

Process Characterisation of Out-of-Autoclave Towpreg Filament Winding for Composite Solid Rocket Motor Casings

Charles Kendall
4815971

*Submitted in partial fulfilment of the requirements for the degree of
Master of Science in Aerospace Engineering.*

To be defended publicly on 9th September, 2026.

Thesis committee:

Dr. Julie Teuwen	Chair
Dr.ir. Daniël Peeters	Supervisor
Dr. Roger Groves	Examiner



Dedicated to my girlfriend Denisy, my family, and my friends.

Particular thanks to Casper Blaszyk, Aashna Patel and Anna Stolk for their help with sample preparation.

Contents

List of Abbreviations	vi
List of Symbols	viii
1 Introduction	1
1.1 Manufacturing Problem	2
2 Literature Review	3
2.1 Application Context: Solid Rocket Motors	3
2.1.1 Composite Casings for Solid Rocket Motors	4
2.2 Pressure Vessel Architecture	5
2.2.1 End-Fitting and Polar Boss Interface	6
2.3 Fundamentals of Filament Winding	7
2.3.1 Prepreg and Towpreg Material Systems	9
2.4 Pressure Vessel Geometry	9
2.5 Autoclave and Out-of-Autoclave Processing	9
2.6 Process Parameters for Optimisation	10
2.6.1 Winding Tension and Speed	10
2.6.2 Cure Cycle	11
2.6.3 Consolidation Pressure	11
2.6.4 Liner Bonding and Surface Preparation	12
2.7 Testing Methods	12
2.7.1 Thermal Characterisation Methods	13
2.7.2 Determination of Fibre and Void Volume Fraction	13
2.7.3 Hoop Sample Testing	13
2.7.4 Beam Testing	14
2.7.5 Pressure Vessel Testing	14
2.8 Failure Prediction	14
2.8.1 Netting Analysis	14
2.8.2 Effect of Void Content on Mechanical Properties.	15
2.8.3 Progressive Failure and Pressure-vessel-specific Failure Modes.	15
2.8.4 Finite-Element Analysis (FEA) Approaches.	15
2.9 Summary and Identified Research Gaps	16
3 Research Aim, Case Study, and Scope	17
3.1 Case Study: Space-Shot Motor	17
3.2 Research Aim and Questions	18
3.3 Thesis Structure	19
4 Experimental Methodology and Campaign	20
4.1 Experimental Campaign Design	20
4.1.1 Candidate Process Parameters and Experimental Scope	20
4.1.2 Tranche 1: Pilot Work	21
4.1.3 Liner Baking Experiment	21
4.1.4 Tranche 2: Screening	21
4.1.5 Decision Gate Between Tranches 2 and 3	22

4.1.6	Tranche 3: Directed Study	22
4.2	Manufacturing	23
4.2.1	Filament Winding Machine	23
4.2.2	Specimen Design	23
4.2.3	Cure Conditions	24
4.3	Testing	24
4.3.1	Split-D Hoop Specimens	24
4.3.2	Strain Measurement	25
4.3.3	Mechanical Test Settings	25
4.4	Post-Processing	25
4.4.1	Optical Microscopy	25
4.4.2	Tensile-Data Reduction	26
5	Results	29
5.1	Reported Dataset and Interpretation Basis	29
5.2	Data Quality and Measurement Validation	29
5.2.1	Evidence Hierarchy and Traceability	29
5.2.2	Slack-Correction Diagnostics	31
5.2.3	Strain-Gauge Data Quality	31
5.2.4	Failure-Detection Sensitivity	32
5.2.5	Microscopy Audit and Representativeness	33
5.2.6	Area Normalisation	33
5.2.7	Statistical Validity	33
5.2.8	Permitted Interpretation	34
5.3	Liner Baking	34
5.4	Microscopy Results	34
5.4.1	Carbon Cross-Sectional Area	35
5.4.2	Void Content	35
5.4.3	Carbon-Phenolic Interface	38
5.5	Hoop Strength Results	38
5.5.1	Maximum Force	38
5.5.2	Terminal Capacity and Area Normalisation	40
5.5.3	Effect of Winding Tension	40
5.5.4	Effect of Shrink-Tape Consolidation	40
5.6	First Detected Damage	42
5.7	Stiffness	42
5.8	Mechanical Response vs Void Content	44
6	Discussion	46
6.1	Research Question 1: Microscopic Features	46
6.2	Research Question 2: Mechanical Response	46
6.2.1	Maximum Force and Area-Based Sensitivities	46
6.2.2	Apparent Stiffness and Gauge Response	47
6.2.3	First Detected Mechanical Discontinuity	47
6.3	Research Question 3: Implications for the Target Motor	47
6.4	Research Question 4: Test-Method Development	47
6.5	Research Question 5: Candidate Process Condition	48
6.6	Under-Resolved Relationships and Directions for Further Investigation	48
6.6.1	Tension Effect on Terminal Force	48
6.6.2	Tape Effect on First Detected Damage	48
6.6.3	Void Reduction at High Tension with Tape	48
6.6.4	Method-Limited Questions	48
6.7	Liner Drying and Blistering	49
6.8	Validity and Generalisability	49
7	Preliminary Motor-Design Implications	50
7.1	Pressure Line Loads	50
7.2	Netting Equilibrium	50
7.3	Force-Based Input and Area Sensitivities	51

7.4	Candidate Layups	51
7.5	Damage-Onset Scenario	51
7.6	Design Decision and Required Validation	52
8	Conclusions and Recommendations	53
A	Split-D Fixture Design	55
A.1	Assembly Drawing	55
A.2	Loading Analysis	55
B	Microscopy Results	57
B.1	Tranche 2 Sections	58
B.2	Tranche 3 Sections	60
C	Use of AI Tools	63
	Bibliography	64

Abstract

Composite casings for solid rocket motors offer significant mass savings over metallic alternatives, but their manufacture using Out-Of-Autoclave (OOA) filament winding processes remains significantly less mature than autoclave-based methods. This thesis investigates how the process parameters involved in the winding and curing operations affect the quality and hoop-direction mechanical response of OOA carbon-towpreg overwraps co-cured onto paper-phenolic liners.

The experimental campaign ran in three tranches: preliminary microscopy specimens, a broad screening tranche, and a directed low/medium/high-tension by tape/no-tape study informed by the results of the screening tranche. Split-D testing was completed for 49 coupons from Tranches 2 and 3, and optical microscopy on fifteen sections supplied the cross-sectional area, the detected void area fraction, the void-size distribution, and a descriptive interface metric.

Maximum force was treated as the primary terminal-capacity endpoint and proved comparatively stable across configurations. Nominal stress was treated only as its fixed rescaling. The stress per unit of measured cross-section varied materially because the cross-section itself varied between configurations, though no statistically significant relationship between processing and cross-sectional area could be resolved. Functional shrink tape at high tension reduced the detected optical void area fraction from approximately 12 % to 9 %, the most consistent processing signal observed in this project. It also produced a significantly improved interface appearance compared to the same configuration without tape, though it was found that this did not produce a statistically resolved improvement in ultimate force. The tape may also have slightly reduced the threshold for first damage to occur. Higher recorded tension was associated with a higher apparent system stiffness but a matching change in the intrinsic modulus of the overwrap was only suggestive. First detected force discontinuities were more sensitive to configuration than terminal force, which indicates that the onset of detectable damage responds to processing choices more strongly than the final load capacity does.

Early testing identified an issue with liner blistering. Drying a full-length phenolic liner removed 0.59 % of its initial mass, consistent with a removable volatile. However, the treatment did not prevent blistering in the directed tranche, indicating that an alternative material or a modified cure cycle may be required to avoid this defect.

The work identifies a promising process window and establishes which measurements are sensitive enough to detect manufacturing effects, but it does not establish robust statistical confidence in the observed trends. The results are therefore suitable for informing the next stage of development, but not for establishing design allowables or production readiness. The detected void area fractions of 9 % to 12 % are significantly higher than would typically be achieved with a standard autoclave processing technique, but it is expected that with the use of better manufacturing processes and equipment it may be possible to reduce the void content to acceptable levels for single-use rocket motors for unmanned systems, though not for human-rated systems or high-cycle fatigue environments.

List of Abbreviations

ABD Extensional–Coupling–Bending Stiffness Matrix.

ABS Acrylonitrile Butadiene Styrene.

AFP Automated Fibre Placement.

AGARD Advisory Group for Aerospace Research and Development.

ANOVA Analysis of Variance.

APCP Ammonium Perchlorate Composite Propellant.

ASTM ASTM International.

BOLE Booster Obsolescence and Life Extension.

CFRP Carbon Fibre Reinforced Plastic.

CLT Classical Lamination Theory.

CNC Computer Numerically Controlled.

COPV Composite Overwrapped Pressure Vessel.

CV Coefficient of Variation.

DARE Delft Aerospace Rocket Engineering.

DASML Delft Aerospace Structures and Materials Laboratory.

DIC Digital Image Correlation.

DSC Differential Scanning Calorimetry.

DTUL Deflection Temperature Under Load.

EUROC European Rocketry Challenge.

FEA Finite Element Analysis.

FPF First Ply Failure.

HALEX High Altitude Launch Experiments.

HDT Heat Deflection Temperature.

ILSS Interlaminar Shear Strength.

IPA Isopropyl Alcohol.

LPF Last Ply Failure.

MEOP Maximum Expected Operating Pressure.

NASA National Aeronautics and Space Administration.

NOL Naval Ordnance Laboratory.

OOA Out-Of-Autoclave.

PET Polyethylene Terephthalate.

PFCC Phenolic Formaldehyde Cotton Canvas.

PFCP Phenolic Formaldehyde Cellulose Paper.

RATO Rocket-Assisted Takeoff.

RGB Red–Green–Blue.

RTM Resin Transfer Moulding.

SD Standard Deviation.

SLS Space Launch System.

SRM Solid Rocket Motor.

TGA Thermogravimetric Analysis.

TMDSC Temperature-Modulated Differential Scanning Calorimetry.

US United States.

USA United States of America.

List of Symbols

- A_{eff} Effective measured carbon-envelope area of the two loaded split-D walls.
- $A_{\text{carbon,wall}}$ Measured carbon-envelope cross-sectional area of one split-D wall.
- F_{max} Maximum recorded force in a split-D test.
- I_{sp} Specific impulse.
- N_x Axial membrane line load.
- N_{θ} Circumferential membrane line load.
- P_c Compaction pressure applied normal to the laminate during winding.
- S Fibre-direction design stress used in the netting calculation.
- T Tow or roving winding tension.
- ΔV Velocity increment.
- α Wind angle.
- η^2 Proportion of variance explained by a factor (effect size).
- σ_{meas} Terminal force divided by the measured effective carbon-envelope area.
- σ_{nom} Terminal force divided by the fixed nominal split-D cross-sectional area.
- b Tow bandwidth.
- d Cohen's d ; standardised mean difference between two groups.
- g_0 Standard Gravity.
- p Unadjusted significance probability; in the motor-sizing calculation, internal pressure.
- q Benjamini–Hochberg-adjusted significance probability.
- r Radius in the winding and cylindrical-shell relations; throughout the statistical results, the Pearson correlation coefficient.
- t_h Combined thickness of the balanced helical plies.
- t_{90} Total thickness of the hoop plies.

Chapter 1

Introduction

Composite Overwrapped Pressure Vessels (COPVs) have been widely applied to a range of use-cases since the 1960s, especially where the structural mass of the pressure vessel must be minimised. This has included in scuba-diving gear, in rail tankers, and in aerospace systems [1]. In particular, COPVs have found wide use in rocketry for the production of high-efficiency Solid Rocket Motors (SRMs). These motors range from the very small — the four-inch diameter STAR 4 motor manufactured by Northrop Grumman, to the very large — the 3.7 m diameter Booster Obsolescence and Life Extension (BOLE) test article for the Space Launch System (SLS). These two motors may be compared in fig. 1.1 and fig. 1.2.



Figure 1.1: STAR 4 solid rocket motor, produced by Northrop Grumman [2]




Figure 1.2: BOLE test article [3]

Composite solid rocket motor casings are commonly reinforced with glass, aramid, carbon, or graphite fibres. At a motor level, published documents indicate a structural mass saving of at least 20 % relative to comparable metallic casings, with corresponding improvements to the propellant mass fraction and overall vehicle payload capacity [4].

These benefits are offset by the need for specialised tooling, controlled material storage and labour-intensive manufacture for all composite motors. In addition, the commonly-used autoclave cure cycle requires the use of a specialised composites autoclave to supply heat and consolidation pressure, the size of which limits the maximum dimensions of the parts that can be processed. Out-Of-Autoclave (OOA) processing instead uses an unpressurised oven and lower-cost consolidation methods, offering greater production scalability while making adequate

laminate consolidation more difficult. Figure 1.3 provides a general comparison of the two cure arrangements for a simple composite plate.




Figure 1.3: Comparison of autoclave and out-of-autoclave processes for the manufacture of a composite part

1.1 Manufacturing Problem

The OOA cure can be seen in fig. 1.3 to have a significantly lower consolidation pressure than the autoclave process. The manufacturing problem is, therefore, to consolidate the laminate before the resin loses its mobility, while allowing trapped air and cure-generated gases to escape without the part deforming, debonding, or otherwise failing. For OOA-cured filament-wound pressure vessels, tow tension and shrink-tape compaction provide practical means of influencing that process; their mechanisms and limitations are reviewed in sections 2.6.1 and 2.6.3.

This thesis reviews the relevant literature (chapter 2), defines the research questions and case study (chapter 3), and then presents the experimental work and its interpretation. The full chapter-by-chapter outline is given at the end of chapter 3.

Chapter 2

Literature Review

This chapter reviews the body of knowledge relevant to the manufacture of solid rocket motor (SRM) casings from Carbon Fibre Reinforced Plastic (CFRP) or other composite materials using out-of-autoclave (OOA) prepreg filament winding processes. The review begins with the application context (section 2.1), establishing the role and performance requirements of SRMs and their composite casings. It then surveys the current state of pressure-vessel architecture (section 2.2), winding fundamentals (section 2.3), and pressure vessel geometry (section 2.4). Autoclave and out-of-autoclave processing is compared (section 2.5) and process-parameter optimisation is discussed (section 2.6). A brief review of the characterisation and analytical methods used to judge the resulting structures is also performed (sections 2.7 and 2.8). The review ends by identifying the unresolved OOA manufacturing questions that motivate the research question and sub-questions in chapter 3.

2.1 Application Context: Solid Rocket Motors

Rocket motors are widely used in aerospace and defence applications. There are three primary types of rocket motor in widespread usage. A solid rocket motor stores its fuel and oxidiser mixed together as a solid mass, as opposed to a liquid rocket where both are in liquid form, or a hybrid in which one component (usually the oxidiser) is in liquid form and one component is in solid form. This means that an ordinary (non-thrust-vectoring) solid rocket motor has no moving parts and is highly reliable, as well as requiring little specialised equipment for operation. This has led solid rocket motors to gain extensive adoption for tactical and strategic missile systems, sounding rockets (rockets used for upper-atmospheric and microgravity research which do not achieve orbit) and Rocket-Assisted Takeoff (RATO) units for aircraft [5]. A typical solid rocket motor is illustrated in fig. 2.1.




Figure 2.1: Structure and layout of a typical solid rocket motor [6]

Tsiolkovsky's equation [7] is a major equation underpinning all classical rocket motors:

$$\Delta V = I_{sp} g_0 \ln \left(\frac{m_0}{m_1} \right) \quad (2.1)$$

In Tsiolkovsky's equation, m_0 is the initial mass of the vehicle, m_1 is the final mass, g_0 is standard gravity, and I_{sp} is the specific impulse of the propulsion system. It can be seen that the velocity increment, ΔV , depends directly on the ratio of initial to final mass. Therefore, with all else being equal, the greater the propellant mass as a proportion of total mass (and the lower the structural mass as a proportion of total mass), the greater the velocity increment that can be achieved. For rocket vehicles, the structural mass fraction is therefore a critical parameter for the performance of the overall system.

In a pump-fed liquid rocket, the large propellant tanks can operate near ambient pressure while only the comparatively small thrust chamber experiences the full combustion pressure. In an SRM, by contrast, the casing encloses the complete propellant grain and its entire internal surface experiences the full combustion pressure. As a result, SRMs typically have a higher structural mass fraction than liquid rockets and are especially sensitive to the structural efficiency of the casing.

Furthermore, the lack of fluid propellants means that the commonly used regenerative cooling system (where propellants are circulated through the engine and nozzle walls to provide cooling) is infeasible. Solid rockets therefore rely on ablative materials for thermal protection from the combustion process. Ablative materials are slowly consumed during operation, forming a char layer that insulates the underlying structure from the heat of combustion. The flame temperature of typical Ammonium Perchlorate Composite Propellants (APCPs) exceeds 2000 K, and so it is necessary to protect the rocket motor casing from the generated heat. The ablative protection is typically applied as a 'thermal liner': a non-structural layer that insulates the casing, and which is usually made from either phenolic resin composite or a modified rubber system. This material is slowly consumed over the course of the motor operation by charring and erosion processes. Rubber or phenolic ablatives are generally chosen as a result of their favourable thermal ablation properties and relatively low density [8]. Figure 2.2 shows this wall layout schematically.




Figure 2.2: Schematic of a cross-section of a solid rocket motor casing exposed to hot combustion gases, such as might occur after local propellant is entirely consumed. The liner chars and erodes, and the char front recedes rightwards into the virgin liner material. The CFRP casing is kept below the temperature at which the resin would degrade.

Historically, the majority of solid rocket motor casings have been manufactured from steel alloys or aluminium alloys [9]. However, there has been a significant and increasing usage of composite materials for the manufacture of high-performance rocket motors, as a result of the favourable specific structural properties of composites.

2.1.1 Composite Casings for Solid Rocket Motors

The use of composite materials for SRM casings is motivated primarily by the substantial structural mass savings they offer compared to metallic casings. The performance of a pressure vessel structure is conveniently characterised by the figure of merit PV/m , where P is the burst pressure, V the internal volume, and m the structural mass of the vessel. For a thin-walled cylindrical vessel, this is proportional to the specific strength of the casing material with a geometry-dependent proportionality factor. High-strength carbon fibre composite

systems offer specific tensile strengths significantly greater than high-strength steel alloys such as D6AC or maraging steel, the materials historically used for metallic SRM casings [5, 9]. These gains translate directly into improved system-level performance via the Tsiolkovsky relation. Table 2.1 compiles representative values.

Table 2.1: Representative longitudinal tensile strength, density, and specific strength of D6AC steel and a high-strength unidirectional carbon/epoxy lamina at 62 % fibre volume fraction

Material	Tensile strength (MPa)	Density (kg m^{-3})	Specific strength (kNm kg^{-1})
D6AC steel [10]	1931	7870	245
Carbon/epoxy (UD) [11]	2139	1600	1337

The development of composite casings for SRMs spans more than six decades, from the glass-epoxy filament-wound second stage of the Polaris A-2 fleet ballistic missile in the early 1960s [12] to the 3 m diameter carbon-epoxy P80FW first-stage motor of the Vega launch vehicle [13]. More recently, the BOLE programme for the National Aeronautics and Space Administration (NASA) SLS will replace steel-cased boosters with filament-wound carbon-fibre composite casings [14]. Lee notes that the shift from steel to composite casings has enabled payload capacity increases in excess of 20 % for equivalent motor envelopes, driven by the reduction in inert structural mass and the consequent improvement in the propellant mass fraction [4].

