CHALMERS
UNIVERSITY OF TECHNOLOGY

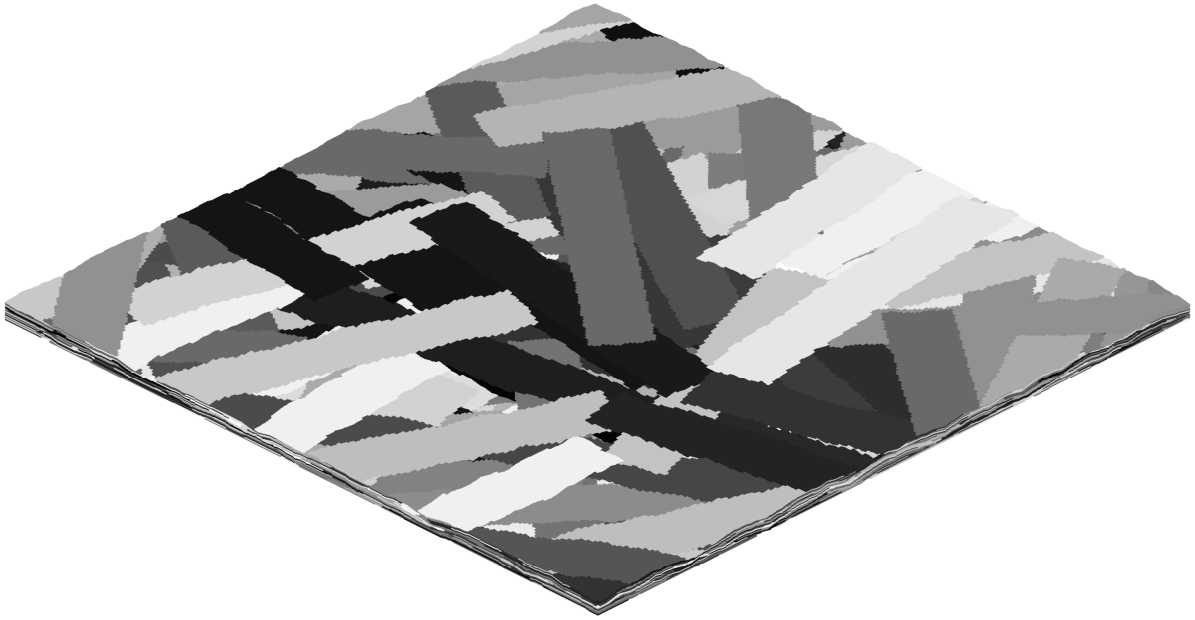

Development of modelling and homogenisation procedures for stochastic tape-based discontinuous composites

Master's thesis in Applied Mechanics

Olle Haglund Nilsson
Jacob Sjöberg

DEPARTMENT OF INDUSTRIAL AND MATERIALS SCIENCE
Chalmers University of Technology
Gothenburg, Sweden 2022

Master's thesis in applied mechanics
2022

Development of modelling and homogenisation procedures for stochastic tape-based discontinuous composites

OLLE HAGLUND NILSSON
JACOB SJÖBERG

CHALMERS
UNIVERSITY OF TECHNOLOGY

Department of Industrial and Materials Science
Division of Material and Computational Mechanics
Chalmers University of Technology
Gothenburg, Sweden 2022

Development of modelling and homogenisation procedures for stochastic tape-based discontinuous composites

Supervisor: Prof. Martin Fagerström, Division of Material and Computational Mechanics.
Examiner: Prof. Martin Fagerström, Division of Material and Computational Mechanics.

Master's Thesis 2022:07
Department of Industrial and Materials Science
Division of of Material and Computational Mechanics
Chalmers University of Technology
SE-412 96 Gothenburg
Telephone +46 31 772 1000

Cover: Generated geometrical model of an stochastic tape-based discontinuous composite.

Typeset in LaTeX
Printed by Chalmers Reproservice
Gothenburg, Sweden 2022

Development of modelling and homogenisation procedures for stochastic tape-based discontinuous composites

OLLE HAGLUND NILSSON
JACOB SJÖBERG
Department of Industrial and Materials Science
Division of Material and Computational Mechanics
Chalmers University of Technology

Abstract

Conventional laminated composites have, due to their high stiffness- and strength-to-weight ratios, shown to be a viable alternative to metals in applications where lightweight structures are of essence. However, the current reach of conventional laminated composites is limited to certain industries and high-end products as a result of their high manufacturing cost. As a way to reduce cost while maintaining high performance, stochastic tape-based discontinuous composites (STBDCs) have emerged, offering a middle ground between the easy manufacturing of short fibre composites and high performance of continuous fibres composites. The irregular mesostructure and high interest from industries have increased the demand for efficient and predictive models (analytical and numerical) describing the mechanical response of this class of material. As a response to the increased interest, this project aimed to develop a method to predict the elastic properties of STBDCs, based on the mesostructural configuration and constituent properties (impregnated tapes and pure matrix).

To replicate the complex material structure, a modular MATLAB code was developed as the core of the project. The MATLAB code takes the tape and plate dimensions, material properties, voxel resolution, etc, as inputs, then builds the material layer-by-layer by placing tapes at random positions and angles until the desired plate thickness has been reached. The developed algorithm allows tapes to form around each other (drape) to replicate the compression part of the real manufacturing process. The outputs of the code are, among other visualisations, a 3D-image of the material structure and Abaqus input files for the coming analysis. Based on the generated material structures, smaller statistical volume elements (SVEs) were extracted and used for the analysis. Nearly 400 SVEs of different dimensions from different samples were extracted and analysed to have a large enough sample size to account for the variability of the material. Computational homogenisation was used to determine the volume averaged in-plane elastic parameters. The homogenisation process was carried out by applying periodic boundary conditions (PBCs) to the SVEs using the Abaqus plug-in EasyPBC. The computational homogenisation was initially focused on a full 3D-model but as a subsequent step, the full 3D-model was reduced to a 2D-model by an intermediate analytical homogenisation process using classical laminate theory. A full 3D- and reduced 2D-model were generated for each SVE sample to allow a comparative analysis. Finally, experimental data of manufactured STBDCs plates was used as a comparison to verify the results of the numerical model.

The study found that the voxel-based 3D-model and the developed methodology can be used to accurately evaluate the elastic properties of STBDCs. More specifically, the generated material samples and methodology used provided reliable results for all studied SVE dimensions, converging to elastic properties within one standard deviation compared to experimental tests. Furthermore, the reduced 2D-model showed a similar accuracy compared to the 3D-model while also requiring significantly less computational power, indicating that this is a reasonable simplification to make.

It should be noted that the developed model is not a perfect representation of reality, simplifications had to be made to stay within the limitations of the project. As a consequence, the fibre volume fraction (FVF) of experimental data could not be reached, making the model slightly under-predict the elastic properties. With further improvements to raise the FVF, the data indicates that the model would produce more accurate results compared to experimental tests, thus being a reliable source material for future industrial use.

Keywords: Stochastic Tape Based Discontinuous Composite, Periodic Boundary Condition, Statistical Volume Element, Voxel, Chopped Fibre Composite, Tow, Homogenisation,

Preface

The master's thesis was conducted as a 30 credits project at Chalmers University of Technology during spring of 2022. Both authors, Olle Haglund Nilsson and Jacob Sjöberg were students at the master's programme Applied Mechanics with previous Bachelor of science degrees in Mechanical Engineering.

Acknowledgements

This thesis work was conducted at the Division of Material and Computational Mechanics at Chalmers University of Technology. We would first and foremost like to thank our supervisor and examiner Professor Martin Fagerström, without you this project would have been only a fraction of what it became. Your expertise, experience and guidance has been an incredible help throughout the project. Furthermore, we want to express our gratitude to everyone at the division. Thank you for opening your arms to us, making us feel like a part of the division. Thanks to all of you, the time during our thesis has not only been a great learning experience, but also a genuinely enjoyable time that we will cherish moving forward.

Table of Contents

Nomenclature

Standard variables

ν	Poisson's ratio
σ	Stress
ε	Strain
φ	In-plane tape angle
D	Plate volume domain
E	Young's modulus
G	Shear modulus
l	Length
$N^t_{\text{vox},h}$	Through thickness number of voxels (tape)
$N^t_{\text{vox},l}$	Number of length voxels (tape)
$N^t_{\text{vox},w}$	Number of width voxels (tape)
R_{vox}	Voxel aspect ratio
t	Thickness
u	Displacement
V	Volume
V_f	Fibre volume fraction
w	Width
X_c	Tape centre position (x-direction)
Y_c	Tape centre position (y-direction)

Subscripts

vox	Voxel
c	Centre
L	Longitudinal
T	Transverse

Superscripts

CA	Cumulative average
e	Element
p	Plate
t	Tape

Abbreviations

FEA	Finite Element Analysis

FVF Fibre Volume Fraction

PBC Periodic Boundary Condition

RP Reference Point

RSA Random Sequential Adsorption

RVE Representative Volume Element

STBDC Stochastic Tape Based Discontinuous Composite

SVE Statistical Volume Element (aperiodic RVE)

1 Introduction

Conventional laminated composites are an advantageous alternative to metals in numerous applications, especially because compared to metals these materials offer better stiffness- and strength-to-weight ratios, corrosion- and fatigue resistance [1]. Given their properties and material structure, continuous composites are particularly suitable for long continuous structures from a manufacturability perspective. Currently these materials are more commonly used in the aircraft industry than in the automotive industry, although a similar carbon fibre reinforced polymer (CFRP) usage in the automotive industry as the aircraft industry could lead to a 54 Mt reduction of CO2 emissions per year due to weight reductions [1, 2]. For reference, a 10% reduction in weight of a personal vehicle leads to a 6-8% reduction in fuel consumption, or for electric vehicles, an increased range [3]. The use of carbon fibre composites in the automotive sector is low and limited to luxury vehicles due to their high manufacturing cost [2]. This project aims to investigate a material that due to its ease of manufacturing and mouldability introduce a viable alternative to metals that can extend the use of CFRP across all industries; stochastic tape-based discontinuous composites (STBDCs) [1].

STBDCs is a class of high performance lightweight discontinuous composites composed of randomly distributed, polymer matrix impregnated, carbon fibre tapes, schematically shown in Figure 1(b) [2, 4]. The term carbon fibre tapes is used to describe flat strands of unidirectional carbon fibres, pre-impregnated with either a thermoplastic or thermoset polymer matrix [5]. The carbon fibre tapes may be sourced as residual scrap from the production of continuous fibre composites which leads to a significant reduction of waste [1] or directly from unidirectional carbon fibre sheets. Due to the discontinuity in the layout of the tapes, STBDCs can be compression moulded to complex geometries, a process where the tapes are placed and press-formed in a mould and then cured under pressure [4, 6]. Compression moulding gives STBDCs great potential for high volume production due to the short cycle time and low cost compared to conventional composites [2, 7].

STBDCs offer a middle ground between the easy manufacturing of short fibre composites and high performance of continuous fibre composites [8]. The key benefits of short fibre composites are that they are easy and fast to produce by either injection- or compression moulding and with that carry a low cost per unit [9]. The downside, however, is that they typically have only a fraction of the stiffness and strength of continuous fibre composites [9]. Continuous fibre composites are stiffer and stronger than short fibre composites but also more costly and difficult to manufacture into as complex shapes [1]. Discontinuous tape-based composites can due to their drapability and discontinuity be used to manufacture complex shapes while maintaining a low cost and high performance [1, 2, 4].

(a) Short fibre composites.

(b) Layup of STBDCs, [4].

(c) Layup of conventional composite laminates.

Figure 1: Schematic layups of (a) short fibre composite, (b) STBDCs, and (c) conventional composite laminates, fibre orientation visualised with black stripes.

Numerous companies globally are actively exploring the potential of implementing STBDCs into the design of new products, a selection of which are presented in Figure 2. Complex metal structures can easily be converted to use STBDC material with little design changes required, whereas large design modifications are typically needed to convert a complex metal part to use continuous laminate composites [10]. As a practical example, the suspension control arm for the Lamborghini concept car Sesto Elemento, shown in Figure 2(a), is a component under high loads. By replacing the aluminium with STBDC, the component became on average 27 % lighter and had a competitive cost compared to its aluminium counterpart [11]. Converting the suspension control arm from aluminium to STBDC only required one major modification; the addition of a third cross member to increase lateral stiffness [11]. A second example is the window frames of the Boeing 787 Dreamliner, shown in Figure 2(b). They used to be made of aluminium but when converted to STBDC, the weight was reduced by nearly 50 % and they gained a superior damage tolerance [4, 12].

(a) Lamborghini Sesto Elemento concept front upper and lower suspension control arms.

(b) Boeing 787 window frame.

Figure 2: Current usage of STBDC materials, [13].

When working with STBDCs it is important to consider that the stochastic nature of the material creates a highly heterogeneous structure containing significant variability of local properties and mesostructure [2, 4]. This results in a material that is locally anisotropic but globally in-plane isotropic [14, 15]. Locally anisotropic refers in this case to that the material properties are direction-dependant on a local scale, even though the average macroscopic response is in-plane isotropic [14]. A drawback of the randomised position and orientation of the tapes is that internal weak spots, e.g. in the form of matrix filled spaces, are distributed throughout the material. These matrix filled spaces occur at two places: i) when tapes lay next to each other without contact, as shown in Figure 1(b), and ii) when one tape drape over another as shown in Figures 7 and 8. Stress concentrations may arise at these matrix rich regions which could cause critical or sub-critical damage to the structure [1]. Due to the irregularities in the meso-scale of STBDCs, no two finished components will ever be identical. This in combination with moulding process sensitivity makes it difficult to characterise general mechanical properties [14]. Thus, increasing the need for predictive analytical as well as numerical models.

1.1 Aim

The main aim of this project has been to develop a method to predict elastic properties of STBDCs. This has then been separated in two main goals: i) to develop a method to generate a voxel-based geometrical representations of material samples, and ii) to evaluate homogenisation strategies for predicting the volume averaged macroscopic properties.

More specifically, for the first goal, the aim has been to develop a modular and multi-functional MATLAB code that based on a set of input parameters outputs a voxel-based numerical model of a flat STBDC plate, along with visualisations of geometrical properties and 3D-models of said plate. For the second goal, the aim has been to investigate, implement, and compare different methods of homogenisation on models of different scale and dimension, namely full 3D models and reduced 2D models of different sizes.

1.2 Method

The method used to reach the goals of the project can be described in three steps. The first step was a literature study to learn more about the real material and capture the state of the art of modelling and analysing STBDCs. This step gave an important understanding of the micro (fibre), meso (tape), and macro (component) scale properties and effects of the STBDC material, as well as procedures to analyse it numerically. The second step was to use the knowledge gained in the literature study and implement it to create a representative voxel-based numerical model of the material. The final step was to transfer the model to the finite element solver Abaqus, evaluate the in-plane material properties, and compare them to experimental data. The last step was done for both a full-scale 3D-model and a reduced 2D-model. The accuracy of the homogenised material properties was then verified by comparing results to experimental data. The method of homogenisation was validated by comparing results of using homogenised material properties to heterogeneous material properties in a practical example.

1.3 Limitations

This project has been conducted as a 30 credit master's thesis, which amounts to roughly 800 hours divided equally between the two authors. Time was not the only limiting factor but also available computational power which limited the size and resolution of the numerical models.

The first limitation is that this project has been conducted purely digitally. No prototypes were built, and no experimental tests were made. For experimental reference data, the project relied solely on published literature.

Regarding the numerical model of the material structure. It was set to be a three-dimensional voxel-based structure with cuboid voxels, no tetrahedral or prismatic representations have been used. The macroscopical structure was set to represent a flat plate, to enable the comparison of results to experimental test data. The tapes were assumed not to deform in the in-plane direction during the compression moulding phase but only in the out-of-plane direction. They did therefore not get wider nor longer than the pre-determined size. From a top view perspective, the tapes had a constant shape independent of draping or other effects of out-of-plane compression. As a result of this, the cross-sectional area of the tapes remained constant.

For the finite element modelling, a few limitations and assumptions were put in place to meet the time deadline. Firstly, the tapes and the matrix were assumed to be perfectly bonded. This meant that adjacent elements with different material properties shared nodes without any additional interface conditions like cohesive contacts. Secondly, only periodic boundary conditions were applied to the model for the homogenisation procedure as they were assumed to give the best response. There was no comparison of periodic boundary conditions to Dirichlet nor Neumann boundary conditions. Thirdly, thermal effects and the resulting residual stresses were not accounted for as they were expected to not affect the prediction of the elastic properties of the material. Fourthly, the model was defined as fully linear elastic, thus without damage or fracture modelling included. Finally, related to the numerical material structure, only cuboid brick elements were used in the 3D finite element analysis, and quadrilateral shell elements for the 2D analysis. This means that the possible implementation of prismatic, tetrahedral, or triangular elements was not explored in this project.

2 Previous work

The following subsections highlights previous work in the field of stochastic tape-based discontinuous composites. The first part focuses on experimental studies, summarising results and briefly explaining the method used in each study. The results gathered from these studies will later be used to verify the accuracy of the FE models. Then, a number of alternative methods for creating and analysing numerical models with different application capabilities will be summarised to provide an overview of the state of the art.

2.1 Experimental testing

In this subsection, three different studies that include experimental testing and manufacturing of STBDCs will be summarised. Test results and physical dimensions of the plates produced in the different studies are presented in Table 1. Mechanical properties of the tapes used in each study are presented in Table 2.

The first study was conducted by Wan and Takahashi in 2016 [14] and was aimed to investigate how the aspect ratio of the tapes affects properties of the final material. Carbon fibre reinforced thermoplastic tapes with a constant width $w_t = 5$ mm and thickness $t_t = 50\ \mu$m were used in the study. The tape length (l_t) for the different plates was varied between 12 mm and 30 mm in 6 mm increments. The tape randomisation method used in the study was essentially to distribute tapes in water and use a strainer to pick them up, thus creating thin sheets. The sheets were then stacked to create STBDC plates. The tensile properties were tested following the JIS K 7073 standard but with an increased specimen width to account for the tape length. The plate fibre volume fraction was evaluated using the ash test where the initial mass and volume of a sample was measured, the thermoplastic matrix was then burnt off in a furnace at a temperature of 500 °C for three hours. Then, with only carbon fibre strands remaining, the mass was re-measured and the fibre volume fraction calculated.

The second study was conducted by Li *et al.* in 2017 [4], where they manufactured and tested STBDCs consisting of two different tape thicknesses $t_t = 164\ \mu$m and $t_t = 285\ \mu$m. MATLAB was used to generate a blueprint of the stochastic tape layout. The tapes were then placed by hand according to the blueprint to form the plates. The systematic tape placement makes the position and angle of every single tape known which allows for geometrical cause and effect studies. For example, how a low local fibre content or tape alignment affect the failure mechanisms. Tensile tests were conducted according to the ASTM-D3039 standard but with increased coupon width to account for the tape length. The fibre volume fractions of the plates were evaluated by digital microstructure analysis. They further noted a presence of voids in the thick tape laminate.

A third study was conducted as a master's thesis by Johansen in 2019 [15]. In this study, the randomised layup was achieved by cutting tapes 1.1 m above a substrate and letting them fall freely through a square channel onto the substrate. The substrate was also rotated every 30 seconds ($\approx$ 300 tapes) to avoid any tendency in the landing position. The weight, and later on volume fractions, of the samples were calculated by matrix digestion using nitric acid. The tensile tests specimens were taken from different parts of the plates and cut at different angles which show some variation in the stiffness and strength results.

Table 1: Data for STBDC plates manufactured and analysed in [14, 4, 15]. † - values read from plot.

Study	Tape size [mm]	t_t [μm]	Plate size [mm]	E [GPa]	$\sigma_{p,\text{fail}}$ [MPa]	$V_{f,p}$ [%]
Wan [14]	12x5	50	-	42.5 ± 2.5†	440 ± 40†	52.2
Wan [14]	18x5	50	-	49 ± 4†	505 ± 15†	55.1
Wan [14]	24x5	50	-	46 ± 4†	480 ± 60†	52.8
Wan [14]	30x5	50	-	47 ± 2†	520 ± 20†	53.1
Li [4]	50x8	164	290x290x2.444	43.41 ± 4.14	290.1 ± 34.9	57.85 ± 3.84
Li [4]	50x8	285	290x290x2.106	39.66 ± 5.06	186 ± 34.8	53.5 ± 5.39
Johansen [15]	40x20	22.5	275x275x0.5	-	-	43.47-49.68

Table 2: Material data for tapes used in [14, 4, 15].

Study	Material	E_L^t [GPa]	E_T^t [GPa]	σ_L^t [MPa]	σ_T^t [MPa]	V_f^t [%]
Wan [14]	TR50s - PA6	240	-	4900	-	-
Li [4]	HexPly-M77	140	9	2400	73	57
Johansen [15]	HS40 - BI018	425	-	4610	-	43.17

It should be noted that in both [4] and [15], the fibre volume fraction of the plate (V_f^p) is higher than the original fibre volume fraction of the tape (V_f^t). Pre-impregnated tapes were used in both these studies and no extra resin was added to the mould. This means that some resin must have bled through the mould during the curing process. Considering that a part of the resin originally bound to the tapes has bled through the mould, and that a part of the resin has migrated to fill resin pockets, it is very likely that the fibre content of individual tapes in the final plate is higher than when originally placed.

2.2 Digital modelling and finite element analysis

Digital models have been created and analysed in terms of elastic properties, failure strength, and failure mechanisms in previous projects. These projects have had different goals and therefore also different methods of implementation. The following paragraphs briefly explain the goals and methodology in five different projects and report the relevant findings.

Li *et al.* [2] wanted to develop a method to analyse large structures while still capturing the local variability of the STBDC material. The proposed methodology utilises an equivalent laminate based analytical model for prediction of stiffness and strength properties of STBDCs. By varying input parameters, a varying set of material properties are computed. These analytically predicted properties are then assigned to different areas (sets of elements) of the larger structure in order to replicate the stochastic variability of the material. The authors used this method to analyse an automotive monocoque and found that the point of initial failure shifts depending on where weak spots are located. This monocoque example shows that the goal of analysing large structures and capturing local effects of the STBDC material was reached.

Selezneva *et al.* [16] developed a 2D model to evaluate in-plane material properties of STBDCs. The authors used a fine grid with square 2D elements to represent a plate or specimen and randomly placed tapes onto the grid. Multiple tapes were allowed to occupy the same element and classical laminate theory was used to evaluate the effective in-plane material properties in each element. The thickness of the plate was pre-defined and constant. In elements where the number of tapes did not amount to the exact plate thickness, the tape thickness was adjusted to compensate, thereby ignoring the presence of resin pockets. The model was used to create and compare specimens with different tape lengths. Abaqus was used for the finite element analysis with pure displacement boundary conditions. The developed model over-predicted the stiffness with roughly 25% for short tape (12 and 25 mm) STBDCs and roughly 6% for long tape (50 mm) STBDCs compared to experimental results. One reason for this over-prediction of stiffness may be the neglection of resin pockets. Another reason why the model over-predicted the modulus for short tape STBDCs was that these are more prone to have tapes oriented out-of-plane, which this model neglected. Any out-of-plane tape variation results in a lower in-plane modulus which the 2D model

was not capable of capturing. Capturing this out-of-plane effect would require a more complex 3D model.

Harper *et al.* [7] developed an architectural model in which individual fibre bundles were modelled as spring systems. The spring system representation of tapes enabled a more advanced tape interaction model. This model could accurately replicate the compression moulding process and the reaction of individual fibre bundles as they contact other fibre bundles. The fibre bundles were therefore allowed to both bend out-of-plane and twist around their longitudinal axis, with bending and shear stiffness represented by different spring stiffnesses. The authors achieved a reasonable fibre volume fraction (>50%) under the assumption that individual fibre bundles had a fibre volume fraction of 60%. They used the embedded element technique in Abaqus where the fibre bundles were discretised using shell elements (STRI65) embedded in solid brick elements (C3D8) representing the matrix. Pure displacement boundary conditions were used in the analysis. The tensile strength and stiffness predictions were both within 5% of experimental results for STBDC fibre volume fractions in the range 40-45%. The difference in stiffness compared to experimental results increased to 30% for fibre volume fractions in the 50-55% range. Another downside to this method was the long model generation time. It took nearly four hours to generate a square plate RVE with side length slightly larger than one tape length and the generation time increased exponentially for larger models. A generation time of 24 hours was reported for a large sized square plate RVE with side length 3.9 times the tape length.

Shah *et al.* [17] wanted to develop a more computationally efficient way of predicting strength and stiffness of STBDC materials. A lot of emphasis was put on developing both a fast model generation procedure and an efficient solution procedure compared to other methods. The tapes were given a prescribed curved shape using spline paths to replicate the out-of-plane distortion, and the tapes were sequentially added to the domain without allowing interactions. This resulted in a tape volume fraction of 61.5% which equates to a fibre volume fraction of roughly 37%. The tapes were then discretised using shell elements embedded in a solid element host mesh representing the matrix material, similar to Harper *et al.* [7]. Also in this case, pure displacement boundary conditions were used in the finite element analysis. The stiffness was underestimated by 12-29% depending on tape length and tape thickness compared to experimental data.

Ryatt and Ramulu [1] set out to develop a model to predict elastic and tensile failure properties of STBDC materials. They added voxelised tapes sequentially to a three-dimensional domain and created digital plates of different thicknesses and widths. I this case, the out-of-plane distortion of tapes was not pre-defined but depended solely on where the tapes were placed. From these digital plates, test coupons of fixed length but varying width were taken out and analysed in Abaqus using reduced integration continuum shell elements (SC8R). They used a twice as coarse mesh for the analysis than the created voxel structure to reduce computational cost and mapped material properties from multiple voxels to a single element. They used pure displacement boundary conditions and found that the computed stiffness was slightly higher (4.2%) than the experimental reference value. The authors found no clear trend in modulus as a function of plate thickness nor width. The evaluated strength of the material was found to be dependent on both plate thickness and width but in agreement with experimental results.

3 Theory

The following section initially covers STBDCs in further detail with regards to manufacturing, microstructure and material properties. Secondly, a thorough introduction on how to geometrically model STBDCs and conduct a FEA is covered in detail in order to ensure a full theoretical background of how the project was conducted, more importantly, why specific choices were made. Finally, a brief overview of the classical laminate theory is covered, introducing key concepts and formulas used in this project.

3.1 Stochastic tape-based discontinuous composites

When it comes to carbon fibre composites, the STBDC class of materials is unique. The part manufacturing process can be described as a mix of the production of short fibre composite parts and continuous fibre composite parts. The resulting microstructure can be structured in layers with subtle hints of the stochastic nature, almost like a continuous fibre composite, but it can also be highly stochastic, mainly depending on the dimensions of the tapes. The effective elastic material properties of the material are typically in-plane isotropic, as long as biases in the randomisation process can be avoided. The stiffness and strength of STBDCs depend not only on the material properties of the constituents but also the mesostructural composition.

