the plies are modeled at the microscale. Figure 5.6 shows both models.

Figure 5.25 shows σ_{zz} along the 0-90 ply interface as a function of the normalized distance from the free edge for both paths shown in Figure 5.22 and all three models considered. For any point with a $|\sigma_{zz}|$ greater than 30% of the peak σ_{zz} along the matrix-rich path, reducing the microscale region to 25% of the ply thickness resulted in a maximum error of about 20% for the near-fiber path and 10% for the matrix-rich path, and reducing the microscale region to 10% of the ply thickness resulted in a maximum error of about 38% for the near-fiber path and 20% for the matrix-rich path. However, no locations with a relatively large error were near the locations of peak stress. Comparing only the peak stress for each path, reducing the microscale region to 25% of the ply-thickness introduced about

1% error, while reducing the microscale region to 10% of the ply-thickness introduced about 6% error.

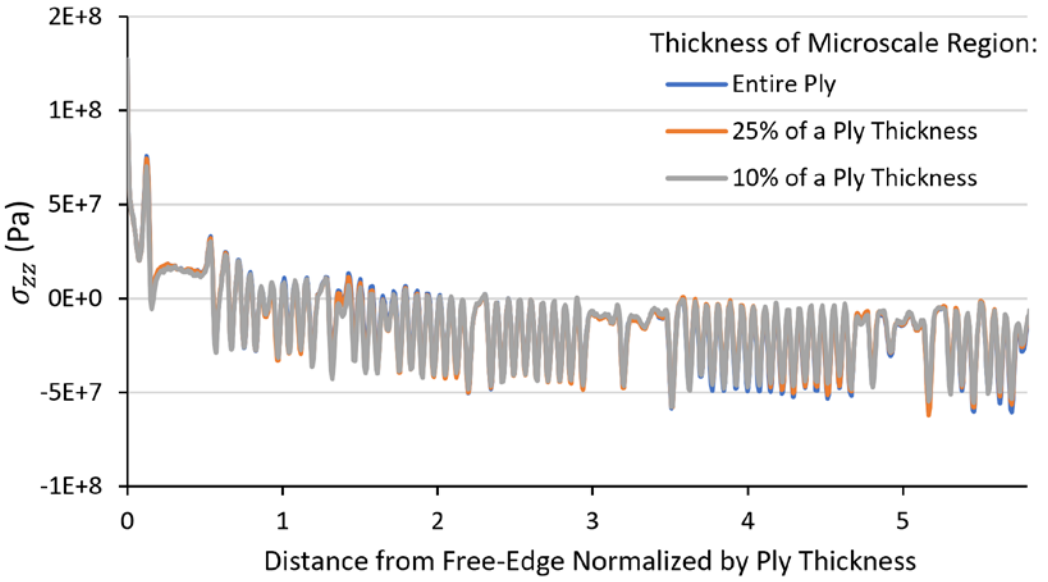

a) Near-fiber path

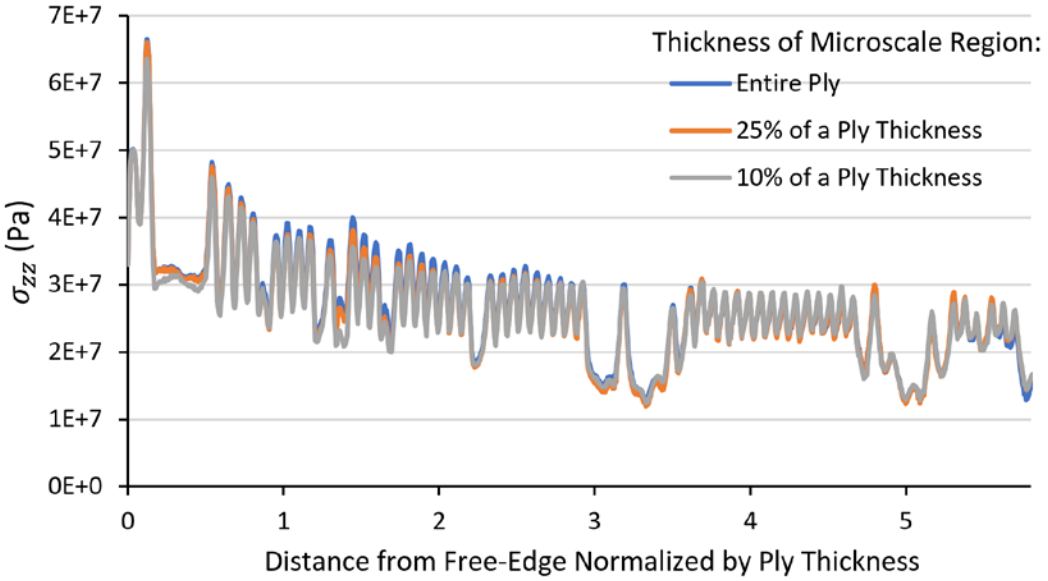

b) Matrix-rich path

Figure 5.25. σ_{zz} plotted along paths parallel to the y-axis versus the distance from the free-edge normalized by the ply thickness for several sizes of microscale regions

5.4.3 Reducing the Microscale Region Normal to the Free-Edge

The previous section explored reducing the microscale region in the 0^0 and 90^0 plies in the direction normal to the ply interface, but this section investigates reducing the microscale in the direction normal to the free-edge, which is along the y-axis. The model from the previous section with 10% of the ply-thickness modeled at the microscale was used as the base-line, which is shown in Figure 5.6, and the microscale region was reduced in the y-direction to one- and two- ply thicknesses from the free-edge. Figure 5.7 shows both configurations.

Figure 5.26 shows σ_{zz} along the 0-90 ply interface as a function of the normalized distance from the free edge for both paths shown in Figure 5.22 and all three models considered. As expected, σ_{zz} quickly dropped to zero where homogeneous properties are used for the 0^0 and 90^0 plies. About two fiber diameters from the transition of modeling fibers and matrix discrete to using homogeneous properties, σ_{zz} matched quite well. For any point within 75% of a ply-thickness from the free-edge, σ_{zz} differed from the baseline about the same amount for either case, with a maximum error of about 8% for the near-fiber path and about 9% for the matrix-rich path. However, when the microscale region was reduced to two or one ply-thickness from the free-edge, the peak σ_{zz} only differed from the baseline by 3% or less.

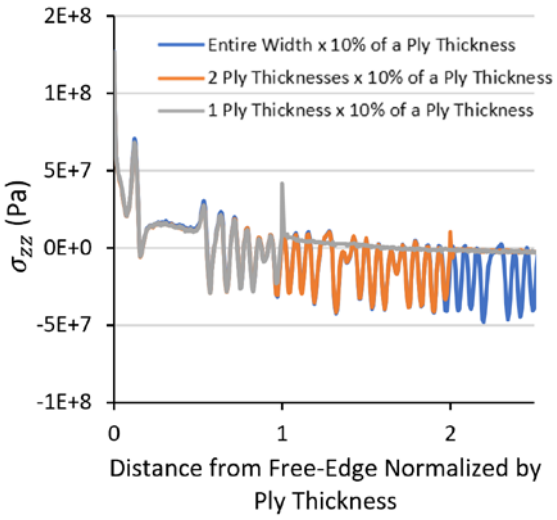

a) Near-fiber path

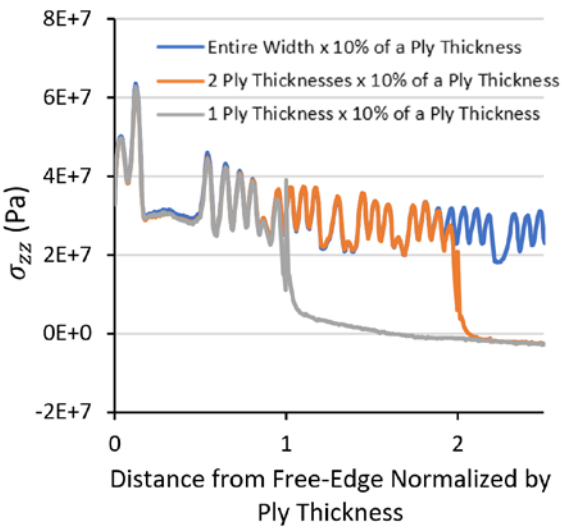

b) Matrix-rich path

Figure 5.26. σ_{zz} plotted along paths parallel to the y-axis versus the distance from the free-edge normalized by the ply thickness for several sizes of microscale regions

5.4.4 Summary

Modeling the heterogeneous microstructure of a $[\pm 45/0/90]_s$ laminate under uniaxial extension had a very significant effect on the predicted interlaminar normal stress along the 0-90 and 90-90 ply interfaces. The interlaminar normal stress varied differently along the ply interface from the free-edge depending on whether the path chosen was near a 90^0 fiber or a matrix-rich region of the 90^0 ply. For a path near a fiber, the peak stress occurred at the free-edge, but for a path near a matrix-rich region, the peak stress occurred near a matrix pocket in the 0^0 ply, which was about two fiber diameters away from the free-edge. The displacements along the boundary of the model closely matched those between the two models, but the interlaminar normal stress predictions were very different.

To allow smaller, more efficient models for future studies, the effect of reducing the region modeled at the microscale was investigated. The microscale region was reduced in the direction normal to the ply interface to 25% and 10% of the ply-thickness, which caused a maximum of 38% and 20% error respectively for any point with a significant stress value and 1% and 6% error respectively for the peak stresses. Reducing the microscale region in the direction normal to the free-edge to be one and two ply-thickness in size did not have a significant effect on the predicted interlaminar normal stress at points within 75% of a ply-thickness of the free-edge. Although some error is introduced, only modeling a small region of the plies as discrete fibers and matrix can drastically reduce the size of the mesh and the computational costs.

5.5 Conclusions

With growing computational resources available to researchers, composite materials can be investigated with far more detail than possible in the past. New capabilities open up the study of novel materials, but it also allows researchers to understand the accuracy and limitations of assumptions that were necessary for past decades. This chapter revisited the classical free-edge problem, which has been exhaustively studied at the mesoscale, with far more detail than done before. Accounting for the microscale heterogeneity removed the singularity at the ply interface and was shown to dramatically affect the predicted interlaminar stresses near a free-edge. The heterogeneous microstructure was shown to also affect the stresses away from the free-edge. The fiber arrangement had an effect on the local stresses, and the more uniform fiber arrangement resulted in lower peak stresses.

Additionally, the interaction of fibers lead to compressive stress concentrations, indicating that matrix rich regions should be avoided near the ply interface to lower the tendency for delamination. Finally, the region needed to accurately predict the microscale stresses near the ply interface was shown to be much smaller then entire ply. Fortunately, future studies can leverage much smaller meshes, allowing larger regions of the laminate or more complex layups to be modeled.

For analyses focused on mesoscale phenomena, treating the plies as homogeneous materials has been shown to work well by many other researchers. However, if microscale phenomena are important, the heterogenous microstructure must be modeled discretely. Many open questions remain for the community on the subject of free-edge induced delaminations, such as understanding where the microscale damage initiates and how it propagates. The layups of laminates have been tailored to mitigate the chances of delaminations, but a deeper understanding will be needed to tailor the microstructure. Future studies should explore the effect of plasticity in the matrix near the free-edge and predict the progression of fiber/matrix debonding and matrix damage.

6 PREDICTING THE PERFORMANCE OF 3D TEXTILES AT THE MESOSCALE*

6.1 Overview

Two-dimensional (2D) fiber/matrix composites, such as laminated composites or 2D textiles, have been extensively studied and routinely used within the aerospace industry. However, three-dimensional (3D) textiles have not achieved the ubiquity of 2D composites, despite several key advantages. Three-dimensional textiles offer improved out-of-plane properties and significantly improved impact resistance over traditional 2D composites, and the 3D architectures of tows allow complex shapes to be manufactured. Additionally, 3D woven composites provide the potential for many multifunctional materials, since functional constituents can be woven into the fabric.

The adoption of 3D textiles has been slowed by complex manufacturing techniques and limited ability to predict performance. The complex geometry makes it difficult to create realistic textile models and a valid mesh for finite element analysis (FEA). Additionally, the resulting system of equations can be relatively expensive to solve due to the large number of degrees of freedom required. Consequently, 3D textiles have been studied with less rigor than 2D composites, and the relatively few works in the literature that do investigate 3D textiles provide limited insight into the complex stress states within the material under static and quasi-static loads. Much of the literature focuses on the dynamic

---

* Parts of this material originally appeared in the Proceedings of the American Society for Composites: Thirty-second Technical Conference, 2017. Lancaster, PA: DEStech Publications, Inc. and the Proceedings of the 21st International Conference on Composite Materials. It is reprinted here with permission of the publisher.

response due to ballistic loads, since 3D textiles are often used for body armor. Furthermore, most works have assumed that the tow architecture has an idealized shape and path, neglecting the complex tow architecture of realistic composites.

This chapter aims to explore the behavior of orthogonally woven composites under different in-plane loads with greater detail than done in the past. The deeper investigation provides insight into how load is distributed throughout the complex tow architecture, locations of critical stress when the material is subjected to different loads, the role of the binders, and the effect of several modeling parameters. Nonidealized textile models were created by simulating the processing using VTMS and meshing the resulting tow geometry using the algorithm described in chapter 2. Two types textiles are considered, a thin 1x1 orthogonal weave, which is a relatively simple 3D textile, and a thicker 2x2 orthogonal weave, which is more realistic to the 3D textiles used in industry. The next section provides a description of the two types of textiles. After that, the investigation for each of the two types of weaves is given in the next two sections, respectively. A summary of the relevant results is provided within each section, but the final section of this chapter provides conclusions from both studies and highlights remaining open questions for future work.

6.2 Configurations

Two types of orthogonally woven textiles are considered in this chapter. This section begins with a description of the textile geometry for each type of textile, followed by a presentation of the meshes created for each textile geometry. Next, the boundary conditions for several load configurations is discussed. The following section describes how the analysis

region is clipped down to a subregion for one of the textile types to avoid boundary effects during visualization and postprocessing. Finally, the material properties used for each type of textile are provided.

6.2.1 Textile Geometry

Two types of orthogonally woven textile composites are considered in this chapter. Orthogonally woven textiles consist of layers of warps and wefts with binder tows weaving through the entire thickness. Binder traverse through the thickness of the textile and travel generally in the x-direction. Warps are relatively straight tows that are approximately aligned with the x-axis, and wefts are relatively straight tows that are approximately aligned with the y-axis. In a 1x1 orthogonal weave, the binder crosses one weft at the top and bottom of the textile before traversing through the thickness of the textile, while in a 2x2 orthogonal weave, the binder crosses two wefts before traversing the thickness. Figure 6.1 shows an example of the both types of textile architectures. Without the binders, the warps and wefts are similar to a crossply laminate but with some space in between tows. The binders provide the 3D component and provide the improved out-of-plane properties, including delamination resistance.

The 1x1 orthogonal weave was considered due to the relative simplicity of the tow architecture. The model contained three rows of wefts and two rows of warps, which is thin compared to a system for a practical application. The binders were assumed to have about the same cross-sectional area as the warps and wefts. This results in a higher binder volume fraction than most textiles in use. Due to the simplicity, the transfer of load is easier to

understand, making the 1x1 orthogonal weave a logical starting point for a detailed analysis of 3D textile composites. The 2x2 orthogonal weave provides a more realistic tow architecture with five rows of wefts and four rows of warps. The binders were about 3.5% of the volume, while the warps and wefts composed 17.2% and 20.9% of the volume respectively. The simulated textile had many similarities to the textile manufactured by T.E.A.M. Inc. and studied by Pankow et. al. in Refs. [63] and [64].

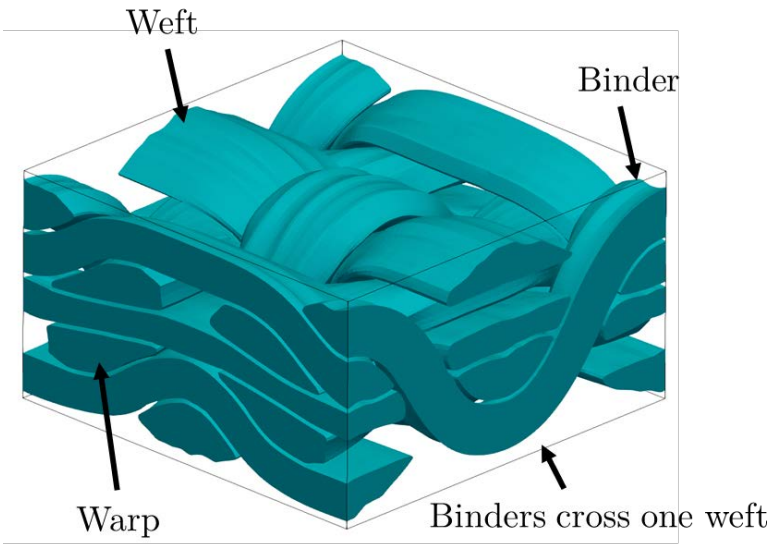

a) 1x1 orthogonal weave

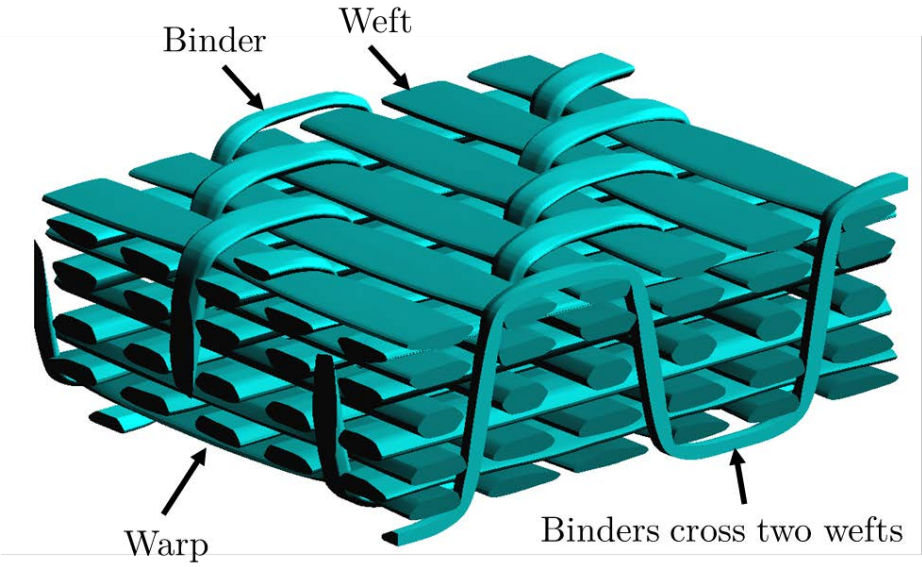

b) 2x2 orthogonal weave

Figure 6.1. Illustration of two types of orthogonally woven textiles considered

A software tool called VTMS was used to create the geometry. [22] The Air Force Research Lab (AFRL) developed VTMS for modeling the effects of processing on the tow geometry, allowing the generation of realistic textile models. It models the effects of processing using "digital chains", which consist of rigid spheres connected by elastic rod elements. Ideally, a digital chain should be used for each fiber within the tow, but this is currently computationally prohibitive. Initially, the digital chains have too much space between them. The displacements at the ends of the digital chains are specified to be zero, and a contact analysis is used to solve for the final configuration. During the contact analysis, a tension is specified within the digital chains and a pair of rigid planes is used to compress the top and bottom of the model. The contact analysis prevents overlap between the rigid spheres of the digital chains, and tension in the chains encourages them to move towards the midplane. It is useful to begin with one digital chain per a tow until much of the unrealistic space between the tows is removed, and afterwards, the single digital chains are replaced with more digital chains. Each single digital chain was replaced with ten digital chains in this study. With multiple digital chains per a tow, the cross-section of the tows can deform during the contact analyses. When the compaction is completed, cross-sectional profiles are created along each tow path by fitting a closed spline around the bundle of digital chains that define the tow cross-section for each sample point along the tow path. Typically, 200 to 500 cross-sectional profiles are sampled for each tow. Once a solid geometry has been created for each tow, there are small interpenetrations between the tows, which are removed by adjusting the geometry of each tow. The algorithm used in

VTMS to remove interpenetrations is not precise enough to allow the creation of a conforming, non-intersecting volume mesh for each tow. Consequently, the cross-sections for each tow are shrunk towards the centroid by a small amount, resulting in some matrix material between each tow. For the final steps in VTMS, the surface geometry of each tow is discretized and clipped to a smaller region to avoid the spurious effects near the edges of the initial model. Refer to chapter 2 for a much more detailed account of how the textile models were created.

Figure 6.2a shows a typical final model with the matrix, along with the dimensions of the model. Figure 6.2b shows the tow architecture more clearly by removing the matrix and highlights the variation of cross-section shapes within the tows. Due to the compact algorithm used by VTMS, the outermost tows experience more deformation than the tows nearest the midplane, which is shown in Figure 6.2b. The extra deformation of the outermost tows is reduced by compacting one layer at a time, but the outermost tows will still experience more deformation than the interior tows. Creating realistic 3D textile models remains an important challenge for the research community. Tools like VTMS and a very similar software called DFMA provide a method for producing non-idealized textile geometries, which is a step in the right direction, but significant work remains to ensure that the geometries closely resemble physical specimens.

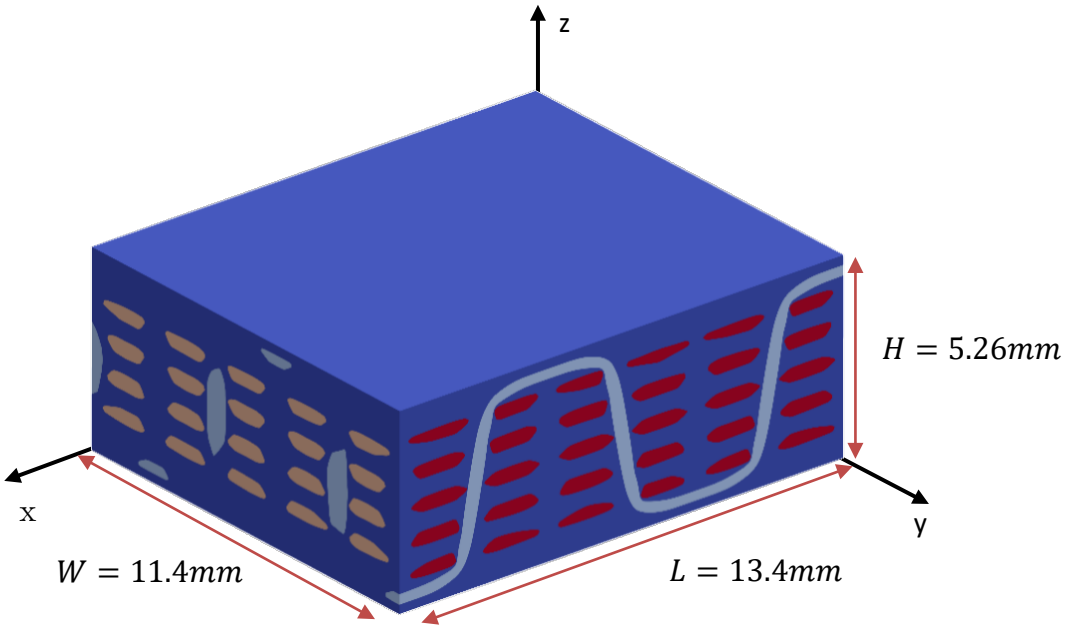

a) Entire analysis region

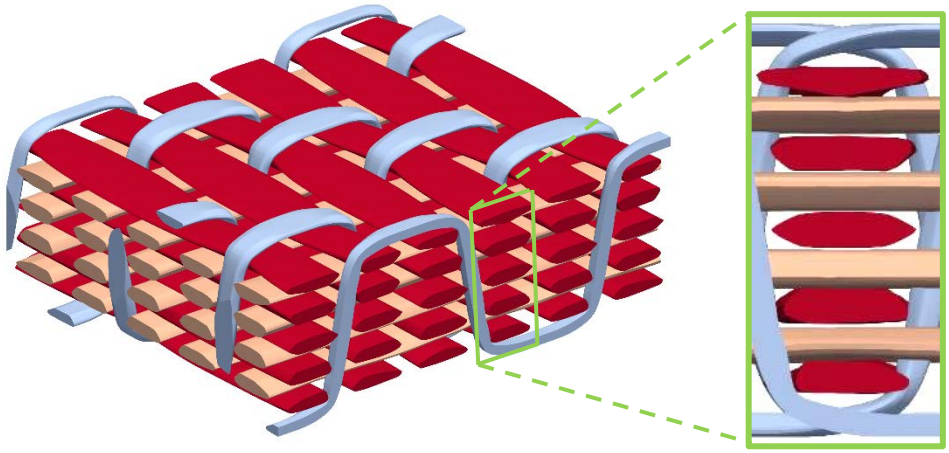

b) Analysis region with the matrix removed to show tows

Figure 6.2. Dimensions of 3D textile model (units shown in meters)

6.2.2 Textile Meshes

This work used a standard FEA formulation as described in chapter 3, which requires a conforming mesh. The VTMS software includes meshing for the independent mesh method [65], but the resulting mesh is not usable herein because it is nonconforming. Instead, an in-house tool was used to create refined surface meshes of the tow geometries. It should be noted that the surface geometry information that is taken from VTMS is faceted, resulting in unrealistic sharp corners between facets. Ideally, the faceted geometry should be smoothed to avoid unrealistic stress concentrations in the matrix material at the sharp corners, but this is not done in this study. At the level of mesh refinement used in this study, it is not clear whether the faceted geometry has a significant effect. The surface meshes were then used to create volume meshes of the tows and matrix using quadratic tetrahedrals, leveraging the general tetrahedral meshing library TetGen [39]. The algorithm used to create a conforming quadratic tetrahedral mesh for the tows and matrix from the tow geometry provided by VTMS is described in much greater detail in chapter 2 of this manuscript.

The in-house meshing tool used a cubic spline fit of the tow centroids to define each tow path. The material coordinate system is determined at each node within a tow mesh by finding the closest point on the spline of the tow path and using the tangent vector of the spline as the fiber direction, which defines the local x-axis. The local y-axis is assigned to be parallel to the cross product between the global z-axis and the tangent vector. The final local z-axis is defined by the cross product of the local x and y-axes, which defines a right-

handed coordinate system. For the analysis of the 1x1 orthogonal weave, tow paths were assumed to be 2D, neglecting the variation of the tow path in 3rd dimension. However, the in-house tool was enhanced to use general 3D splines to fit the tow paths for the study of the 2x2 orthogonal weave.

For the 1x1 orthogonal weave, the effect of mesh refinement is explored. The same geometry was used within VTMS to generate a sequence of three increasingly refined surface geometries, and from those surface geometries, a sequence of FEA meshes were generated. Some information about the sequence of meshes is shown in Table 6.1. It should be noted that because of the complex geometry, as the surface geometry from VTMS is refined the volume of the tows change, which is why the tow volume fraction changes for the meshes in Table 6.1. Figure 6.3 shows the sequence of meshes used for the 1x1 orthogonal weave. As Figure 6.3 reveals, the element size within the mesh was not uniform, with higher refinement existing very near the boundaries of the constituents, which is not ideal but is the result of the particular method used to generate the mesh. The mesh shown in Figure 6.3c is more refined than the majority of 3D textile studies in the literature and is close to the same refinement of the most refined models in the literature, such as those found in Ref. [21].

Table 6.1. Basic statistics of the meshes used in the study

Mesh	# Nodes	# Elements	Tow Volume Fraction
Coarse	61,906	42,306	53.9%
Medium	374,473	266,129	56.1%
Refined	1,962,148	1,436,552	56.7%

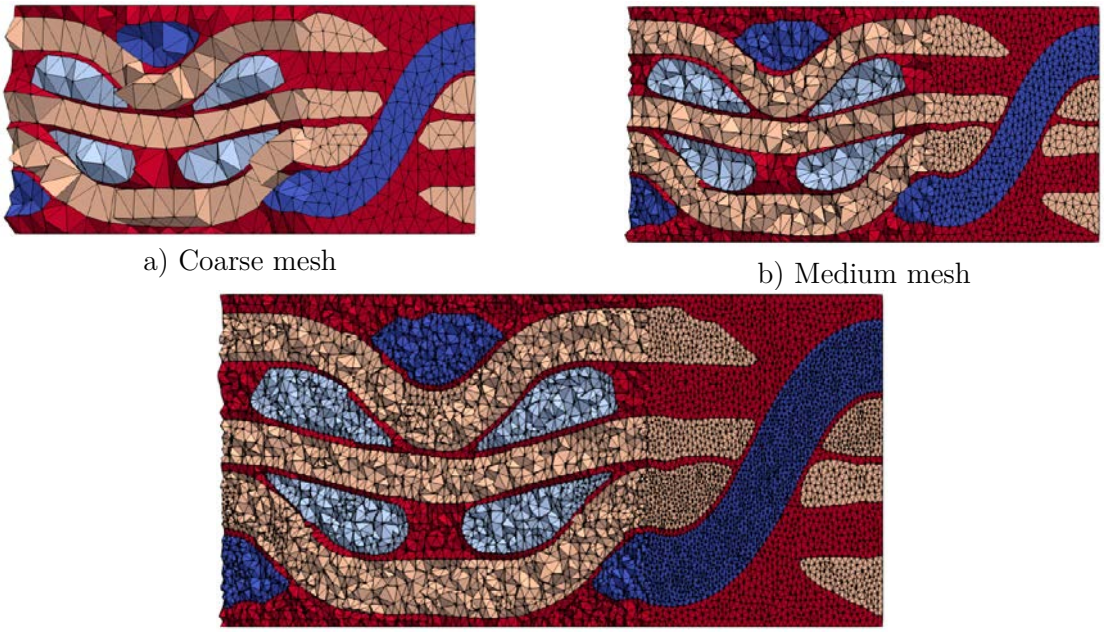

a) Coarse mesh

b) Medium mesh

c) Refined mesh

Figure 6.3. Clipped region of unit-cell showing the mesh refinement of 1x1 orthogonal weave

The mesh used for the 2x2 orthogonal weave was much larger due to the more complex tow architecture and larger analysis region. Figure 6.4 shows the typical mesh refinement. Tows generally have 5 to 10 quadratic tetrahedrals through the thickness of the tows. A higher mesh refinement exists near binder tows. The final mesh used for this paper consists

of 20 million nodes and 15 million quadratic tetrahedrals. It should be noted that the surface geometry that is taken from VTMS is faceted, resulting in unrealistic sharp corners between facets. Ideally, the faceted geometry should be smoothed to avoid unrealistic stress concentrations near sharp corners, but this is not done in this study. At the level of mesh refinement used in this study, it is not clear whether the faceted geometry has a significant effect.

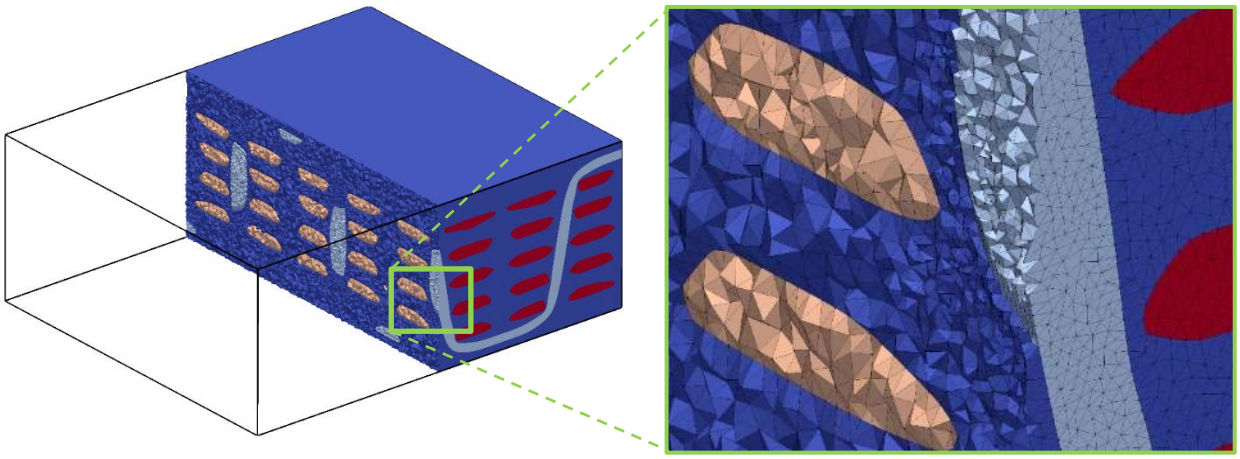

Figure 6.4. Illustration of mesh refinement for 2x2 orthogonal weave

6.2.3 Boundary Conditions

Ideally, periodic boundary conditions would be applied to the textile unit-cell, but the geometry from VTMS is not periodic since it creates a model by virtually simulating textile processing. Consequently, simple boundary conditions are used. Three faces of the analysis region (the $x = 0, y = 0$, and $z = 0$ planes) are assumed to be planes of symmetry. Refer to the coordinate system in Figure 6.2a. This results in the following boundary conditions:

$u(0, y, z) = 0$, $v(x, 0, z) = 0$, and $w(x, y, 0) = 0$. In addition to these boundary conditions, several additional boundary conditions impose the load depending on the configuration, which differs for the two types of textiles considered

For the 1x1 orthogonal weave, uniaxial tension along the global x-axis was the only load considered. To impose this boundary condition, the x-direction displacement on the $x = L$ face of the unit cell was specified such that the unit cell experienced 1% strain, where L is the in-plane dimension of the unit-cell. Figure 6.5 illustrates the boundary conditions used for the 1x1 orthogonal weave. The remaining two planes are traction free. These boundary conditions imply that the specimen is infinitely long but very thin and narrow.

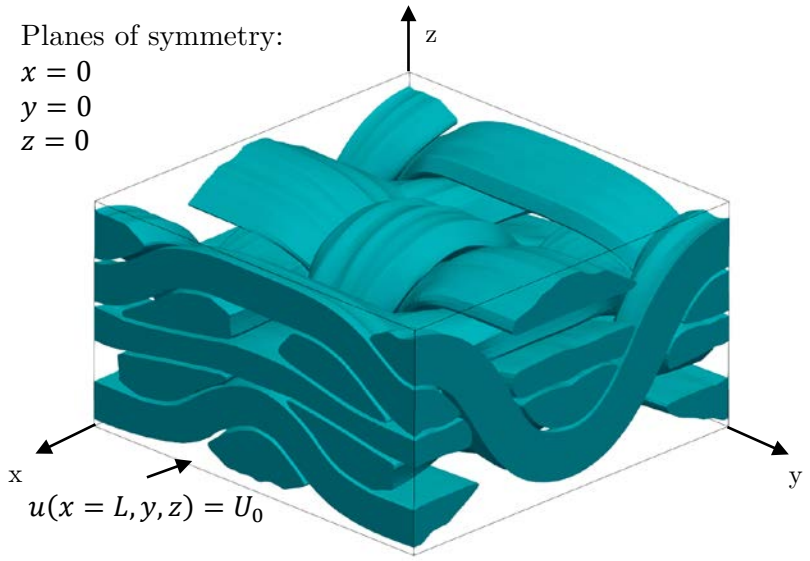

Figure 6.5. Boundary conditions used for 1x1 orthogonal weave

For the 2x2 orthogonal weave, three loading configurations are considered: uniaxial tension along the global x-axis, uniaxial tension along the global y-axis, and in-plane shear.

The loading consists of specified displacements to result in a 1% volume average strain. Specifically, this requires $u(13.4\,mm, y, z) = 0.134\,mm$ for tension along the x-axis, $v(x, y = 11.4\,mm, z) = 0.114\,mm$ for tension along the y-axis, and $u(x, y = 11.4\,mm, z) = 0.114\,mm$ and $v(13.4\,mm, y, z) = 0.134\,mm$ for in-plane shear. All other boundaries are traction free.

