To the Graduate Council:
I am submitting herewith a dissertation written by Thomas Looby entitled "Tokamak 3D Heat Load Investigations Using an Integrated Simulation Framework: HEAT." I have examined the final paper copy of this dissertation for form and content and recommend that it be accepted in partial fulfillment of the requirements for the degree of Doctor of Philosophy, with a major in Nuclear Engineering.

David Donovan, Major Professor

We have read this dissertation
and recommend its acceptance:

Matthew Reinke

Andreas Wingen

Maik Lang

Accepted for the Council:

John Zomchick

Vice Provost and Dean of the Graduate School

To the Graduate Council:
I am submitting herewith a dissertation written by Thomas Looby entitled "Tokamak 3D Heat Load Investigations Using an Integrated Simulation Framework: HEAT." I have examined the final electronic copy of this dissertation for form and content and recommend that it be accepted in partial fulfillment of the requirements for the degree of Doctor of Philosophy, with a major in Nuclear Engineering.

David Donovan, Major Professor

We have read this dissertation
and recommend its acceptance:

Matthew Reinke

Andreas Wingen

Maik Lang

Accepted for the Council:

John Zomchick

Vice Provost and Dean of the Graduate School

(Original signatures are on file with official student records.)

Tokamak 3D Heat Load Investigations Using an Integrated Simulation Framework: HEAT

A Dissertation Presented for the

Doctor of Philosophy

Degree

The University of Tennessee, Knoxville

Thomas Looby

May 2022

© by Thomas Looby, 2022
All Rights Reserved.

Abstract

Reactor class nuclear fusion tokamaks will be inherently complex. Thousands of interconnected systems that span orders of magnitude in physical scale must operate cohesively for the machine to function. Because these reactor class tokamaks are all in an early design stage, it is difficult to quantify exactly how each subsystem will act within the context of the greater systems. Therefore, to predict the engineering parameters necessary to design the machine, simulation frameworks that can model individual systems as well as the interfaced systems are necessary. This dissertation outlines a novel framework developed to couple otherwise disparate computational domains together into a single integrated package for the goal of high fidelity 3D heat load predictions. The framework, called the Heat flux Engineering Analysis Toolkit (HEAT), bridges the gap between plasma physics, engineering, visualization, high performance computing, and more. It is open source and has been used for time varying 3D heat load predictions on 5 tokamaks. The incredible heat loads that will be present in reactor class tokamaks can easily melt the plasma facing components (PFCs) if not properly managed, which can limit performance or even damage the machine. Because these PFCs provide the interface between dozens of reactor systems and the plasma, are expected to operate extremely close to their material limits, employ complicated 3D geometry to survive the hostile conditions inside a tokamak, and can be optimized via engineering design, they are the perfect investigative candidate for an integrated modeling framework such as HEAT. After providing contextual physics and background information, the HEAT architecture will be outlined in detail. An overview of the major investigations to date will be provided, including optical and gyro orbit heat loads on several tokamaks. The novel results that can only be obtained through an integrated simulation framework will be emphasized throughout.

Table of Contents

1 Background Physics **1**

1.1 Magnetic Confinement Fusion 1

1.1.1 Tokamaks 2

1.1.2 The Tokamak Divertor 4

1.2 Models of SOL Heat Flux 6

1.2.1 The 1D NSTX-U Working Group Model 6

1.2.2 The Heuristic Model 8

1.2.3 The Eich Model 9

1.2.4 The Brunner Model 12

1.3 Charged Particle Trajectories in the SOL 14

1.3.1 The Optical Approximation 14

1.3.2 Gaps and Edge Loading 16

1.3.3 The Gyro Orbit Approximation 18

1.3.4 Validity of Neglecting Electric Fields 20

1.3.5 Kinetic Theory Distributions 25

1.4 2D Magnetic Equilibria 27

1.4.1 Coordinate Transformations 27

1.4.2 Flux Surfaces 28

1.4.3 EFIT and GEQDSK Files 31

1.4.4 Power Sharing Between Divertors 35

1.5 Other Related Topics 38

1.5.1 Coil Misalignments and Error Fields 38

1.5.2 Edge Localized Modes 38

1.5.3 0D Heat Flux Approximation 40

1.5.4 0D Heat Flux Width . . . 44
1.5.5 Detached Operation . . . 45

2 Modern PFC Heat Loads and Management 50
2.1 Tokamak PFCs Interface to Many Systems . . . 50
2.2 Expected Tokamak Divertor Heat Loads . . . 51
2.3 Consequences of Poor PFC Heat Management . . . 53
2.4 Managing Divertor Heat Loads . . . 59
2.5 The NSTX-U Divertor . . . 62

3 The Heat flux Engineering Analysis Toolkit (HEAT) 65
3.1 PFC Heat Load Simulation Requirements . . . 65
3.1.1 Software Interfaces . . . 65
3.1.2 Simulation Objectives . . . 68
3.2 Potential Candidate Codes . . . 70
3.3 HEAT Philosophy . . . 70
3.4 HEAT Architecture . . . 72
3.4.1 Concept Architecture . . . 72
3.4.2 Software Architecture . . . 74
3.4.3 Input Files . . . 76
3.5 CAD and Meshes . . . 77
3.5.1 FreeCAD . . . 78
3.5.2 Mesh Algorithms . . . 78
3.5.3 HEAT CAD and PFC Architecture . . . 79
3.6 MHDclass in HEAT . . . 80
3.7 Finding Magnetic Shadows . . . 82
3.7.1 Intersection Checking Algorithms . . . 86
3.7.2 Acceleration Structures . . . 95
3.8 Optical Approximation . . . 100
3.8.1 Optical Heat Flux Calculation . . . 100
3.8.2 Power Balance . . . 102
3.8.3 HEAT Optical Approximation Speed Benchmarking . . . 104
3.9 Gyro-orbit Algorithms . . . 109

3.9.1 Gyro-orbit Macro-particles . . . 110
3.9.2 Sampling the Distribution Functions . . . 112
3.9.3 Macro-particle Power Fractions . . . 119
3.9.4 Gyro Source Planes . . . 122
3.10 Temperature Calculation . . . 124
3.11 HEAT Roadmap . . . 125
3.11.1 Upcoming Physics Modules . . . 126
3.11.2 HEAT in Machine Learning Loops . . . 127

4 HEAT Optical Approximation Results for NSTX-U 130
4.1 Full Lower Divertor Static Discharge . . . 130
4.2 Dynamic Strike Point Sweeping . . . 135
4.3 NSTX-U Discharge 204118 . . . 138
4.4 Using HEAT During NSTX-U Construction . . . 141
4.4.1 OutBoard Divertor Missing Fish-scale . . . 142
4.4.2 OutBoard Divertor Installation Misalignment . . . 144

5 HEAT Gyro Orbit Approximation Results for NSTX-U 148
5.1 Comparison to Existing Research . . . 148
5.2 Power Balance for Gyro Orbits . . . 157
5.3 Convergence Studies . . . 157
5.3.1 Resolution Convergence Studies . . . 158
5.3.2 Power Balance Convergence Studies . . . 161
5.4 NSTX-U Castellation Gyro Orbit Analysis . . . 164
5.5 NSTX-U IBDH Gyro Orbit Predictions . . . 176

6 HEAT Results for Other Tokamaks 179
6.1 ST40 . . . 179
6.2 DIII-D Negative Triangularity . . . 183
6.3 DIII-D DiMES Head . . . 191
6.4 DIII-D Moveable Divertor . . . 194
6.5 STEP / DEMO . . . 194
6.6 SPARC . . . 199

7 Conclusions **200**

7.1 Summary of the Problem 200

7.2 Summary of the Tool 201

7.3 Summary of the Results 202

List of Tables

1.1 Order of GEQDSK variables in text file 33
1.2 GEQDSK file variable names, descriptions, units, and sizes 34

3.1 Speed benchmark for Möller-Trumbore (MT) vs Signed Volume (SV) algorithms on single CPU core with 32 GB of Random Access Memory (RAM). N is number of ROI mesh rays and Nt is number of potential intersection mesh triangle elements. *Matrix* indicates that algorithm performs matrix multiplication with no for loops (aside from those inside numpy). RAM saturates for full matrix implementation of SV, 'SV Matrix'. *Hybrid* indicates that the algorithm collapses one dimension of 3D matrix multiplication into a for loop, to avoid saturating RAM. Note that in HEAT implementation of these algorithms, Hybrid versions will pass for loops across as many cores as possible. 94
3.2 Time benchmarking parameters for a single timestep heat flux calculation on the IBDH and OBD tiles. Maximum edge length is parameter used to define STL mesh, and is not the same as resolution. No. of Ray-Tri Calcs is total number of ray-triangle intersection calculations for all steps up field line. Simulation time includes postprocessing and figure generation. Benchmark performed on a simple 8 logical core (threads) desktop workstation with 32 GB of RAM. 106

List of Figures

1.1 ITER tokamak and associated engineering systems ©ITER Organization https://www.iter.org. 3

1.2 Magnetic field components and plasma current on a (a) single flux surface, and (b) nested toroidal flux surfaces from [4]. 3

1.3 Example SOL and open field lines terminating on machine wall from [4]. 5

1.4 Example X point and divertor targets from [4]. 5

1.5 Equation 1.20 overlaid onto experimental data from [15]. 10

1.6 Scaling of heat flux width as a function of B_{POL} from [15]. 10

1.7 Eich vs Brunner heat flux profiles and Alcator C-Mod diagnostic data from [16] 13

1.8 Extension of figure 1.6 using new C-Mod data and Brunner heat flux profile from [16]. This scaling puts the heat flux widths for ITER and SPARC well below 1mm. 13

1.9 Example magnetic field line (yellow) and helical gyro orbit trajectory (green). The optical approximation assumes particles follow the curved magnetic field lines, while the gyro orbit approximation assumes particles precess about the field lines with finite Larmor radii. 15

1.10 Example of monoblock beveled surfaced from [17]. Shaded regions represent allowable alignment engineering tolerances. 15

1.11 Example of gyration about magnetic field line (blue), and potential for heat loading in tile gaps from [17]. 17

1.12 Example of discrepancy between optical approximation (blue) and gyro-orbit approximation (red), for the inter-ELM and ELM heat fluxes on ITER vertical tile monoblocks from [17]. The plasma facing surface is in the region $0.0 < S_{pol} < 12.0$ and the region outside this domain represents the gaps. 17
1.13 ITER full field ELM ion trajectory with and without electric fields, as well as potential and density, from 1D 3V PIC simulations [17, 25] . . . 21
1.14 Peak temperature difference between including sheath electric fields (PIC) and ignoring electric fields (Ion orbit) from [18]. Including electric fields results in a 10% reduction in peak temperature, implying that using the gyro orbit (ion orbit) approximation is more conservative from an engineering persepective. Results for a steady-state or slow-transient scenario. 21
1.15 Example geometry from PIC simulations performed by Komm [18]. Geometry corresponds to an ITER monoblock that has been misaligned to expose a leading edge. 23
1.16 Comparison of heat flux profiles for optical, ion orbit (from [24]), and SPICE2 PIC code from [18]. Simulations performed on portruding leading edge geometry, where $s_{tor} < 0$ corresponds to the leading edge and $s_{tor} > 0$ corresponds to the top surface. Including ion orbits transfers power from leading edge to PFC surface. Including electric fields enhances this effect. In both cases, power redistribution area is on the order of millimeters wide. 23
1.17 Magnetic equilibrium reconstruction for NSTX-U discharge 204118 at 113ms. Contours of constant pressure, or poloidal flux, can be seen to be nested inside the core. The last closed flux surface (LCFS) overlaid in red. Simplified wall geometry overlaid in green. 30
1.18 Upper/Lower power sharing as a function of δR_{sep} from [33]. 36
1.19 Divertor power sharing as a function of δR_{sep} from [34]. 36

1.20 Left: Poincaré plot and Right: lower horizontal divertor footprint due to error fields in toroidal field coil and poloidal field coil 5 from [36]. Colorbar is normalized poloidal flux. Aspect ratio is not to scale, as figure has been copied and pasted from [36]. 39
1.21 Main plasma ELM energy loss as a function of plasma current for unmitigated ELMs from [39]. 41
1.22 Simple model of SOL adapted from [4]. 42
1.23 Qualitative sketch of detachment levels from [43]. 48
1.24 Divertor heat flux before and after entering detached operation [41]. A reduction in the peak heat flux can be observed, which is critical for future reactor class devices. 49

2.1 Cartoon of some of the systems and physics domains that interface to the PFCs. Physics domains are on the left, engineering domains are on the right. The PFCs provide the interface between many subsystems in the tokamak. 52
2.2 ITER divertor heat flux predictions for axisymmetric case (top), and accounting for 3D PFC geometry (bottom) from [8]. 54
2.3 Expected heat flux in SPARC from [9] assuming 50% radiated power. . 55
2.4 Example of melting tungsten from an intentionally exposed PFC in JET, due to ELM heat loads at near normal angle of incidence [37]. Top frame is unmelted PFC. Increasing shot numbers from top to bottom show the evolution of the melted area. 57
2.5 Example of melting tungsten from an unintended PFC misalignment due to halo current forces from ASDEX-U [46]. Right PFC also has cracks. 58
2.6 Simple cartoon of flux contraction and flux expansion. Purple lines are flux surfaces. If distance between flux surfaces decreases while approaching the divertor, it is considered flux contraction. If distance increases, it is considered flux expansion. Flux expansion distributes the heat load accross a greater surface area. 60
2.7 Simple cartoon of some PFC shapes utilized to manage heat fluxes from [28]. 61

2.8 NSTX-U CAD and abbreviated PFC names: Center Stack Angled Surface (CSAS), Outboard Divertor (OBD), InBoard Divertor Horizontal (IBDH), InBoard Divertor Vertical (IBDV). OBD tile is highlighted in green and shown in exploded side view. Picture of real OBD graphite prototype with castellations and fish-scales as viewed from the front (from center stack). Conical surfaces are locations where heat flux faceting will occur. Figure taken from [28]. 63

2.9 Simulated temperature at 5s discharge length for NSTX-U castellation row from IBDH tile with 8 $\mathrm{MW\,m^{-2}}$ applied uniformly at the PFC surface. NSTX-U graphite engineering limit is 1600°. 63

3.1 Cartoon of some software packages that could be included in a high fidelity heat load prediction. Creating a framework that is capable of interfacing directly to new code packages is desireable, as it enable modular development and does not limit the physics that can be included in the heat load calculations. Codes are from traditionally disparate computational domains, such as computational physics, visualization, engineering solvers, and machine learning. 67

3.2 Tabular comparison between SMITER and HEAT. List is not all-inclusive, and is only valid at current time of publication (Fall 2021). . 71

3.3 HEAT concept flow chart. Note that each box may contain multiple python classes. Interface for 3D plasma magnetic equilibria using M3DC1 / SIESTA is currently under construction so it has been marked as partially implemented. Figure from [28] 73

3.4 HEAT software architecture flow chart. Each box represents a different software module, python class, or input file. Note that some python classes have been omitted for brevity. Third party code such as MAFOT or openFOAM are not included in this flow chart. 73

3.5 Cartoon example of how powerDir is assigned to PFC surfaces. Variable powerDir is a scalar that defines which toroidal direction (ie forward or reverse) power would have to flow to strike a PFC surface, and is defined in equation 3.2 as powerDir $= \mathrm{sign}(\hat{b} \cdot \hat{n}) \times \mathrm{sign}(B_{t0}) \times (-1.0)$ 83

3.6 Shadowed faces are calculated using MAFOT field line traces for a PFC tile in NSTX-U lower divertor. Field line traces are launched from each target mesh face, and if field line intersects another upstream face then heat flux is 0. Left image shows view looking down at divertor from core with ShadowMask=1 when a face is shadowed. Right figure shows close up view of poloidal running gap between two PFC tiles, field line launch point, intersection mesh face, and field line trace. Figure from [28]. . . 85

3.7 Coordinate transformation for Möller-Trumbore ray-triangle intersection checking algorithm from [64]. Note that here $\mathbf{v}_i = \vec{p}_i$ in the discussion. . 91

3.8 Cartoon of phi filter algorithm. Triangles in mesh are filtered according to their toroidal location. Triangles whose vertices are completely outside of the trace step are rejected, whereas triangles who have vertices within the trace step or straddling the trace step are in the passband. Intersection checking will only be performed on triangles in the passband, speeding up the intersection calculation. 96

3.9 Speed comparison for toroidal angle filtering vs no filter. Results for fixed number of intersection mesh elements (68094), with changing ROI mesh Max Edge Length (5mm,4mm,3mm,2mm). Speedup is approximately 6X for all cases. Results shown for NSTX-U single IBDH tile in ROI, 6 tiles in Intersects, and 10 degree trace length. 8 CPU cores with 32GB RAM. 97

3.10 Speed comparison for toroidal angle filtering vs no filter. Results for fixed number of ROI mesh elements (22628), with changing Intersect mesh Max Edge Length (100mm,25mm,10mm,5mm). Speedup increasing with increasing number of potential intersects, and ranges from 2.6X to 11.3X. Increasing number of intersection CAD parts will not significantly increase calculation time, due to toroidal filter. Results shown for NSTX-U single IBDH tile in ROI, 6 tiles in Intersects, and 10 degree trace length. 8 CPU cores with 32GB RAM. 98

3.11 Length of intersection trace versus error % (from equation 3.29) and simulation time, for the lower divertor during NSTX-U discharge 204118 at 1004 ms. Increasing the length of the intersection trace leads to more accurate predictions, at the cost of simulation time. Test performed on an 8 logical core workstation machine. Mesh resolution of 5 mm with intersection resolution of standard. 108

3.12 Guiding center trajectory (blue), particle trajectory (red), and initial particle location (green dot), for two examples with initial gyro phase angle of $\alpha = 0$ (left) and $\alpha = 3\pi/2$ (center). Magnetic field is tangent to circular purple plane at t=0, which has radius equal to 1 gyroradius. The particle can originate anywhere on the circumference of the plane. Gyro phase angles are uniformly distributed from $(0, 2\pi)$, and user selects number of samples, NgP. Right frame shows example samples for NgP=5. 114

3.13 Velocity phase angles (β) represent the energy sharing between perpendicular and parallel components of velocity. Example shown for NvP=4, corresponding to 4 samples uniformly distributed through velocity phase angle space. Note that $\beta = 90°$ and $\beta = 0°$ are not included as they represent the optical approximation and cyclotron motion, respectively. 115

3.14 Example Maxwellian speed distribution with velocity slice samples overlaid on top. Corresponding cumulative distribution function (CDF) on bottom. For this example NvS=3, corresponding to three samples. Red dotted lines represent the boundaries of numerical bins. Bins are plotted as a function of velocity for PDF (upper) plot, and as a function of the CDF for CDF (lower) plot. Bins are chosen such that each bin is equally probable, and velocity samples are placed in center of bins. See text for more discussion on CDF, binning, and sample locations. 115

3.15 Comparison of velocity Maxwellian PDF and CDF for three examples cases with varying number of velocity samples (vSlices). Numerical bin boundaries are shown in red dotted vertical and horizontal lines. Bins are chosen such that each bin is equally probable, and velocity samples are placed in center of bins. 118

3.16 Example of a gyro source plane. Object in light gray at bottom left is PFC of interest in the ROI. Magenta plane in upper right is gyro source plane. Magnetic field line trace and helical gyro orbit trajectory are launched from gyro source plane and strike the PFC of interest. 123
3.17 Example HEAT simulation loops to train machine learning algorithms. Blue loop provides design optimization to change CAD geometry. Red loop creates reduced model for machine control system. 129

4.1 Lower divertor 360° heat flux footprint and magnetic equilibrium from PFC Working Group Memo 010 Case 1.1. Note the various toroidal mode numbers associated with magnetic shadows, which correspond to periodic geometry patterns. Figure from [28]. 131
4.2 Simulated heat flux (top left), temperature at end of discharge (top right), and thermal probe trace (bottom), for Memo 010 Case 1.1. Trace indicates both PFCs exceed the sublimation engineering limit for graphite (1873 K). Figure from [28]. 133
4.3 View of lower divertor (from above) of time evolving temperature profile for 5 second strike point sweep with sweep frequency of 0.1 Hz. Sweep starts on upper tile (OBD), moves with constant velocity, then ends on lower tile (IBDH). Areas in dark red are above the graphite sublimation temperature, 1873 K. 136
4.4 Peak temperature on IBDH tile for various strike point sweep frequencies. PFC is set to 300K at t=0, but t=0 is not plotted. Due to 5s discharge length, 0.1Hz sweep is in a single direction across tile surface, while 1.0Hz and 10.0Hz are oscillatory. Engineering limit temperature for sublimation (Limit) is reached at approximately 0.92s for 0.1Hz, 0.55s for 1.0Hz, and 3.36s for 10.0Hz. 137
4.5 Limiter heat flux and magnetic equilibrium reconstruction during startup for NSTX-U discharge 204118 at 113 ms. User defined heat flux widths and power sharing are provided in figure. Magnetic shadowing creates toroidal asymmetry at the tangent point. Peak heat flux is $\approx$ 2.6 MW m^{-2}. 139

4.6 Lower outer divertor heat flux and magnetic equilibrium reconstruction for NSTX-U discharge 204118 at 1004 ms. User defined heat flux widths and power sharing are provided in figure. Peak heat flux is ≈ 24 MW m^{-2}. Toroidal running gap that was originally designed to be shadowed gets loaded with heat fluxes up to 5.8 MW m^{-2}. 140

4.7 Comparison geometry and associated heat fluxes for 2 Cases. Case 1 (left frames) represents the intended CAD design. Case 2 (right frames) represents the unintentional CAD design that was originally sent to PFC manufacturer. Top frames illustrate the geometry difference - different step sizes for the intercastellation fish-scale between the left and center castellation columns. Bottom frames illustrate the heat fluxes associated with each fish-scale step size. Increasing the step size by 0.4318 mm decreases the peak heat flux and increases the discharge duration by 25%. HEAT was used to find this unintentional CAD design error. CAD was revised and updated with the manufacturer before PFCs were fabricated. 143

4.8 Images demonstrating the NSTX-U OutBoard Divertor (OBD) intended design (left frame) and as built installation (center frame). In the as built installation, the OBD has been shifted vertically towards the plasma by approximately 0.5 inches, exposing the vertical surface to incident heat fluxes. Right frame shows a potential hazardous magnetic equilibrium for this misalignment, where the strike point is directly on the vertical face. 146

4.9 HEAT heat flux predictions for the misaligned OBD geometry as described in figure 4.8. The vertical edge that is no longer shadowed due to the misalignment sees heat fluxes in the 5-10 MW m^{-2} range (top frame). The leading edge sees heat fluxes of nearly 250 MW m^{-2}, although these have a width of only 100 um. The leading edges were also present in the intended PFC design, but the vertical height has been increased due to the misalignment. 147

5.1 Example poloidal (black) and toroidal (red) chords taken across an ITER tungsten monoblock from [17]. 149
5.2 Inter-ELM ion heat flux for optical (blue) and gyro-orbit model (red) from Gunn [17]. 149
5.3 Example geometry from PIC simulations performed by Komm [18]. Geometry corresponds to an ITER monoblock that has been misaligned to expose a leading edge. 150
5.4 Comparison of heat flux profiles for optical, ion orbit (from [24]), and SPICE2 PIC code from [18]. Simulations performed on portruding leading edge geometry, where $s_{tor} < 0$ corresponds to the leading edge and $s_{tor} > 0$ corresponds to the top surface. Including ion orbits transfers power from leading edge to PFC surface. Including electric fields enhances this effect. In both cases, power redistribution area is on the order of millimeters wide. 150
5.5 Top image: CAD representation of test case for comparison to ITER monoblock simulations by Gunn and Komm. PFC of interest is located in the lower NSTX-U divertor, aligned to a castellation in the IBDH tile. The PFC of interest portrudes from the toroidally upstream PFC by 0.3 mm, and is perfectly aligned to adjacent poloidal PFCs with 0.5 mm gaps. PFC of interest extends 10 mm toroidally and 5 mm poloidally. Lower image: HEAT calculates heat fluxes across entire surface, and then mesh heat flux values within 0.25 mm threshold from PFC centerlines are projected to the S_{pol} and S_{tor} planes for plotting. Mesh values to be included for toroidal chord is shown in green, poloidal chord in magenta. 152

5.6 Example HEAT generated poloidal (top) and toroidal (bottom) heat flux predictions plotted along PFC centerlines. Blue trace is ion gyro-orbit approximation heat flux, green trace is optical approximation. Both the toroidal and poloidal profiles show good agreement to the results from Gunn [17] and Komm [18]. For poloidal case, integral of gyro orbit curve is 0.0494 MW m and integral of optical curve is 0.0510 MW m. For toroidal case, integral of gyro orbit curve is 0.1399 MW m and integral of optical curve is 0.1572 MW m. Mesh elements within 0.5mm of chord are plotted. 154

5.7 Gyro orbit (top row) and optical (bottom row) heat fluxes calculated using HEAT to generate figure 5.6. Left column shows outboard side of PFC, right shows inboard side. The colorbar has been capped at 20 MW m^{-2} to show the gradients on the PFC surface and sides. Heat flux is determined using the algorithms described in section 3.9. 156

5.8 Power difference from 5gP5vP5vS (ground truth) heat flux on all mesh elements from geometry described in figure 5.5. Number values in trace names corresponds to number of samples from that distribution, i.e. 2gP3vP4vS would signify 2 samples from gyro phase angle distribution, 3 samples from velocity phase angle distribution, and 4 samples from maxwellian speed distribution. Distributions are discussed in section 3.9.2. Top plot is a rug plot, which plots ticks along 1D coordinate. Bottom plot is histogram. Increasing the samples decreases the spread and skew of the distributions, corresponding to convergence. 159

5.9 Absolute value of interquartile Range (IQR) for each dataset plotted in figure 5.8 subtracted from the ground truth IQR (i.e. $|IQR(N,N,N) - IQR(5,5,5)|$). Convergence in IQR corresponds to converging data spread. Data point at N=5 for (5,5,5) is 0 because IQR(5,5,5) - IQR(5,5,5) = 0 160

5.10 Green trace is the gyro orbit calculation error percentage with respect to the optical calculation. Simulation time for gyro-orbit calculation in blue solid trace. Gyro-orbit calculation consists of running N macro-particle simulations (N=NgP*NvP*NvS), each corresponding to a different sample from the kinetic theory distribution functions. NgP is number of gyro phase samples, NvP is number of velocity phase samples, NvS is number of samples from the Maxwellian speed distribution function. 162

5.11 Example HEAT output for 5 different macro-particle quantities. Three numbers in parentheses indicate number of samples from each distribution: gyro phase angle, velocity phase angle, and Maxwellian speed distribution, (gyroPhase, vPhase, vSlice). These results are plotted in figure 5.10 to illustrate relative error and speed. Increasing the number of samples beyond (3,3,3) does little to improve the integrated power error, but creates a smoother profile. 163

5.12 Divertor geometry for gyro orbit analysis of lower single null steady state plasma. Entire lower divertor is shown on the right. IBDH tile is illuminated in green, and exploded view provided with green outline. A single castellation is provided in lower left figure. Analysis in this section uses single castellation geometry. 166

5.13 Optical and gyro-orbit heat flux predictions for a single castellation in the NSTX-U IBDH tile. Left frames correspond to the optical approximation. Right frames correspond to the gyro-orbit frames. Top and bottom rows each show different vantage points. Primary differences between optical and gyro-orbit approximation are the magnetic shadows and leading edges. 167

5.14 Optical (green) and gyro-orbit (blue) heat flux traces along a 1D coordinate on the castellation surface in toroidal direction. $S_{tor} = 0$ corresponds to the top surface corner. Castellation top surface is called out with an arrow in magenta. Optically shadowed regions are greyed out. Gyro orbit heat flux can access regions that are optically shadowed. 170

5.15 Optical (green) and gyro-orbit (blue) heat flux traces along a 1D coordinate on the castellation surface in poloidal direction. $S_{pol} = 0$ corresponds to the top surface corner. Castellation top surface is called out with an arrow in magenta. Optically shadowed regions are greyed out. Gyro orbit heat flux peaks on corners, due to the ability of the gyro-orbit helical trajectories to wrap around top surface corners and strike vertical side surfaces. 171

5.16 Time evolution of temperature for the NSTX-U castellation corresponding to the heat fluxes in figure 5.13. Top row shows temperatures from optical approximation heat flux, bottom row shows temperatures from the gyro-orbit approximation heat flux. Each column corresponds to a different timestep. Peak temperature for the optical approximation is reached on the faceted optical hot spot corner. Peak temperature for the gyro-orbit approximation occurs on the narrow poloidal facing leading edge. Peak temperatures are within 5% of each other. 173

5.17 Peak temperature comparison for optical (green dotted trace) versus gyro-orbit (blue solid trace) approximations, for $P_{SOL,outer} = 4.9MW$. Heat fluxes correspond to figure 5.13. Temperatures correspond to figure 5.16. For both heat flux models, the peak temperature is reached at the end of the discharge when t=5s. The peak temperature is 2720 K for the optical approximation and 2631 K for the gyro-orbit approximation. Optical trace crosses the temperature limit (1873 K) at 2.21s and gyro-orbit trace crosses the limit at 2.37s. 174

5.18 Heat fluxes for the optical approximation (left) and gyro-orbit approximation (center and right) for the entire NSTX-U IBDH tile in the lower divertor. With the optical approximation, shadows have sharp edges and cover a large fraction of the PFC. One large inter-PFC shadow can be observed on the left edge of the tile, and small inter-castellation shadows can be observed on the left edge of each castellation. Turning on the gyro-orbit approximation with average ion temperature set to 10eV fills in the inter-castellation shadows, but has a minimal effect on the inter-PFC shadows. Increasing the average ion temperature to 100eV increases the gyro-radii of ions, which results in total smearing of the optical approximation shadows. HEAT can perform gyro-orbit analyses on entire PFC geometries directly from engineering CAD. 177

6.1 Tokamak Energy ST40 2.2 divertor CAD (top frame), heat flux (bottom left) and temperature at 150 ms (bottom right). The highly optimized 2.2 divertor creates an extremely non-axisymmetric heat flux footprint. The temperatures for this scenario are well below the enginering limits. 180

6.2 Side by side comparison of ST40 IR experimental data and HEAT predictions for a center stack limited scenario. Good qualitative agreement between the experimental and predicted profiles can be observed, but further quantitative analysis is required to validate the HEAT output. 181

6.3 Normalized poloidal flux (top left) and MHD equilibrium (top right) for the high X-point negative triangularity scenario in DIII-D. Outer strike point is on wall PFCs, which were not designed to handle high heat loads. Bottom frame shows heat flux that corresponds to this magnetic equilibrium and geometry. Plasma parameters are overlaid onto the figure. 184

6.4 High X-point scenario heat flux as viewed from a different vantage point compared to figure 6.3. Top frame illustrates the consequences of no toroidal shaping on these PFCs, where the tile leading edge is exposed to heat fluxes at near normal angles of incidence. Bottom frame shows temperature evolutions for synthetic HEAT probes placed at the leading edges. For both the inner and outer strike point, the engineering limit (1873K) of graphite is reached. Note that extreme temperature for T002A probe is likely higher than reality, a product of the coarse volume mesh chosen for these preliminary workflow scoping studies. Further analysis is required to validate this result. 185

6.5 Normalized poloidal flux (top left) and MHD equilibrium (top right) for the low X-point negative triangularity scenario in DIII-D. Outer strike point is on wall PFCs, which were not designed to handle high heat loads. Bottom frame shows heat flux that corresponds to this magnetic equilibrium and geometry. Plasma parameters are overlaid onto the figure. This scenario has extremely high flux expansion, which results in the self shadowing effect seen at the inner divertor. 187

6.6 Views of the low X-point scenario from different vantage points from 6.5. Due to the high flux expansion of this discharge, the inner strike point sees heat loads on the back side of the PFC, as shown in top frame. Additionally, the rectangular PFC structure creates a faceted optical hot spot inside the gap. 188

6.7 Temperature evolution for the low X-point PFCs, for two thermal probe locations. The top frame represents the temperature evolution in the leading edges, while the bottom frame represents the temperature evolution far from the leading edges. Both of the leading edges reach the engineering limit of graphite (1873K), while the bulk tile (closer to the axisymmetric assumption) maintains temperatures below the engineering limit. Note that extreme temperature for T002A probe is likely higher than reality, a product of the coarse volume mesh chosen for these preliminary workflow scoping studies. Further analysis is required to validate this result. 189

6.8 Divertor Material Evaluation System (DiMES) head and retarding field energy analyzer (RFEA) insert for DIII-D. Top frame shows HEAT calculated heat flux for an initial (not final) design, where the access slit for ions was magnetically shadowed. Final design changed geometry so that access slit was not shadowed. Bottom slide, courtesy of Dr. Jun Ren, shows example HEAT output field line trace and temperature output. HEAT was used by Ren to iterate on optimal DiMES head design. 192
6.9 DIII-D movable divertor analysis courtesy of Dr. Andreas Wingen (ORNL). The movable divertor can operate in an open or closed configuration, which is controlled by actuators. Results here are slides from Dr. Wingen's analysis, showing three potential magnetic equilibria (top), as well as the HEAT output for the center equilibrium (186915). 193
6.10 STEP / DEMO CAD (top panels) and HEAT simulated divertor heat flux (lower panels). Input parameters provided by STEP team. 196
6.11 Example HEAT output for STEP inner limiter. Limiter toroidal chamfer (top left), magnetic field line incident angle (top center) and shadowMask (top right). Bottom panels show heat flux predictions using input parameters provided by STEP team. 197
6.12 Comparison of heat flux prediction provided by SMARDDA (left panel) and HEAT (right panel) for the outboard limiter in STEP. Heat flux footprints are similar but magnitudes differ. SMARDDA output courtesy of Dr. Lorena Richiusa (UKAEA). 198

Chapter 1

Background Physics

This section provides a brief introduction to the physics involved in calculating the heat loads incident upon the tokamak plasma facing components (PFCs). For a full description of tokamak physics one is encouraged to review the references. This synopsis is meant to serve as a reference for the discussion in the following chapters, not as a review of the field in its entirety.

1.1 Magnetic Confinement Fusion

While there are many methods for achieving nuclear fusion, none has been researched to the extent of magnetic confinement thermonuclear fusion. In thermonuclear fusion, a plasma is heated to high temperatures such that the ions have a sufficiently high probability of quantum tunneling through the coulomb barrier [1]. In order to prevent this hot plasma from contacting the walls of the vessel in which it is contained (which would end the reaction), some form of confinement must be employed. Magnetic confinement fusion leverages the fact that charged particles follow trajectories along magnetic field lines to confine this hot plasma in a magnetic bottle. Once a sufficient number of particles are confined at a sufficient temperature for a sufficient period of time, nuclear fusion energy production becomes possible [2].

1.1.1 Tokamaks

One method for magnetic confinement is the tokomak. Figure 1.1 provides an illustration of ITER, the massive tokamak project under construction in France. Today, the tokamak is the leading candidate for magnetic confinement fusion. Originally conceived in the Soviet Union in the middle of the 20th century, the tokamak has been extensively researched by scientists and engineers from all corners of the globe.

Shaped in the form of a toroid (or a donut) this configuration employs toroidal (long way around) and poloidal (short way around) magnetic fields to confine the plasma. These magnetic field components are induced by a combination of current flowing in the plasma (poloidal magnetic field) as well as electromagnets that are spaced at regular poloidal intervals (toroidal magnetic field). The ratio of toroidal to poloidal magnetic field components provides stability to the plasma, and is termed the safety factor, q, which corresponds to the helicity of the field lines [3].

Because of the helical nature of the magnetic field, ionized particles will circumnavigate the machine toroidally many times before circumnavigating poloidally, as they follow the magnetic field lines. As these particles proceed around the machine, most will remain on surfaces of constant magnetic flux. These flux surfaces result in a series of nested toroids, and when taking a poloidal cross section of the tokamak one observes contours of constant poloidal flux, as illustrated in figure 1.2 [4].

Flux surfaces that do not intersect the wall of the machine are termed closed, and the largest closed flux surface is termed the Last Closed Flux Surface (LCFS) or the Seperatrix. Particles that are confined within the LCFS are considered to be in the plasma core, and for the most part will circumnavigate the machine ad infinitum. Yet there is a possibility for core particles to diffuse across the LCFS to the region outside of the LCFS. This region is called the Scrape Off Layer (SOL), and is comprised of open field lines that terminate on plates engineered to withstand the large particle and heat fluxes incident upon them. Figure 1.3 provides a graphical representation of the LCFS and SOL.

External coils may be employed to divert these incident fluxes to controllable locations on the machine wall. The assembly of plates that these fluxes are diverted towards are collectively called the divertor. The current in the external divertor coil

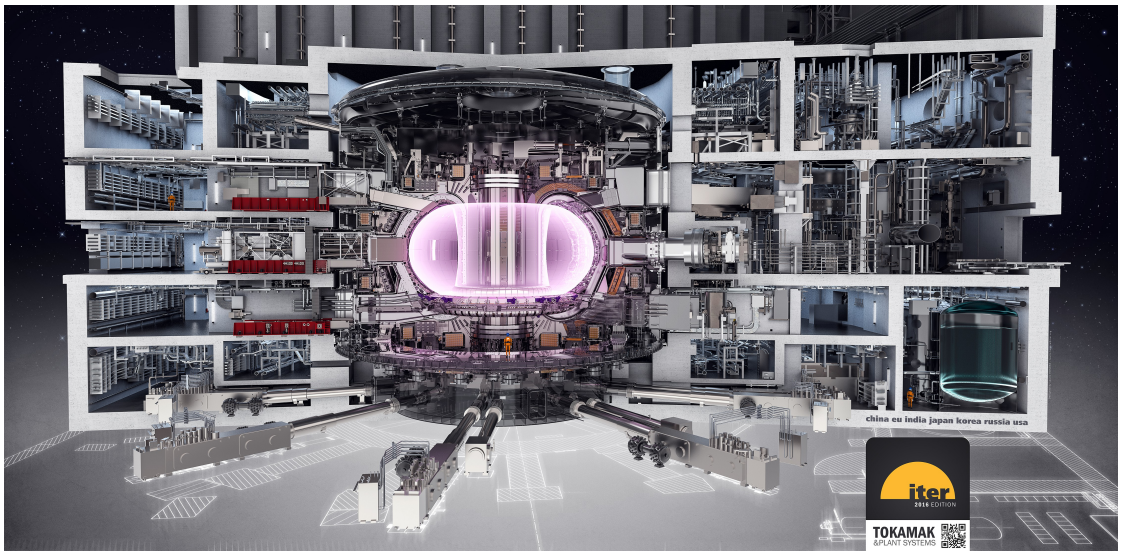


Figure 1.1: ITER tokamak and associated engineering systems ©ITER Organization https://www.iter.org.

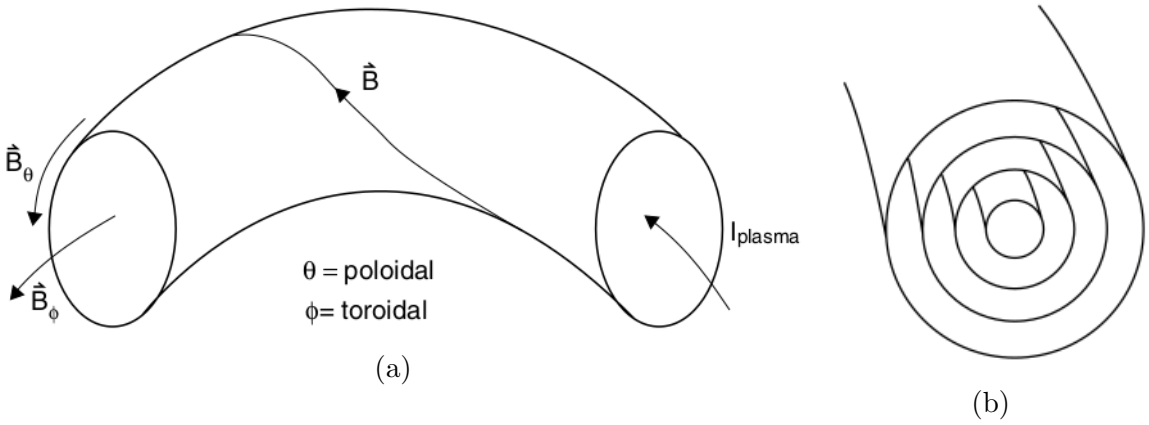


Figure 1.2: Magnetic field components and plasma current on a (a) single flux surface, and (b) nested toroidal flux surfaces from [4].

flows in the same direction as the current in the plasma. As a result, a null in the poloidal magnetic field forms, called the X point. The introduction of the divertor plates/targets and divertor coil creates a plasma flow towards the divertor, caused by thermal and electric gradients in the SOL [4]. This plasma flow is largely parallel to the magnetic field, although it does have a cross-field component. Figure 1.4 illustrates the X point and divertor targets. The plasma confined in the core is called the main plasma. In the SOL, the private flux region is defined as the region below the X-point and the LCFS, while the common flux region is above it. The location where the seperatrix intersects the divertor target plate is called the strike point.

1.1.2 The Tokamak Divertor

The ability for divertor target plates to withstand high heat and particle fluxes will be a key constraint on the size and operational domain of power producing tokamaks [5, 6]. Proper management of these heat and particle fluxes proves to be a difficult engineering challenge. The incredible radiation and heat loads in future tokamaks will require the use of advanced materials for plasma facing components and for structural members, as well as advanced operational techniques to maintain the PFCs within their engineering limits. Preliminary designs for a demonstration reactor (DEMO) indicate that divertor components will see incident heat fluxes up to 20 $\mathrm{MW\,m^{-2}}$ during steady state operation with neutron damage levels of 6-7 dpa per full power year [7]. ITER, the international tokamak project in France, is expected to see divertor incident heat fluxes of nearly 25 $\mathrm{MW\,m^{-2}}$ at low neutral pressures [8], which exceeds the engineering limit of 16 $\mathrm{MW\,m^{-2}}$ [8]. Preliminary projections for SPARC, the tokamak soon to be constructed by private company Commonwealth Fusion Systems, predict incident heat fluxes in excess of 300 $\mathrm{MW\,m^{-2}}$ [9]. These extreme conditions will require innovative designs with regards to selections of materials, geometry, and machine operation, in order to maximize the operational domain while minimizing material degradation. Strike point sweeping, divertor detachment, and PFC optimization, are some example mitigation strategies.

Heat loads in the SOL are a result of particle transport across the separatrix. Heat is transferred through the SOL by conduction, convection, and radiation, to the

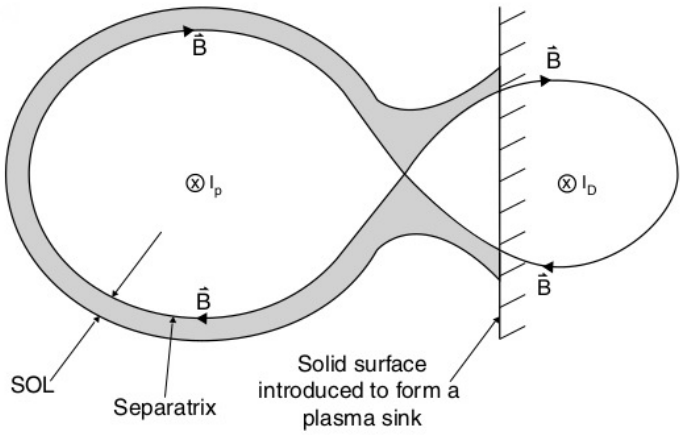


Figure 1.3: Example SOL and open field lines terminating on machine wall from [4].

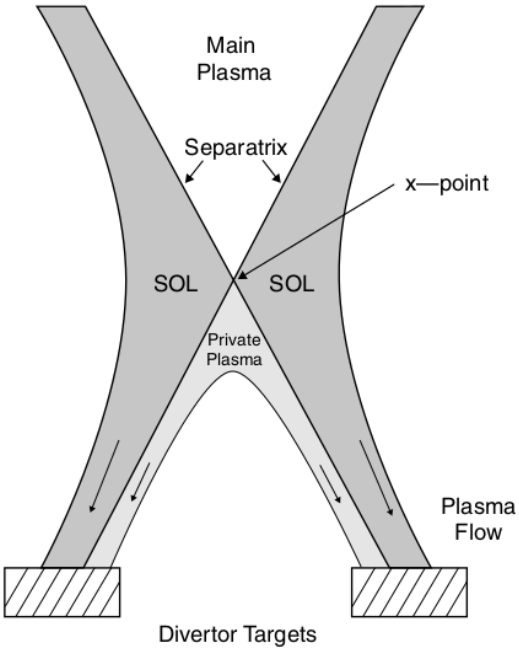


Figure 1.4: Example X point and divertor targets from [4].

divertor target. Heat flows into the divertor target (sometimes referred to as a tile) where it thermally diffuses throughout the tile. Some divertor tiles are actively cooled by coolant loops that run through the tile and carry heat away from the divertor, while some are not. In all cases, thermal gradients within the tile induce stresses that can potentially cause the tile to break if the tile exceeds its fracture limit. Additionally, if the surface temperature for a target material exceeds its limit, then it will sublimate into the main plasma or melt [10].

In order to avoid these failure mechanisms, it is desireable to predict the heat flux incident upon the divertor tile. The primary objective of this research is to create a simulation framework capable of predicting these complicated heat loads to guide machine operation and PFC design. The following sections in this chapter provide some background physics necessary to make these predictions.

1.2 Models of SOL Heat Flux

This sections outlines some of the scrape of layer (SOL) heat flux models that pertain to the analysis in this dissertation.

1.2.1 The 1D NSTX-U Working Group Model

This section follows the method developed by the National Spherical Tokamak eXperiment Upgrade (NSTX-U) PFC Working Group to calculate divertor heat fluxes [11]. To begin, consider the SOL heat flux as being approximately field aligned,

$$\vec{q} \approx \vec{q_{||}} = q_{||} \frac{\vec{B}}{B}. \tag{1.1}$$

Next, assume that there are no heat sources in the SOL,

$$0 = \nabla \cdot \vec{q} \tag{1.2}$$

$$= \left(\nabla \frac{q_{||}}{B}\right) \cdot \vec{B} + \frac{q_{||}}{B} \left(\nabla \cdot \vec{B}\right). \tag{1.3}$$

The second term on the RHS is zero per Gauss's law for magnetism. The result is,

$$0 = \left(\nabla \frac{q_{||}}{B}\right) \cdot \vec{B} \tag{1.4}$$

which implies that the gradient term must be equal to 0 for nonzero magnetic fields,

$$0 = \left(\nabla \frac{q_{||}}{B}\right) \tag{1.5}$$

This result allows some flexibility when selecting a function for $q_{||}$. More specifically, any scalar flux function whose gradient is zero can satisfy 1.4. Mathematically this can be expressed as,

$$\nabla \frac{q_{||}}{B} \cdot \vec{B} = \nabla f(\psi) \cdot \vec{B} = 0 \tag{1.6}$$

$$\rightarrow q_{||} = f(\psi) B \tag{1.7}$$

where ψ represents magnetic flux. This indicates that the scalar flux function, $f(\psi)$, is constant on surfaces orthogonal to $\vec{B}$. In other words, $f(\psi)$ is a function that defines surfaces of constant magnetic flux. To determine the heat flux profile incident upon the target, one can define $q_{||}$ as a function of ψ. Because $f(\psi)$ defines surfaces of constant flux, and ψ is monotonic in the SOL, it is convenient to use ψ as a proxy for the radial coordinate. Choosing $f(\psi) = q_{||0} q(\psi)$ yields the following definition of $q_{||}$ as a function of ψ,

$$q_{||} = q_{||0} B q(\psi). \tag{1.8}$$

where $q_{||0}$ is a constant used to scale the profile to the peak heat flux density value and $q(\psi)$ is the heat flux profile along the ψ coordinate. Because the area of interest is in the SOL, it is convenient to normalize ψ so that $\hat{\psi} = 1$ at the seperatrix,

$$\hat{\psi} \equiv \frac{\psi - \psi_{axis}}{\psi_{sep} - \psi_{axis}} \tag{1.9}$$

$$\text{such that} \quad q_{||} = q_{||0} B \hat{q}(\hat{\psi}). \tag{1.10}$$

Now that $q_{||}$ has been defined, it is possible to determine the heat flux at the outboard mid-plane (omp) and the divertor (div),

$$q_{||omp} = q_{||0}\hat{q}(\hat{\psi})B_{omp} \tag{1.11}$$

$$q_{||div} = q_{||0}\hat{q}(\hat{\psi})B_{div} \tag{1.12}$$

Using the fact that heat flux is conserved in each flux tube (equation 1.3), one can divide $q_{||div}$ by $q_{||omp}$ and solve for the divertor heat flux profile relative to the profile at the outboard midplane,

$$q_{||div} = q_{||omp}\frac{B_{div}}{B_{omp}} = q_{||0}\hat{q}(\hat{\psi})\frac{B_{div}}{B_{omp}} \tag{1.13}$$

This equation maps a heat flux profile, $\hat{q}(\hat{\psi})$, at the midplane to the divertor. Because the incident field lines ($\hat{b} = \vec{B}/B$) are not aligned with the divertor tile surface normal ($\hat{n}_{div}$), an additional factor of $\hat{n}_{div} \cdot \hat{b}$ must be included to account for the field line projection onto the tile,

$$q_{div} = q_{||0}\hat{q}(\hat{\psi})\frac{B_{div}}{B_{omp}}(\hat{n}_{div} \cdot \hat{b}) \tag{1.14}$$

This result is the heat flux profile across the divertor target.

1.2.2 The Heuristic Model

This section follows the heuristic model developed by R. Goldston [12]. In collisional tokamak plasmas grad B and curv B drifts create flow in the vertical direction for both ions and electrons. This flow crosses the separatrix into the SOL. The flow direction is charge dependent, and the plasma compensates for the charge separation with Pfirsch-Schlüter flows [13]. While Pfirsch-Schlüter flows (from the bottom to the top of the tokamak) are one mechanism by which electrical neutrality can be restored, it is also possible for a downwards directed flow to the divertor to compensate for the charge separation. The heuristic model estimates the characteristic width (in the SOL) of this flow to the divertor.

If it is assumed that volumetric sources are nonexistent in the SOL (no gas puffing, etc.) and that the tokamak is operating in the attached regime during H-mode operation, then the following estimate for the drift based SOL width can be made,

$$\lambda = \frac{4a}{\bar{Z}eB_pR}\left(\frac{\bar{A}m_pT_{sep}}{(1+\bar{Z})}\right)^{1/2} \tag{1.15}$$

$$\text{where} \quad \bar{Z} \equiv \frac{n_e}{\sum_i n_i}; \qquad \bar{A} \equiv \frac{\sum_i n_i A_i}{\sum_i n_i}. \tag{1.16}$$

This width corresponds to a channel of flow in the SOL, driven by grad B and curv B drifts across the separatrix. The second assumption is made that anomalous electron thermal diffusion is the source of heat in the channel defined by equation 1.15. The diffusing electron density 'decays' into the flow channel, and is responsible for heat transport. Leveraging this assumption to eliminate T_{sep} in equation 1.15 yields the final expression for the SOL heat flux width,

$$\lambda = 5671 \cdot P_{SOL}^{1/8}\frac{(1+\kappa^2)^{5/8}a^{17/8}B^{1/4}}{I_p^{9/8}R}\left(\frac{2\bar{A}}{\bar{Z}^2(1+\bar{Z})}\right)^{7/16}\cdot\left(\frac{Z_{eff}+4}{5}\right)^{1/8} \tag{1.17}$$

under the assumption that ion magnetic drift dominates the net particle transport, and

$$\lambda = 5671 \cdot P_{SOL}^{1/8}\frac{(1+\kappa^2)^{5/8}a^{17/8}B^{1/4}}{I_p^{9/8}R}\left(\frac{2\bar{A}}{(1+\bar{Z})}\right)^{7/16}\cdot\left(\frac{Z_{eff}+4}{5}\right)^{1/8} \tag{1.18}$$

under the assumption that electron magnetic drift dominates the net particle transport [12].

1.2.3 The Eich Model

This sections follows the work by T. Eich [14, 15]. Using data from a multimachine database of measured SOL heat flux decay widths, Eich et al. created a parametric scaling of heat flux width from machine parameters. JET, DIII-D, ASDEX-U, C-MOD, MAST, and NSTX data was included in the database. The discharges were all H-mode discharges and the heat flux width applies only to the inter-ELM periods.

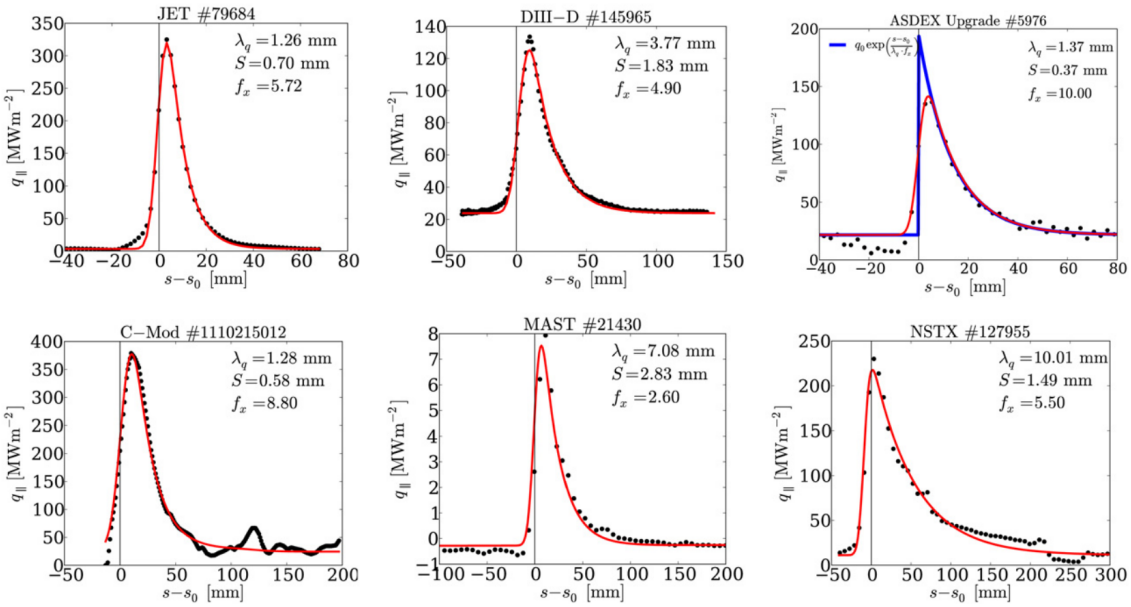


Figure 1.5: Equation 1.20 overlaid onto experimental data from [15].

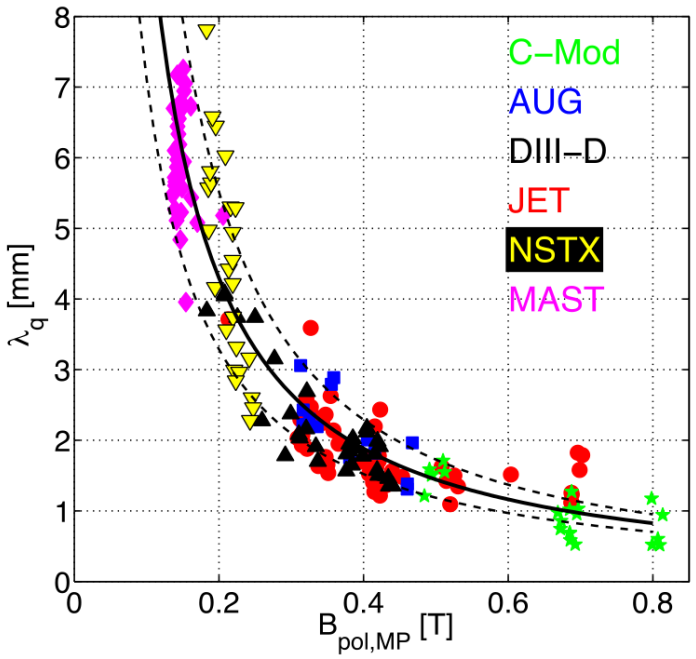


Figure 1.6: Scaling of heat flux width as a function of B_{POL} from [15].

Assuming that the energy transported across the separatrix decays exponentially, one can describe the outboard midplane (omp) heat flux profile with the following relationship,

$$q(r) = q_{||} e^{-\frac{r}{\lambda_q}} \tag{1.19}$$

where $r = R - R_{sep}$, R_{sep} is the major radius at the outboard midplane, and $q_{||}$ is a scalar coefficient that scales the profile to the power crossing the LCFS. After defining this exponential decay at the outboard midplane, a gaussian function of width S is convoluted against $q(r)$. The convolution represents perpendicular energy transport as the heat flux at the outboard midplane, $q(r)$, makes its way to the divertor target, and accounts for heat flux observed in the private flux region. The heat flux profile after this convolution is,

$$q(\bar{s}) = \frac{q_0}{2} \exp\left[\left(\frac{S}{2\lambda_q f_x}\right)^2 - \frac{\bar{s}}{\lambda_q f_x}\right] \operatorname{erfc}\left(\frac{S}{2\lambda_q f_x} - \frac{\bar{s}}{S}\right) + q_{BG} \tag{1.20}$$

where $\bar{s} = s - s_0 = (R_{sep} - R) \cdot fx$, f_x is the magnetic flux expansion, λ_q is the heat flux width in question, and q_{BG} represents background heat flux. Figure 1.5 overlays this profile onto experimental data from each of the six machines in the database. As can be observed, the convolution of exponential and gaussian profiles provides a good fit to the experimental data. Additionally, the ASDEX-U plot contains the original exponential decay. In order to match 1.20 to the experimental data one must establish the value of λ_q. This was achieved via multiple regressions. More specifically, the following equation was utilized to regress the correct value of λ_q,

$$\lambda_q = C_0 B_{tor}^{C_1} P_{SOL}^{C_2} R_{geo}^{C_3} B_{pol}^{C_4} \left(\frac{n}{n_{GW}}\right)^{C_5} \left(\frac{a}{R_{geo}}\right)^{C_6} \tag{1.21}$$

where C_X, $X \in \{0, 1, 2, 3, 4, 5, 6\}$ are all unitless scaling coefficients that scale each respective parameter. Perhaps one of the most interesting results from Eich's regressions is the scaling of λ_q as a function of B_{POL} at the outboard midplane. Figure 1.6 illustates regression 14 from [15] with experimental datapoints from all six machines overlaid.

1.2.4 The Brunner Model

This section follows the work by D. Brunner [16]. The discussion from the previous section, and from [15], indicates that high B_p corresponds to a narrow heat flux width, λ_q. In order to improve understanding of the high B_p on heat flux width, Brunner et al. performed empirical regressions on 300 Alcator C-Mod shots, where B_p reached 1.3T during the final experimental campaign in 2016. While the previous regressions from Eich et al. [15, 14] solely included H-mode discharge data, the C-mod data investigated by Brunner also contained L-mode, EDA H-mode, and I-mode discharges.

In contrast to the Eich heat flux function (equation 1.20) where a single exponential is convoluted with a gaussian, Brunner defines the heat flux profile as a superposition of four different exponential functions, each with a respective domain of validity. This profile can also be convoluted with a gaussian to account for perpendicular diffusion on the way to the target. These four profiles can be broken into a piecewise function corresponding to the private and common flux regions. Mathematically, this can be described by the following form of q,

$$q(s) = \begin{cases} (q_0 - q_{pf})e^{\frac{s}{\lambda_{q,pn}}} + q_{pf}e^{\frac{s}{\lambda_{q,pf}}} & r < 0 \\ (q_0 - q_{cf})e^{-\frac{s}{\lambda_{q,cn}}} + q_{cf}e^{-\frac{s}{\lambda_{q,cf}}} & r \geq 0 \end{cases}$$

where pf refers to the private far SOL, pn refers to the private near SOL, cf refers to the common far region, and cn refers to the common near region. A comparison between the Eich profile (also called the single-λ profile) and the Brunner profile (also called the multi-λ profile) is provided graphically in figure 1.7. While the Eich profile (right plot) suffices to qualitatively describe the heat flux profile for the near SOL in both the common and private flux regions, it deviates from experimental data in the far SOL regions on both sides of the strike point. The Brunner profile (left plot) with four exponential decays shows good agreement across all SOL domains. It appears that for Alcator C-Mod, a multiple e-folding profile is better suited to describe the heat flux profile in the far SOL. Figure 1.8 extends figure 1.6 from [15] to include the higher range of B_p, and compares the single-λ and multi-λ heat flux widths. Figure 1.8 extends beyond ITER's maximum poloidal field (1.2T).

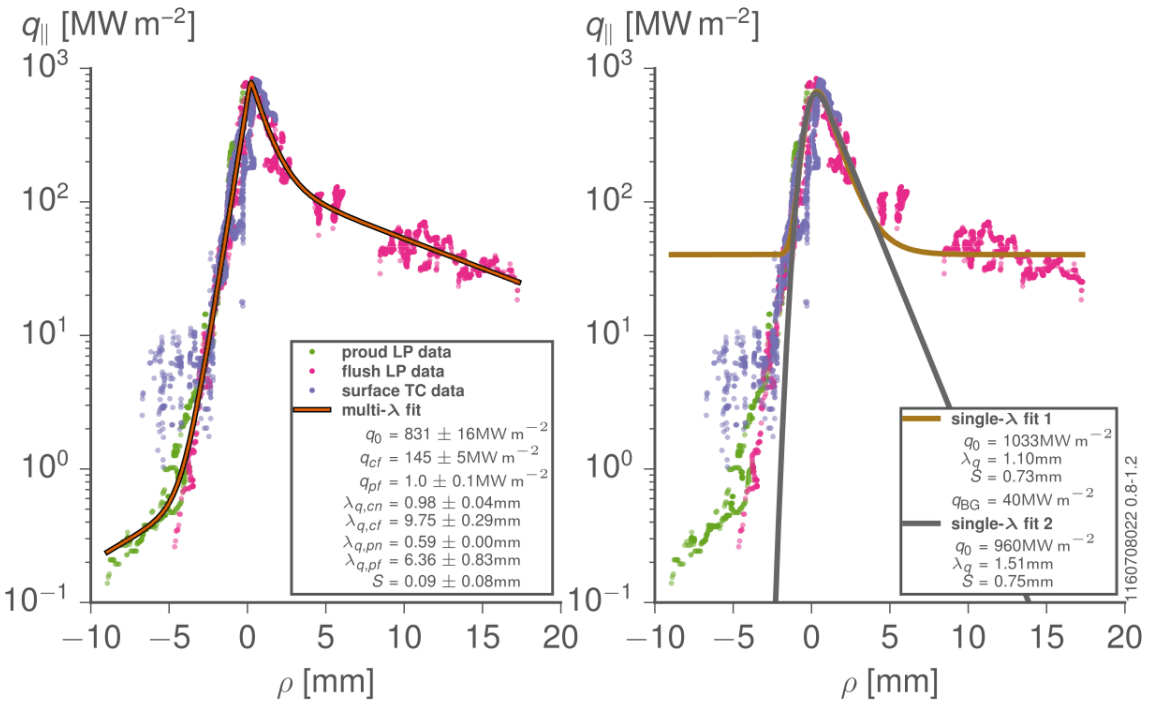


Figure 1.7: Eich vs Brunner heat flux profiles and Alcator C-Mod diagnostic data from [16]

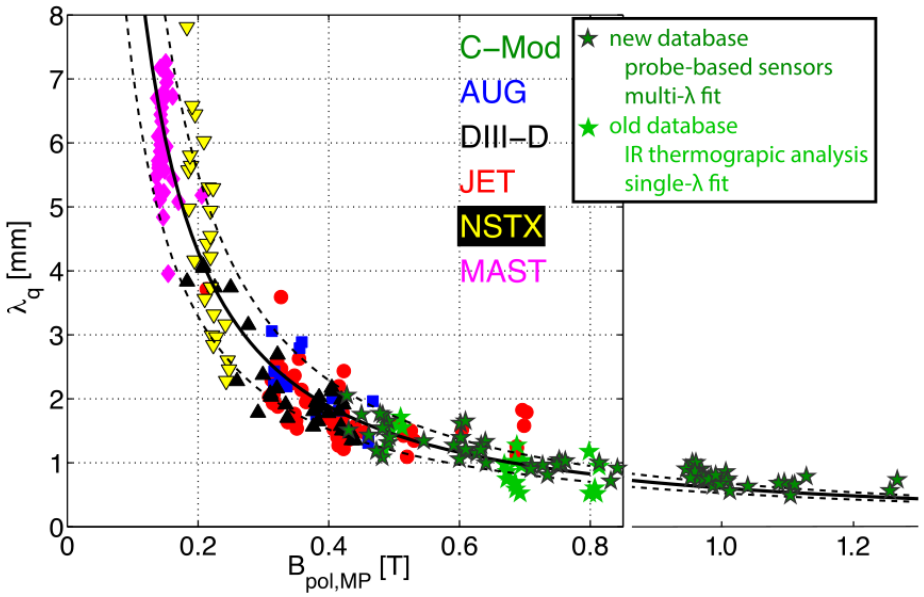


Figure 1.8: Extension of figure 1.6 using new C-Mod data and Brunner heat flux profile from [16]. This scaling puts the heat flux widths for ITER and SPARC well below 1mm.

1.3 Charged Particle Trajectories in the SOL

The motion of charged particles in magnetized plasmas is addressed in this section. First, the simplest assumption of particle motion is presented, the optical approximation. Next, the helical trajectory of charged particles that results from Larmor gyration will be discussed, which is referred to as the gyro-orbit approximation. An overview of gap and edge loads due to gyro-orbit effects from the fusion literature will be presented. Both of these approximations ignore the effects of electric fields and plasma drifts, and the consequences of ignoring electric fields will be discussed. Last, a review of Maxwellian velocity distribution functions will be provided, which serves as the foundation for a kinetic treatment of particle motion.

1.3.1 The Optical Approximation

The simplest model for particle motion in the SOL is called the optical approximation. In the optical approximation, particles are assumed to flow directly along the magnetic field lines, and gyromotion about the field lines is ignored. Figure 1.9 provides an example of an optical approximation trajectory. The magnetic field lines in a tokamak are curved, and particles are assumed to flow directly along these curved field lines, as can be observed by the yellow trace in the figure. This motion along the field lines can be summarized mathematically by the expression,

$$\vec{v} = \vec{v_{||}} = v\frac{\vec{B}}{B} \tag{1.22}$$

where $\vec{v}$ is the particle velocity vector, $\vec{v_{||}}$ is the component of velocity parallel to the magnetic field, v is the scalar particle speed, and $\vec{B}$ is the magnetic field vector. Because power is carried by individual particles, it can similarly be assumed that power flows directly along the magnetic field lines,

$$\vec{q} = \vec{q_{||}} = q_{||}\frac{\vec{B}}{B} \tag{1.23}$$

where $\vec{q}$ is the heat flux vector and q is the scalar heat flux magnitude. The optical approximation is utilized in all of the models of section 1.2, where power is transported through the SOL along the magnetic field lines.