3.1.1 Manufacturing process

The manufacturing process of STBDCs is visualised in Figure 3, where pre-impregnated carbon fibre tapes are compression moulded to form a geometrically complex part. The process starts with generating the stochastic material structure. Some methods of generating a random layup of tapes were discussed in Subsection 2.1, including tape dropping [15] and a method similar to a wet-type paper making process [14]. Another method is to place small batches of tapes in the mould and shuffle it back and forth between each batch to evenly distribute the tapes [8]. Further, the tapes can either be placed directly in the mould or made into sheets before being placed in the mould [4]. The moulding material can also be bought pre-made in the form of sheet moulding compounds (SMCs) [6].

The second part of the process is the compression moulding, which itself consists of four steps [6]. First, the moulding material is placed in a hot mould, either directly or as a sheet. Second, the mould is closed, the pressure increases, and the material is allowed to flow. How much the material flows depends on the material. However, for STBDCs, it has been observed that the material flow is low, at least for flat plates [4]. Third, the material is cured while the pressure and the temperature is maintained. The curing time depends on the resin, part thickness and temperature [6]. Finally, the mould is opened, the part is extracted and left to cool down. The compression moulding process can be completed in down to two minutes if a fast curing matrix material is used [2]. The moulds may include multiple features such as stiffening ribs, holes, and inserts that reduce post-production machining [6].

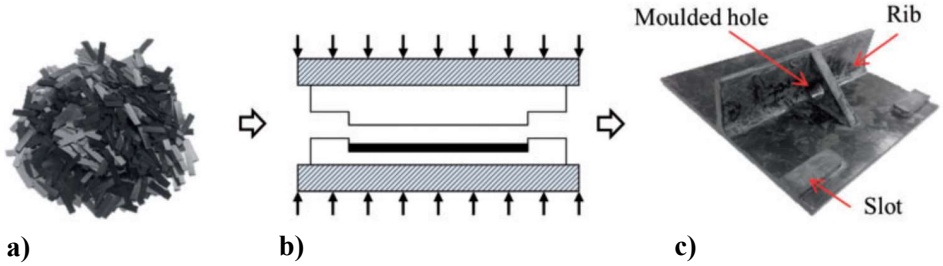

Figure 3: The production process of STBDCs, a) pre-impregnated carbon fibre tapes, b) compression moulding phase, and c) complex STBDC part with multiple features [16].

3.1.2 Microstructure

Depending on the tapes and the manufacturing process, the microstructure may look very different for different STBDCs [8]. Two examples are shown in Figure 4 where short tapes (12x6 mm) and long tapes (50x6 mm) are compared. The microstructure of the plate using short tapes is more irregular, containing more resin pockets and large out-of-plane oriented tapes. The long tape microstructure, on the other hand, is more resembling a continuous fibre laminate in that the tapes are more structured in layers with less out-of-plane variations.

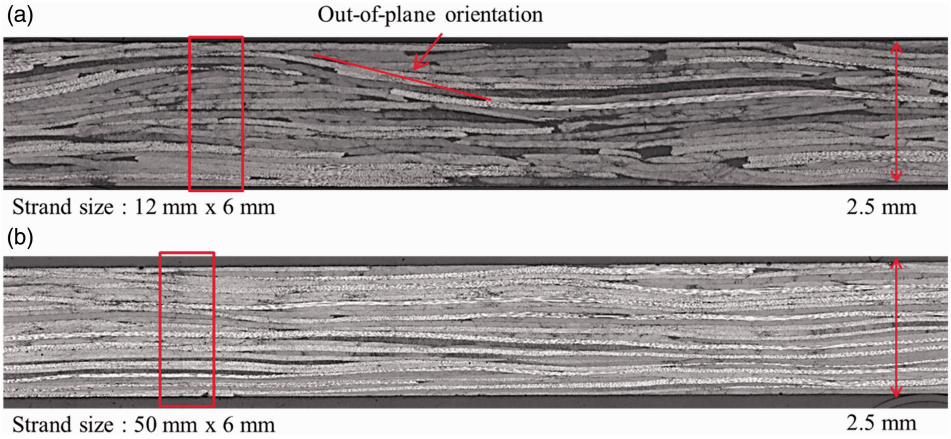

Figure 4: Microstructure images for STBDC materials with (a) 12x6x0.14 mm tapes, and (b) 50x6x0.14 mm tapes, [16].

STBDC plates can have a varying plate thickness despite a fixed number of stacked tapes in a region, an example of this is shown in Figure 5(a). The varying plate thickness is a result of the amount of resin between the layers and how much the tapes are deformed, which stems from the compression phase [15]. Resin pockets are a common characteristic feature in STBDC materials, appearing either where one tape drape over the end of another tape or where two tapes lie next to each other without contact. Figure 5(b) displays a resin pocket in the microstructure.

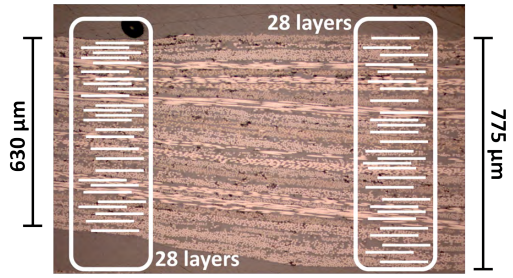

(a) STBDC microstructure image showing thickness difference in a plate despite a fixed number of layers. White lines indicating different layers.

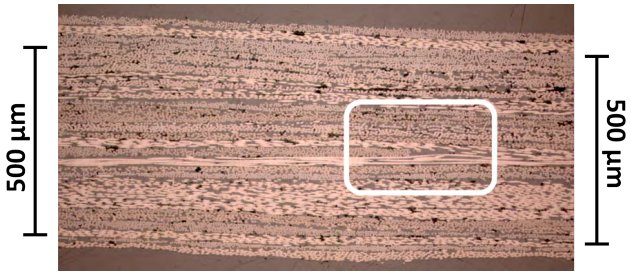

(b) STBDC microstructure image showing a resin pocket.

Figure 5: Microstructure images showing (a) thickness variation, and (b) a resin pocket [15].

Figure 6 shows a comparison of a thin tape ($t_t = 0.164$ mm) and a thick tape ($t_t = 0.285$ mm) STBDC microstructure where multiple characteristic features can be observed. Resin pockets appear in both the thin and the thick tape material, visible in regions 2 and 3, respectively. Further, the varying tape thickness, showcased in Figure 5(a), is also visible in region 1 in Figure 6(a). The final notable characteristic feature is the out-of-plane tape rotation (draping), shown in region 2 in Figure 6(a), where one tape has been pressed over the edge of another tape. Although the microstructure images in Figure 6 only show a small section of the materials, one major difference can be noted between the thin and thick tape STBDCs. The thin tape material shows, in general, a more uniform distribution of tapes with distinct layers and little out-of-plane variations compared to the thick tape material. Region 4 in Figure 6(b) shows this well, where there is no clear boundary between tapes and the layers nor out-of-plane misalignment can be distinguished.

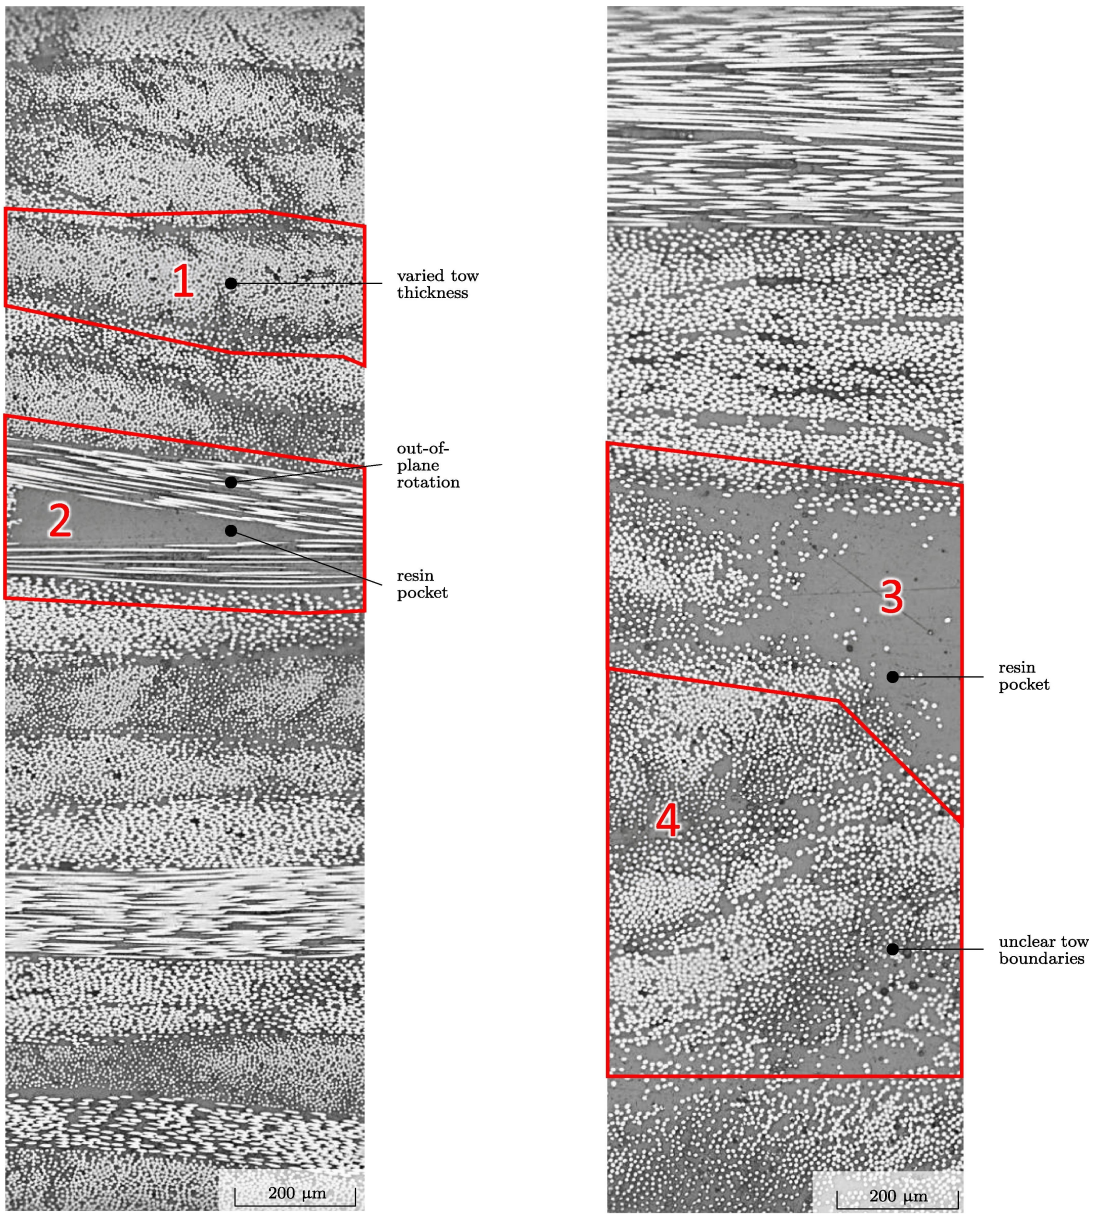

(a) Thin (164 μm) tape STBDCs.

(b) Thick (285 μm) tape STBDCs.

Figure 6: Microstructure images of (a) thin, and (b) thick tape STBDCs [4]. Note that what in this report is called tape is in this image denoted tow.

3.1.3 Material properties

The macroscale material properties of STBDCs are a result of the micro- and mesostructure, which in turn is a result of the tapes and the manufacturing process. Besides the material properties of the tapes and the polymer matrix, there are multiple factors affecting the macroscopical material properties. These factors are mainly the mesostructure of the material and the dimensions of the tapes, but also e.g. the manufacturing process. One effect of the randomised layup of tapes is that the material is locally anisotropic [2, 4]. Globally, on the other hand, the material is in-plane isotropic as long as there is no bias in the tape distribution [14, 15].

The local variability in the material creates both strong and weak spots where the fibre volume fraction is either high or low, where fibres have a preferred direction, and in the presence of resin pockets [4, 8]. Li *et al.* [2] analysed the impact of weak spots in their finite element model for an automotive monocoque, discussed in Subsection 2.2. The authors found that the point of initial failure shift depending on where the local weak spots are located, even in the presence of notable geometrical stress concentrations. A consequence of this is that the region at which the material will start to fail also varies, due to the failure of STBDCs typically occurring at material weak spots. A positive aspect of the inherent weak spots is that they aid in making the material notch insensitive [4], which allows holes and inserts to be used in the final product without compromising its structural integrity.

Another factor that has been shown to affect the properties of the material are the dimensions of the tapes. The mesostructural dependence on tape length and aspect ratio (l^t/t^t) is clearly visible in Figure 4. This effect is not only visual, both the stiffness and the strength of the material increase with increasing aspect ratio of the tapes up to a certain point where the effect plateaus [14]. Li *et al.* [4] experimentally tested STBDCs made of thin tapes with tape thickness $t_t = 0.164$ mm and thick tapes ($t_t = 0.285$ mm) and found that the thick tape plates failed at a lower stress level than the thin tape plates. The earlier failure of the thick tape plates was determined to be partly due to earlier debonding and partly due to a larger microstructural variability than in the thin tape plates. Both the lower aspect ratio and the thickness of the thicker tapes makes them more prone to debonding, i.e. matrix failure between tapes. Also the larger microstructural variability, shown in Figure 6, of the thick tape plates play a role, with a larger presence of out-of-plane oriented tapes and resin pockets. Li *et al.* [4] found that tape pull-out is more dominant for thick tape STBDCs whereas tape fracture is dominant for thin tapes STBDCs. Tape pull-out and tape (fibre) fracture are the two most common failure modes in STBDC materials [4, 8, 15].

As there are multiple factors affecting the STBDC material properties, it is very difficult to state a typical value for e.g. modulus and stiffness. In examples found in literature, the modulus ranges from 20 GPa to 60 GPa, and the strength ranges from 100 MPa to 500 MPa, depending on the constituents and manufacturing process [4].

3.2 Geometrical modelling

The accuracy of the finite element analysis results, in terms of predicted effective material properties, depend greatly on the accuracy of the geometrical model [17]. Therefore, it is important to capture the major mesostructural effects (random tape distribution, tape draping, fibre volume fraction, etc) accurately. On the other end of the spectrum is computational cost, which increases with the level of detail and size of the model. Therefore, simplifications must be made where it is possible and where the impact on the predicted properties is minimal. The following paragraphs describe some geometrical effects needed to be accounted for and methods of doing so in a numerical model.

To create a numerical STBDC model from the ground up, without a physical reference material, there are two general methods to use; the random sequential adsorption technique (RSA) and the Monte-Carlo procedures [18]. In the random sequential adsorption technique, tapes with a random position and angle are sequentially added to a pre-defined domain. When a tape has been accepted in the domain, it is fixed and following tapes are not allowed to intersect previously added tapes. The Monte-Carlo procedures instead add all tapes to a large domain at the beginning. The domain then shrinks to obtain a desired fibre volume fraction by rearranging the tapes. A problem with both these techniques is that without modifications or constraints, the maximum achievable fibre volume fraction is much lower than the real

material [18]. For this project, RSA was chosen due to familiarity and relative ease of implementation. Therefore, the Monte-Carlo procedures will not be covered in more detail.

Tape interactions
Depending on the purpose of the numerical model, different tape interaction conditions may be beneficial to consider. In broad terms there are three different tape interaction categories [17], which are all presented in Figure 7.

For a reduced 2D model, without bending effects considered, the method shown in Figure 7(a) provides an easy implementation. This method allows tapes to intersect and thereby discards the effect of individually modelled tapes, presence of resin pockets, and all out-of-plane variations. Analysis wise for this tape interaction method, one option is to use classical laminate theory (CLT) to evaluate effective material properties where multiple tapes intersect, as done by Selezneva *et al.* [16], described in Subsection 2.2.

A second tape interaction method, shown in Figure 7(b), does not allow intersections between tapes and also discards out-of-plane rotations (draping). This method would result in an inefficiently packed material and a low fibre volume fraction [17]. Therefore, it is not suitable for modelling of STBDCs but may be used as an intermediate step (followed by a Monte-Carlo procedure and a compression phase) or to include bending effects in a reduced 2D model.

Finally, Figure 7(c) shows the method with the highest level of fidelity. It prohibits tape intersections and includes out-of-plane draping. Therefore, the tapes may be more efficiently packed, and a higher fibre volume fraction can be achieved. At the same time, the effect of individually modelled tapes is preserved, and the out-of-plane effects are considered.

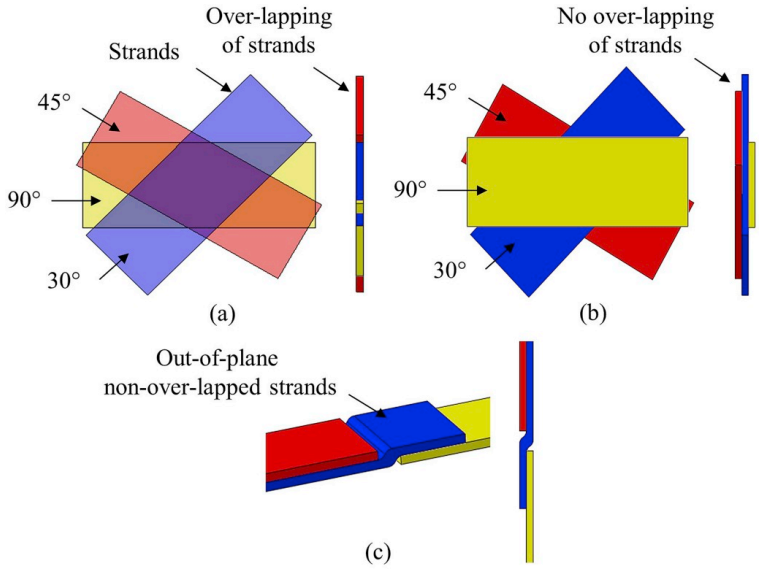

Figure 7: Different tape interaction methods: (a) overlapping strands, (b) no overlapping strands, (c) out-of-plane non-overlapped strands (draping), [17].

Draping
When one tape is placed over the edge of another tape and then cured under pressure, the upper tape will drape over the lower tape until it reaches an obstacle (another tape or a mould surface). This draping is illustrated in Figure 7(c), and shown in microstructure images in Figures 4, 5(b), and 6(a). Depending on the thickness and bending stiffness of the tapes, the draping ratio varies. For a voxel-based structure, the draping ratio (out-of-plane misalignment angle) is based on the height difference between adjacent voxels [1]. The simplest case, a 1:1 (45°) draping ratio without modifications is shown in Figure 8 a). Gentler slopes, e.g. a 2:1 ratio, may be achieved either by implementing a different draping condition, as shown in Figure 8 b). Or by increasing the in-plane dimensions of the voxels while maintaining the same height, as shown in Figure 8 c). An added benefit of the latter is that it also reduces the number of elements.

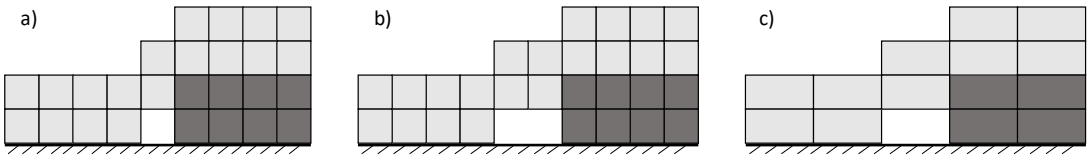

Figure 8: Voxel-based tape draping, light grey tape draping over dark grey tape, a) 1:1 draping, b) 2:1 draping, c) 2:1 draping with increased element size.

Additional constraints for consistent model generation
To avoid problems like non-uniform tape distributions and unrealistically low fibre volume fractions, a number of constraints can be used in the model generation [17]. Although the real STBDC material is in fact produced by a random tape placement, there are some effects that a numerical model cannot capture. Two effects that occur in the real material but are often neglected in computational models are tape sliding and tape deformation. In-plane sliding of tapes was observed, although to a limited extent, in [4] and implies that tapes move from their original position during the curing process. Tape deformation means that the cross-sectional area of the tapes does not remain constant during the curing process. Individual tapes get compressed and laterally expand in certain high-pressure areas resulting in thinner but wider tapes. This compression and lateral expansion was observed in [4, 15] and is one of the causes of the thickness variation in Figures 5(a) and 6(a). Both these effects aid in packing the tapes more efficiently and if neglected in a numerical model, the fibre volume fraction will likely not be representative of the real material.

It is possible that a completely stochastic model generation procedure can generate a realistic material structure, but not consistently as it relies on probability. To increase the consistency of the model generation, additional constraints can be added to the model generation procedure. To uniformly distribute the tapes over the entire plate, the algorithm created in [17] partitions the plate in a number of bins, like a grid. The algorithm then places tapes in the bins systematically, ensuring that no part of the plate has a significantly higher or lower number of tapes. Without this additional constraint, there is just a statistical probability that the tape distribution will be uniform, no guarantee. To ensure that the fibre volume fraction is representative, there are different constraints that can be implemented. The algorithm developed in [1] has the fibre volume fraction as a target variable and iteratively adds tapes to the domain until the desired fibre volume fraction has been reached. The use of fibre volume fraction as a target value inspired the method used in this project, where the plate is built layer-by-layer, without allowing the algorithm to start a new layer before the first layer has a high enough fibre volume fraction. The implementation of the layer-by-layer method is covered in further detail in Subsection 4.1.2

3.3 Finite element analysis

The following subsections cover the fundamental theory on how to conduct an FEA with the aim to evaluate homogenised material properties. Initially, the basic concepts that goes into setting up an FEA are detailed, including choice of elements, number of nodes and rule of integration. Secondly, the concept of homogenisation and the general theory of periodic boundary conditions are explained as well as how to evaluate the effective material properties.

3.3.1 Discretisation

With a voxel-based structure, the obvious choice is to use cuboid brick (hexahedral) elements. It is also possible to use tetrahedral or prismatic elements, but at a higher risk of bad aspect ratios as the cuboid voxel shape would be cut in at least two segments.

In Abaqus, there are a plethora of cuboid elements to choose from, suitable for different problems. For incompressible materials for example, so called hybrid elements (suffix H) are recommended. In this application, regular stress - displacement elements are sufficient but three major choices must be made regarding the discretisation: i) shell or solid elements, ii) number of nodes, and iii) number of integration points and integration method.

Solid or shell elements
There are two options for hexahedral shell elements in Abaqus: continuum shell elements (SC8R), and continuum solid shell elements (CSS8) [19]. These elements look like solid hexahedral elements but behave like, and have the same limitations as, regular shell elements. The main limitation of these elements being that they should only be used for thin structures [19]. For the solid hexahedral elements, there is only one choice for the base configuration (C3D$\bullet\bullet\bullet$) but it may be altered in number of nodes, number of integration points, and method of integration [19]. The solid hexahedral elements are generally not recommended in structures subjected to bending as they are either too stiff or too flexible, due to shear locking and hourglassing, respectively [19].

Number of nodes
Regarding the number of nodes for a brick element, Abaqus offers elements with 8, 20, and 27 nodes [19]. An 8-node brick element has one node in each corner, a 20-node element has 12 additional nodes - one on the centre of each edge, and a 27-node element has 6 additional nodes - one on each face and one additional node in the centre of the element. If the minimal number of nodes is used, the edges and surfaces remain straight during deformation, and the displacement approximation is trilinear. For elements with additional nodes, the edges and faces may obtain a curvature as they deform, and the approximation of the displacement is at least quadratic in all directions.

Integration rule
For the solid hexahedral elements, there are generally three choices regarding the numerical integration for solid elements: full integration, reduced integration, and incompatible modes. Fully integrated elements have just enough integration points for numerical integration to achieve exact integration. For an 8-node hexahedral element, this is $2 \times 2 \times 2 = 8$ integration points. Reduced integration has the suffix R in the element description code and means that the number of integration points is reduced. Reduced integration saves computational power and can, in practice, be used to avoid the over-stiff response of fully integrated elements [20]. Typically, the reduced integration cuboid elements give a worse approximation in bending and are not as capable of capturing surface stresses compared to fully integrated elements [19]. Reduced integration elements are also prone to hourglassing that can make them too flexible, which is why hourglass control is included as a default for these elements. Finally, the incompatible modes elements (suffix I) have been implemented to solve some of the problems related to the regular, fully integrated versions of the same element. The C3D8I element for example uses different shape functions than C3D8 and has additional internal degrees of freedom in order to reduce the shear locking phenomena and thereby increase the accuracy in bending. The incompatible mode elements are more computationally expensive than the fully integrated versions.

Quadrilateral shell elements
For a shell model, the selection of elements is not as wide as for the 3D elements. If the selection is limited to quadrilateral four node elements, excluding triangular elements, it is only one to chose, the S4• • • element with various alterations [19]. In the standard configuration, it is a general purpose shell element with a finite strain formulation, meaning it is suitable for both thick and thin plate structures with large deformations. There is also a reduced integration (suffix R) version with one in-plane integration point instead of four, also including hourglass control [19]. In addition, there are more alterations like the small strain formulations (suffix S) and elements with reduced number of degrees of freedom (suffix 5).

3.3.2 Homogenisation

The macroscale mechanical properties of any material depends on the properties of its subscales [21]. In a practical implementation, it is not possible to model every single grain in a steel or every fibre in a carbon fibre composite. Therefore, the material properties are homogenised to effective properties at a higher scale. The rule of mixtures is an example of analytical homogenisation in which the constituent material properties and volume fractions are used to calculate effective material properties. The rule of mixtures can be used in classical laminate theory to calculate some of the effective lamina properties such as the longitudinal modulus [9]. For more complex and aperiodic structures, analytical models may not be accurate enough and computational homogenisation can be used instead.

The first step of computational homogenisation is to extract or define a representative volume element (RVE) based on the full microstructure as a sampling space. For a periodic microstructure, the RVE is the smallest volume element that by periodicity represents the material [22]. An example of a periodic microstructure and its RVE is shown in Figure 9(a), where the RVE has the domain Ω and boundaries $\delta\Omega^+$ and $\delta\Omega^-$. For completely irregular (aperiodic) microstructures, the RVE should in theory be infinitely large [21]. In practice, infinitely large RVEs cannot be used. Instead, the materials are assumed to be statistically homogeneous, meaning that if the RVE is sufficiently large, the position from where it is extracted will not affect the homogenised results. Since an RVE of an aperiodic microstructure is not representative, they are commonly called statistical volume elements (SVEs) [21]. SVEs do not necessarily have be large enough to give the same material properties irrespective of how they are sampled, instead multiple smaller SVEs can be homogenised to have the same statistical significance as fewer larger SVEs. An example of an aperiodic microstructure and its SVE is shown in Figure 9(b).

The second step of the homogenisation procedure is to carry out the numerical analysis. Different boundary conditions can be used in the homogenisation process, namely: displacement based (Dirichlet), traction based (Neumann), and periodic boundary conditions (PBCs) [23]. In a displacement based homogenisation procedure, the displacement is prescribed to the domain such that it represents a predefined macroscopic strain state, and the resulting macroscopic stress state is evaluated as a volume average [21]. The effective macroscopic elasticity parameters are then evaluated from the resulting relation between macroscopic strain (prescribed) and macroscopic stress (computed). The evaluation of material parameters is further detailed in Subsection 4.2.3.

