CHALMERS
UNIVERSITY OF TECHNOLOGY

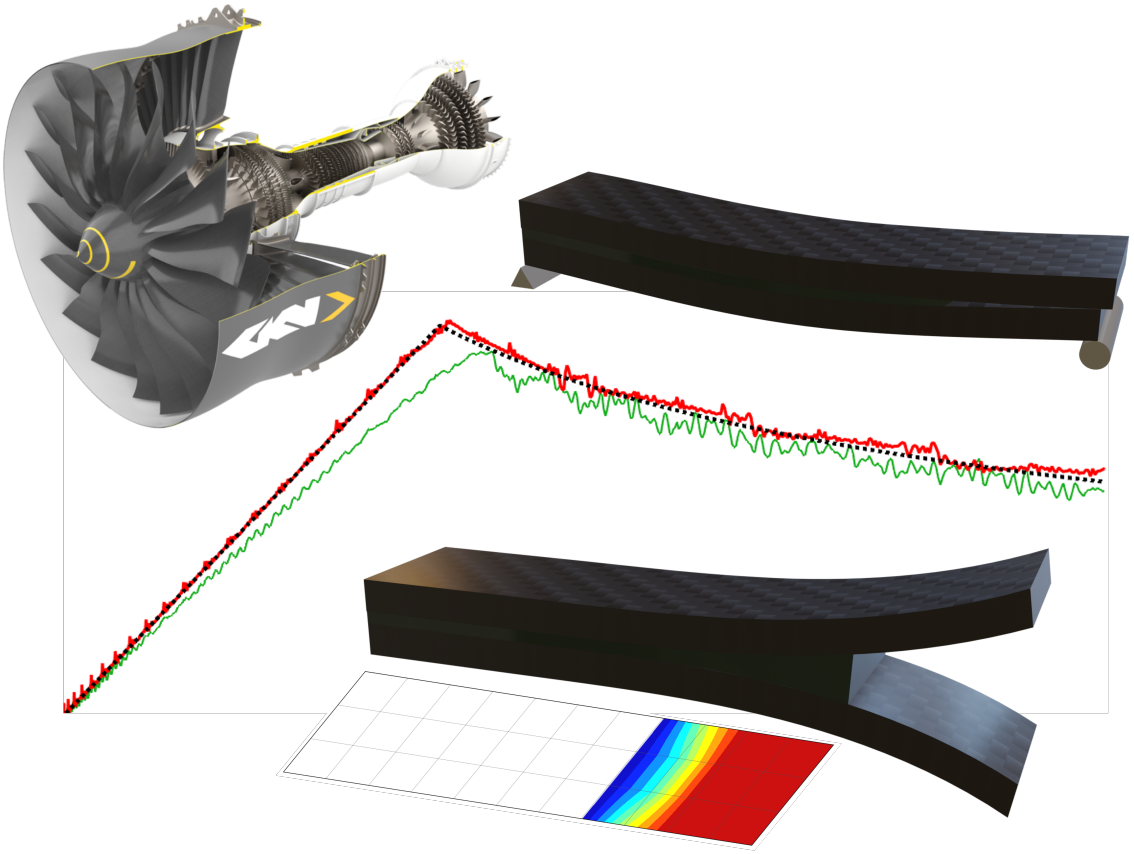

Efficient Simulation of Crack Propagation in Adhesive Bonds

Sustainable Lightweight Structures

Master's Thesis in Applied Mechanics

WALTHER AHL & DANIEL JINNERYD

DEPARTMENT OF INDUSTRIAL AND MATERIALS SCIENCE
CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden 2026
www.chalmers.se

Master's thesis 2026

Efficient Simulation of Crack Propagation in Adhesive Bonds

Sustainable Lightweight Structures

Walther Ahl & Daniel Jinneryd

CHALMERS
UNIVERSITY OF TECHNOLOGY

Department of Industrial and Materials Science
Chalmers University of Technology
Gothenburg, Sweden 2026

Efficient Simulation of Crack Propagation in Adhesive Bonds
Sustainable Lightweight Structures
WALTHER AHL, DANIEL JINNERYD

Supervisor: Martin Fagerström - Chalmers, Carolyn Oddy - GKN Aerospace Sweden AB & Torkel Davidson - Uniso Technologies AB
Examiner: Martin Fagerström - Chalmers Department of Industrial and Materials Science

Master's Thesis 2026
Department of Industrial and Materials Science
Chalmers University of Technology
SE-412 96 Gothenburg
Telephone +46 31 772 1000

Cover: A double cantilever beam with interface damage and force–displacement curve shown, alongside a GKN engine and a end-notched flexure test.

Typeset in LaTeX
Gothenburg, Sweden 2026

Abstract

Accurate yet computationally affordable modelling of interface crack propagation is required for the development and certification of large bonded assemblies. Conventional cohesive zone modelling (CZM) resolves the fracture process zone (FPZ) with traction–separation laws that typically require element edges below 1 mm. This size limitation makes it impractical at industrial relevant scales.
This thesis investigates an energy-release-rate cohesive (ERRC) approach for large-element modelling of adhesive-interface crack propagation. The method combines a virtual crack closure technique (VCCT) propagation criterion with a local cohesive release law, so that crack growth is triggered by the energy release rate while the newly created crack surfaces dissipates the prescribed fracture energy progressively. The main contribution of the work is the reformulation of this approach as a user-defined solid element in LS-DYNA, intended for adhesive interfaces discretised with solid adherends and finite-thickness bondline representation.
The implemented interface is represented by lower–upper nodal pairs, where each pair carries as discrete state: tied, cohesive, or open. In this thesis, two implementation variants were developed: a single-core reference (SCR) implementation and a multi-core capable (MCC) implementation.
For the DCB benchmark, the implemented method reproduced the expected mesh-aligned Mode I crack-growth behaviour. The crack-length evolution followed the corrected beam theory (CBT) reference within the resolution of one discrete interface-element edge, and the interface-energy diagnostics showed that the released pairs followed the intended fracture-energy target. Representative Mode I propagation was obtained using an in-plane interface element length of 4 mm, substantially larger than the sub-millimetre element sizes typically required to resolve a conventional FPZ.
The ENF benchmark confirmed that the discrete nodal-pair formulation can produce a coherent Mode II-dominated propagation chain. However, the global force–displacement response showed pronounced dynamic oscillations, particularly for the MCC implementation. The dissipation diagnostics showed that the implemented damage variable limits the release state, but that the reconstructed path-work quantities can become unreliable during rapid, single-cycle Mode II release events. The ENF results therefore verify important parts of the formulation but also show that further stabilisation and quasi-static assessment are required before the method can be considered robust for Mode II-dominated loading.
Overall, the work demonstrated the feasibility of extending the ERRC method to solid-element adhesive-interface formulation in LS-DYNA. The implementation provides a promising route for efficient large-element simulation of interfacial crack propagation in bonded structures. At the present stage, the method should be interpreted as a proof of concept for regular, mesh aligned benchmark problems rather than a fully general industrial fracture-modelling tool.

Keywords: adhesive, delamination, virtual crack closure technique, cohesive-zone modelling, energy release rate, LS-DYNA, user-defined element, large elements, composite materials.

Acknowledgements

We would first like to express our sincere gratitude to Carolyn Oddy for her guidance, support, and encouragement throughout this thesis. Her dedication, knowledge, and ability to motivate those around her have been truly inspiring. She has consistently pushed us forward and provided valuable insight throughout the work. We are grateful for the time, effort, and commitment she has invested in us. We have great respect for her dedication and the way she leads by example, working harder than anyone else.
We would also like to sincerely thank Martin Fagerström, our examiner. His expertise, sharpness, and commitment to excellence have made working with him both challenging and highly rewarding. He has a remarkable ability to bring out the best in people, and it has been a pleasure to learn from him throughout our master's studies. Beyond being an exceptional teacher, he has made the journey even more enjoyable with his sense of humour. Although he is usually quite stoic during the meetings, he is exceptionally funny when he lets go, his true colours shine bright in the Teams chat.
A heartfelt thank you goes to Torkel and Kajsa for their modelling support.
The wonderful Johannes Främby has provided an amazing backbone for our implementation capabilities.
We gratefully acknowledge Engineers of Sweden for awarding this thesis a scholarship from the Environmental Fund.

Daniel
My utmost gratitude goes to my friends and family, the Jinneryds, for their endless support, joy and shaping me into the person I am today.
I would personally like to thank Walther Ahl, my thesis partner, for the many great years, the strong camaraderie, and the motivation he has given me throughout the entire master's programme.

Walther
To my beloved daughter Elise, you are my everything. To my wonderful fiancé Elin, thank you for all the support during these five years, I love you. To my fantastic grandmother Kathe, who sadly passed away during the thesis, you will be missed. Thank you to the rest of my family for all your support, even though you never ever really understood what the hell I was studying. To Daniel Jinneryd, thank you for five outstanding years, a fantastic thesis period and I look forward to continuing our friendship afterwards.

Walther Ahl & Daniel Jinneryd
Gothenburg, May 2026

Acronyms

The following abbreviations are used throughout this thesis. Each one is also defined on first use in the running text.

BC	Boundary Condition
BK	Benzeggagh–Kenane (Mode-Mixity Criterion)
CBT	Corrected Beam Theory
CFRP	Carbon Fibre Reinforced Polymer
CPU	Central Processing Unit
CZM	Cohesive Zone Modelling
DCB	Double Cantilever Beam (Mode I Delamination Test)
DOF	Degree Of Freedom
ELFORM	Element Formulation Flag (LS-DYNA `*SECTION_SOLID`)
ENF	End-Notched Flexure (Mode II Delamination Test)
EPFM	Elastic–Plastic Fracture Mechanics
ERR	Energy Release Rate
ERRC	Energy-Release-Rate–Cohesive (the method implemented in this work)
FE	Finite Element
FEA	Finite Element Analysis
FEM	Finite Element Method
FPZ	Fracture Process Zone
GKN	Guest, Keen and Nettlefolds (GKN) Aerospace (industrial partner)
HSV	History Variable (LS-DYNA Per-Element Scratchpad)
LEFM	Linear Elastic Fracture Mechanics
LS-DYNA	Commercial Explicit / Implicit FEA Solver by ANSYS
MCC	Multi-Core Capable (Implementation in This Work)
MPI	Message Passing Interface
MPP	Massively Parallel Processing (LS-DYNA Distributed-Memory Mode)
NHSV	Number of History Variables (LS-DYNA Section Parameter)
NIP	Number of Integration Points (LS-DYNA Section Flag)

SCR	Single-Core Reference (Reference Implementation in This Work)
SMP	Shared-Memory Parallelism (LS-DYNA Multi-Thread Mode)
UEL	User-defined Element (LS-DYNA)
UMAT	User-defined MATerial (LS-DYNA)
VCCT	Virtual Crack Closure Technique

Glossary

The following terms are used throughout this thesis. Definitions have been adjusted for reader comprehension and are not intended to replace the formal definitions in the cited literature.

1. **ANSYS**: Finite element analysis software used for engineering simulations; owner of LS-DYNA. ANSYS is owned by Synopsys.
2. **Benzeggagh–Kenane (BK) Criterion**: Empirical mixed-mode fracture criterion that interpolates between Mode I and Mode II toughness as a power law of the local mode-mixity ratio.
3. **Boundary Conditions (BCs)**: Conditions applied to the boundaries of a finite element model to simulate the environment or constraints.
4. **Bulk-Solid Pass**: First parallel phase of an LS-DYNA cycle, in which conventional solid elements assemble their internal forces into the global force array `dm_a`.
5. **Carbon Fibre Reinforced Polymers (CFRP)**: Materials made by combining carbon fibre reinforcements with a polymer matrix.
6. **Cohesive Law**: Traction–separation relation that prescribes the force carried by a debonding interface as a function of opening displacement; typically bilinear with an initial elastic branch and a softening branch terminating at a critical opening.
7. **Cohesive-Zone**: Region ahead of the crack tip in which the interface is still partially intact and a non-zero traction is sustained while opening; in this work the cohesive-zone is collapsed to a single node pair at a time.
8. **Contact Penalty**: Stiff elastic spring activated when an open pair experiences compressive opening, used to prevent unphysical inter-penetration of the released crack faces.
9. **Convergence Study**: Analysis performed to ensure that the numerical solution obtained from a simulation is stable and accurate.
10. **Crack-Extension Area** A_i: Direction-aware area assigned to a tied front pair, defined as the interface surface that will be opened if that pair is released next.
11. **Crack Front**: Location of tied pairs adjacent to at least one open pair; the boundary of the currently delaminated region.
12. **Cracklook**: Non-local synchronisation routine called once per LS-DYNA cycle that resolves nodal-pair ownership across element copies and stages cross-element data into history variables. Registered through `*USER_NONLOCAL_SEARCH`.

13. **Cycle / Time Step**: One increment of the LS-DYNA explicit central-difference integrator.
14. **Damage Variable** D: Scalar in $[0, 1]$ that monotonically degrades the cohesive force. A value of D=0 signifies a fully intact interface while D=1 signifies a fully damaged interface.
15. **Delamination**: Inter-laminar crack propagation between two plies of a laminated composite or in an adhesive bond between two substrates.
16. **Global Force / Nodal-Mass Arrays** : `dm_a` / `dm_xms`. The tied-pair force is reconstructed from `dm_a` and `dm_xms` at the eight UEL nodes after the bulk-solid pass has assembled its contributions.
17. **Energy Release Rate (ERR)**: Rate of mechanical energy released per unit area of new crack surface; denoted G and split by mode into G_I, G_{II}, G_{III}.
18. **ERRC Method**: Hybrid delamination-modelling strategy of Daniel *et al.* [1, 2] that uses VCCT as the propagation criterion and a single-pair cohesive law to dissipate the correct fracture energy as the crack moves through one large element at a time.
19. **Finite Element (FE) Analysis**: Computational method for approximating the behaviour of structures under loading.
20. **Fracture Process Zone (FPZ)**: Region around the crack tip where the actual material degradation takes place; in conventional CZM the FPZ must be discretised by several elements.
21. **Fracture Toughness**: Critical energy release rate at which a crack propagates; mode-dependent (G_{Ic}, G_{IIc}, G_{IIIc}).
22. **Front Pair**: Tied nodal pair located on the current crack front; eligible for VCCT evaluation.
23. **History Variable (HSV)**: Per-element scratchpad LS-DYNA preserves between cycles. The element routine reads only its own slots; cracklook stages cross-element data into them.
24. **Isotropic**: Material with uniform properties that is independent of loading mode/direction.
25. **Linear Elasticity**: Behaviour of materials that return to their original shape after deformation when subjected to loads; based on a linear stress-strain relationship.
26. **LS-DYNA**: Commercial explicit / implicit FEA solver owned by ANSYS used to host the user-defined element implemented in this work.
27. **Mode I / II / III**: Fracture modes corresponding respectively to opening (normal), in-plane shearing, and out-of-plane (tearing) shearing of the crack faces.
28. **Message Passing Interface (MPI)**: A message-passing standard used in parallel computing.
29. **Massively Parallel Processing (MPP)**: LS-DYNA execution mode in which the model is split across MPI ranks.
30. **Multi-Core**: Several SMP threads on the cores of a single CPU or workstation.
31. **Nodal Ownership**: Rule by which only one element of a shared nodal pair is allowed to assemble interface forces in a given cycle.

32. **Tied State** ($s_p = 0$): Initial pair state in which the lower and upper nodes are constrained to move together as a single physical point.
33. **Cohesive State** ($s_p = 1$): Pair state during which the cohesive law degrades the inherited release force toward zero according to the damage variable D.
34. **Open State** ($s_p = 2$): Pair state in which the interface no longer transmits tensile or shear traction; only a compressive contact penalty remains.
35. **Orthotropic**: Material with three mutually perpendicular planes of symmetry, each having distinct mechanical properties.
36. **Outlet Guide Vanes (OGV)**: Components in a turbine engine that provide structural support and guide airflow to improve engine efficiency.
37. **Pair Key $\mathbf{k}_p$**: Pair of sorted global node IDs ($\min(n_p^-, n_p^+)$, $\max(n_p^-, n_p^+)$) used as a unique identifier for the physical interface point shared by several element copies.
38. **Predecessor Pair**: Open neighbouring pair behind the current front pair in the global propagation direction; used to define the local crack-extension direction $\mathbf{d}_p$.
39. **Rank**: One MPP process; owns a part of the model in distributed memory.
40. **Representative Volume Element (RVE)**: Small sample that represents the microstructure of a material.
41. **Shared-Memory Parallelism (SMP)**: LS-DYNA execution mode in which multiple threads run on a single machine and share one memory image; cross-thread data is visible without explicit messaging.
42. **Specific Stiffness**: Stiffness of a material relative to its density.
43. **Transversely Isotropic**: Material with physical properties that are symmetric about an axis that is normal to a plane of isotropy. Five elastic properties are required to characterise it.
44. **User-Element Pass**: Second parallel phase of an LS-DYNA cycle in which user-defined elements (`ELFORM`=101–105) are executed; reads `dm_a` after the bulk-solid pass has finished.
45. **`*USER_NONLOCAL_SEARCH`**: LS-DYNA keyword that registers a per-cycle non-local synchronisation routine; used here to dispatch the cracklook step.
46. **Virtual Crack Closure Technique (VCCT)**: Linear-elastic-fracture-mechanics estimator of the energy release rate from the work needed to close the most recently opened crack increment, evaluated from nodal forces and openings without remeshing.

Contents

List of Figures

List of Tables

1
Introduction

1.1 Background

This chapter motivates the work, reviews the modelling approaches that frame it, and links the effort to its broader sustainability context.

1.1.1 Motivation

The increasing use of composite materials in aerospace structures is driven by the need to reduce structural mass while maintaining stiffness, strength, and damage tolerance. In aircraft and aero-engine applications, weight reduction is directly connected to fuel consumption, emissions, and operating cost. However, the full benefit of composite-intensive structures depends on reliable design, manufacturing, inspection, and simulation methods.

Adhesive bonding is one route toward more weight-efficient composite and multi-material assemblies. Compared with mechanical fastening, bonded joints transfer load over a finite area rather than through discrete fasteners, which can reduce local stress concentrations and avoid damage associated with drilled holes. At the same time, bonded interfaces may fail through interfacial crack-growth mechanisms such as adhesive debonding. The same mechanisms are also relevant to interlaminar delamination in laminated composites, where the fracture-mechanics framework developed in this thesis applies in much the same way. The focus of this work, however, is the interfacial fracture of adhesive bondlines.

For lightweight bonded structures to be introduced with confidence, efficient and reliable numerical methods for interfacial crack propagation are required. This forms part of the motivation for the present work.

1.1.2 Sustainable Lightweight Structures

Reducing the environmental impact of aviation requires a combination of operational, propulsion-related, and structural measures. From a structural perspective, lightweight design remains an important route to lower fuel consumption and associated emissions. According to the International Energy Agency [3], aviation accounted for 2.5 % of global energy-related CO_2 emissions in 2023, motivating continued development of lighter and more efficient aircraft and engine structures.

Carbon-Fibre-Reinforced Polymers (CFRPs) are attractive in this context because they offer high specific stiffness and strength compared with conventional metallic materials. Hagnell and Åkermo [4] and Timmis *et al.* [5] reported that

CFRP structures in aerospace applications can contribute to lifetime CO_2 reductions on the order of 14–20 %, depending on the application and the assumptions used in the life-cycle assessment. These benefits, however, do not come for free. Composite structures are generally more expensive to manufacture, inspect, and certify, and their design requires careful consideration of material anisotropy, damage mechanisms, and process variability. Consequently, lightweight composite design is not only a material-substitution problem. It also requires improved methods for manufacturing, analysis, and verification. This is particularly important for bonded and laminated structures, where damage mechanisms such as delamination and interfacial crack propagation may govern the structural integrity of a component. Reliable predictive methods for interfacial fracture can therefore act as enabling tools for less conservative and more weight-efficient designs.

Aero-engine structures provide a relevant industrial example. Different regions of an engine are exposed to different thermal, mechanical, and environmental requirements; high-temperature regions may require metallic alloys or ceramic matrix composites, whereas colder structural regions offer greater potential for CFRP or hybrid composite–metal solutions. In such applications, replacing metallic concepts with composite or bonded alternatives may reduce mass, but only if the corresponding failure mechanisms can be predicted with sufficient confidence. Fig. 1.1 is a schematic of a representative aero-engine showing the variety of structural regions relevant to lightweight design.

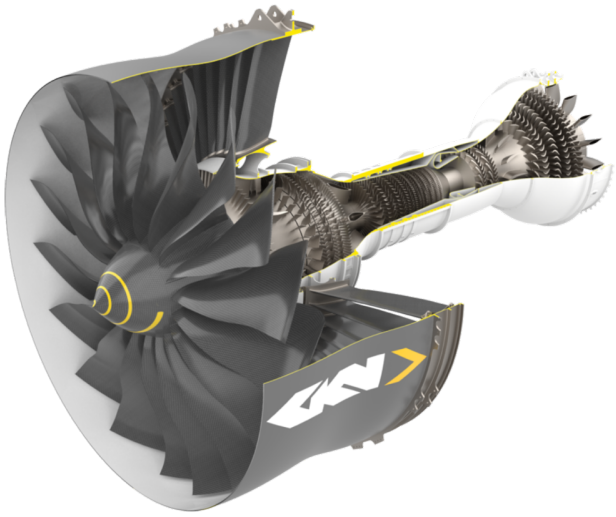

Figure 1.1: Schematic aero-engine structural context illustrating the variety of components and structural regions relevant to lightweight design. Image provided courtesy of GKN Aerospace Sweden AB.

A conservative internal estimate carried out by GKN Aerospace indicates that, for a representative engine, replacing fastener-related mass with bonded solutions yields a mass reduction on the order of 5 kg per engine. For a twin-engined aircraft, this is estimated to reduce emissions by approximately 70 tonnes of CO_2 equivalents. Although indicative, this estimate becomes relevant at engine and fleet level when combined with the broader design freedom enabled by bonded composite structures.

This has the opportunity to translate into reduced material use during manufacturing and reduced fuel burn over an engine's operational life.

1.1.3 Adhesive Bonds

Adhesive bonding is a joining technology in which loads are transferred through an adhesive layer over a finite area rather than through discrete mechanical fasteners. This distributed load transfer can reduce local stress concentrations and enable weight-efficient joints, particularly for thin-walled structures and assemblies with limited access for fastening. In aerospace, adhesive bonding is also closely tied to manufacturing discipline. Surface preparation, process control, inspection capability, and the definition of allowables are treated as integral parts of bonded-joint integrity and qualification practice [6].

From a structural-mechanics perspective, bonded joints are governed by the interaction between joint geometry, adherend stiffness, adhesive properties, and local stress states. Crack initiation and propagation can occur within the adhesive layer, along an interface, or within the adherends; these mechanisms are not mutually exclusive and may compete depending on the material combination, manufacturing quality, environmental exposure, and loading mode. The constitutive law used at the interface must capture mixed-mode crack growth and the verification cases must isolate Mode I and Mode II contributions cleanly. A detailed treatment of crack propagation in adhesively bonded joints using cohesive modelling is given by Noorman [7].

Common bonded joint archetypes include single-lap joints, double-lap joints, strapped joints, stepped joints and doublers; examples of common bonded joint types are shown in Fig. 1.2. The choice of joint type is not only geometric, it is governed by the load-transfer path, the sensitivity to eccentricity and bending, and the relative importance of shear and peel stresses in the adhesive layer.

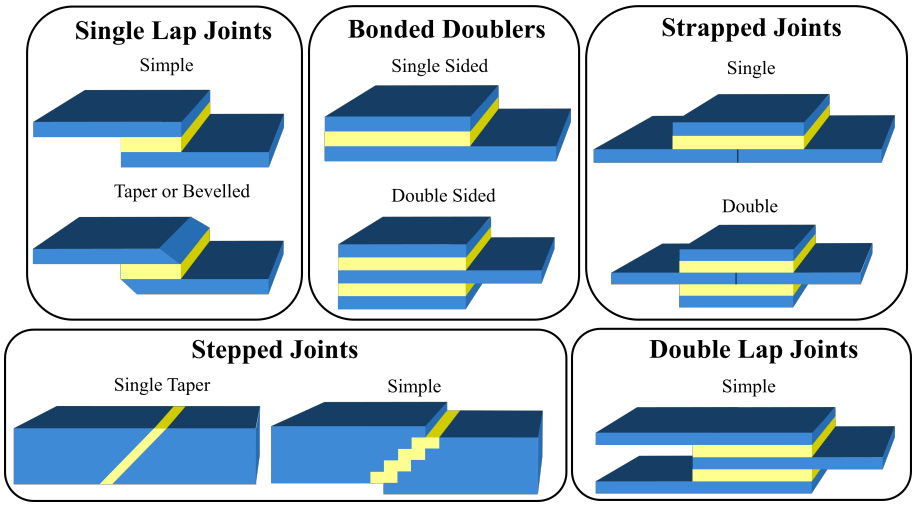

Figure 1.2: Examples of common adhesively bonded joint configurations.

Even when the global loading appears simple, bonded joints frequently develop coupled shear and peel stresses due to eccentricity, local bending, and stiffness mis-

match between adherends. One such example of eccentricity is shown in the fractured lap shear test in Fig. 1.3 corresponding to the simple single lap joint in Fig. 1.2. As a result, crack driving forces at the bond line are commonly mixed-mode. Failure may occur as cohesive failure within the adhesive, adhesive failure along an interface, or adherend-dominated damage, as illustrated schematically in Fig. 1.4 according to ASTM D5573 [8]. For dissimilar adherends, such as metal–composite joints, stiffness mismatch and thermal or environmental effects may further increase the sensitivity of the interface. This is one reason why process control, inspection, and qualification procedures are central in bonded-joint practice [6].

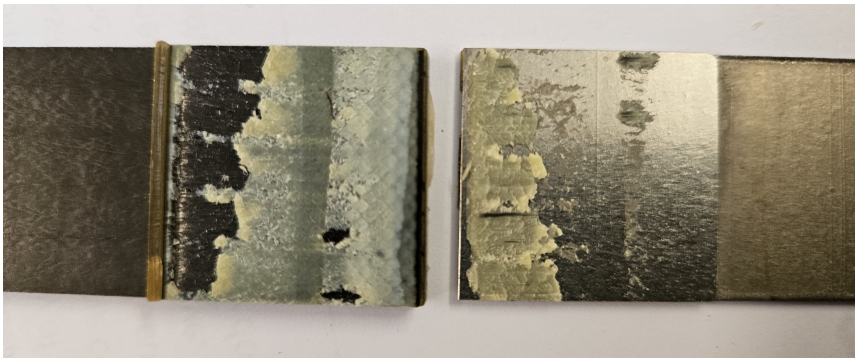

Figure 1.3: Adhesive debonding of a failed lap shear test. Image provided courtesy of GKN Aerospace Sweden AB.

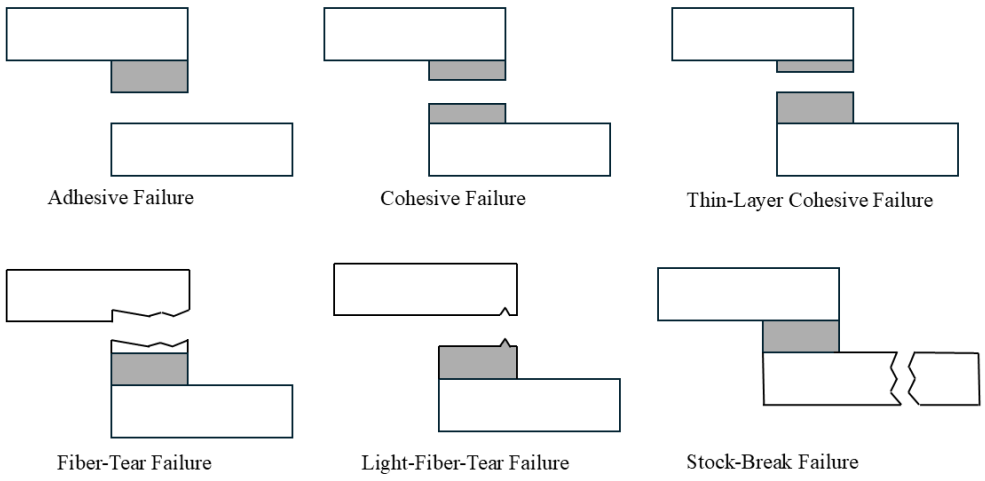

Figure 1.4: Typical failure modes in bonded joints according to ASTM D5573 [8].

The mixed-mode nature of bonded-joint failure motivates the use of standardised coupon configurations as building blocks for modelling and verification. The double cantilever beam (DCB), end-notched flexure (ENF), and mixed-mode bending

(MMB) specimens provide controlled Mode I, Mode II and mixed-mode loading conditions. These are widely used as reference cases for evaluating numerical methods for interfacial crack propagation, Kruger [9] provides one such reference.

