



**Development of Fibre and Particulate Filled  
Aerogel Composites for Subsea Pipeline  
Applications**

A thesis in fulfilment of the requirements for the degree of Doctor  
of Philosophy

By  
Andrew Webley

Department of Mechanical and Aerospace Engineering  
University of Strathclyde  
Glasgow

UK  
2020

## **Declaration of Authenticity and Author's Rights**

This thesis is the result of the author's original research. It has been composed by the author and has not been previously submitted for examination which has led to the award of a degree.

The copyright of this thesis belongs to the author under the terms of the United Kingdom Copyright Acts as qualified by University of Strathclyde Regulation 3.50. Due acknowledgement must always be made of the use of any material contained in, or derived from, this thesis.

Signed:

Date:

## **Acknowledgements**

This research was sponsored by Blueshift International Materials and the Energy Technology Partnership, with additional funding provided by the University of Strathclyde. Dr. Liu Yang and Professor Nash, both from the University of Strathclyde, provided vital insights and support to the research, without which it would not have been possible. Various other researchers at the University of Strathclyde provided experimental data that was used to validate the results and guide the properties to be investigated.

I would also like to thank my family and friends, for their constant support and encouragement.

## Abstract

The aim of this project is to develop an analytical tool through Finite Element Method to understand and guide the design of thermal insulation materials. In particular, fibrous mat and aerogel particle-filled resin are focused in the context of subsea pipe-in-pipe (PiP) application for oil and gas extraction. The former has become part of a novel combination with superinsulating material – aerogel, and found its commercial use in the annulus region of the PiP. The latter is considered to be a form of upgrade to the current nylon-based centraliser stabilising the PiP configuration under hydrolic pressure from the deep seawater.

The use of porous materials (e.g. fibrous matt and foam) as thermal insulation has a long history, A randomly orientated fibre mat is effective insulation due to the lack of a straight heat conduction pathway and relatively small pores in the fibrous structure. Recent development of superinsulating materials involves combining fibrous mat with aerogel, which has resulted in a flexible and less compressible aerogel-fibre blanket with lower thermal conductivity (14 – 20 mw/mK) than typical conventional fibrous counterparts (35 – 50 mw/mK). However, a fundamental understanding of how the fibrous mat interacts the aerogel has not been extensively studied. The use of porous particle-filled resin composites, particularly with aerogel particles, has received less studies in thermal cases and a similar understanding of such integration must be developed in order to inform the design process.

The rationale for carrying out this research was to create a range of simulations that can be easily used to predict the effect of changing properties of these above composites on their thermal and mechanical performance. The use of simulations is less time consuming and allows more parameters to be easily varied to see what the optimal fibre design is for specific applications.

In order to simulate the effects of varying the composition of the mats, an algorithm to generate a random fibre mat was produced, based on the work of Arambakam and Tafreshi [1]. This fibre positioning was done in MathWorks' MATLAB, which was used to generate a script for producing a 3D geometry that can be incorporated in ANSYS APDL. This geometry was then used as the basis for a 3D finite element method model that utilised ANSYS Mechanical to mesh and solve the simulation. The same fibre generation process was used for both the mechanical and the thermal simulations, but with the bounding region being differently shaped to accommodate representing the standard experimental techniques used to measure these values.

Simulations were carried out to investigate the effect of various parameters of the fibre mat. Specifically, the fibre orientation and the volume fraction of the mat were the main parameters of importance, with the orientation both in and out of the plane of heat transfer being investigated. The ratio of straight fibres to sinusoidal fibres was also investigated as a key parameter affecting the heat transfer, with the volume of straight fibres to total fibre volume being used to determine the "straight fibre fraction" of the fibres. The effect of the fibre length and diameter were also investigated, though it was found that on the scales investigated, they had very little effect on the thermal conductivity.

A study into using fully random fibres, where the fibre can be represented by a continuous curve in 3D space that varies in direction throughout the length, was carried out. However, the high level of complexity of this fibre configuration along with the large number of fibre intersections meant that meshing the geometry produced was very difficult. Some successful meshes were produced; however, they included a number of elements that was too large to successfully produce a solution using finite element analysis.

For aerogel particle-filled resin, an effort was made to investigate the effect of dispersing particles of aerogel through a resin matrix, in an effort to reduce the thermal conductivity of the matrix. The results of this showed that it was theoretically possible to achieve significantly lower thermal conductivity. The material properties generated from the FE simulations were then used in modelling the centraliser to determine the insulations effectiveness in pipe in pipe insulation. These simulations included the particle simulation data to guide the addition of aerogel particles in the resin matrix to further reduce the conductivity.

The work carried out in this project created an alternative approach to generating and manipulating geometry (shape and size distribution) and spatial position within a representative volume element (RVE) of composites. This has enabled a micro-scale modelling for obtaining such properties as thermal conductivity and modulus of the fibrous matt in itself as well as the particle-filled resin composites. An excellent agreement was found between the modelling results and experimental data.

# Table of Contents

Acknowledgements.....	2
Abstract .....	3
Table of Contents .....	6
List of Symbols .....	13
List of Figures .....	15
List of Tables .....	23
1 Introduction .....	26
1.1 Background .....	26
1.2 Aim and objectives .....	30
1.3 Outline of thesis.....	31
2 Literature review .....	32
2.1 Composites overview .....	32
2.2 Heat transfer overview.....	33
2.3 Micromechanics overview .....	35
2.3.1 Orthotropic.....	37
2.3.2 Transversely isotropic.....	38
2.3.3 Isotropic.....	38
2.4 Subsea piping overview.....	39
2.4.1 Heat loss in subsea piping thermodynamics.....	39
2.4.2 Radial heat loss.....	39
2.4.3 Axial temperature gradient.....	41

2.4.4	Cost saving as a result of improved insulation .....	41
2.4.5	Precipitation/Phase change issues with heat loss.....	42
3	Thermal Modelling of Geometries Containing A Single Fibre Type.....	45
3.1	Introduction .....	45
3.2	Literature Review .....	45
3.2.1	Straight Fibres in Thermomechanical Modelling .....	45
3.2.2	Simulation of Randomly Orientated Sinusoidal Fibres .....	54
3.2.3	Mathematical Models for Heat Transfer in Porous Insulation .....	55
3.2.4	Heat Transfer by Convection .....	65
3.2.5	Heat Transfer by Radiation.....	67
3.3	Fibre Geometry Generation.....	70
3.4	Simulation Setup .....	79
3.4.1	Fibre Volume Fraction Simulations .....	81
3.4.2	In-Plane Angle Simulations.....	83
3.4.3	Out-of-Plane Angle Simulations.....	86
3.4.4	Bounding Region Size Simulations.....	89
3.4.5	Fibre Diameter Simulations .....	90
3.5	Results and Discussion .....	93
3.6	Experimental Validation of The Thermal Simulations.....	98
3.6.1	Experimental Setup .....	98
3.6.2	Results Comparison .....	99
3.7	Conclusions.....	102



3.8	Future Work .....	103
4	Mechanical Modelling of Geometries Containing a Single Type of Fibre .....	104
4.1	Introduction .....	104
4.2	Literature Review .....	104
4.2.1	Self-consistent Micromechanics Model (Hill).....	105
4.2.2	Rule of Mixtures .....	106
4.2.3	Eshelby and Mori–Tanaka Method .....	107
4.2.4	Shear Lag Model .....	108
4.2.5	Christensen and Waals .....	108
4.2.6	Manera .....	110
4.2.7	Pan.....	110
4.3	Simulation Setup .....	111
4.4	Results and Discussion .....	113
4.4.1	Volume Fraction .....	113
4.4.2	Fibre Orientation.....	114
4.4.3	Fibre Diameter.....	116
4.5	Theoretical Validation of Mechanical Simulations.....	117
4.5.1	Rule of Mixtures .....	117
4.5.2	Christensen and Waals .....	118
4.5.3	Manera .....	119
4.5.4	Pan.....	120
4.5.5	Validation Summary .....	121

4.6	Conclusions.....	121
4.7	Future Work .....	122
5	Modelling of Fibre Geometries Containing Straight and Curved Fibre Types	123
5.1	Introduction .....	123
5.2	Simulation Setup .....	124
5.3	Results and discussion.....	126
5.4	Conclusions.....	128
5.5	Future Work .....	128
6	Modelling of Aerogel/Fibre Composites.....	130
6.1	Introduction .....	130
6.2	Simulation Setup .....	131
6.2.1	Volume fraction .....	133
6.2.2	Orientation.....	136
6.2.3	Bounding Region Size .....	137
6.2.4	Fibre Diameter.....	138
6.3	Results .....	140
6.4	Validation .....	145
6.5	Conclusions.....	150
6.6	Future Work .....	152
7	Modelling of Aerogel Particle Filled Composites.....	153
7.1	Introduction .....	153
7.2	Literature Review .....	153

7.3	Geometry Generation.....	158
7.4	Simulation Setup .....	163
7.4.1	Thermal Simulations.....	163
7.4.2	Mechanical Simulations.....	164
7.5	Results .....	167
7.6	Conclusions.....	178
7.7	Further Research .....	179
8	Macro Modelling of PiP Centraliser/Spacer .....	179
8.1	Introduction .....	179
8.2	Simulation Setup .....	182
8.3	Test 1: Central Loading .....	190
8.4	Test 2: Loading Across the Gap .....	191
8.5	Results and discussion.....	193
8.6	Conclusions.....	195
9	Conclusions .....	197
10	Future Work .....	200
11	References .....	202
A	Derivations.....	210
A.1	Fibre Final Position.....	210
B	Fibre Rotation .....	210
C	Fibre Generation .....	212
C.1	clipped_line.m .....	224

C.2	clip.m.....	226
C.3	curvedfibres.m.....	226
C.4	lineSegmentIntersect.m.....	228
C.5	polint.m .....	229
C.6	randb.m.....	230
C.7	vecleng.m.....	230
C.8	vecleng3d.m.....	230
C.9	vcvf.m.....	231
D	General MATLAB Codes.....	232
D.1	Temperature Dependant Air Thermal Conductivity .....	232
D.2	Boundary Layer Thickness .....	232
D.3	Closed Variable Fibre Geometries.....	232
D.4	Thermal Models 1-4 .....	233
D.5	Comparison Between Conductivity Models for Conduction, Convection and Radiation.....	234
D.6	Convective Heat Transfer Correlation.....	235
D.7	Mean free Path Mechanics Assessment.....	235
D.8	3D open Nonlinear Fibres.....	236
D.9	Disorder Coefficient.....	237
D.10	Model 5 .....	238
D.11	Optimization of Radiation Model.....	239
D.12	A-B Hybrid Model .....	240

D.13	Model using 3 Distinct Conductivities .....	240
D.14	Stiffness of Hybrid Composition.....	241
E	APDL Codes .....	242
E.1	Summing Fluxes in ANSYS.....	242
F	Extended Simulation Setup .....	243
G	Particle Generation .....	244
G.1	randomspheres.m .....	244
G.2	randomspheresDM.m.....	247
G.3	mrrn.m.....	250
G.4	Other MATLAB Codes Required for Particle Generation .....	252
H	DCPD Results.....	253
I	Nylon Results.....	255

## List of Symbols

Symbol	Definition	Unit
GF	Glass fibre	
PET	Polyethylene terephthalate	
RF	Resorcinol-Formaldehyde	
SI	Silica	
$f$	Fibre property	
$a$	Air property	
$x, y, z, t$	Standard position/time values	
T	Temperature	°C
$\rho$	Density	kg/m <sup>3</sup>
E	Energy	
k	Thermal conductivity	W/mK
$k_{f1}$	Thermal conductivity along fibre	W/mK
$k_{f2}$	Thermal conductivity in the axial direction of a fibre	W/mK
q	Heat flux	W/m <sup>2</sup>
Q	Heat flow rate	W
$\varepsilon$	Direct strain	-
$\gamma$	Shear strain	-
$\nu_p$	Poisson's ratio	-
$\sigma$	Direct stress	Pa
$\tau$	Shear stress	Pa
S	General stiffness coefficient	
C	Inverse of S	

---

$v_f$	Volume fraction	-
L	Length	m
d	Fibre diameter	m
Z	Disorder coefficient	
$\sigma_e$	Extinction coefficient	
E	Young's modulus	Pa
K	Bulk modulus	Pa
L	Blanket thickness	m
$C_p$	Specific heat capacity (constant pressure)	J/(K kg)
$C_v$	Specific heat capacity (constant volume)	J/(K kg)
V	Volume	m <sup>3</sup>
P	Pumping power	W
$a$	Anisotropy	
$\cos(\varphi)^2$	Polar orientation parameter	

---

## List of Figures

Figure 2.1: 2D schematic of the heat transfer setting being discussed; where $T_1$ and $T_2$ are different temperatures being applied to the region. ....	34
Figure 2.2: Schematic illustration of PiP insulation, with a simplified 2D geometry on the left, and a more complex, realistic geometry on the right showing multiple pipes bundled together. ....	40
Figure 2.3: Temperature - Viscosity curves of several types of crude oil. These may be used to estimate the pumping power based on an average oil temperature [15]. .....	41
Figure 3.1: Boundary conditions that were used by Arambakam et al. [1] in order to simulate the heat transfer through their fibrous material.....	46
Figure 3.2: Measured thermal (pseudo-) conductivity, $\lambda$ , of organic fibre boards as function of cubed mean temperature, $TR$ ; from [17]. ....	49
Figure 3.3: The three load cases used by Kusenko, Böhm and Drach[22], originally Figure 2 in their paper. (a) is referred to as Load Case 1, which deals with elongation of the fibre(s); (b) is Load Case 2, which deals with direct force across the cross section of the fibre(s); and (c) is Load Case 3, which covers shear across the cross section of the fibre(s). ....	54
Figure 3.4: Comparative graph showing the individual thermal conductivities as a function of volume fraction, for the case of an air-glass fibre blend. ....	57
Figure 3.5: Representative analogies for conductive heat transfer in nonwoven fabrics showing parallel (Model 1) and series (Model 2) style heat transfer for the case of a 2D air glass fibre blend. ....	57
Figure 3.6: The range of possible values of the general Disorder Model, with the upper and lower values matching exactly the values of models 1 and 2 respectively, as $Z$ varies.....	60



Figure 3.7: Comparison of the thermal conductivity of air when it's treated as a continuous and non-continuous fluid at the length scale of the simulations.....	64
Figure 3.8: Two infinite, flat plates at different temperature (black), with a fluid between them (blue) .....	67
Figure 3.9: Comparison of the magnitude of the components of heat transfer through a blanket, with the convection and radiation being treated as conduction equivalents here. ....	68
Figure 3.10: The results from Figure 3.9 zoomed in so that the difference at high temperatures is clearer, with the Solid Conduction removed; note that the y-axis is no longer logarithmic .....	69
Figure 3.11: An example of a straight fibre (left) and curled fibre (right) geometry once it has been imported into ANSYS Workbench.....	71
Figure 3.12: Fibre insertion algorithm, illustrating how a three-dimensional fibre geometry is created by the algorithm .....	75
Figure 3.13: Approximation of a crimped fibre (blue) as three different line segments (purple, dark red) .....	76
Figure 3.14: Visual comparison of the fibre geometries that are possible to be generated .....	78
Figure 3.15: The standard boundary conditions used for thermal simulations .....	80
Figure 3.16: The variation of the conductivity with respect to the volume fraction ...	94
Figure 3.17 : The effect of varying the max in-plane angle on the thermal conductivity .....	95
Figure 3.18: The effect of varying the max out-of-plane angle on the thermal conductivity.....	96
Figure 3.19: The effect of varying the relative fibre length on the thermal conductivity .....	97
Figure 3.20: The effect of varying the fibre diameter on the thermal conductivity ....	97

Figure 3.21: Comparison of the PET experimental and model results.....	100
Figure 3.22: Comparison of the Glass fibre experimental and model results.....	101
Figure 4.1: Illustration of the boundary conditions applied to the mechanical simulations.....	112
Figure 4.2: Young's modulus of the composite at varying volume fractions of glass fibre .....	113
Figure 4.3: Same set of fibres viewed from two different orientations, with the fibres being visibly uniform in one plane. ....	114
Figure 4.4: Young's modulus against random in-plane orientation .....	115
Figure 4.5: Young's Modulus against the diameter of the fibres used. ....	116
Figure 4.6: Comparison of the FEA results with the Rule of Mixtures and Inverse Rule of Mixtures results.....	117
Figure 4.7: Comparison of the FEA results with the 2D and 3D Christensen and Waals model results .....	119
Figure 4.8: Comparison of the FEA results with the Manera model results .....	120
Figure 4.9: Comparison of the FEA results to the 2D and 3D Pan models. ....	121
Figure 5.1: The effect of varying the mixing ratio on the thermal conductivity.....	126
Figure 5.2: The variation of the conductivity based on the volume fraction.....	127
Figure 6.1: Comparison of the derived lower and upper bounds for the thermal conductivity, where Experimental is the experimentally derived blanket conductivity. ....	132
Figure 6.2: Comparison of the mesh used for the 1.5% sample and the 12% sample .....	135
Figure 6.3: Example of the mesh produced on a fibre in the 1.5% VF simulation ..	135
Figure 6.4: Example of the mesh produced on a fibre in the 12% VF simulation ...	136
Figure 6.5: Comparison of the effect changing the volume fraction has on the four composite materials being studied here .....	140

Figure 6.6: The effect of varying the maximum in-plane angle on the thermal conductivity.....	141
Figure 6.7: The effect of varying the maximum out-of-plane angle on the thermal conductivity.....	142
Figure 6.8: Comparison of the effect the number of elements (top) and the volume fraction (bottom) has on the thermal conductivity in the out-of-plane angle simulations .....	143
Figure 6.9: Comparison of the effect of varying the diameter on the four composites investigated. ....	144
Figure 6.10: Effect of varying the relative fibre length on the conductivity .....	145
Figure 6.11: FEA geometry used for the simulation; the box around the fibres was treated as an aerogel/air mix of appropriate conductivity. ....	147
Figure 6.12: Comparison of the Experimental results with those obtained via the FEA using thermal conductivity values derived from the series and parallel models .....	149
Figure 7.1: Temperature dependence of the thermal conductivity of unreinforced aluminium matrix and 40 vol% particulate-SiC-reinforced aluminium matrix composites and unreinforced matrix for a range of values of SiC particle size, from [67]. ....	156
Figure 7.2: MATLAB scatter graph showing the particle positions after generation in 3D, with the positions being defined in micrometres .....	159
Figure 7.3: 3D geometry produced by the algorithm in Design Modeller .....	160
Figure 7.4: Both graphs show a histogram of the particle size so that the distribution of sizes can be seen; the top is at a volume fraction of 10% and the bottom is at a volume fraction of 40%. ....	161
Figure 7.5: Example temperature plot for the particles, with the air around the particles being hidden, the top face being constrained to 30°C and the bottom face being constrained to 10°C. ....	162

Figure 7.6: The geometry for the 20% VF sample. Note the two lines on the bottom left face that allow the centre to be selected as a point. ....	165
Figure 7.7: Effect on conductivity as the number of pores increases; while also investigating how having 4 larger pores and one small one and 4 small pores and 1 large pore change the conductivity. ....	167
Figure 7.8: Comparison of the effect having more pores has on the Thermal Conductivity .....	168
Figure 7.9: Comparison of the simulated thermal conductivities for three particle sizes with the expected upper and lower bounds for the conductivity, which were found using the rule of mixtures models for parallel and series composites .....	169
Figure 7.10: Same comparison as in Figure 7.9 but with the Nielsen Model and the Geometric Mean Model added.....	169
Figure 7.11: This is comparing how the addition of various fractions of aerogel filler to three kinds of matrix reduces the conductivity, with the conductivities being normalised so that the important thing is the change from 0% filler and not the absolute change. ....	171
Figure 7.12: Comparison of the experimental and modelling results for the Isophthalic Polyester material.....	172
Figure 7.13: Comparison of the simulation and experimental results for Polyester 41 .....	173
Figure 7.14: Comparison of the simulation and experimental results for Polyester 41 .....	173
Figure 7.15: Comparison of the simulation and experimental results for Ortho 40.	174
Figure 7.16: Comparison of simulation and experimental data for the carbon fibre composites .....	175
Figure 7.17: Comparison of simulation and experimental data for the aramid composites .....	176

Figure 7.18: Modelling results indicating how the mechanical properties change as the volume fraction changes .....	177
Figure 7.19: Normalised Young's modulus against volume fraction .....	178
Figure 8.1: The load position in the two tests; with Test 1 on the left and Test 2 on the right. The experimental test rig uses a circular arc that covers approximately one quarter of the centraliser to do the test, as opposed to a point load as these images suggest.....	181
Figure 8.2: Image of the failed centraliser, showing that it has broken at the heating wire cut out .....	182
Figure 8.3: Side on profile of the centraliser.....	183
Figure 8.4: Centraliser at an angle so the EHTF cut outs can be seen more clearly; note also the lines on the base that allow the center of those faces to be selected. ....	183
Figure 8.5: Experimental setup for the centraliser testing; showing the circular loading bracket at the top and the inner pipe and partial outer pipe, both of which act as supports for the centraliser during testing. ....	186
Figure 8.6: The bilinear isotropic hardening data used by the DCPD during the simulation .....	188
Figure 8.7: The multilinear isotropic hardening data used for the Nylon during testing .....	188
Figure 8.8: Diagram showing where the 250kN force was applied in Test 1; with the inner pipe also being fixed to no displacement in the normal direction (i.e., towards the centre of the pipe).....	190
Figure 8.9: Diagram showing where the lower centraliser was constrained to not displace in the x-direction to replicate being supported by the lower pipe. ....	190

Figure 8.10: Diagram showing where the force of 250kN was applied to the centraliser and that the inner pipe has been constrained to 0m of displacement in the normal direction.....	191
Figure 8.11: Diagram showing how the other side of the centraliser is constrained, with half being fixed to 0m of displacement in the y-direction, and the bolt flange being fixed fully to prevent rigid body motion .....	192
Figure 8.12: Comparison of the maximum stress in the centraliser (probed value); and the stress distribution at the cut out where the centraliser is thinnest; with the nylon centraliser on the left, and the DCPD one on the right. The slice used for this comparison is not at exactly the same place, which contributes to some of the differences in the geometry and stress in the two images. ....	194
Figure 8.13: The DCPD centraliser in Test 2 load conditions, angled to allow a better view of the overall geometry than in Figure 8.12.....	195
Figure 14: The initial box generation settings (left) and the slice settings (right) for fibre importing.....	243
Figure 15: The deformation (left) and stress (right) in the DCPD centraliser in Test 1 loading conditions using the bilinear material model. Note that the deformation is negative because the x-direction is outwards.....	253
Figure 16: The deformation (left) and stress (right) in the DCPD under Test 1 loading conditions using the trilinear material model. ....	253
Figure 17: The deformation (left) and stress (right) for the DCPD in Test 2 Conditions using the bilinear material model. ....	254
Figure 18: The deformation (left) and stress (right) for the Nylon in Test 1 Conditions using the multilinear material model.....	255
Figure 19: The deformation (left) and stress (right) for the Nylon in Test 1 Conditions using the bilinear material model. ....	255

Figure 20: The deformation (left) and stress (right) for the Nylon in Test 1 Conditions using the trilinear material model. ....256

Figure 21: The deformation (right) and stress (left) in the Nylon under Test 2 load conditions using the multilinear material model. ....256

Figure 22: The deformation (left) and stress (right) for the Nylon in Test 2 Conditions using the bilinear material model. ....257

## List of Tables

Table 3.1: Z-Value comparison table indicating what various Z values mean.....	61
Table 3.2: General thermal properties.....	80
Table 3.3: Geometry and mesh properties of the straight fibre volume fraction models .....	81
Table 3.4: Geometry and mesh properties of the sinusoidal fibre volume fraction models.....	81
Table 3.5: Mesh settings used in both the straight and sinusoidal fibre volume fraction modelling .....	82
Table 3.6: Geometric properties of the straight fibre simulations in the 'In-Plane Angle' Case .....	83
Table 3.7: Geometric properties of the sinusoidal fibre simulations in the 'In-Plane Angle' Case .....	84
Table 3.8: Mesh properties and settings for the straight fibres .....	85
Table 3.9: Mesh properties and settings for the sinusoidal fibres .....	85
Table 3.10: The geometric and mesh properties of the straight fibre simulations ....	86
Table 3.11: The geometric and mesh properties of the sinusoidal fibre simulations	87
Table 3.12: The mesh settings used in the straight fibre simulations.....	88
Table 3.13: The mesh settings used in the sinusoidal fibre simulations .....	88
Table 3.14: The geometric and meshing properties of the straight fibre samples used in the Bounding Region size simulations.....	89
Table 3.15: Mesh settings used in the bounding region size investigation .....	90
Table 3.16: Geometric properties of the straight fibre simulations .....	91
Table 3.17: Geometric properties of the sinusoidal fibre simulations.....	91
Table 3.18: Mesh properties used in the straight fibre simulations .....	92
Table 3.19: Mesh properties used in the sinusoidal simulations.....	93



Table 3.20: Geometric properties for the thermal validation .....	99
Table 3.21: Mesh settings used for the thermal validation.....	100
Table 4.1: Material properties used in the mechanical simulations.....	111
Table 4.2: Boundary conditions used in the mechanical simulations .....	112
Table 5.1: The properties of the geometries used in the mixed fibre modelling simulations.....	124
Table 5.2: Mesh settings used in the mixed fibre modelling simulations.....	125
Table 5.3: The number of elements present in each mixed fibre simulation .....	125
Table 6.1: Aerogel thermal properties .....	<b>Error! Bookmark not defined.</b>
Table 6.2: The volume fraction simulations' geometry parameters.....	133
Table 6.3: Volume fraction mesh settings .....	134
Table 6.4: Element count for the volume fraction simulations .....	134
Table 6.5: Geometry and mesh properties of the in-plane orientation simulations.	136
Table 6.6: The geometry and mesh parameters for the out-of-plane angle simulations .....	137
Table 6.7: Geometric properties for the bounding region size simulations .....	138
Table 6.8: Simulation parameters for the diameter simulations .....	139
Table 6.9: Summarised meshing parameters for the diameter simulations .....	139
Table 6.10: Experimental results for the PET fibre/RF aerogel composite blanket	146
Table 6.11: Mesh settings used in the validation.....	148
Table 6.12: Thermal validation results .....	150
Table 7.1: The properties of the geometries used in the mechanical simulations ..	165
Table 7.2: Mesh settings for the mechanical particle simulations .....	166
Table 7.3: Mesh information for the individual simulations .....	166
Table 8.1: Mechanical material properties for the intial centraliser modelling .....	185
Table 8.2: The mesh properties used for the first set of centraliser simulations ....	186

Table 8.3: Elastic material properties used in the simulation; with the bulk and shear moduli being calculated from them (and hence not shown).....187

Table 8.4: Summarised simulation results; showing that the DCPD generally has bigger stresses and strains than the Nylon does; with the difference in strain in both cases being very significant .....193

# 1 Introduction

## 1.1 Background

Low temperatures on seabed and long transport distances currently pose serious challenges to the oil and gas industry for crude oil extraction in ultra-deep water (that is, water deeper than roughly 1500m). Pipelines used in the operation of oil exploitation need to be thermally insulated to limit heat loss and prevent wax being deposited on the inside of the pipeline. If the temperature inside the pipe gets too low, these deposits can begin to build up and eventually clog the pipeline [2]. Current thermal insulation systems typically involved multi-layered structures made of several materials with substantial thickness in order to meet required insulation and flow assurance. There also exists a monolithic insulation by using pre-moulded rigid foam (e.g. polyurethane) around the inner pipeline. This reduces the complexity of installation process compared to the multi-layered insulation structures but have the disadvantage of performance drop. Nevertheless, all these insulation materials typically can only achieve a thermal conductivity of 25 – 30 mW/mK. This inevitably leads to thick pipes and overall heavy production infrastructures, which cause great difficulties and a tremendous cost in transportation and installation of the pipelines. Being able to reduce the diameter of the pipes required allows a significantly greater amount of pipe to be transported on board a ship at a time and hence for the pipe to be laid faster. Since the hire of ships and crew is one of the big expenses in laying new pipelines, reducing the cost here leads to large savings overall.

Centralises – spacer between inner and outer pipes - must be used in subsea oil and gas applications and they work via an outer pipe that protects the pipe and prevents the sea water degrading the performance of the insulation. In order to maintain the separation between the inner and the outer pipe, centralisers are used periodical, which are typically made of a low conductivity material that is significantly stronger

than the insulation. Unfortunately, this has a trade off in that they are typically around 10 times more heat conducting than the insulation is.

The use of aerogel-based composites is becoming more popular here as a result. The extremely low thermal conductivity and density of aerogel mean that less insulation is required, and lighter piping can be manufactured. As bulk aerogel tends to be very fragile, increasing the flexibility means that manufacturing and transporting a pipe in pipe insulation system becomes easier. This work largely looks at two ways to create such a composite: aerogel/fibre composites, where the aerogel is used as a matrix material around a fibre blanket, and particulate aerogel composites, where particles of aerogel are dispersed through a plastic matrix.

The World's insulation markets are a profitable industrial trade; the global market for thermal insulation materials is worth an estimated US\$ 38 billion and annual projected growth rate around 8%. There currently exists a worldwide need for superinsulation solutions enabled by novel super-insulating materials – aerogels. The aerogels sector, being an emerging market, has seen tremendous annual growth rates from 50 to 70% in recent years and similar developments are expected in the near future. Recent market studies project worldwide sales of aerogel products in excess of US\$ 500 million. The rapid growth and development in improved insulation materials is driven by the political agenda in energy and environment, economic sustainability, and ever-increasing urbanisation.

Unlike conventional insulation materials (for example, polyurethane foam or glass wool), an aerogel is a nano-porous solid material consisting of an open-cell network with numerous exceptional characteristics, including extremely low thermal conductivity ( $\sim 0.015$  W/mK), low bulk density ( $\sim 0.1$  g/cm<sup>3</sup>), and high specific surface area ( $\sim 1000$  m<sup>2</sup>/g). This unique combination of characteristics has enabled NASA, in the late 1990s, to employ aerogels on three outer space missions utilizing their super-

insulating properties and large specific surface areas. Since then, aerogels have attracted a lot of R&D attention and use has now spread into many areas such as oil & gas, construction, automotive, and aerospace sectors. The material currently dominating the aerogel insulation market is silica-based and requires a relatively expensive synthetic route to produce high quality products. Silica-based aerogels also suffer from poor mechanical properties (including low strength and toughness) inhibiting their use in structural applications. Moreover, the fragile silica aerogels tend to break down during handling and shed small dust particles causing health and safety issues. More recently, a new polyimide aerogel has been developed at NASA's Glenn Research Centre and represents a revolutionary advance over the fragile silica aerogels currently on the market. This technology has been exclusively licenced by our industrial sponsor, Blueshift International Materials (Blueshift), who has set the target of developing novel polymer-aerogels for the thermal insulation market.

Table 1.1 summarises the thermophysical properties of commonly seen insulation materials, and the comparison of thermal performance is plotted in Figure 1.1. It is well indicated that vacuum insulation panel (VIP) provides the best thermal insulation performance among all listed materials, with almost ten times better than conventional insulation materials such as EPS or mineral wool. Similar to foamy organic materials such as EPS or extruded polystyrene (XPS), its low thermal stability may limit its use in many applications. Fibrous inorganic materials such as glass wool and rock wool exhibit superior thermal stability with excellent fire resistance grade. However, their thermal insulation performance is lower than the foamy organic materials. Silica aerogels possess excellent thermal insulation and stability; however, one of the significant drawbacks is the cost, as shown in Table 1.2. Nevertheless, superinsulation materials such as VIP and silica aerogels still offer the great potential of dominating the global insulation market in the upcoming future.

Table 1.1 Thermophysical properties of the commonly seen insulation material,  $\lambda$ : thermal conductivity at 20°C, dry material;  $\rho$ : bulk density;  $T_{max}$ : maximum service temperature;  $\sigma$ : compressive stress at 10% deformation except for foam concrete and foam geopolymer

Insulation material	$\lambda$ [mW/m.K]	$\rho$ [kg/m <sup>3</sup> ]	$T_{max}$ [°C]	$\sigma$ [kPa]	Fire resistance grade
VIP	4-8	65-300	90	45-120	-
Silica aerogels	4-20	3-350	750	0-5000	A
PUR-PIR	19-30	25-100	120	100-500	B2
XPS	25-35	20-80	75	150-700	B2
EPS	29-41	10-50	80	60-260	B2
Glass wool	30-46	8-150	500	15-80	A
Rock wool	33-46	13-240	750	15-80	A
Foam glass	38-61	100-200	>400	400-1600	A
Expanded perlite	40-60	50-400	1260-1343	-	A
Foam concrete	80-250	200-700	1000	>100	A
Foam geopolymer	>58	>200	1300	>330	-




Figure 1.1 Thermal conductivity of commonly seen insulation materials (dry, at 20°C)

Table 1.2 Specific material cost of an insulating layer with an R-value of 10 m<sup>2</sup>KW<sup>-1</sup>

Insulation material	$\lambda$ [mW/m.K]	Cost [€/m <sup>2</sup> ]
VIP	8	247
Silica aerogels	12	547
PUR-PIR	25	59
XPS	33	46
EPS	35	32
Glass wool	38	43
Rock wool	40	36
Foam glass	50	50

## 1.2 Aim and objectives

The main aim of this research was to develop predictive models so that the thermal, and some of the mechanical properties, of aerogel containing composite materials could be determined based on their structure. Specifically, randomly arranged fibre-aerogel composites and aerogel particle-based composites were investigated, where the degree of randomness in the fibres could be controlled. The fibres investigated came in two variants: straight fibres and sinusoidal fibres, which could be generated either alone or mixed together. The particles considered were all spherical in shape, though the diameter could be allowed to vary using either normal or uniform distributions, or alternately using a more complex univariate multimodal Gaussian distribution.

The main section of the work was investigating modelling randomly orientated fibrous mats, with a focus on the thermal properties of them and how the structure of the mats affects the conductivity of heat through the mat. The mechanical properties of the mat,

both in terms of the Young's modulus and the overall composite strength, were also considered, albeit in less detail

### **1.3 Outline of thesis**

This chapter serves as an introduction to the aims, objectives and outline of the thesis. The next chapter (Chapter 2) covers a literature review of the physics and engineering underpinning the thesis. The research topics covered by the thesis are discussed in depth in Chapters 3-8. Each of these features a more in-depth literature review looking at chapter specific research. Chapter 3 investigates the thermal performance of single fibre type blankets, with a focus on how the parameters of the blanket affect the thermal conductivity through the blanket. It also contains a significant section underpinning the modelling system used in Chapters 4-6. Chapter 4 covers the mechanical performance of single fibre type blankets, and how the design of the blanket affects the stiffness of it. Chapter 5 covers the effect of mixing the two fibre types together, and is relatively concise as a result. Chapter 6 then goes on to investigate the performance of the blanket when used in aerogel based composites. Chapter 7 introduces aerogel particle filled composites, with a focus on modelling spherical particles. It features a second algorithm to generate a random geometry. Finally, the parameters investigated in the preceding chapters are used to investigate a bulk model of the centraliser, and determine how it behaves. Final conclusions are present in Chapter 9, and some areas of future work are laid out in Chapter 10. Following this are a series of Appendices containing the MATLAB codes used in various stages of the research.



## **2 Literature review**

The literature review presented in this section is intended to provide a brief overview of several common aspects underling all subsequent chapters. There is a more focused literature review specifically dedicated to each chapter in accordance to the topic in that chapeter.

### **2.1 Composites overview**

Composite materials are materials comprised of two or more different materials. For example, glass fibres are commonly used as reinforcements for polymers, where the high tensile strength of the fibres improves the mechanical properties of the resin significantly, while adding minimal cost and weight to the material. Fibres, particles, and platelets are all commonly used to produce composites, but a wide range of different kinds of inclusion exist. Composite materials are common in nature, as seen in materials like wood (cellulose fibres within a matrix of lignin) and bones, composed of a honey-comb like matrix arranged in a range of ways depending on which bone it is [3]. Development of composite materials began in the 1940s, when Owens Corning began patenting various methods to create glass fibres [4-7], before picking up steam in the 1960s when 'Composite Materials' became a field of study on its own [3].

Fibre based composites can be manufactured to be either continuous or discontinuous in nature; where continuous fibres are typically either woven or unidirectional; and discontinuous are typically chopped or formed into a mat [8]. Often continuous fibre composites are constructed into laminates featuring discrete layers, whereas chopped mats are typically found in non-laminated composites [9]. Glass fibres are the most commonly used fibre reinforcement [10], though carbon fibres are very common too. Alternative reinforcements to fibres include plates and particles;

with particles being particularly common in attempts to increase the thermal conductivity of the composite in metal based composites [11].

Aerogels are extremely porous solid gels originally developed by Kistler in 1932 [12], using a process of supercritical drying. Silica aerogels are the most widely researched and commercially available aerogel, though a wide range of other materials have been used including metal foams and hydrocarbons [13]. The high porosity of the aerogel means they have extremely low thermal conductivities, typically below 20mW/mK, and densities in the sub 0.2gcm<sup>-3</sup> range [13].


The most common application for aerogels is thermal insulation, but they have also been used in a wide range of applications from acoustic insulation to catalysts [14].

## **2.2 Heat transfer overview**

This work focusses on the heat transfer through thermal insulation. While some consideration is done for the mechanical performance of it, as the insulation is intended to be used in pipe-in-pipe insulation, understanding the thermal properties prove to be more important. In order to understand the importance of the thermal properties, an understanding how heat transfer is necessary.

There are three main mechanisms for heat transfer: conduction, convection, and radiation. Thermal conduction is an internal mechanism of heat transfer which occurs when particles collide with each other within a body, or through the transfer of electrons in, for example, a metal. Convection occurs within fluids and occurs when the fluid physically moves heat away from the higher temperature location. Radiation is heat transfer carried out by atoms emitting photons of energy which are then reabsorbed.

Figure 2.1: 2D schematic of the heat transfer setting being discussed; where T1 and



$T_2$  are different temperatures being applied to the region.

If we consider a two dimensional system as shown in Figure 2.1, and apply the general thermodynamic energy equation in two dimensions, shown in Equation (1), then it can be shown that the temperature distribution across the purely conductive system is linear. This result will be used several times in the development of theory. Fourier's law will also be used, without proof, in its one and two-dimensional constitutive form.

$$\frac{\partial \rho E}{\partial t} + \frac{\partial \rho u E}{\partial x} + \frac{\partial \rho v E}{\partial y} = \rho S + k_x \frac{\partial^2 T}{\partial x^2} + k_y \frac{\partial^2 T}{\partial y^2} - P \left[ \frac{\partial u}{\partial x} + \frac{\partial v}{\partial y} \right] + \phi \quad (1)$$

The heat transfer due to convection within a fibre composite is generally negligible (Section 3.2.4 discusses this in more depth), since fibrous insulations are usually designed to limit the bulk flow of fluid. As a result, the velocity terms are both approximately 0 and can be neglected to simplify Equation (1). In addition, it can be assumed there is no internal heat source, and that the temperature variance in the  $y$  direction can be neglected. This means that Equation (1) can be simplified to Equation (2). Integrating this produces Equation (3), which in turn can be integrated to produce Equation (4).

$$\frac{\partial^2 T}{\partial x^2} = 0 \quad (2)$$

$$\int \frac{\partial^2 T}{\partial x^2} dx = \frac{\partial T}{\partial x} = C \quad (3)$$

$$\int \frac{\partial T}{\partial x} dx = T(x) = Cx + D \quad (4)$$

By applying boundary conditions to Equations (3) and (4), Equation (5) can be derived.

$$T(x) = T_1 - (T_1 - T_2) \frac{x}{L} \quad (5)$$

Using this knowledge, the temperatures may be related to the heat flux using Fourier's law of 1D conduction:

$$T(x) = T_1 - \left(\frac{q}{K}\right) x \quad (6)$$

These results are used to develop thermal boundary conditions and theory in adiabatic conditions. Therefore, heat transfer within a nonwoven may be approximated by a conduction only model. This model may be decomposed into conduction in fibres, and conduction in interstitial air and solid materials such as aerogel.

### 2.3 Micromechanics overview

The mechanical properties of fibrous composites are primarily influenced by the mechanical properties of the constituent materials. Embedding fibres into a matrix, typically polymer but ceramic and metal are relatively common too, enhances the mechanical properties of the matrix due to the added strength of the fibres.

This section provides a review of the mechanics of fibrous composites, specifically from a micromechanics perspective. Describing the mechanical properties of fibrous composites is challenging due to the anisotropy and inhomogeneity, which mean the strength and stiffness at any point in the specimen will vary depending on which

component the point is in. It can be more practical to identify average properties of the material using a law of mixture which allow it to be treated as a homogeneous material, since this reduces the complexity of the situation significantly. The anisotropy will remain, since they are a factor of the material properties being dependent on the fibre arrangement within the composite, and the fibres are strongest in the transverse direction.

The 3D elastic behaviour of isotropic linear material can be described by Hooke's law using the compliance matrix shown in Equation (7). Inverting this matrix produces the stiffness matrix, shown in Equation (8). This shows how the stresses and strains within the material are related.

$$\begin{bmatrix} \varepsilon_x \\ \varepsilon_y \\ \varepsilon_z \\ \gamma_{yz} \\ \gamma_{zx} \\ \gamma_{xy} \end{bmatrix} = \frac{1}{E} \begin{bmatrix} 1 & -\nu_p & -\nu_p & 0 & 0 & 0 \\ -\nu_p & 1 & -\nu_p & 0 & 0 & 0 \\ -\nu_p & -\nu_p & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 + \nu_p & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 + \nu_p & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 + \nu_p \end{bmatrix} \begin{bmatrix} \sigma_x \\ \sigma_y \\ \sigma_z \\ \tau_{yz} \\ \tau_{zx} \\ \tau_{xy} \end{bmatrix} \quad (7)$$

$$\begin{bmatrix} \sigma_x \\ \sigma_y \\ \sigma_z \\ \tau_{yz} \\ \tau_{zx} \\ \tau_{xy} \end{bmatrix} = \frac{E}{(1 - 2\nu)(1 + \nu)} \begin{bmatrix} 1 - \nu & \nu & \nu & 0 & 0 & 0 \\ \nu & 1 - \nu & \nu & 0 & 0 & 0 \\ \nu & \nu & 1 - \nu & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 - 2\nu & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 - 2\nu & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 - 2\nu \end{bmatrix} \begin{bmatrix} \varepsilon_x \\ \varepsilon_y \\ \varepsilon_z \\ \gamma_{yz} \\ \gamma_{zx} \\ \gamma_{xy} \end{bmatrix} \quad (8)$$

In the anisotropic case, however, the stress and strains vary a lot more: the isotropic case features no interactions between the shear and direct stress, whereas in the anisotropic case, the shear moduli can affect the direct strain and vice versa. This is can be seen in the compliance (Equation (9)) and stiffness (Equation (10)) matrices for anisotropic materials.

$$\begin{bmatrix} \varepsilon_x \\ \varepsilon_y \\ \varepsilon_z \\ \gamma_{yz} \\ \gamma_{zx} \\ \gamma_{xy} \end{bmatrix} = \begin{bmatrix} S_{11} & S_{12} & S_{13} & S_{14} & S_{15} & S_{16} \\ S_{21} & S_{22} & S_{23} & S_{24} & S_{25} & S_{26} \\ S_{31} & S_{32} & S_{33} & S_{34} & S_{35} & S_{36} \\ S_{41} & S_{42} & S_{43} & S_{44} & S_{45} & S_{46} \\ S_{51} & S_{52} & S_{53} & S_{54} & S_{55} & S_{56} \\ S_{61} & S_{62} & S_{63} & S_{64} & S_{65} & S_{66} \end{bmatrix} \begin{bmatrix} \sigma_x \\ \sigma_y \\ \sigma_z \\ \tau_{yz} \\ \tau_{zx} \\ \tau_{xy} \end{bmatrix} \quad (9)$$

$$\begin{bmatrix} \sigma_x \\ \sigma_y \\ \sigma_z \\ \tau_{yz} \\ \tau_{zx} \\ \tau_{xy} \end{bmatrix} = \begin{bmatrix} C_{11} & C_{12} & C_{13} & C_{14} & C_{15} & C_{16} \\ C_{21} & C_{22} & C_{23} & C_{24} & C_{25} & C_{26} \\ C_{31} & C_{32} & C_{33} & C_{34} & C_{35} & C_{36} \\ C_{41} & C_{42} & C_{43} & C_{44} & C_{45} & C_{46} \\ C_{51} & C_{52} & C_{53} & C_{54} & C_{55} & C_{56} \\ C_{61} & C_{62} & C_{63} & C_{64} & C_{65} & C_{66} \end{bmatrix} \begin{bmatrix} \varepsilon_x \\ \varepsilon_y \\ \varepsilon_z \\ \gamma_{yz} \\ \gamma_{zx} \\ \gamma_{xy} \end{bmatrix} \quad (10)$$

Through consideration of the symmetry of the matrix and by relating the constants to the elastic and shear moduli, it can be reduced to 21 elastic constants, shown in Equation (11). Any elastic material with 21 independent constants is defined as fully anisotropic, and the behaviour can be characterised if all of them are known. This requires empirical methods to determine the values of the constants. The matrix can be simplified further if the material possesses some elastic symmetry.

$$C = \begin{bmatrix} C_{11} & C_{12} & C_{13} & C_{14} & C_{15} & C_{16} \\ C_{12} & C_{22} & C_{23} & C_{24} & C_{25} & C_{26} \\ C_{13} & C_{23} & C_{33} & C_{34} & C_{35} & C_{36} \\ C_{14} & C_{24} & C_{34} & C_{44} & C_{45} & C_{46} \\ C_{15} & C_{25} & C_{35} & C_{45} & C_{55} & C_{56} \\ C_{16} & C_{26} & C_{36} & C_{46} & C_{56} & C_{66} \end{bmatrix} \quad (11)$$

### 2.3.1 Orthotropic

Orthotropic materials have three planes of symmetry across their primary axes; and are often a good way to represent unidirectional laminated composites where the material behaviour through the plies, along the fibres and across the fibres are all uniquely different. This allows the stiffness matrix to be reduced to twelve components, as shown in (12). The direct stress components remain coupled to each other; but the shear components are not, with shear strains only being produced from shear stresses.

$$C = \begin{bmatrix} C_{11} & C_{12} & C_{13} & 0 & 0 & 0 \\ C_{12} & C_{22} & C_{23} & 0 & 0 & 0 \\ C_{13} & C_{23} & C_{33} & 0 & 0 & 0 \\ 0 & 0 & 0 & C_{44} & 0 & 0 \\ 0 & 0 & 0 & 0 & C_{55} & 0 \\ 0 & 0 & 0 & 0 & 0 & C_{66} \end{bmatrix} \quad (12)$$

### 2.3.2 Transversely isotropic

If the material has the same properties in only one plane, then the behaviour is transversely isotropic. A single ply of a unidirectional fibre composite is often transversely isotropic, since the properties in the axial direction of the fibres are the same due to the symmetry of the fibres, whereas the transverse direction has significantly different properties. In this case, the stiffness matrix can be reduced to five constants of elasticity, as shown in Equation (13).

$$C = \begin{bmatrix} C_{11} & C_{12} & C_{13} & 0 & 0 & 0 \\ C_{12} & C_{11} & C_{13} & 0 & 0 & 0 \\ C_{13} & C_{13} & C_{33} & 0 & 0 & 0 \\ 0 & 0 & 0 & C_{44} & 0 & 0 \\ 0 & 0 & 0 & 0 & C_{44} & 0 \\ 0 & 0 & 0 & 0 & 0 & C_{11} - C_{12} \end{bmatrix} \quad (13)$$

### 2.3.3 Isotropic

Isotropic materials have properties which are independent of the orientation and the same in any direction. Most bulk materials, like glass and metals, are isotropic in nature. In this case, their elastic behaviour can be described using two elastic constants, as shown in Equation (14).

$$C = \begin{bmatrix} C_{11} & C_{12} & C_{12} & 0 & 0 & 0 \\ C_{12} & C_{11} & C_{12} & 0 & 0 & 0 \\ C_{12} & C_{12} & C_{11} & 0 & 0 & 0 \\ 0 & 0 & 0 & C_{11} - C_{12} & 0 & 0 \\ 0 & 0 & 0 & 0 & C_{11} - C_{12} & 0 \\ 0 & 0 & 0 & 0 & 0 & C_{11} - C_{12} \end{bmatrix} \quad (14)$$

## **2.4 Subsea piping overview**

This section examines in further detail the design considerations that apply to subsea piping. These considerations are what motivated the work carried out here, and in turn, were used to determine what was simulated.

