

©Copyright 2014
Christopher James Hansen

MHD Modeling in Complex 3D Geometries:
Towards Predictive Simulation of SIHI Current Drive

Christopher James Hansen

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

Doctor of Philosophy

University of Washington

2014

Reading Committee:

Thomas R. Jarboe, Chair

Brian A. Nelson

George J. Marklin

Program Authorized to Offer Degree:
Aeronautics & Astronautics

University of Washington

Abstract

MHD Modeling in Complex 3D Geometries:
Towards Predictive Simulation of SIHI Current Drive

Christopher James Hansen

Chair of the Supervisory Committee:
Professor Thomas R. Jarboe
Aeronautics & Astronautics

The HIT-SI experiment studies Steady Inductive Helicity Injection (SIHI) for the purpose of forming and sustaining a spheromak plasma. A spheromak is formed in a nearly axisymmetric flux conserver, with a bow tie cross section, by means of two semi-toroidal injectors. The plasma-facing surfaces of the device, which are made of copper for its low resistivity, are covered in an insulating coating in order to operate in a purely inductive manner. Following formation, the spheromak flux and current are increased during a quiescent period marked by a decrease in the global mode activity. A proposed mechanism, Imposed Dynamo Current Drive (IDCD), is expected to be responsible for this phase of quiescent current drive. Due to the geometric complexity of the experiment, previous numerical modeling efforts have used a simplified geometry that excludes the injector volumes from the simulated domain. The effect of helicity injection is then modeled by boundary conditions on this reduced plasma volume. The work presented here has explored and developed more complete computational models of the HIT-SI device. This work is separated into 3 distinct but complementary areas: 1) Development of a 3D MHD equilibrium code that can incorporate the non-axisymmetric injector fields present in HIT-SI and investigation of equilibria of interest during spheromak sustainment. 2) A 2D axisymmetric MHD

equilibrium code that was used to explore reduced order models for mean-field evolution using equations derived from IDCD theory including coupling to 3D equilibria.

3) A 3D time-dependent non-linear MHD code that is capable of modeling the entire plasma volume including dynamics within the injectors. Although HIT-SI was the motivation for, and experiment studied in this research, the tools and methods developed are general — allowing their application to a broad range of magnetic confinement experiments. These tools constitute a significant advance for modeling plasma dynamics in devices with complex boundary geometries.

TABLE OF CONTENTS

	Page
List of Figures	v
List of Tables	xiii
Chapter 1: Introduction and Background	1
1.1 Magnetic Confinement in Fusion Energy	1
1.1.1 Nuclear Fusion	1
1.1.2 Magnetic Confinement	1
1.1.3 Current Drive in Closed Magnetic Topologies	3
1.2 Helicity Injection and Current Drive in HIT-SI	4
1.2.1 Magnetic Helicity and the Spheromak	4
1.2.2 The HIT-SI Current Drive Experiment	7
1.2.3 Diagnostics	11
1.3 Plasma Modeling of Confinement Concepts	12
1.3.1 Magneto-HydroDynamics	13
1.3.2 Finite Element Methods	15
1.3.3 Existing Simulation Tools for HIT-SI	17
1.4 Open Questions and Motivation	20
1.4.1 Mechanics of Current Drive in HIT-SI	20
1.4.2 Limitations of Current Numerical Models	24
1.4.3 Design Considerations for Future SIHI Experiments	27
1.5 Outline	27
Chapter 2: Numerical Tools	29
2.1 The PSI-TRI Code	31
2.1.1 Grad-Shafranov Equilibrium	32
2.1.2 Marginally Stable Mercier Pressure Profiles	34

2.1.3	Experimental Fitting and Reconstruction	35
2.1.4	DCON Interface	38
2.2	The PSI-TET Code	39
2.2.1	Mesh Construction	40
2.2.2	Geometric Mapping	44
2.2.3	Lagrange Basis Set	46
2.2.4	Nedelec Basis Set	49
2.2.5	Multi-grid preconditioning	59
Chapter 3:	3D Ideal MHD Equilibria	66
3.1	Composite Taylor States and Uniform λ Equilibria	67
3.1.1	Solution Method and Scalability	69
3.1.2	Composite States in HIT-SI	71
3.2	Ideal MHD Equilibrium with Variable λ	72
3.2.1	Flux Surface Identification in 3D	73
3.2.2	Eigenmodes and Plasma Response	77
3.2.3	Solution Method	79
3.2.4	Extension to Non-Zero Plasma Beta	83
3.3	Application to Driven HIT-SI Equilibria	84
3.3.1	Discussion	91
3.4	Discussion and Future Work	92
Chapter 4:	Reduced Order Models for SIHI Design	94
4.1	Current Profiles Produced by IDCD	95
4.1.1	Predicted Equilibria in HIT-SI	96
4.1.2	Evaluating Dynamo Force and Dissipation	97
4.1.3	Predicted Current Evolution	100
4.2	Geometry Optimization and Design for SIHI	102
4.2.1	Predictions for HIT-SI3	103
4.2.2	Finding Matched λ -Profiles	103
4.2.3	Discussion	106
Chapter 5:	3D Non-Linear Hall-MHD	108
5.1	Hall-MHD with Nedelec Elements	109

5.1.1	Vector Sub-space Decomposition	109
5.1.2	Divergence Error with a Nedelec Basis	110
5.2	Verification and Benchmarking	112
5.2.1	Aflven-Whistler dispersion relation	113
5.2.2	Tilt Unstable Spheromak	114
5.3	Multi-Grid Solvers	122
5.3.1	Comments on Scalability	123
5.3.2	Interpolation and Restriction	124
5.4	Modeling The HIT-SI Wall	125
5.4.1	Numerical Formulation with Nedelec Elements	125
5.4.2	Flux Injection and Current Drive	127
5.4.3	Sheath Model	129
5.5	HIT-SI Validation	130
5.5.1	Resolution Convergence Study	132
5.5.2	Internal Magnetic Structure	136
5.5.3	Surface Magnetic Structure (BD)	138
5.5.4	Discussion	155
5.6	HIT-SI Physics Results	156
5.6.1	Injector Dynamics	159
5.6.2	Toroidal Current Scaling	166
5.7	Density Evolution and Transport	170
5.7.1	Full Hall-MHD	171
5.7.2	Initial HIT-SI Results	172
5.8	Future Work	177
Bibliography		180
Appendix A: Mesh Generation Interfaces		189
A.1	T3D	189
A.1.1	HIT-SI Mesh	190
A.2	CUBIT	198
A.2.1	HIT-SI3 Mesh	199

Appendix B: High Order Tetrahedra	205
B.1 Barycentric coordinates	205
B.2 First derivatives	206
B.3 Second derivatives	207
Appendix C: High Order Nedelec Elements	208
C.1 Vertex Elements	209
C.2 Edge Elements	209
C.3 Face Elements	210
C.4 Cell Elements	211
C.5 Generation	212
Appendix D: 3D MHD Formulation	214
D.1 Reduced Hall-MHD	214
D.1.1 Non-Linear Metric	214
D.1.2 Jacobian	215
D.2 Full Hall-MHD	215
D.2.1 Non-Linear Metric	216
D.2.2 Jacobian	217

LIST OF FIGURES

Figure Number	Page
1.1 Coil configuration for a simple tokamak configuration.	3
1.2 2 isolated but linked vacuum flux tubes, yielding a helicity $K = 2\psi_1\psi_2$	5
1.3 2D slice of a spheromak magnetic configuration with a spherical last closed flux surface.	7
1.4 The HIT-SI device showing the flux coils(orange) and voltage coils(green) on each injector. The voltage coil is shown here as a solid core transformer for clarity, however the actual transformer operates with an air core.	8
1.5 I_{tor} and n=1 mode amplitude for a representative high performance shot with dynamic periods marked. Red) Initial period dominated by n=1 injector fields. Blue) spheromak formation associated with Taylor relaxation. Green) Quiescent current growth marked by low fluctuation levels.	9
1.6 Flux tube trace of a force-free injector state from the top(left) and the side(right) illustrating the odd symmetry of imposed fluctuations. . .	10
1.7 Diagram showing the location of surface magnetic diagnostics on the HIT-SI device. Each probe location samples the magnetic field in all three directions.	12
1.8 Convergence of spurious numerical diffusion with different polynomial orders for a benchmark anisotropic thermal conduction problem[1]. . .	18
1.9 Axisymmetric grid used in NIMROD simulations, the colored annuli correspond to normal flux boundary conditions used to approximate the injector volumes.	19
1.10 Comparison of I_{tor} evolution for 4 NIMROD simulations with different background η and physics models to experimental data[2].	20
1.11 Schematic of coaxial helicity injection current drive of a spheromak using a Marshall gun.	21

1.12	Schematic of shear induced in magnetic fluctuations by differential electron flow. In the presence of a gradient in electron velocity, V_e , an applied fluctuation will be sheared from its initial state, δB_f , to a distorted state, δB_d , by the flow.	22
1.13	Comparison of toroidal and poloidal magnetic fields along the HIT-SI internal probe at two different injector frequencies. Shown are probe measurements (stars) and the result of an axisymmetric fit (solid lines) with a linear pressure profile and two- λ current profile.	24
1.14	Comparison of q-profiles from NIMROD simulations, Taylor theory, and axisymmetric fits to experimental data (figure 1.13).	25
2.1	Documentation block	30
2.2	Resulting documentation	30
2.3	PSI-TRI solution of the Taylor state in a future SIHI driven device(HIT-POP), shown is the computational grid and a pseudocolor plot of the poloidal flux	31
2.4	Pressure profile and its derivative for an equilibrium in HIT-SI defined using splines fit to the Mercier criterion.	35
2.5	2D equilibria in HIT-SI showing the difference between zero beta (dashed) and marginal Mercier stability (solid) with constant λ . Only a slight Shafranov shift and modification of the q-profile is seen due to the low β -limit with uniform λ in HIT-SI, $\sim 2.8\%$	36
2.6	Example axisymmetric equilibrium fit to experimental signals. Shown are toroidal (top) and poloidal (bottom) magnetic fields from the internal magnetic probe (red) used as fitting constraints and the model magnetic profiles (blue) for the initial guess (a) and resulting fit (b) equilibria.	37
2.7	Normal displacement for the most unstable eigenmode as determined by DCON for an equilibrium in a cylinder computed by PSI-TRI. . .	38
2.8	Cross-section of HIT-SI injector mouth showing mesh refinement. . .	41
2.9	Child cells produced by uniform refinement a single tetrahedron. . . .	42
2.10	Sub-domains from parallel decomposition of a HIT-SI mesh with 72 MPI tasks. Note that some colors are repeated.	42
2.11	Average number of inter-task connections per MPI task for a cylindrical mesh with different numbers of MPI tasks.	43

2.12	Example mapping from the unit tetrahedron (a) in logical space to the physical cell (b). A quadratic mapping is shown approximating an octant of a sphere with a single cell.	45
2.13	Comparison of grid boundary between linear (a) and quadratic (b) mappings for a small curved feature in the HIT-SI mesh. Shown is the tessellated mesh and $ B $ for the second force-free eigenmode computed using a cubic $H^1(\text{Curl})$ basis on each mesh.	47
2.14	Convergence verification for the Lagrange basis set on a Poisson test case in a unit cube.	48
2.15	Demonstration of interpolation between Lagrange basis representations under geometric (a) and polynomial (b) refinement.	49
2.16	Example scalar (left), H^0 , and vector (right), H^1 , Nedelec elements in 2D corresponding to the lower edge of the triangle. The vector element results from taking the gradient of the scalar function.	53
2.17	Convergence verification for the Nedelec basis set on a vector Poisson test case in a unit cube.	55
2.18	Demonstration of edge interpolation between $H^1(\text{Curl})$ basis representations under geometric refinement.	56
2.19	Demonstration of flux constrained interpolation between $H^1(\text{Curl})$ basis representations under geometric refinement.	57
2.20	Demonstration of node constrained interpolation between H^0 basis representations under geometric refinement.	58
2.21	Diagram of 4 level V-Cycle iteration.	61
2.22	Iteration scaling for a scalar Lagrange tests case with Jacobi preconditioning.	62
2.23	Iteration scaling for a scalar Lagrange tests case with mult-grid preconditioning.	63
2.24	Iteration scaling for a vector $H^1(\text{Curl})$ tests case with Jacobi preconditioning.	64
2.25	Iteration scaling for a vector $H^1(\text{Curl})$ tests case with multi-grid preconditioning.	65
3.1	Magnetic structure of the first two force-free eigenstates in HIT-SI.	68
3.2	Iteration scaling for Taylor state calculation in a cylinder with Jacobi preconditioning.	70
3.3	Iteration scaling for Taylor state calculation in a cylinder with multi-grid preconditioning.	70

3.4	Injector vacuum magnetic field in HIT-SI computed using multi-connected jump method.	72
3.5	Magnetic structure for a Taylor state with flux amplification of 6. . .	73
3.6	Thermal diffusion solve on a Taylor state in a oblate cylinder, comparing the computed surface (red), magnetic flux tubes (gray), and Grad-Shafranov flux surfaces (2D contours).	74
3.7	Comparison of the computed flux coordinate and the location of a large magnetic island in the HIT-SI Taylor state for a range of transport ratios.	76
3.8	Reduction of temperature error with representation order.	76
3.9	Comparison of computed flux surfaces (color) and Poincaré section (black) of the magnetic field in the vicinity of a magnetic island in the HIT-SI Taylor state.	77
3.10	Toroidal flux generation due to eigenvalue resonance in HIT-SI plasma response. Computed flux (red) agrees well with the theoretical scaling (blue).	79
3.11	Force-free field error in plasma response calculation for different generated toroidal fluxes in HIT-SI with fixed injector flux of 1 mWb. . .	80
3.12	Diagram of iteration cycle used in 3D equilibrium solver.	81
3.13	Comparison of flux surfaces between pressureless (dashed) and $\langle\beta\rangle_{vol} = 2.2\%$ (solid) equilibria in a 1x1 cylinder. Flux surfaces for the case with pressure were chosen to match the pressureless surfaces on the outboard side.	85
3.14	Comparison of magnetic field magnitude between pressureless (dashed) and $\langle\beta\rangle_{vol} = 2.2\%$ (solid) equilibria in a 1x1 cylinder.	86
3.15	Ideal MHD equilibrium state in HIT-SI at maximum λ_{Inj} for a flux ratio of 6 with a two- λ model.	88
3.16	Comparison of flux contours and Poincaré section of open field structure on a poloidal slice at a toroidal position of $\pi/2$ for the equilibrium state shown in figure 3.15.	89
3.17	Ideal MHD equilibrium state in HIT-SI at maximum λ_{Inj} for a flux ratio of 20 with a two- λ model.	90
3.18	Maximum λ_{inj} for a range of flux amplifications in HIT-SI with a two- λ model.	92
4.1	Evolution of B_r in time at different major radii along the HIT-SI internal magnetic probe. Note that each radial position is offset by 40 mT to facilitate display on a single plot.	96

4.2	Schematic of decomposition of the expected injector current path due to fast reconnection. The toroidal current (not shown) is oriented in the clockwise direction.	97
4.3	Mean field comparison at a characteristic time during high performance operation in HIT-SI (Shot 122385). The radial position of the separatrix is marked by the orange dashed line.	98
4.4	Magnitude of δB_{\perp}^2 computed from a uniform $\lambda = \lambda_{Taylor}$ force-free equilibrium relative to a flux surface of the axisymmetric Taylor state in HIT-SI. The flux surface is shown from below, with the X-Injector mouths visible as the higher magnitude footprints on the flux surface.	99
4.5	Comparison of IDCD dissipation and surface drive profiles for the empirically predicted mean field equilibrium in HIT-SI.	100
4.6	Path of mean field equilibrium evolution through parameter space in time when using equation 4.4 to model the toroidal current.	101
4.7	Comparison of I_{Tor} between experimental data (black), the IDCD empirical reduced model (red) and the IDCD uniform $\lambda = \lambda_{Taylor}$ force-free model (red) across multiple shots.	102
4.8	Example force-free injector equilibrium state in HIT-SI3 with the 3-phase configuration.	104
4.9	Comparison of IDCD surface drive profiles for HIT-SI (blue), the HIT-SI3 3-phase configuration (solid green), and the HIT-SI3 in-phase configuration (dashed green) to the dissipation profile in steady state.	104
4.10	Comparison of IDCD dissipation and surface drive profiles for mean field equilibria with uniform (dashed) and optimized (solid) λ profiles.	105
4.11	Comparison of flux surface quantities for mean field equilibria with uniform (dashed) and optimized (solid) λ profiles.	106
4.12	Schematic of geometry and profile optimization work flow for SIHI driven devices using IDCD to define drive and profile constraints.	107
5.1	Comparison of propagation frequencies for waves in reduced Hall MHD from numerical tests in PSI-TET (squares) with the analytic dispersion function (solid).	115
5.2	Ideal MHD growth rate of the unstable tilting mode in a cylinder as a function of aspect ratio as computed by Bondeson <i>et al.</i> [3].	116
5.3	Convergence of toroidal flux (a) and current (b) for the unstable cylindrical tilt benchmark at different grid resolutions and polynomial orders.	118

5.4	Convergence of total system energy for the unstable cylindrical tilt benchmark at different grid resolutions and polynomial orders.	118
5.5	Point divergence error (blue, green) compared to magnetic field strength (red) along parallel chords at 0.19ms in the ($N_g = 2$, $np = 2$) simulation from the unstable cylindrical tilt benchmark.	119
5.6	Comparison of divergence error with kinetic energy in the unstable cylindrical tilt benchmark.	121
5.7	Divergence error computed with local jump (a) and global energy (b) metrics for the unstable cylindrical tilt benchmark at different grid resolutions and polynomial orders.	122
5.8	Comparison of solution cost for the HIT-SI validation case with different mesh resolutions and polynomial orders.	124
5.9	Triangle strip used to define the composite basis function for the flux circuit boundary condition on a single injector for HIT-SI simulations.	127
5.10	Slice planes on a single HIT-SI3 injector used to create planar curves and resulting triangle strips for simulating the flux and voltage drive circuits.	128
5.11	Magnetic field early in time in a HIT-SI simulation due to flux injection and current drive boundary conditions on the injector. In-plane field is shown as black arrows and a pseudocolor plot is shaded by the out of plane field strength.	129
5.12	Comparison of I_{Tor} from shot 122385 and PSI-TET validation case at 14.5 kHz.	133
5.13	Comparison of injector current waveforms from shot 122385 and PSI-TET validation case used to align time bases for probe comparisons.	133
5.14	Convergence of toroidal flux evolution in HIT-SI validation case with different mesh resolutions and polynomial orders.	135
5.15	Convergence of toroidal current evolution in HIT-SI validation case with different mesh resolutions and polynomial orders.	135
5.16	Magnetic field comparison along the internal magnetic probe for the HIT-SI validation case.	138
5.17	Signal energies from independent BD analysis of experimental data (shot 122385) and PSI-TET and NIMROD validation cases.	143
5.18	Mode 1 chrono from independent BD analysis of experimental data (shot 122385) and PSI-TET and NIMROD validation cases.	144

5.19	Modes 2 and 3 chronos from independent BD analysis of experimental data (shot 122385) and PSI-TET and NIMROD validation cases. . .	144
5.20	Mode 4 chrono from independent BD analysis of experimental data (shot 122385) and PSI-TET and NIMROD validation cases.	146
5.21	Comparison of mode 4 topo from independent BD analysis of PSI-TET and NIMROD validation cases with experimental data (shot 122385).	147
5.22	Topo correlations between simulation and experimental data (shot 122385) from independent BD analysis.	149
5.23	Chrono correlations between simulation and experimental data (shot 122385) from independent BD analysis.	149
5.24	Mode weights from reference BD analysis of experimental data (shot 122385) and PSI-TET and NIMROD validation cases.	151
5.25	Mode 1 chrono from reference BD analysis of experimental data (shot 122385) and PSI-TET and NIMROD validation cases.	152
5.26	Modes 2 and 3 chronos from reference BD analysis of experimental data (shot 122385) and PSI-TET and NIMROD validation cases. . .	153
5.27	Mode 4 chrono from reference BD analysis of experimental data (shot 122385) and PSI-TET and NIMROD validation cases.	154
5.28	Chrono correlations between simulation and experimental data (shot 122385) from reference BD analysis.	155
5.29	Magnetic field in the X-Z plane from a PSI-TET reduced Hall-MHD simulation before formation of a toroidal current (0.052 ms).	157
5.30	Magnetic field in the X-Z plane from a PSI-TET reduced Hall-MHD simulation following formation of a toroidal current (0.5 ms).	158
5.31	Comparison of flux and current waveforms in the Y-injector from PSI-TET simulations at 53.5 kHz with (a) and without (b) flux control. . .	161
5.32	Comparison of flux coil voltage and current for a PSI-TET simulation at 53.5 kHz with flux control.	162
5.33	Magnetic field in the X-Z plane from a PSI-TET reduced Hall-MHD simulation showing large fluctuations associated with field reversal in the X-injector (1.008 ms).	164
5.34	Magnetic field in the X-Z plane from a PSI-TET reduced Hall-MHD simulation showing smaller fluctuations associated with field reversal in the Y-injector (1.018 ms) when compared to the X-injector (figure 5.33).	165

5.35	Plasma flow velocity in the X-Z plane from a PSI-TET reduced Hall-MHD simulation near peak injector flux (1.008 ms). Flow is visible away from each injector mouth in the in-plane velocity field (black vectors), with an associated rotation in the y-velocity (psuedo-color). More coherent and collimated axial flow is observed near the 0° mouth (right).	167
5.36	Plasma flow velocity in the X-Z plane from a PSI-TET reduced Hall-MHD simulation 1/2 injector cycle after figure 5.35. The collimated axial flows are now observed with the 180° mouth (left).	168
5.37	Comparison of toroidal currents from simulations at different resistivities to model current evolutions assuming impedance scaling for the injector voltage circuit of the form $\sim j/n$	170
5.38	Comparison of I_{Tor} from shot 122385 and PSI-TET full Hall-MHD simulation at 14.5 kHz.	174
5.39	Magnetic field along the internal magnetic probe for a PSI-TET full Hall-MHD simulation of HIT-SI compared to experiment signals (shot 122385).	175
5.40	Ion temperature in the X-Z plane from a PSI-TET full Hall-MHD simulation showing heating and thermal connection through the injectors.	176
5.41	Ion temperature contours in 3D from a PSI-TET full Hall-MHD simulation showing thermal connection through the injectors.	177
A.1	Example of second order Beziér curve(a) and surface(b) showing the control points.	190
B.1	Surfaces of constant ξ_1 in the unit tetrahedron.	206

LIST OF TABLES

Table Number	Page
5.1 Run parameters and settings for cylinder tilt benchmark	117
5.2 Run parameters and settings for HIT-SI validation case	132
C.1 First five Legendre based functions used for H^1 element construction	209

ACKNOWLEDGMENTS

The work presented here would not have been possible without the assistance and supporting work of many people. I would like to thank my adviser Professor Thomas Jarboe for giving me the opportunity to work with the HIT-SI and PSI-Center groups and for guiding me through the world of plasma physics. I would also like to thank Dr. George Marklin, whose teaching and insights over the past several years have been invaluable to my work and education in numerical methods and plasma physics. I also thank Professor Brian Nelson for helping me with editing this document and for all his suggestions throughout this research. Experimental comparisons would not have been possible without the efforts of the past and present HIT-SI staff and graduate students. I wish to especially thank Cihan Akcay, Aaron Hossack, Kyle Morgan, Derek Sutherland, Brian Victor and Jonathan Wrobel for our numerous discussions over the last several years. My education in the world of computational plasma physics would not have been complete without the involvement and insight of the other members of the PSI-Center. I would like to thank Alan Glasser, Eric Held, Jeong-Young Ji, Charlson Kim, Weston Lowrie, Slava Lukin, Eric Meier, Richard Milroy, Peter Norgaard, Uri Shumlak and Carl Sovinec for many helpful discussions. I would also like to thank the faculty and staff of the Aeronautics & Astronautics department at the University of Washington for the opportunity to work and learn in a truly wonderful environment as both an undergraduate and graduate student. Finally, I would like to thank my family for supporting me along this path and in every endeavor I have undertaken over these last 26 years, and Cullen Heaney for keeping me sane, happy and laughing every day of my life.

This research was supported by the United States Department of Energy and some of the simulations presented were performed at the National Energy Research Scientific Computing Center at Lawrence Berkeley National Laboratory.


DEDICATION

To my parents

Chapter 1

INTRODUCTION AND BACKGROUND**1.1 Magnetic Confinement in Fusion Energy***1.1.1 Nuclear Fusion*

As a clean successor to nuclear fission, nuclear fusion has been pursued as a next generation energy source for decades. Nuclear fusion, the process that powers the sun, requires heating hydrogen gas to very high temperatures. High thermal energy ionizes the gas and provides enough kinetic energy for colliding nuclei to overcome the coulomb barrier and fuse. On earth, development of fusion as an energy source focuses on the most readily accessible reaction that involves Deuterium and Tritium, both isotopes of Hydrogen, to produce Helium and energy.

*1.1.2 Magnetic Confinement*

In order to reach the temperatures required for this reaction to occur with reasonable frequency, plasma must be insulated from material surfaces. For this reason, magnetic confinement is studied as a means of reducing thermal conduction from the plasma to the wall. Magnetic fields are effective in confining a plasma by virtue of the Lorentz force, equation 1.2, which acts on charged particles. In the presence of a magnetic field this forces deflects the component of a charged particle's velocity transverse to the field causing them to orbit magnetic field lines as they move through space.

$$\mathbf{F}_\alpha = q_\alpha [\mathbf{E} + \mathbf{v}_\alpha \times \mathbf{B}] \quad (1.2)$$

This force creates an anisotropy in both thermal and particle transport in the plasma by limiting the ability of plasma constituents, ions and electrons, to move in the cross-field direction. By exploiting this fact, closed magnetic configurations can be created in order to confine a plasma and increase its temperature to the energies required for fusion reactions to occur. The level of confinement is determined by the magnetic field energy required to stably contain a given amount of thermal energy, this quantity β is defined in equation 1.3 in a volume averaged sense. Many other normalizations exist, however all appropriate definitions represent the “magnetic efficiency” of a given configuration. Different topologies have their own inherent limit for β .

$$\beta = \frac{\int p \, dV}{\int \frac{B^2}{2\mu_0} \, dV} \quad (1.3)$$

Limitations Due to Instability

Classically, closed magnetic structures have very high degrees of anisotropy that can support large gradients in temperature and pressure. However, in practice instabilities exist that prevent these high β 's from being reached. The upper bound on β is dependent on the specifics of a magnetic configuration, in particular the variation in magnetic energy and tension relative to thermal pressure. This subject is discussed in greater detail in chapter 2.1.2.

The Tokamak

The current leading candidate for magnetic confinement fusion is the Tokamak. Figure 1.1 shows a representative diagram of the arrangement of field coils used to produce a Tokamak configuration. A high toroidal field is produced by large coils that encircle the plasma. Current is then driven in the plasma producing a poloidal field which is supported by field coils on the outboard side of the torus.




Figure 1.1: Coil configuration for a simple tokamak configuration.

The large Toroidal field provides stability to the plasma while the poloidal field created by internal currents provides confinement so that the plasma may be heated.

1.1.3 Current Drive in Closed Magnetic Topologies

As with the Tokamak, most magnetic confinement concepts require current drive inside the plasma to maintain components of the magnetic field produced by internal currents. This creates a difficulty in configurations with closed field topologies as a voltage drop cannot be statically applied as in a DC circuit. Traditionally, transformer action has been used to produce a loop voltage as with the central solenoid shown in 1.1 above. This method is inherently transient as the plasma current, which must be unidirectional and large, is related to the time integrated loop voltage. As a result, the transformer voltage can only be swung in one direction with a range that is limited by physical restrictions on the maximum applied voltage.

For use as a power plant, good confinement must be maintained in or near steady state, eliminating the transformer as a viable current drive method in reactor scenarios. Currently, high energy neutral beam injection or RF based current drive methods

are used for long duration operation and reactor designs. However, these methods suffer from inherently low efficiencies, especially at the high temperatures required for fusion reactions. Therefore, it would be desirable to improve on the efficiency of conventional current drive methods for reactor concepts.

1.2 Helicity Injection and Current Drive in HIT-SI

One method that promises to provide higher efficiencies than present current drive methods is helicity injection and dynamo current drive. This method allows injection of magnetic energy/helicity via transformer current drive in a more convenient and somewhat isolated sub volume of the plasma. Dynamo action, produced by coupling of the magnetic fields and plasma fluid, then transports this energy/helicity into the core plasma where it can replace magnetic energy lost to resistive decay. Much higher efficiency can be achieved with this method as the initial injection of magnetic energy into the system may be done by transformer action.

This method of sustaining plasma currents relies on two concepts: 1) the conservation of magnetic helicity in a magnetized plasma. 2) current penetration due to plasma fluid interaction. Also, note that although “helicity injection” and “dynamo current drive” are often used interchangeably in the context of current drive, they refer to different aspects of a common process.

1.2.1 Magnetic Helicity and the Spheromak

Magnetic helicity is a measure of the self-linkage of magnetic flux and is the best constant of motion in a magnetized plasma. Analytically, the helicity in a volume Ω can be computed using equation 1.4, where A is the vector potential such that $B = \nabla \times A$, when magnetic flux does not pierce the surface bounding Ω .

$$K = \int_{\Omega} A \cdot B \, dV \tag{1.4}$$




Figure 1.2: 2 isolated but linked vacuum flux tubes, yielding a helicity $K = 2\psi_1\psi_2$

The physical interpretation of magnetic helicity as self-linkage or twistedness of magnetic field can be seen by considering the idealized case of two isolated but linked magnetic flux tubes, figure 1.2. If each flux tube has a magnetic flux of ψ_1 and ψ_2 respectively, this integral evaluates to $K = 2\psi_1\psi_2$. This is consistent with the physical picture as flux ψ_1 links ψ_2 and vice versa.

Relative Helicity

When the volume of interest has magnetic flux crossing its boundary, the gauge ambiguity present in the vector potential causes this definition to be ill-defined. In this case, which is of interest when helicity is being injected, a new gauge invariant definition must be introduced[4]. Using a vacuum field, B_{vac} , whose normal component matches that of B everywhere on the boundary surface an appropriate quantity can be defined, called relative helicity, as in equation 1.5.

$$K_{rel} = \int_{\Omega} (A + A_{vac}) \cdot (B - B_{vac}) dV \quad (1.5)$$

Relaxation and the Taylor State

A plasma with sufficient energy and helicity will relax and reconfigure its magnetic topology through reconnection in order to minimize its internal energy while preserving its helicity. This phenomenon is known as Taylor relaxation[5] and has been observed in a wide variety of confinement devices. Fixing the magnetic helicity during magnetic relaxation amounts to a constraint on the topological change of the magnetic fields as energy is dissipated through magnetic reconnection. This effect leads to the process of plasma self-organization which is exploited in the spheromak and Reverse Field Pinch(RFP) confinement concepts to reduce or eliminate external field coils.

The Spheromak

If a magnetized plasma with sufficient energy and helicity is contained in a conducting shell, commonly called a flux conserver, it will relax to a state of minimum energy with fixed helicity – according to Taylor theory[5]. This state is known as the Taylor state and corresponds to a stable configuration which contains no remaining free energy. In a suitable, simply connected, volume this state produces a spheromak[6] with a structure which resembles that of a Tokamak configuration where the toroidal field is produced entirely by internal plasma currents. An example of the magnetic structure of a spheromak[7] in a spherical volume is shown in figure 1.3.

The resulting magnetic field is force-free, supports no pressure gradients, and is determined by the lowest eigenstate of equation 1.6. This equation comes from considering the pressureless form of equation 1.12 and constitutes an eigenvalue problem for “natural” plasma states of a given flux conserver. Where λ is the ratio $\frac{\mu_0 J_{\parallel}}{B}$ which gives a measure of the parallel current on field lines.

$$\mu_0 J = \lambda B \tag{1.6}$$




Figure 1.3: 2D slice of a spheromak magnetic configuration with a spherical last closed flux surface.

The spheromak is a desirable confinement configuration for several reasons, principally: 1) It eliminates the need for external toroidal field coils, which in turn allows a simply connected geometry 2) The additional plasma current provides sufficient heating via ohmic dissipation to eliminate auxiliary heating requirements to reach fusion conditions. The additional plasma current, although beneficial, increases the required current drive power in steady state. This makes efficient current drive a necessity for the spheromak concept.

1.2.2 The HIT-SI Current Drive Experiment

The Helicity Injected Torus with Steady Inductive helicity injection(HIT-SI), figure 1.4, studies steady-state current drive in a simply connected geometry using Steady Inductive Helicity Injection(SIHI)[8, 9, 10]. Toroidal current is driven by means of two semi-toroidal helicity injectors, which inject helical magnetic field into a copper flux conserver. This helical magnetic field then relaxes to produce a spheromak that is grown and sustained by imposed fluctuations applied by continued operation of the




Figure 1.4: The HIT-SI device showing the flux coils(orange) and voltage coils(green) on each injector. The voltage coil is shown here as a solid core transformer for clarity, however the actual transformer operates with an air core.

injectors.

Recent operations have introduced a high performance regime[11] characterized by high current amplification, the ratio of toroidal current in the confinement volume to injector current. A characteristic shot for this regime is shown in figure 1.5. This operation is also marked by a deviation from the previously successful Taylor model for HIT-SI, which modeled the fields as a quasi-static force-free equilibrium. Increased flux and current amplification, with respect to the injector, are of interest for this work in capturing the new physics associated with deviation from the Taylor model.

Steady Inductive Helicity Injection

The injectors on HIT-SI are nominally driven to produce a steady rate of helicity injection. Taking the time derivative of equation 1.5 yields an expression for the rate of change of helicity in a volume[12, 6], equation 1.7. The first term in this equation corresponds to helicity growth or dissipation within the plasma due to internal electric fields $\vec{E} = -\frac{\partial \vec{A}}{\partial t}$ and the second term results from helicity injection due to vacuum




Figure 1.5: I_{tor} and $n=1$ mode amplitude for a representative high performance shot with dynamic periods marked. Red) Initial period dominated by $n=1$ injector fields. Blue) spheromak formation associated with Taylor relaxation. Green) Quiescent current growth marked by low fluctuation levels.




Figure 1.6: Flux tube trace of a force-free injector state from the top(left) and the side(right) illustrating the odd symmetry of imposed fluctuations.

electric and magnetic fields. When the flux through the boundary is fixed the second term can be reduced to a voltage drop, $\Delta V = \int \vec{E}_{vac} \cdot d\vec{l}$, along vacuum flux.

$$\frac{\partial K_{rel}}{\partial t} = -2 \left[\int_{\Omega} \vec{E} \cdot \vec{B} dV - \int_{\Omega} \vec{E}_{vac} \cdot \vec{B}_{vac} dV \right] \quad (1.7)$$

HIT-SI is driven inductively by two sets of coils on each injector as seen in figure 1.4, the flux coil (shown in orange) and the voltage coil (shown in green). The voltage coil is shown here wrapped around an iron core for clarity only and the voltage coil on the machine uses an air core. The flux coil injects locally toroidal magnetic field in each injector while the voltage coil produces an electric field parallel to this, generating helical magnetic field by driving current. This produces helical field in the injector which links back through the main confinement volume, as shown in figure 1.6.

The flux and voltage circuits on each injector are driven in phase with a sinusoidal

wave form in time. The two injectors, labeled x and y, are driven 90° out of phase with each other in time to produce a steady rate of helicity injection. Considering just the external injection term in equation 1.7, the helicity injection rate is shown to be constant, equation 1.8. V_0 and Φ_0 are the loop voltage and flux induced by each injectors voltage and flux coils.

$$\frac{\partial K_{inj}}{\partial t} = 2(V_0 \sin(\omega t) \Phi_0 \sin(\omega t) + V_0 \cos(\omega t) \Phi_0 \cos(\omega t)) = 2V_0 \Phi_0 \quad (1.8)$$

1.2.3 Diagnostics

HIT-SI features an extensive set of magnetic diagnostics, which include an array of 192 magnetic probes mounted in the surface of the flux conserver and a 3-axis internal probe. The surface probes are arranged in poloidal arrays of 16 3-axis probes at 4 different toroidal locations around the device as shown in figure 1.7. In addition to measuring individual magnetic signals each poloidal array is used to construct an amperian loop to compute the local toroidal current passing through each array. The four local currents are then averaged to produce a measurement of the symmetric toroidal current with fast time response[13]. An additional 32 probes are also placed in the diagnostic gap at evenly spaced toroidal locations and may be used for toroidal Fourier mode analysis. The internal magnetic probe consists of three independent stems that may be inserted up to a major radius of 33cm[14]. For this work only the middle stem will be used for comparisons, which consists of 17 3-axis probes each spaced approximately 1.17 cm apart, when measured center to center. Although the internal probe perturbs the plasma, no significant difference has been seen between shots taken with and without the internal probe stems inserted.

Figure 1.7 also contains a reference for the coordinate system used in HIT-SI. Cartesian and daughter cylindrical coordinate systems are used with the z-axis oriented along the symmetry axis of the main flux conserver. The cylindrical coordinate

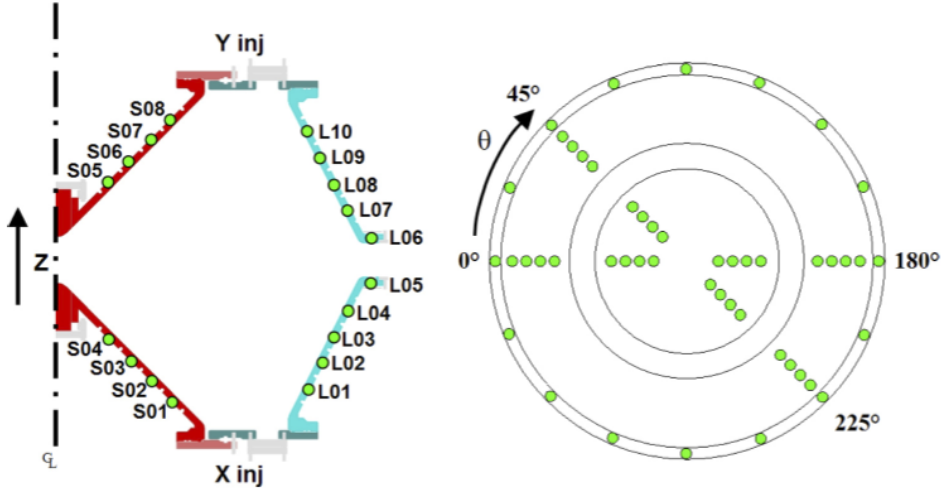


Figure 1.7: Diagram showing the location of surface magnetic diagnostics on the HIT-SI device. Each probe location samples the magnetic field in all three directions.

system (r, θ, z) is then defined with the same z -axis and $\theta = 0$ along the positive x -axis. Each of the two injectors are oriented so that their locally toroidal directions are along their respective axes, x -axis for the X-INJ and y -axis for the Y-INJ. This puts the X-INJ on the “bottom” of the machine, at negative z , and the Y-INJ on the “top”, at positive z . Finally, we define the sign of the i -th injector’s phase as in equation 1.9, where \vec{B}_v^i is the vacuum field at the center of the machine and \hat{x}_i is the unit vector for the corresponding cartesian axis.

$$\text{SIGN} \left(\vec{B}_v^i \cdot \hat{x}_i \right) \quad (1.9)$$

1.3 Plasma Modeling of Confinement Concepts

Magnetic confinement experiments approaching the scale of a fusion reactor are required to validate the physics basis for a commercial device. However, experiments of

this size are costly and require years of development and construction. Additionally, diagnostic access becomes more difficult as the plasma reaches higher performance and temperatures. As a result, the development of predictive computational models for fusion relevant plasmas has become an intense area of research. Validated, accurate, and efficient models of essential plasma phenomena can provide insight into existing experiments as well as guidance in designing future experiments.

1.3.1 Magneto-HydroDynamics

The equations of Magneto-HydroDynamics (MHD) are the most widely studied method of modeling plasma behavior in confinement devices. These equations represent the plasma as a single continuum fluid with the effects of relative electron-ion motion represented by modifications to the magnetic induction equation, Ohm's law. Many variations of the MHD model are used, resulting from further physics simplification. Extended MHD is one of the most complete sets and is governed by equations 1.10a – 1.10d, corresponding to the continuity, momentum, energy and induction equations respectively.

$$\frac{\partial \rho}{\partial t} + \nabla \cdot \rho u = 0 \quad (1.10a)$$

$$\rho \left[\frac{\partial u}{\partial t} + u \cdot \nabla u \right] = J \times B - \nabla (p_i + p_e) - \nabla \cdot \Pi \quad (1.10b)$$

$$\frac{n}{\gamma - 1} \left(\frac{\partial T}{\partial t} + u \cdot \nabla T \right) = -p_i \nabla \cdot u - \nabla \cdot q + Q \quad (1.10c)$$

$$\frac{\partial B}{\partial t} = -\nabla \times \left(-u \times B + \frac{1}{ne} (J \times B - \nabla p_e) + \eta J \right) \quad (1.10d)$$

Where the independent variables are the plasma density $\rho = nm_i$, velocity u , temperature T , and magnetic field B . The remaining quantities are the ion and electron pressures $p_i = nkT$ and $p_e = nkT_e$, viscous stress tensor Π , the heat flux

vector q , heat source density Q , ratio of specific heats γ , and resistivity η . In this single fluid MHD model the plasma is assumed to be quasi-neutral eliminating displacement currents, $\mu_0 J = \nabla \times B$.

