

# SANDIA REPORT

SAND2021-0718

Printed January 25, 2021



Sandia  
National  
Laboratories

# Towards Predictive Plasma Science and Engineering through Revolutionary Multi-Scale Algorithms and Models, Final Report

George R. Laity, Allen C. Robinson, Michael E. Cuneo,  
M. Kathleen Alam, Kristian R. C. Beckwith, Nichelle L. Bennett, Matthew T. Bettencourt,  
Stephen D. Bond, Kyle Cochrane, Louise Criscenti, Eric C. Cyr, Karen De Zetter,  
Richard R. Drake, Evstati G. Evstatiev, Andrew S. Fierro, Thomas A. Gardiner,  
Forrest W. Glines, Ronald S. Goeke, Nathaniel D. Hamlin, Russell Hooper, Jason Koski,  
J. Matthew Lane, Steven R. Larson, Kevin Leung, Duncan A. McGregor, Philip R. Miller,  
Sean M. Miller, Susan J. Ossareh, Edward G. Phillips, Nicholas A. Roberds,  
Charles E. Rose, John N. Shadid, Sidney Shields, Sean C. Simpson, David Sirajuddin,  
Thomas M. Smith, M. Scot Swan, Aidan P. Thompson, Julien G. Tranchida

Sandia National Laboratories  
P.O. Box 5800, MS-1195  
Albuquerque, NM 87185  
grlaity@sandia.gov

Dale R. Welch, Thomas C. Genoni, Carsten Thoma, Alex Russell, Eric D. Watson, David V. Rose  
Voss Scientific, 418 Washington St. SE,  
Albuquerque, NM 87108

John Williams, Azer Yalin  
Department of Mechanical Engineering, Colorado State University,  
2545 Research Blvd,  
Fort Collins, CO 80526

Ryan D. McBride, Trevor J. Smith  
Department of Nuclear Engineering & Radiological Science, University of Michigan,  
2355 Bonisteel Blvd,  
Ann Arbor, MI 48109

Prepared by  
Sandia National Laboratories  
Albuquerque, New Mexico 87185  
Livermore, California 94550

Issued by Sandia National Laboratories, operated for the United States Department of Energy by National Technology & Engineering Solutions of Sandia, LLC.

**NOTICE:** This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government, nor any agency thereof, nor any of their employees, nor any of their contractors, subcontractors, or their employees, make any warranty, express or implied, or assume any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represent that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government, any agency thereof, or any of their contractors or subcontractors. The views and opinions expressed herein do not necessarily state or reflect those of the United States Government, any agency thereof, or any of their contractors.

Printed in the United States of America. This report has been reproduced directly from the best available copy.

Available to DOE and DOE contractors from

U.S. Department of Energy  
Office of Scientific and Technical Information  
P.O. Box 62  
Oak Ridge, TN 37831

Telephone: (865) 576-8401  
Facsimile: (865) 576-5728  
E-Mail: reports@osti.gov  
Online ordering: <http://www.osti.gov/scitech>

Available to the public from

U.S. Department of Commerce  
National Technical Information Service  
5301 Shawnee Road  
Alexandria, VA 22312

Telephone: (800) 553-6847  
Facsimile: (703) 605-6900  
E-Mail: orders@ntis.gov  
Online order: <https://classic.ntis.gov/help/order-methods>



## ABSTRACT

This report describes the high-level accomplishments from the Plasma Science and Engineering Grand Challenge LDRD at Sandia National Laboratories. The Laboratory has a need to demonstrate predictive capabilities to model plasma phenomena in order to rapidly accelerate engineering development in several mission areas. The purpose of this Grand Challenge LDRD was to advance the fundamental models, methods, and algorithms along with supporting electrode science foundation to enable a revolutionary shift towards predictive plasma engineering design principles. This project integrated the SNL knowledge base in computer science, plasma physics, materials science, applied mathematics, and relevant application engineering to establish new cross-laboratory collaborations on these topics.

As an initial exemplar, this project focused efforts on improving multi-scale modeling capabilities that are utilized to predict the electrical power delivery on large-scale pulsed power accelerators. Specifically, this LDRD was structured into three primary research thrusts that, when integrated, enable complex simulations of these devices: (1) the exploration of multi-scale models describing the desorption of contaminants from pulsed power electrodes, (2) the development of improved algorithms and code technologies to treat the multi-physics phenomena required to predict device performance, and (3) the creation of a rigorous verification and validation infrastructure to evaluate the codes and models across a range of challenge problems.

These components were integrated into initial demonstrations of the largest simulations of multi-level vacuum power flow completed to-date, executed on the leading HPC computing machines available in the NNSA complex today. These preliminary studies indicate relevant pulsed power engineering design simulations can now be completed in (of order) several days, a significant improvement over pre-LDRD levels of performance.

## ACKNOWLEDGMENTS

This work was supported by the Laboratory Directed Research and Development (LDRD) Program at Sandia National Laboratories with project title, "Towards Predictive Plasma Science and Engineering through Revolutionary Multi-Scale Algorithms and Models," and project number, 209240. We gratefully acknowledge the Radiation, Electrical, and High Energy Density Science (REHEDS) leadership, the Computing and Information Sciences (CIS) leadership, and the Nuclear Deterrence (ND) leadership for strongly sponsoring this work. We would like to especially thank Keith Matzen and Daniel Sinars for serving as Director Champions for this work, as well as Scott Collis and Charles Barbour for strong institutional support.

We thoroughly enjoyed the strong interactions with our External Advisory Board (John Luginsland, Luis Chacon, Chris Clouse, Brendan Godfrey, Thomas Mehlhorn, James Peery, Charles Seyler, Gregory Van Dyk, John Verboncoeur), who provided critical recommendations to this work, and thank them for their excellent service over the past several years. We also recognize the critical contributions from Richard Macklin and Jeff Dorman for enabling productive team coordination throughout the project. We are grateful for the strong institutional support from the LDRD office, notably Greg Frye-Mason and Leigh Cunningham, that extended for the duration of this project.

The development of the initial version of the Code Comparison Infrastructure (CCI) software environment was supported by the NNSA Advanced Scientific Computing (ASC) program.

We thank Cristiana Di Valentin for her generous help with the DFTB+ code and parameters; Michael Desjarlais, Anastasia Ilgen, and Alan Wright for useful materials science interactions; Mark Hess, Matthew Hopkins, Thomas Hughes, Chris Jennings, and Benjamin Yee provided useful insights into exemplar physics problems; Adam Blake, Keith Cartwright, Sidafa Conde, Marta D'Elia, Chris Moore, Roger Pawlowski, Timothy Pointon, and Nathan Roberts all provided important contributions to ongoing code developments; William Rider provided helpful guidance and suggestions during the launch of this project; Bryan Hughes provided support to set up the GitLab continuous integration environment; Kevin Lind, Chris Mostrom, and Don Voss provided additional support at Voss Scientific; Michael Lee was a key academic parter at Northern Arizona University; Adriana Canavan, Sandra Guthrie, Pauline Head, Tere McCollum, Steve Nickerson, Lauren Shreve, and Keri Trujillo provided critical administrative support.

Sandia National Laboratories is a multi-mission laboratory managed and operated by National Technology and Engineering Solutions of Sandia LLC, a wholly owned subsidiary of Honeywell International Inc. for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525. This report describes objective technical results and analysis. Any subjective views or opinions that might be expressed in the report do not necessarily represent the views of the U.S. Department of Energy or the U.S. Government.

# CONTENTS

<b>Executive Summary</b>	<b>7</b>
<b>1. Introduction</b>	<b>11</b>
<b>2. Desorption Physics</b>	<b>14</b>
2.1. Theoretical Thermal Desorption Physics .....	14
2.1.1. Unpublished Results on Hydrogen Trapping/Diffusion .....	14
2.1.2. MD studies .....	15
2.1.3. Quasi-Equilibrium Water Desorption Barriers .....	19
2.1.4. GCMC studies of water on surfaces and in pores .....	20
2.1.5. H <sub>2</sub> O from reaction between surface oxide and hydrogen inside steel .....	21
2.1.6. Electrode Modeling in ALEGRA .....	22
2.1.7. Technical Gaps and Future Directions .....	22
2.2. Desorption Experiments and Surface Characterization .....	24
2.2.1. Slow Temperature Programmed Desorption of Hydrogen Fired 304L Coupons	25
2.2.2. 100 - 20,000 °C/s Programmed Desorption of Untreated 304L Stainless Steel Coupons .....	26
2.2.3. X-Ray Photoelectron Spectroscopy of 304L Stainless Steel .....	27
2.2.4. Grazing Angle Infrared Spectroscopy of Mineral Oil on Stainless Steel at Low Pressure .....	28
2.2.5. Electrochemical Surface Area Measurements .....	29
2.2.6. Method and Application to Multi-Peak Fitting and Residual Errors of Temperature Programmed Desorption Spectra .....	30
2.2.7. Ultra-Fast Mass Spectroscopy .....	34
2.2.8. Evaluation of a Potential High Speed Mass Spectrometer, "The E-Vader" ..	35
2.2.9. Technical Gaps and Future Directions .....	39
2.2.10. Concluding Remarks .....	40
2.3. Experimental Collaboration with University of Michigan .....	40
<b>3. Engineering Code Development</b>	<b>42</b>
3.1. EMPIRE .....	42
3.1.1. Background .....	43
3.1.2. EMPIRE-PIC .....	43
3.1.3. EMPIRE-FLUID .....	55
3.1.4. EMPIRE-HYBRID .....	79
3.1.5. EMPIRE Advancements Summary .....	88
3.2. CHICAGO Advances to Support Z Pulsed Power Simulation .....	89
3.2.1. CHICAGO Simulations of Z Power Flow including Electrode Physics .....	90

3.2.2. A Fast Hybrid Particle-in-Cell Technique for the Simulation of Pulsed Power Accelerators .....	92
3.2.3. Sensitivity of Current Loss to Water Inventory .....	96
3.2.4. Spherically Expanding Plasma Verification .....	99
3.2.5. CHICAGO Parallelization Strategies – Trilinos/Kokkos Integration .....	102
3.2.6. Implicit Highly-Coupled Single-Ion Hall-MHD Formulation for Hybrid ...	107
3.2.7. Miscellaneous Implicit and Adaptive Particle Management Advances .....	109
<b>4. Model Testing and Evaluation</b>	<b>110</b>
4.1. Reproducible Science Environment .....	110
4.2. Verification Studies .....	112
4.2.1. Thermalization Study .....	112
4.2.2. Neutral Expanding Slab, Cylinder, and Sphere Test Problems .....	115
4.2.3. Surface Joule Heating Modeling Analysis .....	120
4.2.4. Two-Stream Instability .....	120
4.2.5. PIC Statistical Noise Analysis .....	127
4.2.6. Vacuum Fields and Particle Dynamics for MITLs .....	128
4.3. Power Flow Analysis Exemplars .....	129
4.3.1. Planar MITL PIC UQ Study .....	130
4.3.2. Half-O-Lute Study .....	132
4.3.3. Power Flow 18a .....	138
4.3.4. Power Flow 18a on the Sierra HPC Resource .....	142
4.3.5. Power Flow Physics Advancements in CHICAGO .....	149
<b>5. Conclusions and Recommendations</b>	<b>158</b>
<b>References</b>	<b>161</b>
<b>Appendices</b>	<b>173</b>
<b>A. External Publications Generated by this LDRD Project</b>	<b>173</b>

## EXECUTIVE SUMMARY

This report describes the high-level accomplishments from the Plasma Science an Engineering Grand Challenge LDRD at Sandia National Laboratories, a significant cross-laboratory initiative executed over the past several years (FY18-20). Plasmas are commonly found in a variety of natural and laboratory environments, and consequently Sandia encounters plasma phenomena in several mission areas. Examples include (i) pulsed power accelerator technology, (ii) breakdown phenomena encountered in high voltage devices and environments, (iii) x-ray radiation interaction with materials and related effects, (iv) high power RF diode engineering, (v) plasmas that form around spacecraft and other vehicles during atmospheric re-entry, and (vi) high energy density science. The Laboratory has a need to demonstrate predictive capabilities to model plasma phenomena in order to rapidly accelerate several engineering developments in these mission areas. Depending on the specific application, understanding critical phenomena could require spanning a large range of plasma parameters across the kinetic and fluid regimes. In addition, many of these applications require understanding the underlying electrode physics and how those materials and surfaces influence the plasma phenomena.

The purpose of this Grand Challenge LDRD was to advance the fundamental models, methods, and algorithms along with supporting electrode science foundation needed to enable a revolutionary shift towards predictive plasma engineering design principles. This project integrated the SNL knowledge base in computer science, plasma physics, materials science, applied mathematics, and relevant application engineering to establish new cross-laboratory collaborations in these topics. We leveraged multiple code technologies to rapidly develop numerical approaches, grounded through strong coupling to verification and validation techniques. We believe the output from this project has significantly advanced SNL capabilities in high performance computing, multi-physics methods and algorithms, theoretical and experimental electrode science, science-based verification and validation methodology, and computational plasma physics applied to a representative application (e.g., pulsed power).

As an initial exemplar, this project focused efforts on improving multi-scale modeling capabilities that are utilized to predict the electrical power delivery on large-scale pulsed power accelerators. As an electromagnetic wave propagates through the magnetically insulated transmission line (MITL) towards the load, the large electrical current heats the conductors to hundreds of degrees, desorbs material from the conductor surfaces, and generates a dense surface plasma. Through a variety of physical processes, which depend on the surface plasma conditions, magnetic field geometry, and even the downstream load dynamics, a significant portion of the total available current can be lost to particle flows across the anode-cathode (A-K) gap, rather than being delivered to the load. Our lack of understanding is significantly influenced by the difficulty to predictively simulate the complex electrode physics and multi-scale plasma phenomena that exist in the transmission line structures found in these extreme laboratory environments. As a result,

this LDRD was structured into three primary research thrusts that, when integrated, enable these complex simulations: (1) the exploration of multi-scale models describing the desorption of contaminants from pulsed power electrodes, (2) the development of improved algorithms and code technologies to treat the multi-physics phenomena required to predict device performance, and (3) the creation of a rigorous verification and validation infrastructure to evaluate the codes and models across a range of challenge problems.

The rapid desorption of contaminants (e.g., water, hydrocarbons, oxides, trapped gases in the bulk) provide the inventory that fuels undesired plasma generation in pulsed power devices. Historical attempts to model vacuum power flow in these devices often excludes the effects of electrode surface condition, or at best includes simple rate-based desorption models for releasing adsorbed contaminants from the interface. A major goal of the present effort was to explore the underlying electrode physics models in more detail, to improve confidence in extrapolating these models to larger pulsed power accelerators. We applied several molecular modeling methods including density functional theory (DFT) electronic structure calculations, force-field-based molecular dynamics (MD), and Grand Canonical Monte Carlo (GCMC) simulations; these methods were used to explore water and hydrocarbon desorption dynamics as well as estimate initial water inventory on oxide surfaces in relevant vacuum conditions. This work represents the first significant studies applying molecular modeling methods to understand pulsed power electrode physics in the literature. The ALEGRA multi-physics code was applied to explore the extent of electrode heating expected on present and future pulsed power accelerators. These electrode modeling studies were also coupled to several light-lab experiments that were used to determine the initial adsorbate inventories and subsequent desorption rates under fast heating rates. Various techniques were applied including x-ray photo-electron spectroscopy (XPS), Fourier transform infrared (FTIR) spectroscopy, electro-chemical surface area characterization, and temperature programmed desorption (TPD) methods. The experimental methods for performing fast TPD were improved, automated analysis routines were developed to extract model parameters from the TPD experimental data, and initial comparisons to model predictions were performed. Several electrode desorption models were significantly exercised through this process, and ultimately improve confidence in the boundary conditions of inflow material used by the multi-physics plasma simulations discussed in this report.

The second major goal of this project was to improve the fidelity and performance (i.e., efficiency) of multi-physics code technologies that are used to simulate the complex plasma physics observed in pulsed power devices. These efforts are motivated by the long-term need to simulate next-generation pulsed power accelerators with up to a factor of 10x higher electrical energy delivered to targets, which could pose significant computing requirements. Most improvements in this area were targeted for the EMPIRE and/or CHICAGO multi-physics research codes, although a significant amount of verification work was completed in other codes as well. EMPIRE is a next-generation application designed to be highly performant for simulation of plasmas in complex geometries, but prior to this project did not include several physics models required to predict plasma behavior observed in high-current transmission lines. Specific physics models developed and implemented in EMPIRE during this project include (i) multi-level transmission line coupling to the simulation domain, (ii) electrode surface heating and contaminant desorption models, (iii) improvements to implicit algorithms designed to increase performance, and (iv) relativistic fluid representations of key physics models required for power

flow. In addition, a hybrid kinetic-fluid method based on the  $\delta f$  approach was extensively studied and is positioned for future implementation. The CHICAGO research code already included the requisite physics models prior to this project, so efforts were concentrated on improving code stability and performance for large simulation problems of interest. A fast hybrid particle-in-cell technique was explored in CHICAGO, with key performance metrics studied when transitioning between kinetic, multi-fluid, and hybrid techniques during representative simulations of pulsed power devices. Several methods for improving parallelization in CHICAGO were explored including integration with Trilinos/Kokkos and other OpenMP methods; these studies resulted in a significant improvement in computational efficiency. Overall both code technologies are now well-positioned to tackle the largest class of power flow simulations attempted to-date.

The third goal of this project was to create a verification and validation framework to rigorously test and evaluate the various models and codes exercised in this effort. A shared software environment was developed that enables an improved level of understanding of the relevant plasma physics through a disciplined approach of analyst-driven simulation quality checks (e.g. the Code Comparison Infrastructure). This environment consists of a standardized method to deliver access to several code capabilities, as well as utilities to build test problems stored and maintained in a common repository. A variety of test problems were developed across a range of complexities: (i) verification tests consisting of idealized plasma physics problems, (ii) simplified 2D MITL problems designed to test electrode physics models, (iii) moderate complexity multi-level 3D MITL problems for benchmarking code performance, and (iv) large-scale 3D MITL challenge problems representative of actual engineering design calculations. These calculations were completed across a range of high-performance computing (HPC) resources including small-scale shared resources, institutional clusters, and NNSA leadership-class machines. A key deliverable of this project was the initial prototype of an end-to-end analysis pipeline capable of directly integrating multi-scale electrode desorption data into representative MITL physics simulations; this infrastructure can be exercised to provide uncertainty quantification (UQ) metrics of pulsed power design calculations.

Ultimately this GC LDRD project explored a variety of concurrent research activities: improved electrode physics models, cross-code comparison tools, multi-level circuit coupling, HPC scalability, improved algorithms and solvers, mesh analysis routines, UQ infrastructure, and MITL physics analysis. These components were integrated into initial demonstrations of the largest simulations of multi-level power flow completed to-date (e.g., “Power Flow 18a”) using complementary code technologies; these initial studies were completed on the leading HPC computing machines available in the NNSA complex today. These preliminary studies indicate relevant pulsed power engineering design simulations can now be completed in (of order) several days, a effective improvement over pre-LDRD levels of wall-clock performance by at least a factor 30-40x (due to using a combination of PIC and/or multi-fluid approaches).



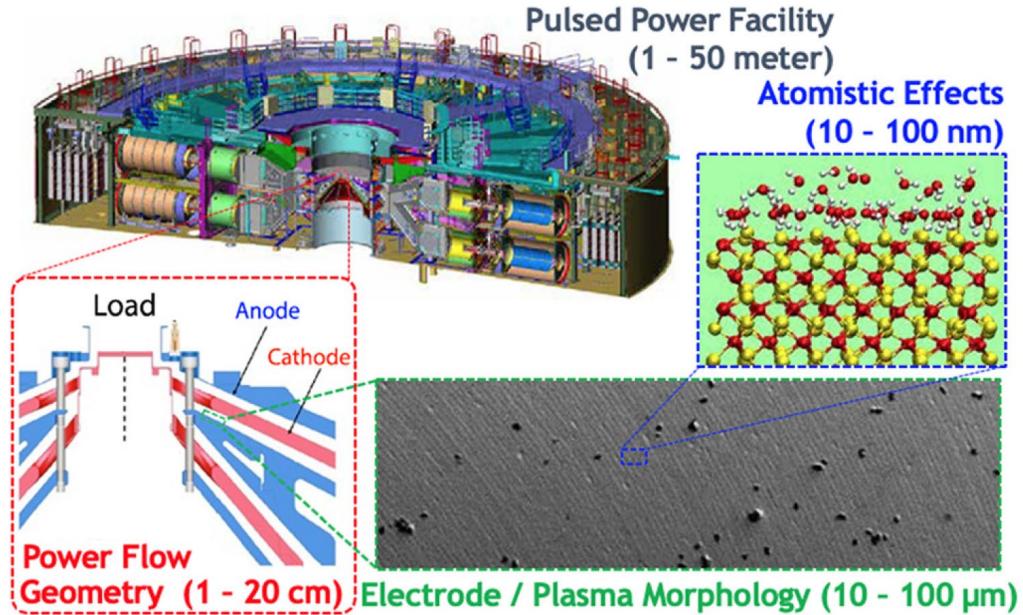
## 1. INTRODUCTION

Plasmas are commonly found in a variety of natural and laboratory environments, and consequently Sandia encounters plasma phenomena in several mission areas. Examples include (i) pulsed power accelerator technology, (ii) breakdown phenomena encountered in high voltage devices and environments, (iii) x-ray radiation interaction with materials and related effects, (iv) high power RF diode engineering, (v) plasmas that form around spacecraft and other vehicles during atmospheric re-entry, and (vi) high energy density science. The Laboratory has a need to demonstrate predictive capabilities to model plasma phenomena in order to rapidly accelerate several engineering developments in these mission areas. Depending on the specific application, understanding critical phenomena could require spanning a large range of plasma parameters across the kinetic and fluid regimes. In addition, many of these applications require understanding the underlying electrode physics and how those materials and surfaces influence the plasma phenomena.

This Grand Challenge LDRD (*Towards Predictive Plasma Science and Engineering through Revolutionary Multi-Scale Algorithms and Models*) was launched to directly address several high-risk science and technology gaps that we believe could have a significant impact on future plasma modeling capabilities at the Laboratory. An over-arching goal of this initiative was to bring together several disparate groups at the Laboratory, and transform those relationships into a more unified SNL community that is positioned to impact future mission opportunities. As a result, this LDRD project integrated the SNL knowledge base in computer science, plasma physics, materials science, applied mathematics, and relevant application engineering to establish new cross-laboratory collaborations in these topics. Given the large mission space for plasma engineering at the Laboratory, we elected to focus our initial efforts into a demonstration for one exemplar mission opportunity: pulsed power accelerator design.

Pulsed power accelerators are used to for a variety of NNSA missions including nuclear survivability, high energy density science, inertial confinement fusion, dynamic materials physics, laboratory astrophysics, and other related weapon assessment science [117]. This LDRD project focused efforts on improving multi-scale modeling capabilities that are utilized to predict the electrical power delivery on the largest scale of pulsed power accelerators. Specifically, as an electromagnetic wave propagates through the magnetically insulated transmission line (MITL) towards the load, the large electrical current heats the conductors to hundreds of degrees, desorbs material from the conductor surfaces, and generates a dense surface plasma. Through a variety of physical processes, which depend on the surface plasma conditions, magnetic field geometry, and even the downstream load dynamics, a significant portion of the total available current can be lost to particle flows across the anode-cathode (A-K) gap, rather than being delivered to the load [42]. Ultimately the atomistic-scale physics responsible for electrode contaminant desorption can have a strong influence on very-large-scale pulsed power accelerator performance [30]; this is

inherently a multi-physics, multi-scale technical challenge (see Figure 1-1 as one example observed today at the Z Pulsed Power Facility [117]).



**Figure 1-1. The spatial scale required to understand Z range from meters to nano-meters. Transmission lines models are sufficient at the largest scales but detailed power flow plasma modeling effects in the convolute arising from contaminants interacting at the meso and atomic scales are required. [117, Figure 16]**

Our lack of understanding is significantly influenced by the difficulty to predictively simulate the complex electrode physics and multi-scale plasma phenomena that exist in the transmission line structures found in these extreme laboratory environments. Plasma modeling codes have historically not been able to simultaneously treat the high and low plasma density regimes due to the computational resources and time-scales required. Instead, systems with low-density plasmas, charge distributions, and high electric field stresses have been treated kinetically [103], while magnetically-driven and charge neutral systems with high-density, collisional plasmas have been treated magneto-hydrodynamically [64, 98, 96]. A major scientific goal is to move closer towards robust capabilities that can bridge these regimes, which could be applicable to a new class of pulsed power design challenges.

As a result, this Grand Challenge LDRD was structured into three primary research thrusts that, when integrated, enable these complex simulations: (1) the exploration of multi-scale models describing the desorption of contaminants from pulsed power electrodes, (2) the development of improved algorithms and code technologies to treat the multi-physics phenomena required to predict device performance, and (3) the creation of a rigorous verification and validation infrastructure to evaluate the codes and models across a range of challenge problems.

The following project report is arranged as follows:

1. Chapter 2 describes the exploration of multi-scale desorption models, electrode conditions, and related experimental characterization;
2. Chapter 3 describes progress in developing improved numerical models and techniques in the EMPIRE and CHICAGO multi-physics codes;
3. Chapter 4 describes our approach to collaboration, cross-code model testing, and the development and sequence of exemplar calculations;

The goal of this report is to provide summaries of key accomplishments of this Grand Challenge LDRD; a large number of references are included, and the reader is directed to more extensive details found in those Sandia technical reports and archival journal articles (including several submitted for publication at the time of this writing).

## 2. DESORPTION PHYSICS

### 2.1. Theoretical Thermal Desorption Physics

This area of the project focused on fundamental science studies of the desorption of contaminants from rapidly heated electrode surfaces under Z-operating conditions. This desorption leads to plasma formation and reduction of current delivered to the load of pulsed power devices. Two early hypotheses were that (A) the main contaminants (water and hydrocarbons) are adsorbed on electrode surfaces, even after pre-Z-shot evacuation at  $\sim 10^{-5}$  Torr; and (B) the main contaminant is hydrogen inside steel which escapes through the steel and surface oxide film into vacuum as  $H_2$  gas. In the course of this project, we also explored a third, composite hypothesis, namely that (C) hydrogen atoms absorbed inside steel diffuse into the surface oxide, reacting with it to yield water which subsequently desorbs. In the pulsed power literature, eight or more monolayers of “water” were often invoked to account for the plasma generated in Z experiments. One of our main goals was to elucidate the main sources and mechanisms associated with desorption of charge-neutral contaminants. *The amount of contaminant inventory present, and the time scale of desorption, combine to determine whether a proposed mechanism [i.e., (A), (B), or (C)] is relevant to Z operations.* By identifying the key mechanisms responsible, improved models can be developed and mitigation strategies can be devised.

We applied molecular modeling methods, include Density Functional Theory (DFT) electronic structure calculations and force-field-based molecular dynamics (MD) and Grand Canonical Monte Carlo (GCMC) simulations, to elucidate water and hydrocarbon desorption dynamics and initial water inventory on oxide surfaces under Z-like vacuum conditions. The ALEGRAcode was also applied to explore the extent of electrode heating expected on large pulsed power accelerators. Several of these topics was or will be published, and brief summaries are provided in this report; additional unpublished work, knowledge gaps, and possible future research directions are also presented.

#### 2.1.1. *Unpublished Results on Hydrogen Trapping/Diffusion*

Desjarlais applied Density Functional Theory-based Molecular Dynamics (DFT/MD or AIMD) to study hydrogen diffusion rate in molten  $Fe_2O_3$  and  $Cr_2O_3$ . The deduced activation barriers associated with diffusion are about 0.6 eV in both cases, small enough to suggest fast hydrogen motion in the oxide. Hydrogen diffusion is likely not a limiting step. Alan Wright applied static, spin-polarized DFT to study hydrogen absorption and diffusion inside face-center-cubic (FCC) iron metal. Even though pure Fe has a different stable lattice structure, austenitic steel made of iron alloys has a FCC structure; hence this lattice was deemed more relevant. The hydrogen atom absorption energy at the most stable octahedral site was predicted to be 2.38 eV relative to an H

atom in vacuum. This value feeds into our subsequent study about hypothesis (C). The octahedral-to-octahedral diffusion barrier was predicted to be 0.76 eV. Wright also found a octahedral-tetrahedral-octahedral diffusion pathway with a lower, 0.31 eV barrier. This suggested that H diffusion inside steel is significantly faster than H diffusion in the surface oxide film, which has a higher diffusion barrier (see above). Similar calculations were performed in the presence of a Ni dopant. The memos associated with this work can be found in the internal GitLab issue #165. We concluded that H diffusion in the electrode does not seem to be a slow process, which supports hypothesis (A) or (C). Significant amount of hydrogen can be found inside steel, so the initial inventory is not an issue unless the steel is specially treated (“hydrogen fired”).

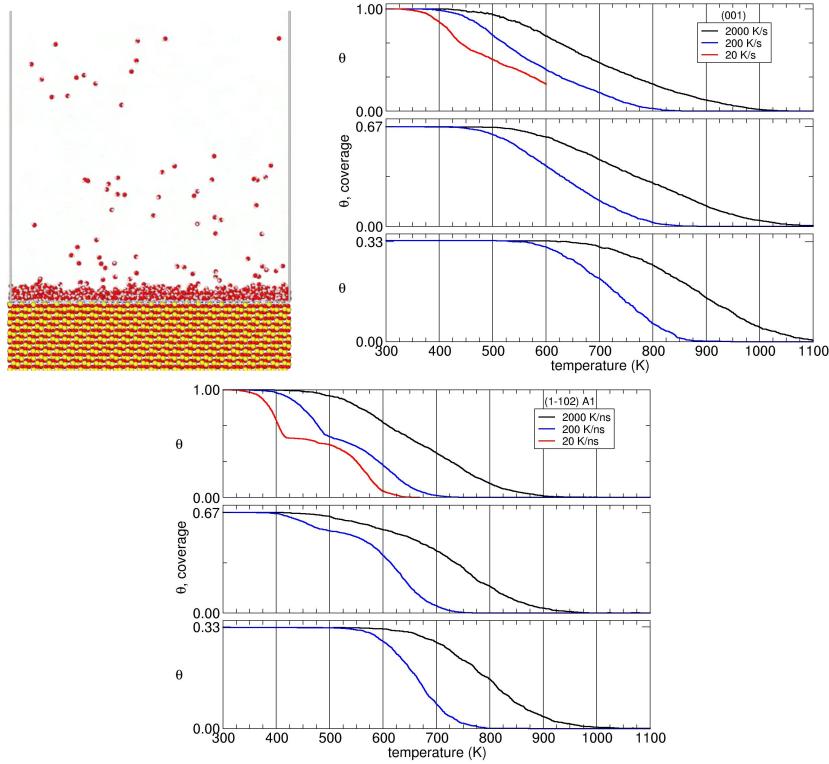
## 2.1.2. **MD studies**

### 2.1.2.1. **Approach and methodology**

Classical molecular dynamics simulation is a common approach for the study of non-equilibrium driven states that are often seen in dynamic extreme environments experiments. It has the time scale ( $\sim 250$  nanoseconds) compatible with Z experiments. The approach has long been used to understand the mechanisms for material response under high-rate mechanical and thermal loading. Here, we apply both non-reactive and reactive methods to study the effect of rapid heating of metal oxide surfaces on adsorbed water and hydrocarbon molecular (physisorbed) films as well as dissociated (chemisorbed) water monolayers (hypothesis (B)). Our primary quantitative observable is the desorption rate of water from these surfaces, but we also investigate more fundamental properties such as binding energies and desorption mechanisms. We investigate three surrogates for stainless steel surfaces, all surface oxide layers. With the non-reactive approach, we use the ClayFF interatomic potential [32] to model iron oxide ( $\text{Fe}_2\text{O}_3$ ), and with the reactive approach, we use ReaxFF [133, 19, 115] to model  $\text{Fe}_2\text{O}_3$ , chromium oxide ( $\text{Cr}_2\text{O}_3$ ) and nickel oxide ( $\text{NiO}$ ). These three oxide layers were identified by experimentalists and selected to best connect with experimental measurements. The two surface contaminants, water and hydrocarbon, were selected because of their likely presence on Z electrode surfaces, the primary application area for this work.

We began the project with a non-reactive approach because it is the most established and efficient method for surface studies of non-dissociative surfaces. However, this approach is only applicable to molecular bound contaminants. In order to study the strongly binding (dissociating) surfaces prevalent on stainless steel surfaces, we employed a reactive potential. These reactive potentials are often much less vetted in these temperature regimes. Thus, significant effort was required to vet these reactive potentials before they could be properly applied to desorption studies.

Classical molecular dynamics simulations were run in LAMMPS [88]. Our primary geometry for these desorption studies was atomically-smooth cleaved surface exposed to vacuum. For each model material, one or more low-energy facets of the ambient phase was studied, a single facet per simulation. The metal oxide substrates ranged in size from  $2\text{nm} \times 2\text{nm} \times 3\text{nm}$  up to  $20\text{nm} \times 20\text{nm} \times 3\text{nm}$  depending on the simulation. See Figure 2-1 for an example geometry. Desorption was studied during high-rate thermal ramps from 300 K to 1300 K, as well as for constant temperatures within this range. The thermal ramps were linear in temperature and



**Figure 2-1.** A snapshot (left) of water desorption from  $\alpha\text{-Fe}_2\text{O}_3$  (left) showing a typical simulation geometry. The desorption profiles as a function of coverage and thermal ramp rate for (center) the (0001) facet and (right) the (1-102) A1 surface. Plots reproduced from Reference [61].

applied at rates from 2000 K/ns down to 20 K/ns. Typical thermal ramps on the Z electrode, due to ohmic heating, are at 10s of K/ns rates. Only the substrates were thermostatted and barostatted to insure that the dynamics of the water were not influenced directly by the driving conditions. Systems were periodic in plane. For the ClayFF simulations, long-range electrostatics were used with a slab dipole correction. A large vacuum space (10 nm) was modeled above the slab. Water molecules which had risen above 0.8 nm from the surface were considered to be desorbed, for counting purposes. In order to prevent water from returning to the surface, molecules were removed from the system when they were 8 nm from the surface.

### 2.1.2.2. Non-reactive ClayFF results: Water Desorption on $\text{Fe}_2\text{O}_3$

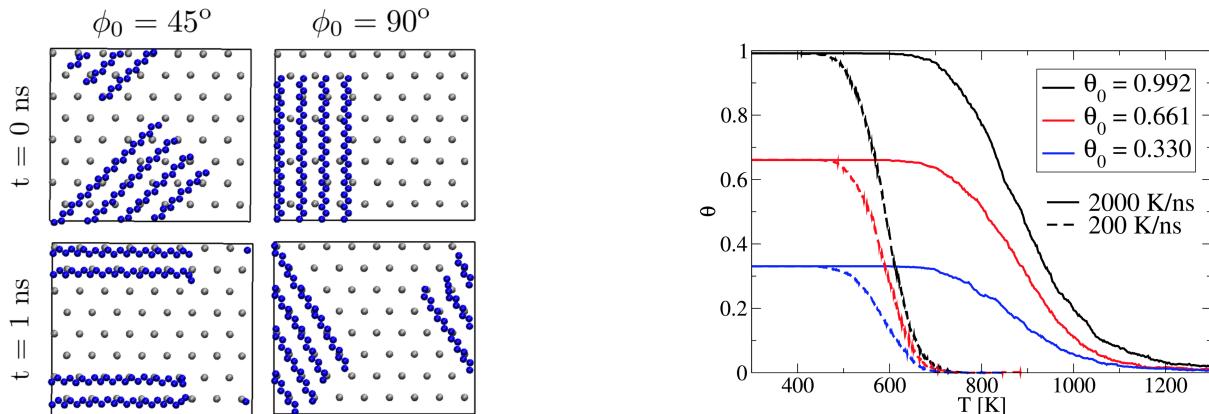
We found that the non-reactive ClayFF force field yielded average zero temperature water binding energies within 0.08 eV of the DFT+U/DF2 electronic structure method. On the  $\alpha\text{-Fe}_2\text{O}_3$  (0001) surface and (1-102) surface, MD simulations revealed significant desorption profile dependence on both the surface coverage and the rate of thermal ramp. From the measurement of the desorption onset temperature we were able to quantify differences in the desorption properties of these surfaces, and use these to parameterize a Temkin reduced model. Surprisingly, the (0001) surface showed significantly greater coverage dependence than the (1-102) C4 surface. We attribute this to the presence of passivating OH groups on this (1-102) surface. This is qualitatively consistent with

the Temkin isotherm expression for effective binding energies, which become reduced with higher  $\theta$ .

Quantitatively, without alterations, Temkin isotherm parameters widely used in the plasma community already references desorption time scales within 50% and 100% of MD data on the  $\text{Fe}_2\text{O}_3$  (0001) and (1 $\bar{1}$ 02) surfaces, respectively. Refitting water binding energies used in Temkin isotherm in the  $\theta \rightarrow 0$  limit to single  $\text{H}_2\text{O}$  desorption MD data yields water desorption rates within 10% of MD data over a large range of initial water content. Full details of this work have been published in Reference [61].

### 2.1.2.3. Non-reactive ClayFF results: Water and Hydrocarbon Desorption on $\text{Fe}_2\text{O}_3$

We simulated and analyzed the desorption of hydrocarbon and hydrocarbon-water mixtures from an  $\alpha\text{-Fe}_2\text{O}_3$  (0001) substrate. We found that the desorption temperature was largely unaffected by the hydrocarbon structure, varying <2% between  $\text{C}_{24}\text{H}_{50}$  paraffin and  $\text{C}_{24}\text{H}_{48}$  naphthene hydrocarbons. Unlike in water, the desorption rate was independent of hydrocarbon coverage. It was dependent on hydrocarbon chain length, due to the increase van der Waals interactions. In the presence of water, the hydrocarbons phase separate on top of the water layer which forms at iron-oxide surface. Figure 2-2 shows an example of the initial and final configuration of hydrocarbons, and an example plot of desorption profiles for the  $\text{C}_{24}$  case, as a function of coverage and rate. Full details of this work have been published in Reference [59].



**Figure 2-2. Hydrocarbon simulations on  $\text{Fe}_2\text{O}_3$ :** (left) visualization demonstrating preferred orientation of polymer 'rafts' on the  $\alpha\text{-Fe}_2\text{O}_3$  (0001) substrate. The top shows initial placement of  $\text{C}_{24}\text{H}$  paraffin chains, bottom image is after 1 ns. (right) Desorption profiles for varying coverages of  $\text{C}_{24}\text{H}_{50}$  paraffin hydrocarbons at various coverages and thermal ramp rates. Plot and figure reproduced from [59].

### 2.1.2.4. Reactive ReaxFF results: $\text{Fe}_2\text{O}_3$ , $\text{Cr}_2\text{O}_3$ and $\text{NiO}$

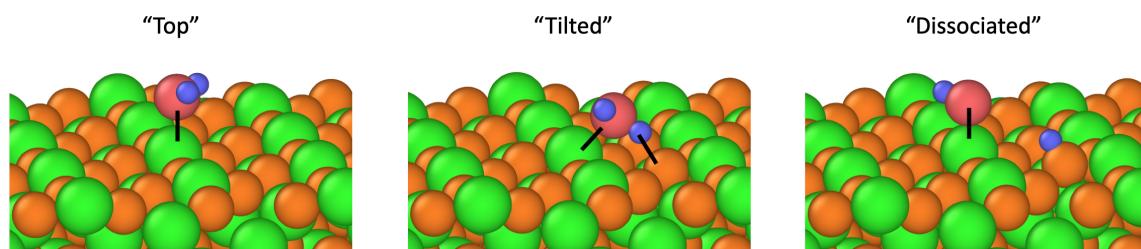
It is known that some  $\text{Cr}_2\text{O}_3$  and  $\text{NiO}$  surfaces dissociate water due to stronger binding to water. Moreover, a combination of molecular water and dissociated water (i.e. surface OH groups) may exist at different temperatures, leading to complex dissociation and recombination behavior.

Accurate modeling of this complex environment requires the ability to capture chemical reactions. ReaxFF is a leading framework for simulating chemistry in classical MD, but most ReaxFF parameterizations are not trained for surface studies at high temperatures. A significant effort went into testing and validating ReaxFF parameterizations for  $\text{Fe}_2\text{O}_3$ ,  $\text{Cr}_2\text{O}_3$  and  $\text{NiO}$  surfaces for this project. While this effort is not complete, significant progress has been made. Our approach was iterative, (1) we began with the evaluation of the substrate and facet stability at high temperatures; (2) the molecular water and dissociated binding energies were calculated and compared to DFT calculations; and, (3) we evaluated the dynamic recombination barrier energies. In most cases, parameterizations failed in either steps 1 or 2. Two recently acquired parameterizations, created specifically for our project, are in step 3 testing, currently.

**Iron Oxide,  $\alpha\text{-Fe}_2\text{O}_3$  (0001):** It was initially believed that  $\text{Fe}_2\text{O}_3$  would be a good starting point for ReaxFF implementation because detailed comparisons could be made to the ClayFF results. Unfortunately, the two potentials in the literature were unsatisfactory for quantitative desorption studies. One predicted far too strong a binding to water, and the other was unstable to thermal ramps. A hybrid combination of the two looked promising, but ultimately failed to model a stable water monolayer at 300 K, due to a weaker than necessary water binding in the tilt orientation.

**Chromium Oxide,  $\alpha\text{-Cr}_2\text{O}_3$  (0001):** Results for  $\text{Cr}_2\text{O}_3$  were far better than for  $\text{Fe}_2\text{O}_3$ . A parameterization from Shin, et al. gave reasonably good substrate lattices, and some binding energies. The potential predicts spontaneous dissociation at the surface. However, the potential did not accurately reproduce the DFT+U minimum energy state, and tended to overpredict desorption energy barriers for both molecular and dissociated water. Shin and van Duin have recently provided a custom parameterization developed specifically to address these shortcomings.

**Nickel Oxide,  $\text{NiO}$  (100) and (111):** Eight  $\text{NiO}$  ReaxFF parameterizations were tested for crystal stability. All failed to reproduce the ambient rock salt structure. As previously stated, Shin and van Duin have recently provided a custom parameterization which not only stabilized the substrate, but is also trained specifically to water surface binding. The potential predicts spontaneous dissociation at the surface. Binding sites and energies are being evaluated for individual water and full/partial monolayers.



**Figure 2-3. Water binding sites as modeled by the ReaxFF potential. The top, tilted and dissociated state energy ordering are being evaluated for a new ReaxFF parameterization.**

### 2.1.2.5. Conclusion and Future Directions

Classical molecular dynamics has proven to be an extremely useful tool for modeling the longer time scales and length scales required for desorption studies. When accurate potentials can be implemented, MD can bridge the gap between quantum calculations and experimental observables. We found that the ClayFF potential has proven to be a quantitatively accurate model for studying water desorption. We were able to apply it to validate the long-standing empirically fit Temkin model using MD, and improve the model fits for water desorption. ClayFF, however, is only applicable for relatively weakly bound molecular water. The promise of a more generalizable reactive model has not been realized within the timeframe of this Grand Challenge LDRD. However, recently developed parameterizations for chromium oxide and nickel oxide are still being evaluated. This effort will continue through an active collaboration with ReaxFF creator Adri van Duin.

### 2.1.3. Quasi-Equilibrium Water Desorption Barriers

A key equation used in plasma codes to describe the contaminant-derived plasma source takes the form

$$\frac{d\theta(t)}{dt} = -k_o(T) \theta(t)^x \exp\left(-\frac{\Delta G^*[\theta(t)]}{k_B T(t)}\right), \quad (2.1)$$

where  $\Delta G^*$  is the activation free energy barrier,  $k_B T$  is the thermal energy at temperature  $T$  in Kelvin (K), the exponent  $x$  describes the kinetics order which is unity for sub-monolayer intact water molecules,  $\theta$  is the fraction of monolayer water coverage, and  $k_o(T)$  is the kinetic prefactor. Although the Z-machine pulse is short, on the order of 250 ns, modeling of the contaminant source term universally made a quasi-equilibrium assumption – which is the origin of the Arrhenius (exponential) term in Equation 2.1. In Reference [61], we applied a standard “Temkin” form  $\Delta G^*(\theta) = E_b(1 - \alpha\theta)$ , without explicit temperature dependence, to compare with MD desorption profiles. Setting  $E_b$  to the T=0 K binding energy yielded real-time desorption profiles which disagreed with direct MD simulations by up to 50%. Much better agreement was obtained when  $E_b$  was treated as an effective parameter, for reasons not understood at the time.

To elucidate the reason and to obtain a more fundamental science-motivated functional form, we conducted explicit determination of  $\Delta G^*$  from equilibrium MD trajectories. In the literature, this is generally performed with enhanced sampling, which is a mature and well-developed area. Hence Leung, Criscenti, and Robinson applied standard potential-of-mean-force technique was applied to calculate  $\Delta G^*$  as a function of  $\theta$  and  $T$  [66]. The predicted temperature dependence of  $\Delta G^*$  was  $\sim 5 k_B T$ , which was larger than initially expected, and reflected the entropy of water vapor released. This large temperature dependence was the reason  $E_b$  could not be approximated by the T=0 K binding energy in Reference [61]. An analytical  $\Delta G^*$  functional form was obtained which, despite some uncertainties, was mostly free from assumptions and was not a fit to results at any temperature ramp rate. (The formalism corresponded to infinitely slow desorption.) This new  $\Delta G^*(\theta, T)$  expression gave water desorption profiles in reasonable agreement with the profiles predicted using the Temkin form fitted in Reference [61] at temperature ramp rates within three orders of magnitudes of the ramp rate at which the fitting was performed in Reference [61].

When the old and new  $\Delta G^*(\theta, T)$  functional forms were used for predictions at ramp rates slower than those, the desorption profiles diverge and the half-width desorption time can differ by a factor of four. This suggested that, if the bench-top experimental (“temperature programmed desorption” or TPD) temperature ramp rate is within 1000 times of MD simulations, reasonable comparisons can be already made. Although based on a specific model surface, this work raised the hope that quantitative comparison between modeling and bench-top TPD does not absolutely require that TPD achieves the fast time scale associated with Z experiments.

The MD simulations described in this and the previous section revealed desorption profiles compatible with plasma generation events. They validated hypothesis (B) discussed in this Chapter’s introduction. Although other computational methods were needed to examine the initial contaminant inventory on electrode surfaces prior to each Z experiment to determine how much plasma can be attributed to hypothesis (B), our work suggested that time scales and activation free energy barriers relevant to Z-operations can be dealt with using existing methods – as long as water adsorbs without dissociating on the surface.

#### **2.1.4. *GCMC studies of water on surfaces and in pores***

As mentioned above, other computational methods, like GCMC, were needed to examine the initial contaminant inventory on electrode surfaces prior to each Z experiment. Using GCMC simulations, Thompson concluded that the average coverage of intact water on flat  $\text{Fe}_2\text{O}_3$  (0001) surfaces at  $T=300$  K and  $10^{-5}$  Torr pressure is below 0.4 monolayers. The coverage on the (1-102) surfaces is close to zero, not counting the  $\text{FeOH}$  groups which are strongly bound [61]. Unpublished GCMC results on  $\text{Cr}_2\text{O}_3$  (0001) surfaces, conducted using a Cr(III) forcefield created by Criscenti, were similar to those on the  $\text{Fe}_2\text{O}_3$  (0001) surface. So far GCMC had not been applied to dissociative water adsorption or hydrocarbons.

Tranchida, Thompson and Lane also conducted an evaluation of initial water coverage and desorption from slit pores. Slit pore models were designed and GCMC and MD simulations were carried out using Sandia’s LAMMPS code. All calculations used a corrected ClayFF interatomic potential. The GCMC results do not predict increased water coverage for slit pores of  $10\text{ \AA}$  and more. For smaller pores, and at pressure of interest for Z experiments ( $0.76 \times 10^{-5}$  Torr), up to 40% water coverage increase can be observed, compared to GCMC surface results. At lower pressures ( $0.76 \times 10^{-7}$  and  $0.76 \times 10^{-8}$  Torr), the slit pores can produce an even larger coverage difference (almost no coverage for flat surfaces, compared with 0.2 to 0.35 ML for 2 and 3  $\text{\AA}$  slit pores, respectively). However, the effective number of adsorbed monolayers on each of the pore surface remains below 1.0. Using the equilibrium water coverage generated by the GCMC calculations as initial conditions, MD simulations were performed at different temperatures and for different slit pore sizes. Initial results indicate a decrease of the desorption rate related to pore depth and width, compared to surface results. Additional calculations will be performed to confirm this trend. For details, see Reference [132].

Finally, Leung performed continuum calculations, analogous to those by Li and Dylla on flat surfaces [67], to explore water desorption dynamics from a model cylindrical pore with water adsorbed on its interior surface. The pores were meant to mimic cavities in the steel oxide

covering the metal. The main conclusion was that water desorption dynamics was not significantly different in pores with reasonable diameters and lengths compared to those on flat oxide surfaces; the desorption time scale was increased by only  $\sim 10\%$ . These findings were described in a memo (internal GitLab issue #165).

These studies suggested that flat oxide surfaces hold only sub-monolayers of water contaminants, and that pores contribute to a small perturbation on desorption kinetics and initial inventory. These deductions are in broad agreement with electrochemical measurements of electrode roughness. While we demonstrated that hypothesis (B) is viable in terms of time scales, and that it contributes to plasma generation, we also showed that (B) does not fully account for the multilayers of inventory needed to explain plasma formation in Z experiments.

### ***2.1.5. $H_2O$ from reaction between surface oxide and hydrogen inside steel***

As mentioned above, our GCMC studies [61], literature estimates [67], and experimental work all suggested that impurities adsorbed on electrode oxide surfaces, even those with pores, are not sufficient to explain the initial water inventory used in plasma codes to analyze Z experiments. Hence we developed hypothesis (C): hydrogen in the steel moves through the surface oxide and attaches to oxygen ions on the surface. If a second hydrogen attaches to the same oxygen, a water molecule is formed; it can now desorb from the surface. The mechanism creates two Fe(II) and removes a O atom from the oxide.

In order to confirm this, we added hydrogen to the system in several different configurations to the surface and/or bulk, optimized the atomic configurations, and calculated the DFT energies which were referenced to the energies of excess H atoms in FCC Fe computed by Wright. This is ongoing work. One definite result we obtained was on the (10-10) surface. This  $Fe_2O_3$  facet is known to contain bridging OH groups where the O atom is coordinated to two surface Fe(III) ions [43]. We found that adding a hydrogen to such a bridging OH group is energetically favorable compared with H inside FCC iron metal. Furthermore, removal of the  $H_2O$  molecule so formed required only  $\sim 1$  eV energy. Assuming this is the  $T=0$  K  $\Delta G^*$ , there is no coverage dependence to desorption on this surface,  $x$  is unity (first kinetic order), and it is appropriate to use the kinetic prefactor temperature dependence deduced in Reference [66], Equation 2.1 predicts that the half-width of the desorption profile occurs at 120 ns in a 250 ns pulse with linear increase of temperature from  $T=300$  K to  $T=1300$  K. At that halfway point,  $T=780$  K. This work is described in Reference [21].

These results suggested that hypothesis (C) indeed generates desorbed  $H_2O$  at time scales relevant to Z-operations. (C) has the added advantage that the hydrogen source is not limited to the immediate surface region, and can therefore account for multilayers of contaminants. This in turn raised questions about the kinetic order of this reaction. This issue, and others, will be addressed in follow-on work.

### **2.1.6. *Electrode Modeling in ALEGRA***

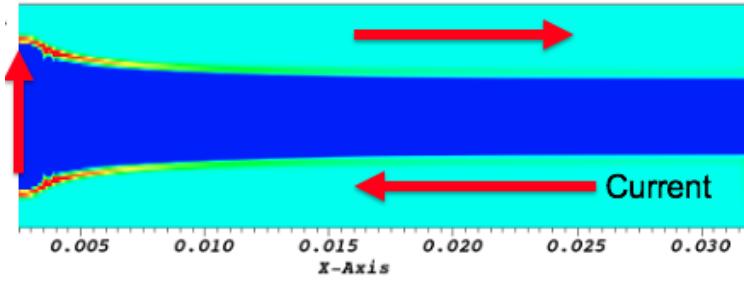
We also conducted more coarse-grained modeling. In order to design a next-generation pulsed power accelerator, accurate material models are required as well as understanding of which regimes of phase space in those models are being accessed. We used the radiation magnetohydrodynamics code ALEGRA[98, 96] to perform two dimensional simulations of current flow through an anode/cathode with a short on axis as shown in Figure 2-4. These simulations represent the existing Z magnetically insulated transmission lines (MITLs) made of stainless steel 304L (a low carbon form of stainless steel) with an A/K gap of three millimeters and 8 cm radius. It was decided that the feedback mechanism of using a circuit models was not needed for these scoping calculations, so the standard Z current profile was used. This profile is representative of that used by inertial confinement fusions experiments, which lasts for approximately 100 ns and peaks at 25 MA. We also compared the MHD results to simulations that used single group (grey) radiation diffusion, as this will allow energy transfer between a hotter near axis short and the MITL's a few centimeters from axis.

The simulations of the 25 MA standard Z current showed that within 5 cm of the axis, the magnetic field compressed the MITLs from  $7.9 \text{ g/cm}^3$  to  $8.5 \text{ g/cm}^3$  and the temperature increased to a few thousand degrees with more compression near axis as expected. The simulation remained cool enough that radiation transport played a minimal role in shifting energy away from axis. We then recompleted the simulations while increasing the current to 60 MA, maintaining the shape of the current pulse as a function of time but increasing the magnitude. In these simulations, the MITLs showed significantly more motion and joule heating. The final density was still a function of radius but near axis increased to  $25 \text{ g/cm}^3$  with significant A/K separation, and as far as 5 cm from the axis was above  $10 \text{ g/cm}^3$ . The temperature showed a similar pattern with the axial temperature being over 20 eV. Farther from axis, the joule heating still melted the MITLs out to 5+ cm. Radiation diffusion did lower the axis temperature and raise the MITL temperature to about 3 cm from the axis. In these simulations we used the SESAME stainless steel equation of state modified by Desjarlais/Mattsson and a preliminary conductivity model created by Mattsson [94] (with a primary focus in the liquid regime near the vapor dome).

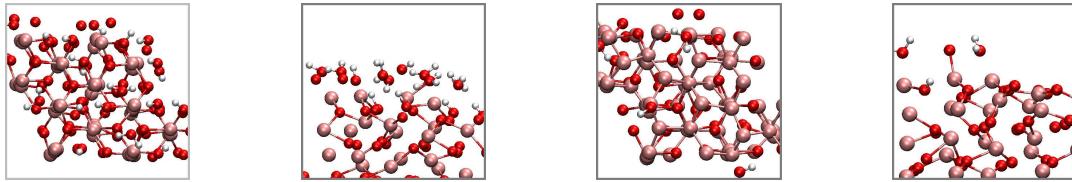
Because these models are critical to future accelerator design, we sent samples to UC San Diego Planetary and Experimental Petrology Laboratory in order to confirm the model accuracy as they have the capability to perform conductivity measurements in diamond anvil cell experiments to approximately 20 GPa. These experiments showed the conductivity model under compression was incorrect. Based on this information we performed magnetic burn-through experiments on the Thor pulsed power driver, which both confirmed the model inadequacy and will be used as prototypes for similar experiments on Z. Based on these results, a full equation of state and conductivity model development plan has been implemented with completion expected in FY2021 and results to be published in FY2021.

### **2.1.7. *Technical Gaps and Future Directions***

Despite the significant advances made in this Thrust, several technical gaps remain. Future molecular modeling should focus on topics which include (i) increased comparison of modeling



**Figure 2-4. Simulation of a stainless steel MITL; current enters along the bottom electrode, travels through the short on axis, and exits along top electrode.**



**Figure 2-5. (a)-(b) Top and side views of DFTB simulation cell at  $t \sim 0$  ps, with 8 intact and 4 dissociated  $\text{H}_2\text{O}$ ; (c)-(d) at  $t \sim 25$  ps, where 2 intact and 1 dissociated  $\text{H}_2\text{O}$  remain.  $T=1300$  K.**

predictions with measurements on real electrodes; (ii) desorption of dissociated water; (iii) carbon and nitrogen contaminants; and (iv) initial inventory of oligomeric species under vacuum.

Regarding (i), comparisons between contaminant desorption modeling and experiments require that both occur on reasonably similar surfaces and at similar time scales. Temporally, MD trajectory lengths are many orders of magnitude lower than bench-top TPD time scales. However, the several orders of magnitude acceleration in TPD temperature ramp rates achieved in other areas of this project and the prediction that TPD only needs to be within 3 orders of magnitude of MD rates to give reasonable comparison [66] meant that this gap was substantially reduced. As for surface specificity: steel surfaces (and other metallic surfaces except those of gold) are covered with oxides. While metals like aluminum are known to exhibit amorphous oxide coatings, transition metals like iron, chromium, and nickel (all common steel components) generally form crystalline oxides on their surfaces. The phases and the facet terminations of the polycrystalline oxides are difficult to characterize, partly because of the small oxide grains. Towards the end of the project, we partly resolved this issue by heating samples (iron coupons) to 650°C in air. This created larger oxygen grains. Ilgen and Goeke obtained experimental evidence that  $\alpha\text{-Fe}_2\text{O}_3$  (0001) and (10-10) facets are present in samples prepared this way, in agreement with the facets chosen in molecular calculations [21]. We advocate that future TPD measurements be conducted on these specially prepared material surfaces to enable quantitative comparisons with molecular modeling.

Regarding (ii): dissociative water adsorption occurs readily on metal oxide surfaces. Demonstrating the reverse process,



would be useful for extracting the kinetic prefactor  $k_o(T)$  (Equation 2.1) under well-defined conditions for dissociative  $\text{H}_2\text{O}$  desorption, analogous to what was done for non-dissociative  $\text{H}_2\text{O}$  desorption [66]. This would involve comparing unconstrained MD simulations with the desorption times computed using Equation 2.1 at different temperatures.

To model dissociated water, we applied the ReaxFF model, which supports splitting of  $\text{H}_2\text{O}$  into its component atoms. So far, ReaxFF had not demonstrated recombination of OH groups to reform  $\text{H}_2\text{O}$  at finite temperature. This work will continue in FY21 under a separate LDRD project. Another viable backup option is the Density Functional Tight-Binding (DFTB) approach [68]. As shown in Figure 2-5, the reaction described by Equation 2.2 spontaneously occurred on the  $\text{Fe}_2\text{O}_3$  (0001) surface in a DFTB simulation within 25 ps at  $T=1300$  K. The advantages of DFTB are as follows: (i) It is relatively accurate in terms of crystal structures;  $\alpha$ - and  $\gamma$ - $\text{Fe}_2\text{O}_3$  crystal slabs appear stable at temperatures up to 1600 K in short MD trajectories; (ii) It explicitly treats spin degrees of freedom, just like DFT; (iii) The water adsorption energetics appear within 20% of DFT/PBE values. The main disadvantage is the computational cost; the DFTB+ code only runs on one node, which means several picoseconds of MD trajectory a day. Furthermore, DFTB+ does not yet permit  $T=0$  K calculation of dissociation barriers, and desorbed water in finite temperature MD simulations had to be removed periodically by hand. Transporting the DFTB parameters to the LAMMPS code may solve these problems in the future.

Regarding (iii), the memo by Criscenti (GitLab issue 165) made succinct recommendations about future DFT calculations to investigate C and N defects. Criscenti conducted a comprehensive search of the impurity-in-steel experimental and molecular modeling literature to help define what should be the future molecular modeling directions associated with defect energetics in the context of Z operations. Finally, regarding (iv), the use of GCMC techniques would be ideal for calculating the equilibrium coverage of hydrocarbon on steel oxide surfaces. Given the number of internal degrees of freedom in  $\text{C}_{24}$ -oligomers, however, such calculations would represent a major investment.

## 2.2. Desorption Experiments and Surface Characterization

This area of the project focused on conducting various light-lab experiments to determine the initial adsorbate inventories and desorption rates during rapid electrode heating at timescales relevant to large pulsed power accelerators, in an attempt to better understand current loss mechanisms, provide improved boundary conditions for multi-physics codes (e.g., CHICAGO and EMPIRE), and move towards predictive plasma models. Previously discussed in Section 2.1,  $\simeq 8$  monolayers of "water" is often invoked in plasma physics simulations, using a Polanyi-Wigner desorption model, to approximate the current losses observed during Z-Machine experiments. However, this assumption overly simplifies the electrode surface conditions and subsequent desorption kinetics. At the onset of this LDRD project, our main objectives were to: (i)

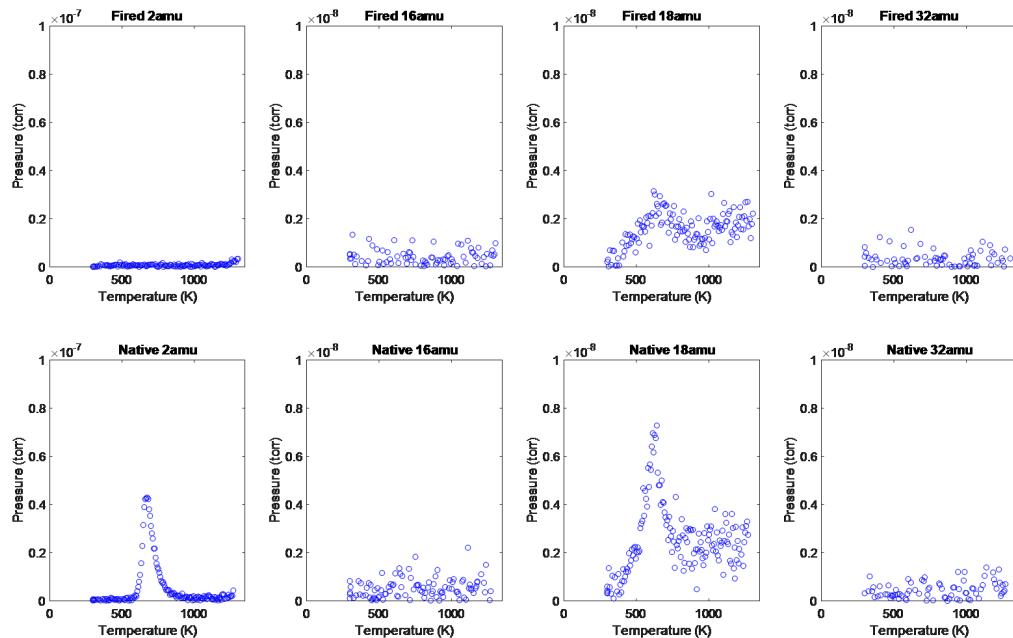
experimentally determine the composition and inventories of adsorbates on real electrode materials at relevant vacuum pressures (i.e.  $10^{-5}$  Torr); (ii) improve the desorption model to include parameters for all desorbates informed by measurements; and (iii) experimentally assess the validity of the Polanyi-Wigner desorption model for thermal ramp rates approaching 2,000 K / 100 ns (i.e. similar to those expected at the Z-Machine).

To help answer (i), we used a combination of Temperature Programmed Desorption (TPD) and X-Ray Photoelectron Spectroscopy (XPS) to help elucidate the composition and initial inventories of adsorbates on 304L stainless steel. Grazing Angle Infrared Spectroscopy was used to assess the presence of vapor-deposited transformer oil that could be present on electrodes found in pulsed power environments. To further inform the surface coverage of contaminants, we assessed the "real" versus geometrical surface area of stainless steel coupons using electrochemical measurements of the double layer capacitance. For (ii), we relied on conventional TPD methods and data analysis techniques to extract the various Polanyi-Wigner parameters. To help improve the data analysis and deconvolve the complex desorption spectra, we developed a multi-peak fitting tool that was ultimately exercised to analyze sample TPD data. As for (iii), we developed an in-house TPD system capable of 100 - 19,000 K/s. This is 2,000x faster than our previous capabilities and 10x faster than other published methods. For a 2,000 K / 100 ns flash desorption capability, we collaborated with Northern Arizona University to conduct a feasibility study for developing a novel ultra-fast mass spectroscopy device. We also collaborated with Colorado State University to evaluate a potential high speed mass spectrometer that may increase our detection speed by several more orders-of-magnitude. The following sections briefly describe this experimental work with additional discussion on technical gaps and future directions.

### **2.2.1. *Slow Temperature Programmed Desorption of Hydrogen Fired 304L Coupons***

Temperature programmed desorption experiments were performed to study surface and bulk inventories of species in native and hydrogen fired stainless steel 304L [21]. In most prior TPD literature, samples are purged of their inventories via in-situ heating and then re-dosed in a controlled fashion to study the energy of desorption for a given species. For many pulsed power applications concerned with off-gassing species, complete removal of surface and bulk contaminants is not as practical. Our approach to these experiments was to investigate desorption of realistic material states (i.e. cleaning method) for pulsed power applications to understand how these inventories could be viewed as contaminants. Initial desorption experiments revealed multi-peak water spectra for native (IPA cleaned) stainless steel 304L coupons, which was intriguing due to the chamber vacuum pressure during measurement ( $10^{-10}$  Torr) prior to desorption. Partial pressure peaks at 200°C and 400°C were common for TPD runs performed at 1°C/sec. At those pressures, loosely bound water should be removed by the vacuum, so additional species were investigated to understand alternative methods for water desorption via recombination. Figure 2-6 shows TPD spectra of a stainless steel 304L coupon in its native state (lower plots) compared to one after a hydrogen firing treatment (upper plots); this is a multi-step cleaning process using high temperature (approx. 800°C), vacuum, and hydrogen gas to prepare metals for vacuum applications and is assumed to largely remove native oxides. Near complete

removal of hydrogen was observed with the hydrogen firing treatment, which also correlated with a water desorption peak reduction. No obvious oxygen peak was observed in either test case. Water desorption estimated via the curve-fitting GUI were 0.44 monolayers for native samples and 0.32 monolayers for hydrogen-fired coupons. While background subtraction was not performed on this data set, TPD experiments without a coupon on the stage produced partial pressures for water over an order-of-magnitude below peak desorption pressures. This data set may support the hypothesis that hydrogen inventories pull oxygen from native oxides during temperature programmed desorption which produces a water signal at the mass detector at temperatures and pressures beyond where one would expect to observe water desorption.



**Figure 2-6. Mass spectra of 2, 16, 18, and 32 amu obtained from (upper frames) hydrogen fired and (lower frames) native stainless steel 304L coupons during temperature programmed desorption at 1°C/sec.**

## 2.2.2. 100 - 20,000 °C/s Programmed Desorption of Untreated 304L Stainless Steel Coupons

As part of this LDRD project, we also developed a custom TPD system capable of thermal ramp rates ranging from 100 - 20,000 °C/s with a near linear profile. This was achieved through the use of a 24 volt DC battery system capable of delivering 4 kA of current to a metal foil resulting in rapid joule heating. This system was initially used to conduct a sweep of desorbates having mass-to-charge ratios of 2 - 44 amu. Further, these measurements were calibrated against pre-calibrated leak standards for each of the respective gases; in the case of mass 28, both a *CO* and *N*<sub>2</sub> standard were used.

