

UC San Diego

UC San Diego Electronic Theses and Dissertations

Title

In-Vessel Top Launch Mirror on the DIII-D Tokamak

Permalink

<https://escholarship.org/uc/item/58g2j7xn>

ISBN

9798191648255

Author

de Boucaud, Nikolai

Publication Date

2026-09-21

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA SAN DIEGO

In-Vessel Top Launch Mirror on the DIII-D Tokamak

A thesis submitted in partial satisfaction of the requirements
for the degree Master of Science

in

Engineering Sciences (Mechanical Engineering)

by

Nikolai de Boucaud

Committee in charge:

Professor Alessandro Marinoni, Chair
Professor Daniel Sievenpiper
Professor George Tynan

2026

Copyright

Nikolai de Boucaud, 2026

All rights reserved

The thesis of Nikolai de Boucaud is approved, and it is acceptable in quality and form for publication on microfilm and electronically.

University of California San Diego

2026

Table of Contents

Title Page	i
Copyright Notice	ii
Thesis Approval Page	iii
Table of Contents	iv
List of Figures	vii
List of Tables	xii
Abstract	xiii
1 Introduction	1
1.1 Nuclear Fusion	1
1.2 Magnetic Confinement Fusion	5
1.3 Plasma Physics of Tokamaks	9
1.3.1 Magnetohydrodynamics	9
1.3.2 Electron Cyclotron Resonance	13
1.4 The DIII-D Tokamak	22
1.4.1 DIII-D Machine Overview	23
1.4.2 The ECH System on the DIII-D Tokamak	27
1.4.3 Top Launchers	36

1.4.4	Legacy Top Launcher Damage	39
1.5	Top Launch Mirror Concept	41
2	Engineering Design	43
2.1	Constraints	43
2.1.1	Beam Size Constraint	44
2.1.2	First Wall Constraints	48
2.1.3	Nominal Mirror Position Determination	50
2.2	Conceptual Study	53
2.2.1	Mirror Material Study	53
2.2.2	Outside Launcher Aiming	58
2.2.3	Sensor Selection	59
2.3	Final Design	60
2.3.1	Optimal Reflection Angle Determination	60
2.3.2	Mirror Mechanical Design	63
2.3.3	Mirror Focusing Design	69
2.3.4	Sensing Design	71
2.3.5	Thermal and Structural Analysis	75
3	Installation and Testing	80
3.1	Installation	80
3.1.1	Thermocouple Installation	81
3.1.2	Laser Alignment	83
3.2	Testing	85
3.2.1	Aiming Verification via Hot Test	85
3.2.2	Thermal Model Calibration and Pulse Extension	89
3.2.3	Plasma Deposition and Polarization Experiment	97
4	Conclusions and Future Work	101

A	104
----------	------------

List of Figures

Figure 1.1:	Average binding energy per nucleon of common nuclei plotted against their atomic number.	2
Figure 1.2:	Basic illustration of a set of toroidally-arranged coils.	6
Figure 1.3:	Definition of the cylindrical coordinate system used for tokamak equilibria.	10
Figure 1.4:	A cross-section of a tokamak equilibrium $\psi(R, Z)$ along a plane in R,Z showing the nested nature of the magnetic flux surfaces.	11
Figure 1.5:	Example of profiles of a plasma. In this case, a relatively low temperature DIII-D shot, No. 200000 at time 1 s, is used.	12
Figure 1.6:	Basic illustration of acceleration of an electron (represented in blue) in a cyclotron orbit by an alternating electric field.	14
Figure 1.7:	Simplified CMA Diagram showing cutoffs in the plasma.	18
Figure 1.8:	A cross-section of a DIII-D plasma equilibrium.	21
Figure 1.9:	The coil systems, vacuum vessel and a representative plasma equilibrium of the DIII-D tokamak.	24
Figure 1.10:	DIII-D vacuum vessel port definitions.	25

Figure 1.11:	A cross-section of an upper null diverted DIII-D shot.	26
Figure 1.12:	Overhead view of a SolidWorks assembly of the DIII-D facility layout.	28
Figure 1.13:	DIII-D vacuum vessel and magnetic coils with the four outside launchers.	29
Figure 1.14:	Angles defining the polar and azimuthal directions for EC beam steering.	30
Figure 1.15:	Cross-section of a DIII-D outside launcher.	31
Figure 1.16:	Definition of the launcher steering angles.	32
Figure 1.17:	Infrared images from the hot test to determine steering limits of a DIII-D launcher.	34
Figure 1.18:	An example of steering limits found via the hot test method.	35
Figure 1.19:	The limitations on launched beam angles (polar/azimuthal) as deter- mined by the hot test.	36
Figure 1.20:	A section of the DIII-D tokamak showing the components making up a top launch system.	37
Figure 1.21:	View of the trajectory of the injected beam from a top launch system.	39
Figure 1.22:	Images of the extracted top launch waveguide.	40
Figure 2.1:	Paraboloid mirror surface of outside launcher focusing mirror.	47
Figure 2.2:	Expansion of the radius of the Gaussian (fundamental mode) beam as it propagates from the exit of the waveguide.	48
Figure 2.3:	Model view of the upper shelf region.	50

Figure 2.4:	Usable region for the top launch mirror.	51
Figure 2.5:	Model view initial concept mirror.	52
Figure 2.6:	Test setup using VNA.	55
Figure 2.7:	Elliptical distribution of power density absorbed by the graphite mirror.	57
Figure 2.8:	Optimization of aiming angles for top launch mirror.	62
Figure 2.9:	Optimal aiming for the top launch mirror.	62
Figure 2.10:	Final mechanical design of the top launch mirror.	65
Figure 2.11:	Standard nut bars used for mounting tiles in the DIII-D vessel. . . .	66
Figure 2.12:	Nut bars used for the top launch mirror.	67
Figure 2.13:	Final design of the tiles surrounding the top launch mirror.	68
Figure 2.14:	The poloidal projection of the top launch mirror's extremities. . . .	69
Figure 2.15:	Propagation of the radius of the Gaussian beam from the exit of the waveguide.	71
Figure 2.16:	Drawing view of a hole in the tile for an embedded thermocouple. . .	72
Figure 2.17:	Model view of the top launch mirror and its surrounding tiles. . . .	73
Figure 2.18:	Graphical User Interface (GUI) view of the arrangement of embedded thermocouples in the top launch mirror.	74
Figure 2.19:	Results of the SolidWorks thermal model applied to the graphite top launch mirror.	77

Figure 2.20: Simulation of plasma power density incident on the surface of top launch mirror.	78
Figure 3.1: Image taken during the machining of the top launch mirror on the CNC mill.	81
Figure 3.2: Wiring path taken in-vessel for the top launch thermocouples.	82
Figure 3.3: Image taken of the side of the top launch mirror during installation.	83
Figure 3.4: Image taken during assembly of the hot test fixture inside the DIII-D vacuum vessel.	87
Figure 3.5: Propagation of the radius of the Gaussian beam from the top launch mirror compared with the measured values.	88
Figure 3.6: Image of the finalized installation of the top launch mirror and its surrounding tiles.	89
Figure 3.7: Thermocouple temperature readings compared with simulated thermocouple readings for a 250 ms pulse.	93
Figure 3.8: Thermocouple temperature readings compared with simulated thermocouple readings for a 6 s pulse.	95
Figure 3.9: Thermocouple temperature readings compared with simulated thermocouple readings for DIII-D shot 202780.	96
Figure 3.10: Experimentally measured wave absorption for polarizations in the X-mode and O-mode.	99
Figure 3.11: Comparison of results for the predicted peak deposition location.	100

Figure 4.1: Model view of the planned installation of eight total top launch mirrors.103

List of Tables

Table 1.1:	Key DIII-D parameters.	23
Table 1.2:	Components Upstream of Outside Launchers	33
Table 3.1:	Parameters tuned to match the thermal model to thermocouple data	91
Table 3.2:	Tuned parameters used in the thermal model.	92
Table A.1:	Measured insertion loss of copper and ATJ graphite miter bends. . .	104

ABSTRACT OF THE THESIS

In-Vessel Top Launch Mirror on the DIII-D Tokamak

by

Nikolai de Boucaud

Master of Science in Engineering Sciences (Mechanical Engineering)

University of California San Diego, 2026

Professor Alessandro Marinoni, Chair

A prototype “top launch” electron cyclotron (EC) heating system has been developed at the DIII-D tokamak that utilizes an in-vessel mirror to reflect an outside-launched microwave beam into the plasma with a nearly-parallel path along the resonance line. The DIII-D EC system currently has eight steerable mirror systems and one of these has been used to test this prototype system. In 2020 and 2022, two top launch systems were installed on DIII-D using

in-vessel waveguides through upper ports on the machine. These systems demonstrated a reactor-relevant technology for driving off-axis current in Advanced Tokamak (AT) scenarios. However, one of the in-vessel waveguides failed and both systems were removed in 2023. This new prototype was designed to replace the previous systems with a reliable alternative at 10% the cost. The prototype in-vessel mirror is made of ATJ graphite to match the material of the DIII-D first wall. The mirror's surface is canted at a fixed angle to reflect the beam at the desired angle and it focuses it inside the plasma. The beam's propagation was measured to closely match the design via an infrared camera and a target setup. The mirror also has thermocouples embedded in it to monitor its temperature. The mirror was predicted to reflect approximately 95% of the power incident on its surface, which was confirmed via thermocouple readings. The successful demonstration of this prototype system will result in its duplication for all steerable mirror systems on DIII-D, which is predicted to double its off-axis current drive capabilities.

Chapter 1

Introduction

To begin the document, I will provide relevant background for the topics involved in the design, installation and testing of the in-vessel top launch mirror system. These background topics include a general discussion of nuclear fusion, magnetic confinement fusion, the concepts of electron cyclotron resonance heating and current drive, the DIII-D tokamak, the system used for electron cyclotron heating and current drive on the DIII-D tokamak, and the previous top launch system that was installed on the DIII-D tokamak.

1.1 Nuclear Fusion

It has been a goal of science and engineering to produce useful energy from nuclear fusion reactions since the early to mid-1900s when it was discovered that stars fuse hydrogen atoms into heavier helium atoms to produce the energy needed to balance their immense gravity. In the star's core, the kinetic energy of atoms is sufficient that atomic nuclei can't maintain the electrostatic bond to their electron pairs and are therefore fully ionized, meaning electrons and nuclei are free to move relatively independently. This state of matter is a gas that, while uncharged overall (i.e. quasineutral), is composed of charged particles that exhibit collective behavior due to long-range electromagnetic interactions. This is known as a plasma. The kinetic energy of particles (temperature) in the star's core is so high, in fact, that if two

nuclei collide, they are likely to have enough speed to overcome their mutual long-range electrostatic repulsion and come close enough that the short-range strong nuclear force pulls them together. This is the process of nuclear fusion. Nuclear fusion is an exothermic reaction for a number of elements, with a general trend toward larger energy product for the lightest elements: hydrogen and its isotopes. A plot showing the average nuclear binding energy per nucleon for a number of elements is given in Fig. 1.1.




Figure 1.1: Average binding energy per nucleon of common nuclei plotted against their atomic number. The data plotted is taken from the AME2020 atomic mass evaluation [1, 2, 3]. Exothermic fusion reactions occur when the total binding energy of the products exceeds that of the reactants.

In order to extract useful energy from the process of nuclear fusion, the necessary conditions must be created to fuse many nuclei using a system that produces more energy than it consumes. A successful machine must therefore create a confined environment, like the core of a star, in which the nuclei have the necessary velocity for fusion to occur. The

conditions satisfying this are typically a plasma. The nuclei in the plasma have random trajectories due to their thermal motion, making head-on collisions rare. However, if a sufficient number density of nuclei are held in this state for a long enough time, fusion reactions become frequent. The figure of merit of fusion reactions in this confined state is described by the energy balance conditions dictated by the triple product of temperature, density, and confinement time [4]. The fusion triple product is derived by balancing the auxiliary power supplied to the plasma with the difference between the power generated by fusion reactions and the power lost by the plasma (through conduction or electromagnetic radiation losses). The triple product is given by

$$nT\tau_E \geq \frac{12}{E_p} \frac{T^2}{\langle \sigma v \rangle}, \quad (1.1)$$

where the left hand side gives the triple product of the number density of the plasma, n , the plasma temperature, T in units of energy (eV), and the confinement time (which measures the rate at which the plasma loses energy), τ_E . The right hand side contains a term for the fusion energy that is transferred to the bulk plasma, as opposed to neutrons and electromagnetic radiation, E_p , and a term describing the reactivity, $\langle \sigma v \rangle$, which is proportional to the probability σ of the reaction under consideration. Most efforts in fusion target the minimum possible value of the right hand side of Eq. (1.1), which uses two different isotopes of hydrogen: ^2H (deuterium) and ^3H (tritium). The probability of fusion reactions is a function of the kinetic energy of the reactants and the reaction between deuterium and tritium has the highest $\langle \sigma v \rangle$ for the lowest value of T . This criterion gives the figure of merit for all confined nuclear fusion concepts, spawning a host of different designs and experimental machines.