A number of existing codes have had success in modeling plasma confinement experiments using MHD models. The NIMROD code is one of the most widely used and has been used extensively on tokamak, spheromak and other confinement devices[15, 16, 17, 18, 19]. These time dependent studies have been used to study dynamic physics, usually associated with some important instability.

Reduced Hall-MHD

Due to limited time and computational resources it is often not desirable to solve the full set of equations for extended MHD. In this case a subset of the equations and/or terms in the equations are omitted. This may be done purely for speed considerations when certain dynamics are not expected to be important or if the effect of a given term is to be analyzed. For this work the reduction of interest is reduced Hall-MHD, which can be recovered by assuming uniform density, temperature, and pressure. This reduces equations 1.10a – 1.10d to the equations 1.11a – 1.11b.

$$\rho \left[\frac{\partial u}{\partial t} + u \cdot \nabla u \right] = J \times B - \nabla \cdot \Pi \quad (1.11a)$$

$$\frac{\partial B}{\partial t} = -\nabla \times \left(-u \times B + \frac{1}{ne} J \times B + \eta J \right) \quad (1.11b)$$

Further reductions can be achieved by: 1) omitting the Hall term, $[\frac{1}{ne} J \times B]$, in equation 1.11b producing reduced resistive MHD. 2) omitting both the Hall term and setting $\eta = 0$ producing ideal MHD.

MHD Equilibrium

As most experiments and future reactors are designed to operate in equilibrium, steady state solutions to these equations are also of interest. The simplest of these states occur in equilibrium with no flow where the ideal MHD equations can be reduced to equation 1.12. This result comes from requiring force balance on the fluid as given by equation 1.10b. Force balance requires pressure to be a flux surface quantity, no variation parallel to B , as the $J \times B$ term acts only in the perpendicular direction. This equation and its application will be discussed further in section 2.1.1 and chapter 3.

$$0 = J \times B - \nabla p \quad (1.12)$$

1.3.2 Finite Element Methods

Finite element methods have had success in a wide variety of fields including structural mechanics and fluid dynamics. A finite element expansion combines the benefits of high order spectral methods, while providing geometric flexibility and reducing the extent of element interaction, preserving sparsity in resulting matrices. Similar to a spectral element representation the solution is expanded as a set of basis functions with corresponding weights. However, with finite elements the domain is discretized first into discrete elements (cells). Basis functions are then defined within each element resulting in only localized interaction between elements and a sparse connectivity graph for the resulting linear systems.

$$u(x) = \sum \omega_i \phi_i(x) \quad (1.13)$$

Finite element and spectral methods do not solve the familiar form of most systems of Partial Differential Equations (PDE), the strong form, shown in equation 1.14a for Poisson's equation. Instead finite element methods solve for the minimum of an

equivalent weak form, shown in equation 1.14b for the same Poisson's equation, with respect to a set of suitable test functions ϕ_j^T . For finite element and spectral methods, the generality of this form is reduced by replacing u with the known basis expansion from equation 1.13. A suitable test function space must also be chosen and many different methods have been developed that result in unique test functions. One of the most common methods, and the method of interest for this work, is the Galerkin method, which chooses the test functions space to be the same as the basis set used to expand the solution. Applying this method and the expansion of u to equation 1.14b yields equation 1.14c, which gives a system of equations for the basis weights that minimize the error of the original equation with respect to the norm induced by equation 1.14b. This may then be rewritten as a linear system, equation 1.14d, noting that the resulting operator L_{ij} is a square matrix, which may be solved numerically.

$$\nabla^2 u = S \quad (1.14a)$$

$$(\phi_j^T, \mathbb{L}(u)) = \int_{\Omega} \phi_j^T (\nabla^2 u - S) dV \quad (1.14b)$$

$$\omega_j \int_{\Omega} \phi_i (\nabla^2 \phi_j) dV = \int_{\Omega} \phi_i S(x) dV \quad (1.14c)$$

$$\begin{aligned} L_{ij} \omega_j &= s_i \\ L_{ij} &= \int_{\Omega} \phi_i (\nabla^2 \phi_j) dV \\ s_i &= \int_{\Omega} \phi_i S(x) dV \end{aligned} \quad (1.14d)$$

Weak Versus Strong Form

It is important to note that although the weak and strong form solutions are expected to converge toward each other as the solution becomes better resolved, the weak form

is not fully equivalent to the original PDE for a given resolution. This is especially true near locations where boundary conditions are applied, restricting the solution space further. Enforcement of these conditions can lead to narrow regions where disagreement persists between the forms, even as resolution is increased. Therefore, one should be careful when analyzing source terms directly, as the effective action of sources in these regions may be strongly influenced by the test function space.

Modeling Confinement

For modeling magnetically confined plasmas, finite element methods are desirable as they can accurately resolve the anisotropies that result from the Lorentz force. Figure 1.8 shows the advantages of high order finite elements for problems with large anisotropies as investigated by Meier[1]. The figure shows the amount of spurious numerical diffusivity resulting from inaccurately resolving the direction of anisotropy in a 3D thermal conduction problem. High order representations yield a significantly lower diffusion error than would be achievable with lower orders due to constraints on problem size.

Existing Codes

Over the years a number of codes have had success in modeling plasma confinement experiments using MHD models and finite element methods. The NIMROD[18, 15, 19, 16, 17], HiFi[20, 21, 22, 23], and M3D[24, 25, 26] codes all employ a finite element representation and are some of the most widely used and successful tools available.

1.3.3 Existing Simulation Tools for HIT-SI

To date, modeling efforts on HIT-SI have focused on the use of the NIMROD code[27, 28, 29, 30, 2]. NIMROD is a high order finite element code that uses a block structured mesh of quadrilaterals to represent the poloidal plane with a spectral Fourier




Figure 1.8: Convergence of spurious numerical diffusion with different polynomial orders for a benchmark anisotropic thermal conduction problem[1].

representation in the toroidal direction (for more information on the code see [15]). To satisfy NIMROD's requirement of an axisymmetric geometry, the injector volumes are removed and their effects are modeled by boundary conditions on the upper and lower plates of the flux conserver, figure 1.9.

In order to approximate the insulating coating on the inner surface of the flux conserver, a thin boundary layer of cells is used with a resistivity that is large with respect to the plasma resistivity, usually a factor of $\sim 10^5$. This method has shown success in modeling the formation and sustainment of a spheromak in HIT-SI and results agree well with experimental observations. Figure 1.10 shows a comparison between modeled current in NIMROD using reduced resistive and reduced Hall MHD and an average of several experimental I_{Tor} measurements at an injector frequency of 14.5 kHz. Resistive NIMROD cases are shown in black, while Hall runs are shown in red, with resistivities corresponding to 7.5 eV and 12 eV electron temperatures




Figure 1.9: Axisymmetric grid used in NIMROD simulations, the colored annuli correspond to normal flux boundary conditions used to approximate the injector volumes.

shown as solid and dashed lines respectively. As can be seen Hall MHD is required to reproduce the salient details of the toroidal current formation and growth, which purely resistive models do not capture. Note that the linear ramp seen in the experimental ensemble is primarily due to loading of the driving circuits in response to plasma breakdown. This causes the power and helicity injection to drop, requiring it to ramp back up, following breakdown ($t \sim 0$) in the experiment, while NIMROD uses a constant injection rate equal to the experimental value late in time.

Although this method has shown good agreement with experimental data, there are some disagreements between computational results and data. This will be discussed more in the following section(1.4). The HiFi code has also been demonstrated on the full HIT-SI geometry[23] using a decaying spheromak, however the flux and voltage drivers were not simulated.




Figure 1.10: Comparison of I_{tor} evolution for 4 NIMROD simulations with different background η and physics models to experimental data[2].

1.4 Open Questions and Motivation

There are currently several open questions relating to the operation of the HIT-SI device and the steady inductive helicity injection current drive technique in general. Present studies are primarily focused on determining the internal characteristics of the generated plasma, including temperature, pressure, and current profiles as well as investigating how magnetic helicity, and thus toroidal current, is transported from the edge region to maintain and grow the spheromak. This section will give a brief overview of these topics and how this work aims to help answer, or develop tools to investigate, those questions.

1.4.1 Mechanics of Current Drive in HIT-SI

Helicity injection current drive has been used on a number of earlier spheromak experiments as well as other device types. However, the exact mechanism of energy transfer




Figure 1.11: Schematic of coaxial helicity injection current drive of a spheromak using a Marshall gun.

and current drive into the sustained plasma has been poorly understood, with Taylor relaxation and dynamo activity associated with reconnection often considered the primary drivers. Previous coaxial helicity injection experiments, figure 1.11, showed evidence of this behavior with high fluctuations and other reconnection signatures during current drive and sustainment[31]. Although this process of internal helicity transport does produce efficient current drive, it is not amenable to good confinement due to the possibility of flux surface breakage during large scale MHD activity.

Imposed Dynamo Current Drive

Observations on the HIT-SI experiment have indicated periods of low fluctuation amplitude while maintaining current growth and sustainment, indicating that the manner in which helicity is injected into the system can impact internal current drive mechanisms as well. This suggests the possibility that HIT-SI is driving current without the global MHD activity seen in previous, coaxial injection, experiments. A model for this new current drive mechanism[32], Imposed Dynamo Current Drive (IDCD), has been proposed. This model does not require the global reconnection activity typically associated with relaxation behavior. This model supports current




Figure 1.12: Schematic of shear induced in magnetic fluctuations by differential electron flow. In the presence of a gradient in electron velocity, V_e , an applied fluctuation will be sheared from its initial state, δB_f , to a distorted state, δB_d , by the flow.

drive not only with low fluctuation levels but also allows for current drive in closed flux configurations.

This method relies on imposed magnetic fluctuations that couple electron flows and resulting currents between a directly driven edge region and the interior of the plasma that remains isolated from the injectors. As shown in figure 1.12, fluctuations will be sheared in the presence of a gradient in the electron velocity, due to the frozen-in condition, resulting in a force to decrease the flow shear. If the flow is maintained at a high velocity in the edge region by direct drive, the force will act to reduce the shear by accelerating electrons in the interior, driving current. The thorough derivation and dynamics involved in this current drive method can be found in reference[32]. This model predicts that when the mean field in HIT-SI is being sustained by the injectors there must exist a force balance on each magnetic surface between the drag force on electrons inside the surface, due to ohmic dissipation, and the current driving force produced by shear in the imposed fluctuations induced by differential electron flows, given by equation 1.15. The volume and surface of integration are defined from the mean field equilibrium. J_{\parallel} , E_{\parallel} , and n are also evaluated from the mean field, while

δB is defined by the applied fluctuations.

$$\int_V ne (\eta J_{\parallel} - E_{\parallel}) dV = - \int_{\partial V} \frac{\delta B_{\perp} \delta B_{\parallel}}{\mu_0} dS \quad (1.15)$$

Pressure Driven Current Penetration

Recent observations during high frequency operation in HIT-SI indicate that plasma pressure is playing an important role in the dynamics. Sustained pressures well in excess of the theoretical β limit have been observed throughout a long portion of the shot, spanning multiple injector cycles. Figure 1.13 shows a comparison of the toroidal and poloidal magnetic fields along the internal probe at two different injector frequencies. Stars correspond to the probe measurements, while solid lines are the result of axisymmetric fits using the PSI-TRI code, section 2.1. Fits indicate a plasma β of 28% at 68.5 kHz, well above the 7.2% limit predicted by the Mercier criterion, section 2.1.2. It has been postulated that a shared mechanism may be responsible for regulating both the pressure and current profiles. Under this model, the pressure profile is maintained by preventing the uncontrolled growth of pressure driven interchange modes through the application of fluctuating boundary fields. The perturbing fields act to break up the growing mode on each injector cycle, forcing the mode to start growing again from a smaller seed, limiting the thermal energy lost due to interchange activity. The remaining, relatively low level, interchange activity that is present acts to flatten the q-profile and thus the λ profile across the equilibrium. When parallel current is maintained at a high level in the edge region it will drive current in the inner region to flatten the overall magnetic shear.

This results in a very similar effect to the imposed dynamo current drive model, but does so through localized breaking of flux surfaces at a rate that is already required to relieve excess thermal energy from the equilibrium. In this sense, the current is transported into the interior as the thermal energy is transported out, producing a balance that does not degrade confinement. Additionally, this process appears




Figure 1.13: Comparison of toroidal and poloidal magnetic fields along the HIT-SI internal probe at two different injector frequencies. Shown are probe measurements (stars) and the result of an axisymmetric fit (solid lines) with a linear pressure profile and two- λ current profile.

to regulate the current profile away from hard MHD stability limits, as shown in figure 1.14, indicating a different current drive process than previously observed in marginally stable coaxial injection experiments. If proven to be correct, this would be a powerful tool for controlling current and pressure profiles in a variety of devices and could be an enabling advance for the RFP and spheromak confinement concepts.

1.4.2 Limitations of Current Numerical Models

The existing NIMROD model of the HIT-SI device developed first by Izzo[27, 28] and later by Akcay[30, 2] have been very successful in providing insight into the important physical mechanisms at play in HIT-SI. However, some noted differences in the internal magnetic field profile as well as the current evolution with high frequency injector operation exist. Investigations have been performed to attempt a reconciliation of these differences, some of which are attributable to non-zero fluid pressure, however

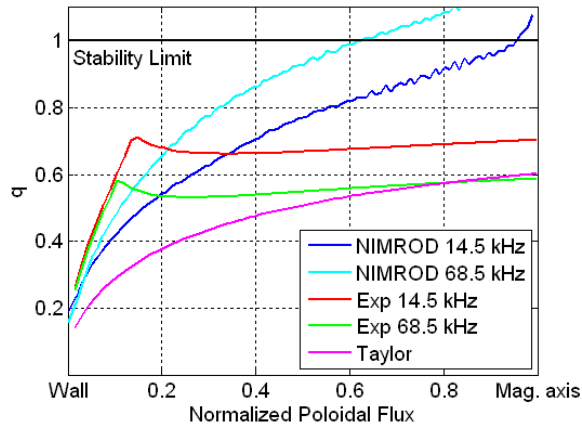


Figure 1.14: Comparison of q -profiles from NIMROD simulations, Taylor theory, and axisymmetric fits to experimental data (figure 1.13).

it was difficult to assess the effects of boundary conditions without a model of the full geometry to compare against.

Removing the injectors to allow use of the NIMROD code requires modeling the impact of physics within those regions as boundary conditions at the interface. The primary effect of the injectors is flux injection and current drive across the flux conserver. In NIMROD the dynamic effect of the injector drive coils are simulated by applying boundary conditions to the magnetic field where the mouths of the injectors would be. These boundary conditions produce a magnetic field at the surface that corresponds to an axisymmetric force-free equilibrium, computed in a volume made by revolving the injector footprints about their center. In practice this produces a field whose parallel current density is too high so the tangential field, which produces normal current at the boundary, is scaled independently of the normal field to match the experimental λ . Subsequently, the magnetic field pushed into the flux conserver is some fixed non-equilibrium state that does not respond to plasma dynamics. In the actual machine, coils produce both the flux and driven current, whose profiles will

then evolve self-consistently due to dynamics. The differences in current drive and formation caused by this approximation are unclear. It is expected that the current path taken by the driven current may be strongly impacted by using the boundary condition approach, especially near the injector mouths.

In addition to imposed magnetic field drive, boundary conditions on plasma fluid variables such as velocity, density, and temperature must be imposed along the injector mouths. Enforcing fixed values on these fields is expected to exhibit a significant effect on their evolution as plasma flows, density pump out, and strong heating are all observed at the injector mouths. In reduced MHD simulations only the velocity is evolved, with all components forced to zero at the boundary, but in HIT-SI strong flows have been observed exiting the injectors. Full MHD simulations with density and temperature evolution applying boundary conditions across injector mouths can be even more troublesome. Anisotropic thermal conduction and heating within the injector volumes should lead to significant temperature variation across the injectors. This will result in different temperature profiles in the confinement volume than with an imposed profile across the injectors, which will act to pin the temperature along field lines connected to the boundary. Similarly, density variation could develop across the injector mouth in the experiment, which cannot be modeled in NIMROD, however the effect of this variation would most likely be less pronounced than for flow and temperature.

Finally, the use of a Fourier representation in the toroidal direction presents challenges for future experiments. This form was chosen for its efficiency in cases when only a relatively small number of toroidal modes are necessary, as is often the case in simulations with primarily axisymmetric fields. For simulations of HIT-SI, and future experiments, the minimum number of toroidal modes is relatively large due to the number of modes required to accurately resolve the injector mouth boundary conditions. For the current two injector configuration this minimum is not very limiting, though the injectors used on HIT-SI3 and projected for HIT-POP designs are much

smaller— requiring the use of more and more toroidal modes, reducing the efficiency of the toroidal representation. The limitation results from the fact that with a purely spectral toroidal representation the resolution can not be locally increased as with a finite element representation in the toroidal direction.

1.4.3 Design Considerations for Future SIHI Experiments

Presently, the understanding of current drive mechanisms in HIT-SI are reaching the point where they can be used to design future experiments. Further investigation of this current drive method in regimes where good confinement can be expected and verified is essential to support development toward an economical fusion power system. Beyond a full code capable of modeling future experimental geometries, it is also helpful to have reduced order models to guide the rough design of future experiments. As part of this work, models were explored to use for predicting the performance of future devices using SIHI current drive techniques. These models require the ability to rapidly test new geometries as well as the integration of various order models to provide speed in the analysis.

1.5 Outline

This work develops numerical tools in order to support further study of the HIT-SI device and performs preliminary investigations using those tools to help address the open questions in the experimental program. The presentation and discussion of this work will proceed as follows: Chapter 2 will introduce the numerical tools developed and their verification and supporting materials drawn on for their development. Chapter 3 will present the development of a method for computing general 3D equilibrium and its application to the HIT-SI device. Chapter 4 will discuss investigation of reduced order models for use as design tools for future SIHI driven spheromaks. Chapter 5 will present the development of reduced and full Hall-MHD codes based on a Nedelec element representation for the magnetic field. This includes the devel-

opment of a self-consistent boundary condition for the insulator coated conducting walls within the Nedelec magnetic discretization. Finally, results from dynamic simulations of the full HIT-SI device using this new code will be presented in comparison to experimental and existing NIMROD simulations.

Chapter 2

NUMERICAL TOOLS

A major motivation for this research was the lack of existing tools capable of modeling devices with complex geometries. As such, a significant amount of work has been devoted to developing the underlying numerical infrastructure to support the new physical models described in the following chapters. This chapter will present brief descriptions and the noteworthy features of each tool. Two independent but related codes were written or extended as a part of this work: 1) PSI-TRI a 2D finite element code based on an unstructured triangular grid. 2) PSI-TET a 3D multi-physics framework for finite element methods on unstructured tetrahedral grids. Both codes are written primarily in FORTRAN using capabilities defined up to the 2003 language standard.

Documentation

In order to facilitate the continued use of these tools, extensive documentation has been added to the code along with instructive examples demonstrating use of the framework and physics modules. Documentation is written in a structured format to allow use of the Doxygen documentation generator, which parses the source files and generates formatted and searchable documentation in the form of a html document. Figure 2.1 shows the *in source* documentation for a representative PSI-TET subroutine, while figure 2.2 shows the resulting documentation, illustrating the ability to include Latex math commands and create linked references within documentation.

```

!-----
! SUBROUTINE: tetmesh_get_jacobian
!-----
!> Compute the jacobian matrix and its determinant for a grid cell
!!
!! Driver function calls mapping specific function depending on mesh order.
!!
!! @param[in] self Mesh containing cell
!! @param[in] i Index of cell for evaluation
!! @param[in] f Logical coordinate in cell [4]
!! @param[out] gop Jacobian matrix \f$ (\frac{\partial x_i}{\partial \lambda_j})^{-1} \f$ [3,4]
!! @param[out] j Jacobian of transformation from logical to physical coordinates
!-----

```

Figure 2.1: Documentation block

```

subroutine tetmesh_mapping::tetmesh_get_jacobian ( type(tet_mesh), intent(in)      self,
                                                    integer(4), intent(in)          i,
                                                    real(8), dimension(4), intent(in)   f,
                                                    real(8), dimension(3,4), intent(out) gop,
                                                    real(8), intent(out)               j
)

```

Compute the jacobian matrix and its determinant for a grid cell.

Driver function calls mapping specific function depending on mesh order.

Parameters	
[in]	self Mesh containing cell
[in]	i Index of cell for evaluation
[in]	f Logical coordinate in cell [4]
[out]	gop Jacobian matrix $(\frac{\partial x_i}{\partial \lambda_j})^{-1}$ [3,4]
[out]	j Jacobian of transformation from logical to physical coordinates

Definition at line 371 of file `tetmesh_mapping.F90`.

Figure 2.2: Resulting documentation




Figure 2.3: PSI-TRI solution of the Taylor state in a future SIHI driven device(HIT-POP), shown is the computational grid and a pseudocolor plot of the poloidal flux

2.1 *The PSI-TRI Code*

The PSI-TRI code is a 2D finite element code used to solve for axisymmetric MHD equilibria. It is based on a FORTRAN port of an existing MATLAB code developed by George Marklin. An unstructured triangular grid is used to discretize the domain, with a Lagrange basis set used to represent solution fields. Currently, up to Quartic representations are supported and higher order basis definitions may be generated automatically using an external Python script. Meshes can be imported from either the T3D[33] or CUBIT[34] mesh generation programs. The latter provides support for parameterizing geometry for sensitivity studies as well as common definition with 3D models for PSI-TET.

Krylov space iterative solvers are used for the solution of linear systems to reduce memory requirements. To accelerate solver convergence and decrease overall run time Multi-Grid preconditioning can be used with lower polynomial orders used as the

coarse representations. The code is also parallelized for shared memory environments using the OpenMP library. Together these methods produce reasonable run time for even the highly resolved cases required for complex profile specifications such as the marginal Mercier profile described below in section 2.1.2. An example of the computational grid for HIT-SI and the resulting Taylor state poloidal flux profile are shown in figure 2.3.

More detailed descriptions of these implementations are provided in chapter 2.2, as the 2D implementations can be viewed as a restricted subset of the 3D implementations.

2.1.1 Grad-Shafranov Equilibrium

In cylindrical coordinates (r, ϕ, z) , where r is the cylindrical radius, ϕ is the toroidal angle, and z is the axial coordinate, an axisymmetric, $\frac{\partial}{\partial \phi} = 0$, magnetic field can be expressed as in equation 2.1. The first term corresponds to the field in the poloidal plane (r, z) while the second term corresponds to the toroidal field scaled by $\frac{1}{r}$. As the magnetic field points in a direction orthogonal to $\nabla\psi$ everywhere ψ is a flux surface variable, meaning magnetic field lines lie on surfaces of constant ψ . For this work ψ is defined as the poloidal magnetic flux contained between a given magnetic surface and the boundary. $F(\psi)$ is an arbitrary function of the poloidal flux that defines both the toroidal flux and poloidal current in the equilibrium.

$$B = \nabla\psi \times \nabla\phi + F(\psi)\nabla\phi \quad (2.1)$$

Using this representation the ideal MHD equilibrium equation 1.12 can be reformulated as equation 2.2[35]. This equation, known as the Grad-Shafranov equation, constitutes a non-linear boundary value equation for ψ . The differential operator applied to ψ on the left hand side of this equation is referred to as the $\Delta^* = \left[r \frac{\partial}{\partial r} \left(\frac{1}{r} \frac{\partial}{\partial r} \right) + \frac{\partial^2}{\partial z^2} \right]$ operator due to its similarity to the Laplacian. In the absence of pressure gradients,

the parallel current density can be defined as $\lambda = \frac{\mu_0 J_{\parallel}}{|B|} = \frac{dF}{d\psi}$. The Taylor state is then the solution to the Grad-Shafranov equation with $F = \lambda\psi$ and $p = 0$, which now defines an eigenvalue problem for ψ . When a non-zero gradient in pressure is present this definition for λ is only true in a surface averaged sense, due to diamagnetic currents.

$$\left[r \frac{\partial}{\partial r} \left(\frac{1}{r} \frac{\partial}{\partial r} \right) + \frac{\partial^2}{\partial z^2} \right] \psi = -\mu_0 r^2 \frac{dp}{d\psi} - F \frac{dF}{d\psi} \quad (2.2)$$

This equation has been well studied in application to many nominally axisymmetric configurations such as the tokamak and spheromak. It is also used extensively for both experimental comparison using fitting codes such as EFIT[36] as well as future experiment design[37].

Solution Method

In order to solve equation 2.2 a Picard iteration is used to handle the non-linearity in conjunction with a finite element expansion of the equation. An updated guess for ψ is computed by evaluating the right hand side of equation 2.2 with the current value ψ^i and then solving the Δ^* operator to compute ψ^{i+1} and an updated scale for F . This process generally converges to a solution in a few non-linear iterations.

The linear system corresponding to the Δ^* operator is solved using a Galerkin finite element method with a nodal basis set. In weak form the operator takes the form 2.3 following integration by parts to replace second derivatives with first order terms only. The surface term is ignored for cases considered so far as the boundary condition $\psi = 0$, corresponding to no external field and a perfectly conducting wall, is assumed. The surface term goes to zero in this case as the test functions must match the boundary condition of the basis set and be equal to zero everywhere on the boundary.

$$\int \frac{1}{r^2} \left[\frac{\partial u^T}{\partial r} \frac{\partial \psi^{i+1}}{\partial r} + \frac{\partial u^T}{\partial z} \frac{\partial \psi^{i+1}}{\partial z} \right] dV - \int \frac{1}{r^2} \left[u^T \frac{\partial \psi^{i+1}}{\partial r} \hat{r} + \frac{\partial \psi^{i+1}}{\partial z} \hat{z} \right] \cdot dS \quad (2.3)$$

The Picard iteration update is then given by equation 2.4.

$$\int \frac{1}{r^2} \left[\frac{\partial u^T}{\partial r} \frac{\partial \psi^{i+1}}{\partial r} + \frac{\partial u^T}{\partial z} \frac{\partial \psi^{i+1}}{\partial z} \right] dV = - \int u^T \left[\mu_0 \frac{dp(\psi^i)}{d\psi^i} - \frac{F(\psi^i)}{r^2} \frac{dF(\psi^i)}{d\psi^i} \right] dV \quad (2.4)$$

Toroidal Flux and Pressure Definition

The PSI-TRI code allows the use of arbitrary functional forms for both F and P by abstracting away the function evaluations from the solution procedure. This allows new profile types to be defined and used without modifying any core subroutines. Currently, a number of general and HIT-SI specific profiles have been implemented for the investigations in Chapter 4. This capability is demonstrated by Figure 2.4, which shows a pressure profile and its derivative defined using cubic splines to approximate marginal stability. Markers indicate locations where the pressure profile is defined from the Mercier criterion, equation 2.5, and fit using splines.

2.1.2 Marginally Stable Mercier Pressure Profiles

The Mercier stability criterion[38] defines a threshold for stability to pressure driven modes at a given flux surface. In order to define equilibria that are marginally stable to this criteria everywhere, a procedure similar to that employed by Mayo and Marklin[39] is used. The Mercier criterion can be evaluated on each flux surface by equation 2.5

$$4\pi^4 (q')^2 + p'V' \left(4\pi^2 q' \left\langle \frac{F}{|\nabla\psi|^2} \right\rangle + V'' \left\langle \frac{B^2}{|\nabla\psi|^2} \right\rangle \right) + (p'V')^2 \left(\left\langle \frac{F}{|\nabla\psi|^2} \right\rangle^2 - \left\langle \frac{r^2}{|\nabla\psi|^2} \right\rangle \left\langle \frac{B^2}{|\nabla\psi|^2} \right\rangle \right) \geq 0 \quad (2.5)$$





Figure 2.4: Pressure profile and its derivative for an equilibrium in HIT-SI defined using splines fit to the Mercier criterion.


where $V' = -2\pi \oint \frac{dl}{B_\theta}$, $q = \frac{F}{2\pi} \oint \frac{B_\theta}{r^2} dl$, $\langle f \rangle = \left(\oint f \frac{dl}{B_\theta} \right) / \left(\oint \frac{dl}{B_\theta} \right)$ and integration is performed around one flux contour. For a given equilibrium, this equation can be used to define p' , such that marginal stability is achieved. However, this pressure profile is not consistent with the equilibrium in general so the equilibrium and pressure profile are successively updated until a self-consistent equilibrium is produced. An example of such an equilibrium for a flat λ profile in the HIT-SI geometry is shown in figure 2.5.

2.1.3 Experimental Fitting and Reconstruction

In addition to allowing the use of general flux functions, the abstraction described in section 2.1.1 facilitates the adjustment of profiles for reconstructing experimental equilibria. PSI-TRI has the ability to adjust profile coefficients to minimize the difference between experimentally observed quantities and equilibrium solutions similar to the EFIT code. PSI-TRI uses the MINPACK library[40] to perform a non-linear minimization of the error between a set of experimental and synthetic diagnostics



(a) Poloidal flux contours



(b) q-profiles

Figure 2.5: 2D equilibria in HIT-SI showing the difference between zero beta (dashed) and marginal Mercier stability (solid) with constant λ . Only a slight Shafranov shift and modification of the q-profile is seen due to the low β -limit with uniform λ in HIT-SI, $\sim 2.8\%$.

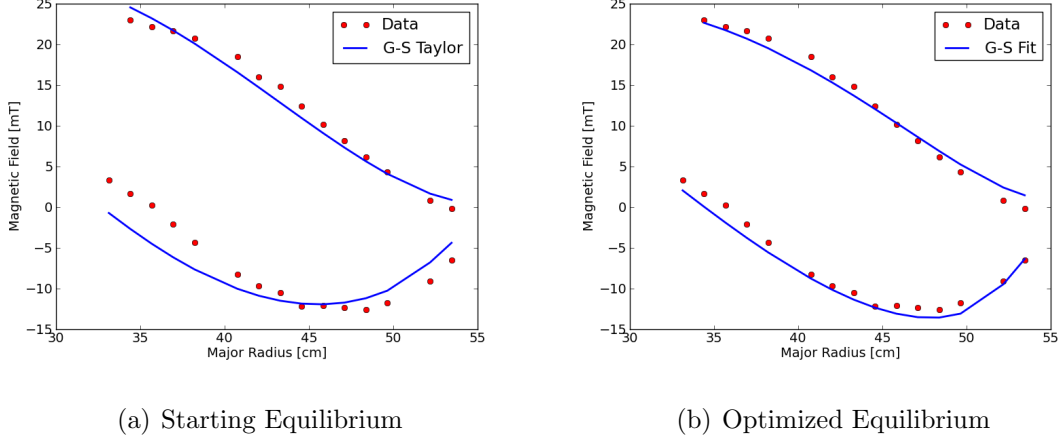


Figure 2.6: Example axisymmetric equilibrium fit to experimental signals. Shown are toroidal (top) and poloidal (bottom) magnetic fields from the internal magnetic probe (red) used as fitting constraints and the model magnetic profiles (blue) for the initial guess (a) and resulting fit (b) equilibria.

used to constrain the equilibrium.

MINPACK uses a Levenberg-Marquardt algorithm to minimize the 2-norm of the error vector $E = \sum_j (S_j^e - S_j^s)^2$, where S^e and S^s are the experimental constraints and simulated signals respectively. Constraints are abstracted away from fitting in a way similar to flux profiles allowing the simple use of signals unique to a given experiment or flux profile case. Figure 2.6 shows an example of this fitting process where internal magnetic probe data are used to constrain the profile. A linear profile in both F' and p was used with the initial guess as a Taylor state, zero pressure and constant F' . The resulting fit, figure 2.6(b), shows much better agreement with the experimental signals when compared to the guess, figure 2.6(a). This provides a useful tool for comparison and analysis of experimental data and verification of model profiles in chapter 4.




Figure 2.7: Normal displacement for the most unstable eigenmode as determined by DCON for an equilibrium in a cylinder computed by PSI-TRI.

2.1.4 DCON Interface

PSI-TRI includes an interface for exporting Grad-Shafranov equilibria to the DCON stability code for analysis. DCON[41] uses a flux coordinate representation to evaluate the generalized Newcomb criterion for stability and provides the ability to analyze a broad range of ideal and non-ideal instabilities. Following calculation of an equilibrium in PSI-TRI the solution may be converted from an R, Z representation to a flux parameterization ψ, θ by tracing. R, Z are computed at evenly spaced positions in the ψ and θ and written out along with the flux surface quantities P, F, q . These data can then be read in to DCON for analysis. Figure 2.7 shows the results of stability analysis for an example unstable equilibrium in a cylinder.

The DCON interface has also been used to validate the calculation of marginally stable equilibria in PSI-TRI using the Mercier criterion. Test cases were exported and compared to DCONs internal Mercier evaluation.

2.2 *The PSI-TET Code*

The PSI-TET code is a 3D finite element framework that supports the development of multi-physics models on a tetrahedral grid. The code features extensible interfaces and mechanics to support constructing and solving linear and non-linear systems of equations resulting from finite element discretization. The design goal was to create a framework that abstracts away as much of the supporting components for multi-physics simulation while maintaining generality in the finite element representation and PDE forms. This ability dramatically simplified the setup and testing of physics models shown in the following sections.

PSI-TET is a massively parallel code supporting both distributed (task) and shared (thread) memory parallelization. Thread level parallelism is supported via the OpenMP library[42], while the Message Passing Interface (MPI) protocol[43] is used for communication between tasks that do not share common memory. To enable the use of large meshes and efficiently utilize parallel computing environments the mesh is decomposed using the grid partitioning library METIS[44]. More information on grid construction is provided in section 2.2.1.

Linear algebra operations are supported through either a native implementation or an interface to the PETSc library[45]. The specifics of common vector/matrix operations are abstracted into a common interface that is back end agnostic. This allows the user to design code that can access a wide variety of solvers through a common and consistent interface. Support for specifying complex solver configurations at run time is also supported through an input with an XML description of the desired solver. Support for the use of XML is provided through the FoX library[46].

Finite element mechanics are also abstracted into common interfaces for flexibility. However, the level of abstraction is less general than the linear algebra interface to maintain full access to the internal representation for optimization considerations. Common operations such as the construction of global indices, weight vectors, and

matrix graphs are handled automatically based on specified geometric ownership for each basis function. Adding new finite element representations is reduced to specifying the basis functions and their geometric ownership by these mechanics. Further description of the specific finite element types currently implemented is given in sections 2.2.3 and 2.2.4. Symmetric quadrature rules in 1, 2, and 3 dimensions derived by Linbo Zhang[47] are also provided for performing weak form integrals and matrix construction.

2.2.1 Mesh Construction

For geometric flexibility PSI-TET uses an unstructured tetrahedral spatial discretization. This provides the ability to quickly generate computational grids from CAD models of experiments by using proven commercial meshing tools. Any generated mesh can be imported into PSI-TET and used through a generic mesh interface, while two specific mesh generation programs also support more complex functionality.

Hierarchical Refinement

In order to support the development and use of Multi-Grid preconditioners PSI-TET supports internal grid refinement. A coarse grid can be imported and subdivided to produce a series of successively finer grids that can be used to define approximate representation levels for geometric Multi-Grid. By subdividing the grid internally a known structured relationship between grid levels is maintained, simplifying level to level interpolation. Each refinement produces a child mesh by subdividing all edges, roughly doubling the resolution, in the parent mesh and adding tetrahedra to match the new node locations. Figure 2.8 shows an example of this refinement procedure for a mesh of the HIT-SI experiment. A slice of the grid is shown at three different refinement levels, illustrating the additional faces and thus cells, created during refinement. This example also demonstrates the capability of the T3D mesh interface to perform refinement will conforming to the CAD boundary.

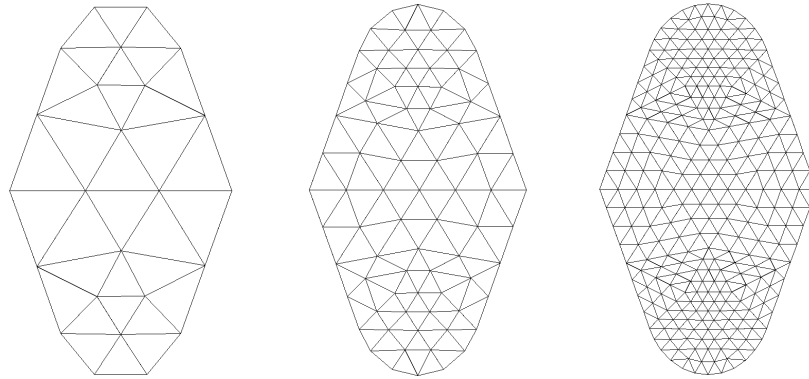


Figure 2.8: Cross-section of HIT-SI injector mouth showing mesh refinement.

Additionally, when a coarse grid is used to define the finer levels global indices can be defined in a consistent manner that allows restart files to be stored and read with a different number of parallel subdomains. As the coarse grid is generally small enough to be constructed on each independent MPI task, a global numbering can be created that is independent of the number of MPI tasks. A similar indexing would be difficult and time consuming to construct for a distributed mesh as it would require sorting distributed lists.

Parallel Decomposition

As mentioned above, the mesh can be decomposed to run on massively parallel systems. This process is performed using general graph partitioning methods available through the METIS library. METIS partitions the mesh into sub-domains by simultaneously balancing the number of cells in each sub-domain and minimizing the number of faces shared between two sub-domains. This process approximately load balances the mesh for matrix construction, volume integration, vector-vector, and matrix-vector operations. An example decomposition for a HIT-SI mesh is shown in figure 2.10 with 72 sub-domains.




Figure 2.9: Child cells produced by uniform refinement a single tetrahedron.




Figure 2.10: Sub-domains from parallel decomposition of a HIT-SI mesh with 72 MPI tasks. Note that some colors are repeated.




Figure 2.11: Average number of inter-task connections per MPI task for a cylindrical mesh with different numbers of MPI tasks.

Additionally, by minimizing the number of shared faces the decomposition provides a reasonably optimized minimization of the required inter-task communication. Inter-task communication is provided by the MPI library and is a significant bottleneck for massively parallel systems due to the relatively long distances between different nodes. As a result minimizing the number of inter-task connections is an important consideration. Figure 2.11 shows the average number of inter-task connections per task as the number of tasks, and therefore sub-domains, is increased. The number of connections initially increases but then plateaus indicating that the number of dependent communications does not continue to grow. Note that the total number of communications required for a global operation does increase, however the majority of these can be performed concurrently and thus do not block each other.

T3D Interface

The T3D[33] mesh generation code provides the ability to generate quality tet meshes from a relatively simple description of the boundary. Rational Bezier Splines (RBS) are used to represent the boundary curves and surfaces, which can be reconstructed

by a simple RBS implementation in PSI-TET. This boundary description supports both the conformal refinement procedure as well as the conformal placement of node points for high order boundary cells. For more information on the specific interface see appendix A.1.

CUBIT Interface

The CUBIT[34] software is a fully featured CAD and mesh generation code developed at Sandia National Labs. A CAD model based on Non-Uniform Bezier Splines (NURBS) is used to represent the geometry, which can be imported from external CAD software or created internally by CUBIT. This representation is the same representation that is used by most commercial CAD programs such as AutoCAD or SolidWorks. PSI-TET uses an interface to the OpenNURBS library[48] to provide the needed functionality to evaluate NURBS geometry. The boundary description supports both the conformal refinement procedure as well as the conformal placement of node points for high order boundary cells. For more information on the specific interface see appendix A.2.

2.2.2 Geometric Mapping

The tetrahedral discretization used in PSI-TET defines a piecewise continuous mapping from logical coordinates of the unit element to physical coordinates. The unit element is a right tetrahedra with its four vertices placed at the origin and at unit length along each of the 3 logical axes respectively. Basis functions are then defined on this element to define the physical coordinates used for simulation, which in turn defines a mapping between the two coordinate spaces. An example where the physical coordinates have a quadratic mapping from logical space is shown in figure 2.12.