The initial results suggested the presence of H, C, O, and possibly N that desorb from the surface as  $H_2$ ,  $CO$ ,  $CO_2$ ,  $H_2O$ ,  $N_2$ , and  $CH_4$ . Similar to the slow TPD experiments, we did not observe any  $O_2$  desorption, which suggests the oxide layer is reduced by either H or C and desorbs as gaseous  $CO$ ,  $CO_2$ , or  $H_2O$ . Table 2-1 summarizes the initial inventories,  $\theta$  given in monolayers, for the first peak of desorption (weakest bound state) from native stainless steel 304L coupons. Here, one monolayer is defined as  $10^{15} molecules/cm^2$ .

m/e	Species	$\beta$ (K/s)	n	$\theta$ (ML)	$v$ (s <sup>-1</sup> )	E (kJ/mol)
2	$H_2$	145	2	16.49	1.85e6	161.93
		344		20.22	7.92e3	127.73
		643		10.83	3.28e5	154.55
18	$H_2O$	145	2	0.10	1.24e6	57.02
		344		0.10	3.33e6	48.40
		643		0.09	3.61e5	42.70
28	$N_2$	145	2	0.43	3.26e4	59.96
		344		0.65	6.05e4	68.12
		643		0.54	1.69e5	77.08
44	$CO_2$	145	2	1.95	4.70e2	44.07
		344		1.85	4.63e2	42.40
		643		2.04	5.02e2	41.23

**Table 2-1. Desorption parameters for the weakest bound state of mass 2, 18, 28, and 44 amu obtained from native stainless steel 304L coupons during rapid temperature programmed desorption at 145 – 643°C/sec.**

These results suggest there is very little water desorption ( $\simeq 0.1$  ML) occurring as a sample is quickly driven to near melt. Alternatively, as much as 11 - 20 ML of hydrogen quickly desorbs under these same conditions.

It should be noted these initial tests also infer the respective desorption parameters for each adsorbate. While we are suspicious of the pre-exponential factor,  $v$ , and corresponding energy of desorption,  $E$ , (i.e. when compared to values extracted using techniques described in literature), these preliminary values can be incorporated into the multi-physics plasma codes as boundary conditions for each desorbate as a proof-of-principle.

### 2.2.3. X-Ray Photoelectron Spectroscopy of 304L Stainless Steel

An XPS study was conducted to help further elucidate the actual inventories of adsorbates on an electrode surface versus those found in the bulk electrode. This measurement technique also has the added benefit of being capable of distinguishing molecular water present on a metal surface. Using similar 304L stainless steel coupons, we were able to model the spectra using Simulation of Electron Spectra for Surface Analysis (SESSA) from NIST to both determine the composition of the adsorbed layers as well as model the substrate intensity to determine the amount present in each layer. As shown in Table 2-2, the molecular water accounts for only 0.3 ML, which is in

good agreement with the slow TPD results. Note that XPS does not have the capability to detect hydrogen, so these results are consistent with those inferred via TPD.

Adsorbed Species	UHV 20 °C	UHV @ 250 °C	Desorbed (20 - 250 °C)
CO	1 ML	0.5 ML	0.5 ML
CO <sub>2</sub>	2 ML	1.7 ML	0.3 ML
H <sub>2</sub> O	0.3 ML	0 ML	0.3 ML
OH	2.2 ML	0 ML	2.2 ML
C (C-C or C-H)	8 ML	6.4 ML	1.6 ML
<b>Total</b>	<b>13.5 ML</b>	<b>8.6 ML</b>	<b>4.9 ML</b>

**Table 2-2. Initial inventories of various adsorbates present on a native 304L stainless steel coupon at room temperature and at 250°C, measured via XPS.**

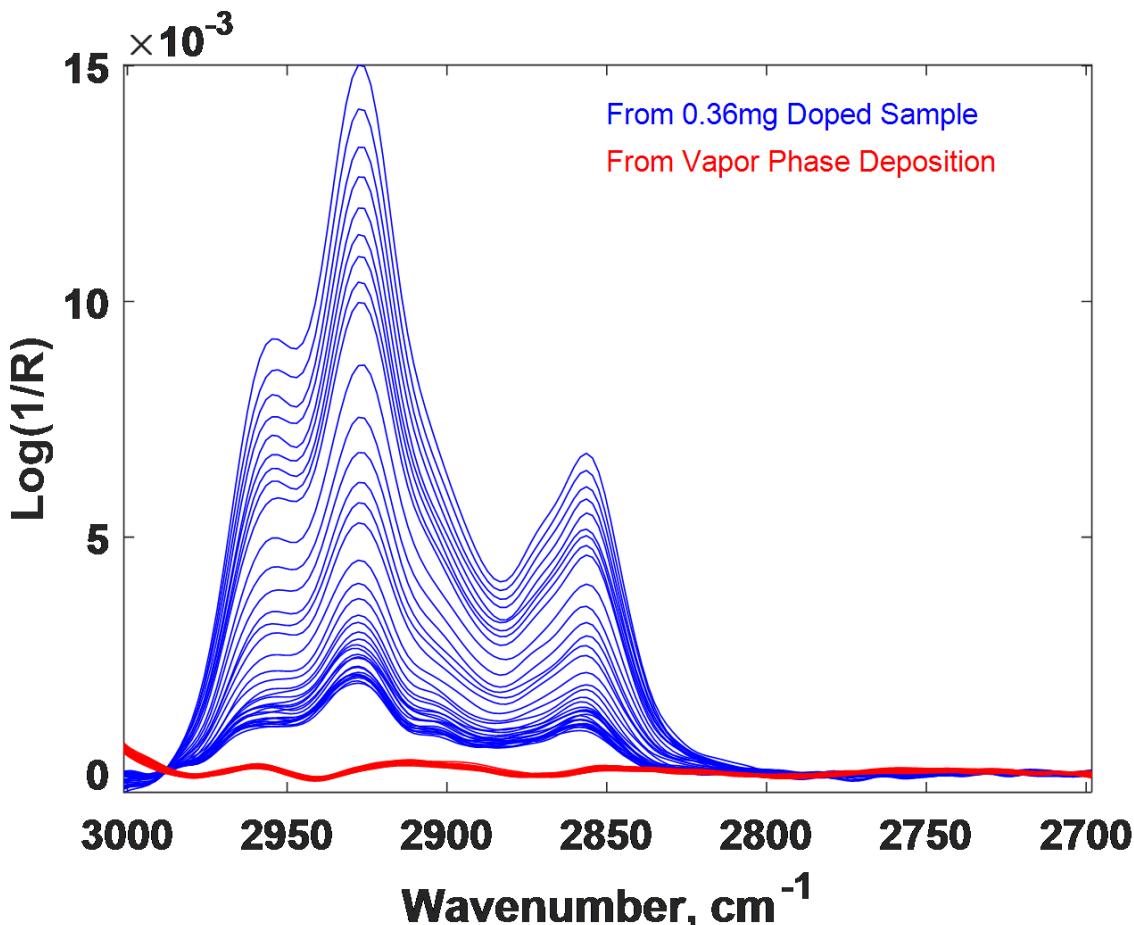
#### **2.2.4. *Grazing Angle Infrared Spectroscopy of Mineral Oil on Stainless Steel at Low Pressure***

Mineral oils and other alternative formulations have historically been used for their electrical insulating properties in transformers, high-voltage (HV) capacitors, and for pulsed power. There are concerns that these oils could also be a source of surface contamination on critical surfaces. For instance, the Z-Machine at Sandia National Laboratories has 36 Marx generators submerged in one million gallons of insulating oil, which is open to the lab environment [109]. There is a concern that any vaporized and redeposited oil may affect experimental behavior, especially if the oil is deposited on the electrode surfaces of the vacuum transmission lines and could subsequently lead to current loss during rapid electrode heating [6, 103, 104].

Infrared Reflectance-Absorption spectroscopy (IRRAS) was used to explore the relative surface concentration of oils deposited on stainless steel surfaces subject to the vacuum environment of  $10^{-5}$  Torr, typically encountered within the pulsed power experiment. IRRAS is also known as infrared grazing angle spectroscopy, transfection and RAIRS (reflection absorption infrared spectroscopy), among others. It is used to examine thin films on surfaces, subjecting the sample to an interrogation infrared beam at an oblique angle. Hoffard notes that contaminant levels as low as  $1 \mu\text{g}/\text{cm}^2$  can be detected with the technique [48]. The light reflected from the surface is highly polarized; thus, enhanced sensitivity is obtained with a polarizer (p-polarization). The oblique angles allow detection of sub-nanometer film thicknesses when using grazing angles greater than 75°.

Using this technique, a simulated surface was monitored within a commercial infrared spectrometer, kept at  $10^{-5}$  Torr in order to simulate vacuum conditions within the Z-Machine [1]. Surfaces doped with a commercial mineral oil were found to show quick loss of the high volatility components, while the low volatility components remained. Experiments to simulate vapor deposition of mineral oil showed no detectable hydrocarbons on the stainless-steel surfaces. Figure 2-7 summarizes these findings in which Fourier Transform Infrared (FTIR) spectroscopy is unable to detect vapor deposited Shell Diala (red) when compared to 0.36 mg of applied to the surface of a 304L stainless steel coupon, which is then characterized under high vacuum as a

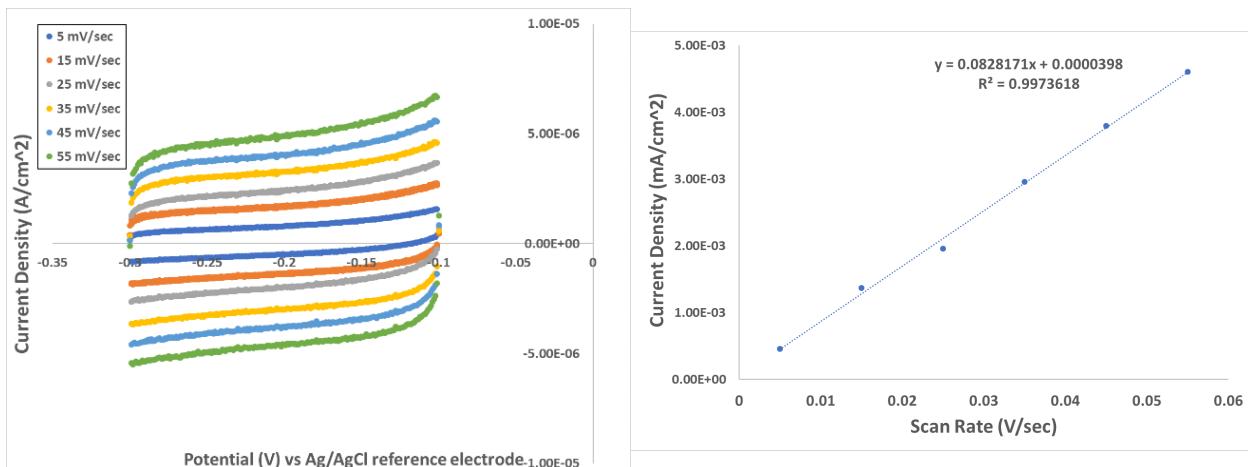
function of time. The most recent results suggest that for parts and samples that have been handled appropriately, little insulating oil should be present due to the environment, even in the presence of large volumes of open-atmosphere oil tanks such as those at the Z-Machine.



**Figure 2-7. Desorption of a stainless steel disk exposed to vapor phase light naphthalic oil (red), compared to desorption of doped sample as a function of time under vacuum (blue).**

### 2.2.5. *Electrochemical Surface Area Measurements*

Another possible uncertainty raised during this LDRD project was the distinction between a "real" versus "geometrical" surface area of a stainless steel electrode. This is of practical importance to help constrain the initial inventories of adsorbates used in multi-physics plasma simulations. Thus, the characterization of true surface area was estimated via electrochemical measurements of double layer capacitance. The electrical double layer is a phenomenon seen at the interface between an electrode and electrolyte in solution where charged species align with the electrode surface. Presumably this layer is present at all surfaces minus pinholes smaller than the double layer thickness. Because steel is not a traditional electrode material for studying surface initiated redox reactions, measurement of this charged layer was performed using a process detailed by Huang [50].



**Figure 2-8. (left) Cyclic voltammograms of stainless steel 304L coupons obtained in 1M KOH at varying scan rates (5-55 mV/sec), compared to a silver/silver chloride reference electrode; (right) Linear fit of current density obtained as a function of voltage scan rate.**

Cyclic voltammograms were used to probe double layer charging via scanning the electrical potential against a silver/silver chloride reference electrode in 1M KOH while recording the resulting current. Electrical potentials were selected to not disrupt the native oxide or induce redox reactions, and multiple scans were performed to allow the resulting charge to stabilize. Scanning the electrical potential across a range of rates (5 –55 mV/sec) increased current with increasing potential which demonstrates a capacitor like behavior. Electrical double layer capacitance was then calculated from the slope of the best fit linear line of scan rate versus current density. Using a modeled value of capacitance for atomically flat steel as a reference, estimated surface area for our native 304L coupons was determined. Ratio of electrochemically to geometrical surface area was estimated to be 1.14. See Table 2-3.

Capacitance of Atomically Flat Steel	0.04	mF/cm <sup>2</sup>
Double Layer Capacitance (DLC)	0.0828171	mF/cm <sup>2</sup>
Calculated Surface Area from DLC	2.0704275	cm <sup>2</sup>
Geometric Surface	1.82	cm <sup>2</sup>
Electrochemical-to-Geometric Surface Area	1.14	

**Table 2-3. Electrochemical-to-geometric surface area ratio calculation generated from electrochemical double layer capacitance measurements.**

## 2.2.6. *Method and Application to Multi-Peak Fitting and Residual Errors of Temperature Programmed Desorption Spectra*

### 2.2.6.1. *Introduction*

Over the course of this LDRD project, we developed various numerical techniques to fit thermal desorption data (experimental and simulated) to extract the material-dependent properties (e.g.,

binding energy, kinetic order of desorption) describing the evolution of the surface desorbate layer [116]. The method presented below assumes that thermal desorption data are taken versus linear temperature ramps,  $T(t)$ , as described by the Polany-Wigner equation:

$$r(\theta) = \frac{-d\theta}{dT} = \frac{1}{\beta} v_0(\theta) \theta^n e^{(-E_{act}/RT)}, \quad (2.3)$$

where  $\theta$  represents the coverage, defined as the number of monolayers (ML) remaining on the surface,  $r(\theta)$  is the removal rate of the layers (ML/s),  $\beta$  is the temperature ramp rate ( $dT/dt$ ),  $v_0$  is the pre-exponential factor of desorption ( $s^{-1}$ ),  $n$  is the kinetic order of desorption taken as integer values (i.e.  $n = 0, 1, 2$ ),  $E_{act}$  is the binding energy (J/mol), and  $R$  is the gas constant (= 8.314 J/mol/K). Often, the activation energy is assumed to have a linear variation with coverage expressed as:

$$E_{act}(\theta) = E_0 - w * \theta, \quad (2.4)$$

where  $E_0$  is activation energy as  $\theta \rightarrow 0$  (J/mol), and  $w$  is the interaction energy coefficient (J/mol/ML). Extraction of properties involves fitting desorption data with the Polany-Wigner equation. We assume the data sets,  $\theta_i(T_i)$ , where  $i$  is the index of data points, have been calibrated (to units of ML/s) and that the temperature ramp is linear (i.e.,  $\beta = \text{const.}$ ). Background subtraction can be performed as part of the analysis or prior to it. To fit a single data-set, we perform least-squares minimization, i.e. we seek to find the vector,  $\vec{X}$ , of best-fit properties (e.g.  $\vec{X} = (E_0, w, v_0)$ ), that minimizes the cost-function,  $C(\vec{X})$ :

$$C(\vec{X}) = \sum_i (\theta(\vec{X}, T_i) - \theta_i(T_i))^2, \quad (2.5)$$

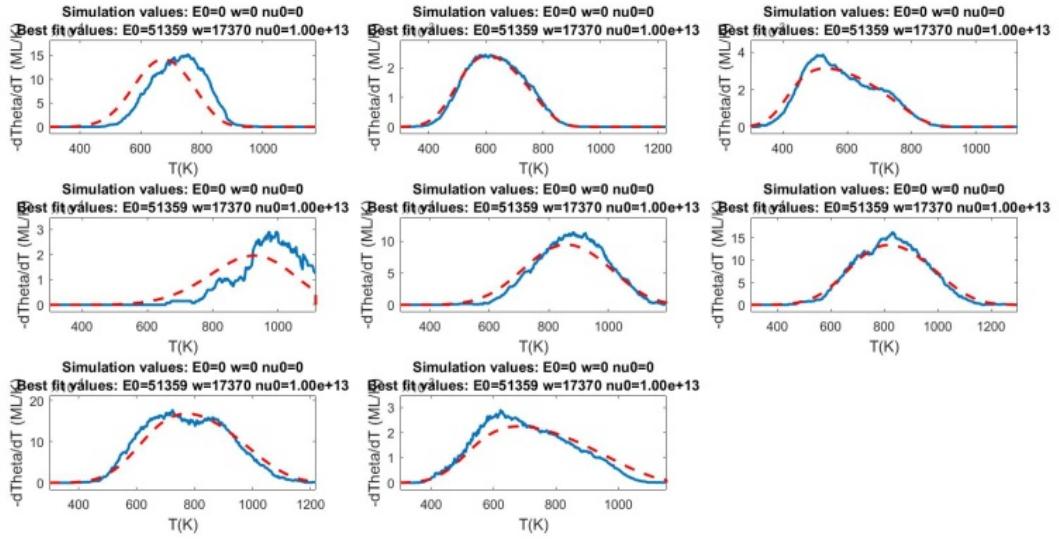
where  $\theta(\vec{X}, T_i)$  is the modeled coverage (# layers) from the Polany-Wigner equation using the vector of fit properties. It is assumed that other properties not included in the vector,  $\vec{X}$ , are specified and held constant. The initial coverage is denoted as  $\theta_0$  (units of ML) and can be specified in advance, treated as a fit parameter, or found by integrating a manually selected portion of the coverage data. The actual minimization is done with a gradient-descent based routine (from MATLAB).

### 2.2.6.2. Demonstration of the Fitting Method

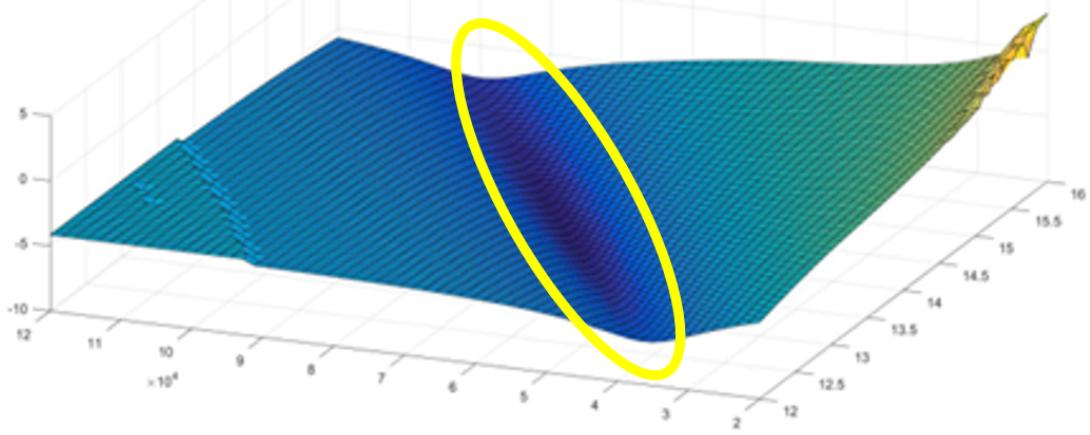
**a. Fitting of Multiple-Data Sets from Molecular Dynamics Simulations:** Here we provide an example of simultaneous fitting of multiple (8) data-sets. The data-sets were generated by Molecular Dynamic (MD) simulation and correspond to a water layer desorbing from  $\alpha$ -hematite with the individual sets being due to different temperature ramps ( $\beta$ ), and different initial coverages ( $\theta$ ). For the fit, it is assumed that the same properties describe each of the data sets and, to fit over the multiple data-sets, the right-hand side of Equation 2.5 is modified to include a second sum over the individual data-sets. The left of Figure 2-9 shows 8 panels of the

raw-data, plotted as the rate of layer removal, i.e.  $r(\theta) = -d\theta/dT$ , and resulting fit for each. In this case, we fit for two free parameters (fitting vector:  $\vec{X} = (E_0, w)$ ) with the minimization providing best fit values:  $E_0 = 51 \text{ kJ/mol}$ ,  $w = 17 \text{ kJ/mol/ML}$  (for fixed  $v_0 = 10^{13} \text{ s}^{-1}$ ).

Ideally we could leave as many of the properties free as possible; however, attempts to also fit for the pre-exponential factor, i.e. to use a fitting vector  $\vec{X} = (E_0, w, v_0)$ , revealed that such a fitting approach is not as robust as desired. More specifically, owing to what we consider as a non-orthogonal (degenerate) nature to the variation of the cost-function with the fit parameters, different combinations of fit parameters yield very similar cost-function minima. To illustrate this point, the right of Figure 2-9 shows the residual, i.e. the value of  $C(\vec{X})$  (summed over the 8 data sets) plotted versus  $v_0$  and  $E_0$ , where for each case the best fit  $w$  was determined. The residual surface does not show a clear minimum point (which would be the case if a unique combination of  $(E_0, v_0)$  gave the best-fit) but rather there are a series of different combinations of  $(E_0, v_0)$  that yield a similar minimum as shown with the yellow ellipse. One possible solution here is to, even if it is just as a reference, to select an agreed upon constant value of  $v_0$  (for consistent use by all involved) and then to fit for  $(E_0, w)$  based on that.



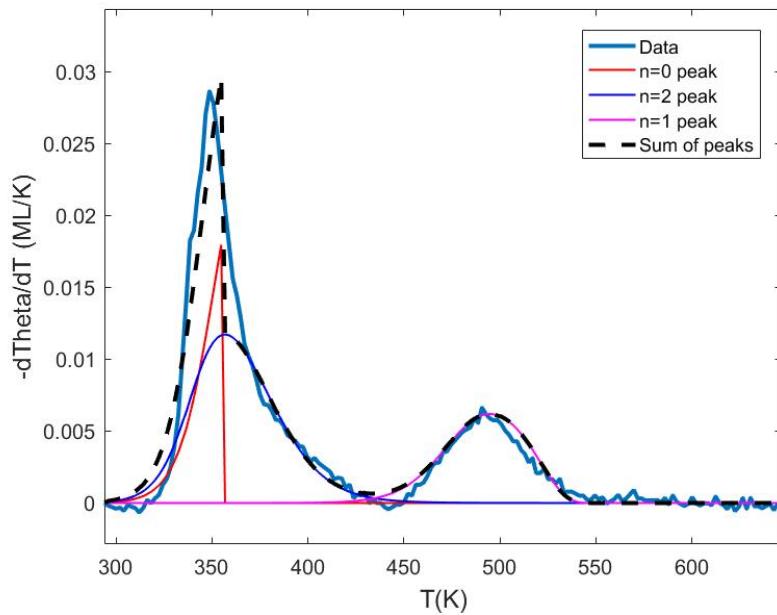
(a)



(b)

**Figure 2-9. a)** Panels of raw and fitted data from MD simulations of a water layer on hematite. **b)** Residual plot showing dependence of residual on  $E_0$  and  $v_0$ . The yellow ellipse illustrates the different  $(E_0, v_0)$  that have similar minima.

**b. Multi-Peak Deconvolution Applied to Experimental Data:** The framework described above can also be used to fit data exhibiting multiple (sometimes overlapped) desorption peaks. In this case, each peak can have its own set of best-fit properties. Each peak is also assumed to work from its own layer inventory. The final dependence of coverage (or rate of layer removal) on temperature is found as the sum of the contributions of individual layers. Figure 2-10 shows an example of a multi-peak fit due to data from a water layer on nickel. Further work is needed to constrain the parameters, but the fit shows reasonable agreement with data when three peaks (of order  $n = 0, 2, 1$  from low to high T) are used.



**Figure 2-10. Example of multi-peak fitting (data is water layer on nickel).**

### 2.2.7. *Ultra-Fast Mass Spectroscopy*

At the onset of this LDRD project, we reviewed the current technology for mass spectrometers and examined their potential with respect to simulating pulsed power applications by temperature programmed desorption with high-speed mass spectrometry (i.e., on the order of  $2000^{\circ}\text{C}/100\text{ ns}$ ) which would allow for direct characterization of flash desorption and comparison to MD simulations on their timescales. The initial intent was to complete a prototype design by the completion of this project. Due to unforeseen circumstances, this portion of the project was removed from the overall scope after an initial design study; however, this work will continue, in-part, under a separate LDRD investment through the Assured Survivability and Agility with Pulsed Power (ASAP) Mission Campaign LDRD. A summary of the initial investigation by Northern Arizona University is provided below [107].

Three commercial-based mass spectrometer designs, i.e., time-of-flight, quadrupole, and magnetic sector mass spectrometer, were evaluated for their potential and limitations with respect to high-speed mass spectrometry. We found that the initial conditions of the ions during

collection, e.g., 3-dimensional position relative to the orifice, point of ionization, trajectory, and initial kinetic energy, pose serious challenges as measurement speed increases. Depending on the exact instrument design, the path variation in critically limits the minimum sampling and prevents any of these designs from achieving nanosecond mass spectral measurements with current technology.

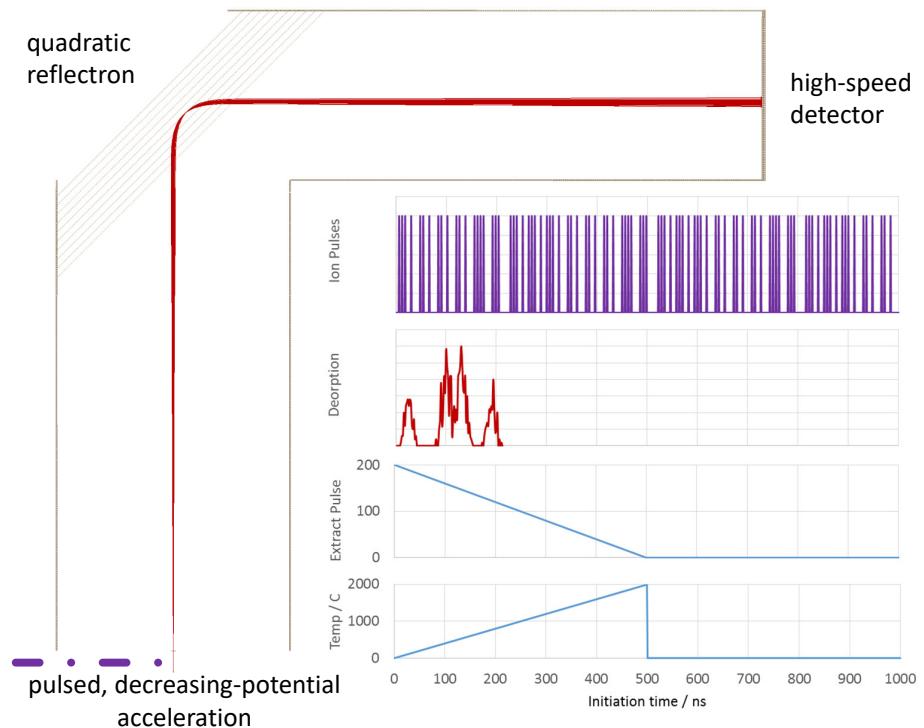
Based on this examination and evaluation, we proposed two additional potential mass spectrometer designs that use laser pulses for ionization to decouple desorption time from the analysis and enable high-speed mass spectrometry. The first of these is based on a Hadamard transform time-of-flight mass spectrometry and uses a pseudorandom pulse pattern to encode the high-speed desorption, and momentum modulation to spread the signal over a measurable time base. The second option is based on a design from the literature that measures one data point per heating ramp. This second design could work with a time-of-flight spectrometer or another type depending on sensitivity. Either of these designs that decouple desorption time from measurement should be capable of nanosecond mass spectrometry measurements with current technology.

Of these, the *Decreasing Amplitude Acceleration, Multiplexed Time-of-Flight Mass Spectrometry (DAAM-TOF-MS)* appeared most promising. This design combines momentum modulation and linear mass spreading used in constant momentum acceleration instruments and multiplexing of pseudorandom pulses used in Hadamard transform instruments. Constant momentum acceleration in TOF-MS uses an acceleration pulse that ends before any ions have left and a quadratic potential reflectron to spread masses (in a select mass range) linearly in detection time. Hadamard transform uses a pseudorandom series of ion pulses to sample a surface more rapidly; the pulse sequence allows deconvolution of overlapping spectra. A conceptual representation is provided in Figure 2-11 and simulated operation and deconvolution are presented in Figure 2-12.

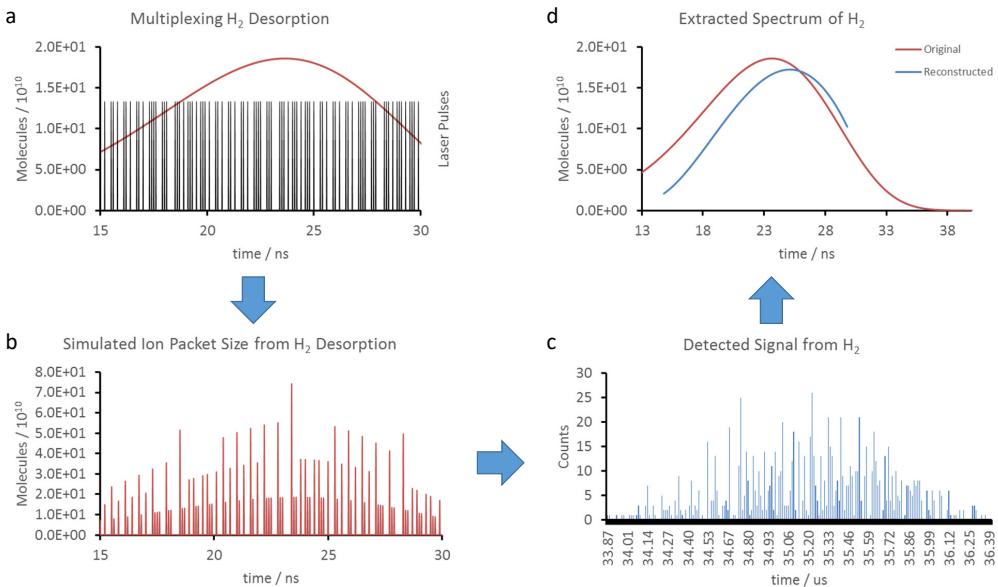
### **2.2.8. *Evaluation of a Potential High Speed Mass Spectrometer, "The E-Vader"***

Using supplemental funding awarded to the GC LDRD by the ASAP Mission Campaign LDRD, we purchased a novel mass spectrometer device and collaborated with Colorado State University to evaluate its overall performance and potential use as a high-speed mass spectroscopy device. The research team led by Dr. John Williams at CSU has been conducting research in the area of plasma diagnostics using electrostatic and magnetic analyzers that yield ion energy and ion charge state distributions, respectively. Recently a small company has combined these two analyzer functions into a compact device (E-Vader) such that simultaneous energy and charge state data can be obtained. Figure 2-13 depicts the actual E-Vader device and its various sub-assemblies.

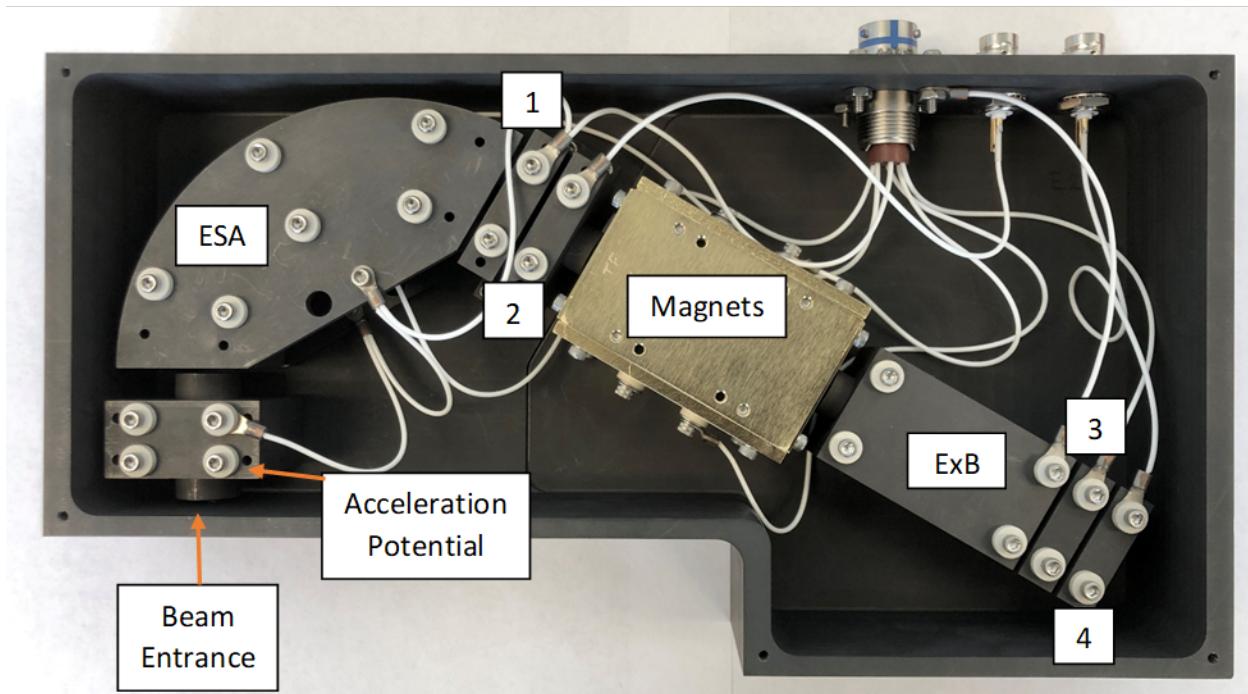
Over the past several months, CSU has been utilizing the E-Vader to collect data on a unique plasma source known as the flamethrower, which produces intense, multi-charged ions in a very localized region that expand from this region to a location where the probe was located. The plasma source can be operated in a steady state or pulsed manner, which allows one to study the expanding plasma in a controlled and systematic manner. The focus of this project was to characterize the ability of the compact energy and charge state analyzer to measure the steady-state ion characteristics and then the temporal ion characteristics of the ions expanding from the highly localized plasma source. The overarching goal of this work was to uncover the



**Figure 2-11. Concept for Decreasing Amplitude Acceleration, Multiplexed Time-of-Flight Mass Spectrometry.** The sample is heated over 500 ns to 2000° C. Over this temperature range,  $H_2$ ,  $H_2O$ , and  $CO$  are expected to desorb from the surface. A laser provides a pseudorandom sequence of 100 ps pulses which multiplex the time data into the pulse-train. As soon as an atom or molecule is ionized, it is acted upon by the acceleration potential. The acceleration potential decreases in time over the duration of the sample heating and reaches zero or a steady-state at the end of heating. Ions are reflected in a quadratic potential reflectron to separate the ions linearly in time based on mass. Each mass can then be deconvolved based on the pseudorandom sequence to find the initial desorption time.



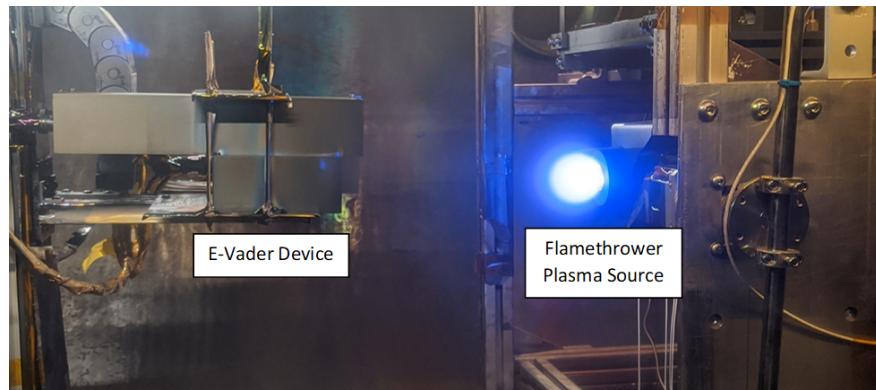
**Figure 2-12. Simulated data from DAAM-TOF-MS for 15-30 ns of  $H_2$  desorbing from stainless steel.** a) A random sequence determines whether the ionizing laser fires or not every 100 ps. b) Each pulse ionizes all of the molecules that have desorbed since the previous pulse. With longer time between pulses, the ion count increases for each pulse. c) The ions from the same mass molecule are imparted momentum that decreases as the sample is heated, so the flight times get spread over 2.5 microseconds. This allows for easier detection and potential detection of all ions that are accelerated. d) Using the pseudorandom sequence the curve can be reconstructed to produce a curve within a few nanoseconds, corresponding to 5-10 degrees, of the original desorption curve.



**Figure 2-13. Inside of the E-Vader device which depicts the combined analyzer functions of the ESA and ExB; ions enter and immediately have the same acceleration potential applied; the ions then travel through the ESA segment and then reach plate 1 which is the ESA ion collimator; ions travel through to plate 2 which also serves as an ESA collector and the ExB collimator; the ions then undergo velocity separation through the permanent magnets and then enter the ExB drift tube; finally at the end of the tube, plate 3 acts as an ExB electron suppressor and plate 4 as the ExB collector.**

fundamental limitations of the temporal resolution of the E-Vader device and systematically minimize their effects on ion energy and charge state measurements such that the device could collect highly resolved temporal information about the energy and mass of particles liberated from a highly unstable plasma source or from the impact of a very short-duration, high-power laser pulse on a surface placed in a vacuum chamber.

At the time of this report, only steady-state data has been collected using the E-Vader to demonstrate that ion energy and  $m/e^-$  can be consistently and accurately characterized. The E-Vader probe was aligned with the plasma source orifice. Subsequently, the data collected using the ESA and ExB correctly identified the Krypton ions. The configuration used for the initial steady-state experiments with the "Flamethrower" is shown in Figure 2-14.



**Figure 2-14. The E-Vader probe is aligned with the flamethrower orifice that is collecting steady-state data of the Krypton plasma.**

During the upcoming period, CSU aims to characterize the fast-transient performance of the E-Vader by rapidly turning off-and-on a plasma source providing the impulse response. This allows for the evaluation of this device as a candidate technology for future designs of "Ultra-fast Mass Spectroscopy" which is discussed in more detail below.

### **2.2.9. *Technical Gaps and Future Directions***

Several significant advancements were accomplished in this project, however technical gaps remain. Future experiments and characterization of "real" electrodes should include:

1. developing advanced diagnostics for flash desorption at timescales closer to regimes directly applicable to pulsed power accelerators;
2. continued coordination with molecular simulations for dissociated and non-dissociated adsorbates and comparison to experiments;
3. characterizing the influence of high-electric fields on the kinetics of desorption;
4. and an in-depth characterization of desorption through melt of the electrode surfaces.

Completion of these various activities serve to improve the electrode boundary conditions used in various simulations, and would provide new tools and capabilities which can be leveraged in the development of next-generation pulsed power accelerators.

### **2.2.10. Concluding Remarks**

Over the course of this Grand Challenge LDRD, the Desorption Theory & Experimental Validation teams have diligently worked to develop new models and experimental capabilities to help assess the desorption kinetics of contaminants released from rapidly heated electrode surfaces. The Desorption Theory team employed a variety of simulation tools (such as DFT, MD, and GCMC) to help understand: the water and hydrocarbon contaminant inventories on electrode surfaces at  $10^{-5}$  Torr; the mechanisms for trapped hydrogen to diffuse through the metal and oxide film, which releases  $H_2$  into the vacuum; and hydrogen reactions with surface oxides and hydroxyls resulting in the formation of water. Likewise, the Desorption Experimental Validation team employed a variety of experiments and analytical methods to: characterize the composition and initial inventories of adsorbates on stainless steel samples; improve the desorption models informed by measurements of "real" materials; and develop new capabilities to validate the Polanyi-Wigner desorption model for fast thermal ramp rates approaching those expected in large pulsed power accelerators. Despite the great progress across several fronts, we believe a considerable amount of research is still necessary to fully understand the desorption kinetics of complex electrode materials found in real pulsed power devices.

## **2.3. Experimental Collaboration with University of Michigan**

As part of this LDRD project, Sandia scientists collaborated with the Prof. Ryan McBride at the University of Michigan, a participating institution in the SNL Campus Executive program. The primary goal of this collaboration was to develop innovative experimental and computational power flow platforms and to acquire new experimental data on UM's 1-MA MAIZE facility. This collaboration was successful in establishing several new capabilities, described in detail in the accompanying project reports [122, 76], while only select results are briefly summarized here.

One key objective for this project was to increase the maximum current of the MAIZE facility, which would allow a greater range of power flow stresses to be applied to scaled experiments on this platform. MAIZE is a linear transformer driver (LTD) cavity [77] consisting of 40 bricks, with each brick an assembly of two high-energy-density capacitors and a triggered gas switch. During this project, each 40-nF General Atomics capacitor was replaced with a 80-nF NWL capacitor; this resulted in a doubling of the energy storage capacity of the machine and an increase in peak current achievable by a factor  $\sqrt{2}$ . In addition, all of the triggered gas switches were replaced with 200-kV low-inductance switches manufactured by L-3 Harris. These upgrades have been implemented and the machine is performing well; we expect to leverage this capability to provide scaled power flow validation data over the coming years.

Several initial power flow studies have been investigated through this initiative. An additively manufactured platform for studying electrode plasma formation while allowing side-on diagnostic

access has been developed; this platform was inspired by strip-line loads utilized by the dynamic materials program at the Z-Machine. This platform has been exercised in an initial set of experiments, where high linear current densities (0.5-0.7 MA/cm) flowing through thin metal foils can cause electrode plasmas to form, which are observed using multi-frame imaging. A detailed description of these first results are planned for an upcoming publication [123]. Finally, a vacuum ultraviolet (VUV) spectroscopy system has been designed and is planned for future implementation at MAIZE; these activities are expected to be supported through post-LDRD investments provided by the SNL pulsed power science program.

## 3. ENGINEERING CODE DEVELOPMENT

### 3.1. EMPIRE

This section focuses on the EMPIRE code and the modifications required to be completed to make this general purpose code suitable for modeling pulsed power phenomena such as vacuum power flow. EMPIRE was designed to be highly performant for simulation of plasma in complex geometries. However, power flow exploration requires specialized modeling capabilities which were added.

Electrical current delivery on large pulsed power accelerators typically begins with energy stored in capacitor banks; these are triggered and the electrical energy pulse is compressed in both space and time until it reaches vacuum regions where plasma is formed and nonlinear effects begin to become critical. In the Z-Machine this region is known as the “inner MITL”, here the energy density (and corresponding electric fields) in this region becomes so high that electrons are emitted from the metal electrode, rapid heating causes material trapped in the metal to be released, and the resulting inventory of material can become ionized creating a high density plasma within the MITL gap. This plasma can reduce electrical power flow efficiency, it can reflect power flowing into the MITL, it can create a conductive path between the anode and the cathode, or it can deposit enough energy to damage the MITL. Modeling these phenomena is a critical requirement in the design of next-generation pulsed power accelerators.

EMPIRE was not designed to model MITLs so several application-specific features were added and verified through this Grand Challenge LDRD initiative, with the base model of EMPIRE developed through ASC/ATDM program support. Specific MITL models added (and now maintained) in the current production version of EMPIRE include:

1. Transmission-line circuit models to simulate the upstream region of the accelerator;
2. Electrode heating and material desorption models;
3. Implicit algorithms to overstep plasma and cyclotron frequencies; and
4. Relativistic fluid representation and supporting boundary models.

The details of these algorithmic improvements are provided in the subsequent sections, including a brief background of the EMPIRE code.

### 3.1.1. *Background*

The ElectroMagnetic Plasma In Realistic Environments (EMPIRE) code is an effort to consolidate and expand on the features available to Sandia in several legacy applications while harnessing the performance potential of next generation computational platforms. EMPIRE was started in earnest in 2015 under the ASC Advanced Technology Demonstration and Mitigation (ATDM) portfolio to provide the ability to model plasma across a range of densities by developing a fully kinetic representation, a fluid representation, and a hybrid approach to bridge those two representations. To make this goal more tractable we made the following assumptions:

1. The kinetic representation would only handle low densities with moderate to low magnetic fields;
2. The fluid representation would have a low bulk velocity and temperature;
3. The hybrid capability will assume a small positive kinetic perturbation on the overall thermalized plasma;
4. The fluid and kinetic perturbation will only be coupled through collisional dynamics and EM fields;

This allows the EMPIRE code to treat the speed of light as the most stringent time-scale, and therefore only require an implicit EM field solve and not implicitly coupled flow properties.

### 3.1.2. EMPIRE-PIC

PIC algorithms are commonly used for the simulation of low density plasma dynamics. This is accomplished by numerically solving the Klimontovich equation that governs the time evolution of a given plasma [86]. A collection of  $N_p$  discrete particles is used to represent the plasma, each particle has an associated position  $\vec{x}$ , velocity  $\vec{v}$ , charge  $q$ , and mass  $m$ .

$$f = \sum_{i=1}^{N_p} f_i = \sum_{i=1}^{N_p} \delta(\vec{x} - \vec{x}_i) \delta(\vec{v} - \vec{v}_i) \quad (3.1)$$

The six-dimensional phase space location for each particle is updated via Newton's Law and the Lorentz force equation, where  $\vec{B}$  and  $\vec{E}$  are the magnetic and electric fields, respectively.

$$\frac{d\vec{x}_i}{dt} = \vec{v}_i \quad (3.2)$$

$$\frac{d\vec{v}_i}{dt} = \frac{q_i}{m_i} \left( \vec{E}(\vec{x}_i) + \vec{v}_i \times \vec{B}(\vec{x}_i) \right) \quad (3.3)$$

These equations can then be assembled into the Klimontovich equation for collisionless particle dynamics [145, 36]:

$$\frac{\partial f_i}{\partial t} + \frac{\vec{v}_i}{m} \cdot \nabla f_i + \frac{q_i}{m_i} \left( \vec{E}(\vec{x}_i) + \vec{v}_i \times \vec{B}(\vec{x}_i) \right) \frac{\partial f_i}{\partial \vec{v}} = 0 \quad (3.4)$$

With initial magnetic and electric fields, this system can be used to fully describe the particle evolution, and Maxwell's equations can be used to couple the charged particle motion to the electric and magnetic fields. They consist of the following: Gauss' Law, the magnetic divergence constraint, Faraday's Law, and Ampère's Law. For clarity, these equations are given below in the form of differential equations. Here  $\rho$  and  $\vec{J}$  are the charge and current densities, and  $\epsilon_0$  and  $\mu_0$  are the permittivity and permeability of free space, respectively.

$$\nabla \cdot \vec{E} = \frac{\rho}{\epsilon_0} \quad (3.5)$$

$$\nabla \cdot \vec{B} = 0 \quad (3.6)$$

$$\frac{\partial \vec{B}}{\partial t} = -\nabla \times \vec{E} \quad (3.7)$$

$$\frac{\partial \vec{E}}{\partial t} = \frac{1}{\mu_0 \epsilon_0} \nabla \times \vec{B} - \frac{1}{\epsilon_0} \vec{J} \quad (3.8)$$

Finally, the particles can be coupled back to Maxwell's equations via the charge and current densities defined below.

$$\rho = \sum_{i=1}^{N_p} q_i f_i \quad (3.9)$$

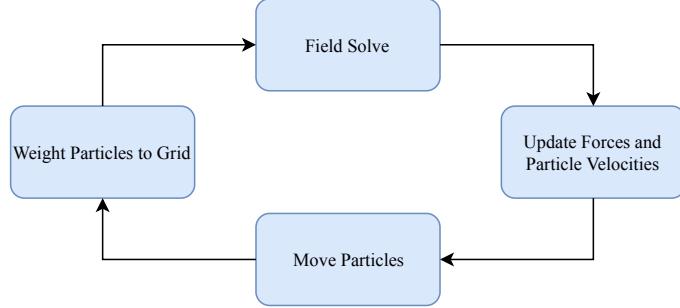
$$\vec{J} = \sum_{i=1}^{N_p} q_i \vec{v}_i f_i \quad (3.10)$$

### 3.1.2.1. The Particle-in-Cell Method

Introduced by Birdsall and Dawson, the PIC method is a well established procedure for modeling the behavior of charged particles in the presence of electric and magnetic fields [10, 33]. Discrete particles are tracked in a Lagrangian frame, while the electric and magnetic fields are stored on stationary points on a fixed Eulerian mesh. Therefore, the algorithm can be thought of as two coupled solvers where one is responsible for updating the electric and magnetic fields, and another updates the particles via the method of characteristics and calculates their charge/current contributions back to the grid. These are referred to as the field solver and the particle mover (sometimes called the particle pusher), respectively. Combining these solvers results in the main time loop of the core PIC algorithm that is composed of four key steps, summarized in Figure 3-1. Clearly, even on the largest computers, we cannot represent every electron an ion, therefore, we model macro particles, each representing a large number of physical particles.

### 3.1.2.2. Solving Maxwell's Equations

In order to obtain the values of the electric and magnetic fields it is necessary to solve Maxwell's equations. EMPIRE has an electrostatic approximation formulation, but that is not sufficient in the context of power flow. We now show in detail the formulation of the electromagnetic problem used in EMPIRE-PIC such that they can be solved via the Finite Element Method (FEM) [55] in



**Figure 3-1. Flow chart summarizing the key components of the PIC algorithm**

space, while finite difference approaches are applied to discretize in time. A unique aspect of the transient Maxwell's equations is the addition of involution conditions satisfied for all time (see Equations (3.5) and (3.6)). To this end, in 3D we use the lowest order Nédélec [85] elements  $\vec{e} \in V_{Curl} \subset H_{Curl}$  (so-called edge elements), and Raviart-Thomas [15] elements  $\vec{b} \in V_{Div} \subset H_{Div}$  (so-called face elements) to approximate the solution of the electromagnetic fields. This discretization ensures that Gauss' law is maintained in the weak form and the divergence constraint on the magnetic field is maintained in the strong form.

$$\frac{\partial B}{\partial t} = -CE \quad (3.11)$$

$$M_E \frac{\partial E}{\partial t} = K_E B - J_{vec} \quad (3.12)$$

where  $B$  and  $E$  represent discrete vectors of coefficients, and

$$M_E = \int_{\Omega} \hat{e}_i \cdot \hat{e}_j d\Omega \quad (3.13)$$

$$K_E = -\frac{1}{\mu_0 \epsilon_0} \int_{\Omega} \hat{b}_j \cdot \nabla \times \hat{e}_i d\Omega \quad (3.14)$$

$$J_{vec} = \int_{\Omega} \vec{J} \cdot \vec{e}_i d\Omega \quad (3.15)$$

$$M_B = \int_{\Omega} \hat{b}_i \cdot \hat{b}_j d\Omega \quad (3.16)$$

$$C = \nabla \times \vec{e}_i \quad (3.17)$$

Here the matrix  $C$  is an edge orientation matrix.

Considering the discretization of the time derivative for Faraday's and Ampère's Laws, EMPIRE-PIC includes backward Euler, Crank-Nicolson (C-N), and Friedman [39] time integration schemes. While Friedman was used in the bulk of the runs, we use C-N for ease of understanding. For simplicity we present the C-N formulation for first-order Maxwell's equations:

$$B^{n+1} + \frac{\Delta t}{2} K_B E^{n+1} = B^n - \frac{\Delta t}{2} C E^n \quad (3.18)$$

$$M_E E^{n+1} - \frac{c_0^2 \Delta t}{2} K_E B^{n+1} = M_E E^n + \frac{c_0^2 \Delta t}{2} K_E B^n - \frac{\Delta t}{\epsilon_0} J^{n+1/2} \quad (3.19)$$

These matrix equations can be solved using a multilevel solver, we use the RefMaxwell path in Trilinos' MueLu package on top of PCG method. Here, we express the current as  $J^{n+\frac{1}{2}}$ , a vector form of  $\frac{1}{\epsilon_0} \int_{\Omega} \vec{J} \cdot \hat{e}_i d\Omega$ . If we assume that charge is represented by a point delta function with some charge  $q$ , we can define the current as shown in Equation (3.20).

$$J_{vec} = \frac{1}{\epsilon_0} \int_{\Omega} \vec{J} \cdot \hat{e}_i d\Omega \quad (3.20)$$

$$= \frac{1}{\epsilon_0} \int_{n\Delta t}^{(n+1)\Delta t} \int_{\Omega} q \vec{u}(t) \delta(\vec{x}(t)) \cdot \hat{e}_i d\Omega dt \quad (3.21)$$

$$= \frac{1}{\epsilon_0} \int_{n\Delta t}^{(n+1)\Delta t} q \vec{u}(t) \cdot \hat{e}_i(\vec{x}(t)) dt \quad (3.22)$$

## Weighting of Fields to Particles

After solving for the field values at the nodes of the simulation grid, we require a method of determining the value of the fields at any specific particle position. In order to do so we weight the field values from the grid nodes to the required location. Assuming that we know the values of the electric and magnetic fields we can evaluate the value of the fields at a given point as shown below by applying the basis function.

$$\vec{E}(\vec{x}_i) = \sum_{j=0}^{N_{edge}} E_j \hat{e}_j(\vec{x}_i) \quad (3.23)$$

$$\vec{B}(\vec{x}_i) = \sum_{j=0}^{N_{face}} B_j \hat{b}_j(\vec{x}_i) \quad (3.24)$$

## Particle Mover

The force felt by a charged particle due to the presence of electric and magnetic fields is described by the Lorentz force equation, shown in Equation (3.25). The particle mover within EMPIRE-PIC is responsible for solving for this force on each particle within the simulation, and subsequently updating the particle velocities and positions. In our method we make use of the well-known Boris algorithm to handle the acceleration due to the electric field, and rotation about the magnetic field [13].

$$\vec{F} = q \left( \vec{E} + \vec{v} \times \vec{B} \right) \quad (3.25)$$

To make the algorithm relativistic we define  $\vec{u} = \gamma \vec{v}$ , where  $\gamma$  is the Lorentz factor. Then the velocity update satisfies

$$\frac{\vec{u}^{n+1/2} - \vec{u}^{n-1/2}}{\Delta t} = \frac{q}{m} \left[ \vec{E}^n + \frac{\vec{u}^{n-1/2} + \vec{u}^{n+1/2}}{2\gamma^n} \times \vec{B}^n \right] \quad (3.26)$$

The Boris method separates the the electric field update from the magnetic rotation to keep the updates linear.

Size	# of Elements	# of Nodes	# of Edges	# of Faces	Particle Count	Particles/element
S	337.0 k	60.0 k	406.0 k	683.0 k	16.0 M	47.5
M	2.7 M	462.0 k	3.2 M	5.4 M	128.0 M	47.8
L	20.7 M	3.5 M	24.4 M	41.6 M	1.0 B	49.5
XL	166.0 M	27.9 M	195.0 M	333.0 M	8.2 B	49.4
XXL	1.3 B	223.0 M	1.6 B	2.7 B	65.6 B	49.2

**Table 3-1. Details of problem sizes used to test EMPIRE-PIC.**

Machine	Nodes	Processor	Accelerator	Compiler(s)
Trinity (Haswell)	9436	2 × Intel Xeon E5-2698v3	-	Intel 18.0.5
Trinity (KNL)	9984	1 × Intel Xeon Phi 7250	-	Intel 18.0.5
Astra	2592	2 × Cavium Thunder-X2 CN9975	-	GCC 7.2.0
Sierra	4340	2 × IBM POWER9 22C	4 × NVIDIA V100	GCC 7.2.0 NVCC 9.2

**Table 3-2. Details of the systems used to collect EMPIRE-PIC performance data.**

### Weighting of Particles to Grid

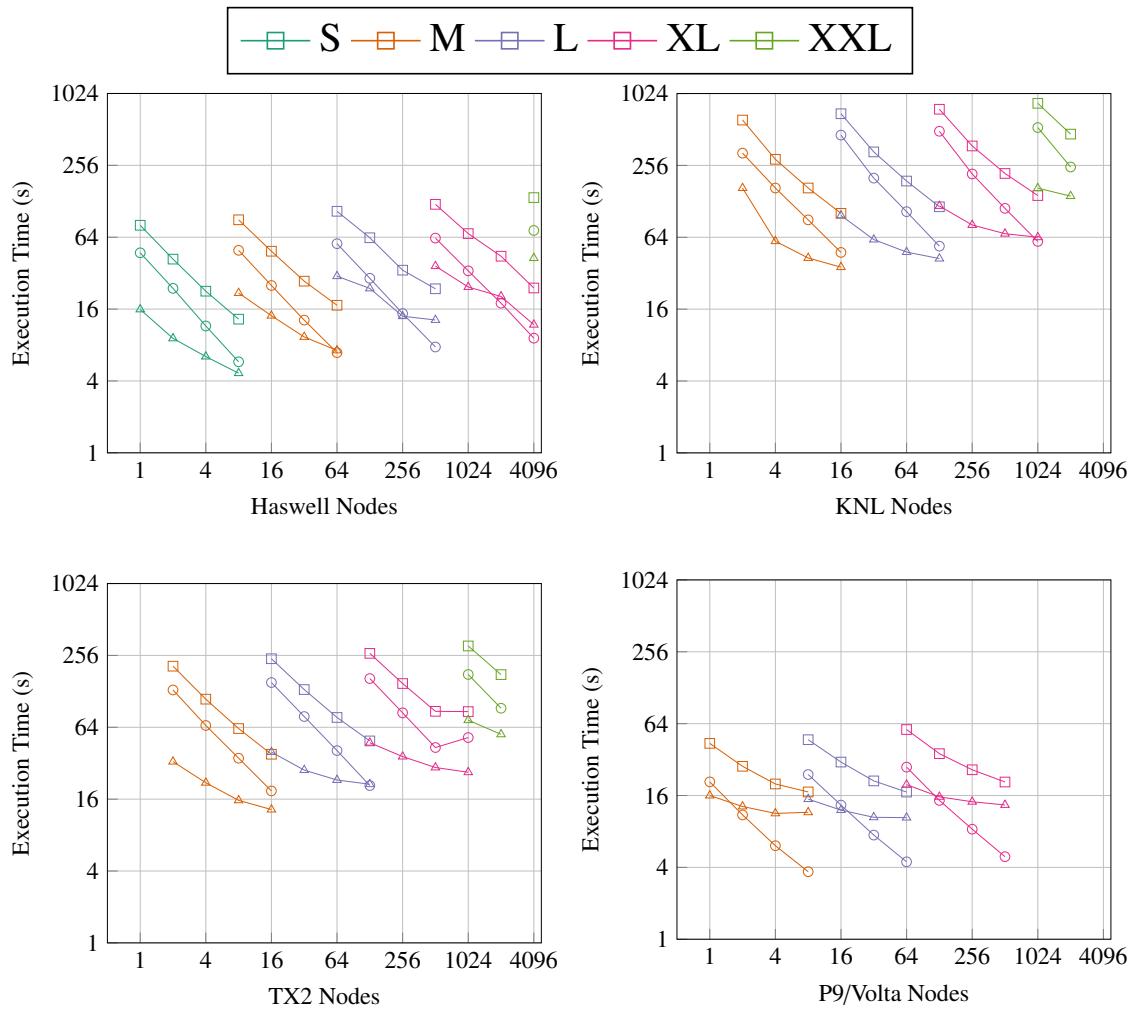
During each time step in the PIC algorithm we must weight the contributions of each particle back onto the grid before we can commence the next field solve. An electromagnetic simulation requires us to evaluate the current Equation (3.20) as shown in Equation (3.27). For simplices, it is sufficient to use a midpoint rule for the integration; however, for higher-order elements we must evaluate this temporal integration with higher-order numerical cubature.

$$\int_{\Omega_j} \vec{J} \cdot \hat{e}_i \, dV = \sum_{k=1}^{N_p} \int_{\Omega_j} \frac{1}{\Delta t} \int_{n\Delta t}^{(n+1)\Delta t} q_k \vec{u}_k \cdot \hat{e}_i \, dV = \sum_{k=1}^{N_p} \Delta t q_k \vec{u}_k \left( \vec{x}_k^{n+1/2} \right) \cdot \hat{e}_i \left( \vec{x}_k^{n+1/2} \right) \quad (3.27)$$

#### 3.1.2.3. Performance Results

We perform routine testing to insure a constant improvement of performance and no regressions. We have a very simple “blob” problem, which consists of a cavity with a quasi-neutral plasma that is allowed to osculate. This approach has the advantage of testing the performance at a large scale in a short time, typically less than an hour. The sizes for the problem are displayed in Table 3-1 and were completed on several different platforms as shown in Table 3-2.

The timing results are given in Figure 3-2. One can see that the code has been run more than 2B grid unknowns and more than 65B particles and is performant across a wide range of systems. The NVidia system is the most performant for these problems, roughly eight times as fast as the KNL hardware, and 4x faster than the more traditional Haswell node.



**Figure 3-2.** EMPIRE-PIC scaling study results for both partitions of Trinity, Astra, and Sierra. Squares, triangles, and circles represent the main time loop, linear solve, and particle kernels respectively. Problem sizes range from small to XXL.

### 3.1.2.4. Thermal desorption boundary condition

The thermal desorption boundary condition injects a temperature-dependent influx of mass on a specified side set. The temperature is computed by EMPIRE's surface temperature diagnostic, which can compute the surface temperature due to Ohmic heating in a manner that is consistent with eddy currents due to the computed magnetic fields. In EMPIRE-PIC, heating due to energetic particles incident on the surface can also be accounted for.

A 1D heating model is applied to each boundary mesh face to compute the surface temperature needed for the thermal desorption boundary condition,

$$\frac{\partial u}{\partial t} = Q(x, t) - \frac{\partial q}{\partial t} \quad (3.28)$$

$$q = -k \frac{\partial T}{\partial x} \quad (3.29)$$

$$T(u) = \int_0^u \frac{du'}{c(u')} \quad (3.30)$$

$$q(x = 0) = 0 \quad (3.31)$$

$$u(x, t = 0) = u_0 \quad (3.32)$$

$$\lim_{x \rightarrow \infty} u(x, t) = u_0 \quad (3.33)$$

Where  $u$  is the internal energy,  $Q$  is the volumetric heating source,  $q$  is the heat flux,  $k$  is the thermal conductivity and  $c$  is the heat capacity. This is a non-linear system of equations and several solver and model options are available to tune for difficult systems that include physics like melting of the heated material. A similar non-linear set of equations govern the diffusion of the magnetic field into temperature-dependent material,

$$\frac{\partial B_z}{\partial t} = \frac{1}{\sigma\mu} \frac{\partial^2 B_z}{\partial x^2} - \frac{1}{(\sigma\mu)^2} \frac{\partial(\mu\sigma)}{\partial x} \frac{\partial B_z}{\partial x} \quad (3.34)$$

$$B_z(x = 0, t) = B(t) \quad (3.35)$$

$$B_z(x, t = 0) = 0, \quad (3.36)$$

where  $\mu$  is the permeability of the material. This surface temperature calculation is technically part of EMPIRE's diagnostics, and the thermal desorption boundary requires the name of an appropriate surface temperature diagnostic.

The thermal desorption flux is specified by providing values for the coefficients of a generalized Arrhenius equation:

$$\Gamma(T) = v T^\alpha \theta^n \exp\left(\frac{-E_a}{k_B T}\right), \quad (3.37)$$

where  $T$  is the surface temperature and  $k_B$  is the Boltzmann constant. The parameters  $\alpha$ ,  $n$ ,  $v$  and  $E_a$  are specified by the user. The quantity  $\theta$  is the surface density of the adsorbed substance. The surface density diminishes as desorption proceeds according to,

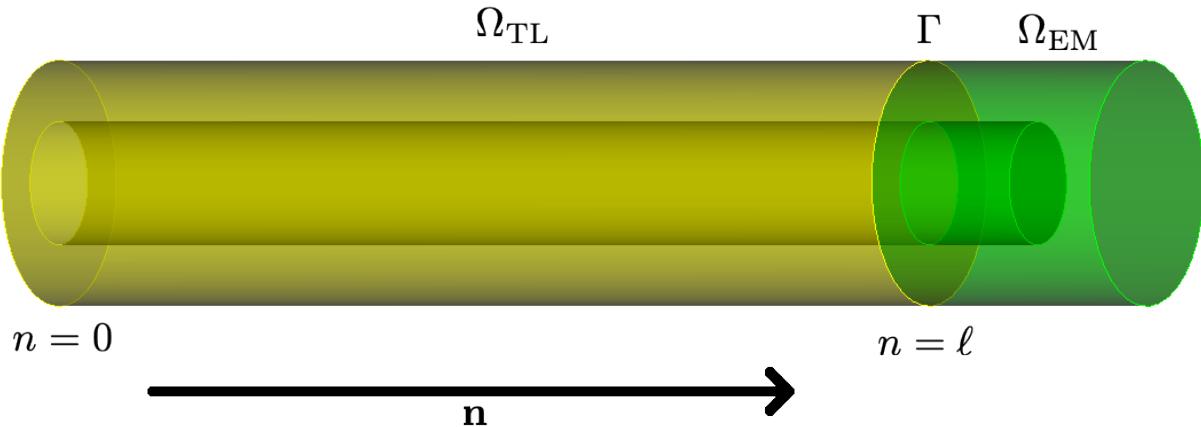
$$\dot{\theta} = -\Gamma(T). \quad (3.38)$$

The total amount of desorption possible is therefore controlled by the initial value of  $\theta$ , which is specified by the user. Note that the user may disable the surface density integration, so that this boundary may be used for other types of thermal emission such as thermionic emission.

### 3.1.2.5. Transmission Line Power Coupling

To enable power flow calculations in EMPIRE a reduced model for the large simple regions of the driver is a practical necessity. We focused on modernizing the transmission line coupling technique used in the QUICKSIVLER code [89].