It's generally necessary to confine a fusion plasma using non-material means, since any contact between the plasma and a cold, non-plasma material will inevitably cool it down and generate an unwanted influx of particles from the container's wall. This is also why the plasma volume must be devoid of any gases, particulates or other matter, called impurities,

apart from the plasma fuel itself. One technique for non-material confinement is to use a pulse of high-power laser radiation surrounding a sphere of solid fuel material, creating a plasma at an extremely high temperature and density and allowing the outer layers of plasma to compress it further. This is generally referred to as inertial confinement due to the inertia of the outer layers of the material being used to compress the core. Although similar in spirit to the core of a star, due to the lack of a gravity-like (i.e. centrally-directed, steady-state) force, this is necessarily not an equilibrated configuration, unlike a star. Another technique for non-material confinement is the use of magnetic fields which, under the Lorentz force (with no electric field),

$$\mathbf{F} = q(\mathbf{v} \times \mathbf{B}), \quad (1.2)$$

maintain a confined region of the charged particles of the plasma. This is done by keeping the magnetic field ($\mathbf{B}$, magnetic flux density) directed away from any material surface, thus resulting in the velocity ($\mathbf{v}$) of any particle with non-zero charge (q) following a circular trajectory in a plane orthogonal to the magnetic field¹. In Eq. (1.2), and throughout the rest of the document, quantities in bold typeface represent vectors in 3-D space. If the strength of the field is sufficiently high, the radius of the circular trajectories of fast particles can be reduced enough that the particles are confined to a desired region in all directions perpendicular to the magnetic field, but notably, particles with velocity purely parallel to $\mathbf{B}$ are unaffected by this force. This circular motion is referred to as cyclotron motion. The equations for the radius and frequency of a cyclotron orbit (called gyroradius and gyrofrequency, respectively) are given by

$$\begin{aligned} r_c &= \frac{v_{\perp}}{\Omega_c}, \\ \Omega_c &= \frac{q|\mathbf{B}|}{m}, \end{aligned} \quad (1.3)$$

¹Note that particles follow circular trajectories in the particular case of a uniform, static magnetic field. Otherwise, the trajectory may be elongated or otherwise perturbed.

where r_c is the gyroradius, Ω_c is the signed angular cyclotron frequency, m is the mass of the charged particle, q is the charge of the particle, and $v_\perp$ is the velocity perpendicular to the applied magnetic field. Given that a plasma is produced by heating a gas with particle velocities in random directions, often following a Maxwellian velocity distribution, all particles in a plasma begin with a mixture of velocity parallel and perpendicular to the magnetic field, causing them to follow helical trajectories. A configuration that maintains the plasma in a confined region using a magnetic field is referred to as magnetic confinement and is generally maintainable in an equilibrium state, meaning that the plasma discharge can be maintained for several energy confinement times τ_E .

1.2 Magnetic Confinement Fusion

The first attempts to produce a magnetically-confined plasma used a series of linearly-spaced magnetic coils that produced a cylindrical region of plasma. The challenge of such a configuration is found at the ends of the cylindrical region, which must be material walls, where the helical particle trajectories must end or be reversed, meaning the plasma is either cooled or unconfined. A typical method of limiting the plasma's leakage from the ends of the cylindrical region is to create a steep gradient in the magnetic field strength at the ends of the cylindrical region, causing the helical trajectories of the particles to switch to the opposite parallel direction. This method was found to be prone to end losses in spite of this, leading researchers to attempt to remove the end regions altogether by joining the opposite sides of the cylindrical region to form a torus. This basic toroidal configuration, along with directional definitions used in the toroidal geometry, is shown in Fig. 1.2. However, due to the close proximity of the inner edges of the coils, compared to the outer edges, the strength of the field is heightened near the inner edge and falls at the outer edge. This continuous gradient in the strength of the field (∇B) creates an effect on the cyclotron motion of the charged particles, in which their orbit is compressed and elongated on the inner and outer side of their orbit,

respectively. This causes the center of a particle's orbit to drift either upward or downward depending on their charge. A charge separation leads to the establishment of an electric field in the plasma that is perpendicular to the magnetic field and causes particles to drift out of confinement very rapidly, irrespective of their charge, due to the Lorentz force with non-zero electric field intensity. In order to prevent the initial drift caused by ∇B , thus preventing the establishment of an electric field, researchers found that a twist of the magnetic field within its toroidal volume, causing the helical cyclotron orbits of charged particles to follow a larger-scale helical trajectory around the torus, neutralized the effect of ∇B . A number of other drift effects are also important to understand for toroidal plasma confinement [4] but the ∇B drift described here makes clear the immediate need for a helical magnetic field. The two dominant methods of achieving this helical twist in the toroidal magnetic field are known as a stellarator and a tokamak.




Figure 1.2: Basic illustration of a set of toroidally-arranged coils. In the isometric projection, the blue arrow indicates the toroidal direction of the magnetic field within the coils and the red arrow indicates the poloidal direction that is everywhere perpendicular to the toroidal direction. A projection in a poloidal plane is shown on the right, indicating the radial and vertical directions.

A stellarator creates a twisted toroidal magnetic field via the geometry of its magnetic

coil configuration. By orienting coils at specific angles, or with specific coil geometry, the magnetic field can be made to follow the desired helical trajectory. This typically results in coil geometries that are difficult to fabricate and plasma geometries that are difficult to model, but ultimately a device that needs only its coils to be energized to allow steady-state confinement to be achieved. In contrast, the tokamak achieves a twisted toroidal magnetic field via an induced toroidal current within the plasma. A plasma, being a volume of charged particles, is a highly conductive medium and can carry a current. Therefore, the principle of electromagnetic induction, in which a magnetic field with a time-varying intensity causes an electric current to be induced, can be used to create a current in the plasma. If this plasma current is carried in a continuous loop around the torus, the magnetic field resulting will encircle it in a direction that is everywhere perpendicular to the externally-imposed toroidal magnetic field - referred to as the poloidal direction. The overall magnetic field is thus twisted similarly to a stellarator but does not depend on the toroidal position. A tokamak can be created by introducing another magnetic coil, in a configuration similar to an electrical transformer, to the simple toroidally-situated coils from Fig. 1.2, using a stack of coils located at the center of the torus as the primary coil of the transformer, called the central solenoid. The central solenoid drives the plasma current inductively, with a current that increases or decreases with time. This geometry is generally simpler to fabricate than that of a stellarator, but is, by nature of its time-varying induced magnetic field, a pulsed device. It's also important to mention the effect of the "bootstrap" current, which is a self-sustained current in the tokamak (also in stellarators) created by a gradient in the plasma's temperature or density. Due to the fact that the bootstrap current is intrinsic to the confined plasma, does not vary with time, and has an additive effect with the driven current, many modern tokamaks take advantage of the bootstrap current by designing plasma scenarios that maximize it. The bootstrap current can make up more than half of the total current in optimized plasmas. The effect won't be expanded on here as the physics is outside the scope of this document.

The toroidal plasma current, in addition to providing the necessary helical twist to the magnetic field for a stable toroidal plasma, produces a Joule heating power (also known as Ohmic heating) in the plasma that's given by

$$P = \eta J^2 V, \tag{1.4}$$

where P is the heating power, J is the current per unit area, η is the electrical resistivity of the material, and V is the volume through which the current flows. This heating power is intrinsic to the tokamak and begins the plasma heating process up to the temperatures required for nuclear fusion. However, the plasma has an important property that distinguishes it from typical conductive materials: the resistivity of a plasma generally reduces with increasing temperature. This is due to the collisionality of the plasma reducing with higher temperature, i.e. the particles within the plasma tend to have fewer direct collisions at higher temperature and instead tend to have glancing collisions with low energy transfer from one particle to another. This leads to charged particles being accelerated without collisions impeding their motion, thus a lower resistivity. Due to this fact, plasma heating beyond 1 keV requires external means of heating. One method of auxiliary heating, titular to this document, uses electromagnetic radiation matching a resonant frequency with the cyclotron orbits of electrons in the plasma and is therefore known as electron cyclotron resonance heating (ECRH or just ECH). Importantly, the heating of particles in the plasma can be performed in a way that preferentially accelerates those with the same charge in a particular toroidal direction, producing a toroidal current. This is known as auxiliary current drive and provides a means of maintaining steady-state confinement in a tokamak plasma after the central solenoid has finished one cycle, or pulse. Electron cyclotron resonance heating that preferentially heats electrons traveling in a particular toroidal direction is known as electron cyclotron resonance current drive (ECCD). These concepts will be expanded on in further sections.

1.3 Plasma Physics of Tokamaks

The tokamak is the leading candidate for a power plant nuclear fusion reactor. The physics of a tokamak is complex and can't be expanded on in complete detail in this document, but some important aspects of a steady-state tokamak plasma are discussed. Additionally, an overview of the mechanism of electron cyclotron heating and current drive is given.

1.3.1 Magnetohydrodynamics

The equilibrated state of a tokamak, that is when the central solenoid is driving a constant plasma current - called a “flat top” phase, is well approximated by the use of magnetohydrodynamic (MHD) theory. In MHD, Maxwell's equations are coupled to a fluid model of the plasma and an equilibrium is reached by balancing the magnetic pressure with the plasma pressure. The resulting equations are given by

$$\begin{aligned}\mathbf{J} \times \mathbf{B} &= \nabla P \\ \nabla \times \mathbf{B} &= \mu_0 \mathbf{J} \\ \nabla \cdot \mathbf{B} &= 0\end{aligned}\tag{1.5}$$

where the plasma current density is given by $\mathbf{J}$, the magnetic flux density is given by $\mathbf{B}$, plasma pressure is given by P and μ_0 is the magnetic permeability of free space. For a tokamak, these equations are most useful if solved in a toroidal coordinate system shown in Fig. 1.3. It can be seen from the first equation that, given a current that is purely toroidal, $\mathbf{J} = J_\phi \hat{\phi}$, the variation of pressure is purely in the poloidal (R, Z) plane. Assuming, further, that there is no variation of these quantities along the toroidal direction, i.e. it is a toroidally symmetric system, the following relationship can be reached

$$\frac{\partial^2 \psi}{\partial R^2} + \frac{\partial^2 \psi}{\partial Z^2} - \frac{1}{R} \frac{\partial \psi}{\partial R} = -\mu_0 R^2 \frac{dP}{d\psi} - \frac{1}{2} \frac{dF^2}{d\psi}$$

$$F = RB_\phi$$

$$\psi = RA_\phi$$
(1.6)

where A_ϕ is the magnetic vector potential in the toroidal direction, B_ϕ is the magnetic flux density in the toroidal direction and $\psi = \psi(R, Z)$ is defined as the stream function. $P = P(\psi)$ and $F = F(\psi)$ are found to be functions of ψ and therefore solving for ψ defines the plasma equilibrium. Equation (1.6) is called the Grad-Shafranov equation [5, 6, 7].




Figure 1.3: Definition of the cylindrical coordinate system used for tokamak equilibria, where the radial direction is denoted by $\mathbf{R}$, the vertical direction is denoted by $\mathbf{Z}$, and the toroidal direction is denoted by ϕ .

The standard form of the Grad-Shafranov equation shows a method to compute the axisymmetric tokamak equilibrium. Axisymmetry refers to symmetry about the Z axis of the tokamak (see Fig. 1.3). This feature is used throughout this work and tokamak physics in general. The magnetic field lines that travel helically around the torus (in some regions closing back on themselves) create nested surfaces, called “magnetic flux surfaces,” growing outward in radius away from a central region where the surface degenerates into a single toroidal line, defined as the magnetic axis. The last closed such surface, called the separatrix, is the edge of the confined plasma, because magnetic field lines beyond this edge impinge on the walls of the device. An example of the nested flux surfaces of a tokamak equilibrium is

shown in Fig. 1.4, for a visual aid.




Figure 1.4: A cross-section of a tokamak equilibrium $\psi(R, Z)$ along a plane in R, Z showing the nested nature of the magnetic flux surfaces. Typically only the right half of this diagram is shown due to the system's symmetry about the Z axis. The equilibrium is constructed via a Python workflow developed previously [8] using a Python package called FreeGS [9] that solves the Grad-Shafranov equation. The separatrix is plotted in red and the wall of the tokamak is represented by the dashed lines. The magnetic axis is located at the region of maximum ψ .

The nested flux surfaces each have a constant value of $\psi(R, Z)$ and since $P = P(\psi)$ and $F = F(\psi)$, the values of P and F are constant as well. Therefore, a value of pressure, current density, temperature and other quantities can be defined as 1-dimensional profiles as functions of ψ normalized to its value at the magnetic axis and at the separatrix. We can define a normalized value of ψ as

$$\rho = \sqrt{\frac{\psi - \psi_{axis}}{\psi_{sep} - \psi_{axis}}} \quad (1.7)$$

where ψ_{axis} and ψ_{sep} are the values at the magnetic axis and the separatrix, respectively. A plot of example profiles of plasma pressure, current density and electron temperature

is provided in Fig. 1.5. Much of the work to develop a tokamak plasma scenario involves properly tuning these profiles to reach desired plasma stability and fusion performance.