6.2.4 Clipped Analysis Region

As mentioned before, simple boundary conditions are used since the textile geometry lacks periodicity. However, the boundary conditions enforce the boundaries of the textile to remain planar, which is expected to induce artifacts near the textile boundaries. Though this issue is not addressed for the simple 1x1 orthogonal weave study, a strategy is adopted for the more realistic 2x2 orthogonal weave study to reduce the boundary effects. A larger section of the textile is modeled, the simple boundary conditions are applied to the larger analysis region, and a subregion is used for post-processing to reduce artifacts due to the boundary conditions. Figure 6.6 shows the full analysis region and the subset of the region used for post-processing. The clipped region used for post-processing is only approximately a full unit-cell, since the textile geometry is non-periodic. The in-plane boundaries the of the clipped analysis region are not planar since the tetrahedral elements are not cut by the clipping process. Any element with a node within the clipped analysis region is kept, while all others are removed, resulting in a rough surface on each clip plane, which can be observed in Figure 6.6b or Figure 6.6c.

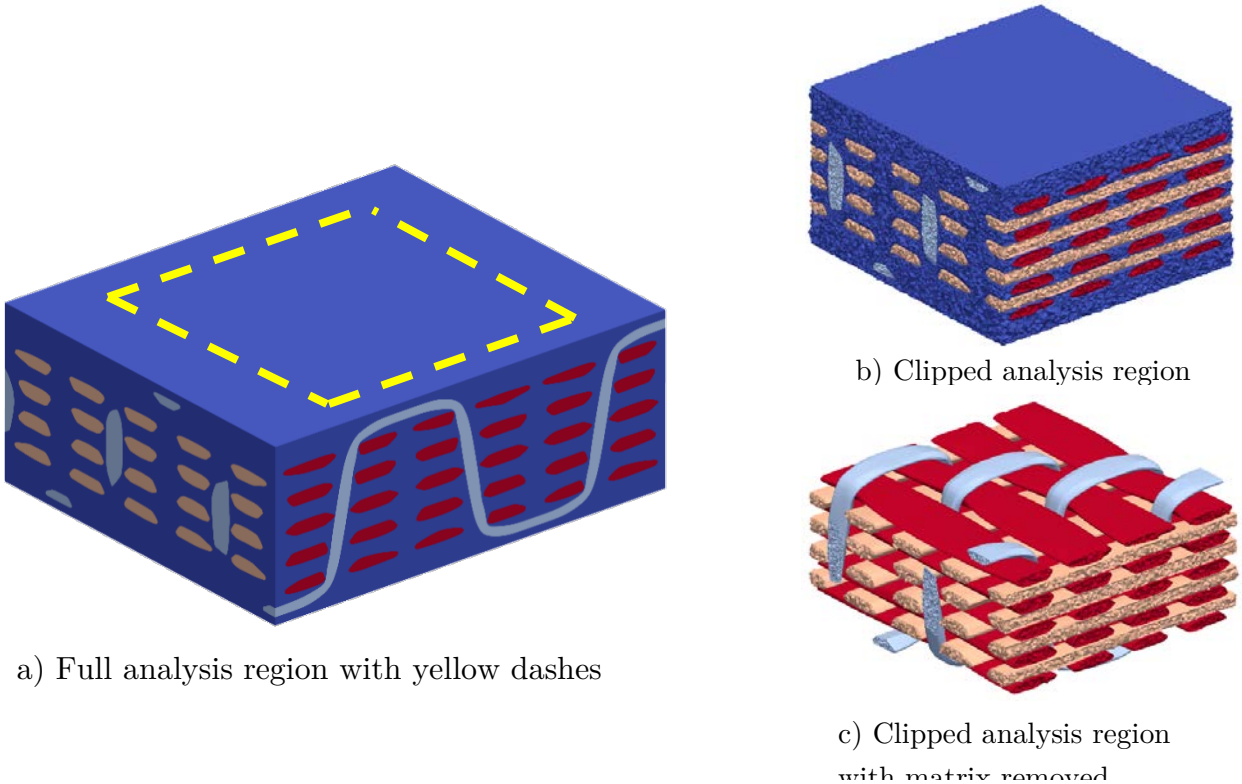

a) Full analysis region with yellow dashes

b) Clipped analysis region

c) Clipped analysis region with matrix removed

Figure 6.6. Illustration of the full analysis region and clipped region used for post-processing

6.2.5 Variation of Cross-Sectional Area

Each tow within VTMS will have a different total volume and a variation of the cross-sectional area along the tow path. This is due to a combination of several factors, including unrealistic space left between the digital chains that was not completely removed during the iterative relaxation process, the way the cross-sectional profiles for each tow are generated from the bundle of digital chains, the interpolation of cross-sectional profiles along the tow path, and the method used for the removal of interpenetrations between tows.

For real textile composites, there is a variation of the cross-sectional area, but this is due to the variation in the amount of matrix material forced out of parts of the tows during the manufacturing process. The same number of fibers should exist at any cross-section along a tow (except for a few broken fibers), so a variation of cross-sectional area would cause a change in the fiber volume fraction. Within VTMS, the variation of cross-sectional area that is caused by the creation of the cross-section profiles and subsequent interpolation along the tow path is not physical in any way. However, the locations where the cross-sectional area varies due to the removal of interpenetrations might correlate to areas where matrix is forced out of the tow during the manufacturing process, but there is no reason to expect the magnitude of the variation to match those seen in actual textile composites. Furthermore, the variation of the cross-sectional area would result in a variation of the material properties along the tow path, since the fiber volume fraction is changing. These concerns should be kept in mind for the future development of textile modeling tools if accurate textile composite models are to be achieved.

For the 1x1 orthogonal weave, a constant set of homogenized tow properties is used for the tows, but the effect of accounting for the variation in the local fiber volume fraction is explored for the 2x2 orthogonal weave. Though the variation of the local fiber volume fraction is not accounted for the 1x1 orthogonal weave, the study of the 1x1 orthogonal weave does explore the variation of the cross-sectional area and how the changing cross-sectional areas relate to the distribution of stresses within each tow. To quantify the variation of the cross-sectional area of each tow, it is useful to define a nominal cross-

sectional area. The volume of each tow is calculated by summing the volume of each element within the tow, which will be denoted by V_i for tow i. The length of each tow, L_i, is using the cubic spline used to fit the tow path. The normalized tow length, $\hat{L}_i$, is defined to be the actual tow length divided by the side length of the unit-cell, which is a characteristic length for the textile.

Since there are sixteen tows within the model used for the 1x1 orthogonal weave, it is helpful to breakdown the tows by type (binders, warps, and wefts), row (in-plane), and layer (out-of-plane and relative to the analysis region) to help with the discussion of particular tows. In the results sections, each type of tow is discussed separately. However, within the plots, the row will be identified by color, and the layer will be identified by line style. This categorization of tows is shown in Figure 6.7. Using this method of categorization, Table 6.2 shows the normalized tow lengths for each tow in the unit-cell.

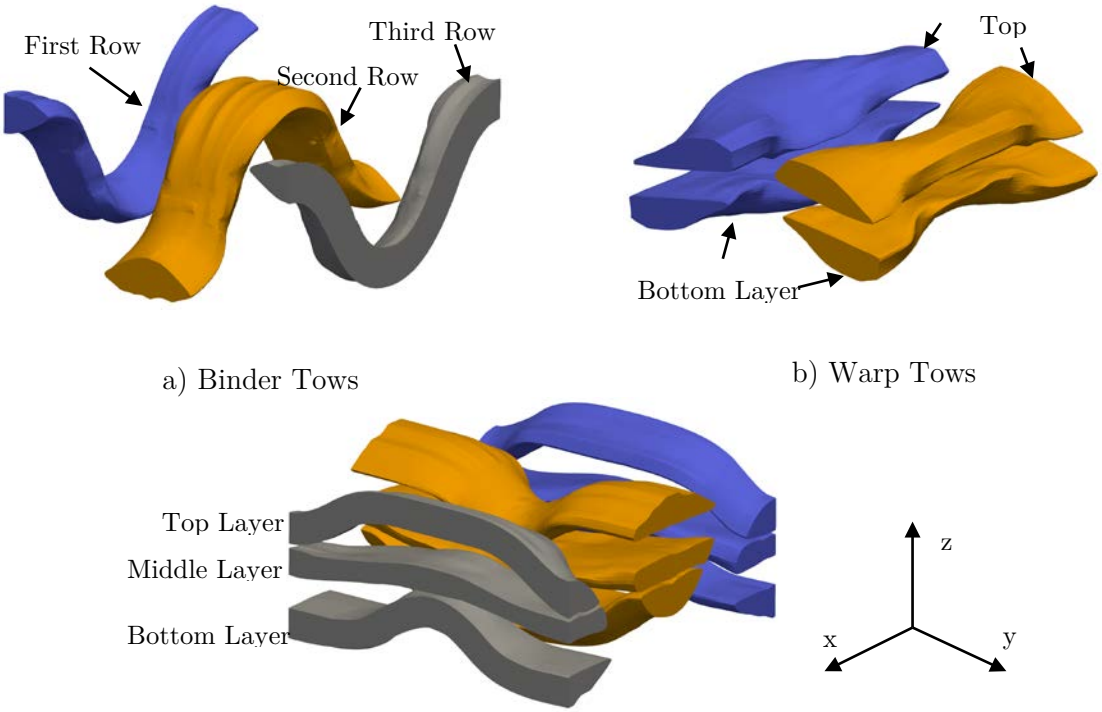

a) Binder Tows

b) Warp Tows

c) Weft Tows and Coordinate Axes

Figure 6.7. Illustration of tow positions and color scheme used in this paper

Table 6.2. Normalized length, $\hat{L}$, of tows by type, row, and layer

a) Binders

Row	$\hat{L}$
1	1.304
2	1.311
3	1.311

b) Warps

	$\hat{L}$	
Row	Top	Bottom
1	1.001	1.002
2	1.001	1.001

c) Wefts

	$\hat{L}$		
Row	Top	Middle	Bottom
1	1.069	1.009	1.074
2	1.065	1.010	1.073
3	1.060	1.007	1.065

Finally, the nominal cross-sectional area for each tow, $\bar{A}_i$, can be defined by the volume of the tow divided by the tow length and is given by Equation (6.1(6.1), where i is a free

index corresponding to the tow and V_i is the volume of the respective tow. The nominal cross-sectional area for each tow can be non-dimensionalized based on the size of the unit-cell using the Equation (6.2), where V_i is the volume of tow i , V is the volume of the unit-cell, and $\hat{L}_i$ is the normalized tow length. Table 6.3 shows the normalized cross-sectional area for each of the tows within the textile.

$$\bar{A}_i = \frac{V_i}{L_i} \tag{6.1}$$

$$\hat{A}_i = \frac{V_i}{V}\frac{1}{\hat{L}_i} \tag{6.2}$$

Table 6.3. Normalized nominal cross-sectional area, $\hat{A}$, of tows by type, row, and layer

a) Binders

Row	$\hat{A}$
1	0.0234
2	0.0458
3	0.0202

b) Warps

	$\hat{A}$	
Row	Top	Bottom
1	0.0451	0.0456
2	0.0450	0.0461

c) Wefts

	$\hat{A}$		
Row	Top	Middle	Bottom
1	0.0194	0.0215	0.0193
2	0.0438	0.0437	0.0446
3	0.0218	0.0200	0.0215

6.2.6 Material Properties

For the 1x1 orthogonal weave, the material properties for both constituents were taken from Ref. [66], which used an analytical homogenization scheme to predict the properties of the tows based on the measured properties of the matrix and graphite fibers. The matrix was assumed to be isotropic with the values: $E = 2.9\,GPa$ and $\nu = 0.3$. The tows were assumed to be homogenous and transversely isotropic with the values: $E_{11} = 221.4\,GPa$, $E_{22} = E_{33} = 12.6\,GPa$, $G_{12} = G_{13} = 7.4\,GPa$, $\nu_{12} = \nu_{13} = 0.34$, and $\nu_{23} = 0.32$. For this initial study of a 1x1 orthogonal weave, linear elasticity was assumed.

For the 2x2 orthogonal weave study, the tows are assumed to consist of IM7 graphite fiber and 5220-4 epoxy resin. In chapter 4 of this manuscript, multiple ensembles of random fiber/matrix models were used to predict the effective tow properties for different fiber volume fractions, and a constitutive model that is a function of the fiber volume fraction was developed, allowing the effective tow properties to be evaluated for fiber volume fractions between 40% and 80%. Chapter 4 contains plots of the engineering properties and effective stiffness matrix terms as a function of volume fraction for the constitutive model for a IM7/5220-4 material system. The study of a 2x2 orthogonal weave utilizes the developed constitutive model to study the effect of accounting for the variation of local fiber volume fraction on the locations and severity of stress concentrations.

Ideally, the strengths would be based on a series of experiments or microscale damage analyses. However, for this paper, the strengths will be assumed to be the same as the tows tested in Ref. [31] and are shown in Table 6.4. It should be noted that the strengths will

not be used to predict the progression of damage in the studies of this chapter. Instead, the strengths are used to assess the severity of stress concentrations, allowing a quantitative comparison of different stress states by dividing each component of stress in the local coordinate system by the respective assumed strength.

Table 6.4. Assumed tow strengths, taken from Ref. [31]

Warps and Wefts		Binders	
X_{11}	3234 MPa	X_{11}	2678 MPa
$X_{22} = X_{33}$	36.4 MPa	$X_{22} = X_{33}$	36.4 MPa
$X_{12} = X_{13}$	53.8 MPa	$X_{12} = X_{13}$	53.8 MPa
X_{23}	61.5 MPa	X_{23}	61.5 MPa

6.3 Analysis of a 1x1 Orthogonal Weave

This study investigates a simple 1x1 orthogonally woven textile, where the binders cross one weft at the top and bottom of the textile before traversing through the thickness of the textile. The aim is to explore how load is distributed throughout a orthogonally woven textile, investigate the effect of mesh refinement, and quantify the variation of the tow architecture and how it relates to locations of stress concentrations. The findings are used to guide the study of a 2x2 orthogonal weave later in this chapter.

The study of the 1x1 orthogonal weave is broken up into four sections. The first section focuses on the effect of mesh refinement on the predicted volume average stresses within the tows. Next, the stress distributions within the tows are explored and the effect of mesh refinement is quantified. The last two results sections investigate the most refined model with greater detail. The third section characterizes the tow architecture, including how the

cross-sectional area of the tows change throughout the model, while the last results section presents the average stresses along the tow paths and correlates the locations of high stresses with patterns in the weave geometry. In the final section of the 1x1 orthogonal wave study, a summary of the results is given.

6.3.1 Volume Average Stresses

The volume average stresses within each constituent provide an initial measure of mesh convergence and insight into how load is transferred within the textile. Table 6.5 shows each component of the volume average stresses in the material coordinate system within each constituent for the refined mesh. The stresses in the material coordinate system are denoted by σ'_{ij}. For the two coarser meshes, the percent difference from the volume average stress of the refined mesh is shown.

Table 6.5 can be used to understand the convergence of the three meshes, but it should be noted that the convergence includes convergence of geometry as well as mesh refinement. The geometry changes with increasing mesh refinement because the geometry can be captured more accurately with additional elements. Unfortunately, for this paper, there was no method implemented for separating these two effects.

For the normal stress along the fiber direction, σ'_{xx}, the coarse mesh predicted a volume average value within 6% for the binders and warps, which are both nominally along the direction of the applied load. The wefts experience a compressive stress in the fiber direction since they run transverse to the applied load and oppose the Poisson contraction of the unit cell. For σ'_{xx} in the wefts, the coarse mesh predicted a volume average stress that was

22.5% larger than the value predicted by the refined mesh. For the transverse normal stresses, σ'_{yy} and σ'_{zz}, within each tow, the medium and coarse mesh predicted a volume average stress that was close to the value predicted by the refined mesh, remaining within 6%. The coarse and medium meshes ranged from moderately under predicting to severely over predicting the volume average normal stresses. However, the coarse and medium meshes generally under predicted the volume average shear stresses for all three types of tows. The volume average shear stresses was close to zero, which would result in a large percent difference of the volume average shear stresses between the different meshes. Consequently, the volume average of the absolute value of shear stresses is used for comparing the meshes. The coarse mesh predicted volume average shear stresses up to 18% lower than the refined much, while the medium mesh differed from the refined mesh by less than 6%. Although the predictions appear to be converging, the refined mesh cannot be said to be converged, and even more refined meshes are needed to determine the converged values precisely. These results for volume average stresses highlights the need for careful mesh refinement in 3D textile models.

Table 6.5. Volume average stresses (in the material coordinate system) in each type of tow for the different meshes

Tow Type	Mesh	$\langle\sigma'_{xx}\rangle$	$\langle\sigma'_{yy}\rangle$	$\langle\sigma'_{zz}\rangle$	$\langle\lvert\sigma'_{xy}\rvert\rangle$*	$\langle\lvert\sigma'_{yz}\rvert\rangle$*	$\langle\lvert\sigma'_{xz}\rvert\rangle$*
Binders	Coarse	-1.1%	-6.0%	-3.2%	-17.7%	-3.7%	-4.4%
	Medium	0.5%	-1.8%	-0.8%	-5.9%	-0.1%	-0.8%
	Refined	**2.37e8**	**-5.70e6**	**2.96e7**	**4.77e6**	**8.17e6**	**3.90e7**
Warps	Coarse	5.6%	-3.8%	23.1%	-17.4%	-3.3%	-4.0%
	Medium	1.6%	-0.7%	16.2%	-5.8%	0.0%	-0.7%
	Refined	**1.57e9**	**8.19e6**	**-1.77e6**	**3.06e6**	**5.25e6**	**2.51e7**
Wefts	Coarse	22.5%	-3.5%	-1.6%	-16.2%	-1.9%	-2.6%
	Medium	7.6%	-0.5%	0.6%	-5.4%	0.4%	-0.3%
	Refined	**-9.23e6**	**7.54e7**	**-2.20e7**	**2.08e6**	**3.57e6**	**1.70e7**

* Note: volume average shear stresses were calculated using the absolute value of the stress, since the volume average shear stress would be close to zero.

In addition to convergence, the volume average stresses shown in Table 6.5 indicate how load is transferred within the textile. Since the warps are aligned with the load direction, they experience the highest level of stress, as expected. In the material coordinate system of the warps, the local z-axis generally aligns with the global z-axis, which is out of plane of the textile. Consequently, the local x- and y-axes generally align within the plane of the textile. Table 6.5 shows that the warps experience tension in the two in-plane directions and compression out of plane. In the warps, the fibers align with the applied load, so the tension along the local x-axis is expected. Regarding the tension within the warps in the direction transverse to the fibers and in the plane of the textile (local y-axis), the unit cell contracts due to the Poisson effect, but the wefts are stiffer than the other constituents in the transverse in-plane direction and have a smaller Poisson's ratio. Consequently, the wefts experience compression while the warps experience tension. The compression in the out of plane direction is due to the binders trying to straighten due to the applied

extensional load in the direction of the binder's undulation. The binders interlock with the wefts and compress the warps and wefts, while the binders experience a tensile stress out of plane. This was confirmed by calculating the volume average stress within the binders in the global z-direction, which was indeed tensile. Additionally, in a study not shown in this paper, the binders were replaced by matrix. For that case, the warps experience tension instead of compression in the out of plane direction, indicating it is the binders trying to straighten that leads to the compressive stresses in the warps. It is important to remember that the volume average stress in the entire unit cell is zero except in the direction of the applied load.

Additionally, the warps experience relatively large shear stresses in every component. Since the warps are carrying most of the load in the direction of the applied load and matrix pockets exist in regions of the textile, the load is generally distributed to the neighboring tows through shear stresses.

Since the wefts lie transverse to the direction of the applied load, they experience a relatively high σ'_{yy}. The compressive stresses for the other two normal components were discussed earlier. Of the three types of tows, the wefts generally experienced a less severe volume average shear stress.

6.3.2 Stress Distribution within Tows

This section investigates the stress distribution of the most severe stress component within each type of tow. Though accurate strengths were not available for the particular material system assumed in this paper, the relatively strengths for the six components of

stress were assumed to be the same as for the material system in Ref. [31]. The strengths will not be used herein to predict failure, but rather, they will only be used to select the most severe stress component to investigate for each type of tow. To select the most severe component of stress, the maximum value for each component of stress was compared against the allowable stress.

Binders

For the binders, σ'_{zz} was clearly the most severe component of stress. Figure 6.8 shows the percentage of the constituent within various stress ranges for each of the mesh refinements. From the column plot in Figure 6.8, even the coarsest mesh (green bars) captures the amount of material at relatively low stresses well. However, it poorly predicts the material at severe tensile and compressive stresses. The binders experience a tensile σ'_{zz} on average, as shown in the figure and Table 6.5, but locally, the stress varies from severe compressive to severe tensile stresses. In the contour plots in Figure 6.8 and subsequent figures, the color mapping is centered on the volume average stress within the constituent, which means that shades of red indicate a stress larger than the volume average, shades of blue indicate a stress lower than the volume average, and gray indicates a stress near the volume average. The binders experience compression where they cross over or under a weft, very near the area of contact between the two tows. A typical location is labeled A in Figure 6.8. Compressive stresses develop in these regions because the binders want to straighten out but are resisted by the wefts. Consequently, compressive stresses form where the binder presses against the wefts. However, within the matrix rich regions that lie between the

rows of wefts, the binders generally experience a tensile σ'_{zz}, as shown in region B of Figure 6.8. In these regions, the local z-axis is nearly parallel to the global x-axis, which is the direction of the applied tensile load. Therefore, the binders experience a tensile σ'_{zz} in these matrix rich regions.

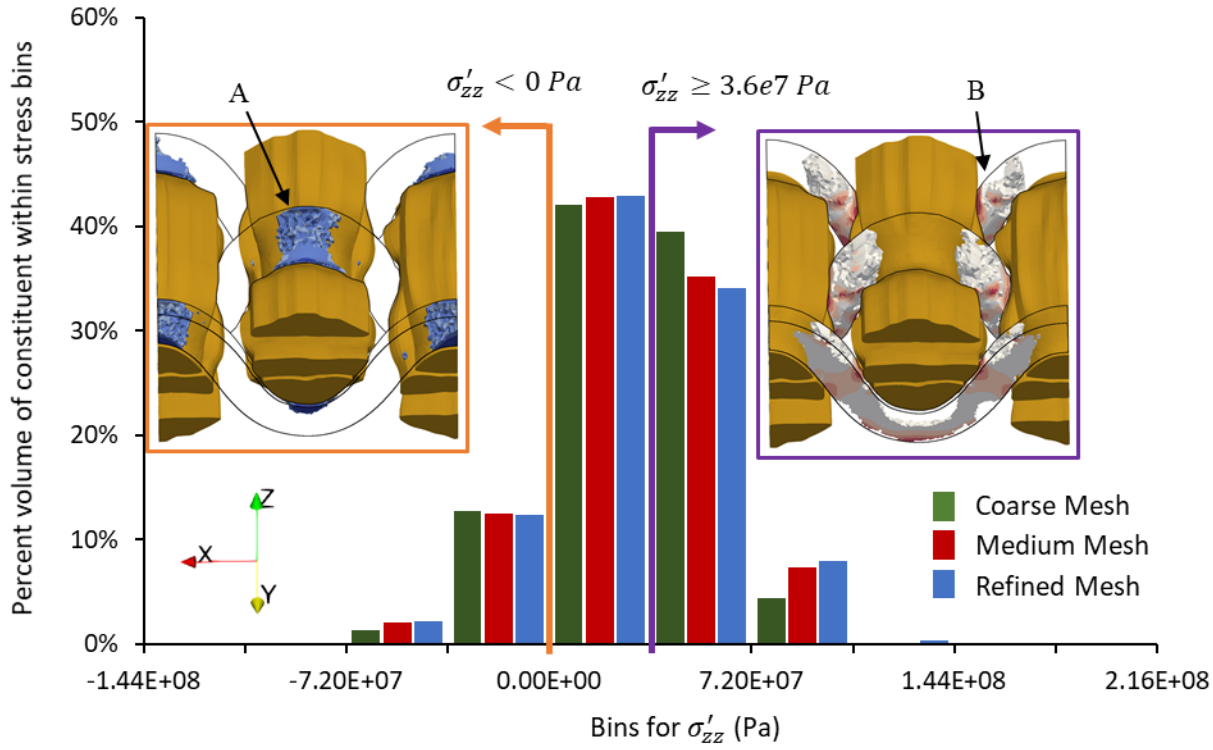

Figure 6.8. Distribution of σ'_{zz} within binders for the different mesh refinements with images showing material below 0 Pa or above 3.6e7 Pa. For the contour plots, gray indicates a stress near the volume average in the constituent, while red and blue indicates a stress higher or lower respectively.

Figure 6.8 focused on the entire distribution of σ'_{zz}, but the highest stresses are of critical importance to understand where damage might initiate. A large tensile σ'_{zz} tends to cause matrix cracks within the tow. Because of this concern, it is important to consider the

highest stresses. Figure 6.9 shows the stress distribution and locations of the material within the tow with a σ'_{zz} higher than 8e7 Pa. Data labels are shown for the refined mesh, since the volume percentage in the last seven bins are difficult to compare due to the very small values. Concerning convergence, the coarse mesh does a very poor job at predicting any material at these high stress levels and only predicted that the stress levels of the first two bins are reached. The medium mesh somewhat captures the stress distribution for the first four bins, but it fails to predict any material reaches the stress levels of the last five bins. As expected, increasing mesh refinement seems to be headed towards convergence, but even further mesh refinement is needed to know if the converged maximum stress is reached with the refined mesh used in this study. The location of the highest σ'_{zz} in the binders is of interest since damage is likely to initiate at those concentrations. The contour plot in Figure 6.9 shows the material with a σ'_{zz} higher than 8e7 Pa. The highest values generally exist in the matrix rich regions between rows of wefts, as shown in the contour plot in Figure 6.9. However, many of the locations of high stresses also lie near relatively sharp edges of the tow. Because of the nature of the geometry generation algorithm in VTMS, each tow's surface geometry is faceted, not smooth. These artifacts can cause higher stresses where the facets form sharp edges, but further study is needed to quantify the effect of the faceted surfaces. Additionally, the smallest stress concentration regions can fit within a sphere that has a diameter that is 1/100th of a tow thickness or less, which is much less than a fiber diameter. Even if the model converges at a higher refinement, at that scale, the heterogeneous microstructure of the tow would need to be considered.

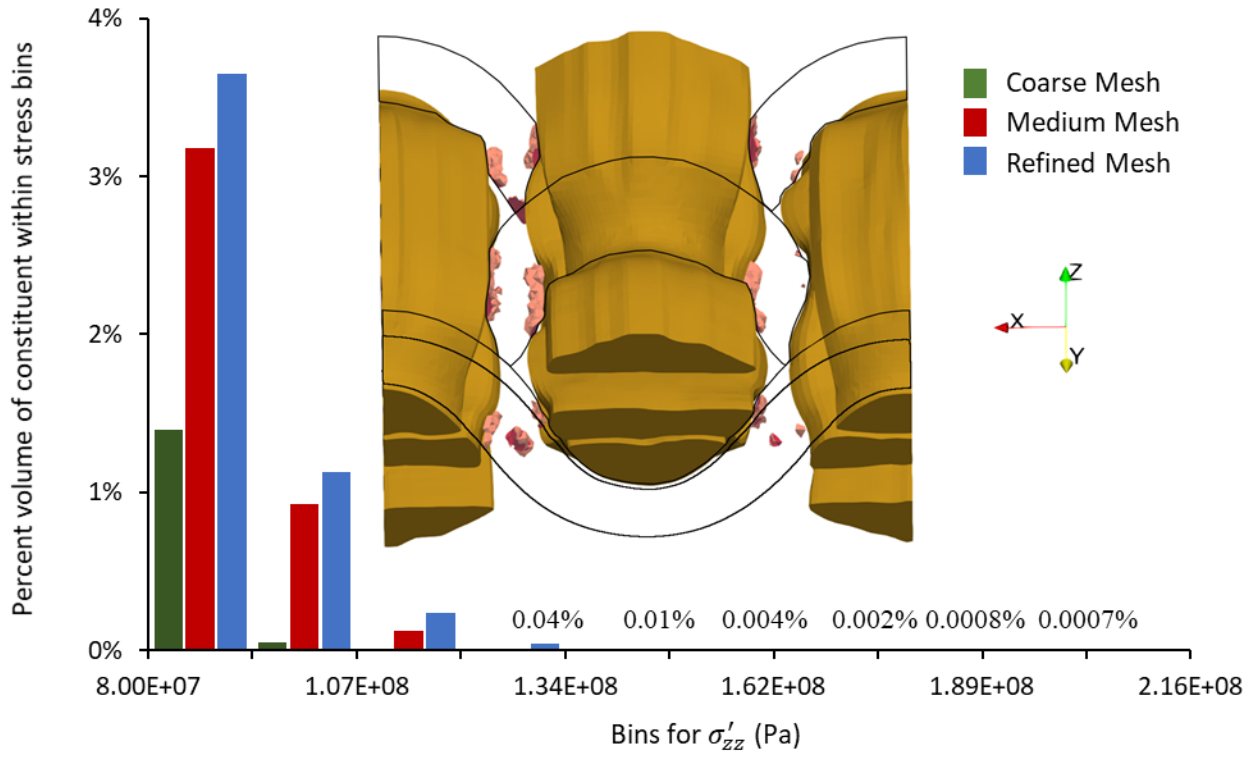

Figure 6.9. Distribution of σ'_{zz} at high stress levels within the binders for the different mesh refinements with an image showing material above 8e7 Pa.

Warps

Using the same methodology as described in the beginning of the previous section, the longitudinal shear stress was selected as the most severe stress component for the warps. The magnitude of the stress along the fiber direction was very high, but the tensile strength of tows along the fiber direction is often at least an order of magnitude greater than the shear and transverse tensile strengths. Of the different types of tows, the warps experiences stress that were least severe compared to the allowable stresses.

The stress by volume histogram is not shown for all values of σ'_{xz} in this paper, but the data showed that all three meshes predicted very similar distributions for low levels of

stress. However, the high stresses are of much greater interest. Since this paper is concerned with the severity of the shear stress, the sign is not important for this discussion, so the absolute value of the shear stresses will be shown. Figure 6.10 shows the percentage of the warp that reach stresses larger than 1.76e8 Pa. A relatively large range of stresses is shown in Figure 6.10, and since the coarser mesh under predicts the amount of volume at the very high stresses, it over predicts the amount of material at a lower stress, represented by the first bin. The coarse mesh poorly predicts the stress distribution for the last eight bins. The medium mesh seems to capture the trend well until the last three or four bins. As with σ'_{zz} in the binders, the maximum stress in the constituent continues to increase as the mesh is refined. It is unclear if the refined model has reached sufficient convergence.

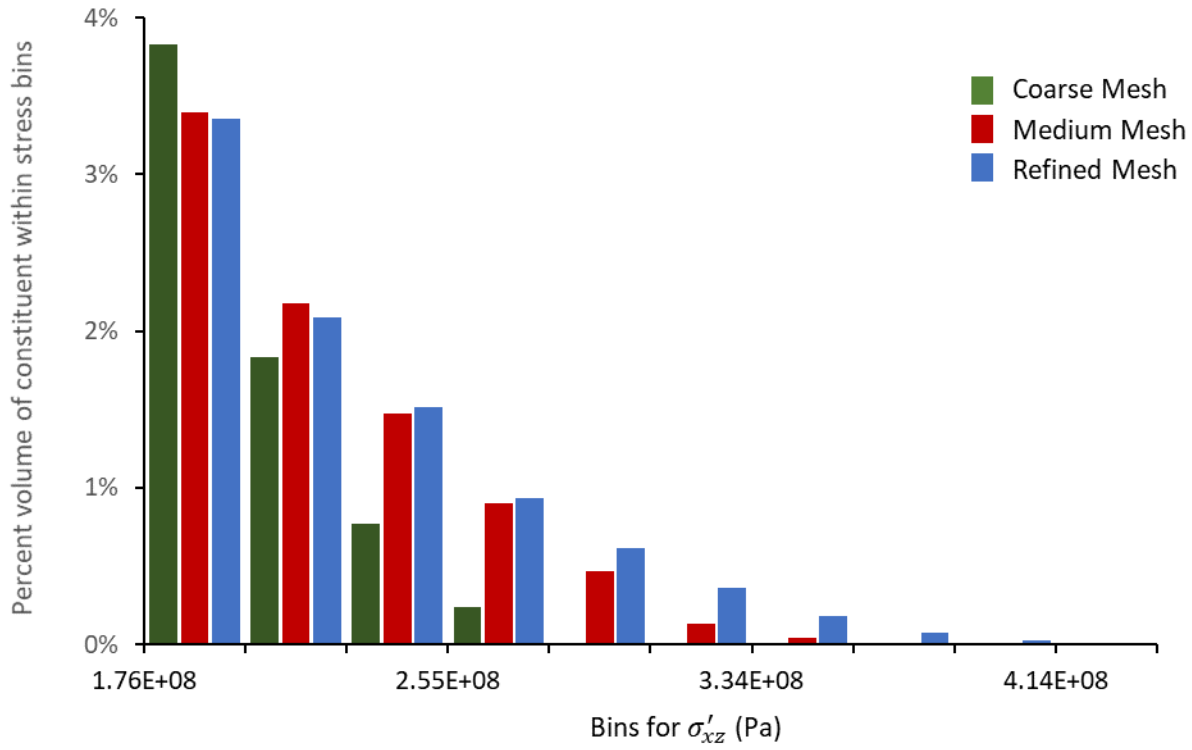

Figure 6.10. Distribution of σ'_{xz} at high stresses within the warps for different mesh refinements

The location of the high shear stresses provide insight into where damage might occur within the warps, though an accurate location of damage initiation cannot be determined without a progressive failure analysis since damage in other constituents will affect the stress distribution in this constituent. Figure 6.11 shows contours for the portion of warps with magnitudes of σ'_{xz} higher than 8e7 Pa, which includes a wider range of stresses than what is shown in the histogram of Figure 6.10. Since the warps generally align with the global coordinate system, the local x-axis is generally along the loading direction and the local z-axis is generally out of plane. Considering this, it seems that load is being transferred in the matrix rich regions to the neighboring binders, which have their greatest out of plane component in the same regions, since the concentrations all remain close to the area where the binders cross the warp paths, as shown in Figure 6.11a. Additionally, the shear stress concentrations are generally closer to the midplane of the unit cell, as shown in Figure 6.11c and Figure 6.11d. However, the high stress seems very localized to the outermost region of the tow, as shown in Figure 6.11c.