While adhesive bonding is often discussed in the context of joining separate components, interfacial failure is equally important within laminated composites. In this case, delamination represents separation between neighbouring plies. Delamination is promoted by the relatively weak interlaminar strength and fracture resistance compared with the in-plane properties of fibre-reinforced laminates. It can also interact with bonded-joint failure in composite adherends, particularly near bond-line edges or regions with high peel and shear stresses. A general composites-mechanics treatment of interlaminar stresses, delamination mechanisms, and laminate-level damage is provided by Agarwal *et al.* [10].

From a modelling viewpoint, adhesive debonding and interlaminar delamination can both be treated as interfacial fracture problems. In both cases, crack growth occurs along an interface and is governed by mixed-mode fracture resistance and energy dissipation. This common interpretation motivates the use of fracture-mechanics-based and cohesive-zone-based methods for both bonded joints and laminated composite interfaces. However, it is once again highlighted here that this work targets interfacial fracture of an adhesive bondline of finite thickness rather than delamination between plies.

1.1.4 Interfacial Crack Propagation Approaches

Progressive interfacial failure can be modelled using several classes of approach, each with distinct trade-offs between physical fidelity, robustness, and computational cost. Following the overview by Noorman [7], the available finite-element approaches for crack-growth analysis may be grouped into four broad families: continuum-mechanics criteria, fracture-mechanics criteria, enriched finite-element methods such as XFEM, and damage-mechanics approaches. The categories differ mainly in (i) how the crack-driving force is represented, (ii) whether a crack path must be predefined, and (iii) how progressively evolving separation is handled in a finite element setting.

In continuum-mechanics approaches, failure is assessed by comparing stresses or strains with critical values. The criteria may include maximum principal stress or strain, shear-stress criteria, or combined criteria. These models are attractive because they are simple and require few additional parameters. In particular they are practical for brittle materials and early-stage screening of failure onset. For cracked geometries, however, stresses near crack tips are affected by singularities and mesh dependence, so peak values become ill-defined and highly sensitive to element size [7].

Fracture mechanics was motivated by the need to address cracks and stress singularities. It uses crack-tip parameters instead of local peak stresses. Linear elastic fracture mechanics (LEFM) relates crack growth to quantities such as the stress intensity factor or the strain energy release rate, while elastic–plastic fracture mechanics (EPFM) extends the framework to ductile behaviour through the J-integral. In finite element analyses, fracture-mechanics-based propagation is often realised

through energy-release-rate criteria. The virtual crack closure technique (VCCT) is a canonical example in which nodal forces and relative displacements around the crack tip are used to compute mode-partitioned energy release rates. The propagation criterion is tied directly to fracture mechanics and the approach can provide high accuracy when the crack front is well defined, but it typically requires an initial crack and additional crack-front bookkeeping for robust progressive growth [9, 7]. However, it should be noted that instant nodal release in an explicit time-integration scheme leads to unstable crack propagation, as discussed by Daniel *et al.* [1].

Extended-finite-element formulations (XFEM) enrich the displacement approximation locally so that discontinuities can be represented without re-meshing. They are conceptually attractive for arbitrary crack growth in a fixed mesh, but are not used in the present work and are mentioned only for completeness.

Damage-mechanics models represent degradation through internal damage variables that progressively reduce stiffness from initiation to failure. Cohesive zone modelling (CZM) can be interpreted as a damage-mechanics approach in which separation across the interface follows a traction–separation law whose area equals the fracture energy. CZM provides a natural framework for progressive debonding and energy dissipation, but accurate results require the fracture process zone (FPZ) to be sufficiently resolved, making predictions sensitive to element size along the interface [7]. Turon *et al.* [11] proposed an engineering strategy that selects cohesive parameters as a function of element size while preserving the target fracture energy, which reduces the dependence on the FPZ resolution. Even with such adjustments, the FPZ along long interfaces remains a computational bottleneck for large industrial structures, especially when several load cases, mesh variants, or optimisation loops are required.

For large structural models, the cost of resolving the FPZ along long interfaces motivates coarse-mesh crack-growth approaches. Daniel *et al.* [1] introduced an efficient energy-release-rate–cohesive (ERRC) framework that uses a VCCT-based propagation criterion together with a nodal cohesive release mechanism. The cohesive release mechanism progressively opens the interface and dissipates the appropriate fracture energy with elements substantially larger than the FPZ. The method was further extended by Daniel *et al.* [2] to handle arbitrarily shaped delamination fronts on large and distorted elements, addressing robustness issues that arise in realistic industrial meshes. These developments define the technical objective of the present thesis: to specialise and implement the ERRC method as an LS-DYNA user-defined element (UEL) for adhesive bondlines. The work targets LS-DYNA explicit, an industrial finite element solver routinely used to simulate the response of structures under fast and severe loading [12], through its UEL interface [13]. The implementation details, including how data are shared between element copies in single- and multi-core execution, are presented in Chap. 3.

1.1.5 Sustainable, Ethical and Environmental Aspects

Sustainability aspects of adhesive bonding extend over the full life cycle of a structure, from manufacturing through operation, maintenance, and end-of-life. In aerospace applications, the operational phase is typically the dominant contributor to life-cycle

impact, so any structural choice that supports lower mass during service translates into reduced fuel burn and emissions. Adhesive bonding contributes indirectly to that goal by enabling lightweight, load-path-efficient, multi-material structures without introducing fastener holes or discrete clamping loads. The benefit is conditional, however, on the bonded structure remaining durable, inspectable, and repairable throughout its service life [14].

The relevance to end-of-life is more nuanced. Durable bonds complicate disassembly and material recovery in multi-material assemblies, and unless separation is designed in from the start, the same durability that is beneficial in service can become a barrier to recycling. Concepts such as controlled debonding aim to support disassembly, rebuilding, and repair without compromising in-service performance. Such concepts are increasingly considered part of the sustainability case for bonded structures [15]. Repairability sits alongside this: bonded composite patches can provide efficient load transfer with less damage to the parent structure than mechanically fastened repairs. Baker [16] showed that bonded fibre-composite patches can be used to repair fatigue-cracked aircraft structures, extending the useful life of a component. Industrial relevance is reflected in aerospace MRO practice, where repair solutions are developed for components such as fan blades, structural cases, and engine structures [17].

Following Ravichandran and Balasubramanian [14] and the IVK circular-economy review [15], design for sustainability in bonded structures can be interpreted through four practical requirements:

- **Design for durability:** reduce peel stresses, local stress concentrations, and environmentally sensitive failure mechanisms to avoid premature debonding.
- **Design for inspectability and repair:** ensure that joint geometry and access allow reliable non-destructive inspection and feasible bonded-repair procedures.
- **Design for process robustness:** select adhesive systems, surface treatments, and manufacturing procedures that are tolerant to realistic variability and reduce scrap or rework.
- **Design for end-of-life:** where feasible, choose joint concepts or adhesive technologies that support controlled separation, rework, or material recovery.

The environmental motivation for lightweight aerospace structures also raises a broader systemic question. Reducing structural mass can lower fuel burn and CO_2 emissions per flight, but per-unit efficiency does not automatically translate into reductions at sector level. If lower operating costs lead to lower fares or increased demand, part of the efficiency gain may be offset by increased air traffic. This phenomenon is commonly discussed as rebound and, in stronger cases, as backfire or the Jevons paradox [18]. Alcott [19] distinguishes between rebound, where efficiency gains are partly offset by increased demand, and backfire, where total consumption increases beyond the initial savings, arguing that efficiency improvements alone cannot guarantee absolute reductions without additional constraints or policy measures. The contribution of the present thesis should therefore be understood as an enabling technology for resource-efficient engineering, not as a complete solution to the environmental impact of aviation.

A separate ethical consideration concerns engineering safety. A user-defined el-

ement (UEL) introduces implementation risk compared with standard, extensively tested solver features. Incorrect assumptions, insufficient verification, or black-box use could lead to misleading predictions of crack propagation and structural integrity. The work therefore emphasises transparent documentation of assumptions, parameter choices, verification cases, and validity limits, including the LEFM basis of the propagation criterion and the need for an initial crack. The implementation is benchmarked against conventional CZM models and analytical reference solutions. The accompanying modelling guidance states recommended use cases, limitations, and known failure modes, so that the method can be applied critically rather than adopted as a black-box tool.

In the context of this work, the sustainability argument is directly connected to modelling. Reliable prediction of interfacial crack propagation supports durable and damage-tolerant bonded structures. It can also reduce overly conservative design margins and repeated test–simulation iterations during development, which supports more efficient use of both material and engineering resources.

1.2 Project Scope, Goals, and Limitations

The purpose of this project is to develop a computationally efficient method to analyse progressive failure along bonded interfaces by specialising the ERRC approach proposed by Daniel *et al.* [1, 2] into an LS-DYNA UEL framework.

1.2.1 Goals

The specific goals of the project are to:

- Implement an LS-DYNA UEL that combines a VCCT-based crack initiation and propagation criterion with a node-based cohesive law to dissipate the correct fracture energy during release, in line with the formulation of Daniel *et al.* [1, 2].
- Ensure robustness on regular meshes across a range of element sizes, so that the method remains usable with elements significantly larger than those required by conventional CZM.
- Verify and benchmark the UEL against established CZM approaches in LS-DYNA on standard fracture coupons (DCB and ENF), including load–displacement response and crack-length evolution.
- Assess computational efficiency (runtime, stable time step, and mesh requirements) relative to traditional CZM on the same fracture coupons.

1.2.2 Limitations

The scope of the work is constrained by the following limitations:

- The method is implemented and verified within the LS-DYNA explicit framework only.
- Linear elastic fracture mechanics (LEFM) is assumed applicable, implying a sufficiently small fracture process zone relative to the structural length scales and a known initial crack.

- Only two verification load cases are considered (DCB and ENF), driven by the available project time (19 weeks).
- The UEL is implemented as a solid (eight-node) interface element. It supports regular meshes with a predefined crack interface, and acts as a proof of concept for further integration with the solver supplier.
- The arbitrary-front capabilities of Daniel *et al.* [2] are not fully included. In particular, the present work does not include the distorted-element correction factors or a general crack-front tracking procedure for arbitrary meshes.
- The method targets bonded interfaces whose bondline thickness is large enough for the underlying tied contact to remain numerically well-conditioned. Very thin bondlines, below 0.001 mm fall outside the validated range.
- The method is implemented using the MPP execution of LS-DYNA, as this was the available solver version with user library support at the start of the project. Consequently, SMP-based shared-memory parallelisation could not be investigated for the proposed implementation, although the implementation structure was designed to allow future multi-core capabilities.

1.2.3 Acknowledgements

The work was conducted as a collaboration between GKN Aerospace, Uniso Technologies, Chalmers University of Technology, and DYNAmore/ANSYS. GKN provided industrial cases and geometries; Uniso contributed LS-DYNA modelling and implementation expertise; Chalmers provided theoretical support and supervision; and DYNAmore/ANSYS supported the UEL implementation within the LS-DYNA user-defined interfaces.

Artificial intelligence has been used to enhance the language and achieve a consistent writing style. Specifically, DeepL [20] has been used for academic-writing assistance, and ChatGPT [21] has been used for text suggestions and error corrections.

2

Theory: Interfacial Fracture Mechanics

The theoretical material in this chapter has three layers. The first is the energetic description of interfacial fracture, with stored energy at the crack front and the mixed-mode interaction deciding when a crack should propagate. The second is the analytical specimen solutions for the standard double cantilever beam and end-notched flexure tests, used as closed-form references. The third is the numerical algorithm that evaluates the ERR and dissipates the stored energy, forming the theoretical basis for the implemented user-defined element in this work.

2.1 Interfacial Fracture and Energy Release Rate

Delamination in adhesive joints is governed by crack propagation along material interfaces. In many structural configurations, such cracks are constrained to pre-defined paths such as adhesive bondlines. Further, the theoretical framework often used to predict crack formation and propagation is linear elastic fracture mechanics (LEFM). Within the framework of LEFM, the bulk material response is assumed linear elastic. Further, it is assumed that all nonlinear processes are confined to a very small region near the crack tip. Under these assumptions, crack growth is characterised by stored elastic energy close to the crack tip, rather than local stress maxima. The central measure of crack driving force is the strain energy release rate, which represents the mechanical energy stored for an incremental increase in crack area [22, 23].

A schematic illustration of the three fundamental fracture modes is shown in Fig. 2.1. These modes form the basis for the energetic mode decomposition discussed later in this section [22, 24]. Note however, that in practice in finite element models, the distinction between Mode II and Mode III is not distinct.

2.1.1 Energy Balance and Griffith-Type Criterion

The energetic description of fracture originates from Griffith's concept that crack growth is governed by changes in the total potential energy of the system. For a cracked body subjected to external loading, the total potential energy, Π, may according to Janssen *et al.* [22] be written as

$$\Pi = U - W, \tag{2.1}$$

where U is the elastic strain energy stored in the body and W is the work performed by external forces. The strain energy release rate is defined as the energy stored for crack extension per unit increase in crack area [22, 23, 24],

$$G = -\frac{d\Pi}{dA}. \tag{2.2}$$

Crack propagation becomes energetically favourable when the available energy equals or exceeds the material resistance to fracture. In LEFM this is typically expressed as

$$G \geq G_c, \tag{2.3}$$

where G_c is the critical energy release rate commonly denoted as fracture toughness [22].

A key result in LEFM is the equivalence between the energy release rate and the stress intensity factor K_α. For linear elastic materials, this relation can be according to [22, 23, 24] written as

$$G = \frac{K_\alpha^2}{E'}, \ \alpha \in \{I, II\} \tag{2.4}$$

where $E' = E$ in-plane stress and $E' = E/(1-\nu^2)$ in-plane strain, where E denotes the Young's moduli and ν Poisson's ratio.

For practical evaluation in fracture test specimens, the energy release rate can also be expressed in terms of the structural compliance. Janssen *et al.* [22] define the compliance as $C = \Delta/F$, where Δ is the load-point displacement and F is the applied load. Using an energy balance argument, the following load-controlled relation is obtained

$$G = \frac{F^2}{2b}\frac{dC}{da}, \tag{2.5}$$

where b is the specimen thickness (out-of-plane), and a is the crack length. This expression forms the basis for compliance calibration procedures used in standard fracture specimens, including the DCB and ENF configurations considered later in this work.

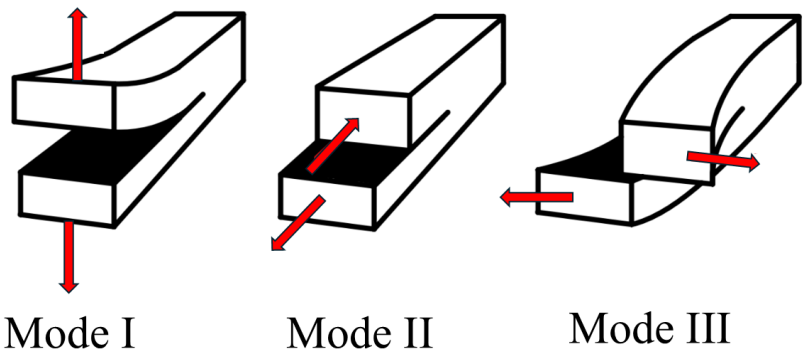

Figure 2.1: Fundamental fracture modes: Mode I (opening), Mode II (in-plane shear) and Mode III (out-of-plane shear).

2.1.2 Mode Decomposition and Mixed-Mode Criteria

Crack loading can occur in three fundamental modes (cf. Fig. 2.1): Mode I (opening), Mode II (in-plane shear), and Mode III (out-of-plane shear) [22, 24]. The total energy release rate is therefore decomposed as

$$G = G_I + G_{II} + G_{III}. \tag{2.6}$$

In mixed-mode fracture, crack growth depends not only on the total available energy but also on the relative contributions of the individual modes. An interaction criterion is therefore required to describe the mixed-mode fracture resistance G_c. One such interaction law is the Benzeggagh–Kenane (BK) criterion. This was the choice adopted by Daniel *et al.* [2], and will also be the one considered here. The criterion is expressed as

$$G_c = G_{Ic} + (G_{IIc} - G_{Ic})\, B^{\eta}, \text{ where } B = \frac{G_{II} + G_{III}}{G} \tag{2.7}$$

where Modes II and III are assumed to have the same critical fracture toughness G_{IIc}. Crack propagation occurs when

$$G \geq G_c(G_I, G_{II}, G_{III}). \tag{2.8}$$

This energetic framework provides the theoretical basis for energy-release-rate-based crack growth formulations and mixed-mode fracture modelling of adhesive debonding employed later in this thesis.

2.2 Corrected Beam Theory for Interfacial Fracture Specimens

Corrected beam theory (CBT) expresses a specimen compliance using beam theory while correcting for (i) transverse shear deformation and (ii) local deformations near the adhesive failure front. The correction is introduced by adding a material-dependent length offset to the measured crack length. The formulation below follows Reeder *et al.* [25], and will be used as an analytic comparison to the results of the implemented UEL.

Extending the formulation in Eq. (2.5), let F denote the applied load, u the measured load-point displacement, b the specimen width, and a the crack length. The compliance is $C = u/F$. The strain energy release rate is written in compliance form as

$$G \;=\; \frac{F^2}{2b}\,\frac{\mathrm{d}C}{\mathrm{d}a} \;=\; \frac{F}{2b}\,\frac{\mathrm{d}u}{\mathrm{d}a}\bigg|_{F=\text{const}}. \tag{2.9}$$

This relation is the starting point used to obtain the closed-form expressions for G once $u(F, a)$ is known from corrected beam kinematics [25].

2.2.1 CBT Crack-Length Correction Parameter

Reeder *et al.* [25] introduces an effective crack length of the form

$$a_{\text{eff}} = f(a, x, h), \tag{2.10}$$

where h is the thickness of one delaminated arm (*i.e.* the total thickness of the specimen is $H = 2h$). The correction parameter x is according to Reeder *et al.* computed from orthotropic elastic constants as

$$x = \sqrt{\frac{E_{11}}{11G_{13}}\left[3 - 2\left(\frac{\Gamma}{1+\Gamma}\right)^2\right]}, \qquad \Gamma = 1.18\,\frac{\sqrt{E_{11}E_{22}}}{G_{13}}. \tag{2.11}$$

Here E_{11} is the modulus along the specimen length direction, E_{22} the transverse modulus, and G_{13} the out-of-plane shear modulus in the arm material representation. A representation of a general specimen is shown in Fig. 2.2, where the directions $1, 2, 3$ are shown alongside total specimen length $2L$, adherend thickness h, initial crack length a_0 and the current crack length a.

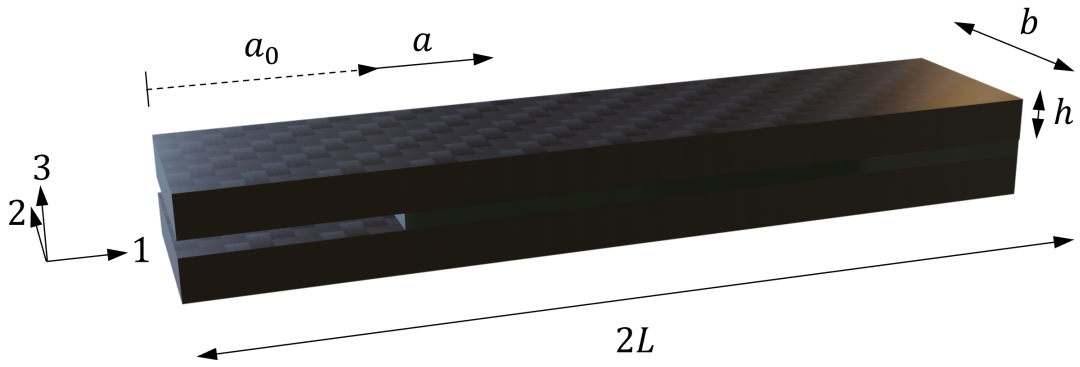

Figure 2.2: Schematic of a general specimen and its characteristic definitions. The specimen is represented with a finite thickness adhesive bondline.

2.2.2 Double Cantilever Beam — Mode I

An example of a DCB sample is shown in Fig. 2.3. A DCB specimen consists of two identical, slender rectangular beams bonded together with a pre-crack or insert at one end, forming a symmetric, double-arm geometry that opens under tensile loading. Reeder *et al.* [25] adopt the Mode I effective crack length $a_{eff,I} = a + xh$ which give the load-point opening displacement

$$u(F, a) = \frac{2(a_{eff,I})^3}{3EI_{\text{del}}}\,F = \frac{2(a + xh)^3}{3EI_{\text{del}}}\,F, \tag{2.12}$$

where EI_{del} is the bending stiffness of one adherend arm. The bending stiffness and second moment of area for a rectangular arm of width b and thickness h is

$$I_{\text{del}} = \frac{bh^3}{12}, \qquad EI_{\text{del}} = E_{11}\,I_{\text{del}} = E_{11}\,\frac{bh^3}{12}. \tag{2.13}$$

Substituting Eq. (2.12) into (2.9) yields the Mode I energy release rate

$$G_I(F,a) = \frac{(a+xh)^2}{b\,EI_{\text{del}}} F^2. \tag{2.14}$$

2.2.3 End-Notched Flexure (3-point) — Mode II

The other geometry that is considered in this work is the ENF specimen. The ENF specimen consists of a rectangular beam with a mid-plane pre-crack, loaded in three-point bending to create a symmetric shear-driven crack propagation. A schematic of the sample is shown in Fig. 2.4. By introducing the Mode II crack-length correction for an ENF specimen as Reeder *et al.* [25] the effective crack length $a_{\text{eff,II}} = a + 0.42\,xh$, along with the mid-point displacement can be expressed as

$$u(F,a) = \frac{3(a_{eff,II})^3 + 2L^3}{96\,EI_{\text{del}}} F = \frac{3(a+0.42\,xh)^3 + 2L^3}{96\,EI_{\text{del}}} F. \tag{2.15}$$

By inserting Eq. (2.15) into Eq. (2.9) the corresponding Mode II energy release rate can according to Reeder *et al.* [25] be calculated as

$$G_{II}(F,a) = \frac{3(a+0.42\,xh)^2}{64\,b\,EI_{\text{del}}} F^2. \tag{2.16}$$

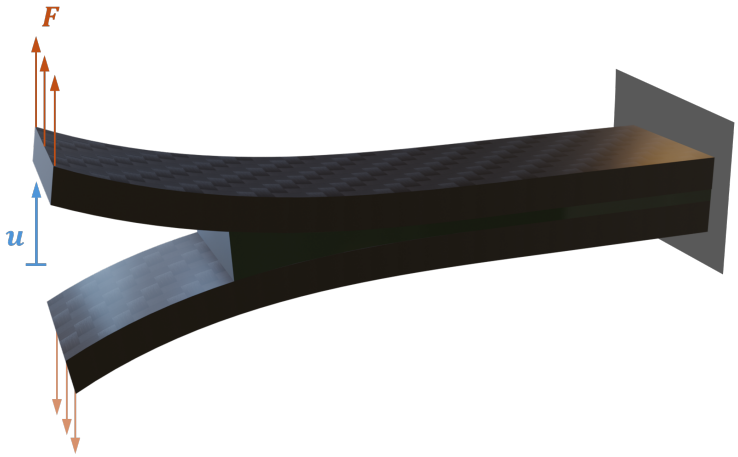

Figure 2.3: Schematic of the double cantilever beam configuration. The specimen is represented with a finite thickness adhesive bondline.

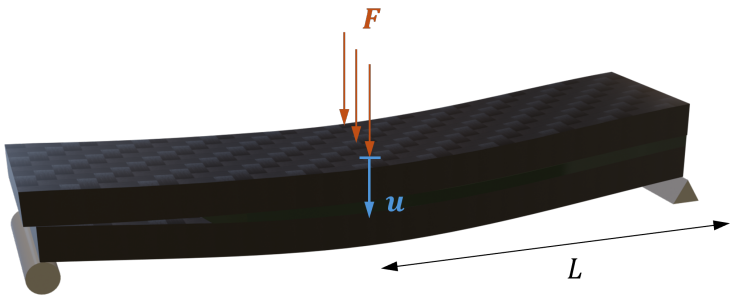

Figure 2.4: Schematic of the end-notched flexure configuration. The specimen is represented with a finite thickness adhesive bondline.

2.3 Interfacial Fracture Modelling

3.1. This section introduces three approaches for modelling interfacial fracture. First, the VCCT is presented together with its limitations in explicit dynamic simulations. Next, the CZM approach is introduced, highlighting its sensitivity to mesh discretisation. Finally, the proposed ERRC method is presented. The ERRC method combines VCCT-based ERR evaluation with cohesive energy dissipation to ensure smooth damage evolution and stable crack propagation in explicit dynamics, even for relatively coarse meshes. This method act as the theoretical basis of the numerical implementation described in Chap. 3.

2.3.1 VCCT

The VCCT is widely adopted as an industrial standard for evaluating the ERR. As discussed by Kruger [26], this is primarily because it provides an explicit separation of the individual fracture modes. This direct modal decomposition, together with the fact that the ERR can be evaluated solely from nodal forces and displacements obtained from a finite element solution, makes VCCT a simple, robust, and computationally efficient crack propagation criterion.

The method is based on LEFM and assumes that the energy released during crack growth is equal to the energy required to close the crack over the same increment. For a regular mesh, when the crack advances one element length l from the initial crack length a_0 with crack tip at node $(i-1)$ to a_0+l with crack tip at node (i), as illustrated in Fig. 2.5, the released energy ΔE is assumed equal to the virtual work required to close the newly created crack surfaces.

Furthermore, it is assumed that a subsequent crack extension from a_0+l (node i) to a_0+2l (node $i+1$) does not significantly alter the mechanical state at the crack tip. Consequently, the displacement at node $(i-2)$ ahead of the crack tip at node $(i-1)$ Fig. 2.5 is equal to the displacement at node $(i-1)$ for a crack tip at node (i) Fig. 2.6 [26].

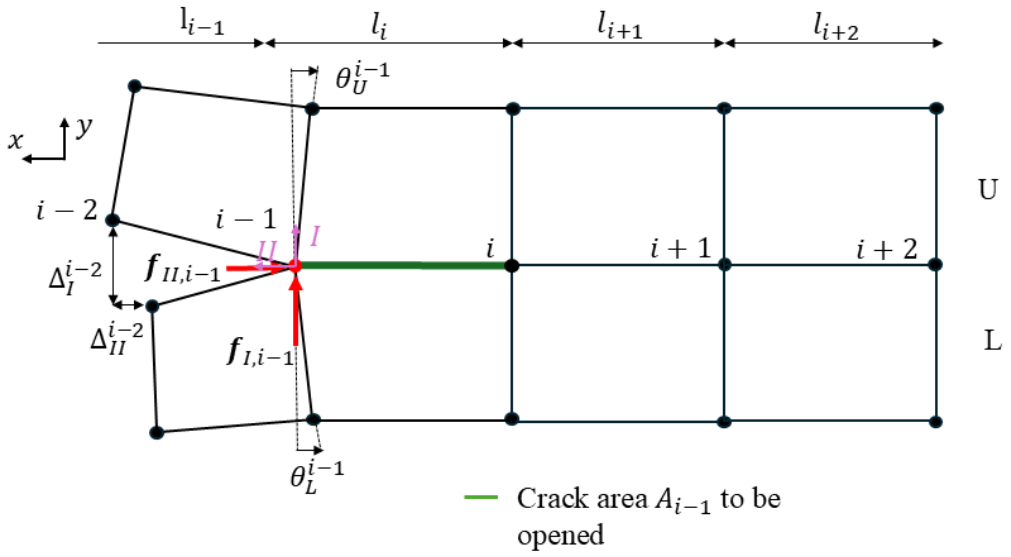

Figure 2.5: A sketch of an initial crack tip configuration.