Depending on if the material response is linear or non-linear, the homogenisation process is different [21]. For a linear material response, the homogenisation can be done once and for all, a priori, and elastic engineering parameters can be determined. For a nonlinear material response, an a priori homogenisation is not sufficient. Instead, one either has to device a two-way coupled system where the homogenisation is nested with the solver [21], or perform off-line so-called virtual testing to generate data to which a macroscale surrogate model can be calibrated.

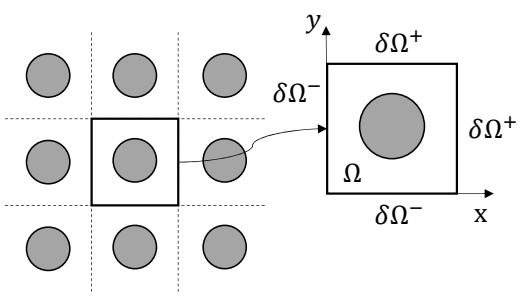

(a) Periodic microstructure and its RVE.

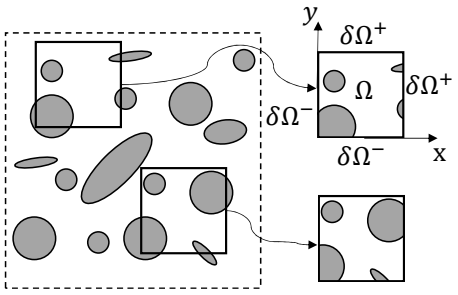

(b) Aperiodic microstructure and its SVEs.

Figure 9: Schematic examples of (a) a periodic microstructure and its RVE, and (b) an aperiodic microstructure and its SVEs.

3.3.3 Periodic boundary conditions (PBCs)

Periodic boundary conditions (PBCs) are a type of boundary condition that allows for the approximation of a large structure by using a smaller representative volume, typically these would be used on structures with a periodic phase arrangement. However, studies such as [23, 24, 25] indicate that for the case of a heterogeneous material with an aperiodic structure, PBCs provide a better prediction of the effective material properties than both Dirichlet and Neumann boundary conditions, assuming that the SVE is not large enough for all three to converge fully. It has also been shown that the use of Neumann boundary conditions provides the lower bound of the mechanical properties, Dirichlet boundary conditions give the upper bound, and periodic boundary conditions generally fall within these bounds, providing better predictions of the effective mechanical properties for aperiodic structures [24]. A visual representation of this relation between the three is shown in Figure 10. It is important to note that the better prediction of PBCs is the general case, there are special cases where PBCs might produce a less correct estimation of the effective mechanical properties for small domains, as shown in [26].

Similar studies have also come to the conclusion that in terms of convergence, in general, periodic boundary conditions are the most efficient compared to the traction (Neumann) and displacement (Dirichlet) boundary conditions [23, 25], this can also be seen in Figure 10 where the PBC reaches the effective value significantly faster.

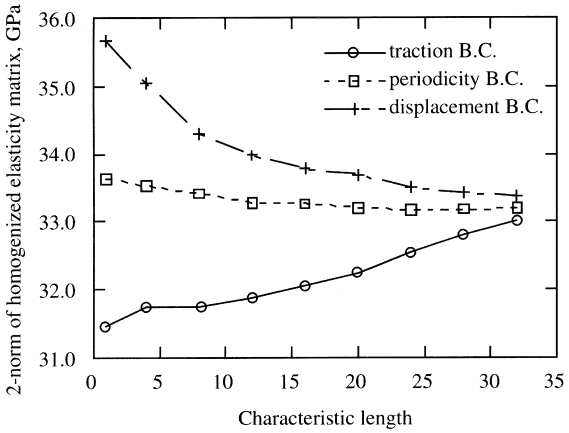

Figure 10: Trend of convergence for homogenised elasticity matrix using different boundary conditions, [25].

As a starting point of implementing periodic boundary conditions on an SVE, consider the domain Ω with boundaries $\delta\Omega$. The external boundaries $\delta\Omega$ are divided into their respective opposing parts $\delta\Omega^+$ and $\delta\Omega^-$, as shown in Figure 9. This boundary division should satisfy the following conditions:

$$\begin{aligned} &\text{(i)} \quad \delta\Omega^+ \cup \delta\Omega^- = \delta\Omega \\ &\text{(ii)} \quad \delta\Omega^+ \cap \delta\Omega^- = 0 \\ &\text{(iii)} \quad \boldsymbol{n}^- = -\boldsymbol{n}^+. \end{aligned} \tag{1}$$

The above conditions say that (i) all pairs of opposing boundary sides compose the total boundary $\delta\Omega$ and (ii) the intersection of two opposing sides should be equal to zero, meaning no intersection. The final condition (iii) says that the normal of two opposing boundaries are equal in opposite directions. Therefore, a cuboid domain is suitable in the application of periodic boundary conditions, preferably with a periodic mesh meaning that the mesh distribution is the same on the parallel faces.

The displacement field $\boldsymbol{u}$ of an SVE submitted to a macroscopic strain can be divided in to two parts: the displacement due to the applied macroscopic strain denoted $\bar{\boldsymbol{u}}$ and a zero-mean fluctuations field $\tilde{\boldsymbol{u}}$, as seen in Equation (2). The periodicity is enforced such that the fluctuating part of two parallel boundaries are equal. This implies that when setting up a displacement constraint between two sides, the only remaining part is the contribution of the linear displacement field (derived from the applied macroscopic strain). This results in the general form of the constraint equation, presented in Equation (3)

$$u_i = \bar{u}_i + \tilde{u}_i = \bar{\varepsilon}_{ij} x_j + \tilde{u}_i, \text{ thus} \tag{2}$$

$$u_i^+ - u_i^- = \bar{\varepsilon}_{ij} x_j^+ - \bar{\varepsilon}_{ij} x_j^- + \underbrace{\tilde{u}_i^+ - \tilde{u}_i^-}_{0} = \bar{\varepsilon}_{ij}(x_j^+ - x_j^-), \tag{3}$$

where $\bar{\varepsilon}_{ij}$ is the macroscopic strain and x_j^+ and x_j^- are the respective coordinate of the material point on opposing boundaries. Thus, $(x_j^+ - x_j^-)$ represents the distance between paired external boundaries. Equation (3) ensures the continuity of the displacement field across the domain. A similar constraint is enforced on the traction as: $\boldsymbol{t}^+ = -\boldsymbol{t}^-$, defined as an anti-periodicity condition. This comes from the definition of the traction $t_i = \bar{\sigma}_{ij} n_j$ which should satisfy condition (iii) in Equation (1) for all pairs $\delta\Omega^+$ and $\delta\Omega^-$. These in turn, also fulfil the Hill-Mandel principle, ensuring that the sub-scale modelling is energetically consistent, and therefore valid. The full theoretical background of the Hill-Mandel principle will not be covered in this thesis, for further reading see [23].

3.3.4 Evaluation of effective material properties

The effective material properties are evaluated using the known/assumed constitutive relation (in Voigt notation) between the macroscopic stress ($\bar{\boldsymbol{\sigma}}$) and strain ($\bar{\boldsymbol{\varepsilon}}$) via the compliance matrix ($\boldsymbol{S}$) as $\bar{\boldsymbol{\varepsilon}} = \boldsymbol{S}\bar{\boldsymbol{\sigma}}$. The material at hand is anisotropic on the lower scale, but in the ideal case quasi-isotropic, or in-plane isotropic, on the macroscale [4]. Therefore, in order to evaluate the effective material properties, an orthotropic constitutive relation is used. This is presented in Equation (4), where $\boldsymbol{S}$ is the orthotropic compliance matrix. As orthotropic materials have material properties that differ along three orthogonal axes, quasi-isotropy is a special case of orthotropic materials. In theory, with the material being ideally quasi-isotropic, the two in-plane components of the elastic stiffness (E_x and E_y), the in plane Poisson's ratios (ν_{xy} and ν_{yx}), and the out-of-plane shear moduli (G_{xz} and G_{yz}) should be identical.

$$\begin{bmatrix} \varepsilon_{xx} \\ \varepsilon_{yy} \\ \varepsilon_{zz} \\ \varepsilon_{yz} \\ \varepsilon_{zx} \\ \varepsilon_{xy} \end{bmatrix} = \underbrace{\begin{bmatrix} \frac{1}{E_x} & -\frac{\nu_{yx}}{E_y} & -\frac{\nu_{zx}}{E_z} & 0 & 0 & 0 \\ -\frac{\nu_{xy}}{E_x} & \frac{1}{E_y} & -\frac{\nu_{zy}}{E_z} & 0 & 0 & 0 \\ -\frac{\nu_{xz}}{E_x} & -\frac{\nu_{yz}}{E_y} & \frac{1}{E_z} & 0 & 0 & 0 \\ 0 & 0 & 0 & \frac{1}{2G_{yz}} & 0 & 0 \\ 0 & 0 & 0 & 0 & \frac{1}{2G_{zx}} & 0 \\ 0 & 0 & 0 & 0 & 0 & \frac{1}{2G_{xy}} \end{bmatrix}}_{\boldsymbol{S}} \begin{bmatrix} \sigma_{xx} \\ \sigma_{yy} \\ \sigma_{zz} \\ \sigma_{yz} \\ \sigma_{zx} \\ \sigma_{xy} \end{bmatrix} \tag{4}$$

Extracting the full set of equations from Equation (4) results in the following relations between the different material properties:

$$\begin{gathered} \varepsilon_{xx} = \frac{\sigma_{xx}}{E_x} - \frac{\nu_{yx}\sigma_{yy}}{E_y} - \frac{\nu_{zx}\sigma_{zz}}{E_z}, \quad \varepsilon_{yy} = \frac{\sigma_{yy}}{E_y} - \frac{\nu_{xy}\sigma_{xx}}{E_x} - \frac{\nu_{zy}\sigma_{zz}}{E_z}, \quad \varepsilon_{yy} = \frac{\sigma_{yy}}{E_y} - \frac{\nu_{xy}\sigma_{xx}}{E_x} - \frac{\nu_{zy}\sigma_{zz}}{E_z}, \\ \varepsilon_{yz} = \frac{\sigma_{yz}}{2G_{yz}}, \qquad \varepsilon_{zx} = \frac{\sigma_{zx}}{2G_{zx}}, \qquad \varepsilon_{xy} = \frac{\sigma_{xy}}{2G_{xy}}. \end{gathered} \tag{5}$$

Using periodic boundary conditions, the applied strain is known for the different load cases. The macroscopic stress state, however, is unknown but may be evaluated as a volume average using Gauss integration [24] as:

$$\bar{\sigma}_{ij} = \frac{1}{\Omega}\int_{\Omega} \sigma_{ij} d\Omega = \frac{1}{\Omega}\sum_{e=1}^{N_\mathrm{e}} \sum_{I=1}^{N_\mathrm{e,int}} \sigma_{ij}(y_I) J(y_I) W(y_I), \tag{6}$$

where $\bar{\sigma}_{ij}$ is the average stress, e denotes element, y_I is the integration point, J is the determinant of the Jacobian, and W is the integration weight. Looping through all elements and integration points of a large RVE is not computationally efficient and the volume average stress may instead be obtained by evaluating the traction at the boundary. Applying the Gauss theorem on Equation (6) yields:

$$\bar{\sigma}_{ij} = \frac{1}{\Omega}\int_{\Omega} \sigma_{ij} d\Omega = \frac{1}{\Omega}\int_{\partial\Omega} \sigma_{ij} n_j x_j d\partial\Omega = \frac{1}{\Omega}\int_{\partial\Omega} t_i x_j d\partial\Omega \approx \frac{1}{\Omega}\sum_{n=1}^{N_\mathrm{node}} F_i x_j, \tag{7}$$

where t_i is the traction, n_j the outwards facing normal, x_j the position, $\partial\Omega$ the boundary, and F_i the nodal force. In practice, the average stress can be evaluated as the total reaction force of a surface divided by the cross-sectional area, as shown in Equation (8), where F_i^tot is the total force and l_i is the SVE side length.

$$\begin{aligned} \bar{\sigma}_{xx} &= \frac{F_x^\mathrm{tot}}{l_y l_z} & \bar{\sigma}_{yy} &= \frac{F_y^\mathrm{tot}}{l_x l_z} & \bar{\sigma}_{zz} &= \frac{F_z^\mathrm{tot}}{l_x l_y} \\ \bar{\sigma}_{xy} &= \frac{F_y^\mathrm{tot}}{l_y l_z} & \bar{\sigma}_{xz} &= \frac{F_z^\mathrm{tot}}{l_y l_z} & \bar{\sigma}_{yz} &= \frac{F_z^\mathrm{tot}}{l_x l_z} \end{aligned} \tag{8}$$

3.4 Classical laminate theory - analytical homogenisation

Analysing a laminated composite with layered plies of known angles allows for the use of classical laminate theory (CLT) in order to evaluate the effective in-plane properties: $\bar{E}_{11}$, $\bar{E}_{22}$, $\bar{G}_{12}$, and ν_{12}, under the assumption that the out-of-plane stress is zero. CLT does not apply to STBDCs on the macroscale, since the material is neither continuous nor structured in layers. But on a sufficiently small scale, the material is actually both structured in layers and continuous. How CLT was used for homogenisation on a smaller scale is presented further in Subsection 4.3. The smaller scale laminate structure of the STBDC material does not only consist of tapes (plies) but could also consist of pure matrix. Therefore, the ply/lamina/pure matrix will from this point be referred to as layers. Knowing the longitudinal, transverse, and shear material properties for each transverse isotropic layer, the stress strain relation can be expressed with the reduced stiffness matrix $\boldsymbol{Q}$, defined through its four independent components. The definition of $\boldsymbol{Q}$ in Voigt notation for a transverse isotropic layer and its respective components are stated in Equations (9) and (10). The subscripts L and T in Equation (10) represent the longitudinal and transverse directions, respectively. Furthermore, since the layers can have different angles and the material properties are direction dependent, the reduced stiffness matrix needs to be transformed to comply with the global coordinate system. The transformation between the local and global coordinate system is done through Equation (9), where $\boldsymbol{T}_1$ is the stress-transformation matrix and $\boldsymbol{T}_2$ is the strain transformation matrix, the components of these will not be stated in this report but can be found in [9].

$$\bar{\boldsymbol{Q}} = \boldsymbol{T}_1^{-1}\boldsymbol{Q}\boldsymbol{T}_2 \quad \text{with} \quad \boldsymbol{Q} = \begin{bmatrix} Q_{11} & Q_{12} & 0 \\ Q_{12} & Q_{22} & 0 \\ 0 & 0 & Q_{66} \end{bmatrix} \tag{9}$$

$$Q_{11} = \frac{E_L}{1-\nu_{LT}\nu_{TL}}, \qquad Q_{22} = \frac{E_T}{1-\nu_{LT}\nu_{TL}}, \qquad Q_{12} = \frac{\nu_{TL}E_L}{1-\nu_{LT}\nu_{TL}}, \qquad Q_{66} = G_{LT} \tag{10}$$

Once the transformed stiffness matrix $\bar{\boldsymbol{Q}}$ has been evaluated for the different possible angles within the system, the laminate stiffness matrix $\boldsymbol{K}$ can be defined as the relation between the moment and stress resultants (M, N), and laminate mid-plane strains and curvatures ($\boldsymbol{\varepsilon}^0, \boldsymbol{k}$). The relation has been defined in Equation (11):

$$\begin{Bmatrix} \boldsymbol{N} \\ \cdots \\ \boldsymbol{M} \end{Bmatrix} = [\boldsymbol{K}] \begin{Bmatrix} \boldsymbol{\varepsilon}^0 \\ \cdots \\ \boldsymbol{k} \end{Bmatrix} = \begin{bmatrix} \boldsymbol{A} & \boldsymbol{B} \\ \boldsymbol{B} & \boldsymbol{D} \end{bmatrix} \begin{Bmatrix} \boldsymbol{\varepsilon}^0 \\ \cdots \\ \boldsymbol{k} \end{Bmatrix}, \tag{11}$$

where $\boldsymbol{A}$, $\boldsymbol{B}$, and $\boldsymbol{D}$ are the *extensional stiffness matrix, coupling stiffness matrix,* and *bending stiffness matrix*, respectively [9]. The evaluation of the $\boldsymbol{A}$, $\boldsymbol{B}$, and $\boldsymbol{D}$ matrices has been stated in Equation (12). The factor $(h_k - h_{k-1})$ denotes the distance between the top and bottom surface of each layer relative to the mid-plane, resulting in the layer thickness.

$$\boldsymbol{A} = \sum_{k=1}^{n} (\bar{\boldsymbol{Q}})_k (h_k - h_{k-1}), \qquad \boldsymbol{B} = \frac{1}{2}\sum_{k=1}^{n} (\bar{\boldsymbol{Q}})_k (h_k^2 - h_{k-1}^2), \qquad \boldsymbol{D} = \frac{1}{3}\sum_{k=1}^{n} (\bar{\boldsymbol{Q}})_k (h_k^3 - h_{k-1}^3) \tag{12}$$

Once all component matrices of the stiffness matrix $\boldsymbol{K}$ have been evaluated, the effective engineering constants can be calculated using the compliance matrix $\boldsymbol{S}$ [16], defined as:

$$\boldsymbol{S} = \begin{bmatrix} \boldsymbol{a} & \boldsymbol{b} \\ \boldsymbol{b} & \boldsymbol{d} \end{bmatrix} = \begin{bmatrix} \boldsymbol{A} & \boldsymbol{B} \\ \boldsymbol{B} & \boldsymbol{D} \end{bmatrix}^{-1}, \quad \text{where} \quad \boldsymbol{a} = \begin{bmatrix} a_{11} & a_{12} & a_{16} \\ a_{21} & a_{22} & a_{26} \\ a_{61} & a_{26} & a_{66} \end{bmatrix}. \tag{13}$$

From this, the in-plane material parameters can be evaluated from the inverse components of the compliance matrix as:

$$\bar{E}_{11} = \frac{1}{a_{11}h}, \qquad \bar{E}_{22} = \frac{1}{a_{22}h}, \qquad \bar{G}_{12} = \frac{1}{a_{66}h}, \qquad \nu_{12} = \frac{-a_{12}}{a_{11}}, \tag{14}$$

where h is the total thickness of the laminate.

4 Methodology

In this section, the methodology used to realise the aim throughout the project are detailed. It starts with an in-depth explanation of the geometrical modelling of the STBDC material, including both an overview of the entire procedure and in-depth details on all the major topics. Then covering the finite element analysis, in a more practical implementation manner. This is followed by a description of how the reduced 2D-model was developed and ends with a design of experiment subsection.

4.1 Geometrical modelling

The first part of the project was to create a voxel-based numerical model, replicating the microstructure of a compression moulded STBDC. A modified RSA technique was used, partly due to the ease of implementation and partly since it is the most commonly used in literature. This section covers the methodology used in order to create the modular MATLAB code with the functionality of generating a realistic material structure. A flowchart outlining the general functionality of the developed MATLAB code is shown in Figure 11.

In short, the flowchart shown in Figure 11 describes the following processes: a main file is used to manage settings such as tape dimensions and toggles such as which plots to plot and if an input file is wanted or not. The geometry is then created using multiple different functions, detailed further in the following subsections. Based on the toggles and settings defined in the main file, the geometrical post-processing is completed and wanted plots produced. The relevant data is then saved in a separate directory for later access. Based on toggles, the sides and top of the domain are cut to a predetermined SVE size to remove edge effects and resin rich top layers. If wanted, an input file is written and saved. Also, a through thickness homogenisation procedure can be done and a reduced 2D-model input file written. Finally, a data check can be run using Abaqus as well as the full simulation. Abaqus Viewer may be automatically opened once the simulation is completed to visualise the results. The entire process can be looped to run multiple simulations in series, without using the Abaqus CAE graphical user interface.

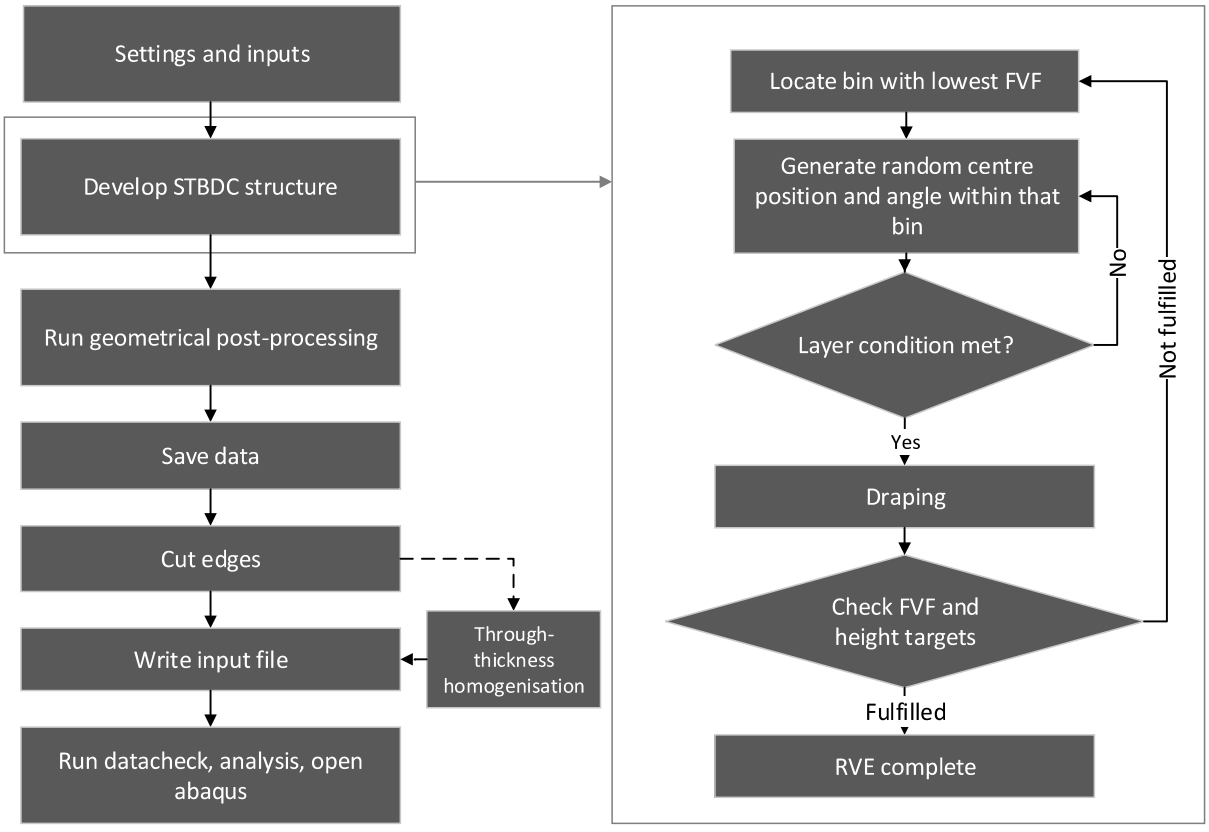

Figure 11: Schematic flowchart of the algorithm used to produce and analyse the digital STBDC material structure.

4.1.1 Definition of domain

For the ease of manipulation and data storage, a 3D array was chosen to virtually depict the full domain of the STBDC plate. For future reference, the indices i, j, k are used to represent the row, column, and layer or y, x, and z axes of the 3D array, respectively. Every tape placed in the domain is stored with an identifiable integer in sequential order (tape index) and the rest of the array is filled with zeros representing the polymer matrix.

In terms of size, the length and width of the domain were set to be equal and defined as a function of the tape length. Similarly, the size of the SVEs was based approximately on the tape length. Doing this, a relation between the size of the SVE and full domain could be maintained, ensuring that it was sufficiently large. The height of the domain depends on the maximum number of stacked tapes, which is a user input parameter. The maximum size of the domain was limited by the available memory of the computer. This inhibits the creation of large structures with high resolution, for a computer with 16 Gb RAM memory, the size of the 3D array was limited to roughly 7000×7000×36 elements. For reference, a tape modelled in high resolution could require 4000 elements along the length of the tape.

The full domain was partitioned twice into two different subdomains, shown in Figure 12. Firstly, the full domain was partitioned by the safety margin l_s to create the safe space subdomain, shown in Figure 12(a). The placement of each tape is governed by its centre position, the centre position may not exceed the boundaries of the internal safe space subdomain. The safety margin was set equal to one tape length, $l_s = l_t$. The rationale for this was to avoid any parts of the tape exceeding the outer boundary which would complicate the programming as MATLAB does not allow calling positions outside the array limits. The second partitioning was to divide the safe space subdomain into a number of so-called bins, which were used to force a more uniform tape distribution (further detailed in Subsection 4.1.2). The layout of these bins are shown in Figure 12(b).

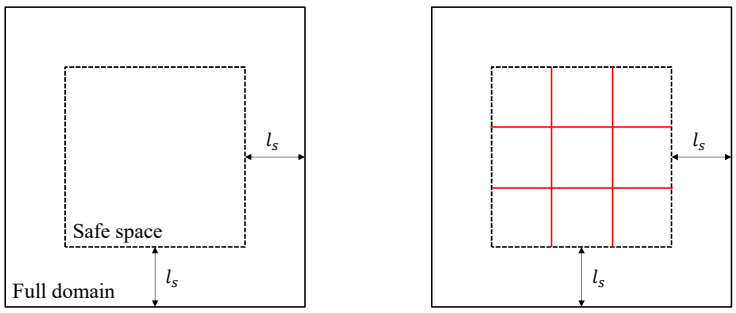

(a) 3D array with subdomain division.

(b) 3D array with subdomain and bin division.

Figure 12: Schematic visualisation of geometric domain and subdomains.

4.1.2 Tape generation and placement

The individual tapes making up the domain, with pre-defined width w^t, length l^t, and thickness t^t, were constructed using cuboid volume pixels (voxels). The amount of voxels used to build up each tape was a result of the user-defined tape dimensions, voxel aspect ratio R_{vox}, and number of voxels in the height direction per tape, $N^t_{\text{vox},h}$ (all input-values to the code). The voxel aspect ratio dictates the in-plane dimensions to thickness relation of the voxels. The number of length voxels ($N^t_{\text{vox},l}$), and number of width voxels ($N^t_{\text{vox},w}$) for each tape were calculated as:

$$N^t_{\text{vox},l} = \text{round}\left(\frac{l_t N^t_{\text{vox},h}}{t_t R_{\text{vox}}}\right), \qquad N^t_{\text{vox},w} = \text{round}\left(\frac{w_t N^t_{\text{vox},h}}{t_t R_{\text{vox}}}\right), \tag{15}$$

where the operator round(•) refers to the MATLAB function with the same name which rounds (•) to the nearest integer value. Increasing the amount of voxels defining each tape results in an increased resolution, providing a more accurate model at the cost of an increased need for computational power.