The design of a composite SRM casing requires the satisfaction of several concurrent requirements. Structurally, the casing must withstand the Maximum Expected Operating Pressure (MEOP) with an appropriate safety factor (giving a final failure pressure typically between 1.15 and 2.25 times MEOP), while maintaining integrity up to the burst pressure. The casing must also be able to withstand the aerodynamic and accelerational loads during flight, and often must also act at the mounting point for fins, interstage assemblies, and thrust vectoring systems. Thermally, the casing must be able to withstand the heat of combustion on the inside, and any potential aerodynamic heating on the outside, assisted by the use of appropriate thermal liners as already discussed. These mechanical, thermal, and interface requirements define constraints against which candidate material systems, laminate architectures, and manufacturing processes must be assessed. Process optimisation must therefore be directed toward satisfying these criteria while maintaining consistent manufacturing quality.

2.2 Pressure Vessel Architecture

Composite pressure vessels exploit the high axial strength and stiffness of unidirectional fibres by winding them along the principal stress directions of the vessel, producing a casing that is significantly stronger per unit mass than a comparable metallic construction. An example of this construction is shown in fig. 2.3.



Figure 2.3: Example COPV, showing the unidirectional carbon fibre bands around the vessel [15]

A solid rocket motor can be regarded as a specialised form of pressure vessel, and thus many of the techniques used in the design and manufacture of generic pressure vessels can also be applied to solid rocket motors.

There are currently five main types of pressure vessel, which are classified numerically [16]:

- I. All-metal construction
- II. Metal liner with a partial composite strengthening wrap
- III. Metal liner with a full composite wrap
- IV. Polymeric liner with a full composite wrap
- V. Linerless full-composite construction

Figure 2.4 distinguishes these architectures by the material and structural role of the liner.




Figure 2.4: Comparison of the structures of different types of pressure vessel [17]

Type IV and type V COPVs provide useful architectural analogues for composite SRM casings. Both carry pressure primarily through layers of unidirectional composite, but only type IV retains a polymer liner. Reported motor-level comparisons indicate mass savings of 20 % or more relative to metallic casing architectures when using a composite casing [4].

An ordinary type IV pressure vessel will commonly use polymeric liners made from polyethylene, polypropylene, Acrylonitrile Butadiene Styrene (ABS) or other thermoplastics [18], which add little mechanical strength to the system but provide an effective leak-proof barrier [19]. A solid rocket motor using a phenolic ablative liner may directly apply the same techniques used to manufacture a type IV pressure vessel, directly overwrapping the phenolic thermal liner with CFRP.

A linerless rocket-motor casing may be manufactured using the same removable-mandrel approach as a type V pressure vessel. If a polymeric thermal liner is subsequently applied, the completed motor is no longer strictly type V; only the load-bearing casing was manufactured by a type-V-like route. In this route, the liner and propellant grain are subsequently loaded into the cured casing. Alternatively, the fibres may be directly wrapped around a pre-manufactured grain assembly, which is useful for minimising residual stresses in the cured propellant grain, though this technique comes with obvious process safety tradeoffs from the use of a highly flammable or explosive mandrel in the winding process.

2.2.1 End-Fitting and Polar Boss Interface

Each dome of a filament-wound SRM casing incorporates a metallic end-fitting, commonly referred to as a polar boss, at the apex of the dome. The polar boss serves two distinct functions. Structurally, the dome closes the polar opening that is created by the standard filament-winding process, and provides a rigid annular surface against which the helical fibre bands terminate during winding. The polar boss also provides the mechanical interface between the composite casing and the other sub-systems of the motor: the forward boss provides the mounting location for the igniter (if present) and sometimes the connection to the rest of the vehicle (otherwise achieved with the use of a skirt that provides this interface), while the aft boss provides the attachment point for the nozzle assembly, thrust vector control system (if present), and sometimes the structural interface with any lower stages or systems.

The polar boss is typically manufactured from a high-strength steel alloy, or titanium or aluminium in small motors. The composite is wound over the boss flange, after which each helical band turns around the polar opening and continues back across the dome. The flange radius and contour determine the local fibre direction and contact area. Pressure load is transferred from the dome fibres into the boss through bearing, friction, and adhesive shear; inadequate contact area concentrates these stresses and can initiate interface failure [20, ch.9]. Figure 2.5 shows a cross-section through a typical polar boss.




Figure 2.5: Schematic of the polar boss region at the apex of a dome, reproduced from Meng et al. [21]. The figure represents one half of the cross section, revolved around a vertical axis to the left of the diagram.

The transition from the composite dome shell to the metallic boss is a critical structural region. The mismatch in stiffness between the carbon- or glass-based composite material and the metallic boss gives rise to interlaminar shear and peel stresses at the interface. The load transfer is primarily achieved by mechanically interlocking the boss and casing through the design of the boss geometry, and the adhesive bond formed between the composite and the boss. The quality of the bond between the casing and the boss is highly critical for the stresses in the casing. Vambol et al. estimated using Finite Element Analysis (FEA) that in cases without bonding, the stresses in the fibre direction can be up to 1.62 times as high as with an effective bond [22]. The shape of the boss is also critical: Vambol et al. also find that optimisation of the thickness, taper angle and flange projection could yield stress reductions of up to 50 %. Many casings also feature shear plies in the boss and skirt attachment construction. These are a thin layer of rubber or film adhesive that is used to minimise the stress discontinuities in the joint [20, ch. 9].

A composite SRM casing typically has unequal polar openings: the forward opening is usually significantly smaller than the aft opening since the nozzle throat is typically far larger than the igniter assembly. This means that the forward and aft domes must be constructed with different minimum winding angles, and have different load distributions. A discussion of composite pressure vessels with unequal openings is out of scope for the current review, since the use of non-geodesic windings to achieve varying wind angles introduces significant complexity both in design and manufacture. Advisory Group for Aerospace Research and Development (AGARD) Lecture Series 150 “Design methods in Solid Rocket Motors” details a specific design methodology for this design scenario [23].

Alternatively, one end of the casing may be left without a dome and the closure directly bonded, riveted or pinned in place; however, this is significantly less structurally efficient than a dual-dome design. The forward closure and the nozzle-to-casing joint are among the most structurally demanding regions of the motor, especially at the bondline between the boss and fibre overwrap. The connection from the boss to the nozzle or closure is in comparison fairly simple, usually comprising a bolted or pinned connection [20].

2.3 Fundamentals of Filament Winding

Filament winding is widely used in the manufacture of composite casings for solid rocket motors [1, 5, 20]. Automated Fibre Placement (AFP) is much less widely used than filament winding, since the overwhelming majority of solid rocket motors (and composite pressure vessels in general) are simple axisymmetric volumes of revolution, and the complexity of the AFP process tends to reserve the use of these machines to scenarios

requiring more complex geometries to be manufactured [24].

The process of filament winding requires the use of a mandrel, around which continuous fibres are wrapped using a specialised machine. The machine typically consists of a few distinct elements:

- A creel holder, which supports the rolls of fibre
- A tensioning system, which controls the tension of the fibres as they are wrapped around the mandrel
- A delivery head, which controls the fibre delivery angle and placement
- A mandrel support and rotation system

A simple arrangement of these components is illustrated in fig. 2.6.




Figure 2.6: Principal components and motions of a basic filament-winding machine

The modern filament-winding machine is Computer Numerically Controlled (CNC), and may be of either standard CNC or robotic arm design, as illustrated by the commercial systems in figs. 2.7a and 2.7b.



(a) Large CNC filament winder by Mikrosam [25]



(b) Robotic filament winding machine by Taniq [26]

Figure 2.7: Commercial filament-winding machine configurations

2.3.1 Prepreg and Towpreg Material Systems

The material system is central to the OOA filament winding process. A prepreg is a composite system that has been pre-impregnated with a partially-cured resin matrix, ready for lay-up without the need for additional matrix material to be added. Two main formats of prepreg are used with filament winding: towpreg and slit-tape prepreg.

Towpreg (a contraction of *tow* and *prepreg*) is manufactured by impregnating a continuous fibre tow with resin, without first calendering and manufacturing a sheet. The tow bundle is commonly referred to by the number of filaments making up a single tow: typically 6K, 12K or 24K (6, 12 and 24 thousand filaments, respectively). This value also loosely dictates the tow width, although several tows may be assembled during winding into one roving of much higher filament count for some applications. Figure 2.8 shows the manufacturing sequence: the tow is spread, impregnated with resin, B-staged (heated until partially cured), and wound onto a bobbin. The finished towpreg is stored under refrigerated conditions to limit further curing. When unwound for processing, the flattened profile of the towpreg can be adjusted by guide rollers and spreader bars to obtain the required tow width before deposition.




Figure 2.8: Stages of preparing a towpreg [27, fig. 1]

Slit-tape prepreg is significantly less common than towpreg in filament winding and is manufactured from a calendered unidirectional prepreg sheet that is cut to a specified nominal width. Szpoganicz et al. identify tighter width control as its principal advantage, while also noting its significantly increased cost relative to towpreg [28]. Slit-tape is normally associated with automated fibre placement and is less commonly used in filament winding.

Two main variants of the filament-winding process are commonly distinguished: wet winding, in which dry fibre is impregnated with liquid resin immediately before winding, and dry winding, in which pre-impregnated fibre such as towpreg is deposited without an in-line resin bath [1]. For thermoset towpreg, the resin is partially cured (B-staged) by the material manufacturer before winding. While wet winding offers a wider variety of resin/fibre combinations and typically lower material cost, it is significantly more challenging from the quality-control perspective: resin content depends on bath temperature, line speed and squeeze-roller pressure, and variation is typically larger than for towpreg [1].

2.4 Pressure Vessel Geometry

Pressure vessels typically comprise a cylindrical section with a dome at each end. Several dome profile options exist, including hemispherical, ellipsoidal and isotenoidal geometries.

Isotenoid domes — in which the fibre stress is theoretically uniform at every point along the meridian — are the mechanically optimal profile for filament-wound vessels and have the advantage that they can be constructed using purely geodesic winding paths, requiring no fibre steering and hence no reliance on friction to maintain path stability [29]. Non-geodesic isotenoidal domes are also possible, enabling the manufacture of designs with unequal polar openings [30]; however, these non-geodesic paths are held in place by friction, making them dependent on the tack and friction properties of the chosen towpreg system.

2.5 Autoclave and Out-of-Autoclave Processing

For filament-wound structures, OOA processing relies on winding tension and shrink-tape consolidation rather than the elevated external pressure available during autoclave cure. Under reduced-pressure processing, air

evacuation, resin flow and impregnation must be completed before gelation, changing the mechanisms governing void formation and the requirements placed on the cure cycle [31]. Figure 2.9 shows the application-specific arrangement used to consolidate a wound pressure vessel with an OOA process.




Figure 2.9: Out-of-autoclave processing of a filament-wound pressure vessel using heat-shrink tape (not to scale)

Void content is a principal separator between the two cure routes. Autoclave pressure helps to keep volatiles in solution and compresses voids that form, while an applied vacuum removes evolved gases [32]. Vacuum-bagged OOA processing is limited to approximately 1 atm of consolidation pressure. Shrink tape applies compaction directly to the wound laminate but does not itself provide negative pressure for removing volatiles and trapped gases.

This combination has made consistent aerospace-quality OOA manufacture difficult. With no elevated vessel pressure, entrapped gases must escape before the resin gels, requiring permeable evacuation pathways such as partially impregnated semipreg or modified towpreg architectures, or more complex double-vacuum-bag arrangements [33].

2.6 Process Parameters for Optimisation

In comparison to a metallic system, the properties of a composite system depend more heavily on the final-stage manufacturing process parameters used to produce the part: the same material applied using two different sets of parameters may vary widely in performance.

The principal process parameters for filament winding with towpreg that are amenable to optimisation are discussed in the following subsections: winding tension, winding speed, winding angle, curing temperature profile, and consolidation method.

2.6.1 Winding Tension and Speed

The recommended values for fibre tension vary widely between studies: Koussios proposes tension ranges of 1 N to 5 N for narrow rovings and 5 N to 15 N for broad rovings [34]. Although these categories are not explicitly defined, they likely correspond to 6–12K and 24K-or-larger tows, respectively. Koussios notes that the roving tension chosen seems to affect the coefficient of friction of the fibre only slightly, but that the tension has to be limited to avoid fibre failure in some cases.

Higher tensions can result in better consolidation and a lower void content, but at the cost of potential damage to the fibres and increased resin squeeze-out. Table 2.2, reproduced from Rosato and Grove's *Filament Winding: its development, manufacture, applications and design* [1], gives an overview of the effect for a range of prepreg resin contents and fibre tension values. The tension values reported in the table range from 26.7 N to 106.8 N. It is also stated by Rosato and Grove that for pre-impregnated fibres, the correct location to apply the tension is at the spool, while for wet-winding processes it is more appropriate to apply the tension to the fibres after the

impregnation process is completed. This ensures that the fibres can be properly impregnated for the wet-winding process, while the approach taken for dry fibres minimises the potential for damage to occur in the transfer between the fibre spool and the mandrel.

Table 2.2: Relationship of glass-roving machine tension to resin content and hoop tensile strength for a 76.2 mm-diameter, 1.524 mm-wall test ring. Values converted from the original United States (US) customary units [1].

Machine tension, N	Resin Content, % by Weight		Hoop tensile strength, MPa
	Prepreg roving before wrap	Cured ring	
26.7	18	18	1744
53.4	18	16	1793
80.1	18	15	1731
80.1	16	16	1793
80.1	22	19	1731
80.1	25	19	1489
106.8	18	15	1627

Pretensioning does not simply add a uniform beneficial stress. Each newly deposited layer compresses the layers beneath it, while relaxation, resin flow, and the changing winding radius redistribute the residual stresses through the thickness [1, 34]. The optimum therefore depends on tow or roving size, bandwidth, mandrel radius, resin state, and winding speed. Tension values cannot be transferred between studies without the accompanying tow count, bandwidth, and vessel radius; the applicable variable is the distributed tow tension and resulting compaction pressure rather than force alone.

Winding tension and speed are coupled process variables: line speed changes the time available for tow spreading, resin flow, and feedback control, while tension changes compaction and the risk of fibre damage or resin squeeze-out. The optimum must therefore be established for the actual towpreg format and machine rather than imported as a force–speed pair from a different winding system.

2.6.2 Cure Cycle

The cure cycle defines the temperature–time profile during thermoset crosslinking and determines the degree of cure, residual stress state, glass transition temperature, and void content of the cured part. For OOA filament-wound casings, all four outcomes are strongly sensitive to cure cycle design. There have been many investigations of specific cure cycle modifications, including those of Kim and Nutt [35], Farooq et al. [36], and González [37], which between them varied out-time and moisture conditioning, dwell temperature and cycle duration: Kim and Nutt showed that a flow-optimised, rapidly heated OOA cycle limits flow-induced porosity and extends the usable out-life of the prepreg by 175 %; Farooq et al. varied the first-dwell temperature (120 °C to 180 °C) of a thermoplastic-toughened epoxy and found the curing history governs the interphase morphology; and González shortened the manufacturer-recommended cycles by up to 72 % while keeping mechanical properties and void content within 4 % of the reference cure. These studies are, however, all highly specific to the resin system being investigated, and it is difficult to draw broad conclusions from them.

2.6.3 Consolidation Pressure

Consolidation pressure is the compressive stress applied perpendicularly to the laminate surface during cure; it is the single parameter that most distinguishes autoclave from OOA processing. Cure in an autoclave at 0.3 MPa to 0.7 MPa of gauge pressure reliably produces void contents below 0.5 % [31], whereas OOA processing must achieve comparable quality with significantly lower consolidation pressures.

Winding tension generates a transient, local compaction pressure as each tow is deposited: for a tow of width b wound at tension T over a cylinder of radius r , the instantaneous compaction pressure is $P_c = T/(r \cdot b)$. For typical values (50 N tension, 100 mm radius, 5 mm bandwidth), this yields approximately 0.1 MPa—comparable to vacuum-bag-only processing, but this pressure is not necessarily maintained through the whole cure.

Shrink-tape consolidation may be used to (partially) address this limitation. A heat-shrinkable polymer film (typically biaxially oriented Polyethylene Terephthalate (PET) or polyimide) is spirally wrapped over the wound mandrel prior to cure. The manufacturer’s free-shrink specification may be of order 5–20%, but a tape constrained by the rigid mandrel and composite overwrap cannot realise that contraction. The prevented thermal recovery

instead generates tensile stress in the tape and a corresponding radial compaction pressure. Bianchi et al. found a 45 % reduction in mean sample thickness when using shrink-tape compared to the uncompacted case, along with an increase in the specific strength of their samples from 35.24 MPa g⁻¹ to 121.66 MPa g⁻¹ [38]. However, any systematic quantification of shrink-tape compaction pressure as a function of tape specifications, tape application parameters, and geometry remains limited in the open literature.

The radial compaction pressure from shrink tape wound circumferentially can be estimated by analogy with a thin-walled pressure vessel in reverse. For a tape of width w with shrink force per unit length t wound with n effective layers over a cylinder of radius r , the force per unit axial length of mandrel is $n t$, giving:

$$P_{\text{tape}} = \frac{n t}{r} \quad (2.2)$$

For a representative PET shrink tape with a maximum shrink force of 1.50 N mm⁻¹ [39], wound with 50 % overlap (giving $n = 2$ effective layers) over a vessel of 120 mm diameter ($r = 60$ mm), eq. (2.2) yields $P_{\text{tape}} \approx 0.050$ MPa, which is roughly half the consolidation pressure achievable with vacuum bagging alone.

2.6.4 Liner Bonding and Surface Preparation

It is also well known that surface roughness and cleanliness have a significant impact on the performance of bonds between composite parts [40]. The bond between the phenolic liner and the composite casing is effectively a co-bonded connection. Yudhanto et al. investigated a range of surface preparation techniques for secondarily bonded connections, the results of which are reproduced in fig. 2.10. The principal result in fig. 2.10 is that no single preparation method maximises every joint metric: methods that increase roughness and mechanical interlocking can improve strength, but contamination removal and surface chemistry remain equally important. It is also clear that significant variance can be expected even for the same preparation method.




Figure 2.10: Effect of surface preparation on relative joint properties [41]

2.7 Testing Methods

The following subsections review the characterisation and testing methods used to evaluate the quality and mechanical performance of filament-wound composite structures. Methods are grouped by their primary purpose: thermal and cure characterisation, compositional analysis, microstructural examination, and mechanical performance evaluation.

2.7.1 Thermal Characterisation Methods

Mulligan et al. compare five complementary thermal-characterisation methods [42]:

Differential Scanning Calorimetry (DSC) measures differential heat-flow rate and can quantify residual cure enthalpy from a small sample, but degree of cure is inferred by comparison with a reference reaction enthalpy and is sensitive to the sample's thermal history.

Temperature-Modulated Differential Scanning Calorimetry (TMDSC) resolves total heat flow into reversing and non-reversing components, helping distinguish heat-capacity changes from kinetic events such as residual cure, although the components do not always correspond uniquely to individual physical processes and the method requires longer, more complex experiments.

Dynamic Mechanical Analysis measures storage and loss moduli and damping under oscillatory loading. It is sensitive to glass transition and temperature-dependent viscoelastic behaviour, but its results depend on specimen preparation, loading mode, oscillation amplitude, frequency, and temperature programme.

Dielectric Analysis measures changes in dielectric response associated with ion and dipole mobility and can therefore monitor cure in situ. Relating this response quantitatively to cure state is material-dependent and requires electrodes in contact with, or embedded near, the resin.

Deflection Temperature Under Load (DTUL) provides a single-point measure of a material's temperature response, taken as the temperature at which a loaded cantilever specimen deflects by a defined strain level. The metric is also known as the heat deflection temperature (Heat Deflection Temperature (HDT)).

2.7.2 Determination of Fibre and Void Volume Fraction

Two destructive methods are in common use for establishing the composition of a cured composite, and with it the void volume fraction.

In Thermogravimetric Analysis (TGA), the sample is heated to burn off the matrix, and the fibre fraction is obtained from the residual mass; Grund et al. discuss the use of this method for fibre-fraction analysis [43].

