


NASA/TM–20260002490



Advances in Entry Systems Modeling: Modeling Summer Visit 2025 Technical Reports

Bruno Dias

*Analytical Mechanics Associates, Inc. at NASA Ames Research Center, Moffett Field, CA 94035,
USA*

Sergio Fraile-Izquierdo

*Analytical Mechanics Associates, Inc. at NASA Ames Research Center, Moffett Field, CA 94035,
USA*

Jeremie Meurisse

*Analytical Mechanics Associates, Inc. at NASA Ames Research Center, Moffett Field, CA 94035,
USA*

Marc Massot

CMAP, CNRS, École polytechnique, Institut Polytechnique de Paris, 91120 Palaiseau, France

Jean Lachaud

*Univ. Bordeaux, CNRS, Bordeaux INP, I2M, UMR 5295, F-33400, Talence, France
Arts et Metiers Institute of Technology, CNRS, Bordeaux INP, I2M, UMR 5295, F-33400
Talence, France*

Thierry Magin

*von Karman Institute for Fluid Dynamics (VKI), Sint-Genesius-Rode, Belgium
Université libre de Bruxelles (ULB), Brussels, Belgium*

National Aeronautics and
Space Administration

Ames Research Center
Moffett Field, California

March 2026

Duy Cuong Dinh
University of Poitiers, France

Frank Richard
University of Poitiers, France

Joanna D'Antoni
Université de Bordeaux, France

Ward Haegeman
DMPE, ONERA, Université Paris Saclay, 91120 Palaiseau, France
CMAP, CNRS, École polytechnique, Institut Polytechnique de Paris, 91120 Palaiseau, France

Giuseppe Orlando
CMAP, CNRS, École polytechnique, Institut Polytechnique de Paris, 91120 Palaiseau, France

Samuel Kokh
Université Paris-Saclay, UVSQ, INRIA, CNRS, CEA, Maison de la Simulation, 91191, Gif-sur-Yvette, France

Loic Gouarin
CMAP, CNRS, École polytechnique, Institut Polytechnique de Paris, 91120 Palaiseau, France

Pierre Matalon
CMAP, CNRS, École polytechnique, Institut Polytechnique de Paris, 91120 Palaiseau, France

Sébastien Dubois
CMAP, CNRS, École polytechnique, Institut Polytechnique de Paris, 91120 Palaiseau, France

Alexandre Hoffmann
CMAP, CNRS, École polytechnique, Institut Polytechnique de Paris, 91120 Palaiseau, France

Josselin Massot
CMAP, CNRS, École polytechnique, Institut Polytechnique de Paris, 91120 Palaiseau, France

Laurent Series
CMAP, CNRS, École polytechnique, Institut Polytechnique de Paris, 91120 Palaiseau, France

National Aeronautics and
Space Administration

Ames Research Center
Moffett Field, California

March 2026

Christian Tenaud
Université Paris-Saclay, CNRS, CentraleSupélec, EM2C, 91190 Gif-sur-Yvette, France

Pierre Bernigaud
DMPE, ONERA, Université Paris Saclay, 91120 Palaiseau, France

Antoine Clout
DMPE, ONERA, Université Paris Saclay, 91120 Palaiseau, France

Daniele Biasone
DMPE, ONERA, Université Paris Saclay, 91120 Palaiseau, France
CMAP, CNRS, École polytechnique, Institut Polytechnique de Paris, 91120 Palaiseau, France

Zoubair Tazakkati
CMAP, CNRS, École polytechnique, Institut Polytechnique de Paris, 91120 Palaiseau, France

Teddy Pichard
CMAP, CNRS, École polytechnique, Institut Polytechnique de Paris, 91120 Palaiseau, France

Pietro Parodi
von Karman Institute for Fluid Dynamics (VKI), Sint-Genesius-Rode, Belgium
Katholieke Universiteit Leuven (KU Leuven), Belgium

Jan Skácel
von Karman Institute for Fluid Dynamics (VKI), Sint-Genesius-Rode, Belgium
Université libre de Bruxelles (ULB), Brussels, Belgium
Masaryk University, Czechia

Federico Bariselli
von Karman Institute for Fluid Dynamics (VKI), Sint-Genesius-Rode, Belgium

Bruno Fontaine
Université de Liège (ULiège), Belgium
Université libre de Bruxelles (ULB), Brussels, Belgium
von Karman Institute for Fluid Dynamics (VKI), Sint-Genesius-Rode, Belgium

National Aeronautics and
Space Administration

Ames Research Center
Moffett Field, California

March 2026

Julien Clotuche
Université de Liège (ULiège), Belgium
Université libre de Bruxelles, Brussels, Belgium
von Karman Institute for Fluid Dynamics (VKI), Sint-Genesius-Rode, Belgium

Koen Hillewaert
Université de Liège (ULiège), Belgium
Cenaero, Belgium
von Karman Institute for Fluid Dynamics (VKI), Sint-Genesius-Rode, Belgium

Sander Holum
von Karman Institute for Fluid Dynamics (VKI), Sint-Genesius-Rode, Belgium
Université libre de Bruxelles (ULB), Brussels, Belgium

Roemer Spreij
von Karman Institute for Fluid Dynamics (VKI), Sint-Genesius-Rode, Belgium
Katholieke Universiteit Leuven (KU Leuven), Belgium

Nicolas Derider
von Karman Institute for Fluid Dynamics (VKI), Sint-Genesius-Rode, Belgium
Université libre de Bruxelles (ULB), Brussels, Belgium

Nicolas Tomić
von Karman Institute for Fluid Dynamics (VKI), Sint-Genesius-Rode, Belgium
Université libre de Bruxelles (ULB), Brussels, Belgium

Alexandre Quintart
von Karman Institute for Fluid Dynamics (VKI), Sint-Genesius-Rode, Belgium
Flying Squirrel, Bourg-St-Pierre, 1964, Switzerland

Diana Martins
von Karman Institute for Fluid Dynamics (VKI), Sint-Genesius-Rode, Belgium
Université libre de Bruxelles (ULB), Brussels, Belgium

Maxime Barbouse
Université de Liège (ULiège), Belgium
Université catholique de Louvain (UCLouvain), Belgium
von Karman Institute for Fluid Dynamics (VKI), Sint-Genesius-Rode, Belgium

National Aeronautics and
Space Administration

Ames Research Center
Moffett Field, California

March 2026

Julien Klauner

Université catholique de Louvain (UCLouvain), Belgium

Université libre de Bruxelles (ULB), Brussels, Belgium

von Karman Institute for Fluid Dynamics (VKI), Sint-Genesius-Rode, Belgium

Pierre Schrooyen

von Karman Institute for Fluid Dynamics (VKI), Sint-Genesius-Rode, Belgium

Joseph El Rassi

von Karman Institute for Fluid Dynamics (VKI), Sint-Genesius-Rode, Belgium

David Henneaux

Cenaero, Belgium

Philippe Chatelain

Université catholique de Louvain (UCLouvain), Belgium

Antoine E. Simon

CMAP, CNRS, École polytechnique, Institut Polytechnique de Paris, 91120 Palaiseau, France

DMPE, ONERA, Université Paris Saclay, 91120 Palaiseau, France

Laurent François

DMPE, ONERA, Université Paris Saclay, 91120 Palaiseau, France

Ali Asad

CMAP, CNRS, École polytechnique, Institut Polytechnique de Paris, 91120 Palaiseau, France

Eric Quémerais

DMPE, ONERA, Université Paris Saclay, 91120 Palaiseau, France

Vincenzo Romano

von Karman Institute for Fluid Dynamics (VKI), Sint-Genesius-Rode, Belgium

Antoine Ghyoot

Université libre de Bruxelles (ULB), Brussels, Belgium

Université de Liège (ULiège), Belgium

von Karman Institute for Fluid Dynamics (VKI), Sint-Genesius-Rode, Belgium

National Aeronautics and
Space Administration

Ames Research Center
Moffett Field, California

March 2026

Ata Omur Başkaya
von Karman Institute for Fluid Dynamics (VKI), Sint-Genesius-Rode, Belgium

Flora Lahouze
Université de Bordeaux, France

Emmanuel Lopez
Université Bordeaux, France

Giuseppe Sciume
Université de Bordeaux, France

Célia Mercader
CANOE, France

Lisa Feuillerat
Airbus, France

Caroline Petiot
Airbus, France

Abderrahman M'jikou
Université de Bordeaux, France

National Aeronautics and
Space Administration

Ames Research Center
Moffett Field, California

March 2026

The use of trademarks or names of manufacturers in this report is for accurate reporting and does not constitute an official endorsement, either expressed or implied, of such products or manufacturers by the National Aeronautics and Space Administration.

Available from:

NASA STI Program / Mail Stop 050
NASA Langley Research Center
Hampton, VA 23681-2199

Contents

1 Preface

2 Porous Media

- 2.1 Theoretical and computational framework for mass conservation in a pyromechanics model.
- 2.2 Modeling the pyrolysis of polymer fibers to produce carbon fibers
- 2.3 Integration of a Compressible Flow Solver into PATO for Strong Fluid–Porous Coupling with ALE Interface Motion Applied to Pore-Scale Pyrolysis
- 2.4 Modelling Radiative Heat Transfer through Cracks in the Char Layer during Wood Pyrolysis
- 2.5 Implementation of competitive pyrolysis mechanisms in PATO

3 Numerical methods

- 3.1 A two-phase flow model with scale separation for the separated-to-disperse phase transition
- 3.2 samurai/ponio: two complementary software libraries for the efficient implementation of adaptive space/time numerical schemes
- 3.3 A splitting-ImEx strategy for convection-reaction-diffusion problems with space-time adaptation and error control.

4 Plasma/Rarefied methods

- 4.1 Simulation of hybrid dense-rarefied flows relying on relaxation schemes and a hierarchy of moment methods
- 4.2 Toward Efficient Simulations of Highly Magnetized Plasma
- 4.3 Study of the Electrode Kinetic Layer in Arc Discharges
- 4.4 Shock Tube Radiation Modeling and Sensitivity Analysis

5 Experimental reconstruction

- 5.1 Holistic Supersonic Plasma Ground Testing
- 5.2 In-Situ 3D Reconstruction in High-Enthalpy Test Facilities

6 Flow/material coupling methods

- 6.1 A coupling strategy of flow and material solvers for modeling of high-enthalpy flows with ablation
- 6.2 Development of a high-order time-adaptive multistep coupling scheme in an HPC library for multiphysics applications

7 Transition and High-speed flows

- 7.1 Investigation of Hypersonic Boundary Layer Transition induced by Wall Protrusions and Freestream Disturbances

1 Preface

This NASA Technical Memorandum presents the collective work carried out during the Modeling Summer Visit (MSV) 2025, a collaborative research program between NASA Ames Research Center and three leading European research institutions: the University of Bordeaux, the von Karman Institute (VKI), and École Polytechnique.

Held from July 7th to August 1st, 2025, the NASA ARC MSV program brought together 42 researchers from Europe to work alongside NASA staff on 22 projects in entry systems modeling. The program was organized by Bruno Dias, Sergio Fraile-Izquierdo, and Jeremie Meurisse from Analytical Mechanics Associates, Inc., in collaboration with Marc Massot (École Polytechnique), Jean Lachaud (University of Bordeaux), and Thierry Magin (VKI). The MSV program traces its origins to an initiative started by Nagi N. Mansour in 2018, with this 2025 edition representing a consolidation and expansion of that original vision.

The projects were organized into six thematic groups: Porous Media, Numerical Methods, Plasma/Rarefied Methods, Experimental Reconstruction, Flow and Material Coupling Methods, and Transition and High-Speed Flows. Throughout the four-week program, participants presented their work in three public presentations (July 7th, July 21st, and August 1st), allowing for continuous feedback and discussion of their evolving research.

Program Outcomes and NASA Benefits

The MSV program accomplished significant technical progress in one month. Beyond the work documented in these reports, the program strengthened research partnerships between NASA and European institutions, creating pathways for continued collaboration on fundamental research applied to entry systems. The visiting researchers applied advanced numerical methods to NASA applications in thermal protection system modeling, plasma simulation, and aerothermodynamics. NASA researchers gained exposure to state-of-the-art methods such as adaptive mesh refinement, kinetic-based boundary conditions, competitive chemical mechanisms, and scale-separated multi-phase flow modeling.

The diverse projects address different application areas of NASA's interests. The porous media projects on competitive pyrolysis mechanisms and polymer fiber modeling can directly improve our modeling capabilities and support the development of next-generation ablative thermal protection systems. The plasma and rarefied flow projects, including shock tube radiation modeling, provide better methodologies to extract chemical mechanisms derived from shock tube experiments. The electrode kinetic layer study also helps us better understand electrode erosion and enhance the boundary conditions used to model the arc jet testing facilities, which will ultimately improve the facility's operational performance. The high-order coupling methods can help to advance NASA's capability to simulate the complex interactions between flowfield and ablating surfaces, essential for accurate heat shield performance predictions and optimizing the design process. The experimental reconstruction project developed techniques to extract three-dimensional surface recession data from ground test facilities, enabling better validation of material response models. Finally,

the transition and high-speed flow projects address boundary layer instability mechanisms that are pertinent to the design of entry vehicles and provide a more comprehensive understanding of the potential impact of this mechanism on the TPS.

Acknowledgements

The success of this program would not have been possible without the support of our NASA Project Hosts: Bruno Dias, Sergio Fraile-Izquierdo, Krishnan Swaminathan Gopalan, Dirk Ekelschot, Andrea Fagnani, Georgios Bellas Chatzigeorgis, Olivia M. Schroeder, Kaelan B. Hansson, and Amal Sahai.

We acknowledge Analytical Mechanics Associates, Inc. (AMA) for their generous support of this program, with special thanks to David Cornelius, Arnaud Borner, and Aida L. Alhadate for handling the organizational and logistical aspects that allowed the research to proceed smoothly. We thank the Entry Systems Modeling (ESM) project supporting the MSV 2025 program, particularly Justin Haskins (Program Manager) and Aaron Brandis (Principal Investigator). Furthermore, we also thank Mairead Stackpole (TSM branch chief) for her steadfast assistance in facilitating the program.

We gratefully acknowledge the funding support that made this program possible: ONERA, Direction of International Affairs; Analytical Mechanics Associates, Inc.; the Entry Systems Modeling Project from NASA; the CIEDS Project Open NUM DEF from École Polytechnique (AID); the HPC@Maths Initiative from the Fondation de l'École Polytechnique; and the Fondation Mathématiques Jacques Hadamard (FMJH).

We thank the guest lecture speakers who shared their expertise: Brett Cruden, Aaron Brandis, Justin Haskins, Dinesh Prabhu, Don Ellerby, Lauren Abbott, Ilenia Battiato (Stanford University), and Irene M. Gamba (University of Texas at Austin). The on-site day on July 16th provided participants with valuable exposure to NASA facilities. We thank the facility tour guides and NASA personnel who escorted participants through the Electric Arc Shock Tube (EAST), Mission Visualization Research Command Center (MVRCC), and Vertical Motion Simulator (VMS).

About This Technical Memorandum

This technical memorandum compiles the final reports submitted by the MSV participants, each summarizing their research contributions during the four-week program. These reports showcase the diverse range of computational and experimental approaches being developed to advance our understanding of atmospheric entry systems, from pyrolysis modeling and radiative heat transfer to plasma dynamics and hypersonic boundary layer transition.

Porous Media

Theoretical and computational framework for mass conservation in a pyromechanics model.

Project Bordeaux 1

Flora Lahouze^{1,2}

Jean Lachaud^{1,2}

¹Univ. Bordeaux, CNRS, Bordeaux INP, I2M, UMR 5295, F-33400, Talence, France

²Arts et Metiers Institute of Technology, CNRS, Bordeaux INP, I2M, UMR 5295, F-33400 Talence, France

November 20, 2025

Abstract

Thermal protection systems (TPS) used during atmospheric re-entry experience coupled chemical, thermal, and mechanical transformations driven by extreme conditions. Accurately predicting their behavior requires a coupled multiphysics model that captures pyrolysis kinetics, heat and mass transfer, and the deformation of the solid matrix. This work presents a conservative formulation of mass conservation embedded within a pyromechanics model for reactive porous materials. The approach integrates the evolving volume computed by the solid mechanics submodel directly into the mass conservation equation, ensuring consistency between morphological evolution and density changes. This formulation is implemented in the Porous material Analysis Toolbox based on OpenFOAM (PATO), which couples pyrolysis kinetics, heat and mass transfer, and stress and strain distributions. The method is validated on a controlled thermal expansion case, demonstrating mass conservation throughout the transformation. This establishes a foundation for future applications under full pyrolysis conditions, with the objective of integrating a generalized pyrolysis model accounting for competitive reaction mechanisms.

Keywords: modeling, pyrolysis, pyromechanics, mass conservation

1 Introduction

During atmospheric re-entry, thermal protection systems (TPS) are exposed to extreme pressure and temperature conditions. These systems protect the payload by undergoing chemical degradation through pyrolysis and surface recession through ablation, thereby limiting the internal temperature rise. Pyrolysis involves predominantly endothermic chemical reactions that break down the material while simultaneously altering the mass, volume, and both the microscopic and macroscopic structure of the TPS.

Accurately capturing this behavior requires a fully coupled treatment of the chemical processes occurring during pyrolysis, heat and mass transfer phenomena, and the material's structural response. To address this need, a dedicated modeling framework, the Porous material Analysis Toolbox based on OpenFOAM (PATO), has been developed for reactive porous materials (Lachaud and Mansour, 2014; Lachaud, Scoggins, et al., 2017). Historically, PATO has combined pyrolysis

kinetics, mass loss, and heat and mass transfer under the assumption of a constant volume. More recently, a solid mechanics module has been integrated to describe morphological evolution, primarily driven by thermal expansion and pyrolysis shrinkage (Lahouze et al., 2025). This approach, called pyromechanics, allows to compute thermally induced stresses.

The present work introduces a methodology to ensure mass conservation within the pyromechanics model to account for volume change. First, the main hypotheses and notations used to derive the model are recalled, followed by the presentation of the methodology to integrate volume changes into the mass conservation equation along with the solid mechanics submodel. Finally, the approach is validated on a controlled thermal expansion case.

2 Hypotheses and notations

The material used in this study is wood, but the presented model is generalized to any type of material,




Figure 1: Schematic of the REV for wood as a deformable material during pyrolysis.

including composite material. Wood is treated as a porous, anisotropic, and deformable material composed of a solid phase, denoted s , and a gas phase, denoted g . The solid phase is represented as a mixture of sub-components. Macroscopic equations are obtained using the volume averaging method (Whitaker, 1986a; Whitaker, 1986b) considering a Representative Elementary Volume (REV) deforming at the intrinsic average velocity of the solid phase, $\langle \mathbf{v}_s \rangle^s$ (Fig. 1). The volume of the wood particle at the initial time is denoted Ω_0 , and the volume of the REV at the initial time is denoted \mathcal{V}_0 . Similarly, the volumes in the current deformed configuration are denoted Ω and \mathcal{V} . The density $\langle \rho \rangle$ of wood, containing gas and N_p solid sub-components, is given by

$$\langle \rho \rangle = \langle \rho_g \rangle + \langle \rho_s \rangle = \epsilon_g \langle \rho_g \rangle^g + \sum_{i \in [1, N_p]} \epsilon_i \langle \rho_i \rangle^i. \quad (2.1)$$

At the macroscopic scale, a configuration of the material system is described by the set of positions of its material points. The position of a point P in the reference configuration C_0 , taken as the initial configuration at $t = 0$, is denoted by the position vector \mathbf{X} . The position P' of the same point in the current configuration C_t is denoted by the vector position \mathbf{x} (Fig. 2). The same coordinate system is used for both the initial and current configurations. The current position of a point in the domain Ω , initially located at \mathbf{X} in Ω_0 , is given by the mapping function

$$\mathbf{x} = \phi(\mathbf{X}, t). \quad (2.2)$$

Based on this relation (2.2), which links the material coordinates \mathbf{X} to the spatial coordinates \mathbf{x} , several equivalent descriptions of motion may be adopted. In the Lagrangian description, the independent variables are (\mathbf{X}, t) , and the unknowns are the three functions $\phi_i(\mathbf{X}, t)$, $i \in [1, 2, 3]$. Since the initial positions \mathbf{X} are known, the displacement of each material point between times t_0 and t is defined as

$$\langle \mathbf{u}(\mathbf{X}, t) \rangle = \mathbf{x} - \mathbf{X}. \quad (2.3)$$

In the Eulerian description, the independent variables are (\mathbf{x}, t) , and the unknowns are the three components of the velocity of the material point that was




Figure 2: Initial and current configurations of the material domain.

at \mathbf{X} and is located at \mathbf{x} at time t . This velocity is defined as

$$\langle \mathbf{v}_s(\mathbf{x}, t) \rangle^s = \frac{\partial \phi(\mathbf{X}, t)}{\partial t} = \frac{\partial \mathbf{x}}{\partial t}. \quad (2.4)$$

3 Solid mass conservation

3.1 Mass conservation

Following the volume averaging method, mass conservation at the macroscopic scale for the solid phase during pyrolysis is written as

$$\frac{\partial(\epsilon_s \langle \rho_s \rangle^s)}{\partial t} + \nabla \cdot (\epsilon_s \langle \rho_s \rangle^s \langle \mathbf{v}_s \rangle^s) = \pi_{\text{tot}}. \quad (3.1)$$

The rigorous formulation of the total mass loss source term, π_{tot} , requires a detailed modeling of the pyrolysis kinetics. In practice, an empirical approach commonly used in the field of pyrolysis consists in exploiting the results of ThermoGravimetric Analysis (TGA) experiments, which provide measurements of the evolution of a small sample's mass as a function of time and temperature.

The present work relies on a pyrolysis mechanism composed of four parallel reactions, each assigned to a sub-component composing the solid phase. The pyrolysis advancement χ_i of sub-component i is modeled according to the Arrhenius law:

$$\frac{d\chi_i}{dt} = \mathcal{A}_i \exp\left(-\frac{\mathcal{E}_i}{RT}\right)(1 - \chi_i), \quad i \in [1, N_p], \quad (3.2)$$

where \mathcal{A}_i is the pre-exponential factor, \mathcal{E}_i the activation energy, R the perfect gas constant and T the temperature. Pyrolysis advancement χ_i , varying from 0 to 1 is defined as the ratio between the mass loss at time t and the the total mass loss of sub-component i :

$$\chi_i = \frac{m_{0,i} - m_i(t)}{m_{0,i} - m_{f,i}}, \quad (3.3)$$

where $m_{0,i}$ is the initial mass of component i , $m_i(t)$ its mass at time t and $m_{f,i}$ its final mass. Eq. (3.3) is rearranged by dividing by $m_{0,i}$ and by introducing

$\mathcal{F}_i = \frac{m_{0,i} - m_{f,i}}{m_{0,i}}$ to obtain an expression for the total solid mass of the particle:

$$m_s(t) = \sum_{i \in [1, N_p]} m_{0,i} (1 - \mathcal{F}_i \chi_i). \quad (3.4)$$

The projection of this global measurement onto the local scale, where the conservation equations are written, relies on the specific experimental conditions of TGA. This experimental framework, based on the small size of the sample and its uniform temperature, allows its volume Ω to be assimilated to the control volume \mathcal{V} of the REV. Since the model is expressed in terms of density, Eq. (3.4) can be rewritten as

$$\begin{aligned} m_s(t) &= \int_{\mathcal{V}} \epsilon_s \langle \rho_s \rangle^s dV \\ &= \sum_{i \in [1, N_p]} \int_{\mathcal{V}_0} \epsilon_{0,i} \langle \rho_{0,i} \rangle^i (1 - \mathcal{F}_i \chi_i) dV_0. \end{aligned} \quad (3.5)$$

with dV an elemental volume in the current configuration and dV_0 an elemental volume in the initial configuration. The variables \mathcal{F}_i and χ_i are constant quantities within the REV. Eq. (3.5) expresses the transformation from a Lagrangian to an Eulerian description and can be written in the simplified form

$$\int_{\mathcal{V}} g(\mathbf{x}, t) dV = \int_{\mathcal{V}_0} G(\mathbf{X}, t) dV_0. \quad (3.6)$$

To perform the integration, we introduce the deformation gradient tensor $\underline{\underline{\mathbf{F}}}$ defined as

$$\underline{\underline{\mathbf{F}}} = \frac{\partial \mathbf{x}}{\partial \mathbf{X}} = \frac{\partial \langle \mathbf{u} \rangle}{\partial \mathbf{X}} - \underline{\underline{\mathbf{I}}}. \quad (3.7)$$

The change of variables in a volume integral between the current configuration and the reference configuration then relies on the following relation:

$$dV = \det(\underline{\underline{\mathbf{F}}}) dV_0, \quad (3.8)$$

where $\det(\underline{\underline{\mathbf{F}}})$ is the Jacobian of the transformation. Substituting Eq. (3.8) into Eq. (3.6) yields

$$\int_{\mathcal{V}_0} g(\mathbf{x}, t) \det(\underline{\underline{\mathbf{F}}}) dV_0 = \int_{\mathcal{V}_0} G(\mathbf{X}, t) dV_0, \quad (3.9)$$

giving

$$g(\mathbf{x}, t) \det(\underline{\underline{\mathbf{F}}}) = G(\mathbf{X}, t). \quad (3.10)$$

Eq. (3.5) then becomes

$$\epsilon_s \langle \rho_s \rangle^s = \sum_{i \in [1, N_p]} \frac{\epsilon_{0,i} \langle \rho_{0,i} \rangle^i}{\det(\underline{\underline{\mathbf{F}}})} (1 - \mathcal{F}_i \chi_i). \quad (3.11)$$

This equation expresses the Lagrangian form of mass conservation, including the volume change. From an Eulerian point of view, mass conservation is written as

$$\frac{dm_s}{dt} = \frac{d}{dt} \int_{\mathcal{V}} \epsilon_s \langle \rho_s \rangle^s dV. \quad (3.12)$$

Using Reynolds's transport theorem,

$$\frac{dm_s}{dt} = \int_{\mathcal{V}} \left[\frac{\partial (\epsilon_s \langle \rho_s \rangle^s)}{\partial t} + \nabla \cdot (\epsilon_s \langle \rho_s \rangle^s \langle \mathbf{v}_s \rangle^s) \right] dV. \quad (3.13)$$

Integrating Eqs (3.1) and (3.5) in Eq. (3.13) it comes

$$\int_{\mathcal{V}} \pi_{\text{tot}} dV = - \sum_{i \in [1, N_p]} \int_{\mathcal{V}_0} \epsilon_{0,i} \langle \rho_{0,i} \rangle^i \mathcal{F}_i \frac{d\chi_i}{dt} dV_0. \quad (3.14)$$

Applying the relation (3.8), it comes

$$\pi_{\text{tot}} = - \sum_{i \in [1, N_p]} \frac{\epsilon_{0,i} \langle \rho_{0,i} \rangle^i}{\det(\underline{\underline{\mathbf{F}}})} \mathcal{F}_i \frac{d\chi_i}{dt}. \quad (3.15)$$

The volume change term, $\det(\underline{\underline{\mathbf{F}}})$, ensures mass conservation. Its evaluation follows the solid mechanics formulation presented in the next section.

3.2 Momentum conservation

In the context of pyrolysis, it would be necessary to adopt a nonlinear formalism both from the kinematic point of view and in the constitutive law of the material, due to potentially large displacements and a damageable plastic behavior. To circumvent this difficulty, the chosen solution is to retain a linear formalism by assuming that the material response is piecewise linear. Moreover, since the thermal problem is inherently transient, this time dependence must be accounted for when analyzing structural evolutions.

The time t of the thermal problem serves as the increment between successive equilibrium states. Any variable $\Psi(t+dt)$ in the current configuration can then be expressed as

$$\Psi(t+dt) = \Psi(t) + \delta\Psi, \quad (3.16)$$

with $\Psi(t)$ being the variable in the previous configuration and $\delta\Psi$ the increment of Ψ between times t and $t+dt$. The momentum conservation at equilibrium can then be written as

$$\nabla \cdot \langle \underline{\underline{\boldsymbol{\sigma}}}(t+dt) \rangle = \nabla \cdot (\langle \underline{\underline{\boldsymbol{\sigma}}}(t) \rangle + \langle \delta \underline{\underline{\boldsymbol{\sigma}}} \rangle) = \mathbf{0}, \quad (3.17)$$

with $\underline{\underline{\boldsymbol{\sigma}}}$ being the Cauchy stress tensor. Since equilibrium at time t must be satisfied, i.e., $\nabla \cdot \langle \underline{\underline{\boldsymbol{\sigma}}}(t) \rangle = \mathbf{0}$, the equation to solve is

$$\nabla \cdot \langle \delta \underline{\underline{\boldsymbol{\sigma}}} \rangle = \mathbf{0}. \quad (3.18)$$

The constitutive law relating stresses to strains during pyrolysis is formulated by considering that thermal expansion and shrinkage due to mass loss occur in solid regions where the temperature is high and decomposition is ongoing. The resulting strain field is not geometrically compatible with the solid regions that have not yet undergone pyrolysis, leading to a mechanically induced strain tensor. Under the assumption of small

perturbations between equilibrium states, the principle of superposition applies, and the total strain $\langle \delta \underline{\underline{\epsilon}} \rangle$ is assumed to be the sum of three contributions:

$$\langle \delta \underline{\underline{\epsilon}} \rangle = \langle \delta \underline{\underline{\epsilon}}^{el} \rangle + \langle \delta \underline{\underline{\epsilon}}^{th} \rangle + \langle \delta \underline{\underline{\epsilon}}^{sh} \rangle, \quad (3.19)$$

where $\langle \delta \underline{\underline{\epsilon}}^{el} \rangle$, $\langle \delta \underline{\underline{\epsilon}}^{th} \rangle$, and $\langle \delta \underline{\underline{\epsilon}}^{sh} \rangle$ are the elastic strain, thermal expansion, and pyrolysis shrinkage, respectively. The linearized strain tensor is expressed in terms of the displacement $\langle \delta \mathbf{u} \rangle$ as

$$\langle \delta \underline{\underline{\epsilon}} \rangle = \frac{1}{2} (\nabla \langle \delta \mathbf{u} \rangle + (\nabla \langle \delta \mathbf{u} \rangle)^T). \quad (3.20)$$

Considering linear elasticity, stresses are related to strains via Hooke's law:

$$\begin{aligned} \langle \delta \underline{\underline{\sigma}} \rangle &= \underline{\underline{\underline{\underline{C}}}} : \langle \delta \underline{\underline{\epsilon}}^{el} \rangle \\ &= \underline{\underline{\underline{\underline{C}}}} : (\langle \delta \underline{\underline{\epsilon}} \rangle - \langle \delta \underline{\underline{\epsilon}}^{th} \rangle - \langle \delta \underline{\underline{\epsilon}}^{sh} \rangle), \end{aligned} \quad (3.21)$$

where $\underline{\underline{\underline{\underline{C}}}}$ is the apparent fourth-order stiffness tensor. In this work, wood is considered as an orthotropic material with three principal directions (Tangential, Radial, and Longitudinal).

To explicitly account for thermal expansion and pyrolysis shrinkage, the REV volume \mathcal{V} is assumed to be a function of temperature and the total pyrolysis advancement τ defined by

$$\tau = \sum_i \frac{\epsilon_{0,i} \langle \rho_{0,i} \rangle^i \mathcal{F}_i}{\sum_i \epsilon_{0,i} \langle \rho_{0,i} \rangle^i \mathcal{F}_i} (1 - \chi_i). \quad (3.22)$$

Pressure effects are neglected in this work, therefore an infinitesimal volume change can be written as

$$\frac{dV(T, \tau)}{\mathcal{V}} = \frac{1}{\mathcal{V}} \frac{\partial \mathcal{V}}{\partial T} dT + \frac{1}{\mathcal{V}} \frac{\partial \mathcal{V}}{\partial \tau} d\tau. \quad (3.23)$$

In Eq. (3.23), the volumetric thermal expansion coefficient β is defined as

$$\beta(T) = \frac{1}{\mathcal{V}} \frac{\partial \mathcal{V}}{\partial T}, \quad (3.24)$$

and, similarly, the volumetric shrinkage coefficient ζ due to pyrolysis is defined as

$$\zeta(\tau) = \frac{1}{\mathcal{V}} \frac{\partial \mathcal{V}}{\partial \tau}. \quad (3.25)$$

Thermal expansion and shrinkage are anisotropic (Hagge and Bryden, 2002; Kollmann and Côté, 1968). Therefore, thermal expansion and shrinkage tensors, $\underline{\underline{\alpha}}$ and $\underline{\underline{\xi}}$, composed of linear coefficients, are introduced such that

$$\alpha_i = \frac{1}{L_i} \frac{\partial L_i}{\partial T} \quad \text{and} \quad \xi_i = \frac{1}{L_i} \frac{\partial L_i}{\partial \tau}, \quad (3.26)$$

where L_i is the length in direction $i \in [T, R, L]$. Assuming that lengths vary linearly with temperature and pyrolysis advancement between equilibrium states, $\underline{\underline{\alpha}}$

and $\underline{\underline{\xi}}$ are constant. In this framework, the strains associated with these phenomena are written as

$$\delta \underline{\underline{\epsilon}}^{th} = \int_{T(t)}^{T(t+dt)} \underline{\underline{\alpha}} dT = \underline{\underline{\alpha}} \delta T, \quad (3.27)$$

$$\delta \underline{\underline{\epsilon}}^{sh} = \int_{\tau(t)}^{\tau(t+dt)} \underline{\underline{\xi}} d\tau = \underline{\underline{\xi}} \delta \tau. \quad (3.28)$$

Finally, the constitutive law reads

$$\langle \delta \underline{\underline{\sigma}} \rangle = \underline{\underline{\underline{\underline{C}}}} : \left(\frac{1}{2} (\nabla \langle \delta \mathbf{u} \rangle + (\nabla \langle \delta \mathbf{u} \rangle)^T) - \underline{\underline{\alpha}} \langle \delta T \rangle + \underline{\underline{\xi}} \langle \delta \tau \rangle \right). \quad (3.29)$$

4 Results and discussion

The conservative form of the mass conservation equation is tested on a controlled thermal expansion case. It consists of a 1D bar made of pine wood subjected to a fixed temperature of 750 K. The thermal expansion coefficient is chosen to be very large ($\alpha = 4 \times 10^{-3} \text{ K}^{-1}$) in order to clearly visualize the difference between the conservative and non-conservative cases. Fig. 3 shows the comparison between the two cases: without the volume variation correction, the mass increases in order to maintain a constant density.

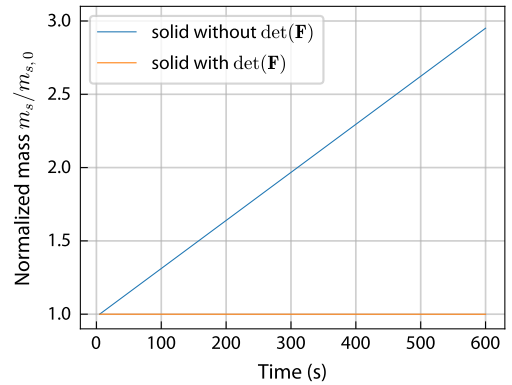


Figure 3: Mass evolution comparison between conservative and non conservative cases.

5 Conclusions

This work presents a rigorous approach to ensuring mass conservation within a pyromechanics modeling framework. More specifically, it integrates the anisotropic volumetric change into the mass conservation equation. This volumetric change is further solved using a solid mechanics framework accounting for elastic strain, thermal expansion and pyrolysis shrinkage. The complete pyromechanics model has been implemented in PATO and the conservative approach has been validated on a simple thermal expansion case.

For future work, a test case under pyrolysis conditions could be considered, but the main objective is

to extend this approach to a pyrolysis model that enables competitive reaction mechanisms accounting for the heating rate.

Acknowledgments

The authors gratefully acknowledge Sergio Fraile and Bruno Dias for their guidance and mentorship during the modeling summer visit at NASA Ames.

References

- Hagge, M. J. and K. M. Bryden (2002). “Modeling the impact of shrinkage on the pyrolysis of dry biomass”. In: *Chemical Engineering Science* 57(14), pp. 2811–2823. ISSN: 0009-2509.
- Kollmann, F. F. P. and W. A. Côté (1968). *Principles of Wood Science and Technology : I Solid Wood*. Springer Berlin Heidelberg, p. 592. ISBN: 9783642879302.
- Lachaud, J. and N. N. Mansour (2014). “Porous-material analysis toolbox based on OpenFOAM and applications”. In: *Journal of Thermophysics and Heat Transfer* 28(2), pp. 191–202.
- Lachaud, J., J.B. Scoggins, T.E. Magin, M.G. Meyer, and N.N. Mansour (2017). “A generic local thermal equilibrium model for porous reactive materials submitted to high temperatures”. In: *International Journal of Heat and Mass Transfer* 108, pp. 1406–1417. ISSN: 0017-9310.
- Lahouze, F., W. Jomaa, C. Métayer, M. Meyer, F. Panerai, and J. Lachaud (2025). “Pyromechanics: A solid mechanics approach to deformation during pyrolysis”. In: *Fuel* 390, p. 134557. ISSN: 0016-2361.
- Whitaker, S. (1986a). “Flow in porous media I: A theoretical derivation of Darcy’s law”. In: *Transport in porous media* 1, pp. 3–25.
- Whitaker, S. (1986b). “Flow in porous media III: deformable media”. In: *Transport in Porous Media* 1, pp. 127–154.

Modeling the pyrolysis of polymer fibers to produce carbon fibers

Emmanuel Lopez^{1,2,3}, Giuseppe Sciume², Jean Lachaud², Célia Mercader³, Lisa Feuillerat¹, and Caroline Petiot¹

¹Airbus

²University of Bordeaux, Institute of Mechanical Engineering

³CANOE

November 1, 2025

Abstract

As part of a project aiming for the development of a process allowing to produce sustainable and cost-effective carbon fibers, this work seeks to understand, model and simulate the continuous charring process used in the industry to transform polymer fibers into carbon fibers. This study is undertaken by adapting an existing pyromechanics model initially developed for wood in the Porous material Analysis Toolbox based on OpenFoam (PATO). The model couples decomposition kinetics, transport phenomena, and solid mechanics via an incremental approach using linear orthotropic elasticity. Preliminary results show the model's ability to simulate the structural evolution of polymer fibers with complex cross sections. This model provides a solid foundation for future improvements, including (visco)plastic behavior.

Keywords: pyromechanics, pyrolysis, PATO, polymer, carbon fiber

1 Introduction

This project is part of a PhD thesis titled "Sustainable and cost effective high performance carbon fiber with lightweight impact," aiming to develop high-performance, lightweight, sustainable, and low-cost carbon fibers. One of the major technological barriers in producing these fibers lies in mastering the carbonization stage [1]. This complex pyrolysis process, carried out under an inert atmosphere and mechanical tension determine the final properties of the fibers [2, 3].

While numerous models describing the chemical kinetics of polymer pyrolysis exist in the literature [4, 5, 6], there is a significant lack of models specifically dedicated to the continuous and controlled process of industrial fiber carbonization. This process is distinguished by the simultaneous application of a precise thermal profile and constant mechanical tension, aimed at orienting the nascent graphitic structure [7, 3].

The main objective of this study is therefore to fill this gap by developing a digital twin focused on the fiber's mechanical response (shrinkage, development of internal stresses) during this critical carbonization stage. Understanding and simulating these phenomena

is essential for optimizing the process, and accelerating the development of competitive carbon fibers [8].

2 Background

Carbon fiber manufacturing is a multi-step process. Initially, a polymer precursor is shaped into continuous filaments using spinning techniques (e.g., wet spinning or dry-jet wet spinning), producing polymer fibers [7, 9]. These fibers are then subjected to an oxidative stabilization step, typically conducted at temperatures between 200°C and 300°C in an oxygen-rich environment [8]. This step is crucial as it transforms the linear polymer chains into a thermally stable ladder-like structure, preventing melting during the subsequent high-temperature carbonization [8].

The carbonization step itself involves heating these precursor fibers in a series of furnaces under an inert atmosphere (usually nitrogen) to very high temperatures, often reaching and exceeding 1000 °C, and potentially up to 1500 °C or more for certain grades [10, 11]. This intense thermal treatment causes the pyrolysis of the polymer: non-carbon elements (hydrogen, oxygen) are eliminated as volatile gases (H₂O, CO, CO₂, CH₄, etc.),

leaving behind an increasingly carbon-rich structure [10, 5]. As the temperature rises, this carbon structure organizes, transitioning from an amorphous state to an increasingly ordered state, with the formation and growth of oriented graphitic domains [10, 11].

A fundamental and distinctive aspect of fiber carbonization is the application of continuous mechanical tension throughout the passage through the furnaces. As pyrolysis progresses, the fiber undergoes significant mass loss and naturally tends to shrink, both axially and radially, due to this mass loss and the atomic rearrangement towards a denser structure [8]. The applied tension primarily counteracts the longitudinal shrinkage. By doing so, it promotes the alignment of the forming carbon chains and graphitic planes along the fiber axis [3, 2]. This tension-induced microstructural orientation is directly responsible for the excellent mechanical properties (high tensile strength and Young's modulus) of carbon fibers in the axial direction [2, 12]. Mastering the complex interplay between the thermal profile (temperatures, heating rates), the inert atmosphere, and the tension level (which may vary along the line) is therefore absolutely crucial for controlling the microstructure and achieving the desired final properties.

3 Methods

The modeling approach adopted combines a pyrolysis reaction model with a mechanical deformation model to simulate the coupled behavior of polymer fibers during thermal treatment under tension.

3.1 Simulation Tool

Simulations are performed using **PATO** (Porous materials Analysis Toolbox based on OpenFOAM). PATO is an open-source toolbox for analyzing multiphase porous reactive materials. It is built upon the OpenFOAM finite volume calculation library and uses Mutation++ library for thermodynamic, transport, and chemistry calculations [13, 6].

3.2 Adaptation of the Pyromechanical Model

This work is built upon the pyromechanical model developed by Flora Lahouze during her PhD for wood pyrolysis [13]. This fundamental model is adapted to the specifics of our use case:

- **Material:** The constitutive model for wood (considered an orthotropic composite of cellulose, hemicellulose, lignin, and water) is simplified to initially represent a cellulose fiber (potentially with residual bound water). This representation may evolve in the future to incorporate other types of polymer precursors. The material properties (thermal, mechanical, kinetic, shrinkage) used initially are derived from Lahouze's model and will need specific calibration for the cellulose fiber under study.

- **Thermal Treatment:** Unlike the "classic" pyrolysis of a bulk sample where temperature results from a complex energy balance (internal conduction, external convection/radiation, reaction heats), we model the continuous industrial carbonization process of fibers by imposing a thermal profile uniformly on the fiber. To simplify and accelerate calculations compared to the real process (several hours with different holding times and heating rates), we apply a linear temperature ramp from ambient temperature (20 °C) up to 1000 °C at a rate of 10 °C/min ($\approx 0.167 \text{ K s}^{-1}$) [14]. The energy conservation equation from the original model is therefore deactivated and replaced by this imposed temperature condition $T(t) = T_0 + \beta t$.

- **Mechanical Conditions:** The mechanical constitutive model is linear elastic (viscoelastic or plastic effects are not taken into account for the moment). To account for large strains an updated lagrangian approach is adopted as the fiber goes under free deformation during the thermal ramp.

3.3 Description of the Model

The proposed model relies on the following components, solved in a coupled manner by PATO using the finite volume method.

3.3.1 Pyrolysis Kinetics and Mass Conservation

The decomposition of cellulose (and potential residual water) is modeled using a simple reaction scheme with N_p pseudo-components (here, $N_p = 2$: cellulose, water). Each component i irreversibly decomposes into gas and a carbonaceous residue ("char") according to a first-order reaction with Arrhenius-type kinetics:

$$\frac{d\chi_i}{dt} = \mathcal{A}_i \exp\left(-\frac{\mathcal{E}_i}{RT(t)}\right) (1 - \chi_i) \quad (3.1)$$

where $\chi_i \in [0, 1]$ is the conversion degree of component i , \mathcal{A}_i is the pre-exponential factor, \mathcal{E}_i is the activation energy, R is the ideal gas constant, and $T(t)$ is the imposed temperature from the ramp.

The total mass loss rate per unit current volume, π_{tot} ($\text{kg m}^{-3} \text{ s}^{-1}$), which represents the gas source, is calculated by summing the contributions from each component:

$$\pi_{\text{tot}}(t) = - \sum_{i=1}^{N_p} \frac{\epsilon_{0,i} \langle \rho_{0,i} \rangle^i}{\det(\mathbf{F})} \mathcal{F}_i \frac{d\chi_i}{dt} \quad (3.2)$$

where $\epsilon_{0,i}$ and $\langle \rho_{0,i} \rangle^i$ are the initial volume fraction and initial intrinsic density of component i , \mathcal{F}_i is the mass fraction of this component that volatilizes (gas yield), and $\det(\mathbf{F})$ is the determinant of the deformation gradient, representing the local volume change.

This rate π_{tot} couples the mass conservation equations for the solid phase (index s) and the gas phase

(index g), averaged over the representative elementary volume (REV):

$$\frac{\partial \langle \rho_g \rangle}{\partial t} + \nabla \cdot (\langle \rho_g \rangle \langle \mathbf{v}_g \rangle^g) = -\pi_{\text{tot}} \quad (3.3)$$

$$\frac{\partial \langle \rho_s \rangle}{\partial t} + \nabla \cdot (\langle \rho_s \rangle \langle \mathbf{v}_s \rangle^s) = \pi_{\text{tot}} \quad (3.4)$$

where $\langle \rho_k \rangle = \epsilon_k \langle \rho_k \rangle^k$ is the apparent density of phase k ($k = s, g$), ϵ_k is its volume fraction, $\langle \rho_k \rangle^k$ its intrinsic density, and $\langle \mathbf{v}_k \rangle^k$ its average intrinsic velocity. Note that $\epsilon_s + \epsilon_g = 1$.

3.3.2 Gas Flow

The flow of gases produced by pyrolysis through the evolving porous structure of the fiber (permeability \mathbf{K} , porosity ϵ_g) is assumed to follow Darcy's law, relating the gas filtration velocity (relative velocity to the solid) to the gas pressure gradient $\nabla \langle p_g \rangle^g$:

$$\epsilon_g (\langle \mathbf{v}_g \rangle^g - \langle \mathbf{v}_s \rangle^s) = -\frac{\mathbf{K}}{\mu_g} \cdot \nabla \langle p_g \rangle^g \quad (3.5)$$

where μ_g is the dynamic viscosity of the gas mixture. The permeability \mathbf{K} is a tensor that can be anisotropic and evolves as the fiber transforms.

3.3.3 Solid Deformation Mechanics (Incremental Approach)

The macroscopic deformation of the solid fiber is considered slow enough to assume quasi-static mechanical equilibrium at each instant:

$$\nabla \cdot \langle \boldsymbol{\sigma} \rangle = 0 \quad (3.6)$$

where $\langle \boldsymbol{\sigma} \rangle$ is the Cauchy stress tensor averaged over the REV. To solve this potentially non-linear problem (large strains, evolving material properties) with a finite volume method, an incremental approach is used. The transformation is discretized into small time steps Δt . At each step, from t to $t + \Delta t$, the displacement increment $\delta \mathbf{u} = \mathbf{u}(t + \Delta t) - \mathbf{u}(t)$ is calculated by solving the equilibrium equation for the stress *increment* $\delta \boldsymbol{\sigma}$:

$$\nabla \cdot \langle \delta \boldsymbol{\sigma} \rangle = 0 \quad (3.7)$$

Within this incremental framework, it is assumed that the total strain increment $\langle \delta \boldsymbol{\epsilon} \rangle = \frac{1}{2}(\nabla \delta \mathbf{u} + (\nabla \delta \mathbf{u})^T)$ can be additively decomposed into a reversible elastic part $\delta \boldsymbol{\epsilon}^{el}$, a part due to thermal expansion/contraction $\delta \boldsymbol{\epsilon}^{th}$, and an irreversible part due to pyrolysis shrinkage $\delta \boldsymbol{\epsilon}^{sh}$:

$$\langle \delta \boldsymbol{\epsilon} \rangle = \langle \delta \boldsymbol{\epsilon}^{el} \rangle + \langle \delta \boldsymbol{\epsilon}^{th} \rangle + \langle \delta \boldsymbol{\epsilon}^{sh} \rangle \quad (3.8)$$

The thermal strain increment is related to the temperature change $\delta T = T(t + \Delta t) - T(t)$ via the thermal expansion tensor $\boldsymbol{\alpha}$: $\delta \boldsymbol{\epsilon}^{th} = \boldsymbol{\alpha} \delta T$. The pyrolysis shrinkage increment is related to the reaction progress (change in total conversion $\delta \tau = \tau(t + \Delta t) - \tau(t)$) via the shrinkage tensor $\boldsymbol{\xi}$: $\delta \boldsymbol{\epsilon}^{sh} = \boldsymbol{\xi} \delta \tau$.