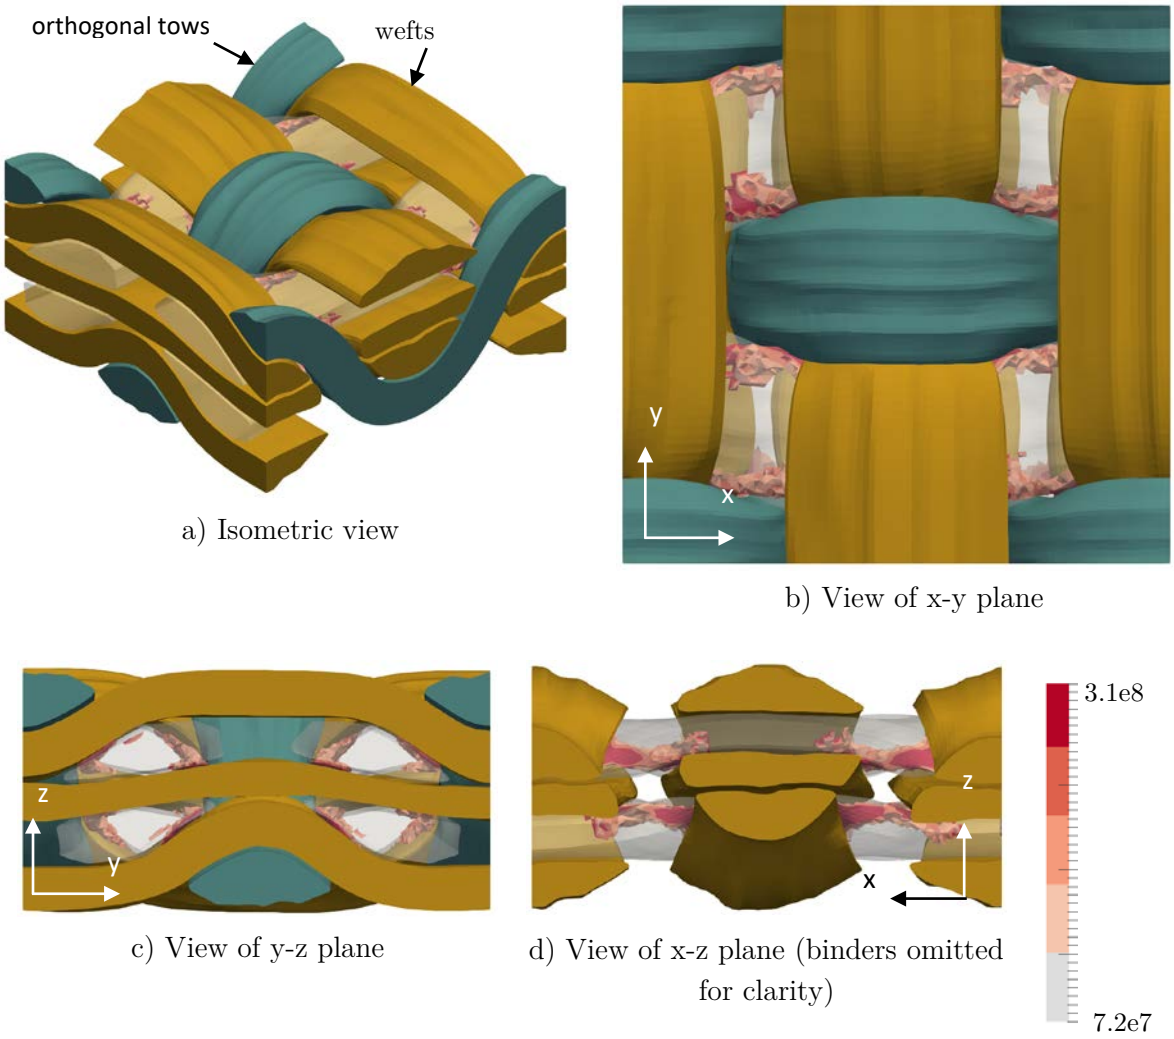

a) Isometric view

b) View of x-y plane

c) View of y-z plane

d) View of x-z plane (binders omitted for clarity)

Figure 6.11. contours for regions of the warps with $abs(\sigma'_{xz}) > 8e7\ Pa$. Binders and wefts are shown for context. The entire warps are also overlaid as a semi-transparent gray solid for context.

Wefts

Since the fibers run transverse to the applied load for the wefts, the transverse normal stress σ'_{yy} for the wefts reached high values, and the allowable for this component of stress

will be significantly less than the allowable in the direction of the fibers. Consequently, σ'_{yy} was the most severe of all the stresses within any type of tow by a large margin.

Figure 6.12 shows a histogram of the amount of material in the wefts within various bins of values for σ'_{yy}. A significant percentage of the material experiences relatively high tensile stresses; however, the tows do not carry that load evenly. The middle layer of wefts carries substantially more of the load than the top and bottom layers, as shown in Figure 6.12b. This is due to the middle layer of wefts being surrounded by two layers of warps, which are carrying most of the load within the textile for the applied load considered in this paper. Whereas, the top and bottom layers of wefts have less avenues for load to be transferred to them with only one nearby layer of warps, and thereby experience lower σ'_{yy}.

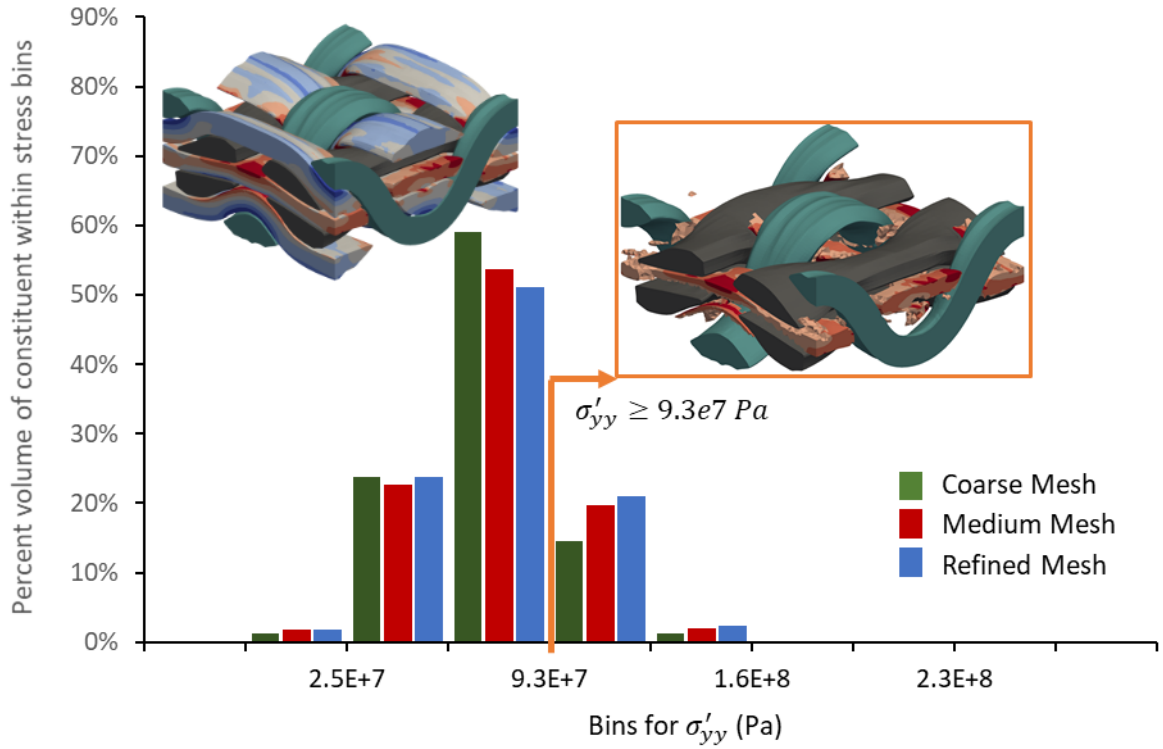

Figure 6.12. Distribution of σ'_{yy} within the wefts. For the contour plots, gray indicates a stress near the volume average, while red and blue indicates a value greater or less than the volume average respectively.

The locations of the highest stress concentrations are important, since damage will likely initiate in these locations. To investigate where damage might occur, the material experiencing a value of σ'_{yy} higher than 1.27e8 Pa, which includes the last five bins shown in Figure 6.12, is shown in Figure 6.13. Since there is a pattern of locations for these concentrations, Figure 6.13a only shows half the unit cell and illustrates two key views that show the locations well. The material within the respective bounding boxes in Figure 6.13a is shown in Figures 8b and 8c, which allows a clear view of the stress concentrations.

Figure 6.13 shows that even though the middle layer of wefts experiences much higher stresses on average, as shown in Figure 6.12b, all three layers through the thickness of the textile experience significant stress concentrations. In the top and bottom layers, stress concentrations occur just before and after a weft crosses an binder, near the outermost edges of the tow cross-section. One such location is highlighted as region A in Figure 6.13b. In the middle layer, stress concentrations occur in the regions where an binder reaches its peak or trough. One such location is highlighted as region B in Figure 6.13b. This leaves a matrix rich region in the center of the textile away from the binder, as shown in Figure 6.13b and Figure 6.13c, and as a result, the wefts carry more of the load locally in the absence of other fiber reinforced constituents nearby.

For the uniaxial extension considered in this paper, it seems that damage will develop initially within all three layers of wefts. However, the middle layer of wefts will likely experience a more rapid progression of damage since the average stress in the layer is significantly higher than the other two.

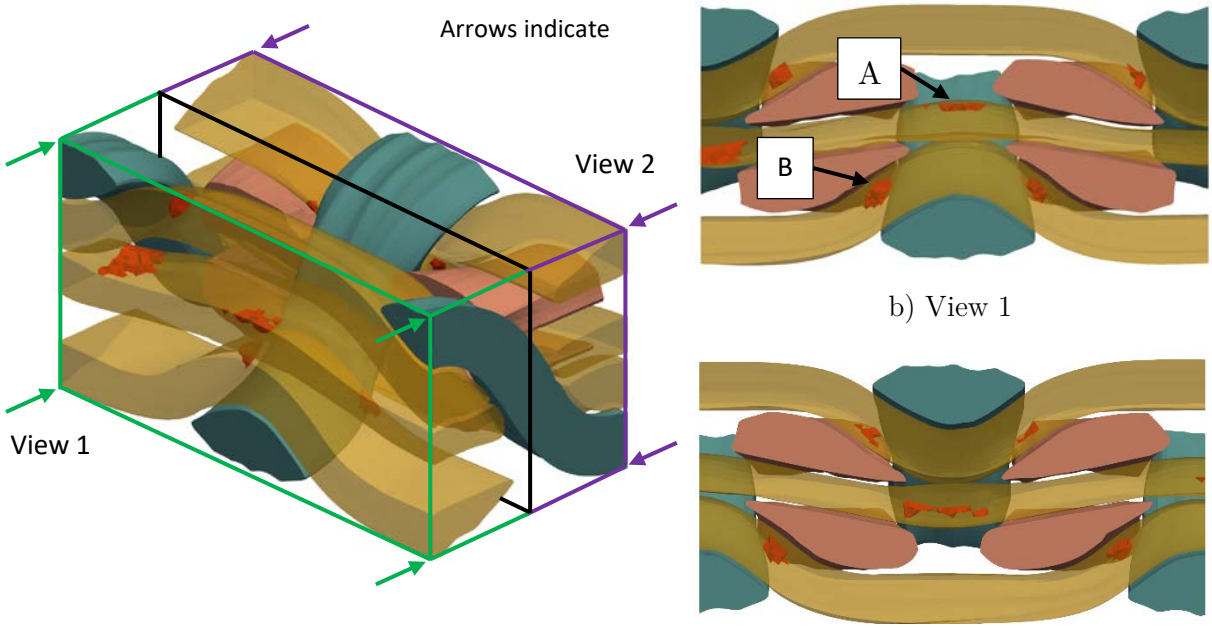

a) Half of the unit-cell and illustration of views

b) View 1

c) View 2

Figure 6.13. Regions of highest σ'_{yy} (highlighted in red) within the wefts

6.3.3 Characterization of Tow Architecture

As described in an earlier section concerning the configurations, the cross-section of the tow can vary along the tow path for the models generated by VTMS, due to a variety of reasons mentioned earlier. This first section aims to quantify and investigate this variation to characterize the tow architecture. The variation of the cross-sectional area is quantified as the difference between the actual cross-sectional area and the nominal value calculated using Equation (6.1). For each type of tow the variation of the cross-sections will be plotted and compared to understand where within the unit-cell that variations occur and how similar the tows are within each type.

Binders

Figure 6.14 shows the variation of the cross-sectional area for the binders. The paths of the binders are also shown for context of where variations occur within the textile. It should be noted that the horizontal axis of the plot is the curvilinear distance along the tow path, s, normalized by the total length of the tow, not the Cartesian coordinate. For Figure 6.14 and the subsequent figures shown in this paper, the color indicates the row of the tow within the unit-cell. Refer to Figure 6.7 for the location of these tows.

The 1x1 orthogonal weave considered is composed of layers of tows nominally parallel to the x- and y-axes and binders that interlock with the outer layers of wefts. In the regions where the binders interlock with the wefts, the cross-sectional areas were the smallest. Conversely, the largest cross-sectional areas occur in the matrix pocket between rows of wefts. During the generation of the geometry within VTMS, the digital chains within each bundle, which represents a tow, begin very spaced out. Tension is applied to each chain and an out-of-plane force compresses the unit-cell to arrive at a relaxed configuration. By visual inspection of the user, it was observed that more space between digital chains remained after the iterative relaxation process in regions where tows do not come in direct contact. Additionally, interpenetrations between the tows occur where the tows come in direct contact, and the method used for removing the interpenetrations results in the removal of some volume of the tows in these regions. The locations of cross-sectional area extrema are related to these steps of the VTMS process for creating the textile geometry.

There are some differences in the variations observed for the three binders in the unit-cell, but the trends are similar. The cross-sectional areas along the binder paths varied within -6.8% and 5.3% of the nominal values. Individual fibers have approximately the same cross-sectional area regardless of the weaving process, so a variation in the cross-sectional area will result in a change in the local fiber volume fraction. From the smallest cross-sectional area to the largest area, the fiber volume fraction would decrease by 11.5%, which is a significant change in the fiber volume fraction and would result in a very significant change in the elastic properties. This study did not account for effect on the material properties, but this result shows that the effect is significant and should be accounted for in future studies.

In addition to the high amplitude changes in cross-sectional area, there are smaller variations of the cross-sectional area with a much higher frequency. These variations are due to a number of factors, including possible error from fitting the chains to obtain the cross-sectional profile, interpolating the profiles, and the two factors already discussed involving removal of interpenetrations and the space between digital chains. This paper does not separate these effects, but understanding what causes these variations is important for creating realistic geometries.

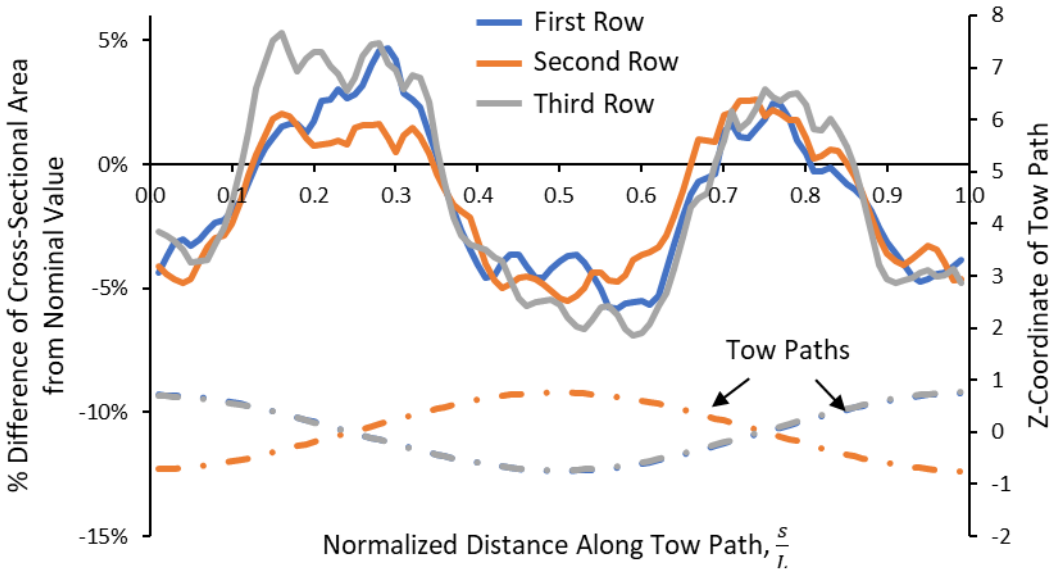

Figure 6.14. Variation of the cross-sectional area for the binders (tow paths shown for context)

Warps

Figure 6.15 shows the variation of the cross-sectional area along key regions labelled for context of where the variations occur. The positions of the tows are identified using the row and layer numbers (see Figure 6.7). The warps exhibited more variation in the cross-sectional area than the binders, varying between -7.92% and 10.08% of the nominal values. This range of variation in the cross-sectional area would result in a 16.4% change in the local fiber volume fraction. The peaks in Figure 6.15 correspond to matrix rich regions in the unit-cell between rows of wefts, and the valleys correspond to regions where the warps cross wefts. Much like the observations for the binders, the digital chains within the fiber bundles had more space between them in the matrix rich regions, resulting in a larger cross-sectional area. In addition to less space existing between the digital chains where the warps

come into contact with the wefts, interpenetrations between the tows occurred in these regions resulting in volume being removed from the tows. For these two reasons, the cross-sectional area was significantly lower where the warps cross the wefts. Though the locations of reduced cross-sectional area might correlate to the characteristics of a physical textile composite, the magnitude of variation of the cross-sectional areas is likely much larger than what would be observed in an actual textile. Detailed experimental measurements are obviously needed.

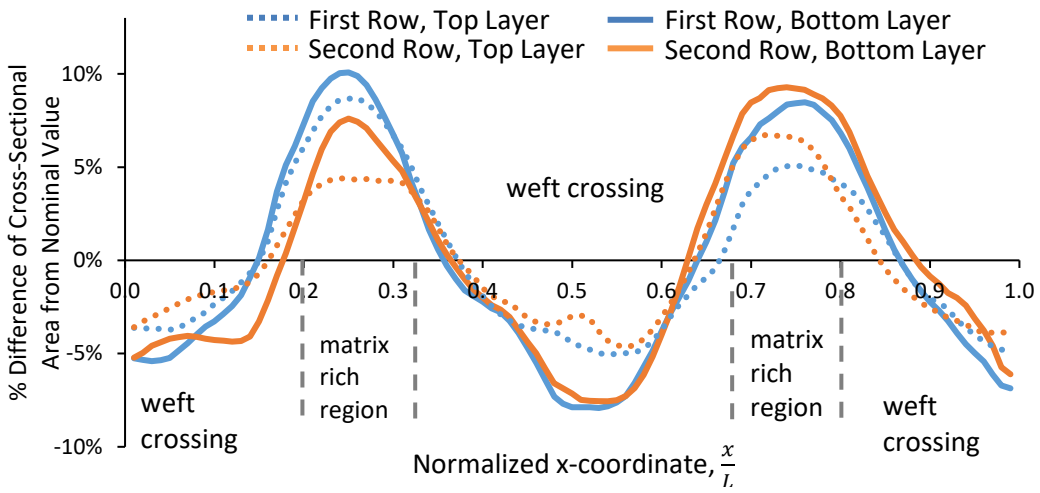

Figure 6.15. Variation of the cross-sectional area for the warps along with a view of unit-cell aligned with horizontal axis of plot for context

Wefts

For this type of tow architecture, the wefts can be separated into two groups: the middle layer of the analysis region and the outer layers of the analysis region. Within the weave, binders interlock with the two outer layers of tows, and since the binders undulate in two

alternating phases, which can be seen in Figure 6.16, all of the wefts will experience a similar response along the tow paths in one of two phases. The first row and third row of the top layer have similar tow paths to the middle row of the bottom layer, while the first row and third row of the bottom layer have similar tow paths to the middle row of the top layer. Between those two groups, the phases differ by half the length of the unit-cell, $0.5L$. Both outer layers of wefts experience a significant undulation due to the interlocking of binders.

The middle layer of wefts in the analysis region experiences relatively little undulation since they are not directly in contact with an binder, which is shown in Figure 6.16. As the binders try to straighten out during the generation of the tow geometry in VTMS, the wefts in the middle layer of the analysis region tend to flatten out. However, all three tows in the middle layer will experience a similar variation of the cross-sectional area in the same phase.

Using this information, the variation of the cross-sectional area is plotted for the two groups of tows, adjusting for the different phases in the outer layers. For every one of the tows in the outer layer, the cross-sectional area is plotted for half the length of the unit-cell starting at the center of when a binder crosses a weft. The paths for two of the wefts are illustrated by the arrows in Figure 6.16b and can be separated three sections: a region where the binder crosses the weft path, labeled A in Figure 6.16a; a region where the warp crosses the weft path, labeled B in Figure 6.16a; and finally, a matrix rich region, labeled C in Figure 6.16a.

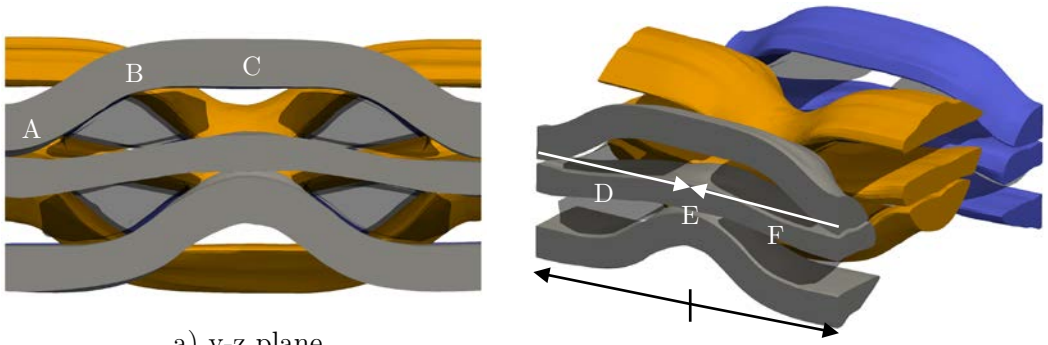

a) y-z plane

b) Isometric view

Figure 6.16. Views of wefts colored by row with labelled regions of interest (warps are shown with 30% opacity for context)

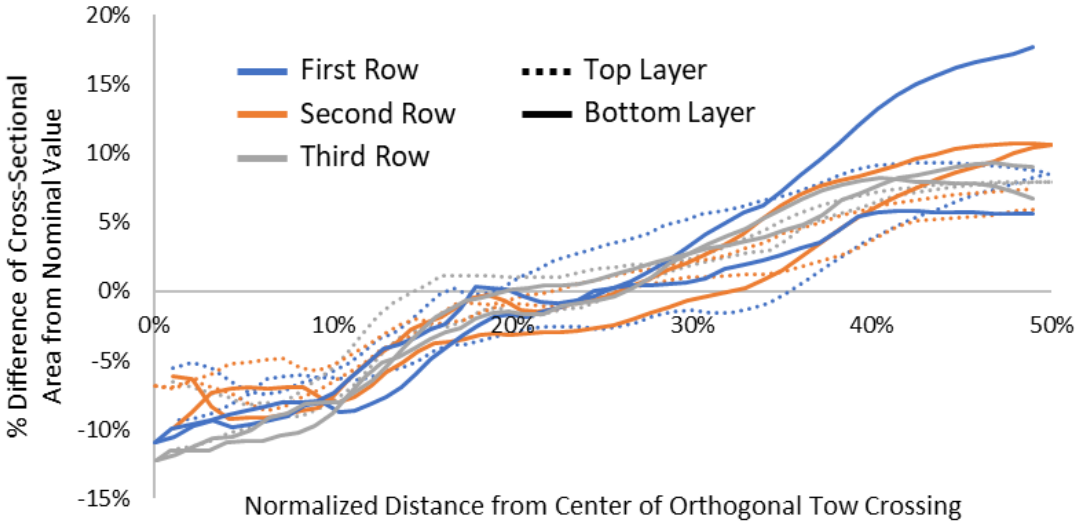

Figure 6.17. Variation of cross-sectional area for outer layers of wefts plotted vs. the distance from the center of a binder crossing normalized by the length of the unit-cell

Figure 6.17 shows the variation of the cross-sectional area for the outer layers of the wefts along the paths described, which adjusts for differences in phase. The positions of the tows given the row and layer denoted in the plot can be found in Figure 6.7. It should be noted that Figure 6.17 will contain two paths for each tow, both starting from an binder crossing. The variation of the cross-sectional area is greater than the variation seen in the orthogonal and warps, falling between -13% and 18% of the nominal values, which would result in a 26.3% change in the fiber volume fraction. For the tow on the first row and bottom layer, the region closest to the origin of the unit-cell significantly differed in cross-sectional area than the other tows, but this is the outlier and the region closest to the origin. The cause for this deviation is unclear. Other than this region and particular tow, the variation of the cross-sectional area for all tows remained within $\pm\, 5\%$. The lowest cross-sectional areas lie in regions where the binders interlock with the wefts, and the largest cross-sectional areas lie within the matrix rich regions. Where the warps cross the wefts, the cross-sectional area is between the values for the other two regions.

Figure 6.18 shows the variation of the cross-sectional area for the middle layer of wefts. In the middle layer of tows, an binder comes close from the top or bottom (alternating every L/2), and the undulation experienced by the tow on the second row will be out of phase by L/2 compared to the other two. However, the effect of binder being nearby on the cross-sectional area should be the same whether it lies above or below the middle layer weft. Therefore, even though the phase of undulation is different, there should be no difference in the variation of the cross-sectional area. Figure 6.18 shows that the general trends are

similar between all three middle layer of tows, but there are some large differences in the magnitude of variation. Even the first and third rows, which experienced the most similar tow paths of the three, had differences of up to 8%. The largest difference between any of the three was about 10%.

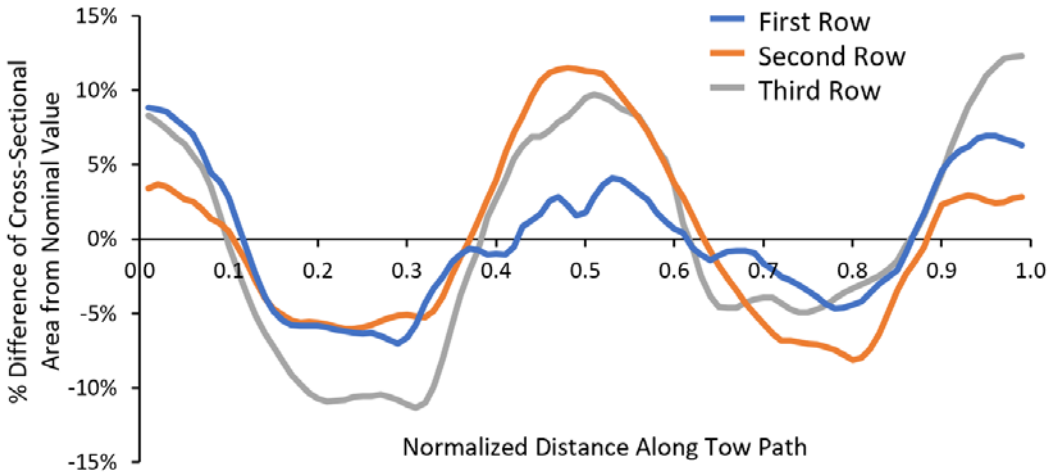

Figure 6.18. Variation of the cross-sectional area for the middle layer of wefts

Overall, the cross-sectional areas varied between -11.3% and 12.3% for the middle layer of y-direction of tows. With the exception of the one outlier in Figure 6.17, the middle layer of wefts experience about the same change in cross-sectional area as the outer layers of wefts, despite the fact that they were relatively straight compared to the outer layers. The amount of undulation does not seem to be important when considering how much the cross-sectional area varies. Considering the cross-sectional area for all the categories of tows, the middle layer of wefts exhibited the largest variations and difference between each other.

6.3.4 Average Stresses Along Tow Paths

A few sections earlier, the most severe stress components for each type of tow (binder, warp, and weft) were identified for the 1x1 orthogonally woven textile model. The most severe component of stress was identified for each type by comparing the extrema of each stress component to a set of assumed strengths. Since damage will generally initiate due to the most severe stress relative to the respective strength of the material, this component of stress is important for understanding where damage will occur within the textile and this information can ultimately be used to design a textile to delay the onset of damage. For each type of tow, this section presents the most severe component of stress along the tow path averaged over the cross-sectional area at each point. Stresses are reported in the material coordinate system, denoted by σ'_{ij}, since failure typically occurs according to modes that correspond to the local material coordinate axes.

Binders

In the binders, σ'_{zz} was the most severe component of stress, which would initiate transverse matrix cracking. Figure 6.19 plots the stress averaged over the cross-section along the tow path for each binder and shows the tow paths for context of where peaks occur. Since symmetry was assumed on the bottom plane of the unit-cell and the top plane was left free, the difference in the phase of undulation makes a difference for stresses, unlike the cross-sectional area. In the figure, there are four curves shown, three of which show the average stresses for each of the binders, while the fourth line (dotted) shows the stresses for the binder in the second row after adjusting for the difference in phase After adjusting for

the phase difference, the stress distribution matched strikingly well for all the tows, despite the first and third row of tows being cut longitudinally by boundaries of the unit-cell.

The highest average stress occurred in the matrix rich regions between rows of wefts, which are labelled as A and B in Figure 6.20. It was shown in Ref. [67] that the highest stress concentrations occur within this matrix rich region, but Figure 6.19 shows that not only do the highest stress concentrations occur in the matrix rich region but also the largest average stresses. The smaller peaks in Figure 6.19 correlate well with the points where the binders start or finish to cross wefts, such as the two positions referred to as C in Figure 6.20. The highest stress concentrations observed in Ref. [67] occurred near these smaller peaks, but this paper shows that the highest average stress occur in the middle of the matrix rich regions, not where the binder begins or finishes crossing a weft. This model treats the tows as homogeneous material as an approximation, so stresses within the very small volumes at the highest stresses cannot be accurately predicted without modeling the fibers discretely. Unfortunately, modeling each fiber within the textile unit-cell would be computationally prohibitive. It is very possible that any noticeable damage at the mesoscale will occur where the average stresses are the highest, not necessarily where the point-wise stresses are highest.

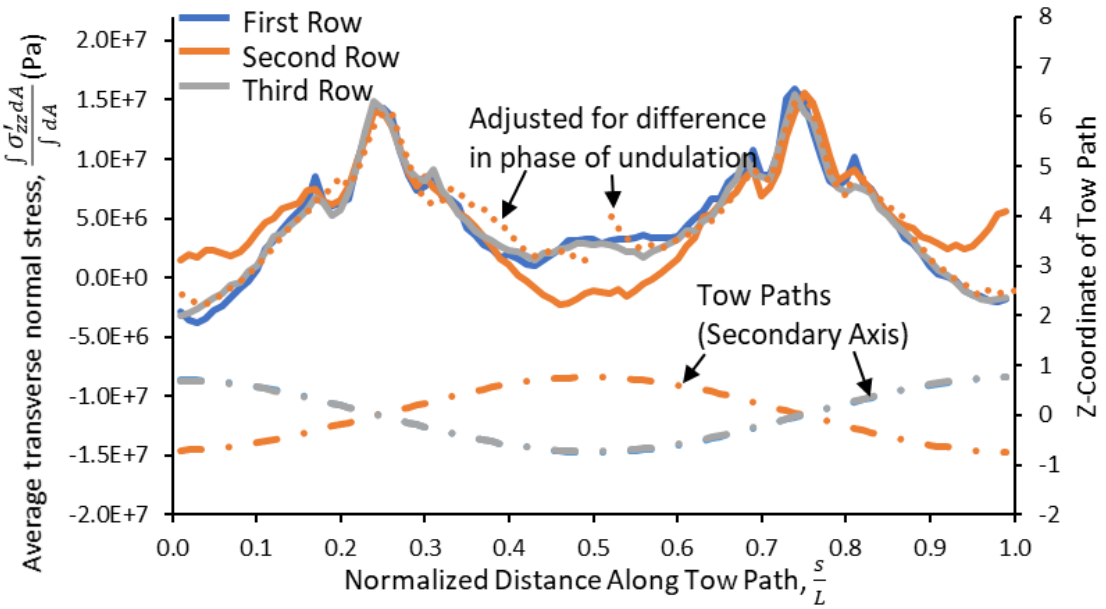

Figure 6.19. σ'_{zz} averaged over the cross-section within the binders plotted along the tow paths

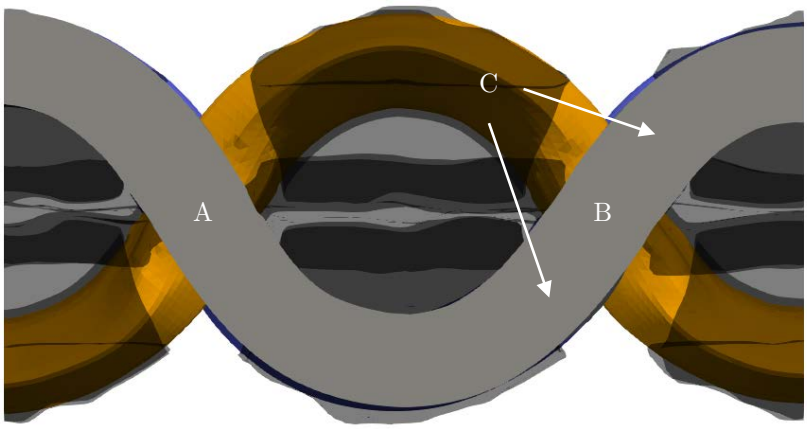

Figure 6.20. Illustration of regions of interest for binders

Warps

In the warps, σ'_{xz} was the most severe component of stress. Figure 6.21 plots σ'_{xz} averaged over the cross-section along the tow path for each of the warps. Figure 6.21 shows that the wavelength of changes in the shear stress is much shorter than the size of the unit-cell. Key segments along the tow paths are labelled to provide some context of where the variations occur, and there is no obvious correlation between the location of peak shear stresses and the surrounding tow geometry.

It may be helpful to remember that the boundary condition applied at $^{x}/_{L} = 0$ results in a zero σ'_{xz}, if the material and global coordinate systems align. The stress value is not quite zero due to a slight misalignment between the two coordinate systems. On the right side at $^{x}/_{L} = 1$, the global shear stress σ_{xz} was zero, but the stress in the material coordinate system was notably different because the local coordinate system did not align well with the global at this point. In VTMS, a larger section of the textile is generated and clipped down to the unit-cell seen in this work, and the warps did not flatten out well in this region.

However, there does not seem to be a correlation between the peak average stresses and the surrounding tows or matrix pockets, Figure 6.22 shows the tow path for one of the warps, revealing that the dominant wavelength in Figure 6.21 matches well with the wavelength of tow undulation. The vertical gridlines align between the two plots, allowing a direct comparison. The average shear stress appears to have several frequencies in the variation, and though the larger variations at low frequencies seems to be heavily influenced

by variations in the tow path, it is still unclear what causes the smaller variations at high frequencies.

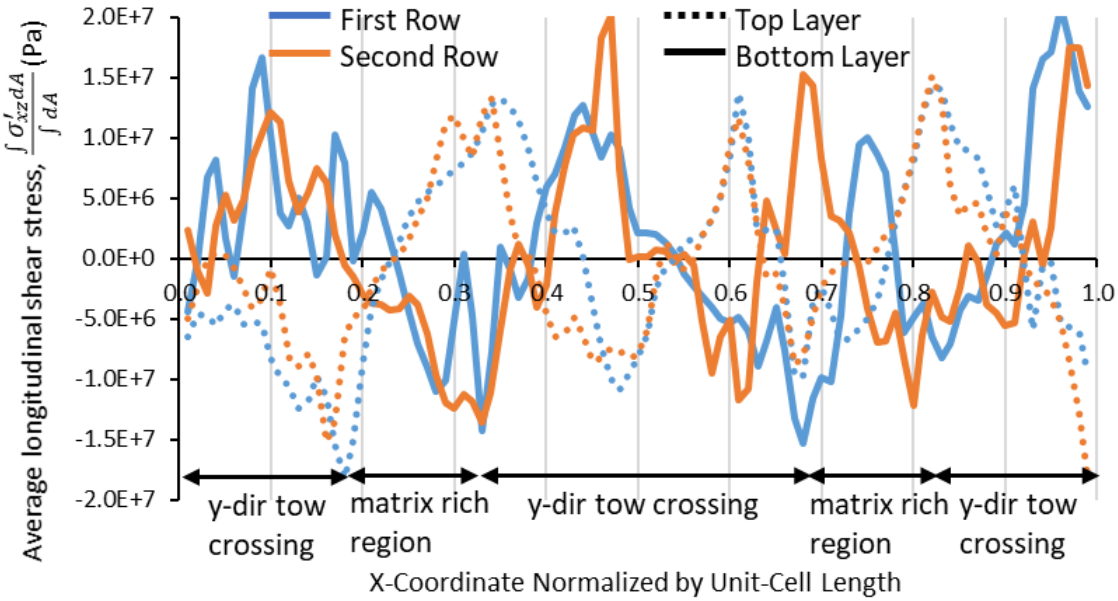

Figure 6.21. σ'_{xz} averaged over the cross-section within the warps

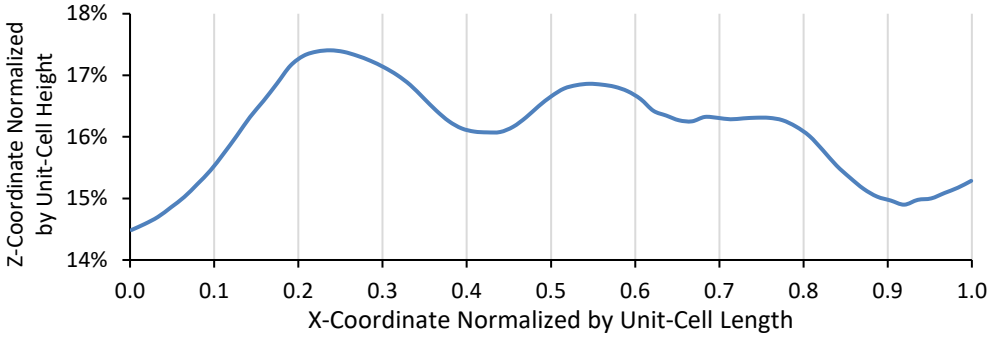

Figure 6.22. Piecewise cubic spline fit for the path of the warp on the first row and top layer

Since σ'_{xz} is the most severe component of stress for the warps and most sensitive to the tow path, the assumption that the tow paths remain approximately 2D appears poor. The in-plane undulation of the centroid has a magnitude that is twice that of the out-of-plane undulation.