The basis functions used to represent modeled fields are also defined on the unit element. Therefore, integrations and derivatives in terms of physical space must be transformed from physical to logical space for evaluation. Integration over the full

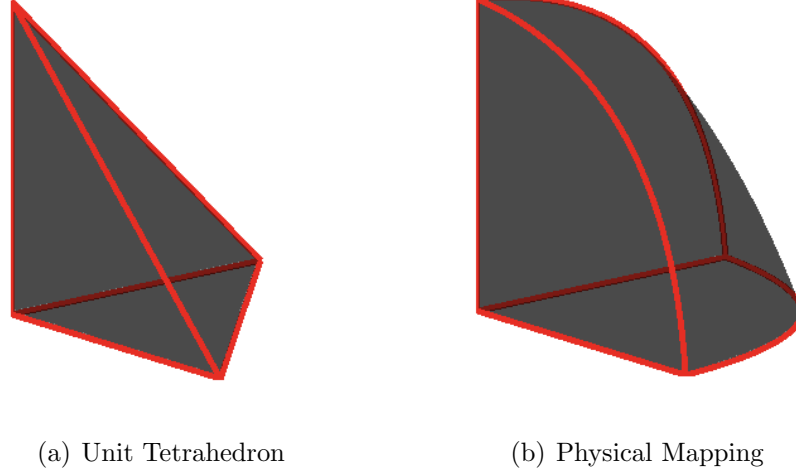


Figure 2.12: Example mapping from the unit tetrahedron (a) in logical space to the physical cell (b). A quadratic mapping is shown approximating an octant of a sphere with a single cell.

domain is broken up in to a discrete sum of continuous intergals within each cell (tetrahedron), which form the piecewise continuous spatial discretization. Integration over a single cell is then evaluated by transforming the integration into logical space as in equation 2.6, where Ω and \vec{x} are the integration domain and vector of coordinates in physical space and Ω' and $\vec{\xi}$ are the integration domain and vector of coordinates transformed into logical space. This transformation introduces the grid Jacobian $|J(\xi)|$, which is the determinant of the grid Jacobian matrix, $J_{ij}(\xi) = \frac{\partial x_i}{\partial \xi_j}$, that represents the mapping between logical and physical space.

$$\int_{\Omega} f(\vec{x})dV = \int_{\Omega'} f(\vec{\xi})\frac{|J(\xi)|}{6}dV' \quad (2.6)$$

Similarly, the mapping can be used to evaluate derivatives with respect to physical coordinates in terms of derivatives of the basis functions with respect to the logical coordinates. Equation 2.7 shows the transformation, where the grid Jacobian matrix

$J(\xi)$ must be inverted to evaluate the desired derivatives in physical space.

$$J_{ij}(\vec{\xi}) \frac{\partial f(\vec{x})}{\partial x_i} = \frac{\partial f(\vec{\xi})}{\partial \xi_j} \quad (2.7)$$

High Order Geometric Mapping

In PSI-TET a nodal definition is used with Lagrange interpolatory polynomials as basis functions that matches the Lagrange basis defined in section 2.2.3. This can be used to define mappings up to 3rd order however in practice only linear and quadratic mappings are commonly used. Relatively small curved features are common in complex geometries, which can cause issues when smooth boundaries are under resolved introducing spurious corners in the mesh. An example of this effect is shown in figure 2.13 where a linear and quadratic physical mapping are compared at the nose cone. The nose cone in HIT-SI has a 1 inch radius and each quadrant of the curve is represented with a single edge in the example mesh. The linear representation under resolves the curvature and introduces a sharp inside corner, which is spurious. By contrast, the quadratic representation produces a smooth mapping across the entire cone. Although this is a small region of the mesh, large gradients exist here in dynamic simulations and a non-smooth boundary here can affect the stability and speed of the entire simulation.

Although both CUBIT and T3D are capable of generating quadratic meshes, PSI-TET generates the additional node points required for high order mappings internally. This is necessary for refined meshes and can be done easily using the same mechanics as used for placing new vertices during geometric refinement.

2.2.3 Lagrange Basis Set

A nodal finite element basis set for scalar fields is defined for use in PSI-TET using evenly spaced nodes and Lagrange interpolatory polynomials. This basis set is relatively simple and its nodal representation allows it to be used for saving solution fields

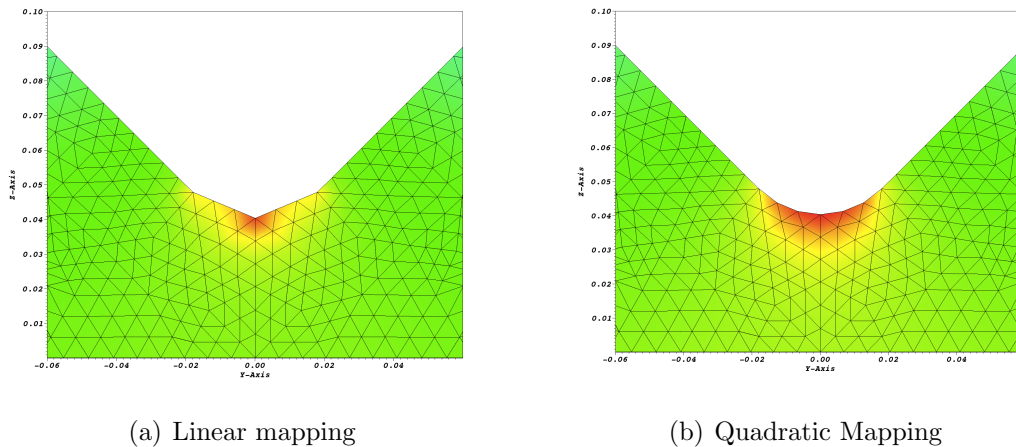


Figure 2.13: Comparison of grid boundary between linear (a) and quadratic (b) mappings for a small curved feature in the HIT-SI mesh. Shown is the tessellated mesh and $|B|$ for the second force-free eigenmode computed using a cubic $H^1(\text{Curl})$ basis on each mesh.

for post-processing in the LLNL developed VisIt visualization software[49]. Additionally, this representation is also the basis used for the logical to physical mapping as described in section 2.2.2. Node points are evenly spaced in each of the independent logical coordinates, though unlike in a hexahedral spatial discretization, this basis does not produce a representation equivalent to construction by direct products of functions in each direction. The highest order term for a given polynomial level is subject to the constraint $i + j + k \leq np$, where i, j, k are the powers of the x, y, z coordinates respectively and np is the polynomial order.

Verification and convergence

In order to verify the element definitions and their expected convergence behavior a static problem is solved with a known analytic solution. A Poisson problem, equation 2.8, is solved in a cube with the Dirichlet boundary condition $\phi = 0$ on the boundary. The cube extends from $[-1,1]$ in each dimension and the source given in equation 2.9

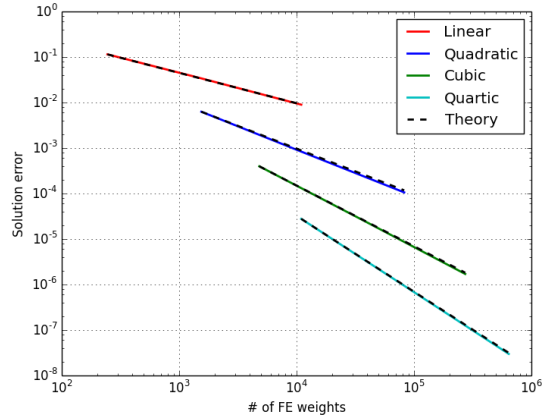


Figure 2.14: Convergence verification for the Lagrange basis set on a Poisson test case in a unit cube.

is used providing a smooth analytic solution.

$$\nabla^2 \phi = S \quad (2.8)$$

$$S = \nabla^2 \cos(\pi x/2) \cos(\pi y/2) \cos(\pi z/2) \quad (2.9)$$

Figure 2.14 shows the convergence behavior of the Lagrange representation across different grid resolutions in PSI-TET for different polynomial orders. Solution error is computed by comparing the calculated ϕ to the analytic solution using an L^2 norm. The error is observed to converge toward zero in a uniform manner at the predicted theoretical rates, shown as black dashed lines for each polynomial order.

Multi-level interpolation

To enable the use of multi-grid solvers as shown in section 2.2.5, operators must be defined to interpolate the weights corresponding to a solution on one basis set to

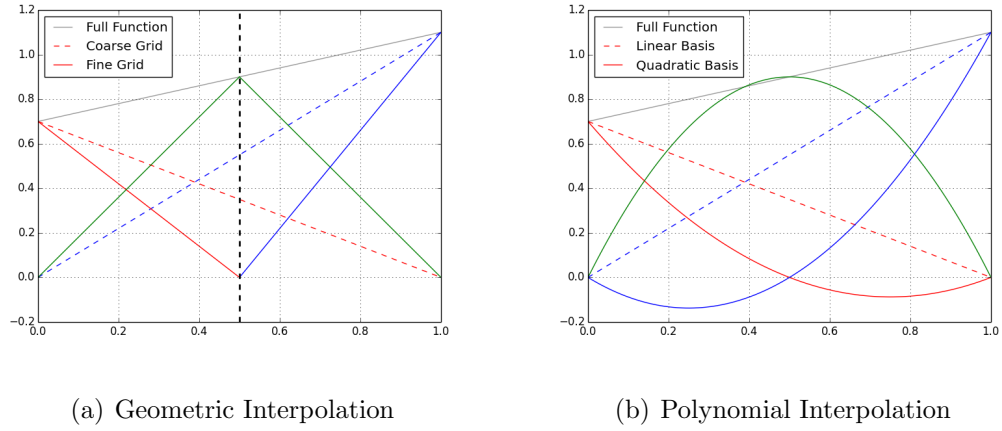


Figure 2.15: Demonstration of interpolation between Lagrange basis representations under geometric (a) and polynomial (b) refinement.

the weights required to reproduce the same solution on the next finest basis set. For the multi-grid method currently used, interpolations between representations are only needed where the two basis sets are either both linear and belong to two adjacent grid refinements in the mesh hierarchy or belong to the same grid but differ in polynomial degree by 1. As the Lagrange basis set is defined by nodal values, the interpolation simply results by sampling the basis function values of the coarser representation on the node points of the fine representation. This is shown for an example linear solution in 1 dimension in figure 2.15 for both geometric and polynomial interpolation.

2.2.4 Nedelec Basis Set

PSI-TET also supports the use of Nedelec finite elements[50] for representing vector fields. Nedelec elements are composed of a sequence of vector valued basis sets that are defined in order to accurately handle the null spaces of vector differential operators such as the curl operator. This property along with the induced continuity of each element space make them desirable for problems where the solution exists outside or

relies on preserving a specific null space. Nedelec elements have been used extensively in modeling wave guides and other electromagnetic wave problems[51, 52, 53, 54], where only the portion of the electric field that supports a curl is of interest.

These elements are used to investigate Ideal MHD equilibria in chapter 3, where the equations to be solved are related to wave guide problems, as well as time dependent MHD in chapter 5, where the elements are used in a new way to enable modeling the HIT-SI device.

Function Spaces

On a given dimensional configuration space a sequence of function spaces are defined based on the induced exterior derivatives for that space. The sequence begins with the space of scalar functions that are dependent on the spatial coordinates and maps to subsequent spaces by the application of successive differential operators. In 3 dimensions this defines a 4 space sequence known as the De Rham Complex, equation 2.10, where H^i is the i -th function space and the differential operators are the familiar gradient, curl, and divergence operators. These spaces in 3 dimensions are commonly referred to as the scalar, vector, psuedo-vector, and psuedo-scalar space respectively.

$$H^0 \xrightarrow{\nabla} H^1 \xrightarrow{\nabla \times} H^2 \xrightarrow{\nabla \cdot} H^3 \quad (2.10)$$

This sequence has the property that the range space of each differential operator exists entirely in the null space of the next operator in the sequence, ex. $\nabla \times (\nabla f) = 0$. Nedelec elements are defined by leveraging this sequence to define basis sets, which accurately preserve this property of adjacent operators. Unique basis sets are defined for each function space by starting with a scalar basis set and applying operators in sequence to generate the subsequent spaces. The basis functions for the H^1 space are defined by taking the gradient of the basis set for the H^0 space. Though as noted above, this only populates the portion of the function space which exists in the null

space of the curl operator, so additional elements must be defined which support a curl. This results in a basis set that spans the full function space while being composed of two distinct subsets that represent the source and null space of the curl operator distinctly, similar to the well known Helmholtz decomposition, equation 2.11.

$$B = \nabla \times A + \nabla \phi \quad (2.11)$$

The term Nedelec element is often used in reference to only the H^1 space, while the elements for the H^2 space are often referred to as Raviart-Thomas elements. For this work the entire sequence will be referred to as Nedelec elements since they are defined and generated as a family. The various function spaces will be distinguished using the H^i notation. PSI-TET currently only implements basis sets for the H^0 and H^1 function spaces, which are used for the following investigation.

Element Definition

PSI-TET uses a representation based on Legendre polynomials presented by Schöberl and Zaglmayr[55]. This representation defines elements to arbitrary polynomial degree allowing high order representations, while optimizing somewhat element orthogonality to improve the condition number of finite element matrices. The elements are a heirarchical basis set as opposed to a nodal set like the Lagrange elements presented above. This means the basis set for a given order is created by taking the basis set from the next lowest order and adding functions without modifying the existing ones.

The basis functions are defined in terms of the Barycentric coordinates of the tetrahedra, λ_i , and contain an implied linkage to the discretization primitives: vertices, edges, faces, and cells. A sample of the element definitions are shown in equations 2.12 and 2.13 to illustrate the way in which the H^1 space is generated from the H^0 space. Functions are associated with a given geometric object if their definition is dependent on the Barycentric coordinates of that entity and no others. For example

edge degrees of freedom, equation 2.13, are dependent only on the position relative to that edge's end points $(\lambda_{e_1}, \lambda_{e_2})$.

Vertex degrees of freedom form the linear representation for scalar fields in the H^0 , while only the gradient subset of the H^1 vector space has basis functions tied to vertices.

$$f_p^0 = \lambda_p \quad (2.12a)$$

$$\vec{f}_p^{1g} = \vec{\nabla} f_p^0 \quad (2.12b)$$

Basis functions tied to edges however are defined for the both the gradient space, equation 2.13b, and the curl space, equation 2.13c. This illustrates the generation of the $H^1(\text{Curl})$ subspace from the H^0 basis set by permuting the order of terms resulting from the gradient of the scalar function. This concept is extended to the face and cell associated basis functions where more permutations are possible.

$$f_e^0 = L_i^s (\lambda_{e_1} - \lambda_{e_2}, \lambda_{e_1} + \lambda_{e_2}) \quad (2.13a)$$

$$\vec{f}_e^{1g} = \vec{\nabla} f_e^0 \quad (2.13b)$$

$$\vec{f}_e^{1c} = \lambda_{e_2} \vec{\nabla} \lambda_{e_1} - \lambda_{e_1} \vec{\nabla} \lambda_{e_2} \quad (2.13c)$$

Continuity Properties

A unique continuity between cells is also induced for each of the function spaces. This is illustrated in figure 2.16, which shows a scalar element and its corresponding child element in the H^1 vector space. The scalar function only attains a non-zero value on the boundary corresponding to its parent edge resulting in C^0 continuity. On the other hand, the vector function has a non-zero normal component on all boundaries of the cell allowing discontinuities in this component between cells. This property of C^0 continuity in only the tangential component of a vector field is one of the reasons these elements are well suited for waveguide problems where the permeability at a




Figure 2.16: Example scalar (left), H^0 , and vector (right), H^1 , Nedelec elements in 2D corresponding to the lower edge of the triangle. The vector element results from taking the gradient of the scalar function.

material interface is discontinuous requiring a jump in only the part of the electric field perpendicular to the interface.

Representation Redundancies

Although the basis subsets for the H^1 space appear to be distinct an overlap is present in the representation of curl-free fields. This overlap is due to the fact that a vector field that satisfies $\nabla \times v = 0$ and $\nabla \cdot v = 0$ can be represented as either the curl of a vector field or the gradient of a scalar. As a result, these types of fields can be represented on either basis subset within a cell. For functions with linear or higher order spatial dependence the cell to cell continuity of the field makes the representation distinct eliminating the ambiguity. However, functions that correspond to the gradient of a linear scalar function do not separate, leaving a singularity in the resulting matrices. This singularity can be eliminated by restricting the representation of these fields to a single subspace, which is most readily achieved by removing the gradient functions associated with vertices from the basis set.

Verification and Convergence

In order to verify the element definitions for the $H^1(\text{Curl})$ space and their expected convergence behavior, a static problem is solved with a known analytic solution. A vector Poisson problem, equation 2.14, is solved in a cube with the Dirichlet boundary condition $\vec{A} \times \hat{n} = 0$ on the boundary, where \hat{n} is the unit normal vector. The cube extends from $[-1,1]$ in each dimension and the source given in equation 2.15 is used providing a smooth analytic solution.

$$\nabla^2 \vec{A} = \vec{S} \quad (2.14)$$

$$\begin{aligned} \vec{S} = \nabla^2 [& \text{Cos}(\pi y/2) \text{Cos}(\pi z/2) \hat{x} \\ & + \text{Cos}(\pi x/2) \text{Cos}(\pi z/2) \hat{y} \\ & + \text{Cos}(\pi x/2) \text{Cos}(\pi y/2) \hat{z}] \end{aligned} \quad (2.15)$$

Figure 2.17 shows the convergence behavior of the Nedelec $H^1(\text{Curl})$ representation across different grid resolutions in PSI-TET for different polynomial orders. Solution error is computed by comparing the computed value of $\nabla \times \vec{A}$ to the analytic solution using an L^2 norm. The error is observed to converge toward zero in a uniform manner at the predicted theoretical rates, shown as black dashed lines for each polynomial order.

Multi-level interpolation: $H^1(\text{Curl})$

As with the Lagrange elements, to enable the use of multi-grid solvers operators must be defined to interpolate the weights corresponding to a solution on one basis set to the weights required to reproduce the same solution on the next finest basis set. In the case of Lagrange elements, this process can be done by simply evaluating the basis functions of one representation at the node points of the other. However, the Nedelec




Figure 2.17: Convergence verification for the Nedelec basis set on a vector Poisson test case in a unit cube.

elements are not a nodal basis set and share a different centering requiring extra care in defining interpolation operators. For interpolation between polynomial levels on the same grid, interpolation is trivial for the $H^1(\text{Curl})$ weights as the basis set forms a hierarchy. As a result, the weights from the coarse level are simply retained and weights corresponding to new basis functions are set to zero.

For interpolation between geometric levels weights on the finer grid are set by imposing a flux constraint on the $\nabla \times$ of the solution. The weights for the linear basis functions of the $H^1(\text{Curl})$ space are equivalent to $\vec{A} \cdot \vec{dl}$, where \vec{A} is the solution field and \vec{dl} is the length of the corresponding grid edge. In order to preserve this quantity the weights of each of the child edges resulting when the grid is refined must be set to $1/2$ the parent edge. Figure 2.18 shows the relationship between the weights of the 6 new edges on the boundary of each face and their parent edges.

Taking Stokes' theorem we can also define the flux of $\nabla \times \vec{A}$ through the parent face as $\sum_i \omega_i$, where ω_i are the function weights for the three basis functions tied to the face's edges. The three new child edges in the center of the face can now be set




Figure 2.18: Demonstration of edge interpolation between $H^1(\text{Curl})$ basis representations under geometric refinement.

in order to maintain the same flux through the parent face as shown in figure 2.19. Finally, the edges created internal to each coarse cell (figure 2.9) are set by computing $\vec{A} \cdot \vec{dl}$ from the coarse solution and length at the center of the edge. The interpolation operator produced from these steps is suitable for preconditioning the $H^1(\text{Curl})$ space and maintains, without mixing, both the solenoidal and low order gradient subspaces of the function space spanned by the linear basis function.

Multi-level interpolation: $H^1(\text{Grad})$

The weights for the $H^1(\text{Grad})$ subspace are also somewhat more complicated to interpolate than the Lagrange basis set. With the $H^1(\text{Curl})$ elements, interpolation between adjacent polynomial order representations on the same grid is trivial, as the basis functions form a hierarchical set. Interpolation from one grid level to another is more complicated however. The $H^1(\text{Grad})$ basis set is equivalent to the gradient of the H^0 scalar basis set, which is also defined in PSI-TET, so the interpolation operator can be thought of either as interpolating a linear vector representation from one grid to another or as interpolating a quadratic scalar representation between the




Figure 2.19: Demonstration of flux constrained interpolation between $H^1(\text{Curl})$ basis representations under geometric refinement.

two grids. The second case is simpler to understand conceptually so the interpolation process will be explained using a scalar function and basis set.

For a nodally centered set of scalar basis functions, interpolation is conceptually the same regardless of polynomial order or grid refinement on the source and destination levels—subject to the constraint that the coarse solution must be fully representable on the finer level. The values of the coarse solution are evaluated at the finer nodal points and used to set the function weights. A similar process can be carried out with two quadratic H^0 scalar representations on adjacent grid levels by defining constraint points that are used to solve for the fine weights in terms of the coarse solution. For simplicity this is done by using the node points for a quadratic Lagrange representation as they provide a suitable number and centering for these constraints. Node points for this polynomial order exist at each vertex and the center of each edge in the grid. Although the full H^0 basis set is not nodal, the linear portion of the space is spanned by a nodal set, corresponding to a weight for each grid vertex, as is required to ensure C^0 continuity. Therefore, these function weights can be interpolated in the same way as shown in figure 2.15(a).




Figure 2.20: Demonstration of node constrained interpolation between H^0 basis representations under geometric refinement.

The edge centered function weights can then be constrained using the remaining constraint points, one at the center of each edge. These constraint points result in fine function weights as defined by equation 2.16, where ω_i^f , ϕ_i^f , and $\vec{\xi}_i$ are the weight, basis function and constraint location respectively for the i -th edge on the fine grid. ω_j^c , ϕ_j^c , ω_k^f , and ϕ_k^f correspond to the weights and basis functions for all other finite element nodes on the coarse and fine grid respectively. In practice only those functions that attain a non-zero value at the constraint point $\vec{\xi}_i$ must be included. As the vertex centered functions are the only other functions on the fine grid that are non-zero at the constraint point, this equation is known and can be evaluated to determine the necessary interpolation operator. An illustration of this interpolation process is shown in figure 2.20 in 1 dimension.

$$\omega_i^f = \frac{\sum_j \omega_j^c \phi_j^c(\vec{\xi}_i) - \sum_k \omega_k^f(\vec{\xi}_i)}{\phi_i^f(\vec{\xi}_i)} \quad (2.16)$$

Multi-level interpolation: Full H^1 space

If the full H^1 vector representation is to be used, an interpolation operator can be defined by combining the $H^1(\text{Curl})$ and $H^1(\text{Grad})$ operators into a single operator. Though complication is introduced depending on how the redundancy discussed in section 2.2.4 is handled. During restriction a portion of the linear gradient subspace on fine levels gets mapped to the 0-th order gradient space on coarse levels. This will not be a problem for the smoother application on the sublevels as the effect will be canceled by the equivalent subspace that is solved for. However, during interpolation this gradient space must be mapped back to the space where these fields are being represented to prevent the build up of a large cancellation term in the solution. This term can grow very large and dominate the solution weights to the point where the desired dynamics reach the level of machine round off, preventing solver convergence.

2.2.5 Multi-grid preconditioning

The PSI-TET code has been designed from the ground up to facilitate and support the use of multi-grid preconditioners when solving linear systems arising from finite element discretization. Multi-grid is a method of preconditioning that has shown near ideal scalability for several types of problems using different discretization methods[56, 57, 58]. Conceptually this method can be thought of as a divide and conquer type method, where different regions of the eigenvalue spectrum for the discretized matrix are solved quasi-independently. Each of these regions can then be solved in a way that converges the solution over that portion of the solution space much more rapidly than is possible when treating the entire matrix. The majority of multi-grid methods attempt to construct sets of nested approximations to the original matrix by successively trimming the largest magnitude end of the eigenvalue spectrum. This is effective as most classic iterative methods, as well as common Krylov subspace methods, converge fastest for solution components that correspond

to the largest eigenvalues in the matrix. By generating a nested set of matrices where only the highest end of the spectrum is removed, solves can be performed which target progressively lower regions of the original matrix's spectrum in a more efficient manner.

Many types of multi-grid algorithms have been developed that employ different methods of coarsening the original matrix, though the majority of these methods fit into one of two categories: 1) Algebraic[59] multi-grid uses direct analysis of the original matrix to attempt to construct an approximate matrix that retains the low end of the eigenvalue spectrum while removing the high end. 2) Geometric[60] multi-grid uses progressively coarser representations for the solution and the problem is discretized on each level, which ideally removes only the high eigenvalue/wavenumber solutions from the representation while retaining the low wavenumber solutions.

PSI-TET supports geometric multi-grid natively, with support for algebraic multi-grid available through the PETSc interface. In PSI-TET coarse representations are constructed by forming a hierarchy of finite element representations. Linear finite element representations are constructed on each grid level produced during geometric refinement, see section 2.2.1. On the finest grid, finite element representations are also constructed up to the desired polynomial order. These provide a continuous sequence of coarsened representations from solution level to the smallest available grid on which approximations to the original matrix are constructed. With approximation levels available the preconditioner can be implemented using the interpolation methods described to transfer solutions between levels. In PSI-TET a V-Cycle type multi-grid iteration is used, which is shown schematically in figure 2.21 where representation levels are shown as green dashed lines, level to level restriction and interpolation are shown as red and yellow arrows respectively, and smoothing is shown as blue circles.

In a V-Cycle iteration smoothing is applied once on each cycle in the sequence per preconditioner application. Smoothing consists of performing an approximate solve to update the current solution on the given level. During the coarsening step, the




Figure 2.21: Diagram of 4 level V-Cycle iteration.

residual is transferred to the next coarsest level using the restriction operator, which is defined as the transpose of the interpolation operator between those levels. During the refinement step the solution is transferred to the next finest level using the interpolation operator. Solvers used to smooth the solution on each level may be the same or different depending on the problem. A common smoother type used in PSI-TET is a Jacobi method with different specified over-relaxation on each level. This method produces good scalability for simple problems, while maintaining the symmetric positive definite (SPD) nature of a matrix to allow the conjugate gradient method to be used as the main solver. Multi-grid in PSI-TET has been found as beneficial in reducing the solution time in a number of cases. Scaling for the benchmark verification problems above is shown in this section, while scaling for individual physics cases is presented in sections 3.1.1 and 5.3.

Lagrange Elements

The convergence benchmark presented in section 2.2.3 was run with and without multi-grid preconditioning to illustrate the performance of multi-grid with a Lagrange




Figure 2.22: Iteration scaling for a scalar Lagrange tests case with Jacobi preconditioning.

representation on a elliptic problem. A Conjugate-Gradient method was used as the main solver for both cases with a Jacobi, diagonal scaling, preconditioner (figure 2.22) and a multi-grid preconditioner (figure 2.23). The multi-grid preconditioner used all available levels for each test case, $N_{levels} = N_g + np$, and a symmetric version of the Jacobi method as the smoother on all levels. Elliptic equations generally scale quite well with multi-grid solvers and that result is reproduced here where the iteration count shows only minimal growth compared to the standard case.

Nedelec Elements

The behavior of multi-grid with Nedelec $H^1(\text{Curl})$ elements is observed to be different than the Lagrange elements. This is somewhat expected as multi-grid generally works best on nodal finite elements due to the importance of diagonal dominance. As with the Lagrange basis we run the convergence benchmark presented in section 2.2.4 with and without multi-grid preconditioning. A Conjugate-Gradient method was used as the main solver for both cases with a Jacobi, diagonal scaling, preconditioner (figure




Figure 2.23: Iteration scaling for a scalar Lagrange tests case with multi-grid preconditioning.

2.24) and a multi-grid preconditioner (figure 2.25). The multi-grid preconditioner used all available levels for each test case, $N_{levels} = N_g + np$, and a symmetric version of the Jacobi method as the smoother on all levels.

The first difference between the Nedelec and Lagrange representations is that the iteration count for the standard iterative solver scales more weakly for the number of weights for the $H^1(\text{Curl})$ test case. This is most likely due to the element definition that was chosen to maintain diagonal dominance with increasing element order. The second difference is that the iteration count with multi-grid preconditioning scales more strongly with the matrix size as the polynomial order is increased than in the Lagrange case. The reason for this difference could be attributed to a number of factors, each of which cannot be easily tested for causality. First, the problem being solved is not fully equivalent to the scalar problem and there is no concrete theory on the exact behavior of multi-grid preconditioning on a given equation system to guide the expected scaling. Second, the property of the basis sets that cause weaker scaling in the standard Jacobi preconditioned case may alter the multi-grid behavior.




Figure 2.24: Iteration scaling for a vector $H^1(\text{Curl})$ tests case with Jacobi preconditioning.

Although the iteration count increases somewhat with polynomial order on a given grid, the iteration count remains fixed when refining the grid. This indicates that the preconditioning is providing a significant speedup, however it is unclear whether or not the solver efficiency is uniform on polynomial versus geometric levels.




Figure 2.25: Iteration scaling for a vector $H^1(\text{Curl})$ tests case with multi-grid preconditioning.

Chapter 3

3D IDEAL MHD EQUILIBRIA

In equilibrium a magnetic field configuration must satisfy the force balance in equation 1.12. Additionally, the fluid pressure and λ are constrained to be constant on each flux surface. Previous efforts have been made to generalize from the well studied axisymmetric case[35]. The VMEC code[61] is the most successful of these methods which allows equilibria in 3D by requiring nested flux surfaces resulting in a simplified equation. An energy minimization is then used to compute equilibrium states. This method has had success in the Stellarator optimization and fitting codes STELLOPT[62] and V3FIT[63]. These methods have also been used more recently to model non-axisymmetric equilibrium states in tokamaks. The IPEC code[64] computes equilibria without the constraint of closed flux surfaces by adding perturbed fields in the linear regime around a known equilibrium. This code has been used primarily to study small deviations from equilibrium in tokamaks with 3D effects[65]. A non-linear extension to VMEC has also been developed into the SIESTA[66] code, which utilizes an energy minimization including non-ideal effects to allow the opening of islands within the equilibrium. Finally, the PIES[67, 68] code uses a flux coordinate grid, constructed via field line tracing, to allow general 3D equilibrium.

A common requirement of all the codes above is that the entire plasma volume must be representable using helical coordinates and only vacuum fields may exist beyond the simulated domain. This precludes the modeling of HIT-SI where the injectors cannot be represented as a contiguous helical system. Although the closed flux separatrix in HIT-SI often meets this criteria, its position and shape are strongly dependent on the plasma response beyond this surface, which must be solved for as

well. In this chapter the development of a general method for computing 3D ideal MHD equilibria is presented. Although HIT-SI is the primary motivation for this work the code developed does not make any geometry dependent simplifications and should be applicable to any other plasma confinement concept. The method of solution and formulation described however is specific to the use of $H^1(\text{Curl})$ elements to represent the vector potential.

3.1 Composite Taylor States and Uniform λ Equilibria

When λ is uniform in the volume, the constraint to be a flux surface variable is satisfied automatically without any need to determine the structure of magnetic surfaces. In this simplified case, equation 1.12 reduces to a relatively straight forward linear equation for the plasma response or an eigenvalue problem when no external fields are present.

Formulation

Eliminating the pressure term from equation 1.12 yields equation 1.6. In order to solve this equation in PSI-TET the magnetic field is first decomposed into an external vacuum component (B_v), which has no currents inside the plasma, and a plasma component (B_p), which contains the plasma response to currents driven on external fields. Using this representation equation 1.6 can be re-expressed as equation 3.1

$$\mu_0 J_p = \lambda (B_p + B_v) \quad (3.1)$$

where $\mu_0 J_p = \nabla \times B_p$. The equation in this form would be difficult to solve using finite element methods due to the poorly conditioned matrix produced by discretizing the curl operator. To alleviate this, the formulation is converted to a vector potential representation for the plasma component of the magnetic field, resulting in equation 3.2 which can be solved more efficiently using PSI-TET. For HIT-SI the case of a




Figure 3.1: Magnetic structure of the first two force-free eigenstates in HIT-SI.

perfectly conducting wall is of interest, which enforces $\vec{A} \times \hat{n} = 0$ and $\vec{B} \cdot \hat{n} = 0$ at the boundary of the plasma. The constraint on normal current due to HIT-SI's insulating coating will also be satisfied automatically by this boundary condition, which can be seen by considering the normal component of equation 3.1, $\vec{J} \cdot \hat{n} = \lambda \vec{B} \cdot \hat{n} = 0$.

$$[\nabla \times \nabla \times - \lambda \nabla \times] A_p = \lambda B_v \quad (3.2)$$

Internal Force-Free States

In the case where no external field is applied, the left hand side of equation 3.2 represents an eigenvalue problem for internal plasma current configurations that are themselves force free. The lowest eigenvalue solution to this equation corresponds to the Taylor state. Magnetic flux traces of the first two eigenstates for HIT-SI are shown in figure 3.1.

3.1.1 Solution Method and Scalability

As the eigensystem is symmetric and positive definite, a non-linear variant of the widely used Conjugate-Gradient (CG) iterative method is used to solve the system. This method is accelerated using a multi-grid approach for both the initial guess and the iterative solve. In order to provide a suitable initial guess, accelerating the non-linear portion of the solve, the eigensystem is solved on each grid level from coarsest to finest. The initial guess for each level is then set by interpolating from the solution computed on the previous level. A random initialization is used for the coarsest level. The iterative solve is then preconditioned using a multi-grid method with a Jacobi iteration used for smoothing on each level. Jacobi iteration is applied in a symmetric way on the up and down portions of the multi-grid to maintain the SPD property of the system required by the CG solver. The discretized version of the $\nabla \times \nabla \times$ operator is used for preconditioning.

To test the scalability of the preconditioner, the lowest eigenvalue solution is computed for a 1x1 cylinder for a sequence of grid resolutions and polynomial representations. Figures 3.2 and 3.3 show the iteration count for the Jacobi and multi-grid preconditioned methods respectively. Overall, the scalability of this eigenvalue solver is very similar to the $H^1(\text{Curl})$ linear system test case shown in section 2.2.5. A slight increase in the scaling of iteration count with system size is seen for the Jacobi case, while a decrease is seen for the scaling with polynomial degree when multi-grid is applied. These differences may be explained by the non-linearity of the system dominating the solution time. In figure 3.3, the iteration count is seen to scale strongly with polynomial order for low grid resolution, while iteration count remains essentially fixed at higher grid resolutions. At low grid resolution, preconditioning of the non-linearity is less effective, as the converged eigenvalue changes significantly with polynomial order.




Figure 3.2: Iteration scaling for Taylor state calculation in a cylinder with Jacobi preconditioning.




Figure 3.3: Iteration scaling for Taylor state calculation in a cylinder with multi-grid preconditioning.

3.1.2 Composite States in HIT-SI

As discussed above, the left hand side of equation 3.2 alone constitutes an eigenvalue equation. As a component of the solution space vanishes on the left hand side at any eigenvalue of this system, a null space is produced in the operator $(\nabla \times \nabla \times -\lambda \nabla \times)$. This introduces the added constraint that equation 3.2 is not solvable if $\lambda = \lambda^i$ is an eigenvalue and the external field couples to the corresponding eigenfunction A_p^i , requiring that equation 3.3 be satisfied at eigenvalues of the plasma response.

$$\int (\nabla \times A_p^i \cdot \lambda B_v) dV = 0 \quad (3.3)$$

In PSI-TET the external vacuum field is expressed as the gradient of a scalar potential, equation 3.4a, that may be multi-valued when the geometry is multiply connected. The vacuum field for a specified normal flux on the boundary, B_n , and internal fluxes passing through topologically distinct loops can be solved for using equations 3.4b and 3.4c - where χ has specified jumps at cut planes to produce the desired internal fluxes.

$$B_v = \nabla \chi \quad (3.4a)$$

$$\nabla^2 \chi = 0 \quad (3.4b)$$

$$\nabla \chi \cdot \hat{n} = B_n \text{ on } \Gamma \quad (3.4c)$$

Magnetic flux density, $|\vec{B}|$, and in plane magnetic field are shown in figure 3.4 for an example vacuum field in HIT-SI with unit flux in a single injector. The HIT-SI flux conserver is designed to exclude normal flux on millisecond time scales enforcing $\vec{B} \cdot \hat{n} = 0$ on the boundary. In HIT-SI the injector vacuum fields do not couple to the Taylor state eigenmode as they are composed of harmonics with odd symmetry only in the toroidal direction ($n = 1, 3, 5, \dots$). As a result, composite states can be computed at the Taylor state eigenvalue $\lambda = \lambda_{Taylor}$, where the toroidal flux or current is a free




Figure 3.4: Injector vacuum magnetic field in HIT-SI computed using multi-connected jump method.

parameter. Figure 3.5 shows a composite Taylor state, $\lambda = \lambda_{Taylor}$ and $p = constant$, with a current amplification of 6, in the range of observed amplifications.

3.2 Ideal MHD Equilibrium with Variable λ

As outlined in the introduction, section 1.2.2, the HIT-SI device drives the injectors with a specified current and flux which in general result in a λ in the injector region well above the Taylor value. Self-organization, towards the Taylor state, which forms the spheromak is expected to drive and maintain the average λ in the volume near λ_{Taylor} . This process leads to a variation in λ from the injector region to the interior, allowing high current in the injector while maintaining the global state near the minimum energy state. Before proceeding it is also useful to break λ into a spatial varying component with unit amplitude $\hat{\lambda}$ and a scale factor $\bar{\lambda}$.

For cases where we allow $\hat{\lambda}$ to vary in space, there is a restriction placed on $\hat{\lambda}$ by equation 1.6. Taking the divergence of this equation yields equation 3.5, which

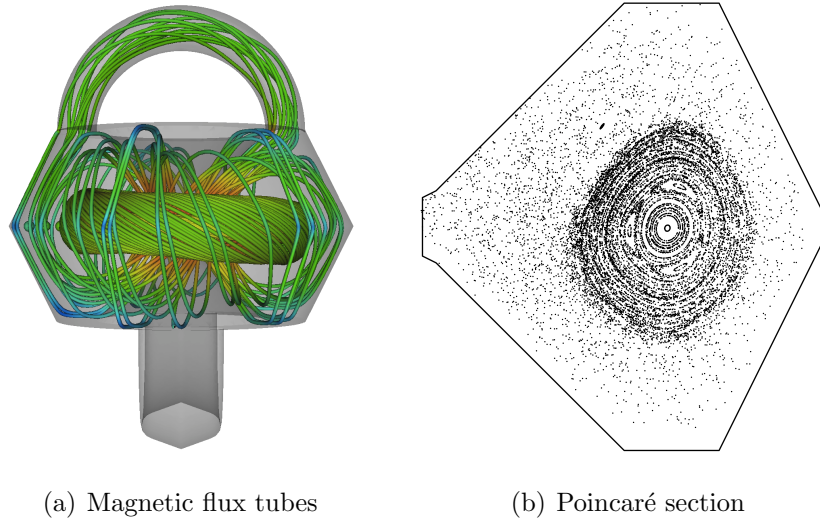


Figure 3.5: Magnetic structure for a Taylor state with flux amplification of 6.

requires that $\hat{\lambda}$ be a flux surface variable as in the Grad-Shafranov case.

$$0 = \nabla \hat{\lambda}(x) \cdot B \quad (3.5)$$

3.2.1 Flux Surface Identification in 3D

In order to set $\hat{\lambda}$, or p in the case of an equilibrium with pressure, the location of flux surfaces must be known on the finite element grid. Traditionally, general 3D methods such as the PIES code have used field line tracing to construct a field aligned grid, where the logical space is analogous to flux coordinates[69]. This method provides flux surface location intrinsically through the grid, however flux coordinates are not amenable to the HIT-SI case as mentioned above. Attempting to map a field line traced surface geometry onto a finite element grid would also introduce issues related to the discrete to continuous mapping.

To produce a suitable flux surface coordinate, the PSI-TET equilibrium solver leverages thermal diffusion with a high degree of anisotropy. The steady state solution




Figure 3.6: Thermal diffusion solve on a Taylor state in a oblate cylinder, comparing the computed surface (red), magnetic flux tubes (gray), and Grad-Shafranov flux surfaces (2D contours).

to thermal diffusion is computed with a large anisotropy aligned with the direction of the magnetic field. Equation 3.6 governs the physics of this model, where \hat{b} is the unit vector along the magnetic field of interest, I is the identity tensor, and χ_{\parallel} and χ_{\perp} are the parallel and perpendicular conductivities respectively.

$$\nabla \cdot \left[\chi_{\parallel} \hat{b} \hat{b} + \chi_{\perp} (I - \hat{b} \hat{b}) \right] \nabla T = S \quad (3.6)$$

The high parallel conductivity drives the resulting temperature to a constant along magnetic field lines. With a positive source S the solution should increase monotonically from the boundary, where a Dirichlet boundary condition is applied, to the magnetic axis. Figure 3.6 shows an example of a resulting temperature surface in comparison to the computed magnetic field lines and analytic poloidal flux contours for an axisymmetric Taylor state. The computed surface agrees well with both the computed magnetic field used to construct the diffusion problem as well as the corresponding poloidal flux contour which can be computed analytically using equation 2.2.