A linear orthotropic elastic constitutive law (adapting the wood assumption) is postulated, relating the

stress increment to the *elastic* strain increment via the stiffness tensor \mathbf{C} (4th order):

$$\langle \delta \boldsymbol{\sigma} \rangle = \mathbf{C}(t) : \langle \delta \boldsymbol{\epsilon}^{el} \rangle = \mathbf{C}(t) : (\langle \delta \boldsymbol{\epsilon} \rangle - \boldsymbol{\alpha}(t) \delta T - \boldsymbol{\xi}(t) \delta \tau) \quad (3.9)$$

The material tensors \mathbf{C} , $\boldsymbol{\alpha}$, and $\boldsymbol{\xi}$, as well as the permeability \mathbf{K} , evolve during pyrolysis. This evolution is modeled by a linear interpolation between the properties of the "virgin" state (φ_v) and those of the "charred" state (φ_c) based on the total pyrolysis conversion $\tau(t) \in [0, 1]$ (where $\tau = 1$ for the virgin state and $\tau = 0$ for the charred state):

$$\varphi(t) = \tau(t) \cdot \varphi_v + (1 - \tau(t)) \cdot \varphi_c \quad (3.10)$$

The total conversion τ is calculated as a weighted average of the individual conversions χ_i based on the initial mass of each component.

After solving Equations (3.7) and (3.9) to find the displacement increment $\delta \mathbf{u}$, the geometric mesh is updated using an ALE (Arbitrary Lagrangian-Eulerian) technique, where the mesh node velocity matches the solid velocity $\langle \mathbf{v}_s \rangle^s = \delta \mathbf{u} / \Delta t$. This allows tracking the large deformations of the fiber.

3.3.4 Geometry and mesh of the fiber

The model geometry is a 50 μm diameter fiber with a hollow, 7-hole geometry (each hole 15 μm in diameter) and a length of 100 μm . Meshing the geometry was performed using the snappyHexMesh utility resulting in a minimum meshsize of 1.5 $\mu\text{m} \times 1.5 \mu\text{m} \times 1.25 \mu\text{m}$ (in the x, y, and z directions, respectively).

4 Results

Figure 1 illustrates the evolution of the stress magnitude (sigma Magnitude) at three key moments of the simulation.

- **At $t = 0$ s (Figure 1a)**, the fiber is in its initial state, before the thermal treatment. The internal stresses are low but not at their lowest (light blue scale), corresponding to the initial resting state.
- **At $t = 1560$ s (Figure 1b)**, the simulation indicates a pyrolysis conversion of approximately 0.5. With the applied heating ramp (10 $^\circ\text{C}/\text{min}$ from 20 $^\circ\text{C}$), this time corresponds to a temperature of about 550 $^\circ\text{C}$. At this moment, we observe that the internal stress in the fiber is maximal. The zones of highest stress (in red, reaching $\approx 3.8 \times 10^6$ Pa) are located mainly in the material between the holes and on the outer periphery of the fiber.
- **At $t = 4200$ s (Figure 1c)**, at the end of the simulation (temperature of 1000 $^\circ\text{C}$), the stress state has returned to a very low level, similar to the initial state (entirely blue).

Visually, the comparison between Figure 1a and Figure 1c also allows us to observe the reduction in the overall diameter of the fiber and the holes, illustrating the radial shrinkage phenomenon simulated by the model.

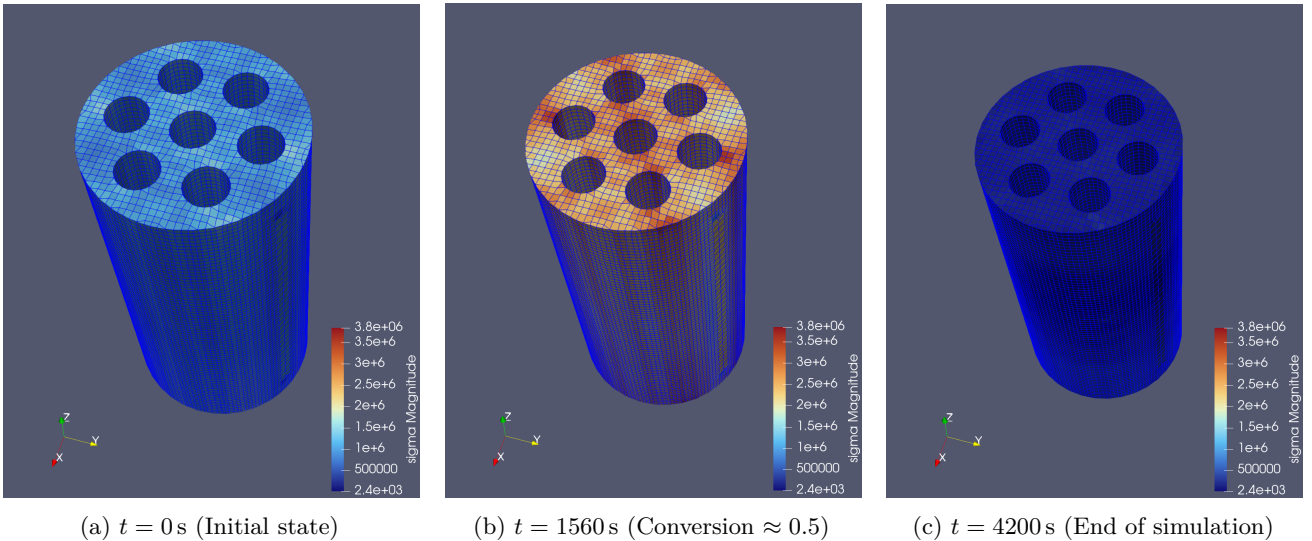


Figure 1: Evolution of the stress magnitude (sigma Magnitude) in the hollow fiber at different times of the pyrolysis simulation (blue: low stress; red: high stress).

5 Discussion

The observation of a significant stress peak at $t = 1560$ s (Figure 1b) is the most striking result. This moment, corresponding to a temperature of approximately 550°C and a reaction conversion of 0.5, is chemically coherent. At this temperature, the polymer structure (cellulose) has undergone its main decomposition and is transitioning predominantly to the "char" state [5, 15]. This transition involves significant atomic reorganization and mass loss, generating shrinkage strains (Equation (3.8)) which, if constrained by the geometry (such as the presence of holes), induce internal stresses.

The fact that the stress is maximal at this precise moment (Figure 1b) suggests that the transition phase to char is the most mechanically critical. Before this, the material is still predominantly polymer; after (Figure 1c), the material is a stable char that no longer undergoes major chemical transformations, and the internal stresses have relaxed (within the framework of this elastic model).

The simulation of stress results clearly identifies zones of high solicitation (between the holes). In a real case, these high-stress zones could correspond to areas of weakening or fragilization during carbonization. The accumulation of internal stress could initiate micro-cracks or structural defects, negatively impacting the final mechanical properties of the fiber [16].

It is, however, essential to recall the limitations of this simulation. The model uses an elastic constitutive law (Equation (3.9)) and simulates free deformation. The actual industrial process, in contrast, occurs under mechanical tension [3, 2], and the material almost certainly exhibits (visco)plastic behavior, especially during the state transition. The stress values obtained here are therefore not quantitatively predictive but serve as powerful qualitative indicators to locate critical points in the geometry during the process.

6 Conclusions

This preliminary work demonstrates the capability of the pyromechanical model, adapted from [13] and implemented in PATO, to qualitatively simulate the carbonization of a polymer fiber with a complex internal geometry. The model successfully couples pyrolysis kinetics, transport phenomena, and solid mechanics.

The simulations have made it possible to:

- Visualize the radial shrinkage of the fiber during the thermal treatment.
- Demonstrate the evolution of internal stresses, with a maximum peak identified at $t = 1560$ s ($\approx 550^\circ\text{C}$).
- Correlate this stress peak to the chemical transition phase to char, which is physically coherent.
- Identify regions of high stress concentration, potential areas of weakening in the real process.

Future work should focus on improving the model to get closer to industrial conditions. The next steps will include calibrating the material parameters (kinetic, mechanical) specific to the fiber under study, integrating more complex mechanical behavior (viscoelasticity or plasticity), and ultimately, simulating the process under axial mechanical tension, to move closer to the targeted digital twin of the carbonization process.

Acknowledgments

The author thanks the organizers of the Modeling Summer Visit 2025 at NASA Ames, as well as mentors and collaborators from the University of Bordeaux, the Von Karman Institute, Polytechnique, and NASA for their guidance. Appreciation is also extended to Dr. Flora Lahouze for the pyromechanical model and to program participants for fruitful discussions.

References

- [1] Hamid Khayyam, Mino Naebe, Alireza Bab-Hadiashar, Farshid Jamshidi, Quanxiang Li, Stephen Atkiss, Derek Buckmaster, and Bronwyn Fox. Stochastic optimization models for energy management in carbonization process of carbon fiber production. *Applied Energy*, 158: 643–655, November 2015. ISSN 03062619. doi: 10.1016/j.apenergy.2015.08.008.
- [2] C. Wilms, G. Seide, and T. Gries. The relationship between process technology, structure development and fibre properties in modern carbon fibre production. *Chemical Engineering Transactions*, 32, 2013. doi: 10.3303/CET1332269.
- [3] Bradley A. Newcomb. Processing, structure, and properties of carbon fibers. *Composites: Part A*, December 2016. doi: 10.1016/j.compositesa.2016.10.018.
- [4] Jing Li, Junhui Gong, and Stanislav I. Stoliarov. Development of pyrolysis models for charring polymers. *Polymer Degradation and Stability*, 115:138–152, May 2015. ISSN 01413910. doi: 10.1016/j.polymdegradstab.2015.03.003.
- [5] Paulo Debiagi, Veronica Piazza, Marco Papagni, Alessandra Beretta, Alessio Frassoldati, and Tiziano Faravelli. Cellulose pyrolysis kinetic model: Detailed description of volatile species. *Proceedings of the Combustion Institute*, 40(1-4):105651, 2024. ISSN 15407489. doi: 10.1016/j.proci.2024.105651.
- [6] Jean Lachaud. *Pyrolysis models and simulation tools across communities : comparative review and common challenges*. Habilitation thesis, University of Bordeaux, Institute of Mechanical Engineering, Talence, France, April 2025.
- [7] Prabhakar V. Gulgunje, Bradley A. Newcomb, Kishor Gupta, Han Gi Chae, Thomas K. Tsotsis, and Satish Kumar. Low-density and high-modulus carbon fibers from polyacrylonitrile with honeycomb structure. *Carbon*, 95:710–714, December 2015. ISSN 00086223. doi: 10.1016/j.carbon.2015.08.097.
- [8] Nishar Hameed, Jordan Sharp, Srinivas Nunna, Claudia Creighton, Kevin Magniez, P. Jyotishkumar, Nisa V. Salim, and Bronwyn Fox. Structural transformation of polyacrylonitrile fibers during stabilization and low temperature carbonization. *Polymer Degradation and Stability*, 128:39–45, June 2016. ISSN 01413910. doi: 10.1016/j.polymdegradstab.2016.02.029.
- [9] Jian Wang, Kaihuang Zheng, Xiaofeng Hu, Zhibin Long, and Kangcheng Chen. Application of a Twin-Screw Extruder with an Ultra-High Length-to-Diameter Ratio in the Dry-Jet Wet Spinning Process of Polyacrylonitrile Nascent Fibers. *Industrial & Engineering Chemistry Research*, 62(21): 8477–8488, May 2023. ISSN 0888-5885, 1520-5045. doi: 10.1021/acs.iecr.3c00509.
- [10] Quan-sheng Ma, Ai-jun Gao, Yuan-jian Tong, and Zuo-guang Zhang. The densification mechanism of polyacrylonitrile carbon fibers during carbonization. *New Carbon Materials*, 31(5): 550–554, October 2016. ISSN 18725805. doi: 10.1016/S1872-5805(16)60031-8.
- [11] A.E. Lewandowska, C. Soutis, L. Savage, and S.J. Eichhorn. Carbon fibres with ordered graphitic-like aggregate structures from a regenerated cellulose fibre precursor. *Composites Science and Technology*, 116:50–57, September 2015. ISSN 02663538. doi: 10.1016/j.compscitech.2015.05.009.
- [12] T. Köhler, R. Brüll, F. Pursche, J. Langgartner, G. Seide, and T. Gries. High strength and low weight hollow carbon fibres. In *IOP Conference Series: Materials Science and Engineering*, volume 254, page 042017, October 2017. doi: 10.1088/1757-899X/254/4/042017.
- [13] Flora Lahouze. *Développement d’un modèle pyromécanique appliqué à la pyrolyse de biomasse*. Phd thesis, Université de Bordeaux, Bordeaux, France, October 2025.
- [14] Yunhao Liang, Michael E. Ries, and Peter J. Hine. Pyrolysis activation energy of cellulosic fibres investigated by a method derived from the first order global model. *Carbohydrate Polymers*, 305:120518, April 2023. ISSN 01448617. doi: 10.1016/j.carbpol.2022.120518.
- [15] Joanna D’Antoni, Amine Ben-Abdelwahed, Cé-cile Gaborieau, Audrey Duphil, Wahbi Jomaa, and Jean Lachaud. Biomass pyrolysis: thermodynamic parameters review and determination through TGA. In *37th International Conference on Efficiency, Cost, Optimization, Simulation and Environmental Impact of Energy Systems (ECOS 2024)*, pages 2042–2053, Rhodes, Greece, 2024. ECOS 2024. ISBN 979-8-3313-0766-0. doi: 10.52202/077185-0175.
- [16] Narayan Shirolkar, Adam Maffe, Edward DiLoreto, Pedro J. Arias-Monje, Mingxuan Lu, Jyotsna Ramachandran, Prabhakar Gulgunje, Kishor Gupta, Jin Gyu Park, Kuo-Chih Shih, Mohammad Hamza Kirmani, Andrew Sharits, Dhriti Nepal, Mu-Ping Nieh, Richard Liang, Thomas Tsotsis, and Satish Kumar. Multichannel hollow carbon fibers: Processing, structure, and properties. *Carbon*, 174:730–740, April 2021. ISSN 00086223. doi: 10.1016/j.carbon.2020.11.077.

Integration of a Compressible Flow Solver into PATO for Strong Fluid–Porous Coupling with ALE Interface Motion Applied to Pore-Scale Pyrolysis

Abderrahman M’jikou^{1,2}

Jean Lachaud^{1,2}

¹Univ. Bordeaux, CNRS, Bordeaux INP, I2M, UMR 5295, F-33400, Talence, France

²Arts et Metiers Institute of Technology, CNRS, Bordeaux INP, I2M, UMR 5295, F-33400 Talence, France

December 4, 2025

Abstract

Porous ablative thermal protection systems (TPS) experience coupled processes of pyrolysis, gas release, and skeleton shrinkage under hypersonic aerothermal loading. The architecture of the Porous material Analysis Toolbox based on OpenFoam (PATO) is designed to allow the coupling of material response (PATO-material class) to external flows (PATO-fluid class). This work develops a strongly coupled framework that integrates the compressible solver `rhoPimpleFoam` in PATO. The approach enforces conservative mass, momentum, and energy exchange at the moving porous/gas interface under a local thermal equilibrium (LTE) assumption for the porous medium and a compressible Navier–Stokes description for the gas. Arbitrary Lagrangian–Eulerian (ALE) mesh motion transfers pyrolysis-induced displacement to the fluid mesh; Laplacian smoothing preserves mesh quality and conformality at the receding interface. The methodology is demonstrated on a pore-scale geometry extracted from high-resolution micro-CT of wood ($\sim 120 \mu\text{m}$ cube). The results reproduce the expected qualitative behavior: inward progression of the pyrolysis front, skeleton recession and morphology change, gas percolation, and blowing that perturbs the near-surface flow. These findings establish the feasibility and robustness of the integration of `rhoPimpleFoam` in PATO to study multiphysics processes in realistic microstructures.

Keywords: Arbitrary Lagrangian–Eulerian (ALE); Strong Coupling; Compressible Navier–Stokes; Porous Media.

1 Introduction

Thermal protection systems (TPS) based on porous ablators undergo pyrolysis, gas release, and skeleton shrinkage under severe aerothermal loads. In this work, we integrate a compressible OpenFOAM [1] fluid solver (`rhoPimpleFoam`) within the Porous material Analysis Toolbox based on OpenFoam (PATO) [2], enabling (i) consistent interface conditions for mass, momentum, and energy, and (ii) Arbitrary Lagrangian–Eulerian (ALE) mesh motion [3] that propagates pyrolysis-driven deformation from the porous domain to the fluid domain. The resulting framework supports conjugate fluid–porous interactions with a moving, deforming interface.

Contributions. (i) A strong-coupling strategy between `rhoPimpleFoam` and the material classes of PATO with interface fields exchanged at each time

step; (ii) an ALE mapping that transfers incremental solid displacement to the fluid mesh with robust Laplacian mesh smoothing; (iii) a consistent set of interface boundary conditions compatible with the local thermal equilibrium model (LTE) in the porous medium [4] and compressible Navier–Stokes in the gas.

2 Background

The architecture of PATO is designed to allow the coupling of material response (material class) to external flows (fluid class) [5]. The material class allows for solving macroscopic porous-media balance equations (mass, momentum, and energy) under local thermal equilibrium (LTE), coupled to solid mechanics and finite-rate pyrolysis kinetics [4, 6]. In this work, OpenFOAM’s `rhoPimpleFoam` solver was integrated as a flow solver; it implements the transient compressible Navier–Stokes

equations with a total-enthalpy/energy formulation using the PIMPLE algorithm [7]. In carbon–phenolic and lignocellulosic ablators, gas produced by pyrolysis percolates through the shrinking skeleton; the resulting interface motion modifies the outer flow and the surface heat/mass transfer [8]. Capturing these interactions motivates a tightly coupled multiphysics algorithm.

3 Methods

We solve two coupled regions: (i) a porous, reactive, deforming solid–gas mixture (PATO-material class) and (ii) an external compressible gas domain (rhoPimpleFoam integrated as a PATO-fluid class). Each region advances its own governing equations; coupling is enforced via physically consistent interface conditions and ALE mesh motion that transmits the porous-domain deformation into the fluid mesh. The fluid solver is integrated within PATO’s run-time control to ensure synchronized time stepping, field exchange, and mesh updates.

3.1 Equations

Since there are two regions, different sets of equations are solved in each region. All equations below are those used in the present work and are accompanied by concise definitions to avoid ambiguity.

Porous medium. *Pyrolysis kinetics:*

$$\frac{\partial \chi_i}{\partial t} = A_i \exp\left(-\frac{\mathcal{E}_i}{R\langle T \rangle}\right) (1 - \chi_i), \quad \forall i \in [1, N_p], \quad (3.1)$$

where N_p is the number of solid pseudo-components; A_i and \mathcal{E}_i are, respectively, the pre-exponential factor and activation energy; χ_i is the conversion (advancement) of component i ; R is the gas constant; and $\langle T \rangle$ is the macroscopic temperature under LTE.

Mass conservation. Gas phase:

$$\frac{\partial}{\partial t} (\varepsilon_g \langle \rho_g \rangle^g) + \nabla \cdot (\varepsilon_g \langle \rho_g \rangle^g \langle \mathbf{v}_g \rangle^g) = \langle \Pi \rangle. \quad (3.2)$$

Here, ε_g is the gas volume fraction (porosity), $\langle \rho_g \rangle^g$ and $\langle \mathbf{v}_g \rangle^g$ are intrinsic gas-phase density and velocity.

Solid phase:

$$\frac{\partial (\varepsilon_s \langle \rho_s \rangle^s)}{\partial t} + \nabla \cdot (\varepsilon_s \langle \rho_s \rangle^s \langle \mathbf{v}_s \rangle^s) = -\langle \Pi \rangle, \quad (3.3)$$

where ε_s is the solid volume fraction, and $\langle \rho_s \rangle^s$, $\langle \mathbf{v}_s \rangle^s$ are intrinsic solid density and velocity.

Pyrolysis mass-loss term:

$$\langle \Pi \rangle = \sum_{i \in [1, N_p]} \frac{\varepsilon_{0,i} \langle \rho_{0,i} \rangle^i}{\det(\underline{\mathbf{F}})} \mathcal{F}_i \frac{d\chi_i}{dt}, \quad (3.4)$$

where $\varepsilon_{0,i}$ and $\langle \rho_{0,i} \rangle^i$ are initial solid volume fraction and virgin density of component i , \mathcal{F}_i is its mass-loss fraction, and $\underline{\mathbf{F}}$ is the deformation gradient [6].

Momentum conservation. Gas phase (Darcy–Biot form [9]):

$$\langle \mathbf{v}_g \rangle^g = - \left(\frac{1}{\varepsilon_g \mu_g} \underline{\mathbf{K}} \right) \cdot \nabla \langle p_g \rangle^g + \langle \mathbf{v}_s \rangle^s, \quad (3.5)$$

with gas viscosity μ_g , intrinsic gas pressure $\langle p_g \rangle^g$, and permeability tensor $\underline{\mathbf{K}}$.

Solid phase (quasi-static linearized form) [6]:

$$\nabla \cdot \langle \delta \underline{\boldsymbol{\sigma}} \rangle = \mathbf{0}, \quad (3.6)$$

and

$$\langle \delta \underline{\boldsymbol{\sigma}} \rangle = \underline{\mathbf{C}} : \left(\frac{1}{2} (\nabla \langle \delta \mathbf{u} \rangle + (\nabla \langle \delta \mathbf{u} \rangle)^T) - \underline{\boldsymbol{\alpha}} \langle \delta T \rangle + \underline{\boldsymbol{\xi}} \langle \delta \tau \rangle \right). \quad (3.7)$$

Here, $\underline{\mathbf{C}}$ is the fourth-order stiffness tensor, $\underline{\boldsymbol{\alpha}}$ the thermal expansion tensor, $\underline{\boldsymbol{\xi}}$ the pyrolysis-shrinkage tensor, $\delta \mathbf{u}$ the incremental displacement, and τ the overall pyrolysis advancement:

$$\tau = \sum_i \frac{\varepsilon_{0,i} \langle \rho_{0,i} \rangle^i \mathcal{F}_i}{\sum_i \varepsilon_{0,i} \langle \rho_{0,i} \rangle^i \mathcal{F}_i} (1 - \chi_i). \quad (3.8)$$

Energy conservation under LTE. We assume $\langle T_g \rangle = \langle T_s \rangle = \langle T \rangle$. The energy equation for both phases reads:

$$\begin{aligned} & \sum_{i \in [1, N_p]} \partial_t (\varepsilon_i \langle \rho_i \rangle^i \langle h_i \rangle^i) + \partial_t (\varepsilon_g \langle \rho_g \rangle^g \langle e_g \rangle^g) \\ & = \begin{cases} \nabla \cdot (\underline{\mathbf{k}} \nabla \langle T \rangle) \\ + \langle \mathbf{v}_g \rangle \cdot \nabla \langle p_g \rangle^g \\ - \nabla \cdot (\varepsilon_g \langle \rho_g \rangle^g \langle h_g \rangle^g \langle \mathbf{v}_g \rangle^g) \\ - \nabla \cdot \sum_{i \in [1, N_p]} \varepsilon_i \langle \rho_i \rangle^i \langle h_i \rangle^i \langle \mathbf{v}_s \rangle^s \end{cases} \end{aligned} \quad (3.9)$$

where $\underline{\mathbf{k}}$ is the effective conductivity tensor (solid conduction, gas conduction, radiative transfer, and dispersion), $\langle h_i \rangle^i$ are solid pseudo-component enthalpies, and e_g is the gas-phase internal energy.

Gas region. In the outer gas domain, we solve the compressible Navier–Stokes equations and the conservative total-energy equation in enthalpy form:

Mass conservation.

The compressible continuity equation for the gas reads

$$\frac{\partial \rho_g}{\partial t} + \nabla \cdot (\rho_g \mathbf{v}_g) = 0, \quad (3.10)$$

where ρ_g is the gas density and \mathbf{v}_g the gas velocity.

Momentum conservation.

The balance of linear momentum is written in conservative form as

$$\frac{\partial (\rho_g \mathbf{v}_g)}{\partial t} + \nabla \cdot (\rho_g (\mathbf{v}_g \otimes \mathbf{v}_g) + p \mathbf{I}) - \nabla \cdot \boldsymbol{\tau}_g = \rho_g \mathbf{g}, \quad (3.11)$$

with p the pressure, \mathbf{I} the identity tensor, $\boldsymbol{\tau}_g$ the viscous stress tensor, and \mathbf{g} the gravitational acceleration.

Total energy (conservative enthalpy form).
The total-energy equation is expressed as

$$\begin{aligned} \frac{\partial(\rho_g(e + K))}{\partial t} + \nabla \cdot (\rho_g \mathbf{v}_g (h + K)) \\ = -\nabla \cdot (k_g \nabla T) + \nabla \cdot (\boldsymbol{\tau}_g \cdot \mathbf{v}_g) + \rho_g \mathbf{g} \cdot \mathbf{v}_g, \end{aligned} \quad (3.12)$$

where e is the specific internal energy, $K = \frac{1}{2} \|\mathbf{v}_g\|^2$ the specific kinetic energy, h the specific enthalpy, T the gas temperature, and k_g the gas thermal conductivity.

3.2 Boundary conditions at the porous/gas interface

The coupling between the two regions is enforced through the following interface conditions applied on Γ , with the unit normal \mathbf{n} directed from the porous medium toward the gas region, consistent with previous formulations [5].

Temperature continuity:

$$[T]_g = [\langle T \rangle]_p, \quad \text{at } \Gamma \quad (3.13)$$

Conductive heat-flux continuity:

$$[-k_g(\nabla T) \cdot \mathbf{n}]_g = [-\underline{\mathbf{k}}(\nabla \langle T \rangle) \cdot \mathbf{n}]_p, \quad \text{at } \Gamma \quad (3.14)$$

Pressure continuity:

$$[p]_g = [\langle p_g \rangle^g]_p, \quad \text{at } \Gamma \quad (3.15)$$

Gas velocity condition (moving interface):

$$[\mathbf{v}_g - \mathbf{u}_{\text{mesh}}]_g = [\varepsilon_g(\langle \mathbf{v}_g \rangle^g \cdot \mathbf{n})\mathbf{n}]_p, \quad \text{at } \Gamma \quad (3.16)$$

These boundary conditions ensure consistent mass, momentum, and heat transfer between the gas and the porous material while preserving thermodynamic compatibility under LTE assumptions [10].

3.3 Moving mesh strategy

Due to pyrolysis, the porous medium deforms and the porous mesh moves following the incremental displacement field $\delta \mathbf{u}$ computed in the solid region using the ALE method. To ensure that the fluid mesh remains consistent with this motion and that no gap appears between the two regions, the following procedure is applied at every time step:

1. Solve the solid momentum conservation equation in the porous region to obtain the incremental displacement field $\delta \mathbf{u}$ caused by pyrolysis.
2. Move the porous mesh by applying this incremental displacement to each cell in the porous region.
3. The motion of the porous mesh creates a temporary gap between the porous and gas meshes; the same incremental displacement is therefore mapped to the fluid-side interface cells so that these boundary cells move and the two meshes remain connected.

4. This interface deformation may generate highly skewed cells in the fluid domain; thus, a Laplacian mesh-smoothing equation is solved in the gas region to propagate the displacement throughout the entire fluid mesh and restore the overall mesh quality.
5. As a result, both the porous and fluid meshes move consistently without gaps, while preserving acceptable cell shape and mesh integrity in both regions.

This strategy ensures that the interface motion remains geometrically consistent and that the fluid and porous meshes remain conformal, allowing the interface boundary conditions to properly enforce mass, momentum, and energy conservation across the evolving interface [3].

4 Results: Pore-scale simulation of pyrolysis

We demonstrate the strong fluid–porous coupling on a pore-scale configuration obtained from a high-resolution micro-CT scan of wood (voxel size $\approx 0.3 \mu\text{m}$). From the segmented dataset, a cubic sub-volume of characteristic size $\approx 120 \mu\text{m}$ was extracted, representing the solid carbonaceous skeleton and its interconnected pore space. At this microscopic scale, the volume includes only a limited number of pores, which facilitates detailed resolution of the fluid–solid interactions during pyrolysis while keeping the computational cost tractable given the fine mesh required and the number of coupled governing equations.

Although pyrolysis is simulated here at the pore scale, the coupling methodology developed in this work remains the same as that used for macroscopic porous ablators. To apply it in this microscopic context, we treat the wood cell walls (the solid microstructure) as a porous medium in the sense of PATO’s-material class governing equations, while the “external” gas region corresponds simply to the gas occupying the pore cavities. In other words, the porous–gas interface appears inside the microstructure itself: the solid skeleton evolves according to the porous-medium model, and the surrounding pore gas plays the role of the external fluid domain. This interpretation allows the full coupling strategy—interface mass-, momentum-, and energy-consistent exchange together with ALE-driven mesh motion—to be exercised directly on a real micro-CT geometry.

Geometry and boundary conditions. The computational domain is a cube with the micro-CT solid skeleton centered inside. At the gas inlet, a uniform velocity u_{inlet} is prescribed with zero-gradient pressure. At the outlet, atmospheric pressure is imposed with zero-gradient velocity. No-slip velocity and zero-gradient pressure are applied on lateral walls. A temperature of 800 K is imposed at the inlet and lateral walls; the interface conditions of Section 3.2 are applied at the




Figure 1: Pore-scale computational domain used for the coupled fluid–porous pyrolysis simulation. The solid structure corresponds to a $\approx 120 \mu\text{m}$ wood microstructural segment reconstructed from high-resolution micro-CT imaging (voxel size $\approx 0.3 \mu\text{m}$), showing the cellular anatomy and internal pore network. This segmented solid skeleton is embedded within a cubic gas domain where the external flow and thermal loads are applied.

porous/gas boundary. Figure 1 illustrates the domain and boundary conditions.

Model parameters. All porous-material thermal, mechanical, transport, and kinetic parameters used in this study are taken directly from the reference dataset in [6], where detailed descriptions, temperature dependences, and tabulated values are provided. This ensures full consistency with experimentally calibrated properties for wood pyrolysis while avoiding unnecessary duplication here. Readers interested in the complete parameter set are referred to [6].

Field evolution. Figure 2 shows solid skeleton deformation over time as a function of the overall pyrolysis advancement τ (with $\tau = 1$ for virgin and $\tau = 0$ for char). Figure 3 shows gas-phase streamlines highlighting pyrolysis-blowing emanating from the receding skeleton and its effect on the external flow.

5 Discussion

The present results show that the coupled approach captures the essential behaviors expected during ablation. As the temperature rises, the pyrolysis front advances inward and the solid skeleton progressively shrinks, leading to a retreat of the porous/gas interface. The gas released inside the material finds its way through the pore network and exits into the outer flow, where it locally disturbs the streamlines (Figs. 2 and 3). This blowing effect near the surface, although qualitatively demonstrated here, is known to influence heat and momentum transfer in ablative thermal protection systems.

The physical models used in this initial study—local thermal equilibrium within the porous medium, a Darcy–Biot formulation for gas transport, and a compressible Navier–Stokes description for the external flow—remain reasonable for the small and highly resistive pore structures studied here. The ALE mesh motion combined with Laplacian smoothing proved effective in maintaining a good mesh quality while the interface deformed, which was essential to preserve conservation at the boundary.

There are, however, clear limitations that need to be addressed. The micro-CT volume is very small ($\sim 120 \mu\text{m}$), which limits the representativeness of the pore network. Material properties were taken from a reference dataset and not specifically calibrated for this specimen. In addition, the current setup neglects multi-species chemistry and radiation, which can become important at higher temperatures. These limitations define the next steps for this effort, together with a broader validation campaign against experimental micro-CT and thermogravimetric data.

6 Conclusions

We presented a strong material-flow coupling that enforces conservative interface conditions under LTE and uses ALE mesh motion to transmit pyrolysis-induced deformation to the gas domain. On a micro-CT wood geometry, the framework produced physically consistent skeleton recession and pyrolysis-driven blowing that alters the outer flow. Next steps include specimen-specific calibration of kinetic/transport laws, temperature- and conversion-dependent permeability/porosity, multi-species gas with interfacial chemistry, and scale-up with adaptive meshing to support predictive TPS simulation.

Acknowledgments

This work was supported in part by the French National Research Agency (ANR) under project ANR-22-CE51-0014-01.

The author would like to thank the organizers of the Modeling Summer Visit 2025 program at NASA Ames Research Center for the opportunity to conduct this work. Additional thanks are extended to mentors and colleagues at the University of Bordeaux, I2M, VKI, École Polytechnique, and NASA for their guidance, feedback, and support throughout the project.

References

- [1] H. Weller, G. Tabor, H. Jasak, C. Fureby, A tensorial approach to computational continuum mechanics using object-oriented techniques, *Computers in Physics* 12 (6) (1998) 620–631. doi: 10.1063/1.168744.
- [2] J. Lachaud, N. N. Mansour, Porous-material analysis toolbox based on openfoam and applications,




Figure 2: Evolution of the porous skeleton during pyrolysis, shown at several levels of reaction progress τ ($\tau = 1$ corresponding to virgin material and $\tau = 0$ to fully charred material). The solid volume progressively shrinks and recedes as the material decomposes, altering pore morphology and gas pathways. Colors represent the local value of τ , highlighting the advancement of the thermochemical decomposition front through the micro-CT derived wood structure.




Figure 3: Velocity streamlines in the gas domain surrounding the deforming porous skeleton, colored by velocity magnitude. Gas generated by pyrolysis percolates through the micro-porous network, exits at the interface, and induces blowing into the external flow.

- [7] T. Holzmann, Mathematics, Numerics, Derivations and OpenFOAM®, release 7.0 Edition, Holzmann CFD, 2019. doi:10.13140/RG.2.2.27193.36960.
- [8] J. Lachaud, HDR Jean Lachaud - Modèles de pyrolyse et outils de simulation dans les différentes communautés, Ph.D. thesis (2024).
- [9] M. A. Biot, Theory of elasticity and consolidation for a porous anisotropic solid, *Journal of Applied Physics* 26 (2) (1955) 182–185. doi:10.1063/1.1721956.
- [10] A. Martin, H. Zhang, K. A. Tagavi, An introduction to the derivation of surface balance equations without the excruciating pain, *International Journal of Heat and Mass Transfer* 115 (Part A) (2017) 992–999. doi:10.1016/j.ijheatmasstransfer.2017.07.078.

Journal of Thermophysics and Heat Transfer 28 (2) (2014) 191–202. doi:10.2514/1.T4382.

- [3] H. Jasak, Z. Tukovic, Automatic mesh motion for the unstructured finite volume method, *Transactions of FAMENA* 30 (2) (2006) 1–20.
- [4] J. Lachaud, J. Scoggins, T. Magin, M. Meyer, N. Mansour, A generic local thermal equilibrium model for porous reactive materials submitted to high temperatures, *International Journal of Heat and Mass Transfer* 108 (2017) 1406–1417. doi:https://doi.org/10.1016/j.ijheatmasstransfer.2016.11.067.
- [5] H. Scandelli, A. Ahmadi-Senichault, F. Richard, J. Lachaud, Simulation of wood combustion in pato using a detailed pyrolysis model coupled to firefoam, *Applied Sciences* 11 (22) (2021). doi:10.3390/app112210570.
- [6] F. Lahouze, W. Jomaa, C. Métayer, M. Meyer, F. Panerai, J. Lachaud, Pyromechanics: A solid mechanics approach to deformation during pyrolysis, *Fuel* 390 (2025) 134557. doi:https://doi.org/10.1016/j.fuel.2025.134557.

Modelling Radiative Heat Transfer through Cracks in the Char Layer during Wood Pyrolysis

D.C. Dinh¹

F. Richard¹

¹University of Poitiers, Pprime UPR3346 CNRS, ISAE-ENSMA, 1 avenue Clément Ader, BP 40109, 86961 Futuroscope Chasseneuil Cedex, France

duy.cuong.dinh@univ-poitiers.fr, franck.richard@univ-poitiers.fr

December 4, 2025

Abstract

Wood pyrolysis is a fundamental process governing the ignition, flame spread, and degradation of lignocellulosic materials in fire scenarios [1]. Accurate modelling of this process requires capturing the complex interactions between heat transfer, mass transport, and chemical decomposition within the evolving porous char layer [2, 3]. While existing pyrolysis models, such as the NASA-developed PATO (Porous material Analysis Toolbox), primarily account for surface radiation and conduction [4], increasing evidence highlights the significant role of in-depth radiative heat transfer through the char [1, 5, 6, 7]. In this work, we extend the PATO framework to incorporate two complementary radiation mechanisms: (i) radiation absorption within the porous char matrix using a Beer–Lambert law formulation; and (ii) radiation transport through crack networks forming in the char layer. These cracks act as preferential pathways that amplify internal heating, thereby influencing pyrolysis dynamics, flame spread, and structural response. The proposed modelling framework provides a more physically realistic representation of heat transfer during wood pyrolysis, enhancing predictive capabilities for fire safety applications.

Keywords: wood pyrolysis, char cracking, heat transfer, radiation, modeling, cone calorimeter

1 Introduction

Wood pyrolysis governs the thermal degradation, ignition, and fire spread of lignocellulosic materials. Accurate modeling is essential for predicting fire dynamics and improving safety engineering practices [8, 9, 10]. The NASA-developed PATO code [4] resolves coupled thermal degradation and transport phenomena in porous materials, typically accounting for surface radiation and conduction. However, experimental evidence indicates that radiation penetrates the porous char layer, altering temperature distribution and pyrolysis front dynamics [11, 12]. Additionally, crack networks in char serve as preferential pathways for heat and mass transfer, significantly enhancing radiation penetration and accelerating pyrolysis [13]. This work extends PATO by implementing: (i) Beer–Lambert law for radiation penetration through porous char, and (ii) a model for radiation transport through crack networks. This combined approach improves heat transfer representation during wood pyrolysis, enhancing predictive capabilities for fire safety applications.

2 Material and Methods

2.1 Experimental setup

Experiments were performed in a cone calorimeter under nitrogen flow (185 L min^{-1} , $< 1 \text{ vol.}\% \text{ O}_2$) to isolate pyrolysis from oxidation [14]. Heat fluxes of 20 and 50 kW m^{-2} were applied via an electrically heated cone, calibrated with a Schmidt–Boelter gauge at 25 mm below the heater. Monte Carlo simulations revealed radial flux non-uniformities (Figure 1); cylindrical samples (80 mm diameter) were used to ensure uniform heating in the central region.

Wood specimens (80 mm diameter, 21 mm thick, $\sim 10 \text{ wt.}\% \text{ moisture}$) were cut with fibers parallel to the exposed surface to control anisotropic heat transfer and crack formation. Samples were mounted in a calcium silicate holder (120 mm diameter, 80 mm height, thermal conductivity $0.08 \text{ W m}^{-1} \text{ K}^{-1}$ to $0.14 \text{ W m}^{-1} \text{ K}^{-1}$) to minimize lateral heat losses.

Measurements included: (i) continuous mass monitoring for mass loss rate (MLR), and (ii) temperature

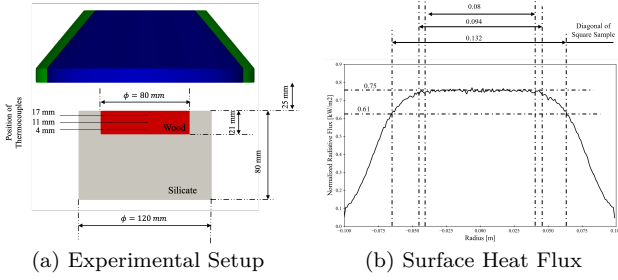


Figure 1: Cone Calorimeter Setup (a) and Surface Heat Flux Distribution (b)

profiles via three K-type thermocouples (0.5 mm diameter) at 4 mm, 11 mm, and 17 mm depths (Figure 1). Mass and temperature tests were performed separately to avoid interference. At 20 kW m^{-2} , minimal cracking occurred; at 50 kW m^{-2} , extensive surface cracking developed.

2.2 Modelling approach

The pyrolysis of wood is modeled as a reactive multi-phase porous medium. We use the volume-averaging method [15, 16, 17, 18] to derive macroscopic conservation laws for heat, mass, and species.

The expressions for conservation equations and property calculations follow the approach described in [3].

2.2.1 Volume averaging and phase definitions

For a variable ψ in phase A (solid S_j or gas G), the phase average and intrinsic phase average are:

$$\begin{aligned} \langle \psi_A \rangle &= \frac{1}{\delta V} \int_{\delta V_A} \psi_A dV, & \langle \psi_A \rangle^A &= \frac{1}{\delta V_A} \int_{\delta V_A} \psi_A dV, \\ \langle \psi_A \rangle &= \epsilon_A \langle \psi_A \rangle^A \end{aligned} \quad (2.1)$$

Here, δV is the total volume of the REV, δV_A is the volume occupied by phase A within the REV, and ϵ_A represents the volume fraction of phase A .


For each solid phase S_j , the volume fraction ϵ_{S_j} , and the porosity ϵ_G for the gas phase G are expressed as:

$$\epsilon_{S_j} = \frac{\delta V_{S_j}}{\delta V}, \quad \epsilon_G = \frac{\delta V_G}{\delta V}, \quad \epsilon_G + \sum_{j=1}^N \epsilon_{S_j} = 1 \quad (2.2)$$

We assume constant total volume during pyrolysis, δV is fixed, only phase fractions ϵ_{S_j} , ϵ_G evolve [4, 3].

2.2.2 Pyrolysis reaction scheme

Wood is decomposed into cellulose, hemicellulose, and lignin, each forming char and gases. The general reaction for each step R_m is:



with $\mu_{j(R_m)}$, $\nu_{G_k(R_m)}$ as stoichiometric coefficients. The Arrhenius rate law is:

$$\langle \dot{\omega}_{S_i} \rangle_{(R_m)} = -k_{(R_m)} \langle \rho_{v0, S_i} \rangle_{(R_m)} \left(\frac{\langle \rho_{S_i} \rangle_{(R_m)}}{\langle \rho_{v0, S_i} \rangle_{(R_m)}} \right)^{n_{(R_m)}} \quad (2.4)$$

where

$$k_{(R_m)} = A_{(R_m)} \exp\left(-\frac{E_a(R_m)}{R\langle T \rangle}\right) \quad (2.5)$$

Product rates:

$$\langle \dot{\omega}_{S_j} \rangle_{(R_m)} = \mu_{j(R_m)} \langle \dot{\omega}_{S_i} \rangle_{(R_m)}, \quad (2.6)$$

$$\langle \dot{\omega}_{G_k} \rangle_{(R_m)} = \nu_{G_k(R_m)} \langle \dot{\omega}_{S_i} \rangle_{(R_m)} \quad (2.7)$$

2.2.3 Mass conservation

Solid species S_i and gas phase G mass balances are:

$$\partial_t \langle \rho_{S_i} \rangle = \langle \dot{\omega}_{S_i} \rangle_{(\text{form})} - \langle \dot{\omega}_{S_i} \rangle_{(\text{dest})} \quad (2.8)$$

$$\partial_t (\epsilon_G \langle \rho_G \rangle^G) + \nabla \cdot (\epsilon_G \langle \rho_G \mathbf{u}_G \rangle^G) = \sum_{m=1}^{N_R} \sum_{k=1}^{N_G} \langle \dot{\omega}_{G_k} \rangle_{(R_m)} \quad (2.9)$$

2.2.4 Momentum conservation: Darcy's law

Gas flow follows Darcy's law:

$$\langle \mathbf{u}_G \rangle^G = -\frac{K}{\mu_G} \nabla \langle p_G \rangle^G \quad (2.10)$$

2.2.5 Gaseous species conservation

Each gas species G_k :

$$\begin{aligned} \partial_t (\epsilon_G \langle \rho_G Y_{G_k} \rangle^G) + \nabla \cdot (\epsilon_G \langle \rho_G Y_{G_k} \mathbf{u}_G \rangle^G) \\ = \nabla \cdot \left(\frac{\epsilon_G}{\eta} \langle \rho_G D_{G_k} \nabla Y_{G_k} \rangle^G \right) + \sum_{m=1}^{N_R} \langle \dot{\omega}_{G_k} \rangle_{(R_m)} \end{aligned} \quad (2.11)$$

2.2.6 Energy conservation

Assuming local thermal equilibrium (LTE), the solid and gas share temperature $\langle T \rangle$. The energy balance is:

$$\begin{aligned} \partial_t (\langle \rho s h_s \rangle + \epsilon_G \langle \rho_G h_G \rangle^G) + \nabla \cdot (\epsilon_G \langle \rho_G h_G \mathbf{u}_G \rangle^G) \\ = \nabla \cdot (k_{\text{eff}} \nabla \langle T \rangle) + \sum_{m=1}^{N_R} \langle \dot{\omega} \rangle_{(m)} \Delta H_{(m)}^0 \\ + \partial_t (\epsilon_G \langle p_G \rangle^G) + \langle \mathbf{u}_G \rangle^G \cdot \nabla (\epsilon_G \langle p_G \rangle^G) + \langle Q_{\text{rad}} \rangle \end{aligned} \quad (2.12)$$

2.2.7 In-depth radiative transfer model


External radiation penetrates the char layer by two mechanisms: (i) absorption in the porous bulk (Beer-Lambert law), and (ii) enhanced penetration via cracks. The total volumetric radiative source is:

$$\langle Q_{\text{rad}} \rangle = \langle Q_{\text{solid}} \rangle + \langle Q_{\text{crack}} \rangle \quad (2.13)$$


Bulk absorption:

$$\langle Q_{\text{solid}} \rangle = \beta I_0 \alpha_{\text{surf}} e^{-\beta x} (1 - f_{\text{crack}}) \quad (2.14)$$

I_0 : incident intensity, α_{surf} : surface absorptivity, β : extinction coefficient, x : depth, f_{crack} : fraction into cracks.



(a) Mass Loss



(b) Mass Loss Rate

Figure 2: Mass Loss and Mass Loss Rate versus Time at 50 kW m⁻²

Crack-mediated radiation:

$$\langle Q_{\text{crack}} \rangle = A_{\text{crack}} I_0 \phi_{\text{crack}}(\tau, x) \alpha_{\tau} e^{-x/\ell_{\text{crack}}} \quad (2.15)$$

$\phi_{\text{crack}}(\tau, x)$: transmissivity, α_{τ} : crack absorptivity, A_{crack} : crack area/volume, ℓ_{crack} : attenuation length.

Effective optical properties:

$$\beta = (1 - \epsilon_{\text{moisture}})[(1 - \tau)\beta_{\text{virgin}} + \tau\beta_{\text{char}}] + \epsilon_{\text{moisture}}\beta_{\text{moisture}}, \quad (2.16)$$

$$\alpha = (1 - \epsilon_{\text{moisture}})[(1 - \tau)\alpha_{\text{virgin}} + \tau\alpha_{\text{char}}] + \epsilon_{\text{moisture}}\alpha_{\text{moisture}} \quad (2.17)$$

where τ : conversion, $\epsilon_{\text{moisture}}$: moisture volume fraction.

Crack parameters:

$$\phi_{\text{crack}}(\tau, x) = \phi_{\text{crack,max}} \max\left\{0, \frac{\tau - \tau_{\text{threshold}}}{1 - \tau_{\text{threshold}}}\right\}^n \times \exp(-x/\ell_{\text{crack}}), \quad (2.18)$$

$$A_{\text{crack}}(\tau, x) = A_{\text{crack,max}} \max\left\{0, \frac{\tau - \tau_{\text{threshold}}}{1 - \tau_{\text{threshold}}}\right\}^n \times \exp(-x/\ell_{\text{crack}}) \quad (2.19)$$

$\tau_{\text{threshold}}$: crack onset, n : sharpness, $\phi_{\text{crack,max}}$, $A_{\text{crack,max}}$: max values.

3 Results and discussion

Some variability was observed between tests due to wood's heterogeneous nature (Figure 2), though overall trends remain consistent and within acceptable limits.

The mass loss curve (Figure 2) shows two phases: rapid decomposition of moisture, hemicellulose, and cellulose (0–1470 s, mass drops to ~30%), followed by slow lignin degradation (1470–3000 s, mass stabilizes at ~25%). The MLR exhibits two peaks (50 s and



(a) At 50 kW m⁻²(b) At 20 kW m⁻²

Figure 3: Temperature along sample depth

1470 s). The first peak corresponds to surface heating and moisture/hemicellulose decomposition, followed by a minimum at 750 s due to char layer formation acting as a thermal barrier. The second peak results from heat accumulation in the insulated bulk, enhanced by crack-mediated radiation at high flux (50 kW m⁻²). After 1470 s, MLR stabilizes at low levels, reflecting slow lignin decomposition.

Temperature profiles (Figure 3) show initial heating slowing near 100 °C due to endothermic moisture evaporation. Between 100–400 °C, the rise is relatively uniform, corresponding to hemicellulose and cellulose pyrolysis. Near 350 °C, temperature acceleration is observed at all depths, attributed to: (i) completion of endothermic reactions, (ii) exothermic lignin decomposition, (iii) thermal insulation at the back face (more pronounced at 11 and 17 mm), and (iv) crack formation enhancing heat transfer (Figure 4). Around 3000 s, temperatures plateau, indicating quasi-thermal equilibrium. Temperature evolution correlates with MLR trends: the initial slowdown matches moisture-driven MLR peak, the sharp rise coincides with major MLR peaks (hemicellulose/cellulose), and the plateau aligns with slow lignin degradation.