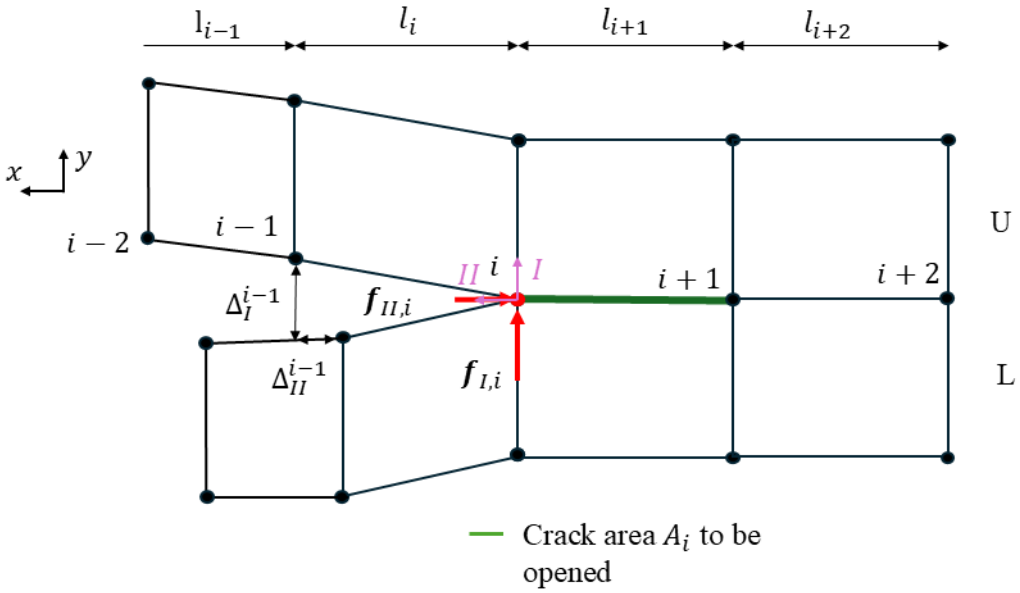

Figure 2.6: A sketch of a progressed crack tip configuration.

For the initial crack tip configuration shown in Fig. 2.5 in 2D, the mode-specific ERRs, G_α for $\alpha = \{I, II, III\}$, are evaluated using three measures. First, there is the nodal forces at crack tip, denoted f_I^{i-1}. Secondly, there is the crack surface area increment, denoted A_{i-1}, which assumes a constant width b, *i.e.* $A_{i-1} = l_i \cdot b$. Finally, the directional nodal displacement behind the crack tip Δ_α^{i-2} is required. As discussed by Daniel *et al.* [2], this allows the mode-specific ERRs to be expressed

as

$$G_I = \frac{1}{2A^{i-1}} f_I^{i-1} \langle \Delta_I^{i-2} \rangle, \tag{2.17a}$$

$$G_{II} = \frac{1}{2A^{i-1}} f_{II}^{i-1} \Delta_{II}^{i-2}, \tag{2.17b}$$

$$G_{III} = \frac{1}{2A^{i-1}} f_{III}^{i-1} \Delta_{III}^{i-2}, \tag{2.17c}$$

where the Macaulay brackets $\langle \cdot \rangle$ defined as $\langle x \rangle = \frac{1}{2}(x + |x|)$ ensure that only physically admissible relative displacements contribute to the fracture energy, *i.e.* that there is no penetration.

To enable the method to be extended to arbitrarily sized elements, recent work Daniel *et al.* [1, 2] refined this approximation. They account for differences between the element length behind the crack tip, l_{i-1}, and the element length ahead of the crack tip, l_i. This is achieved through the introduction of mode-dependent shape correction factors c_I^i, c_{II}^i, and c_{III}^i. Consider an initial crack tip configuration like that shown in Fig. 2.5. Then

$$G_I = \frac{1}{2A^{i-1}} f_I^{i-1} \langle \Delta_I^{i-2} \rangle c_I^{i-1}, \tag{2.18a}$$

$$G_{II} = \frac{1}{2A^{i-1}} f_{II}^{i-1} \Delta_{II}^{i-2} c_{II}^{i-1}, \tag{2.18b}$$

$$G_{III} = \frac{1}{2A^{i-1}} f_{III}^{i-1} \Delta_{III}^{i-2} c_{III}^{i-1}. \tag{2.18c}$$

The correction factors are then given by

$$c_I^{i-1} = \frac{l^i}{\Delta_I^{i-2}} \theta^0 (1 - C) + C^2, \tag{2.19a}$$

$$c_{II}^{i-1} = C, \tag{2.19b}$$

$$c_{III}^{i-1} = C \tag{2.19c}$$

Here, $C = \frac{l_i}{l_{i-1}}$ and θ^0 is the crack tip opening angle approximated as $\theta^0 = \theta_U^0 + \theta_L^0$. These angles are also illustrated in Fig. 2.5. Based on LEFM, if the measured ERR reaches the critical fracture toughness G_c the crack tip is progressed to node i and the process restarts according to new progressed crack tip configuration seen Fig. 2.6

It should be noted that the mesh and crack path are not always aligned with the global coordinate system of the model. For this reason, the forces, f_α, and openings, Δ_α in each mode, are taken in a local interface basis $(\mathbf{e}_I, \mathbf{e}_{II}, \mathbf{e}_{III})$. In this bases the interface normal, the in-plane crack-propagation tangent, and the closing direction are described by $\mathbf{e}_I, \mathbf{e}_{II}$ and $\mathbf{e}_{III}$ respectively. Note also, considering standard vector operations that $\mathbf{e}_{III} = \mathbf{e}_I \times \mathbf{e}_{II}$.

2.3.2 Cohesive Zone Modelling

CZM describes interfacial fracture by equipping a potential crack surface with a traction–separation relation. The interface can therefore transmit tractions while

progressively separating, and crack growth emerges naturally as the traction capacity degrades to zero. A central requirement is energy consistency: the work done by the cohesive tractions during complete separation must equal the critical fracture energy G_c, *i.e.* the area under the traction–separation law. In contrast to a purely LEFM- or VCCT-based crack advance, this provides a mechanism for dissipating fracture energy during progressive separation.

In three dimensions, debonding is represented on an interface surface Γ_d separating two subdomains Ω^+ and Ω^-, as illustrated in Fig. 2.7. Let $\mathbf{u}^U$ and $\mathbf{u}^L$ denote the displacements on the two sides of the interface where U and L denotes upper and lower side, respectively. The displacement jump is defined as

$$[\![\mathbf{u}]\!] = \mathbf{u}^U - \mathbf{u}^L, \tag{2.20}$$

and is resolved in a local orthonormal basis $(\mathbf{n}, \mathbf{s}_1, \mathbf{s}_2)$, according to Section 2.3.1 where $\mathbf{n}$ is the interface normal and $\mathbf{s}_1, \mathbf{s}_2$ span the interface plane as defined in Fig. 2.8. The corresponding separation components are

$$\Delta_I = [\![\mathbf{u}]\!]\cdot\mathbf{n}, \qquad \Delta_{II} = [\![\mathbf{u}]\!]\cdot\mathbf{s}_1, \qquad \Delta_{III} = [\![\mathbf{u}]\!]\cdot\mathbf{s}_2, \qquad \Delta_s = \sqrt{\Delta_{II}^2 + \Delta_{III}^2}. \tag{2.21}$$

In this thesis, only the bilinear traction–separation law is used. This choice is deliberate: the bilinear law provides a transparent relation between physically meaningful parameters such as interfacial strength and fracture toughness, while remaining straightforward to calibrate and implement, as discussed by Noorman *et al.* [7], Turon *et al.* [27] and Davila *et al.* [28]. Considering a single mode, the law is illustrated in Fig. 2.9. It is characterised by an initial penalty stiffness K, a peak traction X, the displacement at damage initiation $\Delta_0 = X/K$, and a final separation Δ_f at complete decohesion,

$$t(\Delta) = \begin{cases} K\,\Delta, & 0 \le \Delta \le \Delta_0, \\ X\left(1 - \dfrac{\Delta - \Delta_0}{\Delta_f - \Delta_0}\right), & \Delta_0 < \Delta \le \Delta_f, \\ 0, & \Delta > \Delta_f. \end{cases} \tag{2.22}$$

The associated fracture energy follows directly from the area under the traction–separation curve,

$$G_c = \int_0^{\Delta_f} t(\Delta)\,\mathrm{d}\Delta = \frac{1}{2} X_0 \Delta_f, \tag{2.23}$$

such that

$$\Delta_f = \frac{2G_c}{X}. \tag{2.24}$$

Hence, for a fixed fracture toughness, the peak traction primarily controls the extent of the softening region, whereas the penalty stiffness sets the initial elastic response and the displacement at damage onset. For mixed-mode loading, the total ERR is defined as the sum of the energy release rate for each mode and the fracture toughness is defined using the BK criterion as introduced in Sec. 2.1.2. For this reason they are not repeated here.

This link between constitutive parameters and the softening region is central, because the cohesive law represents a finite fracture process zone (FPZ) rather than a mathematically sharp crack. In a finite-element model, the FPZ must be resolved by the mesh in order for the degradation process to be captured accurately. As an example, for delamination in laminated composites this is restrictive, since the interlaminar FPZ is often very short. Turon *et al.* [11] showed that the element size should remain smaller than the cohesive-zone length, and that several elements across the FPZ are typically required for accurate results. Later work, by Daniel *et al.* [1, 2] on large-element delamination modelling makes the same point and highlights that, for brittle ply interfaces, the FPZ is often below 1 mm, which explains the high mesh demands of classical CZM.

A useful compact estimate of the FPZ length for pure Mode I and II loading is given by Soto *et al.* [29],

$$l_{\mathrm{FPZ},\alpha} = M_\alpha \frac{E_\alpha G_{\alpha,c}}{X_\alpha^2}, \qquad \alpha \in \{I, II\}, \tag{2.25}$$

where E_α is a directional elastic moduli, $G_{\alpha,c}$ the critical fracture energy, X_α the cohesive strength, and M_α a mode-dependent coefficient. For the pure-mode cases considered here,

$$l_{\mathrm{FPZ},I} = M_I \frac{E_I G_{IC}}{X_I^2}, \qquad l_{\mathrm{FPZ},II} = M_{II} \frac{E_{II} G_{IIC}}{X_{II}^2}. \tag{2.26}$$

The key implication is that the FPZ decreases strongly with increasing cohesive strength,

$$l_{\mathrm{FPZ},\alpha} \propto \frac{1}{X_\alpha^2}, \tag{2.27}$$

so physically high strengths rapidly increase the mesh requirement.

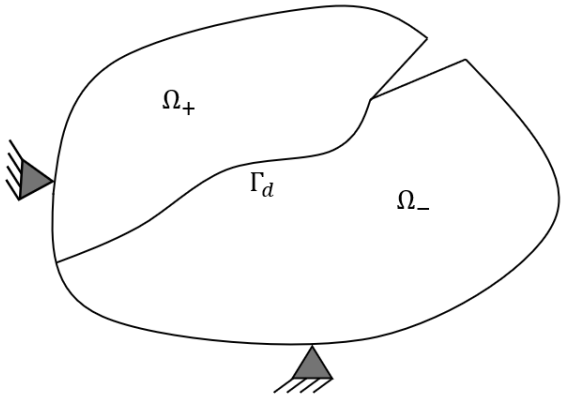

Figure 2.7: Material discontinuity representation.

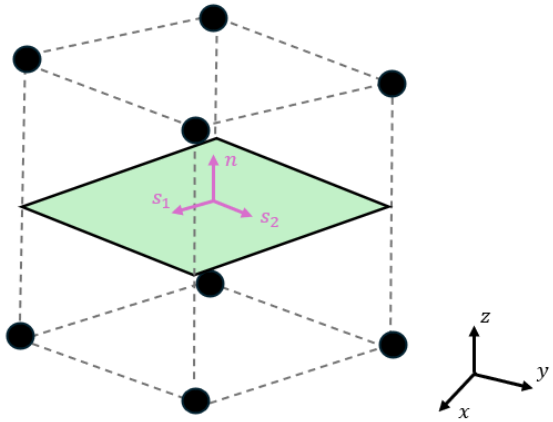

Figure 2.8: Interfacial surface deformation of a three-dimensional CZ respresentation.

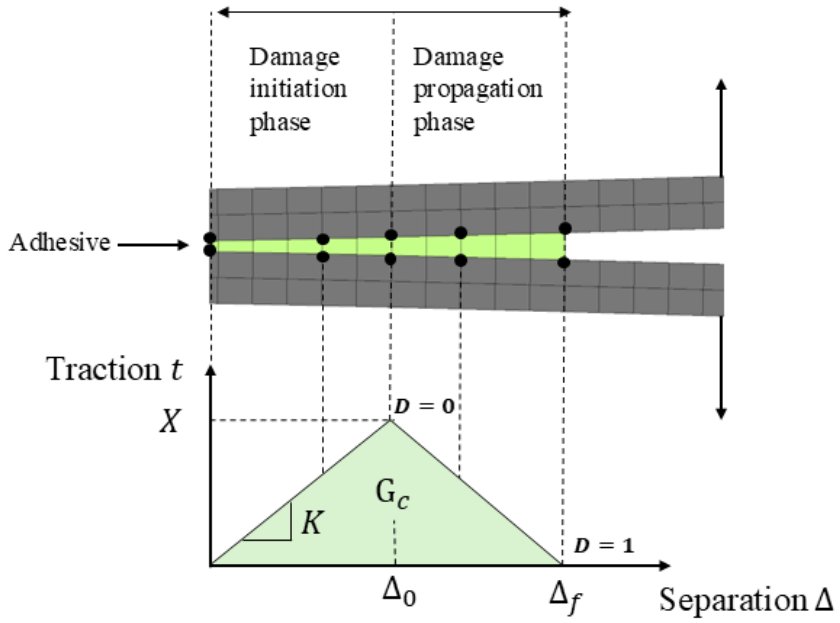

Figure 2.9: Schematic of a FPZ on a coupon with an applied opening (Mode I) load. The adhesive is modelled by the bilinear law. Along the illustrated cohesive-zone the elements representative state in the traction-separation law are marked to visualise the pointwise FPZ progression.

In cohesive-zone modelling of adhesive fracture, the interface stiffness must be chosen such that the undamaged interface remains sufficiently stiff without introducing excessive fictitious compliance into the structural response. At the same time, the interface should not be made unnecessarily stiff, since this may lead to numerical difficulties such as very small stable time increments and oscillatory traction response in explicit analyses. Generalising the method by Turon *et al.* [11] to adhesive fracture, the interface stiffness may be related to the stiffness of the adjacent adherends and their characteristic thickness through

$$K_n = \beta \frac{E_3}{t}, \qquad K_s = \beta \frac{G_{13}}{t}, \tag{2.28}$$

where K_n and K_s denote the normal and tangential (shear) interface stiffnesses, E_3 is the through-thickness elastic modulus of the adjoining adherend, G_{13} is the corresponding shear modulus, t is the characteristic adherend thickness, and β is a dimensionless scaling factor chosen such that the cohesive contribution to the global compliance remains limited.

Eq. (2.28) provides a mechanically motivated stiffness estimate rather than an arbitrary penalty parameter. Its purpose is to ensure that the interface is stiff compared to the surrounding adherands prior to damage, while still avoiding unnecessarily large stiffness values. In this sense, the interface stiffness acts both as a constitutive quantity and as a numerical parameter. A low value of β increases artificial compliance in the undamaged state, whereas a high value reduces this compliance but may result in a more demanding explicit solution. Turon *et al.* [11] show that the equivalent stiffness reduction caused by the interface decreases with increasing β, and that values above approximately $\beta = 50$ lead to less than about 2% stiffness loss in the simplified setting used for their derivation for ply delamination.

In the present work, Eq. (2.28) is used as the theoretical basis for selecting initial cohesive stiffnesses in the benchmark models. The final values are then assessed together with the mesh size, the fracture-process-zone length, and the explicit solution behaviour.

2.3.3 Energy-Release-Rate Cohesive

VCCT in its bare form and conventional CZM as described above are both well established but address different solver regimes. VCCT is well-suited to implicit quasi-static or pseudo-time analyses, where the equilibrium iteration absorbs the abrupt release of a node row once the criterion $G \geq G_c$ is met. Under explicit dynamic time integration, the same instantaneous release injects an impulse into the structure equal to the stored elastic strain energy of the released row. As discussed by Daniel *et al.* [1], without an added dissipation mechanism the crack propagates from node pair i to $i + 1$ in a single step and excites uncontrollable vibrations that can drive an unphysical run-away crack, an example of which is shown in Fig. 2.10.

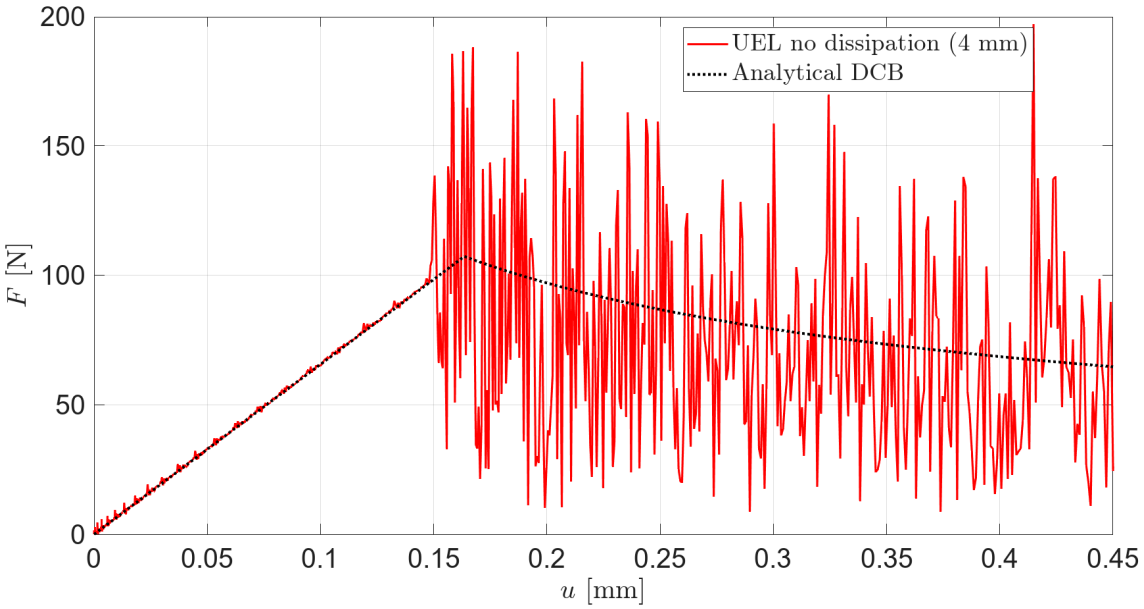

Figure 2.10: VCCT for simulating crack propagation in a double cantilever beam.

Conventional CZM avoids this by softening the interface gradually, but at the cost of requiring elements smaller than the fracture process zone, as discussed in-Section 2.3.2.The hybrid method described next combines the VCCT propagation criterion with a cohesive softening branch and is constructed precisely to give VCCT the missing explicit-dynamics dissipation while preserving its tolerance for elements larger than the FPZ.

The approach adopted by Daniel *et al.* [1, 2] can be viewed as a hybrid between classical fracture mechanics and cohesive modelling, designed specifically to enable delamination growth when the in-plane discretisation is too coarse to resolve a physical fracture process zone. The core separation of roles is:

- **VCCT acts as propagation criterion:** VCCT is used to compute the local energy release rate components (G_I, G_{II}, G_{III}) according to Equation 2.17 at the crack front and to decide whether the crack should advance, through a mixed-mode toughness criterion $G \geq G_c$ (*e.g.* BK from Eq. (2.7)).
- **Cohesive law acts as progressive release:** Once propagation is triggered, a cohesive traction–separation relation is used to progressively release the interface over the newly created crack area such that the dissipated work equals $G_c \Delta A$, avoiding an abrupt node-row release.

A key modelling requirement is **energy consistency**: VCCT provides the driving force and the crack-growth criterion, while the cohesive law provides the dissipation mechanism. The formulation is therefore constructed to avoid double counting of fracture energy by ensuring that the cohesive work associated with an incremental advance ΔA is exactly the target fracture energy $G_c \Delta A$, rather than adding an extra dissipation on top of a pre-released crack. The following mathematical concepts are derived from work by Daniel *et al.*[1]. This method utilises a kinematically tied interface, upon satisfying the growth criterion presented in Subsection 2.3.1, the tying force $\mathbf{f}_{\text{tied}}$ at node pair i are recorded as F_I^0, F_{II}^0 and F_{III}^0, and the constraint is replaced by a nodal cohesive response.

The interface is kept tied by equalising the nodal accelerations of each corresponding pair of upper U and lower L nodes associated with the adherend elements, following the proposed method by Framby *et al.* [30], thus ensuring

$$\ddot{\mathbf{u}}^L = \ddot{\mathbf{u}}^U \quad \Rightarrow \quad \dot{\mathbf{u}}^L = \dot{\mathbf{u}}^U,\ \mathbf{u}^L = \mathbf{u}^U \tag{2.29}$$

Consider an interface nodal pair i as shown in Fig. 2.11. For the upper or lower interfacial node, the internal force $\mathbf{f}_{\text{int},i}$ is corrected with a tied constraint force, $\mathbf{f}^L_{\text{tied},i} = -\mathbf{f}^U_{\text{tied},i} = \mathbf{f}_{\text{tied},i}$ utilising the nodal pair masses m^U and m^L, resulting in the updated nodal force balance

$$m^L_i \ddot{\mathbf{u}}_i = \mathbf{f}^L_{\text{int},i} + \mathbf{f}_{\text{tied},i} = \mathbf{f}^L_i, \tag{2.30}$$

$$m^U_i \ddot{\mathbf{u}}_i = \mathbf{f}^U_{\text{int},i} - \mathbf{f}_{\text{tied},i} = \mathbf{f}^U_i. \tag{2.31}$$

Through algebraic manipulation, the constraint force can be derived as

$$\mathbf{f}_{\text{tied},i} = \frac{m^L_i \mathbf{f}^U_{\text{int},i} - m^U_i \mathbf{f}^L_{\text{int},i}}{m^L_i + m^U_i}. \tag{2.32}$$

The corrected nodal forces can then be expressed as

$$\mathbf{f}^L_i = \frac{m^L_i}{m^L_i + m^U_i} \left(\mathbf{f}^L_{\text{int},i} + \mathbf{f}^U_{\text{int},i} \right), \tag{2.33}$$

$$\mathbf{f}^U_i = \frac{m^U_i}{m^L_i + m^U_i} \left(\mathbf{f}^L_{\text{int},i} + \mathbf{f}^U_{\text{int},i} \right), \tag{2.34}$$

The introduction of the interfacial force can be shown to be fully energy consistent by considering the total work performed while enforcing the kinematic constraint in Eq. (2.29),

$$\mathbf{f}_{\text{tied},i} \cdot \left(\mathbf{u}^L_i - \mathbf{u}^U_i \right) = 0. \tag{2.35}$$

In practice, the nodal cohesive response acts as a spring-like relation between coincident interface nodes applying the nodal forces F_α, with an associated representative area used to enforce the correct energy dissipation, $G_c \Delta A$, and a penalty contact stiffness k_p. Daniel *et. al.* [1] define a degradation of the initial tying forces through a single damage variable D shared by modes I, II and III such that

$$F_I = \frac{|F^0_I|}{\Delta^{\max}_I}(1-D)\left(\langle \Delta_I \rangle + \Delta^{\text{offset}}_I \right) - k_p \langle -\Delta_I \rangle, \tag{2.36}$$

$$F_{II} = \frac{|F^0_{II}|}{\Delta^{\max}_{II}}(1-D)\left(\Delta_{II} + \Delta^{\text{offset}}_{II} \right), \tag{2.37}$$

$$F_{III} = \frac{|F^0_{III}|}{\Delta^{\max}_{III}}(1-D)\left(\Delta_{III} + \Delta^{\text{offset}}_{III} \right). \tag{2.38}$$

here F^0_α denotes the evaluated tied force Eq. (2.32) for either the lower or upper node at the point of $G > G_c$. The initial Δ^{offset}_α offset is added in order to avoid infinite

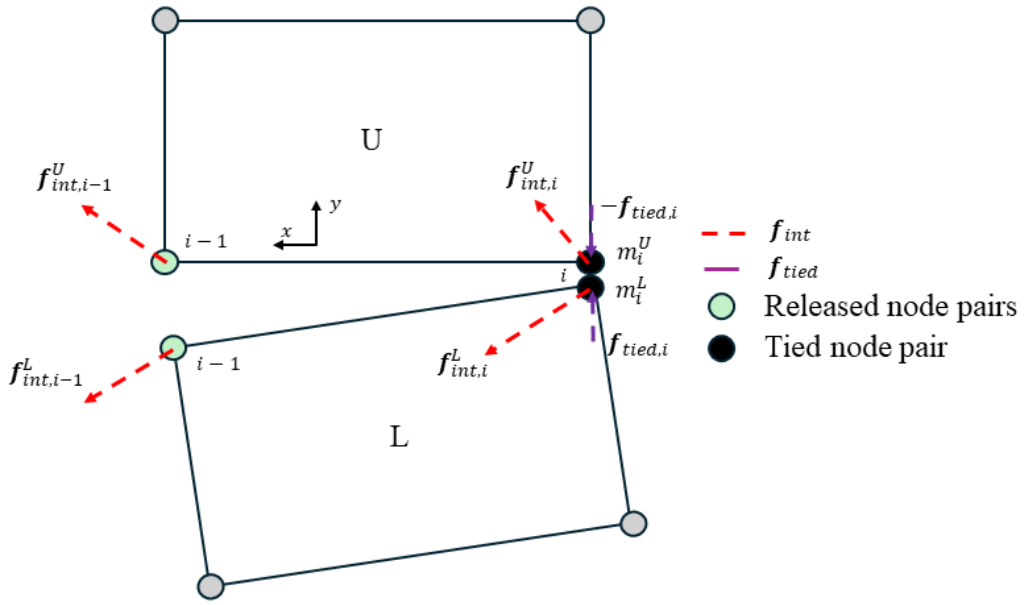

Figure 2.11: Illustration of elemental kinematics for the tied constraint of nodal pair i.

stiffness of the cohesive law and thus infinitely small time step at the initialisation of the cohesive law in an explicit FEM simulation. The initial offset is defined as

$$\Delta_\alpha^{offset} = \frac{F_\alpha^0}{k_\alpha^0}, \qquad \alpha \in \{I, II, III\} \tag{2.39}$$

with k_α^0 is a user defined input. This is acquired by calibrating the stiffness to a desired time step according to mass-spring system of the effective nodal mass as $d\tau = \sqrt{\frac{m_e}{k_\alpha^0}}$. Accounting for the initial offsets, at time τ, $\Delta_\alpha^{\max}$ becomes

$$\Delta_I^{\max} = \max_{\tau\in[0,t]} \langle \Delta_I(\tau) + \Delta_I^{\text{offset}} \rangle \tag{2.40}$$

$$\Delta_{II}^{\max} = \max_{\tau\in[0,t]} (\Delta_{II}(\tau) + \Delta_{II}^{\text{offset}}) \tag{2.41}$$

$$\Delta_{III}^{\max} = \max_{\tau\in[0,t]} (\Delta_{III}(\tau) + \Delta_{III}^{\text{offset}}) \tag{2.42}$$

The damage-driving function is chosen such that each mode contributes in proportion to the ERR-based mode mixity B from Eq. (2.7), which is evaluated at initiation and kept constant during the nodal release

$$d = B_I \frac{\langle \Delta_I^* \rangle}{|\Delta_{fI}|} + B_{II} \frac{\langle \Delta_{II}^* \rangle}{|\Delta_{fII}|} + B_{III} \frac{\langle \Delta_{III}^* \rangle}{|\Delta_{fIII}|}, \text{ where } B_\alpha = \frac{G_\alpha}{G} \tag{2.43}$$

and the damage variable is taken as the maximum attained value,

$$D = \max_{\tau\in[0,t]} \Big(\min(d(\tau), 1) \Big), \tag{2.44}$$

with

$$\Delta_I^* = \langle \Delta_I + \Delta_I^{\text{offset}} \rangle - |\Delta_I^{\text{offset}}|, \tag{2.45}$$

$$\Delta_{II}^* = |\Delta_{II} + \Delta_{II}^{\text{offset}}| - |\Delta_{II}^{\text{offset}}|, \tag{2.46}$$

$$\Delta_{III}^* = |\Delta_{III} + \Delta_{III}^{\text{offset}}| - |\Delta_{III}^{\text{offset}}|. \tag{2.47}$$

The failure openings $\Delta_{f\alpha}$ are defined so that a linear softening from F_α^0 to 0 would consume the modal energy $A_i G_\alpha$:

$$\Delta_{f\alpha} = \frac{2 A_i G_\alpha}{F_\alpha^0}, \qquad \alpha \in \{I, II, III\}. \tag{2.48}$$

A central result is that the cohesive work dissipated during the progressive release is proportional to D,

$$\mathcal{G} = A_i \, G_c \, D, \tag{2.49}$$

so that full release ($D = 1$) dissipates exactly $A_i G_c$ for the newly created area. Hence, VCCT provides the growth trigger and target fracture energy, while the cohesive law provides a controlled mechanism to convert stored elastic energy into fracture work over a finite time, avoiding the instabilities of sudden node-pair releases.

