


Progressive damage modelling of CFRP laminates using a semi-discrete approach in FEM

MSc Thesis

Paula Bustamante Alonso-Lasheras



Progressive damage modelling of CFRP laminates using a semi-discrete approach in FEM

For obtaining the degree of Master of Science in Aerospace Engineering at Delft
University of Technology

by

Paula Bustamante Alonso-Lasheras

Student number: 5279445

Aerospace Structures and Materials
Faculty of Aerospace Engineering, TU Delft

SUPERVISORS

Dr. S.R. Turteltaub · TU Delft

MSc. P. Makiela · German Aerospace Center (DLR)



**Deutsches Zentrum
für Luft- und Raumfahrt**

Preface

Dear reader,

Carbon-fibre-reinforced polymers have given aerospace structures a remarkable freedom in how light and how efficient they can be, and yet the moment one asks not how such a laminate carries load but how it breaks, that confidence gives way to a long list of models, none of them fully to be trusted. This thesis is an attempt to take one of those models, the semi-discrete approach, and ask honestly how much of the failure of an open-hole coupon it can be relied upon to predict. Chasing that question through constitutive laws, mesh-objectivity studies and many long explicit analyses proved as frustrating as it was rewarding, and I am grateful to have been given the room to see it through.

Throughout this journey the guidance of my supervisor, Dr. S.R. Sergio Turteltaub, has been invaluable, both for the freedom to take the work in the direction I found most interesting and for keeping me honest about what the results could and could not claim; for this I am deeply grateful. Thanks go out as well to my supervisor, MSc. Patrick Makiela at the German Aerospace Center whose insight and readiness to discuss the work shaped it throughout, and to the TU Delft High-Performance Computing team, without whose resources none of these simulations would have been possible. Lastly, I would like to thank my family and friends for their continuous love and support, throughout this work.

Paula Bustamante Alonso-Lasheras
Delft, Thursday 27th August, 2026

Contents

Preface	i
Nomenclature	xi
Executive Summary	xvii
1 Introduction	1
1.1 Motivation	1
1.2 Knowledge Gap	1
1.3 Hypothesis and Research Questions	3
1.4 Thesis Outline	3
2 Composite Fundamentals and Failure	5
2.1 Lamina and Laminate Mechanics	5
2.2 Failure Modes and Criteria in CFRP Laminates	7
2.2.1 Intralaminar Failure	8
2.2.2 Interlaminar Failure: Delamination	8
2.2.3 Damage Progression	8
2.2.4 Notch Sensitivity and the OHT Stress Field	10
2.2.5 Complex-Variable (Lekhnitskii–Savin) Formulation	11
2.2.6 Limitations With Damage Prediction	13
2.2.7 Damage-Initiation Criteria	14
2.3 Limitations	16
3 Finite Element Damage Prediction Methods	17
3.1 Continuum Smeared Crack Methods	17
3.1.1 Scalar damage and the mesh-objectivity problem	18
3.1.2 Characteristic length for misaligned, skewed, and non-square elements	20

3.1.3	Mesoscale CDM formulation for composite laminates	23
3.1.4	Inherent limitations and the motivation for discrete extensions	25
3.2	Discrete Crack Methods (DCM)	27
3.2.1	Virtual Crack Closure Technique (VCCT)	27
3.2.2	Cohesive Zone Model (CZM)	28
3.2.3	Extended Finite Element Method and Mesh-Independent Cracking	32
3.2.4	Limitations of discrete methods	33
3.3	Stochastic and Probabilistic Finite Element Approaches	34
3.4	The Semi-Discrete Intersection	34
3.4.1	Meshing strategy	35
4	Finite Element Considerations	37
4.1	Solver Types	37
4.1.1	The Newmark Family	37
4.1.2	Implicit Time Integration	38
4.1.3	Explicit Time Integration	39
4.1.4	Convergence in the Presence of Material Softening	42
4.2	Quasi-Static Modelling	43
4.2.1	Mass	43
4.2.2	Damping	44
4.2.3	Conclusion	46
4.3	Element Types	46
4.4	Boundary Conditions Effects	48
5	Methodology	50
5.1	Problem Definition	50
5.1.1	The Three-Model Ablation	51
5.2	Model Generation and Automation	51
5.3	VUMAT Implementation	58
5.3.1	Intralaminar VUMAT	58
5.3.2	Interlaminar VUMAT	64

5.4	Test setup	64
5.4.1	Quasi Staticness	64
5.4.2	Cohesive Penalty Stiffness	67
5.4.3	Cohesive Mesh Conformity	68
5.4.4	Split and Bulk Dimensions	69
5.4.5	Mesh Convergence	70
5.5	Post Processing	71
6	Results and Discussion	73
6.1	Penalty Stiffness Test	73
6.2	Cohesive Element Partition	75
6.3	Quasi-Static Approximation	76
6.4	Mesh Objectivity	80
6.5	Validation Against Experiments	88
6.6	Out-of-Plane Failure Neglected	89
6.7	Parametric Sensitivity	90
6.8	Ablation Comparison	92
6.8.1	Damage Initiation and Propagation	94
6.8.2	Delamination Patterns	99
7	Conclusions and Outlook	103
7.1	Conclusions	103
7.1.1	SQ1: Contribution of the developed semi-discrete framework	103
7.1.2	SQ2: Numerical objectivity and parameter sensitivity	104
7.1.3	SQ3: Result Fidelity	105
7.1.4	Main Research Question	106
7.2	Limitations and Recommendations for Future Work	107
7.2.1	Mesh Generation	107
7.2.2	Matrix-Split Spacing and Crack Density	107
7.2.3	Temperature Dependency	108
7.2.4	Non-Deterministic Material Properties	108

7.2.5	Constitutive Model	108
7.2.6	Cohesive Element Partitioning	109
7.2.7	Cohesive zone softening	110
7.2.8	In-situ Strength	111
References		112
A Equivalence of the Central Difference Scheme and Newmark with $\beta = 0$		118
B Material Properties		120
C Modal decomposition and stable time increment derivation		121
C.1	Modal Analysis	121
C.2	Stability of the explicit integrator	122

List of Figures

3.1	Crack extensions and crack band widths for a misaligned element. [43]	22
3.2	Bilinear stress–strain law for a single damage mode [41, 81].	24
3.3	Ellipsoidal damage-initiation surface in traction space defined by Equation (3.46). Traction states inside the ellipsoid are purely elastic; a state on the surface triggers initiation. Semi-axes equal the pure-mode interface strengths (illustrative values shown).	29
3.4	Bilinear traction–separation law constructed at the moment of damage initiation [45].	30
4.1	Effect of the damping ratio (ξ) on the free vibration displacement response.	44
4.2	Rayleigh damping ratio versus frequency for $\alpha_R=5.0$ and $\beta_R=0.05$, highlighting the mass-proportional and stiffness-proportional components.	45
4.3	Schematic of hourglass deformation modes in a reduced-integration hexahedral element.	48
5.1	Visualization of the chosen specimen. (a) shows the dimensions of the Open-Hole Tension (OHT) specimen. (b) visualizes the applied boundary conditions, the edge labels, and the reference coordinate system.	51
5.2	Model-generation pipeline. A shared backbone (top and bottom) is common to all three models; Models A, B and C branch at the interlaminar and intralaminar decisions and re-converge before meshing. The dotted box is not used here but implemented as an optional addition.	53
5.3	Material assignment after classification in the semi-discrete model (Model C), shown near the hole for a -45° ply.	54
5.4	Geometric partition set used to locally control mesh density around the hole. Inner region (light), and outer region (dark) covering the rest of the plate.	54
5.5	Fibre-aligned partitioning and resulting mesh for the 0° ply.	55
5.6	Fibre-aligned partitioning and resulting mesh for the 90° ply.	55
5.7	Fibre-aligned partitioning and resulting mesh for the 45° ply.	56
5.8	Fibre-aligned partitioning and resulting mesh for the -45° ply.	56
5.9	Cohesive-layer partitioning conforming simultaneously to the fibre directions of both bounding plies.	56
5.10	Cohesive mesh at a $0^\circ/90^\circ$ interface.	57

5.11 Cohesive mesh at a $\pm 45^\circ$ -involving interface, with a detail of one of the resulting sliver elements near the plate edge.	57
5.12 Adapted cohesive partitioning at a $\pm 45^\circ$ -involving interface, free of the distorted edge elements of Figure 5.11.	58
5.13 Effective-stress update in the ply Vectorised User MATerial subroutine (VUMAT).	59
5.14 Per-increment logic of the intralaminar VUMAT, evaluated at each integration point. Sections 2.2.7 and 3.1.3.	61
5.15 Mapping of a prismatic elements nodes to a hexahedral reference element.	63
5.16 Per-increment logic of the interlaminar (cohesive) VUMAT, evaluated at each integration point. The traction–separation update, quadratic initiation, mixed-mode toughness and damage evolution follow Section 3.2.2.	65
5.17 Stress update in the VUMAT including damping.	67
5.18 Verification configurations for the cohesive penalty stiffness. A two-element coupon is loaded in interface shear ((a) and (c)) and in mode-I opening ((b) and (d)), modelled with the plies tied directly ((a) and (b)) and with a cohesive layer inserted ((c) and (d)).	68
5.19 Free, unpartitioned cohesive mesh	69
5.20 Mesh around the hole at the finest inner element size considered in the convergence study, 0.125 mm, shown for the 0° , 90° and 45° plies.	71
6.1 Two-element coupon in interface shear: axial stress–strain response for the tied reference and three cohesive penalty stiffnesses K	74
6.2 Two-element coupon in mode-I opening: axial stress–strain response for the tied reference and three cohesive penalty stiffnesses K	75
6.3 Stress–strain response of the Cross-Ply (CP) laminate (Model C) with the conforming, ply-strip-partitioned cohesive mesh and with the free-meshed, non-conforming variant.	76
6.4 Stress–strain response of the Quasi-Isotropic (QI) laminate (Model C) with the conforming, ply-strip-partitioned cohesive mesh and with the free-meshed, non-conforming variant.	77
6.5 Kinetic-to-internal energy ratio E_k/E_{in} versus applied strain for the CP laminate at five mass scaling levels (logarithmic scale).	78
6.6 Stress–strain response of the CP laminate for five mass scaling levels.	78
6.7 Kinetic-to-internal energy ratio E_k/E_{in} versus applied strain for the QI laminate at five mass scaling levels (logarithmic scale).	79
6.8 Stress–strain response of the QI laminate for five mass scaling levels.	79
6.9 Mesh objectivity check for Model A (baseline Continuum Damage Mechanics (CDM)), QI laminate: stress–strain response at three in-plane element sizes.	81

6.10 Mesh objectivity check for Model A (baseline CDM), CP laminate: stress–strain response at three in-plane element sizes.	82
6.11 Mesh objectivity check for Model B (cohesive interfaces), CP laminate: stress–strain response at three in-plane element sizes.	82
6.12 Mesh objectivity check for Model B (cohesive interfaces), QI laminate: stress–strain response at three in-plane element sizes.	83
6.13 Mesh objectivity check for Model C (full semi-discrete), CP laminate: stress–strain response at four hole-refined inner element sizes, outer element size fixed at 0.5 mm.	85
6.14 Mesh objectivity check for Model C (full semi-discrete), QI laminate: stress–strain response at four hole-refined inner element sizes, outer element size fixed at 0.5 mm.	86
6.15 Average matrix damage at initiation versus applied strain, Model C, CP laminate, for the same four hole-refined inner element sizes as Figure 6.13.	86
6.16 Average matrix damage at initiation versus applied strain, Model C, QI laminate, for the same four hole-refined inner element sizes as Figure 6.14.	87
6.17 Mesh objectivity check for Model C, CP laminate: stress–strain response with element sizes reduced to 0.25 mm and 0.125 mm.	87
6.18 Mesh objectivity check for Model C, QI laminate: stress–strain response with element sizes reduced to 0.25 mm and 0.125 mm.	88
6.19 Out-of-plane matrix-tension criterion F_{mt33} (recorded, not allowed to degrade), CP laminate, 0° ply, 0.25 mm mesh. The contour range is 0 (dark blue) to 1 (reds), where $F_{mt33} \geq 1$ denotes failure onset	90
6.20 Out-of-plane matrix-tension criterion F_{mt33} (recorded, not allowed to degrade), CP laminate, 90° ply, 0.25 mm mesh. The contour range is 0 (dark blue) to 1 (reds), where $F_{mt33} \geq 1$ denotes failure onset	90
6.21 Split-to-bulk ratio sensitivity, QI laminate: stress–strain response for $w_s/s \in \{1/2, 1/4, 1/6, 1/10\}$	92
6.22 Split-to-bulk ratio sensitivity, CP laminate: stress–strain response for $w_s/s \in \{1/2, 1/4, 1/6, 1/10\}$	92
6.23 Analytical σ_1 stress-concentration field around the hole, CP laminate, 0° and 90° plies.	93
6.24 Analytical σ_1 stress-concentration field around the hole, QI laminate, 0° , 90° and 45° plies.	93
6.25 σ_1 (fibre-direction stress) fields for the CP laminate, 0° and 90° plies.	93
6.26 σ_1 (fibre-direction stress) fields for the QI laminate, 0° , 90° and 45° plies.	94
6.27 Matrix damage variable at initiation for a 0° ply. (Color spectrum with red being fully damaged and blue being undamaged.)	95
6.28 Matrix damage variable after propagation for a 0° ply. (Color spectrum with red being fully damaged and blue being undamaged.)	95

6.29	Matrix damage variable at initiation for a $\pm 45^\circ$ ply. (Color spectrum with red being fully damaged and blue being undamaged.)	96
6.30	Matrix damage variable after propagation for a $\pm 45^\circ$ ply. (Color spectrum with red being fully damaged and blue being undamaged.)	96
6.31	Matrix damage variable at initiation for a 90° ply. (Color spectrum with red being fully damaged and blue being undamaged.)	96
6.32	Matrix damage variable after propagation for a 90° ply. (Color spectrum with red being fully damaged and blue being undamaged.)	97
6.33	CP laminate: matrix damage variable at initiation for a 0° ply. (Color spectrum with red being fully damaged and blue being undamaged.)	98
6.34	CP laminate: matrix damage variable after propagation for a 0° ply. (Color spectrum with red being fully damaged and blue being undamaged.)	98
6.35	CP laminate: matrix damage variable at initiation for a 90° ply. (Color spectrum with red being fully damaged and blue being undamaged.)	98
6.36	CP laminate: matrix damage variable after propagation for a 90° ply. (Color spectrum with red being fully damaged and blue being undamaged.)	99
6.37	QI laminate: delamination at the $45/90$ interface, with points A and B marking matching features in the Model C and Computed Tomography (CT) results.	100
6.38	QI laminate: delamination at the $90/-45$ interface.	100
6.39	QI laminate: delamination at the $-45/0$ interface.	100
6.40	CP laminate: delamination at the $0/90$ interface.	101
7.1	Effect of the cohesive strength reduction on the QI laminate: (a) stress-strain response and (b) cohesive damage variable d_{coh}	110

List of Tables

3.1	Contrasting characteristics of continuum smeared crack methods (Continuum Smeared Crack Methods (CSCM)) and discrete crack methods (Discrete Crack Methods (DCM)).	35
5.1	Groups of solution-dependent state variables carried by the VUMAT between increments. The ply branch uses the intralaminar groups; the cohesive branch uses the interlaminar group.	58
6.1	Wall-clock simulation time for the CP and QI laminates as a function of mass scaling target minimum time increment; all runs on 20 CPU cores.	78
6.2	Experimental failure stress σ_M , CP laminate (German Aerospace Center (DLR), room temperature).	88
6.3	Experimental failure stress σ_M , QI laminate (DLR, room temperature).	88
B.1	Intralaminar material properties and model parameters for 8552-IM7 [76]	120
B.2	Interlaminar material properties for the 8552-IM7 ply interface [76]	120

Nomenclature

Symbols

a	mm	Hole radius (half hole width $2a$)
$\mathbf{A}, \mathbf{B}, \mathbf{D}$	–	Extensional, coupling and bending laminate stiffness matrices (CLT)
a_{av}	mm	Averaging length, average-stress criterion
$\mathbf{A}^v$	N mm^{-1}	System Jacobian / effective tangent stiffness (Newton it. v)
$\mathbf{a}$	mm N^{-1}	Laminate membrane compliance, $\mathbf{a} = \mathbf{A}^{-1}$ (components a_{ij})
A	mm^2	Crack area (increment ΔA); in-plane element area A_{ip}
b	mm	Largest section dimension $\perp$ fibres (Saint-Venant decay)
$\mathbf{B}_{\varepsilon u}$	mm^{-1}	Strain–displacement operator (maps nodal displacement to strain)
$\mathbf{C}$	GPa	Orthotropic stiffness matrix (components C_{ij}); damaged $\mathbf{C}_d$
c_1, c_2	m s^{-1}	Dilatational and distortional wave speeds; element speed c_e
d	mm	Hole diameter
$\mathbf{D}_{\text{damp}}$	N s mm^{-1}	Viscous damping matrix (modal form $\mathbf{D}_m$)
d_{pt}	mm	Characteristic distance, point-stress criterion
D	–	Common denominator in the Lekhnitskii coefficients
d_i	–	Damage variable, mode i (sub-modes $d_{ft}, d_{fc}, d_{mt}, d_{mc}$)
d_f, d_m, d_s	–	Global fibre, matrix and shear damage variables
$\bar{d}_i$	–	Volume-averaged damage variable, mode i
d_{coh}	–	Cohesive (interface) damage variable
E_1, E_2, E_3	GPa	Ply Young's moduli (long., transverse, through-thickness)
$\vec{e}_f, \vec{e}_m$	mm	Fibre and matrix element edge vectors (pairs $\vec{e}_{\text{pair}1,2}$)
E^L	GPa	Effective laminate constants $E_1^L, E_2^L, G_{12}^L, \nu_{12}^L$
E	GPa	Modulus for mode i (E_i); transverse modulus in l_{cz}
E_k, E_{in}	mJ	Kinetic and internal energy
F	N	Airy stress function
$\hat{f}, \hat{m}$	–	Fibre and matrix material-direction unit vectors
F_{tu}	MPa	Unnotched (ultimate) tensile strength
$\mathbf{f}$	N	Nodal force vector: internal $\mathbf{f}^{\text{int}}$, external $\mathbf{f}^{\text{ext}}$, damping $\mathbf{f}^{\text{damping}}$
F_{fi}, F_{mi}	–	Hashin fibre and matrix failure indices ($F_{fi}, F_{mi22}, F_{mi33}$)

$f(\theta)$	–	Puck inter-fibre exposure (failure) function
$F_{x,i}$	N	Crack-front nodal reaction forces $F_{x,i}, F_{y,i}, F_{z,i}$ (VCCT)
G_{12}	GPa	Ply shear moduli G_{12}, G_{13}, G_{23}
g_c^i	N mm ⁻²	Dissipated energy per unit volume, mode i
G	N mm ⁻¹	Strain energy release rate; mode parts G_I, G_{II}, G_{III} , total G_T
G_c	N mm ⁻¹	Fracture toughness / energy, mode i (G_c^i)
G_{Ic}	N mm ⁻¹	Pure-mode critical toughness $G_{Ic}, G_{IIc}, G_{IIIc}$
G_{mc}	N mm ⁻¹	Mixed-mode cohesive fracture toughness
G_{ft}, G_{mt}	N mm ⁻¹	Ply toughnesses $G_{ft}, G_{fc}, G_{mt}, G_{mc}, G_s$
h	mm	Laminate thickness (Ch. 2); plate height / gauge length (Ch. 5)
H	–	Heaviside enrichment function (X-FEM)
$\mathbf{I}$	–	Identity matrix
$\mathbf{K}$	N mm ⁻¹	Global stiffness matrix (parts $\mathbf{K}^{\text{int}}, \mathbf{K}^{\text{ext}}, \mathbf{K}_m$)
k_t	–	Elastic stress-concentration factor
K	N mm ⁻³	Cohesive penalty stiffness (matrix $\mathbf{K}_p$, components K_n, K_s, K_t)
K_b	GPa	Bulk modulus (wave-speed evaluation)
K_L	–	Lekhnitskii orthotropy ratio, $\sqrt{E_1^L/E_2^L}$
K_T	–	Stress-concentration factor (infinite plate / finite width)
L_c	mm	Characteristic (crack-band) length; variants $L_c^m, L_c^f, L_{c,IP}, L_{c,ML}, L_{c,FD}$
l	mm	Specimen span in the loading (x) direction
l_c	mm	Cohesive-element constitutive thickness
L_{Ab}	mm	Abaqus default length, $\sqrt[3]{V}$
l_e	mm	Element length governing the stable time step
L_0	mm	Initial gauge (reference) length
l_{cz}	mm	Cohesive (fracture) process-zone length
$\mathbf{M}_{\text{mass}}$	t	Mass matrix (element form $\mathbf{M}_e$)
m, n	–	$\cos \theta, \sin \theta$ in the ply transformation matrix
$\mathbf{N}, \mathbf{M}$	N mm ⁻¹	Laminate force and moment resultants (CLT)
m_L	–	Lekhnitskii orthotropy parameter, $2\nu_{12}^L - E_1^L/G_{12}^L$
$\mathbf{M}_d$	–	Damage operator; its inverse $\mathbf{M}_d^{-1}$ degrades the effective stress to the nominal stress
M	–	Tan–Heywood exponent (Ch. 2); cohesive-zone parameter (Ch. 3)
$\mathbf{N}_s$	–	Shape-function matrix (shape functions N_i)
n_s	–	Split-bulk pairs skipped at a free edge; tiling counts n_1, n_2
N_e	–	Elements spanning the cohesive process zone
N_{DOF}	–	Number of degrees of freedom
p^f, p^m	–	Crack planes with fibre-normal and transverse-normal
p, q	–	Reduced numerator/denominator of w/l
$p_{\perp\parallel}, p_{\perp\perp}$	–	Puck inclination (friction) parameters

Q	GPa	Reduced / transformed stiffness matrix (components Q_{ij})
q	–	Modal-coordinate vector (scalar q)
R	mm	Hole radius, $R = d/2$ (Lekhnitskii)
r	mm	Radial coordinate from the hole centre
r	N	Residual force vector (implicit solver)
S	GPa ⁻¹	Compliance matrix (damaged S_d ; smeared S_{ij})
s	mm	Bulk spacing between split strips; pitch $s_{\text{tot}} = w_s + s$
S_{12}	MPa	Ply shear strengths S_{12}, S_{13}, S_{23}
s_D	–	Static/dynamic switch (1 dynamic, 0 static)
S_{mt}	–	Shear-damage coupling coefficient
T	–	In-plane stress transformation matrix
t	s	Time; also ply thickness (mm)
t	MPa	Cohesive traction vector, components t_n, t_s, t_t
t_n^0	MPa	Pure-mode interface strengths t_n^0, t_s^0, t_t^0
u	mm	Nodal displacement (velocity $\dot{\mathbf{u}}$, acceleration $\ddot{\mathbf{u}}$)
u_x	mm	Applied axial displacement (maximum u_{max})
u_i, v_i, w_i	mm	Crack-face nodal displacement components (VCCT)
V	mm ³	Element volume (integration-point volume V_i)
v	mm s ⁻¹	Nodal velocity vector
v	–	Newton iteration index
W	mJ	Work of fracture; energies $W_{\text{kin}}, W_{\text{int}}, W_{\text{ext}}$
w	mm	Specimen width (y -direction)
w_s	mm	Matrix-split strip width
X, Y, Z	MPa	Long., transverse, through-thickness strengths (tension t , compression c)
z	mm	Through-thickness coordinate (ply positions z_k)
z_j	–	Lekhnitskii complex variable, $x + \mu_j y$
α	–	Element aspect ratio; cohesive power-law exponent; X-FEM tip index
α_R, β_R	–	Rayleigh mass- and stiffness-proportional damping coefficients
α_j, β_j	–	Real / imaginary parts of the Lekhnitskii root μ_j
β, γ	–	Newmark time-integration parameters
Δ	mm	Cohesive separation vector, components $\Delta_n, \Delta_s, \Delta_t$
Δ_e	mm	Effective separation; onset Δ_e^0 , failure Δ_e^f , peak Δ_e^{max}
Δa	mm	Crack extension per element: $\Delta a_{\text{ML}}, \Delta a_{\text{max}}, \Delta a_{\text{avg}}$
Δt	s	Time increment; critical value Δt_{crit}
δ_{ij}	–	Kronecker delta
ϵ	–	Strain tensor (ϵ_{ij}); mid-plane strain ϵ^0
ϵ_{eq}^i	–	Equivalent strain, mode i (onset $_{,0}$, full damage $_{,f}$)
ϵ_V	–	Volumetric strain

ε	–	Energy-balance tolerance (quasi-static check)
γ_s	–	Element skew angle (crack-band geometry)
γ_{ij}	–	Engineering shear strains, $\gamma_{ij} = 2\varepsilon_{ij}$
$\boldsymbol{\kappa}$	mm^{-1}	Laminate mid-plane curvature vector
κ_j	–	Lekhnitskii potential coefficients ($\kappa_j, \lambda_j, \Lambda_j$)
λ, μ	GPa	Lamé constants ($\mu = G$, shear modulus)
Λ	–	Element eigenvalue (maximum $\Lambda_{\max}$)
λ_i	–	Cohesive mode-energy fractions
μ_j	–	Roots of the Lekhnitskii characteristic equation
μ_a	–	Amplification factor in the stability analysis
η	–	Benzeggagh–Kenane mixed-mode exponent
η_1, η_2, η_3	–	Normalised reference-element coordinates
ν_{12}	–	Ply Poisson's ratios $\nu_{12}, \nu_{13}, \nu_{23}$
ν	–	Compliance-inversion denominator (3-D stiffness)
Ω	mm^3	Body / element (Ω_e) integration domain
φ	–	Circumferential coordinate (Lekhnitskii / Kirsch)
$\boldsymbol{\phi}$	–	Mode-shape vectors ($\boldsymbol{\phi}_i$); modal matrix $\boldsymbol{\Phi}$
Φ_j	–	Lekhnitskii complex potentials
ψ	–	Mesh misalignment angle (per plane ψ^i)
ψ_0	–	Phase angle of a free-vibration mode
ρ	t mm^{-3}	Material density
$\boldsymbol{\sigma}$	MPa	Stress tensor (σ_{ij}, τ_{ij}); effective stress $\tilde{\sigma}$
σ_{eq}^i	MPa	Equivalent stress, mode i (onset $\sigma_{eq,0}^i$)
σ_∞	MPa	Remote / far-field stress (hole-edge stress σ_θ)
$\bar{\sigma}$	MPa	Volume-averaged (loading-direction) stress
σ_M	MPa	Measured / reference (σ_{ref}) failure stress
σ_m	MPa	Mean (hydrostatic) stress; deviatoric change ΔS
σ_{vis}	MPa	Viscous (damping) stress (σ_{elastic} elastic part)
θ	–	Ply orientation / circumferential / action-plane angle
τ_e	MPa	Effective cohesive traction (onset τ_e^c)
τ^0	MPa	Generic interface strength (cohesive-zone length)
σ_n	MPa	Puck action-plane tractions $\sigma_n, \tau_{nt}, \tau_{n1}$
ω	rad s^{-1}	Natural angular frequency (maximum $\omega_{\max}$)
ξ	–	Critical damping ratio
χ	–	Normalised distance from the hole, $(x - R)/(w/2 - R)$ (Tan correction)
ζ_j	–	Lekhnitskii conformal-mapping variable
$\langle \cdot \rangle$	–	Macauly bracket, $\max(\cdot, 0)$
$\hat{(\cdot)}$	–	Unit vector

$\tilde{(\cdot)}$	–	Effective (undamaged) quantity
$\bar{(\cdot)}$	–	Laminate-axis or volume-averaged quantity

Subscripts

1, 2, 3	Ply material axes: fibre, in-plane transverse, through-thickness
x, y, z	Global (specimen) coordinate axes
t	Tension
c	Compression; critical (fracture) value
f	Fibre; value at full damage (e.g. $\varepsilon_{eq,f}$)
m	Matrix
s	Shear
n, s, t	Cohesive normal and two shear directions
I, II, III	Fracture modes (opening, in-plane shear, out-of-plane shear)
ft, fc, mt, mc	Fibre tension/compression, matrix tension/compression damage modes
eq	Equivalent (scalar driving) quantity
0	Value at damage onset (initiation); initial/reference value (L_0)
e	Element quantity; effective (mixed-mode) cohesive quantity
IP	In-plane
ML, FD, CPND, Ab	Characteristic-length variants: mesh-line, fibre-direction, crack-plane-normal, Abaqus
crit	Critical (stable) value
max, avg	Maximum and average value
coh	Cohesive (interface) quantity
ref, M	Reference value; measured (experimental) value
kin, int, ext	Kinetic, internal and external energy
tot	Total (split + bulk pitch)
R	Rayleigh damping coefficient
$n, n+1, n \pm \frac{1}{2}$	Time-step (and half-step) indices

Acronyms

BK	Benzeggagh–Kenane
CDM	Continuum Damage Mechanics
CFL	Courant–Friedrichs–Lewy
CFRP	Carbon Fibre Reinforced Polymer
CLT	Classical Laminate Theory

CP	Cross-Ply
CSCM	Continuum Smeared Crack Methods
CT	Computed Tomography
CZM	Cohesive Zone Modelling
DCM	Discrete Crack Methods
DLR	German Aerospace Center
DOF	Degree of Freedom
DPM	Discrete-Ply Model
EOM	Equation of Motion
FC	Fibre Compression
FE	Finite Element
FEM	Finite Element Method
FSDT	First-Order Shear Deformation Theory
FT	Fibre Tension
GFRP	Glass Fibre Reinforced Polymer
HPC	High-Performance Computing
LEFM	Linear Elastic Fracture Mechanics
MC	Matrix Compression
MT	Matrix Tension
OHT	Open-Hole Tension
QI	Quasi-Isotropic
SD2M	Semi-Discrete Damage Model
SERR	Strain Energy Release Rate
TSL	Traction–Separation Law
VCCT	Virtual Crack Closure Technique
VUMAT	Vectorised User MATerial subroutine
XFEM	eXtended Finite Element Method

Executive Summary

Qualifying weight-efficient Carbon Fibre Reinforced Polymer (CFRP) structures economically depends on being able to predict by simulation, and so reduce the physical testing needed to demonstrate, the point at which a laminate fails rather than merely how stiffly it responds. Establishing that point is hard because the strength of a notched laminate is governed by three damage mechanisms, cracking of the matrix in the plies, separation of the plies at their interfaces, and rupture of the load-bearing fibres, that develop together and in a sequence that itself decides the failure load. The OHT coupon concentrates these mechanisms and their interaction into one reproducible geometry, and is adopted throughout this work as the benchmark against which the modelling is judged.

The framework examined against this benchmark is a semi-discrete one, positioned deliberately between the two established families of failure model. Instead of smearing every crack over the element it occupies, as continuum damage mechanics (CDM) does, or inserting each crack as an explicit discontinuity, as fully discrete methods do, it confines matrix cracking to thin, fibre-aligned strips of ordinary continuum elements while leaving the material between the strips free to be meshed coarsely; delamination, the one mode the plies cannot themselves carry, is added discretely with cohesive elements at every interface. To measure what this construction actually contributes, three models were built from a single automated Python–Abaqus generator and made to differ only in the feature under study: a plain CDM laminate (Model A), the same laminate with cohesive interfaces added (Model B), and the full semi-discrete model carrying both cohesive interfaces and matrix-split strips (Model C). All three draw their constitutive response from a single user material subroutine, which holds the intralaminar ply law and the interlaminar traction–separation law behind a per-material switch and is supported by a dedicated characteristic-length routine (VUCHARLEN), extended here to the distorted hexahedral and prismatic elements that the fibre-aligned split mesh unavoidably produces. The study is carried out on two IM7–8552 layups, a CP and a QI, loaded in open-hole tension.

Before the three models could be compared, every numerical setting not itself under study was fixed by a dedicated study: the cohesive penalty stiffness was raised until the interface was elastically indistinguishable from a tied bond, the quasi-static approximation was confirmed by holding the kinetic energy to a small fraction of the internal energy for every reported run, and the neglected out-of-plane matrix mode was verified to remain well below activation. The most significant, and least anticipated, outcome of this verification concerns mesh objectivity. Because the matrix-split strips give the matrix crack a discrete, localised representation, they reintroduce the process-zone resolution requirement familiar from cohesive-zone modelling, one that the crack-band correction alone does not satisfy and that, unlike for cohesive elements, cannot be relaxed by lowering a strength, because the strip carries a genuine material property. Refining only the region around the hole therefore leaves Model C unconverged; objectivity is recovered only once the resolution requirement is met everywhere the strips are free to propagate, which for the present geometry means essentially the whole plate. The discrete fidelity the method delivers is thus bought at the price of a stricter, and less local, mesh demand than a smeared model imposes.

Against the experimental data, Model C predicts the failure load within one experimental standard deviation for the QI laminate (1.5–2 % under the mean) and within two for the CP laminate (3–4 % over it), in both cases at or within the specimen-to-specimen scatter of the tests. Models A and B both over-predict the measured mean for both laminates. The ablation shows why: the decisive feature is the matrix-split decomposition. By keeping matrix damage within thin strips rather than spreading it across the whole ply, Model C avoids the excessive matrix degradation that otherwise relieves the stress concentration at the hole and lets the 0° fibres survive to an inflated far-field stress. The same feature reproduces the location, orientation and, crucially, the discrete, multi-line character of the matrix cracking, which the single available X-ray CT reference (for the QI laminate) shows across all three ply orientations, whereas Model B renders it only as a smeared patch of the right position but the wrong character. The cohesive interface, by contrast, adds little to the failure load under the in-plane-dominated loading considered here; its value lies in damage-mode completeness, being the only feature by which delamination can be represented at all, and the interlaminar patterns are matched only qualitatively, with neither cohesive model a consistently closer match to the scans.

These outcomes are obtained at a cost that confirms the method's intended position between the continuum and discrete families. The mesh requirement identified above, resolving the process zone over essentially the whole plate, does make Model C more demanding than a locally refined smeared model, but not by enough to change its cost class: it completes in roughly one hour for the CP laminate and at most about nine hours for the most demanding QI case on standard cluster hardware. Although the fibre-aligned split mesh forces very small elements, the mass scaling used to compensate for them keeps this runtime close to that of the far simpler Model B and in the same cost class as a plain CDM analysis, rather than approaching that of a discrete method. For the two carbon/epoxy layups studied, the semi-discrete framework therefore reproduces the load–displacement response and failure load of room-temperature OHT specimens realistically and at tractable cost, together with the discrete character of the matrix damage for the QI laminate against which CT data are available. As qualified below, however, the QI failure load is matched more convincingly than the QI failure mode.

The conclusion is positive but bounded. It rests on two layups of a single material system, one experimental campaign, and a single CT reference available only for the QI laminate. Objectivity of the full model is conditional on the process-zone resolution requirement noted above. The split-to-bulk ratio that sets the width of the matrix-split strips is not yet shown to converge to a limit as it is narrowed, so it must be treated as a potential free parameter of the model rather than a settled discretisation choice; and the reduced cohesive strength used to resolve the interface process zone is implicated in a spurious QI delamination that overrides the split-to-bulk ordering the notch-blunting mechanism predicts, so that the close QI failure-load agreement is reached alongside an interlaminar damage development that is likely premature and a delamination-dominated failure mode, where the literature expects fibre pull-out for this layup. Recalibrating that strength, establishing convergence of the split-to-bulk ratio, improving the mesh generation, and extending the framework to further layups, materials and loadings, and eventually to temperature-dependent and cryogenic behaviour, are the principal directions left for future work.

Introduction

1.1. Motivation

Carbon-fibre-reinforced polymers (CFRP) are used in primary aerospace structures because their high specific stiffness and strength translate directly into weight savings. Realising that benefit in a certified structure, however, requires confidence that the failure of a laminate can be predicted, not merely measured. Qualifying a composite structure relies heavily on physical testing, which is costly and time-consuming. A predictive model that reproduces the strength and the failure process of a laminate allows part of this effort to be carried by analysis rather than testing, reducing the number of physical specimens required and allowing more of the design space to be explored before anything is built. The accuracy demanded of such a model is high, because the quantity of interest is not the elastic response but the load at which the laminate fails and the sequence of events that leads there.

This is difficult precisely because failure of a CFRP laminate is not a single event. It is the interaction of several distinct mechanisms acting at different scales: matrix cracking within the off-axis plies, delamination between plies, and the final fibre failure that exhausts the load-carrying capacity. These mechanisms are coupled, and their ordering governs the strength of the laminate [26, 2]. The OHT coupon is common for studying them together [1]. The hole introduces a stress concentration that localises damage in a known region, yet the response remains rich enough to exhibit matrix cracking, delamination and fibre failure and their interaction in a single, repeatable test. A progressive-damage model that aspires to predictive use must therefore reproduce the OHT response, which is the validation target of this thesis.

1.2. Knowledge Gap

The finite-element methods available to model this progression fall into two families, distinguished by how each represents a crack. Chapter 3 reviews both in detail; the contrast that defines the gap is summarised here. CDM smears the crack over the element volume and represents a crack by degrading

the element stiffness. It is cheap and, through the crack-band correction, objective with respect to mesh size [41, 5]. Two limitations remain, however. Treating the ply as a single homogenized medium, it cannot distinguish a low-energy crack running along the fibres from a higher-energy one running across them; and because damage is evaluated point-by-point from the local stress alone, the softening band localizes along the directions of the mesh rather than the physical fracture plane, so the predicted crack path stays mesh-dependent. Furthermore CDM cannot model a realistic damage pattern as matrix cracks are smeared over the element and do not have a discrete appearance.

Discrete crack methods take the opposite stance, representing a crack as a true displacement discontinuity. They recover the crack opening, the tractions on the crack faces and the energy released directly, and so reproduce the failure process with high fidelity. The cost of this fidelity is twofold: the discontinuity must be placed explicitly, which usually requires the crack plane to be known in advance or a specialised element formulation to introduce it, and the computational effort grows with the number of cracks tracked [49]. For the dense, distributed field of matrix cracks that develops in a laminate, fully discrete modelling becomes expensive and, where crack paths are not known beforehand, complicated to set up [1].

The semi-discrete approach occupies the intersection of these two families. It keeps the low cost of CDM by using only standard continuum elements everywhere, while approaching the fidelity of the discrete methods by embedding the crack geometry into the mesh itself [49]. Matrix cracking is confined to thin strips of elements pre-aligned with the fibre direction, the matrix-split strips, so that the crack band is forced to localize along the physically correct fibre-parallel plane. The splits are modelled with a small width to represent discrete like cracks. The surrounding bulk material, in which matrix cracking is suppressed, may be meshed coarsely. Interlaminar delamination, the one mode that the continuum plies still cannot represent, is added discretely with cohesive elements at every ply interface. The result is a model that is continuum in its elements but approximates discrete matrix cracks.

What remains unresolved is how much this construction actually buys or if its predictive capabilities are significantly better, and whether it delivers on its promise for the OHT problem. The published validation of the semi-discrete method rests almost entirely on the work of the originating group and a small number of carbon/epoxy systems [50, 53, 51] (Chapter 3), so its predictive capability has not been broadly established. Furthermore, the method has not been situated against the very baselines it claims to improve upon in a controlled, like-for-like comparison: the contribution of the cohesive interfaces and of the matrix-split strips, each isolated against an otherwise identical model, has not been quantified.

This thesis addresses that gap directly. It takes the underlying concept of the semi-discrete method, namely the meshing strategy, and builds it from the ground up. It includes a three-model ablation (Chapter 5) in which all three models are built and differ only in the modelling feature under study, so that any difference in the predicted response can be attributed to that feature alone: a baseline CDM laminate, the same laminate with cohesive interfaces added, and the full semi-discrete model with both cohesive interfaces and matrix-split strips. The entire constitutive response, the intralaminar ply law and the interlaminar traction–separation law alike, is implemented in a single user material subroutine, supported by a purpose-written characteristic-length subroutine (VUCHARLEN) that supplies the per-element crack-band length. This routine implements existing methods of calculating characteristic lengths for some hexahedral elements, and is extended to introduce treatment of all hexahedral elements and approximate prismatic elements, that the fibre-aligned split mesh unavoidably produces, for which the standard hexahedral length measure does not apply. All three models are then produced from one

automated model-generation pipeline so that the comparison is genuinely controlled. The models are then verified to a degree that, to the author's knowledge, has not previously been undertaken: the quasi-static approximation, the cohesive penalty stiffness, the split-and-bulk strip dimensions, and the in-plane mesh density are each settled through a dedicated parametric study before the models are compared. The capabilities of the resulting model are then judged by validation against experimentally tested OHT laminates.