Acid digestion is used, similarly to TGA, to establish the fibre volume fraction of the composite [44]. The mass of the sample is first obtained before a strong acid such as hydrochloric acid, sulphuric acid or piranha solution is used to dissolve the matrix of the composite. By measuring the residual mass of the fibres, the fibre mass fraction and resin mass fraction can be obtained. The theoretical densities of the two components can then be used to convert to volume fractions and estimate the void content of the composite.

Neither method indicates the location of the voids within the cured part, and so microscopy is required for this purpose. Cilley indicates that resin removal (e.g. by acid digestion or burnout) is the preferred method for basic void fraction estimation, but that optical microscopy provides the additional benefit of characterising the distribution of voids within the laminate [45]. Furthermore, microscopy also allows the classification of voids according to size and geometry. Its disadvantages are that specimen cutting, mounting, polishing, imaging, and threshold selection are destructive and labour intensive, and that a two-dimensional field of view may not represent a heterogeneous three-dimensional laminate. Bodaghi et al. use microscopy for this purpose in Resin Transfer Moulding (RTM) composites, using MATLAB code to classify voids automatically [46], while Machado et al. propose a machine learning approach to improve void identification as a result of difficulties with traditional coding methods for void detection in composites [47].

The mechanical consequences of void content are discussed in the failure-prediction review below (section 2.8.2).

2.7.3 Hoop Sample Testing

A hoop sample is a short ring wound circumferentially at a winding angle of 90° to the vessel axis, so that its fibres carry the circumferential (hoop) stress produced by internal pressure. Historically, testing of filament-wound composites has extensively used hoop samples for characterisation of material performance. This technique was developed at the Naval Ordnance Laboratory (NOL) in the United States of America (USA), and thus these samples are commonly referred to as "NOL rings". These samples are favoured for their simplicity, low price and ability to yield reproducible data [1, ch.8]. However, since these samples are hoop-wound and loaded in the fibre direction, the failure of the samples is inherently in a fibre-dominated mode. This may limit the utility of hoop samples for the evaluation of matrix properties, where shear-load testing would be more appropriate.

The split-disk test (ASTM International (ASTM) D2290 [48]), referred to throughout this thesis as the split-D test,

measures the hoop tensile strength of a filament-wound ring by splitting a ring specimen across its diameter using a split mandrel setup. The split-D test introduces a bending component at the loading pin locations, which results in an artificially reduced value of hoop strength to be obtained compared to the values obtained from methods that evenly load the hoop. The split-D test provides a direct measure of the hoop tensile performance of the filament-wound laminate, including the effects of fibre alignment, void content, and resin quality on the dominant load-carrying direction of a pressure vessel [20]; the test geometry is shown in fig. 2.11.




Figure 2.11: Schematic of the ASTM D2290 [48] split-D test. Opposing semicircular fixtures separate under tensile loading and place the ring primarily in hoop tension, though local bending and contact stresses arise near the fixture gaps.

2.7.4 Beam Testing

The short-beam shear test (Interlaminar Shear Strength (ILSS), ASTM D2344 [49]) subjects a short-span flexural beam specimen to a three-point loading configuration designed to promote interlaminar shear failure before flexural failure. The apparent ILSS is computed from the peak load and specimen geometry. The test is straightforward, requires small coupon specimens that can be cut from filament-wound rings or panels, and is particularly sensitive to void content and interfacial quality. It is the most commonly reported property in void-content-mechanical-property knockdown studies (see section 2.8).

2.7.5 Pressure Vessel Testing

Because the failure prediction of composite pressure vessels is complex and FEA methods are insufficient, proof testing may be used to screen for defects. However, this does not guarantee future performance [19]. In testing, a test vessel is hydrostatically pressurised beyond its normal operating pressure or to failure, and the failure pressure can then be compared to the MEOP of the system to verify that the vessel meets the required safety factor. Strain gauges or Digital Image Correlation (DIC) can be used to evaluate the strain state of the vessel during the pressurisation. The resulting values can then be correlated with analytical or FEA models.

2.8 Failure Prediction

There are many methods used to predict the failure of composite structures. This section reviews a few of these methods, ranging from simple netting analysis to FEA approaches.

2.8.1 Netting Analysis

For composite pressure vessels, netting analysis is commonly used for preliminary analysis. For this technique, the strength and performance of the matrix are neglected. Roylance gives a derivation of the basic equations [50].

These equations yield the well-known relation that for an idealised thin-wall helically wound pressure vessel with an even internal pressure loading, the ideal wind angle is given by eq. (2.3), which resolves to a wind angle of 54.74°.

$$\tan \alpha = \sqrt{2} \quad (2.3)$$

However, the netting method obviously cannot predict matrix failure, and therefore methods based on classical laminate theory are common.

Nettles presents a full derivation of the equations for classical laminate theory, including the effects of temperature and moisture on the material and the relations between load, stress and strain via the Extensional–Coupling–Bending Stiffness Matrix (ABD) system [51]. The classical laminate theory is then coupled with a failure model to predict the actual performance of the laminate. The most commonly used models include the Hashin model and the Puck model, but there are many conflicting models available, with París and Jackson reporting strong debate over the most physically accurate models [52].

2.8.2 Effect of Void Content on Mechanical Properties.

The relationship between void content and mechanical property degradation is well established for fibre-reinforced epoxy composites. Judd and Wright established that the ILSS decreases by approximately 7 % for each volume percent of voids, up to at least 4 % void content, and that this relationship is approximately independent of fibre type, resin system, or fibre surface treatment [53]. Ghiorse subsequently characterised the effect in carbon/epoxy laminates specifically, reporting that each 1 % increase in void content reduced ILSS by 9.7 % and flexural strength by 10.3 % [54]. Boey and Lye report that “A reduction in void content from 40 % to 10 % increased the flexural strength by nearly three times, and almost doubled the modulus” [55]. Compressive strength and transverse tensile strength are affected to a similar or greater degree, as all three are matrix-dominated properties whose failure is sensitive to stress concentrations around void boundaries. Fibre-dominated properties such as longitudinal tensile strength are less severely affected by moderate void contents.

2.8.3 Progressive Failure and Pressure-vessel-specific Failure Modes.

Classical laminate theory predicts first-ply failure (First Ply Failure (FPF)) — the load at which the weakest ply orientation first reaches its failure stress according to the selected criterion (Hashin, Puck, or Tsai–Wu). For internally pressurised filament-wound vessels, FPF typically occurs in the matrix-dominated transverse or shear mode, manifesting as matrix microcracking. In composite pressure vessels, matrix cracking is not catastrophic: the fibres continue to carry load beyond FPF, and the laminate can sustain additional pressure until Last Ply Failure (LPF), at which point fibre failure initiates a progressive, unstable burst. Progressive failure modelling, in which the stiffness contribution of each failed ply is degraded after FPF and the load is redistributed to the remaining plies, is necessary to predict the burst pressure accurately [56].

The *leak-before-burst* criterion, applicable to composite pressure vessels with a gas-permeable liner or no liner, requires that through-wall matrix cracking (permitting leakage of pressurised gas) is detectable before the burst pressure is reached. This provides a safety benefit by converting a potentially catastrophic burst event into a detectable leak, analogous to the “safe-life” philosophy in metallic pressure vessel design (though this does not eliminate the need for monitoring of the COPV) [19]. For SRM casings with bonded internal insulation, through-wall leakage of hot combustion gases rather than inert pressurising gas is the failure mode of concern, and the leak-before-burst criterion must be evaluated against the thermal and chemical aggression of the combustion environment.

2.8.4 Finite-Element Analysis (FEA) Approaches.

Finite element analysis of filament-wound vessel burst has been demonstrated by Thesken et al. [57] and others using axisymmetric shell elements with angle-dependent orthotropic ply properties derived from Classical Lamination Theory (CLT). Such homogenised orthotropic plies do not resolve shear transfer from matrix to individual filaments. Their predictions are therefore most credible for smooth, continuous lay-ups dominated by in-plane fibre loading and become less reliable near ply terminations, boss interfaces, local waviness, and strongly matrix-dominated load paths. Dome and polar boss regions require more refined modelling to capture the variable thickness and winding angle distributions.

2.9 Summary and Identified Research Gaps

This chapter has reviewed the relevant literature on composite SRM casing design and manufacture, with emphasis on the material system, process parameters, and structural performance criteria pertinent to prepreg filament-wound OOA structures.

The structural case for composite SRM casings over metallic alternatives is well established: specific strength advantages of carbon-epoxy over high-strength steel yield mass savings of 20% or more at the motor level, as demonstrated by programmes from Polaris to the Vega P80FW [12, 13]. The mechanics of filament-wound pressure vessels — dome geometry, geodesic winding paths, netting analysis, and classical laminate theory — are comprehensively treated in the literature, principally by Peters [11, 20], Koussios [34], and Nettles [51]. The polar boss interface and its stress-concentration characteristics have received increasing attention in the finite element literature [22], though experimental validation for SRM-scale geometries remains limited and it appears that empirical relations are still often used in combination with FEA [20].

The central unresolved issue is the reliable manufacture of OOA towpreg-wound structures at the low consolidation pressures available outside an autoclave. Published studies quantify individual effects of cure history, winding tension, and shrink-tape compaction, but their process windows are material-, geometry-, and equipment-specific. They do not establish how winding tension and shrink-tape consolidation interact for a carbon-towpreg overwrap co-cured directly onto a phenolic liner, nor how those choices jointly affect void content, interface quality, stiffness, damage initiation, and ultimate hoop load. This gap leads directly to the research questions in chapter 3.

Chapter 3

Research Aim, Case Study, and Scope

Having identified the manufacturing problem in the literature review, this chapter defines the industrial case study and the proposed motor construction. These are presented as the research aim and questions.

3.1 Case Study: Space-Shot Motor

The test case for the study in this document is a motor currently under development by the High Altitude Launch Experiments (HALEX) team at the student rocketry society Delft Aerospace Rocket Engineering (DARE). Named “Shearwater”, this motor follows a fairly common overall design for sounding rocket and missile motors of this size, using a paper-phenolic liner directly overwrapped with a carbon fibre casing. The motor is intended to be used as a booster for a small dart-type sounding rocket. While the DARE motor is only to be produced in extremely small quantities, the motor design is representative of a class of motors that are produced in larger quantities for both sounding rocket and missile applications. Therefore it is of interest to determine whether the manufacturing process can be optimised to produce motors of this type with sufficient quality and repeatability for industrial production. The “Shearwater” motor has the parameters given in table 3.1.

Table 3.1: Documented target geometry and operating pressure used for the motor-level interpretation

Property	Value
Liner Thickness	2 mm
Liner Inner Diameter	116 mm
Liner Outer Diameter	120 mm
MEOP	11 MPa

The “Shearwater” motor is the follow-on to the “Puffin” rocket motor that was used in DARE’s entry to the European Rocketry Challenge (EUROC) in 2024. This motor had an overall length of 1630 mm, a diameter of 107 mm, around 16 kNs of total impulse, and a propellant mass fraction of about 75 %. A size comparison of the two motors is given in fig. 3.1.

The composite casing of “Puffin” was produced using a wet layup from glass fibre fabric by a convolute wrapping method, and dome structures were not implemented. The manufacturing of the “Puffin” motor was a highly labour-intensive process, requiring 32 to 40 person-hours to manufacture a single motor. There was also significant variance between motors, including large visually observed voids and variances in initial failure pressure, although the final failure pressure in hydrostatic testing was relatively stable. An image of the manufacturing process of the “Puffin” motor is given in fig. 3.2.

It is believed that with the use of filament winding techniques in combination with the knowledge gained from the design and manufacture of the “Puffin” rocket motor, a more efficient and reliable motor can be made. As with “Puffin”, the “Shearwater” motor is intended to use a “leave-in” mandrel manufactured from Phenolic Formaldehyde Cellulose Paper (PFCP) 21. However, in contrast to “Puffin”, the “Shearwater” motor will have two machined phenolic dome ends applied to each end of the main liner tube, enabling the use of overwrapped domes at each end of the motor.




Figure 3.1: Size comparison of the “Shearwater” rocket motor and the proposed layout of the “Puffin” rocket motor



Figure 3.2: Manufacturing of the “Puffin” rocket motor using a wet-layup technique (author on left)

It is hoped that with the use of filament winding and better design analysis, it will be possible to achieve a better propellant mass fraction for the “Shearwater” rocket motor than was achieved for “Puffin”, despite the MEOP being almost four times higher.

3.2 Research Aim and Questions

The research aim is to produce preliminary screening evidence for the manufacturing process parameters of the carbon-overwrapped, phenolic-lined motor design, and to characterise how process variables affect the mechanical response of manufactured test specimens.

The experimental design evolved during the test campaign as variables were discounted in favour of those that were found to be more influential. The final research aim and questions therefore reflect the analytical focus that emerged from the complete campaign rather than the initial questions that were posed before experimentation started. The principal research question is:

How are recorded winding tension and shrink-tape consolidation associated with manufacturing quality and hoop-mechanical response in OOA carbon-towpreg overwraps on paper-phenolic liners?

Five sub-questions were defined to support the primary research question and to provide a framework for the experimental result interpretation:

1. How do observed microscopic features relate to the observed mechanical response, and can these features be used to rescale and interpret the mechanical results?
2. How do the recorded tension and tape conditions relate to apparent stiffness, first detected mechanical failure, maximum force, and gross-area hoop stress?
3. How do the observed mechanical responses inform expected mechanical performance of the target motor?
4. What measurement and data-reduction methods are required to obtain a defensible comparison between configurations from split-D coupons, and what do they limit?
5. Which manufacturing configuration provides the most defensible candidate for subsequent motor development?

3.3 Thesis Structure

Based on the previously defined research questions, the remainder of the thesis is organised as follows:

- Chapter 4 documents the experimental campaign and the manufacturing, testing, and post-processing methods.
- Chapter 5 establishes the interpretation basis, evaluates measurement quality and uncertainty, and reports the microscopy and mechanical results.
- Chapter 6 interprets those results against the research questions, identifies under-resolved relationships, and discusses validity and generalisability.
- Chapter 7 projects the coupon evidence to preliminary motor-design implications and identifies the validation required before those projections can support a motor design.
- Chapter 8 summarises the supported findings, limitations, and recommendations.

Chapter 4

Experimental Methodology and Campaign

This chapter explains the specimen design, characterisation methods, mechanical testing setup and procedures used throughout the experimental campaign. The campaign design and manufacturing configurations defining each tranche of samples are first described, and then the manufacturing, testing and post-processing methods applied to them are discussed.

4.1 Experimental Campaign Design

Initial pilot work was conducted on a small batch of samples to test the winding setup and to provide some initial microscopy calibration data. This initial pilot work is referred to as Tranche 1 of samples.

Following the pilot work, the main experimental campaign was planned with a two-phase design. First, an exploratory screening phase, Tranche 2, would identify which manufacturing factors materially affected the performance of the wound samples. This would then be followed by a directed study of the factors that were retained by the screening pass in Tranche 3. The design of Tranche 3 was only finalised after the results of Tranche 2 were analysed. This chapter therefore presents the campaign chronologically, first setting out the parameters that were considered and the rationale for their selection, and then describing the three experimental tranches in manufacturing order.

4.1.1 Candidate Process Parameters and Experimental Scope

The campaign initially considered a large range of manufacturing parameters for their possible effects on mechanical performance and manufacturing cost. Not every candidate could be varied systematically in the experimental campaign because of time and resource limits. The paragraphs below distinguish the investigated factors from those retained as future process questions.

Rotation of the part during cure. Part rotation was considered as a possible means of limiting gravity-driven resin redistribution during cure. It was not varied systematically in the recorded campaign and is therefore retained as a future process question rather than analysed as an experimental factor.

Consolidation method. Autoclaved parts are typically consolidated with a vacuum bag, which both applies pressure and helps evacuate entrapped gas. Out-of-autoclave parts often instead use shrink tape, which contracts on heating. This provides a consolidation force, but traps any gas that is already present or that evolves during the cure. Shrink tape also provides less consolidation than vacuum bagging. Despite this, since the process under evaluation is out-of-autoclave, the two methods provide more comparable consolidation pressure than they would under autoclave conditions. The manual, skilled overhead of vacuum bagging remains far higher than that of shrink tape. A shrink-tape process would therefore be strongly preferred if the resulting part quality is comparable. However, the difficulty of manufacturing large numbers of samples using a repeatable vacuum-bagging process was too great for the resources available for this study, and so only a simple shrink-tape vs non-shrink-tape comparison was implemented for Tranches 2 and 3.

Winding tension. Higher winding tension can improve consolidation and lower void content, at the cost of potential fibre damage and increased resin squeeze-out. While tension has little direct bearing on manufacturing cost or speed, it was expected to be a primary lever for the mechanical performance of the wound laminate. This parameter was retained for evaluation.

Winding speed. Winding speed directly affects production throughput and may interact with tension, tow placement, and air entrapment. It was identified as a candidate variable but was not varied systematically in the reported tranches owing to the constraints of the available winding machine.

Liner drying. Liner drying was introduced after Tranche 1 revealed blistering in the liner during curing. Tranche 2 introduced it as a variable, but it was not varied in the directed Tranche 3 matrix.

Liner surface preparation. Surface roughness and cleanliness significantly affect the bond between composite parts [40], but advanced surface-preparation steps add cost and time. A comparison between a simple acetone-clean regime and a sand-then-clean regime for the liner/composite interface was therefore of practical interest. It is also known that hand-preparation methods can result in wide performance variance. This parameter was evaluated with a simple hand-sanding test in Tranche 2, but was not retained for Tranche 3.

4.1.2 Tranche 1: Pilot Work

Tranche 1 consisted of preliminary specimens cured for two hours at 170 °C. These specimens were used to develop the manufacturing workflow and to calibrate the optical-microscopy image-processing code. The samples produced were not included in the subsequent tensile testing. Blistering observed in the unbaked liner during cure motivated the liner-drying check described in section 4.1.3, while the observation is reported in section 5.3.

4.1.3 Liner Baking Experiment

Following the blistering of the unbaked Tranche 1 liner, a full-length phenolic mandrel was dried at 80 °C for two hours and weighed before and after drying. The check quantified removable mass consistent with moisture uptake; the measured result and its interpretation are reported in section 5.3. Tranche 2 records additionally preserved the before/after drying masses for configurations C1–C4, while C5 was intentionally not dried. These configurations are described in section 4.1.4.

4.1.4 Tranche 2: Screening


Tranche 2 screened five configurations across three liner tubes. All the samples in Tranche 2 were cured for two hours at 170 °C. The tensioning used one rubber band at its loosest setting, corresponding to an estimated suspended mass of approximately 600 g, or approximately 5.9 N. Table 4.1 records the factors and their levels for each configuration, along with the tube used and the surface preparation applied to that tube.

Table 4.1: Tranche 2 screening configurations

Config.	Tension	Consolidation	Drying	Tube	Surface preparation
C1	No Tension	No Tape	Dried	2	Acetone, then Isopropyl Alcohol (IPA)
C2	Tension	No Tape	Dried	1	Acetone, then IPA
C3	Tension	Failed Shrink Tape	Dried	2	Acetone, then IPA
C4	Tension	Failed Shrink Tape	Dried	1	120-grit sanding; acetone; IPA
C5	Tension	Failed Shrink Tape	Not Dried	3	Acetone, then IPA

The identifiers C1–C5 were the original configuration names used in the records of the testing, and so are retained here for provenance.

In these initial samples, Easy Composites CST-25-10 tape was used for the taped configurations. This tape was discovered to only be rated to 130 °C, below the 170 °C cure temperature, and therefore the tape failed during cure. It may have applied some compression before failure, but the effect of this was judged to be insignificant. Accordingly, C3–C5 failed to provide a valid comparison of functional shrink-tape consolidation. Figure 4.1 shows the applied and failed shrink tape.