Under explicit time integration, the abrupt tied-to-cohesive transition at the release moment excites local oscillations of the freed interface nodes. To stabilise this transient, Daniel *et al.* [2] add a mass-proportional viscous damper at the cohesive node pair. The damping force is linear in the relative velocity of the two interface nodes,

$$\mathbf{F}_\alpha^{\text{damp}} \;=\; c_\alpha \, \dot{\Delta}_\alpha, \qquad \alpha \in \{I, II, III\}, \tag{2.50}$$

with the per-mode damping coefficient set to a fraction ζ of the critical damping of the equivalent single-mode spring–mass system:

$$c_\alpha \;=\; \zeta \, \frac{4}{b} \sqrt{m_{\text{eff}} \, k_\alpha^0}, \qquad m_{\text{eff}} \;=\; \frac{m^L \, m^U}{m^L + m^U}, \qquad b = 1.2, \tag{2.51}$$

where m^L, m^U are the lumped nodal masses of the lower and upper adherand at the active pair. The damping is active only during the cohesive phase, $0 < D < 1$, and is removed at full release ($D = 1$) so that the open crack faces carry only the cohesive constitutive force and contact penalty.

The interface force used in the explicit time integration during the cohesive phase is then the sum of the constitutive cohesive contribution from Eqs. (2.36)–(2.38) and the viscous damping term,

$$F_\alpha \;=\; \underbrace{\frac{|F_\alpha^0|}{\Delta_\alpha^{\max}} (1 - D) \left(\langle \Delta_\alpha \rangle + \Delta_\alpha^{\text{offset}} \right)}_{\text{constitutive (Eqs. (2.36)–(2.38))}} + c_\alpha \, \dot{\Delta}_\alpha \qquad (0 < D < 1). \tag{2.52}$$

The cohesive opening and its rate are modal projections of the relative nodal kinematics,

$$\Delta_\alpha \;=\; \left(\mathbf{u}^U - \mathbf{u}^L \right) \cdot \mathbf{e}_\alpha, \qquad \dot{\Delta}_\alpha \;=\; \left(\dot{\mathbf{u}}^U - \dot{\mathbf{u}}^L \right) \cdot \mathbf{e}_\alpha, \tag{2.53}$$

so that both the constitutive and damping terms in Eq. (2.52) live in the same local per-pair frame as the VCCT trigger.

2.3.4 Local Arbitrary-Front Formulation

Two ingredients still need a discrete definition: the predecessor pair that feeds the VCCT trigger, and the surface area A_i that enters the energy balance.

2.3.4.1 Node-pair representation and corner-cyclic ring

The bonded interface is discretised by eight-node solid interface elements (one element through the bondline thickness). Each element contributes four lower–upper *node-pair* couples, indexed locally by $p \in \{1, 2, 3, 4\}$; in element-local numbering the two adjacent corners of p are

$$\text{neigh1}(p) \in \{2, 3, 4, 1\}, \qquad \text{neigh2}(p) \in \{4, 1, 2, 3\}, \tag{2.54}$$

so that neigh1(p) and neigh2(p) are the two corner-cyclic ring-neighbours of p within that element. A physical node-pair is generally shared by up to four elements; the unique global identifier of a physical pair is the ordered tuple of its two node IDs (the *pair key*). Two elements refer to the same physical pair if and only if their pair keys are identical.

2.3.4.2 Local front rule

Each pair carries a state $s_p \in \{0, 1, 2\}$ (tied, cohesive, open). A tied pair p is marked as a *front pair* whenever at least one of its corner-cyclic ring-neighbours is open. The marking is captured by an integer tag τ_p that records which side(s) opened:

$$\tau_p = \mathbf{I}\Big[s_{\text{neigh1}(p)} = 2\Big] + 2\,\mathbf{I}\Big[s_{\text{neigh2}(p)} = 2\Big] \;\in\; \{0, 1, 2\}, \tag{2.55}$$

where $\mathbf{I}[\cdot]$ is the indicator function. The value $\tau_p = 1$ means only neigh1 is open, $\tau_p = 2$ means only neigh2 is open, and $\tau_p = 0$ means p is not a front pair. The predecessor that enters the VCCT trigger is, by construction of this rule, one of the open ring-neighbours of p.

2.3.4.3 Geometric Predecessor Pre-Identification

Let $\mathbf{e}_{II}$ be the in-plane propagation direction and $\mathbf{e}_{III} = \mathbf{e}_I \times \mathbf{e}_{II}$ the in-plane transverse direction. For the precrack-adjacent pairs of the axis-aligned DCB/ENF coupons used here, $\mathbf{e}_{II}$ is taken as the global $+x$ axis and $\mathbf{e}_{III}$ as the global $+y$ for the purpose of evaluating Eq. (2.57) at initialisation only; the runtime per-pair basis of Eq. (2.61) is built from the geometric step $(\mathbf{m}_p - \mathbf{m}_{\text{pred}(p)})$ and is therefore already per-pair, rotating with the local mesh orientation. The global-axis assumption therefore enters only at init through Eq. (2.57). This method is an attempt to handle arbitrary delamination fronts from a simpler logic, compared to the more complex one employed by Daniel *et al.* [2], and is left for future work. For each ring-neighbour $n \in \{\text{neigh1}(p), \text{neigh2}(p)\}$, define the in-plane offset

$$\boldsymbol{\delta}_n \;=\; \mathbf{m}_n - \mathbf{m}_p. \tag{2.56}$$

The upstream predecessor is the ring-neighbour whose offset has the most negative projection onto $\mathbf{e}_{II}$,

$$\mathrm{pred}(p) \;=\; \underset{n\in\{\mathrm{neigh1}(p),\,\mathrm{neigh2}(p)\}}{\operatorname{argmin}} \boldsymbol{\delta}_n \cdot \mathbf{e}_{II}. \tag{2.57}$$

The selected neighbour is stored as its *pair-key* tuple

$$\mathbf{k}_{\mathrm{pred}(p)} \;=\; \Big(\min(n_p^-, n_p^+),\; \max(n_p^-, n_p^+)\Big)\,\Big|_{n=\mathrm{pred}(p)} \tag{2.58}$$

with n_p^-, n_p^+ the global node IDs of $\mathrm{pred}(p)$.

At runtime the predecessor side is locked to the ring-neighbour whose pair-key matches the stored identifier:

$$\mathrm{iftype}_{\mathrm{eff}}(p) \;=\; k \quad \text{such that} \quad \mathbf{k}_{\mathrm{neigh}k(p)} = \mathbf{k}_{\mathrm{pred}(p)}, \qquad k \in \{1,2\}. \tag{2.59}$$

The same $\mathrm{pred}(p)$ also fixes the in-plane direction of the per-pair basis in Eq. (2.61): predecessor selection and basis construction are one inline-neighbour choice expressed in two contexts.

2.3.4.4 Pair Basis from the Geometric Chain Step

With the predecessor fixed by Section 2.3.4.3, the per-pair orthonormal basis $(\mathbf{e}_I, \mathbf{e}_{II}, \mathbf{e}_{III})$ that projects the VCCT forces and openings is built from the same midpoint geometry. The bond-plane normal is defined as,

$$\mathbf{e}_I \;=\; \frac{(\mathbf{m}_3 - \mathbf{m}_1) \times (\mathbf{m}_4 - \mathbf{m}_2)}{\|(\mathbf{m}_3 - \mathbf{m}_1) \times (\mathbf{m}_4 - \mathbf{m}_2)\|}, \tag{2.60}$$

The in-plane direction is the geometric step from the chosen predecessor toward p, projected out of $\mathbf{e}_I$ to enforce orthogonality:

$$\tilde{\mathbf{e}}_{II} \;=\; (\mathbf{m}_p - \mathbf{m}_{\mathrm{pred}(p)}) - \Big((\mathbf{m}_p - \mathbf{m}_{\mathrm{pred}(p)})\cdot\mathbf{e}_I\Big)\,\mathbf{e}_I, \qquad \mathbf{e}_{II} \;=\; \frac{\tilde{\mathbf{e}}_{II}}{\|\tilde{\mathbf{e}}_{II}\|}, \qquad \mathbf{e}_{III} \;=\; \mathbf{e}_I \times \mathbf{e}_{II}. \tag{2.61}$$

$\mathbf{m}_{\mathrm{pred}(p)}$ is the current midpoint of the neighbour selected by Eq. (2.59); the basis is therefore set by midpoint geometry alone, with no opening-jump dependence.

2.3.4.5 Direction-Aware Crack-Extension Area

The crack-extension area A_i that enters the VCCT denominator in Eq. (2.17) and the cohesive failure openings of Eq. (2.48) is assembled from a directional partition of the elements surrounding p. Let

$$\mathbf{d}_1 \;=\; \mathbf{m}_p - \mathbf{m}_{\mathrm{neigh1}(p)}, \qquad \mathbf{d}_2 \;=\; \mathbf{m}_p - \mathbf{m}_{\mathrm{neigh2}(p)}, \tag{2.62}$$

be the two in-plane direction vectors pointing from each ring-neighbour toward p. Each element e sharing pair p contributes to ahead-pool k when its centroid $\mathbf{c}_e$ (the average of its four pair midpoints) lies on the p-side of neigh$k$:

$$(\mathbf{c}_e - \mathbf{m}_p) \cdot \mathbf{d}_k > 0, \qquad k \in \{1,2\}. \tag{2.63}$$

The contributing strips accumulate as

$$A_{\mathrm{ahead},k} \;=\; \sum_{e\,:\,(\mathbf{c}_e-\mathbf{m}_p)\cdot\mathbf{d}_k>0} A_e \,/\, 2, \qquad k \in \{1,2\}, \tag{2.64}$$

with A_e the bond-plane area carried by element e for pair p. The pool consumed by VCCT is selected by the front tag:

$$A_i = \begin{cases} A_{\mathrm{ahead},1} & \tau_p = 1, \\ A_{\mathrm{ahead},2} & \tau_p = 2, \end{cases} \tag{2.65}$$

The direction $\mathbf{d}_{\mathrm{iftype}_{\mathrm{eff}}(p)}$ that selects the area pool is the same inline-neighbour offset that fixes $\mathbf{e}_{II}$ in Eq. (2.61). The numerical implementation is described in Chap. 3.

3

Numerical Implementation of the ERR–Cohesive Method

The ERR-cohesive formulation was implemented in LS-DYNA as a user-defined discrete solid element that represents the adhesive interface. The implementation used eight-node user-solid elements with finite thickness, in which each lower–upper nodal pair represents a discrete point on the bonded interface as seen in Fig 3.1. Instead of relying on a conventional material routine to update stresses, the user-element directly evaluates the interface forces, nodal-pair states, and history variables required for VCCT-based crack propagation and cohesive energy release.

The implementation follows the theoretical formulation described in Sec. 2.3.3 where the energy release rate is estimated using VCCT and the subsequent crack opening is controlled by a cohesive release law. The present chapter introduces two different implementation procedures suitable for a both LS-DYNA execution modes, a single-thread serial single-core reference (SCR) implementation and the shared-memory parallel version multi-core capability (MCC). The single-core implementation is the serial reference formulation. Ownership of each nodal pair is resolved by a first-to-evaluate flag held in shared `COMMON` storage rather than by the deterministic non-local election used in the multi-core build. The per-pair mechanics, *i.e.* the tied-force reconstruction, the VCCT trigger and the cohesive release law, are identical. It therefore provides the baseline reference against which both the analytical solution and the multi-core implementation are assessed.

This chapter first describes the elemental algorithm used by both procedures (Sec. 3.1 - Sec. 3.2), followed by execution mode specific extensions (Sec. 3.3 - Sec. 3.4).

3.1 Implementation Framework and Interface Element

Each element contains four lower–upper nodal pairs (see Fig. 3.2)

$$(1,5), \quad (2,6), \quad (3,7), \quad (4,8), \tag{3.1}$$

where each pair defines one discrete interface point. The relative displacement between the upper and lower nodes of a neighbouring released pair is used to evaluate the interface opening, while the corresponding nodal force in a tied pair is used in the VCCT evaluation.

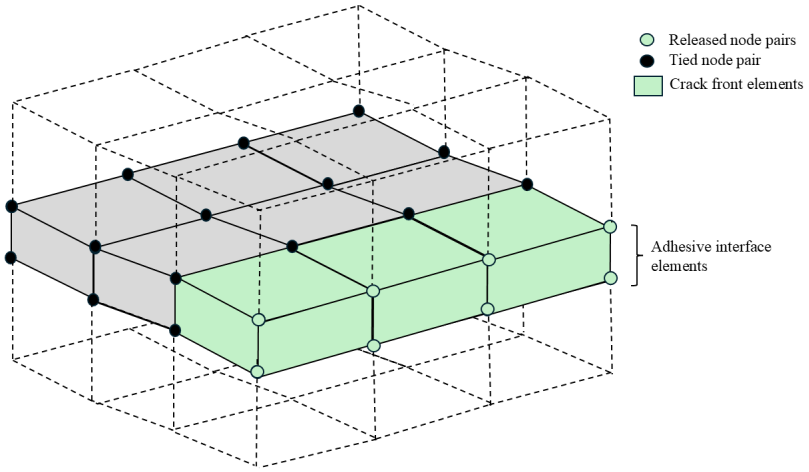

Figure 3.1: Interface with user-defined and conventional elements including crack front elements with tied and released node-pairs.

As explained in the theory chapter, the implementation differs from a conventional cohesive-zone model in that the interface does not initially behave according to a traction–separation law at all integration points. Instead, intact nodal pairs are tied until the local VCCT propagation criterion (Sec. (2.3.1)) is satisfied. Once propagation is triggered, the corresponding nodal pair is transferred to a cohesive release state, where the remaining nodal force is reduced progressively until the pair is fully open.

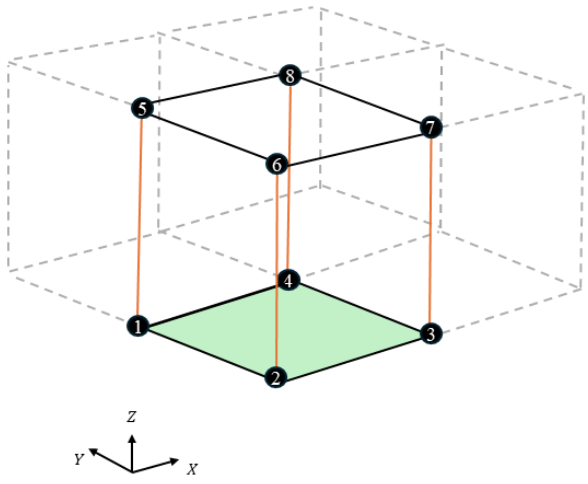

Figure 3.2: Eight-node user-defined interface element used to represent the adhesive layer. Each lower–upper nodal pair defines one discrete interface point at which the tied, cohesive, or open response is evaluated.

This representation allows the crack to advance through large interface elements while still dissipating the correct fracture energy during the release process. The element routine therefore has two main responsibilities. First, it applies the appropriate interface force for each owned nodal pair according to its current state. Second, it updates the history variables that define the state of the interface during

crack propagation.

A complication arises because one physical nodal pair can be present in several neighbouring interface elements. If all element copies were allowed to apply forces independently, the interface force and the dissipated fracture energy would be counted several times. The implementation choice was therefore ensure that only one element is responsible for the mechanics of each physical nodal pair during a solver cycle. How this ownership is handled differs between the two implementation variants presented later in this chapter.

A second complication is that the tied-pair force is reconstructed from the bulk-solid force contributions already assembled by LS-DYNA. The UEL must therefore be evaluated after the surrounding solid elements have written their contributions to the global force array. This solver-ordering requirement is common to both implementation variants and is discussed together with the corresponding LS-DYNA execution modes in Secs. 3.3 and 3.4.

The work was developed in two stages. The energetic logic of the method, including the VCCT trigger, the cohesive release law, and the BK mixed-mode criterion, was first prototyped in two dimensions using MATLAB. The formulation was then transferred to LS-DYNA as the user-defined solid element described in this chapter. The single-core reference (SCR) implementation was developed first and was subsequently extended to the multi-core capability (MCC) implementation by replacing the serial shared-memory bookkeeping with a per-cycle nonlocal synchronisation routine.

3.2 General Logic

The constitutive and topological logic described in this section apply to both execution modes. The single-core and multi-core sections that follow refer back to the equations and definitions introduced here.

3.2.1 Nodal-Pair States and Force Evaluation

Each nodal pair is governed by a local state variable. The state determines whether the pair is intact, undergoing cohesive release, or is fully open. Three states were thus used in the implementation:

$$s_p = \begin{cases} 0, & \text{tied}, \\ 1, & \text{cohesive}, \\ 2, & \text{open}, \end{cases} \tag{3.2}$$

where s_p denotes the state of the nodal pair p.

A tied pair transfers interface forces and can be evaluated for crack propagation if it belongs to the current crack front. A cohesive pair follows the damage evolution law after the VCCT criterion has been satisfied. An open pair does not transfer tensile or shear traction through the adhesive interface. If the released crack faces penetrate, a compressive penalty response may still be applied, but this response is treated as crack-face contact rather than adhesive load transfer.

For tied pairs, the interface force is not computed explicitly inside the user-element, instead it is recovered from the bulk-solid pass that has already run in the same solver cycle. The reconstruction follows the kinematic-constraint derivation given in Sec. 2.3.3, evaluating Eqs. (2.33)–(2.34) once per pair using the lumped nodal masses read from arrays `dm_xms` and the assembled bulk-solid internal forces read from array `dm_a`. The reconstruction is correct only because LS-DYNA's barrier between the bulk-solid pass and the user-element pass, which guarantees that `dm_a` is fully assembled at every UEL node before the user-element reads it. This barrier also prohibits the use of thick shells (TSHELLS) alongside a solids-based UEL since the thick shells are processed the user-solids.

Alg. 1 summarises the element-level logic during one solver cycle.

Algorithm 1: Element-level force evaluation for one nodal pair

```
foreach interface element do
    foreach local nodal pair p do
        if the current element does not own pair p then
            skip force evaluation for pair p
            continue
        // -- Fully open pair --
        if s_p = 2 then
            apply no tensile or shear traction
            apply compressive contact force only if crack-face penetration occurs
            continue
        // -- Cohesive pair --
        if s_p = 1 then
            evaluate cohesive damage evolution
            apply degraded cohesive interface force
            if complete cohesive failure is reached then
                set s_p ← 2
            continue
        // -- Tied pair --
        if s_p = 0 then
            if pair p is marked as a crack-front pair then
                compute local opening (from neighbouring pairs) and interface
                 force components
                evaluate the VCCT propagation criterion
                if the propagation criterion is satisfied then
                    store release force, mode mixity, and crack-extension area A_i
                    set s_p ← 1
                else
                    apply tied interface force
            else
                apply tied interface force
```

3.2.2 Initial Crack Definition

The initial crack was prescribed by assigning different initial states to different nodal pairs before the first load step. In the coupon-level implementation used in this work, the nominal initial crack length a_0 was used to identify the initially released region together with the initial crack-front pairs, assuming a regular mesh and an element-aligned crack front. Elements with centroids x_{ce} located behind the initial crack front a_0 were initialized as released ($s_p = 2$), while nodal pairs located at the initial crack front were marked as tied front pairs ($s_p = 0$), as illustrated in Fig. 3.3. This

provided the first set of nodal forces and relative openings required for the VCCT evaluation.

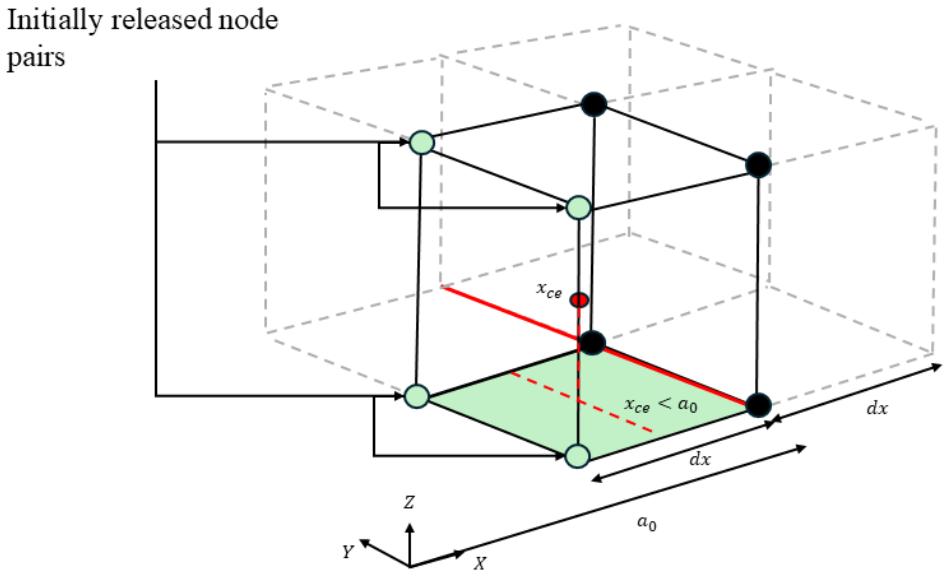

Figure 3.3: Initial crack-front definition in the coupon model. Nodal pairs behind the initial crack length are initialized as open (green), while nodal pairs at the crack front remain tied (black) and are marked for VCCT evaluation.

This initialisation was sufficient for the benchmark specimens considered in this work. However, it is dependent to the coordinate-based definition of the initial crack front.

3.2.3 VCCT Propagation Criterion

For each tied front pair the modal energy release rates G_I, G_{II} and G_{III} are computed from the reconstructed tied interface force and the opening jump of the upstream predecessor (Eqs. (2.17)). The mixed-mode toughness G_c follows the BK interpolation introduced in Section 2.1.2. When $G = G_I + G_{II} + G_{III} \geq G_c$ the pair is transferred from the tied state $s_p = 0$ to the cohesive state $s_p = 1$. At this transition the release-force components F_I^0, F_{II}^0 and F_{III}^0, the mode mixity, and the crack-extension area A_i are stored as elemental history variables per nodal pair. These quantities define the initial condition for the cohesive release law and ensure that the subsequent crack-opening process is linked to the energy release rate evaluated at the instant of propagation.

3.2.4 Cohesive Release Law

After the VCCT criterion has been met, the interface force components are degraded according to the ERRC cohesive law Eqs. (2.36) - (2.38). The damage variable is

updated as the maximum-over-time of the damage-driving function Eq. (2.44) and the failure openings follow Eq. (2.48), so the dissipated work is proportional to the damage variable and equals $A_i G_c$ at full release Eq. (2.49). The damage and cohesive release flow is summarised in Alg. 2.

Algorithm 2: VCCT-triggered cohesive release of an crack-front pair

if *pair p is tied and marked as a crack-front pair* **then**
- compute adjacent local opening components Δ_I, Δ_{II}, and Δ_{III}
- compute local force components F_I, F_{II}, and F_{III}
- evaluate G_I, G_{II}, and G_{III} using the VCCT formulation
- compute $G = G_I + G_{II} + G_{III}$
- evaluate the mixed-mode fracture toughness G_c
- **if** $G \geq G_c$ **then**
  - store release force components F_I^0, F_{II}^0, and F_{III}^0
  - store mode mixity and crack-extension area A_i
  - compute failure openings for the cohesive release law
  - set pair state to cohesive, $s_p \leftarrow 1$
- **else**
  - keep pair tied, $s_p = 0$

if *pair p is cohesive* **then**
- compute current local opening components
- update maximum opening history
- evaluate the damage-driving function d
- update damage variable D
- apply degraded cohesive force components
- **if** $D = 1$ **then**
  - set pair state to open, $s_p \leftarrow 2$

3.3 Single-Core Reference (SCR) Implementation

A lower–upper nodal pair is generally shared between several neighbouring interface elements but maintain separate local history variables. Without proper coordination, multiple element copies could apply tied or cohesive forces to the same nodal pair, leading to duplicated force contributions, incorrect fracture energy dissipation, and inconsistent crack-front evolution. To prevent this, only one element is permitted to evaluate and assemble the interface forces for each physical pair during a given time step.

The single-core implementation enforces this rule with a *first-to-evaluate* principle: a flag associated with each nodal pair that indicates if the nodal quantities have been evaluated. Once a pair has been processed by one element, all non owning neighbouring elements skip it. The flag is shared across elements through LS-DYNA's internal `COMMON`-block memory infrastructure, which stores global state variables accessible by every element in the serial assembly loop. This approach is straightforward to implement and facilitates a quick proof of concept develop-

ment and verification, as it avoids any explicit ownership routine or inter-processor synchronisation.

Following the elemental assembly order, each element reads nodal history variables, evaluates and updates its own nodal states, and stores the result as global variables via writes to `COMMON`-block arrays. The subsequent element repeats the process, skipping shared nodes that have already been processed by the previous element. The flow is illustrated in Fig. 3.4.

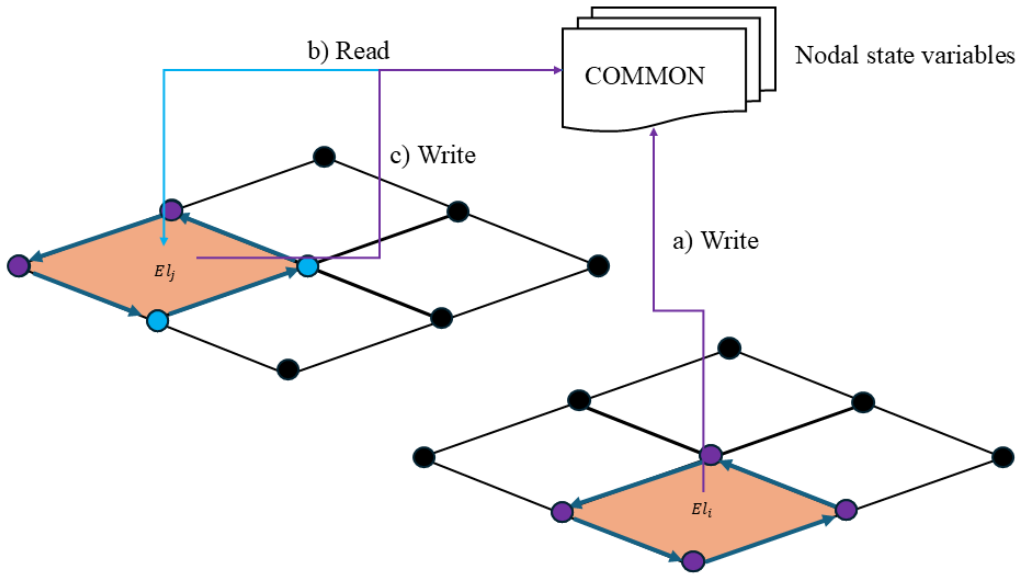

Figure 3.4: Illustration of elemental assembly with intra-element communication and nodal state storage through `COMMON` blocks in the single-core implementation a) the first element evaluating a node, flags the nodal state and stores it in memory b) subsequent elements read the stored nodal states, evaluate their remaining unprocessed nodes, and c) write the updated states to memory.

3.3.1 Single-Core Area Assignment

The single-core implementation assumes mesh regularity with known element area A_e. At initialisation each nodal pair is assigned its own crack-face area depending on its location: interior nodal pairs are assigned the full element area $A_i = A_{el}$, while edge pairs receive half the elemental area $A_j = A_{el}/2$, as illustrated in Fig. 3.5. The assignment is geometric and hard-coded for a regular mesh, it does not adapt to direction-aware crack-extension on distorted or irregular meshes (the directional partition used in the multi-core implementation, Sec. 3.4, removes this restriction). The full algorithm is summarised in Alg. 3.

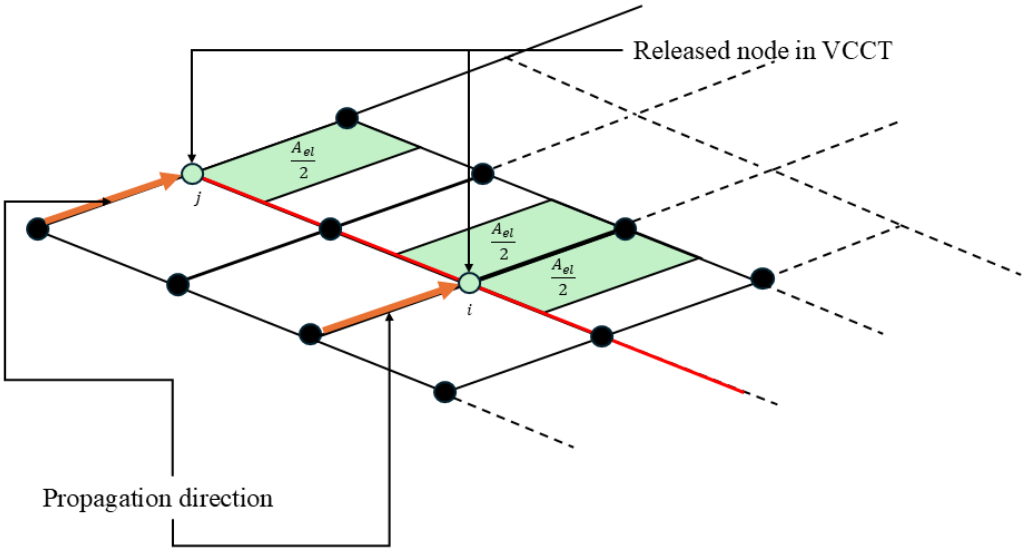

Figure 3.5: Nodal assigned area depending on nodal pair coordinates for the single-core regular-mesh assumption.

