

UC Santa Barbara

UC Santa Barbara Electronic Theses and Dissertations

Title

Simulations of Interface Cracking and Oxidation in Layered Systems

Permalink

<https://escholarship.org/uc/item/0pp8r1kv>

Author

Sehr, Stephen

Publication Date

2020

Peer reviewed|Thesis/dissertation

University of California
Santa Barbara

Simulations of Interface Cracking and Oxidation in Layered Systems

A dissertation submitted in partial satisfaction
of the requirements for the degree

Doctor of Philosophy
in
Mechanical Engineering

by

Stephen Timothy Sehr

Committee in charge:

Professor Matthew R. Begley, Chair
Professor Carlos G. Levi
Professor Robert M. McMeeking
Professor Frank W. Zok

December 2020

The Dissertation of Stephen Timothy Sehr is approved.

Professor Carlos G. Levi

Professor Robert M. McMeeking

Professor Frank W. Zok

Professor Matthew R. Begley, Committee Chair

September 2020

Simulations of Interface Cracking and Oxidation in Layered Systems

Copyright © 2020

by

Stephen Timothy Sehr

This thesis is dedicated to the memories of Gerald and Mary
Dinneen

Acknowledgements

I am grateful for all the people who have helped me over the years. I am especially thankful for my family who have supported me throughout my education. Even though my parents, Timothy and Barbara Sehr, are not engineers, they supported my desire to understand how things worked. I am very grateful for my maternal grandparents, Gerald and Mary Dinneen, for happily encouraging me to go into engineering. I do not know if I would have gone into materials engineering if it weren't for my Grandfather setting up a meeting for me with an engineering professor almost fifteen years ago. Due to that meeting, I was introduced to materials engineering. He may have passed away before I finished this journey, but I know he was happy to have another engineer in the family.

On a more practical level, it would not have been possible to complete this research without all the help provided by those I met at UCSB. Begley Research Group members past and present have helped in innumerable ways throughout the years. A lot of the work I was able to do would not have been possible without Will Pro and Rone Kwei Lim's advancements in cohesive zone modeling. While he was only here for a short time, the computer Phillip Dodge built for his research was crucial to me completing my work. Others such as Eric Clough, Virginia Collier, Connie Dong and Sara Messina have been helpful with discussions on various problems. The ECI staff and Josh Johnson have helped troubleshoot and fix more computer problems than I can count.

Over my nine years here, the different communities I have joined have been incredibly helpful. With the local cycling community, I have spent hours and hours giving my mind a break while (usually) having fun and exploring all sorts of natural beauty. One of the things I looked forward to every week was meeting for lunch with other graduate students and some Capuchin friars. The conversations on just about anything other than engineering were a great distraction.

I cannot thank my wife, Sarah, enough for the love and support she has given me the past few years. She believed in me and was confident I would finish even when I doubted myself. In the past few months especially, she has taken on innumerable small tasks so I could focus on my work. Her knowledge of computer programming helped me think through complex problems of my own. She has always been willing to listen as I try to explain something I'm stuck on and help me move forward. Without her encouragement through what seemed like endless months of frustration, I do not know if I would have had the fortitude to complete my degree. For that I will be eternally grateful.

Curriculum Vitæ

Stephen Timothy Sehr

Education

- 2020 Ph.D. in Mechanical Engineering (Expected), University of California, Santa Barbara.
- 2011 B.S. in Materials Science Engineering, Purdue University

Research Experience

- A. Begley Research Group, University of California, Santa Barbara *Spring 2012-Present*
Research Advisor: Professor Matthew R. Begley
- B. Wiggins Research Group, University of Southern Mississippi *Summer 2010*
Research Advisor: Professor Jeffrey S. Wiggins

Publications

4. S. Sehr, S. Amidi and M.R. Begley. *Interface delamination vs. bulk cracking along wavy interfaces*. **Engineering Fracture Mechanics**. 206. 2019. pg 64-74.
3. J.W. Pro, S. Sehr, R.K. Lim, L.R. Petzold and M.R. Begley. *Conditions controlling kink crack nucleation out of, and delamination along, a mixed-mode interface crack*. **Journal of the Mechanics and Physics of Solids**. 121. 2018. pg 480-495.
2. B.T. Richards, K.A. Young, F. de Francqueville, S. Sehr, M.R. Begley and H.N.G. Wadley. *Response of ytterbium disilicate-silicon environmental barrier coatings to thermal cycling in water vapor*. **Acta Materialia**. 106. 2016. pg 1-14.
1. B.T. Richards, S. Sehr, F. de Francqueville, M.R. Begley and H.N.G. Wadley. *Fracture mechanisms of ytterbium monosilicate environmental barrier coatings during cyclic thermal exposure*. **Acta Materialia**. 103. 2016. pg 448-460.

Professional Experience

- A. Bettis Atomic Power Laboratory, *Engineering Intern*. West Mifflin, PA. *Summer 2009*
Studied residual stresses in bent tubing.
- B. Haynes International, *Engineering Intern*. Kokomo, IN. *Summer 2008*
Analyzed trends in test results and assisted on new product development.

Abstract

Simulations of Interface Cracking and Oxidation in Layered Systems

by

Stephen Timothy Sehr

The use of ceramic materials in propulsion and power generation turbines allows for significant improvements in efficiency by allowing for higher operating temperatures. A central challenge in the development of ceramic systems is the propensity for cracking driven by thermal expansion mismatch and oxidation driven by harsh environmental conditions. Many strategies to improve durability involve geometric or material features that have not been previously considered in simulations of failure, such as non-planar interfaces associated with woven ceramic matrix composites and sprayed coatings, and the impact of oxide formation underneath coatings and near cracks. This work focuses on advances in computational frameworks tailored to address such phenomena, with sufficient speed to conduct parametric studies to quantify important material and geometric interactions.

The use of distributed cohesive zone models holds distinct advantages for simulations of crack evolution in brittle systems, notably the ability to predict crack path evolution (as opposed to the onset of initial advance). However, a significant drawback is that the method is slow compared to conventional finite element models, and numerical stability can be sensitive to the details of implementation. To address these concerns, a systematic study of computational methods has identified methods to improve simulation speeds and stability. Specifically, the use of sub-domains near crack tips provides a basis to limit model size with locally refined meshes, and mitigate the cost of simulating loading that is below the cracking threshold. The numerical accuracy and performance of these sub-

domains has been fully quantified, and used as the basis for novel adaptive remeshing strategies, which are capable of tracking crack tip propagation across significant length-scales.

To illustrate the utility of these advances, a detailed examination of cracking along wavy interfaces was conducted to quantify the potential benefits of crack deflections driven by local geometric features. The cohesive simulation framework is ideal for such studies, since a single simulation can determine whether an interface crack will advance along the interface or through the adjacent bulk material, without a priori assumptions or a broad parameter study involving an enormous range of crack configurations. The simulations on wavy interfaces demonstrate that for specific combinations of interface toughness, bulk toughness and interface waveforms, the far-field loading required to drive cracking along any path can be three times higher than that associated with a flat interface. Regime maps are presented to illustrate crack paths as a function of system properties, which provide useful guidance regarding the potential impact of non-planar interfaces in ceramic systems.

Even if the initial system consists of relatively planar coatings, local oxidation can lead to non-planar features that drive cracking; for example, cracks in environmental barrier coatings provide fast diffusion pathways to exposure the interior of ceramic-based components to reactants that subsequently form local oxide “bubbles”. These local domains drive cracking in the system due to the large volume change associated with the conversion of the monolithic ceramic to oxide: examples include Si being converted to SiO_2 and SiC being converted to SiO_2 . Unlike interface cracking problems where the geometry is presumably constant, oxidation phenomena involves several highly coupled phenomena: transport of reactants (i.e. diffusion), domain evolution (i.e. growth of the oxide), and creep in oxides at high temperature. To address this problem, a multiphysics simulation framework was developed with features that make it amenable to

future integration with cohesive cracking simulations. The framework integrates transport, boundary evolution and time-dependent constitutive descriptions within a single discretization scheme, such that behaviors can be evolved concurrently, with remeshing to account for large changes in system geometry.

The multiphysics framework was then applied to several case studies to gain insight regarding coupling between these phenomena. Oxidation of a bare circular fiber was used to explore the development of tensile stresses on the outer surface of the evolving oxide layer, which likely plays a role in fiber degradation. This study is compared with previous analytical models and shown to be highly accurate. Also, local oxidation at the tip of a crack in an environmental barrier coating was simulated to quantify the effects of crack density and layer diffusivities. A key finding is that local tensile stresses in the underlying substrate (which experiences oxidation) are significant and reach a peak at a critical time. That is, the local oxide domain must be large enough to induce significant stress concentrations, requiring a finite time based on oxide growth kinetics. At longer times, however, creep relaxation intervenes such that local stresses decay. The simulations have important implications for experimental studies of oxidation in ceramic-based components, since cracking may or may not occur based on system parameters and experimental design. Peak stress concentrations have been tabulated as a function of system properties to guide future efforts in this area.

Contents

| | |
|---|-------------|
| Curriculum Vitae | vii |
| Abstract | viii |
| List of Figures | xiii |
| 1 Introduction | 1 |
| 1.1 Motivation and Key Challenges | 1 |
| 1.2 Scope and Key Contributions | 7 |
| 2 Advances in Cohesive Zone Modeling | 11 |
| 2.1 Introduction | 11 |
| 2.2 A Brief Overview of Computational Cost | 16 |
| 2.3 Effects of Cohesive Law Description | 20 |
| 2.4 Interactions of Domain Size and Loading Rate | 28 |
| 2.5 Local Mesh Refinement | 32 |
| 2.6 Adaptive Remeshing | 34 |
| 2.7 Conclusions | 43 |
| 3 Kinking vs. Delamination Along Non-planar Interfaces | 46 |
| 3.1 Introduction | 46 |
| 3.2 Simulation Framework | 52 |
| 3.3 Results | 60 |
| 3.4 Conclusions for Wavy Interfaces | 68 |
| 3.5 Cracking in Wavy Lamellae with Adaptive Remeshing | 70 |
| 4 Multiphysics Framework for Integrated Diffusion, Oxidation and Creep | 76 |
| 4.1 Introduction | 76 |
| 4.2 Transport and Oxidation | 80 |
| 4.3 Mechanical Deformation | 87 |
| 4.4 Module Integration | 90 |
| 4.5 Framework Verification | 99 |

| | | |
|----------|--|------------|
| 4.6 | Simulation Convergence | 106 |
| 4.7 | Conclusion | 111 |
| 5 | Oxidation Under a Protective Coating | 112 |
| 5.1 | Introduction | 112 |
| 5.2 | Oxide Growth Under Barrier Coatings Based on One-Dimensional Diffusion | 118 |
| 5.3 | Multiphysics Simulation Framework | 124 |
| 5.4 | Oxide Growth & Stress Evolution Under Cracked Films | 131 |
| 5.5 | Discussion | 143 |
| 5.6 | Concluding remarks | 151 |
| 6 | Summary and Recommendations for Future Work | 153 |
| A | DCZM User Manual | 158 |
| A.1 | Introduction | 158 |
| A.2 | Installation | 158 |
| A.3 | Input Files | 160 |
| A.4 | Running the Solver | 163 |
| A.5 | Conclusion | 164 |
| B | Multiphysics Solver User Manual | 165 |
| B.1 | Introduction | 165 |
| B.2 | Installation | 165 |
| B.3 | Running the Solver | 166 |
| B.4 | Example Geometries | 167 |
| B.5 | Input Files | 167 |
| B.6 | Summary of Outputs | 168 |
| B.7 | Troubleshooting | 169 |
| | Bibliography | 171 |

List of Figures

| | | |
|-----|---|----|
| 1.1 | Schematic of CMC failure mechanisms. | 3 |
| 1.2 | Schematic illustration of an environmental barrier coating on a woven ceramic matrix composite with different characteristic failure mechanisms. | 4 |
| 2.1 | Plot of the computational cost as a function of the number of elements in the fracture length scale for the simulations used in the mesh resolution study in Chapter 3 [11]. The dashed line shows n_{cz}^3 scaling, which arises from the combined spatial and temporal impact of element size. | 19 |
| 2.2 | Plots of the trapezoidal cohesive law, illustrating (a) the normal traction-separation law with tangential quantities held fixed, and (b) the tangential traction-separation relationship with normal quantities held fixed. Note the characteristic plateau, as well as the fact that normal and tangential behaviors are coupled; this coupling dictates a critical fracture energy that is independent of mode-mixity, commonly referred to as an ideally brittle solid. | 22 |
| 2.3 | Schematic of the bilinear traction separation law for (a) normal and (b) shear tractions. The bilinear traction law does not have the plateau of the trapezoidal law but the behavior is still coupled. | 24 |
| 2.4 | Comparison of the behavior of the (a) trapezoidal and (b) bilinear cohesive laws with an identical mesh and a characteristic simulation from Chapter 3. The plateau of the trapezoidal law leads to many parallel, connected paths of cohesive elements with $\delta > \delta_Y$. The bilinear law does not have a plateau and all displacement is localized in a single path. | 25 |
| 2.5 | Plot of the percent error in applied load leading to fracture in a characteristic wavy interface simulation of Chapter 3 and [11] | 27 |
| 2.6 | Plot of the critical load to fracture as a function of loading rate showing the small domain size allows loading rates to be used that are much higher than typically considered for quasi-static behavior. | 30 |

| | | |
|------|--|----|
| 2.7 | Plots of (a) the critical load and (b) the relative cost for various r/ℓ_{cz} ratios. Both plots are normalized to $r/\ell_{cz} = 10$ and show a representative disk simulation. A course mesh of $n_{cz} \approx 30$ was used resulting in some noise in the critical load. Error bars on the critical load reflect the output frequency of the simulation. An estimate of the cost is shown based on an estimate of the number of elements in the refined zone. | 34 |
| 2.8 | Schematic of the crack tip tracking algorithm. (a) The elastic elements that border newly failed cohesive elements are output from the solver. (b) The elastic elements are paired up across the crack and the failed cohesive elements are found. (c) The failed cohesive elements are then assembled into a connected path and the crack tip advances. | 36 |
| 2.9 | Schematic illustration of the process used to extract the new geometry. (a) The cracked and separated element pairs are found from the results. (b) The cohesive elements joining the paired elements are found and enforced as segments in the new mesh. (c) A new mesh is created enforcing these segments. | 37 |
| 2.10 | Illustration of an old mesh mapped onto a new mesh. The previous crack path and activated cohesive elements are enforced between the two meshes. This enforced path is used to group elements in the new mesh so that correct displacements are applied on either side. | 38 |
| 2.11 | Schematic flowchart of the dynamic meshing algorithm. Once a cracking simulation is run, the new geometry is extracted. A new mesh is then created and equilibrium nodal displacements are found. The nodal velocities from the old mesh are mapped onto the new mesh and the next cracking simulation is run. | 40 |
| 2.12 | Comparisons of crack advance simulations with static refined meshes and dynamically updated meshes (adaptive refinement): (a) the predicted applied load relative to the benchmark, and (b) the cost of the simulations. | 41 |
| 2.13 | Example of the dynamic meshing algorithm tracking an arbitrary number of cracks. The top row shows the initial geometry and mesh. The middle and bottom rows show the algorithm tracking one or multiple cracks depending of the input properties along with the final mesh for each simulation. | 42 |
| 2.14 | Progression of various meshes used showing refined regions growing and shrinking as needed. | 44 |

| | | |
|-----|---|----|
| 3.1 | (a) Schematic illustration of the parent crack, subjected to pure mode I loading, (b) Schematic of the tested geometry with applied mode I displacements on the outer boundary; the radius is assumed to be much smaller than the layer thickness, (c) Close-up view of the region of wavy interface crack; the amplitude and wavelength are presumed to much smaller than the parent crack length, (d) Close-up view of the fine mesh region to capture the transition between interface cracking and kinking illustrating the transition between interface delamination and kinking; the fracture length scale is assumed to be much smaller than the wavelength. | 48 |
| 3.2 | Results for purely elastic calculations (no cohesive zones): (a) energy release rate at the crack tip normalized by that of the parent crack as a function of crack tip position, (b) phase angle (mode-mix) of the interface crack tip as a function of crack tip position. | 49 |
| 3.3 | Map of the amplitudes where crack face contact occurs for various loading phase angles. The boundary curve was determined by interpolating between amplitude and phase angle combinations that resulted in contact and those that did not. | 51 |
| 3.4 | Traction-displacement behavior of the cohesive elements embedded among elements: tangential displacements affect normal traction, and vice versa, such that the work to failure of the elements is constant regardless of loading. | 54 |
| 3.5 | (a) Effect of loading rate on predicted fracture toughness; these results are used to identify loading rate that will reduce computation time while still recovering quasi-static rests. (b) Effect of mesh density on predicted toughness; these results are used to set the size of the elements in the refined zone. | 58 |
| 3.6 | Far field energy release rate associated with crack advance, normalized by the toughness of the interface, for three different toughness ratios Γ_b/Γ_i . Open symbols indicate kinking out of the interface; solid symbols indicate interface cracking. | 61 |
| 3.7 | Apparent toughness needed to drive cracking as a function of the bulk-to-interface toughness ratio Γ_b/Γ_i for an initial crack tip position of (a) $x/L = 0$ and (b) $x/L = 0.375$ | 65 |
| 3.8 | Plots of the predicted transition toughness ratio from Pro et al. [9], He and Hutchinson [42] and simulations for (a) $A/L = 0.25$ and (b) $A/L = 0.5$ as well as (c) the calculated kink angle (from [42]) as a function of crack tip position. | 66 |
| 3.9 | Plot of the maximum load that can be applied before the crack grows unstably for three different toughness ratios Γ_b/Γ_i compared the data from [72]. Open symbols indicate kinking out of the interface; solid symbols indicate interface cracking. | 67 |

| | | |
|------|---|-----|
| 3.10 | Plot of the maximum applied toughness for a given toughness ratio and amplitude. The size of each bubble is scaled to the ratio of the critical toughness G_c to the interface toughness Γ_i . Closed symbols indicate interface failure while open symbols represent kinking. | 69 |
| 3.11 | Schematic illustration of the geometry used for the dinosaur teeth example. At the macroscopic level, a cracked rectangular domain is loaded under displacement control. The width W and height H are significantly larger than the initial crack length a , approximating an isolated crack in a semi-infinite domain. On a smaller length scale there are alternating lamella with wavy interfaces with the primary axis of the wavy interface rotated an angle θ counterclockwise from the plane of the crack. The lamella have thickness h , amplitude A and wavelength L . For reference, the point $x=0$ refers to the midpoint between a trough and a peak of an interface. . . . | 71 |
| 3.12 | An example (a) crack resistance curve with the three crack growth domains marked and (b) the crack path plotted on top of the microstructure with the domains indicated. | 74 |
| 3.13 | Plots of (a) the crack growth resistance for different d/h values showing that increasing the d/h value reduces the needed load for unstable growth. (b) shows the crack speed normalized by the wave speed of material #1 for the same d/h values. | 75 |
| 4.1 | Schematic illustration of oxide growth on silicon carbide. Oxidant molecules like oxygen and water diffuse through the existing oxide, react with silicon carbide and advance the oxide/silicon carbide interface (shown in red) inward. | 79 |
| 4.2 | Plots of the relative stress relaxation for various initial stresses, σ_o at a temperature of 1000°C . The non-linear viscosity is stress dependent and higher initial stresses lead to faster initial relaxation rates. | 79 |
| 4.3 | Flowchart showing the major steps in the multiphysics framework. Finding the new geometry and solving for displacements each have many tasks which are also shown. | 91 |
| 4.4 | Schematic illustration of oxide growth procedure on a subset of the oxide/silicon carbide interface. (a) depicts how the elemental fluxes J_1 and J_2 are averaged to get the average flux at the current node \bar{J} . The normal of the growth interface n_1 and n_2 are averaged to get \bar{n} . (b) shows how the average flux is projected onto the average normal resulting in the perpendicular flux \bar{J}_\perp . (c) shows how the motion of the interface is calculated and applied to the node. This procedure is repeated at all nodes along the boundary. | 93 |
| 4.5 | Plot showing the simulated oxide thickness on a flat surface as a function of time and the thickness predicted by the Deal-Grove model. There is little difference between the simulation and Deal-Grove model. | 101 |

| | | |
|------|---|-----|
| 4.6 | Plot of the calculated stress relaxation with simulation results for a one dimensional compression test. The initial stress of 25.3 GPa is representative of the stress that occurs immediately after oxidation. | 102 |
| 4.7 | Schematic illustration of oxidation on a silicon carbide fiber. At time $t = t_o$ shown in (a), the fiber has a radius of r_o and initial oxide thickness of h_o . At some later time shown in (b), the oxide thickness has grown and the outer diameter has grown due to the volumetric expansion that occurs during oxidation. | 103 |
| 4.8 | (a) Shows the oxide thickness as a function of time on a 3 μm radius fiber. (b) Shows the inward flux normalized by the flux assumed by the Deal-Grove model. The curvature causes the flux to not be constant unlike the Deal-Grove model. | 104 |
| 4.9 | Contour plots of the hoop stress at (a) an early time where the oxide layer is in compression, (b) when the hoop stress at the outer layer is approximately zero, (c) the hoop stress at the outer boundary hits a maximum of around 1 GPa and (d) the tensile layer thickens with time. | 105 |
| 4.10 | Plots showing the hoop stress in the oxide as the simulation progresses. (a) shows the hoop stresses through the thickness of oxide at various times. (b) shows the hoop stress at the outer edge of the oxide as a function of time as well as data from Hay [14] and a previous implementation. | 106 |
| 4.11 | Plots of the deviation in thickness on growth of a round fiber compared to that predicted by the Deal-Grove model as a function of (a) element area for a fixed time increment and (b) a function of time increment for two different element sizes. | 108 |
| 4.12 | Plots of the tensile stress at the outer edge of the fiber as a function of (a) element area with a fixed time increment and (b) time increment with two different element sizes. | 109 |
| 4.13 | Plots of the stress profiles near the oxidation boundary for (a) a constant element size and various time steps and (b) constant time step with various element sizes. | 110 |
| 5.1 | Schematic illustration of an environmental barrier coating (EBC) on a ceramic matrix composite (CMC), with a close-up view of the tip of a crack in the coating. The crack serves as a fast-diffusion pathway that drives local oxidation, which generates high stresses that promote substrate cracking and/or coating delamination. | 113 |

| | | |
|-----|---|-----|
| 5.2 | (a) Schematic diagram of the multi-physics framework that simulates chemical transport, domain evolution driven by oxidation, and creep relaxation in newly formed oxides. Adaptive remeshing is used to enable accurate simulations of large geometric changes. (b) Schematic diagram of different time scales that must be synchronized to integrate various physical processes. In the present work, numerical efficiency is sufficient to allow the use of a single time step for all processes. | 116 |
| 5.3 | (a) Predicted diffusion constants down a narrow channel that combine Knudsen and molecular diffusion mechanisms, for 1% water vapor and 100% molecular oxygen at two temperatures. (b) Prediction of oxide thickness under an intact barrier coating for several coating diffusivities; for $D_b > 10^{-5} cm^2/s$, oxide growth is virtually identical to a bare surface, and controlled by diffusivity in the oxide layer itself. This implies reactant transport down even the most narrow of crack openings is likely never a limiting time scale for the growth of the oxide at the tip of a crack. . . . | 123 |
| 5.4 | (a) Schematic illustration of the idealized geometry used in the current simulations, with transport properties and boundary conditions. (b) Schematic illustration of the mechanical properties and boundary conditions used to predict stress evolution in the system. | 132 |
| 5.5 | Contours of concentration for at three different times, for three different coating diffusivities and relatively large crack spacing (low crack density); low diffusivity in the coating leads to local oxide thickness under the crack that is much greater than that under the intact portion of the coating. . . | 134 |
| 5.6 | Contours of direct stress parallel to the coating interface at three different times for three different coating diffusivities; increasing the coating diffusivity leads to smaller differences between the crack and intact regions of the coating, and smaller tensile stresses in the substrate. The system is subjected to a fixed strain in the x-direction that generates $\sigma_{xx} = 1 GPa$ at the remote boundary. | 135 |
| 5.7 | Contours of direct stress perpendicular to the coating interface at three different times for three different coating diffusivities; increasing the coating diffusivity leads to smaller differences between the crack and intact regions of the coating, and smaller tensile stresses in the substrate. The system is subjected to a fixed strain in the x-direction that generates $\sigma_{xx} = 1 GPa$ at the remote boundary. | 136 |
| 5.8 | Contours of shear stress at three different times for three different coating diffusivities; increasing the coating diffusivity leads to smaller differences between the crack and intact regions of the coating, and smaller tensile stresses in the substrate. The system is subjected to a fixed strain in the x-direction that generates $\sigma_{xx} = 1 GPa$ at the remote boundary. | 137 |

| | | |
|------|--|-----|
| 5.9 | Oxide thickness directly under the crack and under the intact coating (at the symmetry plane between cracks) for several different coating diffusivities. The oxide at the crack location is smaller than bare surface oxidation due to lateral diffusion away from the crack plane. | 139 |
| 5.10 | Oxide thickness directly under the crack, normalized by the thickness predicted by Deal-Grove for an uncoated substrate; four crack spacings are shown for four different coating diffusivity values. The decay with time in the relative thickness arises from two-dimensional diffusion away from the crack tip; for high coating diffusivity, two-dimensional diffusion is effectively eliminated by fast transport through the coating itself. Crack spacings greater than at least one coating thickness play a relatively minor role. | 140 |
| 5.11 | Direct stress distribution in the loading direction at the plane of the crack, for several times and several different coating diffusivity values. Higher coating diffusivity reduces the gradient in oxide thickness from the crack plane to the intact region of the coating, and mitigates stress. | 142 |
| 5.12 | (a) Peak stress experienced in the substrate directly ahead of the local oxide that forms at the tip of the crack as a function of time, for several coating diffusivities and a single crack spacing. (b) Peak stress experienced in the substrate directly ahead of the local oxide that forms at the tip of the crack as a function of time, for several crack spacings and a single diffusivity. (c) Stress concentration in the substrate as a function of time, for several different levels of remote loading and two coating diffusivities. Higher levels of remote loading somewhat mitigate the stress concentration by enhancing creep relaxation in the oxide. | 144 |
| 5.13 | (a) Peak stress concentration in the substrate as a function of crack spacing, for coatings with negligible diffusivity and that quoted in the literature for YbSi_2O_7 . The elastic stress concentration refers to that associated with the crack tip at a distance equal to the initial oxide thickness. (b) Comparison of stress concentrations at 300 hours (which is quite close to the peak concentration) as function of coating diffusivity, for three different levels of remote loading. | 145 |
| 5.14 | Stress in the substrate immediate beneath the moving oxide boundary for three different coating diffusivities and two different values of remote tension. | 147 |
| B.1 | Example single fiber geometry. | 167 |
| B.2 | Example geometry included in this alpha version. | 168 |

Chapter 1

Introduction

1.1 Motivation and Key Challenges

Ceramic composites offer significant opportunities to improve the efficiency of propulsion and power generation turbines, by potentially allowing higher operating temperatures than those currently enabled by coated metallic components. The principle advantage of ceramics is their creep resistance at very high temperatures (e.g. in excess of 1500°C); limitations attributed to the low fracture toughness of monolithic ceramics can be overcome using high strength fibers which limit flaw size, embedded in a ceramic matrix that controls fiber orientation [1,2].

A typical stress-strain curve for a high-quality ceramic matrix composite is shown in Figure 1.1; the underlying mechanisms controlling notable features of the curve are well established [3–5]. The loss of linearity in the material is associated with matrix cracking; the matrix will have embedded flaws (e.g. porosity from processing steps) that are much larger than the fibers and hence fail at much smaller applied loads. When the interface between the fibers and matrix is properly controlled, these composites exhibit significant retained strength even after the matrix cracks, due to the fact that matrix cracks will

deflect into the fiber/matrix interface and avoid propagation of the matrix crack through the fibers. The modulus of the composite after matrix crack is reduced due to load transfer to the fibers near the matrix crack plane; this load transfer is controlled by frictional sliding between the matrix and fibers in the debonded region near the matrix cracks. Ultimately, the fibers fragment at high loads; due to the statistical nature of fiber fracture, fiber failures occur across a broad range of longitudinal locations prior to complete rupture usually precipitated at a matrix crack plane. Thus, a hallmark of high-quality ceramic composites is a tortuous fracture surface comprising fibers that fail away from the matrix plane (at loads between the matrix cracking stress and composite ultimate strength) and subsequently experience ‘pull-out’ from the matrix.

The interface between the matrix and the composite is paramount; interfaces that debond too easily will eliminate any load-sharing capability of the matrix (between matrix crack planes) and increase the likelihood of fiber failure. Interfaces that are too strong (either in the sense that debonding is prohibited or the frictional sliding stress is too high to permit sliding) lead to premature fiber failure, since cracks initiated in the matrix at low loads will propagate through the fibers. To control this behavior, fibers are coated with a thin layer prior to matrix infiltration. While there are many fiber/coating/matrix compositions that exhibit the failure mechanisms outlined above, arguably the most promising system consists of SiC fibers coated with BN and embedded in a SiC matrix [1]. The focus of this work is on this system, although it should be appreciated that many of the studies are relevant to other ceramic composite systems.

While the above mechanical response for unidirectional composites reflects the ideal, in actual service additional factors intervene to alter composite response. First, to control transverse behavior fiber tows are woven and hence include non-planar interfaces between the tows and the matrix. Naturally, controlling the nature of cracking between woven tows and surrounding matrix is important to ensure overall function. Second, operation

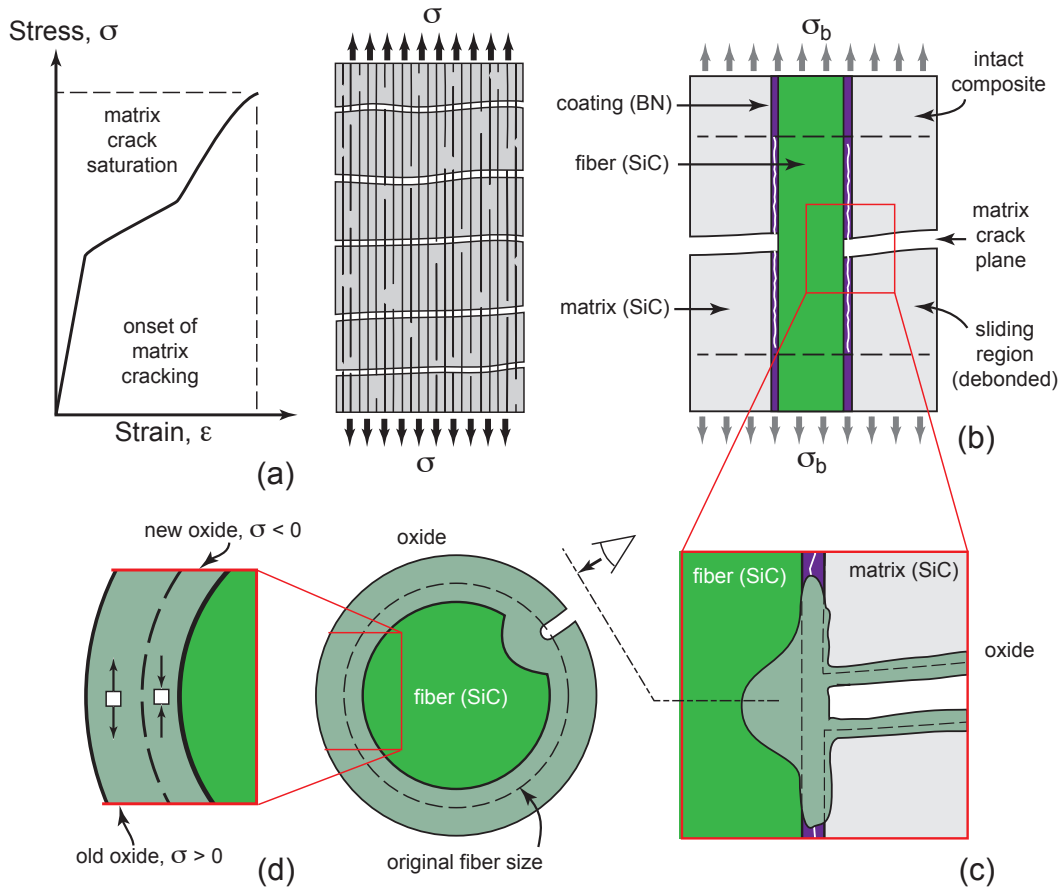


Figure 1.1: Schematic of CMC failure mechanisms.

in harsh environments with water vapor, oxygen and contaminants leads to chemical reactions that degrade composite response. To limit these deleterious, environmental barrier coatings (EBCs) are utilized to eliminate (or at least substantially decrease) the ingress of reactive species [1]. Figure 1.2 provides a schematic overview of composite and coating systems that are emerging as the foundation of robust components. The figure also includes illustrations of numerous degradation mechanisms that motivate the simulations developed in this dissertation.

Arguably, the central design consideration for the composite/coating system is control of the thermal expansion mismatch between the composite (or substrate) and the coating




Figure 1.2: Schematic illustration of an environmental barrier coating on a woven ceramic matrix composite with different characteristic failure mechanisms.

[1, 6]. The combination of high operating temperatures and use of brittle materials implies that even modest differences in coefficients of thermal expansion can produce sufficient stored elastic strain energy to drive a host of cracking mechanisms. The two most prevalent coating failure mechanisms are mud-cracking (cracking transverse to the plane of the coating), and delamination (cracking parallel to the plane of the coating). In and of themselves, these cracks do not degrade the load capacity of the system; rather, they provide pathways for chemical attack in the interior of the composite which in turn can trigger a host of mechanisms that lead to system failure. The principal motivation for the simulation frameworks in this thesis is establishing links between cracking phenomena and subsequent chemical attack; naturally, feedback loops may exist wherein chemical attack promotes subsequent additional cracking.

Two of the most prevalent concerns relating to chemical attack are oxidation (and subsequent volatilization) of the BN coating and oxidation of the SiC fibers, which are associated with numerous deleterious effects [1]. Loss of the BN coating reduces stress transfer to intact sections of the matrix, increasingly the likelihood of fiber fragmentation.

The formation of oxides is particularly troublesome, for two reasons: oxide formation both reduces the volume of load-carrying SiC and introduces misfit stresses that drive cracking. The misfit stress associated with the molar volume change of SiC to SiO₂ is dramatic and much larger than typical component stresses. These stresses drive many of the mechanisms shown schematically in Figure 1.2: for example, continued penetration of mud-cracks that originate in the coating, delamination and premature matrix/fiber failure. Though of a different chemical nature, contamination by mineral depositions on the coating are an important related concern (and drive many aspects of coating selection) [1, 7]; volumetric changes and formation of new phases with mismatched coefficients of thermal expansion can drive similar coupling between chemical and mechanical responses.

The vast majority of analyses of these effects exploit highly idealized geometries, principally the analysis of (i) coatings and reaction products of uniform thickness and (ii) crack geometries that are prescribed a priori [6, 8]. As shown schematically in Figure 1.2, these idealizations likely miss important aspects of local geometry that control the mechanisms underlying failure. These aspects related to local geometric features, such as non-planar (“wavy”) interfaces that arise from fiber tows, coating deposition techniques, and local oxide regions generated by differences in chemical transport down existing cracks, channels formed by coating recession, local penetration by contaminants, etc. Such local features naturally influence the formation and evolution of complex crack networks, which ultimately govern durability. The details of these crack networks and evolution from one crack orientation to another (e.g. mud-cracking to delamination) are critical, in light of the fact that some cracks will be far more benign than others in the contexts of system lifetime. Furthermore, prior work on domain evolution - i.e. the growth of oxide domains or recession fronts - is rather limited due to the fact that traditional analysis methods require broad parameter studies capture the effect of time-dependent changes in local feature geometry. (Similar behaviors also occur upon contamination by

mineral deposits (CMAS); although not a focus of this work, many of the tools designed to address the limitations of prior parameter studies should prove useful in this context as well.)

It is very important to note that these environmentally-driven changes in system geometry occur at elevated temperature, and, due to volumetric changes associated with reaction products, can generate significant stresses even without changes in operating temperature. Conventional wisdom with regards to cracking driven by CTE mismatch between phases is largely based on the assumption that stresses relax at high temperature, and stresses that drive cracking are generated upon cooling. While such behaviors are undoubtedly relevant to the development of the system shown in Figure 1.2, the generation of stress at elevated temperatures involves additional behaviors that influence the time-history of stress in the system. Specifically, cracks may occur at high temperature that serve as the initial damage that precipitates more widespread damage upon cooling. Naturally, creep relaxation at elevated temperature plays an important role in determining whether or not such cracking is likely to occur. In contrast to cooling from a stress-free state, this competition is comparatively unexplored and serves as a central motivation of the frameworks described in this work.

In this broader context, understanding environmentally-driven failure mechanisms of ceramic systems requires new advances in integrated simulation frameworks, which include the ability to: (i) predict the evolution of crack networks, (ii) quantify the role of local, non-planar changes in geometry, (iii) track species transport underlying the formation of new phases, and (iv) define domain evolution and the stresses associated with the formation of new phases at elevated temperatures. In the next section, the scope of efforts to address these challenges is described in conjunction with key contributions from this work.

1.2 Scope and Key Contributions

The broader context outlined in the preceding section provides a rationale for the following activities, which are reasonably correlated with subsequent chapters.

- **Advances in cohesive zone frameworks for simulating complex cracking modes:** Prior work has demonstrated that distributed cohesive zone approaches are effective in capturing crack path evolution, particularly for an isolated crack near interfaces [9]. However, early efforts were limited to tracking a single crack over relatively small distances, due to the requirement that finely focused meshes are needed at the crack tip. Further, the shape of the cohesive law used in cohesive descriptions plays an important role in crack bifurcations; previous treatments utilizing an elastic-perfectly plastic description are prone to delays in crack branching. To address these limitations, methods have been developed to: (i) incorporate softening cohesive descriptions that preserve characteristics of brittle fracture while improving branching behaviors, and (ii) integrate adaptive re-meshing algorithms that translate refined regions as cracks advance, and critically, enable multiple cracks to be tracked simultaneously. This includes both branching and coalescence of different cracks. These advances, in conjunction with speed enhancements that reduce simulation time prior to the onset of crack advance, significantly broaden the applicable range of simulations that predict the evolution of cracking networks. A detailed description of these advances is provided in Chapter 2.
- **Crack path evolution near non-planar interfaces:** The cohesive zone frameworks described above have been utilized to conduct a comprehensive study of crack path selection at or near non-planar bimaterial interfaces. As demonstrated by the study, the topology “wavy” interfaces play an important role in coating adhesion,

biological composites (e.g. dinosaur teeth [10]) and woven composites. Prior simulations have illustrated that cracks along or near such interfaces can either deflect out of the interface to enter adjacent bulk material, or remain “trapped” at the interface; however, there was limited understanding of the quantitative connections between interface shape, the toughness of the interface and adjacent bulk material, and the far-field loading that is required to reach criticality. Toward that end, this work has generated complete regime maps that clearly illustrate combinations of system properties that dictate failure mode (i.e. bulk versus interface cracking). These results illustrate that the geometry of wavy interfaces can provide significant improvements with regards to the magnitude of far-field loading required to drive continued cracking. This work has identified critical values of system properties needed to realize these benefits, which in some cases can imply critical loads that are a factor of three higher than those associated with flat (planar) interfaces. This finding has important implications for coating design and processing targets, and illustrates that interface topology is an important consideration in simulating coating failures. A detailed description of these advances is provided in Chapter 3 (and has appeared in [11]).

- **Multi-physics framework simulating chemo-mechanical behaviors and oxidation of bare fibers:** To address the behaviors outlined in Section 1.1, a simulation framework that integrates multiple physical phenomena has been developed. The core strategy in developing the framework was to utilize numerical descriptions that provide straight-forward communication of chemical behaviors into the thermomechanical analysis framework used to predict cracking. As such, a new finite element framework for tracking species diffusion was generated using triangular elements; this was combined with adaptive remeshing algorithms

developed for the cohesive zone framework, which enable simulations of domain evolution. The full two-dimensional field of species transport allows predictions of oxide growth (or recession, though this feature was not studied in this work) with concomitant remeshing to allow for large changes in geometry. In the present work, Deal-Grove descriptions of oxidation were incorporated within the transport framework to simulate domain evolution as a function of chemical environment [12, 13]. These results are then ported into the thermomechanical framework to simultaneously predict the stresses generated by oxidation. Toward this end, the integration of non-linear creep behaviors into the thermomechanical analysis framework is an important advance that enables predictions of time-dependent stress evolution at high temperature. To benchmark the framework and provide an illustration of its utility, simulations of oxidation on circular fibers were conducted and compared to analytical models in the literature. In this scenario, large compressive stresses are generated in the initial stages of the oxide formation; as early layers of oxide are pushed radially outward by new oxide forming at the reaction front, large geometry changes generate tensile stresses. The advances in adaptive remeshing described above were also incorporated, and illustrate the importance of accounting for large geometry changes. Similarly, as expected, creep plays an important role; the simulations in Chapter 4 illustrates that stress states driving cracking are impacted by multi-axial creep behaviors, indicating the importance of fully-developed numerical solutions (even for planar problems).

- **Internal oxidation beneath EBCs and near intersections of fibers and matrix cracks:** By and large previous studies of the mechanical impact of oxide formation have relied on analyses that ignore the impact of growing oxide on adjacent features, such as oxidation of a single bare fiber [14, 15], the surface of

planar substrates [13] and uniform oxide layers growing within a multilayer [16]. A notable exception is the analysis of oxide within a matrix crack by Xu, et al., who addressed the possibility that oxides wedge open matrix cracks and promote fiber failures. This latter study illustrates the motivation of the study of internal oxidation in this work; the large volumetric changes associated with the formation of oxide in a constrained space clearly plays a role in damage mechanisms associated with internal oxidation. Toward that end, simulations of oxide formation at the tip of a mud-crack in an EBC/substrate system have been conducted to examine the consequences of the formation of local oxide domains in brittle systems. Though idealized, the geometry is also similar to the formation of oxide on a fiber bridging a matrix crack. By varying the relative diffusivity of reactant species down the crack and through the EBC, various oxide domain shapes were generated using the multiphysics framework described in Chapter 4, without any a priori assumptions regarding the spatial distribution of reactant concentrations. A key finding of the work is that the time-scale of oxide growth and subsequent creep relaxation play a critical role in the stresses imposed on adjacent material by oxide formation. Under certain conditions, the peak tensile stresses are reached at short times; this has important implications for the interpretation of testing EBCs and composites subjected to environmental degradation. A complete accounting of the results from this study is provided in Chapter 5.

Chapter 2

Advances in Cohesive Zone Modeling

2.1 Introduction

In cohesive zone modeling of fracture, a traditional finite element framework is modified to include one-dimensional, non-linear cohesive elements that represent material rupture [3, 17–28]. Whereas the traditional elements in the framework are governed by macroscopic continuum material descriptions, the cohesive elements are described by a traction-separation (cohesive) law, which describes the fracture process.

The traction-separation law is defined such that the cohesive element reproduces linear elasticity for loads significantly below a critical strength, in that the effective properties of the combined model reproduce those in the conventional elements themselves. This is essentially enforced by ensuring the initial elastic (linear) portion of the cohesive has stiffness that far exceeds that of the conventional elements. To simulate rupture, the cohesive law is defined to be non-linear; once the peak stress is reached, further deformation results in either constant tractions (akin to perfect plasticity) or decreasing tractions

representing material failure. Continued deformation in the cohesive element ultimately drive the tractions across the adjacent surfaces to zero, and the material (cohesive element) is considered to be fractured. While a great many phenomenological cohesive descriptions exist, they all involve a characteristic peak strength, σ_o , and critical separation, δ_o [29]. The critical fracture energy of the solid, Γ , scales with their product (with a scaling constant depending on the exact shape of the cohesive law.)