Figure 1.5: Example of profiles of a plasma. In this case, a relatively low temperature DIII-D shot, No. 200000 at time 1 s, is used.

Notably, for a stable tokamak plasma, it was found in early experiments that the toroidal magnetic field (B_ϕ) must dominate over the poloidal magnetic field (B_θ). Typically the magnitude of the toroidal magnetic field is a few times larger than that of the poloidal magnetic field [4]. The side of the tokamak closest to the central solenoid is referred to as the “high-field side” of the tokamak while everywhere outboard of the magnetic axis is called the “low-field side” of the tokamak, since the toroidal field coil arrangement results in a larger magnetic field closer to the center (as discussed in 1.2). These are often used interchangeably with the terms “outboard” and “inboard,” respectively.

Additionally, since the Lorentz force only alters the motions of charged particles in the direction perpendicular to the applied magnetic field, their motion parallel to the applied magnetic field is unaffected. In the case of a thermonuclear plasma, the thermal speed of the electrons and nuclei is so large that the plasma state is practically invariant along a magnetic field line. This feature illustrates that the plasma can be perturbed (e.g. heated) at one toroidal location and the perturbation will quickly propagate to every location along the magnetic field line, making the chosen toroidal location irrelevant. This feature, called

axisymmetry, is exploited in every aspect of the device, from heating and current drive to diagnostic instrumentation.

1.3.2 Electron Cyclotron Resonance

The lower mass of the electrons in the plasma, as compared to the nuclei (which are referred to as ions from here on, in keeping with standard plasma physics nomenclature), results in the electrons following cyclotron orbits with higher frequency. There is a separation of approximately three orders of magnitude between the cyclotron frequencies for electrons as compared to ions. The frequency range relevant to modern tokamak experiments tends to be on the order of 100 GHz. This large separation of scales results in their physics being generally separable.

In its most basic form, the effect of an EM wave on an electron in a static magnetic field can be seen through the electric force of the time-varying electric field incident on the electron following a cyclotron orbit. Taking the magnetic field to be along the $\hat{z}$ direction, the cyclotron orbit lies in the x-y plane. If the electric field is polarized along the $\hat{y}$ direction and its frequency and phase are properly selected to exactly match that of the electron's cyclotron frequency, it can be seen that the electric field can be aligned with the direction of the electron on the left half of its orbit and the right half of its orbit. This is illustrated in Fig. 1.6. Further, if the electric field is held to be parallel with the electron's velocity throughout its orbit (as can be achieved for a perfectly in-phase circularly-polarized EM wave) the electric force accelerates the electron throughout its orbit. The cyclotron motion is described by Eq. (1.2) and Eq. (1.3), and the angular electron cyclotron (EC) frequency is shown to be independent of the electron's velocity, resulting in an electron that's continuously energized (or "heated") by the EM wave in this basic model.

In a tokamak plasma, the density of electrons varies spatially and their motion is in a spatially-varying magnetic field, making the dynamics complicated. The problem is typically approached through the use of ray tracing techniques, wherein the wave is propagated in




Figure 1.6: Basic illustration of acceleration of an electron (represented in blue) in a cyclotron orbit by an alternating electric field at a) a time when the electron is on the right half of its orbit and b) a time when the electron is on the left half of its orbit. Here E denotes the electric field vector and F denotes the force on the electron. Note the electron's charge is negative, producing an oppositely directed force to the electric field.

discrete segments and the wave equations are solved for the local plasma parameters for each segment. This method assumes that the plasma's variation is small for length scales on the order of the radius of an electron's cyclotron orbit, which is an accurate assumption for a tokamak with confining magnetic field of the order of a fraction of a tesla, or higher. A more complete description of EC resonance is provided in [10] and a general treatment of waves in plasmas is described in [11], but the relevant physics to provide background to this document will be described here.

Waves in tokamaks can be modeled to first order as linear (i.e. small amplitude) perturbations with Maxwell's equations in a cold, uniform plasma. Particles in the cold, uniform plasma have zero velocity and are uniformly distributed such that the plasma is quasi-neutral (i.e. without any net charge density). Corrections to the cold plasma model to account for a distribution of particle velocities are required in order to accurately account for heating and will be discussed later. The Fourier-transformed Maxwell's equations are given by

$$\begin{aligned}
\mathbf{k} \times \mathbf{B} &= \frac{-4\pi i}{c} \mathbf{j} - \frac{\omega}{c} \mathbf{E} \\
\mathbf{k} \times \mathbf{E} &= \frac{\omega}{c} \mathbf{B} \\
\mathbf{k} \cdot \mathbf{B} &= 0 \\
\mathbf{k} \cdot \mathbf{E} &= -4\pi i \rho_c
\end{aligned} \tag{1.8}$$

where $\mathbf{k}$ is the wave vector, $\mathbf{B} = \mathbf{B}(\mathbf{k}, \omega)$ is the magnetic flux density, $\mathbf{E} = \mathbf{E}(\mathbf{k}, \omega)$ is the electric field intensity, $\mathbf{j} = \mathbf{j}(\mathbf{k}, \omega)$ is the plasma current density, $\rho_c = \rho_c(\mathbf{k}, \omega)$ is the charge density, and ω is the wave frequency. Note that the equations in this section are given in the Centimetre–Gram–Second (CGS) unit system in keeping with convention.

The plasma current density is given by

$$\mathbf{j} = \sum_s q_s n_s \mathbf{v}_s \tag{1.9}$$

where a given species s (replaced by e for electrons or i for ions) has charge q_s , number density n_s , and velocity $\mathbf{v}_s$ in the plasma.

The quantities are then linearized by $A_s = A_0 + \delta A$, where A is replaced by the quantity $\mathbf{B}$, $\mathbf{E}$, $\mathbf{v}_s$ or n_s . The term A_0 gives the unperturbed quantity of A and each term δA gives a perturbation in the quantity. The cold plasma model detailed here assumes $\mathbf{E}_0 = 0$, $\mathbf{v}_{s0} = 0$ (i.e. the particles are stationary), and n_{s0} is constant (i.e. the plasma is uniform).

The equation of motion for particles in the plasma (Fourier transformed) is then given by

$$-i\omega m_s \mathbf{v}_s = q_s \left(\mathbf{E} + \frac{\mathbf{v}_s}{c} \times \mathbf{B}_0 \right) \tag{1.10}$$

where c is the speed of light in a vacuum. Taking Eq. (1.10) with the conservation of charge, given by

$$-\omega \rho_c + \mathbf{k} \cdot \mathbf{j} = 0 \tag{1.11}$$

which, along with Eq. (1.8), yields constitutive relations

$$\begin{aligned}\mathbf{E} + \frac{4\pi i}{\omega} \mathbf{j} &= \underline{\underline{\epsilon}} \cdot \mathbf{E} \\ \mathbf{j} &= \underline{\underline{\sigma}} \cdot \mathbf{E}\end{aligned}\tag{1.12}$$

where $\underline{\underline{\epsilon}}$ is the permittivity tensor, and $\underline{\underline{\sigma}}$ is the conductivity tensor. The relations given in Eq. (1.12), along with Eq. (1.8) and Eq. (1.9), allow for the definition of a linear system of equations for the wave given by

$$\mathbf{k} \times (\mathbf{k} \times \mathbf{E}) + \frac{\omega^2}{c^2} \underline{\underline{\epsilon}} \cdot \mathbf{E} = 0\tag{1.13}$$

By defining the coordinate system such that $\mathbf{B}_0 = B_0 \hat{z}$ and defining θ as the angle between $\mathbf{k}$ and $\hat{z}$, the system is described by

$$\begin{pmatrix} S - |\mathbf{n}|^2 \cos^2 \theta & -iD & |\mathbf{n}|^2 \cos \theta \sin \theta \\ iD & S - |\mathbf{n}|^2 & 0 \\ |\mathbf{n}|^2 \cos \theta \sin \theta & 0 & P - |\mathbf{n}|^2 \sin^2 \theta \end{pmatrix} E = 0\tag{1.14}$$

where the parameters S , D , P , and $\mathbf{n}$ are given by

$$\begin{aligned}S &= 1 - \sum_s \frac{\omega_{ps}^2}{\omega^2 - \Omega_s^2} \\ D &= \sum_s \frac{\omega_{ps}^2}{\omega} \frac{\Omega_s}{\omega^2 - \Omega_s^2} \\ P &= 1 - \sum_s \frac{\omega_{ps}^2}{\omega^2} \\ \mathbf{n} &= \frac{\mathbf{k}c}{\omega}\end{aligned}\tag{1.15}$$

where the plasma frequency for species s , ω_{ps} , and the signed cyclotron frequency of a given species, Ω_s , are given by

$$\begin{aligned}\omega_{ps} &= \sqrt{\frac{4\pi n_s q_s^2}{m_s}} \\ \Omega_s &= \frac{q_s B_0}{m_s c}\end{aligned}\tag{1.16}$$

It's often useful to define the parameters S and D in terms of quantities R and L , representing the right-hand-polarized and left-hand-polarized propagating waves.

$$\begin{aligned}S &= \frac{1}{2}(R + L) \\ D &= \frac{1}{2}(R - L) \\ R &= 1 - \sum_s \frac{\omega_{ps}^2}{\omega(\omega + \Omega_s)} \\ L &= 1 - \sum_s \frac{\omega_{ps}^2}{\omega(\omega - \Omega_s)}\end{aligned}\tag{1.17}$$

Finding the non-trivial solutions to Eq. (1.14) yields a quadratic equation of $|\mathbf{n}|^2$ defining the dispersion relations for arbitrary angle θ . At $\theta = \frac{\pi}{2}$ the dispersion relations reduce to the definitions of the Ordinary mode (O-mode), in which the electric field is parallel to $\mathbf{B}$, and the Extraordinary mode (X-mode), in which the electric field lies in a plane perpendicular to $\mathbf{B}$, which are given by

$$\text{O-mode:} \quad |\mathbf{n}|^2|_{\theta=\pi/2} = P,\tag{1.18}$$

$$\text{X-mode:} \quad |\mathbf{n}|^2|_{\theta=\pi/2} = \frac{RL}{S},\tag{1.19}$$

with values of $|\mathbf{n}|^2$ at oblique incidence ($0 < \theta < \frac{\pi}{2}$) linked to one of two branches of the full dispersion relation by analytic continuation. The polarization at oblique incidence of either the X-mode or O-mode is determined by selection of the desired branch and then projecting the electric field of the wave, at the plasma boundary, onto the eigenvector of Eq. (1.14) in the limit as $\omega_{ps} \rightarrow 0$. This generally corresponds to an elliptical polarization.

The cold plasma model allows for a description of the conditions for which the wave can and can't propagate. The limiting plasma parameters for a given wave to propagate are called cutoffs and are given by cases where $|\mathbf{n}|^2 \rightarrow 0$. From Eq. (1.18), the O-mode cutoff is given by $P \rightarrow 0$, while for the X-mode cutoffs are given by $R \rightarrow 0$ and $L \rightarrow 0$. Using Eq. (1.15), a typical Clemmow-Mullaly-Allis (CMA) diagram for electron cyclotron waves is plotted in Fig. 1.7.




Figure 1.7: Simplified CMA Diagram showing cutoffs in the plasma. Curves labeled *a-c* describe different propagating waves in the plasma. Note that all of the waves shown are assumed to propagate from a region of lower $\mathbf{B}$, proportional to the electron cyclotron frequency Ω_e , to a region of higher $\mathbf{B}$, as is the case for waves throughout this document. Wave *a* shows either an X-mode or O-mode wave propagating to the second harmonic resonance. Wave *b* shows an X-mode wave propagating to a cutoff, illustrating that the first harmonic resonance can't be reached in X-mode in this case. Lastly, wave *c* shows an O-mode wave propagating to the first harmonic resonance.