The placement of each tape inside the domain was decided strictly based on its randomised centre position (X_c, Y_c) and tape angle (φ), defined in Figure 14(b). The centre position was randomised, however, the algorithm has a condition implemented to ensure that the search for an open position was more efficient, and to ensure a uniform tape distribution. For each new iteration, the tape content of each bin (see Figure 12(b)) was evaluated and the tape was randomly placed in the bin with the lowest tape content. This ensured not only that the search area was smaller but that the randomised tape placement was more uniform across the different layers. Another restriction set on the modelling was that the placement of tapes were layer constrained (layer condition in Figure 11), this implies that initially tapes may only be placed in the first layer. Once a user-specified desired volume fraction was achieved for the active layer, tapes may be placed one layer higher. This was then repeated until the user-specified number of layers (plies) had been reached.

The tape angle was pseudo randomised, meaning it was randomly chosen from a set of angles in the range 0°-179° in equal increments determined by the user. An angle increment of 1°, for example, yields the following set of angles to choose from: [0°, 1°, ... , 178°, 179°]. The reason for using 179° instead of 180° as an upper limit was to prevent the average tape angle of all tapes to be skew towards the x-direction. If 180° was used as an upper limit, both 0° and 180° would represent the x-direction whereas only 90° would represent the y-direction. The tape angle and centre position was sufficient to represent a tape with straight edges but to fit the voxel-based architecture, this linear representation was required to be converted to voxel format, as exemplified in Figure 13.

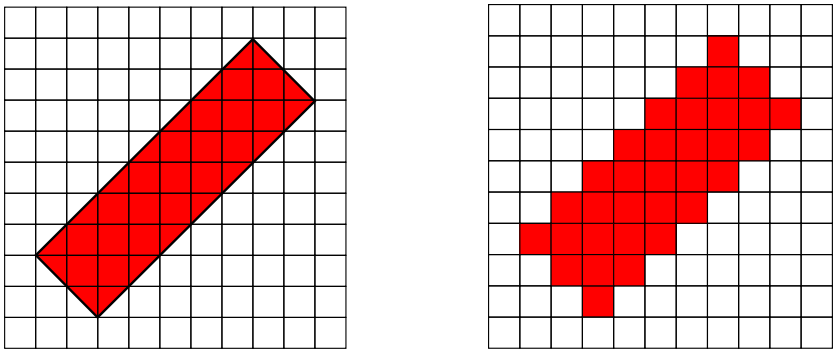

(a) Visualisation of tape discretised with straight edges - linear representation.

(b) Visualisation of tapes with ragged edges - voxel format.

Figure 13: Schematic visualisation of two different tape representations - linear or voxel.

To do the linear to voxel conversion, a Bresenham algorithm created by A. Wetzler was implemented in MATLAB [27]. This is a line drawing algorithm that given two points in a 2-dimensional space provide a pixelated approximation of the line between the two points [28]. Bresenham's algorithm therefore allows for the transformation of angled lines within the domain into its corresponding pixelated representation, this process is visualised in Figure 14(a). The tapes were created by initially using the corner coordinates in combination with Bresenham's algorithm to draw all the edges of one tape, as seen in Figure 14(b). The rest of the tape was then filled in by looping through all the x-positions of the tape and filling it from the minimum to the maximum y-position. One drawback to the voxelisation of the tape was that if the resolution was not high enough, multiple given tape angles could result in identical structures. For example, there was no visual difference between tapes with 0° and 1° tape angles, both were perfectly aligned with the x-axis.

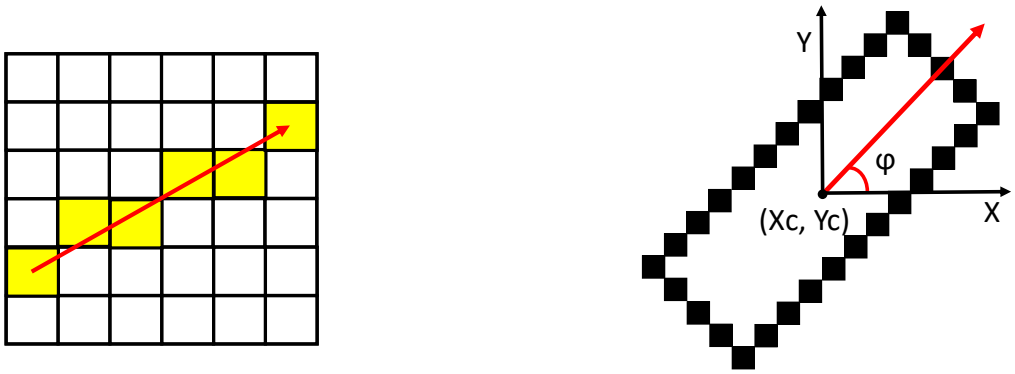

(a) Illustration of the functionality of Bresenham's algorithm.

(b) Visual representation of centre position and tape angle.

Figure 14: Definitions used to define a voxel-based tape

Once the tape position within the domain and the voxel-based tape representation was complete, the entire voxelised tape was placed at the lowest available height, without penetrating another tape already placed and not breaching the height restriction. If this could not be achieved at the given position, the centre position and angle were re-selected until the tape could be placed while not penetrating and staying within the active layers. This step was repeated for a user-defined number of iterations. If the maximum number of iterations was reached, the fibre volume fraction target was lowered. This final step was repeated, either until a possible tape position was found or until the fibre volume fraction target was lowered enough to let a new layer be created.

4.1.3 Draping algorithm

The process of creating a realistic geometric model emulating a STBDC involved allowing the tapes to adapt as they fall on top of each other (form around each other). In order to capture this during the geometry generation, a draping algorithm was implemented, simulating how the real material acts when cured under pressure. The draping algorithm does not allow for any discontinuity of tapes and the draping ratio (out-of-plane angle) is controlled by the thickness to length ratio of the voxels (R_{vox}) as shown in Figure 8(c). If the in-plane dimensions of each voxel is increased, the draping angle is decreased.

The draping algorithm can be divided into three main parts: i) when a suitable in-plane position has been located for a voxelised tape, all voxels are placed at their lowest available position in the z-direction. ii) Once the tape has been placed, a topology matrix is established, describing the highest position of each voxel position of the current tape from a top-view perspective, as seen in the leftmost image in Figure 15. The purpose of this is to locate any discontinuities of the tape, i.e. where the height difference between adjacent voxels is larger than one. iii) The discontinuities are sequentially removed using a search algorithm, looping over all elements in the topology map and calculating the height differences between the current element and its respective neighbours. If any height difference between two neighbouring elements is larger than one, the current element voxel position is moved up one position in the out-of-plane direction (see yellow-marked voxels from left to right in Figure 15). Changing the out-of-plane position of a voxel alters the topology matrix and it has to be updated. Therefore, steps ii) and iii) are repeated until no more changes are made to the out-of-plane position of the tape.

A simplified visualisation of how this draping algorithm works is presented in Figure 15. The leftmost illustration shows how a light grey tape lands on a dark grey tape and takes the lowest possible position. The remaining illustrations in Figure 15 show how the nearby voxels move upward to satisfy the one-to-one draping condition. The evolution of the topology matrix for the light grey tape is shown underneath.

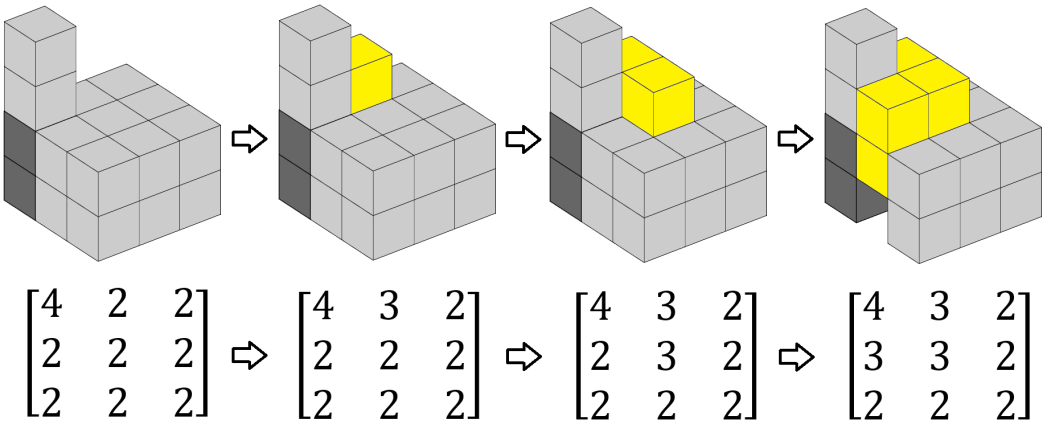

Figure 15: Draping algorithm example, yellow boxes indicating moved voxels. Corresponding topology matrix shown underneath.

The draping ratio was set to 5:1, and achieved by using voxels with a length and width to height ratio (R_{vox}) of 5:1, as schematically shown in Figure 8 (c). This draping ratio was determined by examining microscope images from three different studies and measuring the height and length of the matrix filled crevices. These measurements are shown in Figure 16. Using a 5:1 voxel aspect ratio did not only achieve a more realistic draping ratio but also reduced the number of elements in the resulting models by a factor of 25.

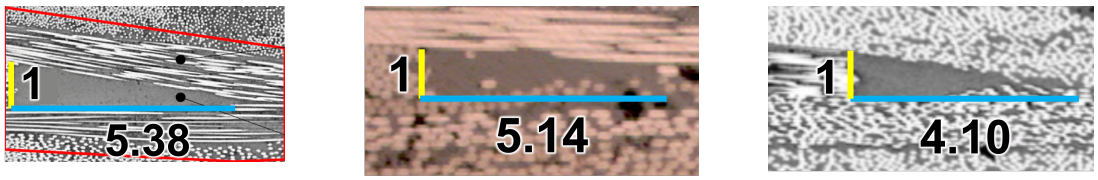

(a) Microstructure from [4], draping ratio 5.38:1.

(b) Microstructure from [15], draping ratio 5.14:1.

(c) Microstructure from [29], draping ratio 4.10:1

Figure 16: Microstructure draping measurements from three different studies.

4.1.4 Geometric material properties

Once a structural model had been established, geometric outputs could be extracted. These were strictly based on the material microstructure and properties of the constituents. The outputs were used not only to gather an increased understanding of the layout of the tapes within the material, but a number of these were also be compared to available experimental data as a comparative analysis. Based on the comparison, conclusions could be drawn on whether the produced model was an accurate depiction of reality or not.

As introduced in Section 4.1.1, the 3D array domain was presented using a somewhat unconventional sign convention. The domain D_{ijk} has three axes: i representing the row or y-axis, j representing the column or x-axis, and k representing the layer or z-axis.

Fibre volume fraction

The global fibre volume fraction was established by evaluating how many positions within the 3D-domain that were occupied by tapes. This was more specifically done by using a binary 3D-array, D_{ijk}^{binary}, where each position of tapes was described as a 1 and matrix as a 0. Therefore, the sum of all elements in D_{ijk}^{binary} equals the number of elements occupied by tapes. To show the in-plane variations of fibre volume

fraction, it was evaluated through the thickness in every in-plane coordinate as:

$$\left(V_f^p\right)_{ij} = \frac{1}{N_k} \sum_{k=1}^{N_k} D_{ijk}^{\text{binary}} \cdot V_f^t, \tag{16}$$

where $(V_f^p)_{ij}$ is an array containing the through thickness fibre volume fraction in each position of the plate, N_k is the number of voxels of the domain in the thickness (k) direction, and V_f^t is the fibre volume fraction of the individual tapes. Then the global plate average fibre volume fraction (V_f^p) could be evaluated as:

$$V_f^p = \frac{1}{N_i N_j} \sum_{i=1}^{N_i} \sum_{i=1}^{N_j} \left(V_f^p\right)_{ij} \tag{17}$$

where N_i and N_j are the number of voxels in the i and j directions, respectively. The in-plane variation of fibre volume fraction was illustrated by plotting $(V_f^p)_{ij}$ in a 2D-plot, like a top-view image.

Average in-plane angles
Due to the stochastic nature of the tape orientation in the material, it was of interest to see how the in-plane angle varies across the structure and how it affects the material properties of the plate. The possible angles are in the range 0°-179°, where a tape with a 0° angle is aligned with the x-axis and a tape with a 90° angle is aligned with the y-axis. To visualise the average tape angle on a linear scale, all angles in the second quadrant (91°-179°) were converted to their respective angle in the first quadrant (0°-90°). For example, 135° was converted to 45° by the following equation: $45° = \arccos(|\cos(135°)|)$.

The average in-plane angle was visualised by looping through the i, j positions in D_{ijk} and taking a through-thickness average of all tape angles in each position (similar to Equation (16)). The result of which was a 2D-array containing the through-thickness average angle in each position. If the active position had no tapes in the k-direction, the value was replaced by $-90°$ to further clarify the visualisation. This means that the range of values varies between 90° (oriented along the y-axis), 0 ° (oriented along the x-axis) and $-90°$ (only matrix through the thickness). Finally, the global average angle was evaluated in a similar manner to Equation (17), where all $-90°$ (matrix) values were excluded.

Level of randomisation
There are three measurements of the level of randomisation used to measure the uniformity of the produced plate and to find potential biases. The level of randomisation was measured for the tape angle and tape centre position. The principle was that the tape angle and centre position should converge towards the mean as the number of tapes were increased. In terms of centre positions, this implies that the mean of all tape centre positions in both x and y should converge towards the middle point in both directions. Similarly, the global average tape angle should converge to 45°, using the quadrant conversion described in the previous paragraph.

Microstructure visualisation
One of the main visual properties of STBDCs is the distribution of tapes between the layers due to the compression moulding. Therefore, a numerical microstructure image was created to enable the comparison to a real material microstructure. To create the microstructure image, a slice from the centre of the D_{ijk} array is taken along the i direction, which results in a new cross-section array C_{ik} containing zeros, representing matrix, and integers (tape indices) for the individual tapes. Then, the cross-section array was modified by looping through all i, k indices and replacing the zeros and tape indices with either -1 for matrix or $\cos(\varphi)$ for the tapes, where φ is the tape angle, which is known for each tape index. This results in a modified cross-section array with values in the range $[-1, 1]$ where 1 represents transverse fibres, 0 represents aligned fibres, and -1 represents matrix. This was plotted as a 2D-image with a black, white, and red colour scheme where black represents transverse fibres (normal to the cross-section), white represents aligned fibres (horizontally aligned with the cross-section), and red represents matrix and is clearly set apart from the rest. The black and white colour scheme was chosen to replicate real microstructure images where transverse fibres tend to have a darker colour and aligned fibres tend to have a lighter colour, see Figure 5 for reference.

4.2 Finite element analysis

The following subsections describe the finite element analysis in a more practical implementation manner, starting with the discretisation (choice of elements), followed by how the periodic boundary conditions were applied. Then the evaluation of elastic properties is covered, and the section ends with an explanation of the Abaqus input file and how it was created. Due to the structure of the input file, the implementation of continuum shell elements and continuum solid shell elements was not feasible in the 3D-model and therefore, omitted.

4.2.1 Discretisation

To determine which element to use for the full 3D-model, two tests were conducted to evaluate the performance of different elements. The first test was for a 25x25x28 voxel structure with randomly oriented tapes. With this random structure, no baseline of what the stiffness should be could be calculated with accuracy. For the second test, a laminated composite was used instead to enable the analytical prediction of homogenised material properties, to use as a baseline. The layup was set to $[90°/0°/0°/90°]$ with a 0.25 mm ply thickness and the voxel structure was set to 25x25x4 voxels in the x, y, and z directions, respectively. Analytical homogenisation of elastic moduli and Poisson's ratio was done based on the method presented in Subsection 3.4. Both tests used material data presented in [2], summarised in Table 2. A summary of the test results is presented in Table 3.

The results show some difference in material properties using different elements for the randomly oriented tape structure where C3D8R elements consistently gives the lowest modulus values and C3D8 elements the highest, with C3D8I elements in the middle. For the laminated composite, all elements give the same results, identical to the analytical solution. The only difference lies in solution time. It should be noted that the laminated composite is not very representative as it is not exposed to the same stress concentrations and bending effects as the randomly oriented structures typically are. In conclusion, the results presented in Table 3 show no clear indication of which element type that is the best for this application, neither which one is the worst. The C3D8I element was therefore chosen based on its enhanced capability to handle bending, even though it seems to be the most computationally demanding.

Table 3: Elastic properties using different elements for a) randomly oriented tapes, b) composite laminate.

a) Randomly oriented tapes, 25x25x28 voxels.								
Element	E_{11} [GPa]	ν_{12}	ν_{13}	E_{22} [GPa]	ν_{21}	ν_{23}	G_{12} [GPa]	Time [sec]
C3D8	26.66	0.324	0.193	32.12	0.391	0.192	13.32	216.67
C3D8R	25.43	0.319	0.197	30.39	0.379	0.196	12.5	214.24
C3D8I	26.42	0.323	0.193	31.76	0.389	0.193	13.15	223.56
b) Laminated composite: $[90°/0°/0°/90°]$ layup, 25x25x4 voxels.								
Element	E_{11} [GPa]	ν_{12}	ν_{13}	E_{22} [GPa]	ν_{21}	ν_{23}	G_{12} [GPa]	Time [sec]
Analytical	74.93	0.041	-	74.93	0.041	-	5.60	-
C3D8	74.93	0.041	0.247	74.93	0.041	0.247	5.60	90.42
C3D8R	74.93	0.041	0.247	74.93	0.041	0.247	5.60	96.43
C3D8I	74.93	0.041	0.247	74.93	0.041	0.247	5.60	89.92

4.2.2 Implementation of periodic boundary conditions

Studies such as [23, 24, 25] indicate that for the case of a heterogeneous material with a non-periodic structure, periodic boundary conditions provide a more accurate prediction of the effective material properties than Neumann or Dirichlet boundary conditions. Therefore, PBCs were adopted in the numerical model for the evaluation of the in-plane effective material properties. The periodic boundary conditions were implemented using the Abaqus CAE plug-in, EasyPBC [22]. The following paragraphs cover how the plug-in EasyPBC works, detailing how the different boundary sets are defined and the set-up of the constraint equations, following the methodology presented in [22].

EasyPBC uses linear multi-point constraints as kinematic constraints (imposing uniform strain) on the degrees of freedom of each pair of nodes connected to opposite faces, edges and vertices [30]. The general definition of this constraint equation has been presented in Equation (3), with the derivation being detailed in Subsection 3.3.3.

The plug-in initially divides the external boundaries into three separate sets: i) The faces, defined as the nodes on each directional face excluding the edge and corner nodes, thus only leaving an inner face region. ii) The edges, which are the sets that contain the nodes situated on the edges of the SVE, excluding the end nodes (corners). iii) The vertices, defined as the 8 corner nodes of the domain.

In EasyPBC, the distinction of the three sets is defined through the individual nodal coordinate values. A visual illustration of this distinction has been included in Figure 17(a), for the front face. The basic idea is that through the use of maximum coordinate values, determine which nodes should be included or not in each set [22]. This principle is then used until all the sets are defined, the result of which can be seen in Figure 17(b).

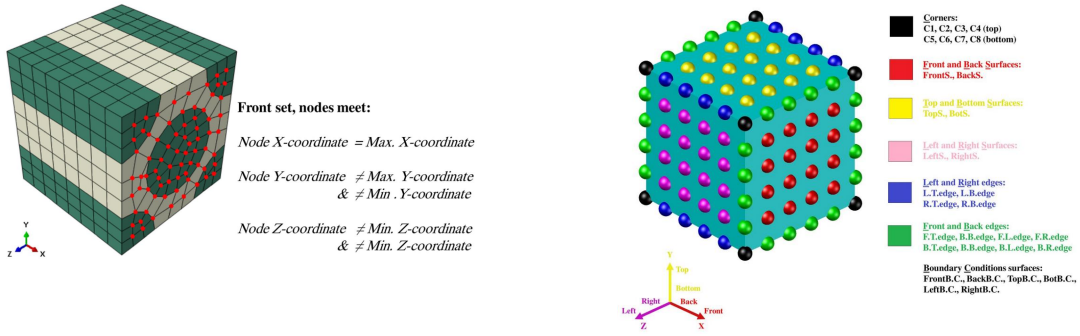

(a) Front face definition using EasyPBC [22].

(b) Illustration of created sets [22].

Figure 17: Set definitions for the application of periodic boundary conditions [22].

For the face sets, the two opposing sides of the SVE were coupled node-wise by pre-defining their relation in terms of displacement (distance between the nodes after deformation). The full set of constraint equations has been defined in Equation (18). This equation covers the 3 translations with α, β representing the respective degrees-of-freedom x, y and z. The difference in terms of displacement between the two opposing sides was set equal to the displacement ($\varepsilon_{\beta\alpha} l_\beta$) of the reference node. l_β represents the initial distance between the faces, also defined as the SVE side-length and $\varepsilon_{\beta\alpha}$ defines the applied strain. Important to note is that the sign of the right-hand side of the equations is dependent on if the boundary in the positive direction is written first or second, for further clarification see Equation (3) and Figure 9.

$$u_\alpha^{\text{FrontS}} - u_\alpha^{\text{BackS}} = \varepsilon_{x\alpha} l_x \quad u_\alpha^{\text{TopS}} - u_\alpha^{\text{BottomS}} = \varepsilon_{y\alpha} l_y \quad u_\alpha^{\text{LeftS}} - u_\alpha^{\text{RightS}} = \varepsilon_{z\alpha} l_z \tag{18}$$

As an example, for the case of applying a strain in the x-direction (ε_{xx}), the respective equations take on the following form:

$$u_{\text{x}}^{\text{FrontS}} - u_{\text{x}}^{\text{BackS}} = \underbrace{\varepsilon_{xx} l_x}_{u_x^{RP}} \quad u_{\text{x}}^{\text{TopS}} - u_{\text{x}}^{\text{BottomS}} = 0 \quad u_{\text{x}}^{\text{LeftS}} - u_{\text{x}}^{\text{RightS}} = 0, \tag{19}$$

where u_x^{RP} is a reference point with applied displacement in the x-direction. For the edges and vertices, the definition of the kinematic equations follows pre-determined paths along the boundaries of the SVE. The constraint relations between edges is separated into three lines which have been illustrated in Figure 18(a).

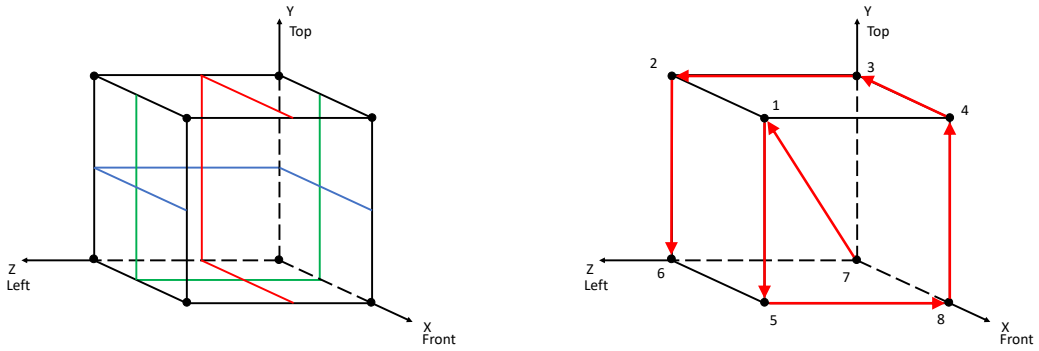

(a) Edge connection for definition of EasyPBC.

(b) Vertex connection for definition of EasyPBC

Figure 18: Illustration of how EasyPBC defines its kinematic equations of a RVE.

Following the red line, defining the relations between the edges produces the following set of equations:

$$u_\alpha^{\text{F.B.edge}} - u_\alpha^{\text{B.B.edge}} = \varepsilon_{x\alpha} l_x \quad u_\alpha^{\text{B.T.edge}} - u_\alpha^{\text{B.B.edge}} = \varepsilon_{y\alpha} l_y \quad u_\alpha^{\text{F.T.edge}} - u_\alpha^{\text{B.T.edge}} = \varepsilon_{x\alpha} l_x. \tag{20}$$

The final node set included the vertices (corners) of the RVE, which follow a pre-specified path similar to the edges, as seen in Figure 18(b), resulting in the following 7 constraint equations:

$$\begin{aligned}
u_\alpha^{\text{C7}} - u_\alpha^{\text{C1}} &= -\varepsilon_{x\alpha} l_x - \varepsilon_{z\alpha} l_z - \varepsilon_{y\alpha} l_y \\
u_\alpha^{\text{C1}} - u_\alpha^{\text{C5}} &= \varepsilon_{y\alpha} l_y \\
u_\alpha^{\text{C5}} - u_\alpha^{\text{C8}} &= \varepsilon_{z\alpha} l_z \\
u_\alpha^{\text{C8}} - u_\alpha^{\text{C4}} &= -\varepsilon_{y\alpha} l_y \\
u_\alpha^{\text{C4}} - u_\alpha^{\text{C3}} &= \varepsilon_{x\alpha} l_x \\
u_\alpha^{\text{C3}} - u_\alpha^{\text{C2}} &= -\varepsilon_{z\alpha} l_z \\
u_\alpha^{\text{C2}} - u_\alpha^{\text{C6}} &= -\varepsilon_{y\alpha} l_y.
\end{aligned} \tag{21}$$

4.2.3 Evaluation of elastic properties

The evaluation of the elastic properties stems from the stress-strain relation stated in Equation (4). The fundamental idea was to apply the strains individually (one at a time) while simultaneously allowing the nodes to move without restrictions in the transverse direction (transverse to active direction of the applied strain), thus removing the macroscopic transverse stress. The result of applying the strain systematically is that the diagonal components of the stiffness matrix can be evaluated separately, as seen in Equation (22), these being the resulting equations from Equation (5) with the removal of transverse stresses. The non-zero stress components were evaluated using Equation (7), simply being the total cross-sectional force divided by the area cross-sectional area of the SVE.

$$E = \frac{\text{axial stress}}{\text{axial strain}}, \quad E_x = \frac{\sigma_{xx}}{\varepsilon_{xx}}, \quad E_y = \frac{\sigma_{yy}}{\varepsilon_{yy}}, \quad E_z = \frac{\sigma_{zz}}{\varepsilon_{zz}}. \tag{22}$$

The same principle applies to Poisson's ratios, which were evaluated column-wise in the stiffness matrix. Therefore, by for example applying a strain in the x-direction (ε_{xx}), one can evaluate the material constants E_x, ν_{xy} and ν_{xz}. The definition of ν_{xy} can be seen in Equation (23), where the transverse strain was evaluated by knowing the transverse movement of the nodes.

$$\nu = \frac{\text{-transverse strain}}{\text{axial strain}}, \quad \nu_{xy} = \frac{-\varepsilon_{yy} E_x}{\sigma_{xx}} = \frac{-\varepsilon_{yy}}{\varepsilon_{xx}}. \tag{23}$$

The evaluation of the shear parameters was done through the same process, applying a shear strain and evaluating the respective modulus, the definition of which can be seen in Equation (24) for the case of G_{xy}.

$$G_{xy} = \frac{\sigma_{xy}}{2\varepsilon_{xy}} = \{2\varepsilon_{xy} = \gamma_{xy} = \varepsilon_{xy} + \varepsilon_{yx}\} = \frac{\sigma_{xy}}{\varepsilon_{xy} + \varepsilon_{yx}}. \tag{24}$$

4.2.4 Abaqus input file

The bridge between the geometrical structure in MATLAB and the finite element model in Abaqus is called an input file. The input file contains all information needed for Abaqus to create a model and run an analysis. This includes node numbering and positions, element types and node connectivity, material properties, boundary conditions, solution procedure, etc. Some parts of the input file are required for it to work, and some are not required but may reduce manual labour in the Abaqus CAE user interface. For this project, the input file was generated using MATLAB, primarily due to all data required to create the input file was already available in MATLAB after the generation of the geometric model. A pseudo-code of a typical input file produced in this project is presented in Listing 1.