Wefts

In the wefts, σ'_{yy} was the most severe component of stress, which also generally aligns with the direction of the applied load. Figure 6.23 plots the average transverse stress along the tow path for the outer layer of wefts, after adjusting for differences in the phase of tow undulation. The trends matched well between all of the outer tows, with no two tows differing by more than 4e6 Pa. The largest stress occurs where the weft crosses an warp, labeled as B in Figure 6.16a. The warps are carrying the most load in the direction of this stress component, so it makes sense that the most load is transferred where the warp comes in the closes proximity. It seems that the average stresses experienced in the outer layers of wefts is largely unaffected by their row or layer but rather dependent on the location relative to the orthogonal and warps.

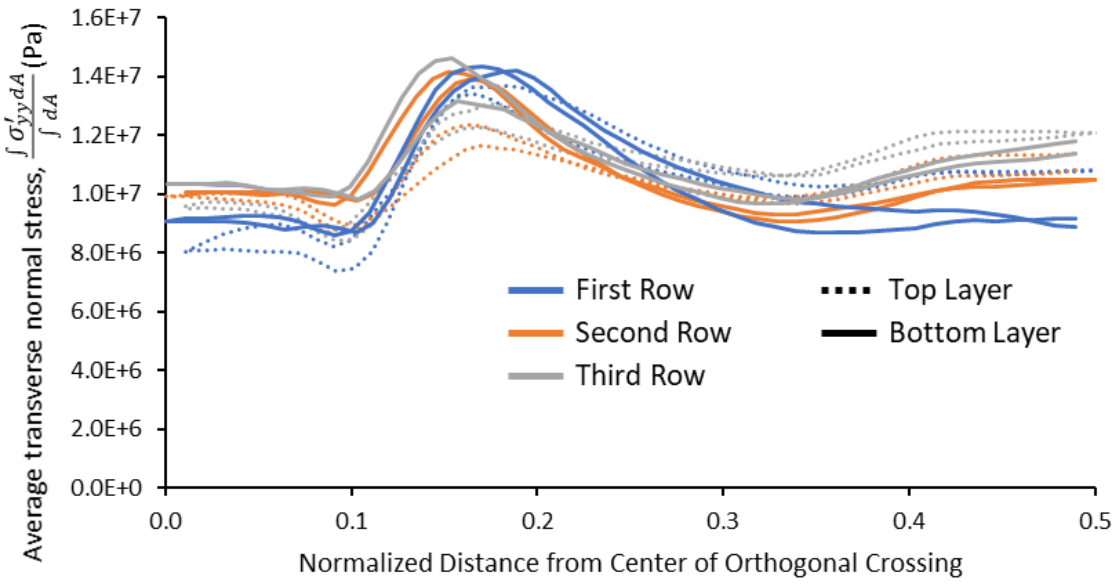

Figure 6.23. Average transverse normal stress within the outer layers of wefts

Figure 6.24 shows the average σ'_{yy} along the tow path for the middle layer of wefts. Unlike the outer layers of tows, the largest average stress occurred near the matrix pockets between where the warps cross the wefts, see Figure 6.16d. Additionally, the middle layer of wefts experienced much higher average stresses than the outer layer of wefts. In Ref. [67], it was observed that stress concentrations occurred in all three layers, but the peak average stress in the middle layer of weft is about 34% higher than the peak value in the outer layers. The observed concentrations were so small in size that the heterogeneous microstructure should be taken into account, so the average stresses might be a more reliable predictor of where noticeable damage will occur at the mesoscale.

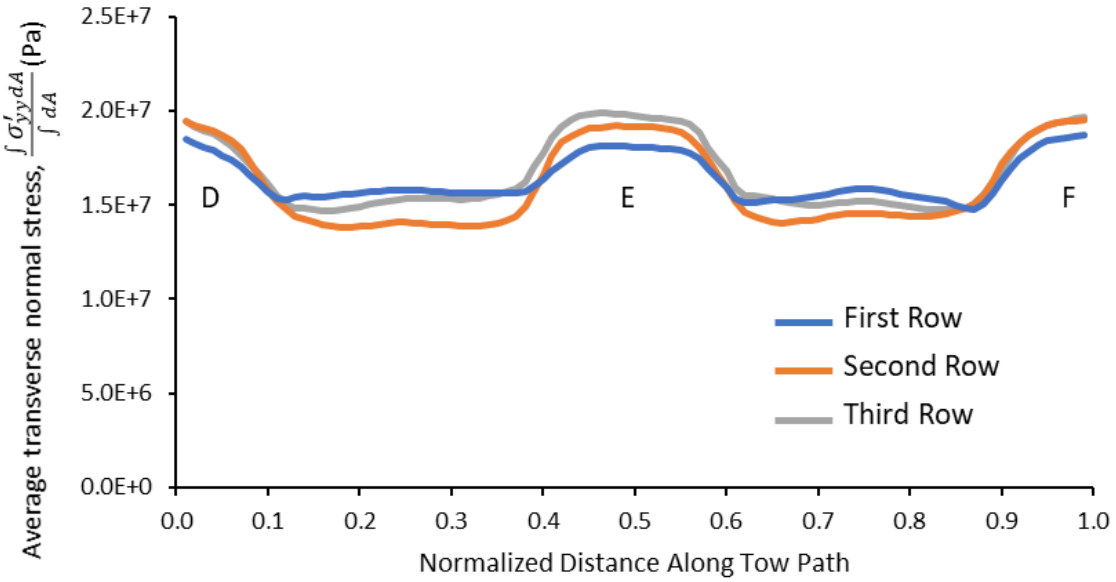

Figure 6.24. Average transverse normal stress within the middle layer of wefts

6.3.5 Locations of Critical Stress in the Matrix

In the previous sections, it shown was shown that the most severe stress concentrations within the binders and wefts occurred where binders come close to the wefts and the most severe stress concentrations within the warps occurred in the gaps between wefts. In the model used in this work, some matrix lies between tows throughout the textile model, since the tows were shrunk to remove overlap between tows. Therefore, the stress concentrations within the matrix unsurprisingly occurred near the severe stress concentrations in the tows, as load is transferred from one tow to another. Figure 6.25 shows the volume of matrix with a von Mises stress, σ_{vm}, greater than 80 MPa with contours of σ_{vm}. The wefts and warps are only shown for half of the unit-cell to show the locations of stress concentrations in the matrix, but the warps are shown semi-transparently in the other half of the unit-cell for context. Figure 6.25a highlights that the stress concentrations within the matrix occurs

where warps and wefts come close to binders. Where warps come close to binders, a severe σ'_{xz} was observed in a previous section, and the transfer of load through shear between warps and binders in this region caused the most severe σ_{vm} concentration in the matrix, as shown in Figure 6.25a. It was shown in a previous section that a severe σ'_{yy} formed in the wefts near binder crossings, but the precise location depended on the layer of wefts. For the middle layer of wefts, the σ'_{yy} concentrations formed where the weft came closest to the binder, and Figure 6.25a shows that stress concentrations occurred in the matrix at these locations, such as point A in the figure. In the top and bottom layers, the σ'_{yy} concentrations in the wefts occurred just before and after the weft crosses the binder, and Figure 6.25a also shows that stress concentrations occurred in the matrix at these locations, such as point B in the figure.

Figure 6.25b and Figure 6.25c shows that the majority of highly stressed matrix material lies between the rows of wefts. Between the rows of wefts, there is a lower volume fraction of tows. Since tension along the global x-axis was considered for this study, the textile can be conceptually divided along the x-axis into sections with alternating higher and lower tow volume fractions. The sections with higher and lower tow volume fractions experience the tensile load in series, so it is expected that the matrix will generally carry more load in sections with a lower tow volume fraction. However, the locations of the most severe stress concentrations in the matrix occurred specifically where load is transferred between tows near stress concentrations in the tows, but it should be noted that the matrix between all of the tows is a result of shrinking all of the tows to remove interpenetrations.

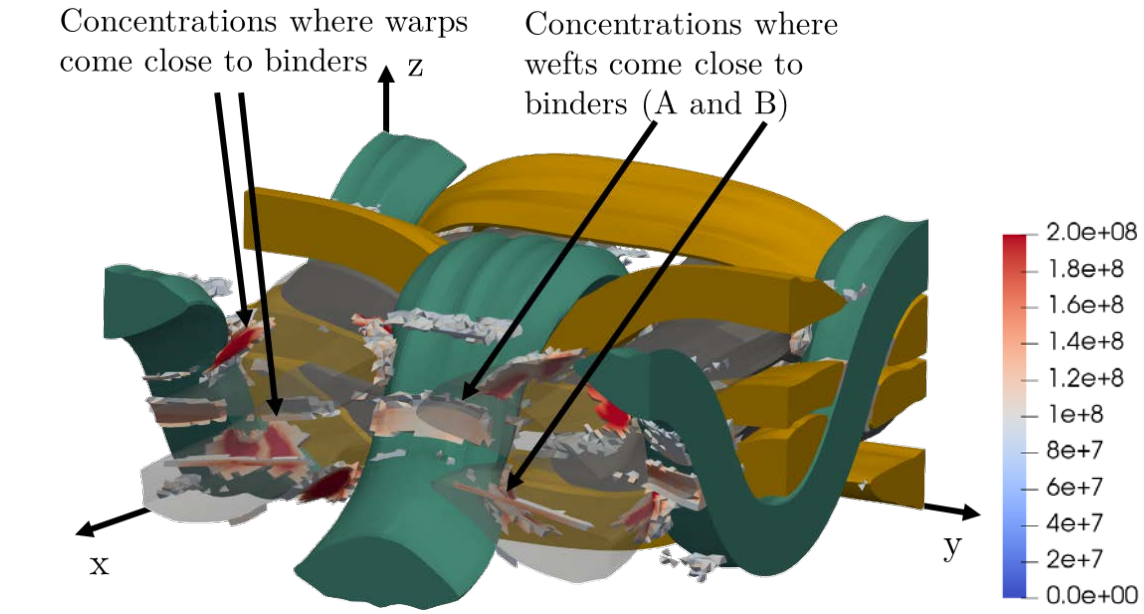

a) View highlighting locations of stress concentrations in the matrix near warps and wefts

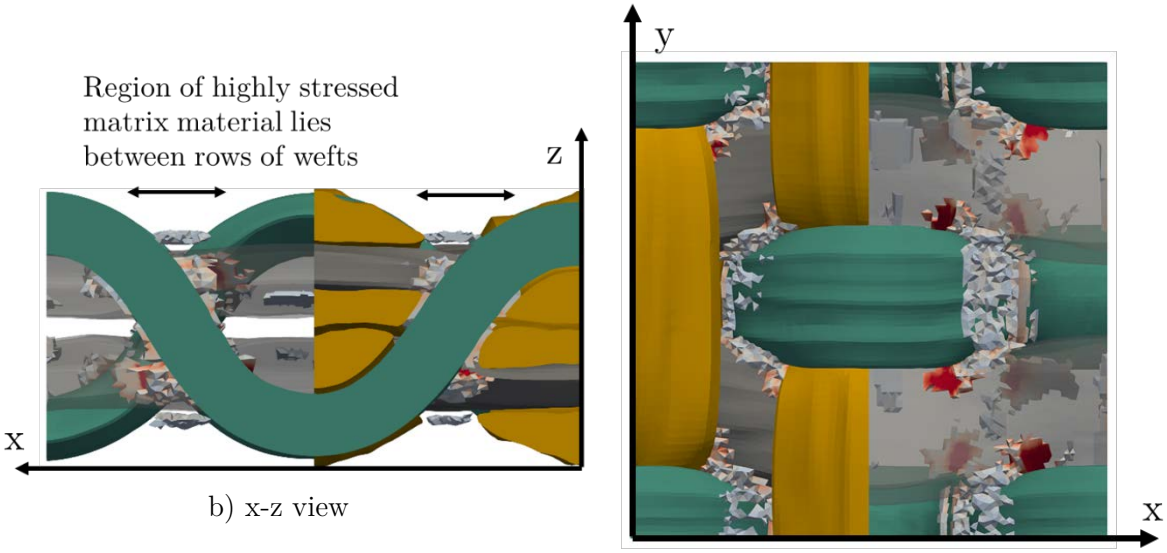

b) x-z view

c) x-y view

Figure 6.25. Volume of mesh with $\sigma_{vm} > 80$ MPa shown with contours of σ_{vm}. Warps, wefts, and binders are shown as solid dark gray, gold, and teal respectively.

6.3.6 Summary

The stresses within an 1x1 orthogonal weave were investigated using a nonidealized finite element model. The effect of mesh refinement on the volume average stresses and distributions of stresses was studied. A much higher refinement than what is typically used in the literature was considered, and yet, mesh convergence was probably not be reached. The coarse and medium meshes, which are closer to meshes usually seen in the literature, capture the stress distributions away from the stress concentrations, but the predictions are poor near the stress concentrations. For one of the volume average stresses, the coarse mesh exhibited a 23% error compared to the refined mesh, which is quite large for a volume averaged quantity. The medium mesh performed better, only differing by 7% or less for the volume average stresses.

Using the most refined model, some insight was gained into understanding the location and component of stress that will likely initiate damage. Within the tows, the transverse normal stress in the wefts, σ'_{yy}, was the most severe, indicating matrix cracking within the wefts would likely be the initial damage within the textile. The load was not distributed evenly among the wefts, with the middle layer of tows carrying more load than the top and bottom layers, but severe σ'_{yy} concentrations occurred in all the layers. Within the warps, the longitudinal shear stress, σ'_{xz}, was the most severe, and within the binders, the transverse normal stress, σ'_{zz}, was the most severe. However, the failure modes of these tows are difficult to predict without a progressive damage analysis, since damage in the wefts could significantly change the stress distribution within the other constituents.

Because of the complexity of the composite, the study included an evaluation of the geometry generated by VTMS, with a particular focus on the cross-sectional area variation. The cross-sectional area of the tows varied significantly along the tow paths, with the wefts exhibiting the most variation. The variation of cross-sectional area would result in up to a 26% change in the fiber volume fraction. The effect of the local fiber fraction variation was not addressed in this study but will be explored for the 2x2 orthogonal weave.

In the binders, the peak average σ'_{zz} occurred near the matrix rich regions between the rows of wefts. In the warps, the average σ'_{xz} appeared sensitive to the undulation of the tow path. This study assumed that the tow paths were approximately 2D, but the in-plane undulation of the tow path was twice the out of plane, which highlights the importance of capturing the 3D nature of the tow paths. For the wefts, the location of the peak average σ'_{yy} depended on the layer considered. For all outer layers of wefts, the distributions of cross-sectional averaged σ'_{yy} were similar, and the largest values occurred where warps cross the weft paths. On the other hand, the middle layer of wefts experienced the largest cross-sectionally averaged stress near the matrix rich regions between the rows of warps. Even though stress concentrations were observed in all three layers, the magnitude of the peak cross-sectionally averaged stress in the middle layer was 34% higher than the peak average stress in the outer of layers of wefts. It is unclear if damage would initiate near the small stress concentrations or where the cross-sectionally average stress is maximum. Finally, the stress concentrations within the matrix occurred near the stress concentrations in the tows, but this may be an artifact of some matrix material existing between the tows throughout

the model to avoid interpenetrations. Creating realistic models and understanding the connections between the tow architecture and the distribution of stress is critical to designing better 3D textile composites, and this study was an early effort.

6.4 Analysis of a 2x2 Orthogonal Weave

This study investigates a 2x2 orthogonally woven textile, where the binders cross two wefts at the top and bottom of the textile before traversing through the thickness. The aim is to explore the locations of critical stress concentrations for a variety of loads and assess the effect of modeling the variation of fiber volume fraction in the tows and plasticity in the matrix.

The study of the 1x1 orthogonal weave in the last section used a model that was thinner than the textiles often used in industry, and the boundary conditions likely introduced artifacts, even in the interior of the model. However, the study showed the importance of accounting for the 3D nature of the tow paths when determining the local material properties and the need to assess the effect of a variation of the local fiber volume fraction. Guided by these conclusions, this study considers a more realistic 3D textile and explores the response under a variety of loads with greater detail. The study of the 2x2 orthogonal weave is broken up into four results sections. The first section characterizes the tow architecture, and the second section investigates the stress distributions within the tows for three load configurations. The third section explores the effect of the binders by replacing the binders with matrix material and observing how the stress distributions change. Finally,

the last section investigates the effect of modeling the variation of the fiber volume fraction and plasticity in the matrix on the locations and severity of stress concentrations.

Throughout this study, stresses in the global coordinate system are denoted as σ_{ij}, while stresses in the local coordinate system are denoted by a prime, σ'_{ij}. To assess the severity of stresses within the tows, components of stress in the local coordinate system are normalized by a nominal strength, which are given in Table 6.4, to yield a normalized stress, which is denoted by a hat, $\hat{\sigma}'_{ij}$.

6.4.1 Characterization of Tow Architecture

Since the textile geometry is the result of simulated processing, the cross-section of the tows can vary along the tow paths due to a variety of reasons. During compaction, the cross-sectional area is likely to be reduced where tows contact other tows. Additionally, the shape of the tows can vary significantly when it changes direction, such as when the binder begins or ends traveling through the thickness of the textile. Characterizing the variation of the cross-sections of the tows is a first step toward understanding how closely the model resembles actual specimens. This section quantifies the variation of cross-sectional area and explains how aspects of VTMS's method for creating textile models led to the variations.

For the model considered in this paper, adjacent binders have a different phase of undulation through the textile, but every fourth binder has the same phase. Every other binder will traverse through the thickness of the textile at the same gap between weft tows, and so a pair of such binders is more likely to have similar cross-sectional areas. Figure

6.26 shows the variation of the cross-sectional area of the first and third binder, which are shaded in the figure as green and purple and labeled 1 and 3. During compaction, the cross-sectional area is reduced where tows come into contact, but the binders do not come into significant contact with other tows where the binders are midway through the thickness of the textile. Consequently, the cross-sectional area at the midpoint through the thickness remains almost unchanged throughout the compaction process and is maximum. The minimum occurs where the binders transition to and from traveling through the thickness. The low cross-sectional area at these points are due to several factors, such as the lower radius of curvature and the tension in the binders causing the binders to strongly contact neighboring wefts. Additionally, it was observed that digital chains tend to pass through each other in these regions, resulting in significant tow interpenetrations. It is unclear which factors are most significant, but the cross-sectional area decreases by 19% to 24% as the binders transition into traveling through the thickness, such as at point A in Figure 6.26. Where the binders are on the top or bottom of the textile, the cross-sectional area varies by at most 18%. It is unclear if this much variation is realistic or merely an artifact of the method used by VTMS to create the textile geometry.

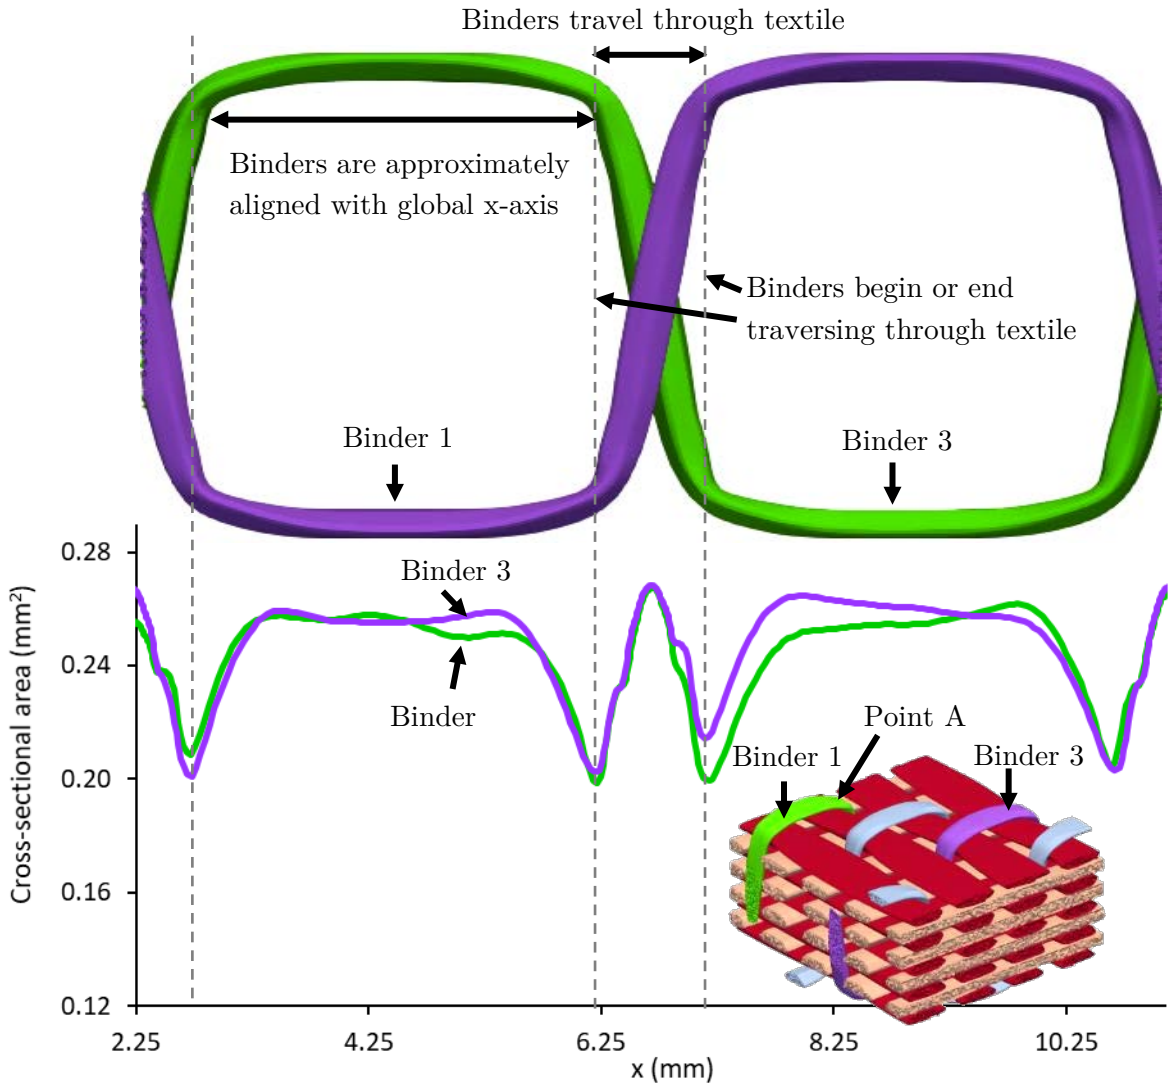

Figure 6.26. Variation of the cross-sectional area in the first and third binders

Figure 6.27 shows the variation of the cross-sectional area in the warp tows. Figure 6.27a shows that the cross-sectional area for the inner two layers of warps (green) remains relatively constant, varying at most by 4.5%. However, Figure 6.27b shows that the warps on the top (purple) and bottom (green) of the textile experience significant changes in cross-

sectional area. There are two types of variation of interest. One type of variation is that the cross-sectional area of a single tow changes along the tow path, which for the outer layer of warps is at most 11%. Another type of variation is that the cross-sectional area varies between tows in the same layer, which is at most 15% for the model presented in this paper. Both types of variations will exist in actual specimens, but comparisons to MicroCT scans would be needed to quantify the accuracy of the model geometry. However, these results do show that the cross-sectional area of tows in a model created by VTMS largely depend on the distance from the boundaries where the compaction planes contact the model. The deformation due to the rigid planes at the top and bottom boundaries does not cause significant deformation well into the interior of the model. Tows near the boundaries will be highly deformed, while tows in the interior will experience little deformation. The variation of the cross-sectional area in the wefts were not shown for conciseness, but they follow the same trend. The outermost layers varied by up to 26%, while the innermost layer varied by less than 2%.

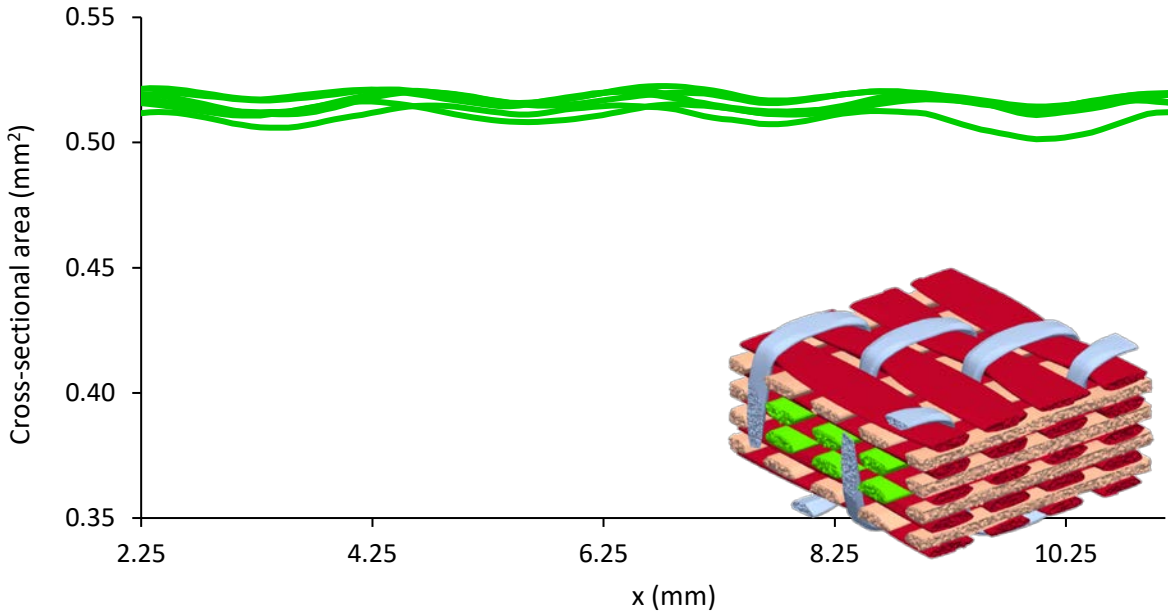

a) Interior warp tows

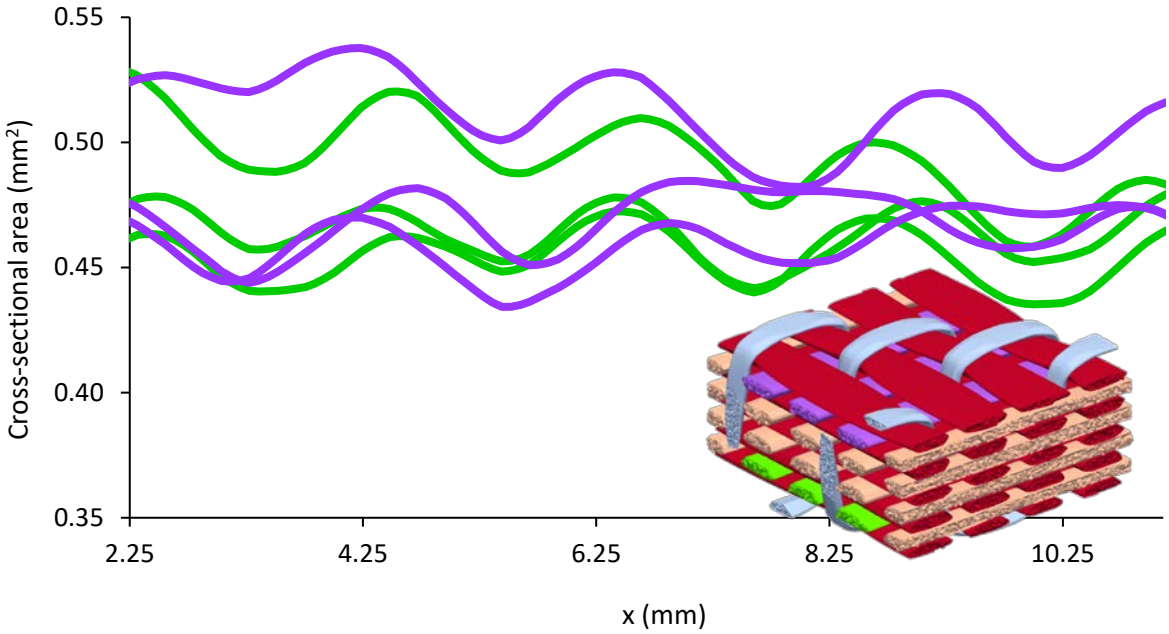

b) Top and bottom layers of warps

Figure 6.27. Variation of the cross-sectional area in the warps

6.4.2 Critical Stresses Under In-Plane Loads Assuming 60% V_f in Tows

The 2x2 orthogonally woven textile model was subjected to three types of load: uniaxial tension along the global x-axis, uniaxial tension along the global y-axis, and in-plane shear. For each load configuration, this section aims to investigate the locations of stress concentrations and explore the local load transfer that caused the concentrations. However, it should be noted that this section assumes the tows remain linear elastic, so the behavior of progressive damage cannot be understood from these results.

Uniaxial Tension Along X-Axis

Figure 6.28 shows the three most severe normalized stresses within the clipped analysis region for an applied volume average strain along the global x-axis, $\langle \varepsilon_{xx} \rangle$, of 1%. The stresses shown are in the local coordinate system, and each stress component is normalized by the respective strength shown in Table 6.4, which provides a measure of the severity of each stress component and is denoted by $\hat{\sigma}'_{ij}$. A larger $\hat{\sigma}'_{ij}$ indicates a more severe σ'_{ij}. For tension along the x-axis, the wefts and binders experience the most severe stresses. The binders experience a severe σ'_{zz} and σ'_{xz}, as shown in Figure 6.28b and Figure 6.28c respectively. The severe σ'_{zz} occurs where the binders travel through the thickness of the textile model, as shown in Figure 6.28b.

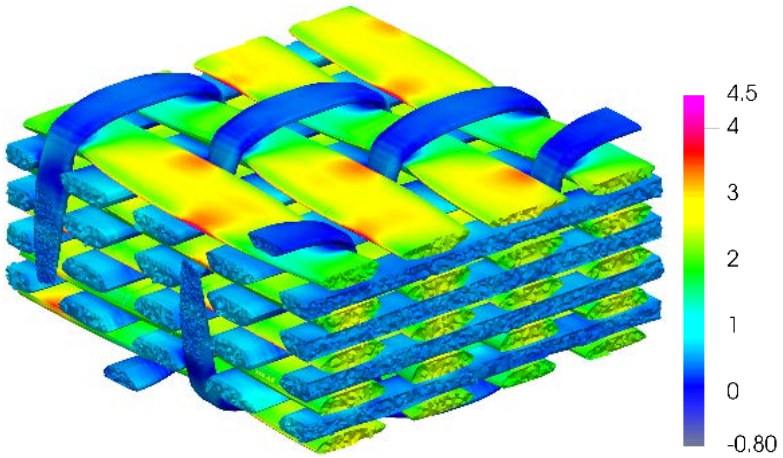

a) $\hat{\sigma}'_{yy}$

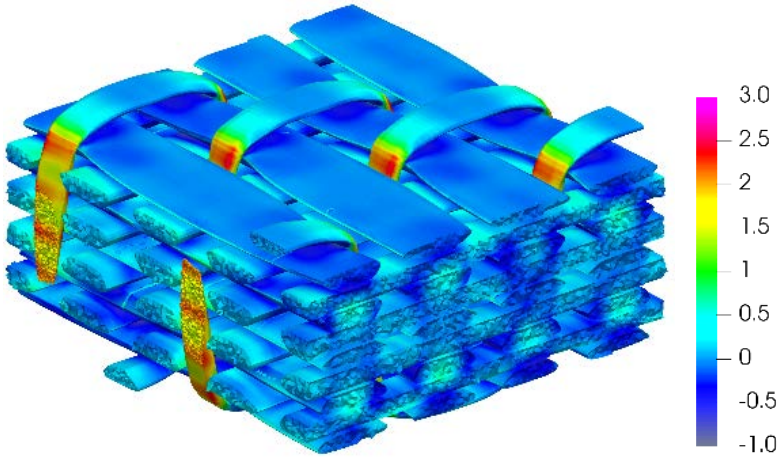

b) $\hat{\sigma}'_{zz}$

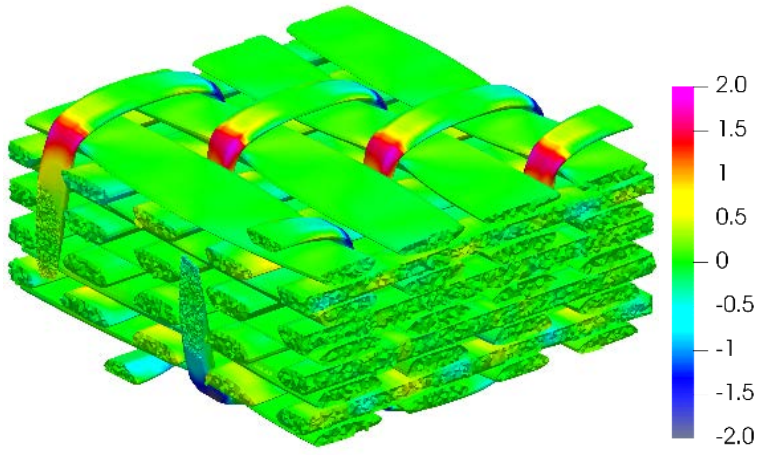

c) $\hat{\sigma}'_{xz}$

Figure 6.28. Contours of three most severe stress components in the local coordinate system normalized by the respective strength for uniaxial tension along the global x-axis

To show the locations of σ'_{zz} concentrations more clearly, Figure 6.29 shows $\hat{\sigma}'_{zz}$ within an x-z slice of one of the binders. In this region, the local z-axis almost aligns with the global x-axis, which is the direction of the applied load, refer to Figure 6.29b. Additionally, there are no wefts in the region where the binders travel through the thickness of the textile. Consequently, the binders take on much of the load via tension along the local z-axis. The peak stresses occur where wefts are nearby to transfer the load to the binders. This component of stress will likely cause transverse matrix cracking within the binders.