In order to describe resonance conditions relevant to EC heating and current drive, such as at the second harmonic (relevant to this document), the cold uniform plasma model so

far introduced requires a correction to model a distribution of non-zero particle velocities: a hot plasma. Typically, a Maxwellian velocity distribution is assumed for particles in a hot plasma, with the form (for the electrons) given by

$$f(\mathbf{v}) = n_e \left(\frac{m_e}{2\pi T_e} \right)^{3/2} \exp \left(-\frac{m_e(v_\perp^2 + v_\parallel^2)}{2T_e} \right) \quad (1.20)$$

where $v_\perp = \sqrt{v_x^2 + v_y^2}$, $v_\parallel = v_z$, and T_e is the electron temperature in units of eV. Plasma heating, which will not be derived rigorously in this document, can be modeled as a diffusion process on a distribution of electron velocities in the plasma. Since $v_\perp \geq 0$, the function $f(v_\perp)$ is a truncated Gaussian curve with its peak at $v_\perp = 0$. Each electron is assumed to have some random gyrophase (phase of its cyclotron orbit) and they exchange energy with the resonant wave depending on their gyrophase relative to the phase of the wave. This gives equal likelihood of an electron being accelerated by the wave as being decelerated by the wave, shifting $v_\perp$ up or down, respectively. The negative gradient given by $\partial f / \partial v_\perp < 0$ requires that there are more particles at lower $v_\perp$ than higher and therefore the symmetric diffusion process causes a net shift of particles to higher $v_\perp$. The assumption is made that a particle that interacts with the wave will not reenter resonance, which is generally held due to irreversible collisional processes with other particles in the plasma. The overall effect describes a plasma wave that damps its energy into resonant electrons and locally flattens their velocity distribution, heating the plasma through diffusion.

The treatment of particle motion relative to the wave requires the use of the Doppler-shifted resonance condition, given by

$$\begin{aligned} \omega &= \frac{\Omega_e \xi}{\gamma} + k_\parallel v_\parallel \\ \gamma &= \left(1 - \frac{|\mathbf{v}|^2}{c^2} \right)^{-1/2} \end{aligned} \quad (1.21)$$

where $\mathbf{k} = k_x \hat{x} + k_y \hat{y} + k_\parallel \hat{z}$, $\mathbf{v} = v_x \hat{x} + v_y \hat{y} + v_\parallel \hat{z}$, ξ is the integer harmonic number (e.g. 1, 2, etc.), and γ is the relativistic factor. The relativistic shift tends to be important in

the high-temperature plasma of a tokamak. A net current is introduced in the plasma by a resonant wave selectively heating electrons with some single-signed (either $+$ or $-$) $v_{\parallel}$ which, as shown by Eq. (1.21), requires some non-zero $k_{\parallel}$. This means that in order to drive current by electron cyclotron resonance, the wave must have some oblique steering angle ($\theta \neq \pi/2$). The theory of electron cyclotron current drive (ECCD) is explained by two competing mechanisms, the Ohkawa current and the Fisch-Boozer current [12, 13] and requires careful treatment of the wave absorption into both $v_{\perp}$ and $v_{\parallel}$ to capture correctly. These theories will not be expanded on in this document.

As stated in 1.3.1, the toroidal magnetic field is dominant in a tokamak and therefore the resonant frequency (and its harmonics) is directly proportional to toroidal magnetic field. The magnitude of the toroidal magnetic field varies proportionally to $1/|\mathbf{R}|$, with $\mathbf{R}$ defining the radial direction given in Fig. 1.3, due to the geometry of the toroidally-situated coils discussed in 1.2. This results in the resonance condition, for any particular frequency of plasma wave, being met on a surface of near-constant radial position in the plasma. The resonant surface reduces to a line in a poloidal cross-section, an example of which is shown in Fig. 1.8. The launch position used throughout this document is exclusively from the low-field side, necessitating the use of the second harmonic resonance to avoid the X-mode cutoff shown in Fig. 1.7. Launching from the high-field side typically requires more complex designs due to the constrained geometry of the inboard coil region.




Figure 1.8: A cross-section of a DIII-D plasma equilibrium (plasma shot 205049) showing the separatrix (red), the EC resonance line (at the second harmonic, blue), contours of $\psi(R, Z)$ (black), and the wall of the tokamak (green). The cross-section shown is only the right half of the tokamak since it is toroidally symmetric, as noted in Fig. 1.4. This image was generated using internal General Atomics codes (EFITviewer).

Fig. 1.7 shows that as the plasma density increases, the wave generally gets closer to cutoff for a given magnetic field strength. As noted in 1.3.1, the density of a tokamak plasma is a function of ψ and generally increases toward the magnetic axis. Typically, in order to reach the resonant surface, the launched EC wave must propagate through a plasma with

steadily-increasing density. If the wave reaches a cutoff prior to absorption, it will be reflected back out of the plasma and could cause damage to sensitive components in the tokamak wall. Additionally, as density increases, wave refraction can become strong enough that the wave's path can curve away from the intended vacuum trajectory and exit the plasma without ever having reached cutoff. This is why it's important to model the path of propagation of the EC wave even if cutoff conditions are not met anywhere inside the plasma.

One of the most desirable features of ECH systems is the control over the angle and position of the wave as it is injected into the plasma. Since, in the limit of $\omega_{ps} \rightarrow 0$ (vacuum), the wave becomes a standard electromagnetic wave, it can be directed from a distance outside of the plasma and smoothly transition to a plasma wave described in this section. This feature is illustrated in Fig. 1.7 by the waves originating at $(\omega_{pe}/\omega)^2 = 0$. Additionally, since the electron cyclotron wave operates at a higher frequency than one relevant to ions, the waveguide (or other transmission line) required to transmit the beam from the source can be much more compact and the wave can be effectively steered using quasi-optical (i.e. approximately optical) components from outside the plasma. This feature is very important for tokamaks since the power density of heating systems needs to be maximized to allow space for other components. The quasi-optical mirror system at the interface of the machine and the transmission line is called a launcher. ECH systems are also known to have highly-localized heating in the plasma, which allows for control of the profiles (temperature, density, current, etc.) of the plasma.

1.4 The DIII-D Tokamak

The DIII-D tokamak is located in San Diego, California and has operated to serve nuclear fusion experiments since the late 1980s. It is a facility owned by the United States Department of Energy and currently operated by the private company General Atomics. The goal of the experimental research programs on DIII-D is to develop plasma scenarios and enabling

technology in order to aid the development of future experimental fusion reactors [14], including ITER [15], SPARC [16], and others. Of particular relevance to this document is the development of the Advanced Tokamak (AT) scheme, to which DIII-D has made significant contributions [17, 18, 19, 20, 21]. The AT concept will be briefly discussed in this overview.

1.4.1 DIII-D Machine Overview

The DIII-D tokamak uses large copper magnetic coils to confine a plasma (typically composed of deuterium) within its toroidal volume. The toroidal volume is encased by a large metal (inconel 625) vacuum vessel for creating a vacuum environment necessary for a deuterium plasma of high purity. The vacuum vessel is protected from direct incidence of plasma by a layer of graphite tiles, composing the “first wall” of the device, for their high melting point and tendency to erode rather than melt with an incident flux of high-temperature plasma [22]. The coil systems of DIII-D, its vacuum vessel, and a representational plasma cross-section are shown in Fig. 1.9. Some important parameters of DIII-D are listed in Table 1.1 for reference [23, 24, 25].

Major radius	1.67 m
Minor radius	0.67 m
Toroidal magnetic field (peak)	2.2 T
Plasma current (peak)	2.5 MA
Neutral beam power (peak)	23 MW
Electron cyclotron heating power (peak)	3.5 MW
Vacuum level	$< 2 \times 10^{-8}$ Torr

Table 1.1: Key DIII-D parameters. The major radius is the distance from the axis of toroidal symmetry to approximately the center of the plasma. The minor radius is the distance from the plasma center to its edge at its midplane. The peak toroidal magnetic field is taken at approximately the center of the plasma region.




Figure 1.9: The coil systems, vacuum vessel and a representative plasma equilibrium of the DIII-D tokamak. A slice is taken from the complete machine to show its cross-section. The geometry is from an internal General Atomics SolidWorks (Dassault Systèmes, 2021) model.

Being that the coils are not superconducting, the stable plasma lifetime is limited by current-handling capabilities of the cooled copper coil systems and their power supplies. The typical pulse duration for a DIII-D plasma is 5-8 seconds for this reason and the typical period of cooldown between plasma pulses is 12-20 minutes.

The tokamak vacuum vessel is tessellated with port holes that provide access to the internal volume through sealed feed-throughs which are used by diagnostics, necessary for sensing and predicting the state of the plasma, and heating and control systems. The ports are situated where space allows between the toroidal and poloidal field coil systems (see Fig. 1.9 for coil definitions). They are split up into poloidal positions dictated by the shape of the vacuum vessel and toroidal positions from the zero degree line, continuing clockwise when

viewed from above. The poloidal and toroidal positions of ports around the vacuum vessel are defined in Fig. 1.10.




Figure 1.10: Left: a slice of the DIII-D vacuum vessel showing the poloidal port positions. Right: top view of the DIII-D vacuum vessel showing the toroidal port positions.

As can be seen in Fig. 1.9, DIII-D uses a cross-sectional plasma shape that is not circular. The ability to shape the plasma's poloidal cross-section is a very important aspect of DIII-D, and most experimental tokamaks, since the shape can vastly change the plasma's performance [26, 27, 21, 28]. The DIII-D tokamak also features an upper and lower divertor, a region of the first wall with increased armor to receive a high localized heat flux from the plasma. By creating a location of a magnetic X-point in the separatrix, defined as a saddle point in $\psi(R, Z)$, see 1.3.1, where

$$\begin{aligned}\frac{\partial\psi}{\partial R} &= 0 \\ \frac{\partial\psi}{\partial Z} &= 0\end{aligned}\tag{1.22}$$

the plasma at the separatrix is allowed to flow to the tokamak wall by following a magnetic

field line that tightly constrains their intersection with the first wall. The divertors are used to improve plasma performance by diverting impurities from outside of the separatrix back to the wall [4]. A DIII-D equilibrium in the “upper null” divertor configuration is plotted in Fig. 1.11 with a cross-section of the graphite first wall surface.




Figure 1.11: A cross-section of an upper null diverted DIII-D shot (126894) showing the separatrix (red), contours of $\psi(R, Z)$ (gray), and the limiting surface given by the graphite first wall (black). This image was generated using internal General Atomics codes (ECpy) that are currently under development.

At DIII-D, the development of the Advanced Tokamak concept has motivated many of the experiments and technology developments at the facility. The main features of AT plasmas include a large fraction of the overall plasma current being produced by the bootstrap current (noted in 1.3.1) and auxiliary current drive (including ECCD), rather than inductive current

drive [21, 17, 18]. High bootstrap current requires a high plasma pressure gradient, which typically causes instabilities to form, but using a technique known as resistive wall stabilization, wherein the wall acts as a feedback for the magnetic flux outside the plasma, these can be mitigated. Wall stabilization can be benefited by a broad current and pressure profile, i.e. a profile with large magnitude away from the magnetic axis, which works synergistically with a high bootstrap current since it is inherently an off-axis current driven by an off-axis pressure gradient. The high pressure in the plasma also increases fusion performance, showing that the AT concept can be an attractive option for a future power plant. AT plasmas benefit from a large off-axis auxiliary current, which can be provided by ECCD systems. The top launch configuration is a system specifically designed to drive off-axis current for AT plasmas.

1.4.2 The ECH System on the DIII-D Tokamak

The DIII-D tokamak facility uses an array of six devices, called gyrotrons, which inject EM waves into the plasma at a frequency of 110 GHz (second harmonic heating), with a total injected power of ≈ 3.5 MW [24, 29]. Each gyrotron produces between 600 kW and 900 kW at its output, using a high-voltage, high-current cathode to emit electrons which travel through a strong axial magnetic field and their cyclotron motion resonates with an electromagnetic mode inside the gyrotron’s cavity. The cavity mode is transformed to an approximately fundamental Gaussian mode at its output. The Gaussian mode is well-approximated by the quasi-optical theory of light and will be referred to as an “electron cyclotron (EC) beam” in keeping with typical nomenclature used in the field. The gyrotrons are located a considerable distance away from the DIII-D device, in order to avoid the perturbative effects of the tokamak magnetic field on their cavity magnetic field, and the beams they produce are transmitted through a series of waveguides that are evacuated and corrugated for maximum transmission efficiency [30, 31]. Each gyrotron beam is transmitted through a dedicated waveguide transmission line, of which there are eight currently available at DIII-D (six of which are connected to gyrotrons), which is made up of segments of 6061 aluminum waveguide

with an inner diameter of 31.75 mm. An overhead view of the transmission line systems is shown in Fig. 1.12. Each of these systems, including the gyrotron, its transmission line, its superconducting magnet, its high-pressure water cooling system, high-voltage power supply, and other auxiliary systems costs around \$10-20 Million.




Figure 1.12: Overhead view of a SolidWorks assembly of the DIII-D facility layout, showing the eight waveguide transmission lines in light green. The DIII-D tokamak hall is shown on the far left and the gyrotron hall is shown on the far right. The transmission lines span a distance of approximately 50 m from the gyrotron vault to the tokamak hall.

At the interface of the waveguide transmission lines and the DIII-D vacuum vessel, there are mechanical fixtures using steerable mirrors for controlling the direction of the beams into the tokamak, called launchers. The ECH system has four launchers, designed and manufactured by Princeton Plasma Physics Laboratory [32], each connecting two waveguide transmission lines to the DIII-D vacuum vessel. The launcher locations are shown in Fig. 1.13. The launchers are symmetric about the $\hat{z}$ axis of DIII-D and each launcher mirror is able to steer EC waves at varying angles, distinguished by their independent polar (i.e. sweeping vertically) and azimuthal (i.e. sweeping horizontally) values. A diagram defining the polar

and azimuthal angles is shown in Fig. 1.14. These existing launchers will be referred to as “outside” launchers throughout this document, due to their location on the outboard extremity of the vacuum vessel and to distinguish them from the “top” launchers to be discussed later.