Algorithm 3: Nodal area assignment routine

```
if nintcy = 0 then
    foreach interface element do
        foreach local nodal pair p do
            Extract nodal pair coordinates x_i, y_i, z_i
            Read nodal initial state from COMMON/nodal_init
            // -- Identify corner nodes --
            if NOT nodal_init then
                if corner node then
                    Store A_i = A_e/2 in COMMON/nodal_area
                else
                    Store A_i = A_e in COMMON/nodal_area
                Store TRUE in COMMON/nodal_init
            else
                // -- Already initialized --
                skip
```

3.3.2 VCCT and Cohesive Release in Single-Core

Crack growth is evaluated only for tied nodal pairs belonging to crack-front elements, these are identified as elements that contain at least one released pair ($s_p = 1$ or $s_p = 2$). For each tied crack-front pair, the interface force and the relative opening of the neighbouring released pair are taken in the global LS-DYNA frame, as illustrated in Fig. 3.6. The VCCT evaluation and the subsequent cohesive release both use the

global components directly, under the assumption of a mesh aligned such that the Mode I, II and III axes coincide with the global z, x, y directions, respectively. The assumption is valid for the DCB and ENF coupons of this work; a basis-projected variant for arbitrary orientations is described in Sec. 3.4.9.

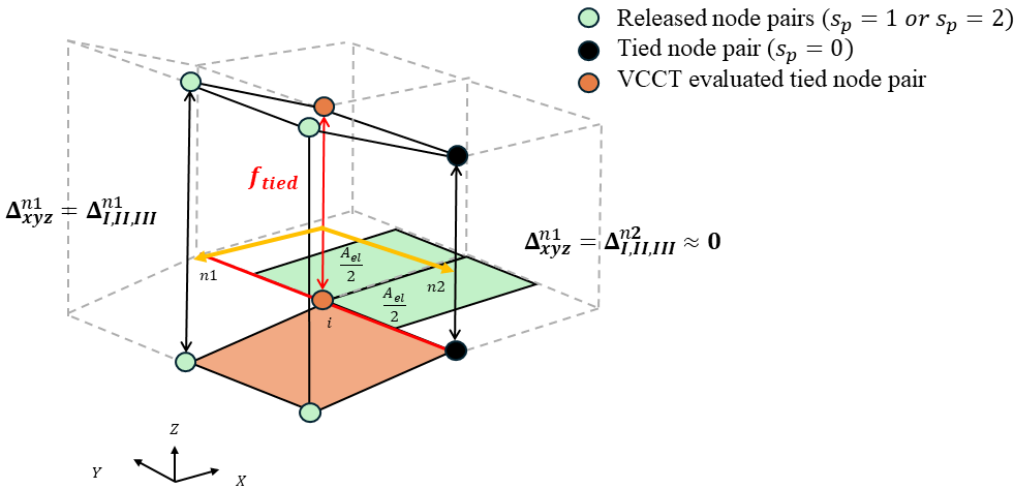

Figure 3.6: VCCT evaluation of a crack-front tied pair using the global components of the neighbouring (n_1 and n_2) released pair's force and opening jump in the single-core implementation.

3.4 Multi-Core Capability (MCC) Implementation

The MCC implementation uses the same nodal-pair mechanics as the single-core reference implementation, but replaces the serial first-to-evaluate ownership rule with a deterministic nonlocal synchronisation routine. This extension is required because, in a shared-memory parallel execution, several element copies may be evaluated by different threads during the same solver cycle. Without a deterministic ownership rule, the force contribution from a shared physical nodal pair could become dependent on thread scheduling.

The nonlocal routine, denoted `cracklook`, is registered through the LS-DYNA keyword `*USER_NONLOCAL_SEARCH`. It is executed once per solver cycle after the UEL pass. For each physical nodal pair, the routine collects the neighbouring elements, selects one owner element, updates the crack-front information, and writes the synchronized state variables back to the history variables of all neighbouring elements. In the following solver cycle, each user element reads only its own history variables and applies mechanics only to the nodal pairs for which it is the owner.

The resulting solver sequence is illustrated in Fig. 3.7. First, the surrounding bulk solid elements assemble their force contributions into LS-DYNA's global force array. Second, the user elements evaluate the interface response using the synchronized

history variables from the previous cycle. Third, `cracklook` performs the nonlocal synchronisation needed for the next cycle. This separation avoids duplicate force application without introducing explicit thread locking inside the UEL routine.

Figure 3.7: One solver cycle with the SMP multi-core implementation. The bulk-solid and user-element passes run in parallel across threads, LS-DYNA's barrier between them guarantees that the bulk contributions to `dm_a` are in place before the user-element reads them. The `cracklook` routine runs once per cycle in serial and stages all cross-element data into history variables for the next cycle (synchronisation).

3.4.1 MPP Execution Scope

The implementation was developed and evaluated using the LS-DYNA massively parallel processing (MPP) executable. In the verification cases considered in this work, the analyses were run using a single Message Passing Interface (MPI) rank. The model was therefore not partitioned across several MPI ranks, and no inter-rank communication of nodal force, mass, state, or ownership data was required.

This distinction is important because the tied-pair force reconstruction relies on access to the assembled nodal force contributions from the surrounding bulk solid elements. In a single-rank run, these contributions are available within the same rank once the bulk-solid pass has completed. In a true multi-rank MPP simulation, however, a nodal pair located at a partition boundary may receive force contributions from elements owned by different MPI ranks. A UEL evaluated on one rank would then not necessarily have direct access to the complete force contribution required for the tied-pair reconstruction.

The present implementation should therefore be interpreted as a deterministic ownership and nonlocal synchronisation strategy demonstrated in the available MPP UEL environment, rather than as a fully validated multi-rank MPP implementation. Extending the method to robust distributed-memory MPP execution would require additional communication of nodal force, mass, state, ownership, and crack-front information across rank boundaries.

Shared-memory parallelism (SMP) is included here only as a useful comparison. In SMP execution, multiple threads share one memory image, so the assembled nodal force contribution is visible to all threads after the bulk-solid pass. In MPP execution, this assumption is not generally valid once the model is partitioned across MPI ranks.

3.4.2 History Variables

The communication between the UEL routine and the nonlocal synchronisation routine is performed through element history variables. Each interface element stores a local copy of the state information for its four nodal pairs. Since one physical nodal pair may be shared by several neighbouring elements, several local history-variable copies may correspond to the same physical interface point.

During the UEL pass (see Fig. 3.7), only the owner element is allowed to update the mechanics of a physical pair. During the subsequent nonlocal synchronisation step, `cracklook` gathers all local copies with the same pair key, reduces the physical state, selects the owner for the next cycle, and writes the synchronized values back to the history variables of all copies. In this way, the UEL routine can remain local in the next solver cycle while still using state information that is consistent across neighbouring elements.

The main history variables used by `cracklook` are summarised in Tab. 3.1.

3.4.3 Pair-Key Identifier

The implementation separates the UEL routine from the nonlocal synchronisation. The user-element routine evaluates its state only for owned nodal pairs, while the nonlocal routine identifies all elements sharing the same physical pair and assigns a unique owner. The physical identity of a nodal pair is defined from the two global node IDs,

$$\mathbf{k}_p = \left(\min(n_p^-, n_p^+), \max(n_p^-, n_p^+)\right), \tag{3.3}$$

where n_p^- and n_p^+ are the lower and upper node IDs of pair p, respectively. Two local pairs are considered copies of the same physical pair only if their pair keys $\mathbf{k}_p$ are identical.

Table 3.1: Main nodal-pair history variables broadcast across elements by `cracklook` in the multi-core implementation.

Quantity	Code variable	Type / values	Purpose
Pair-key	`HKMIN`, `HKMAX`	int tuple $(k_{\min}, k_{\max})$	Physical identity: (min, max) of the two spanning node IDs
Pair state s_p	`HSTAT`	int $\{0, 1, 2\}$	Tied (0), cohesive (1), or open (2)
Owner element	`HOWN`	int (element ID)	Element allowed to apply mechanics
Owner flag	`HISOWN`	flag $\{0, 1\}$	Local force-application flag
Front flag	`HFRONT`	flag $\{0, 1\}$	Marks pair as active crack-front pair
Open-side type	`HFTYPE`	int $\{0, 1, 2\}$	Neighbouring open side
Predecessor pair-key	`HPRED`	int tuple ≥ 0	Pair-key of upstream predecessor; init-filled, release-overwritten
Pair midpoint	`HMIDX`, `HMIDY`, `HMIDZ`	float (coords)	Current bond-plane midpoint; drives inline test and area
Mode mixity	`HBREL`	float $[0, 1]$ (-1 unset)	Stored mode mixity at release
Release force	`HF0I`, `HF0II`, `HF0III`	float (signed)	Force components stored at release
Damage	`HDMG`	float $[0, 1]$	Scalar cohesive damage variable
Local strip area	`HSTRIP`	float ≥ 0	Local topology area contribution
Crack-extension area	`HATRIB`	float > 0	Direction-aware area A_i for VCCT and dissipation

3.4.4 Ownership Logic

Consider an arbitrary interface surface of elements $i, j, k, l, m, n \in \mathbb{N}$, with nodal pairs in the opened ($s_p = 2$), front, and tied states ($s_p = 0$) defined by Eq. (3.2). This topology is illustrated in Fig. 3.8. Ownership is assigned per physical nodal pair through a deterministic priority chain that resolves to a single element each cycle. The priority chain is constructed so that the selected owner can identify the upstream open predecessor from its own corner-cyclic ring of neighbouring pairs, as defined in Eq. (2.54). This is necessary because the element routine needs access to the displacements of nodal pairs contained in the current element. If the predecessor pair lies outside the owning element, the corresponding opening jump cannot be extracted locally for the VCCT evaluation.

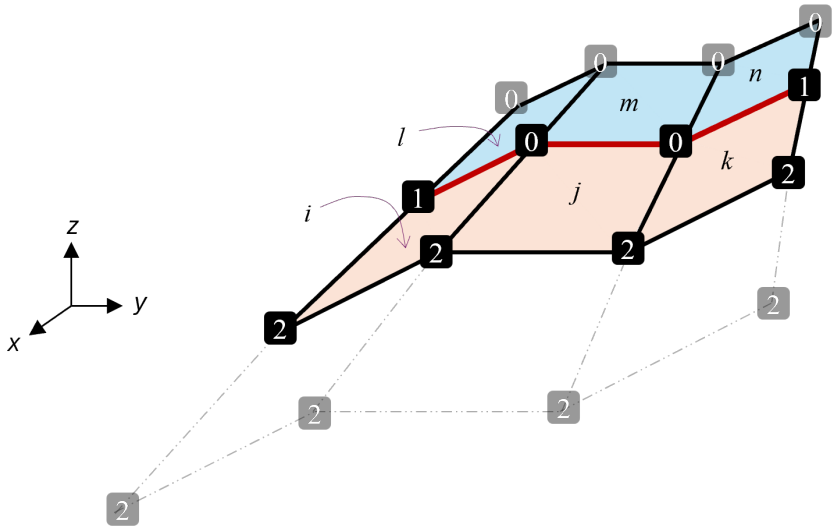

Figure 3.8: Schematic of an interface surface discretised by interface elements. Opened elements, front elements, and tied elements are distinguished by transparent, orange, and blue respectively, to illustrate the crack-front topology. The red line marks the front of the node-pairs.

For a tied front pair p, the preferred owner is the front-active element copy e with the lowest element ID (eID) for which at least one open ring-neighbour $q \in \{\text{neigh}_1(p), \text{neigh}_2(p)\}$ satisfies the inline-and-match acceptance test in Eq. (3.4). The test uses the in-plane offset $\boldsymbol{\delta}_q = \mathbf{m}_q - \mathbf{m}_p$ from Eq. (2.56), the propagation and transverse directions $(\mathbf{e}_{II}, \mathbf{e}_{III})$ defined in Sec. 2.3.4.3, and the pair key $\mathbf{k}_q$ from Eq. (3.3). It combines two conditions:

$$\text{accept}_p(e,q) \iff \underbrace{|\boldsymbol{\delta}_q \cdot \mathbf{e}_{II}| > |\boldsymbol{\delta}_q \cdot \mathbf{e}_{III}|}_{\text{(i) inline}} \;\wedge\; \underbrace{\mathbf{k}_q = \text{HPRED}(p)}_{\text{(ii) match}}. \tag{3.4}$$

Condition (i) rejects ring-neighbours whose offset from p is dominated by the transverse direction. Condition (ii) is an exact pair-key comparison between the candidate neighbour q and the upstream predecessor key pre-identified for p. The directions used in the test are illustrated in Fig. 3.9, where $\mathbf{m}_q$ is the midpoint of the candidate predecessor pair and $\mathbf{e}_{II}$ and $\mathbf{e}_{III}$ are defined by Eq. (2.61).

The initial crack row is handled implicitly. Since the upstream side of the first tied crack-front row is the predefined initial crack region, HPRED(p) is already assigned during `init_hv`. The match condition in Eq. (3.4) is therefore satisfied once `cracklook` detects the corresponding precrack pair as open in the local ring-neighbour topology. This gives the first crack-front row a valid aligned owner.

If no local element copy of p satisfies the inline-and-match test, the ownership rule falls back to a deterministic priority chain. The aligned owner is preferred first. If no aligned owner exists, the lowest-ID front copy (beige elements in Fig. 3.8) is selected. If no front copy exists, the lowest-ID released copy is selected, where a released copy denotes a cohesive or open copy with $s_p \geq 1$. If none of these branches

applies, the lowest-ID copy among all copies is selected:

$$\text{owner}(p) = \begin{cases} \text{owner}_{\text{aligned}}(p) & \text{if a copy passes the inline-and-match gate,} \\ \text{owner}_{\text{front}}(p) & \text{else if a front copy exists,} \\ \text{owner}_{\text{closed}}(p) & \text{else if a closed copy exists,} \\ \text{owner}_{\text{any}}(p) & \text{otherwise.} \end{cases} \tag{3.5}$$

This priority chain ensures that a valid owner always exists. The lower-priority branches act as defensive fallbacks during the precrack-bootstrap cycle and during transient cases where the `HPRED` slot has not yet been populated in all local copies. Since every branch is deterministic, ownership is independent of element evaluation order.

Figure 3.9 illustrates the ownership rule for the subset of elements j, k, m, n around a tied front pair p, with $j < k < m < n$. Elements j and k both contain the open predecessor pair HPRED(p) as a ring-neighbour of p, and can therefore identify p as a front pair. Elements m and n do not contain the predecessor pair in their local ring-neighbour topology and are therefore not preferred owner candidates. Since both j and k are valid aligned candidates, the lowest-ID rule selects j as the owner of the physical pair $\mathbf{k}_p$.

If the current node pair p enters the cohesive state, $s_p = 1$, the owner selected at release is preserved throughout the cohesive and open states. This keeps the cohesive force evaluation tied to the same element reference that initiated the VCCT transition. The owner is also preserved in the open state so that any compressive contact penalty is applied by one element copy only.

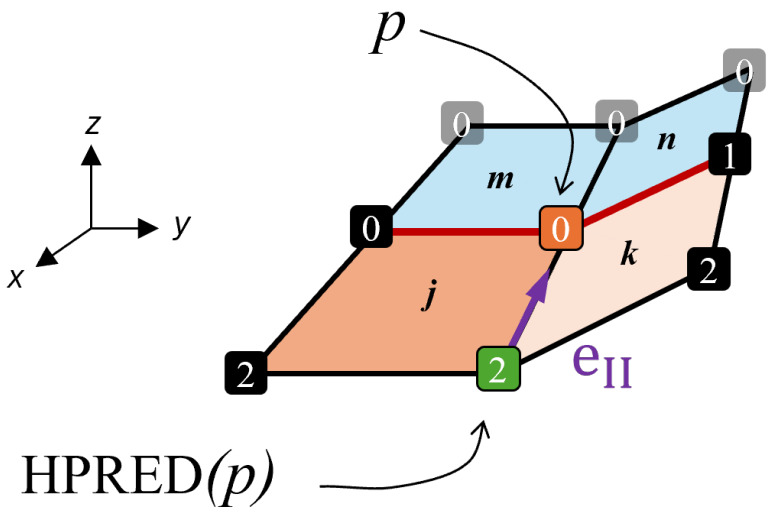

Figure 3.9: Ownership selection for a tied front pair p shared by four element copies j, k, m, n. Elements j and k contain the stored predecessor HPRED(p) as an open ring-neighbour and are therefore aligned owner candidates. Since $j < k$, element j is selected as the owner. Elements m and n do not contain the predecessor as a local ring-neighbour of p and are therefore not preferred candidates.

3.4.5 Predecessor Pre-Identification at Element Initialisation

For a front pair, the VCCT evaluation requires an open predecessor pair from which the relative opening is measured. In the regular DCB and ENF coupon meshes used in this work, the initial crack direction is known in advance. The upstream predecessor of each tied pair can therefore be identified during the initialisation of the history variables.

For each local pair, the two corner-neighbouring pairs are inspected. The neighbour located upstream in the nominal crack-growth direction is stored as a predecessor pair key in `HPRED`. During the analysis, `cracklook` uses this stored key to determine whether an open neighbouring pair is the valid predecessor for VCCT evaluation. If several element copies share the same physical pair, the nonlocal synchronisation step propagates the non-zero predecessor key to all copies.

This initialisation is sufficient for the regular, mesh-aligned DCB and ENF benchmarks considered here. For arbitrary crack fronts or curved bondlines, the predecessor relation should instead be determined from a more general crack-front tracking procedure.

3.4.6 Interface Damping During Cohesive Softening

The viscous damper of Eqs. (2.50) - (2.51) stabilises the released interface during cohesive softening. The damping fraction ζ is supplied as a user-material constant. The Mode-I damping is contact-aware: when the pair is in compressive contact ($\Delta_I < 0$) the penalty stiffness k_p adds in parallel to k_I^0 inside the square root of Eq. (2.51), so the damping coefficient rises automatically when the freed face engages contact.

3.4.7 Front Recompute and Direction-Aware Crack-extension Area

In parallel with the ownership update, `cracklook` recomputes the local front. The front rule introduced in Sec. 3.4.4 is applied per pair after the cross-element state reduction, with the additional inline-and-match gate of Eq. (3.4): a tied pair becomes a front pair whenever at least one of its corner-cyclic ring-neighbours is open and passes both conditions (i) and (ii) of that equation. Lateral open neighbours are rejected by condition (i), since their offset from p is dominated by the transverse direction $\mathbf{e}_{III}$ rather than the propagation direction $\mathbf{e}_{II}$. Open neighbours that are located in the propagation direction but do not match the stored upstream predecessor key HPRED(p) are rejected by condition (ii). Thus, only the open neighbour that is both inline with the local propagation direction and identical to the pre-identified predecessor can activate the front flag for pair p. The integer tag τ_p of Eq. (2.55) records which side(s) passed the gate and is stored in `HFTYPE`(p), where it is used both by the area routine below and by the VCCT branch of the next cycle. The predecessor side that ultimately enters the VCCT trigger is identified directly from the pair-key match condition in Eq. (2.59).

In addition to ownership and the front tag, `cracklook` computes the crack-extension area A_i stored in `HATRIB`(p) for every front pair. The directional partition derived in Section 2.3.4.5 (Eqs. (2.62)–(2.65)) is applied per pair on the owning element using `*USER_NONLOCAL_SEARCH`. The resulting area is illustrated in Fig. 3.10. A_i is held fixed once the pair enters the cohesive state so that the VCCT trigger and the cohesive release law see the same value through the full softening branch.

In parallel with the ownership update, `cracklook` recomputes the local crack front. For each physical nodal pair, the routine first gathers the element copies with the same pair key and synchronises their state by keeping the most advanced state among the copies. Thus, if any copy of the physical pair has become cohesive or open, the same state is written back to all copies before the front is recomputed.

The front rule introduced in Sec. 3.4.4 is then applied to the synchronised pair states. A tied pair becomes a front pair when at least one of its corner-cyclic ring-neighbours is open and satisfies the inline-and-match gate in Eq. (3.4). Condition (i) rejects lateral open neighbours, since their offset from p is dominated by the transverse direction $\mathbf{e}_{III}$. Condition (ii) rejects downstream or otherwise unrelated in-plane open neighbours whose pair key does not match the stored upstream predecessor key HPRED(p). Consequently, a tied pair is marked as front-active only when the open neighbour is both inline with the local propagation direction and identical to the pre-identified predecessor.

The integer tag τ_p of Eq. (2.55) records which side passed the gate and is stored in `HFTYPE`(p). This tag is used by the area routine and by the VCCT branch in the following cycle. The predecessor side that enters the VCCT trigger is therefore identified directly from the pair-key match condition in Eq. (2.59).

In addition to ownership and front status, `cracklook` computes the crack-extension area A_i stored in `HATRIB`(p) for each front pair. The directional partition derived in Sec. 2.3.4.5, Eqs. (2.62)–(2.65), is applied on the synchronized element-copy set associated with the physical pair p. Only contributions located ahead of the front in the selected propagation direction are accumulated (*cf.* A_1 and A_2 in the propagation direction $\mathbf{e}_{II} = \mathbf{d}_1$ in Fig. 3.10). The resulting area is illustrated in Fig. 3.10. Once the pair enters the cohesive state, A_i is held fixed so that the VCCT trigger and the cohesive release law use the same crack-extension area throughout the full softening branch.

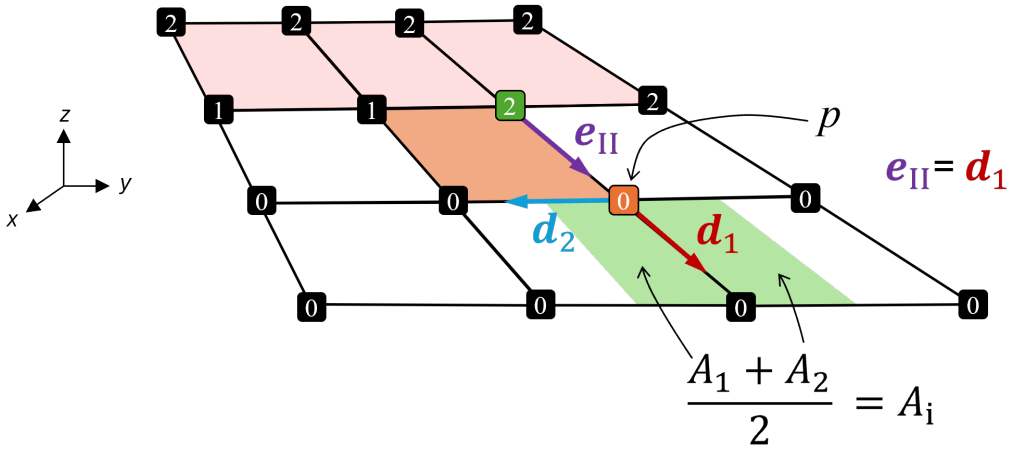

Figure 3.10: Direction-aware crack-extension area associated with a front nodal pair. Only the element contributions located ahead of the crack front are included in A_i.

3.4.8 `Cracklook` Synchronisation Algorithm

The nonlocal synchronisation procedure is summarised in Alg. 4. For every nodal pair p the routine performs seven steps in this order:

1. **Exact-key reduction** — collect all element copies that share the same pair key $\mathbf{k}_p$ from the gathered surrounding list.
2. **State reduction** — reduce s_p to the most advanced state across the copies (open beats cohesive beats tied).
3. **HPRED gather** — reduce $\texttt{HPRED}(p)$ (the upstream-predecessor pair-key tuple) across copies, taking the lowest-eID copy with a non-zero stored key as the consensus value. This propagates the upstream-predecessor identifier from whichever copy resolved it at `init_hv` (Section 3.4.5) or whichever owner wrote it at release time.
4. **Front recompute** — apply the local front rule of Eq. (2.55) together with the inline-and-match gate of Eq. (3.4): p is marked front-active and the tag $\tau_p \in \{1, 2\}$ is incremented only by ring-neighbours that pass both conditions (i) and (ii). The release-time tag is preserved when $s_p \geq 1$.
5. **Owner selection** — the lowest-eID front copy whose open ring-neighbour passed the inline-and-match acceptance test is the aligned owner, if no such copy exists, the fallback chain of Eq. (3.5) resolves a unique owner through the front, closed, then any branches.
6. **Direction-aware area** — when p is tied and front-active, form $\mathbf{d}_1$ and $\mathbf{d}_2$, partition the surrounding elements via Eq. (2.63), select the ahead pool via τ_p (Eq. (2.65)), and apply the never-zero floor.
7. **HSV writeback** — write s_p, $\texttt{HPRED}(p)$, the owner flag, the front flag, τ_p, and A_i into the history variables of every copy of $\mathbf{k}_p$.

This synchronisation step does not apply forces and does not evaluate the VCCT criterion. Its purpose is to make the distributed element copies agree on the physical state of each nodal pair, select the mechanics owner, update the crack-front

description, and provide the direction-aware crack-extension area A_i for the next force-evaluation cycle.

Algorithm 4: Per-cycle nonlocal synchronisation, ownership, and area update (`cracklook`)

foreach *interface element* **do**
  foreach *local nodal pair* p **do**
    form the physical pair key $\mathbf{k}_p = \left(\min(n_p^-, n_p^+), \max(n_p^-, n_p^+)\right)$
    gather all element copies of $\mathbf{k}_p$
    reduce the pair state s_p to the most advanced state among the gathered copies
    reduce the stored predecessor key HPRED(p) across the gathered copies
    recompute the front flag and open-side type τ_p using the synchronised neighbouring states and the inline-and-match gate in Eq. (3.4)
    select the owner using the priority rule in Eq. (3.5)
    if p *is tied and front-active* **then**
      compute the direction-aware crack-extension area A_i
    else
      preserve the previously stored crack-extension area A_i
    synchronise s_p, HPRED(p), owner, front flag, τ_p, and A_i to all copies of $\mathbf{k}_p$

3.4.9 VCCT and Cohesive Release in Multi-core

Crack growth is evaluated only for tied nodal pairs marked as crack-front pairs by `cracklook`. For each owned front pair, the interface force and the relative opening of the chosen upstream neighbour are transformed to the per-pair local crack-front coordinate system illustrated in Fig. 3.11. The same basis is then used consistently for the VCCT evaluation and for the subsequent cohesive release; the per-pair-frame projection replaces the global-axis assumption of the single-core implementation.

The basis $(\mathbf{e}_I, \mathbf{e}_{II}, \mathbf{e}_{III})$ is built on the owning element following Eqs. (2.60)–(2.61): $\mathbf{e}_I$ from the diagonal cross product on the four pair midpoints; $\mathbf{e}_{II}$ from the geometric step $\mathbf{m}_p - \mathbf{m}_{n^\star}$ projected out of $\mathbf{e}_I$, where $n^\star = \mathrm{neigh}_{\mathrm{iftype}_{\mathrm{eff}}(p)}(p)$ is the ring-neighbour whose pair-key matches $\mathtt{HPRED}(p)$; and $\mathbf{e}_{III} = \mathbf{e}_I \times \mathbf{e}_{II}$. The step points from predecessor to p, so F_{II}, Δu_{II} and their cohesive-phase offsets carry a common sign convention, and the cohesive law Eq. (2.37) remains internally consistent.