We consider the domain  $\Omega = \Omega_{\text{TL}} \cup \Omega_{\text{EM}}$ , where  $\Omega_{\text{TL}}$  will be modeled as a transmission line and  $\Omega_{\text{EM}}$  will be modeled with full 3D electromagnetics. The transmission line is long and has uniform cross-section  $\Omega_{\text{TL}} = [0, \ell] \times \Gamma$ . The electromagnetic domain could be complex but shares an interface with the transmission line domain congruent to the cross section  $\Omega_{\text{EM}} \cap \Omega_{\text{TL}} = \Gamma$ . See Figure 3-3 for an example geometry and decomposition of the domain. In the EM domain the



**Figure 3-3.** A prototypical EM-TL domain, in this case a coaxial transmission line terminating in a simple diode. In this case  $\Gamma$  is a circular annulus. The normal ( $\vec{n}$ ) of the surface  $\Gamma$  is the propagation direction of the transmission line.

electromagnetic fields are governed by Maxwell's equations:

$$\begin{cases} \epsilon \frac{\partial \vec{E}}{\partial t} - \nabla \times \mu^{-1} \vec{B} + \vec{J} = \vec{0} & \text{Maxwell-Ampere Law,} \\ \frac{\partial \vec{B}}{\partial t} + \nabla \times \vec{E} = \vec{0}, & \text{Faraday's Law} \\ \nabla \cdot \epsilon_0 \vec{E} - \rho = 0, & \text{Gauss' Law} \\ \nabla \cdot \vec{B} = 0, & \text{Gauss' Law for Magnetism.} \end{cases} \quad (3.39)$$

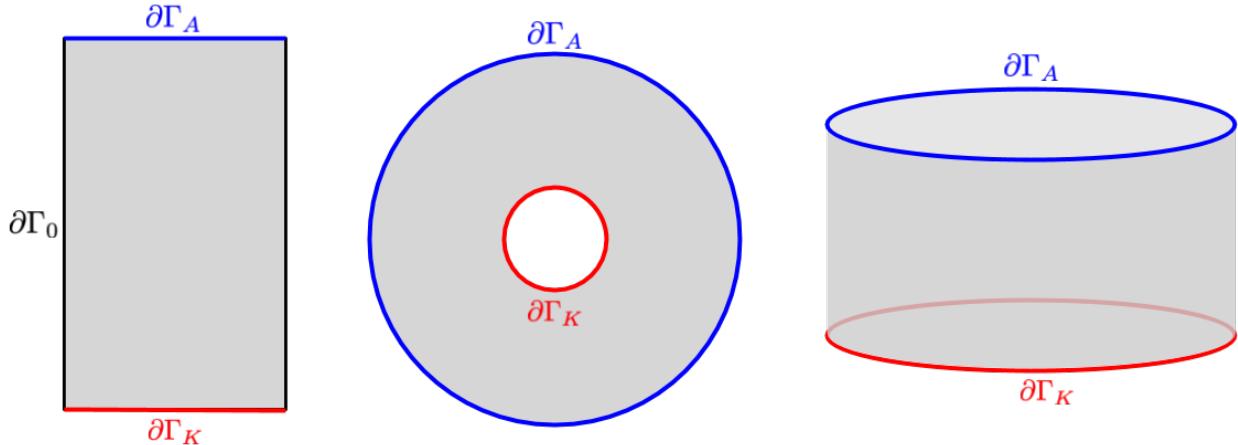
In the TL domain we assume a transverse electromagnetic (TEM) ansatz.

$$\vec{E} = \vec{E}_0(\tau_1, \tau_2)V(n, t), \quad \vec{H} = \vec{H}_0(\tau_1, \tau_2)I(n, t), \quad \vec{E}_0 \perp \vec{H}_0 \perp \vec{n}, \quad (3.40)$$

The voltage  $V$  and current  $I$  have units of V and A respectively while  $\vec{E}_0$  and  $\vec{H}_0$  are the electric and magnetic field responses to a unit voltage or current respectively. Both  $\vec{E}_0$  and  $\vec{H}_0$  have units of  $\text{m}^{-1}$ . The unit electric field  $\vec{E}_0$  is a constant electrostatic profile in the cross-section satisfying

$$\begin{cases} \vec{E}_0 + \nabla_\tau \varphi_0 = 0, \\ \nabla \cdot_\tau \epsilon_0 \vec{E}_0 = 0, \\ \varphi_0|_{\partial\Gamma_K} = 0, \\ \varphi_0|_{\partial\Gamma_A} = 1, \\ \vec{E}_0 \cdot \vec{\xi}|_{\partial\Gamma_0} = 0. \end{cases} \quad (3.41)$$

Here  $\partial\Gamma$  is the union of the anode  $\partial\Gamma_A$ , cathode  $\partial\Gamma_K$  and insulating boundary  $\partial\Gamma_0$ . The vector  $\vec{\xi}$  is the unit normal to  $\partial\Gamma$ . See Figure 3-4 for illustrations of typical cross-section geometries and the corresponding anodes, cathodes, and insulating boundaries.



**Figure 3-4. Several example  $\Gamma$  surfaces (grey) with  $\partial\Gamma_A$  (blue),  $\partial\Gamma_K$  (red),  $\partial\Gamma_0$  (black) labeled. Left shows the cross section of a parallel plate transmission line, middle shows a coaxial transmission line, and right shows a radial line. In the left and middle illustrations  $\vec{n}$  points into the page while for right  $\vec{n}$  points towards to center of the cylinder.**

To enforce continuity of voltage and current, the electric and magnetic fields are related to the voltages and currents at the interface between  $\Omega_{\text{TL}}$  and  $\Omega_{\text{EM}}$  according to the following relations:

$$\begin{cases} CV = \int_\Gamma \epsilon (\vec{E} \times \vec{n}) \cdot (\vec{E}_0 \times \vec{n}) da, & \text{TEM voltage law} \\ LI = \int_\Gamma \mu (\vec{H} \times \vec{n}) \cdot (\vec{H}_0 \times \vec{n}) da, & \text{TEM current law.} \end{cases} \quad (3.42)$$

Here  $\vec{E}$  and  $\vec{H}$  are arbitrary electric and magnetic fields,  $V$  and  $I$  are the voltage and current carried by the TEM mode, and the values  $C$  and  $L$  are referred to as *capacitance per meter* and *inductance per meter* respectively. These relations define the voltage and current by computing the  $L^2(\Gamma)$  inner product of  $\vec{E}$  and  $\vec{H}$  onto the TEM spaces  $\text{span}\{\vec{E}_0\}$  and  $\text{span}\{\vec{H}_0\}$  respectively. The description of variation in  $\Gamma$  along  $[0, \ell]$  for Maxwell's equations on  $\Omega_{\text{TL}}$  is very much analogous to a classical Fourier series representation; although, we consider only the fundamental mode in this manuscript.

In order to couple the EM and TL systems, we need to set self-consistent boundary conditions on  $\Gamma$  for the EM and at  $n = \ell$  for the TL. In addition, we wish to be able to allow the non-TEM part of the electric field to pass through  $\Gamma$  as an out going condition.

The values  $C$  and  $L$  are computed using

$$C = \int_{\Gamma} \epsilon |\vec{E}_0|^2 da, \quad CL = \epsilon \mu. \quad (3.43)$$

The first statement here is clear from the TEM voltage definition while the second assumes the propagation speed in the transmission line is the same for the TEM mode as any other wave. This assumption is reasonable for constant  $\epsilon, \mu$  systems which we will assume for the remainder of the work.

Under the TEM ansatz and defining  $C$  and  $L$  as above, Maxwell's equations can be reduced to the telegrapher's equations in the TL domain,  $\Omega_{\text{TL}}$ :

$$\begin{cases} C \frac{\partial V}{\partial t} + \frac{\partial I}{\partial n} = 0, & \text{voltage evolution equation,} \\ L \frac{\partial I}{\partial t} + \frac{\partial V}{\partial n} = 0, & \text{current evolution equation.} \end{cases} \quad (3.44)$$

c.f. [53] for details on the derivation.

With the volumetric equation established we will now discuss boundary conditions for the system. In the EM domain away from  $\Gamma$  we will assume homogeneous essential or natural conditions, i.e.  $\vec{E} \times \vec{n} = \vec{0}$  or  $\vec{H} \times \vec{n} = \vec{0}$  respectively. In the TL domain we will model the effect of the voltage generator at  $n = 0$  with a Robin condition

$$R_g I(0) = V_{\text{OC}} - V(0). \quad (3.45)$$

Here  $V_{\text{OC}}$  is a prescribed open circuit voltage and  $R_g$  is the resistance of the generator. If we impedance match the generator with the transmission line, i.e.  $R_g = \sqrt{\frac{L}{C}}$ , and  $V_{\text{OC}} = 0$  then this condition collapses to a first order absorbing condition.

To impose the coupling between we impose continuity of voltage TEM voltage between the EM and TL domain by applying the TEM voltage law as constraint on the Maxwell-Ampere law and relax with a Lagrange multiplier. We use the Lagrange multiplier approach to enable outflow of the non-TEM components of the electromagnetic fields. We impose continuity of current with a Neumann condition in the TL domain chosen to guarantee an energy stable coupling between the two domains. The details of this can be found in [79].

### 3.1.2.6. Implicit Algorithms

Implicit PIC algorithms were a new feature added to EMPIRE through this Grand Challenge LDRD. We have implemented two implicit algorithms: an electrostatic version of Lapenta's Energy Conserving Semi-Implicit Method (ECSIM) [63], and the electromagnetic Magnetic Implicit method developed by Voss Scientific [40, 136]. These semi-implicit algorithms can reduce the time to solution over traditional explicit PIC algorithms by allowing for coarse spatial discretizations that under-resolve the Debye length (i.e.  $\Delta x/\lambda_D > 1$ ) and large time-steps that under-resolve the plasma frequency (i.e.  $\omega_p \Delta t > 1$ ). Both of these conditions result in instabilities in explicit PIC algorithms, requiring restrictively fine meshes and small time-steps when dense plasma regions develop in a simulation. On the other hand, implicit methods remain stable when dense plasma regions develop while continuing to resolve only length- and time-scales of interest.

The key idea underlying both implicit algorithms we considered is to advance the electromagnetics and the particle velocities together. This represents the coupling between  $\vec{E}$  and  $\vec{u}$  implicitly, and ensures that no new stability constraint associated with this coupling is introduced. As the algorithms are very similar, we discuss only the Magnetic Implicit method in detail. Consider the Crank-Nicolson discretization of Ampère's law (3.19). Then, instead of advancing  $\vec{u}$  between half-steps as in (3.26), the algorithm advances  $\vec{u}$  from time  $n$  to  $n+1$  along with  $E$ , yielding

$$M_E E^{n+1} - \frac{c_0^2 \Delta t}{2} K_E B^{n+1} = M_E E^n + \frac{c_0^2 \Delta t}{2} K_E B^n - \frac{\Delta t}{\epsilon_0} J^{n+1/2}, \quad (3.46)$$

$$\frac{\vec{u}^{n+1} - \vec{u}^n}{\Delta t} = \frac{q}{m} \left[ \frac{1}{2} \left( \vec{E}^{n+1} + \vec{E}^n \right) + \frac{1}{c} \frac{\vec{u}^n + \vec{u}^{n+1}}{2\gamma^n} \times \vec{B}^{n+1/2} \right]. \quad (3.47)$$

To solve this system, the Magnetic Implicit algorithm breaks it into a predictor and a corrector step. For the predictor step, Faraday's law is advanced explicitly to obtain an approximation of  $\vec{B}^{n+1/2}$  and a guess of  $\vec{E}^{n+1} = 0$  is made. Then (3.47) can be advanced to obtain the predicted velocity  $\vec{u}^*$ , giving rise to a predicted current  $J^*$ . Using the definition of the current and (3.47), the current at the half-step can be written in terms of known data and  $E^{n+1}$  as

$$J^{n+1/2} = J^* + S E^{n+1}, \quad (3.48)$$

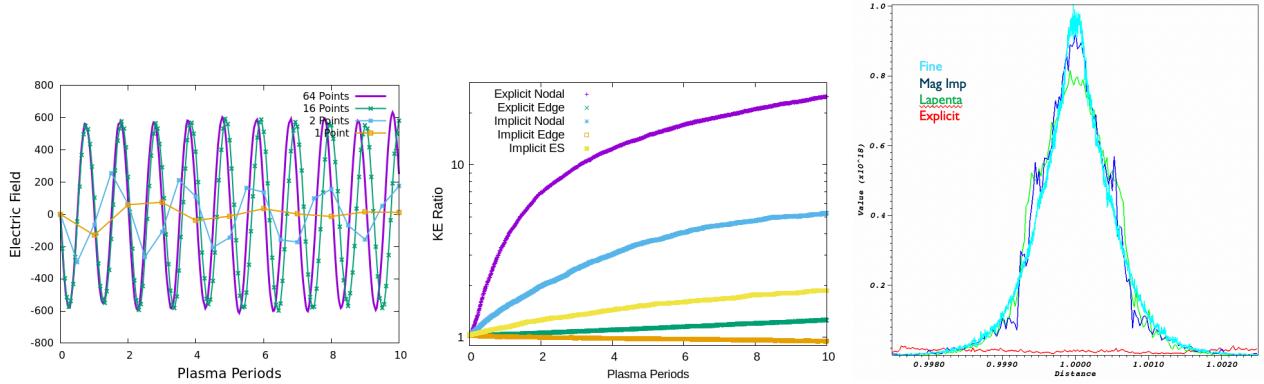
where  $S$  is the susceptibility matrix defined in the finite element setting as

$$S_{i,j} = \sum_{k=1}^{N_p} \int_{\Omega} \frac{1}{2\Delta t} \frac{q_k^2}{\gamma_k^{n+1/2} m_k} \hat{e}_j(\vec{x}_k^*) \cdot T [I - \vec{v}_k^* \vec{v}_k^*] \hat{e}_i(\vec{x}_k^*). \quad (3.49)$$

The operator  $T$  here is a rotation matrix about a scaling of the vector  $\vec{B}^{n+1/2}$ . Given this definition for  $J^{n+1/2}$ , Maxwell's equations are then solved with the augmented version of Ampère's law

$$(M_E + S) E^{n+1} - \frac{c_0^2 \Delta t}{2} K_E B^{n+1} = M_E E^n + \frac{c_0^2 \Delta t}{2} K_E B^n - \frac{\Delta t}{\epsilon_0} J^* \quad (3.50)$$

to obtain  $E^{n+1}$  and  $B^{n+1}$ . A final corrector step of (3.47) is then advanced using the updated electromagnetic fields to evolve the particle velocities and positions. Note that the original



**Figure 3-5. Left: Implicit PIC under-resolving plasma frequency. Middle: Comparison of kinetic energy in methods when under-resolving Debye length. Right: Comparison of methods on an expanding slab problem.**

Magnetic Implicit algorithm has been modified for EMPIRE to work in an unstructured finite element setting. The Magnetic Implicit framework also allows for a modified push that is more resilient to overstepping the cyclotron frequency and a Poisson correction for enforcing charge conservation. These features have not yet been implemented in the production release of EMPIRE as we have focused (in the near term) on the benefits of implicitly representing the coupling of velocity to electric field. The additional costs of this algorithm as compared to explicit PIC include the need for two particle pushes per time-step, the construction of the susceptibility matrix once per time-step, and the need to recompute the preconditioner since the linear system changes every time-step.

Figure 3-5 shows two numerical experiments with EMPIRE that have demonstrated the resilience of the implicit PIC methods to under-resolving the plasma frequency and the Debye length. On the left is a plot of the electric field in a simple plasma oscillation problem for various time-step sizes. This shows that the implicit algorithm captures the plasma frequency well when it is resolved (64 and 16 points per plasma period), with some phase error at 16 points per plasma period. When the plasma frequency is under-resolved (2 and 1 points per plasma period), the implicit method can no longer resolve the frequency but it damps the electric field and remains stable. The explicit PIC algorithm on the other hand goes unstable in this regime. On the right of Figure 3-5, we demonstrate that implicit PIC reduces the numerical heating caused by under-resolving the Debye length. The test problem simulates a thermal plasma in a periodic box with the  $\Delta x/\lambda_D > 1$ . The kinetic energy grows for explicit PIC using both edge- and node-based particle updates (more extremely for the node-based push) as a result of numerical heating. The implicit algorithms all show less heating than their corresponding explicit cases. Finally, the last figure compares the performance of the different algorithms on a simulation of an expanding slab of plasma. The simulation initializes a  $10^{18}$  number density electron-ion plasma in 20% of a 1 cm domain and allows the slab to expand. The methods are compared against a fine explicit simulation that resolves all length- and time-scales. Both Magnetic Implicit and ECSIM show qualitative agreement with the fine simulation while demonstrating some damping. The under-resolved explicit method in the other hand goes immediately unstable and ejects all particles from the domain.

### 3.1.3. EMPIRE-FLUID

#### 3.1.3.1. Relativistic Hydrodynamics for Power Flow

The plasma regimes present within the “inner MITL” in Z-Machine operation are such that electrons can be accelerated to relativistic velocities, i.e. where the Lorentz factor,  $\gamma = 1/\sqrt{1 - v^2/c^2} > 2$ . In turn, this places additional requirements on EMPIRE modeling capabilities that were not envisaged in the original EMPIRE design specification, namely that the fluid representation would have a low bulk velocity. In particular, this assumption meant that EMPIRE-FLUID algorithms were originally designed for Newtonian, rather than relativistic physics; for modeling of the “inner MITL” region in the Z accelerator, this means that electrons within this region would be accelerated to superluminal velocities, in turn corrupting solutions to Maxwell’s equations through (e.g.) the emission of numerical Cherenkov radiation.

To address this shortcoming, we have extended the algorithms within the EMPIRE-FLUID component of EMPIRE to include relativistic effects. Specifically, we have implemented and verified equations for relativistic, compressible hydrodynamics, which are governed by a set of hyperbolic conservation laws with source terms:

$$\partial_t \mathbf{U} + \nabla \cdot \mathbf{F}(\mathbf{U}) = \mathbf{S}(\mathbf{U}) \quad (3.51)$$

where the conserved variables  $\mathbf{U} = [D, \mathbf{M}, E]^T$  are the relativistic density, relativistic specific momentum, and the total energy density, including energy from the rest mass. The flux is

$$\mathbf{F} = \begin{bmatrix} \rho \mathbf{u} \\ \rho h \mathbf{u} \otimes \mathbf{u} + P \mathbf{I} \\ \gamma \rho h \mathbf{u} \end{bmatrix}, \quad (3.52)$$

where we have used the primitive state  $\mathbf{W} = [\rho, \mathbf{u}, P]^T$  containing the rest mass density, the 4-velocity, and pressure and the enthalpy

$$h = \frac{e + P}{\rho} \quad (3.53)$$

where  $e$  is the specific internal energy. The primitive state  $\mathbf{W}$  relates to the conserved state  $\mathbf{U}$  by

$$\mathbf{U} = \begin{bmatrix} \gamma \rho \\ \gamma(e + P)\mathbf{u}/c^2 \\ \gamma^2(e + P) - P \end{bmatrix} = \begin{bmatrix} \gamma \rho \\ \gamma \rho h \mathbf{u}/c^2 \\ \gamma^2 \rho h - P \end{bmatrix} \quad (3.54)$$

where  $\gamma = \sqrt{1 + |\mathbf{u}/c|^2}$  is the Lorentz factor. Note that, unlike in the Newtonian case, the primitive quantities cannot be written in terms of the conserved quantities; instead, a non-linear solver must be employed to obtain the primitive quantities at each step in the update (discussed further below). These equations are closed by an Equation of State; for this work, an ideal equation of state is applied to the relativistic fluid by constraining the enthalpy to

$$h \stackrel{\text{EOS}}{=} c^2 + \frac{\Gamma}{\Gamma - 1} \frac{P}{\rho} \quad (3.55)$$

where  $\Gamma$  is the adiabatic index. This ideal equation of state with a fixed adiabatic index does not capture well an ideal relativistic gas, as the adiabatic index will change as the gas temperature becomes relativistic. For a weakly relativistic gas,  $\Gamma \rightarrow 5/3$  while for a relativistic gas  $\Gamma \rightarrow 4/3$ . Only adiabatic indices below 2 are permissible for physicality.

In Eqn. 3.51, the source term,  $\mathbf{S}(\mathbf{U})$  describes coupling to the electromagnetic fields through the Lorentz force:

$$\mathbf{S} = \begin{bmatrix} 0 \\ \mu\rho(\gamma\mathbf{E} + \mathbf{u} \times \mathbf{B}) \\ \mu\rho\mathbf{u} \cdot \mathbf{E} \end{bmatrix}, \quad (3.56)$$

where we have introduced the charge-to-mass ratio,  $\mu = q/m$ . Compared to the Newtonian case, the relativistic form of the Lorentz force introduces additional complexity due to the presence of the Lorentz factor,  $\gamma$  and the four-velocity,  $\mathbf{u}$ . The crux of this complexity is that the Lorentz force cannot be written simply as a linear function of the *conserved* density and momentum, but instead must be computed through the non-linear relationship between conserved and primitive variables. In addition, the current density that appears in the right-hand side of Maxwell's equation,  $\mathbf{J} = \mu\rho\mathbf{u}$  is also no longer linearly related to the fluid momentum,  $M = \gamma(e + P)\mathbf{u}/c^2$  as in the relativistic case, but must be found using a non-linear solve from the conserved fluid quantities.

## Spatial and Temporal Discretization

EMPIRE-FLUID discretizes the hyperbolic conservation law, Eqn. 3.51 using a discontinuous Galerkin (DG) method in a similar fashion as was proposed by [87] following on the influential sequence [24, 23, 22, 25]). The DG method requires a meshed defined as the subdivision of the domain into non-overlapping hexahedral (3D) or quadrilateral (2D) cells denoted  $\Omega_k \subset \Omega \subset \mathbb{R}^d$ . The approximation of the conserved variables on cell  $\Omega_k$  is written

$$\mathbf{U}(\mathbf{x}) \approx \mathbf{U}^h(\mathbf{x}) = \sum_{k=1} \mathbf{U}_i \phi(\mathbf{x}) \quad \mathbf{x} \in \Omega_k \quad (3.57)$$

where the set  $\{\phi_k(\mathbf{x})\}$  is linearly independent basis that spans a polynomial space of fixed order on element  $\Omega_k$ . Globally,  $\mathbf{U}^h$  is defined as a piecewise polynomial function with discontinuities permitted at cell boundaries. On each cell the approximation solution to Eqn. 3.51 is computed by enforcing the residual is orthogonal to the test space. Practically, after integration by parts, this implies the satisfaction of weak form

$$\int_{\Omega_k} \frac{\partial \mathbf{U}^h}{\partial t} \phi(\mathbf{x}) d\mathbf{x} + \oint_{\partial\Omega_k} \overline{\mathbf{F}(\mathbf{U}^h) \cdot \mathbf{n}} \phi(\mathbf{x}) ds - \int_{\Omega_h} \mathbf{F}(\mathbf{U}^h) \cdot \nabla \phi(\mathbf{x}) d\mathbf{x} = 0, \quad \forall \phi \in \{\phi_k\}. \quad (3.58)$$

The second of these terms corresponds to the flux normal to the surface of an element integrated over the surface of the element. However, the solution at cell interfaces is double-valued as indicated by the overline; one value corresponding to the data inside cell, the other from the neighboring cell. As such, the solution is discontinuous and the flux must be computed using a Riemann solver in a fashion similar to that employed in the finite volume method (for more detail, see e.g. [87]). In EMPIRE-FLUID we have implemented two *approximate* Riemann solvers: HLL and HLLC (see e.g. [82]). Beyond the choice of Riemann solver, the discrete conservation law,

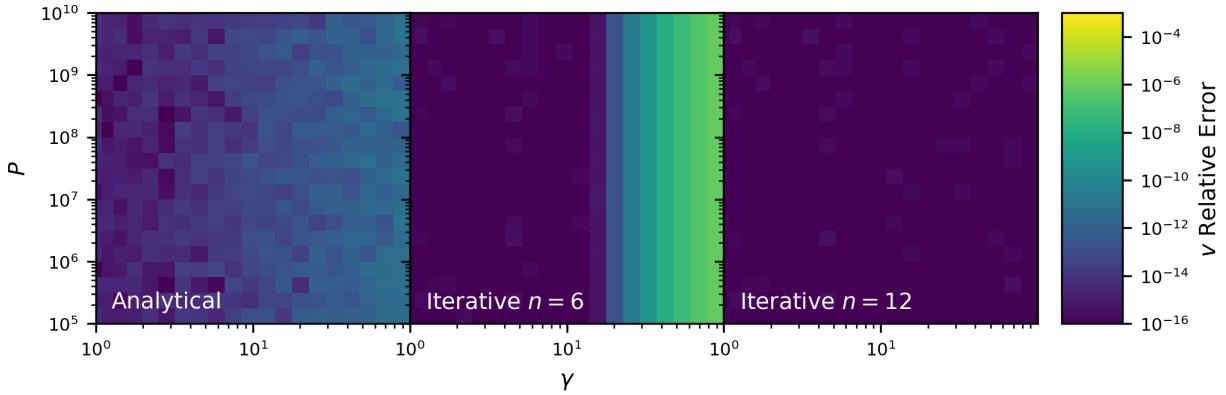
Eqn. 3.58 can admit a range of different basis orders; typically, a first order basis (e.g piecewise constant) will eliminate the contribution of  $\int_{\Omega_h} \mathbf{F}(\mathbf{U}^h) \cdot \nabla \phi(\mathbf{x}) d\mathbf{x}$ , resulting in a scheme equivalent to a first order finite volume discretization. Moving to higher order bases (e.g. piecewise linear, etc.) will introduce the need to provide additional stabilization (e.g. dissipation) at discontinuities and shocks. In EMPIRE-FLUID this is accomplished through the use of a solution limiter, described by [84], which we have extended to handle relativistic flows, as described later in this section.

The discrete form of the conservation law, Eqn 3.58 is advanced using an Implicit-Explicit (IMEX) scheme. For electromagnetic-plasma coupling, these schemes are designed to treat fast timescales in the system (the light wave, plasma frequency and cyclotron frequency) *implicitly*, while the slow advective timescales are treated explicitly. However, for a relativistic fluid where the slow advective timescales can approach the light wave, the utility of separating these timescales is reduced. In particular, we have found that the non-linear coupling between the fluid and electromagnetic solves can eliminate performance gains found for implicit treatment of the plasma and cyclotron frequencies in conditions found in the inner MTL; in future work, we plan to investigate this further by treating the coupling between the two systems in a fundamentally non-linear fashion. For the purpose of the current work, however, we treat the advective, plasma and cyclotron timescales *explicitly* while the fast timescale associated with the light-wave in Maxwell's equations is treated *implicitly*. This latter aspect was found to be necessary to provide sufficient stabilization to the Maxwell discretization discussed in §3.1.2.2. We consider three different IMEX schemes: a first order (1,1,1) scheme; a second order (2,2,2) scheme and a third order (3,2,2) scheme.

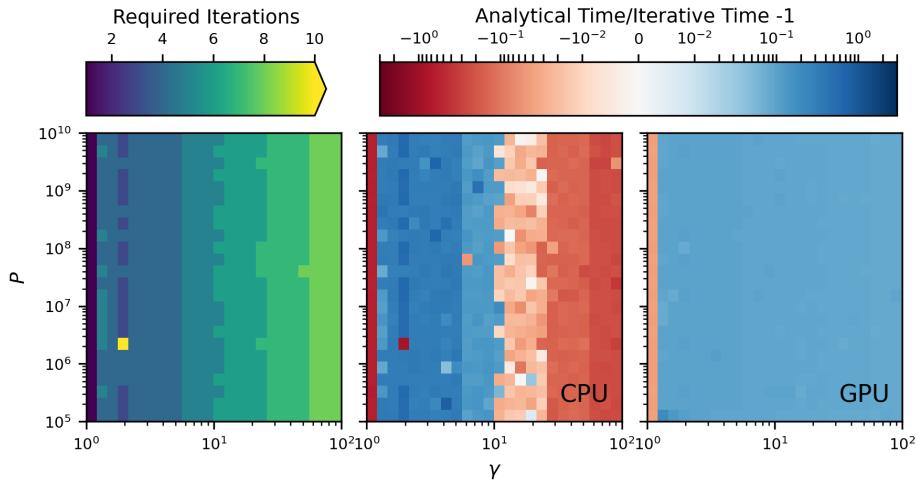
### Conserved to Primitive Conversion for Ideal Gas

Unlike Newtonian hydrodynamics, the relationship between the conserved and primitive variables in relativistic hydrodynamics is non-linear and represents both a significant additional complexity and source of numerical error due to the high order of the polynomials being solved. For highly relativistic flows close to the speed of light, we are often limited by machine precision when representing small changes in the velocity that equate to large changes in the Lorentz factor. To address this requirement, we have implemented two different schemes to recover primitive variables from the conserved state; an analytic solver based on [108] and an iterative solver following the recovery method presented in [95].

Figure 3-6 shows the relative error in the recovered velocity using the analytical method and iterative method using 6 and 12 iterations. The plots are created by applying the methods on a grid of  $25^2$  primitive states with  $D = 1$  and 25 logarithmically spaced pressures from  $10^5$  to  $10^{10}$  and 25 logarithmically spaced Lorentz factors from 1 to 100. Each pair of pressure and Lorentz factor is converted to a conserved state using Eqn. 3.54 that is converted back to a primitive state using one of the two recovery methods. We then compute the relative error of the velocity in the recovered primitive state to the original velocity determined by the Lorentz factor. Overall, we find that the iterative method to recover the primitive variables from the conserved variables is more flexible, robust, accurate, and in some cases faster than the analytical method.



**Figure 3-6.** Map of the error of the analytical method (LEFT) and the error of the iterative method after 6 iterations (MIDDLE) and 12 iterations (RIGHT), where the Lorentz factor varies from 1 to 400 on the  $X$  axis and pressure varies from  $10^5$  to  $10^{10}$  on the  $y$ -axis, and color shows the error in the recovered velocity. White shows where the analytical solver failed to find a solution, due to imaginary numbers encountered during the solve. In all cases shown here  $\rho = 1$  and  $c = 3 \times 10^8$ . The iterative solve is more robust and accurate at high Lorentz factor using 5 iterations than the analytic method.



**Figure 3-7.** The left panel shows the required number of iterations for the iterative primitive recovery method to match the accuracy of the analytical method with logarithmically spaced Lorentz factors from 1 to 100 on the  $x$  axis and logarithmically spaced pressures from  $10^5$  to  $10^{10}$  on the  $y$  axis with colors showing how many iterations are required to match the analytical method on CPUs using the Intel compiler. The middle and right panels show timing comparisons of the iterative and analytical methods on CPUs using Intel and GPUs using CUDA respectively. We time both methods on a  $10^4$  cell mesh with 64 points per cell timing over 200 iterations for each of the primitive states from the left panel. For number of iterations, we use the required number for each architecture, meaning the number of iterations used for the GPU case differs slightly from the CPU case. We use a symmetrical logarithmic color mapping to show the metric Analytical Time/Iterative Time – 1 so that areas in red show the fraction by which the analytical method is faster and blue areas show the fraction by which the iterative method is faster.

We also investigate the performance of the two recovery methods and show our findings in Figure 3-7. Using the same grid of primitive states that we used in Figure 3-6, the left panel of the figure shows the number of iterations of the iterative method needed to recover the velocity to the same or better precision than the analytical method. A purple line on the left hand side shows where both methods can assume a velocity of 0 and avoid most of the computation. Yellow spots at low Lorentz factor show primitive states that the analytic method happened to recover exactly but which the iterative method did not. At higher Lorentz factors the iterative solver requires a few more iterations, but none of the primitives tested in this sweep required more than 10, except for those cases where the analytical solver recovered the velocity exactly.

In the middle and right panels of Figure 3-7, we show comparative timings of the two methods running each of the primitive states from Figure 3-6 on  $10^4$  cells with 64 points per cell, averaging over 200 runs each, with the CPU timing comparison in the middle panel and the GPU timing comparison in the right panel. We show the metric Analytical Time/Iterative Time – 1, where the iterative time is with the number of iterations required to match the analytical accuracy, in order to highlight where the iterative solver is faster – red regions show the fraction by which the analytical method is faster than the iterative method while blue regions show the fraction by which the iterative solver is faster.

## Physicality Forcing Limiter

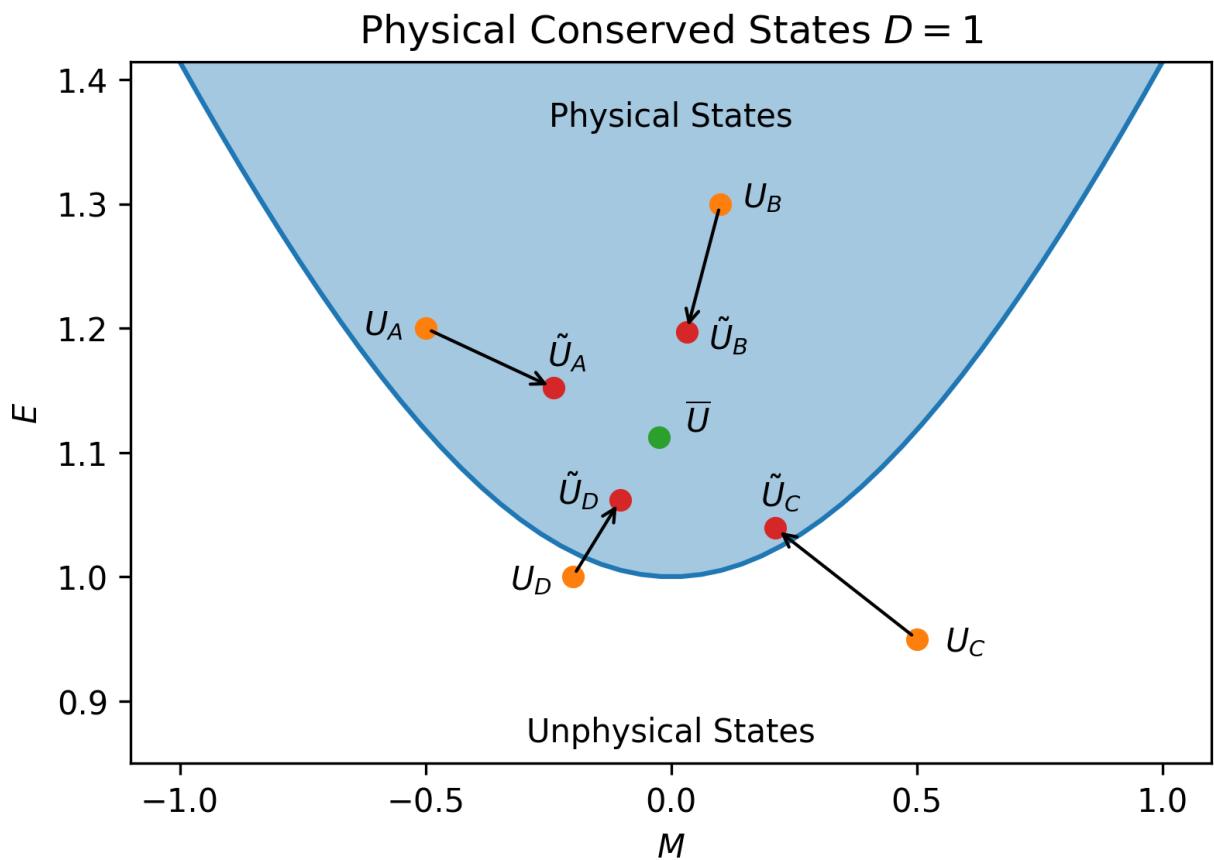
While using zeroth order polynomials (first order accurate scheme) for a relativistic hydrodynamics DG method is guaranteed to produce a physical conserved state after every flux update even with shocks, higher order bases can introduce spurious oscillations and non-physical conserved states within cells around shocks (see [141]). To resolve this issue, a limiter is needed to smooth the solution within a cell. Taking inspiration from the limiter presented in [84], we present here a limiting method that enforces physical conserved states within a cell with a physical volume average.

Following [95] and [141], a conserved state that satisfies

$$D > 0, \quad q(\mathbf{U}) \equiv E/c^2 - \sqrt{D^2 - |\mathbf{M}/c|^2} > 0, \quad (3.59)$$

is a physically admissible state, as long as  $e(\rho, p)$  is continuously differentiable under the chosen equation of state. If a conserved state satisfies Eqn. 3.59, the state can be inverted for a primitive state with positive density and pressure with a velocity less than  $c$ . Note that since Eqn. 3.59 is concave, we can use the same strategies from [84] in a simple limiting procedure to enforce physicality within a DG cell. From a high level, we apply a limiter to average nodal points within a cell towards a physical volume average.

Before limiting cells, we first detect for cells with non-physical nodal points by checking that all conserved states at the nodal points –  $\mathbf{U}_i$  – satisfy Eqn. 3.59. If any point fails, we flag the cell for needed limiting to ensure that all points are physical. We then check that the cell volume average  $\bar{\mathbf{U}}$  of the conserved state satisfies Eqn. 3.59. As long as the cell volume average is physical, a limiter can be found that ensure physicality without changing the global conserved quantities. If



**Figure 3-8.** Graphical demonstration of the physicality preserving limiter, moving the unlimited, some unphysical, orange points to the limited, all physical, red points without changing the cell average.

the cell volume is not physical, then the nodal points cannot be made physical through the limiting procedure without changing the volume average. This procedure is illustrated schematically in Figure 3-8.

### 3.1.3.2. Relativistic Fluid Boundary Conditions for Power Flow

#### Mass Injection Boundary Conditions

We have developed mass injection boundary conditions for both relativistic and Newtonian fluid systems in EMPIRE. These boundary conditions allow the user to specify a time-dependent mass influx through a boundary side set. The mass injected into the system is consistent with the specified mass flux to a high degree of accuracy. The user also specifies an injection velocity and pressure, although it should be noted that the velocity and temperature of the injected substance is not tightly constrained like the mass flux is. Some of the boundary conditions below for thermal desorption and SCL emission are based on this boundary condition. In fact, from a programming point of view, the classes of these other boundaries inherit from the mass injection boundary.

This mass injection boundary condition is a Neumann boundary condition in the sense that it assigns the value of the flux  $\mathbf{F}(\mathbf{U}^h) \cdot \mathbf{n}$  over the corresponding section of the domain boundary. When computing the boundary flux, some considerations must be made so that the boundary condition is well-posed. A naive approach would be to compute the conserved quantities  $\mathbf{u}_{inj}$  that correspond to the specified mass influx  $j(t)$ , injection temperature  $T_{inj}$  and injection velocity  $v_{inj}$ . Then, one could directly compute a boundary flux  $\mathbf{F}(\mathbf{U}_{inj}^h) \cdot \mathbf{n}$  as the flux is a function of  $\mathbf{u}$ . Unfortunately, this boundary condition is ill-posed when the flow is subsonic and the flow will inevitably be subsonic at some point except in special cases. Especially when modeling an electron fluid, which will tend to have a very fast sound speed due to the low mass of an electron. Instead, the boundary flux is specified according to a weak Riemann approach as discussed by Mengaldo et al. [81]. In the weak Riemann approach, a Riemann problem for each quadrature point on the boundary surface is constructed. The state of the *interior* side of the Riemann problem,  $\mathbf{u}_i$ , is determined by the state of the solution at the location of the quadrature point. The state of the *virtual* side,

$$\mathbf{u}_v = \mathbf{u}(\rho_{inj}, -v_{inj}\mathbf{n}, P(T_{inj})), \quad (3.60)$$

is constructed so that the mass influx of charged substance into the system at this point is  $j(t)$ . These two states are then used as input to an HLL approximate Riemann solver,[82, 111] from which we obtain the appropriate boundary flux  $\mathbf{F}(\mathbf{U}^h) \cdot \mathbf{n} = \mathbf{F}_{HLL}(\mathbf{u}_i, \mathbf{u}_v) \cdot \mathbf{n}$ .

The virtual state  $\mathbf{u}_v$  is uniquely determined by the virtual state temperature  $T_{inj}$ , velocity  $v_{inj}$  and density  $\rho_{inj}$ . The user specifies  $T_{inj}$  and  $v_{inj}$  as input deck parameters. The virtual state density  $\rho_{inj}$  must be computed so that the injection influx is the required value. To this end, we define the residual,

$$r(\rho; v_{inj}, T_{inj}) = j - \mathbf{F}_{HLL}(\mathbf{u}_i, \mathbf{u}(\rho, -v_{inj}\mathbf{n}, P(T_{inj}))) \cdot \mathbf{n}. \quad (3.61)$$

We apply the Brent method to find the value  $\rho_{inj}$  for which  $r = 0$ . First the solution is bracketed; a pair of values  $(\rho_{min}, \rho_{max})$  is found such that  $r(\rho_{min}; v_{inj}, T_{inj}) > 0$  and  $r(\rho_{max}; v_{inj}, T_{inj}) < 0$ . Starting at an initial guess of  $j/v_{inj}$ , the value of  $\rho_{max}$  ( $\rho_{min}$ ) is iteratively increased (decreased)

until appropriate values are obtained. Brent's root-finding method [91] is then applied to determine the value of  $\rho_{inj}$ . Because a bracketed root-finding method is used, constraining the search to positive density values is trivial.

### Thermal Desorption Boundary Condition

The thermal desorption boundary condition is built on top of the mass injection boundary condition, as precise control of the amount of desorbed substance introduced into the simulation is required. Note that we have only implemented a thermal desorption boundary condition for relativistic fluids, although it would be straightforward to make a Newtonian fluid thermal desorption boundary condition using the same algorithm.

The thermal desorption boundary condition injects a temperature-dependent influx of mass on a specified side set. The temperature is computed by EMPIRE's surface temperature diagnostic, which can compute the surface temperature due to Ohmic heating in a manner that is consistent with eddy currents due to the computed magnetic fields. See Section 3.1.2.4 for a description of this temperature calculation. Note that the Ohmic heating calculation involves an approximation that makes it only suitable for good electrical conductors: the approximation is to use a perfect electrical conductor boundary condition while in general a resistive wall boundary condition is required. The surface temperature is computed on each mesh face by computed coupled 1D thermal-electrical diffusion into the surface. Temperature dependence of the electrical and thermal properties, including phase changes, can be specified. Expansion/contraction due to heating of the material is not accounted for in the present model, although this would not be extremely difficult to account for in an approximate (linear) manner if required in future applications.

The thermal desorption flux is specified by providing values for the coefficients of a generalized Arrhenius equation, which is the same as that described in 3.1.2.4,

$$j(T) = vT^\alpha \theta^n \exp\left(\frac{-E_a}{k_B T}\right), \quad (3.62)$$

where  $T$  is the surface temperature and  $k_B$  is the Boltzmann constant. The parameters  $\alpha$ ,  $n$ ,  $v$  and  $E_a$  are specified by the user. The quantity  $\theta$  is the surface density of the adsorbed substance. The surface density diminishes as desorption proceeds according to,

$$\dot{\theta} = -j(T). \quad (3.63)$$

The total amount of desorption possible is therefore controlled by the initial value of  $\theta$ , which is specified by the user.

### SCL Boundary Conditions

The space charge limited emission (SCL) of electrons or ions from electrode surfaces is an essential effect in the operation of certain high voltage devices such as magnetically insulated transmission lines. Numerically modeling the operation of these devices is traditionally

accomplished with a PIC method. Monte Carlo methods such as PIC, however, have a number of drawbacks. Especially, for most of these methods the random noise in the computed quantities goes as  $1/\sqrt{N}$  where  $N$  is the number of samples used. Time averaging over an appropriately lengthy interval is required to obtain reasonable approximate computed quantities in many cases.

The SCL boundary condition injects a mass influx of charged substance  $j$  that varies with the component of the electric field normal to the emitting surface  $\mathbf{E} \cdot \mathbf{n}$ ,

$$j(\mathbf{x}, t) = \begin{cases} \frac{q}{|q|} \left( \gamma_0 \mathbf{E} \cdot \mathbf{n} + \gamma_1 \partial_t \mathbf{E} \cdot \mathbf{n} \right) \tanh \left( \frac{t}{\tau_{ramp}} \right), & \text{if } x \geq 1 \\ 0, & \text{if } q \mathbf{E} \cdot \mathbf{n} \geq 0 \end{cases} \quad (3.64)$$

The species charge of the emitted substance is  $q$  (e.g. the charge of an electron or proton). Given an appropriately chosen relaxation parameter  $\gamma_0$  and damping parameter  $\gamma_1$ ,  $\mathbf{E} \cdot \mathbf{n}$  will relax to approximately zero at the emitting surface, and  $j(\mathbf{x}, t)$  will continuously adjust so that the value of  $\mathbf{E} \cdot \mathbf{n}$  remains small as the solution evolves. The SCL boundary condition also requires that the user provide an injection temperature  $T_{inj}$  and injection velocity  $v_{inj}$ . Specifying an appropriate ramp time  $\tau_{ramp}$  enables smoothing of the initial transient for certain initial conditions.

Fluid SCL boundary conditions based on 3.64 have been implemented for both Newtonian fluids and for relativistic fluids. For relativistic fluids, there are two versions, one of which is based on the mass injection boundary condition from Section 3.1.3.2. The other uses a simpler "mass injection" algorithm, where the virtual state density that appears in Eqn. 3.60, is set equal to  $j(\mathbf{x}, t)/v_{inj}$ . The latter boundary condition is deprecated, however it is interesting that this simpler approach to the SCL boundary condition provides reasonable results under many circumstances.

### 3.1.3.3. Verification of the Relativistic Hydrodynamics Solver, Electromagnetic Coupling and Boundary Conditions

To verify the accuracy of the relativistic hydrodynamics scheme, we investigate several standard test problems in 1D and 2D with and without shocks. First, we demonstrate the accuracy of the method evolving five different 1D Riemann problems to high resolution reference solutions generated from a publicly available finite volume code [128]. Next, we demonstrate the schemes ability to handle multi-dimensional shocks through a series of 2D Riemann problems previously established in the literature. Following on from this, we measure the growth rate of the relativistic Kelvin-Helmholtz instability in 2D, which serves to establish the ability of the code to capture shear flows that can play an important role in MITL physics. Finally, we demonstrate coupling to the EMPIRE electromagnetic and transmission line solvers through the use of an ordinary wave test.

#### Riemann Problems in One and Two Spatial Dimensions

First, we investigate the accuracy of the relativistic hydrodynamics method in EMPIRE evolving 1D Riemann problems in order to determine the fidelity to which the scheme is able to capture

shocks. For initial conditions, we use four shocks from [87] which were proposed in [74, 75], and one sod shock tube for a total of five different 1D Riemann problems. For each of the five 1D Riemann problems, we use a  $[0, 1]$  grid with Dirichlet boundary conditions except for test 4, which uses a reflecting boundary condition on the right wall. The tests begin divided into a primitive state on the left  $\mathbf{W}_L = (\rho, v_x, v_y, p)_L$  for  $x \in [0, 0.5]$  and right  $\mathbf{W}_R = (\rho, v_x, v_y, p)_R$  for  $x \in [0.5, 1]$  except for test 4, which begins with a constant primitive state throughout the volume. In all cases,  $v_z = 0$ .

For reference data, we use a  $n = 2^{14}$  solution using a publicly available finite-volume relativistic hydrodynamics algorithm [128]. We run the each 1D Riemann problem with five resolutions in powers of two from  $n = 256$  to  $n = 4096$  with polynomial basis orders 0, 1, and 2 using the HLLC Riemann solver and iterative primitives recovery method. Figure 3-9 shows a log-log plot of the L1 error of the density, longitudinal velocity, and pressure compared to the reference solution along with power fits to the convergence rate and the expected rate of convergence.

Next, we test the robustness of the method evolving interacting shocks in multi-dimensions using the three 2D Riemann problems described in [87]. In each of the three problems, the problem is defined with a  $[-1, 1] \times [-1, 1]$  domain divided into four quadrants with different initial states. Following [87], we denote these states using

$$\mathcal{Q}_1 := [0, 1] \times [0, 1] \quad (3.65)$$

$$\mathcal{Q}_2 := [-1, 0] \times [0, 1] \quad (3.66)$$

$$\mathcal{Q}_3 := [-1, 0] \times [-1, 0] \quad (3.67)$$

$$\mathcal{Q}_4 := [0, 1] \times [-1, 0] \quad (3.68)$$

and denote the initial primitive states in each of these quadrants by  $\mathbf{W}_1, \mathbf{W}_2, \mathbf{W}_3$ , and  $\mathbf{W}_4$  respectively. For all of these Riemann problems, we use an adiabatic index of  $\Gamma = 5/3$ , use  $v_z = 0$  everywhere, and use transmissive boundary conditions on all sides. We evolve each Riemann problem to  $t = 0.8/c$ .

The three tests are defined as follows:

1. This test consists of two interacting shocks and contact discontinuities, forming curved shock fronts. The initial primitive states are

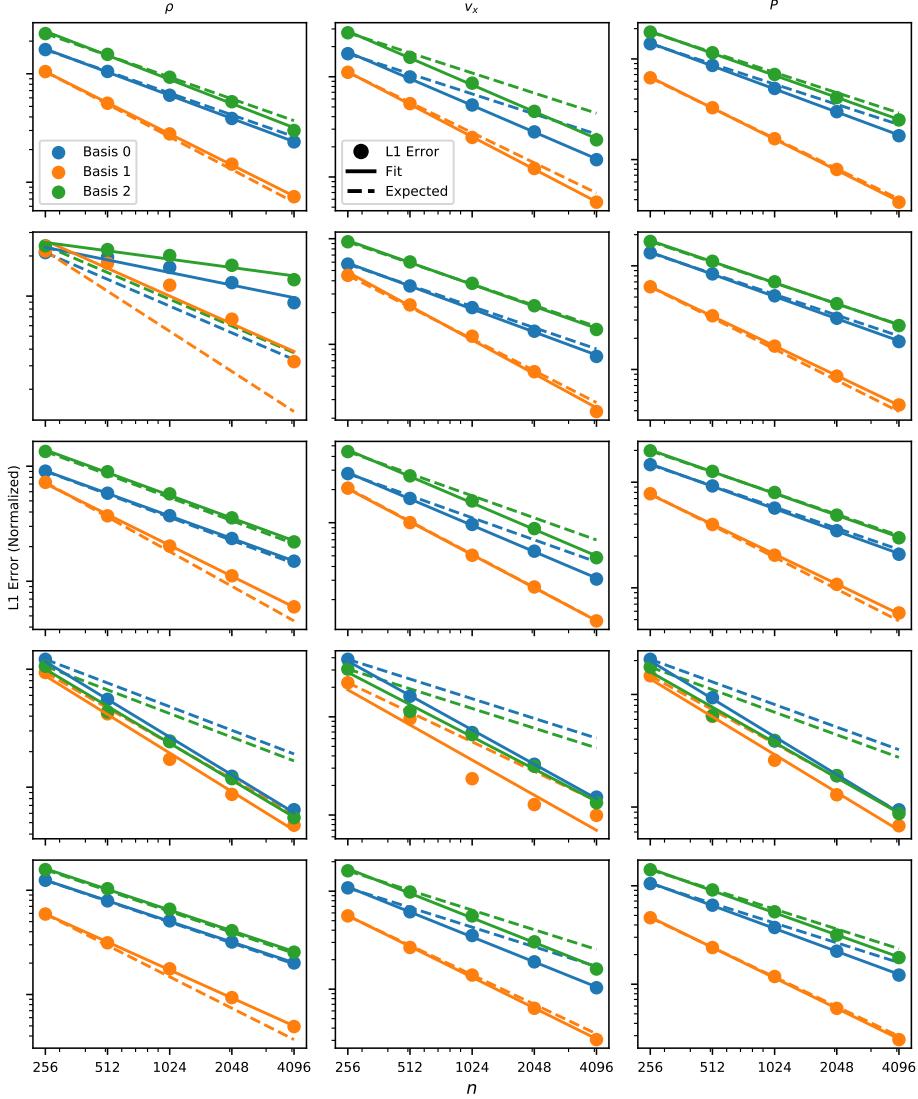
$$\mathbf{W}_1 := (0.035145216124503, 0.0, 0.0, 0.162931056509027c^2) \quad (3.69)$$

$$\mathbf{W}_2 := (0.1, 0.7c, 0.0, 1.0c^2) \quad (3.70)$$

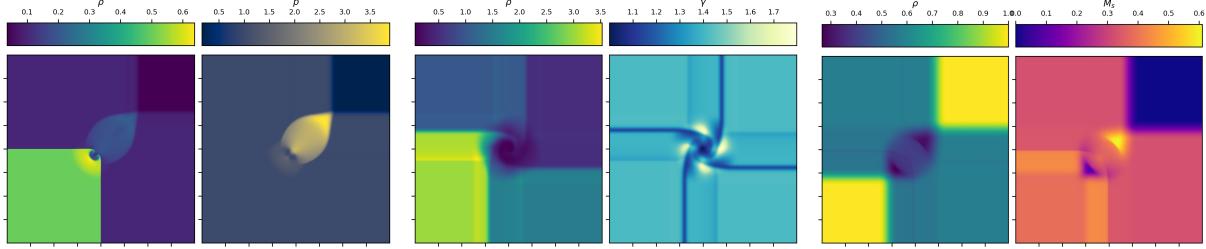
$$\mathbf{W}_3 := (0.5, 0.0, 0.0, 1.0c^2) \quad (3.71)$$

$$\mathbf{W}_4 := (0.1, 0.0, 0.7c, 1.0c^2) \quad (3.72)$$

2. This test consists of four vortex sheets moving outwards from the origin. The initial



**Figure 3-9. Convergence of the L1 error of `EMPIRE` compared to high resolution reference solutions of the four 1D Riemann problems from[87].** From top to bottom, the first row shows a mildly relativistic blast wave, the second a highly relativistic blast wave, the third a blast wave with transverse velocity, the fourth a planar shock reflection, and the fifth a sod shock tube. The columns show from left to right the rest-mass density, the pressure, the velocity, and the Lorentz factor. In each panel we show the L1 error of `EMPIRE` simulations with dots, a fitted convergence rate using logarithmically weighted least squares with a solid line, and the first order convergence rate with a dashed line. We use different colors to denote different basis orders, using blue for basis order 0, orange for basis order 1, and green for basis order 2. Due to the presence of shocks, we expect the L1 error to converge to first order.



**Figure 3-10. From left-to-right, results from EMPIRE-FLUID calculations of 2D Riemann Problems 1, 2, 3. Plots for each test show the rest-mass density in the left panel and the pressure in the right panel at  $t = 0.8/c$  using basis order 2 on a grid with 1024 elements.**

primitive states are:

$$\mathbf{W}_1 := (0.5, 0.5c, -0.5c, 5c^2) \quad (3.73)$$

$$\mathbf{W}_2 := (1.0, 0.5c, 0.5c, 5c^2) \quad (3.74)$$

$$\mathbf{W}_3 := (3.0, -0.5c, 0.5c, 5.0c^2) \quad (3.75)$$

$$\mathbf{W}_4 := (1.5, -0.5c, -0.5c, 5.0c^2) \quad (3.76)$$

3. This test consists of a central rarefaction and has an initial state defined by:

$$\mathbf{W}_1 := (1.0, 0.0, 0.0, 1.0c^2) \quad (3.77)$$

$$\mathbf{W}_2 := (0.5771, -0.3529c, 0.0, 0.4c^2) \quad (3.78)$$

$$\mathbf{W}_3 := (1.0, -0.3529c, -0.3529c, 1.0c^2) \quad (3.79)$$

$$\mathbf{W}_4 := (0.5771, 0.0, -0.3529c, 0.4c^2) \quad (3.80)$$

Results from each of these three test are shown in Figure 3-10. Overall these tests demonstrate the ability of EMPIRE-FLUID to handle relativistic shocks, rarefactions and vorticity in multiple dimensions.

### The Relativistic Kelvin-Helmholtz Instability

To test further test accuracy of the method for evolving 2D discontinuities, we simulated and measured the growth rate of the Kelvin-Helmholtz instability (KHI).

The growth rate of the KHI in 3D special relativistic hydrodynamical flows was explored analytically in [11]. Given the initial perturbation in [11], the change in any variable  $q$  should follow a linear growth rate  $\delta q(t) \propto \exp -i\omega t$ , where the imaginary part of the complex frequency  $\omega$  determines the linear growth rate of the instability.  $\omega$  can be found from the complex phase velocity  $\phi = \omega/(c_s k)$  where  $k$  is the longitudinal perturbation mode. The phase velocity has three

solutions in the linear growth phase: the trivial  $\phi$  mode and two modes found from the solution of

$$\frac{\phi^2}{M^2} = \frac{\mathcal{M}_e + 1 - \beta_e^2 \pm \sqrt{4(\mathcal{M}_e^2(1 - \beta_e^2))^2}}{\mathcal{M}_e^2 + 2\beta_e^2} \quad (3.81)$$

, where only the negative root of  $\phi^2$  is valid and potentially unstable, where  $\mathcal{M}_e$  and  $\beta_e$  are the effective relativistic Mach number and effective velocity relative to the speed of light, defined as

$$\mathcal{M}_e = \mathcal{M} \cos \theta, \quad \beta_e = \beta \cos \theta \quad (3.82)$$

. Here,  $\mathcal{M}$  is the relativistic Mach number  $\mathcal{M} = \gamma M / \gamma_s$  with  $\gamma_s = 1 / \sqrt{1 - c_s^2/c^2}$  and  $\theta$  is the angle between the fluid velocity and wave number projection in the  $x - z$  plane so that

$$\cos \theta = \frac{1}{\sqrt{1 + m^2/k^2}} \quad (3.83)$$

. This perturbation mode is stable as long as

$$\mathcal{M}_e > \sqrt{2}, \quad \mathcal{M}_e > \sqrt{2/(\gamma^2 + \beta_e^2)}. \quad (3.84)$$

We begin by assuming that initial perturbations to the primitive states above and below the interface are proportional to

$$\delta q_{\pm} \propto \exp[i(kx + l_{\pm}y + mz - \omega t)] \equiv \Theta_{\pm}(x, y, z, t) \quad (3.85)$$

where + states refer to perturbation in the region  $y > 0$  and - states refer to perturbations in the region  $y < 0$ . The four-vector  $K_{\pm} = (k, l_{\pm}, m, \omega)$  comprises the spatial and temporal wave vector, where  $k$  and  $m$  can be set arbitrarily but the two wave numbers  $l_{\pm}$  are fixed in order to enforce a continuous pressure at the interface ie.  $\delta p_{+}|_y = 0 = \delta p_{-}|_y = 0 = \delta p$  and the fluid displacement  $\delta \zeta_{\pm}(t, x)$  match at the interface, which in turn means that

$$\delta v_{y,+} / (\omega - kV) \quad (3.86)$$

where the tangential velocity can be found from

$$\delta v_{y\pm} = \frac{cl_{\pm}}{(\omega \mp kU) \rho h \gamma^2} \delta p \quad (3.87)$$

The transverse wave number can be found from

$$l_{\pm} - \gamma k \sqrt{(\phi \mp M)^2 - \left(1 - \mp \frac{\phi \beta^2}{M}\right)^2 - \frac{m^2}{\gamma^2 k^2}} \quad (3.88)$$

Note that  $l_{\pm}$  is complex which means that the sign of  $\Im l_{\pm}$  is arbitrary. In order to enforce that perturbations go to zero as  $y \rightarrow \pm\infty$ ,  $\Im(l_+)$  must be positive and  $\Im(l_-)$  must be negative. This requirement is known as the Sommerfeld radiation condition and must be applied in the frame in which the fluid is at rest [41].

We chose perturbation amplitudes for the density, longitudinal velocity, and pressure relative to their initial values, ie.

$$\delta\rho_{\pm} = A\rho_0\Theta_{\pm}(x,y,z,t) \quad (3.89)$$

$$\delta v_{x,\pm} = AU\Theta_{\pm}(x,y,z,t) \quad (3.90)$$

$$\delta P_{\pm} = AP_0\Theta_{\pm}(x,y,z,t). \quad (3.91)$$

and where  $\delta v_{y,\pm}$  is determined from Eqn. 3.87. We apply the perturbations for density, transverse velocity, and pressure with a discontinuity across  $y = 0$  following

$$r = \begin{cases} r_0 + \delta r_+ & \text{if } y \geq 0 \\ r_0 + \delta r_- & \text{if } y < 0 \end{cases} \quad r \in \rho, v_y, P \quad (3.92)$$

and use  $v_z = 0$  for a 2D KHI problem. For the shear velocity, give a thickness  $a$  to the shearing layer to initiate the mixing of the two layers, i.e.

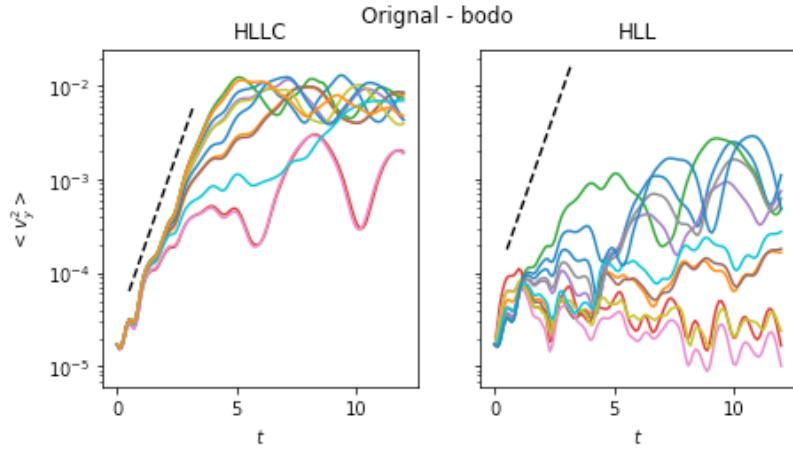
$$v_x = \frac{1}{2} (v_{x,+} - v_{x,-}) [\tanh(y/a) + 1] + u_{x,-}. \quad (3.93)$$

For symmetry in  $y$ , we mirror the perturbation on both sides of  $y = 0$ , placing the interface as described here at  $y = 0.5$  and another interface with a transverse velocity in the opposite direction at  $y = -0.5$ . We chose a shear velocity,  $V = 0.5c$ , adiabatic index  $\gamma = 4/3$  with ideal EOS, initial density  $\rho_0 = 1$ , initial pressure  $P_0 = c^2$ , a perturbation amplitude  $A = 0.05$ , and a shearing layer thickness  $a =$ . We use  $k = 2\pi$  so that the wavelength of the perturbations in  $x$  is 1. Simulations are performed on a  $[-0.5, 0.5] \times [-1.0, 1.0]$  domain using a mesh of square cells with twice as many cells in  $y$  than  $x$ , using mesh sizes in powers of 2 from  $128 \times 256$  to  $2048 \times 4096$  for a total of 5 different mesh sizes. For EMPIRE we test using the HLLC and HLLE Riemann solvers with basis orders 0, 1, and 2.

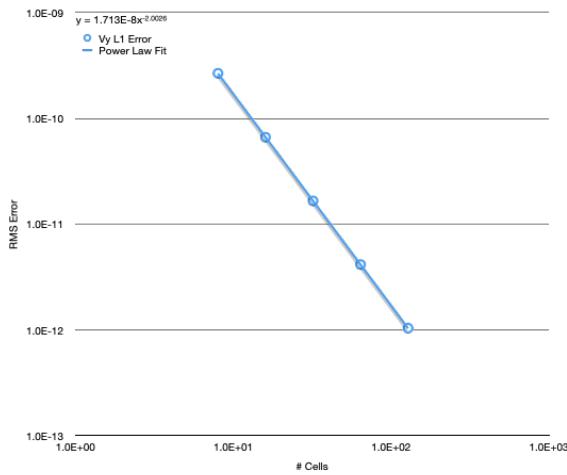
For each test, we start time at  $t = 0$  and run until  $t = 12$  to verify from the growth rate that the transverse velocity perturbations have saturated past the linear growth phase. We use the spatial average  $\langle v_y^2 \rangle$  to measure the linear growth rate of the KHI in all tests. Figure 3-11 shows  $\langle v_y^2 \rangle$  as a function of time. We measure the growth rate from  $t = 0.5$  to  $t = 3.2$  for all simulations. For comparison, we show the analytic growth rate  $\mathfrak{R}(\omega)$ , which depends on the shear velocity  $V$  used in specific perturbation, over the linear growth phase in a dashed black line. While the HLLC-based calculations reproduce the expected analytic growth rate, those based on the HLL-type Riemann solver do not, even for the high order basis; consistent with previous results in the literature, e.g. [5].

### Ordinary Wave Test

Kramer et al. [60] propose a suite of electromagnetic-fluid-plasma verification problems that can be used for numerical verification purposes. These verification problems amount to the propagation of small amplitude disturbances in the plasma and are founded upon the sole requirement that the particle velocities are small compared to the wave phase velocity,  $\mathbf{u} \ll \omega k$ .



**Figure 3-11.** Mean square of the transverse velocity  $\langle v_y^2 \rangle$  of the relativistic 2D Kelvin Helmholtz instability. The left panel shows  $\langle v_y^2 \rangle$  using the HLLC Riemann solver and the right panel shows  $\langle v_y^2 \rangle$  using the HLL Riemann solver. Different colors show different resolutions and different line styles show different orders of polynomial bases. Perturbation growth is greatly diminished when using the HLL solver.



**Figure 3-12.** Convergence of L1 Error norm for the transverse component of the electron three-velocity for the O-Wave convergence test as a function of spatial resolution, with implied temporal refinement. Overall second-order accurate convergence is observed, as expected for the combination of spatial and temporal discretization used in this test.

This same requirement allows an equivalent linearization of the *relativistic* fluid equations, i.e. by specifying that  $\gamma \sim 1$ . Here, we consider one wave in the hierarchy proposed by [60]: a cold ordinary wave, which in a transverse external magnetic field,  $\mathbf{B}_0 = (0, B_0, 0)$  has a dispersion relation:

$$\omega = \sqrt{\omega_{pe}^2 + c^2 k^2} \quad (3.94)$$

where  $\omega_{pe} = \sqrt{\mu_e^2 \rho_e \epsilon_0^{-1}}$  is the electron plasma frequency. For reference, the wave mode solution is:

$$\rho_e(\mathbf{x}, t) = \rho_0 \quad (3.95)$$

$$\mathbf{v}_e(\mathbf{x}, t) = \omega |\mathbf{k}|^{-1} \delta(0, \sin(\mathbf{k} \cdot \mathbf{x} - \omega t), 0) \quad (3.96)$$

$$\mathbf{E}(\mathbf{x}, t) = \mu_e \rho_0 (\epsilon_0 |\mathbf{k}|)^{-1} \delta(1 - \eta^2) (0, \cos(\mathbf{k} \cdot \mathbf{x} - \omega t), 0) \quad (3.97)$$

$$\mathbf{B}(\mathbf{x}, t) = \mathbf{B}_0 + \mu_e \rho_0 (\epsilon_0 \omega)^{-1} \delta(1 - \eta^2) (0, 0, \cos(\mathbf{k} \cdot \mathbf{x} - \omega t), 0) \quad (3.98)$$

where  $\mathbf{u}_e = \gamma_e \mathbf{v}_e$  and  $\gamma_e = 1 / \sqrt{1 - \mathbf{v}_e^2/c^2}$ . The conserved quantities are then initialized according to Eqn. 3.54. The system is initialized on a quasi-one-dimensional mesh, covering a domain  $[-0.5, 0.5] \times [-0.0625, 0.0625] \times [-0.0625, 0.0625]$  with tri-periodic boundary conditions with  $2^n \times 2^{n-2} \times 2^{n-2}$ ,  $n = 3, 4, 5, 6, 7$  zones and  $dt = 2^{-(n-2)} \times 2.4 \times 10^{-3}$  using a second order accurate basis for spatial discretization and a second order strong stability preserving (2,2,2) time integration scheme. The results of this test for the L1 error in the transverse component of the electron velocity are shown in Figure 3-12; overall second order accuracy is observed, as expected for the combination of spatial and temporal discretization used in this test.

### Space-Charge Limited Emission for a Cold Diode

Verification of space-charge limited emission (SCL) driven by the SCL boundary condition is performed through solution of a simple, space charge limited, one-dimensional (1D) relativistic diode problem. This problem considers unipolar flow of an ideal electron fluid,

$$w = \rho c^2 + \Gamma P / (\Gamma - 1) \quad (3.99)$$

$$P = \frac{\rho k_B T}{m} \quad (3.100)$$

where  $\Gamma = 4/3$  is the ratio of heat capacities,  $k_B$  is the Boltzmann constant and  $m$  is the electron mass. The anode plane is  $d = 1$  mm from the cathode plane and the circuit between the anode and cathode maintains a constant charging current  $i = 2.003525 \times 10^9$  A/m<sup>2</sup>. Initially, the electron density is,

$$\rho(x, t = 0) = 10^{17} m_e \text{ kg/m}^3, \quad (3.101)$$

where  $m_e$  is the mass of an electron. The initial electric field is set so that it is consistent with the initial electron density,

$$E(x, t = 0) = \frac{\rho q x}{m_e \epsilon_0}, \quad (3.102)$$

where  $q$  is the charge of an electron and  $\epsilon_0$  is the electric permittivity. The initial velocity of the electron fluid is  $v_0 = 1.876 \times 10^6$  m/s which corresponds to an electron with a kinetic energy of

10 eV and the initial temperature is  $T_0 = 5000$  K. The computational grid spacing is  $dx = 0.25 \mu\text{m}$  and the numerical parameters for the SCL boundary condition are,

$$\gamma_0 = 2000\epsilon_0 \quad (3.103)$$

$$\gamma_1 = 50\epsilon_0^2 \quad (3.104)$$

$$\tau_{ramp} = 1 \text{ ps} \quad (3.105)$$

$$v_{inj} = v_0 \quad (3.106)$$

$$T_{inj} = T_0. \quad (3.107)$$

The electron fluid boundary condition at the anode is a simple "copy out" outflow boundary.