Figure 1.13: Overhead view of a SolidWorks assembly of the DIII-D vacuum vessel and magnetic coils with the four outside launchers as installed in the R+1 row of vessel ports, at toroidal locations: 240°, 255°, 270°, and 285°. Each launcher’s waveguide inputs are shown as the two staggered cylindrical components at the top of the launchers. The naming convention of each waveguide within a launcher is given by L and R (e.g. 270L, 285R), left and right in this view, respectively. The components that interface the transmission line with the outside launcher waveguides are not pictured.

A cross-section of one side of an outside launcher is shown in Fig. 1.15 and consists of a 316L stainless steel corrugated waveguide with an inner diameter of 60.3 mm and an outer diameter of 63.5 mm. Stainless steel, with interior nickel plating for increased electrical conductivity, is used in place of 6061 aluminum due to its higher melting point in the face of the high heat flux from the plasma. At the mouth of the waveguide, the beam transitions from a cylindrical waveguide mode, HE_{11} [30], to a free-space Gaussian mode. Each launcher waveguide opens up to a focusing mirror, which then reflects the beam onto a steerable




Figure 1.14: Angles defining the polar and azimuthal directions for EC beam steering. Left: a top view of DIII-D, showing the definition of the azimuthal angle taken with respect to a line passing through the vessel origin and the ray origin, both projected onto the midplane (defined by the plane normal to the $\mathbf{Z}$ direction at the vessel origin). The polar plane is defined orthogonal to the midplane and coincident with the beam ray. Right: a view of the polar plane, showing the definition of the polar angle taken with respect to a line parallel to the DIII-D central axis (the $\mathbf{Z}$ direction).

mirror. Both mirrors are made of copper alloys for high temperature resistance and high reflectivity. Reflectivity (which gives unity when added to absorptivity) is a function of the surface conductivity of the material and (at normal incidence) is given as a fraction of the incident power by [31, 33]

$$\frac{P_r}{P_i} = \frac{4R_s}{Z_0} \quad (1.23)$$

$$R_s = \sqrt{\frac{\pi Z_0}{\sigma_s \lambda}}$$

where P_r is the power reflected, P_i is the incident power, Z_0 is free space impedance (377Ω), σ_s is the effective surface conductivity, and λ is the free-space wavelength of the beam. At an electromagnetic frequency of 110 GHz, the beam generally behaves as a quasioptical Gaussian beam [34] and its center can be assumed to follow the center of the optical path. Therefore, placing a ray at the center of the beam's path is enough to determine its direction and specular reflections can be calculated using Snell's law. The focusing of the beam is also predicted by quasioptical theory (details of which are provided in 2) and is achieved via

the curved surface of the focusing mirror. The focused beam is directed to the center of the steerable mirror which uses an armature for steering the reflected beam into the plasma. Note that each outside launcher has two of these same geometries to accommodate two waveguide inputs.




Figure 1.15: Cross-section of a DIII-D outside launcher showing the optical path of a ray traveling down the center of the waveguide and being reflected from the focusing mirror and steerable mirror. Note that the launched ray in this example is entirely in the plane of the cross-section.

There are two angles that the steerable mirror is controlled by, called “scan” and “crank”. The crank angle refers to the angular position of the armature that holds the mirror with respect to the horizontal plane of the launcher. The scan angle refers to the angle of the mirror with respect to the armature plane. A pictorial representation of these angle definitions is provided in Fig. 1.16. The mirror angles are controlled remotely via a system of linkages and straight-cut gears connected to DC motors outside of the vacuum vessel. The position of the steering mechanisms is measured via absolute rotary encoders connected to the gear systems. Real-time control (i.e. during a DIII-D plasma pulse) is enabled via a feedback control system implemented for the scan angle.

The mirrors are designed to withstand gyrotron pulses of up to 10 seconds at 1 MW of




Figure 1.16: Definition of the launcher steering angles. The facet angle is given by the angle between z and z_f , defined by the normal of the mirror armature. The tilt angle is given by the angle between z_f and n_m , defined by the normal of the mirror surface. The angle shown is at a positive facet and tilt by convention.

incident power without active cooling [32]. They achieve this by radiative and conductive (using the material's thermal inertia) cooling and the use of copper alloys capable of low EM absorptivity while withstanding high temperature. The focusing mirrors are made of Glidcop Al-15 and the steerable mirrors are made of silver-plated copper chromium zirconium (CuCrZr) with stainless steel ribs for increased structural integrity [35].

Prior to the launcher system, the beam must travel through a number of components necessary for it to safely and reliably receive power from the waveguide transmission line upstream. The necessary components for the complete DIII-D outside launch system, upstream of the launcher, are listed in Table 1.2, in descending order away from the DIII-D

plasma.

Table 1.2: Components Upstream of Outside Launchers

Component	Purpose / Description
Waveguide isolation valve	Isolates the vacuum volume of the tokamak vessel from the waveguide. When open, the vacuum vessel and waveguide transmission line are part of one vacuum volume.
High-voltage electrical break	Electrically isolates the waveguide from the vacuum vessel potential; capable of holding off 20 kV [24].
Waveguide taper	Transitions the beam from a waveguide with an inner diameter of 31.75 mm (used throughout the transmission line) to one of 60.3 mm, as used by the launcher waveguide.
Pump-out tee	Provides ultra-high vacuum pumping through a slotted waveguide open to a vacuum pump chamber. Attached turbo-pump system evacuates the waveguide nearby the vacuum vessel, prior to the isolation valve being allowed to open; capable of pumping below 1×10^{-7} Torr.

Of final important mention are the two polarizers far upstream of the outside launcher and the components listed above. These are mirrors which have grooves machined on their surface at a regular period with the groove depth giving a precise phase delay to a component of the incident wave. The polarizer mirrors are rotatable about their normal axis in order to transform the EC beam's polarization state to an arbitrary elliptical polarization incident on the plasma [36, 37]. The first polarizer acts similarly to an optical quarter-wave plate, having a groove depth of $\frac{\lambda}{8}$ (λ being the free-space wavelength), and transforms the beam's linear polarization (coming from the gyrotron) to an elliptical polarization with its ellipticity dependent on the rotation angle of the polarizer. The second polarizer acts similarly to a half-wave plate, having a groove depth of $\frac{\lambda}{4}$, and controls the inclination of the elliptical polarization. As noted in 1.3.2, accurate polarization control, along with beam steering, is necessary for ECH and ECCD deposition. The polarizer mirrors are not able to be actively controlled during DIII-D shots so their ideal rotation angles must be set prior to the shot.

The allowable steering range of the outside launch systems on DIII-D is limited by the

port bounds surrounding the steerable mirrors. The limits were determined via a visible-light laser injected through the center of the launcher waveguide and subsequently corroborated via a “hot test” procedure. The hot test involves injecting the EC beam at a paper (or sometimes mylar plastic) target placed over a given launcher port, which is observed via an infrared camera. The paper target absorbs a fraction of the power from a short-pulse EC beam passing through it and heats up to $\approx 200^\circ\text{C}$. Paper is used due to its abundance as well as its low thermal conductivity and low thickness leading to a minimal thermal distortion of the measured beam pattern. By varying the launcher mirror steering, the limits on the scan and crank angles can be found by analyzing images captured by the camera for perturbations in the beam radiation pattern caused by the port. Images from the hot test conducted in 2024 are provided in Fig. 1.17.




Figure 1.17: Infrared images from the hot test to determine steering limits for one of the transmission lines of a DIII-D launcher. Images represent the beam trajectory coming out of the port for a specific tilt and facet angle of the steering mirror. The left side image shows a beam coming out of the port unperturbed, the right side image shows the beam perturbed by the port, causing power spread due to its proximity to the port edges.

The hot test is conducted with the vacuum vessel at atmospheric pressure since necessary equipment needs to be placed inside by personnel, prior to the test (all personnel are cleared before pulsing begins). For this reason, each gyrotron is operated at reduced power and a low-pressure infill of dry nitrogen (N_2) gas is supplied to the waveguide transmission line (for

its higher breakdown voltage) through the pump-out tee, to limit arcing. With this method, pulses of up to ≈ 20 ms at ≈ 200 kW of injected power are possible. An example of results from the hot test in 2024 produced the launcher steering limits shown in Fig. 1.18 and the resulting polar and azimuthal steering limits are plotted in the DIII-D geometry in Fig. 1.19.




Figure 1.18: An example of steering limits found via the hot test method (shown for the 270L launcher mirror) showing the limitations on the launcher steering mirror scan and crank angles (in units of degrees).




Figure 1.19: The limitations on launched beam angles (polar/azimuthal) as determined by the hot test. Each colored line represents an extremity of the steering range. In the left image, only one launcher beam is shown (270L) for clarity, although each has some variation about the displayed angles. In the right image, all launcher beams, with a functional gyrotron connected to them at the time of the measurement, are shown. Note that the steering limits for the crank angle vary with the value of scan angle, and vice versa as shown in Fig. 1.18, the limits shown are the extrema of the polar and azimuthal steering angles. These images were generated using internal General Atomics codes (ECpy) that are currently under development.

1.4.3 Top Launchers

DIII-D has previously employed two “top launch” systems which injected EC beams at large polar and azimuthal angles. Rather than having a launcher assembly similar to that shown in Fig. 1.15, the top launchers used a single waveguide that entered the vacuum vessel and opened up to a Glidcop Al-15 (a copper alloy) focusing mirror that was not steerable, making its injection angle fixed. Both of these design differences were due to space limitations in the upper region of the DIII-D vessel. The mirror was surrounded on all sides by the graphite tiles of the first wall, but opened to a hole in the tiles through which the EC beam passed

into the plasma region. A view of the components making up one top launch system is shown in 1.20.




Figure 1.20: A section of the DIII-D tokamak showing the components making up a top launch system, showing the stainless steel (SS) waveguide entering the vacuum vessel and pointing the EC beam at the focusing mirror.

These top launch systems were installed in a configuration without a dedicated gyrotron, instead employing a switchable waveguide junction (called a waveguide switch) to control whether the EC beam travels down its typical launcher path, or diverts to the top launch path. The first top launch system was installed at the vessel port at 300° R+2 (see port definitions in Fig. 1.10 and outside launcher positions in Fig. 1.13 for reference). This system used a waveguide switch placed between the pump-out tee and the waveguide taper. Downstream of the junction there was a high-voltage electrical break, similar to that of the outside launch system, followed by an isolation valve. The waveguide entering the vessel was the same size (31.75 mm inner diameter) as the typical waveguide used for the transmission line, but was made of 316L stainless steel (for the same reason as the outside launcher waveguide) rather than aluminum alloy. This waveguide diameter was chosen due to the size constraint imposed by the hole in the divertor shelf shown in Fig. 1.20, although it was known to have higher power loss and risk of electrical breakdown (arcing) than the larger waveguide employed in

the outside launchers.

The top launchers injected EC beams at a polar angle of 162° and an azimuthal angle of 241° , which are shown in Fig. 1.21 (see Fig. 1.19 for comparison to outside launcher steering angles). The large polar and azimuthal angles, along with the position of the launched beam, were chosen in order to optimize ECCD efficiency in accordance with producing high-performance AT plasmas (see 1.4.1) [38, 39], the plasma physics of which are mostly outside the scope of this document. The ECCD efficiency gain can be explained heuristically by a few effects that work together. First, there is a large Doppler shift (see Eq. (1.21)), due to the large azimuthal launch angle (large $k_{\parallel}$), which allows for the EC beam to interact with high-energy, relativistic electrons. Absorption for high-energy electrons tends to be smaller so a long absorption path is needed to fully absorb the wave. This is achieved through an injection angle that is nearly parallel to the EC resonance surface (rather than passing perpendicular to it, as with outside launch). Efficient current drive is achieved by the presence of a large population of electrons that travel in a consistent toroidal direction without reversing their trajectories (called “passing electrons”) on the high-field side. These parameters result in a list of requirements for the top launch system for optimal ECCD efficiency at the desired off-axis position in the plasma. The requirements are:

1. The beam must travel nearly parallel to the resonance line (at the second harmonic).
2. The beam must resonate on the high-field side of the tokamak.
3. The launched beam trajectory must have large azimuthal angle.
4. The launched beam must be absorbed in the region of $\rho \approx 0.5 - 0.6$.