This thermal diffusion problem is challenging in its own right to solve accurately due to the high anisotropy. Although nested flux regions require a relatively modest anisotropy, island features are only resolved above a threshold in the ratio $\frac{\chi_{\parallel}}{\chi_{\perp}}$ that can be in excess of 10^8 . The specific ratio required is dependent on geometry and determined by a balance between the parallel and perpendicular conduction lengths across the island. Figure 3.7 shows the development of an inflection point and eventually a flat spot in the temperature profile across a magnetic island as $\frac{\chi_{\parallel}}{\chi_{\perp}}$ is increased. Temperature profiles in figure 3.7(b) are taken along the red line indicated in figure 3.7(a). Numerical error acts to artificially increase χ_{\perp} , decreasing this ratio in practice and obscuring small scale features such as these islands. Therefore, it is important to reduce the diffusion error to a level sufficient to model the desired conductivity ratios. The use of high order finite elements has been demonstrated to resolve diffusion problems with high anisotropy[1]. Figure 3.8 shows this effect in PSI-TET by comparing the relative temperature error, equation 3.7, for different element orders using the Taylor state in a cylinder as the guide field. For this test the field is axisymmetric, however as the grid is fairly isotropic, no grid alignment should be expected and the behavior should be characteristic of any field with similar spatial scale (low λ).

$$\epsilon_T = \sqrt{\frac{\int (\nabla T \cdot \hat{b})^2 dV}{\int T^2 dV}} \quad (3.7)$$

The applicability of this method has also been verified on more complicated fields with asymmetries in the toroidal direction, as in figure 3.7. Figure 3.7 shows a comparison of computed temperature surfaces and a Poincaré section for the same Taylor state field in HIT-SI. The island structure near the $q = 1/2$ surface is present in the Taylor state due to asymmetry induced by the injector mouths. Both the nested flux region as well as the boundary and internal structure of the prominent magnetic island are resolved using the diffusion method. These calculations were performed using a mesh with an average edge length of 2.26 cm and quartic elements for both

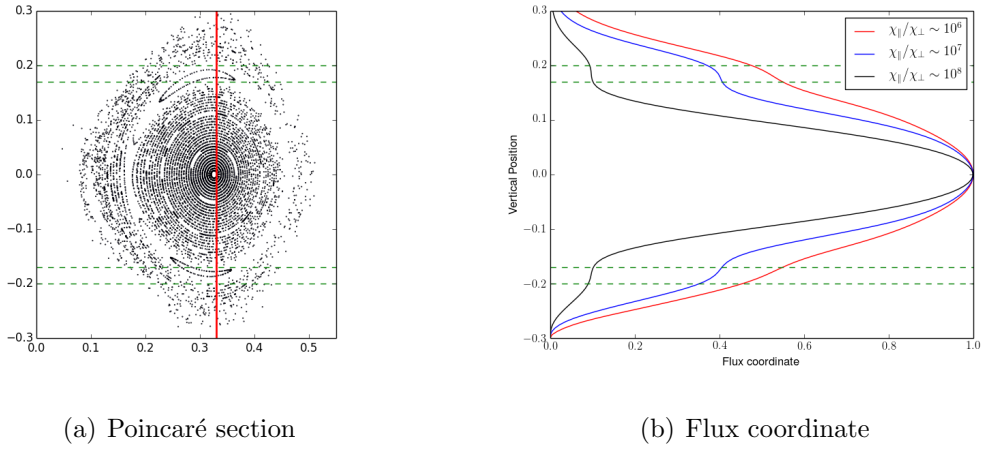


Figure 3.7: Comparison of the computed flux coordinate and the location of a large magnetic island in the HIT-SI Taylor state for a range of transport ratios.




Figure 3.8: Reduction of temperature error with representation order.




Figure 3.9: Comparison of computed flux surfaces (color) and Poincaré section (black) of the magnetic field in the vicinity of a magnetic island in the HIT-SI Taylor state.

the Taylor vector potential and temperature.

3.2.2 Eigenmodes and Plasma Response

The HIT-SI injectors have odd symmetry and produce no net toroidal flux in vacuum or current with uniform λ plasma response, as discussed in section 3.1.2. For the uniform λ case, this resulted in a decoupling where the toroidal current of the spheromak could be scaled independently of the injectors. When λ has variation across flux surfaces, equation 3.3 is no longer satisfied at the lowest eigenvalue due to coupling introduced through the $\hat{\lambda}$ factor in the integral. As a result the solution cannot be decoupled as in the composite Taylor state case by setting the scale to the lowest eigenvalue. Instead a resonance is produced between the externally applied vacuum fields and the plasma response as $\bar{\lambda}$ approaches the lowest eigenvalue of the system. Figure 3.10 shows this effect on the toroidal flux produced as $\bar{\lambda}$ is varied near the eigenvalue.

As the eigenvalue is approached, coupling of the vacuum field term to the lowest eigenmode when solving equation 3.2 causes the magnitude of this eigenmode to

grow toward infinity. The eigenmode of the linear system given by the discretized $(\nabla \times \nabla \times -\sigma \hat{\lambda} \nabla \times)$ operator is the same as from the general eigenvalue problem $\nabla \times \nabla \times A = \bar{\lambda} \hat{\lambda} \nabla \times A$, however the eigenvalues are shifted by $-\sigma$. Therefore, when σ approaches the lowest eigenvalue, $\bar{\lambda}_i$, the eigenvalue for the plasma response approaches 0. The magnitude of a given eigenmode in the solution follows equation 3.8, where M_i is the mode amplitude in the solution, $\bar{\lambda}_i$ is the original eigenvalue from the general eigenvalue problem, and C_i is the coupling coefficient produced by evaluating the integral in equation 3.3. This predicts a $\frac{1}{x}$ scaling for the toroidal flux and current that is observed in the calculations, figure 3.10.

$$M_i = \frac{1}{\lambda_i - \sigma} C_i \quad (3.8)$$

In HIT-SI, the flux surface structure is strongly dependent on the ratio of toroidal to injector flux/current. Figure 3.11 shows this effect by plotting the force-free error, $\int |\nabla \times \nabla \times \vec{A} - \lambda \nabla \times \vec{A}|^2 dV$, for different flux ratios produced by varying the scale of λ when solving equation 3.2. A minimum in the error is seen near a toroidal flux of 6, which was the toroidal flux on the previous iteration where $\lambda(T)$ was computed. This error is primarily due to the fact that $\nabla \cdot \lambda \nabla \times \vec{A}$ is large away from the flux ratio used to compute the magnetic surfaces due to changes in the magnetic topology. Subsequently, it is desirable to hold this value fixed when iterating to compute a new equilibrium state. This imposes that the value of $\bar{\lambda}$ is not a free parameter, but must instead be varied during the field update to maintain the desired ratios. Generally, ratios of interest are relatively large, in the range of 1 – 100, requiring solves very near the eigenvalue. This requires an accurate location for the systems lowest eigenvalue, which varies as flux surfaces and the resulting λ profiles change.

To obtain the eigenvalue a solve is performed along with each field update to compute the current value. The solver used to compute eigenstates in the uniform λ case cannot be used with variable λ however, as this variation causes the system to no longer be SPD, precluding a conjugate gradient solver. Instead, an Implicitly




Figure 3.10: Toroidal flux generation due to eigenvalue resonance in HIT-SI plasma response. Computed flux (red) agrees well with the theoretical scaling (blue).

Restarted Arnoldi (IRA) method[70] is used, which is provided through an interface to the ARPACK library[71]. This method relies on limited subspaces to approximate the eigenvalue spectrum and solver restarts similar to GMRES for continued convergence. Conjugate gradient methods exploit the SPD nature of the system to maintain the full Krylov space up to the current iteration. Although the fixed subspace size used by the IRA method is more general, it is also less efficient, resulting in increased solution cost. Additionally, computing the plasma response very close to the eigenvalue results in a poorly conditioned matrix, as the matrix is approaching singularity. Therefore, computing equilibrium states with high flux ratios is more computationally challenging than low flux ratios.

3.2.3 Solution Method

Solution of general equilibrium states is now possible with methods to locate flux surfaces, allowing surface variables to be set appropriately, and update the magnetic field induced by plasma. In order to compute an equilibrium state an initial guess




Figure 3.11: Force-free field error in plasma response calculation for different generated toroidal fluxes in HIT-SI with fixed injector flux of 1 mWb.

field must be used, either for a uniform λ calculation or a previous nearby equilibrium solve. Using this magnetic field, the flux coordinate is computed by solving the parallel diffusion equation 3.6 and $\hat{\lambda}$ is set and used to update required operators. If resonant coupling is required for the current equilibrium case — as with HIT-SI — the lowest eigenvalue of the plasma response system is computed. Finally, the plasma response field is updated by solving equation 3.2, adjusting the value of $\bar{\lambda}$ if necessary for resonant coupling. The flux coordinate and field updates are then repeated until the solution converges, shown schematically in figure 3.12.

Notes on scalability

The majority of the solvers used in computing updated equilibria can be preconditioned efficiently using multi-grid preconditioners. The parallel diffusion solve is the most troublesome of the solvers used in all cases, due to the stiffness of the operator matrix at high degrees of anisotropy. The operator can even become difficult enough to solve that multi-grid is required for convergence. If numerical round off




Figure 3.12: Diagram of iteration cycle used in 3D equilibrium solver.

error pollutes the low end of the eigenvalue space to a sufficient degree convergence of the conjugate gradient algorithm will stagnate. This process can prevent convergence to a solution even though a conjugate gradient method is theoretically guaranteed to converge with a sufficient number of iterations. Applying preconditioning with Jacobi smoothing allows the iteration to converge, however the iteration count is not scalable in the degree of anisotropy. However, this has not limited the usefulness of the method as the time required remains modest for all cases tested so far. There is a possibility that the scalability could be improved by leveraging more advanced smoothers during preconditioning, although this has not been explored.

When a plasma resonance is relied on for significant field production scalability, concerns arise for both the eigenvalue solve as well as the field update. For the field update, the stiffness can overwhelm even intense preconditioning using multi-grid as the operator approaches a singular matrix, corresponding to when large amplifications of the lowest eigenmode are required. This makes it difficult to compute equilibrium states for these high amplification cases. As these cases were not of interest for

this work, little development was done to improve this. It should be possible to accurately estimate the amplification required analytically from the field coupling and the known eigenvalue and eigenvector. This would allow for a very good initial guess for the plasma response update that may reduce the strain on the linear solve. Additionally, the preconditioner used for this solve currently uses a discretized version of the $(\nabla \times \nabla \times)$ operator and could be improved by using the full force free operator. This would require computing the eigenvalue for each representation level used, in order to keep the value of $\bar{\lambda}$ the appropriate distance from the numerical eigenvalue on each smoother level.

The eigenvalue solve alone has poor scalability and cannot be readily preconditioned due to the nature of the IRA method. In order to improve scalability the IRA solver can also be used in “shift-invert” mode, which can be used to enhance the convergence rate in a localized portion of the eigenvalue spectrum. This mode works by solving the shifted eigenvalue problem given by equation 3.10, where σ is the shift and A and M are operators from the original general eigenvalue problem 3.9, which correspond to the discretized $(\nabla \times \nabla \times)$ and $(\hat{\lambda} \nabla \times)$ operators respectively. Convergence for eigenvalues near σ is enhanced as the corresponding eigenvalues $\nu = 1/(\lambda - \sigma)$ for the shifted system are large in this range.

$$Ax = \lambda Mx \tag{3.9}$$

$$(A - \sigma M)^{-1} Mx = \nu x \tag{3.10}$$

The operator which must be inverted is then the force-free field operator used for the field updates, with σ set modestly close to the previous eigenvalue, usually $\sim 2\%$ away. This produces a stiffer operator to invert on each IRA iteration, though it reduces the number of eigenvalue iterations. In practice, the number of eigenvalue iterations have been observed to stay fixed with increasing resolution at the size of the

Krylov space, typically 20 vectors. This method effectively shifts the stiffness from the eigenvalue solve to the linear solve where it can be preconditioned more readily and efficiently.

3.2.4 Extension to Non-Zero Plasma Beta

The method described above has also been extended to cases with plasma pressure where force balance now requires equation 3.11 be satisfied. This equation still contains the force-free component addressed above, while adding a perpendicular component to the current density ($J_{\perp,\beta}$) that can be readily expressed in terms of the gradient of pressure (∇P) and the magnetic field as equation 3.12. In ideal MHD, the pressure must also be a flux surface quantity, as $\vec{J} \times \vec{B}$ cannot exert a force to balance the pressure gradient parallel to the magnetic field. As a result, P can be set as a function of the flux coordinate T as with λ .

$$\vec{J} \times \vec{B} = \nabla P \tag{3.11}$$

$$\vec{J}_{\perp,\beta} = -\frac{\nabla P \times \vec{B}}{|\vec{B}|^2} \tag{3.12}$$

In order to maintain charge neutrality, $\nabla \cdot \vec{J}_\beta = 0$, a parallel current ($\vec{J}_{\parallel,\beta}$), the Pfirsch-Schlüter current, must also be induced. To enforce that this current be parallel to the magnetic field, it is expressed as in equation 3.13a, where ϕ is some single valued scalar function. The divergence constraint on \vec{J} then takes the form of equation 3.13b, which may be solved using the anisotropic thermal diffusion mechanics addressed above. A contribution to the vector potential (\vec{A}_β) for the pressure driven field can then be computed using equation 3.13c with the same boundary condition of $\vec{A}_\beta \times \hat{n} = 0$.

$$\vec{J}_{\parallel,\beta} = \hat{b}\hat{b} \cdot \nabla \phi \tag{3.13a}$$

$$\nabla \cdot \hat{b}\hat{b} \cdot \nabla \phi = -\nabla \cdot \vec{J}_{\perp,\beta} \quad (3.13b)$$

$$\nabla \times \nabla \times \vec{A}_\beta = \vec{J}_{\parallel,\beta} + \vec{J}_{\perp,\beta} \quad (3.13c)$$

Finally, the pressure induced magnetic field must also be added into the calculation of the force-free plasma response, equation 3.2, to produce a new source term, equation 3.14. The steps to solve equations 3.12, 3.13b, and 3.13c can then be added before the plasma response update to produce a method capable of handling non-zero plasma pressure.

$$[\nabla \times \nabla \times -\lambda \nabla \times] \vec{A}_p = \lambda \left(\vec{B}_v + \nabla \times \vec{A}_\beta \right) \quad (3.14)$$

Example Equilibrium

Due to investigation of dynamic models performed in chapter 5, only preliminary investigation of this method has been performed as part of this work. Figures 3.13 and 3.14 show an example equilibrium computing in 1x1 cylinder with no external field, uniform λ , and a linear pressure profile in the psuedo-flux coordinate T . The resulting equilibrium has a $\langle \beta \rangle_{vol} = 2.2\%$ and exhibits a significant shift in the flux contours, figure 3.13. This example demonstrates the basic functionality of this method and ability of PSI-TET to solve the method as a whole. Further investigation of this method for application to non-axisymmetric devices and comparison with finite β configurations such as stellerators should be performed in the future.

3.3 Application to Driven HIT-SI Equilibria

The process outlined above has been used to perform investigation of driven equilibrium states in HIT-SI. A two- λ model, equation 3.15, was used which transitions from a constant value in the injector field region beyond the field separatrix to a lower value throughout the closed flux region. The separatrix flux value is chosen to match the position of the separatrix, which is typically set to $T_{sep} \sim .01$ for the HIT-SI




Figure 3.13: Comparison of flux surfaces between pressureless (dashed) and $\langle\beta\rangle_{vol} = 2.2\%$ (solid) equilibria in a 1x1 cylinder. Flux surfaces for the case with pressure were chosen to match the pressureless surfaces on the outboard side.





Figure 3.14: Comparison of magnetic field magnitude between pressureless (dashed) and $\langle\beta\rangle_{vol} = 2.2\%$ (solid) equilibria in a 1x1 cylinder.

geometry. Although this value is somewhat empirical, it is not very sensitive to the flux surface geometry due to the extremely low temperature gradient in the injector region when compared to the closed flux region. The same value was found to work well for identifying the separatrix location across all cases tested in HIT-SI. In order to further supplement separatrix location using anisotropic diffusion a boundary condition of $T = 0$ was used on the boundary of the device as well as at the mid-plane of the active injectors. The λ profile shape was chosen due to its simplicity and smooth transition between λ regions. Additionally, this profile matches the profile expected by IDCD theory, section 4.1, during the drive phase in HIT-SI. For these calculations injector vacuum fields were set to a phase during the injector cycle when only a single injector is operating and at peak flux and current, in this case the X-Injector was used.

$$\hat{\lambda}(\hat{T}) = \left(1 + \frac{\alpha}{2}\right) + \frac{\alpha}{2} \tanh\left(30 \left[\hat{T} - T_{sep}\right]\right) \quad (3.15)$$

Equilibrium states were computed for these conditions at a fixed ratio of toroidal to injector flux equal to 6, while varying α to produce a sequence of states with progressively higher λ in the injector region (λ_{Inj}). Each state used the previous equilibrium as the initial condition with the composite Taylor state, figure 3.5, used as the initial condition for the first state. As α was increased, a threshold was encountered beyond which the equilibrium solve no longer converged. This threshold resulted in a maximum attainable λ_{Inj} , which was found to be a function of flux ratio. The final position of the separatrix for the highest λ_{Inj} state at a flux ratio of 6 is shown in 2D on the mid-plane of HIT-SI and in 3D along with magnetic field traces of the injector region in figures 3.15(a) and 3.15(b). Significant toroidal asymmetry is visible in the separatrix position for this case, as is expected from variation in the position of the separatrix in the composite Taylor state at the same flux ratio. The separatrix is still well captured in this final state as shown in figure 3.16 where the flux contours for the final equilibrium state are compared against a Poincaré section of the magnetic



(a) Lambda profile on the midplane

(b) Separatrix surface and injector field lines

Figure 3.15: Ideal MHD equilibrium state in HIT-SI at maximum λ_{Inj} for a flux ratio of 6 with a two- λ model.

field where only field lines in the open, injector linking, region were used.

Figure 3.16 also shows the effect of variation in λ on the equilibrium field structure. Visible in the Poincaré section are three voids just beyond the separatrix. These voids correspond to an $m = 3$ island chain that is also present in the composite Taylor state, figure 3.5. As λ in the injector region is increased, and the resulting current profile becomes more hollow, the island chain moves from its initial position near the middle of the closed flux region to a position just beyond the closed flux region at the highest attainable λ_{Inj} . This migration can be explained as the q at the magnetic axis should increase, due to hollowing of the current profile, causing the $q = 1/3$ surface to move toward the wall, carrying its resonant island structure with it.

A scan was also performed to evaluate the dependence of maximum λ_{Inj} on flux




Figure 3.16: Comparison of flux contours and Poincaré section of open field structure on a poloidal slice at a toroidal position of $\pi/2$ for the equilibrium state shown in figure 3.15.

ratio. Figure 3.17 shows the final position of the separatrix for the highest λ_{Inj} state at a flux ratio of 20, with all other quantities fixed when compared to the flux ratio of 6 calculation. The first thing to note is that the separatrix position and the equilibrium field structure as a whole become more symmetric as the flux ratio is increased. Some asymmetry is still present in this state however and the separatrix location remains well resolved. In addition to being more symmetric, the separatrix also occupies a much larger fraction of the volume than at lower flux ratios, indicating a correspondingly larger ratio of closed to open flux volume. This larger separatrix volume allows high λ_{Inj} to be obtained, as use of the eigenvalue resonance for field amplification requires the volume average of λ to be maintained at $\sim \lambda_{Taylor}$ for all equilibria. When the separatrix volume is large, a small change in that region's — λ which should be stable — will lead to a large change in λ_{Inj} .




Figure 3.17: Ideal MHD equilibrium state in HIT-SI at maximum λ_{Inj} for a flux ratio of 20 with a two- λ model.

3.3.1 Discussion

As the flux ratio is increased further, higher λ_{Inj} can be achieved as expected (figure 3.18). However, λ is still limited to a value below that observed in the experiment — even at flux amplifications much greater than experimentally observed. In HIT-SI the injectors are typically run in excess of $\lambda_{Inj} = 15$ at observed flux ratios peaking around 8. Equilibria on the other hand have been limited to $\lambda_{Inj} < 12.5$ even when the flux ratio is increased to 40. It is unclear if this difference is due to a deficiency in the method or constraints in the physics of ideal MHD equilibrium. From a physics point of view, many of the constraints of ideal MHD are not expected to be satisfied in the experiment — although previous Taylor theory did agree well for low performance operation in helium. The experiment has a relatively high resistivity that would allow variation of λ along field lines, particularly in the injector region. This could allow the current to be carried in a narrower region than uniform λ requires, increasing the separatrix volume. Additionally, the plasma may not be in a quasi-static equilibrium during the drive phase, which would make equation 3.1 a poor approximation to the actual time dependent dynamics. Evidence of large flows, plasma pressure, and other non-static and non-ideal effects are seen in dynamic simulations of HIT-SI using NIMROD and PSI-TET. PSI-TET in particular sees non-uniform λ over the majority of the injector volume during simulations with drive, indicating that transient effects may dominate behavior in the injectors for most of the shot. It is important to note that the presence of dynamics in the actual experiment does not preclude the existence of equilibrium states at similar parameters. However, their presence does indicate that quasi-static models are inappropriate for comparison with at least the current experiment.




Figure 3.18: Maximum λ_{inj} for a range of flux amplifications in HIT-SI with a two- λ model.

3.4 Discussion and Future Work

The methods presented above provide a new way to compute general 3D ideal MHD equilibria. This approach is believed to be more general than previous tools developed to this point. Although direct comparison with HIT-SI did not provide useful comparison, application of this method to other configurations with smaller amplitude fluctuations may be beneficial over existing tools. Tokamaks[72], stellarators[73], and reverse field pinches (RFP)[74] all exhibit significant 3D structure, either due to intrinsic fields or spontaneous generation of asymmetries due to plasma effects. The PIES and SIESTA codes are both capable of general equilibria and have been applied to stellarator cases, with some initial investigation of tokamak states as well. However, these investigations have been limited, at least in published work. The PSI-TET implementation should be applicable to all these geometries and should have improved scalability to high resolution when compared to the HIT-SI case as

none of the other confinement devices operate with reliance on field amplification due to plasma resonance. This reduces the number of solves required to those that can be more effectively preconditioned. Application of PSI-TET to these cases, in comparison to the two existing tools, should be pursued in the future.

Chapter 4

REDUCED ORDER MODELS FOR SIHI DESIGN

For design of future experiments it is helpful to have simplified models and tools that can be used to guide initial design phases without the need for full scale simulation. These types of tools are common place in commercial engineering and are becoming increasingly important in magnetic confinement fusion as research and development programs move to large scale experiments and prototype devices. As discussed in chapter 3, tools for axisymmetric or helical equilibrium investigation are widely used in the tokamak and stellarator communities. However, for HIT-SI and future SIHI driven devices, the applied fluctuations, or some pressure drive profile consistency effects, are expected to be important in determining the equilibrium. To investigate simplified models to capture these effects reduced order models are explored in this chapter for the Imposed Dynamo Current Drive (IDCD)[32] model in HIT-SI. These models involve coupling 3D and 2D equilibrium calculations to predict and/or optimize equilibrium profiles for comparison with observed fields and evaluation of the future SIHI injector configurations.

Imposed Dynamo Current Drive

In previous spheromaks, self-generated fluctuation activity was present during sustainment and was often associated with flux conversion by dynamo activity. Dynamo activity can be analyzed from the perspective of fluctuations around a mean field equilibrium state[31, 18, 19]. In HIT-SI non-axisymmetric fluctuations are imposed by the injector fields leading to the expected Imposed Dynamo Current Drive[32]. IDCD drives current by shear induced in applied fluctuations resulting in a drive

force on a mean flux surface, which balances the total force applied on the plasma within that flux surface. The force that must be balanced at each surface is due to Ohmic dissipation acting on parallel currents and any electric field acting to increase those currents. This force can be expressed as equation 4.1, where the integral is over the volume enclosed by a given flux surface. The balancing drive force on each surface due to shear induced in fluctuations can then be evaluated by integrating over each flux surface. As flow shear only acts on the perpendicular component of δB the drive term can be simplified to a constant factor C_1 , that defines the maximum extent of distortion, and δB_\perp , the component of the fluctuating field orthogonal to mean flux surfaces, only. This gives equation 4.2 that must balance equation 4.1 for IDCD to maintain, or grow by balancing non-zero E_\parallel , the enclosed current. The constraint of force balance on each mean flux surface can in turn be used to determine the expected mean field equilibrium for a given fluctuation profile or conversely to optimize the fluctuation profile to produce a given mean field equilibrium.

$$F_{diss} = \int_V ne (\eta J_\parallel - E_\parallel) dV \quad (4.1)$$

$$F_{drive} = C_1 \int_{\partial V} \frac{\delta B_\perp^2}{\mu_0} dS \quad (4.2)$$

4.1 Current Profiles Produced by IDCD

First, we analyze mean field equilibria that are predicted by the IDCD force balance above for 14.5 kHz operation in HIT-SI. In HIT-SI the fluctuations are observed to increase from the wall to the magnetic axis. Empirical fluctuation amplitude is assessed using oscillations in the radial magnetic field (B_r) measured along an internal probe as shown in figure 4.1. For a symmetric equilibrium state the radial field is zero, making it a good proxy for the perpendicular fluctuation amplitude.




Figure 4.1: Evolution of B_r in time at different major radii along the HIT-SI internal magnetic probe. Note that each radial position is offset by 40 mT to facilitate display on a single plot.

4.1.1 Predicted Equilibria in HIT-SI

As the fluctuation amplitude increases toward the axis the required drive should be exceeded, overpowering the dissipation maintaining flow shear and causing the flow to lock, equation 4.3.

$$\int_V ne (\eta J_{\parallel} - E_{\parallel}) dV < C_1 \int_{\partial V} \frac{\delta B_{\perp}^2}{\mu_0} dS \quad (4.3)$$

This is expected to produce a profile with two constant λ regions and a single step, where the force balance is satisfied, from high edge λ to low interior λ . This profile can be parameterized by λ_{Inj} and the ratio I_{Tor}/I_{Inj} . In order to define the equilibrium the injector current paths are modeled as a symmetric component, which contributes a current to the mean field, and an asymmetric component that produces the fluctuations. Together these components produce an injector current channel which is aligned with the toroidal current, as expected in the presence of rapid reconnection[75]. This




Figure 4.2: Schematic of decomposition of the expected injector current path due to fast reconnection. The toroidal current (not shown) is oriented in the clockwise direction.

is shown schematically in figure 4.2, where the images from left to right show the asymmetric fluctuation, mean field, and reconnected current paths.

The symmetric current contribution should exist in the edge of the mean field equilibrium. The λ in this region is constrained to be equal to the injector value and must also contain the required toroidal current of $1/2I_{inj}$. Together these values fully define the position and height of the jump from edge to interior λ . This equilibrium profile agrees well with experimental observation of the time averaged field in HIT-SI by the internal probe. Comparison of the computed equilibrium magnetic field and data during the steady state portion of a high performance shot is shown in figure 4.3.

4.1.2 Evaluating Dynamo Force and Dissipation

As the HIT-SI flux conserver is nearly symmetric, the mean field equilibrium state can be approximated as an axisymmetric equilibrium that is computed using the PSI-TRI code. The PSI-TRI code allows for the flexibility to easily vary the current and

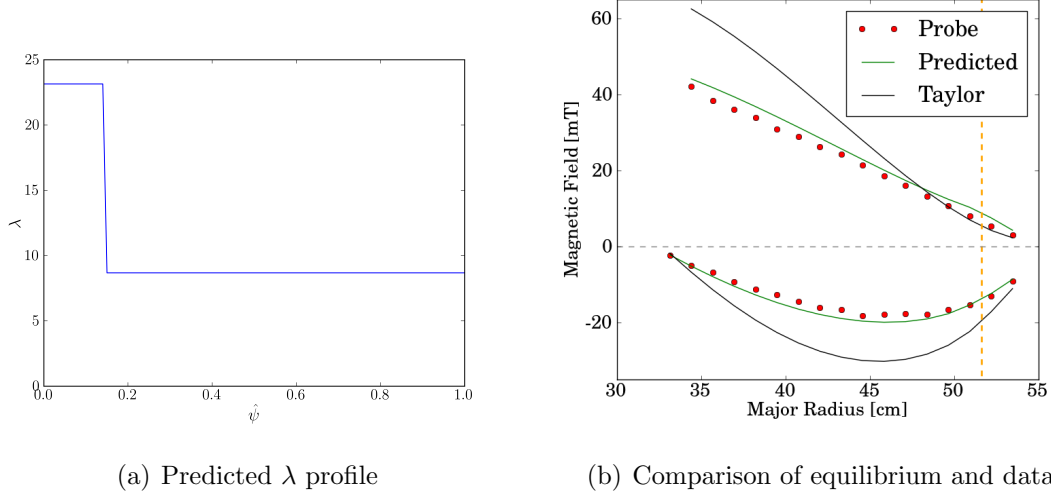


Figure 4.3: Mean field comparison at a characteristic time during high performance operation in HIT-SI (Shot 122385). The radial position of the separatrix is marked by the orange dashed line.

pressure profiles and also supports computing volume and flux surface integrals from the computed equilibria. The dissipation integral can be computed entirely from the mean field equilibrium, where E_{\parallel} is replaced with \dot{A}_{\parallel} to provide a means to balance growth or profile change in the equilibrium.

The drive integral requires the perpendicular component of the fluctuating field, in addition to the flux surface position, to compute. For HIT-SI a uniform $\lambda = \lambda_{Taylor}$ force-free equilibrium is found to agree well with the radial field along the internal magnetic probe, figure 4.1. As this field can be readily computed using the PSI-TET code, section 3.1, it is used to approximate δB_{\perp} in the surface drive term, equation 4.2. The 3D equilibrium field is computed in PSI-TET and then exported on a uniformly spaced toroidal grid. The field can then be read in to the PSI-TRI code and used to evaluate δB_{\perp} during flux surface integration using surface geometry from the axisymmetric equilibrium. An example of the variation of δB_{\perp} on a flux




Figure 4.4: Magnitude of δB_{\perp}^2 computed from a uniform $\lambda = \lambda_{Taylor}$ force-free equilibrium relative to a flux surface of the axisymmetric Taylor state in HIT-SI. The flux surface is shown from below, with the X-Injector mouths visible as the higher magnitude footprints on the flux surface.

surface of the axisymmetric Taylor state in HIT-SI is shown in figure 4.4. The visible “footprints” on the flux surface are due to higher field resulting from collimation near the injector mouths.

Using this model for the applied fluctuations and the experimentally validated mean field profile, we can compare the resulting drive and dissipation terms to check the hypothesis that the interior is over driven. A comparison of the surface drive force (F_{drive}) and the volume dissipation (F_{diss}) is shown in figure 4.5. The profiles are set equal at the location of the step in λ where the force balance is expected to be satisfied. This confirms that the model drive, which agrees with experimental data, appears to exceed the required value from resistive dissipation everywhere in the volume.

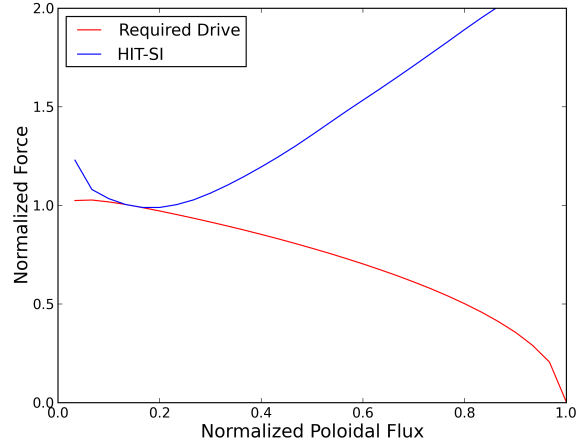


Figure 4.5: Comparison of IDCD dissipation and surface drive profiles for the empirically predicted mean field equilibrium in HIT-SI.

4.1.3 Predicted Current Evolution

The IDCD model is also used to construct a reduced order model for the current evolution in HIT-SI. This has been tested using only experimental inputs and constant assumptions for the fluctuation profile, see reference [32]. Here the process is repeated with the more sophisticated models for the fluctuation amplitude and dissipation electric fields presented above. Rewriting equation 1.15 in terms of only the mean field and making the same assumption as above to replace δB_{\parallel} with δB_{\perp} we get a differential equation for the mean field, equation 4.4. This can be used to model the current evolution through \dot{A} where the integrals are evaluated for the outermost surface where force balance is satisfied. For the predicted profile in HIT-SI, the flow is locked everywhere except for a single surface, corresponding to the mean field separatrix where the integrals are evaluated.

$$\int_V ne \left(\eta J_{\parallel} - \dot{A}_{\parallel} \right) dV = C_1 \int_{\partial V} \frac{\delta B_{\perp}^2}{\mu_0} dS \quad (4.4)$$




Figure 4.6: Path of mean field equilibrium evolution through parameter space in time when using equation 4.4 to model the toroidal current.

This equation can be integrated in time using experimental data for n_e , λ_{Inj} , and I_{Inj} that, along with I_{Tor} , uniquely define the equilibrium profile and magnitude of δB_{\perp} in time. The toroidal current is the only unknown in this set of parameters and is evolved based on the differential equation. As the profile evolves in time the position of flux surfaces and the resulting integrals change as the mean field moves through its defining parameter space. In order to capture this effect terms in equation 4.4 are evaluated to span the expected region of parameter space. Values are then interpolated between known points to provide consistent evolution of individual terms during the time integration. Figure 4.6 shows an example of the path of a shot (blue) in parameter space through the range of computed equilibria (red).

For comparison with experimental data the resistivity is replaced with an experimental value based on the observed resistive decay time, $\eta = \mu_0 / \tau_{L/R} \lambda_{Taylor}^2$. The free coefficient is then adjusted to provide the best agreement for a single shot and held constant for comparison with additional shots. Integration is started from a time after spheromak formation when $I_{Tor}/I_{Inj} > 1$ and IDCD is expected to be the




Figure 4.7: Comparison of I_{Tor} between experimental data (black), the IDCD empirical reduced model (red) and the IDCD uniform $\lambda = \lambda_{Taylor}$ force-free model (red) across multiple shots.

dominant effect. Shown in figure 4.7 is a comparison between experimental data at 14.5 kHz injector frequency, the reduced IDCD model presented in [32], and the coupled Taylor-mean field model presented here. Both the Taylor and reduced models are found to provide good agreement for shots that span a wide range in operational parameters.

4.2 Geometry Optimization and Design for SIHI

With a model for the IDCD force balance investigation and predictions can be made for future experiments using SIHI current drive. The fluctuating magnetic field, using the uniform $\lambda = \lambda_{Taylor}$ force-free assumption, is only dependent on the geometry and λ_{Inj} . Additionally, the mean equilibrium is dependent on the geometry, current ratio, and the current profile, which is related to the perturbing magnetic field. Ideally, the locked equilibrium state that is predicted to exist in HIT-SI should be avoided as the excess fluctuations in the interior may have a detrimental effect on confinement

in larger, hotter experiments. Beyond that, generating the field required to support these high amplitudes in the interior may place unnecessary power requirements on the driving circuits.

4.2.1 Predictions for HIT-SI3

To investigate the ability to tailor fluctuations produced by HIT-SI like injectors the drive and dissipation profiles for the HIT-SI3 experiment are considered for different injector phasing. HIT-SI3 was designed to operate with each of the three injectors phased 120 degrees apart in time as in figure 4.8, like circuits in a 3-phase motor, providing constant helicity injection while driving plasma rotation. The injectors could also be operated such that all flux and current oscillations are in phase, representing the opposite extreme of injector phasing. Using the uniform $\lambda = \lambda_{Taylor}$ force-free assumption surface drive profiles were computed for each phasing and compared to the dissipation and HIT-SI profiles, figure 4.9. These results indicate that it should be possible to use the injector phasing to vary the drive profile predicted by IDCD in HIT-SI3, providing a further test for this current drive theory. Alternatively, if pressure driven effects are the dominant driver for the mean field current profile then phasing should allow for the control of fluctuation amplitude in the interior. By varying the fluctuation amplitude a threshold could be found, where any fluctuation induced transport is at a minimum while maintaining regulation of pressure driven modes as predicted by experimental data.

4.2.2 Finding Matched λ -Profiles

In addition to varying the phase of SIHI injector sets, the geometry and resulting mean field profiles can be adjusted to provide a better match between dissipation and drive forces predicted by IDCD. These factors were considered in analyzing the design of a reactor scale SIHI driven experiment[76], HIT-POW, whose injectors were optimized for the fluctuation profile. The size and spacing of injector mouths were




Figure 4.8: Example force-free injector equilibrium state in HIT-SI3 with the 3-phase configuration.




Figure 4.9: Comparison of IDCD surface drive profiles for HIT-SI (blue), the HIT-SI3 3-phase configuration (solid green), and the HIT-SI3 in-phase configuration (dashed green) to the dissipation profile in steady state.




Figure 4.10: Comparison of IDCD dissipation and surface drive profiles for mean field equilibria with uniform (dashed) and optimized (solid) λ profiles.

set based on constraints for the neutron blanket system as well as assessment of the resulting imposed fluctuations. A configuration was chosen that provided a surface driven profile that resulted in insufficient drive over a large region of the poloidal flux, with excess drive outside of this region. A mean field equilibrium with uniform λ and a pressure profile producing marginal stability with respect to the Mercier criterion was used.

To assess the effects of modifying the mean field profiles the equilibrium was then modified to match the IDCD predicted drive and dissipation terms across the previously under driven portion of the volume. For this optimization the λ profile was adjusted to match the force terms as closely as possible, while maintaining the pressure profile with marginal stability. A comparison of the force terms before and after matching is shown in figure 4.10 over the region where λ was allowed to vary. Matching in this region required hollowing the current profile to increase the rate at which dissipation decreased as a function of flux. The resulting λ and P profiles produced are shown in figure 4.11 in comparison to the unoptimized profiles.

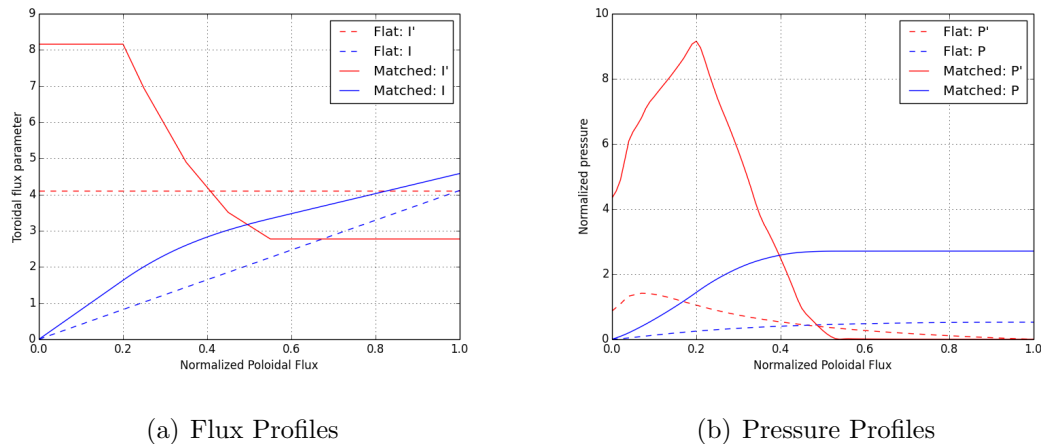


Figure 4.11: Comparison of flux surface quantities for mean field equilibria with uniform (dashed) and optimized (solid) λ profiles.

4.2.3 Discussion

The ability to match the two terms in the IDCD force balance over a significant volume as in this case indicates that optimization of SIHI injector geometry in a more comprehensive way should be possible. Figure 4.12 shows a diagram of the optimization work flow using the steps presented for the HIT-POW case here. An initial geometry can be defined in CAD software that is then used to generate 3D and 2D meshes for the PSI-TET and PSI-TRI using the CUBIT interfaces. The 3D geometry is then used to evaluate the fluctuation fields in 3D, which are imported into PSI-TRI. Mean field equilibrium profiles can then be optimized to match the drive and dissipation forces across some specified region of the flux and used to evaluate some cost metrics to update the initial geometry. This process can then be repeated until the desired performance characteristics are met. An automated process such as this would rely on appropriate metrics and methods to modify the geometry of course, which has not been addressed for this work. Similarly, this process relies on IDCD being an appropriate model for the current drive in future SIHI experiments,

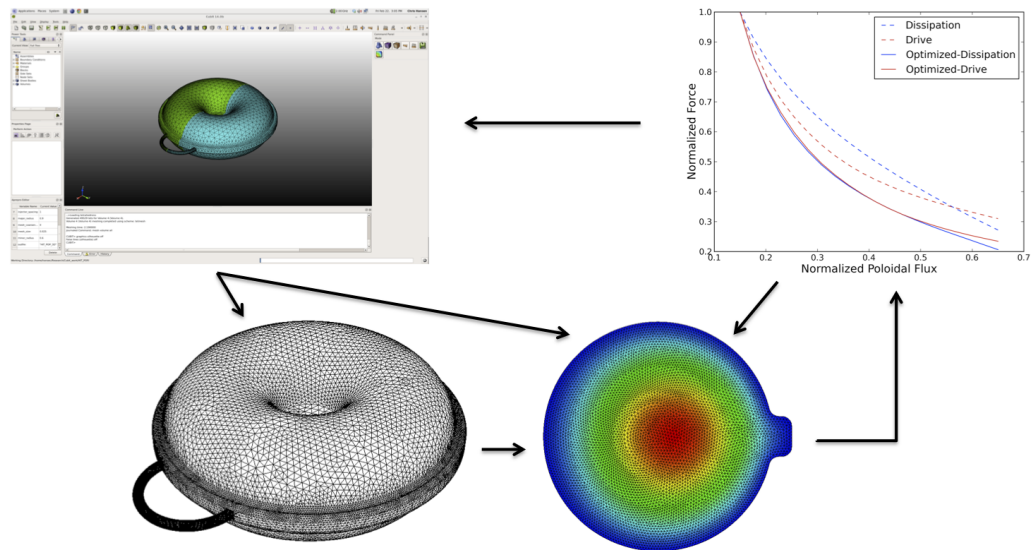


Figure 4.12: Schematic of geometry and profile optimization work flow for SIHI driven devices using IDC'D to define drive and profile constraints.

which is an open question. The general integration of 3D and 2D tools developed and presented here could support a wide range of different models however and should be amenable to future models as well.