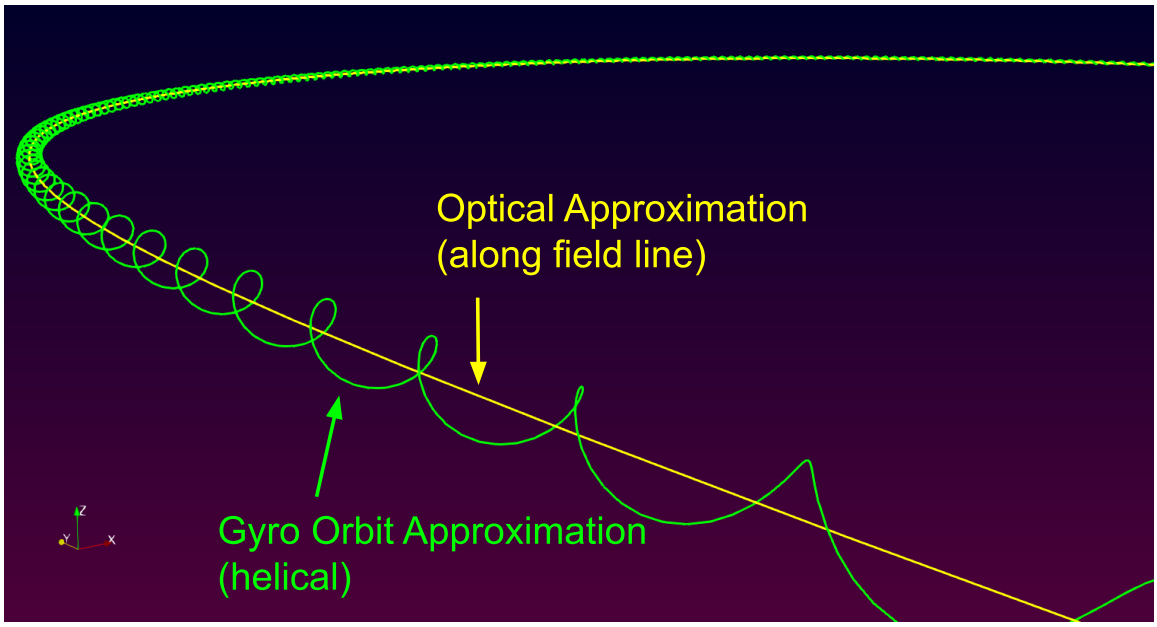


Figure 1.9: Example magnetic field line (yellow) and helical gyro orbit trajectory (green). The optical approximation assumes particles follow the curved magnetic field lines, while the gyro orbit approximation assumes particles precess about the field lines with finite Larmor radii.

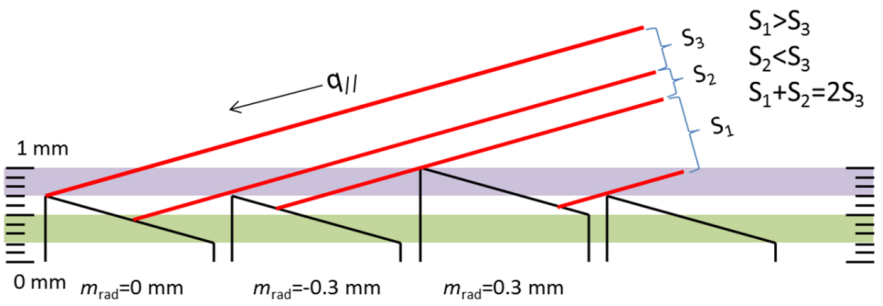


Figure 1.10: Example of monoblock beveled surfaced from [17]. Shaded regions represent allowable alignment engineering tolerances.

1.3.2 Gaps and Edge Loading

A common assumption employed when calculating divertor heat fluxes is the optical approximation, where the effect of ion and electron Larmor gyration about magnetic field lines is ignored. Instead, the trajectory of the guiding center is utilized. The optical approximation provides good qualitative heat flux profiles, but the omission of the Larmor orbits can be significant with respect to heat loading, especially for tiles that have complicated 3D geometries. Additionally, the optical approximation neglects the effect of Debye shielding that occurs in the sheath, which manifests in the form of electric fields capable of altering an ion's trajectory as it approaches the divertor tile. These effects have been studied computationally, with varying degrees of complexity [17, 18, 19].

Gunn et al. performed a comprehensive investigation into the effects of cyclotron / Larmor gyration on divertor heat loading for the tungsten monoblocks in ITER [17]. The ITER divertor tiles are cassettes, consisting of many tungsten monoblocks that extend 28mm in the toroidal direction, 12mm in the poloidal direction, and are 28mm deep. In order to prevent excessive heat loading on the edges of these monoblocks, they are beveled (fish-scaled) to 'shadow' the leading edges from the incident plasma and heat flux. Figure 1.10 provides an illustration of this shadowing effect. Magnetic field lines that would strike the monoblock edge before beveling are intercepted by the beveled surface and prevent heat loading of the edge.

While figure 1.10 provides insight into the shadowing effect of using a beveled tile surface, it neglects the effects of ion cyclotron orbits mentioned above. An ion travelling along the magnetic field line precesses around the magnetic field line at the cyclotron (gyro) frequency, qB/m, while its guiding center follows a trajectory of the magnetic field line. When the divertor tiles consist of beveled surfaces and inter-tile gaps as in figure 1.10, it is possible for the ion to graze the beveled surface, descend into the gap, and strike the vertical tile surface inside the gap. Figure 1.11 from [17] illustrates this concept. The blue line represents the oscillation about the guiding center as the particle approaches the PFC. The oscillatory nature of the trajectory allows the particle to descend into the gap between two monoblocks. The maximum distance a particle may descend into the gap is given by δ, shown in red. Use of the

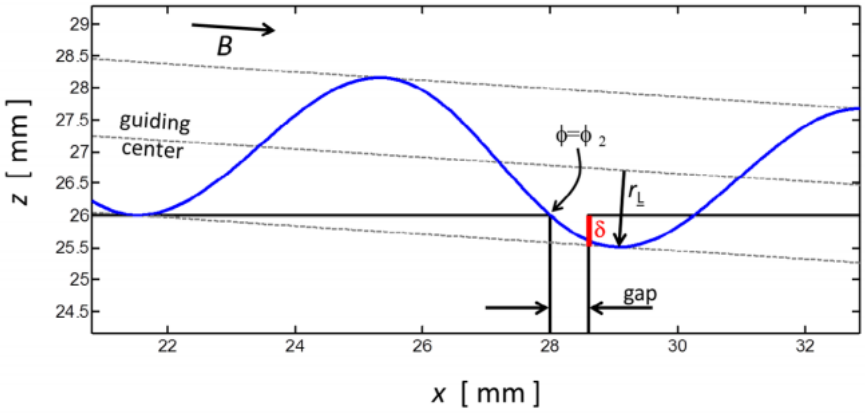


Figure 1.11: Example of gyration about magnetic field line (blue), and potential for heat loading in tile gaps from [17].

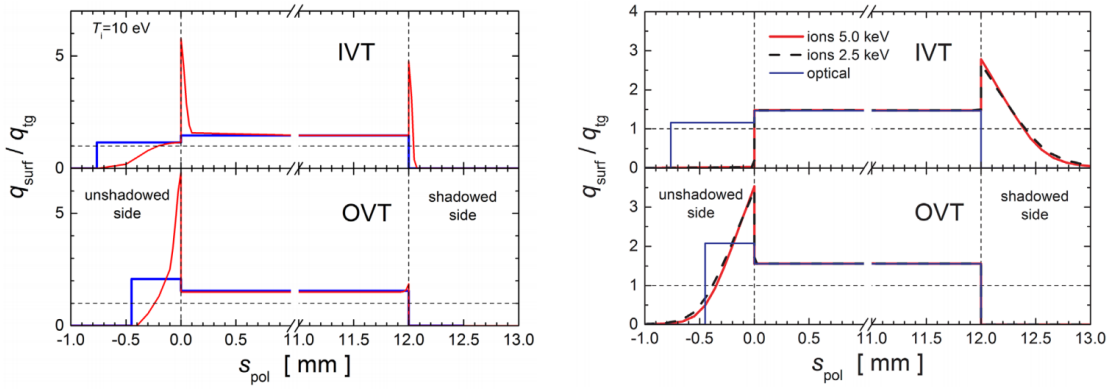


Figure 1.12: Example of discrepancy between optical approximation (blue) and gyro-orbit approximation (red), for the inter-ELM and ELM heat fluxes on ITER vertical tile monoblocks from [17]. The plasma facing surface is in the region $0.0 < S_{pol} < 12.0$ and the region outside this domain represents the gaps.

optical approximation does not take the gyration effect into account, and assumes that all surfaces in a magnetic field line shadow will not see any incident heat flux.

Investigations into gap and edge loading have revealed that heat loads in the gap and on the edges in shadowed regions can exceed the allowable heat flux for a given PFC [17]. Figure 1.12 illustrates the discrepancy between the optical approximation and a simulation that include gyro-orbit effects. The blue lines provide the heat flux estimates for the optical approximations, the red lines and black lines include gyro-orbit effects. There is significant underestimation of the heat flux in the gaps and on the tile edges in the optical approximation. To correctly estimate heat loading on 'shadowed' surfaces it is therefore mandatory to include the effects of gyro-orbits.

To calculate the surface heat flux to any point on a tile surface, q_{surf}, Gunn [17] provides the following equation,

$$q_{surf} = q_{||0} \int_0^{2\pi} \int_0^{\infty} \int_0^{\infty} \frac{1}{2} m_i (v_\perp^2 + v_{||}^2)(\overrightarrow{v} \cdot \hat{n}) f_\varphi f_\perp f_{||} R_E(v_{||}, v_\perp, \phi) H(v_{||}, v_\perp, \phi) \frac{\mathrm{d}\phi}{2\pi} \mathrm{d}v_\perp \mathrm{d}v_{||} \tag{1.24}$$

where $v_{||}$ and $v_\perp$ are the parallel and perpendicular components of velocity, respectively, $f_\varphi, f_\perp, f_{||}$ represent the velocity distribution functions in each direction, $\hat{n}$, is the tile surface normal vector, R_E is a reflection coefficient from equation 2.3 in [20], and H is a heaviside function which is turned on when the ion does not intersect the tile and turned off when an intersection occurs. This equation describes the heat flux that results from particles with helical orbits, and can be used for gyro-orbit calculations.

1.3.3 The Gyro Orbit Approximation

While the optical approximation is a useful tool to perform quick heat load calculations, it overlooks the true interaction between PFC geometry and the helical trajectories of charged particles in magnetic fields. The following derivation follows many of the common fusion textbooks, such as Stacey [21] or Freidberg [22], and is provided to derive a simplified model for gyro orbit heat flux calculations. The Lorentz force governs the particle trajectory in the magnetic field,

$$\vec{F} = q(\vec{E} + \vec{v} \times \vec{B}) \tag{1.25}$$

where $\vec{F}$ is the electromagnetic force, q is the particle charge, $\vec{E}$ is the electric field, $\vec{v}$ is the particle's velocity vector, and $\vec{B}$ is the magnetic field. Neglecting electric fields (see commentary in section 1.3.4) and taking the dot product of $\vec{v}$ with this equation yields,

$$\vec{v} \cdot \vec{F} = m \cdot \vec{v} \frac{d}{dt}(\vec{v}) = \frac{d}{dt}(\frac{1}{2}mv^2) = \vec{v} \cdot q(\vec{v} \times \vec{B}) = 0 \tag{1.26}$$

which indicates that kinetic energy is conserved during the particle's motion. Assuming the magnetic field is aligned to the $\vec{z}$ direction, i.e. $\vec{B} = B_z \hat{e_z}$, enables one to solve for the equations of motion in each direction,

$$\vec{F} = m\ddot{\vec{r}} = m(\ddot{\vec{x}} + \ddot{\vec{y}} + \ddot{\vec{z}}) = q(\vec{v} \times \vec{B}) \tag{1.27}$$

$$= qB(v_y \hat{e}_x + v_x \hat{e}_y + 0\hat{e}_z) \tag{1.28}$$

$$\rightarrow \ddot{x} = \frac{qB}{m} v_y; \quad \ddot{y} = -\frac{qB}{m} v_x; \quad \ddot{z} = 0 \tag{1.29}$$

where qB/m can be identified as the cycloctron frequency, ω_c. The particle velocity can be broken into two components, one which is parallel to the magnetic field and one that is perpendicular to the magnetic field,

$$\vec{v_{||}} = \vec{v} \cdot \hat{e}_z \tag{1.30}$$

$$\vec{v}_\perp = \vec{v} \cdot \hat{e}_x + \vec{v} \cdot \hat{e}_y \tag{1.31}$$

$$|\vec{v}_{||}| = v_{||} = v_z \tag{1.32}$$

$$|\vec{v}_\perp| = v_\perp = [v_x^2 + v_y^2]^{1/2} \tag{1.33}$$

From equations 1.29 and 1.32, one can see that $\dot{z} = v_z = v_{||} =$ constant, which implies that the parallel kinetic energy is also a constant of the motion. Combining this result with the conservation of total kinetic energy from equation 1.26, one can say that both the perpendicular and parallel energies are conserved independently during the motion. This result will be useful when prescribing velocity distribution functions, but one must remember that we are neglecting electric fields and assuming that the magnetic field is static. Further research will extend this model to include these neglected effects.

Equation 1.29 is descriptive of particle motion that gyrates about the magnetic field. This gyration is termed 'gyromotion', and is periodic at the cyclotron frequency, ω_c. Equation 1.29 can be integrated to yield traditional helix form

$$x = \frac{v_\perp}{\omega_c}\sin(\omega_c t + \alpha) + x_0 \tag{1.34}$$

$$y = \mp\frac{v_\perp}{\omega_c}\cos(\omega_c t + \alpha) + y_0 \tag{1.35}$$

$$z = v_{||}t + z_0 \tag{1.36}$$

where the $\mp$ accounts for diamagnetism (- for ions), α is the gyromotion phase angle, and one can define the gyro radius to be,

$$r_g = \frac{v_\perp}{\omega_c} = \frac{mv_\perp}{|q|B} \tag{1.37}$$

Figure 1.9 illustrates the trajectory of the magnetic field lines in yellow and the helical trajectory of an electron gyro-orbit in green. When using the optical approximation, particles will follow the magnetic field line path. If using the optical approximation, the particle's kinetic energy is simply a function of the parallel velocity, $v_{||}$, and is found directly from the heat flux profiles described in the previous sections. Another way of stating this is that the optical approximation is the limit as $v_\perp \to 0$. When using the gyro orbit approximation, finite Larmor radii result in particles that will follow the gyromotion path, and one must account for both the parallel and perpendicular velocities.

1.3.4 Validity of Neglecting Electric Fields

The gyro orbit approximation in section 1.3.3 ignores the contribution of the electric fields to the equation of motion for the ions. This assumption clearly overlooks important physics, especially the effects of electric fields in the sheath. The electric field in the sheath has the potential to alter the ion trajectory as it approaches the PFC surface as shown in figure 10 in [17], included here in figure 1.13. The validity of ignoring these effects is investigated in several PIC code studies [18, 23] as well as in ion gyro orbit analyses for ITER [17, 24]. The consensus seems to be that including

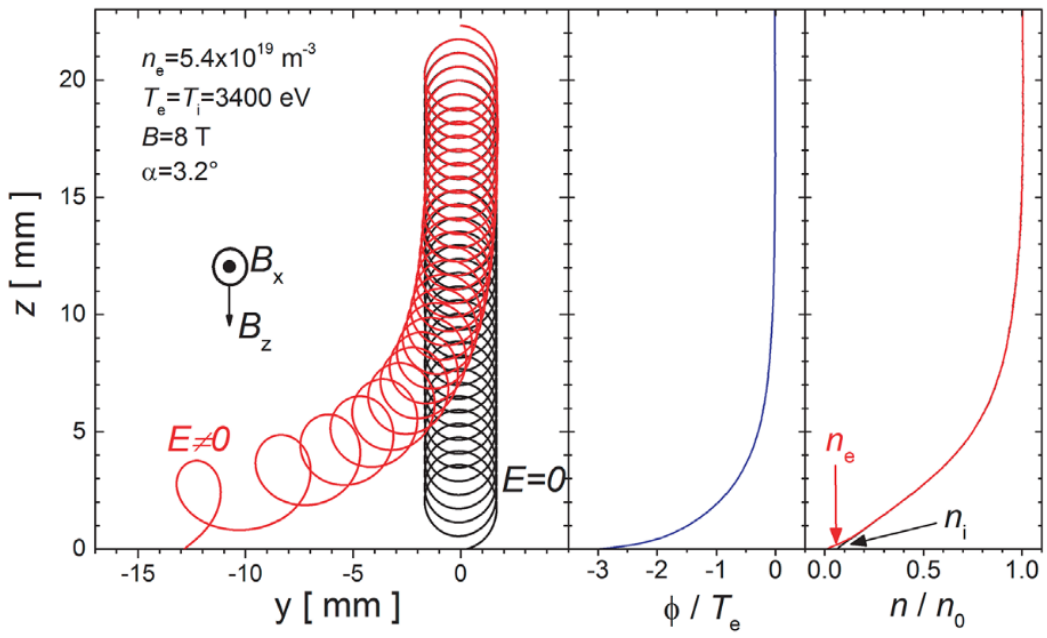


Figure 1.13: ITER full field ELM ion trajectory with and without electric fields, as well as potential and density, from 1D 3V PIC simulations [17, 25]

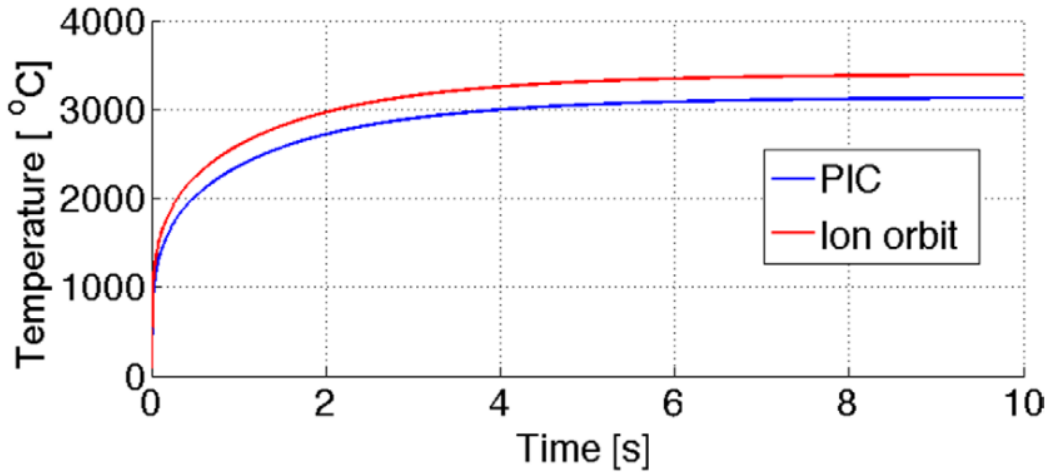


Figure 1.14: Peak temperature difference between including sheath electric fields (PIC) and ignoring electric fields (Ion orbit) from [18]. Including electric fields results in a 10% reduction in peak temperature, implying that using the gyro orbit (ion orbit) approximation is more conservative from an engineering persepective. Results for a steady-state or slow-transient scenario.

electric fields will alter the heat flux profile across the PFC surface, but will have minimal impact on the engineering parameters of interest (ie temperature).

Figure 1.14 from [18] provides a comparison between a PIC code simulation that includes self consistent electric field calculations, and the Ion orbit model from [24] that neglects electric fields. The curves in the figure represent the peak temperature on the ITER monoblock during a steady state simulation (ie no ELMs). At 10s, the difference between the model with electric fields (PIC) and without electric fields (Ion orbit) is around 300 degrees centigrade, corresponding to a 10% temperature difference. Interestingly, the ion orbit model generates a temperature prediction that is higher than the PIC simulation, which is desireable from an engineering factor of safety perspective.

Figure 1.15 provides the geometry included in the PIC simulations performed in [18], and figure 1.16 provides heat flux profiles calculated for the optical approximation, the ion gyro-orbit approximation, and a full PIC calculation that includes electric fields. The case is a steady state scenario with a misaligned ITER Outer Vertical Target monoblock that yields a leading edge located at $s_{tor} = 0$ in the figure, with $s_{tor} < 0$ corresponding to the leading edge face. The Debye length for these simulations ranges from 2-5 um, the angle of incidence is approximately 3 degrees, and the misalignment ranges from 60X to 140X the Debye length (See PG.5 and PG.6 from table 1 in [18]). The optical approximation gives a heaviside function for the heat flux, where the step occurs at the PFC corner. The ion orbit (gyro-orbit approximation) profile shows a smearing effect that distorts the optical approximation results, creating a redistribution of power from the leading edge onto the PFC top surface. The PIC profile further enhances this power redistribution effect, due to the changing gyro orbit pitch (changing $v_{||}$) which results from the electric field. Chapter 3 will outline gyro orbit algorithms that construct this transfer function between optical power loads and ion orbit power loads, but will neglect the electric field effects.

The profiles in figure 1.16 illustrate that the electric fields do indeed play a role in changing the heat flux distribution across the PFC surface. That being said, the objective of the simulation needs to be considered when analyzing these effects. The intent of the investigation outlined herein is to provide high fidelity heat load predictions *as a means of mapping the PFC engineering limits directly to the*

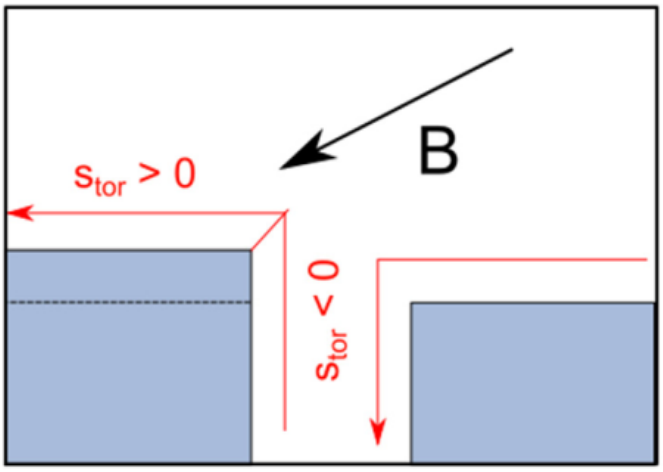


Figure 1.15: Example geometry from PIC simulations performed by Komm [18]. Geometry corresponds to an ITER monoblock that has been misaligned to expose a leading edge.

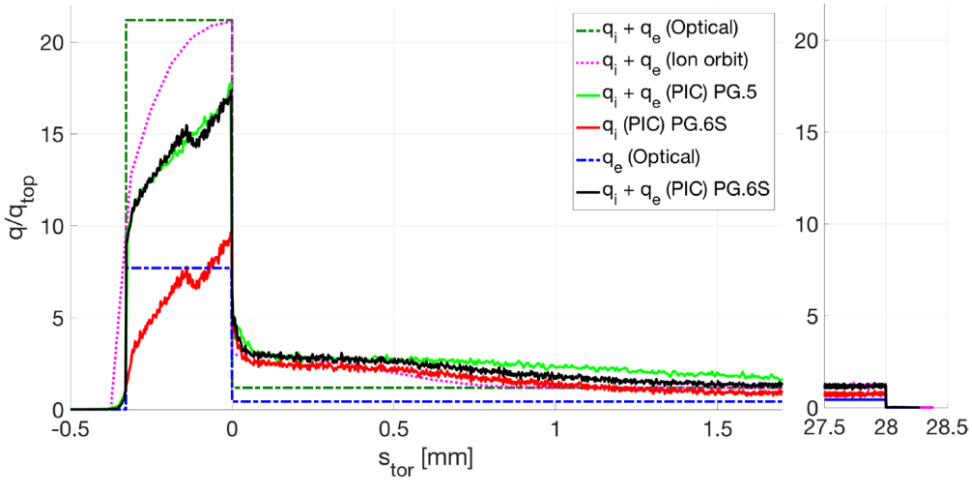


Figure 1.16: Comparison of heat flux profiles for optical, ion orbit (from [24]), and SPICE2 PIC code from [18]. Simulations performed on portruding leading edge geometry, where $s_{tor} < 0$ corresponds to the leading edge and $s_{tor} > 0$ corresponds to the top surface. Including ion orbits transfers power from leading edge to PFC surface. Including electric fields enhances this effect. In both cases, power redistribution area is on the order of millimeters wide.

operational boundary of the tokamak. Arbitrarily high precision is not preferable as it increases computational time and does not change the tokamak operational boundary significantly. For mapping the PFC limits to the tokamak operational domain, the quantity of interest is the temperature in the PFC. When one considers the change in temperature that occurs as a result of incorporating electric fields into the ion trajectory calculation, the change is small (figure 1.14 and [18]). For leading edge cases, neglecting the electric fields actually induces an additional engineering factor of safety into the calculation by slightly overestimating the temperature, which provides for conservative engineering estimates.

Furthermore, the spatial extent of the electric field induced power redistribution must be considered with respect to the simulation domain. In this dissertation, heat loads are calculated on STL meshes as described in chapter 3, where the heat for each element is calculated at the triangular mesh centroid. Obtaining reasonable (<10% error) optical power balance can be achieved on triangular mesh elements whose maximum edge lengths are 5mm, corresponding to a distance of 2-3mm from a centroid to the nearest neighbor mesh element. For the cases when heat load predictions are needed between experimental discharges, the mesh resolution is set to the minimum level required to achieve power balance, in an effort to accelerate the calculation. The aforementioned redistribution of power that arises due to electric field effects happens on length scales ranging from 100s of um to a few mm, as seen in figures 1.16 and 1.12. Because this is oftentimes smaller that the STL resolution, the power can be redistributed from and to the same mesh element when tracing these gyro orbits. In order to simulate the wrapping effect, the STL mesh must be on the order of the length scale of power redistribution for corners and edges, or a gyro radius for flat surfaces. Failure to use a high enough mesh resolution will simply reproduce the optical approximation, even if the gyro orbit approximation or electric fields were turned on.

For inter-shot simulations, running a full ion gyro orbit simulation at the desired micro-meter resolution may not be feasible due to the computation time required. The optical approximation, however, can be simulated in a few minutes as discussed in chapter 3. The optical approximation also tends to generate higher temperatures and heat loads than the gyro orbit approximation or the full electric field treatment, due to the hard shadow boundaries (ie no shadow smearing due to helical trajectories).

For the inter-shot operational objective of machine protection, predictions that yield higher temperatures are likely to be preferred, as the margin for error is higher with a larger factor of safety. Therefore in some cases it may be desireable to use the optical approximation as it provides the most extreme prediction, resulting in a more conservative operational set point.

1.3.5 Kinetic Theory Distributions

While the equations in section 1.3.3 govern the motion of individual particles, it is often convenient to use a kinetic theory approach to quantify macroscopic properties of the plasma. As is typical in kinetic theory, the problem amounts to finding a statistical distribution of the particles as a function of seven variables, $(x, y, z, t, v_x, v_y, v_z)$. In this work we consider the velocity distribution to be Maxwellian. This is sufficient for the first investigations into gyro orbit heat loading, but is not accurate in some situations such as in low collisionality plasmas or when particles are monoenergetic. In these cases, a new functional form of the parallel and perpendicular velocities would be necessary to obtain the true particle velocity distribution. For the results contained herein, it is assumed that the plasma velocity distribution is Maxwellian. Future work will aim to correct this assumption and include the various sources that result in a non-maxwellian velocity distribution. Additionally, future work will incorporate the effects of electric fields, dynamic magnetic fields, drifts, and the plasma sheath, as was discussed in the previous section.

This section follows the derivations and descriptions that can be found in the textbook by Chen [26], as well as derivations by Maxwell that can be found in his 1860 publication [27]. The number of particles per cubic meter, centered at a position $\vec{r}$ at some time t with velocity components within a range between (v_x, v_y, v_z) and $(v_x + dv_x, v_y + dv_y, v_z + dv_z)$ can be written via an integral,

$$n(\vec{r}, t) = \int_{-\infty}^{\infty} dv_x \int_{-\infty}^{\infty} dv_y \int_{-\infty}^{\infty} dv_z f(\vec{r}, \vec{v}, t) \tag{1.38}$$

$$= \int_{-\infty}^{\infty} f(\vec{r}, \vec{v}, t) d^3v \tag{1.39}$$

The function $f(\vec{r}, \vec{v}, t)$ can be normalized such that the integral evaluates to 1, making it a probability distribution,

$$\int_{-\infty}^{\infty} \hat{f}(\vec{r}, \vec{v}, t) d^3v = 1 \tag{1.40}$$

The function $\hat{f}$ is the velocity probability distribution function. Assuming that the macroscopic behavior is steady state and homogoneous in the region of interest reduces the number of variables to $f(v_x, v_y, v_z)$. The normalized Maxwellian distribution for velocity can be derived from statistical thermodynamics, and is written as,

$$\hat{f}_v = (\frac{m}{2\pi kT})^{3/2} \exp\left(\frac{-mv^2}{2kT}\right) \tag{1.41}$$

where m is the mass of the species, k is the Boltzmann constant, T is the temperature, and v is the magnitude of the velocity vector. Because of the spherical symmetry in velocity space, the Maxwellian can be integrated over solid angle, resulting in a probability distribution function of speeds,

$$\hat{f}_S = \int_{-\infty}^{\infty} d\Omega (\frac{m}{2\pi kT})^{3/2} \exp\left(\frac{-mv^2}{2kT}\right) \tag{1.42}$$

$$= 4\pi v^2 (\frac{m}{2\pi kT})^{3/2} \exp\left(\frac{-mv^2}{2kT}\right) \tag{1.43}$$

where Ω is the solid angle differential, $\sin\theta d\theta d\phi$. The Maxwellian distribution for the momentum can be found by using $\vec{p} = m\vec{v}$,

$$\hat{f}_v d^3v = \hat{f}_v \left(\frac{dv}{dp}\right)^3 d^3p = \hat{f}_v \frac{1}{m^3} d^3p = \hat{f}_p d^3p \tag{1.44}$$

$$\rightarrow \hat{f}_p = \left(\frac{1}{2\pi mkT}\right)^{3/2} \exp\left(\frac{-p^2}{2mkT}\right) \tag{1.45}$$

where p is the magnitude of the momentum. Similarly, the energy probability distribution function can be found from the momentum probability distribution by using the relation for energy, $E = \frac{p^2}{2m}$, and exploiting the spherical symmetry by

integrating over solid angle,

$$\hat{f}_p d^3 d\vec{p} = \hat{f}_p 4\pi p^2 dp = \hat{f}_p 4\pi p^2 \frac{dp}{dE} dE = \hat{f}_p 4\pi m\sqrt{2mE} dE = \hat{f}_E dE \tag{1.46}$$

$$\rightarrow \hat{f}_E = 2\sqrt{\frac{E}{\pi}} \left(\frac{1}{kT}\right)^{3/2} \exp\left(\frac{-E}{kT}\right) \tag{1.47}$$

These distribution functions will be referenced later in this dissertation during a discussion of the computational framework devised to simulate the ion gyro motion heat loads.

1.4 2D Magnetic Equilibria

For the investigations present in this dissertation, a toroidally axisymmetric plasma is considered. Per the discussion in [28], simulating an ideal 2D axisymmetric plasma incident upon 3D PFCs eliminates the complexity associated with toroidal variations in the magnetic field, but still captures the response of the 3D geometry. Unless the plasma axisymmetry is significantly perturbed, the 2D plasma approximation is sufficient. For the cases when axisymmetry is perturbed, future work will investigate the change in heat flux that arises from toroidal variation in magnetic field. This section describes the physics relevant to axisymmetric 2D magnetohydrodynamic (MHD) equilibria.

1.4.1 Coordinate Transformations

To begin, let's establish the relevant coordinate transformations. A vector, $\vec{x} = (x, y, z)$, can be described in a cartesian system, or by transforming to a system more convenient to toroidal geometry, the cylindrical coordinate system. Defining the toroidal angle ϕ at any arbitrary location in the tokamak to be,

$$\phi = \text{arctan2}\left(\frac{y}{x}\right) \tag{1.48}$$

the transformation between cartesian coordinates and the tokamak system can be defined by a transformation, $\vec{x}\mathbf{T} = \vec{r}$,

$$\begin{bmatrix} x & y & z \end{bmatrix} \begin{bmatrix} \cos(\phi) & \text{-sin}(\phi) & 0 \\ \sin(\phi) & \cos(\phi) & 0 \\ 0 & 0 & 1 \end{bmatrix} = \begin{bmatrix} r & t & z \end{bmatrix} \tag{1.49}$$

$$\tag{1.50}$$

We can now define the tokamak system magnetic field vector, $\vec{B}$, whose origin is located at a position with toroidal angle ϕ, in terms of its cartesian components,

$$\begin{bmatrix} B_x & B_y & B_z \end{bmatrix} \begin{bmatrix} \cos(\phi) & \text{-sin}(\phi) & 0 \\ \sin(\phi) & \cos(\phi) & 0 \\ 0 & 0 & 1 \end{bmatrix} = \begin{bmatrix} B_r & B_t & B_z \end{bmatrix} \tag{1.51}$$

$$\tag{1.52}$$

This picture can be further simplified by considering the vector components that are orthogonal to, and in the plane of, a 2D plane of constant ϕ. The component of the magnetic field orthogonal to the ϕ plane is called the toroidal component, B_t. The component of magnetic field that lies in the ϕ plane is called the poloidal component, B_p. This can be state mathematically,

$$\vec{B_p} = \vec{B_r} + \vec{B_z} \tag{1.53}$$

$$\vec{B_t} = \vec{B_t} \tag{1.54}$$

$$B_p = |\vec{B_p}| = \sqrt{B_r^2 + B_z^2} \tag{1.55}$$

$$B_t = |\vec{B_t}| \tag{1.56}$$

$$\tag{1.57}$$

1.4.2 Flux Surfaces

MHD equilibrium is achieved by a radial pressure balance and toroidal force balance [22]. Like a gas, a hot plasma will try to expand along the minor axis and uniformly fill its container, resulting in a radially (along the minor axis) directed pressure. Additionally, the toroidal geometry of the tokamak will generate a force directed along

the major axis. Both of these forces threaten to push the plasma into the wall and must be compensated for with electromagnetic forces generated by the confinement system. The equations that govern the magnetic equilibrium are a momentum balance equation, as well as Maxwell's equations,

$$\nabla p = \vec{j} \times \vec{B} \tag{1.58}$$

$$\mu_0 \vec{j} = \nabla \times \vec{B} \tag{1.59}$$

$$\nabla \cdot \vec{B} = 0 \tag{1.60}$$

where p is the pressure, $\vec{j}$ is the current density, μ_0 is the permeability of free space, and $\vec{B}$ is the magnetic field. From these equations it is evident that the current and magnetic field lie on surfaces of constant pressure. Many texts use these equations to derive the Grad-Shafranov equation, and this derivation will not be included here. Chapter 6 from [21] or chapter 11 from [22] are both good sources for the Grad-Shafranov derivation. Because the magnetic field exists on surfaces of constant pressure, and because in a tokamak there is a toroidal plasma current, the surfaces of constant pressure will appear to be nested, as shown in figure 1.17, a real magnetic equilibrium reconstruction from NSTX-U.

Per Stacey [21], we can create a variable that represents these constant flux surfaces, ψ, where $2\pi\psi$ is the poloidal magnetic flux passing through a plane between the magnetic axis and the flux surface,

$$2\pi\psi = \int_S \vec{B}_p \cdot dS \tag{1.61}$$

where S represents the surface in question and B_p represents the poloidal field.

The poloidal flux between a surface, ψ and an adjacent surface, $\psi + d\psi$, can be calculated via the integral,

$$2\pi(\psi + d\psi) = \int_{S+\delta S} \vec{B}_p \cdot dS - \int_S \vec{B}_p \cdot dS \approx B_p 2\pi R dr \tag{1.62}$$

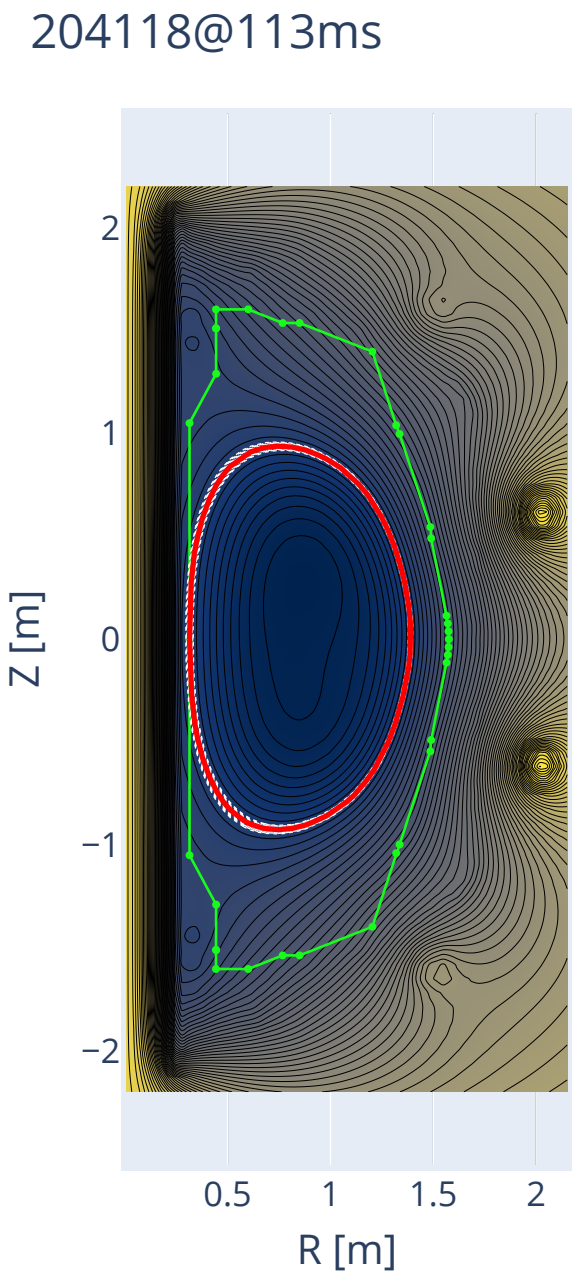


Figure 1.17: Magnetic equilibrium reconstruction for NSTX-U discharge 204118 at 113ms. Contours of constant pressure, or poloidal flux, can be seen to be nested inside the core. The last closed flux surface (LCFS) overlaid in red. Simplified wall geometry overlaid in green.

where R is the major radius of the flux surface. From this result we can define the poloidal field as a function of poloidal flux,

$$RB_p = \frac{d\psi}{dr} = |\nabla\psi| \tag{1.63}$$

$$\rightarrow B_p = \frac{|\nabla\psi|}{R} \tag{1.64}$$

which can be combined with equation 1.53 to yield the radial and vertical components of magnetic field as a function of ψ,

$$B_r = \frac{d\psi}{dz}\frac{1}{R} \tag{1.65}$$

$$B_z = -\frac{d\psi}{dr}\frac{1}{R} \tag{1.66}$$

$$B_t = B_t \tag{1.67}$$

where the negative sign in the definition of B_z preserves the right handed coordinate system. This result implies that if the 2D profile of ψ is known everywhere in the plane of constant ϕ, then the tokamak system components of magnetic field can be determined. A standard file format for saving this ψ profile is known as a GEQDSK file, and its format will be discussed in the next section. The last variable of interest in this section is ψ_N, which is a normalized poloidal flux variable defined as

$$\psi_N = \frac{\psi - \psi_{axis}}{\psi_{lcfs} - \psi_{axis}} \tag{1.68}$$

where ψ_{axis} is the poloidal flux at the magnetic axis and ψ_{lcfs} is the poloidal flux at the last closed flux surface (LCFS).

1.4.3 EFIT and GEQDSK Files

A method for reconstructing the magnetic equilibrium from magnetic diagnostic data was originally provided in the work by Lao [29]. The aim of this section is not to provide the physical derivation in that publication, but rather to provide a reference for the formatting of GEQDSK files, which is file format that contains information about the magnetic reconstruction. It should be noted that the GEQDSK file is not

standardized by a rigorous standard, and as such can differ between organizations. The definitions included in this work reflect those defined at General Atomics by Lao, and may differ from other groups.

The right handed coordinate system (R, ϕ, Z) is used, and due to toroidal axisymmetry the toroidal coordinate, ϕ, can be ignored. The tokamak is divided into a uniform grid of R,Z points. The GEQDSK file, also called a gFile, starts with a single line that has an optional comment, and then three integer values. The first integer corresponds to a case number, the second corresponds to the number of points in the R direction, and the third corresponds to the number of points in the Z direction. All floating point numbers in the GEQDSK file are provided in scientific 'e' notation, where the exponents are always displayed with a sign and two digits (i.e. e+02). The floating point numbers have a single value to the left of the radix and 8 values to the right. Positive values are padded with a space, and negative values are not padded and have a negative sign. The maximum number of floating point columns on any row is five.

```
1 EFITD  05/19/2021   #186902 2300           3  65  65
2  1.70000005e+00 3.20000005e+00 1.69550002e+00 8.39999974e-01 0.00000000e+00
3  1.78785408e+00-3.25682424e-02-2.47078255e-01 1.19898552e-02-1.85712302e+00
4  9.90306312e+05-2.47078255e-01 0.00000000e+00 1.78785408e+00 0.00000000e+00
5 -3.25682424e-02 0.00000000e+00 1.19898552e-02 0.00000000e+00 0.00000000e+00
```

An example of the first five lines from DIII-D shot 186902 @ 2300 ms is provided above. The first line has the comment 'EFITD 05/19/2021 #186902 2300', after which the integers 3, 65, 65, follow, corresponding to case 3, 65 points in the R direction: nw, and 65 points in the Z direction: nh. Lines 2-5 in this example gFile have scalar quantites. Values that are not scalar can either be 1D or 2D profiles. The 1D profiles are functions of a monotonic chord of poloidal flux extending from the magnetic axis to the seperatrix, ie $f(\psi)$, uniformly distributed along nw points. The 2D R,Z profiles are functions of position on the spatial grid, ie f(R[i],Z[j]), and are flattened. The order of the flattening is rows first, columns second, ie [0:nw,0] ... [0:nw,1] ... [0:nw,nh], etc. Additionally, there are several 2D contours defined in the GEQDSK, such as

Table 1.1: Order of GEQDSK variables in text file

Variable Name	Size(i,j)
case	Variable
nw	(1)
nh	(1)
rdim	(1)
zdim	(1)
rcentr	(1)
rleft	(1)
zmid	(1)
rmaxis	(1)
zmaxis	(1)
simag	(1)
sibry	(1)
bcentr	(1)
current	(1)
simag	(1)
xdum	(1)
rmaxis	(1)
xdum	(1)
zmaxis	(1)
xdum	(1)
sibry	(1)
xdum	(1)
xdum	(1)
fpol	(1,nw)
pres	(1,nw)
ffprim	(1,nw)
pprime	(1,nw)
psirz	(nw,nh)
qpsi	(1,nw)
nbbbs	(1)
limitr	(1)
(rbbbs, zbbbs)	(2, nbbbs)
(rlim, zlim)	(2, limitr)

Table 1.2: GEQDSK file variable names, descriptions, units, and sizes

Name	Description	Units	Size (i,j)
case	Identification String	NA	variable
nw	Number of points in R direction	NA	(1)
nh	Number of points in Z direction	NA	(1)
rdim	Radial size of computation box	m	(1)
zdim	Vertical size of computation box	m	(1)
rcentr	R of vacuum toroidal magnetic field bcentr	m	(1)
rleft	Minimum R of computation box	m	(1)
zmid	Center Z of computation box	m	(1)
rmaxis	R of magnetic axis	m	(1)
zmaxis	Z of magnetic axis	m	(1)
simag	Poloidal flux, ψ, at magnetic axis	Weber/rad	(1)
sibry	Poloidal flux at seperatrix	Weber/rad	(1)
bcentr	Vacuum toroidal magnetic field at rcentr	T	(1)
current	Plasma current, I_p	A	(1)
fpol	Poloidal current function $F = RB_t$	m·T	(1,nw)
pres	Plasma pressure	N/m^2	(1,nw)
ffprim	$F \cdot F'(\psi)$	$(m \cdot T)^2/(Weber/rad)$	(1,nw)
pprime	$P'(\psi)$	$(N/m^2)/(Weber/rad)$	(1,nw)
psirz	Poloidal flux at R,Z grid points	Weber/rad	(nw,nh)
qpsi	Safety factor, $q(\psi)$	NA	(1,nw)
nbbbs	Number of LCFS contour points	NA	(1)
limitr	Number of limiter contour points	NA	(1)
rbbbs	R of LCFS points	m	(1,nbbbs)
zbbbs	Z of LCFS points	m	(1,nbbbs)
rlim	R of limiter contour points	m	(1,limitr)
zlim	Z of limiter contour points	m	(1,limitr)

the contour of the machine limiter surface and last closed flux surface (LCFS). These contours are also flattened such that the order is R[0] Z[0] R[1] Z[1] ... R[N] Z[N], where N is the number of contour points. The order of the variables in the GEQDSK format is provided in table 1.1 (note that some variables are printed multiple times). The variables in the GEQDSK are defined in table 1.2.

1.4.4 Power Sharing Between Divertors

Many future reactor class tokamak designs require operation in a double null configuration, consisting of power sharing between four divertors: upper outer, upper inner, lower outer, lower inner [30, 31, 32]. Double null operation enables higher performance operation and reduced machine size, so it is an attractive operating space for future machines. It should be mentioned that while double null operation may appear attractive, there has been no RMP ELM suppression for double null plasmas thus far. Additionally, the empirical scalings of heat flux width will likely need to be modified to account for the difference in attachment levels and particle transport on the inboard vs outboard sides of the machine.

Work on DIII-D on the early 2000s investigated the power sharing between divertors for double null plasmas, resulting in the establishment of a parameter, δR_{sep}, used to quantify the degree of divertor imbalance [33]. More specifically, δR_{sep} is defined as the radial distance between the upper divertor separatrix and the lower divertor separatrix at the outboard midplane. After varying the value of δR_{sep} and recording the associated peak heat fluxes at the upper and lower divertors, the data was fit to a hyperbolic tangent function as illustrated in figure 1.18. As can be observed in the figure, there is a slight offset in x-axis zero crossing, corresponding to a δR_{sep} of approximately 0.25cm, which was attributed to $E \times B$ poloidal drifts. Additionally, power distributed to the outboard divertor was calculated to generally be 2.5 times greater that the power distributed to the inner divertor across the entire range of δR_{sep}, yet in the extreme case approached 10.

More recently, an investigation into double null power sharing on Alcator C-Mod was performed [34]. C-Mod performed a variety of L, I, and H mode discharges where the strike point was swept across langmuir probe arrays in all four divertor

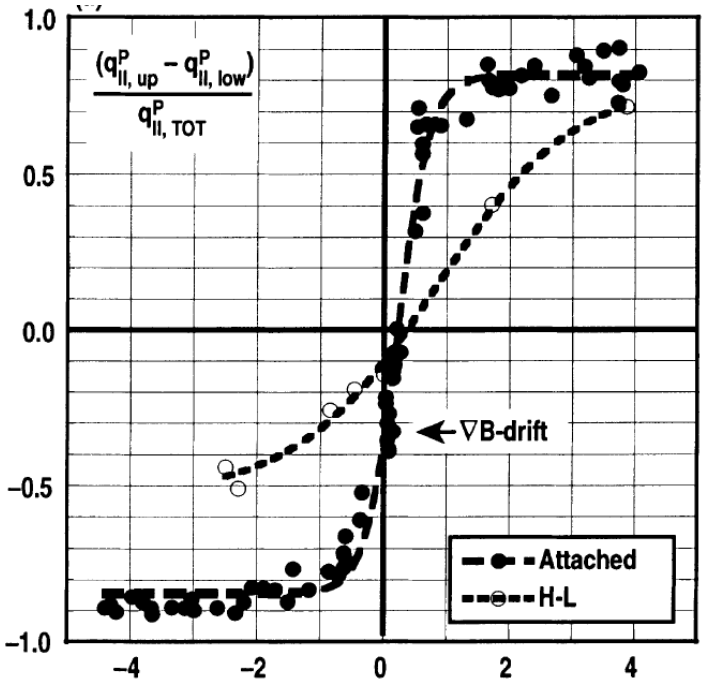


Figure 1.18: Upper/Lower power sharing as a function of δR_{sep} from [33].

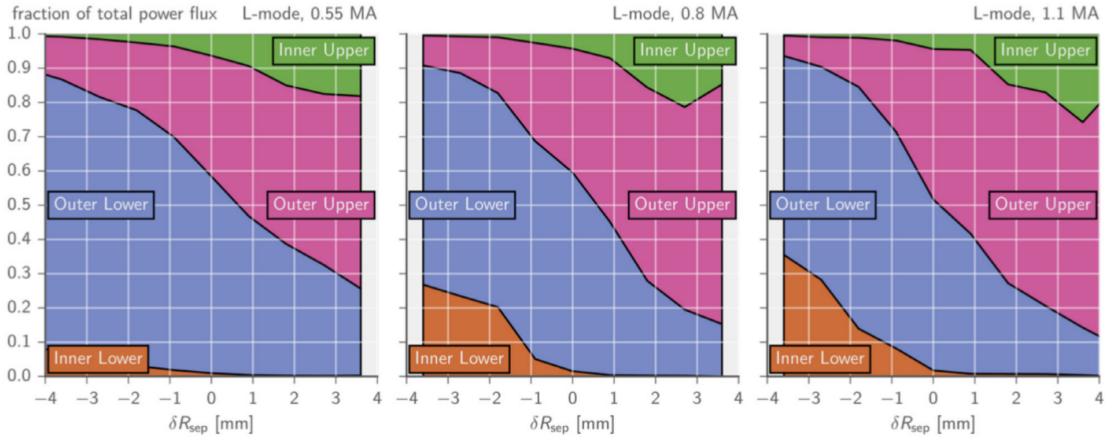


Figure 1.19: Divertor power sharing as a function of δR_{sep} from [34].

locations. Sweeping the strike point enabled reconstruction of the plasma profile across the divertor. For the C-Mod study, δR_{sep} was defined similarly to the DIII-D study mentioned above, and was maintained at a costant value for the duration of each discharge. Various plasma parameters were measured for varying plasma currents during a L-mode discharge, as a function of δR_{sep}. The inner divertor heat flux is relatively symmetric about the strikepoint, while the outer divertor heat flux is asymmetric with an exponential decay occuring in the common flux region. A double e-folding can be observed in the density and temperature profiles, but it is less pronounced in the heat flux profile.

In order to determine the power sharing beween divertors, the measured heat flux profiles were mapped to the midplane, binned in steps of δR_{sep}, and then integrated to find the total power. The relative fractions of distributed power are then calculated as each divertor's integrated power divided by the total integrated power. Figure 1.19 illustrates the results from this calculation. The results indicate that the inner divertors receive a small fraction of the total power when operating in balanced double null ($\delta R_{sep} = 0$) configurations, and δR_{sep} must exceed $\pm$ 1-2 mm before substantial increases in inboard power deposition is observed [34]. Increasing plasma current results in steeper power sharing curves. The DIII-D results of an outer to inner power deposition ratio of approximately 2.5 is confirmed [33].

After determining the fractional power deposition profiles, Brunner developed analytical functions to describe the power sharing between divertors. The following power sharing relationships were developed,

$$\frac{P_{ol}}{P_{ol} + P_{ou}} = \frac{1}{1 + e^{\delta R_{sep}/\lambda_0}} \tag{1.69}$$

$$\frac{P_{ou}}{P_{ol} + P_{ou}} = \frac{1}{1 + e^{\delta R_{sep}/-\lambda_0}} \tag{1.70}$$

$$\frac{P_{il}}{P_{il} + P_{iu}} = \frac{1}{1 + e^{\delta R_{sep}/\lambda_i}} \tag{1.71}$$

where the subscripts ol, ou, and il, correspond to outer lower, outer upper, inner lower, respectively, λ_0 represents the heat flux width on the outboard side, and λ_i represents the inboard heat flux width. These functional forms can be used to reproduce the hyperbolic tangent dependency of power sharing as was described on DIII-D [33].

1.5 Other Related Topics

This section provides some reference to topics that are related to this dissertation, but not mandatory to understanding the analysis included herein. Future work may incorporate these effects in the heat flux predictions.

1.5.1 Coil Misalignments and Error Fields

Non-axisymmetric magnetic fields can be unintentionally created by misalignments in magnets or tiles. These fields are collectively called error fields. Coil misalignments can lead to braking and locking, which can degrade machine performance [35, 36]. Additionally, small tile misalignments can have dire consequences for PFCs, as demonstrated via intentional tungsten tile misalignment on JET that resulted in tile melting and sputtering [37].

Ferraro et al. analyzed the effects of error fields with respect to the divertor heat fluxes on NSTX-U [36]. Error fields can cause toroidal variations in the angle between magnetic field lines and divertor tile surface normals, also called the pitch angle. These asymmetries can result in localized increases in heat flux deposition to the PFCs. The perturbations to the pitch angle were calculated using M3D-C1 to solve for the perturbation equilibrium in three dimensions, and then TRIP3D was utilized to simulate the divertor footprint as shown in figure 1.20. Localized changes in normalized poloidal flux, ψ_N, are shown for a poloidal field shift as compared to a toroidal field tilt. The distortion of the separatrix lobes in the right plot in figure 1.20 demonstrates that divertor footprint changes significantly between the two cases.

1.5.2 Edge Localized Modes

Operation of tokamaks in H-Mode is accompanied by short bursts of plasma ejected across the transport barrier [38]. The enhanced confinement associated with H-mode results in a steep pressure gradient at the pedestal edge, which increases until an MHD instability forms [39]. The instability creates a 'pressure relief valve' in the aforementioned pressure gradient, and plasma spews through into the SOL and eventually onto the divertor targets. After the pressure in the core has been reduced,

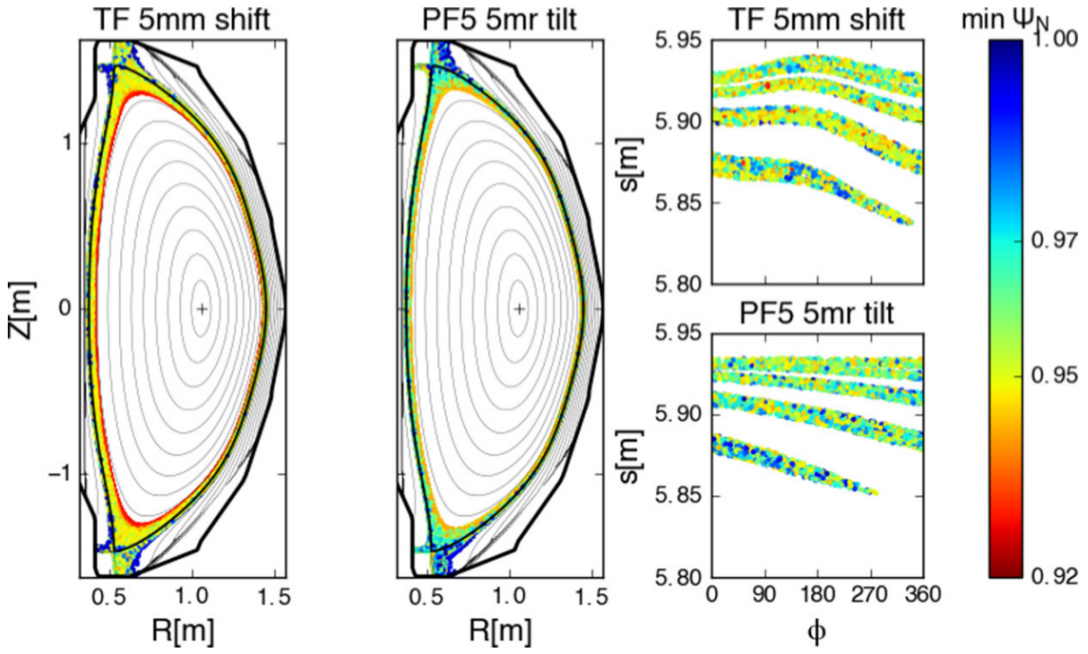


Figure 1.20: Left: Poincaré plot and Right: lower horizontal divertor footprint due to error fields in toroidal field coil and poloidal field coil 5 from [36]. Colorbar is normalized poloidal flux. Aspect ratio is not to scale, as figure has been copied and pasted from [36].

the metaphorical 'valve' closes itself and the gradient is restored until the pressure once again builds. ELMs can transport massive heat loads to the divertor, and figure 1.21 provides a graphical representation of main plasma energy loss, ΔW, as a function of plasma current, I_p [39].

In order to predict the ELM induced heat loading on the divertor, it is necessary to predict the deposited ELM energy on its surface, ε_{div}. Via a multi-machine regression, Eich predicted the parametric scaling of the parallel ELM energy fluence,

$$\varepsilon_{||} = 0.28 \pm 0.14 \frac{MJ}{m^2} \times n_e^{0.75\pm0.15} \times T_{e,ped}^{0.98\pm0.1} \Delta E_{ELM}^{0.52\pm0.16} \times R_{geo}^{1\pm0.4} \tag{1.72}$$

where n_e is in units of $[10^{20} m^{-3}]$, $T_{e,ped}$ in [keV], ΔE_{ELM} in [%], and R_{geo} in [m] [40]. To calculate the heat load onto the divertor tile, one takes equation 1.72 and calculates the dot product with the target surface normal vector,

$$\varepsilon_{||} = \varepsilon_{div} \cdot \sin^{-}1(\alpha) = \varepsilon_{div} \frac{B_\phi}{B_\theta} \cdot f_x \tag{1.73}$$

where α represents the angle between incident field line and divertor surface normal, and f_x represents the flux expansion.

1.5.3 0D Heat Flux Approximation

This section follows chapter five from the book by Stangeby [4]. The simplest divertor heat flux model is 0D, and relates a temperature at a location upstream in the SOL, T_u, to the temperature at the divertor target, T_t. For simplicity, one can ignore the effects of toroidal magnetic curvature and represent the SOL as a simple 1D plasma flow. Figure 1.22 illustrates this 'straightening' of the SOL. Ignoring neoclassical (toridicity) effects is valid so long as there is a high level of collisionality in the SOL. Additionally, we assume that the upstream location is taken halfway between the divertor targets, that there is a single X point (single null), and that electrons and ions are thermally coupled.

In addition to the aforementioned assumptions, each flux tube is assumed to have local conservation of particles. Consequentially, it is assumed that there is no cross field diffusion or flow reversal, so that the plasma flow is 1D. Pressure balance is assumed

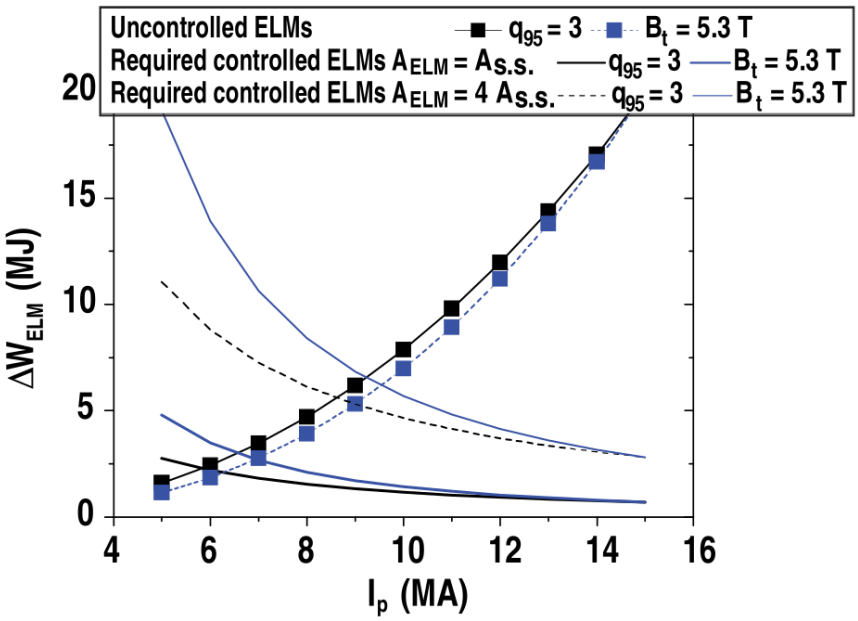


Figure 1.21: Main plasma ELM energy loss as a function of plasma current for unmitigated ELMs from [39].

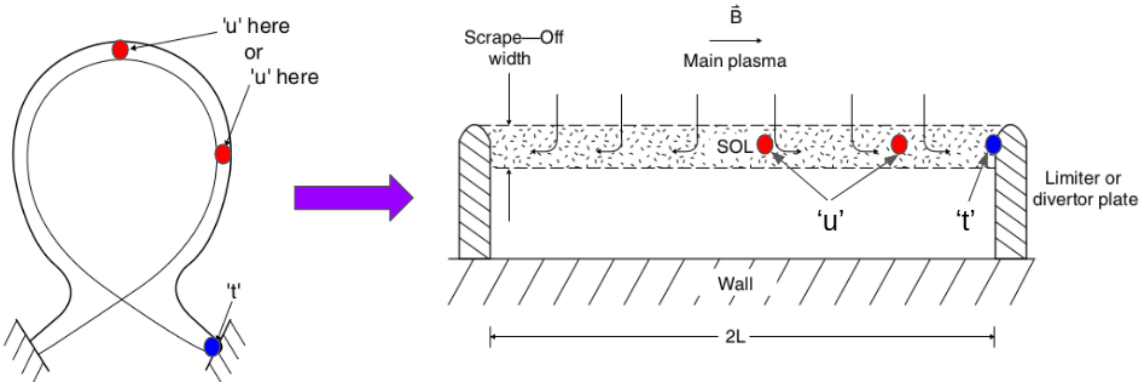


Figure 1.22: Simple model of SOL adapted from [4].

to exist between the static and dynamic pressures. Mathematically this means,

$$p + nmv^2 = \text{constant}. \tag{1.74}$$

where nmv^2 is the dynamic pressure. Application of the Bohm criterion, $v \geq c_s = [(kT_e + \gamma kT_i)/m_i]^{1/2}$, to determine the target velocity yields the following relationship between upstream and target densities and pressures,

$$2n_t T_t = n_u T_u \tag{1.75}$$

It can be assumed that the entirety of the heat flow in the SOL is carried by conduction from the upstream location to the target. This yields the following relationship between upstream and target temperatures,

$$T_u^{7/2} = T_t^{7/2} + q_{||}\frac{7L}{2\kappa_{0e}} \tag{1.76}$$

where $q_{||}$ is the parallel power flux density [$\mathrm{MW\,m^{-2}}$], L is the distance between upstream and target locations, and κ_{0e} is the electron parallel conductivity coefficient. Assuming no power sources or sinks in the SOL, the parallel heat flux density can be written as,

$$q_{||} = q_t = \gamma n_t k T_t c_{st} \tag{1.77}$$

where γ is the heat transmission coefficient. Generally, the upstream density and power are predefined as tokamak operator inputs, while upstream temperature, target density, and target temperature are unknown. Via some subsitutions, the target temperature can be expressed as,

$$T_t = \frac{m_i}{2e}\frac{4q_{||}^2\left(\frac{7q_{||}L}{2\kappa_{0e}}\right)^{-4/7}}{\gamma^2 e^2 n_u^2} \tag{1.78}$$

whereas the target density can be expressed as,

$$n_t = \frac{n_u^3}{q_{||}^2}\left(\frac{7q_{||}L}{2\kappa_{0e}}\right)^{4/7}\frac{\gamma e^2}{2m_i}. \tag{1.79}$$

For the two point model, the heat flux incident upon the divertor tile is solely a function of the upstream conditions in that particular flux tube. Because this simple model depends upon power being emitted from the main plasma into the SOL, a parameter of interest is the heat flux width, λ_q.

1.5.4 0D Heat Flux Width

The heat flux width provides information regarding the decay length of heat that is diffusing across the separatrix into the SOL. This process is caused by perpendicular diffusive transport across the LCFS, so the principles of diffusion can be employed to solve for the variation in heat flux across the SOL. Because of the conservation of particle assumption, all power that diffuses across the seperatrix is transported to the divertor plate. Fick's law describes perpendicular diffusion of particles across the separatrix,

$$q = q_{||} = -X_{\perp}\frac{\partial n_u}{\partial x} \tag{1.80}$$

$$\frac{\partial n}{\partial t} = X\frac{\partial^2 c}{\partial x^2} + Q \tag{1.81}$$

where x is the radial direction (from core) orthogonal to the LCFS, Q is the source term ($Q = 0$ here because of no SOL source assumption), and X is the diffusion coefficient. Using the boundary conditions that particles are not transported very far from the separatrix, $n(x = \infty, t) = 0$, and that initially there are no particles in the SOL, $n(x, t = 0) = 0$, one can convert Fick's second law from a partial differential equation to an ordinary differential equation. Using the substitution, $\theta = n/n_0$ and the similarity transformation,

$$\zeta = \frac{x}{2\sqrt{X\tau_E^{SOL}}}, \tag{1.82}$$

one can convert 1.81 to,

$$\frac{\mathrm{d}^2\theta}{\mathrm{d}\zeta^2} + 2\zeta\frac{\mathrm{d}\theta}{\mathrm{d}\zeta} = 0. \tag{1.83}$$

Using the aforementioned boundary conditions and integrating twice yields,

$$n(x, t) = n_0 \mathrm{erfc}\left(\frac{x}{2\sqrt{X\tau}}\right) \tag{1.84}$$

where τ_E^{SOL} is the scrape off layer energy confinement time. The estimate $x = \sqrt{X\tau_E^{SOL}}$ is useful because erfc(0.5) = 0.5. We are specifically interested in the diffusion length of the heat flux, or the heat flux width λ_q. Setting $x = \lambda_q$ and using the erfc(0.5) = 0.5 estimation yields the following heat flux width,

$$\lambda_{q||} \approx \sqrt{X_{bot}\tau_E^{SOL}}. \tag{1.85}$$

The relationship between the SOL energy confinement time and the power entering the separatrix, P_{SOL}, is

$$P_{SOL} = \frac{3n_u k T_u V_{SOL}}{\tau_E^{SOL}}. \tag{1.86}$$

One finally arrives at the power decay length by using equations 1.76, 1.78, 1.79, 1.77 1.85, and 1.86,

$$\lambda_{q||} = \frac{8^{5/9}\pi^{4/3}}{7^{5/9}}(en_u^{sep}X_\perp^{SOL})^{7/9}\left(\frac{L}{\kappa_0}\right)^{2/9}\left(\frac{B_\theta}{B}\right)^{-2/9}a^{7/9}R^{5/9}. \tag{1.87}$$

Similarly, one may derive the parallel power flux density as,

$$q_{||} = 7^{5/9}8^{-11/9}\pi^{-7/3}\left(\frac{P_{SOL}^2}{en_u^{sep}X_\perp^{SOL}}\right)^{7/9}\frac{\kappa_0^{2/9}}{L^{2/9}(B_\theta/B)_u^{7/9}R^{14/9}a^{7/9}}. \tag{1.88}$$

For a complete review of the physics related to this section, see the book by Stangeby [4]. As was previously mentioned, this chapter provides a synopsis, not a complete review.

1.5.5 Detached Operation

Power carried to the divertor PFCs in tokamak plasmas is transported via electron conduction and convection [41]. Following a similar approach to sections 1.5.3 and 1.5.4, one may describe the power delivered to the target via electron thermal conduction with,

$$q_{||,cond} = \kappa_0 T_e^{5/2}\frac{\mathrm{d}T_e}{\mathrm{d}s} \tag{1.89}$$

where κ_0 is the electron thermal conductivity [4, 41]. For traditional attached plasmas, conduction is the dominant heat transport mechanism. In detached plasmas where

the target temperature is low, thermal transport is dominated by thermal convection, which can be described by the equation,

$$q_{||,conv} = nv_{||}\left(\frac{5}{2}(T_e + T_i) + \frac{1}{2}m_i v_{||}^2 + I_0\right) \tag{1.90}$$

where $v_{||}$ is the parallel fluid velocity and I_0 is the ionization potential (14.8 eV for D) [41]. To enter a detached regime, one increases the SOL collisionality via gas puffing. The increased collisionality causes higher rate coefficients for photonic radiation, charge exchange, recombination, and inelastic collisions, which reduces the conducted power to the divertor PFC, and replaces it with isotropic photonic radiation. These operating regimes are the likely candidates for a future demonstration reactor, where radiated power will need to account for 80%-90% of the total power exhaust [42].

Figure 1.23 provides a picture of detachment, illustrating the attached, partially detached, and fully detached, regimes [43]. Figure 1.24 provides an experimentally derived plot of heat flux across the divertor before and after partial detachment. The dashed line is the heat flux profile after deuterium gas was utilized to detach the plasma from the target. Because the reduction in heat flux is localized at the strike point, this form of detachment is termed partial detachment. If the entire SOL observed a reduction in heat flux, then the divertor would be fully detached. Both figures illustrate a significant reduction in the peak heat flux, and a simultaneous broadening of the heat flux width.