1.3. Hypothesis and Research Questions

The hypothesis underlying this work is that embedding a discrete like representation of matrix cracking and true discrete delamination into an otherwise continuum, ply-level model yields a more realistic prediction of the progressive failure of an OHT laminate than a continuum model alone, while keeping the computational cost tractable. This leads to the following main research question and four sub-questions, the last of which is a conditional extension.

Main research question. To what extent can a semi-discrete progressive-damage model, applied at the ply level in a CFRP laminate, realistically predict the mechanical behaviour, damage evolution, and failure of OHT specimens at room temperature?

SQ1 (contribution). What is the marginal contribution of the cohesive interlaminar modelling and of the embedded matrix-split intralaminar modelling to the predicted response, as isolated by the three-model ablation?

SQ2 (verification). Does the model behave consistently and objectively at the numerical level, with respect to mesh-size objectivity, element-level consistency, and sensitivity to its numerical parameters?

SQ3 (validation). Can the model reproduce the load–displacement response, the failure load, and the observed damage patterns of experimentally tested OHT layups?

1.4. Thesis Outline

The thesis is organised as follows. Chapter 2 establishes the ply- and laminate-level mechanics on which the rest of the work builds: classical lamination theory, the failure modes of a CFRP laminate, the stress field at a hole, the classical notched-strength criteria. It closes by showing that these analytical tools can be used as plausibility checks and preliminary sizing, but fail to capture the multi axial damage state and its evolution, which motivates a numerical treatment.

Chapter 3 reviews the finite-element methods for progressive damage, organised by how each represents a crack: continuum smeared-crack methods, discrete crack methods, and the orthogonal strand of stochastic modelling. It positions the semi-discrete approach as the intersection of the continuum and discrete families.

Chapter 4 develops the finite-element foundations of the model: the explicit and implicit time-integration schemes and why an explicit solver is suited to material softening, the approximation of quasi-static loading, and the choice of element type. Chapter 5 turns from theory to practice, defining the OHT problem, introducing the three-model ablation, and describing the automated model-generation pipeline,

the user material subroutine, and the parametric studies that fix the numerical settings common to all three models.

Chapter 6 presents and discusses the results, from verification through the ablation to validation against the experimental OHT data. The final chapter draws the conclusions and sets out recommendations for future work.

Composite Fundamentals and Failure

Before damage prediction methods for CFRP can be discussed, the fundamentals of CFRP materials have to be established. This chapter therefore establishes the ply- and laminate-level mechanics on which the remainder of the thesis builds. Section 2.1 summarizes classical lamination theory and the effective laminate constants used later. Section 2.2 reviews the failure modes of a CFRP laminate, the anisotropic stress field at a hole, the classical notched-strength criteria whose limitations motivate a progressive-damage approach, and an overview of the ply-level damage-initiation criteria.

2.1. Lamina and Laminate Mechanics

The building block of a laminate is the unidirectional ply, treated as a homogeneous, orthotropic continuum. With the 1-axis aligned with the fibres, the 2-axis transverse in-plane, and the 3-axis through the thickness, the three-dimensional constitutive relation in ply coordinates reads $\sigma = \mathbf{C} \varepsilon$ with the 6×6 stiffness matrix

$$\begin{Bmatrix} \sigma_{11} \\ \sigma_{22} \\ \sigma_{33} \\ \tau_{23} \\ \tau_{13} \\ \tau_{12} \end{Bmatrix} = \begin{bmatrix} C_{11} & C_{12} & C_{13} & 0 & 0 & 0 \\ & C_{22} & C_{23} & 0 & 0 & 0 \\ & & C_{33} & 0 & 0 & 0 \\ & & & C_{44} & 0 & 0 \\ \text{Sym.} & & & & C_{55} & 0 \\ & & & & & C_{66} \end{bmatrix} \begin{Bmatrix} \varepsilon_{11} \\ \varepsilon_{22} \\ \varepsilon_{33} \\ \gamma_{23} \\ \gamma_{13} \\ \gamma_{12} \end{Bmatrix}. \quad (2.1)$$

As seen in Eq. (2.1), Voigt notation is used for the stress and strain vectors, with engineering shear strains

$$\gamma_{ij} = 2\varepsilon_{ij}, \quad i, j = 1, 2, 3. \quad (2.2)$$

This notation will be used throughout this work.

The zero off-diagonal blocks follow from the three mutually orthogonal material symmetry planes of an

orthotropic solid [37]. The nine stiffness coefficients C_{ij} are expressed in terms of the nine engineering constants — longitudinal modulus E_1 , transverse moduli E_2 and E_3 , shear moduli G_{23} , G_{13} , G_{12} , and Poisson's ratios ν_{12} , ν_{13} , ν_{23} — through compliance-matrix inversion as [18]

$$\begin{aligned} C_{11} &= E_1 \frac{1 - \nu_{23}\nu_{32}}{\nu}, & C_{22} &= E_2 \frac{1 - \nu_{13}\nu_{31}}{\nu}, & C_{33} &= E_3 \frac{1 - \nu_{12}\nu_{21}}{\nu}, \\ C_{12} &= E_2 \frac{\nu_{12} + \nu_{32}\nu_{13}}{\nu}, & C_{13} &= E_3 \frac{\nu_{13} + \nu_{12}\nu_{23}}{\nu}, & C_{23} &= E_3 \frac{\nu_{23} + \nu_{21}\nu_{13}}{\nu}, \\ C_{44} &= G_{23}, & C_{55} &= G_{13}, & C_{66} &= G_{12}, \end{aligned} \quad (2.3)$$

with denominator

$$\nu = 1 - \nu_{12}\nu_{21} - \nu_{23}\nu_{32} - \nu_{13}\nu_{31} - 2\nu_{32}\nu_{21}\nu_{13}. \quad (2.4)$$

The six Poisson's ratios are not independent: Maxwell's reciprocal relations require

$$\nu_{32} = \frac{E_3}{E_2} \nu_{23}, \quad \nu_{31} = \frac{E_3}{E_1} \nu_{13}, \quad \nu_{21} = \frac{E_2}{E_1} \nu_{12}, \quad (2.5)$$

leaving nine independent engineering constants. For an undamaged CFRP ply, the assumption of transverse isotropy in the 2–3 plane (the fibre is the axis of symmetry; the matrix is macroscopically isotropic in the transverse plane) further reduces the independent set to five through [18]

$$E_3 = E_2, \quad G_{13} = G_{12}, \quad \nu_{13} = \nu_{12}, \quad G_{23} = \frac{E_2}{2(1 + \nu_{23})}. \quad (2.6)$$

For the laminate-level analysis within Classical Laminate Theory (CLT) the plane-stress condition $\sigma_3 = \tau_{23} = \tau_{13} = 0$ is imposed on each ply, reducing Eq. (2.1) to the familiar reduced stiffness matrix

$$\mathbf{Q} = \begin{bmatrix} Q_{11} & Q_{12} & 0 \\ Q_{12} & Q_{22} & 0 \\ 0 & 0 & Q_{66} \end{bmatrix}, \quad \begin{aligned} Q_{11} &= \frac{E_1}{1 - \nu_{12}\nu_{21}}, & Q_{22} &= \frac{E_2}{1 - \nu_{12}\nu_{21}}, \\ Q_{12} &= \frac{\nu_{12}E_2}{1 - \nu_{12}\nu_{21}}, & Q_{66} &= G_{12}. \end{aligned} \quad (2.7)$$

For a ply oriented at an angle θ to the laminate reference axis, the stiffness in laminate axes is obtained by the standard tensor transformation $\bar{\mathbf{Q}}(\theta) = \mathbf{T}^{-1}\mathbf{Q}\mathbf{T}^{-\top}$ and $\bar{\mathbf{C}}(\theta) = \mathbf{T}^{-1}\mathbf{C}\mathbf{T}^{-\top}$ respectively, where $\mathbf{T}$ is the in-plane stress transformation matrix and the $\mathbf{T}^{-\top}$ factor reflects the engineering-shear (Reuter) convention for the strain vector [37]. Explicitly, with $m = \cos \theta$ and $n = \sin \theta$,

$$\mathbf{T}(\theta) = \begin{bmatrix} m^2 & n^2 & 2mn \\ n^2 & m^2 & -2mn \\ -mn & mn & m^2 - n^2 \end{bmatrix}. \quad (2.8)$$

CLT assembles the ply responses under the Kirchhoff assumption that the laminate normal remains straight and inextensible, so that the strain varies linearly through the thickness, $\boldsymbol{\varepsilon}(z) = \boldsymbol{\varepsilon}^0 + z \boldsymbol{\kappa}$, with $\boldsymbol{\varepsilon}^0$ the mid-plane strains and $\boldsymbol{\kappa}$ the curvatures. Integrating the ply stresses through the thickness yields the

laminate constitutive law relating the in-plane force and moment resultants to the mid-plane deformation,

$$\begin{Bmatrix} \mathbf{N} \\ \mathbf{M} \end{Bmatrix} = \begin{bmatrix} \mathbf{A} & \mathbf{B} \\ \mathbf{B} & \mathbf{D} \end{bmatrix} \begin{Bmatrix} \boldsymbol{\varepsilon}^0 \\ \boldsymbol{\kappa} \end{Bmatrix}, \quad (2.9)$$

where the extensional, coupling and bending stiffness matrices follow from the ply positions z_k as $A_{ij} = \sum_k \bar{Q}_{ij}^{(k)}(z_k - z_{k-1})$, $B_{ij} = \frac{1}{2} \sum_k \bar{Q}_{ij}^{(k)}(z_k^2 - z_{k-1}^2)$ and $D_{ij} = \frac{1}{3} \sum_k \bar{Q}_{ij}^{(k)}(z_k^3 - z_{k-1}^3)$. For the symmetric, balanced laminates the coupling matrix $\mathbf{B}$ vanishes, decoupling membrane and bending response.

It is often convenient to treat the laminate as an equivalent homogeneous orthotropic plate. Inverting the membrane sub-problem, $\boldsymbol{\varepsilon}^0 = \mathbf{a} \mathbf{N}$ with $\mathbf{a} = \mathbf{A}^{-1}$, the effective laminate engineering constants for a laminate of thickness h are

$$E_1^L = \frac{1}{h a_{11}}, \quad E_2^L = \frac{1}{h a_{22}}, \quad G_{12}^L = \frac{1}{h a_{66}}, \quad \nu_{12}^L = -\frac{a_{12}}{a_{11}}. \quad (2.10)$$

These smeared constants are used in Section 2.2.4 to evaluate the anisotropic stress concentration at a hole [38].

Recovery of Ply-Level Stresses

Once the mid-plane strain $\boldsymbol{\varepsilon}^0$ at a point in the laminate is known, from Eq. (2.9) for a given resultant $\mathbf{N}$, or, as used in Section 2.2.4, from the local stress state obtained by treating the laminate as the equivalent homogeneous plate of Eq. (2.10), the Kirchhoff hypothesis guarantees that this same strain is shared by every ply in the stack (for membrane loading, $\boldsymbol{\kappa} = \mathbf{0}$). The stress carried by ply k , still in laminate axes, follows directly from its transformed stiffness,

$$\boldsymbol{\sigma}_k^{\text{lam}} = \bar{\mathbf{Q}}_k \boldsymbol{\varepsilon}^0, \quad (2.11)$$

and its components in material axes are obtained by rotating with the transformation matrix of Eq. (2.8),

$$\begin{Bmatrix} \sigma_{11} \\ \sigma_{22} \\ \tau_{12} \end{Bmatrix}_k = \mathbf{T}(\theta_k) \boldsymbol{\sigma}_k^{\text{lam}}. \quad (2.12)$$

A direct consequence of Eq. (2.12) is that $\boldsymbol{\sigma}_k$ depends only on the ply orientation θ_k and not on its position within the stack: under CLT, all plies of equal angle in a given laminate carry identical in-plane stress. This is not a real behavior of a laminate and only arises under the assumptions of CLT.

2.2. Failure Modes and Criteria in CFRP Laminates

Failure of a CFRP laminate is not a single event but the interaction of several distinct mechanisms acting at the ply and interface scales. Because the matrix is far weaker than the fibres, damage typically initiates as matrix cracking in off-axis plies, redistributes load, and progresses through fibre failure and delamination until the laminate can no longer carry load [38]. The mechanisms are conventionally grouped into

intralaminar failure (within a ply) and interlaminar failure (between plies); the two are coupled, most visibly where matrix cracks reaching a ply interface nucleate delaminations. The following subsections describe these modes, the anisotropic stress field that concentrates them around a hole, the criteria classically used to predict notched strength, and an overview of the ply-level criteria used to predict the onset of damage.

2.2.1. Intralaminar Failure

Intralaminar failure comprises fibre-dominated and matrix-dominated mechanisms, which differ fundamentally in their physics and in the energy they dissipate.

Fibre failure. Fibres carry the majority of the load along their axis. Their failure in plies oriented such that their fibers carry significant load is therefore considered a critical failure mechanism. Under longitudinal tension, once the first fibres fail they redistribute their carried load. These breaks accumulate until a cluster becomes large enough to propagate unstably across the ply. Fiber failure is sudden, and releases a large amount of stored elastic energy [38].

Matrix cracking. Within a ply, load along the fibre direction is carried by the fibres, whereas transverse and shear load must be carried by the matrix and the fibre–matrix interface, which are an order of magnitude weaker. Matrix cracks initiate once the transverse tensile or in-plane shear stress reaches the corresponding ply strength (Y_t , S_{12}). The cracks run parallel to the fibres and arrest at the neighboring ply interfaces, so the ply retains its longitudinal stiffness but loses most of its transverse and shear stiffness. Matrix cracking is usually the first damage mechanism to appear, well below final failure and is therefore considered a sub-critical damage mechanism. Its significance lies in its effect on the ultimate failure strength of the laminate further discussed in Section 2.2.3.

2.2.2. Interlaminar Failure: Delamination

Delamination is the separation of adjacent plies along their interface. It can both be a sub-critical and a critical failure mechanism: first it develops locally and then grows as later described in Section 2.2.3 until it ultimately separates the laminate so extremely that it causes a major loss of load-carrying capacity and therefore fails the specimen catastrophically. Three delamination failure modes exist, opening (mode I), in-plane shear (mode II) and out-of-plane shear (mode III) that can be described using the traction vector on the plane with normal vector pointing in out-of-plane direction so $(\sigma_z, \tau_{xz}, \tau_{yz})$. Each mode has an independent fracture toughness described by the material properties: G_{IC} , G_{IIC} and G_{IIIC} [38].

2.2.3. Damage Progression

In an OHT composite specimen, ultimate failure is frequently preceded by the development of sub-critical damage at load levels significantly below the final failure stress. The process initiates at the edge of the hole, where the linear elastic stress concentration is highest. Damage begins as matrix cracks forming in the fibre direction of each ply, typically at a point tangential to the hole; the exception is the 90° ply, in which the cracking is more extensive rather than confined to a tangential split [25]. Its significance lies in the notch blunting it produces: the stress peaks at the hole edge are redistributed by subcritical damage,

progressively flattening the stress gradient and reducing the stress concentrations in the load-bearing 0° plies [25]. The fibres at the hole edge therefore experience a lower effective stress than linear elastic theory predicts for a given far-field load, and a higher applied far-field stress is required to reach the critical fibre failure strain.

The ultimate failure of OHT composite specimens is determined by a competition between the development of sub-critical damage (matrix cracking and delamination) and the fracture of load-bearing fibers. Depending on the extent of this damage progression before the fibers reach their critical failure point, three distinct failure modes are identified [26].

1. **Brittle Failure:** This mode is characterized by highly limited sub-critical damage development. Catastrophic fiber failure occurs before matrix cracks have the opportunity to propagate significantly across the specimen width. This results in a clean fracture surface perpendicular to the loading direction, where all plies (except the 90° layers) fail by fiber fracture. Because the damage remains localized, the specimen experiences minimal notch blunting leading to significantly lower ultimate strengths than in pull-out failure, nonetheless delayed [26].
2. **Pull-out Failure:** Pull-out failure occurs when sub-critical damage, in the form of intra-ply cracking, has progressed across the full width of the specimen. In this mode, the load-bearing 0° fibers reach their failure strain after the cracks have already crossed the width but before delaminations can spread catastrophically. The result is a failure where off-axis plies pull out from between adjacent layers without failing by fiber fracture. This mode is associated with higher failure strengths because the extensive splitting provides maximum notch blunting, which relieves the local stress concentration at the hole [26].
3. **Delamination Failure:** To understand the delamination failure mode, first its initiation and progression have to be understood. As mentioned previously, the first subcritical damage in an open-hole tension specimen is intralaminar matrix cracking initiating at the hole boundary. Since the 90° plies are oriented perpendicular to the loading direction, matrix cracking is most pronounced in these plies, with numerous transverse cracks initiating around the hole. Cracking in the remaining orientations is markedly less extensive and initiates tangentially to the hole, running parallel to the fibre direction of the respective ply. Owing to the *in-situ* effect, the outer 45° ply is the first non- 90° ply to develop matrix cracks, as its position at the free surface reduces the *in-situ* transverse strength relative to the embedded plies.

As these cracks extend outwards from the hole, cracks in adjacent plies intersect one another. Delamination initiates at these intersections and propagates in-plane, spanning the region bounded by the cracks that caused its initiation, and thereby developing an approximately triangular damage region. Once the cracks reach the free edge of the specimen, a second delamination front initiates at the intersection of the crack with the free edge and propagates inwards towards the hole, forming a corresponding triangular region. The joining of the two delamination fronts constitutes a critical event in the damage progression: at this point the entire region bounded by the matrix cracks of the adjacent plies is delaminated, and no unobstructed in-plane growth direction remains within that interface.

Since both intralaminar matrix cracking and interlaminar delamination constitute fracture of the matrix-dominated material, the matrix cracks bounding the delaminated region provide a continuous damage path into the adjacent interface. Once in-plane propagation within the confined region is exhausted, the delamination front therefore transfers through the ply thickness along these

cracks and continues to propagate within the adjacent interface. Matrix cracks in the subsequent interfaces likewise give rise to delamination initiation at their intersections, although at higher applied loads than in the outer interface, consistent with the *in-situ* effect described above.

This sequence of in-plane propagation followed by through-thickness transfer continues until the delamination reaches the $-45^\circ/0^\circ$ interface. The matrix cracks in the 0° plies run parallel to the loading direction and therefore do not span the specimen width, so the delamination is not bounded in the direction of the grips. In-plane propagation along this interface is consequently unobstructed, and the delamination extends over a substantially greater area than at the preceding interfaces. The delamination front ultimately reaches the specimen grips, resulting in final failure of the specimen before further transfer through the laminate thickness can occur. Unlike the fiber-dominated modes, delamination failure means that any increase in sub-critical damage accelerates the progression toward ultimate collapse, leading to a reduction in specimen strength [26].

The unifying factor across these modes is the extent of sub-critical damage present at the point where the final failure stress is reached. For fiber-dominated failure (brittle and pull-out), increased matrix damage is beneficial because it increases notch blunting and specimen strength. Conversely, for the delamination mode, the same accumulation of damage is detrimental as it accelerates the path toward catastrophic structural separation.

2.2.4. Notch Sensitivity and the OHT Stress Field

The response of a laminate to a hole lies between two limiting idealisations [38]. A perfectly notch-insensitive material, the limiting behaviour of ductile metals, redistributes stress fully by yielding, so that failure is governed only by the reduced net-section area,

$$\sigma_\infty = F_{tu} \left(1 - \frac{2a}{w} \right), \quad (2.13)$$

with F_{tu} the unnotched strength, $2a$ the hole width and w the specimen width. A completely notch-sensitive (ideally brittle) material, by contrast, fails as soon as the peak stress at the hole edge reaches F_{tu} , so that $\sigma_\infty = F_{tu}/k_t$, where k_t is the elastic stress concentration factor. Experimental notched-strength data for CFRP laminates fall between these bounds: even very small holes cause a strength drop towards the notch-sensitive curve and at large hole-to-width ratios the data further approaches it. This intermediate behaviour reflects a finite process zone of matrix cracks, fibre breaks and delaminations at the hole edge that redistributes load to a limited extent [38], precisely the progressive-damage process this thesis sets out to model.

The elastic stress field that drives this process is anisotropic. For an infinite, symmetric orthotropic plate with a circular hole of radius a under remote uniaxial tension σ , the complex-variable elasticity solution of Lekhnitskii [42], generalised for arbitrary cut-outs by Savin [66], gives the circumferential stress at the hole boundary as

$$\sigma_\theta(r = a) = \frac{-K_L \cos^2 \theta + (1 + \sqrt{2K_L - m_L}) \sin^2 \theta}{\sin^4 \theta - m_L \sin^2 \theta \cos^2 \theta + K_L^2 \cos^4 \theta} \sigma, \quad (2.14)$$

where θ is measured from the load direction and the two parameters follow from the effective laminate

constants of Eq. (2.10),

$$K_L = \sqrt{\frac{E_1^L}{E_2^L}}, \quad m_L = 2\nu_{12}^L - \frac{E_1^L}{G_{12}^L}. \quad (2.15)$$

The factor multiplying σ in Eq. (2.14) is the stress concentration factor. Its maximum occurs at $\theta = 90^\circ$, the net-section point on the hole boundary. For a quasi-isotropic laminate $K_L = 1$ and $m_L = -2$, and Eq. (2.14) reduces to $k_t = -\cos^2 \theta + 3 \sin^2 \theta$, recovering the familiar isotropic value $k_t = 3$ at the net section ($\theta = 90^\circ$) and the expected compression $k_t = -1$ at the hole boundary parallel to the load ($\theta = 0^\circ$). For more orthotropic stacking sequences the factor departs from 3, so the peak stress at a hole is layup-dependent. A direct reason why notched response must be assessed per stacking sequence, and the motivation for validating the model against more than one layup.

Equivalent Anisotropic-Plate Idealisation

Equation (2.14) treats the laminate as a single equivalent homogeneous anisotropic plate rather than as a stack of discrete plies. This idealisation rests on the same assumptions already invoked to obtain Eq. (2.10): each ply is a homogeneous, linearly elastic orthotropic sheet in a state of plane stress, the laminate satisfies the Kirchhoff–Love hypothesis of CLT, the stacking sequence is symmetric so that $\mathbf{B} = \mathbf{0}$ and membrane and bending response decouple, and the equivalent plate’s in-plane compliance is set equal to the laminate membrane compliance $\mathbf{a} = \mathbf{A}^{-1}$. Because the resulting hole solution is two-dimensional by construction, the interlaminar stresses that drive delamination at the hole boundary lie outside its scope and must be assessed separately.

2.2.5. Complex-Variable (Lekhnitskii–Savin) Formulation

The onset of failure of any specimen is ultimately governed by its stress state. To validate the plausibility of the later discussed Finite Element Method (FEM) models, the predicted stress state will be compared to an analytical solution of the stress field in an OHT specimen.

An analytical approximation for the stress field can be found using Lekhnitskii’s complex potential formalism [42], as comprehensively derived and discussed by Nguyen-Hoang and Becker [55]. As this is not directly relevant to this work but only a means of validation and lies outside the scope of this thesis, only a brief summary is given.

The symmetric laminate follows the same idealisation as before, a homogeneous orthotropic plate in plane stress with smeared compliances $S_{ij} = a_{ij}$. The in-plane stresses are derived from an Airy stress function F via

$$\sigma_x = \frac{\partial^2 F}{\partial y^2}, \quad \sigma_y = \frac{\partial^2 F}{\partial x^2}, \quad \tau_{xy} = -\frac{\partial^2 F}{\partial x \partial y}, \quad (2.16)$$

which identically satisfy the equilibrium conditions. Following Lekhnitskii’s complex potential formalism [42], the Airy stress function F is found in the form $F = F(x + \mu y)$, which leads to the characteristic equation

$$S_{11} \mu^4 - 2S_{16} \mu^3 + (2S_{12} + S_{66}) \mu^2 - 2S_{26} \mu + S_{22} = 0. \quad (2.17)$$

Its roots occur in conjugate pairs, $\mu_3 = \bar{\mu}_1$, $\mu_4 = \bar{\mu}_2$, so that only μ_1, μ_2 are independent. Each independent root defines a complex variable $z_j = x + \mu_j y$, and the general solution of the governing equation is

obtained by superposing one arbitrary function per variable,

$$F = 2 \operatorname{Re}[F_1(z_1) + F_2(z_2)], \quad (2.18)$$

where taking twice the real part accounts for the conjugate roots and ensures F is real. Introducing the complex potentials as the first derivatives

$$\Phi_j(z_j) = \frac{dF_j}{dz_j}, \quad j = 1, 2, \quad (2.19)$$

the in-plane stresses follow from Eq. (2.16) as

$$\begin{aligned} \sigma_x &= 2 \operatorname{Re}[\mu_1^2 \Phi_1'(z_1) + \mu_2^2 \Phi_2'(z_2)], \\ \sigma_y &= 2 \operatorname{Re}[\Phi_1'(z_1) + \Phi_2'(z_2)], \\ \tau_{xy} &= -2 \operatorname{Re}[\mu_1 \Phi_1'(z_1) + \mu_2 \Phi_2'(z_2)]. \end{aligned} \quad (2.20)$$

For a circular hole of radius $R = d/2$ in an infinite plate under remote tension, the potentials split into a homogeneous part $\Phi_j^{\text{inf},1}$ describing the unnotched plate and a part $\Phi_j^{\text{inf},2}$ cancelling the resulting tractions along the hole edge. The homogeneous part takes the general form

$$\Phi_j^{\text{inf},1}(z_j) = \Lambda_j z_j. \quad (2.21)$$

Writing $\mu_j = \alpha_j + i\beta_j$ and $\Lambda_j = \lambda_j + i\kappa_j$, this system is underdetermined and $\kappa_1 = 0$ may be set arbitrarily, which yields

$$\begin{aligned} \lambda_1 &= \frac{1}{2} \frac{\alpha_2^2 + \beta_2^2}{D} \sigma_\infty, & \lambda_2 &= \frac{1}{2} \frac{\alpha_1^2 - 2\alpha_1\alpha_2 - \beta_1^2}{D} \sigma_\infty, \\ \kappa_1 &= 0, & \kappa_2 &= \frac{1}{2\beta_2} \frac{\alpha_2[\alpha_1(\alpha_1 - \alpha_2) - \beta_1^2] + \alpha_1\beta_2^2}{D} \sigma_\infty, \end{aligned} \quad (2.22)$$

with the common denominator $D = (\alpha_1 - \alpha_2)^2 - \beta_1^2 + \beta_2^2$. The hole correction part is expressed by means of the conformal mapping

$$\zeta_j = \frac{z_j \pm \sqrt{z_j^2 - R^2(1 + \mu_j^2)}}{R(1 - i\mu_j)}, \quad (2.23)$$

where the branch of the square root is chosen such that the stresses remain continuously differentiable and physically admissible along the net-section plane. Denoting the two branches by ζ_j^+ and ζ_j^- , this is achieved by the selection

$$\zeta_j = \begin{cases} \zeta_j^+ & \text{for } 0 \leq \varphi \leq \pi/2 \vee 3\pi/2 < \varphi \leq 2\pi, \\ \zeta_j^- & \text{for } \pi/2 < \varphi \leq 3\pi/2, \end{cases} \quad (2.24)$$

with φ denoting the circumferential coordinate. In the mapped plane the corrective potentials are expressed as,

$$\Phi_1^{\text{inf},2}(z_1) = \frac{C_1 - \mu_2 B_1}{\mu_1 - \mu_2} \zeta_1^{-1}, \quad \Phi_2^{\text{inf},2}(z_2) = -\frac{C_1 - \mu_1 B_1}{\mu_1 - \mu_2} \zeta_2^{-1}, \quad (2.25)$$

with $B_1 = -R \sigma_\infty/2$ and $C_1 = 0$ for the present load case [55]. Superposition of both parts satisfies the traction-free hole boundary while leaving the remote field unchanged, and the derivatives Φ_j' required in

Eq. (2.20) follow analytically from Eqs. (2.21) and (2.25). Evaluating Eq. (2.20) at the hole edge yields the stress concentration factor of the infinite plate,

$$K_T^{\text{inf}} = \frac{\sigma_y^{\text{inf}}(R, 0)}{\sigma_\infty} = 1 + \sqrt{\frac{2}{A_{11}} \left(\sqrt{A_{11}A_{22} - A_{12}^2} + \frac{A_{11}A_{22} - A_{12}^2}{2A_{66}} \right)}, \quad (2.26)$$

with A_{ij} denoting the in-plane laminate stiffness of classical laminate theory.

Finite specimen width is accounted for by the correction factor of Tan [70], which scales the infinite-plate net-section stresses,

$$\sigma_y^{\text{Tan}} = \frac{K_T^{\text{Tan-HW}}}{K_T^{\text{inf}}} \sigma_y^{\text{inf}}, \quad (2.27)$$

where the enhanced Tan–Heywood correction factor is adapted to the photoelasticity-based formula of Heywood [21] and reads [70, 55]

$$\frac{K_T^{\text{Tan-HW}}}{K_T^{\text{inf}}} = \left[\frac{3(1 - d/w)}{2 + (1 - d/w)^3} + \frac{1}{2} \left(\frac{d}{w} \right)^M \left(K_T^{\text{inf}} - 3 \right) \left(1 - \left(\frac{d}{w} \right)^M \right)^6 \right]^{-1}, \quad (2.28)$$

with the infinite-plate stress concentration factor of Eq. (2.26) and the exponent M following from

$$M^2 = \frac{\sqrt{1 + 8 \left[\frac{3(1 - d/w)}{2 + (1 - d/w)^3} - 1 \right]} - 1}{2(d/w)^2}. \quad (2.29)$$

It should be noted that this correction assumes the stress decay of the finite-width problem to be of the same shape as that of the infinite domain. In other words only the stress values change not the shape of the stress contour. Nguyen-Hoang and Becker [55] show that the resulting error remains small in the immediate vicinity of the hole ($0 \leq \chi \leq 0.2$, with $\chi = (x - R)/(w/2 - R)$), where the Tan–Heywood correction yields deviations of about 2 % for quasi-isotropic laminates, whereas deviations grow considerably towards the free edge.

The Isotropic Degeneracy

For quasi-isotropic laminates the characteristic equation (2.17) degenerates to the double root $\mu_1 = \mu_2 = i$, so that the two complex variables are no longer independent and the conformal mapping (2.23) becomes singular. Following Nguyen-Hoang and Becker [55], the corresponding infinite-domain solution of Kirsch [39] can be used instead, whose Airy stress function $F^{\text{inf}} = F^{\text{inf},1} + F^{\text{inf},2}$ reads

$$F^{\text{inf},1} = \frac{\sigma_\infty}{4} (1 + \cos 2\varphi) r^2, \quad F^{\text{inf},2} = \frac{\sigma_\infty}{4} \left[\left(\frac{R^4}{r^2} - 2R^2 \right) \cos 2\varphi - 2R^2 \ln \frac{r}{R} \right], \quad (2.30)$$

from which the stresses follow by differentiation as in Eq. (2.16), expressed in polar coordinates.

2.2.6. Limitations With Damage Prediction

Classical notched-strength prediction couples this stress field to a characteristic-distance criterion. The point-stress criterion of Whitney and Nuismer [78] predicts failure when the local stress at a fixed distance

d_{pt} ahead of the hole reaches the unnotched strength, while the average-stress variant requires the stress averaged over a distance a_{av} ahead of the hole to reach it. The characteristic distances d_{pt} and a_{av} are treated as material constants fitted to test data, with a_{av} commonly of the order of 3.8 mm for carbon/epoxy [38]. These criteria are attractive for preliminary sizing and capture the hole-size effect reasonably for many layups. Their limitations, however, are exactly what motivates the present work: the characteristic distance is an empirical fit with no intrinsic physical meaning, a single fitted length cannot represent the full range of stacking sequences, and the criteria predict only an onset load, they model neither the sequence of failure events nor the interaction between matrix cracking, fibre failure and delamination that determines the final strength and damage pattern.

2.2.7. Damage-Initiation Criteria

Predicting the *onset* of ply-level damage requires a criterion relating the local stress (or strain) state to the ply strengths. A large body of such criteria exists; they are conventionally divided into *mode-independent* criteria, which return a single failure index without identifying the mechanism, and *mode-dependent* criteria, which evaluate the fibre and matrix mechanisms separately [18]. For progressive-damage analysis the distinction matters, because the subsequent stiffness degradation differs by mode; the following overview therefore presents the principal criteria and their governing equations. Throughout, 1, 2, 3 denote the fibre, in-plane transverse and through-thickness ply axes; X_t, X_c, Y_t, Y_c and S_{12}, S_{23} are the longitudinal, transverse and shear strengths. These criteria are surveyed for completeness.

Mode-independent criteria. The non-interactive maximum-stress and maximum-strain criteria compare each component independently with its allowable. The interactive Tsai–Hill and Tsai–Wu criteria introduce coupling through a quadratic failure surface; the Tsai–Wu criterion [73] reads

$$\frac{\sigma_{11}^2}{X_t X_c} + \frac{\sigma_{22}^2}{Y_t Y_c} - \frac{\sigma_{11} \sigma_{22}}{\sqrt{X_t X_c Y_t Y_c}} + \left(\frac{1}{X_t} - \frac{1}{X_c} \right) \sigma_{11} + \left(\frac{1}{Y_t} - \frac{1}{Y_c} \right) \sigma_{22} + \frac{\tau_{12}^2}{S_{12}^2} = 1, \quad (2.31)$$

where the normal-stress interaction coefficient is taken as the standard approximation $F_{12} = -\frac{1}{2}(X_t X_c Y_t Y_c)^{-1/2}$. These criteria are computationally trivial and remain common in industry, but they descend from the von Mises yield condition and so rest on a plasticity analogy that is physically inappropriate for a brittle, heterogeneous composite, and they return no information on the failure mode [18].

Hashin–Rotem. Hashin and Rotem [29] were among the first to separate the mechanisms, proposing independent fibre- and matrix-failure conditions in plane stress. Fibre failure follows a maximum-stress condition $|\sigma_{11}| = X$ (with $X = X_t$ or X_c), while matrix failure is the quadratic interaction

$$\left(\frac{\sigma_{22}}{Y} \right)^2 + \left(\frac{\tau_{12}}{S_{12}} \right)^2 = 1, \quad Y = \begin{cases} Y_t & \sigma_{22} \geq 0 \\ Y_c & \sigma_{22} < 0. \end{cases} \quad (2.32)$$

The criterion is based on logical reasoning and a polynomial fit rather than on the fracture mechanism; it neither resolves a fracture-plane angle nor distinguishes the physics of tensile and compressive matrix failure.

Hashin. Hashin [28] extended the approach to a three-dimensional stress state and four distinct modes, grounding the matrix conditions in the idea that fracture is driven by the stresses on the failure plane. Tensile fibre failure ($\sigma_{11} \geq 0$) is

$$\left(\frac{\sigma_{11}}{X_t}\right)^2 + \frac{\tau_{12}^2 + \tau_{13}^2}{S_{12}^2} = 1, \quad (2.33)$$

compressive fibre failure follows $|\sigma_{11}| = X_c$, and the tensile ($\sigma_{22} + \sigma_{33} \geq 0$) and compressive ($\sigma_{22} + \sigma_{33} < 0$) matrix modes are, respectively,

$$\frac{(\sigma_{22} + \sigma_{33})^2}{Y_t^2} + \frac{\tau_{23}^2 - \sigma_{22}\sigma_{33}}{S_{23}^2} + \frac{\tau_{12}^2 + \tau_{13}^2}{S_{12}^2} = 1, \quad (2.34)$$

$$\frac{1}{Y_c} \left[\left(\frac{Y_c}{2S_{23}} \right)^2 - 1 \right] (\sigma_{22} + \sigma_{33}) + \frac{(\sigma_{22} + \sigma_{33})^2}{4S_{23}^2} + \frac{\tau_{23}^2 - \sigma_{22}\sigma_{33}}{S_{23}^2} + \frac{\tau_{12}^2 + \tau_{13}^2}{S_{12}^2} = 1. \quad (2.35)$$

The matrix conditions are smooth quadratic surfaces rather than action-plane criteria. As Hashin himself noted, the quadratic representation implies that matrix failure always occurs on the maximum transverse shear plane, $\theta = \pm 45^\circ$ [28].

Puck. Puck and Schürmann [62] introduced one of the first thoroughly physically based criteria, treating inter-fibre (matrix) failure with a Mohr–Coulomb hypothesis: fracture is governed solely by the stresses acting on the *action plane*, whose inclination θ is not known in advance. The tractions on a candidate plane rotated by θ about the fibre axis are

$$\begin{aligned} \sigma_n(\theta) &= \sigma_{22} \cos^2 \theta + \sigma_{33} \sin^2 \theta + 2\tau_{23} \sin \theta \cos \theta, \\ \tau_{nt}(\theta) &= (\sigma_{33} - \sigma_{22}) \sin \theta \cos \theta + \tau_{23}(\cos^2 \theta - \sin^2 \theta), \\ \tau_{n1}(\theta) &= \tau_{31} \sin \theta + \tau_{21} \cos \theta, \end{aligned} \quad (2.36)$$

and the inter-fibre failure index is obtained by maximising an exposure function $f(\theta)$ (given in full in [62]) over $-90^\circ \leq \theta \leq 90^\circ$; the orientation that maximises it is the predicted fracture plane. A normal tensile stress σ_n on that plane promotes fracture while a normal compressive stress impedes it, the asymmetry being controlled by empirical inclination parameters $p_{\parallel}$ and $p_{\perp}$. Fibre failure is treated separately with a strain-based fibre condition. The action-plane search gives a strong physical basis at the cost of the iterative angle search and the additional inclination parameters, for which calibration guidelines exist for CFRP and Glass Fibre Reinforced Polymer (GFRP) [62, 18].

Cuntze (Failure-Mode Concept). Cuntze and Freund [14] describe failure through five independent modes (two fibre, three matrix), each governed by a single strength and formulated in terms of stress invariants. The modes are combined into a single failure index by a probabilistic series interaction, avoiding the explicit fracture-plane search of Puck at the price of curve-fitting (rounding) exponents.

LaRC criteria. The LaRC family builds on the Mohr–Coulomb action-plane idea and the concept of in-situ strength. LaRC03 [17] is a two-dimensional set that uses a slit-crack/Mohr–Coulomb treatment for compressive matrix failure and is the first to model compressive fibre failure through fibre kinking in a misalignment frame. LaRC04 [60] extends the criteria to three dimensions, adds shear non-linearity, and removes the friction overestimation of LaRC03 at the cost of an iterative search for the misalignment and kinking angles. LaRC05 [61] replaces the separate matrix conditions with a single action-plane expression (closely related to Puck’s) that distinguishes tensile and compressive matrix failure by the sign of the normal traction, and retains a three-dimensional kinking model for fibre compression; it was among the best-performing criteria in the Second World-Wide Failure Exercise [36, 61].

In-situ strength. The transverse-tensile and in-plane-shear strengths of a ply are not material constants but depend on the ply thickness and on whether the ply is embedded between plies of different orientation, an effect first observed by Parvizi et al. [59]: thinner and embedded plies exhibit higher apparent strengths because the neighbouring plies act as crack stoppers. The in-situ strengths are estimated analytically from a slit-crack fracture analysis [10] and enter the mode-dependent criteria above simply by replacing Y_t and S_{12} with their in-situ values for each ply.

Delamination initiation. Stress-based criteria for the onset of delamination are comparatively few and of limited physical basis, because the interlaminar failure is governed by the through-thickness stress state at a ply interface. Delamination is therefore most commonly handled by a fracture-mechanics treatment with cohesive elements rather than by a ply-level stress criterion [18].