Chapter 5

3D NON-LINEAR HALL-MHD

As discussed in the introduction, previous modeling efforts for the HIT-SI experiment have focused on using the extended MHD code NIMROD. Although these studies have shown good agreement with experimental data there exists some unanswered questions about the approximations used in adapting the HIT-SI geometry for use with NIMROD. To explore these effects the PSI-TET code has been used to develop a 3D reduced Hall-MHD module capable of modeling the full HIT-SI geometry. This code currently incorporates the same physical model as the previous studies of HIT-SI by Akcay[2]. The equations used are equations 5.1a - 5.1c and constitute the reduced Hall MHD model assuming uniform plasma density and temperature, with simplified kinematic viscosity and electron inertia.

$$\rho \left[\frac{\partial u}{\partial t} + u \cdot \nabla u \right] = J \times B + \nabla \cdot \nu \nabla u \quad (5.1a)$$

$$\frac{\partial B}{\partial t} = -\nabla \times \left(-u \times B + \eta J + \frac{1}{ne} J \times B + \frac{m_e}{ne^2} \frac{\partial J}{\partial t} \right) \quad (5.1b)$$

$$\mu_0 J = \nabla \times B \quad (5.1c)$$

These equations have been shown to provide good agreement with experimental data in previous studies using the NIMROD code. As a result they were chosen as the baseline for PSI-TET simulations to provide a well studied comparison to existing results. These comparisons are useful from both a code validation perspective as well as for isolating physical effects due to the addition of injector dynamics.

The following sections will discuss: 1) The implementation of the Hall-MHD model using a mixed Nedelec-Lagrange field expansion. 2) Validation studies performed to verify the correct evolution of equations 5.1 and characterize the scaling and behavior of the solution methods. 3) Comparisons to experimental data and existing NIMROD simulations of HIT-SI for validation purposes. 4) Analysis of previously unstudied physics in the helicity injectors and their coupling to the spheromak volume.

5.1 Hall-MHD with Nedelec Elements

A Nedelec element basis consisting of the full H^1 vector space is used to represent the magnetic field in the present implementation, while a more standard Lagrange scalar representation is used for each component of the velocity field. The primary motivation for this choice was to support a boundary condition for the magnetic field that prevents current from penetrating the boundary. This boundary condition and its structure is discussed further in section 5.4 to follow. In addition to enabling an appropriate boundary condition for HIT-SI this representation has a number of unique properties which are analyzed and discussed below.

5.1.1 Vector Sub-space Decomposition

As presented in section 2.2.4, Nedelec elements provide a natural decomposition of the vector space. This decomposition separates the basis functions that support a curl from those that do not. As a result the magnetic induction equation in weak form takes the rather interesting form, equation 5.2, where v_g, v_g^T and v_c, v_c^T are the basis and test functions for the $H^1(\text{Curl})$ and $H^1(\text{Grad})$ spaces respectively and \dot{b}_i are the corresponding function weights. The term on the right hand side vanishes analytically as $\nabla \times v_g^T = \nabla \times (\nabla \phi^T) = 0$, where ϕ^T is the parent scalar H^0 test function.

$$\begin{pmatrix} \int v_g^T \cdot v_g dV & \int v_g^T \cdot v_c dV \\ \int v_c^T \cdot v_g dV & \int v_c^T \cdot v_c dV \end{pmatrix} \begin{pmatrix} \dot{b}_g \\ \dot{b}_c \end{pmatrix} = \begin{pmatrix} 0 \\ -\int (\nabla \times v_c^T \cdot E) dV \end{pmatrix} \quad (5.2)$$

This implies that only the subset of the basis function weights that support a curl, and thus plasma current, are directly related to the electric field in the plasma. The remaining degrees of freedom are only coupled to the electric field through the H^1 metric matrix present on the left hand side of the equation. Note that equation 5.2 ignores the boundary term generated by moving the curl operator from the electric field to the test functions. For boundary conditions where the surface term cannot be neglected there remains a contribution to the $H^1(\text{Grad})$ rows on the left hand side. For the purposes of the rest of this study however a perfectly conducting boundary is always assumed so that the surface term may be neglected.

Null Space Considerations

The full H^1 vector representation has a null space resulting from linear dependence of zeroth order gradient terms expressed either on the vertex $H^1(\text{Grad})$ or edge $H^1(\text{Curl})$ basis functions, see section 2.2.4. As a result special care must be taken to eliminate this redundancy so that singular matrices are not generated. This is done by setting the vertex $H^1(\text{Grad})$ weights to zero.

5.1.2 Divergence Error with a Nedelec Basis

As the primary evolved variable is the magnetic field, divergence error is of concern. In existing codes that evolve the magnetic field with standard C0 finite element methods special care must be taken to prevent divergence error from effecting the solution. In NIMROD an additional dissipation term[18, 15] is added to the induction equation to “clean” divergence error created by inaccuracy in the null space of the curl operator. This method is successful by maintaining divergence error at a sufficiently low level

so as not to impact the solution without adding too much spurious dissipation in the process. Discontinuous methods, which are locally divergence free, have also been applied to the MHD system[77]. However, in these methods explicit penalty terms, analogous to NIMROD divergence cleaning, must be added to ensure magnetic flux between cells is consistent[78].

In PSI-TET the Nedelec decomposition has an interesting effect on the behavior of divergence error in the magnetic field. If we analyze the $H^1(\text{Grad})$ rows of the induction equation more carefully we find that they actually represent a weak form divergence “constraint”, equation 5.3. In this equation the components of the magnetic field have been combined into the change in the full field \dot{B} and the $H^1(\text{Grad})$ test functions have been expanded to illustrate the equivalent scalar weak form of the integral. This is not a true constraint of course, as the system is already fully defined by the induction equation, but instead indicates that the induction equation itself analytically satisfies $\nabla \cdot \dot{B} = 0$ in a weak sense, when the H^0 basis functions are used as test functions. This constraint is equivalent, for finite element methods, to the constraint seen in divergence preserving[79, 80] finite difference and finite volume schemes, where the divergence is preserved only when evaluated with a suitable stencil[81].

$$\begin{aligned} 0 &= \int_{\Omega} \left(v_g^T \cdot \dot{B} \right) dV = \int_{\Omega} \left(\nabla \phi^T \cdot \dot{B} \right) dV \\ &= - \int_{\Omega} \left(\phi^T \nabla \cdot \dot{B} \right) dV + \int_{\partial\Omega} \left(\phi^T \dot{B} \right) \cdot dS \end{aligned} \tag{5.3}$$

This property results in a different character of divergence error when compared to C0 finite element methods like NIMROD. It also raises a related question about whether the divergence constraint must be satisfied in its strong form or whether satisfying an appropriate weak form can be acceptable. This question would be interesting to explore in greater detail, however for the purposes of this work only investigation sufficient to justify use of this method for simulating HIT-SI was performed.

Time stepping

The widely used Crank-Nicolson time-centered implicit method is used to advance the system of equations in time. Representing the time advance as in equation 5.4, the Crank-Nicolson method takes the form of equation 5.5, where the non-linear function F is evaluated with arguments from the current (n) and next ($n + 1$) time step. This method provides second order accuracy in time, while being unconditionally stable due to its implicit nature. As this is an implicit method the non-linear system must be solved to compute the solution at the next time step. In PSI-TET this is done using a Newton-Raphson method to solve the non-linear equations with a Krylov space iterative method, usually GMRES[82], used to provide an approximate inverse of the system Jacobian. An approximation to the full Jacobian is explicitly formed as a matrix and updated as necessary. A more in depth description of the linear solver is given in section 5.3, while weak form definitions for the non-linear function and Jacobian matrices used in the time advance are provided in appendix D.

$$\frac{\partial u}{\partial t} = F \left(u, x, t, \frac{\partial u}{\partial x}, \dots \right) \quad (5.4)$$

$$\frac{u^{n+1} - u^n}{\Delta t} = \frac{1}{2} \left[F^{n+1} \left(u, x, t, \frac{\partial u}{\partial x}, \dots \right) + F^n \left(u, x, t, \frac{\partial u}{\partial x}, \dots \right) \right] \quad (5.5)$$

5.2 Verification and Benchmarking

The test cases presented in this section were used to verify the reduced MHD implementation and ensure that it reproduced known analytic solutions for simple problems. The Alfvén-whistler dispersion relation for parallel wave propagation on a uniform background in slab geometry as well as the growth rate for a tilt unstable spheromak in a cylindrical vessel are verified to match known solutions. Other test cases have also been performed during development to verify resistive decay rates, momentum advection, and other simple behavior, however these results will not be

presented here. Agreement for these tests provides evidence that the reduced MHD implementation in PSI-TET is effectively reproducing the desired physical system.

5.2.1 *Alfven-Whistler dispersion relation*

The reduced MHD model only supports the propagation of a single distinct wave, the Alfven wave, through interaction of the magnetic field and plasma fluid. When the Hall term is added a new wave[83], the whistler wave, is introduced that is significantly more challenging from a numerical perspective due to its dispersive nature. For parallel propagation, $\hat{k} \cdot \hat{b} = 1$, these two waves are captured by the dispersion relation[84, 85] given in equation 5.6, where the wavenumber (k) is normalized to the ion inertial length $d_i = c/\omega_{p,i}$ and the frequency (ω) is normalized to the ion cyclotron frequency $\Omega_i = eB/m_i$. Note that the presence of the Hall term causes the previously degenerate roots of reduced MHD to separate at high k . The frequency of the shear wave, sometimes called the ion cyclotron wave in Hall MHD, is now limited as it approaches Ω_i . The frequency of the compressional wave on the other hand strongly increases as $\sim k^2$ for high wavenumbers, where it is usually referred to as the whistler wave.

$$\omega = \pm k \sqrt{1 + k^2/4} \pm k^2/2 \quad (5.6)$$

To verify the basic operation of the reduced Hall MHD model in PSI-TET, simulations were performed to determine the dispersion relation numerically and compared to the expected result, figure 5.1. Simulations were performed using a fully ionized deuterium plasma with a density of $N = 1 \times 10^{16} \text{ m}^{-3}$, equilibrium magnetic field of $B_z = .1 \text{ T}$, and low dissipation $\eta = \nu = 10^{-3}$ in a box domain with extent of $[-d_i, d_i]$ in the x and y directions and $[-2\pi d_i, 2\pi d_i]$ in the z direction. A small amplitude perturbation ($\delta B_{max} = 10^{-1} \text{ mT}$) with right hand circular polarization, equation 5.7, was then applied to excite only a single wave number. Although only a single

wavenumber is excited by this perturbation, all wave frequencies of the above dispersion relation are excited. All 4 propagation frequencies are then determined by performing a Fourier analysis of the evolution of B_x at the origin. This method was chosen for simplicity as it did not require initializing eigenmodes of the system, which are complicated by the presence of multiple wave solutions. The resulting frequencies for propagation in the positive direction are shown in figure 5.1, compared to the analytic predictions from equation 5.6 for a range of wavenumbers. Wave dispersion is found to agree well with the expected result, with the largest errors resulting from the FFT frequency analysis.

$$\delta\vec{B} = \sin(kz)\hat{x} + \cos(kz)\hat{y} \quad (5.7)$$

It is worth noting that the results from this test are only meant to verify the evolution of the model equations for a simple linear test case. As such a time step was used, $\Delta t\omega \approx .05$ for the fastest wave frequency, that ensured numerical dispersion introduced by the implicit time advance was small. Similarly, a spatial resolution, 40 elements with polynomial degree of two in the z direction, was used that would capture the shape of the perturbation well. In full simulations the effects of coarse resolution and the time advance will introduce numerical errors for high frequency waves, which have not been addressed here. As this is a novel discretization method it would be worth investigating the detailed behavior of numerical error sources in the future.

5.2.2 Tilt Unstable Spheromak

The second test case provides a more complete test of the implementation by involving a large spectrum of physical waves and modes instead of the carefully chosen modes used in the Alfvén-whistler test. A cylindrical domain is used with major radius of 1 m and a height of 2 m. The initial condition for the magnetic field is set to an




Figure 5.1: Comparison of propagation frequencies for waves in reduced Hall MHD from numerical tests in PSI-TET (squares) with the analytic dispersion function (solid).

axisymmetric spheromak equilibrium. In this geometry the initial condition is ideally unstable to a tilt mode and when perturbed the unstable mode will grow linearly followed by a non-linear saturation and transition of the magnetic configuration to a lower energy non-axisymmetric state. The growth rate for this mode was computed by Bondeson *et al.*[3], as shown in figure 5.2. For an aspect ratio of $L/R = 2$ the growth rate is $\gamma = .1V_a/R$, where V_a is the Alfvén velocity evaluated at the center of the cylinder.

A set of simulations were run with varying grid resolution in order to verify the correct linear growth rate as well as the transition to a non-axisymmetric state. The simulations used deuterium as the ion species with the physical parameters shown in table 5.2.2. Both the equilibrium and the initial perturbation, chosen to produce a net current encircling the y -axis, were computed using the Taylor module in PSI-TET. For this geometry the axisymmetric spheromak corresponds to the 3rd eigenstate and the lowest energy eigenstate was used as a small perturbation to the magnetic




Figure 5.2: Ideal MHD growth rate of the unstable tilting mode in a cylinder as a function of aspect ratio as computer by Bondeson *et al.*[3].

field. The velocity was set to zero at the start of simulations. A perfectly conducting boundary condition was used for the magnetic field, $E \times \hat{n} = \dot{B} \cdot \hat{n} = 0$, while a Dirichlet boundary condition on all components of the velocity was used, $v = 0$. Hall effects were not included for these simulations.

The growth of the unstable mode is shown in figure 5.3 for each resolution and polynomial order compared to the analytic growth rate. Two different grid resolutions are used, $N_g = 1$ with a mean edge length of 10 cm and $N_g = 2$ that is produced by a single refinement with a mean edge length of 5 cm. Shown are toroidal flux and current computed with the symmetry axis oriented along the y -axis to measure the bulk field in the growing mode. Rates converge to the theoretical prediction with increasing resolution, with all simulations except the coarsest within 10%. Similarly, the final toroidal current agrees for all runs. Figure 5.4 shows the evolution of the total system energy, $E_{tot} = \frac{1}{2} \int \left[\frac{1}{\mu_0} \vec{B} \cdot \vec{B} + \rho \vec{v} \cdot \vec{v} \right] dV$, in time. Low dissipation results in a modest decay rate before and after the period of non-linear saturation where large scale dynamics and reconnection are present raising the dissipation through excitation

Table 5.1: Run parameters and settings for cylinder tilt benchmark

Parameter	Simulation
B_0 [mT]	91.6
$E_{\delta B}/E_{B_0}$	10^{-4}
n_e [m^{-3}]	1×10^{19}
η/μ_0 [m^2/s]	1
ν/ρ [m^2/s]	20

of short wavelength modes. These results provide further verification that the reduced resistive MHD implementation in PSI-TET is modeling the physical systems correctly.

Behavior of Divergence Error

During the non-linear saturation phase of the tilt a large spectrum of mode wavelengths are driven. This provides a good test case for analyzing divergence behavior with the Nedeléc discretization in a dynamic simulation. Before considering the time dependent behavior, divergence error at a single time point is assessed by evaluating $h\nabla \cdot B$, where h is the node spacing, on two parallel chords along the x -axis 1 mm above and 1 mm below the mid-plane at 0.19ms in the ($N_g = 2$, $np = 2$) simulation, shown in figure 5.7. These values are compared to the magnetic flux density ($|B|$), evaluated at the same time midway between the two chords. The strong form divergence error exhibits a high frequency variation, at the grid size, with minimal low frequency content. This indicates that the resulting strong form divergence field is produced by a small fraction of the magnetic field with high wavenumber variation in the parallel direction, $|\nabla \cdot B| \sim k\delta B$, where $k \approx 1/h$ is the wavenumber and δB is the fraction of the field with non-zero divergence. Therefore, the metric $h\nabla \cdot B$ is a good measure of the field amplitude supporting strong form divergence error in the

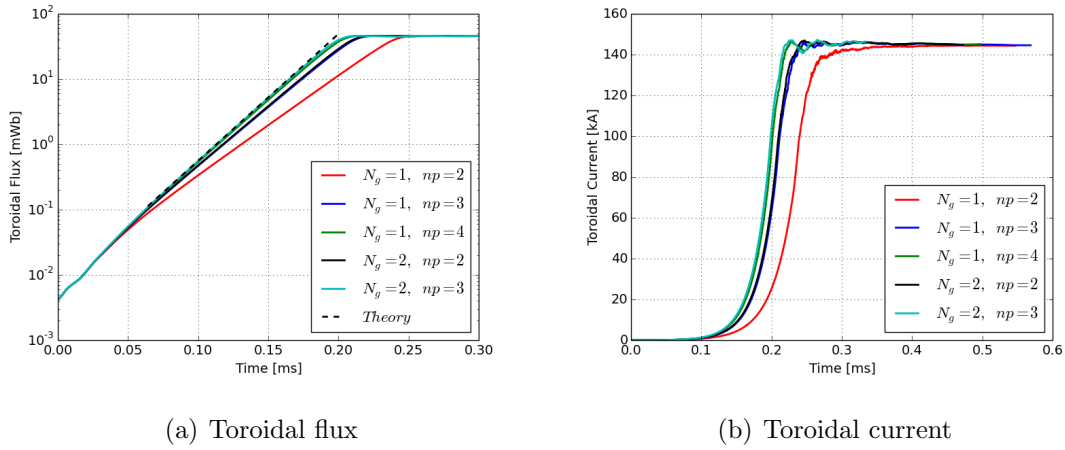


Figure 5.3: Convergence of toroidal flux (a) and current (b) for the unstable cylindrical tilt benchmark at different grid resolutions and polynomial orders.




Figure 5.4: Convergence of total system energy for the unstable cylindrical tilt benchmark at different grid resolutions and polynomial orders.




Figure 5.5: Point divergence error (blue, green) compared to magnetic field strength (red) along parallel chords at 0.19ms in the ($N_g = 2$, $np = 2$) simulation from the unstable cylindrical tilt benchmark.

solution. This picture is consistent with the divergence constraint in weak form, as error at the grid scale will average to zero over sub volumes for each basis function.

In order to investigate the behavior of divergence error throughout time in this benchmark case, two metrics are defined. The first metric, equation 5.8 where h is again the node spacing, provides a measure of the relative energy in the fraction of the field that supports a divergence. From the previous analysis of the strong form divergence error in space the dominant wavenumber is found to be consistent with the finite element node spacing, making this scale a suitable wavenumber normalization. This metric is similar to the divergence error metric used by NIMROD, however it is important to note that in NIMROD this metric is directly related to added dissipation introduced by the divergence cleaning term, which must be kept small. In PSI-TET this metric only has meaning as an approximation of global strong form divergence error.

$$\mathcal{E}_E = \frac{\int (h\nabla \cdot B)^2 dV}{\int (B \cdot B) dV} \quad (5.8)$$

Although the field energy is generally a suitable measure, in the H^1 discretization strong form divergence should be produced by jumps in the field at cell boundaries. As the H^1 representation does not enforce continuity in the normal component of the field, jumps are possible and in general present at cell interfaces. These jumps act as the source of divergence error in the solution, which must be zeroed in a weak sense by the $H^1(\text{Grad})$ subspace. To provide a measure of the field fraction supporting divergence error due to discontinuities, a second metric, equation 5.9, is defined which approximates the RMS jump error between grid cells. The surface integrals are evaluated over all internal faces in the mesh, where $[[B \cdot \hat{n}]]$ is the jump in the magnetic field normal to the face.

$$\mathcal{E}_J = \sqrt{\frac{\int ([[B \cdot \hat{n}]])^2 dS}{\int (B \cdot B) dS}} \quad (5.9)$$

Figure 5.6 shows evolution of the energy metric, equation 5.8, and the square of the jump metric, equation 5.9, over the course of a single simulation, ($N_g = 2$, $np = 2$), along with the kinetic energy. Divergence energy starts out low and grows slowly over the linear portion of the unstable mode growth. An increase in the growth rate of both metrics is seen as the mode reaches the non-linear phase, which can be located by the peak in the kinetic energy. Following saturation, the divergence error decreases back toward the initial level as high wavenumber activity dissipates due to resistive and viscous losses. Note that divergence error stays bounded in time, which is in contrast to C^0 finite element methods where the divergence error will continue to grow in an uncontrolled manner without cleaning. The time behavior is consistent with divergence error resulting from short wavelength mode activity as the peak error is associated with the peak energy dissipation due to high wavenumber activity. The ratio of the two curves also remains relatively fixed throughout most of the shot, indicating that the terms are related solely by a factor. This fixed relationship is further indication that the picture of strong form divergence error at a constant, high




Figure 5.6: Comparison of divergence error with kinetic energy in the unstable cylindrical tilt benchmark.

wavenumber, such as that supported by the grid spacing, is correct.

Finally, the two metrics are compared for each of the different grid resolutions and polynomial orders in figure 5.7. As both metrics measure an error they should decrease, converging toward zero, as the representation is improved. Although this effect is seen the behavior is more complicated, as the two metrics assess the source of and resulting error separately. The jump error metric, equation 5.9 and figure 5.7(a), is seen to decrease monotonically with both increased polynomial degree or grid resolution. The energy metric, equation 5.8 and figure 5.7(b), increases slightly with higher polynomial order, but decreases when the grid is refined. These trends also agree with the picture of divergence error in the H^1 expansion, as the source of the error, discontinuities at cell interfaces, cannot be smoothly zeroed by a polynomial function. This fact would tend to introduce Gibbs phenomena, which may maintain the energy of the error field while locally reducing its magnitude.

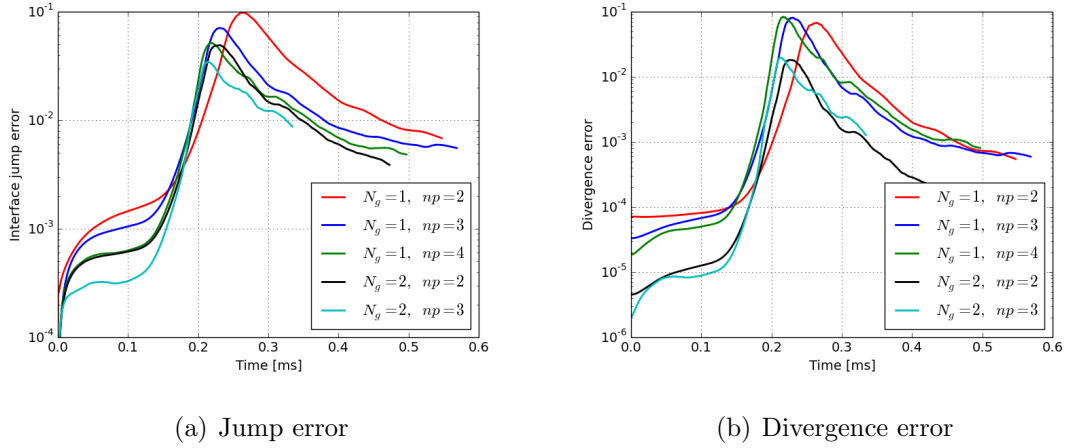


Figure 5.7: Divergence error computed with local jump (a) and global energy (b) metrics for the unstable cylindrical tilt benchmark at different grid resolutions and polynomial orders.

5.3 Multi-Grid Solvers

As with the other solvers and methods discussed to this point the 3D MHD implementation utilizes multi-grid preconditioners to accelerate the inversion of the Jacobian matrix. Magneto-Hydrodynamics is particularly challenging to model due to the large span of time and spatial scales present in the system. As a result the time scale of interest is often much longer, $\sim 10^3 - 10^5$, than the fastest wave speeds. PSI-TET uses an implicit solver to eliminate the numerical CFL condition on the time step, however new limits are introduced by this choice due to stiffness of the linear and non-linear systems that must now be solved. Stiffness in the non-linear iteration is often not a limiting factor due to time step limits imposed to maintain solution accuracy and capture desired short time behavior. The linear system on the other hand can become very difficult to solve, even with a modest time step, due to large variation in wave speeds or other stiffness caused by dynamics. Additionally, as the system size grows, due to increasing resolution or physical extent, the condition number of the Jacobian

matrix increases, providing poor scalability.

The multi-grid preconditioner for reduced Hall MHD is constructed using the same mechanisms used to support single field multi-grid solves as presented in section 2.2.5. The required finite element spaces are constructed on a set of hierarchical representation levels and then used to construct approximations to the system Jacobian matrix on each level. For MHD this requires current solution fields to evaluate the linearized terms in the Jacobian. Currently the full basis representation is used for levels sharing the fine grid and a volume average of solution fields in each cell is used for levels on coarser grids. Different smoother configurations can be easily defined and tested through an XML definition of their structure. The currently used configuration uses a global direct solve on the coarsest level, provided through PETSc, and GMRES on each finer level that is further preconditioned by a block-Jacobi method using a direct solve for each block. Typically, a single block per processor is used to define the block-Jacobi decomposition.

5.3.1 Comments on Scalability

The preconditioner described above has shown good performance with increasing system size on the simulations performed so far. Due to the use of a block-Jacobi method as part of the smoother iteration the solution cost should remain fixed as long as the ratio of degrees of freedom to processors is held constant. Some increase in factorization cost does still exist however as the polynomial order is increased due to decreasing sparsity of the matrix. Required iteration count for convergence scales only very weakly with the system size with this method however. This effect is shown in figure 5.8 where the time advanced through the solution is shown as time steps are applied for the the HIT-SI validation case with different mesh resolutions and polynomial orders. For these simulations the time step was varied to maintain a fixed solution cost of 60 linear iterations per time step. The time simulated is reduced by only $\sim 5\%$ with an increase in system size by a factor of 20 from the smallest to




Figure 5.8: Comparison of solution cost for the HIT-SI validation case with different mesh resolutions and polynomial orders.

largest solve. Solution time for these cases increased slightly due to decreasing matrix sparsity with polynomial order and a practical limit on the number of processor that could be used set by queue wait times, 96 for the $N_g = 0$, $np = 2$ case and 768 for the $N_g = 1$, $np = 3$ case. Although this method does not provide the scalability multi-grid afforded the parabolic test cases for individual fields it has been found to perform significantly better than tested single level preconditioners alone.

5.3.2 Interpolation and Restriction

Interpolation and restriction operators are constructed by merging operators for each of the sub fields being advanced. These operators are constructed for each of the individual finite element representations and representation levels used as described in sections 2.2.3 and 2.2.4. As the full H^1 vector space is being used additional care must also be taken to ensure proper handling of the overlapping gradient subspaces in the Nedelec representation. This is handled in PSI-TET by mapping the 0-th order gradient space from the vertex $H^1(\text{Grad})$ degrees of freedom to the edge $H^1(\text{Curl})$ basis weights following application of the interpolation operator.

5.4 Modeling The HIT-SI Wall

The insulating coating on the interior surface of HIT-SI imposes a constraint that no current may flow into the wall. This not only ensures fully inductive injection of magnetic helicity but also allows the magnetic flux and power to redistribute itself over the boundary. It is expected that the second effect has a significant impact on the formation and current drive in HIT-SI.

$$J \cdot \hat{n} = 0 \tag{5.10}$$

In previous simulations with the NIMROD code this effect is approximated by a thin boundary layer with high resistivity relative to the internal value. To prevent numerical issues associated with this abrupt variation in resistivity it is necessary to introduce a matching boundary layer in the mesh so that the variation can occur within a single cell. This process is relatively straightforward for the 2D structured poloidal mesh used by previous simulations. On a tetrahedral mesh in three dimensions however this is much more complicated and is impractical with the complex geometry where the injectors meet the confinement volume in HIT-SI. It would require a significant amount of time and adjustment by hand to generate such a mesh without adversely affecting the mesh quality. In order to avoid this requirement a new boundary condition has been developed to restrict J_{norm} at the wall by leveraging the vector subspace decomposition provided by Nedelec elements.

5.4.1 Numerical Formulation with Nedelec Elements

Analyzing the physical constrain given by equation 5.10 a constraint on the variation of the magnetic field at the boundary can be extracted, equation 5.11. The requirement that the tangential magnetic field on the boundary must be curl free can be straightforwardly enforced using the Nedelec H^1 representation.

$$\hat{n} \cdot \nabla \times B = \nabla_t \times B_t = 0 \rightarrow B_t = \nabla_t \chi \quad (5.11)$$

A Dirichlet boundary condition can be applied to the curl subspace of the tangential field by removing the $H^1(\text{Curl})$ basis functions that are tied to boundary vertices, edges, and faces. This restricts the basis set, and resulting solution space, to functions that do not support normal current density at the boundary. Variation of the tangential magnetic field is then supported by the $H^1(\text{Grad})$ subspace on the boundary, which by definition satisfies the restricted field defined in 5.11.

As mentioned in section 2.2.4 there is an overlap in the representation of low order gradient fields between the $H^1(\text{Curl})$ and $H^1(\text{Grad})$ basis functions. For the insulated conductor boundary condition this requires that special care be taken for fields with non-zero loop integrals, $\oint \vec{B} \cdot d\vec{l} \neq 0$. Since only gradients of single valued scalar functions can be represented by the $H^1(\text{Grad})$ basis set, loop integrals around closed loops are forced to zero by zeroing all boundary $H^1(\text{Curl})$ weights. To support such fields a constrained subset of the $H^1(\text{Curl})$ basis must be included. A single weight must be added for every closed loop on the boundary that can enclose a unique current. This weight corresponds to a composite basis function that is defined by forcing a string of $H^1(\text{Curl})$ linear edge elements lying on a closed strips of boundary triangles to share a common weight. An example of such a closed triangle strip is shown in figure 5.9 for the flux injection boundary condition in HIT-SI. Although in general these weights may be included as a degree of freedom in the problem all investigations so far have constrained these weights as a part of the boundary condition.

Mesh Construction

Boundary triangle strips can follow any arbitrary closed path, however for convenience cut planes have been used when defining the geometry for the HIT-SI simulations.




Figure 5.9: Triangle strip used to define the composite basis function for the flux circuit boundary condition on a single injector for HIT-SI simulations.

Meshes are created in CUBIT and T3D by hierarchically discretizing CAD vertices, curves, and surfaces. As a result a planar curve in the CAD representation can be used to produce a suitable strip that can be easily identified in the code. Figure 5.10 shows an example of the boundary curves used to build planar triangle strips for the flux and current boundary conditions on a single HIT-SI3 injector.

5.4.2 Flux Injection and Current Drive

By imposing $\oint \vec{B} \cdot d\vec{l} = f(t)$ on independent sets of closed loops on each injector the flux and voltage circuits in HIT-SI can be approximated. This boundary condition amounts to fixing the current enclosed within these loops as a function of time. The effect of the flux circuit is modeled by controlling the loop integral along the locally toroidal direction, fixing the current through the injector handle. This relates to the total coil current in all the flux coils on each injector on the experiment. The effect of the voltage circuit is modeled by controlling the loop integral along the locally poloidal direction, fixing the plasma current through the injector itself. This relates




Figure 5.10: Slice planes on a single HIT-SI3 injector used to create planar curves and resulting triangle strips for simulating the flux and voltage drive circuits.

to imposing the measured injector current in the experiment.

Both the flux and voltage boundary conditions, although approximating the experimental drive, have important differences for the physical device. Both the flux and voltage circuits are tied to resonant tank circuits that act to amplify as well as shape the voltage and current signals from the switching power supplies. This system drives each coil set with a low source impedance resulting in current waveform that may be strongly dependent on the plasma response. The boundary condition used in PSI-TET simulations, which imposes coils currents directly, on the other hand acts with a high impedance. Additionally, the inductance of the flux coil on the real machine and in the simulations, which can be inferred from the relation between the injector flux and flux coil currents, are different. The effects from these differences are discussed in more detail in the HIT-SI physics section 5.6.




Figure 5.11: Magnetic field early in time in a HIT-SI simulation due to flux injection and current drive boundary conditions on the injector. In-plane field is shown as black arrows and a pseudocolor plot is shaded by the out of plane field strength.

5.4.3 Sheath Model

In separate investigations a second boundary condition to simulate the insulated wall in HIT-SI, and more generally any wall with a capacitive sheath, was developed. In order to model a capacitive wall Ampere's law, equation 5.12a, is extended to allow displacement currents in the direction normal to the boundary, equation 5.12b. Where \hat{n} is the boundary unit normal, ϕ is an electrostatic potential on the boundary, and $\frac{\epsilon}{l}$ is the capacitance per unit area of the insulator. This method introduces an additional scalar variable ϕ which must be evolved and provides the tangential electric field caused by charging the insulator on the boundary, $E_t = \nabla\phi$.

$$\nabla \times B = \mu_0 J + \mu_0 \epsilon_0 \dot{E} \quad (5.12a)$$

$$\hat{n} \cdot \nabla \times B = \mu_0 \epsilon \hat{n} \cdot \dot{E} = \mu_0 \frac{\epsilon}{l} \dot{\phi} \quad (5.12b)$$

This method was used in tests with an earlier version of the PSI-TET code that did not support the full element space required for the boundary condition presented above. In that iteration a simplified electron MHD model was used. Although the boundary condition performed well in these tests, the electron MHD proved to be too restrictive to allow the current amplification observed in HIT-SI. This boundary condition has not been tested on the newer physics models, however one advantage to this method is that it would allow capturing the effects of more general plasma sheaths where the current flowing to the wall has a known relationship to the relative voltage. A disadvantage to this method would be the possibility to introduce new, potentially poorly conditioned boundary physics, which may slow down the solution.

5.5 *HIT-SI Validation*

The capabilities presented in this chapter have been used to model the full HIT-SI experimental geometry. This section presents validation of the PSI-TET reduced Hall-MHD model as well as the boundary condition used to simulate the HIT-SI wall. Of particular interest is how this new implementation compares to previous NIMROD studies, which have been used to validate the reduced Hall-MHD model as a good approximation to HIT-SI physics. Recent observations in the experiment as well as numerical studies with NIMROD at high injector drive frequency indicate that the reduced MHD model is not appropriate for all operational ranges. This issue will be discussed more in section 5.7 along with preliminary development of a full MHD model in PSI-TET. However, for the purposes of model validation the behavior of HIT-SI at 14.5 kHz injector frequency has been found to be well captured by reduced MHD.

The validation case was performed with free physical parameters set to closely match experimentally observed values for the baseline comparison shot, shot 122385, shown in table 5.5. This shot was chosen as it has been extensively studied due to its high performance and similarity to other high performance operations at 14.5 kHz.

This shot does not constitute the highest performance achieved in HIT-SI, but is representative of operation in this regime. The flux and voltage waveforms were driven as $\sin(2\pi f_{inj}t)$ and $\sin(2\pi f_{inj}t - \pi/2)$ for the X and Y injectors respectively with a linear ramp in amplitude over the first injector period ($69\mu s$) to the steady state values shown in table 5.5. In Dr. Akcay’s thesis[30] significant attention was paid to the sensitivity of the solution to variation in the physical parameters in table 5.5. Full sensitivity studies have yet to be repeated with PSI-TET, however variation with respect to resistivity has been assessed as presented in section 5.6.2. The validation case was conducted using the “hitsi_rs” T3D mesh (appendix A.1.1) with an average edge length of ~ 4.4 cm for the base mesh. A single mesh refinement was used with a polynomial degree of two for the finite element discretization, leading to a solution mesh with an average node separation of ~ 1 cm. Slightly higher resolution exists in the injectors and near the nose cone due to small geometry features, while slightly lower resolution is presented on the outboard mid plane where low radius of curvature is present. Convergence studies performed to justify this choice as well as discussion are presented in section 5.5.1. A maximum time step of 60 ns was used that is adapted as the simulation advances to maintain a constant number of linear iterations per time step. A multi-grid preconditioner, as discussed in section 5.3, was also used to accelerate solution.

A comparison between the toroidal current in shot 122385 and the simulated toroidal current for the validation case is shown in figure 5.12. Good agreement is seen between the saturation value as well as the time required to reach maximum toroidal current. Variation between the fluctuations on this plot is primarily due to differences in the way each current is calculated. In the experiment the surface probes are used, as discussed in section 1.2.3, while in the simulations a toroidal average is used over the entire volume. For a more direct comparison of oscillations the first chrono of the Bi-Orthogonal Decomposition (BD)[2] analysis using reference topos can be used, figure 5.25. The time bases are synchronized by comparing the model and observed injector

Table 5.2: Run parameters and settings for HIT-SI validation case

Parameter	Experiment (Shot 122385)	Simulation
f_{inj} [kHz]	14.5	14.5
Ψ_0 [mWb]	0.6-1.4	1.2
I_0 [kA]	10-20	21
λ_{inj} [m ⁻¹]	15-25	21
n_e [m ⁻³]	$1 - 4 \times 10^{19}$	1.5×10^{19}
T_e [eV]	6-12	8
η/μ_0 [m ² /s]	25-9	16
ν/ρ [m ² /s]	~ 260	200

current waveforms late in time, as shown in figure 5.13. The simulation is first shifted to align the initial formation of significant toroidal current and then further adjusted to provide the best agreement in the phase between the simulation and experiment with the minimum time change required. For the validation case presented here the simulation was a shifted by 641.5 μs in time to provide the alignments shown in figures 5.12 and 5.13 and used for subsequent comparisons. This shift produces good agreement with both injector waveforms, which are not exactly 90 degrees out of phase as they are in the simulation.

5.5.1 Resolution Convergence Study

In order to have confidence in our simulated results convergence of the current drive behavior was checked with different grid resolutions and polynomial representations. For the parameters used the ion inertial length corresponds to $d_i = c/\omega_{p,i} \approx 8$ cm, the resistive skin depth is $d_\eta = \sqrt{\eta/\mu_0\omega_{inj}} \approx 1.8$ cm, and the electron inertial length is $d_e = c/\omega_{p,e} \approx 1.7$ cm with the enhanced electron mass. This places the average





Figure 5.12: Comparison of I_{Tor} from shot 122385 and PSI-TET validation case at 14.5 kHz.



(a) X-Injector Current

(b) Y-Injector Current

Figure 5.13: Comparison of injector current waveforms from shot 122385 and PSI-TET validation case used to align time bases for probe comparisons.

nodal spacing in the reference simulation on the order of relevant spatial scales in the system. Although this scale is under resolved, activity near the electron inertial and resistive scales like reconnection may not have a strong impact on the evolution of bulk quantities such as the toroidal current. This effect was seen in convergence studies using NIMROD, where the resolution could be coarsened dramatically without affecting the $n=0$ current evolution.

In order to assess the sensitivity of the results to the representation resolution a set of simulations were run with varying mesh size and polynomial orders for the physical parameters used in the validation case, table 5.5. The results of this study are presented in figures 5.14 and 5.15, where the baseline simulation used for experimental comparison corresponds to $N_g = 1$, $np = 2$. Variation in the toroidal current and flux with grid resolution is limited to less than 10% between representations with the exception of the coarsest simulation performed. A larger difference is seen on the toroidal flux than on the toroidal current signals indicating that the mean field current profile may account for the majority of the discrepancy. Oscillations on these signals show much tighter agreement in timing and amplitude between different grids and representations. These results indicate that the qualitative behavior of the simulations is well resolved at the grid resolution used, however the detailed evolution on intermediate time scales may vary. The largest differences seen on the flux and current waveforms appear to be transient as the growth rate varies on a time scale slower than the injector frequency, as seen in figure 5.12. This modulation effect is thought to be due to detailed evolution of the mean field equilibrium through stability space, resulting in temporary drive of unstable modes that may not need to be resolved to a high level. The timing of these variations is also seen to vary in the experimental measurements of toroidal current, leading to a larger spread than observed here. It would be desirable to extend this convergence study for a longer period of time as well as to include additional grid resolutions, however time and resource constraints limited the scope for this study.




Figure 5.14: Convergence of toroidal flux evolution in HIT-SI validation case with different mesh resolutions and polynomial orders.




Figure 5.15: Convergence of toroidal current evolution in HIT-SI validation case with different mesh resolutions and polynomial orders.

5.5.2 Internal Magnetic Structure

The internal magnetic probe on HIT-SI[14] provides information about the magnetic structure along a radial path a fixed toroidal position. Evolution of the poloidal and toroidal components of the magnetic field between shot 122385 of the experiment and the PSI-TET validation case are compared in figure 5.16. In order to prevent clutter in the plot only a subset of probe locations are shown. The signals for each probe are also plotted with a fixed offset between probes of 50 mT and 80 mT for the poloidal and toroidal fields respectively. The locations shown correspond to major radii in cm of 34.4, 38.2, 42.0, 45.8, and 49.6, ordered from bottom to top in the figure.

Overall, very good agreement is seen between PSI-TET and the experiment in both the mean amplitude as well as the fluctuations from the injectors. Phase of injector driven oscillations agree well over the entire shot, with agreement in amplitude when experimental flux and current match the simulated values $\sim 1.5\text{ ms}$. The spheromak formation phase exhibits significantly different fluctuation amplitudes despite reasonably good agreement in the time of formation relative to breakdown in the experiment and the start of the simulation in PSI-TET. This is likely due to differences in the injector ramp up in PSI-TET compared to the experiment. Following breakdown in the experiment a significant drop in the flux and current in the injectors is observed that may cause lower free energy before relaxation, resulting in decreased fluctuation amplitude. Additionally, the steady state injector fields in PSI-TET simulations are ramped up much faster than the experiment where the injector power is seen to increase over a long portion of the toroidal current ramp time. Similar oscillations are also present in NIMROD simulations, which use a PSI-TET like ramp up time for the injectors.