An empirical model for the criteria for detachment as a function of neutral pressure was developed at ASDEX-U by Kallenbach, where scans were performed to determine the level of detachment as a function of deuterium fueling and nitrogen seeding [43]. An extension to this work resulted in a more generic model used for qualifying detached plasmas [44]. The partially detached onset operation point was determined to be given by,

$$\frac{P_{sep}}{R} = \frac{1}{1.3}p_0(1 + f_z c_z) \cdot \frac{\lambda_q}{0.005m} \cdot \left(\frac{R}{1.65m}\right)^{r_z} \times [\mathrm{MWm^{-}1,\ Pa,\ m}] \tag{1.91}$$

where P_{sep} is the power crossing the separatrix in MW, p_0 is the divertor neutral pressure, c_z and f_z are the impurity fraction and radiative efficiency respectively, λ_q

is the heat flux width, and r_z is an exponent used to account for thermal diffusion of λ_q (similar to S in Eich profile). This criteria represents the minimum requirement for detachment to occur. Additional increases to p_0 will result in increased levels of detachment.

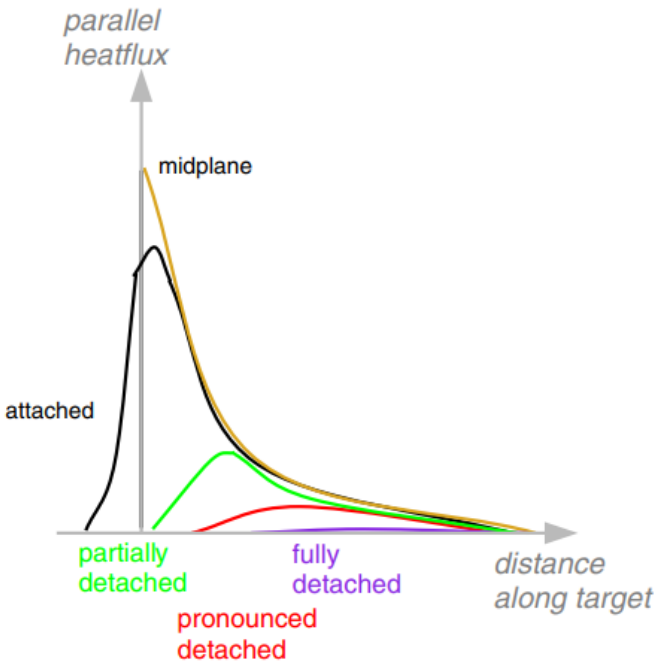


Figure 1.23: Qualitative sketch of detachment levels from [43].

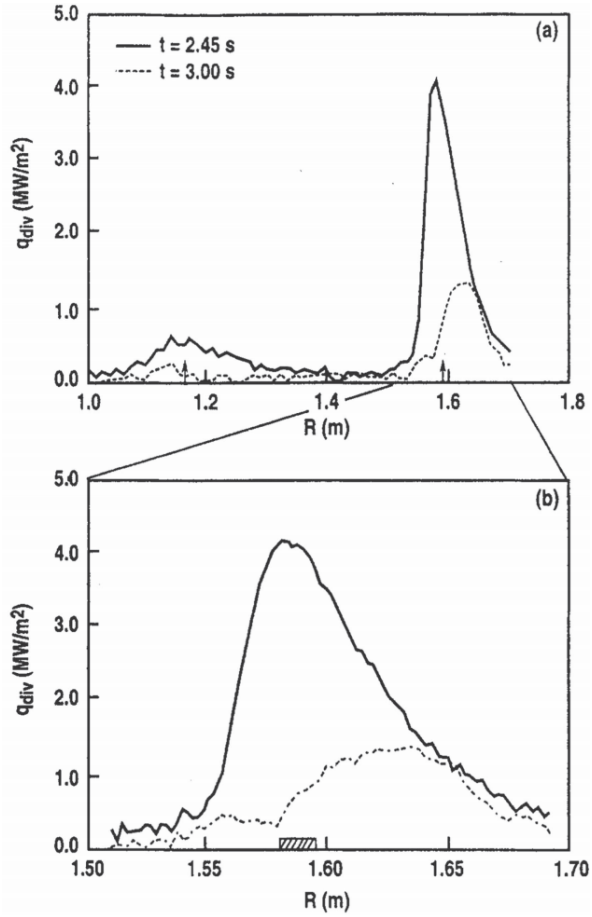


Figure 1.24: Divertor heat flux before and after entering detached operation [41]. A reduction in the peak heat flux can be observed, which is critical for future reactor class devices.

Chapter 2

Modern PFC Heat Loads and Management

The heat loads in tokamaks are some of the most extreme in any engineering domain. For example, the heat flux on the nose cone of a missile is approximately 1 $\mathrm{MW\,m^{-2}}$ [45], whereas the incident heat flux for modern tokamak divertors can exceed 300 $\mathrm{MW\,m^{-2}}$ [9]. Because of these incredible heat loads, reactor class tokamak plasma facing components (PFCs) will operate close to their engineering limits. Designing PFCs that can withstand the harsh environment in the tokamak divertor is a significant engineering challenge. Ultimately, the PFC design will constrain tokamak operations, and must be quantified with high precision to determine the boundary of the operational domain. This chapter discusses the heat loads in modern tokamaks and how they constrain machine performance. An overview of the NSTX-U divertor is then provided, and the heat management techniques employed for its design are highlighted. Chapter 3 discusses a framework designed to simulate the tokamak divertor for high precision heat load predictions.

2.1 Tokamak PFCs Interface to Many Systems

The PFCs in a tokamak serve as the interface between many physical and engineering domains, as shown in figure 2.1. The left boxes (Core, SOL, Sheath) in figure 2.1 provide an illustration of some of the physics domains that are interfaced to the PFC.

Physics that happen in the plasma core will directly impact the PFCs. Some examples are macroscopic MHD equilibrium, radiated power, ELMs, and particle or energy transport across the separatrix. The PFCs can also directly impact the performance in the core via physics that happen in the plasma sheath such as recycling, impurity sublimation, coaxial helicity injection, etc. The power that is radiated on the way to the divertor in the SOL can load PFCs with radiation in detached scenarios, causing changes to the PFC state that can have secondary effects in the core and sheath. Additionally, heat loads that are applied to the PFC will alter its interal state due to temperature dependence of the material properties, which in turn changes the interfaces to each of these physical domains.

The PFCs also provide an interface to many of the engineering systems in the tokamak plant, as illustrated in the right boxes (Vacuum System, Magnet System, Mechanical System, Cooling System) of figure 2.1. The PFCs are in vacuum, so many of the aforementioned physical effects will directly affect the vacuum system, pumping system, neutral gas pumping system, etc. The electromagnetic coils that provide plasma confinement can cause stresses in the PFCs, which can be exacerbated by halo currents coming from the plasma. The PFCs also provide a first wall shield that protects the elecromagnets from some forms of radiation. The PFC movement must be constrained by mechanical fasteners that are part of a greater mechanical system. Changes in the PFC state will induce stresses and strains in the mechanical systems. For tokamaks such as ITER, the tungsten monoblocks will be actively cooled by coolant flowing through the PFCs, which will remove heat but presents challenges for thermal resistance and stress concentrations. Many other engineering systems interface to the PFCs, so must be simulated in parallel to the PFCs to accurately calculate the PFC state. Development of a computational framework that is capable of simulating this interconnectivity with respect to PFC heat loads is described in chapter 3.

2.2 Expected Tokamak Divertor Heat Loads

Modern tokamak divertors will see extreme heat loads when compared to existing machines. Recent heat load projections for ITER indicate that the divertor will see over 20 MW m^{-2} heat fluxes [8]. Figure 2.2 provides a visualization of these expected ITER

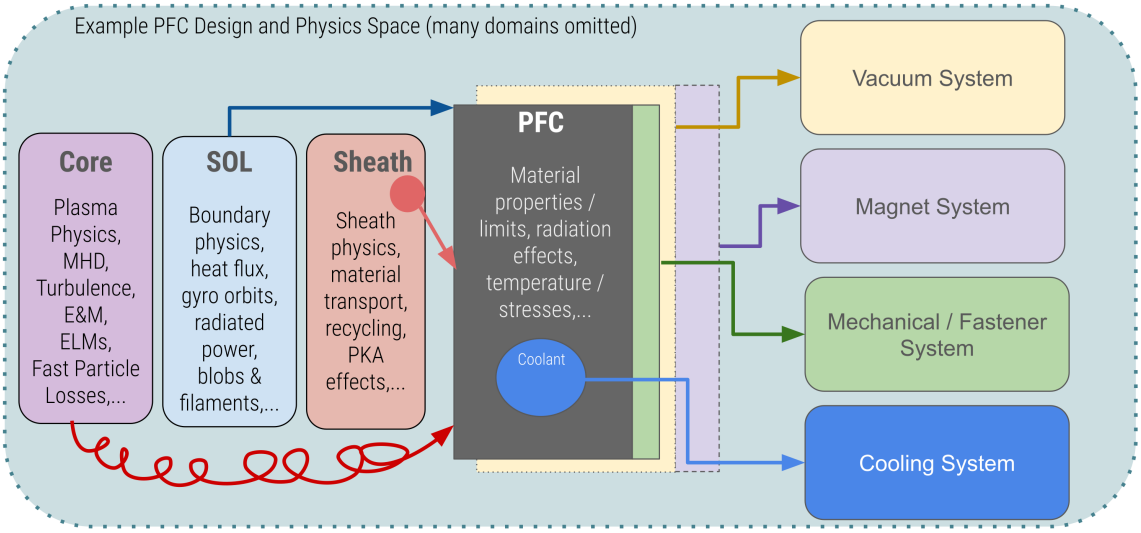


Figure 2.1: Cartoon of some of the systems and physics domains that interface to the PFCs. Physics domains are on the left, engineering domains are on the right. The PFCs provide the interface between many subsystems in the tokamak.

heat loads as a function of neutral pressure, which is a proxy for the level of divertor attachment. Two cases are provided in the figure, the first with an axisymmetric divertor and the second with the true 3D PFC shaping included. The axisymmetric prediction shows heat loads under 20 MW m^{-2}, but when the 3D effects of the PFC geometry are incorporated into the predictions the heat load reaches nearly 25 MW m^{-2} for attached plasmas (low neutral pressure). The ITER tungsten monoblocks have an engineering limit of 16 MW m^{-2} [8]. Sustained heat loads above this level on the ITER CFC-W PFCs will result in tungsten recrystallization, which leads to macro-cracks in the PFC surface, which will likely then lead to subsequent failure modes due to exposed macro-crack leading edges. Because of this engineering limit, the machine operational domain is constrained.

SPARC, the tokamak being constructed by the private company Commonwealth Fusion Systems, will also see incredible divertor heat fluxes. A recent axisymmetric projection for the SPARC divertor puts the heat flux in excess of 300 MW m^{-2} for a 50% radiated power scenario [9]. Figure 2.3 provides a poloidal profile of the heat flux across the divertor target. The figure assumes the heat flux width, λ_q, scales inversely with poloidal field magnitude, as demonstrated in [14], resulting in a λ_q of 0.18 mm. While the material for the SPARC divertor has not been chosen at this time, any material that sees such a high heat flux will likely violate its engineering limits unless properly managed with 3D PFC shaping and appropriate strike point sweeping. This requirement that the strike point location be dynamic places an operational constraint on the tokamak.

2.3 Consequences of Poor PFC Heat Management

In the last section, two examples (ITER and SPARC) of the high heat loads that will be present in modern tokamaks were provided. In both cases, heat management techniques are necessary to ensure the PFCs do not exceed their engineering limits. If these heat loads are not properly managed, then the PFCs will likely either sublimate, melt, or crack due to thermal stresses.

Once the PFC surface reaches a threshold temperature (and pressure), sublimation will begin to occur, where atoms of the bulk material are liberated into the scrape

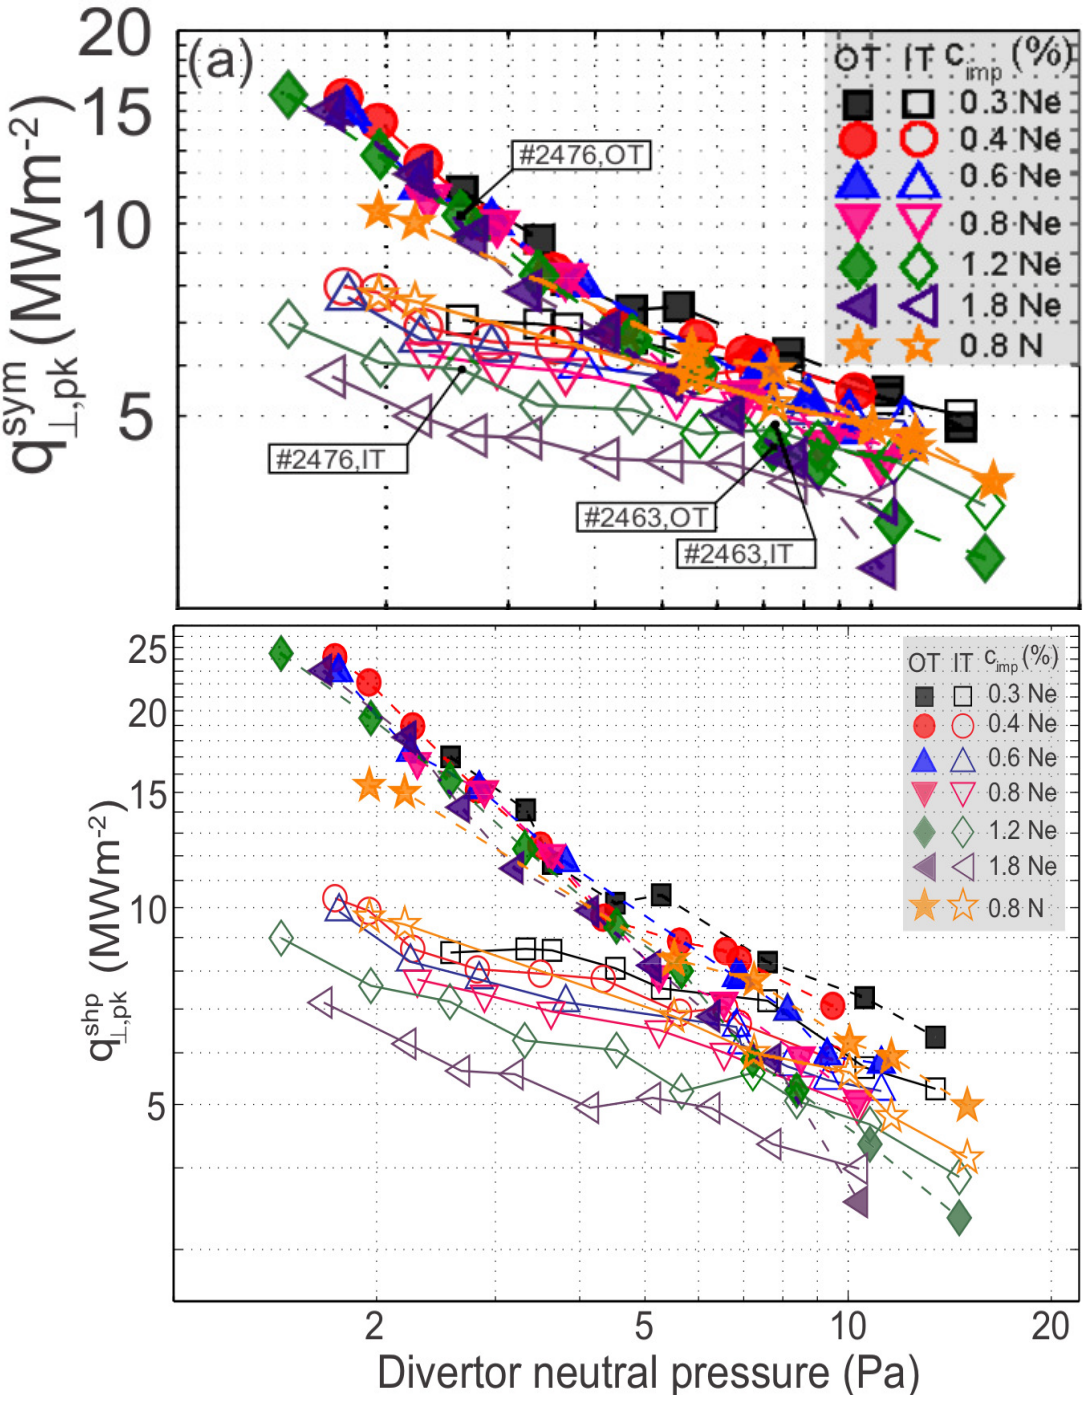


Figure 2.2: ITER divertor heat flux predictions for axisymmetric case (top), and accounting for 3D PFC geometry (bottom) from [8].

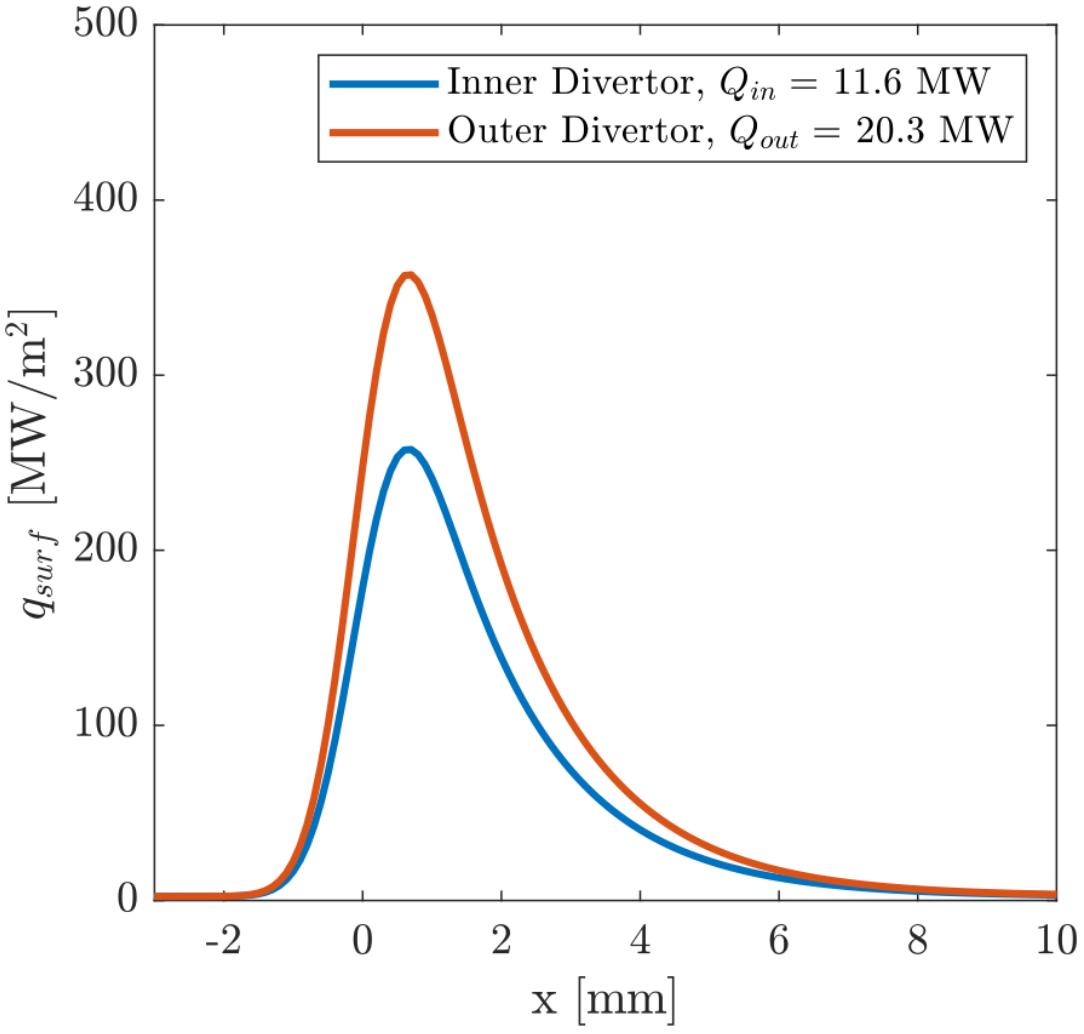


Figure 2.3: Expected heat flux in SPARC from [9] assuming 50% radiated power.

off layer (SOL). Once free of the bulk material, the particles soon become ionized, which allows the electromagnetic forces in the SOL to act upon the particle. These liberated particles can either redeposit in inconvenient locations, or move up the SOL towards the core. If the particle does not move towards the core, it can be redeposited at another location on the PFCs. This redeposition results in minute changes to the PFC geometry, and if a large amount of particles are redeposited then a macroscopic geometry change can ensue. Macroscopic PFC geometry changes can create PFC surfaces that are proud of other surfaces, which acts like a feedback loop for subsequent failure modes. If a particle moves up the SOL towards the core, the charge state of the particle can change multiple times due to collisions with other high energy particles. These charge changes result in photon radiation that carries power away from the plasma to the tokamak wall, as the neutrally charged photons are not confined. Maintaining the PFCs below the sublimation threshold prevents PFC geometry alteration due to particle redeposition, and prevents performance reductions in the core due to photon power loss.

Melting will occur if the PFC material exceeds its melting point, which is a function of (the material pressure and) temperature. Once the PFC surface reaches the melting temperature, the material will change phase to a liquid from a solid. This liquid material will migrate through the divertor and redeposit in unintended locations, macroscopically altering the PFC geometry. After a melted material redeposits in a new location secondary evaporation, sublimation, melting, and fractures may occur, resulting in a positive feedback that can cause further PFC failure. Melting material redeposition usually alters the PFC geometry more than sublimation, so the secondary effects are also usually larger. Maintaining the PFCs below their melt point requires managing the heat loads properly through PFC design and machine operation.

PFC cracking (also called fracturing) occurs due to thermal stresses inside the material that exceed the material yield strength. The PFCs are constrained geometrically by tie downs, mounting hardware, and other rigid interfaces that inhibit free thermal expansion when heat loads are applied. These constraints result in stress concentration zones in the PFCs. Additionally, because the PFCs are not uniformly heat loaded, 3D thermal gradients across the PFC will form, causing thermal stresses due to different amounts of thermal expansion along the gradient. If the PFCs are

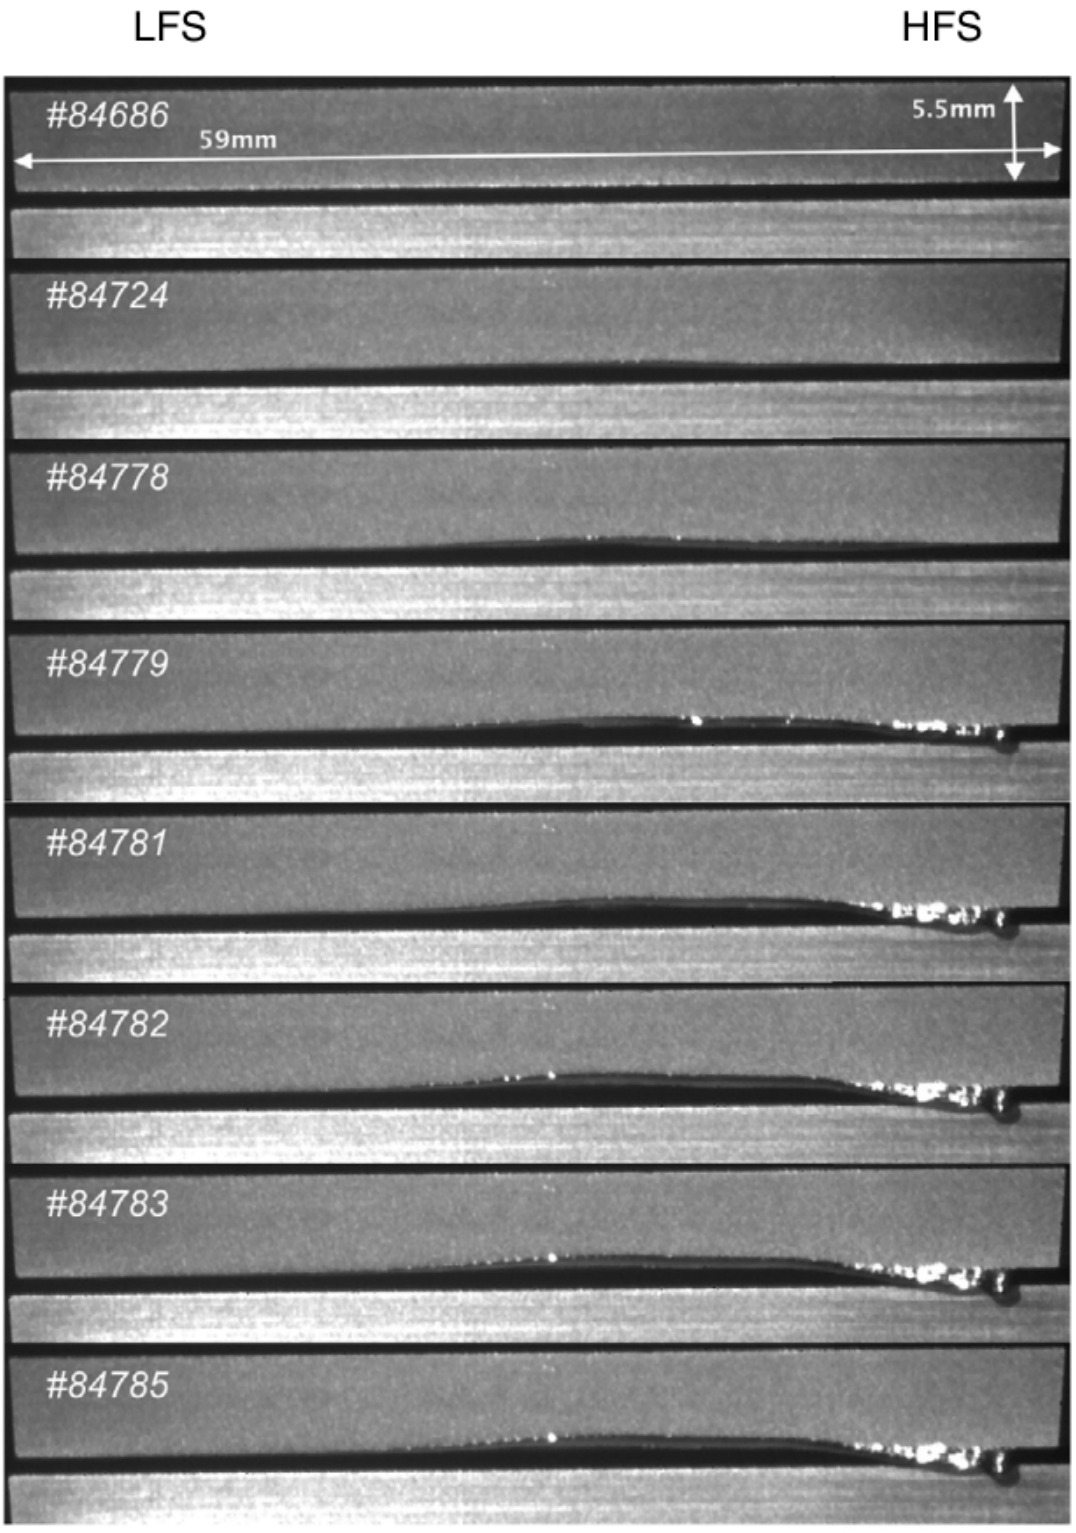


Figure 2.4: Example of melting tungsten from an intentionally exposed PFC in JET, due to ELM heat loads at near normal angle of incidence [37]. Top frame is unmelted PFC. Increasing shot numbers from top to bottom show the evolution of the melted area.

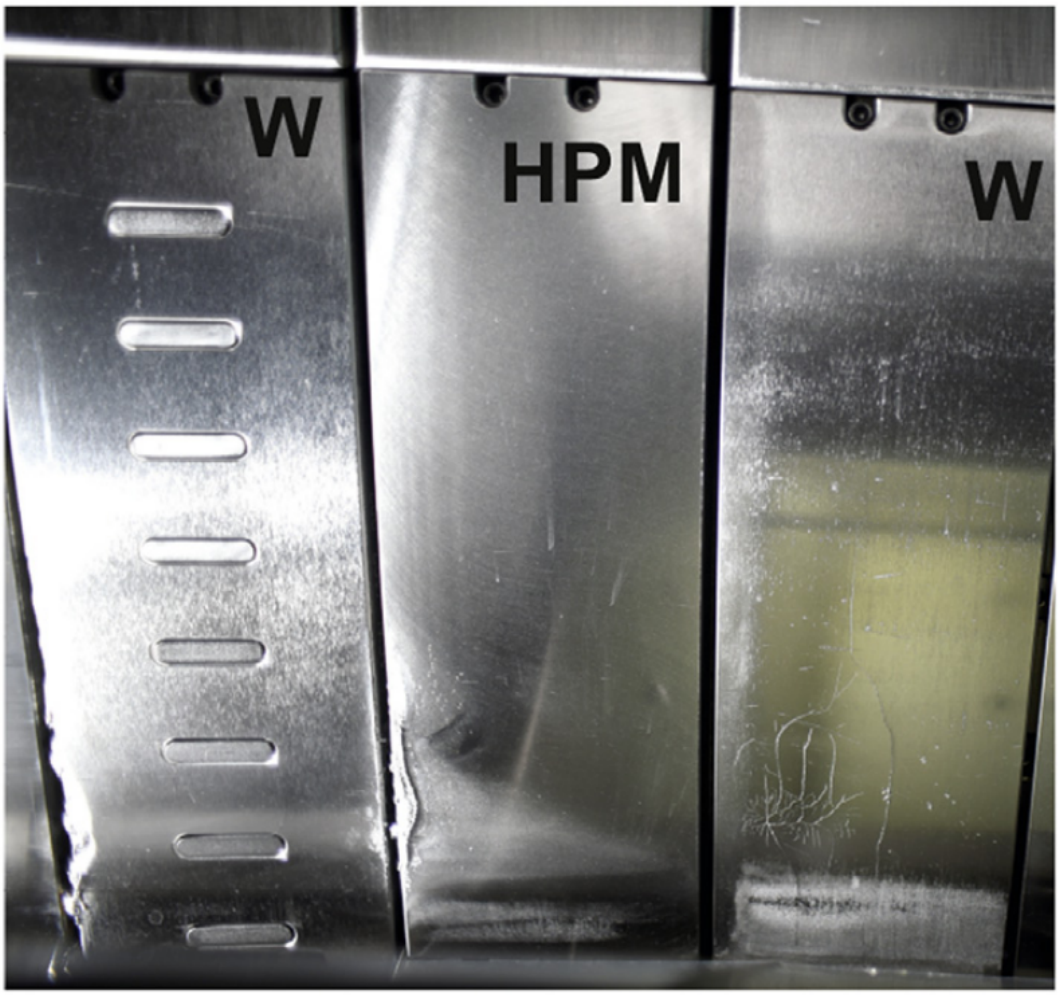


Figure 2.5: Example of melting tungsten from an unintended PFC misalignment due to halo current forces from ASDEX-U [46]. Right PFC also has cracks.

actively cooled, the temperature gradient from the PFC surface to the cooling channel will create thermal stresses. All of these effects create thermal stresses, and when the temperature gradients are sufficiently high the material will fracture. Because the temperature gradient is proportional to the incident heat flux via Fourier's equation, higher heat fluxes will result in larger thermal stresses. Once a PFC cracks, the crack exposes an unintentional leading edge that will be loaded with plasma heat flux, a change to the PFC geometry. Secondary effects of the heat loaded leading edge will be sublimation or melting, both of which reduce machine performance.

There are several examples of these consequences in the tokamak literature. At the Joint European Torus (JET), Tungsten PFCs were intentionally misaligned to observe the effects of leading edges directly perpendicular to the plasma flow in the SOL during transients and ELMs [37, 47]. Figure 2.4 illustrates the observed tungsten melting as a result of the high heat fluxes at angles of near normal incidence. Researchers at the Axially Symmetric Divertor Experiment Upgrade (ASDEX-U) have documented unintended tungsten melting due to PFC misalignments [46]. Figure 2.5 shows three PFCs, tungsten and HPM tungsten heavy alloy, after six months of use. The misalignments were caused by high halo currents in the PFCs that twisted some of the divertor mounting structures, resulting in exposed leading edges that see heat fluxes at near normal angles of incidence. In the far right PFC in figure 2.5 a crack network was discovered that reached through the entire PFC to the back side. In order to accurately predict these heat loads, a simulation framework that incorporates halo currents as well as plasma heat fluxes would be necessary.

2.4 Managing Divertor Heat Loads

Several techniques have been developed to avoid the aforementioned consequences and manage the divertor heat loads. Most modern tokamaks will employ some heat load management strategy, and the strategy must be customized to maximize a specific operational objective. That being said, there is a cost to managing the heat load as each strategy will constrain the tokamak operation differently. Because the divertor heat load happens at the interface between the plasma and the PFCs, the management strategies generally change one of those two systems, the plasma or the PFC. In order

to quantify the trade-offs and design PFCs that operate within their engineering limits, high accuracy time varying heat load predictions are necessary. More specifically, the heat load predictions should incorporate the full 3D geometry of the PFC, as this is the method most commonly used to manage the divertor heat loads.

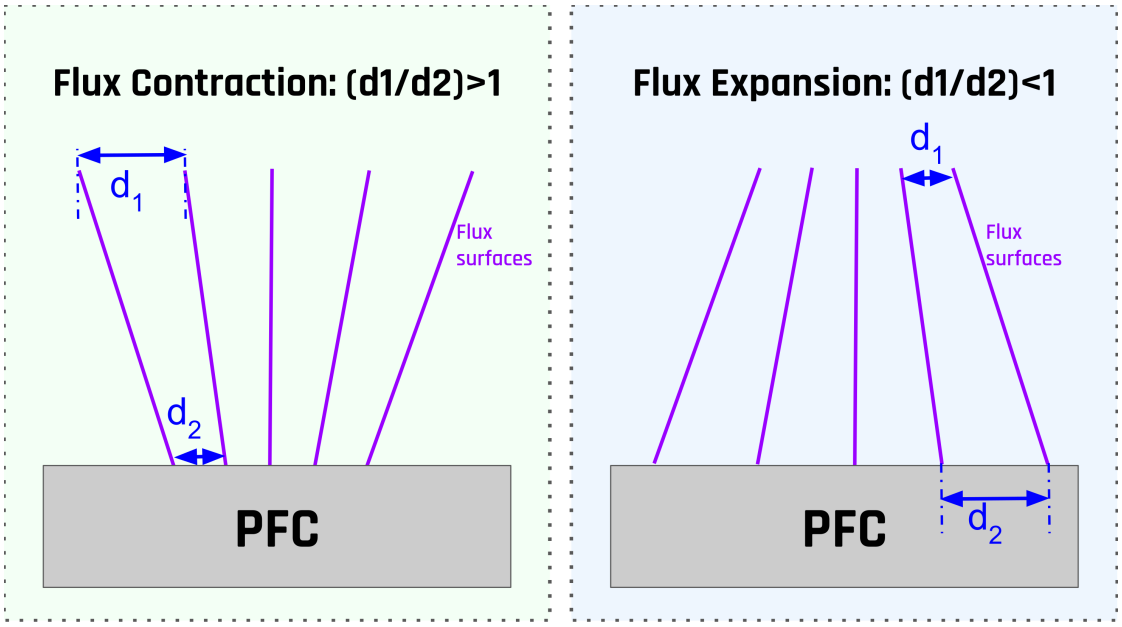


Figure 2.6: Simple cartoon of flux contraction and flux expansion. Purple lines are flux surfaces. If distance between flux surfaces decreases while approaching the divertor, it is considered flux contraction. If distance increases, it is considered flux expansion. Flux expansion distributes the heat load accross a greater surface area.

Changing the plasma amounts to changing the plasma magnetic equilibrium shape, or changing the plasma as a function of time. In both cases, the objective is to increase the surface area that will see a heat load, thereby decreasing the heat fluence. Figure 2.6 illustrates the concept of flux expansion, which is a method used to minimize the heat loads for static magnetic equilibria. Flux expansion is a term that refers to the change in distance between poloidal flux surfaces along any magnetic field line. When discussing heat fluxes to the PFCs, if the radial distance between adjacent flux surfaces increases while approaching the PFC, then that is considered flux expansion. Similarly, if the distance between flux surfaces decreases as the poloidal flux surface approaches the divertor, then that is considered flux contraction. Flux expansion distributes the heat across a larger surface area, thereby reducing the heat flux. Many modern machines operate with high flux expansion to spread the heat over the largest possible surface area. A downside to this effect is that flux expansion generally decreases the magnetic

field line incident angle, which can result in magnetic shadows if the PFC geometry is 3D. To determine the footprint of magnetic shadows, it is necessary to calculate the magnetic shadows. Chapter 3 discusses a tool to predict these magnetic shadows and chapter 4 presents some high flux expansion results for NSTX-U.

Another way to change the plasma to minimize the heat load is by employing a time varying magnetic equilibrium. Sweeping the strike point such that the heat fluxes are distributed across a larger surface area results in a smaller heat fluence. Sweeping the strike point requires advanced divertor coil operation, and may be constrained by the power supplies that perform the strike point sweeping. Finding the highest possible strike point sweep frequency allowed by the divertor coils, the optimum sweep trajectory, and the resulting PFC temperatures and stresses is a complicated analysis, but can be performed with the framework described in chapter 3. Examples of calculating the time varying heat loads due to strike point sweeping, and the optimization of the strike point trajectory, is discussed in chapter 4. Generally, increasing the strike point sweep frequency decreases the peak temperature and increases the allowable discharge length [28].

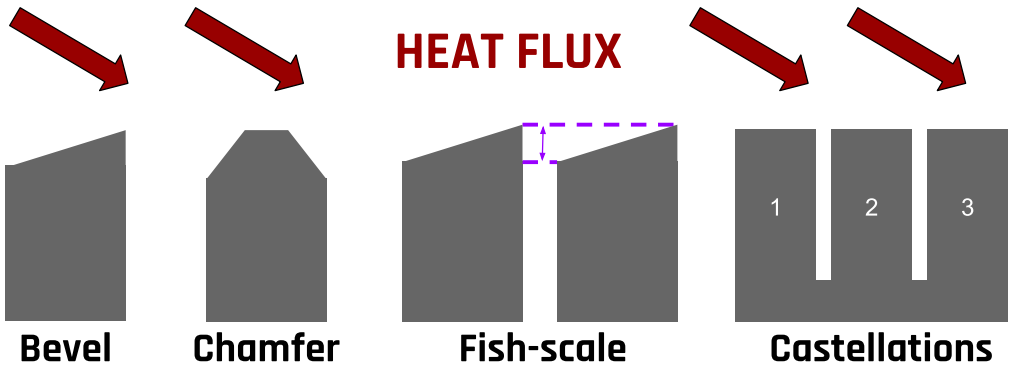


Figure 2.7: Simple cartoon of some PFC shapes utilized to manage heat fluxes from [28].

It is also possible to change the PFC geometry to manage the heat loads. For the toroidally axisymmetric divertor, decreasing the magnetic field line angle of incidence is the preferred method for decreasing the heat load. The axisymmetric divertor can be useful as a thought experiment, but real PFCs are 3D objects in a 3D world, so the axisymmetric assumption is limited in its validity. Even for tiles designed to be toroidally axisymmetric, installation misalignments and forces during the discharge will

violate the axisymmetry and cause the PFCs to vary toroidally. Furthermore, modern PFC designs utilize toroidal shaping to optimize the PFCs for a specific magnetic configuration. Figure 2.7 illustrates some of the common 3D PFC shaping techniques utilized to manage the heat loads [28]. Bevels and chamfers change the incident angle of the magnetic field line. Creating a vertical step between adjacent PFCs is called fish-scaling, and is used to protect leading edges from high angles of incidence. Castellations are formed when intermittent air/vacuum gaps are placed in the PFC, which arrests long range thermal diffusion and stresses. Integrating 3D PFC shaping into the design enables operation with higher parallel heat fluxes, thereby expanding the operational window with respect to parallel heat flux. Yet the bevels, chamfers, and fish-scales, each require tokamak operations to maintain the angle of incidence in a nominal range to be effective at shadowing leading edges. Trade-offs such as this are typical in the PFC design process, and require an ability to accurately predict magnetic shadows on the PFC surface, as well as an understanding of the tokamak operational space with respect to exhausted power and plasma shape [28].

2.5 The NSTX-U Divertor

This dissertation outlines investigations performed using a simulation framework described in chapter 3. This framework has been utilized to simulate heat loads on 5 tokamaks, but the primary investigations have been performed on the National Spherical Torus eXperiment Upgrade (NSTX-U). This section provides some context for the PFC engineering design for NSTX-U, and the text is taken from [28].

The high heat flux regions of the NSTX-U divertor have a castellated graphite design [48, 49], in which narrow slices are orthogonally cut deep into the graphite surface, resulting in void regions between adjacent 30mm wide 'cubes' of graphite. The voids eliminate transverse conduction paths across the tile, thereby reducing long range thermal stresses caused by local hot spots. This results in a temperature limited PFC rather than a stress limited PFC. Additionally, the castellations enable each cube to be used as a calorimeter from which the incident heat flux profile can be reconstructed [10, 50].

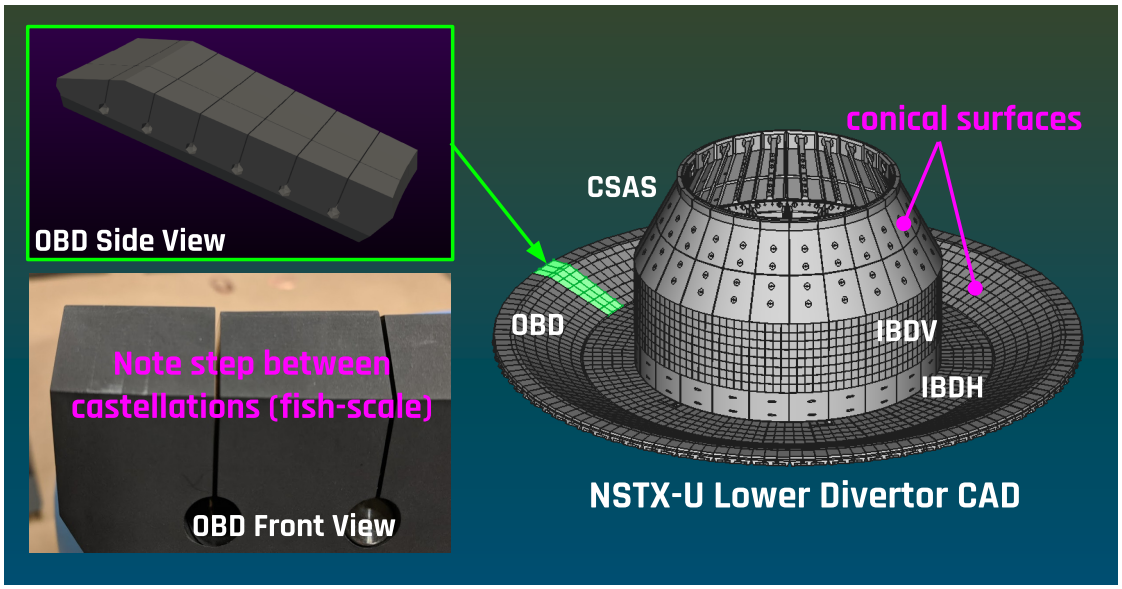


Figure 2.8: NSTX-U CAD and abbreviated PFC names: Center Stack Angled Surface (CSAS), Outboard Divertor (OBD), InBoard Divertor Horizontal (IBDH), InBoard Divertor Vertical (IBDV). OBD tile is highlighted in green and shown in exploded side view. Picture of real OBD graphite prototype with castellations and fish-scales as viewed from the front (from center stack). Conical surfaces are locations where heat flux faceting will occur. Figure taken from [28].

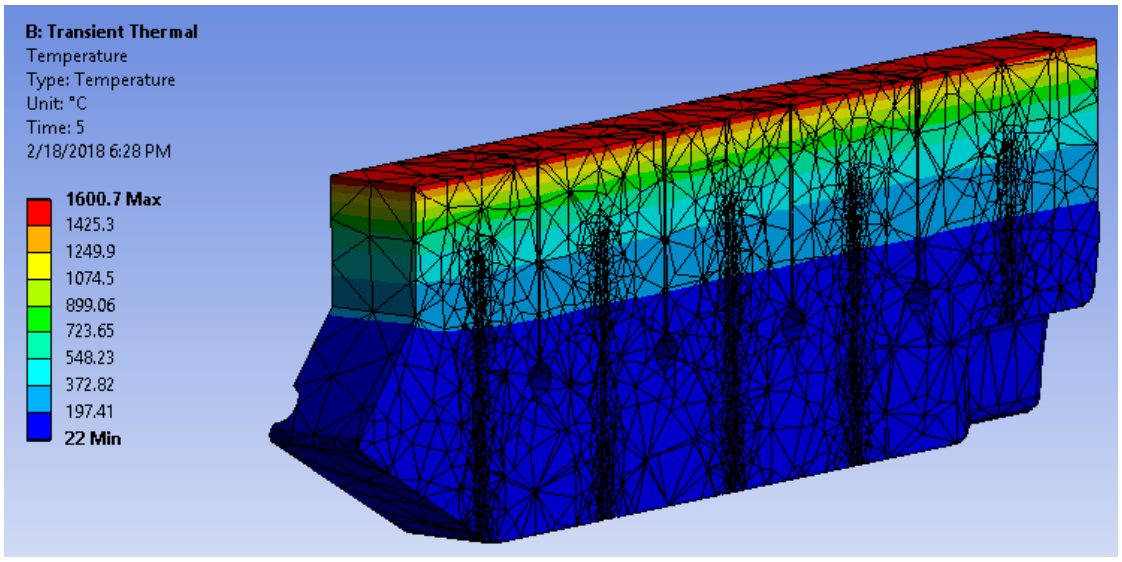


Figure 2.9: Simulated temperature at 5s discharge length for NSTX-U castellation row from IBDH tile with 8 MW m^{-2} applied uniformly at the PFC surface. NSTX-U graphite engineering limit is 1600°.

Figure 2.8 illustrates the lower divertor of NSTX-U, and highlights the 3D features. Four primary tile regions are used to describe the NSTX-U divertor. In the inner divertor, the InBoard Divertor Vertical (IBDV) PFCs are castellated and fish-scaled while the Center Stack Angled Surface (CSAS) PFCs are only fishscaled. In the outer divertor, both the InBoard Divertor Horizontal (IBDH) and OutBoard Divertor (OBD) tiles are castellated and fishscaled. The CSAS and OBD tiles are oriented to a faceted surface designed to match a conical surface and plasma shaping. The 3D nature of the castellations, fish-scales, and faceted surfaces, results in large deviations from the axisymmetric heat flux footprint, and must be accounted for when simulating heat flux to these regions.

The NSTX-U Recovery castellated graphite design assumes a sublimation engineering limit to be 1873 Kelvin, although this is known to be uncertain. Figure 2.9 shows an ANSYS temperature simulation for a single row of castellations (graphite cube) from the NSTX-U IBDH tile. A uniform heat flux of 8 MW m^{-2} was applied across the PFC surface. After 5 seconds, 1873 Kelvin is reached on the material surface. While this analysis is a simplified version of reality because the PFCs will not see a uniform heat flux across the surface, it provides conservative estimate. Furthermore, it illustrates that the machine discharge length will likely be constrained by the engineering limits of the PFCs. In order to accurately ascertain the boundary of the allowable tokamak operational domain with respect to the PFCs, a simulation framework that captures the full 3D geometry, time varying nature of the magnetic equilibrium, divertor physics / models for the heat flux, and temperatures and stresses, is necessary. Chapter 3 outlines the toolkit that was developed to simulate these interdependent systems with high precision.

Chapter 3

The Heat flux Engineering Analysis Toolkit (HEAT)

This chapter details a framework for simulating heat loads in tokamak plasma facing components (PFCs): the Heat flux Engineering Analysis Toolkit (HEAT). First, the need for time varying 3D heat load predictions is reviewed. The existing options for 3D PFC heat load analysis are presented. After a discussion on the development objectives, HEAT is described in detail. Lastly, a roadmap for software and physics development is provided. Some of the text and formulae presented in this section come from a recent HEAT paper [28] published by Looby et al (the author of this dissertation).

3.1 PFC Heat Load Simulation Requirements

3.1.1 Software Interfaces

Section 2.1 described a subset of the physical domains and engineering systems that are directly intefaced to the PFCs. In order to construct a simulation framework that is capable of simulating these complex interconnected systems, software interfaces between various code suites that represent each of these domains / systems is also necessary. Figure 3.1 provides a graphical illustration of some of the software packages that could be included in a PFC heat load calculation. Physics codes that generate MHD equilibria are needed to provide a topology of the plasma. These codes can either be 2D (EFIT, FreeGS) or 3D (M3D-C1, VMEC, SIESTA). Particle tracers

and magnetic field line tracers are needed to ascertain if a PFC is exposed to a heat load (MAFOT, SPIRAL). Radiated power and boundary physics codes are needed to quantify the level of divertor detachement and associated photon fluxes, and can be achieved with a code such as SOLPS. It is desireable to incorporate gettering, recycling, redeposition, into the heat load calculations, and codes such as GITAR can be used to this end.

Once the heat fluxes are determined, calculating engineering parameters of interest is necessary to ascertain the operational boundary of the tokamak and potentially control tokamak operation. Integrating with finite volume packages such as ANSYS or openFOAM allow for the temperatures and stresses in the PFCs to be calculated. The 3D structure of the PFCs can be accounted for by including a parametric CAD model and associated parametric design software. A parametric CAD model enables the HEAT user to adjust the geometry, simulate PFC misalignment, and interface with existing engineering design teams. To control the tokamak using HEAT results, an AI or ML reduced model would be desireable. Incorporating machine learning packages such as TensorFLow or Keras would enable ML algorithms to be trained using HEAT results. The tokamak plasma control system must be interfaced if these reduced models are to be utilized in real time for control. Visualizing data and drawing conclusions from simulation results requires a software suite that can render complex 3D geometry and overlay parameters of interest, such as ParaView. None of these post-processing packages are directly related to plasma physics, but they are necessary within the greater context of PFC simulation.

Interfacing pre-processing, plasma physics, and post-processing, software packages that exist independently is a primary motivation of this dissertation. This type of integrated modeling combines software packages from disparate computational domains together into a single integrated suite. The resulting simulation output is far more powerful than the sum of the individual software components, as it allows the dependence between the various domains described in section 2.1 to be explored in addition to the independent phenomena. The framework described in this dissertation has been applied to predict heat loads in PFCs, but a similar architecture could be applied to many other problems that require interfacing physics to engineering

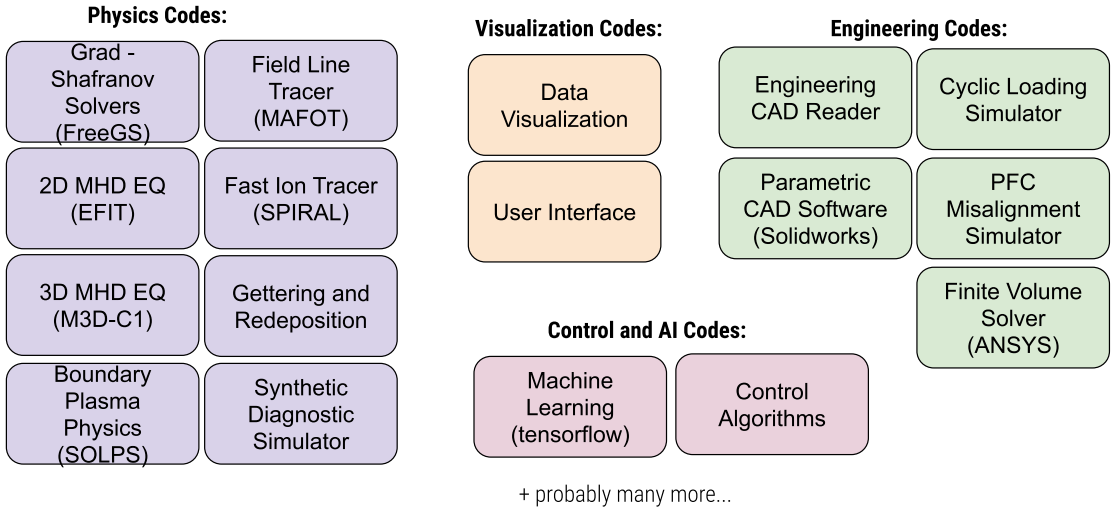


Figure 3.1: Cartoon of some software packages that could be included in a high fidelity heat load prediction. Creating a framework that is capable of interfacing directly to new code packages is desireable, as it enable modular development and does not limit the physics that can be included in the heat load calculations. Codes are from traditionally disparate computational domains, such as computational physics, visualization, engineering solvers, and machine learning.

systems. While the heat load predictions that are contained herein are the focus of the dissertation, the simulation architecture is the hidden gem in this work.

3.1.2 Simulation Objectives

Tokamaks are inherently 3D machines, as they operate in our 3D world. As such, they are subject to 3D phenomena that cannot be captured with the toroidally axisymmetric assumptions that are often used to determine the relevant heat loads. Even for tokamak PFCs that are designed to be perfectly axisymmetric, installation alignments, 3D magnetic field effects, or halo current forces, will likely result in toroidal variation of the PFCs and violate axisymmetry [36, 46]. Furthermore, modern PFC designs intentionally violate axisymmetry as a means of achieving a specific performance objective, as described in section 2.4. Time evolution is also critical to resolving the heat loads in PFCs, as some time varying magnetic equilibria are employed to reduce peak divertor heat loads. Additionally, the temperature profiles in the PFCs are also time varying, and temperature hysteresis will affect how heat diffuses through the tiles. So 3D PFC geometry and time evolution are both critical components to high fidelity heat load predictions.

These needs, and many others, were originally identified by the NSTX-U Recovery Working Group, who was tasked with redesigning the NSTX-U PFCs. During the NSTX-U recovery, many desireable magnetic equilibria were generated and analyzed for viability with respect to the PFC engineering limits, but these analyses were mostly axisymmetric [49, 51]. This PhD dissertation was in part motivated from the need to simulate time varying 3D heat loads. Following on the discussion in the last section, a PFC heat load simulation software suite should be capable of achieving the following objectives, which are taken directly from [28].

- ***Speed***. Quick predictions, in some cases in between discharges, are needed to guide new machine operation. This corresponds to the ability to run a simulation in less than 10 minutes to evaluate prior shots and help determine if the PFC engineering limits will be violated for subsequent plasmas. While clusters and supercomputers are available at most tokamak facilities, running on

a 16 (logical) core workstation was preferable for logistical reasons. A code that could determine if PFC limits were violated in under 10 minutes was desired.

- ***3D calculations***. This topic was discussed in previous sections, but it can be generally stated that using an axisymmetric assumption for heat fluxes induces errors into the prediction. A 3D heat flux can be categorized by two driving causes: 3D PFCs and 3D plasmas. The ability to predict heat loads from 3D plasmas incident upon 3D PFCs was desired. This manuscript will address 3D PFCs coupled to 2D plasmas, and subsequent work will investigate 3D PFCs coupled to 3D plasmas.

- ***FVM compatibility***. The finite volume method (FVM) is commonly utilized in engineering analysis to calculate the temperature and stresses within solids. While these are ultimately the constraining variables when it comes to PFC limits, there are very few interfaces that connect MHD, divertor physics, and FVM, in a single package. An interface between these domains was desired.

- ***Time varying capability***. Many modern tokamaks utilize strike point sweeping to mitigate heat loads, and it is suggested as an off-normal operating condition for a reactor [6]. The ability to input time varying magnetic equilibria, and calculate the associated heat fluxes and temperatures as a function of time was desired.

- ***Open source***. The code needed to be open source and comprised of open source packages, with no legal or intellectual property obligations to a specific machine or organization. The open source philosophy for software development permits contribution from many sources, promotes scientific collaboration, encourages code recycling, and eliminates proprietary legal constraints. It was desired to have a code capable of serving both the public and private sectors, across international and organizational borders.

3.2 Potential Candidate Codes

Before HEAT development began, potential candidate codes that could perform the heat load calculations were reviewed [28]. Ultimately, it was determined that no existing code could achieve all of the aforementioned objectives. A predecessor code, the Surface HEat Flux Tracer (SHEFT) [52], was utilized as a starting point for HEAT development, but HEAT development was largely from scratch. Two potential candidate codes were reviewed to satisfy the design objectives outlined in the last section.

A package developed for ITER called SMITER [53] achieved many of the aforementioned objectives, and it was considered as an alternative to coding HEAT from scratch. Figure 3.2 provides a tabular comparison between HEAT and SMITER. Another code that was investigated to satisfy the objectives was the Plasma Facing Components Flux (PFCFlux) code, which has been used to generate heat flux predictions on the JET ILW, EAST, and EU-DEMO [54, 55, 56]. Both codes were designed to predict heat fluxes for static magnetic equilibria, and to our knowledge have not been utilized in time varying analysis (ie strike point sweeping) for double null plasmas, which was a requirement for HEAT. Additionally, to determine PFC operational limits, a finite volume solver needed to be included in the analysis loop and no such open source solver was included in either candidate code. Lastly, neither SMITER nor PFCFlux have been used to apply 3D plasmas to 3D geometry, which was necessary to predict the heat loads resulting from magnetic coil error fields, toroidal field ripple, and resonant magnetic perturbation (RMP) structures. Both codes seem to be capable in their respective domains, but it was decided that neither could satisfy our objectives. It was therefore decided that the best way to progress was to develop HEAT internally.

3.3 HEAT Philosophy

HEAT is fundamentally open source [28]. Leveraging the power of the open source software community enables rapid software development, and permits organizational agnosticism. In its current form, HEAT is designed to run on Linux. HEAT contains

	SMITER	HEAT
Axisymmetric Plasmas	X	X
Ability to read M3D-C1		X
3D PFCs	X	X
Eich Heat Fluxes		X
Multiple Exponential Heat Fluxes	X	X
Limiters	X	X
Time Varying Analysis		X
Advanced Divertors (ie snowflake)		X
Temperature Simulations		X
Stress Simulations		X
MPI Parallel Simulations	X	X
Radiated Power w/ Raysect		X
Ion Gyro Orbit Tracing		X
Graphical User Interface	X	X
Parametric CAD Design	X	X
Network LAN Capability		X
Windows		X
Mac		X
Linux	X	X
Open Source		X

Figure 3.2: Tabular comparison between SMITER and HEAT. List is not all-inclusive, and is only valid at current time of publication (Fall 2021).

an open source parametric CAD module, FreeCAD [57]. Temperatures and stresses are calculated using OpenFOAM [58], an open source finite volume toolkit. ParaView [59] provides data visualization and analysis and Plotly DASH [60] provides an HTML5 graphical user interface (GUI). While not explicitly mentioned here, many other open source packages can be found in the HEAT source code, which is available to the public on github [61]. HEAT is primarily coded in python3, which makes it easy to code new modules and provides access to the wide array of open source python packages available from the python community.

HEAT's open source nature has made collaboration easy. The framework can be ported to any tokamak, provided there is a magnetic equilibrium GEQDSK and CAD STP file. To date, HEAT has been utilized to simulate heat loads for five tokamaks: NSTX-U, ST40, DIII-D, SPARC, and EU-DEMO. Scientists from three national labs have used HEAT: Oak Ridge National Laboratory, Princeton Plasma Physics Laboratory, and Lawrence Livermore National Laboratory. It is available in an appImage for Linux, which requires running a single executable to run HEAT. Tutorials and youtube videos are linked to the github page [61].

3.4 HEAT Architecture

3.4.1 Concept Architecture

HEAT is a modular code, consisting of six primary modules. These six modules are illustrated graphically in figure 3.3, which is an abstraction of the actual code provided to show HEAT concepts rather than actual software pipelines. The six modules are textually described below.

- ***CAD***: object that reads engineering CAD file (in STP format - see ISO 10303-21) and stores data such as parametric geometry, PFC names, performs surface (STL file format) meshing, and contains other geometry relevant data.
- ***MHD***: object that stores the magnetohydrodynamic (MHD) equilibrium data provided by EFIT equilibrium files, performs field line tracing, calculates MHD related quantities, and handles all magnetic field related operations.

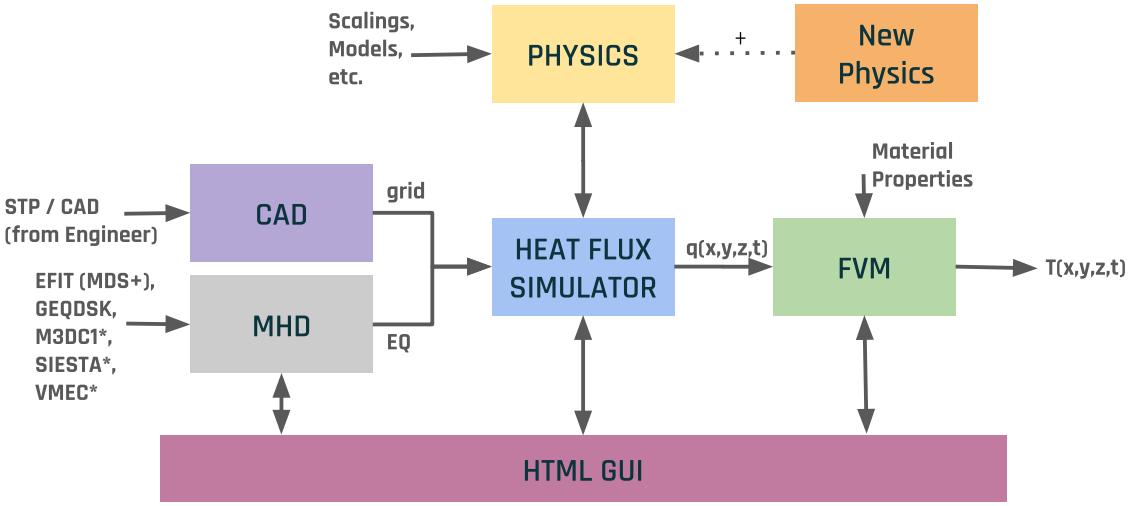


Figure 3.3: HEAT concept flow chart. Note that each box may contain multiple python classes. Interface for 3D plasma magnetic equilibria using M3DC1 / SIESTA is currently under construction so it has been marked as partially implemented. Figure from [28]

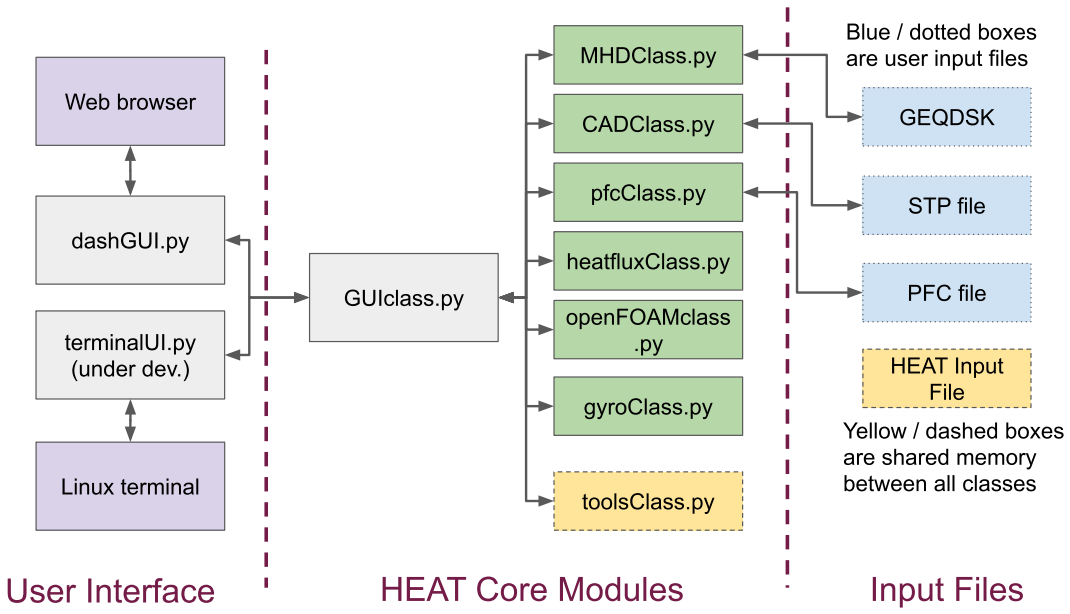


Figure 3.4: HEAT software architecture flow chart. Each box represents a different software module, python class, or input file. Note that some python classes have been omitted for brevity. Third party code such as MAFOT or openFOAM are not included in this flow chart.

- ***PHYSICS***: object that contains user selected empirical regressions, analytical models, and other physical parameters. If additional physics are to be added to HEAT, a submodule would be appended to this module, as illustrated in the figure.

- ***HEAT FLUX SIMULATOR***: object that connects the MHD, CAD, and Heat Flux modules for each PFC of interest and generates a spatio-temporal evolving heat flux profile on PFC surface.

- ***FVM***: object that generates volume meshes from surface meshes, then performs finite volume method (FVM) calculations. Applies surface heat flux to the volume mesh boundary and solves for internal temperature. Generates spatio-temporal evolving temperature profiles.

- ***HTML GUI***: object that provides user with a graphical user interface (GUI) via a web browser. Google Chrome has been used in HEAT developement and testing.

In order to run HEAT, the user must supply a CAD file in the STP (ISO 10303-21) format and magnetic equilibria EFIT reconstructions in the GEQDSK format, as discussed in section 1.4.3 and [29]. The user must also specify some parameters (such as the desired heat flux resolution, profile type, divertor power sharing, etc.) in the GUI or via an input file. If the user loads multiple GEQDSK timesteps into the HEAT GUI, then the HEAT time varying capability automatically is activated and the solution will be found as a function of time. If using the FVM module, then the material properties for the PFC will need to be chosen by the user.

3.4.2 Software Architecture

While figure 3.3 is useful for understanding HEAT concepts, the software pipelines in HEAT are more complicated. Figure 3.4 provides the architecture of some of the actual HEAT python classes (some are omitted for brevity). The figure is divided into three columns, separated by dashed maroon lines. Each column represents a different subdomain of the HEAT software ecosystem. The left column represents the user interface. Currently, HEAT must be run from the HTML5 user interface,

which is constructed with the Dash [60] python3 wrapper around HTML5, CS, JS. The user browses to the IP address of the HEAT session in a web browser to run HEAT. Currently, a terminal based interface is under development in will be released in a future HEAT version. Both the Dash HTML5 interface and the terminal interface connect to an intermediate class object, GUIclass.py.

The HEAT modules in the center column of figure 3.4 are the core modules that perform the heavy lifting in HEAT. The GUIclass object is the primary module that links all other modules together, while the other class objects are secondary modules. Not only does the GUIclass object connect the various HEAT modules, but it also performs the HEAT calculation and steps through time solving for HEAT fluxes and temperatures.