(a) Before cure



(b) After cure

Figure 4.1: Tranche 2 tubes in the cure oven before (left) and after (right) the cure cycle. Tubes 3, 4 and 5 have shrink tape applied. The shrink tape failed during the 170 °C hold. This is evidenced by its crinkled, loose appearance in subfigure (b) compared to its initial shiny appearance in subfigure (a).

4.1.5 Decision Gate Between Tranches 2 and 3

The initial analysis of Tranche 2 was performed in order to determine which factors should be retained for the directed study in Tranche 3. The reasoning is set out here so that the directed design that follows can be read in sequence; the supporting measurements are reported in sections 5.5.1 and 5.5.3.

The analysis of data from the screening tranche showed differences in geometry and mechanical response across the configurations, but their similar terminal forces and confounded factor combinations prevented those differences from being assigned to one manufacturing variable. Recorded winding tension was retained because it was both a controllable consolidation lever and a plausible contributor to the observed configuration-dependent response. Shrink tape was retained because it offered a practical out-of-autoclave consolidation route, while its effect remained unresolved after the Tranche 2 tape exceeded its temperature rating and failed. Tranche 3 therefore replaced the broad screen with a directed comparison of tension level and functional tape. Drying, surface preparation, winding speed, and cure rotation were not treated as factors in that directed matrix and remain outside the isolated Tranche 3 comparison.

4.1.6 Tranche 3: Directed Study

Tranche 3 used a directed 3×2 arrangement of low, medium, and high winding tension with and without functional shrink tape across a single shared liner tube. The taped configurations were compacted with Dunstone 220R shrink tape, nominally 0.0508 mm thick by 25.4 mm wide, rated to 165 °C. This tape remained functional during the cure cycle. All six configurations were wound on the single shared liner, Tube 4, and cured for two hours at 165 °C. A summary of the conditions tested and number of coupons for each Tranche 3 configuration is given in table 4.2.

Table 4.2: Tranche 3 tension conditions; each level included tape and no-tape configurations on Tube 4

Level	Recorded tension setup	Coupons n	
		No tape	Tape
Low	Approximately 0 g	4	5
Medium	Initially ~1 kg (750 g to 1100 g); band snapped, second half treated as 600 g	5	5
High	Initially 1.8 kg to 2.2 kg, later ~1.6 kg	5	5

Because every Tranche 3 coupon is a subsample of the same manufactured tube, the factorial analysis compares regions within Tube 4 and does not establish between-tube manufacturing repeatability.

4.2 Manufacturing

The test articles were manufactured in-house using custom-made equipment. This section describes the winding equipment, specimen geometry, material preparation and cure conditions used throughout the experimental campaign.

4.2.1 Filament Winding Machine

A custom filament winding machine, designed and manufactured by the author and owned by DARE, was used for the manufacture of all specimens in this study. This machine is a four-axis CNC winder, with a rotating mandrel, traversing carriage, cross-carriage and rotating delivery head. It was built from a combination of off-the-shelf and custom components, with the main frame being made from standard 4040 aluminium extrusion. The control board for the machine is a BIGTREETECH Octopus V1.1, which uses a 32-bit ARM Cortex-M4 processor. A render of the winder is shown in fig. 4.2. A custom modification of the open-source Marlin firmware is used for the control of the four axes and the winding process. The filament spool is mounted to a static frame, and the fibre tension is applied with the use of one or two elastic bands, which are stretched around the cardboard spool of the filament bobbin. This allows for a broad tension range to be applied, and the tension can be adjusted by changing the number of bands or their elastic properties. The tension cannot, however, be adjusted dynamically during the winding process, and the tension is therefore assumed to be constant for each specimen.



Figure 4.2: Render of the filament winder used for manufacture of the test samples. Filament is supplied pre-tensioned from the lower left corner of the figure, passing around the redirection rollers and through the delivery arm to the delivery roller, which stands around 2 cm from the mandrel surface. The mandrel is suspended along the axis provided along the upper left extrusion.

4.2.2 Specimen Design

The test specimens were manufactured on PFCP tubes of nominal 100 mm internal diameter and 2 mm thickness. A nominal three-layer overwrap of Teijin Tenax TSTG IT 01 2200007 X01, at a nominal tow width of 6.35 mm, was applied using the filament winding machine described in section 4.2.1. The three-layer wind was based on the maximum expected capacity of the test setup, with the aim of achieving as large a total cross-sectional area as possible to minimise the effects of small variances in the material. The 6.35 mm tow width was the nominal width of the towpreg used, so a strictly 90° wind would inherently lay up this width.

In all cases except a single case to investigate the effect on the overall system performance, the liner was baked for two hours at 80 °C to remove moisture. All liners were prepared with successive wipe-downs with acetone and isopropyl alcohol as the last step before winding to ensure a clean surface for bonding.

4.2.3 Cure Conditions

The cure consisted of a two-hour hold at the defined cure temperature. Tranches 1 and 2 were held at 170 °C for two hours. Tranche 3 was held at 165 °C for two hours, matching rather than exceeding the 165 °C rating of the replacement shrink tape, which was observed to remain intact through the cure (section 4.1.6).

4.3 Testing

All mechanical testing used a split-D hoop configuration, selected for its simple design and manufacture, its ease of use, and relatively high test cadence. This section describes the test method and fixture, the strain instrumentation applied to Tranche 3, and the machine settings common to all specimens.


4.3.1 Split-D Hoop Specimens

The split-D hoop test was selected for material characterisation because it is a relatively simple and low-cost method that can be applied to the as-wound ring geometry. This allows for the direct comparison of the fibre-dominated response between process configurations. The poor characterisation of the matrix-dominated response is a known limitation of the split-D test, but the tests required for matrix characterisation were not feasible within the available time and budget.


In the split-D test, two semicircular inserts support the inside of the ring and are displaced apart until progressive damage and final failure occur (fig. 4.3). The fixture introduces local bending and contact stresses near the two insert gaps, so the test is used as a comparative hoop-performance method rather than as a uniform uniaxial material test.

A custom split-D fixture was designed to accommodate the PFCP liner geometry. Its load-capacity calculation is provided in appendix A. For the area-normalised stress calculation, the liner contribution was neglected and the applied force was attributed wholly to the carbon overwrap. This analytical assumption is supported by the recorded failure sequence reported in section 5.6.

The split-D specimens were prepared by winding a single continuous carbon tow around the liner, translating after each three-layer hoop to allow for many hoops to be created from a single liner tube. The samples were cut from the tube using a Dremel tool, and excess phenolic liner was removed to leave a nominal 6.35 mm wide carbon strip. The translation between hoops resulted in two small sections of non-uniform winding on the first and last layers of each hoop. These could not easily be removed, so were retained in the final specimens. To mitigate the effect of these non-uniform areas, the hoop was placed onto the split-D fixture such that the non-uniform areas would be located away from the split point. Since the maximum stress occurs at the split point, this ensured that the non-uniform areas would not be the first to fail and would not dominate the measured response. The wound sample in the split-D jig before and after tensile testing is displayed in fig. 4.3.



(a) Before testing



(b) After testing

Figure 4.3: Split-D hoop specimen before and after testing, showing the terminal failure state

4.3.2 Strain Measurement

Tranche 2 specimens were not instrumented with strain gauges, so no strain data are reported for that tranche. For Tranche 2, cross-head extension was retained only as a comparative system-displacement measure and is not interpreted as specimen strain, since the cross-head measurement cannot isolate the specimen from the overall system response.

Tranche 3 specimens were instrumented with two Kyowa KFGS-6-230-C1-11 strain gauges, with one gauge bonded to each side of the split-D specimen at the split point. The two channels therefore sampled the carbon-surface response on the opposed load-bearing sides. Each channel was baseline-zeroed independently using pre-peak data below 2 % of the specimen maximum force. A gauge channel was retained when its response between 10 % to 50 % of maximum force had a positive slope, a force correlation of at least 0.80, and a range of at least 0.05 percentage points. Values at or above 3 % absolute strain were treated as saturation or gauge-failure states. The gauge-derived stress–strain slope is reported as an apparent modulus only. The split-D setup induces non-uniform loading, and the gauges are bonded to the outer surface at the split point, where local bending is superimposed on the membrane hoop strain. Figure 4.4 illustrates this superposition as the fixture halves separate and the ring straightens locally at the gap. The recorded gauge strain, therefore, does not represent the section-average strain of a uniformly loaded pressure vessel, and the resulting slopes are treated only as comparative indices of local response between configurations.




Figure 4.4: Schematic of the membrane–bending superposition at the split-D gap. Left: as the fixture halves separate, the ring straightens locally at the gap (exaggerated here for clarity). Right: the resulting outer-surface stress along the horizontal plane at the gap is the superposition of the section-average membrane hoop stress and a local bending stress. This results in a lower measured strain than would be expected from the global mean stress.

4.3.3 Mechanical Test Settings

The mechanical testing settings were consistent across all Tranche 2 and 3 specimens. All specimens were tested on a Zwick 250 kN tensile testing machine available at Delft Aerospace Structures and Materials Laboratory (DASML) at TU Delft. All specimens were tested at room temperature, with a constant cross-head speed of 1.00 mm min^{-1} ($0.01667 \text{ mm s}^{-1}$). Force, cross-head extension, time, and available strain-gauge channels were acquired at 10 Hz.

4.4 Post-Processing

Following testing, microscopy and tensile-data processing were used to derive the geometric, microstructural and mechanical quantities reported in this thesis. This section describes the preparation and analysis procedures applied to the acquired specimens and test records.

4.4.1 Optical Microscopy

Optical microscopy was used to estimate the load-bearing carbon envelope, carbon void area fraction, void-size distribution, and interface bonding quality. Tranche 1 images were used to develop and calibrate the image-processing workflow. For Tranches 2 and 3, one or two polished sections were produced per configuration. Where

two sections were produced, their measured carbon areas were averaged. The mean area or single area was then applied to every tensile coupon in the configuration.

The samples were prepared by using a Dremel grinding tool to excise a small section of the specimen, which was mounted in epoxy and polished using a Struers Tegramin-20 polishing machine.

Microscopy was performed on a Keyence VK-X1000 laser scanning confocal microscope, using a $10\times$ magnification. Both the laser-confocal and the white-light optical imaging modes were trialled in initial testing, with the white-light optical mode being preferred after evaluation. All reported measurements were made on the optical Red-Green-Blue (RGB) images, since the segmentation workflow discriminates carbon from phenolic by colour and the confocal height data did not separate the two materials; the height channels were therefore excluded from the analysis.

An automatic image-segmentation workflow was attempted, but it was found that the epoxy background and the phenolic and carbon did not have sufficiently distinct colouration to allow for a fully automated segmentation. Therefore, the image segmentation was performed manually, with the representative carbon and phenolic areas being labelled by hand.

Void candidates were identified using an automated workflow as locally dark regions within an eroded material mask, excluding the approximately $20\mu\text{m}$ boundary band most vulnerable to illumination and segmentation artefacts. Void candidates within the carbon smaller than $50\mu\text{m}^2$ were rejected. The Keyence microscope provides a scale line in each image, which was used to convert pixel counts to physical areas through the automated image-processing workflow.

The carbon-phenolic interface metric is used descriptively only because damage during specimen preparation and polishing can introduce artefacts that are indistinguishable from true interface delamination.

4.4.2 Tensile-Data Reduction

The final processed dataset contains 49 split-D specimens from Tranches 2 and 3. Force measurement, cross-head extension, time, and available strain-gauge channels were retained.

4.4.2.1 Measured Carbon Area

The nominal thickness of the three-hoop layup was estimated at 0.80 mm. Combining this with the nominal tow width of 6.35 mm gave an area of 5.08 mm^2 for a total (the split runs through two cross sections — one on each side of the hoop) of 10.16 mm^2 nominal area.

The area-normalised calculations used the measured carbon-envelope geometry rather than assuming the nominal dimensions. The effective area was defined as:

$$A_{eff} = 2A_{\text{carbon,wall}} \quad (4.1)$$

The configuration-level mean area from the microscopy sections was applied across all the samples of that configuration. Of note is that the detected voids and any resin-rich regions of the layup remain included in the area: the stress is therefore a gross envelope stress rather than attempting to correct for these factors.

This effective area was used in the gauge-derived apparent-modulus calculation. Because each configuration is represented by only one or two microscopy sections, this is a shared denominator that does not capture the variation between individual specimens. The measured maximum force is therefore the primary response of the sample. Nominal-area stress is reported only as a rescaling of measured force, since the nominal area is common to all specimens, and gross-envelope stress is retained as a secondary measure. An area-independent gauge force-strain slope is used alongside the measured-area apparent modulus.

4.4.2.2 Slack Correction and Apparent Stiffness

Because a split-D ring has no conventional gauge length and because the fixture is not perfectly rigid, cross-head travel is reported as slack-corrected extension rather than strain. The initial cross-head motion was not immediately reflected in the recorded force, consistent with slack in the test setup and the contact load path of the split-D fixture. The amount of slack also differed between specimen configurations and would therefore confound comparisons based on raw cross-head displacement. For each specimen, a least-squares line was fitted to force versus raw cross-head extension over 20 % to 40 % of that specimen's maximum force, using only the

initial loading history before the first maximum (to prevent confounding by any stray post-failure measurements). The fitted zero-force intercept was subtracted from the raw extension. Negative corrected values were retained for traceability but omitted from force-extension plots. A few specimens begin their corrected history at a small negative extension because the extrapolated slack take-up is greater than the earliest recorded travel.

The fitted force–extension slope is an apparent system stiffness containing specimen, fixture, insert-contact, seating and machine compliance, and is not interpreted as a material modulus.

Figure 4.5 illustrates the procedure on a representative specimen. The least-squares line fitted over 20 % to 40 % of $F_{\max}$ is extrapolated back to zero force, and the resulting intercept is subtracted from the raw cross-head travel to give the corrected extension value.




Figure 4.5: Slack-correction procedure on a representative Tranche 3 specimen. The red line is the least-squares fit over the shaded 20 % to 40 % $F_{\max}$ window. Its zero-force intercept (the red dot) defines the slack take-up that is subtracted from the raw extension. The upper axis shows the resulting corrected extension.

4.4.2.3 Partial Failure Detection

The first detected failure was defined algorithmically as the first force drop between one data point and the next at least equal to the larger of 0.4 kN or 2 % of $F_{\max}$, after the specimen had reached 10 % of $F_{\max}$. Candidates separated by no more than 0.2 s were merged. Counts were also evaluated at 3 % and 5 % thresholds because event count was threshold-sensitive. The detector identifies a mechanical failure, but does not attempt to ascribe a physical cause to it (e.g. liner cracking, outer-layer delamination, inner-layer/interface delamination, or carbon rupture). Failure type was noted during the testing when visible to the naked eye, but this was not always possible.

4.4.2.4 Statistical Analysis

Because the Tranche 2 manufacturing factors were varied together rather than independently, no single factor can be isolated from that dataset. Tranche 2 was therefore analysed with a one-way analysis of variance treating configuration itself as the factor, which tests only whether the five configurations differ, not why. The real goal of Tranche 2 was to screen for significant differences between configurations, and to indicate the most interesting factors to consider for Tranche 3.

Tranche 3 followed a 3-by-2 factorial design, with three levels of winding tension and shrink tape either present

or absent. Each term (tension, tape, and their interaction) was tested by comparing the full model against the model with that term removed.

Several metrics were calculated per tranche, so the resulting p -values were adjusted using the Benjamini-Hochberg procedure to control the false-discovery rate. Throughout, a p -value is the probability of observing an effect at least as large as the one found, if the factor in question actually had no effect, and a q -value is the corresponding Benjamini-Hochberg-adjusted p -value, which limits the expected proportion of false positives among all results declared significant. The adjusted q -values are reported alongside the raw p -values, and the adjustment was applied within each tranche separately.

Both the direct-force results and the supporting area-normalised results were included when correcting for the number of statistical tests. If the two disagreed, the direct-force result remained the main conclusion.

Chapter 5

Results

This chapter presents and discusses the outcomes of the experimental campaign of section 4.1. The supporting liner-baking result is given first, followed by the microstructural and mechanical results and their interpretation in the context of the research objectives of section 3.2. Whereas section 4.1 preserves the chronological development from screening to directed study, this chapter is organised thematically by outcome so that Tranche 2 and Tranche 3 evidence can be aggregated for analysis.

5.1 Reported Dataset and Interpretation Basis

The results reported in this chapter are drawn from 49 split-D coupons taken from Tranches 2 and 3. Results from the pilot work (Tranche 1) are not included. Tranche 3 provides the only comparison with functional shrink tape. The constraints this places on interpretation, in particular the confounding within Tranche 2 and the single-tube origin of the Tranche 3 regions, are set out in section 5.2.

The processed force-extension curves obtained using the slack-correction procedure of section 4.4.2.2 are shown in figs. 5.1 and 5.2. Both display the initial response, near-linear loading and progressive damage regimes.

5.2 Data Quality and Measurement Validation

It is important to establish what the recorded measurements are able to support before any process effect is interpreted. This section presents an audit of the recorded measurements and the initial processing that was applied.

5.2.1 Evidence Hierarchy and Traceability

The tensile testing dataset is the primary body of evidence, and consists of 49 tests split between Tranche 2 (20 tests) and Tranche 3 (29 tests). All configurations of Tranche 2 had four tests, while all configurations of Tranche 3 had five except for Low Tension | No Tape, which had four. Each test gives a maximum force measurement, a slack-corrected cross-head history, a fitted system stiffness and a set of algorithmic force-drop metrics. Additionally, each Tranche 3 coupon carries two recorded gauge channels that are subject to the validity rules set out in chapter 4. Repeated coupons were taken from each of the eleven mechanical configurations to quantify the scatter in the process. However, since all Tranche 3 coupons were taken from a single manufactured liner, the residuals measure within-region variability rather than between-liner repeatability. The magnitude of the between-tube component is therefore not estimated by these data, and every Tranche 3 factor result is reported as a within-tube association (table 5.3).

The microscopy dataset is the secondary body of evidence, and consists of 15 final images split between Tranche 2 (6 images) and Tranche 3 (9 images). Seven of the eleven mechanical configurations are represented by a single section and the remaining four by two. The microscopy sections are listed in table B.1.




Figure 5.1: Processed force versus slack-corrected cross-head extension for Tranche 2




Figure 5.2: Processed force versus slack-corrected cross-head extension for Tranche 3

5.2.2 Slack-Correction Diagnostics

The slack-correction procedure and its interpretation limits are defined in section 4.4.2.2, while the fit diagnostics and measured slack offsets are reported here. The group-mean coefficients of determination for the linear fits ranged from 0.9993 to 1.0000 across the eleven groups, which indicates that the fitted slopes are a good representation of the structural response within that interval. None of the 49 fits fell below $R^2 = 0.98$; the single lowest value was $R^2 = 0.9967$ for a Medium Tension | Shrink Tape specimen. Table 5.1 reports the minimum R^2 within each group rather than the mean, as the more conservative diagnostic. The fitted slack take-up had a group mean ranging from 1.1 mm for C3 | Tension | Failed Shrink Tape | Dried to 2.0 mm for Low Tension | Shrink Tape, confirming that raw cross-head displacement would conflate initial seating with the subsequent structural response. Table 5.1 gives the group-by-group diagnostic values for both tranches. This table also shows that the Tranche 2 coupons consistently took up less slack (1.10 mm to 1.34 mm) than those of Tranche 3 (1.72 mm to 1.97 mm).