2.3. Limitations

The analytical framework assembled in this Chapter can locate the onset of failure in a notched laminate but not its evolution. CLT provides the ply stress state, the ply-level criteria of Section 2.2.7 flag the first failure event, and the closed-form notched-strength models of Whitney–Nuismer type in Section 2.2.4 collapse the entire open-hole problem onto a single empirical characteristic distance. This is adequate for preliminary sizing but no more. The OHT response that this thesis sets out to predict is governed by a sequence of interacting, discrete sub-critical events, such as matrix cracks and interlaminar delaminations, that redistribute load, and only then precipitate fibre failure [26, 2]. Capturing progressive damage therefore requires a numerical framework, and FEM is the most promising option. The extension of this concepts to FEM will be presented in the next Chapter.

Finite Element Damage Prediction Methods

This chapter gathers the Finite Element (FE) damage-prediction methods into which the semi-discrete model is placed. The presented options are all on mesoscale, the scale at which the constitutive model of this thesis operates. Constituent-scale micromechanics and the purely analytical strength methods of Chapter 2 are out of scope, except where they motivate a mesoscale choice. Within these bounds, the main differentiation used consistently across independent reviews of the field [49, 71, 80], is on how a crack is represented in the discretization. CSCM, also called CDM, never introduce a crack as a geometric entity. They degrade the stiffness of standard continuum elements so that the effect of cracking is smeared over a volume. DCM, on the other hand, introduce a true displacement discontinuity.

Section 3.1 develops the CSCM family from its scalar-damage foundations. Section 3.2 develops the DCM family. Section 3.3 then introduces the orthogonal strand of stochastic FE modelling, which targets the material variability that drives the progression of damage. Section 3.4 synthesizes these by positioning the semi-discrete method as an intersection of the continuum and discrete with its own costs stated alongside its claimed benefits.

3.1. Continuum Smeared Crack Methods

The smeared approach to progressive damage is the baseline of ply-level finite element analysis. Standard continuum elements retain their mesh topology as damage accumulates, but their stiffness is degraded through scalar internal variables that evolve with the local stress state. No crack surface is introduced, no remeshing or enrichment is required, and any unstructured mesh is admissible; which explains why this class of model became the default for industrial progressive failure analysis of composite laminates [41].

3.1.1. Scalar damage and the mesh-objectivity problem

Each failure mode, fibre tension (Fibre Tension (FT)), fibre compression (Fibre Compression (FC)), matrix tension (Matrix Tension (MT)), and matrix compression (Matrix Compression (MC)), is assigned a scalar damage variable $d_i \in [0, 1]$ that evolves from intact ($d_i = 0$) to fully degraded ($d_i = 1$) [41, 81]. Growth is driven by a mode-specific equivalent strain that collapses the stress state into a single scalar driving measure; the associated stiffness components are reduced accordingly. The constitutive update is purely local: each integration point responds only to its own strain history, with no information exchanged between neighbouring points [41].

Any strain-softening constitutive law creates a numerical problem called mesh dependence. Consider a specimen under tension that is modelled with a uniform grid of elements. As the load increases, one row of elements will reach the strength limit first due to the stress concentrations from geometry, loading, boundary conditions or even just numerical noise, and begin to soften. Once that row softens, its stiffness drops. A softer element attracts more strain under the same applied displacement, so more strain concentrates into that row, which damages it further. Meanwhile the neighbouring rows, now carrying less strain, elastically unload and stop damaging. The result is that all damage localises into a single row of elements, regardless of how many rows are present or how large the mesh is. Refine the mesh, the elements get smaller, the damaged row gets thinner, and less energy is dissipated at failure. The fracture energy the model predicts therefore depends on the mesh, not on the material. Peak load, post-peak response, and damage extent all change with mesh refinement, which is physically wrong [5, 41].

Bazant and Oh introduced a stress-strain based method known as the crack band model that gradually degrades the stiffness after a peak load based on the fracture energy G_c^i of the material [5]. Bazant and Oh state that representing a crack in an element is equivalent to a band of parallel microscopic cracks distributed over the element. The width of the band is the characteristic length L_c of the element, which is a measure of the size of the fracture process zone. The fracture energy G_c^i is the energy W that is dissipated per unit area of crack growth:

$$G_c^i = \frac{W}{A}, \quad (3.1)$$

where A is the area of the crack surface. The area under the stress-strain curve of the softening law gives the energy dissipated per unit volume g_c^i of the damaged element:

$$g_c^i = \int_0^{\varepsilon_{eq,f}^i} \sigma_{eq}^i(\varepsilon) d\varepsilon, \quad (3.2)$$

where $\varepsilon_{eq,f}^i$ is the equivalent strain at which the element is fully damaged and $\sigma_{eq}^i(\varepsilon)$ is the equivalent stress at a given equivalent strain ε . Due to the triangular shape the proposed stress-strain relation exhibits, the relation can be rewritten as

$$g_c^i = \frac{1}{2} \sigma_{eq,0}^i \varepsilon_{eq,f}^i, \quad (3.3)$$

where $\sigma_{eq,0}^i$ is the stress at the onset of damage and $\varepsilon_{eq,f}^i$ is the equivalent strain at which the element is fully damaged. Based on the definition that g_c^i is the energy dissipated per unit volume, the dissipated energy can be written as

$$W = g_c^i V = g_c^i A L_c, \quad (3.4)$$

where V is the volume of the damaged element and L_c is the width of the crack band. Comparing

Equations (3.1) and (3.4), the fracture energy can be expressed as

$$G_c^i = \frac{W}{A} = \frac{g_c^i A L_c}{A} = g_c^i L_c, \quad (3.5)$$

and therefore with Eq. (3.3) as

$$G_c^i = \frac{1}{2} \sigma_{eq,0}^i \varepsilon_{eq,f}^i L_c. \quad (3.6)$$

This can now be rearranged to give the equivalent strain at which the element is fully damaged as

$$\varepsilon_{eq,f}^i = \frac{2 G_c^i}{\sigma_{eq,0}^i L_c}. \quad (3.7)$$

This leads to smaller elements degrading their stiffness slower, effectively resisting the localisation of damage into a single row of elements and restoring the physical independence of the fracture energy from the mesh size [5, 41]. This adjustment has a limit. The failure strain cannot be smaller than the damage onset strain. This condition places an upper bound on the element size based on Eq. (3.7):

$$L_c \leq \frac{2 E_i G_c^i}{(X_i)^2}, \quad (3.8)$$

where E_i is the elastic modulus and X_i is the ply strength for mode i [41, 81]. Elements larger than this bound cannot represent the softening correctly.

The whole correction therefore depends on assigning the right value to L_c . In Abaqus, the characteristic length at an integration point is computed automatically as the cube root of the element volume [16, 18],

$$L_{Ab} = \sqrt[3]{V}. \quad (3.9)$$

This approximation causes a specific problem for thin composite ply elements. A typical ply element has a large in-plane area A_{ip} but a very small thickness t , so the volume is $V = A_{ip} t$. Because t is so much smaller than the in-plane dimensions, the cube root $\sqrt[3]{V}$ is dominated by the thickness and produces a value much smaller than the actual in-plane element size.

This matters because for an in-plane crack, the physically relevant band width is the in-plane element dimension, not the thickness. When Abaqus uses $\sqrt[3]{V}$, it enforces the fracture energy condition correctly for the length scale it computes, but that length scale is wrong for the crack geometry. The softening law ends up calibrated to the ply thickness rather than to the in-plane mesh size, so the results change when the ply thickness changes even if the in-plane mesh and the fracture properties stay the same.

The fix is to replace L_{Ab} with the square root of the in-plane element area. Starting from $V = A_{ip} t$ and $V = L_{Ab}^3$, the in-plane area is

$$A_{ip} = \frac{L_{Ab}^3}{t}, \quad (3.10)$$

and the corrected in-plane characteristic length is [18, 81]

$$L_{c,IP} = \sqrt{A_{ip}} = \sqrt{\frac{L_{Ab}^3}{t}}. \quad (3.11)$$

This recovers the original two-dimensional Bažant–Oh principle, where $L_c = \sqrt{A}$ for a planar element,

and removes the dependence on ply thickness entirely. With this correction applied at every integration point, the dissipated fracture energy is governed by G_c^i and the in-plane mesh size alone, making damage progression mesh-size independent, as confirmed by the mesh sensitivity study in Zhang et al. [81]. However, $L_{c,IP}$ reduces the in-plane geometry to a single area-based scalar: it carries no information about how the element edges are oriented or proportioned relative to the material directions. Its approximation of the crack band width therefore becomes increasingly erred as the ratios of the element edge lengths depart from unity [43], which sets a limitation on the element shape that the next section addresses directly.

3.1.2. Characteristic length for misaligned, skewed, and non-square elements

The corrected length $L_{c,IP}$ of Eq. (3.11) is exact only for an element that is square and whose edges are aligned with the material directions. Leone and Justusson [43] show that the underlying reason is geometric: the crack band width that crack band theory requires must be measured normal to the crack surface, whereas $\sqrt{A_{ip}}$ is a single direction-independent scalar. For a unidirectional ply the crack-surface orientation is fixed by the material, since matrix cracks tend to grow along the fibre material direction $\hat{f}$ and the crack-surface normal is taken parallel to the matrix material direction $\hat{m}$ [43]. An area-based scalar cannot represent a band width defined relative to $\hat{f}$ and $\hat{m}$, and the discrepancy grows once the element is misaligned, stretched, or skewed.

Leone and Justusson [43] quantify the resulting error. For a square element misaligned with respect to the material directions by 45° , using a typical line length across the element instead of the crack band width overestimates the dissipated fracture energy by 41%; not accounting for element aspect ratio underestimates it by 29% and 50% for aspect ratios of two and four, respectively [43]. To remove these errors, the characteristic length must be computed from the element geometry rather than supplied as a single scalar.

Leone and Justusson [43] restrict the derivation to matrix cracks in unidirectional material represented in three-dimensional reduced-integration solid elements with one element through the ply thickness t , where the thickness edges are a perpendicular extrusion of the in-plane geometry. Under these assumptions the through-thickness orientation of the matrix crack is constant, and the problem can be treated as two-dimensional in the fibre–matrix plane [43].

Edge vectors and mesh attributes

Because a general quadrilateral element does not need to have parallel opposite edges, each of the two in-plane directions can be described by the average of the pair of edge vectors opposite to one another [43]:

$$\vec{e}_{pair1} = \frac{1}{2}(\vec{x}_{n2/n1} - \vec{x}_{n4/n3}), \quad (3.12)$$

$$\vec{e}_{pair2} = \frac{1}{2}(\vec{x}_{n3/n2} - \vec{x}_{n1/n4}), \quad (3.13)$$

where $\vec{x}_{ni/nj}$ denotes the position of node i relative to node j . The average edge vector most closely aligned with the fibre material direction is taken as the fibre edge vector $\vec{e}_f$ [43]:

$$\vec{e}_f = \begin{cases} \vec{e}_{pair1}, & \text{if } (\hat{e}_{pair1} \cdot \hat{f}) \geq (\hat{e}_{pair2} \cdot \hat{f}), \\ \vec{e}_{pair2}, & \text{otherwise,} \end{cases} \quad (3.14)$$

with the matrix edge vector $\vec{e}_m$ assigned to the remaining pair and the hat operator denoting a unit vector. Three scalar attributes of the element then describe its geometry relative to the material frame [43]. The mesh misalignment ψ is the angle between the fibre material direction and the fibre edge vector,

$$\psi = \arccos(\hat{e}_f \cdot \hat{f}), \quad (3.15)$$

the element aspect ratio α is the ratio of the fibre edge length to the matrix edge length,

$$\alpha = \frac{\|\vec{e}_f\|}{\|\vec{e}_m\|}, \quad (3.16)$$

and the element skew γ_s is the interior angle between the fibre and matrix edge vectors,

$$\gamma_s = \arccos(\hat{e}_f \cdot \hat{e}_m). \quad (3.17)$$

Crack band width for growth along the mesh lines

For each element in which a matrix crack develops, the crack extends a distance Δa and dissipates fracture energy W equal to the product of the toughness and the newly created surface area $A = t \Delta a$ [43]. From equation (3.1) then follows

$$W = G_c t \Delta a. \quad (3.18)$$

When the crack grows along the mesh lines, the crack extension per failed element is the projection of the fibre edge vector onto the fibre material direction [43],

$$\Delta a_{ML} = \vec{e}_f \cdot \hat{f} = \|\vec{e}_f\| \cos \psi, \quad (3.19)$$

which depends on the misalignment ψ but is independent of the aspect ratio α and the skew γ_s . For a fibre-aligned mesh it reduces to the fibre edge length $\|\vec{e}_f\|$ [43]. From equations (3.4) and (3.5) follows

$$W = \frac{G_c}{L_c} V. \quad (3.20)$$

The volume of the element V can be expressed by the parallelogram spanned by the two edge vectors extruded through the thickness:

$$V = A t = \|\vec{e}_f\| \|\vec{e}_m\| \sin \gamma_s t. \quad (3.21)$$

With this description of the volume equation (3.20) then reads

$$W = \frac{G_c}{L_c} \|\vec{e}_f\| \|\vec{e}_m\| \sin \gamma_s t. \quad (3.22)$$

Substituting Δa_{ML} (equation (3.19)) for Δa in equation (3.18) and then eliminating W using equation (3.22) yields

$$\frac{G_c}{L_c} t \|\vec{e}_f\| \|\vec{e}_m\| \sin \gamma_s = G_c t \|\vec{e}_f\| \cos \psi, \quad (3.23)$$

which, solved for the crack band width, gives the mesh-line characteristic length [43]

$$L_{c,ML} = \frac{\|\vec{e}_m\| \sin \gamma_s}{\cos \psi}, \quad (3.24)$$

with $L_{c,ML}$ denoting characteristic length L_c for the currently discussed case of a crack growing along the mesh lines.

Correction for growth along the fibre direction

Matrix cracks grow along the mesh lines only for small misalignment, below approximately $\psi = 15^\circ$ [43]. At larger misalignment the crack instead advances along the fibre material direction, which requires the periodic failure of “stacked” elements and produces a zig-zag path. More elements fail per unit of crack advance than mesh-line growth assumes, so the model overestimates the dissipated fracture energy by the factor $\Delta a_{ML}/\Delta a_{avg}$, where Δa_{avg} is the average crack extension per failed element [43]. As seen in Fig. 3.1 the maximum crack extension across an element, obtained by applying the law of sines to the element geometry, is [43]

$$\Delta a_{max} = \begin{cases} \frac{\|\vec{e}_m\| \sin \gamma_s}{\sin \psi}, & \text{if } \vec{e}_f \cdot \hat{m} > \vec{e}_m \cdot \hat{m}, \\ \frac{\|\vec{e}_f\| \sin \gamma_s}{\sin(\psi + \gamma_s)}, & \text{otherwise,} \end{cases} \quad (3.25)$$




Figure 3.1: Crack extensions and crack band widths for a misaligned element. [43]

and averaging the crack extension over the element width yields [43]

$$\Delta a_{avg} = \begin{cases} \Delta a_{max} \frac{\vec{e}_f \cdot \hat{m}}{(\vec{e}_f + \vec{e}_m) \cdot \hat{m}}, & \text{if } \vec{e}_f \cdot \hat{m} > \vec{e}_m \cdot \hat{m}, \\ \Delta a_{max} \frac{\vec{e}_m \cdot \hat{m}}{(\vec{e}_f + \vec{e}_m) \cdot \hat{m}}, & \text{otherwise.} \end{cases} \quad (3.26)$$

Correcting the mesh-line band width by this factor defines the fibre-direction characteristic length [43],

$$L_{c,FD} = L_{c,ML} \frac{\Delta a_{ML}}{\Delta a_{avg}}, \quad (3.27)$$

which is the definition implemented in the present work. For rectangular elements the skew terms drop

out and Eqs. (3.26) and (3.27) simplify to [43]

$$\Delta a_{\text{avg}} = \frac{\alpha \|\vec{e}_m\|}{\alpha \sin|\psi| + \cos \psi}, \quad (3.28)$$

$$L_{c,\text{FD}} = \|\vec{e}_m\| (\alpha \sin|\psi| + \cos \psi). \quad (3.29)$$

Equation (3.29) shows explicitly how the misalignment and the aspect ratio enter the band width that $L_{c,\text{IP}}$ omits, and at $\psi = 0$ it returns the matrix edge length independently of α .

Comparison of the definitions

The three definitions $L_{c,\text{IP}}$, $L_{c,\text{ML}}$, and $L_{c,\text{FD}}$ coincide for a fibre-aligned square element and diverge as the mesh departs from it [43]. For an aspect ratio of one, all three are equal at $\psi = 0^\circ$, and $L_{c,\text{ML}}$ and $L_{c,\text{FD}}$ are also equal at $\psi = 45^\circ$; between these values, the error of $L_{c,\text{ML}}$ relative to $L_{c,\text{FD}}$ reaches a maximum of 17% at $\psi = 22.5^\circ$, and the error of $L_{c,\text{IP}}$ relative to $L_{c,\text{FD}}$ reaches a maximum of 29% at $\psi = 45^\circ$, with these relative errors being unaffected by element size when $\alpha = 1$ [43]. Non-unit aspect ratios widen the differences: the error of $L_{c,\text{IP}}$ peaks at $\psi = 0^\circ$ with a magnitude of $\sqrt{\alpha} - 1$, giving 41% at $\alpha = 2$ and 100% at $\alpha = 4$ [43].

Which definition best describes the crack band width depends on how the crack traverses the mesh [43]. For low misalignment and cracks that grow along the mesh lines, $L_{c,\text{ML}}$ provides an exact description of the band width. For larger misalignment, long crack extensions, or long fracture process zones spanning many elements, $L_{c,\text{FD}}$ is the more consistent choice [43].

3.1.3. Mesoscale CDM formulation for composite laminates

The standard engineering realisation for ply-level FE is the formulation of Lapczyk and Hurtado [41], extended to three-dimensional solid elements by Zhang et al. [81]. The constitutive response follows the three stages of the bilinear triangle (Figure 3.2), with each of the four damage modes — fibre tension (FT), fibre compression (FC), matrix tension (MT), and matrix compression (MC) — evolving independently.

Stage 1 — Elastic loading. The ply is characterised by its undamaged orthotropic stiffness $\mathbf{C}$. All stresses are evaluated in the effective (undamaged) space,

$$\tilde{\boldsymbol{\sigma}} = \mathbf{C} \boldsymbol{\varepsilon}, \quad (3.30)$$

until a Hashin-type initiation criterion fires for any mode (see Section 2.2.7).

Stage 2 — Damage initiation. Each mode i is characterised by a scalar equivalent strain that collapses the three-dimensional strain state into a single driving measure [81]:

$$\varepsilon_{eq}^{ft} = \sqrt{\langle \varepsilon_{11} \rangle^2 + \gamma_{12}^2 + \gamma_{13}^2}, \quad \varepsilon_{eq}^{fc} = \langle -\varepsilon_{11} \rangle, \quad (3.31)$$

$$\varepsilon_{eq}^{mt} = \sqrt{\langle \varepsilon_{22} \rangle^2 + \langle \varepsilon_{33} \rangle^2 + \gamma_{12}^2 + \gamma_{23}^2 + \gamma_{13}^2}, \quad \varepsilon_{eq}^{mc} = \sqrt{\langle -\varepsilon_{22} \rangle^2 + \langle -\varepsilon_{33} \rangle^2 + \gamma_{12}^2 + \gamma_{23}^2 + \gamma_{13}^2}, \quad (3.32)$$

where $\langle \cdot \rangle = \max(\cdot, 0)$ is the Macaulay bracket. Tensile modes are driven by positive normal strains, compressive modes by negative ones, and shear strains contribute to all four. When mode i first reaches its

Hashin threshold, the equivalent strain and stress at that instant are locked in as the peak of the triangle:

$$\varepsilon_{eq,0}^i = \varepsilon_{eq}^i|_{\text{init}} \text{ and } \sigma_{eq,0}^i = \sigma_{eq}^i|_{\text{init}}.$$

Stage 3 — Linear softening and energy balance. With the peak $(\varepsilon_{eq,0}^i, \sigma_{eq,0}^i)$ fixed, the failure strain at which mode i is fully degraded follows from requiring the triangular area to equal the fracture energy per unit crack area. The crack-band condition of Equation (3.7) previously shown. The triangle is now fully defined: elastic slope C to the peak, then linear softening to zero at $\varepsilon_{eq,f}^i$ (Figure 3.2).




Figure 3.2: Bilinear stress–strain law for a single damage mode [41, 81].

The per-mode damage variable that traces the softening branch is [41, 81]

$$d_i = \frac{\varepsilon_{eq,f}^i (\varepsilon_{eq}^i - \varepsilon_{eq,0}^i)}{\varepsilon_{eq}^i (\varepsilon_{eq,f}^i - \varepsilon_{eq,0}^i)}, \quad d_i \in [0, 1]. \quad (3.33)$$

Irreversibility is enforced by tracking the peak equivalent strain ever reached, $\bar{\varepsilon}_{eq}^i = \max_{t' \leq t} \varepsilon_{eq}^i(t')$: d_i can only accumulate, never recover.

Global damage variables and stress update. Because tension and compression can both damage the same fibre or matrix direction, global scalars combine the sub-mode variables [81]:

$$d_f = 1 - (1 - d_{ft})(1 - d_{fc}), \quad (3.34)$$

$$d_m = 1 - (1 - d_{mt})(1 - d_{mc}), \quad (3.35)$$

$$d_s = 1 - (1 - S_{mt} d_{mt})(1 - S_{mc} d_{mc}). \quad (3.36)$$

These enter a diagonal operator that degrades the effective stress to the nominal stress used in the equilibrium equations [81]:

$$\sigma = \mathbf{M}_d^{-1} \tilde{\sigma}, \quad \mathbf{M}_d^{-1} = \text{diag}[(1 - d_f), (1 - d_m), (1 - d_m), (1 - d_s), (1 - d_s), (1 - d_s)]. \quad (3.37)$$

An equivalent formulation reaches the same result through the damaged stiffness matrix directly. The undamaged orthotropic compliance $\mathbf{S}$ is modified by replacing the diagonal moduli with their degraded

counterparts, yielding the damaged compliance [81]

$$\mathbf{S}_d = \begin{bmatrix} \frac{1}{(1-d_f)E_1} & \frac{-\nu_{21}}{E_2} & \frac{-\nu_{31}}{E_3} & 0 & 0 & 0 \\ \frac{-\nu_{12}}{E_1} & \frac{1}{(1-d_m)E_2} & \frac{-\nu_{32}}{E_3} & 0 & 0 & 0 \\ \frac{-\nu_{13}}{E_1} & \frac{-\nu_{23}}{E_2} & \frac{1}{(1-d_m)E_3} & 0 & 0 & 0 \\ 0 & 0 & 0 & \frac{1}{(1-d_s)G_{23}} & 0 & 0 \\ 0 & 0 & 0 & 0 & \frac{1}{(1-d_s)G_{13}} & 0 \\ 0 & 0 & 0 & 0 & 0 & \frac{1}{(1-d_s)G_{12}} \end{bmatrix}, \quad (3.38)$$

and inverting gives the damaged stiffness $\mathbf{C}_d = \mathbf{S}_d^{-1}$, used in the stress update $\boldsymbol{\sigma} = \mathbf{C}_d \boldsymbol{\varepsilon}$, this shows the degraded slope C_d shown in Figure 3.2. The two formulations are algebraically equivalent; the M-operator form separates the elastic and damage contributions cleanly, while the $\mathbf{C}_d$ form shows directly how each modulus is reduced by its damage variable.

When any sub-mode damage variable exceeds 0.9999, the element is deleted.

Zhang et al. [81] validated this 3D formulation at multiple scales. At the element level, all four damage modes are correctly characterised with smooth evolution curves, and the model handles non-monotonic tensile–compressive loading correctly, accumulating damage history across load reversals, a capability absent from the two-dimensional Hashin built-in of Abaqus, which overestimates residual stiffness under reversed loading because it cannot simultaneously activate tensile and compressive damage in the same direction [81]. Mesh objectivity was also confirmed: with the in-plane characteristic length of Eq. (3.11), the force–displacement response is consistent across a fourfold range of mesh densities, whereas the two-dimensional built-in shows progressive loss of stiffness degradation rate with refinement [81].

At the laminate level, the formulation was applied to a quasi-isotropic $[90/0/\pm 45]_{3S}$ CFRP laminate with five open-hole specimens spanning hole diameters $d = 2\text{--}10$ mm at constant $w/d = 6$, benchmarked against the experimental data of Camanho, Maimí and Dávila [9]. The 3D model predicted the ultimate failure stress with a mean absolute error of 6.0%, compared with 9.3% for the 2D Abaqus built-in; the 2D model consistently overestimated experiment (delaying the onset of softening) and its inability to accumulate simultaneous tensile and compressive damage in the same direction [81]. Both models reproduced the experimental size effect, whereby ultimate strength decreases monotonically with specimen scale. The predicted failure pattern agreed with the experimentally observed fracture [9, 81].

3.1.4. Inherent limitations and the motivation for discrete extensions

Despite continued development of the CDM framework, fundamental deficiencies remain when the observed fracture involves a strong coupling between intralaminar matrix cracking and interlaminar delamination, or when the dominant failure mode is matrix splitting with pull-out rather than diffuse damage accumulation [64]. These deficiencies trace to two properties shared by virtually all practical CDM implementations.

The first is *locality*: damage is evaluated at each integration point from the stress state at that point alone, without considering the damage state at neighbouring locations [64]. The direction of damage evolution

is therefore governed entirely by the instantaneous local stress distribution. The failure criterion at an individual point may correctly identify the direction of microcrack initiation, but the *sequence* of those initiations — which determines the macroscopic crack path through the redistribution of stress around the growing damage zone — is not tracked. Consequently, the local direction of cracking may be correctly predicted at every point while the macroscopic path is predicted incorrectly [64].

The second is *homogenisation*: the ply is represented as a single homogeneous orthotropic medium, which eliminates the physical distinction between fibres and matrix at the constitutive level. A CDM model therefore cannot distinguish a crack propagating along the fibre direction from one propagating across fibres, even though these two failure modes dissipate very different amounts of energy and have fundamentally different structural consequences [64]. The difficulty is sharpest in shear-dominated failure: consider two ply configurations with the same applied in-plane shear field, one in which the physically expected crack path runs parallel to the fibres (low-energy matrix splitting) and one in which it runs across fibres (higher-energy diffuse matrix cracking). Because the local stress at every integration point is identical, the CDM model predicts an identical sequence of initiation events and hence an identical damage band in both cases [64].

These two fundamental limitations generate several compounding numerical deficiencies. *Mesh-orientation sensitivity* is the most direct consequence: because the softening band must form within existing elements, it localises along the preferred directions of the mesh — typically element edges or element diagonals — rather than the physically correct fracture plane [64]. The crack-band approach addresses mesh-*size* objectivity through Eqs. (3.7)–(3.11), but not mesh-*direction* objectivity. The standard mitigation is to align each ply’s mesh with its fibre direction. Millen, Ullah and Falzon [47] demonstrated across open-hole tension, end-notched tension, low-velocity impact and compression-after-impact specimens that fibre-aligned meshes consistently improve crack-path prediction and failure-load accuracy over conventional structured meshes. However, aligning each ply individually produces non-coincident meshes through the laminate thickness [64].

Shear locking is a further consequence of the smeared-crack kinematics. A CDM crack is represented by degrading the transverse modulus E_2 and shear modulus G_{12} towards zero, which should eliminate stress transfer across the open crack faces. Standard quadrilateral elements, however, couple the in-plane shear strain γ_{12} and the longitudinal strain ε_{11} unless their edges are aligned with the softening band or oriented at 45° to it [64]. When this coupling is present, full shear stress relaxation across the open crack does not occur, and the stress redistribution around the damage zone is inaccurately predicted. As this problem stems from shear locking in the underlying element formulation, other element types that do not suffer from shear locking prevent this issue.

The deepest kinematic limitation is that a single integration point must simultaneously represent the constitutive response of the bulk material on *both* sides of an emerging crack and the evolving crack opening itself. Stress components normal to the crack plane are correctly freed once damage localises, but stress components parallel to the crack plane are shared between the two flanking material volumes through their common integration point, artificially coupling responses that ought to be independent [64]. The natural remedy is to replace the smeared band with two solid elements separated by a cohesive zone, which is precisely the discrete crack formulation developed in Section 3.2.

Together these limitations define the motivation for the semi-discrete model of this thesis. The mesh-orientation and crack-path deficiencies are addressed by the matrix-split elements, which pre-embed the

fibre-aligned fracture plane in the mesh topology so that the crack band is geometrically constrained to propagate in the physically correct direction. The interlaminar kinematic limitation is addressed by cohesive elements placed at each ply interface.

3.2. Discrete Crack Methods (DCM)

Where continuum methods smear a crack over a finite volume, discrete methods represent it as a true displacement discontinuity — an explicit jump in the displacement field resolved at a surface, an enriched set of degrees of freedom, or a pair of phantom nodes. This distinction is not merely conceptual: discrete methods directly recover the crack-opening and sliding displacements, the traction distribution along the crack faces, and the energy released at the crack front, without inferring them from a degraded stiffness tensor. The price is that the discontinuity must be introduced explicitly, and the computational effort scales with the number of discontinuities tracked [49]. Three families are relevant to laminated composite failure analysis and are developed here in order of increasing kinematic generality.

3.2.1. Virtual Crack Closure Technique (VCCT)

The Virtual Crack Closure Technique (Virtual Crack Closure Technique (VCCT)) is grounded in the energy equivalence that the strain energy released when a crack advances by an incremental area ΔA equals the work required to close that same increment back to its original state [65, 67]. Rybicki and Kanninen [65] exploited this equivalence numerically without physically closing the crack: the strain energy release rate (Strain Energy Release Rate (SERR)) is extracted from a single static analysis using the nodal reaction forces at the crack front and the relative displacements of crack-face nodes immediately in the wake.

For a crack propagating in the x -direction with crack-front element of in-plane length Δa and out-of-plane thickness Δz (crack-face area $\Delta A = \Delta a \Delta z$), the mode-separated SERR components are

$$G_I = \frac{1}{2 \Delta A} F_{y,i} (v_i - v_{i^*}), \quad (3.39)$$

$$G_{II} = \frac{1}{2 \Delta A} F_{x,i} (u_i - u_{i^*}), \quad (3.40)$$

$$G_{III} = \frac{1}{2 \Delta A} F_{z,i} (w_i - w_{i^*}), \quad (3.41)$$

where $F_{y,i}$, $F_{x,i}$, $F_{z,i}$ are the nodal reaction forces at the crack-front node i in the normal (y), in-plane shear (x) and out-of-plane shear (z) directions, and $(v_i - v_{i^*})$, $(u_i - u_{i^*})$, $(w_i - w_{i^*})$ are the corresponding relative displacements between the upper and lower crack-face nodes i and i^* in the wake of the crack front [65]. The total SERR is $G_T = G_I + G_{II} + G_{III}$.

Mixed-mode propagation. Delaminations in composites propagate under combined mode I opening and mode II/III sliding, and a mixed-mode fracture criterion is required. The Benzeggagh–Kenane (Benzeggagh–Kenane (BK)) criterion [7], obtained from Double Cantilever Beam (mode I), End-Notched Flexure (mode II) and Mixed-Mode Bending tests, interpolates the critical SERR as

$$G_c = G_{Ic} + (G_{IIc} - G_{Ic}) \left(\frac{G_{II} + G_{III}}{G_T} \right)^\eta, \quad (3.42)$$

where G_{Ic} and G_{IIc} are the pure-mode critical energy release rates, and η is a curve-fit exponent determined from mixed-mode fracture tests. Propagation is predicted when $G_T \geq G_c$. The BK criterion assumes $G_{IIc} = G_{IIIc}$ [7], so modes II and III are grouped as a single shear contribution; the 2D degeneration ($G_{III} = 0$) recovers the original two-mode form of the criterion [7].

Limitations. The SERR expressions of Equations (3.39)–(3.41) require only standard nodal output and no specialised element formulation, so VCCT adds negligible computational overhead to the FE analysis [65, 67]. Its fundamental restriction is that a pre-existing crack is mandatory: VCCT evaluates the work to close an already-open crack face and therefore cannot predict initiation from an intact interface [65, 67]. Furthermore, the self-similar growth assumption breaks down when the delamination front changes shape substantially during loading, and the computed SERR is sensitive to the crack-front element size, requiring mesh-convergence studies to verify accuracy [67]. A separate node-release or moving-mesh procedure must also be implemented to advance the crack front [67]. These constraints motivate the cohesive zone approach.

3.2.2. Cohesive Zone Model (CZM)

Physical basis

Cohesive Zone Modelling (CZM) replaces the singular crack tip of linear elastic fracture mechanics with a finite-length fracture process zone in which the crack faces carry finite, decreasing cohesive tractions as they separate [19, 3, 8]. The concept originates independently with Dugdale [19], who modelled the plastic yield strip ahead of a crack in ductile metals, and Barenblatt [3], who introduced cohesive molecular forces to regularise the brittle crack tip. In both formulations the stress singularity is removed, the process-zone length is set by material properties rather than mesh size, and crack advance is governed by an energy balance rather than a stress-intensity threshold [19, 3]. The critical insight for finite element practice is that this physics translates directly into an interface constitutive law: a traction–separation relationship is prescribed on the crack surface, the area under the law equals the mode fracture energy, and both initiation and propagation are captured in a single formulation [8].

Traction–separation law

A zero-thickness interface element is placed on the crack plane; its nodes sample the upper and lower ply surfaces and relate the traction vector $\mathbf{t} = \{t_n, t_s, t_t\}^T$ to the separation vector $\Delta = \{\Delta_n, \Delta_s, \Delta_t\}^T$ (normal and two shear components) through an uncoupled penalty stiffness [45]:

$$\mathbf{t} = \begin{Bmatrix} t_n \\ t_s \\ t_t \end{Bmatrix} = \begin{bmatrix} K_n & 0 & 0 \\ 0 & K_s & 0 \\ 0 & 0 & K_t \end{bmatrix} \begin{Bmatrix} \Delta_n \\ \Delta_s \\ \Delta_t \end{Bmatrix} = \mathbf{K}_p \Delta, \quad (3.43)$$

where K_n , K_s and K_t are chosen large enough to enforce displacement continuity without introducing artificial compliance; values of order 10^6 N/mm³ are standard for carbon–epoxy systems [8, 67]. Compressive normal separations ($\Delta_n < 0$) carry a contact penalty but do not drive damage. The interface then passes through four stages that together trace the bilinear triangle shown in Figure 3.4.

Stage 1 — Elastic loading and effective scalar measures. Under intact conditions the tractions are purely elastic: $\mathbf{t} = \mathbf{K}_p \Delta$. Because the three components evolve together under mixed-mode loading, they are collapsed into two scalars that serve as the axes of the bilinear triangle in Figure 3.4. The effective mixed-mode separation is [45]

$$\Delta_e = \sqrt{\langle \Delta_n \rangle^2 + \Delta_s^2 + \Delta_t^2}, \quad (3.44)$$

and the effective traction is [45]

$$\tau_e = \frac{t_n \langle \Delta_n \rangle + t_s \Delta_s + t_t \Delta_t}{\Delta_e}, \quad (3.45)$$

where $\langle \cdot \rangle = \max(0, \cdot)$ excludes compressive normal contributions from both quantities.

Stage 2 — Damage initiation. Damage begins when the component tractions satisfy the quadratic nominal-stress criterion [45]:

$$\left(\frac{\langle t_n \rangle}{t_n^0} \right)^2 + \left(\frac{t_s}{t_s^0} \right)^2 + \left(\frac{t_t}{t_t^0} \right)^2 = 1, \quad (3.46)$$

where t_n^0 , t_s^0 and t_t^0 are the pure-mode interface strengths. Geometrically, Equation (3.46) defines an ellipsoid in the three-dimensional traction space (t_n, t_s, t_t) , whose semi-axes are t_n^0 , t_s^0 , t_t^0 . Any traction state whose representative point lies *inside* the ellipsoid leaves the interface undamaged; damage initiates the moment the point first touches the ellipsoidal surface. Figure 3.3 illustrates this with example strength values: a traction state at the origin (zero loading) is entirely safe, whereas a state at the surface, for instance, pure normal traction equal to t_n^0 , marks the onset of degradation.




Figure 3.3: Ellipsoidal damage-initiation surface in traction space defined by Equation (3.46). Traction states inside the ellipsoid are purely elastic; a state on the surface triggers initiation. Semi-axes equal the pure-mode interface strengths (illustrative values shown).

When the criterion is first satisfied, the effective separation and effective traction at that instant are locked in as the peak of the bilinear triangle:

$$\Delta_e^0 = \Delta_e|_{\text{init}}, \quad \tau_e^c = \tau_e|_{\text{init}}. \quad (3.47)$$

Stage 3 — Mixed-mode toughness and triangle geometry. With the peak (Δ_e^0, τ_e^c) established, the mode-I, II and III energy contributions at initiation are

$$G_i = \frac{1}{2} t_i \Delta_i, \quad i = n, s, t, \quad (3.48)$$

and the mode energy fractions $\lambda_i = G_i / (G_n + G_s + G_t)$ characterise the prevailing mode mix at that moment. The mixed-mode fracture toughness G_{mc} is then evaluated from one of two criteria. The

power-law criterion [45] treats all three modes independently:

$$G_{mc} = \left[\left(\frac{\lambda_n}{G_{Ic}} \right)^\alpha + \left(\frac{\lambda_s}{G_{IIc}} \right)^\alpha + \left(\frac{\lambda_t}{G_{IIIc}} \right)^\alpha \right]^{-1/\alpha}, \quad (3.49)$$

where α is a material-fit exponent and G_{Ic} , G_{IIc} , G_{IIIc} are the pure-mode critical toughnesses. The Benzeggagh–Kenane (BK) criterion (Equation (3.42)) is a simpler alternative that groups modes II and III as a single shear contribution through one exponent η , assuming $G_{IIc} = G_{IIIc}$ [7, 8].

Given G_{mc} , the only remaining unknown is the critical separation Δ_e^f at which the interface fully fails. Setting the area of the bilinear triangle equal to the fracture energy gives

$$G_{mc} = \frac{1}{2} \tau_e^c \Delta_e^f \implies \Delta_e^f = \frac{2 G_{mc}}{\tau_e^c}. \quad (3.50)$$

The triangle is now fully defined: it rises linearly from the origin with slope K_p , peaks at (Δ_e^0, τ_e^c) set by the initiation criterion, and descends linearly to zero at Δ_e^f set by the energy balance. Figure 3.4 shows the complete construction.




Figure 3.4: Bilinear traction–separation law constructed at the moment of damage initiation [45].

Stage 4 — Damage evolution. Once the peak is passed, a scalar damage variable $d_{\text{coh}} \in [0, 1]$ degrades all traction components proportionally. To enforce thermodynamic irreversibility, d_{coh} is driven by Δ_e^{max} —the largest effective separation reached in the loading history—rather than the instantaneous value; stiffness lost during loading is not recovered on unloading [45]. The bilinear softening law gives

$$d_{\text{coh}} = \frac{\Delta_e^f (\Delta_e^{\text{max}} - \Delta_e^0)}{\Delta_e^{\text{max}} (\Delta_e^f - \Delta_e^0)}, \quad (3.51)$$

and the degraded component tractions follow as

$$t_i = (1 - d_{\text{coh}}) K_i \Delta_i, \quad \Delta_e^0 < \Delta_e^{\text{max}} < \Delta_e^f. \quad (3.52)$$

When $d_{\text{coh}} \rightarrow 1$ (i.e. $\Delta_e^{\text{max}} \rightarrow \Delta_e^f$), all tractions vanish and the interface is traction-free; the element is then deleted from the simulation.