Consider a pre-crack under Mode I loading, with cohesive elements lying on the plane of crack advance. As the specimen is loaded, the cohesive elements adjacent to the crack surpass their peak strength; the physical distance encompassed by these elements can be thought of as the fracture process zone, or fracture length scale. When the fracture length scale is small in comparison to all other dimensions in the simulation, small-scale yielding is said to persist and crack stability is entirely governed by the critical fracture energy defined by the cohesive law. Critically, in this limit, the precise values of the cohesive strength σ_o and the critical separation δ_o are immaterial, only their product. This implies that for a ceramic, where cohesive strengths are on the order of 10 GPa and critical separations are on the order of 1 nm, equivalent fracture behavior can be achieved with a cohesive strength of 100 MPa and a critical separation of 100 nm – ***provided that the fracture length scale*** is much smaller than any other dimension.

The fracture length scale can be defined in terms of the cohesive parameters to first order as:

$$\ell_{cz} = \frac{E\Gamma}{\sigma_o^2} \quad (2.1)$$

where E is the modulus of the solid. A key feature of this scaling relationship is that equivalent fracture results can be obtained for brittle materials with artificially low cohesive strengths, and the corresponding fracture process zone will be significantly larger.

This provides significant computational advantage; the requisite mesh resolution scales with the size of the fracture length-scale, implying much faster simulations of fracture. The drawback is that if this artifice is pushed too far, the fracture length scale will no longer be smaller than other physical dimensions, such that the conditions of small-scale yielding may be violated. In such cases, crack stability is influenced by the cohesive strength itself and the critical fracture energy, and the results will be non-physical for artificially low cohesive strengths.

It should be noted that in such simulations, the fracture process can only occur at locations and orientations where cohesive elements are present. If the failure plane is known in advance (due to the presence of a weak interface, for example), cohesive elements need only to be included along the anticipated fracture plane to study the stability of that plane [25,30,31]. In this work, cohesive elements are included between all conventional continuum elements, such that any surface indicated by element boundaries in the mesh represents a potential failure plane similar to Xu and Needleman and Pro et al [9, 32]. Prior work has demonstrated that this approach accurately captures the direction of crack kinking out of an interface, provided the mesh is of sufficient resolution and the element boundaries are randomly oriented [9]. In these prior studies of kink crack nucleation, only the direction of crack advance is unknown; the physical location is fixed by the position of the crack tip, and only the onset of crack advance was of interest. This was because the kink crack is unstable once it emerges from the parent crack, and would not regain stability until the crack tip was at a position comparable to large specimen dimensions, which were not considered.

In contrast, when simulating stable or unsteady crack growth in cohesive simulations, cohesive zones must be distributed throughout the domain, with sufficient mesh density (i.e. element sizes well below the fracture length scale) in all orientations and at all points in the mesh. Simply put, including a refined region at the tip of the initial crack with

course exterior regions become inaccurate as the crack propagates into the coarse region. The obvious fix is a refined mesh everywhere, but this can be prohibitively expensive. As such, simulations of stable or quasi-stable crack growth in geometries with multiple length-scales (e.g. wavy interfaces, grains, etc.) are few and far between.

On top of these length scale considerations, simulations of unsteady crack growth can be extremely expensive due to widely varying time-scales associated with crack advance and arrest. Quasi-static loading rates may correspond to millimeters per second or micrometers per second (with respect to boundary velocities that define loading rate); once initiated, crack advance occurs at hundreds to thousands of meters per second (i.e. near the speed of sound). This poses a conundrum: on the one hand, if highly resolved time-stepping is used to capture dynamic fracture events, the initial loading of the specimen to the point of fracture onset is extremely expensive. On the other hand, if the loading rate is artificially increased, inertial effects impact crack advance; the unintentional generation of elastic waves at the loading boundary interfere with the crack tip, such that quasi-static behaviors are not accurately captured. These complications are compounded by the well-established link between acceptable time-step and element size, which must be managed to ensure accurate propagation of stress across the specimen.

In summary, there are a myriad of factors that can make cohesive simulations computationally expensive. Section 2.2 describes the inherent scaling of computational times to shed a quantitative light on the above comments. Sections 2.3-2.6 detail several numerical studies designed to identify approaches to speed up simulations and broaden the bandwidth of problems that can be analyzed. These studies address the following concepts/strategies:

- **Cohesive law description:** While attention is limited to small-scale yielding conditions where the critical fracture energy is the sole governor of crack advance, the

shape of the cohesive law is shown to have a strong impact on computational speed and accuracy requirements. Specifically, crack branching that arises as a consequence of diffuse cohesive failures (or near failures) leads to stochastic variability in simulation results. Section 2.3 illustrates that a bilinear softening cohesive description can mitigate this behavior, and produce precision in critical loads. The associated study involved hundreds of simulations with different mesh densities and different random meshes, and ultimately provides important insight on variability inherent to the method and how to minimize it. This study played a critical role in improving the accuracy of the wavy interface simulations in Chapter 3.

- **Interaction of loading rate and domain size:** The efficiency of fracture simulations can be improved by judicious choice of loading rate and domain size. This relates to the time required to simulate the propagation of elastic waves from the remote boundary to the crack tip. The numerical study in Section 2.4 demonstrates that one can reduce computational time by first solving a quasi-static problem to compute the displacements on a ring surrounding a crack tip, which by-passes the time-stepping needed to propagate a wave to that location. The full dynamic problem is then analyzed with initial conditions over an outer thick annulus obtained from the quasi-static solution. In essence, one truncates the simulation time for initial loading. It should be noted that one cannot simply simulate a smaller specimen, since small-scale yielding solutions may be violated.
- **Local mesh refinement:** After the step size, the number of total elements in the simulation has the strongest influence on the cost of a particular simulation. Naturally, local crack advance requires highly resolved meshes, but there has been relatively little effort placed on quantifying the impacts of local mesh refinement. (A typical approach simply ensures that there are “a sufficient number” of elements

near the crack tip to ensure convergent behavior; however, these have not produced quantitative guidelines for how large the refined region must be in comparison to specimen size.) Understanding the impact of the size of the refined region relative to other features and the fracture length-scale is critical to developing adaptive remeshing strategies.

- **Adaptive remeshing:** Once a strategy for local mesh refinement has been identified, adaptive remeshing becomes possible; that is, the refined region of mesh surrounding a crack tip can be propagated in space as the crack tip advances. This eliminates the need for broadly refined meshes and saves considerable expense, as refined regions are eliminated as the crack tip position evolves. Section 2.6 describes adaptive remeshing algorithms; the central challenge is to map the solution from one mesh to a new mesh while preserving accuracy. The potential utility of the method is illustrated by a fracture near a castellated weak interface, reminiscent of intra-granular cracking in polycrystalline brittle materials. The simulations track multiple crack fronts and capture phenomena such as the emergence of a microcrack ahead of a dominant crack tip.

2.2 A Brief Overview of Computational Cost

Calculations of the precise computational cost of a cohesive zone simulation prior to actually running a simulation can be challenging, as the cost is strongly dependent on the speed and number of processors, optimization of the solver, frequency that data is written to the hard disk, etc. However, recognizing that the number of elements and number of time steps are the dominant factors, one can make a reasonably close approximation of

computational cost as follows. The cost, t_{sim} , is then

$$t_{sim} = C n_{elements} n_{steps} \quad (2.2)$$

where C is a constant that depends on the solver, processor specifics and the number of threads (for parallel implementations), $n_{elements}$ is the number of elements and n_{steps} is the number of time steps. While the requisite number of elements and number of time steps is determined by the problem to be studied, the constant C needs to be determined for a specific computer.

Ideally, C is inversely proportional to the speed of the processors and the number of threads used. Unfortunately, there is a cost associated with increasing the thread count due to overhead caused by the synchronization of threads, passing of information between different threads and any bandwidth bottlenecks between the processor and memory. If too many threads are created, the simulation cost does not decrease even with large simulations. The optimal number of threads varies with the number of elements and needs to be determined experimentally. An exhaustive study of the effects of the number of threads on performance is not presented here, although it was found that for a typical server (Intel Xeon Processors with 8-28 cores each running at 2-3.4 GHz) the optimal number of threads was between four and eight. For the work that follows, it is assumed that 8 threads are used.

Unlike the pre-factor C in eqn. 2.2, the number of elements and number of time steps is dependent on the particular problem of interest. It is well known that elements much smaller than the fracture length scale, ℓ_{cz} , are needed to accurately model fracture. As substantiated by numerous studies [9,11], including the mesh resolution study in Chapter 3 for wavy interfaces, an effective rule of thumb is that the required number of elements, n_{cz} , along a ray in the fracture process zone is at least 10. To achieve equiaxed, randomly

oriented triangles that allow for arbitrary crack advances, the number of elements is proportional to n_{cz}^2 . For dynamic formulations (as considered in this dissertation), the maximum stable time step scales with the element size; in essence, the time step must be smaller than the time required for an elastic wave to propagate across a given element. (Extensive details regarding the relationship between element size and minimum stable time step can be found in [33].) Since the largest stable time step is proportional to $1/n_{cz}$, the number of time steps required for a fixed physical duration of a given simulation is proportional to n_{cz} . With regards to the number of elements, therefore, the total simulation cost scales as:

$$t_{sim} \propto n_{cz}^3 \quad (2.3)$$

Evidence of this scaling is shown in Figure 2.1, which uses data from the simulations in Chapter 3 and [11].

Note that the scaling of simulations with the fewest elements is slightly worse than that implied by cases with the largest number of elements. That is, fitting eqn. 2.2 to data with fewer elements will overpredict the platform-dependent prefactor C . This is due to overhead costs incurred during initialization, such as launching or synchronizing threads, which are fixed and independent of mesh resolution. While the present study does not tabulate these costs (as they are platform-specific), they are a comparatively small contribution when the mesh resolution is sufficient to obtain accurate predictions of crack stability.

While accurate predictions can be guaranteed by a high number of elements in the fracture length, the scaling indicates the cost can quickly become prohibitive. For this reason, it is usually beneficial to conduct a mesh resolution study for the problem at hand, such that one can strike a balance between speed and accuracy. An example can be found in [9] or Chapter 3 [11]. The work in Chapter 3 uses about 100 elements in the

fracture length leading to a simulation cost of around 50 hours as this was determined to be the best balance between cost and accuracy.

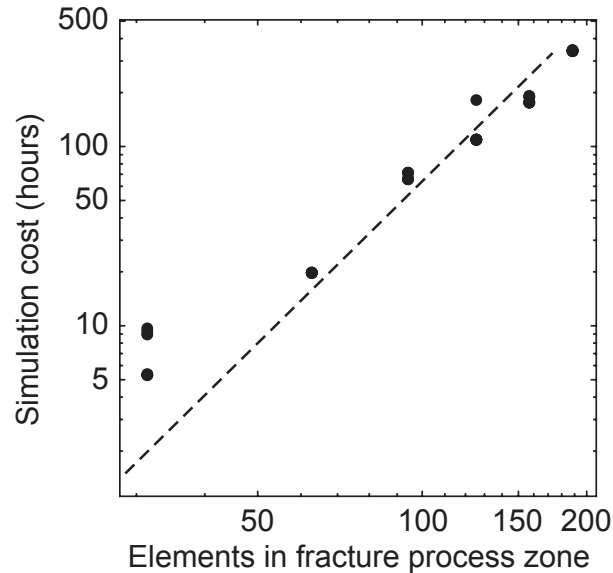


Figure 2.1: Plot of the computational cost as a function of the number of elements in the fracture length scale for the simulations used in the mesh resolution study in Chapter 3 [11]. The dashed line shows n_{cz}^3 scaling, which arises from the combined spatial and temporal impact of element size.

It is worth noting that Figure 2.1 shows the cost for a particular geometry and loading state. Different problems will have slightly different cost curves, but will be similar to the one shown here. For example, for the simulations used to generate Figure 2.1, the mesh consists of a refined mesh near a crack tip, which dominates the number of elements in the model. The size of the refined region is $\sim \ell_{cz}$. Increasing the size of the refined zone size will shift the results to greater simulation costs in proportion to the fraction of elements in the refined region to the total number of elements, while shrinking the size of the refined zone would shift the curve down. Section 2.6 provides additional details about the effects of such local refinements, including their impact on computational cost.

2.3 Effects of Cohesive Law Description

As briefly described in Section 2.1, the details of the cohesive law describing the rupture process (i.e. a loss of traction across adjacent surfaces) are immaterial to crack stability when the fracture length-scale is smaller than any other dimension in the problem. Under such ‘small-scale yielding’ conditions, only the area under the traction-separation curve, or fracture energy, is relevant. That said, even though the macroscopic conditions needed for crack advance may not be affected by the shape of the cohesive law, local details of deformation within the fracture process zone (cohesive elements that are on the verge of rupture) *are* affected by the shape of the cohesive law, and these can have a significant impact on *numerical* stability.

In this section, simulations of crack kinking are used to illustrate potential issues arising from the shape of the cohesive law, using trapezoidal and triangular cohesive descriptions. (For additional details on the impact of cohesive descriptions, see the work of Chandra et al, who summarized many different cohesive descriptions and the types of problems each has been used to solve [29]. The first laws were those proposed by Barenblatt and Dugdale based on interatomic forces or yield stress, respectively. Other laws have been used to tailor the results to specific features of material rupture.) To the author’s knowledge, the impact of the mathematical form of the cohesive law on dynamic crack path selection (e.g. the direction of an emergent kink crack) has not yet been reported.

Motivated by the work of Tvergaard and Hutchinson [34], a trapezoidal cohesive law was used in the work of Pro et al [9] which adopted an identical numerical framework to that utilized in this work. It is worth noting that the unsatisfactory performance of this law in the simulations of Chapter 3 that led to the comparisons in this section. In the

trapezoidal cohesive descriptions, the normal and tangential tractions are given by:

$$\sigma_n = \begin{cases} k\delta_n & , \sqrt{\delta_n^2 + \delta_t^2} \leq \delta_Y \\ k\delta_Y \frac{\delta_n}{\sqrt{\delta_n^2 + \delta_t^2}} & , \delta_Y < \sqrt{\delta_n^2 + \delta_t^2} \leq \delta_o \\ k\delta_n \left(\frac{\delta_Y + \delta_o}{\sqrt{\delta_n^2 + \delta_t^2}} - 1 \right) & , \delta_o < \sqrt{\delta_n^2 + \delta_t^2} \leq \delta_R \\ 0 & , \delta_R < \sqrt{\delta_n^2 + \delta_t^2} \end{cases} \quad (2.4)$$

$$\sigma_t = \begin{cases} k\delta_t & , \sqrt{\delta_n^2 + \delta_t^2} \leq \delta_Y \\ k\delta_Y \frac{\delta_t}{\sqrt{\delta_n^2 + \delta_t^2}} & , \delta_Y < \sqrt{\delta_n^2 + \delta_t^2} \leq \delta_o \\ k\delta_t \left(\frac{\delta_Y + \delta_o}{\sqrt{\delta_n^2 + \delta_t^2}} - 1 \right) & , \delta_o < \sqrt{\delta_n^2 + \delta_t^2} \leq \delta_R \\ 0 & , \delta_R < \sqrt{\delta_n^2 + \delta_t^2} \end{cases} \quad (2.5)$$

where k is the linear elastic stiffness of the initial traction-separation response, δ_Y is the separation at which the cohesive stress becomes constant (and equal to $\sigma_o = k\delta_Y$), δ_o is the separation associated with the onset of the loss of traction, and δ_R is the separation at which the traction is zero, i.e. the rupture separation. Here, the subscript n refers to tractions and displacements acting in the direction normal to the separating surfaces, and t refers to those quantities tangential to the surface. These traction-separation relationships shown graphically in Figure 2.2.

With purely normal or purely tangential separation, the trapezoidal law has a short initial region where the traction increases linearly with separation until a peak stress is reached. The traction is then constant at this plateau stress with additional separation until a short linear decrease. The length of the plateau governs the critical fracture energy of the cohesive law. The cohesive stiffness, k , must be high enough to recover macroscopic elasticity; otherwise, the cohesive elements will introduce additional compliance into the

model and the effective modulus of the continuum will be smaller than that assigned to adjacent conventional elements.

As can be seen in Equations 2.4 and 2.5, the normal and tangential tractions are coupled so that failure occurs when the total displacement is δ_R . In other words, the total work to fracture is always constant regardless of the ratio of normal to tangential displacement. This is shown in Figure 2.2 by the curves, which hold one displacement fixed while varying the other. In Figure 2.2a, increasing the ratio of tangential displacement to the rupture displacement from $\delta_t/\delta_R = 0$ to $\delta_t/\delta_R = 1$ decreases both the slope of the curve as well as its maximum. Similar behavior can be seen in Figure 2.2b while varying δ_n/δ_R . It should be noted that Figure 2.2 shows an artificially low k_e in order to illustrate the shape of the curve. With realistic material properties and useful element sizes, the plateau where $\sigma_n = \sigma_o$ or $\sigma_t = \sigma_o$ is very large. In most cases, the plateau length ($\delta_o - \delta_Y$) is greater than 99% of δ_R .

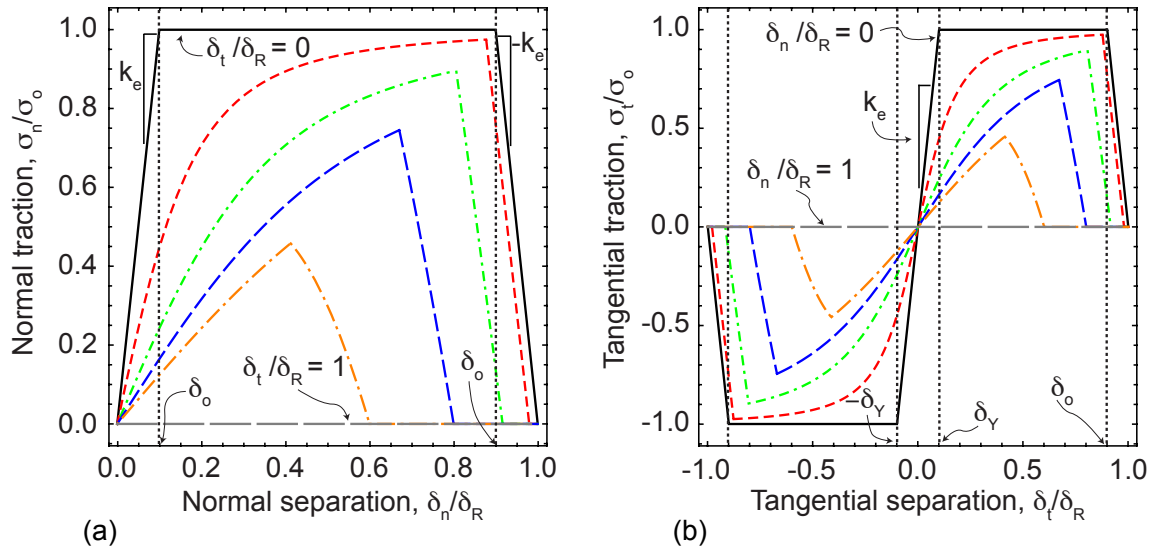


Figure 2.2: Plots of the trapezoidal cohesive law, illustrating (a) the normal traction-separation law with tangential quantities held fixed, and (b) the tangential traction-separation relationship with normal quantities held fixed. Note the characteristic plateau, as well as the fact that normal and tangential behaviors are coupled; this coupling dictates a critical fracture energy that is independent of mode-mixity, commonly referred to as an ideally brittle solid.

The second cohesive law considered in this work is a bilinear description illustrated in Figure 2.3; this description was used in the work described in Chapter 3 and [11], which considers fracture along wavy interfaces. The normal and tangential tractions (σ_n and σ_t respectively) in the bilinear cohesive law are given by:

$$\sigma_n = \begin{cases} k_e \delta_n & , \sqrt{\delta_n^2 + \delta_t^2} \leq \delta_o \\ \frac{k \delta_o \delta_n}{\sqrt{\delta_n^2 + \delta_t^2}} \frac{(\delta_R - \sqrt{\delta_n^2 + \delta_t^2})}{(\delta_R - \delta_o)} & , \delta_o < \sqrt{\delta_n^2 + \delta_t^2} \leq \delta_R \\ 0 & , \delta_R < \sqrt{\delta_n^2 + \delta_t^2} \end{cases} \quad (2.6)$$

$$\sigma_t = \begin{cases} k_e \delta_t & , \sqrt{\delta_n^2 + \delta_t^2} \leq \delta_o \\ \frac{k \delta_o \delta_t}{\sqrt{\delta_n^2 + \delta_t^2}} \frac{(\delta_R - \sqrt{\delta_n^2 + \delta_t^2})}{(\delta_R - \delta_o)} & , \delta_o < \sqrt{\delta_n^2 + \delta_t^2} \leq \delta_R \\ 0 & , \delta_R < \sqrt{\delta_n^2 + \delta_t^2} \end{cases} \quad (2.7)$$

where as before, δ_n and δ_t are the normal and tangential displacements, k_e is the cohesive stiffness, δ_o is the displacement resulting in the peak stress σ_o and δ_R is the rupture displacement. As with the trapezoidal law, the normal and tangential tractions are coupled together so that the total work to failure is the same regardless of the balance of normal and tangential displacements. For pure normal or tangential displacement, the magnitude of the traction linearly increases up to δ_o . After that point, the magnitude decreases linearly until the rupture displacement is reached. Mixing normal and tangential displacements results in lower tractions in those direction while keeping the total area under the curves (i.e. the work to failure) constant.

While the trapezoidal law can be used to accurately predict the macroscopic loading parameters corresponding to the onset of crack kinking as shown by Pro et al [9], the local details of deformation are different from those obtained with the bilinear law; an illustrative example of this is shown in Figure 2.4. The important difference to note is the

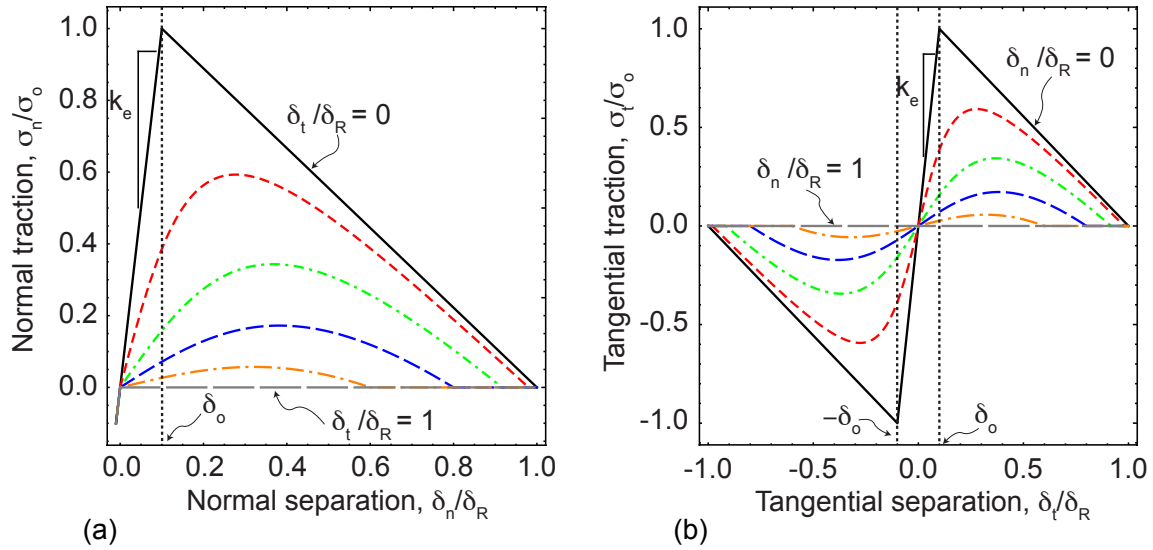


Figure 2.3: Schematic of the bilinear traction separation law for (a) normal and (b) shear tractions. The bilinear traction law does not have the plateau of the trapezoidal law but the behavior is still coupled.

presence of multiple parallel cracks emerging from the crack tip on the left (trapezoidal description), which are absent in the crack tip on the right (bilinear description).

These qualitative differences in the distribution of damage arise from the shape of the cohesive law itself, and can impact the mesh resolution needed to achieve accurate results. The large plateau of the trapezoidal law leads to multiple parallel paths of connected cohesive elements that have all ‘yielded’. This can be seen in Figure 2.4a where it appears there are multiple cracks ahead of the dominant crack tip. The dominant potential crack path is the lower one but a secondary path is seen just above it. Multiple other parallel potential paths can also be seen. These cohesive elements have not yet failed so the crack has not grown. When a cohesive element finally reaches the end of the plateau and the traction starts to decrease, separation stored in other elements can be transferred to the failing element. When this phenomena occurs, a significant amount of energy can be stored in the cohesive elements prior to failing before being released quickly into the failed element. This energy storage artificially increases the apparent

load needed to drive fracture.




Figure 2.4: Comparison of the behavior of the (a) trapezoidal and (b) bilinear cohesive laws with an identical mesh and a characteristic simulation from Chapter 3. The plateau of the trapezoidal law leads to many parallel, connected paths of cohesive elements with $\delta > \delta_Y$. The bilinear law does not have a plateau and all displacement is localized in a single path.

This non-ideal behavior is primarily a consequence of inserting cohesive elements between every pair of elastic elements. By doing this, there is a connected two dimensional network of cohesive elements. Evidence of it can be seen in Figure 4 of [9] for two different loading phase angles. Other work that uses the trapezoidal law but only includes cohesive elements along a single plane, such as [34], does not encounter this. Furthermore, the longer plateaus greatly increase the likelihood of multiple approximately parallel cohesive elements yielding. For the brittle materials of interest, the initial stiffness (i.e. initial slope of the traction curve) must be very high in order to recover elastic behavior leading to the plateau region being orders of magnitude longer than the initial region.

The appearance of the multiple competing cohesive element paths occurs regardless of the element size or loading. Overall, however, shrinking the element size reduces the influence of the phenomena. By reducing the element size, the likelihood of cohesive

elements that are highly stressed near the crack tip and closely aligned with the preferred crack growth direction increases which would also increase the likelihood of this phenomena. However, this is mitigated by the higher prevalence of cohesive elements that are slightly *misaligned* to the preferred growth direction. The energy penalty for the crack to deviate from the preferred path is smaller for denser meshes than it would be for coarser meshes. The misalignment of a cohesive element also means that the loading it experiences is not purely normal, but a combination of normal and tangential displacements. The mix of loading results in a reduction or complete elimination of the plateau. A shorter plateau length leads to localization of deformation in the cohesive element due to the softening that occurs later on in the cohesive law. The localization of deformation mitigates the other potential crack paths and a dominate path forms sooner.

While the artificial toughening caused by this phenomena can be mitigated by decreasing the element size, shrinking the element size has a drastic effect on the computational cost, as shown previously. Another way to mitigate it would be to run multiple simulations with the same mesh density, but different meshes, and average the result. However, this is still expensive and may still artificially increase the load to failure if the phenomena occurs. A different cohesive law without the long plateau would behave better and allow accurate results without resorting to smaller elements.

Unlike the trapezoidal cohesive law, the bilinear law does not have a long plateau stress under any loading conditions. As soon as the peak stress is reached further separation leads to a decrease in tractions. It is much less likely that multiple parallel cohesive elements would pass the peak stress at the same time since the decreasing tractions would lead to localization of the separation. This can be seen in Figure 2.4b which uses the same mesh as Figure 2.4a. Using the same mesh with the bilinear cohesive law results in a localization of the deformation and a single crack path emerging as shown in Figure 2.4b. With the localization of deformation the mesh dependence decreases. Slight mesh

dependence still exists due to the possibility of a discontinuous crack as shown in Figure 2.4b, but the variation in load to failure is much smaller. This can be seen in Figure 2.5. By reducing this variation, the same results can be determined from one simulation instead of many simulations with the trapezoidal law.

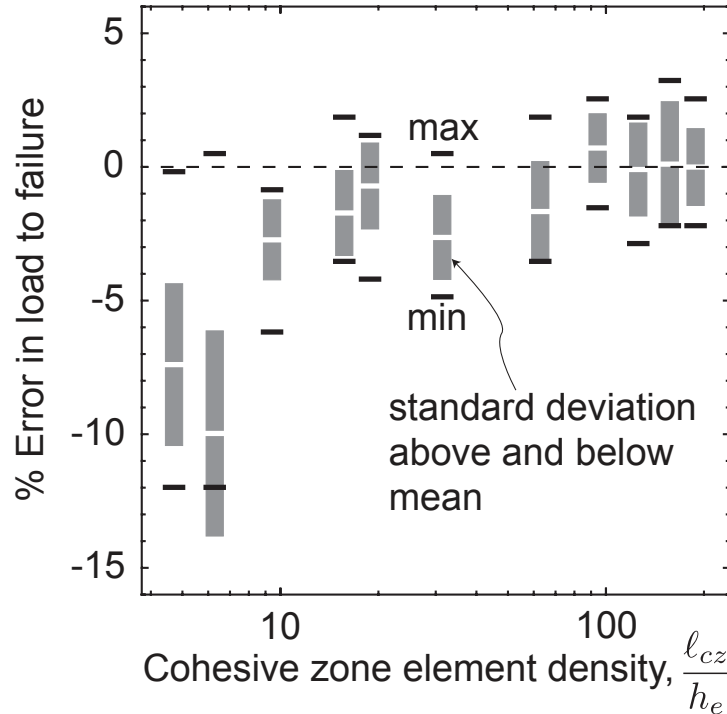


Figure 2.5: Plot of the percent error in applied load leading to fracture in a characteristic wavy interface simulation of Chapter 3 and [11]

It is worth noting that, without careful consideration of the underlying physical scaling, the inputs to the specific cohesive law used also impact the computational cost of the simulation. As mentioned previously, the initial stiffness, k_e , must be high enough to recover elastic behavior at lower stresses. Otherwise, the apparent modulus will be lower than desired. However, if the stiffness is too high, a smaller time step is needed to properly resolve dynamic waves. A good rule of thumb is that the stiffness must be $k_e \approx 50E/h_e$ where E is the Young's modulus and h_e is the characteristic element size. As another example, the peak stress for a truly brittle material, σ_o , should arguably be

10 GPa or greater. As discussed in the previous section, small elements greatly increase the computational cost of the simulation. In order to manage the computational cost by allowing for larger elements, the peak stress used in both Pro et al [9] and the wavy interface work of Chapter 3 [11] is limited to 2 GPa. Pro showed that using a lower peak stress can still yield the same results as long as certain conditions are met (see Figure 3.6 in [33]). This lower peak stress allows for the use of larger elements thereby reducing the computational cost. Along the same lines, the absolute value of the critical fracture energy or toughness Γ can impact simulation time. On the one hand, smaller toughnesses can lead to faster simulations under fixed loading rate because fracture loads and hence simulation times are smaller. On the other hand, the fracture length scale, ℓ_{cz} , scales linearly with the toughness. Since the characteristic element size must be significantly smaller than the fracture length scale and element size greatly impacts the computational cost, smaller toughnesses generally increase the computational cost. These concerns are applicable to both the trapezoidal and bilinear cohesive laws.

The key conclusion from the study of the impact of the cohesive law is that the artificial toughening arising from distributed damage with the trapezoidal law can be mitigated with smaller elements, while the bilinear law appears to avoid this effect completely. Accurate results can be acquired without having to use very small elements or having to run multiple meshes and average the results. Further work in this dissertation is focused only on using the bilinear cohesive law.

2.4 Interactions of Domain Size and Loading Rate

Explicit dynamic simulations of fracture using cohesive elements are favored over implicit or static frameworks because they provide an efficient numerical approach to deal with the non-linear nature of material rupture, and capture inertial effects needed to ac-

curately assess crack stability. That said, they are inherently computationally expensive for brittle materials, due to limitations on time steps required for stability and the fine mesh resolution needed to accurately predict crack advance. In this section, the basic scaling relationships that guide effective computational approaches are reviewed, followed by a discussion of simple modeling practices that limit computational cost.

The Newmark- β algorithm used to solve dynamic equations of motion dictates that stable time-integration is achieved when the numerical time-step is:

$$\Delta t_{max} \leq \frac{\Delta x}{\alpha c} \quad (2.8)$$

where Δx is the size of the smallest element in the domain, $c = \sqrt{E/\rho}$ is the elastic wavespeed in the material (where E is the elastic modulus and ρ is the density), and α is a dimensionless factor on the order of 10-100 that arises as a consequence of the stiffness of the cohesive elements. (Note that the stiffness of the cohesive elements is set to be larger than that of the continuum elements in order to recapture continuum elasticity.)

For accurate predictions of cohesive fracture, $\Delta x = \ell_{cz}/N_{cz}$, where $N_{cz} = 10 - 50$ is the number of elements in the fracture process zone (of size ℓ_{cz}). Even when cohesive parameters are adjusted to artificially increase the size of the fracture process zone (e.g., 1 μm), the maximum allowable step size for brittle materials with $c \geq 10^3 m/s$ is typically on the order of 10^{-7} seconds. This implies that realistic quasi-static loading rates involve billions of time steps, such that significant advantage is gained if artificially high loading rates can be utilized. The upper limit on allowable loading rate is naturally defined by the point at which inertial effects start to influence the response. For a purely elastic domain of characteristic size R , quasi-static response (i.e. not influenced by inertial effects) is ensured when the applied strain rate is significantly less than c/R . (This is equivalent to

saying the velocity at the boundaries must be significantly less than the wavespeed.)

The impact of loading rate and accuracy of the above scaling relationships are illustrated in Figure 2.6, which plots the critical load to fracture for a center-cracked panel versus the applied strain rate (defined as the velocity of the panel divided by its height). With $E = 200$ GPa, $\rho = 4$ Mg/m³ and a domain size of 1.25 mm, the characteristic strain rate for inertial effects, $c/R = 5.6 \times 10^6$. At this loading rate, the critical rupture load is 25% higher than expected for quasi-static loading. At much lower loading rates, say 10^4 , one obtains results virtually identical to the theoretical quasi-static results. Hence, even though the corresponding applied strain rate rate is far greater than generally considered for quasi-static behaviors, the small domain size implies that inertial effects are limited.




Figure 2.6: Plot of the critical load to fracture as a function of loading rate showing the small domain size allows loading rates to be used that are much higher than typically considered for quasi-static behavior.

This example clearly demonstrates that simulation domain size should be kept as small as possible. As the domain size increases, the maximum allowable loading rate de-

creases; slower loading rates imply more time steps to reach the remote strain associated with fracture. Using the above scaling, and the fact that the maximum allowable strain is $\sim 10^2 c/R$, one can demonstrate that the minimum number of time steps scales roughly as $N_{steps} = 15R/\ell_c$, assuming the loading rate is as fast as possible without inertial effects coming into play.

For brittle fracture in a monolithic, isotropic material, the domain size must be at least a factor of ten larger than the fracture process zone size to be in the limit of small-scale yielding. Hence, irrespective of any chosen parameters or normalization, there typically must be at least several thousand timesteps to capture quasi-static loading. This conclusion is based on the assumption that the domain is dynamically loaded with constant strain rate, starting from zero load (i.e. $\epsilon = \dot{\epsilon}t$). Naturally, if an accurate idea of fracture load is known a priori, this can be reduced by running a purely static simulation with a prescribed load known to be below that required from crack advance. The result of the purely static calculation is then used as the initial condition for a full dynamic solution, eliminating the early time steps prior to the onset of fracture.

The implicit cost of larger domains in dynamic fracture simulations is compounded in problems with a characteristic length-scale that is larger than the fracture process zone. The wavy interface problem in Chapter 3 is an example of such a case. Simulations must be much larger than the amplitude and wavelength to capture the effect of a continuous periodic interface with, implying minimal domain sizes that are much, much larger than that required for small-scaling yielding (i.e. brittle fracture). That is, if the wavelength must be at least $10\ell_{cz}$ to capture fracture, and the domain must span at least ten periods, the minimum domain size is $100\ell_{cz}$.

An effective strategy for dealing with this inherent length-scale problem (which cannot be removed via normalization or re-scaling) involves two components. First, locally refined meshes should be utilized at the length scale of the fracture process zone, with

course meshes at the next scale. Strategies for defining local refinements are outlined Section 2.5. Second, purely elastic sub-modeling should be used to identify boundary conditions on a smaller domain that is much larger than the fracture process zone but much smaller than physical boundaries.

For example, one can analyze a full component with tens of wavy interface periods to determine the elastic boundary boundary conditions on an inner annulus that contains only the interface period containing the crack tip. This tiered approach can lead to order of magnitude reductions in computational costs, and is practically unavoidable for problems with microstructural features on the order of several microns or more. Generally sub-domain modeling is clear simply based on the features of interest. The practice of local mesh refinement to reduce the mesh density requires greater attention, and is discussed in the next section. This sets the stage for adaptive remeshing, wherein in the local mesh region near the crack tip is propagated during the simulation to track crack growth that exceeds the initial refined region.

2.5 Local Mesh Refinement

As noted above, the cost of a particular simulation is proportional to the number of elements. One way to keep the number of elements tractable is to only place smaller elements near the fracture process zone and use larger elements elsewhere. In many problems, such as crack kinking, the direction of crack advance is not known beforehand such that an isotropic mesh should be used in a annulus surrounding the crack tip. If there are only small elements near the crack tip and much larger elements everywhere else, the number of elements in a mesh is dominated by the area near the tip. The

number of elements in a mesh can then be approximated by

$$n_{elements} \approx \pi \left(\frac{rn_{cz}}{\ell_{cz}} \right)^2 \quad (2.9)$$

for n_{cz} elements per fracture length scale ℓ_{cz} within a radius r of a crack tip. The size of the refined zone, r , and mesh density defined by r and n_{cz} needs to be large enough to not affect the results but not too large as to drastically increase the simulation cost.

To illustrate the potential benefits of a local refined zone, a pre-cracked disk of size $R = 10\ell_{cz}$ with a uniform mesh with 30 elements along a ray emanating from the crack tip is used as a baseline. A refined zone is then introduced within this disk, with 30 elements along a ray inside the refined region; as the size of the refined region is decreased, a coarse mesh is introduced between the local refined zone and the outer boundary at $R = 10\ell_{cz}$. The critical load to fracture for (using asymptotic mode I crack tip fields along the outer boundary) was then computed for each mesh to explore the accuracy and cost benefit of the refined zone.

The results are shown in Figure 2.7a. The error bars shown on the plot are a result of the output frequency of the simulation and are equal to $\pm 0.25\%$. The computational cost is shown in Figure 2.7b as a function of the size of the refined region and again normalized to the results with $R = 10\ell_{cz}$ (all cases have $n_{cz} = 30$). The results demonstrate that while a large refined region may be more accurate, it is also quite costly. Shrinking the refined region down greatly reduces the simulation cost. Refined regions do not need to be much larger than the fracture length in order to get meaningful results in a reasonable time.

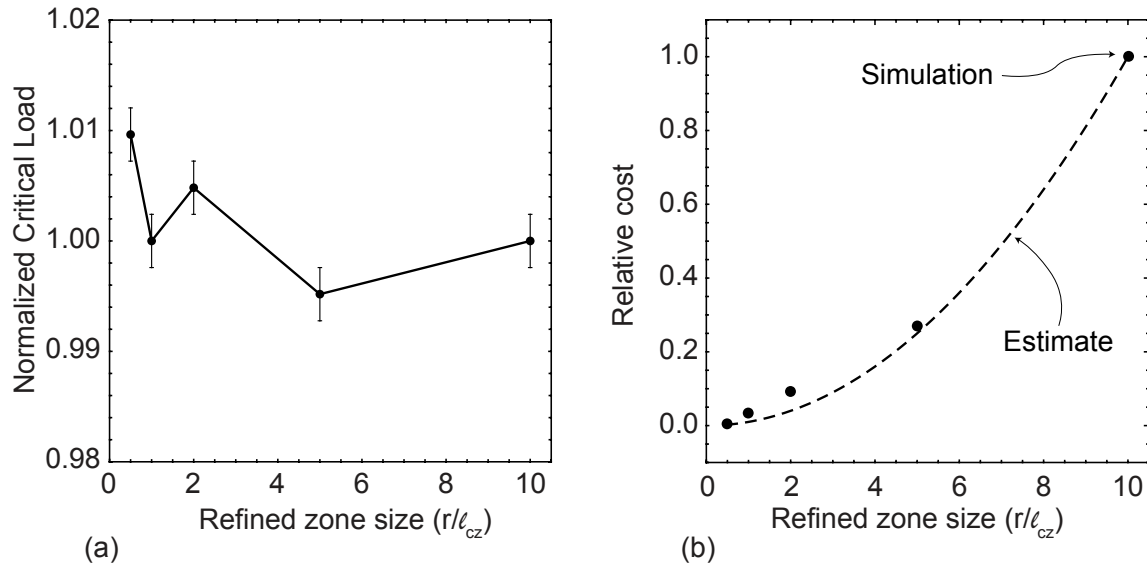


Figure 2.7: Plots of (a) the critical load and (b) the relative cost for various r/ℓ_{cz} ratios. Both plots are normalized to $r/\ell_{cz} = 10$ and show a representative disk simulation. A course mesh of $n_{cz} \approx 30$ was used resulting in some noise in the critical load. Error bars on the critical load reflect the output frequency of the simulation. An estimate of the cost is shown based on an estimate of the number of elements in the refined zone.

2.6 Adaptive Remeshing

The downside to the local mesh refinement discussed in Section 2.5 is that local refinements reduce the distance the crack can grow before remeshing is needed. This is unimportant when only the onset of crack advance is of interest, such as kinking tests of Pro et al [9]. In contrast, to simulate crack growth that spans multiple fracture lengths, refined regions must always be present at the crack tip; previous approaches have simply accepted the associated cost of a larger refined region that spans multiple fracture lengths. Once the crack propagates past the initial fracture process zone, the refined mesh at the position of the original crack tip is no longer necessary and represents computational “baggage” that must be carried through the simulation. As such, if the refined region can be propagated through the simulation to follow the crack tip, the number of elements can be greatly reduced and the cost becomes far more manageable.

To track the crack tip with a small locally refined mesh, a new mesh must be created intermittently as the crack grows, and the information transferred onto it from the old mesh. This includes not only spatial distributions of displacements and velocities, but also geometry information that defines the crack path. In the distributed cohesive zone method described by Pro [9, 33] and used in this work, the crack path is only defined implicitly based on nodal displacements and must be extracted from the simulation output. When a crack grows, as shown in Figure 2.8a, the elements that are connected to a failed cohesive element are identified as in Figure 2.8b. These elements are then paired together to find the failed cohesive elements. The cohesive elements are linked into a continuous path and connected to the previous crack path. The crack tip is then moved shown in Figure 2.8c.

As can be seen in Figure 2.8, there are pairs of continuum elements that are separated but not marked as connected to a failed cohesive element. These cohesive elements between these elements are inelastic, but not yet failed. The displacements in those elements exceed the critical separation δ_o . This information must be accounted for in the new mesh in order to preserve the correct cracking behavior. Thus, the combination of the crack path and partially separated element pairs constitute the relevant geometry that needs to be carried over to the new mesh.

Figure 2.9 shows the progression of determining the new relevant geometry. First, both the elastic elements that are connected to failed cohesive elements as well as those connected to inelastic cohesive elements are identified in Figure 2.9a. Those connected to failed cohesive elements are used to define the new crack path as well as identify cohesive elements that are inelastic. These different regions are highlighted in Figure 2.9b. Once the segments are identified, a new mesh, shown in Figure 2.9c, is created. These failed and inelastic segments are imposed as features of the new mesh to preserve continuity of the fracture process zone between both meshes.




Figure 2.8: Schematic of the crack tip tracking algorithm. (a) The elastic elements that border newly failed cohesive elements are output from the solver. (b) The elastic elements are paired up across the crack and the failed cohesive elements are found. (c) The failed cohesive elements are then assembled into a connected path and the crack tip advances.