The first step of writing the input file was to number all the nodes and store their positions based on the MATLAB generated voxel domain. Then the element connectivity (which nodes belong to which element) was established based on an arithmetic expression. The node usage was verified to ensure that there were no nodes that were used more times than they should. For example, in any cuboid RVE, there should always be eight nodes that are connected to only one element: the corner nodes. Similarly, there are simple expressions to calculate how many nodes should be connected to two, four and eight elements.

As the tapes are transverse isotropic, the material for the tapes was described using linear elastic `engineering constants` in Abaqus. This requires nine material properties (E_{11}, E_{22}, E_{33}, G_{12}, G_{13}, G_{23}, ν_{12}, ν_{13}, ν_{23}) and additionally a description of the material orientation. The in-plane material orientation was set using the known tape angle for each tape. The matrix material is isotropic and therefore described using linear elastic `isotropic` in Abaqus which only requires the stiffness and Poisson's ratio and no material orientation.

Listing 1: Pseudo-code for an Abaqus input file written in this project.

```
*HEADING

*NODE
node number, xpos, ypos, zpos

*ELEMENT, TYPE=C3D8I, ELSET=MATRIX
matrix element number, node1, node2, node3, node4, node5, node6, node7,
    node8

*ELEMENT, TYPE=C3D8I, ELSET=TAPE
tape element number, node1, node2, node3, node4, node5, node6, node7, node8

*SOLID SECTION, ELSET = MATRIX, MATERIAL=MATRIX_MATERIAL
*SOLID SECTION, ELSET=TAPE, MATERIAL=TAPE_MATERIAL, ORIENTATION=ORI

*MATERIAL, NAME=MATRIX_MATERIAL
*ELASTIC
E, nu

*MATERIAL, NAME=TAPE_MATERIAL
*ELASTIC, TYPE=ENGINEERING CONSTANTS
E11, E22, E33, nu12, nu13, nu23, G12, G13, G23

*DISTRIBUTION TABLE, NAME=TAB1
     COORD3D,COORD3D
*DISTRIBUTION, NAME=DIST1, LOCATION=ELEMENT, TABLE=TAB1
, X_axis_xpos_ref, X_axis_ypos_ref, X_axis_zpos_ref, XY_plane_xpos_ref,
    XY_plane_ypos_ref, XY_plane_zpos_ref
element number, X_axis_xpos, X_axis_ypos, X_axis_zpos, XY_plane_xpos_ref,
    XY_plane_xpos_ref, XY_plane_xpos_ref

*ORIENTATION, NAME=ORI, SYSTEM=RECTANGULAR
     DIST1
rotation_axis, additional_rotation
```

4.3 2D-model

Large 3D-models contain an immense amount of information, and therefore also need a lot of computational power to solve. Therefore, the size of the models is restricted by available computational power. Due to this, a methodology allowing for the conversion from 3D to 2D was introduced, inspired by the work of Selezneva *et al.* [16]. The methodology used by Selezneva *et al.* was briefly described in Subsection 2.2. The methodology used in this project is schematically visualised in Figure 19. Initially the 3D-domain was divided into partitions, or voxel columns through the thickness. The material properties and orientation of each voxel in the vertical columns are known, which allows for them to be treated as classical laminates. The effective in-plane material properties for each partition could then be evaluated using classical laminate theory, covered in Subsection 3.4.

In terms of conducting a finite element analysis, the methodology was practically identical to that of the 3D-model, with a small number of changes in terms of the input file and choice of elements. The input file no longer required a definition of z-position and the material type used was instead `lamina`, the plane stress equivalent of engineering constants. Besides a material orientation, the `lamina` material definition requires six elastic material parameters: E_{11}, E_{22}, ν_{12}, G_{12}, G_{13}, and G_{23}. The S4R element was chosen for the 2D-model as it is a reduced integration general-purpose element. General-purpose meaning it is suitable for both thick and thin plates. Reduced integration was chosen over full integration as it should reduce the solution time which was one of the core reasons for creating the 2D-model.

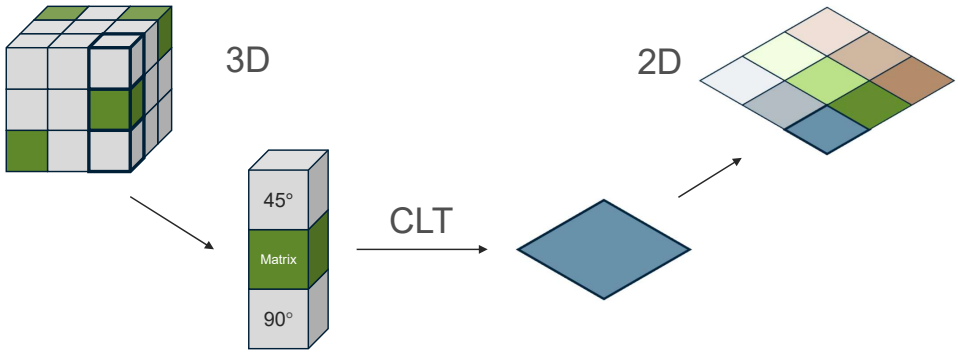

Figure 19: Schematic visualisation of analytical through-thickness homogenisation. Different colours represent different materials or material properties.

4.4 Design of experiment

The experiment selected to be replicated was conducted by Li *et al.* [4], summarised in Section 2.1. In particular, focus was on the thin tape plate they manufactured and tested. This study was selected due to the detailed reporting of methodology, results and input parameters. As discussed in Subsection 4.1.2, the number of elements required to discretise a tape depends on the tape thickness. Therefore, thin tape models have a lower maximum SVE size than thick tape models. The thin tape ($t_t = 0.164$ mm) plates studied by Li *et al.* were considered to have a suitable tape thickness to model sufficiently large SVEs.

The homogenisation procedure was conducted for three different sizes of SVEs, with different number of samples for each SVE size, listed in Table 4, to enable a comparison of results. In addition, for every SVE generated, the homogenised material properties were evaluated using both the full 3D-model and the reduced 2D-model. All SVEs were cut out from the centre of larger models to remove edge effects. Also, the equivalent of four tape thicknesses were removed from the top of each SVE, as if they had been sanded. The top layers of the numerical model typically had a lower fibre volume fraction and by cutting off the top, the global plate fibre volume fraction became more representative of the real material.

The SVE sizes were based on tape size and available computational resources. All SVEs had the height 28 voxels (2.24 mm) and were square in-plane. The medium (110 voxel side length) SVE was chosen as it was the maximum feasible size to run on a desktop computer, and the side length is almost equal to one tape length. From there, the small (55 voxel side length) SVE was selected simply as half the size of the medium SVE, with twice as many samples. Finally, the large size (250 voxel side length) SVE was chosen as it is exactly twice as long as a single tape and as it was reasonably fast to run on the Chalmers Centre for Computational Science and Engineering (C3SE) computer cluster [31].

Table 4: SVE size, original domain size, and number of samples for the three homogenisation cases.

SVE size (voxels, tape dimensions)	**Original plate size (voxels)**	**Number of samples**
55x55x28 $\approx \frac{l^t}{2} \times \frac{l^t}{2} \times 14t^t$	500x500x36	250
110x110x28 $\approx l^t \times l^t \times 14t^t$	500x500x36	125
250x250x28 $= 2l^t \times 2l^t \times 14t^t$	1000x1000x36	10

Although the material data list and input parameters used by Li *et al.* [4] was extensive, it was not complete. Li *et al.* made a numerical model of the material in a connected study [2], summarised in Subsection 2.2, where the input data was even more detailed. Shah *et al.* [17] made a numerical model trying to replicate the material produced by Li *et al.* [4]. The article written by Shah *et al.* had values for the matrix material properties, which Li *et al.* excluded in both [2] and [4]. Even with three articles using the same material, the list of input material data was not complete and some parameters had to be assumed or derived. E_{33}, G_{13} and ν_{13} were directly derived based on the transverse isotropy assumption of the tapes as: $E_{33} = E_{22}$, $G_{13} = G_{12}$ and $\nu_{13} = \nu_{12}$. The Poisson's ratio ν_{23} was assumed to be 0.4, since ν_{23} typically lies between 0.4 and 0.45 [32]. Then the shear modulus G_{23} was calculated as:

$$G_{23} = \frac{E_{22}}{2(1+\nu_{23})}. \tag{25}$$

Table 5 shows the input parameters for tape size and material data used in the model. The dimensions of the tapes used in the FE model were not exactly the same as the dimensions used by Li *et al.* [4]. This was a result of the discretisation, using in-plane square elements. The dimensions of the elements was based on the tape thickness (t^t), the number of elements in the thickness direction ($N^t_{\text{vox},h}$) and the voxel aspect ratio (R_{vox}) as:

$$t_{\text{vox}} = \frac{t^t}{N^t_{\text{vox},h}}, \quad l_{\text{vox}} = w_{\text{vox}} = \frac{t^t R_{\text{vox}}}{N^t_{\text{vox},h}}, \tag{26}$$

where t_{vox} is the element thickness and $l_{\text{vox}} = w_{\text{vox}}$ are the element in-plane side lengths. For the modelled tapes to have the exact same dimensions as the real tapes, the following conditions must be fulfilled:

$$l^t = x_1 l_{\text{vox}}, \quad w^t = x_2 w_{\text{vox}}, \quad t^t = x_3 t_{\text{vox}}, \tag{27}$$

where l^t, w^t, and t^t are the tape length, width, and thickness, respectively, and x_1, x_2, and x_3 are integer values. If x_1, x_2, and x_3 are not integer values, the modelled tapes cannot have the exact same dimensions as the physical tapes and at least one of the dimensions must deviate. In this case, the thickness could be reduced by 0.004 mm for the width and length of the modelled tapes to coincide with the real tapes.

Table 5: Constant input parameters for tape and polymer matrix. † - values taken from [2], †† - values taken from [17], ⋆ - assumed values typical for carbon fibre composites, †⋆ - values assumed by [2].

Tape size		**Moduli [GPa]**			**Shear moduli [GPa]**			**Poisson's ratio**			**Matrix**	
[voxels]	[mm]	$E^{\dagger}_{11}$	$E^{\dagger\star}_{22}$	E_{33}	$G^{\dagger\star}_{12}$	G_{13}	G_{23}	$\nu^{\dagger\star}_{12}$	ν_{13}	$\nu^{\star}_{23}$	$E^{\dagger\dagger}$ [GPa]	$\nu^{\dagger\dagger}$
125x20x2	50x8x0.16	140	9	9	5.6	5.6	3.2	0.34	0.34	0.4	3.45	0.31

5 Results

This section presents the results produced in this project. As a starting point, one of the many generated plate samples, and the results of taking it from a full plate model to an SVE is presented. A meso- and microstructural comparison to a real material sample is made, followed by an overview of the variation in average tape angle and fibre volume fraction across all samples. The section moves on to show the predicted material properties, averaged for all SVEs, followed by a presentation of how the predicted mechanical properties depend on the geometrical properties of each individual SVE. The section is finalised with a numerical experiment where the validity of using homogenised STBDC material properties in a practical implementation was tested.

5.1 Geometrical properties of the numerical STBDC model

The following subsections present the results of the developed geometrical model. In the first part, the visual appearance of the generated structure is illustrated, including how the geometry is modified in order to generate a usable SVE domain. Using the presented geometrical structure, images showcasing the stochastic nature of the material are presented for different geometrical variables such as fibre volume fraction and average tape angle.

A top-view image of a generated geometrical model is presented in Figure 20, the colour scheme is set so that white represents tapes aligned with the x-axis and black represents tapes aligned with the y-axis. The full domain consists of two subdomains, the safe space and the SVE itself. The safe space subdomain was initially introduced to avoid programming related issues (see Subsection 4.1.1) but it also works as a border that removes most of the edge effects. The SVE in Figure 20 has the in-plane dimensions 110x110 voxels, the medium size of the three SVEs that were analysed in this study.

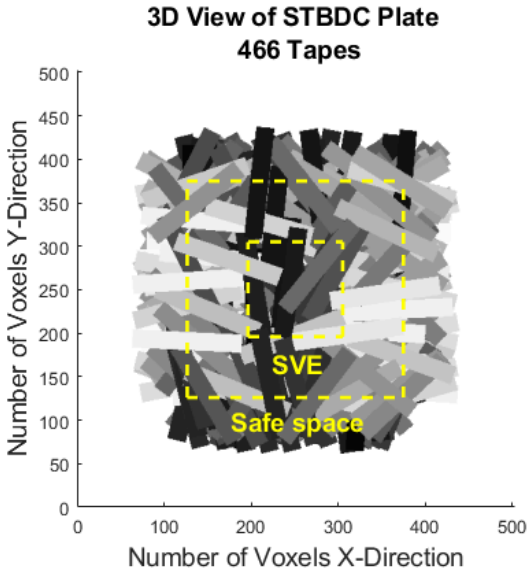

Figure 20: top view of a 3D generated material structure with subdomains visualised.

5.1.1 Results of domain modifications

The result of cutting the full domain along the safe space border and removing the top layers, as explained in Subsection 4.4, is presented in Figures 21-23, showcasing how the domain changes and also how the fibre volume fraction and in-plane tape angle are affected. Removing resin rich top layers does not only increase the FVF but it also cuts some of the tapes, resulting in partial tapes and tapes with holes in the top layers, displayed in Figure 21(b).

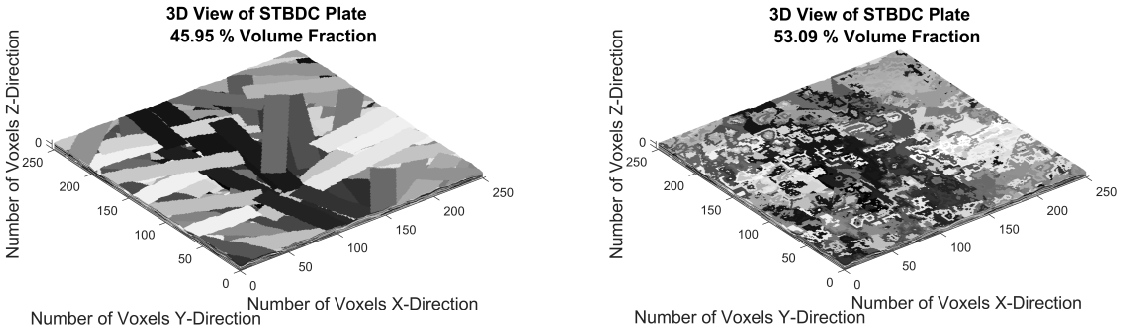

(a) 3D ISO-view of a 250x250x36 STBDC.

(b) 3D ISO-view of a 250x250x28 STBDC.

Figure 21: Visualisation of a computationally generated material structure, (a) unmodified and (b) modified by having top layers removed.

The variation of the through-thickness fibre volume fraction for the stated domain has been presented in Figure 22 with (a) being 36 voxels in height (not modified) and (b) being 28 voxels height-wise (modified). For the unmodified domain the variation of the FVF is evident, with high levels of variation. This variation of FVF is to be expected for these types of materials, although, generally less intense. By removing the top layers, the fibre volume faction increases and becomes more uniform, this is evident when comparing Figure 22(b) to Figure 22(a). Comparing the two figures, it can be also noted that the fibre volume fraction has increased from 45.95% to 53.09%, bringing it closer to the experimental target value of 57.85%.

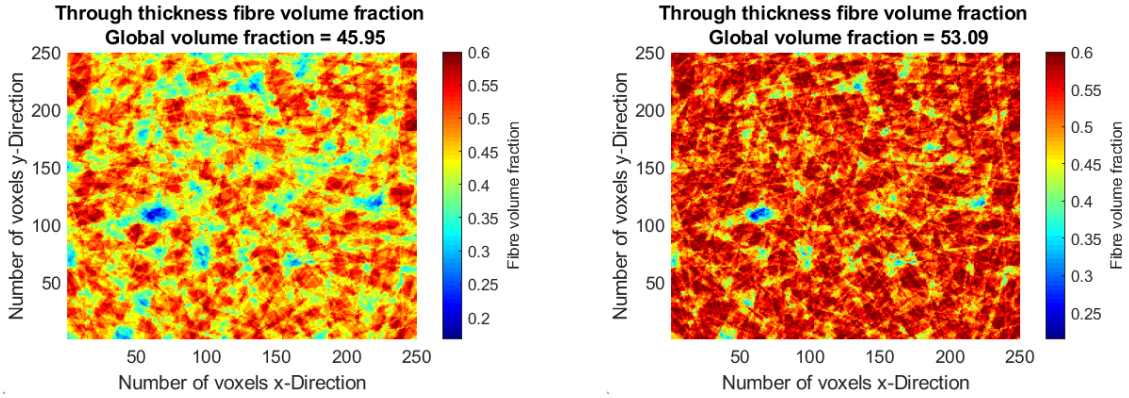

(a) FVF for a 250x250x36 voxel plate.

(b) FVF for a 250x250x28 voxel plate.

Figure 22: Visualisation of through thickness fibre volume fraction distribution, (a) unmodified and (b) modified by having top layers removed.

The change in terms of the average in-plane tape angle across the domain has been presented in Figure 23, for the same domain as the FVF. The difference is hardly noticeable between the two, with the modified domain receiving a slight increase in terms of the global average (for this specific case). The global average increasing is however not always the outcome. The deviation from the original global average tape angle could go either way depending on the layout of the top layers, unlike for the FVF where an increase is almost always the case.

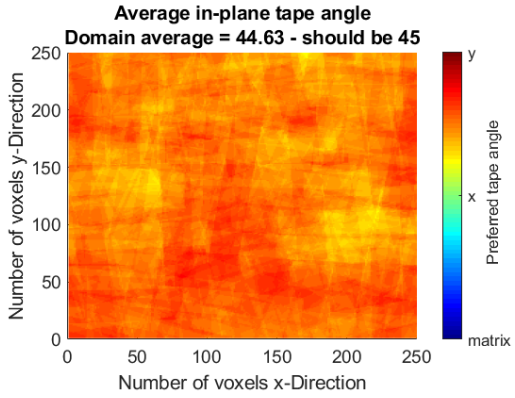

(a) Average in-plane tape angle for a 250x250x36 voxel plate.

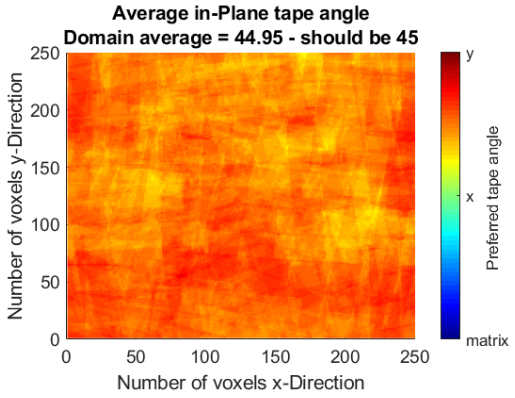

(b) Average in-plane tape angle for a 250x250x28 voxel plate.

Figure 23: Visualisation of through thickness average tape angle distribution, (a) unmodified and (b) modified by having top layers removed.

5.1.2 Mesostructure

The stochastic distribution of the tapes results in the randomised visual appearance that characterise these types of materials. In order to visually examine if the developed algorithm generates a geometry that replicates the material at hand, a visual comparison has been presented in Figure 24. The colours in Figure 24(b) are based on the orientation of the tapes, with black indicating that it is aligned with the y-axis (vertical) and white being aligned with the x-axis (horizontal). The white dots in Figure 24(a) are MATLAB-generated, marking out the boundaries of the tapes, and the colours themselves have no correlation to a geometric property.

Comparing the real and computer-generated mesostructure, the visual appearance is similar between the two, indicating that the generated model is an accurate representation of the real material. The stochastic distribution of the tapes has therefore, been captured by the developed algorithm.

(a) Real material mesostructure [4].

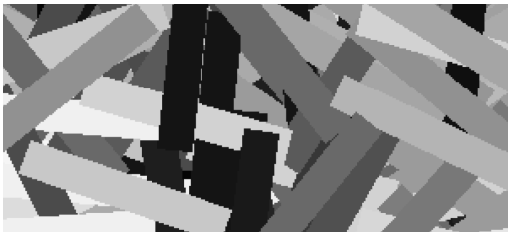

(b) Numerical model mesostructure.

Figure 24: Top-view of a STBDC structure, (a) real material sample, (b) generated voxel-based material sample.

5.1.3 Microstructure

A method of verifying the quality of the generated geometrical model is to compare the computer-generated microstructure to a real material sample microstructure. A microscope image of a real material sample has been presented in Figure 25, where the alignment of the tapes can be distinguished by the their visual appearance and colour. The ones that appear to be constructed out of dots (small circles), are oriented in the transverse direction (normal to the cross-section), and tapes that appear more as lines are aligned along the cross-section.

Microstructure images of the generated geometrical model are presented in Figures 26 and 27, where the red colour represents matrix, and the grey-scale is dependent on the orientation of the tape. If a tape is aligned with the x-direction (aligned with the cross-section), it would appear as white in the figure. The tapes that are instead transverse (aligned in the y-direction, normal to the cross-section), appear as black, and the tapes within these limits receive their respective colour according to the grey colour-spectrum as seen in the colour bars. The difference between Figure 26 and 27 is that the latter is modified with the top layers removed.

Comparing Figure 25 to Figure 26, the generated model does capture the waviness of the structure, although less smooth as a consequence of the voxels and relatively low resolution. The draping ratio of 5:1, meaning that the voxel dimensions are five times larger in-plane than out-of-plane, does appear to provide a realistic representation of the real material. The main issue is the high matrix content at the top, which is not seen in the real plate, as a result the height was modified to remove these layers, resulting in Figure 27.

The removal of 8 voxel-layers (4 tape-layers) from the top resulted in a more accurate representation, as can be seen by comparing Figure 27 to 25. The real material does not contain the high-resin content region at the top; thus this modification appears to be a valid method to generate a more accurate representation of the STBDC microstructure.

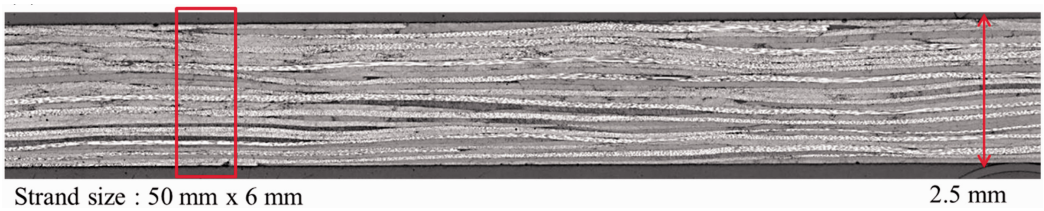

Figure 25: Microscopy image of the cross-section of a manufactured STBDC [16].

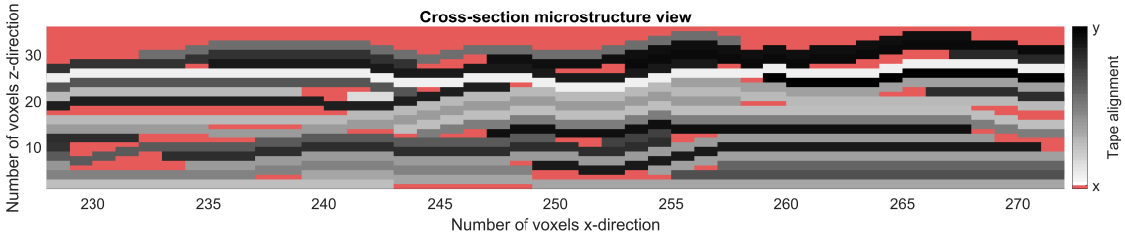

Figure 26: Cross-section image of a computer-generated STBDC, 36 voxels in height.

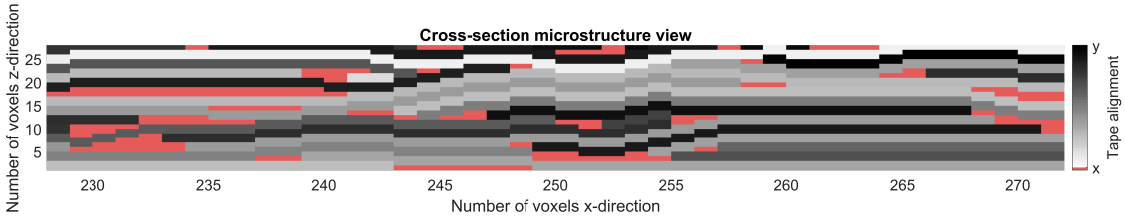

Figure 27: Cross-section image of a computer-generated STBDC, modified from 36 to 28 voxels in height.

5.1.4 Distribution of fibre volume fraction and average tape angle

The distributions of (a) fibre volume fraction and (b) average tape angle are presented as histograms with normal distribution curves in Figures 28, 29, and 30, for the small, medium, and large sized SVEs, respectively. A nearly normal distribution of both average tape angle and fibre volume fraction can be observed for the small SVE with 250 samples, with one outlier in fibre volume fraction. As the SVE size is increased and the sample size reduced, the histograms do not fit the normal distribution curve as well. Further, as the domain increases in size, both the global fibre volume fraction and the average tape angle converge in the sense that the differences between the largest and smallest values decrease. It can be concluded that there is a larger variation in average tape angle than fibre volume fraction for all SVE sizes. Finally, it should be noted that the sample sizes are, at least for the large SVE, too small to draw statistically sound conclusions.

The reference material, made by Li *et al.* [4], did not present results for the average tape angle, but the authors did present results regarding the fibre volume fraction of their plates. The average FVF of the reference material was 57.85%, which means that the small, medium, and large SVEs have an average fibre volume fraction that is 9.1%, 9.5% and 9.7% below the reference average, respectively.