Table 5.1: Slack-correction diagnostic summary for all eleven configurations

Group	n	Minimum R^2	Mean slack (mm)	Standard Deviation (SD) (mm)
<i>Tranche 2</i>				
C1 No Tension No Tape Dried	4	0.9998	1.258	0.072
C2 Tension No Tape Dried	4	0.9998	1.335	0.042
C3 Tension Failed Shrink Tape Dried	4	0.9999	1.105	0.067
C4 Tension Failed Shrink Tape Dried Sanded	4	1.0000	1.153	0.046
C5 Tension Failed Shrink Tape Not Dried	4	0.9999	1.173	0.018
<i>Tranche 3</i>				
Low Tension No Tape	4	0.9999	1.878	0.082
Low Tension Shrink Tape	5	0.9996	1.966	0.126
Medium Tension No Tape	5	0.9999	1.725	0.040
Medium Tension Shrink Tape	5	0.9967	1.761	0.160
High Tension No Tape	5	0.9999	1.885	0.031
High Tension Shrink Tape	5	0.9998	1.791	0.021

5.2.3 Strain-Gauge Data Quality

58 total gauge channels were recorded across the 29 Tranche 3 coupons, and the validity of each channel was assessed against the criteria set out in chapter 4. 54 of the 58 channels satisfied the validity criteria, and every coupon retained at least one valid channel.

The per-configuration counts and the corresponding gauge–cross-head agreement are given in table 5.2.

Table 5.2: Summary of valid strain-gauge channels by configuration. The R^2 is computed per channel against the slack-corrected cross-head extension; the group “Mean” and “Minimum” columns aggregate those per-channel values, and the two gauges are not pre-averaged.

Group	n	Channels		Valid-channel R^2	
		Recorded	Valid	Mean	Minimum
Low Tension No Tape	4	8	7	0.996	0.986
Low Tension Shrink Tape	5	10	10	0.938	0.834
Medium Tension No Tape	5	10	9	0.983	0.875
Medium Tension Shrink Tape	5	10	9	0.978	0.886
High Tension No Tape	5	10	9	0.982	0.868
High Tension Shrink Tape	5	10	10	0.996	0.983
All Tranche 3	29	58	54	0.978	0.834

Among the valid channels, the mean gauge–cross-head R^2 ranged from 0.938 for Low Tension | Shrink Tape to 0.996 for Low Tension | No Tape and High Tension | Shrink Tape. The lowest individual value was 0.834, which indicates that satisfying the force-correlation criterion does not guarantee equally close agreement with cross-head motion.

Given the positioning of the gauges at opposite sides of the split-D setup, the differences between the channels may reflect genuine bending or asymmetry, or equally bonding or instrumentation error. Taking the mean of the two gauges reduces the sensitivity to deviations in a single side, but this cannot eliminate error from one side unless that side is detected as anomalous and therefore removed from the analysis. Furthermore, the apparent stiffness of the carbon hoop as measured by the strain gauges will not be the same as that measured on a standard flat coupon of the same material, as a result of the uneven loading and bending stresses resulting from the specific test setup. Therefore, the gauge-derived slopes can only be given as apparent moduli, rather than true moduli, and are interpreted alongside overall system stiffness rather than as a replacement for it.

5.2.4 Failure-Detection Sensitivity

An automated event detection algorithm was used to attempt to monitor the progressive failure behaviour of the hoop samples. A threshold to detect drops above a certain percentage of $F_{\max}$ was applied. As might be expected, the number of detected events fell as the threshold percentage rose, indicating that event count is a property of the detection rule as much as the specimen. Figure 5.3 quantifies this sensitivity: the mean detected event count fell from approximately 3.5 at the primary threshold to approximately 2 at the 5% threshold in both tranches. In all cases, however, the primary failure event is large enough to be caught by all reasonable values for the drop threshold, and therefore remains useful as a consistently applied marker for the first failure. Figure 5.4 shows a representative specimen, illustrating both the full force-time history and a close-up of the first detected drop.




Figure 5.3: Mean detected event count by drop-detection threshold for Tranches 2 and 3. Error bars show ± 1 standard deviation across specimens within each tranche. The primary threshold is 0.4 kN or 2% of $F_{\max}$ — whichever is larger

The average ratio between the first detected failure and the maximum force varied widely between groups, ranging in Tranche 3 from 0.545 for Medium Tension | Shrink Tape to 0.699 for Low Tension | No Tape, with High Tension | Shrink Tape at 0.681.

Three factors limit the interpretation of the first-event force. First, the scatter between coupons within a group is substantial, and is largest in the low-tension groups. Second, the mode of the observed first failure was not consistent, varying between liner cracking and liner–hoop disbonding. Third, the applied loading case is not truly representative of the loading experienced by an operational rocket motor. The first-event force is therefore not treated as a liner-cracking limit, and is instead reported only as the first detected mechanical discontinuity.




Figure 5.4: Representative force–time history (left) and close-up of the first detected drop (right) for a Medium Tension, No Tape specimen. The dashed line marks the 10 % F_{max} search-start level; the red dot and vertical line mark the first event detected by the primary threshold. Note that the dot does not sit at the max pre-drop force, as it instead detects the first sample where the following drop is larger than the threshold value.

5.2.5 Microscopy Audit and Representativeness

The microscopy workflow returns a carbon envelope area, a detected dark-region fraction, a void count and a set of interface statistics for each final image. The void detection was calibrated by eye, by adjusting the thresholds until the generated overlays agreed with the features visible in the images, and the same settings were then applied to every section. It is to be noted, however, that damage to the sample from the cutting and polishing process, fibre pull-out, and other factors can cause additional dark regions that may have been included in the detected void area, and that the two-dimensional image is not a direct measurement of three-dimensional void volume fraction (which is expected to vary over the circumference of the hoop).

The reported quantity is therefore termed the *detected optical void area fraction*, since it is a two-dimensional image-derived estimator rather than a direct measurement of three-dimensional void volume fraction. Similarly, the total carbon area measured is a gross area, describing the carbon-layer envelope inclusive of detected voids and resin-rich regions, and it should not be read as an area of carbon fibre or as a fibre volume fraction.

Where two sections exist for a configuration, the difference between them offers an indication of the variability throughout the hoop. This variability was particularly pronounced in the Low tension configurations, where the roving wandered significantly during the winding process.

Applying a configuration mean derived from one or two images to four or five coupons introduces a shared geometric denominator, so that coupon-to-coupon force scatter is measured while the corresponding scatter in area is not. The effective precision of any stress comparison is therefore lower than the number of mechanical coupons alone would suggest.

5.2.6 Area Normalisation

The force and area normalisations defined in section 4.4.2 are applied consistently below. Because every coupon in a configuration shares one microscopy-derived area, measured-envelope stress preserves the within-group force scatter but can change the ranking between groups. Force, measured area and both stress representations are therefore shown side by side.

5.2.7 Statistical Validity

Tranche 2 and Tranche 3 were designed to explore different aspects of the manufacturing process, and they consequently have different statistical properties. Tranche 2 cannot isolate the effects of tension, tape, drying, surface preparation or tube identity, and its one-way Analysis of Variance (ANOVA) therefore only describes the configuration separation. Tranche 3 adopts a 3×2 regional arrangement, but because its coupons were all taken from Tube 4, their residuals measure within-region variability on a single tube. The factorial F -tests can identify

separation among those regions under the fitted model, but they cannot strictly demonstrate that the same effect would recur on independently wound tubes. Benjamini–Hochberg adjustment limits the false-discovery rate across the metrics tested, but it is not possible to repair the confounding without supplying the missing manufacturing replication.

For these reasons, p - and q -values are reported alongside effect sizes and descriptive intervals, and a non-significant term is described as *not resolved* rather than as proof of no effect or of equivalence. The single Tranche 3 result that remained resolved after adjustment — the association between recorded tension and apparent system stiffness — is similarly described as a within-tube association.

Effect sizes show the magnitude of the observed differences. They are not used to estimate how many additional coupons would be needed to resolve those differences statistically, because the next experiment requires independently manufactured tubes rather than more coupons from Tube 4.

5.2.8 Permitted Interpretation

These observations are compiled in table 5.3, which sets the principal limitation of each reported quantity against the interpretation that the available evidence will support.

Table 5.3: Measurement limitations and permitted interpretations

Quantity	Principal limitation	Permitted interpretation
Maximum force	Split-D contact load path	Primary comparative terminal capacity of the tested rings
Nominal-area stress	Fixed nominal geometry	Unit conversion of force; no independent statistical evidence
Measured-envelope stress	Shared, sparsely sampled envelope area	Secondary area-efficiency sensitivity analysis
Apparent stiffness	Includes machine, fixture, contact, and seating	Structural/system response of the tested setup
Gauge modulus	Local curved-ring surface measurement	Apparent local response, not a standard coupon modulus
First event	No physical identifier	First algorithmically detected discontinuity
Void content	Unvalidated two-dimensional segmentation	Exploratory detected dark-area fraction
Interface continuity	Preparation-sensitive segmentation	Descriptive interface-contact indicator
Factorial terms	One shared manufactured tube	Within-Tube 4 regional associations

5.3 Liner Baking

A full-length PFCP liner was baked at 80 °C for two hours and weighed before and after the treatment. Its mass fell from 1295.2 g to 1287.5 g, a loss of 7.7 g or 0.59 %. The treatment therefore removes a measurable amount of mass from the as-received material, which is consistent with absorbed moisture or another volatile, which may have contributed to the Tranche 1 blistering. The blistering observed in the unbaked Tranche 1 liner is shown in fig. 5.5.

The lost material was not chemically identified, so this check on its own does not establish the mechanism. However, it was found that the liner used in Tranche 3 also suffered blistering even after the baking treatment was applied, so the mechanism is not fully mitigated by the treatment.

5.4 Microscopy Results

Optical microscopy was performed on polished cross-sections from each configuration in order to measure the carbon cross-sectional area, the void content and the quality of the carbon–phenolic interface. Because the consolidated cross-section differed substantially from the nominal 6.35 mm width, measured area is reported explicitly and used in the measured-envelope sensitivity analysis.



Figure 5.5: Section through a blister in the unbaked Tranche 1 phenolic liner

5.4.1 Carbon Cross-Sectional Area

The cross-sectional area was measured from the polished section of each configuration, taking the mean of the two sections where two images existed, and the areas of the two split-D walls were summed to give the effective area. The results are summarised in table 5.4.

Table 5.4: Measured carbon cross-sectional area by configuration. Nominal area is 10.16 mm^2 for two cross-sections of 6.35 mm width and 0.8 mm thickness.

Configuration	Single-wall area (mm^2)	Effective area (mm^2)	vs nominal
<i>Tranche 2</i>			
C1	4.47	8.95	-12%
C2	4.41	8.82	-13%
C3	4.40	8.81	-13%
C4	5.83	11.65	+15%
C5	6.15	12.29	+21%
<i>Tranche 3</i>			
Low Tension No Tape	4.55	9.10	-10%
Low Tension Shrink Tape	5.16	10.32	+2%
Medium Tension No Tape	5.20	10.40	+2%
Medium Tension Shrink Tape	4.45	8.90	-12%
High Tension No Tape	4.78	9.56	-6%
High Tension Shrink Tape	4.51	9.02	-11%

The nominal 0.8 mm thickness by 6.35 mm width generally exceeds the true carbon cross-section. The exceptions are C4 and C5, which measured substantially larger than nominal, and Low Tension | Shrink Tape and Medium Tension | No Tape, which sat marginally above the nominal value. No consistent effect of tension or tape on area was found in Tranche 3, where the differences between configurations were of the same magnitude as the image-to-image spread within a configuration, given that only one or two sections were available for each. This is despite significant resin squeeze-out being clear in the samples using shrink-tape for compaction, while the untaped samples show no such squeeze-out. This is illustrated in fig. 5.6.

5.4.2 Void Content

A representative raw microscopy image, its hand-segmented carbon region and the corresponding carbon-void detection are shown in fig. 5.7. The complete microscopy dataset is found in appendix B.






Figure 5.6: Resin squeeze-out visible on a shrink-tape-compacted configuration (right) by the darkening of the phenolic liner, compared with an untaped configuration (left) showing no squeeze-out



(a) Raw microscopy image



(b) Hand-segmented carbon region



(c) Detected voids within the same carbon region

Figure 5.7: Microscopy processing example for the High Tension | Shrink Tape configuration. Panel (a) shows the raw microscopy image, (b) shows the hand-segmented carbon mask directly in blue, and (c) shows detected dark features (voids) within that region filled red with yellow rims.

All the void metrics in this thesis are measured within the segmented carbon envelope of the polished sections only; voids in the liner are not part of the measurement, and “void content” below always refers to this carbon-region quantity. The detected void area fraction was measured from the microscopy images, and the results are shown in fig. 5.8. Without tape, void content edged upwards with tension, from 11.7 % to 11.9 % and then 12.1 %. With tape it fell, from 11.7 % to 11.2 % and then 9.0 %, so that the benefit of tape grew with tension and reached approximately three percentage points at high tension. This is the most consistent manufacturing signal in the microscopy data, although it again rests on only one or two sections per configuration.




Figure 5.8: Void area fraction by configuration. Tranche 2 samples are in grey, Tranche 3 in blue. ± 1 sample standard deviation error bars are shown where two images were taken for a configuration. All other samples had one image per configuration.

The corresponding Tranche 2 porosities were 7.9 % for C1, 9.7 % for C2, 8.9 % for C3, 10.0 % for C4 and 10.9 % for C5. Every Tranche 2 value therefore lies below every untaped Tranche 3 value, which is the opposite of what the failed Tranche 2 tape would suggest. The two tranches were wound on different liner tubes and cured at 170 °C and 165 °C respectively, and each configuration rests on only one or two sections, so this offset is not attributed to any process factor and void content is compared within a tranche only.

The void size distribution shown in fig. 5.9 was strongly right-skewed in every configuration. Most voids were between 12 and 30 μm in equivalent diameter, with a tail extending to approximately 100 μm , and tape applied at high tension appeared to significantly thin that tail. The distributions are therefore dominated by small voids, with only a few larger ones present.




Figure 5.9: Void size distribution by configuration.

5.4.3 Carbon–Phenolic Interface

The carbon–phenolic bond line is the mechanical load-transfer boundary between the overwrap and the liner; a poor interface can fail before either surrounding material does. The quality of the interface was therefore assessed as part of the microscopy analysis. The close interface fraction, defined as the proportion of the bond line in which the carbon and phenolic are in close contact, was measured from the microscopy images. It was found that the close interface fraction was near-complete with tape at approximately 0.99 and lower without it, at approximately 0.75 to 0.89, which indicates that tape produces a tighter bond line. Across the six Tranche 3 configurations, mean close-interface fraction was not significantly correlated with mean first-event stress ($r = -0.37$, $p = 0.46$). Applying an arcsine-square-root transformation to the close-interface fraction did not alter this conclusion ($r = -0.39$, $p = 0.44$). Since the interface is the feature most sensitive to sectioning and polishing (and during the sectioning process there were several cases of the carbon separating entirely from the phenolic), the bond-line metric cannot be used to support any conclusion about strength, and is reported as a qualitative observation only.

5.5 Hoop Strength Results

This section presents the measured hoop-strength response, first in terms of maximum force and then using area-normalised terminal capacity. The effects of winding tension and shrink-tape consolidation are subsequently evaluated for the Tranche 3 specimens.

5.5.1 Maximum Force

The individual maximum-force distributions are shown in fig. 5.10. In Tranche 2 the group means occupied the narrow range 19.00 kN to 21.01 kN, whereas the measured carbon envelope area varied from 4.40 mm² to 6.15 mm² per wall. The coefficient of variation between the Tranche 2 configuration means was 3.6 %, while within the groups it ranged from 4.1 % to 10.4 %.

In Tranche 3 the mean maximum forces were 17.68 kN without tape and 17.73 kN with tape at low tension, 20.24 kN and 19.23 kN at medium tension, and 20.65 kN and 18.94 kN at high tension. The within-group coefficient of variation ranged from 5.1 % for High Tension | No Tape to 20.0 % for Low Tension | No Tape. As described in section 5.2, this scatter describes coupons cut from the tested regions and does not quantify the manufacturing uncertainty between tubes.




Figure 5.10: Maximum measured force by configuration. Each box spans the interquartile range, the orange line marks the median and the green triangle the group mean; whiskers extend to the most extreme values within $1.5\times$ the interquartile range, and open circles mark points beyond that range.




Figure 5.11: First-break force by configuration, taken as the force at the first detected abrupt load drop. Each box spans the interquartile range, the orange line marks the median and the green triangle the group mean; whiskers extend to the most extreme values within $1.5\times$ the interquartile range, and open circles mark points beyond that range.

5.5.2 Terminal Capacity and Area Normalisation

Table 5.5 summarises the group means for terminal force, including stress representations and assigned effective area.

Table 5.5: Group mean terminal force, nominal stress, measured-envelope stress and assigned effective area

Configuration	$F_{\max}$ (kN)	σ_{nom} (MPa)	σ_{meas} (MPa)	A_{eff} (mm ²)
<i>Tranche 2</i>				
C1	19.00	1870	2124	8.95
C2	21.01	2068	2382	8.82
C3	20.17	1985	2291	8.81
C4	20.41	2009	1752	11.65
C5	20.34	2002	1654	12.30
<i>Tranche 3</i>				
Low Tension No Tape	17.68	1740	1942	9.10
Low Tension Shrink Tape	17.73	1745	1719	10.31
Medium Tension No Tape	20.24	1992	1946	10.40
Medium Tension Shrink Tape	19.23	1893	2161	8.90
High Tension No Tape	20.65	2032	2160	9.56
High Tension Shrink Tape	18.94	1864	2100	9.02

5.5.3 Effect of Winding Tension

In Tranche 2, maximum force did not resolve a configuration difference (one-way ANOVA $p = 0.40$, $\eta^2 = 0.22$). Every group lay between 19 and 21 kN, with a between-group Coefficient of Variation (CV) of 3.6 %. Measured-envelope stress, by contrast, differed strongly ($p < 0.0001$, $\eta^2 = 0.83$) and its group means had a CV of 15.9 %; this is a secondary normalisation result. C4 and C5 have 15 to 21 % more segmented envelope area but carry no more terminal force, so the load calculated per unit of gross envelope is lower. Ultimate stress and measured carbon area are strongly negatively related at the configuration level ($r = -0.85$, $p < 0.01$, $n = 11$ configurations across both tranches; within Tranche 2 alone, $r = -0.97$, $p = 0.008$, $n = 5$), but this relationship is mathematically coupled, since area is the denominator of the stress. It is therefore clear that terminal force is comparatively stable across the Tranche 2 configurations while the gross envelope geometry (and therefore the terminal stress) is not.

In Tranche 3, no factor resolved maximum force: tension gave $p = 0.091$ ($q = 0.36$, partial $\eta^2 = 0.17$), tape $p = 0.26$ ($q = 0.50$) and the interaction $p = 0.69$ ($q = 0.83$). The low-, medium- and high-tension pooled means were approximately 17.7 kN, 19.7 kN and 19.8 kN. The measured-envelope sensitivity gives a smaller raw tension p -value ($p = 0.029$, $q = 0.30$, partial $\eta^2 = 0.25$), but this does not survive correction and is strongly influenced by the assigned areas. Because the coupons are nested within one tube, neither result demonstrates a repeatable manufacturing effect. Figure 5.12 compares the four group metrics for Tranche 3.

5.5.4 Effect of Shrink-Tape Consolidation

Shrink tape showed no resolved effect on maximum force in Tranche 3 ($p = 0.26$, $q = 0.50$, partial $\eta^2 = 0.050$), nor on measured-envelope ultimate stress ($p = 0.87$, $q = 0.93$).

The high-tension microscopy in section 5.4 nevertheless visually suggests improved consolidation with tape. This did not translate into a measurable strength benefit at the sample sizes tested. A possible explanation is that strength here is dominated by the total cross-section and the degree of consolidation rather than by void content alone.