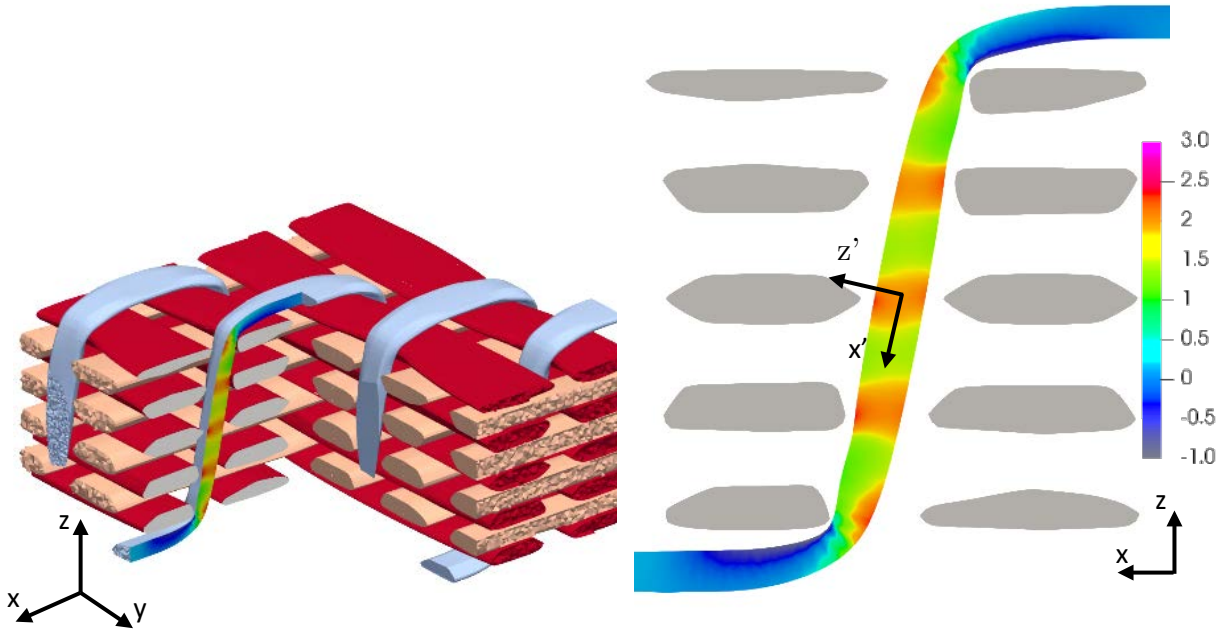

a) Location of x-z cross-section b) Normalized transverse normal stress $\hat{\sigma}'_{zz}$

Figure 6.29. $\hat{\sigma}'_{zz}$ contours for an x-z cross-section of a selected binder illustrating locations of severe σ'_{zz} in binders

In addition to the severe transverse tension as they traverse the thickness of the textile, the binders experience a severe shear stress, σ'_{xz}, where they begin and finish traveling through the thickness of the textile, as shown in Figure 6.28c. The shear stress develops to maintain equilibrium as the binders shift between carrying a significant amount of load at the top and bottom of the textile, where the fibers of the binders are aligned with the load direction, to carrying much less load as they travel through the thickness of the textile.

Though the binders experience severe stresses, the wefts experience the most severe stress, namely $\hat{\sigma}'_{yy}$, within the textile, as shown in Figure 6.28a. For the wefts, the local y-axis is closely aligned with the global x-axis, which is the direction of the applied load for this configuration, so it is expected that the transverse tension in the wefts will be severe. However, the severity of the transverse tension varies dramatically depending on the location within the wefts. Figure 6.30 shows the local and global stress components of interest for an x-z cross-section centered on a binder tow, which is the same location illustrated in Figure 6.29a. Figure 6.30a shows the $\hat{\sigma}'_{yy}$ contours in the wefts and binder in the x-z cross-section. In most locations within the wefts, $\hat{\sigma}'_{yy}$ remains between 2 and 3. At some points when the wefts come close to a binder, $\hat{\sigma}'_{yy}$ reaches values near 4.5, such as point A in Figure 6.30a, but at other points near a weft tow, no stress concentration occurs, such as point B in Figure 6.30a. The common factor for areas of elevated stress is a sharp cross-sectional shape of the weft. When the cross-section is similar to a rectangle with rounded corners, there is no significant stress concentration, but when the cross-section comes to sharper point, a significant stress concentration occurs. The sensitivity of stress

concentrations to the tow cross-sectional shape emphasizes the importance of creating textile models with realistic tow geometries.

There is one situation where the tow shape is less prone to cause a stress concentration for this configuration: where a binder crosses over a weft tow. In this region, $\hat{\sigma}'_{yy}$ in the weft drops to about 0, such as point C in Figure 6.30a. The drop in transverse tension is due to a compressive σ_{xx} (global) in the region where a binder crosses the wefts at the top or bottom of the textile, such as region D highlighted in Figure 6.30b. The compressive stress forms as load is transferred to or from the binder via shear, since the binder is orders of magnitude stiffer in the global x-direction than the surrounding material at the top and bottom.

It should be noted that the locations of stress concentrations for the 1x1 orthogonal weave differed from the 2x2 orthogonal weave. For the 1x1 orthogonal weave, stress concentrations occurred in all layers of the wefts, but the middle layer of wefts carried more of the load than the top and bottom layers. However, for the 2x2 orthogonal weave, the top and bottom layers of wefts carried more load, and severe stress concentrations were limited to the top and bottom layers. For the 2x2 orthogonal weave, the binders remain aligned with the load direction for more of the tow path in the 2x2 orthogonal weave, since they cross two wefts instead of one before traversing the thickness of the textile, which causes the binders to carry more of the load than in the 1x1 orthogonal weave. Since the binders transfer load to the top and bottom layers of wefts, the top and bottom layers experience more severe stresses than the middle layers.

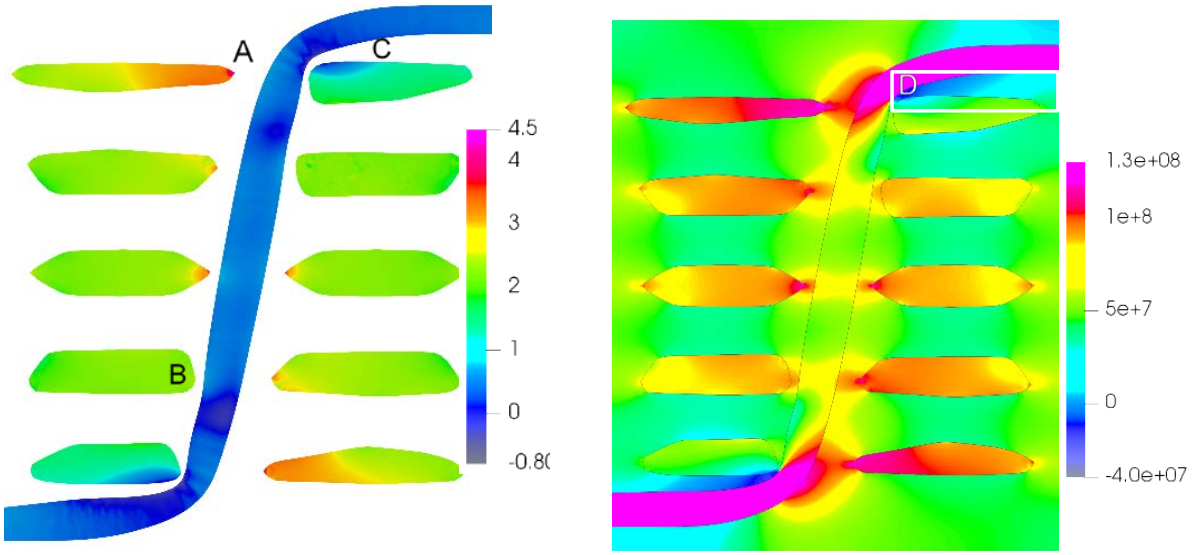

a) Local normalized transverse normal stress $\hat{\sigma}'_{yy}$ b) Global σ_{xx}

Figure 6.30. An x-z cross-section illustrating locations of severe $\hat{\sigma}'_{yy}$ in wefts

Since the matrix is isotropic, the von Mises stress, denoted by σ_{vm}, is often used as scalar measure of the stress, and σ_{vm} will be used herein to explore the stress concentrations within the matrix. However, it should be noted that plastic deformation in the matrix is often dependent on the pressure, so the von Mises stress cannot be directly used to assess possible yielding.

Figure 6.31 shows the volume of matrix (shaded as yellow) with a von Mises stress higher than 80 MPa. The figure clearly shows that the volume of matrix with a von Mises stress higher than 80 MPa is limited to regions where binders begin or end traversing through the thickness of the textile. However, the stress concentrations only occur on the outside of the binders, refer to Figure 6.31. Earlier in this section, Figure 6.30b showed

that the binders carry a significant amount of the applied load when the they align with the load direction at the top and bottom of the textile, see the region just above the label D in the figure. However, when the binders begin to traverse the thickness, the load is transferred to the wefts, see location A in Figure 6.30a. The transfer of load from the binder to the wefts when the binder begins or ends traversing the thickness of the textile causes the stress concentrations in the matrix, as shown in Figure 6.30b. Figure 6.32, shows σ_{vm} contours for a x-z cross-section centered on one of the binders, which highlights that the locations of severe σ_{vm} in the matrix occurs at the same location of σ_{xx} concentrations that were observed in Figure 6.30b.

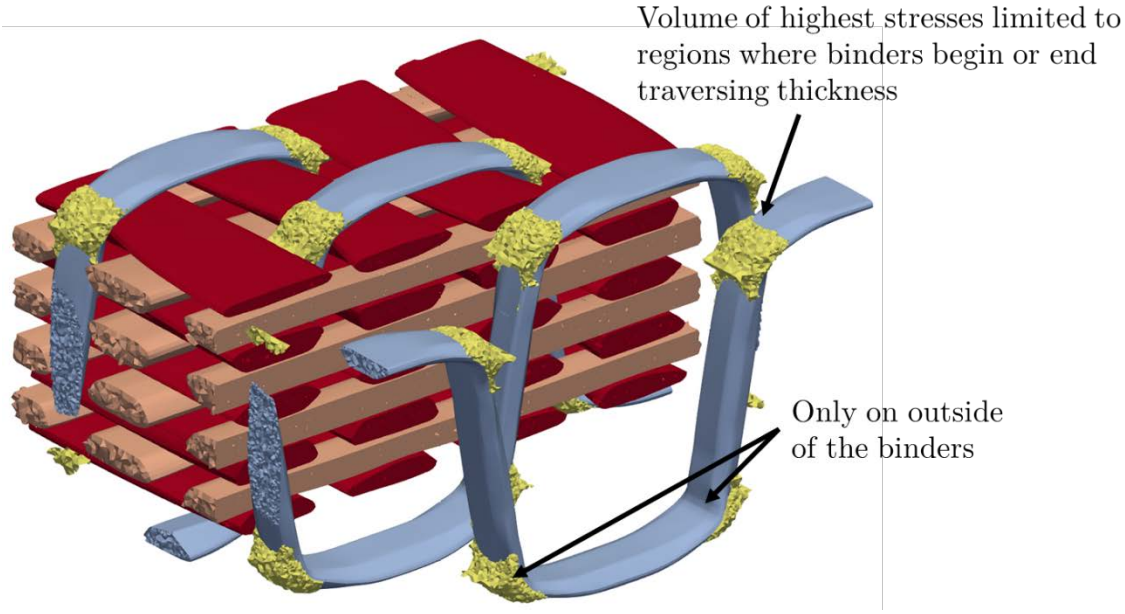

Figure 6.31. Volume of matrix with a $\sigma_{vm} > 8e7\ Pa$ (shown in yellow) with binders and clipped warps/wefts

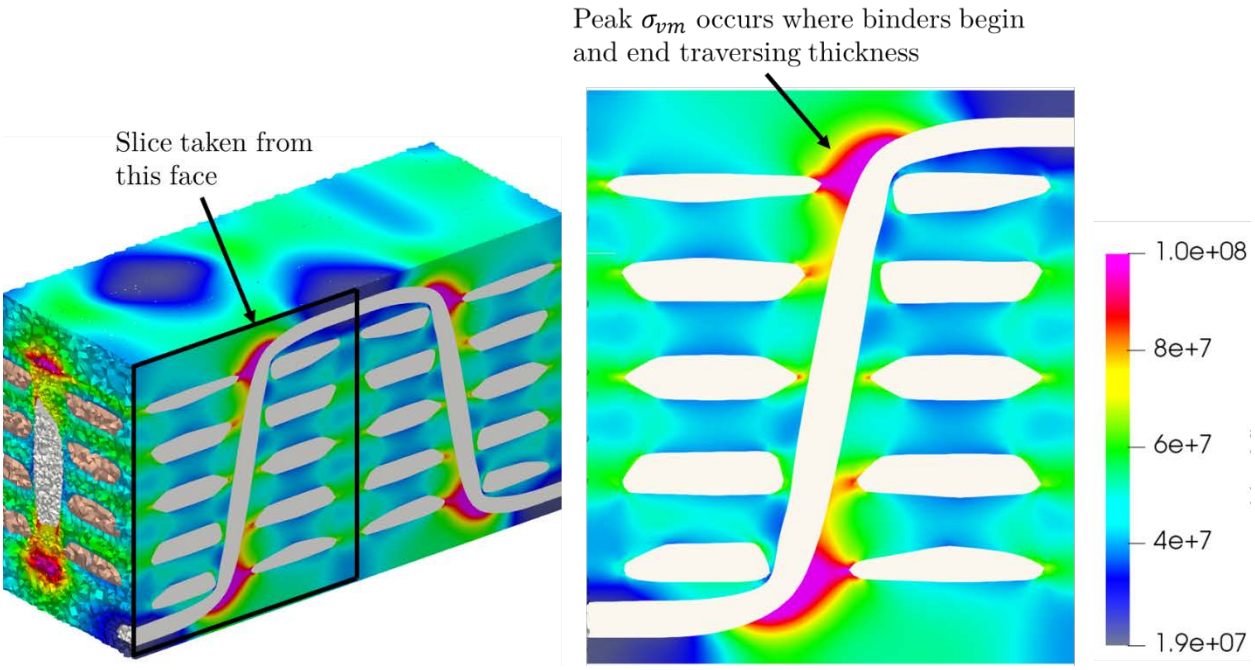

Figure 6.32. Contours of the von Mises stress, σ_{vm}, within the matrix for a slice that highlights the location of critical σ_{vm}

Uniaxial Tension Along Y-Axis

Figure 6.33 shows three most severe normalized stresses within the clipped analysis region for an applied volume average strain along the global y-axis, $\langle \varepsilon_{yy} \rangle$, of 1%. Again, the normalized stresses shown are in the local coordinate system. Overall, Figure 6.33 shows that the warps and binders experience severe transverse tension ($\hat{\sigma}'_{yy}$), while the wefts experience severe longitudinal shear ($\hat{\sigma}'_{xy}$ and $\hat{\sigma}'_{xz}$).

The most severe stress in the textile is $\hat{\sigma}'_{yy}$ in the binders and warps. For both types of tows, the local y-axis is closely aligned with the global y-axis, which is the direction of the load for this configuration. To illustrate where the peak stresses occur in the textile, Figure 6.34 shows $\hat{\sigma}'_{yy}$ for one binder and row of warp tows, along with the local coordinate system

for a point in the binder. The figure shows that the stress concentrations form where warps and binders come closest, such as point A in Figure 6.34. Additionally, stress concentrations form where the binder begins or ends traveling through the thickness of the textile, such as points B in Figure 6.34.

In summary, the transverse normal stress $\hat{\sigma}'_{yy}$ was the most severe component for both tensile load configurations. In both cases, the peak $\hat{\sigma}'_{yy}$ occurred within the tows that are perpendicular to the load. However, when the load is along the global x-axis, the stress concentrations only form within the wefts, but when the load is along the y-axis, the stress concentrations form within the binders and warps. Additionally, the stress concentrations are more severe for the case of tension along the x-axis.

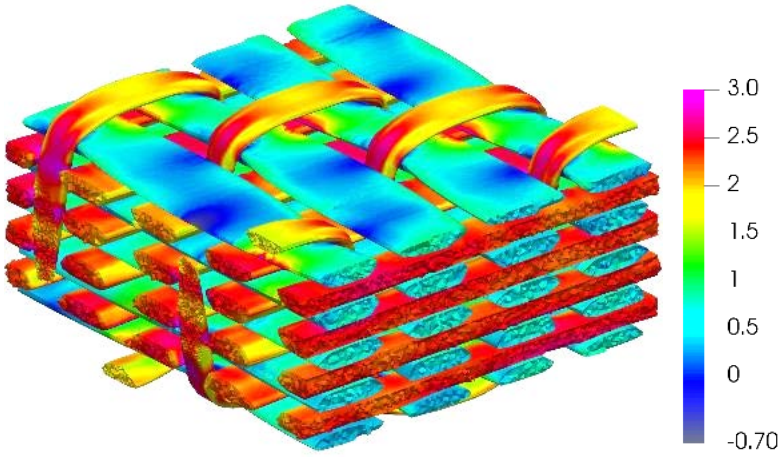

a) Normalized transverse normal stress $\hat{\sigma}'_{yy}$

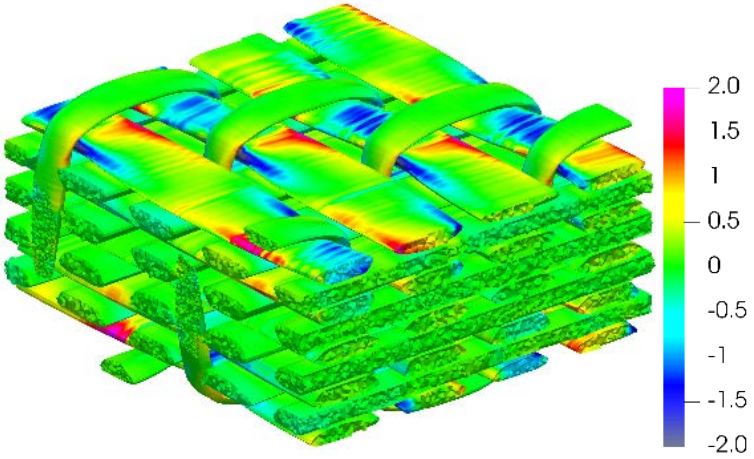

b) Normalized longitudinal shear stress $\hat{\sigma}'_{xy}$

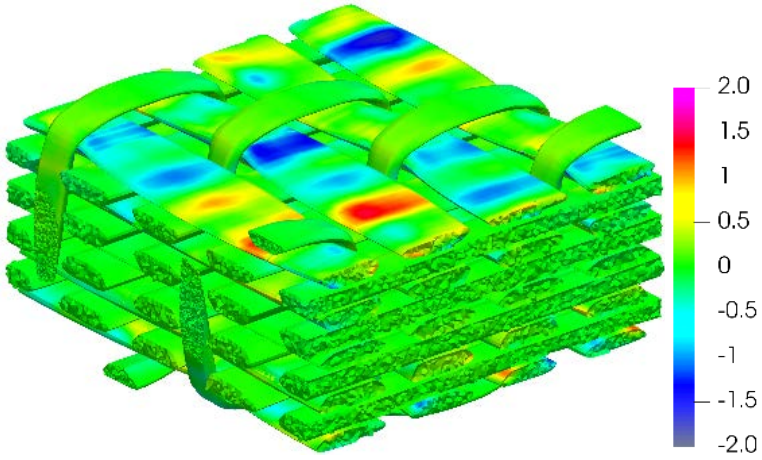

c) Normalized longitudinal shear stress $\hat{\sigma}'_{xz}$

Figure 6.33. Contours of each stress component in the local coordinate system normalized by the respective strength for the case of uniaxial tension along the global y-axis

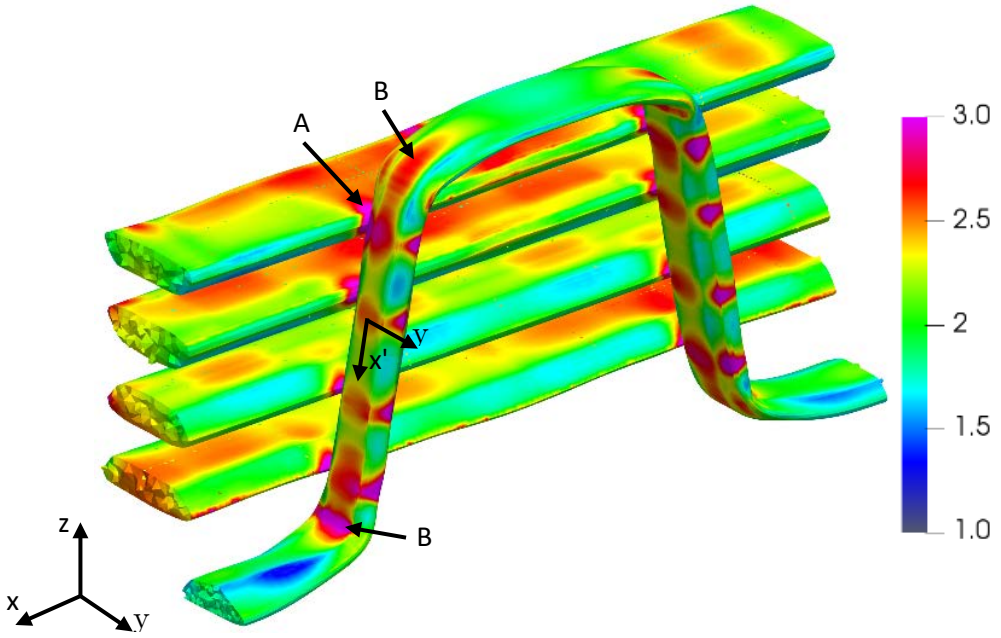

Figure 6.34. $\hat{\sigma}'_{yy}$ for a selected binder and row of warp tows

The second most severe type of stress in the textile is the longitudinal shear stresses, $\hat{\sigma}'_{xy}$ and $\hat{\sigma}'_{xz}$, within the wefts. To more clearly show the locations of severe shear stress, Figure 6.35 shows the magnitude of the normalized longitudinal shear stress, $\hat{\sigma}'_s = \left(\hat{\sigma}'^2_{xy} + \hat{\sigma}'^2_{xz}\right)^{1/2}$, for the volume of wefts with a value of $\hat{\sigma}'_s > 1.2$. The entire weft tows are shown semi-transparently for context of where stress concentrations occur. Figure 6.35 shows that the severe shear stresses in the wefts occur in the top and bottom layers of tows and near locations where binders cross the wefts. It should be noted that the stress concentrations do not occur uniformly throughout the model due to variations in the textile geometry and mesh refinement.

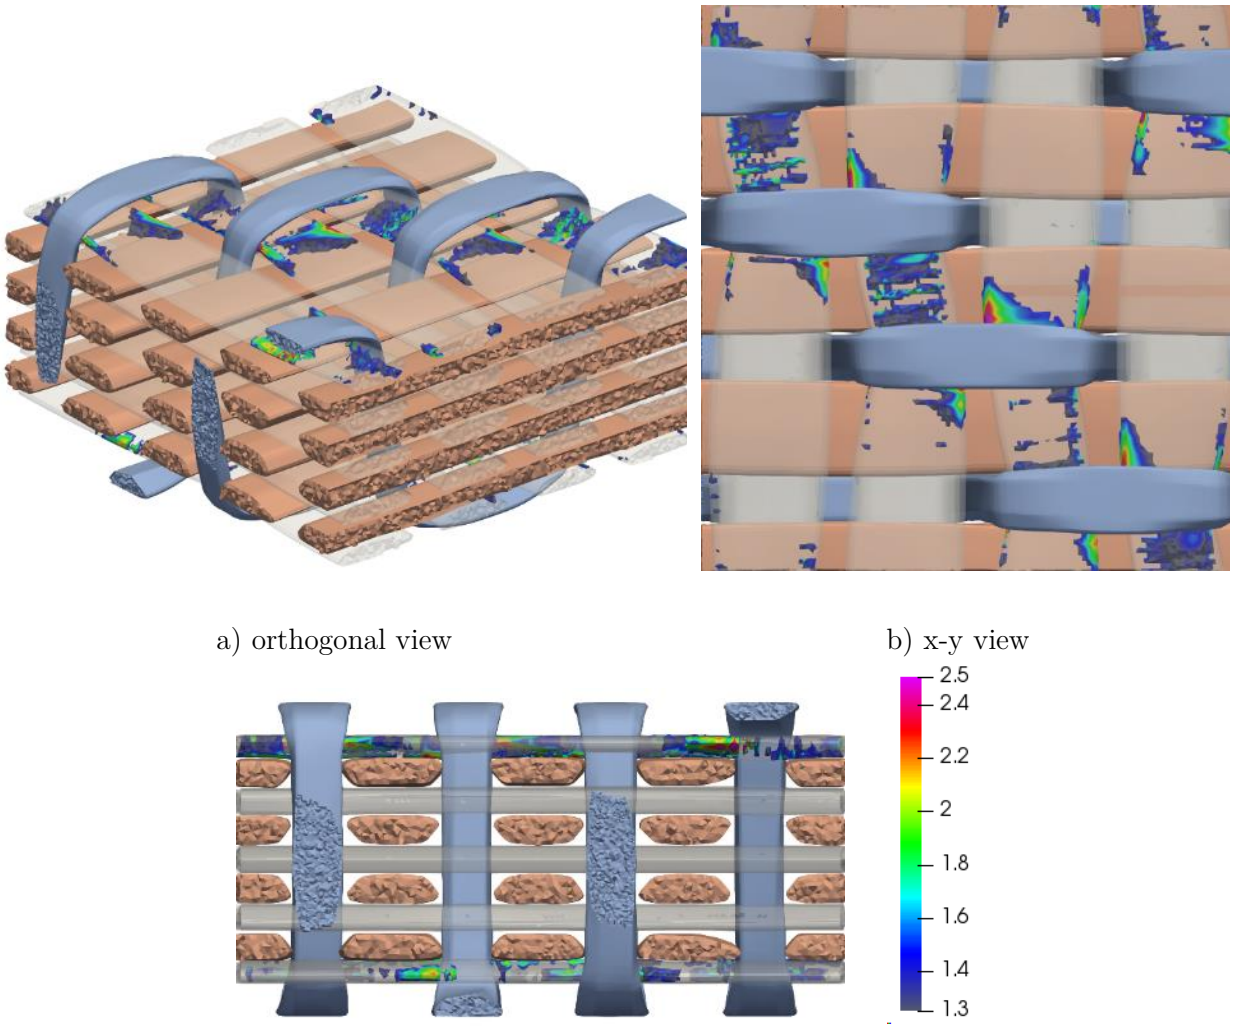

a) orthogonal view

b) x-y view

c) y-z view

Figure 6.35. Magnitude of normalized longitudinal shear stress, $\hat{\sigma}'_s = \left(\hat{\sigma}'^2_{xy} + \hat{\sigma}'^2_{xz}\right)^{1/2}$, for the volume of wefts with a value of $\hat{\sigma}'_s > 1.2$ (entire weft surfaces are shown as semi-transparent)

In summary, the second most severe stress for both tensile load configurations was the longitudinal shear stress, but the locations of severe shear stress differed for the two types of loads. When the load was along the global x-axis, the severe shear stresses developed in

the binders, but when the load was along the global y-axis, the severe shear stresses developed in the wefts. This indicates that the progression of damage might depend on the direction of load relative to the binders' paths.

Recall Figure 6.31 where the volume of matrix with a σ_{vm} above 80 MPa was shown for uniaxial tension along the x-axis. Figure 6.36 shows the volume of the matrix with a σ_{vm} above 80 MPa for uniaxial tension along the y-axis. First, it is clear that the volume of matrix at such a high stress level is significantly smaller for tension along the y-axis then tension along the x-axis. However, the locations of these high stresses were similar between the two loads. In both cases, the matrix with the highest stresses occurred where in regions surrounding where binders begin or end traversing the thickness of the textile.

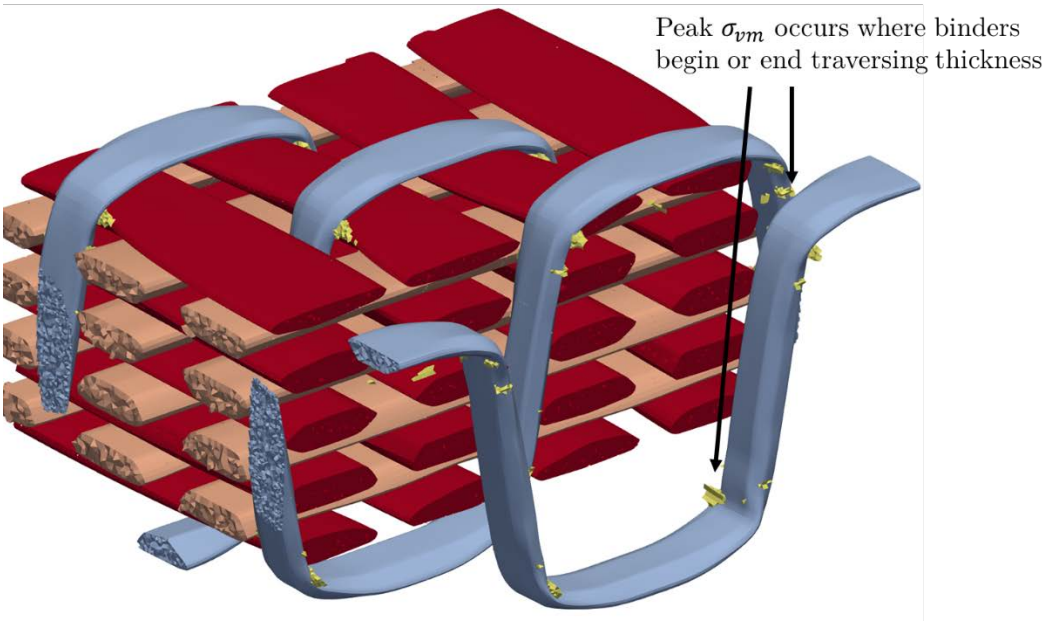

Figure 6.36. Volume of matrix with a $\sigma_{vm} > 8e7\ Pa$ (shown in yellow) with binders and clipped warps/wefts

To understand where a larger percentage of the matrix at high stresses exists, Figure 6.37 shows the volume of matrix with a σ_{vm} over a lower threshold, 60 MPa. With the threshold lower, some differences between tension along the x-axis and tension along the y-axis become clear. Instead of the high stresses localizing around the where binders begin or end traversing the textile thickness as observed for tension along the x-axis, high stresses occur in many more locations for tension along the y-axis, usually between rows of wefts. However, the stress concentrations are limited to regions near binders, just before or after a binder for a path along the y-axis. The reason for this concentration is similar to the explanation given for the matrix stress centration for tension along the x-axis. The transverse Young's modulus of the binders, E_{22}, is significantly larger than the Young's modulus of the matrix. So the binders are taking up much of the load in the region between rows of warps and between rows of wefts, see locations of peak stress in the binder shown in Figure 6.34 for an example. As load is transferred from the binders to neighboring wefts, the matrix experiences a stress concentration.

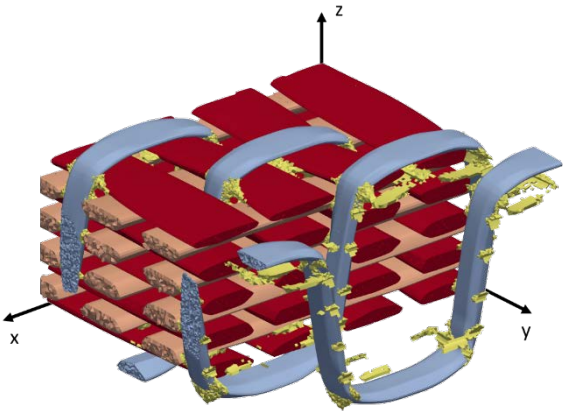

a) Orthogonal view

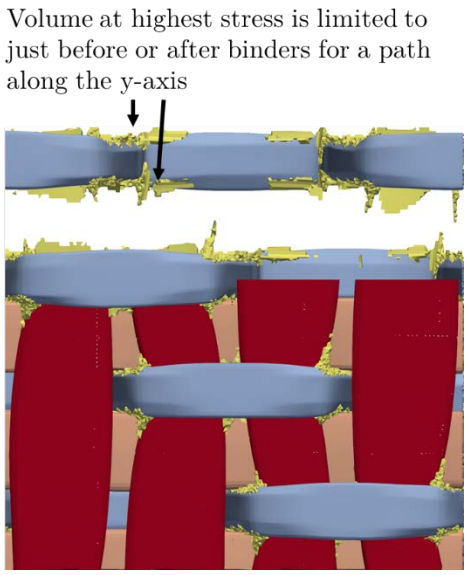

b) Top-down view

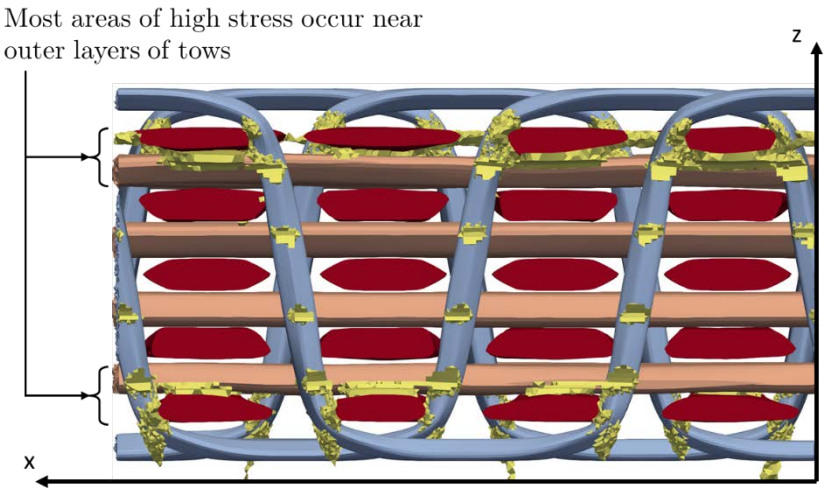

c) View along y-axis

Figure 6.37. Volume of matrix with a $\sigma_{vm} > 6e7\ Pa$ (shown in yellow) with binders and clipped warps/wefts

Considering the peak stresses occur near the edges of binders, Figure 6.38 shows contours of σ_{vm} for slice near one of the binders, refer to Figure 6.38a for the location of the slice. Figure 6.38b shows that the stresses generally become elevated between rows of wefts, with

peak stresses occurring near binders between rows of wefts, which has been discussed. However, Figure 6.38b shows that the stress in the matrix at the bottom of the textile is modestly smaller than stresses near the top of the textile. This could be in part due to the free surface at the top of the textile and in part due to the irregularly shaped tows near the top of the textile. It is unclear which of these two factors is more important.

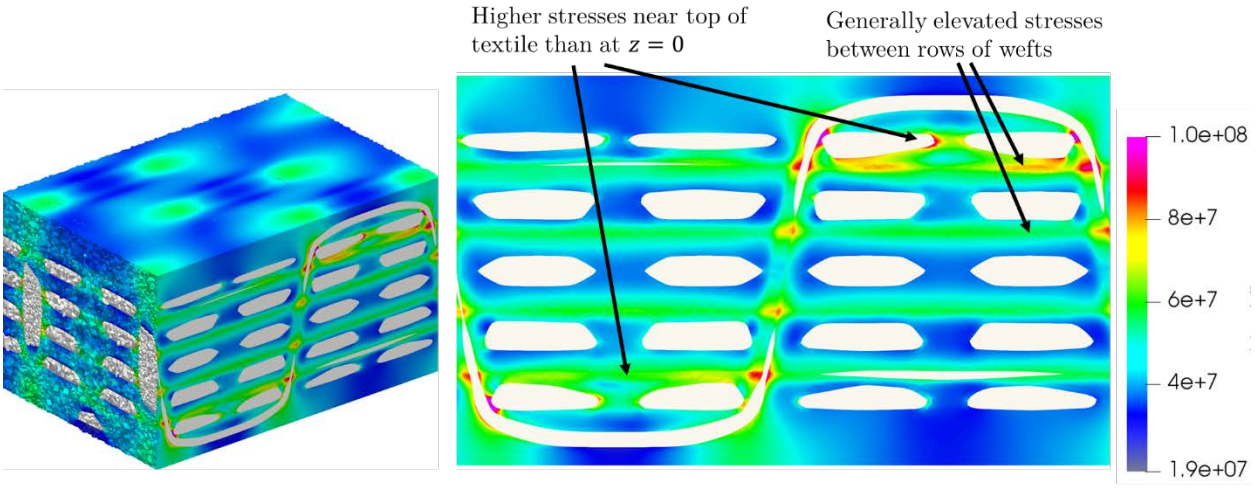

a) Orthogonal view of model showing where slice is taken

b) x-z slice

Figure 6.38. Contours of the von Mises stress, σ_{vm}, within the matrix for a slice that highlights the location of critical σ_{vm}

In-Plane Shear

Figure 6.39 shows three most severe normalized stresses within the clipped analysis region for an applied volume average in-plain engineering shear strain, $\langle \varepsilon_4 \rangle$, of 1%. Note that the engineering shear strain is twice that of the tensorial value, $\langle \varepsilon_{xy} \rangle$. Again, the normalized stresses shown are in the local coordinate system. Overall, Figure 6.39 shows

that the all three types of tows experience a severe longitudinal shear strain ($\hat{\sigma}'_{xy}$). However, only the binders experience a severe transverse shear stress ($\hat{\sigma}'_{yz}$).