Once the new mesh is created, the displacements and velocities from the old mesh must be mapped onto the new mesh. Unlike a traditional finite element mesh, the presence of the cohesive elements dictates that the displacement field is not continuous between elements. For example, if six elements share the same node in the mesh, there are six different sets of displacements and velocities for that one node corresponding to each element present. Far away from any failed or inelastic cohesive elements, the difference between the values is slight. However, along failed or yielded cohesive elements the displacements and velocities can be significantly different.

This implies is easier to consider the mapping procedure on an elemental basis instead of a nodal basis, since the presence of the cohesive elements allows for multiple nodes at

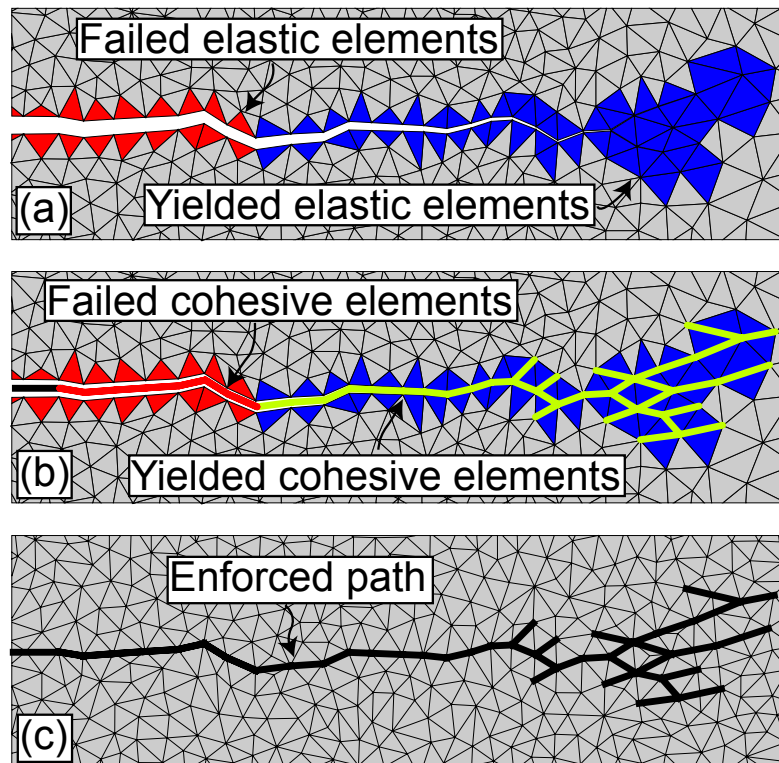


Figure 2.9: Schematic illustration of the process used to extract the new geometry. (a) The cracked and separated element pairs are found from the results. (b) The cohesive elements joining the paired elements are found and enforced as segments in the new mesh. (c) A new mesh is created enforcing these segments.

the same point in space. A graphical representation of how elements in an old and new mesh may align is shown in Figure 2.10. In order to map the displacements in each new element, the old mesh is checked to see if the new element matches an element in the old mesh. If so, the nodal values are taken directly from the nodes of the old element. If there is not a matching element, each node is located within the old mesh. If it exists within an old element but does not match an old node, the displacements and velocities are simply interpolated from the old solution.

If the node falls on an enforced path, such as the crack path, but does not match an old node, it falls on the boundary between two elements in the old mesh. The values for the new node are interpolated from the element in the old mesh that is on the same side

as the element in the new mesh. Finally, if the node is coincident with a node in the old mesh, the displacements and velocities for the new node are averaged from the nodal values in the old mesh. While performing this averaging, if the node is on an enforced path, the averaging only takes into account elements and nodes that are on the same side of the path as the current new element.




Figure 2.10: Illustration of an old mesh mapped onto a new mesh. The previous crack path and activated cohesive elements are enforced between the two meshes. This enforced path is used to group elements in the new mesh so that correct displacements are applied on either side.

Since the presence of the cohesive elements creates small gaps between the elastic elements, the displacement field is not continuous in the mapping procedure described above. This implies the imposed displacement field on the new mesh may not be in equilibrium. This can be corrected through the use of an elastic finite element solver that incorporates the cohesive elements as simply linear springs. During the equilibration step, the displacements along the outer boundary and the enforced path are fixed and only those away from the enforced path are changed. While this does not take the

nonlinear cohesive law into account, this adjusted displacement field is a much more accurate mapping of the prior state. This step is also very cheap to perform.

Once the solution has been mapped to the new mesh, all the pieces are in place for the dynamic meshing algorithm. The general flow of the dynamic meshing algorithm is shown in Figure 2.11. During a dynamic cohesive zone method simulation, the solution is monitored until crack growth reaches the boundary of the refined region surrounding the crack tip. The simulation (time-stepping) is then interrupted and the new crack geometry is extracted. A new mesh is created with the new crack geometry and displacements adjusted as above. Meanwhile, the old velocities are mapped directly onto the new mesh without adjustments. Once both the velocities and displacements are prescribed, time-stepping in the DCZM simulation is restarted. This loop can repeat continuously to allow for arbitrary crack growth of any length.

The dynamic meshing algorithm was implemented with a combination of C/C++ applications and Python scripts. Meshing was performed in a C application which utilized Jonathon Shewchuk's Triangle library for triangulation [35]. The static finite element solver was written in C++. The DCZM solver was also written in C++ and used in the work of Pro et al [9] and the wavy interface studies covered in Chapter 3 and [11]. Python was used to run the various C/C++ applications, map the solution onto new meshes and create input files for the DCZM solver. Incorporating all the various pieces together into one application written in C++ would result in additional, and possibly significant, cost reductions. However, that integration is left for future work. An example of the crack tracking algorithm in use is shown at the end of Chapter 3 studying wavy lamellar structures.

The applied load as a function of crack length for a simple center crack panel is shown in Figure 2.12a for static meshes with different refined region sizes and the dynamic meshing algorithm. The applied load is normalized by the applied load for unstable

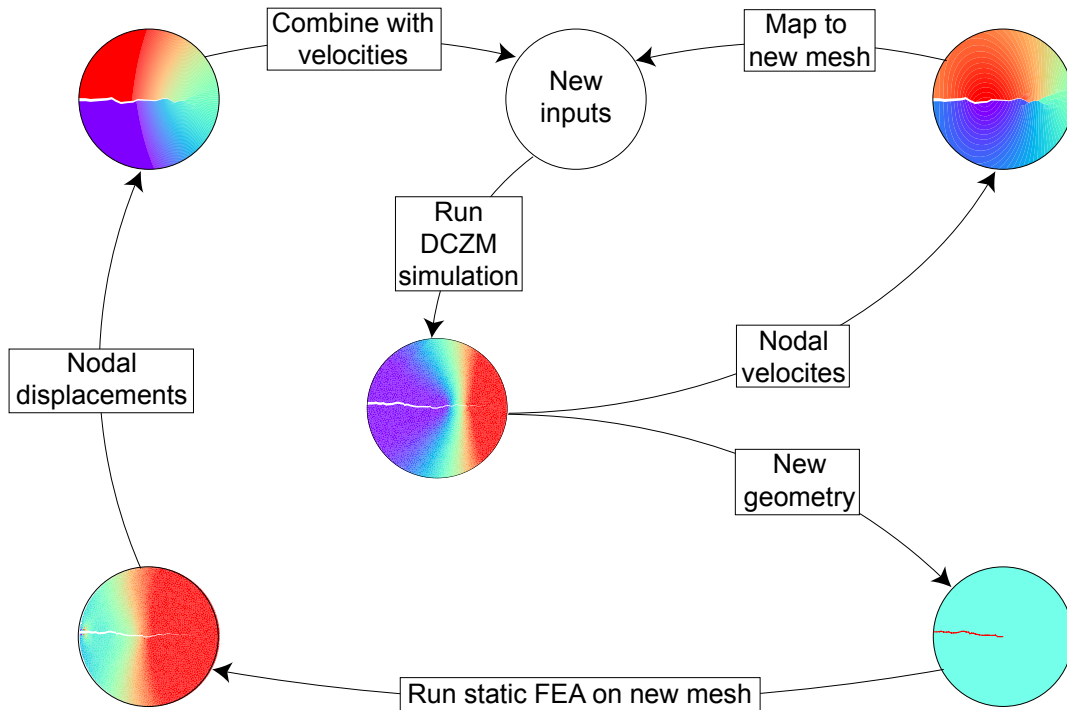


Figure 2.11: Schematic flowchart of the dynamic meshing algorithm. Once a cracking simulation is run, the new geometry is extracted. A new mesh is then created and equilibrium nodal displacements are found. The nodal velocities from the old mesh are mapped onto the new mesh and the next cracking simulation is run.

growth in the mesh with the largest refined region ($r/\ell_{cz} = 10$). Curves for each refined region size end when the crack leaves the refined region to reduce the effects of increasing element size. Refined regions of $r/\ell_{cz} = 0.5$ and $r/\ell_{cz} = 1$ are too small. They increase the required load to drive cracking and do not accurately show the transition to unstable growth. Larger refined regions of $r/\ell_{cz} = 2$ and $r/\ell_{cz} = 5$ show much closer results to the largest refined region of $r/\ell_{cz} = 10$, but do not show the same amount of growth. Using the dynamic meshing algorithm results in errors of only about 1% out to a crack length of nearly $10\ell_{cz}$ even though the moving refined region size was kept at $r/\ell_{cz} = 2$. In other words, each successive mesh had approximately the same number of elements but the center of the refined region moved with the crack tip.

The real benefit of the dynamic meshing algorithm is shown in Figure 2.12b. With

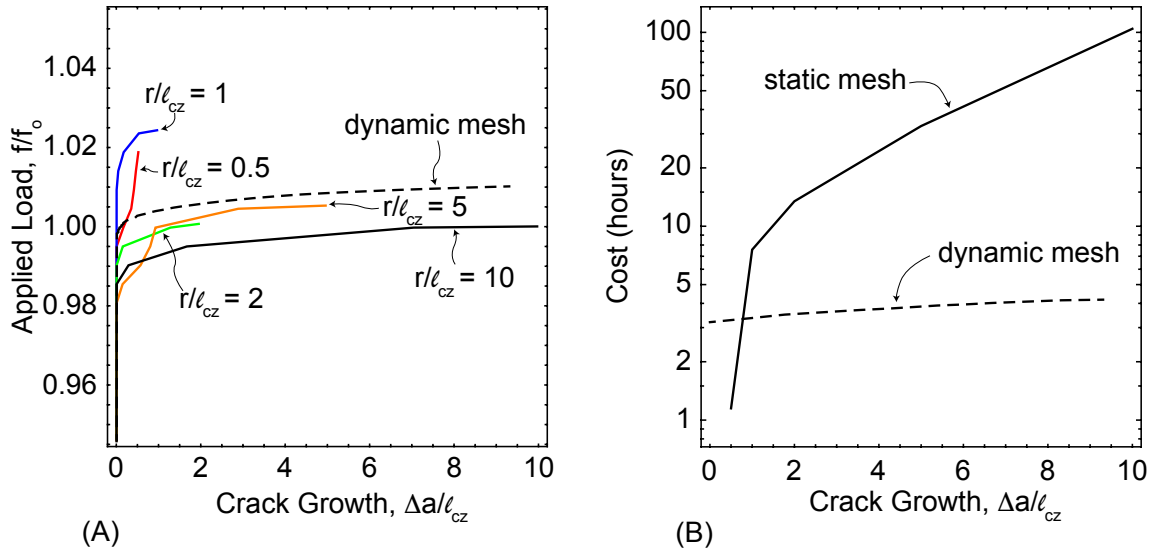


Figure 2.12: Comparisons of crack advance simulations with static refined meshes and dynamically updated meshes (adaptive refinement): (a) the predicted applied load relative to the benchmark, and (b) the cost of the simulations.

static meshes, large refined regions are needed to model significant growth. The large refined regions greatly increase the computational cost. The dynamic meshing algorithm, on the other hand, does not greatly increase in cost with increased crack growth. For problems where the crack path is unknown beforehand, the dynamic meshing algorithm is extremely beneficial. A more complex example of its usage with a single crack is shown in Chapter 3.

While the previous procedure outlines tracking and remeshing around one crack, the process can be easily extended to an arbitrary number of cracks. Figure 2.13 shows an example of the dynamic meshing algorithm with a complex geometry with multiple cracks. In the example, two identical materials are joined along a brittle interface with a saw-tooth profile. The interface toughness between the two materials is a function of position along the interface; horizontal portions of the interface have toughness Γ_{i1} and slanted portions have toughness Γ_{i2} . For this example, the toughnesses of the two materials are set equal, $\Gamma_1 = \Gamma_2$.

An initial crack comes from the left and terminates at the first slanted portion of the interface. The crack is loaded macroscopically under with farfield uniform displacements in order to apply macroscopic mode I conditions and cracks are allowed to form and grow. The interface toughnesses was set so that portions of it were always weaker than the bulk materials. In one case, both the horizontal and slanted portions of the interface were set to be equal but much weaker than the bulk materials. In the second case, only the horizontal segments of the interface were set to be much weaker than the bulk. Likely crack initiation points are known based on geometry. Both cases started with the same initial mesh and conditions.

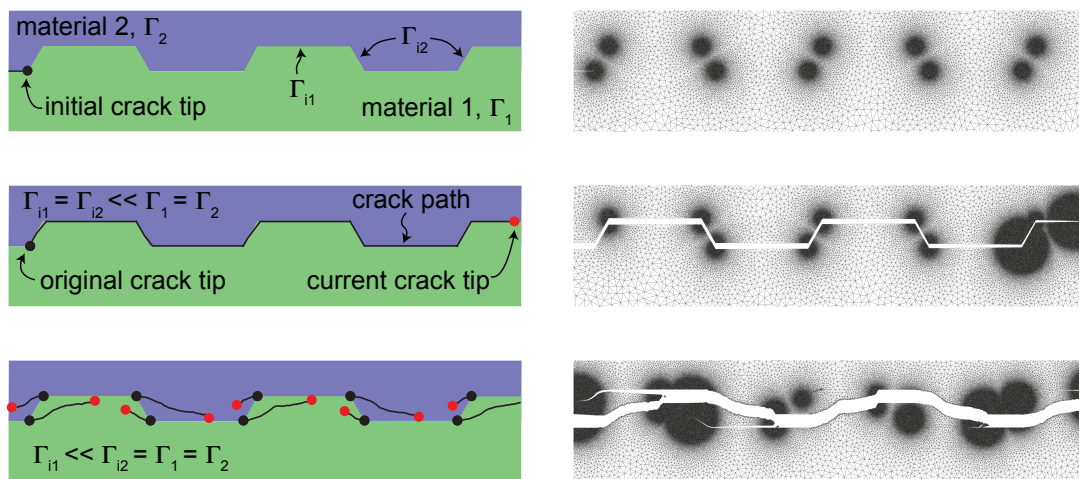


Figure 2.13: Example of the dynamic meshing algorithm tracking an arbitrary number of cracks. The top row shows the initial geometry and mesh. The middle and bottom rows show the algorithm tracking one or multiple cracks depending of the input properties along with the final mesh for each simulation.

After significant time, significant cracking had occurred in both cases. The final crack paths are shown in the left column. When the entire interface is much weaker than the bulk, a single crack grows along it. The mesh for that case shows two larger refined regions near the crack tip and smaller regions behind the tip at corners of the interface. The more interesting case is when only the horizontal portions of the interface are weaker than the bulk. The flat portions of the interface fail ahead of the parent crack. At either

end of the flat segment, cracks grow and can be seen curving to the nearest (failed) flat interface segment. In some cases, the crack impinges on the failed interface which stops further growth.

The different size of the various refined zones in either final mesh are an example of one feature of the algorithm. To further reduce computational cost, the refined region at any crack tip shrinks if the likelihood for further growth diminishes. The process of the refined zones changing in size in order to keep the cost low is shown in Figure 2.14. Initially, all refined zones are the same minimum size. As the simulation progresses, certain refined regions grow as needed. If a particular crack tip is still closed, the refined region does not expand. Once the refined zone has grown to its maximum size, it is stable until the crack closes. If the crack closes, the refined zone shrinks in size in order to reduce the number of elements.

This complex crack evolution would be incredibly expensive to solve with a static mesh. A uniform mesh would be need to be used all along the interface greatly increasing the number of elements and thus the cost. The dynamic meshing algorithms allows for the study of much more complex behaviors that would not be feasible with static meshing techniques.

2.7 Conclusions

Simulations of fracture using dynamic cohesive zone frameworks can be in inherently expensive, due to the need to resolve the fracture process zone and limit time steps to ensure stability. However, there are several computational approaches that are highly effective in reducing costs, commonly leading to decreases between one to two orders of magnitude. These can be summarized as:

- A bilinear cohesive law with decreasing tractions immediately after their peak is




Figure 2.14: Progression of various meshes used showing refined regions growing and shrinking as needed.

more desirable than the trapezoidal law with a large plateau region analogous to perfect-plasticity (used in [9]). The description drives localization of damage that encourages precise selection of advance direction. This advantage plays a critical role in the analysis of crack kinking along wavy interfaces, where distributed damage with the trapezoidal law leads to variability in kinking phenomena.

- The cost of a distributed cohesive zone simulation scales with the linear mesh

resolution raised (i.e. number of elements in a given direction) to the third power, resulting in expensive simulations of long range crack growth. This is a consequence of using a dynamic framework, which has a maximum time step for stability that scales with element size.

- There is an intrinsic link between the physical size of the domain and the cost of the simulations, due to the mesh-dependent time step and the nature of wave propagation. Small time steps demanded by stability imply high costs for physically realistic quasi-static loading rates; this can be offset by artificially increasing the load rates. The upper limit on loading rate is dictated by the elastic wave speed in the material and the domain size, implying that higher loading rates are allowable for smaller domains.
- Regions of local mesh refinement can dramatically reduce cost by obviating the need for a dense mesh everywhere in the simulation domain. Provided the fracture length-scale is resolved with ~ 30 elements, one obtains results that are equivalent to fine meshes over much larger domains.
- The advantages of local mesh refinement can be retained throughout a simulation of crack advance across large length scales, provided the local mesh region is translated with the propagating crack. An algorithm has been provided that accomplishes this with adaptive remeshing; the refined region is propagated based on the extent of cracked interfaces. A new mesh is generated with the new position of the refined region, and the results mapped from the old mesh to the new mesh. This produced a factor 20 speed-up for a test case in which the crack advanced a distance ten times the fracture process zone size.

Chapter 3

Kinking vs. Delamination Along Non-planar Interfaces

3.1 Introduction

The reliability of many multilayered systems is governed by the stability of interfaces; two notable examples are ceramic thermal barrier coatings used to protect metallic engine components [36–38], and environmental barrier coatings [39–41] used to protect ceramic matrix composites. The vast majority of previous studies to date have considered flat (planar) interfaces subject to mixed mode loading, which can lead to either delamination (along the interface) or crack kinking (where the interface crack turns to penetrate the adjacent material) [42–57]. However, many systems, such as those produced with thermal spraying, produce interfaces with non-planar features (see for example [16, 37, 58–64]). Indeed, textile ceramic matrix composites can inherently include wavy, non-planar interfaces arising from the underlying weave [65, 66]. Interestingly, many biological structures also involve wavy interfaces, such as skulls that incorporate compliant seams between adjacent plates of bone [67–71]. These biological systems have inspired the study of

textured interfaces [67].

In many of these systems, the size scale associated with non-planar features is much smaller than other geometric length-scales, i.e. the crack length at failure and the controlling component length scale (such as coating thickness, [6]). The multiscale nature of such problems is shown schematically in Figure 3.1; at the length-scale of the component, the interface appears flat, while at smaller scales the asperities influence behaviors near the tip of an interface crack. Despite their small size, non-planar features alter the local character of crack tip fields; even in instances where the macroscopic ‘parent’ crack can be considered to be planar and subject to pure mode I loading, the problem is inherently mixed mode due to local crack geometry. This is evident from the conventional finite element results shown in Figure 3.2, which depicts the energy release rate and phase angle at the tip of a wavy crack as a function of the crack tip position for several different wave amplitudes (relative to the wavelength). For interfaces with strong mode II toughness, the mode mix arising from wavy interfaces can lead to increases in critical loads needed to drive failure. (See for example [6] for a review.)

Even for ideally brittle interfaces with mode-independent toughness, however, there is a purely geometric effect, as illustrated by Figure 3.2. Under pure mode I loading in the far field, the driving force for interface delamination is reduced when the local crack path turns away from the direction of the parent crack. This leads to an apparent increase in macroscopic toughness, due to the fact that larger far field mode I amplitudes are needed to advance the crack along the interface. That is, the critical far-field energy release, G_c is larger than the interface toughness, Γ_i . This behavior is also apparent in previous studies of crack kinking involving a putative kink crack [42, 45]. Under pure mode I loading of the parent crack, the driving force at the tip of the putative kink crack is smaller than that of the parent crack for all kink angles.

For wavy interfaces, the lower driving forces for certain crack tip positions (where




Figure 3.1: (a) Schematic illustration of the parent crack, subjected to pure mode I loading, (b) Schematic of the tested geometry with applied mode I displacements on the outer boundary; the radius is assumed to be much smaller than the layer thickness, (c) Close-up view of the region of wavy interface crack; the amplitude and wavelength are presumed to be much smaller than the parent crack length, (d) Close-up view of the fine mesh region to capture the transition between interface cracking and kinking illustrating the transition between interface delamination and kinking; the fracture length scale is assumed to be much smaller than the wavelength.

the crack turns away from the far-field mode I direction, analogous to a kink) may lead to crack ‘trapping’, since higher loads may be required to drive the crack past the local

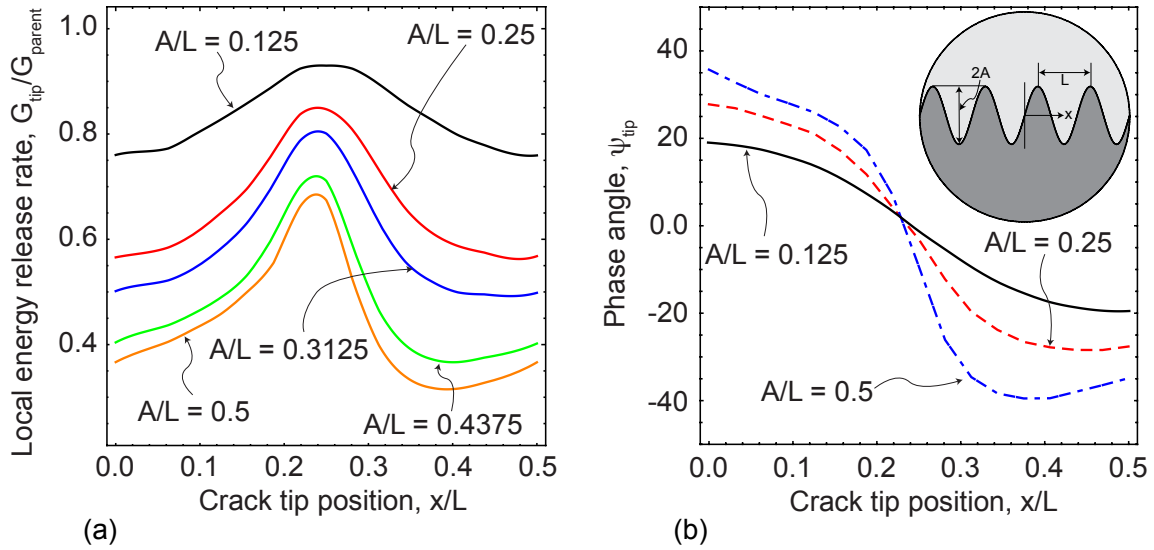


Figure 3.2: Results for purely elastic calculations (no cohesive zones): (a) energy release rate at the crack tip normalized by that of the parent crack as a function of crack tip position, (b) phase angle (mode-mix) of the interface crack tip as a function of crack tip position.

minima shown in Figure 3.2a. [72, 73]. Naturally, crack trapping requires that the crack remains in the interface and not kink into the adjacent bulk material. That is, the bulk toughness Γ_b must be large enough to suppress nucleation of kink cracks off the interface, presumably in the direction close to that of the macroscopic pure mode I field. In this work, we address questions that specifically address this problem: what is the ratio of bulk-to-interface toughness, Γ_b/Γ_i that suppresses kinking, and how does it depend on interface amplitude? What is the maximum possible apparent toughness for a given interface amplitude, given toughness ratio Γ_b/Γ_i ?

The answers to these questions are generated using a *distributed* cohesive element approach, in which cohesive elements are embedded throughout the entire mesh and allow for arbitrary crack paths. The work builds upon previous studies that examine crack stability for wavy interfaces that did not allow for kinking into the bulk [59–61, 72, 74–76]. While qualitatively (and at times, quantitatively) similar, the present work captures the effect of interacting fracture process zones (as opposed to just one along the interface) and

specifically addresses conditions controlling kinking into the bulk material. The present simulations also have much in common with a similar study of planar interface cracks [9], which provides a comprehensive discussion of the numerical details, as well as detailed description of the role of cohesive parameters that control the transition between kinking and delamination.

Naturally, mixed-mode loading in the far-field (i.e. at the component length-scale) often has a significant impact, particularly for large amplitudes that promote contact between asperities in the wake of the crack. In such instances, frictional contact between asperities behind the crack tip provides crack tip shielding that raises the effective macroscopic toughness, by lowering the local driving force at the crack tip [77]. While such behavior is obviously important, we limit the present study to scenarios which do not involve asperity contact, allowing us to isolate the role of interface geometry in the competition between kinking and delamination. The treatment of asperity contact is sensitive to the friction description and the shape of the asperities, and hence represents a considerable expansion in scope. Simplifying the crack behind the crack tip by removing the asperities might alter the results as shown in [72]. In light of this focus, Figure 3.3 provides a map indicating the applicability of the present study; the figure illustrates the range of loading phase angles, ψ , and interface shape that avoid asperity contact in the wake of the crack. Outside of this range, more detailed treatments invoking friction must be invoked, and the results will be far more sensitive to asperity geometry. To determine the applicable range, traditional finite element simulations were run every 5° from a phase angle of -90° to 90° and with an increment in A/L of 0.05. At each phase angle and amplitude combination, multiple initial crack tip locations were tested ranging from $x/L = -0.5$ to $x/L = 0.5$ with an increment of 0.03125. As expected, results were symmetric about a phase angle of $\psi = 0$.

While complete details of the simulation framework are given in the following section,

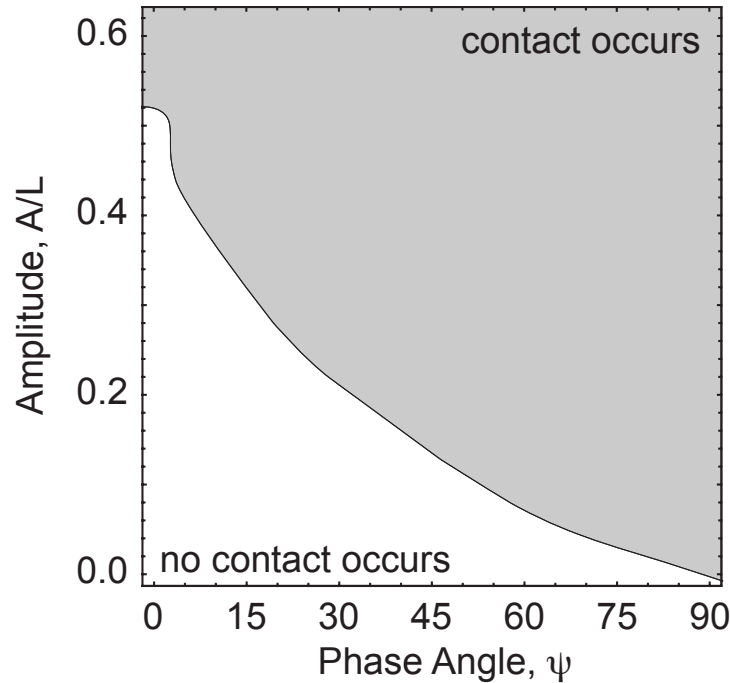


Figure 3.3: Map of the amplitudes where crack face contact occurs for various loading phase angles. The boundary curve was determined by interpolating between amplitude and phase angle combinations that resulted in contact and those that did not.

it is worth emphasizing the broader context of the underlying fracture theory. Cracks emerge as separation between elements, with a traction-displacement relationship (cohesive law) that defines the energetics of the fracture process zone (i.e. the near tip region that experiences inelastic separations). For systems with isotropic toughness and in the limit that the fracture process zone is much smaller than all other length-scales, the shape of the cohesive law is immaterial, with behavior governed by the area under of the traction-separation curve, which represents the intrinsic toughness of the material (or interface). The present approach utilizes this ‘small-scale yielding’ limit, such that the physical mechanisms of rupture are largely immaterial. Various types of processes, including bond rupture, can be represented by the phenomenological cohesive law, whose details play a secondary role to the energy associated with complete separation (i.e. the intrinsic toughness).

For systems with heterogeneous toughness, such as those where the interface toughness is different from that of the bulk, the behavior is somewhat more nuanced. Cohesive separations in one domain (e.g. the bulk) influence separations in another (e.g. the interface). In such instances, the relative strengths of the cohesive laws in both domains can also influence fracture [9]. Motivated by completely brittle systems, we assume that the cohesive strengths (peak stresses) of the bulk and interface are comparable, though the associated toughness may be different. This difference in cohesive parameters has a much smaller influence than that of the toughness difference, provided small-scale yielding conditions are maintained (as they are); furthermore, the absolute scale of the critical separation is immaterial, as the fracture process occurs over length-scales much smaller than other features.

In this context, the present results are applicable to a variety of fracture mechanisms, provided they occur over small length scales and can be reasonably described by equal cohesive (peak) strengths. This largely limits the scope to brittle or quasi-brittle systems, as ductile interfaces (e.g. adhesive layers) involve peak strengths that are much smaller than those associated with bond rupture. For wavy interfaces, systems with a very ductile phase often violate small-scale yielding conditions because the fracture process zone will be comparable to or larger than the asperity size. In all instances in the present simulations, the fracture process zone is much smaller than the asperity amplitude, wavelength or radius of curvature; this is clearly most applicable to brittle systems.

3.2 Simulation Framework

The simulation framework consists of an explicit dynamic finite element method that embeds cohesive laws between every element in the mesh, similar to the approach pioneered by Xu and Needleman [32]. The stiffness of the cohesive law is set such that

linear elasticity is recovered at small strains and controlled by the conventional elements. Crack path evolution corresponds to a loss of traction between element faces and is a natural outcome of the computation. Previous work has demonstrated that for random triangular meshes, arbitrary crack paths (associated with the direction of maximum energy release rate) can be captured without a priori assumptions of growth direction or additional fracture criteria (beyond the material toughness) [9]. An explicit solution technique was used here (as in [32] and [9]), but an implicit solver might also be used if the cohesive laws were modified as described by Gao and Bower [78].

3.2.1 Material Description

The cohesive law is described in terms of a peak stress σ_o , and the cohesive displacement (element separation) associated with that stress, δ_o . For cohesive displacements greater than δ_o , the traction exhibits linear softening, as shown in Figure 3.4. The elastic stiffness of the cohesive law, k_e , is kept large relative to Eh_e , where E is the elastic modulus of the elements themselves and h_e is the element size; this recovers conventional elasticity for the bulk material, in that deformation of the elements themselves dominates the response. The normal and tangential cohesive laws are defined such that the work to fracture of any interface is dictated by $\Gamma = c(\delta_{max})\sigma_o\delta_o$, where $c(\delta_{max})$ is a dimensionless constant of order unity that is weakly influenced by the maximum displacement at which tractions go to zero. Put another way, in small-scale yielding, the material fractures as an ideally brittle solid, without mode-dependence.

The reader is referred to our earlier work [9] for a full discussion of the influence of cohesive parameters, and we provide only a brief summary here. Embedded cohesive




Figure 3.4: Traction-displacement behavior of the cohesive elements embedded among elements: tangential displacements affect normal traction, and vice versa, such that the work to failure of the elements is constant regardless of loading.

elements introduce a fracture length scale given by

$$\ell_{cz} = \frac{E\Gamma}{\sigma_o^2} \quad (3.1)$$

where E is the elastic modulus of the elements, Γ is the work-to-fracture of the cohesive law, and σ_o is the peak strength of the traction-displacement relationship. The cohesive length-scale describes the physical size of the fracture process zone, where separations exceed δ_o . Strictly speaking, the actual length of the fracture process zone according to this is a scalar multiple of eqn. 3.1; however, previous works indicate that this scalar is close to unity for the cohesive description utilized here. If the fracture process zone is small relative to all other dimensions in the model, ℓ_{cz} serves the sole purpose of dictating the mesh required to resolve the behavior in the fracture process zone; i.e. h_e/ℓ_{cz} (where h_e is the size of the elements) must be sufficiently small.

It is important to keep in mind that two independent cohesive laws are needed for the present study, one for the interface and one for the bulk material (taken to be identical

on either side of the interface). With two independent cohesive laws, there are two corresponding fracture length scales, ℓ_{cz}^b for the bulk and ℓ_{cz}^i for the interface. Here, simulations are presented for $\ell_{cz}^b = \ell_{cz}^i$; details regarding the impact of this assumption can be found in [9]. For all simulations, the interface strength was set at $\sigma_o^i = 2GPa$ and the modulus was fixed at $E = 200GPa$. For equal cohesive lengths, $\ell_{cz}^b = \ell_{cz}^i$, and toughness ratios of $1.25 \leq \Gamma_b/\Gamma_i \leq 4$, this implies $1.12 \leq \sigma_o^b/\sigma_o^i \leq 2$, and $1.12 \leq \delta_o^b/\delta_o^i \leq 2$.

Unlike our previous study of kinking with this framework, the present study uses a bilinear cohesive law like those used in [29, 79–81] as shown in Figure 3.4. As with the trapezoidal law used in Pro et al. [9], the normal and tangential tractions are coupled so that the work to failure is the same regardless of the balance of opening and sliding displacements. In preparing the results presented here, we have found that the softening law reduces the emergence of multiple simultaneous possible crack paths and strongly favors the formation of a dominant crack. Presumably, this is because adjacent cohesive zones do not support uniform tractions across a broad range of openings; once the cohesive law exceeds the critical opening, deformation localizes in one of the emergent cracks and a single path quickly becomes dominant. Put another way, the softening law more quickly resolves competitions between multiple possible crack paths. The net effect is that slightly coarser meshes can be used with softening laws; one can obtain identical results with an equivalent trapezoidal law [9], but finer meshes are needed to reduce the possibility of multiple crack paths.

The bulk material (represented with conventional constant strain triangular elements) is described with linear, isotropic elasticity with elastic modulus E , Poisson’s ratio ν , and density ρ , assuming plane strain conditions. The material on either side of the interface has identical properties. The color difference shown in Figure 3.1 is only used to highlight the location of the interface. To speed the computational costs of quasi-static loading conditions, artificial proportional damping is introduced as is typical in conventional finite

element methods [82, 83]. A later section describes the positive impact of this artificial damping with respect to computational speed, and the negligible impact with respect to critical cracking phenomena.

3.2.2 Geometry and Loading

Figure 3.1b and Figure 3.1c illustrate the geometry used in the present simulations; it is assumed that wavelength of the wavy interface L is much smaller than the parent crack, such that the elastic fields surrounding the initial crack tip are governed by conventional planar solutions. On either side the parent crack tip, the interface is defined by $y(x) = A \sin(2\pi x/L)$ for a distance of $3L$ both ahead of and behind the crack tip. Outside of the wavy region of the parent crack, the interface is flat; additional numerical simulations, not presented in this work, as well as previous literature [72] have confirmed that modeling additional periods does not affect the results. We present results for amplitude ratios in the range $0 < A/L < 0.5$, as shown in Figure 3.1c.

The disk surrounding the parent crack tip is loaded by applying time-dependent displacements defined by the asymptotic crack tip field and controlled by the time-dependent, pure mode I remote stress intensity factor $K(t)$. (Locally, the crack tip experiences mixed-mode fields due to interface waviness.) Ramp loading was applied defined by:

$$K(t) = \dot{K}t \quad (3.2)$$

where \dot{K} is the loading rate. On the one hand, slow loading rates are needed to capture quasi-static behaviors; on the other, computational time increases as loading rate decreases when using an explicit dynamic time stepping scheme. The inclusion of small amounts of mass-damping reduces the impact of elastic waves associated with the loading time scale and allows for larger time steps that recover quasi-static behaviors. Compu-

tational considerations and the selection of damping parameters are discussed in the following sections.

3.2.3 Computational Considerations: Time Stepping

The finite element framework reduces to a set of governing equations of the form $[M] \{\ddot{u}\} + [c] \{\dot{u}\} + [K] \{u\} = \{F(t)\}$, where $\{u\}$ is a set of nodal displacements, $[M]$ is an equivalent mass matrix, $[K]$ is the stiffness matrix, and $[c] = \alpha_D [M] + \beta_D [K]$ is the effective damping matrix (constructed from the mass and stiffness matrices). Here, the mass and stiffness damping parameters, α_D and β_D respectively, are not motivated by material behaviors but rather are a numerical artifice intended to damp elastic waves generated by ramp loading. Acceptable damping values for the present application were determined as follows. First, a simple one-dimensional bar problem was analyzed with a uniform mesh and ramp loading, to identify parameters that produce quasi-static behaviors at minimal cost. This numerical study identified values of α_D and β_D that minimized computation time *and* produced uniform stress fields within 3% of the elastic solution at the end of the loading ramp. Values of $\alpha_D = 0.0181/ns$ and $\beta_D = 0.71ns$ with a time step of $\Delta t = 3.75 \times 10^{-5} ns$ achieved this goal with a density of $4g/cm^3$. It should be understood that the time-scale is artificial in the sense that the present results focus on the quasi-static limit.

Subsequently, a characteristic cracking problem with an amplitude of $A/L = 0.25$, initial crack tip position of $x/L = 0$ and a toughness ratio of $\Gamma_b/\Gamma_i = 1.25$ was analyzed with a coarse mesh at various loading rates, \dot{K} . This particular case involves kinking from the interface. Displacements were applied at the disk boundary with loading rates ranging from $\dot{K} = 1 \times 10^{-4} K_o/ns$ to $\dot{K} = 1 \times 10^{-2} K_o/ns$ where $K_o = \sqrt{\bar{E}\Gamma_i}$ is the plane strain toughness of the interface and $\bar{E} = E/(1 - \nu^2)$. The far-field applied energy release

rate required for the kink crack to nucleate, G_c , is shown in Figure 3.5a as a function of loading rate. As the loading rate increases, the critical value of the energy release rate increases due to the presence of mass-damping, which dissipates energy that would otherwise be available to drive crack growth. Since computational time increases with decreasing time step, loading rate of $\dot{K} = 10^{-3}K_o/ns$ was chosen as it is within 5% of the quasi-static limit while providing significantly faster computations. For example, this choice of loading rate is twice as fast as the next slowest loading rate shown in Figure 3.5a.

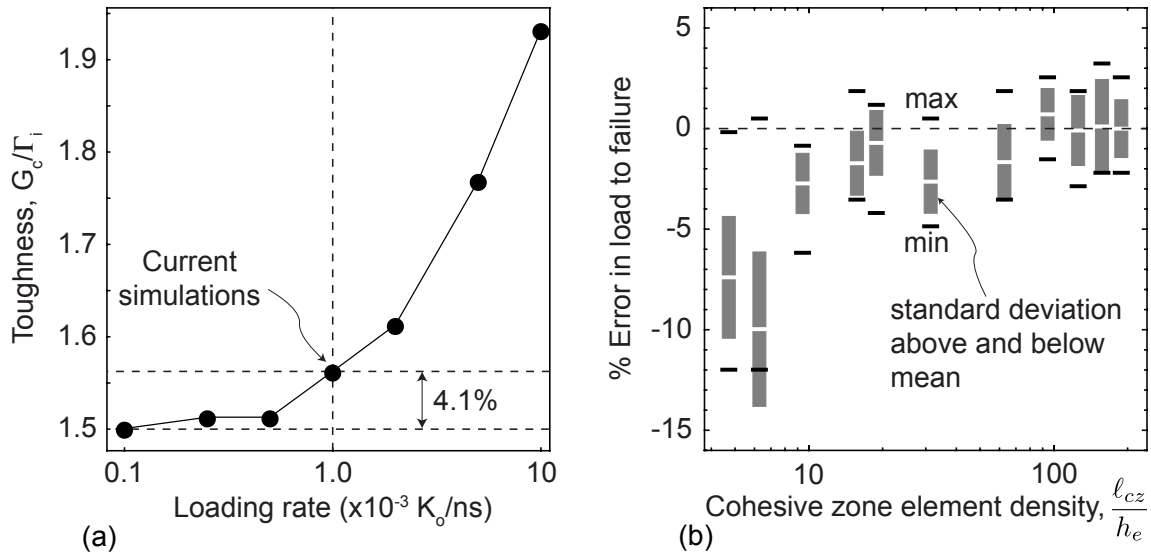


Figure 3.5: (a) Effect of loading rate on predicted fracture toughness; these results are used to identify loading rate that will reduce computation time while still recovering quasi-static results. (b) Effect of mesh density on predicted toughness; these results are used to set the size of the elements in the refined zone.

With these properties and numerical parameters, the simulation cost for a single cracking problem with a resolved mesh (described in the next section) was ~ 60 hours running with 8 threads on a distributed cluster of servers each with two quad core Intel Nehalem processors. The cluster allowed up to 24 concurrent simulations. The approximately three hundred simulations in this study took a total of about one month of wall clock time.

3.2.4 Computational Considerations: Mesh Resolution

Before discussing the mesh resolution needed for the problem, a note of the length scales inherent to both the problem and the method is needed. The fracture length scale ℓ_{cz} must be smaller than any geometric length scale for the small scale yielding approximation to be valid and the simulations to correspond to brittle fracture. In contrast kinking studies of flat interfaces [9], the inclusion of the wavy interface introduces two additional length scales: the wavelength, L , and the radius of curvature of the interface, ρ_c . The present simulations fixed the wavelength to $0.1R$ and vary the amplitude of the interface. For a sinusoidal interface, the minimum of the radius of curvature is given by:

$$\rho_c^{min} = \frac{L^2}{4\pi^2 A} \quad (3.3)$$

Since this is the smallest length scale present in the geometry, the fracture length scale was set such that $\ell_{cz}^i = 0.28\rho_c^{min}$. Smaller cohesive zones had significantly larger computational times and a negligible affect on the results.

Finally, a mesh resolution study was performed with the same geometry as the loading rate study, but with multiple meshes with different number of elements in the fracture length scale. Also, a number of different meshes at the same resolution were created by slightly perturbing the input points to the mesh generator; 10 meshes were studied at each resolution. In each case, the applied energy release rate for the crack to grow a distance of $\sim 0.15\ell_{cz}$ was then computed; the results of this study are shown in Figure 3.5b, where the white lines show the mean of all meshes, the black lines show the minimum and maximum selected from all the meshes, and the gray box shows one standard deviation arising from all meshes. One observes that coarse meshes (except for the two coarsest mesh densities) underestimate the mean from fine meshes by about 4%, while the variability due to the use of values from a single mesh is about 5%. In the present study, we

chose $\ell_{cz}/h_e \sim 94$ to balance accuracy and speed. Overall, when one considers all the results in Figure 3.5, it is reasonable to expect that numerical discretization introduces an uncertainty (variability) of about 5% with regards to the critical far-field energy release rate that drives cracking.