At 50 kW m⁻², significant surface cracks form (Figure 4), necessitating in-depth radiation modeling. Simulations without radiation show strong discrepancies near 350 °C and 1000–1500 s (crack formation regime), while 20 kW m⁻² simulations (minimal cracking) match experiments well.

Simulations using the in-depth radiation model (Section 2.2.7) with parameters $A_{\text{crack,max}} = 500$, $\phi_{\text{crack,max}} = 0.05$, $\tau_{\text{threshold}} = 0.3$, $\ell_{\text{crack}} = 0.005$, $n = 1$ show improved temperature predictions, particularly near crack formation conditions, and better MLR agreement (Figure 2). Remaining discrepancies



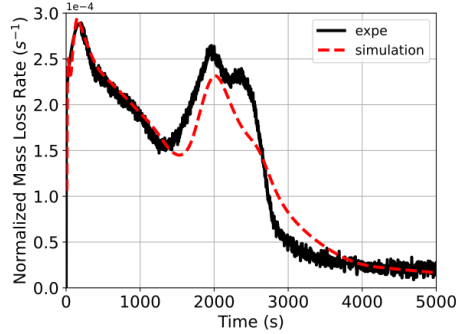


Figure 4: Surface cracks (a,b) and in-depth crack propagation (c) in char layer



(a) Mass Loss Evolution



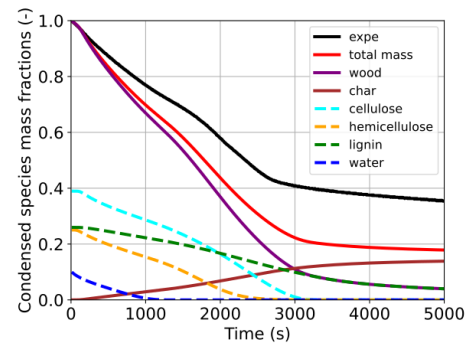
(b) Mass Loss Rate Evolution

Figure 5: Ranzi model predictions vs. Experiment at 20 kW m^{-2}


are attributed to: (i) simplified 1D geometry (3D effects observed in [2, 3]), (ii) simplified chemistry (Figure 5 shows detailed Ranzi's mechanism improves late-stage mass predictions), and (iii) finite-thickness pyrolysis front with concurrent reactions (Figure 7). This demonstrates the model captures main radiation-crack trends; further parameter optimization and 3D simulations with detailed chemistry are needed.

4 Conclusions

This study advances the modelling of wood pyrolysis by extending the PATO code to account for in-depth radiative heat transfer through both porous char and crack networks. The combined Beer–Lambert absorption and crack-mediated radiation transport provide




(a) Parallel Mechanisms

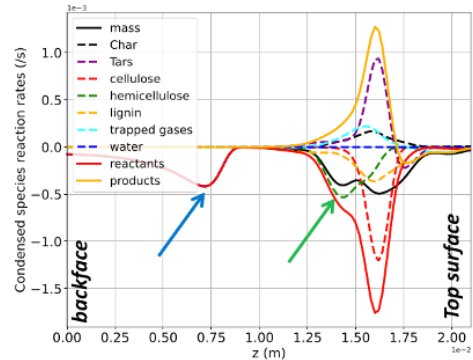


(b) Ranzi Mechanisms

Figure 6: Species Mass Fraction Evolution Comparison



(a) Pyrolysis Front Evolution



(b) Species Reaction Rate

Figure 7: Pyrolysis Front and Species Reaction Rate in Ranzi's Mechanism

a more realistic description of internal heating mechanisms, improving predictions of pyrolysis dynamics and fire behaviour. These developments contribute to more accurate fire safety assessments of wood-based materials and structures under thermal exposure.

Acknowledgments: The authors gratefully acknowledge the support of the NASA Ames Research Center during the 2025 Modeling Summer Visit.

References

- [1] Alastair I Bartlett, Rory M Hadden, and Luke A Bisby. A review of factors affecting the burning behaviour of wood for application to tall timber construction. *Fire technology*, 55(1):1–49, 2019.
- [2] Duy Cuong Dinh, Franck Richard, Benjamin Batiot, and Thomas Rogaueme. A comprehensive model for wood pyrolysis incorporating detailed chemical reactions and porous medium interactions. In *Journal of Physics: Conference Series*, volume 2885, page 012012. IOP Publishing, 2024.
- [3] Duy Cuong Dinh. *Development of a Detailed Approach to Model the Solid Pyrolysis with the Coupling Between Solid and Gases Intra-Pores Phenomena*. PhD thesis, Chasseneuil-du-Poitou, Ecole nationale supérieure de mécanique et d . . . , 2024.
- [4] Jean Lachaud, James B Scoggins, Thierry E Magin, MG Meyer, and Nagi N Mansour. A generic local thermal equilibrium model for porous reactive materials submitted to high temperatures. *International Journal of Heat and Mass Transfer*, 108:1406–1417, 2017.
- [5] Kaiyuan Li, Yanyan Zou, Serge Bourbigot, Jie Ji, and Xianfeng Chen. Pressure effects on morphology of isotropic char layer, shrinkage, cracking and reduced heat transfer of wooden material. *Proceedings of the Combustion Institute*, 38(3):5063–5071, 2021.
- [6] Alekski Rinta-Paavola, Dmitry Sukhomlinov, and Simo Hostikka. Modelling charring and burning of spruce and pine woods during pyrolysis, smoldering and flaming. *Fire Technology*, 59(5):2751–2786, 2023.
- [7] Lingyun Zhang, Yupeng Hu, and Minghai Li. Combined heat transfer mechanisms in the porous char layer formed from the intumescent coatings under fire. *Coatings*, 11(02):200, 2021.
- [8] Colomba Di Blasi. Modeling chemical and physical processes of wood and biomass pyrolysis. *Progress in energy and combustion science*, 34(1):47–90, 2008.
- [9] Eliseo Ranzi, Alberto Cuoci, Tiziano Faravelli, Alessio Frassoldati, Gabriele Migliavacca, Sauro Pierucci, and Samuele Sommariva. Chemical kinetics of biomass pyrolysis. *Energy & Fuels*, 22(6):4292–4300, 2008.
- [10] Kenneth M Bryden, Kenneth W Ragland, and Christopher J Rutland. Modeling thermally thick pyrolysis of wood. *Biomass and Bioenergy*, 22(1):41–53, 2002.
- [11] Michael F Modest and Sandip Mazumder. *Radiative heat transfer*. Academic press, 2021.
- [12] Dominique Baillis and Jean-Francois Sacadura. Thermal radiation properties of dispersed media: theoretical prediction and experimental characterization. In *Radiative Transfer II. Proceedings of the Second International Symposium on Radiation Transfer*. Begel House Inc., 1997.
- [13] Vytenis Babrauskas. Ignition of wood: a review of the state of the art. *Journal of Fire Protection Engineering*, 12(3):163–189, 2002.
- [14] Vytenis Babrauskas, William H Twilley, Marc Janssens, and Shyuitsu Yusa. A cone calorimeter for controlled-atmosphere studies. *Fire and materials*, 16(1):37–43, 1992.
- [15] Stephen Whitaker. Advances in theory of fluid motion in porous media. *Industrial & engineering chemistry*, 61(12):14–28, 1969.
- [16] Stephen Whitaker. The transport equations for multi-phase systems. *Chemical Engineering Science*, 28(1):139–147, 1973.
- [17] William G Gray. A derivation of the equations for multi-phase transport. *Chemical Engineering Science*, 30(2):229–233, 1975.
- [18] Yohan Davit, Christopher G Bell, Helen M Byrne, Lloyd AC Chapman, Laura S Kimpton, Georgina E Lang, Katherine HL Leonard, James M Oliver, Natalie C Pearson, Rebecca J Shipley, et al. Homogenization via formal multiscale asymptotics and volume averaging: How do the two techniques compare? *Advances in Water Resources*, 62:178–206, 2013.

Implementation of competitive pyrolysis mechanisms in PATO

MSV Workshop Report

Joanna D'Antoni^{1,2}

Jean Lachaud^{1,2}

¹Univ. Bordeaux, CNRS, Bordeaux INP, I2M, UMR 5295, F-33400, Talence, France

²Arts et Métiers Institute of Technology, CNRS, Bordeaux INP, I2M, UMR 5295, F-33400

Talence, France

December 4, 2025

Abstract

Pyrolysis, the thermal degradation of material at high temperatures, is crucial in applications from biomass energy to spacecraft thermal protection. Current models in the Porous material Analysis Toolbox based on OpenFoam (PATO) use simplified parallel reactions, limiting predictive accuracy for gas species. This work focuses on implementing complex, competitive pyrolysis mechanisms (e.g., Ranzi) within PATO to address this limitation. We detail the integration methodology using PATO's chemistry framework coupled with the Mutation++ library, adopting a species-based approach with frozen chemistry due to undefined secondary reactions. The implementation was validated in 0D against a benchmarked Python script, showing excellent agreement (final error < 1%). 1D simulations demonstrate the model's capability to capture coupled heat/mass transfer effects and the propagation of the pyrolysis front. This work significantly enhances PATO's predictive capabilities for detailed pyrolysis gas yields.

Keywords: Pyrolysis, PATO, Competitive Mechanisms, Heat Transfer, Mass Transfer.

1 Introduction

Pyrolysis, the thermal degradation of material at high temperatures without oxygen, is a key phenomenon in applications from biomass energy to spacecraft thermal protection systems [1]. Accurate modeling is crucial, and the PATO (Porous material Analysis Toolbox based on Openfoam [9]) platform was developed to simulate the coupled heat and mass transfers in degrading porous media [10].

Current PATO models consist in parallel reactions (typically 3–4) [2]. While computationally efficient, this approach fails to capture the competitive nature of pyrolysis or predict detailed gas species. Indeed, pyrolysis is a complex network of competing reactions (e.g., [3, 4]) where pathways depend heavily on operating conditions. Prior work focused on implementing these mechanisms in a 0D python script [8]. However, this approach ignores geometry and coupled transport phenomena.

The main objective of the work carried out during the MSV workshop was to fill this gap by developing and integrating these competitive reaction mechanisms directly into the PATO framework. This report details the methodology and validation of this implementation.

2 Background and Methods

2.1 Thermochemical model

The implementation relies on OpenFoam's chemistry framework, which solves the species conservation equations for solid and gas phases:

$$\frac{\partial(\epsilon_{s,i}\langle\rho_{s,i}\rangle)}{\partial t} = \mathcal{M}_i\langle\omega_i\rangle, \quad \frac{\partial(\epsilon_g\langle\rho_g\rangle y_i)}{\partial t} = \mathcal{M}_i\langle\omega_i\rangle, \quad (2.1)$$

where ϵ is the volume fraction, ρ the density, y_i the gas mass fraction, \mathcal{M}_i the molar mass, and $\langle\omega_i\rangle$ the net molar production rate. It is computed from the competitive chemical mechanism (e.g., Ranzi [4]) involving N_s species S_i and N_r reactions:

$$\sum_{i \in N_s} \nu_{i,r}^f S_i \xrightleftharpoons[k_r^b]{k_r^f} \sum_{i \in N_s} \nu_{i,r}^b S_i, \quad r \in N_r. \quad (2.2)$$

The total molar production rate of species i is

$$\langle\omega_i\rangle = \sum_{r \in N_r} (\nu_{i,r}^f - \nu_{i,r}^b) \left(k_r^f \prod_{j \in N_s} \langle X_j \rangle^{m_{j,r}^f} - k_r^b \prod_{j \in N_s} \langle X_j \rangle^{m_{j,r}^b} \right), \quad (2.3)$$

where $\langle X_j \rangle$ is the molar density of species j . For pyrolysis, reactions are assumed irreversible ($k_r^b = 0$).

Each forward rate follows an Arrhenius law:

$$k_r^f = A_r \langle T \rangle^{n_r} \exp\left(-\frac{E_r}{R \langle T \rangle}\right), \quad (2.4)$$

where A_r , E_r , and n_r are the pre-exponential factor, activation energy, and temperature exponent.

To manage the thermodynamics, PATO is coupled with the `Mutation++` library. This tool provides two modeling pathways: an Elemental Approach (elements like C, H, O are produced and equilibrated to estimate the species composition) or a Species Approach (the species composition computed within PATO is used directly). The Species Approach can be used with *Finite Rate* or *Frozen* chemistry; in this work, *Frozen* chemistry is adopted due to the absence of defined secondary reactions.

2.2 Development and Validation

The development was performed building on J. Chevalier's prior work on PATO's `MaterialChemistryModel` and `PyrolysisModel` classes [6]. While previous work focused on the `EquilibriumElement` model, this contribution finalized and validated the `SpeciesConservationPyrolysisMaterialChemistryModel`.

Verification was progressive: first, a simplified 5-species mechanism was used to test 0D, 1D, and 2D robustness (including pyromechanical coupling [1]). Secondly, the full Ranzi mechanism [7] was implemented and verified compared with a 0D simulation capable of reproducing Ranzi's results [8]. Finally, a 1D test-case was run and the results are presented here.

3 Results

3.1 0D Verification: TGA Simulation

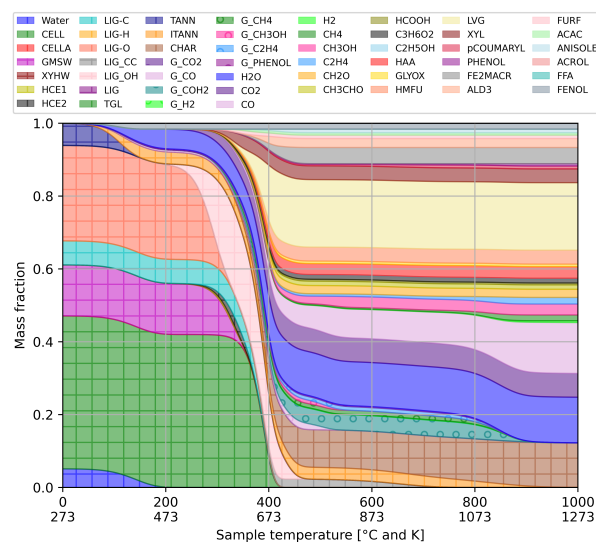
The first step was verifying the `SpeciesConservationPyrolysisMaterialChemistryModel` in a 0D test case (simulating a Thermogravimetric Analysis, TGA). This 0D approach models the material as a single finite-volume cell, assuming a small particle with uniform temperature and composition, neglecting internal heat and mass transfer gradients. A consistent numerical approach was used for the verification. Both the 0D Python benchmark and PATO simulation solve the species equations using an implicit Euler time integration scheme with a fixed time step of 0.5s.

The test case simulated the pyrolysis of pine sawdust (*Pinus radiata*) at 60 K/min using the mechanism and initial composition from Ranzi et al. 2017 [7]. Gas properties (enthalpy and heat capacity) for common species (CO, CO₂, H₂O, CH₄...) are accurately defined, while those for more complex or intermediate species are approximated as Nitrogen for the moment.


This verification was done by comparing PATO results against a standalone 0D Python script, which was previously verified against literature (e.g., Ranzi [7]) and established as a reliable benchmark.

The results are presented in Figure 1. These figures show the evolution of the mass fraction of all species (solids and gases) over time/temperature. Solid species are represented by shaded areas with a hatching pattern, while gaseous species are represented by solid colors (flat shading). Gases considered to be "trapped" (intermediate species like G_COH₂) are indicated by small circular markers overlaid on their colored areas.

An excellent qualitative and quantitative agreement is observed. The evolution of all solid and gaseous species are virtually identical.



(a) 0D Python benchmark.



(b) 0D PATO simulation.

Figure 1: 0D verification of the full Ranzi competitive mechanism at 60 K/min. Comparison between the verified Python script (a) and the new PATO implementation (b).

To quantify this agreement, the absolute error ($\text{MassFraction}_{\text{PATO}} - \text{MassFraction}_{\text{Python}}$) for each species is plotted in Figure 2. The maximum absolute error is around 0.13 (for G_CH₄). Most errors are transient, suggesting slight temporal shifts. The final error stabilizes below 0.01 for all species, confirming the final mass balance accuracy.

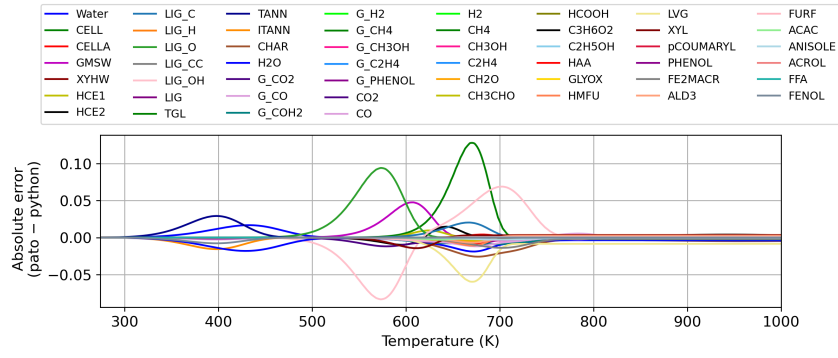


Figure 2: Absolute error (PATO - Python) for each species during the 0D TGA simulation (60 K/min).

3.2 1D Simulation: Coupled Transport and Chemical Effects

A one-dimensional simulation was performed to demonstrate the coupled heat, mass, and species transport phenomena using the `SpeciesConservationPyrolysisMaterialChemistryModel`. The numerical framework relies on an implicit Euler integration scheme with an adaptive time step governed by the pressure evolution. PATO employs a Preconditioned Conjugate Gradient (PCG) solver for solid-phase equations and PBiCGStab for gas-phase transport to ensure stability.

The domain represents a 2.5 cm thick slab discretized with 100 uniformly distributed cells, initially at equilibrium (300 K, 1 atm). Heating is applied via a Dirichlet condition on the top surface (linear ramp of 60 K/min), where gas pressure is fixed at P_{atm} to allow generated volatiles to escape. All other faces are adiabatic and impermeable. Material properties are linearly interpolated between virgin and char states based on degradation. Three probes record the local evolution at the surface (0), mid-depth (1), and near the back face (2).

Thermal and transport results.

Figure 3 gathers the temporal evolution of key variables. The temperature front propagates from the heated surface (>1200 K) toward the back face (≈ 600 K). As heat penetrates, the solid density ρ_s decreases due to devolatilization, confirming strong thermal-chemical coupling. The material transitions from a virgin state (bulk density ≈ 735 kg m $^{-3}$, 51% porosity) to a final char residue of ≈ 95 kg m $^{-3}$, reflecting the specific yield determined by the competitive mechanism.

As volatile species form, internal pressure rises significantly, up to 102800 Pa at the bottom of the material (see Figure 3c). This overpressure drives Darcy flow within the porous matrix, illustrated by the gas velocity accelerating towards the outlet (Figure 3d) and reaching 0.07 m s $^{-1}$ at the end of the simulation. The overall trends demonstrate a physically coherent coupling between heat conduction, pyrolysis kinetics, and gas transport.

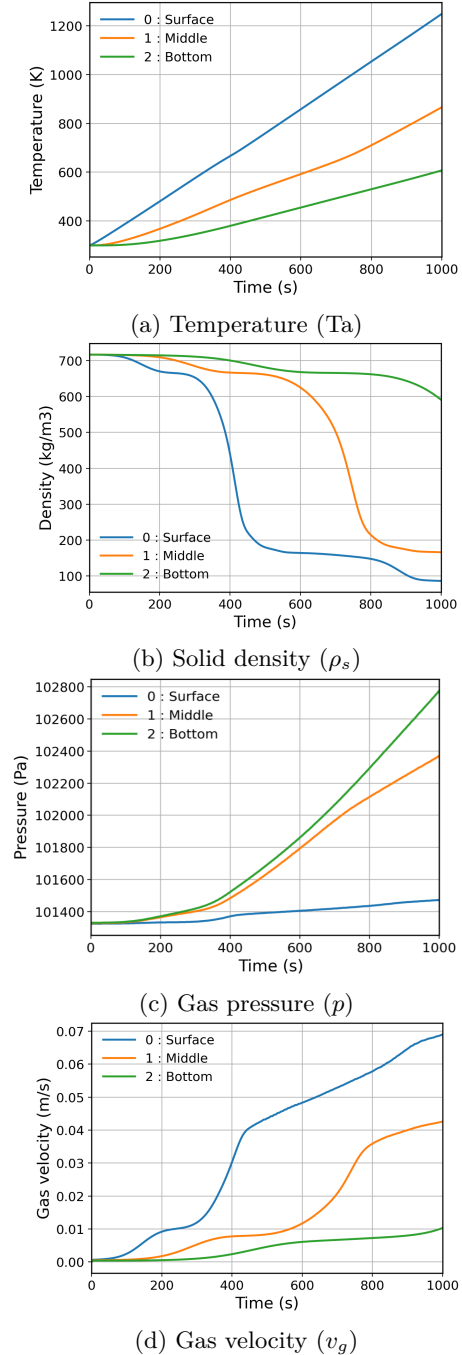


Figure 3: Temporal evolution of key fields at the three probes (surface, middle, back face): temperature, solid density, gas pressure, and gas velocity.

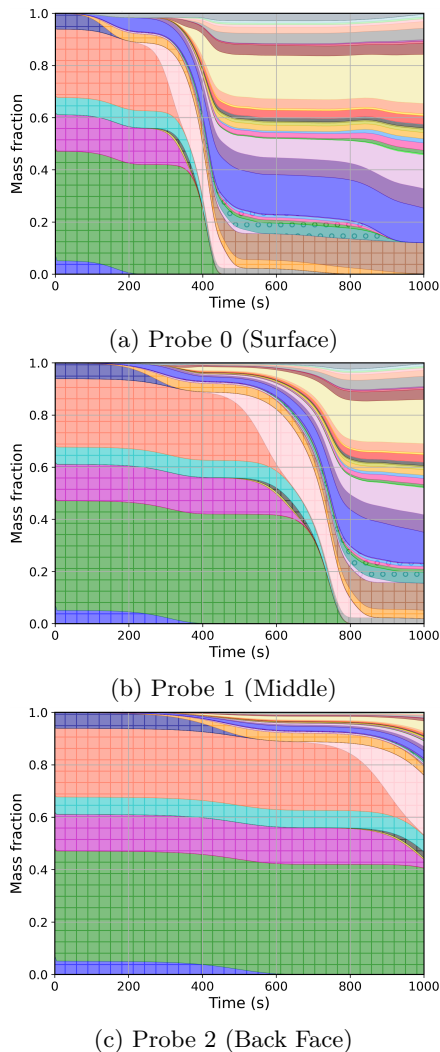


Figure 4: Temporal evolution of species mass fractions at the three probes during the 1D simulation.

Species evolution and chemical pathways.

To analyze the chemical decomposition sequence, Figure 4 presents the time evolution of species mass fractions at the three probes, as obtained from the full Ranzi mechanism.

At the heated surface (probe 0), decomposition begins early as the temperature rises, producing primary volatiles (e.g., water, CO, CO₂) and levoglucosan (LVG). The model correctly captures the transient nature of intermediate lignin and cellulose fragments (LIG, HCE), which appear and are then consumed, converting into char and gases. Capturing these transient species is crucial, as their secondary reactions (not modeled here) can significantly alter the final gas yields. Deeper in the material, at probes 1 and 2, the onset of pyrolysis is significantly delayed due to the slower temperature rise. At probe 1 (mid-depth), major decomposition does not begin until 600 s. At probe 2 (back face), the material barely starts to degrade by 1000 s. This spatial variation of intermediate yields highlights the sensitivity of the competitive reaction pathways to the local temperature history—an effect naturally captured in this coupled 1D framework but absent in purely homogeneous 0D simulations.

4 Conclusions

The work during the MSV 2025 workshop successfully finalized the implementation of the `SpeciesConservationPyrolysisMaterialChemistry Model` within PATO. This represents a significant advancement, enabling the simulation of complex, competitive pyrolysis mechanisms (like Ranzi) using a “Frozen Chemistry” species-based approach, moving PATO beyond simplified parallel reaction schemes.

The model was rigorously verified in 0D against a benchmark, confirming the chemical solver’s accuracy. It also demonstrated physically consistent coupling in 1D, showing robust results at a 60 K/min heating rate. This 1D case proves the model can handle tight coupling between heat conduction, internal gas transport, and multi-species solid-phase chemistry. The implementation significantly enhances PATO’s capability for predicting detailed pyrolysis gas yields, critical inputs for modeling subsequent gas-phase combustion or gas-surface interactions.

Extending the model beyond these initial validations highlighted key challenges. Achieving robust solutions at higher, application-relevant heating rates (e.g., > 1000 K/min) requires further development. These scenarios introduce extreme thermal gradients and very high reaction rates, leading to severe numerical stiffness that challenges solver stability. Moreover, while 2D simulations were robust with simplified mechanisms, the full mechanism’s large species count (over 30 species) causes numerical instabilities and computational overhead, particularly affecting gas-phase convergence.

While gas properties (enthalpy and C_p) for common species (e.g., CO, CO₂, H₂O) are accurately defined, those for more complex intermediates are approximated as Nitrogen. This approximation is non-negligible, as heavier species have different thermochemical properties. It currently limits the model’s ability to capture the full thermal feedback (endothermic/exothermic effects) from pyrolysis, directly influencing the temperature field, reaction rates, and overall energy balance.

Future work will therefore prioritize: (1) resolving numerical stiffness and data-handling issues to ensure robust 2D/3D simulations; (2) improving model stability for high heating rates; (3) implementing realistic gas mixture properties for all species, replacing the N₂ approximation to enable accurate thermal feedback; and (4) conducting thorough experimental validation against gas-phase measurements (e.g., TGA-GC/MS).

Acknowledgments

The author thanks the organizers of the Modeling Summer Visit 2025 program at NASA Ames Research Center, along with the mentors and collaborators from University of Bordeaux, VKI, Polytechnique and NASA. Special thanks to J. Chevalier and G. Bellas Chatzigeorgis, for their guidance and support during this project. This research was funded in part by the French National Research Agency (ANR) under project number ANR-22-CE51-0014-01.

References

- [1] F. Lahouze, W. Jomaa, C. Métayer, M. Meyer, F. Panerai, and J. Lachaud, “Pyromechanics: A solid mechanics approach to deformation during pyrolysis,” *Fuel*, vol. 390, p. 134557, 2025. doi: 10.1016/j.fuel.2025.134557.
- [2] F. Torres-Herrador, V. Leroy, B. Helber, L. Contat-Rodrigo, J. Lachaud, and T. Magin, “Multicomponent pyrolysis model for thermogravimetric analysis of phenolic ablators and lignocellulosic biomass,” *AIAA Journal*, vol. 58, no. 9, pp. 4081–4089, 2020. doi: 10.2514/1.J059423.
- [3] K. A. Trick and T. E. Saliba, “Mechanisms of the pyrolysis of phenolic resin in a carbon/phenolic composite,” *Carbon*, vol. 33, no. 11, pp. 1509–1515, 1995. doi: 10.1016/0008-6223(95)00092-R.
- [4] E. Ranzi, A. Cuoci, T. Faravelli, A. Frassoldati, G. Migliavacca, S. Pierucci, and S. Sommariva, “Chemical kinetics of biomass pyrolysis,” *Energy & Fuels*, vol. 22, no. 6, pp. 4292–4300, 2008. doi: 10.1021/ef800551t.
- [5] A. Anca-Couce, “Reaction mechanisms and multi-scale modelling of lignocellulosic biomass pyrolysis,” *Progress in Energy and Combustion Science*, vol. 53, pp. 41–79, 2016. doi: 10.1016/j.pecs.2015.10.002.
- [6] J. Chevalier, *Modélisation des transferts thermiques dans les matériaux composites pyrolysables*, PhD thesis, University of Bordeaux, 2024.
- [7] E. Ranzi, P. E. A. Debiagi, and A. Frassoldati, “Mathematical modeling of fast biomass pyrolysis and bio-oil formation. Note I: Kinetic mechanism of biomass pyrolysis,” *ACS Sustainable Chemistry and Engineering*, vol. 5, no. 4, pp. 2867–2881, 2017. doi: 10.1021/acssuschemeng.6b03096.
- [8] J. D’Antoni, A. Ben Abdelwahed, W. Jomaa, and J. Lachaud, “Numerical study of biomass pyrolysis: Impact of operating conditions on char yield and quality,” in *26^e Congrès Français de Mécanique*, Metz, France, Aug. 2025. Available at <https://hal.science/hal-05313101>.
- [9] H. G. Weller, G. Tabor, H. Jasak, and C. Fureby, “A tensorial approach to computational continuum mechanics using object-oriented techniques,” *Computers in Physics*, vol. 12, no. 6, pp. 620–631, 1998. doi: 10.1063/1.168744.
- [10] J. Lachaud, *Modèles de pyrolyse et outils de simulation dans les différentes communautés*, Habilitation à Diriger des Recherches (HDR), University of Bordeaux, 2024.

Numerical methods

A two-phase flow model with scale separation for the separated-to-disperse phase transition

Ward Haegeman^{1,2} Giuseppe Orlando² Samuel Kokh³ Marc Massot²
Loïc Gouarin² Pierre Matalon² Sébastien Dubois²

¹DMPE, ONERA,
Palaiseau, France

²CMAP, École polytechnique,
Palaiseau, France

³Maison de la Simulation, CEA,
Gif-sur-Yvette, France

December 4, 2025

Abstract

We present a novel unified two-scale, fully compressible two-phase flow model extending a recently derived isothermal two-scale model to full thermodynamics. In doing so, we identify and overcome key conceptual challenges. The resulting model ensures thermodynamic consistency across scales and broadens the applicability of the original formulation. A numerical proof of concept illustrates the impact of the modeling choices and demonstrates the model's potential.

Keywords: Two-phase flows, Two-scale models, Atomization, Interface regularization, Stationary Action Principle, Irreversible thermodynamics

1 Introduction

Two-phase flows play an important role in a variety of industrial applications, such as the injection and atomization of liquid fuel in combustion chambers [5], or the ablation through spraying of meteoroid melt [3]. The multi-scale nature of such flows prevents the use of Direct Numerical Simulations, as mesh convergence is out of reach [8]. Since different flow regimes can be classified according to the topology of the interface, a large number of models for either separate or disperse regimes have been developed. However, as the aforementioned applications are characterized by the transition or co-existence of different regimes, many efforts have been conducted to couple solvers originally tailored for a particular regime [4, 10]. Still, for such couplings, mass transfer from one solver to another is often mesh-dependent and may induce numerical instabilities, in particular for compressible flows.

Recently, a unified description of separated-disperse phases has been proposed in [12]. It relies on a scale separation wherein the disperse phase is represented by a sub-scale model, independently of the mesh, so as to naturally take into account the multi-scale nature of atomization phenomena. The model is derived by means of the Stationary Action Principle ensuring that it is compatible with the conservation of total energy. Its main novelty is that it allows for a mass transfer

process from the large-scale to the small-scale to be defined as a local dissipative process which regularizes the large-scale interface. This goal is achieved through the modeling of surface tension effects at both scales and by introducing a length-scale threshold. Whenever the interface's local curvature exceeds the prescribed threshold, sub-scale droplets are locally generated thus regularizing the interface. The sub-scale model relies on a kinetic description of the disperse phase and is obtained through a geometric method of moments. Additional moments can be considered in order to account for polydisperse or oscillating droplets [13] thus showing the potentialities of the modeling strategy.

However, for the existing unified two-scale model to be applicable to realistic flows, further modeling efforts are necessary. First and foremost, the model has been derived in the barotropic setting, limiting its application to isothermal flows. Second, a thermodynamic mismatch between scales prevents its extension to non-isothermal flows in a consistent manner and introduces a parasitic term in the inter-scale mass transfer terms.

In this contribution, we analyze the unified two-scale model [12] and identify the key modeling challenges that must be met to tackle realistic flows. This allows for the derivation of a new two-scale model with thermodynamic consistency across scales, supported by a numerical proof of concept. The new modeling approach allows to naturally extend the barotropic

model to complete thermodynamics and thermal non-equilibrium flows, as is briefly outlined in this note.

2 Analysis of the existing unified two-scale model

In order to model the dynamics at different scales, the liquid volume fraction α_{liq} , is divided into a large-scale liquid volume fraction α_ℓ and a small-scale liquid volume fraction α_d representing the disperse phase (e.g. droplets). These satisfy the following saturation constraint $\alpha_\ell + \alpha_d + \alpha_g = 1$, where α_g is the volume fraction of the gas. The effective mass is denoted $m_k = \alpha_k \rho_k$, with $k = \ell, d, g$, and ρ_k the associated density. In addition to its volume fraction, the small-scale liquid droplets d are characterized through their interfacial area density Σ_d , such that $3\alpha_d/\Sigma_d$ represents their Sauter mean radius. In [12], the small-scale droplets are assumed to be incompressible, and it is proposed to rescale the large-scale volume fractions by introducing $\bar{\alpha}_\ell = \alpha_\ell/(1 - \alpha_d)$ and $\bar{\alpha}_g = \alpha_g/(1 - \alpha_d)$, so as to represent the available volume. These satisfy $\bar{\alpha}_\ell + \bar{\alpha}_g = 1$. The unified two-scale model presented in [12] then writes

$$\partial_t m_g + \nabla \cdot (m_g \mathbf{u}) = 0, \quad (2.1a)$$

$$\partial_t m_\ell + \nabla \cdot (m_\ell \mathbf{u}) = R_{m_\ell}, \quad (2.1b)$$

$$\partial_t m_d + \nabla \cdot (m_d \mathbf{u}) = R_{m_d}, \quad (2.1c)$$

$$\partial_t \alpha_d + \nabla \cdot (\alpha_d \mathbf{u}) = R_{\alpha_d}, \quad (2.1d)$$

$$\partial_t \Sigma_d + \nabla \cdot (\Sigma_d \mathbf{u}) = R_{\Sigma_d}, \quad (2.1e)$$

$$D_t \bar{\alpha}_\ell = R_{\bar{\alpha}_\ell}, \quad (2.1f)$$

$$\partial_t (\rho \mathbf{u}) + \nabla \cdot (\rho \mathbf{u} \otimes \mathbf{u} + \bar{p}_{\text{eff}} \mathbf{I} - \sigma \bar{\boldsymbol{\Omega}}) = \mathbf{R}_{\rho \mathbf{u}}, \quad (2.1g)$$

with $D_t \bullet$ denoting the material derivative. Here, eqs. (2.1a–2.1c) represent the mass conservation equations. The source terms on the right-hand side correspond to the inter-scale mass transfer and satisfy $R_{m_\ell} + R_{m_d} = 0$, so as to ensure global mass conservation for each phase. Eqs. (2.1d) and (2.1e) represent the small-scale dynamics. Their source terms are proportional to R_{m_d} . In particular, $R_{\alpha_d} = R_{m_d}/\rho_d$ ensures that $D_t \rho_d = 0$, in agreement with the small-scale incompressibility assumption, while R_{Σ_d} determines the size of the droplets transferred to the small-scale through a parameter $S_{\text{avg}}/m_{\text{avg}}$. The large-scale interface dynamics are governed by (2.1f). As the model assumes velocity equilibrium between scales and phases, the total momentum is governed by eq. (2.1g), where $\rho = m_\ell + m_d + m_g$ represents the total density, $\bar{p}_{\text{eff}} = \bar{\alpha}_\ell p_\ell + \bar{\alpha}_g p_g$ is the effective pressure given by the large-scale mixture pressure, and $\sigma \bar{\boldsymbol{\Omega}}$ is the large-scale capillary tensor which accounts for surface tension following the Continuum Surface Stress model [1], and is given by

$$\sigma \bar{\boldsymbol{\Omega}} = \sigma |\nabla \bar{\alpha}_\ell| \left(\mathbf{I} - \frac{\nabla \bar{\alpha}_\ell \otimes \nabla \bar{\alpha}_\ell}{|\nabla \bar{\alpha}_\ell|^2} \right). \quad (2.2)$$

We refer to the original contribution [12] for details concerning the derivation and the exact expressions of the source terms, we simply recall the main ideas. In

particular, the large-scale volume fraction relaxation source term writes

$$R_{\bar{\alpha}_\ell} = \frac{1}{\varepsilon} \left(p_\ell - p_g - \frac{\sigma}{1 - \alpha_d} \bar{H}_{\text{lim}} \right), \quad (2.3)$$

with $\bar{H}_{\text{lim}} = \max(\bar{H}, H_{\text{max}})$, and where $\bar{H} = -\nabla \cdot (\nabla \bar{\alpha}_\ell / |\nabla \bar{\alpha}_\ell|)$ is the curvature of the large-scale interface, while H_{max} is the inverse of the cut-off length-scale which induces the scale separation. When $\bar{H} \leq H_{\text{max}}$, this source terms represents a relaxation towards the Laplace pressure equilibrium and the remaining source terms are zero. However, when $\bar{H} > H_{\text{max}}$, it triggers the atomization process: the inter-scale mass transfer terms, which are proportional to $\Delta \bar{H} = \bar{H} - \bar{H}_{\text{lim}}$, are activated. These regularize the interface, and the pressure relaxation then yields a Laplace law corresponding to an interface which curvatures do not exceed the threshold H_{max} .

As the model is derived through Stationary Action Principle, it is consistent with the conservation of the total energy, which is the sum of the kinetic energy, the internal energies and the surface energies at both scales, $\sigma \Sigma_d$ and $\sigma |\nabla \bar{\alpha}_\ell|$, as given by the Hamiltonian of the system. In the barotropic case, the total energy serves as mathematical entropy. For this entropy, the source terms are in agreement with the second principle of thermodynamics as they induce a dissipation \mathcal{Q} which writes

$$\begin{aligned} \mathcal{Q} = & -\frac{1 - \alpha_d}{\varepsilon} \left(p_\ell - p_g - \frac{\sigma}{1 - \alpha_d} H_{\text{lim}} \right)^2 \\ & + \frac{\sigma \Delta \bar{H}}{\varepsilon} \left(\sigma \frac{\rho_\ell}{\bar{\alpha}_g} \frac{S_{\text{avg}}}{m_{\text{avg}}} - \frac{\sigma}{1 - \alpha_d} H_{\text{lim}} + (p_\ell - p_g) \bar{\theta} \right) \\ & + \mathbf{u} \cdot \mathbf{R}_{\rho \mathbf{u}}. \end{aligned} \quad (2.4)$$

The first two terms of the second line in eq. (2.4) represent the energy cost of the atomization process, as an amount of interface $\propto H_{\text{lim}}$ is torn off the large-scale ligaments and generates droplet surface $\propto S_{\text{avg}}$. As the atomization process yields an increase of the total surface energy, it must be drawn from other energy reservoirs. This is the role of $\mathbf{R}_{\rho \mathbf{u}}$, chosen so that the two last lines of (2.4) cancel each other and $\mathcal{Q} \leq 0$.

Note that this closure for $\mathbf{R}_{\rho \mathbf{u}}$ causes the model neither to conserve total momentum nor to be Galilean invariant. However, this is only a temporary limitation due to the single velocity assumption, which forces the surface energy increase to be drawn from the kinetic energy. In reality, atomization is triggered by slip effects between the phases; lifting the single-velocity assumption would then allow to draw the energy associated to the atomization process from the energy associated to relative velocity. The modeling of multi-velocity by Stationary Action Principle is an on-going work [7, 6] and is not considered in the present contribution.

Another important aspect of eq. (2.4) is the $\bar{\theta}$ term which is a parasitic inter-scale mass transfer term, related to an enthalpy difference between the liquid at large- and small-scale. Its exact expression can be found in [12], where it is denoted h , and claimed to be negligible. In the barotropic setting, enthalpy differences play

the role of the chemical potential difference to drive phase change. Consequently, this reflects a thermodynamic mismatch between scales, since the inter-scale mass transfer, which is primarily a geometrical process, is partly represented as a phase change in the entropy balance. Moreover, in the two-dimensional air-blasted liquid column simulation presented in [12], the $\bar{\theta}$ term was not accounted for, since deemed negligible, however, for this same test case, it has a clear impact on the inter-scale mass transfer as shown on the evolution of the total sub-scale interfacial area density during the simulation (Figure 1).




Figure 1: Impact of the parasitic inter-scale mass transfer term for the numerical test case of [12].

Finally, the model has been derived in the barotropic setting. Moreover, the incompressibility assumption for the small scale yields some difficulties for the extension to complete thermodynamics. Indeed, in Stationary Action Principle modeling, full thermodynamics is accounted for by means of the internal energy as a function of entropy and density. For an incompressible fluid, however, this is not clearly defined, as internal energy becomes a function of temperature only. Furthermore, the aforementioned thermodynamic mismatch between scales is in part related to the fact that the thermodynamic description of the liquid is not the same at all scales.

However, it turns out that lifting the small-scale incompressibility assumption does not suffice to correct the thermodynamic mismatch. Indeed, if we follow the derivation proposed in [12], but consider instead compressible small-scale inclusions without changing any of the other assumptions, we obtain two Laplace laws, one for each scale. They write

$$p_g - p_\ell + \frac{\sigma}{1 - \alpha_d} \bar{H} = 0, \quad (2.5a)$$

$$p_g - p_d + \frac{2}{3} \sigma \frac{\Sigma_d}{\alpha_d} = -\bar{\alpha}_\ell \frac{\sigma}{1 - \alpha_d} \bar{H}. \quad (2.5b)$$

While eq. (2.5a) is the exact large-scale Laplace law obtained in [12], eq. (2.5b) represents the new small-scale Laplace law. Note, however, that it is physically incoherent as it suggests that droplets are directly affected by large-scale interfaces in their vicinity. Using (2.5a), eq. (2.5b) also writes $\bar{p}_{\text{eff}} - p_d + \frac{2}{3} \sigma \frac{\Sigma_d}{\alpha_d} = 0$. Computations show that the physical inconsistency in eq. (2.5b) is directly related to the introduction of the rescaled volume fractions $\bar{\alpha}$. This rescaling inevitably leads to a "three-fluid" description where the small-scale

liquid may simultaneously mix with the large-scale liquid and the gas, thus the Laplace law between p_d and the large-scale mixture pressure \bar{p}_{eff} .

3 A novel two-scale model

We derive a novel unified two-scale model that extends to full thermodynamics in a thermodynamically consistent manner. First, we present a new two-scale description in the barotropic setting, aiming to develop a fully compressible, thermodynamically consistent framework with the same features as the previous model. We then outline its extension to thermal non-equilibrium flows. We closely follow the methodology adopted in [12], that can be summarized as follows. First, starting from the Lagrangian, which defines the phenomena we account for through their associated energies, we derive a model at pressure equilibrium, which satisfies a Laplace law due to surface tension. Next, we derive its out-of-pressure-equilibrium version. The variational derivation ensures that the two models are connected, as the out-of-pressure-equilibrium model can be equipped with a dissipative relaxation source term that drives it back toward equilibrium consistently with the second law of thermodynamics. Finally, we modify the Laplace law defining the equilibrium by introducing a cut-off length scale to limit the curvature and derive the inter-scale mass transfer terms necessary to preserve the dissipative structure of the modified relaxation process.

First, in order to enforce thermodynamic consistency across scales, we must lift the incompressibility assumption for the small scale. However, as discussed previously, this is insufficient, and to prevent mixing between the droplets and the large-scale liquid, we need to adopt a two-fluid description, where one fluid corresponds to the large-scale liquid while the other fluid is a gas carrying particles — which are the droplets. While this corrects the unphysical small-scale Laplace law, it does not fully remove the thermodynamic mismatch, as a parasitic inter-scale mass transfer, similar to $\bar{\theta}$, also appears. To avoid this and ensure full thermodynamic consistency, we then enforce that liquid particles retain their thermodynamic properties during the atomization process, and therefore we impose $\rho_\ell = \rho_d$. Hence, we discard the small-scale volume fraction equation (2.1d), as it is now determined by $\alpha_d = \alpha_\ell Y_d / Y_\ell$, where Y_d and Y_ℓ are the mass fractions of the liquid at each scale. For the new at-pressure-equilibrium model, we then obtain a single, two-scale Laplace law, which writes

$$\alpha_\ell (p_{\text{liq}} - p_g - \sigma H) + \alpha_d \left(p_{\text{liq}} - p_g - \frac{2}{3} \sigma \frac{\Sigma_d}{\alpha_d} \right) = 0. \quad (3.1)$$

Here, p_{liq} is the pressure of the liquid at all scales. Near the large-scale interface, $\alpha_d \ll 1$, we recover the large-scale Laplace law, while in the gas, $\alpha_\ell \ll 1$, we recover the small-scale Laplace law.

Next, we derive the out-of-equilibrium model so as to obtain the following system in the barotropic setting,

$$\partial_t m_g + \nabla \cdot (m_g \mathbf{u}) = 0, \quad (3.2a)$$

$$\partial_t m_\ell + \nabla \cdot (m_\ell \mathbf{u}) = R_{m_\ell}, \quad (3.2b)$$

$$\partial_t m_d + \nabla \cdot (m_d \mathbf{u}) = R_{m_d}, \quad (3.2c)$$

$$D_t z = R_z, \quad (3.2d)$$

$$D_t \alpha_\ell = R_{\alpha_\ell}, \quad (3.2e)$$

$$\partial_t (\rho \mathbf{u}) + \nabla \cdot (\rho \mathbf{u} \otimes \mathbf{u} + \mathbf{\Pi}) = R_{\rho \mathbf{u}}. \quad (3.2f)$$

Here z is an auxiliary variable used to determine the interfacial area density $\Sigma_d = \rho z / \rho_{\text{liq}}^{2/3}$, while accounting for compressibility [2]. Indeed, in the absence of inter-scale mass transfer, $R_z = 0$, eq. (3.2d) reduces to

$$\partial_t \Sigma_d + \nabla \cdot (\Sigma_d \mathbf{u}) = \frac{2}{3} \Sigma_d \nabla \cdot \mathbf{u} + \frac{2}{3} \Sigma_d \frac{D_t \alpha_d}{\alpha_d}. \quad (3.3)$$

We recover the interfacial area density equation of [11], with an extra term accounting for droplet compressibility. The pressure tensor $\mathbf{\Pi}$ accounts for the phasic pressures as well as the surface tension at both scales.

Owing to the fully compressible description that is thermodynamically consistent across scales, the model may be extended to complete thermodynamics by accounting for the dependence of the Lagrangian on the entropies. Here we extend the model into a two-temperature model for which each phase has its own temperature. The variational principle yields the system in its entropy formulation. We then operate a change of variables to recover the energy form. We consider the phasic energies $\mathcal{E}_g = \mathbf{u}^2/2 + e_g + \sigma(|\nabla \alpha_\ell| + \Sigma_d)/\rho$, and $\mathcal{E}_{\text{liq}} = \mathbf{u}^2/2 + e_{\text{liq}} + \sigma(|\nabla \alpha_\ell| + \Sigma_d)/\rho$. The phasic energy equations then write

$$\begin{aligned} \partial_t (\alpha_g \rho_g \mathcal{E}_g) + \nabla \cdot (\alpha_g \rho_g \mathcal{E}_g \mathbf{u} + \alpha_g \mathbf{\Pi}_g \mathbf{u}) - \mathbf{u} \cdot \mathbf{\Xi} \\ = R_{m_g \mathcal{E}_g}, \end{aligned} \quad (3.4a)$$

$$\begin{aligned} \partial_t (\alpha_{\text{liq}} \rho_{\text{liq}} \mathcal{E}_{\text{liq}}) + \nabla \cdot (\alpha_{\text{liq}} \rho_{\text{liq}} \mathcal{E}_{\text{liq}} \mathbf{u} + \alpha_{\text{liq}} \mathbf{\Pi}_{\text{liq}} \mathbf{u}) + \mathbf{u} \cdot \mathbf{\Xi} \\ = R_{m_{\text{liq}} \mathcal{E}_{\text{liq}}}, \end{aligned} \quad (3.4b)$$

with $\mathbf{\Pi}_g$ and $\mathbf{\Pi}_{\text{liq}}$ the phasic pressure tensors, and $\mathbf{\Xi}$ the kinetic energy exchange term, analogous to that in the single-velocity, two-pressure model [14], but also accounting for the capillary pressures. Similarly as before, the volume fraction relaxation drives the system towards the modified Laplace law, here $R_{\alpha_\ell} = [\alpha_\ell (p_{\text{liq}} - p_g - \sigma H_{\text{lim}}) + \alpha_d (p_{\text{liq}} - p_g - \frac{2}{3} \sigma \frac{\Sigma_d}{\alpha_d})] / \varepsilon$ with $H_{\text{lim}} = \max(H, H_{\text{max}})$. The inter-scale mass transfer terms are chosen so as to obtain the following equation for the total entropy ρs ,

$$\begin{aligned} \partial_t (\rho s) + \nabla \cdot (\rho s \mathbf{u}) = \frac{1}{\varepsilon \alpha_\ell T_{\text{eff}}} \left[\alpha_\ell (p_{\text{liq}} - p_g - \sigma H_{\text{lim}}) \right. \\ \left. + \alpha_d (p_{\text{liq}} - p_g - \frac{2}{3} \sigma \frac{\Sigma_d}{\alpha_d}) \right]^2, \end{aligned} \quad (3.5)$$


where $T_{\text{eff}}^{-1} = Y_g T_g^{-1} + Y_{\text{liq}} T_{\text{liq}}^{-1}$.