It will be of interest to compare the solution to those of simple diode models; we note that an warm-fluid diode model, which takes into account the force due to fluid pressure is presented in the next section. The present case, however, is in a regime where this effect is small and so we proceed with a model which only accounts for electrostatic force. Given a uniform electric current density  $i$  and the flow velocity as a function of the electric potential,  $v(\phi)$ , the diode equation follows directly from Poisson's equation,

$$\phi''(x) = \frac{-i}{\epsilon_0 v(\phi)}, \quad (3.108)$$

where  $\epsilon_0$  is the electric permittivity. This relation holds for a 1D diode operating in a steady state. The boundary conditions for the SCL diode problem are,

$$\phi'(x = 0) = 0 \quad (3.109)$$

$$\phi(x = 0) = 0. \quad (3.110)$$

To obtain the relativistic diode equation, a relativistic expression for the velocity is required,

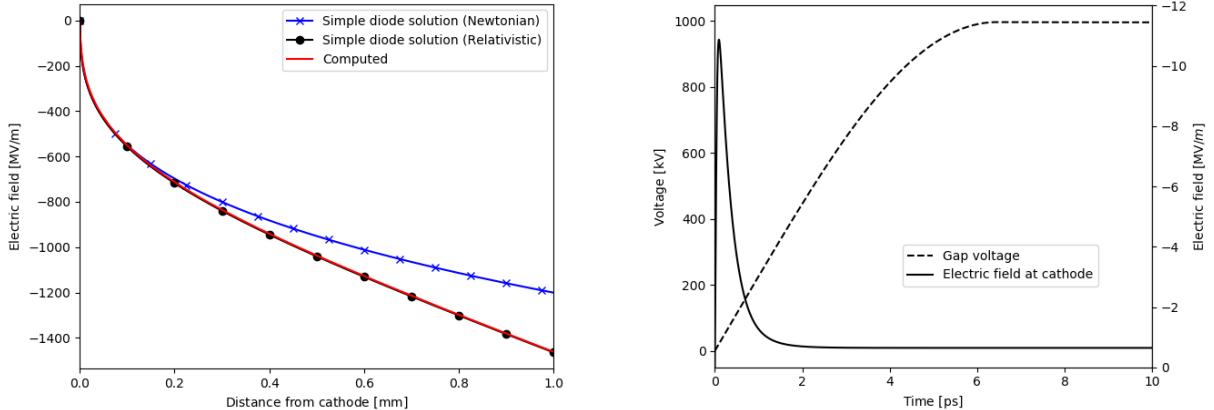
$$v(\phi) = \frac{c \sqrt{\left(\frac{w_0+q\phi}{mc^2}\right)^2 + \frac{2(w_0+q\phi)}{mc^2}}}{1 + \frac{w_0+q\phi}{mc^2}}, \quad (3.111)$$

where  $q$  and  $m$  are the charge and mass of an electron, and  $k_0$  is the kinetic energy of an electron when it is initially injected at the boundary. A non-relativistic model is obtained by using the corresponding Newtonian expression for velocity,

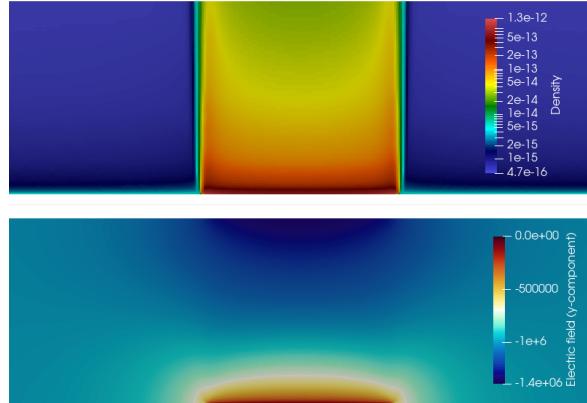
$$v(\phi) = \sqrt{\frac{2(w_0 + q\phi)}{m}}. \quad (3.112)$$

Since the SCL boundary condition requires a non-zero injection velocity, we set  $k_0 = 10$  eV; the solution will differ from the usual Child-Langmuir solution because the initial electron velocity is non-zero.

The steady-state solution to the present diode problem, which utilizes the relativistic fluid SCL boundary condition, is shown in the left-hand panel of Figure 3-13. The AK gap voltage in the simple relativistic diode solution is 1 MV, and the averaged root-mean-square (RMS) difference



**Figure 3-13. Solution for the 1D SCL diode problem. This is a benchmark problem for the relativistic SCL boundary condition. The exact solutions for simple relativistic and simple Newtonian diodes are shown for comparison.**



**Figure 3-14. Steady state solution for the 2D diode problem, computed using the relativistic fluid SCL boundary condition.**

between the fluid diode solution and the simple relativistic diode solution is only 0.43%. The simple non-relativistic diode solution differs from the fluid diode solution by 13.8%.

We also consider a two-dimensional (2D) SCL diode benchmark problem. We use the same electron fluid model as described for the 1D diode benchmark problem. The AK gap is 1 mm and SCL emission of electrons occurs on a strip that is 1 mm in width. There is a strong, uniform magnetic field normal to the electrode plates  $B_y = 1$  T. The magnetic field is intended to keep the electron beam from spreading, so that the current density  $\vec{i} \approx i(x)\hat{e}_y$ , where  $x$ -coordinate surfaces are parallel to the electrode planes and  $\hat{e}_y$  is the unit vector that is normal to the electrode surfaces. This 2D diode problem is similar to the spatially nonuniform current density strip problem considered by Luginsland et al. [69]. The solution is shown in Figure 3-14. It is similar to the result in [69].

### 3.1.3.4. Warm Fluid Diode Verification Study

In this section we describe the verification of EMPIRE-FLUID under mesh refinement for a one-dimensional warm diode. The diode is idealized as an electron emitting cathode and absorbing anode with a potential drop across the gap. The emitted electron fluid has a non-zero velocity and temperature greater than zero which implies a non-zero pressure.

Under steady-state conditions a semi-analytic solution can be derived that includes the pressure force in addition to electrodynamic forces. This problem was chosen because it provided a very good verification problem not only because there is an analytic solution but also because EM and pressure forces must balance to achieve a steady solution. Semi-analytic solutions including pressure gradients for non-relativistic and relativistic warm diodes are derived in Hamlin *et al.* [45] and the non-relativistic solution is summarized here.

In the non-relativistic regime, the following steady-state Euler-Poisson system of equations describe the fluid/electrodynamics of the diode:

$$m_e n_e v_e \frac{dv_e}{dx} = n_e e \frac{d\phi}{dx} - \frac{dP_e}{dx} \quad (3.113)$$

$$\frac{d^2\phi}{dx^2} = \frac{n_e e}{\epsilon_0} \quad (3.114)$$

$$P_e = \frac{k_B T_{e0}}{n_{e0}^{\Gamma-1}} n_e^\Gamma \quad (3.115)$$

$$J_0 = n_e e v_e = \text{constant.} \quad (3.116)$$

Eqn. (3.113) describes momentum transfer. Eqn. (3.114) is Poisson's equation. Eqn. (3.115) is the adiabatic equation of state assumption, where  $\Gamma$  is the adiabatic index,  $T_{e0}$  and  $n_{e0}$  are the injected electron temperature and density at the cathode and  $e$  and  $J_0$  denote the magnitude of electron charge and current density, respectively.

For a given electron current  $J_0$ , Eqns. (3.113) - (3.116) are solved in Mathematica subject to the initial conditions  $v_e(0) = v_{e0}$ ,  $\phi(0) = 0$ ,  $\phi'(0) = -E_0$ , and  $n_e(0) = n_{e0} = J_0/(ev_{e0})$ .

Each quantity in the solution, non-relativistic or relativistic, is then mapped onto a grid of 10,001 evenly-spaced points spanning the gap length  $d$ , and written to a csv file for comparison to the EMPIRE-FLUID solution. While both the non-relativistic and relativistic solutions are formulated using an electrostatic assumption, the solution procedure involves recasting the equations into an initial value problem. This is advantageous from a verification standpoint because EMPIRE-FLUID is solving Maxwell's equations using a semi-implicit solver and the inflow velocity is supersonic, therefore, the analytic solution and EMPIRE-FLUID can use exactly the same inflow variables.

The objective of this study was to verify spatial order of accuracy EMPIRE-FLUID for a coupled EM/fluid problem. The one-dimensional warm diode is parameterized by gap distance  $d = 0.01$  (m), voltage drop  $V_0 = 100$  (V), electron beam energy,  $W_0 = eV_0$  (J), electron temperature  $T_{e0}$ , adiabatic index  $\Gamma = 1.01$  and current density  $J_0 = \frac{16\epsilon ps_0}{9d^2} \sqrt{(2e/m_e)} (W_0/e)^{3/2}$  ( $A/m^2$ ), where  $J_0 = \frac{1}{2} J_{SCL}$ , is one half the space charge limit predicted by Jaffe [54] for the cold diode. This value was chosen to ensure that a steady solution

was obtainable. The velocity of the emitted electrons was determined from the beam energy  $v_{e0} = \sqrt{2W_0/m_e}$  where  $m_e$  is the electron mass and the number density was determined from the current density  $n_{e0} = J_0/(ev_{e0})$ . The adiabatic index was chosen to produce a nearly isothermal flow to prevent the formation of shock waves. The electric field at the emitting surface was obtained from the Jaffe solution.

Three inflow temperatures were chosen  $T_{e0} = 1, 10, 100$  (eV). The EMPIRE-FLUID simulations used IMEX time integration with an explicit time step restriction to march to steady-state, determined to be when the gap voltage reached a steady state value.

A summary of the results will now be presented. Additional results are presented in Smith *et al.* [124]. A comparison of the electric field and electron pressure profiles, normalized by the flow values, with the analytic solution for three temperatures is presented in Figures 3-15. For  $T_{e0} = 1$  eV, the solutions are almost identical to the cold diode solutions [54, 102]. Those solutions [125, 44] were used as comparison checks during the development of the current pressure gradient solution. At  $T_{e0} = 10$  eV, a slight separation between  $T_{e0} = 1$  eV and  $T_{e0} = 10$  eV begins to appear and finally, at  $T_{e0} = 100$  eV the solutions are distinctly different. The black colored curves on these plots represent the analytic solution and apparently, the EMPIRE-FLUID and analytic solutions are in good agreement.

A series of six hierarchical computational grids, where the number of cells in the  $x$ -direction was doubled for each successive mesh, were used to establish spatial convergence properties of EMPIRE-FLUID for this problem. The finite element basis order for the fluid equations was one, therefore, producing a second-order accurate fluid solver. The basis order for Maxwell's equations was piecewise constant and so the expected order was one. A discrete  $L_1$  error norm was defined as

$$\mathcal{E}_{L1} = \frac{1}{N_x} \sum_{i=1}^{N_x} | \langle f(x) \rangle_i - \langle f^h \rangle_i |$$

where  $\langle f(x) \rangle_i$  is the cell average of the analytic solution computed internally using Gauss quadrature of table values interpolated to element nodes, and  $\langle f^h \rangle_i$  are the cell averaged values of the EMPIRE-FLUID solution for cell  $i$ . Figures 3-16 show the relative error profiles ( $\text{error}_i^{\text{rel}} = |\langle f(x) \rangle_i - \langle f^h \rangle_i| / |\langle f(x) \rangle_i|$ ) from solutions on each mesh for the electric field and electron pressure (top row), and the error norms for electric field and electron pressure (bottom row) for  $T_{e0} = 1$  eV. Similar results for electron number density and velocity were observed and these trends continue as the temperature is raised to  $T_{e0} = 10, 100$  eV. These results demonstrate consistent monotonic spatial convergence. Finally, the rates of  $L_1$  cell based error norm convergence for the electric field and electron number density, velocity and pressure at the three temperatures tested, for the three most refined mesh level pairs are presented in Table 3-3. Second order accuracy is predicted for all four quantities. The prediction that the electric field is converging at a rate of two is somewhat surprising. This may simply be due to how the error norms are defined. In a different analysis, the  $L_1$  norms are generated internally using Gauss quadrature to integrate the error within each element. This analysis predicts a rate of two for the primitive flow quantities and a rate of one for the electric field.

This verification exercise produced a formal spatial accuracy verification of the EMPIRE-FLUID multifluid code that included coupled EM/fluid solutions with non-trivial boundary conditions to

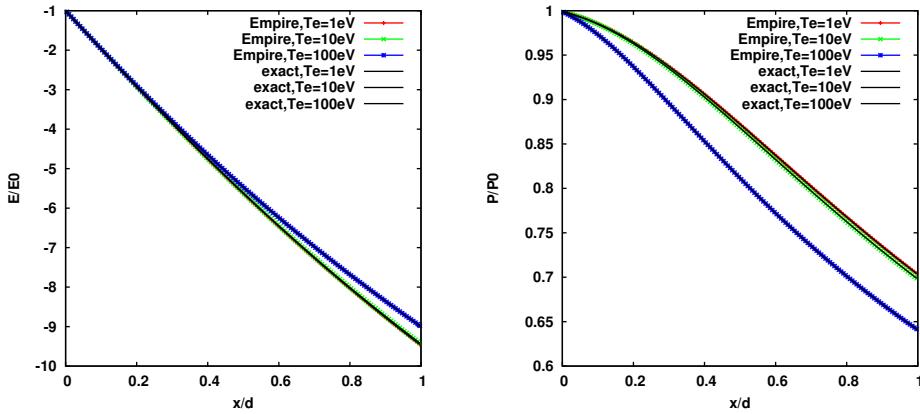


Figure 3-15. Warm diode solutions for electric field and electron pressure for  $T_e = 1, 10, 100 (eV).$

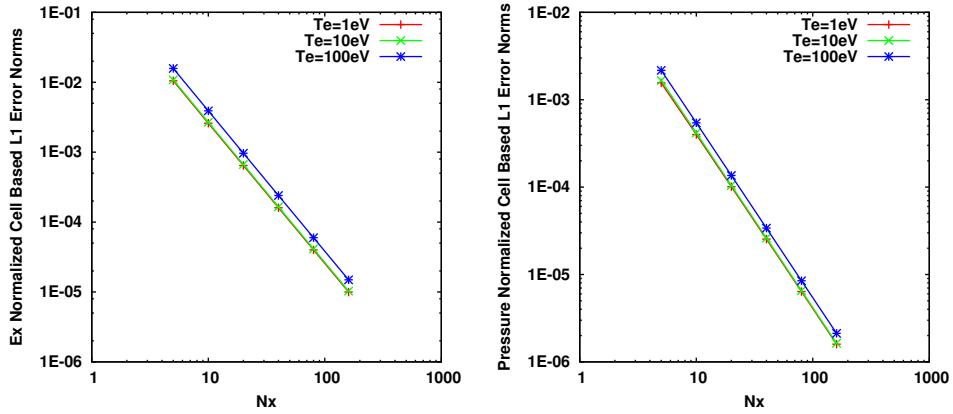


Figure 3-16. Cell based norms (bottom) for electric field and electron pressure for a set of hierarchical mesh levels and  $T_e = 1 (eV).$

Temperature (eV)	Mesh pair	E field	number density	velocity	pressure
1	20-40	2.006	2.000	1.997	1.992
	40-80	2.003	2.000	1.999	1.996
	80-160	2.002	2.000	2.000	1.999
10	20-40	2.006	1.997	2.000	2.000
	40-80	2.003	1.998	2.000	2.000
	80-160	2.002	1.999	2.000	2.000
100	20-40	2.009	2.002	1.996	2.002
	40-80	2.005	2.000	1.997	2.000
	80-160	2.003	2.001	2.000	2.001

Table 3-3. Rates of convergence of the cell based  $L_1$  error norms for electric field, electron number density, velocity and pressure.

expected order of accuracy. The warm diode has become an archival test using the VVTest testing harness, being executed on a regular basis to help prevent code regression. It also laid the foundation for future studies which may include additional plasma phenomena such as inter-species collisional interactions and relativistic fluid velocities.

A relativistic warm diode verification test has recently been created. Preliminary runs have exposed a subtle issue with how the rest mass energy is being handled when the code option to subtract it from the total energy density is used. At the time of writing this report, reasonable convergence rates for primitive variables has been observed. These results have provided valuable feedback for improving the semi-analytic solution as well. Additional work is ongoing to demonstrate convergence of the conserved variables and complete the test. This will be reported in [44] which is currently under preparation.

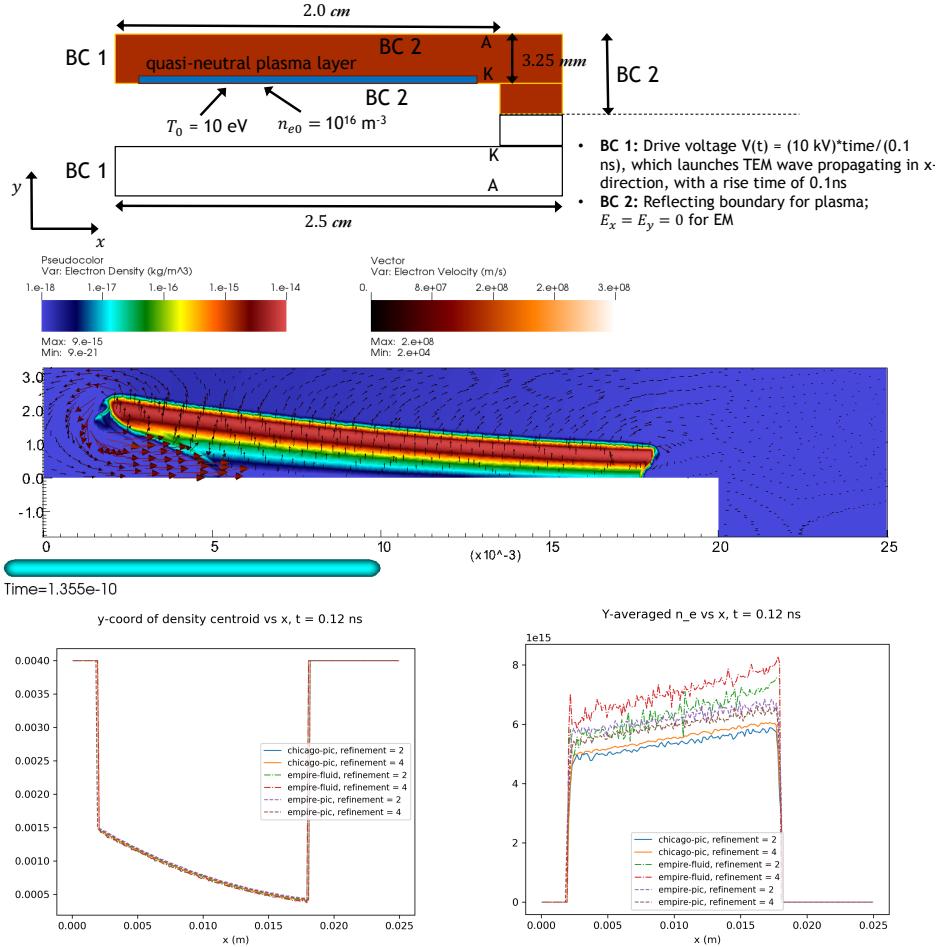
### **3.1.3.5. Planar MITL calculations with Relativistic Fluids and Future Work**

The target application for the relativistic fluid algorithm development conducted under this LDRD is power flow within the “inner MITL” region for the Z-Machine. In the preceding sections, we have described the algorithms that were implemented within the EMPIRE-FLUID code to achieve this goal and a sample of the verification exercises that were conducted in order to build confidence in the implementation of these algorithms within EMPIRE-FLUID.

As a next step in the verification hierarchy, the relativistic fluid algorithm was applied to the evolution of an initial plasma layer located within a simplified planar MITL geometry and the results compared to those produced by the CHICAGO and EMPIRE-PIC codes. The geometry, transmission line drive and results from this calculation are shown in Figure 3-17. The data of this figure suggest that the evolution of the initial quasi-neutral plasma layer is equivalent between the three different codes; this result serves to build confidence that the relativistic fluid algorithms implemented within EMPIRE-FLUID can accurately evolve plasmas found within the inner MITL region.

The relativistic fluid algorithms were then applied to the problem of space-charge limited emission of electrons from the cathode in the planar MITL geometry shown in Figure 3-18 using a Z-relevant drive (shown in the figure). Electrons were injected from the cathode into the MITL with a temperature of 10 eV at Mach 2. Neutralization of the electric field at the cathode surface was observed, followed by dynamic evolution reminiscent of the particle-in-cell calculations described later in this report for this same geometry. Follow-on work from this LDRD will focus on benchmarking the evolution of the electron plasma and performing detailed comparison with results of particle-in-cell calculations.

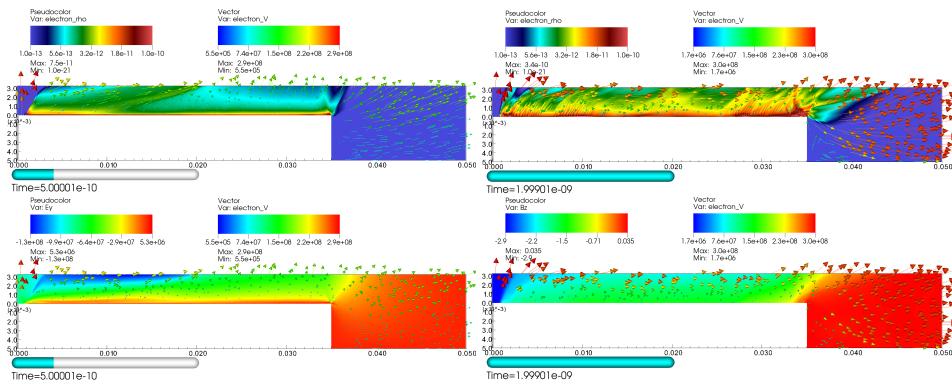
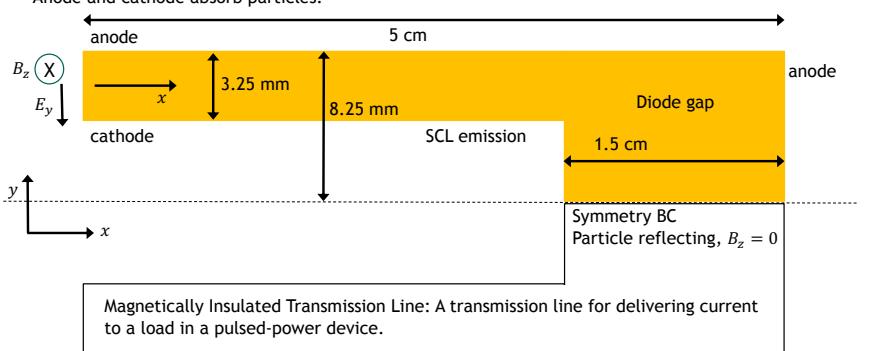
In summary, this LDRD has developed a unique relativistic fluid modeling capability for power-flow applications, providing the foundations for a fast tool that enables designers to explore design space quickly and robustly, before engaging with high-fidelity kinetic particle-in-cell calculations. In addition, the capabilities in this tool provide a robust foundation to unify kinetic and fluid modeling efforts through the EMPIRE-HYBRID component; such calculations could take the form of  $\delta$ -f type approaches such as detailed in the next section, or species differentiated hybridization, where (for example) relativistic electrons are treated using



**Figure 3-17. Evolution of an initial quasi-neutral plasma layer located within a simplified planar MITL geometry (top panel). The center panel shows the density distribution (filled contours) and plasma velocity at late times within the MITL, after peak voltage has been reached. Lower panels show the location of the density weighted centroid of the layer (left panel) and average density within the layer (right panel) at late times compared with CHICAGO and EMPIRE-PIC computations of the same problem at varying levels of refinement.**

- Driven voltage:  $V(t) = (10 \text{ MV}) \frac{t}{10 \text{ ns}}$
- Anode and cathode absorb particles.

Yellow region is the vacuum gap being simulated.



**Figure 3-18.** Evolution of the planar MITL driven by a Z-relevant voltage with space-charge limited emission of electrons from the cathode. The top panel shows the geometry, boundary conditions and drive. The center panels show the electron density distribution and velocities at 0.5ns (left) and 2ns (right), while the bottom panels show the electric field normal to the cathode at 0.5ns (left) and the magnetic field into the plane at 2ns (right).

the fluid modeling capability described here, while slow ion species are treated by kinetic particle-in-cell techniques.

In order to achieve these goals, we have identified a selection of near-term follow-on efforts that are necessary to deliver a robust, performant modeling capability:

- Performance enhancements: in order for the fluid capability to provide designers with a tool fast enough to explore design space, time-to-solution has to be shorter than particle-in-cell. Currently, this is limited by the need to treat the plasma and cyclotron frequencies in an explicit fashion, due to expensive iterations in the conserved-to-primitive variable solver. Future work should focus on mitigating this cost through new non-linear solver techniques, e.g. methods that integrate the non-linear relationship between conserved and primitive quantities in relativistic hydrodynamics.
- Robustness and accuracy enhancements: stabilization techniques for discontinuous Galerkin discretizations (e.g. limiting) can be shown to violate charge-conservation when combined with the Maxwell solvers present within EMPIRE. New techniques (such as local diffusion currents derived from the limiter) should be developed and verified that mitigate this effect; such an activity would have impact across a broad range of activities of interest to Sandia National Laboratories.

In addition to these efforts, further development work focussed on coupling of thermal desorption boundary conditions that account for electron heating of electrode surfaces along with efforts to ensure conservation of total mass, momentum and energy across multiple ion species, electrons and electromagnetic fields will be necessary to accurately capture the behavior of dense plasmas at electrode boundaries. Finally, continuous development and application of new verification problems will be needed to ensure that the results delivered by these algorithms are understood across the range of conditions relevant to power flow in the inner MITL on Z.

### **3.1.4. EMPIRE-HYBRID**

There are a wide class of problems for which the state of the plasma varies in density, temperature, and velocity by many orders of magnitude within the problem domain. The case of power flow and subsequent load dynamics on large pulsed power accelerators is an extreme example of this type of problem.

Traditionally, these sorts of problems are addressed by decomposing the modeling task into a set of domains over which the plasma reasonably falls into one of the traditional categories of a weakly collisional kinetic, two-fluid, or MHD model. The results of these models are then stitched together after the fact, often using reduced models for combining them, e.g. using transmission line models for coupling the electromagnetic wave propagation through the early parts of the MITLs [79] to a PIC code for modeling power flow, or a circuit model for coupling the power flow to an MHD code for modeling load dynamics [64, 98, 96]. There are indications, however, that we are pushing the limits of this approach on Z today and this problem is sure to become significantly more challenging as we look to the future of a higher power machine. These challenges are driving us to consider a new approach to modeling multi-scale plasmas.

In this section we describe research into integrating fluid and kinetic descriptions of a plasma in a self-consistent, coupled fashion. There are a variety of approaches that have been reported in the literature [4, 8, 65, 28, 34, 35, 143]. While each of these approaches has its own set of strengths and weaknesses, the primary challenge for coupling fluid and kinetic descriptions, or even a two-fluid to MHD model, is consistently incorporating the model reduction in the coupling and whether this consistency is maintained by the reduced or higher model.

To see why, consider the case of coupling two-fluid and MHD plasma models using a domain decomposition approach. In this approach, the solution domain is decomposed into two-fluid and MHD sub-domains. To utilize this approach, one must construct a procedure to define the decomposition, adjusting the domain boundaries, transferring plasma back and forth between a two-fluid and MHD model, and finally coupling the evolution of each model across the domain interfaces. This is where one finds the challenges. For example, when the domain boundary moves and one must turn a two-fluid plasma into an MHD plasma, what should be done with the excess charge? Is the MHD model extended to incorporate a finite charge density, or should this excess charge be swept to the domain boundary, consistent with a good conductor? Also, if plasma flows across the boundary between a two-fluid and MHD domain, what is done with the excess charge? From a physical perspective, these should all be small terms if the domain boundaries are selected in accordance with the model reduction approximations, but how do these small corrections affect the numerical stability and solution convergence properties? It is for these sorts of reasons that in this project we decided to pursue a  $\delta f$  approach for coupling fluid and kinetic plasma descriptions.

In the sections that follow, we outline the basic formalism for a  $\delta f$  fluid/kinetic model. At a high level, the reduced (fluid) model is treated as an approximation and the consistency in the models is maintained by the higher (kinetic) model. This choice is possible because both the fluid and kinetic models exist over the entire simulation domain. In the sections that follow, we describe the technical details for implementing this system of equations, present results, and outline future research directions.

### 3.1.4.1. Nonrelativistic Vlasov $\delta f$

In this section we outline the basic  $\delta f$  formalism for the non-relativistic Vlasov equation

$$\partial_t f_\alpha + u \cdot \nabla_x f_\alpha + \left( \frac{q_\alpha}{m_\alpha} \right) (E + u \times B) \cdot \nabla_u f_\alpha = C_\alpha. \quad (3.117)$$

The variable  $f_\alpha(t, x, u)$  is the single-particle distribution function for species  $\alpha$  (electrons, ions, ...) and is a function of time  $t$ , position  $x$ , and velocity  $u$ . The quantity  $m_\alpha$  is the particle mass,  $q_\alpha$  is the particle charge,  $(\mathbf{E}, \mathbf{B})$  are the electromagnetic fields, and  $C_\alpha$  is the collision operator. To simplify the notation, we define the Vlasov operator for species  $\alpha$

$$L_\alpha = \partial_t + u \cdot \nabla_x + \left( \frac{q_\alpha}{m_\alpha} \right) (E + u \times B) \cdot \nabla_u \quad (3.118)$$

such that the Vlasov equation can be written as

$$L_\alpha f_\alpha = C_\alpha. \quad (3.119)$$

The  $\delta f$  formalism follows from decomposing the distribution function

$$f_\alpha = M_\alpha + \delta f_\alpha \quad (3.120)$$

where  $M_\alpha$  represents a Maxwell-Boltzmann and  $\delta f$  represents the nonequilibrium corrections. The Vlasov equation thus becomes

$$L_\alpha \delta f_\alpha + L_\alpha M_\alpha = C_\alpha \quad (3.121)$$

and this is the essence of the  $\delta f$  fluid/kinetic method.

There are many different ways to discretize Eqn. 3.121. In the work presented here, the Maxwellian,  $M_\alpha$ , is a functional of the fluid solution and the correction,  $\delta f$ , is represented with a PIC discretization. The PIC discretization can be constructed following the same procedure that the standard PIC system can be derived from the Vlasov equation, e.g. [38]. Doing so results in a modified PIC system of equations where  $L_\alpha M_\alpha$  results in a term that can either be interpreted as a weight evolution or particle generation term. Alternatively, one could simply use operator splitting to write Eqn. 3.121 as

$$L_\alpha \delta f_\alpha = 0 \quad (3.122)$$

$$\frac{\partial}{\partial t} \delta f_\alpha = C_\alpha - L_\alpha M_\alpha \quad (3.123)$$

and see that the basic structure of this system can be thought of as a standard, collisionless PIC evolution, as expressed by Eqn. 3.122, operator split from  $\delta f$  particle generation (or weight evolution) and collisions as expressed by Eqn. 3.123. To close this system, we couple it to a fluid system of the form

$$\frac{\partial q_\alpha}{\partial t} + \nabla_x \cdot g(q_\alpha) = S(q_\alpha) \quad (3.124)$$

where  $q_\alpha = (\rho_\alpha, \rho_\alpha v_\alpha, \mathcal{E}_\alpha)$ , the vector of conserved variables for species  $\alpha$ ,  $g(q_\alpha)$  is the associated vector of hydrodynamic fluxes,  $S(q_\alpha)$  represents a source term that could include the Lorentz force, collisional momentum and energy exchange, etc. and, of course, this system is closed with some appropriate equation of state.

Hence, at its core, constructing a  $\delta f$  fluid/kinetic method, when using operator splitting, can be seen as solving Eqn. 3.123 as a method for coupling a collisionless PIC system, Eqn. 3.122, with a fluid plasma system, Eqn. 3.124. The Monte Carlo sampling of  $L_\alpha M_\alpha$  in velocity-space and normalization for particle generation is discussed in §3.1.4.3. We discuss the space-time discretization of  $L_\alpha M_\alpha$  using a strong form in §3.1.4.4 and a weak form using the same space-time mesh as the Finite Volume (FV) method in §3.1.4.5. Extending these methods from FV to a Discontinuous Galerkin (DG) method should follow naturally. Before diving into these details, however, in §3.1.4.2 we address some high-level questions regarding the fluid equations and how the evolution of this system affects the  $\delta f$  particle generation in the coupled system.

### 3.1.4.2. High Level $\delta f$ System Considerations

When one sets out to implement the  $\delta f$  system of equations 3.123 - 3.124, there are a number of questions that arise regarding the implementation and the answers to these questions have important implications for the numerical method. Consider the Euler / Boltzmann system in the collisionless limit. Should the fluid system be evolved in terms of the ideal, collisional Euler system or should we interpret the fluid system as a set of moment equations for the kinetic system? The implication of this decision can be readily understood by considering a *gedanken* experiment.

The *gedanken* experiment consists of a hydrodynamic Riemann problem. Described from the perspective of a fluid, the initial conditions consist of a one-dimensional problem in which the velocity is zero, the pressure is constant and the density is piece-wise constant with a discontinuity at  $x = 0$ . This is, of course, a stationary contact discontinuity in the ideal, Euler equations. A consistent description of the distribution function is a Maxwell-Boltzmann that locally matches the mass, momentum and energy density. This is not, however, a stationary solution for the Boltzmann equation, except in the limit of a vanishingly small collisional mean free path.

If we choose to evolve the fluid system using the ideal, Euler system of equations, then, as noted above, the Euler system remains stationary. When we evaluate  $LM$ , however, the “upwind” kinetic flux will have mass, momentum and energy flux “moments” that differ from the ideal Euler fluxes. As a result, the moments of  $\delta f$  generated at the end of the time-step will be non-zero. For a hydrodynamic system, this, in and of itself, is not a big deal. The biggest problem is the fact that the fluid and kinetic solutions are diverging as time evolves. As a result, maintaining  $\delta f/f < 1$  will require adding an additional flux correction process to correct the fluid mass, momentum and energy for these flux differences.

If we consider the same *gedanken* experiment for the 2-fluid / Vlasov system, assuming a neutral plasma in the initial state, we find much the same. The fluid systems are initially stationary while the kinetic systems are not. Evaluating  $LM$  and updating  $\delta f$  again results in mass, momentum and energy density moments for  $\delta f$  that are non-zero. Unlike the neutral gas system, however, these flux differences are associated with a current density correction that must be incorporated into the evolution of the electromagnetic field.

These considerations imply that there are two basic approaches that one could follow when discretizing the  $\delta f$  fluid / kinetic system. In the first approach, one could choose a Riemann solver based on the upwind distribution function for a collisionless kinetic system. This has the benefit that both the kinetic and fluid systems have a consistent description of the distribution function at the interface. As a result,  $\delta f$  particles generated on zones have zero mass, momentum or energy. The associated Riemann solver, used to compute hydrodynamic fluxes, is described as a kinetic flux vector splitting method [92, 72, 142, 70]. One downside of this approach is that in the highly collisional limit, contact discontinuities will diffuse because of the dissipative Riemann flux. The second approach that one could follow is essentially to treat the fluid system (equation 3.124) as a predictor step and the particle generation (equation 3.123) as a corrector step. In this approach, any Riemann solver can be used to advance the fluid system, even one that maintains a stationary discontinuity. Consistency with the kinetic system, however, then dictates that  $\delta f$

particles are generated with finite mass, momentum and energy density. In one method of this type,  $\delta f$  particles are generated on both zones and interfaces at the mid-point of the time-step for a second order method. In §3.1.4.5 we present results from this second approach.

### 3.1.4.3. $\delta f$ Particle Generation

In this section we take up the topic of Monte Carlo sampling of Eqn. 3.123 in the absence of collisions. The  $\delta f$  particle distribution function, dropping the subscript  $\alpha$  for notational simplicity, is

$$\delta f = \sum_{r=1}^{N_T} w_r \prod_{i=1}^D \delta(x^i - x_r^i) \prod_{j=1}^3 \delta(u^j - u_r^j) \quad (3.125)$$

in a  $D$ -dimensional physical space, 3-dimensional velocity space, and with particles numbered by the index  $r$  with weight  $w_r$ , position components  $x_r^i$  and velocity components  $u_r^j$ . Substituting this expression into Eqn. 3.123 and integrating over all velocity space we find

$$\partial_t \left[ \sum_{r=1}^{N_T} w_r \prod_{i=1}^D \delta(x^i - x_r^i) \right] + \int LM d^3 u = 0 \quad (3.126)$$

Integrating over the discretization volume,  $\Omega_k$ , appropriate for  $LM$

$$\partial_t \left[ \sum_{r=1}^{N_p} w_r \right] + \int_{\Omega_k} \int LM d^3 u d^D x = 0 \quad (3.127)$$

where the particles  $r = (1, N_p)$  have  $x_r^i \in \Omega_k$ . Integrating over a time step from  $t^n$  to  $t^{n+1}$  we obtain

$$\sum_{r=1}^{N_p} [w_r(t^{n+1}) - w_r(t^n)] + \int < LM >_k d^3 u = 0 \quad (3.128)$$

where in the last step we've introduced the space-time integral  $< LM >_k$  as

$$< LM >_k = \int_{t^n}^{t^{n+1}} \int_{\Omega_k} LM d^D x dt. \quad (3.129)$$

In the next step, we divide the space-time integral  $< LM >_k$  into the sum of positive and negative contributions

$$\Sigma_k^\pm = \text{Max}(\pm < LM >_k, 0) \quad (3.130)$$

such that

$$< LM >_k = \Sigma_k^+ - \Sigma_k^-. \quad (3.131)$$

With this definition, we can split the particles into those associated with  $\Sigma_k^\pm$  and rewrite Eqn. 3.128 as

$$\sum_{r=1}^{N_p^\pm} [w_r^\pm(t^{n+1}) - w_r^\pm(t^n)] + \int (\pm \Sigma_k^\pm) d^3 u = 0. \quad (3.132)$$

This is the point at which we need to choose whether to interpret this equation in the sense of weight evolution, or particle generation.

Following the particle generation path, these weights represent new  $\delta f$  particles and hence  $w_r^\pm(t^n) = 0$ . Furthermore, if we choose particle velocities at random according to the scaled probability distribution functions  $\Sigma_k^\pm$ , i.e. use importance sampling, then each particle has the same weight [2] and we obtain the final expression

$$w_r^\pm(t^{n+1}) = \left( \frac{-1}{N_p^\pm} \right) \int (\pm \Sigma_k^\pm) d^3 u. \quad (3.133)$$

Note, that selecting particle velocities according to  $\Sigma_k^\pm$  can be accomplished using the rejection method [90] and we can estimate the integrals of  $\Sigma_k^\pm$  using a Monte Carlo approximation. As with initializing a distribution function in PIC simulations, however, enforcing exact conservation is desirable.

Exact conservation can be enforced by referring back to the fluid equations. First, note that the moments

$$\int \langle LM \rangle_k h(u) d^3 u, \quad (3.134)$$

where  $h(u) = m\{1, u, u^2/2\}$ , can be evaluated analytically and reduce to the ideal Euler or two-fluid equations when considering the Boltzmann or Vlasov equations respectively, evaluated on the domain  $\Omega_k$ . Furthermore, if this domain is the same domain used to advance the fluid equations, then these moments are identically zero and we are lead to the the conclusion (and constraint) that the  $\delta f$  particles must be generated with zero net mass, momentum or energy. If the domain  $\Omega_k$  differs from the domain used to advance the fluid system, e.g. a dual-mesh discretization of  $LM$ , or if the moments of the kinetic flux differs from the hydrodynamic fluxes, then the moments of the generated  $\delta f$  particles will generally have a finite value for the net mass, momentum and energy. In either case, note that these moments can also be expressed in terms the generated  $\delta f$  particles resulting in a set of constraint equations which can be satisfied by adding a correction to the particle weights, average and particular velocity components.

Before proceeding, it is worth highlighting the implications just mentioned for the case in which  $\delta f$  particles are generated with finite mass, momentum or energy. For the two-fluid / Vlasov system, this results in the generation of  $\delta f$  particles with a finite charge and current density. This in turn implies that there is an additional current density associated with the flux of  $\delta f$  particles that must be incorporated into Ampère's equation.

#### 3.1.4.4. Strong Form Discretization

In this section we consider a strong form discretization of the term  $LM$  as a reasonable approach for linear problems and a simple approach to begin testing the  $\delta f$  fluid / kinetic method. In particular, we consider the case of a collisionless plasma and in what follows we neglect the species subscript  $\alpha$  for notational simplicity. This analysis begins by considering the Maxwell-Boltzmann distribution function

$$M[\rho, \rho v, \mathcal{E}, u] \quad (3.135)$$

to be a functional of mass density, momentum density, total energy density and velocity. With that in mind, one can calculate the time derivative as

$$\partial_t M = \left( \frac{\partial M}{\partial \rho} \right) \partial_t \rho + \left( \frac{\partial M}{\partial \rho v_i} \right) \partial_t \rho v_i + \left( \frac{\partial M}{\partial \mathcal{E}} \right) \partial_t \mathcal{E}. \quad (3.136)$$

Substituting Eqn. 3.124 for the time-derivatives of the conserved variables, this equation can be written as

$$\begin{aligned} \partial_t M = & \left( \frac{\partial M}{\partial \rho} \right) [-\nabla_x \cdot g(\rho)] \\ & + \left( \frac{\partial M}{\partial \rho v_i} \right) [S(\rho v_i) - \nabla_x \cdot g(\rho v_i)] \\ & + \left( \frac{\partial M}{\partial \mathcal{E}} \right) [S(\mathcal{E}) - \nabla_x \cdot g(\mathcal{E})]. \end{aligned} \quad (3.137)$$

Next, one can show that if  $S(\rho v_i)$  and  $S(\mathcal{E})$  represent the force and work done by the Lorentz force, then

$$\left( \frac{\partial M}{\partial \rho v_i} \right) S(\rho v_i) + \left( \frac{\partial M}{\partial \mathcal{E}} \right) S(\mathcal{E}) + \left( \frac{q}{m} \right) (E + u \times B) \cdot \nabla_u M = 0. \quad (3.138)$$

Thus the final result is

$$\begin{aligned} LM = & \left( \frac{\partial M}{\partial \rho} \right) [u \cdot \nabla_x \rho - \nabla_x \cdot g(\rho)] \\ & + \left( \frac{\partial M}{\partial \rho v_i} \right) [u \cdot \nabla_x (\rho v_i) - \nabla_x \cdot g(\rho v_i)] \\ & + \left( \frac{\partial M}{\partial \mathcal{E}} \right) [u \cdot \nabla_x \mathcal{E} - \nabla_x \cdot g(\mathcal{E})]. \end{aligned} \quad (3.139)$$

The remaining task of calculating the partial derivatives of the Maxwell-Boltzmann distribution is tedious, yet straight-forward.

### 3.1.4.5. Weak Form Discretization

In this section we outline a weak-form discretization of the  $\delta f$  fluid / kinetic equations and present sample results for a neutral gas in the collisionless limit. The fluid solver is a two-step method described in [127] using a local Lax-Friedrichs Riemann solver. Eqn. 3.123 is discretized on the same spatial zones and time-step as the fluid solver and is evaluated in the predictor / corrector sense using a second-order, endpoint approximation. The particle generation and transport steps are integrated using Strang splitting and resolved in three velocity space dimensions.

The first problem we consider is that of a contact discontinuity simulated in one spatial-dimension on the domain  $(0 \leq x \leq 6)$  and three velocity-dimensions. The initial conditions are a piece-wise constant, equilibrium Maxwell-Boltzmann distribution function with a discontinuity at  $x = 3$ . In

this normalized set of units, the particle mass is 1, the velocity is 0, the pressure is 0.6, the ratio of specific heats  $\gamma = 5/3$ , and the mass density

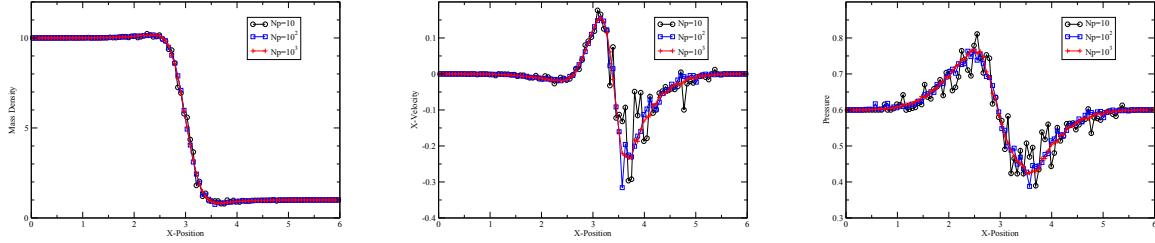
$$\rho = \begin{cases} 10 & \text{if } x < 3, \\ 1 & \text{if } x > 3. \end{cases} \quad (3.140)$$

The analytic solution to the collisionless Boltzmann equation for these initial conditions is

$$f = \begin{cases} M_L & \text{if } (x-3)/t < u_x, \\ M_R & \text{if } (x-3)/t > u_x \end{cases} \quad (3.141)$$

where  $M_{L,R}$  represent the initial Maxwell-Boltzmann in the left and right state and  $u_x$  is the  $x$ -velocity coordinate.

Figure 3-19 presents results for the mass density, velocity and scalar pressure at time = 0.96 including both the fluid and  $\delta f$  particle distributions. The particle generation counts,  $N_p = (10, 10^2, 10^3)$ , are the number of  $\delta f$  particles generated on interfaces and zones at each time-step. The spatial domain is resolved on 100 zones and the time-step remains constant for these calculations. Hence, this figure gives an indication of the self-convergence with respect to particle generation.

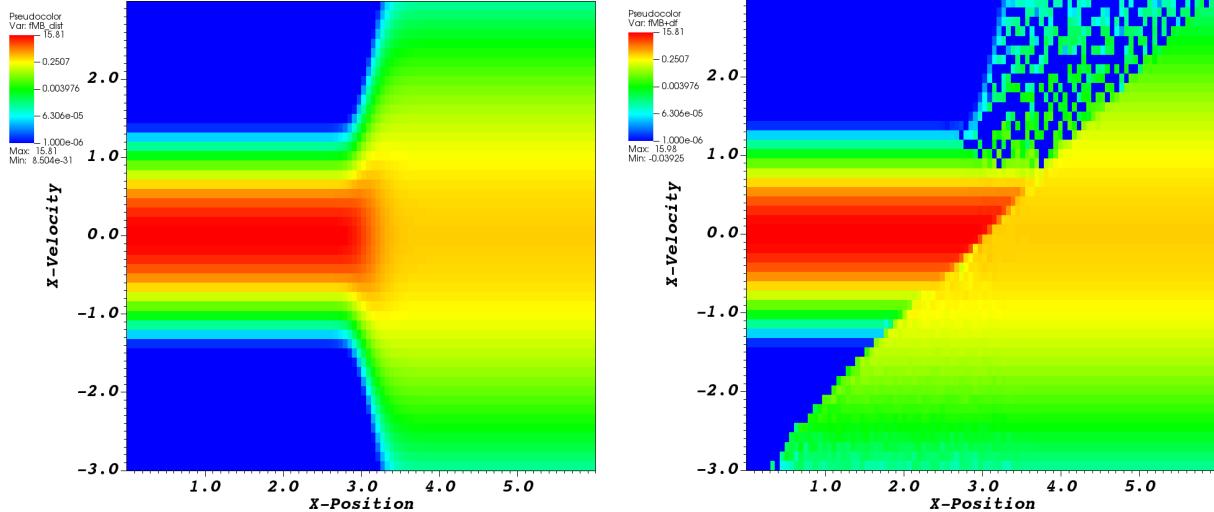


**Figure 3-19. Mass density (left), x-velocity (center) and scalar pressure (right) for the hybrid solution of a contact discontinuity for different particle generation counts.**

Figure 3-20 presents the logarithm of the distribution functions, integrated over transverse velocity coordinates  $(u_y, u_z)$ , as a function of  $(x, u_x)$  for  $N_p = 10^3$ . The fluid distribution function, a Maxwell-Boltzmann, shows a bit of dissipation of the contact discontinuity resulting from the use of the Lax-Friedrichs flux. The fluid plus  $\delta f$  distribution shows a reasonable estimate for the analytic solution in Eqn. 3.141.

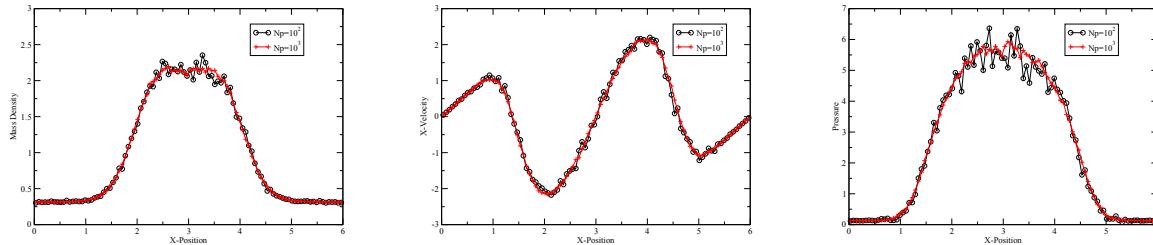
The second problem we consider is that of a supersonic, sinusoidal velocity variation simulated in one spatial-dimension on the domain  $(0 \leq x \leq 6)$  and three velocity-dimensions. The initial conditions are given by a Maxwell-Boltzmann distribution function with a uniform density and pressure and a sinusoidal  $x$ -velocity. In a normalized set of units, the particle mass is 1, the mass density is 1, the transverse velocity components  $v_y = v_z = 0$ , the pressure is 0.6, the ratio of specific heats  $\gamma = 5/3$ , and the  $x$ -velocity,

$$v_x = 4 \sin\left(\frac{2\pi x}{6}\right). \quad (3.142)$$



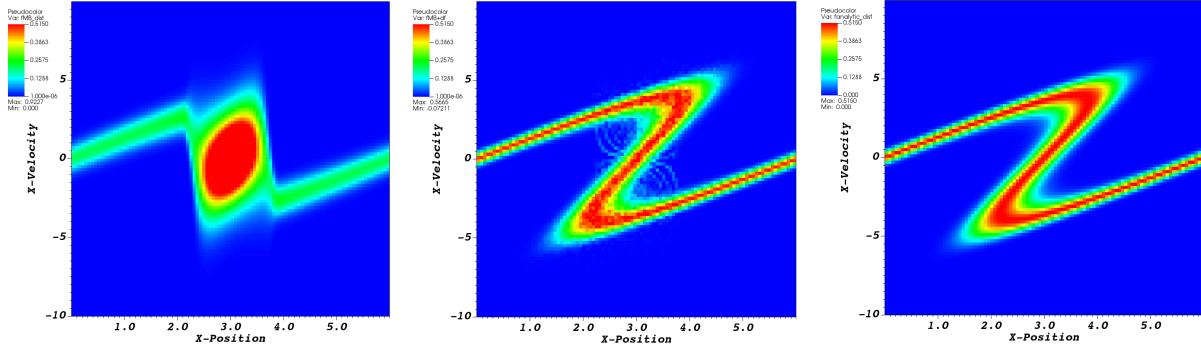
**Figure 3-20. Logarithm of the distribution function for the fluid (left) and fluid +  $\delta f$  (right) for the hybrid solution of a contact discontinuity for  $N_p = 10^3$  particle generation count.**

Figure 3-21 presents results for the mass density, velocity and scalar pressure at time = 0.53 including both the fluid and  $\delta f$  particle distributions. The particle generation counts,  $N_p = (10^2, 10^3)$ , are the number of  $\delta f$  particles generated on interfaces and zones at each time-step. The spatial domain is resolved on 100 zones and the time-step remains constant for these calculations. Hence, this figure gives an indication of the self-convergence with respect to particle generation.



**Figure 3-21. Mass density (left), x-velocity (center) and scalar pressure (right) for the hybrid solution of an initially sinusoidal velocity profile for different particle generation counts.**

Figure 3-22 presents the fluid, fluid plus  $\delta f$  and the analytic solution distribution function, integrated over transverse velocity coordinates  $(u_y, u_z)$ , as a function of  $(x, u_x)$  for  $N_p = 10^3$ . Comparing the fluid plus  $\delta f$  distribution to the analytic solution shows that the hybrid solution is a reasonable estimate for the analytic solution. Meanwhile, the fluid distribution function, a Maxwell-Boltzmann, shows the development of two counter-propagating shock waves and a substantial accumulation of the mass into the region between the shocks.



**Figure 3-22. Distribution function for the fluid (left) and fluid +  $\delta f$  (center) and analytic result (right) for the hybrid solution of an initially sinusoidal velocity profile for  $N_p = 10^3$  particle generation count.**

### 3.1.4.6. Hybrid Summary

In this section we presented an overview of a coupled  $\delta f$  fluid / kinetic method, basic design considerations and highlighted a few numerical results. A basic component of this method is the weight evolution, or generation of  $\delta f$  particles. This characteristic is also common to  $\delta f$  methods for collisions [144], and an efficient particle merge method is critical to making it computationally efficient. The second characteristic that is apparent in the solutions presented here is a secular deviation in the fluid (mass, momentum, energy)-density from the moments of the full distribution function. Adding a flux correction process to minimize this drift will assist in maintaining  $\delta f/f < 1$  and reduce the impact of particle noise apparent in these results.

### 3.1.5. EMPIRE *Advancements Summary*

This Grand Challenge LDRD initiative coexisted with the ASC/ATDM program that desires to build a new multi-physics code from the ground up. While EMPIRE included several features matured through programmatic work, many of the features required for pulsed power problems were excluded from the original EMPIRE scope, and were added through this LDRD project. All the features added to EMPIRE under this LDRD effort are maintained and tested on a daily basis across a wide range of platforms.

Under the PIC component, transmission line coupling, thermal desorption and implicit methods were developed and tested on a range of verification problems and accuracy and traits were characterized. These were tested against full-scale power flow problems, and initial results are encouraging. Several improvements have been identified with the thermal desorption model, which are being addressed using programmatic resources in FY21. EMPIRE is currently being exercised as a design tool for power flow problems, and the capabilities developed in this project are being validated and extended.

A relativistic fluid capability was developed in order to solve the power flow problem with several matching boundary conditions. The solver was verified on a range of test problems, ranging from standard relativistic hydrodynamics benchmarks taken from the literature, electromagnetic wave problems that probe the coupling of the electromagnetic solver to the hydrodynamics algorithm,

space charge limited emission problems taken from the literature and new verification problems for relativistic warm diodes that were developed as part of this LDRD. Finally, the combination of relativistic fluid algorithms, electromagnetic solvers and boundary conditions were demonstrated on a relativistic planar MITL using a Z-relevant voltage drive; this final problem points the way to future development and verification efforts that are being addressed in FY21.

Finally, a  $\delta f$  fluid-particle hybrid approach was researched. This approach was never brought to the level of maturity to remain in the production EMPIRE code environment. However, this work will continue through additional LDRD investment, with the near-term goal to identify and address critical algorithmic deficiencies.

### **3.2. CHICAGO Advances to Support Z Pulsed Power Simulation**

Hybrid implicit particle-in-cell (PIC) algorithms permit the simulation of complex problems involving both kinetic and fluid plasma regimes over large spatial and temporal scales. The Magnetic Implicit algorithm, now able to under-resolve the plasma as well as the cyclotron frequency, is compatible with kinetic, fluid and hybrid techniques within the CHICAGO code. Fluid electrons can be computationally fast where and when fluid assumptions are valid. Additional flexibility is obtained if discrete PIC macroparticles, with velocities advanced by either fluid or kinetic equations, are permitted to dynamically migrate between the two descriptions based on phase space criteria. Ideally, these migrations result in energetic particles treated kinetically and dense thermal plasma particles as a fluid. With an energy-conserving particle advance, resolution of the plasma Debye length is not required for numerical accuracy or stability. For pulsed power applications, the simulation time step is usually constrained by the electron cyclotron frequency not the more restrictive plasma frequency. A new implicit technique permits accurate particle orbits even at highly under-resolved cyclotron frequencies. Thus, greater temporal and spatial scales can be accurately modeled relative to conventional PIC techniques.

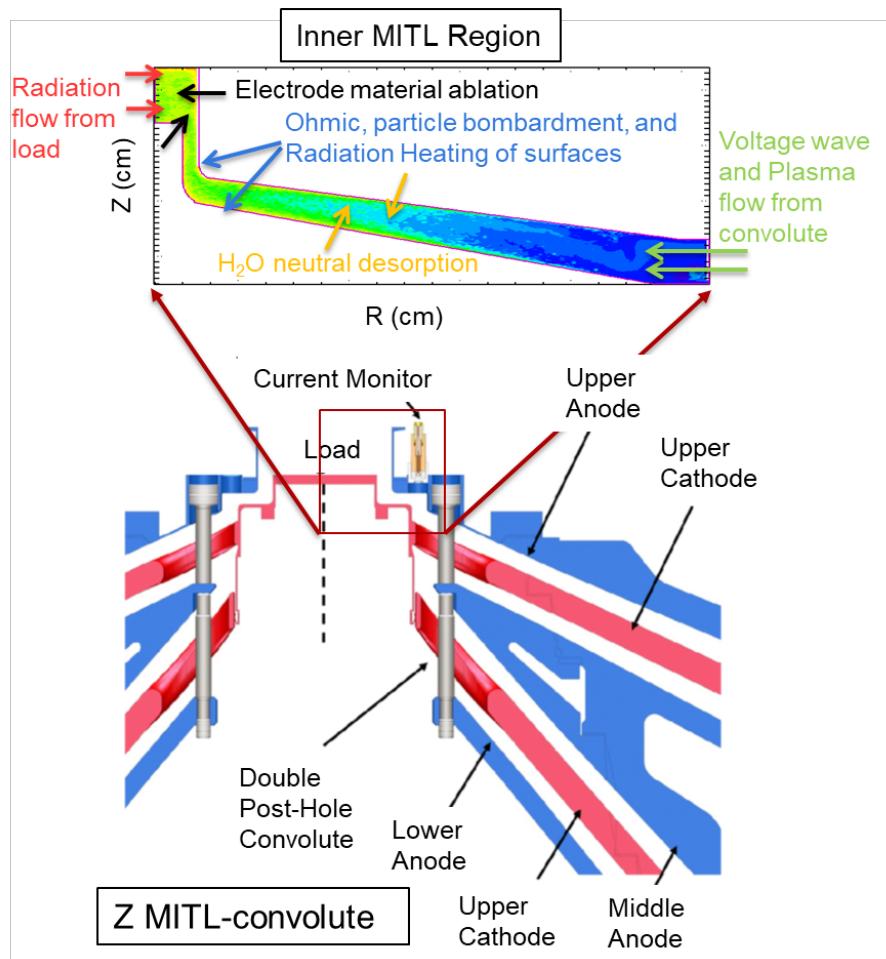
In this report, we describe the hybrid PIC technique and fully electromagnetic, hybrid simulations of plasma evolution and current loss in representative sections of pulsed power accelerators designed for driving a Z-pinch load. The dynamics of electrode heating, electron transport, and surface contaminant evolution are studied in a series of relativistic hybrid-implicit PIC simulations. These dynamics can lead to the diverting of current before reaching the Z-pinch load, thus degrading load performance. We describe the enabling algorithms and compare results from fully kinetic, multi-fluid, and hybrid kinetic-fluid simulations while discussing the computational performance of the three approaches. The key thrust of the work is to identify possible computational acceleration, through hybrid methods, applicable to pulsed power accelerator design.

### 3.2.1. CHICAGO *Simulations of Z Power Flow including Electrode Physics*

Understanding the dynamics of electrode heating, sheath flow, and contaminant plasma evolution is critical for predicting the performance of large pulsed power accelerators such as Z. These dynamics can lead to the loss or diverting of current before reaching the Z-pinch load, which results in degraded load performance. Previous work has focused on current diverted in the upstream magnetically insulated transmission lines (MITLs) and post-hole convolute regions of Z. In these regions, losses were found to scale with load impedance and the system vacuum achievable, with losses calculated to be as high as 1–2 MA. Downstream from the convolute region in Z, accurate current measurements are challenging to achieve, leading to a lack of understanding of the loss mechanisms in the small radius ( $< 3$  cm) MITL feeding the load.

In Reference [136], we presented the first ever 2D fully electromagnetic, fully-kinetic simulations of plasma evolution and current loss in the inner MITL region of Z. This region is defined by a radially converging MITL, which is a feature common to MA-scale Z-pinch accelerators. Shown schematically in Figure 3-23, the electrodes in this region are rapidly heated via mainly Ohmic or skin depth heating. Surface contaminants are liberated and plasmas quickly form as the temperatures exceed 700 K. Instabilities lead to a rapid plasma density fill of the inner MITL and subsequent current loss. The instability growth is likely due to the resistivity of the magnetized electrode plasma. Detailed CHICAGO simulations of the plasma evolution show 1–2 MA current loss in the inner MITL region after the plasma exceeds  $10^{15}$  cm $^{-3}$  density. The plasma density was found to rapidly fill the gap due to an instability in the electrode plasma and electron sheath.

The Magnetic Implicit (MI) algorithm has been improved to combine the implicit time-integration of the direct implicit (DI) technique [26, 140], for large plasma frequencies ( $\omega_p$ ) with the highly accurate orbit integration properties at large cyclotron frequency ( $\omega_c$ ) [40].



**Figure 3-23.** A cylindrical ( $r,z$ ) representation of the MITL-convolute system is shown. Power is delivered from 4 MITLs through a double-post hole convolute to the load. A detail view (top) of the radially-converging inner MITL region is used to illustrate the physical processes which occur as power is delivered to a Z-pinch. The voltage pulse and plasma generated upstream enter the inner MITL as shown on the right. The physical processes mapped by region are upstream charged particle flow, contaminant desorption, and ablation of material from the electrodes heated by ohmic dissipation, charged particles and radiation.

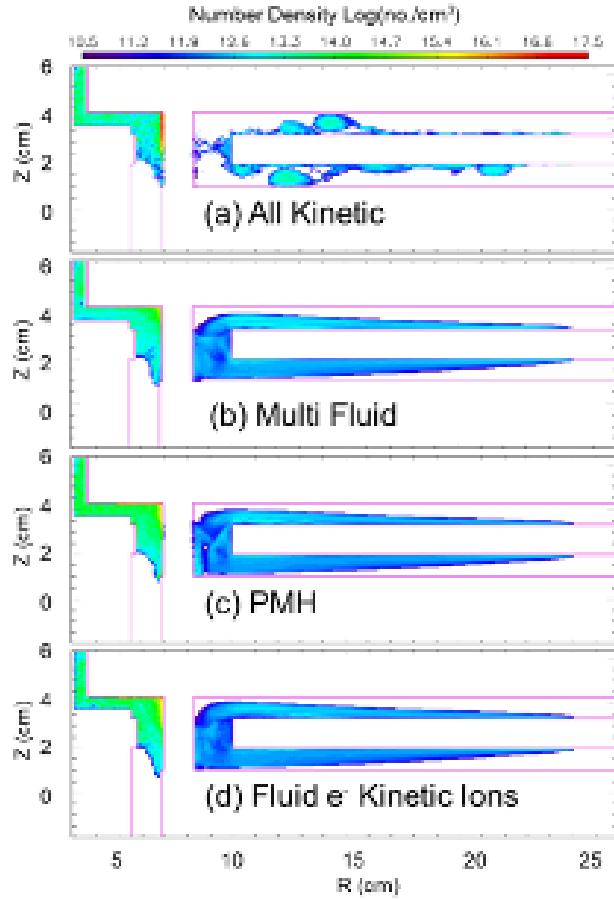
### 3.2.2. *A Fast Hybrid Particle-in-Cell Technique for the Simulation of Pulsed Power Accelerators*

In Reference [137], we present a hybrid-implicit PIC technique to more rapidly simulate the relativistic and thermal particle populations observed in high-power systems. Discussed in our recent paper [136], the new Magnetic Implicit method has demonstrated the efficacy of PIC method for modeling stressing high-density magnetized plasmas for kinetic equation of motion. Here, the goals are to develop an accurate evaluation of the plasma physics and a fast method for evaluating new designs. First discussed in Reference [135], the hybrid technique described here uses macroparticles advanced with both kinetic equations and fluid equations similar to that described by Brackbill and Roppel [14] and recently in Reference [73]. The sum of these representations constitutes the complete 6-dimensional distribution function. Fluid representations are better suited to describing the dense plasma regions near the electrode surfaces while kinetic representations are necessary to describe details of the lower-density sheath flow. It can be advantageous to include both descriptions and/or allow macroparticles to change their governing equations dynamically such that the computationally faster but valid description is used. We refer to this technique in which the original equations governing kinetic and fluid particles are retained as “particle migration hybrid” or PMH. This is distinguished from the  $\delta f$  technique which represents a perturbation from an equilibrium distribution.

We demonstrate the kinetic, fluid and hybrid techniques in two previously-published problems relevant to the 26-MA Z accelerator [129, 78, 110, 105]. In idealized three-dimensional and highly resolved two-dimensional simulations, the currents from three pulsed power transmission lines, operated as MITLs, are added in a convolute to drive a Z-pinch load. Results from the fluid and hybrid techniques are compared to the fully kinetic treatment. While the electron sheath dynamics are only faithfully captured in the kinetic treatment, the current loss during transport is similar in each treatment, indicating the multi-fluid and PMH approaches have utility in these types of power flow simulations.

We completed four simulations of the geometry in Figure 3-24 with representative Z voltage pulses using: (i) fully kinetic, (ii) multi-fluid, (iii) all PMH hybrid, i.e. fluid and kinetic treatments for both electrons and ions, and (iv) fluid electrons with PMH ions. The simulations include water desorption with electron, hydrogen ions and oxygen ions (assumed in +1 state only). The transition criterion for migration from fluid to kinetic for these three charged particles is  $KE > 50$  keV. Once particles transition to kinetic, few will thermalize before being lost to the walls.

All simulations show some small loss (10's of kA) to the upstream post as described in References [103] and [6], where the magnetic null is found (a channel of very small magnetic field due to the symmetry breaking of the post current). The current transported to the load and the difference between the current at the entrance to the MITL and the load current (or loss current) are plotted in Figure 3-25. Until approximately 65 ns, the four simulations calculate the same load and loss currents to within 2% and 5%, respectively. All calculate the same time of initial current loss peak at 53 ns where the voltage reaches 1.5 MV and electron flow current is highest. This initial loss is mainly due to launched electron flow in the outer MITLs that is mostly lost in the convolute. The numerical agreement largely demonstrates that all numerical techniques are sufficiently accurate for modeling the flow current.



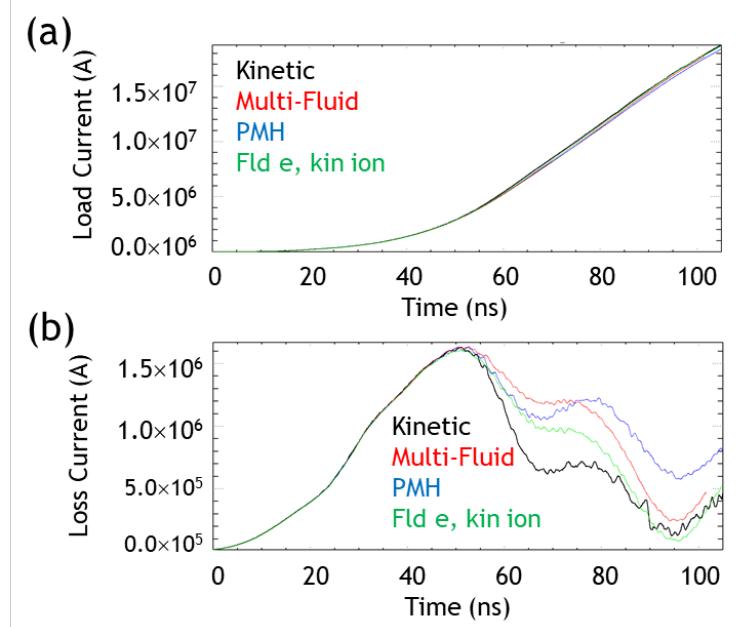
**Figure 3-24. Electron density 90 ns into the current pulse, located in the  $\theta = 0$  plane (center of post), for simulations consisting of (a) all kinetic particles, (b) all multi-fluid particles, (c) all PMH with fluid particles transitioning to kinetic above 50 keV, and (d) fluid electrons with PMH ions; all simulations have a dynamic timestep such that  $\omega_c \Delta t < 9$**

Because the downstream load is a short circuit in this example, the voltages plateau and the current continues to rise after 53 ns causing the electron flow current to fall. All simulations exhibit a smaller second peak in current loss due mainly to dense plasma motion just downstream of the post for  $r < 7$  cm, although the multi-fluid and PMH simulations' peak is larger due to small differences in the motion of the thermal plasma from the electrode. The reason for the difference is likely the larger dispersion seen in the Eulerian fluid description.