All modules in the center column of figure 3.4 are secondary modules that perform specific functions. The MHDClass reads GEQDSK files, performs magnetic field operations, and calls the field line tracer, MAFOT. The CADclass uses FreeCAD to import STEP files, filter and organize part objects, and create surface STL meshes. For each part in the user defined region of interest (ROI), a PFCclass object is created. Each PFCclass object contains the data for that part object, and also initiates the parallel intersection checking algorithms for optical and gyro orbit heat loads. The heatfluxClass performs heat flux calculations and power calculations. It contains empirical regressions, physics models, and functions related to power balance. The openFOAMclass object provides the interface to the openFOAM finite volume simulations. It generates volume meshes from surface STL meshes, copies openFOAM templates to each PFC's data directory, and runs the openFOAM binaries for HEAT. HEAT has a custom openFOAM solver, heatFoam, which solves the 3D heat diffusion equation with temperature dependent thermal properties, using the heat flux as a boundary condition. The gyroClass is responsible for calculations related to gyro orbit heat loads. It samples from velocity distribution functions, and traces helical gyro trajectories. The toolsClass object is a shared memory class that hosts utility functions, as well as the parallelized intersection algorithms that are called by the PFCclass objects. Because it is a shared object, MPI processes can access its memory without the overhead of transmitting data. For more information on the HEAT code, see the github repo [61]

3.4.3 Input Files

The right column in figure 3.4 contains the input files for HEAT. The GEQDSK file is a magnetic equilibrium file that was described in detail in section 1.4.3. The STP file is a CAD file that follows the STEP ISO 10303-21 standard. The PFC file is a comma seperated value (CSV) file that contains the definitions of the region of interest, timesteps of interest, potential intersections, and divertor code. Below is an example of a typical PFC file (future versions may vary):

```
1 #PFC CSV FILE
2 #Comments are allowed
3 timesteps, PFCname, resolution, DivCode, intersectName, excludeName
4 0:5000, Paddle-Frontplate-N_Default<As Machined>002, 5.0, UO, all, none
```

The first uncommented line in the PFC CSV file should always appear as line 3 does: 'timesteps, PFCname, resolution, DivCode, intersectName, excludeName'. Each of these strings correspond to a column in the CSV. The HEAT user creates this file to tell HEAT which PFCs the heat loads should be calculated on. The list of CAD parts that will have a calculation performed on them is referred to as the region of interest (ROI). For the example above, there is a single PFC in the ROI, and its definition is on line 4. Each new line after line 3 represents a part from the CAD file that should be included in the ROI. The PFCname column is the name of the CAD part in the ROI, and can only include one part. The string provided in the PFCname column should directly correspond to a single part name from the STEP file. The PFCname should correspond to a primitive, or a part object, and not an assembly. The timesteps column is a colon (:) seperated list containing start and end times for that PFC. The resolution column contains the ROI mesh resolution for that PFC. The DivCode column is contains tags that assign each PFC in the ROI to a specific divertor in the tokamak. Its values can be LO: Lower Outer, LI: Lower Inner, UO: Upper Outer, or UI: Upper Inner. The intersectName column contains a colon (:) separated list of all parts in the CAD that could potetially cast magnetic shadows onto the PFCname PFC. Parts that cast magnetic shadows upon the PFCs in the ROI are called 'intersects,' and each PFC in the ROI has a unique list of intersects. If the user is unsure about where the magnetic shadows in the CAD may originate, then 'all' can be written in this column to check

against all PFCs in the STEP file. The excludeName column is a colon (:) seperated list of part names that should be excluded from the intersection calculation. This excludeName column is useful when the user defines the intersectName to be 'all', but wants to omit a specific PFC from the list of potential shadow casting PFCs. When not in use, the excludeName column can be set to 'none' or 'NA'. For the example above the part name is Paddle-Frontplate-N_Default<As Machined>002 and the heat loads will be calculated on its surface for any MHD equilibrium timesteps between 0 and 5000 ms. Paddle-Frontplate-N_Default<As Machined>002 is assigned to the lower outer divertor, and all the part objects in the CAD file should be checked as sources of magnetic shadows on its surface.

In addition to the PFC CSV file, the GEQDSK file, and the CAD STP file, the user has the option to upload a 'HEAT Input File' that contains all of the variables that may be chosen on the 'Inputs' tab in the GUI. This enables the user to skip the process of filling out text boxes in the GUI and simply drag and drop all the text box values into the GUI with a single click. The contents of the GUI regularly change with updated HEAT versions, so this file is likely to change and each value will not be presented here. It should be noted, however, that the HEAT Input File also follows the CSV format. It contains two columns. The left column is the name of the HTML5 div ID that corresponds to the text box (as assigned in dashGUI.py), and the right column is the value it should be filled with. A default version of this file can be downloaded from any HEAT release. Again, the file is not mandatory. See the HEAT github and tutorials for more details.

3.5 CAD and Meshes

To capture the heat loading effects associated with 3D PFC geometry, HEAT reads STP files (ISO 10303-21) provided directly from the design engineer. This eliminates any toroidal assumptions about the geometry and enables HEAT to directly connect the magnetic equilibrium to real 3D PFC geometry. This section discusses the HEAT CAD algorithms that include STP file operations, parametric CAD operations, meshing algorithms, and the associated architecture for these modules. Typically in computer graphics for CAD, a primitive is defined as geometric shape. In this discussion,

the primitive is synonymous with the words: part, component, PFC, and tile, and references the basic FreeCAD part object (Part.Feature).

3.5.1 FreeCAD

HEAT utilizes the power of FreeCAD [57], an open source parametric computer aided design (CAD) program, to perform operations on the CAD. From the HEAT simulation perspective, FreeCAD's power lies in its fully featured python3 interface that enables users to perform any graphical user interface (GUI) operation directly in a python console or from a python script. In this way, HEAT can perform any CAD operation that can be performed in a typical GUI based parametric CAD program. Future work will leverage this power to create autonomous design optimization loops, but for now the FreeCAD module is used for basic functions. FreeCAD is available for Linux, Mac, and Windows. On Linux, FreeCAD can be installed from the typical yum or apt repositories, but can also be run as an appImage. The HEAT appImage includes the Linux FreeCAD libraries. HEAT has an entire class dedicated to CAD operations using FreeCAD, CADclass.py, as shown graphically in 3.4.

3.5.2 Mesh Algorithms

Some of the following discussion is from [28]. HEAT uses FreeCAD (section 3.5.1) to convert parametric part objects into STL formatted surface meshes. More specifically, the FreeCAD **Mefisto** mesher is used in this process, which creates a triangle mesh. The STL mesh resolution is determined by constraining the maximum allowable edge length (decreasing the maximum edge length increases the resolution). It should be noted that the resolution is always finer than the maximum mesh edge length, as mesh elements are distributed randomly across the PFC surface. The required mesh resolution is highly dependent upon the PFC geometry, and must be chosen to avoid heat flux aliasing on small features. Section 3.8.3 describes these nuances in more detail. Additionally, if the user is incorporating ion gyro orbit heat loads into the results, then the mesh resolution should be chosen to be smaller than the ion's Larmor radius.

Once the surface STL mesh has been created in HEAT, all calculations performed will utilize the mesh. The mesh center locations and normal vectors are calculated for each mesh triangle in the mesh. Calculations such as heat flux, magnetic field line incident angle, etc., are performed using the mesh centers and mesh normals. The order of the mesh faces in the mesh centers list will directly correspond to the calculated quantities. Magnetic field line traces originate from the mesh centers. Calculation of magnetic shadows (section 3.7) happens using ray-triangle intersection algorithms that use the mesh vertices of all the potential primitives in the CAD.

In addition to the surface STL meshing, HEAT will also create volume meshes for finite volume solvers (FVMs) as a means to calculating temperatures and stresses throughout the PFCs. HEAT uses openFOAM, a robust partial differential equation solver, for finite volume method calculations [58]. Using the STL mesh as a starting point, HEAT will 'snap' a volume mesh to the surface of the STL mesh using the openFOAM **snappyHexMesh** utility. Heat fluxes generated using HEAT will have been calculated at the STL mesh centers, and will need to be interpolated to the volume mesh surface using the openFOAM **timeVaryingMappedFixedValue** algorithm. The user controls the accuracy of the volume mesh snapping process via control in the HEAT GUI. Similarly to the STL surface mesh discussion provided above, the accuracy of the volume mesh should reflect the objective of the simulation. Additionally, the user may take the heat flux output to a third party FVM solver such as ANSYS or COMSOL.

3.5.3 HEAT CAD and PFC Architecture

When running HEAT, the user supplies a STEP file containing the CAD geometry as well as a PFC file containing information regarding a region of interest (ROI), where the heat loads will be calculated (see section 3.4.3). Each line in the PFC CSV input file defines a new part object from the STEP that should be included in the ROI. Additionally, each line in the PFC CSV input file defines the intersects (STEP parts that could potentially be casting magnetic shadows upon the ROI PFC). Both the ROI and the intersects are meshed according the user defined settings in the GUI, using HEAT's CAD module.

The ROI is used to generate several lists. A ROIList object contains the names of all the PFCs in the ROI. A ROIparts object contains a list of all of the primitives in the ROI. The ROImeshes object contains a list of ROI meshes created from ROI primitives. ROInorms and ROIctrs are lists containing all of the mesh normal vector and center locations. Each of these lists is indexed to match the ROIList object, the list of PFCnames taken line by line from the PFC CSV Input file. Additionally, lists are created for all the potential intersection PFCs (PFCs that may cast magnetic shadows, so would be intersected on a field line trace originating at the ROI).

For each ROI object (PFC CSV file line), a PFCclass object is created called 'PFC', using the PFCclass.py from figure 3.4. PFC objects contain a reference to the ROI and intersect meshes for that PFC, the name of the PFC, the mesh normals, mesh centers, mesh areas, magnetic equilibria references, various point clouds (heat flux, incident angle, power direction, etc.), and other information. The PFC objects regularly get passed between other modules to build out this data. In this manner, all the data for each part in the ROI is stored in a single PFC object. The developer can get all the data for a PFC from a single object.

3.6 MHDclass in HEAT

HEAT requires a 2D magnetic equilibrium (EQ) file in GEQDSK format, which is described in section 1.4.3. Loading a GEQDSK file via the HTML5 user interface is the preferred method, but HEAT also has the ability to read from an MDS+ server. This feature is rarely used in HEAT, and is primarily included for future HEAT deployments that live on network clusters at national labs. After the GEQDSK has been loaded into HEAT, the MHDclass.py from figure 3.4 is utilized to create a magnetic equilibrium object in HEAT, called an 'EP' object. While HEAT can read GEQDSK files that have been named in any format, the preferred method is the DIII-D format, g<shotNumber>.<timestep[ms]>, where shotNumber and timestep[ms] are replaced with the shot number and timestep of the GEQDSK.

The MHDclass also contains various tools to manipulate GEQDSK files. The HEAT HTML5 GUI has a tab titled 'gFile Tools', where the user can change the definition of the last closed flux surface (LCFS), change the discharge helicity or toroidal field

direction, interpolate GEQDSK files for strike point sweeps, and more. The tab also has several plots that make it easy to determine if an incoming GEQDSK is formatted correctly.

Magnetic quantities, such as the magnetic field components, are computed at each mesh center location in the ROI. Equations 1.65 - 1.67 are used to calculate the magnetic field across the PFC surface, and this information is stored in the PFC object as described in section 3.5.3. Additionally, information about the magnetic field line incident angle at each mesh center is stored in the PFC object. This quantity is defined as,

$$\text{PFC.bdotn} = \hat{b} \cdot \hat{n} \tag{3.1}$$

where $\hat{b}$ is the normalized magnetic field vector and $\hat{n}$ is the mesh element surface normal vector, both evaluated at the mesh center location. A quantity that is of particular importance in the HEAT calculation is the flow direction of power onto the PFC surface, called 'powerDir'. Every CAD part will have surfaces that are magnetically shadowed, as well as faces that are heat loaded. The powerDir quantity defines the direction that power must flow to strike a face of the PFC, given the magnetic field orientation at that location. Mathematically, powerDir is defined as,

$$\text{PFC.powerDir} = \text{sign}(\hat{b} \cdot \hat{n}) \times \text{sign}(B_{t0}) \times -1.0 \tag{3.2}$$

where B_{t0} is the toroidal field strength [T] at the magnetic axis, and the final -1.0 is used because power flows into the PFC surface (not away from it). The quantity B_{t0} is found directly from the GEQDSK file, and its sign is assumed to be positive when it aligns with the $\hat{\phi}$ unit vector. The $\hat{\phi}$ vector is determined from the right handed $(\hat{R}, \hat{\phi}, \hat{Z})$ coordinate system (such that $\hat{\phi}$ goes counterclockwise when the torus is viewed from above). HEAT uses the powerDir to decide which toroidal direction power would have to flow to be incident upon a mesh surface. The powerDir represents the toroidal direction power flows in the global $(\hat{R}, \hat{\phi}, \hat{Z})$ coordinate system. This variable is a scalar, and will be 1 if power flows in the positive toroidal direction $(+\hat{\phi})$ and -1 if power flows in the negative toroidal direction $(-\hat{\phi})$. In equation 3.2, $\text{sign}(\hat{b}\cdot\hat{n})$ represents the relative orientation between a PFC mesh face and the magnetic field, and $\text{sign}(B_{t0})$ represents the relative orientation between the magnetic field and the global $\hat{\phi}$ coordinate (could

also be represented by sign($\hat{b} \cdot \hat{\phi}$), but this is more computationally expensive). Because powerDir is a mesh quantity (ie it is evaluated for each mesh element rather than an entire PFC object), it can reverse across a PFC surface, depending upon the MHD equilibrium. This is important for discharges where the sign of $\hat{b} \cdot \hat{n}$ reverses across the PFC surface, such as center stack limited discharges.

Of the two possible directions power can flow through a surface mesh element (ie an STL planar triangle), only one will be physical. This is because plasma power flows from outside the PFC solid into it, not from inside the solid to the exterior plasma. The surface normal vector, $\hat{n}$, of each mesh element is defined such that it points to the exterior of the solid, not to the interior. For every mesh element, the power flow direction (opposite the subsequent field line trace direction) will be determined by the magnetic field direction and the orientation of the mesh element normal vector relative to the magnetic field direction. If the mesh normal vector ($\hat{n}$) and the magnetic field normal vector ($\hat{b}$) are aligned ($\hat{b} \cdot \hat{n} > 0$), then power will flow opposite to the magnetic field vector, as shown by the lower (magenta) normal vector in figure 3.5. If the mesh normal vector and the magnetic field vector are not aligned ($\hat{b} \cdot \hat{n} < 0$), then power will flow in the direction of the magnetic field vector, as shown by the top (blue) normal vector in figure 3.5. Switching the plasma B_t direction should change the power flow direction, as is depicted in the comparison between forward (left frame) and reverse (right frame) B_t provided in figure 3.5.

Once the power flow direction for each PFC mesh face has been determined, HEAT will launch magnetic field line traces from each mesh center. As the trace progresses away from the origin mesh center, a ray-triangle intersection checking algorithm will determine if the magnetic field line intersects any other mesh faces from the CAD. If the magnetic field line does intersect another face, then the origin is considered magnetically shadowed, and will see no heat flux. Section 3.7 discusses these intersection checking algorithms in detail.

3.7 Finding Magnetic Shadows

Some PFC faces will be shadowed from the incident heat flux due to the 3D geometry of the PFC. These shadowed regions are referred to as 'magnetic shadows', and

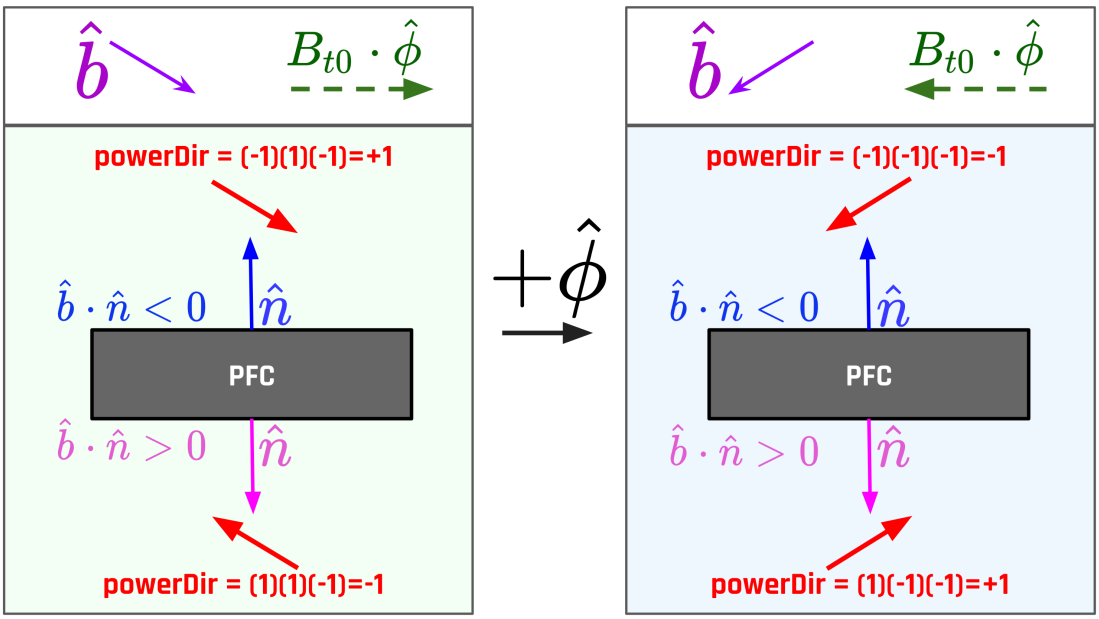


Figure 3.5: Cartoon example of how powerDir is assigned to PFC surfaces. Variable powerDir is a scalar that defines which toroidal direction (ie forward or reverse) power would have to flow to strike a PFC surface, and is defined in equation 3.2 as $\text{powerDir} = \text{sign}(\hat{b} \cdot \hat{n}) \times \text{sign}(B_{t0}) \times (-1.0)$. Positive toroidal direction goes from left to right. Switching B_t direction reverses powerDir.

change with different magnetic equilibria. HEAT has a set of algorithms dedicated to determining where these magnetic shadows lie on the PFC geometry. A field line trace is initiated from every PFC target face, and as the trace progresses away from the surface HEAT determines if any other PFC faces are intersected along its path towards the core. In HEAT the algorithmic process of finding magnetic shadows is referred to as 'intersection checking.' HEAT was designed to make no assumptions about the CAD geometry, and can perform intersection checking on any CAD object in the input CAD file without any user input except for the length of the intersction trace.

The intersection checking magnetic field line traces are generated by a HEAT wrapper around a C++ code developed by A. Wingen called MAFOT [62], which is capable of field line tracing for 2D (axisymmetric) and 3D (non-axisymmetric) plasmas. The MHDclass is used to write control files for MAFOT, as well as perform system calls to the MAFOT binaries. For the results discussed in this manuscript, the 2D field line tracer is sufficient, but it is important to note that MAFOT can already perform field line tracing for 3D plasmas, which is why it was chosen to satisfy the HEAT development objectives. Simulating an ideal 2D (axisymmetric) plasma incident upon 3D PFCs eliminates the complexity associated with toroidal variations in the magnetic field, but still captures the response of the 3D geometry. Unless the plasma axisymmetry is significantly perturbed, the 2D plasma approximation is sufficient. For the cases when axisymmetry is perturbed, future work will enable the HEAT user to investigate the change in heat flux that arises from toroidal variation in magnetic field using the 3D plasma MAFOT module. After receiving the origin coordinates for the field line trace, MAFOT 'walks' along the magnetic field line, saving the new position of the trace every toroidal degree. HEAT creates linear line segments from each of these one degree steps, and checks for an intersection with each face in the potential intersection PFCs. This process is parallelized using an open source implementation of the message passing interface (MPI), and is designed to use all available cores (leaves two cores for overhead). If an intersection is found as the trace progresses from the target to the core, no heat flux is assigned, and the face is assigned a value of 1 in a binary mask variable called ShadowMask. The trace for that face is terminated, and that face is omitted from the remainder of the intersection calculations. An example

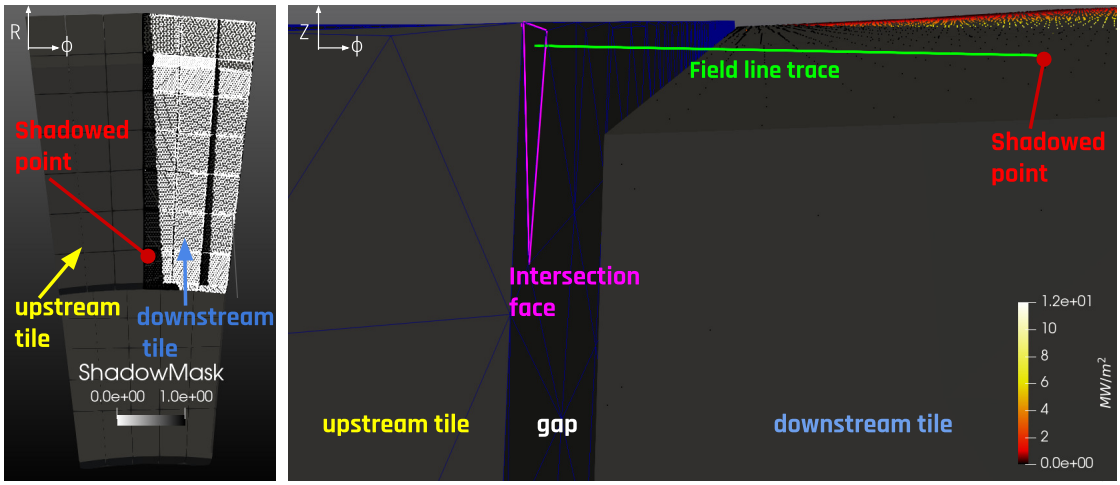


Figure 3.6: Shadowed faces are calculated using MAFOT field line traces for a PFC tile in NSTX-U lower divertor. Field line traces are launched from each target mesh face, and if field line intersects another upstream face then heat flux is 0. Left image shows view looking down at divertor from core with ShadowMask=1 when a face is shadowed. Right figure shows close up view of poloidal running gap between two PFC tiles, field line launch point, intersection mesh face, and field line trace. Figure from [28].

of a ShadowMask on an NSTX-U PFC is given in figure 3.6, along with a visualization of the field line trace, intersection face, and corresponding heat flux.

3.7.1 Intersection Checking Algorithms

HEAT uses a fully parallel implementation of ray-triangle intersection checking. During the intersection checking process, a magnetic field line is launched from each mesh center in the ROI. The magnetic field line is traced in 1 degree steps, and at every step each trace is checked for an intersection with all of the potential Intersect PFC mesh elements (the mesh triangles that could potentially be casting magnetic shadows on the ROI). Each trace can be performed in parallel using HEAT's openMPI implementation of intersection checking. This corresponds to parallelization across the PFC surface in space. A MAFOT trace identifies the start and end point of the magnetic trace for a 1 degree sector, and these points are referred to as $\vec{q_1}$ and $\vec{q_2}$. For the first trace step from the PFC surface, $\vec{q_1}$ is the mesh center. For the remaining steps, $\vec{q_1}$ is the last step's $\vec{q_2}$. Once $\vec{q_1}$ and $\vec{q_2}$ are calculated, a linear line segment is formed between the two points (ie $\vec{q_2} - \vec{q_1}$). It is this line segment that is checked against all potential Intersect mesh triangles via a ray triangle intersection algorithm. There are three ray-triangle intersection checking algorithms in HEAT. The original methods were adaptations of a signed volume (SV) method, which is described in [63]. These were eventually replaced with the Möller-Trumbore (MT) algorithm, which uses a barycentric coordinate transformation [64]. Both the SV and MT algorithms will be reviewed below.

Given below is the definition for signed volume:

$$\text{SigVol}(\vec{a}, \vec{b}, \vec{c}, \vec{d}) = \frac{1}{6}[(\vec{b} - \vec{a}) \times (\vec{c} - \vec{a})] \cdot (\vec{d} - \vec{a}) \tag{3.3}$$

Assuming a single line segment whose start and end points are labeled $\vec{q_1}$ and $\vec{q_2}$, and a mesh triangle whose vertices are labeled $\vec{p_1}$, $\vec{p_2}$, $\vec{p_3}$, an intersection between the

line segment and the triangle can be determined by the boolean variable shadowMask,

$$\text{sign1} = \text{sign}(\text{SigVol}(\vec{q}_1, \vec{p}_1, \vec{p}_2, \vec{p}_3)) \tag{3.4}$$

$$\text{sign2} = \text{sign}(\text{SigVol}(\vec{q}_2, \vec{p}_1, \vec{p}_2, \vec{p}_3)) \tag{3.5}$$

$$\text{sign3} = \text{sign}(\text{SigVol}(\vec{q}_1, \vec{q}_2, \vec{p}_1, \vec{p}_2)) \tag{3.6}$$

$$\text{sign4} = \text{sign}(\text{SigVol}(\vec{q}_1, \vec{q}_2, \vec{p}_2, \vec{p}_3)) \tag{3.7}$$

$$\text{sign5} = \text{sign}(\text{SigVol}(\vec{q}_1, \vec{q}_2, \vec{p}_3, \vec{p}_1)) \tag{3.8}$$

$$\text{test1} = \quad \text{test1!} = \text{test2} \tag{3.9}$$

$$\text{test2} = \quad (\text{sign3} == \text{sign4}) \wedge (\text{sign3} == \text{sign5}) \tag{3.10}$$

$$\text{shadowMask} = \quad \text{test1} \wedge \text{test2} \tag{3.11}$$

where shadowMask is 1 if an intersection occured and 0 otherwise. This algorithm works for individual ray-triangle intersections, but is not well suited to large arrays of line segments and triangles. To use this algorithm on multiple rays and triangles, a custom vectorized version of this algorithm was developed. If $\vec{q}_1$ and $\vec{q}_2$ now define the start and end points of an array of line segments, and $\vec{p}_1$, $\vec{p}_2$, $\vec{p}_3$ define an array of mesh triangle vertices, then all of the quantities in the calculation of equation 3.11 can be manipulated into 3D arrays and the entire calculation can be performed using a single (very large) matrix multiplication. The code to build out the 3D matrices and perform the multiplication is given in python below,

```
1  import numpy as np
2
3  def SigVol(a,b,c,d, ax=1):
4    #NOTE: true signed volume would be multiplied by 1/6
5    return np.sum(np.cross(b-a, c-a) * (d-a), axis=ax)
6
7  #build out 3D arrays
8  Nt = len(p1)
9  q13D = np.repeat(q1[:,np.newaxis], Nt, axis=1)
10 q23D = np.repeat(q2[:,np.newaxis], Nt, axis=1)
11
```

```
12  #Perform SigVol Intersection Check
13  sign1 = np.sign(SigVol(q13D,p1,p2,p3,ax=2))
14  sign2 = np.sign(SigVol(q23D,p1,p2,p3,ax=2))
15  sign3 = np.sign(SigVol(q13D,q23D,p1,p2,ax=2))
16  sign4 = np.sign(SigVol(q13D,q23D,p2,p3,ax=2))
17  sign5 = np.sign(SigVol(q13D,q23D,p3,p1,ax=2))
18  test1 = (sign1 != sign2)
19  test2 = np.logical_and(sign3==sign4,sign3==sign5)
20
21  #1 if we intersected
22  shadowMask = np.logical_and(test1,test2)
```

This matrix multiplication version of the SigVol algorithm is very fast, but saturates RAM very quickly. It is therefore not an ideal algorithm, especially when the ROI contains a high resolution mesh or when there are a large number of intersection PFCs. It is preferable to iterate over one of the 3 dimensions in the aforementioned calculation, which preserves RAM and can be parallelized across multiple CPU cores. Assuming that the arrays $\vec{q}_1, \vec{q}_2, \vec{p}_1, \vec{p}_2, \vec{p}_3$ are all stored in a shared memory object (toolsClass from figure 3.4), the signed volume can be performed using a hybrid parallel-matrix method by calculating the intersection of each trace against all potential intersections in parallel. That is, by passing the index, i, corresponding to a single magnetic field line trace to individual CPU cores, each trace's intersections can be calculated in parallel. The python code is below, and would be included in a multiprocessing loop that iterates over i to pass all traces to CPUs for processing.

```
1  import numpy as np
2
3  def SigVol(a,b,c,d, ax=1):
4    #NOTE: true signed volume would be multiplied by 1/6
5    return np.sum(np.cross(b-a, c-a) * (d-a), axis=ax)
6
7  #build out 3D arrays
8  Nt = len(p1)
9  q13D = np.repeat(q1[i,np.newaxis], Nt, axis=0)
```

```
10    q23D = np.repeat(q2[i,np.newaxis], Nt, axis=0)
11
12    #Perform Intersection Test
13    sign1 = np.sign(SigVol(q13D,p1,p2,p3))
14    sign2 = np.sign(SigVol(q23D,p1,p2,p3))
15    sign3 = np.sign(SigVol(q13D,q23D,p1,p2))
16    sign4 = np.sign(SigVol(q13D,q23D,p2,p3))
17    sign5 = np.sign(SigVol(q13D,q23D,p3,p1))
18    test1 = (sign1 != sign2)
19    test2 = np.logical_and(sign3==sign4,sign3==sign5)
20
21    #return 1 if we intersected
22    if np.sum(np.logical_and(test1,test2)) > 0:
23          shadowMask[i] = 1
24    else:
25          shadowMask[i] = 0
```

The signed volume algorithm was the original algorithm used in HEAT, but it was eventually replaced with a more advanced algorithm that provides a significant speedup. The Möller-Trumbore (MT) algorithm is a ray-triangle intersection algorithm published in 1997 in the Journal of Graphics Tools [64]. The following equations, derivations, and descriptions, are taken from that publication. The MT algorithm uses barycentric coordinates in order to computer the intersection, and requires minimal storage because the plane equation need not be precomputed. A ray can be represented by the equation,

$$\vec{R}(t) = \vec{q_1} + t\vec{D} \tag{3.12}$$

where $\vec{R}(t)$ is the ray, $\vec{q_1}$ is the origin location, $\vec{D}$ is the ray direction and t is the ray length. A point on a triangle can be given in barycentric coordinates by the equation,

$$T(u,v) = (1-u-v)\vec{p_1} + u\vec{p_2} + v\vec{p_3} \tag{3.13}$$

where $(\vec{p_1}, \vec{p_2}, \vec{p_3})$ are the triangle vertices and (u, v) are the barycentric coordinates. The intersection of the ray and the triangle occurs at the point where $\vec{R}(t) = T(u, v)$,

or,

$$\vec{q}_1 + t\vec{D} = (1 - u - v)\vec{p}_1 + u\vec{p}_2 + v\vec{p}_3 \tag{3.14}$$

This formula represents the geometric translation of the triangle to point $\vec{q}_1$, and scaling it such that it is a unit triangle where $\vec{D}$ is aligned with the x-axis, as illustrated in figure 3.7 from [64]. This system of equations can be cast in matrix form,

$$\begin{bmatrix} -\vec{D} & (\vec{p}_2 - \vec{p}_1) & (\vec{p}_3 - \vec{p}_1) \end{bmatrix} \begin{bmatrix} t \\ u \\ v \end{bmatrix} = \vec{q}_1 - \vec{p}_1 \tag{3.15}$$

and using the following definitions,

$$|\vec{A}, \vec{B}, \vec{C}| = -(\vec{A} \times \vec{C}) \cdot \vec{B} = -(\vec{C} \times \vec{B}) \cdot \vec{A} \tag{3.16}$$

$$\vec{E}_1 = \vec{p}_2 - \vec{p}_1 \tag{3.17}$$

$$\vec{E}_2 = \vec{p}_3 - \vec{p}_1 \tag{3.18}$$

$$\vec{T} = \vec{q}_1 - \vec{p}_1 \tag{3.19}$$

$$\vec{P} = \vec{D} \times \vec{E}_2 \tag{3.20}$$

$$\vec{Q} = \vec{T} \times \vec{E}_1 \tag{3.21}$$

the linear system defined in equation 3.15 can be solved using Cramer's rule to yield the final result,

$$\begin{bmatrix} t \\ u \\ v \end{bmatrix} = \frac{1}{\vec{P} \cdot \vec{E}_1} \begin{bmatrix} \vec{Q} \cdot \vec{E}_2 \\ \vec{P} \cdot \vec{T} \\ \vec{Q} \cdot \vec{D} \end{bmatrix} \tag{3.22}$$

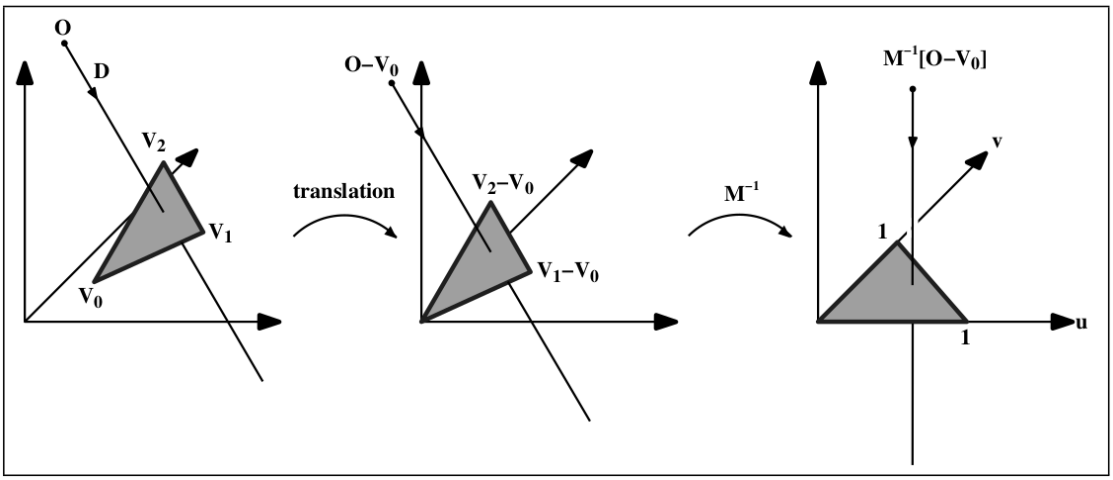


Figure 3.7: Coordinate transformation for Möller-Trumbore ray-triangle intersection checking algorithm from [64]. Note that here $\mathbf{v}_i = \vec{p}_i$ in the discussion.

The line segment intersected the triangle if the following tests all return true,

$$
\begin{aligned}
a &= (\vec{D} \times \vec{E}_2) \cdot \vec{E}_1 \\
\vec{s} &= (\vec{q}_1 - \vec{p}_1) \times \vec{E}_1 \\
l &= \frac{1}{a} \vec{E}_2 \cdot \vec{s} \\
\text{test1} &= a > 0 \quad \vee \quad a < 0 \\
\text{test2} &= u < 0 \quad \vee \quad u > 1.0 \\
\text{test3} &= v < 0 \quad \vee \quad (u + v) > 1.0 \\
\text{test4} &= l < 0 \quad \vee \quad l > |\vec{D}|
\end{aligned}
$$

These condition are needed to test for an intersection between a single ray and a triangle. HEAT uses a custom adapted version of the MT algorithm that has been parallelized, vectorized, and ported to line segments rather than rays. Similar to the previous discussion regarding the Signed Volume algorithm, it is assumed that the arrays $\vec{q}_1, \vec{q}_2, \vec{p}_1, \vec{p}_2, \vec{p}_3$ are all stored in a shared memory object (toolsClass from figure 3.4). The index, i, of a specific magnetic field line trace is passed to a worker CPU in a multiprocessing loop, and a check for intersection with any of the potential Intersects is performed. In Python, the intersection check that is performed on a single CPU using index i is as follows:

```
1   import numpy as np
2
3   D = D_vec[i] / np.linalg.norm(D_vec, axis=1)[i]
4   eps = 0.0
5   h = np.cross(D, E2)
6   a = np.sum(E1*h, axis=1)
7   f=1.0/a
8   s = q1[i] - p1
9   u = f * np.sum(s*h, axis=1)
10  q = np.cross(s,E1)
11  v = f*np.sum(D*q, axis=1)
12  l = f*np.sum(E2*q, axis=1)
```

```
13
14    test1 = np.logical_and( a>-eps, a<eps) #ray parallel to triangle
15    test2 = np.logical_or(u<0.0, u>1.0) #ray inside triangle
16    test3 = np.logical_or(v<0.0, (u+v)>1.0) #ray inside triangle
17    test4 = np.logical_or(l<0.0, l>D) #ray long enough to intersect triangle
18
19    #return 1 if we intersected
20    if np.sum(~np.any([test1,test2,test3,test4], axis=0))>0:
21        shadowMask[i] = 1
22    else:
23        shadowMask[i] = 0
```

Both of the aforementioned algorithms are options in the HEAT code, yet the MT algorithm provides a significant speedup when compared to the Hybrid Signed Volume algorithm. Table 3.1 provides a comparison of the three different aforemention algorithms that are included in HEAT. The speed benchmark for Möller-Trumbore (MT) vs Signed Volume (SV) algorithms was performed on single CPU core with 32 GB of Random Access Memory (RAM). N is number of ROI mesh rays and Nt is number of potential intersection mesh triangle elements. *Matrix* indicates that algorithm performs matrix multiplication with no for loops (aside from those inside numpy). As can be seen in the 'SV Matrix' column in the table, RAM saturates for full matrix implementation of the SV algorithm. *Hybrid* indicates that the algorithm collapses one dimension of 3D matrix multiplication into a for loop, to avoid saturating RAM. Note that in HEAT implementation of these algorithms, Hybrid versions will pass for loops across as many cores as possible, as mentioned earlier in this section.

Table 3.1 indicates that the MT Hybrid algorithm is the preferred method for calculating ray-triangle intersections on large meshes. While there is an advantage to using a full matrix implementation of the SV algorithm for small datasets, most HEAT datasets are large and contain thousands of rays and triangles. Both the SV algorithms scale poorly with the number of calculations to perform, when compared to the MT Hybrid method. For the case when there are 12k rays and 12k triangles in table 3.1, the MT Hyrbid algorithm is 2.3 times faster than the SV Hybrid algorithm,

Table 3.1: Speed benchmark for Möller-Trumbore (MT) vs Signed Volume (SV) algorithms on single CPU core with 32 GB of Random Access Memory (RAM). N is number of ROI mesh rays and Nt is number of potential intersection mesh triangle elements. *Matrix* indicates that algorithm performs matrix multiplication with no for loops (aside from those inside numpy). RAM saturates for full matrix implementation of SV, 'SV Matrix'. *Hybrid* indicates that the algorithm collapses one dimension of 3D matrix multiplication into a for loop, to avoid saturating RAM. Note that in HEAT implementation of these algorithms, Hybrid versions will pass for loops across as many cores as possible.

# Elements		Time [s]		
N	Nt	MT Hybrid	SV Hybrid	SV Matrix
1	1	0.0002	0.000179	0.00019
10	10	0.001098	0.001618	0.000234
100	100	0.009909	0.016214	0.002156
1000	1000	0.156431	0.274099	0.340941
10000	10000	7.684268	15.971446	35.052392
12000	12000	10.989756	25.257792	50.709693
15000	15000	18.184741	49.79649	RAM Overflow
20000	20000	33.621931	93.369408	RAM Overflow

and the 4.6 times faster than the SV Matrix algorithm. For the case when there are 20k rays and 20k triangles, the MT Hyrbid algorithm is 2.8 times faster than the SV Hybrid algorithm, and the SV Matrix algorithm overflows RAM and is killed. The MT Hybrid algorithm is therefore the default algorithm in HEAT, and is used in intersection calculations for the optical approximation as well as the gyro orbit approximation.

3.7.2 Acceleration Structures

A typical HEAT run consists of thousands of magnetic field line traces (or helical gyro orbit traces) launched from the ROI, traced for 10-20 steps (100s of steps), checked for intersections with 10s or 100s of thousands of potential Intersect mesh triangles. A 10 step trace for 10k mesh elements in the ROI and 100k potential intersect mesh elements amounts to N = 10 x 10k x 100k = 1E10 calculations. The intersection checking algorithms described in section 3.7.1 have been vectorized and parallelized to reduce this number, but still requires significant overhead when it comes to compute time. Section 3.1.2 listed *Speed* as a HEAT objective, and required that a modest resolution calculation be performed in under 10 minutes. Despite implementation of the Möller-Trumbore algorithm yielding a 3X speedup, the large number of potential intersection triangles still created a large computational cost, especially for gyro orbit heat load calculations where entire helical trajectories need to be traced. For this reason, means of speeding up HEAT were investigated. The current version of HEAT utilizes an acceleration structure to speed up the intersection calculation.

Many of the potential intersection triangles in the HEAT meshes will be nowhere near a given segment of the magnetic field line trace, and removing these triangles from the intersection calculation provides a significant speedup. This process can be thought of as creating a filter passband, and only triangles in the passband will be utilized in the intersection calculation. This filter can also be viewed as a form of dimensionality reduction, which reduces the 3D dataset to a single dimension. Figure 3.8 provides an illustration of this process. In HEAT, the line segment corresponding to a trace step can be defined by its origin, $\vec{q}_1$, and terminus, $\vec{q}_2$. Equation 1.49 can be used to convert from cartesian coordinates to the tokamak cylindrical coordinates. In equation 1.49, the toroidal coordinate is labeled t, but it is common in the tokamak literature to label

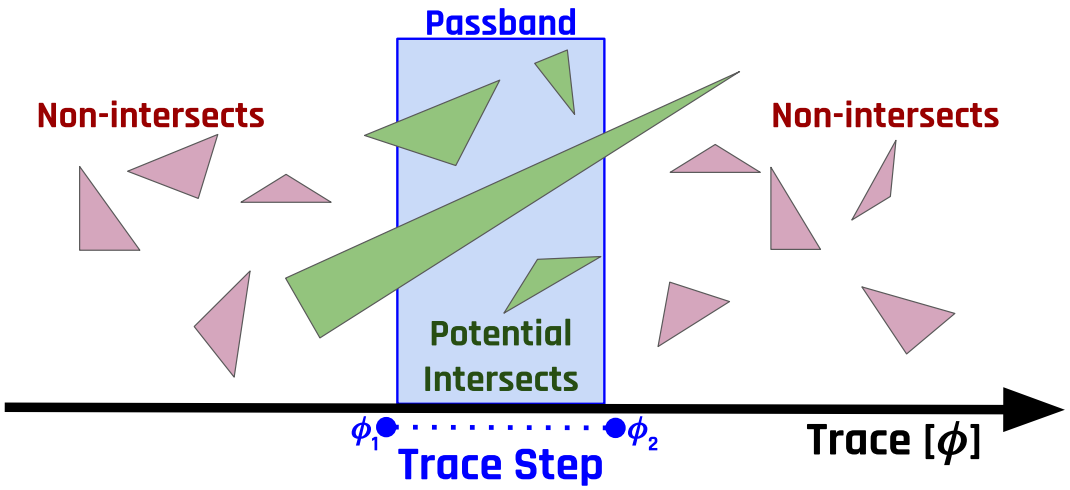


Figure 3.8: Cartoon of phi filter algorithm. Triangles in mesh are filtered according to their toroidal location. Triangles whose vertices are completely outside of the trace step are rejected, whereas triangles who have vertices within the trace step or straddling the trace step are in the passband. Intersection checking will only be performed on triangles in the passband, speeding up the intersection calculation.

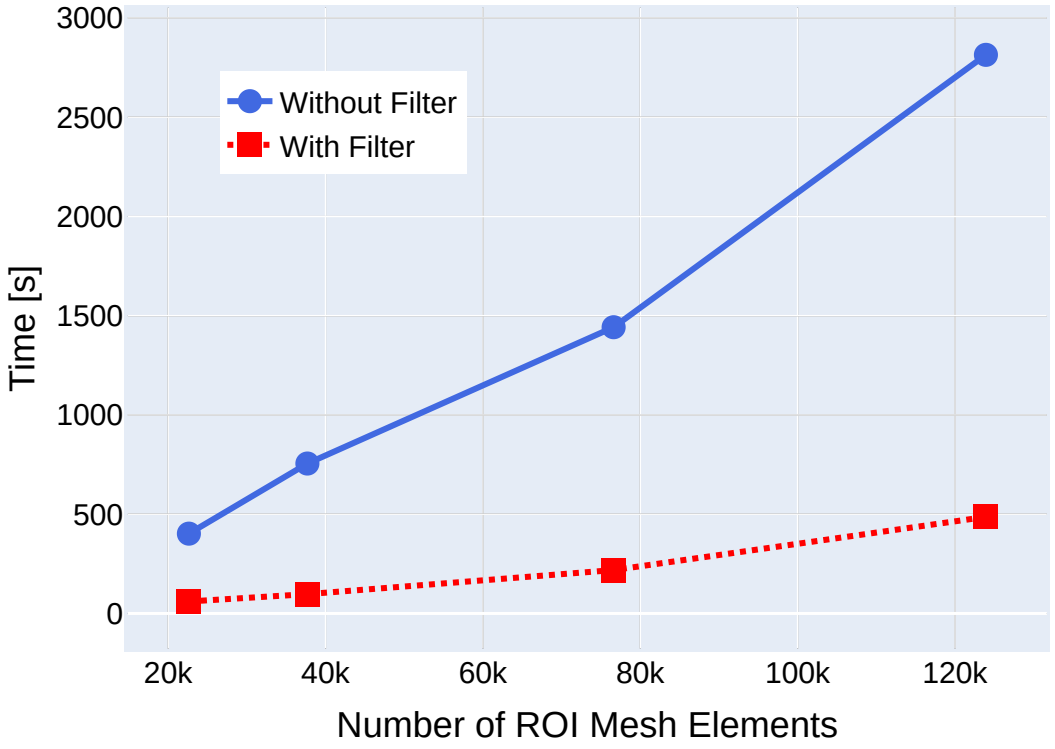


Figure 3.9: Speed comparison for toroidal angle filtering vs no filter. Results for fixed number of intersection mesh elements (68094), with changing ROI mesh Max Edge Length (5mm,4mm,3mm,2mm). Speedup is approximately 6X for all cases. Results shown for NSTX-U single IBDH tile in ROI, 6 tiles in Intersects, and 10 degree trace length. 8 CPU cores with 32GB RAM.

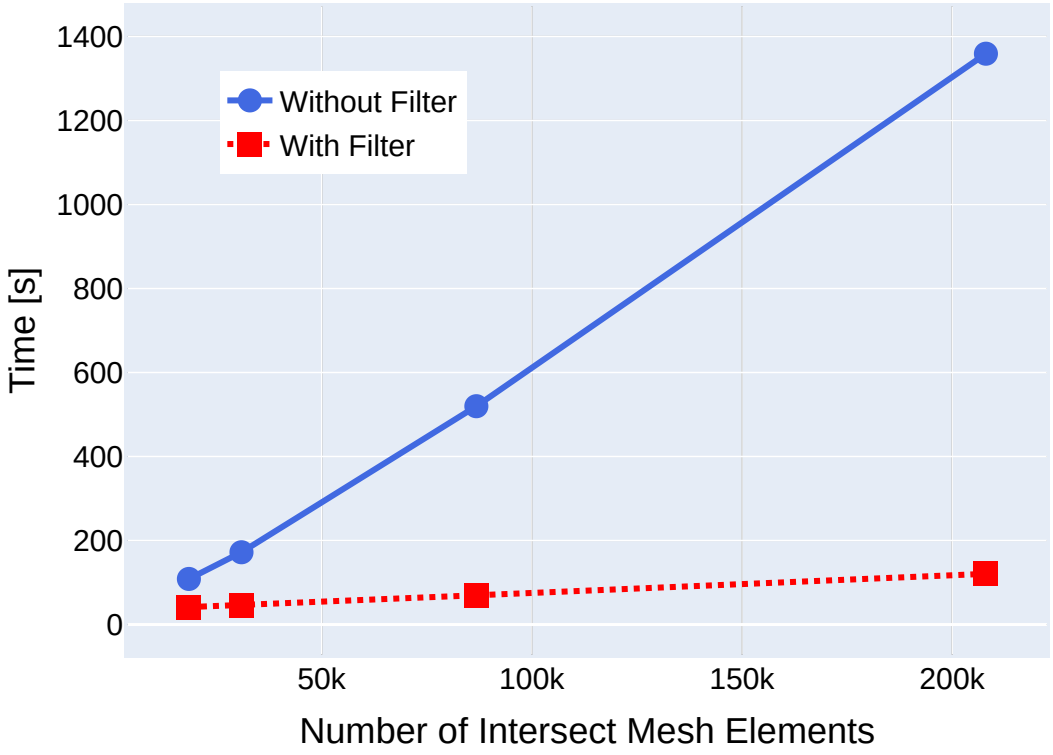


Figure 3.10: Speed comparison for toroidal angle filtering vs no filter. Results for fixed number of ROI mesh elements (22628), with changing Intersect mesh Max Edge Length (100mm,25mm,10mm,5mm). Speedup increasing with increasing number of potential intersects, and ranges from 2.6X to 11.3X. Increasing number of intersection CAD parts will not significantly increase calculation time, due to toroidal filter. Results shown for NSTX-U single IBDH tile in ROI, 6 tiles in Intersects, and 10 degree trace length. 8 CPU cores with 32GB RAM.

this coordinate ϕ. Both $\vec{q}_1$ and $\vec{q}_2$ will have a ϕ, and this coordinate can be utilized for to create a filter passband. Triangles in the mesh with vertices $(\vec{p}_1, \vec{p}_2, \vec{p}_3)$ can be filtered by comparing their toroidal coordinates to $\vec{q}_1$ and $\vec{q}_2$. More specifically, a triangle on the mesh will be rejected, therefore not included in the intersection calculation, if the binary variable 'Reject' below evaluates to true,

$$\text{test0} = (\phi_{p1} < \phi_{q_1}) \quad \wedge \quad (\phi_{p2} < \phi_{q_1}) \quad \wedge \quad (\phi_{p3} < \phi_{q_1}) \tag{3.23}$$

$$\text{test1} = (\phi_{p1} > \phi_{q_2}) \quad \wedge \quad (\phi_{p2} > \phi_{q_2}) \quad \wedge \quad (\phi_{p3} > \phi_{q_2}) \tag{3.24}$$

$$\text{Reject} = \text{test0} \quad \vee \quad \text{test1} \tag{3.25}$$

where ϕ_{q1} is the toroidal angle of $\vec{q}_1$, ϕ_{q2} is the toroidal angle of $\vec{q}_2$, and ϕ_{pi} is the toroidal angle of $\vec{p}_i$ for i=1,2,3. This filter only accepts triangles that have vertices on either side the passband, or triangles that have vertices inside the passband. A significant fraction of the triangles will be omitted from the intersection calculation using this filter, resulting in a sinificant speedup. Figures 3.9 and 3.10 illustrate the speedup obtained using the toroidal angle filtering. Both figures show two traces: one for the HEAT calculation time without toroidal angle filtering and one trace with toroidal angle filtering enabled. Figure 3.9 shows the speedup obtained when holding the number of Intersection mesh elements constant while increasing the ROI mesh resolution. Because the number of Intersect mesh elements does not change, and because the toroidal angle filter acceleration structure is a function of the number of Intersection mesh elements, the speedup is relatively constant at 6X for all ROI resolutions. In figure 3.10, however, the number of ROI mesh elements was held constant while the Intersect mesh resolution was increased. As the number of Intersect mesh elements increases, the speedup from the acceleration structure also increases.

This acceleration structure is designed to enable the HEAT user to load CAD STEP files containing many potential intersections parts. Because the toroidal filter automatically removes excess mesh elements from the calculation, the HEAT user need not worry about which PFCs in a CAD file are relevant, or could potentially cast magnetic shadows on the ROI. Combining this acceleration structure with the HEAT PFC CSV Input file 'all' option creates a completely autonomous method for calculating magnetic shadows, that simultaneously preserves speed and performance.

It should be noted that a similar acceleration structure exists in HEAT for filtering by ψ, the poloidal flux surface coordinate, but the speedup results are nearly identical to the aforementioned toroidal filter results so will not be shown here. Including both the toroidal and poloidal filters is typically only necessary for extremely large HEAT meshes. For moderate meshes, there is little to be gained from using both the toroidal and poloidal filters as using a single filter already rejects enough potential Intersects so that the calculation is not CPU limited, but rather is bandwidth limited (ie front side bus communication limited). If the user desires to use the poloidal filter, the blue toroidal angles (ϕ_1 and ϕ_2) of the trace step in figure 3.8 are replaced with the poloidal coordinates (ψ_1 and ψ_2). When using the poloidal filter, faces can occasionally be misidentified as in the rejection band, which results in extra heat loaded faces in the ROI. This is likely due to the fact that ψ is interpolated across the mesh surface, the mesh is an approximation of the true parametric surface, and there are regions of minimally changing ψ across PFC faces. For this reason, the default setting in HEAT is to calculate the potential Intersect faces using only the toroidal angle filter. For maximum acceleration, the user may select both filters in the HEAT GUI, which is beneficial for extremely large HEAT meshes. However the user must be careful to verify that all the faces that should be shadowed are indeed shadowed. By using the acceleration structures, HEAT has been run on meshes consisting of millions of elements in less than 10 minutes, which aligns with the original HEAT objective for speed.

3.8 Optical Approximation

This section describe the physics and algorithms developed in HEAT to simulate power flowing along magnetic field lines. It follows on background provided in chapter 1, specifically section 1.3.1.

3.8.1 Optical Heat Flux Calculation

In order to accomodate any magnetic and geometric configuration, an algorithm that provides a general physical model for power deposition was conceived and coded into

HEAT. This general model allows a user to prescribe the heat flux profile at an arbitrary origin location, as a function of poloidal flux. The function must be defined at an origin location along a monotonic chord of poloidal flux, usually taken to be the outboard midplane (OMP). Power is mapped from the origin location to the PFC targets along magnetic field lines. Because power is assumed to flow along field lines, cross field diffusion must be included in the user defined heat flux profile. The user defined heat flux profile is assumed to be axisymmetric, although 3D heat flux profiles are on the HEAT road map.

Presently no dissipation in the boundary plasma ($\nabla \cdot \vec{q} = 0$) is assumed, and power is assumed to flow along magnetic field lines ($\vec{q} = q_{\parallel} \frac{\vec{B}}{|B|}$). The mapping from origin location to target location is performed using the following equation, which is an adapted version of equation 1.14,

$$q_{target} = q_{||0} \cdot \hat{q}_{||}(\psi) \cdot \frac{|B_{target}|}{|B_{origin}|} \cdot (\hat{b} \cdot \hat{n}) \cdot (\neg \text{ShadowMask}) \tag{3.26}$$

where

- q_{target} is the surface heat flux at the target location;
- $q_{||0}$ is a scaling coefficient for the parallel heat flux;
- ψ is the poloidal flux at the target location;
- $\hat{q}_{||}(\psi)$ is the normalized user defined heat flux profile evaluated at ψ;
- B_{target} is the magnetic field at the target location;
- B_{origin} is the magnetic field at the location of ψ along the origin chord (OMP);
- $\hat{b}$ is the normalized magnetic field vector at the target;
- $\hat{n}$ is the target surface normal vector;
- $\neg$ is the logical negation (bitwise inversion) operator;
- ShadowMask is a binary variable that is 1 if face is magnetically shadowed.

HEAT allows the user a large amount of flexibility when defining the heat flux profile, $\hat{q}_{||}(\psi)$. Three options have been hard coded into the HEAT user interface for convenience. A 'Gaussian Spreading' option assumes power decays exponentially from the separatrix, and accounts for cross field diffusion by convoluting a gaussian with the exponential, as described in the well known work by Eich [15]. The Gaussian spreading term, S, can be calculated from the scaling by Makowski [65]. A 'Multi-Exponential' option assumes power decays as a superposition of multiple decaying exponentials each with an independent characteristic length as described in work by Brunner [16]. A 'Limiter' option assumes all power is transported in the common flux region as a superposition of two exponentials, as described by Horacek [66]. For each of these options, the characteristic widths (λ_q, S, etc.) are calculated using the aforementioned empirical scalings directly from the magnetic equilibrium, or may be defined explicitly by user input. Lastly, the HEAT user may also define a novel function for $\hat{q}_{||}(\psi)$, in approximately 30 lines of python code.

3.8.2 Power Balance

In order to confirm that the magnetic shadows and heat fluxes were calculated correctly, power balance is used as a post-simulation verification step. The user defines the power crossing the separatrix into the scrape off layer (SOL), P_{SOL}, in the HEAT user interface. After power is transported across the separatrix, further transport into the SOL is defined by the user defined heat flux profile, $\hat{q}_{||}(\psi)$. Presently, it is assumed that the power that crosses the separatrix is conserved and ultimately will be distributed on the tokamak wall along magnetic field lines. The normalized, user-defined, heat flux profile is scaled using the variable $q_{||0}$ along the origin chord (usually OMP) such that the following constraint is satisfied:

$$q_{||0} = \frac{P_{SOL}}{2\pi \int \frac{\hat{q}_{||}(\psi)}{B_{origin}(\psi)} \mathrm{d}\psi} \tag{3.27}$$

Defining $q_{||0}$ in this manner enforces power balance at the origin chord location, but does not guarantee power balance at the target. Power balance at the target will only be achieved if the ShadowMask has been correctly calculated via the intersection checking algorithm. After the simulation is complete, power mapped to each PFC

target face is added to create a summation power variable, P_{sum},

$$P_{sum} = \sum_i (q_{target,i} \cdot A_i) \tag{3.28}$$

where

- i iterates through each mesh face;
- $q_{target,i}$ is the heat flux calculated from equation 3.26 at mesh face i;
- A_i is the area of mesh face i.

P_{sum} is compared against the original user defined P_{SOL}. If the HEAT simulation was accurate the error percentage between P_{sum} and P_{SOL} is usually less than 5%. Mathematically the error percentage can be described by,

$$\text{Error } \% = \frac{|P_{SOL} - P_{sum}|}{P_{SOL}} \cdot 100\% \tag{3.29}$$

Oftentimes the error source is mesh artifacts or a low resolution mesh (sometimes desireable to increase the simulation speed). Increasing mesh resolution of the target and intersection PFCs will generally decrease the error in these cases. Extremely high values of P_{sum} often indicate that some mesh faces that are physically shadowed are not shadowed in simulation. Non-physical CAD geometry or too short of an intersection trace are usually the sources of very high P_{sum}. Ensuring that the CAD model includes all upstream objects that cast magnetic shadows upon the PFC of interest, as well as ensuring that the intersection trace is long enough to reach these upstream objects, will remedy these potential error sources.

It is important to note that the user is responsible for verifying that the magnetic shadows are correct. The power balance is used as an aid in the verification process, not as a definitive metric. The code does not flag specific magnetic shadow profiles as "correct" or "incorrect." This is intentional, as sometimes power balance will be "violated", but the magnetic shadow will still be correct. An example could be if the user only simulated a fraction of the power wetted area in the divertor, in which case the power balance would show a significantly lower P_{sum} than P_{SOL}. Requiring the

user to intelligently verify HEAT's results prevents any hard coded assumptions about the nature of correct shadow profiles.

While the optical approximation is a useful tool to perform quick heat load calculations, it overlooks the true interaction between PFC geometry and the helical trajectories of charged particles in magnetic fields. Section 3.9 describes the HEAT methods for calculating the gyro orbit heat loads, which includes these helical trajectories for ions. The next section describes some speed benchmarking performed for optical approximation calculations.

3.8.3 HEAT Optical Approximation Speed Benchmarking

This section provides a speed benchmark for the optical approximation in HEAT, and outlines relevant computational considerations. All results shown here have been generated on a desktop workstation with 8 logical cores (Intel(R) Xeon(R) W-2123 CPU @ 3.60GHz) and 32 GB of DDR4 RAM. The workstation runs Ubuntu 20.04 Linux, although the HEAT appImage has also been tested on Ubuntu 18.04, Centos 7, and Centos 8. Undoubtedly, HEAT runs faster with more cores and more RAM, but for this benchmarking excercise a modest machine was chosen to reflect the intended HEAT user's computational setup. For complex HEAT runs, it may be necessary to increase the number of CPUs in order to achieve the simulation goal of less than 10 minutes. HEAT is parallelized and will use as many cores as available, but preserves 2 cores for overhead (so it used 6 CPUs out of the 8 available on these tests).

In HEAT, heat flux calculations utilize a surface mesh in STL format (described in section 3.5.2). The STL mesh resolution is determined by constraining the maximum allowable edge length (decreasing the maximum edge length increases the resolution). It should be noted that the resolution is always finer than the maximum mesh edge length, as mesh elements are distributed randomly across the PFC surface. Increasing the mesh resolution increases the computational time, as indicated in table 3.2, which provides the maximum edge length, number of ray-triangle calculations, and simulation time, for 7 different example runs. These runs calculate the heat flux in the lower outer divertor, on the IBDH and OBD tiles (surfaces shown in Figure 4.3), for one timestep. HEAT simulation time is directly dependent upon the number of mesh elements. For

this NSTX-U PFC geometry the simulation time increases stronger than linearly as a function of maximum edge length, but increases linearly with the number of ray-triangle calculations.

The required mesh resolution is highly dependent upon the PFC geometry, and must be chosen to avoid heat flux aliasing on small features. While HEAT's heat flux calculation algorithms were designed to be general and make no assumptions about the geometry, the user must intelligently select a sufficient mesh resolution to resolve the physical features in question. When in doubt, the user can observe the power balance and simulation error to determine if the heat flux is being properly reconstructed. For PFC geometries that do not contain small features, such as the example provided in table 3.2, increasing the mesh resolution does not always increase the simulation accuracy. These results yielded a reconstruction error of $\approx 2\%$ for all resolutions that were sufficiently high to resolve the heat flux decay width (1-5mm).

There is no universally perfect mesh in HEAT. The maximum mesh edge length must be chosen to satisfy the user's intended objective for each application. HEAT is designed to be general, so no assumptions about the intended use case have been incorporated into the code. There is a separate optimum mesh resolution for each separate application. It is up to the HEAT user to determine the optimum mesh resolution for a given application.

For example, one use case may require HEAT to reconstruct the heat flux profile in between experimental discharges in an effort to qualitatively validate the physics models. In such a scenario a high resolution mesh will be undesirable (simulation time increases with the number of mesh elements) as the HEAT simulation may need to be completed in less than five minutes before the next discharge begins. However, a mesh resolution that is too low could result in mesh elements that are larger than the heat flux decay width, and would distribute power inaccurately. In Table 3.2 the 100 mm maximum edge length was too coarse to properly resolve the heat flux decay width, and as a result the Error % was nearly five times larger than the other simulations. A mesh resolution that is sufficiently high to capture the heat flux decay, but low enough to run quickly in between discharges, would be the optimum.

Another example HEAT use case would be utilizing HEAT to optimize inner limiter PFC geometry during startup. In this case, the HEAT simulation may need to be run

Table 3.2: Time benchmarking parameters for a single timestep heat flux calculation on the IBDH and OBD tiles. Maximum edge length is parameter used to define STL mesh, and is not the same as resolution. No. of Ray-Tri Calcs is total number of ray-triangle intersection calculations for all steps up field line. Simulation time includes postprocessing and figure generation. Benchmark performed on a simple 8 logical core (threads) desktop workstation with 32 GB of RAM.

Max Mesh Edge Length [mm]	# Ray-Tri Calcs	Simulation Time [s]	Error %
100	5.59 E 09	82.14	9.22
25	9.47E+09	85.49	4.51
10	2.55E+10	94.14	1.73
5	6.23E+10	146.06	2.62
4	1.16E+11	228.33	2.28
3	1.99E+11	424.05	2.66
2	3.25E+11	728.34	2.23

at a high resolution to capture the heat loads on narrow leading edges for magnetic field lines at shallow angles of incidence. The mesh resolution should be increased to sample the narrow geometric features. However, increasing the mesh resolution will require more computer RAM, CPU cores, and time, so there is a computational limit for the maximum mesh resolution. Additionally, increasing the mesh resolution beyond the level necessary to resolve the geometric features will yield diminishing returns. In Table 3.2, all maximum mesh edge lengths smaller than 5mm yielded similar error %, which illustrates that 5mm was sufficient to sample the geometric features for this PFC design. Balancing the available computational resources with the geometric feature size would be required to find the optimum mesh resolution for this case.

In addition to the number of CPU cores and mesh resolution, the length of the intersection tracing algorithm can change the HEAT simulation time. As described in section 3.7, HEAT traces magnetic field lines from the surface of the PFCs to determine the magnetic shadows in a process called 'intersection checking'. The HEAT user defines the length of these intersection traces in the user interface. Long intersection check traces result in more accurate calculations, but come with the cost of longer simulation times.

Figure 3.11 provides an example of the accuracy improvement observed when increasing the toroidal distance of the intersection trace, for a 5mm maximum mesh edge length simulation. When the intersection trace is set to 0° there is no magnetic shadow calculation, heat flux is prescribed neglecting magnetic shadows, and the error (from equation 3.29) is extremely large. Increasing the intersection check trace length from 0° to 15° reduces the error exponentially as the traces search upstream for divertor PFCs casting magnetic shadows. Increasing the length of the intersection trace beyond 15° does little to reduce the error, as these traces are likely on field lines that have long connection lengths and will ascend out of the divertor on the way to the outboard midplane. Simulation time increases linearly with intersection trace length, as shown by the blue curve and right y-axis. For the example provided in figure 3.11, HEAT can reconstruct the heat flux with less than 10% error in just over two minutes, and with 2% error in under three minutes.

The PFCs used in the simulations shown in figure 3.11 are 15° wide, corresponding to a geometric toroidal mode number of 24. As this is the toroidal mode number

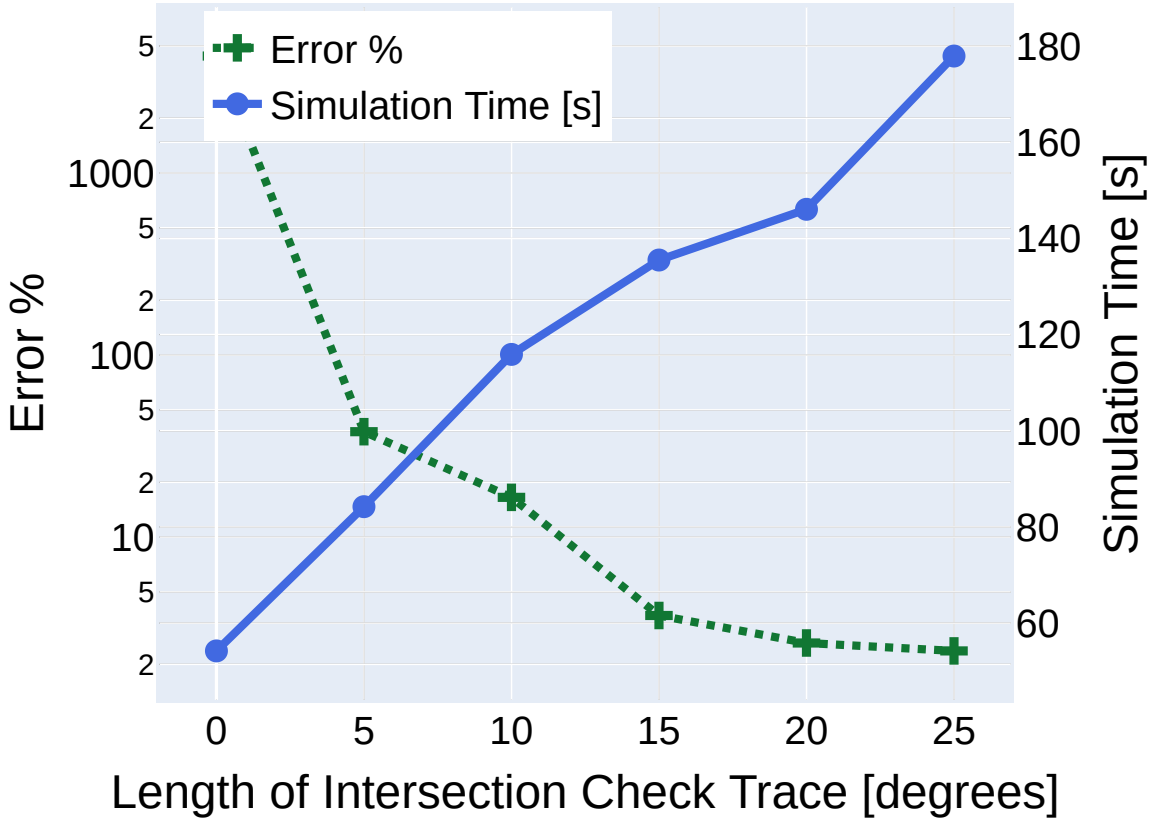


Figure 3.11: Length of intersection trace versus error % (from equation 3.29) and simulation time, for the lower divertor during NSTX-U discharge 204118 at 1004 ms. Increasing the length of the intersection trace leads to more accurate predictions, at the cost of simulation time. Test performed on an 8 logical core workstation machine. Mesh resolution of 5 mm with intersection resolution of standard.

associated with the largest periodic geometric feature in the divertor (ie the width of the PFC tiles themselves), most of the intersections calculated for this MHD equilibrium scenario will occur within the first 15° of the field line trace. Figure 3.11 confirms this notion, as the error improves by less than 1% when the trace is increased from 15° to 25°. This increase in trace distance does little to reduce the error, but comes at the cost of increasing the simulation time by nearly one third. If speed is the objective, then a magnetic field line trace that covers the smallest toroidal mode number is usually sufficient. That being said, increasing the trace distance beyond 15° does reduce the error, even if only by a small amount, as long range gap loading is detected at distances further than those associated with the smallest toroidal mode number. If simulation accuracy is the objective, then the trace distance can be increased until the error converges. The trace distance at which error convergence happens will be different for every PFC geometry and MHD equilibrium.

Finally, a note on running HEAT as a third-party user. Third parties can use HEAT via a Linux appImage available on the github page [61]. This appImage enables users to run a single Linux executable file, and eliminates the need to install dependencies. The appImage has been tested on Ubuntu 18.04, 20.04, and on CentOS 7, 8. A test case has also been provided on the HEAT github page, and a wiki is available to users who wish to install and run this test case. Because HEAT is open source under the MIT license it can be used by anyone. Additionally, because HEAT's graphical user interface is coded in HTML5 and can be accessed via a web browser, it is possible to install HEAT on a dedicated Linux server or cluster and access it from client machines across a local area network (LAN) from any OS with a web browser (Windows, macOS, Android).