The addition of tape was also associated with a great improvement in the bond-line quality, as described in section 5.4, but this did not translate into a measurable strength benefit and cannot be attributed with certainty to the tape since the bond-line metric is confounded by the potential for damage during the sectioning process which is difficult to distinguish from the true bond strength. Figure 5.13 isolates the tape effect by grouping the Tranche 3 results at each tension level, showing both maximum force and first-event force side by side. The tape effect on maximum force is small and inconsistent in direction, whereas the first-event force is consistently lower with tape at low and medium tension.




Figure 5.12: Tranche 3 maximum force, measured effective area, measured-envelope stress and apparent system stiffness. Error bars show ± 1 standard deviation where coupon-level variation exists.




Figure 5.13: Tape-effect comparison for Tranche 3: maximum force (left) and first-event force (right) grouped by recorded tension level, with and without shrink tape. Error bars show ± 1 standard deviation.

The tape results point in different directions for terminal strength and the onset of detected damage. For measured-envelope ultimate stress, the pooled taped and untaped means differ by only approximately -28 MPa, and the factorial tape term is correspondingly small ($p = 0.87$, partial $\eta^2 = 0.001$). For first-event force, however, an additive model that adjusts for tension estimates a tape coefficient of -1.49 kN, corresponding to an earlier first event in the taped specimens; its 95 % confidence interval is -2.99 kN to 0.01 kN. To express this estimate relative to the unexplained scatter in the data, the fitted tape coefficient was divided by the model residual standard deviation of approximately 1.96 kN, giving a residual-standardised effect size of $d \approx -0.76$. The residual standard deviation describes the specimen-to-specimen variation in first-event force remaining after tension has been accounted for. The model therefore predicts a first-event force for taped specimens approximately 0.76 residual standard deviations lower than that for untaped specimens. This effect size is interpreted similarly to Cohen's d , although its denominator is the regression model's residual standard deviation rather than the raw pooled within-group standard deviation. The confidence interval nevertheless includes zero, and the interaction pattern indicates that the reduction occurs principally at low and medium tension rather than at high tension. Independent tube replication would be required before the result can be interpreted as a repeatable tape effect.

5.6 First Detected Damage

As described in section 5.2.4, an automated event detection algorithm was used to attempt to monitor the progressive failure behaviour of the hoop samples. The number of detected events varied according to the exact parameters used in the algorithm, but the first detected event was found to be a robust metric that could be compared across configurations and tranches.

Three mechanisms were identified based on observations of the tensile testing: liner failure, outer-layer delamination and inner-layer or phenolic delamination. These, respectively, refer to the liner cracking, the outer carbon layer delaminating from the inner layers, and a failure at the carbon-phenolic interface. In all cases, the first failure was either liner cracking or phenolic delamination, with liner cracking being overwhelmingly the most common first failure mode. Across the campaign, the first detected event occurred at approximately 0.67 of maximum force. In combination with visual observations during the testing, it was judged that the liner was generally already cracked or had begun to separate from the overwrap before terminal load was reached.

In Tranche 2, the first-event force differed between configurations ($p = 0.00078$, $\eta^2 = 0.70$). The corresponding measured-envelope first-event stress separation was stronger ($p < 0.0001$, $\eta^2 = 0.81$), but is secondary and partly reflects the shared area denominator. The confounding of the Tranche 2 factors prevents either result from being attributed to any one factor.

In Tranche 3, no first-event-force term resolved: tension gave $p = 0.19$ ($q = 0.42$), tape $p = 0.051$ ($q = 0.31$) and the interaction $p = 0.13$ ($q = 0.38$). In the secondary measured-envelope stress analysis, the interaction had $p = 0.042$ but $q = 0.30$; tension ($p = 0.068$, $q = 0.35$) and tape ($p = 0.19$, $q = 0.42$) were also unresolved.

Figure 5.14 shows the primary first-event force beside the measured-envelope first-event stress sensitivity for Tranche 3. Their differing group rankings expose the effect of the assigned area denominator.

5.7 Stiffness

Apparent system stiffness, defined as the derivative of the fitted force versus slack-corrected extension curve and reported in kN mm^{-1} , is independent of area, so these results are unaffected by the change from nominal to measured geometry. Both stiffness measures are summarised by configuration in table 5.6.

In Tranche 2, stiffness differed strongly between configurations ($p < 0.0001$, $\eta^2 = 0.85$), with group means ranging from 8.63 kN mm^{-1} to 10.16 kN mm^{-1} .

In Tranche 3, stiffness increased with tension ($p = 0.00065$, $q = 0.023$), the group means being approximately 8.1 kN mm^{-1} to 8.6 kN mm^{-1} at low tension, 8.6 kN mm^{-1} to 8.8 kN mm^{-1} at medium tension and 9.1 kN mm^{-1} to 9.3 kN mm^{-1} at high tension. Neither tape nor the interaction of tape and tension was resolved.

The stiffness discussed is an apparent structural stiffness of the whole system, comprising the specimen, the fixture, the seating and the machine, and not a material modulus. The area-independent gauge force–strain slope is the primary gauge comparison. Its tension term was unresolved ($p = 0.12$, $q = 0.38$), as were tape ($p = 0.93$, $q = 0.93$) and the interaction ($p = 0.93$, $q = 0.93$). The measured-area apparent modulus had group means of 215 GPa to 289 GPa ; its tension trend was likewise unresolved after Benjamini-Hochberg correction




Figure 5.14: Tranche 3 first detected events for individual coupons, plotted against mean gauge strain at the detected event. Force is the primary endpoint (left); measured-envelope stress is the secondary area-normalised sensitivity (right). Colours distinguish the six tension–tape configurations.

Table 5.6: Group mean apparent system stiffness and gauge-derived apparent modulus by configuration. Strain gauges were fitted to the Tranche 3 coupons only, so no gauge modulus is available for Tranche 2.

Configuration	<i>n</i>	System stiffness (kN mm ⁻¹)		Gauge modulus (GPa)	
		Mean	SD	Mean	SD
<i>Tranche 2</i>					
C1	4	8.63	0.47	—	—
C2	4	9.55	0.18	—	—
C3	4	9.72	0.03	—	—
C4	4	10.16	0.17	—	—
C5	4	9.61	0.10	—	—
<i>Tranche 3</i>					
Low Tension No Tape	4	8.13	0.52	242	31
Low Tension Shrink Tape	5	8.62	0.29	215	58
Medium Tension No Tape	5	8.78	0.15	236	34
Medium Tension Shrink Tape	5	8.57	0.59	270	24
High Tension No Tape	5	9.28	0.23	265	33
High Tension Shrink Tape	5	9.08	0.21	289	39

($p = 0.035$, $q = 0.30$), while tape was unresolved ($p = 0.43$, $q = 0.83$). These values exceed the fibre-direction modulus attainable by this laminate at any realisable fibre volume fraction, and they are calculated on the gross carbon envelope, which includes voids and resin-rich regions in the denominator and therefore biases the result towards lower values. The gauge slope is consequently an index of the local response at the split point rather than a laminate modulus, as set out in section 4.3.2, and only the relative ordering between configurations is interpreted below. That a strong tension effect appears in system stiffness but only a weak one in gauge modulus may imply that the effect lies largely in seating, fixture engagement and pre-stress, or perhaps in a modification to the bending behaviour rather than in the overwrap itself, though this requires further investigation. Of these contributors, only the initial seating take-up is corrected for (by the slack correction); the fixture engagement, pre-stress and the elastic compliance of the fixture and machine remain inside the measured stiffness for every configuration. A genuine change in the effective modulus of the overwrap is not excluded by these data, and would not require the modulus of the carbon fibre to change: higher winding tension straightens the tow and reduces fibre waviness and off-axis scatter, which is expected to raise the effective hoop modulus of the wound laminate for a fixed material modulus. The gauge measures the response of the overwrap as wound, so it cannot separate a change in fibre alignment from a change in the fibre itself.

The system stiffness and the area-independent gauge force–strain slope are moderately related across the 29 coupons ($r = 0.44$, $p = 0.017$), but the centred within-configuration association is unresolved ($r = 0.31$, $p = 0.106$). The six configuration means correlate strongly ($r = 0.85$, $p = 0.030$, $n = 6$), so the raw association is principally evidence of group ordering rather than coupon-level agreement. The measured-area modulus gives a weaker and unresolved correlation between the six means ($r = 0.48$, $p = 0.34$).

5.8 Mechanical Response vs Void Content

Figure 5.15 compares detected void area fraction with maximum force in the left-hand panel and measured-envelope ultimate stress in the right-hand panel.

The primary force comparison provides no evidence that configurations with greater detected void area carried a lower terminal load. The configuration-mean correlations are $r = 0.71$ in Tranche 2 ($p = 0.18$, $n = 5$), $r = 0.17$ in Tranche 3 ($p = 0.75$, $n = 6$), and $r = -0.11$ when all eleven configurations are pooled ($p = 0.75$). These estimates are descriptive because both the mechanical and microscopy observations are clustered by manufactured tube or region.

The correlations with measured-envelope stress in the right-hand panel are more negative ($r = -0.62$ in Tranche 2, $r = -0.29$ in Tranche 3 and $r = -0.40$ pooled). However, measured carbon-envelope area enters the stress denominator and is obtained from the same segmented micrographs as the void area fraction. The right-hand panel therefore cannot isolate a void–strength mechanism. A matrix-dominated test, such as short-beam shear (ILSS, ASTM D2344), combined with independently validated fibre and void volume fractions from resin burn-off or acid digestion, would address the porosity question more directly.




Figure 5.15: Detected optical void area fraction versus maximum force (primary, left) and measured-envelope ultimate stress (secondary sensitivity, right) by configuration. Error bars show the available within-group mechanical and between-image microscopy variation.

Chapter 6

Discussion

Overall, it was found that processing conditions were associated with changes in gross geometry, detected void content, interface appearance, apparent system stiffness and the onset of the detected force discontinuities, whereas maximum force was less sensitive to configuration than several of these intermediate responses. What the evidence shows, then, is that the manufacturing choices influence how the ring seats, consolidates and accumulates damage more clearly than they establish any repeatable change in the terminal capacity of the carbon overwrap.

6.1 Research Question 1: Microscopic Features

How do observed microscopic features relate to the observed mechanical response, and can these features be used to rescale and interpret the mechanical results?

The microscopy results in section 5.4 suggest that tape becomes more effective once higher winding tension has seated the tow, but they do not establish a monotonic process effect on gross area. This mechanism remains a hypothesis until it is reproduced across more sections and independently manufactured regions. This is somewhat different from the results reported by Bianchi et al., who reported reduced sample thickness (with an expected reduction in gross area resulting) [38]. The present campaign also failed to resolve a corresponding benefit in terminal force. The remaining void levels are also consistent with the wider limitation of OOA processing: shrink tape supplies compaction but neither the elevated pressure of an autoclave nor, by itself, a pathway for removing trapped gases before matrix gelling [31–33].

The near-complete close-interface fraction in the taped sections is consistent with external compaction closing the visible carbon–phenolic gap, but no effect on bond strength could be observed in the mechanical response. The current metric does not separate sectioning damage and a genuinely open interface, leaving the interpretation of the interface appearance as a qualitative observation rather than a quantitative measure of bond quality. Since no relationship with the first-event response was resolved either, it can only be used as a descriptive quality indicator.

6.2 Research Question 2: Mechanical Response

How do the recorded tension and tape conditions relate to apparent stiffness, first detected mechanical failure, maximum force, and gross-area hoop stress?

6.2.1 Maximum Force and Area-Based Sensitivities

The force results in section 5.5 do not support a consistent terminal-capacity benefit from tape, while the measured-envelope stress ranking changes with the assigned area. Process selection therefore rests first on force, while the area-normalised stress is useful for asking whether a configuration carries load efficiently, but cannot distinguish improved material performance from a smaller denominator. There may be some utility when considering the overall structural thickness of the final rocket motor, but this is not demonstrated since the sample does not establish how tows in an actual lay-up would interact.

6.2.2 Apparent Stiffness and Gauge Response

The mismatch between the resolved system-stiffness trend and the weaker gauge response is consistent with tension changing ring seating, pre-stress, contact, sample bending or fixture engagement more strongly than it changes the overwrap itself. The present measurements cannot identify which contribution dominates or establish a change in material modulus.

It is expected that the effective modulus of the overwrap would increase with higher tension, since the tow would be laid down with less fibre waviness; increasing fibre waviness has been shown to reduce both tensile strength and modulus [58]. The strain gauges read the overwrap as wound, so they cannot separate a change in fibre alignment from a change in the fibre, and an alignment mechanism and a seating mechanism could predict the same measurement. Fibre alignment would also vary from coupon to coupon through local tow placement, which is consistent with the coupon-level correlation between the two stiffness measures reported in section 5.7. Separating the two would require fibre-angle or waviness measurement on the polished sections, taken alongside a stiffness measurement that is free of fixture compliance. This is expected to be extremely challenging to achieve experimentally.

Additionally, a test setup that provides even loading through the whole ring, rather than a localised bending/membrane stress superposition, would help to close the gap between the measured and true stiffness responses. Rosato and Grove describe a “Hydraulic Tension Ring Test”, which uses hydraulic pressure to load the ring evenly and can provide more representative estimates of hoop tensile strength and modulus than the split-D. However, this test setup requires more complex equipment, careful sealing and maintenance [1].

6.2.3 First Detected Mechanical Discontinuity

The first-event results show that configurations with similar terminal capacity can follow different damage paths. Their design significance depends on identifying the physical event behind each force drop, rather than on obtaining another statistical comparison of the algorithmic event alone.

The observations from the mechanical testing indicate three causes of discontinuity: liner splitting, outer-layer delamination and separation at the inner layer or interface. Of these, the liner splitting was by far the most common cause of first failure. In a solid rocket motor, such a failure could potentially be catastrophic, as it would allow hot combustion products to directly impinge on the carbon overwrap, weakening the resin and causing motor failure. Outer-layer delamination is unlikely to occur in practice, as the hoops had an available “free end” to delaminate, which would not be present in a motor. Separation at the bond line is possibly a concern, but would be far less catastrophic than liner splitting, and the effect may be mitigated by the even loading of the motor as opposed to the uneven load case present in a split-D.

6.3 Research Question 3: Implications for the Target Motor

How do the observed mechanical responses inform expected mechanical performance of the target motor?

The mechanical responses measured on the coupons inform the expected performance of the target motor only through projection, because the test geometry, diameter, dome and boss load paths, and thermal exposure of the motor are not represented in the split-D coupon. That projection is developed in chapter 7.

6.4 Research Question 4: Test-Method Development

What measurement and data-reduction methods are required to obtain a defensible comparison between configurations from split-D coupons, and what do they limit?

The comparison rests on methods developed alongside the campaign itself: a purpose-built split-D fixture sized by a screening capacity check (appendix A), slack correction of the cross-head history and a fitted system stiffness (section 5.7), per-channel strain-gauge validity screening (section 5.2), and an algorithmic force-drop detector whose threshold sensitivity was checked (section 5.2.4). Each method also bounds what can be claimed: the slack-corrected stiffness is a system response rather than a material modulus, the gauges read a local response with bending superimposed, and the first detected event is a discontinuity rather than an identified failure mode.

6.5 Research Question 5: Candidate Process Condition

Which manufacturing configuration provides the most defensible candidate for subsequent motor development?

The evidence supports a trade-off rather than a single optimum. Some applied tension is required for effective tow placement — the low-tension setup was extremely hard to work with and had very poor tow placement, while the medium and high tension settings were far more controllable. This absence of a simple monotonic tension–strength relationship is consistent with the literature, because greater winding tension can improve consolidation while also increasing fibre damage, resin squeeze-out and the redistribution of residual stresses through the laminate [1, 34]. High tension with functional tape is the stronger consolidation candidate, while high tension without tape produced the best terminal-force response on Tube 4. Neither condition is ready for selection without independent tube replication, cure verification, assessment of the suspected delamination and representative-diameter pressure testing. The motor calculation in chapter 7 therefore uses a conservative response range rather than the most favourable group.

6.6 Under-Resolved Relationships and Directions for Further Investigation

The unresolved signals reported in chapter 5 identify where the next experimental effort would have the greatest return. The following subsections focus on the evidence needed to resolve them rather than repeating their statistical estimates.

6.6.1 Tension Effect on Terminal Force

The same split-D geometry as was used in the preceding experiment description is suitable for testing the possible tension effect on terminal force, but the next experiment must replicate winding conditions across independent tubes and define a practically important force difference in advance. This would distinguish a repeatable minimum-tension requirement from region-specific tow placement, while measured-envelope stress should remain a secondary efficiency comparison. It is expected that the use of a split-D setup decreases the terminal force value somewhat from the ideal case, since the bending induced will tend to stress the sample across the layers of the composite, inducing inter-layer delamination and failure. A pressurised-ring test as described in section 6.2.2 would likely provide a more representative terminal-force estimate.

6.6.2 Tape Effect on First Detected Damage

The possible tape-associated reduction in first-event force could represent an important consolidation–damage trade-off, but the algorithm used in this study does not identify what failed. Independent replication should therefore be paired with physical event identification through interrupted testing, synchronised observation, acoustic-emission monitoring or post-test examination. The discussion of whether liner failure is stress- or strain-driven, and therefore whether the use of tape could be increasing the stiffness resulting in earlier failure by inducing more bending follows in section 6.6.4.

6.6.3 Void Reduction at High Tension with Tape

The void reduction for the high-tension setup with tape compared to the non-taped setup can be checked relatively cheaply by sampling more polished sections from independently manufactured regions. Because its benefit did not appear in the fibre-dominated terminal-force result, the use of short-beam interlaminar-shear specimens or ± 45 degree wound tube specimens in tension or compression would test whether it instead improves a matrix- or interlaminar-dominated property. The lack of a corresponding terminal-force increase with void reduction does not contradict published void-knockdown studies, which report pronounced effects on interlaminar shear and flexural properties but do not generally comment on tensile performance [53, 54].

6.6.4 Method-Limited Questions

Two further questions are limited by the test method rather than by sample size, and no realistic number of additional split-D coupons would resolve them.

The porosity–strength relationship couples the measurement with the least sensitivity to porosity (fibre-dominated hoop tension) to the metric with the least validation (optically segmented void area). A more informative test

configuration would pair short-beam interlaminar shear specimens, where the matrix and void content dominate the response, with true volume-fraction measurement by acid digestion or burn-off rather than optical area fraction. This combination would test the mechanistic expectation that higher void content degrades the matrix-dominated response. The burn-off option would be recommended for a future student-led project to avoid the use of strong acids which would result in a significantly more complex safety setup.

The question of whether liner cracking is strain-keyed or stress-keyed is blocked by the absence of physically labelled failure events. The current force-drop detector identifies that a discontinuity occurred but not what failed, so the crack-initiation mechanism cannot be correlated with the mechanical state at the event. In-situ observation—either through a synchronised close-up video record or by acoustic emission source location—would be needed to separate liner damage from overwrap damage in the force record. Given that liner splitting was the most common observed cause of first failure and the most consequential for motor integrity, resolving this question is a priority for any subsequent campaign. It is also to be noted that the loading scenario with the split-D fixture is not representative of the true loading scenario, and that the bending induced by the test setup is likely to result in far earlier liner failure than would be expected on an evenly-loaded pressure vessel. A pressurised-ring test could provide a more representative loading scenario and is therefore expected to reduce the dominance of liner cracking as a primary failure mode (section 6.2.2).

6.7 Liner Drying and Blistering

The four recorded Tranche 2 drying treatments removed between 0.54 % and 0.58 % of the initial liner mass, with a mean of 0.56 %, and the separate full-length liner lost 0.59 %. The consistency of these before and after weighings shows that the specified heating treatment removes a repeatable amount of volatile mass from the as-received phenolic material. The mass balance does not identify what was removed, but it is consistent with absorbed moisture from storage.