Implementations. Two standard FE implementations are in use [8, 67]. Zero-thickness cohesive elements are inserted between the ply meshes; the element nodes sample the upper and lower ply surfaces and the relative nodal displacement drives the Traction–Separation Law (TSL). Surface-based cohesive contact instead assigns the TSL to the interaction between two ply surfaces. Both give a conforming interface when the adjacent plies carry the same surface mesh. The two differ once the plies are meshed differently: the cohesive layer is itself meshed, so it can be built to conform to both ply surfaces, whereas cohesive contact inherits whatever discretisation the ply surfaces happen to have and the interface is then necessarily non-conforming. Both are available in Abaqus and both appear in the laminate literature; the choice between them is revisited when the semi-discrete model’s interlayer meshing strategy is introduced in Section 3.4.1.

Parameter calibration and the cohesive-zone length

Two practical issues dominate CZM parameter selection. The first is the cohesive-zone length: the fracture process zone must be spanned by enough elements to integrate the traction distribution accurately; too few elements underestimate the dissipated energy, producing spurious snap-back in the load-displacement response [74]. Turon *et al.* [74] derive the estimate

$$l_{cz} = M E \frac{G_c}{(\tau^0)^2}, \quad (3.53)$$

where E is the transverse (through-thickness) modulus of the adjacent ply, τ^0 is the interface strength, G_c is the fracture energy, and M is a parameter that varies from 0.21 – 1 depending on the used CZM [74]. Equation (3.53) reveals that l_{cz} scales inversely with the *square* of the interface strength: for carbon–epoxy systems with high interlaminar strengths and relatively low fracture energies, the process zone spans only tens of micrometres [74], demanding element sizes that are impractical at the laminate scale.

The second issue is the consequence of mesh coarseness. Turon *et al.* [74] propose an engineering remedy: deliberately reduce τ^0 below its physical value so that the numerically resolved process zone $l_{cz} = N_e l_e$ spans a chosen number of elements N_e of size l_e [74]. Sources disagree on N_e , with higher values giving more accurate results. For a given element size, there exists a N_e for which increasing it further does not significantly improve the results [74]. Inverting Equation (3.53) gives the adjusted strength as

$$\tau^0 = \sqrt{M \frac{E G_c}{N_e l_e}}. \quad (3.54)$$

Lu *et al.* [45] qualify this remedy by distinguishing *strength-sensitive* from *strength-insensitive* problems: in the former, reducing τ^0 corrupts the predicted initiation load; in the latter it is safe and the coarse-mesh result converges to the correct propagation response. Applying strength reduction indiscriminately thus risks over-predicting the onset load in strength-sensitive geometries, and a geometry check against the criterion of Lu *et al.* is advisable before invoking Equation (3.54). Additional sensitivity to the form of the TSL (bilinear, exponential, trapezoidal) and the choice of mixed-mode propagation criterion is documented by de Oliveira and Donadon [58], and consolidated practical guidance for coarse-mesh laminate simulations is given by Song, Dávila and Rose [67].

CZM for intralaminar cracks and its limit

CZM is naturally suited to delamination, where the crack plane between plies is known in advance. It can also represent intralaminar matrix cracks and fibre-aligned splits by pre-placing cohesive interfaces at anticipated crack locations within each ply. The discrete-ply model (Discrete-Ply Model (DPM)) of Achard *et al.* [1] takes this to its limit: cohesive interfaces are inserted *a priori* along the fibre direction within each ply and between each pair of plies, reproducing the splitting, delamination and fibre-failure interaction observed in OHT tests across four layups with very few calibration parameters. The approach is capable and parameter-efficient, but genuinely off-axis, branching or curved crack paths cannot be captured without re-meshing [1], and the *a priori* placement demands foreknowledge — or a strong prior assumption — of where cracks will propagate. For delamination, where the ply interface is the crack plane by definition, this is no constraint; for general intralaminar cracking it is a fundamental kinematic limitation, motivating the mesh-independent enrichment methods of the next section.

3.2.3. Extended Finite Element Method and Mesh-Independent Cracking

Both VCCT and CZM require the crack plane to be prescribed before the analysis: as a pre-existing SERR path or as a pre-placed interface. For intralaminar matrix cracks whose trajectory is solution-dependent and not known in advance, a family of element technologies introduces the displacement discontinuity independently of the mesh topology [48, 46, 32].

XFEM. The extended finite element method (eXtended Finite Element Method (XFEM)) enriches the standard displacement interpolation with discontinuous functions so that a crack can be represented without remeshing [48]. Moës, Dolbow and Belytschko [48] formalised XFEM through the partition-of-unity framework: the displacement field is written as

$$\mathbf{u}(\mathbf{x}) = \underbrace{\sum_{i \in \mathcal{N}} N_i(\mathbf{x}) \mathbf{u}_i}_{\text{standard}} + \underbrace{\sum_{j \in \mathcal{N}_H} N_j(\mathbf{x}) H(\mathbf{x}) \mathbf{a}_j}_{\text{jump}} + \underbrace{\sum_{k \in \mathcal{N}_{\text{tip}}} N_k(\mathbf{x}) \sum_{\alpha=1}^4 F_\alpha(\mathbf{x}) \mathbf{b}_k^\alpha}_{\text{tip}}, \quad (3.55)$$

where N_i are the standard shape functions, $H(\mathbf{x})$ is the Heaviside step function (+1 on one side of the crack, -1 on the other) with enriched degrees of freedom $\mathbf{a}_j$ for nodes whose support is bisected by the crack, and $F_\alpha(\mathbf{x})$ are near-tip asymptotic functions with degrees of freedom $\mathbf{b}_k^\alpha$ for nodes in the crack-tip neighbourhood [48]. The Heaviside term produces the displacement jump across the crack faces; the tip terms reproduce the $r^{1/2}$ singular strain field consistent with Linear Elastic Fracture Mechanics (LEFM) [48]. Extensions of this framework include a cohesive-segments variant in which the singular near-tip enrichment is replaced by a traction-separation law, coupling propagation to a fracture-energy criterion and removing the need for pre-defined asymptotic fields.

Phantom-node method. A topologically equivalent approach is the overlapping-element formulation of Hansbo and Hansbo [27]: a cracked element is replaced by two superimposed sub-elements, each defined on a subset of the original domain, with phantom nodes that activate when the crack passes through the element. Van der Meer and Sluys [46] adapted this formulation to composite laminates, inserting mesh-independent straight matrix splits governed by a mixed-mode cohesive law. The construction is directly comparable to the matrix-split elements of the semi-discrete model, and the kinematic equivalence

to XFEM, different bookkeeping, same displacement field, has been demonstrated [46].

3.2.4. Limitations of discrete methods

Despite their physical fidelity, each of the three principal discrete methods carries limitations that constrain its use in full progressive failure analysis of composite laminates.

VCCT. Beyond the pre-crack requirement noted above, the crack path in VCCT is entirely determined by the mesh topology: advancement relies on node-release procedures that constrain growth to element edges, and the mesh must be updated continuously as the crack front advances [40]. This strong mesh dependency produces oscillatory behaviour in the computed strain energy release rate and numerical instability in the solver [40]. Furthermore, the $1/\sqrt{r}$ singularity of linear elastic fracture mechanics is known to be incorrect for the fibre–matrix interface in composites [40], restricting VCCT’s physical validity to resin-rich interlaminar interfaces and excluding intralaminar matrix cracking.

CZM. The intrinsic CZM formulation — cohesive elements inserted between all element pairs to allow cracking anywhere — introduces either an artificial change in global stiffness or, for high interface stiffnesses, an ill-conditioned stiffness matrix [74]. This confines the intrinsic formulation in practice to delamination, where the crack plane is known a priori and selective insertion is possible. Beyond this, the CZM involves a large number of parameters (interface strength, penalty stiffness, fracture energy, mode-interaction exponent, and artificial viscosity) that are not all independently measurable [67]. The physical basis of the traction–separation law itself has been questioned: experimental evidence indicates that interfacial damage extends beyond the interface plane, and the measured traction–separation relation follows a trapezoidal shape inconsistent with the standard bilinear assumption [34]. In implicit simulations, convergence difficulties are frequently encountered, particularly when damage evolves rapidly across the process zone [74]. In explicit simulations, the main disadvantage is related to the computational cost that comes with the need of high interface stiffnesses. These lead to higher deformation wave speeds and therefore longer computational times, further explored in Section 4.1.3.

XFEM. XFEM removes the need for a prescribed crack path but does not itself prescribe how the crack front advances; it is therefore coupled with CZM to supply the propagation criterion, inheriting its convergence difficulties while adding the numerical complexity of the enrichment [30]. Commercial implementations impose further constraints: the Abaqus implementation excludes the near-tip singular enrichment terms and cannot represent multiple simultaneous parallel cracks within the same enriched region [16], preventing the simultaneous capture of the distributed matrix-crack fields expected in laminates under OHT loading.

More broadly, the cost constraint is shared by all discrete methods regardless of how the discontinuity is discretised: computational effort grows with the number of tracked discontinuities [49] as each crack requires either enriched nodes, a cohesive zone or similar additional treatment; all adding computational time. Full discrete modelling is viable for a coupon with a handful of dominant cracks but becomes prohibitive when a dense, spatially distributed matrix-crack field must be captured across an entire laminate [49, 1].

3.3. Stochastic and Probabilistic Finite Element Approaches

The methods above are deterministic: one set of material properties yields one response. Real composites are inherently non-uniform, fibre packing and orientation vary, constituent strengths scatter, and voids, ply waviness and non-uniform curing introduce defects, and this variability drives the progression of damage, not merely its mean [52, 68]. Stochastic FE approaches treat material properties, and strengths in particular, as random variables or random fields. The statistical bedrock is the Weibull weakest-link distribution [77], which captures the size dependence of brittle and fibre-dominated strength; the review of Stefanou [69] surveys the broader stochastic-FE landscape, and that of Sriramula and Chryssanthopoulos [68] catalogues the sources of uncertainty in fibre-reinforced composites across scales and the magnitudes of property scatter reported for stiffness and strength. Two solution philosophies exist. Sampling-based Monte Carlo simulation, applied to composite strength as early as Oh [57], propagates a random strength field through repeated analyses to recover a distribution of response. The spectral stochastic FE method of Ghanem and Spanos [23] instead represents input random fields by a Karhunen–Loève expansion (a series in uncorrelated random modes) and the response by polynomial chaos. As the input randomness field is expressed using a Karhunen–Loève expansion, it is a function of an uncorrelated randomness variable per considered random mode. Polynomial chaos allows the output to be expressed as a function of the same randomness variables as the input. While Monte Carlo samples random realisations of the input field before a simulation run, it requires many iterations to infer relevant information of the output randomness field, the spectral stochastic FE method solves for coefficients in the polynomial chaos resulting in a probability distribution with only one simulation run [23]. In turn the spectral stochastic FE method adds unknown coefficients that need to be found making the simulation itself computationally more expensive.

For progressive damage the decisive consequence of strength variability is that it controls the gradual accumulation of matrix cracks. A deterministic model with uniform strength tends to crack everywhere at once when the critical stress is reached, whereas in reality cracks initiate at the weakest locations and multiply as load increases, producing the characteristic rise of transverse crack density with strain. Li, Hallett and Wisnom [44] embed a Weibull weakest-link criterion at the element level within an explicit FE progressive-failure scheme, integrating survival probability over the loaded volume to reproduce both the scatter and the size effect of tensile fibre failure. Stochastic modelling is thus not a competing family to CSCM and DCM but an orthogonal enrichment layered onto either.

3.4. The Semi-Discrete Intersection

The limitations of continuum and discrete methods documented in Sections 3.1.4 and 3.2.4 trace to a common origin: both families commit fully to either a smeared or an explicit representation of the crack, and each commitment carries a distinct set of trade-offs, summarised in Table 3.1. The semi-discrete damage model (Semi-Discrete Damage Model (SD2M)) of Nguyen and Waas [49] navigates this trade-off through a specific meshing strategy that uses only standard continuum elements while embedding the directional specificity of discrete crack methods directly into the mesh topology.

Table 3.1: Contrasting characteristics of continuum smeared crack methods (CSCM) and discrete crack methods (DCM).

	CSCM / CDM	DCM
Crack representation	Smeared over element volume via degrading stiffness	Sharp displacement discontinuity (interface, enrichment, embedded or extra nodes)
Element formulation	Single standard continuum element throughout	Specialised: cohesive, enriched (XFEM), or phantom node
Mesh requirement	Arbitrary mesh possible; refinement or fibre alignment needed for crack sharpness	Crack accommodated by interface placement or enrichment
Crack-path fidelity	Follows the mesh unless aligned / crack-spacing controlled	High; directionality set actively or by interface placement
Computational cost	Low; largely independent of number of cracks	High; scales with the number of discontinuities modelled
Robustness	High; standard solvers, no topology change	Lower; convergence and topology handling more demanding

3.4.1. Meshing strategy

The central device of SD2M is a structured decomposition of each ply into two geometrically distinct domains [49]. The bulk of the ply is meshed with continuum elements of arbitrary shape. In the bulk regions only fibre damage can initiate with the matrix left undegraded. Embedded within the bulk mesh are thin, periodically spaced strips of continuum elements, the matrix-split strips capturing all damage modes, oriented along the fibre direction and meshed with exactly one layer of structured hexahedral elements across their width. The strip width is chosen to be much smaller than the bulk element size, so that crack localisation within a strip produces a near-zero-width softened band that approaches the behaviour of a discrete crack while every element in the model remains a standard continuum element [49].

Failure-mode separation. The two-region geometry enforces a strict separation of failure modes. Matrix-splitting failure, transverse cracking oriented along the fibre direction, is active only within the matrix-split strips; fibre failure is admitted in both the strips and the bulk regions. As a result, the crack band for matrix cracking is geometrically constrained to propagate in the physically correct fibre-parallel direction regardless of the surrounding bulk mesh orientation. The crack-path sensitivity to mesh orientation identified in Section 3.1.4 is therefore addressed not by enrichment or dynamic crack-tracking but by embedding the fracture plane in the mesh geometry at the outset. This resolves the locality limitation of Section 3.1.4, whose most direct symptom is precisely this mesh-orientation sensitivity of the crack path. The homogenisation limitation is likewise resolved: because matrix cracking is admitted only in the physically correct fibre-parallel direction and confined to the strips, and because the crack-band regularisation is scaled by each element’s characteristic length using Section 3.1.2 so that the softening law dissipates the energy proper to the fracture mode represented. Both fundamental limitations of Section 3.1.4 are thereby addressed. This separation also creates a realistic damage pattern, in which matrix cracks have a discrete appearance.

Element sizing and the Bažant limit. For a standard CDM element the crack-band model requires the element size to satisfy the Bažant–Oh limit in all directions associated with active failure planes, as seen

by Equation 3.8 [5]. In SD2M, because matrix failure is excluded from the bulk regions, only the limit in the fibre direction, governing fibre failure, applies to the bulk elements; their transverse dimension is unconstrained [49]. Bulk elements can therefore be considerably coarser than a uniform crack-band mesh would permit. The matrix-split strips are geometrically constrained in their width as they are one element wide for which each of those elements is restricted in their size by the Bažant–Oh limit.

Interlaminar completion. As anticipated above, cohesive elements are adopted here in preference to surface-based cohesive contact. The interface mesh is generated together with the intralaminar mesh, so each cohesive element face can be partitioned to match the discretization of the ply it adjoins. With conforming nodes, tractions are transferred directly through shared degrees of freedom rather than through the node-to-surface projection used by surface-based contact, avoiding the smoothing of the near-front stress field that would otherwise degrade the predicted delamination onset [49].

Computational cost and the discrete crack field. Within the established methods, the semi-discrete model sits at the intersection of CDM and CZM. It approximates the fidelity of an explicit crack while avoiding the computational cost that comes with resolving every matrix crack as a true discontinuity. This cost is avoided because the matrix-split strips do not require the artificially high penalty stiffness that a cohesive interface needs to remain closed before onset; being ordinary continuum elements, they simply carry the same stiffness as the surrounding ply. The strips therefore introduce none of the elevated deformation wave speeds, and hence none of the reduced stable time increment, that the stiff cohesive interfaces of Section 3.2.4 impose in explicit analysis. By the same token, in an implicit standard analysis the strips add no ill-conditioning to the global stiffness matrix, so they avoid the convergence problems that the high penalty stiffness of a cohesive interface provokes. This distinction matters because a realistic transverse failure pattern consists of many densely spaced discrete matrix splits: capturing it with CZM is either prohibitively expensive or convergence-limited, while a CDM smears the damage and misses the discrete character of the pattern altogether. The semi-discrete framework resolves this by producing discrete-like cracks while still admitting a large number of closely spaced splits, precisely because of the computational savings above and, in the implicit case, the non-singular stiffness matrix.

Limitations. The method is not without cost or limitation, and these must be stated as plainly as those of the competing families. Matrix cracking is confined to strips placed a priori along the fibre direction, so genuinely off-axis, curved matrix cracks cannot be represented without re-partitioning, a restriction the enriched methods of Section 3.2.3 are specifically designed to avoid. This is the same restriction faced by CZM in Section 3.2.4: matrix cracks, like cohesive cracks, can only form where the discretisation admits them in advance rather than at an arbitrary emergent location. The strips still demand a structured, size-limited mesh aligned with each ply, adding pre-processing effort relative to a free-mesh; partitioning every ply along its fibre direction and seeding the strips is laborious. Finally, the published validation base is essentially that of a single research group on a small number of toughened and non-toughened carbon/epoxy systems [50, 53, 51], so further validation remains outstanding, a gap this thesis is, in part, intended to help close.

4

Finite Element Considerations

4.1. Solver Types

Finite element analyses require a time-integration strategy that advances the equations of motion from one state to the next. Explicit and implicit methods, together with methods for the solution of equilibrium problems, differ in their stability properties, their computational cost, and their behaviour in the presence of material softening. As shown below, both the implicit and the explicit formulation descend from a single parametrized time integrator; the choice of one parameter decides which of the two is obtained.

4.1.1. The Newmark Family

The semi-discrete equation of motion at step $n + 1$ is

$$\mathbf{M}_{\text{mass}} \ddot{\mathbf{u}}_{n+1} + \mathbf{f}_{n+1}^{\text{int}} = \mathbf{f}_{n+1}^{\text{ext}}, \quad (4.1)$$

where $\mathbf{f}_{n+1}^{\text{int}} = \mathbf{f}^{\text{int}}(\mathbf{u}_{n+1}, \dot{\mathbf{u}}_{n+1}, t_{n+1})$ is the internal forces at step $n + 1$, $\mathbf{f}_{n+1}^{\text{ext}} = \mathbf{f}^{\text{ext}}(t_{n+1})$ the external forces at step $n + 1$, and $\mathbf{M}_{\text{mass}}$ the mass matrix. This leads to a system that relates the new-step displacement, velocity and acceleration. $\mathbf{f}_{n+1}^{\text{int}}$ is not only a function of a displacement and velocity but also a function of time, as materials can change their stiffness and damping properties over time due to effects like plasticity or degradation of the stiffness as consequence of damage in the material. To advance the given system, the relations between $\mathbf{u}_{n+1}$, $\dot{\mathbf{u}}_{n+1}$ and $\ddot{\mathbf{u}}_{n+1}$ are parametrized by a truncated Taylor expansion with two free weights β and γ , as introduced by Newmark:

$$\mathbf{u}_{n+1} = \mathbf{u}_n + \Delta t \dot{\mathbf{u}}_n + \Delta t^2 \left[\left(\frac{1}{2} - \beta \right) \ddot{\mathbf{u}}_n + \beta \ddot{\mathbf{u}}_{n+1} \right], \quad (4.2)$$

$$\dot{\mathbf{u}}_{n+1} = \dot{\mathbf{u}}_n + \Delta t \left[(1 - \gamma) \ddot{\mathbf{u}}_n + \gamma \ddot{\mathbf{u}}_{n+1} \right]. \quad (4.3)$$

Here γ controls the artificial damping introduced by the integrator—the choice $\gamma = \frac{1}{2}$ adds none and makes the method second-order accurate, raising the global error from $O(\Delta t)$ to $O(\Delta t^2)$, while β weights the contribution of the *new-step* acceleration $\ddot{\mathbf{u}}_{n+1}$ to the displacement update.

Both explicit and implicit methods are related to the Newmark family. While an explicit scheme obtains the new displacement directly from known quantities from previous steps, an implicit scheme describes the new displacement as a function of unknown quantities at the new time step. The nature of the scheme is therefore decided through the weight β through the displacement update (4.2), as it couples the solution to the unknown acceleration at the new time step. While $\beta \neq 0$ results in an implicit displacement update, $\beta = 0$ results in an explicit update. Even though γ weighs the contribution of the new-step acceleration to the velocity update (4.3), it can be modified to be used in both explicit and implicit methods as further discussed in Section 4.1.3.

4.1.2. Implicit Time Integration

When $\beta \neq 0$, the displacement update (4.2) is split into a part built entirely from known step- n quantities and a part carrying the unknown acceleration,

$$\mathbf{u}_{n+1} = \underbrace{\mathbf{u}_n + \Delta t \dot{\mathbf{u}}_n + \Delta t^2 \left(\frac{1}{2} - \beta\right) \ddot{\mathbf{u}}_n}_{\tilde{\mathbf{u}}_{n+1} \text{ (known)}} + \beta \Delta t^2 \ddot{\mathbf{u}}_{n+1}, \quad (4.4)$$

which can be solved for the new-step acceleration,

$$\ddot{\mathbf{u}}_{n+1} = \frac{1}{\beta \Delta t^2} (\mathbf{u}_{n+1} - \tilde{\mathbf{u}}_{n+1}). \quad (4.5)$$

Substituting (4.5) into the equation of motion (4.1) eliminates the acceleration and leaves a system in the displacement alone,

$$\frac{1}{\beta \Delta t^2} \mathbf{M}_{\text{mass}} (\mathbf{u}_{n+1} - \tilde{\mathbf{u}}_{n+1}) + \mathbf{f}_{n+1}^{\text{int}} = \mathbf{f}_{n+1}^{\text{ext}}. \quad (4.6)$$

The unknown $\mathbf{u}_{n+1}$ now appears both in the inertia term and inside the nonlinear internal force $\mathbf{f}_{n+1}^{\text{int}}$, so the new-step displacement cannot be written down directly, as each step requires the solution of a coupled generally nonlinear system. This is the defining feature of an implicit method: equilibrium is enforced at the end of the increment, using quantities at the new time step. Writing the residual $\mathbf{r}_{n+1}$ of step $n+1$ as the difference between the left- and right-hand sides of (4.6) yields

$$\mathbf{0} = \mathbf{r}_{n+1} = s_D \frac{1}{\beta \Delta t^2} \mathbf{M}_{\text{mass}} (\mathbf{u}_{n+1} - \tilde{\mathbf{u}}_{n+1}) + \mathbf{f}_{n+1}^{\text{int}} - \mathbf{f}_{n+1}^{\text{ext}}, \quad (4.7)$$

with $s_D = 1$ for a dynamic problem and $s_D = 0$ for a static equilibrium problem, the latter dropping the inertia term entirely, these are nonlinear algebraic equations in the nodal displacements $\mathbf{u}_{n+1}$.

The most widely used and most robust method for their solution is Newton's method (the Newton-Raphson method). A first order Taylor expansion of the residual about the displacement of the current iteration v , gives the linear model

$$\mathbf{A}_{n+1}^v \Delta \mathbf{u}_{n+1}^v = -\mathbf{r}_{n+1}^v, \quad (4.8)$$

where $\mathbf{A}_{n+1}^v = \left. \frac{\partial \mathbf{r}}{\partial \mathbf{u}} \right|_{t=t_{n+1}, i=v}$ is the system Jacobian, called in computational mechanics the effective tangent stiffness matrix. The increment is added to the displacement at the previous iteration, $\mathbf{u}_{n+1}^{v+1} = \mathbf{u}_{n+1}^v + \Delta \mathbf{u}_{n+1}^v$, and the process is repeated until a convergence criterion is met. The Jacobian collects the tangent stiffness matrix of the internal nodal forces $\mathbf{K}^{\text{int}} = \partial \mathbf{f}^{\text{int}} / \partial \mathbf{u}$, the load stiffness matrix of the external nodal forces

$\mathbf{K}^{\text{ext}} = \partial \mathbf{f}^{\text{ext}} / \partial \mathbf{u}$, and for dynamic problems the mass matrix:

$$\mathbf{A}_{n+1}^v = \frac{s_D}{\beta \Delta t^2} \mathbf{M}_{\text{mass}} + \mathbf{K}_{n+1}^{\text{int}^v} - \mathbf{K}_{n+1}^{\text{ext}}. \quad (4.9)$$

The full Newton algorithm evaluates and inverts the Jacobian at every iteration. As this is computationally expensive many programs instead use a modified Newton algorithm, in which the Jacobian is assembled and triangulated only at the beginning of a step or at irregular intervals during it, resulting in a faster but less robust scheme.

One advantage of this implicit method is that, it is unconditionally stable for the static case $s_D = 0$. Considering a dynamic ($s_D \neq 0$) linear transient problem, stability is bound by Newmark which is unconditionally stable for $\beta \geq \frac{\gamma}{2} \geq \frac{1}{4}$. The major restrictions on the implicit time step arise instead from accuracy requirements and from the decreasing robustness of the Newton procedure as the step grows: with a large step the starting iterate may be far from the solution for problems exhibiting strong non-linearities, increasing the possibility that Newton's method fails to converge. In return for their enhanced stability, implicit methods require the solution of nonlinear algebraic equations at each step, the construction of the linear models is often involved, and the storage of the equations requires significant memory.

4.1.3. Explicit Time Integration

The explicit method is a special case of the Newmark equations, obtained by setting $\beta = 0$ [6]. With $\beta = 0$ the new-step acceleration drops out of the displacement update (4.2), which becomes

$$\mathbf{u}_{n+1} = \mathbf{u}_n + \Delta t \dot{\mathbf{u}}_n + \frac{1}{2} \Delta t^2 \ddot{\mathbf{u}}_n. \quad (4.10)$$

The entire right-hand side is now made up of step- n quantities, so $\mathbf{u}_{n+1}$ is obtained directly: no coupled system and no end-of-step equilibrium iteration.

One of the most popular implementations is the central difference scheme with a variable time step [6] using half-steps. The equivalence of this method to the Newmark equations with $\beta = 0$ is shown in Appendix A. The variable time increments are defined using [6]

$$\Delta t_{n+\frac{1}{2}} = t_{n+1} - t_n, \quad t_{n+\frac{1}{2}} = \frac{1}{2}(t_{n+1} + t_n), \quad \Delta t_n = t_{n+\frac{1}{2}} - t_{n-\frac{1}{2}}. \quad (4.11)$$

Using the central difference for the velocity of the next half step and rearranging to the displacement of the next step yields:

$$\mathbf{u}_{n+1} = \mathbf{u}_n + (t_{n+1} - t_n) \dot{\mathbf{u}}_{n+\frac{1}{2}} = \mathbf{u}_n + \Delta t_{n+\frac{1}{2}} \dot{\mathbf{u}}_{n+\frac{1}{2}}. \quad (4.12)$$

Rearranging the central difference of the acceleration of the current step to the velocity of the next half-step yields

$$\dot{\mathbf{u}}_{n+\frac{1}{2}} = \dot{\mathbf{u}}_{n-\frac{1}{2}} + (t_{n+\frac{1}{2}} - t_{n-\frac{1}{2}}) \ddot{\mathbf{u}}_n = \dot{\mathbf{u}}_{n-\frac{1}{2}} + \Delta t_n \ddot{\mathbf{u}}_n. \quad (4.13)$$

The acceleration at the current step $\ddot{\mathbf{u}}_n$ is obtained from the Equation of Motion (EOM) (4.1) for the current step. Rearranging then yields:

$$\ddot{\mathbf{u}}_n = \mathbf{M}_{\text{mass}}^{-1} (\mathbf{f}_n^{\text{ext}} - \mathbf{f}_n^{\text{int}}). \quad (4.14)$$

From equations (4.14), (4.13), and (4.12) the new displacement can be computed. From these new

quantities a new internal load can be obtained using a constitutive law allowing the algorithm to repeat.

The full step follows this algorithm in a loop until the final step is reached:

1. using the last nodal half-step velocities and the last internal force vector the new half-step velocity is determined using equation (4.13);
2. with equation (4.12) the displacement of the new step is determined;
3. at each element integration point the strain increment is recovered from the nodal displacements through the strain–displacement operator, $\Delta \epsilon = \mathbf{B}_{\epsilon u} \Delta \mathbf{u}^{n+1}$;
4. the constitutive update maps the strain to the integration-point stress σ^{n+1} , advancing the damage state variables;
5. the stresses are integrated over each element to form its internal force vector,

$$\mathbf{f}_e^{\text{int}} = \int_{\Omega_e} \mathbf{B}_{\epsilon u}^T \sigma \, dV, \quad (4.15)$$

evaluated by numerical quadrature at the element integration points;

6. the element contributions are assembled into the global internal force vector $\mathbf{f}^{\text{int}, n+1}$

Conditional stability and the stable time step. The price paid for this simplicity is that explicit methods are only conditionally stable: if the time step exceeds a critical value Δt_{crit} the solution grows unboundedly.

The upper bound for the critical timestep is based on the Courant–Friedrichs–Lewy (CFL) condition by Courant, Friedrichs and Lewy (1928) [13]. The CFL condition states that a necessary condition for convergence of a finite difference scheme is that the physical domain of dependence of the underlying differential equation must be contained within the numerical domain of dependence of the scheme itself [13]. In other words if the physical information front travels faster than the mesh spacing allows information to propagate numerically, the numerical solution cannot converge to the true solution of the differential equation. This translates directly to structural explicit FE schemes and imposes a bound on the timestep [79]:

$$\Delta t_{\text{crit}} \leq \min_e \frac{l_e}{c_e}, \quad (4.16)$$

where c_e is the wave speed of the stress wave propagating through an element e and l_e is a characteristic length of that element [15, 79].

In structural dynamics problems the motion of the system further bounds the critical timestep. To find this bound the semi-discrete EOM (4.1) is further analyzed. Assuming viscous damping, expanding the internal force $\mathbf{f}^{\text{int}}$ in the EOM (4.1) yields

$$\mathbf{M}_{\text{mass}} \ddot{\mathbf{u}} + \mathbf{D}_{\text{damp}} \dot{\mathbf{u}} + \mathbf{K} \mathbf{u} = \mathbf{f}^{\text{ext}}. \quad (4.17)$$

Using modal decomposition, Appendix C derives this limit as

$$\Delta t_{\text{crit}} \leq \frac{2}{\omega} \left(\sqrt{\xi^2 + 1} - \xi \right), \quad (4.18)$$

where ω is the highest natural frequency of the structure and ξ the critical damping ratio. For any given natural frequency of a system the resulting function $f(\xi)$ of the upper bound established in Eq. (4.18)

is steadily decreasing, converging to value of zero. That means that higher damping results in smaller timesteps. The undamped case is recovered by setting $\xi = 0$:

$$\Delta t_{\text{crit}} \leq \frac{2}{\omega}. \quad (4.19)$$

This requires the estimation of the largest eigenvalue of the global problem, which is computationally expensive to evaluate at every increment in a large simulation [79]. A practical bound is obtained by working at the element level. Fried showed that the maximum eigenvalue of the assembled global system cannot exceed the largest eigenvalue of any individual element considered in isolation [22], giving the element eigenvalue inequality

$$|\Lambda^{\text{max}}| \leq \max_e \Lambda_{\text{max}}^e. \quad (4.20)$$

The element eigenvalue problems $\mathbf{K}_e \mathbf{u}^e = \Lambda \mathbf{M}_e \mathbf{u}^e$ are small, with their sizes being the number of element Degree of Freedoms (DOFs), and can be solved analytically or allows finding conservative bounds [20]. A common approximation of an elements maximum eigenfrequency, used for example by the Abaqus/Explicit solver [16, 15], is [20]:

$$\omega_{\text{max}} = 2 \frac{l_e}{c_e}. \quad (4.21)$$

This is the maximum eigenfrequency of a one-dimensional bar element with two nodes. When substituted into equation (4.19) the CFL condition, specifically equation (4.16), is recovered. Flanagan and Belytschko [20] proved that this not conservative and underestimates the true maximum eigenfrequency of some elements, leading to the solution to become unstable. Abaqus counteracts this instability by choosing l_e as the shortest distance in the mesh instead of finding a critical length in each element and by imposing a safety factor on the estimated timestep [16]. For a two-dimensional and three-dimensional simulation Abaqus uses a safety factor of $\frac{1}{\sqrt{2}}$ and $\frac{1}{\sqrt{3}}$ respectively [16]. This recovers the conservative recommended maximum frequency for square elements in the two-dimeantional case and for cube elements in the three-dimensional case as seen in Table 2 and Table 3 of Flanagan and Belytschko(1984) [20].

Wave Speed Determination. To use the estimations of the critical time steps above the wavespeed has to be quantified. In an Infinite medium two types of uncoupled waves exist: the dilatational (change in volume) and the distortional (rotational) wave [24]. As an infinite medium does not experience interference with it self though reflections, the wavespeeds derived for such a medium represent the fastest speed of the wave. For an isotropic material each waves has one wave speed [24]

$$c_1 = \sqrt{\frac{\lambda + 2\mu}{\rho}} \quad \text{and} \quad c_2 = \sqrt{\frac{\mu}{\rho}}, \quad (4.22)$$

where c_1 is the speed of the dilatational wave and c_2 the one of the distortional wave. λ and μ are the Lamé constants with μ being the shear modulus G [24] and λ describing the contribution of volumetric changes to the normal stress components. As both μ and λ are positive, equations (4.22) show that the dilatational wave speed c_1 is always higher than c_2 and is therefore the critical speed to be used for the estimation of the critical time step. The bulk modulus K_b relates the mean stress σ_m to the volumetric strain ε_V :

$$K_b = \frac{\Delta \sigma_m}{\Delta \varepsilon_V}. \quad (4.23)$$

It is related to the Lamé constants in the following way:

$$K_b = \lambda + \frac{2}{3}\mu. \quad (4.24)$$

With this, the dilatational wave speed c_1 from Equation 4.22 can be rewritten in terms of the bulk and shear moduli:

$$c_1 = \sqrt{\frac{K_b + \frac{4}{3}\mu}{\rho}} = \sqrt{\frac{K_b + \frac{4}{3}G}{\rho}}. \quad (4.25)$$

For orthotropic and anisotropic materials the dilatational wave speed depends in the material direction. As a simplification Abaqus uses the isotropic expressions. For this it calculates an effective bulk modulus each step by sampling the stress response given the change in volumetric strain that caused it [16]. An effective μ is calculated by sampling the change in deviatoric stress ΔS given the deviatoric strain Δe that caused it, using the relation [16]

$$\Delta S = 2\mu\Delta e. \quad (4.26)$$

From these two effective values the current wave speed can be calculated using equation (4.25). From this wave speed the critical timestep for the next step is estimated as described earlier. In the case of strong non-linearities or highly anisotropic materials this assumption will not be sufficient and the time step has to be adjusted by manually reducing the safety factor.

Eq. (4.25) reveals that higher material densities will yield a lower wave-speed and therefore increase the critical timestep. As the bulk modulus is a measure of the stiffness of the material, Eq. (4.25) also shows that increasing the material stiffness or the shear modulus along a direction that the wave propagates, will lead to an increase in the wave speed in that direction, and effectively lower the timestep whenever the dilatational wave travels along that said direction.

Energy balance. The stability conditions above stem from a linear analysis of the integrator; there is at present no stability theory that covers the full range of nonlinear phenomena encountered in practice. It is therefore possible for instabilities to develop even when the time-step condition is observed, and in nonlinear problems such instabilities are not always readily discernible from the results. Any instability produces a spurious generation of energy that violates energy conservation, so whether stability was maintained can be checked by an energy balance. The internal and external energies are integrated in time (for example by the trapezoidal rule), the kinetic energy is $W_{\text{kin}}^n = \frac{1}{2}(\mathbf{v}^n)^T \mathbf{M}_{\text{mass}} \mathbf{v}^n$, which is why velocities at integer time steps are computed, and energy conservation requires

$$|W_{\text{kin}} + W_{\text{int}} - W_{\text{ext}}| \leq \varepsilon \max(W_{\text{ext}}, W_{\text{int}}, W_{\text{kin}}), \quad (4.27)$$

with ε a small tolerance, typically of order 10^{-2} .

4.1.4. Convergence in the Presence of Material Softening

The introduction of material softening significantly alters the performance and stability of finite element solvers. Softening behavior, characterized by a decrease in stress capacity with increasing strain, manifests mathematically as a negative tangent modulus at the material level [6]. This fundamental change dictates the practical suitability of the time-integration schemes discussed previously.

For implicit time integration, material softening poses severe numerical challenges. The Newton–Raphson method relies on the continuous assembly and inversion of the system Jacobian, $\mathbf{A}_{n+1}^v$. When local material points enter the softening regime, their contribution to the global tangent stiffness matrix, $\mathbf{K}^{\text{int}}$, becomes negative. If a sufficient number of elements soften, the global Jacobian loses its positive-definite property and may become singular or ill-conditioned [6]. Consequently, the iterative solver struggles to find global equilibrium, losing its characteristic quadratic convergence rate. The residual $\mathbf{r}_{n+1}$ fails to decrease, iterates oscillate, and the solution often aborts entirely due to non-convergence [6]. Furthermore, the residual is not continuously differentiable as the softening curve has a triangular shape. For these problems the second order convergence does not apply anymore [6], in the worst case leading to non-convergence.

Conversely, explicit time integration offers inherent robustness when dealing with softening materials [6]. Because the explicit scheme advances the kinematic state without enforcing end-of-step equilibrium, it never requires the assembly or inversion of a global stiffness matrix. The internal forces, $\mathbf{f}_n^{\text{int}}$, are computed directly from the current stress state and assembled purely as a right-hand side vector to evaluate the nodal accelerations (Eq. 4.14). As a result, negative local tangent stiffnesses do not disrupt the algorithmic flow. The solver naturally captures the transient dynamics of the failure process, including the rapid redistribution of stresses and the sudden release of strain energy into kinetic energy, without solver divergence [6]. This structural advantage makes explicit methods the predominant choice for simulating progressive failure, provided the conditional stability criteria and energy balances are strictly monitored.

4.2. Quasi-Static Modelling

To capture true material responses, a load is applied to the specimen, causing accelerations and therefore inducing oscillations due to inertial effects. The data of interest, however, is a material response that is independent of loading speed and representative of the static response of the specimen. To obtain a true material response the loading process has to be quasi-static. A quasi-static process is a dynamic process that happens infinitely slow, allowing the system to reach static equilibrium at all times.

For implicit solvers modeling a quasi-static process is trivial because static equilibrium can be enforced at each step as described in Sec. 4.1.2. As explicit solvers are inherently dynamic, due to the fact that they rely on nodal accelerations (Sec. 4.1.3), a quasi-static process can only be approximated. Different methods exist to achieve a valid approximation.

4.2.1. Mass

Another way to think about a quasi-static process is using the resulting stress wave of a loading event. As the stress wave propagates through the system it reflects at the specimens edges and interferes with itself, eventually reaching a steady state stress field, at which static-equilibrium is reached. The faster the wave propagates, the faster this state is reached. Eq. (4.22) shows that a material with given stiffness properties will exhibit higher wave speeds for a lower material density. Decreasing the density will therefore bring the solution closer to that of the equivalent quasi-static process. As discussed in Sec. 4.1.3 higher wave speeds will lead to a smaller critical timestep and therefore to longer simulation runtimes.

This effect is also observed from the the inertial term in the EOM (4.1). Decreasing the mass of the system will reduce the inertial forces and therefore the oscillations induced by the loading event.

4.2.2. Damping

Another way to obtain results closer to that of a quasi-static process is to introduce damping into the system. Damping is a physical phenomenon that dissipates kinetic energy from the system, for example through internal friction or viscous drag. This leads to a reduction in the amplitude of oscillations and therefore to a faster convergence to a static-equilibrium state.

Damping can be quantified by the damping ratio ξ . A damping ratio of $\xi = 0$ corresponds to an undamped system, while a damping ratio of $\xi = 1$ corresponds to a critically damped system. Critically damped systems return to static equilibrium in the shortest time possible without oscillating. Underdamped systems ($0 < \xi < 1$) oscillate with a decaying amplitude, while overdamped systems ($\xi > 1$) return to static equilibrium without oscillating but more slowly than critically damped systems. [12] Fig. 4.1 visualizes the effect of the damping ratio.