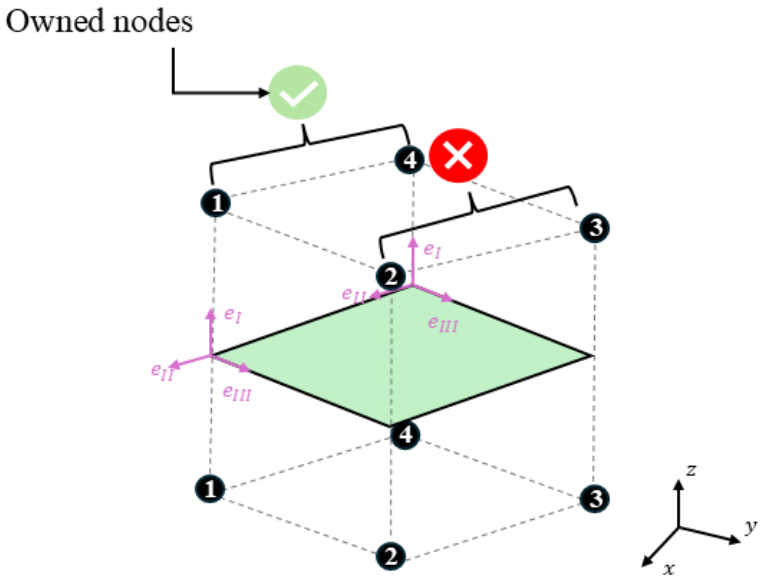

Figure 3.11: Local crack-front coordinate system for an element owning two nodal pairs (pair 1 and 4). For each owned pair, the local mode directions are constructed from the owning elements plane-normal, the crack-growth direction, and their cross product.

4

Numerical Benchmarks

In order to verify the accuracy of the ERRC method, two sets of FEM coupons were established in LS-DYNA. The DCB and ENF coupons were used to compare the global fracture response of the ERRC method against analytical beam-theory solutions (Sec. 2.2) and conventional cohesive-zone modelling (CZM), across a sweep of in-plane mesh discretisations.

4.1 Modelling – LS-DYNA

4.1.1 Element Formulations and Part Ordering in the Keyword Deck

Two element types were used in the bulk–interface stack. The laminate adherends were modelled as linear-elastic orthotropic using eight-node solid elements with one integration point, *i.e.* `ELFORM = -2`. The interface was modelled in one of two ways:

1. A conventional cohesive layer, *i.e.* `*MAT_138` on `ELFORM = 19` cohesive solid elements. This was used as the reference for comparison.`ELFORM = 19` is the standard LS-DYNA eight-node cohesive solid intended for thin (or zero-thickness) interfaces.
2. A user-defined ERR-cohesive element implemented into `ELFORM = 101` and `103`, `NIP=0` with `NHSV = 100` and `LMC = 15`. See Sec. 3.1 for more information.

The adherends were modelled with solid elements rather than as thick-shell (`TSHELL`) elements. Although thick-shells are often computationally efficient for thin laminates, they were not suitable for the present implementation. The ERRC UEL reconstructs the tied-pair force from the assembled nodal force vector `dm_a`. This requires the surrounding adherend contributions to be available before the UEL routine is evaluated. In the LS-DYNA execution order used here, standard solid elements are evaluated before the user-defined solid element, whereas thick-shell contributions are not available to the user element at the required stage of the cycle. The implementation was therefore restricted to solid adherends.

Among the solid parts, the adherend parts must also appear before the UEL interface part in the keyword deck. This ordering ensures that the bulk-solid force contributions have been written to `dm_a` before the user-element pass reads them, cf. Främby [31]. The solver-cycle ordering is described in Chap. 3 and illustrated in Fig. 3.7.

4.1.2 CZM Reference Setup

The conventional CZM reference cases used `*MAT_138` bilinear cohesive elements with the bilinear traction–separation law and BK criteria as described in Sec. 2.3.2. The material model supports both zero-thickness (thin) and finite-thickness (thick) cohesive layers through the `ROFLG` flag. With `ROFLG=0`, the user-specified density is interpreted as mass per unit volume and is appropriate for finite-thickness cohesive elements. With `ROFLG=1`, the density is interpreted as mass per unit area and is used for cohesive elements with zero initial volume [32].

The CZM reference cases in this work used the zero-thickness configuration with `ROFLG=1`. This choice was made to keep the conventional CZM benchmark consistent with the analytical DCB and ENF reference solutions in Sec. 2.2, where the crack is represented as a mid-plane interfacial separation without an additional finite adhesive thickness. The implementation of the UEL suffered from numerical instabilities for smaller elemental thicknesses $t < 0.001$ mm, therefore, it was implemented as a finite-thickness solid interface element with $t = 0.12\,\text{mm}$ to showcase its ability to represent a physical adhesive thickness [33]. The comparison should therefore be interpreted as a comparison between a conventional zero-thickness CZM reference and the proposed finite-thickness user-element formulation, rather than as two geometrically identical adhesive-layer models.

The FPZ for brittle ply interfaces is typically shorter than the element edge used in industrially relevant meshes. For this reason Turon *et al.* [11] have proposed an engineering solution to accommodate large element sizes. In their proposed method, the cohesive strength is reduced[1] such that the FPZ can be sufficiently resolved by the required mesh size. In this work, apart from using the nominal material parameter values, the DCB specimens with mesh sizes of 2×2 and 4×4 mm^2 were analysed with reduced strength values in Mode I (X_I), such that the FPZ length predicted by Eq. (2.25) spanned at least three elements at the chosen mesh size, as seen in Tab. 4.1. Initial interface stiffnesses K_n and K_s were sized from the Turon relation in Eq. (2.28), and presented in Tab. 4.2.

Table 4.1: Comparison between the element length, strength X and the size of the FPZ for the compared DCB modelling strategies.

Method	X_I **(MPa)**	**FPZ (mm)**	**Element length (mm)**
ERR-Cohesive (Present)	–	–	2, 4 and 6
CZM typical X_I	30	0.80	0.25, 2 and 4
CZM reduced X_I	2.05	6.25	2
CZM reduced X_I	0.54	12.5	4

Table 4.2: Interface stiffness used for the different coupon configurations.

Coupon	K_n [MPa/mm]	K_t [MPa/mm]
DCB	33900	15300
ENF	76000	35000

[1]The reduction should be made with caution so that the global results remain largely unaffected

4.1.3 User-Element: Modelling

The ERRC interface was represented by the user-defined solid element described in Sec. 3.1. Since LS-DYNA requires a material card for the user-defined solid section, a dummy `*MAT_ELASTIC` card was assigned to the interface part. This in order for LS-DYNA to calculate an estimated but not needed time-step. The density is allowed to represent a physical adhesive mass. The Young's moduli and Poisson's ratio should be chosen such that the calculated timestep is larger than the one estimated by the solver from adherand bulk elements, model geometry and contacts, in order to not unnecessary slow down the computational time. This is most easily done by equalising the unused Young's moduli to the density value, resulting in a timestep of $dt \approx 1$.

For the MCC implementation, the additional include file `nonlocal.k` dispatched the per-cycle nonlocal synchronisation routine `cracklook` through the LS-DYNA keyword `*USER_NONLOCAL_SEARCH`. The relevant deck parameter is the search radius R. If R is too small, the routine fails to identify all element copies associated with a physical nodal pair, which breaks deterministic ownership and tied-pair force assembly. If R is unnecessarily large, the surrounding-element search becomes more expensive without changing the result.

For the regular meshes used in this work, the search radius was chosen based on the in-plane element size and selected large enough for `cracklook` to identify all element copies associated with a physical nodal pair and its neighbouring topology. The complete deck structure, include-tree layout, keyword cards, and radius-selection procedure are documented in App. B.

The cohesive-release stiffnesses k_α^0 used in the ERRC law were selected according to Daniel *et al.* [1]. The values used in the DCB and ENF models are summarised in Tab. 4.3.

Table 4.3: Cohesive-release stiffnesses k_α^0 used in the ERRC formulation.

	Element Edge (mm)	$k_I^0 = k_{II}^0$ (kN mm^{-1})
DCB	2 mm	62.2
$\alpha = I$	4 mm	248.9
	6 mm	498.3
ENF	2 mm	86.6
$\alpha = II$	4 mm	346.5

4.2 Coupons – DCB & ENF

The DCB and ENF models were constructed such that the interface element representing the adhesive layer could be modelled by both the conventional CZM (`ELFORM = 19`) and the ERRC method implemented in the UEL.

The reaction force was monitored at the loaded node set in order to generate load–displacement curves for verification against analytical CBT solutions in Sec. 2.2. Frictionless surface-to-surface contact was defined between the crack faces to prevent

interpenetration during crack propagation while avoiding artificial shear transfer across the debonded region.

The laminate material properties used in the coupon models are summarised in Tab. 4.4.

Table 4.4: The considered laminate material properties [9]. Configuration according to Fig. 2.2.

Parameter	**DCB**	**ENF**	**Units**
Material	T300/1076	IM7/8552	–
Layup per adherend arm	$[0]_{12}$	$[0]_{12}$	–
E_{11}	139.4	161.0	GPa
$E_{22} = E_{33}$	10.16	11.38	GPa
$G_{12} = G_{13}$	4.6	5.2	GPa
G_{23}	3.54	3.9	GPa
$\nu_{12} = \nu_{31}$	0.3	0.32	-
ν_{23}	0.436	0.45	-
G_{Ic}	0.170	0.212	kJ/ m^2
$G_{IIc} = G_{IIIc}$	0.494	0.774	kJ/ m^2
η	1.62	2.1	-
ρ	1500	1500	kg/ m^2

4.2.1 DCB

The DCB specimen represents a pure Mode I opening configuration. The geometric parameters and laminate layup used in the present study are summarised in Tab. 4.5, and the schematic of the model with prescribed-displacement boundary conditions is shown in Fig. 2.2 and 2.3.

The model consists of two symmetric laminate arms separated by a predefined mid-plane interface. The initial crack of length a_0 (Tab. 4.5) was introduced by leaving the interface unconnected over the pre-crack region and tied (CZM-side) or fully tied (ERRC-side) elsewhere. The adhesive thickness t in the table reflects the two interface options: zero thickness for the CZM reference, and 0.12 mm for the ERRC case following the previous discussion.

Boundary conditions were applied such that the DCB was fully constrained at one end in all translational degrees of freedom. A vertical displacement was prescribed at the corresponding node set at the end of the upper and lower arm, generating symmetric opening of the specimen, cf. Fig. 2.3.

Table 4.5: Geometry and layup of the DCB benchmark specimen (pure Mode I).

Parameter	**Value**
Width b	12 mm
Adherend thickness h	3 mm
Total length $2L$	64 mm
Initial crack length a_0	20 mm
Adhesive thickness t	0 (**CZM**), 0.12 (**ERRC**) mm

4.2.2 ENF

The ENF specimen represents a pure Mode II shear-dominated fracture configuration. The geometric parameters and laminate layup are summarised in Tab. 4.6, and the corresponding model is illustrated in Fig. 2.2 and 2.4. The model consists of a pre-cracked laminate subjected to three-point bending. The initial crack of length a_0 was introduced by not placing any CZM or UEL elements along the predefined mid-plane region.

Simply supported boundary conditions were applied at the two lower support locations. The supports were constrained in the vertical direction while allowing in-plane movement to avoid artificial axial restraint. Vertical displacement was prescribed at the central loading point to generate bending-induced shear at the crack tip, cf. Fig. 2.4.

Table 4.6: Geometry and layup of the ENF benchmark specimen (pure mode II).

Parameter	**Value**
Width B	20 mm
Adherend thickness h	2.25 mm
Half-span L	50 mm
Total length $2L$	100 mm
Initial crack length a_0	34 mm
Adhesive thickness t	0 (**CZM**), 0.12 (**ERRC**) mm

4.2.3 LS-DYNA Models

The LS-DYNA models were created according to configurations seen in Fig. 2.2 - 2.4 for regular in-plane meshes. A preliminary through-thickness discretisation study was performed for the DCB and ENF adherends to obtain a practical compromise between bending stiffness accuracy and computational cost. The study compared the force–displacement response of a simple cantilever beam model against Euler–Bernoulli beam theory. The comparison showed a difference of approximately 14 % in resultant force between the two-element and five-element through-thickness discretisations. This difference was not negligible, but the five-element model produced a disproportionate increase in computational cost. A reduced through-thickness discretisation was therefore adopted as a practical compromise: two solid elements through the total thickness for the DCB model and three solid elements through the total thickness for the ENF model. Examples of the in-plane discretisations are shown in Figs. 4.1 - 4.2.

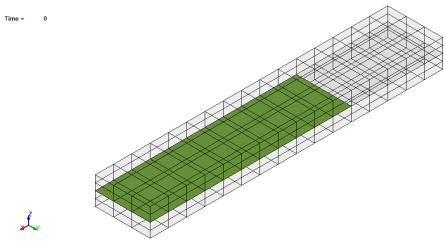

(a) DCB 4x4mm^2 in-plane mesh.

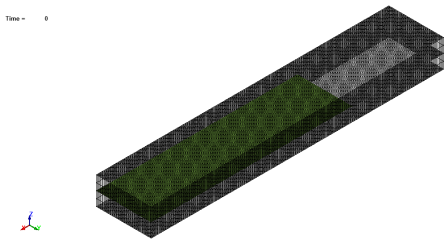

(b) DCB 0.25x0.25mm^2 in-plane mesh.

Figure 4.1: DCB model for two different in-plane discretisations.

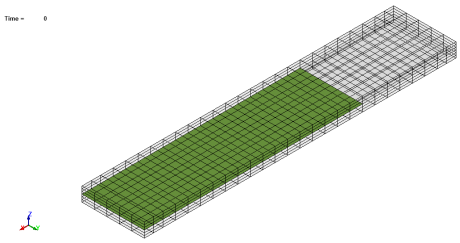

(a) ENF 4x4mm^2 in-plane mesh.

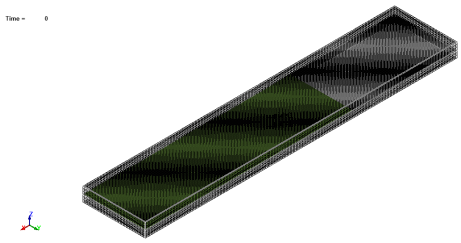

(b) ENF 0.5x0.5mm^2 in-plane mesh.

Figure 4.2: ENF model for two different in-plane discretisations.

In order for the ERRC-method to initiate, the adhesive layer has to be discretised in a way that allows the crack front elements to have a state of open nodes in front of a crack and tied nodes at the crack. This was done by adding a sacrificial element extending from $a_0 - dx$ to a_0 as indicated in Fig. 4.3

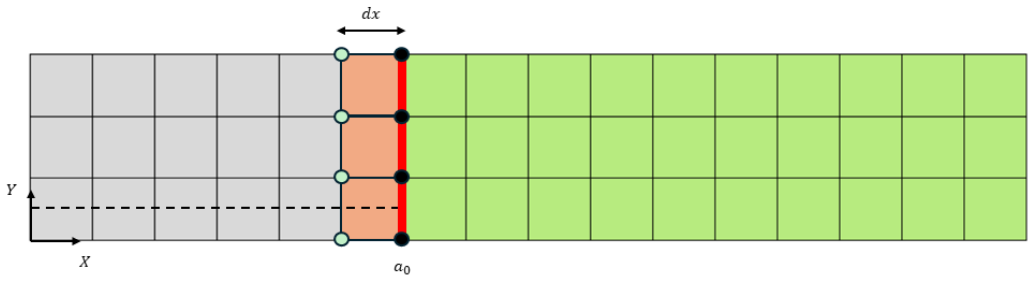

Figure 4.3: Adhesive interface discretisation indicating the use of sacrificial elements (orange) aligned at the crack front in order to initiate the ERRC-method.

4.2.4 Loading and Time Stepping

Loading was applied as a smoothed haversine displacement ramp on the prescribed-motion node set. The haversine shape was used to alleviate the high-frequency transients that a linear ramp could introduce at $t = 0$. The simulations were defined over a period where $t_{end} = 0.01\,\mathrm{s}$.

The resulting average loading rate is faster than the rate typically prescribed in quasi-static delamination analyses, including the values used by Daniel *et al.* [1, 2].

The shorter ramp was selected to keep the implementation study computationally feasible and to allow repeated model iterations within the project time frame. As a global check, the kinetic-to-internal-energy ratio was monitored during the simulations. A low global kinetic-energy ratio indicates that the pre-fracture elastic response is approximately quasi-static. However, this criterion should be interpreted as a global indicator only. It does not guarantee that the local tied-to-cohesive release events remain quasi-static after crack propagation has started.

4.2.5 Verification Metrics

Three metrics were computed per run and per mesh to evaluate the implementation:

1. **Global load–displacement response.** The reaction force at the loaded node set was recorded as a function of the prescribed displacement. The peak load, initial stiffness, and post-peak compliance evolution were compared against the analytical CBT solutions for the DCB and ENF specimens from Sec. 2.2. They were also compared against the conventional CZM simulations.
2. **Crack-length evolution.** The damage variable D along the bond was extracted from the per-pair history in the SMP multi-core approach, and was compared against the analytical CBT crack length $a_{\mathrm{CBT}}(u)$ obtained from the compliance relation in Sec. 2.2.
3. **Computational efficiency**. In order to estimate the computational efficiency of the multi-core implementation the following metrics were used: 1. Total CPU runtime, 2. Average stable timestep and 3. Total number of nodes. The average stable timestep was computed as

$$\Delta t_{\mathrm{avg}} = \frac{T}{N_{\mathrm{cycles}}}$$

where T is the total simulated physical time and N_{cycles} is the total number of explicit integration cycles.

5

Results and Discussion

This chapter presents the numerical results for the conventional CZM, the single-core reference (SCR) implementation, and the multi-core capable (MCC) implementation. The DCB and ENF configurations are evaluated according to the benchmark definitions in Chap. 4. First, the conventional CZM results are used to establish the reference behaviour and the mesh-size sensitivity of a standard CZM approach. The SCR and MCC implementations are then compared against the analytical CBT solutions and the finest CZM references through force–displacement response and crack-length evolution. Finally, the interface energy dissipation is examined for the MCC implementation before the DCB and ENF results are summarised in terms of release behaviour and computational efficiency.

5.1 Cohesive Zone Modelling

The conventional CZM provides the validation reference for the user-element implementation. The DCB force–displacement response, evaluated for three different in-plane discretisations, ranging from 4×4 to $0.25 \times 0.25\,\mathrm{mm}^2$, are compared against the analytical solution in Fig. 5.1a. Note results using the reduced material properties according to Tab. 4.1 are also shown in the figure. The corresponding ENF response, evaluated over four discretisation from 4×4 down to $0.5 \times 0.5\,\mathrm{mm}^2$, is shown in Fig. 5.1b.

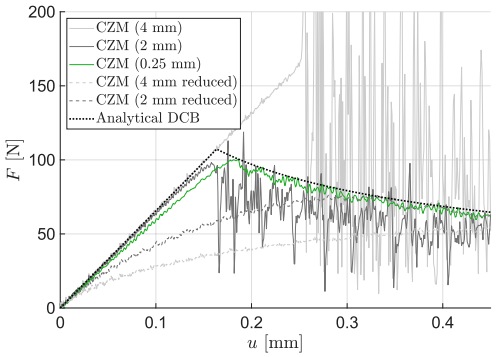

(a) DCB results for ordinary CZM of 4×4, 2×2 and 0.25×0.25 mm^2 in-plane discretisation.

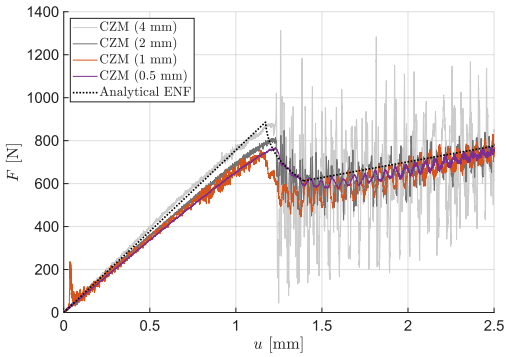

(b) ENF results for ordinary CZM of 4×4, 2×2, 1×1 and 0.5×0.5 mm^2 in-plane discretisation.

Figure 5.1: Comparison of ordinary CZM results for DCB and ENF specimens with different in-plane discretisations.

The nominal CZM results show the expected mesh sensitivity. When the interface element size is too large relative to the fracture process zone, the peak load and propagation response are not captured accurately. Refining the mesh improves the agreement with the analytical response, with the finest nominal CZM models giving the closest reference solutions. The reduced-strength CZM cases improve the coarse-mesh response by artificially increasing the effective FPZ, following the engineering strategy described in Sec. 4.1.2. These cases are therefore useful as coarse-mesh engineering references, but they should not be interpreted as using the physical cohesive strength of the adhesive.

The finest nominal CZM result is retained as the main reference against which the UEL implementation is compared. The reduced-strength CZM cases are used only to illustrate how conventional CZM can be made less mesh-sensitive by modifying the cohesive strength.

A discrepancy in the elastic pre-fracture stiffness should be noted before the comparisons that follow. In an study modelling the DCB and ENF the adherends with a single thick-shell element through-the-thickness, reproduced the analytical pre-fracture stiffness almost exactly. The user-element, however, requires solid adherends, as previously discussed in Chap. 3. The use of solid elements does not recover the bending stiffness as faithfully. The consequence is a systematic under-prediction of the initial elastic stiffness relative to both the analytical solution and the thick-shell model. This offset is purely elastic and it does not affect the dissipated fracture energy but will slightly alter the propagation response. This should be kept in mind when reading the elastic branches of the force–displacement comparisons in the following sections.

5.2 Single-Core Reference Results

Figs. 5.2a and 5.2b present the force–displacement response for the DCB coupon using the single-core implementation. Fig. 5.2a compares three different in-plane discretisations of 6×6, 4×4 and 2×2 mm^2 against the analytical solution. Fig. 5.2b overlays the corresponding responses with the finest CZM reference model with a mesh size of 0.25×0.25 mm^2 and the 4×4 mm^2 reduced stiffness. The single-core DCB response aligns closely with the analytical solution for each in-plane mesh. It accurately captures the elastic pre-fracture region and the fracture point better than the CZM 0.25×0.25 mm^2 reference. The implemented user-element manages to propagate the crack and oscillates in a stable fashion around the analytical solution.

Figs. 5.3a and 5.3b present the corresponding results for the ENF coupon. Two different in-plane discretisations of 4×4 and 2×2 mm^2 are compared against the analytic solution and the finest cohesive-zone reference model with a mesh size of 0.5×0.5 mm^2. As previously discussed, it is challenging to recover the appropriate initial elastic stiffness. The SCR ENF responses for discritisations of 4×4 and 1×1 mm^2 also show difficulties in capturing post-fracture force and have a highly dynamic and oscillatory behaviour after fracture. However, the results do show some decrease in dynamic oscillations for a finer mesh. For clarity in Fig. 5.3b, the force response for the 4×4 mm^2 mesh was filtered using an SAE J211 low-pass filter with a cutoff frequency of 1000 Hz in LS-PrePost, reducing high-frequency numerical oscillations.

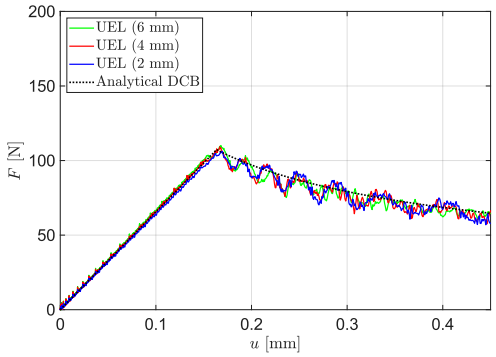

(a) Single-core UEL implementation for 6×6, 4×4 and 2×2 mm^2 in-plane discretisation for a DCB.

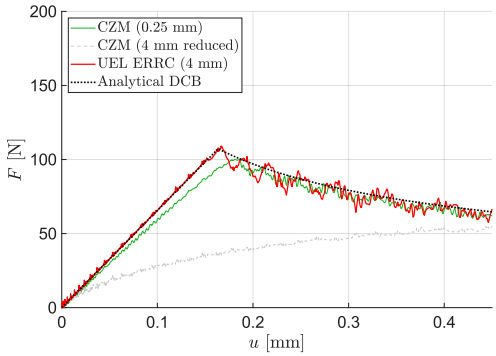

(b) Single-core UEL implementation compared to ordinary CZM for regular mesh DCB.

Figure 5.2: DCB results for the single-core UEL implementation.

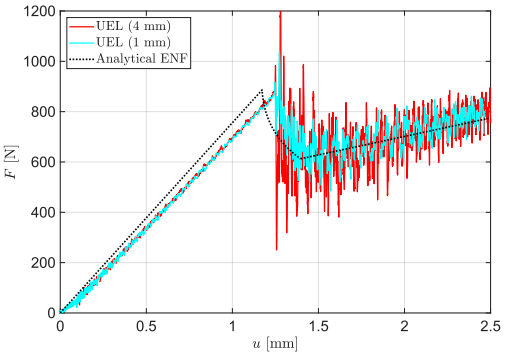

(a) Single-core UEL implementation for $4{\times}4$ and $1{\times}1$ mm^2 in-plane discretisation for an ENF specimen.

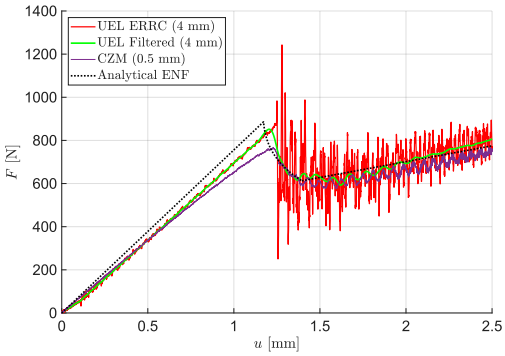

(b) Single-core UEL implementation compared to ordinary CZM for a regular mesh ENF specimen with included filtered response.

Figure 5.3: ENF results for the single-core UEL implementation.

5.3 Multi-Core Capability Results

5.3.1 Force-Displacement

The global force–displacement response of the multi-core UEL implementation is shown in Figs. 5.4-5.5 for both the DCB and ENF samples. For each specimen the response is reported first across a mesh-convergence series. In the case of the DCB specimens, results are shown for 6×6, 4×4 and 2×2 mm^2. For the ENF samples they are shown for in-plane discretisations of 4×4 and 2×2 mm^2. This results section is then concluded with a comparison directly against the finest cohesive-zone CZM reference for both specimen types. An interface damping coefficient of $\zeta = 0.1$ for the DCB models was used while a coefficient of $\zeta = 0.2$ was considered for the

ENF models. The larger damping coefficient for the ENF cases reflect the more aggressive shear-release dynamics.

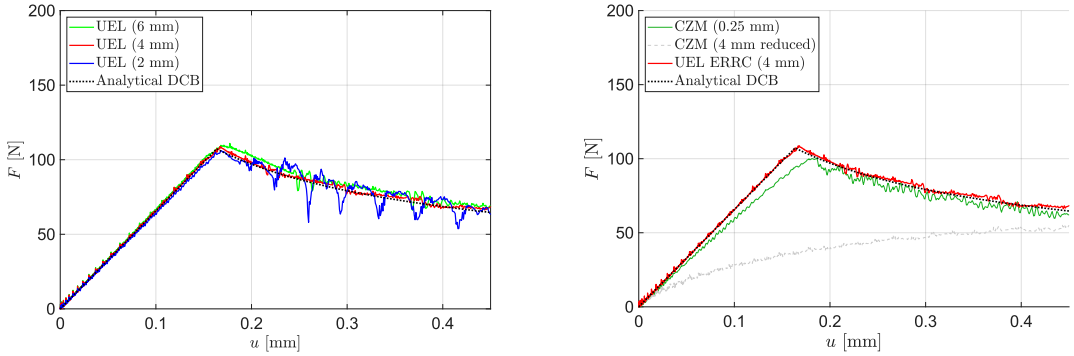

(a) Multi-core UEL implementation for 6×6, 4×4 and 2×2 mm^2 in-plane discretisation for a DCB specimen.

(b) Multi-core UEL implementation compared to ordinary CZM for a regular mesh DCB specimen.

Figure 5.4: DCB results for the multi-core UEL implementation.

The DCB response in Fig. 5.4 follows the expected Mode I behaviour. The response is initially linear, reaches the onset of crack propagation, and then decreases as the crack advances and the specimen compliance increases under displacement control. The $6 \times 6\,\mathrm{mm}^2$ and $4 \times 4\,\mathrm{mm}^2$ meshes give similar global responses and remain close to the analytical and CZM references. The $2 \times 2\,\mathrm{mm}^2$ case, however, contains local force spikes that are not present in the single-core reference in Fig. 5.2. This indicates that the MCC ownership and front-update logic can affect the timing of individual release events even in the otherwise stable Mode I case. The $4 \times 4\,\mathrm{mm}^2$ model is therefore used as the main MCC DCB reference in the crack-length and dissipation analysis below.

The ENF response in Fig. 5.5b captures the general trend of the analytical and CZM references, but the raw force–displacement response contains pronounced high-frequency oscillations after crack propagation starts. These oscillations are interpreted as a dynamic effect of the discrete tied-to-cohesive release process. In Mode II, the interface is shear-loaded and the release of one pair can rapidly change the local force distribution seen by the next front pair. This gives a more abrupt propagation sequence than in the DCB case.