Note, the better agreement in current loss between the all kinetic and fluid electron/PMH ion results shown in Figure 3-25. This behavior suggests that the kinetic ion description is preferable to fluid later in time when the losses are small. While the observed differences in calculated densities between the techniques are interesting, the overall performance of each technique in simulating current transport to drive the load is similar indicating the fluid and PMH approaches have utility in these types of power flow simulations.

The potential for computational speed up with the fluid and hybrid approaches make them

desirable for design level calculations with speed up factors as high as  $6\times$  (relative to the fully-kinetic treatment, depending on the total particle number) observed so far. If one is strict with kinetic particle number, as was the case in the simulations presented here, the multi-fluid speed up was a more modest 20%. These simulations required roughly 20,000 processor-hours on a standard Linux cluster.

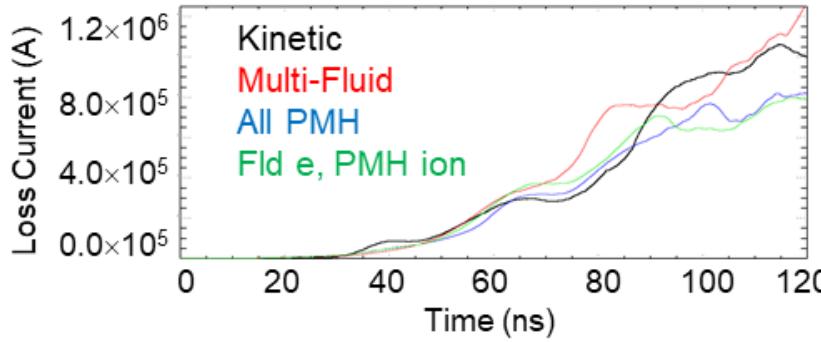


**Figure 3-25.** The (a) load current and (b) loss current from four simulations (Kinetic, Multi-Fluid, PMH, and fluid e with PMH ions) using the geometry in Figure 3-23.

With 2D simulations of the Power Flow 18a experiment, we examine the inner MITL current transport using the four techniques used to simulation the convolute: kinetic, multi-fluid, PMH, and fluid electrons with ion PMH. The inner MITL, modeled as a circuit in the previous section, connects the convolute and load. To model the plasma flow from the convolute at 10 cm radius into the simulation, we inject a linearly rising current of electrons and protons to 500 kA at 110 ns, constant thereafter. In the inner MITL due to higher plasma densities evolving from the more rapidly heating electrodes, simulations must resolve the smaller skin depths as well as the larger cyclotron frequencies found near the load. We find that typically these simulations require 50  $\mu$ m spatial resolution for convergence, but now can be run in only 2 dimensions.

The Z current drive at 10-cm radius rises to 22.6 MA by 120 ns, and the current lost is plotted in Figure 3-26. In each simulation, this loss current rises to roughly 1 MA by 120. The multi-fluid simulation is highest and the two hybrid simulations' losses rise close to the kinetic curve until dropping 100 kA below kinetic after 90 ns. Assuming the kinetic simulation is the most accurate, it's difficult to draw many conclusions about the physics efficacy, although the PMH hybrid techniques give a similar result compared to the kinetic treatment for a longer period.

The relative computation speed of the multi-fluid simulation does have a considerable advantage. The kinetic simulation requires 38,640 processor-hours to run to 120 ns. The fluid simulation was



**Figure 3-26.** The calculated current loss for four inner MITL simulations (Kinetic, Multi-Fluid, PMH, and fluid electrons with PMH ions). The drive current rises to 22.6 MA by 120 ns.

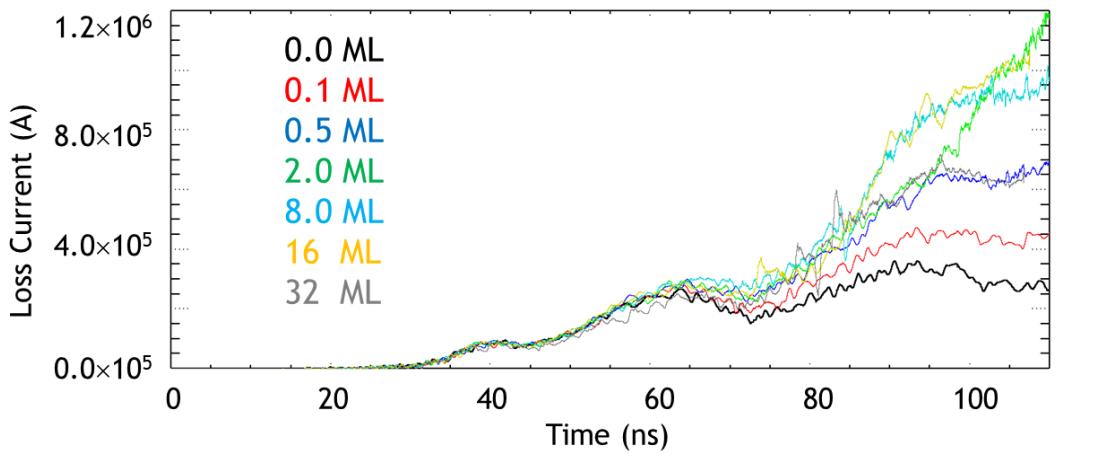
only 11,760 proc-hrs, a  $3.3\times$  speedup. The two hybrid simulations fell in-between at 29,570 (all PMH) and 14,000 proc-hrs. These results combined with the convolute study suggest the multi-fluid technique is a faster, reasonably accurate, alternative to fully-kinetic approaches. Additional aspects of simulation speed-up are discussed in a later section of this report (§3.2.5).

### 3.2.3. ***Sensitivity of Current Loss to Water Inventory***

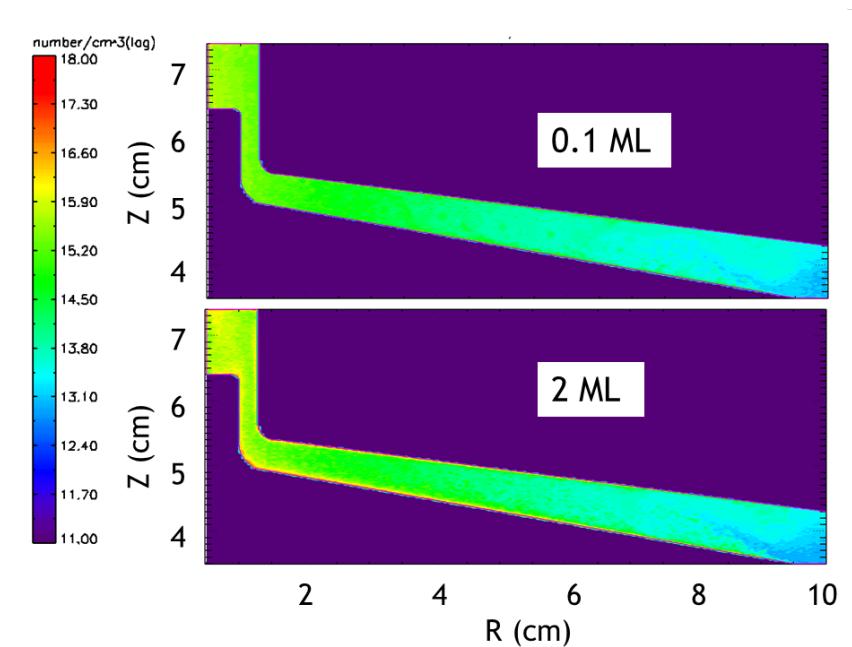
A series of power flow experiments on the Z accelerator were performed to study losses in the system. Recent simulations of the inner MITL in these experiments have shown the effects of strong magnetization can rapidly drive plasma into the gap which leads to a Hall-like current loss mechanism [6]. We repeat the 2D kinetic simulations of these experiments while varying the water molecule inventory on all the surfaces from 0 to 32 monolayers. Here, a monolayer (ML) is defined as  $10^{15}$  molecules/cm<sup>2</sup>. The 8-ML result was described in detail in Reference [136]. The '0 ML' simulation only permitted proton (above 400 C) and electron (above 240 kV/cm electric field) space-charge-limited emission from the surfaces and no water desorption. The inner MITL connects to the post-hole convolute upstream. To model the plasma flow from the convolute at 10 cm radius into the simulation domain of the inner MITL, we inject a linearly rising current of electrons and protons to 500 kA at 110 ns, constant thereafter. We find that typically these simulations require 50  $\mu$ m spatial resolution for convergence. Because of the higher plasma densities at the surface with water inventory, we increased resolution for  $\geq 16$  ML simulations to 25  $\mu$ m. As a convergence check, we compared the results from the two resolutions for the 16 ML case, and the calculated current loss in the simulations agrees to within 5%.

A summary of the current lost between the outer radial boundary of the simulation domain and the entrance to the load is plotted in Figure 3-27. We observe that the losses increase with inventory peaking at 2 ML, then remain nearly constant before falling at 32 ML. The 0.5 ML and 32 ML simulation show a very similar loss curve. The electron density at 110 ns is plotted for the 0.1- and 2 ML simulations in Figure 3-28. The 0.1 ML simulation electron density approaches  $10^{15}$  cm<sup>-3</sup> density by 110 ns with little density remaining on the electrodes. By contrast, the 2-ML simulation has a density in the load region center approaching  $10^{16}$  cm<sup>-3</sup> density with significant density still on the electrode surfaces. The initial expansion of plasma from the surface into the gap has been shown in Reference [136] to be the result of the sheath-resistive plasma instability. The surface plasma from which the electron sheath forms has a perpendicular conductivity reduced by roughly the square of the Hall parameter. As the water inventory of the surface is increased, the increasing plasma density results in a higher collisional rate, lower Hall parameter, and decreasing growth rate of the instability. The resulting reduced density in the gap of the MITL decreases the current loss. This behavior is consistent with the plateau and slight decline in current loss at large inventory.

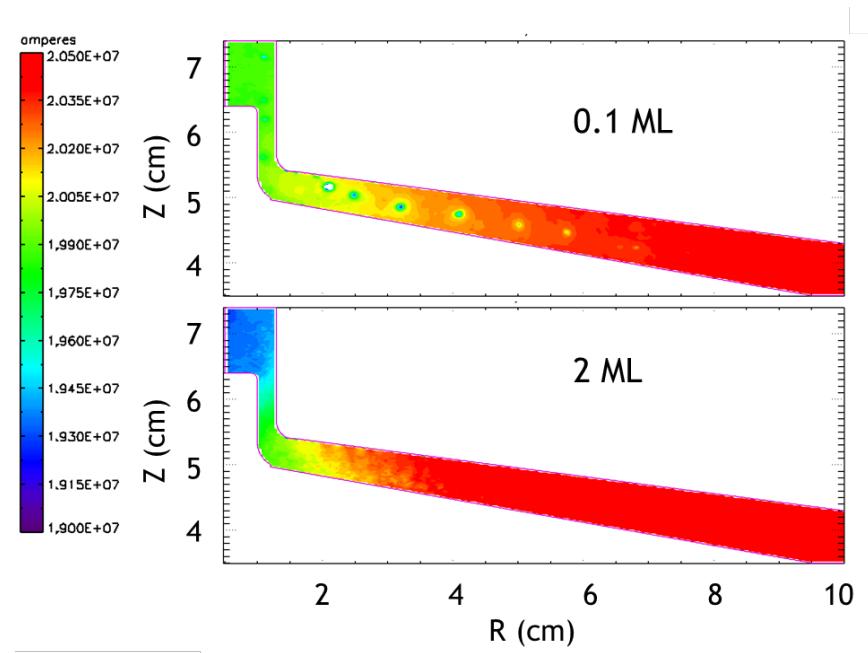
In Figure 3-29, the enclosed current at each position ( $rB_\theta$  product) is plotted at 110 ns for the 0.1- and 2-ML simulations. In the 0.1-ML simulation, most of the current lost is between 1.5 and 5 cm radius in the radial portion of the inner MITL. The small 250-500 kA depressions in enclosed current are the result of vortices lifting off the cathode and eventually striking the anode. These vortices are the result of the changing MITL impedance over this region. It appears that a significant amount of current loss is carried by these features at small inventories. In the 2-ML simulation, the vortices are absent and the current loss is found for  $r < 3$  cm in the regions of the higher plasma densities. The current loss is due to a Hall-like mechanism first addressed in Reference [6] and discussed in detail in Reference [7].



**Figure 3-27. Simulated current loss between the 10-cm radius boundary and just upstream of the load, for several initial water inventory conditions on the electrodes.**



**Figure 3-28. Simulated electron number density at 110ns for the 0.1- and 2-ML initial water inventory conditions on the electrodes.**



**Figure 3-29.** Simulated radially enclosed current (radius-magnetic field product) at 110ns for the 0.1- and 2-ML initial water inventory conditions on the electrodes.

### 3.2.4. Spherically Expanding Plasma Verification

A key process in fluid and kinetic theory is the conversion of thermal energy into kinetic energy. In this section, we discuss this process in both neutral gases and plasma. Hughes and Musson [51] derived a 1D spherical steady state expression of the conversion of thermal energy into “fluid” motion for warm electron (temperature  $T_e$ ) and cold ion (mass  $m_i$ ) populations with an ambipolar field. The dimensionless velocity, position and density are given by

$$\begin{aligned}\hat{v}(r) &= \frac{v(r)}{v(1)} = \frac{v(r)}{c_s}, \\ \hat{r} &= \frac{r}{r(1)} = \sqrt{\frac{1}{\hat{v}}} \exp \frac{\hat{v}^2 - 1}{4}, \\ \hat{n} &= \frac{n}{n(1)} = \frac{\hat{n}(1)}{\hat{r}^2 \hat{v}}\end{aligned}\quad (3.143)$$

where  $c_s = \sqrt{kT_e/m_i}$  is the sound speed. The equation for  $\hat{v}$  is transcendental and must be iterated for solution. The problem is initiated at small radius (1 cm) with  $T_e$  and  $v(1)$  and is followed for 10  $\mu$ s as the particles expand to 10 cm as shown in Figure 3-30. We test the CHICAGO kinetic and fluid algorithms. The ions and neutrals are assumed to have hydrogen mass. In all these cases, we include a simple Spitzer collision operator for kinetic and fluid species. For neutral particles, we assume a hard sphere collision operator for neutral on neutral.

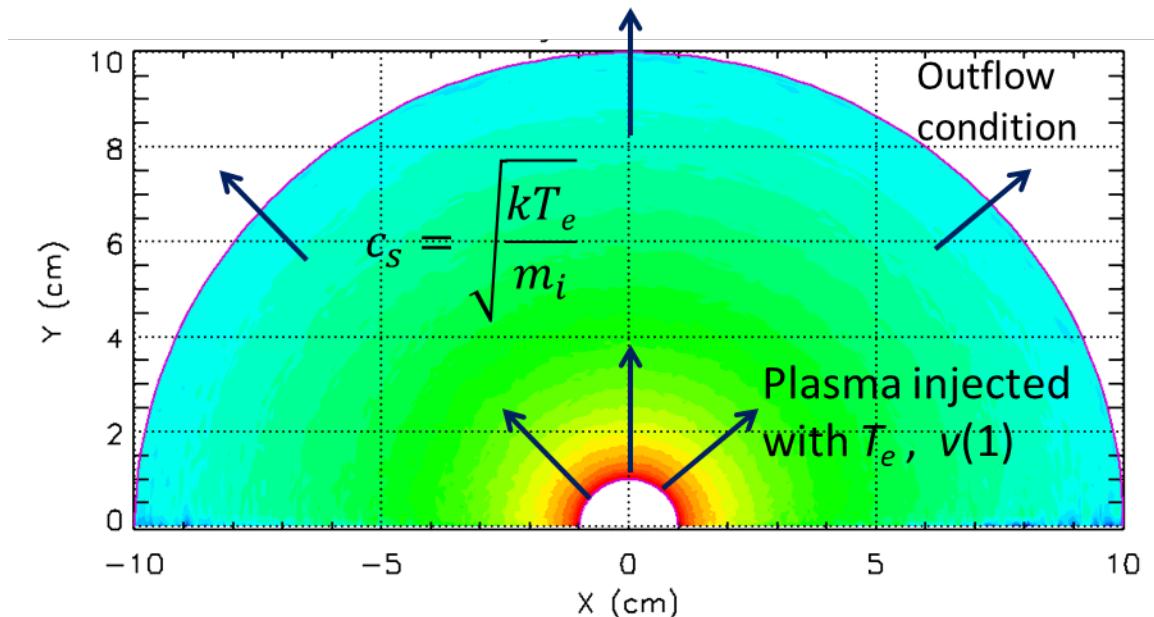
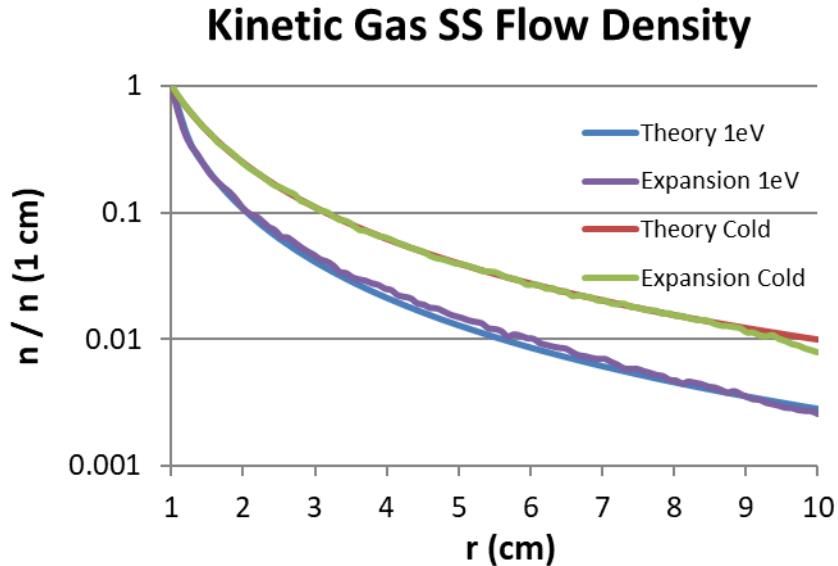


Figure 3-30. The spherical geometry for the steady state theory and simulations. The electron number density is plotted for the 0.1- and 2-ML water inventory kinetic simulations.

### 3.2.4.1. Neutral Gas Expansion

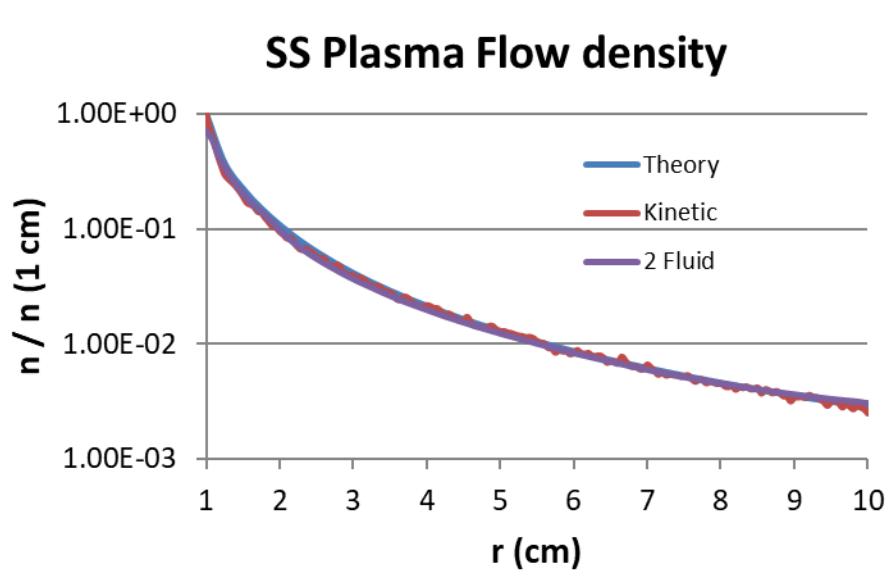
As a first test, we inject neutral gas cold with  $v = c_s$  (1 eV). In this case, the steady state solution simplifies to  $\hat{n} = \hat{n}(1)/r^2$ . In a second test, the neutrals are injected with no net velocity and 1 eV temperature. The 1 eV gas expands to roughly  $v = c_s$  (3 eV) and the density falls more rapidly as seen in Figure 3-31. The Hughes and Musson theory and CHICAGO simulation results are shown for the two example cases and agree quite well.



**Figure 3-31. Steady state gas density with initial conditions of 1 eV and cold with  $v(1 \text{ cm}) = c_s$  (1eV), shown for both CHICAGO simulation and analytical theory.**

### 3.2.4.2. Plasma Expansion with Kinetic and Two-Fluid Simulations

In this section, we include charged particles and an electrostatic field solution to model a  $2.0 \times 10^{11} \text{ cm}^{-3}$  e- H+ plasma with a  $v(1 \text{ cm}) = c_s$  (1 eV) =  $9.82 \times 10^5 \text{ cm/s}$  velocity. We use CHICAGO's energy-conserving implicit technique with fully kinetic and two-fluid (Eulerian) particle treatments in 2D  $r - \theta$  spherical geometry. We use 100 cells in  $\theta$  over the entire  $\pi$  range. In radius, we use 180 cells or 0.05-cm uniform cell size. Note that the plasma collisionless skin depth is 1.2 cm and the Debye length is 0.0017 cm - under resolved by 20 in the simulation. Thus, this is an excellent test of the energy conserving technique. The time step (0.018 ns) adequately resolves the plasma frequency ( $\omega_p \Delta t = 0.42$ ) so the implicit technique is not actually needed. We inject 1 particle/step per transverse cell yielding 5M macroparticles for electrons and ions. In Figure 3-32, we plot the Hughes and Musson theory and CHICAGO simulation results after 10  $\mu\text{s}$  using both the fully kinetic and two-fluid treatments; excellent agreement is observed.



**Figure 3-32. Steady state plasma densities for an initial condition of  $T_e = 1\text{eV}$ ,  $v = c_s$  (1eV), for both CHICAGO kinetic and two-fluid treatments, as well as analytical theory.**

### 3.2.4.3. CHICAGO Verification Summary

The simulation of both neutral gas and plasma has been verified against the analytic theory of Hughes and Musson at steady state. The problem tests the accuracy of the conversion of thermal energy into kinetic or fluid motion in both kinetic and two fluid treatments. The agreement of CHICAGO with the theory is excellent despite under resolving the Debye length by a factor of 20, which stresses the energy conservation of traditional PIC techniques. We have also completed plasma simulations with the new Magnetic Implicit algorithm using electromagnetic field solution and observed equally good agreement.

### **3.2.5. CHICAGO *Parallelization Strategies – Trilinos/Kokkos Integration***

As part of this project, new parallelization strategies have been evaluated and developed within the CHICAGO particle-in-cell code to reduce simulation execution times. Through Trilinos and Kokkos, CHICAGO targets hybrid parallelism (distributed and shared memory) across HPC platforms with several backend programming models. Trilinos is a software framework comprised of packages that offer parallel solver algorithms for multi-physics applications. Kokkos facilitates the underlying parallel operations in supported Trilinos packages and may also be used independently to write explicit parallel sections. Kokkos is a C++ programming model that provides abstractions for data management and parallel execution, enabling developers to write portable code that targets new and emerging HPC systems. Trilinos and Kokkos support an MPI+X execution model with Pthreads, OpenMP and CUDA backend support for shared memory. Trilinos has been integrated within CHICAGO field solutions and Kokkos parallel sections established within particle push procedures.

#### **3.2.5.1. Trilinos Field Solve Integration within CHICAGO**

CHICAGO's existing spatial decomposition scheme divides the simulation into regions and domains, assigning MPI ranks to each division. Additional shared memory parallelism is then utilized within each domain to further divide the work. The PETSc numerical solution library is currently used to solve sparse linear systems in CHICAGO. The current and planned programming model for PETSc is limited to MPI. Our motivation in adopting Trilinos is that it enables CHICAGO to further exploit intra-domain parallelism and effectively reduce field solve execution times. The Trilinos field solve integration primarily incorporates the following Kokkos-enabled packages: Tpetra, Ifpack2, and Belos. After establishing an interface from CHICAGO C code to a separate C++ compilation unit, the Trilinos linear system was created with Tpetra objects. Tpetra Maps define the parallel distribution of matrix rows and columns over processes. There are fundamental differences between PETSc and Trilinos row/column mapping and matrix insertion methods. This impacts the distribution of locally and globally owned indices of the matrix over MPI ranks. The matrix-fill procedures were adjusted accordingly to construct the Tpetra CrsMatrix and MultiVectors for the linear problem. Legacy arrays were replaced with Kokkos View objects whenever applicable to interface with Trilinos object constructors. Ifpack2 preconditioners must then be constructed, initialized, computed and applied to the linear system. Similarly, Belos objects and solver algorithms are selected and parameters specified. Support for run-time parameter passing of preconditioner and solver parameters was added by utilizing Trilinos existing support for XML file input. The sparse linear system is solved with Belos and the solution is moved back into CHICAGO legacy data structures. After which control is returned to the C compilation unit to complete each time step. Kokkos facilitates hybrid parallelism for all Trilinos data structures. The host and device parallel backends are determined during the configuration of the Trilinos library. After compiling CHICAGO against Trilinos, one may specify details of the parallel run through command line parameters and environment variables. This allows users to specify the number of threads or number of GPU devices. Using this info, Trilinos can select appropriate defaults for parallel execution and memory management.

### 3.2.5.2. Trilinos Verification and PETSc Comparison

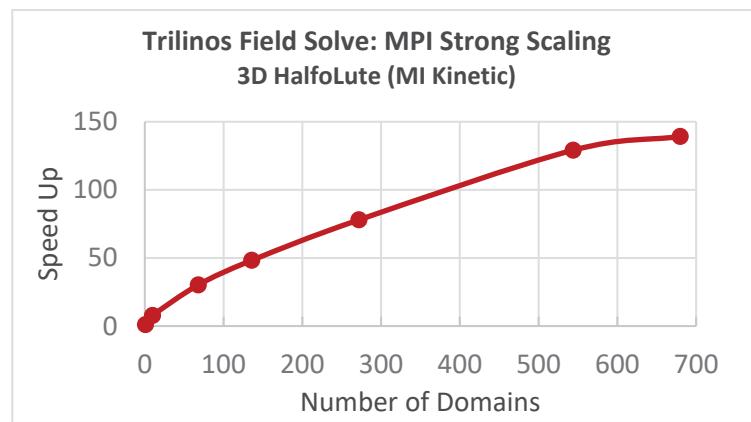
At this time Trilinos has been integrated into CHICAGO to support both the electromagnetic and static field solutions. Computational accuracy was verified through an example 2D inner MITL simulation, involving 5.8 million cells and varying number of domains. This verification was completed first as a cold test (no particles) to eliminate stochasticity. MPI-only execution, as well as OpenMP and CUDA backends were tested. Ifpack2 Jacobi and Additive Schwarz preconditioners were evaluated along with Belos GMRES and stabilized Biconjugate Gradient (BiCGStab) solver algorithms. Identical field calculations results were found between PETSc and Trilinos implementations over a 10ns ( 140,000 timestep) interval. A preliminary evaluation of distributed memory scaling for this same 2D inner MITL simulation was completed and performance compared directly to that of the PETSc field solution. Through this analysis it was found that the additional overhead of Trilinos results in slower execution times overall. However, with more MPI ranks involved the difference in performance is small. At 500 processors PETSc was found to execute approximately 15% faster for this simulation.

### 3.2.5.3. Trilinos Performance Analysis

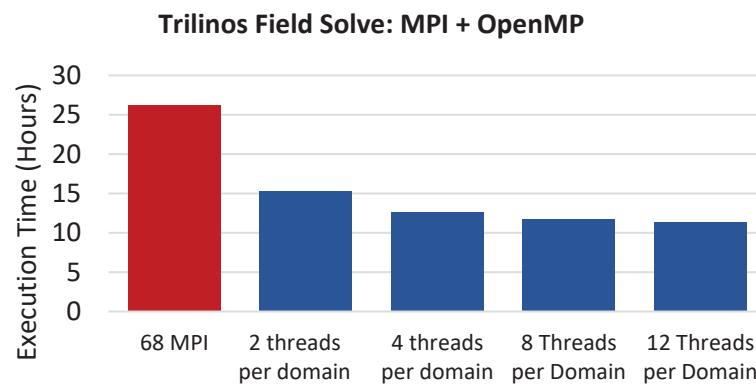
After further refinement of the Trilinos implementation, a more in-depth analysis of both distributed and shared memory parallelism was completed using the 3D half-o-lute simulation. This simulation involves roughly 5.3 million cells. MPI strong scaling was evaluated from 1 to 680 processors for a cold test, to evaluate the performance improvements coming from the field solution alone. MPI parallelization in correspondence with CHICAGO's domain decomposition found 140x speedup at 680 processors. Trilinos strong scaling results for this test case are presented in Figure 3-33.

OpenMP performance was investigated for this identical simulation with a fixed 68 domains and varying number of OpenMP threads assigned to each MPI rank. By merely increasing shared memory parallelism within each domain, it was found to improve performance 2.3x for this 3D half-o-lute simulation. In Figure 3-34, we see the impact on execution time when adding OpenMP threads to further divide the work within each domain, without dividing the simulation space further.

During this LDRD project, CHICAGO was built for Sandia's Vortex compute cluster, which uses a similar architecture to the leadership-class Sierra machine. Vortex is an IBM Power 9 machine with 4 Nvidia Volta V100 GPUs per node. With access to Vortex we have begun larger scale GPU simulations. CHICAGO is currently running on Vortex GPUs with Trilinos/Kokkos support. A preliminary evaluation of GPU field solve performance was conducted using a 48-domain test of the Power Flow 18a simulation, in which the addition of 1 GPU per domain reduced execution time by approximately 25%. Further evaluation is ongoing for Trilinos/Kokkos CUDA execution. The ability to configure CHICAGO for use on a new HPC architecture, while still having the flexibility to efficiently target parallel backends of interest, is a promising demonstration of Kokkos portability.



**Figure 3-33. Trilinos Field Solve MPI Scaling.**



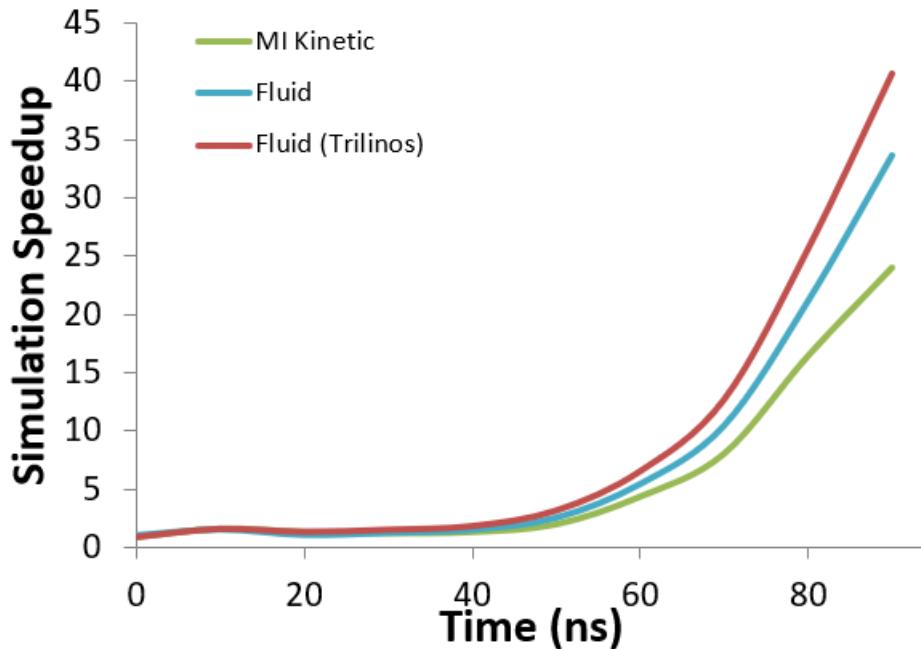
**Figure 3-34. Trilinos Field Solve MPI + Open MP Scaling.**

### 3.2.5.4. Kokkos Particle Advance Integration

Within CHICAGO's MPI domain-decomposition, explicit OpenMP parallel sections were previously established in performance-critical areas. Preliminary CUDA kernels were also developed to run with GPU acceleration in the early stages of this LDRD project. These GPU kernels demonstrated improvements in performance with dense particle populations for smaller-scale idealistic simulations. The OpenMP and CUDA implementations exist as two separate codes written in the context of their respective programming models. Through Kokkos we aim to establish a single code construct that can target all parallel backends of interest and simplify future development. Kokkos allows us to do this by managing data allocations, access patterns and parallelism depending on the execution space. The Kokkos programming model is being incrementally adopted into particle advance functions of CHICAGO. Kokkos interoperability with parallel backends allows for existing OpenMP calls to be compiled and run alongside Kokkos. One can replace OpenMP constructs in CHICAGO with Kokkos parallel sections that run in the host execution space. However, for performance portable code across host and device spaces, the integration of Kokkos is more invasive. Data must be managed through Kokkos in order to be accessible from multiple memory spaces. Numerous legacy C data structures must be refactored into Kokkos-friendly C++ objects. As data structures are modified the access methods of data must be adjusted in nested function calls. These external functions must also be made accessible to the device memory space. Once in place, Kokkos can simplify the management of memory between execution spaces and abstract away the complexity of CUDA kernels from the user. However, existing C code must be refactored substantially to support device execution through Kokkos. At this time several particle advance procedures have been translated from OpenMP to Kokkos parallel constructs through which host execution is supported. Refactoring of data structures and subroutines is underway to facilitate the particle advance on GPUs. With the 3D half-o-lute simulation we have taken a first glance at Trilinos/Kokkos performance improvements using hybrid parallelism (MPI + OpenMP) in both field solve and particle advance components. With an identical number of domains and increasing only the shared memory parallelism, we were able to reach a peak speedup of 40x for the all-fluid 3D half-o-lute simulation. Thereby reducing the overall execution time from 314 hours (single domain), down to 31.65 hours. Figure 3-35 illustrates the speed-up over time for these simulations, run on the Saturn system at Voss Scientific. Table 3-4 summarizes the performance results for the 3D half-o-lute for each algorithmic and software optimization explored.

Run	Total Time (hr)	Average Speedup	Peak Speedup
DI 1 thread	314	1	1
MI 2 thread	42.37	7.4	24
All Fluid 1 thread	37.49	8.4	33.7
All Fluid - Trilinos	31.65	9.9	40.7
Full Hybrid 2 threads	43.65	7.2	24

**Table 3-4. Summary of half-o-lute computational performance.**



**Figure 3-35. 3D half-o-lute Speedup with Time.**

### 3.2.5.5. Conclusion

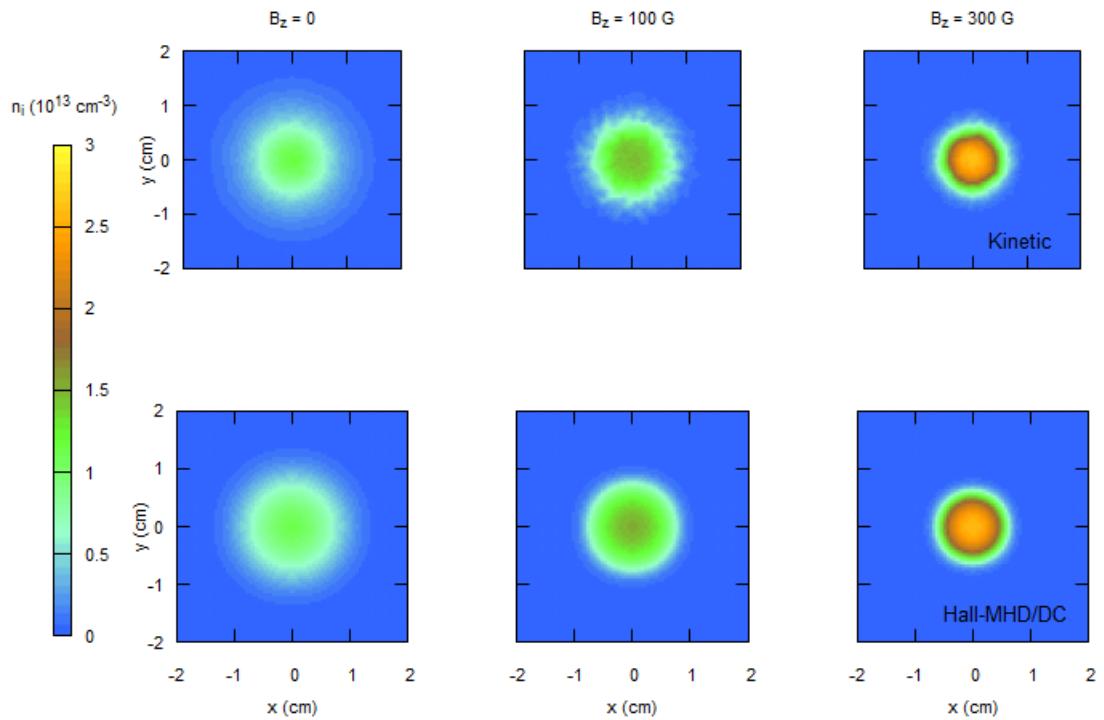
Software development completed under this Grand Challenge LDRD brought new capabilities and advances to multiple areas of CHICAGO. CHICAGO's field solution now supports hybrid parallel execution through Trilinos with OpenMP and CUDA. Both field and particle procedures have adopted Kokkos for performance portability. GPU support for the Kokkos particle advance is in development and will be further evaluated. Meanwhile, additional performance-critical sections in CHICAGO have been considered for incremental conversion to the Kokkos programming model. With a combination of techniques, we have been successful in expanding the resource utilization for a typical power flow simulation from several hundred processors (pre-LDRD) to several thousand processors (post-LDRD) with favorable performance scaling.

### 3.2.6. *Implicit Highly-Coupled Single-Ion Hall-MHD Formulation for Hybrid*

In the standard magnetohydrodynamic (MHD) formulation, the Hall term in the generalized Ohm's law (GOL) can be omitted when the Hall parameter  $H = \omega_{ce}/v_e \ll 1$ , where  $\omega_{ce}$  is the electron cyclotron frequency, and  $v_e$  is the electron-ion collision frequency. This term must be retained in regimes where the Hall parameter is non-negligible. Hall physics has been found to play a role in such phenomena as magnetic field transport in plasma opening switches, and in magnetic recombination processes. Hall physics has also recently been used to explain seemingly anomalous behavior in the magneto-Rayleigh- Taylor instabilities seen in the MagLIF experiment at Sandia National Laboratories. As described in Reference [131], the Hall-MHD model is obtained from the full one-fluid center-of-mass plasma model by assuming quasi-neutrality and by neglecting both electron inertia and displacement current (DC) in Ampere's law. For systems with a large Hall parameter, the Hall term must be retained in the GOL. This description leads to the presence of a whistler wave with a dispersion relation for which  $w$  is proportional to  $k^2$  at large values of  $k$ . This whistler wave behavior is problematic for explicit finite-difference Hall-MHD implementations as the wave speed grows unbounded at short wavelengths, requiring a stringent timestep constraint due to the Courant-Friedrichs-Lowy (CFL) condition:  $\Delta t \leq \Delta x/v$ , where  $v$  is the maximum velocity scale in the problem. This difficulty can be overcome by different techniques. In our work we have implemented a highly-coupled implicit particle-based Hall-MHD algorithm into the hybrid PIC code CHICAGO. The Hall-MHD algorithm has also been extended to allow for multiple ion species, but we restrict our description at this time to the case of a single ion species.

In this new PIC-based Hall-MHD model, DC can either be neglected or retained. The retention of DC also eliminates the quadratic  $k$  dependence of the whistler dispersion relation, but introduces a quasi- electromagnetic (EM) frequency branch. This leads to the introduction of a new frequency scale which must be resolved numerically at high Hall parameters. If DC is neglected, the highly-coupled implicit algorithm described below appears to be uniformly stable despite the quadratic dispersion behavior of the whistler mode. The algorithm consists primarily of a "highly-coupled" grid-based temporal advance, which is used to push ion momentum and electric and magnetic fields. The term "highly-coupled" is here used to denote the use of matrix methods to solve the approximate system of coupled ordinary differential equations which result when both coupling coefficients and terms involving spatial derivatives are treated as constant during a short time interval, i.e., a timestep. This approach allows for the relaxation of the system to the appropriate steady-state result (consistent with the assumption of constant coupling terms and sources) when relevant timescales are under-resolved. As noted above, when DC is retained there are still timestep constraints which occur at large Hall parameter due to magnetic effects. If DC is neglected these constraints are lifted. The highly-coupled matrix solutions obtained in this fashion are then utilized in a leapfrog scheme in which electric and magnetic fields are advanced at integral timestep intervals and particles at half-integral timestep intervals. As a demonstration of the highly-coupled Hall-MHD/DC algorithms we consider the 2D expansion of a plasma column perpendicular to a guide magnetic field. The initial plasma density profile of a fully-ionized hydrogen plasma has a peak density of  $3.1 \times 10^{13} \text{ cm}^{-3}$  in axisymmetric column with an initial width (FWHM) of approximately 1 cm. Initially both electrons and ions are initialized with a uniform temperature of 1 eV. The simulation is conducted in the x-y plane, with the

applied field (if present) in the  $z$  direction. The simulation is run for 400 ns with guide fields of 0, 100, and 300 G, a time duration over which the FWHM of the unmagnetized plasma column approximately doubles. For comparison we have also repeated these simulations using a fully-explicit kinetic PIC treatment. All simulations are performed with a uniform grid for which  $\Delta x = \Delta y = 0.05$  cm. The 2D ion profiles at  $t = 400$  ns are shown in the figure. Kinetic simulations, shown in the top row, were performed with a timestep of  $c\Delta t = 0.02$  cm, which is approximately as large as is allowable due to standard electromagnetic PIC constraints. In the bottom are the results from simulations using the Hall-MHD/DC algorithm with a timestep of  $c\Delta t = 50$  cm (a factor of 2500 greater than for the kinetic simulation). Comparing the kinetic explicit results with Hall-MHD/DC in Figure 3-36, there is generally good agreement in the gross features of the plasma expansion, but there is a visible flute-like drift instability in the magnetized kinetic simulations, which is especially visible in the 100 G case. Such drift instabilities which occur at density gradients are generally driven by charge separation between electrons and ions, are therefore absent in quasi-neutral formulations.



**Figure 3-36. The plasma density after 400 ns is plotted for  $B_z = 0$ , 100 G and 300 G using kinetic particles (top) and the Hall-MHD/DC particles (bottom).**

### **3.2.7. *Miscellaneous Implicit and Adaptive Particle Management Advances***

The initial stages of Z-pinch liner compression involve the rapid acceleration of electrons on the liner surface and inward penetration of the magnetic field. Because the liner is a solid material, standard PIC field solvers are ill equipped to handle these dynamics. As such, a new particle push was devised which is capable of modelling scenarios where  $vdt \gg 1$  without employing any of the assumptions utilized in MHD. Instead of separately advancing particles due to fields and then due to collisions, this new algorithm simultaneously evaluates both by applying damping matrices to the particle motion in the particle push, and then incorporating those matrices into the field susceptibility for consistency. While this algorithm is currently only capable of modeling stationary blocks of material, it can be extended to mobile material as well. This algorithm could potentially also be used to explore areas of physics such as warm dense matter.

Magnetically insulated transmission lines (MITLs) were studied with the purpose of determining the effect of a complex reaction chain on the overall circuital behavior of the MITL, as well as analyzing the dynamics of the multispecies system. The chain studied involved standard ionization, electron-collision induced fragmentation, and dissociative ionization and included the species  $H_2O$ ,  $H$ ,  $H_2$ ,  $O$ ,  $OH$ , and their ions. The desorption of  $H_2O$  was dictated through the Langmuir adsorption model using the Temkin Isotherm, and the desorption of  $H$  was calculated through a simple electric field threshold model with threshold 200 kV/cm. Ionization and electron induced fragmentation were both calculated kinetically, using previously established cross sections. One of the primary points of investigation was to determine whether the assumption that all ionization and reactions occurred near the cathode [135] was accurate. It was determined that by enforcing the ion population to be entirely generated on the cathode, the total charge contribution from  $H^+$  and  $O^+$  was substantially higher, by a factor of about 27x for the former and a factor of 3x for the latter. However, up to the voltages analyzed (300 kV), there was no notable difference to the voltage transmission between the two simulations.

The work on MITLs was assisted by the development of a modification to the current particle management scheme utilized by CHICAGO. The rapid and potentially unpredictable particle generation present in some simulations through ionization and reaction modeling can result in steep density profiles that are difficult to recreate with a macroparticle distribution with low charge weight variance. This modification aids in this by laying down dynamically sized meshes of particles that offer an effectively higher resolution with which to place particles. This technique was tested in multiple scenarios, including the MITL simulation described above and electrical breakdown in an AK gap. Due to a higher success rate, the modified algorithm is capable of supporting a higher ionization sampling rate while maintaining a reasonable number of particles per cell. This results in improved modeling of higher energy electrons most responsible for ionization dynamics. The modification is also shown to be able to enforce a more consistent particle count per cell in MITL simulations.

## 4. MODEL TESTING AND EVALUATION

This project endeavored to formalize a process for careful evaluation of the various codes and algorithms as they were developed and delivered. A core group of infrastructure developers and analysts was concerned with issues involving verification problem development, verification of plasma code including comparison between codes, and exploratory analysis of power flow physics.

Section 4.1 summarizes the usage of the Code Comparison Infrastructure for supporting disciplined comparison between various plasma physics codes. Section 4.2 provides an overview of a set of test problems that were developed for analysis purposes. Section 4.3 surveys a set of problems and results that were achieved along a sequence of power flow exemplar problems.

### 4.1. Reproducible Science Environment

A software environment was developed for the purpose of enabling an improved level of understanding of the relevant plasma physics, as a disciplined approach to analyst driven simulation quality. This environment is termed the Code Comparison Infrastructure (CCI). It consists of a standardized method to deliver access to several code capabilities in the proper Linux environment and utilities; this includes easily built test problems stored (and maintained) in a separate repository. A detailed paper describing this work, as well as lessons learned from this project, has been submitted for publication [97]. One of the interesting aspects of this paper is a probabilistic analysis that shows how the approach and the associated CCI turned out to be so helpful [97]. The approach utilizes modern software development tools to enable the foundational principles for reproducible science and computer modeling. The operating principles are *Verification, Validation, Continuous Integration, Cross Check* and *Continuous Improvement*. The CCI was utilized for many of the examples described in Sections 4.2.1 and 4.3.

The development of this disciplined, repository software-based analysis and infrastructure is motivated by future machine designs; the supporting modeling work is the foundation for not only understanding present Z-Machine experiments and power flow geometries, but also for justifying potential future next-generation pulsed power accelerator designs in a highly defensible and reproducible way.

The present environment is conveniently accessible for executing calculations on several institutional clusters at Sandia, as well as a special set of hardware that contains both shared memory CPU (two Intel Xeon Gold CPUs with 20 physical cores (40 virtual cores) each) and GPU capabilities (four Nvidia Tesla V100 GPUs). The availability and easy accessibility to GPUs in a standard multi-core environment was very helpful.

Foundational to the software environment that enabled the simulation results to be reproducible is the test harness tool called VVTest. Both interactive and automated workflows were used extensively during development. Additional capabilities needed from VVTest were identified and implemented over the course of this project. The primary advance was the addition of general test dependencies, which allows complex verification sequences to be broken out into distinct and independently executable components, followed by a final analysis to be completed (e.g., model comparisons). Also noteworthy was the addition of built-in support for GPU device management. Simulation analysts were provided the option to execute on a GPU-enabled platform, as well as the option to execute multiple concurrent tests, using a common pool of GPUs on the platform.

The driving force behind developing tests for the CCI is a unified environment that allows analysts to write tests/simulations for multiple codes and to have each code feel the same. Part of that standard feel was to allow tests to be seamlessly executed as either CPU or GPU tests without any user intervention and minimal boilerplate in the tests. This goal was fully achieved for all flavors of EMPIRE (EM, PIC, Fluid, and Hybrid). This allowed analysts to quickly scope out tests to evaluate the best platform and architecture to run on. It also allows tests to be easily portable between HPC resources for either performance testing or simply because one resource is idle at the time of need. The present direction of CCI development is to dynamically manage heterogeneous resources (CPUs and GPUs) concurrently. This capability is currently implemented inside VVTest, but has not been widely exercised in our test suite.

The LDRD project repository was utilized to prototype automated inclusion of desorption data fits generated by the TPD experimental group. Snapshots of data and associated TPD analysis code were delivered to the LDRD test repository and placed in a convenient location accessible by the analyst code comparison scripting. An rudimentary Linux-based testing methodology for this software was implemented. A Python-based VVTest methodology for scanning and extracting neutral molecule desorption data for use by the plasma physics codes, in the necessary format for both EMPIRE and CHICAGO, was developed and is in present use; this is discussed with respect to the planar MITL and half-o-lute tests discussed in Section 4.3. The next step in the development of this interface is to encourage the TPD analysis group to sponsor its own Git repository for data storage, associated fitting software, and embedded software quality testing. Taking advantage of software and numerical expertise working highly with the experimental group is recommended as a next step. For example, multiple users with separate analysis repositories could bring in the TPD repository information from the original source repository using Git submodules. Proper management of distributed data sources for access by computational analysts is a key lesson learned applicable for future work.

A software development and tool-centric approach has been generally identified as robust, disciplined, and amenable to multiple applications and situations. Future possibilities for extensions to the CCI include: (1) additional support for access to and comparison with the validation data; (2) robust extensions to developing machine learning and multi-fidelity model approximations; (3) improved direct Git repository access to sources of modeling input data obtained from TPD experiments; and (4) associated data analysis that is expertly maintained. In this regard, proper design of accessible repositories of the correct scope, size, and supporting teams is believed to be important for future system-level success.

## 4.2. Verification Studies

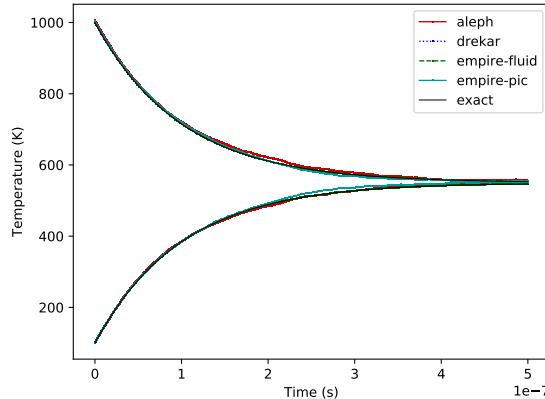
### 4.2.1. *Thermalization Study*

Of significant interest in plasma physics modeling is how a collisionless kinetic fluid model may transition to a collisional fluid model. Fluid model transport coefficients should also agree with kinetic models in the highly collisional limit. As a first step in the process, it is desirable to compare neutral kinetic collision models with fluid transport models to demonstrate that in the highly-collisional limit the same solution is achieved. For this problem we utilized four plasma physics codes: ALEPH ([49]), DREKAR ([113, 83, 112]), EMPIRE-FLUID (Section 3.1.3) and EMPIRE-PIC (Section 3.1.2,[9]).

A Maxwell molecule interaction force relationship has an inverse fourth power potential between two particles and an inverse fifth power in the forces. It has a collision kernel that is exactly integrable. This allows for the development of exact collision frequencies in the case of a fluid model and analytic cross sections for use in kinetic models without prior knowledge of the distribution functions. Therefore, the collision operators are exact and consistent between fluid and kinetic models and thus provide for rigorous verification between fluid and kinetic modeling. We studied a test problem that begins with an initial difference in velocity distribution function for two separate species and is then allowed to approach thermal equilibrium using both kinetic and fluid modeling approaches. The problem specification was taken from [114]. In that report, this problem was used to verify EMPIRE-PIC's and EMPIRE-FLUID's collision operators using an identical problem. Here, we extend this test problem by adding an additional fluid model result using the DREKAR code and an additional kinetic model result using the ALEPH PIC code. For the fluid models, the simulations are effectively zero-dimensional with no advective transport and no spatial derivatives. The computational domain is triply periodic, and the governing equations reduce to ordinary differential rate equations. For the kinetic models, the computational domain is a single-element “box” with doubly periodic boundary conditions. The governing equations remain the same and solutions are represented as particle averages within the cell. The particles are neutral such that no electromagnetic fields are present.

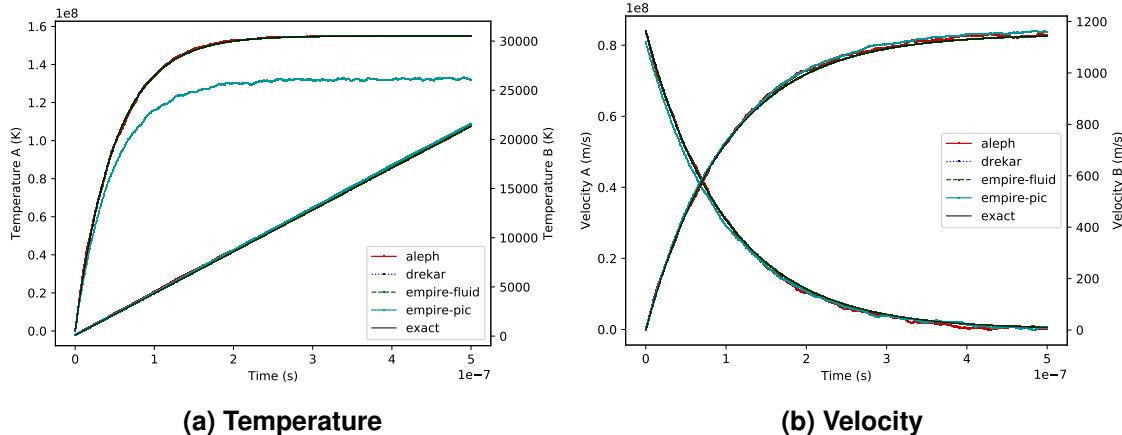
During this study, we initially obtained good agreement for an undrifted scenario between the exact solution, ALEPH, EMPIRE-PIC and EMPIRE-FLUID. However, there were differences with EMPIRE-FLUID outside expected bounds based on the mesh and time step sizes informed by previous rigorous verification convergence studies. Further investigation facilitated by the CCI identified inconsistent scaling of the characteristic time scale with very good agreement achieved once this was fixed. This highlights the value of the CCI in early detection of subtle issues that can easily be missed, especially for codes with inherent stochasticity. We also discovered a difference between EMPIRE-FLUID's and DREKAR's source terms that was later diagnosed as a difference in the treatment of the reduced particle mass. Once this was resolved, all four codes and the exact solution gave consistent solutions for the evolution of temperature as shown in Figure 4-1.

Another example highlighting the value of the CCI is shown in Figures 4-2 and 4-3 for Maxwell molecule results including large and small particle drift velocities. Early in the project, the EMPIRE-PIC code produced good agreement with the exact solution (see [114]) but later produced transient temperature profiles that deviated beyond established acceptable tolerances

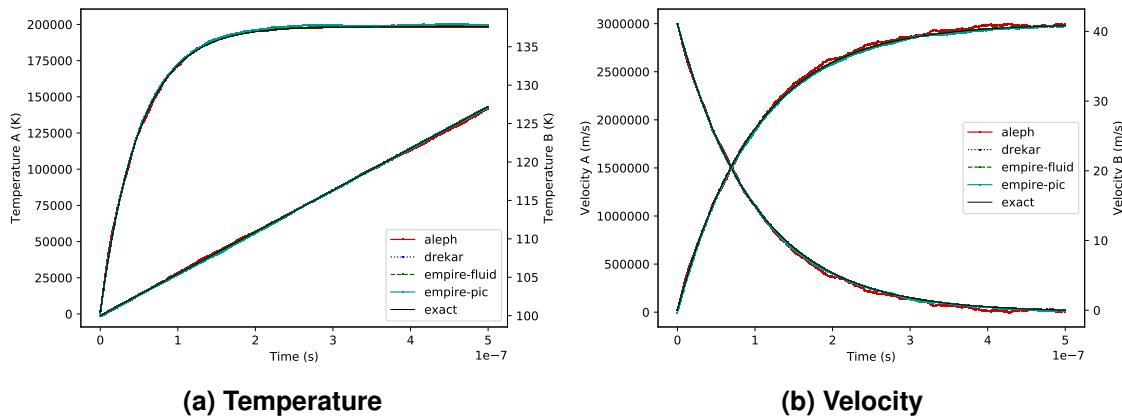


**Figure 4-1. Undrifted Maxwell molecule thermalization temperatures for species A (top curves) and B (bottom curves) compared to exact solution.**

when included in a cross-code comparison as shown in Figure 4-2. Closer examination revealed that EMPIRE-PIC had evolved to include relativistic effects; these fully explained the regression from the previous behavior. In the original test the velocity was  $8.4 \times 10^7 \text{ m s}^{-1}$ , or roughly  $0.28c$  where  $c$  is the speed of light. Reducing the drift velocity to  $.01c$  so that relativistic effects were smaller than other expected contributions to cross-code differences (e.g., time step size, mesh size and particle statistics) restored good agreement between the codes and with the corresponding exact non-relativistic solution, as shown in Figure 4-3. The CCI thus helped identify and clarify the role of the relativistic model in the EMPIRE-PIC code. Appreciating differences, such as these examples demonstrate, is extremely important when choosing among several simulation codes to perform credible analyses, and is one key outcome enabled by the CCI.



**Figure 4-2. High-speed drifted Maxwell molecule thermalization temperatures (left) and velocities (right) for species A and B compared to non-relativistic exact solution at a  $.28c$  drift velocity.**



**Figure 4-3. Drifted Maxwell molecule thermalization temperatures (left) and velocities (right) for species A and B compared to non-relativistic exact solution at a  $.01c$  drift velocity.**

#### 4.2.2. Neutral Expanding Slab, Cylinder, and Sphere Test Problems

In the process of verifying codes and implementing new algorithms, it is extremely important to have recourse to test problems admitting exact solutions. This is especially true during the development of hybrid fluid/kinetic methods for the Vlasov and Boltzmann equations that intrinsically contain both numerical fluid algorithms and kinetic solution methods.

A set of exact solutions and their Python implementations were generated for the purpose of examining and understanding numerical methods [100]. The force-free Vlasov equation

$$\frac{\partial f}{\partial t} + \mathbf{v} \cdot \nabla_{\mathbf{x}} f = 0 \quad (4.1)$$

is solvable exactly by the method of characteristics. For a uniform Maxwellian (equilibrium) slab, cylinder and sphere, these expanding solutions were obtained and analyzed to compute the mass, momentum, pressure and heat flux as a function of space and time. A moment generating function approach was adopted for effectively computing the moments. Expansion problems encompass several challenging physical processes in tandem, and therefore provide a highly relevant solution space for examining the quality of numerical methods.

One immediate benefit of using these exact solutions was to effectively assess noise arising in the kinetic particle methods [101, 97], especially at the expansion front/vacuum interface. The decreasing quality of the numerical higher order moments are potentially problematic for the effectiveness of hybrid  $\delta f$  algorithms that feed back moments from the kinetic solution to guide the closures of the fluid solver.

The kinetic expanding slab problem was further extended from a slab expanding into vacuum to include an infinite outer region with a different initial distribution function and thereby represent expansion of plasma into a lower density background. Further, an Euler fluid version of the expanding slab problem into a background gas was developed by carefully following the classical literature and providing an implementation in Python which resolves the leading shock and expansion fan as well as the reflection of the backward facing release wave off the slab centerline. This solution was used to rigorously verify fluid solver results [100].

The neutral expanding solutions can be used with electromagnetic forces included by implementing multiple species with the same initial distribution functions as long as the total charge sums to zero. Numerical noise can still appear, and thus this problem is a sensitive indicator of the effectiveness of discretized numerical methods and coupling between electromagnetics and plasma transport, since the total electromagnetic force throughout the calculation should be maintained at zero.

Obtaining multifluid solutions to plasma expansions into a vacuum with Eulerian-based formulations are difficult to achieve and require specialized flux functions. A variation of the expanding slab test was designed around the Euler fluid solution mentioned above with non-zero pressure outside of the inner plasma slab in order to evaluate multifluid solutions as well as PIC solutions. The initial plasma contained two species, electrons and positrons of equal mass and opposite charge. The number density in the core was  $n_0 = n_{e0} = n_{p0} = 1.0 \times 10^6 (m^{-3})$ , the temperature was  $T_0 = T_{e0} = T_{p0} = 10,000 K$  and the pressure ratio between the core plasma and

background was 100, the adiabatic index was  $\gamma = 5/3$ . The one-dimensional domain was  $x = [-L_x/2, L_x/2]$  where  $L = 0.02(m)$  and the plasma core extent was  $x_p = [-L_p/2, L_p/2]$  where  $L_p = 0.001(m)$ . The initial velocity was zero. The time step was reduced by one half with each finer resolution mesh to maintain a nearly constant CFL condition, and the plasma frequency was well-resolved at the chosen time step sizes. Figure 4-4 shows the number density profiles for five mesh resolutions,  $r = [1, 2, 4, 8, 16]$ , corresponding to the number of cells in the  $x$ -direction  $N_x = [128, 2048]$ . The multi-fluid solutions are compared with the exact fluid solution, and the PIC codes are compared to the exact kinetic solution at time  $t = 2 \times 10^{-9}(sec)$ . The multiple mesh resolution fluid solutions seem to indicate spatial convergence to the exact solution. Establishing convergence of the kinetic solutions is more difficult due to stochastic noise requiring refinement of the number of particles as well as mesh spacing and time step size. For refinement level  $r$ , the number of particles/cell was  $N_c = 128 * r^2$  for both EMPIRE-PIC and ALEPH. The PIC solutions are approaching the exact solutions under mesh/time step/particle refinement. Formal verification of the two multi-fluid codes and two PIC codes using these exact solutions will be reported in an upcoming publication [100].

In order to assess spatial convergence properties, fluid and PIC codes are considered separately. For DREKAR and EMPIRE-FLUID we compute the convergence of  $L_1$  error norms where the cell error is defined as

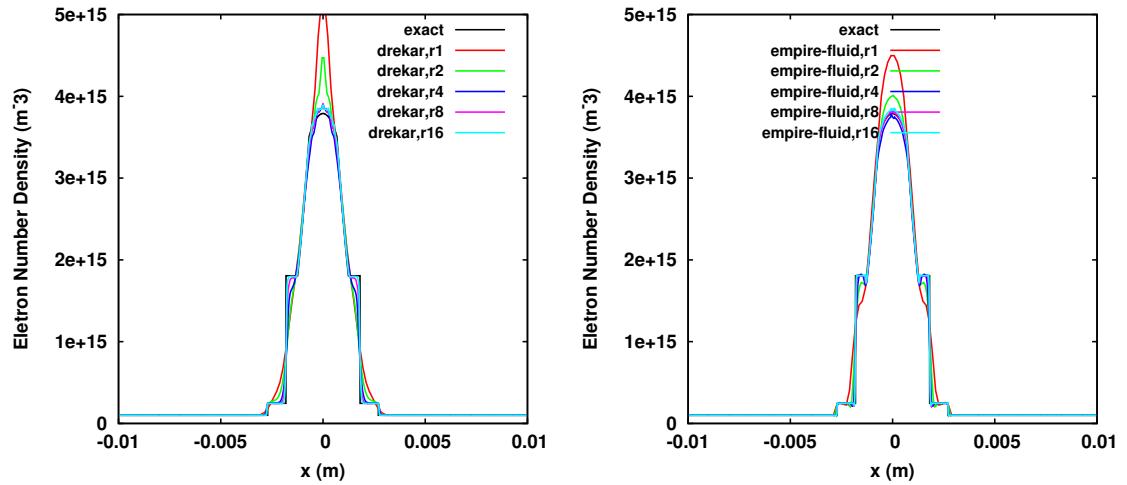
$$\mathcal{E}_i = | \langle f(x) \rangle_i - \langle f^h \rangle_i | \quad (4.2)$$

$\langle f(x) \rangle_i$  is the exact solution integrated over the cell (in this case using the trapezoidal rule rule) and  $\langle f^h \rangle_i$  is the cell averaged simulation value within cell  $i$  and  $f = (n_e, u_e, P_e)$  is a set of quantities of interest. The  $L_1$  error norm is defined as

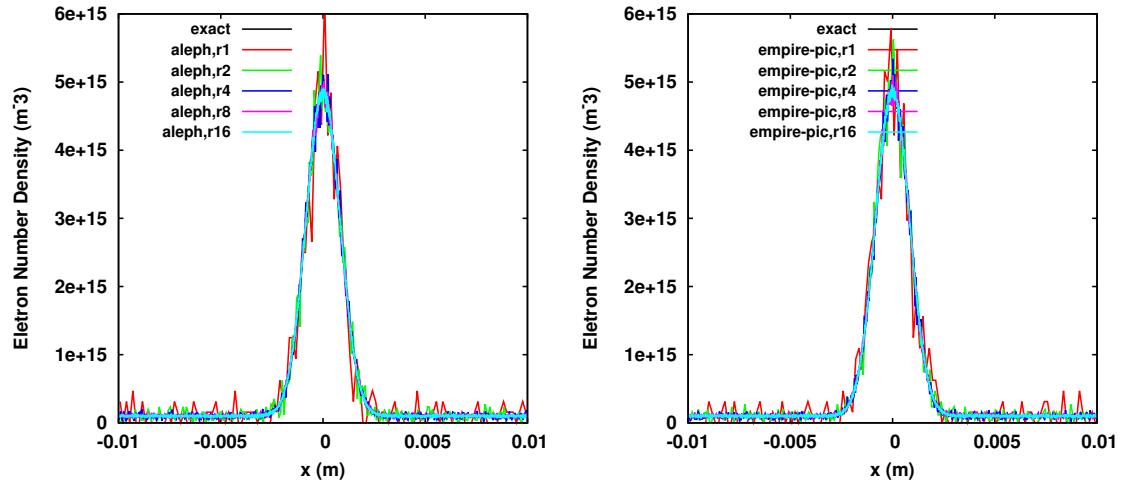
$$\mathcal{E}^{L_1} = \frac{1}{N_x} \sum_i^{N_x} \mathcal{E}_i \quad (4.3)$$

where  $N_x$  is the number of cells.