Note that the surface energy is a mixture quantity and must therefore be distributed among the phases to define \mathcal{E}_g and \mathcal{E}_{liq} . The choice of how to distribute it is not trivial and affects both the structure of the equations and the source terms, in particular fixing the interfacial pressure in the pressure relaxation source terms. The choice considered here gives an interfacial pressure $p_I = Y_{\text{liq}} p_g + Y_g p_{\text{liq}}$ in agreement with the expression determined in [6].

4 Numerical proof of concept

We provide a numerical proof of concept of the new model, implemented in the open-source library Samurai [9]. We consider the test case presented in [12], and adopt a similar numerical strategy to test our new fully compressible model in the barotropic setting. In particular, we solve (3.2) by considering an operator splitting between the convective fluxes, the surface tension operators and the sources terms. For the convective fluxes, we use a second order MUSCL space discretization and time integration is performed using a second order SSP method. The source terms are treated as instantaneous relaxations and a pseudo-time integration technique is used to preserve admissible states.

Figure 2 shows a comparison between the interface location represented by $|\nabla \alpha_\ell|$, without the two-scale modeling 2a and with the two-scale modeling 2b. By generating small-scale droplets (represented in green), the two-scale model regularizes the interface, and mesh convergence can be obtained for the large-scale features. Without the two-scale approach, the interface keeps spreading and, once it is no longer properly resolved by the mesh, it is no longer closed. The extension of the numerical method to the non-isothermal case is an ongoing work.



(a) without inter-scale mass transfer (b) with inter-scale mass transfer

Figure 2: Interface regularization

5 Conclusions

A novel thermodynamically consistent unified two-scale multi-fluid model has been derived. Inter-scale mass transfer relies on a local dissipative process that regularizes the large-scale interface without discarding information concerning the small-scale features. Preliminary numerical results in the barotropic setting show the model's potential.

Acknowledgments

We thank the NASA hosts for their hospitality. This work has been supported by the Fondation Mathématiques Jacques Hadamard through the Junior Scientific Visibility programme, as well as the Agence Innovation Défense, the CIEDS project OPEN-NUM-DEF, and the HPC@Maths Initiative of the Fondation École polytechnique.

References

- [1] J.U. Brackbill, D.B. Kothe, and C. Zemach. “A continuum method for modeling surface tension”. *Journal of computational physics* 100.2 (1992), pp. 335–354.
- [2] R. Di Battista. “Towards a Unified Eulerian Modeling Framework for Two-Phase Flows : Geometrical Small Scale Phenomena and Associated Flexible Computing Strategies”. PhD thesis. Institut Polytechnique de Paris, 2021.
- [3] O.G. Girin. “A Hydrodynamic Mechanism of Meteor Ablation - The Melt-Spraying Model”. *Astronomy & Astrophysics* 606 (2017), A63.
- [4] F. Granger et al. “Interface Area Density Model for Large-Eddy Simulation of Assisted Atomization in Fiber Regime”. *International Journal of Multiphase Flow* 179 (2024), p. 104927.
- [5] H. Grosshans. “Large eddy simulation of atomizing sprays”. PhD thesis. Lund University, 2013.
- [6] W. Haegeman et al. “An All-Topology Two-Fluid Model for Two-Phase Flows Derived through Hamilton’s Stationary Action Principle”. (Submitted, available on HAL 05249139). 2025.
- [7] W. Haegeman et al. “Modelling of Relative Velocity, Velocity Fluctuations and Their Interactions for Two-Fluid Models by Stationary Action Principle”. (Submitted, available on HAL 05224712). 2025.
- [8] J. C. Hoarau et al. “Direct Numerical Simulation of a Subcritical Coaxial Injection in Fiber Regime Using Sharp Interface Reconstruction”. *International Journal of Multiphase Flow* 180 (2024), p. 104974.
- [9] *Hpc-Maths/Samurai*. HPC@Maths. 2025. URL: <https://github.com/hpc-maths/samurai/>.
- [10] C. Le Touze et al. “A Compressible Two-Phase Flow Framework for Large Eddy Simulations of Liquid-Propellant Rocket Engines”. *Applied Mathematical Modelling* 84 (2020), pp. 265–286.
- [11] D. Lhuillier. “Evolution of the volumetric interfacial area in two-phase mixtures”. *Comptes Rendus. Mécanique* 332.2 (2004), pp. 103–108.
- [12] A. Loison et al. “A unified two-scale gas-liquid multi-fluid model with capillarity and interface regularization through a mass transfer between scales”. *International Journal of Multiphase Flow* 177 (2024), p. 104857.
- [13] A. Loison et al. “Small-Scale Interface Dynamic Modelling Based on the Geometric Method of Moments for a Two-Scale Two-Phase Flow Model with a Disperse Small Scale”. *Journal of Fluid Mechanics* 1003 (2025), A27.
- [14] M. Pelanti and K.-M. Shyue. “A Mixture-Energy-Consistent Six-Equation Two-Phase Numerical Model for Fluids with Interfaces, Cavitation and Evaporation Waves”. *Journal of Computational Physics* 259 (2014), pp. 331–357.

samurai/ponio: two complementary software libraries for the efficient implementation of adaptive space/time numerical schemes

Sébastien Dubois¹ Loïc Gouarin¹ Alexandre Hoffmann¹ Josselin Massot¹
Marc Massot¹ Pierre Matalon¹ Laurent Series¹

¹CMAP, CNRS, École polytechnique, Institut polytechnique de Paris

Abstract

We present two complementary open-source C++ libraries for the numerical simulation of multiscale physical systems. `samurai` provides an efficient adaptive Cartesian meshing framework, while `ponio` offers a wide range of time-integration schemes for ordinary differential equations. Together, they enable the rapid implementation and testing of high-performance, adaptive space-time numerical methods for PDEs. Both libraries emphasize modularity, data locality, and modern C++ design. Their capabilities are illustrated through applications developed during the Modeling Summer Visit 2025, including two-phase flows, rarefied gases, and reactive flows for aerospace applications relevant to NASA.

Keywords: adaptive mesh, AMR, adaptive multiresolution, time-stepping schemes, HPC.

1 Introduction

Numerical simulation of complex physical phenomena requires the implementation of efficient numerical schemes for the discretization of both space and time. The optimal choice of numerical methods depends strongly on the problem under consideration, and the landscape of available approaches continues to evolve with the regular emergence of new algorithms. At the same time, limited computational resources and the demand for ever greater accuracy to resolve fine-scale physical features make adaptive methods essential for the efficient treatment of multiscale problems [9]. These observations highlight two complementary needs: (i) to quickly assess existing methods on complex problems such as the ones that can be found in aerospace engineering, and (ii) to easily prototype and test new numerical schemes. To address them, the pair of open-source libraries `samurai`¹/`ponio`² provides efficient implementations of reference numerical techniques for the fast solution of partial differential equations. `samurai` is a parallel Cartesian mesher designed for mesh adaptation, and `ponio` is an ODE solver (method of lines) that offers a wide range of time-stepping methods readily available to the user. Each library can be used independently,

and `samurai` alone enables the development of coupled space-time schemes with mesh adaptation. Both are implemented in modern C++, emphasizing flexibility and ease of integration. In the following, Section 2 presents the main features of `samurai`, Section 3 describes `ponio`, and Section 4 illustrates how both libraries have been applied in the context of the Modeling Summer Visit 2025 to a range of problems, including two-phase flows, rarefied gases, and combustion models.

2 The samurai library

Adaptive mesh refinement (AMR) [4] and multiresolution analysis (MRA) [12, 7] techniques are essential for resolving steep gradients, shocks, or localized features, while saving computational resources and memory trace in domain regions where the solution is smooth. `samurai` is a C++ library built around a data structure specifically designed to efficiently handle adaptive Cartesian meshes. The originality of `samurai`'s approach hinges on a special data structure enabling both AMR and MRA, improving data contiguity and facilitating the search for neighbours in stencil-based computations.

¹<https://github.com/hpc-maths/samurai>

²<https://github.com/hpc-maths/ponio>

2.1 The interval-based data structure

Most adaptive meshing frameworks rely on a quad-tree/octree structure, individual leaves representing a single cell (cell-based refinement) or a uniform block of cells (block and patch-based refinement). For instance, the meshing libraries `p4est` [6] and `AMREx` [18] use octrees. In contrast, the fundamental building block in `samurai` is the *interval*. An interval represents a sequence of identical, contiguous cells along a Cartesian direction, storing only the indices of the first and last cell and implicitly including all cells in between. This process, applied in the x -direction, enables a first step of compression. In a second step, the y -coordinates associated with these x -intervals are also grouped whenever they are contiguous, yielding a list of y -intervals. The association between each y coordinate and its corresponding list of x -intervals is organized in a manner reminiscent of the *compressed sparse row* (CSR) format widely used for sparse matrices. Additional axes (z , etc.) are handled recursively in the same way, which makes the approach dimension-independent. `samurai` is therefore not limited to 3D, but natively handles any dimension.

At each level of refinement, `samurai` maintains such a multidimensional interval structure, and the complete adaptive mesh is obtained as a hierarchy of these level-wise data structures. This approach drastically reduces memory usage and avoids the recursive pointer traversal typical of tree-based structures.

2.2 Field storage and stencil-based computations

When a data field is attached to the mesh, its memory layout is determined by the numbering of the underlying cells, since the stored values will follow the same order in a contiguous memory space. In classical octree-based approaches, cells are typically indexed using a space-filling curve [16] (e.g., Hilbert or Morton ordering). Such orderings improve spatial locality, which benefits cache efficiency, but they do not guarantee strict contiguity along Cartesian directions. This lack of axis-aligned contiguity can hinder vectorized computations, particularly for stencil-based numerical schemes.

In `samurai`, cells are numbered level by level and, within each level, interval by interval along the x -direction. As a result, data fields are stored contiguously along every x -interval, providing both memory alignment and cache-friendly access. This layout is especially advantageous for stencil computations: for instance, given a two-cell stencil to compute fluxes, once the stencil is anchored on two neighbouring cells (displayed either horizontally or vertically), it can slide along the x -direction to vectorize the computation of multiple fluxes, up until the end of the common interval where the stencil has been anchored.

Stencil-based schemes require, for each cell, quick access to its neighbors. A common strategy is to precompute and store these relationships in a correspondence table. While this allows fast lookup, it comes at the

cost of significant additional memory. In `samurai`, neighbor relations are computed on the fly, without storing explicit connectivity. The interval-based data structure provides a set of simple geometric operations—union, intersection, difference, translation, etc.—that make this possible. For example, translating an x -interval upward yields, in a single operation, the upper neighbors of all its cells. Only the index of the first cell in the translated interval needs to be retrieved from the mesh’s internal storage; the rest follow contiguously.

2.3 Mesh adaptation

The interval-based mesh hierarchy in `samurai` supports both AMR and MRA. Since the MSV projects presented here all rely on MRA, we give below a concise and deliberately simplified overview of this technique. Unlike AMR, which is a heuristic refinement strategy, the MRA approach is driven by wavelet compression. It does not require a user-defined refinement criterion. Instead, refinement and coarsening are guided by an intrinsic regularity test. The principle is as follows: projection and prediction operators are defined so that any local polynomial of degree up to a prescribed order is exactly preserved when projected onto coarser cells and then predicted back to the fine cells. This test is applied to the numerical solution in each cell. The discrepancy between the predicted value and the actual value defines the so-called *detail*. By design, the detail will be small if the solution is smooth, and will grow as large as the regularity of the solution is low. Should a cell be coarsened, the detail quantifies the information that cannot be recovered by the prediction operator. A user-defined tolerance specifies the acceptable loss of information, and cells are refined or coarsened to ensure that all details remain below this threshold. Compared to AMR, which is heuristics-driven, MRA presents the advantage of providing rigorous error control, directly tied to the actual smoothness of the solution.

2.4 Flux-based user interface

`samurai` provides dedicated tools for implementing finite volume schemes. In practice, the user only needs to define how a flux is computed at the interface between two cells, given a chosen stencil size around that interface. The implementation is fully dimension-independent: the same flux definition applies in 1D, 2D, or 3D. Mesh traversal is handled internally by `samurai`, which successively provides the stencil values needed for flux evaluation. As a result, the user can write schemes as if working on a uniform mesh, without having to account for level jumps. When such a jump occurs, `samurai` automatically supplies the appropriate stencil values: values predicted into ghost cells on the coarse side and values from real cells on the fine side of the interface. The following code illustrates the implementation of the Laplace operator (the flux is a gradient) using the classical second-order scheme with a two-cell stencil.

```

[] ( FluxValue<cfg>& flux,
     const StencilData<cfg>& data,
     const StencilValues<cfg>& u )
{
    const int L = 0; // left
    const int R = 1; // right
    double dx = data.cell_length;
    flux = (u[R] - u[L]) / dx;
};

```

Interfaces for both conservative and non-conservative fluxes are provided. In the conservative case, a single flux is computed: the contribution is added to the left cell and subtracted from the right one. In the non-conservative case (useful, *e.g.*, in two-phase flow models), two distinct fluxes (incoming and outgoing) must be computed.

The application of MRA requires that all field values involved in flux computations be predicted at the finest level of resolution via the MRA prediction operator. To balance accuracy and cost, however, the prediction level can be chosen by the user. In either case, the flux function itself remains unchanged—the only difference lies in the stencil values provided by `samurai`.

3 The ponio library

`ponio` is a C++ library providing a collection of ready-to-use numerical integration schemes for ordinary differential equations (ODEs). It integrates first-order systems of the form

$$\dot{u} = f(t, u), \quad (3.1)$$

where t is the time variable, u is the unknown state, and f a temporal flux function.

`ponio` also allows users to extend its capabilities by defining custom Butcher tableaux, and provides Python tools for the automatic analysis of numerical methods.

3.1 Automatic method analysis and code generation

Each Runge–Kutta (RK) method is implemented directly from its Butcher tableau, which serves as the sole input. From this tableau, dedicated and optimized C++ code is automatically generated. User-defined tableaux can be imported in JSON format, triggering an analysis that computes the method’s theoretical properties, such as its stability domain, order of convergence, and, when applicable, its stage order. The data of the 60+ Butcher tableaux already imported is available online³, covering a wide variety of classical and modern RK schemes. The same analysis tools are also employed for more advanced numerical schemes, such as extended-stability and IMEX methods.

3.2 Catalogue of methods

A wide range of Runge–Kutta (RK) methods are available, including explicit RK schemes and their associated

³<https://hpc-maths.github.io/ponio/>

Lawson formulations [13], exponential RK methods, and diagonally implicit Runge–Kutta (DIRK) methods. Adaptive time-stepping is supported through embedded RK schemes, such as the Dormand–Prince family for explicit methods and ESDIRK(5,4) for implicit ones.

Several extended-stability methods are implemented: (i) with an arbitrary number of stages: the Runge–Kutta Chebyshev method RKC2 and the Runge–Kutta Legendre methods [15] of first and second orders (RKL1 and RKL2); (ii) with a dynamically adjusted number of stages: ROCK2 [2] and ROCK4 [1], available with both fixed and adaptive time steps. This kind of method is useful for diffusion problems or other problems with very large negative eigenvalues.

Additionally, `ponio` includes operator-splitting approaches such as Strang [8] and Lie splitting, as well as several (IMEX) schemes, offering a flexible framework for stiff and multiscale problems alike.

3.3 PIROCK: a dedicated scheme for reaction-diffusion problems

`ponio` also provides time-integration methods specifically designed for reaction–diffusion systems, among which the most prominent is the PIROCK [3] scheme.

Reaction–diffusion equations present major challenges for numerical integration, owing to the stiffness of the reaction operator and the spectrum of the diffusion operator. They are commonly treated using operator-splitting techniques, where each subproblem is integrated with a dedicated solver. Although effective, such approaches introduce a splitting error.

To overcome this limitation, the PIROCK method combines an extended-stability explicit scheme (ROCK2) for the diffusion operator—avoiding the solution of large linear systems—with a DIRK scheme for the reaction operator, which only requires solving small, independent, local systems. These two integrators are coupled within an IMEX framework, where the stabilization stages are performed independently of the reaction solve, followed by only four coupling stages.

Compared to the original PIROCK scheme, `ponio` extends the method with a flexible mechanism for error estimation. A naïve computation of the norm on an adaptive mesh tends to overestimate the error in the finest cells; `ponio` corrects this by weighting the error according to cell size and, when necessary, normalizing each field component when their magnitudes differ significantly.

3.4 A versatile user interface

Using C++ templates, `ponio` relies directly on the user’s data structures to store the system state (*i.e.* the data type of u in formula (3.1)). This can range from a single scalar (for scalar-valued problems) to any array or vector type, as long as it supports either basic linear operations or random-access. In the context of MSV 2025, `samurai` field types are typically employed as state containers. The example below illustrates how to define the flux function f for the heat equa-

tion $\dot{u} = \Delta u$ using `samurai`, and how to integrate it with `ponio`'s RKC2 solver configured with 20 stages.

```
// problem construction
auto u_ini = samurai::make_scalar_field
    <double>("u", mesh);
using Field = decltype(u_ini);
auto laplacian = samurai::
    make_laplacian_order2<Field>();
auto f = [&](double t, auto&& u, auto& du)
{
    du = laplacian(u);
};

// ponio solver
auto method = ponio::runge_kutta::
    chebyshev::explicit_rkc2<20>();
double T = 1;
double dt = 1e-2;
auto sol_range = ponio::make_solver_range(
    f, method, u_ini, {0., T}, dt);

// time loop
while (it_sol->time < T)
{
    ++it_sol;
    save(it_sol->state);
}
```

4 Application to MSV projects

`samurai` and `ponio` have been used in MSV projects modeling two-phase flows [11], rarefied gases [5], reactive flows [10] and high-order code coupling [17]. Below, we give details on how the `samurai/ponio` frameworks have been leveraged for the first three projects.

4.1 Two-phase flow with scale separation

The two-phase flow model of [14] is extended to non-isothermal regimes by incorporating complete equations of state and accounting for energy exchanges during inter-scale mass transfer. This ensures total energy conservation and consistent entropy production in high-enthalpy flows such as ablation or fuel atomization.

Numerical discretizations of both models have been implemented in the `samurai` framework with Rusanov and HLLC fluxes. Both diffuse and sharp interface approaches are supported. Recent developments focus on improving numerical robustness and validating the fully compressible model through direct comparison between these two interface representations within the same `samurai` environment.

Note that the default MRA prediction operator is built as a constant linear combination of field values, without enforcing any exterior constraint. In particular, when predicting volumic fractions on the finest level, the resulting values may very well be out of admissible bounds (below 0 or above 1). To solve this issue, a custom-made prediction operator has been implemented.

4.2 Hybrid dense-rarefied flows

The simulation of hybrid dense–rarefied flows is addressed, where the transition between Navier–Stokes and kinetic regimes remains difficult to capture with traditional hybrid Navier–Stokes/direct simulation Monte Carlo methods. This work aims to develop a deterministic Eulerian fluid solver capable of spanning both continuous and rarefied regimes, reducing reliance on heuristic coupling between solvers.

The approach is based on the method of moments, which derives reduced fluid models from the Boltzmann equation by tracking macroscopic moments of the velocity distribution instead of individual molecular states. The main challenges include choosing suitable closure relations and ensuring consistent coupling between different moment hierarchies. A relaxation scheme is designed to smoothly connect these models across flow regions. The resulting 2D solver is implemented using the `samurai` library, leveraging its adaptive mesh refinement and error control capabilities to efficiently resolve local rarefaction gradients.

4.3 Multicomponent reactive flows

This work proposes a dual splitting/IMEX strategy for efficiently solving multicomponent reacting flow equations that exhibit strong stiffness and multiscale coupling in space and time. Traditional explicit schemes fail due to stability constraints, while fully implicit ones are too costly.

The approach combines the advantages of adaptive operator splitting with IMEX Runge–Kutta methods. The convection term is integrated using high-order monotonicity-preserving one-step schemes, whereas the reaction–diffusion subsystem is handled with the PIROCK method (cf. Section 3.3), offering stability and error control. This hybrid technique aims to retain the low computational cost of splitting methods while overcoming their limitations in strongly coupled regimes. It is tested on canonical reaction–diffusion systems (ignition and Belousov–Zhabotinsky problems) and will later be on combustion applications with increasing chemical complexity, including premixed and diffusion flames. The implementation relies on adaptive spatial and temporal discretization with `samurai` and `ponio` to achieve optimal accuracy at minimal cost, targeting performance comparable to or better than existing adaptive-splitting strategies.

5 Conclusion

The development of `samurai` and `ponio` stems from a common goal: to provide researchers with reliable, modular tools for designing adaptive numerical solvers. Both libraries emphasize clarity, composability, and reproducibility, making them valuable testbeds for new numerical ideas. Future developments will focus on tighter coupling between spatial and temporal adaptivity, extended parallelization strategies, and the inclusion of additional physical models.

Acknowledgements

The authors want to thank the NASA hosts for their invitation, and acknowledge the financial support of the CIEDS (project OPEN-NUM-DEF) and the Fondation de l'École polytechnique through the HPC@Maths Initiative.

References

- [1] A. Abdulle. Fourth order Chebyshev methods with recurrence relation. *SIAM Journal on Scientific Computing*, 23(6):2041–2054, 2002.
- [2] A. Abdulle and A. A. Medovikov. Second order Chebyshev methods based on orthogonal polynomials. *Numerische Mathematik*, 2001.
- [3] A. Abdulle and G. Vilmart. PIROCK: A swiss-knife partitioned implicit–explicit orthogonal Runge–Kutta Chebyshev integrator for stiff diffusion–advection–reaction problems with or without noise. *Journal of Computational Physics*, 242:869–888, 2013.
- [4] M. J. Berger and P. Colella. Local adaptive mesh refinement for shock hydrodynamics. *Journal of Computational Physics*, 82(1):64–84, 1989.
- [5] P. Bernigaud, A. Clout, D. Biasone, S. Dubois, A. Hoffmann, and M. Massot. Simulation of hybrid dense-rarefied flows relying on relaxation schemes and a hierarchy of moment methods. Technical report, NASA, 2025.
- [6] C. Burstedde, L. C. Wilcox, and O. Ghattas. p4est: Scalable algorithms for parallel adaptive mesh refinement on forests of octrees. *SIAM Journal on Scientific Computing*, 33(3):1103–1133, 2011.
- [7] A. Cohen, S. Kaber, S. Müller, and M. Postel. Fully adaptive multiresolution finite volume schemes for conservation laws. *Mathematics of Computation*, 72(241):183–225, 2003.
- [8] S. Descombes, M. Duarte, and M. Massot. Operator splitting methods with error estimator and adaptive time-stepping. Application to the simulation of combustion phenomena. In R. Glowinski, S. Osher, and W. Yin, editors, *Splitting Methods in Communication, Imaging, Science, and Engineering*, pages 1–13. Springer International Publishing, 2015.
- [9] M. P. Duarte. *Adaptive numerical methods in time and space for the simulation of multi-scale reaction fronts*. Theses, Ecole Centrale Paris, Dec. 2011.
- [10] L. Gouarin, J. Massot, M. Massot, P. Matalon, L. Series, and C. Tenaud. A dual splitting/ImEx strategy for multicomponent reacting flows with space/time adaptation and error control. Technical report, NASA, 2025.
- [11] W. Haegeman, G. Orlando, S. Kokh, M. Massot, L. Gouarin, P. Matalon, and S. Dubois. A two-phase flow model with scale separation for the separated-to-disperse phase transition. Technical report, NASA, 2025.
- [12] A. Harten. Multiresolution algorithms for the numerical solution of hyperbolic conservation laws. *Communications on Pure and Applied Mathematics*, 48(12):1305–1342, 1995.
- [13] J. D. Lawson. Generalized Runge-Kutta processes for stable systems with large Lipschitz constants. *SIAM Journal on Numerical Analysis*, 4(3):372–380, 1967.
- [14] A. Loison, S. Kokh, T. Pichard, and M. Massot. A unified two-scale gas–liquid multi-fluid model with capillarity and interface regularization through a mass transfer between scales. *International Journal of Multiphase Flow*, 177:104857, 2024.
- [15] C. D. Meyer, D. S. Balsara, and T. D. Aslam. A stabilized Runge–Kutta–Legendre method for explicit super-time-stepping of parabolic and mixed equations. *Journal of Computational Physics*, 257:594–626, 2014.
- [16] H. Sagan. *Space-Filling Curves*. Springer New York, 1994.
- [17] A. E. Simon, L. François, A. Asad, E. Quémerais, and M. Massot. Development of a high-order space-time coupling HPC library for multiphysics applications. Technical report, NASA, 2025.
- [18] W. Zhang, A. Myers, K. Gott, A. Almgren, and J. Bell. AMReX: Block-structured adaptive mesh refinement for multiphysics applications. *The International Journal of High Performance Computing Applications*, 35(6):508–526, 2021.

A splitting-ImEx strategy for convection-reaction-diffusion problems with space-time adaptation and error control.

Josselin Massot¹ Marc Massot¹ Laurent Séries¹ Christian Tenaud²

¹CMAP, CNRS, École polytechnique, Institut Polytechnique de Paris, 91120 Palaiseau, France

²Université Paris-Saclay, CNRS, CentraleSupélec, EM2C, 91190, Gif-sur-Yvette, France

Abstract

Advanced numerical strategies for the predictive simulation of multiscale reactive flows with realistic and complex models are mandatory for both computational efficiency or energy resource limitations. Whereas a class of operator splitting techniques have shown such efficiency, they reach their limits when the reaction and diffusion operators are coupled at small scales and also involve splitting errors. In this context, there is a need to introduce a new class of methods that maintain the same level of computational efficiency but avoid the drawbacks of splitting methods. The purpose of the present contribution is to introduce such a dual splitting-ImEx approach and to analyze and assess it on simple one-dimensional and two-dimensional test cases with dynamic mesh adaptation and error control.

Keywords: Reaction-diffusion multiscale PDEs, ImEx and operator splitting scheme, Mesh adaptation

1 Introduction

General convection-reaction-diffusion PDEs can model a wide range of realistic physical phenomena in which multiple processes can interact simultaneously at different time and spatial scales. Such multiscale systems often exhibit stiffness and thus pose challenges for accurate numerical simulation at reasonable computational cost, especially when dealing with complex reactive source terms. This is mainly the case for combustion applications [1], where multiscale phenomena are a common feature and multicomponent reacting flows involve complex chemistry and detailed diffusion transport.

Whereas fully explicit methods fail due to the involved stiffness, implicit methods exhibit good numerical properties for solving such problems; however, they can become too expensive because of the cost of solving large nonlinear systems at each time step, especially in multidimensional problems. To circumvent this issue, [2] has designed a strategy based on time-adaptive operator splitting combined with dynamically adaptive mesh and error control. This strategy has optimal computational cost properties, as the splitting time step is chosen to control splitting errors, decoupling it from stability limits. For instance, regarding combustion applications [3], large splitting time steps compared to the smallest time scale can be used without impairing accuracy. Such an approach, however, can lose its efficiency when high precision is required and when

the coupling between the various operators occurs at the smallest timescale. This might be the case in reactive flows with complex chemistry when the required splitting time step obtained through error control is too small. As an alternative, Implicit/Explicit (ImEx) schemes [4] may be complex to implement and design but offer several attractive features in terms of computational cost. They solve the diffusion terms using an extended stability explicit Runge-Kutta method, while the stiff reaction terms are solved implicitly. However, the framework is usually related to the method of line and based on an ODE approach, whereas the most efficient schemes for convection can involve space-time discretization [5]. Besides, since the time step used in the global approach is based on the stability limit of the schemes, the time step can be artificially limited by advection for convection-dominated problems, such as in [3], which reduces the efficiency of the method by solving the reaction term too often.

The goal of this contribution is to propose a dual method that combines the beneficial features of both approaches to alleviate the limitations of operator splitting for reacting flow applications with complex chemistry, while offering the same impact in terms of limited computational cost and adaptability in space and time with error control. Although the objective of this work is to tackle convection-reaction-diffusion PDEs, the splitting part, where we use convection on the one side resolved by One-Step Monotonicity Preserving methods

[5] and reaction diffusion on the other side, has already been presented in [3]. The present contribution is thus primarily restricted to addressing reaction-diffusion coupling problems. We present here the numerical strategy, along with the libraries that we will rely on, in which the reaction-diffusion PDEs are solved using a time-adaptive ImEx method inspired by [4]. The present strategy is designed to use an optimal time step to solve the reaction-diffusion coupled problem while preserving a large splitting time step to preserve the high computational gain obtained in [3]. Additionally, the ImEx approach is coupled with a multiresolution strategy (MRA) [6, 2] for mesh adaptation, enabling the handling of multiple spatial scales.

This ImEx strategy is evaluated through several numerical tests for reaction-diffusion PDEs, enabling us to highlight the advantages and drawbacks of the strategy in comparison to operator splitting techniques. We present results obtained for the three-equation 1-D reaction-diffusion *Belousov-Zhabotinsky* (BZ) equation to demonstrate the relevance of coupling reaction and diffusion operators with an ImEx method compared to a splitting operator approach. The good behavior of the ImEx approach is assessed on the 2-D BZ problem with dynamic mesh adaptation.

The report is organized as follows. In a second section, we present the numerical strategy and the libraries we will rely on. In a third section, several numerical tests for reaction-diffusion PDEs are presented to clarify the advantages and drawbacks of the ImEx strategy compared to operator splitting techniques.

2 Numerical strategy

Starting from the convection-reaction-diffusion PDEs, we use an operator splitting in order to separate convection on the one-side and reaction-diffusion on the other side, with potentially adaptive splitting [7]. Whereas the convection part is resolved through a couple in time and space numerical strategy [5], the reaction-diffusion part is first semi-discretized in space through a method of line and we end up with an ODE of the form :

$$d_t \mathbf{y}(t) = \mathbf{F}_D(\mathbf{y}(t)) + \mathbf{F}_R(\mathbf{y}(t)), \quad (2.1)$$

where $\mathbf{y}(t) \in \mathbb{R}^{m \times N}$ is the solution vector of the m equations discretized at N points depending on time (t). \mathbf{F}_D , and \mathbf{F}_R are, respectively, the diffusion and the reaction operators. In order to solve this type of ODEs, several time integration strategies can be adopted.

The operator splitting approach used in [7, 2] has shown very interesting properties in terms of computational efficiency, it requires time-stepping adaptation to ensure operator coupling at the proper scale and can lead to very small time steps when the coupling of the two operators is taking place at the finest scale such as in complex chemistry. As an alternative, we study the opportunity to use additive Runge-Kutta ImEx schemes, which allows to eliminate the splitting errors, while preserving the "good properties" of the operator splitting techniques : solving the diffusion with explicit

stabilized RK methods, solving the reaction with an implicit method that allows local source resolution in each computational cell (embarrassingly parallel). In the following, we first present the fully operator splitting strategy based on a Strang integration [3] and an ImEx strategy based on the PIROCK method [4].

2.1 Strang operator splitting strategy

An operator splitting strategy [7, 2] has been employed which allows the use of high-order time integrations that are particularly well-suited to solve diffusion (\mathbf{F}_D) and reaction (\mathbf{F}_R). This approach uses a 2nd-order Strang procedure [8]:

$$\mathbf{y}(t + \Delta t) = \mathbf{F}_R^{(\Delta t/2)} \cdot \mathbf{F}_D^{(\Delta t)} \cdot \mathbf{F}_R^{(\Delta t/2)} \cdot \mathbf{y}(t). \quad (2.2)$$

A 4th-order extended stabilized explicit Runge-Kutta method (ROCK4 [9]) is employed for solving the diffusion operation. A fully implicit Runge-Kutta method (Radau IIA) that is particularly suited for stiff and differential-algebraic problems is used for the numerical integration of the reaction operator.

Although this strategy has optimal properties in terms of computational cost, it faces drawbacks when high precision is required and when the coupling occurs at the smallest timescale, because the splitting time step must be drastically decreased to recover the physical behavior [3]. This particularly occurs in combustion applications at high Damköhler numbers when reaction and diffusion are physically coupled, which can be seen, for instance, in the presence of thermo-diffusive instabilities for flame acceleration or in auto-ignition of diffusive flame with complex chemistry.

2.2 An ImEx fully coupled strategy

We use an additive Runge-Kutta ImEx method originally developed for three operators (PIROCK introduced in [4]), which we adapt for the reaction-diffusion case. It relies on an extended stability domain Runge-Kutta method (ROCK2 [9]) for a diffusion operator \mathbf{F}_D , a Diagonally Implicit Runge-Kutta (DIRK) method for the reaction operator \mathbf{F}_R , which is A - and L -stable.

The resulting method is based on a 5-step procedure:

1. Diffusion stabilization procedure:

$$\begin{aligned} \mathbf{y}^{(0)} &= \mathbf{y}^n \\ \mathbf{y}^{(1)} &= \mathbf{y}^n + \alpha \mu_1 \Delta t \mathbf{F}_D(\mathbf{y}^n) \\ \mathbf{y}^{(j)} &= \alpha \mu_j \Delta t \mathbf{F}_D(\mathbf{y}^{(j-2)}) - \nu_j \mathbf{y}^{(j-1)} - \kappa_j \mathbf{y}^{(j-2)}, \\ & \quad j = 2, \dots, s - 2 + \ell \end{aligned}$$

where μ_j, ν_j, κ_j are the same coefficients as in the standard ROCK2 method with s stages; α is a degree of freedom to relax the stability domain.

2. Finishing procedure for diffusion:

$$\begin{aligned} \mathbf{y}_*^{(s-1)} &= \mathbf{y}^{(s-2)} + \sigma_\alpha \Delta t \mathbf{F}_D(\mathbf{y}^{(s-2)}) \\ \mathbf{y}_*^{(s)} &= \mathbf{y}_*^{(s-1)} + \sigma_\alpha \Delta t \mathbf{F}_D(\mathbf{y}_*^{(s-1)}) \end{aligned}$$

where $\sigma_\alpha = \frac{1-\alpha}{2} + \alpha\sigma$, and σ is the same coefficient as in the standard ROCK2 method with s stages.

3. Stating value for reaction: $\tilde{\mathbf{y}} = \mathbf{y}^{(s-2+\ell)}$.
4. Finishing procedure for reaction and coupling:

$$\begin{aligned} \mathbf{y}^{(s+1)} &= \tilde{\mathbf{y}} + \gamma \Delta t \mathbf{F}_R(\mathbf{y}^{(s+1)}) \\ \mathbf{y}^{(s+2)} &= \tilde{\mathbf{y}} + \beta \Delta t \mathbf{F}_D(\mathbf{y}^{(s+1)}) \\ &\quad + (1 - 2\gamma) \Delta t \mathbf{F}_R(\mathbf{y}^{(s+1)}) \\ &\quad + \gamma \Delta t \mathbf{F}_R(\mathbf{y}^{(s+2)}) \\ \mathbf{y}^{(s+3)} &= \tilde{\mathbf{y}} + (1 - \gamma) \Delta t \mathbf{F}_R(\mathbf{y}^{(s+1)}) \end{aligned}$$

where $\beta = 1 - 2\alpha P'_{s-2+\ell}(0)$ with $P'_{s-2+\ell}$ is the derivative of the stability polynomial of first $s-2+\ell$ stages of ROCK2 method with s stages, this value coming from order relations; $\gamma = 1 - \frac{\sqrt{2}}{2}$.

5. Integration step computation $\mathbf{y}^{(n)} \mapsto \mathbf{y}^{(n+1)}$:

$$\begin{aligned} \mathbf{y}^{n+1} &= \mathbf{y}_*^{(s)} - \sigma_\alpha \left(1 - \frac{\tau_\alpha}{\sigma_\alpha^2}\right) \Delta t \\ &\quad \left(\mathbf{F}_D(\mathbf{y}_*^{(s-1)}) - \mathbf{F}_D(\mathbf{y}^{(s-2)}) \right) \\ &\quad + \frac{1}{2} \Delta t \mathbf{F}_R(\mathbf{y}^{(s+1)}) + \frac{1}{2} \Delta t \mathbf{F}_R(\mathbf{y}^{(s+2)}) \\ &\quad + \frac{J_R^{-\ell}}{2 - 4\gamma} \Delta t \left(\mathbf{F}_D(\mathbf{y}^{(s+3)}) - \mathbf{F}_D(\mathbf{y}^{(s+1)}) \right) \end{aligned}$$

where $\tau_\alpha = \frac{(\alpha-1)^2}{2} + 2\alpha(1-\alpha)\sigma + \alpha^2\tau$ with τ the same coefficient as in the standard ROCK2 method with s stages.

The stating value for reaction resembles an intermediate step of a Lie or Strang operators splitting method. With this trick, we do not need to implicitly involve the reaction at each step of the ROCK2 method, everything happens as if we make a splitting operator method with the diffusion and reaction, but with a stronger coupling which does not involve any splitting error. The matrix $J_R = I - \gamma \Delta t \frac{\partial \mathbf{F}_R}{\partial \mathbf{u}}(\mathbf{y}^{(s)})$ is used for further stabilizing the coupling of the diffusion step with the reaction operator. This trick has been noticed by Shampine and presented in [10].

We explore two choices for the parameters ℓ and α :

- If $\mathbf{F}_R \rightarrow 0$, so we want to take $\ell = 2$ and $\alpha = 1$ to get the classical ROCK2 method. This choice is relevant if the problem is dominated by diffusion.
- If $\alpha = \frac{1}{2P'_{s-2+\ell}(0)}$ so $\beta = 0$, with $\ell = 1$ we can reduce the computing cost.

For this method, we can build an embedded method with error estimate for a variable step size control:

$$\begin{aligned} err_D &= \Delta t \sigma_\alpha \left(1 - \frac{\tau_\alpha}{\sigma_\alpha^2}\right) \left(\mathbf{F}_D(\mathbf{y}_*^{(s-1)}) - \mathbf{F}_D(\mathbf{y}^{(s-2)}) \right), \\ err_R &= \frac{\Delta t}{6} J_R^{-1} \left(\mathbf{F}_R(\mathbf{y}^{(s+1)}) - \mathbf{F}_R(\mathbf{y}^{(s+2)}) \right). \end{aligned}$$

The step size adjustments for adaptive integration are based on the strategy $\Delta t_{\text{new}} = \left(\frac{\text{tol}}{\text{err}}\right)^{1/p} \Delta t$ with an optional safety factor, where p is the order of the method. The global error takes the maximum value as $err = \max(\|err_D\|, \|err_R\|)$. The new time step is the computed with $\Delta t_{\text{new}} = 0.8 \Delta t \left(\frac{\text{tol}}{\text{err}}\right)^{1/2}$.

3 Numerical results

We focus on the BZ system of reaction-diffusion PDEs with three species, a classic example of catalyzed oxidation of an organic species by an acid bromate ion. The system reads:

$$\begin{cases} \partial_t a - D_a \Delta a = \frac{1}{\mu} (-qa - ab + fc) \\ \partial_t b - D_b \Delta b = \frac{1}{\varepsilon} (qa - ab + b(1-b)) \\ \partial_t c - D_c \Delta c = b - c \end{cases}$$

with the following parameters: $\varepsilon = 10^{-2}$, $\mu = 10^{-5}$, $f = 3$, $q = 2 \cdot 10^{-4}$, and diffusion coefficients given by $(D_a, D_b, D_c) = (2.5 \cdot 10^{-3}, 2.5 \cdot 10^{-3}, 1.5 \cdot 10^{-3})$, with homogeneous Neumann boundary conditions.

3.1 The 1-D BZ test case

The computational domain spreads over $x \in [0, 4]$. The initial condition is given by:

$$b(x) = \begin{cases} 0.8 & \text{if } \theta < x_0 \\ q \frac{f+1}{f-1} & \text{else} \end{cases}, \quad c(x) = q \frac{f+1}{f-1} + \frac{\theta}{8\pi f},$$

$a(x)$ is initialized at equilibrium $fc(x)/(q+b(x))$, x_0 is the initial location of the sharp gradient, and

$$\theta = \begin{cases} \arctan(x - x_0) + 2\pi & \text{if } x < x_0, \\ \arctan(x - x_0) & \text{elsewhere.} \end{cases}$$

The solution is a multi-scale traveling wave computed using the ImEx strategy (§ 2.2), with a constant number of substeps in the ROCK2 integration. The L_∞ errors, computed at the final time $t = 0.5$, are plotted versus the integration time step ranging from 5×10^{-5} to 10^{-3} in Figure 1, and compared with the Strang splitting. Let us note that the Strang splitting procedure starts by solving the reaction operator, since it involves the fastest time scales. Both strategies recover the expected




Figure 1: L_∞ errors versus integration time step obtained at $t = 0.5$, for two spatial resolutions $Nx = 201$ (left), and $Nx = 401$ (right). The 2nd-order accuracy is shown with the black dashed line.

2nd-order accuracy. Although the error magnitudes are very similar at the lowest spatial resolution ($Nx = 201$), the ImEx coupling strategy yields lower errors in the asymptotic range of the time step for $Nx = 401$. At the largest integration time steps, the ImEx method exhibits larger errors, while increasing the number of

RadauIIA relative tolerance	Splitting/ImEx elapsed time ratio
10^{-8}	18.0
10^{-6}	10.7
10^{-4}	6.5
10^{-2}	3.6
1.	3.6

Table 1: Elapsed time ratio at a specific time step $\Delta t = 2.5 \times 10^{-4}$, between the splitting strategy for several relative tolerances and the ImEx strategy.

substeps in the ROCK2 integration (compared to $Nx = 201$). This might be attributed to a loss of stability at such large time steps.

The most important disagreement concerns the elapsed time to recover all these solutions. They are plotted versus the integration time step in Figure 2. The ImEx approach is clearly more efficient than the Strang splitting strategy, with elapsed times five to six times lower than the splitting method. The high cost




Figure 2: Elapsed time used for computing solutions up to $t = 0.5$ versus the integration time step, for two spatial resolutions ($Nx = 201$, and $Nx = 401$).

of the splitting strategy can be attributed to the severe tolerance we have used in the RadauIIA solver, where the splitting error should be the main source of error to enable adaptive splitting. For the sake of identifying the impact of such a choice in terms of computational time, we have prescribed the relative tolerance and deduced the absolute tolerance with a ratio of 10^{-3} . We then vary the relative tolerance of the RadauIIA integrator to compare with the ImEx elapsed time at a specific time step ($\Delta t = 2.5 \times 10^{-4}$). The ImEx strategy is very efficient, providing accurate solutions at the expected 2nd-order at a cost at least 4 times lower than the splitting strategy for loose relative tolerance. The elapsed time ratio can increase significantly when the relative tolerance of the RadauIIA is reduced, thus indicating that for a given error level, the ImEx approach is computationally more competitive than the splitting one, without involving splitting errors.

3.2 The 2-D BZ test case

The 2-D BZ test case has also been experienced to demonstrate the capability of the ImEx coupling strategy. The initial condition is given by:

$$b(x, y) = \begin{cases} 0.8 & \text{if } \theta < \frac{1}{2} \\ q \frac{f+1}{f-1} & \text{else} \end{cases}, \quad c(x, y) = q \frac{f+1}{f-1} + \frac{\theta}{8\pi f}$$

where $\theta = \arctan(x/y)$, and $a(x, y)$ is initialized to chemical equilibrium given by: $a(x, y) = \frac{fc(x, y)}{q+b(x, y)}$. The solution is a spiral traveling wave.

Additionally, a new implementation of the ImEx integration has been developed in a C++ open-source software *ponio* [11]. Thanks to this new implementation, the time integration can be coupled with spatial mesh adaptation via a multiresolution strategy using the open-source software *samurai* [12].

We present preliminary results from the new implementation of the ImEx integration, coupled with mesh adaptation. These results obtained with a specific integration time step $\Delta t = 10^{-3}$ are plotted at a final time $t = 1$ in Figure 3, for the Strang splitting strategy, on the left, and the ImEx strategy (§ 2.2), on the right. Both strategies give the same result with the same mesh




Figure 3: Results at time $t = 1$ of 2-D BZ PDE's with multiresolution, and Strang splitting (left) and ImEx coupling strategy (right), with time step $\Delta t = 10^{-3}$.

adaptation, confirming the correct implementation of the ImEx algorithm. However, as in the 1D test case, the ImEx approach is more efficient, with an elapsed time lower than that of the splitting.

4 Conclusion

We have presented a numerical strategy for solving the reaction-diffusion PDEs using an ImEx method inspired by [4]. The three-species Belousov-Zhabotinsky system of PDEs has been used as a test case to highlight the advantages and drawbacks of the strategy compared to traditional operator-splitting techniques. In the 1D case, the two methods have completely different sources of error but yield a comparable level of error. The ImEx strategy proves to be computationally more competitive. However, separating the splitting errors from the integration error of each operator using high-order adaptive methods results in a strong dependence of the time to solution for the splitting method, depending on the tolerance we use for the internal operator integration. Eventually, since we can avoid splitting errors as well as the usual weakness of the splitting method, the ImEx strategy seems to be a very promising technique. We still have to demonstrate that it remains the case when complex chemistry is involved and the ImEx strategy is in the process of being coupled through adaptive splitting to a convection operator for combustion applications.

Acknowledgments

We thank the NASA hosts for their hospitality. This work has been supported by the HPC@Maths Initiative of the Fondation de l'École polytechnique.

References

- [1] T. Echehki, “Multiscale methods in turbulent combustion: strategies and computational challenges”, *Computational Science & Discovery* **2** (2009), no. 1, p. 013001.
- [2] M. Duarte, M. Massot, S. Descombes, C. Tenaud, T. Dumont, V. Louvet, F. Laurent, “New Resolution Strategy for Multiscale Reaction Waves using Time Operator Splitting, Space Adaptive Multiresolution, and Dedicated High Order Implicit/Explicit Time Integrators”, *SIAM Journal on Scientific Computing* **34** (2012), no. 1, p. A76-A104.
- [3] M. Duarte, S. Descombes, C. Tenaud, S. Candel, M. Massot, “Time–space adaptive numerical methods for the simulation of combustion fronts”, *Combustion and Flame* **160** (2013), no. 6, p. 1083-1101.
- [4] A. Abdulle, G. Vilmart, “PIROCK: A swiss-knife partitioned implicit–explicit orthogonal Runge–Kutta Chebyshev integrator for stiff diffusion–advection–reaction problems with or without noise”, *Journal of Computational Physics* **242** (2013), p. 869-888.
- [5] V. Daru, C. Tenaud, “High order one-step monotonicity-preserving schemes for unsteady compressible flow calculations”, *Journal of Computational Physics* **193** (2004), no. 2, p. 563-594.
- [6] A. Cohen, S. M. Kaber, S. Müller, M. Postel, “Fully adaptive multiresolution finite volume schemes for conservation laws”, *Math. Comp.* **72** (2003), p. 183-225.
- [7] S. Descombes, M. Duarte, T. Dumont, V. Louvet, M. Massot, “Adaptive time splitting method for multi-scale evolutionary partial differential equations”, *Confluentes Mathematici* **3** (2011), no. 3, p. 413-443.
- [8] G. Strang, “On the Construction and Comparison of Difference Schemes”, *SIAM Journal on Numerical Analysis* **5** (1968), no. 3, p. 506-517.
- [9] A. Abdulle, A. Medovikov, “Second order Chebyshev methods based on orthogonal polynomials”, *Numerische Mathematik* **90** (2001), p. 1-18.
- [10] E. Hairer, G. Wanner, *Solving Ordinary Differential Equations II*, Springer Berlin Heidelberg, 1996.
- [11] J. Massot, M. Massot, L. Series, “ponio”, 2024, <https://hal.science/hal-04710549>.
- [12] T. Bellotti, L. Gouarin, M. Massot, P. Matalon, “samurai”, Apr 2023, <https://hal.science/hal-04545389>.
- [13] G.-S. Jiang, C.-W. Shu, “Efficient Implementation of Weighted ENO Schemes”, *Journal of Computational Physics* **126** (1996), no. 1, p. 202-228.
- [14] D. S. Balsara, C.-W. Shu, “Monotonicity Preserving Weighted Essentially Non-oscillatory Schemes with Increasingly High Order of Accuracy”, *Journal of Computational Physics* **160** (2000), p. 405-452.
- [15] L. Pareschi, G. Russo, “Implicit–Explicit Runge–Kutta Schemes and Applications to Hyperbolic Systems with Relaxation”, *Journal of Scientific Computing* **25** (2005), no. 1/2, p. 129–155.
- [16] D. Thévenin, S. Candel, “Ignition dynamics of a diffusion flame rolled up in a vortex”, *Phys. Fluids* **7** (1995), no. 2, p. 434-445.