The baking of the liner did not prevent the blistering of Tube 4 in Tranche 3, indicating that the blistering is not solely a result of absorbed moisture. Two remedies remain untried: alternative liner materials such as Phenolic Formaldehyde Cotton Canvas (PFCC), which the author has found to exhibit less cracking in rocket-motor nozzles, and a slower heating and cooling profile to reduce the thermal shock and the associated cracking and blistering.

6.8 Validity and Generalisability

The confidence that can be placed in each result depends on how directly it was measured. Maximum force was measured directly and is therefore the most secure result of the campaign — unfortunately it is also the least sensitive to configuration.

The microscopy-derived quantities carry the uncertainty of a small number of unvalidated sections, and the split-D load path applies a non-uniform load that is not identical to the pressure loading of a motor. Comparisons between configurations are strongest within Tranche 3, where the tension levels were varied deliberately, and weakest within Tranche 2, where several factors changed together; the recorded tension also varied somewhat within each nominal Tranche 3 level. Finally, the coupons of Tranche 3 are subsamples of a single tube, and the test geometry is a hoop of smaller diameter without internal pressure, domes, bosses or thermal exposure compared to the intended motor application.

Taken together, these define the campaign as a screening study, and it succeeds on those terms. It establishes which manufacturing conditions change the response, which measurements are sensitive enough to detect those changes, and where the effort should be directed next; the tension effect on system stiffness and the void-content reduction at high tension with tape are both clear enough to carry forward. Confirming these as process capability or design allowables is the task of the following stage, which requires independently manufactured tubes and full-vessel hydrostatic testing.

Chapter 7

Preliminary Motor-Design Implications

This chapter attempts to project the hoop-coupon results back onto the planned Shearwater rocket motor. It is a sizing exercise based on equilibrium, and not a validated prediction of burst pressure: the calculation models neither the dome, the boss nor the nozzle attachment, and takes no account of thermal degradation, cure variation, the strength of helical coupons or the manufacturing scatter between tubes.

7.1 Pressure Line Loads

For a thin cylindrical shell of internal radius r under pressure p , the membrane line loads per unit axial length are:

$$N_{\theta} = pr \quad (7.1)$$

$$N_x = \frac{pr}{2} \quad (7.2)$$

where N_{θ} is the circumferential line load and N_x is the axial line load. Using the planned 116 mm liner inner diameter and 2 mm wall gives an outer radius of 60 mm at the carbon–liner interface. At the targeted 11 MPa operating pressure,

$$N_{\theta} \approx 0.660 \text{ kN mm}^{-1} \quad (7.3)$$

$$N_x \approx 0.330 \text{ kN mm}^{-1}. \quad (7.4)$$

The interface radius is used here because the carbon overwrap carries the membrane load outside the liner. This is a conservative reference-radius convention rather than an exact treatment of the pressure boundary. The pressure resultant acts on the 58 mm bore radius; using the 60 mm interface radius increases the calculated membrane demand by approximately 3.4 %, and therefore reduces the corresponding pressure-equivalent capacity by approximately 3.3 %. Retaining the interface radius avoids presenting that small geometric refinement as additional structural margin.

7.2 Netting Equilibrium

Let t_{90} be the total hoop-ply thickness and t_h the combined thickness of balanced $+\alpha$ and $-\alpha$ helical plies. With angle α measured from the vessel axis and equal fibre-direction design stress S , netting equilibrium gives

$$N_{\theta} = S(t_{90} + t_h \sin^2 \alpha) \quad (7.5)$$

$$N_x = S t_h \cos^2 \alpha \quad (7.6)$$

These equations assume that the fibres carry the membrane loads, that the matrix maintains the architecture, that the balanced helical pair carries no resultant shear load, and that the cylindrical section is free of local discontinuities.

7.3 Force-Based Input and Area Sensitivities

The split-D force result is the primary experimental input. A typical two-section maximum force near 19.6 kN corresponds to approximately 9.8 kN per loaded wall, or 1.54 kN mm^{-1} once divided by the nominal 6.35 mm band width. That is approximately 2.3 times the operating hoop line load of 0.660 kN mm^{-1} , although the comparison takes no account of the axial loading, the helical architecture, the pressure load path, or any qualification factor.

Netting calculations require a stress-like input, so two conversions are carried alongside this force check. Dividing the group mean force by the fixed nominal area gives configuration means of approximately 1740 MPa to 2070 MPa, while dividing by the measured envelopes gives approximately 1650 MPa to 2380 MPa. The latter range is an area-efficiency sensitivity, not a material allowable. A central screening value of 2000 MPa lies within both ranges, but is neither a characteristic strength nor a statistical lower bound.

7.4 Candidate Layups

The planned winding angle is $\alpha = 41.81^\circ$, based on a geodesic path over the chosen isotenoid dome geometry. The nominal ply thickness is taken as one third of the split-D coupon's nominal 0.8 mm three-layer thickness, giving 0.267 mm per ply. Two candidate layups are considered, with either one or two hoop plies and one balanced pair of $\pm 41.81^\circ$ helical plies. The microscopy-derived three-layer thicknesses ranged from 0.811 mm to 0.936 mm across the measured Tranche 3 configurations. The 0.8 mm nominal value is therefore slightly below every measured configuration and gives a conservative thickness-based capacity estimate within the assumptions of the netting model. For these candidate configurations, the netting calculations give:

Table 7.1: Preliminary cylindrical-section netting sensitivity

Candidate	Nominal thickness	Pressure-equivalent range	Capacity-to-operating-pressure ratio
[90, 90, $\pm 41.81^\circ$]	1.07 mm	17.2 MPa to 20.4 MPa	1.6–1.9
[90, $\pm 41.81^\circ$]	0.80 mm	14.6 MPa to 17.4 MPa	1.3–1.6

These values are very preliminary screening figures, rather than an attempt at a prediction of burst pressure. The calculation assigns the hoop-coupon stress range equally to the hoop and helical plies, although no helical coupons were tested. This assumption is particularly important for the four-ply candidate because its calculated capacity is governed by the axial contribution of the helical pair.

Under the assumptions made, the four-ply candidate remains above a capacity-to-operating-pressure ratio of 1.5 across the stated stress-input sensitivity, and is governed by axial equilibrium. The three-ply candidate is governed by hoop equilibrium and crosses a ratio of 1.5 only near the upper end of the stress range. However, this ratio must not be read as a qualified factor of safety because model bias, temperature, manufacturing variability and independent-tube strength are not included.

The axial-equilibrium constraint governing the four-ply candidate is directionally consistent with metallic motor cases that use a circumferential groove to promote a controlled, axially governed failure rather than uncontrolled hoop rupture. This may represent a preferable failure-mode objective for the final casing. However, the present netting calculation does not model the resulting fracture path from a failure, so that behaviour would require separate detailed analysis and experimental validation.

Overall, this makes [90, 90, $\pm 41.81^\circ$] a more promising preliminary sizing candidate for further evaluation through CLT and FEA.

7.5 Damage-Onset Scenario

The first detected discontinuity commonly occurred near two-thirds of maximum force, although the fraction varied by configuration. Applying that fraction would reduce the pressure-equivalent screening range substantially, but it is not treated as a motor limit because the physical event is unidentified and the candidate motor includes a different load path and helical plies.

7.6 Design Decision and Required Validation

The calculation provides an initial indication that the four-ply candidate layup may be most promising for the Shearwater rocket motor, but it is clear that significant further testing and simulation will be necessary before progressing to the actual motor. Some of the remaining uncertainties are the limited experimental dataset, which fails to provide data for helical layups, the complexity of the motor domes and mechanical interfaces, and the lack of characterisation data for the material over a wide temperature range, as well as considerations regarding the liner material with regard to the prevention of blistering. Any motor design decision must also be made against the applicable proof, burst, thermal and production requirements, which are not established in the current study.

Chapter 8

Conclusions and Recommendations

This thesis set out to characterise an out-of-autoclave carbon-towpreg process co-cured onto a paper-phenolic liner, and to apply the obtained results to a small solid rocket motor casing. The evidence base combines process observations, fifteen final microscopy summaries and forty-nine split-D coupon tests, distributed across a confounded screening tranche and a directed set of regions on one shared tube. The principal research question can therefore be answered only at the level of the structure that was tested: recorded tension and functional shrink tape were associated with microstructural and mechanical differences, but the campaign does not establish repeatable manufacturing effects at the level of a population of tubes.

The single clearest finding is a separation between terminal capacity and the path to failure. Mean maximum forces in Tranche 2 sat in a narrow band near 19 kN to 21 kN despite configurations chosen to drive the response apart, and in Tranche 3 only the low-tension regions fell noticeably below the rest, by roughly 2.5 kN in the group means. Across the same specimens the gross carbon envelope, the detected void content, the apparent system stiffness and the first detected force discontinuity all moved with the manufacturing variables. The processing choices therefore influence how a ring seats, consolidates and accumulates damage more clearly than they set the load at which the carbon finally yields. That distinction matters for a liner-supported motor casing, where manufacturing quality and early damage can govern before the carbon ruptures.

The clearest microscopy contrast appeared at high tension, where functional shrink tape reduced the detected optical void area fraction from approximately 12.1 % to 9.0 %; the pattern is consistent with tape becoming effective once the tow is already seated, but it rests on one or two sections per region and on a segmentation method with no independent validation, so it is suggestive rather than confirmed. Both levels remain high relative to the levels obtainable with the use of an autoclave, so whether this OOA route can reach structurally acceptable void levels at the motor scale is itself still open. It may be that with the use of more advanced manufacturing equipment, improved processing and handling and further development it is possible to reduce the void content to acceptable levels for unmanned systems such as sounding rockets and missiles, though it is not considered likely by the author that levels acceptable for human-rated vehicles could be achieved without significant modifications to the material and processing techniques.

The manufactured gross envelope also varied without a resolved tension or tape pattern. Because that sparsely sampled area strongly changes the calculated stress ranking, maximum force remains the primary terminal-load result and measured-envelope stress a secondary efficiency sensitivity.

Within Tube 4, recorded tension was the only Tranche 3 term to survive correction for multiple comparisons in the coupon-level factorial analysis ($q = 0.023$), with the higher-tension regions exhibiting greater apparent system stiffness. This provides evidence of a statistically resolved within-tube association, but does not establish that the effect is repeatable between independently manufactured tubes. The area-independent gauge force-strain slope did not resolve the same tension effect, while the measured-area apparent modulus is only a secondary sensitivity. Since every coupon of Tranche 3 was cut from a single tube, the campaign was not able to resolve whether the effective modulus of the overwrap itself changes with tension. The application of shrink tape produced no consistent increase in maximum force or first detected failure, even though it improved the microscopic appearance of the interface.

Four Tranche 2 liners lost a mean of 0.56 % of their initial mass during the recorded drying treatment, and the separate full-length liner lost 0.59 %; the removable mass is repeatable, which is consistent with the presence of a volatile compound (likely water) that may contribute to the observed blistering, but the drying treatment failed

to prevent blistering from reoccurring in the medium-tension region of Tube 4.

A force-based line-load check gives a typical split-D capacity approximately 2.3 times the operating hoop line load. Using a nominal ply thickness of 0.8 mm/3, a separate netting sensitivity gives a capacity-to-operating-pressure ratio of approximately 1.6–1.9 at 11 MPa for the preliminary $[90, 90, \pm 41.81^\circ]$ cylindrical-section candidate. The three-ply candidate gives approximately 1.3–1.6 and exceeds 1.5 only near the upper end of the stress-input range. These results retain a slightly conservative interface-radius convention and transfer the hoop-coupon stress range to untested helical plies. They are therefore conditional screening projections, not factors of safety, proof levels, burst ratings or design allowables.

Among the tested configurations, High Tension | Shrink Tape gives the clearest consolidation signal, while High Tension | No Tape gives the highest and least variable terminal-force result in Tranche 3; these are candidate directions for subsequent validation.

The next campaign should concentrate on the possible tension effect on terminal force, the tape-associated shift in damage onset and the high-tension consolidation signal. These questions require independently wound tubes and physically identified damage events; the porosity question instead requires a matrix-dominated test paired with validated volume-fraction measurements. The present campaign therefore supports continued development, but not production readiness or design allowables.


Nevertheless, the campaign establishes a practical basis for that development: the co-cured liner–overwrap concept was manufactured and tested, the hoop coupons sustained terminal loads above the operating line-load estimate, and the results identified high tension as a promising route to greater stiffness and high tension with shrink tape as the clearest route to improved consolidation. These findings narrow the process space and define the evidence that must now be obtained through replicated manufacture, validated porosity measurements and representative pressure testing. The work therefore advances the concept from an uncharacterised manufacturing proposition to a test-supported candidate for focused motor-scale validation.

Appendix A

Split-D Fixture Design

The split-D fixture used in this campaign was designed and built in-house. This appendix records the assembly drawing and the loading analysis performed before testing began.

A.1 Assembly Drawing



A.2 Loading Analysis

The following calculations are a first-order analysis of the principal possible failure modes of the testing jig. The values given are therefore preliminary estimations of the true failure loading.

For all the following calculations, a yield tensile strength of 230 MPa was assumed for 303 stainless steel [59].

The net tensile section at the through-bolt comprises four 6.9 mm × 7.5 mm sections. Its combined area is:

$$A_{\text{net}} = 4 \cdot 6.9 \cdot 7.5 = 207 \text{ mm}^2, \quad (\text{A.1})$$

giving a tensile-yield estimate of:

$$F_{\text{net}} = \sigma_y A_{\text{net}} = 230 \cdot 207 = 47.6 \text{ kN}. \quad (\text{A.2})$$

The total projected bearing area in the two 7.5 mm side plates is

$$A_b = 2 \cdot 16 \cdot 7.5 = 240 \text{ mm}^2, \quad (\text{A.3})$$

which gives the first-order bearing-yield estimate (conservatively assuming bearing yield is the same as tensile yield):

$$F_b = \sigma_y A_b = 230 \cdot 240 = 55.2 \text{ kN}. \quad (\text{A.4})$$

This value is not intended as a true estimation of bearing failure, but a very conservative check.

Each M16 fastener is nominally loaded in double shear. The threaded portion passes through both shear planes. For an M16 property class 8.8 bolt with the shear plane passing through the threaded portion, the design shear resistance is 60.3 kN per shear plane [60]. The double-shear resistance is therefore

$$F_{v,\text{thread}} = 2(60.3 \text{ kN}) = 120.6 \text{ kN}. \quad (\text{A.5})$$

The plate net-section estimate of 47.6 kN therefore governs over the bearing-yield estimate and the fastener design shear resistance in this screening calculation.

The jig was deliberately designed so that the holes would be far from the edges or well-supported by the presence of the sample to minimise any potential chances of tear-out failure.

The values given here are only first-order estimations, but were sufficiently far above the achievable values for the tested samples that testing proceeded as planned. A full FEM analysis including the effects of material hardening, ageing and fatigue would be recommended for a jig that is planned to be in use for a significantly greater lifespan than the one described here.

Appendix B

Microscopy Results

The fifteen final polished sections provided the raw microscopy record from which the configuration porosity and interface estimates were derived. Each preserved image carries its source filename, the pixel scale, a refinement flag identifying whether the segmentation parameters were adjusted from the calibrated defaults, the carbon and phenolic mask areas, the detected carbon-void counts and area fractions, and the interface-gap statistics.

Table B.1 summarises the numerical results for all fifteen sections, and the following pages reproduce the raw optical appearance alongside the segmentation and thickness-axis overlay for each configuration, so that every reported quantity can be checked back against the image from which it was derived.

Table B.1: Summary of the fifteen polished-section measurements. Carbon area is the single-wall gross envelope; void fraction is the detected optical dark-area fraction within that envelope; close-interface fraction is the proportion of the carbon–phenolic boundary classified as closed.

Configuration	Section	Area (mm ²)	Void (%)	Close (%)
<i>Tranche 2</i>				
C1 No Tension No Tape	Beta	4.39	5.7	65.2
	Alpha	4.56	10.1	58.8
C2 Tension No Tape	—	4.41	9.7	91.5
C3 Tension Failed Tape	—	4.40	8.9	93.5
C4 Tension Failed Tape Sanded	—	5.83	10.0	99.3
C5 Tension Failed Tape Not Dried	—	6.15	10.9	92.7
<i>Tranche 3</i>				
High Tension No Tape	1	4.70	12.0	84.9
	2	4.85	12.2	79.2
High Tension Shrink Tape	1	4.51	9.0	99.8
Low Tension No Tape	1	4.55	11.7	74.9
Low Tension Shrink Tape	1	5.16	11.7	99.8
Medium Tension No Tape	1	5.97	10.9	89.8
	2	4.44	13.0	87.4
Medium Tension Shrink Tape	1	4.40	9.3	98.8
	2	4.50	13.1	99.0

B.1 Tranche 2 Sections




Figure B.1: C1 | No Tension | No Tape | Dried — Section Beta. Raw optical appearance (left) and segmentation overlay with thickness axis (right). Detected void fraction 5.7 %; close-interface fraction 65.2 %.




Figure B.2: C1 | No Tension | No Tape | Dried — Section Alpha. Detected void fraction 10.1 %; close-interface fraction 58.8 %.




Figure B.3: C2 | Tension | No Tape | Dried. Detected void fraction 9.7 %; close-interface fraction 91.5 %.




Figure B.4: C3 | Tension | Failed Shrink Tape | Dried. Detected void fraction 8.9 %; close-interface fraction 93.5 %.




Figure B.5: C4 | Tension | Failed Shrink Tape | Dried | Sanded. Detected void fraction 10.0 %; close-interface fraction 99.3 %.




Figure B.6: C5 | Tension | Failed Shrink Tape | Not Dried. Detected void fraction 10.9 %; close-interface fraction 92.7 %.

B.2 Tranche 3 Sections




Figure B.7: High Tension | No Tape — Section 1. Detected void fraction 12.0 %; close-interface fraction 84.9 %.




Figure B.8: High Tension | No Tape — Section 2. Detected void fraction 12.2 %; close-interface fraction 79.2 %.




Figure B.9: High Tension | Shrink Tape. Detected void fraction 9.0 %; close-interface fraction 99.8 %.




Figure B.10: Low Tension | No Tape. Detected void fraction 11.7 %; close-interface fraction 74.9 %.




Figure B.11: Low Tension | Shrink Tape. Detected void fraction 11.7 %; close-interface fraction 99.8 %.




Figure B.12: Medium Tension | No Tape — Section 1. Detected void fraction 10.9 %; close-interface fraction 89.8 %.




Figure B.13: Medium Tension | No Tape — Section 2. Detected void fraction 13.0 %; close-interface fraction 87.4 %.




Figure B.14: Medium Tension | Shrink Tape — Section 1. Detected void fraction 9.3 %; close-interface fraction 98.8 %.




Figure B.15: Medium Tension | Shrink Tape — Section 2. Detected void fraction 13.1 %; close-interface fraction 99.0 %.

Appendix C

Use of AI Tools

AI assistants (GLM-5.2, Kimi K3, Claude Opus 5, and GPT-5.6 Sol) were used during this work to support the statistical analysis, the processing of the microscopy images, and drafting and editing of text. All outputs were reviewed and validated by the author.