3.3 Results

Figure 3.6 illustrates the far field energy release rate that must be applied to drive cracking, as a function of crack tip position and for several toughness ratios Γ_b/Γ_i . Open symbols indicate that crack advance occurs via kinking, while filled symbols indicate that the crack's initial advance is along the interface. Note that for a given property set, cracks will advance along the interface for some crack tip positions, while others will lead to kinking. Figure 3.6a illustrates that for close to equal toughness and low amplitudes, the crack remains trapped on the interface with only a small difference between the critical far-field energy release rate and the intrinsic toughness. As the amplitude increases, kinking becomes more prevalent, and occurs for all crack tip positions: the results near $x/L = 0.25$ correspond to a crest, such that interface advance is indistinguishable from kinking.

Kinking is suppressed (i.e. the crack is ‘trapped’) if the toughness of the bulk is much higher than the interface, as illustrated in Figures 3.6b and 3.6c. Further, the far field energy release rate needed to drive interface cracking is much higher than the intrinsic toughness. This is a result of the decrease in local driving force when the crack turns away from the parent crack direction, as seen in Figure 3.2a. Indeed, the inverse of the results in Figure 3.2a (which neglect any kinking damage) is also plotted in Figure 3.6c for an amplitude of $A/L = 0.5$, and is in excellent agreement. This is a useful check on the cohesive zone results; when kinking is suppressed, one recovers identical results to

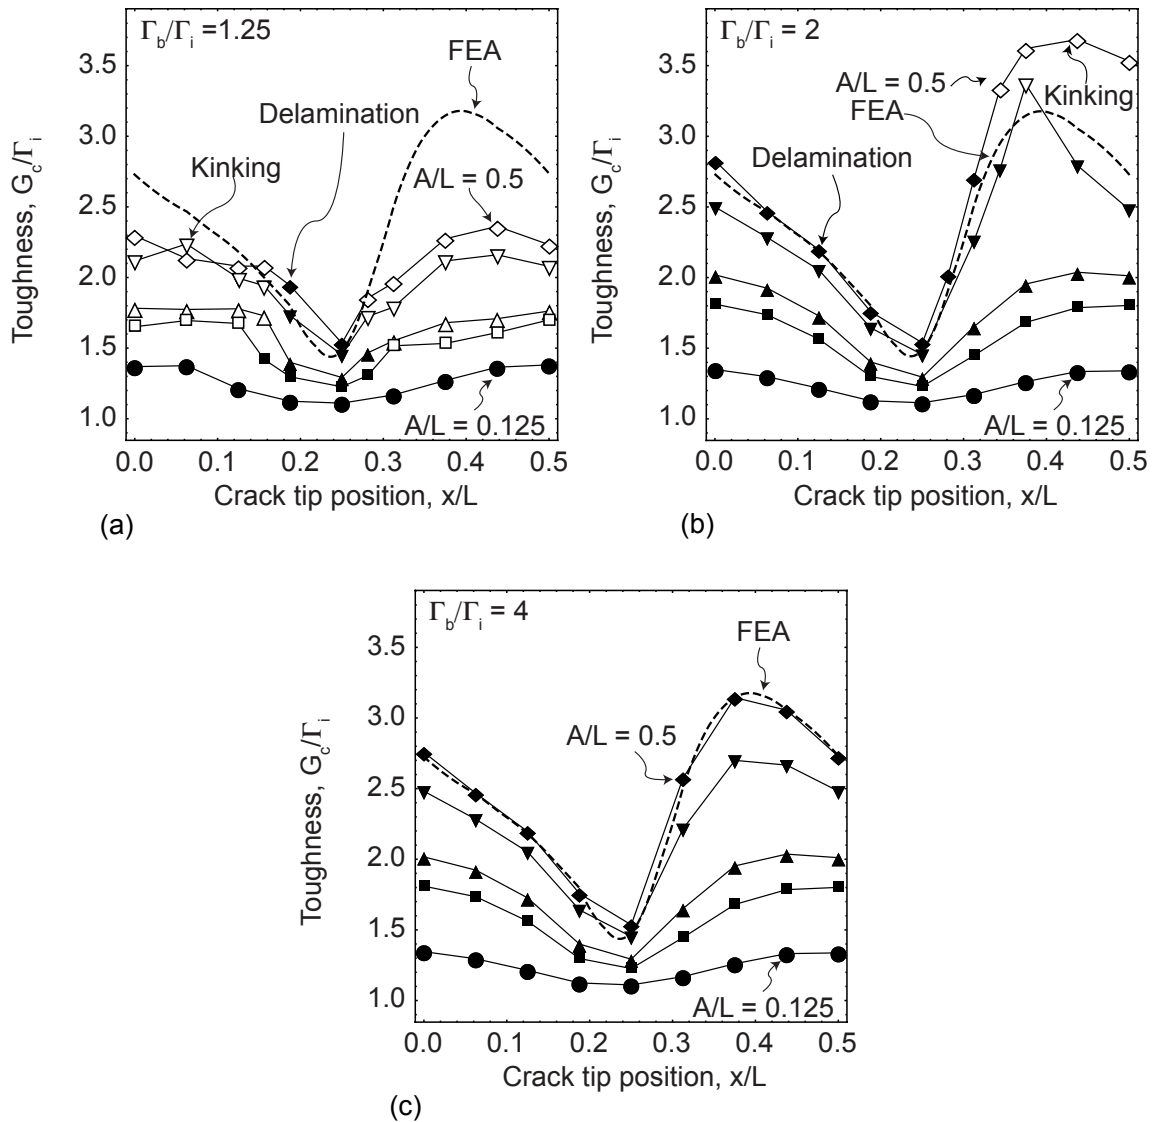


Figure 3.6: Far field energy release rate associated with crack advance, normalized by the toughness of the interface, for three different toughness ratios Γ_b/Γ_i . Open symbols indicate kinking out of the interface; solid symbols indicate interface cracking.

conventional finite elements. It is interesting to note that if this estimate is used, one completely misses the failure mode at lower toughness ratios, as can be illustrated by superposing the conventional ERR result onto Figure 3.6a.

These results shed important new insight as to whether or not crack ‘trapping’ along the interface can be effectively achieved; obviously, if the crack can kink into the bulk,

the crack is no longer trapped. For example, consider the results in Figure 3.6b for $A/L = 1/2$ (a strongly wavy interface crack) and suppose the crack tip starts at $x/L = 0$, the mid-point between a trough and a crest. If the far-field energy release rate G is roughly $G = 2.8\Gamma_i$, the crack advances along the interface. Since the critical value of far-field G needed to advance the crack is less than this value for all crack positions $x/L < \sim 0.32$ - *and kinking does not occur* - the crack advances to $x/L \approx 0.32$ and arrests, since the required driving force to advance the crack at that position is $G > 2.8\Gamma_i$. Provided $G < 3.5\Gamma_i$, the crack neither advances along the interface nor kinks. The crack is effectively arrested or ‘pinned’ at that position as well as ‘trapped’ along the interface. However, if $G > 3.5\Gamma_i$, a kink crack nucleates and is no longer ‘pinned’. Hence, the crack is only effectively ‘trapped’ or ‘pinned’ for $G < 3.5\Gamma_i$. Note that the kink crack nucleates at $G = 1.75\Gamma_b$, a consequence of the difference in the local orientation of the crack tip and the parent crack, which influences mode-mix at the local crack tip and controls kink nucleation [9].

The results in Figures 3.6a-c illustrate that crack trapping is most effective for large amplitudes and high toughness ratios, as one expects; when the amplitude is small or the bulk toughness is comparable to the interface, there is a narrow range of far field energy release rates that are relevant. Consider the results in Figure 3.6a, for the case where the pre-existing crack somewhat luckily starts at a crest, $x/L = 0.25$. The crack can be advanced and arrested only if the far-field energy release rate G is above the critical condition and is never increased by more than $\sim 20\%$; otherwise, kinking occurs. As the bulk toughness is increased, the window of ‘quasi-stable’ far-field G (where the crack will advance but subsequently arrest) also increases, as seen in Figure 3.6b.

An interesting feature of wavy cracks with large amplitudes (relative to the wavelength) is that the driving force is asymmetric with respect to crack tip position near the crest of the wave. That is, the driving force is higher as the crack tip approaches a crest,

as compared to the crack tip leaving the crest of the wave. This asymmetry occurs even when the local crack tip angle (tangent to the interface at the crack tip) has the same magnitude relative to the parent crack. This asymmetry is observed in the conventional FEA results (Figure 3.2), the present cohesive zone predictions (Figure 3.6), and similar calculations (which do not allow for kinking) published previously (e.g. [72], Figure 10). This may be somewhat counterintuitive, since the local crack tips have positions with equal distance to the macroscopic crack plane, and the orientation of the crack tips and the curvatures of the cracks are equal. It should be noted that a different definition of x/L was used in this work compared to [72].

This behavior may be understood as follows. The behavior behind the crack tip strongly impacts the influence of the crack tip position; absent partial interface delamination, the shape of the interface ahead of the crack plays no role since the material is isotropic. For cracks approaching the crest of a wave, the full asperity in the wake is almost fully unloaded and its influence diminishes as the crack moves into the new asperity. That is, the crack approaching the tip of the wave (on the uphill slope) is entering a nearly fully stressed asperity, with the previous (largely unloaded) asperity far removed. On the other hand, for a crack tip that passes the crest of the wave (on the downhill slope), the stress in the asperity is reduced by the presence of the crack, which somewhat reduces the strain energy available to drive the crack. Simply put, the stress field in the ‘active’ asperity containing the crack dominates the response; once the crack passes the crest the small amount of strain energy in the ‘active’ asperity that feeds the crack is reduced by crack advance. This hypothesis is supported by the fact that the effect disappears for small amplitudes; in this case, a large portion of *any* asperity in the wake of the crack is stress free and plays little role.

Wavy cracks increase the critical far-field energy release rate (of the parent crack) needed to drive either kinking or delamination; as evident from Figure 3.6, this increase

is optimal for intermediate toughness ratios (e.g. $\Gamma_b/\Gamma_i \sim 2$). This is because the highest critical energy release rate is obtained when there is a competition between kinking and delamination. This is clearly illustrated in Figure 3.7, which depicts the critical energy release rate for the parent crack as a function of toughness ratio. Two different crack locations are shown, roughly corresponding to the locations with the steepest interface and peak toughening from Figure 3.6 ($x/L = 0$ and $(x/L = 0.375)$, respectively). One can see that the critical energy release rate for kinking increases with toughness until a critical value is reached; at this critical value, kinking damage is completely suppressed, and the failure mode switches to delamination. After this point, the increase in apparent toughness is completely controlled by the interface toughness and the amplitude ratio. Figure 3.7 illustrates that nearly optimal toughening is obtained for $\Gamma_b/\Gamma_i \sim 2$, with an increase of about a factor of 3-4 in apparent toughness. The optimal toughening obtained for $\Gamma_b/\Gamma_i \sim 2$ has only a slight dependence on the two crack tip positions. The value of this transition is further explored below. The optimal toughening is a consequence of active competition between fracture process zones at the interface and in the bulk; note that the effective toughness is roughly the sum of both fracture toughness values. It should be noted that the increase in apparent toughness that results from competing fracture mechanisms can depend on the ratio of the cohesive strengths of the bulk and interface; this is discussed in extensive detail in [9] (which covers kinking from a planar interface crack) and the references included in that work.

Since the optimal toughening occurs near the transition between kinking and delamination, predicting the transition toughness would be useful for different geometries (where there is no crack face contact). The results of [42] and [9] can be used along with the phase angles shown in Figure 3.2b to see if either presents a valid prediction of the transition toughness ratio. The results of this analysis can be seen in Figure 3.8 for $A/L = 0.25$ (a) and $A/L = 0.5$ (b). In this figure, the transition toughness ratio was determined to

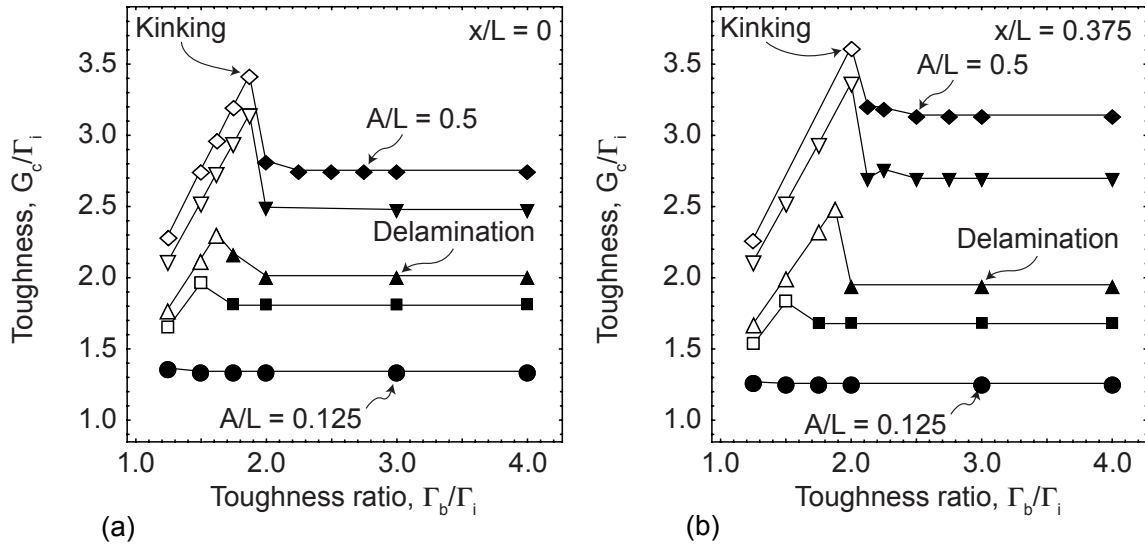


Figure 3.7: Apparent toughness needed to drive cracking as a function of the bulk-to-interface toughness ratio Γ_b/Γ_i for an initial crack tip position of (a) $x/L = 0$ and (b) $x/L = 0.375$.

be the average of the highest ratio resulting in kinking and the lowest ratio resulting in delamination. This means there is an error in the transition toughness ratio of ± 0.125 . Kinking is expected at a certain crack tip position if the toughness ratio is below the curve and delamination is expected if it above. The He and Hutchinson model greatly under-predicts the transition, which is consistent with the work of [9] which showed the He and Hutchinson's model does not accurately predict the transition toughness ratio for flat cracks under mixed mode loading conditions (Figure 8 in [9]). Using the results from Pro et al., the transition toughness ratio can be fairly accurately predicted only with the local mode mixity. Differences between both models and the simulation results around $x/L = 0.25$ can be explained by the curvature of the interface at that point. The difference in the instantaneous angle of the interface and the preferred kink angle is very small near the peak of a wave. Also, at $x/L < 0.25$, the interface curvature is a close approximation of the ideal kink angle, as shown in Figure 3.8c. Delamination is then more preferred than the predictions from Pro et al. suggest.

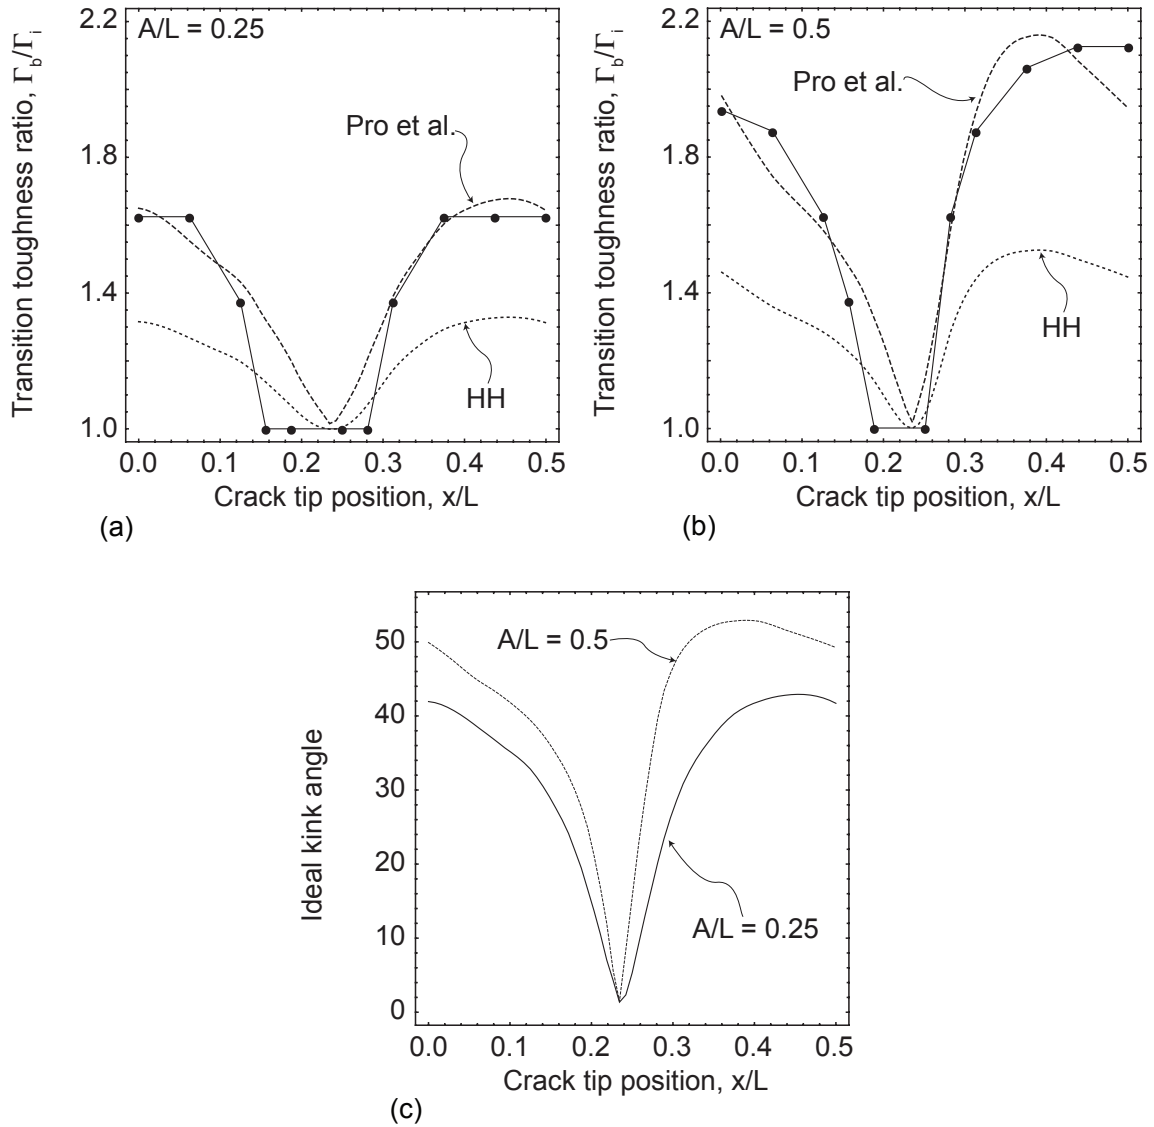


Figure 3.8: Plots of the predicted transition toughness ratio from Pro et al. [9], He and Hutchinson [42] and simulations for (a) $A/L = 0.25$ and (b) $A/L = 0.5$ as well as (c) the calculated kink angle (from [42]) as a function of crack tip position.

The increase in toughness due to the competition between kinking and delamination can be used to increase the critical far-field energy release rate that drives unstable cracking. It is assumed that crack is always unstable after kink nucleation; delamination is unstable if the crack does not arrest *and* does not kink prior to arrest. As shown in Figure 3.9, the increase in critical far-field energy release rate is dependent on both

the amplitude and toughness ratio. Even at a low toughness ratio of $\Gamma_b/\Gamma_i = 1.25$, the critical energy release rate increases with increased amplitude even when kinking happens at most amplitudes. Increasing the toughness ratio from 1.25 to 2 (an increase of 60%) increases the critical energy release rate by almost 100%. Further increasing the toughness ratio so that the crack does not kink decreases the critical far-field energy release rate, G_c , but shows nearly identical results to [72] (which did not consider kinking). This provides additional evidence of the accuracy of the present calculations. It is clear that competitions between kinking and delamination can be a useful mechanism to toughen a rough interface, but only if the bulk is sufficiently tougher than the interface (but not too much tougher).

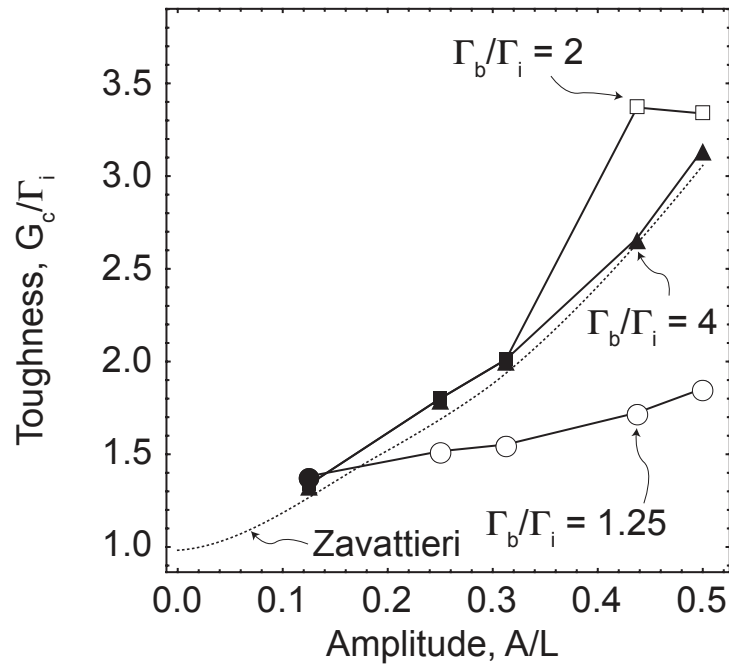


Figure 3.9: Plot of the maximum load that can be applied before the crack grows unstably for three different toughness ratios Γ_b/Γ_i compared the data from [72]. Open symbols indicate kinking out of the interface; solid symbols indicate interface cracking.

Arguably, the central question regarding the development of robust coating systems is, “for a given toughness ratio and interface amplitude, what is the maximum energy

release rate that avoids cracking entirely (for *any* crack tip position)?” The answer to this question is provided in Figure 3.10, which plots the critical energy release rate to drive unstable cracking as function of toughness ratio Γ_b/Γ_i and amplitude ratio A/L . In Figure 3.10, the size of each bubble is scaled to the critical toughness and closed bubbles indicate delamination while open bubbles indicate kinking. This is a more complete view of the results than that shown in Figure 3.9, showing the critical failure mode across the tested parameter space. One can see immediately that the optimal apparent toughness occurs near the kinking/delamination boundary. Increases in the amplitude ratio increase the apparent toughness, with only weak dependence on the toughness ratio. One can observe that for large amplitudes, staying on the kinking side of the transitions produces the biggest gains, as suggested by Figure 3.7. The difference in toughness is lower than that suggested by Figure 3.7 because it is a comparison of the highest delamination critical energy release rate compared to the lowest kinking energy release rate for a given amplitude. This slight difference can be seen in Figures 3.6a and 3.6b.

3.4 Conclusions for Wavy Interfaces

The work described provides new insights regarding the competition between kinking and delamination for a wavy interface crack:

- The embedded cohesive zone approach is an effective tool to consider competitions between brittle crack mechanisms, provided sufficient mesh resolution is utilized and one takes care to utilize parameters for which the fracture process zone is much smaller than all other length scales.
- When delamination is heavily favored (e.g. for high toughness ratios Γ_b/Γ_i), the represent results are in broad agreement with previous studies of wavy interfaces




Figure 3.10: Plot of the maximum applied toughness for a given toughness ratio and amplitude. The size of each bubble is scaled to the ratio of the critical toughness G_c to the interface toughness Γ_i . Closed symbols indicate interface failure while open symbols represent kinking.

that ignore the possibility of kink nucleation. In this limit, increasing waviness increases the critical far-field energy release rate of the parent crack for all crack tip positions.

- For moderate toughness ratios, e.g. $\Gamma_b/\Gamma_i \leq 2$, kink nucleation can occur; generally speaking, this limits the apparent toughening to values approximately 70% less than would be obtained if kinking were suppressed.
- Maximal toughening is obtained for specific toughness ratios that produce a competition between kinking and delamination; the benefit of this competition is roughly 10-70%, corresponding to small amplitudes to large amplitudes respectively. The toughness ratio that maximizes apparent toughening increases slightly with amplitude but generally falls in the range $1.5 \leq \Gamma_b/\Gamma_i \leq 2$ for $0.25 \leq A/L \leq 0.5$.

3.5 Cracking in Wavy Lamellae with Adaptive Remeshing

Many biological systems, such as dinosaur teeth [10], consist of lamellae that are non-planar, as shown schematically in Figure 3.11. Based on Brandon Krick’s observations at Lehigh, crack patterns exhibit interesting deflections (kinking) behaviors based on the local microstructure. For example, Krick indented dinosaur teeth to measure properties and induce indent cracks that would then propagate; in many instances, the cracks would propagate through many lamellae before making abrupt turns. The results from Section 3.4 illustrate that kinking behaviors near wavy interfaces are strongly influenced by the location of the crack tip along the wave: e.g. near the crest, trough or mid-points.

These observations imply conventional “static mesh” finite element analysis of crack stability will be exceedingly expensive. To span the range of possible crack/microstructure interactions, crack orientation would have to varied alongside variations in crack tip locations throughout the domain (in two directions) to span the range of possible/crack microstructure interactions. In contrast, the adaptive remeshing algorithm described in Chapter 2 is ideally suited to simulating these interactions: an initial crack can be placed in a few locations and allowed to evolve as desired, including the propagation across lamellae followed by kinking when the crack hits a local feature with the “right” orientation.

To gain some initial insight into Krick’s observations and illustrate the power of the adaptive remeshing algorithm, a limited study was conducted using lamellae elastic properties obtained from Krick’s indentation experiments, and geometry identified by Krick’s optical micrographs. Specifically, these studies suggest that the modulus difference between the layers is relatively small: $E_2/E_1 \sim 1.25$. The hardness difference between the layers is only slightly larger: $H_2/H_1 = 1.5$. The amplitude (A) to wavelength (L)




Figure 3.11: Schematic illustration of the geometry used for the dinosaur teeth example. At the macroscopic level, a cracked rectangular domain is loaded under displacement control. The width W and height H are significantly larger than the initial crack length a , approximating an isolated crack in a semi-infinite domain. On a smaller length scale there are alternating lamella with wavy interfaces with the primary axis of the wavy interface rotated an angle θ counterclockwise from the plane of the crack. The lamella have thickness h , amplitude A and wavelength L . For reference, the point $x=0$ refers to the midpoint between a trough and a peak of an interface.

ratio ranged between $0.05 < A/L < 0.3$, while the observed thickness (h) range was $0.07 < h/L < 0.175$.

Motivated by these observations, a parameter study was conducted with $A/L = 0.175$ and $h/L = 0.255$; the crack angle θ shown in Figure 3.11 was varied from 0 to 90° and various crack tip positions relative to the lamellae were considered. To simplify the study, the cohesive strengths of the layers was set to be $\sigma_o^2/\sigma_o^1 = 1.5$; the toughness of each layer was adjusted so that the corresponding fracture process zone sizes were the same throughout the material. This implies $\Gamma_2/\Gamma_1 = 1.8$, and $\Gamma_i/\Gamma_1 = 1.4$. The upshot of this is that the stiffer layer is 20% stiffer and 80% tougher than the more compliant, weaker layer; further the remote stress needed to advance an isolated crack in a monolithic specimen of material #2 is 50% higher than if the specimen were made from material #1.

The outer geometry consisted of an edge crack panel with height and widths of at least 50 lamellae wavelengths, while the initial crack length spanned at least 15 lamellae for orientations greater than 45° . The displacement at the top edge of the panel was increased linearly with time (constant velocity) loading, at rates well below the quasi-static threshold identified in Chapter 2. This last point is important, since it implies all velocities along the boundaries are at least two orders of magnitude smaller than the wave speed of the material; further, one should note that purely quasi-static response would indicate that crack growth is stable under far-field displacement control.

The full set of simulations exhibited surprisingly consistent behavior that was independent of the initial crack orientation and initial crack tip position. In essence, the cracks always ran straight ahead with only minor influences of the domain waviness, regardless of the initial angle of the lamellae to the initial crack orientation. The load to initiation was always virtually identical to the material surrounding its tip. While propagating, the stress needed to drive failure as a function of crack length (a proxy for

the macroscopic R-curve behavior) jumps when moving from the brittle to the tougher phase. In essence, no crack deflection along interfaces or kinking was experienced.

This universality is undoubtedly a consequence of dominant far-field mode I loading, the limited difference in elastic properties between the layers, and the limited difference in fracture toughness both between the layers and along the interface. However, there are some interesting behaviors of crack propagation, unrelated to propagation direction, that are also common to all simulations, and these are the focus of the remainder.

Figure 3.12a shows the level of applied stress (inferred from boundary displacements while assuming uniform strain ahead of the crack) as a function of crack tip position. This is, in essence, the R-curve behavior for the microstructure that is illustrated in Figure 3.12b. The corresponding microstructure that is encountered as the crack advances can be inferred from Figure 3.12b. Clearly, the applied load that drives initial growth jumps from that associated with cracking in a uniform brittle material (#1), to something that is not quite that associated with cracking in a uniform tough material (#2). As the crack moves from layer to layer, the transition is not abrupt, indicating alterations to the crack tip fields by the advance layers. The results in Figure 3.12a raise the question: why isn't the resistance of the tougher phase ever experienced fully? It is interesting to note that it is never fully reached, regardless of whether crack runs through a long portion of the tough phase (for certain combinations of crack orientation or initial crack tip position).

While there are likely a myriad of reasons, examination of crack dynamics after the onset of fracture exhibit interesting trends that are consistent across all experiments with this property set. The key insight from this study is shown in Figure 3.13a and 3.13b; Figure 3.13a depicts the critical stress needed to advance the crack as a function of crack length, for several different initial crack tip positions. The crack always starts in the more brittle phase, and the distance from the initial crack tip to the next interface is d . When $d = 0$, the crack tip starts on the interface between the brittle and tough phases,

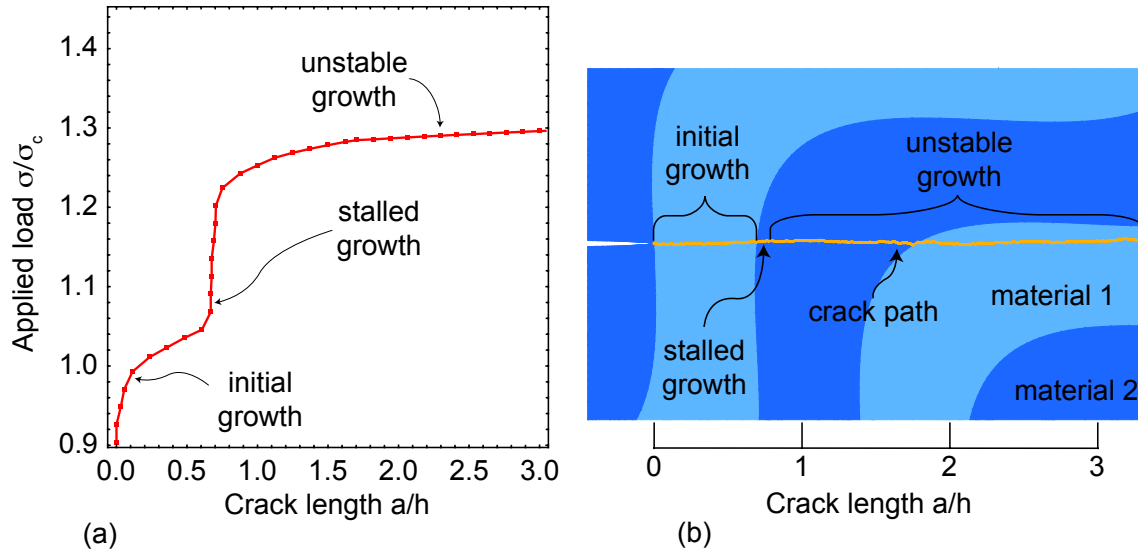


Figure 3.12: An example (a) crack resistance curve with the three crack growth domains marked and (b) the crack path plotted on top of the microstructure with the domains indicated.

propagating immediately into the tougher phase. The corresponding crack speed as a function of crack length is shown in Figure 3.13b. It is clear that as the crack tip starts further from the tougher phase, the plateau resistance is a smaller the percentage of expected resistance.

Figure 3.13b suggest that this is a dynamic effect; when the crack in the brittle layer advances, it accelerates and reaches the next layer with finite speed which is far greater than the velocity of the remote boundary. Hence, while the far field remote loading is quasi-static, the local crack tip fields are undoubtedly influenced by inertial effects of the propagation crack tip. These inertial effects add to the far-field loading and drive the crack forward, even though the elastic contribution by itself is insufficient (based on conventional elastostatic analysis). From this, it seems clear that the percentage of maximum resistance obtained during an experiment will be impacted by differences in wave speed between the layers (i.e. relative density), as well as all of the elasticity and fracture parameters fixed in this study.

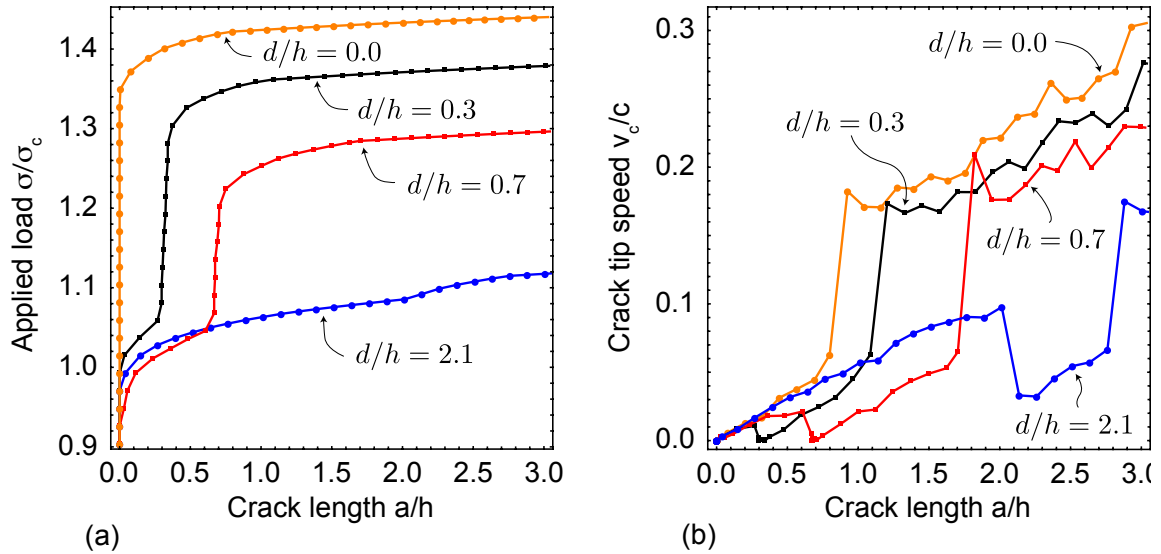


Figure 3.13: Plots of (a) the crack growth resistance for different d/h values showing that increasing the d/h value reduces the needed load for unstable growth. (b) shows the crack speed normalized by the wave speed of material #1 for the same d/h values.

This will be a focus of future work, which will broaden the range of density, moduli and toughness mismatch between the layers. It should be emphasized that such simulations are intractable without efficient adaptive remeshing algorithms that limit the cost of the simulation by translating the local refined mesh with the propagating crack. Conducting an alternative study with static meshes that have refined regions throughout the domain of crack advance (which is limited in this case because the crack direction is relatively fixed) would take multiple years to run.

Chapter 4

Multiphysics Framework for Integrated Diffusion, Oxidation and Creep

4.1 Introduction

High temperature oxidation leads to a host of mechanisms that degrade mechanical performance, including: (*i*) outright erosion (loss of material due to oxidation followed by oxide volatilization), (*ii*) the generation of elevated stress driven by the volumetric expansion associated with oxide formation, and (*iii*) oxide cracking during cooling, which provides additional pathways for reactant ingress and/or propagates into the adjacent material. These behaviors can involve strong coupling between transport (governing the flux of reactants to oxide growth boundaries), domain evolution (describing changes in shape of oxide domains as time progresses), and stresses (driven by the volumetric expansion of the oxide and mitigated by creep).

The framework described in this chapter consists of computational modules for dif-

fusion, domain evolution and stress analysis. The principal contribution of the work in this chapter relates to their integration, which involves coordination of solution variables across time and space: e.g. concentrations, flux arriving at the growth boundary, displacements, eigenstrains associated with oxide expansion, and creep strains. As will be illustrated, the individual modules utilize well established numerical techniques.

The choice of numerical schemes chosen for the framework was guided by two factors. First and foremost, the schemes utilized here provide a facile pathway to integrate this framework with a model that tracks crack path evolution, i.e. that described in Chapter 2 and applied in Chapter 3. The present framework utilizes the same triangular mesh to facilitate direct mapping of relevant solution variables to that framework. Further, it uses the same mesh generation algorithms needed to track large changes in geometry; in future work, crack tracking with local refined meshes can be combined with oxide domain evolution. Second, because the modules utilize conventional finite element descriptions, future modifications (to material behavior, geometry, boundary conditions, etc.) will be straightforward. While other numerical schemes may offer advantages for any one portion of the framework – e.g. level set methods to track domain evolution – they pose integration challenges that outweigh their isolated benefit. Based on the learnings of constructing this framework, future recommendations for different numerical schemes are provided.

The principle physical motivation for the framework in this chapter pertains to oxidation behaviors in SiC ceramic matrix composites, utilized in gas turbines that are exposed to prevalent amounts of water vapor and gaseous hydrogen. That said, it should be appreciated that the modules can be easily adapted to other material systems, e.g. with suitable changes in material descriptions, the relevant species to be tracked, and the description of oxide growth.

In a typical application environment for SiC composites, temperatures can exceed

800°C, such that silicon carbide oxidizes to form silica glass SiO_2 . Both water vapor and oxygen may play an important role in oxidation; to simplify the presentation of the framework, only oxygen is assumed to be present. (Again, additional reactants and reaction products can be tracked by expanding the numerical implementation of the relevant modules.)

The key behaviors are illustrated in Figure 4.1. Oxidants are able to diffuse from a gaseous environment, through existing oxide and react with underlying silicon carbide. The molar volume of silica is much larger than silicon carbide, which results in a large expansion upon oxidation that can produce elastic stresses (immediately after oxide formation) that are greater than 20 GPa in the newly formed oxide. However, at the relevant temperatures of formation, silica undergoes non-linear viscoelastic creep. Silica is typically assumed to behave like a Maxwell material akin to an elastic spring and viscous dashpot in series with a non-linear, stress dependent viscosity.

The relaxation behavior for various initial stresses are shown in Figure 4.2 for a representative temperature of 1000°C. A significant amount of relaxation occurs within the first few seconds, but the amount of relaxation depends on the initial stress σ_o . Extreme initial stresses on the order of $\sigma_o \geq 10$ GPa, such as those caused by the volumetric expansion associated with oxidation, relax on the order of a few seconds. More moderate stresses, like those caused by remote loading of components on the order of 1 GPa, relax over a larger time scale due to the strongly non-linear component of the relevant creep law. The majority of the relaxation, regardless of the level of the initial elastic stress, occurs within minutes.

The remainder of this chapter describes how the framework incorporates relevant behaviors as follows. In Section 4.2, the diffusion problem of reactant transport and subsequent oxidation growth is described. Section 4.3 describes how creep is incorporated into the mechanical analysis module. Module integration is outlined in Section 4.4. To

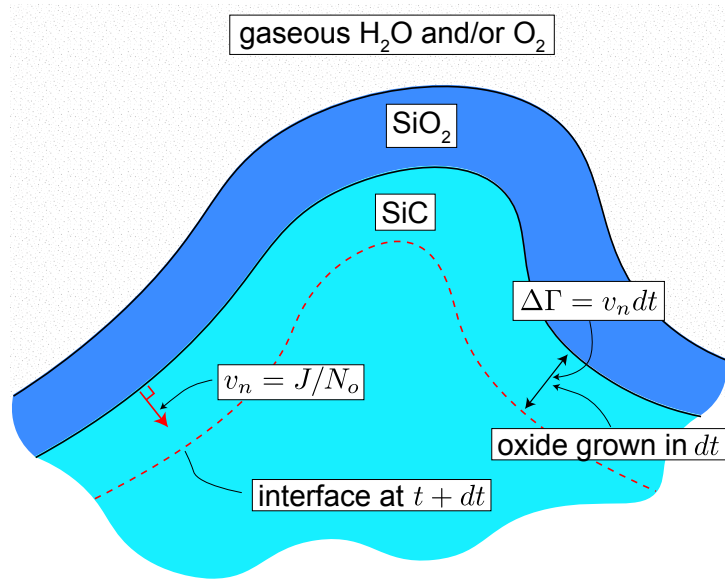


Figure 4.1: Schematic illustration of oxide growth on silicon carbide. Oxidant molecules like oxygen and water diffuse through the existing oxide, react with silicon carbide and advance the oxide/silicon carbide interface (shown in red) inward.




Figure 4.2: Plots of the relative stress relaxation for various initial stresses, σ_o at a temperature of 1000°C. The non-linear viscosity is stress dependent and higher initial stresses lead to faster initial relaxation rates.

illustrate coupled behaviors, Section 4.5 analyzes an oxide layer growing on a circular SiC fiber cross-section, and compares the results to previous quasi-analytical results. Finally, the framework's convergence behaviors, which provide practical guidelines on numerical parameters, is summarized in Section 4.6.

4.2 Transport and Oxidation

One of the key challenges in predicting oxidation phenomena is to couple transport of the oxidant molecules to the reaction front to evolution of the oxide domain, defined by motion of the oxide/substrate boundary. In broad terms, this involves computing the flux of the oxidant species to the domain boundary and computing the velocity of the evolving boundary based on the rate of reactant consumption. To illustrate this coupling, the Deal-Grove model of time-dependent oxidation on a flat surface is first described to outline the key underlying assumptions. The assumptions of this model are then generalized to non-planar surfaces to provide the mathematical descriptions of domain boundary motion used in the present work.

While Deal and Grove first studied oxide growth on silicon in one dimension [12], others have since shown that their approach can be applied to silicon carbide [14,15]. A common assumption when applying the Deal-Grove oxidation model to silicon carbide is to ignore any by-products of the oxidation. As such, the current framework implementation tracks a single diffusing species.

In Deal and Grove's one dimensional model for silicon, diffusion and reaction are combined into a unified model. The model relies on the assumption that the fluxes of oxidant from the atmosphere into the oxide, through the oxide and at the oxide/silicon interface are all equal at any instant in time. It is worth noting that any transients are ignored, which is not necessarily the case in the current numerical framework.

In the Deal-Grove model flux from the atmosphere into the oxide, J_1 , is

$$J_1 = \beta (C_{atm} - C_o) \quad (4.1)$$

where β is the gas-phase transport coefficient, C_{atm} is the concentration of the oxidant in the ambient atmosphere and C_o is the concentration at the free edge of the oxide. The flux through the oxide, J_2 , is described by Fick's first law:

$$J_2 = -D_{eff} \frac{dC}{dx} \quad (4.2)$$

with an effective diffusion coefficient in the oxide D_{eff} which takes into account the effects of localized ionic charges. Since diffusion through the oxide is assumed to be at steady state, the flux through the oxide is constant and can be expressed as

$$J_2 = D_{eff} \frac{C_o - C_i}{h_{ox}} \quad (4.3)$$

where C_i is the concentration at the interface. Finally, the flux at the oxide/silicon interface is

$$J_3 = k_s C_i \quad (4.4)$$

where k_s is the surface reaction rate.