Plasma/Rarefied methods

Simulation of hybrid dense-rarefied flows relying on relaxation schemes and a hierarchy of moment methods

P. Bernigaud¹ A. Clout¹ D. Biasone^{1,2} S. Dubois² A. Hoffmann²
M. Massot²

¹DMPE, ONERA, Université Paris Saclay, 91120 Palaiseau, France

²CMAP, CNRS, École polytechnique, Institut Polytechnique de Paris, 91120 Palaiseau, France
December 4, 2025

Abstract

Rarefied gas flows are encountered in various aerospace applications, including hypersonic flight, atmospheric re-entry, and rocket propulsion. These flows often span a wide range of rarefaction levels, posing a challenge to the selection of suitable numerical models. While the Navier–Stokes equations efficiently describe continuum regimes, their assumptions fail as the Knudsen number increases. Conversely, the Direct Simulation Monte Carlo (DSMC) method accurately resolves rarefied regimes but becomes computationally prohibitive near the continuum limit. In this work, we propose a fluid model that efficiently solves flows with varying levels of rarefaction, from the continuum up to transitional regime. The fluid model is derived applying the method of moments and formulating a relaxation scheme that enables the coupling of moment systems of different orders without the need for an explicit interface. The proposed approach is evaluated on the benchmark case of a two-dimensional expanding jet, representative of high-altitude rocket plumes, compared with DSMC results and assessed.

Keywords: rarefied flow, moment method, relaxation scheme, numerical simulation, DSMC

1 Introduction

Rarefied gas flows occur in many aerospace applications, such as hypersonic flight, atmospheric re-entry, and rocket propulsion. Their study is particularly challenging when rarefaction varies strongly, as in high-altitude rocket plumes, where the dense nozzle flow rapidly expands into a highly rarefied environment.

The degree of rarefaction is characterized by the Knudsen number (K_n), defined as the ratio of the molecular mean free path to a characteristic length. For $K_n \ll 1$, the flow is continuous and well described by the Euler or Navier–Stokes equations. When $K_n \gg 1$, continuum assumptions fail, and the gas must be modelled using the Boltzmann equation. Between these limits lies the transitional regime, where $K_n \approx 1$.

A widely used approach to solve the Boltzmann equation is the Direct Simulation Monte Carlo (DSMC) method, popularized by Bird [1]. This stochastic technique estimates the molecular distribution function by tracking representative particles undergoing binary collisions. DSMC has proven highly successful but its computational cost increases significantly near the

continuum regime due to the large number of simulated particles. To overcome this limitation, hybrid Navier–Stokes/DSMC frameworks have been developed, solving the continuum regions with a Navier–Stokes solver and the rarefied zones with DSMC. Although such hybrid strategies are now mature and applied to industrial configurations [2, 3], their main challenge lies in the accurate placement of the coupling interface, which must be located where the Navier–Stokes equations remain marginally valid, yet within a sufficiently rarefied region to keep DSMC computations tractable.

Motivated by these limitations, we propose a moment-based fluid formulation capable of bridging the gap between continuum and rarefied regimes. Although moment methods are well established for rarefied flows [4, 5], we improve efficiency by formulating a relaxation scheme. This approach enables the coupling of moment systems of different orders without the need for an explicit interface, allowing the model to adapt its level of description to the local degree of rarefaction. The method is demonstrated on the benchmark case of a two-dimensional expanding jet, representative of high-altitude rocket plumes. While this study focuses

on a single-species gas, the approach has the potential to be extended to multi-species systems.

2 Numerical method

2.1 Moment method

In order to capture flow features in rarefied regions, it is proposed that a reduced fluid model be derived by employing the method of moments. Rather than trying to estimate the density distribution function of the molecules directly, as would a DSMC solver, this deterministic method focuses on the computation of moments of the distribution. Indeed, engineers are most often interested in statistically averaged quantities such as density, flow velocity or temperature.

We derive a system of equations on moments from the Boltzmann equation. In the interest of conciseness, calculations are shown in 1D space dimension and 1D in molecule velocity phase space (1D1V). In this setting, the Boltzmann equation reads

$$\frac{\partial f}{\partial t} + \xi \frac{\partial f}{\partial x} = \mathcal{C}(f, f), \quad (2.1)$$

where x and t are the space and time dimensions, $f(\xi, x, t)$ is the density distribution function of the gas, ξ is the 1D peculiar velocity of a given molecule, and $\mathcal{C}(f, f)$ a collision operator. Multiplying Eq. (2.1) by the family $(\xi^k)_{0 \leq k \leq n-1}$ and integrating over the velocity phase space, one obtains

$$\partial_t \langle \xi^k f \rangle + \partial_x \langle \xi^{k+1} f \rangle = \langle \xi^k \mathcal{C}(f, f) \rangle \quad (2.2)$$

where $\langle \cdot \rangle$ denotes the integration over \mathbb{R} . Introducing $\mu_k(x, t) = \langle \xi^k f \rangle$ the k -th moment of f and S_k the collision term, a system of n conservation equations for the moments is expressed as

$$\partial_t (\mu_k) + \partial_x (\mu_{k+1}) = S_k. \quad (2.3)$$

Noting that the equation for μ_k depends on μ_{k+1} it is necessary to include an additional equation to close the system. This equation takes the form of a relationship providing an estimate \hat{f} of the unknown distribution f from the n moments at our disposal,

$$\hat{f}(\xi, x, t) = \mathcal{F}(\mu_0, \mu_1, \dots, \mu_{n-1}) \quad (2.4)$$

and under the constraint that \hat{f} realizes the moments $(\mu_k)_{0 \leq k \leq n-1}$, *i.e.*,

$$\forall k \in \llbracket 0, n-1 \rrbracket, \quad \mu_k = \langle \xi^k \hat{f} \rangle. \quad (2.5)$$

The system is then closed approximating the unknown moment μ_n by

$$\mu_n \approx \langle \xi^n \hat{f} \rangle. \quad (2.6)$$

The main difficulty encountered when designing a moment method is the choice of the reconstruction function \mathcal{F} . Historically, the pioneering work of Grad [6] proposed to express \hat{f} as a sum of Hermite polynomials, but this approach suffers from severe drawbacks.

Namely, it has been shown this approach results in an hyperbolic system only near equilibrium. Since then, the development of closures for the moment system has been an active field of research, developing new closures based on the use of quadratures [7, 8, 9], entropy minimization problems [10, 11, 5], or modified Grad-like expansions [12, 13]. At the present time, no consensus has been reached in the community designating a specific closure as the best course of action. Comparing methods is an important research subject, but it is beyond the scope of the present report.

In this study, we adopt the classical Grad closure [6] for its simplicity and low computational cost, as the focus here lies on the development of the relaxation scheme rather than on the closure problem itself. The assumed shape of the distribution follows the ansatz

$$\hat{f} = \mathcal{M}(\rho, u, \theta) \sum_{\alpha=0}^{n-1} f_\alpha \theta^{-\alpha/2} \text{He}_\alpha \left(\frac{\xi - u}{\sqrt{\theta}} \right), \quad (2.7)$$

where ρ is the local density, u is the flow velocity, and θ is the specific thermal energy, derived from the first three moments. He_α denotes the Hermite polynomial of order α , and $(f_\alpha)_{0 \leq \alpha \leq n-1}$, local coefficients to be determined to ensure the compatibility condition Eq. (2.5). Finally, $\mathcal{M}(\rho, u, \theta)$ is the equilibrium Maxwellian distribution of parameters ρ , u and θ , expressed as

$$\mathcal{M}(\rho, u, \theta) = \frac{\rho}{\sqrt{2\pi\theta}} \exp \left(-\frac{1}{2} \frac{|\xi - u|^2}{\theta} \right). \quad (2.8)$$

The last missing piece in our model is the expression of the collision operator $\mathcal{C}(f, f)$. We consider the Bhatnagar-Gross-Krook (BGK) operator [14], modelling the collisional process as a relaxation toward equilibrium,

$$\mathcal{C}_{\text{BGK}}(f, f) = \frac{1}{\tau} (\mathcal{M}(\rho, u, \theta) - f), \quad (2.9)$$

Note that this simplifying assumption enforces a unit Prandtl number. Future work will consider more realistic collision models, such as ES-BGK [15, 16] or Boltzmann collisions.

2.2 Relaxation scheme

To efficiently capture flows with spatially varying rarefaction levels, the governing equations should adapt their level of detail to the local regime, retaining more moments in rarefied regions and fewer in dense ones. This is achieved through a relaxation scheme that couples moment systems of different orders without requiring an explicit interface.

Relaxation schemes have been widely used to simplify the numerical solution of complex systems [17, 18, 19]. Boileau *et al.* [20] applied this concept in spray modelling to couple the Euler and pressureless gas dynamics equations. Extending this idea, we use a similar strategy to couple two moment systems of different orders sharing a common Grad-type closure.

Suppose we have a *full system* of moment equations including moments up to order N , which we aim at

coupling with a *reduced system* including moments up to order $n < N$. We denote by \mathcal{V} and \mathcal{U} the state vectors of the reduced and full system, expressed as

$$\mathcal{V} = (\mu_0, \dots, \mu_{n-1}) \quad (2.10)$$

$$\mathcal{U} = (\mathcal{V}, \mu_n, \dots, \mu_{N-1}), \quad (2.11)$$

and we define the flux vectors $G(\mathcal{V})$ and $F(\mathcal{U})$ for the reduced and full systems,

$$G(\mathcal{V}) = (\mu_1, \dots, \langle \xi^n \hat{f}_n \rangle) \quad (2.12)$$

$$F(\mathcal{U}) = (\mu_1, \dots, \langle \xi^N \hat{f}_N \rangle), \quad (2.13)$$

where \hat{f}_n (resp. \hat{f}_N) is the estimated distributions provided by the closure based on the first n (resp. N) moments. The reduced and full systems are then expressed as

$$\partial_t \mathcal{V} + \partial_x G(\mathcal{V}) = S(\mathcal{V}) \quad (2.14)$$

$$\partial_t \mathcal{U} + \partial_x F(\mathcal{U}) = S(\mathcal{U}), \quad (2.15)$$

where $S(\mathcal{U})$ and $S(\mathcal{V})$ denote the collision operator.

To produce a relaxation scheme, we introduce the relaxation operator

$$\mathcal{R}(\mathcal{U}) = (\mathcal{R}_0, \dots, \mathcal{R}_N), \quad (2.16)$$

where

$$\mathcal{R}_k = \begin{cases} 0, & k \in \llbracket 0, n-1 \rrbracket \\ \langle \xi^k \hat{f}_n \rangle - \mu_k, & k \in \llbracket n, N-1 \rrbracket. \end{cases} \quad (2.17)$$

We then propose to solve the following relaxed system,

$$\partial_t \mathcal{U} + \partial_x F(\mathcal{U}) = S(\mathcal{U}) + \lambda \mathcal{R}. \quad (2.18)$$

The parameter $1/\lambda$ represents a characteristic relaxation time controlling the transition between the two levels of description. In the limit $\lambda \rightarrow 0$, the relaxation term vanishes and the full N -moments system (2.15) is recovered. Conversely, for $\lambda \rightarrow \infty$, the higher-order moments with index $k > n$ are projected on values provided by the closure implied by the first n moments, yielding the reduced n -moment system (2.14).

By letting $\lambda = 0$ in rarefied regions to recover the full system and $\lambda \rightarrow \infty$ in dense regions to recover the reduced one, substantial computational savings are achieved by advancing the high-order moments ($k > n$) only in cells where the full system is required.

This relaxation scheme can be generalized to higher velocity dimensions by defining full and reduced systems based on the moment index sets \mathcal{I}_N and \mathcal{I}_n with $\mathcal{I}_n \subset \mathcal{I}_N$. The relaxation operator is then

$$\mathcal{R}_k = \begin{cases} 0, & \alpha \in \mathcal{I}_n \\ \langle \xi^\alpha \hat{f}_n \rangle - \mu_\alpha, & \alpha \in \mathcal{I}_N \setminus \mathcal{I}_n. \end{cases} \quad (2.19)$$

with ξ the multi-dimensional molecular velocity.

It is to be noted that this relaxation scheme is stable only if the *sub-characteristic condition* is satisfied [21],

meaning that the characteristic speeds of the full system are interlaced with those of the reduced one. Chen *et al.* [22] introduced a strong condition to enforce this stability, namely the dissipation of an entropy in the relaxed system. For a system of moment equations, compliance with this condition depends solely on the chosen closure. In this study, Grad's closure is adopted as a proof of concept. Since it dissipates entropy only near equilibrium, the method should be applied with caution in strongly non-equilibrium regimes. Future work will use this relaxation scheme in conjunction with any hyperbolic closure for improved stability.

3 Results

The proposed relaxation scheme is implemented in C++ leveraging the adaptative mesh refinement library *Samurai*¹. It is tested on a 2D3V toy problem representative of the high-altitude expansion of a single species rocket plume.

The computational domain is a square of size $L = 0.2$ m. Two parallel argon streams enter from the left boundary with identical velocity and temperature but different densities, described by Maxwellian distributions of parameters

$$\rho(x=0, y) = \begin{cases} 3.5 \times 10^{-3} \text{ kg/m}^3, & 0 \leq y < \delta \\ 3.5 \times 10^{-4} \text{ kg/m}^3 & \delta \leq y \leq L \end{cases}$$

$$\mathbf{u}(x=0, y) = (300 \text{ m s}^{-1}, 0 \text{ m s}^{-1})$$

$$\theta(x=0, y) = 10^4 \text{ m}^2/\text{s}^2,$$

where $\delta = 5$ mm denotes the interface between the two streams. The denser jet mimics a rocket plume, while the upper flow represents the surrounding rarefied atmosphere. For argon, these conditions correspond to $T = 48$ K, with pressures of $P = 35$ and 3.5 Pa in the dense and rarefied regions, respectively. Neumann conditions are imposed at $x = L$ and $y = L$, while a symmetry condition is enforced at $y = 0$.

Using δ as the characteristic length, the local Knudsen number ranges from $\text{Kn} = 10^{-3}$ to 10^{-2} . This limited variation keeps the flow near equilibrium, ensuring the Grad closure remains hyperbolic and the relaxation scheme stable.

The relaxation time τ for the BGK collision operator, is computed based on the local state as $\tau = \mu/P$. The dynamic viscosity μ is computed employing the variable hard sphere model [1]

$$\mu = \mu_{\text{ref}} \left(\frac{T}{T_{\text{ref}}} \right)^{s^\omega} \quad (3.1)$$

$$\mu_{\text{ref}} = \frac{15 \sqrt{\pi m_{\text{Ar}} k_b T_{\text{ref}}}}{2(5 - 2\omega)(7 - 2\omega) \pi d_{\text{ref}}^2},$$


with reference temperature $T_{\text{ref}} = 273$ K, reference diameter $d_{\text{ref}} = 4.17 \times 10^{-10}$ m, temperature index $\omega = 0.81$ and argon molecular mass $m_{\text{Ar}} = 6.63 \times 10^{-26}$ kg.

Fluxes are computed using a first-order Rusanov scheme. The full system retains moments up to third


¹<https://hpc-math-samurai.readthedocs.io/en/latest/>

order (including the heat flux as an explicit variable), whereas the reduced system is limited to second-order moments, corresponding to the anisotropic Euler equations. The computational grid consists of one million cells, and results are compared against a reference DSMC simulation using eight million cells.

The density field computed with both approaches is shown in Fig. 1. Both methods are in fairly good agreement. During expansion, the initially dense jet undergoes rapid rarefaction, followed by a recompression along the centreline further downstream. An oblique region of increased density appears at the interface where the two flow streams interact.



(a) Relaxed method of moments.




(b) Reference DSMC.


Figure 1: Density field.

A more challenging quantity to capture is the heat flux. Its x -wise component is shown in Fig. 2. Both approaches reproduce a similar field topology, with a pronounced oblique band of negative heat flux at the collision zone, indicative of strong heating, and a region of moderate positive flux below, corresponding to gas cooling during expansion. Differences in magnitude arise from the coarser mesh and first-order flux used in the moment method.

A simulation performed without activating the re-



(a) Relaxed method of moments.



(b) Reference DSMC.

Figure 2: Heat flux along e_x .

laxation scheme produces identical results, while its inclusion reduced the computational cost by approximately 40%. Although this test case is of moderate difficulty, comparable performance gains are anticipated in more demanding configurations with stronger non-equilibrium effects. Achieving this, however, requires replacing Grad's closure with a hyperbolic one to maintain the stability of the relaxation scheme.

4 Conclusions

In this work, we introduced a relaxation-based framework for coupling moment systems of different orders without requiring an explicit interface. The method was evaluated on a moderately rarefied test case, demonstrating substantial computational savings. Future work will focus on incorporating higher-order fluxes and extending the formulation to multi-species flows. By employing hyperbolic closures, the stability of the scheme will be assessed in more demanding configurations featuring stronger non-equilibrium effects and larger variations in K_n .

Acknowledgments

The authors gratefully acknowledge the NASA ARC–TSM Branch and Dr. Bruno Dias for their hospitality. Financial support from ONERA’s International Affairs Directorate, HPC@Maths and CIEDS is also acknowledged

References

- [1] G. A. Bird, *Molecular gas dynamics and the direct simulation of gas flows*. Oxford university press, 1994.
- [2] V. Charton, A. Awad, and J. Labaune, “Optimisation of a hybrid NS-DSMC methodology for continuous-rarefied jet flows,” *Acta Astronautica*, vol. 195, pp. 295–308, 2022.
- [3] A. Clout, A. Langenais, D. Yann, M. Luc, and J. Labaune, “Hybrid NS-DSMC simulation of a full scale solid rocket motor reactive exhaust at high altitude,” in *Joint 10th EUCASS 9th CEAS Conference*, 2023.
- [4] J. Koellermeier and M. Torrilhon, “Two-dimensional simulation of rarefied gas flows using quadrature-based moment equations,” *Multiscale Modeling & Simulation*, vol. 16, no. 2, pp. 1059–1084, 2018.
- [5] S. Boccelli, P. Parodi, T. E. Magin, and J. G. McDonald, “Modeling high-Mach-number rarefied crossflows past a flat plate using the maximum-entropy moment method,” *Physics of Fluids*, vol. 35, no. 8, 2023.
- [6] H. Grad, “On the kinetic theory of rarefied gases,” *Communications on pure and applied mathematics*, vol. 2, no. 4, pp. 331–407, 1949.
- [7] R. McGraw, “Description of aerosol dynamics by the quadrature method of moments,” *Aerosol science and Technology*, vol. 27, no. 2, pp. 255–265, 1997.
- [8] C. Chalons, R. O. Fox, F. Laurent, M. Massot, and A. Vié, “Multivariate Gaussian Extended Quadrature Method of Moments for Turbulent Disperse Multiphase Flow,” *Multiscale Modeling & Simulation*, vol. 15, no. 4, pp. 1553–1583, 2017.
- [9] R. O. Fox and F. Laurent, “Hyperbolic quadrature method of moments for the one-dimensional kinetic equation,” *SIAM Journal on Applied Mathematics*, vol. 82, no. 2, pp. 750–771, 2022.
- [10] C. D. Levermore, “Moment closure hierarchies for kinetic theories,” *Journal of Statistical Physics*, vol. 83, pp. 1021–1065, 1996.
- [11] M. Abdelmalik and E. Van Brummelen, “Moment closure approximations of the boltzmann equation based on φ -divergences,” *Journal of Statistical Physics*, vol. 164, no. 1, pp. 77–104, 2016.
- [12] H. Struchtrup and M. Torrilhon, “Regularization of Grad’s 13 moment equations: Derivation and linear analysis,” *Physics of Fluids*, vol. 15, no. 9, pp. 2668–2680, 2003.
- [13] Z. Cai, “Moment method as a numerical solver: challenge from shock structure problems,” *Journal of Computational Physics*, vol. 444, p. 110593, 2021.
- [14] P. L. Bhatnagar, E. P. Gross, and M. Krook, “A model for collision processes in gases. I. Small amplitude processes in charged and neutral one-component systems,” *Physical review*, vol. 94, no. 3, p. 511, 1954.
- [15] L. H. Holway, “New statistical models for kinetic theory: methods of construction,” *The physics of fluids*, vol. 9, no. 9, pp. 1658–1673, 1966.
- [16] P. Andries, P. Le Tallec, J.-P. Perlat, and B. Perthame, “The Gaussian-BGK model of Boltzmann equation with small Prandtl number,” *European Journal of Mechanics-B/Fluids*, vol. 19, no. 6, pp. 813–830, 2000.
- [17] F. Coquel and B. Perthame, “Relaxation of energy and approximate Riemann solvers for general pressure laws in fluid dynamics,” *SIAM Journal on Numerical Analysis*, vol. 35, no. 6, pp. 2223–2249, 1998.
- [18] C. Berthon, M. Breuß, and M.-O. Titeux, “A relaxation scheme for the approximation of the pressureless Euler equations,” *Numerical Methods for Partial Differential Equations: An International Journal*, vol. 22, no. 2, pp. 484–505, 2006.
- [19] F. Bouchut, C. Chalons, and S. Guisset, “An entropy satisfying two-speed relaxation system for the barotropic euler equations: application to the numerical approximation of low mach number flows,” *Numerische Mathematik*, vol. 145, no. 1, pp. 35–76, 2020.
- [20] M. Boileau, C. Chalons, and M. Massot, “Robust numerical coupling of pressure and pressureless gas dynamics equations for eulerian spray DNS and LES,” *SIAM Journal on Scientific Computing*, vol. 37, no. 1, pp. B79–B102, 2015.
- [21] S. Jin and Z. Xin, “The relaxation schemes for systems of conservation laws in arbitrary space dimensions,” *Communications on pure and applied mathematics*, vol. 48, no. 3, pp. 235–276, 1995.
- [22] G.-Q. Chen, C. D. Levermore, and T.-P. Liu, “Hyperbolic conservation laws with stiff relaxation terms and entropy,” *Communications on Pure and Applied Mathematics*, vol. 47, no. 6, pp. 787–830, 1994.

Toward Efficient Simulations of Highly Magnetized Plasma

Z. Tazakkati¹, T. Pichard¹, J. Massot¹, M. Massot¹

¹CMAP, École Polytechnique, Palaiseau, France
December 18, 2025

Abstract

This work is a numerical study of strongly magnetized multi-species plasmas. Due to the multiscale nature of the problem, accurate and reliable simulations of such a system are computationally expensive. We present model based on moments of a gyrokinetic equations for the motion of electrons. This approach filters out the fast cyclotron frequencies while preserving the essential dynamics. We introduce a numerical adapted to this model when coupled with the ions dynamics and an induced electric field through Poisson equation. The effectiveness of this numerical model is illustrated through a simplified 2D test case.

Keywords: Multi-species plasma, Highly-magnetized, Gyrokinetic, Gyromoments

1 Introduction

The accurate modeling of multi-species plasmas under strong magnetic fields is a major challenge in plasma physics. The main difficulties lies in their multiscale nature, as such plasmas involve physical phenomena with time and length scales varying by several orders of magnitude. This results in systems that are stiff and difficult to solve numerically with a sufficient accuracy and at a reasonable computational cost (see e.g. [4]). First, the considered setup includes an intense magnetic field. As they are charged and they have a low mass, the electrons undergo very rapid rotations around magnetic field lines. These gyrations define two other small scales: the Larmor radius which is the radius of the circular motion of the electrons around the magnetic field lines, and the cyclotron frequency of these rotations. A second temporal disparity arises from the mass ratio between electrons and heavy plasma species: Due to their low inertia, electrons have a quick response to the system dynamics, while heavy species respond on much longer timescales. A third disparity is due to the small size of the Debye length compared of the characteristic length of the domain, which corresponds to a quasi-neutral regime. However, such configurations make commonly appear charge separation near the boundary of the domain, so-called electrostatic sheath. Its thickness is of the order of the Debye length, i.e. very thin compared to the system global dimensions (see e.g. [3]), but it prevents from treating the last two disparity at the continuous level. Therefore, extend-

ing the work in [1], we propose a global strategy to treat the different characteristic scales of such a plasma. We employ a gyromoment approach for the electrons, derived from gyrokinetic theory, to capture the rapid dynamics in a reduced formulation. This method filters out the fastest cyclotron oscillations while preserving the essential transport at larger scales. This removes the stiffness associated with the Larmor radius and cyclotron frequency. This type of strategy is established in the field of fusion plasmas (see e.g. [5]), but remains largely unexploited in cold plasmas. The stiffness related to the Debye length and the inertia of electrons are treated through an adapted numerical approach. We demonstrate how the approach previously developed in [1] extends and is physically enhanced through the introduction of a dynamic coupling with an ion population. We illustrate this strategy on a simplified 2D configuration. Such configurations highlight common challenges encountered in various applications such as for Hall effect thrusters, arcjets or plasma discharges. Thus, the key achievement of this study is the formulation of a comprehensive multiscale strategy that bridges electron gyromoment modeling and ion dynamics to accurately and efficiently capture magnetized plasma behavior.

2 Modeling

Consider a reduced model consisting of a 2D fluid of ion–electron coupled through electromagnetic interac-

tions describing the $(\mathbf{e}_x, \mathbf{e}_z)$ -plane. The magnetic field $\mathbf{B}^* = B_0^* \mathbf{e}_z$ is assumed to be constant, uniform and imposed along the z -axis. Meanwhile, the electric field $\mathbf{E}^* = \bar{\mathbf{E}}^* + \hat{\mathbf{E}}^*$ has a dominant contribution $\bar{\mathbf{E}}^* = E_0^* \mathbf{e}_y$, also constant, uniform and imposed along the y -axis, orthogonal to the 2D simulation plane. The remaining contribution $\hat{\mathbf{E}}^* = -\nabla_{\mathbf{x}^*} \phi^*$ has a smaller magnitude, and is given by Poisson's equation. This setup can be seen as an idealized modeling of a Hall thruster acceleration zone, as described in [1].




Figure 1: Representation of the field geometry

Electrons are described using the gyro-moment approximation developed in [1], which retained four moments. This approximation is derived from the hydrodynamic limit of a gyrokinetic equation. Starting from the Boltzmann equation in the regime of high electromagnetic fields, we take the asymptotic limit in which the electromagnetic fields magnitude tends to infinity. This procedure filters out fast cyclotron gyrations while preserving the essential dynamics of magnetized electrons. Then, using an entropy-based closure, we obtain a reduced fluid formulation for the electrons. This type of reduced modeling, based on the averaging of cyclotronic motion, is well-established in highly magnetized tokamak's fusion plasmas, and several codes rely on such approximations (see e.g. [2]). However, our modeling differs from these existing approaches, as it focuses on a different physical regime, characterized by both strong electric and magnetic fields. Ions are modeled using the classical Euler system of equations, derived from the hydrodynamic limit of the Boltzmann equation in a highly collisional regime. Coupling between ions and electrons is provided solely through the Lorentz force in those equations and through Poisson's equation. The resulting non-dimensional system reads as follows:

$$\begin{cases} \partial_t U^e + \frac{1}{\varepsilon} \nabla_{\mathbf{x}} \cdot \mathbf{F}^e = \frac{1}{\varepsilon} S^e, \\ \partial_t U^i + \nabla_{\mathbf{x}} \cdot \mathbf{F}^i = S^i, \\ -\Delta_{\mathbf{x}} \phi = \frac{\rho^i - \rho^e}{\chi}. \end{cases} \quad (2.1)$$

Two dimensionless scaling parameters have been introduced: The parameter $\varepsilon \ll 1$ corresponds to the square root of the mass ratio between electrons and ions, quantifying the difference in inertia between the two species. Electrons, being much lighter, react on much shorter timescales than ions. The parameter $\chi \ll 1$ represents the normalized Debye length squared.

Its small value reflects the fact that the Debye length is much smaller than the system's characteristic length scale. Physically, this corresponds to a quasi-neutral plasma. Note that, under realistic boundary conditions, a charge separation between the electrons and the ions develops at the boundary, leading to a local breakdown of quasi-neutrality. The vectors of unknowns U^e and U^i involved in (2.1) are described below.

Electrons

The electron state vector is defined as

$$U^e = (\rho_e, \rho_e u_z, \rho_e T_x, \rho_e \mathcal{E}_z)^T,$$

where ρ_e is the electron density, u_z the mean velocity in the z -direction, T_x the temperature in the x -direction, and \mathcal{E}_z the electron energy in the z -direction. The corresponding fluxes in the x - and z -directions are

$$F_x^e = (\rho_e, \rho_e u_z, \rho_e T_x, \rho_e \mathcal{E}_z)^T,$$

$$F_z^e = (\rho_e u_z, \rho_e \mathcal{E}_z, \rho_e T_x u_z, \rho_e u_z (3\mathcal{E}_z - 2u_z^2))^T.$$

The electron source term reads

$$S^e = (0, -\rho_e \tilde{E}_z, \varepsilon \rho_e (T_{\text{eq}} - T_x), -2\rho_e u_z \tilde{E}_z + \varepsilon \rho_e (T_{\text{eq}} - T_z))^T.$$

Note that the source term is stiff only because of the electric field of order $\mathcal{O}(1/\varepsilon)$, unlike the relaxation term relating the two directional temperatures of order $\mathcal{O}(1)$.

The model for electrons consists of a constant drift in the x -direction, i.e. the $\mathbf{E} \times \mathbf{B}$ direction, and a classical Euler model in the z -direction. As detailed in [1], this results from the limit of a dominant and constant electromagnetic field tending to infinity.

For this reason, only the velocity u_z in the z -direction is considered. In the x -direction, the constant drift velocity v_D is not an independent variable.

Ions

The gyrokinetic approximation applies not for the ions dynamics. This approximation canceled the y -velocity in the electron dynamics, while it is non-zero in the ion dynamics. Therefore, even if we are interested in a 2D model, the electromagnetic field makes the dynamics of ions intrinsically three-dimensional. We consider a 3D model for the ions with energy distributed over the three spatial degrees of freedom. However we assume that the mean velocity u_y in the y -direction is fully prescribed. This assumption allows the system to be reduced to 2D equations. Accordingly, the ions are modeled using a four-moment Euler formulation, where the total energy is given by

$$\mathcal{E}^i = \rho^i (5T^i + |\mathbf{u}^i|^2).$$

Here, the ion state vector consists of the density, momentum components in the x - and z -directions, and the total energy

$$U^i = (\rho^i, \rho^i u_x^i, \rho^i u_z^i, \mathcal{E}^i)^T.$$

The corresponding fluxes in the x - and z -directions are defined as

$$F_x^i = (\rho^i u_x^i, \rho^i ((u_x^i)^2 + T^i), \rho^i u_x^i u_z^i, u_x^i (\mathcal{E}^i + 2\rho^i T^i))^T,$$

$$F_z^i = (\rho^i u_z^i, \rho^i u_x^i u_z^i, \rho^i ((u_z^i)^2 + T^i), u_z^i (\mathcal{E}^i + 2\rho^i T^i))^T.$$

Finally, the ion source term, which accounts for the influence of electric and magnetic fields, reads

$$S^i = (0, \rho^i (\tilde{E}_x + B u_y^i), \rho^i \tilde{E}_z, 2\rho^i (u_x^i \tilde{E}_x + u_z^i \tilde{E}_z))^T.$$

The scaling of (2.1) shows that the gyrokinetic approach reduces only part of the system stiffness by filtering out cyclotron frequencies, and some remain at the fluid scale due to the very multiscale nature of the problem, albeit to a lesser extent. This residual stiffness is addressed through the design of a specific numerical strategy.

3 Numerical Methods

The system (2.1) can be viewed as a particular two-fluid Euler-Poisson system, in a regime where the Debye length is small and the electron inertia is very low. There is an extensive literature on this topic (see e.g [6, 7, 3]) devoted to the development of asymptotic-preserving (AP) numerical schemes suitable for this type of system. In this context, we developed in [1] a numerical scheme that solves the system (2.1) considering only the electrons, with the ions treated as a fixed background. The ion population is solved here dynamically. Thus, the stiff electron flux and the electric field are treated implicitly using the same strategy developed in [1]: the linearity of the transport in the x -direction allows for a simpler implicit treatment of these fluxes. An upwind flux in the x -direction and a HLL flux in the z -direction are used. The electron dynamics is solved together with the Poisson equation using an iterative fixed-point method. This allows us to effectively handle the electric field in an implicit manner. For the ions, the absence of stiffness (associated with their lower inertia) enables us to compute the fluxes explicitly, using a HLL solver in both directions, while treating implicitly the source term, using the result given by the electron-Poisson solver.

4 Test-case and numerical results

Consider a 2D domain $[-\pi, \pi] \times [-\pi, \pi]$ with periodic boundary conditions. The domain is discretized using 400 points in the x -direction and 100 points in the z -direction. This anisotropic discretization is necessary for the stability of the numerical scheme, as discussed in [1]. The ion population is initialized with a homogeneous state, defined by

$$\rho^i = 0.5, \quad u_x^i = u_z^i = 0, \quad u_y^i = 1, \quad T^i = 1.$$

The electrons are initialized with the same configuration except for a density given by

$$\rho_e(x, z) = 0.5 + 0.1 \exp(-x^2) \exp(-z^2) + \lambda,$$

where λ is a constant that guarantee the global electro-neutrality. The other quantities are initialized homogeneous

$$u_z = 0, \quad T_x = T_z = 1.$$

The small parameters are $\varepsilon = \chi = 10^{-2}$ and $B = 1$. This test case aims at studying the response of a quasi-neutral electron-ion plasma in which the electrons are forced to drift at a constant velocity in one direction, while the global density equilibrium is perturbed by a localized spatial accumulation of electronic charge. We observe, in a simplified case, how a local loss of electro-neutrality evolves under a realistic, multiscale physical regime, which is typically challenging and hardly addressed in the literature. The total computation time is 251 minutes using a Python implementation running sequentially on an M2 Pro CPU. The code is available at the following GitLab repository: <https://gitlab.labos.polytechnique.fr/zoubair.tazakkati/toward-efficient-simulations-of-highly-magnetized-plasma>

5 Discussion and Perspectives

Fig 2 shows the time evolution of the ion and electron densities. Two characteristic timescales can be distinguished in the dynamics observed. The first is a rapid early phase (before $t \approx 0.1$), beginning with an excess of electron density at the center of the domain. During this step, the electrons respond quickly. Under the combined effect of the pressure gradient and the Lorentz force, density waves propagate in the z -direction and are reflected toward the center of the domain by the periodic boundaries. Simultaneously, all the electrons drift globally in the x -direction at a constant speed. However, the ions hardly move and their density remains nearly uniform. Then, in a second timescale (from $t \approx 0.1$), the ions start to respond more significantly. They gradually follow the motion of the electrons in the x -direction, and their density variations become more pronounced, illustrating the delayed response due to their larger mass and inertia. These two timescales highlight the difference in mobility between electrons and ions. Electrons react almost instantaneously to forces, while ions adapt more slowly, leading to complex coupled $x - z$ dynamics. Over time, these interactions produce periodic patterns in electron and ion density. This simulation indicates that the model can capture small-scale phenomena while maintaining relatively low computational costs. The dynamics of small cyclotronic oscillations are implicitly accounted for through the use of a gyromoment model. Overall, these numerical results suggest that the proposed scheme remains stable and effective in regimes involving small parameters. However, such a numerical analysis is postponed to future work to fully confirm this robustness.

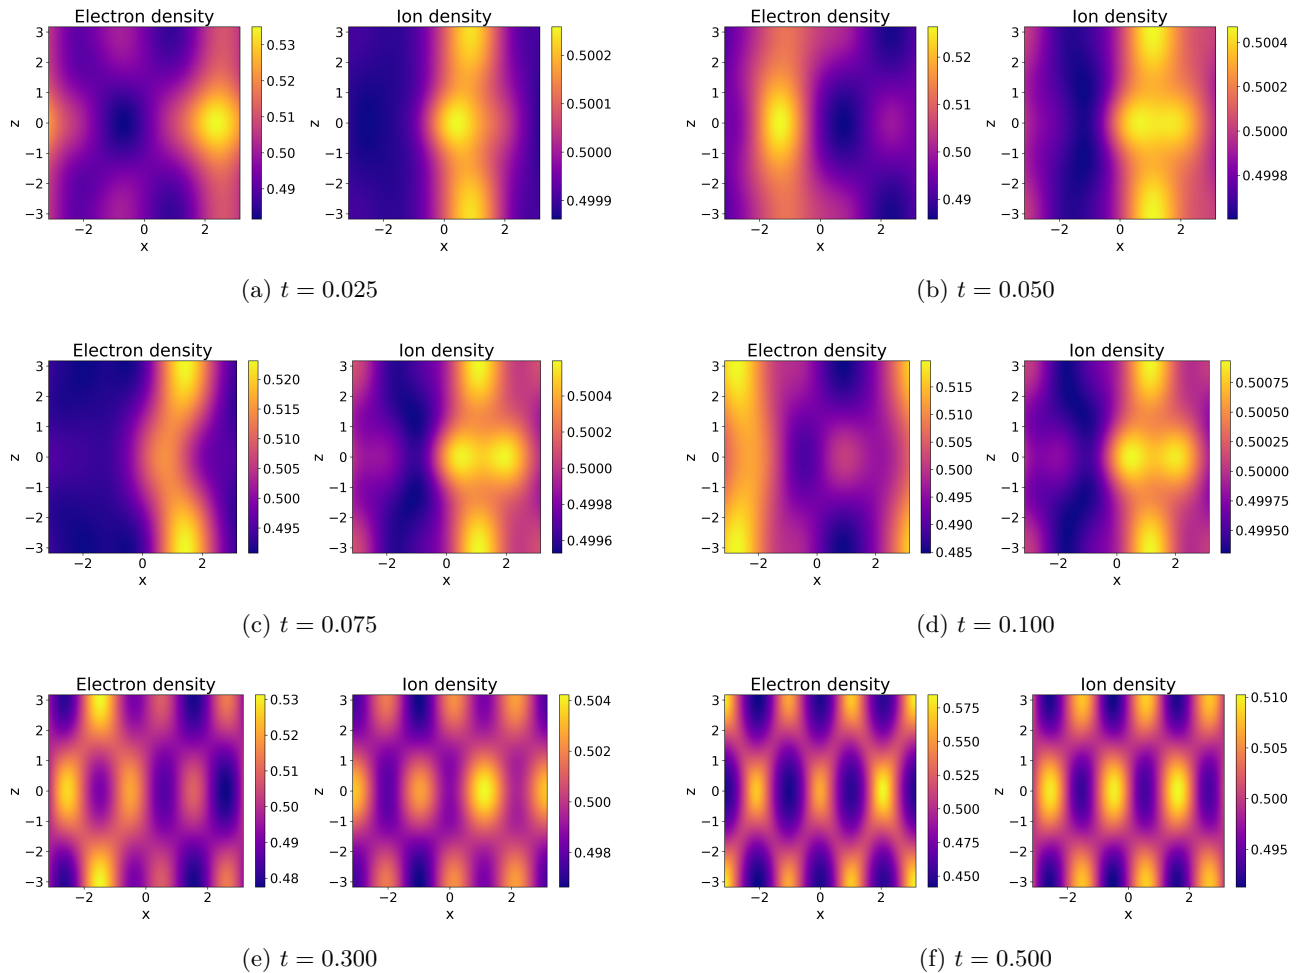


Figure 2: Evolution of the electron and ion densities.

6 Conclusion

This study illustrates the potential of our unified approach to model strongly magnetized, multi-species plasmas. In a simplified 2D configuration, the method captures key aspects of the coupled dynamics of different species across multiple scales. These results suggest that the model can handle complex multiscale regimes. While further validation and extensions are needed, the approach appears promising for more realistic simulations, such as those involving the sheath and applications to space propulsion.

Acknowledgments

The authors thank Louis Reboul, Julien Clotuche, Thierry Magin, Irene Gamba and Krishnan Gopalan for fruitful discussions on plasma modeling, transport phenomena, and numerical methods related to Hall thrusters. This work has been supported by the Fondation Mathématiques Jacques Hadamard through the Junior Scientific Visibility programme, as well as the Agence Innovation Défense, the CIEDS project OPEN-NUM-DEF, and the HPC@Maths Initiative of the Fondation École Polytechnique.

References

- [1] Z. Tazakkati et al, A gyromoment approach for electron dynamics in low-temperature $\mathbf{E} \times \mathbf{B}$ plasmas of Hall thrusters. *submitted*, 2025.
- [2] V. Grandgirard et al, Global full-f gyrokinetic simulations of plasma turbulence. *Plasma Physics and Controlled Fusion*, vol. 49, no. 12B, p. B173, Nov. 2007.
- [3] L. Reboul, Development and analysis of efficient multi-scale numerical methods, with applications to plasma discharge simulations relying on multi-fluid models. *Ph.D. Thesis, Institut Polytechnique de Paris*, 2024. <http://www.theses.fr/2024IPPAX134/document>
- [4] F. Deluzet et al, Numerical methods and macroscopic models of magnetically confined low temperature plasmas. *Kinetic and Related Models*, vol. 16, no. 5, pp. 624–653, 2023.
- [5] A. C. D. Hoffmann et al, Gyrokinetic simulations of plasma turbulence in a Z-pinch using a moment-based approach and advanced collision operators. *Journal of Plasma Physics*, vol. 89, no. 2, p. 905890214, 2023.

- [6] P. Crispel et al, An asymptotic preserving scheme for the two-fluid Euler–Poisson model in the quasineutral limit. *Journal of Computational Physics*, vol. 223, no. 1, pp. 208–234, 2007.
- [7] A. Alvarez Laguna et al, An asymptotic preserving well-balanced scheme for the isothermal fluid equations in low-temperature plasmas at low pressure. *Journal of Computational Physics*, vol. 419, p. 109634, 2020.

Study of the Electrode Kinetic Layer in Arc Discharges

Development of Kinetic-Based Boundary Conditions
for the Electrodes of the NASA Ames Arc Heater

Pietro Parodi^{1,2} Jan Skácel^{1,3,5} Federico Bariselli¹ Bruno Dias⁴ Thierry Magin^{1,5}

¹Aeronautics and Aerospace Department, von Karman Institute for Fluid Dynamics, Belgium

²Department of Mathematics, KU Leuven, Belgium

³Department of Plasma Physics and Technology, Masaryk University, Czechia

⁴Analytical Mechanics Associates, NASA Ames Research Center, USA

⁵Aero-Thermo-Mechanics Department, Université libre de Bruxelles, Belgium

December 4, 2025

Abstract

The NASA Ames arc heater facility operates by heating a working gas at temperatures of the order of 10 000 K through a DC electric arc sustained in a constrictor tube between two copper electrodes. Recent simulations using fluid methods were conducted to investigate the mechanisms of arc instability. These were able to capture the main features of the discharge; however, due to the simplicity of the boundary conditions used, they did not faithfully represent the attachment of the arc to the electrodes, which critically determines their erosion and premature failure. With this work, we aim to study the physics in the thin layer where the arc attaches to the electrodes through 1D particle-based kinetic simulations. This method allows us to include relevant physical mechanisms, such as electrode evaporation, thermionic emission, and gas ionization and excitation, and to investigate their effects. We find that a large voltage drop can develop in the cathode layer of the discharge, which also serves as the ionization source sustaining the arc. We also find evidence that the metal evaporating from the electrode can strongly influence the plasma discharge due to its low excitation energy threshold, providing a strong pathway for energy exchange between electrons and heavy particles. Finally, we provide a strategy for using small-scale simulations to improve the boundary conditions in the fluid simulations of the complete arc.

Keywords: electric arcs, arcjet, plasma, evaporation, particle methods, PIC, DSMC, Monte Carlo

1 Introduction

Arc-heater facilities at NASA Ames are used to test spacecraft Thermal Protection Systems (TPS) under relevant atmospheric re-entry conditions. The high enthalpy plasma flow in such devices is typically modeled using Magnetohydrodynamics (MHD). Recently, the ARCHES code, a state-of-the-art parallel, unstructured, finite volume solver based on the OpenFOAM library, has been used to simulate the arc heater [1]. The study highlighted fluid-dynamic instabilities arising from the interaction between the injected cold gas and the hot plasma, as well as various plasma instabilities, including kink and sausage modes. Such simulations, however, considered simplified boundary conditions at the discharge electrodes and neglected kinetic effects in the bulk of the fluid. At length scales comparable to

the mean free path, within the Knudsen layer, particle velocity distributions can deviate from the Maxwell-Boltzmann distribution. Similarly, at scales on the order of the Debye length, in the plasma sheath, the plasma's quasi-neutrality condition is not satisfied because the surrounding electrons do not promptly screen positive charges.

The objective of this work is to study these phenomena at the mesoscopic level under conditions relevant to the NASA arc heater and, eventually, to extract boundary conditions for the macroscopic flow variables of the MHD model. First, we aim to simulate a 1D domain using particle-based kinetic methods that include the electrode kinetic region. We will use the hybrid Particle-in-Cell – Direct Simulation Monte Carlo – Monte Carlo Collision (PIC-DSMC-MCC) method, implemented in the VKI code PANTERA [2]. Some examples of similar

studies of the near-electrode region, relevant for applications including high-power circuit breakers, pulsed plasma thrusters, metal welding, and discharge lamps, can be found in [3, 4, 5, 6, 7, 8].

The heat flux at the arc attachment points leads to electrode material evaporation, which can diffuse in the gaseous phase, ionize, or condense on the surface [4]. Therefore, we will initially consider evaporation from the metallic surface. First, in Section 3.1, where the evaporated vapor is the only gas present, and then in Section 3.2, where the metal evaporates in a background of argon at 1 atm of pressure. We will then consider the plasma discharge, first with an imposed gas background in Section 3.3, and then including coupling to the background gas and metal vapor in Section 3.4.

Analysis of Characteristic Scales

First, it is helpful to determine the approximate discharge properties in the cathode region of the NASA Ames arc heater. From the simulations shown in [1], a gas pressure of the order of 5-10 bar is estimated. The plasma temperature in the bulk of the arc is approximately 10 000 K, decreasing to 3 000 K in the electrode region. The number density of the background gas can be expected to be approximately 10^{25} m^{-3} . The total current of the arc is of the order of 5 kA. The mean free path for collisions with the background is $\lambda \approx 6 \times 10^{-8} \text{ m}$, and the Debye length, at an ionization degree of 0.1%, is $\lambda_D \approx 7 \times 10^{-8} \text{ m}$. We therefore estimate that the kinetic layer will have a thickness of a few micrometers.

2 Methods

The PIC-DSMC-MCC code PANTERA [2, 9], developed at VKI, is used for all kinetic simulations performed in this work. All simulations are performed in 1D-3V on a uniform grid. The results shown in Sections 3.2 and 3.3 use Monte Carlo collisions [10, 11] with an imposed background. Otherwise, all species are simulated using particles. We use species-specific macroparticle weights when there is a large density disparity between species, for instance, between plasma and neutral particles. The collisional procedure we use in this case is adapted from that of Araki and Martin [12].

Evaporation of copper is modeled by injecting from a half-Maxwellian velocity distribution at the wall temperature T_w from a reservoir with a density corresponding to the vapor pressure at T_w . At 3 000 K, this is $p_v \approx 200 \text{ Pa}$ for copper. We consider that all copper atoms reaching the surface recondense.

Thermionic emission is modeled by emitting electrons from a half-Maxwellian distribution at the temperature of the wall T_w , with current density j_e determined by the Richardson-Laue-Dushman equation, including the Schottky effect.

$$j_e = A_G T_w^2 \exp\left(-\frac{W - \Delta W}{k_B T_w}\right), \quad (2.1)$$




Figure 1: Ratio of the net evaporation mass flow rate \dot{m}_{vap} to the emitted mass flow rate \dot{m}_e as a function of the Mach number at the edge of the Knudsen layer.

where $W = 4.55 \text{ eV}$ is the work function of copper, $\Delta W = [q_e^3 E / (4\pi\epsilon_0)]^{1/2}$ is the change in the potential barrier due to the electric field on the surface of the cathode, $A_G = \lambda_R 4\pi m_e k_B^2 q_e / h^3$ is a universal constant and $\lambda_R \approx 0.5$ is a correction factor for metal surfaces.

3 Results and Discussion

In the following sections, we first consider the electrode evaporation separately, and then include the plasma discharge.

3.1 Evaporation of Copper

As a first step in modeling the electrode, we consider the case of evaporation of the copper electron into the pure copper vapor it generates. At the solid boundary, we impose the evaporation rate. At the boundary facing the fluid, the “piston” boundary condition is used, as described by Frezzotti in [13], which imposes a drift velocity U_∞ for the evaporating gas. This works by reflecting particles specularly on an imaginary surface that is constantly moving at a velocity U_∞ . Collisions using the Variable-Soft-Sphere (VSS) model are considered in the gas phase. Collision parameters for copper are taken from Gosma *et al.* [14]. A domain of the size of approximately 100λ is simulated.

We find an excellent agreement between our results and those of Dias [15]. We identify a region of thermal nonequilibrium near the evaporating wall, where the temperature is higher in the direction parallel to the wall than in the perpendicular direction. In Figure 1, we plot the fraction of the net evaporating mass flow (emitted minus condensing) to the emitted mass flow, also in agreement with the results of Dias [15]. This measurement is an example of a macroscopic relationship that could be used as a boundary condition.