Since the line of resonance is generally perpendicular to the R coordinate in the tokamak (see Fig. 1.3), the first requirement results in a beam that must be launched from the top or bottom of the vacuum vessel, the top of DIII-D being the more convenient choice, mechanically. The second requirement brings the launch position inboard of the magnetic




Figure 1.21: View of the trajectory of the injected beam from a top launch system. Left: poloidal projection, right: toroidal midplane projection.

axis and the third requirement results in the steering of the beam shown in Fig. 1.21. The fourth requirement is due only to the desired current density profile for AT plasmas.

The legacy top launch systems demonstrated a doubling of the off-axis ECCD efficiency over the outside launch systems [39] and operated successfully after this initial result for a number of years, until one of the waveguides failed due to arcing in 2023.

1.4.4 Legacy Top Launcher Damage

Due to a runaway arcing failure, both top launch systems were fully removed from the tokamak since a complete mechanical redesign was needed to avoid subsequent failures. The arcing damage (exhibited by melting of the internal corrugations of stainless steel waveguide) appeared to originate from the end of the waveguide closest to the plasma. The damage propagated up more than 75% of the length of the 2 m stainless steel waveguide and its severity




Figure 1.22: Images of the extracted top launch waveguide. On the left is an image taken from outside the waveguide, showing its warped cross-section and discolored interior. On the right is an image taken using a borescope, showing intense damage to the internal corrugations

diminished with increasing distance from the plasma. The waveguide was also noticeably bent along its length and warped into an oval along its cross-section, likely due to softening of the material by its temperature raising significantly. Images of the extracted waveguide are shown in Fig. 1.22. The resulting damage points to a run-away arcing condition being created near the end of the waveguide closest to the plasma, which likely worsened with each EC pulse through the waveguide, causing subsequent arcing to become more likely.

The runaway arcing conditions are believed to be due to a combination of factors, none of which stands out as the single dominant cause. The smaller, 31.75mm, internal diameter of the waveguide has a higher power density for the same input power, increasing the intensity of the electric field inside, resulting in a higher likelihood of arcing. Additionally, the opening of waveguide and its mirror were necessarily close to the divertor region of the tokamak and also very nearby the second harmonic resonance layer, both of which were given by the requirements in 1.4.3. Lastly, apart from the waveguide, the focusing mirror had some melting damage on its stainless steel mount and the graphite gas break surrounding it, outgassing from which may have increased arcing probability, though it is unclear which damage preceded which.

The damage to the waveguide did not cause secondhand damage to any other components in the DIII-D vessel, but the trajectory of the top launch beam was surely affected due to the waveguide failure. This was noticed from DIII-D experiments involving top launch ECCD, which saw deviation from the designed launch angle based on comparison of computational models with experimental data. The damage to the waveguide was only finally confirmed during attempted back-calibration of the top launcher’s angle, using a laser, which actually could not be completed due to the significant bend angle of the waveguide.

Following the waveguide failures, no plans were in place to install new top launch systems on DIII-D due to the unclear path to exclude the possibility of a similar failure. An investigation was mounted into requirements for modification of the upper port (through which the waveguide previously penetrated) to support a larger diameter waveguide, but this work was likely to be prohibitively costly.

1.5 Top Launch Mirror Concept

From the testing done in 2024 to determine the port limitations (see 1.4.2), it was found that the steering limitations at the lower end of the values of polar angle (i.e. when the beam is aimed the highest in the vacuum vessel) were not due to the beam striking the port, but were due instead to a mechanical limitation on the launcher steering linkage. This mechanical limitation varied slightly for each launcher but all of them had a minimum polar angle of $\approx 60^\circ$. The polar minima were typically disregarded as unusable since the plasma doesn’t fill the region of the tokamak where the launched beam would pass, as is shown in Fig. 1.19. However, I realized that if a beam were aimed at or near its polar minimum, it could pass out of its port unperturbed and strike the first wall at a similar (R, Z) location as that of the previous top launch mirror. In doing so, the beam could propagate in a near vacuum outside of the separatrix of the plasma, strike a mirror-like feature embedded in the first wall, and be reflected along a similar path as the previous top launch system, reviving the possibility of

top launch ECCD.

Once introduced, it was immediately obvious that the same failure mode suffered by the legacy system would not be possible due to the complete lack of a waveguide at the top launch (R, Z) position. Additionally, the concept would require a much lower capital cost than either of the legacy systems. The legacy systems each required many complex components in order to mimic the implementation of the outside launchers, including a waveguide switch, a high-voltage electrical break, a waveguide isolation valve, additional waveguide, and many more ancillary components (all of which were rated for ultra-high vacuum and >1 MW of EC beam power). The total cost of design and installation of a system such as that was in the range of \$1-2M. In contrast, the concept I proposed would use the outside launcher systems as-is and require only an additional mirror and some sensors. The bulk of the cost would be in the design of the system. Since the legacy systems didn't feature a dedicated gyrotron, and neither would the new system, neither includes the cost of the gyrotron in the total cost.

Of further benefit, the concept could be replicated at low cost to every existing outside launcher system, a feature that the legacy systems could not exploit due to the lack of available vacuum vessel ports and difficulties with routing waveguide to them.

After an initial feasibility scoping, the concept was approved for an investigation under General Atomics' internal research funding. The mirror system's design, installation and initial testing will be disclosed in this document. Chapter 2 describes the design including requirements, constraints and engineering decisions made. Chapter 3 describes the installation and testing results. Finally, Chapter 4 summarizes the document and presents immediate future plans stemming from this work.

Chapter 2

Engineering Design

2.1 Constraints

First, it's worth mentioning that the time constraint placed on the project was a significant contributor to decisions made throughout its design and scoping. The initial design process was begun in April 2024 and the completed system had to be installed by the end of February 2025 at the latest (ten months), or be delayed to early 2027, due to the maintenance schedule of the DIII-D tokamak. The result of the time constraint was a general tendency to err on the side of proven technology rather than risk increased design time required for unproven technology. Engineering decisions that were influenced by this include using a fixed mirror structure (disallowing top launch mirror steering), keeping the design constrained to the existing DIII-D first wall geometry, using proven sensing technology (thermocouples), and many more. The project was allocated a budget of \$50,000 for the initial investigation and an additional \$95,000 following its approval for a complete design. A further discussion of the design constraints guiding the design of the top launch system follows.

To initialize the design of the top launch mirror, some basic constraints had to be taken into consideration. As discussed in 1.4, DIII-D (and all other tokamaks) uses a vacuum vessel to maintain a pure plasma fuel. This vacuum vessel is capable of reaching pressures as low

as 2×10^{-8} Torr and therefore any component installed inside it must be compatible with ultra-high vacuum constraints. These include the use of fasteners that do not trap a volume of air beneath them and a material that has a low outgassing rate. The plasma also imposes a large heat flux handling requirement on any features that exist in the first wall of the tokamak. The heat flux incident on a given section of the first wall varies poloidally, but is generally taken to be 500 kW/m^2 for non-divertor regions at DIII-D [24]. It is also imprudent to use any ferromagnetic material inside the tokamak, due to the extremely strong magnetic forces that would act on it.

As noted in 1.4.2, the EC system at DIII-D uses 110 GHz gyrotrons in the 1 MW class, though their power levels are generally lower than this value and transmission line losses bring the injected power down further [29]. An upper bound for the incident power for any component placed in the path of an EC beam is therefore taken to be 1 MW. This means material selection is typically limited to materials capable of withstanding high temperatures with low absorption at 110 GHz, and high thermal and/or electrical conductivity. This is especially stringent for components that are not actively cooled.

Beyond these items, the most considerable constraints come from the size of the EC beam that propagates to the mirror, the geometric constraint from the plasma volume, and the hardware considerations of the first wall of the tokamak. These constraints, combined with the necessary reflection plane for the beam along the top launch path, come together to create a final usable volume for the top launch mirror and a nominal position from which to carry out the final engineering design.

2.1.1 Beam Size Constraint

The expansion of the beam along its path places strong constraints on the geometric design of the top launch mirror system. It was necessary to fully account for the beam's propagation from the outside launcher system to the top launch mirror in order to determine the necessary size of the mirror. This section also develops a workflow used for the quasioptical design of

the mirror in later sections.

Using the optical path detailed in 1.4.2, the principles of quasioptical beam propagation are applied for a 110 GHz EC beam. In this formalism, propagation of the beam starts at the exit of the waveguide and a fundamental¹ Gaussian beam profile is assumed under the paraxial approximation of the EM wave [34]. Focusing mirrors are treated using the thin lens approximation (i.e. they have no dimension along the axis of propagation) assuming normal incidence. This process is known to closely approximate beam propagation in the far field for the frequency range around 110 GHz and an overview of the process of propagation is detailed here. After the propagation is determined, it is transformed onto the launcher optical path to retrieve the beam radius as it propagates through space.

A Gaussian beam is described as an EM wave that has a peaked intensity on its axis of propagation and which has an intensity that decreases with radius with a Gaussian (or bell curve) distribution. A Gaussian beam is given by the equation

$$u(r) = u_0 \exp\left(-\frac{r^2}{w_0^2}\right) \quad (2.1)$$

where $u(r)$ is the wave amplitude as a function of radial distance (r), u_0 is the central wave amplitude and w_0 is the initial beam radius (also known as the beam waist radius). A Gaussian distribution has finite values which extend to $r \rightarrow \infty$ so the beam radius is conventionally set by the value of r where $u(r) = u_0/e$, where e is Euler's number. The propagation of the Gaussian beam is then determined by the beam radius as a function of position along its axis of propagation

$$w(z) = w_0 \left[1 + \left(\frac{\lambda z}{\pi w_0^2} \right)^2 \right]^{1/2} \quad (2.2)$$

where z is the coordinate along the direction of propagation, w_0 is the minimum value of the beam radius along its propagation, also known as the beam “waist”, and λ is the

¹The fundamental mode is assumed throughout this document when Gaussian beams are discussed.

wavelength of the EM wave. The beam waist for an open-ended waveguide is located at its opening and is given by 32.17% of the diameter of the waveguide. This gives a good match of the wave inside the corrugated circular waveguide in the HE_{11} mode to the free-space Gaussian mode [30, 40]. With a thin lens in the path of propagation, it's useful to define $w(z)$ in the alternate form

$$w(z) = \sqrt{\frac{\lambda}{\pi \operatorname{Im}(q(z)^{-1})}} \quad (2.3)$$

where q is defined as the complex beam parameter. Defining the region upstream of the lens as region 1 and downstream as region 2, the region 1 propagation starting from $z = 0$ has complex beam parameter defined by

$$\begin{aligned} \frac{1}{q(z)} &= -\frac{z - j\zeta_c}{z^2 + \zeta_c^2} \\ \zeta_c &= \frac{\pi w_0^2}{\lambda} \end{aligned} \quad (2.4)$$

where ζ_c is known as the ‘‘Rayleigh range’’ and j is the imaginary unit. The complex beam parameter in region 2 is then given by

$$q(z_o) = \frac{j\zeta_c(1 - \frac{z_o}{f}) + d_{in} + z_o(1 - \frac{d_{in}}{f})}{1 - \frac{d_{in}}{f} - j\frac{\zeta_c}{f}} \quad (2.5)$$

where z_o is the distance along the propagation from the lens position ($z_o = 0$ at the position of the lens), f is the focal length of the lens, and d_{in} is the distance from the beam waist in region 1 to the lens. The launcher's focusing mirror surface is a small section of an off-axis paraboloid (i.e. a 3-D rotation of a parabola). In the thin lens approximation, an off-axis paraboloid has a focal distance that can be approximated by

$$f \approx \rho_p = \frac{f_p}{\cos^2(\theta_i)} \quad (2.6)$$

where ρ_p is the effective focal length of the paraboloid and θ_i is the incident angle of the

beam on the surface of the paraboloid at the intersection point. The beam is assumed to pass through the focus of the paraboloid. These definitions are shown pictorially in Fig. 2.1.




Figure 2.1: Left: Definition of the paraboloid surface of the outside launcher's focusing mirror, with dimensions in cm. The incident angle relative to the mirror normal is $\theta_i = 50^\circ$, the paraboloid focus is $f_p = 28.65$ cm and the effective focus is $\rho_p = 69.30$ cm. Right: View of the outside launcher mirror with the paraboloid cut taken from its flat surface.

A Python script was written to compute the radius of the beam as a function of distance along its propagation using Eqs. (2.3), (2.4) and (2.5) and the result is shown in Fig. 2.2. The expansion of the beam is unperturbed by the flat steering mirror in this model. In addition to determining the size of the beam at the top launch mirror, this result was used to define the power density of the beam incident on the top launch mirror's surface as well as to design the focusing contour of the top launch mirror itself, both of which will be expanded on in subsequent sections.




Figure 2.2: Expansion of the radius of the Gaussian (fundamental mode) beam as it propagates from the exit of the waveguide (at $z = 0$), reflects off of the focusing mirror (at $z = 5.1$) and reflects off of the steering mirror (at $z = 20.9$). The positions of the two mirrors along the path of propagation are indicated by the red dashed lines. Region 1 is to the left of the focusing mirror and region 2 is to its right.