3.9 Gyro-orbit Algorithms

This section describes the algorithms developed in HEAT to simulate particles following helical gyro-orbit trajectories. It follows background provided in chapter 1, specifically section 1.3.3. The gyro orbit algorithm outlined in this section follows the helical trajectories of ions; electrons are assumed to have negligibly small Larmor radii. The algorithm neglects the effects of electric fields. A discussion on the validity of this

assumption is provided in section 1.3.4, with the conclusion that the redistribution of power arising from the electric fields is small.

3.9.1 Gyro-orbit Macro-particles

If the user has turned on the gyro-orbit approximation switch, then the magnetic field will be traced upstream for a user defined number of steps from the center of each mesh element in the ROI. This reversed magnetic field line trace will be the forward guiding center trace of a helical gyro-orbit trajectory for a set of *macro-particles*. Macro-particles are used in HEAT to represent the behaviour of a chunk of particles from the velocity and spatial domains, without having to simulate every single particle. In the limit that the number of macro-particles approaches the true density of the plasma, the particles represent actual ions. However, in practice the number of macro-particles can usually be many orders of magnitude smaller than the true number of ions and still capture the physics with a tolerable amount of error. As the helical trajectories of the macro-particles are followed in the forward power flow direction (toward the PFC), the trace is checked to see if it intersects with any mesh element in the CAD. If there is an intersection, then ion gyro orbit power will be assigned to that mesh element.

For every mesh element in the ROI, j, HEAT calculates a pseudo-optical power. The pseudo-optical power can be interpreted as the power that a mesh element would receive if the ions were to follow purely optical (along magnetic field line) trajectories. The pseudo-optical power is calculated according to the optical approximation and shadow finding algorithms defined in other sections in this chapter. As mentioned in the last paragraph, a reverse magnetic field line trace for each mesh element, j, becomes the forward guiding center trajectory for the macro-particles. Once the ion gyro orbit trajectories for all macro-particles are traced and intersecting mesh elements are located, the pseudo-optical power for mesh element j is reassigned to the mesh elements, i, that were intersected along the helical trajectories of the macro-particles. Because each mesh element, j, will need to have many macro-particles traced to sufficiently sample the velocity and spatial domains, it is necessary to partition the pseudo-optical power into fractional components when reassigning to intersecting mesh element i. Sampling the velocity and spatial distribution functions will be discussed

in section 3.9.2. Constructing the fractional components for power reassignment will be discussed in section 3.9.3. The basic algorithm is summarized in enumerated form below for convenience,

1. Pseudo-optical heat flux is calculated at ROI surface mesh centers
2. Velocity and spatial distributions are partitioned per user specifications
3. Magnetic field is calculated at ROI mesh centers
4. Gyro orbit quantities are calculated ($\omega_c, v_\perp, v_{||}$, etc.) at mesh centers
5. Magnetic field is traced upstream for user specified distance from each mesh center
6. Macro-particle helical trajectories are followed downstream for each mesh center
7. Potential intersects are checked along trajectory
8. Fractional power for each macro-particle is calculated via distributions
9. If intersect is found, optical power for that mesh center is multiplied against fractional power and added to intersect mesh element.

For reference in subsequent sections, HEAT defines the thermal velocity to be,

$$v_{th} = \sqrt{\frac{2kT}{m}} \tag{3.30}$$

where k is the Boltzmann constant, T is the plasma temperature, and m is the particle mass. For ion gyro orbit simulations the user must define the plasma temperature, $T0$. The plasma temperature can either define a single value across the entire divertor or the user can upload a CSV point cloud with a temperature profile across the PFC surface. The user also must define the ratio of power sharing between electrons and ions, the length of the gyro orbit trace, the ion effective mass (charge assumed to be that of hydrogen ion, Z=1), and the number of discretization steps per helical period. Lastly, the user defines the number of samples to be taken from the distribution functions, as described in section 3.9.2.

3.9.2 Sampling the Distribution Functions

Newton's law and the Lorentz force equation (equation 1.25) can be solved to yield the equations of motion (equation 1.29) for a particle in a magnetic field with no electric field, which are repeated here for convenience,

$$\ddot{x} = \frac{qB}{m} v_y; \quad \ddot{y} = -\frac{qB}{m} v_x; \quad \ddot{z} = 0$$

where qB/m can be identified as the cycloctron frequency, ω_c, and v_i represents a velocity component in the i direction. These equations of motion can be integrated to yield the particle trajectories,

$$x = \frac{v_\perp}{\omega_c} \sin(\omega_c t + \alpha) + x_0 \tag{3.31}$$

$$y = \mp \frac{v_\perp}{\omega_c} \cos(\omega_c t + \alpha) + y_0 \tag{3.32}$$

$$z = v_{||} t + z_0 \tag{3.33}$$

where t represents time, (x_0, y_0, z_0) represents the guiding center location, α represents the gyro orbit phase angle, the $\mp$ accounts for diamagnetism (- for ions), and $v_\perp$ can be defined by the following energy sharing relationship,

$$v^2 = v_{||}^2 + v_\perp^2 \tag{3.34}$$

$$v_{||} = v \cdot \sin(\beta) \tag{3.35}$$

$$v_\perp = v \cdot \cos(\beta) \tag{3.36}$$

$$\tag{3.37}$$

where v represents the particle instantaneous speed, $v_\perp$ represents the component of velocity perpendicular to the magnetic field, $v_{||}$ represents the component of the velocity parallel to the magnetic field, and β is called the velocity phase angle. Equations 3.31 - 3.36 are the equations necessary to solve for individual macro-particle trajectories. These five equations contain three unknown variables: (α, β, v), which, when chosen, constrain the family of potential trajectory solutions to a single trajectory. Each of the unknown variables must be selected by sampling from a distribution function that

corresponds to the physics of that particular variable. The user specifies the number of discrete samples that are to be taken from each of these distributions, where N_{gP} is the number of gyro phase angle samples, N_{vP} is the number of velocity phase angle samples, and N_{vS} is the number of speed (sometimes called velocity in this text) samples. Each of these distributions is discussed below.

The gyro phase angle represents the orientation of the perpendicular velocity vector with respect to the magnetic field line, and is represented by α in equations 1.34 - 1.35. The gyro phase angle ranges from (0,2π), and samples are taken such that they are uniformly distributed between these bounds.

Figure 3.12 provides some examples of gyro-orbit helical trajectories and the associated gyro phase angles. The left and center frames of the figure both show guiding center trajectories (blue) and particle gyro-orbit trajectories (red), for a fixed helix (ie fixed speed, radius, and pitch). The purple plane depicted in both frames has a radius equal to a single gyro radius, and the magnetic field is normal to it at t=0. For this example, the particle (green dot) can start its trajectory from anywhere on the circumference of the purple plane, corresponding to varying gyro phase angle, α. Two example gyro phase angles are shown. Additionally, the right frame of figure 3.12 provides an example distribution of gyro phase angles for 5 samples (NgP=5). In HEAT, the user specifies the number of samples to be taken from the uniform gyro phase distribution.

The velocity phase angle defines the relationship between the parallel and perpendicular velocity components, $v_{||}$ and $v_{\perp}$, and is represented by β in equations 1.34 - 1.35. For a given total velocity magnitude, v, there are infinite combinations of $v_{||}$ and $v_{\perp}$. In order to constrain $v_{||}$ and $v_{\perp}$, a unique combination of these components that satisfies the relationship, $v = \sqrt{v_{||}^2 + v_{\perp}^2}$, must be chosen. This constraint represents the energy sharing between the parallel and perpendicular velocity components.

Figure 3.13 shows an example distribution for a case when the user has requested 4 velocity phase samples, NvP=4. Neglecting particles that flow backwards up the magnetic field line ($v_{||} > 0$) and using the fact that the plasma is diamagnetic ($v_{\perp} > 0$), leaves only the first quadrant of velocity space. Therefore the velocity phase angle ranges from (0,π/2). Samples are taken uniformly between these bounds, but the bounds themselves are omitted because if $\beta = \pi/2$ the trajectory reduces to the optical

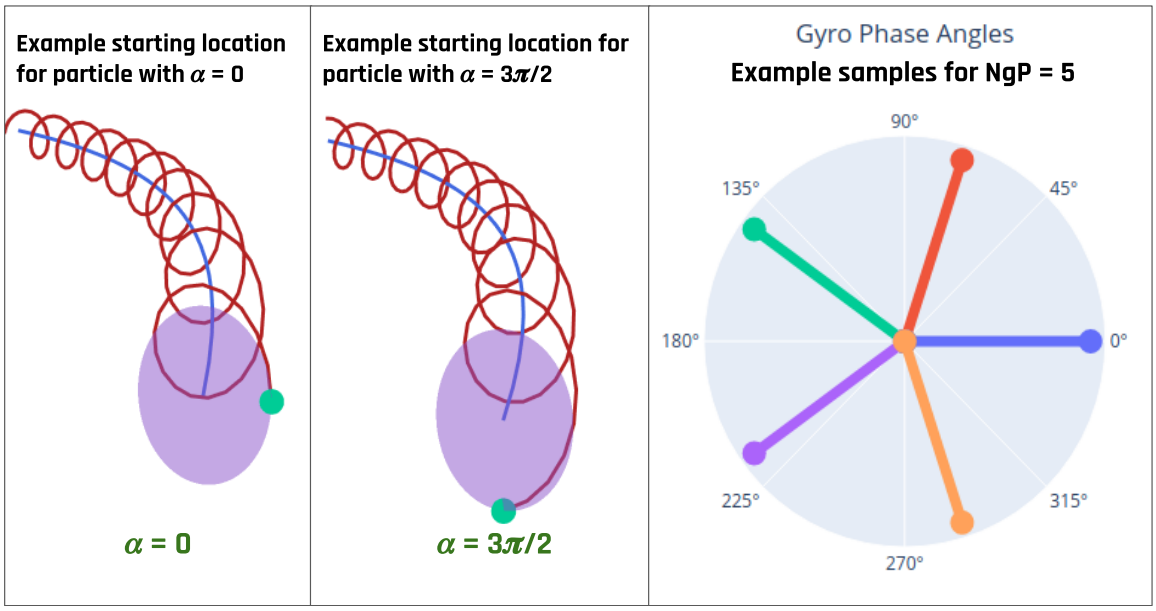


Figure 3.12: Guiding center trajectory (blue), particle trajectory (red), and initial particle location (green dot), for two examples with initial gyro phase angle of $\alpha = 0$ (left) and $\alpha = 3\pi/2$ (center). Magnetic field is tangent to circular purple plane at t=0, which has radius equal to 1 gyroradius. The particle can originate anywhere on the circumference of the plane. Gyro phase angles are uniformly distributed from $(0, 2\pi)$, and user selects number of samples, NgP. Right frame shows example samples for NgP=5.

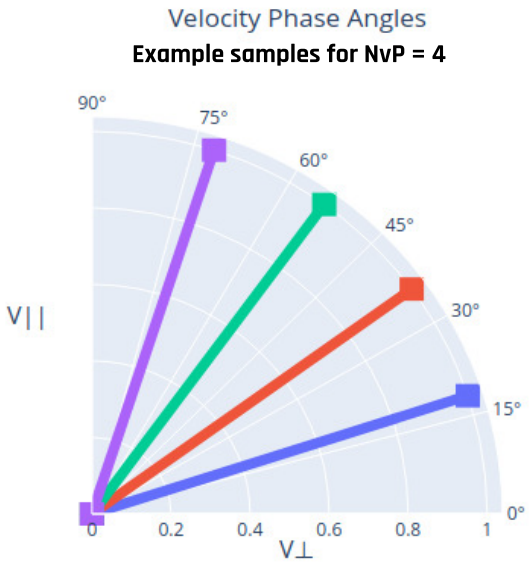


Figure 3.13: Velocity phase angles (β) represent the energy sharing between perpendicular and parallel components of velocity. Example shown for NvP=4, corresponding to 4 samples uniformly distributed through velocity phase angle space. Note that $\beta = 90°$ and $\beta = 0°$ are not included as they represent the optical approximation and cyclotron motion, respectively.

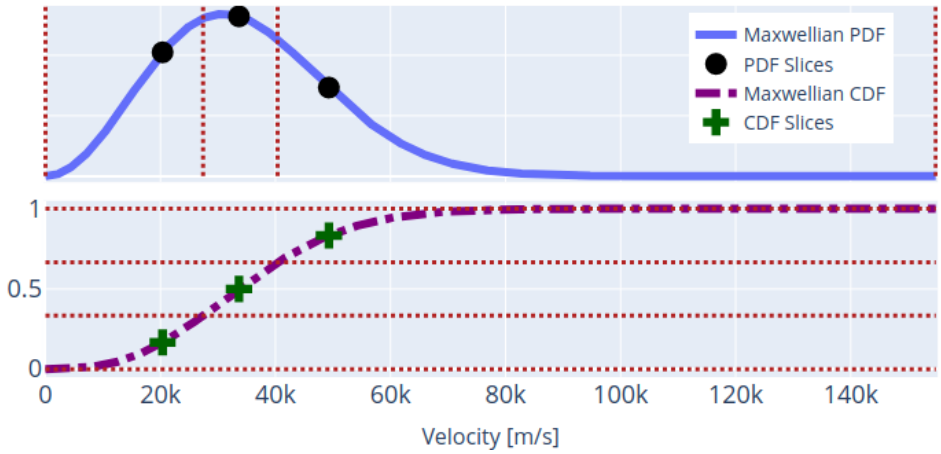


Figure 3.14: Example Maxwellian speed distribution with velocity slice samples overlaid on top. Corresponding cumulative distribution function (CDF) on bottom. For this example NvS=3, corresponding to three samples. Red dotted lines represent the boundaries of numerical bins. Bins are plotted as a function of velocity for PDF (upper) plot, and as a function of the CDF for CDF (lower) plot. Bins are chosen such that each bin is equally probable, and velocity samples are placed in center of bins. See text for more discussion on CDF, binning, and sample locations.

approximation and for $\beta = 0$ the trajectory is purely cyclotron motion. Omitting the bounds results in NvP+2 samples, where only the middle NvP samples are used.

Far upstream from the PFCs, the distribution of particle speeds is assumed to be Maxwellian. HEAT uses the Maxwellian speed distribution function, as described by equation 1.42, in order to calculate the total speed, v. The distribution is numerically truncated at $5 \times v_{th}$. Figure 3.14 provides an example Maxwellian speed distribution function, and 3 samples, also called velocity slices or vSlices, overlaid on top of this function. For this example, the user has selected 3 velocity slices, NvS=3.

The Maxwellian speed distribution is sampled such that each sample has approximately equal probability. This is achieved by splitting the cumulative distribution function (CDF) of equation 1.42 into NvS bins, and making each bin equally probable. The bin boundaries are thus spaced at regular intervals throughout the CDF, which has a total range of (0,1). The CDF interval width can be calculated as 1/NvS. The vSlices (HEAT velocity sample locations) are the velocities that correspond to the center point of these CDF bins. The CDF bin boundaries are called the cdfBounds, the locations of the samples in the CDF are called the cdfSlices, and the velocity samples are called vSlices. For the example provided in figure 3.14 with NvS=3, the bin boundaries would yield cdfBounds=[0.0, 0.333, 0.667, 1.0], and the CDF sample locations would be cdfSlices=[0.167, 0.5, 0.833]. For a 10eV deuterium ion this corresponds to velocity bins at vBounds=[0, 27400, 40379, 154770] m/s and samples located at vSlices=[20375, 33659, 49281] m/s (remember that the PDF is truncated at 5 vThermals). Note that the middle slice in figure 3.14 is skewed to the right of vThermal (peak of Maxwellian) as it represents the location where the CDF=0.5, not the location corresponding to vThermal (CDF value for vThermal is ≈ 0.43).

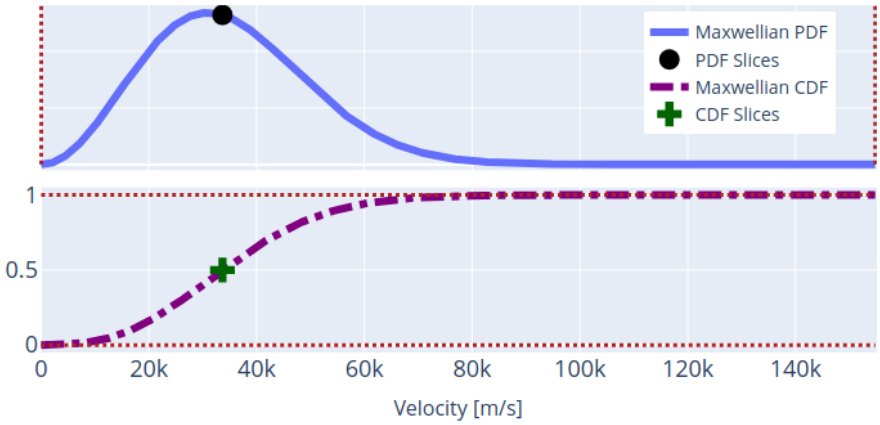
Figure 3.15 provides three additional examples of sampling the Maxwellian distribution function, where NvS is varied. The top figure shows NvS=1, the middle figure shows NvS=2, and the bottom figure shows NvS=5. In each figure, the PDF plot and the CDF plot are stacked vertically for comparison. The PDF plot shows the bins plotted as a function of the x (velocity) axis, whereas the CDF plot shows the bins plotted as a function of the y (CDF) axis. Taking a single sample, as shown in the top figure, results in a bin that is the width of the entire CDF (ie 0 to 1). The vSlice

is placed in the center of this bin, which corresponds to a CDF value of 0.5, as seen on the y-axis on the CDF plot. This CDF value corresponds to a speed of 33659 m/s. The middle figure in figure 3.15 illustrates how an even number of samples results in vSlices that straddle the PDF peak (vThermal). For the middle plot with NvS=2, there are two CDF bins and the cdfBounds=[0.0, 0.5, 1.0]. Placing a velocity sample in the center of each of these bins results in cdfSlices=[0.25, 0.75]. These cdfSlices correspond to vSlices=[24094, 44357] m/s, which straddle the peak of the distribution function. When the number of bins is increased to NvS=5, the samples provide relatively good coverage of the PDF when compared to the other examples, but will come at the cost of longer simulation times. An example for NvS=5 is shown in the bottom figure.

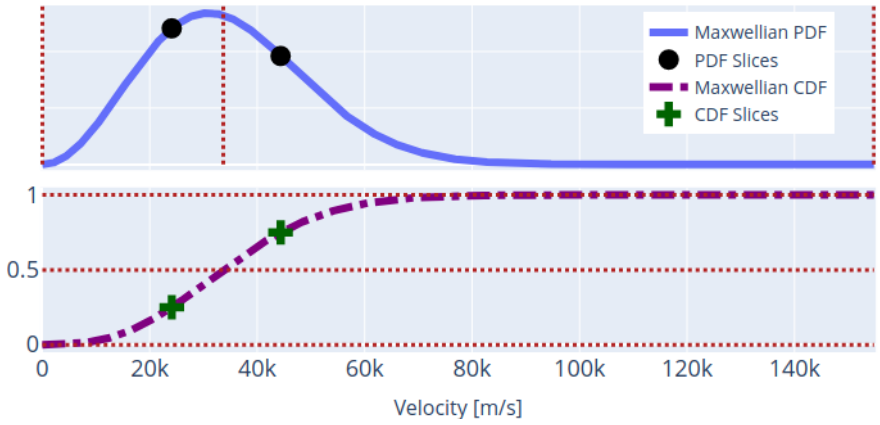
Various binning approaches have been experimented with in HEAT over the course of gyro-orbit module development, but they all yield similar results with respect to power depostion and temperature analysis, so the exact algorithm is not extremely important so long as the sampling is approximately uniform in the velocity CDF space. The current algorithm uses evenly spaced CDF bin boundaries, although this could change in future versions of the code. Regardless of the binning technique utilized, the power assigned to each macro-particle, or sample, should be calculated by integrating the energy PDF between the boundaries. This is described in section 3.9.3.

Once samples from the gyro phase angle (α), the velocity phase angle (β), and the particle speed (v), have been chosen, the macro-particle trajectories can be calculated. Each sample from each of the aforementioned domains (α, β, v) represents a single macro-particle, or a single helical trajectory, that will be traced. Each helical period is discretized into a user defined number of steps, called N_gyroSteps. Line segments are stitched between these discrete step positions, and the parallel ray-triangle intersection algorithms described in section 3.7.1 are used to determine if the helical trace intersects with any of the CAD mesh in the Intersect List. Once an intersection is found, the power that will be assigned to it is also determined from these distribution functions. The next section describes the power fractions that are used to reassign power on the mesh.

(a) Example vSlices and cdfSlices for NvS = 1



(b) Example vSlices and cdfSlices for NvS = 2



(c) Example vSlices and cdfSlices for NvS = 5

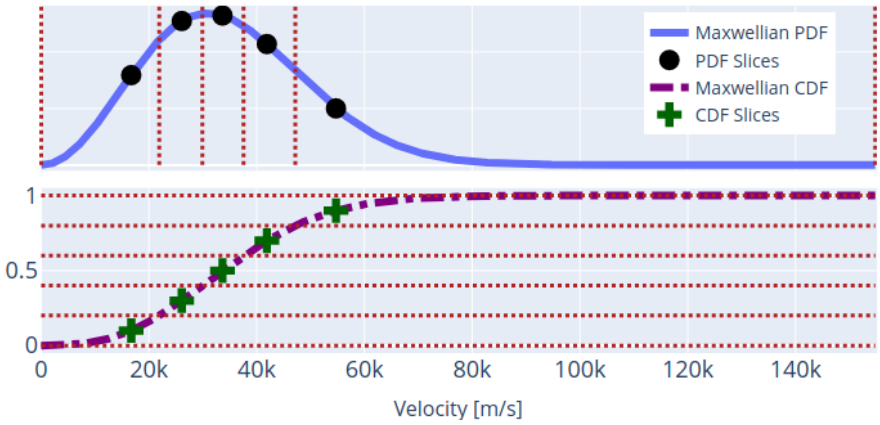


Figure 3.15: Comparison of velocity Maxwellian PDF and CDF for three examples cases with varying number of velocity samples (vSlices). Numerical bin boundaries are shown in red dotted vertical and horizontal lines. Bins are chosen such that each bin is equally probable, and velocity samples are placed in center of bins.

3.9.3 Macro-particle Power Fractions

HEAT samples the distribution functions as described in section 3.9.2 to define macro-particles, then follows the trajectories of these macro-particles looking for intersections with mesh elements. Each macro-particle is traced and the mesh element index of the intersection location is recorded in a matrix called intersectRecord. After all macro-particles have been traced, the psuedo-optical power for each particle is reassigned to the corresponding intersection mesh element with an appropriate weighting factor to reflect the numerous macro-particles. More specifically, the ion gyro orbit power to any mesh element, i, is described by the following equation,

$$P_{i,gyro} = \sum_{jklm} (P_{j,optical} \cdot \text{ionFrac} \cdot \text{gPFrac}_{j,k} \cdot \text{vPFrac}_{j,l} \cdot \text{vSFrac}_{j,m} \cdot \text{intersectRecord}_{i,jklm}) \tag{3.38}$$

where

- i is the mesh element gyro power is being reassigned to
- j is the mesh element pseudo-optical power is being reassigned from via macro-particle klm
- k corresponds to the macro-particle gyro orbit phase angle, α
- l corresponds to the macro-particle velocity phase angle, β
- m corresponds to the macro-particle velocity slice (speed) from the Maxwellian
- $P_{j,optical}$ is the power that would have been assigned to element j using the optical approximation
- ionFrac is the fraction of power carried by ions (as opposed to electrons)
- $\text{gPFrac}_{j,k}$ is the fraction of power prescribed to the k gyro phase angle for distribution j
- $\text{vPFrac}_{j,l}$ is the fraction of power prescribed to the l velocity phase angle for distribution j
- $\text{vSFrac}_{j,m}$ is the fraction of power prescribed to the m velocity slice for distribution j

- intersectRecord$_{i,jklm}$ is a binary matrix that is 1 if macro-particle $jklm$ strikes mesh element i

Because the gyro phase angle and velocity phase angles are both sampled uniformly from their respective distributions, the power fractions for these variables is defined by,

$$\text{gPFrac} = \frac{1}{N_{gP}} \tag{3.39}$$

$$\text{vPFrac} = \frac{1}{N_{vP}} \tag{3.40}$$

whereas the power fraction for each sample (vSlice) from the Maxwellian speed distribution function is determined by energy. First, for each speed sample (vSlice) from the Maxwellian, the corresponding velocity bin is determined using the CDF of the distribution defined in equation 1.42. Note that every mesh element can potentially have a different temperature, so has a unique distribution function. The Pythonic algorithm to determine the boundaries of the velocity bins from the CDF boundaries appears as follows,

```
1  import numpy as np
2
3  #generate the CDF from an assumed pdf (v_pdf)
4  #  and corresponding 1D speed array (v):
5  v_cdf = np.cumsum(v_pdf[1:])*np.diff(v)
6
7  #create bspline interpolators for the cdf inverse
8  inverseCDF = interp1d(v_cdf, v, kind='linear')
9
10 #CDF location of bin boundaries
11 cdfBounds = np.linspace(0,v_cdf[-1],N_vSlice+1)
12
13 #makes vSlices center of these bins in CDF space
14 cdfSlices = np.diff(cdfBounds)/2.0 + cdfBounds[:-1]
15
16 #vSlices are pdf sample locations (@ bin centers)
17 vSlices = inverseCDF(cdfSlices)
```

```
18
19 #velocity bin boundaries
20 vBounds = inverseCDF(cdfBounds)
```

This results in the vBounds variable, which contains velocity bin boundaries for each speed sample (vSlice) from the Maxwellian distribution function, and is directly linked to the CDF boundaries described in section 3.9.2. As is shown in the python code, transforming the CDF boundaries to velocities results in the vBounds. The energy fractions are determined by converting the velocity bins to energy bins, integrating the bin to find the bin energy, then dividing the energy in each bin by the total energy for the entire distribution function. For a given mesh element, i, the Pythonic algorithm for determining the vSFrac (also called energyFrac) is as follows

```
1  import numpy as np
2
3  #energy distribution function
4  f_E = lambda x: 2 * np.sqrt(x / np.pi) * (T0[i])**(-3.0/2.0) * np.exp(-x /
       T0[i])
5
6  #energy integrals
7  for j in range(N_vSlice):
8      Elo = 0.5 * (mass_eV/c**2) * vBounds[j]**2
9      Ehi = 0.5 * (mass_eV/c**2) * vBounds[j+1]**2
10     energyIntegrals[i,j] = integrate.quad(f_E, Elo, Ehi)[0]
11 energyTotal = energyIntegrals[i,:].sum()
12 #energy fractions
13 for j in range(N_vSlice):
14     energyFracs[i,j] = energyIntegrals[i,j] / energyTotal
```

The mathematical definition of the energy fraction for each velocity slice is given as follows,

$$\text{vSFrac}_m = \frac{\int_{Ea}^{Eb} f_{E,j}(E)dE}{\int_0^{\infty} f_{E,j}(E)dE} \tag{3.41}$$

where vSFrac_m is the velocity slice fraction for sample m from the Maxwellian distribution function, Ea and Eb are the energy bin boundaries for a specific vSlice, and $f_{E,j}(E)$ is the Maxwellian for mesh element j defined by equation 1.47.

Once the fractions for gyro phase angle (gPFrac), velocity phase angle (vPFrac), and velocity slice (vSFrac), have been determined, the total power for the macro-particle can be found. The portion of equation 3.38 inside the summation represents the power carried by macro-particle (k, l, m). The summation serves to calculate the power superposition of all of the macro-particles launched from all the source mesh elements. The process for creating gyro-orbit power sources is described in section 3.9.4.

3.9.4 Gyro Source Planes

The preceding gyro orbit approximation discussion has described a process where pseudo-optical power is reassigned to new mesh elements based upon kinetic theory distribution functions. Originally, the pseudo-optical heat loads were calculated on the same PFCs that the gyro orbit heat loads were calculated upon, but it was soon discovered that the helical nature of the gyro orbits resulted in some power that would be scraped off and reassigned far upstream. In order to avoid simulating a large toroidal fraction of the tokamak, a new method was necessary to prevent the power in the ROI from being reassigned to mesh elements outside of the ROI.

This method involves creating a gyro source plane - a 2D plane placed at a location upstream that will serve as the source for the gyro orbit power. The pseudo-optical heat flux is calculated on the gyro source plane at the user defined ROI heat flux resolution, and then the helical trajectories are traced directly down to the PFCs in the ROI, which is why the plane has been called a source. Figure 3.16 provides a graphical demonstration of the gyro source plane. The magneta plane in the upper right corner of the figure is the gyro source plane, and the grey object in the lower left corner is the PFC of interest. The pseudo-optical power will be calculated on the gyro source plane, then magnetic field line traces will be calculated from each mesh center in the gyro source plane. These magnetic field line traces that flow downstream towards the PFC of interest are the guiding centers for the gyro orbit helical trajectories, which are

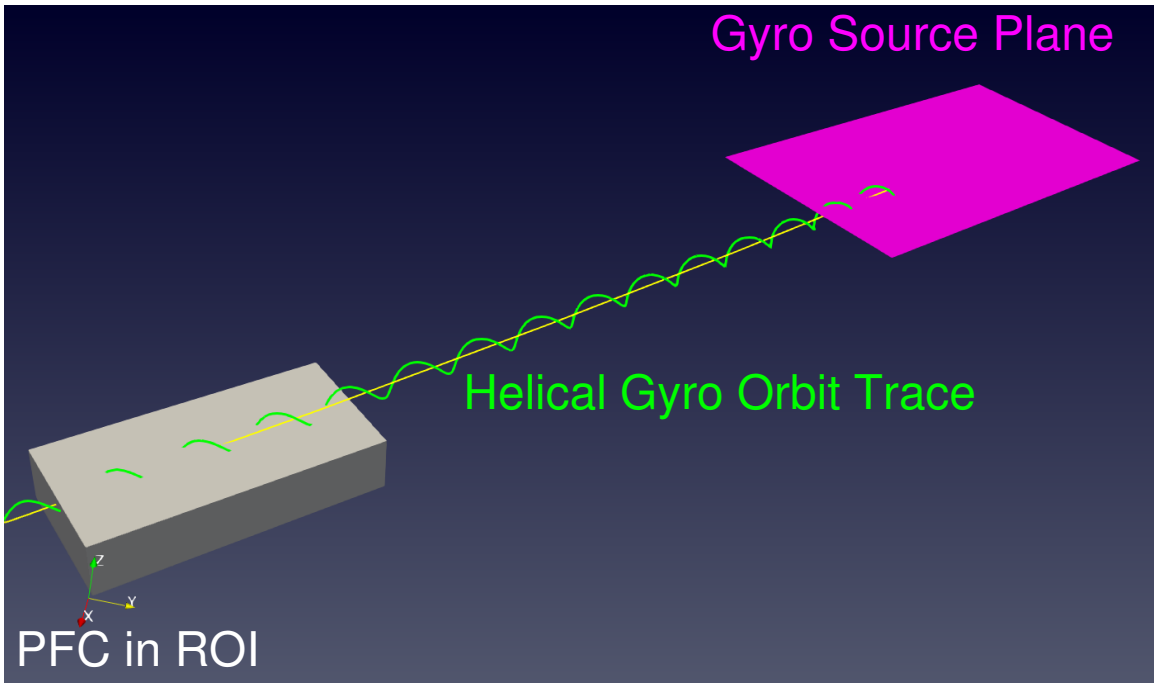


Figure 3.16: Example of a gyro source plane. Object in light gray at bottom left is PFC of interest in the ROI. Magenta plane in upper right is gyro source plane. Magnetic field line trace and helical gyro orbit trajectory are launched from gyro source plane and strike the PFC of interest.

calculated using the discussions in preceding sections. When the helical traces intersect a PFC in the ROI, the power is reassigned.

Using gyro source planes enables the gyro orbit calculation to be performed much faster than if the power were being reassigned from the PFC in the ROI to itself. This is because the shallow angles of incidence require tracing helical trajectories far upstream to the location where they will be scraped off. By creating a gyro source plane that is spatially closer to the PFC in the ROI, the trace distance can be reduced, thereby speeding up the HEAT calculation. The user must include the gyro source plane in the CAD STEP file, and change the PFC input file appropriately. The gyro source plane will not affect the optical calculation, as the user places the name of the gyro source plane in the PFC input file excludeName column for the PFCs in the ROI. Another benefit to using a gyro source plane is that power balance can be utilized as a test of accuracy. Similar to the discussion in section 3.8.2, power balance will only be achieved if the gyro orbit trajectories and associated shadows were calculated correctly by the intersection checking algorithm.

3.10 Temperature Calculation

HEAT uses openFOAM (C++), a robust partial differential equation toolkit, to perform finite volume method (FVM) calculations. A volume mesh is snapped to the PFC surface mesh using the openFOAM **snappyHexMesh** utility. Temperature throughout the PFC is initially set to a fixed temperature, and the temperature gradient on the surface is determined by combining the HEAT calculated heat flux and Fourier's equation for heat conduction in 1D,

$$\nabla T = -\frac{q_{target}}{\kappa(T)} \tag{3.42}$$

where

- T is the temperature;
- q_{target} is the HEAT calculated heat flux on the surface;
- $\kappa(T)$ is the temperature dependent thermal conductivity.

The temperature gradient on the PFC surface is interpolated in space and time using the openFOAM utility **timeVaryingMappedFixedValue**. Finding the temperature throughout the PFC now amounts to solving the temperature dependent heat equation with a Neumann boundary condition. A custom openFOAM solver called 'heatFoam' was developed to constrain the heat equation with the Neumann boundary condition, interpolate the temperature gradient in space and time, interpolate the temperature dependent thermal properties with respect to temperature, iteratively step through time solving the heat equation, and write the data to files accessible to HEAT. The heat equation is explicitly given here,

$$\frac{\partial T}{\partial t} = \alpha(T)\nabla^2 T; \qquad with \quad \alpha(T) = \frac{\kappa(T)}{\rho c_p(T)} \tag{3.43}$$

where

- $\alpha(T)$ is the temperature dependent thermal diffusivity;
- ρ is the density;
- $c_p(T)$ is the temperature dependent specific heat capacity.

In HEAT's current revision, it is assumed that power is conducted to the target and no radiation occurs between PFCs. For simple temperature calculations this is sufficient, but future work will aim to add complexity to this model. Because the heatFoam solver is a custom openFOAM solver, it is possible to change the solver as desired with minimal coding. Future work will add radiation from the plasma, inter-PFC blackbody radiation, stresses, and multi-region solutions, as described in the next section.

3.11 HEAT Roadmap

HEAT development is ongoing and this sections aims to highlight some of the upcoming features that will be added to HEAT. It also highlights potential machine learning loops that can be utilized for design optimization and machine control. Some text in this section is taken from [28].

3.11.1 Upcoming Physics Modules

While there are a wide range of physical phenomena and tokamak operational constraints that can be investigated with the toolkit in its current form, several advanced physics modules are currently being developed to augment the heat flux predictions. A radiated power module that can calculate photon power radiated from the plasma to the PFCs is already under development. All the results presented in this paper assume that there is no volumetric dissipation in the SOL, $\nabla \cdot \vec{q} = 0$. While this assumption is useful for ascertaining the maximum heat loads for modern tokamaks, many future tokamaks will operate in a partially detached regime in order to keep the PFCs within their allowable heat flux limits, as described in section 1.5.5. For high power density tokamaks, radiated power can be multiple MW m^{-2}, and loads plasma facing surfaces and sub-surface features differently from surfaces loaded by heat conducted along open field lines. For example, PFC mounting hardware in deep recesses that are magnetically shadowed may become heat loaded because they are not line-of-sight shadowed. This module will incorporate power radiated from the plasma, as well as inter-PFC blackbody radiation. These radiation sources will generally increase the heat flux and temperatures in the PFCs they are incident upon. The module will accept an input file describing the axisymmetric radiated power profile throughout the SOL, and will distribute the corresponding photon radiation to the engineering CAD PFCs. Self consistency with the conducted power can be enforced by using plasma fluid simulation results or by subtracting the radiated power defined heueristically via integrating the losses along a field line.

Incorporating radiation sources into HEAT is likely to change the output for specific simulation scenarios, but not significantly for the results presented herein. Because the simulations presented in this dissertation are fully attached plasmas, the power radiated from the plasma will be small compared to the power transported via ions and electrons. Additionally, the peak heat flux in attached plasmas is likely to be a result of ions and electrons, not photons. Including inter-PFC blackbody radiation would likely increase the global PFC temperature in the aforementioned results slightly, but the short time-scale transient thermal responses would be relatively unperturbed. However, if one were to simulate a full experimental day with multiple discharges, including blackbody

radiation would be necessary to capture temperature ratcheting effects that occur on multiple hour timescales.

A filament model that can investigate heat loads from SOL turbulence and Edge Localized Modes (ELMs) is under development. The present HEAT capability focuses on heat that is transferred along field lines in a stationary plasma, but convective instabilities like ELMs exhibit competition between energy draining along the field and energy radially propagating in the far SOL. Examples of this work have been shown as part of the ITER first wall design [67]. Being able to investigate power deposited on structures that are in the shadow between limiters will be useful in designing RF antenna structures and diagnostics for high power density tokamaks. The heat flux conducted to divertor PFCs during an ELM can also be captured on average using empirical scalings, similar to the inter-ELM heat flux width [40]. Section 1.5.2 described ELMs.

Finally, a module that incorporates 3D magnetic fields will be added to HEAT. The current version of HEAT assumes an axisymmetric plasma. Simulating an ideal 2D (axisymmetric) plasma incident upon 3D PFCs eliminates the complexity associated with toroidal variations in the magnetic field, but still captures the response of the 3D geometry. Unless the plasma axisymmetry is significantly perturbed, the 2D plasma approximation is sufficient. For the cases when axisymmetry is perturbed, future work will enable the HEAT user to investigate the change in heat flux that arises from toroidal variation in magnetic field using the 3D plasma MAFOT module. These toroidal variations can be intentional, such as when using Resonant Magnetic Perturbations to control ELMs. They can also be unintentional, such as the error fields due to coil installation misalignments and described in section 1.5.1. Because the MAFOT field line tracer already has an interface to many 3D plasma codes (M3D-C1, SIESTA, VMEC), adding the 3D capability to HEAT should be greatly simplified.

3.11.2 HEAT in Machine Learning Loops

Not only does the upcoming HEAT development roadmap include adding new physics to HEAT, but it also includes using HEAT to generate datasets for artificial intelligence (AI) or machine learning (ML) algorithms. Because HEAT predicts the 3D PFC heat

loads, and because it can predict engineering parameters that will constrain machine operation (ie temperature), HEAT trained ML algorithms would be ideal candidates for both design optimization and real time control. Figure 3.17 illustrates how HEAT can be utilized in the loop for ML algorithm generation and training. First, a database that covers the physics, operational, and geometry, domains is generated with HEAT. This database can then be used to train artificial neural networks as described in [50], or other ML algorithms.

The Engineering Design loop in figure 3.17 uses HEAT output to optimize CAD geometry. Because HEAT already has the capability to modify CAD geometry using FreeCAD, optimizing a specific CAD feature amounts to gradient descent as a means to minimize a user defined parameter of interest (ie heat flux or temperature). This loop does not need to be controlled by an artificial intelligence (AI) or ML algorithm, but for high dimensional problems where many geometrical features can be manipulated, an AI/ML approach may accelerate the optimization process by navigating through design space.

The Control loop shown in figure 3.17 uses the HEAT database to train a reduced model AI/ML algorithm. HEAT maps plasma MHD equilibrium and CAD geometry to parameters such as temperature and stress in the PFCs. These parameters are important to the tokamak operator because they directly constrain the machine operational domain. Creating a ML reduced model of the HEAT algorithm that can run in sub-ms timescales would enable 3D PFC effects to be included in real time control of the tokamak, directly from plasma parameters. This loop could be used to either 1) shutdown the machine if a PFC operational limit is violated or 2) adapt the MHD equilibrium to minimize the PFC heating.

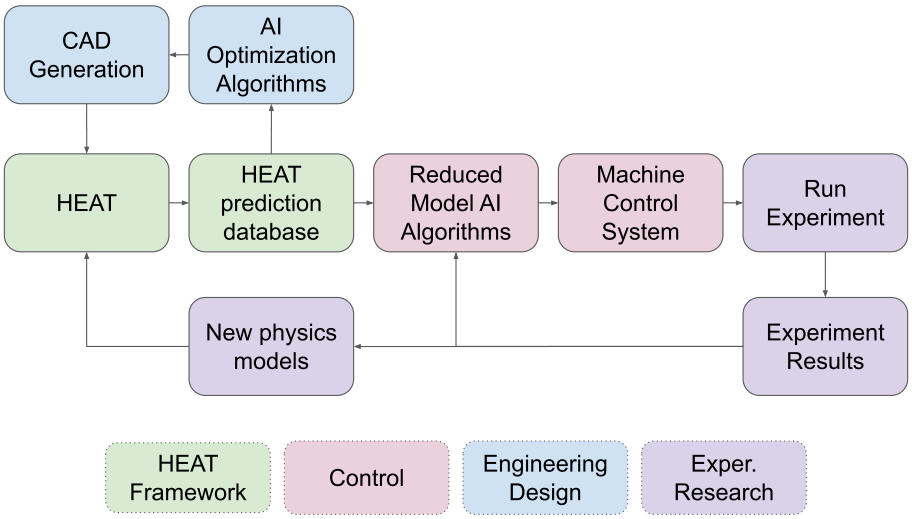


Figure 3.17: Example HEAT simulation loops to train machine learning algorithms. Blue loop provides design optimization to change CAD geometry. Red loop creates reduced model for machine control system.

Chapter 4

HEAT Optical Approximation Results for NSTX-U

This section provides optical approximation results for NSTX-U. First, a full 360 degree static analysis is performed for the lower divertor. Another lower divertor static analysis is performed that couples the heat flux outputs to temperature predictions. A time varying strike point sweeping analysis quantifies the effects of changing the strike point sweep frequency on the machine operational domain. Magnetic equilibria from 2016 are used to generate heat load predictions for diverted and limited discharges. And finally, examples of using HEAT during the NSTX-U construction process are included. The primary finding of this section is that HEAT can be utilized to directly ascertain the time varying boundary of the tokamak operational domain as a function of the PFC engineering limits.

4.1 Full Lower Divertor Static Discharge

In order to achieve 5 second, high power, stationary discharges, NSTX-U utilizes high poloidal flux expansion to mitigate the heat load. These high poloidal flux expansion discharges often have angles of incidence less than 1°, which can result in magnetic shadows cast by PFCs more than 10° upstream. HEAT was used to simulate these shadows, heat fluxes, and temperatures in the PFCs and evaluate if these discharges

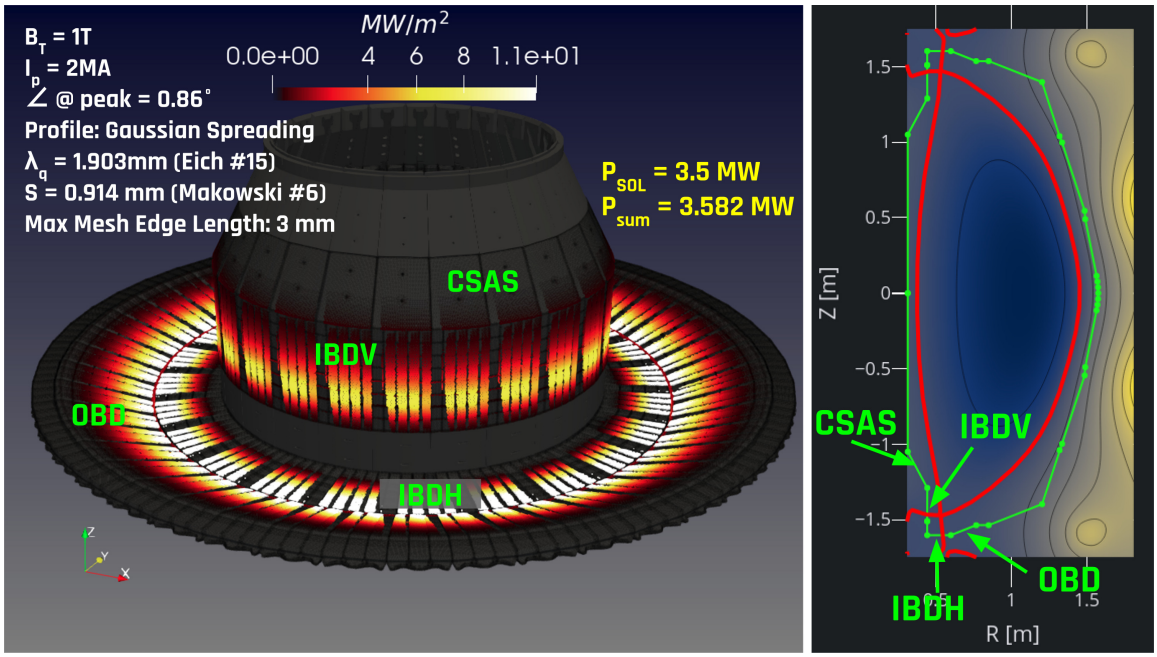


Figure 4.1: Lower divertor 360° heat flux footprint and magnetic equilibrium from PFC Working Group Memo 010 Case 1.1. Note the various toroidal mode numbers associated with magnetic shadows, which correspond to periodic geometry patterns. Figure from [28].

can be run for the intended 5 second duration without violating the PFC temperature limits.

Figure 4.1 provides a visualization of the magnetic equilibria and associated heat flux footprint for a PFC working group case, see Memo 010 Case 1.1 [68]. This case was generated with 1 T toroidal field (B_t), 1 MA plasma current (I_p), and 3.5 MW of power conducted to the lower divertor (30% directed to the inner divertor and 70% directed to the outer divertor). Using these parameters the heat flux width (λ_q) is calculated using Eich's regression 15 [15] to be 1.903 mm. The gaussian spreading term (S) is calculated to be 0.914 mm from Makowski's figure 6 [65] assuming a Greenwald density fraction of 0.6. After finding magnetic shadows and mapping power to the lower divertor using the Gaussian Spreading option in HEAT, the reconstructed power (P_{sum}) is calculated to be 3.582 MW, corresponding to a simulation error of approximately 2.4%.

Figure 4.1 illustrates the complex nature of the magnetic shadows. These shadows, caused by geometric features such as fish scales, repeat toroidally. Several different toroidal mode numbers can be observed in the shadow pattern, corresponding to different fish-scale step sizes in the CAD geometry. It should be noted that while HEAT makes it possible to calculate the full 360° heat flux footprint for every discharge, as shown in figure 4.1, it is usually preferable to decrease the computational time by simulating a smaller toroidal section, as shown in figure 4.2.

Figure 4.2 provides a visualization of two PFC tiles in the lower divertor as viewed from above. The heat flux pattern is observed to be non-axisymmetric, with wetted area fractions that vary poloidally. The two magnetic shadow mode numbers are caused by two different fish-scale step sizes. The narrow shadow bands that repeat toroidally on every castellation are a result of the smaller inter-castellation fish-scale step size. The large band of no heat flux on the left-most row of castellations on both tiles is a result of a large fish-scale between PFC tiles. This inter-tile fish-scale must be large to allow for PFC fabrication and installation tolerances. For this magnetic equilibrium, including magnetic shadows in the heat flux calculation decreases the heat flux loaded surface area by more than 50% when compared to an axisymmetric PFC.

The temperature at the end of this discharge is plotted alongside the heat flux pattern, and the colorbar has its upper bound set to the assumed temperature limit of graphite, 1873 K. Temperature dependent material properties are those of Sigrafine

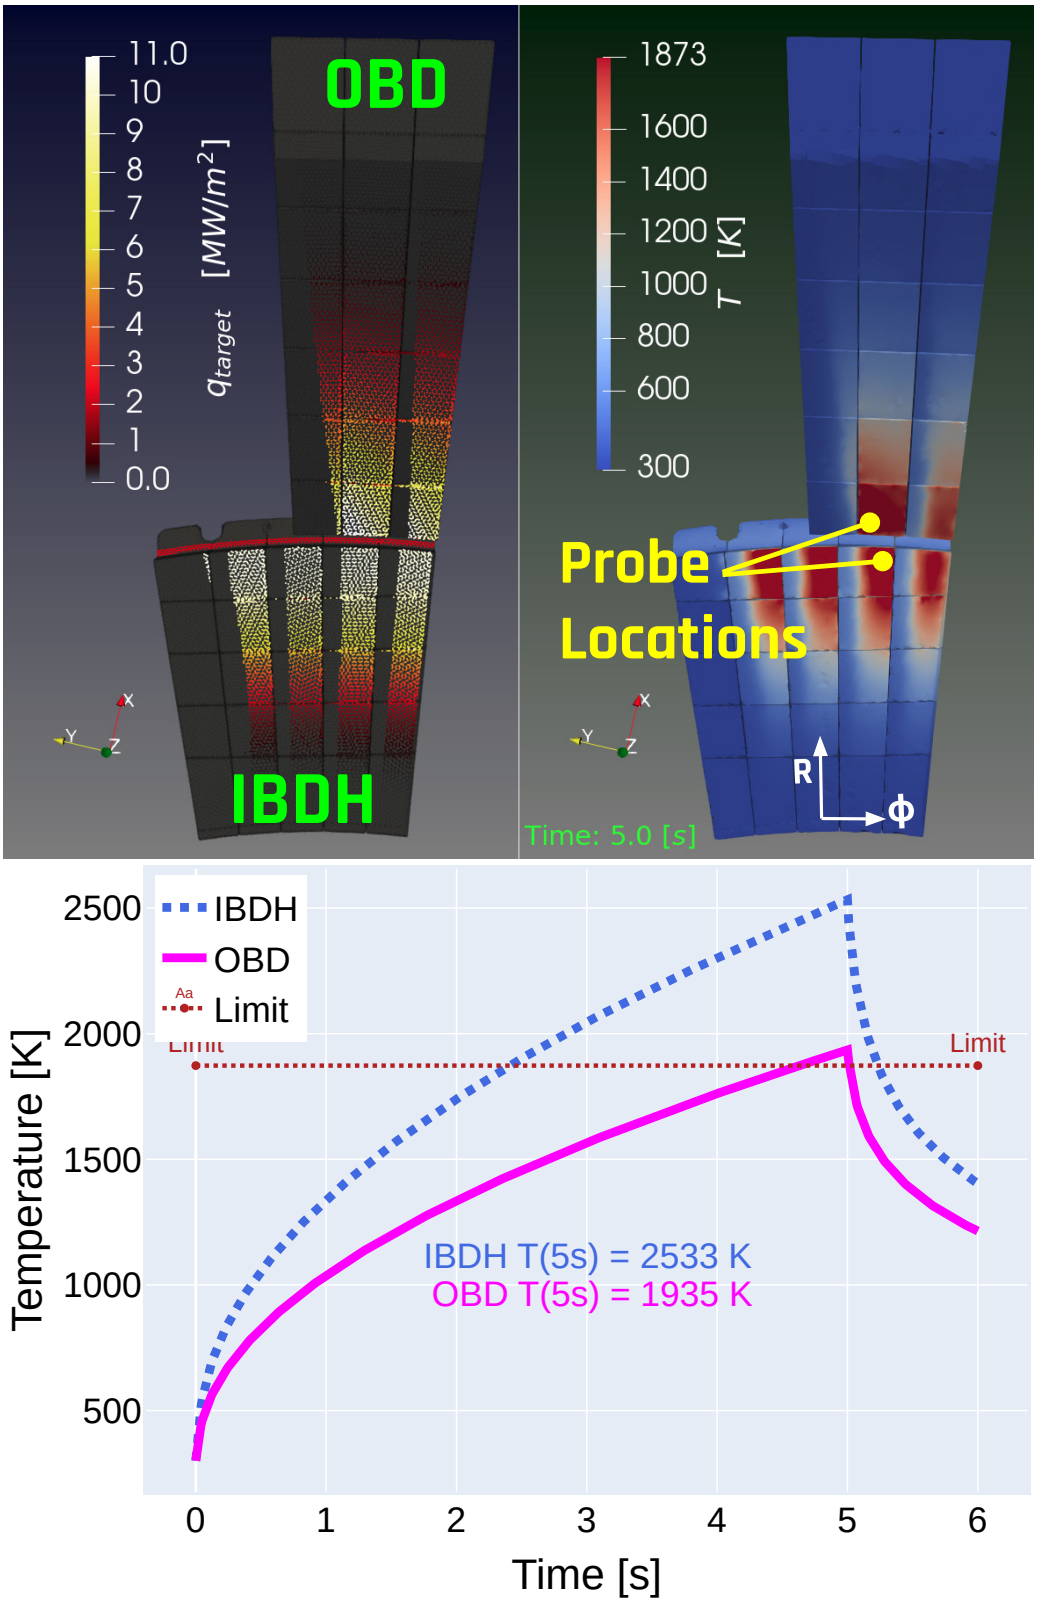


Figure 4.2: Simulated heat flux (top left), temperature at end of discharge (top right), and thermal probe trace (bottom), for Memo 010 Case 1.1. Trace indicates both PFCs exceed the sublimation engineering limit for graphite (1873 K). Figure from [28].

R6510 isostatically pressed graphite. Both PFC tiles have areas that are in excess of the assumed temperature limit at 5 seconds, and would likely be sublimating small amounts of graphite into the SOL. Using HEAT's temperature probe feature, synthetic thermal probe signals can be generated using the finite volume solver. Two synthetic thermal probes are overlaid onto the temperature visualization in figure 4.2, and time evolving temperature traces are provided for each location in the lower plot of the figure. The probes are placed on the PFC surface at the location of highest calculated heat flux. The temperature probe traces confirm that both PFCs are above the engineering limit at the end of the discharge.

HEAT's unique ability to couple magnetic equilibria, divertors physics, engineering CAD, and thermal analysis, enables it to map out the machine operational space directly as a function of the engineering limits. The NSTX-U graphite PFCs are thermally limited by sublimation at 1873 K, so running this discharge for the intended 5 second duration will exceed this limit. The IBDH tile crosses the sublimation temperature at approximately 2.43 seconds into the discharge. HEAT is unable to evaluate the effect of sublimination on the plasma (and determine associated reduction of heat flux), but results of HEAT temperature evolution could be coupled to standard boundary plasma codes as a 2D or 3D carbon source.

Preventing the PFCs from reaching the temperature limit can be achieved through several methods, such as decreasing the shot duration, lowering the parallel heat flux, increasing the flux expansion, or sweeping the strike point. Decreasing the shot duration is the simplest method, but would require halving the duration to be effective. Lowering the parallel heat flux, $q_{||}$, can be achieved by either changing the power crossing the separatrix, P_{SOL} (which in turn limits core performance), or by stimulating and controlling radiation on open field lines. Increasing the flux expansion will result in more complicated magnetic shadows, and can redistribute power in unforeseen ways. Sweeping the strike point requires dedicated divertor coil operation, including power supplies to drive the fast current variation. Each of these methods comes with trade-offs that constrain the tokamak operational space differently. HEAT can be used to explore these trade-offs, from an operational or tokamak design vantage point, and find a method to minimize the side-effects or reactor cost. The next section provides an example of using a strike point sweep to lower the peak PFC temperature.

4.2 Dynamic Strike Point Sweeping

For some high power NSTX-U scenarios, a stationary discharge will push the PFCs above the engineering limit temperature for exposure times of ≈1 second. In these cases sweeping the strike point across the PFC surface can reduce the peak temperature, thereby increasing the allowable discharge length. Several strike point sweep magnetic equilibria scans were developed by the NSTX-U PFC Working Group, and have been recently analyzed with HEAT to determine the feasibility of strike point sweeping with respect to the PFC engineering limits. PFC Working Group Memo 010 Case 2 Scan 4 [68] is one such strike point sweep generated with 1 T toroidal field (B_t0), 2 MA plasma current (I_p), and 3.5 MW of power conducted to the lower divertor (P_{SOL}). The sweep is achieved by stitching together six different magnetic equilibria, each with identical plasma parameters but with divertor coils adjusted to provide a different strike point location. Each of these equilibria are separate time-independent ISOLVER simulations, and future work is to use time-dependent equilibrium solvers which include vessel currents. A visualization of temperature during a constant velocity 5s sweep is shown in figure 4.3, corresponding to a strike point sweep frequency, f_{sweep}, of 0.1 Hz. At the beginning of the discharge, the strike point is on the OBD, and as time evolves the strike point moves onto the IBDH. The colorbar in the figure is bounded at 1873 K, the graphite (sublimation) engineering limit temperature. At the end of the discharge, a large section of the IBDH tile is above this temperature limit, as represented by dark red.

Figure 4.4 shows the time history of peak temperature in the IBDH tile for three strike point sweep frequencies: 0.1 Hz, 1 Hz, and 10 Hz. Generally, increasing f_{sweep} decreases the peak temperature. At 0.1 Hz, the strike point moves slowly from the OBD to the IBDH and is not oscillatory. Because this f_{sweep} is slow and the PFCs are inertially cooled, the PFC saturates with heat quickly and the sublimation temperature is reached soon thereafter. Sublimation begins at 0.92 s and a peak temperature of 4064 K is reached by the end of the discharge. Increasing f_{sweep} to 1 Hz results in a shorter time to sublimation because the strike point quickly moves onto the PFC of interest, IBDH, rather than lingering on the outboard tile (as in the 0.1 Hz case). For 1 Hz the PFC begins sublimating at 0.55 s, and reaches a peak temperature of

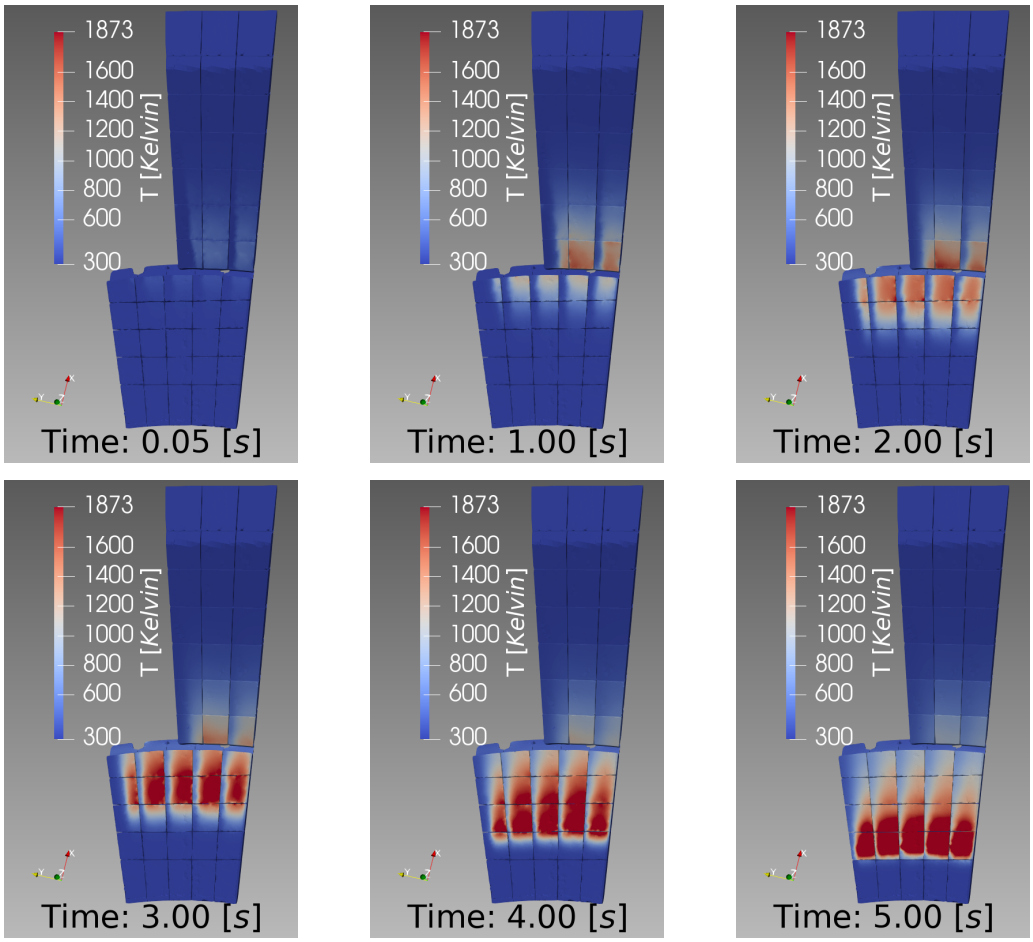


Figure 4.3: View of lower divertor (from above) of time evolving temperature profile for 5 second strike point sweep with sweep frequency of 0.1 Hz. Sweep starts on upper tile (OBD), moves with constant velocity, then ends on lower tile (IBDH). Areas in dark red are above the graphite sublimation temperature, 1873 K.

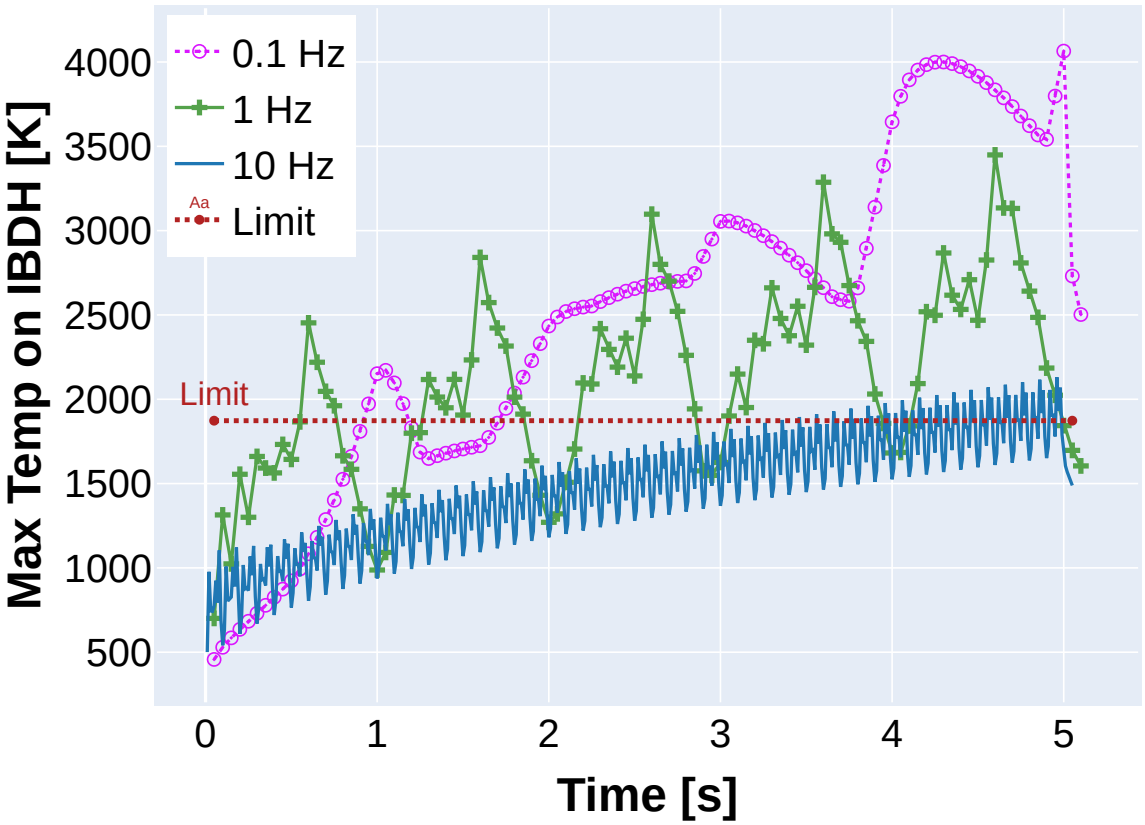


Figure 4.4: Peak temperature on IBDH tile for various strike point sweep frequencies. PFC is set to 300K at t=0, but t=0 is not plotted. Due to 5s discharge length, 0.1Hz sweep is in a single direction across tile surface, while 1.0Hz and 10.0Hz are oscillatory. Engineering limit temperature for sublimation (Limit) is reached at approximately 0.92s for 0.1Hz, 0.55s for 1.0Hz, and 3.36s for 10.0Hz.

3449 K. Sweeping the strike point at 10 Hz maintains the PFC below the sublimation temperature until 3.36 s, corresponding to a 6x improvement over the 1 Hz case. The peak temperature is also reduced significantly for the 10 Hz case, to 2131 K.

Both 0.1 Hz and 1 Hz have limited viability in real NSTX-U operation due to the short allowable discharge times. A viable magnetic configuration must sweep the strike point at a frequency sufficiently high to enable the PFC to inertially cool between oscillation periods. This constrains the operational domain to a minimum f_{sweep}, which must be balanced against the divertor coil voltage limitations. HEAT is equipped to iterate on magnetic equilibria, CAD, and f_{sweep}, to find a nominal configuration that mitigates the heat loads while simultaneously preserving the high power intention of the discharge. Mapping out the tokamak operational space directly as a function of the PFC engineering limits provides an interface between plasma physics and the PFC design.

4.3 NSTX-U Discharge 204118

To test HEAT's ability to do post-shot reconstruction, it has been applied to the 2016 NSTX-U discharge 204118. More details about this plasma, as well as early NSTX-U operations, can be found in [69]. This discharge begins as a center stack limited discharge, and then progresses to a lower single null diverted configuration. In this section, HEAT predictions for both configurations will be presented.

Figure 4.5 shows a center stack limited heat flux profile for discharge 204118, which occurs during the discharge ramp up. The magnetic equilibrium is taken directly from the NSTX-U MDS+ tree, and the PFC geometry has been taken directly from Recovery-era CAD designs. While HEAT can utilize empirical regressions for limited discharges, in this example the near and far SOL heat flux widths and ratios have been prescribed per the user's definitions: the near SOL has a decay width of 3 mm and carries 90% of P_{SOL}, whereas the far SOL has a decay width of 1 mm and carries 10% of P_{SOL}. The power crossing the separatrix, P_{SOL}, has been defined at 2 MW. Using a mesh resolution of 2 mm, the HEAT reconstructed power, P_{SUM}, is calculated to be 1.928 MW, corresponding to 3.6% error.

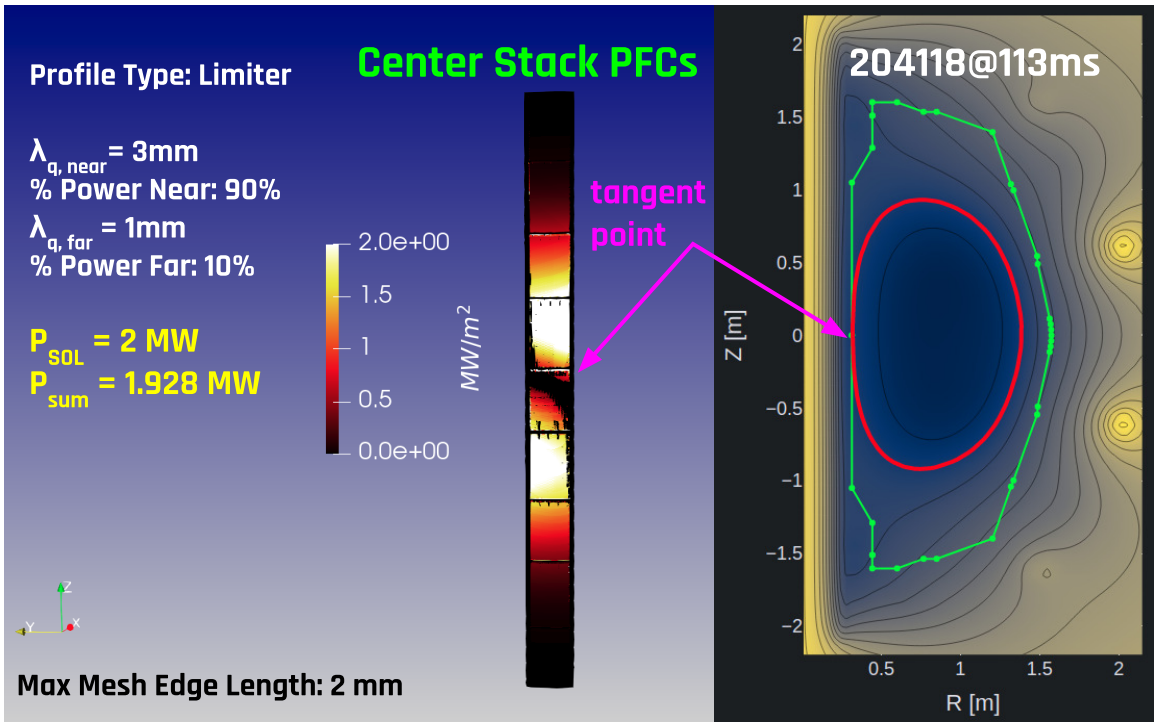


Figure 4.5: Limiter heat flux and magnetic equilibrium reconstruction during startup for NSTX-U discharge 204118 at 113 ms. User defined heat flux widths and power sharing are provided in figure. Magnetic shadowing creates toroidal asymmetry at the tangent point. Peak heat flux is $\approx 2.6\ \mathrm{MW\,m^{-2}}$.