3.2 Evaporation Into a Background Gas

Here, we consider the case in which the copper electrode evaporates into a background of argon gas under

static conditions. In this case, the evaporating copper acts essentially as a trace species diffusing into the background gas. As in the previous case, we impose the evaporating flux at the wall and a “piston” boundary condition at the fluid-facing boundary. We perform simulations in a domain of roughly 20λ with a background argon gas at 1 atm and 3000 K, the same as the wall temperature. We perform simulations with piston velocities U_∞ ranging from 30 to 1500 m/s. The results, shown in Figure 2, indicate that the evaporating copper essentially obeys Fick’s law of diffusion. In 1D, assuming copper is at the limit of infinite dilution, a linear density profile is expected, as shown in the figure. The results suggest that the evaporating copper is everywhere in thermal equilibrium with the background gas. The net evaporating flux appears to be insensitive to the piston velocity, with a value of $j_{\text{Cu}} = 0.127 \text{ kg m}^{-2}\text{s}^{-1}$. For piston velocities above 500 m/s, almost no particle is reflected into the domain, and the results are essentially unaffected, as shown by the curves at 500 and 1500 m/s.




Figure 2: Number density of copper for the case in which it evaporates into a background of argon at 1 atm and 3000 K.

These results suggest that the correct way to impose an evaporation boundary condition for the arc heater would be to assume (near-)thermal equilibrium of copper with the background gas in the “kinetic” layer, and that whatever concentration gradient is imposed at the fluid boundary would be maintained throughout the layer.

3.3 Plasma Discharge With a Fixed Background

As a first step towards simulating the electrode kinetic layer, we model a 1D plasma discharge that includes thermionic emission from the cathode. This was inspired by microdischarge simulations such as that of Zhong *et al.* [16]. A potential difference is applied across the domain boundaries via an external circuit comprising a series ballast resistor and a parallel capacitor. A background of argon at 1 atm and at the temperature of the cathode T_w is imposed, corresponding to a density $n_{\text{Ar}} = 3.19 \times 10^{24} \text{ m}^{-3}$. MCCs are

Table 1: Elastic and inelastic processes included in the simulation of the microdischarge.

Name	Process	E_a
IZ1	$e^- + \text{Ar} \rightarrow e^- + \text{Ar}^+ + e^-$	15.8 eV
EX1	$e^- + \text{Ar} \rightarrow e^- + \text{Ar}^*$	11.5 eV
EL1	$e^- + \text{Ar} \rightarrow e^- + \text{Ar}$	0
EL2	$\text{Ar}^+ + \text{Ar} \rightarrow \text{Ar}^+ + \text{Ar}$	0

computed between the plasma and the background gas (see Table 1).

The results, averaged at steady-state, are shown in Section 3.3. The cathode and anode sheaths can be distinguished by inspecting the potential drop across each. Since the size of the discharge is limited, the quasineutral plasma column is limited to the center of the discharge at $0.4 < x/L < 0.5$. This value of current density produced a discharge with a plasma density of the order of 10^{18} m^{-3} , so with a rather low ionization degree of $\approx 10^{-6}$. The temperature plot shows that ions equilibrate with the background neutrals through elastic collisions. Electrons, on the other hand, do not exchange sufficient energy with the neutrals within the domain to equilibrate with the heavies. Godyak [17] proposed a simplified expression for the electron temperature relaxation length λ_e that accounts for all electron energy loss mechanisms in a discharge. We report it here as written by Chabert and Braithwaite [18],

$$\lambda_e = \lambda_{\text{en}} \left[\frac{2m_e}{m_i} + \frac{\nu_{ee}}{\nu_m} + \frac{2}{3} \left(\frac{q_e \varepsilon_{\text{exc}}}{k_B T_e} \right) \frac{\nu_{\text{exc}}}{\nu_m} + \frac{2}{3} \left(\frac{q_e \varepsilon_{\text{iz}}}{k_B T_e} \right) \frac{\nu_{\text{iz}}}{\nu_m} \right]^{-1/2}, \quad (3.1)$$

where λ_{en} is the electron-neutral mean free path. The first term in parentheses is due to elastic collisions. The second term arises from electron-electron (Coulomb) collisions. The third and fourth terms are due to electron energy loss by inelastic collisions (ionization and excitation). Here, we do not show the term due to energy loss at the boundaries. The elastic term for argon is of the order of 10^{-5} , and the Coulomb collision term is also typically considered negligible. However, this should also be verified in the conditions of interest. At the temperature of $\approx 5 \text{ eV}$ measured in this case, the terms due to ionization and excitation are also small ($\nu_{\text{iz}}, \nu_{\text{exc}} \ll \nu_m$), since the electron temperature is well below the ionization and first excitation thresholds $\varepsilon_{\text{exc}} = 11.5 \text{ eV}$ and $\varepsilon_{\text{iz}} = 15.8 \text{ eV}$. However, metal vapors, which typically have much lower thresholds for electronic excitation, could play a role in equilibrating electrons. This motivated the simulation in the following section.

3.4 Plasma Discharge Coupled to the Neutral Background

To investigate the role of excitation reactions in energy transfer between electrons and heavy particles, we in-




Figure 3: Number densities, electric potential, and temperatures in the plasma discharge with imposed background.

clude the lowest-energy excitation reaction for copper, at $\varepsilon_{\text{exc}} = 1.5$ eV. We then need to include a mechanism for energy transfer from the excited copper to the other heavy species. This requires a coupled solution of the plasma and neutral flow. Therefore, we also simulate the background argon and copper as particles, with the weight of Ar neutrals increased by a factor of 2000 and that of Cu by a factor of 10.

The results are shown in Figure 4. In this case, a distinct quasineutral plasma column forms, where electrons cool to a temperature below 1 eV. However, while ions further equilibrate to the neutrals, electrons remain out of thermal equilibrium.

3.5 Formulation of the Boundary Conditions

In the previous sections, we presented the simulation at the kinetic level of conditions representative of the layer in proximity to the arc heater electrodes. The initial goal of formulating MHD boundary conditions would be to run simulations over a matrix of conditions to extract functional relationships between the parameters. Although we did not perform this task, we now have a better sense of which variables could serve as inputs and outputs. In the simulations, one must fix the wall temperature T_w , the driving voltage ΔV , the resistance of the external circuit R , the temperature(s) T_f and neutral density $n_{n,f}$ at the fluid interface, and the density of copper vapor $n_{\text{Cu},f}$ (or equivalently its diffusion velocity). Electrons must be adequately reflected to avoid the formation of an artificial anode. The simulation in turn will provide the

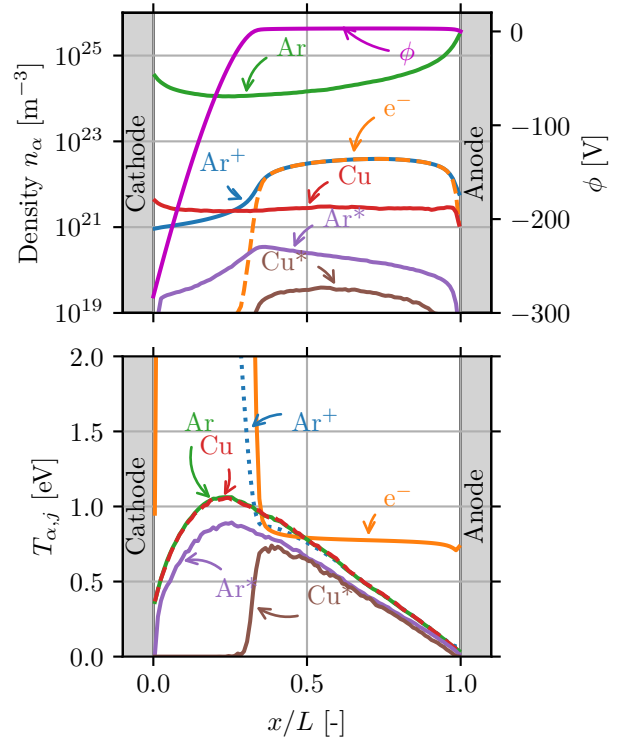


Figure 4: Number densities, electric potential, and temperatures in the plasma discharge with neutral-plasma coupling.

electrode evaporation rate \dot{m}_{vap} , (electron) arc current at the fluid interface $j_{e,f}$ and heat fluxes at the wall \dot{q}_w and at the fluid interface \dot{q}_f . These will generally differ due to inelastic collisional processes taking place in the kinetic layer. Notice that such a model could also be coupled to a material solver to yield self-consistently the electrode temperature based on \dot{q}_w .

4 Conclusions and Future Work

In this work, we investigated the kinetic layer near the electrodes in high-pressure arc discharges by performing increasingly complex particle-based simulations. The results showed that the background gas fundamentally alters the dynamics of evaporation of the electrode material with respect to the “strong blowing” case considered by Dias [15]. In turn, the metal vapors play an important role in electron-heavy energy transfer through excitation reactions and could alter the plasma dynamics in the electrode layer. These findings suggest that evaporation and plasma attachment are coupled and should be investigated as such. For arc simulations, we considered a setup analogous to a microdischarge between two electrodes. It would be convenient to consider only the region between an electrode and the bulk of the discharge. However, formulating the proper boundary conditions for the fluid side proved challenging and is left for future work.

Acknowledgments

The work of PP was supported by a travel grant from the VKI Alumni Association and by a Strategic Basic PhD fellowship (ref. 1S24022N) from the Research Foundation–Flanders (FWO). JS’s work was supported by Erasmus+ funding through the VKI. We gratefully acknowledge Dr. Nagi N. Mansour for insightful discussions and his enthusiasm for the project, which have been a source of motivation.

References

- [1] J. B. Meuris, A. Alvarez Laguna, M. Panesi, and N. N. Mansour, “Three-dimensional unsteady model of arc heater plasma flow,” *Aerospace Science and Technology*, vol. 123, p. 107465, 2022.
- [2] P. Parodi, S. Boccelli, F. Bariselli, and T. E. Magin, “Pantera: A PIC-MCC-DSMC software for the simulation of rarefied gases and plasmas,” *SoftwareX*, vol. 31, p. 102244, 2025.
- [3] N. A. Almeida, M. S. Benilov, and G. V. Naidis, “Unified modelling of near-cathode plasma layers in high-pressure arc discharges,” *Journal of Physics D: Applied Physics*, vol. 41, no. 24, p. 245201, 2008.
- [4] I. I. Beilis, “Kinetics of plasma particles and electron transport in the current-carrying plasma adjacent to an evaporating and electron emitting wall,” *IEEE Transactions on Plasma Science*, vol. 34, no. 3, pp. 855–866, 2006.
- [5] J. Huo, Y. Wang, and Y. Cao, “3D computational study of arc splitting during power interruption: the influence of metal vapor enhanced radiation on arc dynamics,” *Journal of Physics D: Applied Physics*, vol. 54, no. 8, p. 085502, 2020.
- [6] P. Sarrailh, L. Garrigues, G. J. M. Hagelaar, J. P. Boeuf, G. Sandolache, and S. Rowe, “Sheath expansion and plasma dynamics in the presence of electrode evaporation: Application to a vacuum circuit breaker,” *Journal of Applied Physics*, vol. 106, no. 5, p. 053305, 2009.
- [7] S. A. Self and L. D. Eskin, “The boundary layer between electrodes and a thermal plasma,” *IEEE Transactions on Plasma Science*, vol. 11, no. 4, pp. 279–285, 1983.
- [8] I. L. Semenov, I. V. Krivtsov, and U. Reisgen, “Numerical study of the anode boundary layer in atmospheric pressure arc discharges,” *Journal of Physics D: Applied Physics*, vol. 49, no. 10, p. 105204, 2016.
- [9] J. Skacel, P. Parodi, G. Gangemi, F. Bariselli, Z. Bonaventura, and T. Magin, “Simulation of space platform charging in Very Low Earth Orbit using stochastic particle methods,” 2025.
- [10] V. Vahedi and M. Surendra, “A Monte Carlo collision model for the particle-in-cell method: applications to argon and oxygen discharges,” *Computer Physics Communications*, vol. 87, no. 1, pp. 179–198, 1995. Particle Simulation Methods.
- [11] P. Parodi and F. Petronio, “Step-by-step verification of particle-in-cell Monte Carlo collision codes,” *Physics of Plasmas*, vol. 32, no. 1, p. 013902, 2025.
- [12] S. J. Araki and R. S. Martin, “Interspecies fractional collisions,” *Physics of Plasmas*, vol. 27, no. 3, p. 033504, 2020.
- [13] A. Frezzotti, “Kinetic theory description of the evaporation of multi-component substances,” in *Proceedings of the 20th International Symposium on Rarefied Gas Dynamics, Beijing, China, 1996*.
- [14] M. Gosma, K. Swaminathan Gopalan, S. Subramaniam, and K. Stephani, “Recommended direct simulation Monte Carlo collision model parameters for planetary entry and related applications,” *Physics of Fluids*, vol. 37, no. 3, p. 036156, 2025.
- [15] B. Dias, *Thermal ablation and radiation modeling of meteor phenomena*. PhD thesis, Université catholique de Louvain, 2020.
- [16] Y. Zhong, H. Wu, X. Li, J. Gao, W. Jiang, Y. Zhang, and G. Lapenta, “Numerical characterization of the breakdown process of DC-driven micro-discharges sustained by thermionic emission,” *Journal of Physics D: Applied Physics*, vol. 55, no. 21, p. 215203, 2022.
- [17] V. A. Godyak, “Nonequilibrium EEDF in gas discharge plasmas,” *IEEE Transactions on Plasma Science*, vol. 34, no. 3, pp. 755–766, 2006.
- [18] P. Chabert and N. Braithwaite, *Physics of Radio-Frequency Plasmas*. Cambridge University Press, 2011.

Shock Tube Radiation Modeling and Sensitivity Analysis

Radiation Modeling of H_2/He Entry Flows within an Uncertainty Quantification Framework

Bruno Fontaine^{1,2,3} Julien Clotuche^{1,2,3} Federico Bariselli¹ Koen Hillewaert^{1,3} Thierry Magin^{1,2}

¹Aeronautics and Aerospace Department, von Karman Institute for Fluid Dynamics, Belgium

²Aero-Thermo-Mechanics Department, Université libre de Bruxelles, Belgium

³Aerospace and Mechanical Engineering Department, Université de Liège, Belgium

October 31, 2025

Abstract

Future missions such as the Uranus Orbiter and Probe will require robust modeling tools to predict the extreme thermal environments encountered during atmospheric entry at velocities approaching 25 km/s. Under these conditions, radiative heating can contribute significantly to the total heat flux, making accurate radiation modeling essential for thermal protection system design. This work focuses on developing a predictive framework for the radiative output of H_2/He entry flows and on identifying the chemical processes that most strongly influence radiative heating. Radiance predictions are computed by a one-way coupling of a flow model with detailed chemistry to NASA's NEQAIR radiation software. A variance-based sensitivity analysis was performed to quantify the contribution of reaction rate uncertainties to variations in radiative output. Model results show the model's ability to reproduce EAST experimental data, and a preliminary sensitivity analysis identifies H_2 dissociation and electron-impact excitation as the groups of reactions with the most significant influence on radiative intensity, establishing the foundation for future calibration of reaction rates using shock tube experiments.

Keywords: ice giant, reentry, state-to-state, collisional-radiative model, plasma

1 Introduction

The study and exploration of Ice Giants has been identified as a priority in the National Academy of Sciences' Decadal Survey [1]. Planned missions such as the Uranus Orbiter and Probe [2] reflect this growing interest. Among other goals, this mission would provide in-situ measurements of Uranus's atmosphere, providing invaluable data about the formation of our solar system. Designing a probe for entry into an outer planet's atmosphere poses significant challenges due to the extremely high entry velocities, which can reach about 25 km/s for Uranus. Under such flow conditions, it is essential to quantify the different heating mechanisms acting on the probe. Shock tube tests have shown that, under these conditions, radiative heating can contribute up to 20% of the peak heating [3], highlighting the need for accurate radiation modeling.

Extensive experimental campaigns [3–6] have provided valuable data for model validation. Several research groups have proposed approaches to predict radiative heat flux in these flow regimes [7–12]. However, no consensus yet exists on the dominant mechanisms, and reaction rate coefficients can vary widely across

sources. For instance, Colonna *et al.* [7] proposed a 9-species model emphasizing the role of H_3^+ , H_2^+ , and H^- ions opening new paths for the generation of free electrons, while Carroll *et al.* [11] achieved similar agreement with EAST measurements [3] using a 6-species model excluding these ions, but boosting reaction rates of electron-driven processes within their uncertainty bounds.

These modeling discrepancies highlight the need to identify dominant reactions in H_2/He entry flows, and to determine accurate reaction rates consistent with experimental data. This work, therefore, aims to build a predictive model of the radiative output of Ice Giant entry flows and to perform a sensitivity analysis to identify the reactions that most strongly affect it. Ultimately, these developments lay the groundwork for future calibration of reaction rates within an uncertainty quantification framework.

This report is structured as follows: Section 2 describes the radiation model and sensitivity analysis methodology; Section 3 presents simulation results and preliminary sensitivity analysis; and Section 4 discusses the findings and outlines future work.

2 Methods

In this section, we describe the modeling methodology for predicting spectrally resolved radiance based on pre-shock flow conditions and reaction rate coefficients. The model consists of a fully coupled flow and chemistry model, with its outputs passed to the radiation model via one-way coupling. We then describe the methodology for the preliminary sensitivity analysis to identify the reactions with the highest impact on the radiance predictions.

2.1 Flow and chemistry model

The flow model is a multispecies, two-temperature, 1D steady state Euler model for thermal relaxation past a shock wave [13,14]. The rotational temperature, T_r , is assumed equal to the heavy translational one, T_t , while the vibrational temperature, T_v , equal to the electronic temperature, T_e , and to the free-electron temperature, T_{fe} . This model offers the advantage of being an inexpensive set of ordinary differential equations that can be solved with a space-marching method, while still capable of capturing complex chemistry and thermal non-equilibrium effects. The initial conditions are determined from the Rankine-Hugoniot jump relations based on the specified inflow conditions. The equations solved are mass conservation for each species, momentum conservation, total and internal energy conservation:

$$\frac{d}{dx}(\rho_i u) = \omega_i, \quad i \in \mathcal{S} \quad (2.1)$$

$$\frac{d}{dx}(\rho u^2 + P) = 0 \quad (2.2)$$

$$\frac{d}{dx}(\rho H u) = \mathcal{P} \quad (2.3)$$

$$\frac{d}{dx} [u(\rho e^v + \rho e^e + \rho_e e_e^t)] + P_e \frac{du}{dx} = \Omega^{\text{in}} \quad (2.4)$$

where x (m) is the axial position, ρ_i (kg/m³) is the density of species $i \in \mathcal{S}$ in the set of species considered, ρ (kg/m³) the total density, ω_i (kg/(m³s)) the net production rate of species i , u (m/s) is the flow velocity, P (Pa) the pressure, H (J/kg) the total enthalpy, \mathcal{P} (W/m³) the volumetric heat source term, e^v , e^e , e_e^t (J/kg) the components of the internal energy, respectively the vibrational, electronic and free electron translational specific energies, P_e (Pa) the electron pressure, and Ω^{in} (W/m³) the volumetric power transferred from internal energy to translational energy.

The species considered in the model are H₂, H₂⁺, H₃⁺, H, H⁺, H⁻, He, He⁺ and e⁻, defining the species set \mathcal{S} , as proposed by Colonna *et al.* [7]. A state-to-state (StS) approach is used for the first 10 energy levels of atomic hydrogen H($n = 1 - 10$), treated as pseudo-species, to allow for non-Boltzmann distributions of excited states as observed experimentally [3].

The collisional reactions taken into account include dissociation of H₂ by heavies and electrons [15,16], dissociative recombination [17], electron- and heavy-impact ionization for H₂ and helium [16,18], electron-impact electron detachment [7], charge exchange [7],

electron- and heavy-impact ionization from H levels [19] and electron- and heavy-impact excitation for H levels [19]. For each reaction, the forward rate coefficient is parametrized by a modified Arrhenius formula

$$k_{f,r}(T) = A_r T^{\eta_r} \exp\left(-\frac{T_{a,r}}{T}\right), \quad r \in \mathcal{R}_{\text{coll}} \quad (2.5)$$

where A_r and η_r are the fitting parameters, and $T_{a,r}$ is the activation temperature. The rate-controlling temperature, T , is reaction-type-dependent, consistent with Park's two-temperature approach. The backward reaction rate coefficients $k_{b,r}$ are computed from the equilibrium constant of the reaction $K_{\text{eq},r}$.

In addition to the collisional reactions, radiative decay H(u) \rightarrow H(l) + $h\nu$, $u \in \{2, \dots, 10\}$, $l < u$ is accounted for, adding 45 reactions. The Einstein coefficients A_{ul} (s⁻¹) are extracted from the NIST Atomic Spectral Database [20]. Radiation trapping is accounted for by reducing the effective decay frequency with local transition-dependent escape factors Λ_{lu} (-) between 0 and 1, according to the Mewe empirical formula [21]. The rate coefficient for radiative decay reactions therefore reads

$$k_{f,r} = \Lambda_{lu} A_{ul}, \quad r \in \mathcal{R}_{\text{rad}}. \quad (2.6)$$

Since absorption is accounted for through the escape factors, no backward rate coefficients are considered for radiative decay reactions: $k_{b,r} = 0$, $r \in \mathcal{R}_{\text{rad}}$. The model just described is implemented in the code BRODERS⁺⁺ [14], which relies on the library MUTATION⁺⁺ [22] for all the thermo-chemical models.

2.2 Radiation model

The flow and chemistry models described in the previous sections provide densities of each species in \mathcal{S} along x after the shock, including the excited states of atomic hydrogen. These outputs, as well as the evolution of the two temperatures, are given as input to NASA's line-by-line radiation code NEQAIR *v15.3* [23]. This one-way coupling approach assumes non-local radiation effects are negligible; only local reabsorption is accounted for by the escape factors in the chemistry model. NEQAIR provides spatially and spectrally resolved intensity maps $I(x, \lambda)$ (W/(cm²μm sr)) for selected spectral regions. Each line is broadened based on the physical flow properties, then convolved with the instrumental response functions in both the spatial and spectral dimensions.

The chosen wavelength regions correspond to those measured experimentally in the EAST facility [3]. In particular, we focus on the Blue (323 to 497 nm) and Red (653 to 659 nm) spectral regions, capturing the H-Balmer lines.

2.3 Sensitivity analysis

A sensitivity analysis was conducted to quantify the influence of chemical kinetic parameters on the simulated radiative quantities of interest (QoI). The analysis focused on the Arrhenius pre-exponential factors A_r .




Figure 1: Evolution of the species densities after the shock, in the conditions of Test 56 shot 7.

Sampling of the input parameter space was performed using Saltelli's extension of the Sobol sequence [24], as implemented in the SALIB Python library [25, 26]. Total Sobol sensitivity indices [27] were then computed to evaluate the effects of the parameters on the model outputs.

In total, 12 QoI were defined to capture the spatial evolution of the Balmer- α , β , γ , and δ lines:

$$\text{QoI}_{i,j} = \int_{\lambda_{0,i}-\Delta\lambda}^{\lambda_{0,i}+\Delta\lambda} \int_{x_{0,j}}^{x_{1,j}} I(x, \lambda) dx d\lambda \quad (2.7)$$

where $\lambda_{0,i}$ (nm) are the wavelengths associated with the transitions of interest, and $x_{0,j}$ and $x_{1,j}$ (cm) are the spatial bounds of integration. For each line, we define three consecutive sets of bounds to capture the onset of radiation. For the Balmer- α line, these bounds are $(x_{0,j}, x_{1,j}) \in \{(1, 2), (2, 3), (3, 4)\}$ cm, while for all other lines we use $(x_{0,j}, x_{1,j}) \in \{(2, 3), (3, 4), (4, 5)\}$.

Due to the computational cost of a full-scale analysis, the sensitivity analysis was restricted to grouped reactions by type rather than to each reaction, enabling identification of the most influential reaction groups within the project timeframe. The reactions were therefore separated into six groups: H_2 dissociation, charge exchange, electron-impact ionization, heavy-impact ionization, electron-impact excitation, and heavy-impact excitation. The analysis of variance was carried out on a scaling factor applied to the A_r coefficients of all reactions within the same group, thereby exploring the space around the reference values. The logarithms of these multiplication factors are the values sampled for the sensitivity analysis, spanning two orders of magnitude in each direction around the reference value for all reaction types. In total, 14 336 samples were used for the sensitivity analysis.

3 Results

3.1 Shock tube radiation predictions

Figures 1 and 2 show typical species densities and temperature evolutions behind the shock, as predicted




Figure 2: Evolution of the two temperatures after the shock, in the conditions of Test 56 shot 7.




Figure 3: Boltzmann plot of the densities of $\text{H}(i)$ at three positions after the shock, in the conditions of Test 56 shot 7.

by the flow and chemistry model of Equation (2.1) to Equation (2.4). These results were obtained for the flow conditions of shot 7 in the Test 56 campaign in the EAST facility [3], with pre-shock flow velocity $u_\infty = 27.8$ km/s, pressure $P_\infty = 0.5$ Torr, and room temperature $T_\infty = 291$ K.

Focusing on Figure 1, we see a region of a few centimeters far from chemical and thermal equilibrium. The dissociation of H_2 is nearly instantaneous. In this simulation, the creation of H^+ ions primarily occurs via charge exchange, $\text{H} + \text{H} \rightleftharpoons \text{H}^+ + \text{H}^-$. We see that the presence of negative H^- ions balances the charge. Most electrons are then produced by electron detachment from these negative ions, until their density is high enough to drive the chemistry to a fast equilibrium. The progressive evolution to equilibrium is also illustrated in Figure 3 by the densities of excited states of atomic hydrogen, which start from a highly out-of-equilibrium configuration to a perfectly aligned Boltzmann distribution at the electron temperature.

Figure 4 (a,b) shows the spatially and spectrally resolved radiative intensity predictions computed by NEQAIR, when fed the results discussed above, for the Blue and Red spectral regions measured experimentally, capturing the H-Balmer series of lines. Figure 4 (c,d) shows the experimental intensity measurements




Figure 4: Radiative predictions and comparison to experiment for Test 56 shot 7. **(a)** Blue region, synthetic. **(b)** Red region, synthetic. **(c)** Blue region, experimental. **(d)** Red region, experimental.

for Test 56, shot 7 [3], in the same spectral regions. We see that they compare well, both qualitatively and quantitatively: the onset of radiation occurs a few centimeters behind the shock and persists; the experimental broadening is also well represented; and the absolute intensities are comparable.

3.2 Preliminary sensitivity analysis

The total Sobol indices representing the relative influence of the different reaction groups were computed for the QoIs defined in Equation (2.7). Figure 5a shows the results for a given window in the x direction (from 2 to 3 cm after the shock), for the considered H-Balmer lines. From these results, we see that varying the electron-impact excitation and dissociation rates has by far the most significant impact on the radiative predictions. Figure 5b shows similar results, but highlights the impact of each reaction type on QoI of the Balmer- α line defined by different x intervals. Further from the shock, the impact of H_2 dissociation rates is lower, while electron-impact excitation becomes increasingly important.

4 Discussion

The results of Section 3 show that our model has the potential to replicate the experimental measurements. However, high uncertainties remain on the reaction rates, which vary widely depending on the source [11, 28], with the model defined in Section 2, very different results can be obtained depending on how the reaction rates are tuned, as illustrated in Figure 6 where three runs of the model were performed in the same flow conditions (Test 56 shot 7), but for different sets of reaction rates, respectively by boosting or slowing some reaction groups. This further justifies the need for careful calibration of these rates and is also a promising

first result, showing that the model is suited for this task as we approach the experimental results.

The results of the sensitivity analysis show how this approach can help identify reactions that affect the model's radiative output. However, such an approach using reaction groups has limitations, as it does not allow discrimination between individual reactions.


5 Conclusions

This work presented the development of a predictive model for radiative heating in H_2/He entry flows relevant to future ice giant missions by combining flow, chemistry, and radiation models, which can provide good agreement with the experiment. This work also lays the groundwork for future model calibration with a variance-based sensitivity analysis, identifying the reaction groups whose uncertainties have the highest impact on the predicted radiative output. This represents a key step toward reducing the current discrepancies in chemical mechanisms observed across the literature and achieving predictive consistency with experiment in radiative heating models for outer planet entries.


Future work will focus on performing a large-scale sensitivity analysis of individual reactions and identifying the coefficients that require the most careful calibration. Bayesian inversion will then be carried out to calibrate the most critical reaction rate coefficients using high-quality EAST experimental data, providing a calibrated kinetic model suitable for mission design applications.

Acknowledgments

BF is an F.R.S.-FNRS Research Fellow. JC is part of the Space4ReLaunch project, supported by the SPW Économie Emploi Recherche of the Walloon Region, under grant agreement no. 2210181. The sponsors



(a) H-Balmer lines integrated between $(x_0, x_1) = (2, 3)$ cm.



(b) Balmer- α line at different (x_0, x_1) intervals.

Figure 5: Total Sobol indices for the considered reaction groups (Dis: dissociation, CX: charge exchange, ELz: electron-impact excitation, HLz: heavy-impact ionization, EEx: electron-impact excitation, HEx: heavy-impact excitation).




Figure 6: Spectrally integrated radiation intensity after the shock ($x_{\text{shock}} = 0$) for three tested sets of reaction rates, compared to the experimental results of Test 56 shot 7.

of the research stay at NASA Ames are VKI, ULiège, ULB, and the F.R.S.-FNRS. We also want to thank Dr. Brett Cruden, Dr. Bruno Dias, Dr. George Chatzigeorgis, Dr. Kaelan Hansson, Dr. Andrea Fagnani, and Dr. Amal Sahai for the insightful discussions throughout the stay.

References

[1] National Academies of Sciences, Engineering, and Medicine, Division on Engineering and Physical Sciences, Space Studies Board, and Committee on the Planetary Science and Astrobiology Decadal Survey. *Origins, Worlds, and Life: A Decadal Strategy for Planetary Science and Astrobiology*

2023-2032. National Academies Press, Washington, D.C., 2023.

[2] A. Simon, F. Nimmo, and R. C. Anderson. Journey to an Ice Giant System: Uranus Orbiter and Probe. Technical report, NASA, 2021.

[3] B. A. Cruden and D. W. Bogdanoff. Shock Radiation Tests for Saturn and Uranus Entry Probes. *Journal of Spacecraft and Rockets*, 54(6):1246–1257, 2017. Publisher: American Institute of Aeronautics and Astronautics.

[4] B. A. Cruden and A. C. Tibère-Inglesse. Impact of Trace CH_4 on Shock Layer Radiation in Outer Planet Entry. In *AIAA SCITECH 2024 Forum*, Orlando, FL, 2024. American Institute of Aeronautics and Astronautics.

[5] C. M. James, D. E. Gildfind, R. G. Morgan, S. W. Lewis, and T. J. McIntyre. Experimentally Simulating Giant Planet Entry in an Expansion Tube. *Journal of Spacecraft and Rockets*, 57(4):656–671, 2020. Publisher: American Institute of Aeronautics and Astronautics.

[6] J. Steer, P. Collen, A. Glenn, C. Hambidge, L. J. Doherty, M. McGilvray, T. Sopek, S. Loehle, and L. Walpot. Commissioning of Upgrades to T6 to Study Giant Planet Entry. *Journal of Spacecraft and Rockets*, pages 1–18, 2024.

[7] G. Colonna, L. D. Pietanza, and A. Laricchiuta. Ionization kinetic model for Hydrogen-Helium atmospheres in hypersonic shock tubes. *International Journal of Heat and Mass Transfer*, 156:119916, 2020.

[8] A. J. Erb, T. K. West, and C. O. Johnston. Investigation of Galileo Probe Entry Heating with Coupled Radiation and Ablation. *Journal of Spacecraft and Rockets*, 57(4):692–706, 2020.

- [9] J. Coelho and M. Lino Da Silva. Aerothermodynamic analysis of Neptune ballistic entry and aerocapture flows. *Advances in Space Research*, 71(8):3408–3432, 2023.
- [10] A. Meini. Development of Aerothermodynamic and Radiative Simulations for Hypersonic Atmospheric Entry into Ice Giants. Master’s thesis, Università di Pisa, von Karman Institute for Fluid Dynamics, Pisa, Italy, 2024.
- [11] A. T. Carroll, G. Blanquart, A. M. Brandis, and B. A. Cruden. State-Specific Kinetic Modeling for Predictions of Radiative Heating in H₂/He Entry Flows. In *AIAA SCITECH 2024 Forum*, Orlando, FL, 2024. American Institute of Aeronautics and Astronautics.
- [12] D. Ninni, F. Bonelli, G. Colonna, A. Laricchiuta, and G. Pascazio. Assessment of hybrid Macroscopic/State-to-State model for numerical simulation of Ice Giant orbit insertion. *International Journal of Heat and Mass Transfer*, 249:127188, 2025.
- [13] T. E. Magin, L. Caillault, A. Bourdon, and C. O. Laux. Nonequilibrium radiative heat flux modeling for the Huygens entry probe. *Journal of Geophysical Research: Planets*, 111(E7), 2006.
- [14] S. Boccelli, F. Bariselli, B. Dias, and T. E. Magin. Lagrangian diffusive reactor for detailed thermochemical computations of plasma flows. *Plasma Sources Science and Technology*, 28(6):065002, 2019.
- [15] J. G. Kim, O. J. Kwon, and C. Park. Master equation study and nonequilibrium chemical reactions for h + h₂ and he + h₂. *Journal of Thermophysics and Heat Transfer*, 23(3):443–453, 2009.
- [16] L. P. Leibowitz and T.-J. Kuo. Ionizational Nonequilibrium Heating During Outer Planetary Entries. *AIAA Journal*, 14(9):1324–1329, 1976.
- [17] M. Furudate, I.-S. Jeung, and S. Matsuyama. *Nonequilibrium Calculation of Flowfield over Galileo Probe*.
- [18] D. Rapp and P. Englander-Golden. Total cross sections for ionization and attachment in gases by electron impact. i. positive ionization. *The Journal of Chemical Physics*, 43(5):1464–1479, 1965.
- [19] H. W. Drawin. Influence of atom-atom collisions on the collisional-radiative ionization and recombination coefficients of hydrogen plasmas. *Zeitschrift für Physik A Hadrons and nuclei*, 225(5):483–493, 1969.
- [20] A. Kramida, Y. Ralchenko, J. Reader, and NIST ASD Team. Nist atomic spectra database (version 5.12), 2024.
- [21] R. Mewe. Relative intensity of helium spectral lines as a function of electron temperature and density. *British Journal of Applied Physics*, 18(1):107–118, 1967.
- [22] J. B. Scoggins, V. Leroy, G. Bellas-Chatzigeorgis, B. Dias, and T. E. Magin. Mutation++: Multi-component Thermodynamic And Transport properties for IONized gases in C++. *SoftwareX*, 12:100575, 2020.
- [23] A. M. Brandis and B. A. Cruden. NEQAIR v15.0 Release Notes.
- [24] A. Saltelli, P. Annoni, I. Azzini, F. Campolongo, M. Ratto, and S. Tarantola. Variance based sensitivity analysis of model output. Design and estimator for the total sensitivity index. *Computer Physics Communications*, 181(2):259–270, 2010.
- [25] J. Herman and W. Usher. SALib: An open-source python library for sensitivity analysis. *The Journal of Open Source Software*, 2(9), 2017.
- [26] T. Iwanaga, W. Usher, and J. Herman. Toward SALib 2.0: Advancing the accessibility and interpretability of global sensitivity analyses. *Socio-Environmental Systems Modelling*, 4:18155, 2022.
- [27] I. M. Sobol. Global sensitivity indices for nonlinear mathematical models and their Monte Carlo estimates. *Mathematics and Computers in Simulation*, 55(1-3):271–280, 2001.
- [28] P. Reynier, G. D’Ammando, and D. Bruno. Review: Modelling chemical kinetics and convective heating in giant planet entries. *Progress in Aerospace Sciences*, 96:1–22, 2018.

Experimental reconstruction

Holistic Supersonic Plasma Ground Testing

Sander Holum^{1,2} Roemer Spreij^{1,3}

¹von Karman Institute for Fluid Dynamics ² Université libre de Bruxelles ³ Katholieke
Universiteit Leuven
December 4, 2025

Abstract

Characterization efforts for the supersonic mode of the VKI Plasmatron Inductively Coupled Plasma wind tunnel have so far been limited by the available numerical tools, with no single tool that models all key phenomena including chemical and thermal non-equilibrium (CNEQ/TNEQ) and the magnetohydrodynamic interaction in the plasma torch. This project establishes an initial simulation pipeline based on a chaining of three CFD codes: the VKI ‘ICP’ code, CFD++ and VKI ‘Stagline’. Simulation results clearly point to the presence of CNEQ and thermal freezing in the nozzle, which fits with experiment. The flow is found to be close to rarefied. Rebuilding of heat flux and stagnation pressure measurements remains a challenge, but it is clear that ground-to-flight extrapolation using the traditional subsonic LHTS method and standard (50mm) intrusive probes is not possible due to the non-equilibrium in front of the probe. Moving to larger probes to promote equilibrium could be a solution, leading to equivalent flight conditions around 70km altitude and 8km/s velocity.

Keywords: Ground testing, plasma wind tunnel, ground to flight extrapolation, flight relevant testing, supersonic jet, plasma flow, non-equilibrium flow

1 Introduction

The use of ground testing is crucial to the study of atmospheric entry. Plasma wind tunnels, which include arc-jet and Inductively Coupled Plasma (ICP) facilities reproduce the high-temperature hypersonic shock layer encountered in flight. The von Karman Institute for Fluid Dynamics (VKI) uses the ‘Plasmatron’ ICP facility for this purpose. The test methods used in this facility have in the past primarily revolved around subsonic stagnation point testing. In this context, the Local Heat Transfer Simulation (LHTS) methodology is employed to relate ground test to flight conditions and vice-versa (see Section 2). LHTS is based on the idea of reproducing exactly the stagnation point boundary layer. This must be contrasted with methods that rely on direct duplication of heat flux and stagnation pressure, without necessarily reproducing the exact stagnation point environment. The latter are often employed at arc-jet facilities, including those of NASA Ames Research Center [1, 2]. However, supersonic nozzles have come into operation at the VKI Plasmatron that both extend the operating range of the facility and enable flat plate (off-stagnation point) testing with higher shear [3]. Two main questions arise: what is the

exact nature of the flow coming out of these nozzles, and is it still possible to relate the ground test conditions to flight in the case of supersonic stagnation point testing? This work attempts to answer these questions.

2 Background

The LHTS framework, as outlined by Kolesnikov and applied at VKI [4, 5, 6] states that, assuming the flow is in Local Thermodynamic Equilibrium (LTE) at the boundary layer edge, three parameters need to be matched between flight and ground to reproduce the stagnation point boundary layer. They are the stagnation pressure p_s , the stagnation enthalpy h_s , and the (inviscid) stagnation point velocity gradient $\frac{du_e}{dx}|_s = \beta_s$. The assumption of LTE, meaning thermal and chemical equilibrium conditions at the boundary layer edge, is usually fulfilled for subsonic VKI Plasmatron operation (except for low pressures around 15 mbar) [7, 8]. For the supersonic operation, two complications appear. Firstly, chemical and thermal non-equilibrium (CNEQ, TNEQ) effects must be considered in the nozzle and after, which is supported by recent spectroscopic measurements that suggest high levels of CNEQ and TNEQ

after the nozzle exit and the probe location [9].

Secondly, the flow may become rarefied during the expansion through the nozzle. The key parameter to monitor is the Knudsen number, the ratio of the mean free path λ and a relevant length scale L . In a nozzle flow, there is no clear way to define a fixed L , and λ changes significantly during the expansion. Candler and Bird [10] define the Gradient Length Local (GLL) Knudsen number, with λ calculated locally and L defined using the gradient of a strongly varying flow property, such as density:

$$Kn_{GLL} = \frac{\lambda}{L} = \frac{\lambda}{\rho} \left| \frac{\partial \rho}{\partial x} \right| \quad (2.1)$$

With objects in the flow, like probes, one can use the shock stand-off distance as a length scale [11]. To capture the CNEQ/TNEQ effects, this work uses a finite-rate 11 species air model (N, O, NO, N₂, O₂, N⁺, O⁺, NO⁺, N₂⁺, O₂⁺, e⁻) and Park's two-temperature (2T) model [12, 13, 14]. The dissociation reactions are modeled using Arrhenius laws with forward rates K :

$$K = CT_a^n e^{-\frac{H_0}{kT_a}} \quad (2.2)$$

where the controlling temperature $T_a = \sqrt{T_{ev}T_{rt}}$ is the geometric average of the electronic-vibrational temperature T_{ev} and the rotational-translational temperature T_{rt} .

3 Simulation methodology

The aim is to simulate the flow in the supersonic VKI Plasmatron, including plasma torch, nozzle and chamber, with and without the presence of an intrusive stagnation point probe. As of now, a code able to capture all the complex phenomena present in the flow (compressibility, high-temperature CNEQ/TNEQ, magneto-hydrodynamics in the ICP torch) is not available to us. As such, the domain has been divided into three regions with a different code used for each to be able to reproduce the relevant phenomena in the flow regions where they occur. Coupling between the regions and codes is done manually. These are described in the following sections.

3.1 Simulation of plasma torch with VKI 'ICP' Code

The in-house VKI 'ICP' code based on the COOLFluid platform is used to solve the plasma torch and the electromagnetic coupling between the flow and the induction coil [15]. However, the ICP code is limited in that it can only model LTE chemistry, making it unsuitable to handle the CNEQ/TNEQ flow in the nozzle and jet. To reproduce a test condition, the mass flow rate in the torch is imposed and the numerical electrical power (different from the measured electrical power listed in Table 1 because of the unknown coupling efficiency between coil and plasma [16]) is adjusted until the numerical reservoir pressure in the torch P_{res} matches the

measured value. Because the subsonic flow in the nozzle contraction affects the flow field in the torch itself, the simulation domain goes up to and slightly beyond the nozzle throat. The outputs of the ICP simulation are the flow profiles at the start of the nozzle contraction, which are passed to the CFD++ simulation. A similar approach using a parabolic inlet profile at the throat of the nozzle, defined using macro properties from the reservoir, has been used by NASA Ames for the arc-jets [11]. Both the ICP and CFD++ simulations are 2D-axisymmetric around the jet centerline to reduce computational cost.

3.2 Simulation of nozzle and chamber with CFD++

The code chosen to solve the flow through the nozzle and in the test chamber is CFD++ from Metacomp Technologies Inc., due to prior knowledge of the code and its relative ease of use and robustness. The CFD++ simulation domain with boundary conditions (BC's) is shown in fig 1. The inlet profiles of pressure, temperature and chemical composition are taken from ICP simulations, as described in sec 3. The nozzle is modeled as an isothermal wall (350 K). The top of the chamber is an inviscid wall and is placed far enough away not to affect the jet development. The outlet is a supersonic outlet with imposed back pressure equal to the experimental static chamber pressure. The left chamber side is a blowing BC, where constant static pressure, x-velocity, temperature, and chemical composition are imposed. The blowing stabilizes the simulation. Park's 2T model is used with an 11-species air model, described in sec 2. The tabulated chemical kinetics are extracted from the Mutation++ thermodynamic library [17].

3.3 Simulation of intrusive probe stagnation line with VKI 'Stagline' code

The simulations are meant to replicate the experimental conditions in the Plasmatron chamber, that feature intrusive heat flux and stagnation pressure measurements. Experimental cold wall heat flux measurements were taken with a 50 mm diameter hemispherical probe with a copper water-cooled calorimeter, experimental stagnation pressure is measured with a pitot probe of the same geometry. The experimental values are listed in Table 1. Considering the available computational resources, the intrusive probe simulations were not done in CFD++, but rather with a fast pseudo-1D code 'Stagline' based on the Dimensionally Reduced Navier-Stokes Equations (DRNSE) [18, 19]. These are valid on the stagnation line towards a sphere. The Stagline simulations, with a domain extending 10 cm from the probe wall, take as inlet BC the CFD++ output on the centerline at the probe location: T_{rt} and T_{ev} , the species composition, and the flow velocity. Stagline is also coupled with the Mutation++ library and uses the same 2T model. However, an issue with the handling of charged species in Mutation++, believed to be due to a bug, prevented the use of the 11-species




Figure 1: Example CFD++ solution Mach number contours including BC conditions. Lower line is a symmetry BC. The dotted line figure is an approximation of the experimental probe location in the jet.

air model. Instead, the composition was converted to a 5-species model (O_2 , N_2 , NO , O , N) by adding the densities of the ions to those of the corresponding neutrals. The wall boundary condition is set to replicate the cooled calorimeter surface: 350 K imposed temperature with surface recombination treated using the ‘gamma’ model defining surface recombination probabilities $\gamma_N = \gamma_O = \gamma$ equal to the ratio of the flux of impinging atomic species divided by the flux of recombined species [20]. The value of γ is both hard to determine and has a strong effect on the computed heat flux [16]. This work uses the pressure-dependent γ model defined by Viladegut & Chazot [21], with the gamma values interpolated at the measured p_s .

4 Results and discussion

4.1 Plasma jet

Figure 4 shows the value of Kn_{GLL} along the centerline of the cases considered in Table 1. The red line shows the limiting value of Kn_{GLL} suggested by Boyd et. al. [10] above which the continuum flow assumption breaks down, and the shaded areas roughly indicate the different flow regimes. There is only one region where the Knudsen number barely exceeds the limit: in the compression region of the flow, after the probe location. The use of CFD methods to characterize the flow is therefore deemed appropriate in the region of interest: before the probe and in the nozzle.

To look at the effect of CNEQ and TNEQ, T_{rt} , T_{ev} , and the mass fractions of O and N are plotted along the jet centerline in Figure 2. The temperatures show a large deviation from thermal equilibrium developing through the nozzle (>4000 K difference), as expected. The compression shock is visible as an increase in T_{rt} after the probe location. T_{ev} stays frozen from the exit of the nozzle, to the end of the domain. The squares in the plot are spectroscopic data measuring a supposed ‘LTE’ temperature from the O777 spectral line, see Holum et. al [9]. This spectral line from atomic oxygen is sensitive to the electronic temperature, and the CFD simulations show this staying high along the jet. Hence, the 2T model is crucial to match the experimental result. The CNEQ is also clear as atomic oxygen stays fully dissociated, and atomic nitrogen stays almost above 50% mass fraction to the end of the jet even with low T_{rt} , and the geometric averaged temperature driving the chemical kinetics in Park’s 2T model is below 4000 K, where an LTE calculation (Mutation++) gives less than 5% atomic oxygen.

4.2 Intrusive probe simulation

The outputs of the Stagline reproduction of the ground tests are listed on the right of Table 1. The left plots of Figure 3 show the species composition and temperature profiles for case 3 of Table 1. It is clear that the assumption of LTE, necessary for the application of the LHTS method, does not hold. The species composition at the boundary layer edge is far from equilibrium, and there is no equilibration between T_{rt} and T_{ev} . This is due to the small size of the heat flux probe, which is reflected in the large β_s values. Comparing the intrusive measurements with the Stagline reproduction in Table 1, the heat flux results fluctuate around the measured values. The inconsistency could be attributed to the unknown value of γ . The stagnation pressure is consistently overestimated by Stagline. Measurement uncertainty for \dot{q} and p_s are typically taken as 10% [22]. Modelling assumptions, like non-catalytic nozzle walls, may account for higher stagnation pressure as less energy is removed from the flow in the nozzle. Simulations with other catalytic assumptions, like a super-catalytic condition for O and N recombination as in Stewart et. al., should be attempted [11]. Ground to flight extrapolation with LHTS is performed by deriving the flight velocity V_∞ from h_s ; altitude from the V_∞ and p_s ; and the equivalent flight radius R_{eq} from Modified Newtonian theory, or another approximate pressure distribution method. For a more in-depth overview see [5]. Due to the lack of LTE, applying this method directly to the cases of Table 1 results in flight conditions that poorly reproduce the boundary layer and heat flux. LTE can be promoted in the ground test with a larger probe. The right plots of Figure 3 shows the results of the ground to flight extrapolation performed by simulating a fictional probe with $R_{eq} = 1m$ in the Plasmatron. The flow is clearly closer to LTE and the resulting boundary layer is well duplicated. This shows also in the close match of the simulated heat flux (Table 2). The velocities lie roughly in the 7.5-8.5 km/s range, and the altitudes are in a small range between 69-73 km, reflecting the near-rarefied state of the flow (Section 4.1). The flight R_{eq} is close to that of the ground test (typically not the case for subsonic Plasmatron operation). Increasing the R_{eq} of a Plasmatron probe while limiting the base diameter R_b can be done by using flat-faced probes, with a theoretical limit of $R_b \approx 0.3R_{eq}$ [23]. Additionally, the increased shock stand-off distance requires placing the probe farther back in the chamber and blockage issues may arise. Further study is needed to see if a realistic configuration can be found that allows for LHTS ground-to-flight extrapolation in the supersonic Plasmatron.