Figure 4.1: Effect of the damping ratio (ξ) on the free vibration displacement response.

In newtonian fluids the damping force is proportional to the velocity of the system, which is known as viscous damping. Viscous damping however can also be used to model damping in solids [75]. While it is only an approximation of the physical damping mechanisms in solids, it can be used to reduce oscillations as the actual dynamic response of the system is not of interest. Viscous damping can be represented using the damping matrix $\mathbf{D}_{\text{damp}}$ adding a damping term to the EOM (4.1):

$$\mathbf{M}_{\text{mass}}\ddot{\mathbf{u}} + \mathbf{D}_{\text{damp}}\dot{\mathbf{u}} + \mathbf{K}\mathbf{u} = \mathbf{f}^{\text{ext}}. \quad (4.28)$$

One way to define the damping matrix $\mathbf{D}_{\text{damp}}$ is using a Rayleigh damping model, which is a linear combination of the mass matrix $\mathbf{M}_{\text{mass}}$ and the stiffness matrix $\mathbf{K}$ [63]:

$$\mathbf{D}_{\text{damp}} = \alpha_R \mathbf{M}_{\text{mass}} + \beta_R \mathbf{K}, \quad (4.29)$$

where α_R and β_R are user-defined coefficients that control the damping behavior of the system. The same damping matrix $\mathbf{D}_{\text{damp}}$ will have different damping ratios ξ depending on the frequency of the motion.

The relationship between the damping ratio ξ and the Rayleigh coefficients α_R and β_R is given by [12]:

$$\xi = \frac{\alpha_R}{2\omega} + \frac{\beta_R \omega}{2}, \quad (4.30)$$

where ω is the frequency of the motion. Mass weighting α_R is therefore more effective at low frequencies as it's contribution to the damping ratio is inversely proportional to the frequency, while stiffness weighting β_R is more effective at high frequencies with its contribution being directly proportional to the frequency. This is seen in Fig. 4.2. Numerical effects or shock events such as fractures live at high frequencies, that will stay undamped if no stiffness weighting or no damping is used.




Figure 4.2: Rayleigh damping ratio versus frequency for $\alpha_R=5.0$ and $\beta_R=0.05$, highlighting the mass-proportional and stiffness-proportional components.

Thomson (1865) [72] and Voigt (1892) [75] showed that the the total stress in a viscoelastic solid can be expressed as the sum of an elastic and a viscous stress component:

$$\sigma = \sigma_{\text{elastic}} + \sigma_{\text{viscous}}. \quad (4.31)$$

The viscous stress component σ_{viscous} is the stress that integrated over the volume of the system results in the damping force $\mathbf{D}_{\text{damp}}\dot{\mathbf{u}}$ in the EOM (4.1). In FEM this can be expressed as

$$\mathbf{f}^{\text{damping}} = \int_{\Omega} \mathbf{B}_{\epsilon u}^T \sigma_{\text{viscous}} dV = \mathbf{D}_{\text{damp}}\dot{\mathbf{u}}, \quad (4.32)$$

where $\mathbf{B}_{\epsilon u}$ is the strain-displacement operator and Ω is the volume of the system. Splitting the damping force using rayleigh damping as in Eq. (4.29) into a mass and stiffness proportional component yields $\mathbf{f}^{\text{damping}} = \alpha_R \mathbf{M}_{\text{mass}}\dot{\mathbf{u}} + \beta_R \mathbf{K}\dot{\mathbf{u}}$. As the global stiffness matrix $\mathbf{K}$ is the volume integral of the material stiffness matrix $\mathbf{C}$ multiplied by the strain-displacement operator $\mathbf{B}_{\epsilon u}$,

$$\mathbf{K} = \int_{\Omega} \mathbf{B}_{\epsilon u}^T \mathbf{C} \mathbf{B}_{\epsilon u} dV, \quad (4.33)$$

and the strain rate $\dot{\epsilon}$ is related to the nodal velocities $\dot{\mathbf{u}}$ through the strain-displacement operator $\mathbf{B}_{\epsilon u}$ as

$$\dot{\epsilon} = \mathbf{B}_{\epsilon u} \dot{\mathbf{u}}, \quad (4.34)$$

the stiffness proportional component of the damping force can be expressed as

$$\beta_R \mathbf{K} \dot{\mathbf{u}} = \beta_R \int_{\Omega} \mathbf{B}_{\epsilon u}^T \mathbf{C} \mathbf{B}_{\epsilon u} \dot{\mathbf{u}} dV = \int_{\Omega} \mathbf{B}_{\epsilon u}^T (\beta_R \mathbf{C} \dot{\epsilon}) dV. \quad (4.35)$$

Comparing Eq. (4.32) and Eq. (4.35) shows that the stiffness proportional component of the damping force can be expressed as a viscous stress component σ_{viscous} as

$$\sigma_{\text{viscous}} = \beta_R \mathbf{C} \dot{\epsilon}. \quad (4.36)$$

With

$$\mathbf{M}_{\text{mass}} = \int_{\Omega} \rho \mathbf{N}_s^T \mathbf{N}_s dV, \quad (4.37)$$

where ρ is the material density and $\mathbf{N}_s$ is the shape function matrix, the mass proportional damping can be written as

$$\alpha_R \mathbf{M}_{\text{mass}} \dot{\mathbf{u}} = \alpha_R \int_{\Omega} \rho \mathbf{N}_s^T \mathbf{N}_s \dot{\mathbf{u}} dV = \int_{\Omega} \mathbf{N}_s^T (\alpha_R \rho \dot{u}(x)) dV. \quad (4.38)$$

As this integral is not in the form of Eq. (4.32) mass proportional damping can not be expressed as a viscous stress component σ_{viscous} . Therefore, using rayleigh damping on a constitutive level is only possible when using stiffness proportional damping ($\alpha_R = 0$).

Two caveats exist when using viscous damping to approximate a quasi-static process. As seen in Eq. (4.18) higher damping ratios will lead to smaller critical timesteps and therefore longer simulation runtimes. Additionally, a damping coefficient that is too high will lead to a higher total stress in the system as seen in Eq. (4.31), which will lead to an overestimation of the effective stiffness of the system when analyzing the resulting stress-strain response. When using stress based failure criterions, an overestimated stiffness can also lead to failure predictions that are not representative of the actual material behavior. Eq. (4.31) also shows that high loading speeds will increase the viscous stress component, making the effect of the elevated effective stiffness more pronounced.

4.2.3. Conclusion

The combination of mass reduction and viscous damping can be used to approximate a quasi-static process in an explicit solver. In both cases, the critical timestep will be reduced, leading to longer simulation runtimes. Mass or density reduction will lead to a lower dynamic response, while viscous damping leads to the dissipation of already existing oscillations. A trade-off has to be found to ensure correct results while ensuring that the simulation runtime is not excessive.

4.3. Element Types

The element type determines which stress and strain components are available as primary solution quantities at integration points, and therefore what constitutive models and failure criteria can be evaluated during the analysis. Three element families are relevant to mesoscale composite modelling in Abaqus: shell elements, continuum shell elements, and solid continuum elements [16, 4].

Shell elements. General-purpose shell elements (S4, S4R, S8R) are built on CLT or First-Order Shear Deformation Theory (FSDT) (First-Order Shear Deformation Theory) kinematics and the laminate stacking sequence is encoded as a material definition rather than as geometric thickness. Assembling multiple shell layers in the thickness direction is not recommended for this element class [76, 16], which precludes a layer-wise representation in which each ply occupies its own element with a distinct material definition and orientation. Shell elements are therefore not considered further.

Continuum shell elements. Continuum shell elements (SC8R) combine the three-dimensional nodal topology of solid elements with the reduced kinematic assumptions of a shell. Element thickness is derived from the nodal geometry rather than specified as a material property [76]. These elements resolve in-plane, bending, and shear-dominated deformations accurately, and can provide the through-thickness stresses σ_{33} , τ_{13} , τ_{23} ; however, these transverse components are recovered via equilibrium equations in post-processing rather than being computed from the constitutive law at integration points [16]. For the present model, in which cohesive zone tractions and three-dimensional intralaminar failure criteria must be evaluated at integration points during the explicit time-integration loop, transverse stresses that are not part of the primary constitutive solution are not sufficient.

Solid continuum elements. Solid continuum elements (C3D8, C3D8R, C3D20R) impose no kinematic assumptions on the through-thickness displacement field. The complete three-dimensional stress state, including σ_{33} , τ_{13} , τ_{23} , emerges directly from the constitutive update and is available at every integration point, satisfying the requirements of three-dimensional failure criteria and cohesive zone driving forces without additional recovery steps [76]. Solid elements perform best when element aspect ratios are near unity, a condition not met for individual composite plies; however, with one element per ply in the thickness direction the accuracy degradation is modest and is outweighed by the advantage of a fully resolved stress state [76].

Among the available solid formulations, the eight-node linear brick with reduced integration and hourglass control (C3D8R) provides the best ratio of accuracy to computational cost [76]. Reduced integration evaluates element stiffness at a single central point rather than at eight integration points [16]. Full integration (C3D8) is strictly disadvantageous here: it introduces shear locking under bending-dominated states [16], so it is both more expensive than C3D8R and less accurate [76]. Quadratic elements (C3D20R) recover accuracy but at almost an eightfold increase in computational cost, which scales unfavourably across the many plies of a composite laminate [76]. The single integration point of C3D8R, however, cannot detect certain spurious deformation patterns, known as hourglass or zero-energy modes, where the element distorts without registering any strain energy [16]. In simple terms, the element can “fold” in an unphysical way that the solver is blind to, as seen in Figure 4.3. Hourglass control is implemented in two ways, using relaxed stiffness where Abaqus introduces a small artificial stiffness to counteract, or enhanced hourglass controls where the solver is adapted so that hourglass could never happen in the first place [16]. As it is the default, relaxed stiffness is used.




Figure 4.3: Schematic of hourglass deformation modes in a reduced-integration hexahedral element.

The C3D8R element, with one element per ply in thickness, is therefore adopted for all ply elements in the present model.

4.4. Boundary Conditions Effects

The displacement boundary conditions imposed at the loaded ends of a specimen do more than introduce the load: the kinematic constraints applied there disturb the local stress field in a thin region adjacent to the constrained faces. This is inconsequential in a purely elastic analysis, where only the far-field response is of interest, but it is not inconsequential for a progressive-damage model. Such a model evaluates damage initiation and softening from the local stress state at every integration point, so any stress that the boundary condition introduces artificially is read by the model as a physical driving force for failure.

Consider a coupon loaded in tension along its axis, denoted direction 1, with the transverse directions 2 (in-plane, across the width) and 3 (through-thickness) perpendicular to it. The load can be reacted at the fixed end in two limiting ways: by constraining only the axial displacement, $u_1 = 0$, leaving the transverse directions free; or by additionally constraining the transverse displacements, $u_1 = u_2 = u_3 = 0$, a fully clamped end. The two differ only in whether the cross-section is free to contract at the loaded face. It is important to note that the loaded end and the fixed end have to be consistent in there constraints.

Restraint of transverse contraction. Under uniaxial tension the stress state in the gauge region is $\sigma_{11} > 0$ with all other components zero. By the constitutive relation this axial stress is accompanied by a transverse contraction,

$$\varepsilon_{22} = -\nu_{12} \varepsilon_{11}, \quad \varepsilon_{33} = -\nu_{13} \varepsilon_{11}, \quad (4.39)$$

where ν_{12} and ν_{13} are the major Poisson's ratios. Constraining only u_1 at the reacted face leaves this contraction free to develop, so the boundary condition is consistent with the uniaxial state that the test is intended to reproduce. Constraining u_2 and u_3 as well forces $\varepsilon_{22} = \varepsilon_{33} = 0$ at the face while the adjacent interior continues to contract according to (4.39). The resulting incompatibility is equilibrated by transverse reaction stresses (σ_{22} , σ_{33}) and, where the contraction is non-uniform across the section, by transverse shear (τ_{12} , τ_{13}). A clamped end thus replaces the intended uniaxial state by a multiaxial one confined to a boundary effect at the face. This is the finite-element analogue of the well-documented experimental observation that the gripped region of a tensile coupon is under a combined longitudinal, transverse and shear state rather than pure tension, and that the restraint of Poisson contraction at the grip is a primary cause of premature failure there rather than in the gauge section [26].

Consequence for a progressive-damage model. Spurious transverse stresses of modest magnitude in absolute terms can satisfy a matrix or interlaminar initiation criterion at the constrained face, starting damage that has no physical counterpart in the test. Once initiated, this artificial damage softens the material locally and can perturb the load path and the energy balance of the simulation, contaminating the quantities the model is intended to predict. Leaving the transverse directions free at the loaded face removes this driving force at the source.

Spatial extent of the disturbance. Whether such a boundary layer reaches the region of interest is governed by Saint-Venant's principle: at a sufficient distance from the loaded end the stresses depend only on the resultant of the applied load and become independent of the exact manner in which it is imposed [4]. For an isotropic material this distance is of the order of the cross-sectional dimension, so that end effects decay within roughly one specimen width. This estimate cannot be carried over to a carbon fibre laminate. The routine application of Saint-Venant's principle is not justified for highly anisotropic or composite materials, because self-equilibrated end loads decay far more slowly than in the isotropic case [11, 31]. The characteristic decay length scales with $b\sqrt{E_1/G_{12}}$, where b is the largest dimension perpendicular to the fibres. As E_1 is fibre dominated and G_{12} matrix dominated this ratio can get much higher than in an isotropic media, so that end effects persist over a distance of the order of several specimen widths [31]. A boundary-layer disturbance in a CFRP coupon is therefore not confined to a thin region at the face: if the gripped length is short it can extend towards the hole, and the hole must consequently be placed several widths away from the loaded ends for its stress field to be free of boundary-condition artifacts.

Isolating the region of interest. Two measures follow from the preceding arguments. First, the end constraints can be limited to the degrees of freedom strictly required to react the load and to suppress rigid-body motion, with the transverse directions left free, so that the boundary layer is minimised at its source. Second, because a residual disturbance cannot be eliminated entirely and decays slowly in a composite, a common expedient is to assign the material in a strip adjacent to the loaded faces a purely linear-elastic response with the damage model deactivated. The constitutive damage and softening laws are then active only in the interior region, where the stress field is physically meaningful, so that boundary-condition artefacts cannot satisfy a failure criterion and start damage that would propagate into the region of interest. Such an elastic strip carries the load without altering the predicted response in the gauge and hole region, since by the decay argument above that region is, once placed sufficiently far from the loaded faces, insensitive to the exact end treatment [31].

5

Methodology

This chapter turns from theory to practice and describes how the model was actually built and exercised. The chapter follows the order in which a study is set up. Section 5.1 defines the problem and then introduces the three-model ablation (Section 5.1.1). Section 5.2 describes the automated pipeline through which all three are generated from a single parametrised program, and Section 5.3 the user material subroutine that supplies their constitutive response. Section 5.4 fixes the numerical settings common to all three models through dedicated parametric studies, and Section 5.5 closes the chapter by describing how the raw simulation output is condensed into the stress–strain curves used throughout Chapter 6. The section order mirrors that of the results in Chapter 6, so that the two chapters can be read in parallel from the simplest model to the most complex.

5.1. Problem Definition

Section 1.1 motivated the OHT coupon as the validation target of this thesis because it exercises matrix cracking, delamination and fibre failure together in a single, repeatable test [1]. The same property makes it the natural choice for the three-model ablation introduced below (Section 5.1.1): the hole imposes a stress concentration, so the stress state around it is steep and multi-axial rather than the near-uniform state of an unnotched coupon. It is exactly under such a gradient that the three models are expected to diverge, since a smeared CDM description and a discrete, crack-like one only differ where damage is localised; a specimen without a stress raiser would give the additions of Models B and C little opportunity to distinguish themselves from the Model A baseline. The choice of specimen is therefore not tied to a narrower research question of its own, but follows from wanting a geometry demanding enough to discriminate between the three modelling fidelities as well as to serve as a realistic validation case.

The research questions formulated in Section 1.3 are tested on an OHT specimen with a hole diameter $d = 6.35$ mm, a plate width $w = 50$ mm, and plate height $h = 32$ mm as seen in Fig. 5.1 (a) with the origin of the defined reference coordinate system being in the centre of the hole and at half thickness of the layup. Edges are labeled by their direction as seen in Fig. 5.1 (b).




Figure 5.1: Visualization of the chosen specimen. (a) shows the dimensions of the OHT specimen. (b) visualizes the applied boundary conditions, the edge labels, and the reference coordinate system.

Two layups are chosen:

1. Cross Ply (CP): $[0^\circ/90^\circ]_{4S}$
2. Quasi Isotropic (QI): $[45^\circ/90^\circ/-45^\circ/0^\circ]_{2S}$

with a ply thickness of $t = 0.125$ mm. The symmetry of the layups allows to model only half of the specimen, with a symmetry boundary condition applied at the mid-plane.

Two displacement boundary conditions are applied to the plate as seen in Fig. 5.1 (b). The displacement in x direction u_x on E_{x-} is fixed to zero. A tensile test is then performed by smoothly increasing u_x at E_{x+} from 0 to $u_{\max} = 2$ mm. All other DOFs are unconstrained.

The chosen material is IM7-8552 carbon/epoxy. The properties of this material are shown in Appendix B.

5.1.1. The Three-Model Ablation

In order to carry out this analysis three variations of the described specimen are introduced. The first model constructed is the simplest one, being a standard CDM model. This model will be referred to as Model A. It represents a baseline to show the improvements due to the additions in the other models.

The full semi-discrete model can benefit from the use of cohesive elements, to ensure optimal connectivity. To isolate the effects of the cohesive layer, Model B extends Model A by using CZM to model delaminations, implemented as cohesive elements at ply interfaces. Model C then adds the differentiation between bulks and splits and represents the full semi-discrete model critical to this thesis.

5.2. Model Generation and Automation

The three models introduced in Section 5.1.1 are not assembled by a single parametrised PYTHON program that drives Abaqus through its scripting interface, in keeping with the requirement for a fully automated model-generation workflow. The automation of Models A, B and C is the instrument used to isolate the marginal contribution of each discrete addition (Section 5.1.1), the three models must be identical in every respect that is not under study. Generating them from the same baseline code guarantees that the specimen geometry, the discretisation, the material data, the boundary conditions, and the solver settings

are shared, so that any difference in the predicted response can be attributed to the modelling features being ablated rather than to incidental differences in how the models were built.

The generator is organised as one shared build sequence with two branch points, as shown in Figure 5.2. The backbone, common to all three models, defines the specimen, discretises it, applies the load and boundary conditions, and submits the explicit analysis; the model-specific features branch off the backbone at the interlaminar and intralaminar.

Shared build sequence. The backbone first creates the geometry one ply at a time, each ply extruded to the ply thickness t and stacked in the assembly in the order of the layup. Building the laminate ply-by-ply rather than as a single block is what later allows each ply to carry its own fibre orientation and, in Model C, its own fibre-aligned partitioning. Every ply is assigned the orthotropic IM7–8552 stiffness and a material orientation rotated to its layup angle, and is discretised with a single eight-node reduced-integration solid element (C3D8R) through the thickness, the element selected with its hourglass control, previously explained in Section 4.3.

The safeguard against boundary-condition artefacts set out in Section 4.4 is realised at this stage. The loaded edges are constrained only in the axial direction (Section 5.1), leaving the transverse directions free so that the cross-section can contract; this avoids the boundary layer of transverse and shear stresses that a fully clamped end would impose by restraining Poisson contraction.

Model A: baseline CDM. In the simplest model the plies are bonded directly: adjacent ply surfaces are tied, so the interface is perfectly rigid and delamination cannot be represented. Every element of every ply is a homogeneous continuum carrying the full smeared damage law of Section 3.1.3, with both fibre and matrix modes active throughout. Model A is therefore a conventional three-dimensional CDM laminate and serves as the baseline against which the two discrete additions are measured.

Model B: cohesive interfaces. Model B takes the first branch: a cohesive layer is inserted between every adjacent ply pair to represent delamination by the cohesive-zone law of Section 3.2.2. Each interface is realised as a single layer of thin elements that carry the traction–separation response through the same VUMAT (Section 5.3) rather than through surface-based cohesive contact, the conforming-element route being preferred over contact for the reasons given in Section 3.4.1. The cohesive layers are connected to them by tie constraints. The plies themselves are unchanged from Model A, so the only difference relative to the baseline is the presence of the interlaminar law.

Model C: full semi-discrete. Model C takes the second branch and adds the bulk/split decomposition that defines the semi-discrete method (Section 3.4). Here the cohesive layers use a conforming mesh, so that interlaminar tractions transfer through shared nodes. Each ply is partitioned into thin strips aligned with its own fibre direction and spaced by the bulk spacing, as shown for every ply orientation of the laminate in Figures 5.5–5.8 below. Every resulting cell is classified as a split cell (inside a strip) or a bulk cell (between strips). The two populations are then given different material assignments: split cells receive a material in which matrix damage is active, while bulk cells receive one in which it is suppressed; fibre damage remains active in both. The two assignments are distinguished by colour in Figure 5.3, shown near the hole for a -45° ply. The crucial point is that this is the only mechanism by which the semi-discrete




Figure 5.2: Model-generation pipeline. A shared backbone (top and bottom) is common to all three models; Models A, B and C branch at the interlaminar and intralaminar decisions and re-converge before meshing. The dotted box is not used here but implemented as an optional addition.

behaviour is introduced. The split/bulk distinction is realised purely by geometric partitioning and material assignment, with the constitutive switch read by the VUMAT. No special element formulation is used: every element in Model C is the same standard continuum element as in Models A and B, in line with the defining premise of the semi-discrete approach (Section 3.4). The strips are meshed with one structured element across their width, the width being restricted by the Bažant–Oh limit of Eq. (3.8), while the bulk elements, freed of matrix failure, may be coarser (Section 3.4.1). The strip to bulk ratio are not fixed a priori but are treated as a parameter and investigated in Section 5.4.




Figure 5.3: Material assignment after classification in the semi-discrete model (Model C), shown near the hole for a -45° ply.

Independently of the split/bulk partitioning, a second, purely geometric partition set is introduced around the hole, shown in Figure 5.4: an inner region hugging the hole boundary within which a smaller element size can be prescribed, and an outer region covering the remainder of the plate. Because the stress concentration and the resulting damage initiation and growth in an OHT specimen are localised at the hole, refining the whole plate to the element size needed to resolve the hole would add elements, and therefore computational cost, without improving the accuracy of the prediction. Bounding the refined region to the vicinity of the hole confines this cost to where it is needed. The element size used in each region is not fixed and is one of the settings fixed by the mesh convergence study of Section 5.4.5; the figures below all use an outer element size of 0.7 mm and an inner, hole-refined element size of 0.3 mm, a combination used throughout for illustration and revisited later.




Figure 5.4: Geometric partition set used to locally control mesh density around the hole. Inner region (light), and outer region (dark) covering the rest of the plate.

The split and bulk partitions are cut directly into this hole-refined geometry, so the two partition sets interact: a strip crossing from the outer into the inner region must still mesh cleanly across the boundary between the two element sizes, and the elements near the hole boundary must not degenerate into slivers. Meeting both requirements while keeping the mesh structured and free of excessively distorted or badly

skewed elements required additional partition lines beyond the strips and the bulk/split boundary themselves. These partitions are not derived from a closed-form rule; they were obtained by trial and error, generating the mesh at each candidate partitioning, inspecting the resulting element quality around the hole, and adjusting the partition layout until the distortion was removed. The partitioning shown for each ply orientation below is the outcome of that process and is what is used throughout the rest of this work.

The four ply orientations of the laminates of Section 5.1 require different partition layouts, since the strip direction rotates with the fibre angle relative to the fixed hole and plate boundary. For the 0° and 90° plies the strips run parallel, respectively perpendicular, to the loading direction and can be carried across the full plate width without difficulty, so their partitioning spans the whole model (Figures 5.5 and 5.6). For the 45° and -45° plies, the partitions are not carried out to the plate edges (Figures 5.7 and 5.8). Carrying an off-axis strip partition all the way to a free edge forces elements at that edge into highly skewed parallelogram shapes, which produces small, distorted elements there and a correspondingly small stable time step, without any benefit, as an OHT specimen is designed so that failure initiates at the hole rather than at the edges. Stopping the off-axis partitions short of the edges therefore avoids this distortion and its computational cost while leaving the split/bulk resolution exactly where it is needed, around the hole.

Each of Figures 5.5–5.8 presents one ply orientation and is organised the same way: the fibre-aligned partition of the ply (left), the resulting mesh of the full model (centre), and a detail view of the mesh in the hole-refined region (right), so that the partitioning can be traced directly through to the mesh it produces.




Figure 5.5: Fibre-aligned partitioning and resulting mesh for the 0° ply.




Figure 5.6: Fibre-aligned partitioning and resulting mesh for the 90° ply.




Figure 5.7: Fibre-aligned partitioning and resulting mesh for the 45° ply.




Figure 5.8: Fibre-aligned partitioning and resulting mesh for the -45° ply.

A conforming mesh at the interlaminar boundary means that every cohesive layer must share nodes with both plies it bounds, not only the ply above or the ply below, as proposed by Nguyen et al. [53]. Since the two bounding plies carry different fibre orientations in the chosen laminates, the cohesive partition cannot simply repeat a single ply's strip pattern: it must carry the fibre-aligned strip family of each bounding ply, so that its nodes line up with the split/bulk boundaries on both faces at once. For the $0^\circ/90^\circ$ interfaces of the cross-ply laminate this produces the regular grid partition of Figure 5.9 (a); at every interface involving a $\pm 45^\circ$ ply, the corresponding strip family is carried at that angle, as in Figure 5.9 (b).




Figure 5.9: Cohesive-layer partitioning conforming simultaneously to the fibre directions of both bounding plies.

Meshed directly, the $0^\circ/90^\circ$ partition produces the clean, regular mesh of Figure 5.10. The $\pm 45^\circ$ partition, however encounters a problem, where a $\pm 45^\circ$ strip meets a $0^\circ/90^\circ$ strip and meets the plate edge, at a shallow angle, the cohesive elements there are compressed into extremely thin, highly skewed slivers. Figure 5.11 shows the resulting mesh together with a detail of one such element near the edge. These sliver elements are not merely low-quality but numerically unusable: already so distorted at the start of the

analysis, they fail the Abaqus/Explicit element-distortion checks well before any physically meaningful strain is reached, while the surrounding material is still elastic.




Figure 5.10: Cohesive mesh at a $0^\circ/90^\circ$ interface.




Figure 5.11: Cohesive mesh at a $\pm 45^\circ$ -involving interface, with a detail of one of the resulting sliver elements near the plate edge.

To remove this failure, the $\pm 45^\circ$ partition is stopped before meeting the edge of the plate while the $0^\circ/90^\circ$ is allowed to span the full length. In order to make this adaptable for all plate sizes and split and bulk dimensions the $\pm 45^\circ$ partition is implemented following,

$$n_s \geq \frac{\sqrt{2}}{2} \frac{w}{s_{\text{tot}}}, \quad (5.1)$$

where $n_s \in \mathbb{N}$ is number of split-bulk pairs to be skipped, w is the the span of the plane in y direction, and s_{tot} is the distance between two split strips orthogonal to the fibre direction. This is visualized in the adapted partition of Figure 5.12 (a). Since interlaminar damage in an OHT specimen is expected to concentrate at the hole in the same way intralaminar damage does, conforming connectivity may only be necessary there, and the conforming mesh imposed over the rest of the plate may be adding cost without adding accuracy. The resulting mesh, Figure 5.12 (b), is free of the distorted edge elements and is the partitioning used for every $\pm 45^\circ$ -involving cohesive interface throughout the work.




Figure 5.12: Adapted cohesive partitioning at a $\pm 45^\circ$ -involving interface, free of the distorted edge elements of Figure 5.11.

Even adapted, the conforming requirement still forces the cohesive mesh to resolve two fibre-aligned strip families simultaneously which results in small elements which make the computations significantly more expensive than the free-meshed, non-conforming cohesive layer of Model B. However, whether this increase in cost is paid for in result accuracy is examined as a parametric variant in Section 5.4.3.

5.3. VUMAT Implementation

The constitutive response of every element population in all three models is supplied by a single user material subroutine (VUMAT). This subroutine combines separate algorithms for the plies and the cohesive layers, both the intralaminar ply law and the interlaminar cohesive law. A per-material switch selects the relevant branch at run time. The bulk/split distinction of Model C is likewise a flag read from the material definition. The subroutine is called once per increment, per material card, and loops through all integration points of that material. It receives the strain increment for the current iteration, and returns the updated stress, while all path-dependent information is carried forward in solution-dependent state variables. The state variables retained between increments are grouped in Table 5.1.

Table 5.1: Groups of solution-dependent state variables carried by the VUMAT between increments. The ply branch uses the intralaminar groups; the cohesive branch uses the interlaminar group.

Group	Purpose
Current stress and accumulated strain	carry the integration-point state forward and provide output
Failure indices (FT, FC, MT, MC)	values of the initiation criteria of the current increment
Damage-onset quantities	equivalent strain and equivalent stress locked in for each mode at initiation
Damage variables	per-mode ply variables $d_{ft}, d_{fc}, d_{mt}, d_{mc}$ or the cohesive variable d_{coh}
Cohesive triangle quantities	mode-mixed energies, toughness, and onset and failure separations (cohesive branch)
Element-deletion flag	marks the element for removal once fully degraded
Characteristic lengths	frozen per-mode fibre and matrix lengths (see below)

5.3.1. Intralaminar VUMAT

The ply branch implements the three-dimensional smeared-damage law of Section 3.1.3; its per-increment logic is shown in Figure 5.14. The undamaged orthotropic stiffness is assembled in every increment

from the engineering constants according to Eq. (2.3), even though it does not change. This is done for future implementations in which the material properties may vary, such as temperature dependency. Each subsequent call accumulates the strain increment into the total strain, following the small strain assumption, and forms the effective (undamaged) stress, following the linear path sketched in Figure 5.13. The nominal stress σ_{old} stored from the previous increment is first mapped back up onto the undamaged response, removing the degradation it carried:

$$\tilde{\sigma}_{\text{old}} = \mathbf{M}_d \sigma_{\text{old}}, \quad (5.2)$$

where $\mathbf{M}_d$ is the damage operator, the inverse of the degradation operator $\mathbf{M}_d^{-1}$ of Eq. (3.37). The elastic stress increment $\mathbf{C} \Delta \epsilon$ is then added along that undamaged line to give the new effective stress $\tilde{\sigma}_{\text{new}}$. It is this effective stress that drives the failure criteria, and, once the damage variables have been updated as described below, it is degraded back down to the nominal stress $\sigma_{\text{new}} = \mathbf{M}_d^{-1} \tilde{\sigma}_{\text{new}}$ that is returned to the solver. This up-and-down process is exactly how the triangle shown in Figure 3.2 is constructed.




Figure 5.13: Effective-stress update in the ply VUMAT.

The Hashin failure criterion is chosen for this work, as it is the simplest criterion, in terms of implementation, that separates matrix and fiber failures as seen in Section 2.2.7. This separation is necessary for the differentiation of split and bulk elements. The choice of damage initiation criterion is not the focus of this work, and other criteria could be used, as long as this separation exists.

In Section 2.2.7, the Hashin failure criterion considers crack orientation. In this implementation a fixed crack plane assumption is made with three crack planes with crack plane normals aligned with the fibre direction and the two transverse directions. This way the Hashin criterion simplifies to three independent criteria, one for each crack plane:

$$\begin{aligned} F_{fi} &= \left(\frac{\sigma_{11}}{X_i} \right)^2, \\ F_{mi22} &= \left(\frac{\sigma_{22}}{Y_i} \right)^2 + \left(\frac{\tau_{12}}{S_{12}} \right)^2 + \left(\frac{\tau_{23}}{S_{23}} \right)^2, \\ F_{mi33} &= \left(\frac{\sigma_{33}}{Z_i} \right)^2 + \left(\frac{\tau_{13}}{S_{13}} \right)^2 + \left(\frac{\tau_{23}}{S_{23}} \right)^2, \end{aligned} \quad (5.3)$$

where i is either t for tension or c for compression, X_i , Y_i , Z_i , S_{12} , S_{13} and S_{23} are the material strengths in the respective directions, $F_{fi} \geq 1$ is the fibre failure criterion, and $F_{mi22} \geq 1$ and $F_{mi33} \geq 1$ are the

two matrix failure criteria. The crack plane with crack plane normal aligned with the through-thickness direction is not considered, neglecting the possibility of intralaminar delaminations. This simplification is made on the basis of the expectation that major delamination events will occur at the ply interfaces. These delaminations are captured by the cohesive elements in Models B and C. Nguyen and Waas [50] have shown that a Hashin criterion with a fixed crack plane assumption and neglected intralaminar delaminations gives reasonable results with all performed tests falling within an error to experimental mean less than one standard deviation.

A further simplification made throughout this work is that the in-situ effect surveyed in Section 2.2.7 is not modelled: the transverse-tensile strength Y_t and the in-plane-shear strength S_{12} are treated as ply-independent constants rather than raised for thin and embedded plies. Because the embedded plies are thereby assigned a lower apparent strength than they would have in reality, their matrix cracks are predicted to initiate prematurely, and more sub-critical matrix damage forms at a given load. Following the competition set out in Section 2.2.3, this shifts the predicted response in opposite directions for the two families of failure mode. For the fibre-dominated modes (brittle and pull-out) the additional matrix cracking increases notch blunting, relieving the stress concentration at the hole and delaying fibre failure to a higher applied load. For the delamination mode the same earlier and more extensive matrix cracking brings interlaminar damage forward and gives the delaminations more time to spread along the matrix cracks that bound them, accelerating the progression to failure. In particular, the in-situ ordering by which the outer interface delaminates first and each subsequent interface only at a higher load no longer holds, so the delamination reaches the $-45^\circ/0^\circ$ interface, and therefore final failure, sooner.

For every considered mode that has reached its failure criterion, the equivalent strain and equivalent stress are computed from Eqs. (3.31)–(3.32). The values at the instant of initiation are stored once as the peak of the bilinear law and not overwritten thereafter. The scalar damage variable for the mode is then advanced along the linear softening branch by Eq. (3.33), with the failure strain fixed by the crack-band seen in Eq. (3.7) using the appropriate per-mode characteristic length, discussed below, and constrained to grow monotonically. The semi-discrete restriction enters here: fibre damage is evaluated in every element, whereas matrix damage is advanced only where the split flag is set, in the matrix-split strips of Model C, and everywhere in Models A and B. The per-mode variables are combined into the fibre, matrix and shear damage variables of Eq. (3.36), which degrade the effective stress back to the nominal stress returned to the solver (Eq. (3.37)), as previously described. When both modes reaches near-complete degradation, specifically 0.9999, the element is flagged for deletion. The original semi-discrete modelling paper [50] deletes the element when either the fibre or matrix damage variable reaches near-complete degradation. However, an element that failed in a matrix mode can still carry load in the fibre direction, and vice versa. As matrix damage is confined to the splits in the semi-discrete model, this effect is not highly pronounced. In model A and B, however, the effect is more pronounced as matrix damage can occur anywhere in the model, leading to premature failure of the specimen.

This number is chosen to be as close to 1 as possible in order to avoid a reduction in the fracture toughness caused by premature element deletion. Additional guards delete elements whose strains or volume change have become physically meaningless, to protect the stability of the explicit step.




Figure 5.14: Per-increment logic of the intralaminar VUMAT, evaluated at each integration point. Sections 2.2.7 and 3.1.3.

Characteristic Length The crack-band regularisation of Section 3.1.1 requires a characteristic length, and the discussion of Section 3.1.2 showed that for a misaligned or non-square element this length must be computed from the element geometry. The derivation of Section 3.1.2 is for cracks running along the

fibre direction. As matrix damage and fibre damage follow independent softening laws, they require separate characteristic lengths. The characteristic length of the crack plane with crack plane normal in in-plane transverse direction is referred to as L_c^m and used for matrix damage. This is the length the equations of Section 3.1.2 have been derived for. Applying the same construction to the crack with crack plane normal in the fibre direction gives the characteristic length L_c^f for fibre damage. The presented quantities have to be redefined to consider both crack planes. The crack plane P with crack plane normal in the fibre direction will be referred to as P^f and the crack plane with crack plane normal in the in-plane transverse direction will be referred to as P^m . j denoting f for P^i with $i = m$ and $j = m$ for P^i with $i = f$ yields the following generalized definitions for a crack plane P^i :

$$\psi^i = \arccos(\hat{e}_j \cdot \hat{j}) \quad \text{from Eq. (3.15),} \quad (5.4)$$

$$\Delta a_{\text{ML}}^i = \vec{e}_j \cdot \hat{j} = \|\vec{e}_j\| \cos \psi^i \quad \text{from Eq. (3.19),} \quad (5.5)$$

$$L_{\text{c,ML}}^i = \frac{\|\vec{e}_i\| \sin \gamma_s}{\cos \psi^i} \quad \text{from Eq. (3.24),} \quad (5.6)$$

$$\Delta a_{\text{max}}^i = \begin{cases} \frac{\|\vec{e}_i\| \sin \gamma_s}{\sin \psi^i}, & \text{if } \vec{e}_f \cdot \hat{m} > \vec{e}_m \cdot \hat{m}, \\ \frac{\|\vec{e}_j\| \sin \gamma_s}{\sin(\psi^i + \gamma_s)}, & \text{otherwise,} \end{cases} \quad \text{from Eq. (3.25),} \quad (5.7)$$

$$\Delta a_{\text{avg}}^i = \begin{cases} \Delta a_{\text{max}}^i \frac{\vec{e}_j \cdot \hat{i}}{(\vec{e}_j + \vec{e}_i) \cdot \hat{i}}, & \text{if } \vec{e}_j \cdot \hat{i} > \vec{e}_i \cdot \hat{i}, \\ \Delta a_{\text{max}}^i \frac{\vec{e}_i \cdot \hat{i}}{(\vec{e}_j + \vec{e}_i) \cdot \hat{i}}, & \text{otherwise,} \end{cases} \quad \text{from Eq. (3.26),} \quad (5.8)$$

$$\text{and } L_{\text{c,CPND}}^i = L_{\text{c,ML}}^i \frac{\Delta a_{\text{ML}}^i}{\Delta a_{\text{avg}}^i} \quad \text{from Eq. (3.27),} \quad (5.9)$$

where the subscript CPND (Crack Plane Normal Direction) replaced FD (Fibre Direction) of Eq. (3.27).

Section 3.1.2 introduces two scenarios with two different characteristic lengths depending on the structure of the mesh. For perfectly fibre-aligned elements, both introduced equations result in the same characteristic length. For low misalignment of approximately $\psi \leq 15^\circ$ the crack will propagate along the mesh lines and Eq. (5.6) is used to compute the characteristic length. For higher misalignments the crack in the crack plane of interest will propagate along the respective direction of a crack with that crack plane in a zig-zag pattern and Eq. (5.9) has to be used to correct the characteristic length from Eq. (5.6). Depending on the misalignment of the element the according variant will be chosen:

$$L_c^i = \begin{cases} L_{\text{c,ML}}^i & \text{if } \psi^i \leq 15^\circ, \\ L_{\text{c,CPND}}^i & \text{otherwise.} \end{cases} \quad (5.10)$$

Furthermore, the edge vector computation as presented in Eqs. (3.12) and (3.13) assumes that all elements are hexahedrons that can be represented as quadrilaterals extruded in the out-of plane transverse direction. As this does not hold for an arbitrary mesh another method was found. First all vectors describing the edges are found from the nodal coordinates of the element. Because of node ordering conventions Abaqus uses, all of these vectors can be assigned an edge in an equivalent undistorted reference element in normalized coordinates of the same type as seen in Fig. 5.15 (a). For hexahedral elements all vectors corresponding to edges in the reference element pointing in the same direction are averaged by building

their sum and dividing by the number of vectors. This way three edge vectors are found that are assigned to the material directions using the scalar product similar to Eq. (3.14).