### **2.4.1 Heat loss in subsea piping thermodynamics**

The heat loss in subsea piping is a complex process which involves a range of fluid behaviours, from the effect of the movement of the sea around the pipeline to the effect of phase changes within the oil itself, including the effects of precipitation building up inside the pipeline. As a simplification, the properties looked at were considered to be temperature independent, as the heat loss can be assumed to take place within a liquid phase only. This assumption is plausible since the liquid phase typically have significant larger diameter than the wall of the pipe.

### **2.4.2 Radial heat loss**

As discussed in Section 2.2 the heat loss of a subsea pipe may be examined using Fourier's law in radial form. By performing an energy balance on a circular oil control volume, the rate of change of temperature in the internal fluid per unit length may be directly related to the geometry. A more realistic scenario of a dual walled subsea pipe is considered, in which the inner wall contains the fluid, and the external wall supplies the pressure boundary for the sea. The cavity between the internal and external walls is filled with thermal insulation.

The rate of change of temperature per unit length may be found using Equation (15) [15].



$$\frac{dT}{dL} \approx \frac{2}{L_p \rho_o \rho C_p} \left( \frac{T_1 - T_4}{\frac{1}{h_1} + r_1 \left( \frac{\ln\left(\frac{R_2}{R_1}\right)}{K_1} + \frac{\ln\left(\frac{R_3}{R_2}\right)}{K} + \frac{\ln\left(\frac{R_4}{R_3}\right)}{K_3} + \frac{1}{R_4 h_2} \right)} \right) \quad (15)$$

It is assumed, in the derivation of this equation, that pipelines are sufficiently long that all heat loss is eventually dissipated radially. Therefore, there is no need for a second coupling equation to account for heat loss in the axial direction.




Figure 2.2: Schematic illustration of PiP insulation, with a simplified 2D geometry on the left, and a more complex, realistic geometry on the right showing multiple pipes bundled together.

This geometry is based on the Subsea 7 S.A. cable bundle geometry, simplified into a single polar setup, shown in Figure 2.2.

### 2.4.3 Axial temperature gradient

Based on the rate of change of temperature per unit length, the final temperature of the crude oil may be found using Equation (16), assuming a linear temperature distribution – shown to be appropriate in Section 2.4.2.

$$T_5 = -\frac{dT}{dL}L + T_1 \quad (16)$$

### 2.4.4 Cost saving as a result of improved insulation




Figure 2.3: Temperature - Viscosity curves of several types of crude oil. These may be used to estimate the pumping power based on an average oil temperature [15].

Using the temperatures calculated previously, the required pumping power may be found using crude oil viscosity curves, shown in Figure 2.3, and Equation (17) [16].

$$P = Q\Delta P = \left(\pi \frac{D^2}{4} \bar{v}\right) \left(\frac{32\mu L \bar{v}}{D^2}\right) = 8\pi L \bar{v}^2 \mu(T_1, T_5) \quad (17)$$

With a functional estimate of the pumping power required for a particular insulation, the cost benefit of that particular analysis across a well lifetime may be estimated.

A basic cost model, shown in Equation (18), accounts for the reduced pump power consumption cost, the reduced cost associated with a lower failure rate, and the potential for increased capital costs associated with a well overhaul, to measure the lifetime influence of improved insulation [16].

$$\Delta E = t_w P_L (\varepsilon P L_P P_E + \Delta F_r F_E) - C_C \quad (18)$$

#### **2.4.5 Precipitation/Phase change issues with heat loss**

As mentioned in the previous section, any major heat loss to pumped crude oil not only increases the viscosity, but also causes chemical issues, which may increase the pipeline and pump failure rate. These issues are summarised below.

##### *2.4.5.1 Solids deposition*

Solid deposition occurs when the temperature drops below the hydrate/wax transition temperature and hydrates change phase, attaching to a pipe inner wall. In any single instance this is not problematic; however, it quickly leads to a build-up of other deposits such as grit, scale, and asphalt.

This build up can lead to blockages and create a turbulent flow around the deposition vortex, accelerating erosion. Additionally, depending on the composition of the deposited hydrates, this can also lead to accelerated corrosion of the internal pipe wall.

##### *2.4.5.2 Hydrate deposition*

Hydrates are hard crystalline, water based, ice-like substances, with a structure which is dependent on the relative gas composition. Circumstances leading to their formation are complex to predict. Crucially, hydrate formation leads to an increased pressure drop and a reduced delivery of water at outlet flow. This leads to greater pumping loads and increased erosion on pump components and pipe walls.

The circumstances under which hydrate formation generally occurs are high-pressure, relatively low-temperature, high water concentration (>10%), in a completely turbulent highly salient flow. As a result, the most effective way to reduce hydrate deposition is to reduce the heat loss.

#### *2.4.5.3 Wax and paraffin deposition*

Wax and paraffin deposits may be hard or soft, depending on the relative composition of the extracted crude. Unlike hydrate deposition, which typically occurs at the end of a wells cycle, wax/paraffin deposits occur during the initial oil extraction. They are characterised by a large pressure drop and can be formed in crude oil and associated liquid condensates.

There are a limited number of theoretical models to predict wax and paraffin deposition conditions. Instead, experimental tests to determine the Wax Appearance Temperature (WAT) / cloud point of stock oil must be performed. A typical WAT of crude oil is around 50°C, therefore keeping the working fluid at a temperature in excess of this is beneficial.

#### *2.4.5.4 Scale deposition*

Scale deposits are generally rigid mineral deposits and occur due to solubility changes of brine undergoing temperature or pressure change. Normal scale deposits consist of calcium, magnesium, and salt scales. There are also scales due to the incompatibility of waters, for example injected seawater or aquifer water, which generally causes sulphate scaling. However, these are less dependent on temperature compared to those due to solubility changes.

## 2.5 Summary

It has shown that there is an emerging demand for high-performance insulation materials that can be adopted in a wide range of applications. The current high-

performance insulation strategies are constrained by a few options commercially available including vacuum insulated panel and rigid porous plastic. Aerogels present an exciting avenue in the field of passive heat insulation but their implementation has not been extensively explored in different applications. The deep-sea operations (e.g. oil & gas extraction) have provided a timely opportunity to study key material design parameters underpinning some novel concept of aerogel composites including fibre-reinforced aerogel and aerogel particle filled polymer.

A number of common technical aspects involved in the development of aerogel composites for thermal applications. It spans from basic thermodynamic framework to design principles in anisotropic materials. It is can be seen that the literature in this area is quite scarce because in part aerogel composite is still a largely under development and material internal structure is complicated by dissimilar materials constitution with a significant amount of nano-to-micro scale pores randomly dispersed throughout the solid phase.

## **3 Thermal Modelling of Geometries Containing A Single Fibre Type**

### **3.1 Introduction**

This chapter investigates the properties of single fibre type blends of composite blankets. The aim is to create a modelling tool with the ability to manipulate individual fibre characteristics (e.g. size, shape, and composition) as well as the fibre content (e.g. fibre volume fraction) in order to form a defined fabric microstructure. Such fabric microstructure can then be used to investigate the effect of fibre-related parameters on fabric thermal and mechanical performance. This work looks at two kinds of fibre: straight fibres and sinusoidal fibres.

These kinds of fibre geometry are often used as thermal insulation, as the relatively large pores within them are good at capturing a large volume of air, and preventing said air from undergoing significant bulk movement, thus preventing convective heat transfer within the fibre mat [16].

This chapter deals with each fibre type being individually modelled, with a mixture of the two being considered in Chapter 5; and the fibres surrounded by aerogel is investigated in Chapter 6. The thermal conductivity is being investigated, with a focus on predicting the optimal fibre properties to minimise the thermal conductivity while minimising the impact on the mechanical properties. In this work, the straight fibres are generally modelled as being made of glass fibre, while the sinusoidal fibres are modelled as being made of PET fibres.

### **3.2 Literature Review**

#### **3.2.1 Straight Fibres in Thermomechanical Modelling**

Work on predicting the properties of fibrous insulation has been ongoing, with a large focus on deriving analytical and empirical relations between the properties,

particularly the volume fraction, of a fibrous mat and the thermal conductivity of it, since the 1950s. While initially the focus was on deriving empirical and analytical relations which could be calculated fairly easily; the increase in available computing power meant that in recent years the focus has been more on numerical methods that produce more data and a deeper insight into the behaviour of composite materials.




Figure 3.1: Boundary conditions that were used by Arambakam et al. [1] in order to simulate the heat transfer through their fibrous material

The work of R. Arambakam et al in designing a simple method of simulating fibres [1] was used extensively here. They developed an algorithm that allows the generation of fibres in random orientations, with the option of controlling the angle of the fibres and the volume fraction of fibres added to the geometry. This allowed a three dimensional geometry such as that shown in Figure 3.1 to be generated in a computationally inexpensive manner. They then carried out simulations of the heat transfer through the material, with it being kept between two flat plates and symmetry conditions being used on the boundaries of the material. This was done using ANSYS Fluent, a computational fluid dynamics program, with a focus on the way the heat flows through the air surrounding the fibres and also the conduction through the fibres.

Arambakam et al never looked into the mechanical properties of fibrous mats produced in this manner, which is likely to be a function of their fibre generation method producing the fibres in isolation and lacking a mechanism to prevent rigid body movement occurring when forces are applied.

A huge range of parameters have to be considered to accurately model fibre-based porous composites, including the direction, porosity, and conductivity of the fibres as well as the influence of air, water and ventilation on wet fabrics, as well as the anisotropy of the fibres. Woo et al [16] derived an analytic formulation that attempts to take all these properties into account:

$$k_c = k_a \sin(\varphi P_i)^2 + k_{f2} \cos\left(\frac{\varphi \alpha v_f}{1 + \alpha}\right)^2 + \frac{\sin(\varphi(1 - P_i)^2)^2}{\frac{v_f}{k_{f1}} + \frac{(1 - v_f)(1 + \alpha)}{k_a}} \quad (19)$$

$$P_i = \left(1 - \frac{8v_f}{\pi} + \left(\frac{8}{\pi}\right)^2 \frac{v_f^2 \alpha}{(1 + \alpha)^2}\right)^{\frac{L}{2d}} \quad (20)$$

where  $k_c$  is the thermal conductivity through the fabric,  $k_a$  is the thermal conductivity of air,  $k_{f1}$  and  $k_{f2}$  are the thermal conductivities of the fibres along the across their axes,  $v_f$  is the fibre volume fraction,  $\alpha$  is the anisotropy factor,  $L$  is the fabric thickness, and  $d$  is the fibre diameter, and  $\varphi$  is associated with the polar orientation parameter.

While this provides a comprehensive approach that includes the factors of a range of properties, it also includes a number of terms which are difficult to measure, including the thermal conductivity across the axis of a fibre.

They did some comparisons of the experimental and predicted values, and obtained a good level of agreement, as long as they assumed the polar orientation parameter and anisotropy factor values were 0.06 and 1.5, respectively. These values were never justified, and come with the assumption that as a result, they were chosen due



to them providing accurate results over being representative of the experimental samples.

The work of Fricke et al. [17] covers calculating the radiative component of the thermal conductivity of thin, low density test specimens typically manufactured using fibre boards. The fibre boards have densities of the order of  $10 \text{ kg/m}^3$ , and are of a thickness/layup/design so that the heat transfer through them is largely done by two components: the infrared radiation and the conduction through the air gaps (with the solid conduction through the fibres being basically negligible: this is a good approximation for fibre boards but not fibre mats in general. The photon diffusion approximation breaks down, which means that it appears that direct radiation transfer between the boundaries of the material can happen (it appears that the “photon diffusion approximation” implies that there is no net internal radiation transfer as there is too much material in the way, leading to the heat just flowing around randomly rather than in a general “hot” to “cold” way, and so when the assumption breaks down, it is because there is a significant amount of internal heat transfer via radiation).

The authors specifically investigated the behaviour of organic fibres, with the spectral transmission and reflection measurements occurring at wavelengths between 2 and  $40 \mu\text{m}$  in size.

The thermal conductivity for these low density materials is not considered a material property because of its dependence on the geometry (specifically, the thickness of the test specimen affects the thermal conductivity so that the thermal conductivity is not just dependant on the materials the test specimen is made of; and the emissivity of the test machine is also important). As such, it is common to refer to the experimental thermal conductivity as the “apparent conductivity” or “pseudo-conductivity”.

The aim of the paper was to come up with an analytical expression to determine the conductivity of the fibre board. The radiation parameters were obtained using “calorimetric” experiments, with a mean test temperature between 285 and 305K (10-30°C), with thicknesses from 33 to 100mm and densities between 8 and 25kg/m<sup>3</sup>. The emissivity of the surfaces was 0.92, although they also ran the tests using foil covered specimens with an emissivity of 0.06. The “special design and calibration of the heat flow meter apparatus” ensured the experimental error was small enough that the thermal resistance was within 0.5% of the certified value for the reference material they used to ensure the apparatus was correctly calibrated. As a result, the expected accuracy is to within 2% of the actual value of thermal conductivity here.

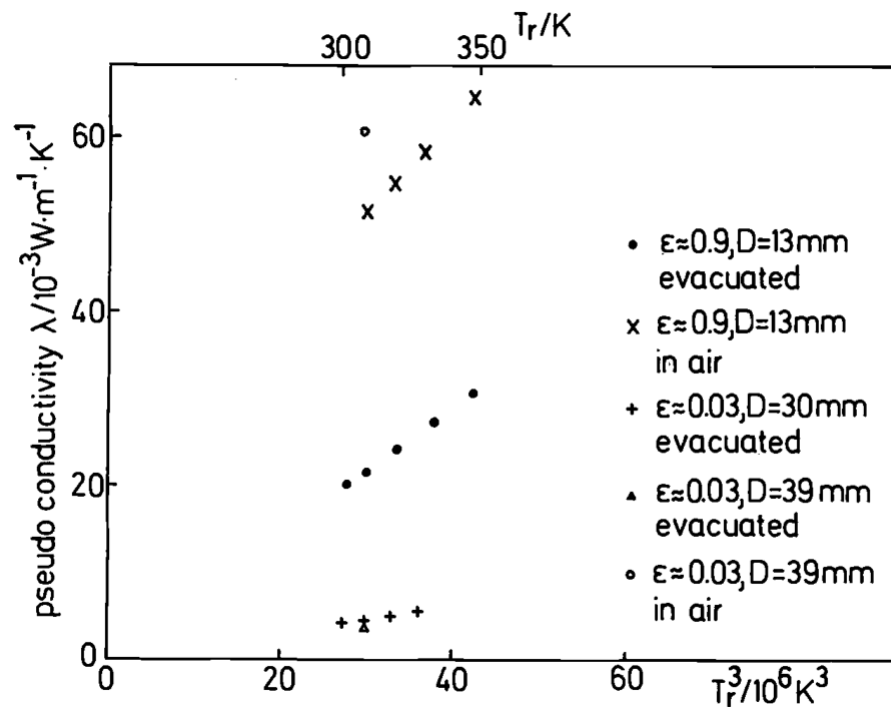


Figure 3.2: Measured thermal (pseudo-) conductivity,  $\lambda$ , of organic fibre boards as function of cubed mean temperature,  $T_R$ ; from [17].

Figure 3.2 shows a comparison of the pseudo-conductivity for a range of test specimens, and it shows that there is a massive increase in the pseudo-conductivity when the test specimen is used in an air atmosphere as opposed to an evacuated atmosphere. The results suggest that fibre diameter is not important when the specimen is in an evacuated state but is important when dealing with test specimens in an air atmosphere. They also appear to have significant temperature dependence, with a variation from around 0.05 to 0.065 occurring over a temperature range of ~50 K (using standard thermal conductivity units). The large increase in the conductivity which occurs from adding in air, which is larger than the conductivity of air, is likely to be due to the “coupling between radiation and air/solid conduction heat transfer” through an interaction between phonon and photon.

The total extinction coefficient was determined based on the relative scattering and absorption cross-sections for the fibres, as well as the porosity of the fibre board and the fibre diameter. The specific extinction can also be determined by taking the extinction coefficient and dividing it by the density. The cross sections depend on the complex index of refraction (which in turn depends on the fibre diameter and the wavelength of light) as well as the angle of the fibres relative to the incoming radiation. Apparently, fibres typically scatter well in the “forward” direction, and so it is common to consider the effect of the fibre angle by scaling the extinction cross section by an anisotropy factor. Sometimes, the intensity function (which appears to describe how what the light hitting the specimen is doing in terms of intensity and is a function of the direction) is normalised with respect to the albedo.

Since the fibre direction in the test specimens is likely to be randomly orientated, all these aspects have to be integrated across every possible direction, which makes this process complex to carry out. The effective specific extinction coefficient varies a significant amount as the wavelength changes (roughly between 7 and 24 m<sup>2</sup>/kg).

Fricke et al did calculate the Rosseland mean spectral extinction coefficient, but they noted that there was some uncertainty as to how accurate this was since the fibre diameter seemed to have some variation. Specifically, they found that the diameter tended to increase near fibre crossings, and so they measured it to be either between 30 and 50 micrometres or between 35 and 45 micrometres.

They also studied the effect of measuring the spectral extinction coefficient using optical methods. It mentions that this method is most important when the refractive index of the material being tested is unknown. It depends on knowing the effective optical thickness of the material, which can theoretically be determined from the optical thickness and the anisotropy factor. The optical thickness is defined as the natural log of the ratio of incident to transmitted radiation on the material.

In 2008, Flourey, Carson and Pham [18] carried out a Finite Element Method simulation investigating both square particles and square fibres to determine the thermal conductivity. They investigated both fully random fibre orientations and loosely aligned in the z-direction fibre orientations (where the fibres are only randomly orientated in two dimensions and not all three). They demonstrated that FEM simulations could produce results with a good agreement to theoretical models, but their focused on using it for simulating food items meat that no experimental validation exists. While this method is useful for modelling composites in general, adapting it to use differently shaped fibres would be necessary since most conventional fibre composites use cylindrical fibres.

Wang et al [19] developed a novel theoretical method based on existing models to predict the conductivity of composites under all conditions. It allows a mix of composite structures to be analysed using a single equation. However, the cost of this was that an extensive derivation process is required to figure out how to relate the structural conductivities to each other. It also necessitates assuming either a pure

parallel or pure series component for fibre-based composites: this can be somewhat compensated for by assuming there is a weighting of the two and the true case is somewhere in between. The largest gap presented here is that the new model is only ever compared to the theoretical models it was partially derived from and not to experimental data. However, since the theoretical models are well established and their accuracy is generally known, this should be sufficient to assume the model is relatively accurate.

While all the methods presented so far allow the thermal conductivity to be predicted if material properties are known, they do not allow specific investigation of how heat flows through a blanket and can obscure insights related to particular fibre blanket features that can significantly impact the thermal conductivity.

The work of M. Wang et al [20] allowed for the creation of simple, 2D lattices representing a plane through a composite material. This allows for the creation of relatively computationally simple thermal simulations to predict how heat flows through composites, both featuring highly orientated fibres and very random fibre mats. However, the 2D nature means that the effect of the fibre orientation in the plane of the heat transfer direction cannot be studied. They found that the system would produce symmetrical conductivities as long as the orientation was completely random: namely, biasing it so that fibres are more likely to lie along the x-axis than the y-axis (for example) would in turn cause a higher conductivity in the x-direction than the y-direction.

While their work should produce accurate results, they did not carry out a validation process here: neither experimental nor theoretical methods were used to compare with their results. This is understandable since it is hard to create a purely 2D experimental setup; but a lot of simple thermal models (including the series/parallel models and models derived from them) allow the calculation of thermal conductivity

in 2D planes relatively easily as long as e.g., the porosity (a key factor in this simulation) is known.

P. Kanouté et al [21] carried out a comprehensive review of multiscale models for predicting how 3D composite structures behave mechanically at both the micro and macroscales. While they exclusively focus on mechanical behaviour of composites in the review, they highlight a number of approaches to simulations that are informative here. Specifically, they validate the approach chosen on some levels, since modelling the microstructure is determined to be a standard method for determining the properties of the material. While a true multiscale modelling approach was not carried out here; the model created could be used to feed data directly into a macroscale model.

Kuksenko et al [22] investigated simulating the mechanical behaviour of wavy fibres. A hyperelastic model was used to describe the material behaviour present in the fibres, where test examples were not intended to represent real fibres. Their findings show that the fibre radius has no significant effect on the mechanical response of the composite, with a possible exception for when the properties of the fibre are very different than the matrix. However, their findings are mostly focussed on straightening the fibres and not the effect the fibres have on the material properties of the composite.

Towards the end, they consider the effect of including multiple fibres, in two different regular arrangements and a random periodic arrangement. They find that, at least in considerations of how the fibres are stretching, the arrangement of fibres has little effect in general, except if the fibres are relatively similar to the matrix mechanically. However, they do not run any simulations of the effect the arrangement of fibres has in Load Case 2 or 3 (see Figure 3.3) which are the two load cases where the response of the fibre is likely to be more affected by the fibre arrangement present.




Figure 3.3: The three load cases used by Kusenko, Böhm and Drach[22], originally Figure 2 in their paper. (a) is referred to as Load Case 1, which deals with elongation of the fibre(s); (b) is Load Case 2, which deals with direct force across the cross section of the fibre(s); and (c) is Load Case 3, which covers shear across the cross section of the fibre(s).

The only validation of their results comes from a comparison to a micromechanical model created by Karami et al [23], which uses a different hyperelastic model for simulating the fibres than Kusenko et al did. Kusenko et al switched to using the same hyperelastic model to ensure their results could be more easily compared to the work of Karami et al, but it raises questions of how much of an impact the hyperelastic model they used for all of their results is having here. It is worth noting that Karami et al validated their model off experimental data and found a good fit to the data. The work of Karami et al only considers one fibre arrangement: the hexagonal distribution which Kusenko considered as one of their two periodic arrangements.

### 3.2.2 Simulation of Randomly Orientated Sinusoidal Fibres

The work of Fisher et al [24], in characterising fibre waviness in nano-tube reinforced composites shows very clear results indicating that the waviness of the fibres was

causing a significant degradation in the mechanical properties of the composite, particularly when the fibres are much stronger than the matrix material is.

These works on the material properties of wavy fibres have significant implications to the current work and suggest that using sinusoidal fibres for greater thermal conductivity performance will in turn have a significant impact on the mechanical performance of the composites. None of these cases are definitive here: hyperelastic models and nanofibers do not, necessarily, correspond to the fibres under study; but the results are very suggestive in this case.

The work of Yu et al [25] investigated the percolation threshold in both straight and wavy fibres. In this research, percolation would be a significant detriment: it would imply the existence of a path through the composite that would allow efficient heat transfer through it, and in turn would cause the thermal conductivity to increase significantly. Their findings indicate that the percolation threshold is higher the wavier the fibres are, which in turn implies that wavy fibres have some value in allowing a higher volume fraction to be reached before percolation occurs and there is a path in the heat transfer direction entirely through the fibres.

### **3.2.3 Mathematical Models for Heat Transfer in Porous Insulation**

There are five practical models which may be used to estimate the thermal conductivity of a nonwoven or porous insulation (although there are other less applicable models):

- Parallel
- Series
- Clayton
- Pilling
- Disorder



Parallel and Series models provide physical upper and lower bounds on the thermal conductivity, directly from basic thermodynamic theory. However, these models are least accurate to describe effective thermal conductivity of a geometry that does not have different constitutive elements arranged in either parallel or series. Clayton and Pilling are industrial standards, whilst the Disorder Model allows for any possible intermediate value to be taken between the theoretical upper and lower limits.

Nonwoven models are generally expressed as a function of volume fraction (Equation (21)). For the case of multiple fibre types, a second function, composition, is defined (Equation (22)).

$$v_f = \frac{V_{\text{Fibre}}}{V_{\text{Total}}} = 1 - \frac{V_{\text{Air}}}{V_{\text{Total}}} \quad (21)$$

$$\theta = \frac{V_{\text{Glass-Fibre}}}{V_{\text{Glass-Fibre}} + V_{\text{PET-Fibre}}} \quad (22)$$

where  $v_f$  is the fibre volume fraction,  $V_{\text{Fibre}}$  and  $V_{\text{Air}}$  are the fibre volume and air volume in the fabric,  $V_{\text{Total}} = V_{\text{Fibre}} + V_{\text{Air}}$ , and  $\theta$  = glass fibre volume fraction in a hybrid fabric composed of both glass and PET fibres.

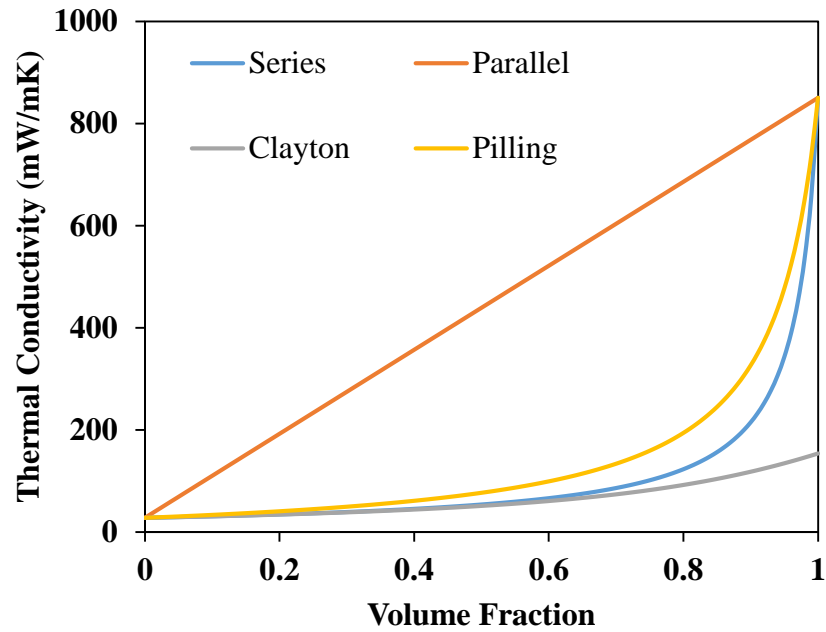


Figure 3.4: Comparative graph showing the individual thermal conductivities as a function of volume fraction, for the case of an air-glass fibre blend.

### 3.2.3.1 Model 1: Parallel Model




Figure 3.5: Representative analogies for conductive heat transfer in nonwoven fabrics showing parallel (Model 1) and series (Model 2) style heat transfer for the case of a 2D air glass fibre blend.

Figure 3.5 illustrates the difference between the parallel and series models of heat transfer: in this case, parallel means the fibres are parallel to the heat transfer direction, analogously to parallel resistors, and the heat flow is split between the fibres and the matrix. In the series case, the heat transfer has to alternate between the fibres and the matrix. Both of these cases are simplifications of real fibre geometries,

The Parallel model, shown in Equation (23), simplifies the geometry of a nonwoven significantly and acts as an upper limit on the heat transfer present, as the fibre alignment means that the majority of the heat transfer can occur through the most thermally conductive component of the composite.

$$k = k_f v_f + k_a (1 - v_f) \quad (23)$$

### 3.2.3.2 Model 2: Series Model

The Series model, shown in Equation (24), behaves as an effective lower bound for the heat transfer through the composite material, as it requires the heat to pass through all of the least thermally conductive component of the composite. It is, in general, more accurate than the parallel model [26], though it still acts as an oversimplification of the material in most cases [27].

$$k = \frac{k_f k_a}{k_a v_f + k_f (1 - v_f)} \quad (24)$$

It is possible to combine the Series and Parallel model using a ratio to indicate how much the composite is like both, using for example Equation (25). Typically, these coefficients would be derived experimentally, though it is possible to derive them theoretically.

$$k_{\text{Hybrid}} = A[k_f v_f + k_a (1 - v_f)] + B \left[ \frac{k_f k_a}{k_a v_f + k_f (1 - v_f)} \right] \quad (25)$$

In Equations (23-25),  $k$  represents thermal conductivity and subscript  $f$  and  $a$  designate fibre and air respectively.  $v_f$  is fibre volume fraction and  $A$  and  $B$  are volume percentages of parallel and series arrangements respectively in a hybrid model. These terms also apply to the following equations including Equations (26-29).

### 3.2.3.3 Model 3: Clayton Model

The Clayton model, shown in Equation (26) [28] was derived based on the formation of a unit cell of a carbon composite material. It provides a good approximation of the thermal conductivity of non-woven models, and performs similarly conservatively to the Series model, particularly at low thermal conductivities. At higher volume fractions, it becomes more similar to the parallel model, representing that the higher level of fibre content means the heat transfer can proportionally be more through the fibres [28].

$$k = \frac{k_a}{4} \left[ \sqrt{(1 - v_f)^2 \left( \frac{k_f}{k_a} - 1 \right)^2 + 4 \left( \frac{k_f}{k_a} \right)} - (1 - v_f) \left( \frac{k_f}{k_a} - 1 \right) \right]^2 \quad (26)$$

### 3.2.3.4 Model 4: Pilling Model

The Pilling Model, shown in Equation (27) [29], was derived using an analogy between elastic behaviour and thermal behaviour, and is generally accurate at low volume fractions. It is, however, a conservative model that tends to underestimate the conductivity. It was designed for carbon fibre reinforced composites at temperatures between 80 and 270K and is reasonably accurate in that condition.

$$k = k_f \left[ \frac{(1 - v_f) + (1 + v_f) \left[ \frac{k_f}{k_a} \right]}{(1 + v_f) + (1 - v_f) \left[ \frac{k_f}{k_a} \right]} \right] \quad (27)$$

### 3.2.3.5 Model 5: Disorder Model




Figure 3.6: The range of possible values of the general Disorder Model, with the upper and lower values matching exactly the values of models 1 and 2 respectively, as  $Z$  varies.

The Disorder model, shown in Equation (28), uses a statistical coefficient,  $Z$ , to take into account the effect the fibre orientation has on the thermal conductivity. This is quantified in Figure 3.6, using the standard component conductivities in Table 3.2). The meaning of values for the coefficient are indicated in where  $\beta$  is the fibre orientation with respect to the direction of heat transfer.

Table 3.1.

$$k = k_a + \frac{(k_f - k_a)}{1 + \frac{v_f}{1 - v_f} \left( 1 + Z \left( \frac{\frac{k_f}{k_a} - 1}{\frac{k_f}{k_a} + 1} \right) \right)} \quad (28)$$

The Z coefficient is found by assessing the relative contributions of Models 1 and 2 [30, 31]. Models 1 and 2 are for  $k_{\beta=90}$  and  $k_{\beta=0}$  respectively.

$$Z \approx \left( 1 + \left( \frac{k_{\beta=90}}{k_{\beta=0}} - 1 \right) \sin^2 \beta \right) \quad (29)$$

where  $\beta$  is the fibre orientation with respect to the direction of heat transfer.

Table 3.1: Z-Value comparison table indicating what various Z values mean

Z Value	Equivalent Angle	Meaning
0	0°	Fibres are orientated parallel to heat flux (as in the Parallel model)
2/3	30°	Fibres are randomly distributed uniformly in 3D
5/6	45°	Fibres are distributed equally between parallel and perpendicular to the heat flux direction
1	90°	Fibres are orientated perpendicularly to the heat flux (as in the Series Model)

### 3.2.3.6 Limitations of temperature independent modelling

Conduction theory generally assumes that the thermal conductivity of a material is independent of its temperature. This assumption is often accurate for ceramic and metallic solids, particularly when the temperature difference is small, but it is not generally valid for gasses [32].

This is because the lack of strong inter molecular bonding within gasses means that increasing temperatures allows a significantly larger amount of atomic movement, and in turn causes the heat transfer through the gas to increase. This can be quantified using the kinetic theory of gases combined with Fourier's constitutive law for one-dimensional conduction.

As a result, the thermal conductivity of a gas is a function of two independent variables, as shown in Equation (30): the mean particle speed,  $\overline{V_{\text{Molecule}}}$  (defined in Equation (31)), and the mean free path of the gas,  $\gamma_{\text{MFP}}$ .

$$k_g = \frac{N\overline{V_{\text{Molecule}}}\gamma_{\text{MFP}}C_v}{3Na} \quad (30)$$

$$\overline{V_{\text{Molecule}}} = \sqrt{\frac{3TR}{m_m}} \quad (31)$$

where T is the temperature, R is the gas constant,  $m_m$  is the mass of molecule, and N is the Avogadro number.

The manner by which the mean free path affects the thermal conductivity will be discussed further in the next section. However, both the mean free path and the mean particle speed are functions of the temperature.

For this reason, it is generally incorrect to ignore the temperature dependent effects present within the simulations, particularly if there is a significant temperature difference across the simulation. In practice, working with exact relationships like above is impractical, and tabulated data can be used instead.

In this work, the tabulated data was curve fitted to produce an average thermal conductivity for the air. Several correlations [33, 34] exist which allow for a continuous value to be determined (with an accuracy of  $\pm 0.15\%$ ), one of which is shown in Equation (32). A true, correctly weighted average may be found using integration, as shown in Equation (33).

$$k(T) = \frac{2.3340 \cdot 10^{-3} \sqrt{T^2}}{164.54 + T} \quad (32)$$

$$k_{\text{average}} = \left[ \frac{1}{T_{\text{Upper}} - T_{\text{Lower}}} \right] \int_{T_{\text{Lower}}}^{T_{\text{Upper}}} \left[ \frac{2.3340 \cdot 10^{-3} \sqrt{T^3}}{164.54 + T} \right] dT \quad (33)$$

A short MATLAB script has been written to perform this integration to provide the true average thermal conductivity used in simulations, which produced an air conductivity value of 27.1 mW/mK.

### 3.2.3.7 Mean Free Path

As shown in Equation (30), the thermal conductivity of a gas is a function of the mean free path. In macroscopic models, this effect may be neglected; however, the scale of individual fibres is extremely small and so it may be of interest here. In general, the limit of continuum mechanics (namely, the ability to treat gas as a bulk material) for air is around 0.1  $\mu\text{m}$ . Since the diameter of a fibre is around 8-14  $\mu\text{m}$  in the simulations, this effect should be negligible but was briefly investigated to ensure it was [35].

The application of kinetic theory, mean free path mechanics and molecular dynamics means that it can be shown the thermal conductivity is a function of the mean collision distance, as shown in Equations (34) and (35) [36];

$$k(\text{mfp}) = \frac{k_{\infty}}{1 + \frac{4C_p}{C_p + C_v} \frac{\text{Kn}}{\text{Pr}}} \quad (34)$$

$$\text{Kn} = \frac{\gamma_{\text{mfp}}}{l_{\text{collision}}} \quad (35)$$

Estimating the collision in these conditions is difficult, so an average empirical relationship for fibrous materials was used, partially based on the Disorder model (Section 3.2.3.5), which is shown in Equation (36). Once this is known, the mean free path can be calculated using Equation (37).

$$l_{\text{collision}} = Z \frac{d_{\text{fib}}}{v_f} \quad (36)$$



$$\gamma_{\text{mfp}} = \frac{k_b T}{\sqrt{2} \pi d_m^2 P(T)} \quad (37)$$

This allows Equation (34) to be used to determine the thermal conductivity of bulk air at microscopic scales, with correlations being used for the specific heat capacity, specific enthalpy, pressure and Prandtl number.

This simplification assumes the thermal accommodation is negligible, which is true as long as the gas is not at low pressures [37, 38]. Since the blanket can be assumed to be at standard atmospheric pressure, this assumption is accurate.

Figure 3.7 shows the comparison of the thermal conductivity of air when determined using Equation (32) and when determined using Equation (34). It can be seen that the mean free path effects on this scale are, as predicted, negligible. This means it is appropriate to treat the air as a continuum in this work.




Figure 3.7: Comparison of the thermal conductivity of air when it's treated as a continuous and non-continuous fluid at the length scale of the simulations

### 3.2.4 Heat Transfer by Convection

Thermal insulation in general attempts to minimise the convective heat transfer within it. Traditional convection (both natural and forced) relies on there being a bulk movement of air, and inhomogeneous fibre structures are designed to limit this as much as possible to reduce the heat transfer condition to just conduction and radiation.

Convection was briefly investigated to ensure that it could be neglected. Two methodologies were utilised here: one using order of magnitude analysis and one using an assessment of whether a boundary layer could form upon the fibres [39].

#### 3.2.4.1 Order of Magnitude Analysis

In order to demonstrate that convection was negligible, a comparison with the conduction was carried out to show that convective heat transfer was less than 10% of the magnitude of the conductive heat transfer. In this case, the Nusselt number is the property of interest, as it compares the relative size of both mechanisms of heat transfer.

Evaluating a theoretical Nusselt number directly is challenging without good experimental data, which means that heat transfer correlations were used instead. This is made more challenging as they are typically for relatively high Rayleigh numbers; whereas the lack of directed flow within a composite blanket usually means the Rayleigh number would be low: particularly when coupled with the extremely small scales being investigated.

The selected correlation, shown in Equation (38) [33], only requires knowledge of the Grashof and Prandtl numbers. The Grashof number was calculated using Equation (74) [40], which represents the worst-case geometry arrangement: a vertically orientated fibre which is the full length of the bounding region. In this case, the wall

temperature was 80°C and the bulk temperature was 4°C, with the Prandtl number being that of air (0.711).

$$\text{Nu} = \frac{4}{3} \left[ \frac{\text{Gr}}{4} \right]^{\frac{1}{4}} \left[ \frac{0.75\sqrt{\text{Pr}}}{(0.609 + 1.221\text{Pr}^{0.5} + 1.238\text{Pr})^{\frac{1}{4}}} \right] \quad (38)$$

$$\text{Gr} = \frac{g\beta_R(T_w - T_\infty)l^3}{\nu^2} \quad (39)$$

This produces a Nusselt number of 0.22, which implies that in the worst case the convection within the simulation is about 1/5<sup>th</sup> of the conduction. Since this worst case is generally to be avoided, it can be assumed the convection is, in most cases, significantly smaller and hence can be treated as being negligible.

#### 3.2.4.2 Boundary Layer Formation

While the previous method indicates that the convection is negligible, it works under the assumption that convection can occur within the material. However, this is not necessarily the case. If the convective boundary layer cannot form, due to the lack of space within the composite, then convection cannot occur within the material [41].




Figure 3.8: Two infinite, flat plates at different temperature (black), with a fluid between them (blue)

If the case of two infinite plates separated by an interstitial fluid, as shown in

Figure 3.8, is considered, where  $T_1$  and  $T_2$  are at different temperatures, then convection between the two plates can only occur if a convective boundary layer can be formed. As such, as long as it can be demonstrated that the boundary layer thickness,  $L_B$ , is comparable in size to the distance between fibres,  $L_C$ , it can be determined that convection cannot occur.

$$N_1 = \frac{L_B}{L_C} > 5 \quad (40)$$

It can be shown that for small lengths, the velocity boundary layer in a natural convection problem is linear and proportional with its vertical position along the body<sup>1</sup>:

$$L_B = L_{Fib} \left[ \frac{3.93}{\sqrt{Pr}} \right] \left[ \frac{0.952 + Pr}{Gr} \right]^{\frac{1}{4}} \quad (41)$$

$$L_C = Z \frac{d_{fib}}{\alpha} \quad (42)$$

$$N_1 = \frac{L_B}{L_C} = 21 \quad (43)$$

Based on these two methodologies convection is considered completely negligible in the analysis.

### 3.2.5 Heat Transfer by Radiation

#### 3.2.5.1 Order of Magnitude Analysis

Radiative heat transfer is dependent on the view factor, and also has a much more significant dependence on the temperature than conduction and convection have, since the temperature difference is to the fourth power. In order to demonstrate that

---

<sup>1</sup> Note that whilst in mean free path assessments, 'worst case' was a lowest value, in this case the 'worst case' is the highest value and therefore, it is the highest value that has been used.

radiation can be neglected, an order of magnitude analysis was carried out to determine how big its effect is in comparison with the conductive heat transfer.

For glass fibre, the emissivity of the surface is approximately 0.85 [42], and the view factor was treated as being the same as in the case of a flat surface enclosed within a larger surface, that is, 0.2. This allows Equation (44) to be used to determine the radiative conductivity, and how large it is in comparison to the air conductivity.

$$N_2 = \frac{k_a}{k_r} = \frac{k_a}{\left[ \frac{\sigma \cdot \Delta x \cdot (T_1^2 + T_2^2) \cdot (T_1 + T_2)}{\frac{1}{F_{1-2}} + 2 \left( \frac{1 - \varepsilon}{\varepsilon} \right)} \right]} \approx 85 \quad (44)$$

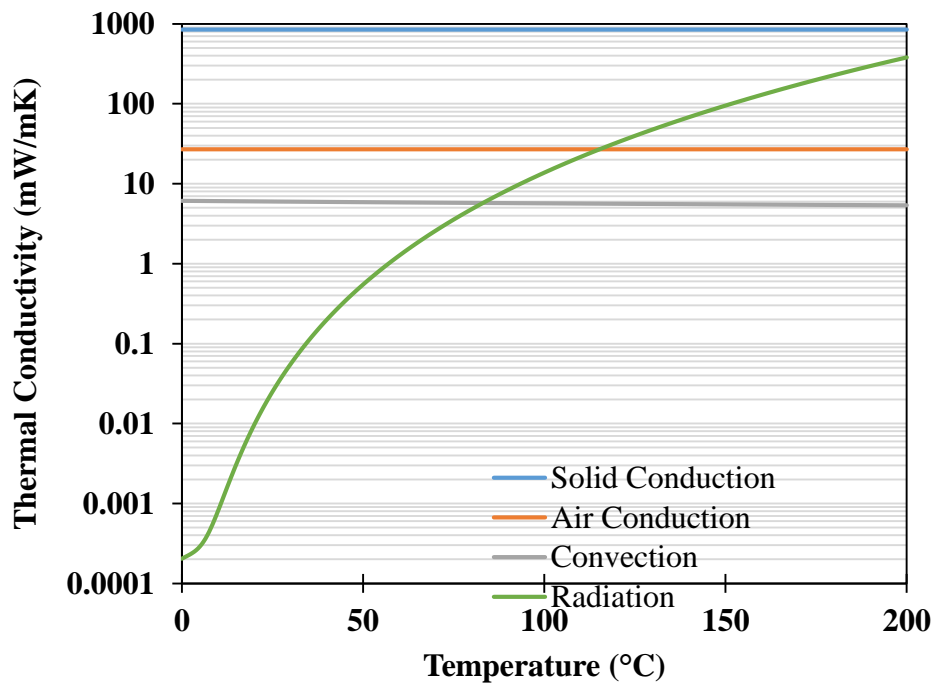


Figure 3.9: Comparison of the magnitude of the components of heat transfer through a blanket, with the convection and radiation being treated as conduction equivalents here.