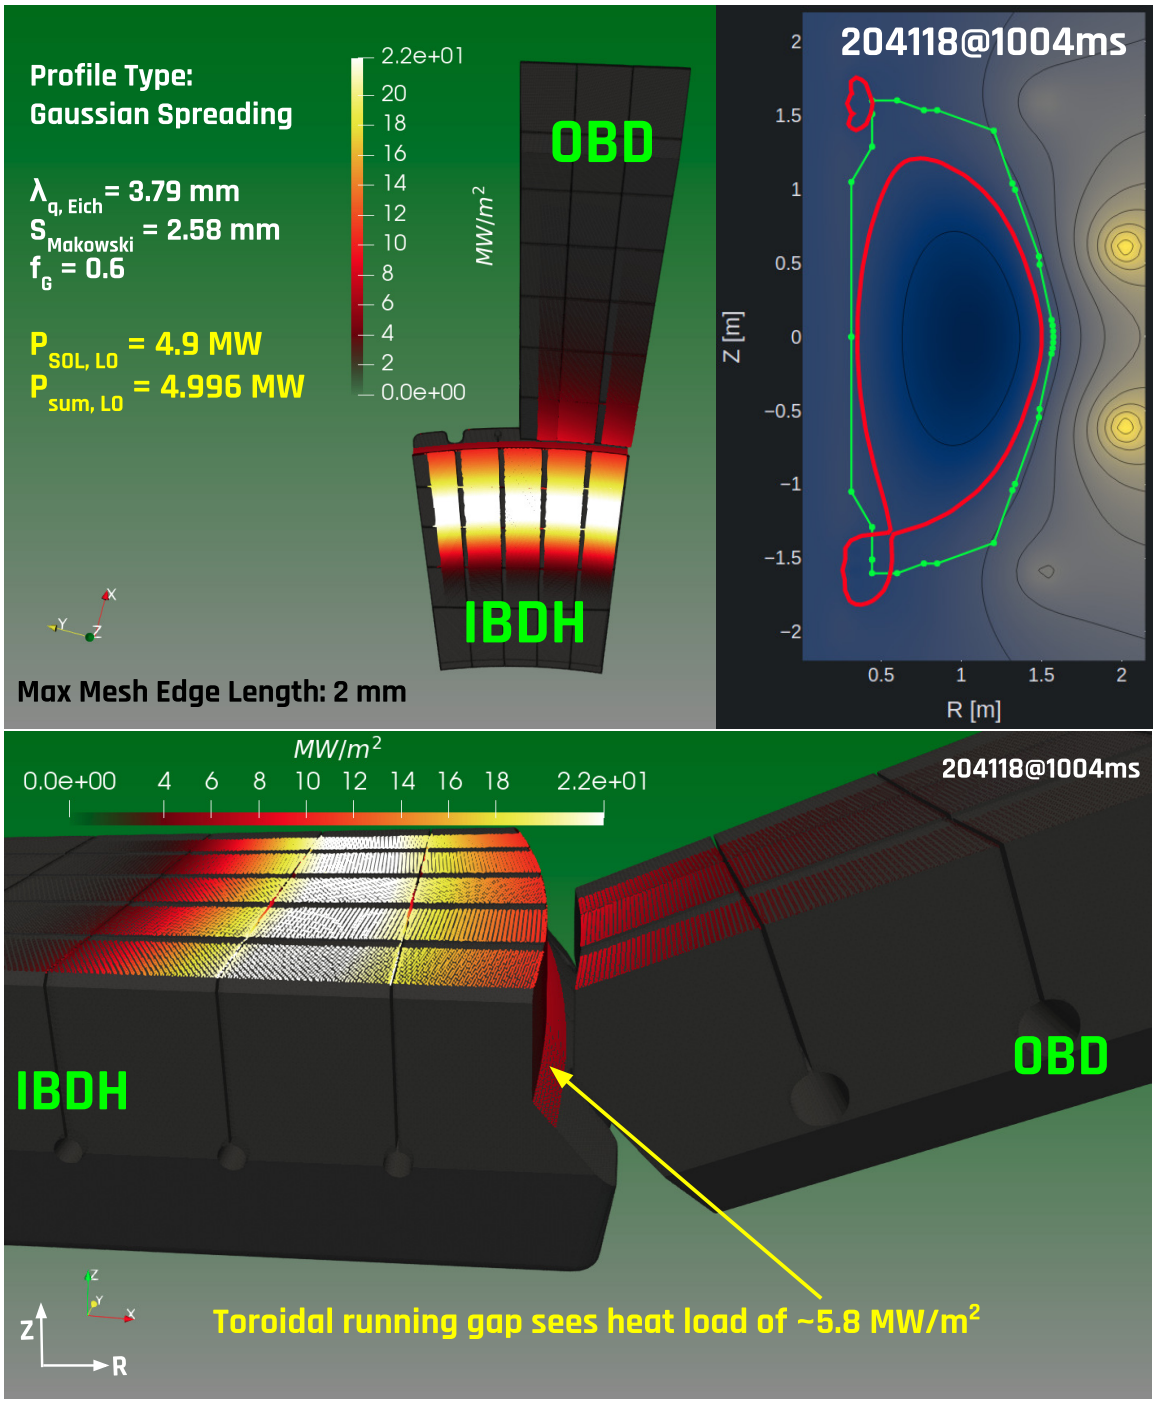


Figure 4.6: Lower outer divertor heat flux and magnetic equilibrium reconstruction for NSTX-U discharge 204118 at 1004 ms. User defined heat flux widths and power sharing are provided in figure. Peak heat flux is ≈ 24 MW m^{-2}. Toroidal running gap that was originally designed to be shadowed gets loaded with heat fluxes up to 5.8 MW m^{-2}.

The limiter point, where the magnetic field lines are perfectly tangential to the PFC surface, is called out in the figure. At this point the heat flux is analytically expected to be null, as the $(\hat{b} \cdot \hat{n})$ component of equation 3.26 returns 0. Experiment has indicated that there is a smearing effect at the tangent point, as shown in work by Stangeby [70], but this effect is neglected in figure 4.5 and will be addressed in future publications where gyro orbit heat flux and cross-field heat flux from filaments will be explored. The two off axis heat flux peaks occur as a result of competing effects due to $(\hat{b} \cdot \hat{n})$ and $\hat{q}_{||}(\psi)$ in equation 3.26. Power flows from opposite directions on either side of the limiter point, resulting in the sigmoidal magnetic shadow that appears oriented in the vertical direction.

After ramp up, discharge 204118 progresses into a lower single null diverted configuration. Figure 4.6 shows two views of the lower outer divertor heat flux, as well as the magnetic equilibrium reconstruction at 1004 ms. The heat flux peak at the strike point corresponds to nearly 24 MW m^{-2}. At these levels of heat flux, some plasma control or shaping would need to be employed to prevent sublimation if the discharge would be sustained for the 5 seconds that NSTX-U is designed to reach. The magnetic equilibrium recontruction indicates that poloidal flux surfaces are nearly perpendicular to the IBDH, enabling particles to descend into toroidal running gaps as illustrated in the lower frame of the figure. Heat loads of 5.8 MW m^{-2} can be observed in the toroidal running gap that separates the IBDH from the OBD. In order to reduce this heat flux, the strike point could be moved or the magnetic field could be adjusted to include a larger radial component at the gap.

4.4 Using HEAT During NSTX-U Construction

During the initial design phases of the NSTX-U Recovery Project, HEAT was not yet developed. In order to ensure that the PFCs were maintained within their allowable engineering limits, conservative requirements were created that incorporated a factor of safety. After the development of HEAT, many of the initial scoping scenarios were revisited to determine the validity of using these conservative requirements. The previous sections in this chapter outlined several of these HEAT analyses. In addition to revisiting the original scenario studies, HEAT was used to validate the engineering CAD

design. This section provides two examples of how HEAT predictions were utilized during the NSTX-U construction process to alter the PFC design in real time.

4.4.1 OutBoard Divertor Missing Fish-scale

During the analysis provided in the preceding sections, HEAT generated heat flux footprints for most of the heat loaded tiles in the NSTX-U divertor. The castellated graphite design (see section 2.5) was employed to protect the PFC leading edges from high angle of incidence ($\hat{b}\cdot\hat{n}$ from equation 3.26) heat fluxes. During the HEAT analysis for some of the PFC design scoping scenarios, a missing fish-scale was detected on the OutBoard Divertor (OBD) tile. Figure 4.7 provides a visualization of the OBD tile with Case 1 (left panels) corresponding to the intended PFC design and Case 2 (right panels) correponding to the unintended CAD design that was sent to the manufacturer for production.

In figure 4.7, Case 1 highlights the fish-scale step between the left and center castellations (top left frame). This fish-scale step is approximately 0.5 mm in height, and produces an identical shadowed region on each castellation (bottom left frame). For Case 1, the leading edge of the center castellation is protected from the incoming heat fluxes as can be observed by the fact that it resides in the shadowed region of the PFC. Case 1 represents the intended engineering design, where the leading edges are protected and the shadows appear on each castellation column.

Case 2 in figure 4.7 represents the actual engineering CAD drawing that was provided by the NSTX-U CAD engineering team. The fish-scale is highlighted and the step size is 0.0762 mm (top right frame). This small fish-scale step size produces almost no shadow in the center castellation column, as can be observed in the lower right frame of the figure. This missing shadow results in the leading edge on the center castellation column receiving a non-negligible amount of heat flux. Additionally, for magnetic scenarios with steeper angles of incidence, there would likely be a very hot leading edge. As it turns out, the small fish-scale step size in Case 2 was an unintentional drafting error; the design was intended to have an identical fish-scale step size between each castellation column, as shown in Case 1. This design error went undetected until HEAT analysis was performed on the PFCs and the 'incorrect' heat

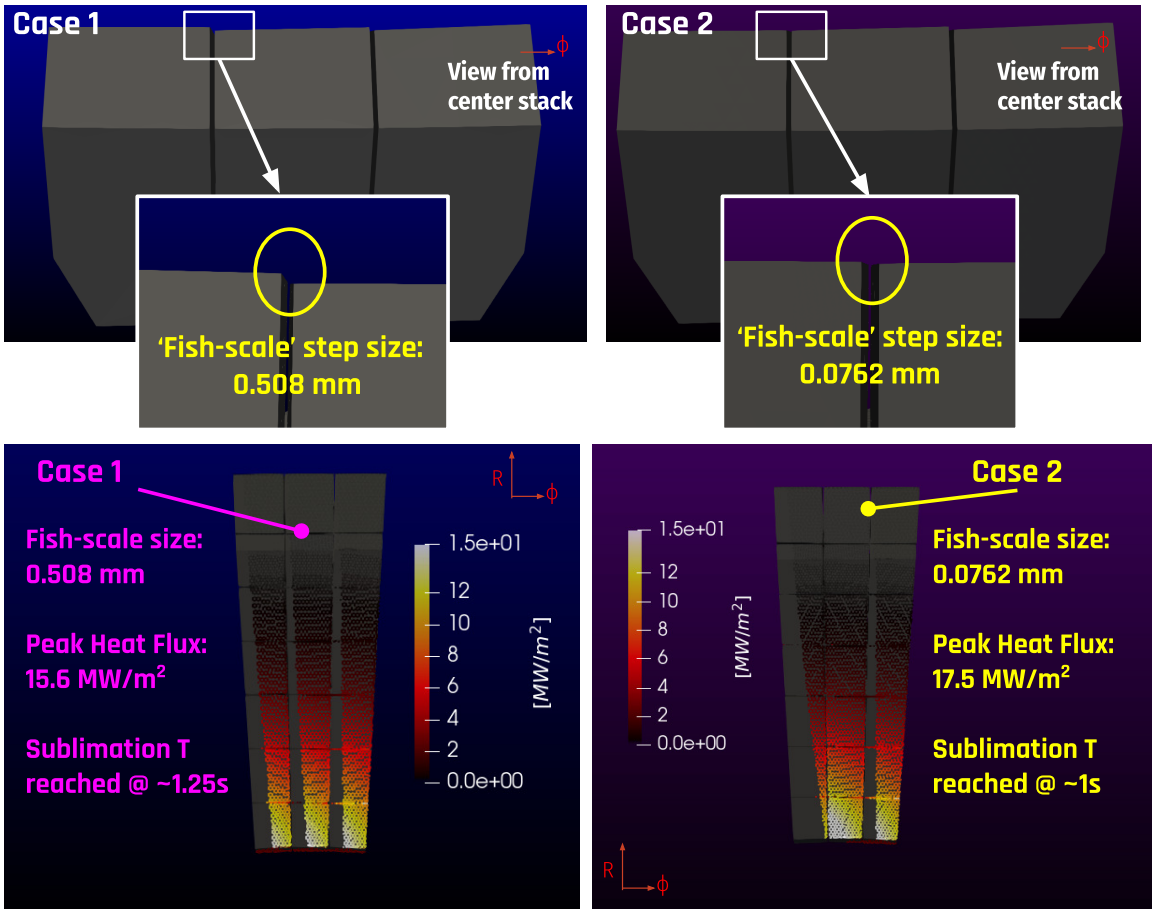


Figure 4.7: Comparison geometry and associated heat fluxes for 2 Cases. Case 1 (left frames) represents the intended CAD design. Case 2 (right frames) represents the unintentional CAD design that was originally sent to PFC manufacturer. Top frames illustrate the geometry difference - different step sizes for the intercastellation fish-scale between the left and center castellation columns. Bottom frames illustrate the heat fluxes associated with each fish-scale step size. Increasing the step size by 0.4318 mm decreases the peak heat flux and increases the discharge duration by 25%. HEAT was used to find this unintentional CAD design error. CAD was revised and updated with the manufacturer before PFCs were fabricated.

flux footprint was observed in the output. Luckily, this design error was found before the PFC was manufactured, so the engineering team was able to update the design and send the correct drawings to the manufacturer.

This difference between Case 1 and Case 2 from figure 4.7 provides an interesting opportunity to quantify the sensitivity of heat fluxes and temperatures to a single geometric feature, the inter-castellation fish-scale height. The lower frames in figure 4.7 provide some values that correspond to this analysis. The heat flux footprint in Case 1 generates a peak heat flux of approximately 15.6 $\mathrm{MW\,m^{-2}}$, while Case 2 generates a heat flux of approximately 17.5 $\mathrm{MW\,m^{-2}}$. This difference in heat fluxes can be attributed to the heat flux faceting effect. The lower left corner of the center castellation column in figure 4.7 is where the heat flux faceting effect occurs. In Case 1, this corner is protected by the fish-scale shadow provided by the upstream castellation. In Case 2, however, the missing shadow enables heat flux to reach the corner and the increasing angle of incidence in the shadow generates a heat flux maximum. If the temperatures that arise from the heat fluxes in both cases are simulated using the HEAT thermal analysis tools, the engineering sublimation limit is reached at 1.25 s for Case 1 and at 1.0 s for Case 2. It can therefore be concluded that increasing the fish-scale step size from 0.0762 mm to 0.508 mm has the potential to extend the tokamak steady state discharge length by 25% for this PFC. The mass difference between the OBD PFCs in Case 1 and Case 2 is approximately 5 grams of carbon. That such minute changes to the PFC geometry can have such a profound impact on the operational boundary (i.e. discharge length) of the tokamak validates the need for simulation tools such as HEAT.

4.4.2 OutBoard Divertor Installation Misalignment

After the PFC fabrication process was completed, the divertor tiles were installed in the NSTX-U divertor. According to the NSTX-U engineering department, there was a manufacturing error associated with the vacuum vessel. This manufacturing error had a cascade effect on other components in NSTX-U, one of these components being the OutBoard Divertor tiles. More specifically, the OBD tiles were installed 0.5 inches closer to the plasma than in the intended design. Figure 4.8 provides

images that highlight the differences between the intended design and the as built PFCs in the left two frames. The OBD tile edge is elevated from the InBoard Divertor Horizontal (IBDH) edge, resulting in a surface that is exposed directly to the plasma. The right frame of the figure provides an illustration of a NSTX-U magnetic equilibrium that could be potentially hazardous with this misalignment error. After this misalignment error was discovered, HEAT was used to determine the negative impact of this alignment error for the aforementioned magnetic equilibrium.

The results of the HEAT predictions for this misaligned OBD PFC are provided in figure 4.9. The top frame provides a wide angle view, and the bottom frame provides a closeup view that shows the heat loads on the leading edges. The OBD surface of interest is called out in the figure. The misalignment exposes a larger area of this vertical surface to heat fluxes, but this surface is curved to match the magnetic field lines so the heat fluxes remain in the 5-10 $\mathrm{MW\,m^{-2}}$ range. The misalignment elevated the PFC and subjected a larger surface area to this heat flux, but the magnitudes and spatially integrated power are relatively unperturbed due to the PFC curvature.

The leading edge, however, presents a potential problem for NSTX-U. On the leading edge, a 246 $\mathrm{MW\,m^{-2}}$ heat load can be observed in the lower frame in figure 4.9. The leading edge only barely shows up for the HEAT predictions at 500 um maximum mesh edge length, and further analysis quantified the leaded edge to be approximately 100 um wide. Despite the extreme peak heat flux, this narrow leading edge will likely not push the PFCs past their engineering limits as the spatially integrated deposited energy is low. For pulsed discharges under 10s in length, the thermal mass of the PFC is sufficiently large to absorb the heat incident upon this leading edge. It should also be noted that the heat flux on this leading edge would have been present under the intended design as well, but the vertical height of the leading edge would have been reduced. These HEAT predictions enabled the NSTX-U construction to continue, despite this misalignment error. During operations, extra care will be taken to ensure that equilibria like the one in figure 4.8 are avoided.

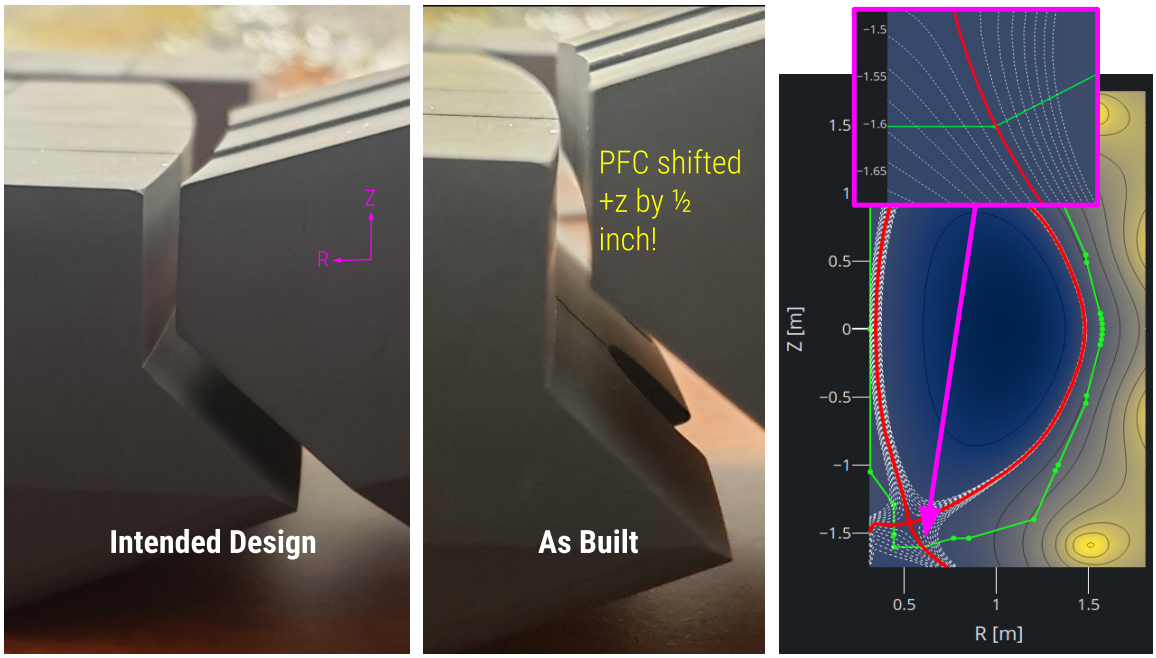


Figure 4.8: Images demonstrating the NSTX-U OutBoard Divertor (OBD) intended design (left frame) and as built installation (center frame). In the as built installation, the OBD has been shifted vertically towards the plasma by approximately 0.5 inches, exposing the vertical surface to incident heat fluxes. Right frame shows a potential hazardous magnetic equilibrium for this misalignment, where the strike point is directly on the vertical face.

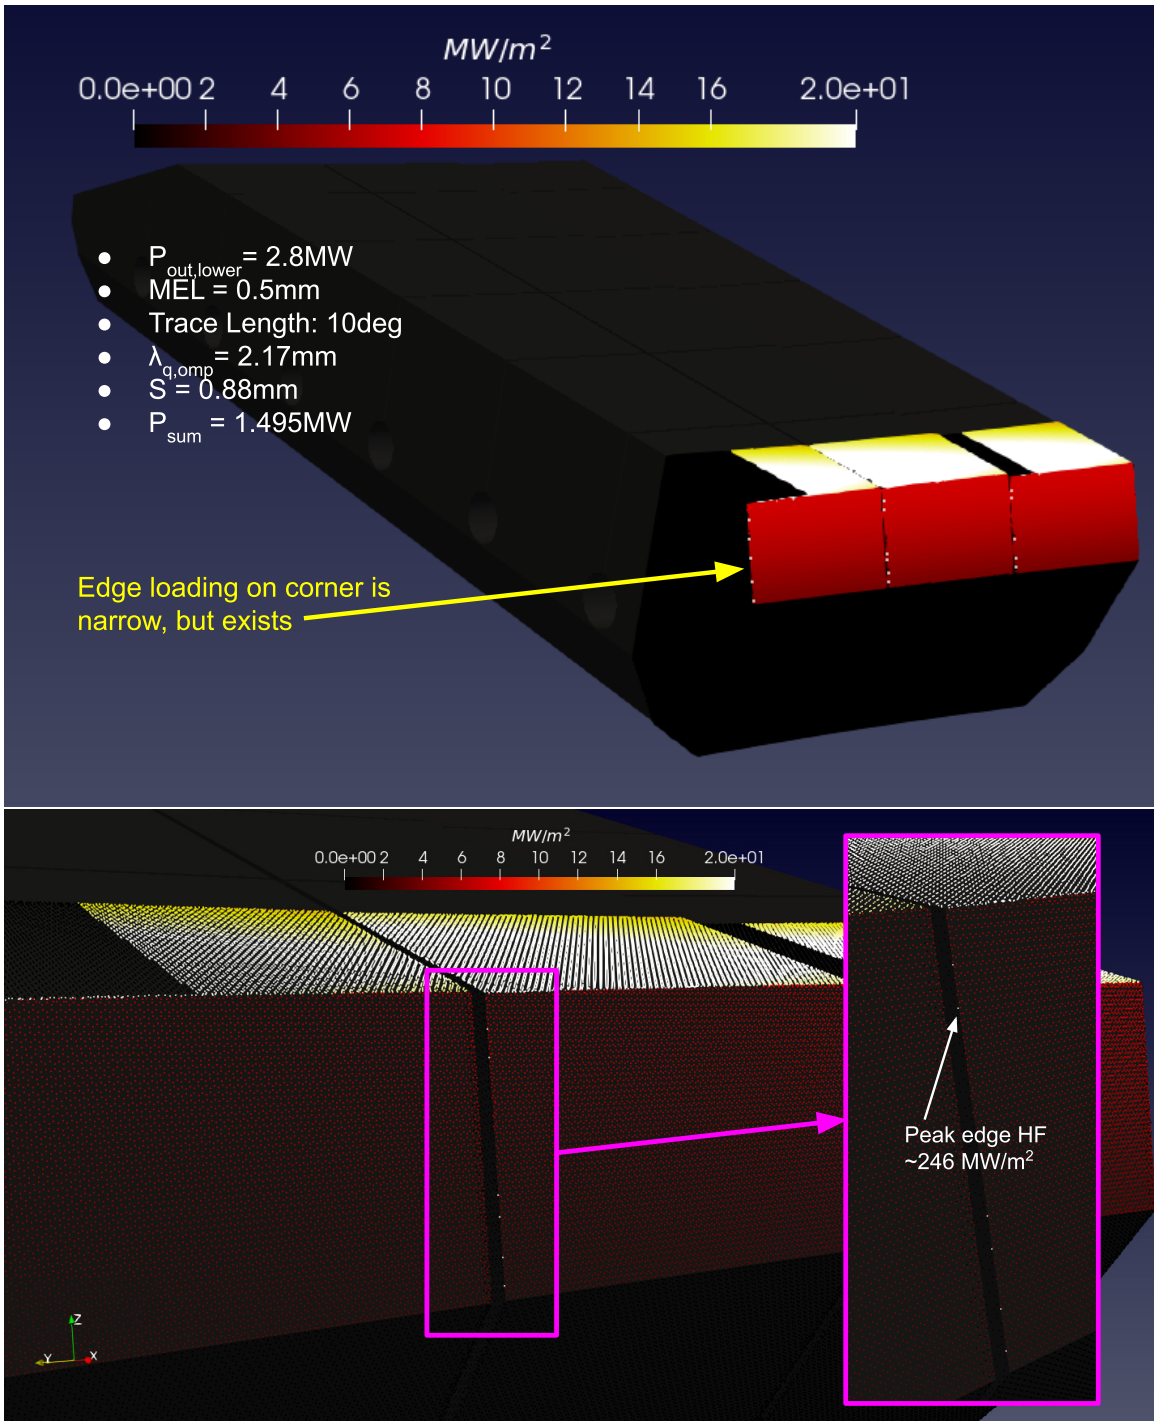


Figure 4.9: HEAT heat flux predictions for the misaligned OBD geometry as described in figure 4.8. The vertical edge that is no longer shadowed due to the misalignment sees heat fluxes in the 5-10 MW m^{-2} range (top frame). The leading edge sees heat fluxes of nearly 250 MW m^{-2}, although these have a width of only 100 um. The leading edges were also present in the intended PFC design, but the vertical height has been increased due to the misalignment.

Chapter 5

HEAT Gyro Orbit Approximation Results for NSTX-U

This chapter provides some HEAT results for gyro orbit approximation heat fluxes. Because this is the most technically advanced module in HEAT, an entire chapter has been devoted to the discussion. First, a simple geometry is chosen. Next, this simple geometry is compared to previous research performed on the ITER tungsten monoblocks [24, 18]. The simple geometry is then used in a series of convergence studies. Finally, the gyro orbit HEAT module is applied to the NSTX-U castellated graphite PFCs, providing the first analysis of its kind for NSTX-U.

5.1 Comparison to Existing Research

Section 1.3.2 outlined some existing research that has been conducted with respect to ion gyro orbit heat loads. This section benchmarks HEAT gyro-orbit output against work done in support of ITER. More specifically, the inter-ELM HEAT gyro orbit heat flux predictions are compared to published research by the ITER team.

Gunn has published several examples of using the 'ion gyro-orbit' approximation to predict heat loads for the ITER tungsten monoblocks [24, 17]. Figure 5.2 shows an ITER inter-ELM heat load incident upon a tungsten monoblock at the ITER inner and outer vertical targets. The heat flux profile is plotted as a function of the poloidal chord shown in figure 5.1, which begins at the corner. In both the inner vertical target (IVT)

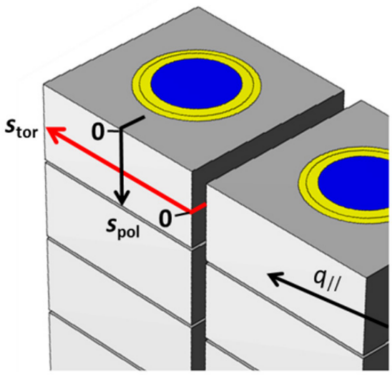


Figure 5.1: Example poloidal (black) and toroidal (red) chords taken across an ITER tungsten monoblock from [17].

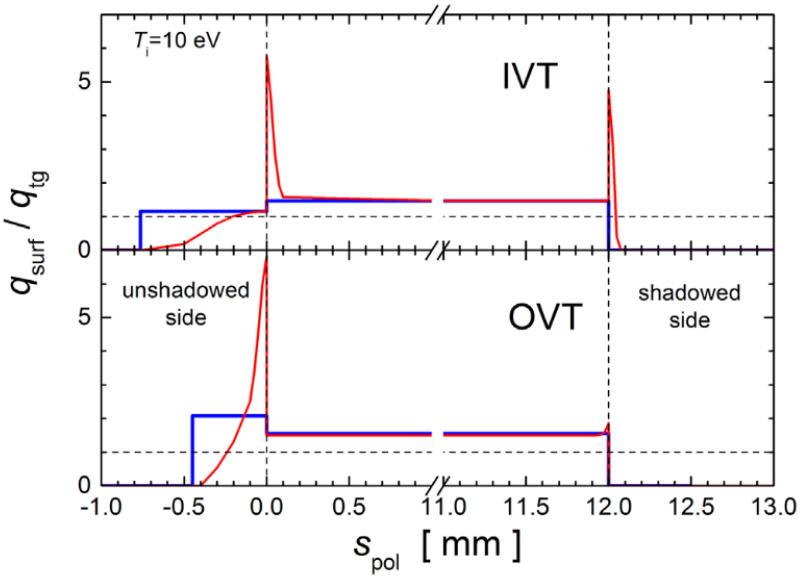


Figure 5.2: Inter-ELM ion heat flux for optical (blue) and gyro-orbit model (red) from Gunn [17].

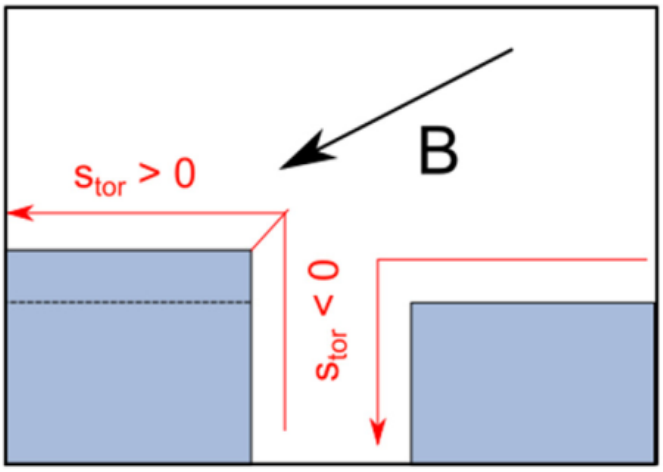


Figure 5.3: Example geometry from PIC simulations performed by Komm [18]. Geometry corresponds to an ITER monoblock that has been misaligned to expose a leading edge.

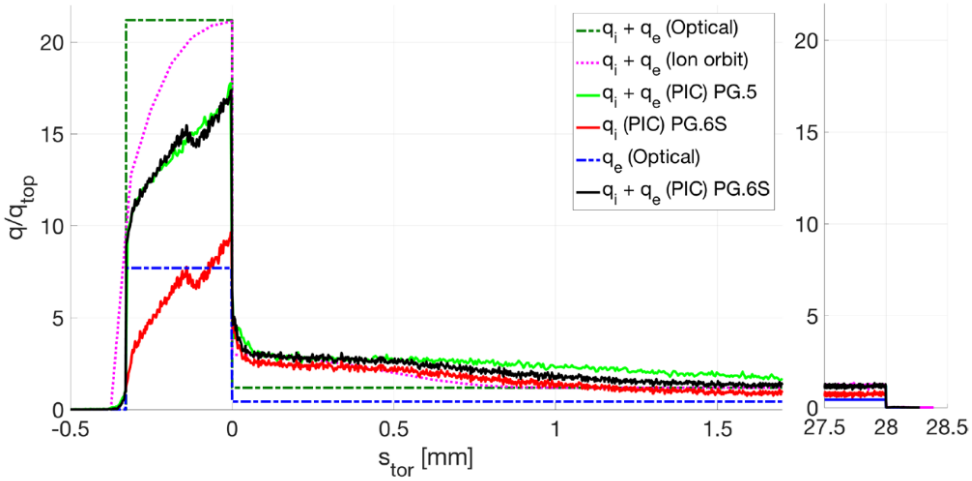


Figure 5.4: Comparison of heat flux profiles for optical, ion orbit (from [24]), and SPICE2 PIC code from [18]. Simulations performed on portruding leading edge geometry, where $s_{tor} < 0$ corresponds to the leading edge and $s_{tor} > 0$ corresponds to the top surface. Including ion orbits transfers power from leading edge to PFC surface. Including electric fields enhances this effect. In both cases, power redistribution area is on the order of millimeters wide.

and the outer vertical target (IVT), optical approximation power is 'redistributed' around the PFC edges by the helical trajectories of the ions, but the integrated power is approximately the same. The peak heat flux is greatly increased in the gyro-orbit approximation, as power that is transported around the PFC corner loads a narrow region.

The Gunn calculations ignore electric field effects, and as a means of validating this assumption a series of self consistent particle in cell (PIC) analyses that included electric fields were performed by Komm [18]. Figure 5.3 provides a visual depiction of the geometry, where a poloidally running leading edge has been exposed to the incident heat flux. Assuming a steady state inter-ELM scenario for this geometry, figure 5.4 plots the heat flux along a toroidally running chord, and compares the optical approximation, the Gunn ion orbit approximation, and the PIC code with electric fields. The optical approximation is observed to have the highest heat load, as the leading edge is loaded with a constant heat flux. The ion-orbit and PIC code results provide lesser heat loads, as some of the optical power is 'redistributed' around the corner of the leading edge onto the top PFC surface.

As a means of validating the HEAT gyro-orbit approximation and associated algorithms against the aforementioned results for ITER, a test case was developed in HEAT. Figure 5.5 illustrates the geometry utilized for these test runs. The PFC of interest extends 5 mm in the poloidal direction and 10 mm in the toroidal direction. These dimensions were chosen because they are sufficiently large to quantify the gyro orbit power redistribution, but small enough to run quickly on a desktop workstation. In the toroidal direction, the PFC of interest portrudes from the upstream PFC by 0.3 mm, which matches the step height in Komm [18]. The adjacent poloidal PFCs are spaced 0.5 mm apart, yet their surfaces are perfectly aligned with the PFC of interest as in Gunn's analyses [17]. Because this test case matches the geometry toroidally in Komm's PIC analysis and poloidally in Gunn's ion-orbit analysis, a comparison to both can be performed from a single set of HEAT results. The divertor plasma temperature for the HEAT runs was set to 10eV, which produces slightly smaller gyro-radii than Gunn (10eV) and Komm (approx. 5eV) due to the magnetic field difference, although the HEAT algorithms sample uniformly in velocity phase space, which will result in a variety of gyro-radii.

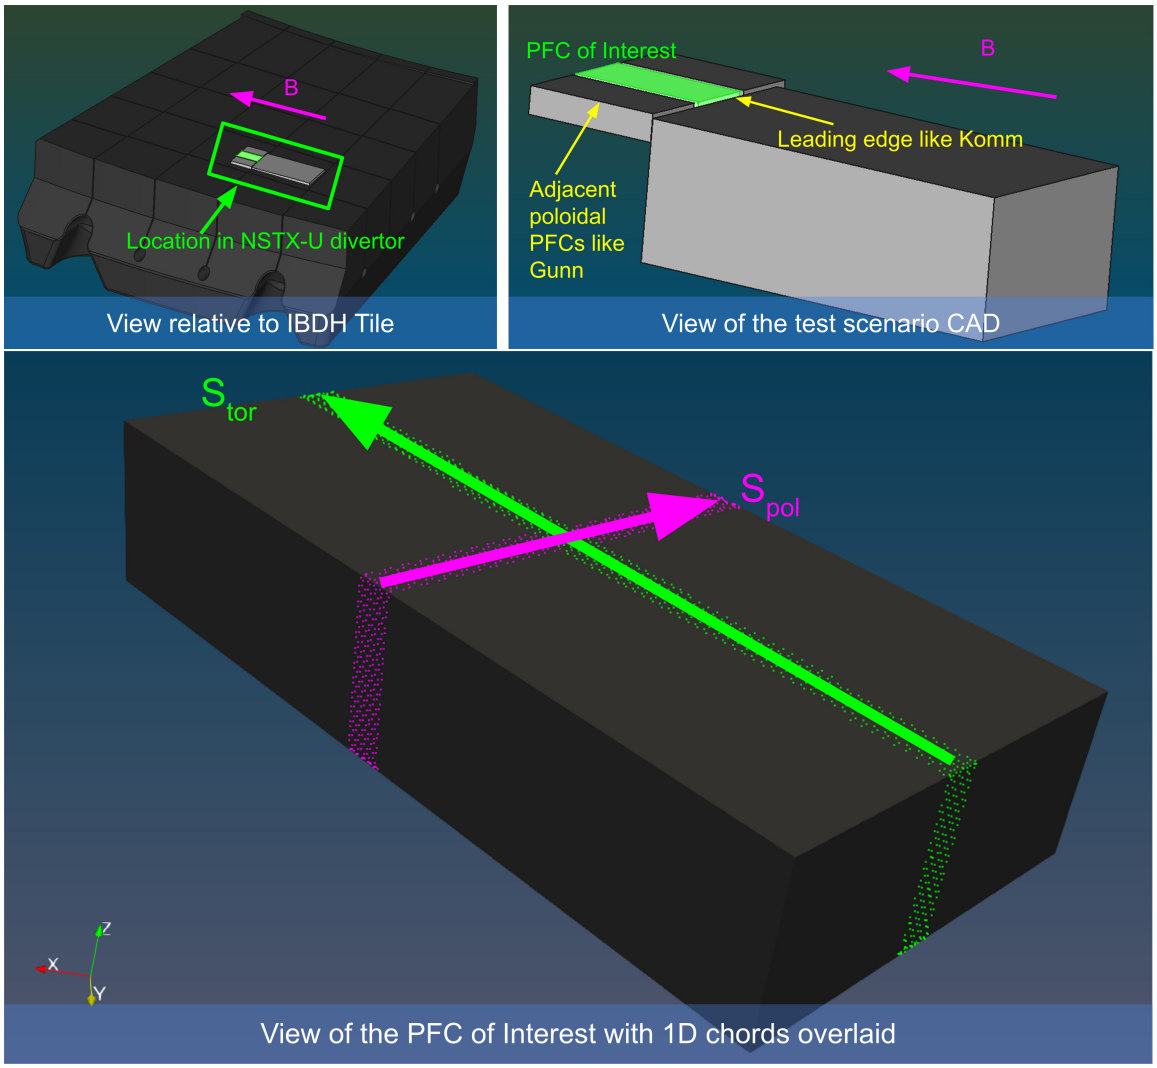


Figure 5.5: Top image: CAD representation of test case for comparison to ITER monoblock simulations by Gunn and Komm. PFC of interest is located in the lower NSTX-U divertor, aligned to a castellation in the IBDH tile. The PFC of interest portrudes from the toroidally upstream PFC by 0.3 mm, and is perfectly aligned to adjacent poloidal PFCs with 0.5 mm gaps. PFC of interest extends 10 mm toroidally and 5 mm poloidally. Lower image: HEAT calculates heat fluxes across entire surface, and then mesh heat flux values within 0.25 mm threshold from PFC centerlines are projected to the S_{pol} and S_{tor} planes for plotting. Mesh values to be included for toroidal chord is shown in green, poloidal chord in magenta.

Because HEAT predicts the heat loads over the entire 3D PFC surface, the 3D heat fluxes must be collapsed to a 1D profile in order to compare with the results shown in figures 5.2 and 5.4. To this end, an algorithm for generating heat flux profiles along a 'chord' of the PFCs has been developed in HEAT. First, a plane is constructed that will define the chord. For these comparisons the planes are taken to be the centerlines of the PFCs in the longitudinal and transverse directions, as shown in figure 5.5. Next, all points that are further than a user specified threshold (here +/- 0.25 mm) are ommited. The remaining points are then projected orthogonally down to the plane. Figure 5.5 shows the points that are included in the heat flux profile for the toroidal (green) and poloidal (magenta) directions. The 1D profile (chord) is constructed by defining the PFC corner, as illustrated in 5.1, and then plotting the heat fluxes as one walks the PFC surface away from the corner. Due to the non-uniform mesh, there is some noise in this method, but it is generally sufficient to compare against the aforementioned results for ITER.

Figure 5.6 shows the 1D heat flux profiles as a function of the poloidal and toroidal chords. The poloidal plot (top) can be directly compared to figure 5.2. The poloidal plot shows an optical leading edge with heat fluxes at grazing angles of incident for $S_{pol} < 0$ and with a uniform heat flux across the PFC surface. The Ion Gyro profile (blue) in the figure peaks on both corners of the PFC, where the helical trajectories of ions enables them to wrap around the corners. On the right corner, ions that descend into the toroidally running gap between adjacent PFCs are scraped off by the leading edge, resulting in the left peak. On the left corner, ions whose guiding center trajectories descend to just above the PFC surface wrap around the corner and strike the shadowed edge of the PFC, resulting in the right peak. In both cases, the interaction between the gap width between adjacent PFCs and the gyro radius size is what defines the magnitude of the power that strikes the PFC edges. These results match the Gunn results shown in figure 5.2, although the plasma scenarios are different (ITER vs. NSTX-U).

The toroidal plot (bottom) in figure 5.6 can be directly compared to figure 5.4. In this plot, the PFC of interest is elevated 0.3 mm above the toroidally upstream PFC, resulting in a leading edge that is exposed at near normal angle of incidence to the magnetic field lines. This results in an extreme optical heat flux (green) on the

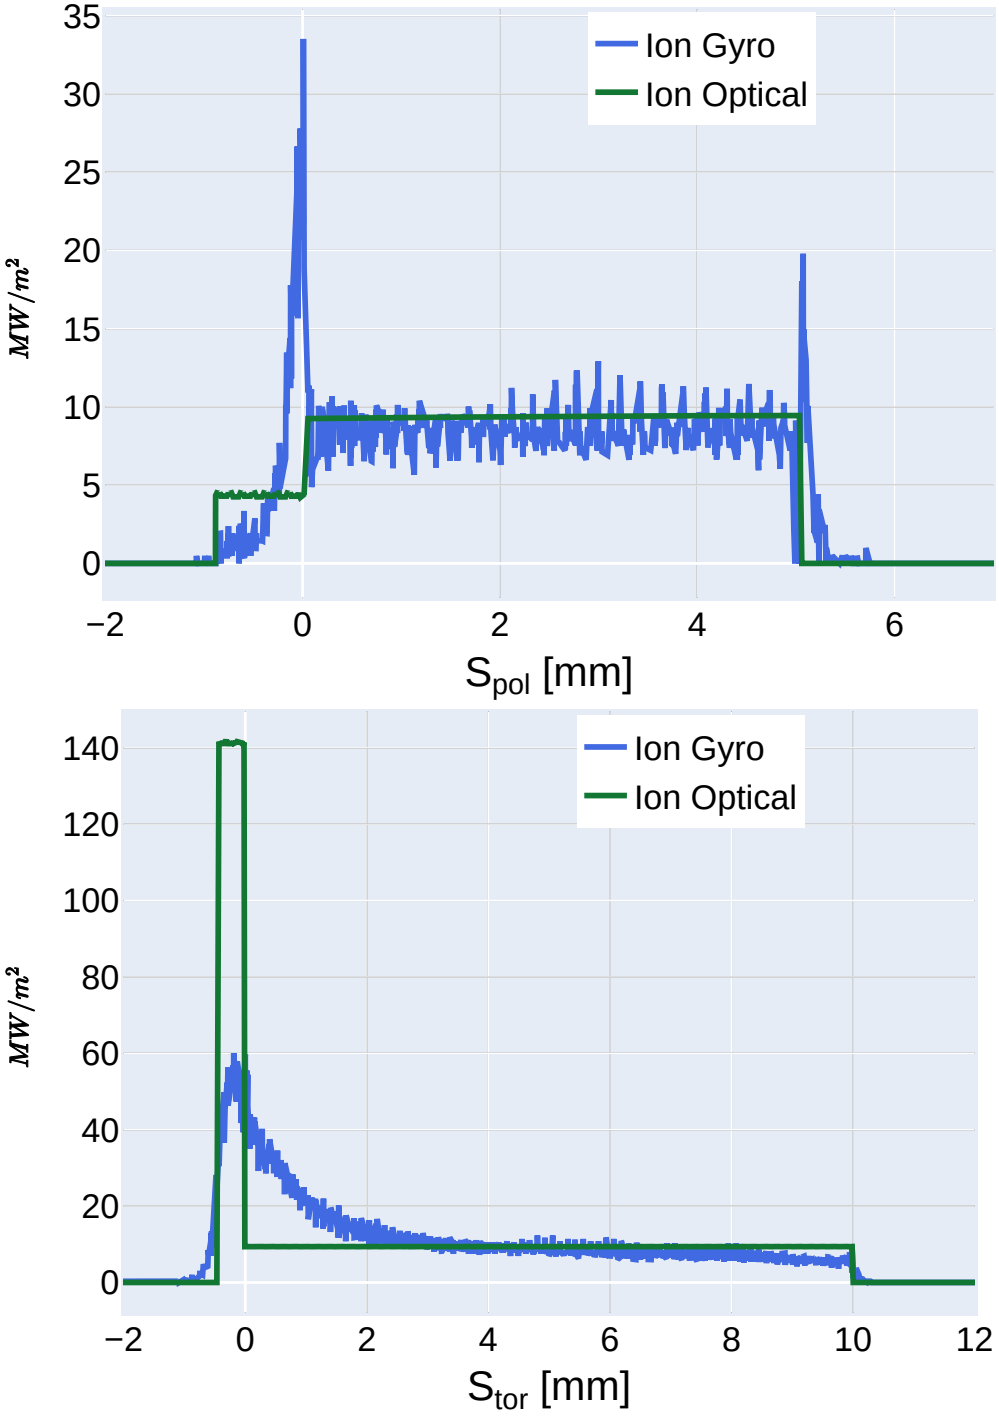


Figure 5.6: Example HEAT generated poloidal (top) and toroidal (bottom) heat flux predictions plotted along PFC centerlines. Blue trace is ion gyro-orbit approximation heat flux, green trace is optical approximation. Both the toroidal and poloidal profiles show good agreement to the results from Gunn [17] and Komm [18]. For poloidal case, integral of gyro orbit curve is 0.0494 MW m and integral of optical curve is 0.0510 MW m. For toroidal case, integral of gyro orbit curve is 0.1399 MW m and integral of optical curve is 0.1572 MW m. Mesh elements within 0.5mm of chord are plotted.

leading edge ($S_{tor} < 0$) that exceeds 200 MW m^{-2}. On the top surface the optical heat flux is uniform. On the exposed leading edge the Ion Gyro-orbit heat flux (blue) is much smaller than the optical heat flux due to the helical pitch of the ion trajectories. The helical motion enables ions with certain gyro phase angles to gyrate up onto the top surface of the PFC, even when their guiding center strikes the leading edge. As S_{tor} increases along the PFC surface, the gyro-orbit heat load decays as particles that have 'gyrated up' are scraped off. The red trace in figure 5.4, which includes electric field effects, can be compared to the blue trace in figure 5.6 which is the HEAT gyro-orbit approximation. The results for the toroidal chord profiles for the two results qualitatively agree, despite the differences in scenarios (ITER vs. NSTX-U).

These results demonstrate that the 3D heat load predictions generated by HEAT can be collapsed to 1D to reproduce results similar to pre-existing research. This excercise is useful to benchmark the validity of the HEAT results, but does not leverage the full capabilities of the HEAT ion gyro-orbit module. Predicting the time varying ion gyro-orbit 3D heat flux profile incident upon fully-featured engineering CAD geometry is a novel analysis that has not yet been performed in any existing research until now using HEAT. To demonstrate this capability the 3D heat flux on the PFC of interest used to generate figure 5.6 is shown in figure 5.7. The figure provides four views, where the top row is the gyro orbit heat flux, the bottom row is the optical heat flux, the left column is the outboard side of the PFC, and the right column is the inboard side of the PFC. The colorbar has been capped at 20 MW m^{-2} to show the gradients on the PFC surface and sides. The heat flux peaks on the corners from figure 5.6 can be seen in white on both the inboard and outboard corners of the PFC. Additionally, the large band of white on the toroidally upstream side of the PFC can be attributed to ions with gyro phase angles that enable them to gyrate up onto the top surface of the PFC. This is effectively a power redistribution from the leading edge to the PFC surface. As described in the previous discussion, the peak heat load occurs on the poloidally running leading edge that is exposed to the magnetic field line at near normal angles of incidence, and the optical approximation predicts higher peak heat loads at this location. Subsequent sections in this chapter will apply this type of analysis to the complicated divertor geometries of NSTX-U.

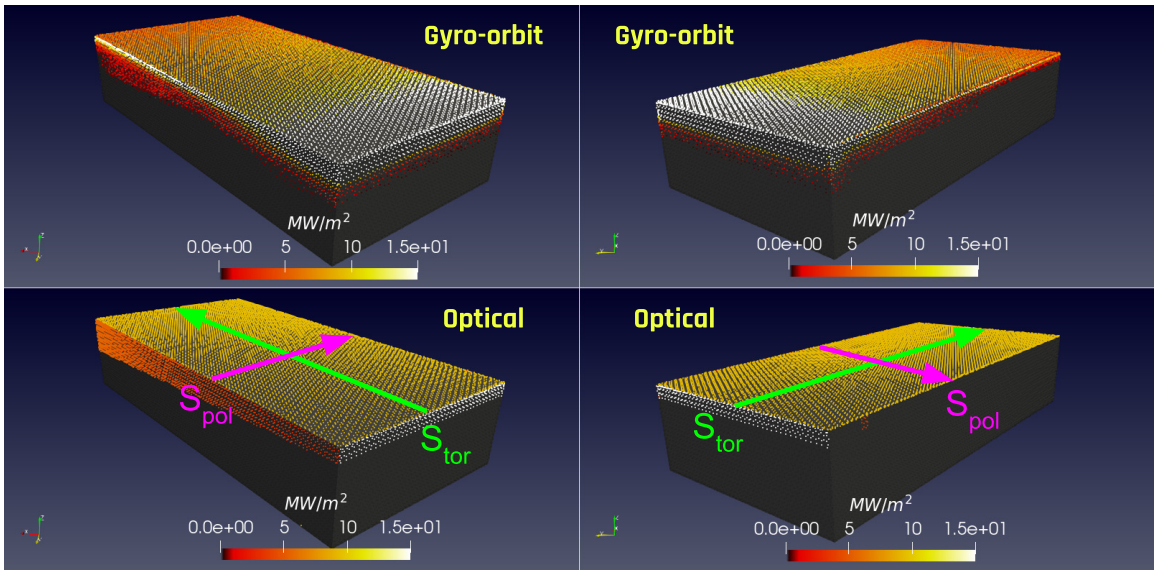


Figure 5.7: Gyro orbit (top row) and optical (bottom row) heat fluxes calculated using HEAT to generate figure 5.6. Left column shows outboard side of PFC, right shows inboard side. The colorbar has been capped at 20 $\mathrm{MW\,m^{-2}}$ to show the gradients on the PFC surface and sides. Heat flux is determined using the algorithms described in section 3.9.

5.2 Power Balance for Gyro Orbits

Section 3.8.2 described how energy balance is utilized in HEAT to check the results from optical approximation calculations. This approach allows the HEAT user to determine the allowable accuracy of the simulation, and tune the input variables to achieve that accuracy. The gyro orbit module also uses energy balance to check the heat flux output. The integrated power across an entire PFC surface will only yield power balance if the gyro-orbit trajectory shadows have correctly been calculated. Additionally, for a magnetic field that does not vary along a PFC, the optical and gyro-orbit power should be related by the fractional power carried by electrons and ions, respectively. Yet while the integrated power across a PFC surface may be identical between the optical approximation and the gyro orbit approximation, the heat flux profiles can be very different. The geometry determines the difference between the optical heat flux profile and the gyro-orbit heat flux profile. Geometries that employ many edges, toroidal shaping, and small features, tend to have a stronger difference between the optical and gyro-orbit profiles than smooth axisymmetric surfaces.

5.3 Convergence Studies

As a means of validating the computational integrity of the gyro-orbit model in HEAT, various convergence studies were performed. This section outlines some, but not all, of these convergence studies that served to validate the HEAT gyro-orbit predictions. Each of the studies performed demonstrates convergence, and taken together provide confidence in the HEAT output. As discussed in chapter 3, HEAT uses macro-particles to simulate gyro-orbit heat loading. Macro-particles are used to represent the behaviour of a chunk of particles from the velocity and spatial domains without having to simulate every single particle. The number of macro-particle simulations is a user input in HEAT, and this section describes convergence studies performed to determine how many macro-particle simulations are necessary to converge to a 'ground truth' solution. This analysis was performed by incrementally increasing the number of macro-particles while observing the change to the output heat flux profile.

5.3.1 Resolution Convergence Studies

The distribution of heat fluxes across the PFC surface should converge to a 'ground truth' solution as the number of macro-particles from the distribution functions (see section 3.9.2) is increased. This was confirmed by running HEAT five times, each with N samples from each of the three distributions (N,N,N) - gyro phase angle, velocity phase angle, and maxwellian speed - where N is an integer 1 through 5. This corresponds to N^3 particles per run (per mesh element).

The results of this study are plotted in figure 5.8, which displays a rug plot (top) and histogram plot (bottom). The (5,5,5) run was taken to be the ground truth, and runs (1,1,1) through (4,4,4) are plotted as different traces in the figures. Each datapoint, i, corresponds to the difference between the heat flux assigned to mesh element i in the (N,N,N) run compared to the (5,5,5) ground truth run. As expected, the results converge as the number of macro-particles is increased. More specifically, increasing the number of macro-particles reduces the data spread (interquartile range) and skew. Other studies were performed that increased the number of samples up to 12 (i.e. 12gP12vP12vS) and the convergence increases for these cases as well, although they are not shown here.

One particular measure of convergence is the interquartile range (IQR). IQR represents the statistical dispersion of the dataset, and is defined as the difference between the 75th and 25th percentiles. Figure 5.9 provides a plot of the absolute difference between the IQR values for each of the datasets from figure 5.8 and the IQR data for the ground truth (5,5,5) case. Convergence in this IQR difference corresponds to convergence in the amount of dispersion for each dataset. Figure 5.9 shows that the IQR difference converges, which indicates that increasing the number of HEAT macro-particles converges to the ground truth solution (N=5 is ground truth and IQR difference = 0). Because there can be multiple datasets with the same IQR, this does not represent unique convergence. However when taken with the analysis from figure 5.8 and the next section on power balance convergence, the family of solutions is constrained to be unique.

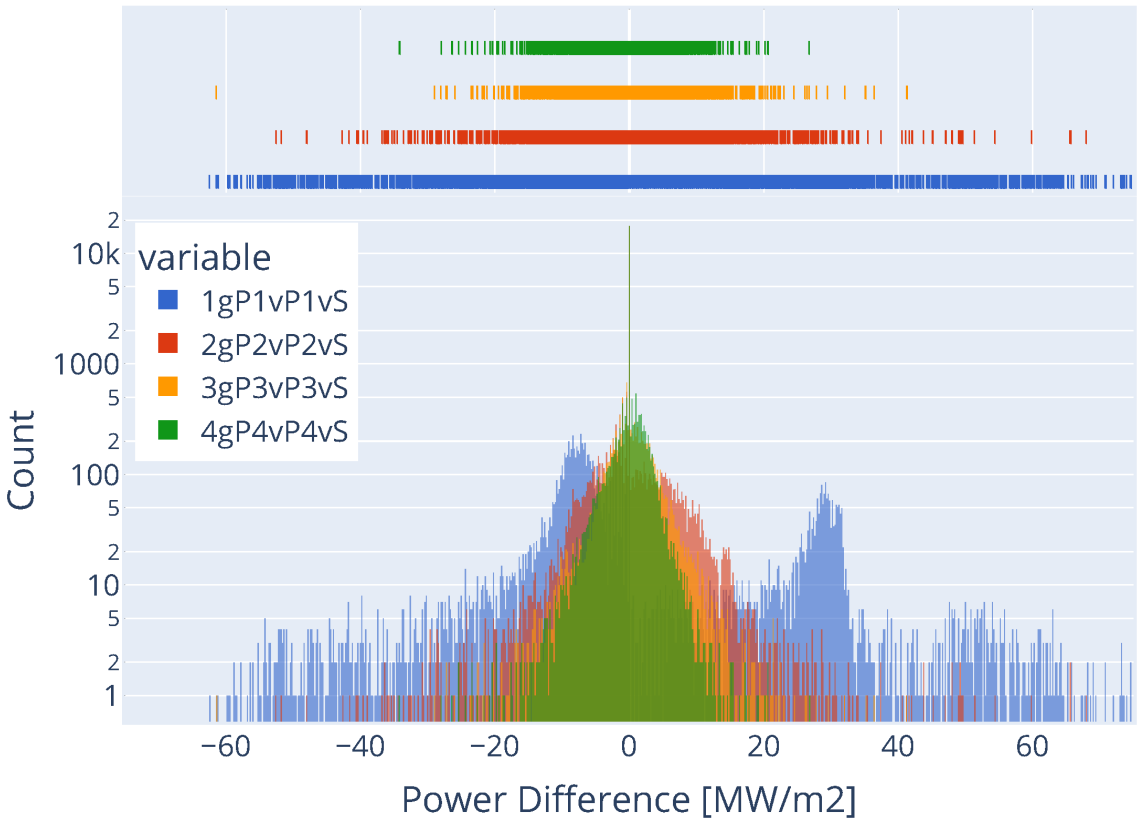


Figure 5.8: Power difference from 5gP5vP5vS (ground truth) heat flux on all mesh elements from geometry described in figure 5.5. Number values in trace names corresponds to number of samples from that distribution, i.e. 2gP3vP4vS would signify 2 samples from gyro phase angle distribution, 3 samples from velocity phase angle distribution, and 4 samples from maxwellian speed distribution. Distributions are discussed in section 3.9.2. Top plot is a rug plot, which plots ticks along 1D coordinate. Bottom plot is histogram. Increasing the samples decreases the spread and skew of the distributions, corresponding to convergence.

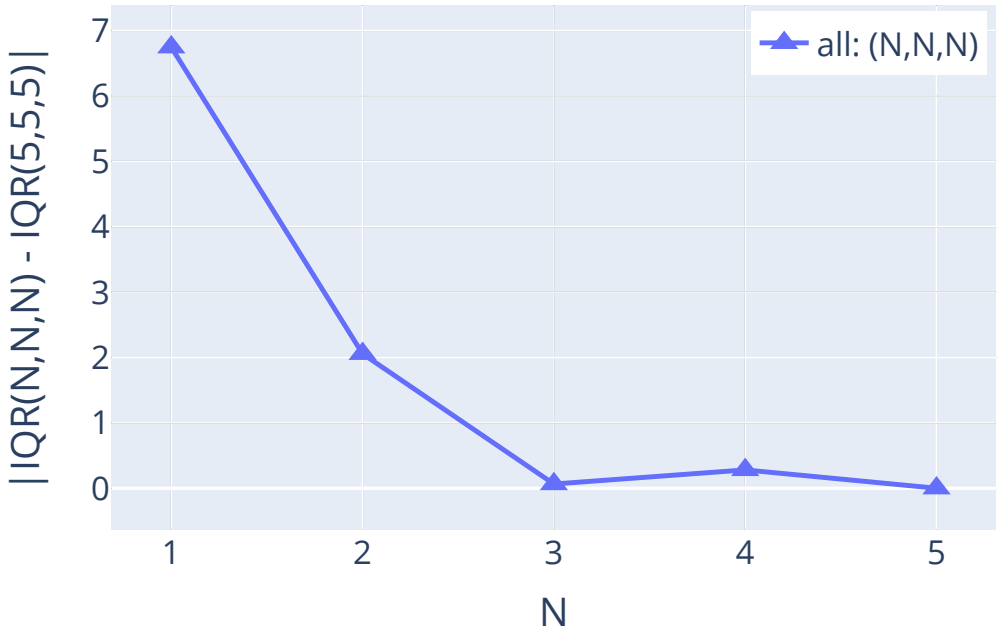


Figure 5.9: Absolute value of interquartile Range (IQR) for each dataset plotted in figure 5.8 subtracted from the ground truth IQR (i.e. $|IQR(N,N,N) - IQR(5,5,5)|$). Convergence in IQR corresponds to converging data spread. Data point at N=5 for (5,5,5) is 0 because IQR(5,5,5) - IQR(5,5,5) = 0

5.3.2 Power Balance Convergence Studies

For PFCs that are poloidally narrow, the magnetic field will not change greatly across the PFC surface. If a gyro orbit source plane is used as the source of gyro orbit power on a poloidally narrow PFC, the optical and gyro orbit integrated power can be compared directly to determine power balance for the gyro orbit profile. Figure 5.5 provides an example geometry for NSTX-U that has a poloidally thin PFC in the region of interest. This geometry was utilized to study the convergence of the gyro orbit heat load solutions. Section 3.9.2 provides an introduction to the HEAT algorithms for sampling from the kinetic distribution functions for the gyro-orbit module. A macroparticle's velocity components and spatial location will be determined by sampling these distributions. To review, the gyro phase angle represents the particle's initial location at the gyro-orbit source plane, and the variable NgP is used to describe the number of samples taken from this uniform distribution. The velocity phase angle represents the energy sharing between perpendicular and parallel velocity components, and the variable NvP is used to describe the number of samples taken from this uniform distribution. Last, the variable NvS is used to describe the number of samples pulled from a Maxwellian speed distribution that corresponds to the total velocity magnitude. More details on these distributions and the equations that describe them can be found in chapter 3.9.2.

A HEAT gyro orbit run consists of running N (N=NgP*NvP*NvS) macro-particle simulations, where each simulation corresponds to a different sample from the three dimensionsal kinetic domain (gyro phase, velocity phase, speed maxwellian). As the number of samples is increased, the simulation results will converge, corresponding to a convergence to the true gyro orbit heat flux profile. For poloidally narrow PFCs, the integrated gyro orbit power should also converge to the optically integrated power. Figure 5.10 illustrates this convergence to the optically integrated power, for 5 HEAT runs. The power across the entire PFC was integrated for both the gyro orbit and optical heat flux profiles, and the error percentage is defined as,

$$\text{Error } \% = 100\% \cdot \frac{|P_{gyro} - P_{optical}|}{P_{optical}} = 100\% \cdot \frac{\sum_i |q_{gyro,i} - q_{optical,i}| \cdot A_i}{\sum_i q_{optical,i} \cdot A_i} \tag{5.1}$$

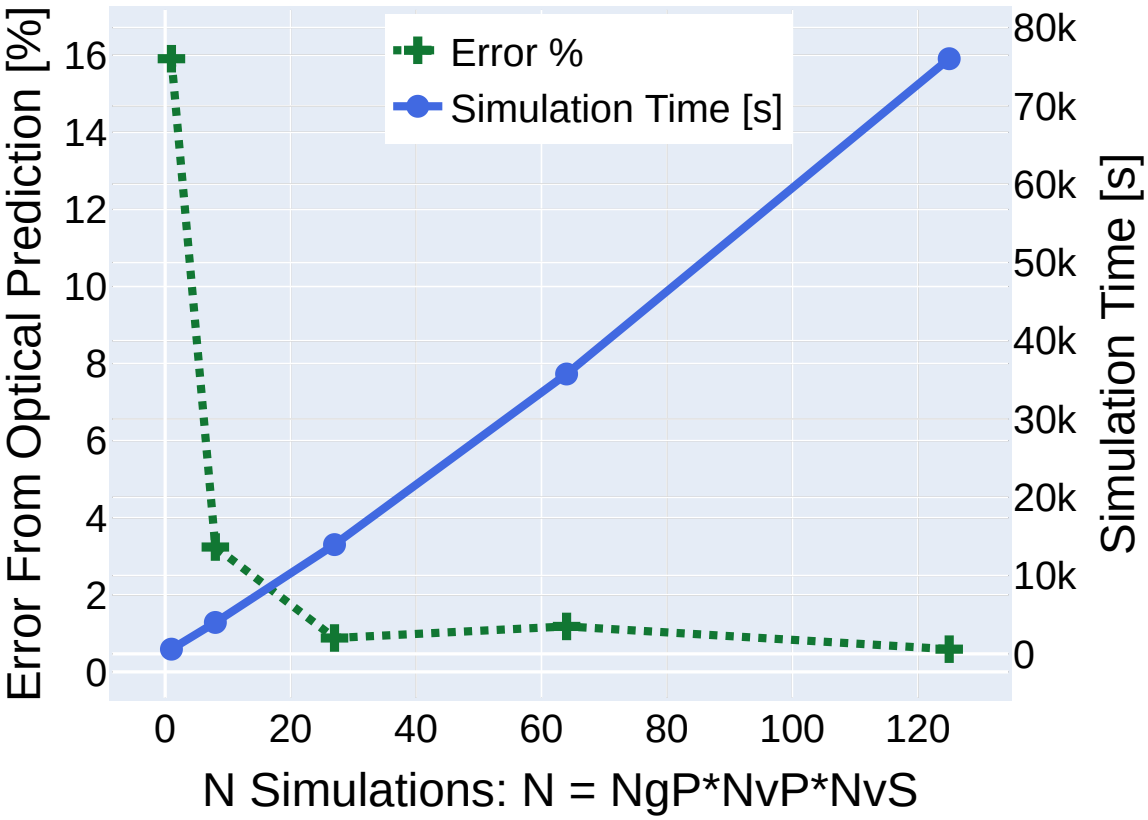


Figure 5.10: Green trace is the gyro orbit calculation error percentage with respect to the optical calculation. Simulation time for gyro-orbit calculation in blue solid trace. Gyro-orbit calculation consists of running N macro-particle simulations (N=NgP*NvP*NvS), each corresponding to a different sample from the kinetic theory distribution functions. NgP is number of gyro phase samples, NvP is number of velocity phase samples, NvS is number of samples from the Maxwellian speed distribution function.