A low-pass filtered force response is included in Fig. 5.5b to show the underlying load–displacement response. The filtered response follows the expected trend more clearly. The increased damping coefficient, $\zeta = 0.2$, reduces the oscillation amplitude but does not remove the dynamic response at the loading rate used here. The energy consequences of this behaviour are examined separately in Sec. 5.3.3.

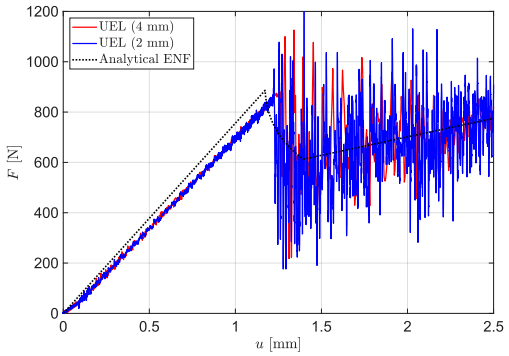

(a) Multi-core UEL implementation for 4×4 and $2\times2\,\mathrm{mm}^2$ in-plane discretisation for an ENF specimen.

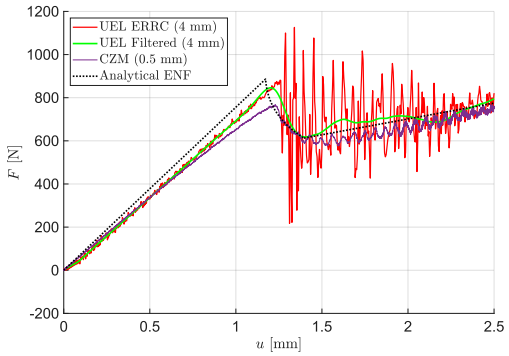

(b) Multi-core UEL implementation compared to ordinary CZM for a regular mesh ENF specimen with included filtered response.

Figure 5.5: ENF results for the multi-core UEL implementation.

5.3.2 Crack-Length Evolution

The bond-plane damage field for the DCB $4 \times 4\,\mathrm{mm}^2$ coupon is shown in Fig. 5.6 at six prescribed displacement levels, $u \in \{0.4, 0.5, 0.7, 0.8, 0.9, 1.0\}$ mm. Each panel shows the per-pair damage variable D on the bond plane. Fully tied regions are left uncoloured, while partially damaged and fully released regions are coloured according to their damage level. The black dashed line marks the analytical CBT crack length $a_{\mathrm{CBT}}(u)$, obtained from Eq. (2.12).

The black dashed line in each panel marks the analytical CBT crack-tip position $a_{\mathrm{CBT}}(u)$ obtained by inverting Eq. (2.12). At each visualised displacement level the analytical prediction falls inside the numerical cohesive-zone and sits roughly half an element edge ahead of the fully-released ($D \to 1$) boundary. Further this consistent offset can be attributed to the row-by-row discrete release inherent to the mesh-aligned bond.

The DCB crack front follows the analytical reference closely throughout the sampled displacement range. The fully released boundary, $D \to 1$, remains behind the analytical crack-tip position by approximately the width of the active cohesive-release row. This offset is expected, since the implemented method releases one discrete row of nodal pairs at a time. The offset does not grow during propagation, which indicates that the VCCT trigger and cohesive release law advance the crack at a rate consistent with the CBT reference.

Equally important is that the DCB crack front is row-uniform. Across the width $y \in [0, 12]$ mm all except one panel shows a single contiguous cohesive-zone of constant x-position. Pure Mode I bending produces a near-symmetric opening jump along the bond normal at every front pair, so every row in a given x-column reaches the VCCT trigger with only one, or a few cycles in-between. The edge nodal pairs reach the VCCT criteria before the midline ones (cf. Fig. 5.6e). The pair basis is constructed from the geometric node-position offset of Eq. (2.61) and is therefore unambiguous on the regular coupon mesh independently of the in-plane loading

state. The row-uniform, CBT-consistent front in Fig. 5.6 confirms that the VCCT–cohesive coupling is sound when the loading is Mode I-dominant. However, the VCCT is initiated at the edges across the width. This is in contrast to Daniel *et al.* [2] which saw the opposite triggering order. This could come from the difference in crack-path algorithm where the current method, introduced in Sec. 3, has no damage smearing capability in contrast to [2].

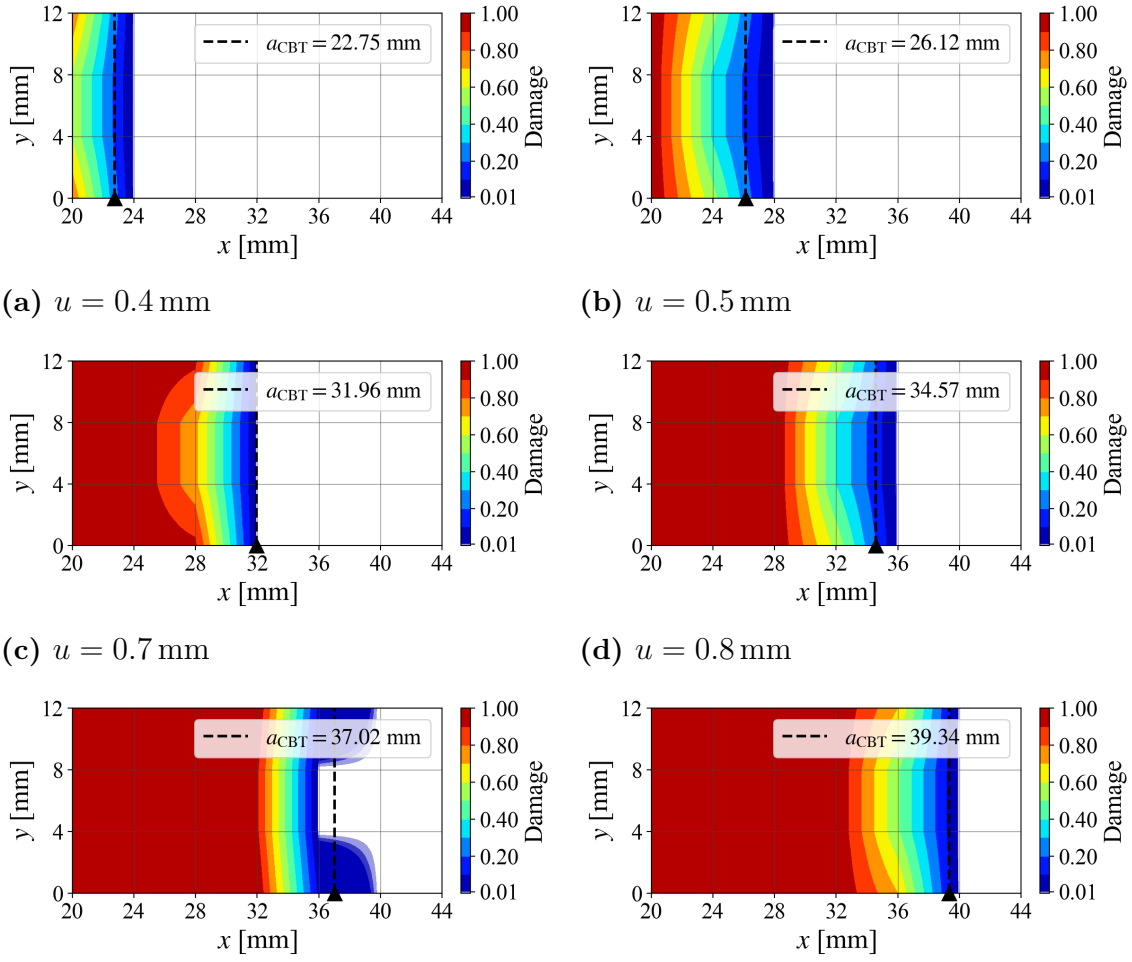

(a) $u = 0.4\,\text{mm}$ **(b)** $u = 0.5\,\text{mm}$

(c) $u = 0.7\,\text{mm}$ **(d)** $u = 0.8\,\text{mm}$

(e) $u = 0.9\,\text{mm}$ **(f)** $u = 1.0\,\text{mm}$

Figure 5.6: Bond-plane damage field on the DCB 4×4 mm^2 coupon at six vertical displacements levels. Black dashed line marks the analytical corrected beam theory crack length $a_{\text{CBT}}(u)$

The bond-plane damage field for the ENF $4 \times 4\,\text{mm}^2$ coupon is shown in Fig. 5.7 at six prescribed displacement levels, $u \in \{1.20, 1.35, 1.395, 1.50, 2.00, 2.50\}$ mm. For the first three displacement levels, the analytical CBT crack length $a_{\text{CBT}}(u)$ from Eq. (2.15) is included as a reference. At larger displacements the predicted crack length reaches the midspan region, where the standard ENF CBT expression is no longer admissible as a direct crack-length reference. The analytical line is therefore omitted from the later panels.

Three pre/at-midspan panels and three post-midspan panels are shown so that

the front advance can be tracked both within and beyond the regime where the analytical reference (Eq. 2.15) is admissible.

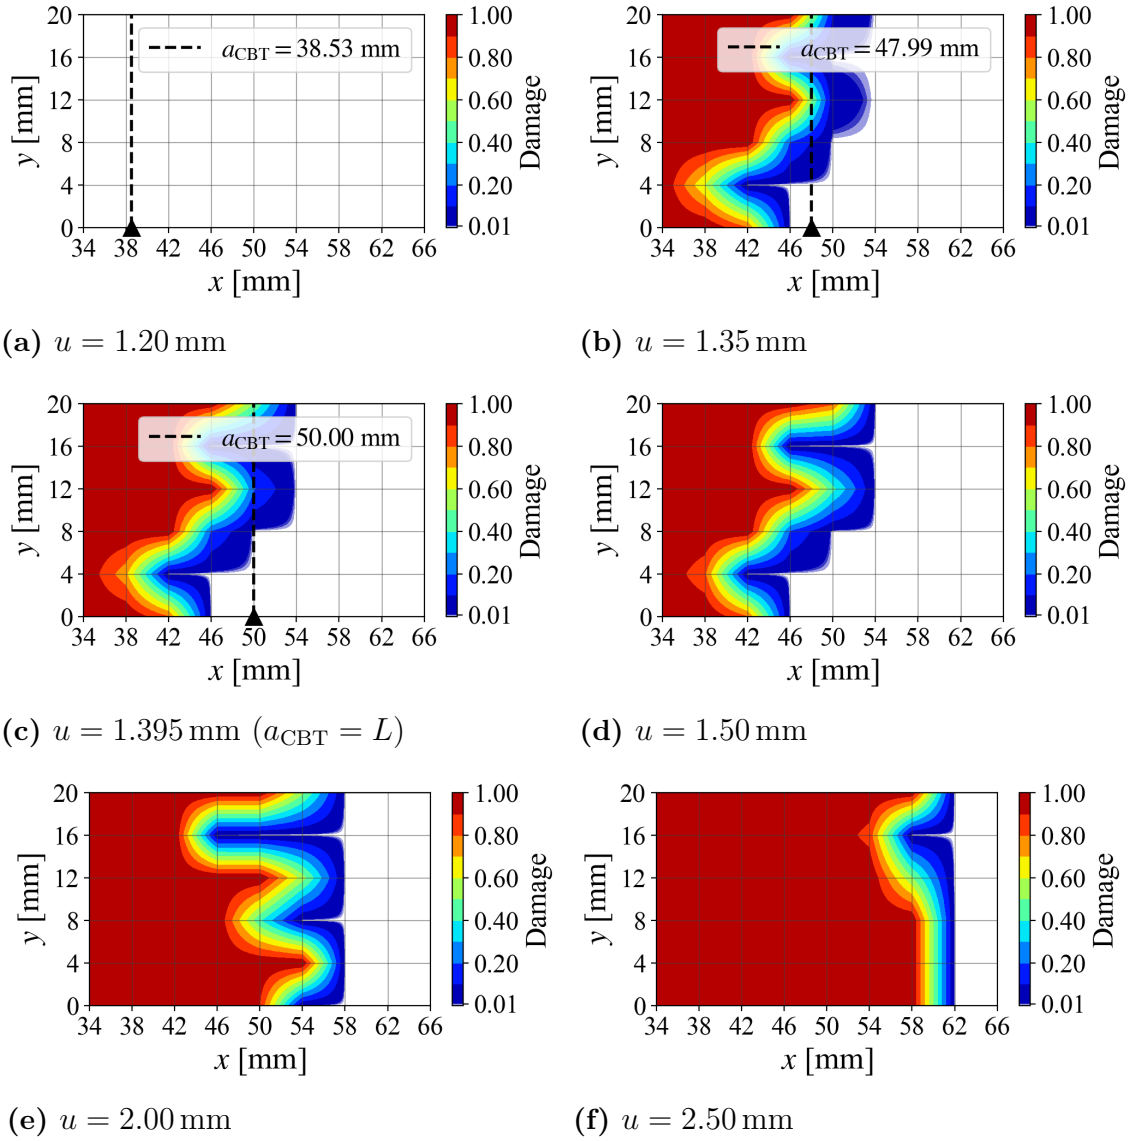

(a) $u = 1.20\,\text{mm}$ **(b)** $u = 1.35\,\text{mm}$

(c) $u = 1.395\,\text{mm}$ $(a_{\text{CBT}} = L)$ **(d)** $u = 1.50\,\text{mm}$

(e) $u = 2.00\,\text{mm}$ **(f)** $u = 2.50\,\text{mm}$

Figure 5.7: Bond-plane damage field on the ENF 4×4 mm^2 coupon at six vertical displacements levels. The black dashed line marks the analytical corrected beam theory crack length $a_{\text{CBT}}(u)$ given by Eq.((2.15)) for $a_{\text{CBT}} \leq L$. Past the midspan, the CBT solution is no longer admissible. The analytical line is therefore omitted in the later panels.

The ENF model produces a coherent propagation chain in the intended in-plane direction. The crack front advances through the interface without clear evidence of lateral jumping between neighbouring rows, which indicates that the inline-and-match gate in Eq. (3.4) enforces the intended predecessor relationship. The main discrepancy is therefore not the direction of propagation, but the dynamic character of the release process.

Across the specimen width, the front develops asymmetrically, as seen in Fig. 5.7.

One free edge advances ahead of the other during the early stages of propagation. As the opposite side catches up, the front becomes jagged before becoming more uniform again at larger displacements. This behaviour is not expected for the idealised ENF benchmark and should not be interpreted as a physically validated crack front shape. Instead, it indicates that the current Mode II implementation is sensitive to the timing of local release events and requires further stabilisation.

5.3.3 Interface Energy Dissipation

The force–displacement and damage-field results show that the DCB and ENF cases differ mainly in how the crack front evolution is predicted. The DCB case produces a nearly row-uniform Mode I crack front, whereas the ENF case produces a more oscillatory Mode II response with an asymmetric and locally jagged front. To separate the prescribed fracture-energy dissipation from the dynamic release effects, the interface work was evaluated per nodal pair for the MCC $4 \times 4\,\mathrm{mm}^2$ DCB and ENF models.

For a released nodal pair i, the logged modal force $F_{i,\alpha}$ denotes the interface force component acting in local mode α. This force contains both the cohesive contribution from the ERRC release law and the viscous damping contribution used to reduce dynamic oscillations. The cohesive part is therefore obtained as $F_{i,\alpha} - F_{i,\alpha}^{\mathrm{damp}}$, while $F_{i,\alpha}^{\mathrm{damp}}$ is integrated separately to quantify the work removed by damping.

Let $\Delta_{i,\alpha}$ denote the modal opening of pair i, where $\alpha \in \{I, II, III\}$. The force–opening path followed by pair i in mode α, from the tied-to-cohesive transition up to cycle c, is denoted by $\Gamma_{i,\alpha}(c)$. The work contributions for pair i were then evaluated as

$$W_i^{\mathrm{coh}}(c) = \sum_{\alpha \in \{I,II,III\}} \int_{\Gamma_{i,\alpha}(c)} \left(F_{i,\alpha} - F_{i,\alpha}^{\mathrm{damp}} \right) \mathrm{d}\Delta_{i,\alpha}, \tag{5.1}$$

$$W_i^{\mathrm{damp}}(c) = \sum_{\alpha \in \{I,II,III\}} \int_{\Gamma_{i,\alpha}(c)} F_{i,\alpha}^{\mathrm{damp}} \, \mathrm{d}\Delta_{i,\alpha}. \tag{5.2}$$

In the numerical post-processing, the path integrals were evaluated by trapezoidal summation over the logged force–opening history between the tied-to-cohesive transition cycle and the current solver cycle c. This notation is used because the ENF force–opening paths are not necessarily monotonic and the work depends on the logged path, not only on the initial and final opening values.

This work reconstruction is expected to represent the intended cohesive dissipation only when the local opening path is sufficiently well resolved. This is the case if the opening is essentially monotonic during release, or if any unloading and reloading increments are small enough that the logged force–opening path captures the corresponding elastic unloading. Under these conditions, a fully released pair should dissipate approximately $G_c A_i$, apart from the work removed by damping. If a fully released pair gives a reconstructed work value far from $G_c A_i$, this indicates that the local release process is too dynamic for the current implementation and post-processing to capture the expected dissipation accurately.

The total reconstructed work for pair i was defined as

$$W_i^{\mathrm{tot}}(c) = W_i^{\mathrm{coh}}(c) + W_i^{\mathrm{damp}}(c). \tag{5.3}$$

The corresponding cumulative interface work was obtained by summing over all released pairs,

$$W^{\mathrm{coh}}(c) = \sum_{i \in \mathcal{R}(c)} W_i^{\mathrm{coh}}(c), \tag{5.4}$$

$$W^{\mathrm{damp}}(c) = \sum_{i \in \mathcal{R}(c)} W_i^{\mathrm{damp}}(c), \tag{5.5}$$

$$W^{\mathrm{tot}}(c) = W^{\mathrm{coh}}(c) + W^{\mathrm{damp}}(c), \tag{5.6}$$

where $\mathcal{R}(c)$ is the set of nodal pairs that have entered the cohesive state by cycle c.

Fig. 5.8 shows the cumulative interface work for the DCB and ENF cases. In the DCB case, the cohesive work follows the expected $G_c A$ step response closely, while the damping contribution remains small. Here, A denotes the total released crack-extension area accumulated over the released nodal pairs. The small difference at the end of the DCB simulation is caused by the last active front pairs still undergoing cohesive release. Overall, this indicates that the released elastic energy is converted primarily into the intended cohesive fracture work.

The ENF case behaves differently. The damping contribution is more pronounced, which shows that part of the released energy is absorbed by the viscous damper during the local dynamic response. This is consistent with the oscillatory force–displacement response in Fig. 5.5b. The cumulative work therefore supports the same interpretation as the force response: the ENF crack front can propagate through the interface, but the Mode II release process is not sufficiently quasi-static.

The same behaviour is visible at the level of individual nodal pairs, as shown in Fig. 5.9. Here, r denotes the release order of the nodal pair. The DCB response is nearly uniform: most released pairs reach approximately $W_i^{\mathrm{tot}}/(G_c A_i) = 1$, with only a small damping contribution. This agrees with the cumulative response and confirms that the DCB release is close to the intended cohesive energy dissipation.

The ENF response is less uniform. The first released pairs dissipate approximately the intended amount of work, but several later releases deviate from the expected level. These deviations are important because, for a fully released pair under a well-resolved quasi-static release path, the reconstructed work should be close to $G_c A_i$. Values far below or above this level indicate that the local Mode II release is dynamically affected. The pair-wise work therefore does not merely explain the oscillations, it shows that the oscillations also reduce the accuracy of the reconstructed energy dissipation.

Some ENF pairs show reconstructed cohesive work close to zero even though they have reached the fully released state $D = 1$. Other ENF pairs show reconstructed work above the nominal level $G_c A_i$. These cases should not be interpreted as the damage variable allowing more than one complete release of the same pair. The implemented damage law is bounded by $D \leq 1$, so once $D = 1$ the pair has reached the complete-failure state. However, the work deviations show that reaching $D = 1$ is not sufficient to guarantee a clean quasi-static energy response.

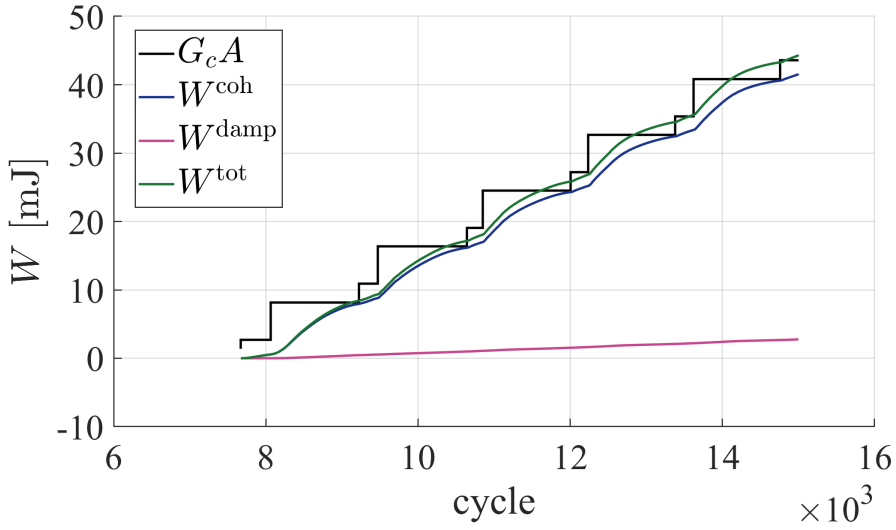

(a) DCB

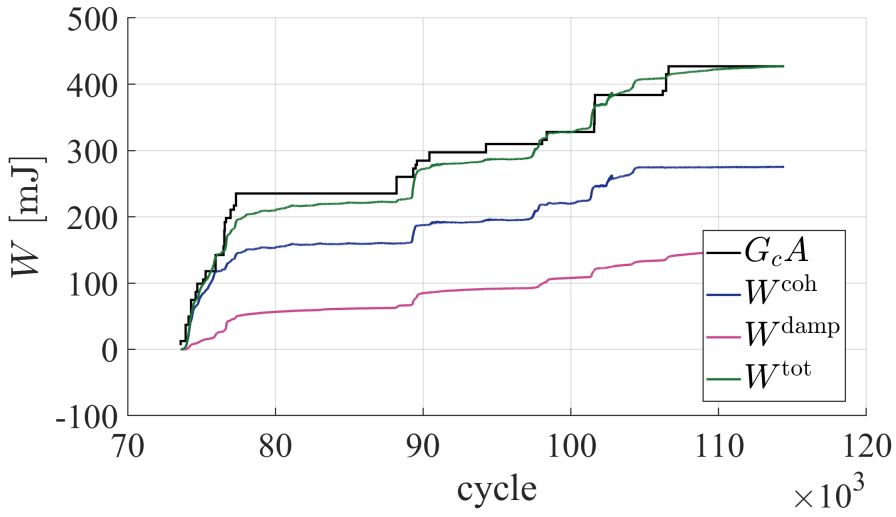

(b) ENF

Figure 5.8: Cumulative cohesive (W^{coh}), damping (W^{damp}) and total (W^{tot}) interface work for the MCC $4 \times 4\,\mathrm{mm}^2$ DCB and ENF models. The expected fracture work is compared with the cohesive work, damping work, and total interface work.

The relevant issue is the local force–opening path during release. In the DCB case, the Mode I opening evolves gradually and the cohesive softening process is resolved over several explicit cycles. The reconstructed work therefore remains close to the expected $G_c A_i$ value. In the ENF case, the transition from tied to cohesive can occur over only a few explicit cycles while the neighbouring interface is still vibrating. The opening path then becomes strongly oscillatory, and the work is split between cohesive work, damping, kinetic energy, and contact interaction. Under these conditions, the reconstructed path work no longer represents the ideal monotonic cohesive dissipation reliably.

The dissipation results therefore support the global comparison between the DCB and ENF benchmarks. The DCB case is the successful reference case: the crack front remains regular, the global force response is smooth, and the interface work

is dominated by cohesive fracture work close to the expected G_cA level. The ENF case confirms that the implementation can propagate a Mode II-dominated crack front, but the larger damping contribution and the pair-wise work deviations show that the release process is not sufficiently quasi-static. The ENF result should therefore be interpreted as dynamically problematic and not as a fully validated energy-dissipation result.

This conclusion is also consistent with the release behaviour. In the DCB case, the release process is stable and the propagation criterion is reached under a smooth Mode I opening response. In the ENF case, local vibrations can cause the VCCT criterion to be satisfied during a transient force state. A pair may therefore enter the cohesive state at a locally elevated energy-release-rate value. Although the damage variable still limits the pair to the complete-failure state $D = 1$, the subsequent energy redistribution is not purely cohesive: part of the released energy appears as local kinetic energy and part is removed by damping. Reducing this dynamic excitation is therefore a central requirement for future Mode II improvement.

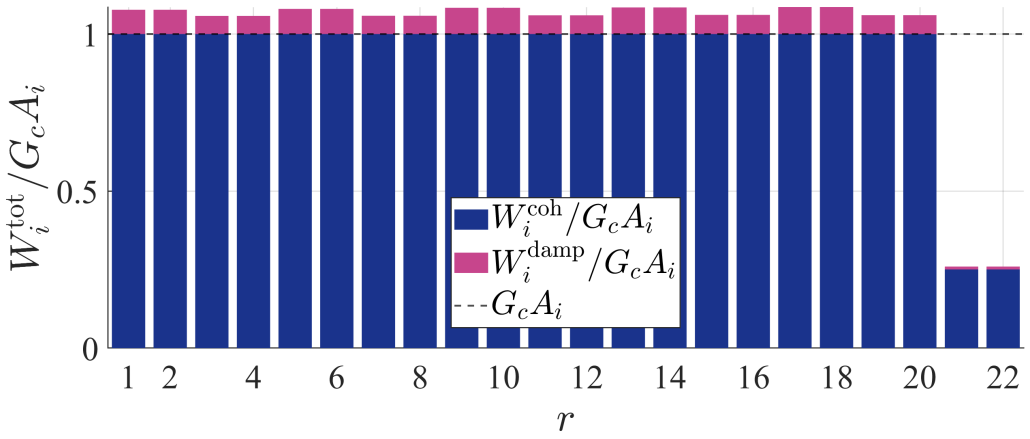

(a) DCB

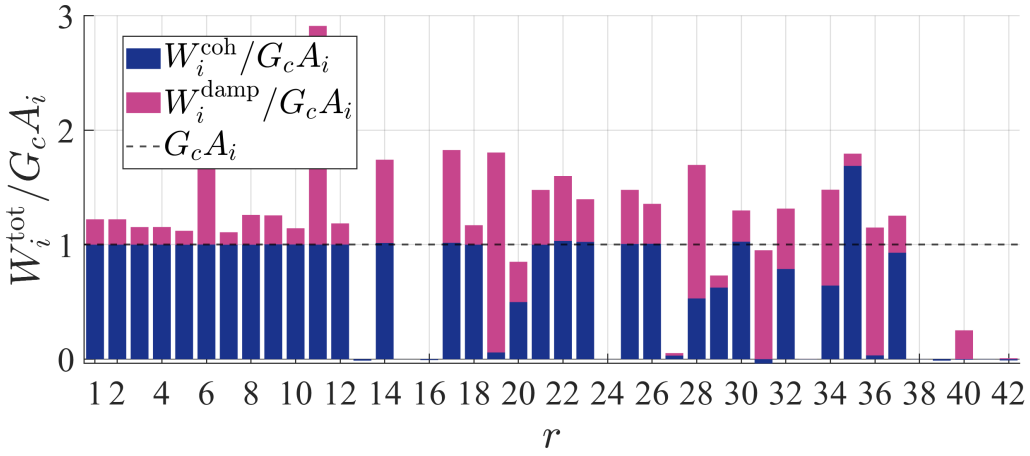

(b) ENF

Figure 5.9: Total interface work per released pair, normalised by G_cA_i. The bars separate cohesive work and damping work.

5.3.4 Efficiency Estimate

In order to assess the higher potential efficiency for the industrially feasible multi-core implementation, efficiency comparisons were not performed using identical interface discretisations. Instead, each method was evaluated using the coarsest mesh density capable of reproducing the analytical fracture response based on the results from Sec. 5.3. This allows the comparison to reflect the practical computational efficiency of each formulation rather than purely per-element performance. The efficiency estimates for both the DCB and ENF results can be seen in Tabs. 5.1 and 5.2.

Table 5.1: DCB computational efficiency benchmark comparing CZM and the proposed UEL formulation.