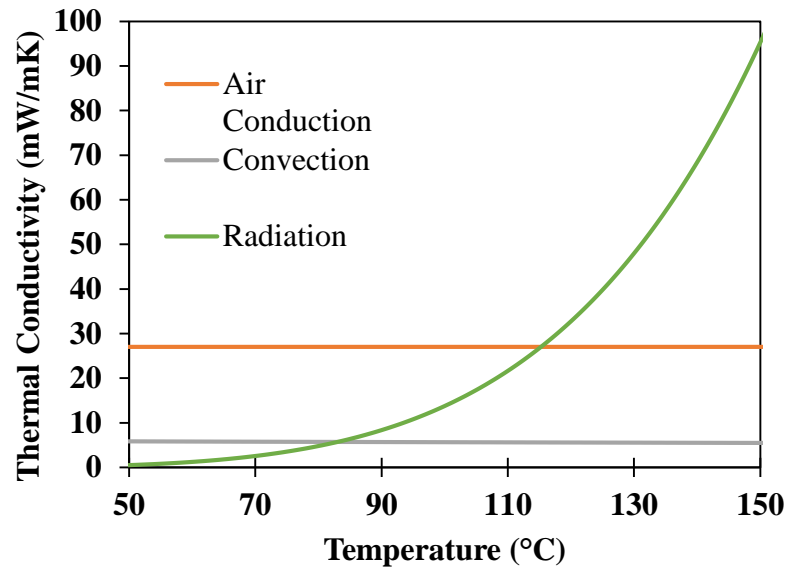


Figure 3.10: The results from Figure 3.9 zoomed in so that the difference at high temperatures is clearer, with the Solid Conduction removed; note that the y-axis is no longer logarithmic

Figure 3.9 shows the relative effective thermal conductivities of convection and radiation, compared with the average thermal conductivity of air and glass fibres. This is a strongly temperature dependent phenomena which does not become significant until approximately 85°C, as can be seen more clearly in Figure 3.10.

### 3.2.5.2 The Diffusion Approximation

It can be assumed that thermal radiation can only travel short distances inside a fibre before being scattered or absorbed, and so that bulk fibres are optically thick (i.e., that the amount of transmitted thermal radiation through either glass fibres or PET fibres is negligible).

This assumption means the diffusion approximation for radiation can be used to allow the radiative heat flux to be represented in terms of Fourier's Law of Conduction [43,

44], as shown in Equation (45). Since the fibres are optically thick, we can assume the refractive index (N) is approximately equal to 1, which means that it can be simplified to Equation (46).

$$q_r = -k_r \frac{dT}{dx} = -\left(\frac{16 \sigma_e N^2 T^3}{3 \sigma_R}\right) \frac{dT}{dx} \quad (45)$$

$$q_r = -\left(\frac{16 \sigma_e}{3 \sigma_R} T^3\right) \frac{dT}{dx} \quad (46)$$

In this equation,  $\sigma_R$  is the Rosseland mean extinction coefficient, sometimes referred to as the Extinction coefficient or the attenuation coefficient. The calculation of the Rosseland coefficient is beyond the scope of this project, however, it is assumed to be comparable to that of Carbtex fibres because of the similar geometric arrangement and surface finish. Also, Carbtex fibres much like insulation-based glass fibres are optically thick due to the low-grade manufacturing process and so this assumption is considered valid for the current applications [30].

$$N_3 = \frac{k_a}{k_r} = \frac{k_a}{\frac{16 \sigma_e}{3 \sigma_R} T^3} \approx 200 \quad (47)$$

### 3.3 Fibre Geometry Generation

The first step in the fibre generation process is to define a fixed region of space to generate the fibres within, and a fixed number of fibres to attempt to insert into the region of space. The next step is to determine what kind of fibre to insert: this is done

by comparing the ratio of straight fibres to total fibres against the desired ratio of straight fibres<sup>2</sup>.

Next, the position of one end of the fibre and two orientations are generated, with the position being bounded to be close to but not necessarily within the fixed region. This choice was made to ensure that the fibre distribution through the region was more even and not clustered on the centre. The two orientations define the “in-plane” and “out-of-plane” orientation, where the out-of-plane orientation controls how vertical the fibres are, and the in-plane angle controls their angle in the horizontal plane. In this case, vertical is the heat transfer direction, and horizontal is perpendicular to that. The fibre diameter and length are then generated, and the fibre orientation is used to determine where the end of the fibre is in 3D space, as described in Appendix A.1.




Figure 3.11: An example of a straight fibre (left) and curled fibre (right) geometry once it has been imported into ANSYS Workbench (x direction: horizontal and y direction: vertical)

---

<sup>2</sup> There's a line in the code that needs to be changed if a purely straight or purely curled fibre simulation is required: "if gfvol/totvol <= gfr || isnan(gfvol/totvol) == 1" inserts a straight fibre first; whereas "if gfvol/totvol < gfr" inserts a curled fibre first



The calculation of the x position of the fibre always produces a positive x value, which in turn would bias the fibre direction and make the mat less random, so the sign of the x-position of the final point is randomly flipped approximately half the time.

All of these parameters are assigned based on a uniform distribution between two fixed values, which can be the same if no variation in parameters is desired, for example, if every fibre in the simulation should have the same diameter. While the effect of it was not investigated, it is also easy to change the uniform distribution into a normal distribution, where the variation is centred on a chosen mean value and the expected variation from the mean is used to control how wide the distribution is.

The initial height of the fibre is very high and significantly outside the bounding box, and the fibre is then lowered into it until it reaches another fibre. This is done to ensure that the fibre starts completely out of the bounding region and can be successfully lowered into it without being already touching or underneath a different fibre. This process is done by iterating through the previously placed fibres and finding out which one is the shortest vertical distance from the new fibre. This is done by treating the existing fibre and the new fibre as being two, two dimensional lines, and then computing where the two lines intersect. Once this is done, linear interpolation can be used to determine what the z-positions in each fibre at the intersection point is. Once this is known, trigonometry is used to calculate how much of the vertical distance between the fibres is within each fibre, meaning the total distance between the fibres can be found as the difference in the z-positions minus the distance within the fibres (since the centre lines of the fibres should not intersect, merely the outside of the fibres). Once the minimum distance between the new fibre and an existing one is known, the fibre is lowered by the vertical distance so that they touch, and then slightly

further to ensure the fibres intersect<sup>3</sup>. If there are no existing fibres underneath the new fibre, then it can be instead lowered until one end of the fibre touches the bottom of the bounding region. Note that in the case of a horizontal fibre, the whole fibre volume will end up touching the bottom of the bounding region in this case.

The final stage of the fibre generation process is to clip the fibre. For the horizontal fibres in the simulation, the Liang-Barsky algorithm is used (since it is simpler and faster to run, but it is also a 2D algorithm and does not cope with the fibres having variation in 3D) [45, 46], while for the fully 3D fibres, the Kodituwakku-Wijeweera method [46] was used instead. This method is relatively fast, but unfortunately relies on a gradient computation for the lines which, in fibres which are flat, produces a division by 0 error and as a result, it only works for fibres where the ends are at different x, y and z positions.

The clipping process is carried out to ensure that the fibres remain within the desired region; since the initial point can be generated outside the region and the final point's orientation is not constrained to be in the direction of the region, it is possible for a fibre to be either partially or fully outside the bounding region. If the fibre is fully outside it, it is marked for deletion, whereas if it is partially outside the region, it is clipped to be within the bounding region. The algorithm behind clipping process is essentially to find the interception between individual fibres with the bounding region and remove the segment of fibre outside the bounding region while keeping the rest of the fibre inside the region. After clipping, the length of the fibre is recomputed, partially for

---

<sup>3</sup> Overlapping the fibres in this manner is a significant difference from the physical reality of a fibre mat, but was done because two fibres touching provides a very small contact area that is extremely prone to simulation errors due to the nature of finite element analysis; and a general underestimation of the thermal conductivity of the sample by adding many bottlenecks into the geometry that would in reality have larger contact areas due to fibre movement.

calculating the volume of the fibre for estimating the volume fraction and partially to ensure the fibre has not become very short. In general, a length less than  $1/10^{\text{th}}$  of the initial fibre length was considered very short<sup>4</sup>, and any fibre meeting this criterion was deleted. This is done because the meshing algorithm used considered such fibres to be more important than usual and tended to create a high density mesh on them, while in reality such fibres are not of much importance to the heat transfer in the mat and are unrepresentative of a mat with relatively long fibres.

Once the mathematical positioning has been conducted, the fibre matrix in MATLAB is processed into a valid APDL script that allows a 3D geometry to be produced. For each straight fibre, the APDL script follows a simple process:

1. Move the working plane to the initial point of the fibre
2. Rotate the working plane so the z-axis points along the fibre
3. Add a cylinder there with a length and diameter equal to the fibre length/diameter

Once this is done, APDL produces an IGES file representing the entire fibre geometry that can be imported into ANSYS Workbench and solved with appropriate boundary conditions applied.

---

<sup>4</sup> The definition of “very short” can be customised very easily to allow for investigation of shorter fibre simulations or simulations featuring very long fibres more easily.

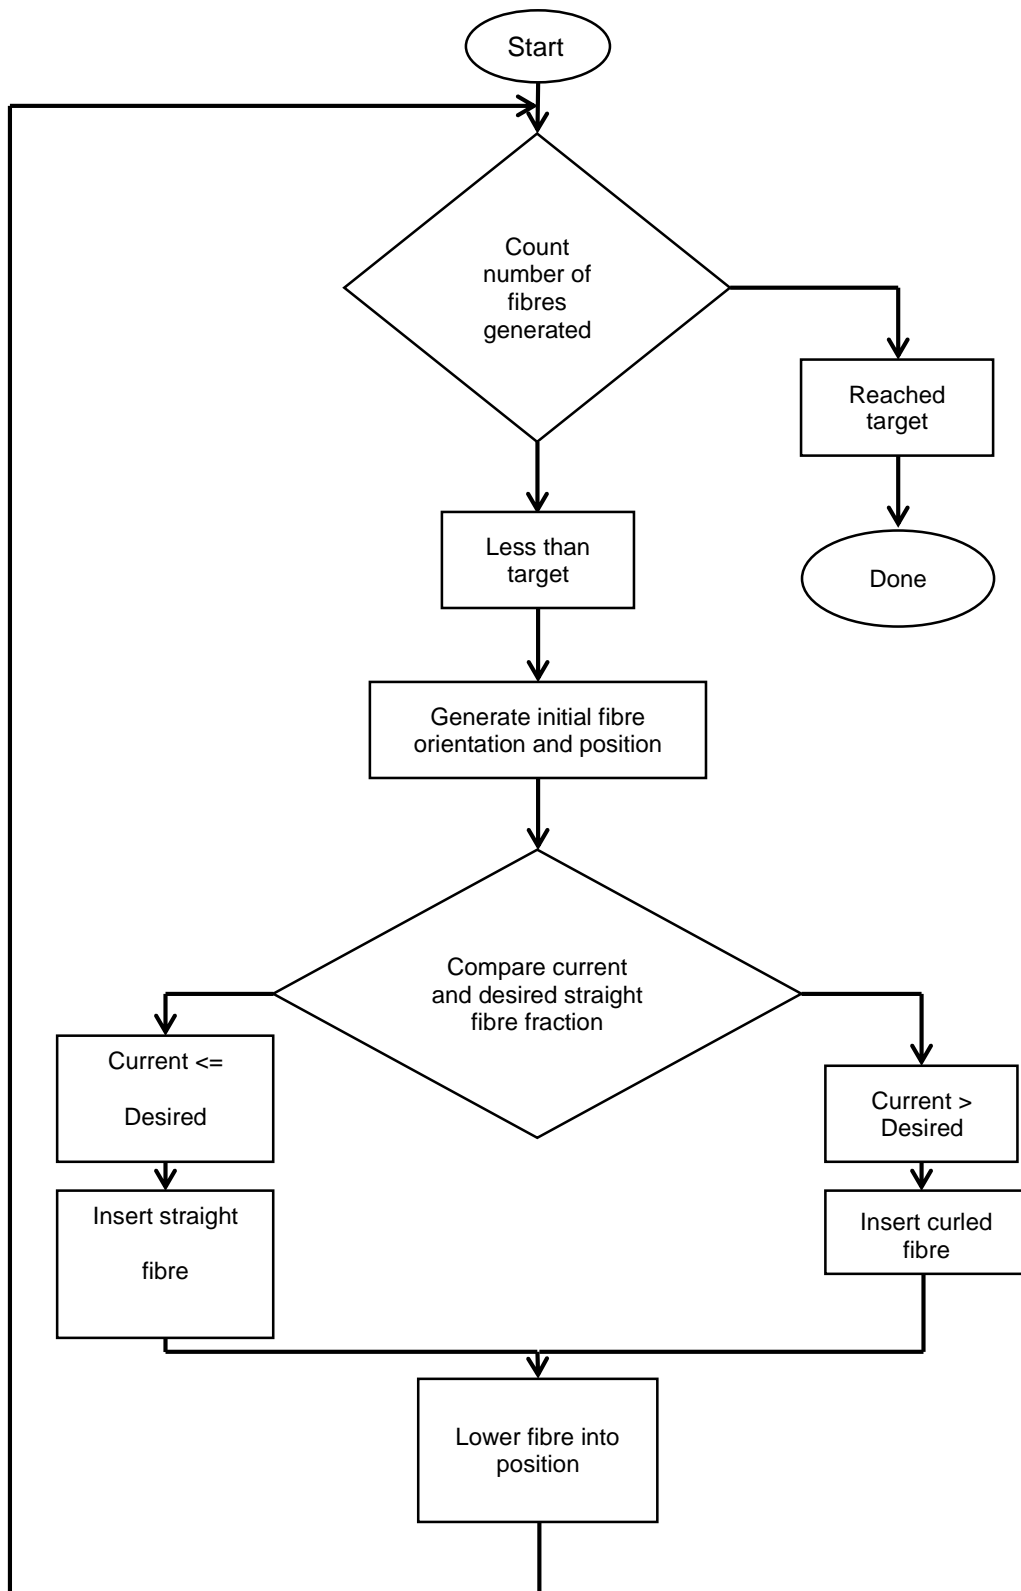


Figure 3.12: Fibre insertion algorithm, illustrating how a three-dimensional fibre geometry is created by the algorithm

For sinusoidal fibres, the process was carried out a very similar, albeit with some key differences. The first difference is that a “fake diameter” was used throughout to represent the sinusoidal fibre as a single cylinder with a radius equal to the amplitude of the fibres. This was done as the “fake diameter” allows the positioning of the sinusoidal fibre to be calculated as if it was a straight fibre, including using it as if it was the actual fibre diameter; since it means that the positioning of the crimps does not need to be known before the fibre can be placed. The main reason for doing this is to simplify the generation process: the clipping algorithms can shorten the fibre, so placing the fibre before crimping it means that partial crimps do not need to be considered. The trade off of this simplification is that it means the sinusoidal fibres effectively take up the entire “fake diameter” space, and so they cannot be packed as efficiently. As this algorithm is mainly looking at thermal conductivity, and particularly reducing it, this was deemed to be an acceptable trade-off.

The other major fibre generation difference with the sinusoidal fibres is that the fibres are more complex to generate a 3D computer model of since they cannot be represented as a simple cylinder. The current system basically calculates where the turning points of the fibre are and then produces a spline between points representing them. Once this is done, a circle is swept along the centre line to produce a 3D sinusoidal fibre with an approximately constant radius.




Figure 3.13: Approximation of a crimped fibre (blue) as three different line segments (purple, dark red)

Since the actual shape of the fibre is not fully determined until it is rendered in APDL, calculating the exact volume fraction in MATLAB is not possible. Instead, the volume of a fibre, and in turn the volume fraction, can be estimated by treating each fibre as a series of straight line segments, as shown in Figure 3.13. Each straight-line segment extends from the centre line of the fibre by the “fake diameter” and is considered to have a diameter equal to the true diameter of the fibres. As a result, the total fibre volume can be approximated as a number of cylinders based on the number of “crimps” present in the fibre.

A visual comparison of how the orientation of the fibres for both fibre types, and mixed fibres, is shown in

Figure 3.14. The left most column represents fibres that are close to unidirectional, the middle column is random in-plane but “flat” out of plane; and the final column is fully random fibres. The first two rows contain the straight and sinusoidal fibres respectively while the final row is a mixture of the two fibre types. These results demonstrate how fabric architecture can be manipulated through the developed the codes and the capability can represent the simplest fibre format such as regular fibre packing in woven fabrics as well as the more complex pattern found in nonwoven fabrics.

The MATLAB codes used to generate the geometry are included in Appendix C, with the subsections containing required functions for the main algorithm to run.

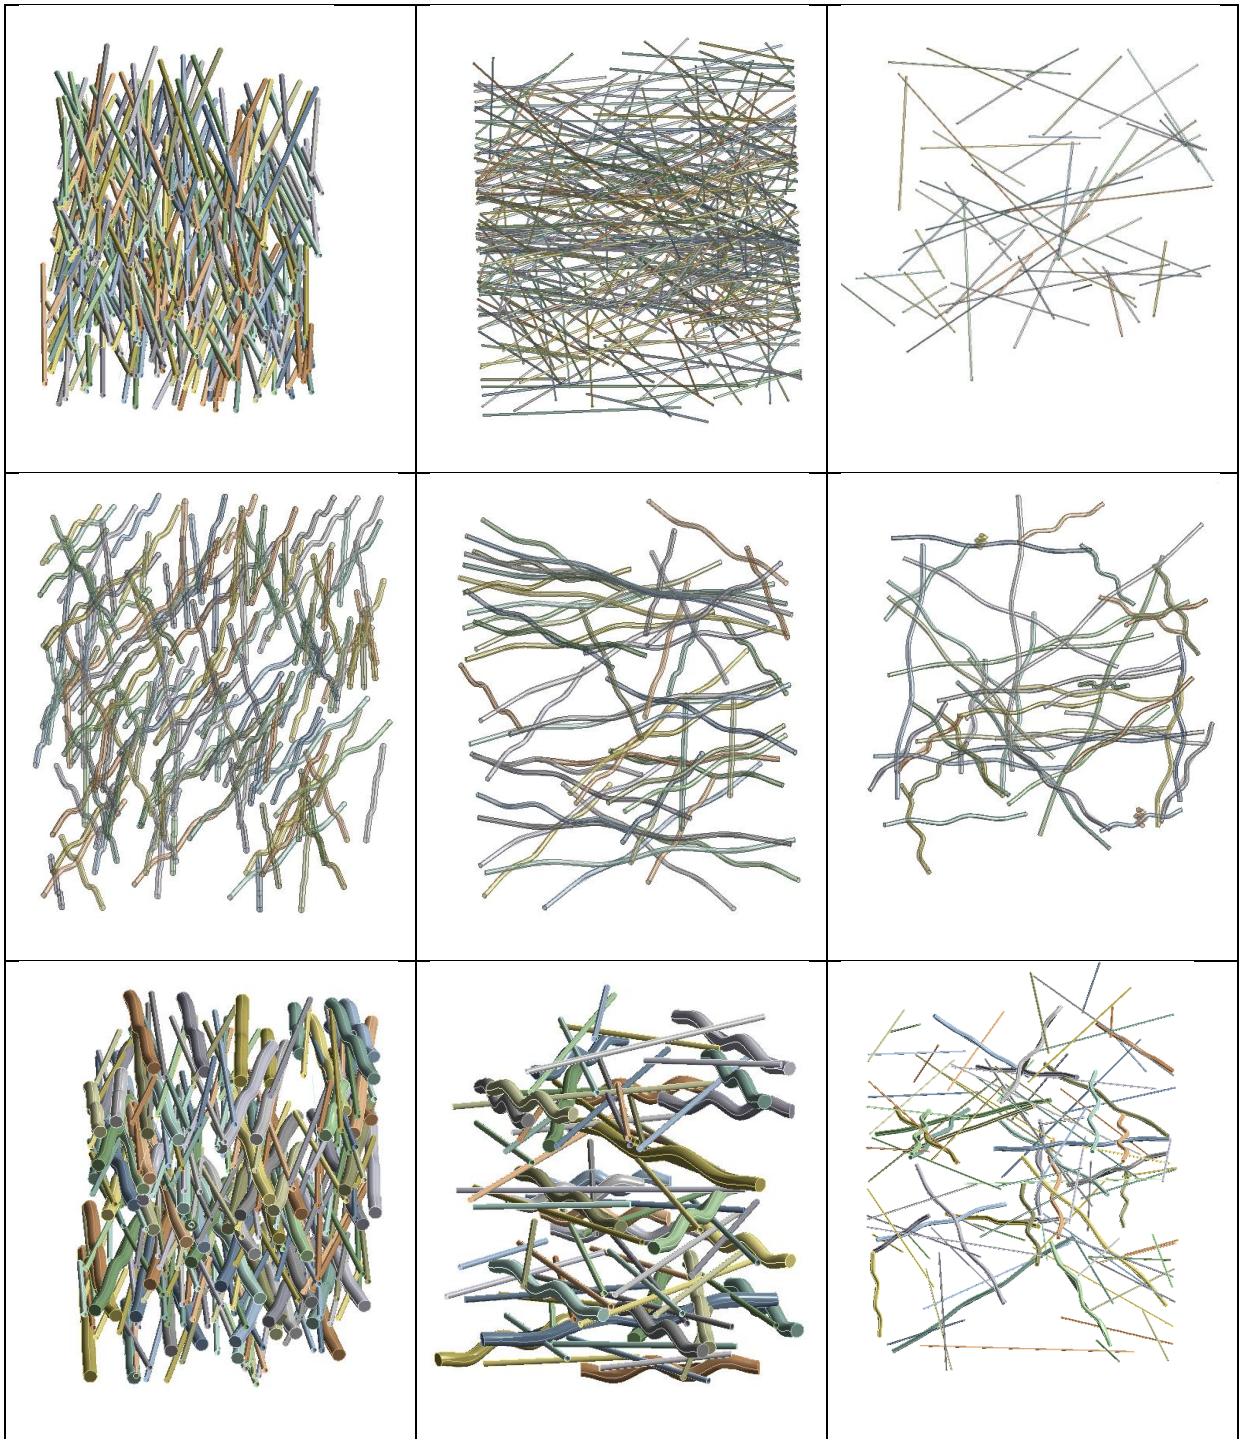


Figure 3.14: Visual comparison of the fibre geometries that are possible to be generated

### 3.4 Simulation Setup

The first step in simulating fibre geometries is to import the individual geometry files into DesignModeler. Generally, a box is created at this stage to embed the fibres in (and to represent a matrix material: throughout this work, this is typically either air, aerogel or a polymer matrix), with a recommendation to ensure the box is significantly larger than the fibre geometry being generated<sup>5</sup>. The fibres can then be inserted, one type at a time, using the 'Slice Material' option (which creates a hole in the matrix which corresponds with the location of the fibres; using 'Add Material' would overlap the fibre and the matrix). Once one fibre type has been inserted, the fibres can be grouped together into one part to allow material properties to be assigned to it as a whole, instead of requiring each fibre to have properties assignments individually. After the fibres have been successfully inserted, the bounding region can be resized to the intended size for it.

It can be worth setting the 'Fluid/Solid' property for the bounding region to 'Fluid', as that causes ANSYS Mechanical to render it semi-transparently instead of fully opaquely, and does not have any other effects on the simulation except if ANSYS Fluent is to be utilised.

Once the geometry has been generated, DesignModeler can be closed. The materials to be used in the simulation can be defined at this stage using 'Engineering Data'; typically the properties shown in Table 3.2 are used throughout, though some Chapters use alternative properties.

---

<sup>5</sup> Very occasionally, the geometry generated is larger than expected, or close to intersecting the sides of the box, and DesignModeler can fail to import them as a result.



Table 3.2: General thermal properties

Material	Isotropic Thermal Conductivity (mW/mK)
Air	24.2
Glass Fibre	850
PET Fibre	150

The fibre/fibre and fibre/air contact settings used throughout this work assume that perfect contact exists, that is, that there is no interfacial thermal resistance. While this is just an approximation of the interface that is not fully accurate, it does appear to not cause a significant amount of inaccuracy in the simulations, particularly since most are occurring at a fairly low volume fraction which means the interface is a relatively very small component of the simulation. It also means that complex experimental measurements are not required to determine what the interfacial properties are.




Figure 3.15: The standard boundary conditions used for thermal simulations

The boundary conditions are shown in Figure 3.15; where the top and bottom are treated as being kept at a fixed temperature of 80°C and 4°C respectively. The other faces of the bounding region are treated as perfect insulators; so that no heat can transfer across them. These temperatures were chosen as they are reflective of the conditions the actual material would be exposed to in subsea pipe-in-pipe insulation.

The thermal validation carried out later uses equivalent boundary conditions, albeit with alternative temperatures applied. This validation uses the same mechanism to generate a geometry and the same contact settings.

### 3.4.1 Fibre Volume Fraction Simulations

Table 3.3: Geometry and mesh properties of the straight fibre volume fraction models

<b>Volume Fraction (%)</b>	<b>Bounding Region Size (mm)</b>	<b># of Fibres</b>	<b># of Elements (solid70)</b>
<b>1.47</b>	6.09	128	3,094,083
<b>2.92</b>	3.04	59	948,423
<b>5.74</b>	1.52	31	331,588
<b>13.23</b>	0.72	13	189,653

Table 3.4: Geometry and mesh properties of the sinusoidal fibre volume fraction models

<b>Volume Fraction (%)</b>	<b>Bounding Region Size (mm)</b>	<b># of Fibres</b>	<b># of Elements</b>
<b>1.35</b>	4.14	16	518,147
<b>2.83</b>	3.19	18	403,909
<b>4.83</b>	2.04	6	339,986
<b>9.59</b>	2.20	33	540,435

The properties used to generate a geometry for the single fibre type volume fraction simulations are shown in Table 3.3 (straight) and Table 3.4 (sinusoidal). The inexact volume fraction here is a consequence of the volume fraction that is generated being uncontrolled: it varies randomly depending on the fibre positioning, which in turn depends on the allowable fibre orientation and the size of the bounding region most, though the fibre length and diameter contribute in addition.

In turn this explains why the bounding region size was changed through these simulations: in order to alter the volume fraction significantly, the bounding region size needs to be changed to control how densely packed the fibres end up becoming. The orientation was kept the same as it was believed to have a larger influence on the heat transfer through the simulation: while the bounding region does affect how many fibres are inside the geometry, it should not affect the results directly as long as there are sufficient fibres within the region to ensure it is representative. However, the orientation controls how common fibre/fibre contact is, which in turn is representative of how easily heat can be transferred through the system.

Table 3.5: Mesh settings used in both the straight and sinusoidal fibre volume fraction modelling

<b>Property</b>	<b>Value</b>
<b>Relevance</b>	40
<b>Size Function</b>	Adaptive
<b>Relevance Centre</b>	Medium
<b>Transition</b>	Slow
<b>Span Angle Centre</b>	Medium

The mesh settings used in this simulation are shown in Table 3.5. These settings were, broadly, used throughout as they produced a good trade off were the number

of elements was large enough to ensure the simulation was accurate while not being so large that the computational time became excessive. In some cases, either due to the geometry making the meshing process harder, or due to these settings producing excessive numbers of elements, these settings were changed.

### 3.4.2 In-Plane Angle Simulations

The next set of simulations that were carried out were the 'In-Plane Angle' ones. In this set, the fibres were changed from being approximately unidirectional (low maximum in-plane angle) to being fully random in 2D (high maximum in-plane angle).

Table 3.6: Geometric properties of the straight fibre simulations in the 'In-Plane Angle' Case

<b>Max In-Plane Angle (°)</b>	<b>Volume Fraction (%)</b>	<b>Bounding Region Size (mm)</b>	<b># of Fibres</b>
<b>360</b>	2.48	0.41	148
<b>180</b>	2.60	0.401	142
<b>90</b>	2.46	0.42	140
<b>45</b>	2.63	0.401	150
<b>45</b>	3.41	0.401	206
<b>36</b>	3.41	0.416	209
<b>36</b>	4.32	0.406	241
<b>30</b>	3.97	0.42	244

Table 3.7: Geometric properties of the sinusoidal fibre simulations in the 'In-Plane Angle' Case

<b>Max In-Plane Angle (°)</b>	<b>Volume Fraction (%)</b>	<b>Bounding Region Size (mm)</b>	<b># of Fibres</b>
<b>360</b>	0.05	4.2	83
<b>180</b>	0.05	4.0001	78
<b>90</b>	0.05	4.15	67
<b>45</b>	0.05	4.0001	86
<b>45</b>	0.06	4.1	118
<b>36</b>	0.05	4.0001	98
<b>36</b>	0.07	4.0001	125
<b>30</b>	0.07	4.0001	111

The geometric properties used in these simulations are shown in Table 3.6 and Table 3.7. There are some duplicate rows here as the packing density (number of fibres per volume) increased, which was caused by the orientation being narrowed, and so the volume fraction became larger. Specifically, this occurs because randomly orientated fibres take up comparatively more space than closely orientated ones.

A sample at 0° was not included purely because the algorithm to generate fibres was not designed for unidirectional fibre geometry building, and the distance between fibres was miscalculated there which led to a large degree of fibre/fibre overlap throughout the entire length of the fibres. This also caused the volume fraction present to become very large, in addition to the number of fibres present, which in turn meant that solving such a case was not practical.

Table 3.8: Mesh properties and settings for the straight fibres

<b># Of Elements</b>	<b>Relevance</b>	<b>Size Function</b>	<b>Relevance Center</b>	<b>Transition</b>	<b>Span Angle Center</b>
210,777	0	Adaptive	Coarse	Fast	Coarse
3,730,796	-80	Adaptive	Coarse	Slow	Fine
1,591,158	0	Adaptive	Medium	Slow	Medium
195,099	0	Adaptive	Coarse	Fast	Coarse
4,917,188	40	Curvature	Coarse	Fast	Medium
1,658,673	-40	Curvature	Coarse	Fast	Medium
302,546	0	Adaptive	Coarse	Fast	Coarse
307,902	0	Adaptive	Coarse	Fast	Coarse

Table 3.9: Mesh properties and settings for the sinusoidal fibres

<b># Of Elements</b>	<b>Relevance</b>	<b>Size Function</b>	<b>Relevance Center</b>	<b>Transition</b>	<b>Span Angle Center</b>
296,893	0	Adaptive	Coarse	Fast	Coarse
480,706	-80	Adaptive	Medium	Slow	Medium
263,429	0	Adaptive	Coarse	Fast	Coarse
296,803	-80	Adaptive	Coarse	Fast	Coarse
420,736	-80	Adaptive	Coarse	Fast	Coarse
283,917	0	Adaptive	Coarse	Fast	Coarse
496,249	0	Adaptive	Coarse	Fast	Coarse
439,411	0	Adaptive	Coarse	Fast	Coarse

The mesh settings used in the In-Plane Angle simulations are shown in Table 3.8 and Table 3.9. The high number of elements produced, particularly for the sinusoidal fibres, means that coarser mesh settings were used than in the Volume Fraction simulations to ensure that a similar number of elements was used. Some of the sinusoidal simulations used an alternative size function as the adaptive one could not create a mesh upon their surfaces.

### 3.4.3 Out-of-Plane Angle Simulations

These simulations aimed to investigate the effect that increasing the verticality of the fibres would have on the thermal conductivity of the samples. In this case, the maximum out of plane angle was varied while the minimum angle was kept at 0. This means that a distribution of out of plane angles was created. In effect, this is a study of how increasing the 3D randomness affects the thermal properties, as the 2D randomness was left at the maximum.

Table 3.10: The geometric and mesh properties of the straight fibre simulations

<b>Max Angle (°)</b>	<b>Volume Fraction (%)</b>	<b>Bounding Size (mm)</b>	<b>Region # of Fibres</b>	<b># of Elements</b>
0	3.42	0.201	54	877,490
30	3.73	0.205	52	1,229,918
36	3.99	0.201	54	574,339
45	3.97	0.204	57	255,918
60	3.78	0.206	50	803,660
90	2.54	0.201	33	677,732
120	2.37	0.201	32	361,114
180	2.32	0.205	36	967,429

Table 3.11: The geometric and mesh properties of the sinusoidal fibre simulations

<b>Max Angle (°)</b>	<b>Volume Fraction (%)</b>	<b>Bounding Size (mm)</b>	<b>Region # of Fibres</b>	<b># of Elements</b>
0	3.98	4	68	1,346,157
30	3.92	4	62	224,160
36	4.07	4	67	238,540
45	4.01	4	64	244,036
60	2.57	4	41	149,541
90	2.46	4	46	1,214,401
120	2.49	4	18	431,835
180	2.58	4	39	409,088

The properties used to generate the simulation geometries are shown in Table 3.10 and Table 3.11. A change in the volume fraction exists as the packing density becomes much smaller when the fibres are more random as fibres at a vertical angle effectively take up more space within the geometry. Compromising by allowing the volume fraction to change was the most effective way to allow this study to be completed, particularly since the volume fraction simulations mean that its effect on the results is well known.



Table 3.12: The mesh settings used in the straight fibre simulations

<b>Max Angle (°)</b>	<b>Relevance</b>	<b>Size Function</b>	<b>Relevance Center</b>	<b>Transition</b>	<b>Span Angle Center</b>
0	100	Uniform	Fine	Fast	None
30	100	Uniform	Fine	Slow	None
36	60	Uniform	Fine	Fast	None
45	0	Adaptive	Coarse	Slow	Coarse
60	60	Uniform	Fine	Slow	None
90	100	Uniform	Fine	Fast	None
120	40	Uniform	Fine	Fast	None
180	100	Curvature	Fine	Fast	Coarse

Table 3.13: The mesh settings used in the sinusoidal fibre simulations

<b>Max Angle (°)</b>	<b>Relevance</b>	<b>Size Function</b>	<b>Relevance Center</b>	<b>Transition</b>	<b>Span Angle Center</b>
0	100	Adaptive	Fine	Fast	Coarse
30	0	Adaptive	Coarse	Fast	Coarse
36	0	Adaptive	Coarse	Fast	Coarse
45	0	Adaptive	Coarse	Fast	Coarse
60	0	Adaptive	Coarse	Fast	Coarse
90	-60	Proximity	Coarse	Fast	Coarse
120	60	Proximity	Coarse	Slow	Fine
180	60	Adaptive	Medium	Fast	Coarse

The mesh settings used in these simulations are shown in Table 3.12 and Table 3.13. These, in general, use less refined mesh settings than the standard values, as the

standard setting produced a mesh with too many elements to be solved without taking excessive time. In some cases, alternative meshing functions had to be used due to the geometry proving challenging for the standard settings to mesh.

### 3.4.4 Bounding Region Size Simulations

This set of simulations aimed to investigate the effect that changing the relative size of the fibres to the size of the bounding region had on the thermal conductivity.

Table 3.14: The geometric and meshing properties of the straight fibre samples used in the Bounding Region size simulations

<b>Bounding Region Size (mm)</b>	<b>Volume Fraction (%)</b>	<b># of Fibres</b>	<b># of Elements</b>
<b>2</b>	2.92	26	360502
<b>2.5</b>	2.98	44	642175
<b>3</b>	2.95	97	918620
<b>3.5</b>	2.99	150	713589
<b>4</b>	2.99	211	932408
<b>5</b>	2.99	387	706804
<b>6</b>	2.95	651	637520

The properties used to generate a geometry for the bounding region size simulations are shown in Table 3.14. The volume fraction was controlled here, as well as the generation length of the fibres (though, due to the clipping present within the simulation, the actual fibre length will vary randomly), so that is largely an investigation of how increasing the number of fibres in the simulation affects the results.

Table 3.15: Mesh settings used in the bounding region size investigation

<b>Relevance</b>	<b>Size</b>	<b>Relevance</b>	<b>Transition</b>	<b>Span</b>	<b>Angle</b>
	<b>Function</b>	<b>Center</b>		<b>Center</b>	
40	Adaptive	Medium	Slow	Medium	
40	Adaptive	Medium	Slow	Medium	
0	Adaptive	Medium	Slow	Medium	
-40	Adaptive	Medium	Slow	Medium	
-20	Adaptive	Medium	Slow	Medium	
0	Adaptive	Medium	Fast	Medium	
-40	Adaptive	Medium	Fast	Medium	

The settings used to generate the mesh are shown in Table 3.15; with the other parameters being kept as default. The decrease in relevance is done to compensate for the increase in fibres causing an increase in the number of elements, partially to ensure that each simulation had roughly the same number of elements and partially to prevent the simulations having too many elements to be solved.

### 3.4.5 Fibre Diameter Simulations

The effect of the fibre diameter was investigated as it is one of the comparatively simple parameters to vary within a fibre blanket, since it can be adjusted by using alternative fibres in the composite.

Table 3.16: Geometric properties of the straight fibre simulations

<b>Diameter (<math>\mu\text{m}</math>)</b>	<b>Volume Fraction (%)</b>	<b>Bounding Region Size (mm)</b>	<b># of Fibres</b>	<b># of Elements</b>
4	2.41	0.20	112	600224
6	2.45	0.20	46	391707
8	2.49	0.40	116	645747
10	2.49	0.40	71	88178
12	2.50	0.40	49	345778
14	2.28	0.61	77	513974
20	2.35	0.92	71	259419
25	2.44	0.91	49	123673
30	2.46	1.00	48	1050811

Table 3.17: Geometric properties of the sinusoidal fibre simulations

<b>Diameter (<math>\mu\text{m}</math>)</b>	<b>Volume Fraction (%)</b>	<b>Bounding Region Size (mm)</b>	<b># of Fibres</b>	<b># of Elements</b>
12	1.46	0.40	23	107714
14	1.32	0.42	15	75157
16	1.42	0.64	27	119736
18	1.42	0.62	16	93966
20	1.45	0.62	17	78281
22	1.45	0.62	15	64088
30	1.40	0.97	21	75099

The properties used to generate geometries for the diameter simulations are shown in Table 3.16 and Table 3.17. The volume fraction was controlled to ensure that the

simulations would be as similar as possible, which in turn meant that the bounding region had to be changed as the thicker diameter fibres could not be packed as efficiently. The number of fibres varies largely as a result of the randomness inherent in the system, and because less fibres are needed to achieve the same volume fraction in the same volume. There is not a strong downwards trend in the number of fibres present because the total simulation volume is increasing.

Table 3.18: Mesh properties used in the straight fibre simulations

<b>Diameter (<math>\mu\text{m}</math>)</b>	<b>Relevance</b>	<b>Relevance Centre</b>	<b>Transition</b>	<b>Span Angle Center</b>
4	60	Fine	Fast	Coarse
6	100	Fine	Fast	Coarse
8	100	Fine	Fast	Coarse
10	-20	Coarse	Fast	Coarse
12	100	Fine	Fast	Coarse
14	100	Fine	Fast	Coarse
20	100	Coarse	Fast	Coarse
25	100	Coarse	Fast	Coarse
30	60	Fine	Slow	Medium

Table 3.19: Mesh properties used in the sinusoidal simulations

<b>Property</b>	<b>Value</b>
Relevance	100
Size Function	Adaptive
Relevance Center	Coarse
Transition	Fast
Span Angle Center	Coarse

The mesh settings used for this simulation are summarised in Table 3.18 and Table 3.19; with unmentioned parameters being kept as default. In particular, the straight fibre simulations used the Adaptive Size Function.

### **3.5 Results and Discussion**

In order to investigate the effect a range of parameters has on the thermal conductivity; each parameter was varied, and a simulation was run to determine what the conductivity was.

Specifically, the effects of the volume fraction; the straight fibre fraction; the fibre length (relative to the size of the representative volume); the fibre diameter and the fibre orientation were investigated, which are primary contributors to the conductivity of the mat.

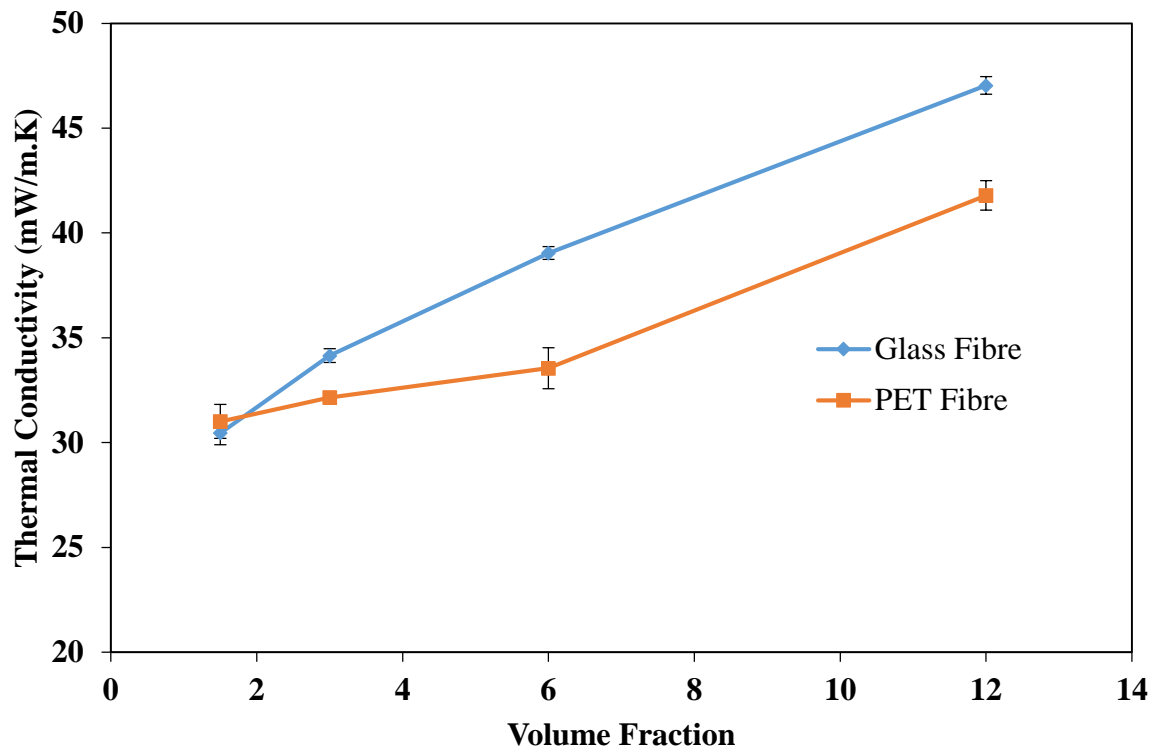


Figure 3.16: The variation of the conductivity with respect to the volume fraction

From Figure 3.16, it can be seen that increasing the volume fraction tends to cause the thermal conductivity to increase. The PET fibres appear to have a slower increase in the conductivity than the glass fibres do, which appears to be because of the shape of the fibres affecting how efficiently they are packed. At very low volume fractions, the effects of convection and radiation would become more prominent and allow a higher level of heat transfer through the blanket, which is not considered in the model. Specifically, at the lower volume fractions, the size of the pores become larger and the bulk movement of air begins to occur. Radiation becomes more important at lower volume fractions simply because the radiation can more easily pass through the fibrous mat with fewer fibres in the way. Volume fractions higher than 12% were not

considered, because the results indicate that increasing the volume fraction further would just lead to the conductivity increasing further.