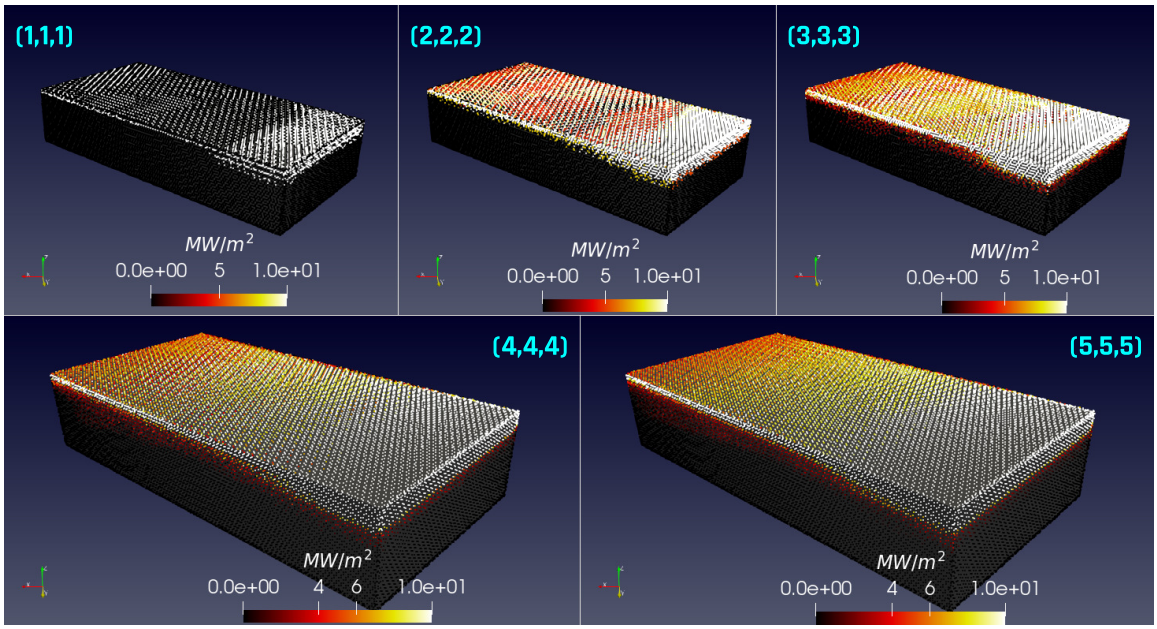


Figure 5.11: Example HEAT output for 5 different macro-particle quantities. Three numbers in parentheses indicate number of samples from each distribution: gyro phase angle, velocity phase angle, and Maxwellian speed distribution, (gyroPhase, vPhase, vSlice). These results are plotted in figure 5.10 to illustrate relative error and speed. Increasing the number of samples beyond (3,3,3) does little to improve the integrated power error, but creates a smoother profile.

where q represents the power, A represents mesh element area, and the summation over i represents the summation over all mesh elements. The results in figure 5.10 indicate that the integrated power converges to the optically integrated power when NgP=NvP=NvS=3, which corresponds to 27 samples from the three dimensional kinetic domain. Increasing the number of samples in the domain beyond this number does little to decrease the error percentage. It should be noted that these results are for a 10 eV plasma, with a maximum mesh element length of 0.1 mm. The mesh size must be considered when analyzing the convergence of the gyro orbit heat flux profiles, as it can greatly change the convergence rate. For example, if the combination of plasma temperature and magnetic field strength yield a gyro radius of 1 mm, then a 20 mm PFC mesh will not be able to resolve the differences between optical and gyro orbit profiles. In other words, the mesh size must be chosen so that for the smallest gyro radius in the kinetic domain, aliasing will not occur. In addition to the HEAT gyro orbit error, figure 5.10 also provides the simulation times for each of the HEAT runs. As expected, the total simulation time increases linearly with the number of macroparticle traces.

Figure 5.11 provides the 3D heat flux profiles that correspond to the aforementioned data from figure 5.10. Each image has a set of numbers in parentheses, corresponding to the number of samples from each dimension of the macro-particle space (gyro phase angle, velocity phase angle, Maxwellian speed distribution). Increasing the number of macro-particles from 1 to 27, in the figure (1,1,1) to (3,3,3), greatly changes the heat flux footprint. Increasing the number of macro-particles beyond (3,3,3) does little to qualitatively change the heat flux profile on the PFC surface, but does generate a smoother profile. Combining the results from figures 5.10 and 5.11 indicates that (3,3,3) is sufficient to capture the 3D heat flux profile with reasonable fidelity. Increasing from (3,3,3) to (5,5,5) takes approximately 6X longer to run the calculation, and provides a marginally better heat flux profile.

5.4 NSTX-U Castellation Gyro Orbit Analysis

Chapter 4 provided optical approximation analysis for NSTX-U. This section provides gyro-orbit approximation predictions for a steady state lower single null diverted

discharge for the NSTX-U castellated PFCs, and compares these predictions to the optical approximation. The results provided here are for the lower divertor, specifically the InBoard Horizontal Divertor (IBDH) tile, which employs a castellated graphite design as described in section 2.5 and illustrated here in figure 5.12. Each individual castellation is protected from near normal angle of incidence heat loads on the leading edges by a toroidal fish-scale. There is no poloidal fish-scale, as the poloidally-facing edges do not see angles of incidence as high as the toroidally-facing edges. Figure 5.12 provides a visualization of the CAD for a single castellation, and calls out its location in the lower divertor.

For a 10 eV (average ion temperature, see equation 3.30) steady state deuterium plasma, a side by side comparison between the optical and gyro orbit approximation is provided in figure 5.13. The castellation of interest described in figure 5.12 is the CAD that heat fluxes are overlaid onto in figure 5.13. Four frames are provided in the figure, where the left column is the optical approximation and the right column is the gyro orbit approximation. The top row of the figure provides a view of the castellation from a downstream vantage point, while the bottom row provides a view of the castellation from the upstream vantage point. Each castellation surface is labeled with two coordinates, S_{tor} which extends in the toroidal direction, and S_{pol} which extends in the poloidal direction. The S variables begin at $S = 0$ on the top surface corner at each respective location. The colorbar has been capped at 15 MW m^{-2}.

The optical approximation heat flux predictions (left column) in figure 5.13 demonstrate some of the characteristic effects of the optical approximation. The top surface consists of a magnetic shadow on the toroidally upstream leading edge, caused by the toroidal fishscale on the upstream (not-pictured) castellation. The heat loaded region on the top of the castellation has a relatively uniform heat flux. The magnetic shadow has a sharp edge, which is characteristic of the optical approximation, and is caused by the relatively 'straight' magnetic field lines when compared to the small 0.5 mm inter-castellation gaps. This shadow reduces the heat loaded area by approximately 15%, which increases the peak heat flux on the uniform region of the PFC. Additionally, an optical hot spot occurs on a poloidally-facing leading edge. This hot spot occurs where the rectangular PFC surface and the curved magnetic field line interact to generate a faceted heat flux on the downstream PFC corner. Because the

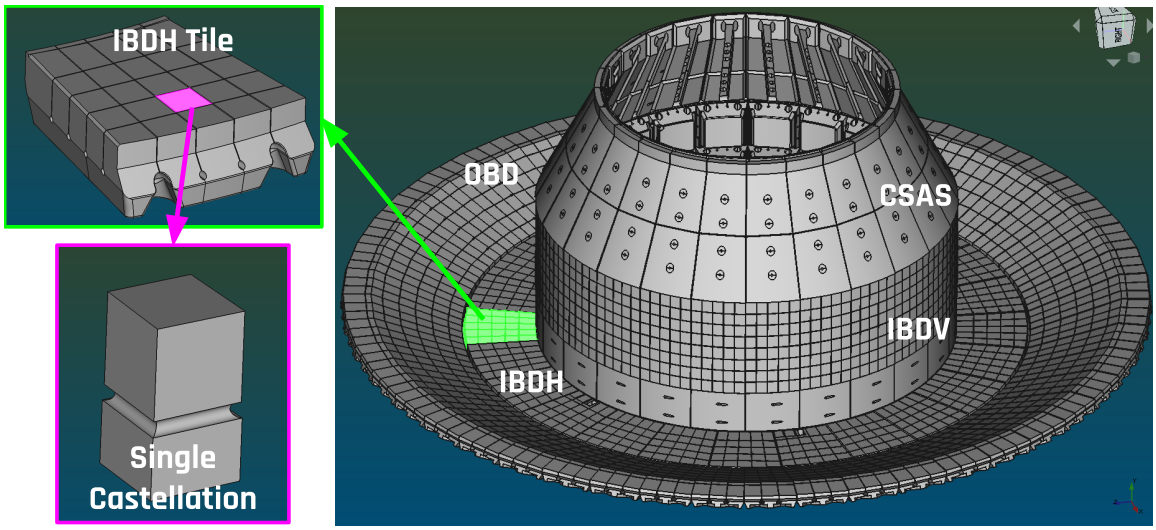


Figure 5.12: Divertor geometry for gyro orbit analysis of lower single null steady state plasma. Entire lower divertor is shown on the right. IBDH tile is illuminated in green, and exploded view provided with green outline. A single castellation is provided in lower left figure. Analysis in this section uses single castellation geometry.

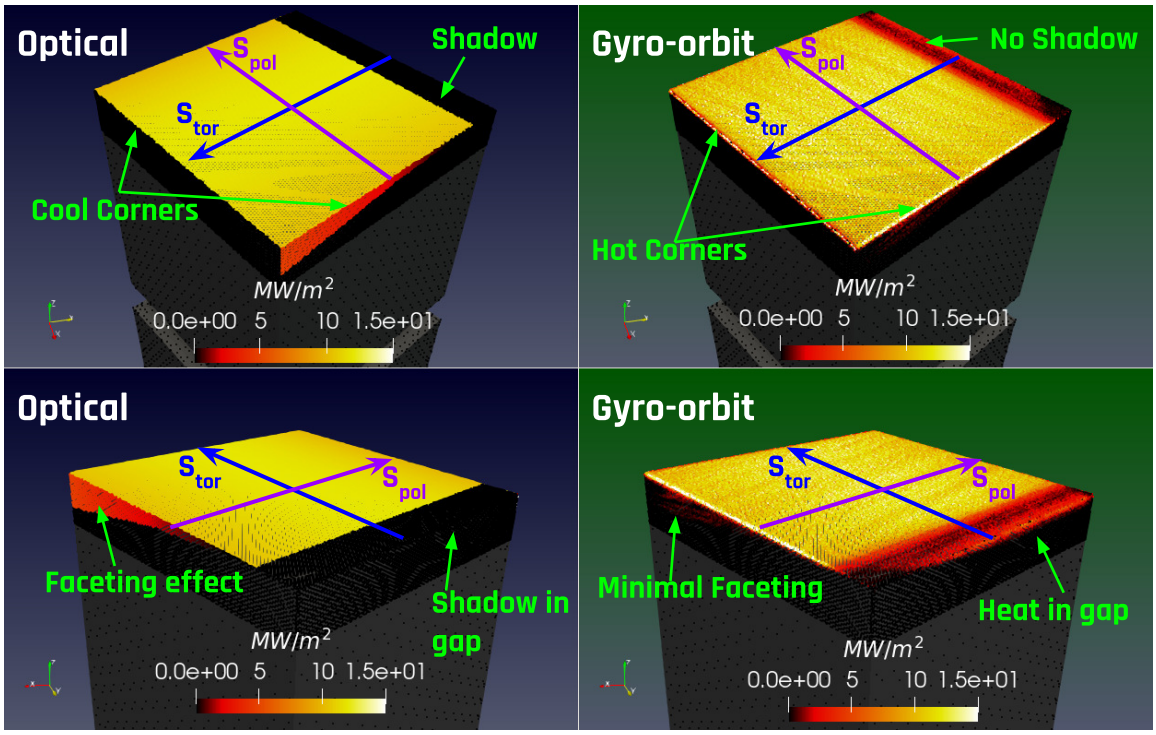


Figure 5.13: Optical and gyro-orbit heat flux predictions for a single castellation in the NSTX-U IBDH tile. Left frames correspond to the optical approximation. Right frames correspond to the gyro-orbit frames. Top and bottom rows each show different vantage points. Primary differences between optical and gyro-orbit approximation are the magnetic shadows and leading edges.

optical power does not have a gyrating component, it can descend into the gaps between castellations where the curvature of the magnetic field line along the flat PFC surface results in a changing $\hat{b} \cdot \hat{n}$ (from equation 3.26) along the PFC surface. The changing $\hat{b} \cdot \hat{n}$ causes the peak of the heat flux facet to occur at the PFC corner.

The gyro-orbit approximation heat flux predictions (right column) in figure 5.13 differ from the optical approximation predictions in the shadows and on leading edges. The optical shadow described in the preceding paragraph is almost completely smeared out by the gyro-orbit heat fluxes. This is because the helical trajectories of ions enable them to dip into the magnetic shadow and strike the optically shadowed surface. The gyro-radius and gyro-period are relatively small compared to the fish-scale step size and inter-castellation gap size, which is why these trajectories are capable of accessing the optical shadow. The heat fluxes in this shadow are generally less than 5 $\mathrm{MW\,m^{-2}}$, but they do increase the plasma wetted area of the castellation top surface, thereby reducing the peak heat loads. The corners in the gyro-orbit predictions are hot. The bottom right frame of 5.13 illustrates that the heat flux on the castellation poloidally-facing corner is nearly 19 $\mathrm{MW\,m^{-2}}$, which is the highest heat load on the entire PFC surface. This hot corner is caused by ions whose helical trajectories enable them to wrap around the castellation corner and strike the vertical edge of the castellation, as well as the horizontal top surface, creating an amplified heat load on the corner.

The optical hot spot that occurs as a result of the faceting effect is absent in the gyro-orbit prediction. Optical power that strictly follows the magnetic field line is able to descend into the gap between castellations and strike the flat castellation edge. Gyro-orbit power, however, is scraped off on the top corners of castellations before descending down into the gaps. A certain subset of macroparticles, with specific helix pitch and gyro-radii, will be able to descend into the gap. In this manner, the gap widths can be interpreted as a bandpass particle filter, but only a small fraction of the gyro-orbit distribution of particles will be in the passband, resulting in the dissapearance of the optical hot spot. This phenomenon could also be used as a diagnostic, in a manner similar to a Katsumata probe [71]. Additionally, on the toroidally facing leading edge, a heat load appears in the gyro-orbit approximation. This heat load on the vertical

castellation face arises because some helical trajectories can dip into the poloidally-running gap and strike the toroidally-leading edge of the castellation despite the fish-scale designed to protect this edge. While the heat fluxes in this gap are low (2-5 $\mathrm{MW\,m^{-2}}$), they are non-negligible. This example illustrates that despite the toroidal fish-scale, power can access regions designed to be protected and should be considered when designing the PFCs.

While viewing the heat flux over an entire castellation is useful for gaining intuition into the gyro-orbit heat fluxes, analyzing a 1D chord in both the toroidal and poloidal directions, as described in section 5.1, is useful for quantitatively providing insight into the heat loads. Analysis in 1D also enables the NSTX-U castellated graphite design to be compared to similar work performed for ITER by Gunn [17] and Komm [18], as described in sections 1.3.2 and 5.1. The 1D profiles used for the castellation analysis are depicted in figure 5.13 as S_{tor} and S_{pol}, corresponding to the toroidal and poloidal directions, respectively. A range of mesh elements on either side of the S_{tor} and S_{pol} coordinate lines in figure 5.13 are 'collapsed' into a single 1D coordinate. The heat flux profiles can then be plotted in 1D, as depicted in figures 5.14 and 5.15.

Figure 5.14 illustrates the heat fluxes along the 1D coordinate, S_{tor}, depicted in in figure 5.13. $S_{tor} = 0$ corresponds to the top surface corner and the castellation top surface is called out with an arrow in magenta. Optically shadowed regions are greyed out. At $S_{tor} = 0$, the optical heat flux (green trace) is zero, due to the magnetic shadow cast by an upstream castellation fish-scale. At approximately $S_{tor} = 4$mm, the magnetic shadow ends and the optical heat flux maintains approximately 12 $\mathrm{MW\,m^{-2}}$ across the remainder of the castellation top surface. The gyro-orbit heat flux (blue trace), however, can access regions that are optically shadowed. For this example the gyro-orbit heat fluxes begin on the $S_{tor} < 0$ surface, which is designed to be protected by a toroidal fish-scale. In the optical shadow, the gyro-orbit heat flux peaks on the $S_{tor} = 0$ corner, and is non-zero thereafter. In the center of the castellation, the heat fluxes are similar between the optical and gyro-orbit approximations.

Figure 5.14 illustrates the heat fluxes along the 1D coordinate, S_{pol}, depicted in in figure 5.13. $S_{pol} = 0$ corresponds to the top surface corner and the castellation top surface is called out with an arrow in magenta. Optically shadowed regions are greyed out. The heat fluxes in the center of the castellation are similar between the optical and

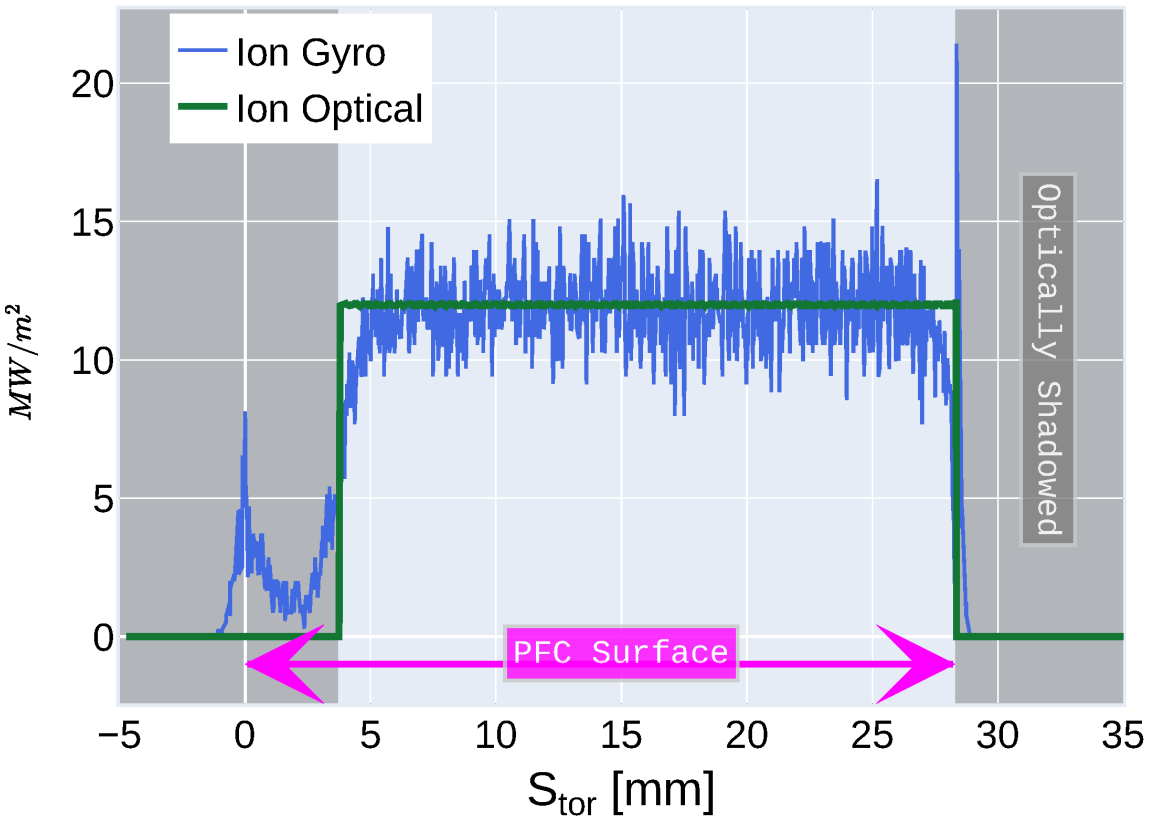


Figure 5.14: Optical (green) and gyro-orbit (blue) heat flux traces along a 1D coordinate on the castellation surface in toroidal direction. $S_{tor} = 0$ corresponds to the top surface corner. Castellation top surface is called out with an arrow in magenta. Optically shadowed regions are greyed out. Gyro orbit heat flux can access regions that are optically shadowed.

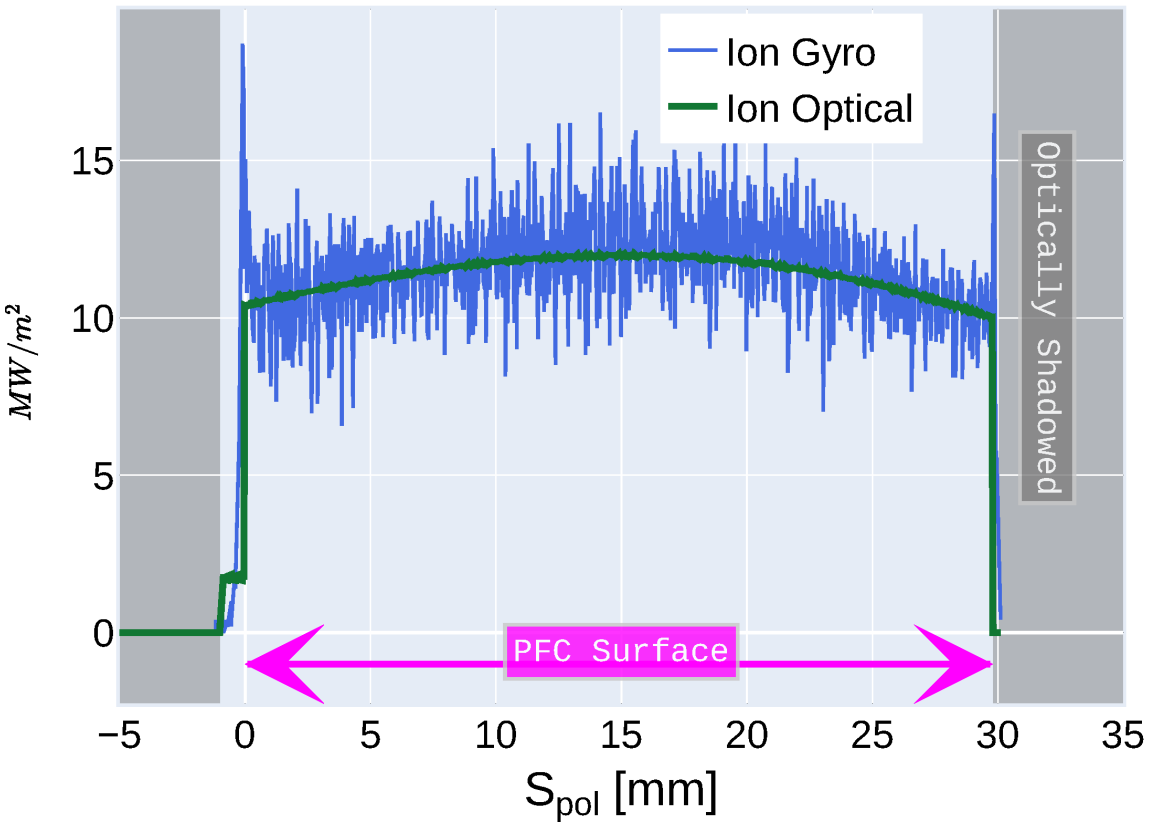


Figure 5.15: Optical (green) and gyro-orbit (blue) heat flux traces along a 1D coordinate on the castellation surface in poloidal direction. $S_{pol} = 0$ corresponds to the top surface corner. Castellation top surface is called out with an arrow in magenta. Optically shadowed regions are greyed out. Gyro orbit heat flux peaks on corners, due to the ability of the gyro-orbit helical trajectories to wrap around top surface corners and strike vertical side surfaces.

gyro-orbit approximations for the poloidal chord. The primary difference between the profiles is on the corners, where $S_{pol} = 0$ and $S_{pol} = 28mm$. The gyro-orbit heat flux on the corners is far higher than the optical approximation, because the helical trajectories of ions enables them to wrap around the castellation corners. This spike in heat flux is nearly 2X the optical approximation for the poloidally upstream corner, and peaks out at nearly 19 MW m^{-2}. This high heat flux on the castellation poloidally-facing corners is a characteristic feature of the gyro-orbit heat flux model.

In order to quantify how the gyro-orbit heat fluxes will impact the PFCs with respect to their engineering limits, a thermal analysis has been performed using the heat fluxes depicted in figure 5.13. Figure 5.16 provides a visualization of the 3D temperature profile for the optical and gyro orbit approximations. The NSTX-U PFCs are temperature limited rather than stress limited, so temperature is the quantity of interest for this discussion. Section 3.10 provides a description of how HEAT generates temperature profiles. Figure 5.16 provides six snapshots of the castellation temperature. Each column of snapshots corresponds to a unique point in time, and three time points have been chosen to illustrate the temporal evolution of the castellation temperature. The top row in each column uses the optical approximation heat flux as a boundary condition for the finite volume mesh, and the bottom row in each column uses the gyro-orbit approximation heat flux.

The discharge simulated in figure 5.16 lasts for 5s. At the onset, the PFC is assumed to be in thermal equilibrium at 300 Kelvin. Once the discharge commences, heat begins to load the top surface and castellation edges as illustrated in figure 5.13. At 0.3 s, which corresponds to the left column in figure 5.16, the optical approximation temperature begins to peak on the corner that is associated with the faceted optical hot spot. Simultaneously, the gyro orbit temperature peaks on the narrow poloidal leading edge that is exposed to the helical ion gyro orbits. In the center column of figure 5.16, at 2.65s, the castellation surfaces begin to show temperatures in excess of 1700 Kelvin. A large continuous region of the optical castellation surface is above this temperature at 2.65s, whereas the gyro castellation has two distinct regions of high temperature (on the edge and in the center). The optical hot spot that arises from the faceting effect in the optical approximation drives a faster temperature increase than the narrow helical leading edge observed in the optical approximation, despite the helical leading edge

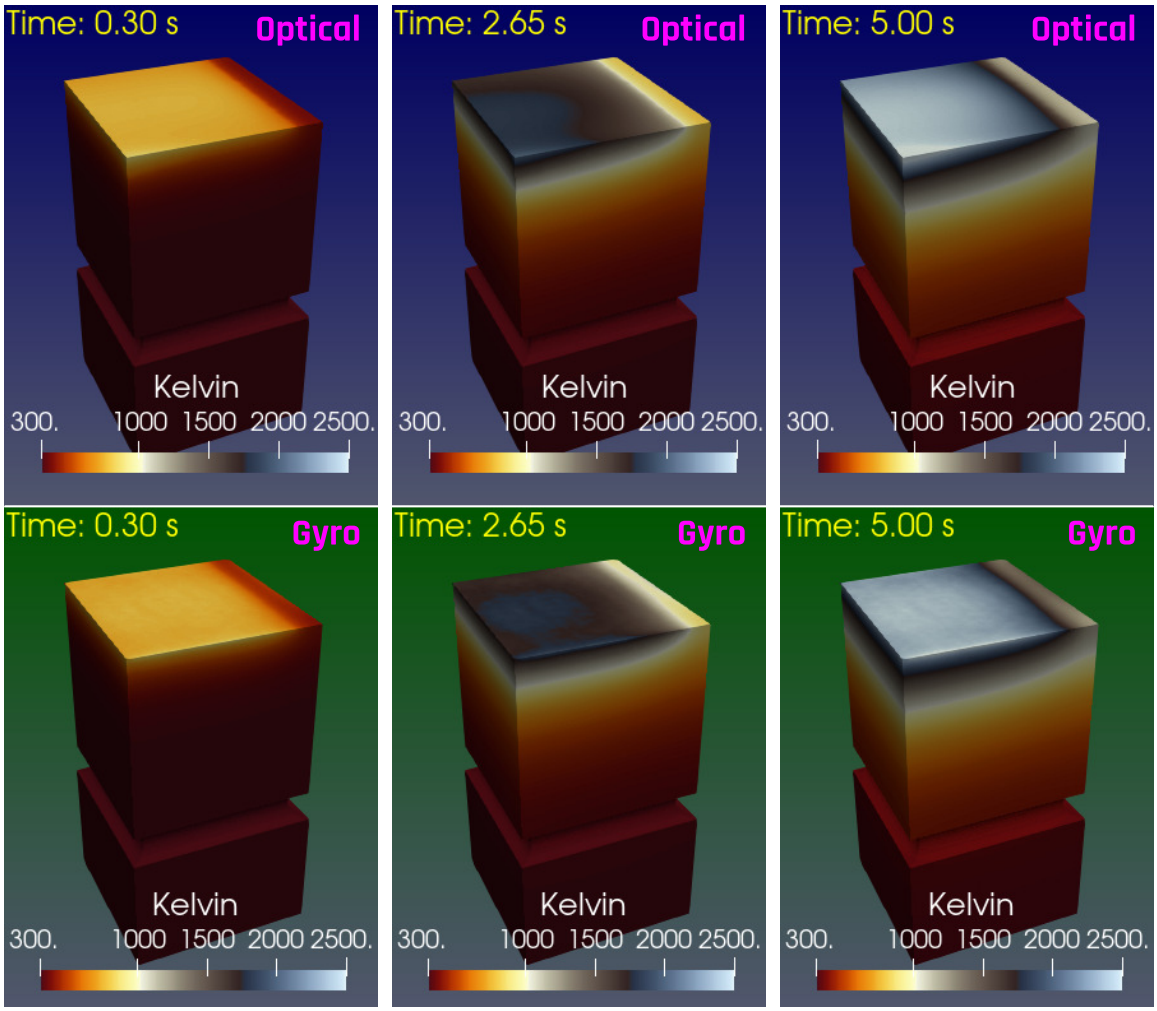


Figure 5.16: Time evolution of temperature for the NSTX-U castellation corresponding to the heat fluxes in figure 5.13. Top row shows temperatures from optical approximation heat flux, bottom row shows temperatures from the gyro-orbit approximation heat flux. Each column corresponds to a different timestep. Peak temperature for the optical approximation is reached on the faceted optical hot spot corner. Peak temperature for the gyro-orbit approximation occurs on the narrow poloidal facing leading edge. Peak temperatures are within 5% of each other.

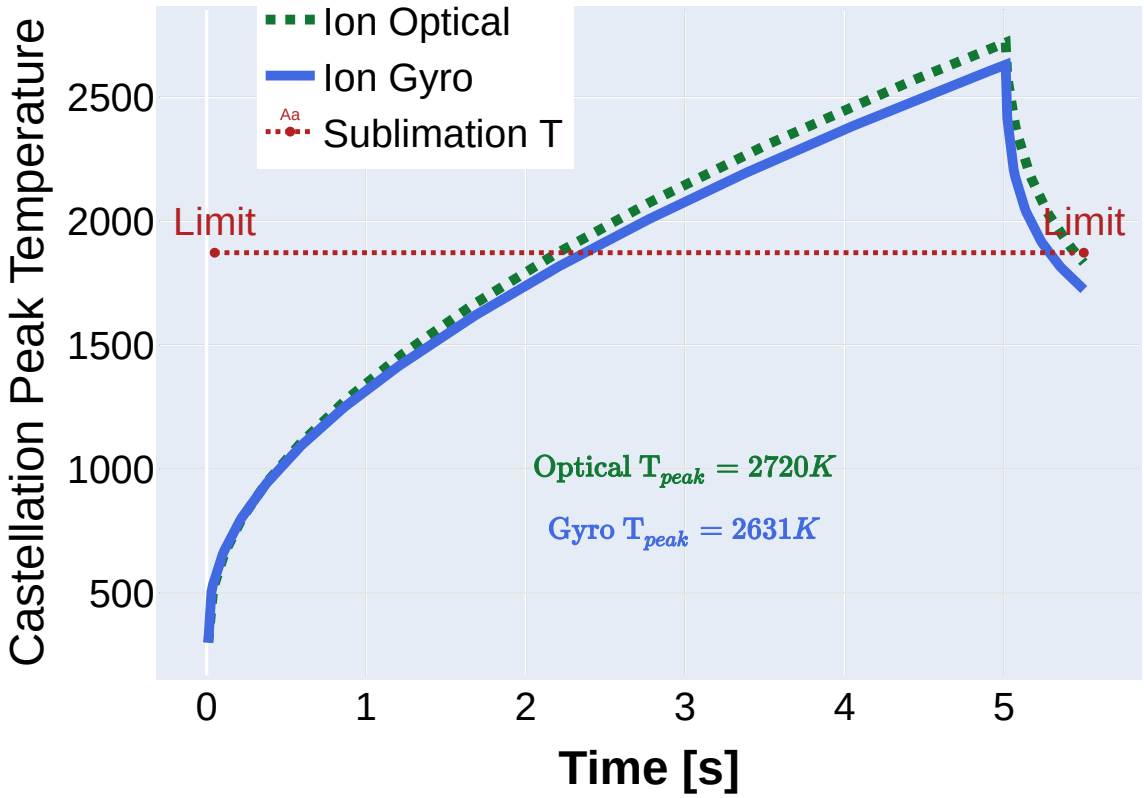


Figure 5.17: Peak temperature comparison for optical (green dotted trace) versus gyro-orbit (blue solid trace) approximations, for $P_{SOL,outer} = 4.9MW$. Heat fluxes correspond to figure 5.13. Temperatures correspond to figure 5.16. For both heat flux models, the peak temperature is reached at the end of the discharge when t=5s. The peak temperature is 2720 K for the optical approximation and 2631 K for the gyro-orbit approximation. Optical trace crosses the temperature limit (1873 K) at 2.21s and gyro-orbit trace crosses the limit at 2.37s.

having a higher peak heat flux. At the final timestep in figure 5.16 (right column), both the heat flux models have temperature profiles that have exceeded 2500 K. The optical hot spot in the optical approximation reaches 2720 K, and the gyro-orbit helical leading edge reaches 2631 K. Both of these simulations exceeded the engineering limit of NSTX-U graphite, which is 1873 K.

In figure 5.17, the peak temperatures over the entire castellation for the optical (green dotted trace) and gyro-orbit (blue solid trace) approximations are plotted. These traces represent the maximum temperature in any volume mesh cell over the entire castellation at each timestep. The engineering temperature limit has also been overlaid onto the plot at 1873 K. The simulated discharge consists of a 5s pulse followed by a 500ms cooldown period. At approximately 2.21s, the optical approximation temperature exceeds the engineering limit. At approximately 2.37s, the gyro orbit trace also exceeds the engineering limit. At 5.01s, both traces reach their maximum temperature for the entire discharge. The optical heat flux generates a temperature of 2720 K, and the gyro-orbit heat flux generates a temperature of 2631K. Both traces are within 5% of each other at the peak.

It is somewhat surprising that the gyro-orbit heat flux does not generate a higher peak temperature than the optical heat flux, given that the gyro-orbit heat flux is nearly 19 $\mathrm{MW\,m^{-2}}$ on the poloidally-facing leading edge. The optical heat flux generates an optical hot spot on the vertical faceted corner, but it has a far lower magntidude of 5 $\mathrm{MW\,m^{-2}}$. The poloidally-facing leading edge that is loaded by the gyro-orbit heat flux is extremely narrow, extending approximately 0.2 mm on either side of the corner. While the peak of this heat flux is 19 $\mathrm{MW\,m^{-2}}$, its spatial extend is small, resulting in a small spatially integrated deposited energy. The optical heat flux extends nearly 3 mm down the side of the castellation, so although the optical hot spot heat flux is smaller, the spatially integrated deposited energy is higher. In other words, the bulk material can diffuse the energy from the narrow gyro-orbit leading edge more easily than the energy from the wider optical hot spot. Additionally, the magnetic shadow cast onto the toroidal leading edge by the upstream castellation fish-scale reduces the heat flux wetted area by approximately 15% for the optical approximation. For the gyro-orbit approximation, however, this shadow gets filled in by ions with helical trajectories that can load the surface inside the magnetic shadow, thereby increasing the plasma wetted

area and decreasing the peak heat load on the top surface. The combination of these effects results in a lower peak temperature for the ion gyro-orbit approximation.

Simulating the gyro-orbit heat fluxes takes far longer than simulating the optical heat fluxes. For situations when speed is the primary motive, such as physics operation or when generating a large number of HEAT results for a database, using the gyro-orbit approximation may not be feasible. In these cases, the optical approximation calculation can be completed quickly (usually within a few minutes for medium resolutions), and can be utilized as a conservative estimate of the heat flux. Because the peak temperature using the optical approximation will generally be higher than the gyro-orbit approximation, the engineering limits will be reached earlier when using the optical heat flux. That being said, there are certain situations when the gyro-orbit calculation is imperative. These include simulations that have geometric features that are resonant with gyro-orbits, or when the gyro orbit trajectories load components that are not designed to tolerate heat fluxes. It is up to the HEAT user to intelligently determine when the gyro-orbit calculation can be omitted, and when it is mandatory.

5.5 NSTX-U IBDH Gyro Orbit Predictions

The previous section outlined some of the gyro-orbit effects that arise for a single castellation. While this is useful to ascertain the heat loads local to a single castellation, it neglects the macroscopic effects of gyro-orbit heat fluxes on an entire PFC. It also neglects the full power of HEAT, which can analyze heat loads on entire 3D CAD geometries. This section is designed to accompany the previous section, and describe the macroscopic effects of the ion gyro-orbit heat fluxes on entire PFCs. The NSTX-U InBoard Divertor Horizontal IBDH tile has been chosen for this example, as it will see some of the highest heat loads of the entire machine. Figure 5.12 provides a visualization of the IBDH tile and where it resides in the NSTX-U divertor.

Figure 5.18 provides a visual comparison between the optical approximation and two gyro-orbit predictions for the entire IBDH tile in the lower NSTX-U divertor. HEAT requires the user to input the average plasma temperature [eV] at the PFCs (see equation 3.30), and uses this temperature to generate a Maxwellian speed distribution as described in sections 1.3.5 and 3.9. Changing the input temperature

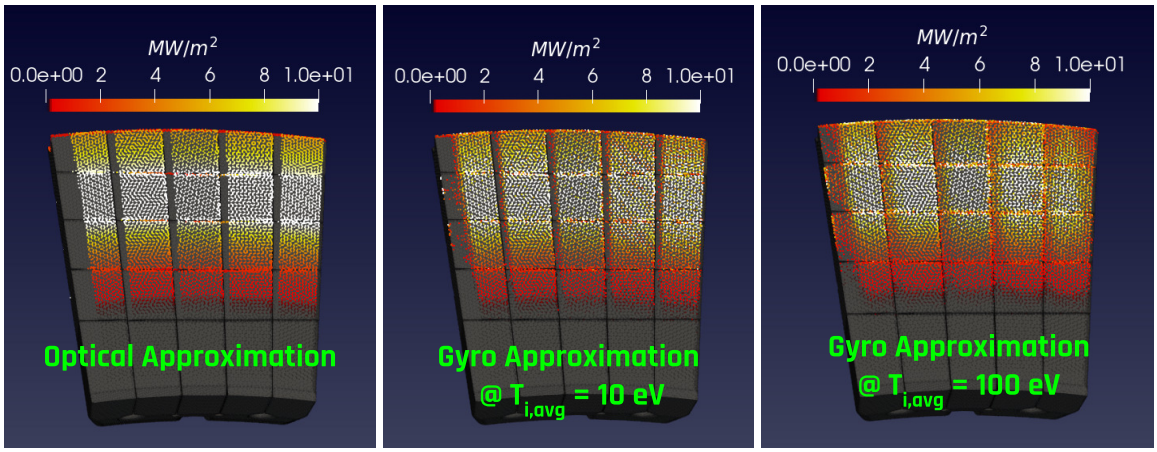


Figure 5.18: Heat fluxes for the optical approximation (left) and gyro-orbit approximation (center and right) for the entire NSTX-U IBDH tile in the lower divertor. With the optical approximation, shadows have sharp edges and cover a large fraction of the PFC. One large inter-PFC shadow can be observed on the left edge of the tile, and small inter-castellation shadows can be observed on the left edge of each castellation. Turning on the gyro-orbit approximation with average ion temperature set to 10eV fills in the inter-castellation shadows, but has a minimal effect on the inter-PFC shadows. Increasing the average ion temperature to 100eV increases the gyro-radii of ions, which results in total smearing of the optical approximation shadows. HEAT can perform gyro-orbit analyses on entire PFC geometries directly from engineering CAD.

will change the gyro-radii of particles, and increasing the temperature increases the average distribution of gyro-radii. This dependence of gyro-radii on temperature is important when calculating the gyro-orbit heat flux profiles, as the gyro-radius of an ion determines the amount it can 'dip' into an optical magnetic shadow. Figure 5.18 shows the heat loads corresponding to the optical approximation and two plasma ion input temperatures.

The optical approximation (left frame) in figure 5.18 has a heat flux footprint that is characterized by alternating areas of heat flux and shadows, which correspond to the geometric features of the IBDH tile. The narrow bands of shadows that appear at the left side of every column of castellations is due to the inter-castellation fish-scale. The inter-castellation fish-scale is small, so these shadows are narrow. The wide shadow at the left side of the IBDH tile is a result of the inter-PFC fish-scale, which is designed to protect the leading edges on adjacent IBDH tiles. As discussed in the preceding sections, the helical trajectories of ions when the gyro-orbits are accounted for enables them to access optically shadowed regions. This can be observed in the center and right frames in figure 5.18, which correspond to input temperatures of 10eV and 100eV, respectively.

For the 10eV case (center frame), the stark shadow lines of the optical approximation begin to be smeared out by ions with finite Larmor radii. The inter-castellation fish-scale shadows are approximately halfway filled in, while the inter-PFC fish-scale shadows are approximately one quarter filled. At 100eV average ion temperature (right frame), the inter-castellation fish-scale shadows have completely disappeared, and the majority of the inter-PFC fish-scale shadow has disappeared. The average ion gyro-radii at 100eV is sufficiently large to smear out almost all of the optical shadows on the PFC, which increases the PFC plasma wetted area. This increase in plasma wetted area serves to enhance the PFC performance, as the same amount of spatially integrated power is distributed over a larger surface area. Figure 5.18 does not focus on the small gyro-orbit effects on the edges (as was performed in the previous sections), but rather on the macroscopic PFC performance. For the macroscopic PFC analysis, the helical trajectories of the ions serves to spread the heat load across a surface area that is inaccessible via the optical trajectory.

Chapter 6

HEAT Results for Other Tokamaks

The primary objective of this chapter is to provide a summary of some of the non-NSTX-U work that has been performed using HEAT. So far, HEAT has been used for heat load predictions on 5 tokamaks, and some of these results are included here.

6.1 ST40

Collaborations between the U.S. National Laboratories (ORNL, PPPL) and the private company Tokamak Energy have enabled HEAT predictions for the ST40 tokamak. These collaborations are particularly exciting as they may provide experimental validation of HEAT on a highly non-axisymmetric divertor in the coming years. To assist the ST40 team, HEAT was used to generate heat flux and temperature predictions for the Programme 2.2 divertor. Because this divertor is highly optimized for specific magnetic configurations, it is also highly nonaxisymmetric. Predictions on non-axisymmetric divertors are uniquely suited to HEAT analysis, as the full 3D CAD geometry and magnetic structure combine to create complex 3D heat flux footprints.

Figure 6.1 provides an illustration of the ST40 2.2 divertor CAD geometry (top frames) as well as the simulated heat flux (lower left frame) and temperature (lower right frame) for a target ST40 magnetic scenario. The toroidal shaping, fish-scales, bevels, and chamfers, are used in the 2.2 divertor to protect leading edges from potentially detrimental heat loads. The heat flux footprint is highly non-axisymmetric due to these features yet there are no leading edges or optical hot spots. The 'blocky' structure of the heat flux footprint is caused by the small PFC tiles casting independent

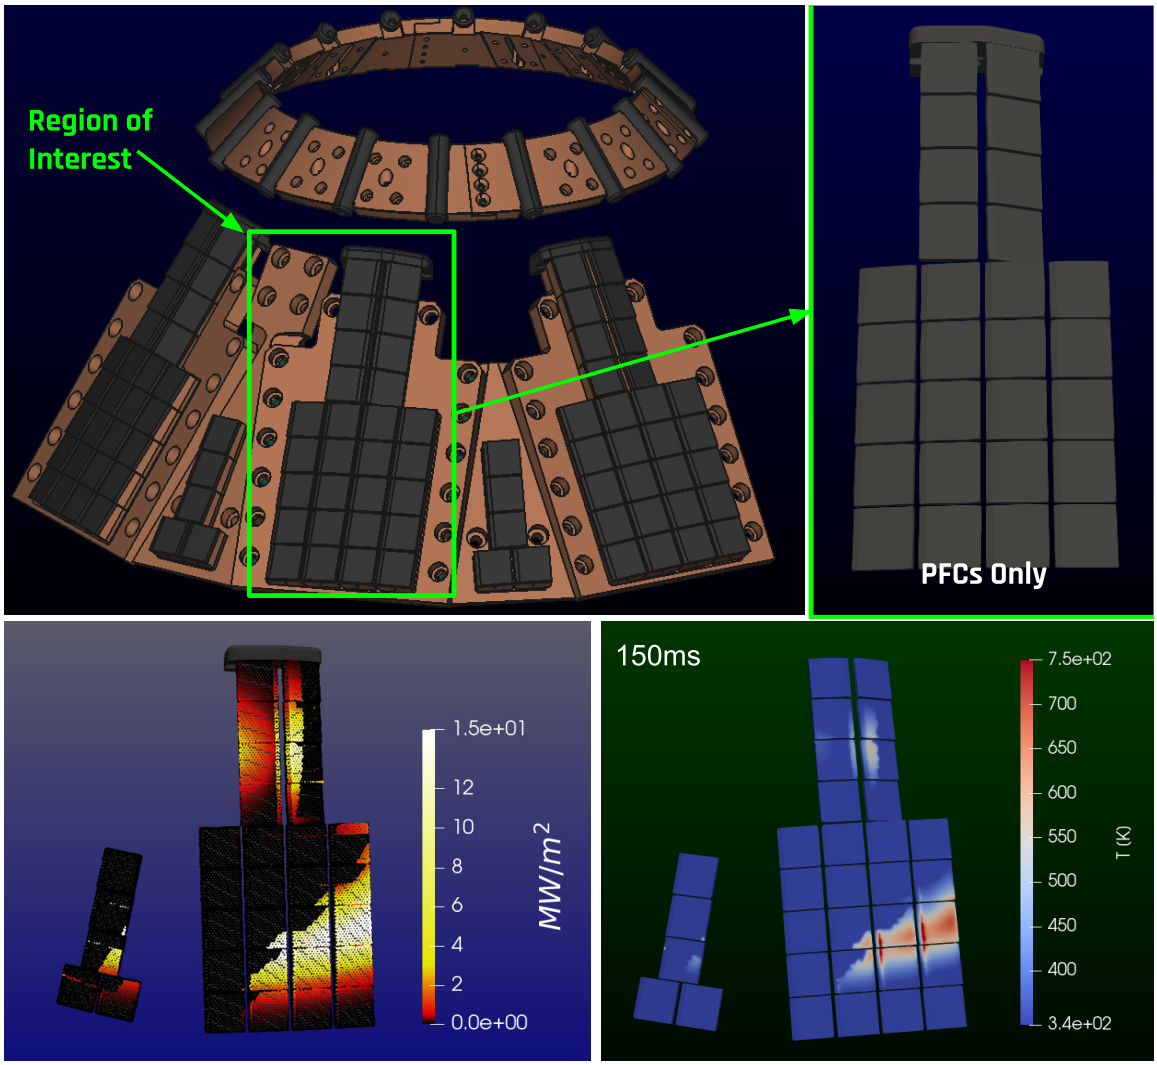


Figure 6.1: Tokamak Energy ST40 2.2 divertor CAD (top frame), heat flux (bottom left) and temperature at 150 ms (bottom right). The highly optimized 2.2 divertor creates an extremely non-axisymmetric heat flux footprint. The temperatures for this scenario are well below the enginering limits.

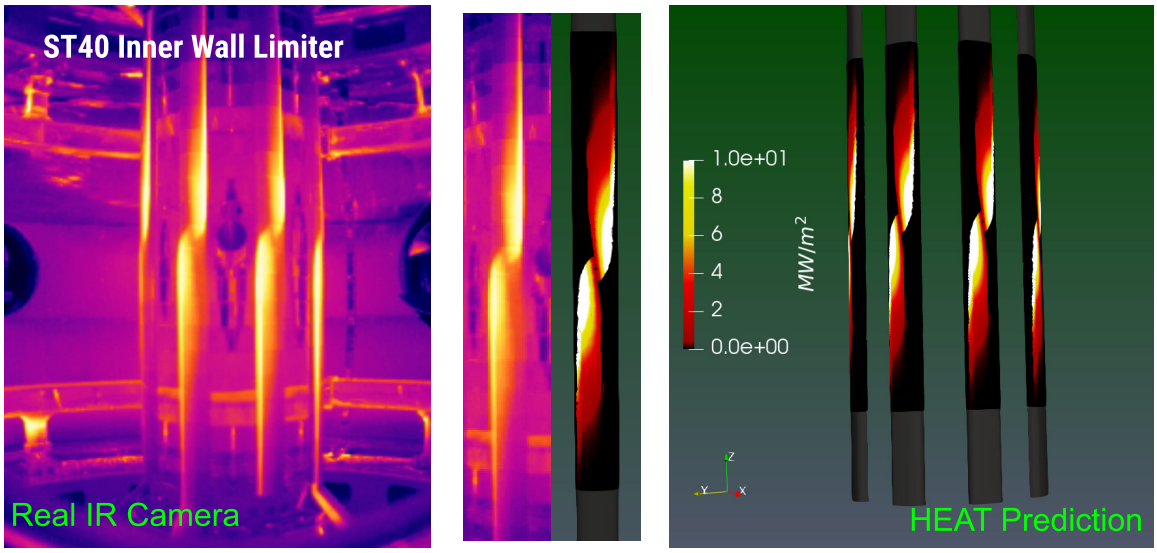


Figure 6.2: Side by side comparison of ST40 IR experimental data and HEAT predictions for a center stack limited scenario. Good qualitative agreement between the experimental and predicted profiles can be observed, but further quantitative analysis is required to validate the HEAT output.

magnetic shadows. The maximum heat flux observed for this HEAT prediction is 15 $\mathrm{MW\,m^{-2}}$ for a P_{SOL} of 1 MW. Because the pulse lengths for these discharges are relatively short (100-200 ms), it is unlikely that the PFCs will exceed the temperature or stress engineering limits. The temperature provided in the figure shows the peak temperature to be 750 K at 150 ms, which is well below the sublimation temperature for molybdenum as described in [20].

Additionally, HEAT has been used to predict limiter heat fluxes for ST40. This analysis enables a side by side comparison with real ST40 IR data, as illustrated in figure 6.2. The left panel in the figure shows the IR camera image of the center stack limiters, which repeat toroidally around the machine. The right frame in the figure shows the heat fluxes predicted by HEAT for a corresponding limited magnetic scenario. The heat flux footprints for the experimental data and the HEAT predictions indicate that the magnetic shadows calculated with the HEAT field line tracers also appear in the experimental IR data. In the center frame of the figure, a single limiter from both the IR data and the HEAT prediction are placed side by side. The limiter tangent point can be observed in both figures as a null point in the heat flux located approximately halfway up the limiter (magnetic axis). The leading edges are observed on the right side of the upper limiter, and the left side of the lower limiter, caused by the power flow direction reversal across the LCFS contact point. The leading edge heat flux flares away and decays in the toroidal direction due to the changing $\hat{b} \cdot \hat{n}$ and the geometric shape of the limiters.

The HEAT flux footprints are qualitatively similar, although a rigorous quantitative analysis has not yet been performed. It is expected that in the coming year a quantitative side by side comparison between the ST40 experimental data and the HEAT predictions will be performed for the limiter and divertor. Additionally, an ion gyro-orbit analysis will be performed and compared against experiment. For now, these results indicate that the HEAT predictions are indeed providing heat flux footprints that can be observed in experiment. Chris Marsden, HEAT user at Tokamak Energy, has provided valuable feedback on HEAT algorithms and identified critical bugs.

6.2 DIII-D Negative Triangularity

The DIII-D tokamak team has been exploring the benefits associated with operating the tokamak with negative triangularity. The confinement advantages of negative triangularity are outside the scope of this dissertation, but the heat loads that arise from such operations can be studied with HEAT. Figure 6.3 provides images of the geometry, magnetic equilibrium, and heat flux, for one of the simulated negative triangularity discharges - the so-called high X-point configuration. The outer strike point has been placed on the outer wall PFCs, which were not designed to handle the high 16.9 MW m^{-2} heat loads that are observed in the bottom frame of figure 6.3. It should be noted that a 70/30 ratio has been utilized for power sharing between the outer/inner targets. This ratio is not grounded in any scientific evidence, and will be a subject of investigation during the negative triangularity campaigns, but suffices for initial scoping studies like the one included herein. Similarly, the value of the heat flux width has been taken from Eich's empirical scaling [15], but this is also unvalidated for negative triangularity and will be a topic of scientific inquiry during the experimental campaign. The objective of these HEAT runs was to obtain general qualitative results, not explicit quantitative ones.

The heat flux profile that arises from this 'high X-point' scenario is significant for both divertor locations. The inner divertor exhibits a faceting effect as the $\hat{b}\cdot\hat{n}$ changes along the PFC surface. The gradient in heat flux associated with this faceting effect on the inner target can be observed in the lower frame of figure 6.3. The inner target heat flux peak is 8.3 MW m^{-2} on the top PFC surface, and moderate flux expansion distributes the heat flux across half of the PFC. The outer divertor exhibits less faceting, a narrower profile, and a higher peak. This is due to the different orientations of the PFCs with respect to the magnetic field lines, as well as a lesser flux expansion at the outer target.

The top frame in figure 6.4 provides a close-up view of the heat flux profiles provided in the lower frame of figure 6.3, but from a different toroidal vantage point. These PFCs are not toroidally shaped, and as such the toroidally-facing leading edges are exposed to near normal angle of incidence heat fluxes. Both the inner and outer targets have exposed leading edges in this configuration. The inner target leading edge heat flux is

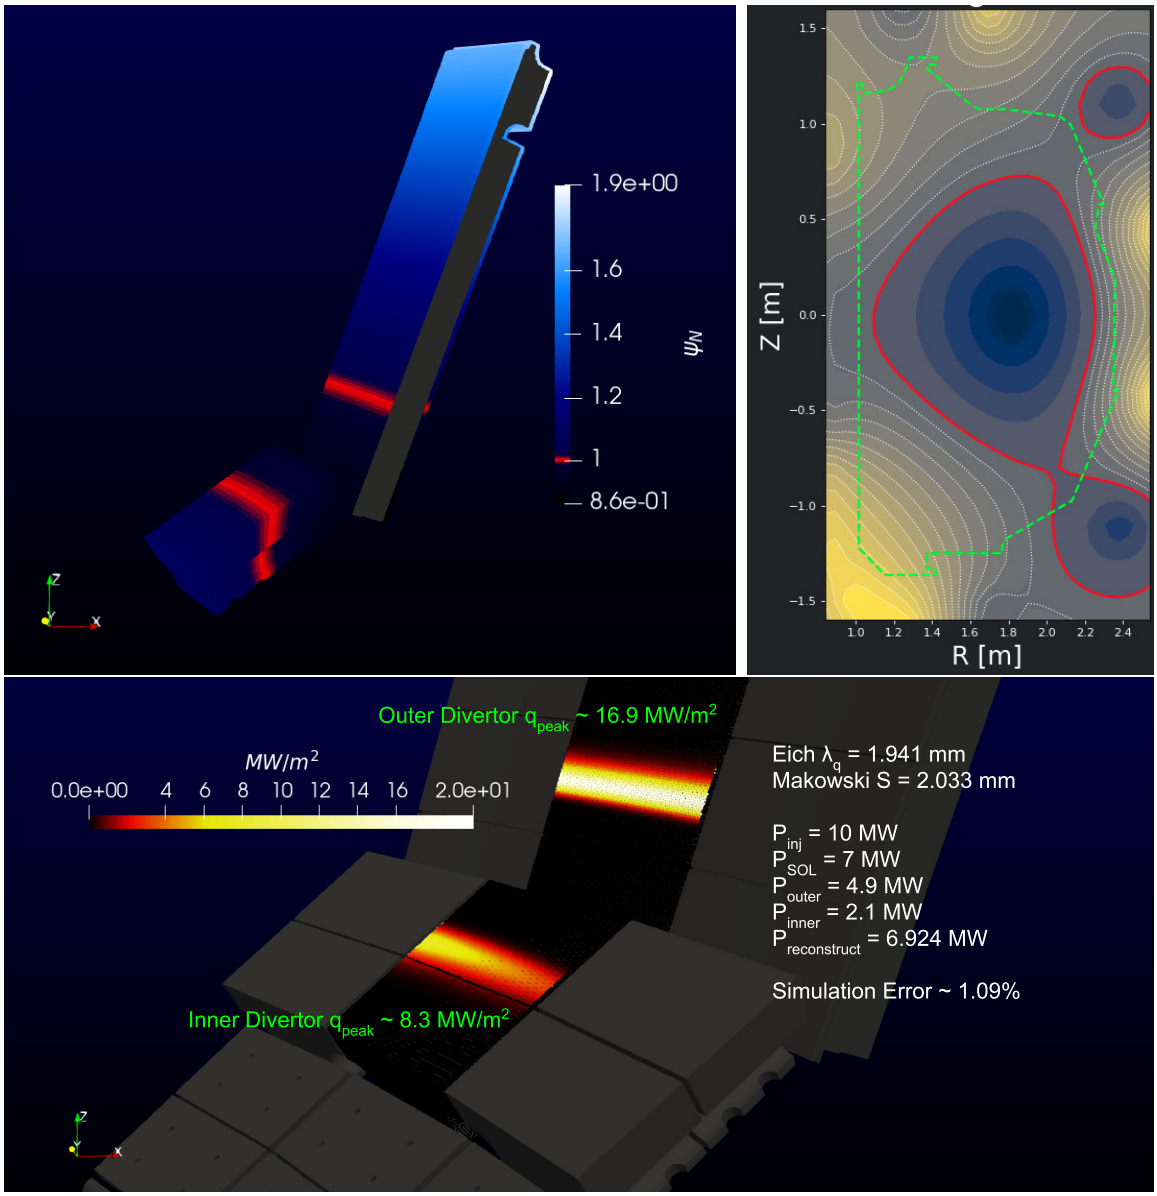


Figure 6.3: Normalized poloidal flux (top left) and MHD equilibrium (top right) for the high X-point negative triangularity scenario in DIII-D. Outer strike point is on wall PFCs, which were not designed to handle high heat loads. Bottom frame shows heat flux that corresponds to this magnetic equilibrium and geometry. Plasma parameters are overlaid onto the figure.

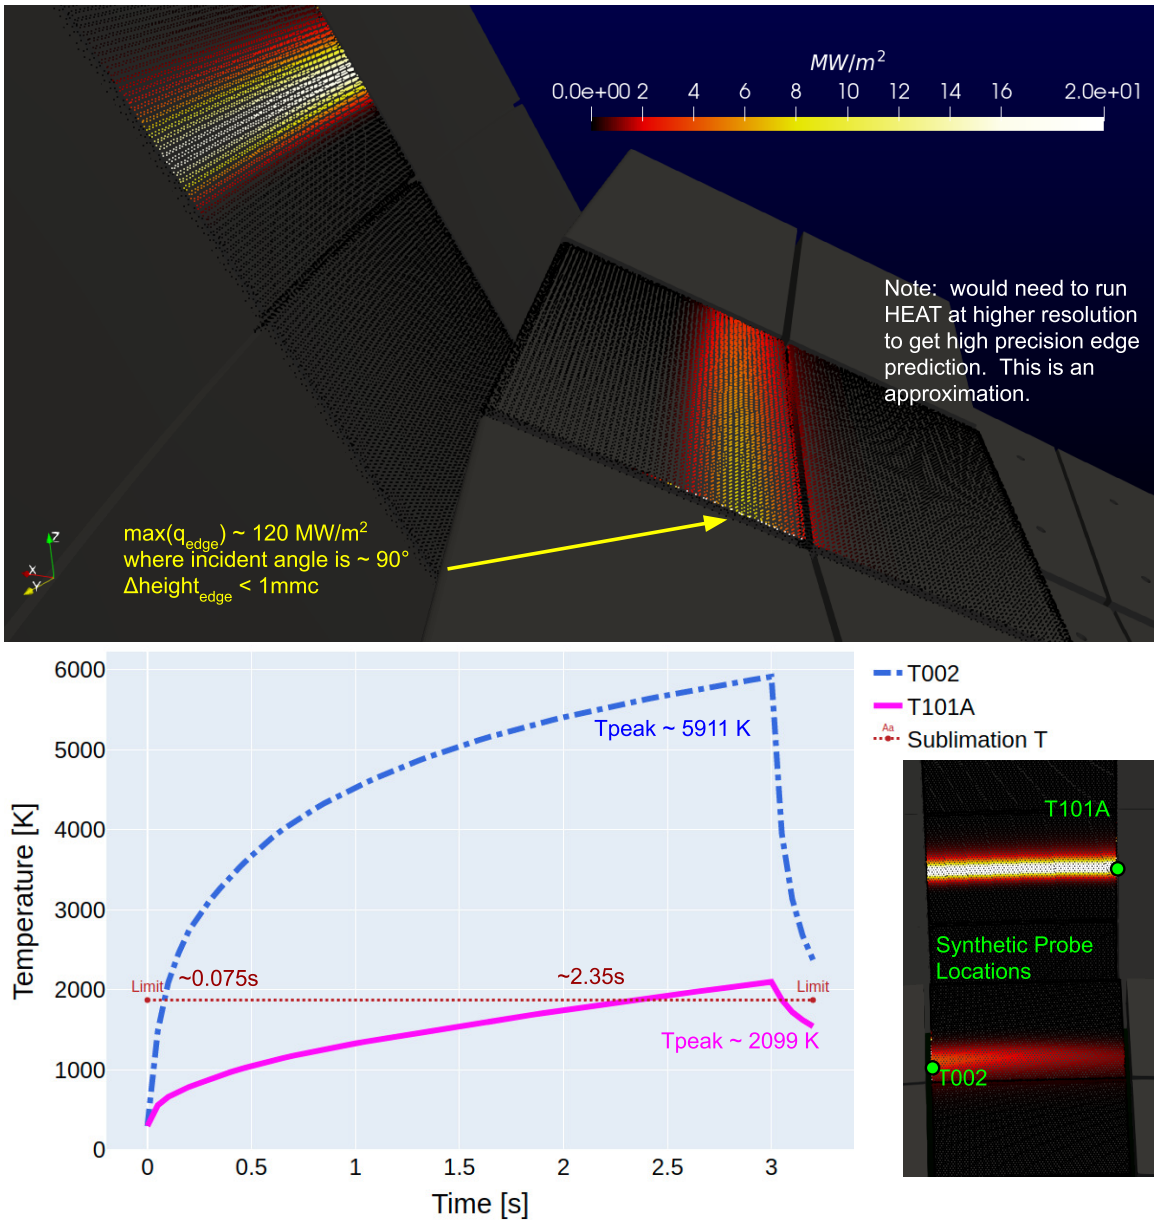


Figure 6.4: High X-point scenario heat flux as viewed from a different vantage point compared to figure 6.3. Top frame illustrates the consequences of no toroidal shaping on these PFCs, where the tile leading edge is exposed to heat fluxes at near normal angles of incidence. Bottom frame shows temperature evolutions for synthetic HEAT probes placed at the leading edges. For both the inner and outer strike point, the engineering limit (1873K) of graphite is reached. Note that extreme temperature for T002A probe is likely higher than reality, a product of the coarse volume mesh chosen for these preliminary workflow scoping studies. Further analysis is required to validate this result.

called out in the top frame of figure 6.4. On the corner, the angle of incidence is nearly 90 degrees, and the peak heat load is approximately 120 $\mathrm{MW\,m^{-2}}$. The width of this leading edge is less than 1mm, which is relatively narrow compared to the height of the PFCs, but the high heat flux magnitude easily pushes the PFCs beyond an example engineering limit of graphite, 1873K.

The lower frame of 6.4 provides time evolution traces for two synthetic probes placed at the leading edges, as well as a depiction of the probe locations overlaid onto the geometry. The inner target probe (in the figure labeled T002) crosses the engineering limit almost immediately. This is because the high flux expansion at this target results in a wider leading edge exposed to the 120 $\mathrm{MW\,m^{-2}}$ heat flux. Note that extreme temperature for T002A probe is likely higher than reality, a product of the coarse volume mesh chosen for these preliminary workflow scoping studies. Further analysis is required to validate this result. The outer target probe (in the figure labeled T101A), does not cross the engineering limit until 2.35s. This is because the outer target has a lesser flux expansion, thereby the leading edge is narrower. It should also be noted that these HEAT runs were not high resolution (ie <1mm) so the leading edge heat loads may have some numerical / mesh artifacts.

The second magnetic configuration investigated with HEAT for the DIII-D negative triangularity modeling effort was a 'low X-point' scenario. The MHD equilibrium, poloidal flux surfaces, and heat flux predictions for this configuration is provided in figure 6.5. The strike point locations for this scenario are similar to those in the high X-point scenario, but the flux expansion is greater for both targets. The lower frame in figure 6.5 shows that the peak heat flux on the top PFC surface is approximately 10.2 $\mathrm{MW\,m^{-2}}$ at the outer strike point and 5.15 $\mathrm{MW\,m^{-2}}$ at the inner strike point. The inner strike point has a significant amount of flux expansion, and the heat flux covers nearly both of the PFCs in the inner divertor. Additionally, at the inner strike point an interesting magnetic shadow appears at the upstream corner of the PFC. This magnetic shadow is due to the self-shadowing effect, where a flat PFC has a shadowed region because the spatial extent of its flat surface is long enough that $\hat{b}\cdot\hat{n}$ (see equation 3.26) changes sign across the surface. The inner target also exhibits a strong heat flux faceting effect, which can be observed in the gradient of the heat flux in the toroidal direction.

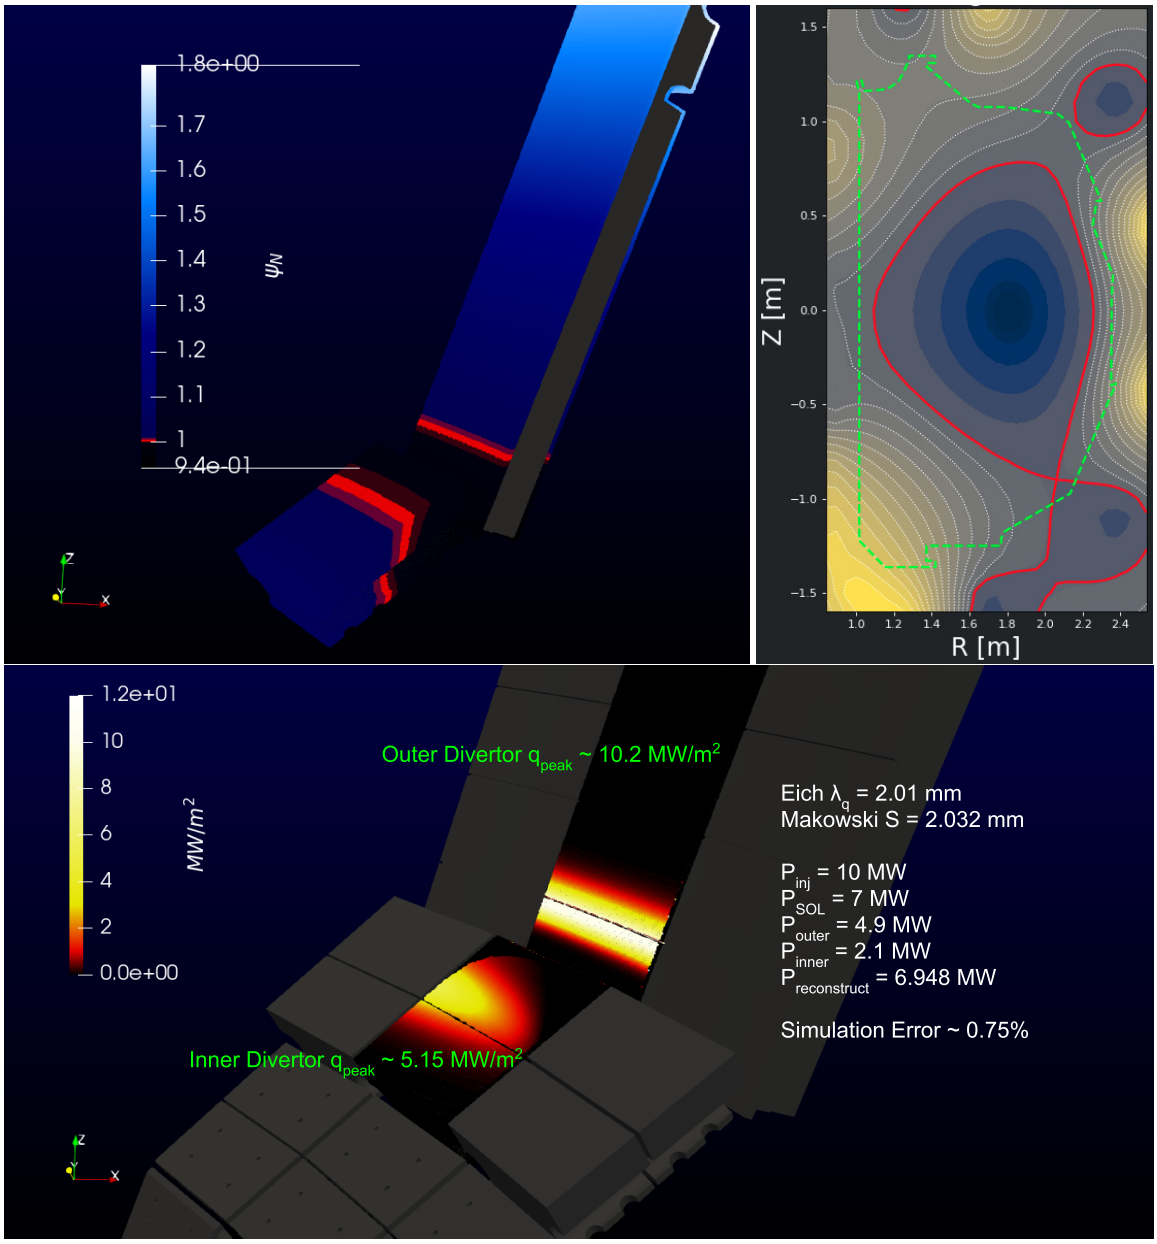


Figure 6.5: Normalized poloidal flux (top left) and MHD equilibrium (top right) for the low X-point negative triangularity scenario in DIII-D. Outer strike point is on wall PFCs, which were not designed to handle high heat loads. Bottom frame shows heat flux that corresponds to this magnetic equilibrium and geometry. Plasma parameters are overlaid onto the figure. This scenario has extremely high flux expansion, which results in the self shadowing effect seen at the inner divertor.