Representative norms are plotted for DREKAR and EMPIRE-FLUID in Figure 4-5. These norms are normalized by the initial number density  $n_0$ , sound speed ( $c_{s0} = \sqrt{\gamma k_B T_0 / m_e}$ ) and pressure  $P_0 = n_0 k_B T_0$ . We also present the rates of convergence in Table 4-1 along with plots in Figure 4-6 of fits to a power law error model of the form  $\mathcal{E}^{L_1} = C(R)^p$  for refinement level  $R$  using the  $n_e$  field results. The parameters  $C$  and  $p$  are fit using a least-squares regression to the data (points equally weighted) with  $p$  representing an overall rate of convergence when refining time step, mesh size and number of particles consistently together. For the kinetic codes, EMPIRE-PIC and ALEPH, the number of particles per cell was increased as the square of increasing refinement level to resolve mean statistics (moments) at the same order as the lower of spatial or temporal convergence. Figure 4-6 indicates the PIC codes behave as expected for the  $n_e$  result, i.e.  $p \approx 1.0$ . It should be noted that higher order moments (such as pressure) show markedly lower rates of convergence due in large part to the requirement of significantly more particles per cell to resolve higher order statistics. This has implications for hybrid methods that may require feedback of such moments from PIC treatments to fluids. The fluids codes are both converging the error at a rate less than one. Comparing these rates to what is observed from Euler solutions for Sod's shock tube, where a combined rate of order one for the shock and order  $p/(p+1)$  for the contact, where  $p$  is the polynomial order of the scheme, one for both codes, these results are in reasonable

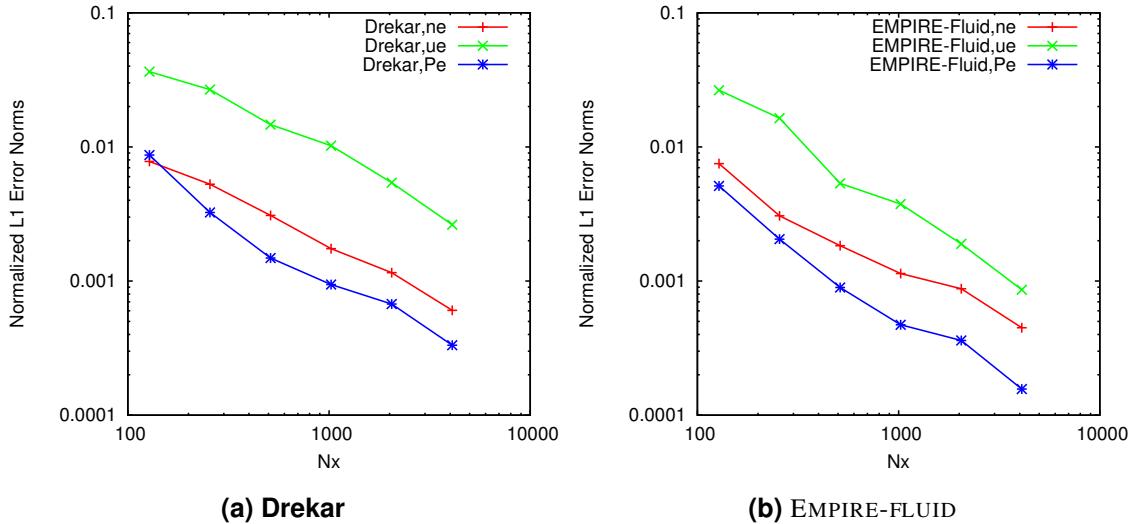


**(a) Multi-fluid number density profiles compared to the exact solution with mesh refinement**



**(b) PIC number density profiles compared to the exact solution with mesh refinement**

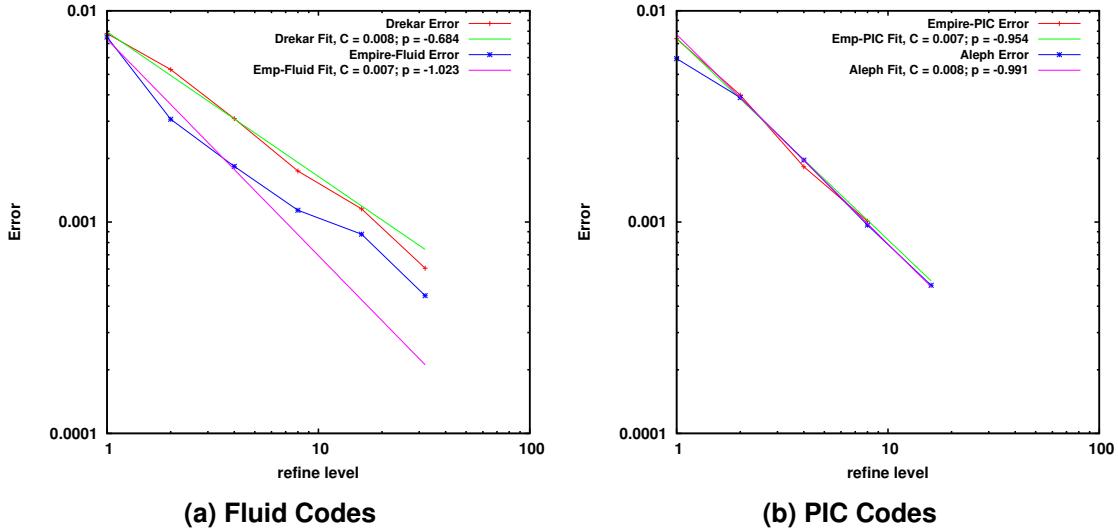
**Figure 4-4. Neutral expanding slab with non-zero background pressure.**



**Figure 4-5. Relative  $L_1$  error norms for different mesh resolution both DREKAR and EMPIRE-FLUID. Norms are normalized by  $n_0, c_{s0}$  and  $P_0$  for  $n_e, u_e, P_e$  respectively.**

Fluid Code	Mesh Pair ( $N_x$ )	Number Density ( $1/m^3$ )	Velocity ( $m/s$ )	Pressure ( $kg/m/s^2$ )
DREKAR	128-256	0.5655	0.4448	1.4200
	256-512	0.7709	0.8675	1.1290
	512-1024	0.8231	0.5221	0.6561
	1024-2048	0.5952	0.9147	0.4828
	2048-4096	0.9332	1.0410	1.0210
EMPIRE-FLUID	128-256	1.2920	0.6919	1.3180
	256-512	0.7402	1.6150	1.1970
	512-1024	0.6901	0.5110	0.9241
	1024-2048	0.3765	0.9870	0.3873
	2048-4096	0.9665	1.1370	1.204

**Table 4-1. Rates of convergence of the cell based  $L_1$  error norms for electron number density, velocity and pressure.**



**Figure 4-6. Log-log plot of Error vs. refinement level  $R$  for  $n_e$  results together with least-square fits to the error model  $\mathcal{E}^{L_1} = C(R)^p$**

Code	Field	$C$	$p$
DREKAR	$n_e$	0.008	-0.684
	$u_x$	0.019	-0.643
	$P$	0.001	-1.268
EMPIRE-FLUID	$n_e$	0.007	-1.023
	$u_x$	0.014	-0.941
	$P$	0.001	-1.238
ALEPH	$n_e$	0.008	-0.991
	$u_x$	0.152	-1.014
	$P$	0.002	-0.055
EMPIRE-PIC	$n_e$	0.007	-0.954
	$u_x$	0.176	-1.149
	$P$	0.003	-0.147

**Table 4-2. Fitting parameters to the error model  $\mathcal{E}^{L_1} = C(R)^p$  based on various fields (moments),  $n_e$ ,  $u_x$ ,  $P$**

agreement. Given the difficulty due to a very low back pressure, as well as the presence of numerically generated electromagnetic "noise", the results are quite acceptable.

#### **4.2.3. *Surface Joule Heating Modeling Analysis***

Kinetic modeling codes like EMPIRE and CHICAGO must implement heating models to compute an effective surface temperature. This temperature is used to drive model fits for desorption physics detailed in Chapter 2.

One of the physical effects causing temperature rise is Joule heating in the accelerator wall. A simple closed formula model that is dependent on a linear polynomial field history has been historically utilized [57, 58, 103]. During this project EMPIRE implemented a discrete numerical version of the magnetic field history diffusion model in order to capture non-linear property effects, but at potentially significant computational expense that may require additional load balancing. Intermediate approaches were investigated that maintain the underlying constant property exact solution model but fully support a time-dependent history in a compact way. The effect of taking into account this history and its potential effect on desorption release timing, including both surface contaminant desorption as well as melt release, was investigated and published [99]. It appears likely that the effect of improved history modeling may be most validatable for contaminant releases associated with electrode melt.

This study is an example of the detailed analysis required to fully understand, verify, and eventually validate all the boundary condition models entering into a multi-physics pulsed power modeling code.

#### **4.2.4. *Two-Stream Instability***

The two-stream instability (discovered by Bohm and Gross in 1949 [12]) has become a staple verification problem for plasma simulation codes (e.g. [18, 20, 3, 29, 93, 106, 46]). Its name derives from the instability that arises in a quasi-neutral plasma involving two counter-streaming populations of the same or opposite charge. These interpenetrating beams produce density fluctuations in each other that conspire to amplify, rather than remedy, the original disturbances leading to spatial bunching. The growing electric field resulting from these mutually reinforced perturbations further interacts with the drifting charges, where the energy transfer can be shown to be directed from the particles to the wave, on average leading to exponential growth in the electric field, i.e. instability [126]. Linearized theory predicts this growth continues unhinged; however, kinetic simulations (which include nonlinear effects) reveal a saturation in the long-time limit after the wave has grown large enough to trap a majority of the particles participating in the energy exchanges (cf. Figures 4-7, 4-8). This steady-state is described by undulations in the electrostatic energy, where potential energy in the wave is interchanged with kinetic energy of trapped particles (and vice versa) in a manner that is completely analogous to a mechanical pendulum. It can be shown that the period of these undulations roughly corresponds to the bounce period of trapped particles [120], and so the late-time behavior is dominated by nonlinear effects.

The problem is particularly appealing as a verification problem in the code comparison infrastructure (CCI) for the following reasons:

1. it has a straightforward figure of merit to measure simulation success, i.e. the deviation in the simulated growth rate of the electric field versus the exact value from linearized theory;
2. it has a linear regime in early times where fluid, kinetic, and hybrid codes should all agree;
3. it has a nonlinear regime in late times where fluid codes break down yet where kinetic/hybrid codes should maintain their fidelity; and
4. it presents an ideal testbed for hybrid modeling approaches (cf. Section 3.1.4) where both its accuracy and performance can be scrutinized against conventional models (fluid or kinetic).

We have designed and installed a VVTest structure on the GC LDRD GitLab for the following codes: ALEPH, DREKAR, EMPIRE-PIC, EMPIRE-FLUID, and EMPIRE-HYBRID whose principal differences as applied to this problem are summarized in Table 4-3.

We choose to model two non-relativistic counter-flowing electron beams with speeds  $u_b = 0.2c$  for our verification problem. A specific test is defined by choices for the electron beam density and beam temperature (Table 4-4). A domain length  $L$  is chosen to be the wavelength

$\lambda = 2\pi / (\sqrt{3/8}\omega_{pe}/u_b)$  of the fastest growing wave in the cold fluid case [126]. This sets up a large enough domain so that any unstable waves arising in our warm simulations fit inside the domain and their growth can be monitored.

The electric field arising from spatial bunching and interaction with drifting electrons is primarily in the direction of the beam. Therefore, the problem is one-dimensional in that a preferred direction exists. The waves evolve independent of boundaries so that periodic conditions for both the fields and particles/fluid moments are appropriate. When periodic field conditions are unavailable, we have found that enforcing homogeneous Dirichlet conditions on the electric potential does not interfere with the instability physics at least for electrostatic codes such as ALEPH. Whenever a one-dimensional simulation is not supported by a code, a thin transverse dimension with periodic boundaries is created using the minimum number of cells permitted by the code. The spacing in this perpendicular direction is chosen to be small or otherwise equal to that in the preferred direction. The periodicity in the perpendicular direction ensures that any particles developing oblique velocity components that happen to exit the domain do so only to

Code	Paradigm	Dimension	Fields/solver	Field BCs	Particle BCs
ALEPH [49]	PIC	1	ES/CG explicit	Dirichlet	periodic
DREKAR [113, 83, 112]	fluid	3	EM/CG implicit	periodic	periodic
EMPIRE-PIC [9]	PIC	2	EM/CG explicit	periodic	periodic
EMPIRE-FLUID (§3.1.3)	fluid	2	EM/DG explicit	periodic	periodic
EMPIRE-HYBRID	PIC/fluid	3	EM/CG explicit	periodic	periodic

**Table 4-3. ES/EM = electrostatic/magnetic, (D)CG = (dis)continuous Galerkin.**

return and continue on their path along the preferred direction with minimal consequence. As the electric fields are strongly in one direction, these events are probably not significant but is an example of how a wider range of codes can access a given problem even without the ability to model the specific prescription given “on paper” is not strictly met. Each code settles on its own grid spacing and time stepping decisions based on code-specific stability constraints as well as considering physical demands from resolving the instability wavelength, electron Debye lengths, and plasma frequencies. As concerns the spatial discretization, we use 100 cells to cover the domain in the preferred direction for all codes except EMPIRE-HYBRID which uses only 25 cells. This coarser mesh is employed purely for cost savings as the method is more expensive and also requires a 3D grid with a minimum of 2 cells in each periodic direction. Therefore the problem is simulated on a  $25 \times 2 \times 2$  grid. While both of these grids cannot resolve the Debye length in any of the tests, we do not see evidence that its consequences, e.g. numerical grid heating, corrupts the results which is probably due to the short durations being simulated (cf. Table 4-4). For more rigorous pursuits, such as convergence analysis, we would need to take this concern more seriously, but for the present purposes of baselining the codes and to contribute to nightly regression tests we find these decisions economical while still producing even excellent agreement with theory at times (e.g. ALEPH). As concerns the time stepping decisions, we take 50 steps per plasma period for the PIC codes, for the EM fluid codes we choose a step based on a satisfying the speed-of-light  $CFL = 0.9$ , and for hybrid we discover a much stricter step than that anticipated ( $CFL \leq 0.1/\sqrt{3}$ ).

Parameter	values
electron beam density $n_b$	$1e+14 \text{ m}^{-3}, 1e+15 \text{ m}^{-3},$
electron beam temperature $T_e$	$1e+04 \text{ K}, 1e+05 \text{ K}, 1e+06 \text{ K}$
electron beam speed $u_b$	$0.2c = 59958491.6 \text{ m/s}$
domain length $L$	$2\pi/(\sqrt{\frac{3}{8}} \frac{\omega_{pe}}{u_b}) \text{ m}$
duration $T$	$5 \cdot \tau_{pe} \text{ periods}$

**Table 4-4. Shared parameters in every test.** The electron plasma period  $\tau_{pe} = 2\pi/\omega_{pe}$  where the angular frequency  $\omega_{pe} = \sqrt{2n_{eb}e^2/(m_e\epsilon_0)}$  corresponds to both beams. The domain length  $L$  is one wavelength  $L = 2\pi/k_{max}$  in extent which corresponds to the fastest growing wave with rate  $\gamma_{max} = \max_k \text{Im } \omega$  whose wave number is  $k_{max}$  according to cold fluid theory.

The VVTest structure interfaces with each code through a parameterizable input file courtesy of APREPRO, orchestrates all simulation runs, performs any post-processing analysis, and generates summary tables of the results along with code-specific (or cross-code) plots. For the full group testing defined in Table 4-4, we show the test results summary in Table 4-5 which is written out automatically through VVTest. As this table also reports timings for each run, it is prudent to delineate the differences between the codes and paradigms (fluid vs. kinetic vs. hybrid) from at least a top-level so that these timings can be understood in context.

The particle-in-cell (PIC) codes include calculations which fluid codes do not (e.g. particle movers), and vice versa (e.g. Riemann solvers). ALEPH [49]/EMPIRE-PIC (§ 3.1.2) interleave ES/EM field solves with particle pushing steps (e.g. Figure 3-1). As concerns ALEPH, the problem is solved on a 1D domain decomposed into 100 cells on which the electrostatic potential

code	density [m <sup>-3</sup> ]	temperature [K]	$\frac{\gamma_{sim} - \gamma_{theory}}{\gamma_{theory}} \times 100\%$	cpu time [min / proc]
ALEPH	5.000e+14	1.000e+04	-0.3845577	2.934 min / 1 proc
DREKAR	5.000e+14	1.000e+04	-2.0471102	2.853 min / 32 proc
EMPIRE-FLUID	5.000e+14	1.000e+04	-1.3731070	4.180 min / 8 proc
EMPIRE-HYBRID	5.000e+14	1.000e+04	-2.9976224	12.77 min / 16 proc
EMPIRE-PIC	5.000e+14	1.000e+04	-0.2712892	0.766 min / 8 proc
ALEPH	5.000e+14	1.000e+05	-1.1259112	3.120 min / 1 proc
DREKAR	5.000e+14	1.000e+05	-1.8927217	2.919 min / 32 proc
EMPIRE-FLUID	5.000e+14	1.000e+05	-0.6077596	4.398 min / 8 proc
EMPIRE-HYBRID	5.000e+14	1.000e+05	-2.3942097	13.07 min / 16 proc
EMPIRE-PIC	5.000e+14	1.000e+05	-1.0699148	0.806 min / 8 proc
ALEPH	5.000e+14	1.000e+06	0.3406601	3.093 min / 1 proc
DREKAR	5.000e+14	1.000e+06	-1.5695018	3.081 min / 32 proc
EMPIRE-FLUID	5.000e+14	1.000e+06	-0.1273656	4.433 min / 8 proc
EMPIRE-HYBRID	5.000e+14	1.000e+06	-8.2564916	13.25 min / 16 proc
EMPIRE-PIC	5.000e+14	1.000e+06	-4.7598521	0.782 min / 8 proc
ALEPH	5.000e+15	1.000e+04	-0.1563418	3.056 min / 1 proc
DREKAR	5.000e+15	1.000e+04	-1.9199136	2.327 min / 32 proc
EMPIRE-FLUID	5.000e+15	1.000e+04	-1.2011811	2.168 min / 8 proc
EMPIRE-HYBRID	5.000e+15	1.000e+04	-6.6460535	10.21 min / 16 proc
EMPIRE-PIC	5.000e+15	1.000e+04	-0.1154789	0.868 min / 8 proc
ALEPH	5.000e+15	1.000e+05	-0.8445391	3.053 min / 1 proc
DREKAR	5.000e+15	1.000e+05	-2.0046369	2.401 min / 32 proc
EMPIRE-FLUID	5.000e+15	1.000e+05	-0.0790751	2.520 min / 8 proc
EMPIRE-HYBRID	5.000e+15	1.000e+05	-5.5822316	10.30 min / 16 proc
EMPIRE-PIC	5.000e+15	1.000e+05	-0.7328495	0.806 min / 8 proc
ALEPH	5.000e+15	1.000e+06	0.3329511	2.949 min / 1 proc
DREKAR	5.000e+15	1.000e+06	-1.7132287	3.275 min / 32 proc
EMPIRE-FLUID	5.000e+15	1.000e+06	-0.1018737	2.145 min / 8 proc
EMPIRE-HYBRID	5.000e+15	1.000e+06	-2.6139827	10.78 min / 16 proc
EMPIRE-PIC	5.000e+15	1.000e+06	-4.4408987	0.802 min / 8 proc

**Table 4-5. (Two-stream instability) Summary output from VVTest group analysis. The tests are defined by electron beam density, electron beam temperature and the results are reported in terms of the relative error of the quantity of interest (growth rate), and cpu time over number of processors (proc).**

is solved by the continuous Galerkin finite element method (CG-FEM) applied to the Poisson equation. Particle trajectories are integrated according to a Leapfrog scheme where we have decided on 40,000 computational particles to resolve the electron dynamics over 250 time steps (50 points per electron plasma period). We point out that ALEPH is a relatively mature code whose routines are highly optimized. It also solves the electrostatic (ES), rather than the electromagnetic (EM), problem and is the only code in this section that employs a 1D grid. Therefore, its high performance seen here is a reflection not only of its maturity, but also because its problem size is smaller. As concerns EMPIRE-PIC, we model EM fields by a continuous Galerkin (CG) finite element representation and employ a Crank-Nicolson time integrator also over 250 time steps but on a  $100 \times 2$  structured grid (quad elements). We use 20,000 computational particles for each of the two electron streams and the Boris scheme to integrate the particle trajectories. The implementation of this particle mover has been optimized so it is generally expected to be faster than, e.g. explicit EMPIRE-FLUID solver performance, but not faster than ALEPH as its problem size is overall larger (e.g. 2D versus 1D, EM versus ES). For EMPIRE-FLUID, we use 8 processors per simulation to evolve EM fields and the two 5-moment Euler fluids (electron streams) by basis order one discontinuous Galerkin (DG) finite elements using a second-order strong-stability preserving Runge-Kutta (SSP-RK2) explicit time integrator equipped with a simple flux limiter [84]. We take 22,631 steps ( $n_b = 0.5\text{e}+15 \text{ m}^{-3}$ ) or 48,157 steps ( $n_b = 0.5\text{e}+16 \text{ m}^{-3}$ ) depending on the test (Table 4-4). The number of steps is larger in the higher density case since the domain size  $L \sim u_b/\omega_{pe} \sim n_b^{-1/2}$  is smaller, but we use the same number of cells  $N_c$  in the preferred direction ( $N_c = 100$ ). Therefore, these smaller cell widths  $\Delta x = L/N_c$  in the higher density case require a smaller time step  $\Delta t$  to meet the same stability constraint (CFL  $\equiv c\Delta t/\Delta x \leq 1$ ). We point out that not only are the number of steps significantly more in the EMPIRE-FLUID runs compared to EMPIRE-PIC simulations, but its solver has not been optimized to the extent of EMPIRE-PIC implementations. Thus, longer run times are actually expected as compared to EMPIRE-PIC. The current setup for EMPIRE-HYBRID models one PIC stream and one fluid stream where the EM fields are represented as CG finite elements and otherwise we carry over the same solver settings from EMPIRE-PIC and EMPIRE-FLUID for each component. One important difference is that a two-step nonlinear time integrator is currently required which is known to be more expensive than the Crank-Nicolson scheme which is used by the EMPIRE-PIC simulations in this section. For these reasons, EMPIRE-HYBRID is more expensive than EMPIRE-PIC and EMPIRE-FLUID in its current form, and this is reflected in the timing data in Table 4-5. For DREKAR multifluid simulations, the full three-dimensional fluid species equations are solved for two opposing electron streams using a  $100 \times 1 \times 1$  cell mesh. The electric field is obtained from the ES equation. The spatial discretization utilizes CG FEM with algebraic flux correction stabilization. A diagonally implicit RK time integrator (DIRK 1) is used and the entire system of equations is solved using a monolithic iterative Newton method. In general the monolithic solver approach used by DREKAR applied to this problem takes between three to six times longer to solve compared to EMPIRE-FLUID using an explicit time integration.

An example plot output by the group analysis routine is shown in Figure 4-8. This example corresponds an electron beam density  $n_e = 5 \times 10^4 \text{ m}^{-3}$ , with beam temperature  $T_e = 10^5 \text{ K}$ . The time histories of the electrostatic energies for each code shown in this plot reveals visually similar growth rates (slopes), but a quantitative comparison is required for further scrutiny. Further, in a reproducible science (RS) approach, we require a recipe for calculating this growth rate  $\gamma$  that is

both reproducible and hands-free (automated). In published works, the calculation of the growth rate is often done expediently rather than systematically: a time interval is selected where the electrostatic energy appears to exhibit linear growth, and the corresponding growth rate is calculated by some means that is not always described. The first step of eyeballing a correct range introduces an unacceptable degree of variability in the intended calculation; an unbiased eye that auto-selects the optimal range will be required to make this reproducible. Meanwhile, the second step can be done in a variety of ways, so a standardized (“best”) procedure should be agreed upon.

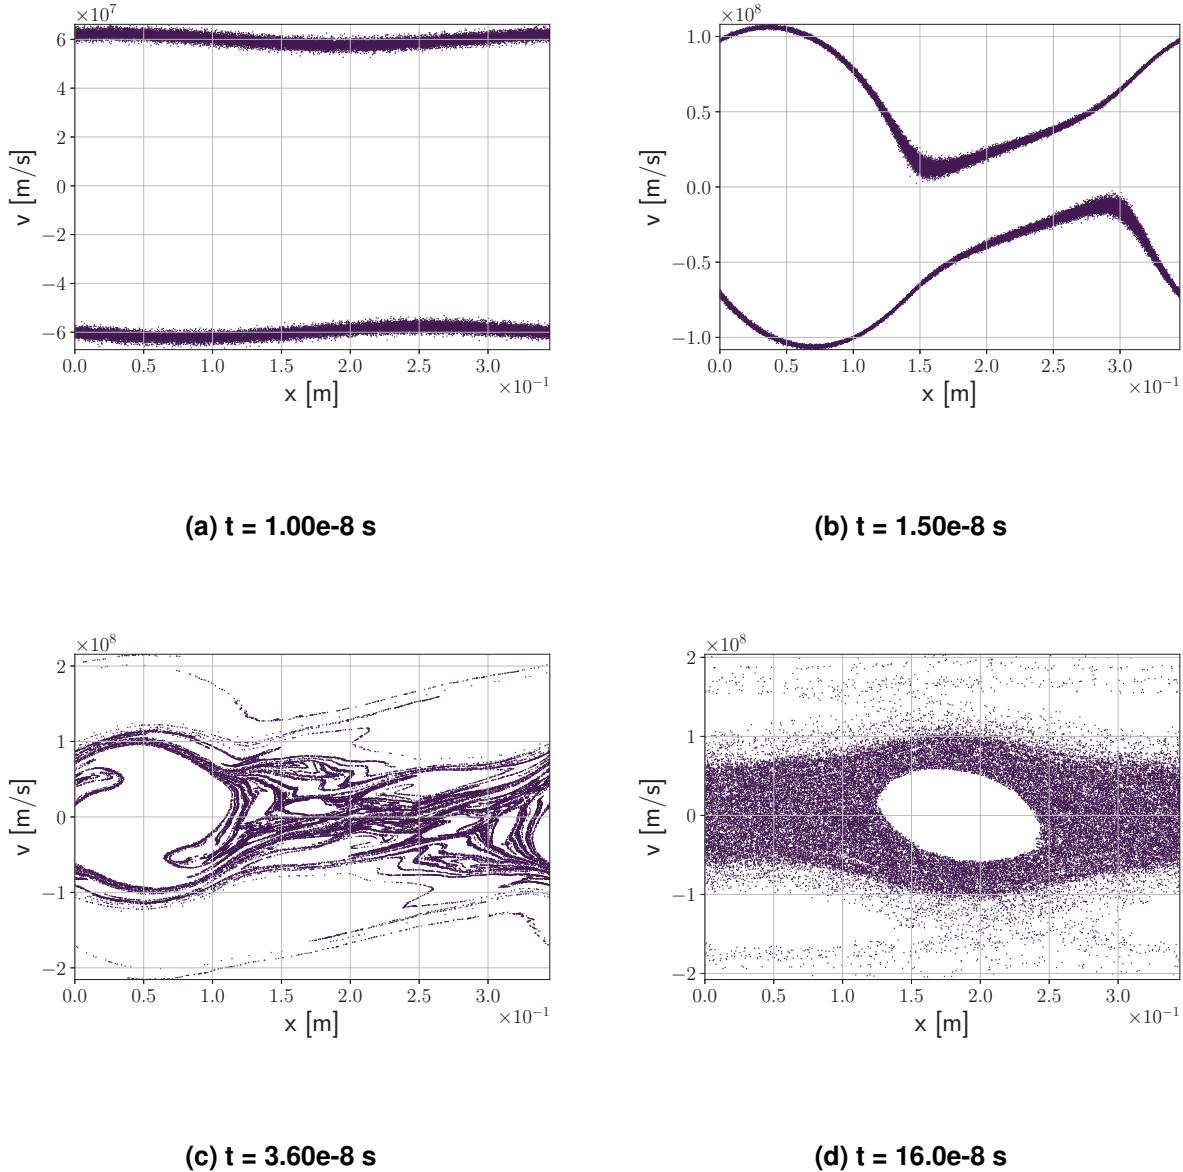
To answer the first issue, the arbitrariness introduced by using the eyeball norm can be removed by replacing this with an actual norm. The coefficient-of-determination ( $R^2$ ) or the chi square ( $\chi^2$ ) statistic are usual candidates. These goodness-of-fit parameters indicate how well a null hypothesis (e.g. linear growth) describes the trends in the underlying data (e.g. the data in one interval versus another). It is straightforward to design a sweeping algorithm to determine which subinterval of data (out of all possible subintervals) achieves the best goodness-of-fit statistic with respect to the model of exponential growth<sup>1</sup>. This is how we engineer the unbiased eye required to systematically determine the interval describing the linear regime. While both statistics provide both reproducible and automated determination of a single best subinterval inside the full interval, in our experience we have found that the chi square statistic is more reliable in discovering the correct interval ( $\approx 1 \sim 2\%$  improvement over  $R^2$ ) for measuring linear growth in the sense of being closer to the theoretical value. Thus, we use the chi square statistic as the appropriate judge for this task.

As concerns the second point, the most common approach in published works is apparently to fit a linear function to the logarithm of the data whereby the fit parameter describing the slope is the desired growth rate. In our experience, we have found that fitting a linear function to the logarithm of the data introduces  $\lesssim 3\%$  propagated error into the final calculation as compared to fitting the same raw data to exponential function. Thus, we fit the electrostatic energy histories to an exponential in our test analysis. More details are provided in the internal report [121].

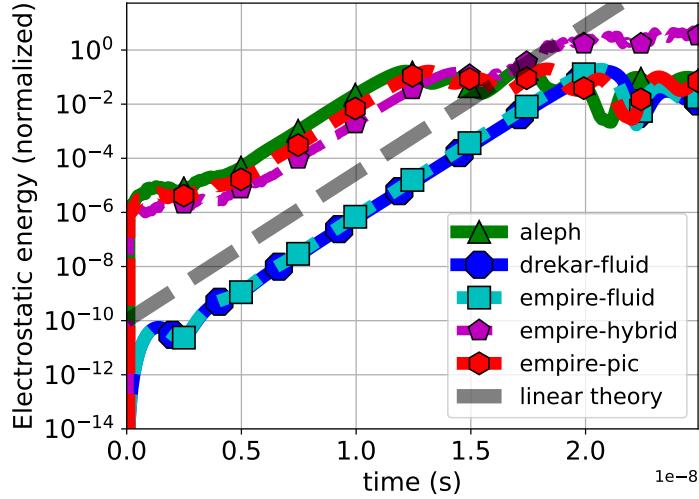
We are currently completing a repurposing of this test that pursues convergence studies for each code in the linearized growth rate. This more rigorous goal is a work in progress, but has already informed a number of small changes to the above prescription to make the calculations much more exact. First, the growth rates for each code are not anticipated to all converge to the same answer. The PIC codes (ALEPH and EMPIRE-PIC) should converge to the growth rate described by the linearized kinetic equations, the fluid codes (DREKAR and EMPIRE-FLUID) should converge to the rate contained in the 5-moment (warm) fluid equations, and the hybrid simulations should converge to the growth rate associated with the system of coupled kinetic and fluid equations. Secondly, the domain length  $L$  should be set to wavelength  $\lambda = 2\pi/k$  associated with the solution  $\omega = \omega(k)$  from the dispersion relations for each code. This requires a

---

<sup>1</sup>note that a “linear growth rate” in plasma physics parlance refers to a growth rate derived from a linearization by writing each quantities  $q$  as a combination of an averaged (“zeroeth order”) term  $q_0$  with a fluctuating (“first order” or “linear”) term  $q_1 = \tilde{q}\exp(ikx - i\omega t)$  for some  $\tilde{q} \in \mathbb{R}$ , i.e.  $q = q_0 + \tilde{q}\exp(ikx - i\omega t)$ . Therefore, the growth rate  $\gamma$  identified in the complex-valued frequency  $\omega = \omega_r + i\gamma$  where  $i = \sqrt{-1}$  strictly means exponential growth,  $\sim \exp(\gamma t)\exp(ikx - i\omega_r t)$ . Of course, it is also the case that a linear growth rate appears linear on semilog axes.



**Figure 4-7. (Two-stream instability) phase space plots of EMPIRE-PIC particle distributions are shown at key times throughout the problem. Panel (a) shows the growth of a perturbation in the electron distribution which corresponds to a kickstart the linear growth of the electric field wave. The electric field continues to grow until it is large enough to trap a majority of the particles in the wave frame as in (b) which quenches further growth. Particle bounce orbits in the wave frame appear as closed trajectories in phase space that encircle  $v = 0$ . Therefore, the picture that develops is the formation of a so-called hole in phase space as seen in (c) which partially, though never completely, fills in due to trapped particle orbits in the long-time limit.**



**Figure 4-8. (Two-stream instability)**  $n_e = 5 \times 10^4 \text{ m}^{-3}$ ,  $T_e = 10^5 \text{ K}$ . The time history of the electrostatic energy is plotted for each simulation. A linear region, that generally does not have the same start or end times from code to code is clearly seen on the semilog axes. The growth rate (slope) of these lines is the quantity of interest for the two-stream instability test. The grey dotted line corresponds to linearized theory, and we point out that it is only the slope that is meaningful.

two-dimensional parameter search to find the maximum growth rate  $\gamma_{max} = \max_k |\text{Im } \omega(k)|$  across all wave numbers  $k$ . Third, looking to the total energy content as we do here is sufficient for the sake of a cross-code comparison on a common problem, but the total energy contains information from all modes. A rigorous convergence analysis must be completed by tracking the linear term, i.e. first mode of the electric field. We have since included additional VVTest structure to allow for this branched test and specialized post-processing (e.g., Fourier analysis). Finally, the convergence of statistical methods (e.g., PIC and hybrid) require robust approaches to reveal convergence. Therefore, we are looking to utilities such as STREEQ for help with this task.

#### 4.2.5. PIC Statistical Noise Analysis

The particle-in-cell (PIC) method is the most widely used numerical approach for simulating fully-kinetic plasmas in three spatial dimensions (3D). Albeit a tremendous success, its major drawback is high numerical noise, much higher compared to that of fluid simulations. The high numerical noise not only may obscure the simulated physics but may also negatively affect coupling of fluid and kinetic algorithms in some approaches, for example the  $\delta f$  method explored during this LDRD.

It is known that this noise scales as  $\sqrt{1/N_p}$ , where  $N_p$  is the total number of macro-particles in the simulation. Therefore, increasing  $N_p$  typically leads to decreased noise in the simulation. However, it is also clear that such approach to reducing the noise quickly becomes prohibitive in large 3D PIC simulations.

A novel statistical analysis was developed [37], in large part during this LDRD, that quantitatively describes noise, and more generally error, in particle-based numerical methods. The analysis pivots on the trade-off between two contributions to the total error: that of bias, which results from the finite size of the computational particles, and that of variance (noise) due to the finite number of computational particles. Optimizing between these two contributions is a process known as the bias-variance optimization (BVO). The major conclusion is that it is possible to decrease noise in PIC simulations by multiple times by carefully choosing the macro-particles' width. A general approach for creating custom particle shapes is proposed and guidelines for practical application of the BVO is described.

Following the guidelines outlined in Reference [37] would allow to improve the fidelity of numerical simulations by multiple times, while keeping the computational load unchanged. Alternatively, the run time could be decreased by a corresponding amount, thus making for a much quicker turnaround in the design of pulsed power experiments.

#### **4.2.6. Vacuum Fields and Particle Dynamics for MITLs**

Magnetically insulated transmission lines (MITL) are used to transport electrical energy to the load on pulsed power accelerators, such as the Z-Machine. The section just before the load, the inner MITL, is subjected to extreme magnetic and electric fields on the order of hundreds of Tesla and tens of megavolts per meter. Under such conditions, the transmission line electrodes are heated (by multiple mechanisms) to hundreds of degrees Celsius and desorb impurities, which then ionize and form a plasma. This plasma can evolve in space and time across the anode-cathode gap, becoming a major channel of parasitic current loss. Understanding and simulating the effects of this plasma is a major scientific goal of the pulsed power sciences program.

Because of their complex geometrical shapes, as well as the physics involved, the design of MITL structures is completed predominantly via the use of multi-physics modeling tools. The software suite **EMPIRE** now includes the multi-physics models necessary to complete such design calculations. A major effort is the verification of this (and other) code against analytically tractable problems. A small class of simplified geometrical MITL shapes lend themselves to analytical treatment. Two such shapes have been explored: the radial cylindrically symmetric MITL and the spherically curved cylindrically symmetric MITL [47]. For these geometries, we have been able to calculate exact electric and magnetic fields in a regime relevant to the Z-Machine for current pulses on the order of 100 ns.

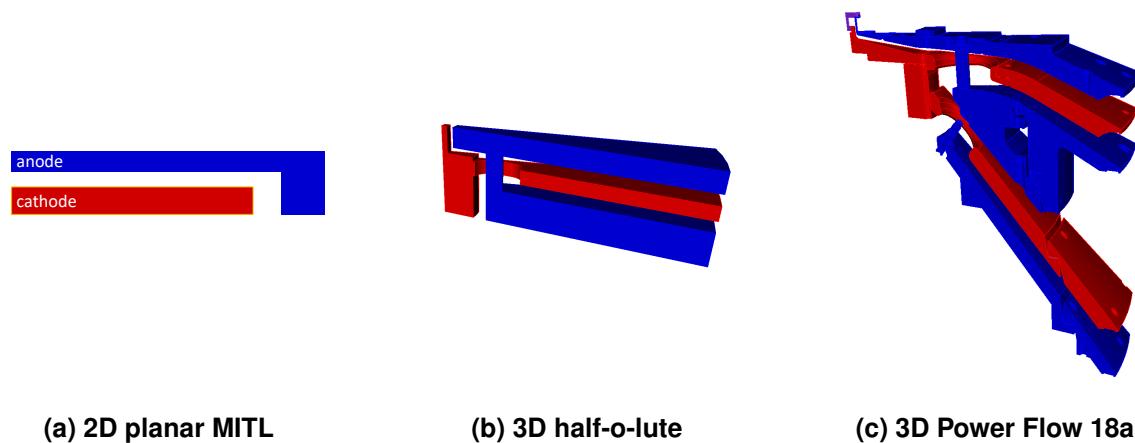
In Reference [47] we study the fields and particle (electron) trajectories in these two geometries. Since the fields are described analytically, the trajectories can be resolved numerically with very high temporal and spatial accuracy. A simplified description in terms of particle drifts is also proposed. The initial results indicate the electron trajectories are well described by particle drifts. The cases when the drift trajectory description fails have been addressed as well. When using drifts, simulation times can be reduced by at least three orders of magnitude.

Put in perspective, this work is a first step towards providing a specialized but highly accurate numerical tool that can be used for code verification. At the same time, after including further

relevant physics such as space charge effects, thermal desorption, and space charge limited emission (SCL), it can be used for quick and efficient design of pulsed power experiments based on this initial set of simplified geometries (the spherical MITL is a particularly good candidate), thus providing a testbed for physical model validation as well.

### 4.3. Power Flow Analysis Exemplars

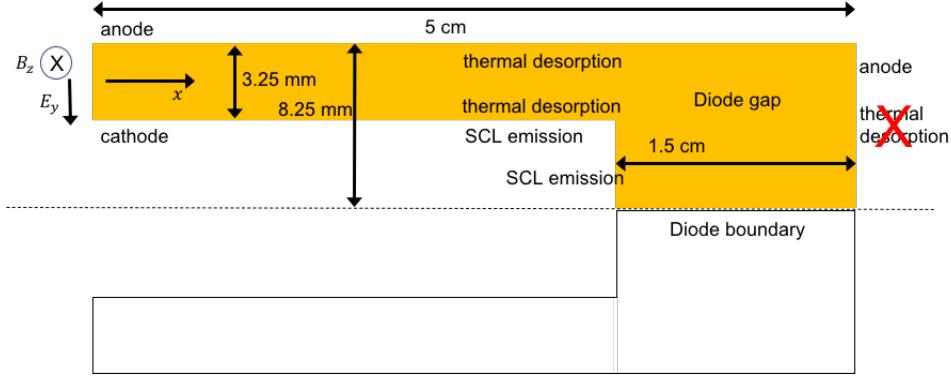
A sequence of pulsed power exemplar problems were outlined (Figure 4-9) to provide a set of stretch goals to evaluate integration progress during this GC LDRD. The first of these goals was to implement a two-dimensional MITL geometry test problem that is quickly executed, and could be used to understand space charge limited (SCL) and desorption contaminant (TPD) boundary conditions. Another goal was to demonstrate that the fitted TPD data parameters collected as described in Section 2.2 could be automatically retrieved and utilized in the CCI-based test problem parameter scans. The results of this study are described in Section 4.3.1.



**Figure 4-9. The sequence of pulsed power exemplar problems pursued during this GC LDRD, shown in order of increasing complexity from left-to-right.**

The next goal was the so-called "half-o-lute" study described in Section 4.3.2. This idealized, three-dimensional benchmark test problem includes half of the incoming transmission lines in a typical pulsed power accelerator, simplified in a wedge representation of the multi-level post-hole convolute. This problem does require a connection with a fairly-sophisticated transmission-line circuit model. Results accomplished in EMPIRE on CPU-based parallel machines are described.

The final goal was to assess the state-of-the-art multi-physics codes on a three-dimensional wedge representation of a typical engineering-quality design simulation (Power Flow 18a), and to investigate scalability and performance. This problem also requires a connection with a realistic Z-Machine transmission-line circuit model as well as several key boundary condition models. In Section 4.3.3 we provide a comparison of the simulations completed for this challenge problem in EMPIRE along-side CHICAGO results.



**Figure 4-10. Planar MITL geometry; the yellow region is the simulation domain.**

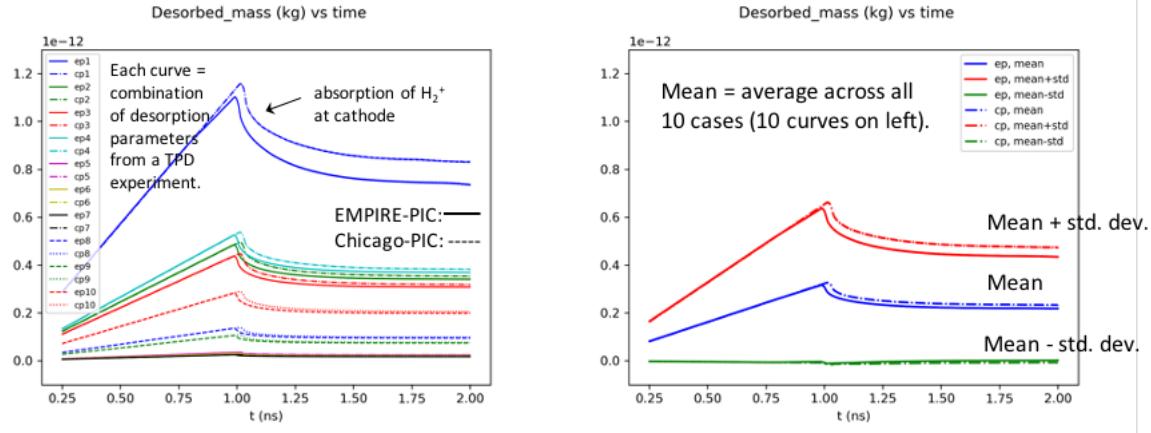
#### 4.3.1. Planar MITL PIC UQ Study

The 2D planar MITL is the simplest of the geometries used for cross-code verification of power flow modeling capabilities being implemented into EMPIRE-PIC [9, 16]. This verification effort is a crucial step in the eventual hybridization of the EMPIRE-PIC and EMPIRE-FLUID codes in a multi-scale model designed to capture the wide range of plasma densities and time scales inherent in power flow modeling. The verification effort was automated within the Code Comparison Infrastructure, and consisted of a series of comparisons between EMPIRE-PIC and CHICAGO [130, 136].

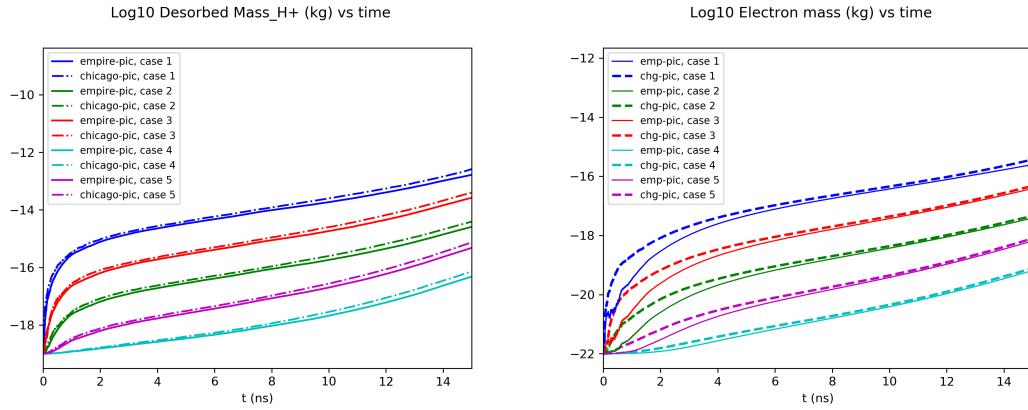
We also performed uncertainty quantification (UQ) studies of the sensitivity of desorption, ionization, and power flow losses to uncertainty in experimentally determined desorption parameters. These verification and UQ studies are both central to developing predictive capability for power flow modeling. Performing these studies in a simplified 2D planar geometry lays the groundwork for extending these analyses to higher-fidelity 3D geometries. A high-level summary of the work is provided below with additional details provided elsewhere [44]. The planar MITL simulation domain is shown in Figure 4-10.

The first study was a UQ analysis of thermal desorption parameters obtained from best-fits to experimental data from a series of ten temperature programmed desorption (TPD) experiments. Thermal desorption was limited to the horizontal surface of the anode. To enable a direct comparison of the desorption models, ions were desorbed directly from the surface, using the desorption parameters for the neutral species. The results in Figures 4-11 are for desorption of  $H_2^+$ . Each case is a set of best-fit desorption parameters from a TPD experiment. EMPIRE-PIC and CHICAGO agree to within 10% on the desorbed mass. Differences between the codes are likely due to differences in the timing of the downward acceleration of the desorption front from the anode and its absorption by the cathode. Both codes agree on substantial variation in the desorbed mass across the ten test cases, giving rise to a standard deviation which is comparable to the mean, where some cases correspond to negligible desorption losses and others to substantial losses (i.e., a substantial fraction of the steady-state electrode current). Similar results were obtained for desorption of  $H^+$  and  $O^+$  ions representing water fragmentation.

A UQ analysis was then performed for desorption and ionization of neutral H from the anode and



**Figure 4-11. Mass of desorbed  $H_2^+$  vs time, comparing EMPIRE-PIC and CHICAGO for each of ten TPD experiments (left), along with the mean and standard deviation of the desorbed mass (right).**

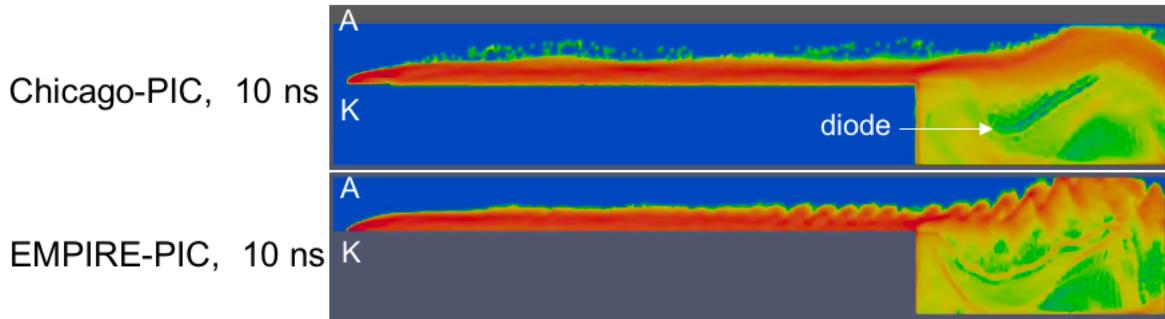


**Figure 4-12. Time dependence of  $H^+$  (left) and electron mass (right) for each of five cases (combinations of desorption parameters), where the solid curves are EMPIRE-PIC results, and the dashed curves are CHICAGO results.**

cathode, sweeping across desorption parameters historically found in several pulsed power references such as [31, 71, 103]. Figures 4-12 show the time-dependence of  $H^+$  ion and electron mass produced from ionization of desorbed H. The two codes differ by  $\sim 33\%$  in  $H^+$  mass and  $\sim 15 - 20\%$  in electron mass, which is greater than the fractional difference of  $\sim 10\%$  of the previous study above. This is most likely due to differences between the two codes in the ionization modeling, the influence of which could be assessed with a solution verification study to see whether the agreement improves upon mesh and time step refinement.

As with the previous UQ study, both codes agree on the high sensitivity of desorbed mass to the desorption parameters, though these losses correspond, in both codes, to a vanishingly small fraction of the driven current. The total current, and therefore the resulting current losses, are noisier in CHICAGO than EMPIRE-PIC; this is under further investigation.

An EMPIRE-CHICAGO PIC verification effort was also performed for modeling of SCL emission from the cathode, with the effect of desorption removed. EMPIRE-PIC and CHICAGO exhibit qualitative agreement in the evolution of electron emission over a broad range of densities. The



**Figure 4-13. Comparison of SCL electron densities in the 2D planar MITL between EMPIRE-PIC and CHICAGO.**

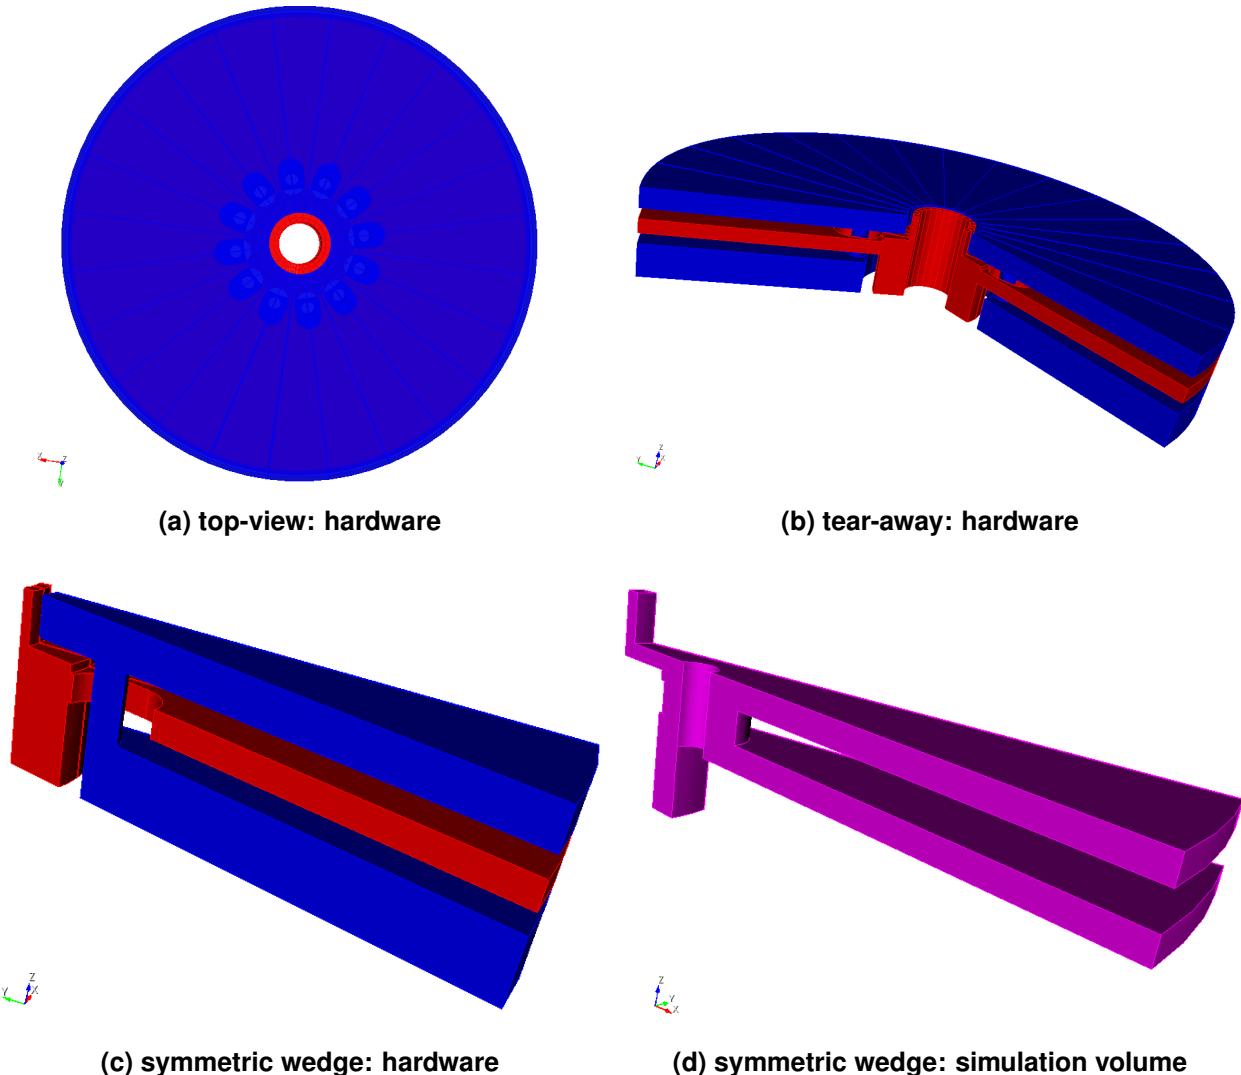
level of agreement remains to be quantified. Noise was observed, starting at about 8 ns, in the time histories of quantities such as electron mass, kinetic energy, current, and voltage. This noise is likely due in part to turbulent behavior in the diode, seen in both codes. In EMPIRE-PIC the noise is somewhat greater, most likely due to oscillations in the electron sheath starting at  $\sim 6$  ns; see Figure 4-13 for one example snap-shot acquired at 10 ns. It remains to be determined whether these oscillations are physical or numerical. Their presence in EMPIRE-PIC and not CHICAGO is likely due to the broader range of high frequencies which are damped in the CHICAGO time advance compared to EMPIRE-PIC.

The planar MITL verification effort has identified both areas of agreement and disagreement between EMPIRE-PIC and CHICAGO, outlined above, in the modeling of thermal desorption, ionization, and SCL emission; all of these models are needed for a predictive power flow model. Initial work has been completed, though more work remains, on quantifying and understanding the discrepancies in terms of model differences between the two codes. The UQ studies have demonstrated a working automated pipeline for analyzing the sensitivity of power flow relevant loss metrics to uncertainties in experimentally-derived input.

#### 4.3.2. *Half-O-Lute Study*

The next level of complexity in the pulsed power modeling pipeline is a half-way point between the the simplicity of a single MITL (Section 4.3.1) and the full complexity of an engineering-quality geometry (Section 4.3.3). This exemplar problem is called a *half-o-lute*, which is a portmanteau named on the basis that it models roughly “half of the convolute” geometry in the Z-Machine. The *convolute* is a combination region wherein current paths from the individual MITLs join together into a single *inner MITL* that delivers the total current to the load. In the Z-Machine, inclined MITLs (A,B,C, and D) from the outer region join through a double post-hole convolute (DPHC) (Figure 4-21) device consisting of 12 double anode posts arranged azimuthally around the machine center; these electrically connect four anode electrodes belonging to the four outer MITLs through holes in the cathodes. The cathode current path indents around the posts and joins together on the down-stream side. The half-o-lute idealizes the system (Figure 4-14) by simplifying the geometry and reducing the number of outer transmission lines such that a single post-hole convolute (PHC) joins two radial transmission lines (A,B) to a

vertical coaxial MITL, and takes the input of the combined currents from C+D lines in the full geometry. Using the same driving voltages from the BERTHA circuit model (Figure 4-15) as in a full geometry, the idealized pulsed power system contains all the physics of the Z-Machine, but without the geometric complexity of the full device. As a result, the half-o-lute presents an opportunity to interrogate device performance (e.g., full-physics models, current loss mechanisms, etc.) in a reduced problem space with modest computational cost as compared to the real device.



**Figure 4-14. (Half-o-lute)** The electrode configurations (blue = anodes, red = cathodes) and simulation volume are shown; In (a), a transparent view reveals the 12 post-holes that combine into a single-post hole convolute array; In (b), a tear-away view is used to show how the radial transmission lines (A,B) join with the vertical coaxial (C+D) transmission lines in the PHC to form a single inner MITL; The minimum symmetric volume is modeled as a 15° wedge whose hardware is shown in (c), and which corresponds to a simulation volume (space in-between) as in (d).

The code-comparison infrastructure (CCI) has been instrumental to the success of developing the EMPIRE half-o-lute model by providing a platform for cross-code benchmarking against the

mature half-o-lute CHICAGO model, which has formed the basis of several previous studies [103, 71]. The CHICAGO half-o-lute model was generously shared with the EMPIRE team courtesy of Voss Scientific, who also provided up-to-date CHICAGO executables of the test problem. This allowed what became the main use of the CCI for this problem in the initial stages: (1) sandbox investigations for troubleshooting inconsistencies that arose in the EMPIRE (2) detailed simulations that had not previously appeared in the more simplified 2D planar MITL geometry, and (3) to vet individual capabilities that had not found previous opportunity to be verified such as the 2-to-1 branching feature of 1D transmission lines and a proof-of-concept for translating the lumped BERTHA parameters to an equivalent distributional transmission line parameters for EMPIRE.

Verifying the emerging capabilities in EMPIRE through a pipeline of increasing complexity for power flow exemplars (e.g., the planar MITL handing off capabilities to the half-o-lute modeling effort when sufficiently matured) played an indispensable role in identifying and resolving initial issues in the code, as well as revising modeling decisions for power flow systems. However, there were corner cases that were only observed when deploying these capabilities on a 3D geometry with more complex transmission line physics. Therefore, this modeling effort played a key role in troubleshooting new issues that appeared in 3D, before making another hand-off of the matured code features for use in the significantly more complex engineering-quality models<sup>2</sup> (Section 4.3.3). In particular, the EM wave propagation and circuit models required frequent interactions between the analyst and development teams. In some cases, conversations concerning the underlying methods were explained, and those assumptions better appreciated by the analyst led to improved modeling decisions that fixed an apparent issue. In other cases, actual code bugs were discovered that were resolved throughout this iteration process.

By the end of this LDRD project, we completed a set of baseline results for the half-o-lute exemplar that includes the transmission line circuit models, EM wave physics, and electron emission models. Additionally, a general VVTTest structure has been installed on the GC LDRD GitLab that allows toggling between tests (circuits vs. EM vs. electron emission) and various levels of refinement therein with automated post-processing that compares CHICAGO with EMPIRE results. We describe a selection of the major results from these tests below. Our experience in iterating on this model suggests the differences observed (e.g., currents near the load) are probably attributed to the resolution in modeling decisions. For example, grid spacing was found to be crucial to calculate the fields correctly. The demand for resolution is further increased in regions such as the final vertical coaxial MITL leading up to the load where a  $1/r$  magnetic field requires many cells to faithfully resolve the variation at these small values  $r \in [3.0, 3.6] \text{ cm}$ , perhaps as many as 20 tetrahedral elements in the case of EMPIRE. Otherwise, we find more noise in the particle emission case for EMPIRE-PIC as compared to CHICAGO, but observe the results are clearly consistent with each other. This being stated, we find the baselines here between EMPIRE and CHICAGO simulations for “cold” (EM-only) versus “hot” (particle emission cases) encouraging and are well-positioned to iterate on this for further improvement.

In Figure 4-15, we show the 1D BERTHA circuits which drive the 3D electromagnetic half-o-lute

---

<sup>2</sup>e.g. for a case where cathodes are emitters of space-charge limited electrons due to high field stresses: a half-o-lute mesh of 19,988,575 elements was used (min = 250  $\mu\text{m}$ , max = 1000  $\mu\text{m}$  tetrahedral side lengths); Power Flow 18a required at least 44,024,030 elements (min: 375  $\mu\text{m}$ , max = 2000  $\mu\text{m}$  tetrahedral side lengths).

domain typical of all simulations in this Section. In the cold case, we find exceptional agreement in the voltages (Figures 4-16(a,b)), the electric fields (Figure 4-17), and the currents leading up to the convolute (Figures 4-16(c,d)). As mentioned, the differences of the load currents appear to be resolvable with a higher resolution for the grid and is currently being revised.

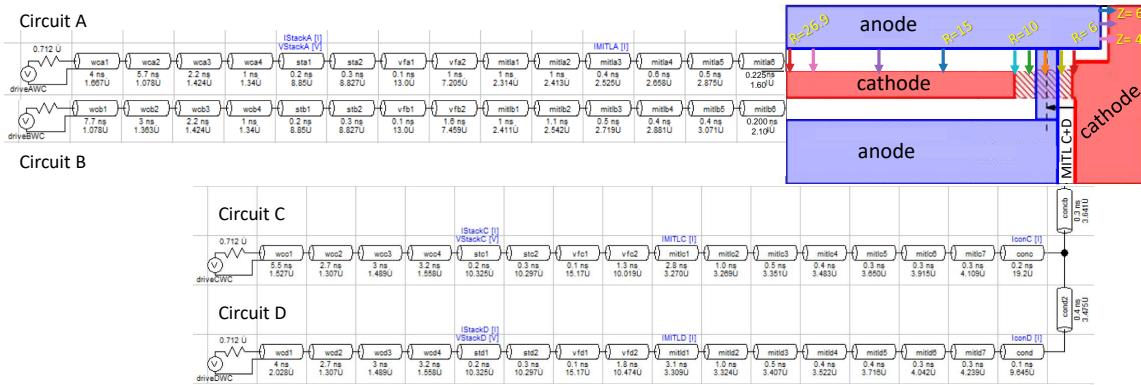
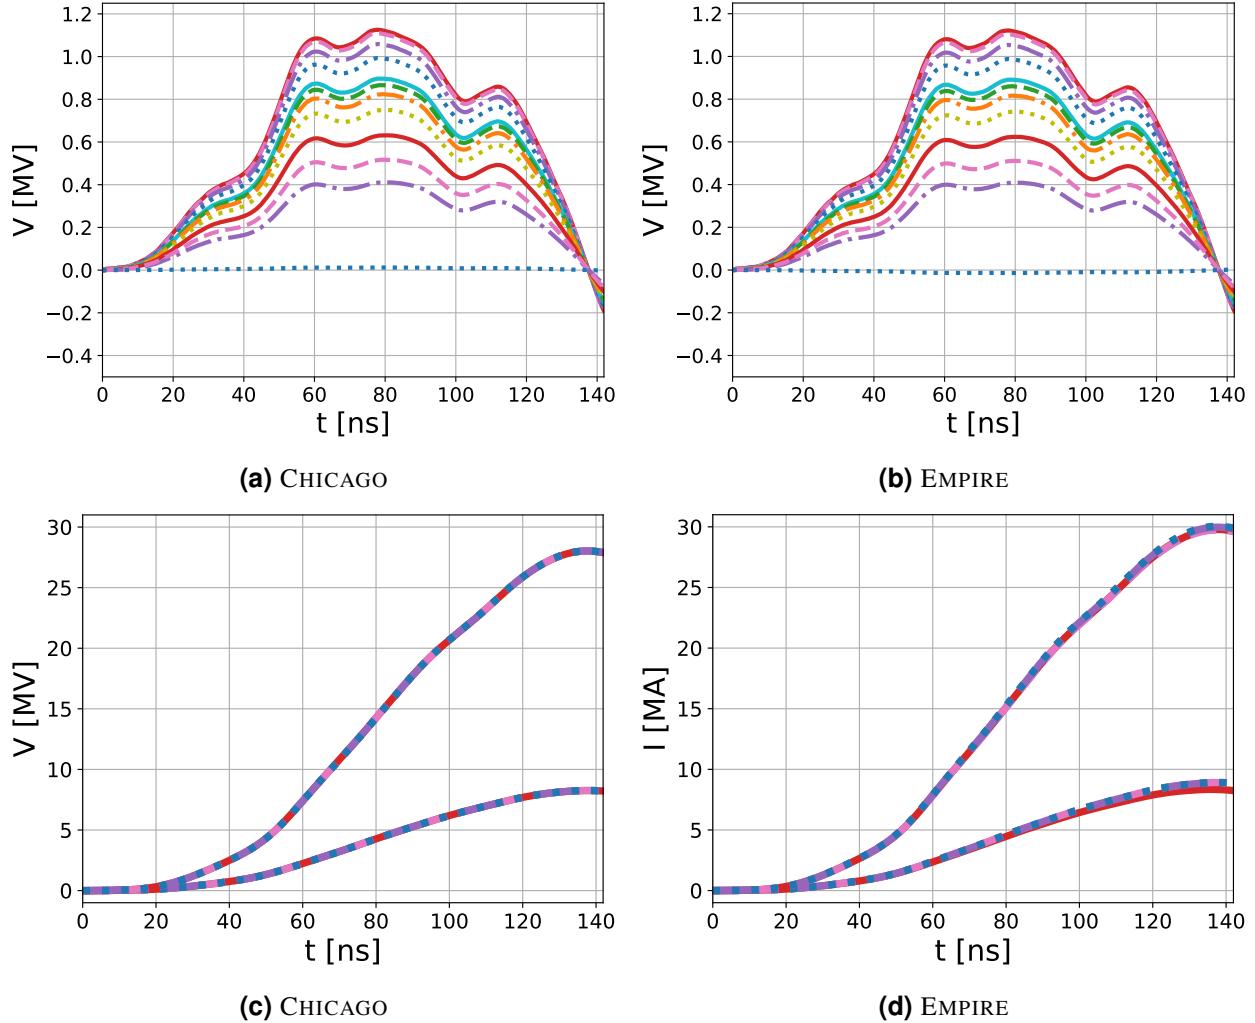


Figure 4-15. Half-o-lute diagram showing how the 1D transmission lines couple to the input side sets of the half-o-lute domain (illustration adapted [103]); the colored arrows on the top level are referenced in subsequent figures; all dimensions are in cm.

Next, we present results from initial “hot” simulations which include field-emitted electron emission through a space-charge limited emission model. Overall, we observe consistent solutions in the voltages (Figure 4-18), magnetic fields and similar electron flows (Figure 4-19) in the CHICAGO versus the EMPIRE-PIC simulations. More discussion and results are available in the internal report for this exemplar problem [118]. As mentioned, the noise present in the hot simulations for EMPIRE-PIC should be diminished with more judicious (higher resolution) decisions for the mesh and the emission model parameters (e.g. emitting more particles per face will reduce the noise). Improving the results is currently being pursued in continuing work.

Note the distinct vortex structure formed in the convolute (Figure 4-19) is also observed in the Power Flow 18a simulations (Figure 4-29). This signature is interesting to observe in both systems as it gives credence to the argument that the idealized convolute system does describe not only the same physics but similar features observed in the full engineering-quality geometry.