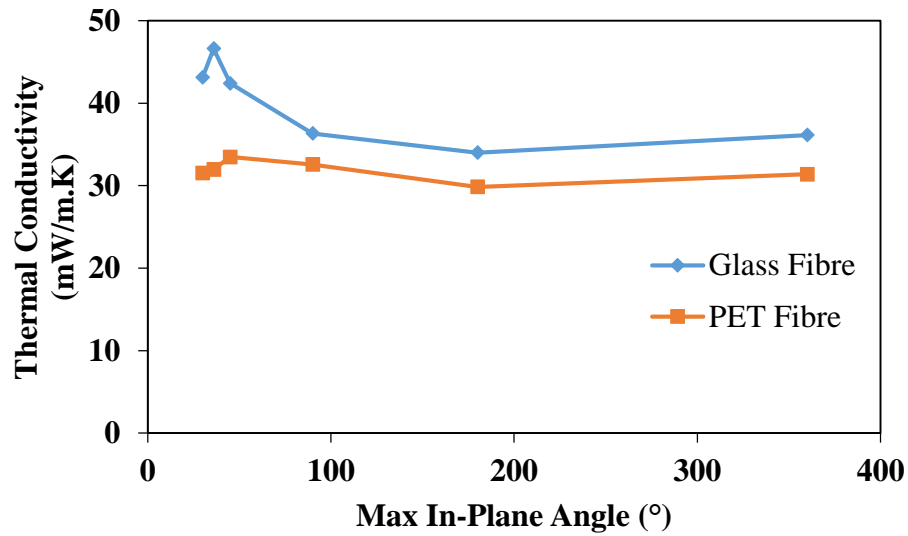


Figure 3.17 : The effect of varying the max in-plane angle on the thermal conductivity

The in-plane orientation appears to not have any significant effect on the conductivity until the fibres become close to unidirectional in orientation, according to the results shown in Figure 3.17. This is largely down to how efficient the solid conduction is, with a high range of angles meaning that in general, heat has to transfer horizontally along the fibre for a significant distance before it can transfer vertically upwards, since the contact points tend to be further apart. However, once the fibres become approximately unidirectional, the optimum heat transfer path allows a much more vertical transfer of heat through the fibre mat.



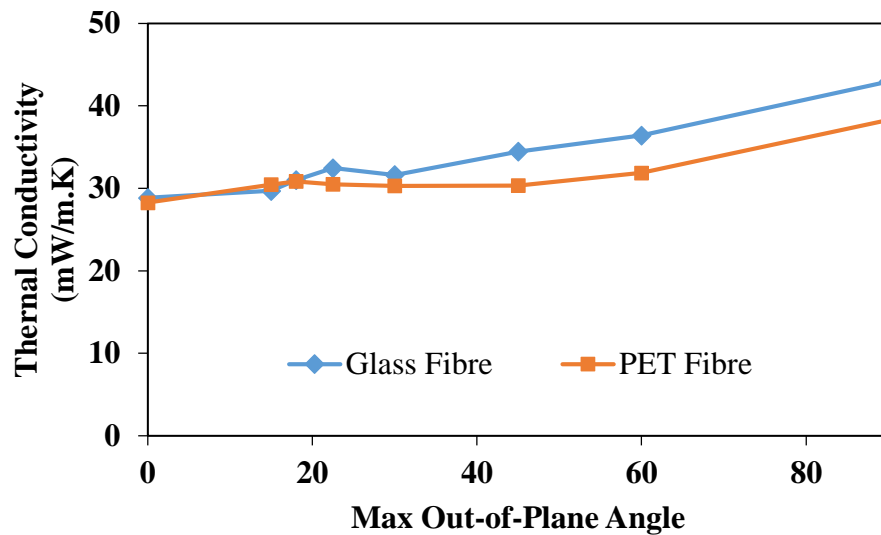


Figure 3.18: The effect of varying the max out-of-plane angle on the thermal conductivity

The out-of-plane orientation has a much more significant effect on the thermal conductivity as shown in Figure 3.18. This is largely because more vertical fibres allow the heat to transfer easily upwards through them, which in turns causes the conductivity of the material to increase. The results suggest that a small level of out-of-plane fibres can be beneficial however, which is probably because it leads to the formation of larger air voids that offset the solid conductivity. At higher angles, the more vertical fibres lead to such a large increase in solid conductivity that larger void sizes are less relevant.

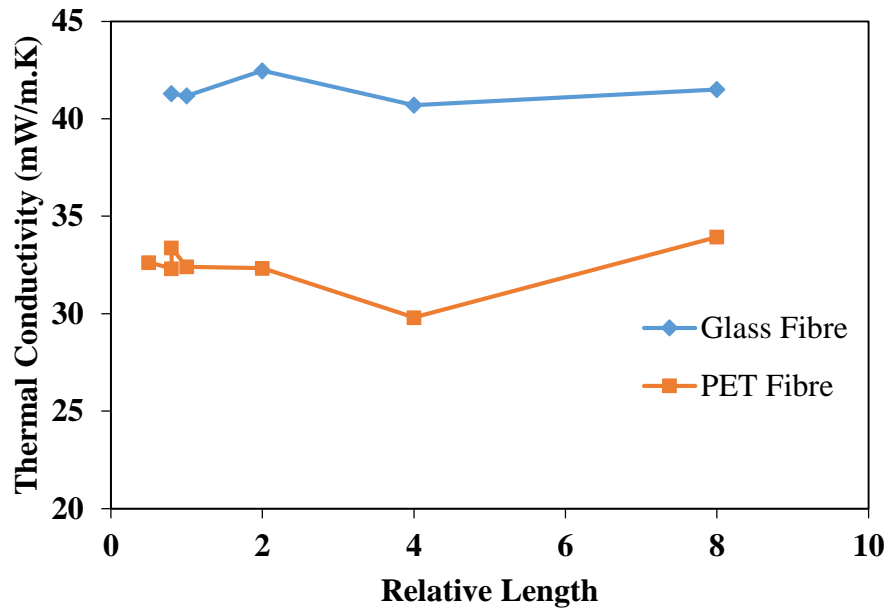


Figure 3.19: The effect of varying the relative fibre length on the thermal conductivity

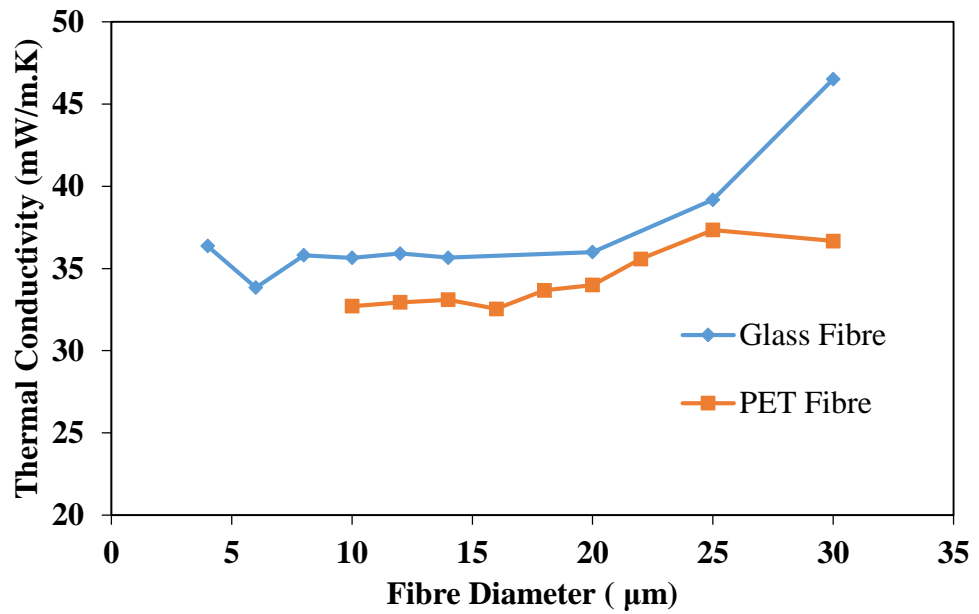


Figure 3.20: The effect of varying the fibre diameter on the thermal conductivity

The fibre length relative to the size of the representative volume and the fibre diameter both appear to have no significant effects on the conductivity, as shown in Figure 3.19 and Figure 3.20. This is expected because while larger fibres can transfer heat more effectively through them, the volume fraction has been conserved and so fewer fibres are present in the test region. This reduction in absolute fibre numbers means that there is less contact between the fibres and the optimum heat transfer path is less efficient. It looks like increasing the diameter further for both fibre types would lead to a large increase in the conductivity, but part of this is likely to be a stacking issue with the fibre geometry. Once the diameter starts getting too large, the amount of fibres that can fit into the representative volume begins to drop to the point where too few fibres are within it for the results to be representative.

### **3.6 Experimental Validation of The Thermal Simulations**

The numerical model was validated against results obtained using a Netzsch Heat Flow Meter, for the case of a pure glass fibre blanket (straight fibre fraction of 1), and a pure PET blanket (straight fibre fraction of 0). Specifically, the thermal conductivity value for a range of mean temperatures (-20-90) and a fixed temperature difference was measured using a heat flow meter for both a PET blanket and a glass fibre blanket, and this was compared to a model with appropriate fibre properties and dimensions so that the accuracy could be determined.

#### **3.6.1 Experimental Setup**

A 30 cm by 30 cm sample of needle punched glass fibres, and an equivalently sized sample made of PET fibres, was measured in a Netzsch 436 Lambda Heat Flow Meter with the mean temperature being varied from -20°C to 70°C and being measured in 10°C intervals, and with a temperature difference of 20°C between the

upper and lower plates. The heat flow meter uses two parallel plates a known distance apart and with a fixed temperature at each plate, and measures the heat flux based on the electric potential in a thermocouple at each allowing the overall heat transfer through the sample to be obtained, and hence the thermal conductivity to be output.

### 3.6.2 Results Comparison

In order to accurately create a computer model of the experiment, the fibre geometry that is generated has to accurately represent the fibre geometry used in the experiment. As a result, the diameter and length of both the PET fibres and the glass fibres was measured, as well as the orientation of the fibres. The heat flow meter needs the dimensions of the sample as well as the mass and calculates the density from this. Since the density of a fibre and the density of air are known, the volume fraction of the fibre mat may be calculated so that a model with the same volume fraction can be used for simulation purposes. Since the representative volume is thinner than the test specimen, the temperature difference was scaled so that the heat transfer experienced by the sample was the same as for the experimental sample, assuming the experimental sample was a bulk material.

Table 3.20: Geometric properties for the thermal validation

<b>Material</b>	<b>Volume Fraction (%)</b>	<b>Bounding Region Size (mm)</b>	<b># of Fibres</b>	<b># of Elements</b>
<b>GF</b>	6.16	0.13	41	844,840
<b>PET</b>	7.17	0.22	35	953,117

Table 3.21: Mesh settings used for the thermal validation

Material	Relevance	Size Function	Relevance Center	Transition	Span Angle Center
GF	60	Uniform	Fine	Slow	-
PET	-20	Proximity and Curvature	Fine	Fast	Fine

The same geometry and mesh settings, generated using the properties shown in Table 3.20 and Table 3.21, was used for all the data points so that the only difference in the models was the temperature gradient applied.

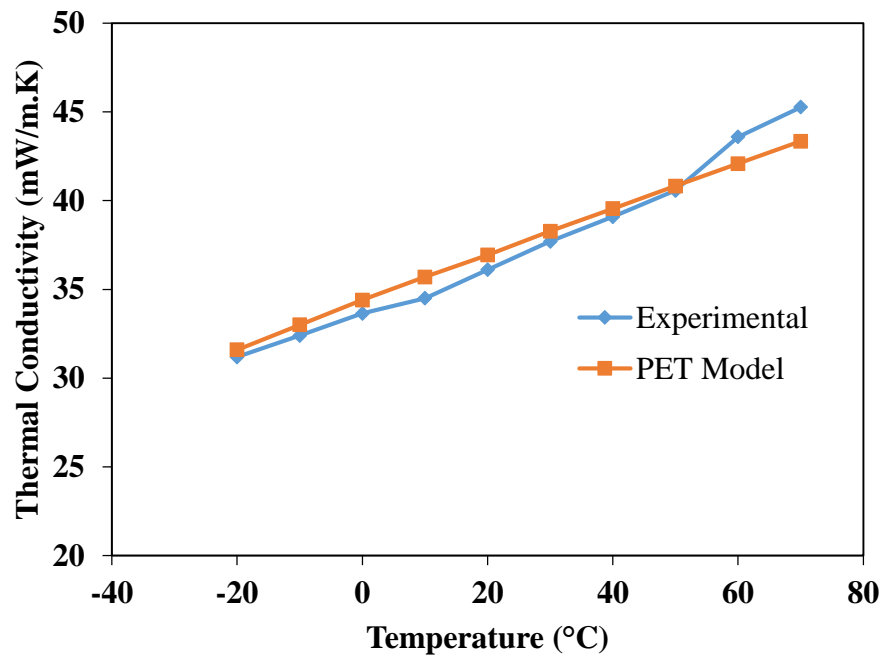


Figure 3.21: Comparison of the PET experimental and model results




Figure 3.22: Comparison of the Glass fibre experimental and model results

The results of the modelling for the PET fibres are shown in Figure 3.21; from which it can be seen that the model is representing the experiments to within 7%. While the absolute values are only partially accurate, the model is accurately demonstrating the trends present in the material. This suggests that the main inaccuracy in the PET model is that the conductivity of the fibres has been overestimated. The slight divergence from the trend that the experimental results show at the higher temperature results suggests that radiation heat transfer may be beginning to have a larger effect, which is not being taken into account by the either model, but appears to be having a larger effect in the PET model. The results of the glass fibre modelling are shown in Figure 3.22; from which it can be seen that there is a great deal of accuracy present in the pure glass fibre model. The model appears to be following the trends accurately.

### 3.7 Conclusions

The validation carried out here suggests that the model, and the modelling approach, are likely producing accurate results for both the fibre blankets. There is a high level of agreement between the model data and the experimental data, and in general both sets of data follow the same trends.

In turn, this suggests but does not prove that the thermal conductivity values produced are accurate, and the trends within are likely to be so. As a result, in general a low volume fraction composite is better for thermal properties but worse mechanically; and the optimum depends on what is desired most. The orientation generally causes the material to be stiffest when the fibres are in the direction that forces are applied in; whereas the fibres being orientated in the direction of heat transfer is undesirable in thermal insulation purposes.

The diameter of the fibres, at least over the range investigated here, appear to have little effect on the thermal or mechanical properties of the composite. It is possible that this occurs as a result of the comparatively small range of diameters investigated, or because the model is not taking some nuance of the conductivity into account that is important here.

The simulations produced by the model are generally computationally light to solve, though the random nature of the geometries produced here mean that it is possible for fibres to be close to contacting in a way that produces thin sections that make the computation time large. The lack of control over parameters such as the volume fraction in the simulation also means that it can be difficult to produce geometries when a specific setup is required; especially if the out-of-plane orientation is allowed to produce large values.

### **3.8 Future Work**

This work did not investigate the effect that the interface between the fibres and the matrix (either the air or the epoxy resin) had on the properties. In both thermal and mechanical cases, this could have a large effect when fibre volume fraction is high, as it controls how easily heat can flow between the two materials and how well the stress in the material can be transferred from the matrix to the fibres. For most of cases studied in this project, the fibre volume fraction is in the low range and the effect of the fibre-matrix interface on fabric thermal conductivity is probably not significant. Nevertheless, future work should be carried out to confirm this.

The effects of convection and radiation were both believed to be negligible and not included here, but there is room for a further study to identify whether there are any cases where this is not true and, in turn, add in mechanisms to allow it to be considered when necessary. Theoretically, this could perhaps be predicted at the geometry generation stage, since it would likely be a property of the geometry instead of the materials used, though it would also increase the computation time.



## **4 Mechanical Modelling of Geometries Containing a Single Type of Fibre**

### **4.1 Introduction**

In order to fully understand the trade-offs of using low thermal conductivity materials in PIP insulation, it is important to also understand the mechanical behaviour of the material. For example, a stiffer blanket requires less centralisers to ensure the inner and outer pipe remain separate, and in turn this means the blanket does not need to be as low in thermal conductivity.

This chapter primarily deals with exploring the mechanical properties of a single fibre type and focusses in turn on predicting the Young's Modulus of the geometry. The effect of the orientation, the fibre diameter and the volume fraction are all investigated here, with validation against numerical models to demonstrate the effectiveness of the simulations.

The simulation methodology is very similar to that of Chapter 3, with the main change being in how the actual simulations are carried out, and the shape of the geometry being generated. The materials being looked at here are the same still: straight fibres representing glass fibres and sinusoidal fibres which represent PET fibres.

### **4.2 Literature Review**

In ductile, homogeneous materials, there is often a distinct yield strength: a point where the elastic deformation applied to the material becomes plastic instead, resulting in permanent deformation. Typically, this is approximated as occurring at 0.2% strain, particularly in metals.

Composite materials behave differently: relatively small applied loads can have permanent effects due to local stress concentrations acting on the fibres and the

distribution of fibres changing upon loading. As a result, composite materials are especially susceptible to stiffness and strength degradation.

There are two main approaches take to analysing composite materials: the micromechanical and the macromechanical [47]. The micromechanical approach focuses on fibres at an individual level and involves modelling the fibre/matrix interaction. This extends towards describing damage as the changes which take place within these materials and the interaction between the components. To contrast, the macromechanical approach aims to treat the composite as if it were a homogeneous material using averaged properties, with a single strength value being used instead to define failure. Young's Modulus Predictions

#### 4.2.1 Self-consistent Micromechanics Model (Hill)

Hill [48] began the development of composite micromechanics by approximating the engineering constants. Hill showed the elastic constants for transversely isotropic materials were independent from the geometry of the composite and could be related by universal relations. The generalised stiffness matrix was created by assuming perfect alignment of continuous fibres within an embedded phase in a composite model:

$$C = \begin{bmatrix} n & l & l & 0 & 0 & 0 \\ l & k+m & k+m & 0 & 0 & 0 \\ l & k+m & k+m & 0 & 0 & 0 \\ 0 & 0 & 0 & y & 0 & 0 \\ 0 & 0 & 0 & 0 & \mu & 0 \\ 0 & 0 & 0 & 0 & 0 & \mu \end{bmatrix} \quad (48)$$

If the system has two planes of symmetry,  $n$  and  $l$  are independent of  $k$ . If the material is transversely isotropic,  $y$  and  $m$  are equal, and only three independent composite moduli exist.

#### 4.2.2 Rule of Mixtures

Being able to derive the stiffness and strength properties from the constituent properties is desirable as it means the mechanical properties can be easily determined theoretically without requiring extensive prototyping and experimentation. For continuous, straight fibres, the 'rule of mixtures' can be used (Equation (49)); which reduces the overall Young's modulus of the material down to the modulus of the components and the ratio of the fibre to matrix.

$$E_n = v_f E_f + v_m E_m \quad (49)$$

When the load is transverse to the fibre direction, the 'inverse rule of mixtures', as shown in Equation (50), is a more accurate approximation of the Young's modulus of the materials.

$$\frac{1}{E_t} = \frac{v_f}{E_f} + \frac{v_m}{E_m} \quad (50)$$

While neither rule of mixtures model is fully accurate, they do provide a reasonable upper and lower bound for the actual composite modulus, particularly in the case of randomly orientated fibres which typically behave somewhere between the two cases.

Halpin [49] improved on these results by incorporating experimental results into the equations, shown in Equations (51) and (52).

$$E_n = \frac{E_m(1 + \eta v_f)}{(1 - \eta v_f)} \quad (51)$$

$$\eta = \frac{\left(\frac{E_f}{E_m} - 1\right)}{\left(\frac{E_f}{E_m} + \xi\right)} \quad (52)$$

The equation for the modulus in any direction can be found through the use of an experimentally derived fitting parameter,  $\xi$ .

$$E_t = \frac{E_m(1 + \xi\eta v_f)}{(1 - \eta v_f)} \quad (53)$$

This same relation can be used to relate the overall shear modulus to the component shear moduli. The Poisson's ratio can also be derived using it.

#### 4.2.3 Eshelby and Mori–Tanaka Method

Eshelby derived this model to allow the mechanical characterisation of elliptical inclusions within an infinitely large matrix [50]. Assuming that the stress-strain field within the inclusions was uniform, analytical solutions can be found to allow the functions to be solved easily.

Mori and Tanaka revised the Eshelby model to consider the volume fraction, and the effect of the mechanical properties of the inclusion, which allowed it to be utilised in fibre composites [51]. Specifically, an elliptical inclusion cannot transfer a stress through it, whereas a fibre can, so the Eshelby model requires modification to include a mechanism to model the interfacial stresses and the stress transfer.

The Mori-Tanaka method approximates the inhomogeneity of the composites through the use of an average matrix stress across the entire infinite geometry. The matrix stress can be determined from the loading boundary conditions at an inhomogeneity, and the effective mechanical properties of the composite can be found by relating the average fibre strain to the average matrix strain. This is done using the strain concentration tensor in Equation (54).

$$A = \left( I + \mathbb{R}(P_m - P_f) \right)^{-1} \quad (54)$$

$$\bar{P} = P_m + v_f(P_f - P_m):A \quad (55)$$

#### 4.2.4 Shear Lag Model

Cox [52] developed a shear-lag model to estimate the composite modulus when the fibres are orientated axially to the loading direction. It assumes the stress forms in the fibre entirely from the shear stress at the interface between the fibre and the matrix. In turn, it assumes the shear stress at an arbitrary point in the matrix is the same as that at an arbitrary point in the fibre which is the same depth into the fibre as the first point.

The parameter  $\bar{\eta}$ , which can be found using Equation (56), describes the behaviour of the shear and normal stress relative to the position in the fibre. Using this, and the parameter  $\lambda$ , found using Equation (57), allows the composite moduli to be found via Equation (58).

$$\bar{\eta} = \sqrt{\frac{2G_m}{r_f^2 P_f \ln\left(\frac{R}{r_f}\right)}} \quad (56)$$

$$\lambda = 1 - \frac{\tanh\left(\frac{\bar{\eta}L}{2}\right)}{\frac{\bar{\eta}l}{2}} \quad (57)$$

$$P = \lambda v_f P_f + v_m P_m \quad (58)$$

The fibre arrangement can be considered in the shear lag model using the function  $R$ , defined in Equation (59).

$$R = r_f \sqrt{\frac{T_R}{v_f}} \quad (59)$$

#### 4.2.5 Christensen and Waals

Christensen and Waals [53] developed a model to predict the elastic modulus of a composite material when it consists of short, randomly orientated fibres. The volume

fraction is the main property that controls the elastic modulus here, with the fibre distribution being of lesser importance. This means that treating it at a macroscopic level, instead of dealing with the micromechanics, is valid.

This theory was derived from the stress-strain constitutive equations by determining an average value for the ratios of the stress-strain constituents. The resulting equation relates the overall stiffness to the components Young's, bulk and shear moduli and the Poisson's ratio.

$$E = \frac{[E_{11} + K_{23}(4v_{p_1}^2 + 8v_{p_1} + 4)][E_{11} + K_{23}(4v_{p_1}^2 - 4v_{p_1} + 1) + 6(\mu_{12} + \mu_{23})]}{3[2E_{11} + K_{23}(8v_{p_1}^2 + 12v_{p_1} + 7) + 2(\mu_{12} + \mu_{23})]} \quad (60)$$

Simpler versions of this can be derived via the use of experimental results and by plotting Equation (60) to allow the creation of best fit models.

By plotting this equation and with the aid of experimental results the theory can be stripped down to two more manageable forms. One is for random fibre geometries in two dimensions, shown in Equation (61).

$$E \approx E_f \frac{v_f}{6} + E_m(1 + v_f) \quad (61)$$

Another is for random fibre geometries in three dimensions, shown in Equation (62).

$$E \approx E_f \frac{v_f}{6} + E_m(1 + v_f[1 + v_{p_m}]) \quad (62)$$

The advantage of these equations is that they only rely on the volume fraction of the composite, and generally well-known material properties. The main downside is that assumptions used in the derivation of these equations mean they are only reliable for fibre volume fractions below 20%; though this is not a significant impediment in this work due to the low volume fractions typically looked at.

#### 4.2.6 Manera

Manera [54] derived a model based on Puck's equations, using the method of least squares to simplify them, and eliminating terms that are generally negligible. The result of this is the model shown in in Equation (63).

$$E = v_f \left( \frac{16}{45} E_f + 2E_m \right) + \frac{8}{9} E_m \quad (63)$$

Due to the assumptions made to eliminate terms, it's only suitable in specific cases: namely, when the volume fraction is between 10% and 40%; when the resin matrix has a Young's Modulus between 2 and 4GPa; and when the fibres are between 5 and 30 microns long and have an aspect ratio of more than 300. These assumptions are generally applicable in this work, albeit the volume fraction is often lower than the minimum value.

#### 4.2.7 Pan

Pan [55] derived an equation to determine the modulus which utilised the area fraction of the fibres,  $A_f$ . This parameter can be used to help describe the orientation of the fibres, which is important in general and particularly in the case of random, short fibre composites. The generalised form is shown in Equation (64).

$$E = E_f \Omega(\theta, \Phi) v_f + E_m [1 - \Omega(\theta, \Phi) v_f] \quad (64)$$

Where in a particular direction the probability density function is defined as:

$$\frac{\Omega(\theta, \Phi)}{(\theta, \Phi)} = \frac{A_f}{v_f} \quad (65)$$

For the case of fibres which are random in 2D (i.e., aligned in the 3<sup>rd</sup> dimension), Equation (66) can be used; while Equation (67) is valid for fully random fibre geometries.

$$E = E_f \frac{v_f}{\pi} + E_m \left( 1 - \frac{v_f}{\pi} \right) \quad (66)$$

$$E = E_f \frac{v_f}{2\pi} + E_m \left(1 - \frac{v_f}{2\pi}\right) \quad (67)$$

To summarise, it can be seen from the literature above that a range of micromechanical models have been developed to assess material stiffness of fibre-matrix composites. However, it appears that none takes consideration of individual fibres. This is quite understandable given the stiffness is an effectively bulk effect typically determined by the microstructure. However, the property such as material strength is often determined by local structural details and would benefit from a more detailed modelling approach considering individual fibre architecture.

### 4.3 Simulation Setup

Table 4.1: Material properties used in the mechanical simulations

<b>Material</b>	<b>Young's Modulus (GPa)</b>	<b>Poisson's Ratio</b>	<b>Tensile Strength (MPa)</b>	<b>Fibre Diameter (<math>\mu\text{m}</math>)</b>
Glass Fibre	72	0.21	1950	6
PET Fibre	2.8	0.37	80	14
Epoxy Matrix	3.9	0.35	45.7	-

The mechanical properties used throughout the simulations are listed in Table 4.1. As the simulations are intended to allow comparison between the geometric properties, the accuracy of the material properties here is not hugely important. As such, generic properties for the three materials were used in place of ones corresponding to specific composite compositions.



The fibre alignment for the mechanical simulations was constrained so the fibres were all roughly orientated in the direction the stress was applied in.




Figure 4.1: Illustration of the boundary conditions applied to the mechanical simulations

Table 4.2: Boundary conditions used in the mechanical simulations

Label	Parameter
A	4N force
B	Displacement support restricting movement in the same direction as A
C	Fixed support to prevent rigid body motion

The boundary conditions used in these simulations are shown in Figure 4.1 and Table 4.2. The 4N force chosen here was largely an arbitrary choice designed to be a

relatively small load on the sample. The fixed support is in the centre of the face to ensure that it does not introduce a moment around the bottom of the sample; and the movement in the tensile direction is fixed to 0 to replicate.

## 4.4 Results and Discussion

### 4.4.1 Volume Fraction

The first analysis that should be the most significant factor contributing to the mechanical properties is the volume fraction. The boundary region was reduced in stages and the number of fibres chosen appropriately to control the volume fraction. Results were taken from ANSYS to Microsoft Excel to calculate and plot the composite's Young's Modulus and tensile yield strength.




Figure 4.2: Young's modulus of the composite at varying volume fractions of glass fibre

There is a moderate amount of variation in the results, but a trend line does show an overall rise of stiffness as the volume fraction increases. It is difficult to go beyond the upper and lower values of volume fraction used. Attempting to achieve higher volume fractions produces too few fibres; since the geometry has to be further reduced in size to force them closer together on average; whereas decreasing the volume fraction further means a much larger increase in the number of fibres present (as the bounding region becomes larger). This causes a corresponding increase in the computing power required.

The equation of the trend line in Figure 4.2 suggests that at a volume fraction of 0%, the modulus of the simulation would be around 4.45GPa. The epoxy used has a modulus of 3.94GPa, which suggests that the simulation is behaving relatively accurately.

#### 4.4.2 Fibre Orientation

The effect of the fibre orientation was also investigated. Specifically, the effect of varying the fibre angle in one direction was studied (since the geometry is symmetrical, the actual variation here would be the same regardless of which direction it is, as long as only one direction was varied).




Figure 4.3: Same set of fibres viewed from two different orientations, with the fibres being visibly uniform in one plane. The arrow indicates the direction the force is applied in.

The through-plane fibre orientation is kept uniform and the in-plane is varied in steps of 15 degrees. The volume fraction was kept near 5% but not exactly as the code cannot be controlled that way. Each angle was repeated for 3 results and an average value was plotted.

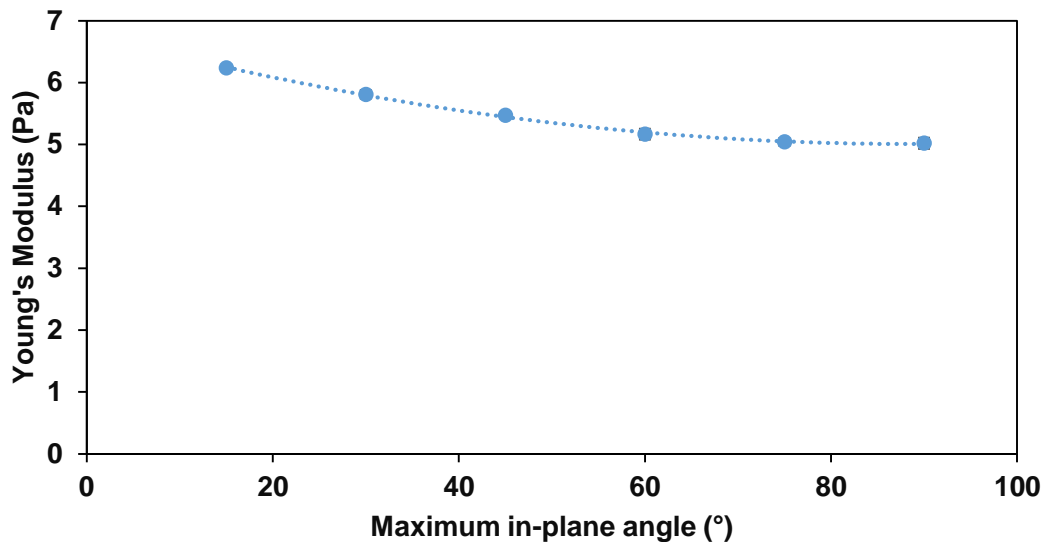


Figure 4.4: Young's modulus against random in-plane orientation<sup>6</sup>

There is a downwards trend in the results which appears to level off at 90°, which is to be expected as increasing the angle beyond this point should not change the fibre distribution as the fibres would already be at their most random.

There is a downwards trend seeming to level off at 90 degrees as it should since the plane cannot become more random than that. The trend line fits well to a polynomial trend line but there is no known reason for this type of relationship. The randomness

---

<sup>6</sup> Error bars indicating one standard deviation are present here; the high degree of similarity between each sample means they are not visible.

lowers the stiffness, and this is expected as fewer fibres are in the prime orientation of directly facing the load.

#### 4.4.3 Fibre Diameter

The fibre diameter was changed from 8 microns to 16 microns and variation in volume fraction was minimised.

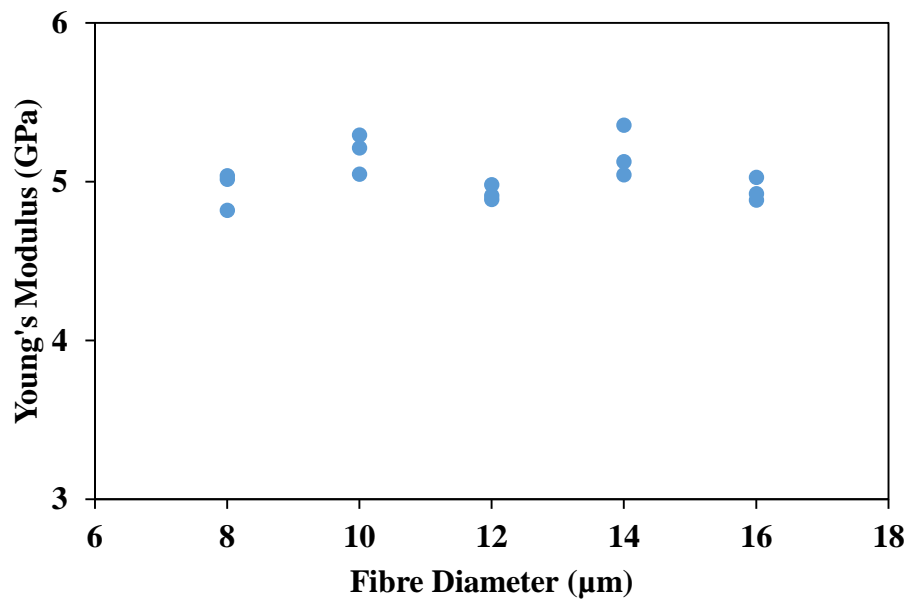


Figure 4.5: Young's Modulus against the diameter of the fibres used.

There is no trend here, so results suggest fibre diameter alone does not affect the modulus of elasticity. Variation here is suspected to be due to the difference in volume fraction as it could not be kept completely constant. It was limited from 4.9% to 5.1% in MATLAB, but this resulted in a true range of around 4.5% to 5.5%. No apparent change in fibre diameter results is useful as it provides another method of controlling the volume fraction.

## 4.5 Theoretical Validation of Mechanical Simulations

To verify if the model is providing a reasonable output the volume fraction results are compared individually to the theoretical models previously mentioned. The theoretical models were added to excel and increments of volume fraction were inputted then plotted.

### 4.5.1 Rule of Mixtures

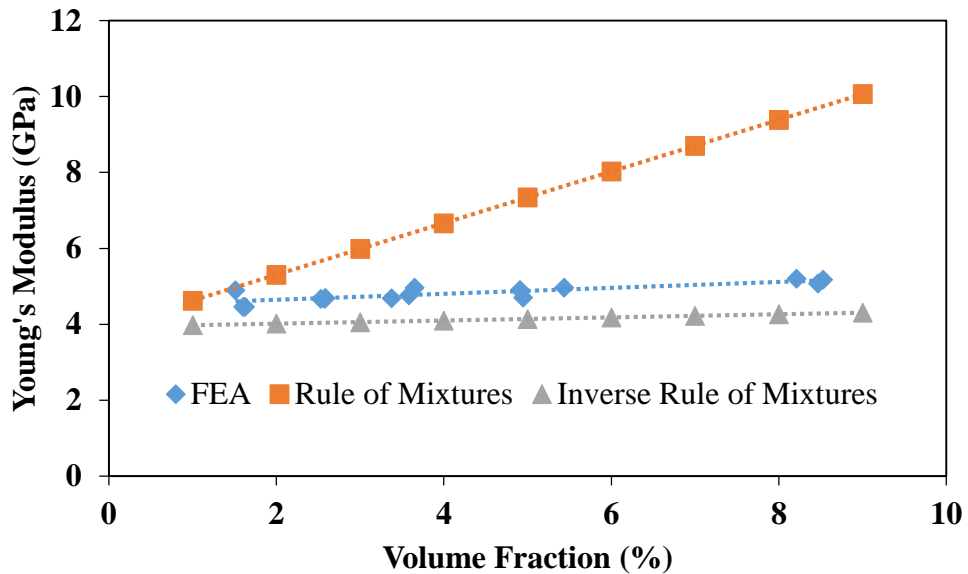


Figure 4.6: Comparison of the FEA results with the Rule of Mixtures and Inverse Rule of Mixtures results

The first comparison was to the two Rule of Mixtures models, shown in Figure 4.6, which behave as upper and lower bounds for the Young's Modulus value. It can be seen here that the FEA results fall within these bounds, which suggests that the model is behaving accurately. The gap between the bounds is, particularly at higher volume fractions, fairly large and so the model fitting within them does not guarantee accuracy here. This is due to the fact that at higher fibre volume fraction, the fibre architecture

is subjected to more “artificial” alterations brought about by the fibre generation process.

The model appears to be beginning to exceed the upper bound at very low volume fractions, though the difference between the bounds at these values is minimal and it may just be due to the random nature of the simulations.

A large degree of agreement is present in the trend of the FEA results and the Inverse Rule of Mixtures results, even if the FEA results are shifted upwards; which is unexpected as the random orientation here should be closer to the standard rule of models geometry.

#### 4.5.2 Christensen and Waals




Figure 4.7: Comparison of the FEA results with the 2D and 3D Christensen and Waals model results

The FEA results in Figure 4.7 are showing reasonable agreement with the Christensen and Waals model, particularly at the higher volume fractions. This is a good sign as the Christensen and Waals model is supported by experimental evidence, and the close agreement here implies the FEA results are behaving accurately. There is some overestimation at lower conductivities, though the cause is uncertain. It seems likely that the simulations are overestimating the Modulus by overestimating the effect of the fibres, as the trend here implies the simulation modulus would be above the young's modulus of the epoxy resin at a volume fraction of 0.

#### **4.5.3 Manera**

Manera's equation has also been plotted. There is less agreement here than with Christensen and Waals other than the region of intersection at about 4% volume fraction. Remembering that Manera's equations are supposed to be valid in the range of 10% to 40% is unexpected as the results suggest increasing divergence before 10% is even reached.






Figure 4.8: Comparison of the FEA results with the Manera model results

#### 4.5.4 Pan

Finally, Pan's theory is compared to the FEA data points. Pan's correlation seems better than Manera but not quite as close as Christensen and Waals. Pan could also be useful to compare to fibre orientation if more time were spent understanding the statistical theory underlying it.




Figure 4.9: Comparison of the FEA results to the 2D and 3D Pan models.

#### 4.5.5 Validation Summary

Overall, the theory suggests that the FEA simulation is producing realistic data. The range is limited however, and more data especially at higher volume fractions would be key to properly testing the FEA validity as it seems as though the gradient of the FEA trend line is not high enough, assuming the theories are accurate.

#### 4.6 Conclusions

The Finite Element data is producing a good fit against most of the theoretical models it was tested against, particularly at the relatively low volume fractions considered. That said, the data was only modelled at low volume fractions, and so there may be more of a discrepancy at higher volume fractions.

It can be seen that mechanically the orientation and volume fraction are the dominant parameters here, where they have a more significant effect on the fibre properties than the other factors such as fibre diameter and length.

#### **4.7 Future Work**

The effect of the fibre properties on the strength of composites was not investigated here either. While the effect on Young's Modulus can be predicted relatively easily, the effect on the strength requires the development of a robust failure criterion that can adequately determine when failure occurs within the system.

## 5 Modelling of Fibre Geometries Containing Straight and Curved Fibre Types

### 5.1 Introduction

In order to determine the effect of having a mixture of straight and curled fibres within the simulation, a series of simulations to evaluate the effect of varying the “straight fibre fraction” was done. The straight fibre fraction is defined as indicated in Equation (68). The volumetric ratio was used here so that the ratio could be more easily calculated within the MATLAB code: while specifying a density for both fibre types could be done to allow the fibre masses to be calculated, it was decided to not introduce an additional two variables for simplicity.

$$\text{Straight fibre fraction} = \frac{\text{Straight fibre volume}}{\text{Total fibre volume}} \quad (68)$$

In order to calculate the total fibre volume, the volume of each individual fibre has to be calculated. This is relatively straightforward for the straight fibres, which can be calculated using Equation (69), which is just the formula for calculating the volume of a cylinder. The exact volume of the curled fibres was not able to be easily approximated since the exact shape of the spline is determined by APDL and not the MATLAB generating code; however it was approximated using Equation (70). This approximates the volume of the curled fibre as being composed of  $n$  different cylinders extending from the centre line to a distance equal to the “fake diameter”, which is the amplitude of the fibres.

$$V_{SF} = \frac{\pi}{4} \cdot D_{SF}^2 \cdot L_{SF} \quad (69)$$

$$V_{CF} = \frac{\pi}{4} \cdot D_{CF}^2 \cdot n \sqrt{\left(\frac{L_{CF}}{n}\right)^2 + D_{\text{Fake}}^2} \quad (70)$$

## 5.2 Simulation Setup

The same algorithm outlined in Chapter 3 was utilised here, where the straight fibre fraction was varied as desired to setup a series of geometries that could be used to investigate the effect that blending the fibres together has on the overall conductivity. Three simulations were carried out here: two to investigate blends of two geometries of fibre blended together when they are all the same material, and one to investigate a blend of two different fibre types and materials.

The value generated in MATLAB is an approximation since the volume of the curled fibres is not precisely known until it has been generated in APDL. The fibre ratio can be controlled to some degree, though the stochastic nature of the fibre generation means that the volume of the fibres varies depending on the positioning of the fibre with respect to the boundary region and the fibre length.

The properties of the geometry used are shown in Table 5.1. The fibre ratio here has been calculated from the actual volumes of fibres in the simulation; and in general, is consistent with the desired fibre ratio; with only 0.35 being somewhat different. This discrepancy is due to a curled fibre failing to generate properly, and the simulation putting an additional straight fibre in. The in-plane orientation for the fibres was allowed to vary fully; while the out-of-plane orientation was constrained to  $0^\circ$ , so that all the fibres were perpendicular to the direction of heat transfer. This is more representative to typical insulation fabrics while keeping geometrical simplicity without compromising the accuracy.

Table 5.1: The properties of the geometries used in the mixed fibre modelling simulations

<b>Fibre Ratio</b>	<b>Volume Fraction (%)</b>	<b>Boundary (mm)</b>	<b>Size</b>	<b># of Straight Fibres</b>	<b># of Curled Fibres</b>
--------------------	----------------------------	----------------------	-------------	-----------------------------	---------------------------

<b>0</b>	2.34	3.15	0	34
<b>0.35</b>	2.20	3.15	16	35
<b>0.5</b>	2.32	0.31	27	35
<b>0.71</b>	2.27	0.31	37	23
<b>1</b>	2.49	0.30	52	0

As in Chapter 3, the contact settings were kept to their default, so that fibre/matrix contact was treated as having infinite conductivity. This also means that any fibre/fibre contact that exists has the same assumption.

The mesh settings used throughout the mixed fibre simulations are shown in Table 5.2. These are based on the standard ones used throughout this work to ensure the results throughout are consistent.

Table 5.2: Mesh settings used in the mixed fibre modelling simulations

<b>Setting</b>	<b>Value</b>
<b>Relevance</b>	0
<b>Size Function</b>	Adaptive
<b>Relevance Centre</b>	Medium
<b>Transition</b>	Slow
<b>Span Angle Centre</b>	Medium
<b>Automatic Mesh Based Defeaturing</b>	On
<b>Defeature Size</b>	Default

These settings produced a mesh corresponding with the numbers in Table 5.3; where the same geometry and mesh were used for all three simulation types.

Table 5.3: The number of elements present in each mixed fibre simulation

Fibre Ratio	Number of Elements	Number of Nodes
0	835103	1260736
0.3	752023	1123559
0.5	899729	1333972
0.7	677312	1098433
1	578414	1020947

### 5.3 Results and discussion




Figure 5.1: The effect of varying the mixing ratio on the thermal conductivity

The results of the simulation are shown in Figure 5.1. The results at 0 and 1 make sense: when the simulation is pure straight fibres, the mixed and GF results are the same because they both have the same conductivity; and the same is true when the simulation only contains curled fibres. The general trends also make sense: as the mixed fibres becomes more dominated by straight fibres, it becomes closer to the

glass fibre result, whereas at low fibre ratios, the results are much more similar to the PET fibre results.

The relatively small variation in the thermal conductivity here is interesting: the PET has a conductivity that's approximately 5.5 times lower than the glass fibres conductivity; and so the expectation would be that changing this conductivity would have a much larger impact on the thermal conductivity than it does. The comparatively low change is likely due to the relatively low volume fraction that the simulations are taking place at: since the fibres are only taking up approximately 2.5% of the volume of the material, their effect on the conductivity is minimal compared to the effect of the conduction through the air around the fibres.




Figure 5.2: The variation of the conductivity based on the volume fraction

The strange results are unlikely to be a result of the variation in the volume fraction here; plotting the thermal conductivity against the volume fraction instead of the fibre



ratio, as shown in Figure 5.2, does not show a strong trend: for all three materials, the conductivity is approximately constant regardless of the volume fraction. The geometry generation was controlled to attempt to keep the volume fraction as similar as possible throughout the simulations, and so there is not a wide range of volume fractions here. The relative lack of correlation here implies that the fibre ratio is having a large effect on the results, since the volume fraction and the fibre ratio are the two big parameters expected to affect the heat transfer through the simulated geometry.

## **5.4 Conclusions**

The results show that, from a thermal perspective, blending the fibre types does not typically reduce the conductivity beyond that which could be obtained with a single fibre type. The validation done implies that as the single fibre blankets are being modelled accurately, combining them is likely also producing accurate results.

The different materials are behaving somewhat differently here, though the two pure fibre types are following the same general trend. The mixed fibre simulation is following a different trend, which is to be expected as it has to bridge from the lower conductivity PET fibre results to the higher conductivity glass fibre as the relative ratio of those two changes.

## **5.5 Future Work**

A more in-depth look at the effect of mixing the fibre types could be carried out to investigate whether the orientation of the fibres has a more significant effect when the fibres are non-straight, for example. The results obtained so far indicate that the overall composite is likely to behave very similar to the individual components.

This work only looked at a mix of two fibre types, but theoretically the principles could be expanded to a blend of multiple types of straight/sinusoidal fibres. Indeed, the geometry generation process the algorithm developed here uses can easily be

changed to allow alternative fibre geometries, including square or hexagonal cross sectioned fibres. This could be done with both the straight and the sinusoidal fibres.