Considering the lower longitudinal shear strength of the tows compared to their longitudinal tensile strength, it is expected that the tows would experience a severe longitudinal shear stress when the textile is subjected to an in-plane shear load. Figure 6.39a shows that the warps and wefts both experience similar high values of $\hat{\sigma}'_{xy}$ but with opposite sign. To highlight the locations of high $\hat{\sigma}'_{xy}$ within the warps and wefts, Figure 6.40 shows contours of $\hat{\sigma}'_{xy}$ for two longitudinal slices along the tows. The warps experienced an elevated stress between rows of wefts, and similarly, the wefts experienced an elevated stress between rows of warps. In the gaps between rows of warps or wefts, there is a lower tow volume fraction, and since the matrix has a much lower shear stiffness, the tows carry more of the load. Figure 6.41 shows that the binders experience an elevated $\hat{\sigma}'_{xy}$ along the top and bottom of the textile. In these regions, the binders behave much like warps with the fibers aligned with the global x-axis. Near where the binders begin or end traversing the textile thickness, the longitudinal shear stress decreases, and as shown in Figure 6.39b the binders begin to experience a critical $\hat{\sigma}'_{yz}$.

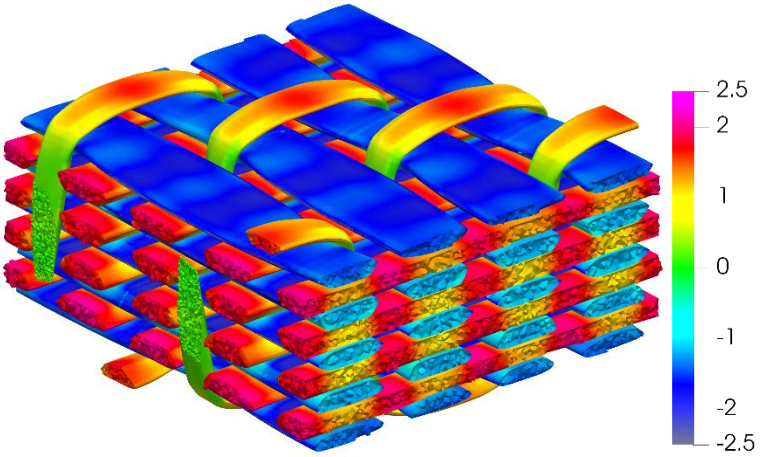

a) $\hat{\sigma}'_{xy}$

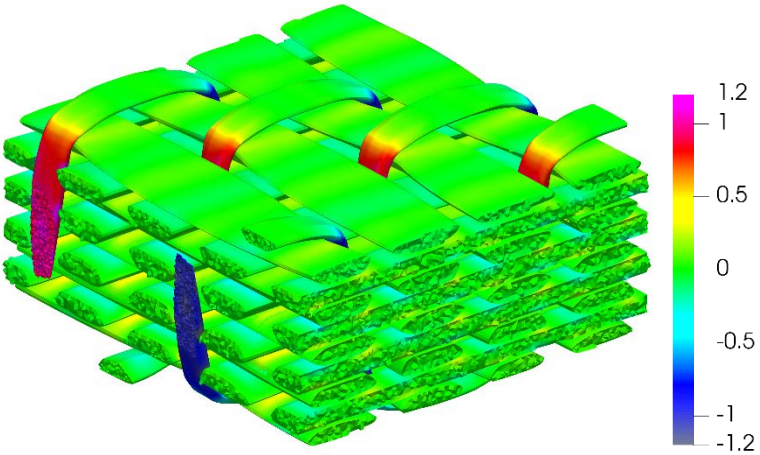

b) $\hat{\sigma}'_{yz}$

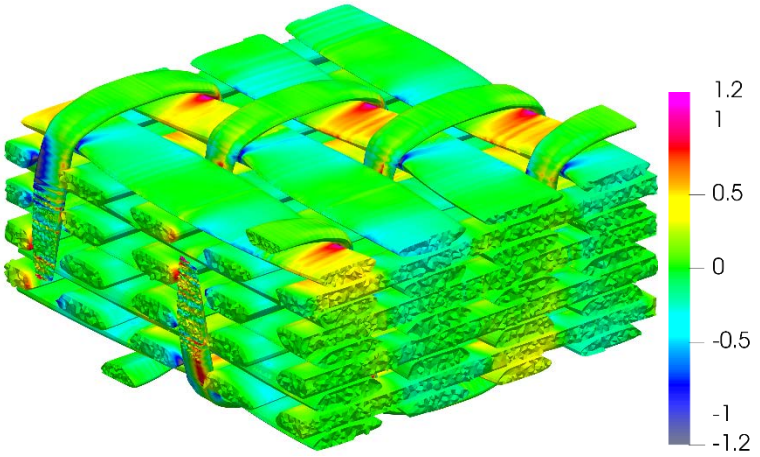

c) $\hat{\sigma}'_{yy}$

Figure 6.39. Contours of three most severe stress components in the local coordinate system normalized by the respective strength for in-plane shear

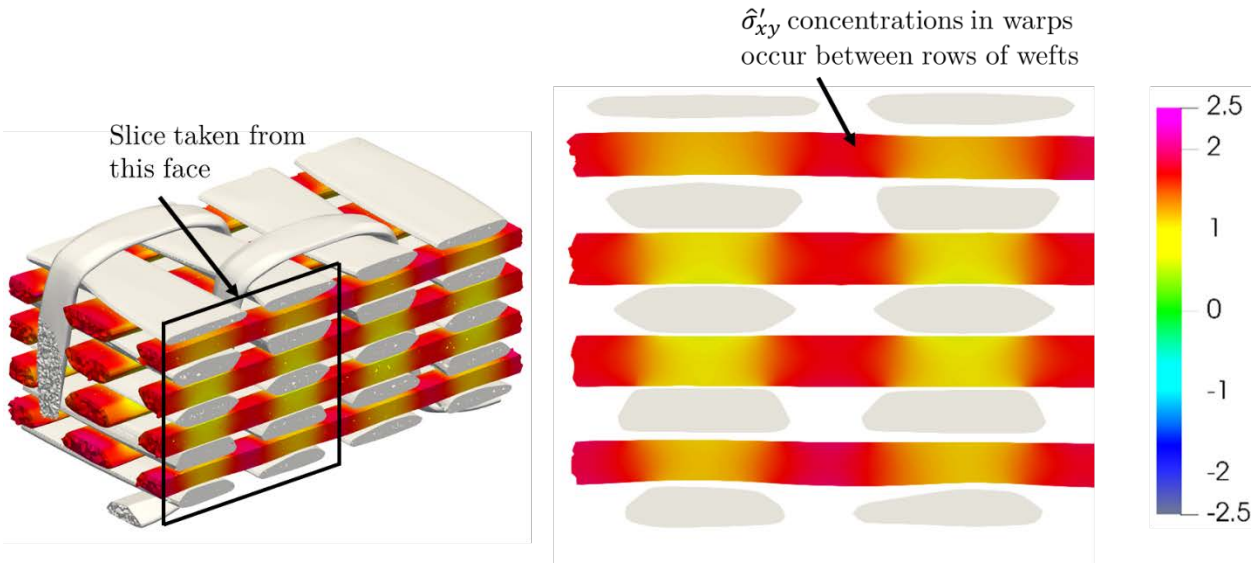

a) Warps

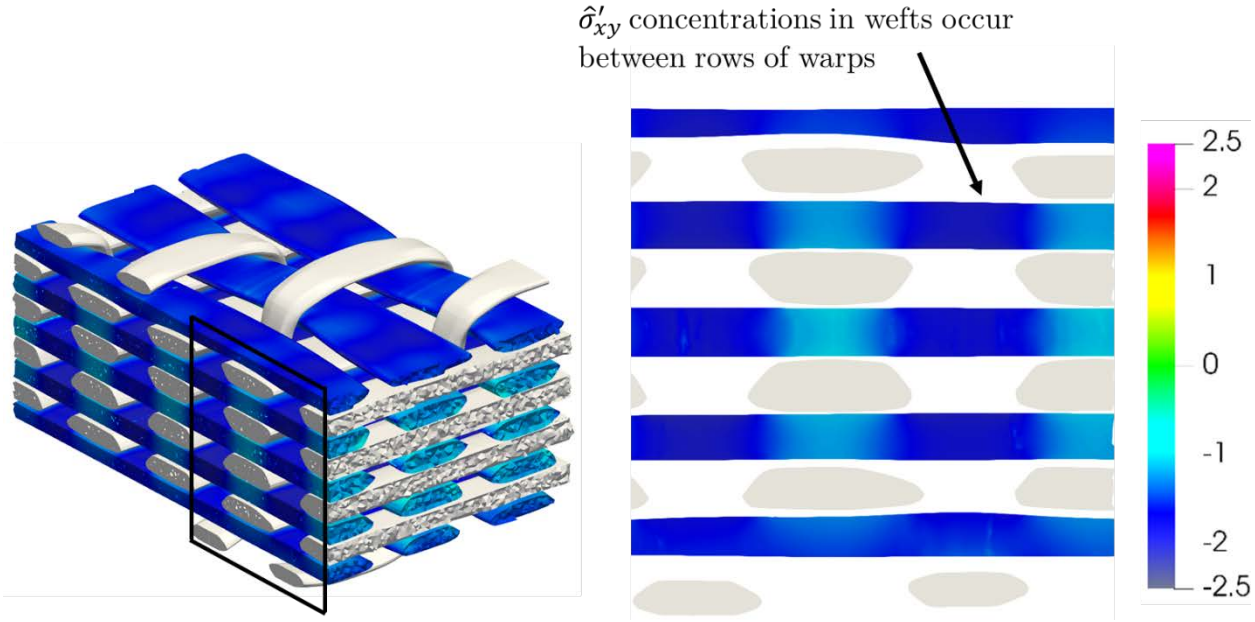

b) Wefts

Figure 6.40. Contours of $\hat{\sigma}'_{xy}$ within warps and wefts showing locations of elevated stress

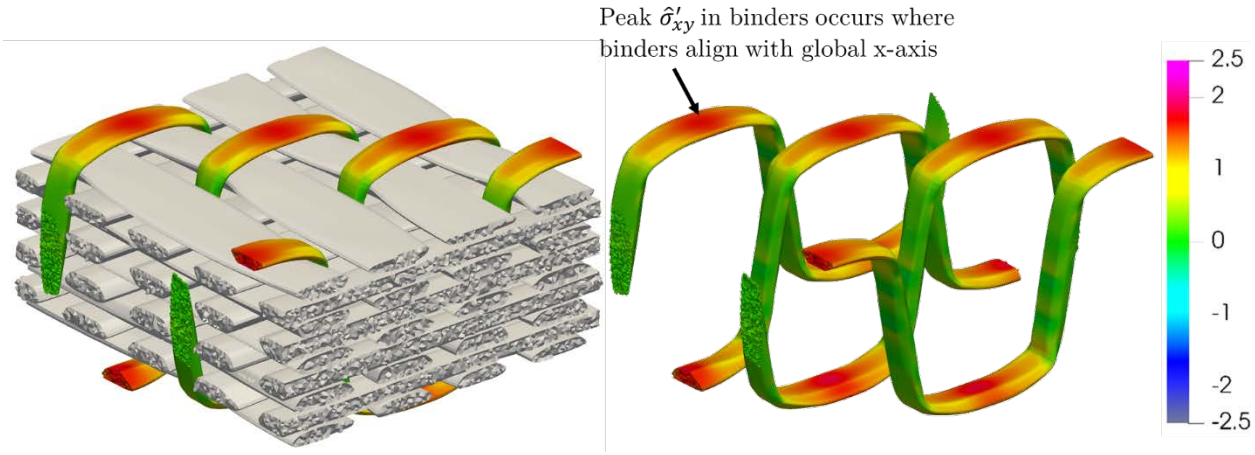

Figure 6.41. Contours of $\hat{\sigma}'_{xy}$ within binders showing locations of elevated stress

As discussed, the binder experiences a critical longitudinal shear stress along the top and bottom of the textile, but Figure 6.42 shows that the binders experience a critical $\hat{\sigma}'_{yz}$ as they traverse through the thickness of the textile. As they traverse the thickness, the local z- and y-axes align approximately align with the global x- and y-axes, respectively. Additionally, there transverse shear modulus of the tows, G_{23}, is much higher than the shear modulus of the matrix, so the binders carry more of the load in this region.

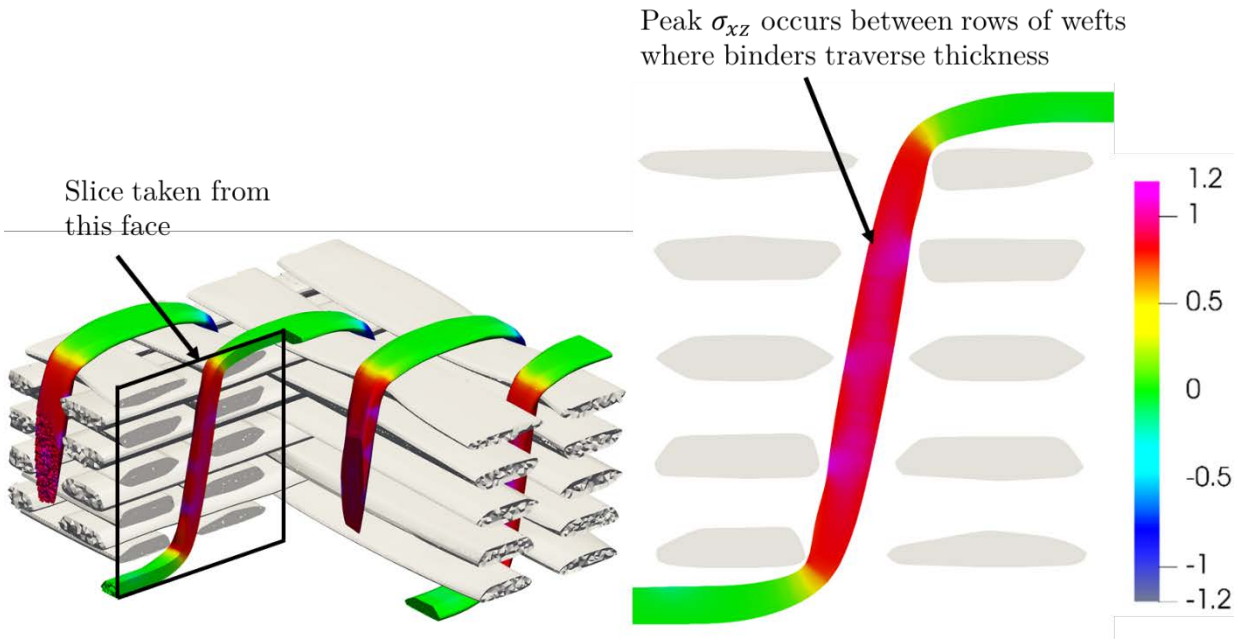

a) Location of slice b) Location of peak $\hat{\sigma}'_{yz}$

Figure 6.42. Contours of $\hat{\sigma}'_{yz}$ within the binders showing locations of elevated stress

Figure 6.39c showed that severe $\hat{\sigma}'_{yy}$ concentrations occurred in the wefts, but Figure 6.43 shows that severe $\hat{\sigma}'_{yy}$ concentrations also occurs within the binders as they transverse through the thickness of the textile and within the warps. For all three types of tows, the concentrations occur where another tow comes close. Due to the in-plane shear load, the tows pivot around each other, creating a tensile stress on one side of the contact and a compressive stress on the other side of the contact. For the wefts, it occurs where binders cross over the outer layer of wefts, as shown in Figure 6.43a. For the warps, it occurs where binders come close while traversing the thickness of the textile, as shown in Figure 6.43b.

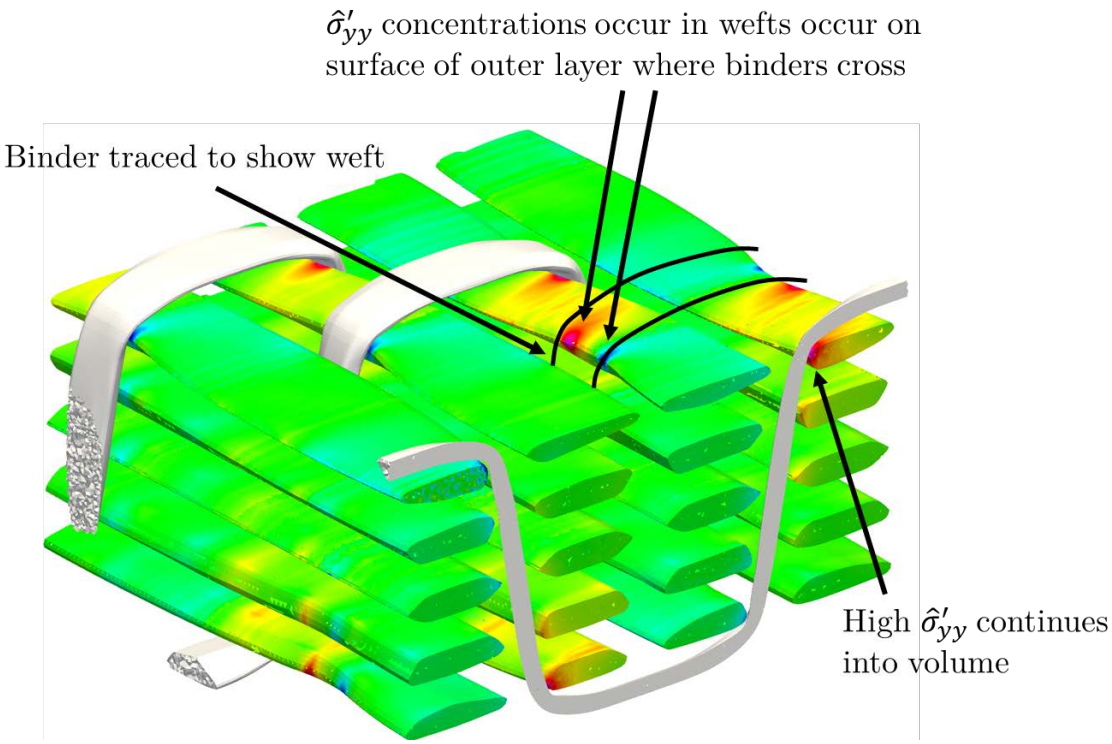

a) Peak σ'_{yy} in wefts

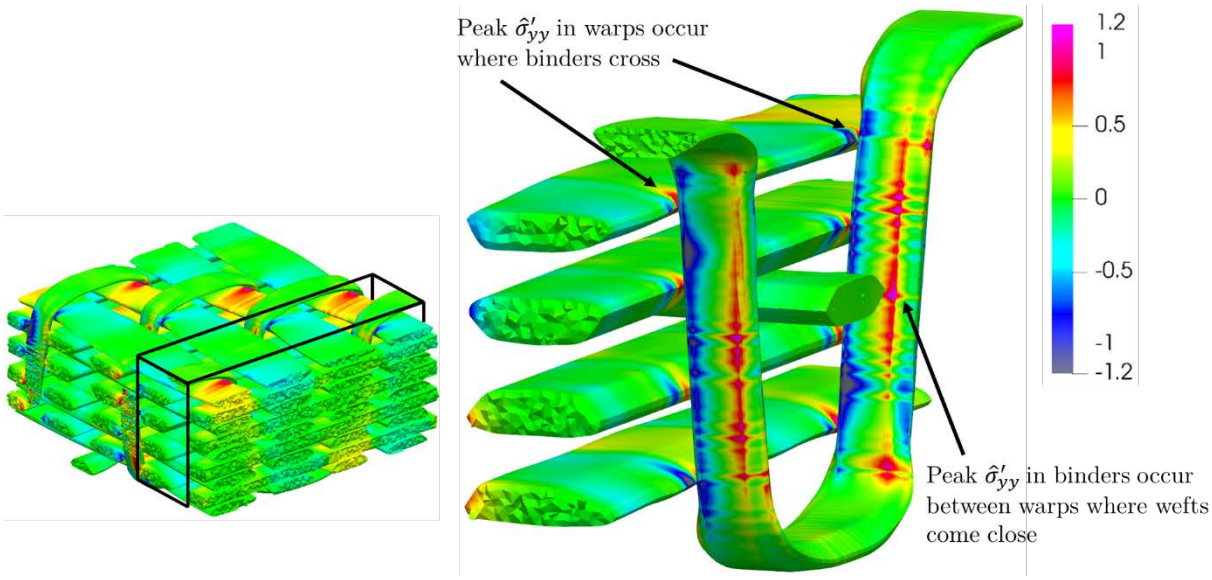

b) Peak σ'_{yy} in binders and warps (too fix the orientation thing on left and solid color weft)

Figure 6.43. Contours of $\hat{\sigma}'_{yy}$ within all three types of tows showing the locations of stress concentrations

Similar to what was shown for the other two load configurations, Figure 6.44 shows the volume of the matrix with a σ_{vm} above 80 MPa. A greater volume of the matrix experienced stress above the threshold than the tensile loads discussed previously. The volume of matrix at a high stress was limited to the gap between the warps, regardless of whether a binder traversed the through the thickness within the gap. If a binder lies within the gap, the severe shear stress in the binder, σ_{yz}, is transferred to the surrounding warps and wefts through the matrix. When there is no binder traversing the thickness, the matrix experiences a high stress as the shear stress is transferred from one row of warps to another and one row of wefts to another. To support this explanation of load transfer, Figure 6.45 shows that the von Mises stress peaks precisely where the binder comes close to the warps and wefts. Additionally, the right gap of Figure 6.45b shows the elevated stress in the gap between rows of wefts and warps where a binder does not traverse the thickness.

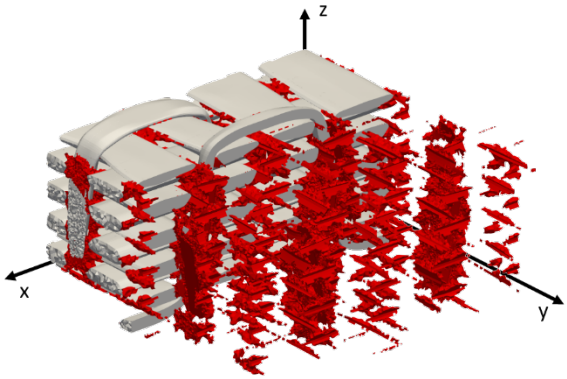

a) Orthogonal view

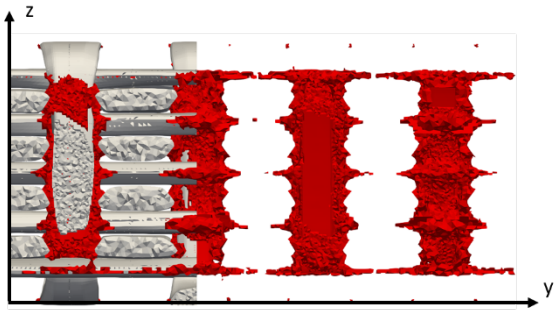

b) View along x-axis

Volume of matrix at a relatively high stress is limited to gap between warps and wefts

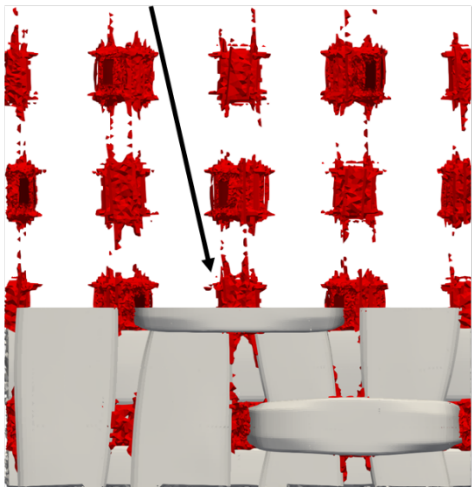

c) View along z-axis

Figure 6.44. Volume of matrix with a $\sigma_{vm} > 8e7\ Pa$ (shown in red) with a clipped section of the tows (shown in gray) for in-plane shear

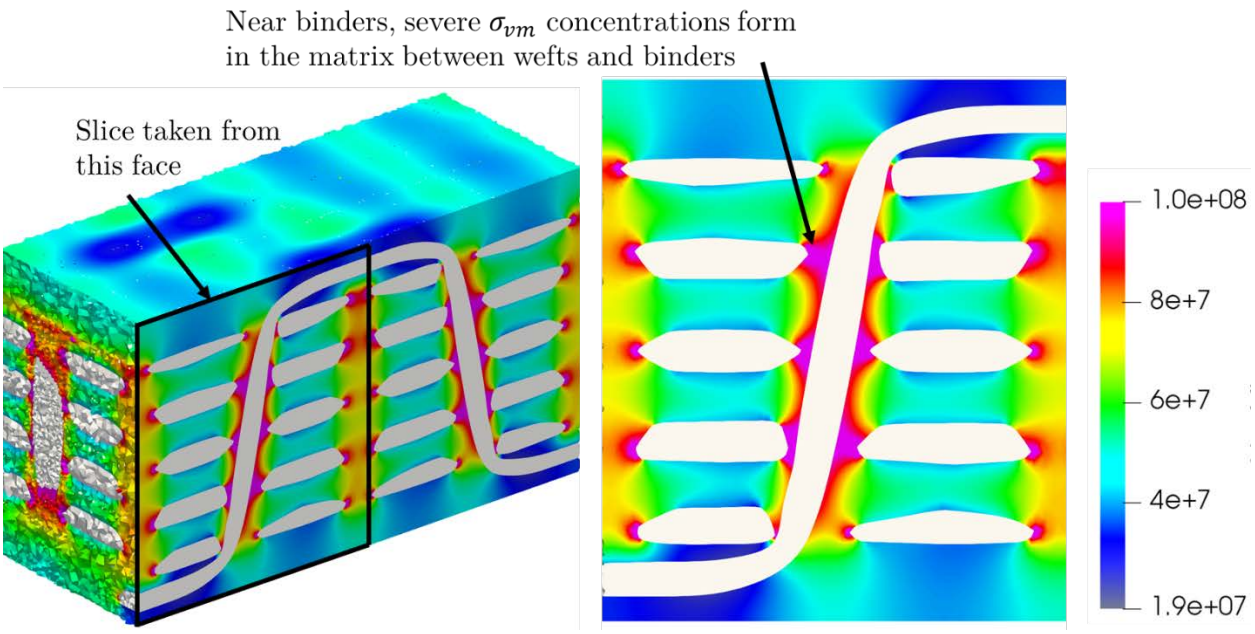

a) Slice along length of a binder

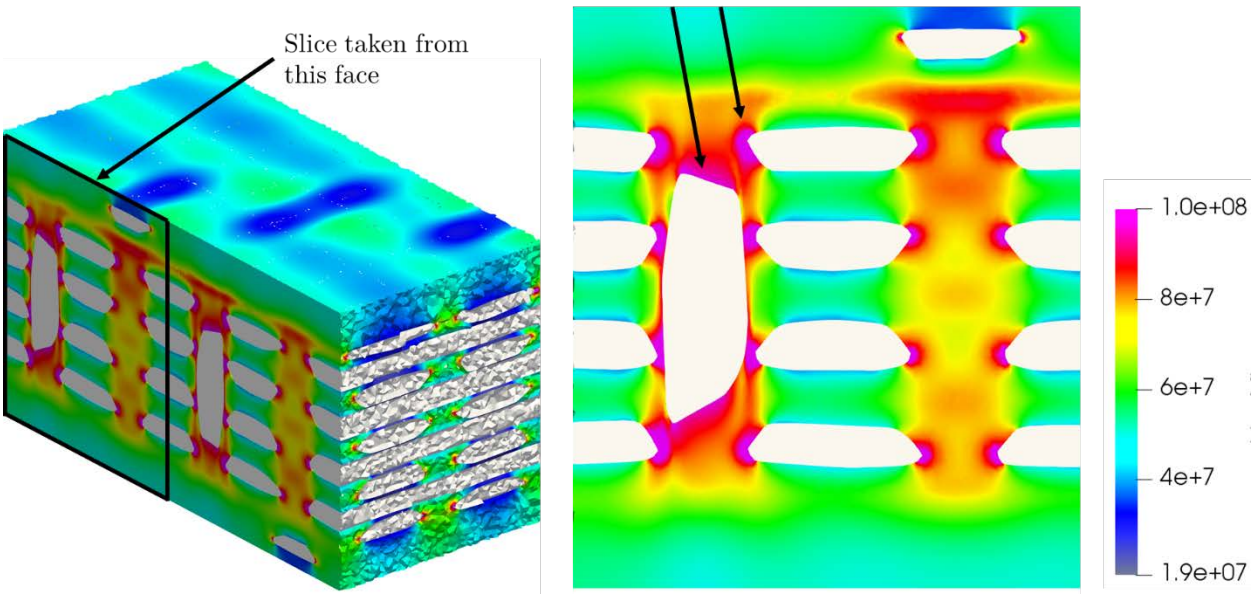

b) Slice perpendicular to binders

Figure 6.45. Contours of the von Mises stress, σ_{vm}, within the matrix for twp slice that highlights the location of critical σ_{vm}

Where load is transferred to and from the binders to neighboring warps and wefts, the cause for the stress contractions can be explained through load transfer. However, Figure 6.46 shows that the cross-sectional shape may be important for the magnitude the stress concentration reaches. The matrix experience much more severe stress concentrations near weft cross-sections that come to a shaper point. An example of this is shown in Figure 6.46c. This highlights an artifact of faceted tow surfaces.

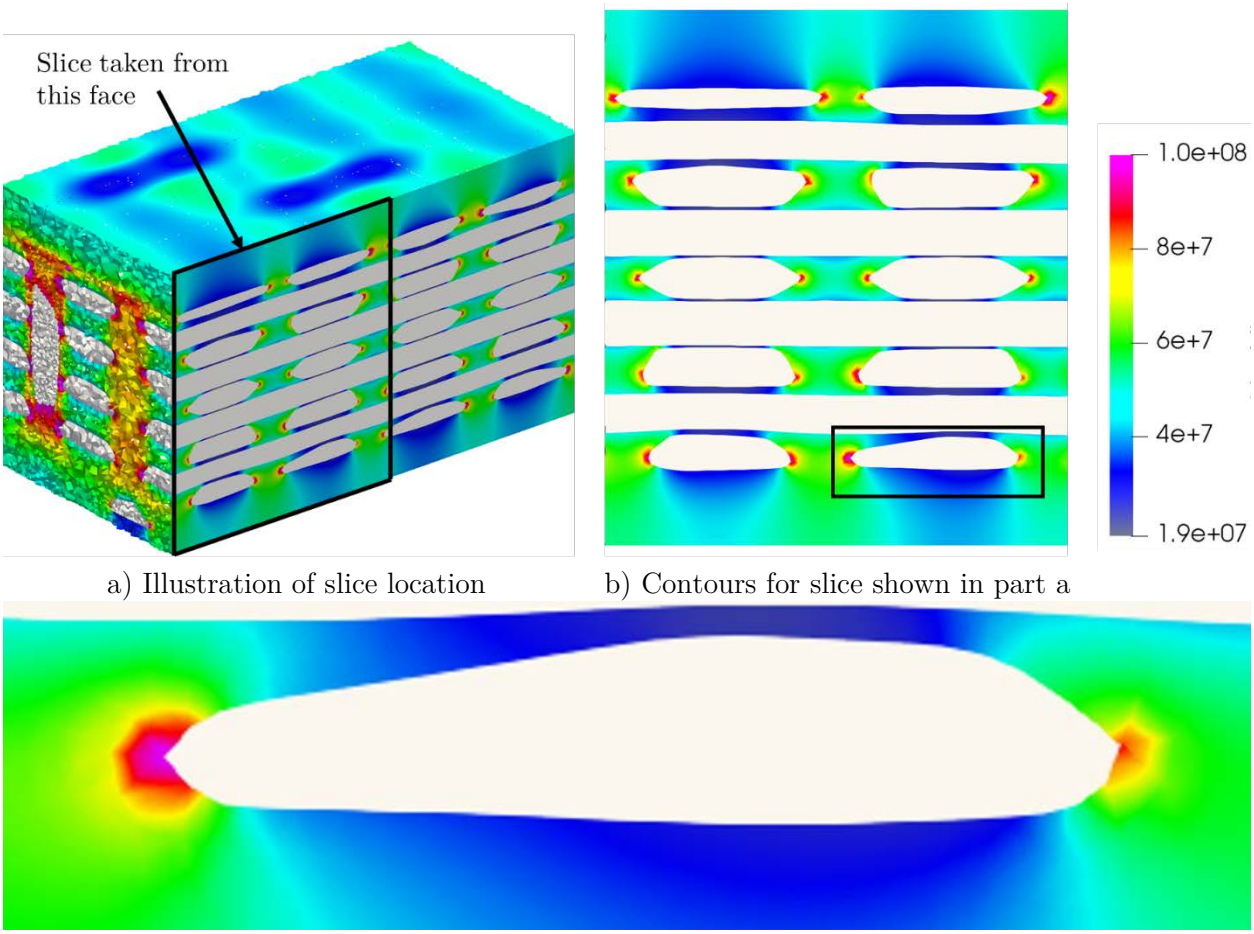

a) Illustration of slice location

b) Contours for slice shown in part a

c) Zoom-in of region boxed in part b

Figure 6.46. Contours of the von Mises stress, σ_{vm}, showing stress concentrations near the edges of wefts

6.4.3 Effect of Binders

Within the complex architecture of the 3D textile, the binders bind the layers of the warps and wefts and have been shown to have a significant effect on the damage resistance of the textile, especially to delamination. However, in the discussion of local stress states, there are several factors to consider, including the existence of the binders and waviness in the tows. The binders can carry load within the textile, possibly affecting the locations of the stress concentrations due to the redistribution of load. To understand the effect binders have on the locations of stress concentrations, this section considers the cases where the binders are assumed to be matrix material instead of orthotropic tows. It should be noted the tow geometry remains unchanged, so the effect of the tows on the tow geometry still exists, leaving waviness in the warps and wefts.

Uniaxial Tension Along X-Axis

Under uniaxial tension along the x-axis, the transverse tensile stress, σ'_{yy}, in the wefts was shown to be the most severe component of stress within the textile. When the binders were replaced with matrix material, the σ'_{xx} within the warps increased slightly, and σ'_{yy} within the wefts decreased significantly. Figure 6.47 shows a comparison of the $\hat{\sigma}'_{yy}$ contours showing effect of binders for tension along x-axis. The peak stress in the wefts decreased by about 50%. When the binders were replaced with matrix, the distribution of stress within the textile changed dramatically, with the warps carry a very large amount of the load and resulted in σ'_{xx} in the warps becoming the most severe stress. Additionally, the

other stress components investigated when the binders included, σ'_{zz} and σ'_{xz}, become far less significant.