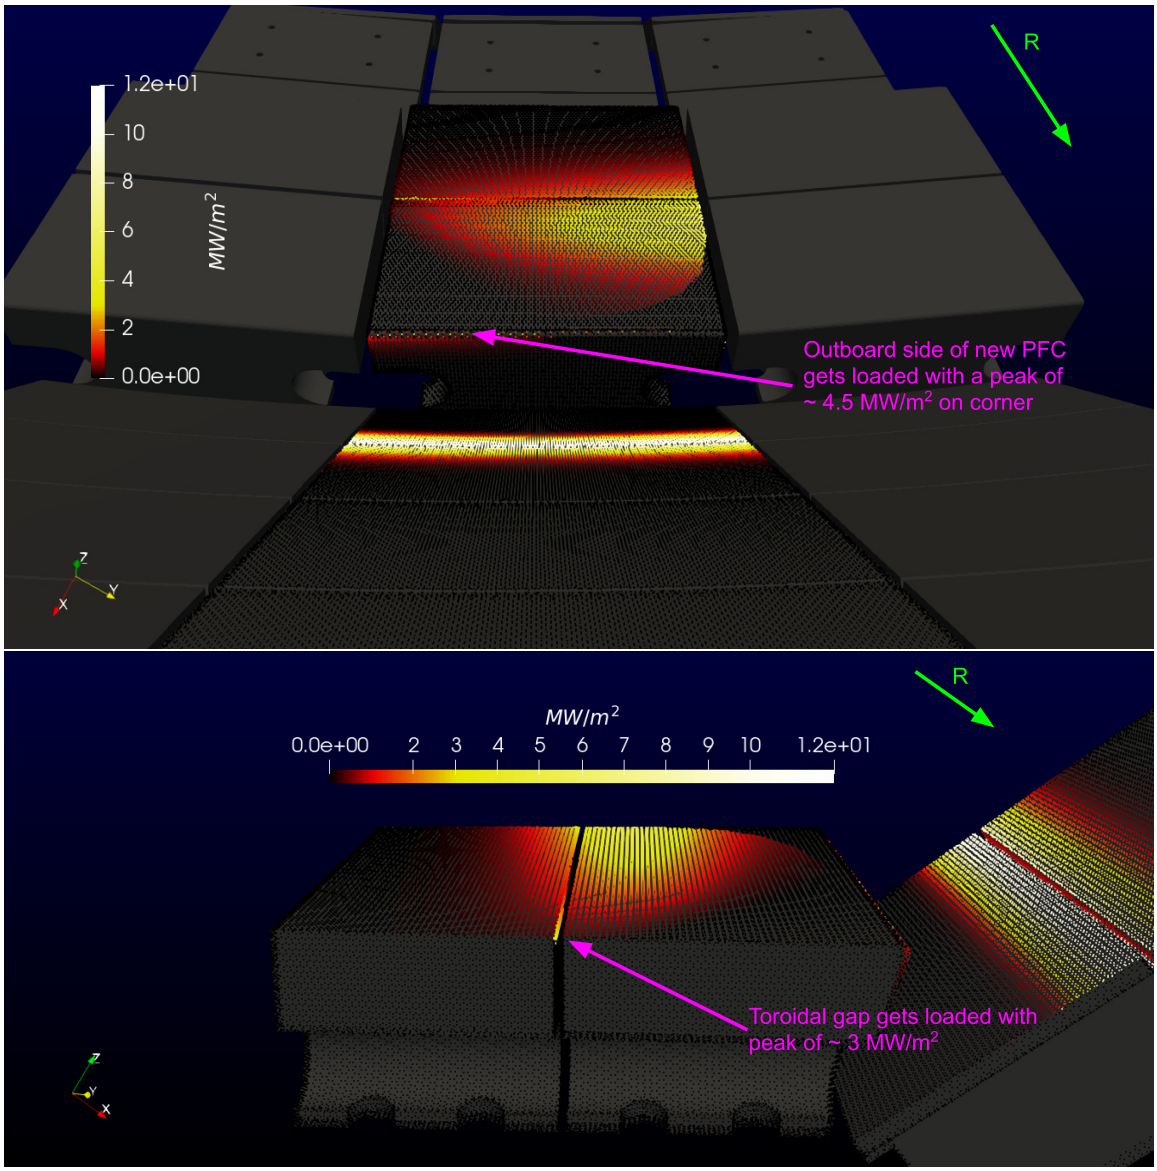


Figure 6.6: Views of the low X-point scenario from different vantage points from 6.5. Due to the high flux expansion of this discharge, the inner strike point sees heat loads on the back side of the PFC, as shown in top frame. Additionally, the rectangular PFC structure creates a faceted optical hot spot inside the gap.

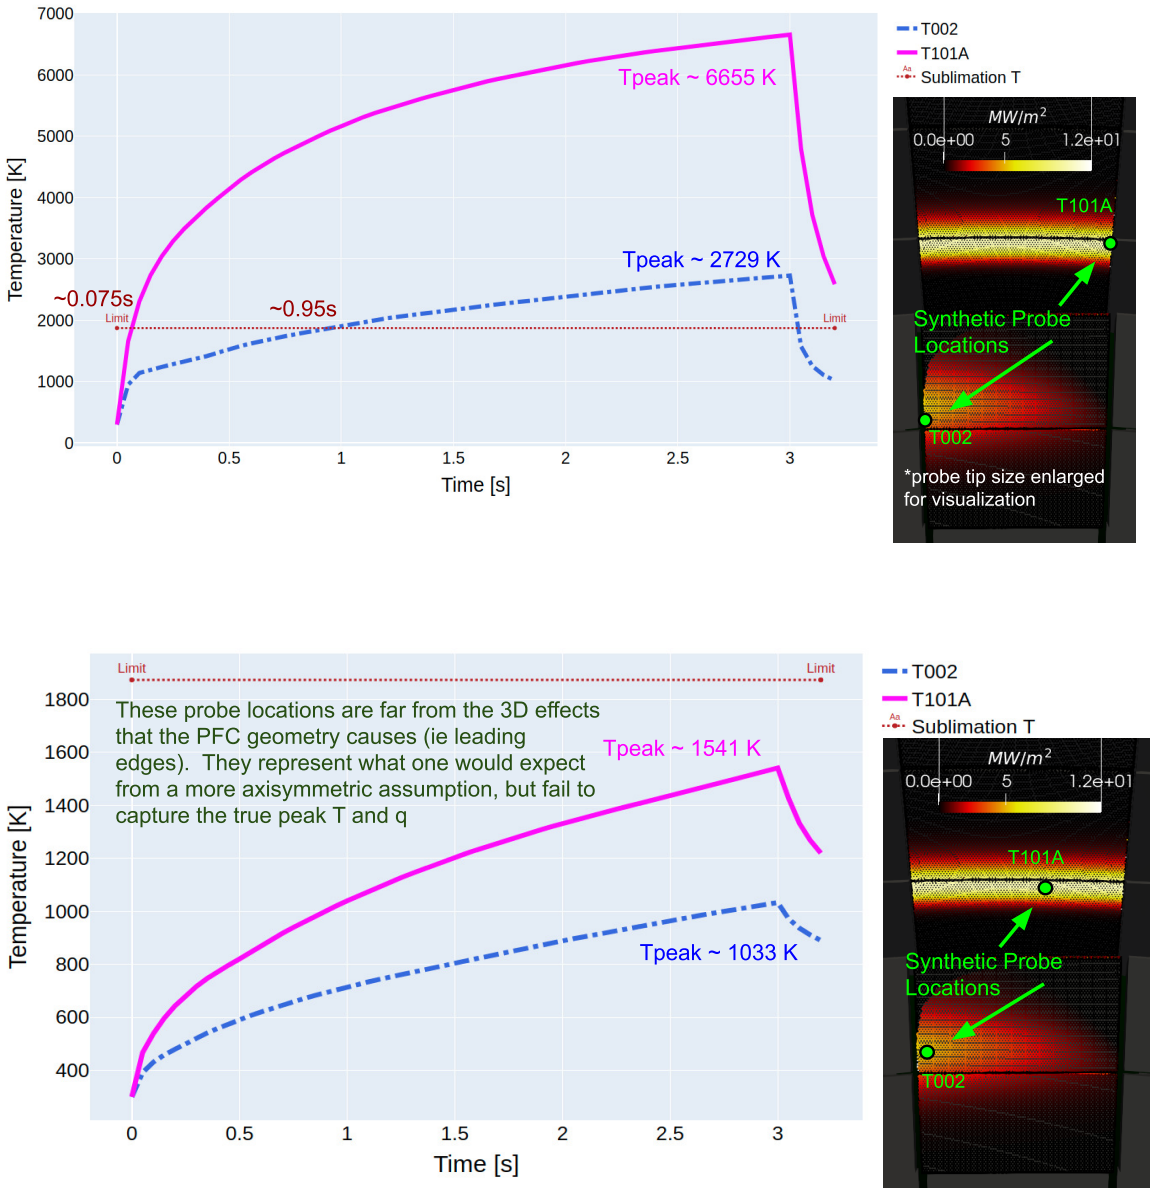


Figure 6.7: Temperature evolution for the low X-point PFCs, for two thermal probe locations. The top frame represents the temperature evolution in the leading edges, while the bottom frame represents the temperature evolution far from the leading edges. Both of the leading edges reach the engineering limit of graphite (1873K), while the bulk tile (closer to the axisymmetric assumption) maintains temperatures below the engineering limit. Note that extreme temperature for T002A probe is likely higher than reality, a product of the coarse volume mesh chosen for these preliminary workflow scoping studies. Further analysis is required to validate this result.

Figure 6.6 provides two different views of the heat flux associated with the configuration provided in figure 6.5. The top frame shows a vantage point looking from the outboard side of the PFCs looking in towards to center stack. The outboard edge of the inner divertor is loaded with heat fluxes up to 5 $\mathrm{MW\,m^{-2}}$. This poloidally-facing leading edge is an optical hot spot, that appears as a result of the heat flux faceting effect. This faceting effect occurs because the flat PFC surface is not toroidally symmetric with respect to the magnetic field lines, which results in a changing $\hat{b}\cdot\hat{n}$ across the PFC surface. The poloidally-facing leading edge inside the toroidally-running gap between the two divertors in the inner divertor also shows evidence of an optical hot spot due to heat flux faceting. Because this leading edge is inside a gap, some of the power is scraped off by the poloidally upstream PFC, so the heat fluxes are lesser than the aforementioned poloidally-facing leading edge.

Figure 6.7 provides the temperature evolution for two probes located in the divertor for the low X-point configuration. The top frame illustrates the temperature evolution for the leading edges in the plot, and the probe locations in the image. Note that extreme temperature for T002A probe is likely higher than reality, a product of the coarse volume mesh chosen for these preliminary workflow scoping studies. Further analysis is required to validate this result. When compared to the high X-point scenario, the peak edge temperature for this scenario is higher for both divertors, as the larger flux expansion leads to a wider leading edge. Both the inner and outer divertor leading edges exceed the example engineering limit of graphite overlaid on the plot. In the lower frame of figure 6.7 the probes are placed at locations (relatively) far from the leading edges, but still within the peak of the heat flux faceting effect. The probe locations are provided visually in the image to the right of the figure. Placing the probe at a location that is far from the leading edges (ie in the middle of the PFC) provides a temperature evolution that is more aligned with an axisymmetric assumption, as the 3D geometry of the PFCs plays the largest role at the edges, corner, and gaps between tiles. Using this probe location results in lower temperatures, as can be observed in the plot in the lower frame of figure 6.7. The peak temperature on the PFC edges (top frame) is approximately 4.3 X higher than the the peak temperature far from the PFC edge (bottom frame). Even though the DIII-D divertor is designed without toroidal

shaping, the 3D effects of the PFCs and inter-PFC gaps results in significant deviations from the axisymmetric heat loads.

The simulations presented in this section correspond to early CAD designs of the DIII-D negative triangularity campaign. Subsequent work on the negative triangularity CAD design was completed by Dr. Filippo Scotti (LLNL) who is an independent HEAT user. The CAD geometry was revised to address the issues outlined in this section.

6.3 DIII-D DiMES Head

Another HEAT collaboration on DIII-D involved creating predictions for a Divertor Material Evaluation System (DiMES) head that contained a retarding field energy analyzer (RFEA). This work is being led by Dr. Jun Ren (UTK), in collaboration with the DIII-D team.

The top frame of figure 6.8 illustrate an early draft of the DiMES head CAD design, as well as the RFEA insert, The HEAT predicted optical heat flux is overlaid onto the CAD geometry for both the DiMES head and the RFEA insert. The edge of this DiMES head version was loaded with high heat fluxes; later versions remedied this with a modified chamfer that was slightly rotated. More importantly, the slit in the RFEA insert is observed to be completely magnetically shadowed from incident heat fluxes. The DiMES head in this work is designed to measure the energy of incoming particles and the associated heat fluxes by allowing some particles through this slit, so it is imperative that the slit not be shadowed from incident particle fluxes. Later versions modified the design to remedy this shadow, as depicted by the lower slide in figure 6.8. Dr. Ren, who is an independent user of HEAT, iteratively ran the simulations in the lower slide to find the optimum DiMES head and RFEA insert CAD design for his objective. This example demonstrates how HEAT can be utilized for design optimization of divertor diagnostics, as well as how HEAT can be used by third party users.

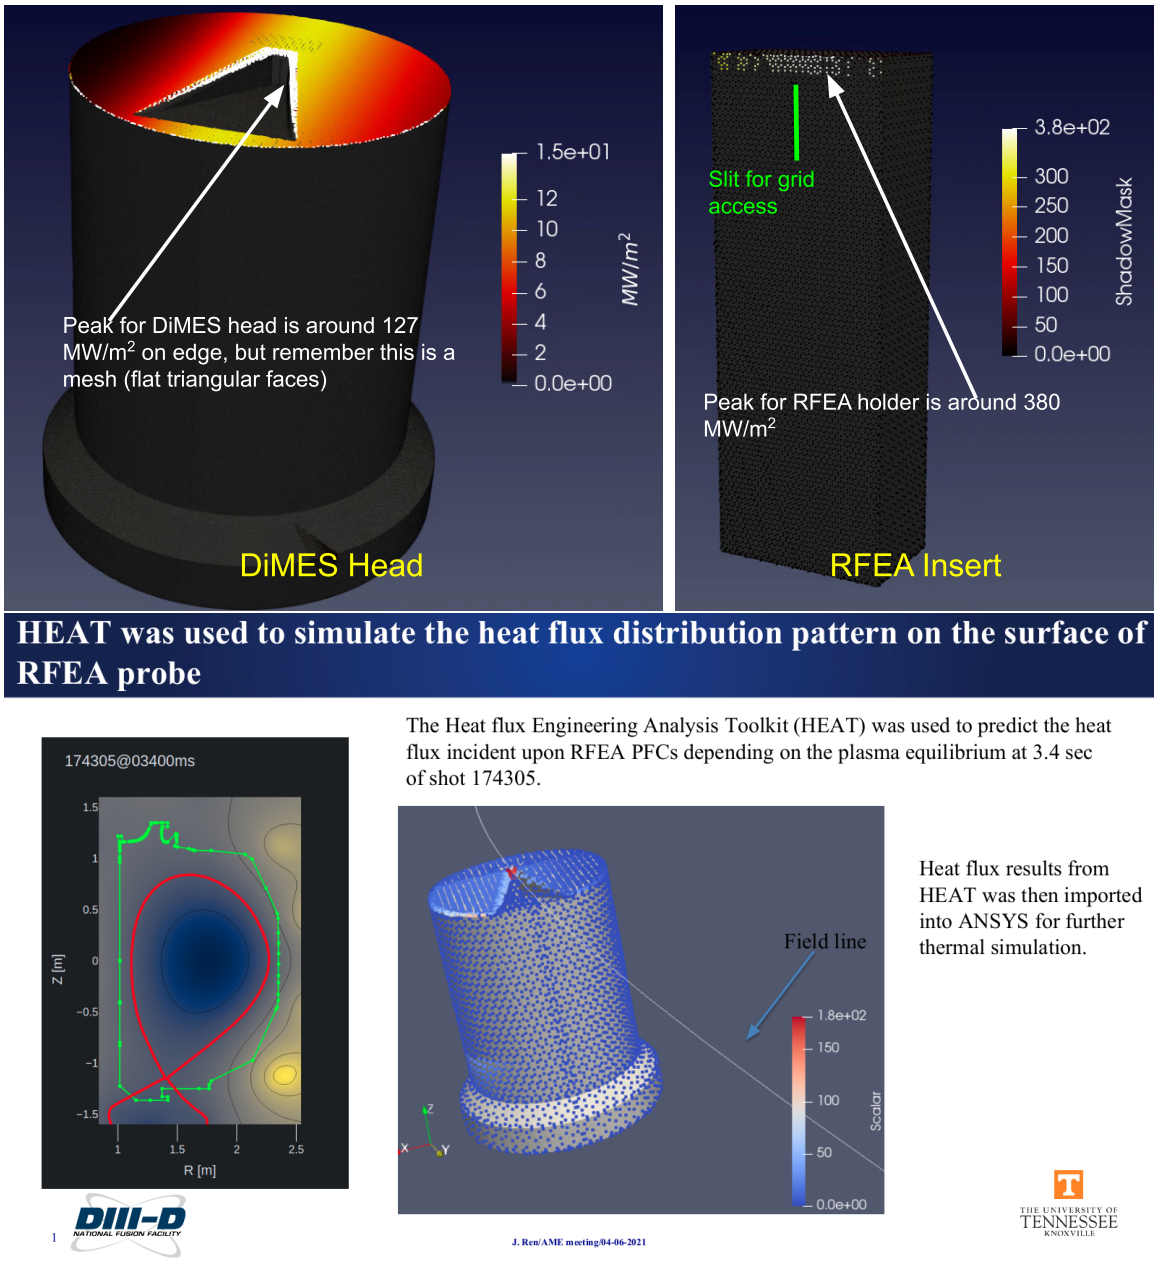


Figure 6.8: Divertor Material Evaluation System (DiMES) head and retarding field energy analyzer (RFEA) insert for DIII-D. Top frame shows HEAT calculated heat flux for an initial (not final) design, where the access slit for ions was magnetically shadowed. Final design changed geometry so that access slit was not shadowed. Bottom slide, courtesy of Dr. Jun Ren, shows example HEAT output field line trace and temperature output. HEAT was used by Ren to iterate on optimal DiMES head design.

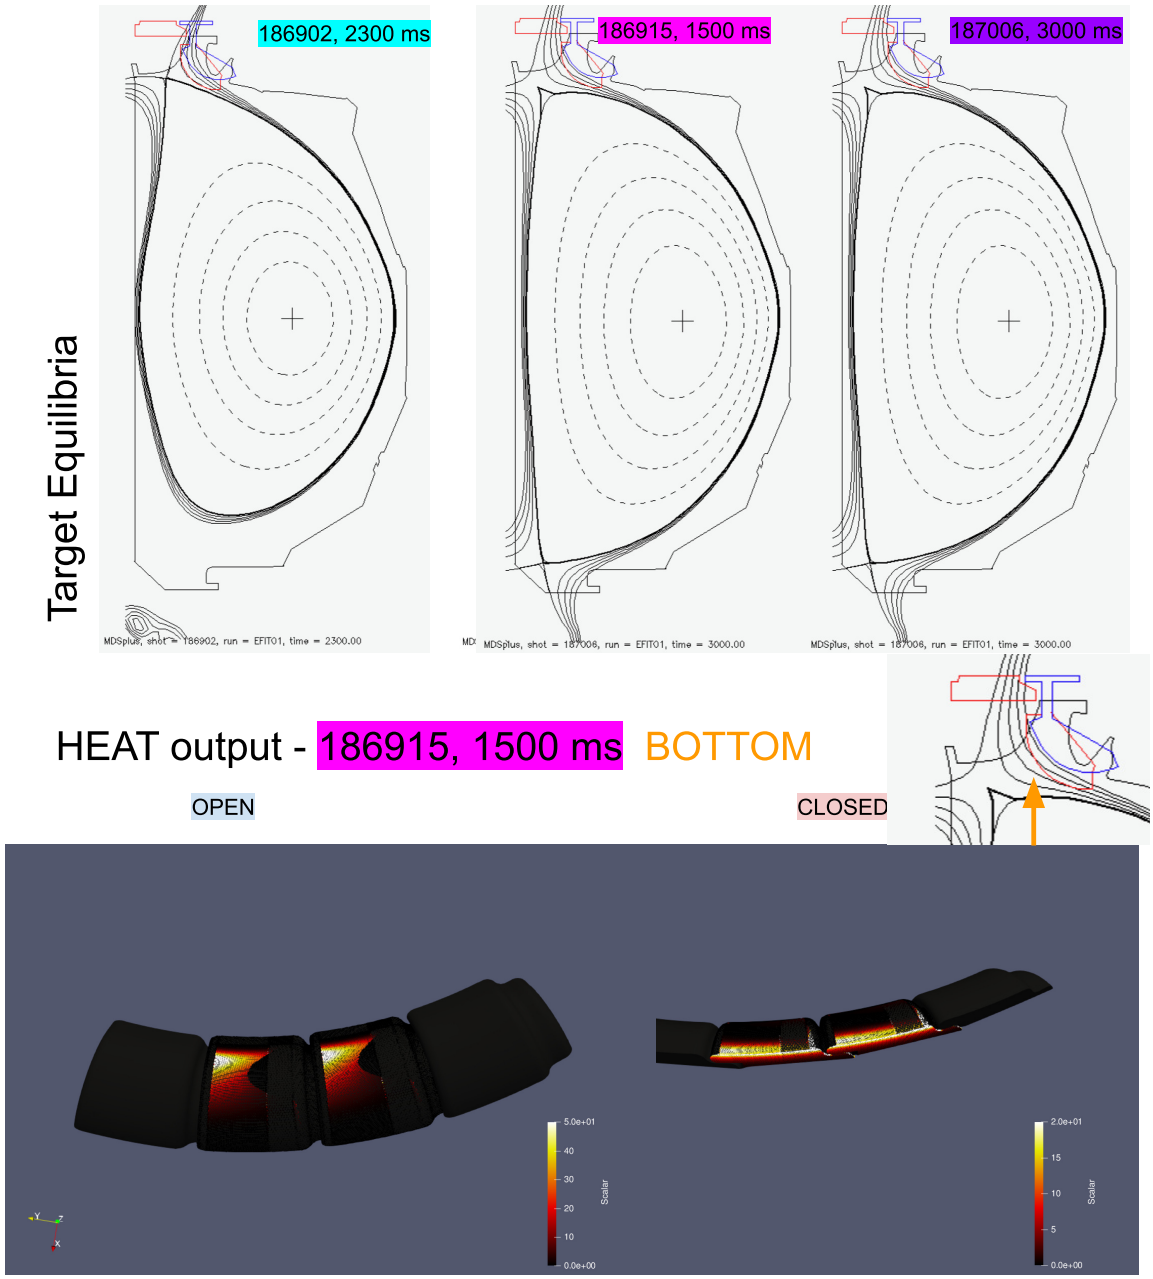


Figure 6.9: DIII-D movable divertor analysis courtesy of Dr. Andreas Wingen (ORNL). The movable divertor can operate in an open or closed configuration, which is controlled by actuators. Results here are slides from Dr. Wingen's analysis, showing three potential magnetic equilibria (top), as well as the HEAT output for the center equilibrium (186915).

6.4 DIII-D Moveable Divertor

The DIII-D tokamak team is in the early design phase for a movable divertor, and has utilized HEAT to generate some predictions for these designs. This divertor will have actuators capable of rotating the PFCs into two configuration: open and closed. Figure 6.9 is comprised of two slides provided courtesy of Dr. Andreas Wingen (ORNL), who is the PI responsible for running HEAT for the movable divertor. Three MHD equilibria are provided in the top frame of the figure, and HEAT output for shot 186915 is provided in the lower frame. When the divertor PFCs are in the open configuration a large magnetic shadow extends into the center of each PFC, as a result of the upstream PFCs curvature. Additionally, there is a strong faceting effect observed when the divertor is in the open configuration, which sees a peak heat flux in excess of 50 $\mathrm{MW\,m^{-2}}$. When closed, the heat loads are more concentrated in a single poloidal location, but across the entire toroidal extent of the PFC. This work is ongoing at DIII-D, and the mention in this text is not meant to be quantitative analysis, but instead is provided to illustrate the wide variety of HEAT use cases.

6.5 STEP / DEMO

The EU-DEMO team has been working on initial design studies for the Spherical Tokamak for Energy Production (STEP). This tokamak is primarily in the concept stage, and some team members from EU-DEMO have been working on its design. During some of these early design studies, HEAT was evaluated as a potential candidate to use for predicting the 3D heat loads in STEP. The work presented here uses CAD and MHD equilibrium from DEMO, but the intention was to transition to STEP CAD and MHD EQ after the initial workflow was established. Because the DEMO team has already used SMITER (see section 3.2) for heat flux predictions, this enabled a side by side comparison between the SMITER and HEAT output. This analysis was not extremely involved, but there is general agreement between the two codes.

Figure 6.10 provides visualization of the DEMO geometry in the upper frame. A toroidal sector of 22.5 degrees is show in the upper left and the divertor is shown in the upper right. In the lower frame the heat fluxes for the divertor are provided. The real

engineering CAD for STEP / DEMO has not been designed yet, so the CAD included for the divertor is created from the EFIT (rlim,zlim) wall. HEAT contains tools to build STEP and STL files directly from the EFIT, which were used here. The lower left frame of figure 6.10 provides a snapshot of the divertor heatflux for a heat flux width of 50 mm, which was requested by the STEP team. The lower right frame shows a zoomed in view of these heat fluxes in the divertor. The numbers for the heat flux width (λ_q) and the power sharing ratios were provided by the STEP team, rather than taken from any empirical regression or model. For such a large heat flux width the peak heat flux is a mere 2 $\mathrm{MW\,m^{-2}}$.

The HEAT output for the inner limiter of DEMO is shown in figure 6.11. The limiter CAD is only surface features (ie no volume or part objects), but was provided by the STEP team. The top left image provides a view of the toroidal chamfer included in each limiter. The top center frame provides the field line incident angle ($\hat{b}\cdot\hat{n}$), which changes sign across the chamfer. The top right frame provides the shadowMask for this limiter design, and the regions that are heat loaded are indicated in white. The toroidal chamfer serves to protect the limiter in both toroidal directions. The bottom frame in figure 6.11 shows the HEAT heat flux predictions for the inner limiter and divertor. The heat loaded regions correspond exactly to the shadowMask results, and the divertor heat fluxes peak at 1.22 $\mathrm{MW\,m^{-2}}$ for this high value of λ_q (50mm per STEP team).

The predictions provided for STEP / DEMO in this section are primarily included to demonstrate that HEAT can be ported to machines that are far outside of the sizes and shapes of traditional tokamaks. The size of DEMO makes heat flux calculations on a mesh extremely time consuming, but the HEAT acceleration structure make the computation relatively fast. The results that were provided should not be taken as ground truth, as the inputs for heat flux width, power sharing, etc., are unknown. An example comparison between SMARDDA, which is the package under the hood of the SMITER code, and HEAT is provided in figure 6.12. The SMARDDA output is courtesy of Dr. Lorena Richiusa (UKAEA) The two codes show nearly identical heat flux footprints, but the magnitudes of the heat fluxes on the leading edges are different. Further analysis would be needed to further benchmark these two codes and ascertain the source of this magnitude difference, but the initial workflows are now established

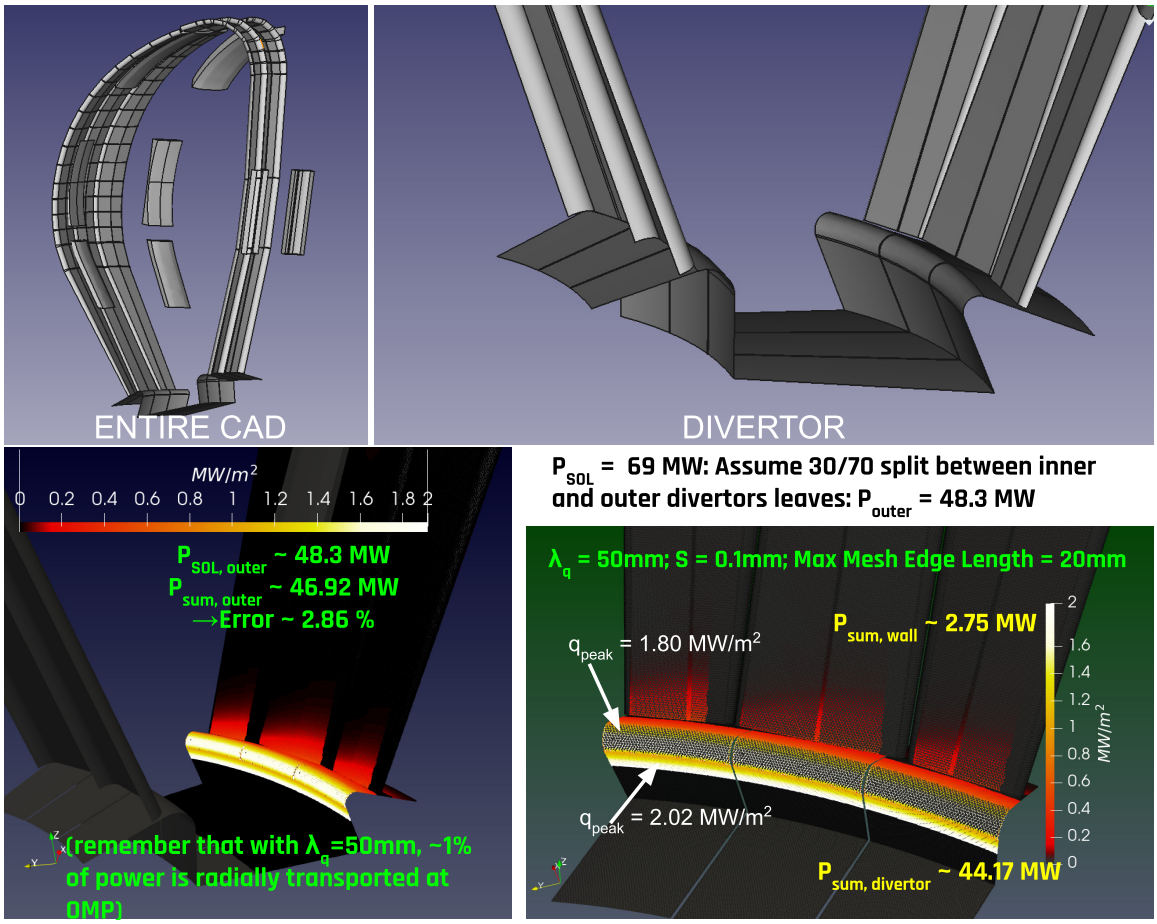


Figure 6.10: STEP / DEMO CAD (top panels) and HEAT simulated divertor heat flux (lower panels). Input parameters provided by STEP team.

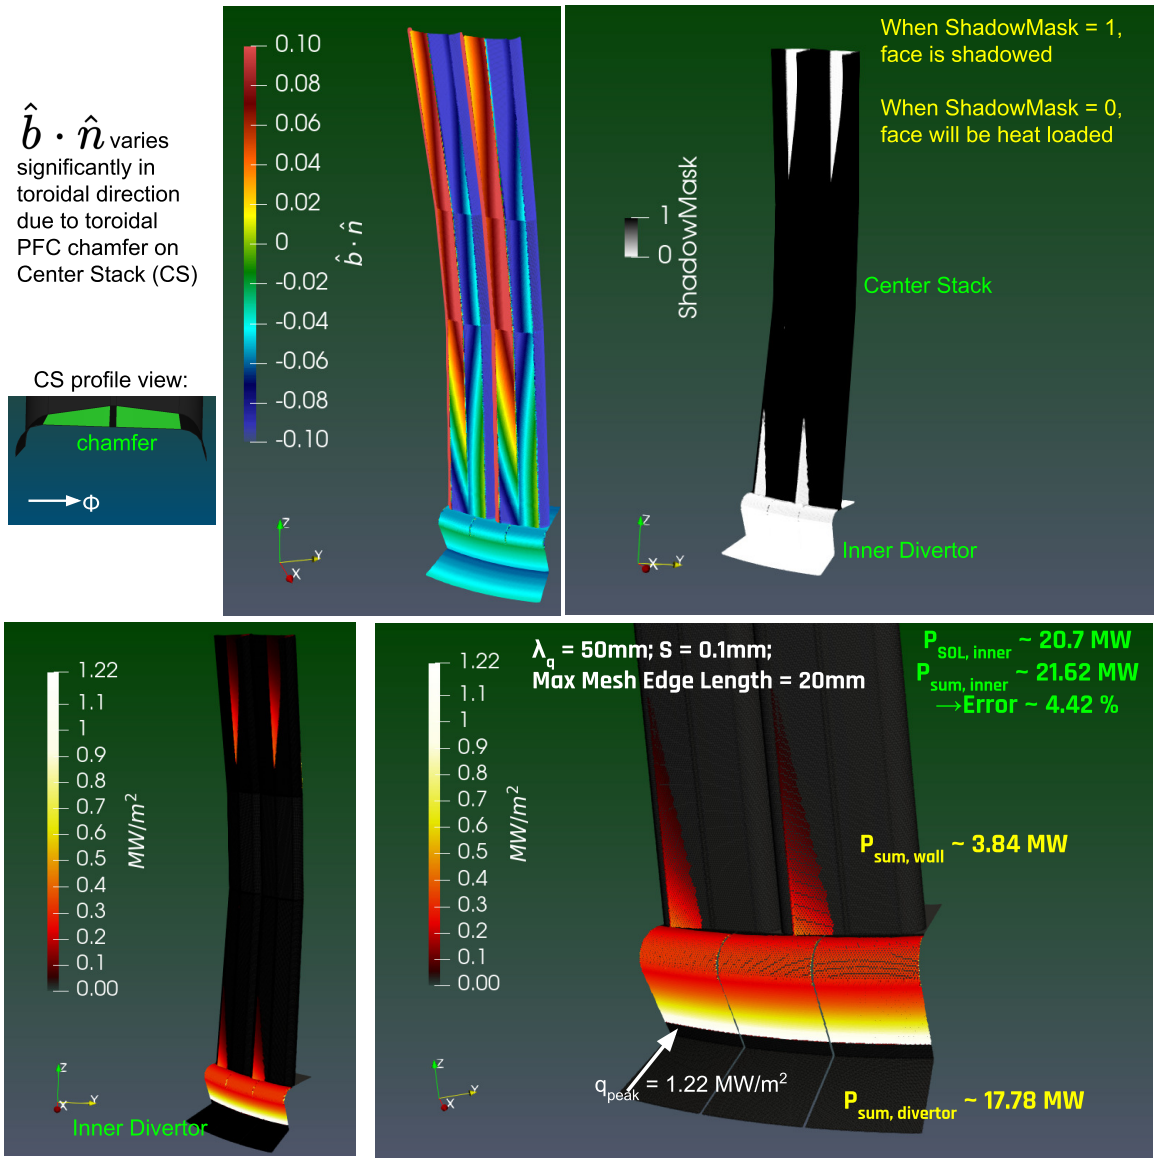


Figure 6.11: Example HEAT output for STEP inner limiter. Limiter toroidal chamfer (top left), magnetic field line incident angle (top center) and shadowMask (top right). Bottom panels show heat flux predictions using input parameters provided by STEP team.

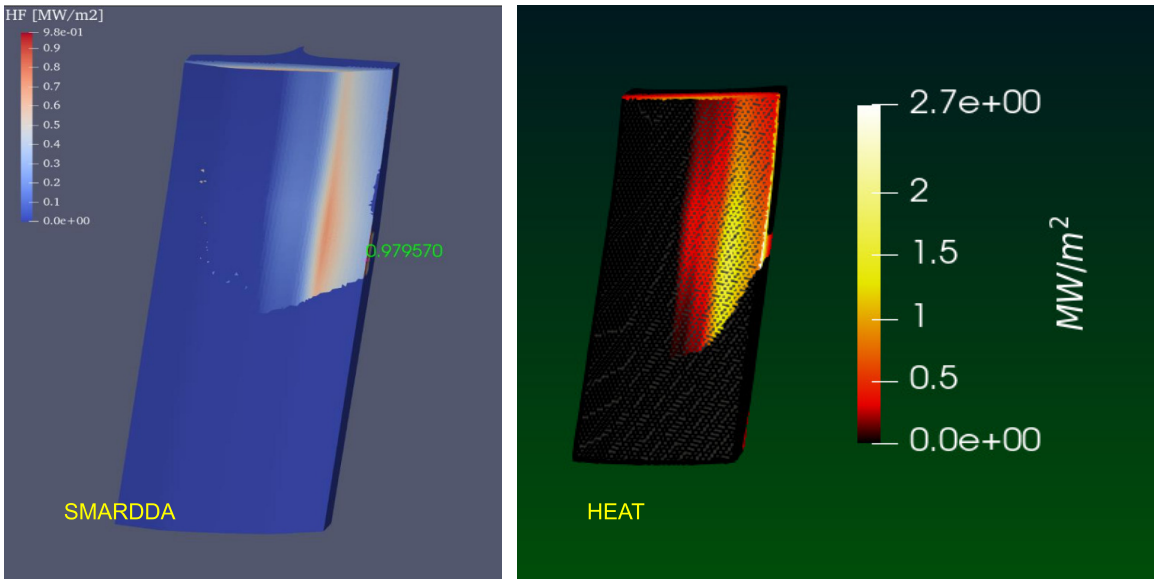


Figure 6.12: Comparison of heat flux prediction provided by SMARDDA (left panel) and HEAT (right panel) for the outboard limiter in STEP. Heat flux footprints are similar but magnitudes differ. SMARDDA output courtesy of Dr. Lorena Richiusa (UKAEA).

to make this comparison. It should be noted that Dario Vaccaro (UKAEA) is a HEAT user or STEP, and has recently been creating HEAT predictions directly for STEP (ie not DEMO).

6.6 SPARC

HEAT has been utilized to predict heat loads for the net energy gain SPARC tokamak being developed by Commonwealth Fusion Systems. Because the SPARC divertor is a reactor relevant divertor, it is considered intellectual property and cannot be shared in this dissertation. No results for SPARC will be included herein, but the work has been extensive. This subsection was included simply to mention that there is ongoing HEAT work at SPARC.

Chapter 7

Conclusions

This dissertation provided an investigation into the nature of 3D heat loads in tokamak divertors. In order to predict the heat loads observed in a tokamak, a complex simulation framework that connects plasma physics to the 3D engineering geometry was developed: the Heat flux Engineering Analysis Toolkit (HEAT). Not only is the toolkit novel in its ability to perform time varying 3D heat load analysis, but many of the underlying algorithms are also novel. The toolkit has been utilized to investigate a wide variety of physical phenomena in five tokamaks: NSTX-U, DIII-D, ST40, STEP, and SPARC. Many of these investigations provided new insights into the complex interaction between divertor plasmas and PFC geometry that can only be understood through integrated modeling frameworks. To conclude the dissertation, this section provides a synopsis of the need for 3D heat load predictions, HEAT, and the findings the toolkit has yielded.

7.1 Summary of the Problem

The PFCs in tokamaks will see some of the most incredible heat loads of any engineering domain. As such, it is imperative that these heat loads are characterized, quantified, and investigated, as discussed in chapter 2. The PFCs serve as the interface between most of the systems in the tokamak plant and the plasma. Because of the complex nature of the tokamak system, changes to interfaced subsystems that are seemingly negligible may directly impact the PFCs. Poor heat load management in high power tokamaks can result in performance reductions, or in catastrophic PFC failure which

(in the worst cases) can affect machine operation. These consequences demand precise engineering techniques as a means of ensuring that the PFCs remain within the allowable engineering domain. The two philosophies for controlling the heat loads are a) changing the plasma, or b) changing the PFCs. In most high power machines a combination of a) and b) are incorporated.

Modern tokamak designs employ 3D shaping in order to protect the PFCs from high heat and particle fluxes. These 3D shapes result in complex interactions between plasma heat fluxes that are transported down the Scrape Off Layer (SOL) and the PFC geometry. Bevels, chamfers, fish-scales, and castellations, are all common methods used to prevent the PFCs from becoming overheated. Additionally, plasma methods such as increasing the flux expansion and sweeping the strike point also may decrease the heat loads. Because changes in the plasma can have very different heat flux outcomes depending upon the PFC geometry, simulating both the plasma and the PFC geometry simultaneously is mandatory. Up until now, no such tools existed that could couple time varying plasmas directly to the PFCs for time varying thermal analyses. The HEAT code was developed during the course of this dissertation to fill this gap.

7.2 Summary of the Tool

The Heat flux Engineering Analysis Toolkit (HEAT) was developed to simulate time varying 3D heat loads on PFCs, as described in chapter 3. HEAT couples magnetohydrodynamic (MHD) equilibria directly to the tokamak engineering CAD. To do this HEAT generates mesh objects and calculates quantities of interest (i.e. poloidal flux, $q_{\perp}$, etc.) at each mesh element in the CAD. In order to determine if a mesh element is shadowed, HEAT traces magnetic field lines upstream from the mesh elements and checks for intersections with other mesh objects along the way. The algorithms used for this process are adapted versions of ray-triangle intersection checking algorithms, which are novel in their vectorization and parallelization. Additionally, coordinate transformation mesh filter algorithms were invented as acceleration structures for the HEAT calculation.

Optical power is calculated by mapping each mesh element onto a corresponding poloidal flux surface, and then assigning power using empirical regressions. Gyro

orbit power is calculated by sampling from the velocity space distribution functions, tracing the helical trajectories of macro-particles onto the mesh elements, and then redistributing power defined in optical empirical regressions. Power balance is used as a primary method of validating the HEAT output. Temperature is calculated by using the calculated heat fluxes as boundary conditions on a finite volume mesh. The finite volume mesh is generated using internal HEAT algorithms that snap a volume mesh to the surface mesh. The time evolution of temperature can be calculated as a macroscopic quantity across the entire PFC volume, or by specifying the location of an idealized temperature probe. The thermal properties are all temperature dependent.

Many other helper functions are available in the toolkit. The user may stitch together magnetic equilibria into strike point sweeps, redefine the last closed flux surfaces to be tangential to limiters, and redefine other magnetic equilibria parameters directly in the HEAT user interface. HEAT can be used to trace magnetic field lines and gyro orbit helical trajectories from any user defined location. HEAT generates output in csv format, as well as visualization output in the .vtk format which can be imported directly to paraview. HEAT embraces an open source philosophy, and the underlying software architecture is modular and hackable. Last, HEAT can be used by Linux users via an appImage, and by macOS and windows users via virtual machines or over a network connection to a HEAT server via a web browser.

7.3 Summary of the Results

HEAT predictions couple engineering CAD directly to time varying magnetic equilibria. The optical approximation is generally found to be the most conservative prediction with respect to the PFC engineering limits. The optical approximation NSTX-U results provided in chapter 4 illustrate how the heat loads incident upon the divertor can easily push the PFCs past their engineering limits if not properly managed. The 3D shaping employed to protect the PFCs results in complex heat flux footprints that can only be discovered by mapping the PFCs directly into flux coordinates via a field line tracer. Additionally, it is found that strike point sweeping can directly constrain the allowable discharge length, and that HEAT can be utilized as a transfer function that maps strike point sweep frequency directly to allowable discharge duration. Additionally,

HEAT can be used during machine construction to predict heat loads that arise from unintended misalignments or manufacturing errors in tangentially related systems.

The gyro-orbit approximation predictions provided in chapter 5 provide a first-of-its-kind time varying analysis by coupling helical charged particle trajectories directly to engineering CAD. These results are novel as they are connected directly to the 3D engineering CAD to generate 3D heat flux profiles from the gyro-orbit approximation. Other analyses of this type have included simplifications to the CAD by assuming the problem is 2D, by reducing the CAD model to simple geometries, or by eliminating the time varying nature of the discharge. The methods used in HEAT also are well suited to including sheath electric fields, although that is not included in this dissertation. The results indicate that gyro-orbit heat fluxes will deviate strongly from the optical heat fluxes where the CAD geometry includes features such as edges, gaps, and corners. There can be a significant heat flux enhancement associated with these features, but the difference in PFC thermal response for pulsed discharges is negligible. The helical trajectories of particles also serve to spread the heat loads into optically shadowed regions, thereby decreasing the PFC temperature rise because the power is spread over a larger surface area.

Chapter 6 provides several examples of utilizing HEAT to optimize CAD designs and magnetic scenarios, design diagnostics, and generate predictions for complex nonaxisymmetric divertors. Highlights across 4 additional tokamaks are provided to illustrate the general nature of HEAT and its models. Application of HEAT for a wide variety of use cases is the theme of these analyses.

Bibliography

[1] G. Gamow, "Quantum theory of the atomic nucleus," *ZP*, vol. 51, p. 204. 1

[2] J. D. Lawson, "Some criteria for a power producing thermonuclear reactor," *Proceedings of the Physical Society. Section B*, vol. 70, no. 1, pp. 6–10. 1

[3] J. Wesson, *Tokamaks*. Claredon Press, 3rd ed. 2

[4] P. C. Stangeby, *The plasma boundary of magnetic fusion devices*. Plasma physics series, Institute of Physics Pub. ix, xi, 2, 3, 4, 5, 40, 42, 45

[5] G. Federici, R. Kemp, D. Ward, C. Bachmann, T. Franke, S. Gonzalez, C. Lowry, M. Gadomska, J. Harman, B. Meszaros, C. Morlock, F. Romanelli, and R. Wenninger, "Overview of EU DEMO design and r&d activities," vol. 89, no. 7, pp. 882–889. 4

[6] M. Siccinio, G. Federici, R. Kembleton, H. Lux, F. Maviglia, and J. Morris, "Figure of merit for divertor protection in the preliminary design of the EU-DEMO reactor," *Nuclear Fusion*, vol. 59, p. 106026, Oct. 2019. 4, 69

[7] C. Linsmeier, M. Rieth, J. Aktaa, T. Chikada, A. Hoffmann, J. Hoffmann, A. Houben, H. Kurishita, X. Jin, M. Li, A. Litnovsky, S. Matsuo, A. von Müller, V. Nikolic, T. Palacios, R. Pippan, D. Qu, J. Reiser, J. Riesch, T. Shikama, R. Stieglitz, T. Weber, S. Wurster, J.-H. You, and Z. Zhou, "Development of advanced high heat flux and plasma-facing materials," vol. 57, no. 9, p. 092007. 4

[8] R. Pitts, X. Bonnin, F. Escourbiac, H. Frerichs, J. Gunn, T. Hirai, A. Kukushkin, E. Kaveeva, M. Miller, D. Moulton, V. Rozhansky, I. Senichenkov, E. Sytova, O. Schmitz, P. Stangeby, G. De Temmerman, I. Veselova, and S. Wiesen, "Physics

basis for the first ITER tungsten divertor," *Nuclear Materials and Energy*, vol. 20, p. 100696, Aug. 2019. xi, 4, 51, 53, 54

[9] A. Q. Kuang, S. Ballinger, D. Brunner, J. Canik, A. J. Creely, T. Gray, M. Greenwald, J. W. Hughes, J. Irby, B. LaBombard, B. Lipschultz, J. D. Lore, M. L. Reinke, J. L. Terry, M. Umansky, D. G. Whyte, S. Wukitch, and the SPARC Team, "Divertor heat flux challenge and mitigation in SPARC," *Journal of Plasma Physics*, vol. 86, p. 865860505, Oct. 2020. xi, 4, 50, 53, 55

[10] T. K. Gray, N. Allen, M. L. Reinke, G. Smalley, D. L. Youchison, R. Ellis, M. A. Jaworski, T. Looby, M. Mardenfeld, and D. E. Wolfe, "Integrated plasma facing component calorimetry for measurement of shot integrated deposited energy in the NSTX-U," *Rev. Sci. Instrum.*, p. 6, 2018. 6, 62

[11] J. Menard, "NSTX-U PFC Working Group Memo 004." https://nstx.pppl.gov/DragNDrop/Working_Groups/PFCR/memos/PFCR-MEMO-004-01.pdf. Accessed: 2019-08-14. 6

[12] R. Goldston, "Heuristic drift-based model of the power scrape-off width in low-gas-puff h-mode tokamaks," vol. 52, no. 1, p. 013009. 8, 9

[13] W. M. Stacey, *Fusion Plasma Physics*. Wiley, 1 ed. 8

[14] T. Eich, B. Sieglin, A. Scarabosio, W. Fundamenski, R. J. Goldston, A. Herrmann, and ASDEX Upgrade Team, "Inter-ELM power decay length for JET and ASDEX upgrade: Measurement and comparison with heuristic drift-based model," vol. 107, no. 21, p. 215001. 9, 12, 53

[15] T. Eich, A. W. Leonard, R. A. Pitts, W. Fundamenski, R. J. Goldston, T. K. Gray, A. Herrmann, A. Kirk, A. Kallenbach, O. Kardaun, A. S. Kukushkin, B. LaBombard, R. Maingi, M. A. Makowski, A. Scarabosio, B. Sieglin, J. Terry, A. Thornton, and E. Contributors, "Scaling of the tokamak near the scrape-off layer H-mode power width and implications for ITER," *Nucl. Fusion*, p. 8, 2013. ix, 9, 10, 11, 12, 102, 132, 183

[16] D. Brunner, B. LaBombard, A. Kuang, and J. Terry, "High-resolution heat flux width measurements at reactor-level magnetic fields and observation of a unified

width scaling across confinement regimes in the Alcator C-Mod tokamak," *Nuclear Fusion*, vol. 58, p. 094002, Sept. 2018. ix, 12, 13, 102

[17] J. P. Gunn, S. Carpentier-Chouchana, F. Escourbiac, T. Hirai, S. Panayotis, R. A. Pitts, Y. Corre, R. Dejarnac, M. Firdaouss, M. Ko, W. Zhao, and G. Zhong, "Surface heat loads on the ITER divertor vertical targets," *Nucl. Fusion*, p. 36, 2017. ix, x, xvii, xviii, 15, 16, 17, 18, 20, 21, 148, 149, 151, 154, 169

[18] M. Komm, J. Gunn, R. Dejarnac, R. Pánek, R. Pitts, and A. Podolník, "Particle-in-cell simulations of the plasma interaction with poloidal gaps in the ITER divertor outer vertical target," *Nuclear Fusion*, vol. 57, p. 126047, Dec. 2017. x, xvii, xviii, 16, 20, 21, 22, 23, 24, 148, 150, 151, 154, 169

[19] W. Hu, C. Sang, Z. Sun, and D. Wang, "Effects of the divertor tile geometries and magnetic field angles on the heat fluxes to the surface," vol. 116, pp. 5–9. 16

[20] R. Langley, J. Bohdansky, W. Eckstein, P. Mioduszewski, J. Roth, E. Taglauer, E. Thomas, H. Verbeek, and K. Wilson, "Data compendium for plasma-surface interactions," vol. 24, pp. S9–S117. 18, 182

[21] W. Stacey, *Fusion Plasma Physics*. Wiley-VCH, 2nd ed. 18, 29

[22] J. P. Freidberg, *Plasma Physics and Fusion Energy*. Cambridge University Press, 2007. 18, 28, 29

[23] R. Dejarnac, A. Podolnik, M. Komm, G. Arnoux, J. Coenen, S. Devaux, L. Frassinetti, J. Gunn, G. Matthews, R. Pitts, and JET-EFDA Contributors, "Numerical evaluation of heat flux and surface temperature on a misaligned JET divertor W lamella during ELMs," *Nuclear Fusion*, vol. 54, p. 123011, Dec. 2014. 20

[24] J. Gunn, S. Carpentier-Chouchana, R. Dejarnac, F. Escourbiac, T. Hirai, M. Komm, A. Kukushkin, S. Panayotis, and R. Pitts, "Ion orbit modelling of ELM heat loads on ITER divertor vertical targets," vol. 12, pp. 75–83. x, xvii, 20, 22, 23, 148, 150

[25] J. P. Gunn, "The influence of magnetization strength on the sheath: Implications for flush-mounted probes," *Physics of Plasmas*, vol. 4, pp. 4435–4446, Dec. 1997. x, 21

[26] F. F. Chen, *Introduction to Plasma Physics and Controlled Fusion*. Cham: Springer International Publishing, 2016. 25

[27] J. C. Maxwell, "Illustrations of the dynamical theory of gases - Part I. On the motions and collisions of perfectly elastic spheres.," *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, vol. 19, pp. 19–32, Jan. 1860. 25

[28] T. Looby, M. Reinke, A. Wingen, J. Menard, S. Gerhardt, T. Gray, D. C. Donovan, E. A. Unterberg, J. Klabacha, and M. Messineo, "Tokamak 3D Heat Load Predictions with HEAT," *Fusion Science and Technology*. Publisher: Taylor & Francis. xi, xii, xiii, xv, 27, 61, 62, 63, 65, 68, 70, 73, 78, 85, 125, 131, 133

[29] L. Lao, H. St. John, R. Stambaugh, A. Kellman, and W. Pfeiffer, "Reconstruction of current profile parameters and plasma shapes in tokamaks," *Nuclear Fusion*, vol. 25, pp. 1611–1622, Nov. 1985. 31, 74

[30] K. Kim, K. Im, H. Kim, S. Oh, J. Park, S. Kwon, Y. Lee, J. Yeom, C. Lee, G.-S. Lee, G. Neilson, C. Kessel, T. Brown, P. Titus, D. Mikkelsen, and Y. Zhai, "Design concept of k-DEMO for near-term implementation," vol. 55, no. 5, p. 053027. 35

[31] B. Sorbom, J. Ball, T. Palmer, F. Mangiarotti, J. Sierchio, P. Bonoli, C. Kasten, D. Sutherland, H. Barnard, C. Haakonsen, J. Goh, C. Sung, and D. Whyte, "ARC: A compact, high-field, fusion nuclear science facility and demonstration power plant with demountable magnets," vol. 100, pp. 378–405. 35

[32] F. Najmabadi, "Spherical torus concept as power plants—the ARIES-ST study," vol. 65, no. 2, pp. 143–164. 35

[33] T. W. Petrie, C. M. Green, A. W. Leonard, M. A. Mahdavi, T. H. Osborne, M. J. Scha, D. M. Thomas, W. P. West, D.-D. Team, S. L. Allen, M. E. Fenstermacher, C. J. Lasnier, G. D. Porter, N. S. Wolf, J. G. Watkins, and T. L. Rhodes, "The

effect of divertor magnetic balance on h-mode performance in DIII-d," p. 5. x, 35, 36, 37

[34] D. Brunner, A. Kuang, B. LaBombard, and J. Terry, "The dependence of divertor power sharing on magnetic flux balance in near double-null configurations on alcator c-mod," vol. 58, no. 7, p. 076010. x, 35, 36, 37

[35] M. Nave and J. Wesson, "Mode locking in tokamaks," vol. 30, no. 12, pp. 2575–2583. 38

[36] N. Ferraro, J.-K. Park, C. Myers, A. Brooks, S. Gerhardt, J. Menard, S. Munaretto, and M. Reinke, "Error field impact on mode locking and divertor heat flux in NSTX-u," vol. 59, no. 8, p. 086021. xi, 38, 39, 68

[37] J. Coenen, G. Arnoux, B. Bazylev, G. Matthews, A. Autricque, I. Balboa, M. Clever, R. Dejarnac, I. Coffey, Y. Corre, S. Devaux, L. Frassinetti, E. Gauthier, J. Horacek, S. Jachmich, M. Komm, M. Knaup, K. Krieger, S. Marsen, A. Meigs, P. Mertens, R. Pitts, T. Puetterich, M. Rack, M. Stamp, G. Sergienko, P. Tamain, V. Thompson, and JET-EFDA Contributors, "ELM-induced transient tungsten melting in the JET divertor," *Nuclear Fusion*, vol. 55, p. 023010, Feb. 2015. xi, 38, 57, 59

[38] F. Wagner, G. Becker, K. Behringer, D. Campbell, A. Eberhagen, W. Engelhardt, G. Fussmann, O. Gehre, J. Gernhardt, G. v. Gierke, G. Haas, M. Huang, F. Karger, M. Keilhacker, O. Klüber, M. Kornherr, K. Lackner, G. Lisitano, G. G. Lister, H. M. Mayer, D. Meisel, E. R. Müller, H. Murmann, H. Niedermeyer, W. Poschenrieder, H. Rapp, H. Röhr, F. Schneider, G. Siller, E. Speth, A. Stäbler, K. H. Steuer, G. Venus, O. Vollmer, and Z. Yü, "Regime of improved confinement and high beta in neutral-beam-heated divertor discharges of the ASDEX tokamak," vol. 49, no. 19, pp. 1408–1412. 38

[39] A. W. Leonard, "Edge-localized-modes in tokamaks," vol. 21, no. 9, p. 090501. xi, 38, 40, 41

[40] T. Eich, "ELM divertor peak energy fluence scaling to ITER with data from JET, MAST and ASDEX upgrade," *Nuclear Materials and Energy*, p. 7, 2017. 40, 127

[41] A. W. Leonard, "Plasma detachment in divertor tokamaks," vol. 60, no. 4, p. 044001. xi, 45, 46, 49

[42] V. A. Soukhanovskii, "A review of radiative detachment studies in tokamak advanced magnetic divertor configurations," vol. 59, no. 6, p. 064005. 46

[43] A. Kallenbach, M. Bernert, M. Beurskens, L. Casali, M. Dunne, T. Eich, L. Giannone, A. Herrmann, M. Maraschek, S. Potzel, F. Reimold, V. Rohde, J. Schweinzer, E. Viezzer, M. Wischmeier, and the ASDEX Upgrade Team, "Partial detachment of high power discharges in ASDEX upgrade," vol. 55, no. 5, p. 053026. xi, 46, 48

[44] A. Kallenbach, M. Bernert, R. Dux, F. Reimold, M. Wischmeier, and ASDEX Upgrade Team, "Analytical calculations for impurity seeded partially detached divertor conditions," vol. 58, no. 4, p. 045013. 46

[45] A. Dillmann, G. Heller, M. Klaas, H.-P. Kreplin, W. Nitsche, W. Schröder, E. H. Hirschel, W. Schröder, K. Fujii, W. Haase, B. Leer, M. A. Leschziner, M. Pandolfi, J. Periaux, A. Rizzi, B. Roux, and Y. I. Shokin, eds., *New Results in Numerical and Experimental Fluid Mechanics VII*, vol. 112 of *Notes on Numerical Fluid Mechanics and Multidisciplinary Design*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2010. 50

[46] R. Neu, H. Maier, M. Balden, R. Dux, S. Elgeti, H. Gietl, H. Greuner, A. Herrmann, T. Höschen, M. Li, V. Rohde, D. Ruprecht, B. Sieglin, and I. Zammuto, "Results on the use of tungsten heavy alloys in the divertor of ASDEX Upgrade," *Journal of Nuclear Materials*, vol. 511, pp. 567–573, Dec. 2018. xi, 58, 59, 68

[47] J. W. Coenen, G. F. Matthews, K. Krieger, D. Iglesias, P. Bunting, Y. Corre, S. Silburn, I. Balboa, B. Bazylev, N. Conway, I. Coffey, R. Dejarnac, E. Gauthier, J. Gaspar, S. Jachmich, I. Jepu, C. Makepeace, R. Scannell, M. Stamp, P. Petersson, R. A. Pitts, S. Wiesen, A. Widdowson, K. Heinola, A. Baron-Wiechec, and JET Contributors, "Transient induced tungsten melting at the Joint European Torus (JET)," *Physica Scripta*, vol. T170, p. 014013, Dec. 2017. 59

[48] S. P. Gerhardt, "OVERVIEW OF THE NSTX-U RECOVERY PROJECT PHYSICS AND ENGINEERING DESIGN," p. 8, 2018. 62

[49] J. Fang, R. Ellis, and P. Titus, "Thermal and Structural Analysis of Outboard Divertor Row 1&2 Tiles for NSTX-U," *IEEE Transactions on Plasma Science*, vol. 48, pp. 1468–1473, June 2020. 62, 68

[50] T. Looby, M. Reinke, D. Donovan, T. Gray, M. Messineo, and A. Khodak, "Convolutional Neural Networks for Heat Flux Model Validation on NSTX-U," *IEEE Transactions on Plasma Science*, vol. 48, pp. 3–13, Jan. 2020. 62, 128

[51] A. Khodak, D. Loesser, M. Messineo, A. Brooks, M. Jaworski, D. L. Youchison, and J. Cook, "Design and Analysis of High Heat Flux Plasma-Facing Components for NSTX Upgrade," *IEEE Transactions on Plasma Science*, pp. 1–7, 2021. 68

[52] A. Wingen, M. L. Reinke, T. Looby, and D. M. Orlov, "Non-axisymmetric heat flux patterns on tokamak divertor tiles," in *Proceeding to the 46th EPS conference on plasma physics*, vol. 1040, (Milan), p. 5, July 2019. 70

[53] L. Kos, R. A. Pitts, G. Simič, M. Brank, H. Anand, and W. Arter, "SMITER: A field-line tracing environment for ITER," *Fusion Engineering and Design*, vol. 146, pp. 1796 – 1800, 2019. 70

[54] M. Firdaouss, V. Riccardo, V. Martin, G. Arnoux, and C. Reux, "Modelling of power deposition on the JET ITER like wall using the code PFCFLux," *Journal of Nuclear Materials*, vol. 438, pp. S536–S539, July 2013. 70

[55] B. Zhang, M. Firdaouss, X. Gong, A. Ekedahl, X. Peng, and X. Zhang, "Study of power load pattern on EAST divertor using PFCFlux code," *Fusion Engineering and Design*, vol. 107, pp. 58–63, June 2016. 70

[56] J. Gerardin, M. Firdaouss, F. Maviglia, W. Arter, T. Barrett, and M. Kovari, "Simplified heat load modeling for design of DEMO discrete limiter," *Nuclear Materials and Energy*, vol. 20, p. 100568, Aug. 2019. 70

[57] "FreeCAD webpage." https://www.freecadweb.org/. Accessed: 2021-01-02. 70, 78

[58] G. Chen, Q. Xiong, P. J. Morris, E. G. Paterson, A. Sergeev, and Y.-C. Wang, "OpenFOAM for Computational Fluid Dynamics," *Notices of the American Mathematical Society*, vol. 61, p. 354, Apr. 2014. 70, 79

[59] "ParaView webpage." https://www.paraview.org/. Accessed: 2021-01-02. 72

[60] "Plotly Dash webpage." https://plotly.com/dash/. Accessed: 2021-01-02. 72, 74

[61] "HEAT github." https://github.com/plasmapotential/HEAT. Accessed: 2021-01-02. 72, 75, 109

[62] A. Wingen, T. E. Evans, and K. H. Spatschek, "High resolution numerical studies of separatrix splitting," *Nucl. Fusion*, p. 9, 2009. 84

[63] A. Kensler and P. Shirley, "Optimizing Ray-Triangle Intersection via Automated Search," in *2006 IEEE Symposium on Interactive Ray Tracing*, (Salt Lake City, UT, USA), pp. 33–38, IEEE, Sept. 2006. 86

[64] T. Möller and B. Trumbore, "Fast, minimum storage ray-triangle intersection," *Journal of Graphics Tools*, vol. 2, no. 1, pp. 21–28, 1997. xiii, 86, 89, 90, 91

[65] M. A. Makowski, D. Elder, T. K. Gray, B. LaBombard, C. J. Lasnier, A. W. Leonard, R. Maingi, T. H. Osborne, P. C. Stangeby, J. L. Terry, and J. Watkins, "Analysis of a multi-machine database on divertor heat fluxes," *Physics of Plasmas*, vol. 19, p. 056122, May 2012. 102, 132

[66] J. Horacek, R. A. Pitts, J. Adamek, G. Arnoux, J.-G. Bak, S. Brezinsek, M. Dimitrova, R. J. Goldston, J. P. Gunn, J. Havlicek, S.-H. Hong, F. Janky, B. LaBombard, S. Marsen, G. Maddaluno, L. Nie, V. Pericoli, T. Popov, R. Panek, D. Rudakov, J. Seidl, D. S. Seo, M. Shimada, C. Silva, P. C. Stangeby, B. Viola, P. Vondracek, H. Wang, G. S. Xu, and Y. Xu, "Multi-machine scaling of the main SOL parallel heat flux width in tokamak limiter plasmas," p. 12, 2016. 102

[67] M. Kočan, R. Pitts, S. Lisgo, A. Loarte, J. Gunn, and V. Fuchs, "Modelling ELM heat flux deposition on the ITER main chamber wall," *Journal of Nuclear Materials*, vol. 463, pp. 709–713, Aug. 2015. 127

[68] S. Gerhardt and M. Reinke, "Heat fluxes on the inboard divertor horizontal and near outboard divertor regions (rows 1 and 2)," Sep 2017. 132, 135

[69] J. E. Menard, J. P. Allain, D. J. Battaglia, F. Bedoya, R. E. Bell, E. Belova, J. W. Berkery, M. D. Boyer, N. Crocker, A. Diallo, F. Ebrahimi, N. Ferraro, E. Fredrickson, H. Frerichs, S. Gerhardt, N. Gorelenkov, W. Guttenfelder, W. Heidbrink, R. Kaita, S. M. Kaye, D. M. Kriete, S. Kubota, B. P. LeBlanc, D. Liu, R. Lunsford, D. Mueller, C. E. Myers, M. Ono, S. A. Sabbagh, O. Schmitz, F. Scotti, Y. Sechrest, C. H. Skinner, D. R. Smith, V. Soukhanovskii, T. Stoltzfus-Dueck, H. Yuh, Z. Wang, I. Waters, N. Bertelli, A. Bhattacharjee, D. Brennan, R. Buttery, A. Capece, G. Canal, J. Canik, C. S. Chang, D. Darrow, L. Delgado-Aparicio, C. Domier, S. Ethier, T. Evans, J. Ferron, M. Finkenthal, R. Fonck, K. Gan, D. Gates, I. Goumiri, T. Gray, J. Hosea, D. Humphreys, T. Jarboe, S. Jardin, M. A. Jaworski, B. Koel, E. Kolemen, S. Ku, R. J. L. Haye, F. Levinton, N. Luhmann, R. Maingi, R. Maqueda, G. McKee, E. Meier, J. Myra, R. Perkins, F. Poli, T. Rhodes, J. Riquezes, C. Rowley, D. Russell, E. Schuster, B. Stratton, D. Stutman, G. Taylor, K. Tritz, W. Wang, B. Wirth, and S. J. Zweben, "Overview of NSTX Upgrade initial results and modelling highlights," *Nucl. Fusion*, p. 18, 2017. 138

[70] P. Stangeby, C. Pitcher, and J. Elder, "The nature of plasma fluxes to surfaces nearly tangential to the magnetic field," *Nuclear Fusion*, vol. 32, pp. 2079–2089, Dec. 1992. 141

[71] M. Cada, Z. Hubicka, P. Adamek, J. Olejnicek, S. Kment, J. Adamek, and J. Stockel, "A modified Katsumata probe—Ion sensitive probe for measurement in non-magnetized plasmas," *Review of Scientific Instruments*, vol. 86, p. 073510, July 2015. 168