The mechanical properties were not investigated and examining the trade-off between the thermal and mechanical performance could be investigated. This would allow a better understanding of when, and why, mixed fibre blankets like this should be utilised.

While more complex fibre shapes were briefly investigated over the course of this work, truly random fibre geometries were never simulated, and adding that in would allow a greater degree of flexibility in the fibre composite modelling performed here. The difficulty with that is that meshing the very complex geometries of a truly random fibre is challenging and not something that commercial meshing algorithms are designed to handle. Indeed, meshing can be a challenge even with these comparatively simple fibre geometries, since the random nature of the fibre geometries means that they can end up being close to contacting or the edge of the bounding region, for example.

## **6 Modelling of Aerogel/Fibre Composites**

### **6.1 Introduction**

In this chapter, composites consisting of straight fibres mixed with a range of aerogel types are modelled, and the results of this discussed. The aim is to explore the creation of flexible, aerogel composites where the fibre reinforcement helps overcome the inherent brittleness of aerogel materials.

Polyimide aerogels were investigated at the request of the sponsoring company, Blueshift International Materials, who commercially produce both powder and stock shapes from them. They are a novel polymer based aerogel originally developed by NASA for use in high temperature, flexible insulation [56]. Resorcinol-formaldehyde (RF) aerogels are the benchmark of polymer aerogels, with a low conductivity and large body of research behind them; with silica aerogels being the benchmark for all aerogels, with more research and a lower conductivity than polymer aerogels have. Composites using Silica, Resorcinol-Formaldehyde and Polyimide based aerogels were examined. Silica is the industry standard aerogel and has been extensively studied; but is prone to producing hazardous dust which is also hydrophobic and unpleasant to handle without using gloves. RF aerogels are easier to manufacture and produce less dust. Polyamide aerogels have high temperature stability and are good in high temperature applications. While they are all different materials, all three feature the very low thermal conductivities typically associated with aerogels, with silicon aerogels typically having the lowest conductivity and polyamide aerogels the highest. The cost of this is that they all have the same highly porous structure which causes them to be mechanically weak.

The addition of a fibre blanket helps to improve this by providing a flexible scaffolding matrix that the aerogel structure can be supported by, with a minimal impact on the actual thermal conductivity of the material. This allows more mechanically stable

aerogel systems to be used in cases where minimising the thermal conductivity is not the only requirement.

In subsea applications, this makes assembling the PiP easier as the insulation layer can be wrapped around the pipe; whereas solid aerogel needs to be precisely shaped and machine to fit the pipe without having air gaps.

## 6.2 Simulation Setup

The effect of varying a range of fibre properties on the three different kinds of aerogel was investigated. Silicon aerogel, with a thermal conductivity of 17mW/mK, and Polyamide aerogel, with a thermal conductivity of 30mW/mK, were modelled, as well as two speculative conductivity values for resorcinol-formaldehyde.

The speculative values were calculated from the measured thermal conductivity of a RF aerogel/PET fibre composite blanket, which were believed to represent the upper and lower bound for the thermal conductivity. These values were obtained by rearranging the Series (Equation (24)) and Parallel (Equation (23)) models to allow the aerogel conductivity to be determined. This produced Equation (71) for the Series model, and Equation (72) for the Parallel model.

$$k_a = \frac{k_c - v_f k_f}{1 - v_f} \quad (71)$$

$$k_a = \frac{1 - v_f}{\frac{1}{k_c} - \frac{v_f}{k_f}} \quad (72)$$

Experimental thermal conductivity data for an RF aerogel/PET fibre composite was produced. The RF aerogel used 50% calcium carbonate and 50% sodium carbonate as a catalyst, with a resorcinol to catalyst ratio of 650, and 30% solids during the sol-gel process (i.e., it was 70% deionised water). The measurement of the conductivity was done on a 30cm by 30cm blanket using a Netzsch 436/3/1E Heat Flow Meter.




Figure 6.1: Comparison of the derived lower and upper bounds for the thermal conductivity, where Experimental is the experimentally derived blanket conductivity.

From this, and knowledge of the conductivity of the PET blanket, an approximation of the aerogel conductivity can be produced so that the material can be treated as a two-phase system. An upper and lower bound for this conductivity were calculated, the results of which are shown in Figure 6.1.

In this case, the series model is tending to produce a lower component conductivity than the parallel model is as a result of the series model tending to overestimate the conductivity. In turn, for the series model to produce the same conductivity value as the parallel model, it needs a lower component conductivity. The experimental conductivity is the blanket conductivity and not specifically the aerogel conductivity here; and ‘aerogel conductivity’ is something of a misnomer as it is actually a combination of the conductivity of the aerogel and the air in the pores inside the blanket.

Various fibre properties were simulated to test their effect on the aerogel conductivity. Specifically, the volume fraction and the fibre diameter, length and orientation were considered, with Chapter 7 looks at the effect that varying the aerogel particles have on the composite material when the matrix is considered to be a bulk material.

### 6.2.1 Volume fraction

The volume fraction being used here is the volume of the fibres in the composite: the rest of the composite is assumed to be an aerogel/air mixture. Four volume fractions were simulated for the four material properties: 1.5%, 3%, 6% and 12%. This was chosen to give a broad overview of the behaviour of the system. The same geometry generation method was used here as laid out in Chapter 3, which in turn means that the same limitations that prevent high volume fraction geometries being generated then are preventing higher volume fractions here. The parameters used for the geometry are shown in Table 6.1: the size of the bounding region is being varied here as the volume fraction is not a directly controlled parameter and instead depends on how well the fibres are packed. In turn, decreasing the size of the bounding region causes the number of fibres to decrease.

Table 6.1: The volume fraction simulations' geometry parameters

<b>Volume Fraction</b>	<b>Bounding Region Side Length (mm)</b>	<b># of Fibres</b>
<b>1.5%</b>	6.07	125
<b>3%</b>	3.06	67
<b>6%</b>	1.52	34
<b>12%</b>	0.72	18

The contact settings were kept to the default ones, which produce no interfacial layer and no interfacial thermal resistance. The mesh settings used throughout with one

exception are listed in Table 6.2; with the one exception being the ‘Relevance’ for the 1.5% sample, which was set to -100. This change was entirely done to limit the number of elements present in the simulation, as a relevance of 40 produced too many elements to be ran on the computer hardware available.

The consequence of using the same mesh settings for all four geometries is that the number of elements varies significantly through the four simulations, as shown in Table 6.3. This was deemed to be acceptable as it means the actual mesh sizing between the simulations was more consistent, instead of having significantly smaller elements on the smaller geometries, since the fibre size is the same in all the simulations.

Table 6.2: Volume fraction mesh settings

<b>Mesh Setting</b>	<b>Value</b>
<b>Relevance</b>	40
<b>Size Function</b>	Adaptive
<b>Relevance Centre</b>	Medium
<b>Smoothing</b>	Medium
<b>Transition</b>	Slow
<b>Span Angle Centre</b>	Medium

Table 6.3: Element count for the volume fraction simulations

<b>Volume fraction</b>	<b># of Elements</b>	<b># of Nodes</b>
1.5%	1148135	1797502
3%	948423	1858412
6%	331588	735710

12%	189653	360028
-----	--------	--------

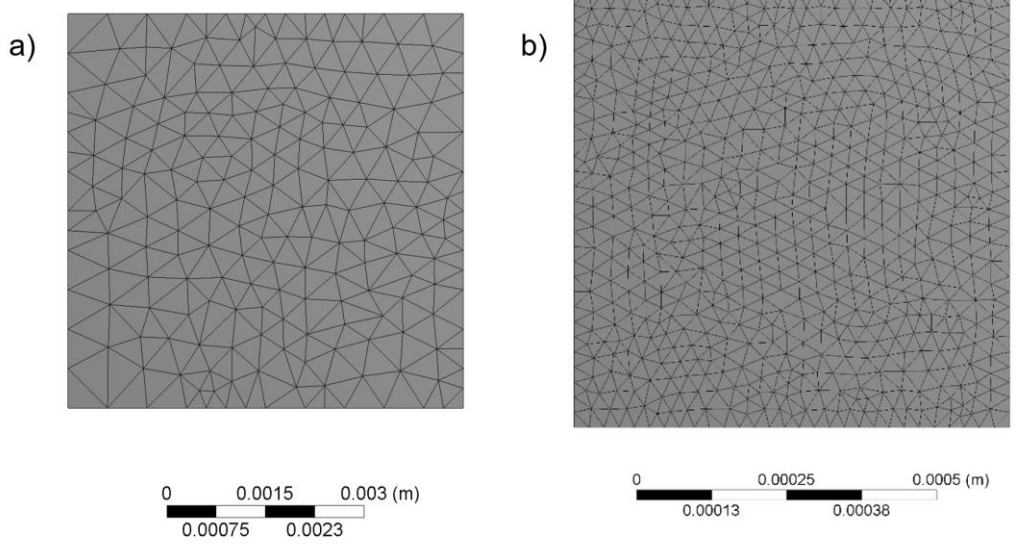


Figure 6.2: Comparison of the mesh used for the 1.5% sample (a) and the 12% sample (b)




Figure 6.3: Example of the mesh produced on a fibre in the 1.5% VF simulation






Figure 6.4: Example of the mesh produced on a fibre in the 12% VF simulation

As can be seen in




Figure 6.2, Figure 6.3 and Figure 6.4, the mesh used on the geometries was kept as similar as possible. The stochastic nature of the geometry means that the mesh cannot be exactly the same in every simulation, since the fibre geometry itself changes, but the algorithm to produce a mesh was kept tuned to the same parameters as far as possible to ensure that the results are comparable.

### 6.2.2 Orientation

Both the in-plane orientation and the out-of-plane orientation were investigated here. The results are, as in previous examinations, defined by the maximum angle allowed, with the minimum being kept at  $0^\circ$ . The parameters used to generate a geometries for in-plane and out-of-plane orientated fibres are shown in Table 6.4 and Table 6.5 respectively.

The maximum out-of-plane orientation allowed was 90°, since this is when fibres become vertical, and increasing it further does not allow more variation due to symmetry.

Table 6.4: Geometry and mesh properties of the in-plane orientation simulations

<b>Maximum Angle (°)</b>	<b>Volume Fraction (%)</b>	<b>Side (µm)</b>	<b>Length</b>	<b># of Fibres</b>	<b># of Elements</b>
<b>30</b>	6.87	400		244	307902
<b>36</b>	6.53	400		241	1144043
<b>36</b>	6.90	400		209	302546
<b>45</b>	4.94	400		187	1049099
<b>45</b>	6.55	400		150	1246000
<b>90</b>	4.85	400		140	997491
<b>180</b>	4.84	400		142	833542
<b>360</b>	4.84	400		148	210777

Table 6.5: The geometry and mesh parameters for the out-of-plane angle simulations

<b>Maximum Angle (°)</b>	<b>Volume Fraction (%)</b>	<b>Side (µm)</b>	<b>Length</b>	<b># of Fibres</b>	<b># of Elements</b>
<b>0</b>	4.06	200		54	464377
<b>15</b>	3.94	200		52	1136727
<b>18</b>	4.03	200		54	513210
<b>22.5</b>	4.09	200		57	495630
<b>30</b>	3.92	200		50	459856
<b>30</b>	2.57	200		37	253911

<b>45</b>	2.45	200	33	281767
<b>90</b>	2.42	200	36	324144

### 6.2.3 Bounding Region Size

In order to investigate whether the fibre length had any effect on the simulation results, the bounding region size was varied. This changes the relative size of the bounding region to the fibres and is equivalent to changing the fibre length in effect. The parameters used for this are listed in Table 6.6.

Table 6.6: Geometric properties for the bounding region size simulations

<b>Side Length (<math>\mu\text{m}</math>)</b>	<b>Volume Fraction (%)</b>	<b># of Fibres</b>	<b># of Elements</b>
2	3.00	26	360502
2.5	3.02	46	642175
3	3.04	107	918620
3.5	3.01	166	713589
4	3.04	228	932408
5	3.02	424	706804
6	2.97	694	637520

### 6.2.4 Fibre Diameter

The diameter simulations were carried out using random in-plane only fibre geometries. The key geometry generation properties are shown in Table 6.7. Increasing the side length was necessary to ensure that the overall volume fraction within the simulation remained constant.

Table 6.7: Simulation parameters for the diameter simulations

<b>Diameter (<math>\mu\text{m}</math>)</b>	<b>Volume Fraction (%)</b>	<b>Side (<math>\mu\text{m}</math>)</b>	<b>Length</b>	<b># of Fibres</b>	<b># of Elements</b>
<b>4</b>	2.52	200		116	600224
<b>6</b>	2.50	200		51	391707
<b>8</b>	2.50	400		118	645747
<b>10</b>	2.50	400		74	88178
<b>12</b>	2.51	400		49	345778
<b>14</b>	2.51	600		79	513974
<b>20</b>	2.51	900		74	259419
<b>25</b>	2.46	900		51	123673
<b>30</b>	2.51	1000		49	1050811

The mesh generation properties were kept as consistent as possible, though the relevance was varied to ensure that the number of elements did not become excessive and prevent the simulation being solved. The parameters used are shown in Table 6.8, with the other mesh settings being kept at their default values.

Table 6.8: Summarised meshing parameters for the diameter simulations

<b>Mesh Setting</b>	<b>Value</b>
<b>Size Function</b>	Adaptive
<b>Relevance Centre</b>	Fine
<b>Smoothing</b>	Medium
<b>Transition</b>	Fast
<b>Span Angle Centre</b>	Coarse

### 6.3 Results




Figure 6.5: Comparison of the effect changing the volume fraction has on the four composite materials being studied here (note the results for silica and RF lower bound are overlapping)

The summarised results for the volume fraction simulations are shown in Figure 6.5. The difference between the four simulated materials is behaving much as would be expected: the lower bound for the RF aerogel thermal conductivity and the Silica aerogel thermal conductivity are both very similar, and the polyimide conductivity is the highest of the four. The trends within the data make sense too: at higher volume fractions, there is more high conductivity fibres for the heat to flow through, and also more fibre/fibre contact so that the heat transfer path can be more efficient.

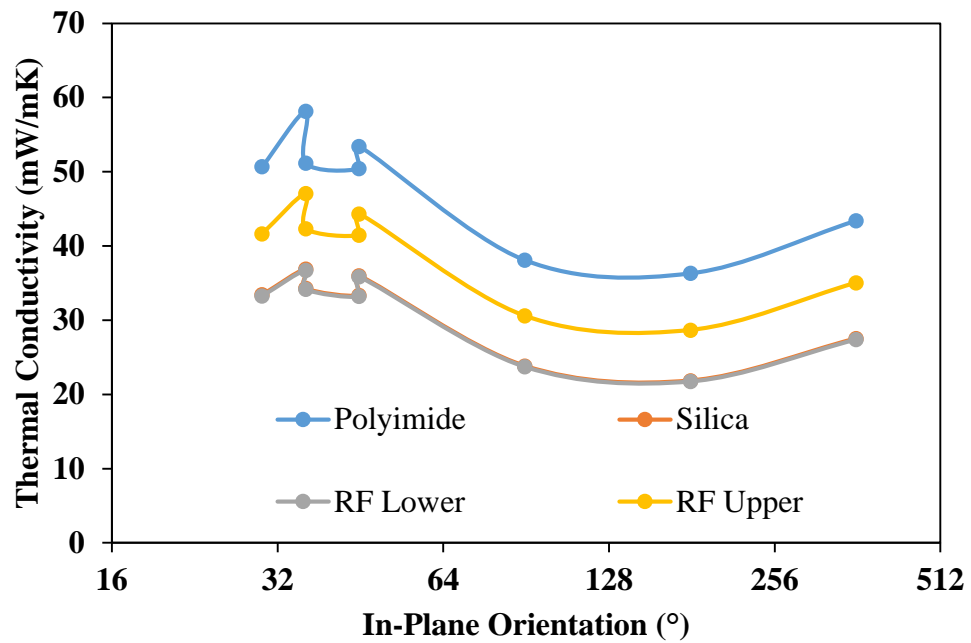


Figure 6.6: The effect of varying the maximum in-plane angle on the thermal conductivity (note the results for silica and RF lower bound are overlapping)

The investigation into the effect of the orientation on the thermal conductivity in general produces similar conclusions as with pure fibre blankets in Section 3. The thermal conductivity in the in-plane results, shown in Figure 6.6, is highest at low angles due to the high level of ordering present in the geometry. Above approximately 90°, the conductivity remains more constant as the increase in max allowable angle is covered by symmetry and does not strongly add to the randomness of the fibres.

The general trend of the results is as expected again, with the materials with higher conductivities producing higher overall thermal conductivities, and the conductivity dropping as the increased randomness causes the optimal heat transfer path through the material to become longer.




Figure 6.7: The effect of varying the maximum out-of-plane angle on the thermal conductivity





Figure 6.8: Comparison of the effect the number of elements (top) and the volume fraction (bottom) has on the thermal conductivity in the out-of-plane angle simulations



While the expected behaviour for the out-of-plane simulations is that the thermal conductivity will increase as the maximum angle increases, since the fibres have a higher conductivity than the matrix and them being increasingly in-plane should allow for an easier heat transfer through them, the results shown in Figure 6.7 indicate that, in general, the opposite trend occurs. This is not easily explainable as an error caused by, for example, the volume fraction changing in the simulation, as Figure 6.8 shows that neither the volume fraction nor the number of elements changing are producing a significant trend which could reasonably be biasing the results. The results behaving unexpectedly is unlikely attributed to a single parameter but is perhaps indicative of deeper issues within the simulation process.




Figure 6.9: Comparison of the effect of varying the diameter on the four composites investigated. (note the results for silica and RF lower bound are overlapping)

The effect the fibre diameter has on the simulations appears to be largely negligible here, based on the results in Figure 6.9. While a slight downwards trend does exist,

there is not enough data to call it significant. The individual materials behave much as expected: with the highest conductivity aerogel also producing the highest overall thermal conductivity.




Figure 6.10: Effect of varying the relative fibre length on the conductivity

The effect of varying the size of the generation region is shown in Figure 6.10; where it can be seen that it causes the overall conductivity to increase. Since this corresponds to an increase in the fibre length, it is plausible that this occurs because it leads to an increase in fibre/fibre contact and hence more efficient heat transfer paths through the geometry than in a smaller region.

#### 6.4 Validation

In order to validate the model, a geometry representative of an actual PET fibre/RF aerogel composite was created.

The same heat flow meter used in the previous section to obtain an overall conductivity was also used to produce a set of data relating the temperature to the thermal conductivity of the sample. This was done as the temperature is capable of

being adjusted without having to produce specifically designed fibre blankets; whereas any other property of the composite blanket would require a significant amount of time to produce due to the length of the sol-gel process.

The standard testing conditions for this were a compressive load of 13.8kPa, and a temperature difference of 20°C between the hot and cold plates of the heat flow meter.

The test was done at the mean temperatures indicated in Table 6.9.

Table 6.9: Experimental results for the PET fibre/RF aerogel composite blanket

<b>Mean Temperature (°C)</b>	<b>Thermal Conductivity (mW/mK)</b>
<b>-20</b>	17.54
<b>-10</b>	18.24
<b>0</b>	19.25
<b>10</b>	20.06
<b>20</b>	21.05
<b>30</b>	21.93
<b>40</b>	22.83
<b>50</b>	23.79
<b>60</b>	24.80
<b>70</b>	25.99




Figure 6.11: FEA geometry used for the simulation; the box around the fibres was treated as an aerogel/air mix of appropriate conductivity.

The geometry used to model the aerogel composites is shown in Figure 6.11, and consists of a 200 $\mu$ m sided matrix containing 23 curled fibres. The matrix was given properties consistent with either the parallel or series model derived thermal conductivity, shown in Figure 6.1. The boundary between the fibres and the simulated aerogel was treated as having infinite thermal conductivity, that is to say, the interfacial layer poses no resistance to heat transfer through it.

As the simulations were done in ANSYS Mechanical, the exact element types were not specified for the meshing process. The mesh settings used are listed in Table 6.10, with all other mesh settings being kept as the default settings. This produced a mesh containing 745k elements, with 1.15 million nodes, with approximately 74% of them representing the aerogel.

Table 6.10: Mesh settings used in the validation

Parameter	Value
Relevance	-40
Size Function	Proximity and Curvature
Smoothing	Low
Transition	Fast
Span Angle Center	Fine

In order to ensure the heat flow through the much thinner FEA sample was the same as the heat flux through the larger blanket, Fourier's law of conduction was applied, and the two heat fluxes equated to each other, which is shown in Equation (73). Since the conductivity of the experimental blanket can be assumed to be equivalent to the conductivity of the simulation, that term can be neglected; and since the size of the simulation sample is already known, the temperature difference across it to produce the same mean temperature as the experimental data can be found.

$$k_e \frac{\Delta T_e}{L_e} = k_s \frac{\Delta T_s}{L_s} \quad (73)$$

This calculation indicates that a temperature difference of 0.658°C across the sample would be representative. This means that to replicate a mean temperature of 20°C in the experimental data, the upper and lower temperature constraints in the model would be 20.329 °C and 19.671°C, respectively. The sides of the box were treated as being perfectly insulated, which is generally unrepresentative but should not introduce significant errors as the principal heat transfer direction is vertically.




Figure 6.12: Comparison of the Experimental results with those obtained via the FEA using thermal conductivity values derived from the series and parallel models

The results of the simulation are shown graphically in Figure 6.12, and in tabulated form in Table 6.11. While both models are overestimating the conductivity, the series model conductivity is significantly more accurate and reaches values within 10% of the experimental values. The parallel model conductivity is a lot less accurate, only being within approximately 40% of the experimental conductivity on average. This is somewhat to be expected: the derivation of it in Figure 6.1 implies that, as it is much closer to the blanket conductivity, it is overestimating the conductivity of the aerogel component alone. This is born out in part by the series model conductivity following the trend of the overall conductivity better both as an individual component and also when used in the full model of the blanket conductivity.

Table 6.11: Thermal validation results

Mean Temperature (°C)	Thermal Conductivity (mW/mK)		
	Experimental	Series	Parallel
<b>-20</b>	17.54	19.48	26.30
<b>-10</b>	18.24	20.10	27.63
<b>0</b>	19.25	21.11	29.30
<b>10</b>	20.06	21.87	30.76
<b>20</b>	21.05	22.81	32.41
<b>30</b>	21.93	23.69	33.94
<b>40</b>	22.83	24.55	35.50
<b>50</b>	23.79	25.49	37.12
<b>60</b>	24.8	26.50	38.79
<b>70</b>	25.99	27.75	40.66

## 6.5 Conclusions

The results of the simulations behave, generally, much as expected. The validation results also suggest a high degree of confidence in the results: in particular, the trends present here are likely to be accurate even if the actual overall conductivity values are not. The validation results however, imply that the results should, in general, be within around 10% of the actual conductivity, though the nature of stochastically generated geometries means that there will be some variation between the results of each geometry even with the same variables behind the generation.

The lower bound for the RF aerogel and the Silica aerogel consistently produced very similar results throughout this: this is largely due to the thermal conductivity of both the materials being very similar but it is a good sign that the model is not doing

anything hugely unexpected. The polyimide aerogel and the upper limit for the RF aerogel both behaved as expected in general too; with the polyimide aerogel producing the highest overall thermal conductivity.

In general, the results here are mirroring the same behaviour seen in Section 3; though this very simple model is neglecting some important parameter such as how well the fibre and the aerogel will bond together to produce a composite. Unfortunately, this is extremely complicated to predict at the microscale level and would likely require carrying out molecular dynamics simulations to fully consider. This means it is hard to tell exactly how much of an effect the fabric blanket layup is having on the aerogel formation without carrying out extensive experimental testing or simulations to identify it.

The validation results are behaving as expected here due to the way the two speculative values for the RF conductivity were derived. Since the Parallel model tends to overestimate the total conductivity, it will in turn tend to produce a lower component conductivity than the Series model. This is required because the same conductivities in both models would produce a higher conductivity in the Parallel model, so to get the same conductivity out of the Parallel model, the component conductivities have to be lower; and since the PET conductivity is fixed here, it means the RF aerogel conductivity in the Parallel model is the one that has to be lower.

The experimental conductivity of the RF aerogel not being between the two models is somewhat surprising even if the series model appears to be accurately representing the trend in the conductivity change. It is possible that a miscalibration in the temperature data for the HFM is causing issues here with the mean temperature being slightly off. This would not have a huge effect on the experimental data but would affect the simulation data more due to the comparatively smaller temperature range that it takes place over.



## 6.6 Future Work

This is a surface level look into the effect of combining aerogels and fibre blankets as a composite. The next chapter delves more rigorously into modelling aerogels than this one does; but there is room to combine the upcoming work with the fibre model to produce a more accurate simulation of an aerogel/fibre composite. This would allow for a better understanding of the importance of the way the structure of the aerogel affects the properties of the composite, and also allow investigation into how well the fibre/aerogel bond together, and what effect this has on the thermal conductivity of the whole.

Only the thermal properties, and only those featuring one kind of fibre, were investigated here: a more thorough examination of the trade-off between the mechanical and thermal properties of the blanket, including investigation of the flexibility of the blanket and the effect that flexing it has on the thermal conductivity, could be undertaken. This would likely be most effectively done via experimental testing, but a dynamic physical model could allow it to be done via FEA or MD simulations.

## **7 Modelling of Aerogel Particle Filled Composites**

### **7.1 Introduction**

The influence of particles dispersed through a solid matrix was investigated. The purpose was to investigate whether they would be able to lower the conductivity of the matrix significantly and also to investigate whether this would in turn cause the mechanical properties of the matrix to be lowered significantly.

Specifically, the particles under consideration here were manufactured from a larger block and sieved into a specific size, which means that while the shape may be roughly spherical, there is a distribution of particle sizes present. As a result, accurately modelling their size requires a model that can take this diameter distribution into account.

### **7.2 Literature Review**

Research into embedding particles into a matrix has been studied extensively in some regards. The experimental measurement of the properties of these composites has been carried out for a wide range of materials: from cement [58] to high density polyethylene [59, 60] to silica aerogel based composites [61]. A focus on increasing the thermal conductivity is often seen (so that electrically insulating heat sinks can be produced) [62].

R.C. Progelhof, J.L. Throne and R.R. Ruetsch together carried out a relatively comprehensive review of theoretical models and empirical correlations that can be used to determine the conductivity of composites [63]. Their work shows that in general, simple theoretical models can produce accurate results in some conditions but that most are unsuited to covering every possible case. However, their work only considers composites where the matrix conductivity is lower than the particle

conductivity; namely ones where increasing the thermal conductivity of the composite is the desired behaviour.

Their work partially relies on the experimental data obtained by D.W. Sundstrom and Y. Lee [65], which they believe to be accurate to about 3%. This was obtained using a thermoconductometer, which measures the conductivity of the sample based on how quickly heat is transferred from a hot liquid at the bottom through a silver plate, the sample, an upper plate and a second, different liquid. Since the temperature of the liquids can be determined based solely on when they carry out phase changes, the temperature difference is known. The conductivity can then be correlated to how long it takes the upper liquid to produce a set amount of vapour. The main concern produced by this is that the indirect nature of the conductivity measurement (since it is determined by a correlation with known data rather than being able to be determined directly) suggests a possible level of inaccuracy.

The work of Han, Wood et al [60] is illustrative as to the properties that affect the thermal behaviour of a particle based composite. They experimentally measured the thermal conductivity of a small range of materials, and varied properties of the particles. One of their significant findings was that the size of the particles did not appear to be impactful on the conductivity significantly, with the volume fraction being a much more significant impactor. Further their findings suggest that a large interfacial mismatch in the mechanical modulus of the particles (which they correlate with the speed of sound through the material, which in turn is related to the ease with which “phonons” of thermal energy can transfer through it [66]) caused large thermal contact resistances. This is suggested because more mechanically dissimilar materials had a lower increase in thermal conductivity despite the particles having comparable or higher thermal conductivities. While this is definitely indicative of this being the case, it seems like Han et al lacks sufficient evidence to definitively prove it.

However, they were investigating the addition of particles to increase the thermal conductivity rather than decrease it, which means some uncertainty exists as to how the trends affect materials when the matrix is more thermally conductive than the particles are.

The results showing that size of the particles not affecting the conductivity contradicts the work of D.P.H. Hasselman and K.Y. Donaldson [67], who found a significant effect on the conductivity of an aluminium matrix reinforced with silicon-carbide particles. Their results showed that smaller particles caused the thermal conductivity to decrease, with the aim of increasing the conductivity of the aluminium matrix. Figure 7.1 is of particular interest as it shows that adjusting the particle size from 28 $\mu\text{m}$  to 0.7 $\mu\text{m}$  caused the conductivity to drop almost in half at low temperatures. It appears that the largest particle size has a more significant temperature correlation than the smallest particle size, with the smallest behaving almost the same as the matrix with respect to the conductivity variation.




Figure 7.1: Temperature dependence of the thermal conductivity of unreinforced aluminium matrix and 40 vol% particulate-SiC-reinforced aluminium matrix composites and unreinforced matrix for a range of values of SiC particle size, from [67].

It is worth noting, however, that Han et al were working with composites at a volume fraction of 10% to 30% when comparing the effect of changing the particle size, whereas Hasselman and Donaldson were working with composites at a volume fraction of 40%. This higher volume fraction has the potential to lead into significant amounts of percolation: where there are enough particles that conduction between the particles begins to have an important effect. At low volume fractions, it can be assumed the particles are not contacting, or that the contact is relatively unimportant to the overall thermal conductivity. However, once the percolation threshold has been reached, it becomes more accurate to consider the particles as a single component rather than a series of discrete particles.

Ghosh, Singh and Maity [68] looked into the thermal performance of fibre blankets made of jute fibres experimentally using a variation on a heat flow meter design. They found that for crimped fibres, increasing the level of “crimp percentage”, defined as shown in Equation (74), led to the thermal conductivity dropping.

$$\text{Crimp\%} = \frac{\text{Extended length} - \text{Initial Length}}{\text{Initial Length}} \times 100 \quad (74)$$

However, the method of crimping used would produce random, uncontrolled crimping as opposed to regular crimping. While this means it is somewhat indicative that this can help improve the thermal performance of a periodically crimped fibre, it is not definitive. They also investigated the effect of porosity on the conductivity and found

that increasing porosity generally led to an increase in conductivity. The range of porosities investigated was relatively small: from 95% to 98%, but there was a significant increase in conductivity from 0.378 to 0.414 W/mK. The minimum conductivity was also not at the lowest porosity: 95.43% has a conductivity of 0.398W/mK, with 96.35% porosity being the 0.378W/mK value. This trend is interesting: theoretically, increasing the porosity here should decrease the fibre content and hence decrease the thermal conductivity. It is possible that the high levels of porosity present here mean that convection can more easily occur within the pores.

The Nielsen model [63, 64] is shown in Equation (75):

$$\frac{k_c}{k_m} = \frac{1 + ABv_p}{1 - B\psi v_p} \quad (75)$$

Where A, B and  $\psi$  are defined as:

$$A = K_{E-1} \quad (76)$$

$$B = \frac{\frac{k_p}{k_m} - 1}{\frac{k_p}{k_m} + A} \quad (77)$$

$$\psi = 1 + v_p \left( \frac{1 - \phi_m}{\phi_m^2} \right) \quad (78)$$

$\phi_m$  is the maximum packing efficiency the particles can have; and is used here as a measure of how efficiently the particles are packed in here. The generalised Einstein coefficient,  $K_E$ , is determined by the type of particle system that is being modelled. For randomly dispersed particles,  $\phi_m = 0.637$  and  $A = 1.5$  [64].

The Geometric mean model can be used to determine the overall thermal conductivity using Equation (79):

$$k_c = 2^{v_p \log_2(k_p) + (1-v_p) \log_2(k_m)} \quad (79)$$

### 7.3 Geometry Generation

The generation for this follows a very similar pattern to the fibre generation method. In order to make generating the geometry easier and the simulations to be carried out faster, the particles were assumed to be spherical in shape and the validation of this assumption depends on how aerogel powders are synthesised. The first step in generating a dispersed particle geometry was to define a bounding region into which the particles would be placed.

Next, a particle was generated with a random initial position and a defined radius. The distance from the particle to the edge of each side of the bounding region was tested to ensure the particle was fully within the bounding region. If it was, then it was tested against all the previously positioned particles to confirm that it was not overlapping with any of them, which was done iteratively until either all the previous particles had been tested or it was found to be overlapping with one particle.

Finally, the particle was saved, and the loop started over until the volume fraction of particles had exceeded the desired volume fraction. At this point, a scatter plot of where the particles are can be generated using MATLAB, as shown in Figure 7.2. It should be noted that the diameter of the circle is not representative of the diameter of the particle. The MATLAB codes used to generate a particle-based geometry are included in Appendix G.




Figure 7.2: MATLAB scatter graph showing the particle positions after generation in 3D, with the positions being defined in micrometres

Because the particle geometry is less complex than the fibre geometry, it is exported as an ANSYS Design Modeller script which works by creating arcs of the same radius as the circles and then revolving them to produce a full sphere. Each revolve is a “slice material” operation, which means that a hole in the resin is produced where the spheres are to ensure that the two materials are not overlapping. An example of the geometry produced this way is shown in Figure 7.3, where the particles have been generated with a normal distribution for the diameter.






Figure 7.3: 3D geometry produced by the algorithm in Design Modeller

This method is relatively simple to carry out from a computational perspective; since the longest calculation is an application of Pythagoras's Theorem in 3D in order to work out the distance between the centres of two spheres; which can then be compared to the two radii to determine whether they overlap. Spheres overlapping were found to be a significant problem as ANSYS had difficulties meshing the interface between the two spheres and the matrix, since the geometry here creates a very thin matrix section at the point of intersection.





Figure 7.4: Both graphs show a histogram of the particle size so that the distribution of sizes can be seen; the top is at a volume fraction of 10% and the bottom is at a volume fraction of 40%.

While in general, the diameter of every particle was fixed to one single value, an investigation of a normal distribution centred on 50µm diameter particles was done to investigate the effect of the volume fraction. This led to the encountering of a distribution issue at higher volume fractions. Specifically, the larger particle sizes become progressively harder to insert as the box becomes more full of spheres, which is exacerbated by preventing them from touching to prevent singularities occurring at the very thin points in the matrix between at the point where the spheres stop contacting. This means that the smaller particles become easier to insert in the geometry and as a result, are more likely to be successfully inserted into it, which in turn leads to the geometry having a higher proportion of smaller particles than desired, as can be seen in Figure 7.4.




Figure 7.5: Example temperature plot for the particles, with the air around the particles being hidden, the top face being constrained to 30°C and the bottom face being constrained to 10°C.

Unfortunately solving this is not straightforward using the current method of particle generation, as the highly random nature of particle insertion means that it is very hard to ensure that space is kept for larger particles to be inserted without in turn biasing the distribution towards more regular patterns than desired. It is possible to generate a histogram of the particle size distribution which allows this issue to be detected once the initial MATLAB script has finished running. The actual volume fraction at which the distribution begins to change from the desired distribution is dependent on the size of the spheres compared to the size of the bounding region, with relatively small spheres having a higher volume fraction obtainable before this becomes an issue. Similarly, this can be an issue when a fixed diameter is specified since it is possible for the geometry to become full so that no more spheres of that size can be inserted, which in turns means that the code will never complete. This can be limited by counting how many iterations have passed since the last successful particle was positioned and exiting once it reaches a certain amount, such as 1000.

## **7.4 Simulation Setup**

### **7.4.1 Thermal Simulations**

The same cuboidal geometry was used here as in Chapter 3, to ensure a large surface area for the heat to transfer through. The radius of the particles was kept the same within each geometry, and the same within each set of simulations (except the simulations looking at the effect the radius has on the thermal conductivity of the samples).

The thermal simulations for the particle-based composites looked at four matrix materials: Epoxacast, Polyester 41 (Poly 41), Isophthalic Polyester (Iso) and Orthophthalic Polyester 40 (Ortho 40). These materials were chosen by Blueshift

International Materials, with them all having both good mechanical properties and relatively low thermal conductivities.

The particles used in this simulation were made out of AeroZero aerogel and were intended to replicate the effect of grinding and sieving a block of AeroZero and distributing the resulting particles in a composite material. Various sieve levels can be used to produce a range of different particle size distributions, which can be replicated using complex particle size distributions. These distributions were not used here so that the effects of varying the fibre parameters could be more easily identified, since the variation would cause a larger degree of 'noise' in the results obtained.

#### **7.4.2 Mechanical Simulations**

The samples used for the mechanical simulations, as in Chapter 4, were shaped to replicate standard tensile test specimens for composite materials, as shown in Figure 7.6. The properties used to generate these samples are given in Table 7.1. The volume fractions are lower than anticipated here because some of the particles generated by MATLAB ended up failing to generate in ANSYS. For the 1% sample, this was only two particles, but as the number of particles in the simulation increased, the number that failed to generate reached 69 in the 20% sample. This increase is likely down to the much larger number of particles present in the simulation, since it means problematic geometry features such as particles overlapping, and particle/bounding region contact are more likely to occur.

The particle radius was kept the same in the tests here to minimise the effect that the packing has and eliminate the effects of the size distribution being slightly different in each sample.




Figure 7.6: The geometry for the 20% VF sample. Note the two lines on the bottom left face that allow the centre to be selected as a point.

Table 7.1: The properties of the geometries used in the mechanical simulations

<b>Volume Fraction (%)</b>	<b>Particle Diameter (mm)</b>	<b>X-Length (mm)</b>	<b>Y-Length (mm)</b>	<b>Z-Length (mm)</b>	<b># of Particles</b>
<b>0</b>	0.05	0.50	1.00	5.00	0
<b>0.96</b>	0.05	0.50	1.00	5.00	46
<b>4.94</b>	0.05	0.50	1.00	5.00	236
<b>9.72</b>	0.05	0.50	1.00	5.00	464
<b>18.56</b>	0.05	0.50	1.00	5.00	886

The mesh settings used for this simulation are shown in Table 7.2. Due to the large variation in the number of particles present in the simulation, the ‘Relevance’ value (which controls how small the elements are) was varied across the simulations. The

other values were kept the same to ensure that the meshes for each simulation were as close to each other as possible for random geometries to be.

Table 7.2: Mesh settings for the mechanical particle simulations

<b>Property</b>	<b>Value</b>
<b>Size Function</b>	Adaptive
<b>Relevance Centre</b>	Medium
<b>Transition</b>	Slow
<b>Span Angle Center</b>	Medium
<b>Automatic Mesh Based Defeaturing</b>	On
<b>Defeature Size</b>	Default

The mesh produced for each simulation is quantified in Table 7.3. The 0% volume fraction sample contains far fewer elements than the other as it is topologically equivalent to a standard tensile test specimen and contains no spherical inclusions.

Table 7.3: Mesh information for the individual simulations

<b>Volume Fraction (%)</b>	<b>Relevance</b>	<b># of Elements</b>	<b># of Nodes</b>
<b>0</b>	0	1280	6761
<b>0.96</b>	100	699501	959019
<b>4.94</b>	60	654198	942598
<b>9.72</b>	0	948564	1390076
<b>18.56</b>	-100	5590608	8169857

## 7.5 Results




Figure 7.7: Effect on conductivity as the number of pores increases; while also investigating how having 4 larger pores and one small one and 4 small pores and 1 large pore change the conductivity.

The first thing investigated here was the effect that simply including a series of pores within a geometry had on the results. In this case, a very simple set of test cases were produced to allow the effect of just changing the size of the pores while keeping the pore volume constant was considered. This led to the results shown in Figure 7.7; where it can be seen that larger pores tend to produce higher overall thermal conductivity than smaller pores, and in turn implying that adding more particles can cause the thermal conductivity to drop.






Figure 7.8: Comparison of the effect having more pores has on the Thermal Conductivity

This is further confirmed by the data shown in Figure 7.8, which shows the trend is relatively linear as the number of pores included in the 2D simulation is increased. This leads to the hypothesis that, in general, including spherical inclusions could reduce the overall thermal conductivity in a composite.




Figure 7.9: Comparison of the simulated thermal conductivities for three particle sizes with the expected upper and lower bounds for the conductivity, which were found using the rule of mixtures models for parallel and series composites




Figure 7.10: Same comparison as in Figure 7.9 but with the Nielsen Model and the Geometric Mean Model added

The results of the simulation for a range of particle diameters are shown in Figure 7.9. The Upper and Lower bounds are based on the parallel and series models respectively (see Chapter 3). The actual particle diameters are having less of an impact here than the volume fraction is, with all the simulations producing pretty similar results at the same volume fraction. The actual results appear to be behaving as expected here; particularly when compared to the additional theoretical models present in Figure 7.10. There is a good general agreement between the simulations, the geometric mean model, and the Nielson model, with the geometric mean model appearing more accurate at lower volume fractions and the simulation results more accurate at higher volume fractions.

Despite this agreement, the trends in the simulated data do not follow the trends present in any of the theoretical models well: the curvature is most reminiscent of the lower bound, while the other three models are much more linear in nature. It is, however, tailing off quicker than the lower bound is, so while the curve is more similar, it is not the same.




Figure 7.11: This is comparing how the addition of various fractions of aerogel filler to three kinds of matrix reduces the conductivity, with the conductivities being normalised so that the important thing is the change from 0% filler and not the absolute change.

A further comparison of the effect the material properties have in the simulations is given in Figure 7.11. The data here has been normalised so that the relative conductivity indicates what fraction of the matrix conductivity the overall conductivity is: namely, a value of 1 implies the overall conductivity and the matrix conductivity are the same; and a value of 0.5 implies the overall conductivity is half of the matrix conductivity.

The trend of all three simulations fits most closely with the lowest geometric mean conductivity, and indeed, while the simulations do show some variation between the conductivities tested, the difference is a lot more minimal than it is for the geometric mean model.

In order to further validate the modelling results, they were compared to experimental results produced using the same experimental setup described in Chapter 3. Four materials were used initially: two 3 mm thick isophthalic polyester blankets, one filled with particles to a weight fraction of 10%; and a second experiment using two such blankets to give a total thickness of 6mm for the samples. It can be seen that there is a high degree of correlation between the modelling results and the experimental results, particularly for the thicker sample. This discrepancy can be partially explained by the relative thinness of the 3 mm sample: at this level, the surface resistance in the experimental results is expected to be playing a significant part in causing the thermal conductivity of the sample to appear lower than it is.




Figure 7.12: Comparison of the experimental and modelling results for the Isophthalic Polyester material




Figure 7.13: Comparison of the simulation and experimental results for Polyester 41




Figure 7.14: Comparison of the simulation and experimental results for Polyester 41

Simulation results for the Polyester 41 are shown in Figure 7.13 and Figure 7.14. As before, the model is producing results very close to the experimental results; the general trend is not quite exact here, but it is similar. In particular, the second chart shows more variation in the trends in the modelled data than the experimental one, but it does accurately represent the bilinear relationship between the temperature and the conductivity shown in the experimental data.

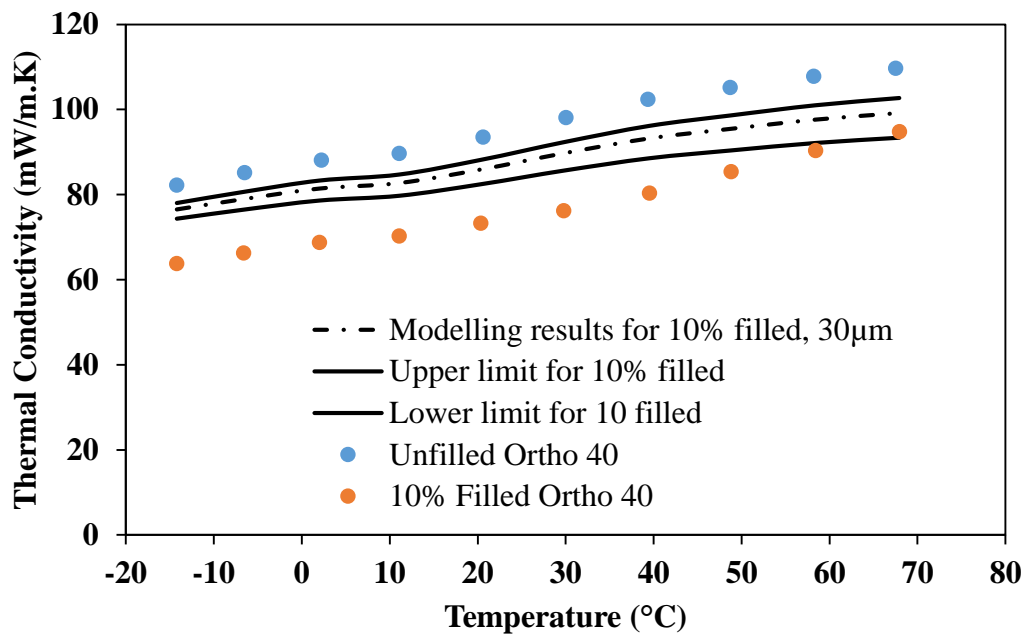


Figure 7.15: Comparison of the simulation and experimental results for Ortho 40

As with the isophthalic polyester, the orthophthalic polyester results are relatively close, when the experimental and modelling results are compared (as shown in Figure 7.15). The simulation is continuing to overestimate the thermal conductivity here; though at higher temperatures, this is less of an overestimate.