Method	Mesh	CPU Time [s]	Cycles	Δt_{avg} [s]	Nodes
CZM Fine	0.25×0.25 mm^2	77776	460556	2.17×10^{-8}	75558
CZM Reduced	2×2 mm^2	205	76824	1.30×10^{-7}	1386
UEL	4×4 mm^2	100	75000	1.34×10^{-7}	408

Table 5.2: ENF computational efficiency benchmark comparing CZM and the proposed UEL formulation.

Method	Mesh	CPU Time [s]	Cycles	Δt_{avg} [s]	Nodes
CZM Fine	0.5×0.5 mm^2	28971	248610	4.01×10^{-8}	49446
UEL	4×4 mm^2	1042	165505	6.04×10^{-8}	972

Compared to the fine CZM discretisation, the reduced stiffness CZM model achieved an approximately 379× reduction in computational runtime but severly fails to predict the correct response, while the proposed UEL formulation achieved an approximately 778× reduction. This means that, compared to the reduced-stiffness CZM formulation, the UEL approach remained approximately 2× faster while maintaining accurate agreement with the analytical response. The reduced-stiffness CZM and UEL formulations exhibited nearly identical average stable timesteps, indicating that the dominant efficiency improvement of the UEL primarily originates from the significantly reduced interface discretisation and overall model size rather than time-step stability improvements alone.

In the case of the ENF sample, the fine CZM discretisation achieved an approximately 27.8× reduction in computational runtime in comparison to the proposed UEL formulation. It should also be noticed it simultaneously reduced the total number of nodes from 49446 to 972. The UEL formulation exhibited an almost identical average stable time step compared to the fine CZM model, indicating that the reduced computational time is a result mainly from the reduced model size.

6

Conclusions

The present work showed that the ERRC method of Daniel *et al.* [1, 2], originally developed for efficient large-element delamination modelling, can be specialised into a solid-interface formulation for adhesive bondlines and implemented as a user-defined solid element in LS-DYNA. In the implemented formulation, the bonded interface was represented by lower–upper nodal pairs, where each pair could be in one of three states: tied, cohesive, or open. This state-based representation allowed the interface to remain tied before crack growth, to transfer into a cohesive release state once a VCCT-based propagation criterion was satisfied, and to become fully separated once the damage variable reached $D = 1$.

In contrast to a conventional CZM, the method does not require the physical FPZ to be resolved by several small interface elements. This confirms the main practical motivation of the method: interfacial crack propagation can be represented using larger in-plane interface elements than would normally be required by a mesh-resolved CZM approach, provided that the release process remains sufficiently stable.

The DCB benchmark showed the most robust behaviour. For the 4×4 mm^2 interface discretisation, the crack front propagated in a row-uniform manner and the crack-growth response was consistent with the expected Mode I behaviour. The analytical CBT crack length was recovered within approximately one discrete interface-element edge, which is consistent with the resolution limit of the discrete-pair formulation. The dissipation diagnostics also showed that the released pairs followed the intended fracture-energy target, with only a small damping contribution. These results show the implementation captures the intended ERRC mechanics for a pure Mode I, mesh-aligned crack-growth problem. However, the $2 \times 2\,\text{mm}^2$ MCC case showed local force spikes not present in the SCR solution, indicating that the nonlocal ownership and front-update logic can still affect the timing of individual release events.

The ENF benchmark was more challenging. The implementation front-activation logic produced a coherent Mode II-dominated propagation chain in the intended in-plane direction, but the crack front developed asymmetrically across the specimen width rather than advancing as a single straight front. More importantly, the global force–displacement response contained pronounced high-frequency oscillations superimposed on the expected response. The MCC implementation showed a more dynamic response than the SCR implementation.

The dissipation analysis showed that these oscillations also affect the reconstructed energy response. The implemented damage law limits the release state through $D \leq 1$, so a nodal pair cannot release beyond the complete-failure state. However, for the fastest Mode II releases, the tied-to-cohesive transition can occur

over only a few explicit cycles while the surrounding interface is still vibrating. In such cases, the force–opening path is no longer a well-resolved quasi-static softening path, and the reconstructed work quantities W_i^{coh}, W_i^{damp}, and W_i^{tot} may deviate from the expected $G_c A_i$ level. These deviations should therefore be interpreted as evidence that the present ENF release is dynamically contaminated and not yet energy-accurate in a quasi-static sense.

The ENF result therefore verifies that the method can advance a shear-dominated crack front, but it does not yet provide a sufficiently smooth quasi-static Mode II response. This is the main remaining limitation of the current implementation.

A key contribution of this work was the extension of the ERRC concept to a solid user-element formulation for finite-thickness adhesive interfaces. This enables adhesive bondlines to be represented using solid adherends, where a finite-thickness solid discretisation is a natural modelling choice. To support this formulation, a deterministic nodal-pair ownership strategy and a nonlocal synchronisation routine were proposed to prevent duplicate force application for physical nodal pairs shared by several element copies.

Overall, the work supports the feasibility of implementing the ERRC method as a solid-element adhesive-interface formulation in LS-DYNA. The strongest verified result is the regular Mode I DCB benchmark, where large-element crack propagation was obtained with substantially coarser interface discretisation than required by conventional CZM. The ENF benchmark confirms that the same implementation can propagate a Mode II-dominated crack front, but also identifies the main remaining challenge: controlling the local dynamics of the tied-to-cohesive release process. The current implementation should therefore be regarded as a proof of concept for regular, mesh-aligned benchmark problems rather than as a fully general industrial fracture-modelling framework.

7
Future Work

Several extensions are required before the proposed adhesive-interface formulation can be applied broadly to industrial fracture simulations.

First, the method should be extended to irregular meshes and arbitrarily shaped crack fronts. The present implementation was verified only for regular, mesh-aligned DCB and ENF coupons, hence the implementation should be extended with the correction factors allowing irregular meshes and the arbitrary-front concepts proposed by Daniel *et al.* [2]. The current predecessor identification is based on the regular coupon topology and should be replaced by a more general crack front tracking procedure.

The current initialisation procedure also assumes an initially straight, mesh-aligned crack. Future work should therefore introduce a more general initial crack definition, for example through prescribed node sets or interface-front data exported from pre-processing tools. For curved bond surfaces, the local pair basis should also be made consistent across all element copies associated with the same physical pair, for example by averaging or otherwise synchronising the basis in the nonlocal routine. Further, the LS-DYNA implementation should be made more robust for thin adhesive layers. Internal LS-DYNA checks for `ELFORM=101-105` evaluate the element volume to detect distorted elements and collapsed stable time increments. These checks limited the range of bondline thicknesses that could be represented by the current formulation. In addition, the tied-pair force reconstruction became sensitive for very small adhesive thicknesses, where stable kinematic enforcement could not be achieved below the validated thickness range. Addressing this limitation may require direct collaboration with LS-DYNA developers.

The Mode II release dynamics should be investigated systematically. The ENF results showed that some tied-to-cohesive transitions occur over very few explicit cycles, leading to high-frequency force oscillations, asymmetric crack front evolution, and unreliable reconstructed path-work measures. Although the implemented damage variable is bounded by $D \leq 1$, the sampled path-work quantities can under- or over-count the intended monotonic fracture work during rapid release events. Future studies should therefore vary the loading rate, kinetic-to-internal-energy ratio, cohesive damping coefficient, contact penalty stiffness, and crack-face contact formulation. The aim should be to separate physical Mode II instability, ERRC-related release dynamics, and model-specific contact artefacts. The dissipation diagnostics should also be strengthened. In the present work, the damage-based quantity $G_c A_i D_{\text{final}}$ was used as the primary measure of intended fracture-energy release, while the reconstructed path-work quantities were interpreted as diagnostics of dynamic release. Future work should add solver-level energy bookkeeping for the

cohesive and damping contributions, preferably written directly from the user element during the release event. This would reduce the ambiguity introduced by cycle-sampled force–opening reconstruction and make it easier to distinguish fracture work, damping work, kinetic energy, and contact work.

The `cracklook` routine is dispatched through the LS-DYNA keyword `*USER_NONLOCAL_SEARCH`, which requires the user to define a search-radius parameter for each interface. Selecting this parameter is a non-trivial part of the model setup (App. B.2.2). If the radius is too small, the routine fails to identify all elements associated with a shared physical pair, compromising deterministic ownership and tied-pair force assembly. If the radius is unnecessarily large, the nonlocal search becomes more expensive. Future work should therefore include an automatic search-radius estimate based on the local interface element size and expected neighbouring topology.

The parallel execution scope should also be expanded. The present verification cases were evaluated in the available MPP UEL environment using single-rank execution. A fully distributed-memory MPP implementation would require communication of nodal force, mass, state, ownership, and crack front information across MPI-rank boundaries. Shared-memory SMP execution should also be evaluated when the required user-library support is available, since the implementation was designed with multi-core capability in mind.

The method should also be extended beyond predefined cracks. The present implementation requires an initially open region and a predefined crack path. For more general adhesive-joint simulations, crack initiation should be introduced through a stress-, strain-, or energy-based criterion that can transfer an initially tied pair into the cohesive state without a prescribed precrack. This would allow the same framework to simulate both crack initiation and crack propagation.

Finally, additional benchmark cases should be introduced. Mixed-mode bending (MMB) should be used to verify the mixed-mode fracture criterion and the transition between Mode I- and Mode II-dominated release. More representative three-dimensional adhesive-joint geometries should then be considered. A systematic computational-cost comparison against conventional CZM should also be repeated for these cases, including element count, stable time step, and mesh-size sensitivity.

Bibliography

[1] P. M. Daniel, J. Främby, M. Fagerström, and P. Maimí, "An efficient err-cohesive method for the modelling of delamination propagation with large elements," *Composites Part A: Applied Science and Manufacturing*, vol. 167, p. 107423, 2023, doi: 10.1016/j.compositesa.2022.107423.

[2] ——, "A method for modelling arbitrarily shaped delamination fronts with large and distorted elements," *Engineering Fracture Mechanics*, vol. 306, p. 110193, 2024, doi: 10.1016/j.engfracmech.2024.110193.

[3] International Energy Agency (IEA), "Aviation," [Online]. Available: https://www.iea.org/energy-system/transport/aviation, 2024, Accessed: 2026-05-07.

[4] M. K. Hagnell, S. Kumaraswamy, T. Nyman, and M. Åkermo, "From aviation to automotive — a study on material selection and its implication on cost and weight efficient structural composite and sandwich designs," *Heliyon*, vol. 6, no. 3, p. e03716, 2020, doi: 10.1016/j.heliyon.2020.e03716.

[5] A. J. Timmis, A. Hodzic, L. Koh, M. Bonner, C. Soutis, A. W. Schäfer, and L. Dray, "Environmental impact assessment of aviation emission reduction through the implementation of composite materials," *The International Journal of Life Cycle Assessment*, vol. 20, no. 2, pp. 233–243, 2015, doi: 10.1007/s11367-014-0824-0.

[6] ESA-ESTEC, *Space engineering – Adhesive bonding handbook*, European Cooperation for Space Standardization (ECSS), 2011, eCSS-E-HB-32-21A, 20 March 2011.

[7] D. C. Noorman, "Cohesive zone modelling in adhesively bonded joints: Analysis on crack propagation in adhesives and adherends," Master's thesis, Delft University of Technology, 2014.

[8] ASTM International, "Standard practice for classifying failure modes in fiber-reinforced-plastic (frp) joints," ASTM International, West Conshohocken, PA, USA, Standard ASTM D5573-99(2019), 2019.

[9] R. Krueger, "A summary of benchmark examples to assess the performance of quasi-static delamination propagation prediction capabilities in finite element codes," *Journal of Composite Materials*, vol. 49, no. 26, pp. 3297–3316, 2015, doi: 10.1177/0021998314561812.

[10] B. D. Agarwal, L. J. Broutman, and K. Chandrashekhara, *Analysis and Performance of Fiber Composites*, 3rd ed. Hoboken, NJ: John Wiley & Sons, 2006.

[11] A. Turon, C. G. Dávila, P. P. Camanho, and J. Costa, "An engineering solution for mesh size effects in the simulation of delamination using cohesive zone models," *Engineering Fracture Mechanics*, vol. 74, pp. 1665–1682, 2007, doi:

10.1016/j.engfracmech.2006.08.025.
[12] ANSYS, Inc., "LS-DYNA General," [Online]. Available: https://lsdyna.ansys.com/knowledge-base/general/, 2023, Accessed: 2025-05-06.
[13] ——, "LS-DYNA User Defined Interfaces," [Online]. Available: https://lsdyna.ansys.com/user-defined-materials-and-user-defined-interfaces/, 2023, Accessed: 2025-05-06.
[14] B. Ravichandran and M. Balasubramanian, "A comprehensive review on sustainability evaluation of joining methods for engineering materials," *Materials Today Sustainability*, vol. 31, p. 101151, 2025, doi: 10.1016/j.mtsust.2025.101151.
[15] Industrieverband Klebstoffe e.V. (IVK), "Adhesives in the circular economy," Industrieverband Klebstoffe e.V., Tech. Rep., Aug. 2023.
[16] A. Baker, "Bonded composite repair of fatigue-cracked primary aircraft structure," *Composite Structures*, vol. 47, no. 1–4, pp. 431–443, 1999, doi: 10.1016/S0263-8223(00)00011-8.
[17] GKN Aerospace, "Repair Solutions," [Online]. Available: https://www.gknaerospace.com/markets-solutions/aftermarket-services-mro/repair-solutions/, 2026, Accessed: 2026-05-07.
[18] W. S. Jevons, *The Coal Question: An Inquiry Concerning the Progress of the Nation, and the Probable Exhaustion of Our Coal-Mines*, 2nd ed. London, UK: Macmillan and Co., 1866.
[19] B. Alcott, "Jevons' paradox," *Ecological Economics*, vol. 54, no. 1, pp. 9–21, 2005, doi: 10.1016/j.ecolecon.2005.03.020.
[20] DeepL SE, "DeepL Write," [Online]. Available: https://www.deepl.com/write, 2026, Large language model/writing assistant. Accessed: 2026-05-09.
[21] OpenAI, "ChatGPT," [Online]. Available: https://openai.com/chatgpt, 2026, Large language model. Accessed: 2026-05-09.
[22] M. Janssen, J. Zuidema, and R. J. H. Wanhill, *Fracture Mechanics*, 2nd ed. Delft, The Netherlands: Delft University Press, 2002.
[23] A. Ekberg and H. Andersson, *Non-linear Fracture Mechanics*. Chalmers University of Technology, 2025, course compendium (PDF provided by the author/instructor).
[24] N. E. Dowling, *Mechanical Behavior of Materials: Engineering Methods for Deformation, Fracture, and Fatigue*, 4th ed. Harlow, England: Pearson, 2013.
[25] J. R. Reeder, K. Demarco, and K. S. Whitley, "The use of doubler reinforcement in delamination toughness testing," *Composites Part A: Applied Science and Manufacturing*, vol. 35, no. 11, pp. 1337–1344, 2004, doi: 10.1016/j.compositesa.2004.02.021.
[26] R. Kreuger, "Virtual crack closure technique: History, approach, and applications," *Applied Mechanics Reviews*, vol. 57, no. 2, pp. 109–143, 2004, doi: 10.1115/1.1595677.
[27] A. Turon, P. P. Camanho, J. Costa, and C. G. Dávila, "A damage model for the simulation of delamination in advanced composites under variable-mode loading," *Mechanics of Materials*, vol. 38, pp. 1072–1089, 2006, doi: 10.1016/j.mechmat.2005.10.003.
[28] C. G. Dávila, P. P. Camanho, and M. F. de Moura, "Mixed-mode decohesion

elements for analyses of progressive delamination," in *Proceedings of the 42nd AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conference*. AIAA, 2001, paper AIAA-2001-1486, Seattle, WA, USA, April 16–19, 2001, doi: 10.2514/6.2001-1486.

[29] A. Soto, E. González, P. Maimí, A. Turon, J. Sainz de Aja, and F. de la Escalera, "Cohesive zone length of orthotropic materials undergoing delamination," *Engineering Fracture Mechanics*, vol. 159, pp. 174–188, 2016, doi: 10.1016/j.engfracmech.2016.03.033.

[30] J. Främby, M. Fagerström, and J. Karlsson, "An adaptive shell element for explicit dynamic analysis of failure in laminated composites part 1: Adaptive kinematics and numerical implementation," *Engineering Fracture Mechanics*, vol. 240, p. 107288, 2020, doi: 10.1016/j.engfracmech.2020.107288.

[31] J. Främby, "Private communication regarding LS-DYNA implementation," Private communication, Mar. 2026, Synopsys (ANSYS).

[32] ANSYS, Inc., *LS-DYNA R14 Keyword Manual, Volume II*, ANSYS, Inc., 2023, created 24 February 2023.

[33] GKN Aerospace, "Internal confidential documents GKN Aerospace," Internal documents, February 2026.

[34] ANSYS, Inc., *LS-DYNA R14 Keyword Manual, Volume I*, ANSYS, Inc., 2023, updated 28 November 2023.

A

Material Cards

A.1 Composite Material Cards

```
*KEYWORD
*MAT_ORTHOTROPIC_ELASTIC_TITLE
DCB - T300/1076
$#   mid       ro        ea       eb       ec      prba     prca     prcb
       2   1.5E-9  139000.0  10160.0 10160.0    0.0218   0.0218    0.436
$#  gab       gbc       gca      aopt         g      sigf
 4600.0    3540.0    4600.0       0.0       0.0       0.0
$#   xp        yp        zp        a1        a2        a3     macf     ihis
    0.0       0.0       0.0       0.0       0.0       0.0        1        0
$#   v1        v2        v3        d1        d2        d3     beta      ref
    0.0       0.0       0.0       0.0       0.0       0.0      0.0      0.0
*END
```

LS-DYNA Keyword A.1: DCB – T300/1076

```
*KEYWORD
*MAT_ORTHOTROPIC_ELASTIC_TITLE
ENF/MMB - IM7/8552
$#     mid        ro        ea        eb        ec      prba      prca
      prcb
         21.50000E-9  161000.0   11380.0   11380.0    0.0226    0.0226
     0.450
$#     gab       gbc       gca      aopt         g      sigf
    5200.0    3900.0    5200.0       0.0       0.0       0.0
$#      xp        yp        zp        a1        a2        a3      macf
      ihis
       0.0       0.0       0.0       0.0       0.0       0.0         1
         0
$#      v1        v2        v3        d1        d2        d3      beta
       ref
       0.0       0.0       0.0       0.0       0.0       0.0       0.0
       0.0
*END
```

LS-DYNA Keyword A.2: ENF/MMB – IM7/8552

A.2 Cohesive Zone Material Cards

```
*KEYWORD
*MAT_COHESIVE_MIXED_MODE_TITLE
CZM
$#  mid     ro     roflg  intfail      en      et   gic   giic
   138 1.5E-12         1      1.0 33867.0 15333.0  0.17  0.494
$#     xmu         t         s       und       utd     gamma
      1.62      30.0      30.0       0.0       0.0       1.0
*END
```

LS-DYNA Keyword A.3: Standard cohesive material card used in the DCB

```
*KEYWORD
*MAT_COHESIVE_MIXED_MODE_TITLE
CZM
$#  mid     ro     roflg  intfail      en      et   gic   giic
   138 1.5E-12         1      1.0  76000.0 35000.0  0.212  0.774
$#     xmu         t         s       und       utd     gamma
      2.10      30.0      60.0       0.0       0.0       1.0
*END
```

LS-DYNA Keyword A.4: Standard cohesive material card used in the ENF

```
*KEYWORD
*MAT_COHESIVE_MIXED_MODE_TITLE
CZM
$#  mid     ro     roflg  intfail      en      et   gic   giic
   138 1.5E-12         1      1.0 33867.0 15333.0  0.17  0.494
$#     xmu         t         s       und       utd     gamma
      1.62      0.54      0.54       0.0       0.0       1.0
*END
```

LS-DYNA Keyword A.5: Reduced-strength cohesive material card for 4 mm mesh

```
*KEYWORD
*MAT_COHESIVE_MIXED_MODE_TITLE
CZM
$#  mid     ro     roflg  intfail      en      et   gic   giic
   138 1.5E-12         1      1.0 33867.0 15333.0  0.17  0.494
$#     xmu         t         s       und       utd     gamma
      1.62      2.05      2.05       0.0       0.0       1.0
*END
```

LS-DYNA Keyword A.6: Reduced-strength cohesive material card for 2 mm mesh

B

User-Element: Modelling

The implemented user-element needs only two new keyword files compared with a conventional LS-DYNA deck: `parts.k`, which carries the section card for the user-element together with the laminate plies, materials, nodes and elements, and `nonlocal.k`, which dispatches the per-cycle non-local synchronisation routine *crack-look*. Both files are described below.

B.1 Overview

The deck structure used in this work follows a modular include layout:

```
*KEYWORD
*INCLUDE
parts.k          $ sections, materials, parts, nodes, elements
*INCLUDE
nonlocal.k       $ *USER_NONLOCAL_SEARCH for cracklook
*INCLUDE
bc_sets.k        $ boundary conditions, loads, history nodes
*INCLUDE
controls.k       $ time stepping, contact, output databases
*END
```

LS-DYNA Keyword B.1: Top-level include tree.

The two files that are specific to the user-element are `parts.k` and `nonlocal.k`. The remaining files (`bc_sets.k`, `controls.k`) are conventional LS-DYNA setup and are therefore not described here.

The order in which the parts appear inside `parts.k` is significant: the bulk laminate solids must precede the user-element interface part in the keyword stack so that bulk-solid contributions are written into `dm_a` before the user-element pass reads them (Section 3.4). The explicit ordering is shown in the listing in Section B.2.1.

B.2 Submodules

B.2.1 parts.k

Module `parts.k` carries everything that defines the bulk–interface element stack: the section cards, the materials, the part declarations, and the nodes and elements. The two section cards used are an orthotropic constant-stress hex (`ELFORM`

= -2) for the laminate plies and the user-element section (ELFORM = 103, NIP = 0, LMC = 15, NHSV = 100) for the interface. The user-element section also carries the cm(1)..cm(14) parameter vector inline, since the user-element reads the cm slots for its constitutive parameters (Section 3.1).

```
$ --- 1. Materials ------------------------------------------------------
*MAT_ENHANCED_COMPOSITE_DAMAGE_TITLE   $ MID 2: laminate plies
Composite_T300_1076
   2, ...                              $ E11, E22, ..., G_Ic, G_IIc

*MAT_ELASTIC_TITLE                     $ MID 3: placeholder for UEL
    section
Elastic_UEL_placeholder
   3,  1.500E-12,  210.0E-12,  0.3,  0.0,  0.0,  0,  0

$ --- 2. Sections --------------------------------------------------------
*SECTION_SOLID                         $ SECID 300: laminate (ELFORM = -2)
   300, -2

*SECTION_SOLID_TITLE                   $ SECID 200: user-element
UEL_interface
$# secid  elform   aet  unused  unused  unused  cohoff  gaskett
     200     103     0                                 0.0       0.0
$#   nip   nxdof   ihgf   itaj    lmc   nhsv   xnod   unused
       0      0      0      0     14    100      0
$#    p1     p2     p3     p4     p5     p6     p7     p8
    1E+4     20  0.170  0.494    4.0  62.2E+3 62.2E+3  0.0
$#    p9    p10    p11    p12    p13    p14
    1.60   0.10      2    335      0    1.0

$ --- 3. Parts (ORDER MATTERS: bulk solids first, UEL last) ------------
*PART
Ply1_lower
   1,  300,  2,  0,  0,  0,  0,  0    $ part 1, secid 300, mid 2

*PART
Ply2_upper
   3,  300,  2,  0,  0,  0,  0,  0    $ part 3, secid 300, mid 2

*PART
UEL_interface_t012mm
   2,  200,  3,  0,  0,  0,  0,  0    $ part 2, secid 200, mid 3

$ --- 4. Nodes and elements ---------------------------------------------
*NODE
   ...
*ELEMENT_SOLID_ORTHO                    $ laminate plies
   ...
*ELEMENT_SOLID                         $ UEL interface
   ...
```

LS-DYNA Keyword B.2: Schematic structure of `parts.k` showing the UEL section card and the explicit bulk-before-UEL part ordering. Numerical values are taken from the DCB 2 mm coupon.

The user-element section must use `NIP = 0` so that LS-DYNA does not invoke a built-in stress update; the response is supplied entirely by the user routine. `LMC = 15` reserves the fifteen `cm` slots that the section card carries inline. `NHSV = 100` reserves the per-element history-variable budget; the user-element uses 56 of these for the per-pair quantities (pair-keys, ownership flags, front type, area, damage, release-force components) listed in Sec. 3.4.2 (Table 3.1).

The placeholder `*MAT_ELASTIC` (MID 3) on the user-element section card supplies LS-DYNA with the density needed for the time-step computation and is not used for the constitutive response, since the user routine overrides the element forces directly.

B.2.2 nonlocal.k

The non-local synchronisation routine *cracklook* is dispatched by the LS-DYNA keyword `*USER_NONLOCAL_SEARCH` [34]. The keyword needs to be present in the deck for the synchronisation to be invoked at all, and it must be set up per interface, since the routine acts on the user-element parts that are explicitly listed on the keyword card. The relevant parameters in the LS-DYNA cards are summarised below.

```
*USER_NONLOCAL_SEARCH
$ SSID, ASID, STYPE, ATYPE, R, SF1, SF2, SF3
         2         2         1         1       4.5       1.0       1.0
      1.0
$ Card 2:
$ NFREQ, VOLTYPE, UTYPE, NUCONST, NUELHSV
         1         1       101         8        56
$ Card 2.1:
$ User constants P1..P8 (currently unused, all zero)
       0.0       0.0       0.0       0.0       0.0       0.0       0.0
      0.0
$ Card 2.2:
$ Gather UE-HSV locations 1..56 from surrounding user-elements
$  1.. 4 = KEY_MIN(1..4)
$  5.. 8 = KEY_MAX(1..4)
         1         2         3         4         5         6         7
        8
$  9..12 = OWNER_EID(1..4)
$ 13..16 = IS_OWNER(1..4)
         9        10        11        12        13        14        15
       16
$ 17..20 = MID_X(1..4)
$ 21..24 = MID_Y(1..4)
        17        18        19        20        21        22        23
       24
$ 25..28 = MID_Z(1..4)
```

```
$ 29..32 = STRIP_AREA(1..4)
        25        26        27        28        29        30        31
       32
$ 33..36 = FRONT_FLAG(1..4)
$ 37..40 = FRONT_TYPE(1..4)
        33        34        35        36        37        38        39
       40
$ 41..44 = AREA_TRIB(1..4)
$ 45..48 = PAIR_KEYMAX(1..4)
        41        42        43        44        45        46        47
       48
$ 49..52 = HSTATLOC(1..4)
$ 53..56 = HBRELLOC(1..4)
        49        50        51        52        53        54        55
       56
```

LS-DYNA Keyword B.3: The `*USER_NONLOCAL_SEARCH` card used by the present implementation. Card 1 carries the part-set IDs and the search radius `R`; Card 2 sets the call frequency, the volume/user-subroutine identifiers and the number of UE history-variable slots to gather (`NUELHSV`); Card 2.1 lists eight user constants `P1..P8` (unused in this work); Card 2.2 lists the `NUELHSV` HSV slot indices gathered from surrounding user-elements each cycle.

Search radius. The `R` parameter on the first card defines the spatial radius within which neighbouring element copies are considered candidates for the same physical nodal pair. The radius selection is a non-trivial part of the deck setup:

- If `R` is too small relative to the in-plane element edge, neighbouring element copies of the same physical pair are not found, the deterministic owner-selection in cracklook is incomplete, and tied-pair forces are either over-counted (multiple element copies all assemble) or zero (no copy is selected as the owner). The simulation runs without warning but the bond-line response is wrong.
- If `R` is too large, the surrounding-list scan visits more candidate elements per pair than necessary. The runtime cost grows with the number of UEL elements within `R` of any given pair midpoint, so a generous radius noticeably slows down the simulation without changing the result.

A practical choice is a approx. 1.5-2 times the in-plane element edge, enough to robustly cover the immediate neighbour element copies and the nearest predecessor along the propagation direction, but not so generous that distant elements participate in the surrounding-list scan. If the radius is chosen to small, the interface does not provide deterministic ownership and the tied condition will not work.

Per-interface accounting. Each physical interface in the model needs its own `*USER_NONLOCAL_SEARCH` card listing the user-element parts that belong to that interface. The implementation does not auto-detect interfaces; every bond line is registered explicitly. For coupons with a single delamination plane this is one card; multi-interface laminates need one card per interface, and the per-pair history-variable budget (`NHSV` ≤ 100 per element) applies independently to each interface.

DEPARTMENT OF INDUSTRIAL AND MATERIALS SCIENCE
CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden
www.chalmers.se

CHALMERS
UNIVERSITY OF TECHNOLOGY