5 Conclusions

During the MSV, a first simulation pipeline has been established to simulate the whole supersonic Plasmatron operation, including the ICP torch, supersonic nozzle expansion, and shock layer environment in front




Figure 2: T_{ev} (-), T_{rt} (:), ρ_O (o) and ρ_N (x) along centerline for CFD++ simulations. Vertical lines, nozzle exit (left), and the probe location when inserted (right). Y-axis, right: Temperature [K], left: Mass fraction [-]

Operating conditions				Intrusive measurements		Stagline results (% error w.r.t. ground test)				
ID	\dot{m} g/s	P_{el} kW	$P_{chamber}$ Pa	P_{torch} mbar	\dot{q} MW/m ²	p_s mbar	\dot{q}_{SL} MW/m ²	$p_{s,SL}$ mbar	$h_{s,SL}$ MJ/kg	$\beta_{s,SL}$ 1/ms
1	6	300	236	125	3.92	28.9	3.02 (-23.0%)	33.6 (16.3%)	29.4	101.4
2	6	400	238	138	4.20	30.8	3.68 (-12.4%)	37.9 (23.1%)	34.1	109.2
3	10	300	344	165	3.31	39.1	3.77 (-13.9%)	49.8 (27.4%)	31.0	103.6
4	10	400	346	179	4.96	44.7	4.18 (-15.7%)	54.7 (22.4%)	34.6	109.2
5	10	500	346	196	5.72	48.7	4.53 (-20.8%)	58.8 (20.7%)	37.2	114.4
6	12	400	416	201	3.12	47.7	4.63 (48.4%)	63.6 (33.3%)	35.6	110.7
7	12	500	408	215	5.33	54.0	4.96 (-6.9%)	67.2 (24.4%)	37.9	115.4

Table 1: Table of the experimental Plasmatron conditions simulated using ICP, CFD++ and Stagline (SL).




Figure 3: Left: Stagnation line profiles of case 3, with a 50 mm diameter hemispherical probe. Dashed lines in mole fractions plot indicate equilibrium species composition. Right: Stagnation line profiles of same ground test conditions with $Re_{eq} = 1$ m, and equivalent flight conditions (Table 2).




Figure 4: Average Kn_{GLL} along flow centerline for all CFD++ simulations. The shaded region indicates the max and min value across all simulations. Vertical line shows probe location (0.377 m).

Ground test ($Re_{eq} = 1$ m)		Flight conditions (% error w.r.t. ground test)			
ID	$\dot{q}_{SL,ground}$ MW/m ²	$\dot{q}_{SL,flight}$ MW/m ²	Velocity km/s	Altitude km	$Re_{eq,flight}$ mm
1	0.537	0.572 (6.5%)	7.67	72.6	1025
2	0.653	0.698 (6.9%)	8.25	72.7	1024
3	0.669	0.701 (4.8%)	7.88	70.2	1028
4	0.735	0.768 (4.5%)	8.32	70.3	1027
5	0.778	0.808 (3.9%)	8.63	70.3	1025
6	0.802	0.831 (3.6%)	8.44	69.4	1028
7	0.836	0.864 (3.3%)	8.71	69.4	1026

Table 2: Equivalent flight conditions for the test conditions of Table 1 with a simulated $Re_{eq} = 1$ m probe. Heat flux computed with a hot wall ($T_w = 2000$ K) boundary condition.

of the measurement probes (pressure and heat flux). CNEQ is shown to be present throughout the flow field, and persists after the bow shock, affecting boundary layer chemistry and heat flux. This invalidates the LHTS ground-to-flight extrapolation methodology. To promote LTE at the boundary layer edge, a bigger probe has been simulated, showing chemistry closer to equilibrium and leading to equivalent flight conditions around 70 km altitude and at LEO entry velocities,

with issues of practical implementation of such a probe remaining. The simulation pipeline shows promise in replicating the physical phenomena occurring during a supersonic Plasmatron test. Still, the difference between experiments and simulations is quite high for p_s , \dot{q} , and freestream temperature. Some are due to experimental error, especially freestream temperature. However, there are many improvements to be made in the supersonic plasma jet modeling. In the reservoir (ICP simulations), the LTE assumption is known to be restrictive in the plasma generation [8]. The effect of catalycity both on the probe and in the nozzle should be explored. Moving to a single code to model the entire domain could improve accuracy.

Acknowledgments

Both authors wish to thank the organizers and hosts of the MSV at NASA Ames: Bruno Dias & Jeremie Meurisse, as well as all the organizers at VKI, École Polytechnique and Université de Bordeaux. We thank in particular our NASA mentors for the fruitful discussion during the stay, including Georgios Bellas-Chatzigeorgis and Andrea Fagnani.

S. Holum would like to thank DEFRA Belspo for funding (Contract no. 4000142160/23/D/BL).

R. Spreij is supported by Fonds Wetenschappelijk Onderzoek Vlaanderen (FWO) fellowship no. 1SHAL24N, and gratefully acknowledges additional support for this research stay from the FWO (grant no. K230325N).

References

- [1] Jay Grinstead, George Raiche, and Tahir Gökçen. Test Engineering for Arc Jet Testing of Thermal Protection Systems: Design, Analysis, and Validation Methodologies. In *25th AIAA Aerodynamic Measurement Technology and Ground Testing Conference*, San Francisco, California, June 2006. American Institute of Aeronautics and Astronautics.
- [2] Alireza Mazaheri, Walter E. Bruce, Nathaniel J. Mesick, and Kenneth Sutton. Methodology for Flight-Relevant Arc-Jet Testing of Flexible Thermal Protection Systems | Journal of Spacecraft and Rockets. *Journal of Spacecraft and Rockets*, 51(3), May 2014.
- [3] Joseph Elrassi, Sander Holum, Loïc Sombaert, Alan Viladegut, Louis Walpot, Olivier Chazot, and Bernd Helber. Upgraded vki plasmatron capabilities with supersonic nozzles for extended material characterization methods. *Sampe Journal*, 60, 03 2024.
- [4] A. F. Kolesnikov. Conditions of simulation of stagnation point heat transfer from a high-enthalpy flow. *Fluid Dynamics*, 28(1):131–137, January 1993.
- [5] P. F. Barbante and O. Chazot. Flight Extrapolation of Plasma Wind Tunnel Stagnation Region Flowfield. *Journal of Thermophysics and Heat Transfer*, 20(3):493–499, July 2006.
- [6] Alessandro Turchi, Juan J. Matesanz Saiz, Thierry E. Magin, and Olivier Chazot. Duplication of hypersonic stagnation-region aerothermochemistry and gas-surface interaction in high-enthalpy ground testing. *Experiments in Fluids*, 62(11):238, October 2021.
- [7] Yann Cressault, Ph Teulet, X Baumann, and A Gleizes. Non-equilibrium phenomena in thermal plasmas. *Plasma Research Express*, 2(4):043001, 2020.
- [8] Marco Panesi, P Rini, and O Chazot. Numerical simulation of non-equilibrium plasma flows. *Project Report*, 23:5, 2005.
- [9] Sander Holum, Ian Ballou, Bernd Helber, Alessandro Parente, and Olivier Chazot. Flow characterization of supersonic plasma flow using optical emission spectroscopy. ESA FAR 2025, Arcachon, France, May 2025.
- [10] Iain D. Boyd, Gang Chen, and Graham V. Candler. Predicting failure of the continuum fluid equations in transitional hypersonic flows. *Physics of Fluids*, 7(1):210–219, January 1995.
- [11] David Stewart, Tahir Gokcen, and Yih-Kanq Chen. Characterization of Hypersonic Flows in the AHF and IHF of NASA Ames Arc-Jet Facilities. In *41st AIAA Thermophysics Conference*.
- [12] Chul Park. Assessment of a two-temperature kinetic model for dissociating and weakly ionizing nitrogen. *Journal of Thermophysics and Heat Transfer*, 2(1):8–16, January 1988.
- [13] Chul Park. Assessment of two-temperature kinetic model for ionizing air. *Journal of Thermophysics and Heat Transfer*, 3(3):233–244, July 1989.
- [14] Chul Park. The Limits of Two-Temperature Kinetic Model in Air. In *48th AIAA Aerospace Sciences Meeting Including the New Horizons Forum and Aerospace Exposition*, Orlando, Florida, January 2010. American Institute of Aeronautics and Astronautics.
- [15] Andrea Lani, N. Villedieu, K. Bensassi, Lilla Koloszar, Martin Vymazal, Mehmet Sarp Yalim, and Marco Panesi. Coolfluid: An open computational platform for multi-physics simulation and research. *21st AIAA Computational Fluid Dynamics Conference*, 01 2013.
- [16] Andrea Fagnani. *Development of Measurement Techniques and Study of the Aerothermal Response of Space Debris Materials to Atmospheric Entry Plasmas*. PhD thesis, Vrije Universiteit Brussel, 2023.
- [17] James B. Scoggins, Vincent Leroy, Georgios Bellas-Chatzigeorgis, Bruno Dias, and Thierry E. Magin. Mutation++: Multicomponent thermodynamic and transport properties for ionized gases in c++. *SoftwareX*, 12, 2020.
- [18] Arno Klomfass and Siegfried Müller. Calculation of stagnation streamline quantities in hypersonic blunt body flows. *Shock Waves*, 7(1):13–23, February 1997.
- [19] A. Munafò and T. E. Magin. Modeling of stagnation-line nonequilibrium flows by means of quantum based collisional models. *Physics of Fluids*, 26(9):097102, September 2014.
- [20] R. Goulard. On Catalytic Recombination Rates in Hypersonic Stagnation Heat Transfer | Journal of Jet Propulsion. *Journal of Jet Propulsion*, 28(11), November 1958.
- [21] A. Viladegut and O. Chazot. Empirical Modeling of Copper Catalysis for Enthalpy Determination in Plasma Facilities. *Journal of Thermophysics and Heat Transfer*, 34(1):26–36, 2020.
- [22] Enrico Anfuso, Andrea Fagnani, Aurélie Belle-mans, Dieter J. De Baere, and Olivier Chazot. Considerations on a Free-Stream Characterization Procedure for Plasma Jet of the VKI Plasmatron Facility. In *AIAA SCITECH 2025 Forum*, AIAA SciTech Forum, Orlando, Florida, USA, January 2025. American Institute of Aeronautics and Astronautics.

- [23] Ernest V. Zoby and Edward M. Sullivan. Effects of Corner Radius on Stagnation-Point Velocity Gradients on Blunt Axisymmetric Bodies. *Journal of Spacecraft and Rockets*, 3(10):1567–1567, 1966.

In-Situ 3D Reconstruction in High-Enthalpy Test Facilities

Modeling Summer Visit 2025

Nicolas Derider^{1,2} Nicolas Tomić^{1,2} Alexandre Quintart^{1,3}

¹von Karman Institute for Fluid Dynamics (VKI), Sint-Genesius-Rode, Belgium

²Université libre de Bruxelles (ULB), Brussels, Belgium

³Flying Squirrel, Bourg-St-Pierre, 1964, Switzerland

Abstract

Ground testing of materials used for atmospheric entry requires reliable methods to monitor the evolution of sample geometry during exposure to high enthalpy flows. Existing approaches at NASA Ames Research Center and von Karman Institute for Fluid Dynamics mainly rely on side view imaging for real time two dimensional tracking or on ex-situ three dimensional scanning before and after testing. These techniques do not capture transient deformation or localized recession that may develop during a test. This work presents a proof of concept for an in-situ three dimensional reconstruction system intended for use in high-enthalpy plasma facilities. Several open source and commercial structure from motion tools were evaluated and a complete calibration and reconstruction pipeline was implemented on the NASA ARCTRON facility. The system was tested first on a sample inside the chamber and then during an operating run. The study shows that real time three dimensional monitoring is feasible when cameras are carefully arranged and pre calibrated and when additional surface features are introduced to ensure robust image matching. The results provide a first demonstration of three dimensional reconstruction under relevant conditions and outline the steps needed for future integration of this capability into routine facility operations.

Keywords: in situ 3D reconstruction, structure from motion, photogrammetry, high-enthalpy testing, thermal protection systems, plasma wind tunnels

1 Introduction

Ground testing of Thermal Protection Systems (TPS) and other materials is essential for validating spacecraft performance during atmospheric entry. At NASA Ames Research Center (ARC), the Arc Jet Complex hosts several facilities capable of reproducing these extreme environments. As mission requirements evolve toward higher fidelity and faster turnaround, the need for *in-situ* 3D reconstruction arises, enabling real-time geometric monitoring of samples under high-enthalpy testing conditions.

Currently, surface evolution during these tests is typically characterized using either (1) side-view imaging for real-time 2D profile tracking, or (2) ex-situ 3D scanning of the samples before and after testing, often complemented by simple recession-length measurements. While effective for global recession assessment, these methods lack the temporal and spatial resolution needed to capture transient deformation and localized

ablation phenomena occurring during exposure. An automated *in-situ* 3D reconstruction pipeline would therefore provide valuable real-time insight into material behavior under high-enthalpy conditions.

During the NASA Ames summer visit, a proof-of-concept *in-situ* 3D reconstruction system was developed for ground testing in high-enthalpy facilities such as the ArcJet and Plasmatron. While similar work has been demonstrated at the DLR Stuttgart facility, such capabilities are not yet implemented at NASA Ames (ArcJet, ARCTRON) or at the VKI Plasmatron, motivating this study.

The methodology followed three main steps: (1) performing 3D reconstruction outside the chamber to benchmark available techniques, including open-source libraries (*OpenCV*, *COLMAP*, *Meshroom* [1]) and commercial software (*Agisoft Metashape*, *RealityScan*); (2) designing and installing the imaging setup on the ARCTRON facility in non-test conditions to validate the pipeline; and (3) deploying the system during an actual

test, reconstructing sample geometry in real time with plasma exposure.

By addressing key challenges in imaging, calibration, and reconstruction under extreme environments, this work extends earlier photogrammetric recession-measurement approaches such as those demonstrated in [2], and shows the feasibility of real-time 3D monitoring of TPS samples during arc-jet testing. The following sections describe the calibration and reconstruction methodology, preliminary results, and perspectives for integrating this capability into routine facility operations.

2 State of the Art

2.1 Previous work

Efforts to characterize recession and ablation of TPS materials during ground testing have historically relied on two complementary approaches: *ex-situ* 3D scanning before and after exposure, and *in-situ* video-based imaging during the test. Both methods aim to quantify surface evolution under high-enthalpy conditions, but they differ in resolution, temporal coverage, and feasibility across facilities.

At NASA Ames, the STARscan system was developed as a multi-camera photogrammetry rig to reconstruct high-resolution 3D meshes of samples before and after ArcJet testing [3, 4]. It uses turntable-based acquisition and structure-from-motion to achieve sub-millimeter precision, and is routinely used to quantify recession and deformation. However, its workflow is inherently *ex-situ* and does not provide time-resolved evolution.

In parallel, the arcjetCV tool was designed to analyze the evolution of the shock location and sample edge position from video recordings of ArcJet tests, enabling time-resolved tracking of sample boundaries and in-plane recession [5]. It incorporates lens-distortion correction and edge-following algorithms to quantify geometric evolution directly during the test. While arcjetCV filled a critical gap in enabling *in-situ* analysis, it remains restricted to 2D edge tracking and does not support full 3D shape reconstruction.

Outside of NASA, multiple groups have explored real-time or near-real-time photogrammetric reconstruction in harsh aerothermal environments. At the DLR Institute of Space Systems in Stuttgart, Löhle and collaborators demonstrated stereo and photogrammetric tracking in plasma wind tunnels [6], showing that 3D reconstruction is feasible under luminous plasma exposure using specialized optical configurations tailored to the PWK facilities.

Similarly, Yee et al. developed a dual-camera photogrammetric setup capable of reconstructing 3D surface morphology during oxyacetylene OTB testing [7]. Their system produced time-resolved 3D meshes at a cadence of one second using Agisoft Metashape, enabling the extraction of recession rates and demonstrating sensitivity to camera settings and illumination. Although their test environment differs from arc-jet facilities,

their results further confirmed the potential for in-test 3D reconstruction of ablative materials.

To our knowledge, such *in-situ* 3D reconstruction capabilities have not yet been implemented at the NASA Ames ArcJet Complex or at the VKI Plasmatron. This motivated the present effort to integrate SfM pipelines with in-test acquisition setups adapted to these facilities, bridging the temporal resolution of arcjetCV with the geometric fidelity of STARscan.

2.2 Structure-from-Motion

Structure from Motion (SfM) is a photogrammetric computer vision technique that simultaneously determines both the 3D coordinates of scene points and the positions and orientations of the cameras that captured the images. It generally works in the following key steps, [8, 9]:

- **Image acquisition:** Capture multiple images with good overlap and varying viewpoints.
- **Feature detection and matching:** Detect and match key features across images using algorithms such as SIFT, SURF, or ORB.
- **Camera internal parameters:** Camera intrinsics \mathbf{K} are either known from prior calibration or estimated via monocular calibration. The intrinsic matrix encodes the camera’s internal geometry, including focal lengths, principal point coordinates, and possible skew. They are used to convert the fundamental matrix \mathbf{F} into the essential matrix:

$$\mathbf{E} = \mathbf{K}^T \mathbf{F} \mathbf{K},$$

which is then used for camera pose estimation.

- **Camera poses estimation:** Determine the extrinsic parameters using the fundamental and essential matrices \mathbf{F} and \mathbf{E} . The extrinsic parameters define the camera’s position and orientation in space relative to the scene.
- **Scene reconstruction:** Compute the 3D coordinates of the scene points by triangulation using the known camera poses and intrinsics.
- **Bundle Adjustment:** Refine both intrinsic and extrinsic parameters by minimizing the total reprojection error over all cameras and 3D points:

$$\arg \min_{\{\mathbf{R}_i, \mathbf{t}_i, \mathbf{K}_i, \mathbf{X}_j\}} \sum_{i,j} \|\mathbf{x}_{ij} - \pi(\mathbf{K}_i, \mathbf{R}_i, \mathbf{t}_i, \mathbf{X}_j)\|^2,$$

where \mathbf{x}_{ij} is the detected 2D location of 3D point \mathbf{X}_j in image i , and $\pi(\cdot)$ projects 3D points into the image. The optimization adjusts the 3D points and camera parameters to minimize the distance between observed points and their reprojected positions.

2.3 Software

Multiple open-source software and libraries were first considered namely: *OpenCV*, *Meshroom*, *COLMAP*. But it was then decided to use commercial softwares their easiness of use when calibrating the system and robustness to the quality of the image data set. The commercial software explored are *Agisoft Metashape* and *RealityScan*

3 Methodology

3.1 Setup

The imaging setup consisted of multiple industrial cameras equipped with fixed focal-length lenses selected to maximize field of view overlap while maintaining sufficient spatial resolution on the sample surface. Custom camera mounts were designed and 3D printed to rigidly adapt the cameras to the existing optical ports of the ARCTRON facility. The cameras were positioned to maximize angular separation, improving triangulation, while preserving sufficient overlap between neighboring views to facilitate feature matching. Since most reconstructions were performed outside plasma conditions, artificial lighting was needed to ensure uniform illumination and sufficient texture for feature matching.




Figure 1: ARCTRON facility CAD.

3.2 Calibration

The calibration procedure is performed prior to the test using an image dataset of a calibration target with abundant features, acquired from a set of cameras that typically exceeds the number used for the test reconstruction. This approach improves calibration accuracy by providing a larger and more consistent dataset than what can be obtained during the test itself. In contrast, performing calibration in real time during experiments generally yields less reliable results, as the imaging conditions (e.g., the presence of plasma or reduced feature visibility) are less favorable.

After performing the calibration in *Agisoft*, the software exports a text file containing the calibration data for several camera rigs, this happens because each set of images obtained with multiple cameras is treated as a separate rig. A Python script was developed to

parse the exported file, extract the calibration information for each rig, and save it as an individual text file. In contrast, the calibration workflow in *RealityScan* is more straightforward: an *.xmp* file is generated for each image and named identically to the image file. To apply these calibration files, simply move them into the test-image folder and rename them so that they match the corresponding test images, *RealityScan* will then automatically apply them when importing the image set for reconstruction. In both software packages, the rig yielding the smallest reprojection error is selected for the final reconstruction. From this rig, the calibration parameters of a chosen subset of cameras are extracted and applied to the reconstruction process.

3.3 Reconstruction

Once the cameras are calibrated, the reconstruction process begins by importing the corresponding calibration files into the selected software (*Agisoft* or *RealityScan*). The intrinsic parameters, distortion coefficients and relative camera poses are kept fixed, ensuring that the reconstruction relies solely on image correspondences rather than re-estimating geometry during alignment.

The reconstruction pipeline follows the regular SfM workflow. Both *Agisoft* and *RealityScan* provide quality-assessment tools such as reprojection error and confidence maps to evaluate reconstruction quality.

4 Results & Conclusions

4.1 In-situ calibration & reconstruction

The two main objectives of the group for the month were, first, to perform an in-chamber reconstruction in the ARCTRON facility, which required installing the setup and applying the reconstruction pipeline described in the previous paragraphs to objects placed inside the chamber, and second, to carry out a reconstruction during an actual test. In the latter case, the setup is calibrated before the test, and during the test, sets of images from all cameras are captured simultaneously. A reconstruction is then performed for each image set, providing a three-dimensional representation of the sample at each acquisition time.

In theory, reconstructions during a test present two main challenges. The first is the need for all cameras to capture and store images simultaneously, since the sample evolves during the test. The second arises from the presence of plasma between the cameras and the sample, which must somehow be filtered out. A reconstruction was successfully performed during a test, but since the facility was only operating at 1 kW and the magnetic coil, intended to concentrate the plasma and achieve higher local heating rates, had not yet been installed, the plasma was barely visible on the images. As a result, it did not interfere with the reconstruction, and no filtering solutions could be tested in-situ. Furthermore, synchronization of the cameras was unnecessary because the sample did not evolve over time. Consequently, the test reconstruction did not provide

new insights, and the remainder of this section focuses on comparing two in-chamber (not during a test) reconstructions performed with *Agisoft Metashape* and *RealityScan*.

It should also be noted that several samples were 3D printed with shapes similar to those used in real ablation experiments to provide realistic test objects. The only sample for which results are presented in the following paragraphs is a wedge with a hole on its upper surface, shown in Figure 2. This geometry is similar to samples used in NASA experiments, where it is of particular interest because it allows the evolution of the hole to be monitored during testing. The results shown in this section are reconstructions of the top surface of the wedge including the hole, which is flat.




Figure 2: 3D model of the sample.

Setup installation procedure


There are several key aspects to consider when positioning cameras in a facility for 3D reconstruction:

- The overlap between camera fields of view must be high enough for the reconstruction software to match a sufficient number of common features. This is not necessarily easy to achieve depending on the optical accesses of a given facility.
- Lighting should be uniform across the object. This is particularly important for in-chamber reconstructions not performed during a test, as during real tests the plasma around the object will act as the main light source. Additionally, all cameras should ideally share identical settings, especially exposure. Adjusting exposure individually to equalize brightness between cameras can cause inconsistencies when more than two or three cameras are used. It is therefore simpler to keep the same exposure on all cameras and adjust the lighting instead.
- The entire object should be in focus, which requires lenses with a large depth of field and adequate zoom so that the object occupies most of the frame. Since achieving a large depth of field involves closing the iris, strong and uniform lighting becomes even more critical.

Calibration & Reconstruction

As shown in Figure 3, a sheet of paper with printed features was attached to the surface of the model with tape for both calibration and reconstruction, since the SfM software could not detect enough matching points on the raw 3D-printed material due to its reflective surface. For calibration, a total of 12 images of the


sample were used, corresponding to 2 sets of images taken with the 6 cameras installed in the facility. For each of the 2 image sets the sample was placed in a different position. For the reconstruction, the goal was to use a single set of 4 images corresponding to the subset of cameras providing the best overlap, but initially no reconstruction could be obtained despite the printed features. To address this, a laser projecting an array of green points was introduced, as shown in Figure 3b.



(a) Features added on the sample during calibration. (b) Features added on the sample during reconstruction.


Figure 3: Features added on the sample for calibration and for reconstruction.

The reconstructions obtained with 4 images with *Agisoft Metashape* and *RealityScan* are shown in Figure 4 and Figure 5.



(a) Reconstruction with *Agisoft Metashape* with textures. (b) Reconstruction with *Agisoft Metashape* without textures.

Figure 4: Reconstructions obtained with *Agisoft Metashape*.



(a) Reconstruction with *RealityScan* with textures. (b) Reconstruction with *RealityScan* without textures.

Figure 5: Reconstructions obtained with *RealityScan*.

Accuracy of the results

Due to time constraints, it was not possible to properly assess the measurement accuracy of the setup. However, both SfM software tools provide confidence maps for the reconstructions, shown in Figure 6, which indicate if the number of features successfully matched at each location on the reconstructed surface was sufficient or not.

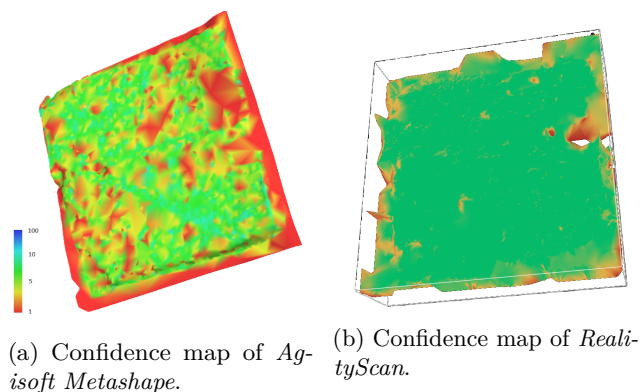


Figure 6: Confidence maps of both SfM softwares for their respective reconstructions.

It is immediately apparent that the software’s confidence in the reconstruction, and therefore its accuracy, is higher in regions where more features can be matched. For *Agisoft Metashape*, the areas surrounding the green laser dots on the sample’s surface correspond to high-confidence regions, while confidence decreases farther from the dots. Additionally both software tools struggled to achieve reliable reconstruction in the right-hand region where blue tape covered part of the sheet of paper. Finally, the confidence map produced by *RealityScan* appears at first glance to be more optimistic, however this should be interpreted with caution, as the software seems overly confident and exhibits clear artifacts near the edges of the reconstructed surface, particularly in the upper-right edge of the reconstruction.

5 Future work

The results presented in this study demonstrate that in-situ 3D reconstruction inside a high enthalpy facility is feasible, provided that cameras are carefully arranged, pre-calibrated, and supported by sufficient surface features to ensure robust image matching. Several limitations remain, however, and these outline the next steps for continuing this work.

A first priority is to reproduce and extend these experiments at the VKI Plasmatron. The facility offers different optical access and higher plasma luminosity compared to the NASA ARCTRON setup. Demonstrating reconstruction robustness under these conditions will be essential for validating the approach across multiple high-enthalpy environments. In particular, future work will investigate the impact of stronger plasma

emission on feature detection, evaluate potential filtering strategies, and test reconstruction stability when the sample undergoes significant deformation.

A second direction will be to repeat the in-situ tests at NASA Ames once the ARCTRON nozzle is operational. Concentrating the plasma column will increase local heat flux and provide more representative ablation conditions, enabling a more realistic assessment of reconstruction performance and the influence of plasma-induced optical distortion.

On the software side, additional reconstruction tools will be evaluated to improve performance, robustness, and processing speed. A newly released version of *Meshroom* (August 2025) includes upgraded feature-detection pipelines, improved depth-map computation, and more advanced GPU acceleration. Future work will benchmark this new release against *Agisoft Metashape* and *RealityScan*, both for in-chamber and in-test reconstructions. Integrating the Meshroom pipeline into an automated acquisition system may further enable near-real-time reconstruction.

Finally, subsequent steps will involve developing a fully automated workflow including camera triggering, synchronized data capture, real-time preprocessing, and live monitoring of reconstruction quality. Achieving this level of automation will be necessary for routine operational use in plasma wind tunnels and arc-jet facilities.

Overall, the results of this study form the basis for a broader effort to establish reliable in-situ 3D reconstruction capabilities at NASA Ames, VKI, and other high-enthalpy test facilities.

Acknowledgments

The authors would like to thank Dr. Magnus A. Haw, Sebastián V. Colom, Dr. Joseph El Rassi, Dr. Bruno Dias and Professor Thierry Magin for their valuable guidance and support. This research is supported by Université libre de Bruxelles (ULB), the von Karman Institute (VKI), Internal Research and Development (IRAD) funds from the Thermal Protection Materials Branch, by the Entry System Modeling (ESM) project and the Mars Sample Return (MSR) project.

References

- [1] Carsten Griwodz, Simone Gasparini, Lilian Calvet, Pierre Gurdjos, Fabien Castan, Benoit Maujean, Gregoire De Lillo, and Yann Lanthony. Alicevision meshroom: An open-source 3d reconstruction pipeline. In *Proceedings of the 12th ACM Multimedia Systems Conference (MMSys '21)*. ACM Press, 2021.
- [2] Edward Schairer and James Heineck. *Photogrammetric Recession Measurement of Ablative Materials During Arc-Jet Testing*.
- [3] Alexandre Quintart, Magnus A. Haw, and Sebastian Colom. New tools for automating arcjet sample

- recession tracking and analysis. Technical Report 20240013872, NASA Ames Research Center, 2024. Presentation at the 14th Ablation Workshop, Laurel, MD, November 5–7, 2024.
- [4] Magnus A. Haw, Alexandre M. Quintart, and Sebastian V. Colom. STARSscan: Spatial Targeting and Alignment Rig for Scanning. Technical Report 20240013702, NASA Ames Research Center, 2024. Poster presented at the 14th Ablation Workshop, Laurel, MD, November 5–7, 2024.
- [5] Alexandre M. Quintart, Magnus A. Haw, and Federico Semeraro. arcjetcv: Open-source software to analyze material ablation. *Journal of Spacecraft and Rockets*, 62(5):1644–1653, 2025.
- [6] Stefan Loehle, Tina Staebler, and Thomas Reimer. *Photogrammetric Surface Analysis of Ablation Processes in High Enthalpy Air Plasma Flow*. Atlanta, Georgia, USA, June 2014.
- [7] Colin M. Yee, Evangelos Y. Maragoudakis, Kyle Woo, Noah Langton, Tyler T. Joyce, and Joseph H. Koo. *Design of a Photogrammetric System for Real Time Surface Morphology Analysis*.
- [8] Johannes L. Schönberger and Jan-Michael Frahm. Structure-from-motion revisited. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4104–4113, 2016.
- [9] Klaus Häming and Gabriele Peters. The structure-from-motion reconstruction pipeline – a survey with focus on short image sequences. *Kybernetika*, 46(5):926–937, 2010.

Flow/material coupling methods

A coupling strategy of flow and material solvers for modeling of high-enthalpy flows with ablation

Martins, Diana^{1,2} Borbouse, Maxime^{1,3,4} Klauner, Julien^{1,2,4} Schrooyen, Pierre¹
Henneaux, David⁵ Chatelain, Philippe⁴ Hillewaert, Koen^{1,3,5} Magin, Thierry^{1,2}

¹von Karman Institute for Fluid Dynamics ²Université Libre de Bruxelles

³Université de Liège ⁴Université catholique de Louvain ⁵Cenaero

December 4, 2025

Abstract

This work demonstrates a staggered coupling strategy between a flow solver and a material solver to model high-enthalpy flows with ablation. The approach is validated using a test case based on the Zoram material, a carbon-phenolic composite, under experimental conditions from the VKI Plasmatron facility. The coupling leverages the PreCICE library for data exchange, handling the mapping and interpolation of the solution at the interface. The flow solver, ArgoDG, computes the high-enthalpy aerothermal environment as well as recession, while the material solver, PATO, models the thermal and mechanical response of the ablative material, including thermal response and pyrolysis. Results show that the coupled simulation captures the interaction between the flow and material, with the interface evolving dynamically over time. While the coupled approach overestimates stagnation point temperature and recession compared to experimental data, potential causes such as oxidation modeling and inflow conditions are discussed. This work establishes a foundation for future studies on more complex materials and coupling strategies.

Keywords: hypersonics, ablation, recession, discontinuous Galerkin, high-enthalpy flows, coupling strategy, thermal protection systems

1 Introduction

During atmospheric entry, spacecraft encounter extreme aerodynamic heating that poses severe constraints on structural integrity. To withstand such harsh conditions, Thermal Protection Systems (TPS) are employed. These often use ablative materials that undergo pyrolysis and surface ablation at high temperatures, absorbing the incoming heat through thermal decomposition and material removal [13].

Predicting the response of ablative TPS materials for design and sizing purposes remains a significant challenge. This process requires highly detailed modeling of the aerothermal environment, the interaction between gas and surface, and the internal thermal response of the materials. Additionally, since the high-enthalpy flow surrounding the capsule and the material behavior are strongly dependent, accurate prediction of the ablative heat flux must account for the evolving surface shape during material recession, as demonstrated in the work of Gokcen et al. [10].

The numerical modeling of ablation is commonly

divided into four distinct strategies. The first involves using a flow solver with an ablative boundary condition to directly impose heat and mass balances [2, 6, 15, 25]. In the second approach, the primary focus is on studying the material's response by using approximated flow correlations applied at the surface boundary. This method employs the B' formalism, which represents a coupled mass and energy flux-balance condition derived from a control volume analysis at the interface between the external flow and the material surface [17]. This formulation acts as a blowing correction factor, accounting for the effect of surface mass injection into the boundary layer on convective heat transfer at the surface. This method has been used over the past year for TPS design, and is implemented in ablation codes such as FIAT [5], PATO [12] and KATS [28]. To overcome the limitations of the approximated boundary conditions, the third approach establishes a staggered coupling between the two previous models. Different coupling methodologies have been proposed, each adapted to a specific flow environment and material behaviour, and differing in time-stepping depending on the explicit

or implicit nature of the coupling and on the parameters that are exchanged [4, 8, 21, 23, 24]. Building on these efforts, Schroeder et al. [1, 20] developed ARES, a coupling framework designed to link the flow and material solvers, accurately managing the time and spatial exchange of data during the simulation. Finally, the last strategy integrates numerical tools capable of simultaneously resolving the flow and material domains in a unified framework [9, 22]

The present work consists of coupling two simulation tools to address the two main aspects of ablative material modeling: the aerothermal environment, handled by the ArgoDG [11] solver, and the thermochemical material response with the Porous-Material Analysis Toolbox based on OpenFOAM code, PATO [12]. Data exchange between the two solvers is performed through the preCICE library [7], which enables interpolation, mapping and temporal coupling of the boundary quantities at the interface of the two domains. This is accomplished by solving the surface mass balance within the flow solver and the surface energy balance within the material solver. As a result, the material surface recession is determined by the flow based on finite-rate chemistry, while the material code computes the surface temperature and the pyrolysis gas injection rate, following an explicit approach.

This paper will first describe the numerical coupling methodology. The developed methodology will then be validated against an experimental test case performed in the inductively coupled Plasmatron facility at the von Karman Institute, focusing on the ablative carbon-phenolic Zeram material [26]. Finally, conclusions and perspectives for further improvements will be provided.

2 Methods

The method is based on two solvers. *ArgoDG*, a flow solver based on discontinuous Galerkin method and immersed boundary, able to simulate high-enthalpy flows with thermo-chemical non-equilibrium and ablation [3, 11], is used to solve the flow environment around the probe. *PATO*, a material solver based on the finite volume method and on the framework of OpenFOAM [27], able to simulate the thermo-chemical and mechanical response of ablative materials [12], is used to solve the material response, including recession and pyrolysis.

2.1 Boundary conditions

Reaching convergence requires a conservative and consistent model, which implies a well-posed setup and correctly implemented boundary conditions. Therefore, the equations solved at the interfaces between the two solvers need to be chosen with care. In this work, the mass balance is solved on the flow solver side by imposing, as a Neumann boundary, the mass fluxes J_i of species i as the right-hand side of the balance. It is given by

$$J_i = -\frac{\rho_i \dot{m}_{\text{blow}}}{\rho} + \hat{\omega}_i + \dot{m}_g C_{i,\text{pyro}}, \quad (2.1)$$

with

$$\dot{m}_{\text{blow}} = \underbrace{\sum_i \hat{\omega}_i}_{\dot{m}_c} + \dot{m}_g, \quad (2.2)$$

ρ_i being the density of species i , ρ the total density, \dot{m}_{blow} the total blowing rate, $\hat{\omega}_i$ the production rate of species i from chemical reactions in the flow, \dot{m}_c the mass blowing rate from char oxidation and \dot{m}_g the mass blowing rate from pyrolysis and $C_{i,\text{pyro}}$ the mass fraction of species i in the pyrolysis gas. The latter is computed under the assumption of chemical equilibrium with constant elemental fractions. The mass blowing rate from pyrolysis, \dot{m}_g , is provided by PATO. The pyrolysis gas flow within the porous material is calculated using Darcy's law, with gas motion driven by pressure gradients and temperature. The wall normal velocity is imposed inside ArgoDG as a Dirichlet condition and corresponds to the total blowing velocity $v_{\text{blow}} = \dot{m}_{\text{blow}}/\rho$. On PATO's side, the energy balance is computed as

$$q_{\text{conv,CFD}} + \dot{m}_c h_c + \dot{m}_g h_g - \sigma \varepsilon_w (T_w^4 - T_\infty^4) - q_{\text{cond}} = 0, \quad (2.3)$$

with

$$q_{\text{conv,CFD}} = -k \nabla T + \sum h_i J_i - h_w (\dot{m}_c + \dot{m}_g), \quad (2.4)$$

$q_{\text{conv,CFD}}$ being the convective heat flux computed on the flow solver side, h_c the enthalpy of the char oxidation products, h_g the enthalpy of the pyrolysis gas, σ the Stefan-Boltzmann constant, ε_w the wall emissivity, T_w the wall temperature, T_∞ the ambient temperature, q_{cond} the conductive heat flux inside the material, k the thermal conductivity, h_i the enthalpy of species i , and h_w the enthalpy of the gas at the wall. The heat flux $q_{\text{conv,CFD}}$ and the ablation mass blowing rate, \dot{m}_c , are provided by ArgoDG as constant input values at the surface and are imposed in PATO as Neumann boundary conditions. As a result, surface temperature, T_w , is retrieved from PATO, and then passed to ArgoDG, where it is imposed as a Dirichlet boundary condition.

2.2 Time-stepping strategy

The two solvers communicate through the software *PreCICE* [7], which ensures the mapping and interpolation of the solution at the interface between them. The time iterations are launched and handled in PATO, which computes the material response for a given time step, knowing the pressure, the heat flux, the ablative blowing rate and the shear stress at the interface. The latter exchanges the position of the interface, the surface temperature, and the pyrolysis blowing rate with ArgoDG. ArgoDG then computes the flow solution in a steady-state manner. Once converged, ArgoDG sends the necessary boundary conditions back to PATO. A schematic view of the coupling is proposed in Figure 1. These steps are repeated until the final time is reached, supposing that ablation, recession, and pyrolysis are fully resolved.

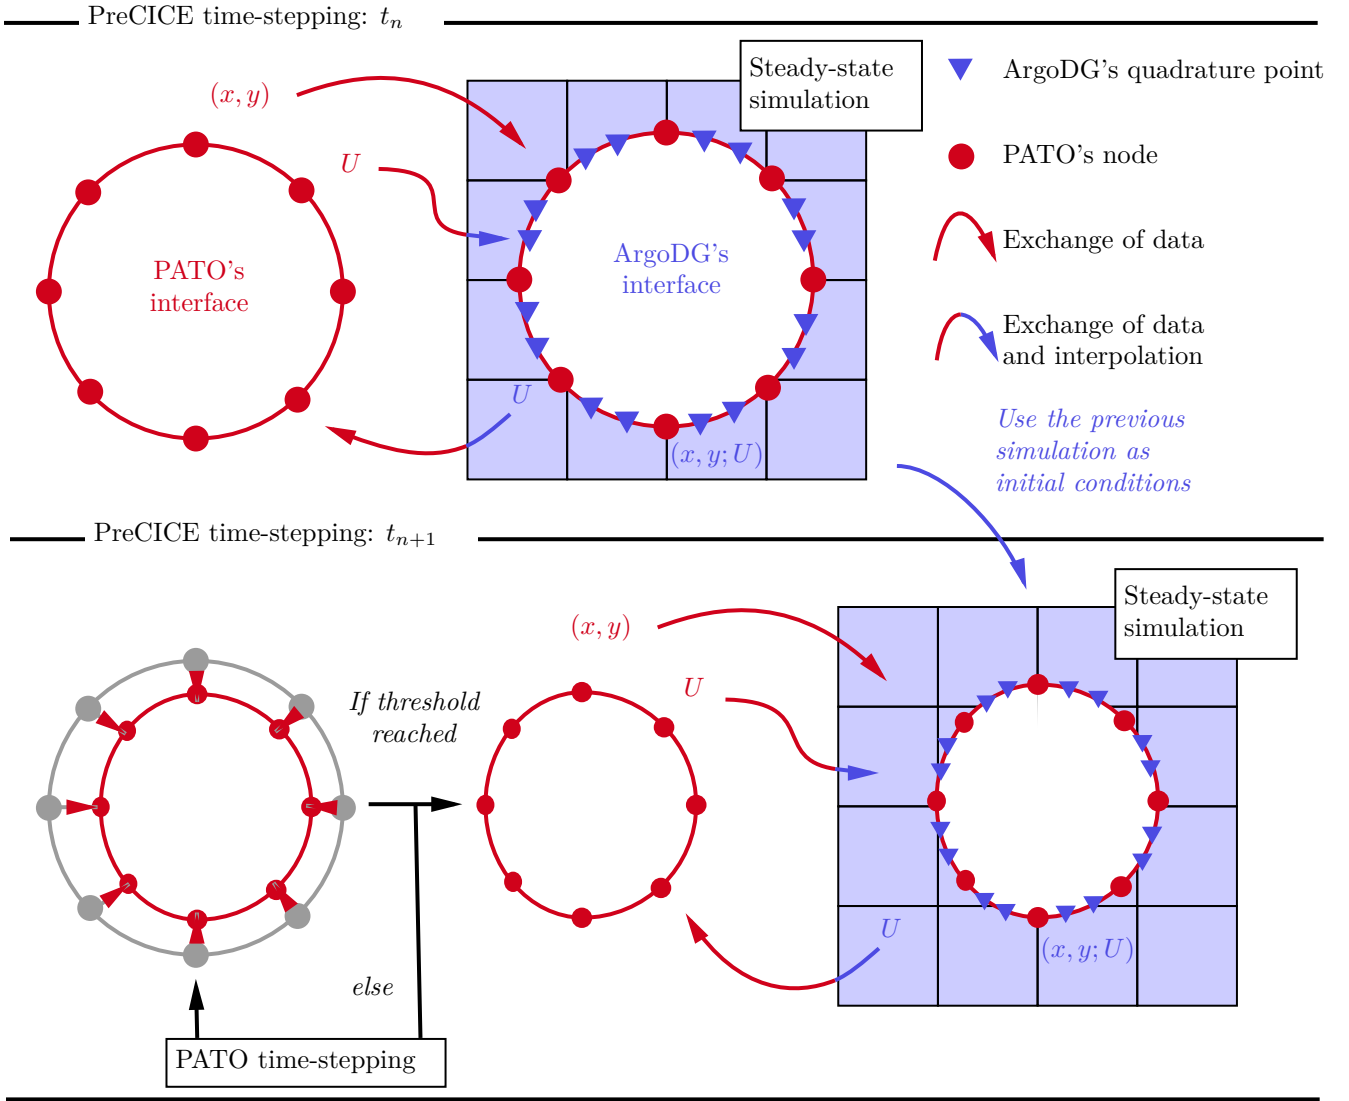


Figure 1: Coupling algorithm between ArgoDG and PATO using PreCICE.

2.2.1 Thresholds and exchange criteria

The boundary conditions imposed in PATO remain constant during a time step. To avoid excessive calls to ArgoDG, which would be computationally expensive, the latter is only called when certain thresholds are reached. In this work, we impose a threshold on the increases in temperature and the recession at the stagnation point. The temperature threshold is $\Delta T = 50$ K, and for recession is $\Delta \delta_{\text{recc}} = 0.32$ mm. The temperature threshold is low enough to correctly consider the transitional phase of the temperature, and the recession threshold allows the flow to be updated during the established phase. These two thresholds remain wide, to avoid too many calls to ArgoDG, but can be reduced to gain precision.

3 Results

3.1 Numerical setup

A verification and validation test case is conducted to assess the correctness of the coupling implementation. The latter corresponds to a case from an experimental campaign on a 25 mm radius hemispherical sample made of carbon-phenolic Zeram [29], conducted in the VKI's high-enthalpy Plasmatron under the project AblanTIS [26]. Numerical reproduction of these experiments with PATO are also available in the AblanTIS project and will serve as verification. These previous numerical simulations were performed in PATO using flow approximations at the boundary, based on the B' method, as described in [16]. This approach is required when the material solver is used in standalone mode. The selected test case is Exp-2-x-Z.

The numerical setup is presented in Figure 2, with the corresponding boundary conditions. The initial temperature for the material is $T = 300$ K. In this case, the pressure profile along the sample is kept

constant as $P_w = 5417$ Pa. For the flow, the inlet conditions are constant and evaluated based on the boundary layer edge total enthalpy given by the rebuilt free-stream testing conditions [26]. The chosen mixture for the flow is composed by 12 species: $N, O, NO, N_2, O_2, C, CO, CO_2, CH_4, H_2, OH, H_2O$, as taken in [14] for study of charring ablator. For the chemical rates of the mass balance, Equation 2.1, only oxidation is chosen and follows the Park model [18] where the recombination probability β_i is of Arrhenius type and depends on the temperature as

$$\beta_i = \gamma_i \exp(-E_a/(RT_w)), \quad (3.1)$$

with $\gamma_i = 0.63$ a recombination probability constant, E_a the Arrhenius activation energy and R the universal gas constant where $E_a/R = 1160$ K.

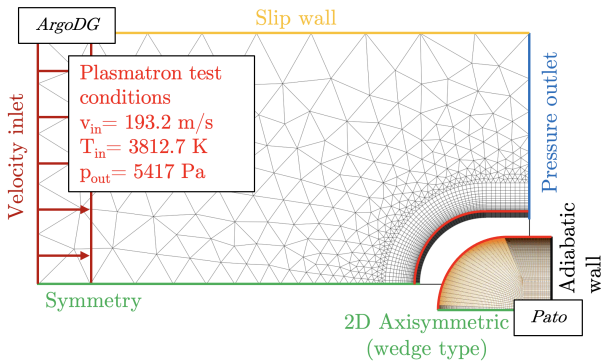


Figure 2: Numerical setup for the Zoram test case.

3.2 Simulation results and discussion

At first, the coupling works as expected. The exchange is performed correctly using the second-order interpolation scheme from preCICE each time the threshold on temperature or on recession is reached. The deformed interface from the material solver is correctly retrieved in ArgoDG and rebuilt using the immersed interface method, allowing the flow quantities to be updated with respect to the new geometry. The stagnation point temperature and recession are presented in Figures 4 and 5, respectively, together with the results obtained using only PATO, only ArgoDG, and the experimental data. Note that the bump in temperature appearing at approximately $t = 55$ s for the PATO simulation corresponds to the moment when the mass blowing rate of the pyrolysis gas becomes negative, meaning the solver considers that the material is absorbing pyrolysis gas from outside, which rapidly increases the surface temperature. However, since this phenomenon is not physical and there is no production of pyrolysis gas inside the flow, ArgoDG ignores the pyrolysis contribution in the mass balance whenever the associated blowing rate becomes negative. When comparing the different numerical methods, all the temperature estimations are of the same order of magnitude as the experimental data, but they all over-predict it. The same observation applies to the




Figure 3: Temperature fields of the flow and the material for time $t = 1, 50, 100$ and 150 s. The blue line corresponds to the interface evolving with time, while the translucent blue one corresponds to the interface at initial time.

recession at the stagnation point and so to the mass blowing rate from ablation.

However, the temperature from ArgoDG and the coupling simulation is even more overestimated than with only PATO. It has been observed that oxidation is too high, and that reducing it lowers both the temperature and the recession. This may be influenced by the finite-rate model, although it uses the Park model, which is widely used in the literature and should not be adjusted. Another possible cause could be the inlet conditions, which are assumed to be constant rather than following a torch profile. Moreover, the imposed inflow temperature and velocity are determined based on an approximated reconstruction from matching the mixture enthalpy at the boundary layer edge with the one evaluated from free-stream reconstruction, where the latter reconstruction is another source of error. Furthermore, in [26], that conducts the numerical reproduction of this test case using PATO, also point out the over-estimation of both temperature and recession and could not find apparent reasons. They express their uncertainty about the correct evaluation of the oxygen content in the free-stream conditions, or a quicker recombination of atomic oxygen in the boundary layer against the one computed by the gas-chemistry model in the flow, as well as pyrolysis gas injection that could have provided oxidation protection to the downstream.

4 Conclusions and perspective

This work details the implementation of a coupling framework between ArgoDG, the flow solver, and PATO, the material response solver. The exchange of coupled boundary conditions is enabled using the preCICE library through an explicit coupling approach.

A hemispherical 25 mm radius probe made of carbon-phenolic Zoram material exposed to the high-




Figure 4: Comparison of the temperature at stagnation point with respect to time for the coupling, the material solver, the flow solver and the experimental values. The red dots for the coupling curve corresponds to the calls to the flow solver to update the quantities for the material solver.

enthalpy Plasmatron flow is simulated using the developed methodology. The results indicate that the exchange is correctly achieved whenever the temperature or recession threshold is reached. Data is accurately transferred and preserved through the preCICE second-order nearest-neighbour interpolation scheme. When comparing the numerical with the experimental results, we observe an overestimation of both temperature and recession. While further investigations should be performed, this discrepancy may be attributed to uncertainties in the free-stream oxygen content, simplified inflow boundary conditions, or limitations of the finite-rate gas chemistry model and pyrolysis gas injection effects. In addition, uncertainties in the material database can also contribute to the observed deviations. Future work should include a sensitivity analysis to assess the influence of such parameters.