An investigation in the thermal performance of layered composites was also, briefly, looked at. Specifically, carbon fibre and aramid fibres were investigated at the behest of Blueshift International Materials, where the materials were constructed from a laminate containing either 0, 9 or 18 layers of AeroZero film within it. This film is a low conductivity, flexible film that is approximately 5mm thick, which is added to the laminate to lower its overall conductivity.




Figure 7.16: Comparison of simulation and experimental data for the carbon fibre composites






Figure 7.17: Comparison of simulation and experimental data for the aramid composites

The results of these simulations are shown in Figure 7.16 and Figure 7.17. The modelling results are pretty close to the experimental results in both cases, which is generally predicted for the much more developed field of laminate composite simulation. This is particularly true for the carbon fibre composite, with slight discrepancies existing for the aramid, particularly at the highest and lowest temperatures. This discrepancy is approximately 2% at most, and generally much smaller, indicating the mode is accurately predicting the thermal conductivity.




Figure 7.18: Modelling results indicating how the mechanical properties change as the volume fraction changes

Following the heat transfer modelling, a series of mechanical simulations were carried out so that the trade-off between the thermal properties and the mechanical properties could be identified. Figure 7.18 shows how the four main materials Young's Moduli are being affected by the change in volume fraction. The trend for all four materials is very similar and much as expected: introducing the less mechanically stiff aerogel particles into the system causes the overall Young's Modulus to drop, and more of them causes this effect to be larger.




Figure 7.19: Normalised Young's modulus against volume fraction

The normalised data in Figure 7.19 shows that the significant difference in the Young's Modulus behaviour between the samples is largely down to the matrix material having different Modulus values; since the normalised data shows that all the materials are following the same general trend. This is to be expected due to the relatively simplistic way the simulation is setup, since the contact between the particles and the matrix is being treated as perfect: in the actual materials, the behaviour at the interface would be more complex as, for example, partial debonding can occur, and the stress transfer at the interface would be less complete.

## 7.6 Conclusions

The comparison of the simulation results in conjunction with the experimental results implies a high degree of confidence in the thermal conductivity simulations. As the analytical models for the mechanical properties show a high degree of agreement with the simulation results, it is likely that the Young's Modulus predictions are accurate too.

From this we can see that the Polyester 41 sample is producing a good compromise: the Young's Modulus values for it in Figure 7.18 are the highest of the four materials investigated at all the volume fractions while the thermal conductivity results for it in Figure 7.13 are some of the lowest thermal conductivity results present here.

The trade-off for the volume fraction is more difficult to resolve it highly depends on just how much more important low thermal conductivity is than mechanical performance.

## **7.7 Further Research**

As with the fibre modelling, the interface between the particles and the matrix was considered to be perfectly conductive and bonded: namely, there could be no matrix/particle separation under loading, and the heat transfer across the boundary was assumed to be the same. Realistically, neither of these assumptions are true, and a more complex modelling approach could be taken to handle the interface more accurately. The difficult here is experimentally determining appropriate values to model the interface.

This research also considered the particles of aerogel to be perfect spheres. This is an approximation due to the complexity of the particle shape, and a better approximation could be done by allowing for example elliptical particles, or particles of a fully random 3d geometry. Characterising the particle shape can be done to allow a more representative particle distribution to be used.

# **8 Macro Modelling of PiP Centraliser/Spacer**

## **8.1 Introduction**

The investigation in the last chapter has shown that the thermal conductivity of a polymer resin can be reduced by incorporating aerogels. In particular, the effect of

aerogels on thermal and mechanical property of the polymer resin can be well described by FEA and some analytical models found in the literature. It is worth noting that all the models presented insofar do not take into account of any change of aerogel in itself when it is mixed with liquid resin and subsequently immobilized after the resin is solidified. This is a crucial assumption and its validity will very much depend on how aerogels/resin composites are processed. It is perfectly possible for the resin to infiltrate the pores in the aerogel during mixing if the resin has a low viscosity, aerogel surface energy is higher and the resin, and shear force is too high. If pores can be preserved beyond the mixing stage, there is also a possibility of pores collapse due to thermal residual stress related to resin curing cycle. Clearly, these are all relevant to the final structure-performance of aerogel/polymer composites and would require a detailed investigation. However, it would also need a comprehensive experimental programme to properly study these above factors largely associated with material manufacturing and processing. Hence, this is out of scope of this current project, which will assume an ideal scenario where no pores filled by the resin or damaged by thermal residual stresses during resin cure cycle.

It follows that the optimum material and conditions were found in Chapter 7 will be used to model the thermal and mechanical performance of a centraliser in this chapter. This allowed the predictions from that chapter to be validated, but it also means that their results can be validated through the use of experimental testing on a physical centraliser. This testing was primarily focussed on the mechanical performance of the centraliser, as the centraliser ended up failing under the expected mechanical load conditions. The current centraliser is made of casted nylon and a wide range of other alternative polymers filled with aerogel (e.g. AeroZero) as in the last chapter will be used to compare the nylon counterpart in this chapter.

The experimental testing done on the centraliser used three load cycles. The first two are to load it to the desired loads, 200 kN and 350 kN, keep it there for five minutes and then release the load. Between cycles, five minutes was granted to allow the centraliser to relax. The loading speed used was 1mm/min. The third cycle follows the same steps at 400 kN except after the load has been applied for 5 minutes, an axial load was applied and increased until the centraliser slides onto the outer pipe; at which point the load is released. The force is applied in two locations, shown in Figure 8.1, creating two load cases that are referred to as Test 1 and Test 2 throughout this document.




Figure 8.1: The load position in the two tests; with Test 1 on the left and Test 2 on the right. The experimental test rig uses a circular arc that covers approximately one quarter of the centraliser to do the test, as opposed to a point load as these images suggest.

Initial simulation work carried out only simulated Test 1; and the results of that simulation, based on the experimentally derived material properties for the DCPD, indicated that the centraliser would be capable of withstanding the required load. With this, the clearance was given to progress on to experimental testing using DCPD

centralisers manufactured by Subsea 7. The centraliser passed the Test 1 conditions at all three load levels. However, during the final loading cycle of Test 2, the centraliser fractured at the EHTF cut-out, as shown in Figure 8.2. A second centraliser was tested using the same process, which was stopped at the end of the second load cycle. At this point, a level of buckling and cracking had already started to happen, and the third load cycle was not done.



Figure 8.2: Image of the failed centraliser, showing that it has broken at the heating wire cut out

In order to determine whether the unexpected failure occurring during Test 2 was due to the new design of the centraliser, or the DCPD not being strong enough to withstand the necessary loading conditions, a series of simulations were conducted. The geometry used in the simulation was the same as the centraliser as shown Figure 8.1, with an internal supporting steel ring and an externally applied force in the location of the applied load in both tests.

## 8.2 Simulation Setup




Figure 8.3: Side on profile of the centraliser




Figure 8.4: Centraliser at an angle so the EHTF cut outs can be seen more clearly; note also the lines on the base that allow the centre of those faces to be selected.

Two sets of simulations were carried out here: the first was before the centraliser experiments were run and the second was a more realistic version designed to copy the experiments and investigate why it failed. The first set of simulations looked at a wide range of materials in order to investigate what the optimum composition was and were carried out at a lower overall force than the second set of simulations.

Materials in the first set of simulations can be divided into two categories: particle-based composites, and laminated composites. In the first case, AeroZero particles



were dispersed throughout a matrix randomly, as in the simulations in Chapter 7. In the second, laminated glass fibre, carbon fibre and aramid samples were modelled, with a variable number of thin sheets of AeroZero between layers. The materials used here were selected by Blueshift International Materials, and generally were expected to be good compromises between the thermal requirements for the centraliser and the mechanical performance required to ensure it can keep the inner and outer pipe separate through the life span of the pipeline.

The properties used for the first set of simulations are shown in Table 8.1. The samples with volume fractions listed contain AeroZero particles dispersed through them, with 0% and 10% being included for both so that the effect adding AeroZero had on the composites could be understood better.

Table 8.1: Mechanical material properties for the initial centraliser modelling

<b>Material</b>	<b>Volume Fraction of AZ (%)</b>	<b>Young's Modulus (GPa)</b>	<b>Poisson's Ratio</b>
Hydrex	0	3.25	0.35
Hydrex	10	3.39	0.35
Polyester Isophthalic	0	3.79	0.35
Polyester Isophthalic	10	3.7	0.35
Polyester 41	0	4.95	0.35
Polyester 41	10	5.51	0.35
Polyester 41	50	3.2	0.35
Epoxacast	0	3.94	0.35
Epoxacast	10	4.55	0.35
0 Layers Glass	-	3.62	0.35
9 Layers Glass	-	4.21	0.35
18 Layers Glass	-	3.61	0.35
Glass	-	26.56	0.35
Carbon	-	78.5	0.35
Aramid	-	21.41	0.35

These simulations were carried out using the mesh settings listed in Table 8.2; which produced a mesh containing 1,462,551 elements for the centraliser geometry.

Table 8.2: The mesh properties used for the first set of centraliser simulations

Property	Value
Relevance	100
Size Function	Proximity and Curvature
Relevance Center	Fine
Transition	Slow
Span Angle Center	Fine



Figure 8.5: Experimental setup for the centraliser testing; showing the circular loading bracket at the top and the inner pipe and partial outer pipe, both of which act as supports for the centraliser during testing.

The boundary conditions for the second simulation set were chosen to replicate the experimental setup, which is shown in Figure 8.5, as closely as possible. It can be

seen that the force is applied directly to the centraliser across an area that is approximately one quarter of the diameter of the centraliser: with the outer pipe supporting a large amount of the lower centraliser. In both tests simulated, the outer pipe was represented as a displacement boundary condition constraining part of the exterior of the centraliser to have no displacement in the y-direction. The inner pipe was simulated as being “flexible” (meaning it was allowed to deform under loading), while also being constrained to have no displacement in the normal direction (i.e., towards the centre line of the cylinder).

Table 8.3: Elastic material properties used in the simulation; with the bulk and shear moduli being calculated from them (and hence not shown)

<b>Material</b>	<b>Poisson’s Ratio</b>	<b>Young’s Modulus (GPa)</b>
<b>Nylon</b>	0.3	3.71
<b>DCPD</b>	0.3	2.55

The elastic properties of both materials are shown in Table 8.3; however since the centraliser has a large force applied to it, which causes a significant amount of plastic deformation to occur, isotropic hardening data was added as well to allow this behaviour to be more accurately modelled.




Figure 8.6: The bilinear isotropic hardening data used by the DCPD during the simulation

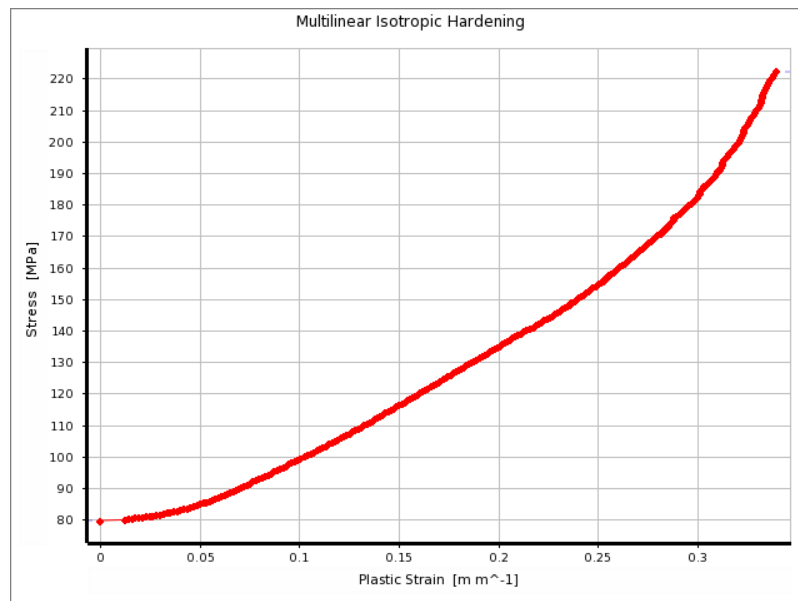


Figure 8.7: The multilinear isotropic hardening data used for the Nylon during testing

Several isotropic models were used here, with a final settling on using a bilinear model for both materials. Some investigation was done in using a multilinear model for the

nylon, with the model being shown in Figure 8.7; however, it was decided to not use this as the main model as the DCPD experience significant stress reduction after yield which ANSYS is unable to handle as a multilinear model. A bilinear model was instead used, where the initial linear stage represented by the Young's modulus was used to the yield point, and then a second modulus value was used to represent the change in material behaviour at this point. For DCPD, this is shown in Figure 8.6; and for Nylon, it is shown in Figure 8.7.

Subsea 7 used a model which, after the material reached yield, levelled off so that the stress was constant as the strain increased. This was tested with both materials under both test conditions.

The multilinear model is likely to be the most accurate model used here, with the caveat that it is also influenced by the effect of the testing machine compressing instead of the test samples. This effect is likely to be significant only at very high stresses where all the methods will have similar issues. However, it could not be used for the DCPD since the experimental results for that have a significant stress relaxation occurring immediately after yield. ANSYS Mechanical is incapable of modelling this kind of behaviour using multilinear isotropic hardening behaviours, since it is unable to determine what the appropriate strain to match the stress is when multiple strains are possible. This simplification also means that the model will produce significantly inaccurate results when the stress in the material exceeds approximately 130MPa, since the actual DCPD material begins to rapidly increase in stress here and the simulation is assuming a much slower increase in stress.

### 8.3 Test 1: Central Loading




Figure 8.8: Diagram showing where the 250kN force was applied in Test 1; with the inner pipe also being fixed to no displacement in the normal direction (i.e., towards the centre of the pipe)

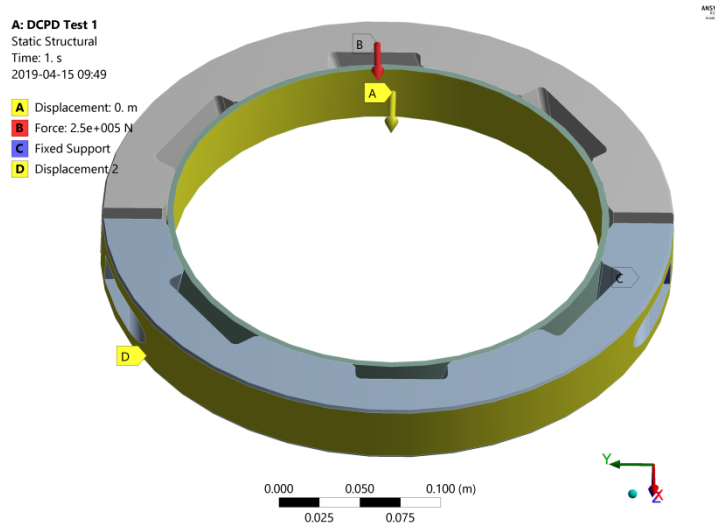


Figure 8.9: Diagram showing where the lower centraliser was constrained to not displace in the x-direction to replicate being supported by the lower pipe.

In this loading case, the outer face of the lower centraliser half was fixed so that it could not move in the x-direction, while a force of 250kN was applied to a large section

of the outer face of the upper centraliser, as shown in Figure 8.8 and Figure 8.9. The bolts were treated as “Flexible”, meaning that they could deform and take stresses; whereas the nut and washer attached to the bolt were “Rigid” because their behaviour should not be significant in this test. As this test is mostly looking at the effect on the centraliser, the bolt tightening was not considered, and neither were the threads on either the bolt or the nut.

#### 8.4 Test 2: Loading Across the Gap




Figure 8.10: Diagram showing where the force of 250kN was applied to the centraliser and that the inner pipe has been constrained to 0m of displacement in the normal direction






Figure 8.11: Diagram showing how the other side of the centraliser is constrained, with half being fixed to 0m of displacement in the y-direction, and the bolt flange being fixed fully to prevent rigid body motion

The inner pipe, bolts, nuts and washers were treated the same as in Test 1, with the exception of the bolt on the opposite side of where the load was applied, which was treated as “Rigid” in order to reduce the complexity of the simulation and help prevent rigid body motion occurring. In this case, the force of 250kN was applied in the y-direction, with the inner pipe being fixed to no displacement in the normal direction, and the bottom side of the centraliser being constrained to no displacement in the y-direction. The inner face of each side of the bolt cut out was also fully fixed to prevent rigid body movement occurring in either centraliser. These conditions are shown in Figure 8.10 and Figure 8.11.

## 8.5 Results and discussion

Table 8.4: Summarised simulation results; showing that the DCPD generally has bigger stresses and strains than the Nylon does; with the difference in strain in both cases being very significant

Material	Model	Test 1		Test 2	
		Max Stress (MPa)	Max Displacement (mm)	Max Stress (MPa)	Max Displacement (mm)
Nylon	Multilinear	166	2.70	222	8.85
	Bilinear	160	2.37	240	6.28
	Trilinear	87	3.62	N/A	N/A
DCPD	Bilinear	159	2.88	227	6.27
	Trilinear	99	3.11	N/A	N/A

The simulation results have been summarised in Table 8.4. The multilinear and bilinear models for the nylon are producing relatively similar results, which suggest that the bilinear model is a reasonably accurate way to represent the material properties here. The difference between Test 1 and Test 2 is huge here: it is very apparent that Test 2 is a significantly harder to past test; and both materials are experiencing a large amount of stress that is very likely to cause them to fail in this load condition.




Figure 8.12: Comparison of the maximum stress in the centraliser (probed value); and the stress distribution at the cut out where the centraliser is thinnest; with the nylon centraliser on the left, and the DCPD one on the right. The slice used for this comparison is not at exactly the same place, which contributes to some of the differences in the geometry and stress in the two images.

In Figure 8.12, the difference in the stresses at the region of the centraliser where the cut out for the heating wires and the cut out for the bolt to be attached are shown. Both centralisers are experiencing a similar stress distribution, with low stresses on either side of the bolt cut out, and a high stress band in seen on the far side of the centraliser at the cut out there.

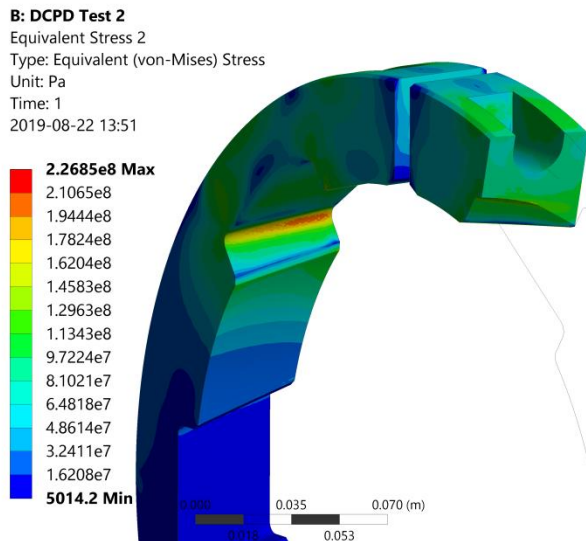


Figure 8.13: The DCPD centraliser in Test 2 load conditions, angled to allow a better view of the overall geometry than in Figure 8.12

## 8.6 Conclusions

The results indicate that the DCPD is likely to be stronger in Test 2 load conditions than the nylon is. The bilinear model used in both simulations treats the plastic deformation as being linear which is only approximately true, and only approximately true for relatively small stresses (up to around 130 MPa). Due to the stress relaxation in the DCPD under experimental conditions, this second linear section ends up indicating a higher amount of strain occurs for the same stress. This is reflected in the very similar amounts of displacement experienced in both materials while the DCPD also experiences significantly less stress.

It is clear that Test 2 goes significantly beyond the point where the experimental data remain reliable when the equipment fixture itself has a dominate effect on the measured stiffness. Using material data in such a large strain range will artificially cause higher stress level and in turn pre-mature failure. Thus, the material data after yield is likely to lead to underestimated structure integrity. Nevertheless, this is in turn

leads to the conclusion that the simulation is showing both materials would fail under the provided load conditions as observed in the actual tests.

The trilinear stress-strain behaviour adopted in the simulations further this: the successful solving of Test 1 is indicative that both materials would pass this test. The inability of Test 2 to solve is indicative that neither the Nylon nor the DCPD would pass this test since it is clear the stress is high enough to take it into the constant stress zone, which in this model is generally considered to be a failure.

The simulation data is showing a larger amount of maximum deformation than the experimental results have. However, it is unclear whether this is because of the simulation exaggerating the deformation or because the deformation of the actual centraliser is only measured at a single location. This means it is not necessarily an accurate measure of the full deformation experienced in the centraliser. It is also plausible that the simulated inner pipe is deforming more than the inner pipe in the experimental setup, and hence allowing the centraliser to deform more in the simulations. This does imply that the stress and strain values reported here may be inaccurate, but since both materials are experiencing the same load conditions in the simulation, the results can be interpreted comparatively at minimum since the effect of the inner pipe is standardised between them.

## 9 Conclusions

Chapter 1 has shown that there is an emerging demand for high-performance insulation materials that can be adopted in a wide range of applications. The current high-performance insulation strategies are constrained by a few options commercially available including vacuum insulated panel and rigid porous plastic. Aerogels present an exciting avenue in the field of passive heat insulation but their implementation has not been extensively explored in different applications. The deep-sea operations (e.g. oil & gas extraction) have provided a timely opportunity to study key material design parameters underpinning some novel concept of aerogel composites including fibre-reinforced aerogel and aerogel particle filled polymer.

Chapter 2 has summarised a number of common technical aspects involved in the development of aerogel composites for thermal applications. It spans from basic thermodynamic framework to design principles in anisotropic materials. It is can be seen that the literature in this area is quite scarce because in part aerogel composite is still a largely under development and material internal structure is complicated by dissimilar materials constitution with a significant amount of nano-to-micro scale pores randomly dispersed throughout the solid phase.

Chapter 3 has shown that a robust simulation method has been produced, through which the behaviour of two fibre types can be independently modelled and produce results that, per the validation carried out, appear to be accurate within 10% in general, though it is unclear how accurate this is when extrapolated over the full range of parameters that can be varied.

The optimal thermal performance here was achieved when the fibres where at a low volume fraction, with a large degree of in-plane randomness but no out-of-plane randomness. The diameter and fibre length both contributed relatively little to the

thermal properties, though it is possible that the simplicity of the simulation here is eliding their effect.

Chapter 4 indicates that increasing the volume fraction of the composites improves their mechanical performance: meaning that the trade off between the mechanical and thermal performance is very dependent on the volume fraction of the composites. To a lesser extent, the orientation negatively effects the mechanical performance, with more fibres being aligned in the direction the force is applied in being beneficial mechanically, but detrimental thermally.

Chapter 5 investigated the effect mixing the straight and curled fibres together had on the overall properties of the composite. It was found that the curled fibres performed worse thermally than the straight fibres did, though it is believed that the mechanical performance can be increased when sinusoidal fibres are introduced. This is caused by the greater degree of mechanical interlinking that adds further rigidity to the fibre composite.

Chapter 6 investigated how basic fibre/aerogel composites behave, indicating that in general, the same design considerations identified in Chapters 3 and 4 apply to aerogel/fibre composites too.

Chapter 7 demonstrated that aerogel particles could be reasonable approximated using spherical inclusions. In addition, the effect of increasing the volume fraction and adjusting the diameter distribution of the inclusions was simulated, showing that the diameter distribution does not have a huge effect on the results but that the volume fraction does. In particular, a higher volume fraction leads to better thermal performance but degraded mechanical performance, which is much as expected.

Chapter 8 investigated bulk modelling of the centraliser design, and showed that the failure of the centraliser was likely a combination of the new material being weaker

than the old material and the new centraliser design featuring more significant stress concentrations than the old design, leading to the failure experienced in the experimental testing. It is believed, though not proven, that the reduction in the material properties was merely a contributing factor and not the root cause of the failure.



## 10 Future Work

The results obtained from this project have constructed an initial framework with the ability to manipulate individual fibres spatially. This has enabled the possibility to investigate several key fibre parameters to form a variety of fabric architectures ranging from regular fibre packing typically found in high-performance fibre reinforced composites to more chaotic fibre arrangement in nonwoven fibrous insulation products. Consequently, this work opens up a great many opportunities for analytical and design operations in both cases.

Future work can build on this to further improve the modelling approach and expend the capability of this computational tool for wider exploitation. The following areas desire more studies.

- 1) The algorithm and codes developed in this work have some limitations. For instance, the method of generating fibres in a defined volume has an intrinsic issue of tightly controlling fibre volume fraction. Although this does not necessarily present a problem of varying fibre content as seen in this project, it can make parametric study more consistent by strictly controlling the fibre content as the most influential material parameter for final property.
- 2) There are some simplifications made throughout this work. For example, curled fibres are set to follow a mathematically scripted pattern. In reality, this is most likely not to be the case based on the fact of how these fibres are manufactured. Future work on this can take on a more statistical approach and bring in some random element to fibre shapes.
- 3) With respect to the aerogel particle filled polymer, the aerogel modelling relied on either bulk aerogel or spherical particle approximations, whereas realistically aerogel composites would not strongly resemble either of these

cases. As a result, a more in-depth treatment of aerogel particle filled polymer could be carried out by implementing a distribution of particle size and shape.

- 4) In addition, a lot of the mechanical investigation was focussed on the Young's modulus of the samples; but a deeper investigation into failure could be carried out. This would involve determining what appropriate failure conditions are in the samples. However, this is difficult to achieve in all cases, for example: in the particle filled composites case, a single particle getting crushed does not necessarily mean that the entire composite fails, particularly if the aim is good thermal performance, in which case a more complex simulation that accurately models the behaviour of the aerogels would be necessary so that a coupled thermal/mechanical simulation could be carried out and the point where the mechanical degradation of the material causes a significant increase in the thermal properties can be found.
- 5) There are, indeed, some cases here where the Young's Modulus was not investigated: in particular, the mixed fibre composites. These simulations could be carried out to fill in the data and determine better what the trade-off is in that case.

## 11 References

- [1] Arambakam, R., Vahedi Tafreshi, H., and Pourdeyhimi, B., 2013, "A simple simulation method for designing fibrous insulation materials," *Materials & Design*, 44, pp. 99–106.
- [2] Sarmiento, R. C., Ribbe, G. A. S., and Azevedo, L. F. A., 2004, "Wax Blockage Removal by Inductive Heating of Subsea Pipelines," *Heat Transfer Engineering*, 25(7), pp. 2-12.
- [3] Chawla, K. K., 2019, "Introduction," *Composite Materials: Science and Engineering*, Springer International Publishing, Cham, pp. 3-6.
- [4] Games, S., 1941, "Method of producing glass fibers," Google Patents.
- [5] Games, S., 1949, "Method and apparatus for making fine glass fibers," Google Patents.
- [6] Games, S., 1944, "Fibrous glass product," Google Patents.
- [7] Games, S., 1947, "Apparatus for manufacturing fibrous glass," Google Patents.
- [8] Campbell, F. C., 2010, *Structural composite materials*, ASM international.
- [9] Gibson, R. F., 2016, *Principles of composite material mechanics*, CRC press.
- [10] Yang, L., Sáez, E. R., Nagel, U., and Thomason, J. L., 2015, "Can thermally degraded glass fibre be regenerated for closed-loop recycling of thermosetting composites?," *Composites Part A: Applied Science and Manufacturing*, 72, pp. 167-174.
- [11] Chawla, K. K., 2012, "Metal Matrix Composites," *Composite Materials: Science and Engineering*, Springer New York, New York, NY, pp. 197-248.

- [12] Pierre, A. C., 2011, "History of Aerogels," Aerogels Handbook, M. A. Aegerter, N. Leventis, and M. M. Koebel, eds., Springer New York, New York, NY, pp. 3-18.
- [13] Koebel, M., 2015, "Aerogel-based thermal superinsulation: an overview," Journal of Sol-Gel Science and Technology, 63(3), pp. 315-339.
- [14] Hrubesh, L. W., 1998, "Aerogel applications," Journal of Non-Crystalline Solids, 225, pp. 335-342.
- [15] PetroWiki, "Oil Viscosity," Webpage, [https://petrowiki.org/Oil\\_viscosity](https://petrowiki.org/Oil_viscosity).
- [16] Woo, S. S., Shalev, I., and Barker, R. L., 1994, "Heat and Moisture Transfer Through Nonwoven Fabrics," Textile Research Journal, 64(3), pp. 149-162.
- [17] Fricke, J., Caps, R., Hummer, E., Doll, G., Arduini, M. C., and De Ponte, F., 2015, "Optically Thin Fibrous Insulations," ASTM.
- [18] Flourey, J., Carson, J., and Pham, Q. T., 2008, "Modelling Thermal Conductivity in Heterogeneous Media with the Finite Element Method," Food and Bioprocess Technology, 1(2), pp. 161-170.
- [19] Wang, J., Carson, J. K., North, M. F., and Cleland, D. J., 2006, "A new approach to modelling the effective thermal conductivity of heterogeneous materials," International Journal of Heat and Mass Transfer, 49(17-18), pp. 3075-3083.
- [20] Wang, M. a. H. J. a. Y. J. a. P. N., 2007, "Lattice Boltzmann modeling of the effective thermal conductivity for fibrous materials," International Journal of Thermal Sciences, 46(9), pp. 848--855.
- [21] Kanouté, P., Boso, D. P., Chaboche, J. L., and Schrefler, B. A., 2009, "Multiscale Methods for Composites: A Review," Archives of Computational Methods in Engineering, 16(1), pp. 31-75.

- [22] Kuksenko, D., Böhm, H. J., and Drach, B., 2018, "Effect of micromechanical parameters of composites with wavy fibers on their effective response under large deformations," *Advances in Engineering Software*, 121, pp. 206-222.
- [23] Karami, G., Grundman, N., Abolfathi, N., Naik, A., and Ziejewski, M., 2009, "A micromechanical hyperelastic modeling of brain white matter under large deformation," *J Mech Behav Biomed Mater*, 2(3), pp. 243-254.
- [24] Fisher, F. T., Bradshaw, R. D., and Brinson, L. C., 2003, "Fiber waviness in nanotube-reinforced polymer composites—I: Modulus predictions using effective nanotube properties," *Composites Science and Technology*, 63(11), pp. 1689-1703.
- [25] Yu, Y., Song, S., Bu, Z., Gu, X., Song, G., and Sun, L., 2013, "Influence of filler waviness and aspect ratio on the percolation threshold of carbon nanomaterials reinforced polymer nanocomposites," *Journal of Materials Science*, 48(17), pp. 5727-5732.
- [26] Daryabeigi, K., Knutson, J. R., and Cunnington, G. R., 2007, "Measurement of Heat Transfer in Unbonded Silica Fibrous Insulation and Comparison with Theory," 29th International Thermal Conductivity Conference (ITCC) Birmingham, AL; United States, p. 11.
- [27] Pathak, S. V., 2013, "Enhanced Heat Transfer in Composite Materials," Ohio University.
- [28] Clayton, W. A., Fabish, T. J., Lagedrost, J. F., and Boeing Co Seattle Wa Aerospace Systems, D. I. V., 1969, "THERMAL CONDUCTIVITY OF PHENOLIC-CARBON CHARS."

- [29] Pilling, M. W., Yates, B., Black, M. A., and Tattersall, P., 1979, "The thermal conductivity of carbon fibre-reinforced composites," *Journal of Materials Science*, 14(6), pp. 1326-1338.
- [30] Bourbigot, S., Duquesne, S., Vannier, A., and Delobel, R., 2008, "Nonwoven as heat barrier: Modeling of the efficiency of Carbtex fibers," *Journal of Applied Polymer Science*, 108(5), pp. 3245-3255.
- [31] Karamanos, A., Papadopoulos, A., and Anastasellos, D., "Heat transfer phenomena in fibrous insulating materials."
- [32] McQuillan, F., Culham, J., and Yovanovich, M., "Properties of Dry Air as One Atmosphere, Microelectronics Heat Transfer Lab," Rept. UW/M HTL, 8406.
- [33] Rattner, A., and Bohren, J., 2008, "Heat and Mass Correlations," School of Engineering and Applied Science of the University of Pennsylvania: Philadelphia, PA, USA.
- [34] Montgomery, R. B., 1947, "VISCOSITY AND THERMAL CONDUCTIVITY OF AIR AND DIFFUSIVITY OF WATER VAPOR IN AIR," *Journal of Meteorology*, 4(6), pp. 193-196.
- [35] Zuev, A. V., and Prosuntsov, P. V., 2014, "Model of the Structure of Fibrous Heat-Insulating Materials for Analyzing Combined Heat Transfer Processes," *Journal of Engineering Physics and Thermophysics*, 87(6), pp. 1374-1385.
- [36] Grinchuk, P. S., 2014, "Contact Heat Conductivity Under Conditions of High-Temperature Heat Transfer in Fibrous Heat-Insulating Materials," *Journal of Engineering Physics and Thermophysics*, 87(2), pp. 481-488.

- [37] Rader, D. J., Castaneda, J. N., Torczynski, J. R., Grasser, T. W., and Trott, W. M., 2005, "Measurements of thermal accommodation coefficients," Sandia National Laboratories.
- [38] Goodman, F. O., and Wachman, H. Y., 1967, "Formula for Thermal Accommodation Coefficients," *The Journal of Chemical Physics*, 46(6), pp. 2376-2386.
- [39] Haji-Sheikh, A., and Vafai, K., 2004, "Analysis of flow and heat transfer in porous media imbedded inside various-shaped ducts," *International Journal of Heat and Mass Transfer*, 47(8), pp. 1889-1905.
- [40] Schlichting, H., and Gersten, K., 2017, "General Properties of the Equations of Motion," *Boundary-Layer Theory*, Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 83-99.
- [41] Jaluria, Y., 1980, "Natural convection, heat and mass transfer," Unknown Host Publication Title, Pergamon Press.
- [42] Schweiz, M. V., 2012, "Table of Emissivity of various surfaces," Schaffhausen: Mikron Instrument, nd Print.
- [43] Larkin, B. K., and Churchill, S. W., 1959, "Heat transfer by radiation through porous insulations," *AIChE Journal*, 5(4), pp. 467-474.
- [44] Qashou, I., Tafreshi, H. V., and Pourdeyhimi, B., 2009, "An Investigation of the Radiative Heat Transfer through Nonwoven Fibrous Materials," *Journal of Engineered Fibers and Fabrics*, 4(1), p. 155892500900400103.
- [45] Liang, Y.-D., and Barsky, B. A., 1984, "A New Concept and Method for Line Clipping," *ACM Transactions on Graphics*, 3(1), pp. 1-22.

- [46] Kodituwakku, R., Wijeweera, K. R., and P.Chamikara, M. A., 2012, "An Efficient Line Clipping Algorithm for 3D Space," *International Journal of Advanced Research in Computer Science and Software Engineering*, 2(5), pp. 96-101.
- [47] Talreja, R., and Singh, C. V., 2012, *Damage and Failure of Composite Materials*, Cambridge University Press, Cambridge.
- [48] Hill, R., 1965, "A self-consistent mechanics of composite materials," *Journal of the Mechanics and Physics of Solids*, 13(4), pp. 213-222.
- [49] Halpin, J. C., 1969, "Stiffness and Expansion Estimates for Oriented Short Fiber Composites," *Journal of Composite Materials*, 3(4), pp. 732-734.
- [50] Eshelby, J. D., 1957, "The determination of the elastic field of an ellipsoidal inclusion, and related problems," *Proceedings of the royal society of London. Series A. Mathematical and physical sciences*, 241(1226), pp. 376-396.
- [51] Mori, T., and Tanaka, K., 1973, "Average stress in matrix and average elastic energy of materials with misfitting inclusions," *Acta Metallurgica*, 21(5), pp. 571-574.
- [52] Cox, H. L., 1952, "The elasticity and strength of paper and other fibrous materials," *British Journal of Applied Physics*, 3(3), pp. 72-79.
- [53] Christensen, R. M., and Waals, F. M., 1972, "Effective Stiffness of Randomly Oriented Fibre Composites," *Journal of Composite Materials*, 6(3), pp. 518-535.
- [54] Manera, M., 1977, "Elastic Properties of Randomly Oriented Short Fiber-Glass Composites," *Journal of Composite Materials*, 11(2), pp. 235-247.
- [55] Pan, N., 1996, "The Elastic Constants of Randomly Oriented Fiber Composites: A New Approach to Prediction," *Science and Engineering of Composite Materials*, p. 63.



- [56] Meador, M. A. B., Nguyen, B. N., Guo, H., Vivod, S., He, Z., Malow, E., and Silva, R., "Improvements to the Synthesis of Polyimide Aerogels," NASA Website.
- [57] Schwan, M. a. R. L., 2013, "Flexibilisation of resorcinol-formaldehyd aerogels," *Journal of Materials Chemistry A*, 1(43), p. 13462.
- [58] Benazzouk, A., Douzane, O., Mezreb, K., Laidoudi, B., and Quéneudec, M., 2008, "Thermal conductivity of cement composites containing rubber waste particles: Experimental study and modelling," *Construction and Building Materials*, 22(4), pp. 573-579.
- [59] Zhou, W., Qi, S., Li, H., and Shao, S., 2007, "Study on insulating thermal conductive BN/HDPE composites," *Thermochimica Acta*, 452(1), pp. 36-42.
- [60] Han, Z., Wood, J. W., Herman, H., Zhang, C., and Stevens, G. C., "Thermal Properties of Composites Filled with Different Fillers," *Proc. Conference Record of the 2008 IEEE International Symposium on Electrical Insulation*, pp. 497-501.
- [61] Xie, T. a. H. Y.-L. a. H. Z.-J., 2013, "Theoretical study on thermal conductivities of silica aerogel composite insulating material," *International Journal of Heat and Mass Transfer*, 58(Issues 1–2), pp. 540--552.
- [62] Tekce, H. S., Kumlutas, D., and Tavman, I. H., 2007, "Effect of Particle Shape on Thermal Conductivity of Copper Reinforced Polymer Composites," *Journal of Reinforced Plastics and Composites*, 26(1), pp. 113-121.
- [63] Progelhof, R. C., Throne, J. L., and Ruetsch, R. R., 1976, "Methods for predicting the thermal conductivity of composite systems: A review," *Polymer Engineering & Science*, 16(9), pp. 615-625.
- [64] Nielsen, L. E., 1974, "The thermal and electrical conductivity of two-phase systems," *Industrial & Engineering chemistry fundamentals*, 13(1), pp. 17-20.

- [65] Sundstrom, D. W., and Lee, Y.-D., 1972, "Thermal conductivity of polymers filled with particulate solids," *Journal of Applied Polymer Science*, 16(12), pp. 3159-3167.
- [66] Weber, E. H., 2001, "Development and Modeling of Thermally Conductive Polymer/Carbon Composites Dissertation," PhD PhD, Michigan Technological University.
- [67] Hasselman, D. P. H., Donaldson, K. Y., and Geiger, A. L., 1992, "Effect of Reinforcement Particle Size on the Thermal Conductivity of a Particulate-Silicon Carbide-Reinforced Aluminum Matrix Composite," *Journal of the American Ceramic Society*, 75(11), pp. 3137-3140.
- [68] Ghosh, S., Singh, S., and Maity, S., 2019, "Thermal Insulation Behavior of Chemically Treated Jute Fiber Quilt," *Journal of Natural Fibers*, pp. 1-13.

## A Derivations

### A.1 Fibre Final Position

The initial point of each fibre is initially assumed to be (0,0,0) to simplify the calculation; the end point is then sifted relative to (0,0,0) based on the position of the initial point. The length of the fibre in 3D can be calculated entirely from the position of the final point in this case, as per Pythagoras's theorem in three dimensions:

$$l = \sqrt{x^2 + y^2 + z^2} \quad (80)$$

The in plane ( $\theta_{IP}$ ) and out-of-plane ( $\theta_{OOP}$ ) angles can be related to the final point positions using two-dimensional trigonometry as shown:

$$y = x \cdot \tan(\theta_{IP}) \quad (81)$$

$$z = x \cdot \tan(\theta_{OOP})$$

Substituting these two identities in allows us to rearrange Equation (80) into a form that allows us to calculate  $x$  as shown:

$$x = \sqrt{\frac{l^2}{1 + \tan(\theta_{IP})^2 + \tan(\theta_{OOP})^2}} \quad (82)$$

This calculation always produces a positive  $x$  value, when in reality a negative  $x$  value is just as acceptable and as a result, at the end of the final position calculation, there's a 50% chance that the  $x$ -coordinate will be flipped and become negative. Once the  $x$  value is known, Equation (81) can be used to calculate the  $y$  and  $z$  values, which are then added to the initial ( $x, y, z$ ) co-ordinates to get the actual position of the final point.

## B Fibre Rotation

Some work was one on rotating the fibres upon reaching this point. Specifically, the point of contact was considered to be a "pivot point" and the length of the fibre on

each side of the pivot computed. The longer side was then lowered until it collided with either another fibre or the bottom of the bounding region. This presented some problems with significant fibre overlap occurring at the pivot location; and the calculation for where the fibre hits the bottom of the bounding region turned out to be extraordinarily complex and required solving using Maxima, an open source program for doing symbolic maths in. The solution provided by it can be calculated using MATLAB code fairly easily despite the complexity. The fibre overlap was a more major factor in not using the code, because it made actually generating a fibre mat in APDL very challenging as it was not well suited to this level of complexity. Also no consensus was reached on the appropriate action to take when the fibre was both too high in the geometry for the rotation to allow it to touch the bottom of the region but also did not have a fibre below it, meaning that the current model just leaves it where it is instead of rotating it.

## C Fibre Generation

%% WARNING: this script overwrites the csv files for the fibres silently

% Introduction

%

% This code will generate a fibre distribution inside a 3D cuboid bound between two points. All the variables are defined in this script, so it takes no arguments and produces a range of outputs. It creates two main outputs: a CSV file parsable by APDL so that APDL can reproduce the fibre distribution; and a 3D scatter plot showing each of the fibres as a different line. The fibre matrix and other set variables are accessible once the script has finished running. This scatter plot doesn't include the diameter of the fibres, so checking whether excessive fibre overlap is occurring is difficult to do until the geometry is imported into APDL.

% Data on each fibre is stored in a separate row of the "fibre" matrix. This stores the following information: the fibre number, the fibre diameter, the [x,y,z] co-ordinate of each end point of the fibre and the length of the fibre. The fibre number is constant (i.e., the 206th fibre to be generated will always be numbered 206) and so it doesn't always correspond to number of the row in the fibre matrix. The diameter also never changes; but the positions will (almost certainly) change during clipping and when moving the fibre. Also, the length can change if the fibre is shortened by clipping.

% The generation process follows a four step process: 1) Generate a fibre initially, with a random length, diameter and single endpoint. The orientation is randomised and the other end point is found using the length of the fibre and two angles. 2) Find the closest fibre vertically below the new fibre, and move the new fibre to be in contact with it. 3) Clip the fibre to the bounding box, and recalculate the length. 4) Delete any fibres completely out of the box or which are too short.

% In order to reduce the number of fibres checked for intersection, a quadrant system is used. Specifically, for each fibre the intersection points with the x and y axis are calculated, and then checked to see which quadrants the fibre is in. If the fibre is wholly in one quadrant, then this is also detected. This is one of many really cool things the code does.

% Once the generation process has finished, we proceed to produce an APDL file to generate a full 3D geometry of the fibres, which can be used to run simulations on the fibres. This file is saved in one of two places. If the folder the script is in was added to MATLAB's path, then it will be saved in the \$USER/Documents/MATLAB folder. If the folder the script is in is set as the folder (the 'Change Folder' option), then it will be saved in that folder. Rerunning the script will overwrite this file, as long as it isn't currently in use by another program.