Bibliography

- [1] D. V. Rosato and C. S. Grove. *Filament Winding: Its Development, Manufacture, Applications, and Design*. 1st. Interscience Publishers, 1964. 360 pp. ISBN: 978-0-470-73338-7.
- [2] *Northrop Grumman Propulsion Products Catalog*. URL: <https://cdn.northropgrumman.com/-/media/wp-content/uploads/NG-Propulsion-Products-Catalog.pdf?v=1.0.0> (visited on 14/01/2026).
- [3] *Space Launch System Solid Rocket Booster* - NASA. Section: Humans in Space. URL: <https://www.nasa.gov/reference/space-launch-system-solid-rocket-boosters/> (visited on 14/01/2026).
- [4] Tae-Ho Lee. "Review of the Composite Materials Application to the Solid Rocket Motor Cases". In: *Composites Research* 25.3 (2012), pp. 82–89. ISSN: 2288-2103. DOI: 10.7234/kscm.2012.25.3.082. URL: <https://koreascience.or.kr/article/JAK0201220962920320.page> (visited on 04/03/2026).
- [5] George P. Sutton and Oscar Biblarz. *Rocket propulsion elements*. 7th ed. New York: John Wiley & Sons, 2001. 751 pp. ISBN: 978-0-471-32642-7.
- [6] Kelly Mathesius and R. Hansman. "Manufacturing Methods for a Solid Rocket Motor Propelling a Small, Fast Flight Vehicle". MSc thesis. Massachusetts Institute of Technology, 05/06/2019. URL: https://www.researchgate.net/publication/333640120_Manufacturing_Methods_for_a_Solid_Rocket_Motor_Propelling_a_Small_Fast_Flight_Vehicle.
- [7] K. Tsiolkovskiĭ and Mikhail Iakovlevich Marov. *Selected works*. 2nd ed. OCLC: 271810261. Moscow: Nauka, 2006. ISBN: 978-5-02-034126-5.
- [8] S. E. Twichell and R. B. Keller. *Solid rocket motor internal insulation*. NTRS Author Affiliations: Hercules, Inc., Wilmington, NASA Lewis Research Center NTRS Report/Patent Number: NASA-SP-8093 NTRS Document ID: 19770023227 NTRS Research Center: Legacy CDMS (CDMS). 01/12/1976. URL: <https://ntrs.nasa.gov/citations/19770023227> (visited on 13/02/2026).
- [9] Alain Davenas. *Solid rocket propulsion technology*. 1st English ed. Oxford New York: Pergamon Press, 1993. ISBN: 978-0-08-040999-3.
- [10] Latrobe Specialty Steel Company. *Lescalloy D6AC VAC-ARC High Strength Alloy Steel Data Sheet*. Latrobe, PA: Latrobe Specialty Steel Company, 09/2008.
- [11] S. T. Peters. *Handbook of Composites*. Second Edition. Boston, MA s.l: Springer US Imprint: Springer, 1998. 1118 pp. ISBN: 978-0-412-54020-2 978-1-4615-6389-1. DOI: 10.1007/978-1-4615-6389-1.
- [12] R. A. Fuhrman. "The Fleet Ballistic Missile System: Polaris to Trident". In: *Journal of Spacecraft and Rockets* 15.5 (1978). eprint: <https://doi.org/10.2514/3.57317>, pp. 265–286. ISSN: 0022-4650. DOI: 10.2514/3.57317. URL: <https://doi.org/10.2514/3.57317> (visited on 05/03/2026).
- [13] *Vega qualification flight VV01*. 06/02/2012. URL: https://download.esa.int/docs/VEGA/Vega_PressKit_06-02-2012_EN.pdf (visited on 05/03/2026).
- [14] Mark E. Tobias et al. "Booster Obsolescence and Life Extension (BOLE) for Space Launch System (SLS)". In: Institute of Electrical and Electronics Engineers. NTRS Author Affiliations: Northrop Grumman Innovation Systems NTRS Report/Patent Number: M18-7142 NTRS Document ID: 20190002126 NTRS Research Center: Marshall Space Flight Center (MSFC). New York, NY, United States, 02/03/2019. URL: <https://ntrs.nasa.gov/citations/20190002126> (visited on 05/03/2026).
- [15] *Hexcel, HyPerComp debut Type 4 COPV for aerospace, space*. 20/02/2026. URL: <https://www.compositesworld.com/news/hexcel-hypercomp-debut-type-4-copv-for-aerospace-space> (visited on 22/02/2026).
- [16] Lyazid Bouhala et al. "Advancement in the Modeling and Design of Composite Pressure Vessels for Hydrogen Storage: A Comprehensive Review". In: *Journal of Composites Science* 8.9 (09/2024), p. 339. ISSN: 2504-477X. DOI: 10.3390/jcs8090339. URL: <https://www.mdpi.com/2504-477X/8/9/339> (visited on 13/02/2026).
- [17] Alih John Eko et al. "A review of type IV composite overwrapped pressure vessels". In: *International Journal of Hydrogen Energy* 109 (14/03/2025), pp. 551–573. ISSN: 0360-3199. DOI: 10.1016/j.ijhydene.2025.

- 02.108. URL: <https://www.sciencedirect.com/science/article/pii/S0360319925006731> (visited on 15/12/2025).
- [18] Alexander Air, Md Shamsuddoha and B. Gangadhara Prusty. "A review of Type V composite pressure vessels and automated fibre placement based manufacturing". In: *Composites Part B: Engineering* 253 (03/2023), p. 110573. ISSN: 13598368. DOI: 10.1016/j.compositesb.2023.110573. URL: <https://linkinghub.elsevier.com/retrieve/pii/S1359836823000768> (visited on 15/12/2025).
- [19] Pat B. McLaughlan, Scott C. Forth and Lorie R. Grimes-Ledesma. *Composite Overwrapped Pressure Vessels, A Primer*. NTRS Author Affiliations: NASA Johnson Space Center, Jet Propulsion Lab., California Inst. of Tech. NTRS Report/Patent Number: S-1046 NTRS Document ID: 20110008406 NTRS Research Center: Johnson Space Center (JSC). 01/03/2011. URL: <https://ntrs.nasa.gov/citations/20110008406> (visited on 12/02/2026).
- [20] S.T. Peters, ed. *Composite Filament Winding*. ASM International, 01/09/2011. ISBN: 978-1-62708-338-6. DOI: 10.31399/asm.tb.cfw.9781627083386. URL: <https://dl.asminternational.org/handbooks/book/153/Composite-Filament-Winding> (visited on 04/02/2026).
- [21] Zhiwen Meng et al. "Analysis of the Boss Structure of Type IV Composite Vessel for a High-Pressure Hydrogen Tube Trailer". In: *Sustainability* 16.12 (2024). ISSN: 2071-1050. DOI: 10.3390/su16125098. URL: <https://www.mdpi.com/2071-1050/16/12/5098>.
- [22] Oleksii Vambol et al. "Impact of Metal Inserts on Stress Distribution in the Polar Regions of Filament-Wound Composite Pressure Vessels". In: *Integrated Computer Technologies in Mechanical Engineering - 2024*. Ed. by Oleksii Lytvynov, Volodymyr Pavlikov and Dmytro Krytskyi. Cham: Springer Nature Switzerland, 2025, pp. 478–490. ISBN: 978-3-031-94852-7. DOI: 10.1007/978-3-031-94852-7_40.
- [23] Organisation du traité de l'Atlantique nord (NATO/OTAN), ed. *Design methods in solid rocket motors*. Rev. version 1988. AGARD lecture series 150. Neuilly-sur-Seine: AGARD, 1988. ISBN: 978-92-835-0454-2. URL: <https://apps.dtic.mil/sti/tr/pdf/ADA199356.pdf> (visited on 05/03/2026).
- [24] Dirk H. -J. A. Lukaszewicz, Carwyn Ward and Kevin D. Potter. "The engineering aspects of automated prepreg layup: History, present and future". In: *Composites Part B: Engineering* 43.3 (01/04/2012), pp. 997–1009. ISSN: 1359-8368. DOI: 10.1016/j.compositesb.2011.12.003. URL: <https://www.sciencedirect.com/science/article/pii/S1359836811005452> (visited on 13/02/2026).
- [25] *Mikrosam filament production line supports Amargo tank production*. 20/02/2026. URL: <https://www.compositesworld.com/news/mikrosam-filament-production-line-supports-amargo-tank-production> (visited on 22/02/2026).
- [26] *Robot Configurations for Filament Winding*. TANIQ. URL: <https://composites.taniqu.com/robot-configurations> (visited on 22/02/2026).
- [27] Kumar C. Jois et al. "Towpreg manufacturing and characterization for filament winding application". In: *Polymer Composites* 45.9 (2024). eprint: <https://aspepublications.onlinelibrary.wiley.com/doi/pdf/10.1002/pc.28311>, pp. 7893–7905. ISSN: 1548-0569. DOI: 10.1002/pc.28311. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1002/pc.28311> (visited on 05/03/2026).
- [28] Eduardo Szpoganicz et al. "Comparative analysis of NOL-ring tensile strength in towpreg and slit-tape for filament winding: Influence of resin viscosity, tack, and consolidation". In: *Composites Science and Technology* 264 (03/05/2025), p. 111123. ISSN: 0266-3538. DOI: 10.1016/j.compscitech.2025.111123. URL: <https://www.sciencedirect.com/science/article/pii/S0266353825000910> (visited on 05/03/2026).
- [29] Lei Zu et al. "Development of Filament Wound Composite Isotensoidal Pressure Vessels". In: *Polymers and Polymer Composites* 22.3 (03/2014), pp. 227–232. ISSN: 0967-3911, 1478-2391. DOI: 10.1177/096739111402200301. URL: <https://journals.sagepub.com/doi/10.1177/096739111402200301> (visited on 15/12/2025).
- [30] Lei Zu, Sotiris Koussios and Adriaan Beukers. "Design of filament-wound isotensoid pressure vessels with unequal polar openings". In: *Composite Structures* 92.9 (08/2010), pp. 2307–2313. ISSN: 02638223. DOI: 10.1016/j.compstruct.2009.07.013. URL: <https://linkinghub.elsevier.com/retrieve/pii/S0263822309002487> (visited on 15/12/2025).
- [31] T. Centea, L. K. Grunenfelder and S. R. Nutt. "A review of out-of-autoclave prepregs – Material properties, process phenomena, and manufacturing considerations". In: *Composites Part A: Applied Science and Manufacturing* 70 (01/03/2015), pp. 132–154. ISSN: 1359-835X. DOI: 10.1016/j.compositesa.2014.09.029. URL: <https://www.sciencedirect.com/science/article/pii/S1359835X14003108> (visited on 05/03/2026).
- [32] Tom Cender et al. "Void reduction during out-of-autoclave thermoset prepreg composite processing". In: International SAMPE Technical Conference. 01/06/2014.
- [33] Okunzuwa Austine Ekuase et al. "A Review on the Out-of-Autoclave Process for Composite Manufacturing". In: *Journal of Composites Science* 6.6 (13/06/2022), p. 172. ISSN: 2504-477X. DOI: 10.3390/jcs6060172. URL: <https://www.mdpi.com/2504-477X/6/6/172> (visited on 15/12/2025).

- [34] Sotiris Koussios. *Filament winding: a unified approach*. OCLC: 60520507. Delft: Delft University Press, 2004. ISBN: 978-90-407-2551-7.
- [35] D. Kim and S. R. Nutt. "Effective cure cycle development via flow optimization and advanced cure environments". In: *Advanced Manufacturing: Polymer & Composites Science* 6.3 (01/09/2020). _eprint: <https://doi.org/10.1080/20550340.2020.1815276>, pp. 164–172. ISSN: 2055-0340. DOI: 10.1080/20550340.2020.1815276. URL: <https://doi.org/10.1080/20550340.2020.1815276> (visited on 06/03/2026).
- [36] Ujala Farooq et al. "Effect of a Dwell Stage in the Cure Cycle on the Interphase Formation in a Poly(ether imide)/High Tg Epoxy System". In: *ACS Applied Polymer Materials* 3.12 (10/12/2021), pp. 6111–6119. DOI: 10.1021/acsapm.1c00956. URL: <https://doi.org/10.1021/acsapm.1c00956> (visited on 06/03/2026).
- [37] José Antonio González et al. "Towards time-reduced cure cycles of epoxy resins for mass production of composites maintaining the thermo-mechanical properties". In: *Journal of Thermal Analysis and Calorimetry* 150.2 (01/01/2025), pp. 1205–1215. ISSN: 1588-2926. DOI: 10.1007/s10973-024-13442-8. URL: <https://doi.org/10.1007/s10973-024-13442-8> (visited on 06/03/2026).
- [38] Iacopo Bianchi et al. "Effect of Heat-Shrinkable Tape Application on the Mechanical Performance of CFRP Components Obtained by a Filament Winding Process". In: *Journal of Composites Science* 8.12 (16/12/2024). ISSN: 2504-477X. DOI: 10.3390/jcs8120535. URL: <https://www.mdpi.com/2504-477X/8/12/535> (visited on 06/03/2026).
- [39] Dunstone. *Hi-Shrink Tape (PTFE-Coated on Both Sides) 25.4 mm, Roll 91.4 m*. URL: <https://www.r-g.de/en/art/390450-D> (visited on 25/08/2026).
- [40] Eileen O. Kutscha et al. "Contamination and Surface Preparation Effects on Composite Bonding". In: Society for the Advancement of Materials and Process Engineering. NTRS Author Affiliations: Boeing Research and Technology, NASA Langley Research Center NTRS Report/Patent Number: NF1676L-26032 NTRS Document ID: 20170006189 NTRS Research Center: Langley Research Center (LaRC). Covina, CA, United States, 22/05/2017. URL: <https://ntrs.nasa.gov/citations/20170006189> (visited on 12/02/2026).
- [41] Arief Yudhanto, Marco Alfano and Gilles Lubineau. "Surface preparation strategies in secondary bonded thermoset-based composite materials: A review". In: *Composites Part A: Applied Science and Manufacturing* 147 (01/08/2021), p. 106443. ISSN: 1359-835X. DOI: 10.1016/j.compositesa.2021.106443. URL: <https://www.sciencedirect.com/science/article/pii/S1359835X21001640> (visited on 13/02/2026).
- [42] David Mulligan, Samuel Gnaniyah and Graham Sims. *Thermal Analysis Techniques for Composites and Adhesives*. Teddington, Middlesex, UK, 03/2003.
- [43] Dominik Grund, Manuel Orlishausen and Iman Taha. "Determination of fiber volume fraction of carbon fiber-reinforced polymer using thermogravimetric methods". In: *Polymer Testing* 75 (01/05/2019), pp. 358–366. ISSN: 0142-9418. DOI: 10.1016/j.polymertesting.2019.02.031. URL: <https://www.sciencedirect.com/science/article/pii/S0142941818320774> (visited on 18/02/2026).
- [44] *Standard Test Methods for Constituent Content of Composite Materials*. 03/2022. DOI: 10.1520/D3171-22. URL: <https://store.astm.org/d3171-22.html> (visited on 18/02/2026).
- [45] E. Cilley, D. Roylance and N. Schneider. "Methods of Fiber and Void Measurement in Graphite/Epoxy Composites". In: *Composite Materials: Testing and Design (Third Conference)*. Composite Materials: Testing and Design (Third Conference). ASTM International, 01/01/1974, pp. 237–249. DOI: 10.1520/STP35492S. URL: <https://store.astm.org/stp35492s.html> (visited on 13/02/2026).
- [46] M. Bodaghi et al. "Experimental characterization of voids in high fibre volume fraction composites processed by high injection pressure RTM". In: *Composites Part A: Applied Science and Manufacturing* 82 (01/03/2016), pp. 88–99. ISSN: 1359-835X. DOI: 10.1016/j.compositesa.2015.11.042. URL: <https://www.sciencedirect.com/science/article/pii/S1359835X15004534> (visited on 18/02/2026).
- [47] João M. Machado et al. "Automatic void content assessment of composite laminates using a machine-learning approach". In: *Composite Structures* 288 (15/05/2022), p. 115383. ISSN: 0263-8223. DOI: 10.1016/j.compstruct.2022.115383. URL: <https://www.sciencedirect.com/science/article/pii/S0263822322001830> (visited on 13/02/2026).
- [48] *Standard Test Method for Apparent Hoop Tensile Strength of Plastic or Reinforced Plastic Pipe*. 2019. DOI: 10.1520/D2290-19A. URL: <https://store.astm.org/d2290-19a.html> (visited on 03/08/2026).
- [49] *Standard Test Method for Short-Beam Strength of Polymer Matrix Composite Materials and Their Laminates*. 08/2022. DOI: 10.1520/D2344_D2344M-22. URL: https://store.astm.org/d2344_d2344m-22.html (visited on 03/08/2026).
- [50] David K. Roylance. *Netting Analysis for Filament-Wound Pressure Vessels*. Technical Note AMMRC TN 76-3. Watertown, MA: Composites Division Army Materials and Mechanics Research Center, 08/1976. URL: <https://web.mit.edu/roylance/www/netting.pdf> (visited on 01/03/2026).
- [51] Alan T. Nettles. *Basic mechanics of laminated composite plates*. NTRS Author Affiliations: NASA Marshall Space Flight Center NTRS Report/Patent Number: NASA-RP-1351 NTRS Document ID: 19950009349

- NTRS Research Center: Legacy CDMS (CDMS). 01/10/1994. URL: <https://ntrs.nasa.gov/citations/19950009349> (visited on 03/03/2026).
- [52] Federico Paris and Karen E. Jackson. *A Study of Failure Criteria of Fibrous Composite Materials*. NAS 1.26:210661. NTRS Author Affiliations: George Washington Univ., NTRS Document ID: 20010035883 NTRS Research Center: Langley Research Center (LaRC). 01/03/2001. URL: <https://ntrs.nasa.gov/citations/20010035883> (visited on 06/03/2026).
- [53] N. C. W. Judd and W. W. Wright. "Voids and their effects on the mechanical properties of composites – an appraisal". In: *SAMPE Journal* 14.1 (1978), pp. 10–14.
- [54] S. R. Ghiorse. "Effect of void content on the mechanical properties of carbon/epoxy laminates". In: *SAMPE Quarterly* 24 (01/01/1993). ADS Bibcode: 1993SAMPQ..24...54G, pp. 54–59. issn: 0036-0821. URL: <https://ui.adsabs.harvard.edu/abs/1993SAMPQ..24...54G> (visited on 05/03/2026).
- [55] F. Y. C. Boey and S. W. Lye. "Void reduction in autoclave processing of thermoset composites: Part 1: High pressure effects on void reduction". In: *Composites* 23.4 (01/07/1992), pp. 261–265. issn: 0010-4361. doi: 10.1016/0010-4361(92)90186-X. URL: <https://www.sciencedirect.com/science/article/pii/001043619290186X> (visited on 05/03/2026).
- [56] David W. Sleight. *Progressive Failure Analysis Methodology for Laminated Composite Structures*. NTRS Author Affiliations: NASA Langley Research Center NTRS Report/Patent Number: NAS 1.60:209107 NTRS Document ID: 19990052628 NTRS Research Center: Langley Research Center (LaRC). 01/03/1999. URL: <https://ntrs.nasa.gov/citations/19990052628> (visited on 05/03/2026).
- [57] John C. Thesken, Pappu L. N. Murthy and Leigh Phoenix. "Composite Overwrap Pressure Vessels: Mechanics and Stress Rupture Lifting Philosophy". In: American Society of Mechanical Engineers. NTRS Author Affiliations: Ohio Aerospace Inst., NASA Glenn Research Center, Cornell Univ. NTRS Document ID: 20070022369 NTRS Research Center: Glenn Research Center (GRC). United States, 23/04/2007. URL: <https://ntrs.nasa.gov/citations/20070022369> (visited on 18/02/2026).
- [58] Cong Zhao et al. "An Experimental Study of the Influence of in-Plane Fiber Waviness on Unidirectional Laminates Tensile Properties". In: *Applied Composite Materials* 24 (12/2017). doi: 10.1007/s10443-017-9590-z.
- [59] MakeItFrom.com. *Annealed 303 Stainless Steel*. 30/05/2020. URL: <https://www.makeitfrom.com/material-properties/Annealed-303-Stainless-Steel> (visited on 23/02/2026).
- [60] EurocodeApplied team. *Table of Design Properties for Metric Steel Bolts M5 to M39—Eurocode 3*. URL: <https://eurocodeapplied.com/design/en1993/bolt-design-properties> (visited on 23/02/2026).