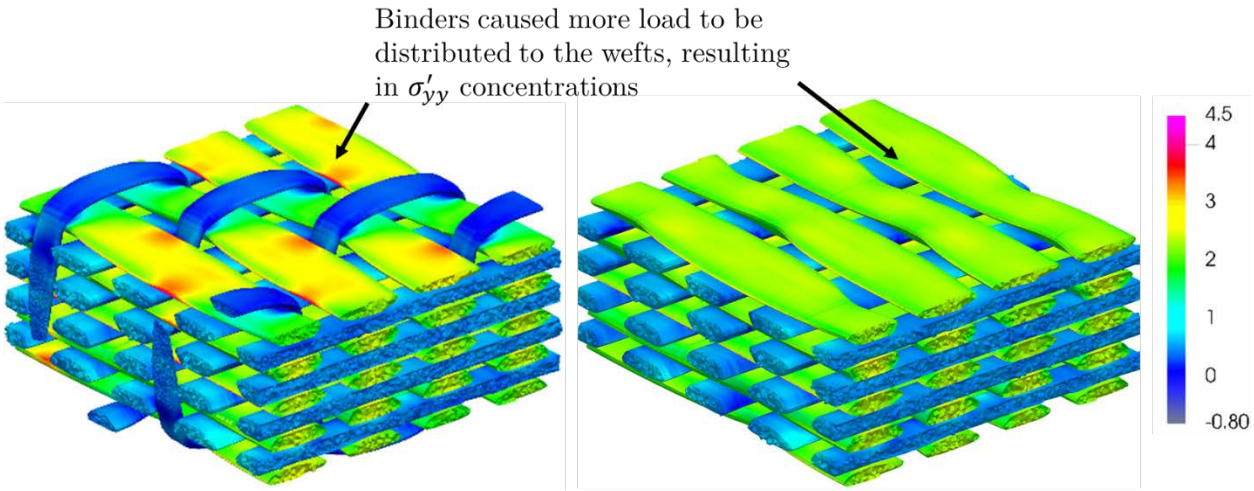

Figure 6.47. Comparison of $\hat{\sigma}'_{yy}$ (σ'_{yy} normalized by the respective strength) contours showing effect of binders for tension along x-axis

Figure 6.48 shows the top-down view of the warps with contours for σ'_{xx} after the binders has been replaced with matrix. The figure shows that the longitudinal stress varies significantly, due to the waviness in the warps. Because the longitudinal stiffness of the tows is so much larger than the transverse stiffness, the predicted longitudinal stress is very sensitive to waviness in the tows aligned with the load direction. This highlights the importance of accounting for tow waviness within the complex tow architecture of 3D textiles.

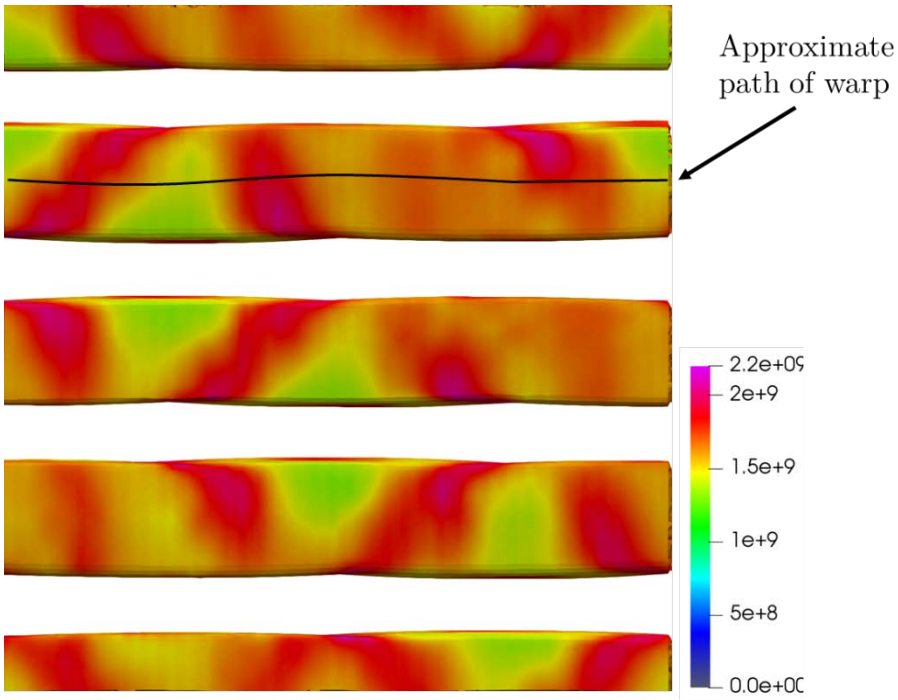

Figure 6.48. View along z-axis of warps with contours for σ'_{xx}

In-Plane Shear

Under in-plane shear, the longitudinal shear stress, σ'_{xy}, in the warps and wefts was shown to be the most severe component of stress. When the binders were removed and replaced with matrix material, the magnitude of σ'_{xy} at the locations of highest stress decreased by 23%,, and the locations of σ'_{xy} concentrations remained the unaffected. To illustrate that eh locations of critical stress remain unchanged, Figure 6.49 shows a comparison of σ'_{xy} contours for slices within the wefts and warps for the results shown in Figure 6.40 (left) and the case of the binders replaced with matrix (right).

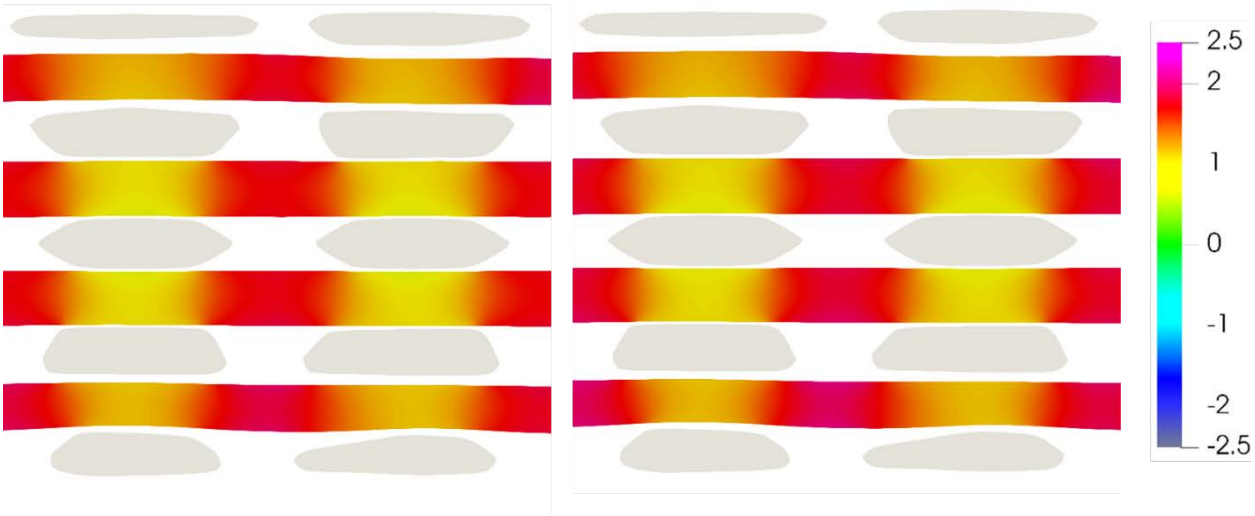

a) Warps

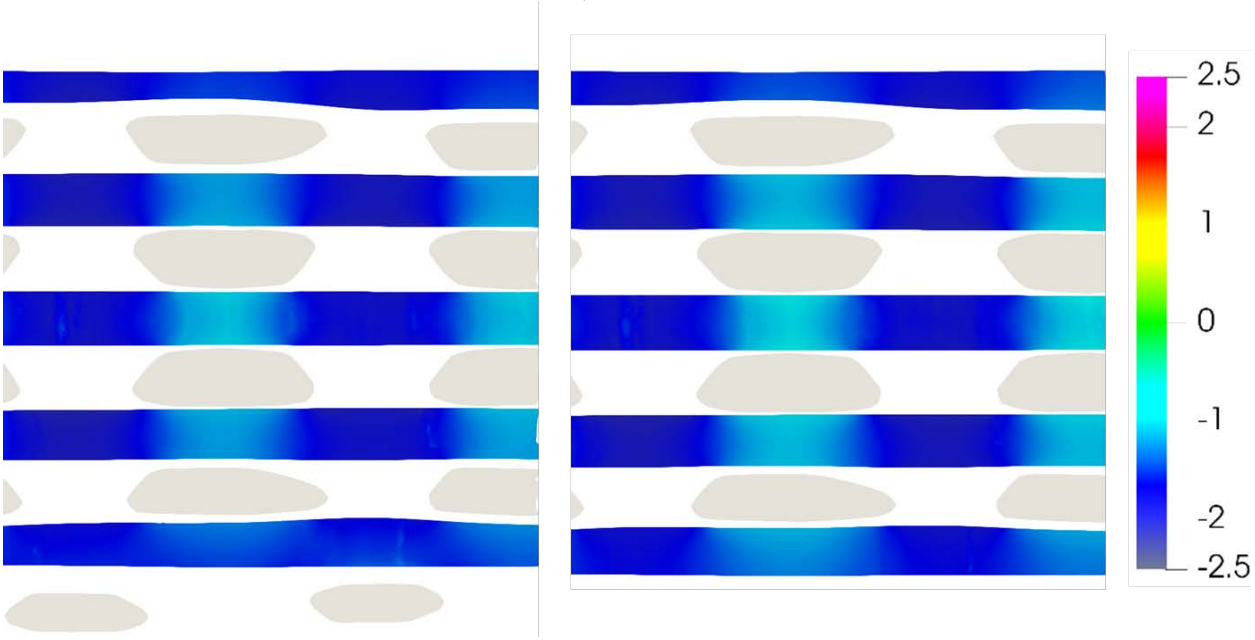

b) Wefts

Figure 6.49. Comparison of σ'_{xy} contours for slices within the wefts and warps for the results shown in Figure 6.40 (left) and the case of the binders replaced with matrix (right)

Based on Figure 6.39c, the σ'_{yy} concentrations could be due to either the waviness or the redistribution of load by the binders. Figure 6.50 shows that the elevated stresses remain in the wefts after the binders have been removed, but the magnitude of σ'_{yy} does significantly decrease. The highest values of σ'_{yy} occurred near the location of binders crossed the wefts when binders were included, and these concentrations are significantly reduced when the binders are removed. Across the rest of the wefts, the elevated stress remain when the binders were replaced with matrix. Figure 6.51 shows that the locations of elevated σ'_{yy} occur due to waviness in the wefts. The waviness in the tow paths would be ignored by idealized models, which reinforces the merit of nonidealized 3D textile models due to the complex tow architectures.

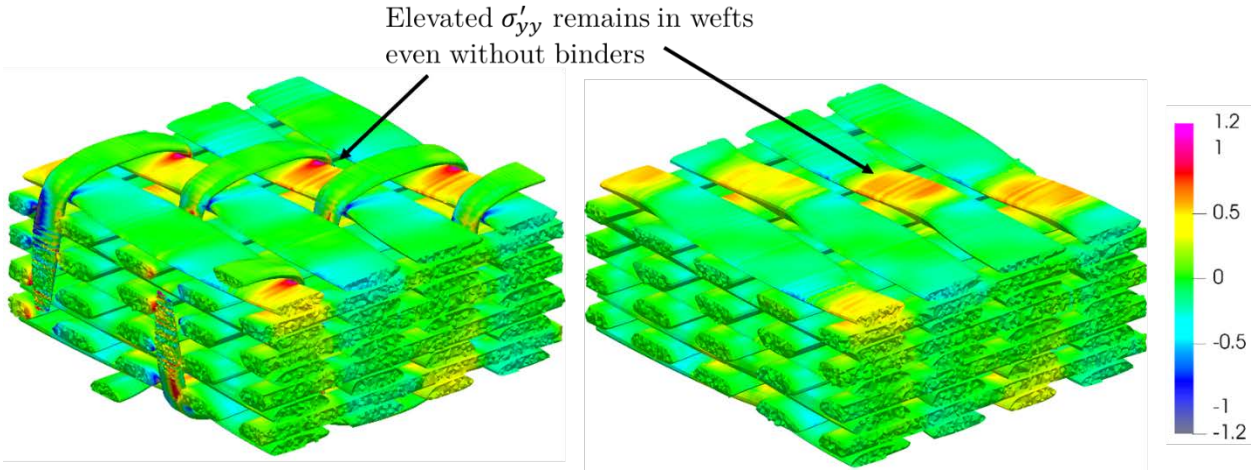

Figure 6.50. Comparison of $\hat{\sigma}'_{yy}$ (σ'_{yy} normalized by the respective strength) contours showing effect of binders for in in-plane shear

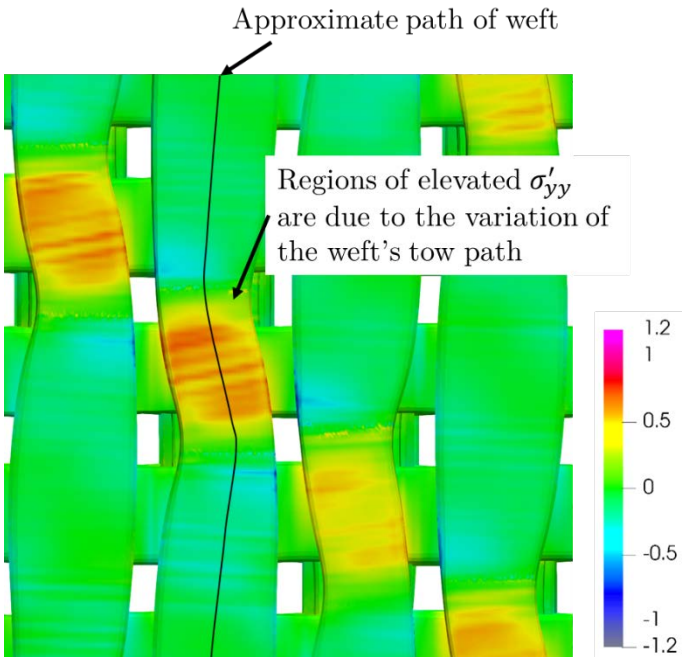

Figure 6.51. Top-down view of $\hat{\sigma}'_{yy}$ contours illustrating the effect of tow waviness on the observed stress concentrations

6.4.4 Effect of Modeling Parameters for Uniaxial Tension Along X-Axis

There are many different model parameters that could be changed in these studies. This section explore the effect of considering two: accounting for the variation of the fiber volume fraction within the tows and accounting for plasticity in the matrix. In the previous results of this chapter, uniaxial tension along the x-axis was shown to result in the most critical stresses at an applied 1% for the configurations considered. Consequently, tension along the x-axis configuration was selected for this section.

Effect of Modeling the Variation of Local Fiber Volume Fraction

As discussed in chapter 2, the digital chain method used within VTMS to create the nonidealized textile models resulted in a tow architecture with significant variation in the cross-sectional areas along the tow paths. In reality, tows experience some amount of variation in the cross-sections, but about the same number of fibers exist for any slice, which causes the fiber volume fraction to vary along the tow path. The range of variation of cross-sectional area in the model indicated that the variation of fiber volume fraction could be significant. In the previous sections, the fiber volume fraction was assumed to be a constant 60%. The effect of accounting for the variation of the fiber volume fraction will be assessed in this section using the distribution of V_f shown in Figure 6.52, which was calculated using the variation of the cross-sectional area using a method discussed in chapter 2. It should be noted that the fiber volume fractions were determined on an element basis, so one value of V_f will be used for all quadrature points within an element.

In chapter 4, a series of many microscale analyses were used to characterize how the elastic properties of the tows vary with respect to the fiber volume fraction. The result of that work was a constitutive law based on spline fits of the mean effective properties for 18 discrete volume fractions. The properties exhibited some anisotropy and variation, which is expected with random fiber arrangements, but the mean properties were smooth and forced to be transversely isotropic to avoid introducing more complicity, since the effect of local anisotropy within tows has not been well characterized. Using the spline fits of the

smooth effective properties, the linear elastic stiffness matrix can be created for each element within the tows using given the local V_f.

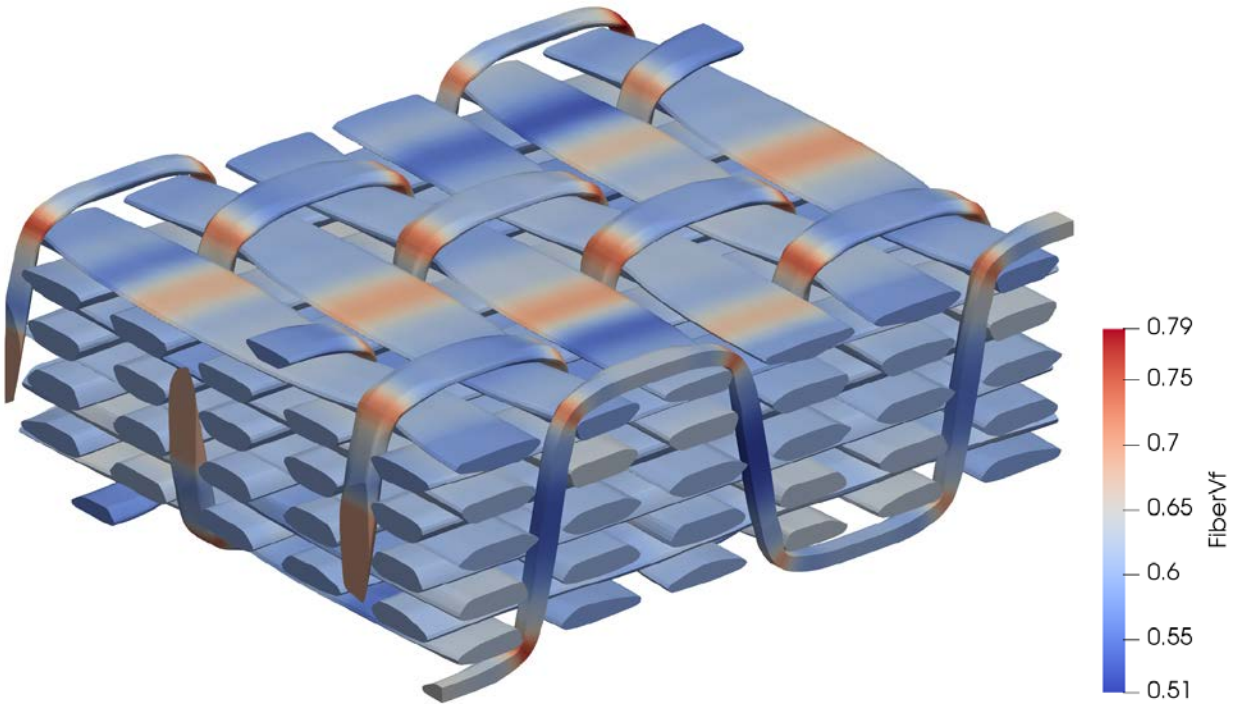

Figure 6.52. Contours of local fiber volume fraction in the tows with a nominal fiber volume fraction of 0.6

Figure 6.53 shows a side-by-side comparison of the σ'_{yy} contours for the cases where a constant V_f (left) and the varying V_f (right) is used. The figure shows that accounting for the variation of the local V_f within the tows results in a significantly larger region of severe σ'_{yy} within the tows. Recalling the contours in Figure 6.52, the locations of the worsened stress highlighted in Figure 6.53 occurs where the weft has a higher V_f. The higher V_f

causes the transverse stiffness of the wefts to increase locally, leading to the stiffer region carrying more of the load. Figure 6.54 shows contours of the percent change of σ'_{yy} within the wefts for regions with a σ'_{yy} higher than 40 MPa. Other regions, including two binders for context, are shown in black. The threshold was chosen to avoid a misleading plot due to large percent changes in regions near zero stress. The figure shows that there is a 20% increase in the magnitude of the stress in the region experiencing a severe σ'_{yy}. Since the locations of the most severe stress concentrations in the model increased by 20% when the variation of the fiber volume fraction was accounted for, it could be important to account for the variation of V_f if the local stresses or progression of damage are to be accurately predicted.

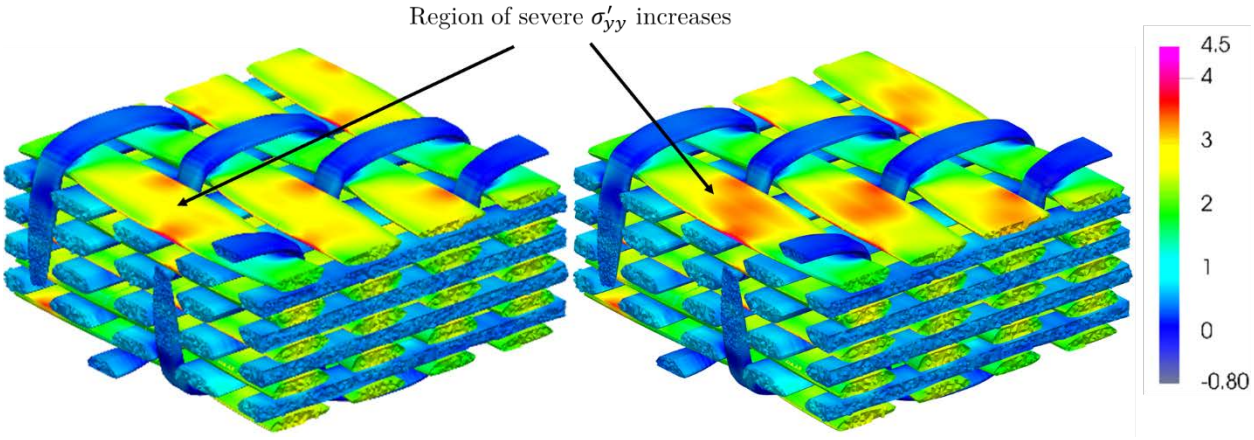

a) Constant V_f of 60% b) Varying V_f

Figure 6.53. Comparison of $\hat{\sigma}'_{yy}$ (σ'_{yy} normalized by the respective strength) contours between the two cases where tows have a constant and varying fiber volume fraction

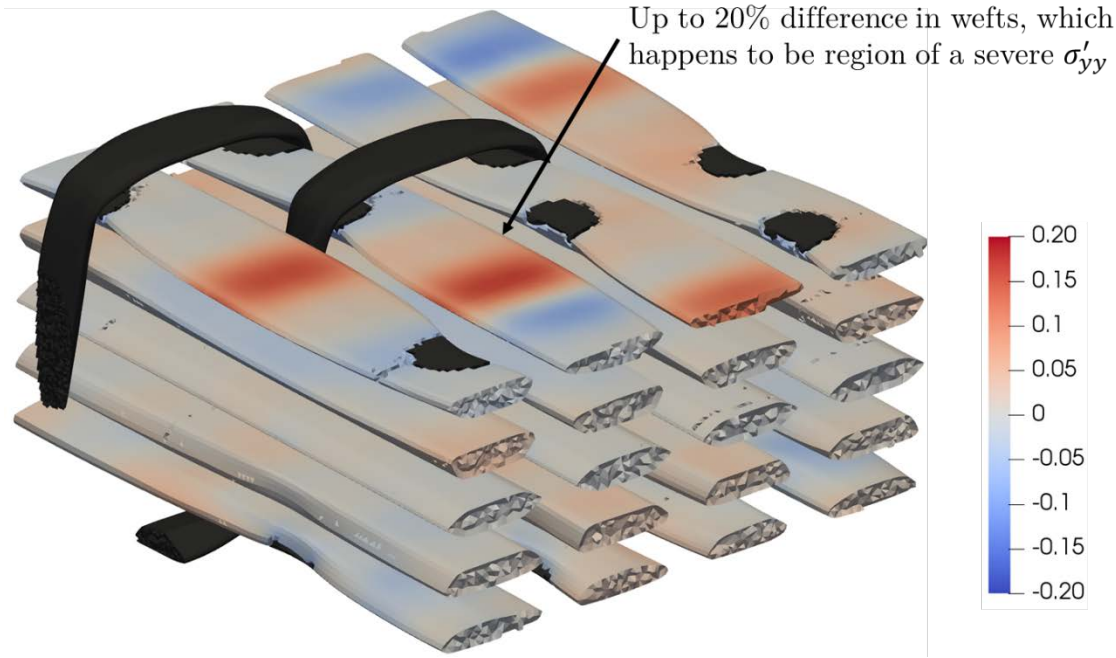

Figure 6.54. Contours of the percent change of σ'_{yy} for the volume experiencing a $\sigma'_{yy} > 40\,MPa$ (all other region, including two binders for context, are shown in black)

Figure 6.55 shows a side-by-side comparison of the σ'_{zz} contours for the cases where a constant V_f is used and the varying V_f is accounted for. Unlike, σ'_{yy}, the figure shows that accounting for the local variation of V_f results in no noticeable change in the contours. Similarly, Figure 6.56 shows a side-by-side comparison of the contours for a slice of the binders containing the peak σ'_{zz} stresses. Some very small difference can be observed, but the contours remain almost unchanged due to account for the variation of V_f. To quantify the change in σ'_{zz} Figure 6.57 shows the percent change for the volume of the binders experiencing a σ'_{zz} higher than 40 MPa, with the remaining volume shown as black. Where the peak σ'_{zz} stresses occur, there is about a 5% difference, and the region of with an

evaluated σ'_{zz} changed between 5% and 10%. Accounting for the local variation of V_f had a much less significant effect on σ'_{zz} as σ'_{yy}, which was the most severe component of stress.

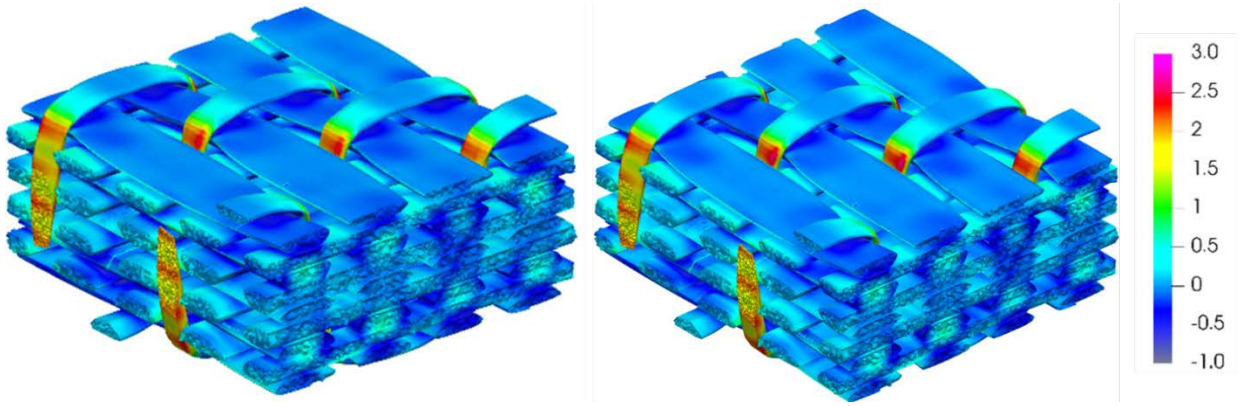

a) Constant V_f of 60% b) Varying V_f

Figure 6.55. Comparison of $\hat{\sigma}'_{zz}$ (σ'_{zz} normalized by the respective strength) contours for the two cases where tows have a constant and varying fiber volume fraction

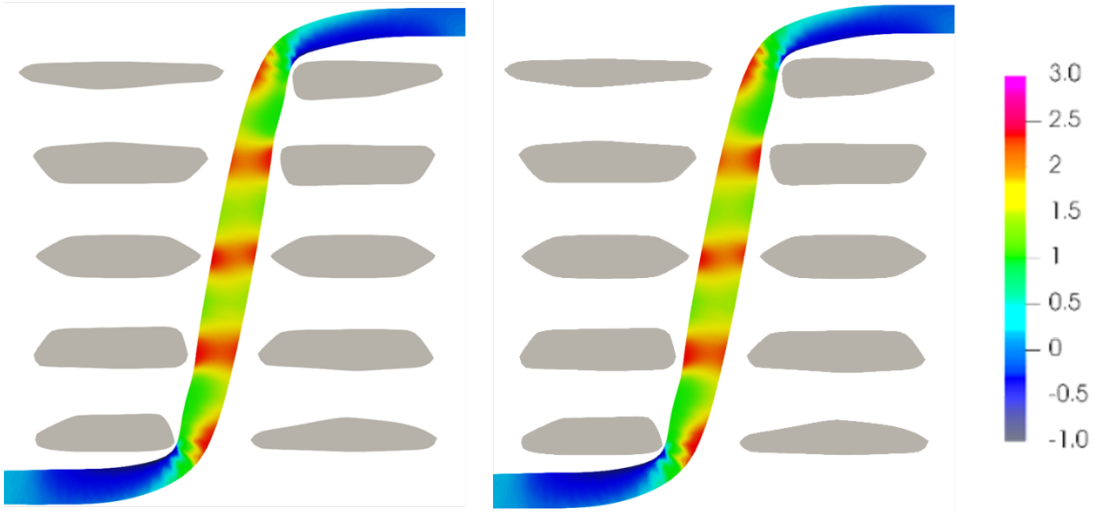

a) Constant V_f of 60% b) Varying V_f

Figure 6.56. Comparison of $\hat{\sigma}'_{zz}$ contours for an x-z cross-section of a selected binder for the two cases where tows have a constant and varying fiber volume fraction

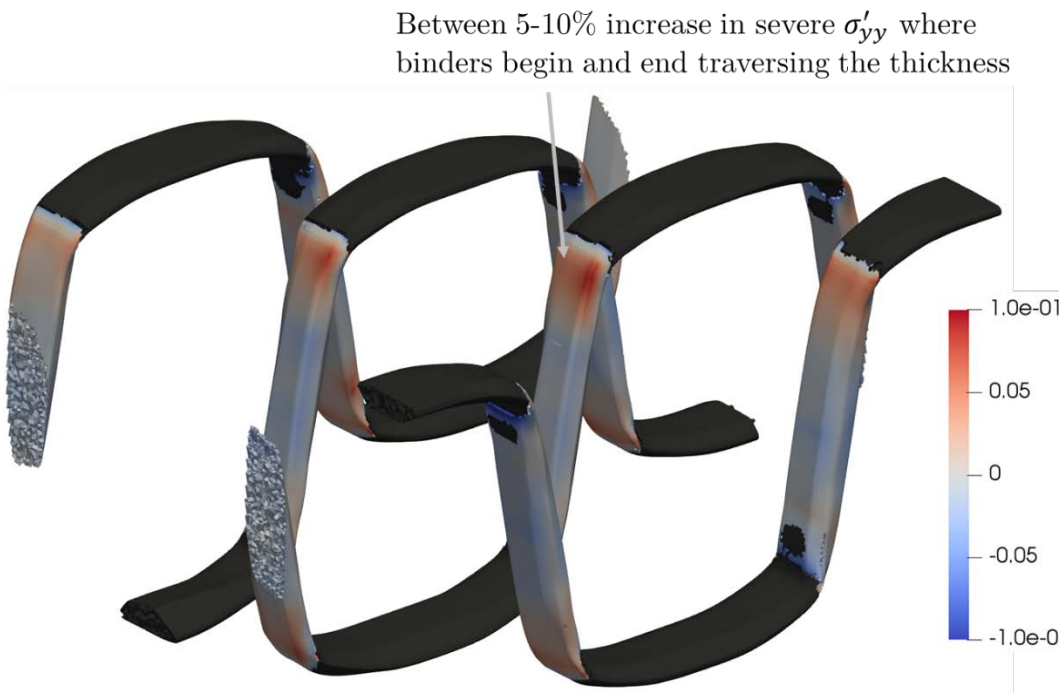

Figure 6.57. Contours of the percent change of σ'_{zz} for the volume experiencing a $\sigma'_{zz} > 40\,MPa$ (all other region are shown in black)

Within the matrix, accounting for the variation of V_f within the twos had no significant effect on the regions experiencing high stress. To illustrate the non-effect, Figure 6.58 shows a side-by-side comparison of the volume of matrix with a σ_{vm} greater than 80 MPa, which shows no discernable difference.

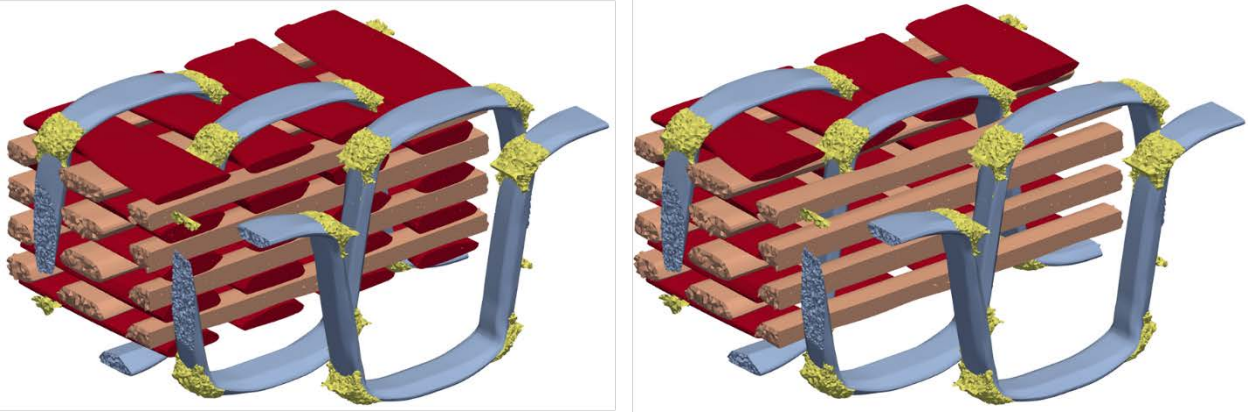

a) Constant V_f of 60% b) Varying V_f

Figure 6.58. Comparison of volume of matrix with a $\sigma_{vm} > 8e7\ Pa$ (shown in yellow) with binders and clipped warps/wefts for the two cases where tows have a constant and varying fiber volume fraction

Effect of Plasticity in the Matrix

In the previous linear elastic analyses, plasticity in the matrix was ignored. For tension along the x-axis, the locations stress concentrations in the matrix occurred in the same locations of severe σ'_{xz} in the binders and near the areas in the wefts with a severe σ'_{yy}. Plastic deformation within the matrix may lead to a relaxation of the stress at these critical locations. To study the effect of plasticity in matrix on the severe stresses within the tows, a model with a parabolic yield criterion and a strain hardening law, which is described in chapter 3, was used. For this study of the effect of plasticity in the matrix, the matrix in the neat matrix pocket of the textile was assumed to be an 8552 epoxy resin with the assume properties shown in Table 6.6.

Table 6.6. Assumed elastic and plastic properties for neat matrix pocket of textile

Elastic Properties	
E	4.7 GPa
ν	0.35

Plastic Properties	
ν_p	0.3
σ_{yt}	65 MPa
σ_{yc}	120 MPa
$H_C = H_T$	200 MPa
$\eta_C = \eta_T$	0.2

For the case of tension along the x-axis resulting in a 1% volume average strain, $\langle \varepsilon_{xx} \rangle$, Figure 6.59 shows contours of the effective plastic strain in the matrix, which shows that most of the plastic deformation occurs where binders begin or end traversing through the thickness of the textile. As expected, the locations of plastic deformation match well with the regions of the matrix experiencing high stress in linear analysis, see Figure 6.31. Figure 6.60 shows the contours of the effective plastic strain for a x-z cross-section centered on a binder, the same location used in Figure 6.32. The figures shows that most of the plastic deformation occurs very close to the binder, which experiences about 1% effective plastic strain. Considering the boundaries result in a 1% volume average strain, 1% effective plastic strain locally indicates significant localization of strain.

Figure 6.59. Contours of the effective plastic strain in the matrix (binders shown in gray)

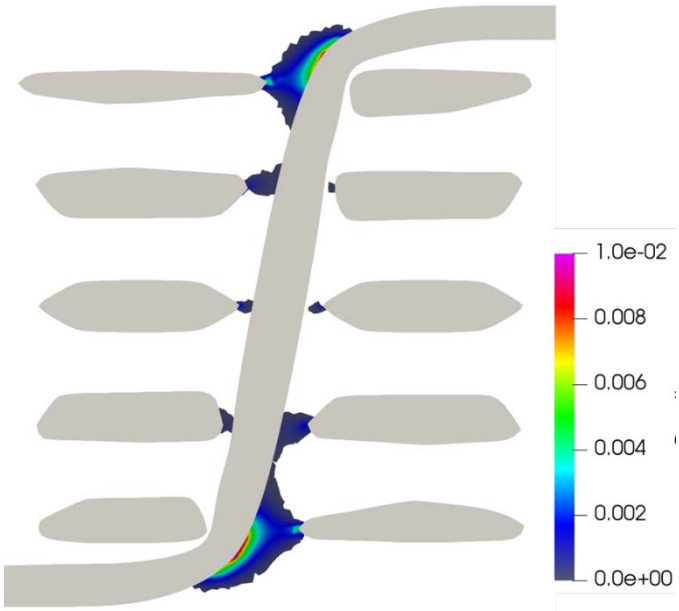

Figure 6.60. Contours of effective plastic strain in matrix for an x-z cross-section centered on a binder

To understand if the plastic deformation affects the locations of severe stress in the tows, Figure 6.61b shows the percent change of σ'_{yy} within the tow for a selected region shown in Figure 6.61a. As Figure 6.61a shows, the critical locations of σ'_{yy} occur with the wefts, but the wefts experience no significant change in σ'_{yy} when plastic deformation occurs. Since the effective plastic strain is so localized to the binders, which was shown in . Figure 6.60, the stress concentrations in the wefts remain unaffected.