% At the end of the generation process, the volume fraction is also calculated using two different methods. The first is the 'Volume Counting' method, where the volume of each fibre is added up and the ratio of the fibre volume to bounding region volume is calculated. The second is the 'Monte Carlo' method, where

```

% random points are generated and then tested to see if they're on a fibre. The
first method is very quick to calculate, but tends to overestimate the volume
% fraction, since some fibre overlap is created to improve the accuracy of
thermal insulation. The 'monte Carlo' method is much slower, but more accurate.
The
% difference between the two methods can be as much as 10 times.

% The script uses a range of functions originally written by other people, which
were modified to work better for our uses.
% 'lineSegmentIntersect.m' was originally produced by U. Murat Erdem, and
released under the BSD license. The original source is found here:
% http://uk.mathworks.com/matlabcentral/fileexchange/27205-fast-line-segment-intersection
% It has been modified to use only two line segments, instead of a series of line
segments, and also so that it calculates the distance in the z-direction
% between two cylinders which intersect in the x-y direction. Also, the output has
been changed to only be either the distance or 0 rather than a range of
% information about the intersection.

% 'clipped_line.m' was originally produced by Peter Krushe, and can be found
here: http://www.mass-communicating.com/code/2013/05/12/line-clipping.html
% This is a matlab implementation of the Liang-Barsky clipping algorithm. The
version used doesn't plot the results; and also outputs a single vector of the
% results instead of two vectors.

%Generation conditions
%-----
-----
% Basically just so I can easily check whether I've run the script to test changes
etc.
disp(datestr(now))

% Fibres to generate; includes deleted fibres but not vertical fibres.
numfibres = 3000;

% Vertical fibres to generate. Due to the largely uncontrolled manner of fibre
generation, the actual number of vertical fibres will be less.
vertfibres = 0;

totfibres = numfibres+vertfibres;

% Boolean tolerance for APDL.
btol = 2e-8;

% Tolerance for MCVF calculation. 1e-5 seems to be a good level to have this at,
since the results look accurate and don't take much time to obtain.
mctol = 1e-5;

%[smallest,biggest] for next two arrays
% fibrediameterrange = [70e-6, 70e-6];
% fibrelengthrange = [2e-2 2e-2];

% [mean value, %deviation] for next two arrays; used to generate a normal
distribution centred on first value with a standard deviation equal to the second

```

% value percent of the first (e.g, [10, 1] would generate a normal distribution at 10, with a standard deviation of 1% of 10, or 0.1).

```
gffibrediameterrange = [8e-6, 0];
```

```
petfibrediameterrange = [8e-6, 0];
```

% Multiply actual diameter by this to account for crimping; basically we spend a lot of time treating the PET fibres as being the fake diameter so that we can % ignore the curved shape of the fibres, which is hella complicated, until we actually need to make them curved.

```
fakepetdiameterinc = 10;
```

% Number of crimps. Less realistic, easier coding.

```
crimps = 10;
```

% Ratio of glass fibres to total fibres. Note that since some fibres are deleted, the actual ratio will vary. It's possible that one kind of fibre is easier to % fit into the bounding region than the other, so this value may be very inaccurate but this is usually only the case when very few fibres are generated.

% Note that you may need to manually edit the code if this is either 0 or 1 to ensure the right fibre type is inserted first. Otherwise, it doesn't matter which % is first.

```
gfr = 1;
```

% Planeangle is the range of allowed in-plane angles. Specifically, this is the angle between the x-axis and the fibre in the x-y plane.

% Specified as two angles.i

```
planeanglerange = [0 pi*2];
```

% Inclination is the out-of-plane angle of each fibre. Measured in radians. Specifically, an inclination of 0 is 'horizontal' (i.e, z is constant along the fibre length), and an inclination of pi

% is 'vertical' (i.e, x,y are the same and only z varies along fibre length).

% Specified as an angle and a percent standard deviation.

```
inclination = [0 0];
```

% 3D cuboid region of the form [x1, y1, z1, x2, y2, z2]; specifying diagonally opposite corners. Note that the first point has to be [0,0,0] for the block % calculation to produce an air block enclosing the fibres.

```
boundingregion = [0,0,0,4.5e-3,4.5e-3,4.5e-3];
```

```
flips = 0;
```

% Note that fibres are deleted if they are too short compared to the first number. This means that if it is much larger than the size of the bounding box,

% most/all of the fibres will be deleted.

```
brr = 2;
```

```
gffibrelengthrange = [brr*boundingregion(5), 5];
```

```
% gffibrelengthrange = [brr*20, 0];
```

```
petfibrelengthrange = gffibrelengthrange;
```

% Calculate the actual standard deviation rather than % deviation.

% % deviation is used because it makes having some variation easier to do when varying parameters since it doesn't need to be recalculated.

```

% The actual deviation is more useful though, since we can use it as an RNG
parameter.
gfdiamstd = gffibrediameterrange(1)*(gffibrediameterrange(2)/100);
gflengthstd = gffibrelengthrange(1)*(gffibrelengthrange(2)/100);
petdiamstd = petfibrediameterrange(1)*(petfibrediameterrange(2)/100);
petlengthstd = petfibrelengthrange(1)*(petfibrelengthrange(2)/100);

% Minimum allowable fibre length. Short fibres tend to force a higher density
mesh (at least on that fibre) while also not contributing much to the overall heat
% transfer (since they're likely to only be in contact with one other fibre).
% Currently, we just use one minimum length; but we can expand to two if that is
necessary.
% minlength = min(gffibrelengthrange(1),petfibrelengthrange(1))/5;
% minlength = gffibrelengthrange(1)/4;
minlength = boundingregion(6)/2;

% Fibre generation is centered on the centre of the bounding region, so we
ensure that the density doesn't drop off at the edges by generating fibres over a
% larger area than the size of the x-y plane of the bounding region. This controls
how much larger that area is.
fuzzfactor = 0.5;

% Ensure that the fibre is definitely completely out of the bounding region, and
that all the generated fibres will be once they're stacked.
% This is mostly useful when the fibres are horizontal, so that the Liang-Barsky
clipping algorithm is used, since the 3D clipping algorithm will clip in the
% z-direction as well.
initialheight = 10*boundingregion(6);

%Preinitialise fibre and deletion matrices for maximum speed. Note that the
deletion matrix is initially set to the maximum possible size, since the number of
%fibres to be deleted is unknown.
fibre=zeros(totfibres,10);
% fibretop=zeros(totfibres,12);
deletion = zeros(totfibres,1);

% Pre-allocate quadrants for speed.
quadrant1 = zeros(totfibres,1);
quadrant2 = zeros(totfibres,1);
quadrant3 = zeros(totfibres,1);
quadrant4 = zeros(totfibres,1);

% Pre-calculate quadrant boundaries
qbx = (boundingregion(4)-boundingregion(1))/2;
qby = (boundingregion(5)-boundingregion(2))/2;

origfibre = zeros(size(fibre,1),size(fibre,2));
lengdel = 0;
clipdel = 0;

gfvol = 0;
totvol = 0;

gfcount = 0;
totcount = 0;

```



```

impossiblecount = 0;
screwyrotation = 0;
% Generate geometry
%-----
-----
for iter=1:totfibres
    % Randomly generating the diameter and length based on the min/max
    diameter and length.

    % diameter = randb(fibrediameterrange(1),fibrediameterrange(2));
    % length = randb(fibrelengthrange(1),fibrelengthrange(2));
    % Generate two initial angles, indicating the angle between x and y,z.
    initialanglexy = randb(planeanglerange(1), planeanglerange(2));

    % We want to generate the specified number of fibres using the standard fibre
    generation parameters (so, specified inclination, slightly larger bounding
    % region). We also want to generate a small number of vertical fibres to
    represent the needle punched fibres.
    if iter <= numfibres
        % Generate fibres initially either inside or just outside the boundingregion,
        so that some fibres can start outside and enter the region. Weighted so that
        % the centre of the bounding region should correspond with the average
        initial position. Both PET and Glass can be generated this way.
        initialposition = [randb(-
        fuzzfactor*boundingregion(4),(1+fuzzfactor)*boundingregion(4)),randb(-
        fuzzfactor*boundingregion(5),(1+fuzzfactor)*boundingregion(5))];
        initialanglexz = randb(inclination(2),inclination(1));
        if gfvol/totvol <= gfr || isnan(gfvol/totvol) == 1
        %   if gfvol/totvol < gfr
            maxdiameter = abs(gfdiamstd*randn(1)+gffibrediameterrange(1));
            length = gflengthstd*randn(1)+gffibrelengthrange(1);
            % 0 ==> same as max, not that the real diameter is 0. This is just
            because PET fibres are complicated and I use two diameters to describe them.
            realdiameter = 0;
            gf = 1;
        else
            realdiameter = abs(petdiamstd*randn(1)+petfibrediameterrange(1));
            length = petlengthstd*randn(1)+petfibrelengthrange(1);
            maxdiameter = fakepetdiameterinc*realdiameter;
            gf = 0;
        end
        % The end point is constrained by two angles and a length. The relation
        between the three distances and the length is found using pythagoras's theorem,
        and basic
        % trigonometry can be used to related the angle between x and y with the
        lengths. This gives a system of three equations which can be solved
        simultaneously
        % for x,y and z
        a = 1+tan(initialanglexy)^2+tan(initialanglexz)^2;
        finalposition(1) = sqrt(length^2/a);
        finalposition(2) = finalposition(1)*tan(initialanglexy);
        finalposition(3) = finalposition(1)*tan(initialanglexz);

        % The x-position is always positive. So, generate a random float between 0
        and 1 and if it's above 0.5, then we flip the sign. This means that approximately

```

```

    % 50% of the time, the x-position becomes negative. The other values can
    be negative if need be.
    if rand(1) >= 0.5
        finalposition(1) = -finalposition(1);
    end

    %The final position is currently relative to (0,0,0). Add in the initial position
    so that we can get accurate results.
    finalposition = [finalposition(1)+initialposition(1)
    finalposition(2)+initialposition(2) finalposition(3)+initialheight];
    else
        % These are the only different lines between the two fibres, since we want
        to add the needlepunch fibres in the same way; and also clip them to be in
        % the box as well. This ensures that the fibre doesn't clip through other
        fibres as well.
        % Needlepunch fibres are verticle, so force them to be within the bounding
        region in the x-y direction.
        initialposition =
        [randb(boundingregion(1),boundingregion(4)),randb(boundingregion(2),bounding
        region(5)),boundingregion(6)+initialheight];
        % Verticle fibre ==> fixed angle of pi/2.
        initialanglexz = pi/2;
        realdiameter = 0;
        length = gflengthstd*randn(1)+gffibrelengthrange(1);
        finalposition = [initialposition(1:2), initialposition(3)+length];
        maxdiameter = abs(gfdiamstd*randn(1)+gffibrediameterange(1));
        gf = 1;
    end
    % Add all the fibre data together to produce the fibre data, without calculating
    where the fibre is vertically

    fibre(iter,:)= [iter,maxdiameter,initialposition(1),initialposition(2),initialheight,final
    position(1:3),length,realdiameter];

    if fibre(iter,10) == 0
        fibretheta = 90-atan2((fibre(iter,8)-fibre(iter,5))/(fibre(iter,6)-fibre(iter,3)));
        a = fibre(iter,2)/2*sind(fibretheta);
        b = fibre(iter,2)/2*cosd(fibretheta);
    end

    % Figure out where the fibre crosses the quadrant boundaries. Note that we
    go really far out the bounding region to ensure that we can split everything into
    % four quadrants accurately; which is basically because this needs a finite
    distance but we want infinite separation.
    [~,yint] = polyxpoly(fibre(iter,[3,6]),fibre(iter,[4,7]), [qbx qbx], [-
    100*boundingregion(5), 100*boundingregion(5)]);
    [xint,~] = polyxpoly(fibre(iter,[3,6]),fibre(iter,[4,7]), [-100*boundingregion(4),
    100*boundingregion(4)], [qby, qby]);

    % No intersection with quadrant walls --> Wholly in one quadrant. Compare
    x,y co-ordinates with quadrant boundaries to determine which quadrant.
    if isempty(yint) && isempty(xint)
        if fibre(iter,3) > qbx
            if fibre(iter,4) > qby
                quadrant1(iter) = 1;
            else
                quadrant4(iter) = 1;
            end
        end
    end

```

```

        end
elseif fibre(iter,4) > qby
    quadrant2(iter) = 1;
else
    quadrant3(iter) = 1;
end
else
    % Only two checks are required since a straight line can't cross both the
    +ve and -ve sides of one axis.
    % If it crosses the +ve y axis, then it must be in quadrants 1 and 2. If -ve y-
    axis, then 3 and 4.
    if yint > qby
        quadrant1(iter) = 1;
        quadrant2(iter) = 1;
    elseif yint < qby
        quadrant3(iter) = 1;
        quadrant4(iter) = 1;
    end
    % +ve x ==> quadrants 1 and 4. -ve x ==> quadrants 2 and 3.
    if xint > qbx
        quadrant1(iter) = 1;
        quadrant4(iter) = 1;
    elseif xint < qbx
        quadrant2(iter) = 1;
        quadrant3(iter) = 1;
    end
end
end
% Set the minimum distance to a large number and closestfibre to 0 so that
any closer fibres are counted and the right fibre is used for the height
% calculation
mindistance = 3000000;
closestfibre = 0;
% % This loop figures out where the closest fibre is, if it exists, and then
positions the next fibre above it
for x = 1:iter-1

    % Check whether the fibres intersect (2D, so we're ignoring the fact that
    one fibre will be much higher than the other)
    % First, figure out if the two fibre share at least one common quadrant. If
    not, then they can't intersect.
    if (quadrant1(x) == 1 && quadrant1(iter) == 1) || (quadrant2(x) == 1 &&
    quadrant2(iter) == 1) || (quadrant3(x) == 1 && quadrant3(iter) == 1) ||
    (quadrant4(x) == 1 && quadrant4(iter) == 1)
        if fibre(iter,5) == fibre(iter,8)
            [y,xint,yint,zint] =
lineSegmentIntersect(fibre(iter,3:8),fibre(x,3:8),fibre(iter,2)/1.8,fibre(x,2)/1.8);
            elseif fibre(x,10) == 0 && fibre(iter,10) == 0
                [y,xint,yint,zint] =
lineSegmentIntersect(fibre(iter,3:8),fibre(x,3:8),fibre(iter,2)/1.8,fibre(x,2)/1.8);
            elseif fibre(x,10) ~= 0 && fibre(iter,10) == 0
                [y,xint,yint,zint] =
lineSegmentIntersect(fibre(iter,3:8),fibre(x,3:8),fibre(iter,2)/1.8,fibre(x,10)/1.8);
            elseif fibre(x,10) == 0 && fibre(iter,10) ~= 0
                [y,xint,yint,zint] =
lineSegmentIntersect(fibre(iter,3:8),fibre(x,3:8),fibre(iter,10)/1.8,fibre(x,2)/1.8);
            elseif fibre(x,10) ~= 0 && fibre(iter,10) ~= 0

```

```

        [y,xint,yint,zint] =
lineSegmentIntersect(fibre(iter,3:8),fibre(x,3:8),fibre(iter,10)/1.8,fibre(x,10)/1.8);
    end
    else
        % No shared quadrants ==> no intersection.
        y = 0;
    end
    if y ~= 0
        %The previous if statement should determine whether there's an
intersection (non-zero y) or not;
        if y < mindistance
            % Fibre x is the current closest fibre. So now we record which one it is,
and the distance to it.
            closestfibre = x;
            mindistance = y;
            pivx = xint;
            pivy = yint;
            pivz = zint;
        end
    end
end
end
% No closest fibre
if closestfibre == 0
    % x is set as the minimum fibre height.
    x = boundingregion(3)+fibre(iter,2)/2;
    % Figure out which end of the fibre is closer to being at the minimum
height, then how far it has to move to get there
    [~,tmp] = min(fibre(iter,[5,8]));
    if tmp == 1
        change = fibre(iter,5) - x;
    else
        change = fibre(iter,8) - x;
    end
    % Finally, move both ends of the fibre the same distance. This should
preserve the fibre gradient while translating it as far downwards as allowed
    fibre(iter,5) = fibre(iter,5) - change;
    fibre(iter,8) = fibre(iter,8) - change;
else
    % Mindistance tells us how close the nearest fibre is vertically. Move the
new fibre by the separation distance.
    fibre(iter,5) = fibre(iter,5)-mindistance;
    fibre(iter,8) = fibre(iter,8)-mindistance;
    pivratio =
(vecleng3d([fibre(iter,3:5),pivx,pivy,pivz])/(vecleng3d(fibre(iter,3:8))));
end

origfibre(iter,:) = fibre(iter,:);

% Applying boundary region
%-----
-----

[fibre(iter,:), deletion] = clip(fibre(iter,:),[boundingregion(1:3)-
boundingregion(4)/100,boundingregion(4:5)+boundingregion(4)/100
boundingregion(6)],minlength,iter,deletion);

```

```

% Move deleted fibres away from the region of interest. This prevents other
fibres being stacked on top of them,
if deletion(iter) ~= 0
    if vecleng3d(fibre(iter,3:8)) < minlength
        lengdel = lengdel+1;
    else
        clipdel = clipdel+1;
    end
    fibre(iter,3:8) = 90000;
    % Also set it to be not in any of the four quadrants. This ensures we never,
ever compute the distance to the deleted fibre.
    quadrant1(iter) = 0;
    quadrant2(iter) = 0;
    quadrant3(iter) = 0;
    quadrant4(iter) = 0;
else
    if gf == 1
        newfibrevol = pi*(fibre(iter,2)/2)^2*fibre(iter,9);
        gfvol = gfvol+newfibrevol;
    else
        newfibrevol =
pi*(fibre(iter,10)/2)^2*8*sqrt((fibre(iter,9)/8)^2+(fibre(iter,2)/2)^2);
    end
    totvol = totvol+newfibrevol;
    %     if gf == 1
    %         gfcount = gfcount+1;
    %     end
    %     totcount = totcount+1;
    % Recalculate the lines
    fibretheta = 90-atan2((fibre(iter,8)-fibre(iter,5))/(fibre(iter,6)-fibre(iter,3)));
    % Check fibre type and use the right fibre diameter
    if fibre(iter,10) == 0
        a = fibre(iter,2)/2*sind(fibretheta);
        b = fibre(iter,2)/2*cosd(fibretheta);
    else
        a = fibre(iter,10)/2*sind(fibretheta);
        b = fibre(iter,10)/2*cosd(fibretheta);
    end
end

% Display the number of iterations every 1000 iterations, so we can check
progress is occurring.
% Note that the time to do 1000 iterations increases each time, since there
are more fibres to compute distance to.
if mod(iter,1000) == 0
    fprintf('%i\n',iter);
end
end

% For the 3D clipping algorithm, we clip to above the top of the bounding box. So
this just makes sure the fibre is completely within the bounding box.
for iter = 1:size(fibre,1)
    % [fibre(iter,:),deletion] =
clip(fibre(iter,:),[boundingregion(1:3)+fibre(iter,2)/2,boundingregion(4:6)-
fibre(iter,2)/2],minlength,iter,deletion);
end

```

```

    [fibre(iter,:),deletion] =
clip(fibre(iter,:),boundingregion(1:6),minlength,iter,deletion);
end

% Remove all 0 entries from the deletion matrix so that we don't break the script
(fibre(0,:)) is undefined because matlab counts from 1)
%fprintf('Length Deletion: %i; Clip Deletion: %i\n',lengdel,clipdel);
deletion = deletion(deletion~=0);

% Actually delete all the rows we don't need.
fibre(deletion,:) = [];

% Abort if there are no fibres now. Only happens if the length means that fibres
can't fit into the boundingregion without being short.
if size(fibre,1) == 0
    return
end

deletion = zeros(size(fibre,1),1);
for iter = 1:size(fibre,1)
    fibre(iter,:) = clip(fibre(iter,:),[boundingregion(1:3)+fibre(iter,2)/2,
boundingregion(4:6)-fibre(iter,2)],minlength,iter,deletion);
    if vecleng3d(fibre(iter,3:8)) < minlength
        deletion(iter) = iter;
    end
end
% Clear out the deleted fibres.
deletion = deletion(deletion~=0);
fibre(deletion,:) = [];

% % Same as above, but done by adding up the fibre volume and finding the
ratio of that to the bounding region volume.
dupsvf2 = vcvf(fibre,boundingregion);

% Write results to file
% -----
% -----
%
% Open file in loop with delay, so if the file is in use, the script will wait until it's
available again.
% Also displays how many times it loops so that we can be confident that it's not
looping infinitely.
fid = -1;
count = 0;
while fid < 0
    % Open a file for editing. 'fid' is the file id, and allows us to interact with the
file.
    % The file is saved in the directory shown just above the editor window in
MATLAB. This is usually either the folder the script is in, or the
Documents\MATLAB
    % folder.
    fid = fopen('1gf.csv','wt');
    if fid < 0
        % If the file didn't open, wait 5 seconds and try again.
        pause(5)
    end
end

```

```

    count = count+1;
    if count ~= 1
        % Print the wait count so that it's easy to tell if there's an issue opening the
file.
        fprintf('Count: %i\n',count);
    end
end

% End and clear any currently started simulations, and move into preprocessing.
fprintf(fid,'finish\n');
fprintf(fid,'/CLEAR\n');
fprintf(fid,'/PREP7\n');

% Wrap it in a do loop, set to loop just once. This allows us to exit the script more
easily.
% fprintf(fid,'*do,yellow,1,1,1\n');

% The easiest way to create a cylinder in an arbitrary direction in APDL is to
change the working plane to the desired cylinder direction and then just specify
% the radius and length of the cylinder. Since multiple fibre additions are used
(at least 2: one to initially add them, then they get subtracted from a block
% and then re-added, but if the block is split into multiple blocks, then there can
be a large number of fibre additions), a new fibre matrix is set up so we can
% easily just read off the fibre matrix. This should also improve the run time,
since everything only needs to be computed once.
neofibre = zeros(size(fibre,1),8);
for k=1:size(fibre,1)
    % Horizontal fibres
    if fibre(k,5) == fibre(k,8)
        angles = [atand((fibre(k,7)-fibre(k,4))/(fibre(k,6)-fibre(k,3)))
atand((fibre(k,8)-fibre(k,5))/(fibre(k,7)-fibre(k,4))) atand((fibre(k,6)-
fibre(k,3))/(fibre(k,8)-fibre(k,5)))];
        elseif fibre(k,3) == fibre(k,6) && fibre(k,4) == fibre(k,7)
            angles = [0 0 0];
        else
            angle1 = angle(fibre(k,6)-fibre(k,3),fibre(k,7)-fibre(k,4));
            newpos = xyleng([fibre(k,6)-fibre(k,3),fibre(k,7)-fibre(k,4),0,0]);
            angle2 = angle(newpos, fibre(k,8)-fibre(k,5));
            angles = [angle1 angle2 0];
        end
        neofibre(k,:) = [fibre(k,3:5),fibre(k,2)/2,fibre(k,9),angles];
    end

gfcoun = 0;
avert = 0;
for k = 1:size(neofibre,1)
    if fibre(k,10) == 0
        % To ensure the working plane is just translated from (0,0,0), we also have
to specify where the x-axis and y-axis is pointing. This is done relative to the
% new position, so we just add one to the x position and y position for each.

fprintf(fid,'wplane,1,%.8f,%.8f,%.8f,%.8f,%.8f,%.8f,%.8f,%.8f\n',neofibre(k,1:
3),neofibre(k,1)+1,neofibre(k,2:3),neofibre(k,1),neofibre(k,2)+1,neofibre(k,3));
        % Now rotate the working plane so that the z-axis points along the fibre
length
        fprintf(fid,'wprota,%.f,%.f, %f\n',neofibre(k,6:8));
        % No realdiameter ==> glass fibre, so make a nice cylinder

```

```

fprintf(fid,'cylind,%.12f,0,%.12f,0,0,0\n',neofibre(k,4),neofibre(k,5));

% Count how many fibres are vertical.
% This is accurate as long as inclination = [0 0] (i.e, all non-vertical fibres
are horizontal.
% This is a horrible, terrible asumption in a lot of cases, but works when I
need it to.

if fibre(k,5) ~= fibre(k,8)
    avert = avert+1;
end

    gfcoun = gfcoun+1;
end
end

% Select everything again. Just to make sure it all visible and stuff.
fprintf(fid,'vsel,all\n');

fprintf(fid,'igesout,1gf,iges\n');

gfratio = gfvol/totvol;
% gfratio = gfcoun/size(fibre,1);
fprintf('GF ratio: %.6f; VF: %.6f; |f: %i; impossiblecount: %i, screwedup:
%i\n',gfratio,dupsvf2,avert,impossiblecount,screwrotation);

fclose('all');

fid = -1;
count = 0;
while fid < 0
    % Open a file for editing. 'fid' is the file id, and allows us to interact with the
file.
    % The file is saved in the directory shown just above the editor window in
MATLAB. This is usually either the folder the script is in, or the
Documents\MATLAB
    % folder.
    fid = fopen('2petf.csv','wt');
    if fid < 0
        % If the file didn't open, wait 5 seconds and try again.
        pause(5)
    end
    count = count+1;
    if count ~= 1
        % Print the wait count so that it's easy to tell if there's an issue opening the
file.
        fprintf('Count: %i\n',count);
    end
end

% End and clear any currently started simulations, and move into preprocessing.
fprintf(fid,'finish\n');
fprintf(fid,'/CLEAR\n');
fprintf(fid,'/PREP7\n');

nkp = 1;
lkp = 1;

```



```

akp = 1;
tmp = curvedfibres(fibre,neofibre,crimps,fid,nkp,lkp,akp,angles);
% Reset working plane to (0,0,0) and standard directions.
fprintf(fid,'wpstyl,defaln');

% Select everything again. Just to make sure it all visible and stuff.
fprintf(fid,'vsel,all\n');

% Write the geometry to apdl.iges. This will be saved in the directory you've
chosen as a working directory in APDL.
fprintf(fid,'igesout,2petf,iges\n');

% Close the file
fclose('all');

% Create a 3d plot of the fibres so that things like the uniformity of the fibre
distribution, and how much fibre clustering is occurring, can be seen visually.
% Note that this doesn't show the fibre diameter, so it's not that useful for
checking for appropriate levels of intersection.
figure(1)
plot3(fibre(:,[3,6]),fibre(:,[4,7]),fibre(:,[5,8]));

%
% figure(2)
% plot3(fibre(:,[3,6]),fibre(:,[4,7]),fibre(:,[5,8]));

% For debugging purposes, it can be handy to see how the fibres change once
clipped or moved downwards.
% figure(2)
% plot3(origfibre(:,[3,6]),origfibre(:,[4,7]),origfibre(:,[5,8]));
% title('Original Fibre');

```

## C.1 clipped\_line.m

*Note that the following MATLAB code needs to be saved as 'clipped\_line.m' and located in a path that MATLAB can find. This is an implementation of a 2D Clipping Algorithm used by the main fibre generation algorithm.*

```

function out = clipped_line(vx1, vx2, v1, v2)
% Liang-Barsky Algorithm for line clipping. See
% http://en.wikipedia.org/wiki/Liang%E2%80%93Barsky\_algorithm
%
% Take a 2-D line from vx1 to vx2, and clip to rectangle with corners at
% v1, v2.
%
% Inputs must be vectors of length 2.
% The output is a vector of length 4 of the form [x1 y1 x2 y2], where (x1,y1) and
% (x2,y2) are the end points of the line after clipping
% Or NaN in the case where the whole line is outside the rectangular clipping
area

```

```

if length(vx1) ~= 2 || length(vx2) ~= 2 || length(v1) ~= 2 || length(v2) ~= 2
    error('All inputs must be vectors of length 2.');
```

```

end

x0 = vx1(1);
x1 = vx2(1);
y0 = vx1(2);
y1 = vx2(2);

x_min = min(v1(1), v2(1));
x_max = max(v1(1), v2(1));
y_min = min(v1(2), v2(2));
y_max = max(v1(2), v2(2));

out = nan(1,4);

dx = x1 - x0;
dy = y1 - y0;

p = [ -dx;
      dx;
      -dy;
      dy ];

q = [ x0 - x_min; % negative => left of window
      x_max - x0; % negative => right of window
      y0 - y_min; % negative => below window
      y_max - y0; % negative => above window
    ];

if ~isempty(find(p == 0 & q < 0, 1, 'first'))
    return;
end

u1 = 0;
u2 = 1;

for i = 1:4
    if p(i) < 0
        u1 = max(u1, q(i)/p(i));
    end
    if p(i) > 0
        u2 = min(u2, q(i)/p(i));
    end
end

if u1 > u2 % line is outside
    return;
end

out(1) = x0 + dx*u1;
out(2) = y0 + dy*u1;

out(3) = x0 + dx*u2;
out(4) = y0 + dy*u2;
end

```

## C.2 clip.m

Note that the following MATLAB code needs to be saved as 'clip.m' and located in a path that MATLAB can find. This is an implementation of a 3D Clipping Algorithm used by the main fibre generation algorithm.

```
function [newfibre,deletion] = clip(fibrer,br,minleng,iter,deletion)
    if fibrer(5) >= br(6) && fibrer(8) >= br(6)
        deletion(iter) = iter;
        clipf = nan(6);
    elseif fibrer(5) == fibrer(8)
        clipf = [clipped_line(fibrer(3:4),fibrer(6:7),br(1:2),br(4:5)) 0 0];
        clipf = [clipf(1:2) fibrer(5) clipf(3:4) fibrer(8)];
    elseif fibrer(4) == fibrer(7)
        clipf = [clipped_line(fibrer([3,5]),fibrer([6,8]),br([1,3]),br([4,6])) 0 0];
        clipf = [clipf(1) fibrer(4) clipf(2:3) fibrer(7) clipf(4)];
    elseif fibrer(3) == fibrer(6)
        clipf = [clipped_line(fibrer(4:5),fibrer(7:8),br(2:3),br(5:6)) 0 0];
        clipf = [fibrer(3) clipf(1:2) fibrer(6) clipf(3:4)];
    else
        clipf = clip3d(fibrer(3:5),fibrer(6:8),br(1:3),br(4:6));
    end
    %clip3d returns NaN if the whole fibre is out the bounding region, so it can be
    safely deleted. Actually, it needs to be deleted to prevent it interfering
    %with future results
    if isequal(isnan(clipf),zeros(size(clipf,1),size(clipf,2))) ~ = 1
        deletion(iter) = iter;
    else
        fibrer(3:8) = clipf;
        %Fibre is likely shorter now, so update the length value. Saves iterating
        through them again later.
        fibrer(9) = vecleng3d(fibrer(3:8));
    end
    % Remove very short fibres since it tends to lead to a high mesh density on
    that fibre, and also not much heat transfer in the fibre
    if fibrer(9) <= minleng
        deletion(iter) = iter;
    end
    newfibre = fibrer;
end
```

## C.3 curvedfibres.m

Note that the following MATLAB code needs to be saved as 'curvedfibres.m' and located in a path that MATLAB can find. This is the code which converts the sinusoidal fibres from cylinders into sinusoidal shapes in APDL.

```
function fb = curvedfibres(dupfibre,neofibre,crimps,fid,nkp,lkp,akp,angles)
```

```

% In general, when I talk about a direction, I mean relative to the rotated axis. So
'x-direction' is assuming z points along the fibre. Note that the
% x-direction isn't really more constrained than that...
% Loop through the matrix again

```

```

maxcrimplen = max(dupfibre(:,9))/crimps;

```

```

for k=1:size(dupfibre,1)
    if dupfibre(k,10) ~=0

```

```

fprintf(fid,'wplane,1,%f,%f,%f,%f,%f,%f,%f,%f,%f\n',neofibre(k,1:
3),neofibre(k,1)+1,neofibre(k,2:3),neofibre(k,1),neofibre(k,2)+1,neofibre(k,3));

```

```

    fprintf(fid,'wprota,%f,%f,%f\n',neofibre(k,6:8));

```

```

    fprintf(fid,'csys,4\n');

```

```

    fprintf(fid,'k,%i,0,0,0\n',nkp);

```

```

    nkp = nkp+1;

```

```

    nkpo = nkp;

```

```

    crimps = ceil(dupfibre(k,9)/maxcrimplen);

```

```

    if crimps < 6

```

```

        crimps = crimps+3;

```

```

    end

```

```

    crimplen = neofibre(k,5)/(crimps-1);

```

```

    diam = dupfibre(k,2)/2;

```

```

    for l=2:crimps-1

```

```

        if l == 2

```

```

            fp = [diam,0,crimplen*(l-1)];

```

```

        end

```

```

        fprintf(fid,'k,%i,%f,%f,%f\n',nkp,diam,0,crimplen*(l-1));

```

```

        nkp=nkp+1;

```

```

        diam = -diam;

```

```

    end

```

```

    fprintf(fid,'k,%i,%f,%f,%f\n',nkp,0,0,neofibre(k,5));

```

```

    fprintf(fid,'ns_=%i\n',nkpo-1);

```

```

    fprintf(fid,'nf_=%i\n',nkp);

```

```

    fprintf(fid,'nl_=%i\n',nkpo);

```

```

    fprintf(fid,'*cfopen,spl_ine,mac\n');

```

```

    fprintf(fid,'*cfwrite,FLST,3,nl_,3\n');

```

```

    fprintf(fid,'*do,i,ns_,nf_-1\n');

```

```

    fprintf(fid,'*cfwrite,FITEM,3,i_\n');

```

```

    fprintf(fid,'*enddo\n');

```

```

    fprintf(fid,'*cfwrite,BSPLIN, ,P51X\n');

```

```

    fprintf(fid,'*cfclose\n');

```

```

    fprintf(fid,'spl_ine.mac\n');

```

```

    angles = [0 90-atand((fp(3)-neofibre(k,3))/(fp(2)-neofibre(k,2))) atand((fp(1)-
neofibre(k,1))/(fp(2)-neofibre(k,2)))]];

```

```

    fprintf(fid,'cyl4,%f,%f,%f\n',0,0,dupfibre(k,10)/2);

```

```

    fprintf(fid,'vdrag,%i,,,,,%i\n',akp,lkp);

```

```

    nkp = nkp+9;

```

```

    lkp = lkp+13;

```

```

    akp = akp+6;

```

```

    end
    fb = 0;
end

```

#### C.4 lineSegmentIntersect.m

*Note that the following MATLAB code needs to be saved as 'lineSegmentIntersect.m' and located in a path that MATLAB can find. This is the code that determines whether two fibres would meet in the x-y plane, and then determines the amount of z displacement required to make that happen.*

```

% XY1 and XY2 give the two lines being tested. XY1 = [x1 y1 z1 x2 y2 z2] and
the same for XY2, so that each line is defined by two end points.
% Deals specifically with segments, so intersection only occurs if the intersection
point is between the end points

```

```

% Obtained from: http://uk.mathworks.com/matlabcentral/fileexchange/27205-fast-line-segment-intersection/content/lineSegmentIntersect.m

```

```

function [distance, xint,yint,zint] = lineSegmentIntersect(XY1,XY2,R1,R2)
%% Prepare matrices for vectorized computation of line intersection points.
%-----

```

```

%Get the x,y,z co-ordinates

```

```

X1 = XY1(1);
X2 = XY1(4);
Y1 = XY1(2);
Y2 = XY1(5);
Z1 = XY1(3);
Z2 = XY1(6);
X3 = XY2(1);
X4 = XY2(4);
Y3 = XY2(2);
Y4 = XY2(5);
Z3 = XY2(3);
Z4 = XY2(6);

```

```

% Find the direction vector from each x set

```

```

X4_X3 = (X4-X3);
Y1_Y3 = (Y1-Y3);
Y4_Y3 = (Y4-Y3);
X1_X3 = (X1-X3);
X2_X1 = (X2-X1);
Y2_Y1 = (Y2-Y1);

```

```

% Magic

```

```

numerator_a = X4_X3 * Y1_Y3 - Y4_Y3 * X1_X3;
numerator_b = X2_X1 * Y1_Y3 - Y2_Y1 * X1_X3;

```

```

denominator = Y4_Y3 * X2_X1 - X4_X3 * Y2_Y1;

% Gradients of parametric line between end points; if u_a and u_b are between 0
and 1, then there's a point of intersection. Using the known x and y values and
these gives the x,y position of the intersection point
u_a = numerator_a / denominator;
u_b = numerator_b / denominator;

% If INT_B = 0, no points of intersection. Otherwise, it's at INT_X, INT_Y
if (u_a >= 0) && (u_a <= 1) && (u_b >= 0) && (u_b <= 1)

    % Compute point of intersection and check whether it exists
    xint = X1+X2_X1*u_a;
    yint = Y1+Y2_Y1*u_b;

    % Find the corresponding height using linear interpolation (No error since
they're lines)
    z1 = polint(X1,X2,xint,Z1,Z2);
    z2 = polint(X3,X4,xint,Z3,Z4);
    zint = z2;

    % Cylinders are hard. This should give the vertical distance inside the
cylinder.
    fl1 = R1*cosd(atan2((Z2-Z1)/(sqrt((Y2-Y1)^2+(X2-X1)^2))));
    fl2 = R2*cosd(atan2((Z4-Z3)/(sqrt((Y4-Y3)^2+(X4-X3)^2))));

    % The centrelines cross when fibre 1 is moved downwards by the difference
between their z co-ordinates at the point of intersection.
    % However, fibres have a diameter, so we need to take into account how
much of the difference is contained within the fibres.
    distance = abs(z2-z1)-((fl1+fl2)*1.4);
else
    % No points of intersection ==> No min distance.
    distance = 0;
    xint = 0;
    yint = 0;
    zint = 0;
end
end

```

## C.5 polint.m

*Note that the following MATLAB code needs to be saved as 'polint.m' and located in a path that MATLAB can find. This is the code that finds the z co-ordinate where the two fibres meet based on linear interpolation along the fibre length (point-line-intercept).*

```

function z3 = polint(x1,x2,x3,z1,z2)
    d1 = x3-x2;

```

```

d2 = z2-z1;
d3 = x2-x1;
z3 = z2+(d1*d2)/d3;
end

```

## C.6 randb.m

*Note that the following MATLAB code needs to be saved as 'randb.m' and located in a path that MATLAB can find. This code uniformly generates a random number between two values.*

```

%Input two numbers and return a uniform random number between the two
numbers.
function z = randb(low,high)
    z = (high-low)*rand(1)+low;
end

```

## C.7 vecleng.m

*Note that the following MATLAB code needs to be saved as 'vecleng.m' and located in a path that MATLAB can find. This code calculates the length of a 2D line segment from the endpoints.*

```

% Find the length of the vector from (x1,y1) to (x2,y2).
function l = vecleng(x1,y1,z1,x2,y2,z2)
    l = sqrt((x1-x2).^2+(y1-y2).^2+(z1-z2).^2);
end

```

## C.8 vecleng3d.m

*Note that the following MATLAB code needs to be saved as 'vecleng3d.m' and located in a path that MATLAB can find. This code calculates the length of a 3D line segment from the endpoints.*

```

% Find the length of the vector from (x1,y1,z2) to (x2,y2,z2).
function l = vecleng3d(x)
    l = sqrt((x(1)-x(4)).^2+(x(2)-x(5)).^2+(x(3)-x(6)).^2);
end

```

## C.9 vcvf.m

*Note that the following MATLAB code needs to be saved as 'vcvf.m' and located in a path that MATLAB can find. This code calculates the volume of a fibre, taking into account whether it's cylindrical or sinusoidal in shape.*

```
function vf = vcvf(fibre,br)
    fibrevol = 0;
    for k = 1:size(fibre,1)
        if fibre(k,10) == 0
            fibrevol = (pi/4*(fibre(k,2)^2)*fibre(k,9))+fibrevol;
        else
            fibrevol =
pi*(fibre(k,10)/2)^2*8*sqrt((fibre(k,9)/8)^2+(fibre(k,2)/2)^2)+fibrevol;
        end
    end

    brvol = prod(br(4:6));

    vf = fibrevol/brvol*100;
end
```



## D General MATLAB Codes

### D.1 Temperature Dependant Air Thermal Conductivity

```
clc
clear all
%Thermal Conductivity of Air Vs Temperature
T_Lower=4; % Degrees C
T_Upper=80; % Degrees C
T=0:1:84;
k=((2.334*10^(-3)).*(T+273).^(-1.5))./(164.54+T+273.15);
plot(T,k)
hold on

plot([T_Lower,T_Lower],[0.018,0.034], '--r')
plot([T_Upper,T_Upper],[0.018,0.034], '--r')
title('Thermal Conductivity of Air')
xlabel('Temperature (C)', 'FontSize',14, 'fontname', 'Helvetica') %Create x label

ylabel('Instantaneous Thermal Conductivity
(W/mK)', 'FontSize',14, 'fontname', 'Helvetica') %create y label
%legend('Model of Fibres Oriented Parallel to Heat Direction', 'Model of Fibres
Oriented Perpendicular to Heat Direction') %create legend and content of legend
grid 'on' %add gridlines

set (gca, 'XTick', [-100:5:140])
%set (gca, 'YTick', [0:0.001,0.05])
K_Ave=(k(T_Upper)+k(T_Lower))./2;
```

### D.2 Boundary Layer Thickness

```
%Boundary Layer Thickness Natural Convection
clc
clear all
close all
x=0:5e-6:2e-4;
Gr=0.04;
Pr=0.711;
t=x*(3.93./sqrt(Pr))*(((0.952+Pr)./(Gr)).^0.25);
plot(x,t)
```

### D.3 Closed Variable Fibre Geometries

```
clc
clear all
close all

for i=1:10
A=2*(rand(1,15)-0.5); %note the normalised random number to prevent
clustering around y=Kx
h=10;
Theta=0:1/10:10;
```

```

x=(1/2.5)*(A(1,1)*cos(Theta)+A(1,2).*sin(Theta)+A(1,3).*sin(Theta).*cos(Theta)
+A(1,4).*sin(Theta).*sin(Theta)+A(1,5).*cos(Theta).*cos(Theta));
y=(1/2.5)*(A(1,6)*cos(Theta)+A(1,7).*sin(Theta)+A(1,8).*sin(Theta).*cos(Theta)
+A(1,9).*sin(Theta).*sin(Theta)+A(1,10).*cos(Theta).*cos(Theta));
z=(1/2.5)*(A(1,11).*cos(Theta)+A(1,12).*sin(Theta)+A(1,13).*sin(Theta).*cos(Theta)
+A(1,14).*sin(Theta).*sin(Theta)+A(1,15).*cos(Theta).*cos(Theta))
%z=cos(Theta)+sin(Theta)

plot(x,y)

grid on
hold on
end

xlabel('X Direction','FontSize',14,'fontname','Helvetica') %Create x label
ylabel('Y Direction','FontSize',14,'fontname','Helvetica') %create y label
zlabel('Z Direction','FontSize',14,'fontname','Helvetica') %create z label
grid 'on' %add gridlines
set(gca,'XTick',[-1:0.2:1])
set(gca,'YTick',[-1:0.2:1])
set(gca,'ZTick',[-1:0.2:1])

```

#### D.4 Thermal Models 1-4

%Thermal Conductivity Parallel Perpendicular Comparison

```

clear all
close all
clc

%V_Fibre=1;
%V_Total=10;

K_Air=0.0278;
K_Fibre=0.85;

%Alpha=(V_Fibre./V_Total);
Alpha=0:0.01:1;

K_Paralell=(Alpha*K_Fibre+(1-Alpha)*(K_Air));
K_Planar1=(K_Air*K_Fibre)/(K_Air*Alpha+K_Fibre*(1-Alpha));
K_Planar2=K_Air*((sqrt((1-Alpha).^2*((K_Fibre./K_Air)-1).^2+4*(K_Fibre/K_Air))-
(1-Alpha)*(K_Fibre./K_Air-1))./2);
K_Planar3=K_Air.*((1-Alpha)+(1+Alpha).*(K_Fibre./K_Air))./((1-
Alpha).*(K_Fibre./K_Air)+(1+Alpha));
%K_Mean=A*K_Paralell+B*K_Planar;

plot(Alpha,K_Paralell,Alpha,K_Planar1,Alpha,K_Planar2,Alpha,K_Planar3)
xlabel('Volumetric Ratio of Fibres to Total
Volume','FontSize',14,'fontname','Helvetica') %Create x label
ylabel('Mean Thermal Conductivity (W/mK)','FontSize',14,'fontname','Helvetica')
%create y label
legend('Theoretical Paralel Model', 'Theoretical Serise Model','Clayton
Model','Pilling Model') %create legend and content of legend
grid 'on' %add gridlines

```

```
%set (gca,'XTick',[0:0.1:1])
%set (gca,'YTick',[0:0.1:1])
hold on
```

## D.5 Comparison Between Conductivity Models for Conduction, Convection and Radiation

```
clc
clear all
close all

T_2=0:1:200;
T_1=74+T_2;
K_Air=0.027;
K_Fibre=0.85;

L=2*10^-4;
Sigma=5.67*10^-8;
Pr=0.711;
mu=1.9*10^-5;
T_Film=(T_1+T_2)./2;
B=1./(T_Film+273);
g=9.81;
Gr=(g.*B.*(T_1-T_2)*L.^3)./(mu.^2)
%Ra=Pr*Gr;
%Nu=(0.825+0.387*((Pr*Gr).^(1./6))./(((1+(0.492./Pr).^(9/16)).^(8/27)))));
g_Pr=(0.75*Pr.^(0.5))./((0.609+1.22*(Pr.^0.5)+1.238.*Pr).^0.25);
Nu=(4./3)*((Gr./4).^0.25)*g_Pr;
H=Nu.*(K_Air./L);
F=0.2;
e=0.75;

K_Conv=H*L;
K_Rad=((Sigma).*(pi*2*8e-6).*(T_1.^2+T_2.^4).*(T_1+T_2))/(2.*((1-e)./e)+1./F);

K_Tot=K_Air+K_Fibre+K_Conv+K_Rad;
K_Air_Norm=K_Air./K_Tot;
K_Fibre_Norm=K_Fibre./K_Tot;
K_Conv_Norm=K_Conv./K_Tot;
K_Rad_Norm=K_Rad./K_Tot;

%plot(T_2,K_Air_Norm,T_2,K_Fibre_Norm,T_2,K_Conv_Norm,T_2,K_Rad_Norm)
%plot(T_2,K_Air_Norm)
plot(T_2,(K_Air),T_2,(K_Fibre),T_2,(K_Conv),T_2,(K_Rad))
%plot(T_2,K_Conv)
xlabel('Upper Temperature (C)','FontSize',14,'fontname','Helvetica') %Create x
label
ylabel('Equivalent Conductivity (W/mK)','FontSize',14,'fontname','Helvetica')
%create y label
legend('Air, K Value', 'Fibre, K Value', 'Convection, K Value', 'Radiation, K Value')
%create legend and content of legend
set (gca,'YScale','log');
grid 'on' %add gridlines
%set (gca,'XTick',[0:1])
%set (gca,'YTick',[0:0.1:1])
```