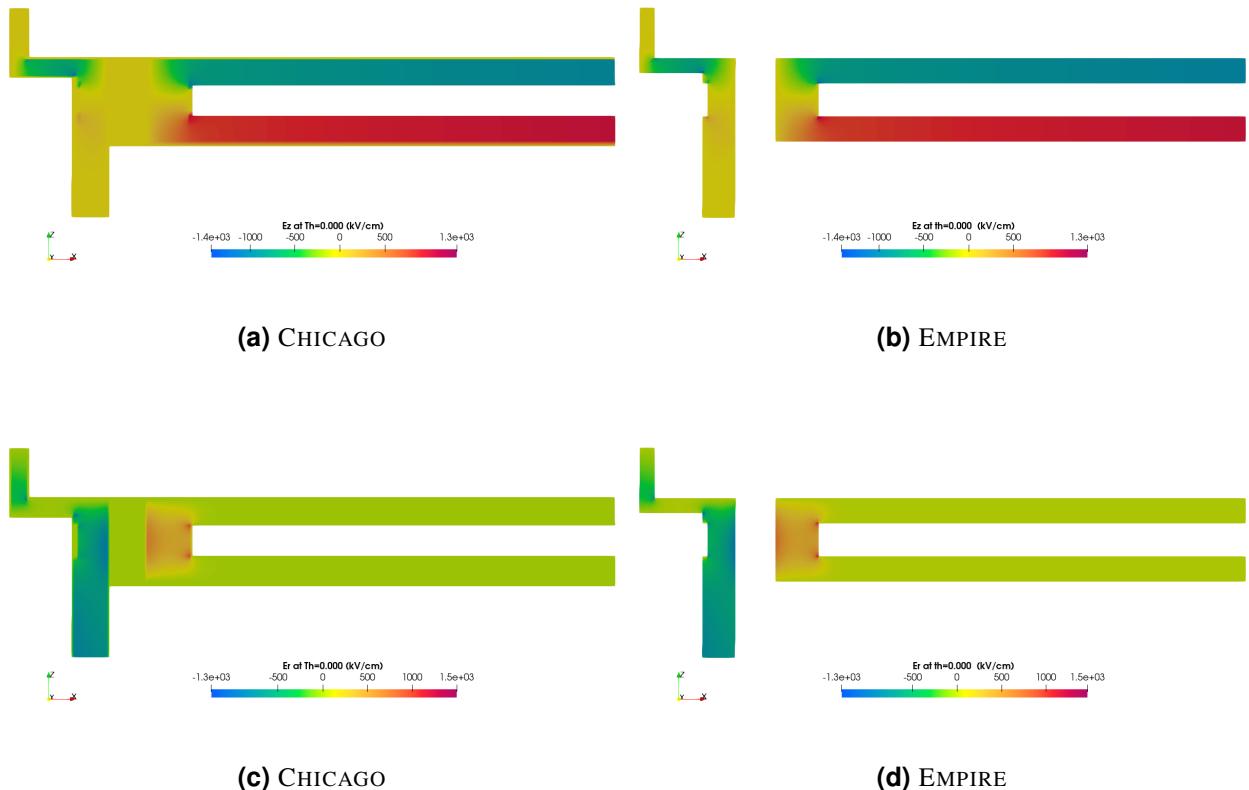
After iterating on the models to improve the above agreement further, our next steps include incorporating self-consistent surface heating by both Ohmic and particle impacts, along with thermal desorption physics of neutral water molecules off surfaces undergoing fragmentation such that electrode plasmas are modeled. This immediate next-step requires combining several emerging capabilities in EMPIRE-PIC and ensuring these models not only work individually with well-understood verification tests, but also that they work together on the half-o-lute model without unanticipated consequences. This LDRD project also enabled the training to handle complex modeling tasks in a systematic way, which incidentally mirrors the “agile” working structure that has been implemented in various teams throughout this project. For example, larger tasks are broken down into smaller verification problems; these problems are designed within the CCI and installed in the nightly testing suite when ready to protect against code regress. As the



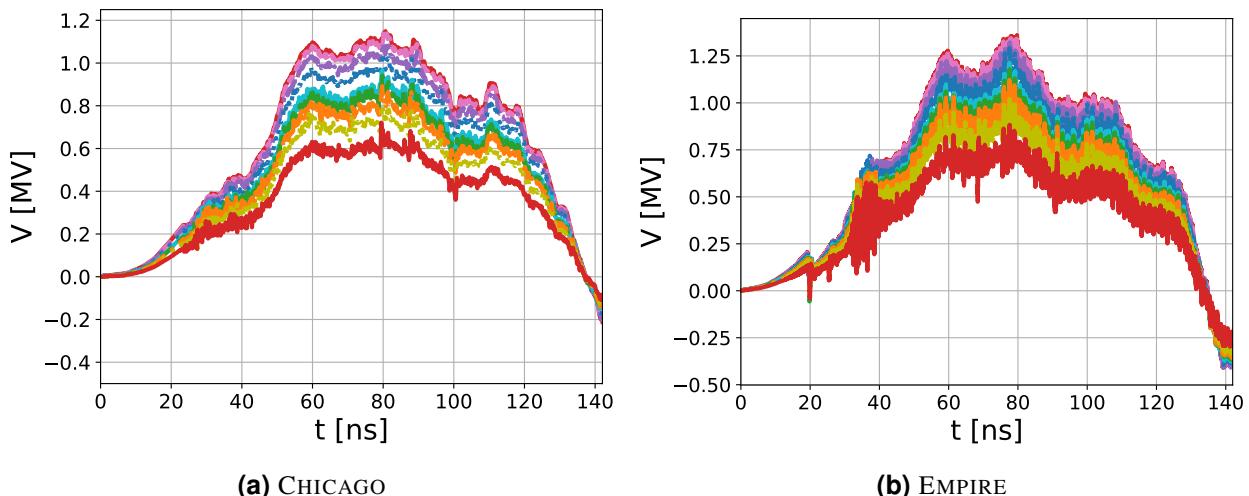
**Figure 4-16. (Half-o-lute: cold) cross-code comparisons of voltage, or  $\int_a^k E \cdot dl$ , diagnostics (panels (a,b)) and current, or  $\oint B \cdot dl$ , diagnostic (panels (c,d)) measured at the various locations marked in panel (a) along the top level are shown. The  $\int_a^k E \cdot dl$  paths are from anode ( $a$ ) to cathode ( $k$ ) as marked in diagram 4-15, and the  $\oint B \cdot dl$  paths are circles whose starting and ending points intersect the tail of those same marked arrows. The colors of the marked arrows in diagram 4-15 label the different curves in these plots. The results correspond to two overlapping sets of current measurements. The lower-valued currents are all measured in the outer MITL A whereas the higher currents are the combined currents from each MITL as measured in the inner MITL leading up to load. We have verified in test cases that the bottom set of EMPIRE curves will overlap when the grid resolution is increased.**

individual components fall into place, the full problem is pieced together and iterated until any remaining issues associated with the combination of modules is assessed.

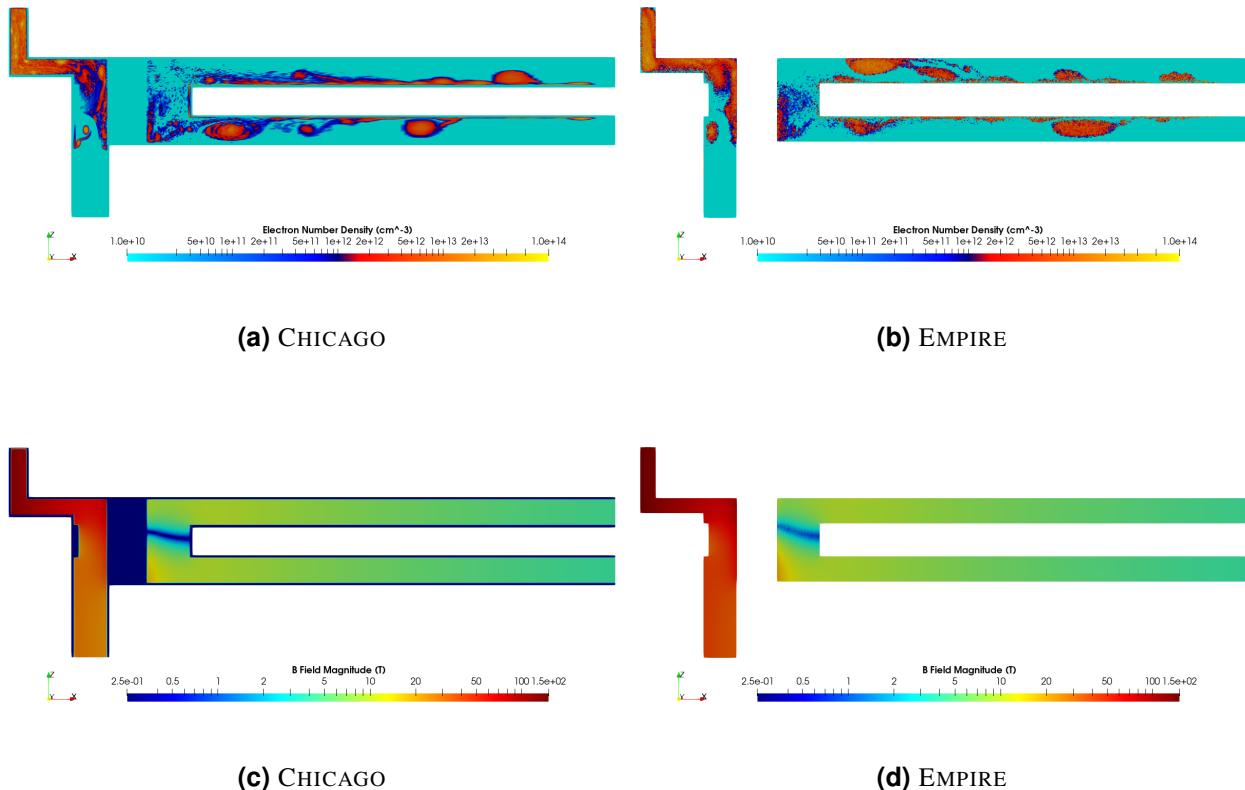
An additional investigation, currently in progress, extends the uncertainty quantification (UQ) work completed for the planar MITL 4.3.1 to the more complex half-o-lute test problem. This task intends to characterize the effects of measurement uncertainty, and the resulting desorption parameter fits derived from those temperature programmed desorption experiments, on power flow metrics such as current delivered to the load. The UQ infrastructure from the planar MITL



**Figure 4-17. (Half-o-lute: cold)** Electric field components  $E_r$  and  $E_z$  at  $t = 110$  ns are shown in the  $\theta = 0^{\text{circ}}$  plane. Note that the CHICAGO slices are drawn with the boundary cells and the anode post included (conductors), therefore these cells are colored at  $E_r = 0$  (or  $E_z = 0$ ).  $t = 0$  corresponds to 18 ns into the pulse.



**Figure 4-18. (Half-o-lute: hot)** cross-code comparisons of voltages ( $\int_a^k E \cdot dl$ ), diagnostics measured at the marked arrows in diagram 4-15.



**Figure 4-19. (Half-o-lute) cross-code comparison of electron number densities at  $t = 90$  ns and the magnetic field magnitudes in the  $\theta = 0^\circ$  plane. Note that the CHICAGO slice is drawn with the boundary cells and the anode post included (conductors), therefore these cells are colored at  $n_e = 0$  and  $B = 0$ , respectively.  $t = 0$  corresponds to 18 ns into the pulse.**

problem 4.3.1 has been incorporated into the half-o-lute, and efforts are currently underway to improve the computational efficiency of EMPIRE-PIC for UQ studies incorporating desorption and ionization physics with realistic BERTHA circuit drives at the MITL A, MITL B, and MITL C+D ports. This UQ analysis will then be compared between EMPIRE-PIC and CHICAGO.

#### 4.3.3. Power Flow 18a

The highest level of complexity in the pipeline of pulsed power exemplar problems (Figure 4-9) is an actual engineering-quality Z-Machine power flow geometry. In this section, we describe results from EMPIRE-PIC simulations of the Power Flow 18a design, an example shot series conducted at the Z-Machine in March 2018. A short description of large pulsed power accelerator topologies, such as the Z-Machine, is provided below.

The 33-meter diameter Z-Machine delivers megajoules of combined electrical energy stored in 36 individual pulsed power modules (Marx generators) at its outer perimeter through a set of radially-convergent parallel transmission lines towards a common load at the machine

center [129]. The 2,160 capacitors ( $2.6 \mu\text{F}$  each) housed in these Marx generators discharge in series with characteristic rise time on the order of  $\mu\text{s}$ . Several stages are required to compress and shape the pulse so that a  $\sim 100 \text{ ns}$  rise time electrical pulse arrives at the target location [105, 77]. At a high level (Figure 4-20), the accelerator launches a pulse from each module in the azimuthal array of oil-insulated Marx generators towards 36 corresponding intermediate storage capacitors. This second stage is discharged by controlled laser-triggered switches programmed according to pulse shaping demands of the experiment. Subsequently, the individual pulses enter the water-insulated region consisting of pulse-forming lines (PFLs) and output transmission lines, mediated by a series of self-breaking water switches to achieve additional pulse sharpening. The pulses combine at the water convolute, and are subsequently distributed among four *magnetically-insulated* transmission lines. The pulses combine in a *post-hole convolute* device to form a single transmission line that delivers the total machine current (up to 26 MA) to an experimental load. In these final vacuum regions, local temperatures incur significant ramp rates (10 K/ns), thermal desorption of neutrals is significant, magnetic fields reach up to 300 T, electric fields reach MV/m levels, and electrons are strongly field-emitted from the electrodes.

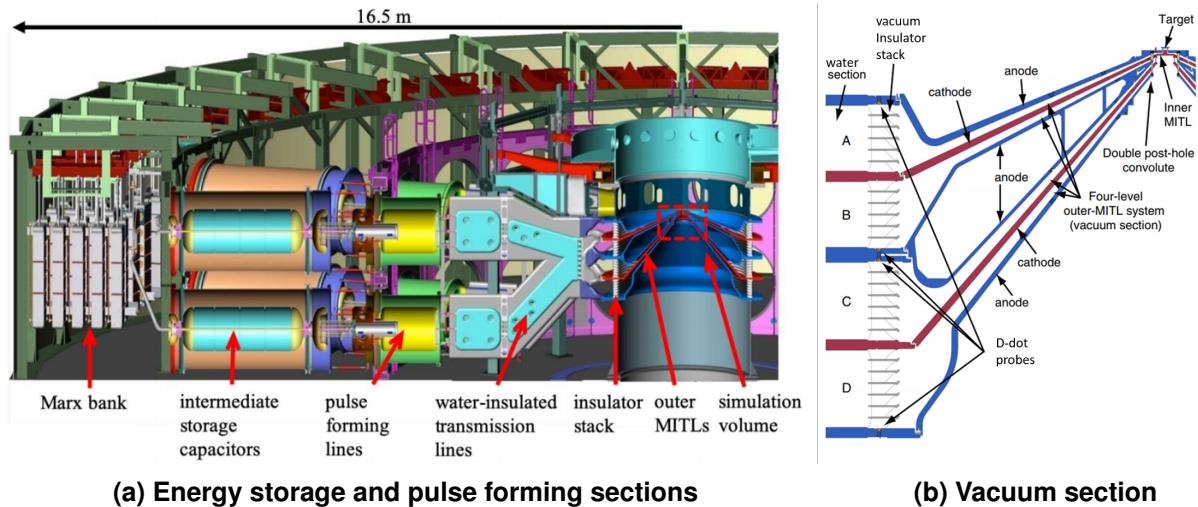


Figure 4-20. High-level cross-section of the Z-Machine [134]

Electrical power flow can be associated with a very large scales in time ( $\mu\text{s}$  discharge lengths versus THz electron cyclotron frequencies), space (transmission line lengths (m) versus electron gyroradius ( $\mu\text{m}$ )), and velocity (thermal speeds of outgassed neutrals versus the light-speed EM wave propagation in each medium), as well as highly coupled physics. This complex system presents an opportunity to push the boundaries of next-generation code capabilities, and also a unique possibility to validate simulations on a well-characterized highly coupled multi-physics system against real measurements. As mentioned in Section 3.1, this LDRD project enabled several specific capabilities needed for pulsed power applications to be implemented as part of EMPIRE-PIC code development.

The configuration in Figure 4-20 is modeled by coupling low-cost 1D transmission lines (which can accurately describe well-behaved regions of the full Z-Machine circuit) with a detailed 3D electromagnetic PIC model employed only for the final  $\sim 25 \text{ cm}$  of vacuum transmission lines; this area is where crucial 3D effects such as field-emitted electrons, relativistic dynamics, and

their self-consistent interplay with the electromagnetic pulse must be modeled in full detail. The transmission line models are translated from *reduced* BERTHA circuit developed by Hutsel et al. [52]; this implementation is detailed in a companion report [119]. The 1D transmission lines employed range from 4.5 to 5.4 meters and physically extends only to the water convolute. The lines are driven by an open circuit waveform derived from the full model to retain the forward voltage information in these reduced transmission lines. Note that neglecting the circuit beyond water-insulated lines is justified, as everything upstream from the water convolute is essentially decoupled from the load during the time scales of interest. This reduced BERTHA circuit is consistent with the full end-to-end model developed in [52], but significantly reduces the simulation time required<sup>3</sup> without consequence.

The 3D geometry and its corresponding unstructured mesh are produced using the in-house meshing software, CUBIT, and parameterized through APREPRO, (e.g. different mesh settings); simple front-end interfacing is preserved by exposing useful input parameters (e.g. cathode cell spacing) to the user through a VVTest companion file. All tests are stored in the GitLab repository and are push-button reproducible courtesy of VVTest.

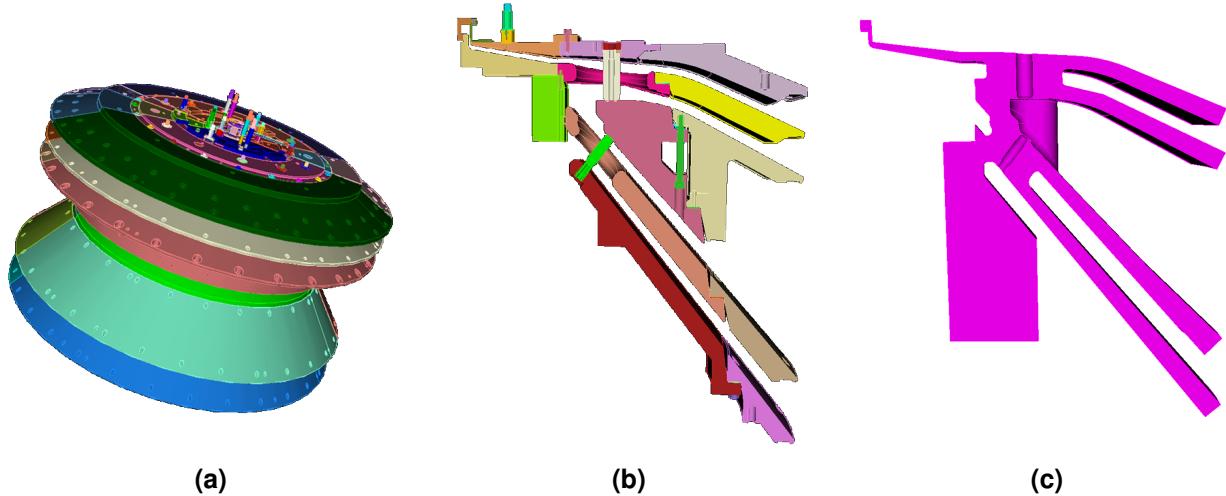
We describe a selection of results from high-fidelity simulations using EMPIRE for several test cases. In the following, we re-center the time axis so that  $t = 0$  corresponds to 18 ns, which corresponds to 0.225 ns before the first traveling wave launched from 1D transmission lines (TL) reaches the 1D TL/3D EM domain boundary<sup>4</sup>. Cross-code comparisons with CHICAGO were helpful to establish baseline agreement and to characterize the electromagnetics of the system before including additional physics. Measurements have been compared to simulations when possible (e.g. measurements provide fair comparisons only when the simulation models contain sufficient physics to approximate the full device physics).

First, we establish the comparison between CHICAGO and EMPIRE by simulating an electromagnetic only (“cold”) case. Useful comparisons include line integrals  $\int \mathbf{E} \cdot d\mathbf{l}$  across A-K gaps and  $\oint \mathbf{B} \cdot d\mathbf{l}$  around either electrode, which serve as proxies for the voltages and electrode surface currents, respectively. Electromagnetic field distributions in constant planes (e.g.  $\theta = 0^\circ$ ) are also useful as a visual check to ensure qualitative agreement. E and B line integral comparisons at the port boundaries, and also at the start of the inner MITL ( $R = 11$  cm measured from the machine center), are shown in Figures 4-22 and 4-23. For this cold case, the relevant length scales are constrained only by the need to track high frequency content in EM waves propagating and developing in the 3D domain. We obtain reasonable agreement with CHICAGO with a coarse EMPIRE mesh of just 953,841 elements (cells with sidelengths as large as 2.5 mm in the outer regions). Using an implicit Runge-Kutta time integrator at 7.5 ps per step (speed-of-light CFL: min = 0.46, average = 0.9, maximum = 11.85), we progress forward to 160 ns in less than 5.5 hours on 1 node / 4 GPUs. Additional refinement did not produce noticeable improvement in these metrics.

---

<sup>3</sup>consider the cost savings obtained by omitting the  $\sim 1.5 \mu\text{s}$  discharge time characteristic of the Marx generators versus simulating the  $\sim 150$  ns electrical pulse characteristic to the vacuum power flow section.

<sup>4</sup>Significant cost savings are afforded [79] by delaying the turn-on of 3D EM calculations until the a sufficient threshold is exceeded (e.g. when the voltage at the 1D TL/3D EM boundary is greater than 1e-16). CHICAGO has implemented this cost-saving measure and EMPIRE-PIC has also recently introduced this feature.

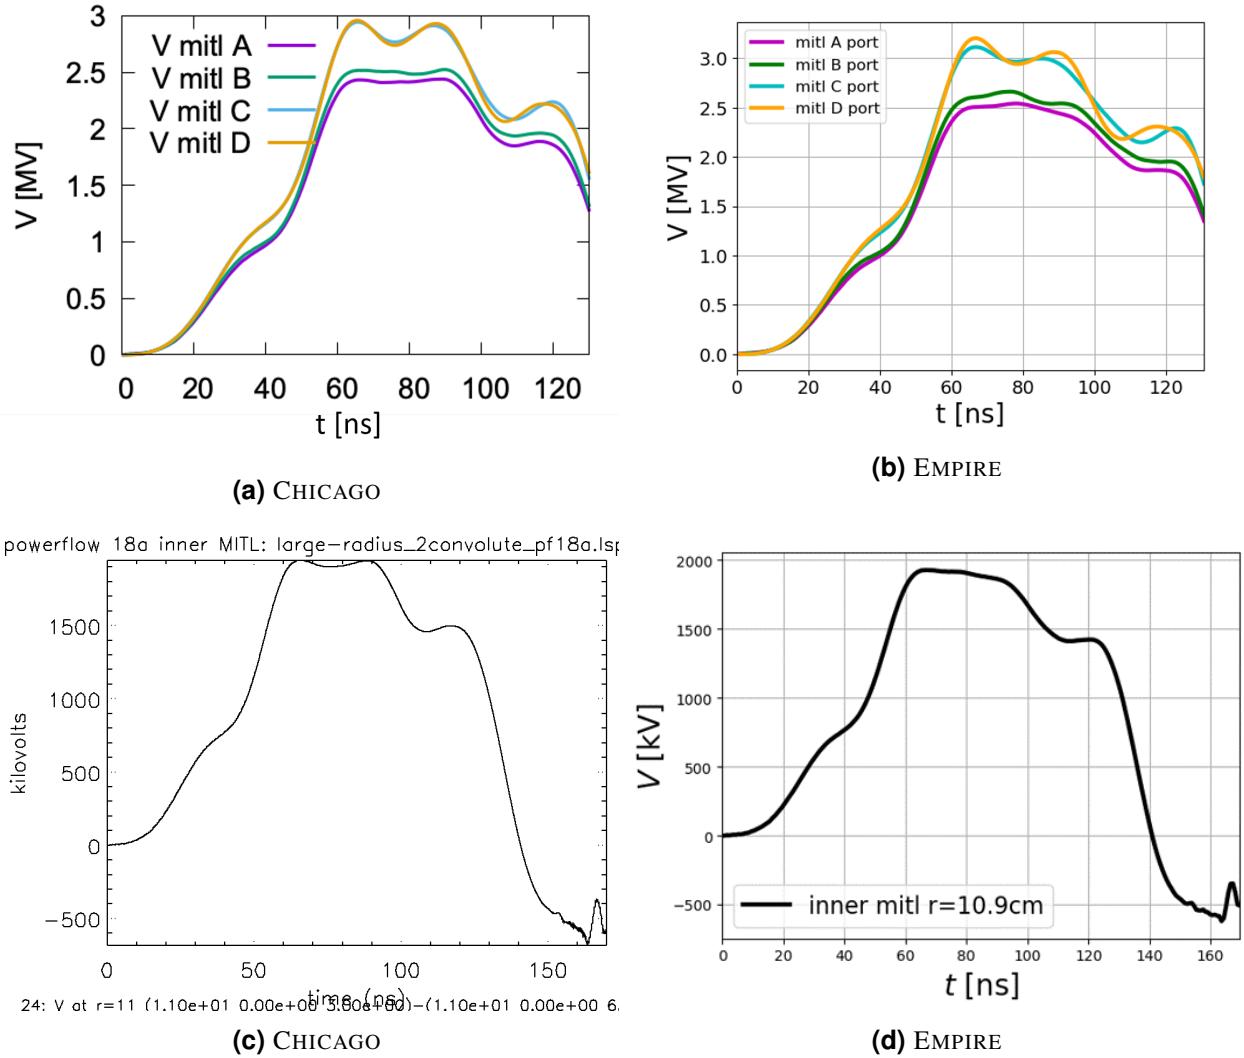


**Figure 4-21.** CUBIT was used to (a) import the full engineering CAD drawing of the hardware, (b) cut out and de-feature a 15 degree symmetric wedge of this hardware, and (c) extract the simulation volume (the “space in-between”). Note that we include a portion of the empty space below the lowest level (MITL D) to track or otherwise verify that no significant activity takes place below the transmission line system.

With this baseline established, we next address the more significant task of including particles (“hot” simulations). Field-emitted electrons have been introduced into EMPIRE-PIC simulations using a space-charge limited emission model. With particles, the mesh constraints are more significant as the details of thin magnetic insulation layers (cycloidal electron flows) in the outer MITLs must be visible from the grid perspective (Section 4.3.5.1). We refine our mesh so that tetrahedral cells near the cathode (emitting) surfaces have nominal sidelength of  $375 \mu\text{m}$  to resolve at least the coarse details of the insulated electron flow. We further take advantage of the fact that EMPIRE can accommodate unstructured meshes to reduce the overall cell count by gradually increasing the cell sizes away from emitting surfaces. Thus, the mesh is characterized with minimum  $375 \mu\text{m}$  cell widths at the cathode up to a maximum of  $1000 \mu\text{m}$  width near the non-emitting (anode) surfaces and within the interior volume. This 44M element completed a 160-ns simulation using 3.42 ps time steps in less than 4.85 hours on 64 nodes / 256 GPUs; this averaged approximately 15M particles throughout in the domain.

Even with this coarse resolution, we find encouraging agreement with CHICAGO particle simulations that include both electron emission and electrode plasma desorption. Since CHICAGO has one additional emission mechanism, we observe an (expected) lower load current in the CHICAGO results versus the EMPIRE-PIC (Figure 4-25); otherwise we observe good similarity in integrated diagnostics such as the conductor currents (e.g. Figure 4-26) and 3D EM fields (Figures 4-27).

Figure 4-28 shows electron density alongside a 2D slice of the B-field magnitude. In (a), the detailed physics of the electron flow can be observed, which reveals key features in a post-hole convolute system. Recall that electrical power flows from right-to-left in this diagram. The electron density reveals insulated regions in each of the outer transmission lines that are at least as thin as 375 microns (this layer is one cell width), though higher resolution simulations from

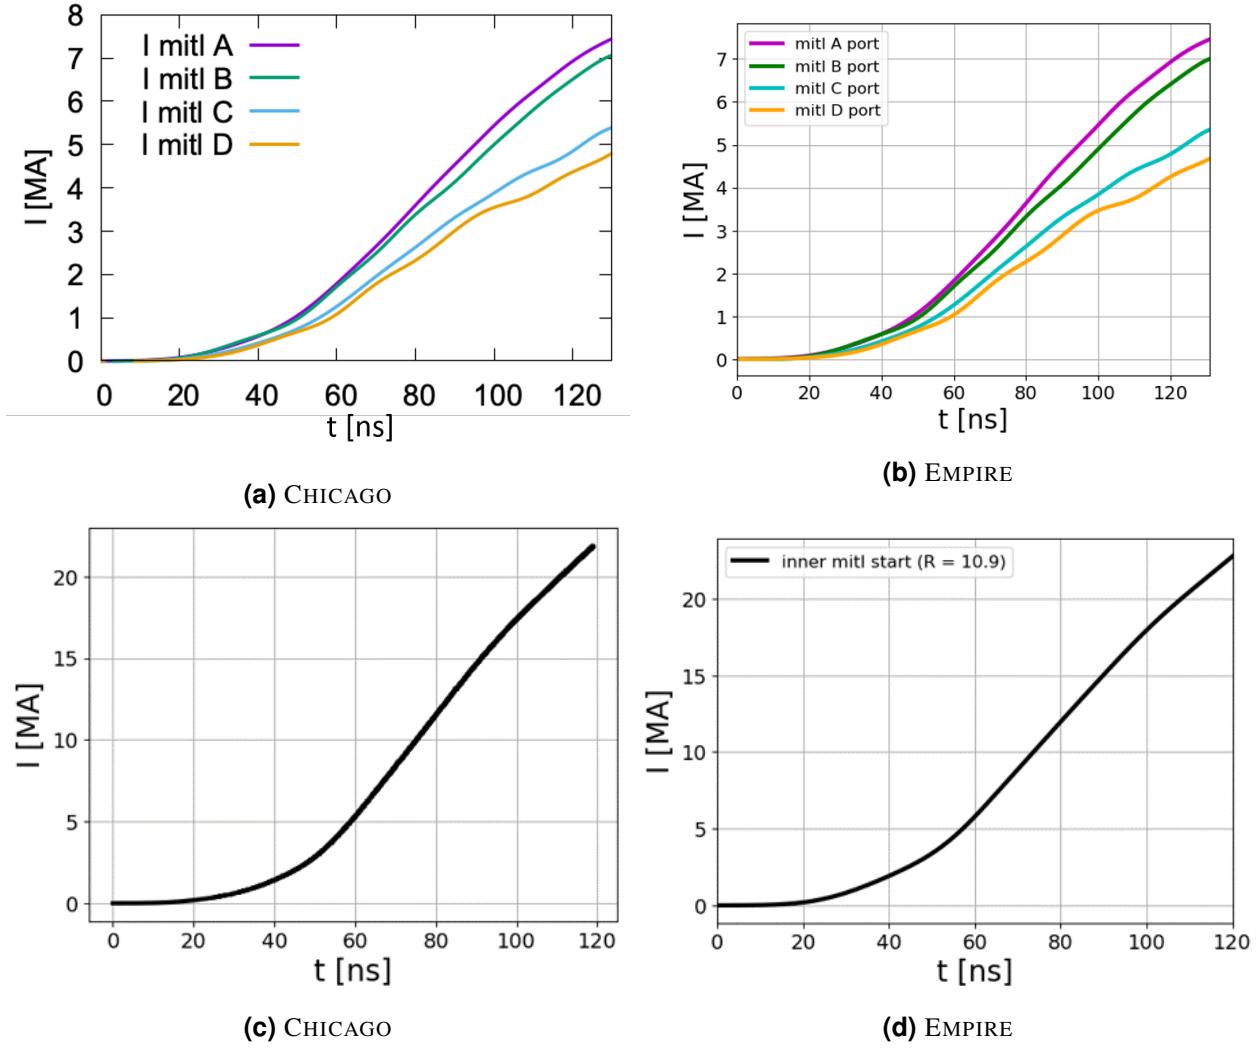


**Figure 4-22. (Power Flow 18a: cold) cross-code voltage comparisons at the outer MITL ports and at the start of the inner MITL ( $R \approx 11$  cm from machine center);  $t = 0$  corresponds to 18 ns into the pulse.**

CHICAGO indicate this layer can be a thin as 100  $\mu$ m when thermal desorption is included [6]. Further downstream, the emitted electrons drift into the convolute region, e.g. at the magnetic field nulls as seen in (b). These trajectories follow  $\mathbf{E} \times \mathbf{B}$  drifts shown in Figure 4-29, which trace out the coherent vortex structures that carry electrons from the cathode side before the convolute to the anode on the other side in inner MITL. After the convolute, a single inner MITL delivers the combined current towards the load. This non-neutral region generates vortex flow also observed in previous CHICAGO works [136, 6].

#### 4.3.4. Power Flow 18a on the Sierra HPC Resource

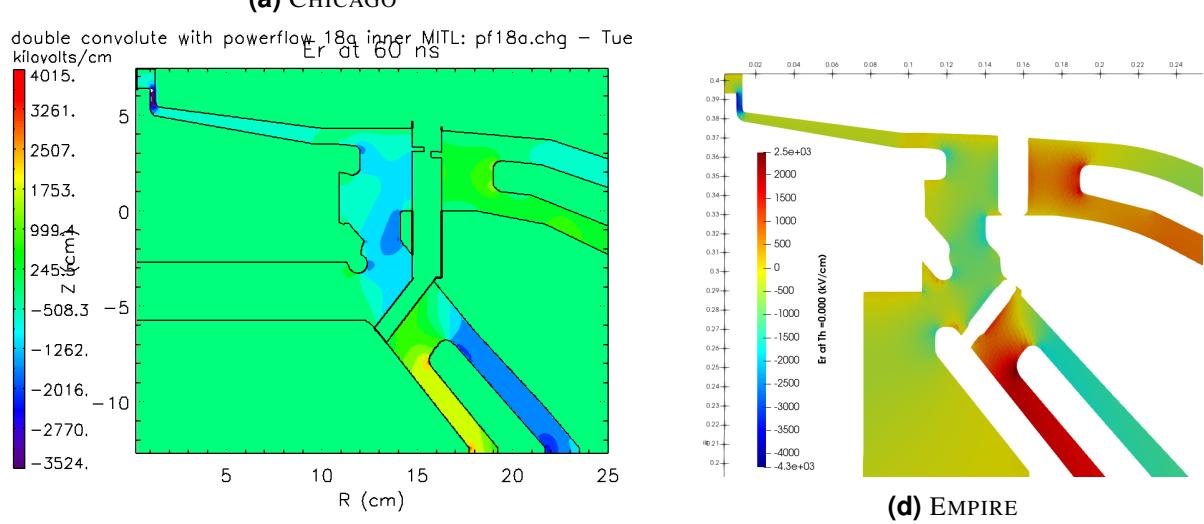
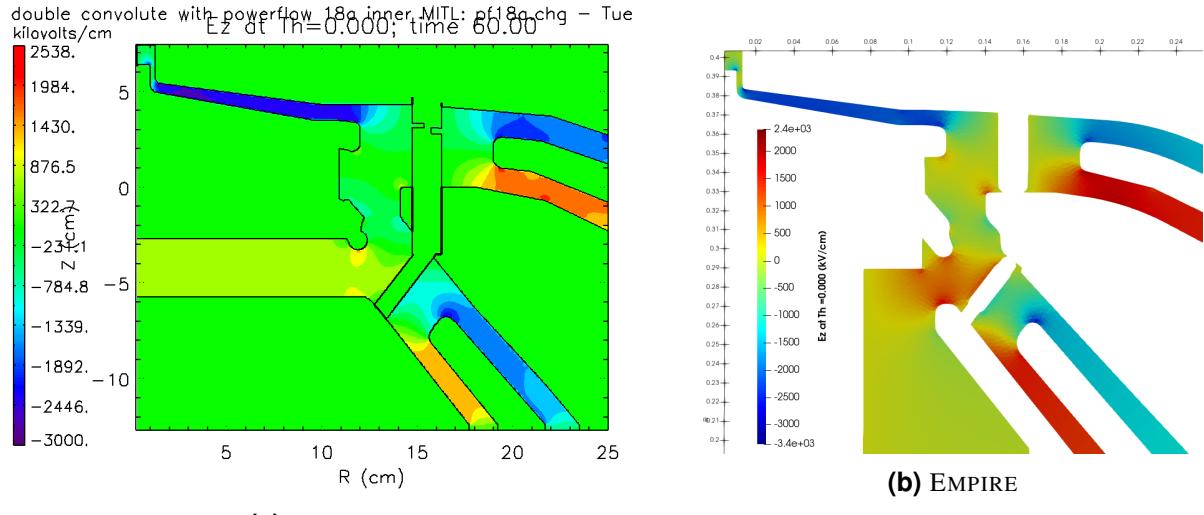
Previous studies have indicated that modeling the electrode plasmas formation processes, e.g. resulting from contaminant desorption, is critical for predicting integrated load current. The final



**Figure 4-23. (Power Flow 18a: cold) cross-code current comparisons at the outer MITL ports and at the start of the inner MITL ( $R \approx 11$  cm from machine center);  $t = 0$  corresponds to 18 ns into the pulse.**

step towards this *full physics* model from the baseline established in Section 4.3.3 requires not only several related physics capabilities working in tandem (surface heating, thermal desorption, and particle fragmentation models), but also additional code features to reduce the computational cost to capture these complex phenomena as much as possible. Invariably, significant computing resources may be required to execute these full physics simulations within the intended target of several days. Therefore, we aim to next generation platforms (NGPs) for additional performance gains by leveraging GPU-based parallel architectures in the tri-lab advanced technology systems (ATS). In this section, we describe several key lessons learned obtained from initial large-scale EMPIRE runs on LLNL’s Sierra (ATS-2) system under the ATCC9 campaign.

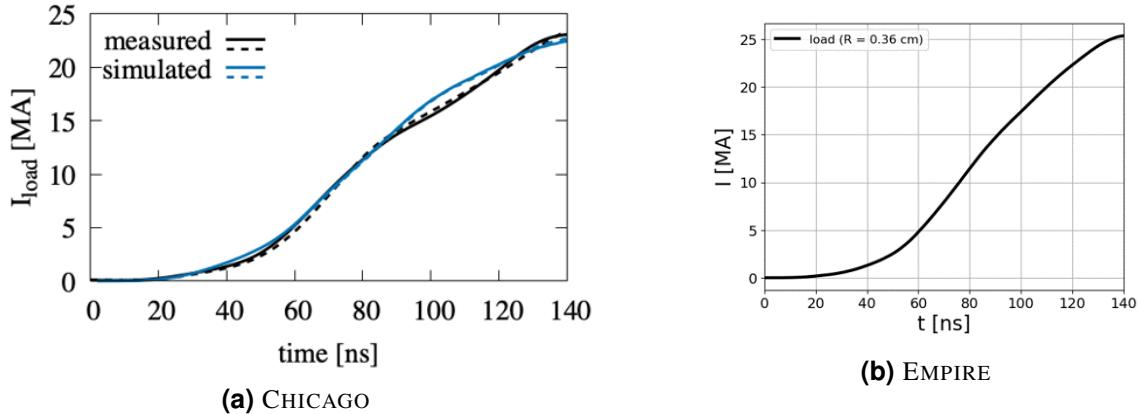
During this LDRD project, two performance tests were conducted on Sierra: (1) a full physics test built on top of the space-charge limited electron emission mesh used in Section 3.1.2.4 to assess the costs introduced by the new physics models, and (2) a scaling test to gauge code performance



(c) CHICAGO

**Figure 4-24. (Power Flow 18a: cold) Electric field components  $E_r$  and  $E_z$  at  $t = 60$  ns are shown in the  $\theta = 0^\circ$  plane;  $t = 0$  corresponds to 18 ns into the pulse.**

on a mesh meeting the expected minimum resolution requirements to sufficiently resolve all salient features in a full run. For the latter simulation, we carry forward the SCL electron emission case so that a clear comparison can be made between coarse versus fine resolutions. The full physics simulation builds on the space-charge limited emission Power Flow 18a model. Electrode surface heating is incorporated through a self-consistent heating model described in Section 3.1.2.4, which takes into account both Ohmic heating (magnetic diffusion) and incident particle fluxes. The most significant electrode contaminant, water, is emitted into the system from cathode and anode surfaces using the thermal desorption model defined in Section 3.1.2.4, with parameters chosen to match Temkin isotherm model documented in a recent CHICAGO study [6]. Ionization is accomplished by a staged fragmentation model producing  $\text{H}_2\text{O} \rightarrow 2\text{H}^+ + \text{O}^+ + 3\text{e}^-$  plasma through direct simulation Monte Carlo (DSMC) collisions defined with a mean free path of one cell width from the emitting surface. The simulation was performed on a coarser mesh

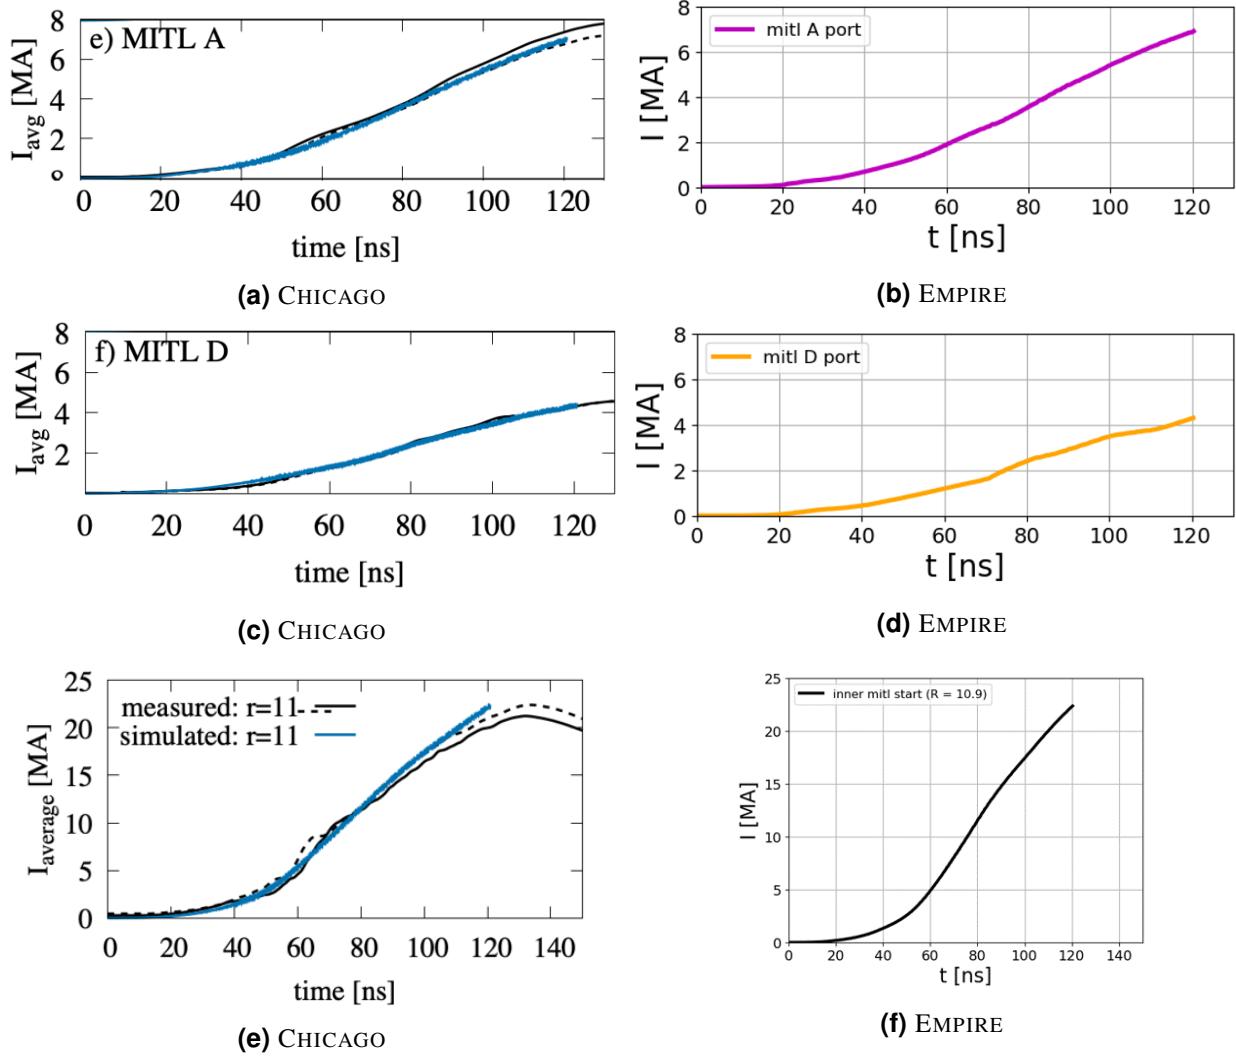


**Figure 4-25. (Power Flow 18a: hot) currents at the load,  $R \approx 0.36$  cm from the machine center; In (a), measured currents are plotted in black (solid and dotted) from two experiments in the Power Flow 18a series;  $t = 0$  corresponds to 18 ns into the pulse.**

(44M elements, mesh spacing: 375  $\mu$ m minimum, 1000  $\mu$ m maximum) than what is required to sufficiently resolve the sheaths at these higher densities, but gives information on increased costs associated with introducing these additional physics models compared to the baseline simulations described in Section 4.3.3. The simulation executed using 128 nodes / 512 GPUs of Sierra and reached 55% completion after a 12-hour period. A post-mortem analysis largely confirms the challenges expected in undertaking a simulation on this scale. First, the time taken per full step steadily increased throughout the simulation, slowing to an undesirable level of seconds per step (cf.  $\sim 10^5$  steps over a full simulation). This is primarily due to a growing inventory of particles throughout, which exceeded 6.5B in total by the point maximum walltime was reached. Particle management strategies involving merge algorithms to reduce their overall count and load balance schemes, to more evenly distribute the work across processors, are therefore needed to handle the increasing particle inventory resulting from multiple creation mechanisms including electron emission, thermal desorption, and fragmentation. Second, the fact that the walltime was reached before the end of simulation is an indication that a restart capability is needed to maximize utilization of shared resources. This is an emerging feature in EMPIRE. Finally, the fragmentation model introduces significant cost that can be streamlined by opting for programmed fragmentation events (auto-fragmentation) rather than statistically realized events a DSMC engine. This optimization is the subject of future work.

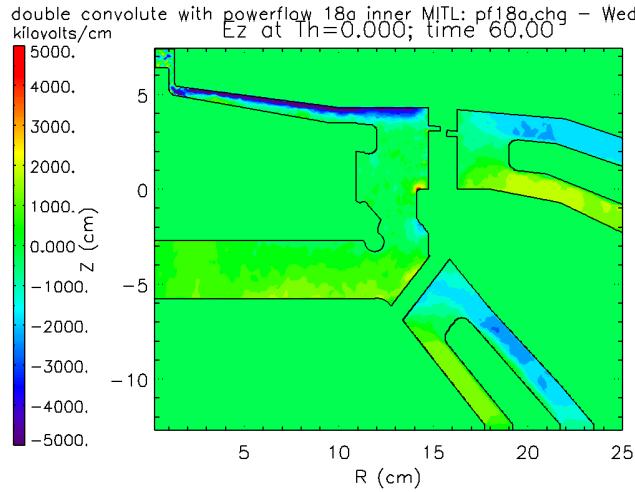
Based on prior works [136, 6], we anticipate resulting plasma densities on the order of  $10^{16} - 10^{17}$  cm $^{-3}$ , magnetic fields up to 300 T, and sheaths as thin as 100  $\mu$ m. Resolving the plasma and cyclotron frequencies at these conditions would restrict time stepping to 0.05 ps and 0.01 ps, respectively, requiring  $\sim 14$ M time steps to complete a 140 ns simulation. Implicit time stepping schemes to overstep the plasma and cyclotron frequencies are therefore requirements to make the full problem tractable. Two implicit algorithms are currently supported in EMPIRE (Section 3.1.2.6), including the magnetic implicit algorithm originally developed by Voss Scientific; this can relax the cyclotron condition by a factor of 9 [40, 136]. Implicit time stepping will also be leveraged in future simulations.

The  $\lesssim 100$   $\mu$ m resolution requirement on the mesh motivated a final Sierra run near the end of

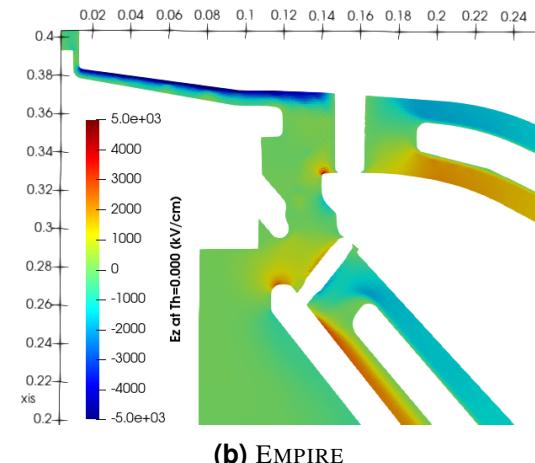


**Figure 4-26. (Power Flow 18a: hot) cross-code comparison of simulated currents at the outer MITL A, MITL D, and at the inner MITL start ( $R \approx 11$  cm from the machine center); In the left-panels, measurements from two experiments in the Power Flow 18a series are plotted in black (solid and dotted);  $t = 0$  corresponds to 18 ns into the pulse.**

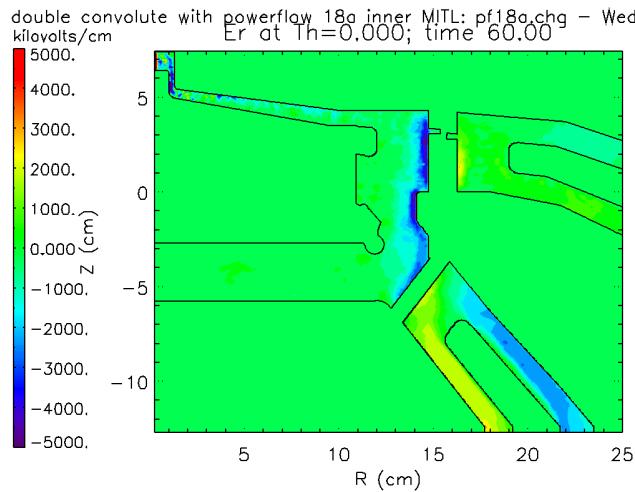
this LDRD project to gauge code performance on a mesh meeting these requirements. PERCEPT was used to refine the 44M element mesh under uniform refinement, producing a 2.8B element mesh meeting the minimum mesh size requirement of  $93.75 \mu\text{m} < 100 \mu\text{m}$  elements near the cathode emitters and a maximum of  $250 \mu\text{m}$  elements in the volume. We note that this uniform refinement slightly decreases the mesh quality, which challenges the solvers more significantly; a more strategic refinement can likely produce a mesh meeting the minimum requirements with an overall lower element count; this will be pursued for use in future simulations. This Sierra run reached only 30.5 ns into the 160 ns simulation due to platform-dependent I/O problems encountered at the time of our run. Unfortunately, this indicates that a significant particle inventory was not created throughout the simulation (yet), so the timing data is primarily associated with EM solves. However, the simulation utilized 2000 nodes / 8000 GPUs on Sierra and averaged a time per full step at 0.344 seconds. This is an encouraging timing



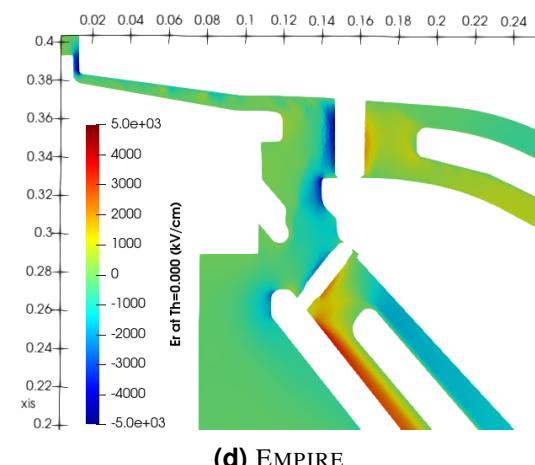
(a) CHICAGO



(b) EMPIRE



(c) CHICAGO

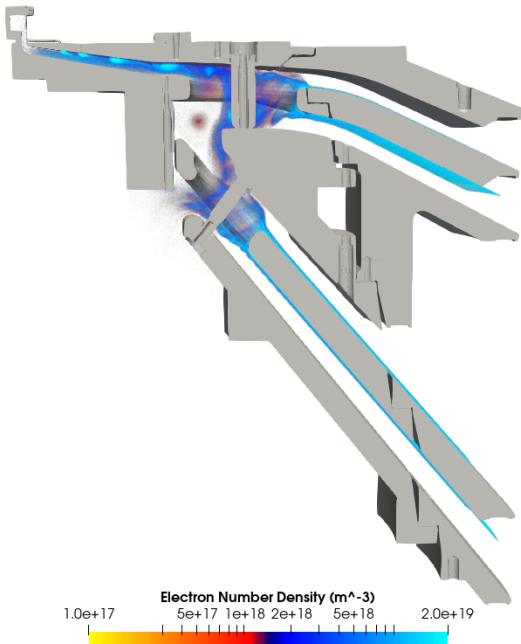


(d) EMPIRE

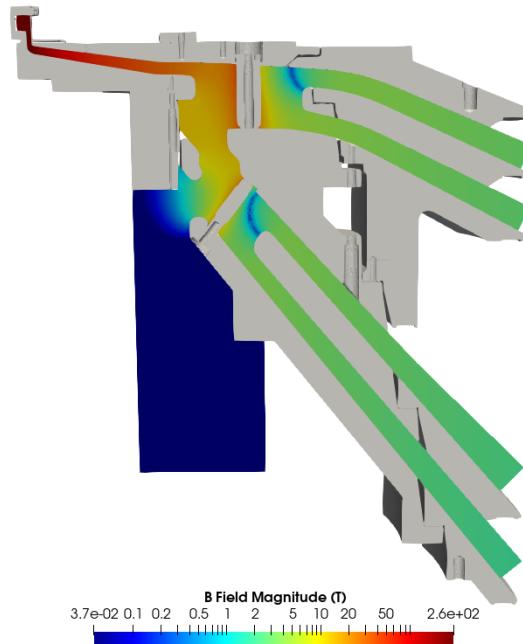
**Figure 4-27. (Power Flow 18a: hot) Electric field components  $E_z$  and  $E_r$  are shown at  $t = 60$  ns in the  $\theta = 0^\circ$  plane;  $t = 0$  corresponds to 18 ns into the pulse.**

given that the problem size was a tighter fit to the resources than optimal. Using a higher node count should reduce this timing by roughly a factor of 2, or 0.172 seconds per full step. From prior experience, the increased cost of including particles is roughly twice that of a cold simulation. If this timing can be maintained at-scale then a ballpark estimate for a full physics simulation over 140 ns at  $9 \times 0.01$  ps steps (1,555,555 steps) suggests turnaround times could be as short as  $2 \times 0.172 \frac{\text{seconds}}{\text{step}} \times 1,555,555 \text{ steps} = 6.19$  days. The factor of  $\approx 9$  over the 0.01 ps constraint from the electron cyclotron frequency is afforded from using the magnetic implicit time stepping algorithm.

Comparing this with the pre-LDRD level-of-performance of an  $\sim 150$  days indicates the possibility for a significant improvement for power flow design simulation capability. While this has not been demonstrated, if this figure for turnaround time can be achieved, then EMPIRE power flow simulations could be executed fast enough to affect the design cycle for future Z-Machine

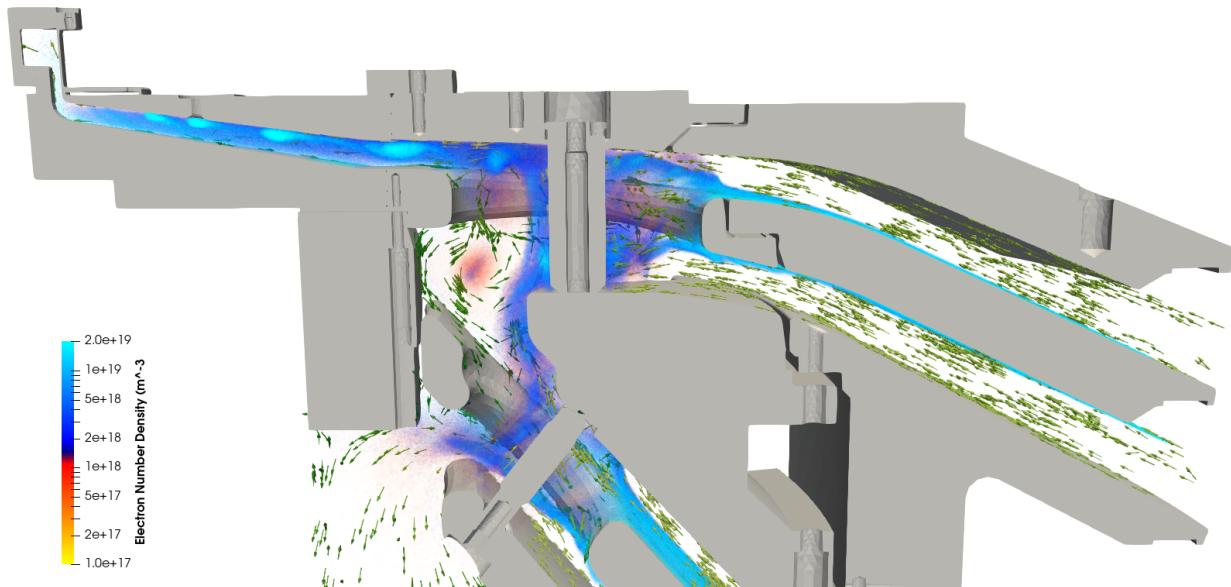


(a) electron number density



(b) B field magnitude

**Figure 4-28.** (EMPIRE-PIC Power Flow 18a: hot,  $t = 93.262$  ns) The electron number density (a) and magnetic field magnitude (b); several features of electron flow in a post-hole convolute system are exhibited in (a), for example strong insulation of emitted electrons in the outer (right) regions, electron injection into the convolute where the magnetic field topology changes (e.g. through the nulls shown in (b)), and the resulting vortex flows; In both renderings, the physical hardware has been included to aid interpretation.



**Figure 4-29.** (Power Flow 18a: hot,  $t = 93.262$  ns) The electron number density during a typical Z-Machine pulse;  $E \times B$  glyphs are drawn to show the dominant single-particle drifts;  $t = 0$  corresponds to 18 ns into the pulse.

experiments. We believe this is an exciting development. Further work to achieve a full physics power flow EMPIRE simulation on a leadership-class HPC machine (e.g., Sierra) is ongoing.

### 4.3.5. **Power Flow Physics Advancements in CHICAGO**

In this section, we describe simulations conducted in CHICAGO that demonstrate code performance in Z-relevant geometries and energy densities. We first describe resolution requirements as an aid to code analysts. These criteria are applied in 2D and 3D simulations of experimental configurations fielded (or planned) for the Z-Machine.

#### 4.3.5.1. **Resolution Requirements for Momentum Conservation in Kinetic Plasma Simulations**

Kinetic CHICAGO simulations use a magnetic implicit particle and field solution [40, 136] which simultaneously relax the time constraints of the cyclotron and plasma frequencies as well as the electromagnetic Courant limit, enabling  $\omega_c \Delta t > 1$ ,  $\omega_p \Delta t > 1$ , and  $c \Delta t \gg 1$ . An adaptive time step and particle management algorithm [139] enable both the time step and macroparticle count to adjust to the particle densities generated in the electrode plasma model [136].

In addition to relaxing the temporal resolution, the energy-conserving implicit particle solution allows the cell size ( $\Delta x$ ) to exceed the Debye length

$$\lambda_D = \sqrt{\frac{\epsilon_0 k T_e}{n_e e^2}}.$$

(Momentum-conserving particle advance algorithms traditionally result in numerical heating of the plasma until  $\Delta x \sim \lambda_D$ .) The artificial collisionality inherent in an energy-conserving algorithm is reduced by a cloud-in-cell treatment [10], in which a particle's effective density is uniformly distributed within its grid cell [138].

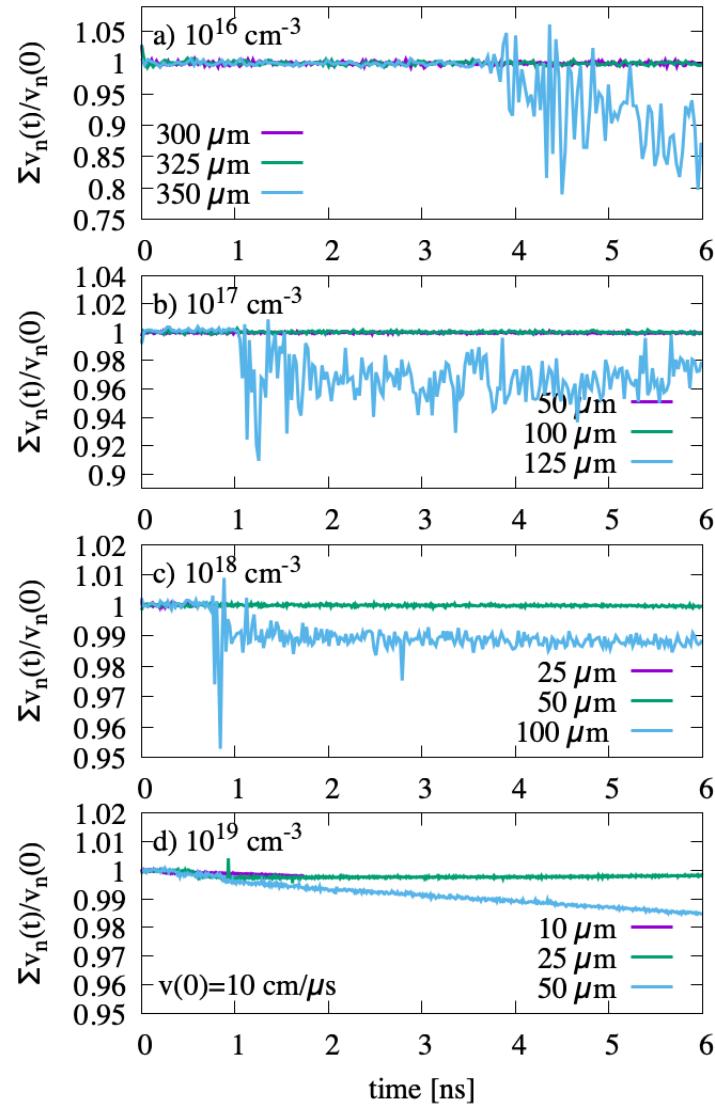
For electrode plasmas of  $10^{16} - 10^{19} \text{-cm}^{-3}$  density with temperatures below 100 keV, the electron collisionless skin depth,  $l_s = c/\omega_p$ , exceeds  $\lambda_D$  and is relevant for current carrying sheaths. The spatial resolution is then dictated by  $l_s$  and by the wavelengths of key instabilities. Reference [138] noted that the required resolution of a 3-eV/ $10^{15} \text{-cm}^{-3}$ -density plasma is  $\Delta x = 0.5l_s$  for an explicit solution and  $\Delta x = 7.2l_s$  for an implicit solution.

To test the spatial resolution required for the MITL electrode plasmas, we analyzed momentum conservation in the plasma as cell size is increased. For this study, a simple 2D ( $x, z$ ) simulation is initialized with a drifting  $H^+$  plasma and no external forces. Periodic boundaries are used to preclude surface losses. The time step used 0.5 – 0.8 times the Courant limit to remove temporal resolution as a limiting factor.

For relevance to electrode plasmas, the primary plasma drift velocity ( $v_d$ ) considered is 10 cm/ $\mu$ s and the density ranges from  $10^{16}$  to  $10^{19} \text{ cm}^{-3}$ . The net momentum of the plasma electrons is monitored over a brief 6-ns drift period. The simulation is repeated with increasing  $\Delta x$  to

determine the value above which momentum is not conserved, or  $\sum v_{d,n}(t) \neq \sum v_{d,n}(0)$ . The tolerance used here is  $|\sum v_{d,n}(0) - \sum v_{d,n}(t)| / \sum v_{d,n}(0) > 0.01$ .

The value of  $\sum v_{d,n}(t) / \sum v_{d,n}(0)$  is plotted for a range of  $\Delta x$  for a  $10^{16}\text{-cm}^{-3}$ -density plasma in Figure 4-30a. In the simulation with  $\Delta x = 325 \mu\text{m}$ , or  $\sim 6l_s$ , electrons escape the quasi-neutral plasma region. The electric field stresses generated by the redistributed space charge exceed 10 kV/cm and provide a feedback that accelerates the particle losses from the plasma.

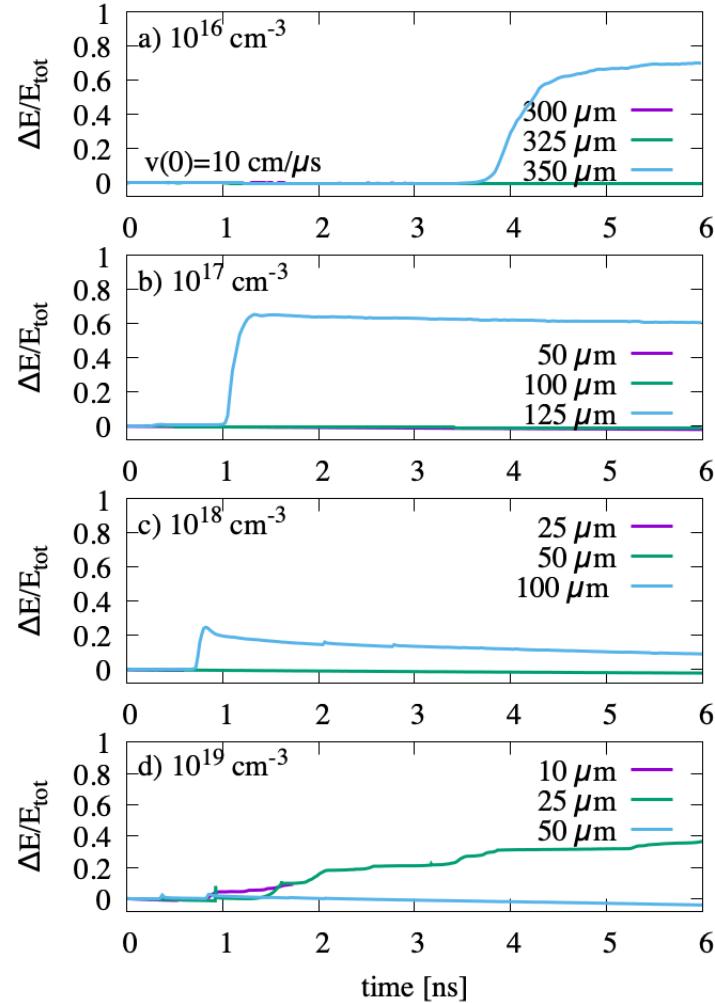


**Figure 4-30. Momentum conservation,  $\sum(v_n(t)/\sum v_n(0))$ , as a function of time and grid resolution for drifting plasmas with densities of a)  $10^{16} \text{ cm}^{-3}$ , b)  $10^{17} \text{ cm}^{-3}$ , c)  $10^{18} \text{ cm}^{-3}$ , and d)  $10^{19} \text{ cm}^{-3}$ .**

This feedback mechanism likely occurs widely in kinetic simulations. Space-charge separations occur randomly in kinetic plasmas but may be restored on a  $1/\omega_p$  time scale with sufficient resolution. When the simulation is under-resolved, particles near the plasma boundaries escape and the numerical instability may grow.