Prismatic elements are treated as hex elements with two edges being collapsed into one as seen in Fig. 5.15. When computing all vectors of the element the collapsed vector is assigned to the two separate edges of the hexahedral reference element. The assignment is then the same as for a hexahedral element. Even though this is an approximation for edge vectors representative of the prismatic element, a better approximation could be to average the two triangle sections and calculate two lengths as the average crack extend in the two material directions and the third by calculating the average edge vector from the edges connecting the two triangular parts of the element. As the chosen mesh is hex dominant the first variant is chosen so that the same algorithm can be used for both hexahedral and prismatic elements.




Figure 5.15: Mapping of a prismatic elements nodes to a hexahedral reference element.

The chosen method is formalized as follows. The nodes in the reference element can be represented in normalized coordinates η_1, η_2 and η_3 with $\eta_1, \eta_2, \eta_3 \in [-1, 1]$ with the normalized coordinate systems origin being the reference elements centroid. All nodes of an 8-node reference element fulfill $\eta_1, \eta_2, \eta_3 \in \{-1, 1\}$ resembling a cube. Based on Abaqus node index convention, each reference node $i \in 1, \dots, 8$ with normalized coordinate η^i is assigned a node $j \in 1, \dots, 8$ or $j \in 1, \dots, 6$ (for a hexahedral or prismatic element respectively) in the real element with coordinate x^j . A pair (i, j) is referred to using the index k representing the pair (k, j) . j_k then represents the node j belonging to the pair k . The average vector in one direction is found by summing all nodal coordinates multiplied with the component of the normalized coordinate belonging to the direction of the average to be found. The result is then divided by the number of vectors or equivalently half of the number of nodes of the reference element. An edge vector in direction $l \in \{1, 2, 3\}$ can therefore be found as

$$\vec{e}_l = \frac{2}{n_{\text{nodes}}} \sum_{k=1}^8 \eta_l^k x^{j_k}, \quad (5.11)$$

where n_{nodes} is the number of nodes in the reference element. This addition allows a possible extension of the third failure plane as well as support for prismatic elements that can be found in some hex-dominated meshes.

The characteristic lengths are calculated using the VUCHARLEN subroutine that, on the first increment

determines the characteristic lengths once per element and hands the values to the VUMAT. They are then frozen in state variables and reused for the remainder of the analysis, so the geometric computation is paid for only once.

5.3.2. Interlaminar VUMAT

The cohesive branch implements the bilinear traction–separation law of Section 3.2.2 on the thin cohesive elements that form the interface layers of Models B and C; its logic is shown in Figure 5.16. The element strains are interpreted as interface separations through the constitutive thickness as:

$$\varepsilon = \frac{\Delta L}{l_c} = \frac{\Delta}{l_c}, \quad (5.12)$$

where ΔL is the change in length that, for a crack appearing in the element, would extend the original length l_c by the length of the crack separation Δ . The elastic tractions follow from the penalty stiffness of Eq. (3.43). The three components are collapsed into the effective separation and effective traction of Eqs. (3.44) and (3.45), in which the Macaulay bracket removes the compressive normal contribution, and the quadratic nominal-stress criterion of Eq. (3.46) is tested.

While the criterion is below unity the interface is elastic and the linear tractions are returned unchanged. When it is first reached, the mode-mixed energies and the mixed-mode toughness from the power law of Eq. (3.49) are evaluated, and the bilinear triangle is fixed by recording the onset separation and traction and the failure separation $\Delta_e^f = 2G_{mc}/\tau_e^c$ of Eq. (3.50); these onset quantities are stored once and held constant. The scalar interface damage variable then follows the monotonic softening law of Eq. (3.51) and degrades the tractions through Eq. (3.52). When the damage variable approaches unity the element is flagged for deletion, at which point the interface is traction-free and the plies have fully delaminated. The returned stress tensor populates only the through-thickness normal traction and the two transverse shear tractions, consistent with the interpretation of the element as an interface.

5.4. Test setup

Before the three models can be compared, a number of numerical settings have to be fixed that are not part of the modelling features under study but nonetheless influence the result. Each is settled here through a dedicated parametric study, in which the setting is varied in isolation and chosen as the most economical value that no longer changes the response. Five such settings are addressed in turn: the approximation of quasi-static loading (Section 5.4.1), the cohesive penalty stiffness (Section 5.4.2), the conformity of the cohesive mesh to the ply partitioning (Section 5.4.3), the split and bulk strip dimensions of the semi-discrete mesh (Section 5.4.4), and the in-plane mesh density (Section 5.4.5). Because the explicit cost of each model is sensitive to all five, every study closes with the computational trade-off it entails. With these settings fixed, Section 5.5 describes how the resulting simulations are post-processed into the stress–strain curves used from here on.

5.4.1. Quasi Staticness

The OHT coupon is loaded quasi-statically, but the explicit solver is inherently dynamic and can only approximate that limit (Section 4.2). The two levers available for the approximation are the material




Figure 5.16: Per-increment logic of the interlaminar (cohesive) VUMAT, evaluated at each integration point. The traction–separation update, quadratic initiation, mixed-mode toughness and damage evolution follow Section 3.2.2.

density and the addition of damping, and their theoretical effect on the response are established in Section 4.2. When working with user defined materials in Abaqus/Explicit three mechanisms exist to approximate a quasi-static process.

The first and most straight forward method is to set a density such that computational time as well as the quasi-static approximation are satisfactory. As the critical time step is often restricted by a few small elements in the mesh, increasing the overall material density to raise the critical time step is inefficient as larger elements would unnecessarily add inertia making the process less quasi-static than required.

While the material density sets a baseline, mass scaling scales up only the density of element that cause a critical time step smaller than a user defined minimum. As sufficiently small elements do not contribute much to the overall inertia of the system, mass scaling allows to significantly reduce computational time without a high trade off. If applied to aggressively mass scaling will add to much inertia and the solution will be too different to that of an equivalent quasi-static process. Mass scaling therefore has to be carefully tuned to find the optimal operating point.

The third method is to add damping to the system. As discussed in Section 4.2 damping can be implemented as rayleigh damping, which is a combination of mass and stiffness proportional damping. As user subroutines only allow responses on constitutive level, only stiffness proportional damping can be implemented in the VUMAT, due to mass proportional damping requiring the actual velocity of the nodes, while stiffness proportional damping can be implemented using the strain rate instead. When implementing damping in the VUMAT this sets a limitation on the type of damping that can be applied and the quality of the damping effects: as seen in Fig. 4.2, stiffness proportional damping is a linear function of frequency, effective at high frequencies. If lower frequencies are to be damped this limitation will lead to an excessively high damping ratio at higher frequencies, which will lead to over-damping and smaller time steps, as the timestep depends on the damping ratio at the highest natural frequency of the system, as seen in Eq. (4.18). To implement this damping, the stress added though damping as of Eq. (4.31) and Eq. (4.36) is computed from the strain rate multiplied by the damaged stiffness matrix and a damping factor β_R that is tuned to achieve the desired damping behavior. The strain rate is computed from the strain increment from the previous increment to the current increment, divided by the time step in between these. The strain increment as well as the time step are provided by Abaqus to the VUMAT subroutine as part of the input arguments of the subroutine. When reporting the new stress to Abaqus the stress due to damping is added to the stress computed from the constitutive model. As Abaqus cannot distinguish between the instantaneous damping stress and the persistent stress from the constitutive model, the old stress reported by Abaqus as an input argument to the VUMAT subroutine cannot be used for the next increment, as it would falsely assume the damping stress to be persistent, leading to premature damage initiation, element deletion, and overall an elevated effective stiffness. Instead the stress excluding the damping stress is stored in a state variable and used as the old stress state in the next increment. The additional stress due to damping is then only added to the stress reported to Abaqus leaving the actual stress state unaffected. The stress update with damping is visualized in Fig. 5.17.




Figure 5.17: Stress update in the VUMAT including damping.

To obtain results similar to an equivalent quasi-static response density, mass scaling, and damping are adjusted together until the response is acceptably close and the simulation runtime is reasonable. The degree of quasi-staticity is monitored from the energy history of the whole simulation: the kinetic energy relative to the total (internal) energy is plotted over the complete load history, and the approximation is judged adequate when the kinetic energy stays a small fraction of the total throughout. The density reduction and damping are tuned against this ratio, the accepted settings being the least aggressive ones that keep it below the chosen bound.

5.4.2. Cohesive Penalty Stiffness

The penalty stiffness K of the traction–separation law (Eq. (3.43)) is not a measured material property but a numerical device: before damage initiates it exists only to enforce displacement continuity across the interface (Section 3.2.2). Inserting a cohesive layer to predict delamination must not, by itself, alter the elastic stiffness of an otherwise identical configuration. Too low a value introduces an artificial compliance that softens the laminate spuriously; too high a value is unnecessary and, as shown below, expensive. The purpose of this study is to fix K relative to the ply stiffness at the smallest value that leaves the elastic response unchanged.

Because a cohesive element transmits only the three interface tractions, one normal and two shear, that is modes I, II and III, the verification is run in the two loadings sketched in Figure 5.18: a mode-I opening across the interface, and a shear loading that is closer to the in-plane-dominated stress state actually

seen in the OHT specimen. Each loading is modelled twice on the same two-element coupon: once with the plies tied directly, which carries no interface compliance and provides the reference elastic stiffness (Figure 5.18, top), and once with a cohesive layer inserted between them (Figure 5.18, bottom). The penalty stiffness is then varied, expressed as a multiple of the ply through-thickness stiffness, and the elastic stress-strain slope of the cohesive model is compared with the tied reference in both loadings. The accepted value is the smallest K for which the two slopes coincide in both opening and shear.




Figure 5.18: Verification configurations for the cohesive penalty stiffness. A two-element coupon is loaded in interface shear ((a) and (c)) and in mode-I opening ((b) and (d)), modelled with the plies tied directly ((a) and (b)) and with a cohesive layer inserted ((c) and (d)).

Computational Cost The penalty stiffness also sets the cost of the cohesive layer. The thin interface elements already carry a small through-thickness dimension l_e , and raising K raises their effective stiffness; through the dilatational wave speed of Eq. (4.25) a stiffer element propagates stress waves faster, and through the stable-time-step bound of Eq. (4.16) (Section 4.1.3) the increased wave speed and the small l_e together drive the critical time step down, lengthening the run. The penalty stiffness must therefore be kept as low as the elastic-invisibility requirement above allows, rather than set arbitrarily high.

5.4.3. Cohesive Mesh Conformity

The conforming cohesive partitioning of Model C takes the fibre-aligned strip pattern of both bounding plies, even after the edge adaptation of Figure 5.12 this leaves the cohesive layer considerably more finely and more irregularly meshed than a free, unpartitioned mesh would be. That cost is only justified if the conforming connectivity it buys is actually needed.

This is tested by building a variant of Model C in which the cohesive layer is not partitioned to the ply strips at all: Abaqus is left free to mesh it with its own paving algorithm, unconstrained by the fibre-aligned partition, giving the mesh of Figure 5.19. Because this mesh no longer shares nodes with the plies it bounds, the tractions are transferred by interpolation across non-matching nodes instead of directly. The free-meshed variant is compared against the conforming, adapted mesh of Section 5.1.1 on the same specimen and loading, both in the predicted stress-strain response and in the computational

cost of the analysis, to establish whether the conforming cohesive mesh is actually required for an accurate prediction or whether it can be replaced by the cheaper free mesh without a material loss of accuracy.




Figure 5.19: Free, unpartitioned cohesive mesh

Computational Cost The two cohesive meshing strategies trade robustness for cost in opposite directions. The conforming, strip-partitioned mesh guarantees direct nodal load transfer at every point of the interface but forces small, sometimes irregular elements, each contributing to a smaller stable time step through Eq. (4.16). The free mesh removes that constraint and can be built with fewer, larger, more regular elements away from the hole, raising the stable time step; its cost is instead paid at the tie constraint, whose interpolation between non-matching meshes is less robust in explicit time integration than a truly conforming interface [67]. Whether the resulting saving is large enough, and the tie-constraint interpolation accurate enough, to prefer the free mesh over the conforming one is decided by the comparison of this study.

5.4.4. Split and Bulk Dimensions

The semi-discrete mesh of Model C is controlled by two geometric parameters: the width w_s of the fibre-aligned matrix-split strips and the bulk spacing s at which they are repeated. Not every spacing is admissible. The spacing that places w_s or s on an edge, produces degenerate sliver elements that cannot be meshed cleanly. The generator guards against this by suppressing any partition plane that would fall on the boundary, so that only the discrete set of spacings that tile the region without a strip landing on its edge can be used. To find the space between two splits orthogonal to the fibre direction $s_{\text{tot}} = w_s + s$ such that the corners of the plate land exactly on the interface between a split strip and a bulk region, the span of the specimen in y -direction w and the span of the specimen in x -direction l both have to result in a natural number when divided by s_{tot} projected onto the x or y unit vector:

$$\frac{w}{s_{\text{tot}}} = n_1 \wedge \frac{l}{s_{\text{tot}}} = n_2, \quad \text{with } n_1, n_2 \in \mathbb{N}. \quad (5.13)$$

Therefore, any natural number pair (n_1, n_2) that fulfills

$$\frac{w}{n_1} = \frac{l}{n_2}, \quad (5.14)$$

results in an admissible solution for s_{tot} with

$$s_{\text{tot}} = \frac{w}{n_1} = \frac{l}{n_2}. \quad (5.15)$$

Let p/q be the reduced rational number belonging to the fraction w/l , then Eq. (5.14) yields

$$n_1 = k p, \quad n_2 = k q, \quad \text{with } k \in \mathbb{N}. \quad (5.16)$$

Substituting this in Eq. (5.15) results in the equation giving all viable values for s_{tot} :

$$s_{\text{tot}} = \frac{w}{k p} = \frac{l}{k q}. \quad (5.17)$$

Nguyen and Waas [49] set the bulk to be 4 times the ply thickness and the split to be 10% of this, as the split always has to be significantly smaller. That said no justification is given. Testing the effect of the variations in dimension is the motivation for the study.

Computational Cost The split and bulk dimensions trade resolution against cost directly: a smaller strip width or a smaller bulk spacing produces more, and smaller, elements, which raises both the number of degrees of freedom and, through the reduced element length l_e in Eq. (4.16), the number of explicit increments required. The coarser bulk regions permitted by the semi-discrete scheme (Section 3.4.1) are what keep this cost bounded.

5.4.5. Mesh Convergence

The crack-band regularisation of Section 3.1.1 makes the dissipated fracture energy independent of element size through Eqs. (3.7) and (3.11), but mesh-size objectivity of the dissipated energy does not by itself guarantee that the global response is converged. The elastic stress concentration around the hole and the steep gradients that drive initiation both continue to depend on the in-plane mesh density, so the predicted peak load and damage pattern can still shift with refinement. The mesh convergence study below therefore refines the in-plane element size directly and confirms that the stress–strain response and the failure load stabilise.

For Model C, the study follows the inner/outer partitioning of Figure 5.4 and is carried out in two stages. The inner, hole-refined region is converged first, since it must resolve the stress concentrations at the hole it is the most critical. With the outer element size held fixed at 0.5 mm, the inner element size is refined through 0.5 mm, 0.375 mm, 0.25 mm and 0.125 mm, and the coarsest of the three at which the stress–strain response and failure load no longer change appreciably is retained. With the inner size fixed at that converged value, the outer element size is then varied in the second stage, from 0.5 mm up to the 1 mm, and again the coarsest value that leaves the response unchanged is retained. Figure 5.20 shows the effect of the inner refinement on the mesh directly, for the finest inner element size considered, 0.125 mm, on the 0° , 90° and 45° plies.




Figure 5.20: Mesh around the hole at the finest inner element size considered in the convergence study, 0.125 mm, shown for the 0°, 90° and 45° plies.

Models A and B carry no matrix-split strips, so the same convergence question is checked more simply for them: rather than the two-stage, inner/outer refinement above, a single, uniform in-plane element size is applied over the whole model, without a hole-refined inner region, and swept through 1 mm, 0.75 mm and 0.5 mm for both laminates; the coarsest size at which the stress–strain response and failure load stabilise is retained for each laminate.

Computational Cost Mesh refinement penalizes the explicit cost twice over. Halving the in-plane element size both multiplies the number of elements, and thus the internal-force work of Eq. (4.15), which scales with the element count, and shortens the stable time step of Eq. (4.16) in proportion to l_e , so the total cost grows faster than the element count alone. The converged mesh is therefore chosen as the coarsest discretization at which the response no longer changes appreciably.

5.5. Post Processing

For every frame of the quasi-static step, the stress field component aligned with the loading direction was weighted by the corresponding integration point volume field to obtain a single, volume-weighted average stress describing the mechanical state of the specimen:

$$\bar{\sigma}(t) = \frac{\sum_i \sigma_i(t) V_i(t)}{\sum_i V_i(t)},$$

where the sum runs over all integration points i of the model (integration points with a non-finite stress value, e.g. due to element deletion, are excluded), $\sigma_i(t)$ is the stress component along the loading axis and $V_i(t)$ its associated integration point volume. Because elements in the ply regions carry a local material orientation that can differ from element to element, the raw stress tensor field was first rotated into a fixed global coordinate system (defined once as an Abaqus datum coordinate system) before extracting the loading-direction component; for the cohesive interface elements, whose local coordinate system already coincides with the physically meaningful normal/shear traction directions, the untransformed stress components were used directly.

The corresponding strain was obtained from the displacement history of the loaded node set E_{x+} (Section 5.1), taken along the loading direction and normalized by the initial gauge length $L_0 = h$ of the

specimen,

$$\varepsilon(t) = \frac{u(t)}{L_0},$$

so that plotting $\bar{\sigma}(t)$ against $\varepsilon(t)$ for every extracted frame yields the engineering stress–strain curve of the model.

6

Results and Discussion

This chapter presents the numerical results obtained with the three-model framework introduced in Chapter 5. The presentation follows the same building-up logic used to construct the models themselves: Sections 6.1 to 6.4 verify the individual modelling choices, the cohesive penalty stiffness, the cohesive mesh partitioning, the quasi-static approximation, mesh objectivity, and split and bulk parametric analysis. in isolation, before Section 6.5 compares the resulting predictions against room-temperature open-hole tension experiments. Section 6.8 closes the chapter with the ablation comparison proper, isolating the contribution of the cohesive interface and the matrix-split elements by comparing the stress fields, damage initiation and propagation, and delamination patterns predicted by Models A, B and C against one another and against experimental X-ray CT data.

Because the room-temperature open-hole tension data against which the models are validated (Section 6.5) were supplied by the German Aerospace Center (DLR) and may not be reported directly, every stress-valued result in this chapter, model and experiment alike, is normalised by the mean experimental CP failure stress σ_{ref} of Table 6.2. The same reference value is used throughout for both laminates, so normalised results remain directly comparable to one another and to experiment; only their value in physical units is withheld.

6.1. Penalty Stiffness Test

The verification introduced in Section 5.4.2 loads the two-element coupon of Figure 5.18 in mode-I opening and in interface shear, comparing the plies tied directly against a cohesive layer of varying penalty stiffness K inserted between them.

Interface shear. Figure 6.1 shows the axial response for $K \in \{10^3, 10^4, 10^5\}$ MPa/mm against the tied reference. All three coincide with the tied response across the full loading range, indistinguishable even at the lowest penalty stiffness tested. This follows directly from the loading: the coupon is extended purely along the fibre direction with no through-thickness component except for lateral contraction, so nearly no

traction develops across the interface and the axial stiffness is set entirely by the plies, independent of K . This idealisation does not carry over to the OHT specimen: the stress redistribution around the hole can introduce interlaminar shear there, so the interface can still be loaded and drive delamination even though this isolated test shows no sensitivity to K at all.

Mode-I opening. Figure 6.2 shows the opposite behaviour. Loaded normal to the interface, the applied extension is carried directly across the cohesive layer, and the response is markedly sensitive to K : at $K = 10^3$ MPa/mm the assembly is visibly more compliant than the tied reference, and the response stiffens monotonically as K increases. Even at $K = 10^7$ MPa/mm, where the effective modulus reaches approximately 10.1 GPa against a the modulus of 10.9 GPa, the response remains approximately 7.3 % below the tied reference. This is consistent with delamination being a genuinely load-bearing mode under mode-I-dominated loading, it contributes to the total compliance in loading direction, so the penalty stiffness directly sets the effective elastic stiffness in the loading direction of the assembly.

Accepted value. Because the mode-I response converges between $K = 10^5$ MPa/mm and $K = 10^7$ MPa/mm, raising K further only adds computational cost (Section 5.4.2) without closing the residual gap to the tied reference. $K = 10^5$ MPa/mm is therefore adopted as the cohesive penalty stiffness for Models B and C throughout the rest of this work.




Figure 6.1: Two-element coupon in interface shear: axial stress–strain response for the tied reference and three cohesive penalty stiffnesses K .




Figure 6.2: Two-element coupon in mode-I opening: axial stress–strain response for the tied reference and three cohesive penalty stiffnesses K .

6.2. Cohesive Element Partition

Section 5.4.3 compares the conforming, ply-strip-partitioned cohesive mesh used by Model C against a free-meshed variant, in which the cohesive layer is left unpartitioned and tied to the plies through non-matching-mesh interpolation instead of shared nodes. Both variants are run on the full OHT specimen, for both laminates, and compared in stress–strain response (Figures 6.3 and 6.4) and in computational cost. Both models are set to a mass scaling of 10^{-6} s and a split to bulk ratio of 1/6 this is revised and studied in later sections, but for now this is the standard.

CP laminate. Figure 6.3 shows the two meshing strategies coincide throughout the elastic range and reach very similar peak stress, around 1.009 for the not partitioned and 1 for the partitioned, at $\varepsilon_x \approx 0.008$. The responses only begin to separate past the peak: the free-meshed variant collapses to the post-failure oscillation almost immediately, while the conforming mesh sustains a serrated, near-peak stress for a short additional strain interval before it too collapses. Once the specimen has fully failed both variants converge again to the same low-amplitude ringing, and the two failure strains differ only marginally.

QI laminate. Figure 6.4 shows an even closer agreement: the two curves are visually indistinguishable through the elastic region and through the gradual, damage-driven softening that precedes final failure, and differ slightly in the peak stress round 0.87 for the not partitioned and 0.85 for the not partitioned case. The post-failure residual stress level and its oscillation are matched as well.

Discussion. Given the difference in computational cost, roughly two days per analysis for the conforming mesh against a few hours for the free-meshed variant (Section 5.4.3), this level of agreement shows that the conforming cohesive partitioning contributes very little to the predicted response in the present configuration. This is consistent with how the specimen is loaded: the applied displacement is purely axial (Section 5.1), and the penalty-stiffness verification of Section 6.1 already showed that a cohesive

interface carries essentially no load, and is therefore insensitive to how it is meshed, once the applied strain has no through-thickness or transverse-shear component. Some interlaminar shear is still generated locally by the stress redistribution the hole introduces, and delamination is present in every simulation, but it remains a secondary failure mode: it is intralaminar, fibre-dominated failure that sets the peak load and drives its evolution, and for that the cohesive mesh topology is essentially irrelevant. This suggests that, for the present OHT configuration, the conforming cohesive partitioning can be replaced by the cheaper free mesh without a material loss of accuracy. Nixon-Pearson and Hallet record that the main failure mode for the QI-layup chosen in this work is pull-out [56], using the same hole to width ratio and material, further strengthening the made decision. In agreement to that, Nguyen and Waas find that using SD2M with the here chosen QI-layup results in pull-out failure and observe minimal delamination [54].




Figure 6.3: Stress–strain response of the CP laminate (Model C) with the conforming, ply-strip-partitioned cohesive mesh and with the free-meshed, non-conforming variant.

6.3. Quasi-Static Approximation

This study is carried out for Model C. As the most computationally demanding of the three models (Section 5.1.1), it is the one for which the trade-off between runtime and dynamic-inertia error is most consequential, and the one for which mass scaling is actually needed to reach a feasible runtime; Models A and B are treated separately below. The kinetic-to-internal energy ratio E_k/E_{in} is tracked over the full load history to assess the degree of quasi-staticity, as established in Section 5.4.1. Five mass scaling levels, parameterised by the target minimum stable time increment from 10^{-10} s to 10^{-6} s, are evaluated for each laminate. A larger target increment implies a shorter natural time step being scaled up, meaning more mass is added and the simulation becomes more dynamic; a smaller target imposes less mass but increases the runtime.

CP laminate. Figure 6.5 shows E_k/E_{in} on a logarithmic scale. The 10^{-6} scaling holds the ratio near unity throughout, indicating an inertia-dominated response; the corresponding stress–strain curve in Figure 6.6 confirms this, as no failure occurs and the laminate loads linearly beyond the expected failure




Figure 6.4: Stress–strain response of the QI laminate (Model C) with the conforming, ply-strip-partitioned cohesive mesh and with the free-meshed, non-conforming variant.

strain. The 10^{-7} scaling reduces the ratio to around 10^{-2} but remains elevated and produces an artificially smoothed post-failure response with a higher apparent peak stress. The three finer scalings (10^{-8} , 10^{-9} , 10^{-10}) track together at $E_k/E_{in} \approx 10^{-3}$ throughout the elastic phase. The transient spike at the failure strain ($\epsilon_x \approx 0.0085$) is an artefact of the fracture event and does not represent a loss of quasi-staticness in the load-carrying phase. Their stress–strain curves are mutually converged in both elastic stiffness and peak stress.

QI laminate. The QI laminate imposes a more demanding test because the progressive, sequential failure, matrix cracking, interlaminar damage, and final fibre fracture, generates several distinct dynamic events across a wider strain range. Figure 6.7 shows that 10^{-6} and 10^{-7} remain clearly dynamic throughout. The 10^{-8} and 10^{-9} scaling maintain an acceptable ratio during elastic loading but the kinetic energy rises to roughly 10^2 and 10^1 times the internal energy after the major failure event at $\epsilon_x \approx 0.024$ and $\epsilon_x \approx 0.028$ respectively. This indicates dynamic crack propagation rather than a quasi-static process. The 10^{-10} scaling remains within good bounds during the elastic phase and show only moderate transient spikes at each failure event, the kinetic energy rising up to only 10^{-1} times the internal energy. The process remains quasi static until final failure at $\epsilon_x = 0.004$. The stress–strain curves for these (Figure 6.8) agree closely in peak stress and post-peak character; the 10^{-6} and 10^{-7} cases produce qualitatively different responses with lower peak stresses and inertia-broadened softening.

Selected settings and computational cost. For the CP laminate, scalings of 10^{-8} and finer yield a converged, quasi-static response; the most economical of these is used. For the QI laminate the stricter criterion imposed by progressive failure requires 10^{-10} or finer. The wall-clock simulation times for each level are listed in Table 6.1; the selected settings are those that meet the quasi-static criterion at the lowest runtime.

All simulations were performed on 20 CPU cores on the TUDelft High-Performance Computing Cluster. The runtimes confirm the strong sensitivity to the mass scaling level: successive reductions of one decade




Figure 6.5: Kinetic-to-internal energy ratio E_k/E_{in} versus applied strain for the CP laminate at five mass scaling levels (logarithmic scale).




Figure 6.6: Stress-strain response of the CP laminate for five mass scaling levels.

Table 6.1: Wall-clock simulation time for the CP and QI laminates as a function of mass scaling target minimum time increment; all runs on 20 CPU cores.

Target $dt_{\min}$ (s)	CP [hh:mm]	QI [hh:mm]
10^{-6}	0:15	0:30
10^{-7}	0:20	0:37
10^{-8}	1:05	1:36
10^{-9}	3:46	3:27
10^{-10}	10:10	9:17




Figure 6.7: Kinetic-to-internal energy ratio E_k/E_{in} versus applied strain for the QI laminate at five mass scaling levels (logarithmic scale).




Figure 6.8: Stress-strain response of the QI laminate for five mass scaling levels.

in $dt_{\min}$ multiply the CP runtime, so choosing the least aggressive acceptable setting is decisive for a feasible parametric study. In absolute terms, however, the runtimes are strikingly low for a three-dimensional explicit progressive-damage analysis of a composite laminate. The accepted CP setting completes in just over one hour; even the most demanding case, the QI laminate at 10^{-10} , is roughly nine hours. For a model with concurrent intralaminar and interlaminar failure across multiple plies, runtimes of this order confirm the computational viability of the semi-discrete approach and make the full ablation and validation campaign practical on standard engineering compute hardware.

Models A and B. Models A and B are run purely as the baseline and serve as comparison against Model C (Section 5.1.1), so the same runtime is not at a premium for them the way it is for the full semi-discrete model: their simulations are already inexpensive without any mass scaling. For these two models it is therefore enough to lower the material density directly until E_k/E_{in} meets the same quasi-static criterion used above, rather than tuning a mass-scaling schedule; the resulting runtimes remain short, so the more elaborate mass-scaling study performed for Model C is not required for Models A and B.

Since the selected mass scalings keep E_k/E_{in} well below 5 % throughout the elastic and damage phases, the stiffness-proportional damping implemented in the VUMAT (Section 5.4.1) was not activated for the present OHT models. The capability is retained in the subroutine as a lever for configurations in which more aggressive mass scaling is unavoidable, for instance, meshes with locally very small elements, larger specimens, or highly dynamic failure events, and is relevant to future work involving such cases. Every other result reported in this chapter, for all three models and both laminates, has been checked against the same E_k/E_{in} criterion and confirmed to stay below the 5 % bound throughout loading. These energy histories are not reproduced individually for every subsequent result, as doing so would be redundant, but the quasi-static character of every reported result has been verified in this way.

6.4. Mesh Objectivity

Both branches of the VUMAT (Section 5.3) share the same crack-band regularization of Section 3.1.1: a per-element characteristic length is computed once by VUCHARLEN and frozen in a state variable (Section 5.3). Since Models A, B and C are all driven by this same subroutine, mesh objectivity is checked by refining the in-plane element size uniformly and comparing the resulting stress–strain response. For Models A and B, three element sizes are compared, 1, 0.75 and 0.5 mm, for both the CP and QI laminates. For Model C, since the semi-discrete mesh follows the two-stage, hole-refined convergence scheme of Section 5.4.5; the results shown below correspond to the first of these two stages, in which the outer element size is held fixed at 0.5 mm and only the hole-refined inner element size is varied.

Model A: baseline CDM. Figure 6.9 shows the QI laminate: the elastic slope and the peak stress, 0.72–0.75 and under 5 % apart, coincide for all three element sizes and remain close through the sharp post-peak drop. The response only separates in the extended post-failure region: the residual-stress plateau drifts apart with mesh size, rising toward 0.4 for the coarsest mesh but falling toward 0.32 for the finest by $\varepsilon_x = 0.04$, so objectivity holds at initiation and peak but is lost in the long-tail residual response. Figure 6.10 shows the CP laminate is less well behaved: the elastic response is again shared by all three sizes, but the peak stress now spreads over 1.02–1.12, about 10 %, and the peak strain shifts with it, from

$\varepsilon_x \approx 0.009$ for the coarser meshes to $\varepsilon_x \approx 0.0115$ for the finest, i.e. the finer the mesh, the higher and later the peak, however, there is an approximate convergence for 0.75 mm and 0.5 mm. The response settles into a comparatively narrow band immediately after failure, before diverging again over the extended post-peak failure range in the same manner as the QI laminate.

Model B: cohesive interfaces. Figure 6.12 shows the QI laminate reproduces the same pattern as Model A: all three element sizes coincide through the elastic range and the peaks being very close, 0.69–0.72, separating only after failure. Figure 6.11 shows the CP laminate is again more mesh-sensitive: the elastic slope is shared up to $\varepsilon_x \approx 0.007$, after which the coarsest mesh, 1 mm, peaks earliest and lowest, at 0.95 around $\varepsilon_x \approx 0.0085$, while the two finer sizes, 0.75 mm and 0.5 mm, track one another closely and reach a higher, later peak of about 1.07 near $\varepsilon_x \approx 0.011$, within 1 % of each other. The trend runs in the same direction as Model A: refining the mesh raises and delays the peak, and here the two finest sizes essentially converge, leaving the coarsest, 1 mm mesh as the outlier.

Discussion. The pattern is broadly consistent across both models: the QI laminate is close to mesh-objective up to and including the peak load for both Model A and Model B, while the CP laminate is more mesh-sensitive in both cases, with a finer mesh raising and delaying the peak. For Model B the two finest sizes, 0.75 mm and 0.5 mm, converge to within 1 % of one another, so the remaining spread is best read as the coarsest, 1 mm, mesh simply being too coarse for the CP laminate. A spread of this size between the coarsest mesh and the converged finer pair is still within acceptable bounds for the present study.




Figure 6.9: Mesh objectivity check for Model A (baseline CDM), QI laminate: stress–strain response at three in-plane element sizes.




Figure 6.10: Mesh objectivity check for Model A (baseline CDM), CP laminate: stress-strain response at three in-plane element sizes.




Figure 6.11: Mesh objectivity check for Model B (cohesive interfaces), CP laminate: stress-strain response at three in-plane element sizes.




Figure 6.12: Mesh objectivity check for Model B (cohesive interfaces), QI laminate: stress–strain response at three in-plane element sizes.

Model C: full semi-discrete. Figures 6.13 and 6.14 show the same check for Model C, at the four hole-refined inner element sizes 0.5, 0.375, 0.25 and 0.125 mm described above. For the CP laminate the four sizes coincide through the elastic range but separate sharply from there: the peak stress climbs from 1.009 at $\epsilon_x \approx 0.008$ for the coarsest inner size to 1.37 at $\epsilon_x \approx 0.014$ for the finest, a spread of roughly 35 %, with every refinement step still producing a visibly different, higher and later peak. The QI laminate repeats the same behavior, with the peak rising from 0.87 to 1.14, about 31 %. Unlike Models A and B, none of the four curves for either laminate show any sign of levelling off: the finest size tested, 0.125 mm, is not converged either, only the closest to it of the sizes considered.

This is a striking result, since Model C is driven by exactly the same VUMAT, the same crack-band regularisation and the same VUCHARLEN characteristic-length construction (Section 5.3) as Models A and B, and the spread above is two to three times larger than the worst case already seen for either of them on the CP laminate. The crack-band correction guarantees that the fracture energy dissipated by a single damaged element is independent of that element’s size (Section 3.1.1). For Models A and B, the main distinction is that matrix damage is smeared over the whole ply, with no attempt to represent the crack as a spatially localised band, so there is no process zone for the mesh to resolve beyond what the crack-band correction already accounts for. Model C is different by construction: the matrix-split strips exist precisely to give the matrix crack a discrete, spatially localized representation, similarly to how cohesive zone modeling aims to represent cracks. This resemblance in modelling strategies, reintroduces the same resolution requirement previously seen in cohesive-zone modelling (Section 3.2.2).

This fixes an upper bound on the split element size from the physical process-zone length of the matrix crack, evaluated from Eq. (3.53), using the ply’s matrix properties. Previously in Section 3.2.2 this requirement could be relaxed by using the adjustment of Eq. (3.54) and lowering the cohesive-interface strength, however, in this case the split material’s strength is a real material property that cannot be reduced to relax the requirement. Evaluating Eq. (3.53) with $M = 0.88$ separately for matrix tension and matrix compression gives process-zone lengths of approximately 0.75 mm in tension and 0.093 mm in compression; the compression limit is far more restrictive and would force an impractical element size, so

the working hypothesis adopted throughout this work is that matrix tension and shear dominate the failure process at the hole, contrary to matrix compression. This is checked against the damage patterns obtained later in this chapter, and only the tensile limit is enforced.

Fixing the element size at 0.75 mm, however, only guarantees that a single element fits inside the process zone, which is not the same as resolving it. The cohesive-zone literature of Section 3.2.2 is consistent that a single element per process zone under-integrates the traction distribution across it, and typically recommends at least about three elements spanning l_{cz} , with some sources recommending as many as ten for an accurate result [74]. Applied to the tensile process-zone length found above, a minimum of three elements implies an element size no larger than $0.75 \text{ mm}/3 \approx 0.25 \text{ mm}$. None of the coarser sizes tested above, 0.5 mm and 0.375 mm meet this criterion; only the two finest inner sizes, 0.25 mm and 0.125 mm, do.

The mechanism behind this is intuitive once stated. With too few elements spanning the process zone of the crack, the element sitting at the crack tip cannot share the load gradually with several neighbouring elements the way a fine mesh would: it alone must carry essentially the full traction gradient across the zone, and is pushed to its own local damage threshold correspondingly sooner. In a fine mesh, by contrast, the peak traction is distributed smoothly over many integration points, so any single point reaches its critical stress only relatively late in the loading history, and the zone unloads gradually as damage spreads element by element rather than being concentrated in one location. A crack resolved by too few elements therefore initiates, and loses its load-carrying capacity, earlier than the same crack resolved finely, which is exactly the trend seen in Figures 6.13 and 6.14: the coarser inner sizes reach their peak at both a lower stress and an earlier strain than the finer ones.

Figures 6.15 and 6.16 confirm this directly by tracking the average matrix damage at initiation for the same four inner element sizes. For both laminates, the 0.25 mm and 0.125 mm curves, the two sizes meeting the three-element criterion above, coincide almost exactly throughout, while the 0.375 mm and, more markedly, the 0.5 mm curves initiate earlier and reach a visibly higher average damage over the same strain range. The average damage variables are calculated as the damage variable of the individual elements averaged over the volume:

$$\bar{d}_i = \frac{\sum_{k=1}^{N_e} d_i^{(k)} V^{(k)}}{\sum_{k=1}^{N_e} V^{(k)}}, \quad (6.1)$$

where $d_i^{(k)}$ is the local damage variable for damage mode i of element k with a volume of $V^{(k)}$ and a total number of elements N_e . While a small damage area with high damage will give the same average as a higher area with lower damage variable in the elements, it is still a measure of damage initiation and gives insight into the severity of the damage in the model.

This, however, only converges the response local to the hole: all four curves above share the same fixed outer element size of 0.5 mm, well above the 0.25 mm requirement, and the global stress–strain curves of Figures 6.13 and 6.14 still fail to converge even once the inner size satisfies the criterion. The explanation is that damage initiates at the hole but does not stay there: once the crack in a matrix-split strip has initiated within the fine, hole-refined inner region, it continues to propagate along the strip into the coarser outer region as loading proceeds, and that region, at 0.5 mm, does not meet the three-element criterion. Because the the matrix-split strips, run the full width or length of the plate rather than being confined to the hole

vicinity, an outer mesh that does not meet the same resolution requirement corrupts the prediction just as effectively as an under-resolved inner mesh would, only later in the loading history. For Model C damage must be measured correctly not just where it initiates, but everywhere it is free to propagate, and for the matrix-split strips that means essentially everywhere in the model, not only at the hole. This is confirmed by refining the outer mesh together with the inner one, rather than holding it fixed at 0.5 mm as above: the outer element size is set to 0.25 mm, the largest size still meeting the three-element criterion derived above, and the inner, hole-refined size is further reduced to 0.125 mm. Figure 6.17 shows the resulting CP response: the two curves are now nearly indistinguishable through the elastic range and coincide almost exactly at the peak, 1.029 against 1.04 at $\varepsilon_x \approx 0.008$, a striking contrast to the 35 % spread obtained above when only the inner mesh was refined and the outer size left at 0.5 mm. The two curves separate only in the post-failure snap-back, which is expected: once the load-bearing crack has fully propagated, the drop and its subsequent oscillation are governed by dynamic timing rather than mesh resolution, and fall outside the load-bearing range of interest here. Figure 6.18 shows the same check repeats for the QI laminate: the elastic slope again coincides exactly, and the peak stress now agrees, 0.86 against 0.87 at $\varepsilon_x \approx 0.011$, against the 31 % spread obtained with only the inner mesh refined; the two curves separate only after failure, where the finer mesh sustains a longer serrated post-peak plateau before both settle into the residual response. This confirms the working hypothesis above for both laminates: it is the mesh everywhere the matrix-split strips are free to propagate, not only the region immediately at the hole, that must satisfy the process-zone resolution requirement of Eq. (3.53) for Model C to be mesh-objective.




Figure 6.13: Mesh objectivity check for Model C (full semi-discrete), CP laminate: stress-strain response at four hole-refined inner element sizes, outer element size fixed at 0.5 mm.