By setting $J = J_1 = J_2 = J_3$ and eliminating C_i and C_o the total flux J is

$$J = \frac{k_s C_{atm}}{1 + k_s/\beta + k_s h_{ox}/D_{eff}} \quad (4.5)$$

The growth of the oxide layer is controlled by the differential equation

$$\frac{dh}{dt} = \frac{J}{N_o} \quad (4.6)$$

where N_o is the number of oxidant molecules incorporated into a unit volume of oxide. Solving this differential equation with the initial condition

$$h_{ox} = h_o \quad (4.7)$$

at time $t = 0$ yields

$$h_{ox}^2 + A_{DG}h_{ox} = B_{DG}(t - \tau) \quad (4.8)$$

which is the Deal-Grove model for oxide growth. The coefficients A_{DG} , B_{DG} and τ are

$$A_{DG} = 2D_{eff} \left(\frac{1}{k_s} + \frac{1}{\beta} \right) \quad (4.9)$$

$$B_{DG} = 2D_{eff} \frac{C_o}{N_o} \quad (4.10)$$

$$\tau = \frac{h_o^2 + A_{DG}h_o}{B_{DG}} \quad (4.11)$$

The constants A_{DG} and B_{DG} are based on the material properties and τ is an offset in time based on the initial thickness of oxide. It should be noted that the gas-phase transport coefficient is much greater than the effective diffusion coefficient such that the oxide is saturated at the free edge and J_1 reaches a constant value. Oxide growth is then limited by diffusion through the oxide and reaction of new silicon carbide. This simplification is used in the numeric modeling.

As the Deal-Grove model is one dimensional, it assumes the flux is unidirectional across any surface in the model. For more complex geometries where diffusion occurs in multiple directions, oxide growth is nonetheless a function of the flux of oxidant molecules

to the growth interface and the reaction rate, k_s . The velocity of the oxide growth interface v_Γ relative to the free edge depends on local fluxes normal to the interface, as in:

$$v_\Gamma = \frac{J_n}{N_o} \quad (4.12)$$

where J_n is the flux perpendicular to the growth interface Γ . For a small enough time increment dt , oxide growth must be linear. The spatial motion of the oxide boundary, defined as $\Delta\Gamma$, is then:

$$\Delta\Gamma = v_\Gamma \Delta t = \frac{J dt}{N_o} \quad (4.13)$$

Thus, the boundary motion at any given time can be calculated from the flux of oxidant that reaches the reaction surface. The flux that reaches the surface can be calculated from the concentration profile from a diffusion analysis. Therefore, evolution of the oxide boundary is controlled by the solution to the diffusion problem, which yields the flux through the boundary at any given time.

Here, bulk diffusion through a solid is assumed to obey Fick's second law which implies

$$\frac{\partial C}{\partial t} = D \nabla^2 C \quad (4.14)$$

where C is the concentration at a point in space at time t and ∇ is the gradient operator. Here, finite element formulations are utilized to generate numerical solutions for arbitrary geometries. The finite element approximation for the diffusion equation can be written as

$$[M] \{\dot{C}\} + [K_{diff}] \{C\} = \sum_i^n \{Q_i\} \quad (4.15)$$

where

$$[M] = \int_V [N]^T [N] dV \quad (4.16)$$

$$[K_{diff}] = \int_V D [B]^T [B] dV \quad (4.17)$$

$$\{Q_i\} = \int_{S_i} q_i [N]^T dS_i \quad (4.18)$$

with $[N]$ is the matrix of element shape functions, $[B]$ is the matrix of the derivatives of the shape functions and q_i is an imposed flux along surface S_i . It should be noted here that it is assumed that the diffusivity D is isotropic and spatially uniform.

When SiC oxidizes, there is an outward flux on the interface between the silicon carbide and silica that provides the oxidant reacting with the silicon carbide. The outward flux at any point along the interface can be computed as $q_{rxn} = k_s C$ where k_s is the surface reaction rate and C is the concentration at that point. Since the concentration at any point can be interpolated as $C = [N] \{C\}$, the nodal equivalent flux on the reacting boundary S_{rxn} can be written as:

$$\{Q_{rxn}\} = \int_{S_{rxn}} k_s [N]^T [N] \{C\} dS_{rxn} \quad (4.19)$$

This expression for flux at the reaction boundary, however, leads to an implicitly defined system, as the concentrations $\{C\}$ exist on both sides of the equation. As will be explained shortly, an explicit system is desired in order to find the solution quickly. It must hold that there is some equivalent diffusivity matrix $[K_{rxn}]$ relating the concentrations to the reaction flux:

$$[K_{rxn}] \{C\} = \{Q_{rxn}\} \quad (4.20)$$

This equivalent diffusivity matrix is then

$$[K_{rxn}] = \int_{S_{rxn}} k_s [N]^T [N] dS_{rxn} \quad (4.21)$$

The finite element equations in the absence of any other imposed flux is then obtained by combining Equations 4.15, 4.20 and 4.21:

$$[M] \{\dot{C}\} + ([K_{diff}] - [K_{rxn}]) \{C\} = 0 \quad (4.22)$$

This system is fully explicit and can be solved with simple time-stepping. While here are many different techniques to solve a dynamic system of equations, this work adopts a simple updating scheme. Over a short enough time increment, Δt , change in concentration must be linear. The time derivatives of the concentration, $\{\dot{C}\}$, are then:

$$\{\dot{C}\} = \frac{\{\dot{C}^{t+\Delta t}\} - \{\dot{C}^{t_0}\}}{\Delta t} \quad (4.23)$$

Taking this into account, the finite element equations become:

$$\left(\frac{1}{\Delta t} [M] + [K_{diff}] - [K_{rxn}] \right) \{C^{t_0+\Delta t}\} = \frac{1}{\Delta t} [M] \{C^{t_0}\} \quad (4.24)$$

with the concentration at the previous time increment moved to the right hand side. For simplicity, the right hand side can be defined as:

$$\{R\} = \frac{1}{\Delta t} [M] \{C^{t_0}\} \quad (4.25)$$

As such, the final form of the diffusion equations is

$$\left(\frac{1}{\Delta t} [M] + [K_{diff}] - [K_{rxn}] \right) \{C^{t_o+\Delta t}\} = \{R\} \quad (4.26)$$

At first glance, the transformation from an implicit system to an explicit system may seem unnecessary. However, the explicit formulation with simple updating yields finite element equations of the form $[A] \{x\} = \{B\}$. Many fast, well-established equation solvers can be used to solve this system directly for each time increment.

The equations laid out above are general and can be used for multiple materials in contact with each other, provided one assumes perfect diffusivity between materials. (Otherwise, additional boundary conditions need to be imposed using standard techniques.) Diffusion through any medium, including gases, can be modeled provided the domain is meshed and effective properties are assigned. For example, narrow cracks can be modeled by assigning an effective diffusivity that combines molecular and Knudsen diffusion defined by:

$$D_c = \frac{D_m D_K}{D_m + D_K} \quad (4.27)$$

where the Knudsen diffusivity, D_K and the molecular diffusivity, D_m , are given by:

$$D_K = \frac{2\delta}{3} \sqrt{\frac{8RT}{M}}; \quad D_m = \frac{0.0018583 T^{3/2}}{P\sigma^2\Omega_{AB}} \left(\frac{1}{M_{air}} + \frac{1}{M} \right)^{1/2} \quad (4.28)$$

with δ as the crack opening, R is the ideal gas constant, T is the temperature, M_{air} is the molecular weight of the air within the diffusion channel, M is the molecular weight of the diffusing species, σ is the average collision diameter, Ω is the collision integral, and P is the pressure. This will be used in the studies of barrier coatings in Chapter 5.

4.3 Mechanical Deformation

The finite element discretization described in the previous section can also be used to solve for the thermomechanical response of the entire domain, including newly formed oxide and remaining substrate. In this section, the details of the numerical formulation needed to account for creep is described, with an emphasis on how creep strains are handled. Ignoring accelerations, the principle of virtual work states:

$$\int_V \{\delta\epsilon\} [\sigma] dV = \int_S \{\delta\tilde{u}\} \{f\} dS \quad (4.29)$$

for a set of tractions $\{f\}$. The approximate displacements $\{\tilde{u}\}$ are an interpolation of the nodal displacements $\{u\}$ using the same shape functions $[N]$ as used in the diffusion problem by

$$\{\tilde{u}\} = [N] \{u\} \quad (4.30)$$

The total strain at any point in time is the superposition of strains due to elastic deformation (ϵ_{el}), inelastic deformation such as creep (ϵ_{cr}), any transformation strains (ϵ_{tr}) and any strain implicitly stored in the description of the mesh (ϵ_m). The latter arises when the role of accumulated strain (from creep, transformation or elastic deformation) is accounted for by updating the geometry. With this in mind, the principle of virtual work can be written as:

$$\int_V \{\delta u\} [B]^T [E] \{\epsilon_{el}\} dV = \int_S \{\delta u\} [N]^T \{f\} dS \quad (4.31)$$

or, noting that $\epsilon = \epsilon_{el} + \epsilon_{cr} + \epsilon_{tr} + \epsilon_m$

$$\int_V \{\delta u\} [B]^T [E] \{\epsilon\} dV = \int_S \{\delta u\} [N]^T \{f\} dS + \int_V \{\delta u\} [B]^T [E] \{\epsilon_{tr}\} dV + \int_V \{\delta u\} [B]^T [E] \{\epsilon_{cr}\} dV + \int_V \{\delta u\} [B]^T [E] \{\epsilon_m\} dV \quad (4.32)$$

where $[B]$ is the matrix of shape function derivatives and $\{\epsilon\} = [B] \{u\}$. Since the statement must be true for all virtual displacements $\{\delta u\}$ the finite element equations can be expressed as

$$[K] \{u\} = \{F\} + \{F_{tr}\} + \{F_{cr}\} + \{F_m\} \quad (4.33)$$

where

$$[K] = \int_V [B]^T [E] [B] dV \quad (4.34)$$

$$\{F\} = \int_S [N]^T \{f\} dS \quad (4.35)$$

$$\{F_{tr}\} = \int_V [B]^T [E] \{\epsilon_{tr}\} dV \quad (4.36)$$

$$\{F_m\} = \int_V [B]^T [E] \{\epsilon_m\} dV \quad (4.37)$$

$$\{F_{cr}\} = \int_V [B]^T [E] \{\epsilon_{cr}\} dV \quad (4.38)$$

The tractions are determined by loading or boundary conditions. The various strain components contribute to the nodal force vector. The transformation strains $\{\epsilon_{tr}\}$ are eigenstrains that result from the volumetric expansion of silicon carbide into silica. They are handled similarly to any thermal expansion that may occur. The simulation framework stores the deformation history as stored strains $\{\epsilon_m\}$ by updating the mesh.

Creep strains $\{\epsilon_{cr}\}$ are calculated using the constitutive law for the oxide. The

common assumption is that creep in silica follows a non-linear Maxwell viscoelastic model [14,15]. In the Maxwell model, the strain rate is

$$\dot{\epsilon}_{ij} = \frac{1}{E} [(1 + \nu) \dot{\sigma}_{ij} - \nu \dot{\sigma}_{kk} \delta_{ij}] + \frac{1}{2\eta} \left(\sigma_{ij} - \frac{1}{3} \sigma_{kk} \delta_{ij} \right) \quad (4.39)$$

for a Young's modulus of E , Poisson's ratio of ν , stress rate of $\dot{\sigma}_{ij}$, viscosity of

$$\eta = \eta_o \frac{\tau/\tau_c}{\sinh(\tau/\tau_c)} \quad (4.40)$$

and the Kronecker delta, δ_{ij} . The viscosity is dependent on a reference viscosity of

$$\eta_o = 3.8 \times 10^{-14} \exp\left(\frac{712\text{kJ/mol}}{RT}\right) \text{ Pa} \cdot \text{s} \quad (4.41)$$

as well as an equivalent shear stress, τ and a critical shear stress, τ_c . The critical shear stress is slightly dependent on temperature but is taken to be a constant of 100 MPa.

The equivalent shear stress is taken to be

$$\tau = \left[\frac{1}{2} s_{ij} s_{ij} \right]^{1/2} \quad (4.42)$$

where s_{ij} is the deviatoric stress defined as

$$s_{ij} = \sigma_{ij} - \frac{1}{3} \sigma_{kk} \delta_{ij} \quad (4.43)$$

It should be noted that the strain rate can be rewritten as

$$\dot{\epsilon}_{ij} = \frac{1}{E} [(1 + \nu) \dot{\sigma}_{ij} - \nu \dot{\sigma}_{kk} \delta_{ij}] + \frac{1}{2\eta} s_{ij} \quad (4.44)$$

which makes it clear the creep is controlled by the deviatoric stresses. Hay states that

the creep only occurs via the shear. [14]. A form of the Maxwell model for equivalent shear stress relaxation rate is

$$\dot{\tau} = -\frac{E\tau}{2(1+\nu)\eta(\tau)} \quad (4.45)$$

which has the solution

$$\tau = 2\tau_c \tanh^{-1} \left(\tanh \left(\frac{\tau_o}{2\tau_c} \right) \exp \left(-\frac{Et}{2\eta_o(1+\nu)} \right) \right) \quad (4.46)$$

It should be noted that [14] has some typographical errors in the expression for τ that have been corrected above. Other stresses, such as σ_{11} , σ_{22} and σ_{12} are taken to relax proportionally to τ/τ_o . This relaxation is different from some literature but is used for its simplicity.

4.4 Module Integration

In this section, the integration of this transport analysis and the stress analysis outlined in Section 4.3 is described with an emphasis on the underlying algorithms. The multiphysics framework solves for growth and stresses through an iterative time stepping scheme. Each time increment consists of a number of tasks, which are illustrated in Figure 4.3 and broad summarized as follows. Start by assuming that the concentration distribution is known at time t for a mesh defined at the same time. The concentration at time $t + dt$ can be computed immediately, as it is disconnected from the stress analysis. However, to track the domain and stress evolution, a specific sequence is needed due to their interdependency. This sequence is as follows; at time t , the next physical behavior that occurs is that the oxide interface advances incrementally, because the reaction kinetics are much faster than mechanical response. Hence, the first task is to advance

this boundary, and impose volumetric strains in the region that has been converted from SiC to oxide. These imposed volume strains immediately generate stresses, which deform the solid and change total displacements. That is, after advance, the immediate change in displacements from volume expansion is computed assuming creep strains are fixed. Once the stresses are found from a static analysis with fixed creep strains, the creep strain rates are computed for that instant at time, and the creep strains are updated to $t + dt$. Creep strain increments that underlie that relaxation process are added to previous totals to retain the history of creep strains. The updated creep strains are then used to update stress relaxation over the increment, which computes the final nodal locations at time $t + dt$. Hence, at the end of performing all tasks, one has obtained the concentration distribution at time $t + dt$, the accrued creep strains throughout the system at time $t + dt$, and current deformed state at time $t + dt$. The problem is then remeshed to account for the motion of the domain boundary during this increment. This process is then repeated.

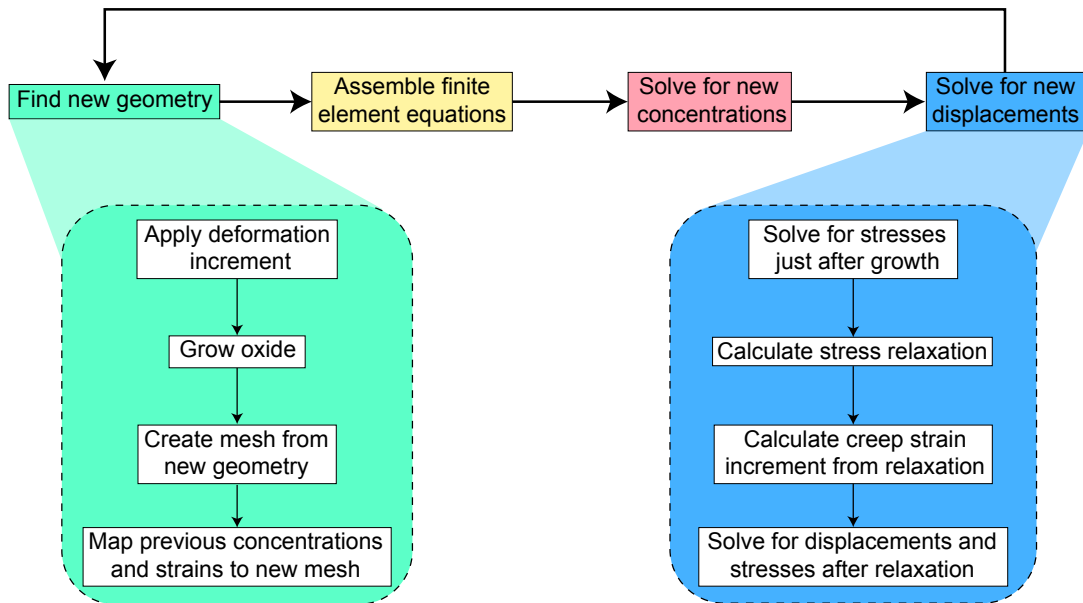


Figure 4.3: Flowchart showing the major steps in the multiphysics framework. Finding the new geometry and solving for displacements each have many tasks which are also shown.

During the updates associated with the time increment, the motion of the interface converts elements from SiC to SiO₂. The geometry in the simulation is described by a series of adjacent but non-overlapping polygons. Each polygon represents a distinct continuous material domain with uniform material properties. It should be noted that any cracks, such as those in the system described in Chapter 5, are explicitly meshed and are considered a continuous medium with a negligible elastic modulus and a prescribed diffusivity. Each polygon is defined as a clockwise ordered list of coordinate points defining the boundary.

The first step during each time increment is to find the new geometry. For the first time increment, or $t = 0$, the new geometry is input into the framework and nothing more needs to be done. When $t > 0$, however, the displacement increment found in the previous time increment must be applied to the geometry used in that time increment. The displacement increment is applied to both the material polygons and the previous nodal positions. The stored strains associated with the previous mesh are also updated with the displacement increment. This places the entire solution from the previous step into the previous mesh.

Once the previous displacement solution is applied to the previous mesh, the oxide polygon is grown based on the concentration solution. The oxide growth procedure is shown in Figure 4.4. The growth of the oxide polygon starts by identifying the points of the oxide polygon that are along the boundary defining the oxide silicon carbide interface. For each of these points, the elements connected to the point that also lay on the boundary are identified. The fluxes of these elements and the normal vectors of the sides along the boundary are averaged as shown in Figure 4.4a. The average flux \bar{J} of the connected elements is then projected onto the average of the normal vectors \bar{n} , providing an averaged normal flux \bar{J}_\perp shown in Figure 4.4b. The boundary motion $\Delta\Gamma$ is then calculated from the averaged normal flux, the number of oxidant molecules incorporated

into a unit volume of oxide N_o , the time increment Δt and a proportional constant f :

$$\Delta\Gamma = f \frac{\bar{J}_\perp \Delta t}{N_o} \quad (4.47)$$

The proportional constant f is a function of the expansion strain associated with oxidation of silicon carbide into silica and the Poisson's ratio of silica and accounts for the expansion eigenstrains applied in the mechanics module as well as the plane strain conditions imposed. Another way to think of this nuance is that the boundary is moved by the amount of silicon carbide consumed.

This process is repeated for all the points on the oxide polygon that also lay on the boundary with silicon carbide. The remainder of the points of the oxide polygon are unchanged during this procedure. When the mechanics module is used, the expansion strain is assigned to the entire oxide, but the stored strains in the new oxide layer are zero. Since the expansion strain is not accounted for in the stored strain, the new oxide layer expands.




Figure 4.4: Schematic illustration of oxide growth procedure on a subset of the oxide/silicon carbide interface. (a) depicts how the elemental fluxes J_1 and J_2 are averaged to get the average flux at the current node \bar{J} . The normal of the growth interface n_1 and n_2 are averaged to get \bar{n} . (b) shows how the average flux is projected onto the average normal resulting in the perpendicular flux \bar{J}_\perp . (c) shows how the motion of the interface is calculated and applied to the node. This procedure is repeated at all nodes along the boundary.

While the previous discussion on oxide growth only detailed how the points on the oxide polygon are shifted, the same movements are also applied to the points on the

silicon carbide polygon that fall along the boundary with the oxide polygon. For more simple geometries consisting of just silicon carbide and silica such as those used in the verifications in the next section, the new geometry is completely defined. In more complex geometries like the system of a cracked protective coating on top of a silicon carbide substrate with a thin oxide layer between that is studied in Chapter 5 some further work must be done involving the polygons for the protective coating and the crack.

Due to the large expansion that occurs with oxidation, the oxide may extrude up the crack. When this occurs, the oxide polygon may end up slightly overlapping the coating polygon. This overlap is very small relative to other dimensions and involves very little material. In order to reduce computational headaches caused by the overlap, the points on the oxide polygon that fall within the coating are removed. The polygon for the crack is recreated from the oxide and coating polygons as well at the boundary with the atmosphere that is defined at the beginning of the simulation. The points on each of the polygons and boundary that do not lay on another are used to create the crack polygon.

The polygons that define the geometry, either those directly input into the framework for $t = 0$ or those found when $t > 0$, are used to create the mesh for the current time increment. The open source Triangle triangulation library was used to mesh the geometry with three-noded constant strain (or flux) elements [35]. This library is freely available and both fast and reliable. Delaunay triangulations with tens of thousands of elements can be found in under a second and can be directly integrated into a C++ application.

The polygons that define the different material regions are resolved to a list of points and segments which are then input into Triangle. This ensures that domain boundaries are preserved in the mesh, i.e. no elements will be split between domains. The Triangle subroutine allows for fine control of element size. This functionality was used to maintain a constant average element size within the oxide domain and a gradient in size moving away from the silicon carbide/oxide interface. By controlling the element size in this way,

the number of elements is greatly reduced which speeds up computation time considerably while still keeping small elements in the regions of importance. While the base mesh is completed by Triangle, additional work must be done before the mesh can be used to assemble the finite element matrices.

For all but the first time increment, nodal and elemental data from the previous mesh must be mapped onto the new mesh. Mapping nodal quantities such as concentration between meshes is trivial; the shape functions for the previous element provide the interpolation functions necessary to calculate a nodal value in the new mesh. The interpolated value \tilde{c} at a point (x, y) within an element is

$$\tilde{c}(x, y) = \sum_{i=0}^2 N_i(x, y) c_i \quad (4.48)$$

where N_i is the shape function corresponding to node i . To find the interpolated value, the element in the old mesh that holds the point of interest must be found. In the case where a point is on the boundary of one or multiple elements, any element can be used to interpolate the value since field variables are continuous along elemental boundaries. No other work is necessary to map nodal values to the new mesh. It should be noted here that the only nodal value mapped between meshes is concentration. While the displacement increment found in the previous time increment is a nodal value, it is mapped via the elemental stored strains.

Mapping the elemental values of creep and stored strains is more complicated than concentration. Interpolation cannot directly happen with elemental values. The elemental strains must first be converted to equivalent nodal strains. Before the strains can be, an equivalent nodal creep must be found for each node. If a node is connected to n

elements, an average for an equivalent nodal strain, $\{\tilde{\epsilon}_j\}$, can be defined as

$$\{\tilde{\epsilon}_j\} = \frac{\sum_{i=1}^n A_i \{\epsilon_j^i\}}{\sum_{i=1}^n A_i} \quad (4.49)$$

where A_i is the area of element i connected to node n and strain type j (either creep or stored). Once elemental data is converted to equivalent nodal values on the old mesh, the nodal values for the new mesh can be individually interpolated using the same technique as the concentrations. The elemental creep and stored strains for an element in the new mesh are then recovered by averaging the equivalent values at the nodes that define the element.

This averaging technique does not perfectly recreate the strain state in the new mesh because linear elements are used and strain is discontinuous between elements in the old mesh. This has the effect of smoothing out high gradients in strains. While technically an error that can be alleviated with increasing mesh density, this has the benefit of smoothing out any local numerical noise that arises from boundary perturbations generated from the motion of nodes along the boundary. The effect of this artificial smoothing on convergence is discussed in Section 4.6. It should be noted here that nodes that are on boundaries between materials are handled slightly different. For those nodes, there are equivalent nodal creep and stored strains for each material. Only elements of the old mesh in the specific material are used to determine the equivalent nodal strains for that material.

After the new mesh is created and the data from the previous mesh is mapped if needed, the finite element matrices are assembled. There are two sets of finite element matrices to be assembled: diffusion and mechanical deformation. The definition of the elemental matrices for three-noded constant strain/flux triangular elements are as follows.

For constant strain triangles, the elemental matrix $[m]$ works out to

$$[m] = \frac{A_{el}}{12} \begin{bmatrix} 2 & 1 & 1 \\ 1 & 2 & 1 \\ 1 & 1 & 2 \end{bmatrix} \quad (4.50)$$

and $[k_{diff}]$ is

$$[k_{diff}] = A_{el} D [B]^T [B] \quad (4.51)$$

for an element with area A_{el} . There is not a similar elemental version of the $[K_{rxn}]$ matrix as the integration is done over the surface. Instead, it is done for each elemental edge that falls along the reaction front. For example, if element nodes 1 and 2 are along the reaction front but node 3 is not, the elemental $[K_{rxn}]$ is

$$[k_{rxn}] = \frac{k_s L_e}{6} \begin{bmatrix} 2 & 1 & 0 \\ 1 & 2 & 0 \\ 0 & 0 & 0 \end{bmatrix} \quad (4.52)$$

where L_e is the distance between nodes 1 and 2.

The right hand side of the diffusion finite element equations $\{r\}$ contain the concentration at the beginning of the time increment. For constant flux elements, this is

$$\{r\} = \frac{A_{el}}{12\Delta t} \begin{bmatrix} 2 & 1 & 1 \\ 1 & 2 & 1 \\ 1 & 1 & 2 \end{bmatrix} \begin{Bmatrix} c_0 \\ c_1 \\ c_2 \end{Bmatrix} \quad (4.53)$$

where $\{c_0 \ c_1 \ c_2\}^T$ are the concentrations at the nodes.

Assembly of the elemental matrices is done through a normal process and is not outlined here. It is assumed the atmospheric concentration is high enough to lead to

saturation at exposed edges. With this assumption, diffusion from the atmosphere into material (or into a crack) is assumed to be perfect. This simplifies the diffusion problem and transformed the boundary condition into a fixed concentration along the edge exposed to the atmosphere. The concentration constraints are applied via Lagrange multipliers to the assembled global matrix.

In mechanical deformation, three-noded elements are referred to as constant strain. The stiffness matrix for a constant strain triangular element is

$$[k] = A_{el} [B]^T [E] [B] \quad (4.54)$$

where A_{el} again is the area of the element, $[B]$ is the matrix of shape function derivatives and

$$[E] = \frac{E}{(1-2\nu)(1+\nu)} \begin{bmatrix} 1-\nu & \nu & 0 \\ \nu & 1-\nu & 0 \\ 0 & 0 & \frac{1-2\nu}{2} \end{bmatrix} \quad (4.55)$$

assuming plane strain conditions, a Young's modulus E and Poisson's ratio ν . Neglecting any applied traction, the right hand side is

$$\{f_{tr}\} + \{f_{cr}\} + \{f_m\} = A_{el} [B]^T [E] (\{\epsilon_{tr}\} + \{\epsilon_{cr}\} + \{\epsilon_m\}) \quad (4.56)$$

where $\{\epsilon_{tr}\} = \epsilon_T \{1 \ 1 \ 0\}^T$ for linear transformation strain of ϵ_T , creep strains of $\{\epsilon_{cr}\} = \{\epsilon_{11}^{cr} \ \epsilon_{22}^{cr} \ \epsilon_{12}^{cr}\}^T$ and stored strains of $\{\epsilon_m\} = \{\epsilon_{11}^m \ \epsilon_{22}^m \ \epsilon_{12}^m\}^T$. When the matrices are first assembled, the creep strains mapped from the previous mesh are used. Creep that occurs during the current time increment is handled separately.

Like the diffusion matrices, the elemental matrices are assembled into the global matrices using traditional techniques. Displacement constraints are applied via Lagrange

multipliers to the global matrices.

After assembly of the matrices is completed, the solution process begins. Solving for the concentrations at the end of the time increment is straightforward. The UMFPACK linear solver is used with the diffusion finite element equations to get the concentrations at the end of the time increment. That is all that is needed to complete the diffusion problem on this time increment.

Solving the mechanical deformation problem is more involved and requires solving the full system twice. The first solution of the mechanical deformation problem finds the displacement step without any additional creep allowed during the *current* time increment. In other words, it finds the stress state immediately after growth occurs but before relaxation can occur. Once this displacement step occurs, the elemental stresses are computed. The expected stress relaxation for each element is then found and the increment in creep strains is calculated from the relaxation. The increment in creep strains is added to the current creep strains and the right hand side of the mechanical deformation problem is assembled again. The mechanical deformation problem is solved a second time. This second solution finds the displacement increment after relaxation. This multistep solution technique was also used by Hay in his study of the oxidation of silicon carbide fibers [14]. Time may then be advanced and the next iteration started by growing the oxide polygon again.

4.5 Framework Verification

While some prior work exists that combines oxidation with stress development of silicon or silicon carbide, differences in implementations prohibit a direct comparison. Other implementations use slightly different assumptions about creep and the handling of large displacements. In lieu of a full verification, two primary phenomena were in-

independently verified: oxide growth was compared in one dimension to the Deal-Grove model and stress relaxation was compared to the Eyring model for various temperatures. In addition, the combined effect of oxidation and creep behaviors during oxidation of a single fiber was compared to results of Xu et al. and Hay [14, 15].

The first verification focused on oxide growth. A flat silicon carbide substrate with an initial thin oxide layer on top was simulated. Boundary conditions were applied such that all growth and deformation occurred in one dimension, emulating the conditions assumed in the Deal-Grove model derivation. The implementation matches the Deal-Grove model quite closely for a flat surface as shown in Figure 4.5. The simulation framework slightly under-predicts growth but the error is less than 1% out to almost 150 hours and an oxide thickness of about 600 nm (with Deal-Grove parameters of $A_{DG} = 18.12 \text{ nm}$ and $B_{DG} = 0.75 \text{ nm}^2/\text{s}$). The gas-phase transport coefficient β was assumed to be infinite leading to direct relations between the Deal-Grove parameters and the material properties used in the simulation framework. Dry oxygen was used as the oxidant which corresponds to $C_o = 5.5 \times 10^{-5} \text{ mol}/\text{nm}^3$ and $N_o = 22.5 \text{ mol}/\text{nm}^3$. The effective diffusivity used in the simulation framework was $D_{eff} = 1.54 \times 10^5 \text{ nm}^2/\text{s}$ and the surface reaction rate was $k_s = 1.7 \times 10^4 \text{ nm}/\text{s}$. A time increment of $dt = 500 \text{ s}$ was used. A mesh and time increment resolution study was performed to reduce spatial and temporal discretization errors. An example convergence study is shown in a following section.

While in this benchmark case the flux throughout the system is constant at any given time, that is only a consequence of the geometry and conditions applied to match the Deal-Grove model. In general, the flux throughout the system can vary, as will be seen on oxidation of a curved surface. In other words, the assumption that the flux is constant in the Deal-Grove model was not enforced in the simulation but was a natural result.

Stress relaxation is accurately captured at multiple temperatures with this technique, as shown in Figure 4.6. In these tests, the domain was prevented from displacing in one




Figure 4.5: Plot showing the simulated oxide thickness on a flat surface as a function of time and the thickness predicted by the Deal-Grove model. There is little difference between the simulation and Deal-Grove model.

direction and an expansion strain equivalent to that associated with the oxidation of silicon carbide was applied at time $t = 0$. This results in an initial elastic compressive stress of 25.3 GPa. In effect, this domain is the new layer of oxide that forms on silicon carbide. The stress observed in the constrained direction was then extracted. The theoretical stress as predicted by the Maxwell model and the simulation results for temperatures ranging from 800°C to 1200°C show excellent agreement. It should be noted here the speed at which the high initial stress is initially relaxed.

A final, more comprehensive study was conducted of oxidation of a single fiber, similar to the work of Hay [14]. A quarter cross section of a single fiber was modeled to exploit symmetry. Boundary constraints were applied to adequately model a full fiber. Plane strain was assumed with no out of plane strain. The geometry is shown in Figure 4.7. Figure 4.7a shows the beginning of the simulation where the fiber has an initial radius of r_o and initial oxide thickness of h_o . After the oxide has grown in Figure 4.7b, the total radius ($r + h$) is greater than the initial radius due to oxide expansion. This study




Figure 4.6: Plot of the calculated stress relaxation with simulation results for a one dimensional compression test. The initial stress of 25.3 GPa is representative of the stress that occurs immediately after oxidation.

used an initial silicon carbide fiber radius of $r_o = 3 \mu m$ and an initial oxide thickness of $h_o = 5 nm$. The temperature was set to $1000^\circ C$ and the fiber exposed to dry air. A mesh and time increment resolution study was performed to minimize errors. A time increment of 500 s was used and a total time of about 140 hours was simulated. This allowed a good comparison to one of the cases studied by Hay [14].

The oxide thickness on a curved radius of $3 \mu m$ is shown in Figure 4.8a compared to the Deal-Grove model. As is clearly shown, the growth on a curved surface is marginally faster than Deal-Grove predicts for a flat surface. The Deal-Grove model assumes that the flux is constant throughout the system, including through the thickness of oxide. As shown in Figure 4.8b, this is not strictly true for a curved surface. In this figure, the inward flux at each time is normalized by the flux assumed in the Deal-Grove model for an oxide of that thickness. At the interface between the silicon carbide and silica, the flux is greater than the Deal-Grove model while it is lower at the free edge. This increased flux at the interface results in the marginally faster growth seen in Figure 4.8a. This is




Figure 4.7: Schematic illustration of oxidation on a silicon carbide fiber. At time $t = t_o$ shown in (a), the fiber has a radius of r_o and initial oxide thickness of h_o . At some later time shown in (b), the oxide thickness has grown and the outer diameter has grown due to the volumetric expansion that occurs during oxidation.

consistent with results from studies of oxide growth on curved silicon surfaces [84, 85]. In those previous works, numerical models were developed to predict oxide thicknesses on curved silicon surfaces of varying radii.

Qualitatively, the hoop stress results shown in Figure 4.9 follow the results of Hay [14]. At an early time shown in Figure 4.9a, the entire oxide layer is in compression. After some more growth, the hoop stress in the outer layer vanishes, as shown in Figure 4.9b. Further growth and expansion leads to the tensile hoop stresses on the outer edge seen in Figure 4.9c. Continued growth, like that in Figure 4.9d, the thickness of the tensile layer increases. Hay showed similar behavior; the tensile hoop stress on the outer edge is caused by the expansion of newer oxide layers underneath pushing out older layers that have already relaxed.

The hoop stresses through oxide at various times are shown in Figure 4.10a. The hoop stresses at all times show similar behavior; there is compression near the increases

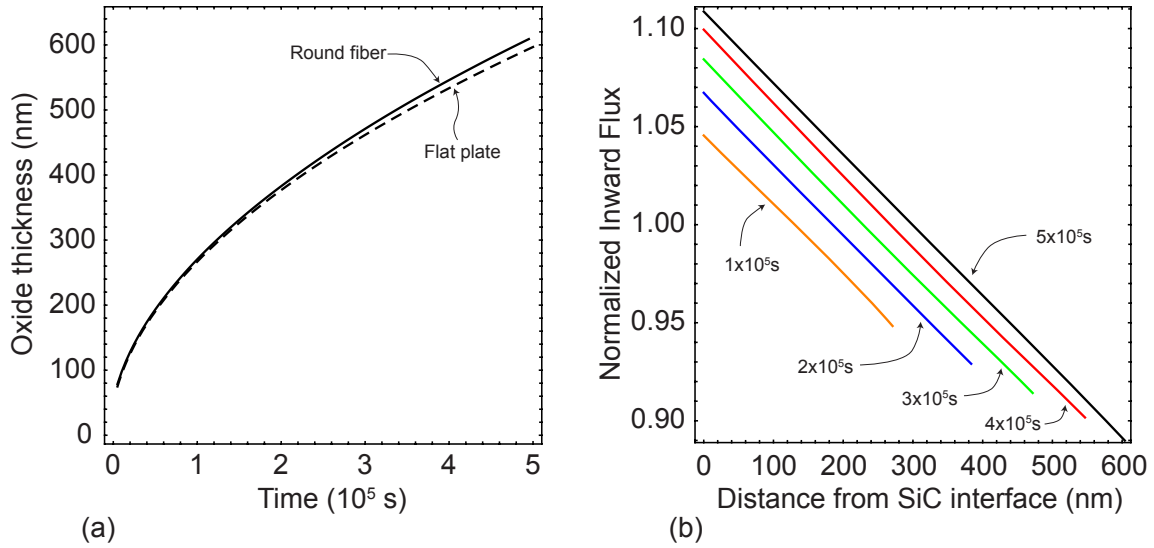


Figure 4.8: (a) Shows the oxide thickness as a function of time on a $3 \mu\text{m}$ radius fiber. (b) Shows the inward flux normalized by the flux assumed by the Deal-Grove model. The curvature causes the flux to not be constant unlike the Deal-Grove model.

that decreases in magnitude. Later times show the stress switching into tension and increasing in magnitude. The outer layers of oxide at any time are in a similar stress state so that later times (thicker oxides) shows a significant tensile layer at about 750 MPa. The stress does not rise above this level, even as the oxide thickens, as shown in Figure 4.10b. The stress does not further increase past this point due to creep.

As can be seen in Figure 4.10b, the results from the current framework do not quantitatively match those from Hay [14]. The discrepancy is due to subtle differences between Hay's model and the framework used here. First, stress relaxation is handled slightly differently. Hay calculates relaxation for the hoop stresses by decreasing them proportionally to the relaxation in the effective shear stress and the difference between the hoop and radial stresses. Radial stresses are not relaxed at all. Hay assumes that the radial stresses are small enough that appreciable relaxation would not occur. In this work, all stress components are relaxed proportionally to the effective shear stress and no assumption is made about stress components not relaxing.




Figure 4.9: Contour plots of the hoop stress at (a) an early time where the oxide layer is in compression, (b) when the hoop stress at the outer layer is approximately zero, (c) the hoop stress at the outer boundary hits a maximum of around 1 GPa and (d) the tensile layer thickens with time.

The second, and arguably more important, difference between Hay and this work is how large deformations are handled. Both works assume small displacements when relating deformation to strain. Hay references back to the undeformed, initial configuration; when geometric changes are small, this is entirely acceptable. However, the large geometric changes associated with oxide growth inherently implies large deformations; i.e. an incremental change such as relaxation should be referenced to the most recent geometry. In the present formulation, the current geometry is updated using the displacement at each time increment, and this is used to compute incremental changes. Small strain kinematics can still be used during an update, since the displacement increment is small relative to the current deformed configuration. During development, an alternate version of the framework was created that continuously referenced the undeformed configura-

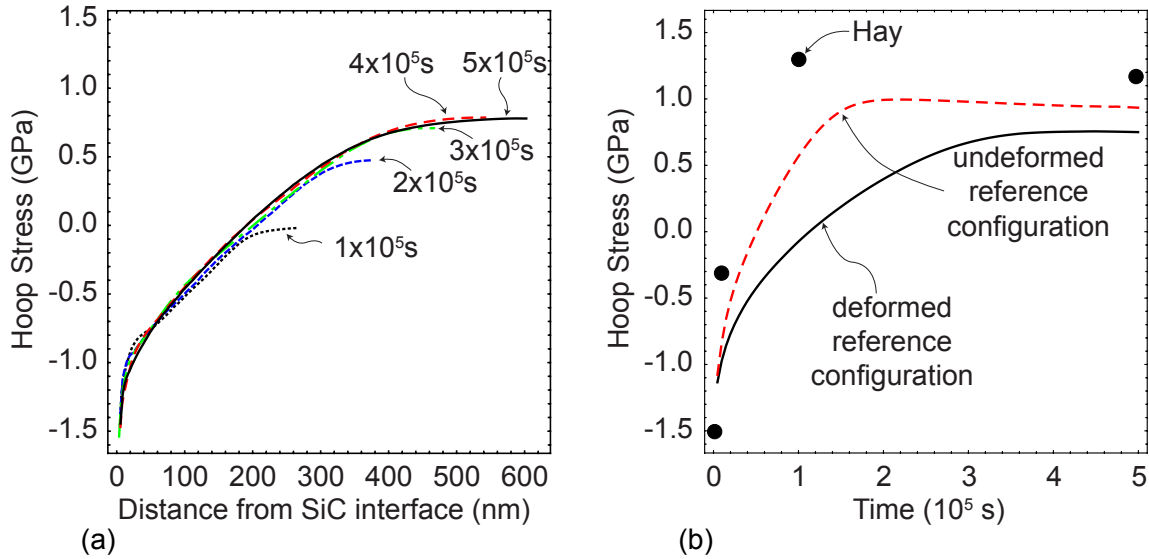


Figure 4.10: Plots showing the hoop stress in the oxide as the simulation progresses. (a) shows the hoop stresses through the thickness of oxide at various times. (b) shows the hoop stress at the outer edge of the oxide as a function of time as well as data from Hay [14] and a previous implementation.

tion, as done by Hay. The hoop stresses at the outer edge of a fiber for that version are shown in Figure 4.10b and are much closer to the stresses Hay found. Remaining differences arise from Hay’s neglect of radial stresses in the multi-axial creep law, which are accounted for in the present framework.

4.6 Simulation Convergence

In the present multiphysics framework, the time-step and element size have strongly interconnected impacts on numerical performance that are different from those associated with isolated modules. Independently, each module exhibits the well-established convergence behaviors expected for finite element frameworks. Once integrated, the overall numerical performance of framework is impacted by several factors.

Remeshing requires mapping data to the new mesh, which unavoidably introduces spatial smoothing since the mapping involves averaging between elements in the old mesh.

This effect is most pronounced for creep strains (and hence stress), which are constant within each three-noded element. While decreasing element size alleviates smoothing effects, it also limits the allowable time step in the diffusion analysis. (Decreasing element size while holding the time step fixed leads to decreasing accuracy in the diffusion analysis.) As the time step is decreased, the increment in boundary motion decreases, and in turn, the mesh resolution of newly formed oxide decreases. In essence, simultaneously decreasing mesh size and time step can improve the inaccuracies associated re-mapping only up to a certain point; eventually, it leads to artificial smoothing due to repeated mappings and spatial variations in stress arising from a coarser mesh in the new oxide.

On the one hand, the effects of spatial smoothing during remapping and spatial variations along the domain boundary are negligible with respect to the diffusion/growth behavior. Spatial smoothing arising from remapping concentrations is negligible since concentrations are continuous along element boundaries (being nodal variables); the mapping is exact. Spatial variations along the boundary when considered relative to the dimensions of the oxide domain are minimal. This is illustrated in Figure 4.11, which provides the convergence behavior of the location of the growth interface for the fiber problem discussed in the previous section. For a fixed time step (Figure 4.11a), the position of the interface (averaged along its length) is highly accurate and insensitive to element size. For a fixed element size (Figure 4.11b), the boundary motion is highly accurate and insensitive to time step, provided it is not so large that a first-order integration scheme becomes inaccurate. (That is, for smaller elements, larger time steps can violate the assumption that the flux is approximately linear over the time interval; e.g. 5 nm^2 elements and time step of 5×10^3 secs).