Throughout the shot some differences in fluctuation amplitude and character are also seen in the poloidal field particularly near the mean field magnetic axis, which is defined as the zero crossing of the time averaged poloidal field. However, PSI-

TET simulations show much better agreement, than NIMROD, with the experiment in the behavior of poloidal fluctuations field in this region. The broad frequency spectrum and random events present in experimental data are reproduced by PSI-TET, which is in contrast to the steady, fixed frequency fluctuations seen in this region for NIMROD simulations. Very good agreement is even seen in the detailed behavior of some transient events, such as poloidal fluctuations on the two lowest probes signals in figure 5.16(a) just after 1.5 ms. PSI-TET appears to see more dynamic activity than experimental data throughout the shot, however agreement in transient events indicates that the physical nature of these fluctuations is largely captured. Experimental data also exhibits a broader frequency spectrum and lower fluctuation amplitude than PSI-TET simulations on the outer poloidal field probes, indicating less direct coupling of fluctuations to the injector frequency.

Finally, a difference exists between the steady state values of both toroidal and poloidal field when the toroidal currents are comparable indicating a difference in the mean field equilibrium in the simulation. This effect could be due to a different toroidal mode spectrum, current profile, the effect of non-zero pressure in the experiment, or some combination of the three. Surface probe data indicates that the mean field is relatively symmetric, making it unlikely that differences in locked mode amplitudes would cause this effect. Current profile and pressure effects are suspected to cause the discrepancy, motivating the development of full Hall MHD models with plasma pressure as discussed in section 5.7. Axisymmetric fits to experimental data, figure 1.13, indicate the existence of pressure at higher injector frequencies. Although it is not seen in fits pressure may still be present at lower injector frequencies at a sufficient level to impact the mean field profile, but low enough to not be distinguishable during fitting. Additionally, it is thought that pressure may have a role in current drive even at low levels, as the Mercier β is HIT-SI in low, leading to current profile differences. A similar deviation in the time averaged field is seen in NIMROD simulations with reduced MHD, further pointing to pressure effects as the possible

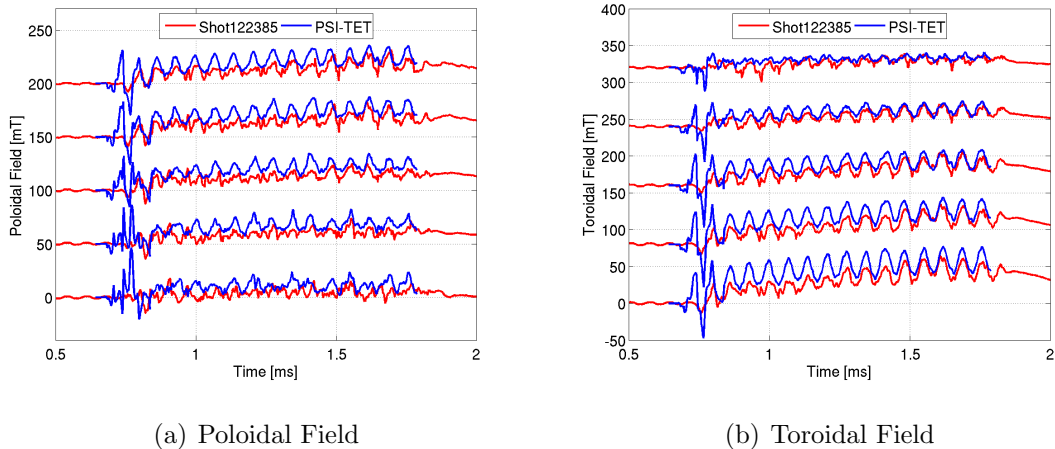


Figure 5.16: Magnetic field comparison along the internal magnetic probe for the HIT-SI validation case.

explanation.

5.5.3 Surface Magnetic Structure (BD)

The surface magnetic probe array on HIT-SI[13] provides a rich set of diagnostic signals for comparison with simulations. For the simulations these diagnostic signals must be extracted from solutions for comparison, requiring the placement of synthetic probes in the models. Synthetic versions of the surface probes are placed in PSI-TET as close as possible to the actual probe locations, while requiring that the synthetic probe stay on the domain boundary. The magnetic field is then reconstructed to these locations using the full basis expansion. As the PSI-TET mesh is based on an accurate CAD model of HIT-SI the resulting locations are very close to the original probes. Synthetic probes in NIMROD are placed poloidally at the closest finite element node to the actual probe location, while requiring that the node be outside the resistive layer. The toroidal position is set exactly by reconstructing the Fourier components of the solution accordingly. As both the PSI-TET and NIMROD grids

do not include the diagnostic gap for computational convenience, the gap probes are excluded from the following analysis. This reduces the working set to 64 probe locations with toroidal and poloidal fields at each location. These probes provide information about the toroidal and poloidal variation of fluctuations and mean field states with high frequency response. However, the number of probes also poses a challenge for comparison due to the large number of individual signals. Standard methods for decomposing or simplifying the data, such as Fourier analysis, are also not well suited to the surface probe array. As the probes are not evenly spaced in the toroidal direction and the poloidal cross-section is not circular, complex exponentials used in Fourier mode analysis provide a poor basis. To address these difficulties a new method, bi-orthogonal decomposition, is being explored on HIT-SI for comparing magnetic probes signals across different data sets[30, 2]. This method has shown promise as a tool for understanding experimental mode activity as well as for defining metrics for validating simulations against experimental data.

Bi-Orthogonal Decomposition for Signal Based Mode Decomposition

Bi-Orthogonal Decomposition (BD)[86] is a method for decomposing experimental signals into principal components for further analysis. In order to use this method the data must first be arranged into a signal matrix $Y_{i,j} = y_i(t_j)$, where individual diagnostic signals(y_i), sampled at discrete time points t_j , are stored as columns. This results in an $N \times M$ matrix where N is the number of diagnostic signals and M is the number of time points. Note that all signals must share a common time base in this representation. The signal matrix is then decomposed using Singular Value Decomposition(SVD), equation 5.13, producing sets of singular values and vectors that correspond to the spatial and temporal dependence of principal modes in the diagnostic signals. Each set is contained in one of the factored matrices, ϕ for the spatial vectors (topos), and Ψ for the temporal vectors (chronos). Mathematically these modes correspond to the strongest mappings between the vector space represented by

the time points t_j and the vector space represented by the diagnostic signals y_i . By convention these matrices are constructed to be unitary matrices, where the individual singular vectors are orthogonal to each other and have unit magnitude. Weights, amplitudes, for individual modes are stored in the diagonal matrix A . As the matrix Y is not square in general, only $K = \text{MIN}(N, M)$ modes exist with both spatial and temporal singular vectors. The remaining vector space in the larger dimension does not produce a mapping and can be neglected during analysis. In addition to the matrix notation a mode based notation will also be used in the following analysis, where the singular vectors and weights for the k -th mode are labeled as $\vec{\phi}_k$, $\vec{\Psi}_k$, and A_k .

$$Y = \phi A \Psi^T \quad (5.13)$$

As BD produces a small set of important modes from a large signal space it can be used to analyze a large number of similar signals simultaneously by picking out dominant structures in the data. This makes BD attractive as a method for comparing data from large diagnostic arrays either between experimental data[87], simulation results[88], or both[2]. BD analysis is presented here comparing mode shapes and correlations between experimental data and NIMROD and PSI-TET simulations. Two different approaches are considered using first independent SVD decomposition of signal matrices from each source and then SVD decomposition of a single reference source, from a characteristic experimental shot, which is then used to decompose and compare against the remaining signal matrices. As mentioned above a consistent time base is required for BD analysis, not only for signals within a single data source, but also between data sources. To ensure consistent analysis simulation data sets were first shifted in time to match the experimental oscillations as discussed above. The NIMROD injector waveforms trail PSI-TET by 90° in phase, resulting in a shift as above of $641.5 \mu\text{s}$ for PSI-TET and $572.5 \mu\text{s}$ for NIMROD. Signals are then resampled for a sampling window of 14 injector periods starting at $740 \mu\text{s}$ with a sampling

frequency of 1 MHz. This ensures phase shifts or sampling differences do not affect the SVD comparisons.

Independent BD Comparisons

The first method of comparison between data sets is the method used by Akcay[30, 2] for comparison of NIMROD and experimental data. Here this analysis is repeated with PSI-TET data as well to compare the level of agreement between the two codes and the experiment. Signal matrices are constructed from surface magnetic probe data for all probe locations and field directions for experimental data (shot 122385), synthetic probe data from a NIMROD validation simulation with 12 eV equivalent resistivity, and synthetic probe data from the PSI-TET validation case. BD analysis is performed on each data set independently resulting in dominant modes that are then compared to assess agreement.

The resulting signal “energy” spectrum, A_k^2 , from each of the decompositions are compared in figure 5.17. For magnetic signals these energies are directly related to the true energy of each mode integrated over the length of the sampling window. Both PSI-TET and NIMROD reproduce the qualitative shape of the experimental energy spectrum, including a knee in the spectrum around the fifth mode. From previous BD analysis of HIT-SI it is known that the first mode corresponds to the spheromak mean field state, while the second and third modes correspond primarily to the directly driven injector fluctuations. This can be seen from the corresponding chrono of the first mode, figure 5.18, which matches the toroidal current and the second and third mode chronos, figures 5.19(a) and 5.19(b), which oscillate at the injector frequency. The tail of the energy spectrum, for high mode number, shows better agreement for PSI-TET simulations than NIMROD. The exact behavior and dynamics of these modes are not as well known as the more dominant modes so it is unclear if this difference is significant however. Some differences do exist between PSI-TET and NIMROD simulations that may explain better agreement in

this region of the energy spectrum. First, as PSI-TET models the internal injector dynamics the forcing of the high- k turbulence spectrum may be more similar to that of the experiment. This would be the case if reconnection and reversal dynamics in the injectors are playing a significant role in driving smaller wavelength fluctuations. A second explanation could be the difference in spatial discretizations between the two codes. As NIMROD uses a spectral representation in the toroidal direction, which for these cases has lower toroidal resolution than the PSI-TET grid, the high- k toroidal portion of the spectrum may be damped to a larger extent than in PSI-TET simulations. Similarly, the anisotropy in resolution between the poloidal and toroidal directions may be affecting the drive and damping of modes with high- k in all directions. The effect of injector dynamics seems a more plausible explanation, although this forcing is difficult to extract from the complete dynamics to assess its effect. Representation differences may be playing a role, which could be explored by comparing BD energies of NIMROD calculations with different toroidal resolutions, however if this is the case it would also indicate that the tail has little effect on the evolution of the mean field equilibrium. NIMROD simulations agree well in this respect and are well converged in toroidal resolution.

A discrepancy exists in the first energy, primarily due to differences in the amplitude and evolution of toroidal current in the simulations when compared to the experiment. This is seen in the time behavior for the first mode, figure 5.18, where the amplitudes for the simulations are larger than the experimental amplitude over the majority of the comparison time. A similar effect is also seen for the injector associated modes, figures 5.19(a) and 5.19(b), where the injector field amplitudes are larger over most of the comparison window due to the ramp time caused by the circuit response in the experiment. As these two modes primarily consist of the injector frequency and are nominally 90° out of phase they are typically associated with a rotating mode tied to the injector drive. However, the amplitude of the injector fields increase significantly over the course of the shot, while only the amplitude of the




Figure 5.17: Signal energies from independent BD analysis of experimental data (shot 122385) and PSI-TET and NIMROD validation cases.

second mode from BD is seen to increase. Growth in a single mode cannot capture growth in a rotating structure like the injectors are expected to produce, indicating that higher modes must play a role in reproducing the driven fluctuations. Similarly, the chronos from the NIMROD and PSI-TET data sets vary over the course of the shot despite constant injector amplitude, indicating a plasma response component beyond a pure injector signal. Overall, agreement is good however for both PSI-TET and NIMROD with the experiment in the time dependence of the second and third BD modes. Both codes also capture the asymmetry present in the signal energy of the second and third mode, however neither code matches the ratio seen in the experiment.

The fourth BD mode is the first mode that is not dominantly associated with a known physics phenomenon in HIT-SI. This mode has a significantly lower energy than the three modes discussed thus far, which may indicate that it plays a more subtle role in dynamics. In the IDCD model the injector fluctuations themselves (2nd






Figure 5.18: Mode 1 chrono from independent BD analysis of experimental data (shot 122385) and PSI-TET and NIMROD validation cases.



(a) Mode 2 Chrono



(b) Mode 3 Chrono

Figure 5.19: Modes 2 and 3 chronos from independent BD analysis of experimental data (shot 122385) and PSI-TET and NIMROD validation cases.

and 3rd modes) are responsible for current drive. As a result small differences between data sets in these modes, due to the lack of sheared flow coupling for example, may result in large differences in lower modes where higher relative energy is required to retain a given absolute amplitude. These lower modes may also be important in capturing dynamics that are essential to, but not directly involved in current drive. Redistribution of flux and current profiles following injector reconnection for example may be separable as a mode structure in BD, but required to produce the conditions for current drive. Figure 5.20 shows the temporal behavior of the fourth mode as primarily oscillatory with a dominant frequency approximately twice the injector frequency and a small offset early in time. This mode also does not appear to be associated with a rotating mode as it does not have an associated frequency pair at similar amplitude like modes 2 and 3. Both PSI-TET and NIMROD decompositions have a corresponding mode with similar frequency behavior, however the NIMROD mode has a phase shift with respect to the experimental and PSI-TET chronos that agree. Figure 5.20(b) shows the 3 chronos over a small time window with the phase offset on the positive portion of the oscillation clearly visible. Although differences exist in the time evolution, the signal energies for this mode agree extremely well across data sets. As this mode is still tied strongly to a multiple of the injector frequency it is likely that it is playing a role in the direct injector fluctuations, perhaps as a plasma response to, or redistribution of, the applied fields. It may also be due to asymmetries in the injector drive that are captured in PSI-TET by modeling the injector volumes, see section 5.6.1. Oscillation at this frequency are predicted by the IDCD model, where rapid reconnection of injector current causes local toroidal current to vary as $\frac{1}{2} \left(|I_{x\text{-inj}}| + |I_{y\text{-inj}}| \right)$ [14] and are observed on the line averaged density signal in the experiment.

The analysis thus far has considered the signal energy and temporal evolution of modes with equal rank in the BD analysis of each data set. There is no requirement from BD alone however that these modes have the same spatial structure even if

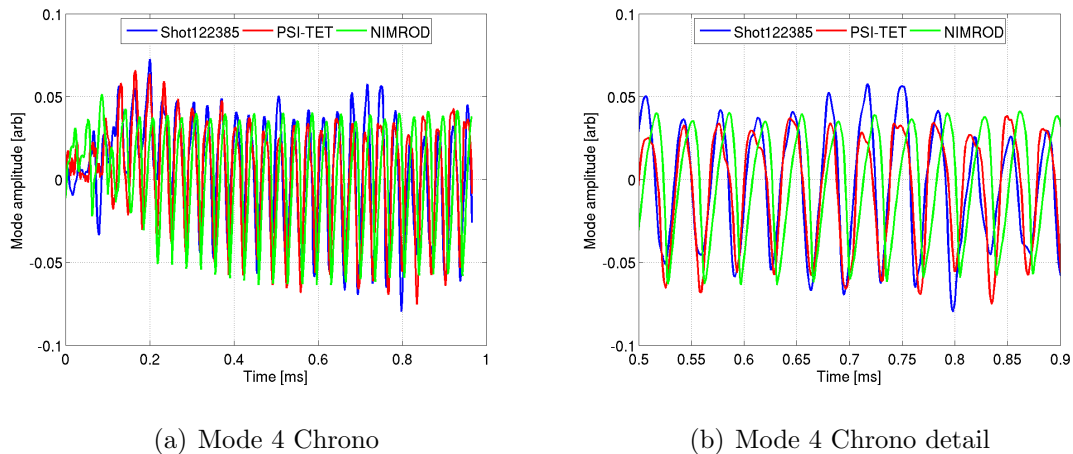


Figure 5.20: Mode 4 chrono from independent BD analysis of experimental data (shot 122385) and PSI-TET and NIMROD validation cases.

their weights and time behavior agree. Figure 5.21 shows the spatial variation of the magnetic field for the fourth mode for both PSI-TET and NIMROD in comparison to the experimental data. The magnetic field structure shown is reconstructing by mapping the topo signal for the fourth mode back to the probe locations and directions used to generate it. The poloidal angle is defined such that the angle increases from zero at the nose cone of the machine toward the Y-injector, placing location “S05” at the lowest poloidal angle and “S04” at the largest poloidal angle, as shown in figure 1.7. The mode exhibits a roughly 1, 2 mode structure, where toroidal variation with $n = 2$ symmetry is seen near the X-injector side at the top of the plot and a more symmetric shape in the toroidal direction is seen near the Y-injector side. This structure shape is resonant with the $q = 1/2$ surface inside the Taylor state for HIT-SI where a large island is present in the equilibrium, figure 3.7. This correlation may indicate that this island structure and its rational surface are being coupled to by the injector drive. Due to energy level of this mode though it is difficult to compare with existing Fourier diagnostics, although some modest $n = 2$ activity is seen in




Figure 5.21: Comparison of mode 4 topo from independent BD analysis of PSI-TET and NIMROD validation cases with experimental data (shot 122385).

the experiment. Agreement is fair between PSI-TET and the experiment with less correlation seen for the NIMROD data, consistent with the temporal comparison. The PSI-TET model is able to capture the toroidal variation seen near the X-injector and some other qualitative features that NIMROD does not. This may also be attributable to the $q = 1/2$ surface coupling as the island at this surface in the Taylor state is introduced by the asymmetry of the injector mouths on the boundary and as a result would likely not be present in the NIMROD simulations without the appropriate driver.

To assess the agreement of topos or chronos between two data sets, Y^1 and Y^2 , in a more quantitative manner a correlation metric is defined, given by equation 5.14. To evaluate correlation between chronos instead of topos the ϕ 's are replaced with Ψ 's in this equation. This metric simply evaluates how well aligned the singular vectors are with respect to each other by taking their dot product. As the individual singular vectors have unit magnitude this correlation will vary in the range $[0, 1]$ where 0 indicates the singular vectors are orthogonal, complete disagreement, and 1 indicated

the singular vectors are identical, complete agreement.

$$C_{i,j}^{1,2} = \phi_i^1 \cdot \phi_j^2 \quad (5.14)$$

This metric is used to analyze the topos of the first 5 modes of the PSI-TET and NIMROD data sets for correlation with the first five modes of the experimental data set in figure 5.22. Correlation is computed for each of the five modes in the simulation data set (colored bars) against the same modes from the experimental reference data set (axis position). The same method is also used to analyze the chronos of the first 5 modes of the PSI-TET and NIMROD data sets for correlation with the first five modes of the experimental data set in figure 5.23. Strong agreement is seen with the experiment for both PSI-TET and NIMROD simulations both spatial and temporal structures for the first 3 modes. Combining this with the previous temporal analysis this confirms that the behavior of the 3 dominant modes is well captured by both codes. Slightly more cross coupling is seen in modes 2 and 3 for the PSI-TET simulations pointing to possibly higher phase mixing relative to the NIMROD results which are slightly more aligned with the experiment. The most significant difference is strong correlations in the 4 mode between PSI-TET and experimental data sets—a feature that is not well captured by NIMROD simulations. For the fifth mode and beyond, not shown here, the correlations become dominated by random distributions in the cross terms indicating that higher modes in this analysis do not agree. This conclusion may be misleading however. Since singular value decomposition is a non-linear process the resulting modes in the tail regions, where mode energies are more closely spaced, may be more susceptible to random mixing and reordering. For this range of the spectrum the relative energies are still considered to be a valid comparison but detailed mode correlation is less clear due to unknown sensitivities.

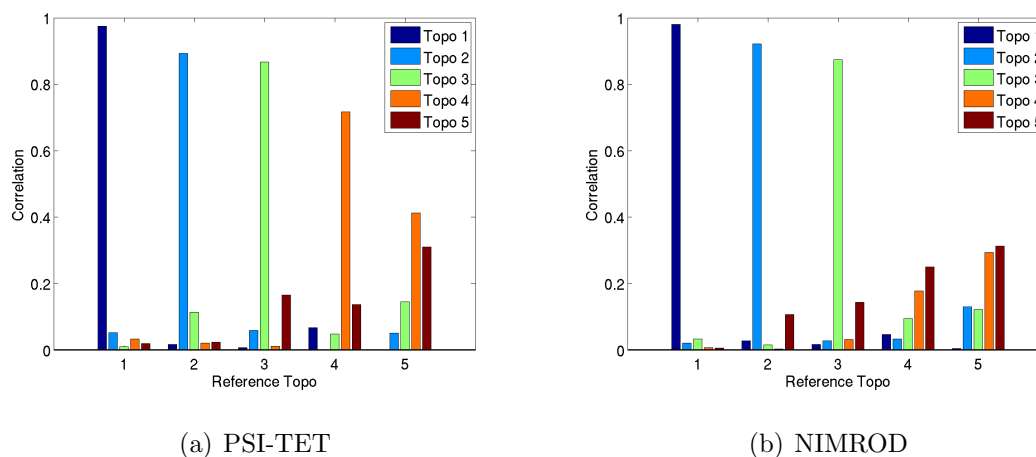


Figure 5.22: Topo correlations between simulation and experimental data (shot 122385) from independent BD analysis.




Figure 5.23: Chrono correlations between simulation and experimental data (shot 122385) from independent BD analysis.

Shared Basis Comparisons

To address the concerns of ambiguity in correlation metrics from independent BD analysis a new decomposition method has been developed using the decomposition of a reference data set to analyze the remaining data sets. A data set is first chosen as a reference and decomposed using BD. The resulting singular vectors can then be multiplied against the other signal matrices, equation 5.15, to produce the projected topo and chrono matrices, $\bar{\phi}$ and $\bar{\Psi}$. The individual mode vectors are then normalized to produce a projected weight matrix, \bar{A} . If the reference data set is used as the signal matrix to project the original BD modes are retrieved by this method, equation 5.16. The projection order also has an effect on the weights in the \bar{A} matrix, so the weights produced from projection to retrieve $\bar{\phi}$ and $\bar{\Psi}$ will be different.

$$\begin{aligned}\bar{A}\bar{\Psi}^T &= (\phi^{ref})^T Y \\ \bar{\phi}\bar{A} &= Y\Psi^{ref}\end{aligned}\tag{5.15}$$

$$A^{ref} (\Psi^{ref})^T = (\phi^{ref})^T \left\{ \phi^{ref} A^{ref} (\Psi^{ref})^T \right\} = (\phi^{ref})^T Y^{ref}\tag{5.16}$$

With this method, non-linearity in the SVD operation is removed from the comparison and the resulting decompositions and correlations, which now involve $\bar{\phi}_k$ and $\bar{\Psi}_k$, have a slightly modified interpretation. The decomposition for the non-reference data sets can be thought of as a Fourier like decomposition, where the basis functions are not sines and cosines but dominant structures and waveforms identified from the reference data. If an experimental data set is used then the decomposition can be considered to adapt its basis set to the important dynamics in the experiment, as defined by signal energy. This interpretation should be used only as a conceptual aide however, as the requirement that all modes be orthogonal in time and space from SVD may restrict the signal energy in unknown ways.

The BD analysis above was repeated using the shared basis method presented




Figure 5.24: Mode weights from reference BD analysis of experimental data (shot 122385) and PSI-TET and NIMROD validation cases.

here with the experimental data set (shot 122385) used as the reference. Projection was done using the reference topologies to retrieve temporal signals of dominant spatial structures from the experimental data. The resulting weights from this analysis are shown in figure 5.24. Agreement in the first 3 mode energies is very similar to the independent BD results, figure 5.17, with a noticeable decrease in the signal energy for the second and third modes. With this method the tail of the energy spectrum yields much poorer agreement between simulations and the reference data. This effect is expected as relatively small amounts of energy that are missing from higher energy modes must be made up by large changes in the energy of lower energy modes. Agreement through the first 5 modes is comparable to the level seen in independent analysis.

Analyzing the shape of the projected chronos produced for the first mode, figure 5.25, shows a noticeable difference in the waveform from independent analysis. The oscillations in the amplitude exhibit a character that agrees much better with the observed oscillations than independent chronos. Late in time, when both injector amplitude and toroidal currents are comparable for all 3 cases, the amplitude, phase, and




Figure 5.25: Mode 1 chrono from reference BD analysis of experimental data (shot 122385) and PSI-TET and NIMROD validation cases.

structure of the oscillations from the PSI-TET simulations match the experimental oscillations well. NIMROD also shows better agreement in this regard, however the convex shape seen in PSI-TET and the experiment is not reproduced. As this mode is generally associated with the mean field equilibrium this change from independent to shared basis analysis provides further indication that the mean field structure is different between simulations and experiment.

The second and third chronos, figure 5.26, show little change from the independent analysis. The signal energies however, show a significant reduction relative to the independent decomposition. The PSI-TET energies drop to comparable or below the experimental values despite a higher injector amplitude over the majority of the shot. NIMROD energies drop a similar amount but stay above the experimental spectrum as would be expected from the difference in injector strength. Although these modes contain primarily injector frequency in the temporal wave forms, these changes indicate that they are not solely due to, or fully capture, injector driven fields. In fact deviations in these first 3 modes are likely the cause of the higher tail energies in PSI-TET and NIMROD simulations as small differences in these modes

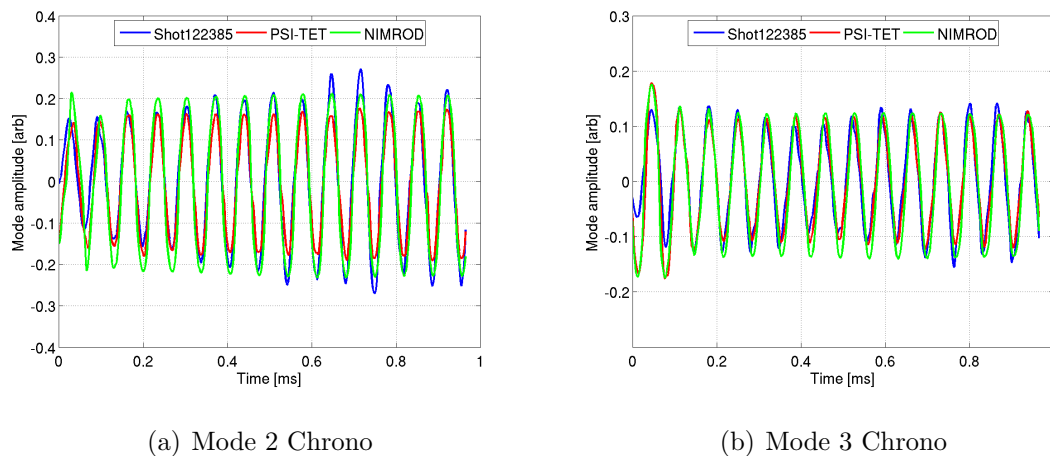


Figure 5.26: Modes 2 and 3 chronos from reference BD analysis of experimental data (shot 122385) and PSI-TET and NIMROD validation cases.

are canceled by higher wavelength modes.

The projected chrono for the fourth mode, figure 5.27, shows a more significant difference from the independent analysis. A persistent offset is seen in both the PSI-TET and NIMROD data that is not present in the experimental data. PSI-TET still shows good agreement in phase and frequency with the experiment, while NIMROD now exhibits additional frequencies as well as the phase offset seen previously. The existence of an offset in simulation data points to a locked mode, but indications of this have not been seen in other analysis. The locked mode may be a portion of the mean field equilibrium in the simulations that is oscillatory in the experiment. A possibility that is supported by the spatial structure of the mode as discussed above. In this case the low relative weight of mode 4 compared to mode 1 would make the structure difficult to distinguish as locked using additional methods such as Fourier decomposition.

Agreement in the correlation metric, which now takes the form of equation 5.17, remains fairly consistent with the independent analysis, however an increase in the

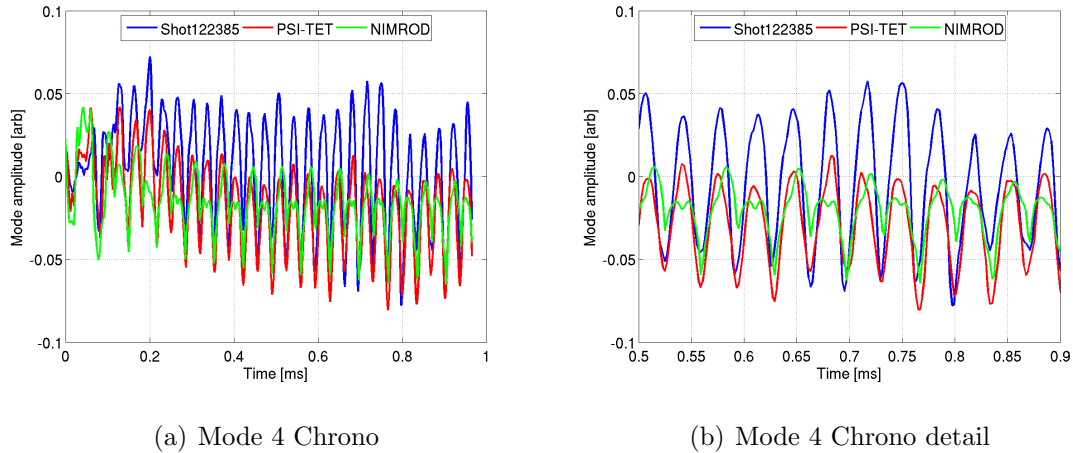


Figure 5.27: Mode 4 chrono from reference BD analysis of experimental data (shot 122385) and PSI-TET and NIMROD validation cases.

cross coupling is seen across the board. This cross coupling is likely due to removing the influence of frequency on the independent decompositions. When BD is performed on each data set independently the constraint of orthogonality in the chrono set as well as the energy ordering will tend to separate slower frequency behavior in the same way the spatial structure is separated. With the shared basis method however the frequency separation only occurs in the reference data, as a result the differences between spatial and temporal structure between data sets end up being combined in this metric. This can be seen most noticeably in the correlation of the time behavior in chrono 4 from PSI-TET and NIMROD with the first chrono of the reference shot. The first mode in the independent decompositions, and the reference decomposition, does not have oscillatory behavior. In independent decomposition this fact would force the majority of spatial structure with an offset in time into this mode. Therefore, the large cross correlation seen in figure 5.17 shows up as a smaller error in the topo and chrono correlations in figures 5.22 and 5.23 for the first mode in each data set.




Figure 5.28: Chrono correlations between simulation and experimental data (shot 122385) from reference BD analysis.

$$C_{i,j}^{ref,2} = \Psi_i^{ref} \cdot \bar{\Psi}_j^2 \quad (5.17)$$

5.5.4 Discussion

The comparisons presented above demonstrate a high level of agreement between the PSI-TET simulations of HIT-SI and the experimental observations from shot 122385. This level of agreement is very similar to the agreement seen in NIMROD validation simulations, indicating that the qualitative behavior of the experiment is captured by both models. Semi-quantitative agreement is also found in many of the diagnostics compared, consistent with NIMROD validation simulations as well. In addition to this baseline agreement there are several areas where better agreement is achieved with the PSI-TET model. Qualitatively, the observed poloidal signal from the internal magnetic probe shows a better match, in terms of the oscillation behavior, with PSI-TET results. Oscillations observed on the first mode chrono from shared basis BD analysis also show better agreement for the PSI-TET simulations than NIMROD

results. Quantitatively, the surface magnetic probe structure correlates well with the spatial and temporal evolution of an additional BD mode for PSI-TET results when compared to NIMROD. Better agreement is also seen in the toroidal current scaling, which will be discussed in the next section (5.6.2). Overall, PSI-TET appears to provide better agreement with the observations from the experiment. However, these differences are small and do not indicate that the NIMROD approximations have a significant impact on modeling of HIT-SI using reduced Hall MHD. Differences are still present in both PSI-TET and NIMROD simulations with the mean field equilibrium as well as current amplifications at higher injector frequencies. Additional physics is likely needed to capture the remaining dynamics of HIT-SI in a quantitative manner.

5.6 *HIT-SI Physics Results*

With the validation presented in the previous section complete, new physical features introduced by modeling the dynamics in the injector volumes and their drive can be investigated. This section will present and discuss features observed in simulations of HIT-SI that are related to the injectors as well as the injector impedance in PSI-TET derived by scanning resistivity.

Injector Coupling

By including the injectors in the simulation domain the connection of the injector fields to the magnetic object in the confinement volume can be analyzed. Early in time the confinement volume is dominated by an $n = 1$ object that shares the symmetry of the injectors, as discussed in the introduction. Figure 5.29 shows the magnetic field structure from a HIT-SI simulation in the X-Z plane during this time. The X-injector is shown at a time of peak flux and current. The entire magnetic field structure is $n = 1$ symmetric with the injector magnetic field linking into an aligned structure in the confinement volume that closely resembles the $n = 1$ eigenmode shown in figure 3.1(b). This structure rotates in time to align itself with the active

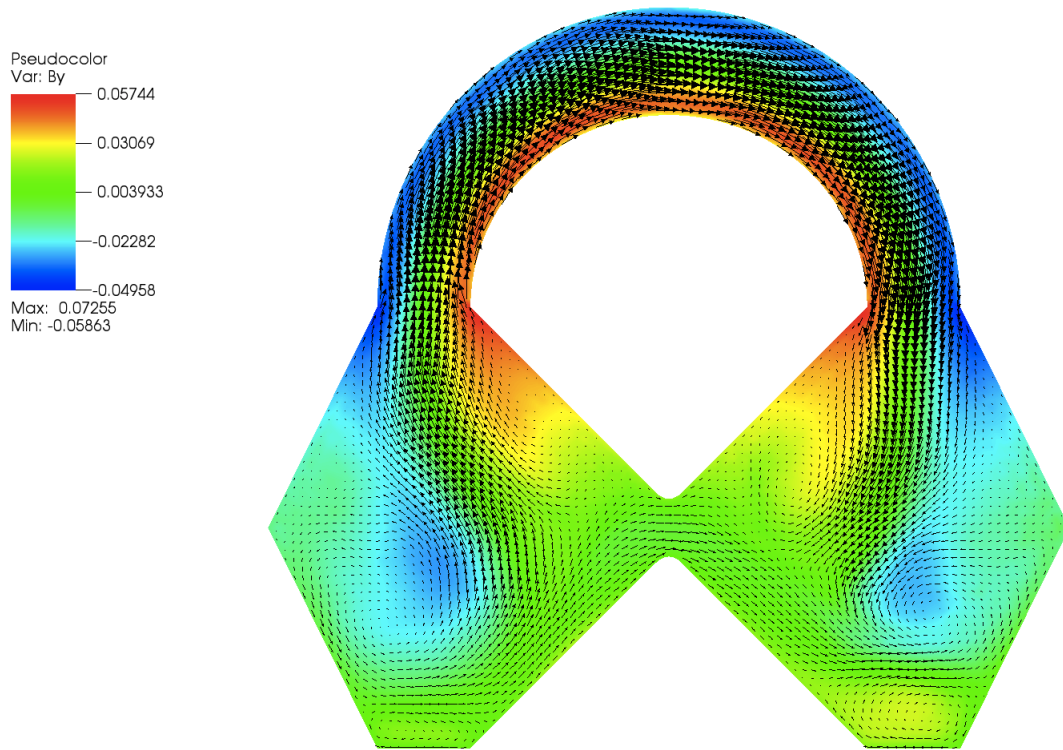


Figure 5.29: Magnetic field in the X-Z plane from a PSI-TET reduced Hall-MHD simulation before formation of a toroidal current (0.052 ms).

injector, as seen in experimental data and NIMROD simulations.

The $n = 1$ symmetric object in the confinement volume is continuously built up by the injectors. However, this mode is unstable and relaxes, once it reaches sufficient free energy, to the desired spheromak equilibrium with $n = 0$ symmetry. As the spheromak energy is increased the injector fields are forced to align with this object by virtue of reconnection, producing a new state when the X-injector is at its peak in the cycle, shown in figure 5.30. The spheromak object is clearly visible in this figure as a large $n = 0$ structure in the y -component of the magnetic field. Magnetic field from the injector is forced to align with the spheromak as it exits the injector, causing




Figure 5.30: Magnetic field in the X-Z plane from a PSI-TET reduced Hall-MHD simulation following formation of a toroidal current (0.5 ms).

the injector driven flux and current channel to “lean” against the inner and outer walls of the flux conserver. This configuration, which was expected during design of HIT-SI[89] and observed in NIMROD simulations, is conducive to current drive in the spheromak by IDCD. As the main magnetic structure is now largely symmetric it does not rotate, but instead the injector channels reconnect during reversal of the injector field to align with the spheromak on the opposite side.

As expected due to the qualitative agreement these results are largely the same as in NIMROD simulations. Small recirculation regions in the magnetic field are seen in NIMROD simulations near the corners of the grid opposite the injector mouth that

are not seen in PSI-TET. This may be due to differences in the surface boundary condition, if the resistive layer is not allowing magnetic flux from the injector to distribute in the same way as PSI-TET. However, these regions are small and may not be significant in the formation and current drive dynamics. Reconnection location and activity during reversal shows a more significant difference due to the additional volume available in PSI-TET simulations. This will be discussed in section 5.6.1.

5.6.1 *Injector Dynamics*

Adding the injector volumes to the simulations introduces a new and rich set of dynamics in these regions that have not been modeled previously. These regions turn out to be some of the more numerically challenging regions due to dynamics associated with field reversal that occurs 2 times per injector period in each injector. The coils driving the injectors always push field in from the boundary, meaning that reversal occurs by injecting magnetic flux and current in the opposite direction of the current fields. As the new flux is injected the existing flux is compressed into a smaller and smaller channel in the center of the injector until the current gradients drive instability at the interface and the interior flux breaks up into islands and reconnects. Figure 5.33 shows an example of the magnetic field structure midway through the reconnection process in the X-injector. Magnetic flux in the confinement volume but linked to the injectors appears to reconnect to the spheromak object, with which it is already well aligned except at the injector mouths.

This reversal process drives small scale fluctuations and large flows, requiring a smaller time step to capture. Although the reconnection layer is not resolved in these simulations, the time step restriction to capture the dynamics during this period remains noticeable. The primary reconnection phase spans less than 3 μs in 14.5 kHz simulations, making it much faster than other non-linear phenomena. This time scale is consistent with the Sweet-Parker[90, 91] reconnection time, $\tau_{sp} = \sqrt{\tau_a \tau_{L/R}} \approx 2.5 \mu\text{s}$, computed for current sheets in the injector just before the onset of reconnection. At

this time the injector flux reverses on a scale of ≈ 4 cm with average field strength in the reversal channel of ≈ 50 mT, yielding a resistive decay time of $\tau_{L/R} \approx 50 \mu s$ and an Alfvén transit time across the channel of $\tau_a \approx 0.2 \mu s$. This entire process is inherently different from the injector fields applied in NIMROD simulations, which vary in amplitude but not shape at the injector mouths over the course of a period. The additional experimental agreement seen with PSI-TET simulations compared to NIMROD simulations may be primarily due to this difference. However, due to the stiffness of the dynamics associated with modeling this behavior the added cost is significant.

Flux Circuit Loading and Injector Phasing

The effect of plasma response in the injectors also has an effect on the phase between the flux driver and the resulting injector flux. In PSI-TET simulations the action of the flux coil is modeled by imposing a boundary condition of $\oint \vec{B} \cdot d\vec{l}$ on the surface of each injector. This constraint amounts to enforcing a total current in the flux coils, shown as orange windings in figure 1.4. The loop voltage providing the flux injection is then dependent on the inductive screening and resistive load of the plasma on the circuit, which depends strongly on internal dynamics. As a result, as the injector driving frequency is increased a phase shift is introduced between the coil currents and the resulting injector flux, indicating an increase in resistive loading during portions of the cycle. This effect is shown in figure 5.31 where the flux and current signals for a simulation with $f_{Inj} = 53.5$ kHz are out of phase. When the flux and current are no longer in phase, regions of the injector volume are being driven with opposite signs of helicity causing negative power and helicity injection. This reduces the efficiency of the injectors in a way inconsistent with the experiment.

To prevent this a flux control algorithm has been added using a simple PID type control method. A desired flux waveform (Ψ_D) is specified and the actual flux signal (Ψ_S) is computed during the simulation to determine a signal error (ϵ). The flux




Figure 5.31: Comparison of flux and current waveforms in the Y-injector from PSI-TET simulations at 53.5 kHz with (a) and without (b) flux control.

coil boundary condition is then scaled based on the error and its history as $\Psi_A = \omega_P \epsilon + \omega_I \int \epsilon dt + \omega_D \frac{d\epsilon}{dt}$, where ω_P , ω_I , and ω_D are three gains that are adjusted to provide the best performance. This method was chosen for its simplicity and is widely used in control systems in a variety of fields. When flux control is applied the flux signal is returned to the desired phase and waveform as shown in figure 5.31(a). Driving a sine waveform for the flux requires significant deviation from this waveform in the coil current, shown in figure 5.32, which is consistent with increasing resistive effects due to plasma screening. The new current waveform also highlights the change in character of the circuit load over the course of an injector cycle. A reducing phase shift between the voltage and current, indicating an increase in dissipation in the plasma load, is seen as the injector flux is reversed. The phase shift then returns following the reconnection process in the injector as the flux increases with fixed circuit current. This effect is also seen in the injector flux without control (figure 5.31(b)), where the slope stagnates before reconnection yields a uniform sign of magnetic flux in the injector. Therefore, flux compression preceding the reconnection process is likely




Figure 5.32: Comparison of flux coil voltage and current for a PSI-TET simulation at 53.5 kHz with flux control.

playing a role in modifying the plasma response and producing the observed phase shifts.