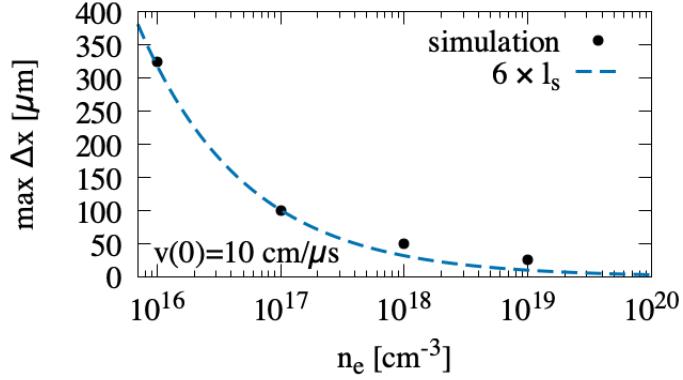
The study of momentum conservation is repeated for plasma densities of  $10^{17}$ ,  $10^{18}$ , and  $10^{19} \text{ cm}^{-3}$  with results shown in Figure 4-30b, 4-30c, and 4-30d, respectively. An estimate of the minimum resolution required is provided by the second largest grid size plotted for each density.

The analysis is repeated replacing momentum conservation with energy, where energy conservation is defined as  $\Delta E/E_{\text{tot}}$ . The minimum resolution as a function of plasma density are plotted in Figure 4-31 and are consistent with the momentum results.



**Figure 4-31. Energy conservation,  $\Delta E/E_{\text{tot}}$ , as a function of time and grid resolution for drifting plasmas with densities of a)  $10^{16} \text{ cm}^{-3}$ , b)  $10^{17} \text{ cm}^{-3}$ , c)  $10^{18} \text{ cm}^{-3}$ , and d)  $10^{19} \text{ cm}^{-3}$ .**

The minimum grid size for each density is plotted in Figure 4-32. The drop in cell size with density follows a trend line defined as six times the electron collisionless skin depth ( $l_s = c/\omega_p$ ). We caution that this scale factor is likely related to our chosen  $v_{\text{drift}}$  for these simulations, however a general relationship between grid resolution and  $l_s$  also appears likely.



**Figure 4-32. Minimum grid resolution as a function of plasma density with  $v_{drift} = 10 \text{ cm}/\mu\text{s}$ . The results scale as six times the electron collisionless skin depth ( $l_s = c/\omega_p$ ).**

#### 4.3.5.2. Resolution Requirements for Insulated Free Current Flow

A reliable method for determining the minimum resolution required for MITL flow is to compare simulation results to the transmission line theories of Creedon [27] and Mendel and Rosenthal [80]. For laminar sheath flow, both theories relate the voltage drop across the transmission line ( $V_0$ ) to the anode current ( $I_a$ ), the cathode boundary current ( $I_c$ ), and vacuum impedance ( $Z_0$ ). The relation derived from Reference [27] is

$$V_0 = 0.511 \left( \gamma_m - 1 + (\gamma_m^2 - 1)^{1/2} \left[ \frac{I_c Z_0}{I_a Z_\alpha} - \ln[\gamma_m + (\gamma_m^2 - 1)^{1/2}] \right] \right) \text{ [V].} \quad (4.4)$$

where  $Z_\alpha = 60 \Omega$ ,  $I_\alpha = (2\pi)(\mu_0 c)(m_e c^2)(e) \approx 8500 \text{ A}$ , and  $\gamma_m = I_a/I_c$  is also the relativistic factor at the extent of the sheath. This results in the MITL operating at an impedance ( $Z_{op}$ ) that is less than  $Z_0$ .

Agreement between  $\int \mathbf{E} \cdot d\mathbf{l}$  and  $V_0$  from Equation 4.4 for MITL flow may be regarded as a test of simulation accuracy, even for systems with a dynamic  $Z_{op}$  [17]. We determined the impact of spatial resolution on simulation accuracy using a 2D coaxial MITL operating at 40 kA. Using the difference between  $\int \mathbf{E} \cdot d\mathbf{l}$  and Equation 4.4 as the metric, the simulations converged to the accurate solution as the spatial resolution increased from 0.5 to 0.25 mm. Increasing the temporal resolution alone had no effect.

The emission rates for electric-field-stress and thermal desorption are governed by the space-charge-limited (SCL) current density, the non-relativistic form of which is [62]:

$$j_{CL} = \frac{4\epsilon_0}{9} \left( \frac{2Ze}{m} \right)^{1/2} \frac{V_0^{3/2}}{d^2}, \quad (4.5)$$

where  $V_0$  is the voltage drop,  $d$  is the AK gap width,  $m$  is the particle mass, and  $Z$  is its charge state. An exact relativistic solution was derived in Reference [56], the series expansion of which

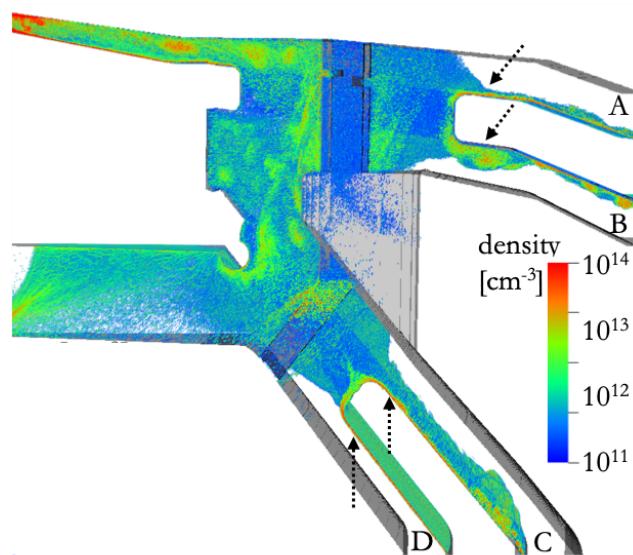
may be written

$$j_{rel\ CL} = \frac{2mc^3\epsilon_0}{Zed^2} \left[ \left( \frac{ZeV}{mc^2} + 1 \right)^{1/2} - 0.8471 \right]^2, \quad (4.6)$$

The magnitude of the electron current liberated from the cathode in SCL emission is defined by Equation 4.5, after which the electrons are accelerated by  $E_r$  in the grid cell adjacent to the conductor. As the resolution decreases,  $\Delta E_r$  in each cell increases and the electron begins sampling a local field value that is too large. The electrons are transported to larger radius than they would with proper resolution. The sheath thus becomes broader,  $Z_{op}$  is reduced, and the anode and cathode boundary currents are increased.

Resolution scans of a coaxial MITL demonstrate that this artificial sheath broadening will appear as a smaller effective AK gap, reducing  $Z_{op}$ , and will exhibit noise-induced turbulent radial motion. The simulation will thus appear to predict flow instabilities and current transport across the gap where none may actually exist.

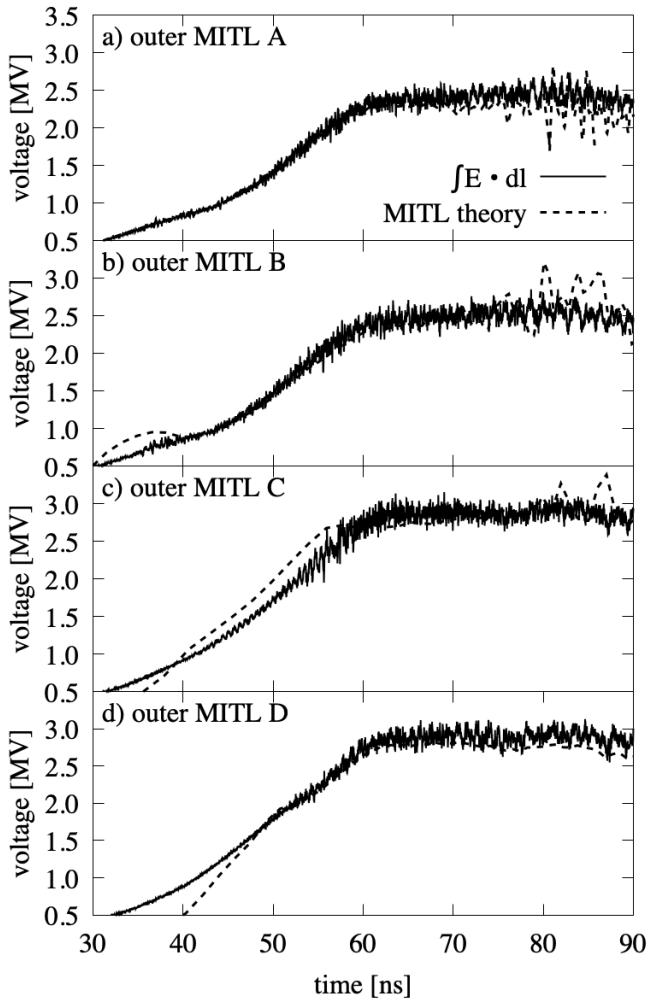
The test for agreement between  $\int \mathbf{E} \cdot d\mathbf{l}$  and Equation 4.4 should be conducted for each simulation grid condition. An example for the 3D Power Flow 18a geometry, shown in Figure 4-33, is plotted in Figure 4-34 for the four outer MITLs. The probe locations ( $I_a$ ,  $I_c$ , and  $V_0$ ) are marked by arrows in Figure 4-33.



**Figure 4-33. Electron density contours in the 31-cm convolute for the Power Flow 18a geometry during the electrical pulse rise; Arrows mark the locations in each outer MITL where Equation 4.4 is compared to the simulation voltage in Figure 4-34.**

#### 4.3.5.3. Example Power Flow Physics Results

Large-scale 3D simulations of the Power Flow 18a, Power Flow Validation 19a, and Large Convolute 19a (e.g., 31-cm convolute) geometries were executed on Sandia's compute clusters.

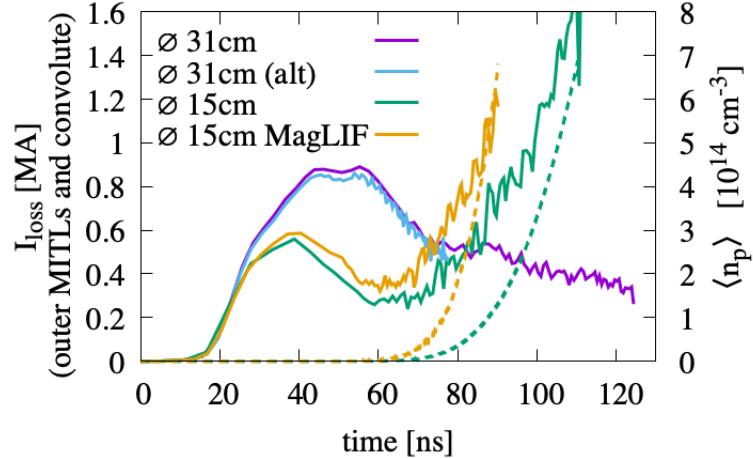


**Figure 4-34. The voltage drop across the four outer MITL gaps from the simulated electric fields compared to pressure balance theory of Equation 4.4.**

The 3D grids use symmetry boundaries to limit the models to a  $15^\circ$  wedge in azimuth with 8-10 divisions in  $\theta$ . Very high resolution simulations that span  $45^\circ$  in azimuth were executed on the LANL Trinity (ATS-1) platform for massive parallelization. These high-resolution simulations build on the results presented here by making comparisons to new diagnostics in the inner MITL and convolute regions. These comparisons will be presented in a future report.

The 3D simulations demonstrate that current loss in the convolute region is a result of initial uninsulated electron emission in the outer MITLs and, later, plasma expansion near the magnetic nulls [6]. Insulated flow ends at the convolute where the fields ( $B_\theta, E_{r,z}$ ) change sign. The combined losses ( $I_{loss}$ ) in the outer MITLs and convolute are plotted in Figure 4-35 for four example configurations fielded on the Z-Machine.

We explored an enhanced ion current as a loss mechanism in the inner MITL, but found that although the electric fields are enhanced due to electron space charge, cross gap transport is



**Figure 4-35. The current loss in the outer MITLs and convolute for two 31-cm and two 15-cm convolute configurations; The two 31-cm examples have similar load inductances while the two 15-cm differ by 80%; Dashed lines represent the mean proton densities in the 15-cm convolutes, matched by color.**

inhibited by the strong magnetic fields.

To estimate when and where ions are able to cross the A-K gap and contribute to current loss (without anomalous transport), we estimate the radius below which the ions are magnetically insulated. Using the current pulse to determine  $B_\theta(r, t)$  and the voltage ( $V(t)$ ) at  $r = 5$  cm to estimate the non-relativistic ion velocity, the ion Larmor radius equals the gap width ( $r_L = d$ ) at the MITL radius

$$r_{max} = \frac{e}{m} \frac{\mu_0 I(t)}{2\pi} \left( \frac{m}{2ZeV(t)} \right)^{1/2} d. \quad (4.7)$$

This is plotted in Figure 4-36 for  $H^+$  and  $O^+$  for a 7-mm gap. Also plotted is the maximum radius for thermal emission from Joule heating:

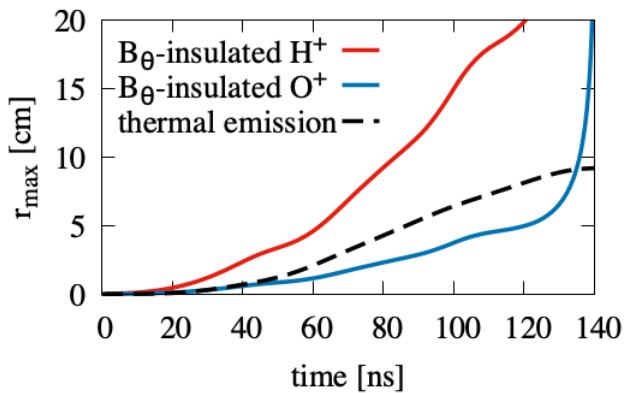
$$\Delta T_J(t) \approx \frac{\vartheta \mu_0 H^2(t)}{2c_v}, \quad (4.8)$$

This shows that for  $H^+$ , as with electrons, the conditions for insulation exist prior to emission. And although  $O^+$  is not insulated at emission, the surface area over which  $O^+$  current loss may occur reduces in time as  $j_{CL}$  is increasing.

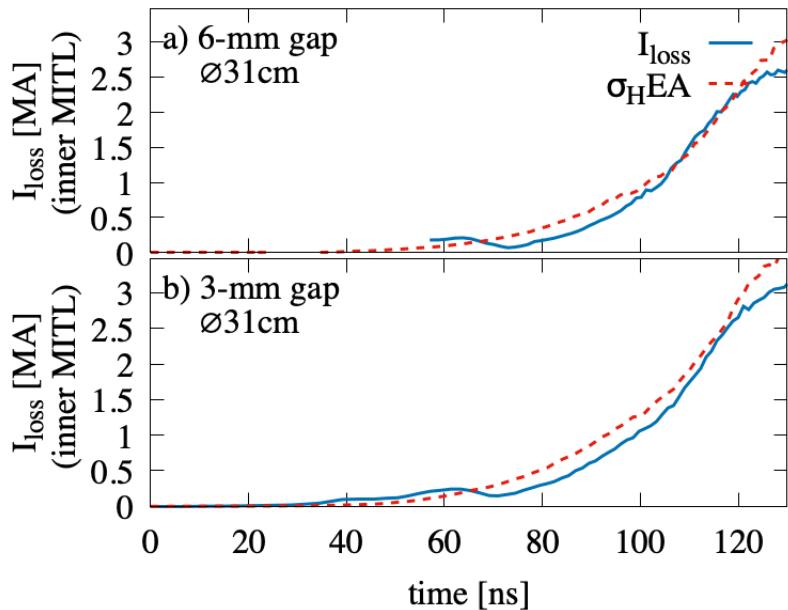
The Hall-conductivity related current loss has been explored in greater detail than in Reference [6]. We refer to the cross-term conductivity:

$$\sigma_\perp \frac{\omega_c}{v_{ei}} = \frac{n_e e^2 \omega_c}{m_e (v_{ei}^2 + \omega_c^2)}, \quad (4.9)$$

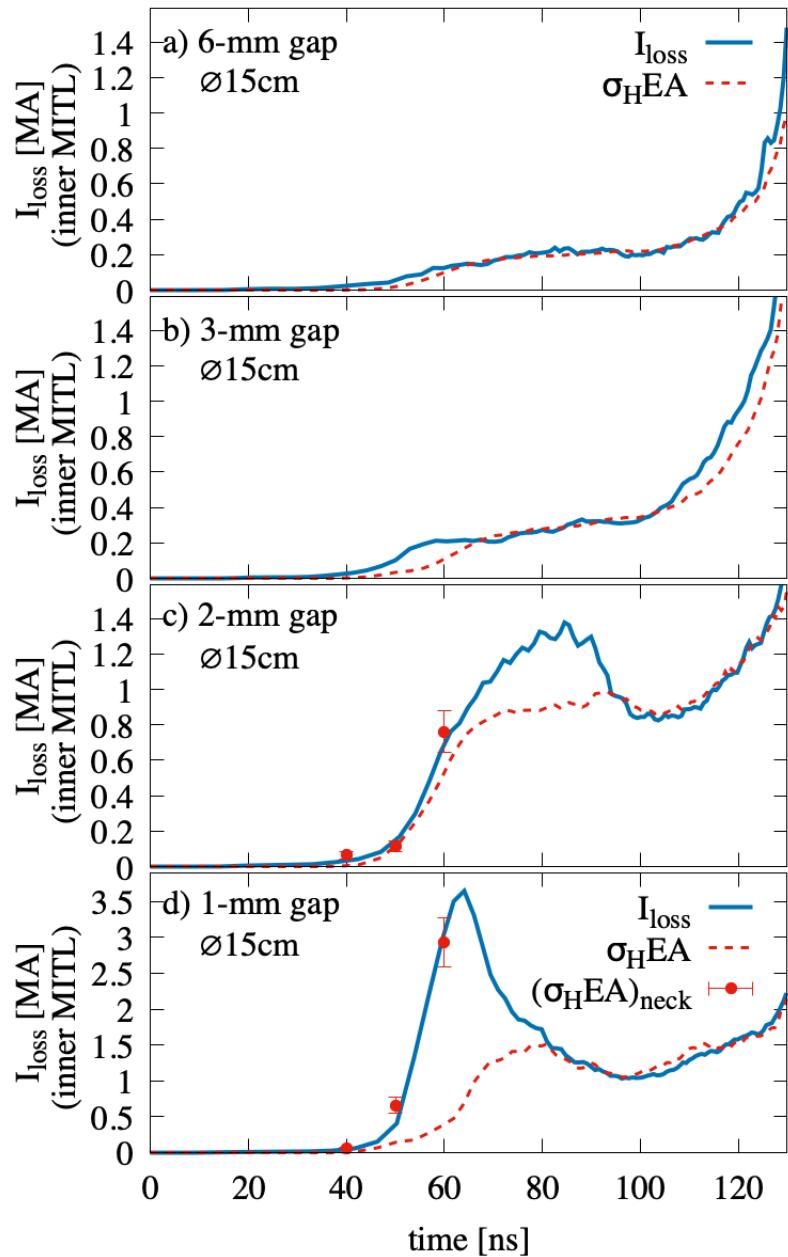
as the Hall conductivity  $\sigma_H$ . The comparisons of current loss with a volumetric Hall-related loss,  $\sigma_H EA$ , for example configurations (Power Flow 18a and Power Flow Validation 19a) are plotted in Figures 4-37 and 4-38. The early loss is a function of plasma at 1 cm radius.



**Figure 4-36.** The maximum radius at which ions are magnetically insulated as a function of time (Equation 4.7) for a 7-mm gap; The water constituents  $H^+$  and  $O^+$  are plotted in red and blue, respectively; The maximum radius for thermal emission from Equation 4.8 for  $\Delta T = 400$  K (in black) shows thermally emitted  $H^+$  are insulated.



**Figure 4-37.** Current loss compared to an estimate of Hall-conductivity loss for two example Power Flow 18a inner MITLs.



**Figure 4-38. Current loss compared to an estimate of Hall-conductivity loss for four example Power Flow Validation 19a inner MITLs with different A-K gaps.**

## 5. CONCLUSIONS AND RECOMMENDATIONS

This Grand Challenge LDRD (*Towards Predictive Plasma Science and Engineering through Revolutionary Multi-Scale Algorithms and Models*) was launched to directly address several high-risk science and technology gaps that will have a significant impact on future plasma modeling capabilities at the Laboratory. Ultimately this GC LDRD project explored a variety of concurrent research activities: improved electrode physics models, cross-code comparison tools, multi-level circuit coupling, HPC scalability, improved algorithms and solvers, mesh analysis routines, UQ infrastructure, and MITL physics analysis. While a large variety of topics have been explored, we believe the most significant accomplishments of this GC LDRD include:

1. First application of molecular dynamics (MD) and density function theory (DFT) techniques to quasi-validate multi-scale models of desorption mechanisms relevant to pulsed power electrode performance;
2. Demonstration of rapid temperature programmed desorption (TPD) validation experiments at a factor  $10^4$  faster heating rates than was available with pre-LDRD capabilities, coupled with automated analysis routines, and with preliminary scoping of even faster capabilities required to approximate Z-Machine heating rates (e.g.  $10^3$  further improvement is recommended to begin converging with MD scales);
3. Identification of surface chemistry as a potentially significant contributor in electrode contaminant release and subsequent current loss mechanism;
4. Exploratory measurements indicating a large increase in contaminant desorption is observed during electrode melt, the effects of which could be an important mechanism to explore for next-generation pulsed power accelerators;
5. Demonstration of significant improvements in multi-physics code technologies including algorithms, HPC scaling, and treatments resulting in an effective factor of 30-40x improvement over pre-LDRD performance in terms of practical design capability (e.g. demonstrated on a representative power flow benchmark test problem);
6. Maturation of EMPIRE pulsed power models to a level (or very near level) of design readiness and execution speed comparable to CHICAGO;
7. Integration of Trilinos/Kokkos solvers in the CHICAGO multi-physics research code, providing platform portability for next-generation HPC architectures;
8. Exploration of hybrid kinetic-fluid methods that highlight key algorithm challenges, and are positioned for future implementation in EMPIRE;

9. Significant theory and code methods developed to enable space-charge-limiting boundary models in relativistic fluid codes, such as EMPIRE-FLUID, which could lead to alternative hybrid kinetic-fluid approaches in EMPIRE;
10. Identification of key similarities between kinetic and multi-fluid power flow simulations completed in CHICAGO, which has already informed spin-off approaches and engineering design capabilities;
11. Completion of a robust cross-code comparision infrastructure for verification and validation activities consisting of multiple codes, models, and test problems positioned to enable a defensible power flow science approach to next-generation pulsed power;
12. First end-to-end demonstration of a rigorous uncertainty quantification (UQ) approach to pulsed power simulation using an example model design problem;
13. Demonstration of the largest fully-integrated power flow simulations completed on the leading HPC computing machines avialable in the NNSA complex today, with preliminary studies indicating relevant pulsed power engineering design calculations can now be completed in (or order) several days;

The goal of this report was to provide summaries of key accomplishments of this Grand Challenge LDRD; a large number of references are included, and the reader is directed to more extensive details found in those Sandia technical reports and archival journal articles; at the time of this writing more than 20 journal manuscripts have been published (or submitted) as a result of this LDRD project, with several more in development; the reader is referred to Appendix A for the list of publications.

In addition, an over-arching goal of this initiative was to bring together several disparate groups at the Laboratory, and transform those relationships into a more unified SNL community that is positioned to impact future mission opportunities. We believe this effort has been extremely successful in this goal, and more than 15 new early-career scientists, post-docs, and students have been hired into the Laboratory (across several S&T Centers) using direct support provided by this LDRD project. This is a key legacy of this initiative.

Looking to the future, many of the capabilities matured through this LDRD project are positioned for direct use in the core pulsed power science programs. EMPIRE and CHICAGO both achieved significant improvements as a result of this LDRD project, and it is critical both design tools continue to be exercised in complementary fashion through a stressing validation regime (e.g., against Z-Machine experiments designed to identify and constrain key power flow physics phenomena). Opportunities for continued improvement are available in both code technologies, particularly in the execution and performance of large-scale design calculations on leadership-class HPC machines (e.g., Trinity, Sierra); these investigations could be opportunities for strong collaboration between the NNSA Office of Experimental Sciences and the NNSA Advanced Scientific Computing programs. Several of the less mature (or higher risk) initiatives are already finding support through follow-on LDRD projects, either through the Radiation, Electrical, and High Energy Density Sciences (REHEDS) portfolio, or as part of the recently launched Assured Survivability and Agility with Pulsed Power (ASAP) LDRD misison campaign. All of these activities will be critical to demonstrating new multi-physics modeling tools, and

ultimately a defensible understanding of power flow physics phenomena, that can be leveraged to predict the performance of the next-generation of pulsed power accelerators for the stockpile science mission.

## REFERENCES

- [1] K. Alam and S. Simpson. Grazing angle infrared spectroscopy of mineral oil on stainless steel at low pressure. *to be submitted to Journal of Vacuum Science and Technology A*, 2021.
- [2] A. Y. Aydemir. A unified Monte Carlo interpretation of particle simulations and applications to non-neutral plasmas. *Physics of Plasmas*, 1(4):822–831, 1994.
- [3] J. W. Banks and J. A. Hittinger. A new class of non-linear, finite-volume methods for Vlasov simulation. *IEEE Transactions on Plasma Science*, 38(9):2198 – 2207, 2009.
- [4] D. C. Barnes, J. Cheng, and S. E. Parker. Low-noise particle algorithms for extended magnetohydrodynamic closure. *Physics of Plasmas*, 15(5):055702, 2008.
- [5] K. Beckwith and J. M. Stone. A second-order Godunov method for multi-dimensional relativistic magnetohydrodynamics. *The Astrophysical Journal Supplement Series*, 193:6, 2011.
- [6] N. Bennett, D. R. Welch, C. A. Jennings, E. Yu, M. H. Hess, B. T. Hutsel, G. Laity, J. K. Moore, D. V. Rose, K. Peterson, and M. E. Cuneo. Current transport and loss mechanisms in the Z accelerator. *Physical Review Accelerators and Beams*, 22:120401, 2019.
- [7] N. Bennett, D. R. Welch, D. V. Rose, G. Laity, J. K. Moore, and M. E. Cuneo. Magnetized particle transport in multi-ma accelerators. *to be submitted to Physical Review Accelerators and Beams*, 2021.
- [8] M. Bennoune, M. Lemou, and L. Mieussens. Uniformly stable numerical schemes for the Boltzmann equation preserving the compressible Navier–Stokes asymptotics. *Journal of Computational Physics*, 227(8):3781 – 3803, 2008.
- [9] M. Bettencourt and et. al. EMPIRE-PIC: A performance portable un-structured particle-in-cell code. SAND Technical Report in preparation, 2021.
- [10] C. K. Birdsall and A. B. Langdon. *Plasma Physics via Computer Simulation*. Institute of Physics Publishing, 1991.
- [11] G. Bodo, A. Mignone, and R. Rosner. Kelvin-Helmholtz instability for relativistic fluids. *Physical Review E*, 70:036304, 2004.
- [12] D. Bohm. *The Characteristics of Electrical Discharges in Magnetic Fields*. McGraw-Hill, 1949.
- [13] J. Boris. Relativistic plasma simulation: Optimization of a hybrid code. In *Proceedings of the Fourth Conference on Numerical Simulation of Plasmas*, pages 3–68, 1971.

- [14] J.U. Brackbill and H.M. Ruppel. FLIP: A method for adaptively zoned, particle-in-cell calculations of fluid flows in two dimensions. *Journal of Computational Physics*, 65(2):314 – 343, 1986.
- [15] F. Brezzi and M. Fortin. *Mixed and Hybrid Finite Element Methods*, volume 15. Springer Science & Business Media, 2012.
- [16] D. Brown and M. Bettencourt. Higher-order particle representation for particle-in-cell simulations. *submitted to Journal of Computational Physics*, 2020.
- [17] N. Bruner, T. Genoni, E. Madrid, D. Rose, D. Welch, K. Hahn, J. Leckbee, S. Portillo, B. Oliver, V. Bailey, and D. Johnson. Modeling particle emission and power flow in pulsed-power driven, nonuniform transmission lines. *Physical Review Special Topics Accelerators and Beams*, 11(4):040401, 2008.
- [18] C. Z. Cheng and G. Knorr. The integration of the Vlasov equation in configuration space. *Journal of Computational Physics*, 22:330–351, 1976.
- [19] K. Chenoweth, A. Van Duin, and W. Goddard. Reaxff reactive force field for molecular dynamics simulations of hydrocarbon oxidation. *The Journal of Physical Chemistry A*, 112(5):1040–1053, 2008.
- [20] A.J. Christlieb, R. Krasny, J.P. Verboncoeur, J.W. Emhoff, and I.D. Boyd. Grid-free plasma simulation techniques. *IEEE Transactions on Plasma Science*, 34(2):149–165, 2006.
- [21] K. Cochrane, A.G. Ilgen, R. Goeke, A. Wilson, and K. Leung. Elucidating hydrogen reaction-induced water desorption from oxide-passivated metal surfaces for plasma applications. SAND Technical Report in preparation, 2021.
- [22] B. Cockburn, S. Hou, and C-W. Shu. The Runge–Kutta local projection discontinuous Galerkin finite element method for conservation laws. IV. the multidimensional case. *Mathematics of Computation*, 54(190):545–581, 1990.
- [23] B. Cockburn, S-Y. Lin, and C-W. Shu. TVB Runge–Kutta local projection discontinuous Galerkin finite element method for conservation laws III: one-dimensional systems. *Journal of Computational Physics*, 84(1):90–113, 1989.
- [24] B. Cockburn and C-W. Shu. TVB Runge–Kutta local projection discontinuous Galerkin finite element method for conservation laws. II. general framework. *Mathematics of Computation*, 52(186):411–435, 1989.
- [25] B. Cockburn and C-W. Shu. The Runge–Kutta discontinuous Galerkin method for conservation laws V: multidimensional systems. *Journal of Computational Physics*, 141(2):199–224, 1998.
- [26] B. Cohen, A. B. Langdon, and A. Friedman. Implicit time integration for plasma simulation. *Journal of Computational Physics*, 46(1):15 – 38, 1982.
- [27] J. M. Creedon. Relativistic Brillouin flow in the high  $v/\gamma$  diode. *Journal of Applied Physics*, 46:2946, 1975.

[28] A. Crestetto, N. Crouseilles, and M. Lemou. Kinetic/fluid micro-macro numerical schemes for Vlasov-Poisson-BGK equation using particles. *Kinetic & Related Models*, 5(4):787, 2012.

[29] N. Crouseilles, M. Mehrenberger, and E. Sonnendrücker. Conservative semi-Lagrangian schemes for Vlasov equations. *Journal of Computational Physics*, 229(6):1927 – 1953, 2010.

[30] M. E. Cuneo. The effect of electrode contamination, cleaning and conditioning on high-energy pulsed-power device performance. *IEEE Transactions on Dielectrics and Electrical Insulation*, 6(4):469–485, 1999.

[31] M. E. Cuneo, P. R. Menge, D. L. Hanson, W. E. Fowler, M. A. Bernard, G. R. Ziska, A. B. Filuk, T. D. Pointon, R. A. Vesey, D. R. Welch, J. E. Bailey, M. P. Desjarlais, T. R. Lockner, T. A. Mehlhorn, S. A. Slutz, and M. A. Stark. Results of vacuum cleaning techniques on the performance of LiF field-threshold ion sources on extraction applied-B ion diodes at 1-10 TW. *IEEE Transactions on Plasma Science*, 25(2):229–251, 1997.

[32] R. Cygan, J-J. Liang, and A. Kalinichev. Molecular models of hydroxide, oxyhydroxide, and clay phases and the development of a general force field. *The Journal of Physical Chemistry B*, 108(4):1255–1266, 2004.

[33] J. Dawson. Particle simulation of plasmas. *Reviews of Modern Physics*, 55:403–447, 1983.

[34] G. Dimarco and R. Loubere. Towards an ultra efficient kinetic scheme. Part II: The high order case. *Journal of Computational Physics*, 255:699 – 719, 2013.

[35] G. Dimarco, L. Mieussens, and V. Rispoli. An asymptotic preserving automatic domain decomposition method for the Vlasov–Poisson–BGK system with applications to plasmas. *Journal of Computational Physics*, 274:122 – 139, 2014.

[36] T. Dupree. Kinetic theory of plasma and the electromagnetic field. *Physics of Fluids*, 6(12):1714–1729, 1963.

[37] E. G. Evstatiev, J. M. Finn, B. A. Shadwick, and N. Hengartner. Noise and error analysis and optimization in particle-based kinetic plasma simulations. *submitted to Journal of Computational Physics*, 2020.

[38] E.G. Evstatiev and B.A. Shadwick. Variational formulation of particle algorithms for kinetic plasma simulations. *Journal of Computational Physics*, 245:376 – 398, 2013.

[39] G. Friedman, G. Friedman, A. Gutsol, A. Shekhter, V. Vasilets, and A. Friedman. Applied plasma medicine. *Plasma Processes and Polymers*, 5(6):503–533, 2008.

[40] T. C. Genoni, R. E. Clark, and D. R. Welch. A fast implicit algorithm for highly magnetized charged particle motion. *The Open Plasma Physics Journal*, 3:36, 2010.

[41] R. Gerwin. Stability of the interface between two fluids in relative motion. *Review of Modern Physics*, 40:652–658, 1968.

[42] M.R. Gomez, R.M. Gilgenbach, M.E. Cuneo, C.A. Jennings, R.D. McBride, E.M. Waisman, B.T. Hutsel, W.A. Stygar, D.V. Rose, and Y. Maron. Experimental study of current loss and plasma formation in the Z machine post-hole convolute. *Physical Review Special Topics Accelerators and Beams*, 20(1):010401, 2017.

[43] H. Guo and A. Barnard. Surface structure and environment-dependent hydroxylation of the nonpolar hematite (100) from density functional theory modeling. *Journal of Physical Chemistry C*, 115:23023, 2011.

[44] N. Hamlin, A. Robinson, and E. Evstatiev. Cross-code verification and desorption uncertainty analysis of power-flow physics in a 2-D planar MITL. *to be submitted to Physical Review Accelerators and Beams*, 2021.

[45] N.D. Hamlin, K. Beckwith, T.M. Smith, and N. Roberds. Generalized fluid models for SCL emission. *to be submitted to IEEE Transactions on Plasma Science*, 2021.

[46] R.E. Heath, I.M. Gamba, P.J. Morrison, and C. Michler. A discontinuous Galerkin method for the Vlasov–Poisson system. *Journal of Computational Physics*, 231(4):1140 – 1174, 2012.

[47] M. H. Hess and E. G. Evstatiev. The vacuum fields and associated particle dynamics within cylindrically symmetric MITLs near a load. *to be submitted to Physical Review Accelerators and Beams*, 2021.

[48] T. Hoffard. Grazing-angle fourier transform infrared spectroscopy for surface cleanliness verification. Technical report, Naval Facilities Engineering Service Center, Port Hueneme, USA, 2003.

[49] R. Hooper, Z. Eckert, T. Hall, J. Boerner, J. Pacheco, and A. Grillet. Aleph: Highly scalable unstructured PIC-DSMC low temperature plasma code. Poster presentation (SAND2020-13413C) at the 47th IEEE International Conference on Plasma Sciences, December 6-10, 2020.

[50] X. Huang, S. Chang, W.S.V. Lee, J. Ding, and J.M. Xue. Three-dimensional printed cellular stainless steel as a high-activity catalytic electrode for oxygen evolution. *Journal of Materials Chemistry A*, 5, 2017.

[51] T.P Hughes and L.C. Musson. Spherical expansion of a dense plasma: test of ambipolar field model in Aleph. Technical Report SAND2014-19336, Sandia National Laboratories, Albuquerque, NM, 2014.

[52] B. T. Hutsel, P. A. Corcoran, M. E. Cuneo, M. R. Gomez, M. H. Hess, D. D. Hinshelwood, C. A. Jennings, G. R. Laity, D. C. Lampa, R. D. McBride, J. K. Moore, A. Myers, D. V. Rose, S. A. Slutz, W. A. Stygar, E. M. Waisman, D. R. Welch, and B. A. Whitney. Transmission-line-circuit model of an 85-TW, 25-MA pulsed-power accelerator. *Physical Review Accelerators and Beams*, 21:030401, 2018.

[53] John David Jackson. *Classical Electrodynamics*. American Association of Physics Teachers, 1999.

[54] G. Jaffé. On the currents carried by electrons of uniform initial velocity. *Physical Review*, 65(3):91–98, 1944.

[55] J. Jin. *The Finite Element Method in Electromagnetics*. Wiley-IEEE Press, 3rd edition, 2014.

[56] H.R. Jory and A.W. Trivelpiece. Exact relativistic solution for the one-dimensional diode. *Journal of Applied Physics*, 40:3924, 1969.

[57] H. Knoepfel. *Pulsed High Magnetic Fields*. North-Holland Publishing Company, 1970.

[58] H. Knoepfel. *Magnetic Fields: A Comprehensive Theoretical Treatise for Practical Use*. Wiley, 2008.

[59] J. Koski and J.M. Lane. Hydrocarbon and water desorption from iron-oxide surfaces using molecular dynamics. *AIP Conference Proceedings - Shock Compression of Condensed Matter*, 2020.

[60] R. Kramer, E. Cyr, S. Miller, E. Phillips, G. Radtke, A. Robinson, and J. Shadid. A plasma modeling hierarchy and verification approach. Technical Report SAND2020-3576, Sandia National Laboratories, 2020.

[61] J.M.D. Lane, K. Leung, A.P. Thompson, and M.E. Cuneo. Water desorption from rapidly-heated metal oxide surfaces-first principles, molecular dynamics, and the Temkin isotherm. *Journal of Physics: Condensed Matter*, 30:465002, 2018.

[62] I. Langmuir. The effect of space charge and residual gases on thermionic currents in high vacuum. *Physical Review*, 2:450, 1913.

[63] G. Lapenta. Exactly energy conserving semi-implicit particle in cell formulation. *Journal of Computational Physics*, 334:349 – 366, 2017.

[64] R. W. Lemke, M. D. Knudson, A. C. Robinson, T. A. Haill, K. W. Struve, J. R. Asay, and T. A. Mehlhorn. Self-consistent, two-dimensional, magnetohydrodynamic simulations of magnetically driven flyer plates. *Physics of Plasmas*, 10(5):1867–1874, 2003.

[65] M. Lemou. Relaxed micro–macro schemes for kinetic equations. *Comptes Rendus Mathematique*, 348(7):455 – 460, 2010.

[66] L. Leung, L.J. Criscenti, and A.C. Robinson. Quasi-equilibrium predictions of water desorption kinetics from rapidly-heated metal oxide surfaces. *Journal of Physics: Condensed Matter*, 32:335101, 2019.

[67] M. Li and H.F. Dylla. Modeling of Water Outgassing from Metal Surfaces (III). *Journal of Vacuum Science and Technology*, 13:1872–1878, 1995.

[68] H. Liu, E. Bianchetti, and C. Di Valentin. Insight into the interface between  $\text{Fe}_3\text{O}_4$  (001) surface and water overlayers through multiscale molecular dynamics simulations. *Journal of Chemical Physics*, 152:124771, 2020.

[69] J.W. Luginsland, Y.Y. Lau, R.J. Umstattt, and J.J. Watrous. Beyond the Child–Langmuir law: A review of recent results on multidimensional space-charge-limited flow. *Physics of Plasmas*, 9(5):2371–2376, 2002.

[70] S. H. Lui and K. Xu. Entropy analysis of kinetic flux vector splitting schemes for the compressible Euler equations. *Zeitschrift für angewandte Mathematik und Physik ZAMP*, 52(1):62–78, 2001.

[71] E. A. Madrid, D. V. Rose, D. R. Welch, R. E. Clark, C. B. Mostrom, W. A. Stygar, M. E. Cuneo, M. R. Gomez, T. P. Hughes, T. D. Pointon, and D. B. Seidel. Steady-state modeling of current loss in a post-hole convolute driven by high power magnetically insulated transmission lines. *Physical Review Special Topics Accelerators and Beams*, 16:120401, 2013.

[72] J.C. Mandal and S.M. Deshpande. Kinetic flux vector splitting for Euler equations. *Computers & Fluids*, 23(2):447 – 478, 1994.

[73] S. Markidis, V. Olshevsky, C. Sishtla, S. Chien, E. Laure, and G. Lapenta. PolyPIC: The Polymorphic-Particle-in-Cell method for fluid-kinetic coupling. *Frontiers in Physics*, 6:100, 2018.

[74] J. Martí and E. Müller. Numerical hydrodynamics in special relativity. *Living Reviews in Relativity*, 6(1), 2003.

[75] José María Martí and Ewald Müller. Grid-based methods in relativistic hydrodynamics and magnetohydrodynamics. *Living Reviews in Computational Astrophysics*, 1(1):3, 2015.

[76] R. McBride, T. Smith, P. Campbell, J. Woolstrum, G. Dowhan, S. Miller, A. Shah, B. Sporer, and N. Jordan. University of Michigan GC LDRD Year 3 Summary. Archived in project GitLab repository issue #124, October 1, 2020.

[77] R. D. McBride, W. A. Stygar, M. E. Cuneo, D. B. Sinars, M. G. Mazarakis, J. J. Leckbee, M. E. Savage, B. T. Hutsel, J. D. Douglass, M. L. Kiefer, B. V. Oliver, G. R. Laity, M. R. Gomez, D. A. Yager-Elorriaga, S. G. Patel, B. M. Kovalchuk, A. A. Kim, P. . Gourdain, S. N. Bland, S. Portillo, S. C. Bott-Suzuki, F. N. Beg, Y. Maron, R. B. Spielman, D. V. Rose, D. R. Welch, J. C. Zier, J. W. Schumer, J. B. Greenly, A. M. Covington, A. M. Steiner, P. C. Campbell, S. M. Miller, J. M. Woolstrum, N. B. Ramey, A. P. Shah, B. J. Sporer, N. M. Jordan, Y. Y. Lau, and R. M. Gilgenbach. A primer on pulsed power and linear transformer drivers for high energy density physics applications. *IEEE Transactions on Plasma Science*, 46(11):3928–3967, 2018.

[78] D.H. McDaniel, D. H. McDaniel, M. G. Mazarakis, D. E. Bliss and J. M. Elizondo, H. C. Harjes, H. C. Ives, D. L. Kitterman, J. E. Maenchen, T. D. Pointon, S. E. Rosenthal, D. L. Smith, K. W. Struve, W. A. Stygar, E. A. Weinbrecht, D. L. Johnson, and J. P. Corely. In J. Davis, editor, *Proceedings of the 5th International Conference on Dense Z-Pinches*, page 23, 2002.

[79] D. A. McGregor, E. Phillips, D. Sirajuddin, and T. Pointon. Variational, stable, and self-consistent coupling of 3D electromagnetics to 1D transmission lines in the time domain. *submitted to Journal of Computational Physics*, 2020.

[80] C. W. Mendel and S. E. Rosenthal. Modeling magnetically insulated devices using flow impedance. *Physics of Plasmas*, 2:1332, 1995.

[81] G. Mengaldo, D. De Grazia, F. Witherden, A. Farrington, P. Vincent, S. Sherwin, and J. Peiro. A guide to the implementation of boundary conditions in compact high-order methods for compressible aerodynamics. *7th AIAA Theoretical Fluid Mechanics Conference*, page 2923, 2014.

[82] A. Mignone and G. Bodo. An HLLC Riemann solver for relativistic flows – I. Hydrodynamics. *Monthly Notices of the Royal Astronomical Society*, 364(1):126–136, 2005.

[83] S. T. Miller, E. C. Cyr, J. N. Shadid, R. M. J. Kramer, E. G. Phillips, S. Conde, and R. P. Pawłowski. IMEX and exact sequence discretization of the multi-fluid plasma model. *Journal of Computational Physics*, 397:108806, 2019.

[84] S. Moe, J. Rossmanith, and D. Seal. A simple and effective high-order shock-capturing limiter for discontinuous Galerkin methods. *arXiv*, 2015.

[85] J-C. Nédélec. Mixed finite elements in  $\mathbb{R}^3$ . *Numerische Mathematik*, 35(3):315–341, 1980.

[86] D. R. Nicholson. *Introduction to Plasma Theory*. Krieger Publishing Company, 1992.

[87] J. Núñez-de la Rosa and C-D. Munz. Hybrid DG/FV schemes for magnetohydrodynamics and relativistic hydrodynamics. *Computer Physics Communications*, 222:113–135, 2018.

[88] S. Plimpton. Fast parallel algorithms for short-range molecular dynamics. *Journal of Computational Physics*, 117(1):1–19, 1995.

[89] T. Pointon and D. Seidel. Current loss in the vacuum section of the refurbished z accelerator. In *2009 IEEE Pulsed Power Conference*, pages 1159–1164. IEEE, 2009.

[90] W. Press, S. Teukolsky, W. Vetterling, and B. Flannery. *Numerical Recipes in C (2nd Ed.): The Art of Scientific Computing*. Cambridge University Press, USA, 1992.

[91] W. Press, S. Teukolsky, W. Vetterling, and B. Flannery. *Numerical Recipes (3rd Ed.): The Art of Scientific Computing*. Cambridge University Press, 2007.

[92] D.I Pullin. Direct simulation methods for compressible inviscid ideal-gas flow. *Journal of Computational Physics*, 34(2):231 – 244, 1980.

[93] J-M. Qiu and C-W. Shu. Conservative semi-Lagrangian finite difference WENO formulations with applications to the Vlasov equation. *Communications in Computational Physics*, 10(4):979–1000, 2011.

[94] T. Rene and M. Desjarlais. Equation of state and electrical conductivity of stainless steel. *SAND2004-5253*, 2004.

[95] G. Riccardi and D. Durante. Primitive variable recovering in special relativistic hydrodynamics allowing ultra-relativistic flows. *International Mathematical Forum*, 3(42):2081 – 2111, 2008.

[96] A. Robinson, T. Brunner, S. Carroll, R. Drake, C. Garasi, T. Gardiner, T. Haill, H. Hanshaw, D. Hensinger, D. Labreche, R. Lemke, E. Love, C. Luchini, S. Mosso, J. Niederhaus, C. Ober, S. Petney, W. Rider, G. Scovazzi, O.E. Strack, R. Summers, T. Trucano, V.G. Weirs, M. Wong, and T. Voth. ALEGRA: An arbitrary Lagrangian-Eulerian multimaterial, multiphysics code. *Proceedings of the 46th AIAA Aerospace Sciences Meeting*, 2008.

[97] A. Robinson, R. Drake, M.S. Swan, N. Bennett, T. Smith, and R. Hooper. A software environment for effective reliability management for pulsed power design. *submitted to Reliability Engineering and System Safety*, 2021.

[98] A. Robinson and C. Garasi. Three-dimensional z-pinch wire array modeling with ALEGRA-HEDP. *Computer Physics Communications*, 164(1):408 – 413, 2004.

[99] A. Robinson and A. Porwitzky. Pulsed power accelerator surface Joule heating models. *Physics of Plasmas*, 27(10):103103, 2020.

[100] A. Robinson, T. Smith, R. Hooper, and M.S. Swan. Exact, PIC and multifluid solutions to an expanding neutral plasma. *to be submitted to Physics of Plasmas*, 2021.

[101] A. Robinson, M.S. Swan, R. Hooper, T. Smith, R. Drake, and N. Bennett. Neutral expansion exact solutions for plasma code verification. Poster presentation (SAND2019-10067C) at the International Conference on Numerical Simulation of Plasmas, September 3-5, Santa Fe, New Mexico, 2019.

[102] A. Rokhlenko and J.L. Lebowitz. Linear stability of electron flow produced by field emission. *Journal of Applied Physics*, 113(063304), 2013.

[103] D. V. Rose, E. A. Madrid, D. R. Welch, R. E. Clark, C. B. Mostrom, W. A. Stygar, and M. E. Cuneo. Computational analysis of current-loss mechanisms in a post-hole convolute driven by magnetically insulated transmission lines. *Physical Review Special Topics Accelerators and Beams*, 18:030402, 2015.

[104] D. V. Rose, D. R. Welch, T. P. Hughes, R. E. Clark, and W. A. Stygar. Plasma evolution and dynamics in high-power vacuum-transmission-line post-hole convolutes. *Physical Review Special Topics - Accelerators and Beams*, 2008.

[105] D. V. Rose, D. R. Welch, E. A. Madrid, C. L. Miller, R. E. Clark, W. A. Stygar, M. E. Savage, G. A. Rochau, J. E. Bailey, T. J. Nash, M. E. Sceiford, K. W. Struve, P. A. Corcoran, and B. A. Whitney. Three-dimensional electromagnetic model of the pulsed-power Z-pinch accelerator. *Physical Review Special Topics Accelerators and Beams*, 13:010402, 2010.

[106] J. Rossmanith and D. Seal. A positivity-preserving high-order semi-Lagrangian discontinuous Galerkin scheme for the Vlasov–Poisson equations. *Journal of Computational Physics*, 230(16):6203–6232, 2011.

[107] K. Rustin, C. Loo, J. Klingenberg, S. Almesaied, and M. Lee. Potential for high-speed mass spectrometry for pulsed-power applications. *Report from Northern Arizona University*, 2018.

[108] D. Ryu, I. Chattopadhyay, and E. Choi. Equation of state in numerical relativistic hydrodynamics. *The Astrophysical Journal Supplement Series*, 166(1):410–420, 2006.

[109] M. E. Savage, K. N. Austin, B. T. Hutsel, R. J. Kamm, G. R. McKee, W. A. Stygar, P. E. Wakeland, N. R. Wemple, and W. M. White. Pulsed power performance of the Z machine: Ten years after the upgrade. In *IEEE International Pulsed Power Conference*, 2018.

[110] M. E. Savage, L. F. Bennett, D. E. Bliss, W. T. Clark, R. S. Coats, J. M. Elizondo, K. R. LeChien, H. C. Harjes, J. M. Lehr, J. E. Maenchen, D. H. McDaniel, M. F. Pasik, T. D. Pointon, A. C. Owen, D. B. Seidel, D. L. Smith, B. S. Stoltzfus, K. W. Struve, W. A. Stygar, L. K. Warne, J. R. Woodworth, C. W. Mendel, K. R. Prestwich, R. W. Shoup, D. L. Johnson, J. P. Corley, K. C. Hodge, T. C. Wagoner, and P. E. Wakeland. In *Proceedings of the 16th IEEE Pulsed Power and Plasma Science Conference, Albuquerque, NM*, page 979, 2007.

[111] V. Schneider, U. Katscher, D.H. Rischke, B. Waldhauser, J.A. Maruhn, and C-D. Munz. New algorithms for ultra-relativistic numerical hydrodynamics. *Journal of Computational Physics*, 105(1):92–107, 1993.

[112] J. N. Shadid, E. M.M. Crockatt, T.M. Smith, , S. Conde, R.P. Pawlowski, A. Rappaport, and T.A. Gardiner. Towards multifluid multiphysics continuum plasma simulation for modeling magnetically-driven experiments on Z. Sandia Technical Report SAND2019-12553, Sandia National Laboratories, Albuquerque, NM, 2019.

[113] J. N. Shadid, E. C. Cyr, E. G. Phillips, R. Pawlowski, M. Bettencourt, K. Cartwright, P. Lin, R. Kramer, J. Niederhaus, G. Radtke, and A. Robinson. Towards scalable and efficient solution of full Maxwell electromagnetics - multifluid plasma systems. Presented at NECDC, 2016.

[114] S. Shields, D.S. Jensen, B.M. Medina, T.C. Powell, T.D. Pointon, and C.H. Moore. The thermalization verification problem for EMPIRE-PIC and EMPIRE-Fluid. Sandia Technical Report SAND2019-9621, Sandia National Laboratories, Albuquerque, NM, 2019.

[115] Y. Shin, H. Kwak, A. Vasenkov, D. Sengupta, and A. van Duin. Development of a ReaxFF reactive force field for Fe/Cr/O/S and application to oxidation of butane over a pyrite-covered Cr<sub>2</sub>O<sub>3</sub> catalyst. *Acs Catalysis*, 5(12):7226–7236, 2015.

[116] S. Simpson and A. Yalin. Alternative methods for analyzing temperature programmed desorption spectra of complex substrates with random adsorbates. *to be submitted Journal of Surface Science*, 2021.

[117] D. B. Sinars, M. A. Sweeney, C. S. Alexander, D. J. Ampleford, T. Ao, J. P. Apruzese, C. Aragon, D. J. Armstrong, K. N. Austin, T. J. Awe, A. D. Baczewski, J. E. Bailey, K. L. Baker, C. R. Ball, H. T. Barclay, S. Beatty, K. Beckwith, K. S. Bell, J. F. Benage, N. L. Bennett, K. Blaha, D. E. Bliss, J. J. Boerner, C. J. Bourdon, B. A. Branch, J. L. Brown, E. M. Campbell, R. B. Campbell, D. G. Chacon, G. A. Chandler, K. Chandler, P. J. Christenson, M. D. Christison, E. B. Christner, R. C. Clay, K. R. Cochrane, A. P. Colombo, B. M. Cook, C. A. Coverdale, M. E. Cuneo, J. S. Custer, A. Dasgupta, J.-P. Davis, M. P.

Desjarlais, D. H. Dolan, J. D. Douglass, G. S. Dunham, S. Duwal, A. D. Edens, M. J. Edwards, E. G. Evstatiev, B. G. Farfan, J. R. Fein, E. S. Field, J. A. Fisher, T. M. Flanagan, D. G. Flicker, M. D. Furnish, B. R. Galloway, P. D. Gard, T. A. Gardiner, M. Geissel, J. L. Giuliani, M. E. Glinsky, M. R. Gomez, T. Gomez, G. P. Grim, K. D. Hahn, T. A. Haill, N. D. Hamlin, J. H. Hammer, S. B. Hansen, H. L. Hanshaw, E. C. Harding, A. J. Harvey-Thompson, D. Headley, M. C. Herrmann, M. H. Hess, C. Highstrete, O. A. Hurricane, B. T. Hutsel, C. A. Jennings, O. M. Johns, D. Johnson, M. D. Johnston, B. M. Jones, M. C. Jones, P. A. Jones, P. E. Kalita, R. J. Kamm, J. W. Kellogg, M. L. Kiefer, M. W. Kimmel, P. F. Knapp, M. D. Knudson, A. Kreft, G. R. Laity, P. W. Lake, D. C. Lamppa, W. L. Langston, J. S. Lash, K. R. LeChien, J. J. Leckbee, R. J. Leeper, G. T. Leifeste, R. W. Lemke, W. Lewis, S. A. Lewis, G. P. Loisel, Q. M. Looker, A. J. Lopez, D. J. Lucero, S. A. MacLaren, R. J. Magyar, M. A. Mangan, M. R. Martin, T. R. Mattsson, M. K. Matzen, A. J. Maurer, M. G. Mazarakis, R. D. McBride, H. S. McLean, C. A. McCoy, G. R. McKee, J. L. McKenney, A. R. Miles, J. A. Mills, M. D. Mitchell, N. W. Moore, C. E. Myers, T. Nagayama, G. Natoni, A. C. Owen, S. Patel, K. J. Peterson, T. D. Pointon, J. L. Porter, A. J. Porwitzky, S. Radovich, K. S. Raman, P. K. Rambo, W. D. Reinhart, G. K. Robertson, G. A. Rochau, S. Root, D. V. Rose, D. C. Rovang, C. L. Ruiz, D. E. Ruiz, D. Sandoval, M. E. Savage, M. E. Sceiford, M. A. Schaeuble, P. F. Schmit, M. S. Schollmeier, J. Schwarz, C. T. Seagle, A. B. Sefkow, D. B. Seidel, G. A. Shipley, J. Shores, L. Shulenburger, S. C. Simpson, S. A. Slutz, I. C. Smith, C. S. Speas, P. E. Specht, M. J. Speir, D. C. Spencer, P. T. Springer, A. M. Steiner, B. S. Stoltzfus, W. A. Stygar, J. Ward Thornhill, J. A. Torres, J. P. Townsend, C. Tyler, R. A. Vesey, P. E. Wakeland, T. J. Webb, E. A. Weinbrecht, M. R. Weis, D. R. Welch, J. L. Wise, M. Wu, D. A. Yager-Elorriaga, A. Yu, and E. P. Yu. Review of pulsed power-driven high energy density physics research on Z at Sandia. *Physics of Plasmas*, 27(7):070501, 2020.

- [118] D. Sirajuddin and M. Bettencourt. Half-o-lute exemplar for the code-comparison infrastructure (CCI) testing environment. SAND Technical Report in preparation, 2021.
- [119] D. Sirajuddin, M. Bettencourt, and N. Bennett. Power Flow 18a exemplar for the code-comparison infrastructure (CCI) testing environment. SAND Technical Report in preparation, 2021.
- [120] D. Sirajuddin and W. Hitchon. A truly forward semi-Lagrangian WENO scheme for the Vlasov-Poisson system. *Journal of Computational Physics*, 392, 05 2019.
- [121] D. Sirajuddin, S. Shields, and T. Smith. A two-stream instability exemplar for cross-code verification. SAND Technical Report in preparation, 2021.
- [122] T. Smith, P. Campbell, G. Dowhan, S. Miller, A. Shah, B. Sporer, N. Jordan, and R. McBride. University of Michigan GC LDRD Year 2 Summary. Archived in project GitLab repository issue #124, September 30, 2019.
- [123] T.J. Smith, P.C. Campbell, G. V. Dowhan, N. M. Jordan, M. D. Johnston, G. R. Laity, , and R. D. McBride. Power flow platform development for plasmas in university scaled, high power transmission lines at the 1-MA MAIZE facility. *submitted to Review of Scientific Instruments*, 2021.

[124] T.M. Smith, K. Beckwith, N.D. Hamlin, and N. Roberds. Verification of plasma physics codes. Technical report SAND2020-7927C, Sandia National Laboratories, Albuquerque, New Mexico 87185, August 2020.

[125] T.M. Smith, T.D. Pointon, K.L. Cartwright, and W.J. Rider. EMPIRE-PIC code verification of a cold diode. Technical report SAND2019-9384, Sandia National Laboratories, Albuquerque, New Mexico 87185, August 2019.

[126] T. Stix. *Waves in Plasmas*. American Institute of Physics, 1992.

[127] J. Stone and T. Gardiner. A simple unsplit godunov method for multidimensional MHD. *New Astronomy*, 14(2):139 – 148, 2009.

[128] J. Stone, K. Tomida, C. White, and K. Felker. The Athena++ adaptive mesh refinement framework: Design and magnetohydrodynamic solvers. *The Astrophysical Journal Supplement Series*, 249(1):4, 2020.

[129] W.A. Stygar, M.E. Cuneo, D.I. Headley, H.C. Ives, R.J. Leeper, M.G. Mazarakis, C.L. Olson, J.L. Porter, T.C. Wagoner, and J.R. Woodworth. Architecture of petawatt-class Z-pinch accelerators. *Physical Review Special Topics Accelerators and Beams*, 10:030401, 2007.

[130] C. Thoma, D. R. Welch, R. E. Clark, D. V. Rose, and I. E. Golovkin. Hybrid-PIC modeling of laser-plasma interactions and hot electron generation in gold hohlraum walls. *Physics of Plasmas*, 24(6):062707, 2017.

[131] C. Thoma, D. R. Welch, and D. V. Rose. Implicit highly-coupled single-ion Hall-MHD formulation for hybrid particle-in-cell codes. *Computer Physics Communications*, 261, 2021.

[132] J. Tranchida, J.M.D. Lane, and A. Thompson. Water coverage and desorption from hematite [0001] slit pores. (*in preparation*), 2021.

[133] A. Van Duin, S. Dasgupta, F. Lorant, and W. Goddard. ReaxFF: a reactive force field for hydrocarbons. *The Journal of Physical Chemistry A*, 105(41):9396–9409, 2001.

[134] E. M. Waisman, R. D. McBride, M. E. Cuneo, D. F. Wenger, W. E. Fowler, W. A. Johnson, L. I. Basilio, R. S. Coats, C. A. Jennings, D. B. Sinars, R. A. Vesey, B. Jones, D. J. Ampleford, R. W. Lemke, M. R. Martin, P. C. Schrafel, S. A. Lewis, J. K. Moore, M. E. Savage, and W. A. Stygar. Voltage measurements at the vacuum post-hole convolute of the Z pulsed-power accelerator. *Physical Review Special Topics Accelerators and Beams*, 17:120401, 2014.

[135] D. Welch, D. Rose, N. Bruner, R. Clark, B. Oliver, K. Hahn, and M. Johnston. Hybrid simulation of electrode plasmas in high-power diodes. *Physics of Plasmas*, 16(12):123102, 2009.

[136] D. R. Welch, N. Bennett, T. C. Genoni, D. V. Rose, C. Thoma, C. Miller, and W. A. Stygar. Electrode contaminant plasma effects in  $10^7$ -A Z pinch accelerators. *Physical Review Accelerators and Beams*, 22:070401, 2019.

- [137] D. R. Welch, N. Bennett, T. C. Genoni, C. Thoma, and D. V. Rose. Fast hybrid particle-in-cell technique for pulsed-power accelerators. *Physical Review Accelerators and Beams*, 23:110401, 2020.
- [138] D. R. Welch, D. V. Rose, B. V. Oliver, E. Schamloglu, K. Hahn, and J. E. Maenchen. Transport of a relativistic electron beam in gas and plasma-filled focusing cells for X-ray radiography. *Physics of Plasmas*, 11(2):751–760, 2004.
- [139] D.R. Welch, T.C. Genoni, R.E. Clark, and D.V. Rose. Adaptive particle management in a particle-in-cell code. *Journal of Computational Physics*, 227(1):143 – 155, 2007.
- [140] D.R. Welch, D.V. Rose, B.V. Oliver, and R.E. Clark. Simulation techniques for heavy ion fusion chamber transport. *Nuclear Instruments and Methods in Physics Research A*, 464(1–3):134 – 139, 2001.
- [141] K. Wu and H. Tang. Physical-constraint-preserving central discontinuous Galerkin methods for special relativistic hydrodynamics with a general equation of state. *The Astrophysical Journal Supplement Series*, 228(1):3, 2016.
- [142] K. Xu, L. Martinelli, and A. Jameson. Gas-kinetic finite volume methods, flux-vector splitting, and artificial diffusion. *Journal of Computational Physics*, 120(1):48 – 65, 1995.
- [143] B. Yan. A hybrid method with deviational particles for spatial inhomogeneous plasma. *Journal of Computational Physics*, 309:18 – 36, 2016.
- [144] B. Yan and R. Caflisch. A Monte Carlo method with negative particles for Coulomb collisions. *Journal of Computational Physics*, 298:711 – 740, 2015.
- [145] Yu. L. Klimontovich. *The Statistical Theory of Non-Equilibrium Processes in a Plasma: International Series of Monographs in Natural Philosophy*, volume 9. Elsevier, 2013.

## APPENDIX A. External Publications Generated by this LDRD Project

1. Alam et al., Grazing Angle Infrared Spectroscopy of Mineral Oil on Stainless Steel at Low Pressure (drafted), *Journal of Vacuum Science and Technology A* [1]
2. Bennett et al., Current Transport and Loss Mechanisms in the Z Accelerator, *Physical Review Accelerators and Beams* [6]
3. Bennett et al., Magnetized Particle Transport in Multi-MA Accelerators (drafted), *Physical Review Accelerators and Beams* [7]
4. Brown et al., Higher-Order Particle Representation for Particle-in-Cell Simulations (submitted), *Journal of Computational Physics* [16]
5. Evstatiiev et al., Noise and Error Analysis and Optimization in Particle-based Kinetic Plasma Simulations (submitted), *Journal of Computational Physics* [37]
6. Hamlin et al., Cross-code Verification and Desorption Uncertainty Analysis of Power Flow Physics in a 2-D Planar MITL (drafted), *Physical Review Accelerators and Beams* [44]
7. Hamlin et al., Generalized Fluid Models for SCL Emission (drafted), *IEEE Transactions on Plasma Science* [45]
8. Hess et al., The Vacuum Fields and Associated Particle Dynamics within Cylindrically Symmetric MITLs near a Load (drafted), *Physical Review Accelerators and Beams* [47]
9. Koski et al., Hydrocarbon and Water Desorption from Iron-oxide Surfaces using Molecular Dynamics (submitted), *AIP Conference Proceedings - Shock Compression of Condensed Matter* [59]
10. Lane et al., Water Desorption from Rapidly-heated Metal Oxide Surfaces: First Principles, Molecular Dynamics, and the Temkin Isotherm, *Journal of Physics: Condensed Matter* [61]
11. Leung et al., Quasi-Equilibrium Predictions of Water Desorption Kinetics from Rapidly-heated Metal Oxide Surfaces, *Journal of Physics: Condensed Matter* [66]
12. McBride et al., A Primer on Pulsed Power and Linear Transformer Drivers for High Energy Density Physics Applications, *IEEE Transactions on Plasma Science* [77]
13. McGregor et al., Variational, Stable, and Self-Consistent Coupling of 3D Electromagnetics to 1D Transmission Lines in the Time Domain (submitted), *Journal of Computational Physics* [79]
14. Miller et al., IMEX and Exact Sequence Discretization of the Multi-Fluid Plasma Model, *Journal of Computational Physics* [83]

15. Robinson et al., A Software Environment for Effective Reliability Management for Pulsed Power Design (submitted), *Reliability Engineering and System Safety* [97]
16. Robinson et al., Pulsed Power Accelerator Surface Joule Heating Models, *Physics of Plasmas* [99]
17. Robinson et al., Exact, PIC, and Multi-Fluid Solutions to an Expanding Neutral Plasma (drafted), *Physics of Plasmas* [100]
18. Simpson et al., Alternative Methods for Analyzing Temperature Programmed Desorption Spectra of Complex Substrates with Random Adsorbates (drafted), *Journal of Surface Science* [116]
19. Smith et al., Power Flow Platform Development for Plasmas in University-scaled, High Power Transmission Lines at the 1-MA MAIZE Facility (submitted), *Review of Scientific Instruments* [123]
20. Thoma et al., Implicit Highly-Coupled Single-Ion Hall-MHD Formulation for Hybrid Particle-in-Cell Codes, *Computer Physics Communications* [131]
21. Tranchida et al., Water Coverage and Desorption from Hematite (0001) Slit Pores (drafted) [132]
22. Welch et al., Electrode Contaminant Plasma Effects in  $10^7$ -A Z-pinch Accelerators, *Physical Review Accelerators and Beams* [136]
23. Welch et al., Fast Hybrid Particle-in-Cell Technique for Pulsed Power Accelerators, *Physical Review Accelerators and Beams* [137]

## DISTRIBUTION

### Hardcopy—Internal

Number of Copies	Name	Org.	Mailstop
1	George R. Laity	1659	1195
1	L. Martin, LDRD Office	1910	0359

### Email—External [REDACTED]

Name	Company Email Address	Company Name
John Luginsland, Ph.D.	john.luginsland@confluentsciences.com	Confluent Sciences
Luis Chacon, Ph.D.	chacon@lanl.gov	Los Alamos National Laboratory
Chris Clouse, Ph.D.	clouse1@llnl.gov	Lawrence Livermore National Laboratory
Col. Greg Van Dyk, Ph.D.	gregory.van_dyk.1@us.af.mil	U.S. Air Force Nuclear Weapons Center
Brendan Godfrey, Ph.D.	brendan.godfrey@ieee.org	Lawrence Berkeley National Laboratory
Thomas A. Mehlhorn, Ph.D.	tamehlh@uw.edu	U.S. Naval Research Laboratory
Charles Seyler, Ph.D.	ces7@cornell.edu	Cornell University
John Verboncoeur, Ph.D.	johnv@egr.msu.edu	Michigan State University

### Email—Internal [REDACTED]

Name	Org.	Sandia Email Address
George R. Laity	1659	grlaity@sandia.gov
Allen C. Robinson	1443	acrobin@sandia.gov
Michael E. Cuneo	1650	mecuneo@sandia.gov
Technical Library	1911	sanddocs@sandia.gov



**Sandia  
National  
Laboratories**

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.