On the other hand, spatial smoothing during remapping and spatial variations along the domain boundary have a much bigger influence on stress evolution. First, consider the effect of remapping, using the stress at the outer edge of the fiber oxidation problem

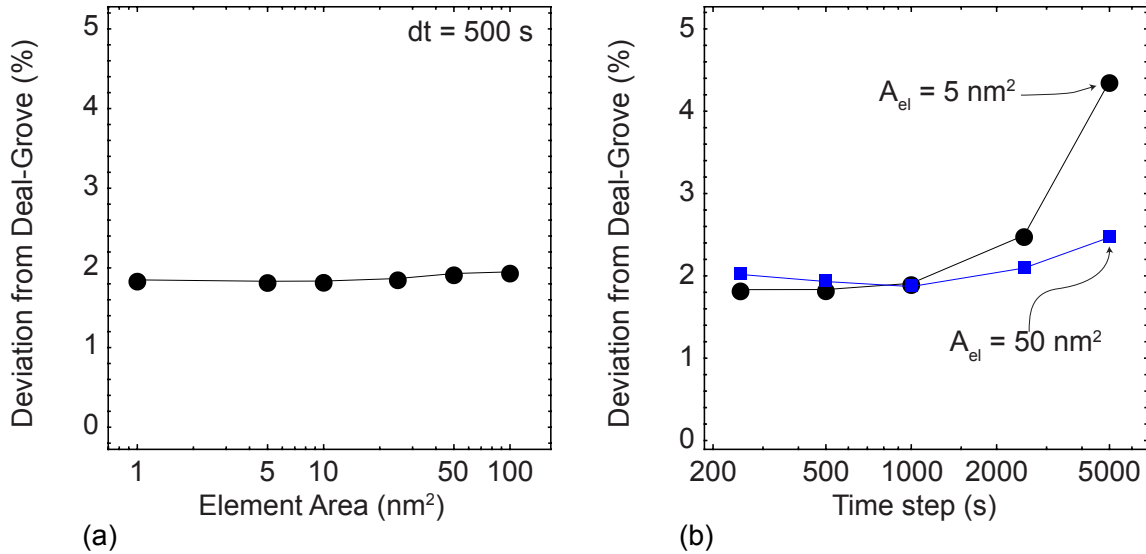


Figure 4.11: Plots of the deviation in thickness on growth of a round fiber compared to that predicted by the Deal-Grove model as a function of (a) element area for a fixed time increment and (b) a function of time increment for two different element sizes.

described in the previous section. Figure 4.12 illustrates convergence behaviors associated with adjusting element size and time step. For a desirable time step that avoids errors in the boundary motion (i.e. the diffusion analysis), Figure 4.12a illustrates that the stress at the outer edge of the fiber clearly asymptotes to a fixed value independent of the element size. However, if the time step is decreased too far for a given element size, spatial smoothing during remapping can artificially smooth the stress distribution, as shown in Figure 4.12b.

Second, consider the impact of the size of the time step on spatial variations in the sliver of new oxide formed at the interface during the time step. For fixed mesh size, increasing the time step increases the number of elements in the new oxide domain. Regardless of how many elements are swept up during domain advance, the boundary will always cut a row of elements, since the boundary must be advanced based on the current mesh. This cutting of elements introduces spatial variations along the interface because elements that fall mostly inside the new oxide will experience significant creep

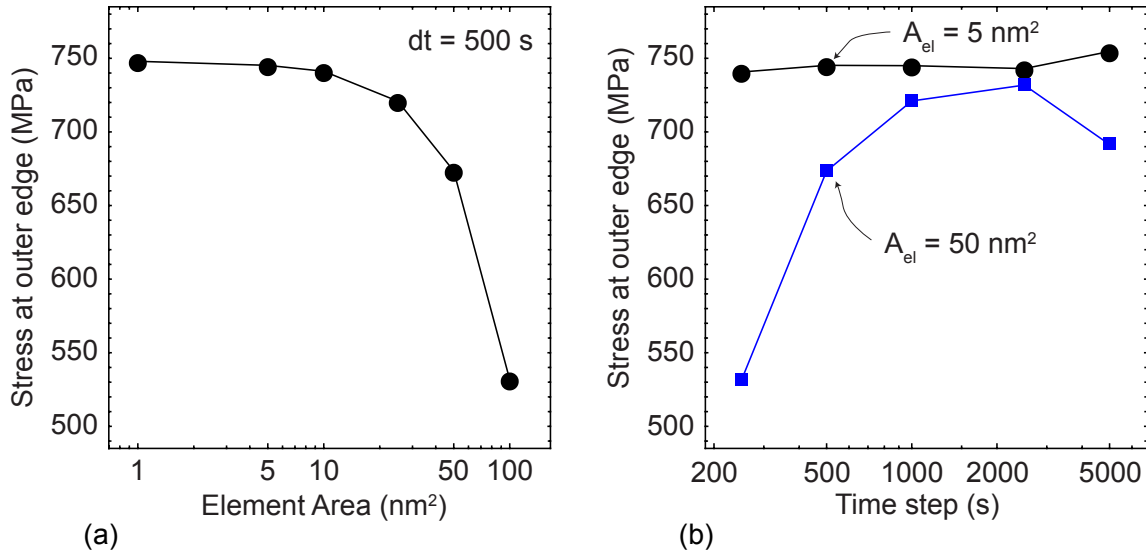


Figure 4.12: Plots of the tensile stress at the outer edge of the fiber as a function of (a) element area with a fixed time increment and (b) time increment with two different element sizes.

strains, while those that fall mostly inside the SiC will have no creep strains. Thus, the distribution of creep strain increments along the interface will have larger spatial variations than those surrounded by identical material.

This behavior is inherent to the present formulation and can only be alleviated by more sophisticated remeshing that is based on tracking of two interfaces, one at the previous time step and one based on the growth increment, such that the updated mesh includes only elements that fall entirely on one-side of the interface. It is interesting to note that the problem arises from the interface itself, not the physical size of the new oxide domain: taking larger time steps to advance the interface multiple elements does not avoid cutting those along the boundary, and has proven ineffective in mitigating spatial variations in creep strains and hence stress. It is worth noting that a complete parallel study on the use of higher order elements (e.g. six-noded triangles) only partially alleviates these behaviors, while leading to increases in computational costs.

That said, the spatial influence of this interface behavior is confined to thin slivers

on either side of the boundary that scale with the mesh resolution. This is illustrated in Figure 4.13, which depicts stress profiles near the domain boundary at various instances in time. The position is shifted such that $x = 0$ always falls on the interface. Ahead of the boundary, the stress distribution is relatively unimpacted by the cutting of elements along the interface. For a fixed element size, increasing the time step leads to far more noise near the interface. Holding the time step fixed and increasing the area increases the distance smoothing occurs and decreases the stress in the substrate near the interface. The practical implication of this is that stress quantities should be extracted from the simulation only at locations that are a fixed distance away from the interface, i.e. at a distance that spans at least several elements. While the physical extent of this “no go zone” near the interface can be limited to some extent by increasing mesh resolution, the concomitant adjustment in time-step (demanded in the diffusion analysis) mitigates the effectiveness due to dramatic increases in computational speed, as well as noise introduced during infinitesimally small evolutions of the oxide domain.

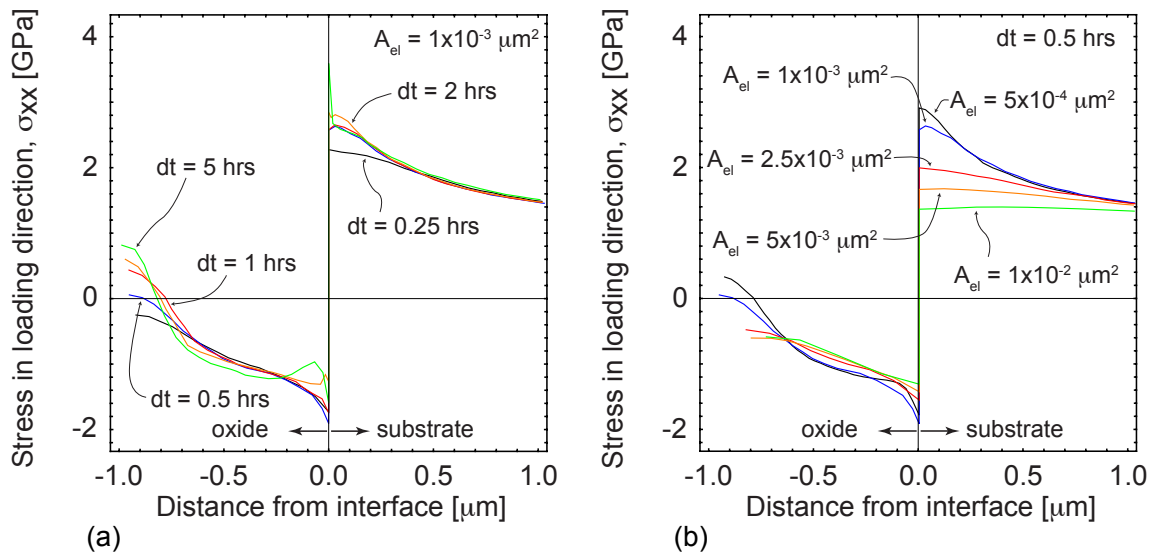


Figure 4.13: Plots of the stress profiles near the oxidation boundary for (a) a constant element size and various time steps and (b) constant time step with various element sizes.

In summary, the selection of time step and mesh resolution in the coupled framework involves quantifying competing effects that arise from module integration, which do not arise when the modules are used in isolation. These competing effects are associated with discretization limits along the interface itself. Figures 4.11-4.12 provide illustrations of the behaviors that guide selection of numerical parameters, as discussed further in Chapter 5.

4.7 Conclusion

The multiphysics framework described in this chapter combines oxidation growth with stress evolution in order to study complex oxidation problems. The framework can replicate one-dimensional oxidation and match the Deal-Grove model despite not enforcing a priori all of its underlying assumptions. Oxidation and stress results from the more complex case of a bare silicon carbide fiber exposed to dry air are consistent with previously published work, but differ due to different assumptions regarding stress state and relaxation. Convergence behaviors for the integrated modules are more nuanced than those arising from the use of isolated modules, due to coupling that occurs at the domain boundary. For a given case, a convergence study is needed to identify time increment and element size that balance computational speed, the impact of remeshing, and the presence of spatial variations along the boundary.

Chapter 5

Oxidation Under a Protective Coating

5.1 Introduction

Ceramic composites offer significant opportunities to allow higher operating temperatures than those currently enabled by coated metallic components, which translates into impactful performance gains for power generation and propulsion systems [1, 86, 87]. One of the most promising material systems consists of fine-grained SiC fibers (with high strength and creep resistance) embedded within a SiC matrix [1]. In many operating environments, the presence of oxygen and water vapor can lead to oxidation and volatilization of SiC, which must be managed to ensure component durability [1, 88]. Environmental barrier coatings (EBCs) are the prevailing strategy to protect CMC components from such chemical attack (see Figure 5.1), by slowing or eliminating the transport of reactive species to the underlying load-bearing structure [16, 58, 89, 90].

Rare earth silicates have emerged as a leading candidate for EBC systems on SiC composites, as they offer relatively good resistance to chemical attack and have coeffi-




Figure 5.1: Schematic illustration of an environmental barrier coating (EBC) on a ceramic matrix composite (CMC), with a close-up view of the tip of a crack in the coating. The crack serves as a fast-diffusion pathway that drives local oxidation, which generates high stresses that promote substrate cracking and/or coating delamination.

cients of thermal expansion (CTEs) that are relatively close to that of SiC [16,58,89,90]. Matching the CTE of the coating to that of the coating is critical, since misfit stresses (from deposition and/or exposure high temperatures) promote delamination of the coating and/or penetrating cracks. As shown schematically in Figure 5.1, penetrating cracks then serve as fast diffusion pathways for reactants that drive internal degradation, which may induce further cracking and progressive degradation mechanisms. For example, amorphous oxide typically forms at elevated temperature; even at high temperature,

the volume changes associated with oxide formation may drive cracking, which open new pathways for fast diffusion. In some operating conditions, crystallization and phase transformations in the oxide trigger additional dimensional changes that drive cracking, as shown in Figure 5.1 [90].

While the ideal solution is the development of fully-dense coatings that simply avoid cracking to maintain a hermetic seal, this either requires extremely thin coatings (which suffer from shorter lifetimes due to volatilization by water vapor) or reduced operating temperatures that arguably defeat the purpose of using ceramic systems. Assuming that some degree of cracking may be inevitable, the question naturally arises: what degree of cracking is acceptable? A qualitative answer is obvious: the degree that limits internal oxidation over the lifetime of the component and avoids premature component failure. Quantitative answers, however, require new insight regarding the connections between chemical transport, the evolution of interior reaction domains (e.g. newly formed oxides or cavities created by volatilization) and the resulting stresses that drive cracking. Such connections require an integrated modeling approach that simultaneously addresses species transport, evolution of domain boundaries driven by local reactions, and stress analysis (including the possibility of creep).

In this work, we utilize an integrated multiphysics simulation framework (shown schematically in Figure 5.2) that integrates these components to simulate a vertically cracked EBC bonded to a SiC substrate exposed to water vapor. Relevant parameters are chosen assuming oxygen is chosen as the reactant species, although it should be appreciated that water vapor under certain conditions causes much faster growth of SiO₂ (oxide) on exposed SiC surfaces [89, 91, 92]. The principal driver of oxidation is largely immaterial to the present contributions for reasons that will become clear; it is worth emphasizing the multiphysics simulation framework is materials agnostic and comparable simulations can be easily conducted for other compositions. Evolving oxide domains

are arguably the central culprit driving mechanical degradation of such systems, owing to the large molar volume change that generates large compressive stresses in the oxide. In and of themselves, even for planar oxide films (of uniform thickness), these stresses can generate significant strain energy that drive coating delamination [16]. For localized oxidation, where the oxide domain has an irregular shape, the large compressive stresses generated in the oxide can induce large tensile stresses in the adjacent material [14, 15]. While this behavior is examined in this chapter for cracked EBCs, it also has critical implications for oxidation of SiC fibers that bridge matrix cracks in CMCs. Simply put, the formation of localized oxide domains within a number ceramic systems plays a central role in degradation mechanisms that limit high temperature performance.

It should be emphasized that the mechanisms described above may be deleterious even at high temperature, albeit to an unknown degree (a point addressed in this chapter). This is in stark contrast to conventional wisdom in design of multilayers for high temperatures, which typically assumes that misfit stresses are relaxed at high temperature and are generated upon cooling, such that cracking occurs at low temperature. The origin of damage that may occur at high temperature is comparatively unexplored. Naturally, relaxation also occurs in the oxide at high temperature; hence, it is critical to include creep relaxation in the stress analysis to judge the competing effects of volume change and inelastic relaxation. Although the focus of this work is behavior at elevated temperature, cooling can trigger additional damage, notably phase transformations in crystalline silica (if present) that involve further large (contractile) volume changes. The present calculations for high temperature behavior can be viewed as the starting point for such studies by identifying the misfit strains that are retained at the onset of cooling. These misfit strains are critical to assessing cracking during cooling.

The idealized geometry considered here representing a cracked barrier coating on a planar substrate is shown at the bottom of Figure 5.1. The fast transport channel

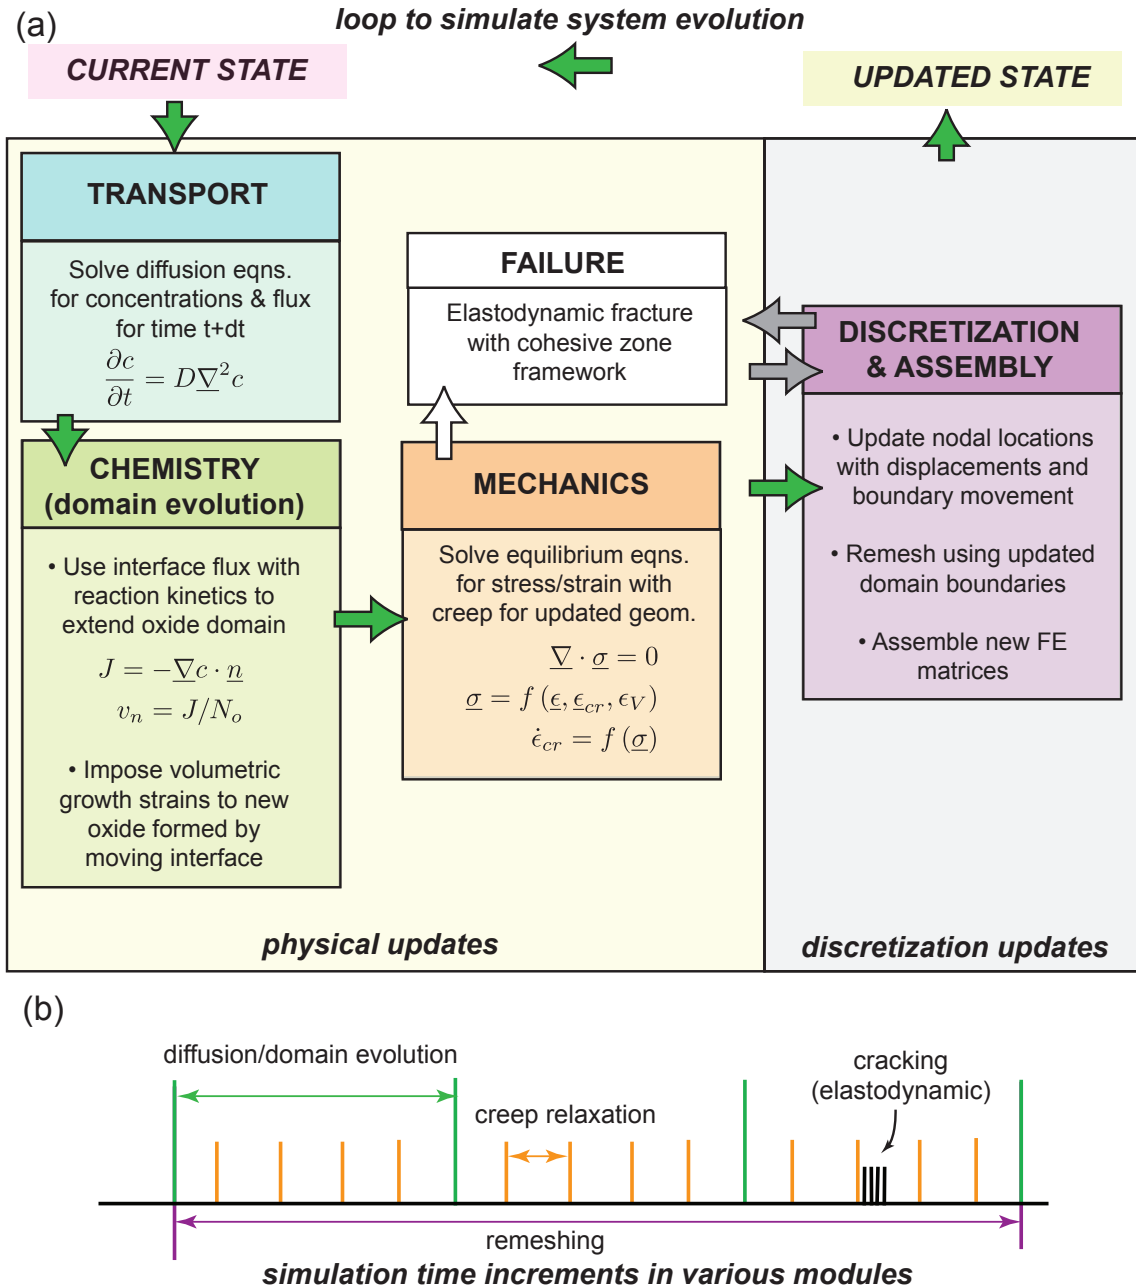


Figure 5.2: (a) Schematic diagram of the multi-physics framework that simulates chemical transport, domain evolution driven by oxidation, and creep relaxation in newly formed oxides. Adaptive remeshing is used to enable accurate simulations of large geometric changes. (b) Schematic diagram of different time scales that must be synchronized to integrate various physical processes. In the present work, numerical efficiency is sufficient to allow the use of a single time step for all processes.

representing the crack has a uniform, fixed width set as one of the parameters of the model. Coatings with larger CTE than SiC (e.g. ytterbium and yttrium monosilicates), deposited at temperatures above operating temperatures, are prone to cracking upon cooling. Upon reheating, the crack will remain open provided the effective deposition temperature is above the operating temperature. This picture is likely overly simplistic, since other mechanisms can intervene to ensure cracks remain fast diffusion pathways regardless of specific misfit strains, as crack waviness and small amounts of inelastic deformation may wedge the crack open. Hence, more sophisticated treatments of channel openings based on cracking phenomena are difficult to specify, since they will vary significantly with the choice of EBC material, the reference temperature associated with zero misfit stress (typically taken as an effective deposition temperature), and any inelastic relaxation in the coating or the adjacent material.

Here, the dominant role of the crack opening is simply to dictate the flux of reactant species at the location of the crack, which will be different to diffusion through the intact coating. Diffusion through the coating is allowed for two reasons: (i) many deposition methods result in porosity and micro-cracking that likely decreases the barrier performance of fully dense base materials [93–95], and (ii) some amount of porosity in the coating may be beneficial, as it reduces the strain energy in the system driving failure. Diffusion down the open channel is modeled using an effective continuous medium with zero elastic modulus and a diffusivity calculated from a series approximation based on gaseous molecular diffusion and Knudsen diffusion down a confined channel. As will be demonstrated, the most deleterious scenario corresponds to an isolated crack (with crack spacing much greater than the coating thickness) with a wide channel, i.e. large crack opening displacements. This creates large gradients in oxide thickness that set up significant tensile stresses in the substrate.

The remainder of this chapter is organized as follows; Section 5.2 reviews the clas-

sical Deal-Grove model for time-dependent oxide growth and generates the corresponding model for the case with a coating over the oxide. This provides context for the multiphysics simulation framework described in Section 5.3. (Additional details of this framework are provided in Chapter 4.) Section 5.4 presents results from parametric case studies of the cracked EBC problem, elucidating the evolution of the oxide domain at the tip of the crack and associated stresses. Section 5.5 discusses the implications of these results in the context of EBC failure mechanisms, and potential areas of future study.

5.2 Oxide Growth Under Barrier Coatings Based on One-Dimensional Diffusion

Diffusion through multilayers is a complex problem, due to differences arising from the atomic state of the diffusing species, material-dependent saturation concentrations, and permeability differences across material interfaces. An elegant and comprehensive model for transport through an EBC to form an underlying oxide has been provided by Sullivan fully describing such effects [13]. In this work, we take a simplified version of their approach, neglecting differences in saturation concentration and permeability between layers (which can produce concentration jumps across interfaces). The rationale for this is twofold; first, the simplified approach reduces the number of parameters in the simulations, many of which have not been rigorously established for many materials of interest. Second, while many enhancements can lead to important changes regarding the kinetics of oxide growth, those changes are unlikely to impact the interaction between local oxidation and stress evolution. It should be understood more sophisticated treatments of diffusion through multilayers may be required to rigorously establish quantitative agreement with oxide formation experiments.

The simplified model presented next (as in the work of Sullivan [13]) is an adaption of the classical Deal-Grove growth model, with modifications that address the presence of a coating on top of the oxide. Transport of the reactant is at steady-state, which implies a constant flux through the system. In the Deal-Grove model, which considers only transport through the growing oxide layer, the corresponding descriptions of flux are: (*i, surface*) $J = \beta(c_\infty - c_s)$, where J is the flux, β is a adsorption coefficient that depends on atmospheric conditions, c_∞ is the concentration of the reactant in the atmosphere, and c_s is the concentration of at the surface of the oxide, (*ii, film*) $J = D(c_s - c_i)/h$, where D is the diffusivity of the film, h is the film thickness, and c is the concentration at the reaction front, and (*iii, interface*) $J = k_s c_i$, where k_s is the interface reaction constant. Since all three fluxes are equal, c_s and c_i can be eliminated from the three equations to produce a single equation for the flux at the interface, $J(\beta, D, k_s, c_\infty, h)$. This defines the rate of change of film thickness as $\dot{h} = J/N_o$, where N_o is number of reaction product molecules per unit volume and the dot denotes the time derivative.

The Deal-Grove analysis is extended here to consider oxide growth as an interlayer between a barrier coating and substrate. Differences in saturation concentrations in the layers are neglected. Defining D_b as the diffusivity of the barrier coating, D_{ox} as the diffusivity of the oxide layer, and h_b as the thickness of the coating (with h retained as the thickness of the oxide), the model for a coating on top of the oxide that follows Deal-Grove assumptions produces the following:

$$J = N_o \dot{h}(t) = \frac{D_b c_\infty D_{ox} \beta k_s}{h_b D_{ox} \beta k_s + D_b D_{ox} \beta + D_b D_{ox} k_s + D_b \beta k_s h(t)} \quad (5.1)$$

In the limit that the adsorption coefficient (controlling the difference in concentration between the atmosphere and exposed surface) is large, this can be simplified to the

following:

$$N_o \dot{h}(t) = \frac{c_\infty D_b D_{ox} k_s}{D_b D_{ox} + D_{ox} h_b k_s + D_b k_s h(t)} \quad (5.2)$$

This can be re-casted in the classical form of the Deal-Grove growth law, using the following parameters:

$$\dot{h}(t) = \frac{B}{A(1 + \alpha) + 2h(t)} \quad (5.3)$$

where

$$A = \frac{2D_{ox}}{k_s}; \quad B = \frac{2D_{ox}c_\infty}{N_o}; \quad \alpha_b = \frac{k_s h_b}{D_b} \quad (5.4)$$

The parameters A and B are identical to those in the Deal-Grove model, with α_b as an additional dimensionless parameter that reflects the impact of the barrier coating overlaying the oxide. The solution to eqn. (3), for an initial oxide thickness defined by h_o , is:

$$h(t) = \frac{1}{2} \left[\sqrt{4Bt + [A(1 + \alpha) + 2h_o]^2} - A(1 + \alpha) \right] \quad (5.5)$$

where again, $\alpha = 0$ for an uncoated substrate. Increasing α leads to decreases in oxide thickness $h(t)$ at any given time. Thus, increasing α by increasing coating thickness h_b or decreasing its diffusivity leads to inhibited oxide growth under the EBC. Note that A has a dimensions of length (thickness); for $A \gg h_o$, where h_o is the initial thickness, the thickness scales as $h(t) \propto t$. At longer times, when $h(t) \ll A$, the thickness scales as $h(t) \propto \sqrt{t}$. The initial thickness h_o may represent a native oxide layer that forms prior to coating deposition.

The time to reach a critical oxide thickness under the EBC can be determined via eqn. (5); this is given by

$$t_{cr} = \frac{(h_{cr} - h_o)(h_{cr} + h_o + A(1 + \alpha))}{B} \approx \frac{h_{cr}^2}{B} + \frac{h_{cr}A(1 + \alpha)}{B} \quad (5.6)$$

where h_{cr} is the oxide thickness deemed critical. The approximate form corresponds to neglecting any initial (native) oxide. A more insightful form for the time to reach critical oxide thickness is:

$$t_{cr} = \frac{N_o h_{cr}}{2k_s C_\infty} \left(1 + \frac{k_s h_{cr}}{D_{ox}} + \frac{k_s h_b}{D_b} \right) \quad (5.7)$$

The first term in the parentheses corresponds to cases with small critical thickness *and* small coating thickness, where the critical time is reaction-controlled. The second two terms dominate cases with larger critical thickness, and reflect diffusion-limited growth controlled by either the coating or transport through the oxide itself. Note that, in eqn. (7), the reaction constant k_s cancels when the second two terms dominate: this is a reflection of diffusion-controlled growth.

For typical turbine operating conditions related to oxidation of SiC, the latter two terms are on the order of 10-1000; for example, the second term (reflecting diffusion through the oxide) is ~ 50 for a critical thickness of $2 \mu m$, while the third term (reflecting diffusion through the coating) is ~ 2500 for a low diffusivity (dense) coating that is $25 \mu m$ thick. Thus, for diffusion-limited growth under barrier coatings, the time to reach a critical oxide thickness should be well approximated by:

$$t_{cr} = \frac{N_o h_{cr}^2}{2D_{ox} C_\infty} \left(1 + \frac{h_b}{h_{cr}} \frac{D_{ox}}{D_b} \right) \quad (5.8)$$

This provides a simple dimensionless parameter to estimate the impact of including a

barrier coating.

The above analysis also provides some insight regarding the presence of a crack in the barrier coating. Assume transport occurs only in the direction of the crack and lateral diffusion is negligible, even in the oxide. With this approximation, the above barrier analysis holds, only with the parameter $\alpha_c = k_s h_b / D_c$, where D_c is the effective diffusivity of the crack. (That is, the crack behaves as barrier layer with enhanced diffusivity.) Diffusion down the narrow crack can be treated using a combination of molecular and Knudsen diffusion, defined by:

$$D_c = \frac{D_m D_K}{D_m + D_K} \quad (5.9)$$

where the Knudsen diffusivity, D_K and the molecular diffusivity, D_m , are given by:

$$D_K = \frac{2\delta}{3} \sqrt{\frac{8RT}{M}}; \quad D_m = \frac{0.0018583 T^{3/2}}{P\sigma^2\Omega_{AB}} \left(\frac{1}{M_{air}} + \frac{1}{M} \right)^{1/2} \quad (5.10)$$

with δ as the crack opening, R is the ideal gas constant, T is the temperature, M_{air} is the molecular weight of the gas within the diffusion channel, M is the molecular weight of the diffusing species, σ is the average collision diameter, Ω is the collision integral, and P is the pressure. For oxygen at 1000°C and a crack opening of 1 μm , the diffusivity down the crack is within the diffusivity of free space for openings of 10 μm or greater; it drops to about one-half of the molecular diffusivity for an opening of 1 μm and 10% of the free space diffusivity for openings of 0.05 μm .

Figure 5.3a shows the predicted diffusivity in the crack as a function of crack opening, assumed to be constant along the crack. (The validity and implications of this assumption is discussed in Section 5.5 of this chapter.) Over a range of relevant conditions, the crack opening at 1 μm is within a factor of three of the molecular diffusivity, indicating the

crack does not inhibit diffusion very much relative to open atmosphere. Crack openings less than $1\ \mu\text{m}$ can have a dramatic effect; however, even at such small openings, the effective diffusion constant is many orders of magnitude faster than those associated with EBCs. Figure 5.3b shows time-dependent oxide thickness under an intact coating for several plausible coating diffusivities [16,96]. For a $25\ \mu\text{m}$ thick coating, any coating with a diffusivity above $O[10^{-6}\text{cm}^2/\text{s}]$ will exhibit the same oxide growth as a bare surface. With this in mind, and in light of the values shown in Figure 5.3a for the crack, the oxide thickness at the tip of the crack – neglecting any lateral diffusion – should be very close to that observed for an uncoated substrate. In essence the characteristic time controlling transport down the crack is much smaller than that controlling transport through the coating, such that diffusivity in the crack plays a negligible role. This has been verified using simulations described in Section 5.4.

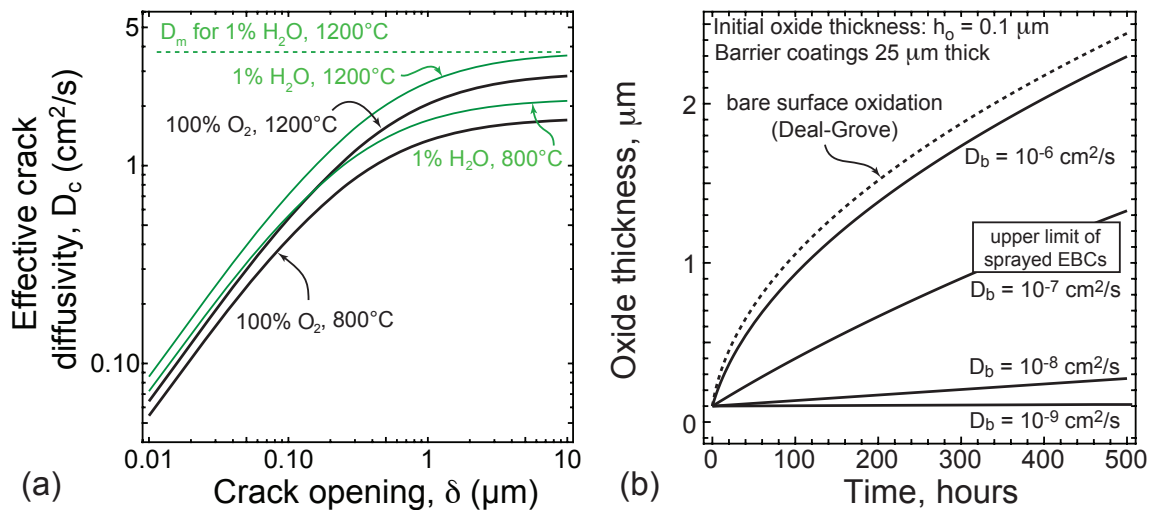


Figure 5.3: (a) Predicted diffusion constants down a narrow channel that combine Knudsen and molecular diffusion mechanisms, for 1% water vapor and 100% molecular oxygen at two temperatures. (b) Prediction of oxide thickness under an intact barrier coating for several coating diffusivities; for $D_b > 10^{-5}\text{cm}^2/\text{s}$, oxide growth is virtually identical to a bare surface, and controlled by diffusivity in the oxide layer itself. This implies reactant transport down even the most narrow of crack openings is likely never a limiting time scale for the growth of the oxide at the tip of a crack.

Since the diffusivity in free space is so much larger than that of the oxide or coating,

the growth of the oxide will be much larger directly under the crack, as compared to a remote location with an intact barrier. As will be demonstrated, this local penetration of the oxide produces local stresses that are much higher than those in fully intact systems. An approximate measure of the extent of local penetration is obtained using eqn. (5); consider the limit where diffusion-controlled growth is operative. Neglecting any initial oxide thickness, the ratio of oxide thickness under the crack to that under the intact barrier is given by:

$$\frac{h_{ox}^{cr}}{h_{ox}^b} = \frac{\sqrt{\tau + (1 + \alpha_c)^2} - (1 + \alpha_c)}{\sqrt{\tau + (1 + \alpha_b)^2} - (1 + \alpha_b)} \quad (5.11)$$

where $\tau = k_s c_\infty t / N_o$ is the dimensionless measure of time. Note that at sufficiently large times, the thickness of the oxide under the crack will be identical to that under the intact barrier coating; in this limit the oxide has thickened to the point that diffusion through the oxide itself is the controlling factor. This limit will correspond to extremely thick oxides that are deleterious for many reasons, and as such is not of significant practical interest to the present study.

5.3 Multiphysics Simulation Framework

A schematic overview of the multiphysics finite element framework utilized in this work is shown in Figure 5.2a. The framework combines three different physics modules: (i) the analysis of transport to determine the spatial distribution of reactant species that enter through the top of the coating and diffuse through the coating, down the crack and through the oxide layer, (ii) the prediction of the evolving oxide domain shape based on flux to the oxide/substrate interface, and (iii) the analysis of stresses that evolve through the system, assuming the coating and the substrate are elastic and the oxide

domain creeps according to a non-linear viscoelasticity description. These three physics modules are combined with a discretization module that performs on-the-fly remeshing that facilitates accurate prediction of stresses near the moving boundary that defines oxide growth. The colored modules in Figure 5.2a have been integrated with previous cohesive element frameworks for simulating crack evolution; however, to limit the scope of this chapter, related simulations involving oxide-driven fracture is left for future work. In this section, we summarize the essential features of the modules in Figure 5.2a used in the present simulations and key aspects of their integration. Complete details of the mathematical descriptions and numerical implementation were provided in Chapter 4.

5.3.1 Overview of Material Descriptions and Oxide Growth Law

The transport module in Figure 5.2a solves the standard linear diffusion equation $\partial c/\partial t = D\nabla^2 c$, where c is the concentration of reactant, and D is the diffusivity; the diffusivity is assumed to be isotropic but different in each domain, i.e. the barrier coating (D_b), the oxide (D_{ox}), and the crack (D_c). The diffusion equation is solved using a conventional finite element approach, with first-order time integration. The crack opening is modeled as a fixed-width channel that is explicitly meshed to obtain two-dimensional concentration distributions governed by D_c . Diffusion is assumed to be perfect between material domains; the concentrations are continuous across material interfaces.

Once the transport module has solved for the concentration profile, the chemistry module takes over to compute the oxide growth. The velocity of the oxide growth front is controlled by the flux normal to the interface. The oxide is advanced by marching along the existing interface and determining the outward flux normal to the interface at all points. The Deal-Grove model for one dimension assumes the rate of change in

thickness in oxide is

$$\dot{h}(t) = \frac{J}{N_o} \quad (5.12)$$

where J is the flux of oxidant molecules reaching the reaction front and N_o is the number of oxidant molecules consumed to create a unit volume of oxide. While this relation is for one dimensional, unconstrained growth, it is useful for developing a general oxide growth rule. In general, the velocity of the oxide front v_n is

$$v_n = \gamma \frac{J}{N_o} \quad (5.13)$$

with the proportionality constant a function of the oxide's Poisson's ratio and the ratio of the molar volumes of oxide and silicon carbide, $\gamma = f(\nu, \Omega)$. The proportionality constant takes into account the volumetric expansion that occurs concurrent with oxidation as well as the plane strain approximation used within the framework.

In addition to the molar volume of silica being significantly larger than silicon carbide, silica exhibits non-linear viscoelastic relaxation in shear. An Eyring model is used to approximate this relaxation. The Eyring model states that the shear strain rate is

$$\dot{\tau} = -\frac{E\tau}{2(1+\nu)\eta(\tau)} \quad (5.14)$$

with an elastic modulus of E , Poisson's ratio of ν , effective shear stress τ and stress dependent viscosity of

$$\eta = \eta_o \frac{\tau/\tau_c}{\sinh(\tau/\tau_c)} \quad (5.15)$$

for a reference viscosity of η_o and critical shear stress τ_c . The critical shear stress is slightly dependent on temperature but is taken to be constant of 100 MPa. The effective

shear stress is

$$\tau = \left[\frac{1}{2} s_{ij} s_{ij} \right]^{1/2} \quad (5.16)$$

where s_{ij} are the deviatoric stresses. The equivalent shear stress at any time t after an initial stress of τ_o is

$$\tau = 2\tau_c \tanh^{-1} \left(\tanh \left(\frac{\tau_o}{2\tau_c} \right) \exp \left(-\frac{Et}{2\eta_o(1+\nu)} \right) \right) \quad (5.17)$$

The stress components σ_{11} , σ_{22} and σ_{12} are assumed to relax proportionally to τ/τ_o . Creep is only allowed to occur within the oxide domain; all other materials are assumed to be completely elastic.

The mechanics module in Figure 5.2a solves for the increment in displacement over the time step. Similar to the transport module, the mechanics module utilizes a finite element approach but in a multistep solution technique. Similar to the work of Hay [14], the displacements are solved for twice: once to determine the state at the beginning of the time step immediately after oxide growth and a second time to find the displacements at the end of the time step after relaxation occurs. The strain state is a combination of the creep strain and deformation history and is applied through equivalent nodal forces. Between the two solution steps, the creep strain is incremented based on the stress state at the beginning of the time step and the relaxation law. In other words, the stresses are determined immediately after oxide growth, they are relaxed, the change in creep strain is computed and the displacements are found with this relaxation.

5.3.2 Module Integration and Synchronization

The reactant concentrations and associated fluxes outputted from the transport module are used as inputs in the chemistry module, which evolves the boundary of the oxide

domain according to Deal-Grove growth models described in the previous section. That is, the oxide/substrate interface is translated in space to simulate oxide growth according to the condition $v_n = k_s \underline{\nabla} c \cdot \underline{n}$, where v_n is the velocity of the growing interface normal to the interface, k_s is the reaction constant, and \underline{n} is the surface normal of the interface. This corresponds to a flux boundary condition; all reactant molecules that arrive on the interface are consumed in the oxidation process to generate a new layer oxide. As such, the diffusivity of the layer below the oxidation front is immaterial.

After the boundary advances (i.e. a new layer oxide is created), the geometry is passed to the mechanics module; this module imposes growth strains in the new oxide film (generated by the motion in the chemistry module) and solves for creep-mediated stress relaxation. At any point in time, the output from the transport and mechanics modules can be passed to a discretization module that remeshes all domains to ensure the mesh is appropriate for a given domain size and shape. This module maps field variables from the old mesh to the new mesh using techniques described in Chapter 4.

In general, the characteristic time scales of physical processes addressed by the framework can be very different, as illustrated in Figure 5.2b. Since the current implementation invokes the assumption that flux to the oxidation front controls domain evolution, the time constants for transport and boundary motion are the same. Creep relaxation in the oxide, however, occurs over much shorter time scales, since the oxide near the boundary experiences extremely high stress due to large volume changes generated by oxide formation. That is, a typical time for relaxation is on the order of seconds, as opposed to the minutes or hours that may be needed for meaningful oxide growth. (As illustrated in Figure 5.2b, cracking may involve even smaller time steps to capture elastodynamic crack advance.)

Conceptually, one can use different time steps in each module provided there is synchronization between the global times marking the beginning and end of each pass

through entire framework. (E.g., one can use smaller time steps to simulate rapid creep relaxation provide the total time increment of all the steps combined matches that used in the transport module.) Remeshing can be done with arbitrary frequency throughout the simulation, since this simply involves mapping one spatial discretization at fixed time to another. In the present work, the synchronization between modules is rendered moot by using the same time increment in each module; transport, domain evolution and stress relaxation are computed using identical time increments. The disparity between creep and oxidation time scales is handled by developing a non-linear analytical solution to accurately predict relaxation over large time scales. Simply put, a higher order update to creep relaxation is used to ensure accurate stresses are obtained even over large time steps.

5.3.3 Overview of Numerical Implementation and Convergence

All three modules utilize the same mesh of triangular elements with three nodes; both concentration and displacements are interpolated linearly in each element. These low order elements were used for several reasons. First, there are existing open-source codes for automatic meshing of these elements that have been highly optimized and allow for extremely rapid mesh generation. This is a critical advantage for problems that simulate large motions of boundaries that generate strong gradients. In such cases, refined meshes are needed to capture these gradients regardless of its position; using a fixed mesh that was not tied to domain evolution would require exceedingly fine meshes throughout the entire domain. (This is true even for more elegant boundary tracking approaches such as level set or phase-field methods: while boundary stability is more favorable in such methods, they do not obviate the need for a refined mesh near the boundary, unless strong gradients are not present.) Second, constant strain triangles are generally

preferred for cases involve highly nonlinear material response and strong gradients, as higher order elements can be prone to convergence issues. Third, the advantages are currently exploited in highly parallelized codes with distributed cohesive elements to simulate cracking, such that the transport and stress analysis framework can be used to generate inputs to additional modules for simulating damage.

Motion of the oxide growth front is tracked as follows; for simplicity, the process is explained here for growth in one dimension. At the start of each increment (defined by time t) the interface is defined by a node with an oxide element on one side and substrate element on the other. At the end of the increment (time $t + dt$) the front is advanced by the distance $\Delta x = v_n dt$. The concentration distribution is used to calculate v_n as previously described. This converts a region from SiC to SiO₂ and is not tied to element size; the converted region may be a small fraction of an element or multiple elements. The conversion corresponds to assigning oxide properties and the volumetric strain associated with the conversion to the region for the following time step.

The elemental matrices and nodal force vectors for the mechanics module are formulated with the oxide domain as it exists at the beginning of the time increment, i.e. prior to any growth. The initial solution in the mechanics module is found to determine the state at time t immediately after the oxide grows. After relaxation occurs, the nodal force vectors are reformulated with the additional creep strain that occurs over the time increment dt . The displacements are then calculated for time $t + dt$, but before any additional oxide growth is allowed.

Prior to the next increment in time, the positions of the nodes along the interface are updated according to both the boundary movement and any mechanical displacements. This updated collection of points along the interface are then used to define the growth boundary at time $t + dt$ (by interpolating between the nodal positions obtained at the end of the step). A new mesh is automatically created using the updated boundary; this

can be generated so that the average element size along the boundary remains fixed even if the line length of the boundary increases due to growth. Field variables (such as creep strains) from the old mesh are mapped onto the new mesh, finalizing the transition from a solution at time t to the new time $t + dt$.