2.1.2 First Wall Constraints

While it is undesirable to perturb the existing cross-sectional profile of the DIII-D first wall (i.e. the limiting surface of the poloidal cross-section) with the introduction of any new system that enters the vacuum vessel, the strict requirement for how perturbative the design could be was not known in its initial stages. Necessarily, the existing tiles of DIII-D limited the total plasma volume, but the extent to which their bounding boxes were truly the limiting surface of the plasma needed some study. The hard cutoff for the protrusion of the mirror was certainly at the separatrix of the DIII-D plasma. DIII-D is able to maintain many cross-sectional plasma shapes and shaping of the plasma is an important aspect of many experiments, further motivating a design that didn't significantly alter the limiting surface. During the development of the top launch mirror, the “shape and volume rise” (SVR) divertor was installed [41, 42], which allowed for plasma shapes with a high value

of positive triangularity (i.e. the radial ($\mathbf{R}$) distance between the centroid of the plasma cross-section and the point of highest elevation of the separatrix is almost as large as the horizontal minor radius of the plasma) and a larger overall poloidal cross-section. Together, these characteristics formed a separatrix that extended nearby the intended region for the top launch mirror, constraining the top launch mirror more than any other DIII-D plasma shape. For this reason, a poloidal cross-section of one of these shapes was taken as the cutoff for the poloidal extremities of the top launch mirror. A DIII-D cross-section showing this is given in Fig. 2.4. Importantly, the final top launch mirror design would need to account for the high plasma heat flux immediately outside the separatrix, to avoid damage. This is expanded on in 2.3.5.

Alongside this comes the constraint of the mounting surface available for installation of tiles comprising the first wall. The tiles in the region of the “upper shelf” are attached to a water-cooled plate that is stood off from the vacuum vessel. The plate has threaded studs that are welded to its bottom surface for each tile to mount to. This plate would be the mounting location of the top launch mirror as well and form another hard cutoff for the mirror’s location. The plate has multiple rows of tiles on it which span its toroidal circumference. In an effort to avoid making large, costly changes to the existing tile layouts, fitting the tile into one of these rows was considered highly preferable. The tile row constraint didn’t impose a hard cutoff like the plasma volume constraint, but was adhered to, preferentially. The mounting plate and rows of tiles are shown in Fig. 2.3.




Figure 2.3: Model view of the tiles (in gray) and water-cooled mounting plate (in purple) composing the upper shelf region. The mounting plate is installed to the vacuum vessel using the brackets at the top. Row A, B, and C were considered potential mounting positions for the top launch mirror.

2.1.3 Nominal Mirror Position Determination

The range of possible launch angles, as determined during the calibration mentioned in 1.4.2, varied slightly depending on which outside launcher mirror would be used. Therefore, the launcher mirror with the largest polar steering range in the direction of the intended mirror position was selected for the initial design study. This, combined with the computed propagation of the EC beam, provides a final hard cutoff for the usable volume in the vessel. The combination of the hardware constraints, beam constraints, and the plasma boundary come together to form the final usable volume for the top launch mirror. The usable volume is therefore a toroidally symmetric region in the upper shelf of the vessel. A cross-section of the geometric constraints and the usable volume is given in Fig. 2.4.

Most important to the successful operation of the mirror system is the constraint on the angle of the EC beam after it is reflected from the surface of the mirror. The exact value of the ideal reflection angle is based on maximizing the efficiency of ECCD in AT




Figure 2.4: Shown on the left is a poloidal projection of each of the constraining elements, including the minimum polar angle of the beam (as well as its expansion with propagation distance), the SVR plasma region, and the extremities of the water-cooled mounting plate. Shown on the right is a poloidal projection of the usable region produced by these constraints, taken together.




Figure 2.5: Model view initial concept mirror placed in row ‘B’ of the upper shelf tiles of DIII-D. The envelope of the EC beam is shown in translucent green, with graphite tiles (various colors) surrounding the black top launch mirror.

plasmas (see 1.4.3) and is a function of the position of the launched beam itself. For this reason, the launched angle of a legacy top launch system was used for this initial scoping study, as this was thought to be within a reasonable range of the (yet to be computed) ideal angle. The launched ray therefore used a 162° polar angle and a 241° azimuthal angle for the initial study. This reflected angle, combined with the incident angle coming from the outside launcher, provides the required angle of the mirror surface for a specular reflection. By positioning various mirror geometries into the 3-D model of the upper shelf region, it was found that the mirror would necessarily be installed in the tile row marked ‘B’ in Fig. 2.3, as a position in row ‘A’ would protrude outside of the usable volume. With this as the nominal mirror position, and the feasibility of the geometry determined, more detailed engineering design could begin. The initial concept of the top launch mirror that was spawned from these constraints is shown in Fig. 2.5. The position of the top launch mirror along the axis of propagation was found to be about 100 cm, giving a beam radius incident on the mirror of about 4.5 cm.

2.2 Conceptual Study

Prior to establishing a final design, some further conceptual studies were performed in order to fill the remaining informational gaps and further detail the intended operation of the top launch mirror. These include:

1. Determination of the material makeup of the mirror.
2. Determination of the reliability and precision of the outside launcher system's aiming to the mirror center.
3. Determination of the sensor selection for the mirror.

2.2.1 Mirror Material Study

For maximal reflectivity at 110 GHz, as mentioned in 1.4.2, the most desirable mirror material is a metal alloy with high electrical conductivity (e.g. copper alloys like those used for the outside launcher mirrors). However, as mentioned in 2.1.2, the mirror's implementation was constrained to avoid heavily altering the makeup of the first wall of DIII-D. This constraint carried over to the material choice, since it would be undesirable to introduce a new source of impurities, other than carbon from the graphite of the wall, to the plasma. The outside launcher mirrors, meanwhile, are not generally in direct contact with the plasma for sustained periods and thus don't introduce copper impurities to it. For this reason, if it could be deemed a capable reflector for an EC beam, the same material as DIII-D's first wall would be the most desirable selection for the mirror. The first wall of DIII-D is made of ATJTM graphite, which is a grade of graphite preferred for its high strength, density and machinability [43]. It was therefore necessary to determine graphite's behavior as a reflector for a 110 GHz EC beam to ensure it would not overheat and result in a structural failure.

At DIII-D, many millimeter-wave systems assume an absorptivity (see definition in 1.4.2) of graphite of 8%–10%, while an absorptivity of $\approx 5\%$ is reported in [44] for 140 GHz and

theoretical calculations using Eq. (1.23) find an absorptivity of 2%–3%².

Due to this large range of graphite absorptivity figures, and the lack of many published findings on it, a study was conducted using a vector network analyzer (VNA) to help determine what value could be used for thermal modeling. The VNA setup is shown in Fig. 2.6 and consists of the two VNA heads, two tapers which transform the wave from (to) a fundamental mode in a rectangular waveguide (WR8 size) to (from) an HE_{11} mode in a 31.75 mm circular corrugated waveguide, and a miter bend which reflects the beam at a 90° angle. The miter bend uses a mirror at a 45° angle to the incident and reflected beams. The setup was first used to measure the loss from a copper miter bend mirror, as a reference for the measurement, then an ATJTM graphite mirror whose dimensions matched those of the copper mirror. The measurement setup is known to have a systematic error owing to the use of the miter bend and its HE_{11} mode conversion [30, 36]. Due to this setup, the loss measurements can be separated into a component of power lost in the tapered mode converters (due to mode conversion) and a component of power lost in the mirror itself. Separating these components would require more sophisticated analysis and the overall loss measurements were taken to be a reasonable upper bound for the loss due to a given material’s reflectivity.

The mean loss (taken over multiple measurements with the miter bend in the E- and H-plane) across the graphite miter bend was found to be 4.8%, which closely matched the result found in [44]. Full results of the testing can be found in Table A.1 in the appendix. The graphite miter bend was tested with and without surface polishing and the results showed no measureable difference, which was a positive outcome since polishing the mirror’s surface would add to its manufacturing time. This figure of 4.8% was then used to determine the thermal characteristics of a mirror made of graphite. This was done using a simplified thermal model in the commercial software SolidWorks, in which an approximately Gaussian distribution of thermal power was applied to the surface of a rudimentary model of the mirror in order to determine the maximum pulse length for a given gyrotron. A surface temperature

²Note that these figures assume an incident angle normal to the surface and the absorptivity can be assumed to be $\approx 2\%$ higher for a 45° incident angle based on [44]




Figure 2.6: Test setup using VNA. The tapers shown transform WR8 waveguide to 31.75 mm with a highly pure HE_{11} mode output, and vice-versa. The test mirror is either a copper miter bend mirror (used throughout DIII-D ECH transmission lines) or a graphite disk matching the mirror's geometry.

limit was set at 2000°C , based on a general safety limit used at DIII-D [22]. More precise thermal modeling is reported in 2.3.5, but was not necessary for the initial design scoping. The following assumptions were made in this thermal model:

1. The absorptivity remains constant, at 4.8%, with temperature.
2. The beam is in a 100% fundamental Gaussian mode, incident on the surface.
3. The incident beam carries 900 kW of power.
4. A constant 500 kW/m^2 is incident on the surface of the mirror to model heat flux from the plasma (far away from the divertor).

The first assumption is valid for the temperature range of $0\text{--}2000^{\circ}\text{C}$ based on the change in electrical resistivity as a function of temperature, since, for graphite, resistivity initially drops as temperature rises, up to about 1000°C and then rises linearly back to its nominal (room temperature) value of $12 \times 10^{-6} \Omega\text{ m}$ near 2000°C [45]. Thermal modeling assuming

this value of electrical resistivity remains constant is therefore overly conservative, which was desirable for the study. The second assumption was a practical decision taken to avoid complicated power distributions and remains accurate for a beam with a reasonable match to a free-space Gaussian mode ($>\approx 80\%$). The assumption is commonly used in quasioptical design. The third assumption was based on the maximum power of any single gyrotron system at DIII-D, though was known to be an overestimate due to transmission line loss [29]. Setting the incident power at 900 kW would provide a reasonable “worst-case scenario” for the mirror and a factor of safety of about two for the simulation since it was approximately double the injected power of the gyrotron system intended to be used for the top launch mirror. The fourth assumption is valid for DIII-D’s graphite armor tiles outside of the divertor region, as noted previously [24].

The thermal model was based on the power distribution across a Gaussian beam, see 2.1.1, which results in the equation for the power across a circular aperture of radius r

$$P(r) = P_0 \left[1 - \exp \left(-2 \left(\frac{r}{w(z)} \right)^2 \right) \right], \quad (2.7)$$

where the total incident power is given by P_0 and the beam radius as a function of position along the propagation axis, as defined in 2.1.1, is given by $w(z)$ [34]. The power absorbed by the graphite is then found by multiplying Eq. (2.7) by the approximate absorptivity of 4.8%. However, since the incident beam’s propagation axis is not normal to the surface of the mirror, but instead approximately 45° from normal, the power across an elliptical aperture must be used. This is given by a projection of the circular aperture onto a plane at 45° , which results in the absorbed power across an elliptical aperture of major radius a and minor radius b

$$P(a, b) = (0.048) P_0 \left[1 - \exp \left(-2 \left(\frac{ab}{a_0(z) b_0(z)} \right) \right) \right] \quad (2.8)$$

where the elliptical projection of the beam radius as a function of position along the




Figure 2.7: Elliptical distribution of power density absorbed by the graphite mirror at 4.8% absorptivity, split into ten discrete regions for the SolidWorks thermal simulation.

propagation axis has major radius $a_0(z)$ and minor radius $b_0(z)$. A Python script was written to split up the beam's power into ten different elliptical apertures, using Eq. (2.8), and a constant heat flux (power per unit area) was applied to the surface of the graphite mirror for a given pulse duration in SolidWorks' thermal simulation suite. The material of the mirror was defined using an internal General Atomics material setting for ATJ graphite and its initial temperature was set to 21 °C. In order to have a safety margin in the engineering design, there was no cooling or thermal sink applied in this model; the mirror was assumed to be in a complete vacuum with no ability to transfer heat either through conduction or radiation. The power density for each elliptical aperture is plotted in Fig. 2.7.

The results of the thermal model showed that the mirror would stay below the maximum temperature of 2000 °C for about 5.5 seconds. Given that a DIII-D plasma pulse typically doesn't exceed 8 seconds, the pulse limit was deemed acceptable. All of the parameters used in these assumptions were subject to some variation, so it was decided that the pulse length of the EC beam would be gradually extended (from about 50 ms, up to full pulse length) while monitoring the temperature of the mirror via thermocouple temperature sensors (see

2.3.4).

2.2.2 Outside Launcher Aiming

In SolidWorks, the geometry of DIII-D’s vacuum vessel, graphite first wall, and outside launchers was assembled and the steering angle for the beam was found by tracing a ray to be incident with the center of the nominal mirror location determined in 2.1.2. This yields the required outside launcher aiming (see 1.4.2), which was found to be about 63.6° and 191.6° polar and azimuthal, respectively. These angles would need to be measured precisely in situ (see 3.1.2), and they would also change somewhat as the design of the mirror matured and its center shifted slightly, as will be discussed in subsequent sections.