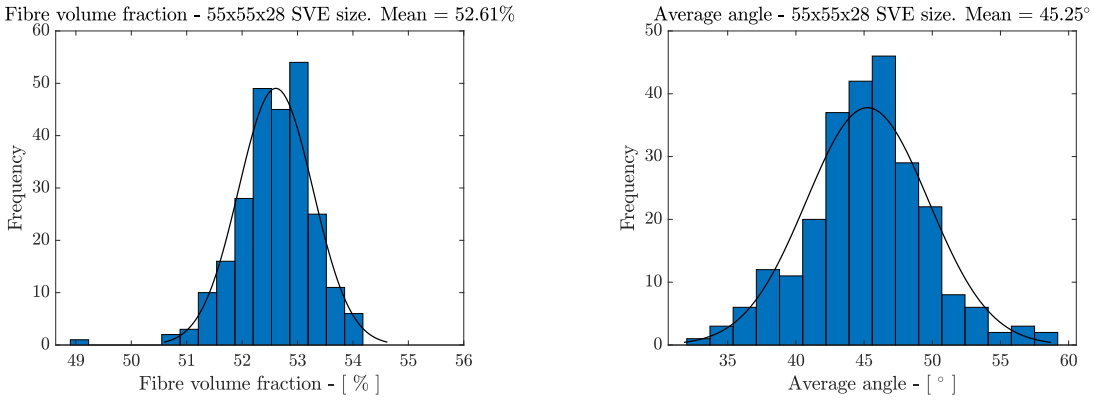

(a) Plate fibre volume fraction distribution.

(b) Average angle distribution.

Figure 28: Histogram with normal distribution curve of (a) fibre volume fraction distribution, and (b) average angle distribution, for the 55x55x28 voxel SVEs.

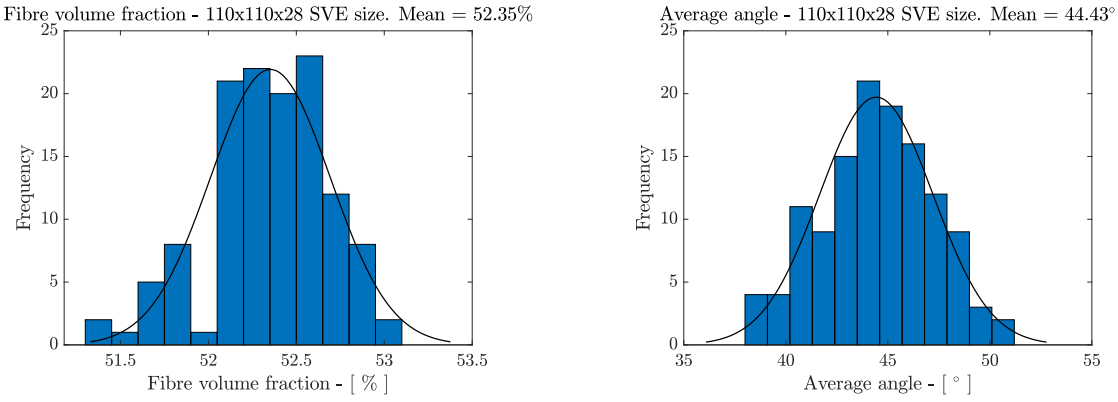

(a) Plate fibre volume fraction distribution.

(b) Average angle distribution.

Figure 29: Histogram with normal distribution curve of (a) fibre volume fraction distribution, and (b) average angle distribution, for the 110x110x28 voxel SVEs.

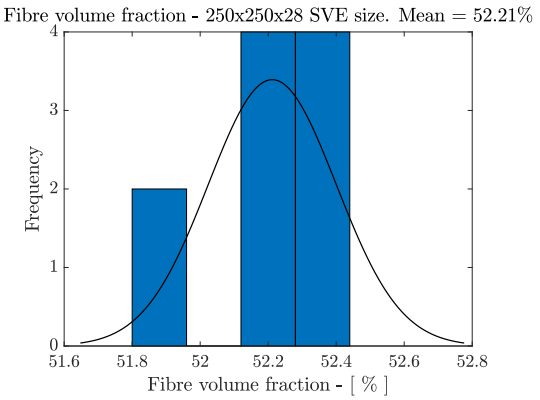

(a) Plate fibre volume fraction distribution.

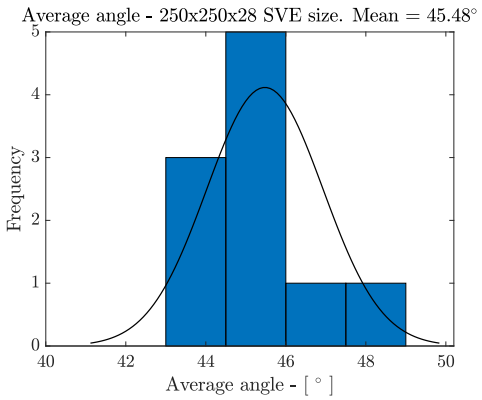

(b) Average angle distribution.

Figure 30: Histogram with normal distribution curve of (a) fibre volume fraction distribution, and (b) average angle distribution, for the 250x250x28 voxel SVEs.

5.2 Homogenised material properties

The homogenised in-plane material properties were computed for the three different SVE sizes, with different number of samples for the different sizes. Smaller SVEs statistically have a larger variability whereas larger SVEs have less variability. Therefore, a larger number of realisations are required for smaller models to compensate for the larger variability. The average (for all samples) homogenised results are summarised in Tables 6 and 7 for the full 3D-model and the reduced 2D-model, respectively. In these tables, most values are written in the form $\mu \pm S$ where μ is the mean and S is the standard deviation, calculated as

$$\mu = \frac{1}{N}\sum_{i=1}^{N} a_i, \qquad S = \sqrt{\frac{1}{N-1}\sum_{i=1}^{N} |a_i - \mu|^2}, \tag{28}$$

where N is the number of observations and a_i the observed variable. In the tables, also the solution time has been included in terms of core hours, calculated as the number of cores used multiplied by the time Abaqus required to solve the problem. The solution times are only given as approximate values, mainly because they were calculated based on a reduced sample size, but also since it was seen that the solution time could vary depending on if other processes were running on the computer during solution process.

It should be noted in Tables 6 and 7 that the standard deviation decreases with increasing SVE size for all measured properties, even though the sample size is decreased. Another observation to point out is that the standard deviation is consistently larger for the reduced 2D-model than for the full 3D-model, even though these were made for identical SVEs. Further, the reduced 2D-model consistently produces slightly lower average values of all material properties. The largest difference between the two models regarding stiffness appears for the large SVEs, where E_{11}^{2D} is 2.13% lower than E_{11}^{3D}. For the shear modulus, this percentual difference increases and the maximum difference is found for the medium size SVEs, where G_{12}^{2D} is 3.96% lower than G_{12}^{3D}. Further, the Poisson's ratios are also consistently lower for the 2D-model, the largest difference occurs for the medium sized SVEs, where ν_{12}^{2D} is 4.47% lower than ν_{12}^{3D}. Finally, the total solution time, i.e. the number of CPU hours used to run all samples, varies drastically between the different cases. A single large 3D-model homogenisation process requires more CPU hours than all the small 3D-models combined. This effect is not at all as evident for the reduced 2D-models.

Table 6: Averaged homogenised results - full 3D-model.

SVE size [Voxels]	**55x55x28**	**110x110x28**	**250x250x28**
Number of samples (N)	250	125	10
$\boldsymbol{E}_{11}$ **[GPa]**	40.19 ± 6.36	41.31 ± 4.06	39.77 ± 2.41
$\boldsymbol{E}_{22}$ **[GPa]**	40.83 ± 6.20	40.01 ± 3.98	40.95 ± 1.60
$\boldsymbol{G}_{12}$ **[GPa]**	15.25 ± 1.43	14.95 ± 0.85	15.36 ± 0.41
$\boldsymbol{\nu}_{12}$ **[-]**	0.308 ± 0.062	0.301 ± 0.040	0.290 ± 0.013
$\boldsymbol{\nu}_{13}$ **[-]**	0.305 ± 0.024	0.308 ± 0.016	0.312 ± 0.005
$\boldsymbol{\nu}_{21}$ **[-]**	0.313 ± 0.061	0.291 ± 0.036	0.299 ± 0.023
$\boldsymbol{\nu}_{23}$ **[-]**	0.303 ± 0.023	0.312 ± 0.014	0.309 ± 0.009
Average tape angle [degrees]	45.25 ± 4.49	44.43 ± 2.78	45.48 ± 1.45
Domain FVF [%]	52.61 ± 0.67	52.35 ± 0.34	52.21 ± 0.19
Solution time [CPUs × h]	≈ 0.7	≈ 4.2	≈ 246.1
Total time [CPUs × h × N]	175	525	2461

Table 7: Averaged homogenised results - reduced 2D-model.

SVE size	**55x55x28**	**110x110x28**	**250x250x28**
Number of samples (N)	250	125	10
$\boldsymbol{E}_{11}$ **[GPa]**	39.84 ± 6.78	40.59 ± 4.21	38.92 ± 2.63
$\boldsymbol{E}_{22}$ **[GPa]**	40.55 ± 6.68	39.42 ± 4.30	40.17 ± 1.69
$\boldsymbol{G}_{12}$ **[GPa]**	15.02 ± 1.64	14.36 ± 0.92	14.76 ± 0.39
$\boldsymbol{\nu}_{12}$ **[-]**	0.295 ± 0.066	0.287 ± 0.041	0.282 ± 0.015
$\boldsymbol{\nu}_{21}$ **[-]**	0.300 ± 0.065	0.279 ± 0.040	0.292 ± 0.025
Average tape angle [degrees]	45.25 ± 4.49	44.43 ± 2.78	45.48 ± 1.45
Domain FVF [%]	52.61 ± 0.67	52.35 ± 0.34	52.21 ± 0.19
Solution time [CPUs × h]	≈ 0.2	≈ 0.5	≈ 2.7
Total time [CPUs × h × N]	50	62.5	27

A visual comparison between the numerically predicted Young's moduli E_{11} and E_{22}, and the experimental data is presented as bar charts in Figure 31. The reference interval is given as the upper and lower bounds based on experimental data ($E = 43.41 \pm 4.14$ GPa) [4]. The numerically predicted E_{11}, shown in Figure 31(a), is on average 7.62% lower than the reference mean but 2.12% above the lower bound. Similarly, the Young's modulus E_{22}, shown in Figure 31(b), is on average 7.11% lower than the reference mean but 2.68% larger than the lower bound, with no samples falling outside the reference interval. E_{11} based on the average of the large 2D-SVEs is the only comparable homogenised value that falls outside the lower bound by 0.88% (rightmost bar in Figure 31(a)).

The shear modulus G_{12} and the Poisson's ratios ν_{12} and ν_{21}, plotted in Figure 32, were not measured in the experimental study used for comparison and therefore, these figures have no reference interval to compare with. What can be noted in these figures is the fact that the 2D-models consistently produce values of lower magnitude, although, the percentual difference is small.

In both Figure 31 and 32, the earlier note that the standard deviation decreases with increasing SVE size is clearly visible. The second note that the standard deviation of the homogenised properties is increased for the 2D-models compared to the 3D-models is not as visible, since the difference is comparatively small.

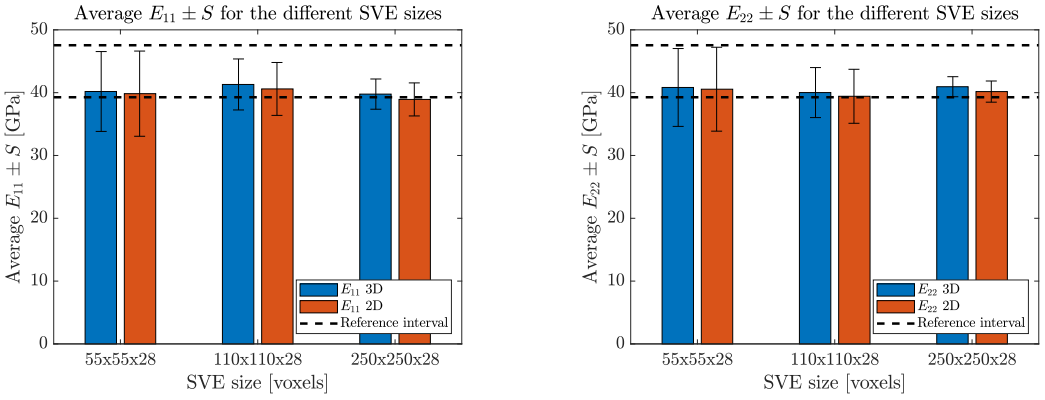

(a) Average $E_{11} \pm S$ for the different SVE sizes.

(b) Average $E_{22} \pm S$ for the different SVE sizes.

Figure 31: Bar charts of (a) average $E_{11} \pm S$ and (b) average $E_{22} \pm S$ for all SVE sizes, both 2D and 3D.

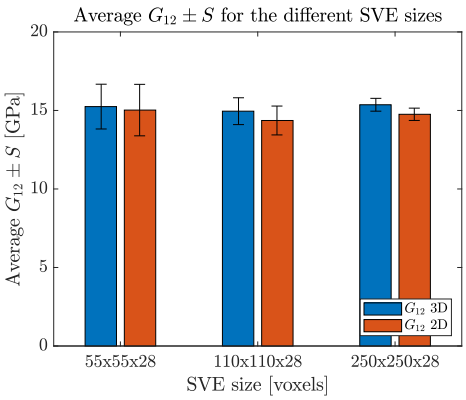

(a) Average $G_{12} \pm S$ for the different SVE sizes.

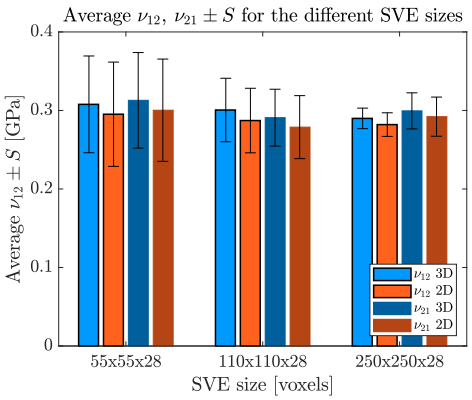

(b) Average ν_{12}, $\nu_{21} \pm S$ for the different SVE sizes.

Figure 32: Bar charts of (a) average $G_{12} \pm S$ and (b) average ν_{12}, $\nu_{21} \pm S$ for all SVE sizes, both 2D and 3D.

5.2.1 Convergence study

The use of different SVE sizes with a specific number of realisations for each, naturally introduces a convergence study in order to see if there is a notable improvement as the size of the domain is increased. Therefore, in Figures 33-35, the convergence for the different SVE sizes has been presented as a cumulative average (CA). The evaluation of the cumulative average was initiated by having all data sets in chronological order as the data was gathered. The first point of the cumulative average was then evaluated using only the first data point, the next average included the first and second value of the data set and so on until the last data point has been reached. The formula for evaluating the CA for a given a data set is defined as:

$$\mu_n^{\text{CA}} = \frac{1}{n} \sum_{i=1}^{n} a_i, \quad n = 1, 2, 3, ...N, \tag{29}$$

where μ_n^{CA} is the cumulative average for data point n, a_i is the observed variable, and N is the number of observations.

In Figures 33(a)-35(a), the CA has been evaluated for E_{11} and E_{22}, using both the 3D- and the 2D-model for each of the SVE sizes, respectively. The primary observation made is that the values for both moduli converge and result in a smaller deviation between each other as the number of data points accounted for increases. A noteworthy observation is that E_{22} is only smaller than E_{11} for the 110x110x28 domain. For the other two SVE sizes, E_{22} converges to larger values than E_{11}. The reason for this appears to be related to the global average tape angle for the entire data set, meaning that if the average angle across all data points is below 45°, E_{11} will be larger and if its above 45°, E_{22} becomes the larger of the two. There is no bias in the generation of the structure that should generate a consistent difference between the two, meaning that either of the two can be larger or smaller than the other after converging. However, in a perfect scenario where the global average tape angle converges to 45°, the Young's moduli E_{11} and E_{22} would likely converge to the same value. In this case, there is a difference between the converged E_{11} and E_{22}, indicating that the geometric representation is not perfect although, generates results within reasonable margins. The maximum difference between E_{11} and E_{22} appears for the large 2D-model where E_{22} is 3.2% larger than E_{11}. The largest difference between 2D- and 3D-models also appear for the large SVE, where E_{11}^{2D} is 2.1% lower than E_{11}^{3D}

In Figures 33(b)-35(b), the cumulative average has been evaluated for all effective material properties resulting from the homogenisation and normalised against their respective final value, wherefore all converge to unity. This has also been done for both 2D and 3D, indicating how many realisations are required for the properties to converge within a reasonable error margin. It can be observed that the cumulative average is well within 5% margin of error compared to the final value at around 100 realisations for the 55x55x28 voxel SVE, 60 for 110x110x28, and 9 for 250x250x28. Although the data

set for the large SVE does not contain a large number of data points, it appears to converge significantly faster than the smaller SVEs.

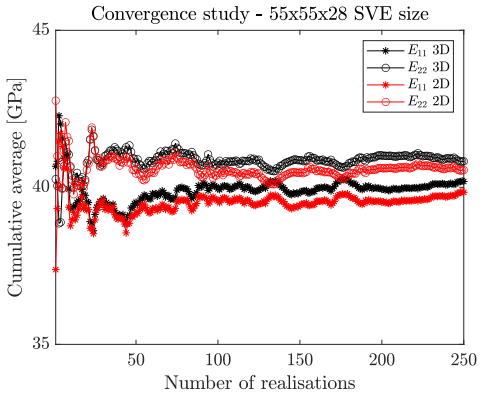

(a) Cumulative average for E_{11} and E_{22} for the 2D and the 3D-model.

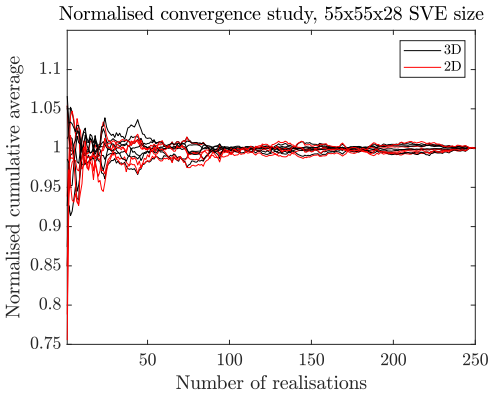

(b) Normalised cumulative average for all homogenised material properties for the 2D- and the 3D-model.

Figure 33: Convergence study for the 55x55x28 voxel SVE, (a) cumulative average for Young's moduli, (b) normalised cumulative average for all material properties.

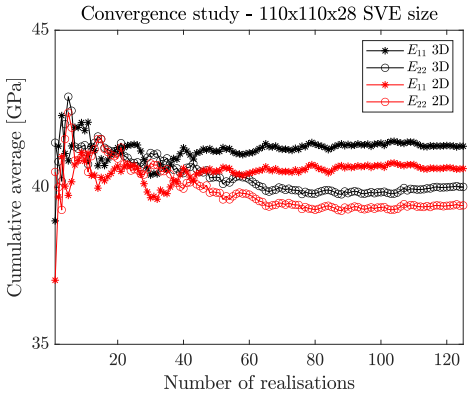

(a) Cumulative average for E_{11} and E_{22} for the 2D and the 3D-model.

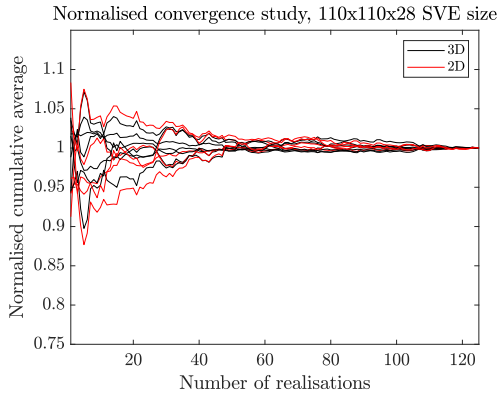

(b) Normalised cumulative average for all homogenised material properties for the 2D- and the 3D-model.

Figure 34: Convergence study for the 110x110x28 voxel SVE, (a) cumulative average for Young's moduli, (b) normalised cumulative average for all material properties.

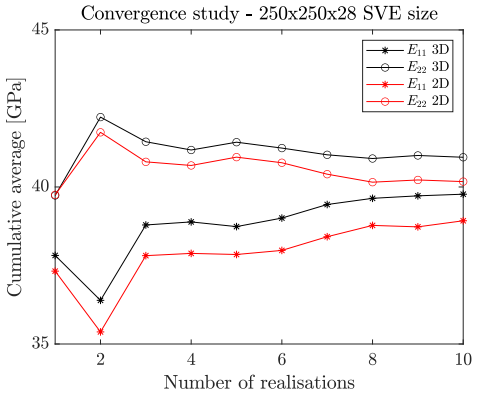

(a) Cumulative average for E_{11} and E_{22} for the 2D and the 3D-model.

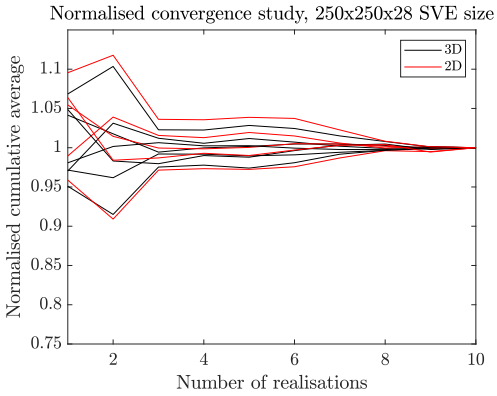

(b) Normalised cumulative average for all homogenised material properties for the 2D- and the 3D-model.

Figure 35: Convergence study for the 250x250x28 voxel SVE, (a) cumulative average for Young's moduli, (b) normalised cumulative average for all material properties.

5.2.2 Stress and strain response

Analysing the stress or strain distributions in the SVEs was not a focus in this project. However, these are results of the homogenisation process and therefore included. The top and bottom views of a medium-sized SVE, displayed in the Abaqus CAE graphical user interface, is shown in Figures 36 and 37, where (a) shows the undeformed state and (b) shows the ε_{11} strain distribution in a deformed state with an equivalent applied strain $\varepsilon_{11} = 0.2$. The effect of removing the top layers is clearly visible when comparing the two figures. The bottom side of the undeformed plate, shown in Figure 37(a) has clearly separated tapes, distinguishable from each other. On the contrary, the top side of the undeformed plate, shown in Figure 36(a), has no distinguishable tapes. It is an irregular mix of tapes and matrix. The same effect is visible in the deformed state, shown in Figures 36(b) and 37(b). The strain magnitude appears to vary between different tapes in Figure 37(b), the reason for this apparent variation is that the ε_{11} strain is plotted based on the local tape orientation coordinate system, instead of the global coordinate system. On the top-side instead, the strain only seem to vary between tapes and matrix, which is reasonable since the tapes are between 2.6 and 39 times stiffer than the matrix material, depending on the tape orientation.

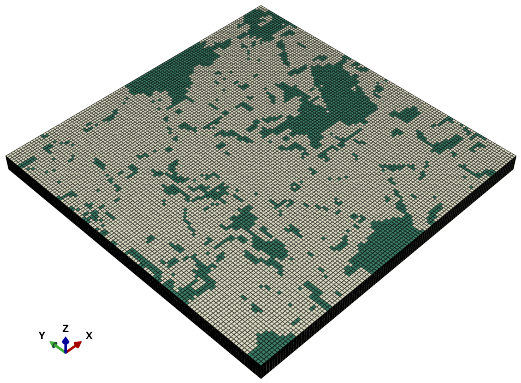

(a) Undeformed STBDC plate top-view. White - tape, green - matrix.

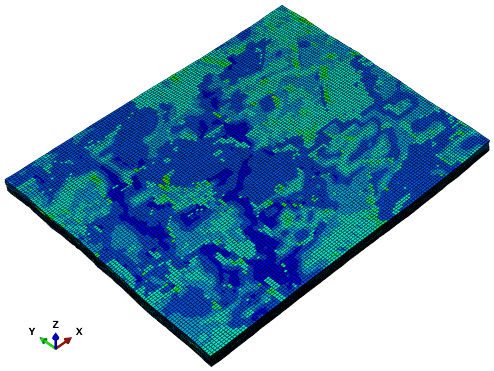

(b) Deformed STBDC plate top-view. ε_{11} strain distribution.

Figure 36: Top-view of an STBDC plate, (a) undeformed, (b) deformed with applied strain $\varepsilon_{11} = 0.2$.

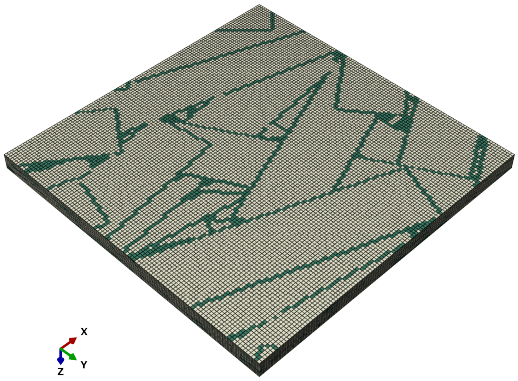

(a) Undeformed STBDC plate bottom-view. White - tape, green - matrix.

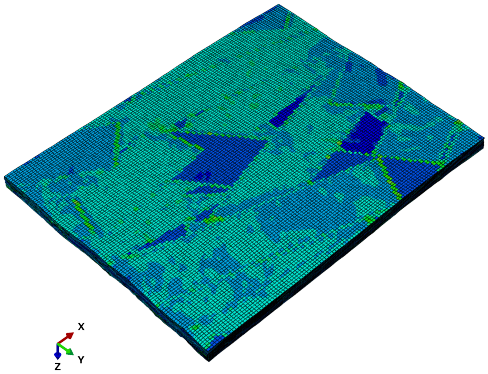

(b) Deformed STBDC plate bottom-view. ε_{11} strain distribution.

Figure 37: Bottom-view of an STBDC plate, (a) undeformed, (b) deformed with applied strain $\varepsilon_{11} = 0.2$.

The σ_{11} stress distribution in the deformed plate is displayed in Figure 38, where (a) shows the bottom view and (b) shows the top view. Note that the stress distribution, like the strain distribution, is based on the local tape coordinate system, not the global coordinate system. The stress distribution seems to follow the strain distribution to a large extent, which is expected. Areas with large strain result in areas with high stress, except for where the stiffness is low. Another similarity to the strain distribution is that, also in this case, the individual tapes are distinguishable from each other on the bottom side but not on the top side. Although, the stress distribution seems to have lower gradients (less abrupt distinctions between material boundaries) than the strain distributions. Therefore, the material boundaries are less pronounced in the stress plots than in the strain plots.

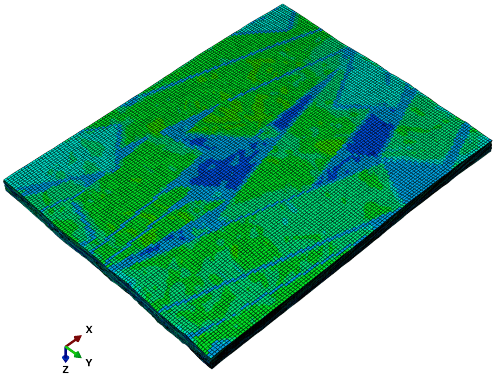

(a) Deformed STBDC plate bottom-view. σ_{11} stress distribution.