The convergence of the transport and mechanics modules was verified individually with respect to both spatial discretization and time stepping. This included establishing the convergent behavior of the non-linear analytical solution used to predict relaxation over relatively large time intervals. The chemistry module that predicts motion of the oxidation front was verified using Deal-Grove analytical solutions for both planar and cylindrical surfaces. Using the transport module to predict the spatial distribution of concentration, the evolution of oxide domain was proven to be independent of mesh density and time step and in complete agreement with analytical solutions. The coupling between the oxidation front and the mechanics module (i.e. creep) was also shown in Chapter 4 to be independent of mesh density and time step, and shown to reproduce the spatio-temporal distribution of stress obtained elsewhere. Finally, the convergence of the re-meshing algorithm was verified by altering the re-meshing frequency in relation to the time step (i.e. re-meshing every time step or multiples of the time step) for various mesh densities.

5.4 Oxide Growth & Stress Evolution Under Cracked Films

The multiphysics framework described in the preceding section was used to the idealized problem shown in Figure 5.4. In all cases, the coating thickness is $25 \mu m$ and the crack opening is modeled as a narrow channel with fixed width of $1 \mu m$. The model

comprises one period of an array of cracks separated by the distance s . (For efficiency, only one half of the period shown in Figure 5.4 is discretized with appropriate symmetry conditions; depictions of the actual geometry obfuscate the model by compressing the image.) A thin initial oxide layer with uniform thickness of $h_o = 0.1 \mu\text{m}$ is included to start the calculation.

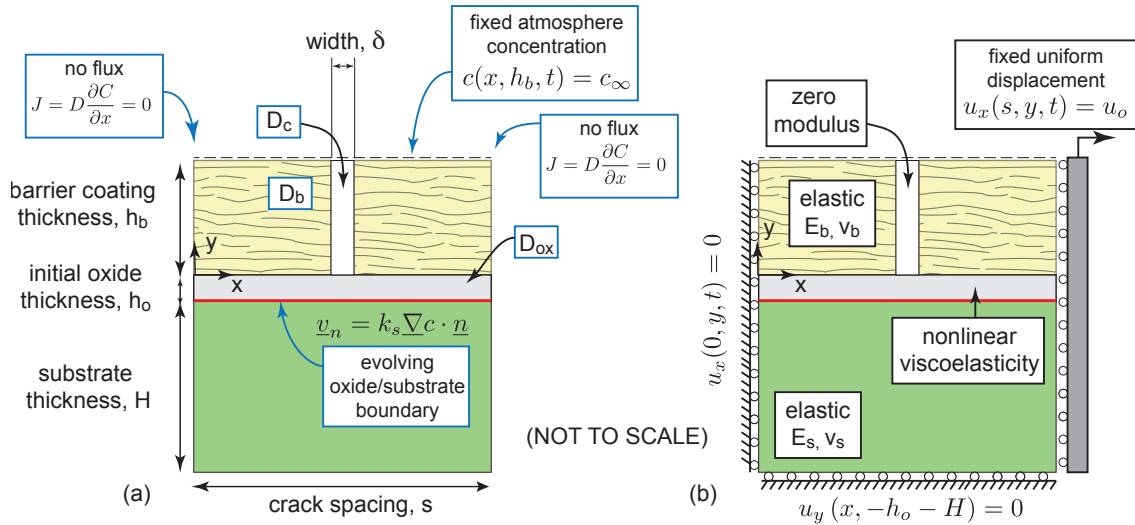


Figure 5.4: (a) Schematic illustration of the idealized geometry used in the current simulations, with transport properties and boundary conditions. (b) Schematic illustration of the mechanical properties and boundary conditions used to predict stress evolution in the system.

Only the bottom interface shown in red in Figure 5.4, between the initial oxide layer and the substrate, evolves during the simulation due to oxidation of the underlying substrate. The channel representing the crack is modeled as a continuous medium with zero elastic modulus and a diffusivity calculated from the width of the channel, as described in Section 5.2. A fixed concentration of reactant is applied to the top surface, with zero flux conditions on the vertical boundaries as dictated by symmetry. Table 5.1 provides a summary of the range of properties considered in this work.

The idealizations of the geometry shown in Figure 5.4 are largely motivated by the

the fact they limit computational times and enable a broad parameter study of coupling between transport, reactions and stress. The physical implications of several obvious approximations in this approach are discussed in the next section, following illustrations of the resulting behavior.

| | Coating | Channel | Oxide, SiO ₂ | Substrate, SiC |
|-------------------------|---------------------------|------------------------------|--|----------------|
| Thickness, h | $25 \mu m$ | $\delta = 1 \mu m$ | $h_o = 0.1 \mu m$ | $500 \mu m$ |
| Elasticity, E, ν | 180 GPa, 0.27 | 0, 0 | 70 GPa, 0.17 | 400 GPa, 0.35 |
| Creep, η_o, τ_c | n/a | n/a | $\eta_o = 620 \times 10^4 \text{ GPa} \cdot s$ $\tau_c = 100 \text{ MPa}$ | n/a |
| Diffusivity, D | 0, 0.1, 1, 10 $\mu m^2/s$ | $1.72 \times 10^8 \mu m^2/s$ | $0.73 \mu m^2/s$ | n/a |
| Channel spacing, s | n/a | 25, 50, 100, 200 μm | n/a | n/a |

Table 5.1: A summary of properties used in the simulations; the channel represents the crack, while the oxide thickness refers to the initial uniform layer thickness at the start of the simulation. Creep is only allowed in the oxide. The oxide growth law dictates that all reactants are consumed at the interface, such that diffusivity in the substrate (SiC) is irrelevant. Properties are computed at $T=958^\circ\text{C}$ assuming oxygen as the reactant; the associated reaction rate constant is $k_s = 10 \mu m/s$. The concentration held at the top of the coating is $c_\infty = 5.5 \times 10^4 / \mu m^3$ which corresponds to the solubility limit of the oxide.

The spatial distributions of concentration and stress are shown in Figures 5.5-5.8, at three different times and for three different coating diffusivity values. In all cases shown, the remote applied stress is $\sigma_a = 1 \text{ GPa}$. There are several features of these results that are characteristic of system behavior across the entire range of the parameter space considered here.

To begin, the distribution of reactant and resulting shape of the oxide domain is dominated by the diffusivity of the coating (Figure 5.5). At low coating diffusivity, the oxide directly under the crack is localized to the location of the crack, and adopts a semi-cylindrical shape arising from radial diffusion away from the crack tip. The boundary of the oxide domain is shown by the contour with zero concentration, since all reactant is consumed at the growing interface. (The boundaries are more crisply visible in the stress plots.) At high coating diffusivity, transport through the coating reduces this effect and leads to more uniform oxide thickness. Figure 5.6 illustrates that thick oxide regions




Figure 5.5: Contours of concentration for at three different times, for three different coating diffusivities and relatively large crack spacing (low crack density); low diffusivity in the coating leads to local oxide thickness under the crack that is much greater than that under the intact portion of the coating.




Figure 5.6: Contours of direct stress parallel to the coating interface at three different times for three different coating diffusivities; increasing the coating diffusivity leads to smaller differences between the crack and intact regions of the coating, and smaller tensile stresses in the substrate. The system is subjected to a fixed strain in the x-direction that generates $\sigma_{xx} = 1 \text{ GPa}$ at the remote boundary.




Figure 5.7: Contours of direct stress perpendicular to the coating interface at three different times for three different coating diffusivities; increasing the coating diffusivity leads to smaller differences between the crack and intact regions of the coating, and smaller tensile stresses in the substrate. The system is subjected to a fixed strain in the x -direction that generates $\sigma_{xx} = 1 \text{ GPa}$ at the remote boundary.




Figure 5.8: Contours of shear stress at three different times for three different coating diffusivities; increasing the coating diffusivity leads to smaller differences between the crack and intact regions of the coating, and smaller tensile stresses in the substrate. The system is subjected to a fixed strain in the x-direction that generates $\sigma_{xx} = 1 \text{ GPa}$ at the remote boundary.

localized under the channel produce locally high stresses in the substrate; creep in the oxide layer mitigates the stress in the newly formed oxide. With low stress in the oxide, the local oxide domain under the crack effectively serves as a notch in the underlying substrate, with stress concentrations proportional to the depth of the oxide under the crack relative to that under the intact coating.

The stresses normal to the coating/substrate interface are significantly lower, as seen in Figure 5.7; this is a consequence of the absence of remote loading in this direction. (Indeed, without any remote loading, *all* stresses are significantly lower, due to creep relaxation in the oxide.) At the end of the coating next to the channel, the interface stresses are compressive due to the fact that the oxide expands upwards and pushes on the edge of the coating. This compressive interface stress is balanced by tensions that develop along the interface to the side of the channel. This likely has important implications for coating delamination, which may initiate at high temperature. The shear stresses shown in Figure 5.8 are similar in magnitude to the direct stresses acting perpendicular to the loading direction.

Since the stresses are governed by the relative depth of the oxide near the crack (as compared to that under the intact coating), it is worth examining the nature of oxide growth prior to considering the evolution of stress in the system. Figure 5.9 illustrates the time-dependence of the oxide thickness directly under the crack, and under the intact coating (at the symmetry plane between cracks). Two features have important implications.

First, the depth of the oxide directly under the crack is close that expected for an uncoated substrate, regardless of the diffusivity of the coating. This is a consequence of the behavior shown in Figure 5.3, which demonstrates diffusion down the crack is orders of magnitude faster than diffusion through either the coatings or the oxide itself. The thickness under the crack is somewhat smaller than for a bare surface because lateral




Figure 5.9: Oxide thickness directly under the crack and under the intact coating (at the symmetry plane between cracks) for several different coating diffusivities. The oxide at the crack location is smaller than bare surface oxidation due to lateral diffusion away from the crack plane.

diffusion detracts from the amount of reactant reaching the crack front at the symmetry plane. I.e., it is the difference between one-dimensional diffusion in a uniform system and two-dimensional diffusion from a point source. This is further illustrated in Figure 5.10, which depicts the ratio of the oxide depth under the crack to that of one-dimensional diffusion. The difference becomes larger as time progresses because the radial diffusion path length becomes larger. It is interesting to note from Figure 5.10 that the crack spacing is largely immaterial, indicating the cracks act as isolated diffusion sources.

Second, the difference between crack growth under the crack and under the intact coating is a strong function of coating diffusivity; larger coating diffusivity is closer to oxidation of an uncoated surface and hence the thickness difference is smaller. This has important implications for stress, since flatter oxide fronts (more uniform oxide thickness) are associated with smaller stress concentrations arising from shape. From Figure 5.10, it is clear that irrespective of crack density, there is a critical change in behavior at a coating diffusivity of approximately $D_b = 1 \mu\text{m}^2/\text{s}$; for higher values, the oxide thickness

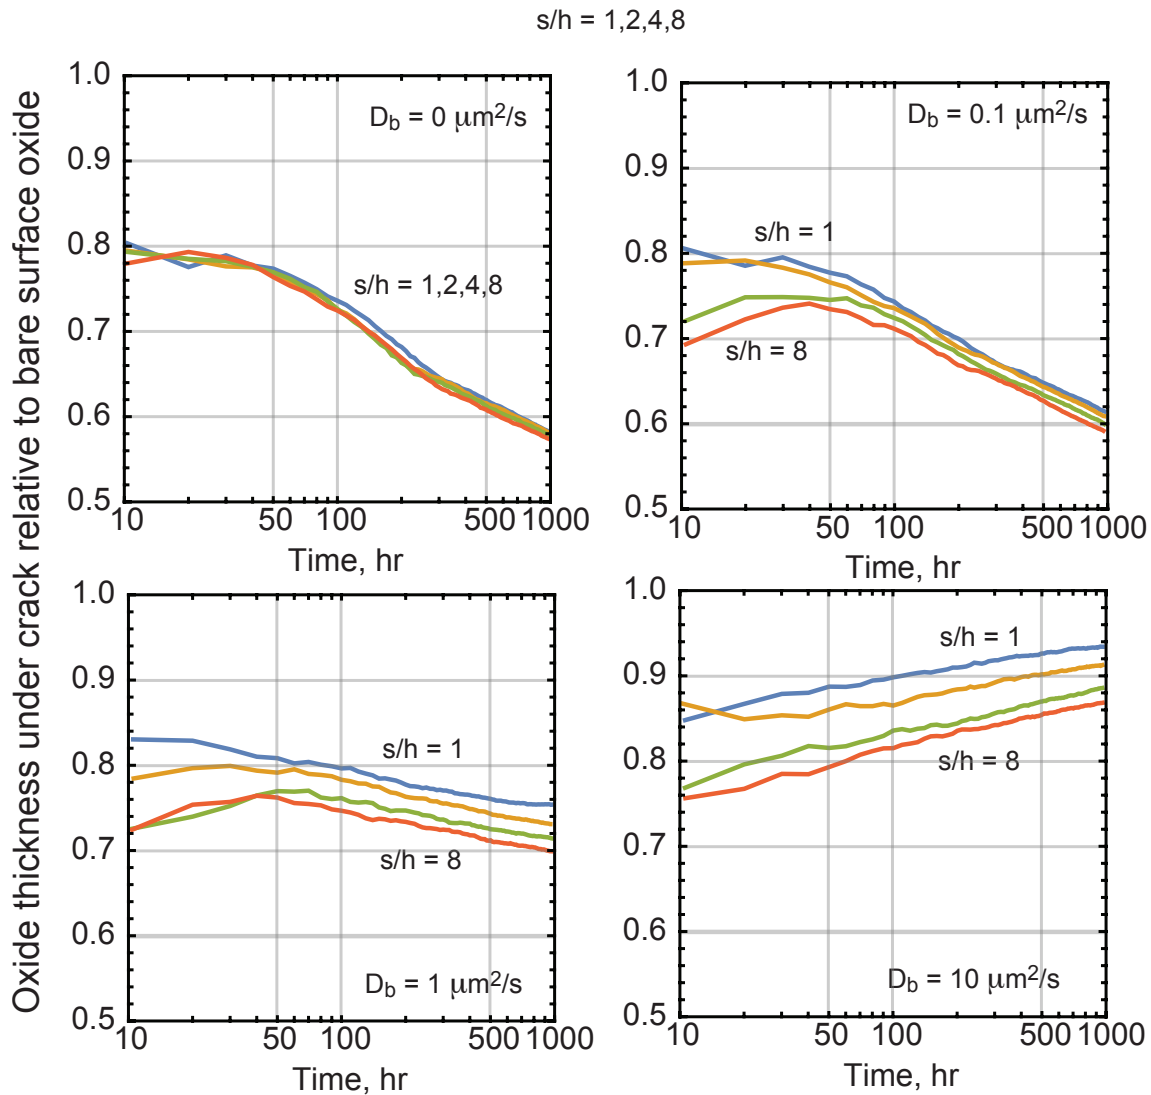


Figure 5.10: Oxide thickness directly under the crack, normalized by the thickness predicted by Deal-Grove for an uncoated substrate; four crack spacings are shown for four different coating diffusivity values. The decay with time in the relative thickness arises from two-dimensional diffusion away from the crack tip; for high coating diffusivity, two-dimensional diffusion is effectively eliminated by fast transport through the coating itself. Crack spacings greater than at least one coating thickness play a relatively minor role.

in reasonably close to that associated with an uncoated surface. It is worth noting that the reported range for candidate EBC materials falls in the range $D_b = 0.1 - 10 \mu\text{m}^2/\text{s}$, with the upper end associated with porosity arising from processing. Naturally, this

critical threshold depends on the coating thickness, which is discussed further in the next section.

The implications of these differences in oxide thickness (under the crack versus under the coating) with regards to stress evolution are shown in Figures 5.11-5.13. The focus of these figures is on the tensile stress that develops in the substrate directly under the thickness part of the oxide domain. Unlike the evolution of oxide geometry, crack spacing and the level of remote stress play an important role.

Figure 5.11 plots the stress in the loading direction (parallel to the coating/substrate interface) at several instances in time, for three different coating diffusivity values. There are two central features of the stress distribution in all simulations. *(i)* The stress in the newest oxide, adjacent to the oxidation front, are highly compressive; creep, however, dramatically mitigates the magnitude of these stresses. The $O[20\text{ GPa}]$ stress that arises from volumetric expansion is relaxed over time scales on the order of seconds, which is far smaller than the propagation of the front. *(ii)* A large tensile stress develops in the substrate, immediately adjacent to the oxide domain, as required to balance the compressive stresses in the oxide and maintain equilibrium along this plane. Arguably, the critical aspect of behavior shown in the results in Figure 5.11 is that the peak stress in the substrate reaches a peak at intermediate times when the coating diffusion is low, indicating that oxide growth controls mechanical response. However, for higher coating diffusivity, the highest stress occurs is at the onset of the simulation (i.e. from the stress concentration of the channel), such that any deleterious effects of oxide growth is offset by creep.

Figure 5.12 illustrates the time-dependence of the peak stress computed in the substrate. The noise seen in Figures 5.12a and b are a consequence of averaging element stresses from the constant strain triangles; this noise can be mitigated by higher resolution studies that come at great computational expense but do little other than reduce

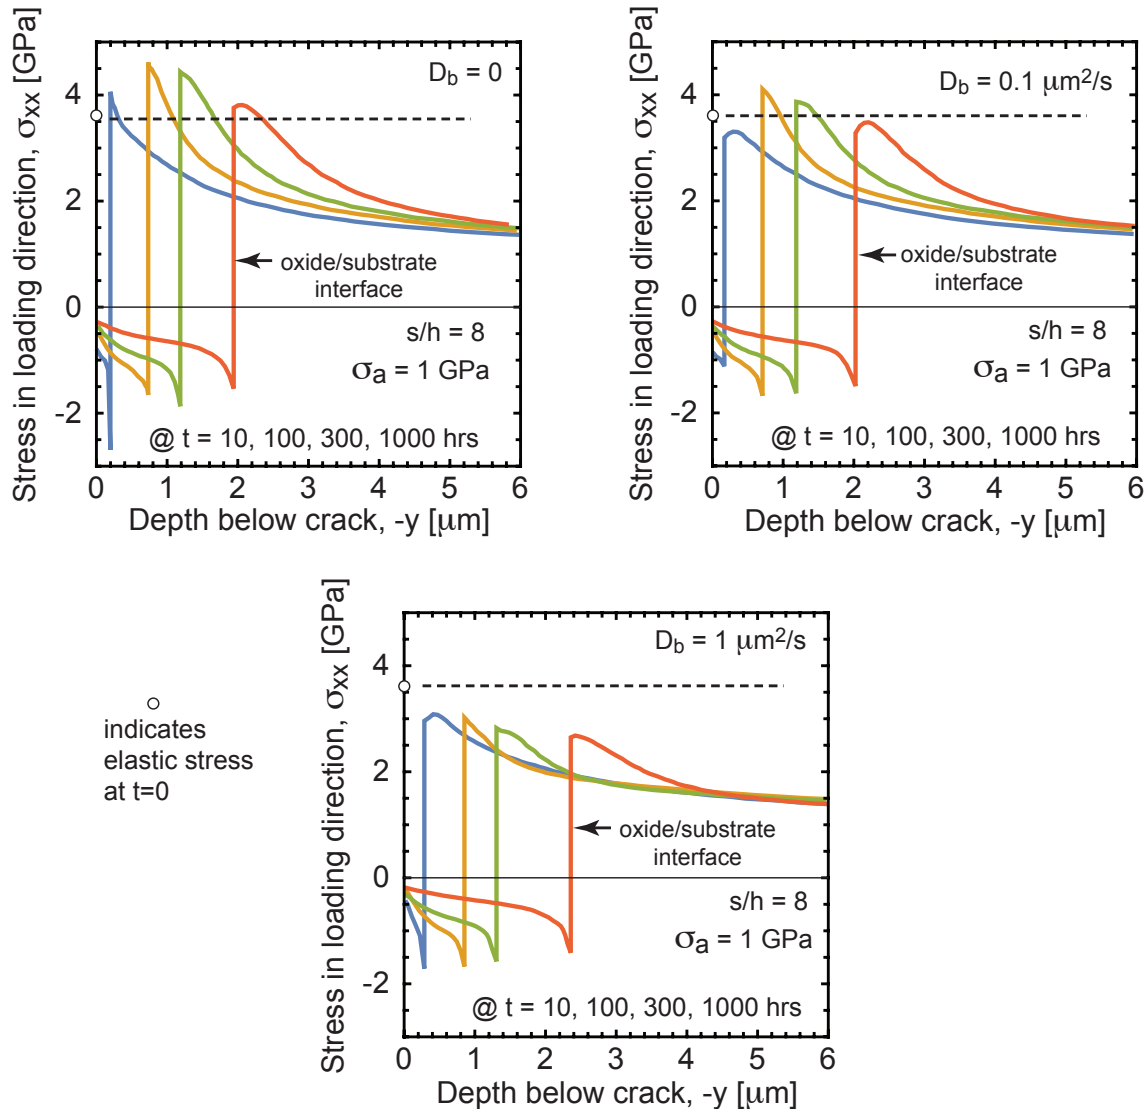


Figure 5.11: Direct stress distribution in the loading direction at the plane of the crack, for several times and several different coating diffusivity values. Higher coating diffusivity reduces the gradient in oxide thickness from the crack plane to the intact region of the coating, and mitigates stress.

the noise. The open circles in Figures 5.12a and b reflect the values at the start of the simulation, and are controlled by the stress singularity at the corner of the channel and the initial oxide thickness. At very short times, before significant oxide growth, the maximum stress in the substrate falls due to creep relaxation in the initial oxide that is present, as illustrated by the inset to Figure 5.12a. Once the oxide thickness under

the crack increases by a factor of 3-5, oxide growth becomes the controlling factor and stresses may again rise. Whether or not stress in the substrate rises or falls after the initial loading is controlled by the shape of the oxide front (i.e. the difference in oxide thickness between the crack and symmetry plane), which is in turn controlled by coating diffusivity.

Note from Figure 5.12b and Figure 5.12c the crack spacing and remote loading have an important impact on the stress behavior, unlike their negligible role in the formation of the oxide. This is a consequence of the fact that stresses induced by the presence of the crack are significant over a length-scale controlled by the coating thickness, as opposed to the oxide thickness. (That is, at higher crack density, the initial stresses in the substrate at the bottom of the oxide are lower.) The stress concentration is slightly lower at high levels of remote loading, presumably due to the reduction in misfit strain arising from oxide growth.

The effect of crack spacing and remote loading is illustrated in Figure 5.13, which tabulates stress concentration factors as a function of crack spacing, applied loads and diffusivity in the coating. The key implication of these results is that for higher levels of coating diffusivity, the accelerated growth of the oxide under the crack is not sufficient to elevate stresses beyond those experienced at the outset of the simulation. At high levels of diffusivity in the coating, the initial stresses due to the presence of the crack, and any elevation due to differential oxide growth, are completely mitigated by creep in the oxide.

5.5 Discussion

The results clearly illustrate that the largest driver of high stress in the substrate is the geometry change associated with enhanced oxide thickness directly under the crack. Geometry changes are controlled by diffusion differences between the location of the crack

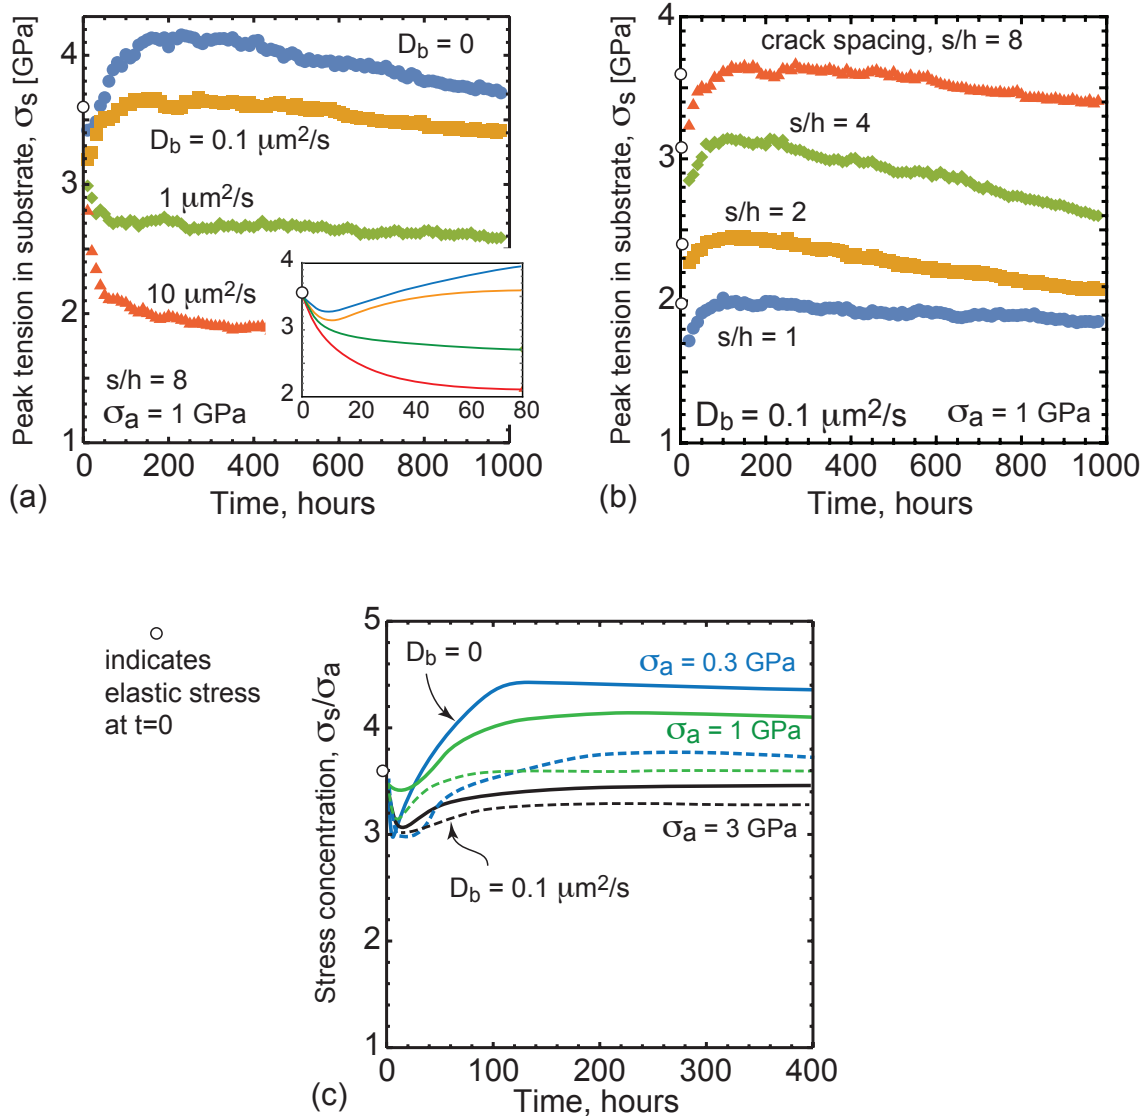


Figure 5.12: (a) Peak stress experienced in the substrate directly ahead of the local oxide that forms at the tip of the crack as a function of time, for several coating diffusivities and a single crack spacing. (b) Peak stress experienced in the substrate directly ahead of the local oxide that forms at the tip of the crack as a function of time, for several crack spacings and a single diffusivity. (c) Stress concentration in the substrate as a function of time, for several different levels of remote loading and two coating diffusivities. Higher levels of remote loading somewhat mitigate the stress concentration by enhancing creep relaxation in the oxide.

and under the intact coating; in turn, since diffusion down the crack is by comparison instantaneous, these are governed by the difference in diffusion through the oxide formed

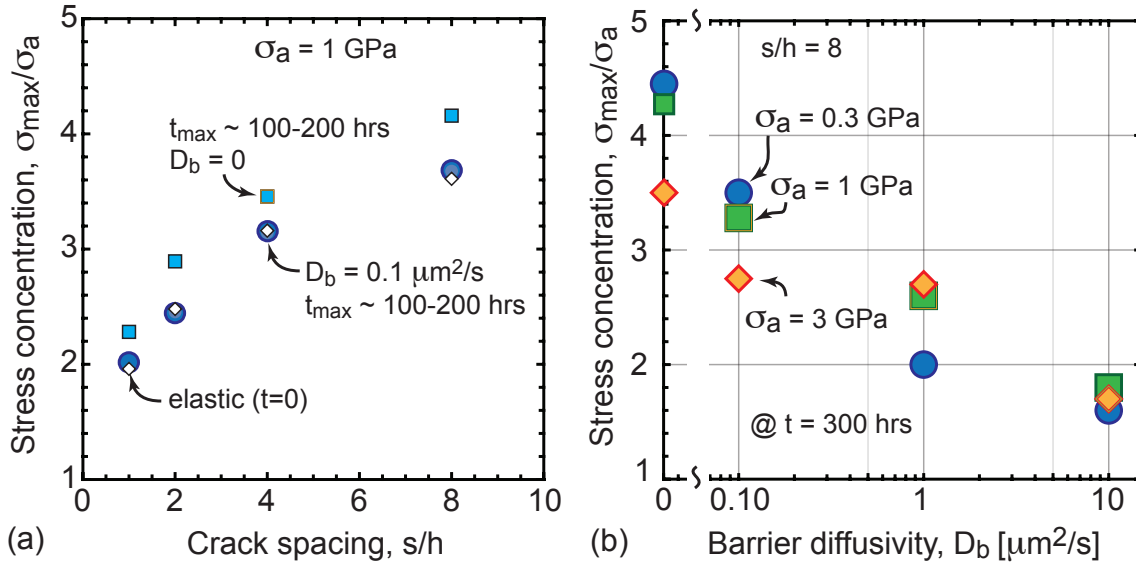


Figure 5.13: (a) Peak stress concentration in the substrate as a function of crack spacing, for coatings with negligible diffusivity and that quoted in the literature for YbSi_2O_7 . The elastic stress concentration refers to that associated with the crack tip at a distance equal to the initial oxide thickness. (b) Comparison of stress concentrations at 300 hours (which is quite close to the peak concentration) as function of coating diffusivity, for three different levels of remote loading.

at the tip of the crack (akin to an oxide on a bare surface) and that formed under the coating. Phenomena that increase diffusion through the oxide itself, or slow diffusion through the coating (e.g. a thicker coating), generally produce larger tensile stresses that could drive damage in the substrate even at elevated temperatures. In light of the results for the case with no diffusion through the coating, however, there is clearly an upper limit to the geometric changes that generate stress; in the limit that the transport through the oxide is much much greater than the coating, the local oxide adopts a semi-cylindrical shape whose radius evolves with time. For very thick substrates (such as those considered here), the stress concentration associated an equiaxed “oxide notch” in the substrate is independent of depth; hence, stress concentrations in the substrate become nearly constant as time progresses.

It should be noted that when the thickness of the oxide under the crack is similar to

that under the intact coating – as occurs for coatings with poor barrier properties – the stresses in the substrate are not appreciably changed from that associated with remote loading. In the present simulations, the stress in the substrate directly under the crack and the oxide layer starts an elevated value, due to the stress concentration associated with that of the crack; this value depends strongly on the initial oxide thickness, i.e. the distance from the corner singularity at the bottom of the crack. Over time, as the oxide thickens, the peak stress in the substrate decreases because the top of the substrate moves away from the singularity point and local variations in oxide thickness do not intervene to elevate substrate stress.

Given the large volumetric change associated with oxide formation, it may be tempting to conclude that this expansion itself – i.e. the large compressive stresses in the oxide – drive large tensile stresses in the substrate. While this may be a contributing factor, the results shown in Figure 5.14 suggest this effect plays a second role, if any. In Figure 5.14, the stress in the substrate adjacent to the oxidation front is shown as a function of time for two low levels of remote loading. Cases with small coating diffusivity and large geometry changes result in high tensile stresses; cases without geometry change (i.e. high diffusivity, $D_b = 1 \mu\text{m}^2/\text{s}$) exhibit *compressive* stresses immediately adjacent to the oxide interface. This strongly suggests that the results presented earlier for elevated stresses can not be viewed simply as the superposition of stresses induced by oxide growth and those associated with remote loading; if the large tensile stresses reflected in Figures 5.11-5.13 were generated by the oxide regardless of remote loading, they should remain tensile (and arguably quite high) even at low levels of remote stress. This is clearly not the case; in Figure 5.14, high stresses are only generated for cases that exhibit significant geometry change, and those stresses scale with remote loading. The notion that substrate stresses are dominated by geometry change and remote loading and not by oxide expansion is further evidenced by the results in Figure 5.13b, which demonstrates that scaled stresses

(i.e. the stress concentration factor) are relatively insensitive to level of remote loading.




Figure 5.14: Stress in the substrate immediate beneath the moving oxide boundary for three different coating diffusivities and two different values of remote tension.

That said, it is not fair to conclude that the expansion of the oxide plays *no* role in the generation of stress in the substrate. It is simply that this role is mitigated by the presence of creep in the oxide, which quickly reduces the stress over the majority of the oxide domain. That is, extremely high compressive stresses due to oxide formation persist over relative short times, and hence are confined to a highly localized sliver of oxide adjacent to the growth front, i.e. the newly formed oxide. The stress in the oxide bubble directly under the crack does contribute to the stress in the substrate to some degree; this is evidenced by the fact that higher remote stresses decrease the stress concentration of the bubble. It is reasonable to infer that expansion of the elastic medium surrounding the bubble accommodates part of the oxide expansion; the more oxide expansion is accommodated, the less it contributes to substrate stresses. This inference is supported by the fact that stress concentrations in the substrate are 4-5 for low levels of remote loading, but only 3-4 for higher levels of loading. Even for the same oxide geometry, increasing remote loading decreases the contribution of oxide expansion to the overall

stress concentration.

Despite the extensive number of simulations in the current parameter study (e.g. variations in coating diffusivity, crack spacing and applied loading), there are number of issues remaining that warrant additional study. Some of this can likely be reduced by considering the present results in terms of dimensionless variables that account for competing effects. For instance, the influence of coating thickness can likely be re-interpreted by adjusting the diffusivity of the coating, and vice versa. The characteristic transport time through the coating is defined by h_b^2/D_b , such that results for thicker coatings (e.g. $100 \mu m$) can be inferred by examining results for lower diffusivities, which will have the same transport times. Similarly, changes in oxidation rate (i.e. the Deal-Grove parameters controlling oxidation front velocities, and imposed boundary concentrations) are likely well-captured by adapting the ratio of oxidation rate to transport through the oxide. Since diffusion down the crack will be significantly faster regardless of such parameters, the only influence of the crack is the width of the opening relative to the crack spacing, as this sets the area fraction of substrate surface that is exposed to rapid diffusion. However, the crack opening of $1 \mu m$ is never more than 4% of the crack spacing; for most purposes, it can be reasonably thought of as a point-source for the oxidant.

As such, with regards to time scales associated with diffusion and oxidation, the present results can be expected to provide applicable insights across a broad range of possibilities. Future work should focus on a more rigorous study that identifies the dimensionless parameters that define system response, with an emphasis on the range of behaviors observed from this study. Naturally, there may be a host of enhancements to the diffusion analysis that may be warranted, such as a jump in saturation concentration between layers (assumed here to be negligible), the influence of adsorption coefficients on the outer surface, etc.

Arguably, the greatest limitation of the present parameter study is the consideration

of a fixed creep properties, i.e. those computed at 958°C. Oxide creep properties are a strong function of temperature and composition, such that future study is strongly warranted to evaluate the competition of creep time scales and oxidation time scales. Interestingly, oxidation rates peak at intermediate temperatures [97], while creep rates will be lower. This strongly suggests that the present calculations may represent the *lower* bound for stresses that develop as a result of oxide growth. Future simulations should focus on a study of the interaction between creep rates and oxide growth kinetics by varying dimensionless parameters involving oxide growth rates and characteristic creep rates as defined by the viscosity of the creep law. In scenarios where creep rates are limited but oxide growth rates are similar, one can anticipate larger effects from oxide volume expansion due to the fact creep relaxation will be less pronounced.

The present study has important potential implications for embrittlement of CMCs at the fiber level, although the inferences should be qualitative rather than quantitative in nature. On the one hand, the cracked coating is similar to CMC matrix cracks serve as fast diffusion pathways that enable attack of internal fibers; the present geometry is similar to that of a “unit cell” of a bridged matrix crack with a single fiber. On the other hand, there are two critical differences that require further simulations to address the bridged crack problem. First, the relevant bridged crack unit cell is obviously axisymmetric, as opposed to the plane strain case considered here. This will likely have critical differences in light of that creep relaxation plays a critical role and is strongly influenced by multi-axial stress states. Further, the fiber problem likely involves oxide thickness that is comparable to the “substrate” dimension, i.e. the fiber radius. This will change the nature of how the elastic SiC phase accommodates the oxide volume expansion and subsequent crepe behavior. Naturally, it will also imply a far greater impact of the “notch” formed by local oxide penetration, since comparable oxide thickness will produce significant increases in net section stress, i.e. the stresses in the SiC fiber at the matrix

crack plane.

Finally, while the focus of the present work has been on the behavior that occurs at elevated temperatures while the oxide is forming, it is clear that cooling will induce additional behaviors that are known to be deleterious. Cooling can lead to fracture in the oxide, which may penetrate into the adjacent SiC layer. This cracking is driven not only by thermal expansion mismatch, but also by phase changes in the silica upon cooling that have significant volume contractions. (This is only relevant if it forms in the crystalline phase at higher temperatures (e.g. greater than 1200°C.) This leads naturally to a *cyclic* damage mechanism: oxide growth drives cracking upon cooling, which upon re-heating serve as fast diffusion pathways to continue oxidation. The next cooling cycle extends the cracks and the process is repeated.

In this context, the present study makes an important contribution by establishing the “baseline” of residual stress present at the start of a cooling period; it is clear from the calculations that assuming the oxide is stress free at elevated temperatures — as is typical of analyses that assess the likelihood of cracking during cooling [16] — is not justified. While creep deformation is significant, it does not alleviate the compressive strain in the oxide at elevated temperatures. This will lower the crack driving force arising from cooling, as since any compressive stress remaining in the oxide will help offset tensile misfit strains. Indeed, the present analysis provides insight into a disconnect between experimental observations of damage and simple estimates of crack driving forces; if the system was stress-free at temperature and then subject to the full misfit strain upon cooling, cracking driving forces would likely be high enough drive cracking after short exposure times [16]. It is clear that simulations of cooling, starting from the present calculations, will generate important insights regarding environmental damage to EBCs and CMCs.

5.6 Concluding remarks

This study of local oxidation near fast diffusion pathways through protective barriers revealed several important insights that will be useful in developing durable environmental barrier coatings. The main insights can be summarized as follows:

- For realistic crack openings, even those at the nanoscale, diffusion down cracks is much faster than diffusion through protective coatings, even for relatively poor barriers with relatively high diffusivities. As such, cracks can be viewed as “point sources” for environmental reactants, and oxide growth at these locations is quantitatively similar to that which occurs for bare surfaces.
- The difference in reactant supply at the location of the crack and that under the intact portion of the coating drives the formation of an oxide “bubble” that adopts a semi-cylindrical shape when the difference is large. The depth of the oxide bubble is less than oxide thickness on bare surfaces, due to diffusion of the reactant in all directions once it exits the crack tip. Nevertheless, a simple and fair approximation to the geometry of the bubble corresponds to the intersection of a half-circle with a radius equal to the bare surface oxide thickness, and the oxide thickness that would form under an intact coating.
- This geometry change to the underlying substrate is the principal factor in the evolution of tensile stresses that are likely a significant contributor to substrate damage. Stress concentrations in the substrate range from 3 (for high remote loading) to 5 (for low remote loading).
- Creep in the oxide layer significantly mitigates the misfit stresses arising from the volume expansion triggered by oxide formation, and reduces the region of high

compressive stresses to a narrow sliver adjacent to the moving oxide front. Nevertheless, compressive stresses in the oxide persist over long times and will likely affect what happens upon cooling.

- The interaction of oxide formation (i.e. volume expansion), remote loading and creep is nuanced, due to the fact that higher remote loading can both drive faster creep and provide some strain accommodation for local oxide “bubble”. Nevertheless, the stress concentration in the substrate is dominated by the geometry change; in effect, the oxide bubble creates a semi-cylindrical notch.
- Additional simulations for behaviors at high temperature are warranted, particularly those that explore the interplay between reaction kinetics and creep relaxation; the facts that oxide growth rates can be faster and creep relaxation slower at intermediate temperatures indicates such simulations will be critical to understand the role of ‘pecking’ damage in CMCs.
- Future analyses of cooling are critical to understanding the evolution of environmentally-driven damage; the present calculations suggest that conventional analyses that assume the system is stress-free when computing crack driving forces will not be accurate. In that regard, the present simulations provide reasonable estimates for the reference state that defines stresses after cooling.

Chapter 6

Summary and Recommendations for Future Work

This dissertation provides important insights into both stress evolution and cracking of brittle ceramic multilayers. Tools and techniques were also presented that allow other problems to be solved that were previously not possible due to computational cost or the multiple physical phenomena that are involved. The contributions in this dissertation support the following conclusions and recommendations for future work:

Advances in cohesive zone modeling

- The distributed cohesive zone method is a powerful tool to study arbitrary crack growth but simulations can be very costly. The proper selection of a cohesive law decreases mesh dependence and increases precision allowing for accurate results to be obtained faster. Excluding near the crack tip, the simulation is mostly elastic and this can be exploited to greatly reduce the computational cost by reducing the effective domain size. Further cost improvements in performance can be had by controlling element size and restricting refined meshes to only near crack tips. An advanced adaptive and dynamic remeshing algorithm was developed that promises

further reductions in computational cost.

- The increases in computational performance allow for the study of problems that were previously prohibitively expensive. The dynamic remeshing algorithm allows for the tracking of multiple crack tips. Dynamic crack growth that includes branching is possible in a reasonable time. An arbitrary number of cracks and how they interact with each other and internal geometric features are easily studied. An interesting application would be to study how cracks interact with internal interfaces, such as between a coating and substrate or between grains. With the ability to track multiple simultaneous cracks, interesting behavior such as spallation of a coating could be observed.
- While the current reduce in the computational cost allow for many more types of problems to be studied, the current implementation of the dynamic remeshing algorithm has significant limitations. Currently, it is implemented with a combination of C++ and Python and reliability can be a problem as the various components interface with each other. Future improvements could include incorporating all components into a unified application and to improve reliability. Many of the techniques that were employed in the multiphysics framework to transfer data between meshes can be applied to the dynamic remeshing algorithm as well.

Cracking near wavy interfaces

- Through the use of the distributed cohesive zone method, the transition from delamination along to kinking off of wavy interfaces was found. It was also determined that while wavy interfaces increase the apparent toughness, the ratio of the toughnesses of the materials controls how much toughening occurs. If the toughness of the bulk material is too low relative to the toughness of the interface, kinking can occur. This kinking generally results in a lower toughening effect unless

there is competition between the two modes. Crack growth through wavy lamellar structures can lead to partial delamination along interfaces and kinking between delamination sections. The lamellar structure increases the apparent toughness of the material. This behavior suggests patterning interfaces may be beneficial by preventing delamination and spallation.

- The current study was focused on macroscopic mode I loading with no material mismatch on either side of the interface. Only a single case with material mismatch was tested but the effects of material mismatch were not fully explored. Likewise different loading conditions were not investigated. The current work also avoided larger amplitudes where crack face contact occurs. Crack face contact may further enhance the toughening behavior.
- Expanding the work by looking at material mismatch as well as orientation of the crack relative to the interface would help answer questions about the behavior of channel cracks in protective coatings. With the dynamic meshing algorithm, more complex scenarios could be studied to determine how a channel crack impinges on an interface and either deflects or penetrates into the substrate. Further extensions could be made by including the effects of thermal residual stress.

Multiphysics framework

- An integrated computational framework was developed that combined the physics of diffusion, oxidation, creep and elastic deformation into a unified solver. The framework explicitly solves for the boundaries between materials as oxidation occurs. This multiphysics framework was shown to recreate one dimensional oxide growth on silicon carbide and match the classic Deal-Grove model as well as accurately model stress relaxation at multiple temperatures. Further simulations on

a curved silicon carbide substrate illustrated how the hoop stresses in the growing oxide layer transition from compressive to tensile stress.