Figure 6.14: Mesh objectivity check for Model C (full semi-discrete), QI laminate: stress-strain response at four hole-refined inner element sizes, outer element size fixed at 0.5 mm.




Figure 6.15: Average matrix damage at initiation versus applied strain, Model C, CP laminate, for the same four hole-refined inner element sizes as Figure 6.13.




Figure 6.16: Average matrix damage at initiation versus applied strain, Model C, QI laminate, for the same four hole-refined inner element sizes as Figure 6.14.




Figure 6.17: Mesh objectivity check for Model C, CP laminate: stress-strain response with element sizes reduced to 0.25 mm and 0.125 mm.




Figure 6.18: Mesh objectivity check for Model C, QI laminate: stress–strain response with element sizes reduced to 0.25 mm and 0.125 mm.

6.5. Validation Against Experiments

With a mesh-objective prediction established above, the peak stress reached by Model C can now be compared against experimental open-hole tension failure data. Six room-temperature coupons of each laminate were tested at the German Aerospace Center (DLR); the measured failure stress σ_M for each specimen is listed in Tables 6.2 and 6.3.

Table 6.2: Experimental failure stress σ_M , CP laminate (DLR, room temperature).

Specimen ID	σ_M [MPa]
PP.B2-498-RT-01	1.043
PP.B2-498-RT-02	0.994
PP.B2-498-RT-03	0.96
PP.B2-498-RT-04	0.995
PP.B2-498-RT-05	0.993
PP.B2-498-RT-06	1.013
Mean value	1
Standard deviation	0.027

Table 6.3: Experimental failure stress σ_M , QI laminate (DLR, room temperature).

Specimen ID	σ_M [MPa]
PP.B2-495-RT-01	0.87
PP.B2-495-RT-02	0.897
PP.B2-495-RT-03	0.885
PP.B2-495-RT-04	0.91
PP.B2-495-RT-05	0.863
PP.B2-495-RT-06	0.89
Mean value	0.886
Standard deviation	0.017

For the CP laminate, the mean experimental failure stress is 1 with a standard deviation of 0.027, against the two mesh-converged Model C peak stresses of 1.029 (0.125 mm) and 1.04 (0.25 mm), over-predictions of about 3 and 4 % respectively, both within two experimental standard deviation. For the QI laminate the mean failure stress is 0.886 with a standard deviation of 0.017, against 0.86 (0.25 mm) and 0.87 (0.125 mm), under-predictions of about 2 and 1.5 %, both within one experimental standard deviation. That the two mesh sizes agree with experiment almost equally well for both laminates is itself a useful check: it is precisely because the 0.25 mm and 0.125 mm meshes are already converged with each other (Section 6.4) that they also converge to essentially the same comparison against experiment. This level of agreement, within or just outside one experimental standard deviation for both laminates and both mesh sizes, is a strong initial validation of Model C's predicted failure load.

Repeating the same comparison for Models A and B shows a different pattern from Model C: for the CP laminate, both models predict a peak stress higher than the experimental mean, 1.02 Model A and 1.07 for Model B. For the QI, similarly, both predict a peak stress higher than the experimental mean, 0.91 Models A and 0.95 for Model B. All using the results from the smallest element size.

This over-prediction is attributed to the excessive matrix damage exhibited by Models A and B, also seen in Section 6.8.1. A greater extent of matrix damage is assumed to accumulate over the loading history, and the consequence follows the notch blunting argument of Section 2.2.3: the excess splitting relieves more of the stress concentration at the hole than is physically present, so the 0° fibres reach their failure strain at an artificially high far-field stress.

6.6. Out-of-Plane Failure Neglected

Section 5.3.1 neglects the crack plane normal to the through-thickness direction, on the assumption that ply-interface delamination, captured by the cohesive elements, accounts for the dominant out-of-plane failure mode. To check this assumption, the corresponding matrix-tension criterion F_{mt33} is recorded at every integration point without being allowed to drive damage, using the converged 0.25 mm mesh. Figures 6.19 and 6.20 show the CP laminate: in the 0° ply the criterion is elevated only in a few small, scattered regions away from the hole, and in the 90° ply it is essentially inactive everywhere; in neither case does it approach unity, so no element would actually be degraded by this mode. The QI laminate shows even less activation and is not shown. The corresponding out-of-plane matrix-compression criterion F_{mc33} is recorded in the same way and is essentially zero everywhere for both laminates, and is therefore not shown either. This supports the assumption of Section 5.3.1: the neglected out-of-plane mode is negligible both in extent and severity.




Figure 6.19: Out-of-plane matrix-tension criterion F_{mt33} (recorded, not allowed to degrade), CP laminate, 0° ply, 0.25 mm mesh. The contour range is 0 (dark blue) to 1 (reds), where $F_{mt33} \geq 1$ denotes failure onset




Figure 6.20: Out-of-plane matrix-tension criterion F_{mt33} (recorded, not allowed to degrade), CP laminate, 90° ply, 0.25 mm mesh. The contour range is 0 (dark blue) to 1 (reds), where $F_{mt33} \geq 1$ denotes failure onset

6.7. Parametric Sensitivity

The split-to-bulk ratio w_s/s introduced in Section 5.4.4 is swept over four values, $1/2$, $1/4$, $1/6$ and $1/10$, for both laminates; a larger ratio widens the matrix-active split strips relative to the matrix-suppressed bulk, so more of the ply is allowed to sustain matrix damage. Figure 6.21 shows the QI laminate: peak stress rises and failure occurs progressively later as the ratio is reduced from $1/2$ to $1/10$. Figure 6.22 shows the opposite trend for the CP laminate: the $1/2$ ratio reaches the highest peak stress and $1/10$ the lowest. This is consistent with the Model A and Model B behaviour of Section 6.5: widening the split fraction toward $1/2$, where matrix damage is least restricted, moves the laminate toward the Model A and Model B limits observed previously. For the QI laminate a different behaviour is seen. Quantitatively, the CP laminate reaches a peak stress of approximately 1.06 at the $1/2$ ratio, falling between the Model A and Model B predictions reported earlier, while the QI laminate reaches approximately 0.74 at the same ratio, very different to the previously obtained results of Model A and Model B.

Interpretation through notch blunting. Section 2.2.3 set out that the accumulation of sub-critical damage delays fibre failure. A larger split fraction places more matrix-active material in the highest-stress region around the hole, so it should trigger this damage more and, when the final failure is fibre-dominated, push both the peak stress and the failure strain to their highest values as $w_s/s \rightarrow 1/2$.

CP laminate: mechanism confirmed. For the CP laminate this plays out as expected. In Figure 6.22, the point at which each curve first departs from the shared linear slope shifts earliest for $w_s/s = 1/2$ and progressively later as the ratio is tightened, exactly what the wider matrix-active fraction predicts. Because delamination is not a significant failure mode for this laminate under the present, purely axial loading (Sections 6.1 and 6.2), nothing intervenes between matrix damage onset and final fibre failure, so the delay carries through cleanly: $w_s/s = 1/2$ reaches the highest and latest peak of the four ratios tested.

QI laminate: delamination breaks the sequence. Figure 6.21 shows the same trend: matrix damage again accumulates most extensively for $w_s/s = 1/2$, visible as the largest deviation from the linear part of the curve. By the mechanism above this should also delay fibre failure the most, so $w_s/s = 1/2$ would be expected to reach the highest peak. Instead the opposite is observed: $1/2$ fails earliest and at the lowest peak stress. This is attributed to delamination: unlike the CP laminate, the QI laminate develops interlaminar damage ahead of final fibre failure, so the interface loses load-carrying capacity before the matrix-relief mechanism can fully delay the fibre-dominated failure, and the expected ordering is overridden. Section 2.2.3 notes exactly this sign change when the failure mode is one of delamination rather than fibre failure.

Convergence of the split-to-bulk ratio. As the split fraction is narrowed, the matrix-active strips increasingly approximate a spatially localised discrete crack, so the response might be expected to converge toward a limiting curve as w_s/s is reduced further. For the CP laminate the response is still shifting slightly between the $1/6$ and $1/10$ ratios; the predictions do, however, remain acceptably close to the experimental peak stress throughout, varying from 1.06 to 0.97. For the QI laminate the delamination interaction described above introduces an additional, ratio-dependent effect that makes any underlying convergence harder to isolate. Establishing whether a converged response is reached at a smaller split fraction is left for future work. Should no such limit exist, the split-to-bulk ratio would constitute a free parameter of Model C rather than a discretisation choice, which is a limitation of the present formulation.




Figure 6.21: Split-to-bulk ratio sensitivity, QI laminate: stress–strain response for $w_s/s \in \{1/2, 1/4, 1/6, 1/10\}$.




Figure 6.22: Split-to-bulk ratio sensitivity, CP laminate: stress–strain response for $w_s/s \in \{1/2, 1/4, 1/6, 1/10\}$.

6.8. Ablation Comparison

While the main objective of this thesis is to predict the damage behavior of the specimen, first the foundation of the model has to be verified. First the stress contour plot of all models are compared against the Lekhnitskii solution for an infinite plate with a finite width correction as presented in Section 2.2.5. The simulation results are then scaled to match a point on the plate, by dividing all simulation result values by the simulation result value at that point and are then multiplied by the concentration factor the analytical model predicts. Similarity of the other points will prove initial plausibility of the stress field that eventually initiates failure.

Figures 6.23 and 6.24 show the underlying analytical σ_1 (fibre-direction stress) fields used as this reference, for each ply orientation and both laminates, using a remote stress value of $\sigma_\infty = 1$, so the color scale reads directly as a local stress-concentration factor.




Figure 6.23: Analytical σ_1 stress-concentration field around the hole, CP laminate, 0° and 90° plies.




Figure 6.24: Analytical σ_1 stress-concentration field around the hole, QI laminate, 0° , 90° and 45° plies.

The comparison for both layups can be seen in Fig. 6.25 and 6.26.




Figure 6.25: σ_1 (fibre-direction stress) fields for the CP laminate, 0° and 90° plies.




Figure 6.26: σ_1 (fibre-direction stress) fields for the QI laminate, 0° , 90° and 45° plies.

While all models seem to be very close to one another and to the analytical result, it can be seen that Model C produces a stress contour that looks less smooth than Model A and Model B. This is most likely a consequence of the quality of the mesh of the model. Skewed elements and high aspect ratios can cause alterations in the stress field. All stress contours however are in good agreement with the theoretical solution, proving fundamental plausibility exists.

6.8.1. Damage Initiation and Propagation

The stress-field comparison above confirms that all three models reproduce a physically plausible elastic stress state around the hole; whether that stress state develops into a realistic damage pattern is checked separately here, for the QI laminate, against the X-ray computed tomography (CT) characterisation of open-hole tension specimens by Nixon-Pearson and Hallett [56]. Their baseline sub-laminate configuration, $[45/90/-45/0]_{2S}$, is the exact stacking sequence used for the QI laminate throughout this work, so the comparison below is made against a same-layup experimental reference. Furthermore, the same material system, IM7/8552, is used in both the experiments and the simulations.

For each ply orientation, the figures that follow show the predicted matrix damage from Model C on the left, the corresponding X-ray CT image from [56] in the centre, and the predicted matrix damage from Model B on the right. Model A is omitted: Sections 6.1 and 6.2 already established that delamination is not a primary failure mode for the present loading, so the ply-level matrix damage predicted by Models A and B, which differ only in their treatment of the interface, is visually very similar, and Model B is shown as representative of both. Two damage states are compared for each orientation: an early, near-initiation state, and a more developed state in which the matrix damage have propagated outward from the hole.

0° ply splits. Figure 6.27 shows the earliest matrix damage in the 0° ply forming as two short, tangential arcs directly above and below the hole, in the same location and orientation as the CT scan. Model B instead shows the same location smeared into a shallow, diffuse patch rather than a distinct line, consistent with its lack of a discrete crack representation. As loading continues, Figure 6.28 shows the Model C splits fanning outward from the hole into a widening set of parallel, slightly tilted in 45 direction, horizontal lines extending well beyond its immediate vicinity; the same fanned, multi-line pattern and extent is present in the CT image. Model B's damage remains a single, confined patch, it also exhibits the 45 direction propagation, but to a lesser extent.




Figure 6.27: Matrix damage variable at initiation for a 0° ply. (Color spectrum with red being fully damaged and blue being undamaged.)




Figure 6.28: Matrix damage variable after propagation for a 0° ply. (Color spectrum with red being fully damaged and blue being undamaged.)

45° ply splits. Figure 6.29 compares the 45° matrix splits, which appear together as patterns radiating from the hole along 45° diagonals. For Model C it is important to focus on the red damaged lines, since those are the truly cracked splits. Matching the radiating pattern reported for the +45° in the CT scan. Model B again places damage in the same diagonal orientation in a similar way, but as a thick, fully saturated bands rather than the finely spaced multi-line pattern seen in both Model C and the CT scan. At the more developed stage (Figure 6.30), the Model C splits extend further along the diagonals, it maintains the 45° overall damage pattern, again consistent with the CT image; Model B's bands are also consistent propagating in a very similar way.




Figure 6.29: Matrix damage variable at initiation for a $\pm 45^\circ$ ply. (Color spectrum with red being fully damaged and blue being undamaged.)




Figure 6.30: Matrix damage variable after propagation for a $\pm 45^\circ$ ply. (Color spectrum with red being fully damaged and blue being undamaged.)

90° ply splits. Figure 6.31 shows the 90° splits initiating to either side of the hole, oriented along the fibre direction and perpendicular to the loading axis, again matching the CT scan in both location and orientation. Model B produces two short, smeared patches at the same locations. At the more developed stage (Figure 6.32), the Model C splits have propagated to span almost the full height of the specimen, that widen away from the hole; the CT image at a comparable stage shows the same behavior, with 90° splits extending the full gauge length away from the hole. Model B develops in a similar way but with a broad, diffused, wedge-shaped damage zone.




Figure 6.31: Matrix damage variable at initiation for a 90° ply. (Color spectrum with red being fully damaged and blue being undamaged.)




Figure 6.32: Matrix damage variable after propagation for a 90° ply. (Color spectrum with red being fully damaged and blue being undamaged.)

Discussion. Across all three orientations, Model C reproduces the location, orientation and, importantly, the discrete, multi-line character of the matrix-split pattern observed by X-ray CT, while Model B's smeared damage representation still captures the correct location and orientation of the same failure, but as expected not its discrete character. This is the qualitative counterpart to the mesh-objectivity findings of Section 6.4: the same matrix-split strips that reproduce the CT damage pattern here are exactly the feature that introduces the additional mesh sensitivity discussed there.

CP laminate. No X-ray CT reference is available for the CP layup, so the comparison here is between Model C and Model B only; Model A is again omitted for the same reason as above. The same two damage states, initiation and propagation, are compared for the 0° and 90° plies, the only two orientations present in the CP laminate.

Figures 6.33 and 6.34 show the 0° ply: Model C again initiates as short, tangential splits at the hole edge before fanning outward into several parallel horizontal lines, the same pattern seen for the QI laminate above. Model B places damage in the same location and orientation but, being a smeared representation, spreads it into thick, fully saturated horizontal bands that already extend across most of the specimen width at the initiation stage.

Figures 6.35 and 6.36 show the same comparison for the 90° ply, with the same story: Model C produces discrete vertical splits flanking the hole that extend to nearly the full specimen height with propagation, while Model B smears the equivalent damage into a wide band directly at the hole and, at the propagated stage, a broad, diffuse region of lower-level damage spreading through much of the surrounding material. In both orientations Model B locates and orients the damage correctly; it simply does not resolve it into discrete cracks the way Model C, and the CT scans above, do.

The initiation figures for both orientations compare the two models at the same load increment, which makes just how aggressive this smeared response is directly visible: at that increment Model C shows only a small nick of damage confined to the elements right at the hole edge, while Model B has already driven a band of elements to full saturation across most of the specimen width. This is a direct consequence of the same lack of a discrete crack path noted above: without matrix split elements to concentrate the strain into a narrow, resolved band, the damage in Model B grows across the wider continuum region essentially at once rather than spreading element-by-element along a crack, so the same physical process reaches a fully damaged state far earlier in the loading history than it does in Model C.




Figure 6.33: CP laminate: matrix damage variable at initiation for a 0° ply. (Color spectrum with red being fully damaged and blue being undamaged.)




Figure 6.34: CP laminate: matrix damage variable after propagation for a 0° ply. (Color spectrum with red being fully damaged and blue being undamaged.)




Figure 6.35: CP laminate: matrix damage variable at initiation for a 90° ply. (Color spectrum with red being fully damaged and blue being undamaged.)




Figure 6.36: CP laminate: matrix damage variable after propagation for a 90° ply. (Color spectrum with red being fully damaged and blue being undamaged.)

6.8.2. Delamination Patterns

The interlaminar damage predicted by Model C is now compared in the same way as the matrix patterns of Section 6.8.1: against the X-ray CT dataset of Nixon-Pearson and Hallett [56] for the QI laminate, and against Model B only for the CP laminate, for which no CT reference is available. Model A is again omitted, since it contains no cohesive layer and therefore cannot represent delamination at all. The QI stacking sequence contains three distinct ply interfaces, $45/90$, $90/-45$ and $-45/0$, all compared below at a common, propagated load stage; the CP laminate has only the single $0/90$ interface, repeated through the thickness.

$45/90$ interface. Figure 6.37 compares the delamination front at the $45/90$ interface. Model C predicts a diagonal band of fully delaminated material running from the hole edge to the top right and bottom left, widening away from the hole. Two points of interest, marked A and B, are highlighted in both the Model C, CT images and Model B: the outlier patch of the upper diagonal (A) and the narrow damage initiation zone still connected to the front at to the hole edge (B). Both features are visible at matching locations in the CT scan and Model B, and Model C. Model C reproduces the same diagonal band in the same location and orientation, but as a broader, smoother region than the ragged, locally varying front seen in the CT scan.

$90/-45$ interface. Figure 6.38 shows the $90/-45$ interface. Model C predicts a thicker, hourglass-shaped delamination band pinched at the hole, consistent in shape and overall -45 propagation direction with the CT scan. However, the CT scan has a much narrower pattern running through the hole. Model B instead shows the delamination front, while much narrower, the -45 propagation direction is not as present.

$-45/0$ interface. Figure 6.39 shows the $-45/0$ interface. Model C predicts the delaminated band through the hole in the -45 diagonal; the CT image, shows a similar damaged region at the hole in a comparable direction. Model B produces a less regular pattern, only a very slight diagonal pattern is seen, and is also narrower when compared to the CT scan and Model C. A distinct feature of the CT scan is an undamaged band encircling the hole that separates the two diagonal delamination lobes, which is not reproduced by Model C or Model B, where the band runs continuously through the hole.




Figure 6.37: QI laminate: delamination at the 45/90 interface, with points A and B marking matching features in the Model C and CT results.




Figure 6.38: QI laminate: delamination at the 90/-45 interface.




Figure 6.39: QI laminate: delamination at the -45/0 interface.

CP laminate. Figure 6.40 compares the single 0/90 interface for the CP laminate; no CT reference is available, so only Model C and Model B are shown. Model C predicts a narrower and more sharply defined delaminated band directly through the hole, propagating vertically. Model B shows an overall similar propagation direction through the hole, however, the delamination has already spread across most of the specimen width, with only a diffused gradient rather than a sharp front.




Figure 6.40: CP laminate: delamination at the 0/90 interface.

Discussion. Unlike the matrix-split comparison of Section 6.8.1, where Model C’s discrete crack representation is a genuine modelling advantage over Model B’s smeared damage, both models here use the same cohesive-zone formulation for the interface itself. There is therefore no equivalent hierarchy of physical fidelity between the two for delamination, and the results above bear this out: neither model is a consistently closer match to the CT scans across the three QI interfaces. At the 45/90 interface, both Model C and Model B reproduce the same points A and B seen in the CT scan. At the 90/−45 interface, Model C captures the −45° propagation direction seen in the CT scan but over-predicts its width, while Model B is narrower but does not reproduce that propagation direction as clearly, each model capturing a different aspect of the same experimental pattern rather than one strictly outperforming the other. At the −45/0 interface, Model C matches the CT result closely in both direction and extent, while Model B is narrower and less regular. For the CP laminate no CT reference exists at all, so the difference between Model C’s narrow band and Model B’s wide, diffuse front cannot be checked against experiment; given the mixed picture found for the QI laminate, neither model’s prediction should be assumed the more accurate of the two.

Both of these shortcomings, the over-wide bands of Model C and the diffuse fronts of Model B, can most likely be traced to the same locality limitation that continuum damage models suffer from (Section 3.1.4): because no integration point is aware of the damage state at its neighbours and carries only its own history, the delaminated region grows along the mesh lines rather than along its physically correct direction. The damage-progression sequence of Section 2.2.3 makes this concrete, as the delamination should build up at the intersections of the matrix cracks of the adjacent plies and grow within the confined region those cracks bound. Partitioning the cohesive layer as originally intended (Section 6.2) would align the element geometry, and with it the mesh lines, with both adjacent fibre directions, so that the delaminations could initiate at the matrix-crack intersections and propagate in the physically correct directions. The undamaged band around the hole at the −45/0 interface (Figure 6.39) is a direct instance of this: the matrix splits of the 0° ply build tangentially to the hole, and a properly confined delamination would follow those splits, leaving the strip of interface between them and the hole intact, exactly the undelaminated band seen in the CT scan. Because this confinement is not captured here, Model C and B miss the band and reproduces only the rough propagation direction. The earlier conclusion that this partitioning can be omitted because it barely shifts the failure point is therefore incomplete: the failure load is indeed almost unchanged, but the delamination pattern is predicted incorrectly. In the worst case a wrongly predicted pattern could produce the wrong failure mode altogether, and with it a very different

failure load and post-peak behaviour. For the present specimens the stress–strain response happened to be almost insensitive to this, so it does not matter much here, but for other geometries and layups the same error could be far more consequential, and the additional computational cost of the partitioned interface would then be justified.

Conclusions and Outlook

This chapter summarizes the key findings of this thesis, addresses the research questions posed in Chapter 1, and outlines the limitations of the current work along with recommendations for future research directions.

Chapter 6 showed that the proposed semi-discrete framework shows good agreement with experimental data for both CP and QI laminates, capturing the key failure mechanisms, stress-strain responses, and failure patterns and are a significant improvement over standard CDM approaches in terms of fidelity without adding significant computational cost. As hypothesized the semi-discrete framework is the intersection between CDM and DCM approaches, with significant computational advantages over DCM approaches, while retaining the ability to capture the discrete nature of ply-level failure.

7.1. Conclusions

The conclusions drawn from the results presented in Chapter 6 allow all research questions posed in Chapter 1 to be answered. Each sub-research question is addressed below, with reference to the relevant findings and evidence. From that the main research question is answered, and the overall significance of the work is discussed.

7.1.1. SQ1: Contribution of the developed semi-discrete framework

The first sub-question asked what the cohesive interlaminar modelling and the embedded matrix-split intralaminar modelling each contribute to the predicted response, isolated by the three-model ablation of Section 5.1.1. Because Models A, B and C are generated from the same pipeline and differ only in the feature under study, the two contributions can be read off cleanly, and they turn out to be of very different character.

The cohesive interface, the single addition separating Model B from the baseline Model A, contributes little to the global response of the OHT coupon under the purely axial loading considered here. This holds

true for the chosen ply-thickness and layups but may not hold for other specimen configurations. The penalty-stiffness verification of Section 6.1 and the mesh-conformity study of Section 6.2 show that, in this case, an interface loaded almost entirely in-plane carries essentially no traction, so that the elastic stiffness and the peak load are set by the plies and are nearly insensitive to how, or even whether, the cohesive layer is meshed. Consistent with this, adding cohesive interfaces does not improve the failure-load prediction: for both laminates Model B over-predicts the experimental mean, and does so slightly more than Model A rather than less (Section 6.5). The contribution of the cohesive layer is therefore not one of failure-load accuracy but of damage-mode completeness: it is the only feature that allows delamination to be represented at all, and while that mode is secondary for the CP laminate, it becomes an active mode for the QI laminate, where it governs the split-to-bulk ordering of the parametric study (Section 6.7) and produces the interface patterns compared against X-ray CT in Section 6.8.2.

The matrix-split decomposition, the single addition separating Model C from Model B, is by contrast the decisive contribution of the framework, and the one that answers the hypothesis of Section 1.3. It acts on both quantities of interest. On the failure load, it is the feature that corrects the over-prediction shared by Models A and B: by confining matrix damage to thin, fibre-aligned strips rather than smearing it over the whole ply, Model C suppresses the premature, excessive matrix degradation that artificially relieves the stress concentration at the hole and lets the 0° fibres survive to an inflated loading stress (Section 6.5). On the damage evolution, it is the feature that reproduces the discrete, multi-line character of the matrix cracking seen in the CT data of Nixon-Pearson and Hallett [56], which Model B, lacking any discrete crack representation, captures only in location and orientation and not in character (Section 6.8.1). The two additions are thus not comparable in weight: the cohesive layer completes the set of representable failure modes, whereas the matrix-split strips are what make the prediction of both the failure load and the damage pattern realistic.

7.1.2. SQ2: Numerical objectivity and parameter sensitivity

The second sub-question asked whether the model behaves consistently and objectively at the numerical level. The answer is a qualified yes: the model is objective and consistent, but only once a resolution requirement stronger than that of a conventional CDM model is recognised and enforced, and one of its geometric parameters is not yet shown to be a pure discretisation choice.

The individual numerical settings each behaved consistently and were fixed by dedicated studies: the cohesive penalty stiffness converged to a value that leaves the elastic response indistinguishable from a tied interface (Section 6.1), the quasi-static approximation was verified against the kinetic-to-internal energy ratio for every reported result (Section 6.3), and the neglected out-of-plane matrix mode was confirmed to remain far below activation (Section 6.6). Models A and B are close to mesh-objective up to and including the peak load, as the crack-band regularisation intends.

The central and least expected finding of the verification concerns mesh objectivity of Model C. Because the matrix-split strips give the matrix crack a discrete, spatially localised representation, they reintroduce the same process-zone resolution requirement familiar from cohesive-zone modelling, and one that the crack-band correction alone does not satisfy (Section 6.4). Crucially, the strength reduction that relaxes this requirement for cohesive elements cannot be used here, because the strip carries a real material strength that must not be altered. Refining only the hole region therefore fails to converge Model C; objectivity is recovered only when the process-zone criterion is met everywhere the matrix-split strips

are free to propagate, which for the present geometry means essentially the whole plate. Once that stronger requirement is enforced, the response converges cleanly for both laminates. The discrete fidelity that answers SQ1 is bought at the price of a more demanding, and less local, mesh requirement than a smeared model imposes.

The one setting that could not be fully settled is the split-to-bulk ratio (Section 6.7). For the CP laminate its effect is understood through the notch-blunting mechanism and the predictions stay close to experiment across the range tested, but the response is still shifting slightly at the finest ratios, and for the QI laminate the delamination interaction, that is further discussed under Section 7.2, masks any underlying convergence. Whether the response converges to a limit as the split fraction is narrowed is therefore left open. Until that is established, the split-to-bulk ratio must be treated as a potential free parameter of the model rather than a settled discretisation choice, and this is the principal reservation attached to the otherwise consistent numerical behaviour of the framework.

7.1.3. SQ3: Result Fidelity

The third sub-question asked whether the model reproduces the failure load and the observed damage of experimentally tested layups. On the failure load, Model C agrees with experiment: its predicted peak stress lands within one experimental standard deviation for the QI laminate and within two for the CP laminate, so the small offset from the measured mean is no larger than the specimen-to-specimen variation of the tests themselves (Section 6.5). Models A and B do not match experiment as closely: for both laminates they predict a peak stress above the measured mean. This closer agreement of Model C is a direct result of the matrix-split decomposition identified under SQ1.

On the damage evolution the fidelity is high for intralaminar cracking and more modest for delamination. For matrix cracking, Model C reproduces the location, orientation and the discrete, multi-line character of the splits seen by X-ray CT across all three ply orientations of the QI laminate (Section 6.8.1), the same feature that Model B captures only as a smeared patch. The elastic stress field that precedes this damage was independently checked against the Lekhnitskii solution and found plausible for all three models (Section 6.8). For delamination, the picture is deliberately reported as mixed: across the three QI interfaces neither Model C nor Model B is a consistently closer match to the CT scans, each capturing a different aspect of the same experimental pattern, and Model C over-predicts the delamination width at one interface while matching another closely (Section 6.8.2). The fidelity claimed for the framework is therefore specific: it is strongest exactly where the semi-discrete addition acts, on the matrix-split cracking and the failure load it governs, and only qualitative for the interlaminar mode that both cohesive models share. For Model C this could be a direct consequence of the unpartitioned cohesive interface studied in Section 6.2, as delaminations occur at cracks in the adjacent plies as described in Section 2.2.3. With partitioned interfaces all split elements are directly tied to elements in the cohesive layer allowing localized delaminations. Unpartitioned, the delaminations are smeared over a larger area leading to less physical delamination patterns.

A second dimension of fidelity is computational, and it is here that a promise of the semi-discrete approach is met. The method exists to approximate the discrete fidelity of a discrete crack model without paying the cost that such models incur (Section 1.2), and the results bear this out. Model C reaches a discrete-like matrix-damage fidelity and the failure-load accuracy above while completing in roughly one hour for the CP laminate and at most about nine hours for the most demanding QI case on standard

High-Performance Computing (HPC) cluster hardware (Section 6.3). The matrix-split strips that give the model its fidelity do force very small elements, and with them a more aggressive mass scaling to keep the explicit step affordable. Because that mass scaling compensates for the finer split mesh, the runtime remains close to that of the far simpler continuum-based Model B rather than moving into a higher cost class. The extra fidelity is obtained at a cost comparable to that of a plain CDM analysis rather than at the cost of a discrete method. For a three-dimensional explicit progressive-damage analysis with concurrent intralaminar and interlaminar failure, runtimes of this order are strikingly low, and they are what make the whole ablation and validation campaign practical on standard engineering hardware. This is the clearest confirmation that the model occupies its intended position between the continuum and discrete families: discrete-like fidelity at continuum cost.

7.1.4. Main Research Question

The main research question asked to what extent a ply-level semi-discrete progressive-damage model can realistically predict the mechanical behaviour, damage evolution and failure of OHT specimens at room temperature. Taken together, the sub-questions support a positive but bounded answer. For the two carbon/epoxy layups tested, the model reproduces the failure load within one to two experimental standard deviations, reproduces the discrete matrix-damage patterns observed by X-ray CT, and does so at a computational cost that keeps the approach practical, thereby confirming the hypothesis of Section 1.3 that embedding a discrete-like representation of matrix cracking and a true discrete delamination into an otherwise continuum ply-level model yields a more realistic prediction than a continuum model alone.

Seen against the two families from which it is drawn, the framework meets the design intent set out in Chapter 3. The locality of CDM, which ties the crack path to the mesh, is removed by the fibre-aligned matrix-split strips that force the crack band onto the physically correct plane; the homogenisation limitation is approached as closely as a ply-level representation allows, by admitting matrix cracking only in its physically correct direction, confining it to the strip, and scaling the crack-band dissipation through the element characteristic length returned by the VUCHARLEN utility so that the released energy matches the fracture mode represented. The limitations that remain are essentially those the approach shares with CZM, chiefly that the crack plane must be fixed in the mesh in advance, while the single deficiency the method was designed to remove, the computational cost of an explicit discrete representation, is the one it demonstrably removes. In that specific sense the target the framework set out to meet is met.

The qualifier “to what extent” must, however, be answered honestly rather than optimistically. The demonstration rests on two layups of a single material system, one experimental campaign, and, for the damage patterns, a single CT reference available only for the QI laminate; the CP damage patterns and, more importantly, the failure loads of both layups are the quantities on which the strongest claims rest. Objectivity of the full model is conditional on a process-zone resolution requirement that is more demanding than a smeared model’s and that must be met over the whole region the splits traverse. The split-to-bulk ratio is not yet shown to be a convergent discretization parameter, and for the QI laminate the delaminations override the expected split-to-bulk ordering (Section 7.2). Within these bounds, and specifically for room-temperature open-hole tension of the layups studied, the semi-discrete framework predicts mechanical behaviour, damage evolution and failure realistically and at tractable cost; extending that conclusion to other layups, materials and loadings, remains the subject of future work.

7.2. Limitations and Recommendations for Future Work

While the semi-discrete framework developed in this thesis has demonstrated significant improvements in modeling composite laminate behavior, several limitations must be acknowledged. These limitations stem from the assumptions made during model development, the scope of validation cases, and the inherent complexities of composite materials.

7.2.1. Mesh Generation

While the mesh generated as described in Section 5.2 is sufficient for the current study, it exhibits multiple problems that limit its applicability to more complex geometries and loading conditions. The first one is that, in the mesh refinement at the hole introduces two unintentionally refined strips through the entire ply as seen in Fig. 5.5 (b) and Fig. 5.6 (b), which will unnecessarily increase the computational cost of the simulation, and visibly increases the elements aspect ratio. Fig. 5.7 (b) and Fig. 5.8 (b) show that the mesh of the $\pm 45^\circ$ plies contain lot of skewed elements. Furthermore, both meshes contain a lot of elements with aspect ratio greater than unity. Modeling the partitioning or otherwise altering the mesh generation could resolve these issues and is recommended for future work.

One inherent limitation of the semi-discrete meshing strategy is that strips only have an aspect ratio of unity if the element size is chosen to be the width of the strip. As the strips have to be thin enough to approximate discrete matrix cracks, this width is very small in all variation of the model. With all split strips being surrounded by bulk regions, the entire mesh is bound to this element size, making the method computationally far more expensive. Creating the split strips and bulk regions as separate parts would allow different element sizes by tying the parts together, but this would introduce interpolation errors at the tying interface and therefore does not necessarily lead to an improvement in the accuracy of the results and may even worsen them.

Furthermore, as the split elements approximate discrete cracks, the process zone of the matrix crack has to contain enough elements to correctly capture the matrix degradation and the stress redistribution around the crack. The process zone of the crack therefore represents an upper bound on the length of the strip and therefore the element size. While cohesive elements can circumvent this issue by reducing the interface strength as described in Section 3.2.2, artificially lowering the transverse strength would change a physical material property as split strips model the actual material. This leads to an inaccurate prediction of the stress field, peak stress, and possibly the damage evolution of the laminate.

7.2.2. Matrix-Split Spacing and Crack Density

A limitation the semi-discrete framework inherits from cohesive-zone modelling is that the plane on which a matrix crack may form must be fixed in the mesh in advance (Section 3.4). In the present scheme the matrix-split strips are placed at a prescribed, periodic spacing s_{tot} (Section 5.4.4), which fixes not only where matrix cracks can form but also the maximum crack density the model can represent. That spacing was set from meshing admissibility, and the parametric study of Section 6.7 varied the split-to-bulk ratio, that is the strip width, at a spacing chosen. The crack density itself was therefore essentially neglected. In reality the transverse-crack density is not a fixed quantity but evolves with load and depends on ply thickness, material and stress state; a periodic, load-independent placement cannot reproduce this evolution.

This is, however, exactly the setting in which the framework's cost advantage is most valuable. As argued in Section 3.4, a realistic transverse failure pattern consists of many densely spaced discrete splits, and the semi-discrete scheme can admit a large number of them without the prohibitive cost or convergence penalty that the same density would impose on a cohesive representation. The framework is therefore well placed to represent dense, and potentially non-uniform, crack fields; it is the choice of where and how densely to place the splits, not the cost of doing so, that limits it here.

A dedicated study of the split spacing, and thus of crack density, is therefore recommended, ideally informed by the observed crack density, for instance from the X-ray CT data, or by a statistical crack-density model. Placing the strips at physically motivated, and where appropriate non-uniform, spacings would let the model exploit its cost advantage to capture the evolving crack density. Predicting where and how densely matrix cracks form is, however, an open problem that is still the subject of active research in its own right, so it falls outside the scope of this thesis and would constitute a separate research effort rather than a minor extension of the present work.

7.2.3. Temperature Dependency

The current model is set up to run at a fixed temperature. As material properties change with varying temperature the stiffness matrix becomes a function of temperature and has to be adapted in each increment. The current algorithm of the VUMAT subroutine builds the stiffness matrix per step per integration point, a simple addition would therefore be a lookup table for all varying material properties. The values can then be interpolated to find an approximation for the material response at the given temperature. A more rigorous approach would be, to imbed the temperature dependency into the constitutive equations. This however requires more research effort and will, compared to the lookup table, raise the complexity and potentially the computational cost, whereas the lookup uses the same VUMAT algorithm, adding only the interpolation which is negligible, and in both cases the thermal simulation itself. Depending on the required fidelity both options are viable. Furthermore, thermal thermo-mechanical coupling adds thermal stresses changing the damage behavior of the specimen.

7.2.4. Non-Deterministic Material Properties

The current model is fully deterministic. Due to slight variations in environmental conditions, or through imperfections present in the specimen during manufacturing, such as resin voids, material properties change, making the failure of any specimen inherently non-deterministic. A stochastic FEM framework as introduced in Sec 3.3 would therefore lead to results closer to reality.

7.2.5. Constitutive Model

Characteristic Length

While the approximation of the characteristic length of hexahedral elements is derived from average crack extends through the element, the characteristic length of a prismatic element is found by treating it as a hexahedral element with two edges collapsed into one. This will result in the crack extend being overestimated, leading to an underestimation of the characteristic length, and in turn a higher failure strain using crack band. A recommendation for future work is to instead find the average crack extend of a prismatic reference element directly similar to

Bi-linear softening curve

This implementation uses a bi-linear softening curve, with three considerations to make.

The first part of the curve is modeled as a linear elastic material law. This comes with the assumption that the stress state before failure depends entirely on the strain at that moment. A real CFRP however also depends on the loading history due to microstructure changes of the material. Due to its good accuracy one widely used theory when modeling CFRPs materials is the Schapery theory [35]. It uses a thermodynamic based law to model non-linearities of the material [35].

The second consideration is the failure criterion determining the peak stress at which the material starts to soften. The failure criterion used here is the simplest criterion that separates matrix and fibre failure and therefore sufficient in the scope of this thesis. Different failure criteria however perform significantly better. The Cuntze criteria for example results in a significant improvement on the failure point prediction with similar computational complexity. Both methods however lack physical basis and are based on empirical parameters. Failure criteria with more physical could be used, trading accuracy for computational efforts. [18]

The second part of the curve is the linear softening. The assumption made is that growing damage degrades the material stiffness linearly. This however is not given for a real material and the stiffness-damage relation is not linear. Finding a more rigorous empirical model or a model with physical basis would improve the outcome.

7.2.6. Cohesive Element Partitioning

The conforming, ply-strip-partitioned cohesive mesh was compared against the cheaper free-meshed variant in Section 6.2, where the two were found to give an essentially identical stress-strain response and peak load for both laminates. On that basis the unpartitioned free mesh was adopted throughout, as it reproduces the global response at a fraction of the cost. That comparison, however, was made on the global response alone and under the assumption that delamination is a secondary failure mode; it did not weigh the fidelity of the delamination pattern itself.

For the QI laminate that assumption proved too strong: delamination is an active mode there, and neither meshing variant matches the measured interlaminar patterns consistently (Section 6.8.2). The reason, examined in that section, is the locality of the underlying continuum formulation (Section 3.1.4), which lets the interface damage spread along the available mesh directions instead of the ones the physics dictates. A partitioned interface removes that freedom by carrying the fibre orientation of both neighbouring plies into the cohesive mesh itself, giving the delamination the intersecting matrix cracks of Section 2.2.3 to nucleate on and a physically meaningful direction to follow; the free mesh, tied through interpolation, has neither and broadens the damage instead.

The decision to drop the partitioning was thus made on too narrow a basis. It is safe as far as the peak load goes, which the two meshes agree on, but it does not extend to the delamination pattern, which the free mesh gets wrong. That mismatch was inconsequential for the coupons studied here because their global response barely depended on it, yet a misjudged interlaminar pattern can in principle tip the analysis into the wrong failure mode, and from there into a markedly different strength and post-peak response. Confirming which mesh is preferable therefore calls for a direct comparison against the X-ray CT interlaminar data before either is treated as the default; for geometries and layouts less well characterised

than these, paying for the partitioned interface is the safer choice. This is left for future work.