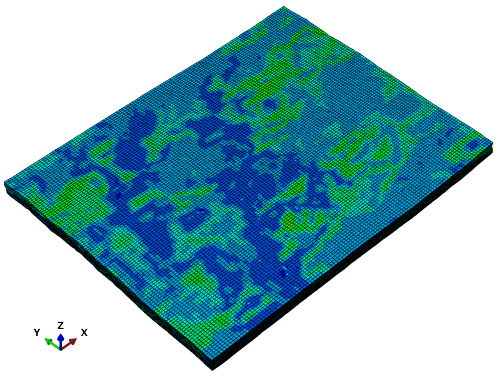

(b) Deformed STBDC plate top-view. σ_{11} stress distribution.

Figure 38: Stress distribution in an STBDC plate subjected to an applied strain ε_{11}=0.2. (a) Bottom-view, (b) top-view.

5.3 Material properties versus geometrical properties

The distribution, or rather change of FVF and average tape angle across different realisations for the same size of SVE (presented in Subsection 5.1.4) raises the question: how do these variations affect the effective material parameters? The relation between the average tape angle and the two Young's moduli E_{11} and E_{22} has been presented in Figures 39-41 for the different SVE sizes, respectively. Similarly, the relation between the average modulus $\bar{E}$ and fibre volume fraction has been presented in Figures 42(b), 43(b), and 44(b), with $\bar{E}$ being defined as:

$$\bar{E} = \frac{E_{11} + E_{22}}{2}. \tag{30}$$

In all these figures, the colour scheme is black and red where black represents results for the full 3D-model and red represents results for the reduced 2D-model.

A linear regression curve has been drawn in these figures, using the `polyfit` function in MATLAB, which fits the best curve to two (x, y) data sets using a least squares procedure. To evaluate the strength of the correlation, the Pearson correlation coefficient was used, through the function `corrcoef` in MATLAB, and it is included in the legend of each figure. The Pearson correlation coefficient (r) for two data sets a_i and b_i with N observations and mean μ_a and μ_b, respectively, is calculated as

$$r = \frac{\sum_{i=1}^{N}(a_i - \mu_a)(b_i - \mu_b)}{\sqrt{\sum_{i=1}^{N}(a_i - \mu_a)^2 \sum_{i=1}^{N}(b_i - \mu_b)^2}} \tag{31}$$

and results in values in the range $r \in [-1, 1]$ [33]. A Pearson coefficient of 1 or -1 implies a perfect correlation, that the data points lie exactly on a straight line with a positive or negative slope, respectively. A value of zero or close to zero implies that there is no correlation between the two data sets.

5.3.1 Young's modulus versus global average tape angle

The relation between the global average tape angle and stiffness in different directions, shown in Figures 39-41, give a clear indication that, as expected, the stiffness depends on the average tape angle. Included in these figures are also the reference stiffness interval based on experimental data [4]. There is a strong correlation between stiffness and average tape angle, $|r| > 0.8$ for all cases. It should be remembered that the way that the global average tape angle is calculated (explained in Subsection 4.1.4), it should converge to 45° for a completely random distribution, with values below 45° showing a tendency to the x-direction and values above 45° a tendency to the y-direction. Figures 39(a), 40(a), and 41(a) show that the stiffness in the x-direction (11) increases when tapes are aligned with the x-axis and that it decreases when tapes are aligned with the y-axis. Similarly, Figures 39(b), 40(b), and 41(b) show that for the stiffness in the y-direction (22), the stiffness increases when tapes are aligned with the y-axis and decreases when the tapes are aligned with the x-axis. This is a result of the tapes being transversely isotropic, and stiffer in the fibre direction.

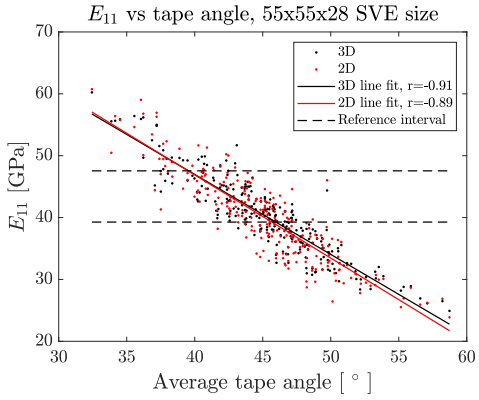

(a) E_{11} against the global average tape angle.

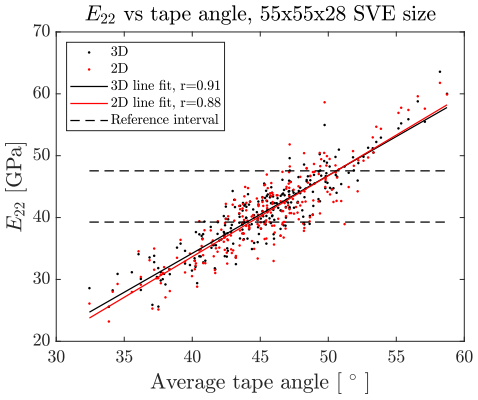

(b) E_{22} against the global average tape angle.

Figure 39: (a) E_{11} and (b) E_{22} against the global average tape angle for the 55x55x28 voxel domain.

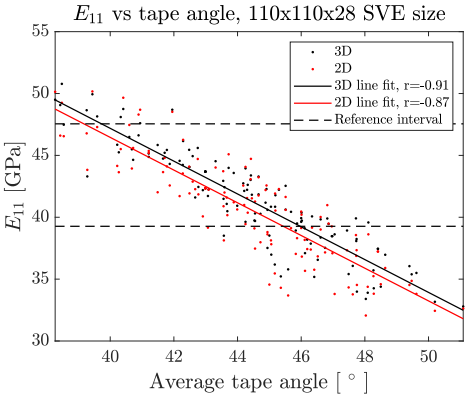

(a) E_{11} against the global average tape angle.

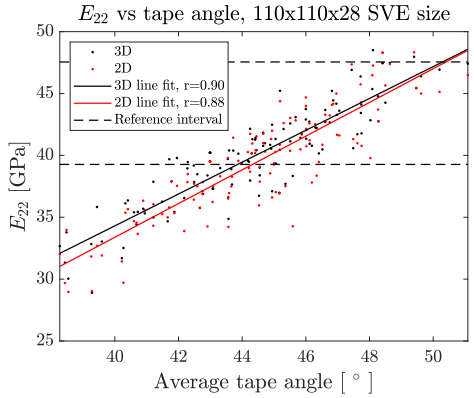

(b) E_{22} against the global average tape angle.

Figure 40: (a) E_{11} and (b) E_{22} against the global average tape angle for the 110x110x28 voxel domain.

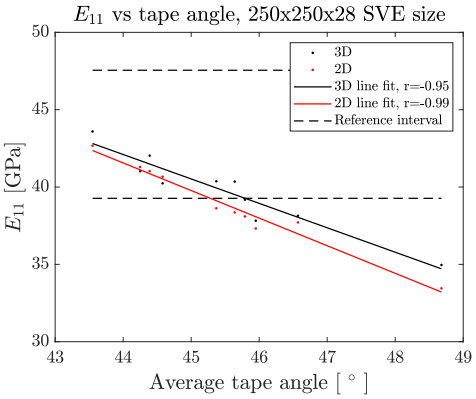

(a) E_{11} against the global average tape angle.

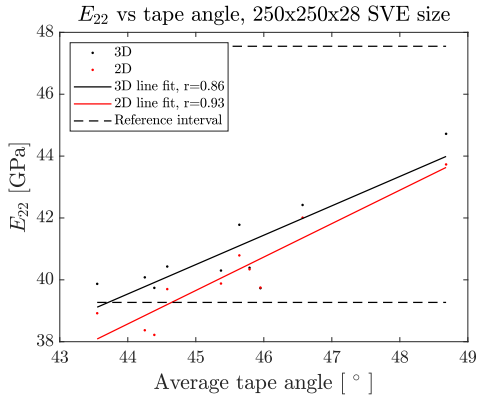

(b) E_{22} against the global average tape angle.

Figure 41: (a) E_{11} and (b) E_{22} against the global average tape angle for the 250x250x28 voxel domain.

5.3.2 Shear modulus versus Young's modulus

The comparison between shear modulus G_{12} and the average Young's modulus $\bar{E}$ (see Equation (30)) is presented in Figures 42(a)-44(a) for the small, medium and large SVEs, respectively. When the shear modulus increases, the average stiffness modulus decreases. One logical reason for this is that this relation is controlled by the orientation of the tapes in the domain, meaning that the direction of the tape alignment is the determining factor. For example, if one tape has the orientation 0° and another has the orientation 90° (like a '+'-sign), the average angle is 45°. If the tapes instead have the orientations 45° and 135° (like a '×'-symbol), the average is also 45°, using the method presented in Subsection 4.1.4. The two types of tape alignment would result in either Young's moduli E_{11} and E_{22} or G_{12} being dominant. As a result, generating the relation seen in Figures 42(a), 43(a) and 44(a), where, as the shear modulus increases, the average Young's modulus $\bar{E}$ decreases.

These results are dependent on the definition of the coordinate system, meaning that if it is rotated by 45°, the new alignment would flip the results for that specific geometrical material structure (as a result of the axes of the coordinate system being aligned as a × instead of a +). However, the important observation is that this shows how the structure will be influenced based on the direction of the applied load. In terms of the strength of the correlation between G_{12} and $\bar{E}$, the Pearson's coefficient indicates that for the 3D models, there is a strong correlation ($|r| > 0.8$), but for the 2D models it is not as definite, as seen by the lower r-values.

5.3.3 Young's modulus versus fibre volume fraction

It is expected that as the fibre volume fraction increases in a carbon fibre composite, so does the stiffness. Therefore, the average stiffness $\bar{E}$ of the small, medium, and large SVEs have been plotted against the global fibre volume fraction in Figures 42(b), 43(b), and 44(b), respectively. The comparison of $\bar{E}$ and the global fibre volume fraction does not indicate a strong correlation between the two smaller SVEs, with the highest correlation coefficient being 0.31. For the larger 250x250x28 voxel sized domain, there appears to be a strong correlation for both the 3D- and 2D-model, indicating that as the FVF increases, so does the modulus, which is to be expected. Even though this could be a correct finding, the lack of data points does not make it enough to make a confident prediction in terms of the relation between the two. An increase of data could continue the same strong correlation or make the plot converge toward a similar r-value as the smaller SVEs. This is a perplexing result, as the stiffness should increase with increasing fibre volume fraction. The reason why there is no strong correlation between the two may be that there is too little variation in the fibre volume fraction, the scatter plots are very grouped and therefore difficult to fit a curve to.

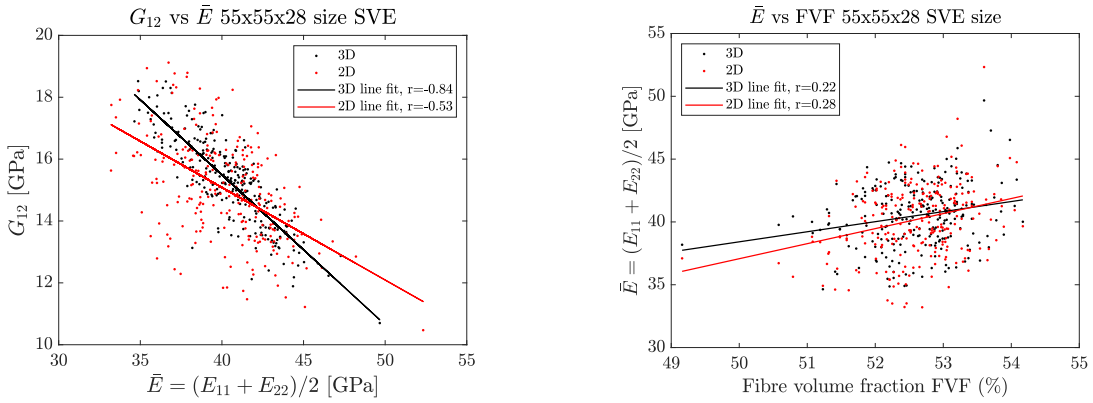

(a) G_{12} against $\bar{E}$.

(b) $\bar{E}$ against FVF.

Figure 42: Correlation between (a) G_{12} and $\bar{E}$, (b) $\bar{E}$ and FVF for the 55x55x28 voxel SVE.

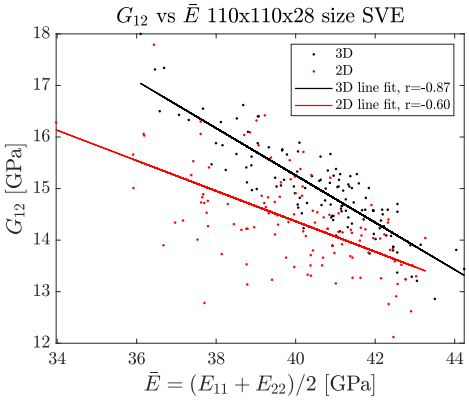

(a) G_{12} against $\bar{E}$.

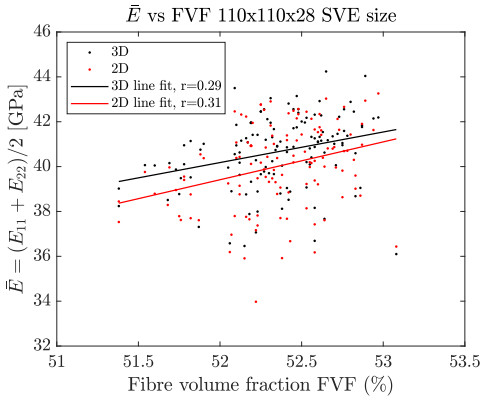

(b) $\bar{E}$ against FVF.

Figure 43: Correlation between (a) G_{12} and $\bar{E}$, (b) $\bar{E}$ and FVF for the 110x110x28 voxel SVE.

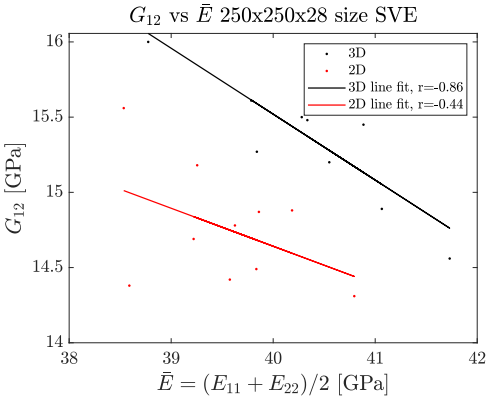

(a) G_{12} against $\bar{E}$.

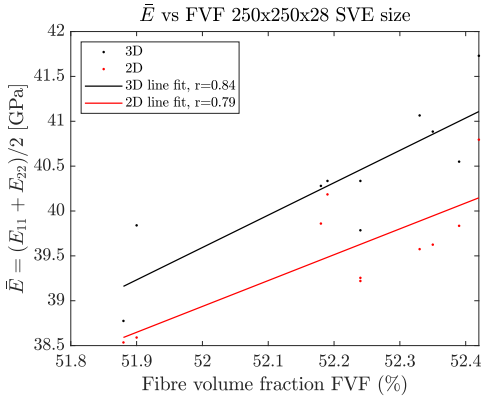

(b) $\bar{E}$ against FVF.

Figure 44: Correlation between (a) G_{12} and $\bar{E}$, (b) $\bar{E}$ and FVF for the 250x250x28 voxel SVE.

5.4 Application to a practical structure

To validate the premise of using homogenised effective material properties for the STBDC material, a test was conducted. A plate with clamped edges and a uniformly applied pressure on the top surface (magnitudes specified in Table 8), in the negative z-direction, was selected as the practical structure and load case. The test was conducted with one sample from each SVE size (small, medium, and large), and for both the full 3D-model and the reduced 2D-model, see Table 8. The most representative SVE, in terms of percentual difference from the averaged material properties, from each SVE size was used in the test. The most representative samples were number 221, 25, and 10 for the small, medium, and large SVEs, respectively. The heterogeneous plate models were already created during the homogenisation process, but the homogeneous models had to be produced based on an "empty" mesh (only containing nodal coordinates and element connectivity). The homogeneous models were given the material properties stated in Tables 6 and 7, for the 3D-models and the 2D-models, respectively. Equation (25) was used to calculate G_{23} and G_{13} for the homogenised plate models. The out-of-plane stiffness E_{33} was assumed to be equal to the transverse tape stiffness (9 GPa), for the homogenised 3D-models.

The von Mises stress distribution in the deformed plates is presented in Figures 45 and 46, where Figure 45 shows the results for the medium sized plate with a heterogeneous material structure, and Figure 46 shows the same results but for a homogeneous material structure. The von Mises equivalent stress is not generally applicable for anisotropic materials, but used here as a tool to show all stress components in one image. The first thing to note here is the stress magnitude. The maximum von Mises stress in the heterogeneous plate is roughly five times larger than in the homogeneous plate. It is evident that there are notable stress concentrations in the heterogeneous plate, that do not exist in the homogeneous model. The second important part to note is the stress distribution. The stress concentrations in the heterogeneous model increases the spectrum in the stress plot, which makes lower magnitude stresses difficult to differentiate. Therefore, the heterogeneous plate is shown again in Figure 47, but with the same limits on the colour spectrum as the homogeneous plate. The stress distribution in the heterogeneous plate is irregular, whereas in the homogeneous plate, the stress distribution is more symmetric.

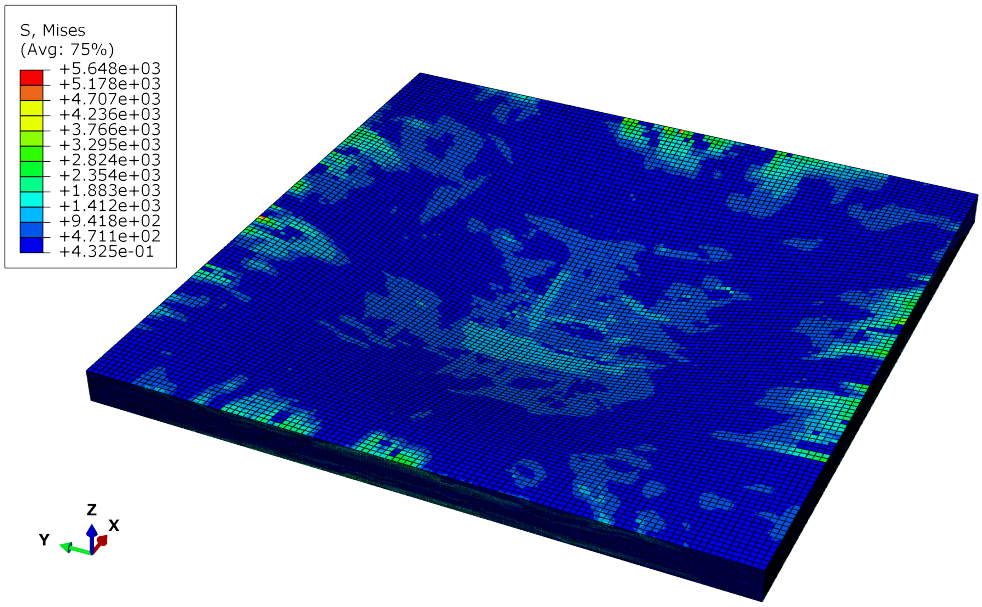

Figure 45: Von Mises stress distribution in a clamped heterogeneous plate (110x110x28 voxels) subjected to 10 [MPa] pressure.

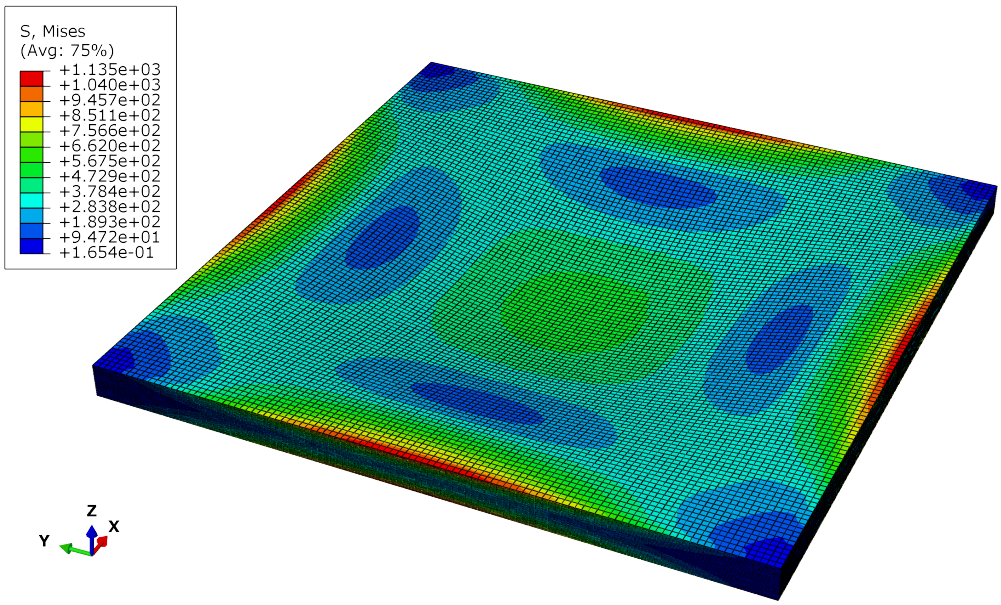

Figure 46: Von Mises stress distribution in a clamped homogeneous plate (110x110x28 voxels) subjected to 10 [MPa] pressure.

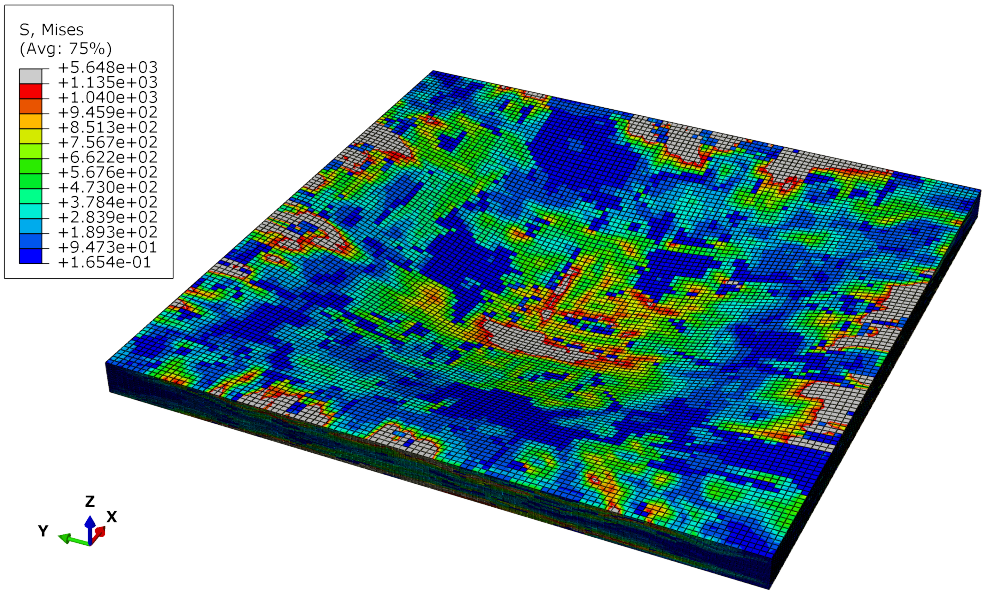

Figure 47: Von Mises stress distribution in a clamped heterogeneous plate (110x110x28 voxels) subjected to 10 [MPa] pressure, plotted with a limited colour spectrum. Grey elements fall outside limited spectrum.

The result of the practical implementation tests are summarised in Table 8 for all different plates tested, with index A, B, C, and D for the different types of models. There seems to be a trend that the heterogeneous 3D-models (A) have the highest von Mises stress and the largest deflection (displacement in negative z-direction) for all plate sizes. The models with homogenised material properties (B, D) and the 2D-models with heterogeneous material (C) have similar values, both regarding maximum stress and maximum deflection, compared to the heterogeneous 3D-models. Another evident trend in Table 8 is that the percentual difference in maximum deflection between the heterogeneous 3D-models and the other models decreases with increasing plate size. The average maximum deflections of the homogeneous models (B, D) and the heterogeneous 2D-models (C),

$$\bar{u}^{\max}_{\mathrm{BCD}} = \frac{u^{\max}_{\mathrm{B}} + u^{\max}_{\mathrm{C}} + u^{\max}_{\mathrm{D}}}{3}, \tag{32}$$

is 36.7%, 17.5%, and 11.7% lower than the deflection of the heterogeneous 3D-models ($u^{\max}_{\mathrm{A}}$), for the small, medium and large plates, respectively. Keep in mind that what is here called a large plate is only 100x100x2.24 mm, which could be small in relation to a real application. If the trend that the displacement results of the 2D-models converge to the same results as the heterogeneous 3D-models with increasing plate size continues, the 2D-models could be very useful in industry due to the drastically reduced solution time compared to the 3D-models. Regarding the maximum von Mises stress, the same trend does not exist. The maximum von Mises stress of the heterogeneous 3D-model ($\sigma^{\max}_{\mathrm{VM,A}}$) is 5.14, 4.63, and 5.99 times larger than the average of the other models ($\bar{\sigma}^{\max}_{\mathrm{VM,BCD}}$), for the small, medium, and large plates, respectively. The homogeneous 3D-models (B) and the 2D-models (C, D) simply cannot capture the material-induced stress concentrations that exist in the heterogeneous 3D-model (A).

Table 8: Results of practical implementation simulations.

Plate size [voxels]	**Applied pressure [MPa]**	**Model**	**Max deflection ($u^{\max}$) [mm]**	**Max Von Mises stress ($\sigma^{\max}_{\mathrm{VM}}$) [GPa]**
55x55x28	50	A - Heterogeneous 3D	0.66	7.11
		B - Homogeneous 3D	0.41	1.52
		C - Heterogeneous 2D	0.42	1.45
		D - Homogeneous 2D	0.41	1.18
110x110x28	10	A - Heterogeneous 3D	1.43	5.65
		B - Homogeneous 3D	1.18	1.14
		C - Heterogeneous 2D	1.14	1.34
		D - Homogeneous 2D	1.22	1.18
250x250x28	1	A - Heterogeneous 3D	3.55	3.85
		B - Homogeneous 3D	3.04	0.58
		C - Heterogeneous 2D	3.25	0.80
		D - Homogeneous 2D	3.13	0.54

6 Discussion

The produced results indicate that the work conducted has been successful with regards to the set objective, generating predictions close to that of experimental data. However, the generated material samples are not perfect representations of reality. Therefore, the following subsections aim to discuss possible improvements that could be implemented to further improve the developed methodology. Potential drawbacks relating to the chosen method of geometry modelling and homogenisation will be further elaborated here.

6.1 Geometrical model

The developed algorithm, generating a geometrical model of an STBDC appears to provide a qualitatively accurate virtual representation of the real material. However, upon further analysis, there are areas where the algorithm fails to produce a fully realistic model, which could be improved.