The present study provides a basis for extending the coupling methodology to more complex materials, such as cork, whose swelling and shrinkage during ablation introduce dynamic surface evolution that must be taken into account. Future advancements to improve both efficiency and accuracy focus on enhancing the current exchange strategies, such as implicit coupling and high-order spatial interpolation methods. The latter can be accomplished by using the CWIPI coupling library [19], instead of preCICE.

Acknowledgments

The present research of M. Borbouse and J. Klauner is part of the Space4ReLaunch project, supported by the SPW Economie Emploi Recherche of the Walloon




Figure 5: Comparison of the recession of the stagnation point with respect to time for the coupling, the material solver and the experimental values. The red dots for the coupling curve corresponds to the calls to the flow solver to update the quantities for the material solver.

Region, under the grant agreement n°2210181. The PhD of D. Martins has the financial support of the ‘Fonds de la Recherche Scientifique (F.R.S. -FNRS)’ for the FRIA grant with reference 40022234. M. Borbouse and J. Klauner express their sincere gratitude to the Université de Liège and Université catholique de Louvain for the additional university grants that facilitated the conduct of this research.

The present research benefited from computational resources made available on Lucia, the Tier-1 supercomputer of the Walloon Region, infrastructure funded by the Walloon Region under the grant agreement n°1910247.

The authors warmly acknowledge the NASA Ames Research Center and the Analytical Mechanics Associates, in particular Bruno Dias, Giorgios Bellas, Sergio Fraile, for having organized this summer program, and Olivia M. Schroeder for their comments on boundary conditions. Special thanks to all the team members present on-site: Université de Bordeaux—particularly the PATO team for their invaluable assistance in acquiring the source code of PATO—, École Polytechnique de Paris and ONERA.

References

- [1] Georgios Bellas Chatzigeorgis, Olivia Schroeder, Joseph C. Schulz, Grant E. Palmer, and Eric C. Stern. Coupled fluid dynamics and material response simulations for nitrogen arc jet flows. In *AIAA AVIATION FORUM AND ASCEND 2024*, page 3815, 2024.
- [2] Daniele Bianchi, Francesco Nasuti, and Emanuele Martelli. Navier-stokes simulations of hypersonic

- flows with coupled graphite ablation. *Journal of Spacecraft and Rockets*, 47(4):554–562, 2010.
- [3] Corentin Carton de Wiart. *Towards a Discontinuous Galerkin solver for scale-resolving simulations of moderate Reynolds number flows, and application to industrial cases*. PhD thesis, Université catholique de Louvain, May 2014.
- [4] Yih-Kanq Chen and Tahir Gökçen. Implicit coupling approach for simulation of charring carbon ablators. *Journal of Spacecraft and Rockets*, 51(3):779–788, 2014.
- [5] Yih-Kanq Chen and Frank S. Milos. Ablation and thermal response program for spacecraft heatshield analysis. *Journal of Spacecraft and Rockets*, 36(3):475–483, 1999.
- [6] Yih-Kanq Chen and Frank S. Milos. Navier-stokes solutions with finite rate ablation for planetary mission earth reentries. *Journal of Spacecraft and Rockets*, 42(6):961–970, 2005.
- [7] G. Chourdakis, K. Davis, B. Rodenberg, M. Schulte, F. Simonis, B. Uekermann, G. Abrams, H. J. Bungartz, L. Cheung Yau, I. Desai, K. Eder, R. Hertrich, F. Lindner, A. Rusch, D. Sashko, D. Schneider, A. Totounferoush, D. Volland, P. Vollmer, and O. Z. Koseomur. preCICE v2: A sustainable and user-friendly coupling library [version 2; peer review: 2 approved]. *Open Research Europe*, 2(51), 2022.
- [8] Brett A. Cruden. On the fluid dynamics boundary condition in ablating or blowing flows. In *AIAA Aviation 2023 Forum*, page 3614, 2023.
- [9] Bruno Dias, Aleksander Zibitsker, Jeremie Meurisse, and Nagi Mansour. Towards a flow-material unified solver for heatshield modeling. 01 2023.
- [10] Tahir Gökçen, Yih-Kanq Chen, Kristina A. Skokova, and Frank S. Milos. Computational analysis of arc-jet stagnation tests including ablation and shape change. *Journal of Thermophysics and Heat Transfer*, 24(4):694–707, 2010.
- [11] Koen Hillewaert. *Development of the discontinuous Galerkin method for high-resolution, large scale CFD and acoustics in industrial geometries*. PhD thesis, Université catholique de Louvain, February 2013.
- [12] Jean Lachaud and Nagi N. Mansour. Porous-Material Analysis Toolbox Based on OpenFOAM and Applications. *Journal of Thermophysics and Heat Transfer*, 28(2):191–202, April 2014.
- [13] Bernard Laub and Ethiraj Venkatapathy. Thermal protection system technology and facility needs for demanding future planetary missions. In *planetary probe atmospheric entry and descent trajectory analysis and science*, volume 544, pages 239–247, 2004.
- [14] Claudio Miccoli, Alessandro Turchi, Pierre Schrooyen, Domenic D’Ambrosio, and Thierry Magin. Detailed modeling of cork-phenolic ablators in preparation for the post-flight analysis of the garman re-entry cubesat. *Aerotecnica Missili & Spazio*, 100:207–224, 06 2021.
- [15] Frank S. Milos and Daniel J. Rasky. Review of numerical procedures for computational surface thermochemistry. *Journal of thermophysics and heat transfer*, 8(1):24–34, 1994.
- [16] Alberto Padovan et al. An extended b’ formulation for ablating-surface boundary conditions. *International Journal of Heat and Mass Transfer*, 218:124770, 2024.
- [17] Alberto Padovan, Blaine Vollmer, Francesco Panerai, Marco Panesi, Kelly A Stephani, and Daniel J Bodony. An extended b’ formulation for ablating-surface boundary conditions. *International Journal of Heat and Mass Transfer*, 218:124770, 2024.
- [18] Chul Park, Richard L. Jaffe, and Harry Partridge. Chemical-kinetic parameters of hyperbolic earth entry. *J. Thermophys. Heat Transf.*, 15(1):76–90, January 2001.
- [19] A. Refloch, B. Courbet, A. Murrone, Philippe Villedieu, Claire Laurent, P. Gilbank, J. Troyes, Lionel Tessé, G. Chaineray, J. Dargaud, Eric Quémerais, and F. Vuillot. CEDRE Software. *AerospaceLab Journal*, March 2011.
- [20] Olivia Schroeder, Joseph Brock, Joseph C. Schulz, Georgios Bellas-Chatzigeorgis, Prakash Shrestha, Grant Palmer, Eric Stern, and Graham V. Candler. Ares: A coupling methodology for ablation modeling. In *AIAA SCITECH 2024 Forum*, page 1477, 2024.
- [21] P. Schrooyen, B. Dias, A. Fagnani, A. Turchi, B. Helber, L. Walpot, and T.E. Magin. ABLATION MODELING OF LIQUEFYING MATERIALS. Heilbronn, Germany, 2022.
- [22] Pierre Schrooyen, Joffrey Coheur, Alessandro Turchi, Koen Hillewaert, Philippe Chatelain, and Thierry Magin. Numerical simulation of a non-charring ablator in high enthalpy flows by means of a unified flow-material solver. In *47th AIAA Thermophysics Conference*, page 3352, 2017.
- [23] Pierre Schrooyen, Alessandro Turchi, Koen Hillewaert, Philippe Chatelain, and Thierry E. Magin. Two-way coupled simulations of stagnation-point ablation with transient material response. *International Journal of Thermal Sciences*, 134:639–652, December 2018.
- [24] John M. Thornton, Dinesh K. Prabhu, Jeremie B. Meurisse, Arnaud Borner, Joshua D. Monk, and Brett A. Cruden. Coupling Heatshield Response and Aerothermal Environment for Mars Entry via Surface Gas Blowing. In *AIAA SCITECH 2023*

Forum, AIAA SciTech Forum. American Institute of Aeronautics and Astronautics, January 2023.

- [25] Alessandro Turchi, Daniele Bianchi, Francesco Natusi, and Marcello Onofri. A numerical approach for the study of the gas–surface interaction in carbon–phenolic solid rocket nozzles. *Aerospace Science and Technology*, 27(1):25–31, 2013.
- [26] Alessandro Turchi, Bernd Helber, Francisco Torres Herrador, Thierry Magin, Luke Chipperfield, F. Pascon, T. Eekelen, and H. Ritter. Ablative-material numerical-test international series (ablantis): An experimental/numerical effort to support the validation of material thermal-response tools. 09 2019.
- [27] H. G. Weller, G. Tabor, H. Jasak, and C. Fureby. A tensorial approach to computational continuum mechanics using object-oriented techniques. *Comput. Phys.*, 12(6):620–631, November 1998.
- [28] Haoyue Weng and Alexandre Martin. Development of a universal solver and its application to ablation problems. In *47th AIAA Thermophysics Conference*, page 3355, 2017.
- [29] C. Zuber, T. Reimer, B. Esser, A. Gülhan, G. Herdrich, and N. Biller. Development of the low-density phenolic-based fibrous ablator ZURAM-K. *J. Spacecr. Rockets*, 58(2):415–424, March 2021.

Development of a high-order time-adaptive multistep coupling scheme in an HPC library for multiphysics applications

A. E. Simon^{1,2} L. François¹ A. Asad² E. Quémerais¹ M. Massot²

¹DMPE, ONERA, Université Paris-Saclay, 91120 Palaiseau, France

²CMAP, CNRS, École polytechnique, Institut Polytechnique de Paris, 91120 Palaiseau, France

Abstract

The multistep coupling scheme has recently been introduced as a new temporal coupling strategy for multiphysics simulations involving multiple specialized monophysics solvers. It enables high-order temporal accuracy and dynamic adaptation of the coupling time step. This work presents the first implementation of the multistep coupling scheme, building on the open-source parallel spatial interpolation library CWIPI. A multidimensional solid-solid conjugate heat transfer test case demonstrates the capabilities of the approach.

Keywords: multiphysics simulation, high order, multistep code coupling, time adaptation, error estimate, conjugate heat transfer

1 Introduction

Complex systems often involve multiphysics interactions as well as multiple temporal and spatial scales. To simulate such systems, it is advantageous to rely on existing specialized and validated solvers for each monophysics subsystem – the so-called partitioned approach – rather than developing a new monolithic solver. However, numerically coupling several specialized solvers remains a challenging task. Two main issues arise: one related to spatial coupling and the other to temporal coupling. On the spatial side, the difficulty lies in accurately and efficiently exchanging information between non-conforming meshes that may overlap or remain disjoint. A number of libraries have already been developed to tackle this task, *e.g.* CWIPI [1] or PRECICE [2]. On the temporal side, the challenge is to let the solvers evolve in a coherent manner, dealing with their differing time scales and time integration techniques, while ensuring both the stability and accuracy of the overall coupled solution. Simple techniques, *e.g.* Jacobi/parallel-staggered, or Gauss-Seidel/serial coupling are traditionally used, where the various solvers exchange at discrete coupling time points the required information (*e.g.* heat flux, temperature), referred to as coupling variables. The exchanged quantities are assumed constant until the next coupling point. They however lack dynamic

time step adaptation, and cannot go beyond second-order accuracy [3]. Various monolithic-like approaches have also been proposed, *e.g.* multirate schemes or additive Runge-Kutta schemes, offering interesting numerical properties, at the cost of being much more intrusive [4], [5]

The multistep coupling scheme introduced in [6], [7] is a newly proposed partitioned approach that uses a polynomial approximation of the temporal evolution of coupling variables to achieve a high order of accuracy in time. It could be interpreted as a generic adaptive high-order time-domain decomposition method, with explicit or implicit variants. The latter is useful to improve the stability of a coupling, at the expense of being iterative in nature. It has initially been implemented in the Python demonstrator RHAPSOPY¹, showing promising results. This contribution aims at presenting a first implementation of the multistep coupling scheme within the open-source C++ library CWIPI, which already provides a space interpolation framework for coupled simulations in a massively parallel context. Dynamic adaptation of the coupling time step based on error estimates is also discussed. Finally, the capabilities of the proposed approach are assessed on a conjugate heat transfer test-case between two solids using solvers based on the SAMURAI [8] and PETSC [9] libraries.

¹github.com/hpc-maths/Rhapsopy

2 The multistep coupling scheme

The multistep coupling scheme is briefly recalled. The approach is similar to staggered schemes, but instead of keeping the coupling variables frozen between each coupling time points, their evolution is prescribed by high-order polynomial predictions based on past values.

The time-continuous coupled problem, after all required spatial discretizations have been applied on each monophysics submodel, is described by the set of $N \geq 2$ monophysics problems:

$$\begin{cases} d_t y_i(t) = \mathcal{F}_i(y_i(t), u_i(t), t) & (2.1) \\ u_i = h_i(y_1, \dots, y_M) & (2.2) \end{cases}$$

with $i \in \llbracket 1, N \rrbracket$ the index of the subsystem, u_i the coupling variables (*e.g.* wall temperature or heat flux), which a priori depend on all state vectors through the Lipschitz functions h_i .

Let k be a non-negative integer. Over a coupling step $[t_n, t_{n+1}]$, the system solvers integrate in parallel or sequentially their internal dynamics described by ODE (2.1), where $u_i(t)$ is replaced by $\hat{u}_i^n(t)$ the degree- k polynomial extrapolation defined by the past points $u_i^{n-k} \dots u_i^n$ at times $t_{n-k} \dots t_n$. The resulting multistep scheme is explicit. An implicit variant can be formulated by incorporating the yet unknown coupling variables u_i^{n+1} at time t_{n+1} in the definition of the interpolation polynomial $\hat{u}_i^n(t)$. This leads to a nonlinear system, which can be solved iteratively, for instance by means of Picard iterations. The iterative procedure is initialized with an explicit prediction to get a first value for u_i^{n+1} and stopped once the coupling variables have converged. In both variants, the order of convergence in time is $k + 1$.

One advantage of this scheme is that the solvers can use any integration technique for their subsystem, potentially with substepping. Their time schemes should only be at least of the same order as the coupling scheme, and care should be taken in the code to evaluate the polynomial approximations at each internal substep (*e.g.* each stage of a Runge-Kutta method).

A first natural and cost-free estimate ϵ_i of the coupling error over a single step can be obtained by comparing the coupling variables updated by the solvers after integration over $[t_n, t_{n+1}]$ with the initial explicit extrapolation $\hat{u}_i^n(t_{n+1})$:

$$\epsilon_i = \frac{|u_i^{n+1} - \hat{u}_i^n(t_{n+1})|}{\text{atol} + \text{rtol} \cdot |u_i^{n+1}|} \quad (2.3)$$

For given absolute and relative tolerances atol and rtol , this error estimate can be used to compute an optimal time step Δt_{opt} given in Equation (2.4) that ensures the estimated relative error is not larger than 1:

$$\Delta t_{\text{opt}} = \beta \Delta t_n \cdot \left[\max_i \epsilon_i \right]^{-\frac{1}{k+1}} \quad (2.4)$$

with β a safety factor, typically set to 0.85. A second error estimate, referred to as *embedded*, can be constructed by comparing with a higher-order or lower-order coupled solution. This however roughly doubles the cost of each coupling step.

The estimate based on comparing the extrapolated and updated values is well suited to the explicit case, but may limit the implicit coupling scheme to use time steps within the stability limit of the explicit variant, since the extrapolation is inherently explicit. The embedded estimate may thus be better suited for an implicit coupling.

The multistep coupling framework has been introduced in CWIPI during the Modelling Summer Visit. The library handles internally the construction of the polynomials and the error estimate. A small set of new functions are available to steer the temporal coupling, which each coupled code should call in a specific order in their internal time loop. For the space interpolation, CWIPI already requires each code to provide a memory buffer where the interpolated data to be received can be written. Upon calling the new function `CWP_Predict(t)` at each internal substep, these buffers are updated with the corresponding evaluation of polynomial approximations.

The implicit variant does not require any additional function call from the codes, everything being handled internally by CWIPI.

3 A conjugate heat transfer test case

To demonstrate the applicability of the coupling strategy with CWIPI, we consider in this section the heat transfer between two solids, as an academic test case representative of typical applications where diffusion is the dominating process near the coupling interface.

3.1 Modelling

The configuration of the test case is depicted in Figure 1. The left domain $[-L_{s1}, 0] \times [0, H]$ is subjected to a Dirichlet boundary condition on the left side $T(x = -L_{s1}, y, t) = T_H(y) = T_H \cdot (H - y) / H$ and on the bottom side $T(x, y = 0, t) = T_H(x) = T_H \cdot (L_{s1} - y) / L_{s1}$, and a zero-flux Neumann boundary condition on the top side. The right domain $[0, H] \times [0, H]$ is subjected to a Dirichlet boundary condition on the right side $T(H, y, t) = T_C$, a zero-flux Neumann boundary condition on the top and bottom sides. The coupling condition between both domains at $x = 0$ at the continuous level is obtained by imposing the continuity of the temperature and the normal heat flux, assuming both domains have the same conductivity: We set $H = L_{s2} = 1$ and $L_{s2} = 5 \cdot L_{s1}$.

Within both solid domains, the temperature field is governed by the heat equation, with α_{s1} (resp. α_{s2}) the diffusivity of the left solid domain (resp. right):

$$\frac{\partial T_{si}}{\partial t} = \alpha_{si} \Delta T_{si} \quad i = 1, 2 \quad (3.1)$$

In order to reduce the number of parameters of the system, the equations have been put in dimensionless form, using $\theta = (T - T_C) / (T_H - T_C)$, $\tau = t / t_{ref}$ where $t_{ref} = H^2 / \alpha_{s2}$, $\bar{x} = x / H$ and $\bar{y} = y / H$. Denoting the




Figure 1: Conjugate heat transfer configuration

dimensionless Laplacian operator as $\bar{\Delta}$, the dimensionless equations read:

$$\begin{cases} \frac{\partial \theta_{s1}}{\partial \tau} - \frac{\alpha_{s1}}{\alpha_{s2}} \bar{\Delta} \theta_{s1} = 0 \\ \frac{\partial \theta_{s2}}{\partial \tau} - \bar{\Delta} \theta_{s2} = 0 \\ \frac{\partial \theta_{s2}}{\partial x} \Big|_{x=0^+} = \frac{\partial \theta_{s1}}{\partial x} \Big|_{x=0^-} \end{cases} \quad (3.2)$$

3.2 Spatial and temporal discretizations

Both solid models are implemented in separate C++ files within the framework provided by the Adaptive Mesh Refinement (AMR) library SAMURAI [8], which provides an efficient and flexible spatial discretization framework through adaptive structured grids. Mesh adaptation is disabled to ensure a constant position and a constant number of coupling variables during the simulation, guaranteeing the temporal coherence of their polynomial representations. Properly handling mesh adaptation along with polynomial-in-time predictions of the coupling variables will be tackled in future work. The left domain is discretized at level 2 (*i.e.* with a 4×20 Cartesian grid) and the right domain at level 3 (8×8 Cartesian grid). The equations are discretized with a cell-centered finite-volume approach, using a second-order centered scheme for the Laplacian.

Standard Neumann-Dirichlet coupling conditions are used at the interface: the left solid prescribes its surface temperature profile as a Dirichlet condition to the right solid, while the latter provides a Neumann condition (heat flux) to the left solid. This conditions are suited for $\alpha_{s1}/\alpha_{s2} < 1$ and should otherwise be reversed. The coupling variable vector u_i is then composed of the discrete heat flux values interpolated on the interface grid of the left solver, and the discrete temperature values interpolated on the interface grid of the right solver.

For temporal integration, both solvers use the implicit Butcher tableau of the 5th-order IMEX additive Runge-Kutta scheme ARK25(4)8L[2]SA₂ from [10], with a single substep per coupling step. The linear algebra is handled by the PETSC library [9].




Figure 2: Convergence analysis of the multistep coupling scheme

3.3 Convergence analysis

The theoretical convergence order of the multistep coupling scheme are now verified by part of the heat transfer transient with a constant coupling step size Δt , and studying the convergence of the coupled result as $\Delta \tau$ is decreased. The first few coupling steps are however inherently of low order, since not enough past data is available. Hence, their error contribution is mitigated by starting the simulation with $\Delta \tau_0 = \Delta \tau / 10^4$, then increasing it by a factor of 1.5 at each new coupling time step (thus 22 initial steps before reaching the target $\Delta \tau$). Simulations use a ratio $\alpha_{s1}/\alpha_{s2} = 1$ and are performed up to the dimensionless time $\tau_f = 4 \cdot 10^{-2}$, roughly corresponding to the middle of the transient. For the implicit variant, the fixed-point iterations are stopped once the relative error between two consecutive iterates is below 10^{-10} . Figure 2 illustrates the convergence of the L^1 -error on all coupling variables at time τ_f relative to a very fine reference solution. The obtained curves confirm the expected orders of convergence.

The ratios between the error constants of the explicit and implicit schemes provided in [7] is well verified by comparing the errors for coupling time step $\Delta \tau = 10^{-4}$ which is, for all convergence graphs, in the asymptotic regime. Compared to the theoretical error ratios given in brackets, the ratios observed are, for order 1 to 5: 1.0 (1.0), 5.0 (5.0), 9.8 (9.0), 15.3 (13.2), 19.0 (17.6).

In order to compare the efficiency of the explicit and implicit multistep coupling schemes, work-precision diagrams can be considered. The cost is evaluated in terms of iterations, *i.e.* the number of calls to both monophysics solver. Figure 3 shows the obtained diagrams. It can be seen that the higher-order schemes are extremely beneficial when looking for an accurate solution.

In the implicit variant, the iterative fixed-point procedure used at each step logically leads to a higher cost per step compared to the explicit variant. In the same spirit as Adams predictor-corrector methods for ODEs [11, Sect. 5.4.2], one may expect a gain in accu-




Figure 3: Cost in terms of the relative L^1 error

racy thanks to the lower error constants of the implicit scheme. However, it turns out that this does not counterbalance the additional cost, unless seeking very high accuracy. This may be explained by the fact that our chosen coupled test case is not stiff, even though the choice $\alpha_{s1}/\alpha_{s2} = 1$ is typically expected to yield the stiffest Neumann-Dirichlet coupling. The main reason for the lack of stiffness is the coarse mesh resolution around the interface.

3.4 Time adaptation analysis

The error estimates from Section 2 may be used to dynamically adapt the coupling time step. To study the performance of the dynamic time adaptation, simulations are run with the explicit scheme with the extrapolation error estimate from Equation (2.3), and the implicit scheme with the embedded estimate. We use the previous test case, as well as a stiffer variant obtained by increasing the diffusivity ratio to $\alpha_{s1}/\alpha_{s2} = 10^3$, without reversing the coupling conditions. In both cases, we simulate up to $\tau_f = 0.04$ s.

Figure 4 presents the time step histories generated by the 5th-order time adaptive multistep coupling scheme for both cases. The explicit variant (denoted “exp. extrap.” in the legend) is compared against implicit coupling employing a higher-order embedded error estimate (“imp. emb.”). The error tolerances for time step adaptation are $\text{atol} = \text{rtol} = 10^{-3}$, and the fixed-point convergence tolerance for the implicit coupling is set to 10^{-4} . Although the implicit scheme consistently selects larger time steps than the explicit scheme for both diffusivity ratios, it is actually more costly in the non-stiff case, requiring 126 iterations in total, compared to 51 for the explicit variant.

The advantage of the implicit strategy however clearly emerges in the stiffer case: the final time is reached in 35 time steps (with no rejected steps), for a total of 243 iterations (approximately 7 iterations per step on average). By contrast, the explicit scheme requires 903 accepted steps, in addition to 209 rejected steps, resulting in a total cost of 1112 iterations. The implicit coupling scheme is able to use the same time




Figure 4: Time step histories for the adaptive 5-th order multistep coupling scheme

steps in both configurations, while the explicit variant is plagued by its stability limit, clearly identified by the plateau in coupling time step values, which has been verified to be almost independent of the error tolerance, as is typical for stability-bound schemes.

4 Conclusions

This work demonstrates the successful implementation of the explicit and implicit multistep coupling schemes within the CWIPI library, extending its capabilities from spatial interpolation to temporal coupling, thus constituting a complete space-time coupling library. The proposed approach combines several key advantages: it achieves high-order temporal accuracy, it maintains a minimally-intrusive partitioned framework, and it enables adaptive time stepping.

The conjugate heat transfer test case validates the implementation and demonstrates high-order convergence in time for both explicit and implicit variants. The ability to dynamically adapt the coupling time step based on generic error estimates opens the door to efficient coupled simulations with heterogeneous time scales.

Future work will focus on exploring more complex multiphysics applications, *e.g.*, fluid-structure interaction or fluid-solid conjugate heat transfer with combustion and radiative heat transfer. Other work will be centered on the improvement of the implicit coupling, by relying on convergence acceleration methods, *e.g.* quasi-Newton approaches, for the fixed-point problem.

Acknowledgments

We thank our NASA hosts for their kind hospitality. The main support of ONERA, Fondation de l’École polytechnique through Initiative HPC@Maths and the Exa-MA project (Methods and Algorithms for Exascale, PEPR NumPEX, Grant Number: ANR-22-EXNU-0002) is gratefully acknowledged.

References

- [1] E. Quémerais. “Cwipi: Coupling with interpolation parallel interface,” ONERA. [Online]. Available: <https://github.com/onera/cwipi>
- [2] G. Chourdakis et al., “Precice v2: A sustainable and user-friendly coupling library,” *Open Research Europe*, vol. 2, p. 51, 2022.
- [3] C. Farhat and M. Lesoinne, “Two efficient staggered algorithms for the serial and parallel solution of three-dimensional nonlinear transient aeroelastic problems,” *Computer methods in applied mechanics and engineering*, vol. 182, no. 3-4, pp. 499–515, 2000.
- [4] V. Kazemi-Kamyab, A. Van Zuijlen, and H. Bijl, “A high order time-accurate loosely-coupled solution algorithm for unsteady conjugate heat transfer problems,” *Computer Methods in Applied Mechanics and Engineering*, vol. 264, pp. 205–217, 2013.
- [5] S. Kang, A. Dener, A. Hamilton, H. Zhang, E. M. Constantinescu, and R. L. Jacob, “Multirate partitioned runge–kutta methods for coupled navier–stokes equations,” *Computers & Fluids*, vol. 264, p. 105964, 2023.
- [6] L. Francois and M. Massot, “Multistep interface coupling for high-order adaptive black-box multi-physics simulations,” in *10th edition of the International Conference on Computational Methods for Coupled Problems in Science and Engineering*, CIMNE, 2023.
- [7] A. E. Simon, L. François, and M. Massot, “High-order multistep coupling: Convergence, stability and pde application,” *Comptes Rendus. Mécanique*, vol. 353, no. G1, pp. 1159–1184, 2025.
- [8] T. Bellotti, L. Gouarin, B. Graille, and M. Massot, “Multiresolution-based mesh adaptation and error control for lattice boltzmann methods with applications to hyperbolic conservation laws,” *SIAM Journal on Scientific Computing*, vol. 44, no. 4, A2599–A2627, 2022.
- [9] S. Balay et al., *PETSc Web page*, <https://petsc.org/>, 2025. [Online]. Available: <https://petsc.org/>
- [10] C. A. Kennedy and M. H. Carpenter, “Higher-order additive runge–kutta schemes for ordinary differential equations,” *Applied numerical mathematics*, vol. 136, pp. 183–205, 2019.
- [11] U. M. Ascher and L. R. Petzold, *Computer methods for ordinary differential equations and differential-algebraic equations*. SIAM, 1998.

Transition and High-speed flows

Investigation of Hypersonic Boundary Layer Transition induced by Wall Protrusions and Freestream Disturbances

Vincenzo Romano Antoine Ghyoot Ata Onur Başkaya
von Karman Institute for Fluid Dynamics, Belgium
December 4, 2025

Abstract

Superficial irregularities and environmental disturbances can trigger the boundary layer (BL) laminar-to-turbulent transition, resulting in increased wall heat flux, which in turn accelerates the ablation phenomena. The strong coupling between BL transition with ambient fluctuations and ablated surface could be critical during atmospheric reentry. In this context, we studied the BL behavior of the canonical hypersonic blunt cone configuration when subjected to surface irregularities, i.e. cavity and step, and imposed freestream noise. Our numerical experiments were based on the H3 conditions, the hypersonic blowdown Mach 6 wind tunnel at the von Karman Institute. The effect of the surface imperfections have been investigated by using the Linear Stability Theory (LST) to compute the amplitude of the main modal instabilities (1st and 2nd modes). A Direct Numerical Simulation (DNS) using the Discontinuous Galerkin method, instead, was employed to gain understanding on the effect of the noise. Results shows that the cavity does not affect much the BL, whereas the step strongly amplifies the 2nd mode. The DNS, instead, shows good agreement with LST as long as the linear growth remains dominant.

Keywords: Hypersonics, Boundary Layer, Laminar-to-Turbulent Transition, Roughness, Noise

1 Introduction

Hypersonic boundary layer (BL) laminar-to-turbulent transition is a crucial phenomenon during any hypersonic flight. When it occurs, wall heat transfer and skin friction rises, increasing the drag of the spacecraft [1]. It also influences the flow separation location. This alters the maneuverability and the dynamic stability of the spacecraft, leading to major unsteady loads along the spacecraft's surface while inducing large vibrations to the hull and payload [2]. Naturally, the BL transition phenomenon has a strong influence on the design and performance of any high-speed vehicle. Accurate prediction of its occurrence and impact is therefore critical, but inherently difficult since it is highly sensitive to surface roughness and environmental noise [3]. In this regard, the current work analyzed the effect of surface imperfections and freestream noise over the canonical hypersonic blunt cone configuration considering the H3 conditions, the blowdown Mach 6 wind tunnel at von Karman Institute (VKI). The effect of surface roughness, particularly a cavity and a step along the cone wall, is investigated with the Linear Stability Theory (LST), which gives information on the behavior of the

main modal instabilities (1st and 2nd modes). The effect of the noise, instead, has been studied with a two-dimensional Direct Numerical Simulation (DNS) using the high-order Discontinuous Galerkin (DG) method, which shed light on the receptivity mechanism across the shock as well. In detail, the present report is organized as follows: section 2 describes the configuration and the conditions, as well as LST with the numerical method used for the base flow computations and how the DNS has been forced; section 3 shows the results of the stability analysis as well as of the noise injection. Finally, section 4 draws the conclusions and comments on possible future investigations.

2 Methods

This section details the blunt cone configuration and the freestream conditions considered, the LST, the employed numerical code, and how the perturbations are introduced in the DNS.

2.1 Configuration and conditions

The geometry used is the canonical hypersonic blunt cone. A nose radius of $R_N = 0.126$ mm, and a cone

angle of $\theta = 7$ deg with an axial length of $L_{\text{cone}} = 265$ mm have been considered. For the wall, the standard no-slip condition is applied along with an isothermal temperature $T_w = 295$ K. The inlet conditions are the one of the VKI H3 facility [4], shown in Table 1.

$\bar{M}_\infty [-]$	$\bar{p}_\infty [\text{Pa}]$	$\bar{T}_\infty [\text{K}]$	$Re_{\infty, \text{unit}} [m^{-1}]$
6	1358.69	60.975	18×10^6

Table 1: VKI H3 freestream conditions.

The air is assumed to behave as a calorically perfect gas. The dynamic viscosity is computed with Sutherland's law and the thermal conductivity is retrieved by the Prandtl number, which is kept constant at $Pr = 0.7368$. For the cavity and step configurations, a size of 0.98 mm has been investigated, located at 39 mm from the stagnation point along the cone axis. In the following, (x, y) are the axial and radial coordinate, whereas (ξ, η, ζ) are the streamwise, the wall normal and spanwise coordinate, respectively.

2.2 Linear Stability Theory (LST)

In the following, $\mathbf{Q} = (\rho, \rho\mathbf{v}, \rho E)$ is the state-vector of fluid-dynamic system (density, momentum, energy). The LST [5] considers writing the state vector as the sum of the steady-state laminar solution (base flow) $\bar{\mathbf{Q}}_b$, with the assumption that the perturbation is much smaller than the base flow (small perturbation):

$$\mathbf{Q} = \bar{\mathbf{Q}}_b + \mathbf{Q}' \quad \mathbf{Q}' \ll \bar{\mathbf{Q}}_b \quad (2.1)$$

Substituting Eq. 2.1 into Navier-Stokes (NS) equations, performing a Taylor expansion and neglecting the high-order terms, we get the linearized NS equations. Then, assuming that the base flow is locally parallel, i.e. $\bar{\mathbf{Q}}_b = \bar{\mathbf{Q}}_b(\eta)$, and a wave-like perturbation, the following stability equations are obtained:

$$\frac{\partial \mathbf{Q}'}{\partial t} = \mathbf{J} \mathbf{Q}' \quad \mathbf{Q}' = \hat{\mathbf{q}}(\eta) e^{i(\alpha\xi + \beta\zeta - \omega t)} \quad (2.2)$$

where $\mathbf{J} = \partial \text{NS} / \partial \mathbf{Q}|_{\bar{\mathbf{Q}}_b}$ is the Jacobian operator of the linearized NS operator. Therefore, performing a BL-LST analysis means solving the linear dynamics of the perturbations. Also, α and β are the streamwise and spanwise wavenumbers, respectively. Here, $\omega = 2\pi f$ is the angular frequency, where f is the frequency. Inserting \mathbf{Q}' in the linearized NS, the analysis leads to an eigenvalue problem for the Jacobian operator. To solve the spatial stability problem, f and β are specified and real-valued, and the α is the complex eigenvalue unknown. Instead, $\hat{\mathbf{Q}}(\eta)$ is the complex eigenfunction, which gives the mode shape through the BL. The nature of instability depends on the sign of the imaginary part of the streamwise wavenumber ($\Im(\alpha)$). The opposite of $\Im(\alpha)$ is referred to by *the streamwise growth rate*. It could be positive, negative, or null, corresponding to a damped, amplified, or neutral wave, respectively. The integration of $-\Im(\alpha)$ gives the well known N -factor, which is a measure of the instability amplitude:

$$N = - \int_{\xi_0}^{\xi} \Im(\alpha(\chi)) d\chi \quad (2.3)$$

where x_0 is the neutral point, so the first point where $\Im(\alpha) = 0$. As described above, LST is a linear formulation and thus does not provide information about where transition occurs, but it gives insights on the instability behavior. In detail, the LST has been run with the VKI in-house stability tool VESTA [6].

2.2.1 Base flow computation

The base flows for VESTA simulations have been computed with the hypersonic code US3D [7], developed by the University of Minnesota. US3D is a cell-centered Finite Volume solver, and it has been run using a 2nd order MUSCL (Monotonic Upstream-Centered Scheme for Conservation Laws) scheme. To perform the US3D simulation the mesh has been aligned to the shock thanks to the HEROES VKI-internal library [8]. Moreover, for the smooth wall configuration, the US3D base flow has also been compared with the one obtained by the ArgoDG code. ArgoDG is a high-order solver based on DG spatial discretization [9]. In this work, p-adaption has been used to force a zero-interpolation order for the elements encompassing the shock to preserve the positivity of the scheme [10], and to have higher resolution inside the BL (p=4). To alleviate the decrease of accuracy in the shock region, mesh refinement tailored to the shock position is performed.

2.3 Receptivity of acoustic waves

The BL receptivity to tunnel-like noise has also been investigated. Two kinds of free-stream disturbances, fast acoustic waves and slow acoustic waves, with a single frequency $f = 450$ kHz. These perturbations were modeled as time dependent inlet boundary conditions where the simulation has been restarted from the steady-state solution. In case of \pm signs in relevant component of the following expressions, the top sign represents the fast acoustic wave and the bottom sign represents the slow acoustic waves. Moreover, the subscript (∞) denotes far-field quantities, while the notations ($\bar{\cdot}$) and (\cdot') represent mean flow quantities and the perturbations, respectively. Then, the instantaneous quantities at the inlet boundary due to the disturbances at time t can be defined as

$$\mathbf{q}_\infty = \bar{\mathbf{q}}_\infty + \sum |\mathbf{q}'| e^{i[k_x(x-a) + k_y(y-b) - \omega t]}, \quad (2.4)$$

where \mathbf{q} is the vector of the non-conservative primitive variables $[p, u, v, T]$, and k_x, k_y are the streamwise and wall-normal component of the wave number vector \mathbf{k} . The perturbations are considered with an angle of incidence of zero such that $k_x = k_\infty$ and $k_y = 0$, where $k_\infty = \frac{M_\infty \omega}{\bar{u}_\infty (M_\infty \pm 1)}$ is the absolute wave number. The source of the perturbations is located far from the body at $(a, b) = (-L_{\text{cone}}, 0)$. The amplitudes of the perturbations, denoted as ($|\cdot'|$), are defined as follows:

$$\begin{aligned} |p'| &= \varepsilon \bar{\rho}_\infty \bar{u}_\infty^2, & |u'| &= \pm \varepsilon \bar{M}_\infty \bar{u}_\infty, \\ |T'| &= \varepsilon (\gamma - 1) \bar{M}_\infty^2 \bar{T}_\infty, & |v'| &= 0. \end{aligned} \quad (2.5)$$

In Eqs. 2.5, the amplitudes of the perturbations are characterized by the small parameter ε . The employed

amplitude is $\varepsilon = 1 \cdot 10^{-5}$, which ensured the linearity of the receptivity process.

3 Results

The laminar base flows are shown first, then the results of LST and finally the outcomes of the noise injection.

3.1 Laminar base flows

The base flows comparison of the three configurations simulated with US3D is shown in Fig. 1. The Mach contour is plotted in color whereas the numerical schlieren is plotted in grayscale.




Figure 1: US3D base flow of the three configurations.

In detail, from the above picture we can notice that the cavity introduces a weak shock which is then enhanced by the presence of the step. Fig. 2 shows, the code-to-code comparison between ArgoDG and US3D on the laminar base flow of the smooth configuration. Apart from minor differences attributed to the different numerical schemes employed, the results are in a good agreement between the solvers.




Figure 2: ArgoDG-US3D base flow comparison for the smooth wall configuration

3.2 Boundary layer stability results

The N -factor has been computed for both 1st and 2nd modes over all the frequency range dictated by the considered conditions. Fig. 3 and Fig. 4 show the envelope evolution and the contour plot of the N -factor, respectively. In detail, in the contour plot of the smooth wall configuration (Fig. 4, right), the white lines represent the neutral curves, i.e. the boundaries between the stable and unstable regions. Here, we can notice that $f = 450$ kHz is the most amplified frequency at the end of the cone: that is the reason why this frequency has been used to force the DNS for the noise investigation. Looking at the envelopes (Fig. 3), the considered cavity does not affect the instabilities since the envelopes match with the ones of the smooth wall. Probably the considered size is too small to affect the BL under this freestream conditions. In opposite, the step configuration amplifies the 2nd mode right after its location. For the 1st mode, we found out that the step stabilizes the flow. However, these results require further analysis, as the presence of the step disrupts the computation and reduces the amplitude of the 1st mode.




Figure 3: Comparison of the N -factor envelopes for the 1st and 2nd instability modes.

3.3 Noise injection

The impact of fast and slow acoustic waves on the development of the 2nd mode has been investigated using high-order DNS computations. The simulation has been performed on the smooth wall configuration for a total runtime of 1 ms. In what follows, harmonics will be denoted as $f_{(i,0)}$ where i corresponds to the level of the harmonic, $f_{(1,0)} = 450$ kHz being the frequency of the perturbations that are injected at the inlet. The amplitude of a perturbation in Eq. 2.5 at a given harmonic can therefore be designated as $|\cdot|_{(i,0)}^*$. In relevant cases, the superscript (*) represents the normalization by the maximum amplitude corresponding to the harmonic referred in the subscript. If no harmonic is indicated in the subscript, the normalization is based on the main harmonic $f_{(1,0)}$. Fig. 5 shows the amplitude of the pressure perturbations on the wall. First, one can observe destructive and constructive interactions at a




Figure 4: Contour plot of the N -factor in the frequency- ξ plane of the three configurations.

relatively constant amplitude between the two types of perturbations from the nose region up to $\xi = 165$ mm. Amplification of the perturbations amplitude related to the main harmonic $f_{(1,0)}$ can be observed at $\xi = 165$ mm. The first higher harmonic $f_{(2,0)} = 900$ kHz is generated at $\xi = 190$ mm, then followed by the excitation of the higher harmonics at $\xi = 215$ mm showing nonlinearities in the flow-field. The generation of higher harmonics is likely due to self-interactions of 2nd mode disturbances where first level interactions would generate the first higher harmonic [11]. Then, interactions between lower and high harmonics would lead to second level interactions that would generate even higher harmonics (i.e. $f_{(3,0)} = 1350$ kHz, $f_{(4,0)} = 1800$ kHz, $f_{(5,0)} = 2250$ kHz).




Figure 5: Amplification of the main harmonic $f_{(1,0)}$ and generation of the higher harmonics along the cone. The amplitude of the pressure fluctuations is considered.

The presence of the 2nd mode instability is confirmed by the well-known observed rope-like structures inside the BL, as well as strong wall pressure fluctuations. Both are shown in terms of T' in Fig. 6.




Figure 6: 2nd mode rope-like structures obtained from the temperature fluctuations along the cone.

Fig. 7 compares the normalized amplitudes of \mathbf{q} along the wall normal direction that are obtained by LST and by the DNS. Note that the profiles from LST were computed based on the baseflow provided by US3D and not ArgoDG, so the comparison may not be exact. Nevertheless, good agreements are obtained between the two results until $\xi = 210$ mm. Further downstream, strong nonlinearities are observed (see Fig. 7, right column). Indeed, nonlinearities cannot be captured by LST, which instead can be considered by non-linear stability methods, such as Nonlinear Parabolized Stability Equations (NPSE) [12].

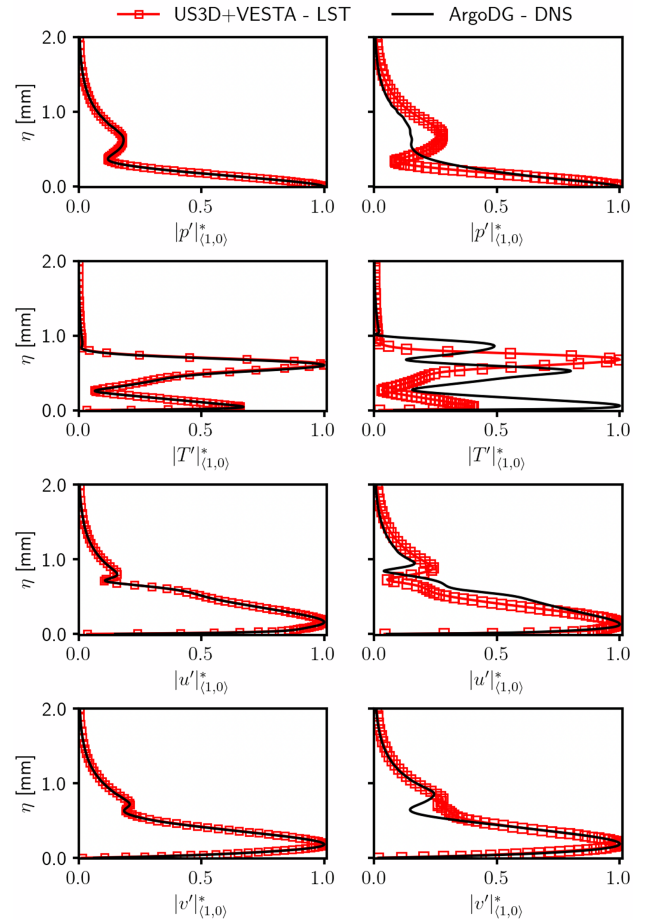


Figure 7: Comparison of fluctuations between LST and DNS at the stations $x = 210$ mm (left column) and $x = 260$ mm (right column).

4 Conclusions and future works

The current report had a double goal: investigate the effect of the roughness and freestream noise on hypersonic BL transition of a blunt cone under the VKI H3 conditions.

- For the former, the effect of a cavity and a step of size 0.98 mm along the cone wall have been investigated on the main modal instabilities (1st and 2nd modes) with LST. Results show that the considered cavity does not significantly affect the BL, most probably its size is too small, whereas the step highly amplifies the 2nd mode instability.
- For the latter, a DNS of a smooth wall cone has been forced with the slow and fast acoustic waves at a frequency given by the previous LST analysis. Results capture the rope-like structures, as well as the wall pressure fluctuations, typical of the 2nd mode instability. Moreover, the DNS results are aligned with those of LST until nonlinearities start to be dominant.

However, future investigations could be done on both sides. Indeed, LST does not give any answer on the non-modal instabilities (streaks) that could rise from the surface imperfections. The next phase of this DNS study, instead, will be to extend it to a 3D configuration to analyze the instability growth until breakdown to turbulence, as well as future analysis should be done on the choice of wave amplitudes.

Acknowledgments

All the authors gratefully thank Dr. Pierre Schrooyen, Dr. David Henneaux, Prof. Koen Hillewaert, and Prof. Thierry Magin for their guidance during the Modeling Summer Visit. Yet, the authors also acknowledge Dr. Dirk Ekelschot for the fruitful discussions during the stay at NASA.

V. Romano intends also to thank his supervisors, Prof. Olivier Chazot, Prof. Thierry Magin, and Prof. Anabel del Val, for all the advice and suggestions, as well as Dr. George Bellas for the discussion on the US3D code. V. Romano is supported by the Fonds de la Recherche Scientifique (F.R.S.-FNRS) of Belgium, under the FRIA grant with reference FC 58087.

A. Ghyoot also expresses his thanks to Ludovic Taguema, from Cenaero, for his advices and suggestions regarding the mesh generation and the data visualizations. The research of A. Ghyoot is part of the Space4ReLaunch project, which is supported by the SPW Economie Emploi Recherche of the Walloon Region, under grant agreement no. 2210181. Computational resources have been provided by the Consortium des Équipements de Calcul Intensif (CÉCI).

A. O. Başkaya sincerely thanks Prof. Stefan Hickel for his guidance and acknowledges the support of Fonds voor Wetenschappelijk Onderzoek Vlaanderen (FWO) through the short-travel grant no. K220125N.

References

- [1] AV Fedorov. “Laminar Turbulent Transition in a Hypersonic Boundary Layer”. In: *Contract 61708* (1997), 96–W0196.
- [2] Tony C Lin. “Influence of laminar boundary-layer transition on entry vehicle designs”. In: *Journal of Spacecraft and Rockets* 45.2 (2008), pp. 165–175.
- [3] Steven P Schneider. “Hypersonic laminar-turbulent transition on circular cones and scramjet forebodies”. In: *Progress in Aerospace Sciences* 40.1-2 (2004), pp. 1–50.
- [4] Guillaume Grossir, Davide Masutti, and Olivier Chazot. “Flow characterization and boundary layer transition studies in VKI hypersonic facilities”. In: *53rd AIAA Aerospace Sciences Meeting*. 2015, p. 0578.
- [5] Leslie M Mack. *Boundary-layer stability theory*. Jet Propulsion Laboratory, 1969.
- [6] Fabio Pinna. “VESTA toolkit: a software to compute transition and stability of boundary layers”. In: *43rd AIAA Fluid Dynamics Conference*. 2013, p. 2616.
- [7] Graham V Candler et al. “Development of the US3D code for advanced compressible and reacting flow simulations”. In: *53rd AIAA Aerospace Sciences Meeting*. 2015, p. 1893.
- [8] Michele Capriati, Domenico Massari, and Pierre Schrooyen. “Development of HEROES: Hybrid Mesh Adapter for Hypersonic”. In: *3rd International Conference on Flight Vehicles, Aerothermodynamics and Re-entry (FAR)*. 2025.
- [9] Koen Hillewaert. “Development of the discontinuous Galerkin method for high-resolution, large scale CFD and acoustics in industrial geometries”. Promotor: Jean-François Remacle. PhD thesis. Université catholique de Louvain, 2013.
- [10] Schrooyen Pierre et al. “Towards high-order discretization method for hypersonic flow”. In: *Proceedings of HiSST: 2nd International Conference on High-Speed Vehicle Science & Technology*. 2022.
- [11] Christoph Hader and Hermann Fasel. “Nonlinear stages of transition and breakdown to turbulence: Direct numerical simulations”. In: *Computation and Analysis of Turbulent Flows – Turbulence and Transition in Supersonic and Hypersonic Flows*. Ed. by Johan Larsson and Xiaolin Zhong. Academic Press, 2026, pp. 283–383.
- [12] Ludovico Zanus and Fabio Pinna. “Stability Analysis of Hypersonic Flows in Local Thermodynamic Equilibrium Conditions by Means of Non-linear PSE”. In: June 2018.