In the experiment, deviation in the coil currents from a pure sine wave to the same extent is not observed. However, an increase in the flux circuit dissipation is observed as injector frequency is increased. This difference is likely due to dynamics that are not included in the reduced Hall MHD model but are likely involved in the injector operation. Ionization near gas puff locations in the injectors may lead to low temperature or density regions where resistivity is increased, reducing plasma screening. Low density regions also allow the plasma in that region to be accelerated more easily, moving the flux in faster by advecting it with the fluid. Further investigation is needed to determine if this discrepancy negatively effects the agreement of PSI-TET simulations at high frequency.

Reconnection Asymmetry

In PSI-TET simulations the injectors exhibit asymmetry in the strength of the fluctuations associated with field reversal. For the HIT-SI validation case presented

above the X-injector is seen to have much larger fluctuations, figure 5.33, than the Y-injector, figure 5.34. Both plots show the magnetic field at approximately the peak of reconnection activity in each injector with similar scales for the field vectors and pseudocolor, set by the equilibrium field magnitude. However, only the fluctuations in the X-injector are clearly visible on this scale, while the Y-injector fluctuations are much less pronounced. This asymmetry remains throughout the entire shot, with the X-injector fluctuations stronger at both reversal times in each period. Which injector has the stronger reconnection does not appear to be dependent on the direction of toroidal current. However, this has only been tested for a single case and has not been investigated thoroughly.

The reconnection within a single injector also shows asymmetric behavior with respect to the injector mouths. Island structures in figure 5.33 are more concentrated near the injector mouth where flux was flowing out at start of reversal. This holds true at the opposite reversal time where the island structures are clustered more toward the left injector mouth in figure 5.33. The connection to the confinement volume may act to stabilize or destabilize the field more at one injector mouth than the other, which may delay the reconnection, leading to the asymmetry. This explanation is consistent with PSI-TET results as the direction of the clustering is found to be dependent on the direction of the toroidal current.

Flow Asymmetry

There also exists an asymmetry in the flow observed near the injector mouths. At times of peak flux for a given injector, flows away from the injector mouths are produced, with the strongest flow observed at the mouth where flux is entering the injector. At this mouth the flow opposes the direction of current in the same region. This is shown in figures 5.35 and 5.36 for the X-injector at a time in the period near peak flux and current. Figure 5.35 corresponds to shortly after figure 5.33 in time, while figure 5.36 is a half period later in time. A poloidal flow is seen near each injector




Figure 5.33: Magnetic field in the X-Z plane from a PSI-TET reduced Hall-MHD simulation showing large fluctuations associated with field reversal in the X-injector (1.008 ms).




Figure 5.34: Magnetic field in the X-Z plane from a PSI-TET reduced Hall-MHD simulation showing smaller fluctuations associated with field reversal in the Y-injector (1.018 ms) when compared to the X-injector (figure 5.33).

mouth, with a strong axial flow out of the mouth where injector reconnection was strongest on the previous field reversal. A weaker axial flow is observed at the opposite injector mouth where the magnetic flux from the injector is turned more sharply to follow the wall. These flows are also seen in NIMROD simulations indicating that it is likely due to effects outside the injector. In PSI-TET simulations the flow primarily develops outside of the injector, which is in agreement with $\vec{J} \times \vec{B}$ forces outside the injector driving the flow. Investigation is currently underway to compare flow profiles in PSI-TET against IDS measurements of flow velocity in the experiment.

5.6.2 Toroidal Current Scaling

In the experiment the impedance of the voltage coil circuit is found to consist of only a real impedance, which scales linearly as the ratio j/n in the confinement volume[32]. This scaling is determined by comparing the coil voltage and current waveforms throughout shot with the toroidal current and line average electron density. In simulations the density is fixed, so the impedance should only depend on the toroidal current. The impedance can be determined from a helicity balance where Ohmic dissipation from the spheromak equilibrium is the dominant dissipation. For both PSI-TET and NIMROD the injector flux and current are fixed so the rate of helicity injection depends only on the impedance. Helicity dissipation is due to resistive decay of the spheromak equilibrium with a characteristic time, $\tau_K = \frac{\mu_0}{2\eta\lambda^2}$, that is known from resistivity and λ_{Taylor} . In steady state this gives the helicity balance in equation 5.18, where Z_{Inj} is the voltage coil impedance, Ψ_{Inj} is the quadrature injector flux, I_{Inj} is the quadrature injector current, and $K_{Tor} \sim I_{Tor}^2$ is the magnetic helicity associated with the toroidal current. In HIT-SI simulations, this helicity balance is only expected to hold when the toroidal current has reached saturation.

$$Z_{Inj}\Psi_{Inj}I_{Inj} = \frac{K_{Tor}}{\tau_K} \quad (5.18)$$




Figure 5.35: Plasma flow velocity in the X-Z plane from a PSI-TET reduced Hall-MHD simulation near peak injector flux (1.008 ms). Flow is visible away from each injector mouth in the in-plane velocity field (black vectors), with an associated rotation in the y-velocity (pseudocolor). More coherent and collimated axial flow is observed near the 0° mouth (right).




Figure 5.36: Plasma flow velocity in the X-Z plane from a PSI-TET reduced Hall-MHD simulation 1/2 injector cycle after figure 5.35. The collimated axial flows are now observed with the 180° mouth (left).

In previous studies with NIMROD the background resistivity was scanned to determine the dependence of the peak toroidal current on resistivity. If $Z_{Inj} \sim j/n$ then the toroidal current should scale as $I_{Tor} \sim \eta^{-1}$. However, NIMROD simulations with reduced Hall MHD found a scaling of $I_{Tor} \sim \eta^{-0.6}$, lower than the scaling expected from experimentally observed Z_{Inj} . In PSI-TET a stronger scaling is seen, which is more consistent with the experiment. However, as this impedance scaling also acts within a single simulation the time to peak toroidal current increases significantly beyond the resistive decay time due to the increase in power and helicity injection. This makes using equation 5.18 at current saturation impractical as the simulated time required increase with reductions in η . To determine the impedance scaling in a shorter amount of time the full evolution of the spheromak helicity is considered. Equation 5.19 provides an evolution equation for the spheromak helicity under the assumptions above for source and dissipation and adds the assumption that the injector impedance scales like a power of the toroidal current.

$$\dot{K}_{Tor} = C_1 I_{Tor}^x \Psi_{Inj} I_{Inj} - \frac{K_{Tor}}{\tau_K} \quad (5.19)$$

This equation is then used to model the growth of the toroidal current in time for 3 PSI-TET simulations at different resistivities ($\eta = [25, 16, 12.5]$). The ODE is initialized for each simulation just after toroidal current formation and evolved in time and compared to the simulated currents in figure 5.37. The same value for the free coefficient (C_1) is used for all 3 cases and was adjusted to provide the best fit across all shots. Finally, the power of the impedance scaling (x) was adjusted for each shot to provide the best fit. The resulting impedance scaling was found to be $Z_{Inj} \sim I_{Tor}^{1 \pm 0.01}$, which is in agreement with the experimentally observed scaling and the IDCD model[32].

Injector impedance is likely a function of the plasma response shaping injector current paths in the device leading to a magnetic configuration that supports imposed dynamo current drive. IDCD activity strongly couples dissipative loading due to




Figure 5.37: Comparison of toroidal currents from simulations at different resistivities to model current evolutions assuming impedance scaling for the injector voltage circuit of the form $\sim j/n$.

currents in the interior of the plasma to the edge region, where injector current is driven, resulting in the observed impedance scaling[32]. As PSI-TET models the injector coupling, this improved agreement may indicate that the impedance effect requires redistribution of the current profile in the injector that is not possible in the NIMROD model. Another explanation for this difference may be the asymmetries in the injector reconnection and flow fields discussed above. This is less likely however as these effects are limited to small, low power portions of the injector cycle.

5.7 Density Evolution and Transport

The remaining discrepancies between reduced Hall MHD models of the experiment using PSI-TET and NIMROD point to the fact that additional physics beyond the reduced model are important in HIT-SI. This is further indicated by recent experimental observation of large Shafranov shifts, which are indicative of confined pressure, in the mean field profile at high injector drive frequency. To investigate additional physics associated with variation and transport of plasma density and thermal energy

the PSI-TET reduced MHD model has been extended to include the full set of Hall MHD equations. This section will briefly introduce the initial development of this model in PSI-TET in addition to early investigation of HIT-SI using this model.

5.7.1 Full Hall-MHD

In deriving the reduced model used in the previous section, the plasma density and temperature were assumed to be uniform and constant. Without this assumption, the system of equations for the full Hall MHD model must be used. The implementation in PSI-TET uses the set of equations 5.20a - 5.20e, which retain the assumption of kinematic viscosity and the inclusion of electron inertia from the reduced model. This systems assumes rapid equilibration between the electron and ion fluids ($T = T_i = T_e$) such that the total fluid pressure is given by $P = 2nkT$. To close the system a simple Fourier-like heat flux vector is defined (eq. 5.20f) based on the ion temperature gradient with anisotropic coefficients for the parallel (χ_{\parallel}) and perpendicular (χ_{\perp}) conduction relative to the magnetic field direction. Energy transfer from magnetic and kinetic to thermal energy due to collisional effects is included through a heat source term (eq. 5.20f) for viscous and Ohmic sources. Particle diffusion is also added to the continuity equation for numerical stability.

The introduction of the continuity and energy equations supports additional wave behavior and dynamics that were not present in the reduced model. In particular plasma sound waves and their field coupled magneto-sonic counterparts break the symmetry of Alfvén wave propagation in the parallel and cross-field directions as discussed in section 5.2.1. This variation in wave speed with propagation direction introduces new sources of stiffness into the problem. In addition to new wave behavior many of the coefficients are temperature and density dependent. Although presently the resistivity is the only coefficient that varies during a simulation, as a function of temperature. The detailed weak forms of the non-linear function and Jacobian matrix used in evolving this system in PSI-TET are provided in appendix D.2.1. The multi-

grid preconditioner developed for reduced MHD still performs well with the addition of these non-linearities and stiffness sources. However, as the number of degrees of freedom per cell is increased by the additional fields being modeled the solution cost is increased over the reduced MHD model.

$$\frac{\partial n}{\partial t} + \nabla \cdot (nu) = D\nabla^2 n \quad (5.20a)$$

$$\rho \left[\frac{\partial u}{\partial t} + u \cdot \nabla u \right] = J \times B - \nabla (2nkT) + \nabla \cdot \nu \nabla u \quad (5.20b)$$

$$\frac{n}{\gamma - 1} \left[\frac{\partial T}{\partial t} + u \cdot \nabla T \right] = -nkT \nabla \cdot u - \nabla \cdot \vec{q} + Q/2 \quad (5.20c)$$

$$\frac{\partial B}{\partial t} = -\nabla \times \left(-u \times B + \eta J + \frac{1}{ne} (J \times B - \nabla p_e) + \frac{m_e}{ne^2} \frac{\partial J}{\partial t} \right) \quad (5.20d)$$

$$\mu_0 J = \nabla \times B \quad (5.20e)$$

$$\vec{q} = -n \left[\chi_{\parallel} \hat{b}\hat{b} + \chi_{\perp} (I - \hat{b}\hat{b}) \right] \cdot \nabla T \quad (5.20f)$$

$$Q = \eta J^2 + \nu (\nabla u)^T : \nabla u \quad (5.20g)$$

5.7.2 Initial HIT-SI Results

This model has been used to perform initial investigation of the effects of density and temperature variation in HIT-SI. The validation simulation presented in section 5.5 was repeated as closely as possible with the full MHD model. Injector parameters and viscosity were kept the same as in the validation case with the same boundary conditions for the magnetic field and plasma velocity. Resistivity with a Spitzer[92] scaling in temperature was used of the form $\eta = \eta_{ref} (T_{ref}/T)^{3/2}$, with $\eta_{ref} = 25$

and $T_{ref} = 6$. A Dirichlet boundary condition was used for the the density and temperature equations with uniform boundary values of $n = 1.5 \times 10^{19} [\text{m}^{-3}]$ and $T = 1.5 \text{ eV}$ respectively. Fixed thermal conduction coefficients of $\chi_{\parallel} = 10^4 [\text{m}^2/\text{s}]$ and $\chi_{\perp} = 200 [\text{m}^2/\text{s}]$ were used, which correspond approximately to the values given by a Braginskii model[93] for $B = 20 \text{ mT}$, $n = 1.5 \times 10^{19} [\text{m}^{-3}]$, and $T = 6 \text{ eV}$ using the enhanced electron mass. Finally, a particle diffusion coefficient of $D = 2 \times 10^3 [\text{m}^2/\text{s}]$ was used to prevent the density from going negative during the simulation. This value was found by adjusting D to find the lowest value where negative density was not introduced by dynamics and is comparable to the value used for the same purpose in NIMROD simulations of SSPX[19]. Ohmic heating was the only heat source included for this run to limit the temperature near the experimentally observed level. The electron pressure contribution to Ohm's law was also neglected in this initial run, but has now been implemented in the model.

Toroidal current for this simulation is shown in figure 5.38 compared to the experimental current from shot 122385. As the heating term increases the temperature quickly, the resistivity drops accordingly to an average value below the constant value used in the validation case. This reduced dissipation causes a much faster rise in the toroidal current than was seen in the validation case. The average temperature at the end of the simulation is $\sim 15 \text{ eV}$, which is above the experimentally observed peak temperatures. However, the simulation does not include energy losses due to neutral ionization and radiative effects that are expected to limit temperatures seen in the experiment. Additionally, the assumption of fixed thermal conductivities may be artificially limiting thermal conduction to the wall in areas where magnetization is low. As the plasma temperature continued to grow beyond the experimental range the simulation was stopped at $\sim 0.35 \text{ ms}$. The Braginskii model is currently being implemented to provide more accurate thermal transport for future simulations.

To account for the difference in current growth times the full MHD simulation is shifted further in time to match the toroidal currents for comparison of internal mag-




Figure 5.38: Comparison of I_{Tor} from shot 122385 and PSI-TET full Hall-MHD simulation at 14.5 kHz.

netic signals. The simulation is shifted by an additional 4 injector period compared to the validation case for a final shift of $986.3 \mu\text{s}$. Magnetic field signals along the internal probe for this simulation are shown in figure 5.39, compared to the experimental signals from shot 122385. The level of agreement in the phase and amplitude for the toroidal signals is similar to the reduced MHD case. Although, slightly better agreement is seen in the mean amplitude near the magnetic axis, bottom traces in the figure. The most significant change from the reduced case is seen in the poloidal signals, where fluctuations through much of the volume are decoupled to a large extent from the injector period. This effect is seen on the experimental signals, but was not captured by reduced MHD simulations with NIMROD and only partially captured by previous PSI-TET simulations. The amplitude of the mean field also appears to be more consistent near the experimentally observed magnetic axis. Longer simulations are needed and underway to properly assess the differences between the two models used on HIT-SI and determine if better quantitative agreement is achieved for magnetic signals when plasma transport effects are included.

When temperature and density are allowed to evolve, variation in these quantities




Figure 5.39: Magnetic field along the internal magnetic probe for a PSI-TET full Hall-MHD simulation of HIT-SI compared to experiment signals (shot 122385).

is expected to develop across and through the injectors. This effect is one of the major motivations for modeling the injector and confinement volumes in a coupled manner. In PSI-TET simulations heating and thermal transport is seen in the injector leading to temperatures on the injector mouths that are significantly higher than the wall temperature. Figure 5.40 show the temperature in the X-Z plane at the of the PSI-TET full MHD simulation presented above. Thermal connection through the injectors along with heating in the injector current channels are clearly visible in this plot. The three dimensional behavior of the plasma temperature is shown in figure 5.41, which contains two orthogonal views of the 15 eV temperature surface at the same time in the simulation. This effect highlights the need to model the full plasma volume in HIT-SI to accurately capture the dynamics of transport in a consistent manner.




Figure 5.40: Ion temperature in the X-Z plane from a PSI-TET full Hall-MHD simulation showing heating and thermal connection through the injectors.




Figure 5.41: Ion temperature contours in 3D from a PSI-TET full Hall-MHD simulation showing thermal connection through the injectors.

5.8 Future Work

The development and investigation performed as part of this work have resulted in a validated code capable of simulating complex geometries in a user friendly and efficient manner. PSI-TET is now being used to perform in depth investigation of the physical mechanisms involved in SIHI current drive, that were not feasible with existing tools. Current study is focused on simulations of high frequency injector operation and full Hall MHD simulations to determine the nature and effect of observed plasma pressure. High frequency operation in HIT-SI has produced the best performance to date, however poorer agreement with simulations is also observed at these frequencies. A frequency scan is currently under way for PSI-TET with early simulations at 53.5 kHz showing a similar deviation as NIMROD using the reduced MHD model. At higher injector frequency the effect of plasma response on the injector field waveform

will also be more pronounced in PSI-TET, due to the dependence of the resistive skin depth and the power required to drive injector fields on frequency. This may change the way that PSI-TET reduced MHD simulations behave when compared to NIMROD simulations where the fields have a relatively fixed profile.

Early investigation with the full MHD model has shown some improved agreement with the experiment, but additional investigation is needed to confirm this and study the effect in more detail. Resolution studies will need to be repeated with the new physical system to ensure that results are converged spatially. The proposed effect of pressure driven current drive will likely involve small spatial scales that may require increased resolution over the reduced simulations to capture. Variable thermal transport coefficients are being added to the full MHD model with the Braginskii form to capture magnetization effects in low field regions. Even with these improved transport physics, the heating terms may raise the plasma temperature higher than the experimentally observed values. Adjustment of transport coefficients and sink terms will likely be necessary to capture the effect of neutrals and radiation on the temperature profile. A more consistent solution to the lack of neutral effects is to implement the model of Shumlak and Meier[94] that is currently being implemented in NIMROD. This is being considered for PSI-TET as it would allow neutral effects, which are important to many small laboratory experiments such as HIT-SI, in a coupled manner.

The flexibility of the PSI-TET code to model complex geometries is also motivating its use on other experiments. Work is currently under way to add the ability to model solid conducting structures along with the plasma model. This will support the study of plasma interaction with passive or active in vessel conducting plates that are common in many tokamak devices. These structures often have breaks and other features in the toroidal and poloidal directions that make them difficult to capture with existing codes. With this new capability PSI-TET will be used to study in vessel conducting plates for their linear and non-linear stabilization properties as well

as torque and destabilization induced by resistive drag.

BIBLIOGRAPHY

- [1] E.T. Meier, V.S. Lukin, and U. Shumlak. Spectral element spatial discretization error in solving highly anisotropic heat conduction equation. *Computer Physics Communications*, 181(5):837 – 841, 2010.
- [2] Cihan Akcay, Charlson C. Kim, Brian S. Victor, and Thomas R. Jarboe. Validation of single-fluid and two-fluid magnetohydrodynamic models of the helicity injected torus spheromak experiment with the NIMROD code. *Physics of Plasmas*, 20(8):082512, 2013.
- [3] A. Bondeson, G. Marklin, Z. G. An, H. H. Chen, Y. C. Lee, and C. S. Liu. Tilting instability of a cylindrical spheromak. *Physics of Fluids (1958-1988)*, 24(9):1682–1688, 1981.
- [4] J. Finn and T. Antonsen. Magnetic Helicity: what is it and what is it good for? *Comments on Plasma Physics and Controlled Fusion*, 9(111), 1985.
- [5] J. B. Taylor. Relaxation of toroidal plasma and generation of reverse magnetic fields. *Phys. Rev. Lett.*, 33(19):1139–1141, Nov 1974.
- [6] T R Jarboe. Review of spheromak research. *Plasma Physics and Controlled Fusion*, 36(6):945, 1994.
- [7] M.N. Rosenbluth and M.N. Bussac. MHD stability of spheromak. *Nuclear Fusion*, 19(4):489, 1979.
- [8] T. R. Jarboe, W. T. Hamp, G. J. Marklin, B. A. Nelson, R. G. O’Neill, A. J. Redd, P. E. Sieck, R. J. Smith, and J. S. Wrobel. Spheromak formation by steady inductive helicity injection. *Phys. Rev. Lett.*, 97:115003, Sep 2006.
- [9] D.A. Ennis, B.S. Victor, J.S. Wrobel, C. Akcay, T.R. Jarboe, G.J. Marklin, B.A. Nelson, and R.J. Smith. New understandings and achievements from independent-injector drive experiments on HIT-SI. *Nuclear Fusion*, 50(7):072001, 2010.
- [10] T.R. Jarboe, C. Akcay, M.A. Chilenski, D.A. Ennis, C.J. Hansen, N.K. Hicks, R.Z. Aboul Hosn, A.C. Hossack, G.J. Marklin, B.A. Nelson, R.G. O’Neill, P.E.

- Sieck, R.J. Smith, B.S. Victor, J.S. Wrobel, and M. Nagata. Recent results from the HIT-SI experiment. *Nuclear Fusion*, 51(6):063029, 2011.
- [11] B. S. Victor, T. R. Jarboe, A. C. Hossack, D. A. Ennis, B. A. Nelson, R. J. Smith, C. Akcay, C. J. Hansen, G. J. Marklin, N. K. Hicks, and J. S. Wrobel. Evidence for separatrix formation and sustainment with steady inductive helicity injection. *Phys. Rev. Lett.*, 107:165005, Oct 2011.
- [12] P. M. Bellan. *Spheromaks: A practical application of magnetohydrodynamic dynamos and plasma self-organization*. London: Imperial College Press, 2000.
- [13] J.S. Wrobel. A study of HIT-SI plasma dynamics using surface magnetic field measurements. *PhD Dissertation, University of Washington*, 2011.
- [14] B.S. Victor. Effects of density control on the internal plasma dynamics and current drive in HIT-SI. *PhD Dissertation, University of Washington*, 2012.
- [15] C.R. Sovinec, A.H. Glasser, T.A. Gianakon, D.C. Barnes, R.A. Nebel, S.E. Kruger, D.D. Schnack, S.J. Plimpton, A. Tarditi, and M.S. Chu. Nonlinear magnetohydrodynamics simulation using high-order finite elements. *Journal of Computational Physics*, 195(1):355 – 386, 2004.
- [16] R. D. Milroy, C. C. Kim, and C. R. Sovinec. Extended magnetohydrodynamic simulations of field reversed configuration formation and sustainment with rotating magnetic field current drive. *Physics of Plasmas*, 17(6):062502, 2010.
- [17] V.A. Izzo, E.M. Hollmann, A.N. James, J.H. Yu, D.A. Humphreys, L.L. Lao, P.B. Parks, P.E. Sieck, J.C. Wesley, R.S. Granetz, G.M. Olynky, and D.G. Whyte. Runaway electron confinement modelling for rapid shutdown scenarios in DIII-D, Alcator C-Mod and ITER. *Nuclear Fusion*, 51(6):063032, 2011.
- [18] J. M. Finn C. R. Sovinec and D. del Castillo-Negrete. Formation and sustainment of electrostatically driven spheromaks in the resistive magnetohydrodynamic model. *Physics of Plasmas*, 8(2):475–490, 2001.
- [19] E. B. Hooper, B. I. Cohen, H. S. McLean, R. D. Wood, C. A. Romero-Talamas, and C. R. Sovinec. NIMROD resistive magnetohydrodynamic simulations of spheromak physics. *Physics of Plasmas*, 15(3):032502, 2008.
- [20] A.H. Glasser and X.Z. Tang. The SEL macroscopic modeling code. *Computer Physics Communications*, 164(13):237 – 243, 2004. Proceedings of the 18th International Conference on the Numerical Simulation of Plasmas.

- [21] V. S. Lukin. Computational study of the internal kink mode evolution and associated magnetic reconnection phenomena. *PhD Dissertation, Princeton University*, 2008.
- [22] E.T. Meier. Modeling plasmas with strong anisotropy, neutral fluid effects, and open boundaries. *PhD Dissertation, University of Washington*, 2011.
- [23] W.B. Lowrie. Development and application of a multi-block high order finite element modeling code as an engineering design tool. *PhD Dissertation, University of Washington*, 2011.
- [24] W. Park and E.V. Belova. Plasma simulation studies using multilevel physics models. *Physics of Plasmas*, 6(5):1796, 1999.
- [25] N. M. Ferraro and S. C. Jardin. Finite element implementation of Braginskii's gyroviscous stress with application to the gravitational instability. *Physics of Plasmas*, 13(9):092101, 2006.
- [26] N. Ferraro. Non-ideal effects on the stability and transport of magnetized plasmas. *PhD Dissertation, Princeton University*, 2008.
- [27] V. A. Izzo. Three-dimensional magnetohydrodynamic simulations of the HIT-SI spheromak experiment. *PhD Dissertation, University of Washington*, 2004.
- [28] V. A. Izzo and T. R. Jarboe. Three-dimensional magnetohydrodynamic simulations of the helicity injected torus with steady inductive drive. *Physics of Plasmas*, 12(5):056109, 2005.
- [29] C. Akcay. Extended MHD simulations of the HIT-SI experiment with the NIMROD code. *General Exam, University of Washington*, 2010.
- [30] C. Akcay. Extended Magnetohydrodynamic Simulations of the HIT-SI Spheromak Experiment with the NIMROD Code. *PhD Dissertation, University of Washington*, 2013.
- [31] R.D. Wood, D.N. Hill, E.B. Hooper, S. Woodruff, H.S. McLean, and B.W. Stallard. Improved operation of the SSPX spheromak. *Nuclear Fusion*, 45(12):1582, 2005.
- [32] T.R. Jarboe, B.S. Victor, B.A. Nelson, C.J. Hansen, C. Akcay, D.A. Ennis, N.K. Hicks, A.C. Hossack, G.J. Marklin, and R.J. Smith. Imposed-dynamo current drive. *Nuclear Fusion*, 52(8):083017, 2012.

- [33] Daniel Ryppl. T3D Mesh Generator, 2004. Department of Mathematics, Czech Technical University.
- [34] CUBIT Development Team. CUBIT: Geometry and mesh generation toolkit, 2013. <http://cubit.sandia.gov/>.
- [35] Harold Grad and Hanan Rubin. Hydromagnetic equilibria and force-free fields. *Journal of Nuclear Energy*, 7(3-4):284 – 285, 1958.
- [36] L.L. Lao, H. St. John, R.D. Stambaugh, A.G. Kellman, and W. Pfeiffer. Reconstruction of current profile parameters and plasma shapes in tokamaks. *Nuclear Fusion*, 25(11):1611, 1985.
- [37] R. Aymar, V.A. Chuyanov, M. Huguet, Y. Shimomura, ITER Joint Central Team, and ITER Home Teams. Overview of ITER-FEAT - The future international burning plasma experiment. *Nuclear Fusion*, 41(10):1301, 2001.
- [38] C. Mercier. Equilibrium and stability of a toroidal magnetohydrodynamic system in the neighbourhood of a magnetic axis. *Nuclear Fusion*, 4(3):213, 1964.
- [39] R. M. Mayo and G. J. Marklin. Numerical calculation of Mercier beta limits in spheromaks. *Physics of Fluids (00319171)*, 31(6):1812, 1988.
- [40] J. J. Moré, B. S. Garbow, and K. E. Hillstrom. User Guide for MINPACK-1. *ANL-80-74*, Argonne National Laboratory, 1980.
- [41] Alan H. Glasser. The direct criterion of newcomb for the stability of an axisymmetric toroidal plasma. *Los Alamos Report LA-UR-95-528*, 1995.
- [42] OpenMP Architecture Review Board. OpenMP Application Program Interface Version 3.1, 2011.
- [43] The MPI Forum. MPI: A Message Passing Interface, 1993.
- [44] George Karypis and Vipin Kumar. A fast and high quality multilevel scheme for partitioning irregular graphs. *SIAM J. Sci. Comput.*, 20(1):359–392, December 1998.
- [45] Satish Balay, Jed Brown, Kris Buschelman, William D. Gropp, Dinesh Kaushik, Matthew G. Knepley, Lois Curfman McInnes, Barry F. Smith, and Hong Zhang. PETSc Web page, 2012. <http://www.mcs.anl.gov/petsc>.

- [46] FoX Development Team. *FoX: A Fortran library for XML*, 2013.
- [47] Liu Linbo Zhang Tao Cui Hui. A set of symmetric quadrature rules on triangles and tetrahedra. *Journal of Computational Mathematics*, 27(1):89 – 96, 2009.
- [48] Robert McNeel and Associates. openNURBS Initiative, 2012. <http://www.rhino3d.com/opennurbs>.
- [49] VisIt Development Team. Visit: Visualization tool, 2013. <https://wci.llnl.gov/codes/visit/>.
- [50] J.C. Nédélec. Mixed finite elements in \mathbb{R}^3 . *Numerische Mathematik*, 35:315–341, 1980.
- [51] Chen Greif, Dan Li, Dominik Schtzau, and Xiaoxi Wei. A mixed finite element method with exactly divergence-free velocities for incompressible magnetohydrodynamics. *Computer Methods in Applied Mechanics and Engineering*, 199(45-48):2840 – 2855, 2010.
- [52] R.D. Graglia, D.R. Wilton, and A.F. Peterson. Higher order interpolatory vector bases for computational electromagnetics. *Antennas and Propagation, IEEE Transactions on*, 45(3):329 –342, mar 1997.
- [53] W. Bangerth, R. Hartmann, and G. Kanschat. deal.II – a general purpose object oriented finite element library. *ACM Trans. Math. Softw.*, 33(4):24/1–24/27, 2007.
- [54] J.P. Webb. Hierarchical vector basis functions of arbitrary order for triangular and tetrahedral finite elements. *Antennas and Propagation, IEEE Transactions on*, 47(8):1244 –1253, aug 1999.
- [55] Joachim Schöberl and Sabine Zaglmayr. High order nédélec elements with local complete sequence properties. *COMPEL: The International Journal for Computation and Mathematics in Electrical and Electronic Engineering*, 24(2):374–384, 2005.
- [56] M Deville and E Mund. Chebyshev pseudospectral solution of second-order elliptic equations with finite element preconditioning. *Journal of Computational Physics*, 60(3):517 – 533, 1985.
- [57] M. Deville and E. Mund. Finite-element preconditioning for pseudospectral solutions of elliptic problems. *SIAM Journal on Scientific and Statistical Computing*, 11(2):311–342, 1990.

- [58] J.J. Heys, T.A. Manteuffel, S.F. McCormick, and L.N. Olson. Algebraic multigrid for higher-order finite elements. *Journal of Computational Physics*, 204(2):520 – 532, 2005.
- [59] K. Stben. A review of algebraic multigrid. *Journal of Computational and Applied Mathematics*, 128(12):281 – 309, 2001. Numerical Analysis 2000. Vol. VII: Partial Differential Equations.
- [60] Rahul S. Sampath and George Biros. A parallel geometric multigrid method for finite elements on octree meshes. *SIAM Journal on Scientific Computing*, 32(3):1361 – 1392, 2010.
- [61] S. P. Hirshman and J. C. Whitson. Steepest-descent moment method for three-dimensional magnetohydrodynamic equilibria. *Physics of Fluids*, 26(12):3553–3568, 1983.
- [62] A Reiman, G Fu, S Hirshman, L Ku, D Monticello, H Mynick, M Redi, D Spong, M Zarnstorff, B Blackwell, A Boozer, A Brooks, W A Cooper, M Drevlak, R Goldston, J Harris, M Isaev, C Kessel, Z Lin, J F Lyon, P Merkel, M Mikhailov, W Miner, G Neilson, M Okamoto, N Pomphrey, W Reiersen, R Sanchez, J Schmidt, A Subbotin, P Valanju, K Y Watanabe, R White, N Nakajima, and C Nhrenberg. Physics design of a high- β quasi-axisymmetric stellarator. *Plasma Physics and Controlled Fusion*, 41(12B):B273, 1999.
- [63] James D. Hanson, Steven P. Hirshman, Stephen F. Knowlton, Lang L. Lao, Edward A. Lazarus, and John M. Shields. V3FIT: a code for three-dimensional equilibrium reconstruction. *Nuclear Fusion*, 49(7):075031, 2009.
- [64] Allen H. Boozer Jong-kyu Park and Alan H. Glasser. Computation of three-dimensional tokamak and spherical torus equilibria. *Physics of Plasmas*, 14(5):052110, 2007.
- [65] Jong kyu Park, Allen H. Boozer, Jonathan E. Menard, Andrea M. Garofalo, Michael J. Schaffer, Richard J. Hawryluk, Stanley M. Kaye, Stefan P. Gerhardt, Steve A. Sabbagh, and NSTX Team. Importance of plasma response to nonaxisymmetric perturbations in tokamaks. *Physics of Plasmas*, 16(5):056115, 2009.
- [66] S. P. Hirshman, R. Sanchez, and C. R. Cook. SIESTA: A scalable iterative equilibrium solver for toroidal applications. *Physics of Plasmas (1994-present)*, 18(6):–, 2011.

- [67] A. Reiman and H. Greenside. Calculation of three-dimensional MHD equilibria with islands and stochastic regions. *Computer Physics Communications*, 43:157–167, December 1986.
- [68] D. Monticello M. Drevlak and A. Reiman. PIES free boundary stellarator equilibria with improved initial conditions. *Nuclear Fusion*, 45(7):731, 2005.
- [69] Allen H. Boozer. Plasma equilibrium with rational magnetic surfaces. *Physics of Fluids (1958-1988)*, 24(11):1999–2003, 1981.
- [70] W. E. Arnoldi. The principle of minimized iterations in the solution of the matrix eigenvalue problem. *Q. Appl. Math.*, 9(17):17–29, 1951.
- [71] RB Lehoucq, DC Sorensen, and C Yang. ARPACK Users Guide: Solution of Large-Scale Eigenvalue Problems with Implicitly Restarted Arnoldi Methods (SIAM, Philadelphia, 1998). *The software and this manual are available at URL <http://www.caam.rice.edu/software/ARPACK>*, 1989.
- [72] M. Bcoulet, E. Nardon, G. Huysmans, W. Zwingmann, P. Thomas, M. Lipa, R. Moyer, T. Evans, V. Chuyanov, Y. Gribov, A. Polevoi, G. Vayakis, G. Federici, G. Saibene, A. Portone, A. Loarte, C. Doebert, C. Gimblett, J. Hastie, and V. Parail. Numerical study of the resonant magnetic perturbations for Type I edge localized modes control in ITER. *Nuclear Fusion*, 48(2):024003, 2008.
- [73] Y. Feng, F. Sardei, J. Kisslinger, P. Grigull, K. McCormick, and D. Reiter. 3D edge modeling and island divertor physics. *Contributions to Plasma Physics*, 44(1-3):57–69, 2004.
- [74] R. Lorenzini, E. Martines, P. Piovesan, D. Terranova, P. Zanca, and M. Zuin. Self-organized helical equilibria as a new paradigm for ohmically heated fusion plasmas. *Nature Physics*, 5:570–574, Aug 2009.
- [75] J. S. Wrobel, C. J. Hansen, T. R. Jarboe, R. J. Smith, A. C. Hossack, B. A. Nelson, G. J. Marklin, D. A. Ennis, C. Akcay, and B. S. Victor. Relaxation-time measurement via a time-dependent helicity balance model. *Physics of Plasmas*, 20(1):012503, 2013.
- [76] D.A. Sutherland, T.R. Jarboe, K.D. Morgan, M. Pfaff, E.S. Lavine, Y. Kamikawa, M. Hughes, P. Andrist, G. Marklin, and B.A. Nelson. The Dynomak: An advanced spheromak reactor concept with imposed-dynamo current drive and next-generation nuclear power technologies. *Submitted to Fusion Engineering and Design*, 2014.

- [77] Fengyan Li and Chi-Wang Shu. Locally divergence-free discontinuous Galerkin methods for MHD equations. *Journal of Scientific Computing*, 22-23(1-3):413–442, 2005.
- [78] T. Warburton and Mark Embree. The role of the penalty in the local discontinuous Galerkin method for Maxwells eigenvalue problem. *Computer Methods in Applied Mechanics and Engineering*, 195(2528):3205 – 3223, 2006.
- [79] Wenlong Dai and Paul R. Woodward. A simple finite difference scheme for multidimensional magnetohydrodynamical equations. *Journal of Computational Physics*, 142(2):331 – 369, 1998.
- [80] Guang-Shan Jiang and Cheng chin Wu. A high-order WENO finite difference scheme for the equations of ideal magnetohydrodynamics. *Journal of Computational Physics*, 150(2):561 – 594, 1999.
- [81] Gbor Tth. The $\nabla \cdot B=0$ constraint in shock-capturing magnetohydrodynamics codes. *Journal of Computational Physics*, 161(2):605 – 652, 2000.
- [82] Youcef Saad. A flexible inner-outer preconditioned gmres algorithm, 1993.
- [83] D. D. Schnack, D. C. Barnes, D. P. Brennan, C. C. Hegna, E. Held, C. C. Kim, S. E. Kruger, A. Y. Pankin, and C. R. Sovinec. Computational modeling of fully ionized magnetized plasmas using the fluid approximation. *Physics of Plasmas*, 13(5):056110, 2006.
- [84] S. M. Mahajan and V. Krishan. Exact solution of the incompressible hall magnetohydrodynamics. *Monthly Notices of the Royal Astronomical Society: Letters*, 359(1):L27–L29, 2005.
- [85] C.R. Sovinec and J.R. King. Analysis of a mixed semi-implicit/implicit algorithm for low-frequency two-fluid plasma modeling. *Journal of Computational Physics*, 229(16):5803 – 5819, 2010.
- [86] T. Dudok de Wit, A.L. Pecquet, J.C. Vallet, and R. Lima. The biorthogonal decomposition as a tool for investigating fluctuations in plasmas. *Physics of Plasmas*, 1(10):3288–3300, 1994.
- [87] C Nardone. Multichannel fluctuation data analysis by the singular value decomposition method. application to MHD modes in JET. *Plasma Physics and Controlled Fusion*, 34(9):1447, 1992.

- [88] Bernd R. Noack, Konstantin Afanasiev, Marek Morzynski, Gilead Tadmor, and Frank Thiele. A hierarchy of low-dimensional models for the transient and post-transient cylinder wake. *Journal of Fluid Mechanics*, 497:335–363, 2003.
- [89] T. R. Jarboe. Steady inductive helicity injection and its application to a high-beta spheromak. *Fusion Science and Technology*, 36:85–91, July 1999.
- [90] P. A. Sweet. The Neutral Point Theory of Solar Flares. In B. Lehnert, editor, *Electromagnetic Phenomena in Cosmical Physics*, volume 6 of *IAU Symposium*, page 123, 1958.
- [91] E. N. Parker. Sweet’s mechanism for merging magnetic fields in conducting fluids. *Journal of Geophysical Research*, 62(4):509–520, 1957.
- [92] L. Spitzer. *Physics of fully ionized gases*. Interscience Publishers, 1962.
- [93] S. I. Braginskii. Transport Processes in a Plasma. *Reviews of Plasma Physics*, 1:205, 1965.
- [94] E. T. Meier and U. Shumlak. A general nonlinear fluid model for reacting plasma-neutral mixtures. *Physics of Plasmas (1994-present)*, 19(7):–, 2012.
- [95] SymPy Development Team. *SymPy: Python library for symbolic mathematics*, 2013.

Appendix A

MESH GENERATION INTERFACES

A.1 T3D

T3D[33] is a licensed mesh generation program written by Dr. Daniel Rypl. It uses a proprietary but simple format to define CAD geometry in terms of rational Beziér curves and surfaces. PSI-TET provides an internal interface for parsing these input files and reconstructing the CAD geometry for use in geometric refinement and high order geometric mappings. Rational Beziér curves, equation A.2, are defined by a set of control points (\vec{x}_i) and weights (ω_i) which multiply Bernstein polynomials (eq. A.1) in the parametric coordinate (u). Rational Beziér curves are commonly used in computer modeling due to their ability to exactly represent conic sections and other useful shapes. An example of a second order Beziér curve and surface are shown in figure A.1.

$$B_i^j(t) = \binom{j}{i} t^i (1-t)^{j-i}, \quad i = 0 \dots j \quad (\text{A.1})$$

$$\vec{r}(u) = \frac{\sum_{i=0}^n \vec{x}_i \omega_i B_i^n(u)}{\sum_{i=0}^n \omega_i B_i^n(u)} \quad (\text{A.2})$$

Rational Beziér surfaces, equation A.3, are defined by combining the representations of their boundary curves, with orthogonal parametric coordinates(u, v), along with additional control points and weights for the interior of the surface if necessary.

$$\vec{r}(u, v) = \frac{\sum_{i=0}^n \sum_{j=0}^m \vec{x}_{ij} \omega_{ij} B_i^n(u) B_j^m(v)}{\sum_{i=0}^n \sum_{j=0}^m \omega_{ij} B_i^n(u) B_j^m(v)} \quad (\text{A.3})$$

In PSI-TET curve and surface reconstruction up to 3rd order Beziér curves and




Figure A.1: Example of second order Beziér curve(a) and surface(b) showing the control points.

surfaces is supported, however due to input file availability only 2nd order surfaces have been tested. New points are located on reconstructed curves and surfaces by minimizing the weighted distance between specified constraint points. This method has been found to be robust even with degenerate and complex surfaces.