The geometrical properties, such as global fibre volume fraction and global average tape angle, provide an indication of how representative the distribution of tapes within each sample is. For the global average tape angle of each sample, there were deviations around the desired 45°. But, in terms of convergence for a set of samples, the average ended up very close to 45°. For the three SVE sizes, the largest deviation was found for the 110x110x28 SVEs with an average tape angle of 44.43°, 1.3% lower than 45°. This indicates that, in terms of angles, there appears to be no consistent bias in the algorithm; it generates material samples both above and below the mid-point value 45°. However, there are a number of potential cases where the placement of tapes is not fully stochastic but rather a consequence of the current constraints of the algorithm, which will be further detailed in the following subsections.

The global fibre volume fraction varied between 49.16% and 54.17% for all samples, with the vast majority laying around 52%-53%. This indicates that there is no issue in terms of consistency with the placement and distribution of the tapes throughout the structure. However, comparing to experimental data, the resulting values are below what would be expected for this type of material (based on the chosen input parameters). The target value based on the reference material [4] was a FVF of 57.85%. The largest deviation from this target appeared for the 250x250x28 voxel SVEs with an average FVF of 52.21%, approximately 10% lower than the reference material. This would suggest that, in terms of fibre volume fraction, there is no problem related to consistency or an integrated bias within the model generation procedure. Instead, it appears to be connected to the algorithm currently not being able to generate fully realistic replicas of the real material, with a realistic fibre volume fraction.

6.1.1 Deforming and sliding tapes

The low FVF is a result of the current algorithm not being able to fully capture the microscopic structure of an STBDC. One process that occurs during the manufacturing of these materials is that, during the compression phase, tapes can deform and thereby become thinner and larger in-plane. This is not currently captured by the algorithm, as tapes are not allowed to change length, width or thickness during the creation of the virtual material sample. There is a possibility that introducing a geometrical adaptation of the tapes would allow the tapes to more naturally fill up empty space, thus increasing the FVF of the material sample.

Furthermore, the layer condition, covered in Section 4.1.2, was mainly implemented to avoid the tapes piling up in certain areas of the domain. If the compression phase is modelled as two mould halves closing and the tapes being allowed to deform, the piles would be subjected to a higher pressure and therefore be forced to deform, naturally becoming packed more tightly. The addition of allowing the tapes to be compressed and more tightly packed in the structure, might therefore, be beneficial and alleviate this issue of pile-buildups as a result of the tapes being more adaptable.

Another process that the model does not capture is sliding of the tapes, which means that tapes are not allowed to adjust their position to better fit into an open space. Adding a sliding functionality could possibly allow for less tapes to be ignored during the placement procedure due to breaching the layer-constraint, and result in a more densely packed structure with a higher FVF. Possibly, combining tape

sliding with the tapes being allowed to deform would result in significantly more accurate representations of the real material.

6.1.2 Tape-resolution

The geometry of the tapes was developed as a voxel-structure, being constructed out of cuboids (volume pixels). This produces a choice of resolution for each tape, meaning how many voxels are used to represent each tape. One known consequence of using too few voxels is that Bresenham's algorithm creates the same voxelised tape representation for multiple angles (small changes of angle results in the same geometry). For the resolution used to produce the results presented in Section 5, there was no geometrical difference between tapes with 0° and 1° degree angles. Thus, not fully capturing the full stochastic nature of the material. The choice of resolution is therefore an optimisation problem between quality and speed, but there are known consequences of lowering it (less realistic tapes, same geometry for different angles, etc.) which should be taken into consideration. The impact of these effects are unknown and could be something worth looking into.

6.1.3 Selective tape placement

The implementation of the layer condition (to avoid pile build-ups) involves an iterative process that may prevent tapes to be placed in the domain if they breach the current maximum allowed height. A consequence of this is that, if two tapes are parallel to each other leaving space for one tape between them, a tape will in some cases not be able to be placed there without breaching the height limit, unless it is aligned with the other two, as shown in Figure 48. The left figure illustrates that, due to the tapes being on top of each other, the height limit has been breached. Thereby, the tape will not be placed. In the right figure, since the tape is aligned and not on top of the others (thus not breaching the height limit), it would be placed. The result of this would be that some areas produce a biased tape orientation, which would not be present if all tapes were placed completely random. The frequency at which this happens during the generation of a material sample is unknown. However, it is something that would happen if the situation appeared and as a result, might negatively impact the elastic material property prediction.

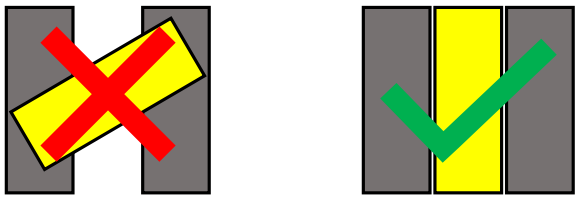

Figure 48: Illustration of selective tape placement.

6.1.4 Removal of top layers

The removal of the top layers was implemented as a way to increase the overall fibre volume fraction and to achieve a more homogeneous distribution of tape and matrix across the sample. This modification resulted in an increase of the fibre volume fraction, also increasing the stiffness of the sample compared to the unaltered version of the same sample. However, tapes at the top layers were cut and weakened as a result of this modification, as shown in Figure 21(b). Although it could not be determined whether this weakening of the top layers had a negative impact on the homogenised stiffness, or to what degree, this could be one factor in why the model under-predicts the stiffness of the real material.

6.1.5 Under-stiff draping

The algorithm that creates the draping effect works under one specific condition: all voxels in a tape should take the lowest possible position, as long as the vertical distance between adjacent voxels is not larger than one. This definition results in an issue with regards to how stiff the tapes are when it comes to draping into small crevices. In Figure 49, a specific scenario is presented where this condition can lead to unrealistic results, where the tapes drape into small crevices. This issue is also a consequence of the resolution. If a larger number of voxels would be used through the thickness, the problem would be mitigated (since each voxel would be thinner compared to the full thickness). Another way to avoid under-stiff draping would be to implement an additional condition in the algorithm.

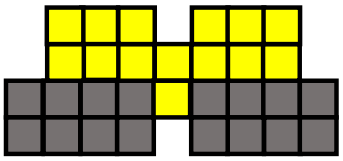

Figure 49: Illustration of draping into a crevice.

6.2 Homogenised material properties

The homogenisation procedure using periodic boundary conditions on the aperiodic SVEs appears to have generated reasonable results. Compared to experimental data, E_{22} and E_{11} were, on average, 7.11% and 7.62% lower than the reference mean respectively, but 2.68% and 2.21% above the lower bound, respectively. In terms of using periodic boundary conditions, other studies [23, 24, 25] suggest that the way it has been implemented in this thesis is a fully reasonable method.

The full 3D-model was reduced to a 2D-model using classical laminate theory. Comparing the results between the two approaches, it was shown that the 2D-model provided a slightly lower prediction of the elastic material properties. Possible reasons for the lower prediction could be the use of reduced integration, which is covered in further detail in subsection 6.2.3

For both methods, the relation between Young's moduli and the fibre volume fraction is shown to be a weak correlation for the small and medium SVEs, but much stronger for the large SVEs. The exact reason for this is unknown, and the weak correlation may simply be a consequence of the high variability in predicted values for the smaller SVE sizes. Also, as the SVE size increases so does the strength of the correlation between Young's moduli and the fibre volume fraction. But, the number of data points also decreases with increasing SVE sizes, implying that if more data is produced, the strength of the correlation could worsen or improve.

6.2.1 Scaling of input material parameters

It was noted in experimental tests [4, 15] that the fibre volume fraction of the plate was higher than the initial fibre volume fraction of the tapes. This means that some matrix material must have bled through the mould. In addition, some of the matrix migrate from the tapes to fill resin pockets. Therefore, it must be so that the fibre volume fraction of the individual tapes increases during the manufacturing process. To what extent this happens depends on the degree of resin bleeding and the concentration of resin pockets. With an increase in tape fibre volume fraction, the input material parameters of the tapes should be scaled accordingly.

The lack of information regarding the amount of resin bleeding and resin pockets in the experimental data hindered the scaling of material parameters in this project. An educated guess of the FVF increase could have been made and implemented, resulting in a stiffer plate response. However, it was determined that rather than scaling the results, making them more "accurate", the unaltered material parameters should be used to better understand how to realistically model STBDCs. Thereby showcasing where the current version of the algorithm lacks realism.

6.2.2 Variation between SVE sizes

The measure of statistical significance is difficult for these types of materials where the micro- and mesostructural variability is high. A simple and logical measure to use is the standard deviation of the sample properties. The standard deviation increases with deviations from the sample mean and decreases with the number of samples, see Equation (28). Therefore, it is expected that large sample sizes with large variability and small sample sizes with small variability have similar standard deviations.

It was expected that, with the number of samples decreasing with increasing SVE size, the standard deviation would be similar for the different respective properties. This was, however, not the case. Perhaps basing the sample size on the SVE side length between the small and medium SVEs was a mistake. The variability of some properties, e.g. the average tape angle, are expected to have a direct dependence on the volume (or at least in-plane area) of the SVEs. Therefore, it might have been wiser to base the number of samples on the volume of the SVE instead of the side length. In that case, the total volume (number of samples times SVE volume) would be the same for the different SVE size. Based on the total volume of the large SVE, the medium and small SVEs should have 52 and 207 samples, respectively, for the total volume to be approximately constant. These numbers were surpassed by far in the test but still, the standard deviation varied between the SVEs, see Tables 6 and 7. Thereby disproving this argument.

It seems like the number of samples required for the different SVE sizes to have a similar standard deviation cannot be determined in advance but must be tested. The standard deviation could be used as a measurement of when the required number of samples have been reached for the different SVE sizes.

6.2.3 Reduced integration

One of the observations made when comparing the elastic properties of the full 3D- and the reduced 2D-model, was that the predicted properties for the 2D-model were consistently lower and more varying than that of the 3D-model. The one big difference between the two, apart from that the 2D-model has an intermediate analytical homogenisation step, is that the reduced 2D-model uses reduced integration. Reduced integration means that each element uses one less integration point in each direction compared to fully integrated elements. One known consequence of using reduced integration which may be a reason for the lower values, is that these elements tend to be too flexible (under-stiff), which would explain the difference between the 2D- and the 3D-models. A smaller study using fully integrated elements for the 2D-model could determine if the reduced integration was the cause of the lower stiffness, but due to the time limitations it was not included in this thesis.

6.2.4 Convergence of Young's moduli

The fully converged E_{11} and E_{22} should result in the same value in the absence of an implemented bias in the model. The possibility for either of the two being larger than the other should be equally likely. The important thing to note, however, is that E_{11} is the largest out of the two for the 110x110x28 SVE size, whereas and for the other two, E_{22} is larger. This would indicate that for different sets of data, either of the two moduli can be larger or smaller, and it is strictly dependant on a statistical probability. All SVEs were generated using the same code, so as long as there is no specific effect of the 110x110x28 dimensions that have been overlooked, this indicates that there is no built-in bias in the model. To confidently claim that there is no built-in bias, further analysis would have to be conducted. This further analysis could also include studying the number of realisations that would have to be done for the different SVE sizes, resulting in the two moduli converging to the same value.

6.2.5 Appearance of tape bands

The appearance of bands in the domain could be a possible factor affecting the resulting elastic material properties, without giving any indication in global average tape angle or fibre volume fraction. Bands are tapes aligned in a specific direction going across the domain as one long ribbon, schematically shown in Figure 50. Whether or not there is a specific reason for these to appear is unknown. It was however noted that when they did appear, at least for the 2D-models, the result deviated from the standard correlation between e.g. global average angle and Young's moduli. For the normal case (without band formation), the global average tape angle would indicate which Young's modulus that would be the largest (seeing what direction the bias is towards). However, in the presence of bands, this appeared not to be the case. For example, say that the schematic plate in Figure 50 has a representative fibre volume fraction and an average in-plane tape angle of 45°. If the plate is loaded in tension in the vertical direction, the horizontal tape band will act as a weak link due to the tapes being more flexible in the transverse direction. Therefore, the resulting effective stiffness of the entire plate will be based on a single weak link, resulting in a lower stiffness in the vertical direction.

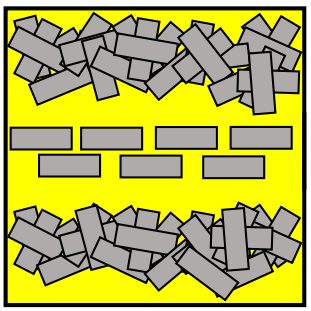

Figure 50: Illustration of the appearance of tape bands.

6.3 Practical implementation of homogenised material properties

The practical implementation, presented in Subsection 5.4, was aimed to investigate the difference between and impact of using a heterogeneous material structure and a homogeneous material structure in a model. The test was conducted for a rather simple, yet common, structure and load case: a fixed plate subjected to a uniform pressure. It was seen in the test that the 2D-models (heterogeneous and homogeneous) and the homogeneous 3D-model could not capture the material-induced stress concentrations observed in the heterogeneous 3D-model (as expected). These material-induced stress concentrations may have been exaggerated by the voxel structure, having sharp corners in the material interfaces. Although, it would likely be a good idea to intentionally include some material weak spots in a homogenised model to capture these stress concentrations. In particular if this kind of model is to be used in industry. Further, the 2D-models and the homogeneous 3D-model showed a trend to converge towards the heterogeneous 3D-model, in terms of maximum deflection, with increasing plate size. This might be a result of a more uniform material structure in the large plate compared to the smaller plates. However, if the trend continues so that the difference between the 2D-models and the heterogeneous 3D-model is reduced, the 2D-models can have some real potential. Especially if intentional material weak spots are introduced to also capture the material-induced stress concentrations in the heterogeneous 3D-model.

7 Conclusions

The aim of this thesis was to develop a method to predict the elastic properties of STBDCs. Firstly, to develop an algorithm to create a virtual representation of the material at hand, secondly, to evaluate homogenisation strategies for predicting the averaged macroscopic properties. It can be concluded that the voxel-based material samples generated from the developed algorithm, in combination with periodic boundary conditions, does produce reliable predictions for the elastic properties of STBDCs.

The characteristic structure of STBDCs has been captured using the developed algorithm. Tapes are stochastically placed in the domain, with an average tape angle that at most deviate 1.3% from the sought 45° target. Furthermore, the distribution of tapes throughout the cross-section has been accurately replicated with respect to experimental samples, showing a clear resemblance in both structure and distribution of material. However, the geometrical models fail to reach the fibre volume fraction of real material samples, as a result of not being able to capture effects such as resin bleeding, sliding of tapes, compression of tapes, etc. Whether or not this would correct the slight under-prediction of the effective elastic parameters is unknown, but the data indicates that finding effective methods of increasing the FVF would be a reasonable continuation of the model development.

The implementation of periodic boundary conditions on the 3D-model results in accurate predictions of the elastic properties, within one standard deviation of the experimental mean. The Young's moduli E_{11} and E_{22} were on average 7.37% lower than the reference mean, but 2.40% above the lower bound. The established relationships between the geometric properties (fibre volume fraction and global average tape angle) and elastic properties are reasonable, capturing how these in practice would affect the structural integrity of the material. Whether or not this can be used to purposely manipulate the parameters to get specific characteristic out of the material has not been concluded, although this might be an area of further possible research. Furthermore, it can be concluded that the methodology used to generate the reduced 2D-model based on the 3D-geometry and classical laminate theory is a reasonable method to implement to drastically lower the computational time. The predictions of elastic properties for the reduced model are slightly lower (at most 2.1%) compared to the full 3D-model, but deemed sufficiently accurate considering the computational savings. Possible improvements can be made through full integration or increased mesh resolution, but in terms of comparison to the experimental data, the improvement would have to be made on the development of the 3D-model, which the reduced 2D-model is based upon.

8 Future work

The future work of the STBDC material should be aimed towards advancing its use in industry. Based on the work conducted in this project and findings in other studies, there are a few interesting paths to take.

Regarding finite element analysis of structural components, the only path to capture the variability of the material and have a reasonable solution time is to use something like the reduced 2D-model developed in this project. Although it is not perfect, it enables the analysis of large structures. A different approach for the reduced 2D-model is to, instead of only using effective material properties evaluated based on the laminate stiffness matrix ($\boldsymbol{A}$), use also the coupling stiffness matrix ($\boldsymbol{B}$) and the bending stiffness matrix ($\boldsymbol{D}$) in the finite element model. This method was actually implemented in this project and showed a lot of promise for real load cases as it captures the bending effects better, but was not analysed in detail and had to be omitted from the report due to time limitations. Further improvements and detailed analysis of this alternative reduced model is a topic for future work.

A problem with the stochastic material for industrial applications is that it is just that - stochastic. Without control, the variability of the material structure will inevitably result in components that do not meet the specifications and thereby risk to experience premature failure. This project found strong correlations between geometrical features and elastic properties, other studies have found similar connections between geometrical features and failure mechanisms. An interesting focus for future work would be to experimentally test and validate these connections between mechanical and geometrical properties. In production, cameras and image/pattern recognition software could then be used to ensure that geometrical effects that cause the material to be too flexible and fail prematurely are avoided. Either by simply discarding parts where these geometrical effects appear, or using targeted tape placement to mitigate the trailing mechanical penalties.

In an academical or research setting, there are some topics of future work that can be investigated. First of all, this could be to develop a material structure generator that is both fast and generates models representative of the real material. The numerical model developed in this project is somewhat of a middle-ground, reasonably fast and fairly accurate. Based on the literature study, other material generation algorithms tend to the extremes. They are either very fast of very accurate. Furthermore, the voxel-based structure should be abandoned to give way for a conforming mesh alternative. The voxel-based structure may be good for predicting elastic properties, but it has very clear limitations when it comes to evaluating strength and failure mechanisms. For example, the stress concentrations induced by the sharp corners of the elements at the material interface, and the ragged edges of angled tapes, prevent the prediction of effects like tape pull-out. A final focus for future work is, again, to further improve the reduced 2D-model. The computational savings of these models are so large compared to the full 3D-model, that they cannot be ignored. If the reduced models are developed to be sufficiently accurate, they could prove to be a very useful tool in the industry.

References

[1] J. Ryatt and M. Ramulu, "Prediction of tensile failure of stochastic tow-based discontinuous composites via mesoscale finite element analysis," *Composite Structures*, vol. 279, pp. 1–16, 2022, ISSN: 0263-8223. DOI: https://doi.org/10.1016/j.compstruct.2021.114769.

[2] Y. Li, S. Pimenta, J. Singgih, C. Ottenwelter, S. Nothdurfter, and K. Schuffenhauer, "Understanding and modelling variability in modulus and strength of tow-based discontinuous composites," 21st International Conference on Composite Materials, Xi'an, China, Aug. 2017.

[3] W. Joost, "Reducing vehicle weight and improving u.s. energy efficiency using integrated computational materials engineering," *JOM*, vol. 64, pp. 1032–1038, Sep. 2012, ISSN: 1543-1851. DOI: https://doi.org/10.1007/s11837-012-0424-z.

[4] Y. Li, S. Pimenta, J. Singgih, S. Nothdurfter, and K. Schuffenhauer, "Experimental investigation of randomly-oriented tow-based discontinuous composites and their equivalent laminates," *Composites Part A: Applied Science and Manufacturing*, vol. 102, pp. 64–75, Nov. 2017, ISSN: 1359-835X. DOI: https://doi.org/10.1016/j.compositesa.2017.06.031.

[5] S. B. Visweswaraiah, M. Selezneva, L. Lessard, and P. Hubert, "Mechanical characterisation and modelling of randomly oriented strand architecture and their hybrids – a general review," *Journal of Reinforced Plastics and Composites*, vol. 37, no. 8, pp. 548–580, 2018, ISSN: 0731-6844. DOI: https://doi.org/10.1177/0731684418754360.

[6] C. Park and W. Lee, "3 - compression molding in polymer matrix composites," in *Manufacturing Techniques for Polymer Matrix Composites (PMCs)*, ser. Woodhead Publishing Series in Composites Science and Engineering, S. G. Advani and K.-T. Hsiao, Eds., Woodhead Publishing, 2012, pp. 47–94, ISBN: 978-0-85709-067-6. DOI: https://doi.org/10.1533/9780857096258.1.47.

[7] L.T. Harper, C. Qian, R. Luchoo, and N.A. Warrior, "3d geometric modelling of discontinuous fibre composites using a force-directed algorithm," *Journal of Composite Materials*, vol. 51, no. 17, pp. 2389–2406, 2016. DOI: https://doi.org/10.1177/0021998316672722.

[8] M. Selezneva and L. Lessard, "Characterization of mechanical properties of randomly oriented strand thermoplastic composites," *Journal of Composite Materials*, vol. 50, no. 20, pp. 2833–2851, 2016, ISSN: 0021-9983. DOI: https://doi.org/10.1177/0021998315613129.

[9] B.D. Agarwal, L.J. Broutman, and K. Chandrashekhara, *Analysis and performance of fibre composites (4th ed.)* 111 River Street, Hoboken, NJ 07030, USA: John Wiley and Sons, 2018, ISBN: 9781119389989.

[10] G.P. Thomas, "Compression molded composites: Processes, benefits and applications," *AZO MATERIALS*, Feb. 10, 2014. [Online]. Available: https://www.azom.com/article.aspx?ArticleID=10665 (visited on Mar. 17, 2022).

[11] P. Feraboli, F. Gasco, B. Wade, S. F. Maier, A. Masini, and L. DeOto, "Lamborghini forged composite ® technology for the suspension arms of the sesto elemento," 2011.

[12] NORDAM Group LLC. "OEM design & manufacturing." (2015-05-18), [Online]. Available: https://www.nordam.com/products-services/commercial-aircraft (visited on Mar. 22, 2022).

[13] S. Pimenta. "Discontinuous composites: A route for manufacturability, damage tolerance and sustainability." (Jul. 2020), [Online]. Available: http://www.bristol.ac.uk/media-library/sites/cabot-institute-2018/events/Day%202%20-%20Soraia%20Pimenta.pdf (visited on Mar. 17, 2022).

[14] Y. Wan and J. Takahashi, "Tensile properties and aspect ratio simulation of transversely isotropic discontinuous carbon fiber reinforced thermoplastics," *Composites Science and Technology*, vol. 137, pp. 167–176, 2016, ISSN: 0266-3538. DOI: https://doi.org/10.1016/j.compscitech.2016.10.024.

[15] M. Johansen, "Manufacturing and Characterisation of Ultra-Stiff Composite Material," M.S. thesis, Chalmers University of Technology, Gothenburg, Sweden, 2019.

[16] M. Selezneva, S. Roy, S. Meldrum, L. Lessard, and A. Yousefpour, "Modelling of mechanical properties of randomly oriented strand thermoplastic composites," *Journal of Composite Materials*, vol. 51, no. 6, pp. 831–845, 2017, ISSN: 0021-9983. DOI: https://doi.org/10.1177/0021998316654748.

[17] S.Z.H. Shah, R.S. Choudhry, and S. Mahadzir, "A new approach for strength and stiffness prediction of discontinuous fibre reinforced composites (dfc)," *Composites Part B: Engineering*, vol. 183, pp. 1–12, Feb. 2020, ISSN: 1359-8368. DOI: https://doi.org/10.1016/j.compositesb.2019.107676.

[18] Y. Pan, L. Iorga, and A. A. Pelegri, "Numerical generation of a random chopped fiber composite rve and its elastic properties," *Composites Science and Technology*, vol. 68, pp. 2792–2798, Jun. 2008, ISSN: 0266-3538. DOI: https://doi.org/10.1016/j.compscitech.2008.06.007.

[19] Dassault Systemes, *SIMULIA User Assistance - Abaqus*, 2020.

[20] N. Ottosen and H. Petersson, *Introduction to the finite element method*. Edinburgh Gate, Harlow, Essex, CM20 2JE, England: Prentice Hall, 1992, ISBN: 978-0-13-473877-2.

[21] R. Jänicke, F. Larsson, and K. Runesson, *Computational homogenization in material mechanics*, Gothenburg, Sweden, Sep. 2018, Lecture notes for the course Material Mechanics.

[22] S.L. Omairey, P.D. Dunning, and S. Sriramula, "Development of an abaqus plugin tool for periodic rve homogenisation," *Engineering with Computers*, vol. 35, no. 2, pp. 567–577, 2018, ISSN: 1435-5663. DOI: https://doi.org/10.1007/s00366-018-0616-4.

[23] V.D. Nguyen, E. Béchet, C. Geuzaine, and L. Noels, "Imposing periodic boundary condition on arbitrary meshes by polynomial interpolation," *Computational Materials Science*, vol. 55, pp. 390–406, 2012, ISSN: 0927-0256. DOI: https://doi.org/10.1016/j.commatsci.2011.10.017.

[24] W. Tian, L. Qi, X. Chao, J. Liang, and M. Fu, "Periodic boundary condition and its numerical implementation algorithm for the evaluation of effective mechanical properties of the composites with complicated micro-structures," *Composites Part B: Engineering*, vol. 162, pp. 1–10, 2019, ISSN: 1359-8368. DOI: https://doi.org/10.1016/j.compositesb.2018.10.053.

[25] K. Terada, M. Hori, T. Kyoya, and N. Kikuchi, "Simulation of the multi-scale convergence in computational homogenization approaches," *International Journal of Solids and Structures*, vol. 37, no. 16, pp. 2285–2311, 1998, ISSN: 0020-7683. DOI: https://doi.org/10.1016/S0020-7683(98)00341-2.

[26] I. Rodrigues Lopes, "Multi-scale modelling and analysis of multi-phase solids using second-order computational homogenisation at finite strains with parallel computing," Ph.D. dissertation, Jun. 2019.

[27] A. Wetzler, *Bresenham optimized for matlab*, Jun. 2010. [Online]. Available: https://se.mathworks.com/matlabcentral/fileexchange/28190-bresenham-optimized-for-matlab.

[28] Z. Chen, C. Wu, and H. T. Tsui, "A new image rectification algorithm," *Pattern Recognition Letters*, vol. 24, no. 1, pp. 251–260, 2003, ISSN: 0167-8655. DOI: https://doi.org/10.1016/S0167-8655(02)00239-8.

[29] P. Feraboli, E. Peitso, F. Deleo, T. Cleveland, M. Graves, and P. Stickler, "Characterization of discontinuous carbon fiber/epoxy systems for aerospace applications: Part i," Apr. 2012. DOI: https://doi.org/10.2514/6.2008-1939.

[30] A.R. Melro, P.P. Camanho, F.M. Andrade Pires, and S.T. Pinho, "Micromechanical analysis of polymer composites reinforced by unidirectional fibres: Part ii – micromechanical analyses," *International Journal of Solids and Structures*, vol. 50, no. 11, pp. 1906–1915, 2013, ISSN: 0020-7683. DOI: https://doi.org/10.1016/j.ijsolstr.2013.02.007.

[31] Chalmers centre for computational science and engineering (C3SE). "C3SE homepage." (2022-05-23), [Online]. Available: https://www.c3se.chalmers.se/ (visited on May 24, 2022).

[32] M. Fagerström, *Lecture notes: Composite mechanics*, Gothenburg, Sweden, 2019, Lecture notes for the course Composite Mechanics.

[33] R. Veitch, S. Zuyev, and A. Muratov, *Introduction to probability and mathematical statistics*, Gothenburg, Sweden, Apr. 2017, Lecture notes for the course Mathematical statistics.