7.2.7. Cohesive zone softening

The interface strength is not used at its physical value but is reduced through Equation (3.54) so that the process zone is resolved by a prescribed number of elements N_e . This reduction is the source of an unexpected result reported in Section 6.7: for the QI laminate delamination develops ahead of final fibre failure and becomes the critical mode, overriding the split-to-bulk ordering that the notch-blunting mechanism predicts and that the CP laminate follows cleanly. This is unexpected as multiple sources agree that the chosen QI specimens failure mode should be pull-out not delamination. Because lowering τ^0 shortens the interface's elastic range and brings interlaminar initiation forward, the reduction is a plausible cause of the premature delamination rather than a genuine physical feature of the laminate.

To test this, the QI laminate was re-run with the interface strength left at its physical value, i.e. without the reduction of Equation (3.54), using a cohesive-zone element size of 0.5 and a split-to-bulk ratio of 1/6. The two responses are compared in Figure 7.1. The cohesive damage variable d_{coh} in Figure 7.1 (b) makes the effect explicit: with the strength reduction it begins to grow early and accumulates steadily throughout loading, whereas at the physical strength it stays close to zero until very late, so interlaminar damage is strongly delayed.




Figure 7.1: Effect of the cohesive strength reduction on the QI laminate: (a) stress–strain response and (b) cohesive damage variable d_{coh} .

It must be stressed that in the stress–strain response of Figure 7.1 (a), the two responses cannot be compared on peak stress or failure point. The process-zone requirement of a minimum number of elements N_e still applies to the cohesive elements at the physical strength, but no adaptation is made to satisfy it; recovering a mesh-converged result would require either tuning N_e or refining the cohesive element size. Moreover, the element size of 0.5 used here is coarse, and Section 6.4 shows that for Model C the number of elements has a large effect on the predicted response. The peak stress in Figure 7.1 (a) should therefore not be read as a quantitative prediction; the comparison serves only to test whether the cohesive strength reduction is a plausible explanation for the delamination observed for the QI laminate, which it supports.

All results in Chapter 6, including the mesh-objectivity and split-to-bulk studies, were obtained with $N_e = 10$. As noted in Section 3.2.2, sources disagree on the appropriate N_e , and the value directly sets the reduced strength through Equation (3.54). It is therefore possible that a different N_e , giving a less

aggressive strength reduction, would not trigger the spurious delamination, and a calibration of N_e against the onset of interlaminar damage for this geometry is recommended before the reduction is relied upon. It should also be noted that Model B, which uses the same cohesive interfaces and the same strength reduction, does not exhibit this premature delamination; this warrants further investigation, as it suggests an interaction between the semi-discrete interlayer meshing and the reduced interface strength rather than a property of the cohesive formulation alone.

A direct consequence is that the split-to-bulk parametric study of Section 6.7 should be repeated once the interface strength is treated appropriately, either at its physical value or with a recalibrated N_e . With the delamination interaction removed, the QI laminate is expected to recover the same monotonic ordering as the CP laminate, in which peak stress and failure strain rise as the split fraction is widened toward $w_s/s \rightarrow 1/2$. Re-running the sweep under these conditions would also allow the convergence question raised under Section 6.7 to be assessed for the QI laminate free of the ratio-dependent delamination effect that currently masks it, and thus to establish whether a converged response is reached at a smaller split fraction. This is left for future work.

7.2.8. In-situ Strength

Throughout this work the ply transverse and shear strengths Y_t and S_{12} are held at fixed values, ignoring the *in-situ* elevation that thin and interior plies gain from the constraint of their neighbours (Section 2.2.7). An interior ply is consequently allowed to crack at a lower stress than it truly would, and the model builds up more matrix damage early in the load history than the specimen actually sustains.

That surplus of early cracking pulls the two failure families in opposite directions (Section 2.2.3). Where fibres govern, the extra splitting blunts the notch more effectively and postpones fibre fracture; where delamination governs, it seeds interlaminar separation earlier and leaves it more room to grow, leading to premature collapse. The staggering of interlaminar onset from the outer interface inward is lost as well, letting the damage arrive at the critical 0° interface earlier than it should.

This makes the missing in-situ strength a candidate explanation, alongside the reduced cohesive strength of Section 7.2.7, for why the QI laminate delaminates prematurely and fails in the wrong mode. In contrast to the cohesive strength, though, no confirmatory run was carried out here, so it stands only as a hypothesis and not a demonstrated cause. Either way, the in-situ effect is a genuine physical mechanism setting when matrix cracking begins, and hence shaping the entire damage sequence; incorporating it, by scaling Y_t and S_{12} with ply thickness and embedment, is recommended for future work.

References

- [1] Victor Achard et al. “Discrete ply modelling of open hole tensile tests”. In: *Composite Structures* 113 (2014), pp. 369–381. DOI: 10.1016/j.compstruct.2014.03.031.
- [2] Hongchen Bao and Guangyan Liu. “Progressive failure analysis on scaled open-hole tensile composite laminates”. In: *Composite Structures* 150 (2016), pp. 173–180. DOI: 10.1016/j.compstruct.2016.05.017.
- [3] G. I. Barenblatt. “The mathematical theory of equilibrium cracks in brittle fracture”. In: *Advances in Applied Mechanics*. Vol. 7. 1962, pp. 55–129. DOI: 10.1016/S0065-2156(08)70121-2.
- [4] Klaus-Jürgen Bathe. *Finite Element Procedures*. 2nd ed. Watertown, MA: K. J. Bathe, 2014. ISBN: 9780979004957.
- [5] Zdeněk P. Bažant and Byung H. Oh. “Crack band theory for fracture of concrete”. In: *Matériaux et Constructions* 16.3 (1983), pp. 155–177. DOI: 10.1007/BF02486267.
- [6] Ted Belytschko et al. *Nonlinear Finite Elements for Continua and Structures*. 2nd ed. Chichester, UK: John Wiley & Sons, 2014. ISBN: 9781118632703.
- [7] M. L. Benzeggagh and M. Kenane. “Measurement of mixed-mode delamination fracture toughness of unidirectional glass/epoxy composites with mixed-mode bending apparatus”. In: *Composites Science and Technology* 56.4 (1996), pp. 439–449. DOI: 10.1016/0266-3538(96)00005-X.
- [8] P. P. Camanho and C. G. Dávila. *Mixed-Mode Decohesion Finite Elements for the Simulation of Delamination in Composite Materials*. Tech. rep. NASA/TM-2002-211737. NASA Technical Memorandum. NASA, 2002.
- [9] P. P. Camanho, P. Maimí, and C. G. Dávila. “Prediction of size effects in notched laminates using continuum damage mechanics”. In: *Composites Science and Technology* 67.13 (2007), pp. 2715–2727. DOI: 10.1016/j.compscitech.2007.02.005.
- [10] P. P. Camanho et al. “Prediction of In Situ Strengths and Matrix Cracking in Composites Under Transverse Tension and In-Plane Shear”. In: *Composites Part A: Applied Science and Manufacturing* 37.2 (2006), pp. 165–176. DOI: 10.1016/j.compositesa.2005.04.023.
- [11] I. Choi and C. O. Horgan. “Saint-Venant’s Principle and End Effects in Anisotropic Elasticity”. In: *Journal of Applied Mechanics* 44.3 (1977), pp. 424–430. DOI: 10.1115/1.3424095.
- [12] Anil K. Chopra. *Dynamics of Structures: Theory and Applications to Earthquake Engineering*. 5th. Pearson, 2017.
- [13] R. Courant, K. Friedrichs, and H. Lewy. “Über die partiellen Differenzengleichungen der mathematischen Physik”. In: *Mathematische Annalen* 100 (1928), pp. 32–74.
- [14] R. G. Cuntze and A. Freund. “The Predictive Capability of Failure Mode Concept-Based Strength Criteria for Multidirectional Laminates”. In: *Composites Science and Technology* 64.3–4 (2004), pp. 343–377. DOI: 10.1016/S0266-3538(03)00218-5.


- [15] Dassault Systèmes Simulia Corp. *Abaqus 2016 Documentation*. Dassault Systèmes. Providence, RI, USA, 2016.
- [16] Dassault Systèmes Simulia Corp. *Abaqus 6.14 Analysis User's Guide*. Dassault Systèmes Simulia Corp. Providence, RI, 2014.
- [17] C. G. Dávila, P. P. Camanho, and C. A. Rose. "Failure Criteria for FRP Laminates". In: *Journal of Composite Materials* 39.4 (2005), pp. 323–345. doi: 10.1177/0021998305046452.
- [18] B. R. van Dongen. "Progressive Damage Modelling of FRPs Using a Blended Stress-Strain and Fracture Mechanics Approach in FEM". MA thesis. Delft University of Technology, 2017.
- [19] D. S. Dugdale. "Yielding of steel sheets containing slits". In: *Journal of the Mechanics and Physics of Solids* 8.2 (1960), pp. 100–104. doi: 10.1016/0022-5096(60)90013-2.
- [20] D. P. Flanagan and T. Belytschko. "Eigenvalues and stable time steps for the uniform strain hexahedron and quadrilateral". In: *Journal of Applied Mechanics* 51.1 (1984), pp. 35–40. doi: 10.1115/1.3167594.
- [21] E. K. Frankl. "Designing by Photoelasticity. R. B. Heywood. Chapman & Hall, London 1952. 414 pp. 65s. net." In: *The Journal of the Royal Aeronautical Society* 57.511 (1953), pp. 472–473. doi: 10.1017/S000192400012202X.
- [22] I. Fried. "Bounds on the extremal eigenvalues of the finite element stiffness and mass matrices and their spectral condition number". In: *Journal of Sound and Vibration* 22.4 (1972), pp. 407–418.
- [23] Roger G. Ghanem and Pol D. Spanos. *Stochastic Finite Elements: A Spectral Approach*. New York: Springer-Verlag, 1991. ISBN: 9781461230946. doi: 10.1007/978-1-4612-3094-6.
- [24] Karl F. Graff. *5.1 Waves Types*. 1975. URL: <https://app.knovel.com/hotlink/pdf/rcid:kpWMES0002/id:kt00B4FSK1/wave-motion-in-elastic/wave-types>.
- [25] Stephen R. Hallett, Wen-Guang Jiang, and Michael R. Wisnom. "The Effect of Stacking Sequence on Thickness Scaling of Tests on Open Hole Tensile Composite Specimens". In: *AIAA Journal* (2008). Cited in related literature as 'in press' in 2008/2009; check final volume/issue for your specific bibliography style.
- [26] Stephen R. Hallett et al. "The open hole tensile test: a challenge for virtual testing of composites". In: *International Journal of Fracture* 158.2 (2009), pp. 169–181. doi: 10.1007/s10704-009-9333-8.
- [27] A. Hansbo and P. Hansbo. "A finite element method for the simulation of strong and weak discontinuities in solid mechanics". In: *Computer Methods in Applied Mechanics and Engineering* 193.33–35 (2004), pp. 3523–3540. doi: 10.1016/j.cma.2003.12.041.
- [28] Z. Hashin. "Failure Criteria for Unidirectional Fiber Composites". In: *Journal of Applied Mechanics* 47.2 (1980), pp. 329–334. doi: 10.1115/1.3153664.
- [29] Z. Hashin and A. Rotem. "A Fatigue Failure Criterion for Fiber Reinforced Materials". In: *Journal of Composite Materials* 7.4 (1973), pp. 448–464. doi: 10.1177/002199837300700404.
- [30] R. Higuchi, T. Okabe, and T. Nagashima. "Numerical simulation of progressive damage and failure in composite laminates using XFEM/CZM coupled approach". In: *Composites Part A: Applied Science and Manufacturing* 95 (2017), pp. 197–207. doi: 10.1016/j.compositesa.2016.12.026.
- [31] C. O. Horgan. "Saint-Venant End Effects in Composites". In: *Journal of Composite Materials* 16.5 (1982), pp. 411–422. doi: 10.1177/002199838201600506.

- [32] E. V. Iarve et al. "Mesh-independent matrix cracking and delamination modeling in laminated composites". In: *International Journal for Numerical Methods in Engineering* 88.8 (2011), pp. 749–773. doi: 10.1002/nme.3195.
- [33] Daniel J. Inman. *Engineering Vibration*. 4th. Pearson, 2014. ISBN: 978-0132871693.
- [34] Meisam Jalalvand et al. "Energy dissipation during delamination in composite materials — An experimental assessment of the cohesive law and the stress-strain field ahead of a crack tip". In: *Composites Science and Technology* 134 (2016), pp. 115–124. doi: 10.1016/j.compscitech.2016.08.001.
- [35] M. Jamshidi and M.M. Shokrieh. "On the Schapery nonlinear viscoelastic model: A review". In: *European Journal of Mechanics - A/Solids* 108 (2024), p. 105403. ISSN: 0997-7538. doi: <https://doi.org/10.1016/j.euromechsol.2024.105403>.
- [36] A. S. Kaddour and M. J. Hinton. "Maturity of 3D Failure Criteria for Fibre-Reinforced Composites: Comparison Between Theories and Experiments: Part B of WWFE-II". In: *Journal of Composite Materials* 47.6–7 (2013), pp. 925–966. doi: 10.1177/0021998313478710.
- [37] Christos Kassapoglou. *Design and Analysis of Composite Structures: With Applications to Aerospace Structures*. 2nd ed. Aerospace Series. Chichester, UK: John Wiley & Sons, 2013. ISBN: 9781118401606.
- [38] Christos Kassapoglou. *Modeling the Effect of Damage in Composite Structures: Simplified Approaches*. Aerospace Series. Chichester, UK: John Wiley & Sons, 2015. ISBN: 9781119013211.
- [39] G. Kirsch. "Die Theorie der Elastizität und die Bedürfnisse der Festigkeitslehre". In: *Zeitschrift des Vereines Deutscher Ingenieure* 42 (1898), pp. 797–807.
- [40] Ronald Krueger. "Virtual crack closure technique: History, approach, and applications". In: *Applied Mechanics Reviews* 57.2 (2004), pp. 109–143. doi: 10.1115/1.1595677.
- [41] Ireneusz Lapczyk and Juan A. Hurtado. "Progressive damage modeling in fiber-reinforced materials". In: *Composites Part A: Applied Science and Manufacturing* 38.11 (2007), pp. 2333–2341. doi: 10.1016/j.compositesa.2007.01.017.
- [42] S. G. Lekhnitskii. *Anisotropic Plates*. Translated from the Russian. New York: Gordon and Breach, 1968. ISBN: 9780677206707.
- [43] Frank A. Leone and Brian P. Justusson. "Effects of characteristic element length on fracture energy dissipation in continuum damage mechanics models". In: *Journal of Composite Materials* 55.24 (2021), pp. 3551–3566. doi: 10.1177/00219983211023790.
- [44] Xiangqian Li, Stephen R. Hallett, and Michael R. Wisnom. "A finite element based statistical model for progressive tensile fibre failure in composite laminates". In: *Composites Part B: Engineering* 45.1 (2013), pp. 433–439. doi: 10.1016/j.compositesb.2012.08.001.
- [45] X. Lu et al. "On cohesive element parameters and delamination modelling". In: *Engineering Fracture Mechanics* 206 (2019), pp. 278–296. doi: 10.1016/j.engfracmech.2018.12.009.
- [46] F. P. van der Meer and L. J. Sluys. "A phantom node formulation with mixed mode cohesive law for splitting in laminates". In: *International Journal of Fracture* 158.2 (2009), pp. 107–124. doi: 10.1007/s10704-009-9344-5.
- [47] S. L. J. Millen, Z. Ullah, and B. G. Falzon. "On the importance of finite element mesh alignment along the fibre direction for modelling damage in fibre-reinforced polymer composite laminates". In: *Composite Structures* 278 (2021), p. 114694. doi: 10.1016/j.compstruct.2021.114694.

- [48] N. Moës, J. Dolbow, and T. Belytschko. "A finite element method for crack growth without remeshing". In: *International Journal for Numerical Methods in Engineering* 46.1 (1999), pp. 131–150. doi: 10.1002/(SICI)1097-0207(19990910)46:1.
- [49] Minh Hoang Nguyen and Anthony M. Waas. "A novel mode-dependent and probabilistic semi-discrete damage model for progressive failure analysis of composite laminates – Part I: Meshing strategy and mixed-mode law". In: *Composites Part C: Open Access* 3 (2020), p. 100073. doi: 10.1016/j.jcomc.2020.100073.
- [50] Minh Hoang Nguyen and Anthony M. Waas. "A novel mode-dependent and probabilistic semi-discrete damage model for progressive failure analysis of composite laminates – Part II: Applications to unnotched and open-hole tensile specimens". In: *Composites Part C: Open Access* 3 (2020), p. 100071. doi: 10.1016/j.jcomc.2020.100071.
- [51] Minh Hoang Nguyen and Anthony M. Waas. "Effect of stacking sequence and interlayer toughening in ± 45 laminates". In: *Journal of the Mechanics and Physics of Solids* 181 (2023), p. 105458. doi: 10.1016/j.jmps.2023.105458.
- [52] Minh Hoang Nguyen and Anthony M. Waas. "Efficient and Validated Framework for Probabilistic Progressive Failure Analysis of Composite Laminates". In: *AIAA Journal* 60.9 (2022), pp. 5500–5520. doi: 10.2514/1.J061499.
- [53] Minh Hoang Nguyen and Anthony M. Waas. "Modeling delamination migration in composite laminates using an enhanced semi-discrete damage model (eSD2M)". In: *International Journal of Solids and Structures* 236–237 (2022), p. 111323. doi: 10.1016/j.ijsolstr.2021.111323.
- [54] Minh Hoang Nguyen and Anthony M. Waas. "Numerical study on the effects of ply stacking and interlayer toughening on the failure modes and response of OHT laminates". In: *Composites Science and Technology* 234 (2023), p. 109935. ISSN: 0266-3538. doi: <https://doi.org/10.1016/j.compscitech.2023.109935>. URL: <https://www.sciencedirect.com/science/article/pii/S0266353823000283>.
- [55] M. Nguyen-Hoang and W. Becker. "Open holes in composite laminates with finite dimensions: structural assessment by analytical methods". In: *Archive of Applied Mechanics* 92.3 (2022), pp. 1101–1125. ISSN: 1432-0681. doi: 10.1007/s00419-021-02095-w.
- [56] O. J. Nixon-Pearson and S. R. Hallett. "An investigation into the damage development and residual strengths of open-hole specimens in fatigue". In: *Composites Part A: Applied Science and Manufacturing* 69 (2015), pp. 266–278. doi: 10.1016/j.compositesa.2014.11.013.
- [57] K. P. Oh. "A Monte Carlo Study of the Strength of Unidirectional Fiber-Reinforced Composites". In: *Journal of Composite Materials* 13.4 (1979), pp. 311–328. doi: 10.1177/002199837901300405.
- [58] Lucas Amaro de Oliveira and Maurício Vicente Donadon. "Delamination analysis using cohesive zone model: A discussion on traction-separation law and mixed-mode criteria". In: *Engineering Fracture Mechanics* 228 (2020), p. 106922. doi: 10.1016/j.engfracmech.2020.106922.
- [59] A. Parvizi, K. W. Garrett, and J. E. Bailey. "Constrained Cracking in Glass Fibre-Reinforced Epoxy Cross-Ply Laminates". In: *Journal of Materials Science* 13.1 (1978), pp. 195–201. doi: 10.1007/BF00739291.
- [60] S. T. Pinho et al. *Failure Models and Criteria for FRP Under In-Plane or Three-Dimensional Stress States Including Shear Non-Linearity*. Tech. rep. NASA/TM-2005-213530. NASA Langley Research Center, 2005.

- [61] S. T. Pinho et al. "Material and Structural Response of Polymer-Matrix Fibre-Reinforced Composites". In: *Journal of Composite Materials* 46.19–20 (2012), pp. 2313–2341. doi: 10.1177/0021998312454478.
- [62] A. Puck and H. Schürmann. "Failure Analysis of FRP Laminates by Means of Physically Based Phenomenological Models". In: *Composites Science and Technology* 58.7 (1998), pp. 1045–1067. doi: 10.1016/S0266-3538(96)00140-6.
- [63] John William Strutt Rayleigh. *The Theory of Sound*. 2nd. Vol. 1. Macmillan, 1894.
- [64] C. A. Rose, C. G. Dávila, and F. A. Leone. *Analysis Methods for Progressive Damage of Composite Structures*. Technical Memorandum NASA/TM-2013-218024. NASA, 2013.
- [65] E. F. Rybicki and M. F. Kanninen. "A finite element calculation of stress intensity factors by a modified crack closure integral". In: *Engineering Fracture Mechanics* 9.4 (1977), pp. 931–938. doi: 10.1016/0013-7944(77)90013-3.
- [66] G. N. Savin. *Stress Concentration around Holes*. Oxford: Pergamon Press, 1961.
- [67] Kyongchan Song, Carlos G. Dávila, and Cheryl A. Rose. "Guidelines and Parameter Selection for the Simulation of Progressive Delamination". In: *2008 Abaqus Users' Conference*. Conference paper; no DOI. Newport, RI, 2008.
- [68] Srinivas Sriramula and Marios K. Chryssanthopoulos. "Quantification of uncertainty modelling in stochastic analysis of FRP composites". In: *Composites Part A: Applied Science and Manufacturing* 40.11 (2009), pp. 1673–1684. doi: 10.1016/j.compositesa.2009.08.020.
- [69] George Stefanou. "The stochastic finite element method: Past, present and future". In: *Computer Methods in Applied Mechanics and Engineering* 198.9–12 (2009), pp. 1031–1051. doi: 10.1016/j.cma.2008.11.007.
- [70] S. C. Tan. "Finite-Width Correction Factors for Anisotropic Plate Containing a Central Opening". In: *Journal of Composite Materials* 22.11 (1988), pp. 1080–1097. doi: 10.1177/002199838802201105.
- [71] T. E. Tay et al. "Progressive Failure Analysis of Composites". In: *Journal of Composite Materials* 42.18 (2008), pp. 1921–1966. doi: 10.1177/0021998308093912.
- [72] William Thomson. "On the elasticity and viscosity of metals". In: *Proceedings of the Royal Society of London* 14 (1865), pp. 289–297.
- [73] S. W. Tsai and E. M. Wu. "A General Theory of Strength for Anisotropic Materials". In: *Journal of Composite Materials* 5.1 (1971), pp. 58–80. doi: 10.1177/002199837100500106.
- [74] A. Turon et al. "An engineering solution for mesh size effects in the simulation of delamination using cohesive zone models". In: *Engineering Fracture Mechanics* 74.10 (2007), pp. 1665–1682. doi: 10.1016/j.engfracmech.2006.08.025.
- [75] Woldemar Voigt. "Ueber die innere Reibung der festen Körper, insbesondere der Krystalle". In: *Abhandlungen der Königlichen Gesellschaft der Wissenschaften zu Göttingen* 36 (1890), pp. 3–47.
- [76] Oliver Völkerink. *Simulation-driven design of bonded joints in fibre composite aircraft structures using progressive damage analyses*. Forschungsbericht DLR-FB-2022-24. Braunschweig, Germany: Deutsches Zentrum für Luft- und Raumfahrt (DLR), 2022. doi: 10.57676/hya4-2a09.
- [77] W. Weibull. "A Statistical Distribution Function of Wide Applicability". In: *Journal of Applied Mechanics* 18.3 (1951), pp. 293–297. doi: 10.1115/1.4010337.
- [78] J. M. Whitney and R. J. Nuismer. "Stress Fracture Criteria for Laminated Composites Containing Stress Concentrations". In: *Journal of Composite Materials* 8.3 (1974), pp. 253–265. doi: 10.1177/002199837400800303.

-
- [79] T. Willmann, M. Schilling, and M. Bischoff. "Data-based estimation of critical time steps for explicit time integration". In: *International Journal for Numerical Methods in Engineering* 126 (2025), e7666. doi: 10.1002/nme.7666.
- [80] M. R. Wisnom. "Modelling discrete failures in composites with interface elements". In: *Composites Part A: Applied Science and Manufacturing* 41.7 (2010), pp. 795–805. doi: 10.1016/j.compositesa.2010.02.011.
- [81] Yichen Zhang, Wim Van Paepegem, and Wouter De Corte. "An Enhanced Progressive Damage Model for Laminated Fiber-Reinforced Composites Using the 3D Hashin Failure Criterion: A Multi-Level Analysis and Validation". In: *Materials* 17.21 (2024), p. 5176. doi: 10.3390/ma17215176.



Equivalence of the Central Difference Scheme and Newmark with $\beta = 0$

The central difference scheme of Section 4.1.3 is formulated in half-step velocities $\dot{\mathbf{u}}_{n\pm\frac{1}{2}}$, whereas the Newmark equations (4.2)–(4.3) are written in on-step quantities $\dot{\mathbf{u}}_n, \ddot{\mathbf{u}}_n$. To show that both generate the same displacement sequence $\{\mathbf{u}_n\}$, it suffices to introduce a consistent on-step velocity and verify that the two Newmark updates follow, with $\beta = 0$ and $\gamma = \frac{1}{2}$.

Time-grid relations. Because $t_{n+\frac{1}{2}} = \frac{1}{2}(t_{n+1} + t_n)$ is the midpoint of the interval $[t_n, t_{n+1}]$, each half-step point splits its interval evenly:

$$t_{n+\frac{1}{2}} - t_n = \frac{1}{2} \Delta t_{n+\frac{1}{2}}, \quad t_n - t_{n-\frac{1}{2}} = \frac{1}{2} \Delta t_{n-\frac{1}{2}}. \quad (\text{A.1})$$

Adding the two contributions that make up $\Delta t_n = t_{n+\frac{1}{2}} - t_{n-\frac{1}{2}}$ gives

$$\Delta t_n = \frac{1}{2} (\Delta t_{n+\frac{1}{2}} + \Delta t_{n-\frac{1}{2}}). \quad (\text{A.2})$$

On-step velocity. We define the on-step velocity as the half-step velocity advanced to the on-step point by the on-step acceleration $\ddot{\mathbf{u}}_n$,

$$\dot{\mathbf{u}}_n := \dot{\mathbf{u}}_{n-\frac{1}{2}} + (t_n - t_{n-\frac{1}{2}}) \ddot{\mathbf{u}}_n = \dot{\mathbf{u}}_{n-\frac{1}{2}} + \frac{1}{2} \Delta t_{n-\frac{1}{2}} \ddot{\mathbf{u}}_n, \quad (\text{A.3})$$

which is the natural on-step counterpart of the staggered scheme. Eliminating $\dot{\mathbf{u}}_{n-\frac{1}{2}}$ between (A.3) and the central-difference velocity update (4.13), and using (A.2), yields the forward relation

$$\dot{\mathbf{u}}_{n+\frac{1}{2}} = \dot{\mathbf{u}}_{n-\frac{1}{2}} + \Delta t_n \ddot{\mathbf{u}}_n = \dot{\mathbf{u}}_n + \frac{1}{2} \Delta t_{n+\frac{1}{2}} \ddot{\mathbf{u}}_n. \quad (\text{A.4})$$

Equation (A.4) states that the half-step velocity is the on-step velocity projected forward across the first half of the interval $[t_n, t_{n+1}]$ with the on-step acceleration; it is the bridge between the two formulations.

Displacement update. Substituting (A.4) into the central-difference displacement update (4.12) gives

$$\mathbf{u}_{n+1} = \mathbf{u}_n + \Delta t_{n+\frac{1}{2}} \dot{\mathbf{u}}_{n+\frac{1}{2}} = \mathbf{u}_n + \Delta t_{n+\frac{1}{2}} \dot{\mathbf{u}}_n + \frac{1}{2} \Delta t_{n+\frac{1}{2}}^2 \ddot{\mathbf{u}}_n, \quad (\text{A.5})$$

which is precisely the Newmark displacement update (4.2) evaluated at $\beta = 0$, i.e. the explicit update (4.10). The new-step acceleration $\ddot{\mathbf{u}}_{n+1}$ does not appear, confirming the explicit character of the scheme.

Velocity update. Applying the definition (A.3) at step $n + 1$ and using (A.1) recovers the on-step velocity at the new step from the same half-step velocity,

$$\dot{\mathbf{u}}_{n+1} = \dot{\mathbf{u}}_{n+\frac{1}{2}} + (t_{n+1} - t_{n+\frac{1}{2}}) \ddot{\mathbf{u}}_{n+1} = \dot{\mathbf{u}}_{n+\frac{1}{2}} + \frac{1}{2} \Delta t_{n+\frac{1}{2}} \ddot{\mathbf{u}}_{n+1}. \quad (\text{A.6})$$

Inserting the forward relation (A.4) for $\dot{\mathbf{u}}_{n+\frac{1}{2}}$ gives

$$\dot{\mathbf{u}}_{n+1} = \dot{\mathbf{u}}_n + \frac{1}{2} \Delta t_{n+\frac{1}{2}} \ddot{\mathbf{u}}_n + \frac{1}{2} \Delta t_{n+\frac{1}{2}} \ddot{\mathbf{u}}_{n+1} = \dot{\mathbf{u}}_n + \Delta t_{n+\frac{1}{2}} [(1 - \gamma) \ddot{\mathbf{u}}_n + \gamma \ddot{\mathbf{u}}_{n+1}] \quad \text{with } \gamma = \frac{1}{2}, \quad (\text{A.7})$$

which is the Newmark velocity update (4.3). The two half-intervals carry the equal weight $\frac{1}{2} \Delta t_{n+\frac{1}{2}}$ because the midpoint splits $[t_n, t_{n+1}]$ symmetrically (A.1); the central difference scheme therefore corresponds to $\gamma = \frac{1}{2}$, the value that introduces no artificial damping and makes the method second-order accurate.

Remark on the algorithm. Although $\ddot{\mathbf{u}}_{n+1}$ appears in (A.7), no coupled solve is required. The half-step velocity (A.4) and the displacement update (A.5) depend only on step- n data, so $\mathbf{u}_{n+1}$ is obtained directly; the acceleration $\ddot{\mathbf{u}}_{n+1}$ then follows from the EOM (4.14) evaluated at $n + 1$, and the on-step velocity $\dot{\mathbf{u}}_{n+1}$ in (A.7) is merely a post-processing recovery. Because $\beta = 0$, it never feeds back into the displacement, and the scheme remains explicit. Equations (A.5) and (A.7) together establish that the variable-step central difference scheme is the Newmark method with $\beta = 0$ and $\gamma = \frac{1}{2}$.

B

Material Properties


Table B.1: Intralaminar material properties and model parameters for 8552-IM7 [76]

Elasticity	E_1 171.4 GPa	$E_2 = E_3$ 9.1 GPa	$G_{12} = G_{13}$ 4.5 GPa	G_{23} 3.0 GPa	$\nu_{12} = \nu_{13}$ 0.32	ν_{23} 0.45
Strength	X_t 2500.4 MPa	X_c 1200.1 MPa	Y_t 50.8 MPa	Y_c 253.0 MPa	S_{12} 91.1 MPa	
Fracture Toughness	G_{ft} 133.3 N/mm	G_{fc} 61.0 N/mm	$G_{mt} = G_{mc}$ 0.24 N/mm	G_s 0.74 N/mm		

The interlaminar cohesive-element properties used for the ply interfaces in Models B and C are collected in Table B.2. The penalty stiffness K is the value found in Section 6.1, and the mixed-mode fracture is governed by a power-law criterion with exponent α .

Table B.2: Interlaminar material properties for the 8552-IM7 ply interface [76]

Penalty stiffness	K 10^5 MPa		
Strength	t_n^0 60.0 MPa	t_s^0 90.0 MPa	t_t^0 90.0 MPa
Fracture Toughness	G_{Ic} 0.29 N/mm	G_{IIc} 0.63 N/mm	G_{IIIc} 0.63 N/mm
Mixed-mode (power law)	α 1.0		



Modal decomposition and stable time increment derivation

The analysis follows the definition of numerical stability: two solutions of the same system with the same forcing but slightly different initial data are compared, and the method is stable if their difference does not grow.

C.1. Modal Analysis

The EOM (4.17) describes a coupled system in which every DOF interacts with all others through $\mathbf{M}_{\text{mass}}$, $\mathbf{D}_{\text{damp}}$, and $\mathbf{K}$. Studying the free, undamped behaviour first reveals the natural modes of vibration, the spatial patterns and frequencies on which the full dynamic model is built.

The Eigenvalue Problem

Free, undamped vibration ($\mathbf{f} = \mathbf{0}$, $\mathbf{D}_{\text{damp}} = \mathbf{0}$) reduces (4.17) to

$$\mathbf{M}_{\text{mass}}\ddot{\mathbf{u}} + \mathbf{K}\mathbf{u} = \mathbf{0}. \quad (\text{C.1})$$

Seeking harmonic solutions $\mathbf{u}(t) = \boldsymbol{\phi} \sin(\omega t + \psi_0)$, with $\boldsymbol{\phi} \in \mathbb{R}^{N_{\text{DOF}}}$ a time-independent shape vector, ω the angular frequency, and ψ_0 a phase angle, and substituting into (C.1) yields the generalized eigenvalue problem [4]:

$$\mathbf{K}\boldsymbol{\phi} = \omega^2 \mathbf{M}_{\text{mass}}\boldsymbol{\phi}. \quad (\text{C.2})$$

This admits N_{DOF} solution pairs $(\omega_i^2, \boldsymbol{\phi}_i)$. Because $\mathbf{M}_{\text{mass}}$ is symmetric positive-definite and $\mathbf{K}$ symmetric positive-semi-definite, all eigenvalues are real and non-negative [4, 33]. Their positive square roots give the natural frequencies $\omega_i \geq 0$, ordered $\omega_1 \leq \omega_2 \leq \dots \leq \omega_{N_{\text{DOF}}}$. The eigenvectors $\boldsymbol{\phi}_i$ are the mode shapes, the deformation patterns at which the structure vibrates freely at frequency ω_i .

Orthogonality and Mass Normalisation

The mode shapes are orthogonal with respect to both system matrices [33]:

$$\phi_i^T \mathbf{M}_{\text{mass}} \phi_j = 0, \quad \phi_i^T \mathbf{K} \phi_j = 0, \quad i \neq j. \quad (\text{C.3})$$

Since (C.2) fixes each ϕ_i only up to a scalar, we mass-normalise via $\phi_i^T \mathbf{M}_{\text{mass}} \phi_i = 1$, which with (C.3) gives [33]:

$$\phi_i^T \mathbf{M}_{\text{mass}} \phi_j = \delta_{ij}, \quad \phi_i^T \mathbf{K} \phi_j = \omega_i^2 \delta_{ij}, \quad (\text{C.4})$$

with δ_{ij} the Kronecker delta.

Modal Superposition

Collecting the mass-normalised mode shapes as columns of the modal matrix $\Phi \in \mathbb{R}^{N_{\text{DOF}} \times N_{\text{DOF}}}$, the orthogonality relations (C.4) become

$$\Phi^T \mathbf{M}_{\text{mass}} \Phi = \mathbf{I}, \quad \Phi^T \mathbf{K} \Phi = \mathbf{K}_m = \text{diag}(\omega_1^2, \dots, \omega_{N_{\text{DOF}}}^2). \quad (\text{C.5})$$

As the mode shapes span $\mathbb{R}^{N_{\text{DOF}}}$, any physical displacement is a superposition [33]:

$$\mathbf{u}(t) = \Phi \mathbf{q}(t), \quad (\text{C.6})$$

where $\mathbf{q}(t) \in \mathbb{R}^{N_{\text{DOF}}}$ collects the modal coordinates. Substituting (C.6) into (4.17), pre-multiplying by Φ^T , and using (C.5) gives

$$\ddot{\mathbf{q}} + \mathbf{D}_m \dot{\mathbf{q}} + \mathbf{K}_m \mathbf{q} = \Phi^T \mathbf{f}, \quad \mathbf{D}_m = \Phi^T \mathbf{D}_{\text{damp}} \Phi. \quad (\text{C.7})$$

C.2. Stability of the explicit integrator

Uncoupled modal equation.

In (C.7) the mass and stiffness terms are diagonal, but $\mathbf{D}_m = \Phi^T \mathbf{D}_{\text{damp}} \Phi$ is in general full. It is diagonal precisely when $\mathbf{D}_{\text{damp}}$ is a linear combination of $\mathbf{M}_{\text{mass}}$ and $\mathbf{K}$ (proportional or Rayleigh damping, $\mathbf{D}_{\text{damp}} = \alpha_R \mathbf{M}_{\text{mass}} + \beta_R \mathbf{K}$), in which case $\mathbf{D}_m = \text{diag}(2\xi_i \omega_i)$ and the system decouples completely into N_{DOF} scalar equations. Since stability can then be assessed one mode at a time, we consider a single representative mode with modal coordinate q , natural frequency ω , and critical damping ratio ξ ; its free response ($f = 0$) obeys

$$\ddot{q} + 2\xi\omega \dot{q} + \omega^2 q = 0. \quad (\text{C.8})$$

The half-step lag.

When the central difference method integrates (C.8), the velocity-dependent (damping) term lags by a half step, because evaluating it at the current half-step velocity is what keeps the update fully explicit—requiring the solution of no equations. The discrete modal equation is therefore

$$\frac{q_{n+1} - 2q_n + q_{n-1}}{\Delta t^2} + 2\xi\omega \frac{q_n - q_{n-1}}{\Delta t} + \omega^2 q_n = 0. \quad (\text{C.9})$$

The characteristic equation.

A linear difference equation has exponential solutions, so set $q_n = \mu_a^n$, the discrete analogue of the continuous $e^{\bar{\gamma}t}$ with t replaced by $n\Delta t$. The perturbation stays bounded if and only if μ_a lies within the unit circle, $|\mu_a| \leq 1$ (equality admissible only for a simple root). Substituting $q_n = \mu_a^n$, cancelling μ_a^n , and multiplying by Δt^2 gives

$$\mu_a^2 + \mu_a(g + h - 2) + (1 - g) = 0, \quad g = 2\xi\omega\Delta t, \quad h = \omega^2\Delta t^2. \quad (\text{C.10})$$

Stability conditions via the z-transform.

Rather than locate the roots in the μ_a -plane directly, the z-transform $\mu_a = (1 + z)/(1 - z)$ maps the unit-disc interior $|\mu_a| < 1$ onto the left half-plane $\text{Re}(z) < 0$. The roots of $a\mu_a^2 + b\mu_a + c = 0$ then satisfy $|\mu_a| \leq 1$ when the Hurwitz conditions

$$a - b + c \geq 0, \quad 2(a - c) \geq 0, \quad a + b + c \geq 0 \quad (\text{C.11})$$

hold. With $a = 1$, $b = g + h - 2$, $c = 1 - g$,

$$a + b + c = h = \omega^2\Delta t^2 \geq 0 \quad (\text{C.12})$$

is automatic;

$$2(a - c) = 2g = 4\xi\omega\Delta t \geq 0 \quad (\text{C.13})$$

holds for non-negative damping $\xi \geq 0$; and

$$a - b + c = 4 - 2g - h = 4 - 4\xi\omega\Delta t - \omega^2\Delta t^2 \geq 0 \quad (\text{C.14})$$

is the binding condition on the time step.

Solving for the critical step.

Condition (C.14) is a quadratic in $\omega\Delta t$; at equality,

$$\omega\Delta t = -2\xi \pm 2\sqrt{\xi^2 + 1}. \quad (\text{C.15})$$

The negative root gives a negative step and is discarded, so the positive root yields

$$\Delta t_{\text{crit}} = \frac{2}{\omega} \left(\sqrt{\xi^2 + 1} - \xi \right), \quad (\text{C.16})$$

with the undamped case recovered at $\xi = 0$:

$$\Delta t_{\text{crit}} = \frac{2}{\omega}. \quad (\text{C.17})$$

For a multi-mode system the most restrictive mode governs, so $\omega = \omega_{\text{max}}$, the maximum frequency of the linearized system. Any non-negative damping thus reduces the stable step relative to the undamped bound: the half-step lag in the velocity makes damping tighten rather than relax the stability limit, and

this holds both for explicit damping such as $\mathbf{D}_{\text{damp}}^v \mathbf{v}$ and for damping arising within a material law.