- The current implementation of the framework does suffer from some limitations. Notably only being able to track a single oxidation front within a single material limits what problems can currently be modeled. Further, the way the explicit material boundaries are defined and strain information mapped from one mesh to another introduces errors. In particular, the mapping procedure artificially smoothes data in a way that can not be directly controlled. While the smoothing helps reduce noise, it also may remove highly stressed regions.
- The current limitations can be addressed with future improvements to the framework. The addition of the capability to track multiple oxidations fronts or oxidation through multiple materials would allow other problems of interest to be studied. For example, a matrix crack in a silicon carbide CMC could be modeled with oxidation of the silicon carbide matrix, boron nitride coating and silicon carbide fiber. This scenario is of great interest as CMC usage increases in turbine engines. Further improvements such as changes to the mapping procedure or using elements that provide continuity in strain may help reduce the introduction of errors.

Internal Oxidation

- The multiphysics framework was used to study the effects of oxidation on silicon carbide under a protective coating. Diffusivity of an oxidant down a crack in a protective coating is significantly faster than through the coating. The relative difference in diffusion rates under a crack versus under an intact coating lead to the formation of an oxide “bubble” which is comparable in size to that predicted by the Deal-Grove model. The formation of this bubble causes a stress concentration

to form in the substrate at high temperatures. These stress concentrations may lead to substrate damage. Creep within the oxide layer greatly diminishes the compressive stress caused by oxidation so only highly compressive stresses are seen near the reaction front. Lower compressive stresses are seen throughout the oxide, however, that will likely affect behavior upon cooling. Applied loading interacts with creep in a nuanced way and higher applied loads do not result in larger stress concentrations.

- The current work was restricted to a single temperature and did not study how diffusion and reaction kinetics may change as the specimen is heated or cooled. Likewise, residual thermal stresses were ignored but could affect the stress state. Since no cooling was allowed, strains caused by differing thermo-mechanical properties or phase changes in the oxide were not addressed. Also, only the oxidant was tracked and any products from oxidation other than silica were ignored. Other work has shown both creep and diffusivity to depend on impurity concentrations (such as carbon).
- The limitations of the current work lead to many additional studies that could be performed. With the current framework, studies at varying constant temperatures can address the interplay between reaction kinetics and creep. The addition of cooling would also help to determine what stress states evolve, especially as the oxide undergoes a phase transformation. Varying creep and diffusion properties would help elucidate the effects of impurities within the oxide. It would also be of interest to look at full thermal cycling.

Appendix A

DCZM User Manual

A.1 Introduction

This document covers how to use the distributed cohesive zone method solver by covering installation, outlining the input files, how to run a simulation and the output files that are produced. The background of the method and a rigorous explanation on the requirements on various inputs are not included, but can be found elsewhere. Likewise, how certain input files are created is not outlined in detail, just the structure that the files need to be in. Some tips on their creation are included, though. The first version of the code was written by Will Pro and then optimized in v2.x by Rone Kwei Lim. Further slight modifications, notably the addition of a bilinear cohesive law, was done by Stephen Sehr. The name Will used was *evolver-x* (for reasons unknown) but Rone referred to this solver as UEC. Final versions of all of Rone's codes were uploaded to a Github repository, but the bilinear law might not be a part of this upload. The bilinear law is present on the versions of the source code present on the Begley group servers (*brg1* and *brg2*) under the *brfiles* user and the GPU servers labeled as v2.07.

A.2 Installation

The source code consists of fifteen header files and three main files. The header files are those from Will and those from Rone's research group. The required header files are

- *arrays.h*
- *block.h*
- *Command.h*
- *declare.h*
- *defs.h*

- *elementProp.h*
- *helper_threads.h*
- *helper.h*
- *initialize.hpp*
- *interface.h*
- *Parameter.h*
- *process.hpp*
- *rand.h*
- *Thread.hpp*
- *Vector.h*

and the main files are

- *checksum.cpp*
- *inverse.cpp*
- *main.cpp*

Each header file will not be fully explained, as they are quite complex and obtusely written. The only exception to this is *interface.h*. This header file implements the cohesive law to be used in the simulation. As the code is written now, the law cannot be changed on the fly and must be recompiled to switch laws. Furthermore, restrictions on inputs limit cohesive law parameters to three. The original law implemented was the trapezoidal law, but this was switched by Stephen to a bilinear law that used the same three inputs. The trapezoidal cohesive law is implemented in the functions *utz_normal* and *utz_tangential* while the bilinear law is in the functions *bilinear_normal* and *bilinear_tangential*. To switch between the two laws, the definition of *sep_x2* must be changed in the function *InterfaceForce* and the function calls switched in the function *Force*. An additional cohesive law could be added similarly

Of the three main files, the bulk of the solution is computed within *main.cpp*. The file *inverse.cpp* computes a matrix inverse and when compiled, creates *inverse.o*. Finally, *checksum.cpp* is a checksum function used to verify the output files. The source is compiled by navigating to the correct directory and calling *make*. There are a number of dependencies, but are not spelled out here.

A.3 Input Files

A total of eight different input files are needed for the solver to run. An explanation of each input file can be found below.

A.3.1 Nodes

The nodes file contains a simple listing of the (x,y) positions of each node in the mesh with one point on each line. Numbering is done implicitly starting at row 0 and the coordinates are comma separated.

A.3.2 Elements

The elements file contains a simple listing of the three nodes that define an element with one element per line and each node number separated by a comma. The node number listed corresponds to the implicitly numbered node list stored in the nodes file. Each element is listed in counterclockwise order, but there is no importance to the first node in each element.

A.3.3 Element Properties

The element properties file contains a simple listing of the elemental properties to be assigned to each element. Each line is a different element and the properties are comma separated. The properties needed for each element are

- A. E_x : Young's modulus in x direction
- B. E_y : Young's modulus in y direction
- C. ν_{xy} : Poisson's ratio in xy plane
- D. ν_{yz} : Poisson's ratio in yz plane
- E. ν_{zy} : Poisson's ratio in zy plane
- F. ν_{xz} : Poisson's ratio in xz plane
- G. G_{xy} : Shear modulus in xy plane
- H. ρ : Density
- I. α_x : Coefficient of thermal expansion in x direction
- J. α_y : Coefficient of thermal expansion in y direction
- K. α_z : Coefficient of thermal expansion in z direction

- L. *region*: Region label for the element, can be used for interface
- M. α_d : Mass damping parameter
- N. β_d : Stiffness damping parameter
- O. *orientation*: Orientation of the material relative to the global x axis

For isotropic materials, many of the parameters are repeated and the *orientation* is set to 0. Determination of the damping parameters is problem and property dependent and is covered elsewhere. The region label is mostly used with the interface file.

A.3.4 Initial Conditions

For the vast majority of simulations, the initial conditions file is a blank file. If a simulation is being restarted or some other initial condition imposed, this file contains a list of the displacements and velocities to be applied to each node. It is organized as a list of entries with one entry per line. Each entry contains six comma separated values: the element number, the local node number, the x displacement, the y displacement, the x velocity and the y velocity. The local node number is number locally to the element, ranging from 0 to 2. Initial conditions can be applied to any number of nodes and can be written in any order.

A.3.5 Parameters

The final input file is the parameter file which holds various parameters that the simulation needs to run. The parameters in the file are

- `stepSize`: Time step size
- `totalSteps`: Total number of time steps:
- `dumpInterval`: the number of time steps before the code saves an output file
- `rseed1`: a random number seed
- `rseed2`: a random number seed
- `displacementAmp`: a scaling factor to be applied to displacement results
- `valueAmp`: a scaling factor to be applied to other output values
- `planeType`: 1 for plane strain, 2 for plane stress
- `threadCount`: the number of threads to be used

A.3.6 Interface

The interface file contains a snippet of C code that is used to define the cohesive properties. Upon a simulation being run, the interface file is read in and compiled into a custom function. This allows for quite a bit of control over the cohesive properties. In the most simple form, four quantities must be defined: *stiffness*, *strength*, *toughness* and *contact_k*. The first three correspond to the three parameters of the cohesive law and the last is the stiffness of a cohesive element upon contact after it breaks. Much more complex functions can be written. As mentioned earlier, the region label is useful in the interface file because it makes it easy to define cohesive properties in different mesh regions as well as along the boundary between the regions.

Some built in variables can be used within the function. At this time, the known variables are *vertex1*, *vertex2*, *center1* and *center2* where the *vertex* variables are the vertexes along the cohesive element and the *center* variables are the center points of the elements on the cohesive element. Each variable is a coordinate pair and contains both *x* and *y* positions. For example, the *x* position of *vertex1* is given by *vertex1.x()*.

Likewise, there are some built in functions that can be used to create random variation in the cohesive properties. The main function for a random number is *rand.genRandNormal* which takes ten inputs: the mean, standard deviation and 8 seeds. To determine these 8 seeds, the two functions *getSmallerPoint* and *getBiggerPoint* can be used with the vertices and centers to come up with the same four points for a cohesive element. Then taking both the *x* and *y* components of these four points yields the eight seeds needed. When the seeds are the same, the returned number will be the same, so it will be consistent. Many different ways of controlling the cohesive properties could be devised such as the orientation of the cohesive element or its distance from some point in space.

A.3.7 Loading

The loading file is a list of loads to be applied to specific elements and nodes. Each line specifies a specific constraint to apply and contains four entries: the element number, the local node number, the loading type and a loading tag. The element and local node numbers are those used in the initial conditions file. The loading type is an integer from 1 to 5. Types 1 and 2 are displacements and 3 and 4 are forces in the *x* and *y* directions, respectively. Type 5 is temperature control. If simulating a temperature change, the constraint only needs to be applied to one node in each element. The final entry is an integer tag value that can be used in the loading function file.

A.3.8 Loading Function

The loading function file is like the interface file in that it is a snippet of C code that is compiled when the simulation runs. In this snippet, a *loadValue* is assigned that can depend on the load type, element or node indices, nodal position, loading tag or simulation time.

A.4 Running the Solver

The solver takes a total of 22 input arguments with the file names requiring their full paths:

- Custom directory: a directory used to save the compiled interface and loading function files
- Node file
- Elements file
- Element properties file
- Initial conditions file
- Parameters file
- Interface file
- Loading file
- Loading function file
- Stress11 file: mathematica format output, pass */dev/null*
- Stress12 file: mathematica format output, pass */dev/null*
- Stress22 file: mathematica format output, pass */dev/null*
- Yield count file: mathematica format output, pass */dev/null*
- Kinetic energy file: mathematica format output, pass */dev/null*
- Strain energy file: mathematica format output, pass */dev/null*
- Total kinetic energy file: a listing of the total kinetic energy at each output time step
- Total strain energy file: a listing of the total strain energy at each output time step
- Velocity file: a listing of the nodal velocities at all nodes for all output time steps
- Acceleration file: a listing of the nodal accelerations at all nodes for all output time steps
- Checksum file
- VTK file prefix: prefix for the main output files that contains the stresses, displacements, kinetic and strain energies

The main output files are a series of sequentially numbered files created with the VTK file prefix and saved as a *vtu* file. This file standard is read by ParaView or other programs. It contains all the output data except of the velocities and accelerations. The other output files were specifically used by Will and Rone, but not ideal nor space efficient.

A.5 Conclusion

While not fully exhaustive, this working document hopes to hold some information about running the distributed cohesive zone solver. It is my hope that a reader has gained some knowledge about running a simulation but it is understood that further questions may arise.

Appendix B

Multiphysics Solver User Manual

B.1 Introduction

This document explains how to install and use the alpha version of the multiphysics framework developed by Stephen Sehr at UCSB. The framework is limited to a single geometry exposed to oxygen but the user may change geometric and material parameters. In this version, the source code is not available. To minimize compatibility issues, the framework executable is distributed as a Docker image. The only requirements are:

- A. Docker Desktop to run the framework
- B. ParaView to view the results

Both pieces of software have versions for Windows, macOS and various Linux distributions. Docker is a system to run software within a “container” instead of natively on a host computer. The container is essentially a small virtual machine used to emulate the operating system the executable needs to run. A container is run from an image file that includes all necessary files and applications.

B.2 Installation

- A. Download and install Docker Desktop (www.docker.com/products/docker-desktop)
- B. Download and install ParaView (www.paraview.org/download/)
- C. Download and unzip the archive named `mp_solver.zip` which contains the Docker image file (as a tar file), an example input file and a copy of these instructions.
- D. Open a terminal window and load the image with the command:
`docker image load -i path/to/image/file`

The multiphysics framework should now be installed and ready to be run.

B.3 Running the Solver

B.3.1 Examples

To run, the multiphysics framework needs an input file that defines certain geometric and material properties. An example input file is located within the `example` directory. To run the example simulation, open a terminal window and enter the command

```
docker run -v <PATH>:/data mp_solver:0.9
```

where `<PATH>` is the full path to either of the example geometry directories (either `fiber_example` or `EBC_example`). Either simulation should complete within a few minutes. The results will be written to the correct `example` directory and can be viewed with ParaView.

Each simulation should exist within its own directory and the input file should be named `input.txt` within that directory. If that is the case the simulation is run by the command

```
docker run -v /path/to/simulation/directory:/data mp_solver:0.9
```

The `-v` option maps `/path/to/simulation/directory` on the host computer to `/data` within the container.

B.3.2 Options for `docker run`

Many optional flags are available for the `docker run` command. The reader is directed to the Docker web documentation (or `docker run --help` for a full listing). A few of the flags that may be useful are:

- `-v`: Maps a directory on the host computer to a specifically created directory within the container. The path on the host computer must be the full path, not a relative path from the current working directory.
- `--rm`: Removes the container once it has finished.
- `-d`: Detaches the running container from the current terminal. This allows multiple simulations to be started from the same terminal window.

It should be noted that `-v` and `-d` used a single dash while `--rm` uses two.

B.3.3 Arbitrary Input File

Input files named other than `input.txt` can be used by passing the name of the input file as an argument at the end of the `docker run` command. For example, if the input file is named `new_input.txt` the run command is

```
docker run --rm -d -v <PATH>:/data mp_solver:0.9 new_input.txt
```

with the recommended options.

B.4 Example Geometries

B.4.1 Bare Fiber

The first example geometry is a single uncoated silicon carbide fiber. For simplicity, only a quarter of the cross section of the fiber is modeled. Appropriate boundary conditions are applied to approximate a full fiber. The outer edge of the fiber is a thin, existing oxide which is exposed to dry air.




Figure B.1: Example single fiber geometry.

B.4.2 Oxidation Under an EBC

The second example geometry included in this package is shown in Figure B.2. A silicon carbide substrate has a protective coating applied with a thin initial oxide layer between the silicon carbide and the protective layer. The top edge of the protective layer is exposed to oxygen at a constant temperature. A through-thickness crack exists in the protective layer that acts as a fast diffusion pathway.

B.5 Input Files

The input file is a simple text file of space separated keywords and values. The first keyword must be `geometry` followed by a value of `fiber` or `EBC`. If this first line is not present an error will be thrown and the simulation terminated. The parameters that may be altered in a simulation are shown in Table B.1 along with their default values. Each default is overridden by including a line in the input file such as

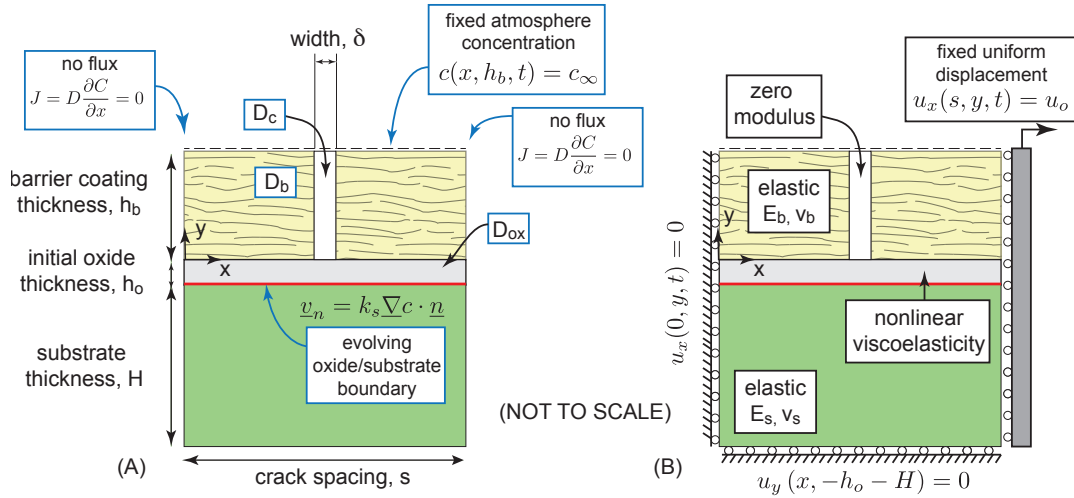


Figure B.2: Example geometry included in this alpha version.

keyword value

The example input files show how to override the defaults for some different parameters.

B.6 Summary of Outputs

With default parameters, the multiphysics framework outputs three types of files to the `base_path` directory:

- `log.txt`: Output log primarily showing when an output file is saved
- `errors.txt`: Error log file that will be blank if there are no errors
- `file_prefix*.vtu` where `*` is an integer: Series of result output files that can be read by ParaView.

Three additional output file types are optional and are primarily used for debugging:

- `file_prefix_<region>*.vtu`: ParaView compatible file for each geometric domain (crack, oxide, substrate or coating)
- `file_prefix_segments*.vtu`: ParaView compatible file with the points and line segments used to create the triangulation
- `timing.txt`: Timing log file showing specific time information for various parts of each step

If an error occurs, the `segments` and `<region>` files are written with `ERROR` in the name in order to help diagnose the source of the problem. Errors are discussed below.

Table B.1: Table of default parameters

| Keyword | EBC default | Fiber default | Description | Source |
|---------------------|--------------------------------|----------------------------------|---|-----------------------------------|
| base.path | ./ | ./ | Path to the directory where output files are written relative to the simulation directory | |
| file_prefix | EBC_ | fiber_ | File prefix for the output files | |
| crack_width | 1 μm | - | Crack opening | |
| crack_spacing | 25 μm | - | Horizontal distance between crack centerlines | |
| coating_thickness | 25 μm | - | Thickness of protective coating | |
| fiber_radius | - | 3 μm | Fiber radius | |
| D_coating | 0.1 $\mu\text{m}^2/\text{s}$ | - | Diffusivity of protective coating | |
| substrate_thickness | 500 μm | - | Thickness of the silicon carbide substrate | |
| ox.t | 0.1 μm | 0.005 μm | Initial thickness of oxide layer | |
| num_growth_steps | 2000 | 99 | Total number of time steps in simulation | |
| dump_int | 20 | 10 | Number of time steps between outputs | |
| dt | 1800 seconds | 500 | Time step size | |
| min_area | 0.001 μm^2 | $5 \times 10^{-6} \mu\text{m}^2$ | Minimum element size | |
| refine_radius | 0.1 μm | 0.1 μm | Controls the transition from small to large elements | |
| D_SiC | 0 $\mu\text{m}^2/\text{s}$ | 0 $\mu\text{m}^2/\text{s}$ | Diffusivity of silicon carbide substrate | |
| E_SiC | 400 GPa | 400 GPa | Young's modulus of silicon carbide substrate | Hay (2012) <i>J. App. Physics</i> |
| nu_SiC | 0.35 | 0.35 | Poisson's ratio of silicon carbide substrate | |
| rxn_rate_SiC | 10 $\mu\text{m}/\text{s}$ | 10 $\mu\text{m}/\text{s}$ | Surface reaction rate | |
| E_coating | 180 GPa | - | Young's modulus of protective coating | |
| nu_coating | 0.27 | - | Poisson's ratio of protective coating | |
| D_SiO2 | 0.725 $\mu\text{m}^2/\text{s}$ | 0.725 $\mu\text{m}^2/\text{s}$ | Diffusivity of the oxide | Norton (1961) <i>Nature</i> |
| eps_SiO2 | 0.3 | 0.3 | Linear expansion strain associated with oxidation | |
| E_SiO2 | 70 GPa | 70 GPa | Young's modulus of oxide | Hay (2012) <i>J. App. Physics</i> |
| nu_SiO2 | 0.17 | 0.17 | Poisson's ratio of oxide | Hay (2012) <i>J. App. Physics</i> |
| D_crack | set by crack width | - | Diffusivity of the crack | |
| Temperature | 1000°C | 1000°C | Temperature | |
| point_equal | 1e-6 μm | $1 \times 10^{-6} \mu\text{m}$ | Tolerance for distance measurements | |
| applied_strain | 0 | - | Applied horizontal strain | |
| applied_stress | 0 GPa | - | Applied horizontal stress | |
| output_polygons | false | false | Output material polygons (debugging) | |
| output_segments | false | false | Output points and segments used to create the mesh (debugging) | |
| output_timing | false | false | Output timing log (debugging) | |

B.7 Troubleshooting

The multiphysics framework includes limited error handling and does not check for invalid inputs. Most computational errors are first observed when the displacement or concentration finite element equations are solved. If the results include non-numeric values (such as NaN), the solver outputs an error message and saves debugging information before attempting to back up to the previous saved output. It was found that this sometimes allowed the simulation to proceed. If the simulation still failed at the same step, it exits showing an error message.

In general, the solver is fairly stable when provided valid inputs. Few, if any, crashes were observed with the single fiber geometry. The majority of crashed simulations occur with the EBC geometry are preceded by geometric instabilities where multiple domains meet. Specifically, the instabilities typically occur at the junction between the crack, oxide and protective coating domains. The source of the instability is not currently

known but is more likely to occur by increasing either the crack spacing or the applied stress. It was observed that a simulation that fails might be able to be successfully run by slightly altering the `crack_spacing`, `crack_width` or `D_coating` by less than 1% allowed the simulation to run to completion.

Bibliography

- [1] F. W. Zok, “Ceramic-matrix composites enable revolutionary gains in turbine engine efficiency,” *www.ceramics.org — American Ceramic Society Bulletin*, vol. 95, no. 5, 2016.
- [2] T. E. Steyer, “Shaping the future of ceramics for aerospace applications,” *International Journal of Applied Ceramic Technology*, vol. 10, no. 3, pp. 389–394, 2013.
- [3] M. D. Thouless and A. G. Evans, “Effects of pull-out on the mechanical properties of ceramic-matrix composites,” *Acta Metallurgica*, vol. 36, no. 3, pp. 517–522, 1988.
- [4] W. A. Curtin, “Theory of Mechanical Properties of Ceramic-Matrix Composites,” *Journal of the American Ceramic Society*, vol. 74, no. 11, pp. 2837–2845, 1991.
- [5] E. B. Callaway and F. W. Zok, “Tensile response of unidirectional ceramic mini-composites,” *Journal of the Mechanics and Physics of Solids*, vol. 138, p. 103903, 2020.
- [6] M. R. Begley and J. W. Hutchinson, *The Mechanics and Reliability of Films, Coatings and Multilayers*. Cambridge: Cambridge University Press, 2017.
- [7] C. G. Levi, J. W. Hutchinson, M. H. Vidal-Sétif, and C. A. Johnson, “Environmental degradation of thermal-barrier coatings by molten deposits,” *MRS Bulletin*, vol. 37, no. 10, pp. 932–941, 2012.
- [8] L. Freund and S. Suresh, *Thin Film Materials: Stress, Defect Formation and Surface Evolution*. Cambridge: Cambridge University Press, 2008.
- [9] J. W. Pro, S. Sehr, R. K. Lim, L. R. Petzold, and M. R. Begley, “Conditions controlling kink crack nucleation out of, and delamination along , a mixed-mode interface crack,” *Journal of the Mechanics and Physics of Solids*, vol. 121, pp. 480–495, 2018.
- [10] G. M. Erickson, S. M. Kuhn-Hendricks, and B. A. Krick, “Fracture Property Experimentation On Hadrosaurid Dinosaur Wavy Enamel Reveals EnergyRobbing Crack Deflection and Channeling To Localize Damage: A Rare Case of MammalianLike Dental Sophistication In Reptiles,” *Journal of the Federation of American Societies for Experimental Biology*, vol. 31, no. S1, pp. 251.5–251.5, 2017.

- [11] S. Sehr, S. Amidi, and M. R. Begley, "Interface delamination vs. bulk cracking along wavy interfaces," *Engineering Fracture Mechanics*, vol. 206, pp. 64–74, oct 2019.
- [12] B. E. Deal and A. S. Grove, "General relationship for the thermal oxidation of silicon," *Journal of Applied Physics*, vol. 36, no. 12, pp. 3770–3778, 1965.
- [13] R. M. Sullivan, "Reformulation of oxide growth equations for oxidation of silicon bond coat in environmental barrier coating systems," *Journal of the European Ceramic Society*, vol. 39, no. 16, pp. 5403–5409, 2019.
- [14] R. S. Hay, "Growth stress in SiO₂ during oxidation of SiC fibers," *Journal of Applied Physics*, vol. 111, p. 063527, mar 2012.
- [15] W. Xu, F. W. Zok, and R. M. McMeeking, "Model of Oxidation-Induced Fiber Fracture in SiC/SiC Composites," *Journal of the American Ceramic Society*, vol. 97, no. 11, pp. 3676–3683, 2014.
- [16] B. T. Richards, K. A. Young, F. De Francqueville, S. Sehr, M. R. Begley, and H. N. Wadley, "Response of ytterbium disilicate-silicon environmental barrier coatings to thermal cycling in water vapor," *Acta Materialia*, vol. 106, pp. 1–14, 2016.
- [17] M. S. Hu, M. D. Thouless, and A. G. Evans, "The decohesion of thin films from brittle substrates," *Acta Metallurgica*, vol. 36, no. 5, pp. 1301–1307, 1988.
- [18] M. D. Drory, M. D. Thouless, and A. G. Evans, "On the decohesion of residually stressed thin films," *Acta Metallurgica*, vol. 36, no. 8, pp. 2019–2028, 1988.
- [19] S. Li, J. Wang, and M. D. Thouless, "The effects of shear on delamination in layered materials," *Journal of the Mechanics and Physics of Solids*, vol. 52, no. 1, pp. 193–214, 2004.
- [20] S. Li, M. D. Thouless, A. M. Waas, J. a. Schroeder, and P. D. Zavattieri, "Use of a cohesive-zone model to analyze the fracture of a fiber-reinforced polymer-matrix composite," *Composites Science and Technology*, vol. 65, no. 3-4, pp. 537–549, 2005.
- [21] S. Li, M. D. Thouless, A. M. Waas, J. a. Schroeder, and P. D. Zavattieri, "Use of mode-I cohesive-zone models to describe the fracture of an adhesively-bonded polymer-matrix composite," *Composites Science and Technology*, vol. 65, no. 2, pp. 281–293, 2005.
- [22] S. Li, M. D. Thouless, A. M. Waas, J. a. Schroeder, and P. D. Zavattieri, "Competing failure mechanisms in mixed-mode fracture of an adhesively bonded polymer-matrix composite," *International Journal of Adhesion and Adhesives*, vol. 26, no. 8, pp. 609–616, 2006.

- [23] J. P. Parmigiani and M. D. Thouless, “The roles of toughness and cohesive strength on crack deflection at interfaces,” *Journal of the Mechanics and Physics of Solids*, vol. 54, no. 2, pp. 266–287, 2006.
- [24] S. Li, M. D. Thouless, A. M. Waas, J. a. Schroeder, and P. D. Zavattieri, “Mixed-mode cohesive-zone models for fracture of an adhesively bonded polymer-matrix composite,” *Engineering Fracture Mechanics*, vol. 73, no. 1, pp. 64–78, 2006.
- [25] J. P. Parmigiani and M. D. Thouless, “The roles of toughness and cohesive strength on crack deflection at interfaces,” *Journal of the Mechanics and Physics of Solids*, vol. 54, no. 2, pp. 266–287, 2006.
- [26] J. P. Parmigiani and M. D. Thouless, “The effects of cohesive strength and toughness on mixed-mode delamination of beam-like geometries,” *Engineering Fracture Mechanics*, vol. 74, no. 17, pp. 2675–2699, 2007.
- [27] R. B. Sills and M. D. Thouless, “The effect of cohesive-law parameters on mixed-mode fracture,” *Engineering Fracture Mechanics*, vol. 109, pp. 353–368, 2013.
- [28] R. B. Sills and M. D. Thouless, “Cohesive-length scales for damage and toughening mechanisms,” *International Journal of Solids and Structures*, vol. 55, pp. 32–43, 2015.
- [29] N. Chandra, H. Li, C. Shet, and H. Ghonem, “Some issues in the application of cohesive zone models for metal ceramic interfaces,” *International Journal of Solids and Structures*, vol. 39, pp. 2827–2855, 2002.
- [30] J. L. Strom and J. P. Parmigiani, “Transition of crack path at bi-material interfaces,” *Engineering Fracture Mechanics*, vol. 115, pp. 13–21, 2014.
- [31] J. W. Foulk, G. C. Johnson, P. A. Klein, and R. O. Ritchie, “On the toughening of brittle materials by grain bridging: Promoting intergranular fracture through grain angle, strength, and toughness,” *Journal of the Mechanics and Physics of Solids*, vol. 56, no. 6, pp. 2381–2400, 2008.
- [32] X. P. Xu and A. Needleman, “Numerical Simulations of Fast Crack-Growth in Brittle Solids,” *J. Mech. Phys. Solids*, vol. 42, no. 9, pp. 1397–&, 1994.
- [33] J. W. Pro, *Distributed cohesive zone methods for modeling fracture in brittle coating systems & composites*. PhD thesis, University of California Santa Barbara, 2016.
- [34] V. Tvergaard and J. W. Hutchinson, “The relation between crack growth resistance and fracture process parameters in elastic-plastic solids,” *Journal of the Mechanics and Physics of Solids*, vol. 40, no. 6, pp. 1377–1397, 1992.
- [35] J. R. Shewchuk, *Triangle: Engineering a 2D quality mesh generator and delaunay triangulator*, vol. 1148. Springer-Verlag, 1996.

- [36] D. R. Clarke and C. G. Levi, “Materials Design for the Next Generation Thermal Barrier Coatings,” *Annual Review of Materials Research*, vol. 33, no. 1, pp. 383–417, 2003.
- [37] N. P. Padture, M. Gell, and E. H. Jordan, “Thermal barrier coatings for gas-turbine engine applications.,” *Science (New York, N.Y.)*, vol. 296, no. 5566, pp. 280–4, 2002.
- [38] D. R. Clarke, M. Oechsner, and N. P. Padture, “Thermal-barrier coatings for more efficient gas-turbine engines,” *MRS Bulletin*, vol. 37, no. 10, pp. 891–898, 2012.
- [39] K. N. Lee, “Key Durability Issues With Mullite-Based Environmental Barrier Coatings for Si-Based Ceramics,” *ASME 1999 International Gas Turbine and Aeroengine Congress and Exhibition*, vol. 4, pp. 2–8, 1999.
- [40] K. N. Lee, “Current status of environmental barrier coatings for Si-based ceramics,” *Surface and Coatings Technology*, vol. 133-134, pp. 1–7, 2000.
- [41] I. Spitsberg and J. Steibel, “Thermal and Environmental Barrier Coatings for SiC/SiC CMCs in Aircraft Engine Applications,” *International Journal of Applied Ceramic Technology*, vol. 1, no. 4, pp. 291–301, 2005.
- [42] M.-Y. He and J. W. Hutchinson, “Kinking of a Crack Out of an Interface,” *Journal of Applied Mechanics*, vol. 111, 1989.
- [43] M.-Y. He, A. Bartlett, A. G. Evans, and J. W. Hutchinson, “Kinking of a Crack out of an Interface: Role of In-Plane Stress,” *Journal of the American Ceramic Society*, vol. 74, 1991.
- [44] M.-Y. He and J. W. Hutchinson, “Kinking of a crack out of an interface: tabulated solution coefficients,” tech. rep., Harvard University Report MECH-113A, Cambridge, MA, 1989.
- [45] J. W. Hutchinson and Z. Suo, “Mixed Mode Cracking in Layered Materials,” 1991.
- [46] B. A. Bilby, G. E. Cardew, and I. C. Howard, “Stress intensity factors at the tips of kinked and forked cracks,” in *Analysis and Mechanics*, pp. 197–200, Elsevier, 1978.
- [47] B. A. Bilby and G. E. Cardew, “The crack with a kinked tip,” *International Journal of Fracture*, vol. 11, no. 4, pp. 708–712, 1975.
- [48] R. J. Nuismer, “An energy release rate criterion for mixed mode fracture,” *International Journal of Fracture*, vol. 11, no. 2, pp. 245–250, 1975.
- [49] C. H. Wu, “Elasticity problems of a slender Z-crack,” *Journal of Elasticity*, vol. 8, no. 2, pp. 183–205, 1978.

- [50] C. H. Wu, “Maximum-energy-release-rate criterion applied to a tension-compression specimen with crack,” *Journal of Elasticity*, vol. 8, no. 3, pp. 235–257, 1978.
- [51] C. H. Wu, “Fracture under combined loads by maximum-energy-release-rate criterion,” *Journal of Applied Mechanics*, vol. 45, no. September 1978, 1978.
- [52] C. H. Wu, “Explicit asymptotic solution for the maximum-energy-release-rate problem,” *International Journal of Solids and Structures*, vol. 15, pp. 561–566, jan 1979.
- [53] K. Hayashi and S. Nemat-Nasser, “Energy-release rate and crack kinking under combined loading,” *Journal of Applied Mechanics*, vol. 48, no. September 1981, pp. 520–524, 1981.
- [54] A. Azhdari and S. Nemat-Nasser, “Energy-release rate and crack kinking in anisotropic brittle solids,” *Journal of the Mechanics and Physics of Solids*, vol. 44, pp. 929–951, jun 1996.
- [55] K. K. Lo, “Analysis of branched cracks,” *Journal of Applied Mechanics*, vol. 45, no. December 1978, pp. 797–802, 1978.
- [56] B. Karihaloo, L. Keer, and S. Nemat-Nasser, “Crack kinking under nonsymmetric loading,” *Engineering Fracture Mechanics*, vol. 13, pp. 879–888, jan 1980.
- [57] H. C. Tankasala, V. S. Deshpande, and N. A. Fleck, “Crack kinking at the tip of a mode I crack in an orthotropic solid,” *International Journal of Fracture*, vol. 207, no. 2, pp. 181–191, 2017.
- [58] B. T. Richards, S. Sehr, F. De Francqueville, M. R. Begley, and H. N. Wadley, “Fracture mechanisms of ytterbium monosilicate environmental barrier coatings during cyclic thermal exposure,” *Acta Materialia*, vol. 103, pp. 448–460, 2016.
- [59] M. Ranjbar-Far, J. Absi, G. Mariaux, and F. Dubois, “Simulation of the effect of material properties and interface roughness on the stress distribution in thermal barrier coatings using finite element method,” *Materials and Design*, vol. 31, no. 2, pp. 772–781, 2010.
- [60] M. Ranjbar-Far, J. Absi, G. Mariaux, and D. S. Smith, “Crack propagation modeling on the interfaces of thermal barrier coating system with different thickness of the oxide layer and different interface morphologies,” *Materials and Design*, vol. 32, no. 10, pp. 4961–4969, 2011.
- [61] M. Ranjbar-Far, J. Absi, and G. Mariaux, “Finite element modeling of the different failure mechanisms of a plasma sprayed thermal barrier coatings system,” *Journal of Thermal Spray Technology*, vol. 21, no. 6, pp. 1234–1244, 2012.

- [62] H. X. Zhu, N. A. Fleck, A. C. F. Cocks, and A. G. Evans, “Numerical simulations of crack formation from pegs in thermal barrier systems with NiCoCrAlY bond coats,” *Materials Science and Engineering A*, vol. 404, no. 1-2, pp. 26–32, 2005.
- [63] C.-H. Hsueh, J. A. Haynes, M. J. Lance, P. F. Becher, M. K. Ferber, E. R. Fuller, S. a. Langer, W. C. Carter, and W. R. Cannon, “Effects of Interface Roughness on Residual Stresses in Thermal Barrier Coatings,” *Journal of the American Ceramic Society*, vol. 82, pp. 1073–1075, 1999.
- [64] K. W. Schlichting, N. P. Padture, E. Jordan, and M. Gell, “Failure modes in plasma-sprayed thermal barrier coatings,” *Materials Science and Engineering: A*, vol. 342, no. 1, pp. 120–130, 2003.
- [65] M. Blacklock, J. H. Shaw, F. W. Zok, and B. N. Cox, “Virtual specimens for analyzing strain distributions in textile ceramic composites,” *Composites Part A: Applied Science and Manufacturing*, vol. 85, pp. 40–51, 2016.
- [66] M. N. Rossol, T. Fast, D. B. Marshall, B. N. Cox, and F. W. Zok, “Characterizing in-plane geometrical variability in textile ceramic composites,” *Journal of the American Ceramic Society*, vol. 98, no. 1, pp. 205–213, 2015.
- [67] I. A. Malik and F. Barthelat, “Bioinspired sutured materials for strength and toughness: Pullout mechanisms and geometric enrichments,” *International Journal of Solids and Structures*, pp. 1–16, 2018.
- [68] A. Maloul, J. Fialkov, D. Wagner, and C. M. Whyne, “Characterization of craniofacial sutures using the finite element method,” *Journal of Biomechanics*, vol. 47, no. 1, pp. 245–252, 2014.
- [69] T. Miura, C. A. Perlyn, M. Kinboshi, N. Ogihara, M. Kobayashi-Miura, G. M. Morriss-Kay, and K. Shiota, “Mechanism of skull suture maintenance and interdigitation,” *Journal of Anatomy*, vol. 215, no. 6, pp. 642–655, 2009.
- [70] B. Coats and S. S. Margulies, “Material properties of human infant skull and suture at high rates,” *Journal of Neurotrauma*, vol. 23, no. 8, pp. 1222–1232, 2006.
- [71] K. Tai, M. Dao, S. Suresh, A. Palazoglu, and C. Ortiz, “Nanoscale heterogeneity promotes energy dissipation in bone,” *Nature Materials*, vol. 6, no. 6, pp. 454–462, 2007.
- [72] P. D. Zavattieri, L. G. Hector, and A. F. Bower, “Determination of the effective mode-I toughness of a sinusoidal interface between two elastic solids,” *International Journal of Fracture*, vol. 145, no. 3, pp. 167–180, 2007.
- [73] V. Janarthanan, P. D. Garrett, R. S. Stein, and M. Srinivasarao, “Adhesion enhancement in immiscible polymer bilayer using oriented macroscopic roughness,” *Polymer*, vol. 38, no. 1, pp. 105–111, 1997.

- [74] B. W. Li, H. P. Zhao, Q. H. Qin, X. Q. Feng, and S. W. Yu, “Numerical study on the effects of hierarchical wavy interface morphology on fracture toughness,” *Computational Materials Science*, vol. 57, pp. 14–22, 2012.
- [75] L. Yang and J. Qu, “Fracture mechanics parameters for cracks on a slightly undulating interface,” *International Journal of Fracture*, pp. 79–91, 1993.
- [76] S. Vajpayee, K. Khare, S. Yang, C. Y. Hui, and A. Jagota, “Adhesion selectivity using rippled surfaces,” *Advanced Functional Materials*, vol. 21, no. 3, pp. 547–555, 2011.
- [77] A. G. Evans and J. W. Hutchinson, “Effects of non-planarity on the mixed mode fracture resistance of bimaterial interfaces,” *Acta Metallurgica*, vol. 37, no. 3, pp. 909–916, 1989.
- [78] Y. F. Gao and A. F. Bower, “A simple technique for avoiding convergence problems in finite element simulations of crack nucleation and growth on cohesive interfaces,” *Modelling and Simulation in Materials Science and Engineering*, vol. 12, no. 3, pp. 453–463, 2004.
- [79] 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, no. 10, pp. 1665–1682, 2007.
- [80] Z. H. Jin and C. T. Sun, “Cohesive zone modeling of interface fracture in elastic bi-materials,” *Engineering Fracture Mechanics*, vol. 72, no. 12 SPEC. ISS., pp. 1805–1817, 2005.
- [81] Y. Yan and F. Shang, “Cohesive zone modeling of interfacial delamination in PZT thin films,” *International Journal of Solids and Structures*, vol. 46, no. 13, pp. 2739–2749, 2009.
- [82] K.-J. Bathe, *Finite Element Procedures*. 2014.
- [83] R. D. Cooke, D. S. Malkus, M. E. Plesha, and R. J. Will, *Concepts and Applications of Finite Element Analysis*. John Wiley & Sons, 2002.
- [84] D.-B. Kao, J. P. McVittie, W. D. Nix, and K. C. Saraswat, “Two-dimensional thermal oxidation of silicon. II. Modeling stress effects in wet oxides - Electron Devices, IEEE Transactions on,” *IEEE Transactions on Electron Devices*, vol. 35, no. 1, pp. 25–37, 1988.
- [85] L. O. Wilson and R. B. Marcus, “Oxidation of Curved Silicon Surfaces,” *Journal of the Electrochemical Society*, vol. 134, no. 2, pp. 481–490, 1987.

- [86] H. Ohnabe, S. Masaki, M. Onozuka, K. Miyahara, and T. Sasa, "Potential application of ceramic matrix composites to aero-engine components," *Composites Part A: Applied Science and Manufacturing*, vol. 30, no. 4, pp. 489–496, 1999.
- [87] J. H. Perepezko, "The hotter the engine, the better," *Science*, vol. 326, no. 5956, pp. 1068–1069, 2009.
- [88] N. S. Jacobson, E. J. Opila, and K. N. Lee, "Oxidation and corrosion of ceramics and ceramic matrix composites," *Current Opinion in Solid State and Materials Science*, vol. 5, pp. 301–309, 2001.
- [89] K. N. Lee, "Protective Coatings for Gas Turbines 4.4.2-1 Introduction," in *The Gas Turbine Handbook* (R. Dennis, ed.), pp. 419–437, Department of Energy, 2006.
- [90] B. T. Richards, M. R. Begley, and H. N. Wadley, "Mechanisms of Ytterbium Monosilicate/Mullite/Silicon Coating Failure during Thermal Cycling in Water Vapor," *Journal of the American Ceramic Society*, vol. 98, no. 12, pp. 4066–4075, 2015.
- [91] M. Maeda, K. Nakamura, and T. Ohkubo, "Oxidation of silicon carbide in a wet atmosphere," *Journal of Materials Science*, vol. 23, no. 11, pp. 3933–3938, 1988.
- [92] E. J. Opila, "Variation of the Oxidation Rate of Silicon Carbide with Water-Vapor Pressure," *Journal of the American Ceramics Society*, vol. 82, no. 3, pp. 625–636, 1999.
- [93] S. H. Leigh, C. K. Lin, and C. C. Berndt, "Elastic response of thermal spray deposits under indentation tests," *Journal of the American Ceramic Society*, vol. 80, no. 8, pp. 2093–2099, 1997.
- [94] C.-J. Li and A. Ohmori, "Relationships Between the Microstructure and Properties of Thermally Sprayed Deposits," *Journal of Thermal Spray Technology*, vol. 11, no. September, pp. 365–374, 2002.
- [95] T. A. Cruse, B. P. Johnsen, and A. Nagy, "Mechanical properties testing and results for thermal barrier coatings," *Journal of Thermal Spray Technology*, vol. 6, no. 1, pp. 57–66, 1997.
- [96] B. S. Good, "Kinetic Monte Carlo Simulation of Oxygen Diffusion in Ytterbium Disilicate," *MRS Advances*, pp. 1203–1208, 2016.
- [97] C. E. Ramberg, G. Cruciani, K. E. Spear, and R. E. Tressler, "Passive-oxidation Kinetics of High-Purity Silicon Carbide from 800 to 1100C," *Journal of the American Ceramics Society*, vol. 79, no. 11, pp. 2897–2911, 1996.