Similarly, Figure 6.62 shows the percent change for σ'_{xz} from when linear elasticity is assumed to the case with plasticity modeled. Since the plastic deformation localizes near locations of severe σ'_{xz}, σ'_{xz} is reduced by a modest 3% at the stress concentrations.

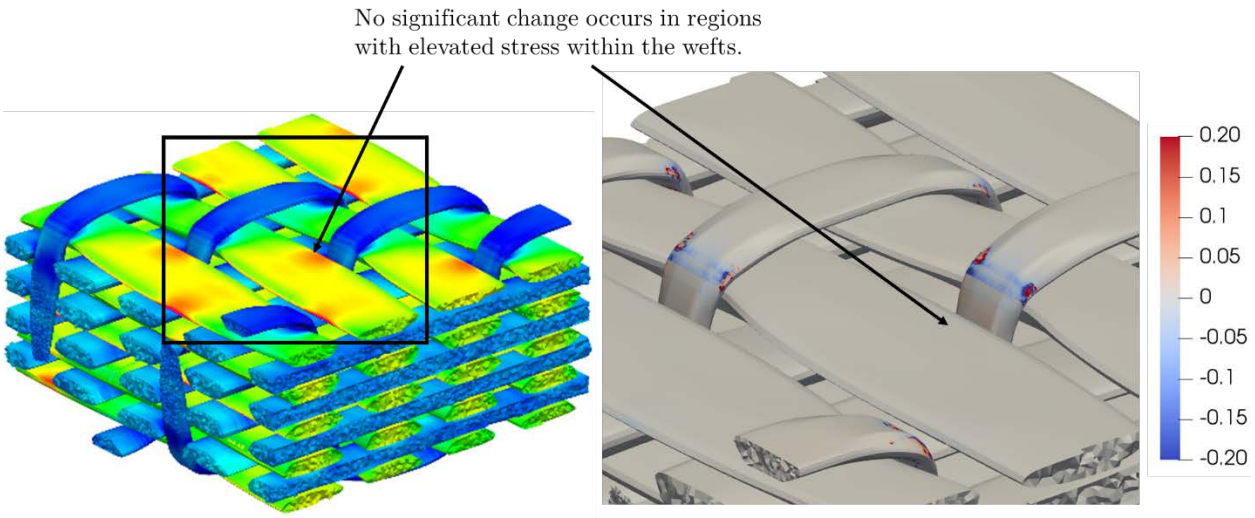

a) Contours of severe σ'_{yy} in the tows (Figure 6.28a)

b) Percent change of σ'_{yy} in the tows

Figure 6.61. Contours of the percent change of σ'_{yy} between the case when plasticity in the matrix pocket is modeled and linear elasticity is assumed for a selected region

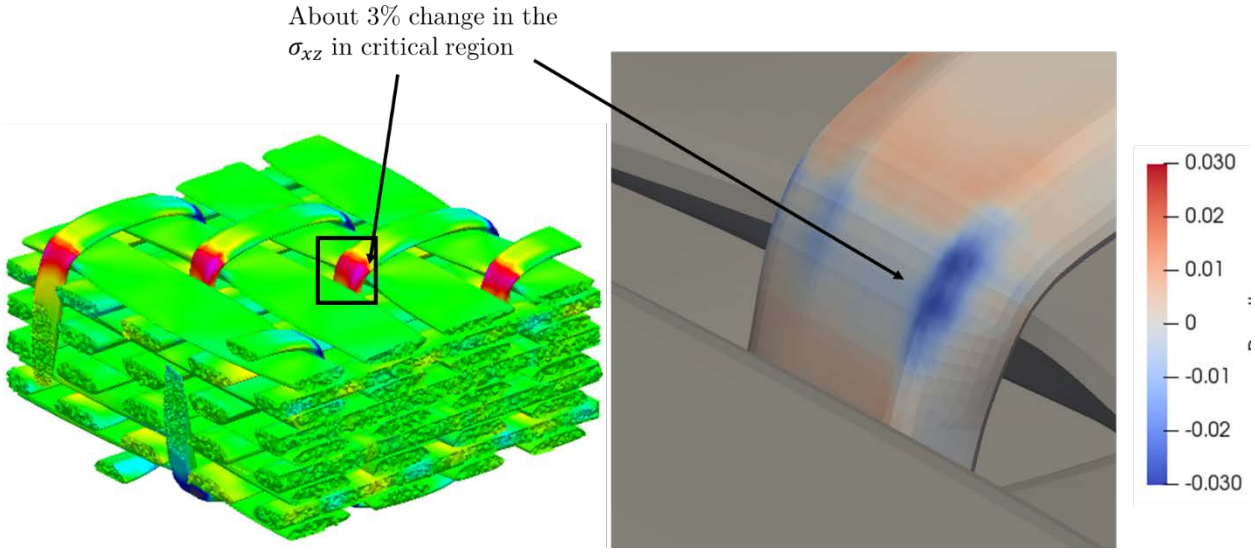

a) Contours of severe σ'_{xz} in the tows (Figure 6.28c)

b) Percent change of σ'_{xz} in the tows

Figure 6.62. Contours of the percent change of σ'_{xz} between the case when plasticity in the matrix pocket is modeled and linear elasticity is assumed for a selected region

6.4.5 Summary

Leveraging the newly developed finite element framework, a 2x2 orthogonal weave was studied using a large non-idealized 3D textile model. Due to the complexity of 3D textiles, the geometry of the model was characterized by quantifying the variation in the cross-sectional area along the tow paths. It was shown that the cross-sectional area in tows near the top and bottom boundary of the model exhibited large amounts of variation up to 26%. On the other hand, tows near the middle of the weave exhibit relatively little variation. In reality, tows experience some amount of variation in the cross-sections, but about the same number of fibers exist for any slice, which causes the fiber volume fraction to vary along the tow path. The range of variation of cross-sectional area in the model indicated that the variation of fiber volume fraction could significantly affect the locations of stress concentrations.

The locations of severe stresses were investigated for three configurations: uniaxial tension along the global x-axis and uniaxial tension along the global y-axis, and in-plane shear. Under tension in the x-direction, the wefts experienced the most severe stress in the textile, exhibiting a high transverse tension along the local y-axis, σ'_{yy}. The severe σ'_{yy} will likely lead to initial matrix cracking in the wefts. It is important to note that the largest σ'_{yy} in the wefts developed when wefts came close to a binder and the cross-sectional shape of the weft came to a sharp corner. The sensitivity of the stress concentrations to the tow shape highlight the need to create textile models with realistic tow architectures if the progression of damage is to be accurately predicted. Additionally, the binders exhibited a

severe transverse tension and longitudinal shear, namely σ'_{zz} and σ'_{xz}, near where they traverse through the thickness of the textile. Within the matrix, a severe von Mises stress developed where the binder begins or ends traversing the thickness of the textile, near the location of the severe σ'_{xz} in the binders. Consequently, the matrix could plastically deform in this region, relieving some of the shear stress in the binders. However, unless matrix cracking in the wefts relieves the σ'_{zz} concentrations in the binders, the binders are likely to experience matrix cracking early in the progression of damage within the textile.

For the case of uniaxial tension along the global y-axis, transverse tension, σ'_{yy} in the warps and binders was the most severe type of stress in the textile. This is very similar to the previous configuration, though the severe stress is in different types of tows. Similarly, transverse matrix cracking is likely to occur in the binders and warps initially. However, the second most severe type of stress is very different from the previous configuration. When the load is transverse to the binders, a severe longitudinal shear stress develops in the wefts. Though the initial type of damage expected matches for the tension along the x- and y-directions, namely transverse matrix cracking, the next most severe types of stress were fundamentally different between the two cases due to the presence of binders. Without the binders, the severe shear stress is not expected to develop, showcasing the need for understanding what role binders play in distributing load within complex 3D textile geometries. The stress within the matrix was shown to be less critical for the case of tension along the y-axis compared to tension along the x-axis.

For the case of in-plane shear, severe shear stresses developed within the tows, as expected. The initial damage within the textile is likely to be within all three types of tows due to the longitudinal shear stress, σ'_{xz}. All three types of tows were shown to also experience severe transverse normal stresses where tows tend to pivot around each other as the textile undergoes shear deformation. However, the stress concentrations were quite small and localized. Within the gaps between warps and wefts, the binders were shown to experience a severe transverse shear stress, σ'_{yz}, and the matrix was shown to experience a high level of von Mises stress, σ_{vm}. The matrix is likely to plastically deform in this region, which may reduce the σ'_{yz} concentration in the binders.

To gain insight into the how binders redistribute the load and contribute to stress concentrations within the textile, the binders were replaced with matrix material and the stress distributions were compared to the case where the binders were included. Under tension along the x-axis, which is aligned with the warp tows, the stress concentrations within the wefts, which was the most severe component of stress, were reduced significantly. It was shown that the binders transferred the load from the warps to the wefts, creating the severe σ'_{yy} within the wefts. Additionally, the warps exhibited a significant variation of stress within the tows, which is due to the waviness of the tow. Thought the warps are nominally straight, the simulated processing of the 3D tow architecture resulted in waviness within the warp tows. Since the longitudinal modulus is much higher than the transverse modulus, the stresses within the warps were very sensitive to small changes in the fiber direction. Under in-plane shear, the binders had significantly less of an effect on the stress

distribution. The most critical stress concentrations remained in the same locations with or without the binders. However, the results showed that the waviness within the wefts significantly affected the local stresses along the tow paths under the in-plane shear. From experiments, it has been shown that binder delay the onset of delamination, so it is likely that binders have a greater effect during the later stages of the progression of damage.

Due to the large variation of the cross-sectional area of the tows of the textile model, the effect of accounting for the variation of the local fiber volume fraction was investigated. It was shown that accounting for a varying V_f noticeably affected the σ'_{yy} concentrations within the wefts, which was the most severe stress. Consequently, ignoring this feature of textile models will likely lead to an over prediction of the initial damage within the textile. However, the varying V_F had little effect on the severe stresses within the binders (5% to 10%) and no noticeable effect on stresses within the matrix.

A plasticity model was used to study the effect of plastic deformation in the matrix on the stress concentrations within the tows. Though up to 1% effective plastic strain is predicted near the locations where binders begin or end traversing the textile thickness, the σ'_{yy} concentrations within the wefts remained unaffected, due to how plastic deformation localized to small regions. However, the plastic deformation did lower the σ'_{xz} concentrations within the binders, but only by 3%.

6.5 Conclusions

Creating realistic 3D textile models and predicting the response of the textiles with high-fidelity remains a challenge for the community, but this work took a step forward. The

study of a 1x1 orthogonal weave highlighted the importance of accounting for the 3D nature of the tow paths when determining the local material properties, considering a larger analysis region to avoid artifacts due to boundary conditions, using significantly refined meshes, and the need to assess the effect of a variation of the local fiber volume fraction. Guided by these conclusions, the study of a 2x2 orthogonal weave used a more refined mesh and modeled a larger region of the textile, providing more detail than any study in the existing literature.

The locations and severity of stress concentrations within the textile for three types in-plane loads were assed, highlighting the complex stress states that develop within the tow architecture. Based on the linear elastic analyses, failure within the textile is predicted to occur primarily within the warps and wefts for the in-plane loads considered. However, the binders were shown to have a significant effect on the stress distributions. In the case of uniaxial tension along the x-axis, the binders worsened regions of critical stress within the wefts and slightly reduced the stress within the warps. Therefore, it is expected that the binders actually expedite transverse microcracking within the wefts, but it might slightly delay the onset of fiber breakage within the warps. The binders had a less significant effect for the case of in-plane shear, though it should be noted that they likely become significant as damage progresses. When the binders were replaced with matrix, the waviness of the tows were shown to greatly affect the local stress distributions, highlighting the importance of more accurate nonidealized textile models. The variation of the cross-sectional area in the tows was shown to be significant for predicting the stresses within the 3D textile models

used, but accounting for the variation of the local fiber volume fraction did not change the locations of critical stresses. If a study aims to predict the strength of a textile, the variation of the fiber volume fraction will be important to consider, though it should be noted that most important effect may be how the tow strengths vary with local fiber fraction, which was not addressed in this study. Accounting for plasticity in the matrix did affect some of the stress concentrations in the tows by up to 3%, but the most severe stresses remained unaffected.

For all configurations considered in this study, it was shown that the shape of tows can be important, the faceted surfaces of the tows in the models introduced severe stress concentrations, and adequate mesh refinement remains a challenge even with HPC resources. Further investigation will be needed with very large models to determine the converged stress state when tows are modeled as homogeneous textiles. However, due to the very small volumes of material experiencing the extreme stresses, it is unclear at this point if treating the tows as homogeneous orthotropic materials is valid, especially near the surfaces of the tows. This work has pushed the state of the art for modeling 3D textiles, but as effort is given to more accurately predict the response within 3D textiles, many new challenges are introduced. Fortunately, with growing computational resources, it is becoming possible to address these complications if the finite element framework used by researchers continues to evolve.

7 CONCLUSIONS

Using a scalable finite element framework that was developed for the multiscale investigation of fiber/matrix composites, new insights into the behavior of laminated and woven composites were discovered. For two types of laminated composites, the classical free-edge problem was revisited with the heterogeneous microstructure directly modeled. For two types of orthogonally woven textile composites, nonidealized textile models under a variety of configurations were simulated, providing insight into how load is distributed throughout the complex tow architecture, locations of critical stresses, and the effect of several modeling parameters. These studies have showcased the advantage of leveraging high-performance computing to predict the response of fiber/matrix materials and provided a step towards designing next generation composites, such as multifunctional materials, from the microscale up possible.

This work began with the development of the finite element framework, which made the studies of laminated and woven composites possible. To create the finite element models, new algorithms were developed. A method for creating random fiber matrix RVEs with more general shapes was developed, which allowed random fiber arrangements with higher fiber volume fractions to be reliably generated. The method also provides the capability to generate microscale models of complex regions within larger models, such as the cross-section of a tow. Leveraging the fiber/matrix model generation tool and third-party libraries, a tool was developed for creating laminate composite models where the heterogeneous microstructures of any 0^0 and 90^0 plies within the layup can be modeled

directly. For textile composites, a combination of VTMS and newly developed algorithms allowed the creation of nonidealized 3D textile models. A description of these preprocessing tools was provided in this manuscript. To allow the finite element analysis of the large laminated and textile models, a new finite element code was developed. A description of the details of the architecture and implementation was outside of the scope of this manuscript, but effective design paradigms, useful third-party libraries, and lessons learned for the development of a FEA code that leverages HPC were discussed. To enable the visualization of the large models, a combination of a new file format, a custom ParaView plugin for reading the data in parallel, and a Python script framework for post-processing the data were developed. Unfortunately, all of the developments could not be discussed in detail in this manuscript, but the aim of this manuscript was to document what is required to create a scalable FEA framework for the prediction of fiber/matrix composites and provide insights to guide the future development of similar frameworks.

Within multiscale analyses, the microscale can be accounted for directly by modeling the heterogeneous microstructure within a mesoscale analysis or by using reduced order models. Both strategies are used within this manuscript. It was observed that the nonidealized 3D textile models created by the preprocessing tools exhibited a significant variation of the cross-sectional area along tow paths. Within a carbon fiber composite, the area of carbon fibers remains relatively constant throughout a tow, so significant changes in the cross-sectional areas indicates that the fiber volume fraction varies significantly. A strategy for capturing the effect of varying local fiber volume fractions was proposed based

on spline fits of the mean properties from multiple ensembles of fiber arrangements, each representing a particular fiber volume fraction. The newly developed method for creating fiber/matrix RVEs allowed periodic hexagonally shaped RVEs with random fiber arrangements to be generated for relatively high fiber volume fractions. Using mean-field homogenization of 18 ensembles each with 50 different fiber arrangements under 6 different load configurations, distributions of effective properties for a range of fiber volume fractions were calculated. Using the mean effective properties of each ensemble, a constitutive law that is a function of the fiber volume fraction was developed. Enabled by the framework, the process of predicting the mesoscale properties as a function of volume fraction was completely automated given the constituent properties. The constitutive law was later used to account for the variation of the local fiber volume fraction in 3D textile models, creating a multiscale method of accounting for the underlying microstructure of the tows. This strategy worked well for this problem and is expected to be useful for many other material properties, such as thermal conductivity and moisture absorption.

The classical free-edge problem, which has been exhaustively studied at the mesoscale, was revisited with far more detail than done before by directly modeling the heterogeneous microstructure of the plies. Accounting for the microscale heterogeneity removed the singularity at the ply interface and was shown to dramatically affect the predicted interlaminar stresses near a free-edge. The heterogeneous microstructure was shown to also affect the stresses away from the free-edge. The fiber arrangement affected the local stresses, and it was shown that a more uniform fiber arrangement results in lower peak stresses. The

microscale Poisson's ratios in the matrix and fibers were shown to significantly affect how the stress fields near several fibers interact. For typical constituent properties, the interaction of fibers lead to compressive stress concentrations. The results indicate that matrix rich regions should be avoided near the ply interface to lower the tendency for delamination. Though direct simulation of the microstructure provides the highest fidelity possible within a finite element analysis, it is impractical to use the large models to study the effect of microstructural parameters. Consequently, the region needed to accurately predict the microscale stresses near the ply interface was investigated and shown to be much smaller then entire ply.

For two different orthogonally woven composites, the locations of critical stress concentrations and the distribution of stresses were investigated, providing a more detailed understanding of how load is transferred throughout the complex tow architecture. The study highlighted the importance of mesh refinement, showing the most refined meshes in the literature are not sufficient to accurately predict local stresses. For the largest models in this study, the locations of critical stress under the different loads qualitatively matched the locations of damage expected for the orthogonal tow architecture. The binders were shown to significantly affect the distribution of load when the textile experience tensile loads along the warp direction. Under in-plane shear, a smaller effect was observed, but the binders are expected to become more significant as damage progresses based on experiments. However, waviness in the tows due to the presence of binders was shown to be significant for both tension along the warps and in-plane shear, indicating the importance of an realistic

textile models. Accounting for variation in the tow volume fraction was shown to increase the stresses in critical regions. Therefore, it is expected that accounting for this detail is very important for progressive damage analyses if the tows exhibit significant variation of the cross-sections. Though the region of matrix at a high stress occurred near some stress concentrations within the tows, modeling the plasticity within the matrix did not significantly affect the stresses within the tows due to how the plastic deformation localized to small regions.

The studies presented in this manuscript have revealed some new insights into the distributions of stress within fiber/matrix composites, but they also highlighted many open questions that require future research. For the subject of free-edge induced delaminations, the layups of laminates have been tailored to mitigate the chances of delaminations in the past, but a deeper understanding of where the microscale damage initiates and how it propagates will be needed to tailor the microstructure and effectively integrate functional constituents. Compressive stress concentrations form where fibers come close to each other, so the critical stresses in the regions between plies occurs near matrix pockets. However, the fiber/matrix interface is known to be a critical location for damage to initiate in fiber/matrix composites. Future studies should explore whether delaminations initiate in the matrix pockets away from fibers or along the fiber/matrix interfaces by predicting the progression of fiber/matrix debonding and matrix damage. Future work should also explore the effect of plasticity in the matrix near the free-edge, since accounting for plasticity may significantly change the locations of stress concentrations.

For predicting the response of 3D textiles, the literature should consider more refined meshes if the locations of stress concentrations are going to be accurately predicted, highlighting the need to leverage HPC. Variations of cross-sectional area and waviness in the tow architecture was shown to significantly affect the predicted stresses within the tows, but detailed studies are needed to quantify how realistic the nonidealized models are to actual specimens. A qualitative comparison is often provided in the literature, but due to the sensitivity of the stresses to the tow architecture, it will be important to quantify how well the models compare to reality. This study investigated the stresses that develop within a 3D textile with greater detail than has been done before, but significant insight will be needed to guide the optimization of textile architectures and the effective integration of functional constituents for the next generation of advanced fiber/matrix composites.

REFERENCES

[1] A. Esser-Kahn, P. Thakre, H. Dong, J. Patrick, V. Vlasko-Vlasov, N. Sottos, J. Moore and S. White, "Three-Dimensional Microvascular Fiber-Reinforced Composites," *Advanced Materials,* vol. 23, p. 3654–3658, 2011.

[2] R. Gibson, "A review of recent research on mechanics of multifunctional composite materials and structures," *Composite Structures,* vol. 92, pp. 2793-2810, 2010.

[3] R. Vaia and J. Baur, "Adaptive composites," *Science,* vol. 319, no. 5862, pp. 420-421, 2008.

[4] D. Hartl, G. Huff, H. Pan, L. Smith, R. Bradford, G. Frank and J. Baur, "Analysis and Characterization of Structurally Embedded Vascular Antennas Using Liquid Metals," in *SPIE Sensors and Smart Structures Technologies for Civil, Mechanical, and Aerospace Systems*, Las Vegas, NV, 2016.

[5] Y. Li and N. Goulbourne, "Numerical simulations for microvascular shape memory polymer composites," *Smart Materials and Structures,* vol. 24, 2015.

[6] G. Whitesides, "The origins and the future of microfluidics," *Nature,* vol. 442, no. 27, pp. 368-373, 2006.

[7] S. Rawal, D. Barnett and D. Martin, "Thermal Management for Multifunctional Structures," *IEEE Transactions on Advanced Packaging,* vol. 22, no. 3, pp. 379-383, 1999.

[8] N. J. Pagano, "Free edge stress fields in composite laminates," *International Journal of Solids and Structures,* vol. 14, pp. 401-406, 1978.

[9] I. S. Raju, J. D. Whitcomb and J. G. Goree, "A new look at numerical analyses of free-edge stresses in composite laminates," 1980.

[10] J. D. Whitcomb, I. S. Raju and J. G. Goree, "Reliability of the finite element method for calculating free edge stresses in composite laminates," *Computers & Structures,* vol. 15, no. 1, pp. 23-37, 1982.

[11] F. Feyel and J.-L. Chaboche, "FE2 multiscale approach for modelling the elastoviscoplastic behaviour of long ®bre SiC/Ti composite materials," *Comput. Methods Appl. Mech. Engrg.,* vol. 183, pp. 309-330, 2000.

[12] A. Nosier and M. Maleki, "Free-edge stresses in general composite laminates," *International Journal of Mechanical Sciences,* vol. 50, pp. 1435-1447, 2008.

[13] S. Chung, "Effects of Interlaminar Stress Gradients on Free Edge Delamination in Composite Laminates," Drexel University, 2003.

[14] N. Rajender, V. Rajasekhar and K. Vijay, "Micromechanical Modeling of Fiber/Epoxy Unidirectional Laminates Using FEA," *The International Journal Of Engineering And Science,* vol. 4, no. 9, pp. 54-75, 2015.

[15] H. Lin, L. Brown and A. Long, "Modelling and Simulating Textile Structures using TexGen," *Advanced Materials Research,* vol. 331, pp. 44-47, 2011.

[16] I. Verpoest and S. Lomov, "Virtual textile composites software WiseTex: Integration with micro-mechanical, permeability and structural analysis," *Composites Science and Technology,* vol. 65, pp. 2563-2574, 2005.

[17] E. Potter, S. Pinho, P. Robinson, I. Iannucci and A. McMillan, "Mesh generation and geometrical modelling of 3D woven composites with variable tow cross-sections," *Computational Materials Science,* vol. 51, pp. 103-111, 2012.

[18] Y. Wang and X. Sun, "Determining the geometry of textile preforms using finite element," in *5th Annual Technical Conference for Composites*, 2000.

[19] G. Zhou, X. Sun and Y. Wang, "Multi-chain digital analysis in textile," *Journal of Composites Science and Technology,* vol. 64, pp. 239-244, 2003.

[20] L. Huang, Y. Wang, Y. Miao, D. Swension, Y. Ma and C.-F. Yen, "Dynamic Relaxation Approach with Periodic Boundary Conditions in Determining the 3-D Woven Textile Micro-geometry," *Composite Structures,* vol. 106, pp. 417-423, 2013.

[21] A. Drach, B. Drach and I. Tsukrov, "Processing of fiber architecture data for finite element modeling of 3D woven composites," *Advances in Engineering Software,* vol. 72, p. 18–27, 2014.

[22] E. Zhou, D. Mollenhauer and E. Iarve, "A Realistic 3D Textile Geometric Model," in *Proceedings for teh 17th ICCM Conference*, Edinburgh, 2009.

[23] S. D. Green, A. C. Long, B. Said and S. R. Hallett, "Numerical modelling of 3D woven preform deformations," *Composite Structures,* vol. 108, pp. 747-756, 2014.

[24] B. Wintiba, B. Sonon, K. Kamel and T. Massart, "An automated procedure for the generation and conformal discretizationof 3D woven composites RVEs," *Composite Structures,* vol. 180, pp. 955-971, 2017.

[25] F. Stig and S. Hallstrom, "Spatial modelling of 3D-woven textiles," *Composite Structures,* vol. 94, pp. 1495-1502, 2012.

[26] P. Tan and G. Steven, "Modelling for predicting the mechanical properties of textile composites—A review," *Composites Part A: Applied Science and Manufacturing,* vol. 29, no. 11, pp. 903-922, 1997.

[27] A. Prodromou, S. Lomov and I. Verpoest, "The method of cells and the mechanical properties of textile composites," *Composite Structures,* vol. 93, pp. 1290-1299, 2011.

[28] A. Hallal, R. Younes and F. Fardoun, "Review and comparative study of analytical modeling for the elastic properties of textile composites," *Composites: Part B,* vol. 50, pp. 22-31, 2013.

[29] S. Dai and P. Conningham, "Multi-scale damage modelling of 3D woven composites under uni-axial," *Composite Structures,* vol. 142, pp. 298-312, 2016.

[30] M. Saleh, G. Lubineau, P. Potluri, P. Withers and C. Soutis, "Micro-mechanics based damage mechanics for 3D orthogonal woven composites: Experiment and numerical modelling," *Composite Structures,* vol. 156, pp. 115-124, 2016.

[31] S. D. Green, M. Y. Matveev, A. C. Long, D. Ivanov and S. R. Hallett, "Mechanical modelling of 3D woven composites considering realistic unit cell geometry," *Composite Structures,* vol. 118, pp. 284-293, 2014.

[32] S. Joglekar and M. Pankow, "Modeling of 3D woven composites using the digital element approach for accurate prediction of kinking under compressive loads," *Composite Structures,* vol. 160, pp. 547-559, 2017.

[33] D. Kuksenko, B. Drach and I. Tsukrov, "Prediction of Damage Initiation and Simulation of Damage Propagation in 3D Woven Composites During Processing," in *American Society of Composites*, West Lafayette, IN, 2017.

[34] M. Saleh, "Recent advancements in mechanical characterisation of 3D woven composites," *Mechanics of Advanced Materials and Modern Processes,* vol. 3, no. 12, 2017.

[35] T. Huang, Y. Want and G. Wang, "Review of the Mechanical Properties of a 3D Woven Composite and Its Applications," *Polymer-Plastics Technology and Engineering,* vol. 57, no. 8, pp. 740-756, 2018.

[36] M. K. Ballard, W. R. McLendon and J. D. Whitcomb, "The influence of microstructure randomness on prediction of fiber properties in composites," *Journal of Composite Materials,* vol. 48, no. 29, pp. 3605-3620, 2014.

[37] W. R. McLendon, "A Multiscale Framework for the Characterization of Damage in Textile Composites Under Thermomechanical Loads," Texas A&M University, College Station, TX, 2013.

[38] J. Shewchuck, "Triangle: Engineering a 2D quality mesh generator and Delaunay triangulator," *Applied Computational Geometry,* vol. 1148, pp. 203-222, 1996.

[39] H. Si, "TetGen, a Delaunay-Based Quality Tetrahedral Mesh Generator," *ACM Transactions on Mathematical Software,* vol. 41, no. 2, 2015.

[40] J. L. Blanco-Claraco and R. Pranjal, "nanoflann," [Online]. Available: https://github.com/jlblancoc/nanoflann.

[41] M. Kraft, "libkdtree++," [Online]. Available: https://github.com/nvmd/libkdtree.

[42] G. Bernstein, "Cork Boolean Library," [Online]. Available: https://github.com/gilbo/cork.

[43] C. D. Chapman, "Prediction of moduli and strength of woven carbon-carbon composites using object-oriented finite element analysis," Texas A&M University, 1997.

[44] B. Fiedler, M. Hojo, S. Ochiai, K. Schulte and M. Ando, "Failure behavior of an epoxy matrix under different kinds of static loading," *Composites Science and Technology,* vol. 61, pp. 1615-1624, 2001.

[45] A. A. Benzerga, X. Poulain, K. A. Chowdhury and R. Talreja, "Computational Methodology for Modeling Fracture in Fiber-Reinforced Polymer Composites," *Journal of Aerospace Engineering,* vol. 22, no. 3, pp. 296-303, 2009.

[46] A. R. Melro, P. P. Camanho, F. M. Pieres and S. T. Pinho, "Micromechanical analysis of polymer composites reinforced by unidirectional fibres: Part I – Constitutive modelling," *International Journal of Solids and Structures,* vol. 50, pp. 1897-1905, 2013.

[47] A. R. Melro, "Analytical and Numerical Modelling of Damage and Fracture of Advanced Composites," University of Porto, Porto, Portugal, 2011.

[48] R. von Mises, "Mechanik der festen Körper im plastisch deformablen Zustand," *Journal of Mathematical Physics,* vol. 1, pp. 582-592, 1913.

[49] N. W. Tschoegl, "Failure surfaces in principal stress space," *Journal of Polymer Part C: Science:Polymer Symposia,* vol. 32, no. 1, pp. 239-267, 1971.

[50] E. Neta, D. Peric and D. Owen, Computational Methods for Plasticity: Theory and Applications, John Wiley & Sons Ltd., 2009.

[51] J. Huang and D. V. Griffiths, "Return Mapping Algorithms and Stress Predictors for Failure Analysis in Geomechanics," *Journal of Engineering Mechanics,* vol. 135, no. 4, pp. 276-284, 2009.

[52] M. K. Ballard and J. D. Whitcomb, "Effective use of cohesive zone-based models for the prediction of progressive damage at the fiber/matrix scale," *Journal of Composite Materials,* vol. 51, no. 5, pp. 649-669, 2017.

[53] F. P. van der Meer, "Micromechanical validation of a mesomodel for plasticity in composites," *European Journal of Mechanics A/Solids,* vol. 60, pp. 58-69, 2016.

[54] A. Wongsto and S. Li, "Micromechanical finite element analysis of unidirectional fibre-reinforced composites with fibres distributed at random over the transverse cross-section," *Composites: Part A,* vol. 36, no. 9, pp. 1246-1266, 2005.

[55] K. P. Brennan and D. E. Walrath, "Analysis of Underlying Organization in Randomized Fiber Arrays," *Journal of Composite Materials,* vol. 43, no. 21, pp. 2441-2453, 2009.

[56] M. R. Garnich, R. S. Fertig III, E. M. Anderson and S. Deng, "Micromechanics of Fatigue Damage in Unidirectional Polymer Composites," in *53nd AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference Proceedings*, Honolulu, HI, 2012.

[57] D. H. Mollenhauer, E. V. Iarve, K. Hoos, M. Swindeman and E. Zhou, "Simulation of Mode I Fracture at the Micro-Level in Polymer Matrix Composite Laminate Plies," in *53nd*

AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference Proceedings, Honolulu, HI, 2012.

[58] W. R. McLendon and J. D. Whitcomb, "Micro-scale analysis for the prediction of strength under biaxial thermomechanical load," in *Proceedings of the American Society for Composites*, 2012.

[59] A. Gherissi, F. Abbassia and A. Zghal, "A Comparative Study of Three Different Microscale Approaches for Modeling Woven Composite Material," *Journal of Material Sciences & Engineering,* vol. 16, no. 1, 2016.

[60] N. J. Pagano, G. A. Schoeppner and R. Kim, "Steady-state cracking and edge effects in thermo-mechanical transverse cracking of cross-ply laminates," *Composite Science and Technology,* vol. 58, pp. 1811-1825, 1998.

[61] R. B. Pipes and N. J. Pagano, "Interlaminar stresses in composite laminates under uniform axial extension," *Journal of Composite Materials,* vol. 4, no. 4, pp. 538-548, 1970.

[62] M. K. Ballard and J. D. Whitcomb, "Effect of heterogeneity at the fiber/matrix scale on predicted free-edge stresses for a [0/90]s laminated composite," in *21st International Conference on Composite Materials*, Xi'an, China, 2017.

[63] M. Pankow, A. Quabili, S. Whittie and C. Yen, "Tensile and Laterally Confined Compression Properties of Various 3-Dimensional (3-D) Woven Composites," Army Research Laboratory, 2014.

[64] M. Pankow, A. Quabili, S. Whittie and C. Yen, "Rail Shear and Short Beam Shear Properties of Various 3-Dimensional (3-D) Woven Composites," Army Research Laboratory, 2016.

[65] E. V. Iarve, D. H. Mollenhauer, E. G. Zhou, T. Breitzman and T. J. Whitney, "Independent mesh method-based prediction of local and volume average fields in textile composites," *Composites Part A: Applied Science and Manufacturing,* vol. 40, no. 12, pp. 1880-1890, 2009.

[66] I. Tsukrov, M. Giovinazzo, K. Vyshenska, H. Bayraktar, J. Goering and T. Gross, "Comparison of two approaches to model cure-induced microcracking in three-dimensional woven composites," in *ASME 2012 International Mechanical Engineering Congress & Exposition*, Houston, Texas, 2012.

[67] M. K. Ballard, J. S. McQuien and J. D. Whitcomb, "Analysis of three-dimensional woven composites," in *21st International Conference on Composite Materials*, Xi'an, 2017.

[68] A. Henderson, "ParaView guide, a parallel visualization applicaiton," Kitware Inc., 2007.

[69] S. Soghrati, A. R. Najafi, J. H. Lin, K. M. Hughes, S. R. White, N. R. Sottos and P. H. Geubelle, "Computational analysis of actively-cooled 3D woven microvascular composites using a stabilized interface-enriched generalized finite element method," *International Journal of Heat and Mass Transfer,* vol. 65, pp. 153-164, 2013.

[70] I. Tsukrov, B. Drach, H. Bayraktar and J. Goering, "Modeling of Cure-Induced Residual Stresses in 3D Woven Composites of Different Reinforcement Architectures," *Key Engineering Materials,* Vols. 577-578, pp. 253-256, 2014.

[71] M. K. Ballard and J. D. Whitcomb, "Prediction of Tow Architecture and Stress Distributions for a 3D Woven Composite," in *American Society of Composites*, West Lafayette, Indiana, 2017.

[72] E. Jones, T. Oliphant, P. Peterson and Others, *SciPy: Open source scientific tools for Python,* 2001.

[73] S. Eskandari, F. M. Pires, P. P. Camanho and A. T. Marques, "Intralaminar damage in polymer composites in the presence of finite fiber rotation: Part II – Numerical analysis and validation," *Composite Structures,* vol. 151, pp. 127-141, 2016.

[74] S. Eskandari, F. M. Pires, P. P. Camanho and A. T. Marques, "Intralaminar damage in polymer composites in the presence of finite fiber rotation: Part I – Constitutive model," *Composite Structures,* vol. 151, pp. 114-128, 2016.

[75] P. Camanho, A. Arteiro, A. Melro, G. Catalanotti and M. Vogler, "Three-dimensional invariant-based failure criteria for fibre-reinforced composites," *International Journal of Solids and Structures,* vol. 55, pp. 92-107, 2015.

[76] A. Drach, B. Drach and I. Tsukrov, "Processing of fiber architecture data for finite element modeling of 3D woven composites," *Advances in Engineering Software,* vol. 72, pp. 18-27, 2014.