## D.6 Convective Heat Transfer Correlation

```
%Correct Evaluation of HTC for convection
clc
clear all
close all
Gr=0.04;
Pr=0.711;
gpr=0.75.*(Pr.^1/2)./(0.609+1.221.*sqrt(Pr)+1.238.*Pr)
Nu=(4./3).*((Gr/4).^(1/4)).*gpr
```

## D.7 Mean free Path Mechanics Assessment

```
%conductivity as a function of mean free path
clc
clear all
close all

T=287:1:367;

K_T=((2.334.*10^-3).*T.^(1.5))./(164.54+T);
CP=1030.5-0.19975.*T+(3.9734.*10.^-4).*T.^2;
CV=CP./1.4;
Pr=0.7;

KB=1.38.*10.^-23;
DC=(0.8.*364+0.2.*346).*1e-12;

P_T=1.0125.*(CP-CV).*T;

L_MFP=KB.*T./(sqrt(2).*pi.*DC.*DC.*P_T);

Z=0.3;
D=8e-6;
f=0.1;
L_Col=Z.*(D/f);
Kn=L_MFP./L_Col;

K=K_T./(1+((4.*CP)./(CP+CV)).*(Kn./Pr));
plot(T,K_T,T,K)

title('Thermal Conductivity')
xlabel('Temperature (K)','FontSize',14,'fontname','Helvetica') %Create x label
ylabel('Thermal Conductivity (W/mK)','FontSize',14,'fontname','Helvetica')
%create y label
legend('Bulk Gas', 'Molecular Dynamics Prediction') %create legend and content
of legend
grid 'on' %add gridlines
set(gca,'XTick',[0:10:1000])
set(gca,'YTick',[0:0.001:3])
```

## D.8 3D open Nonlinear Fibres

**%Creation of Normalised Nonlinear Fibres**

```
clear all
close all
clc
format long
N=50; %
c=10

for c=1:N
TH=0:0.1:pi;
A=2*(rand(100,100)-0.5);
L=4

x=((2e-
4)./L)*(((A(1,1).*sin(TH)+A(2,1).*cos(TH)+A(3,1).*sin(TH).*cos(TH)+A(4,1).*(cos(T
H)).*(cos(TH))+A(5,1).*(sin(TH)).*(sin(TH))))+2.5);
x(x<0)=[];
x(x>2e-4)=[];
L_x=length(x);

y=((2e-
4)./L)*(((A(6,1).*sin(TH)+A(7,1).*cos(TH)+A(8,1).*sin(TH).*cos(TH)+A(9,1).*(cos(T
H)).*(cos(TH))+A(10,1).*(sin(TH)).*(sin(TH))))+2.5);
y(y>2e-4)=[];
y(y<0)=[];
L_y=length(y);

z=((2e-
4)./L)*(((A(11,1).*sin(TH)+A(12,1).*cos(TH)+A(13,1).*sin(TH).*cos(TH)+A(14,1).*(
cos(TH)).*(cos(TH))+A(15,1).*(sin(TH)).*(sin(TH))))+2.5);
z(z>2e-4)=[];
z(z<0)=[];
L_z=length(z);

C=[L_x,L_y,L_z];
K=min(C);

x=x(1,1:K);
y=y(1,1:K);
z=z(1,1:K);

%Volume of fibres
for i=1:K-1
F_Dia=8*(10^-6);
V_elements=[1:K];
L_elements=[1:K];

L_elements(i)=sqrt((x(i)-x(i+1))^2+(y(i)-y(i+1))^2+(z(i)-z(i+1))^2);
V_elements=(pi./4)*(F_Dia^2)*L_elements;
V_tot=sum(V_elements);
FIBRE_MATRIX(4,i,:)=(V_elements(i));
end
```

```

for i=1:K
Fibre_Matrix(1,i,c)=x(i);
Fibre_Matrix(2,i,c)=y(i);
Fibre_Matrix(3,i,c)=z(i);
Fibre_Matrix(4,i,c)=V_tot;
end

plot3(x,y,z)
hold on

end

xlabel('x Direction (m)','FontSize',14,'fontname','Helvetica') %Create x label
ylabel('y Direction (m)','FontSize',14,'fontname','Helvetica') %create y label
zlabel('z Direction (m)','FontSize',14,'fontname','Helvetica') %create y label
grid 'on' %add gridlines

% for i=1:K-1
%F_Dia=8*(10^-6);
%V_elements=[1:K];
%L_elements=[1:K];

%L_elements(i)=sqrt((x(i)-x(i+1))^2+(y(i)-y(i+1))^2+(z(i)-z(i+1))^2);
%V_elements=(pi./4)*(F_Dia^2)*L_elements;
%V_tot=sum(V_elements);
%FIBRE_MATRIX(4,i,:)=(V_elements(i));
%end

%FIBRE_MATRIX=zeros(4,K,N);

```

## D.9 Disorder Coefficient

```

clc
clear all
close all

TH=0:pi./1000:pi./2;
ALPHA=0.001;
KA=0.0285;
KF=0.85;
K_1=(KA.*KF)/(KA.*ALPHA+KF.*(1-ALPHA));
K_2=KF.*ALPHA+KA.*(1-ALPHA);
K=(1+((K_1./K_2)-1).*(sin(TH)).^2);
plot(TH,K)
hold on

TH=0:pi./1000:pi./2;
ALPHA=0.05;
KA=0.0285;
KF=0.85;
K_1=(KA.*KF)/(KA.*ALPHA+KF.*(1-ALPHA));
K_2=KF.*ALPHA+KA.*(1-ALPHA);
K=(1+((K_1./K_2)-1).*(sin(TH)).^2);
plot(TH,K)
hold on

```

```

TH=0:pi./1000:pi./2;
ALPHA=0.1;
KA=0.0285;
KF=0.85;
K_1=(KA.*KF)/(KA.*ALPHA+KF.*(1-ALPHA));
K_2=KF.*ALPHA+KA.*(1-ALPHA);
K=(1+((K_1./K_2)-1).*(sin(TH)).^2);
plot(TH,K)
hold on

TH=0:pi./1000:pi./2;
ALPHA=0.5;
KA=0.0285;
KF=0.85;
K_1=(KA.*KF)/(KA.*ALPHA+KF.*(1-ALPHA));
K_2=KF.*ALPHA+KA.*(1-ALPHA);
K=(1+((K_1./K_2)-1).*(sin(TH)).^2);
plot(TH,K)
hold on

TH=0:pi./1000:pi./2;
ALPHA=0.7;
KA=0.0285;
KF=0.85;
K_1=(KA.*KF)/(KA.*ALPHA+KF.*(1-ALPHA));
K_2=KF.*ALPHA+KA.*(1-ALPHA);
K=(1+((K_1./K_2)-1).*(sin(TH)).^2);
plot(TH,K)
hold on

TH=0:pi./1000:pi./2;
ALPHA=1;
KA=0.0285;
KF=0.85;
K_1=(KA.*KF)/(KA.*ALPHA+KF.*(1-ALPHA));
K_2=KF.*ALPHA+KA.*(1-ALPHA);
K=(1+((K_1./K_2)-1).*(sin(TH)).^2);
plot(TH,K)
hold on

title('Statistical Average Angle Correlation')
xlabel('Angle (Rads)','FontSize',14,'fontname','Helvetica') %Create x label
ylabel('Z Coefficient','FontSize',14,'fontname','Helvetica') %create y label
legend('Alpha=0.001', 'Alpha=0.05', 'Alpha=0.1', 'Alpha=0.5', 'Alpha=0.7',
'Alpha=1=0') %create legend and content of legend
grid 'on' %add gridlines
set(gca,'XTick',[0:0.1:2])
set(gca,'YTick',[0:0.1:3])
axis([0 1.6 0 1.1])

```

## D.10 Model 5

%Orientation Model

```

clc
clear all
close all
A=0:0.005:1;
K_Air=0.0242;
K_Fibre=0.85;
B=K_Air./K_Fibre;

for i=0:5
    Z=0.2*i;

K_Ave=K_Fibre*(1+((B-1)./(1+(A./(1-A)).*(1+Z.*((B-1)./(B+1))))));
plot(A,K_Ave)
hold on
end

%plot(T_2,K_Conv)
title('Orientation Dependant Model','FontSize',16,'fontname','Helvetica')
xlabel('Volume Fraction','FontSize',14,'fontname','Helvetica') %Create x label
ylabel('Equivalent Conductivity (W/mK)','FontSize',14,'fontname','Helvetica')
%create y label
legend('Z=0', 'Z=0.2', 'Z=0.4', 'Z=0.6', 'Z=0.8', 'Z=1') %create legend and
content of legend
grid 'on' %add gridlines
%set (gca,'XTick',[0:1])
%set (gca,'YTick',[0:0.1:1])

```

## D.11 Optimization of Radiation Model

```

clc
clear all
close all

x=0.000016:((2e-4)./500000):2e-4;
T_1=(4+(3.85*10^5).*x);

r=8e-6;
h=x./r;
F_1_2=(sqrt(h.^2-4)-h+2*asin(2./h))./(2*pi);
E=0.85;
sigma=5.67E-8;

q=(sigma).*((T_1+273).^4-(277).^4)./((1./F_1_2)+2.*((1-E)./E));

plot(x,q)
hold on

x=0:((2e-4)./500000):0.000016;
T_1=(4+(3.85*10^5).*x);
F_1_2=-1.9801e+04.*(x-0.000016)+0.18168219358181;
q=(sigma).*((T_1+273).^4-(277).^4)./((1./F_1_2)+2.*((1-E)./E));
plot(x,q)

title('Radiative Optimisation')

```



```

xlabel('Position Along Bounding Region (m)', 'FontSize', 14, 'fontname', 'Helvetica')
%Create x label
ylabel('Radiative Heat Flux (W/m^2)', 'FontSize', 14, 'fontname', 'Helvetica')
%create y label
legend('Tabulated View Factor', 'Interpolated View Factor') %create legend and
content of legend
grid 'on' %add gridlines
set (gca, 'XTick', [0:3e-4./20:3e-4])
set (gca, 'YTick', [0:0.5:10])

```

## D.12 A-B Hybrid Model

%Script Estimating the Thermal Conductivity System And Range Of Values

```

clear all
close all
clc
syms A B
[A,B] = meshgrid(0:0.1:1,0:0.1:1);

V_Fibre=1;
V_Total=10;

K_Air=0.0242;
K_Fibre=0.85;

Alpha=(V_Fibre./V_Total);

K_Paralell=(Alpha*K_Fibre+(1-Alpha)*(K_Air));
K_Planar=(K_Air*K_Fibre)./(K_Air*Alpha+K_Fibre*(1-Alpha));

K_Mean=A*K_Paralell+B*K_Planar;
contour(A,B,K_Mean,20);
[C,h]=contour(A,B,K_Mean,10);
clabel(C,h)
xlabel('Value of A', 'FontSize', 14, 'fontname', 'Helvetica') %Create x label
ylabel('Value of B', 'FontSize', 14, 'fontname', 'Helvetica') %create y label
grid 'on' %add gridlines
set (gca, 'XTick', [0:0.1:1])
set (gca, 'YTick', [0:0.1:1])

```

## D.13 Model using 3 Distinct Conductivities

%Triple model

```

A=0:0.01:1;

K_A=0.0245;
K_G=0.85;
K_P=0.15;

for B=0:5
    K_1=(K_P.*K_G)./((K_G.*0.2.*B)+K_P.*(1-0.2.*B));
    K=K_1.*K_A./(K_1.*A+(1-A).*K_A);
    plot(A,K)
end

```

```

    hold on
end
title('Model 2, Extension to 3 Materials')
xlabel('Volume Ratio (Air/total)','FontSize',14,'fontname','Helvetica') %Create x
label
ylabel('Thermal Conductivity (W/mK)','FontSize',14,'fontname','Helvetica')
%create y label
legend('B=0', 'B=0.2', 'B=0.4', 'B=0.6', 'B=0.8', 'B=1') %create legend and
content of legend
grid 'on' %add gridlines
set (gca,'XTick',[0:0.1:2])
set (gca,'YTick',[0:0.1:3])

```

## D.14 Stiffness of Hybrid Composition

```

clc
clear all
close all

```

```

B=0:0.01:1;
E_PET=4.3;
E_GF=64;
P=0.71;

```

```

for i=1:5
    A=0.1.*i
    E_B=A.*(E_PET.*(1-B)+(E_GF.*P).*B);
    semilogy(B,E_B)
hold on
end

```

```

xlabel('Composition','FontSize',14,'fontname',' Helvetica') %Create x label
ylabel('Youngs Modulus (MPa)','FontSize',14,'fontname','Helvetica') %create y
label
legend( '\alpha=0.1','\alpha=0.2','\alpha=0.3','\alpha=0.4','\alpha=0.5') %create
legend and content of legend
grid 'on' %add gridlines
%set (gca,'XTick',[0:3e-4./20:3e-4])
%set (gca,'YTick',[0:0.5:10])
axis([0,1,.43,25])

```

## **E APDL Codes**

### **E.1 Summing Fluxes in ANSYS**

```
RESUME  
SET, LAST  
inres, all  
etable, flux, tf, sum, avg  
etable, flux2, tf, x, avg  
etable, flux3, tf, y, avg  
etable, flux4, tf, z, avg  
ssum, flux
```

## F Extended Simulation Setup

Details View	
[-] Details of Box1	
Box	Box1
Base Plane	XYPlane
Operation	Add Material
Box Type	From Two Points
Point 1 Definition	Coordinates
<input type="checkbox"/> FD3, Point 1 X Coordinate	-1 m
<input type="checkbox"/> FD4, Point 1 Y Coordinate	-1 m
<input type="checkbox"/> FD5, Point 1 Z Coordinate	-1 m
Point 2 Definition	Coordinates
<input type="checkbox"/> FD6, Point 2 X Coordinate	1 m
<input type="checkbox"/> FD7, Point 2 Y Coordinate	1 m
<input type="checkbox"/> FD8, Point 2 Z Coordinate	1 m
As Thin/Surface?	No

Details View	
[-] Details of Import1	
Import	Import1
Source	D:\Hard Drive Documents\GFR\IGES\gf05.iges
Base Plane	XYPlane
Operation	Slice Material
Solid Bodies	Yes
Surface Bodies	No
Line Bodies	No
Simplify Geometry	No
Simplify Topology	No
Heal Bodies	Yes
Clean Bodies	Normal
Stitch Surfaces	Yes
Tolerance	Normal
Replace Missing Geometry	No
Target Bodies	All Bodies

Figure 14: The initial box generation settings (left) and the slice settings (right) for fibre importing

## G Particle Generation

### G.1 randomnesspheres.m

*The following MATLAB code is the initial algorithm used to generate spherical particles, and the script that should be run to generate a geometry. This implementation uses a normal distribution centred on a single radius value with a % variation (i.e, 10% variation means most of the particles are within 10% of the radius selected).*

```
% This script generates a set of randomly positioned spheres within a bounding
region
% The spheres shouldn't intersect with the edges of the bounding region, are
randomly distributed through it
%
% The script produces two outputs: 4sphere.csv, which allows the spheres to be
imported into APDL and a scatter plot showing the locations of the spheres
% Note that the script assumes (and tells APDL!) the sizes are in microns!

%rngstate = rng;

% Mean sphere diameter and % variation from it, 0 means uniform diameters
sphererad = [20 10];
sphereradstd = sphererad(1)*(sphererad(2)/100);

% bounding region, note it's assumed to be a cuboid from an (x,y,z) point to a
second (x,y,z) point
boundingregion = [0 0 0 10e-1 10e-1 10e-1];
dbx = boundingregion(4)-boundingregion(1);
dby = boundingregion(5)-boundingregion(2);
dbz = boundingregion(6)-boundingregion(3);

vbr = abs(dbx*dby*dbz);
% The % volume fraction aimed for, controls how long the script runs for
% Note that since each sphere has a volume, the script will generate a volume
fraction slightly
% higher than the desired value
% WARNING!!! If this is too high, the script may never complete, and higher
values --> longer run times
desiredvf = 5;

% Set the current volume fraction to 0 so the while loop runs and produces an
accurate volume fraction
actualvf = 0;
vsphere = 0;

% Set loop iteration counter to 1 so we know how many particles we've added
etc!
iter = 1;
```

```

sphere = zeros(1,5);
% Sphere generation loop!!!
while actualvf < desiredvf
    delete = 0;
    x = 1;
    % generate a temporary particle that we can do stuff to before deciding
    % whether to keep it or not!
    tempparticle = [iter,
randb(boundingregion(1),boundingregion(4)),randb(boundingregion(2),boundinr
egion(5)),randb(boundingregion(3),boundingregion(6)),
abs(sphereradstd*randn(1)+sphererad(1))];

    if abs(tempparticle(2) - boundingregion(1)) <= tempparticle(5)
        delete = 1;
    elseif abs(tempparticle(2) - boundingregion(4)) <= tempparticle(5)
        delete = 1;
    elseif abs(tempparticle(3) - boundingregion(2)) <= tempparticle(5)
        delete = 1;
    elseif abs(tempparticle(3) - boundingregion(5)) <= tempparticle(5)
        delete = 1;
    elseif abs(tempparticle(4) - boundingregion(3)) <= tempparticle(5)
        delete = 1;
    elseif abs(tempparticle(4) - boundingregion(6)) <= tempparticle(5)
        delete = 1;
    else
        while x <= size(sphere,1) && delete == 0
            dist =
vecleng(tempparticle(2),tempparticle(3),tempparticle(4),sphere(x,2),sphere(x,3),
sphere(x,4));
            if dist < tempparticle(5)+sphere(x,5)
                delete = 1;
            end
            x = x+1;
        end
    end

    if delete ~= 1
        sphere(iter,:) = tempparticle;
        vsphere = vsphere+4/3*pi*tempparticle(5)^3;
        actualvf = vsphere/vbr*100;
        iter = iter+1;
    end
end

%Open a file to write DesignModeller commands into; starts with a 4 just so it
shows up higher when alphabetically sorted
fid = fopen('1sphere.js','wt');

% End and clear any currently started simulations, and move into preprocessing.
fprintf(fid,'function plane1SketchesOnly(p){\n');
fprintf(fid,'p.Plane = agb.GetActivePlane();\n');
fprintf(fid,'p.Origin = p.Plane.GetOrigin();\n');
fprintf(fid,'p.XAxis = p.Plane.GetXAxis();\n');
fprintf(fid,'p.XAxis = p.Plane.GetYAxis();\n');

% Move the working plane to the centre of the sphere, and then create an
appropriately sized sphere centred there.

```

```

q = 1;
for k=1:size(sphere,1)
    % Make a new plane centered on the sphere with a predictable name
    fprintf(fid,'p.pl%i = agb.PlaneFromPlane(p.Plane);\n',k);
    fprintf(fid,'p.pl%i.AddTransform(agc.XformGlobalXOffset,%f);\n',k,sphere(k,2));
    fprintf(fid,'p.pl%i.AddTransform(agc.XformGlobalYOffset,%f);\n',k,sphere(k,3));

fprintf(fid,'p.pl%i.AddTransform(agc.XformGlobalZOffset,%f);\n\n',k,sphere(k,4));

    % Start a new sketch for the spline in the new plane and give it a generic
name
    fprintf(fid,'p.Sk%i = p.pl%i.newSketch();\n',q,k);
    fprintf(fid,'p.Sk%i.Name = "Sketch%i"\n',q,q);
    fprintf(fid,'with(p.Sk%i){\n',q);
    q = q+1;
    % Add a spline that gets revolved into a sphere
    fprintf(fid,'p.ea%i = ArcCtrEdge(0,0,%f,0,%f,0);\n}\n\n',k,-
sphere(k,5),sphere(k,5));
    % Start a new sketch for the line in the new plane and give it a generic name
    % <Doing both in the same sketch leads to the revolve failing
    fprintf(fid,'p.Sk%i = p.pl%i.newSketch();\n',q,k);
    fprintf(fid,'p.Sk%i.Name = "Sketch%i"\n',q,q);
    fprintf(fid,'with(p.Sk%i){\n',q);
    q = q+1;
    % Line to revolve the spline around!
    fprintf(fid,'p.l%i = Line(%f,0,%f,0);\n}\n\n',k,-sphere(k,5),sphere(k,5));
end
fprintf(fid,'p.Sk%i = p.Plane.newSketch();\n',q);
fprintf(fid,'p.Sk%i.Name = "Sketch%i"\n',q,q);
fprintf(fid,'with (p.Sk%i){\n',q);
fprintf(fid,'p.Line12 = Line(0,0,0,%f);\n',boundingregion(5));
fprintf(fid,'p.Line13 =
Line(0,%f,%f,%f);\n',boundingregion(5),boundingregion(4),boundingregion(5));
fprintf(fid,'p.Line14 =
Line(%f,%f,%f,0);\n',boundingregion(4),boundingregion(5),boundingregion(4));
fprintf(fid,'p.Line15 = Line(%f,0,0,0);\n}\n',boundingregion(4));

fprintf(fid,'return p;\n}\n');
fprintf(fid,'var XYPlane = agb.GetXYPlane();\n');
fprintf(fid,'agb.SetActivePlane(XYPlane);\n');
fprintf(fid,'var ps1 = plane1SketchesOnly(new Object());\n');
fprintf(fid,'var ext1 = agb.Extrude(agc.Add, ps1.Sk%i, agc.DirNormal,
agc.ExtentFixed, %f,agc.ExtentFixed, 0.0, agc.No, 0.0,
0.0);\n',q,boundingregion(6));

for k = 1:size(sphere,1)
    fprintf(fid,'var rev%i =
agb.Revolve(agc.Slice,ps1.Sk%i,ps1.l%i,agc.DirNormal,360,0,agc.Yes,0,0);\n',k,2*
k-1,k);
end
fprintf(fid,'agb.Regen();\n');

% Close the file
fclose('all');

fprintf('Actual volume fraction: %f\n',actualvf);

```

```

figure(1)
scatter3(sphere(:,2),sphere(:,3),sphere(:,4))

disp(datestr(now))

```

## G.2 randomspheresDM.m

*This is an example alternate generation code for the particle modelling. This version uses a more complex random number generation process to produce a particle distribution based on a multimodal distribution. As a result, it needs a matrix with two rows to define the particle radius: the first is the modal peaks and the second is the variance (square of the standard deviation).*

```

% This script generates a set of randomly positioned spheres within a bounding
region
% The spheres shouldn't intersect with the edges of the bounding region, are
randomly distributed through it
%
% The script produces two outputs: 4sphere.csv, which allows the spheres to be
imported into APDL and a scatter plot showing the locations of the spheres
% Note that the script assumes (and tells APDL!) the sizes are in microns!

% rngstate = rng;

% Mean sphere diameter in first row, variance in second row, min/max radius as
minr and maxr
sphererad = [16 32 32; 5 7 30];
minr = 5;
maxr = 78;

% spherediamrange = [50 100];
% spherestd = spherediamrange(1)*(spherediamrange(2)/100);

% bounding region, note it's assumed to be a cuboid from an (x,y,z) point to a
second (x,y,z) point
boundingregion = [0 0 0 600 600 600];
dbx = boundingregion(4)-boundingregion(1);
dby = boundingregion(5)-boundingregion(2);
dbz = boundingregion(6)-boundingregion(3);

vbr = abs(dbx*dby*dbz);
% The % volume fraction aimed for, controls how long the script runs for
% Note that since each sphere has a volume, the script will generate a volume
fraction slightly
% higher than the desired value
% WARNING!!! If this is too high, the script may never complete, and higher
values --> longer run times
desiredvf = 5;

```



```

% Set the current volume fraction to 0 so the while loop runs and produces an
accurate volume fraction
actualvf = 0;
vsphere = 0;

% Set loop iteration counter to 1 so we know how many particles we've added
etc!
iter = 1;
deleteinrow = 0;

sphere = zeros(1,5);
% Sphere generation loop!!!
while actualvf < desiredvf
    delete = 0;
    x = 1;
    radius = mmrn(sphererad(1,:),sphererad(2,:),1,'limit',[minr maxr]);
    % generate a temporary particle that we can do stuff to before deciding
whether to keep it or not!
    % tempparticle = [iter,
randb(boundingregion(1),boundingregion(4)),randb(boundingregion(2),boundinr
egion(5)),randb(boundingregion(3),boundingregion(6)),
abs(spherestd*randn(1)+spherediamrange(1))];
    tempparticle = [iter,
randb(boundingregion(1),boundingregion(4)),randb(boundingregion(2),boundinr
egion(5)),randb(boundingregion(3),boundingregion(6)), radius];

    if abs(tempparticle(2) - boundingregion(1)) <= tempparticle(5)
        delete = 1;
    elseif abs(tempparticle(2) - boundingregion(4)) <= tempparticle(5)
        delete = 1;
    elseif abs(tempparticle(3) - boundingregion(2)) <= tempparticle(5)
        delete = 1;
    elseif abs(tempparticle(3) - boundingregion(5)) <= tempparticle(5)
        delete = 1;
    elseif abs(tempparticle(4) - boundingregion(3)) <= tempparticle(5)
        delete = 1;
    elseif abs(tempparticle(4) - boundingregion(6)) <= tempparticle(5)
        delete = 1;
    else
        while x <= size(sphere,1) && delete == 0
            dist =
vecleng(tempparticle(2),tempparticle(3),tempparticle(4),sphere(x,2),sphere(x,3),
sphere(x,4));
            if dist < tempparticle(5)+sphere(x,5)
                delete = 1;
            end
            x = x+1;
        end
    end

    if delete ~= 1
        sphere(iter,:) = tempparticle;
        vsphere = vsphere+4/3*pi*tempparticle(5)^3;
        actualvf = vsphere/vbr*100;
        iter = iter+1;
        deleteinrow = 0;
    else

```

```

        deleteinrow = deleteinrow+1;
        mod(deleteinrow,1000) == 0 && fprintf('deleted: %i; current vf: %.2f;
spheres: %i\n',deleteinrow,actualvf,size(sphere,1));
    end
end

%Open a file to write DesignModeller commands into; starts with a 4 just so it
shows up higher when alphabetically sorted
fid = fopen('2sphereDM.js','wt');

% End and clear any currently started simulations, and move into preprocessing.
fprintf(fid,'function plane1SketchesOnly(p){\n');
fprintf(fid,'p.Plane = agb.GetActivePlane();\n');
fprintf(fid,'p.Origin = p.Plane.GetOrigin();\n');
fprintf(fid,'p.XAxis = p.Plane.GetXAxis();\n');
fprintf(fid,'p.YAxis = p.Plane.GetYAxis();\n');

% Move the working plane to the centre of the sphere, and then create an
appropriately sized sphere centred there.
q = 1;
for k=1:size(sphere,1)
    % Make a new plane centered on the sphere with a predictable name
    fprintf(fid,'p.pl%i = agb.PlaneFromPlane(p.Plane);\n',k);
    fprintf(fid,'p.pl%i.AddTransform(agc.XformGlobalXOffset,%f);\n',k,sphere(k,2));
    fprintf(fid,'p.pl%i.AddTransform(agc.XformGlobalYOffset,%f);\n',k,sphere(k,3));

    fprintf(fid,'p.pl%i.AddTransform(agc.XformGlobalZOffset,%f);\n\n',k,sphere(k,4));

    % Start a new sketch for the spline in the new plane and give it a generic
name
    fprintf(fid,'p.Sk%i = p.pl%i.newSketch();\n',q,k);
    fprintf(fid,'p.Sk%i.Name = "Sketch%i"\n',q,q);
    fprintf(fid,'with(p.Sk%i){\n',q);
    q = q+1;
    % Add a spline that gets revolved into a sphere
    fprintf(fid,'p.ea%i = ArcCtrEdge(0,0,%f,0,%f,0);\n}\n\n',k,-
sphere(k,5),sphere(k,5));
    % Start a new sketch for the line in the new plane and give it a generic name
    % <Doing both in the same sketch leads to the revolve failing
    fprintf(fid,'p.Sk%i = p.pl%i.newSketch();\n',q,k);
    fprintf(fid,'p.Sk%i.Name = "Sketch%i"\n',q,q);
    fprintf(fid,'with(p.Sk%i){\n',q);
    q = q+1;
    % Line to revolve the spline around!
    fprintf(fid,'p.l%i = Line(%f,0,%f,0);\n}\n\n',k,-sphere(k,5),sphere(k,5));
end
fprintf(fid,'p.Sk%i = p.Plane.newSketch();\n',q);
fprintf(fid,'p.Sk%i.Name = "Sketch%i"\n',q,q);
fprintf(fid,'with (p.Sk%i){\n',q);
fprintf(fid,'p.Line12 = Line(0,0,0,%f);\n',boundingregion(5));
fprintf(fid,'p.Line13 =
Line(0,%f,%f,%f);\n',boundingregion(5),boundingregion(4),boundingregion(5));
fprintf(fid,'p.Line14 =
Line(%f,%f,%f,0);\n',boundingregion(4),boundingregion(5),boundingregion(4));
fprintf(fid,'p.Line15 = Line(%f,0,0,0);\n}\n',boundingregion(4));

fprintf(fid,'return p;\n}\n');

```

```

fprintf(fid,'var XYPlane = agb.GetXYPlane();\n');
fprintf(fid,'agb.SetActivePlane(XYPlane);\n');
fprintf(fid,'var ps1 = plane1SketchesOnly(new Object());\n');
fprintf(fid,'var ext1 = agb.Extrude(agc.Add, ps1.Sk%i, agc.DirNormal,
agc.ExtentFixed, %f,agc.ExtentFixed, 0.0, agc.No, 0.0,
0.0);\n',q,boundingregion(6));

for k = 1:size(sphere,1)
    fprintf(fid,'var rev%i =
agb.Revolve(agc.Slice,ps1.Sk%i,ps1.l%i,agc.DirNormal,360,0,agc.Yes,0,0);\n',k,2*
k-1,k);
end
fprintf(fid,'agb.Regen();\n');

% Close the file
fclose('all');

fprintf('Actual volume fraction: %f\n',actualvf);

figure(1)
scatter3(sphere(:,2),sphere(:,3),sphere(:,4))

disp(datestr(now))

```

### G.3 mmrn.m

*Note that the following MATLAB code needs to be saved as 'mmrn.m' and located in a path that MATLAB can find. This is the code that the alternate radius distribution uses to generate a more complex set of random numbers.*

```

function [X,gmix] = mmrn(mus,vars,N,varargin)
% UMGRN = univariate multimodal Gaussian random number generator. The
% function generates random variables from a mixture of N Gaussian
% distribution.
%
% [X,gmix] = UMGRN(mus,vars,N,'opt1','val1',...)
%
% Input: - mus = 1xN vector of means,
%        - vars = 1xN vector of variances,
%        - N = the number of random numbers,
%
% Output: - X = 1xN vector of random numbers with its probability density
function is gmix.
%        - gmix = the probability density function, i.e.:
%              gmix(X) = 1/N * (G(X,mu(1),vars(1)) + G(X,mu(2),vars(2)) + ... +
%              G(X,mu(N),vars(N)))
%        where G(X,mu,s) is the Gaussian function with mean=mu and
variance=s.
%
% Optional arguments:
% - 'limit', [a b]. Defines data range of the output X.
%   Default is computed optimally to cover 99% of the p.d.f. gmix.
% - 'with_plot', 0 | 1. Plot also the p.d.f. and the generated random numbers.

```

```

% Default is 1.
%
% Author: Avan Suinesiaputra (avan.sp@gmail.com)
% Fadillah Tala (fadil.tala@gmail.com)

% default values
lim = [];
with_plot = 0;

% get optional arguments
for i=1:2:length(varargin)
    if( strcmpi(varargin{i},'limit') )
        lim = varargin{i+1};
    elseif( strcmpi(varargin{i},'with_plot') )
        with_plot = varargin{i+1};
    else
        error('Unknown plot. ');
    end
end

% calculate the optimal limit if lim is empty
if( isempty(lim) )
    lim = [min(mus-3*vars) max(3*vars+mus)];
else
    % check the limit
    if( lim(1)>=lim(2) )
        error('Error in the data_range argument. ');
    end
end

% [a,b] = lim
a = lim(1);
b = lim(2);

% check length of mus and vars
if( length(mus)~=length(vars) )
    error('Inequal length of mus and vars arguments. ');
end

% NOTE: Since the output is a sum of p.d.f, we must make sure that the output
gmix must be
% a density (integrates to one). Thus each individual Gaussian function is
weighted by
% 1/M, where M = the number of Gaussian function.

% build the mixture of Gaussian function
gstr = 'exp(-(x-(%f)).^2 / (2*(%f)^2)) / ((%f)*sqrt(2*pi))';
gmix = sprintf('1/%d * (0',length(mus));
for i=1:length(mus)
    gmix = strcat(gmix,+',',sprintf(gstr,mus(i),vars(i),vars(i)));
end
gmix = inline(strcat(gmix,')'),'x');

% THE REJECTION METHOD

% the inequality to be hold: f(x) <= c*g(x) forall x.

```

```

% we can determine c = maximum of gmix, i.e.
1/(min(vars)*length(mus)*sqrt(2*pi)),
%          g(x) = the uniform p.d.f., i.e. g(x) = 1 forall x, and
%          f(x) is the gmix(x), the mixture of gaussian p.d.f.

% calculate the constant c:
c = 1/(min(vars)*length(mus)*sqrt(2*pi));

X = zeros(1,N);
for i=1:N

    % generate xi until it is not rejected
    ok = 0;
    while( ~ok )

        xi = a + (b-a)*rand; % 1. Generate a random number X from uniform dist.
        r = c / gmix(xi);    % 2. Calculate ratio: r = c*g(x) / f(x)
        u = rand;           % 3. Generate a uniform random number

        ok = (u*r) < 1;    % 4. Accept if (u*r) < 1

    end

    % put xi in the output array X
    X(i) = xi;

end

% PLOT if with_plot=1
if( with_plot )
    figure;
    hold on;
    set(gca,'ColorOrder',circshift(get(gca,'ColorOrder'),-1));
    xi = a:0.01:b;
    g = inline(sprintf('exp(-(x-mu).^2 / (2*s^2)) /
((%d)*s*sqrt(2*pi))',length(mus)), 'x', 'mu', 's');
    for i=1:length(mus) Gx(i,:) = g(xi,mus(i),vars(i)); end
    plot(xi,Gx,'LineStyle',':');
    plot(xi,gmix(xi),'b');
    plot(X,zeros(size(X)),'b. ');
end

```

#### G.4 Other MATLAB Codes Required for Particle Generation

The particle generation algorithms also require access to the `veclen.m` code (Section C.7) and the `randb.m` code (Section C.6).

## H DCPD Results

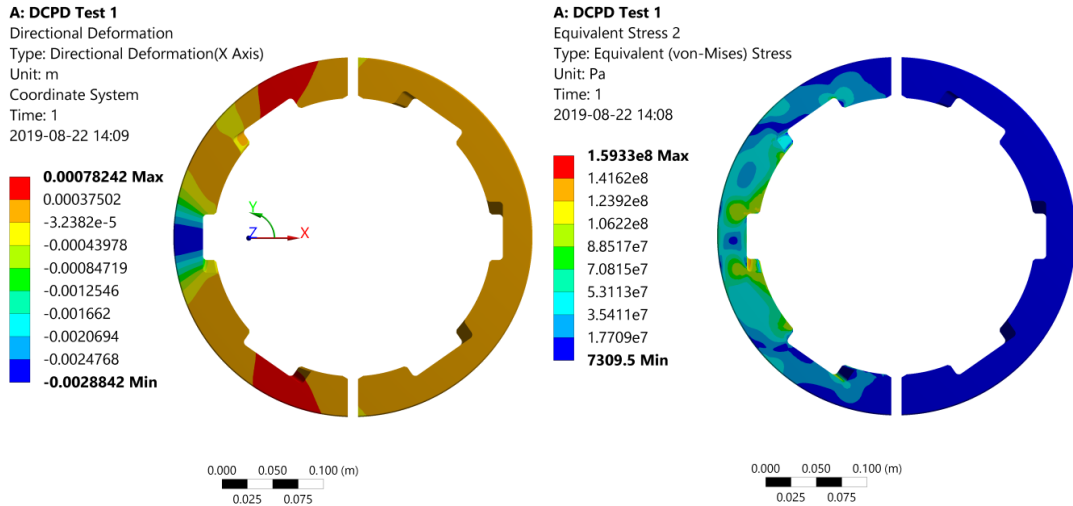


Figure 15: The deformation (left) and stress (right) in the DCPD centraliser in Test 1 loading conditions using the bilinear material model. Note that the deformation is negative because the x-direction is outwards.

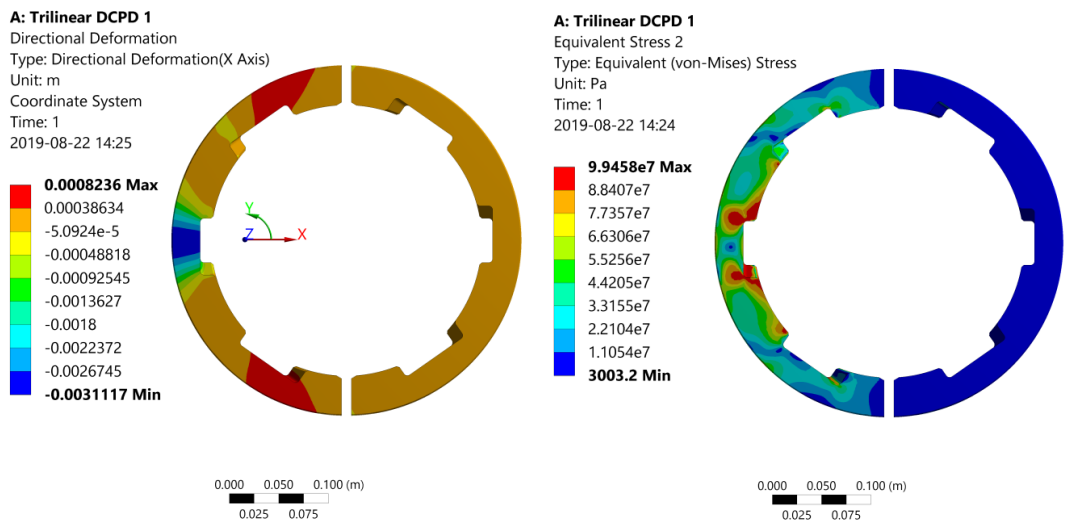


Figure 16: The deformation (left) and stress (right) in the DCPD under Test 1 loading conditions using the trilinear material model.

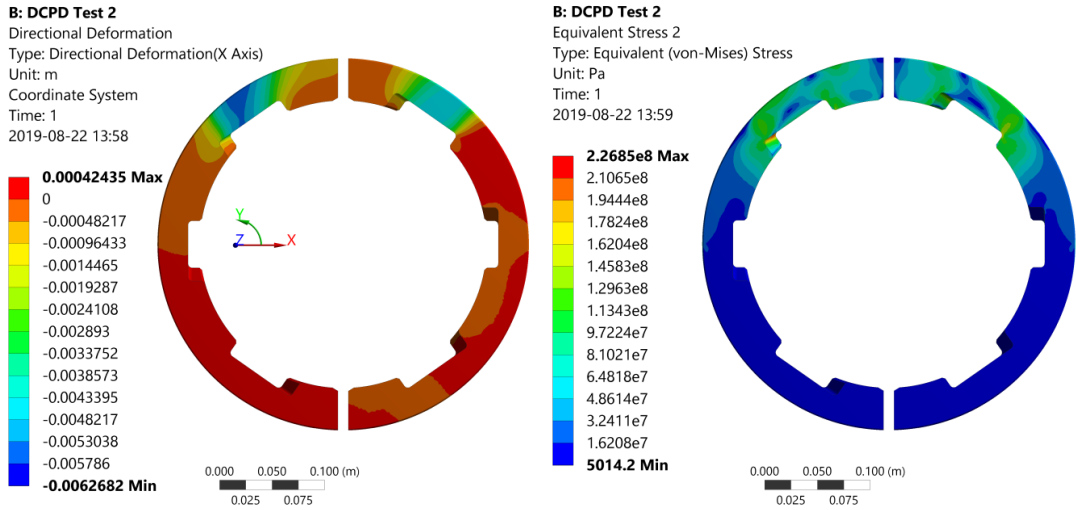


Figure 17: The deformation (left) and stress (right) for the DCPD in Test 2 Conditions using the bilinear material model.

# I Nylon Results




Figure 18: The deformation (left) and stress (right) for the Nylon in Test 1 Conditions using the multilinear material model.




Figure 19: The deformation (left) and stress (right) for the Nylon in Test 1 Conditions using the bilinear material model.



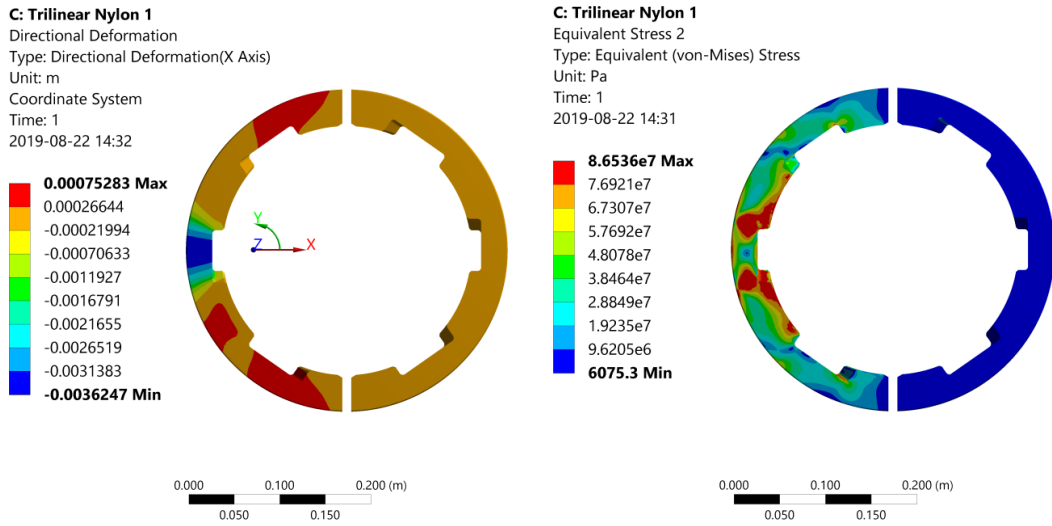


Figure 20: The deformation (left) and stress (right) for the Nylon in Test 1 Conditions using the trilinear material model.

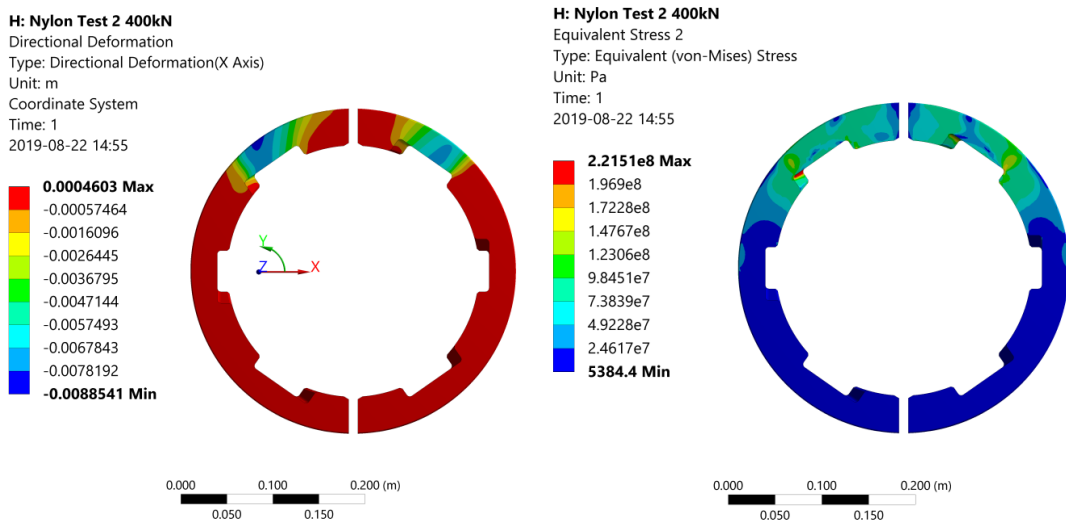



Figure 21: The deformation (right) and stress (left) in the Nylon under Test 2 load conditions using the multilinear material model.

**F: Bilinear Nylon Test 2**  
Directional Deformation 2  
Type: Directional Deformation(X Axis)  
Unit: m  
Coordinate System  
Time: 1  
2019-08-22 14:41



**F: Bilinear Nylon Test 2**  
Equivalent Stress 2  
Type: Equivalent (von-Mises) Stress  
Unit: Pa  
Time: 1  
2019-08-22 14:39




Figure 22: The deformation (left) and stress (right) for the Nylon in Test 2 Conditions using the bilinear material model.