A.1.1 HIT-SI Mesh

The mesh definition used for time dependent MHD runs in chapter 5 is given below. The CAD definition is slightly simplified from the full geometry by removing the diagnostic gap and small features on the injector annuli. Meshes may be generated from this definition using the T3D software as described in the PSI-TET documentation. To reduce the size of the CAD geometry file only 1/4 of the domain is specified and meshed. The full domain is then generated in PSI-TET through reflection across the x-z and y-z planes.

Injector Jump Planes

Injector jump planes are defined for the “hitsi.rs” mesh using the PSI-TET circular cut plane convention as follows.

X-injector:

Flux circuit

$$\vec{r}_{hpc} = (0., 0., -0.6)$$

$$\vec{r}_{hcv} = (-0.5, 0., 0.)$$

Voltage circuit

$$\vec{r}_{hpc} = (0., 0., -0.3)$$

$$\vec{r}_{hcv} = (0., 0.3, 0.)$$

Y-injector:

Flux circuit

$$\vec{r}_{hpc} = (0., 0., 0.6)$$

$$\vec{r}_{hcv} = (0., -0.5, 0.)$$

Voltage circuit

$$\vec{r}_{hpc} = (0., 0., 0.3)$$

$$\vec{r}_{hcv} = (0.3, 0., 0.)$$

Input File

#####

hitsi_rs

#

HIT-SI geometry without the diagnostic gap

#####

vertex 1 xyz .0 .0 -.04034

```

vertex 2 xyz .01796 .0 -.04778
vertex 4 xyz .26691 .0 -.29673
vertex 5 xyz .39247 .0 -.29673
vertex 9 xyz .53942 .0 .0
vertex 10 xyz .0 .0 .0

vertex 11 fixed vertex 1
vertex 12 xyz .0 .01796 -.04778
vertex 14 xyz .0 .26691 -.29673
vertex 15 xyz .0 .39247 -.29673
vertex 19 xyz .0 .53942 .0

vertex 43 xyz .29189 .09293 -.29673
vertex 45 xyz .36995 .09293 -.29673

vertex 70 xyz .0 .0 -.55513
vertex 73 xyz .0 .09293 -.58862
vertex 75 xyz .0 .09293 -.66668
vertex 78 xyz .0 .0 -.70016

curve 1 order 3 vertex 1 2 output yes size def * .5
polygon 1 xyz .01052 .0 -.04034 weight 0.923879533
curve 2 vertex 2 4 output yes size def
curve 4 vertex 4 5 output yes size def
curve 5 vertex 5 9 output yes size def
curve 9 vertex 9 10 output yes size def
curve 10 vertex 10 1 output yes size def * .5

```

```

curve 11 order 3 vertex 11 12 output yes size def * .5
polygon 1 xyz .0 .01052 -.04034 weight 0.923879533
curve 12 vertex 12 14 output yes size def
curve 14 vertex 14 15 output yes size def
curve 15 vertex 15 19 output yes size def
curve 19 vertex 19 10 output yes size def

curve 21 order 3 vertex 1 11 output yes size def
polygon 1 xyz .0 .0 -.04034 weight 0.70710678118655
curve 22 order 3 vertex 2 12 output yes size def
polygon 1 xyz .01796 .01796 -.04778 weight 0.70710678118655
curve 24 order 3 vertex 4 14 output yes size def
polygon 1 xyz .26691 .25840 -.29673 weight 0.70710678118655
curve 25 order 3 vertex 5 15 output yes size def
polygon 1 xyz .39247 .40343 -.29673 weight 0.70710678118655
curve 29 order 3 vertex 9 19 output yes size def
polygon 1 xyz .53942 .55602 .0 weight 0.70710678118655

curve 42 vertex 4 43 output yes size def
curve 44 order 3 vertex 43 45 output yes size def
polygon 1 xyz .33092 .21436 -.29673 weight 0.34202014332567
curve 46 vertex 45 5 output yes size def

curve 72 vertex 70 73 output yes size def
curve 74 order 3 vertex 73 75 output yes size def
polygon 1 xyz .0 .21436 -.62765 weight 0.34202014332567
curve 76 vertex 75 78 output yes size def
curve 78 vertex 70 78 output yes size def

```

```

curve 80 order 3 vertex 4 70 output yes size def
polygon 1 xyz .26691 .0 -.55513 weight 0.70710678118655
curve 83 order 3 vertex 43 73 output yes size def
polygon 1 xyz .29189 .09293 -.58862 weight 0.70710678118655
curve 85 order 3 vertex 45 75 output yes size def
polygon 1 xyz .36995 .09293 -.66668 weight 0.70710678118655
curve 88 order 3 vertex 5 78 output yes size def
polygon 1 xyz .39247 .0 -.70016 weight 0.70710678118655

surface 31 curve 1 22 11 21 output yes size def * .5
polygon 1 1 xyz .01052 .01052 -.04034 weight 0.653281483
surface 32 curve 2 24 12 22 output yes size def
surface 35 curve 5 29 15 25 output yes size def

surface 72 curve 42 80 72 83 output yes size def
surface 74 curve 44 83 74 85 output yes size def
polygon 1 1 xyz .33092 .21436 -.62765 weight 0.24184476264798
surface 76 curve 46 85 76 88 output yes size def
surface 78 curve 4 88 78 80 output yes size def

patch 1 normal 0 -1 0 boundary curve 1 2 4 5 9 10 output \
yes size def
patch 2 normal -1 0 0 boundary curve -11 -12 -14 -15 -19 -10 \
output yes size def
patch 4 normal 0 0 -1 boundary curve -42 -44 -46 -25 14 24 \
output yes size def
patch 5 normal 0 0 -1 boundary curve 42 44 46 -4 output yes \

```

```

size def
patch 7 normal -1 0 0 boundary curve -78 72 74 76 output yes \
size def

vertex 101 xyz 0. .0 .04034
vertex 102 xyz 0. .01796 .04778
vertex 104 xyz 0. .26691 .29673
vertex 105 xyz 0. .39247 .29673
vertex 109 fixed vertex 19
vertex 110 fixed vertex 10

vertex 111 fixed vertex 101
vertex 112 xyz .01796 0. .04778
vertex 114 xyz .26691 0. .29673
vertex 115 xyz .39247 0. .29673
vertex 119 fixed vertex 9

vertex 143 xyz .09293 .29189 .29673
vertex 145 xyz .09293 .36995 .29673

vertex 170 xyz .0 0. .55513
vertex 173 xyz .09293 0. .58862
vertex 175 xyz .09293 0. .66668
vertex 178 xyz .0 0. .70016

curve 101 order 3 vertex 101 102 output yes size def * .5
polygon 1 xyz .0 .01052 .04034 weight 0.923879533
curve 102 vertex 102 104 output yes size def

```

```
curve 104 vertex 104 105 output yes size def
curve 105 vertex 105 109 output yes size def
curve 110 vertex 110 101 output yes size def * .5

curve 111 order 3 vertex 111 112 output yes size def * .5
polygon 1 xyz .01052 .0 .04034 weight 0.923879533
curve 112 vertex 112 114 output yes size def
curve 114 vertex 114 115 output yes size def
curve 115 vertex 115 119 output yes size def

curve 121 order 3 vertex 101 111 output yes size def
polygon 1 xyz .0 .0 .04034 weight 0.70710678118655
curve 122 order 3 vertex 102 112 output yes size def
polygon 1 xyz .01796 .01796 .04778 weight 0.70710678118655
curve 124 order 3 vertex 104 114 output yes size def
polygon 1 xyz .26691 .25840 .29673 weight 0.70710678118655
curve 125 order 3 vertex 105 115 output yes size def
polygon 1 xyz .39247 .40343 .29673 weight 0.70710678118655

curve 142 vertex 104 143 output yes size def
curve 144 order 3 vertex 143 145 output yes size def
polygon 1 xyz .21436 .33092 .29673 weight 0.34202014332567
curve 146 vertex 145 105 output yes size def

curve 172 vertex 170 173 output yes size def
curve 174 order 3 vertex 173 175 output yes size def
polygon 1 xyz .21436 .0 .62765 weight 0.34202014332567
curve 176 vertex 175 178 output yes size def
```

```

curve 178 vertex 170 178 output yes size def

curve 180 order 3 vertex 104 170 output yes size def
polygon 1 xyz .0 .26691 .55513 weight 0.70710678118655
curve 183 order 3 vertex 143 173 output yes size def
polygon 1 xyz .09293 .29189 .58862 weight 0.70710678118655
curve 185 order 3 vertex 145 175 output yes size def
polygon 1 xyz .09293 .36995 .66668 weight 0.70710678118655
curve 188 order 3 vertex 105 178 output yes size def
polygon 1 xyz .0 .39247 .70016 weight 0.70710678118655

surface 131 curve 101 122 111 121 output yes size def * .5
polygon 1 1 xyz .01052 .01052 .04034 weight 0.653281483
surface 132 curve 102 124 112 122 output yes size def
surface 135 curve 105 29 115 125 output yes size def

surface 172 curve 142 180 172 183 output yes size def
surface 174 curve 144 183 174 185 output yes size def
polygon 1 1 xyz .21436 .33092 .62765 weight 0.24184476264798
surface 176 curve 146 185 176 188 output yes size def
surface 178 curve 104 188 178 180 output yes size def

patch 10 normal 0 0 1 boundary curve -9 19 29 output yes size def
patch 11 normal -1 0 0 boundary curve 101 102 104 105 19 110 \
output yes size def
patch 12 normal 0 -1 0 boundary curve -111 -112 -114 -115 -9 -110 \
output yes size def
patch 14 normal 0 0 1 boundary curve -142 -144 -146 -125 114 124 \

```

```

output yes size def
patch 15 normal 0 0 1 boundary curve 142 144 146 -104 output yes \
size def
patch 17 normal 0 -1 0 boundary curve -178 172 174 176 output yes \
size def

region 1 boundary surface 72 74 76 78 boundary patch -5 7 size def

region 2 boundary surface 31 32 35 boundary patch 1 2 4 5 10 size def

region 3 boundary surface 131 132 135 boundary patch -10 11 12 14 15 \
size def

region 4 boundary surface 172 174 176 178 boundary patch -15 17 size \
def

```

A.2 CUBIT

CUBIT[34] a licensed mesh generation program developed at Sandia national labs. It is available free of charge to DOE and other federally funded research groups and to the general community via a commercial version called TRELIS. A proprietary, commercial package is used internally to represent CAD geometry and supports the import of a wide range of commonly used CAD transfer file formats. Due to additional capabilities and user friendliness this program is the preferred mesh generator for PSI-TET moving forward. As with T3D, PSI-TET provides an internal interface for reconstructing source CAD geometry for use in geometric refinement and high order geometric mappings. For CUBIT this is provided through the OpenNURBS library[48], which supplies subroutines to import the CAD representation and evaluate the component curves and surfaces. To generate the required geometry file along with

the generated mesh the "refine parallel" command is used in CUBIT, which requires version 14.0+. New points are located on the boundary using the same method as T3D meshes.

A.2.1 HIT-SI3 Mesh

A script for generating a suitable HIT-SI3 mesh using CUBIT is provided below, along with information required for running a time dependent MHD simulation with injector drive. The CAD description of the plasma volume for HIT-SI3 is provided in the file "hit-si3-plasma.sat", which was created by John Rogers. Following import the geometry is simplified from the full geometry by removing small features on the injector annuli and an unnecessary surface break in the diagnostic gap. Meshes may be generated from this definition using the CUBIT software using the script below.

Injector Jump Planes

Injector jump planes are defined for the HIT-SI3 mesh using the PSI-TET circular cut plane convention as follows. Injectors are numbered 1-3 in the positive toroidal direction with injector 1 oriented along the y-axis at 0° .

Injector 1:

Flux circuit

$$\vec{r}_{hpc} = (0., 0.27, 0.548)$$

$$\vec{r}_{hcv} = (6.9, 0., 0.)$$

Voltage circuit

$$\vec{r}_{hpc} = (0., 0.27, 0.3451)$$

$$\vec{r}_{hcv} = (0., 6.66, 0.)$$

Injector 2:*Flux circuit*

$$\vec{r}_{hpc} = (-0.2338, -0.135, 0.548)$$

$$\vec{r}_{hcv} = (-3.4500, 5.9756, 0.)$$

Voltage circuit

$$\vec{r}_{hpc} = (-0.2338, -0.135, 0.3451)$$

$$\vec{r}_{hcv} = (-5.7677, -3.3300, 0.)$$

Injector 3:*Flux circuit*

$$\vec{r}_{hpc} = (0.2338, -0.135, 0.548)$$

$$\vec{r}_{hcv} = (-3.4500, -5.9756, 0.)$$

Voltage circuit

$$\vec{r}_{hpc} = (0.2338, -0.135, 0.3451)$$

$$\vec{r}_{hcv} = (5.7677, -3.3300, 0.)$$

Construction Script

```
## Reset CUBIT environment
```

```
reset
```

```
undo off
```

```
## Import CAD geometry and orient to z-axis
```

```
import acis "hit-si3-plasma.sat" attributes_on separate_bodies
```

```
rotate Volume 1 angle 270 about X include_merged
Volume 1 scale .01

## Simplify geometry
# Remove small surfaces on upper large cone
create frustum height .24065357 radius 0.52457541 top .40342946
move Surface 56 midpoint location curve 22 include_merged

# Remove small surfaces on lower large cone
create frustum height .24065357 radius 0.52457541 top .40342946
rotate Volume 3 angle 180 about X include_merged
move Surface 59 midpoint location curve 27 include_merged

# Remove unnecessary separate surfaces in gap
create Cylinder height 5 radius .35
create Cylinder height 0.059944 radius .557
subtract volume 4 from volume 2 3 5
unite volume 1 2 3 5

# Fillet gap corners
modify curve 29 blend radius .03
modify curve 43 blend radius .03

## Create injector cut planes
# Isolate cuts to injector region
create surface rectangle width 1 zplane
move Surface 78 midpoint location curve 133 include_merged
webcut volume 1 with plane surface 78
```

```
delete Surface 78
```

```
# Flux cut plane for injector 1  
create surface rectangle width 1 xplane
```

```
# Flux cut plane for injector 2  
create surface rectangle width 1 xplane  
rotate Surface 84 angle 120 about Z include_merged
```

```
# Flux cut plane for injector 3  
create surface rectangle width 1 xplane  
rotate Surface 85 angle -120 about Z include_merged
```

```
# Generate cuts from planes  
webcut volume 1 with plane surface 83  
webcut volume 11 with plane surface 84  
webcut volume 1 with plane surface 85  
delete Surface 83 84 85  
unite volume 1 12
```

```
# Voltage cut plane for injector 1  
create surface rectangle width 1 yplane  
move Surface 144 normal to surface 144 distance .27 include_merged
```

```
# Voltage cut plane for injector 2  
create surface rectangle width 1 yplane  
rotate Surface 145 angle 120 about Z include_merged  
move Surface 145 normal to surface 145 distance .27 include_merged
```

```
# Voltage cut plane for injector 3
create surface rectangle width 1 yplane
rotate Surface 146 angle -120 about Z include_merged
move Surface 146 normal to surface 146 distance .27 include_merged

# Generate cuts from planes
webcut volume 11 13 with plane surface 144
webcut volume 1 17 with plane surface 145
webcut volume 18 19 with plane surface 146
delete Surface 144 145 146

## Generate mesh
# Join all subregions
imprint all
merge all

# Set mesh type
volume all scheme Tetmesh
set tetmesher interior points on
set tetmesher optimize level 3 optimize overconstrained off sliver off
set tetmesher boundary recovery off

# Set size parameters and generate mesh
volume all size .02
volume 7 size .02
mesh volume all
```

204

```
## Export mesh and geometry file  
refine parallel fileroot 'hit_si3' overwrite no_execute
```

Appendix B

HIGH ORDER TETRAHEDRA

B.1 Barycentric coordinates

For simplex elements — triangles in 2D and tetrahedra in 3D — it is useful to define a set of non-orthogonal coordinates that are symmetric with respect to the element vertices, barycentric coordinates. These coordinates, also called vertex areas or vertex volumes, are defined in terms of the center of mass of weights placed at each vertex. In a triangle for example, the position $(1/2, 0, 1/2)$ is the point matching the center of mass for a configuration of equal masses placed at vertices 1 and 3. In triangles and tetrahedra, the coordinate value (ξ_1) relative to a vertex V_1 for a point P can also be thought of in terms of the area or volume contained by a new element formed with point P and the other vertices $V_{i \neq 1}$ of the original element. Figure B.1 shows an example of surfaces created by keeping ξ_1 fixed in the unit tetrahedron, where V_1 is placed at the origin. The coordinates are weighted such that the sum is fixed to one, as given by equation B.1 for tetrahedra. From this equation it is clear that these coordinates are not independent, with one redundant dimension that can be defined in terms of the other three. This fact also carries through to derivatives in this coordinate system, which satisfy equation B.2. In PSI-TET, logical space is defined in terms of barycentric coordinates on the unit element to enable definition of basis functions in a symmetric fashion. In orthogonal right-handed cartesian coordinates (x_i) the first vertex is placed at the origin, with the remaining vertices placed at unit length along each of the cartesian axes ($V_{i+1} = \hat{x}_i$).

$$\xi_1 + \xi_2 + \xi_3 + \xi_4 = 1 \tag{B.1}$$




Figure B.1: Surfaces of constant ξ_1 in the unit tetrahedron.

$$\frac{\partial f}{\partial \xi_1} + \frac{\partial f}{\partial \xi_2} + \frac{\partial f}{\partial \xi_3} + \frac{\partial f}{\partial \xi_4} = 0 \quad (\text{B.2})$$

B.2 First derivatives

Using the coordinates defined above, derivatives of functions in logical space can be transformed to derivatives in physical space by applying a transformation. The derivative of a function in physical space is defined as equation B.3, when physical coordinates are the desired function depend on logical coordinates. The matrix defining this mapping is the grid Jacobian matrix, $J = \frac{\partial x_k}{\partial \xi_i}$, and can be inverted to produce a mapping between logical and physical derivatives, equation B.4. With barycentric coordinates this matrix is not square, however it can be transformed in to an equivalent square matrix by utilizing the coordinate dependence (eq. B.2) to map derivatives in one logical coordinate to the others. Once inverted, equation B.2 is used again to map the matrix back to the full coordinate space allowing use of derivatives in barycentric coordinates.

$$\frac{\partial f}{\partial \xi_i} = \frac{\partial x_k}{\partial \xi_i} \frac{\partial f}{\partial x_k} \quad (\text{B.3})$$

$$\frac{\partial f}{\partial x_k} = \left(\frac{\partial x_k}{\partial \xi_i} \right)^{-1} \frac{\partial f}{\partial \xi_i} \quad (\text{B.4})$$

B.3 Second derivatives

A similar process can be performed for second derivatives, where the mapping of derivatives in physical space is now given by equation B.5. For second derivatives the mapping involves the grid Hessian, $H = \frac{\partial^2 x_k}{\partial \xi_i \partial \xi_j}$, when the spatial mapping has curvature and the matrix of derivative couplings, $\frac{\partial x_k}{\partial \xi_j} \frac{\partial x_l}{\partial \xi_i}$. Inverting this expression produces equation B.6, which can be used to evaluate second derivatives in physical space from basis functions and the physical mapping.

$$\frac{\partial^2 f}{\partial \xi_i \partial \xi_j} = \frac{\partial x_k}{\partial \xi_j} \frac{\partial x_l}{\partial \xi_i} \frac{\partial^2 f}{\partial x_k \partial x_l} + \frac{\partial^2 x_k}{\partial \xi_i \partial \xi_j} \frac{\partial f}{\partial x_k} \quad (\text{B.5})$$

$$\frac{\partial^2 f}{\partial x_k \partial x_l} = \left(\frac{\partial x_k}{\partial \xi_j} \frac{\partial x_l}{\partial \xi_i} \right)^{-1} \left[\frac{\partial^2 f}{\partial \xi_i \partial \xi_j} - \frac{\partial^2 x_k}{\partial \xi_i \partial \xi_j} \frac{\partial f}{\partial x_k} \right] \quad (\text{B.6})$$

Appendix C

HIGH ORDER NEDELEC ELEMENTS

This appendix provides definitions for the H^0 and $H^1(\text{Curl})$ finite element basis functions derived by Schöberl and Zaglmayr[55]. Note that the definitions below do not match the referenced paper exactly as it contains several typos and an unorthodox notation that has been corrected here. The validity of the definitions has been verified by the test in section 2.2.4 as well as other test cases. These basis sets are based on scaled Legendre polynomials, defined in equation C.1, and scaled integrated Legendre polynomials, defined in equation C.2), where $P_i(x)$ is the i -th Legendre polynomial. The first five functions resulting from these definitions are presented in table C.1. When a basis set is generated the number of functions required depends on the desired polynomial degree, indicated in the definitions below as np .

$$l_s(s, t, i) = t^i P_i(s/t) \tag{C.1}$$

$$L_s(s, t, i) = t^i \int_{-1}^{s/t} P_{i-1}(x) dx \tag{C.2}$$

The basis definitions below also use the shorthand $(\nabla \xi_k)$ to represent the column of the grid Jacobian matrix corresponding to the k -th logical coordinate, equation C.3.

$$\nabla \xi_k = \frac{\partial \xi_k}{\partial x_i} \tag{C.3}$$

As the $H^1(\text{Grad})$ space is defined by taking the gradient of the H^0 space its definition is not provided here, however these functions along with the curl of the

Table C.1: First five Legendre based functions used for H^1 element construction

i	$l_s(s, t, i)$	$L_s(s, t, i)$
1	s	$s + t$
2	$\frac{1}{2}(3s^2 - t^2)$	$\frac{1}{2}(s^2 - t^2)$
3	$\frac{s}{2}(5s^2 - 3t^2)$	$\frac{s}{2}(s^2 - t^2)$
4	$\frac{1}{4}\left(\frac{35}{2}s^4 - 15s^2t^2 + \frac{3}{2}t^4\right)$	$\frac{1}{4}\left(\frac{5}{2}s^4 - 3s^2t^2 + \frac{1}{2}t^4\right)$
5	$\frac{s}{8}(63s^4 - 70s^2t^2 + 15t^4)$	$\frac{s}{8}(7s^4 - 10s^2t^2 + 3t^4)$

$H^1(\text{Curl})$ basis functions are provided by the generator script, section C.5. The later can being computed using the vector identity C.4.

$$\nabla \times (\phi \vec{A}) = \phi \nabla \times \vec{A} + \nabla \phi \times \vec{A} \quad (\text{C.4})$$

C.1 Vertex Elements

For mesh vertices only a single scalar basis function is defined in terms of its barycentric coordinate(ξ) and represents linear variation.

H^0 Scalar

$$f(\xi) = \xi \quad (\text{C.5})$$

C.2 Edge Elements

For mesh edges scalar basis functions are defined for $\{i \geq 2, i \leq np\}$. A single vector basis function for the $H^1(\text{Curl})$ space is also defined for each edge. Basis functions are dependent on the barycentric coordinates of two end points of the edge (ξ_1, ξ_2) . In

PSI-TET edges are oriented so that the global index of end point two is larger than the global index of end point one.

$$u_i(\xi_1, \xi_2) = L_s(\xi_1 - \xi_2, \xi_1 + \xi_2, i) \quad (\text{C.6})$$

H^0 Scalar

$$f_i(\xi_1, \xi_2) = u_i \quad (\text{C.7})$$

$H^1(\text{Curl})$ Vector

$$\vec{f}(\xi_1, \xi_2) = \xi_1 \nabla \xi_2 - \xi_2 \nabla \xi_1 \quad (\text{C.8})$$

C.3 Face Elements

For mesh faces scalar basis functions are defined for $\{i \geq 2, j \geq 0, i + j \leq np - 2\}$. Two types of vector basis functions for the $H^1(\text{Curl})$ space are also defined for each face. Basis functions are dependent on the barycentric coordinates of the three corner points of the face (ξ_1, ξ_2, ξ_3) . In PSI-TET faces are oriented so that the global index of the corner points increase for 1-3. Equation C.6 is also reused from the edge element definition, with the new coordinates (ξ_1, ξ_2) from the face.

$$v_j(\xi_1, \xi_2, \xi_3) = \xi_3 l_s(\xi_3 - \xi_1 - \xi_2, \xi_1 + \xi_2 + \xi_3, j) \quad (\text{C.9})$$

H^0 Scalar

$$f_{i,j}(\xi_1, \xi_2, \xi_3) = u_i v_j \quad (\text{C.10})$$

$H^1(\text{Curl})$ Vector

$f_{i,j}^1$ is defined for $\{i \geq 2, j \geq 0, i + j \leq np - 1\}$.

$$\begin{aligned}
\vec{f}_{i,j}^1(\xi_1, \xi_2, \xi_3) &= \left[v_j \frac{\partial}{\partial \xi_1} u_i - u_i \frac{\partial}{\partial \xi_1} v_j \right] \nabla \xi_1 \\
&+ \left[v_j \frac{\partial}{\partial \xi_2} u_i - u_i \frac{\partial}{\partial \xi_2} v_j \right] \nabla \xi_2 \\
&+ \left[v_j \frac{\partial}{\partial \xi_3} u_i - u_i \frac{\partial}{\partial \xi_3} v_j \right] \nabla \xi_3
\end{aligned} \tag{C.11}$$

$f_{i,j}^2$ is defined for $\{i = 2, j \geq 0, i + j \leq np - 1\}$.

$$\begin{aligned}
\vec{f}_{i,j}^2(\xi_1, \xi_2, \xi_3) &= \xi_2 v_j \nabla \xi_1 \\
&- \xi_1 v_j \nabla \xi_2
\end{aligned} \tag{C.12}$$

C.4 Cell Elements

For mesh cells scalar basis functions are defined for $\{i \geq 2, j, k \geq 0, i + j + k \leq np - 2\}$. Three types of vector basis functions for the $H^1(\text{Curl})$ space are also defined for each cell. Basis functions are dependent on the barycentric coordinates of the four corner points of the cell $(\xi_1, \xi_2, \xi_3, \xi_4)$. Cells functions do not require a globally consistent number so the local vertex ordering is used. Equations C.6 C.9 are also reused from the edge and face element definitions, with the new coordinates (ξ_1, ξ_2, ξ_3) from the cell.

$$w_k(\xi_1, \xi_2, \xi_3, \xi_4) = \xi_4 P_k(\xi_4 - \xi_1 - \xi_2 - \xi_3) \tag{C.13}$$

H^0 Scalar

$$f_{i,j,k}(\xi_1, \xi_2, \xi_3, \xi_4) = u_i v_j w_k \tag{C.14}$$

$H^1(\text{Curl})$ Vector

$f_{i,j,k}^1$ is defined for $\{i \geq 2, j, k \geq 0, i + j + k \leq np - 2\}$.

$$\begin{aligned}
\vec{f}_{i,j,k}^1(\xi_1, \xi_2, \xi_3, \xi_4) = & \left[v_j w_k \frac{\partial}{\partial \xi_1} u_i - u_i w_k \frac{\partial}{\partial \xi_1} v_j + u_i v_j \frac{\partial}{\partial \xi_1} w_k \right] \nabla \xi_1 \\
& + \left[v_j w_k \frac{\partial}{\partial \xi_2} u_i - u_i w_k \frac{\partial}{\partial \xi_2} v_j + u_i v_j \frac{\partial}{\partial \xi_2} w_k \right] \nabla \xi_2 \\
& + \left[v_j w_k \frac{\partial}{\partial \xi_3} u_i - u_i w_k \frac{\partial}{\partial \xi_3} v_j + u_i v_j \frac{\partial}{\partial \xi_3} w_k \right] \nabla \xi_3 \\
& + \left[v_j w_k \frac{\partial}{\partial \xi_4} u_i - u_i w_k \frac{\partial}{\partial \xi_4} v_j + u_i v_j \frac{\partial}{\partial \xi_4} w_k \right] \nabla \xi_4
\end{aligned} \tag{C.15}$$

$f_{i,j,k}^2$ is defined for $\{i \geq 2, j, k \geq 0, i + j + k \leq np - 2\}$.

$$\begin{aligned}
\vec{f}_{i,j,k}^2(\xi_1, \xi_2, \xi_3, \xi_4) = & \left[v_j w_k \frac{\partial}{\partial \xi_1} u_i - u_i w_k \frac{\partial}{\partial \xi_1} v_j - u_i v_j \frac{\partial}{\partial \xi_1} w_k \right] \nabla \xi_1 \\
& + \left[v_j w_k \frac{\partial}{\partial \xi_2} u_i - u_i w_k \frac{\partial}{\partial \xi_2} v_j - u_i v_j \frac{\partial}{\partial \xi_2} w_k \right] \nabla \xi_2 \\
& + \left[v_j w_k \frac{\partial}{\partial \xi_3} u_i - u_i w_k \frac{\partial}{\partial \xi_3} v_j - u_i v_j \frac{\partial}{\partial \xi_3} w_k \right] \nabla \xi_3 \\
& + \left[v_j w_k \frac{\partial}{\partial \xi_4} u_i - u_i w_k \frac{\partial}{\partial \xi_4} v_j - u_i v_j \frac{\partial}{\partial \xi_4} w_k \right] \nabla \xi_4
\end{aligned} \tag{C.16}$$

$f_{i,j,k}^3$ is defined for $\{i = 2, j, k \geq 0, i + j + k \leq np - 2\}$.

$$\begin{aligned}
\vec{f}_{i,j,k}^3(\xi_1, \xi_2, \xi_3, \xi_4) = & \xi_2 v_j w_k \nabla \xi_1 \\
& - \xi_1 v_j w_k \nabla \xi_2
\end{aligned} \tag{C.17}$$

C.5 Generation

Basis functions can be generated automatically to any order using the Python script at “src/utilities/basis_functions/nedelec.py” included with the PSI-TET source code. These basis functions can then be used to extend the polynomial order available in PSI-TET, which is currently limited to Quintic and Quartic elements for the H^0 and $H^1(\text{Curl})$ spaces respectively. This script also requires the SymPy[95] symbolic manipulation package to be installed.

The functions can also be defined using the recurrence relation for Legendre polynomials, equation C.18a with $P_0(x) = 1$, $P_1(x) = x$, and the formula relating integration to the polynomial definitions, equation C.18b.

$$(n + 1) P_{n+1}(x) = (2n + 1) P_n(x) - n P_{n-1}(x) \quad (\text{C.18a})$$

$$\int P_n(x) dx = \frac{P_{n+1}(x) - P_{n-1}(x)}{2n + 1} \quad (\text{C.18b})$$

Appendix D

3D MHD FORMULATION

Provided below are the weak form definitions of the non-linear functions and approximate Jacobians used in the time dependent MHD formulations. Weak forms are broken up by test function set. u_c and u_g are the $H^1(\text{Curl})$ and $H^1(\text{Grad})$ basis functions for the magnetic field. u_v is the vector Lagrange basis functions for the velocity field created by expanding each component in a scalar Lagrange basis. u_n and u_T are the scalar Lagrange basis sets for the density and temperature respectively. Finite element expansions are assumed for all evolved variables in the integrals. The non-linear functions are expressed as $(u^{n+1} - \frac{\Delta t}{2} F^{n+1}) \rightarrow F^+$ and $(u^n + \frac{\Delta t}{2} F^n) \rightarrow F^-$. Finally, terms marked as red in the Jacobian matrices are currently omitted.

D.1 Reduced Hall-MHD

Presented below are the weak forms of the non-linear metric function and Jacobian matrix resulting from time discretization of equations 5.1a – 5.1c.

D.1.1 Non-Linear Metric

$$\begin{aligned}
 F_{u_c}^{\pm} &= \int u_c^T \cdot \vec{B} + \frac{m_e}{ne^2} \nabla \times u_c^T \cdot \vec{J} dV \\
 &\pm \frac{\Delta t}{2} \int \nabla \times u_c^T \cdot \left[-\vec{v} \times \vec{B} + \eta \vec{J} + \frac{1}{ne} \vec{J} \times \vec{B} \right] dV
 \end{aligned} \tag{D.1a}$$

$$F_{u_g}^{\pm} = \int u_g^T \cdot \vec{B} dV \tag{D.1b}$$

$$\begin{aligned}
F_{u_v}^\pm &= \int u_v^T \cdot \vec{v} dV \\
&\pm \frac{\Delta t}{2} \int \left\{ u_v^T \cdot \left[\vec{v} \cdot \nabla \vec{v} - \frac{1}{\rho} \vec{J} \times \vec{B} \right] \right. \\
&\left. + \frac{\nu}{\rho} \nabla u_v^T : \nabla \vec{v} \right\} dV
\end{aligned} \tag{D.1c}$$

D.1.2 Jacobian

$$\begin{aligned}
J_{u_c} &= \int \left[u_c^T \cdot \delta \vec{B} + \frac{m_e}{ne^2} \nabla \times u_c^T \cdot \delta \vec{J} \right] dV \\
&- \frac{\Delta t}{2} \int \nabla \times u_c^T \cdot \left[-\delta \vec{v} \times \vec{B} - \vec{v} \times \delta \vec{B} + \eta \delta \vec{J} \right. \\
&\left. + \frac{1}{ne} \left(\delta \vec{J} \times \vec{B} + \vec{J} \times \delta \vec{B} \right) \right] dV
\end{aligned} \tag{D.2a}$$

$$J_{u_g} = \int u_g^T \cdot \delta \vec{B} dV \tag{D.2b}$$

$$\begin{aligned}
J_{u_v} &= \int u_v^T \cdot \delta \vec{v} dV \\
&+ \frac{\Delta t}{2} \int \left\{ u_v^T \cdot \left[\delta \vec{v} \cdot \nabla \vec{v} + \vec{v} \cdot \nabla \delta \vec{v} \right. \right. \\
&\left. \left. - \frac{1}{\rho} \left(\delta \vec{J} \times \vec{B} + \vec{J} \times \delta \vec{B} \right) \right] \right. \\
&\left. + \nu \nabla u_v^T : \nabla \delta \vec{v} \right\} dV
\end{aligned} \tag{D.2c}$$

D.2 Full Hall-MHD

Presented below are the weak forms of the non-linear metric function and Jacobian matrix resulting from time discretization of equations 5.20a – 5.20g.

D.2.1 Non-Linear Metric

$$\begin{aligned}
F_{u_c}^\pm &= \int \left\{ u_c^T \cdot \vec{B} + \frac{m_e}{ne^2} \nabla \times u_c^T \cdot \vec{J} \right\} dV \\
&\mp \frac{\Delta t}{2} \int \nabla \times u_c^T \cdot \left[-\vec{v} \times \vec{B} + \eta \vec{J} \right. \\
&\quad \left. + \frac{1}{ne} \left(\vec{J} \times \vec{B} - kT \nabla n \right) \right] dV
\end{aligned} \tag{D.3a}$$

$$F_{u_g}^\pm = \int u_g^T \cdot \vec{B} dV \tag{D.3b}$$

$$\begin{aligned}
F_{u_v}^\pm &= \int u_v^T \cdot \vec{v} dV \\
&\pm \frac{\Delta t}{2} \int \left\{ u_v^T \cdot \left[\vec{v} \cdot \nabla \vec{v} - \frac{1}{nm_i} \vec{J} \times \vec{B} \right. \right. \\
&\quad \left. \left. + \frac{2k}{m_i} \left(\nabla T + T \frac{\nabla n}{n} \right) \right. \right. \\
&\quad \left. \left. - \frac{\nu}{n^2 m_i} \nabla n \cdot \nabla \vec{v} \right] + \frac{1}{nm_i} \nu \nabla u_v^T : \nabla \vec{v} \right\} dV
\end{aligned} \tag{D.3c}$$

$$\begin{aligned}
F_{u_n} &= \int u_n^T n dV \pm \frac{\Delta t}{2} \int \left\{ u_n^T \left(n \nabla \cdot \vec{v} + \vec{v} \cdot \nabla n \right) \right. \\
&\quad \left. + D \nabla u_n^T \cdot \nabla n \right\} dV
\end{aligned} \tag{D.3d}$$

$$\begin{aligned}
F_{u_T} &= \int u_T^T T dV \pm \frac{\Delta t}{2} \int u_T^T \left\{ \vec{v} \cdot \nabla T + (\gamma - 1) kT \nabla \cdot \vec{v} \right. \\
&\quad \left. - (\gamma - 1) \frac{\nabla n}{n} \left[\chi_{\parallel} \hat{b}\hat{b} + \chi_{\perp} \left(I - \hat{b}\hat{b} \right) \cdot \nabla T \right] \right. \\
&\quad \left. + \frac{\gamma - 1}{2n} \left(\eta \vec{J} \cdot \vec{J} + \nu (\nabla \vec{v})^T : \nabla \vec{v} \right) \right\} \\
&\quad + (\gamma - 1) \nabla u_T^T \cdot \left[\chi_{\parallel} \hat{b}\hat{b} + \chi_{\perp} \left(I - \hat{b}\hat{b} \right) \cdot \nabla T \right] dV
\end{aligned} \tag{D.3e}$$

D.2.2 Jacobian

$$\begin{aligned}
J_{u_c} = & \int \left[u_c^T \cdot \delta \vec{B} - \frac{m_e}{ne^2} \left(\frac{\delta n}{n} \nabla \times u_c^T \cdot \vec{J} + \nabla \times u_c^T \cdot \delta \vec{J} \right) \right] dV \\
& - \frac{\Delta t}{2} \int \nabla \times u_c^T \cdot \left\{ -\delta \vec{v} \times \vec{B} - \vec{v} \times \delta \vec{B} + \eta \delta \vec{J} \right. \\
& - \frac{1}{ne} \left[\frac{\delta n}{n} (\vec{J} \times \vec{B}) - (\delta \vec{J} \times \vec{B} + \vec{J} \times \delta \vec{B}) \right. \\
& \left. \left. - k \left(\delta T \nabla n + T \nabla \delta n - T \frac{\delta n \nabla n}{n} \right) \right] \right\} dV
\end{aligned} \tag{D.4a}$$

$$J_{u_g} = \int u_g^T \cdot \delta \vec{B} dV \tag{D.4b}$$

$$\begin{aligned}
J_{u_v} = & \int u_v^T \cdot \delta \vec{v} dV \\
& + \frac{\Delta t}{2} \int \left\{ u_v^T \cdot \left[\delta \vec{v} \cdot \nabla \vec{v} + \vec{v} \cdot \nabla \delta \vec{v} + \frac{1}{n^2 m_i} \delta n \vec{J} \times \vec{B} \right. \right. \\
& - \frac{1}{n m_i} (\delta \vec{J} \times \vec{B} + \vec{J} \times \delta \vec{B}) \\
& + \frac{2k}{m_i} \left(\nabla \delta T + \delta T \frac{\nabla n}{n} + T \frac{\nabla \delta n}{n} - T \frac{\delta n \nabla n}{n} \right) \\
& \left. \left. - \frac{\nu}{n^2 m_i} \delta n \nabla u_v^T : \nabla \vec{v} + \frac{\nu}{n m_i} \nabla u_v^T : \nabla \delta \vec{v} \right\} dV
\end{aligned} \tag{D.4c}$$

$$\begin{aligned}
J_{u_n} = & \int u_n^T \delta n dV + \frac{\Delta t}{2} \int u_n^T \left(\delta n \nabla \cdot \vec{v} + n \nabla \cdot \delta \vec{v} \right. \\
& \left. + \delta \vec{v} \cdot \nabla n \vec{v} \cdot \nabla \delta n \right) + D \nabla u_n^T \cdot \nabla \delta n dV
\end{aligned} \tag{D.4d}$$

$$\begin{aligned}
J_{u_T} = & \int u_T^T \delta T \, dV + \frac{\Delta t}{2} \int u_T^T \left\{ \delta \vec{v} \cdot \nabla T + \vec{v} \cdot \nabla \delta T \right. \\
& + (\gamma - 1) \left(k \delta T \nabla \cdot \vec{v} + k T \nabla \cdot \delta \vec{v} \right) \\
& - (\gamma - 1) \left(\frac{\nabla \delta n}{n} - \frac{\nabla n}{n^2} \delta n \right) \left[\chi_{\parallel} \hat{b} \hat{b} + \chi_{\perp} \left(I - \hat{b} \hat{b} \right) \cdot \nabla T \right] \\
& + \frac{\gamma - 1}{n} \left(\frac{\delta n}{2n} \eta \vec{J} \cdot \vec{J} + \eta \vec{J} \cdot \delta \vec{J} \right. \\
& \left. + \frac{\delta n}{2n} \nu (\nabla \vec{v})^T : \nabla \vec{v} + \nu (\nabla \vec{v})^T : \nabla \delta \vec{v} \right\} \\
& + (\gamma - 1) \nabla u_T^T \cdot \left[\chi_{\parallel} \hat{b} \hat{b} + \chi_{\perp} \left(I - \hat{b} \hat{b} \right) \cdot \nabla \delta T \right] dV
\end{aligned} \tag{D.4e}$$

VITA

Christopher James Hansen

Bachelor of Science in Aeronautics & Astronautics

University of Washington

June 2009

Doctor of Philosophy

University of Washington

March 2014