The reliability of the outside launcher’s aiming is partly determined via the software used for controlling the scan and crank angles of the mirrors. Hardware uncertainties and hysteresis effects (i.e. having different behavior in one direction, compared to the other) would be measured during installation, but the recent launcher calibration (see 1.4.2) provided measurements for the accuracy of the launcher hardware for the design. The launcher control software uses a feedback loop with DC motors and rotary encoders which report the scan and crank angles of the launcher mirrors in units of “counts.” The control loop is programmed to an accuracy of ± 10 counts, such that the setting is deemed acceptable if the scan and crank counts are each within that range of their respective commanded values. Due to the gearing of the outside launcher steering systems, the crank angle has a much finer adjustment than the scan angle, since the scan angle is optimized for fast, real-time steering feedback control. For the crank angle, this results in a much larger span (in counts) from one end of the steering range to the opposite side. The uncertainty of ± 10 counts in the crank control corresponds to a negligible error in the polar and azimuthal angle of the launched beam. Meanwhile, an uncertainty of ± 10 counts in the scan angle corresponds to angular uncertainties of about $\pm 0.25^\circ$ in the polar angle and $\pm 0.04^\circ$ in the azimuthal angle³. The uncertainty of $\pm 0.25^\circ$

³Note that, although the scan control typically corresponds to the polar beam angle and the crank control

in the polar angle results in a shift of about ± 0.5 cm. Given the incident beam radius of ≈ 4.5 cm, this was deemed acceptable.

2.2.3 Sensor Selection

Due to the constraint imposed on the maximum graphite surface temperature in 2.2.1, it was necessary to record the temperature of the mirror. DIII-D lacks an infrared camera capable of capturing the top launch mirror in its field of view. This meant sensing the temperature of the surface of the mirror and surrounding tiles would require another solution. Installation of a rudimentary infrared camera system capable of viewing the top launch mirror was initially considered as a part of the top launch mirror system installation, but was abandoned due to the high cost. Other camera systems on the DIII-D tokamak typically use costly mechanical periscope systems and require a dedicated port for the viewing of the desired region of the vacuum vessel. This was outside of the scope of the top launch mirror installation and a rudimentary camera system was found to be unable to adequately view the top launch mirror.

Likewise, an initial investigation was conducted into the use of surface eroding thermocouples (SETCs), which are a type of thermocouple that have a tip that is designed to be eroded by the high heat flux of the plasma while offering a high-accuracy measurement of the surface temperature of the tile into which they are installed [41]. The concept was unfeasible for the top launch mirror due to the long lead-time for the purchase of graphite-tipped SETCs and the uncertainty in their ability to handle and accurately record temperature of the incident high-power EC beam.

There are many tiles in DIII-D's first wall that use a thermocouple embedded into the face of the tile opposite the plasma, which, although this requires a thermal model to determine the temperature throughout the tile (mainly at its surface), was considered the most feasible concept to replicate. Additionally, the outside launchers at DIII-D use type E (chromel-constantan) thermocouple temperature sensors to understand the temperature ratcheting

typically corresponds to the azimuthal beam angle, they are not completely decoupled from one another.

behavior of the mirrors and linkages after successive pulses, so an existing electronics rack was available with the necessary hardware for data logging inside the DIII-D machine hall. For these reasons, embedded type E thermocouples were selected as the sensing hardware for the top launch system.

2.3 Final Design

With the constraints and feasibility established, the final design would refine the design created in the conceptual study. The final design required the determination of the optimal steering angle for the EC beam as reflected from the top launch mirror in order to finalize the mechanical design of the mirror. With that complete, the final mechanical design of the mirror and surrounding tiles could be completed, followed by the detailed mirror focusing design and the sensing setup. Finally, more complete thermal and mechanical analyses were completed to ensure the mirror's integrity was maintained.

2.3.1 Optimal Reflection Angle Determination

The aiming of the top launch beam was determined via an optimization of the current drive efficiency to achieve maximum off-axis current drive for a reference AT plasma scenario, a process that was done by my General Atomics colleagues [38, 39, 46]. The process uses a ray tracing software called TORAY [47] that uses the parameters of the EC beam starting from its initial position and direction injected into the plasma. It then models the initial beam as a bundle of rays with a Gaussian distribution of ray intensity (see 2.1.1), and matches the spread of the rays to the far-field Gaussian beam propagation. Notably, since it assumes a far-field beam propagation, TORAY cannot properly model a beam waist located in the plasma like a quasioptical code, such as TORBEAM [48], can. At the time of development, the TORBEAM code had not been fully integrated with the DIII-D geometry, so the analysis could not be completed using TORBEAM on the short timescale necessary for the completion of the design

of the top launch mirror. This meant that the optimization of the top launch mirror aiming was limited to a somewhat columnated beam assumption. Modeling of the optimal focusing parameters is discussed in 2.3.3. The parameters included in the aiming optimization were the polar and azimuthal angles of the launched beam, offering a 2-dimensional contour plot of the current drive efficiency as a function of the aiming angles. The optimization used DIII-D shot 179173 (a relevant AT plasma) during its flat-top phase, which has an on-axis toroidal magnetic flux density of $B_\phi = 1.6$ T. The study began with a range of polar and azimuthal angles, optimizing for the maximum current drive at the desired $\rho = 0.5$ contour. Once the region was located, it was further scanned for a minimum deposition width (i.e. the localization of the driven current across ρ), which typically has a Gaussian shape in ρ . This allows the current drive to be highly localized at the desired value of normalized poloidal flux: $\rho = 0.5$. The results of the study are plotted in Fig. 2.8 with an overplot of peak deposition position in ρ and the peak driven current. The plot shows a scan of azimuthal angles from 226° to 236° and polar angles of 140° to 150° . For the desired peak absorption location of $\rho = 0.5$, the total current driven is ≈ 33 kA per MW of EC beam power with little sensitivity to changes in the aiming angles, which is desirable in case of small misalignments, since the contour of ECCD is very similar to the contour of ρ in which it absorbs. The optimization resulted in the final top launch aiming used for the mirror reflection: a polar angle of 146° and an azimuthal angle of 231° , which are shown in Fig. 2.9




Figure 2.8: Left: Contour plot of aiming angles (polar, azimuthal) for the nominal top launch mirror position vs the driven current (I_{ECCD} normalized to 1 MW of injected EC power). Right: Contour plot of aiming angles (polar, azimuthal) for the nominal top launch mirror position vs the value of ρ at which peak absorption occurs. The desired $\rho = 0.5$ contour is plotted in red. Data courtesy of X. Chen (General Atomics).




Figure 2.9: View of the designed injection path of the top launch beam after its reflection off the top launch mirror. Left: poloidal projection, right: toroidal midplane projection.

2.3.2 Mirror Mechanical Design

The geometry of the mirror was designed entirely around the required angle of the beam incident on its surface and reflected from it, which uniquely determines a flat plane in 3-D space using a specular reflection. From the flat surface, the incident Gaussian beam (as determined in 2.1.1) was projected onto the surface of the plane, resulting in an elliptical mirror face. The determination of the position of the center of the mirror face, and therefore the incident angle, required considerations of the allowable geometry of the machined part, including the thickness required for mounting hardware and obstructions in the mounting plate in DIII-D. Therefore an iterative process was required. For each iteration, the mirror's center was placed at an initial location, the beam was projected onto its surface, and the geometry required for mounting it was assembled. The angle of the mirror surface was particularly stringent to design around because of the thickness required for the mounting hardware. Likewise, if the thickness was increased to give more space for mounting hardware, the furthest protruding surface might exit the usable volume defined in 2.1.3, depending on the size of the mirror face. This process would shift the center of the beam slightly, resulting in another iteration. The initial mirror center was taken from the position determined in 2.1.2, which was also used to determine the optimal top launch beam angle in 2.3.1. During this process, it was assumed that the optimal top launch angle would change only slightly, despite the repositioning of the mirror center (and therefore the beam's initial position). Since the adjustments were relatively small, the iterative technique was able to reach a workable geometry, which was then used for a second optimization (repeating 2.3.1) to yield the final optimal top launch aiming. The angle of the mirror surface was then slightly altered to match the optimal aiming angle.

Another caveat to this design is, as mentioned in 2.1.1, the Gaussian beam's radius used is that dictated by convention, which contains $\approx 86.5\%$ of the total power of the Gaussian beam. In order to capture as much of the power of the beam on the surface of a quasioptical component, the component is typically sized a few times larger than the conventional Gaussian

beam. However, a mirror of that size couldn't fit into the usable volume described in 2.1.3, so the size of the mirror was another free parameter in the iterative design process detailed above. The larger the mirror's face, the more it had to protrude from the existing DIII-D first wall, but the more power it would be capable of effectively steering. After several series of iterations, the beam radius projection used to define the mirror surface was $1.35w$ where w is the radius of the Gaussian beam at the top launch mirror's position along its propagation. This results in $\approx 97\%$ of the power in the Gaussian beam being incident on the surface of the top launch mirror, leaving the other $\approx 3\%$ of the power to be scattered by graphite tiles surrounding the top launch mirror. With an injected power of 900 kW, this leaves ≈ 30 kW of EC power to be dispersed throughout the DIII-D vacuum vessel volume, which was deemed acceptable. The mirror geometry was finalized and given to the DIII-D "machine integration" group so that its geometry could be standardized with their existing armor tile designs, with review provided by their machine shop for ease of manufacture [49]. The resulting final mechanical design of the top launch mirror is shown in Fig. 2.10. Some notable features of the design include the elliptical geometry conforming to the shape of the incident beam, as dictated by its trajectory, as well as the plane determining the "conical cut" (see 2.3.3), which is normal to the plane of the elliptical surface and parallel to its major axis. This final mechanical design resulted in beam reflection angles of 145.9° and 230.9° polar and azimuthal, respectively, in the SolidWorks assembly, which were a close match to the optimal angles found in 2.3.1.




Figure 2.10: Final mechanical design of the top launch mirror. On the left is an elevated view from its side, showing the cutouts for mounting hardware and shiplaps, which are regions designed for adjacent tiles to overhang. On the right is a front view, showing the elliptical shape of the mirror surface.

2.3.2.1 Mirror Mounting Design

Graphite is a relatively brittle material, making its mounting to the hard metal walls of the vessel a delicate process. Most tiles at DIII-D follow a standard method of mounting that uses a nut bar that secures the tile from a pocket machined into the tile that provides a “foot” that contacts the vacuum vessel of DIII-D (or other mounting surface) [20]. An image depicting this mounting design is shown in Fig. 2.11. With this scheme, the stainless steel nut bar is positioned away from the plasma, preserving the high heat flux from the plasma for the graphite surfaces alone. Holes in the front face of the tiles allow a hex key tool to secure the nut bar to the studs that are welded to the vacuum vessel or, in the case of the tiles in the region around the top launch mirror, to the water-cooled plate (see Fig. 2.3). The nut bars typically use a semi-circular cross-section, with a matching circular cutout in the tile, along with layers of GraFoil [®] [50] (compliant layers of graphite) sheet underneath

the nut bars to allow for their mounting without stress concentrations. Pre-loaded spring washers are also typically used on the nut bars to ensure they don't become loose. GraFoil sheets are also typically placed underneath the tile feet to limit stress concentrations from the metal wall and to provide a large heat transfer surface.



Figure 2.11: Standard nut bars used for mounting tiles in the DIII-D vessel. Image depicts nut bars installed on the water-cooled plate in row B (see 2.1.2), with tiles uninstalled.

The top launch mirror's reflecting surface was very close to the positions where holes would typically be drilled for the hex key tool to access the nut bars. In order to avoid EM power from passing into mounting holes, and to avoid perturbing the beam with an irregular hole in the mirror surface, the nut bars used for the top launch mirror were designed to be installed from the side (rather than the top) with a crescent wrench. Their overall thickness was also reduced and their cross-section was flattened in order to limit the thickness of the void needed for them, which allowed for necessary reduction of the mirror's protrusion from the first wall. This design influenced the iterative mechanical design discussed in 2.3.2. The top launch mirror's nut bars are shown in Fig. 2.12. Since the nut bars needed to be secured

from the side, the top launch mirror needed to be secured in place prior to any surrounding tiles to ensure side access.




Figure 2.12: Nut bars used for the top launch mirror. On the left is an elevated view of the full nut bar and on the right is a cross-section of a nut bar. Note the flattened, low-profile design and the hex nuts (orange) able to be fastened by an open-ended wrench, as compared with the traditional style seen in Fig. 2.11. The spring washers (green) are made of 302 stainless steel and the flat plate (dark gray) is made of 304 stainless steel, as are the hex nuts. The GraFoil is shown underneath the nut bars in light gray in the left image.

2.3.2.2 Surrounding Tiles

The tiles surrounding the top launch mirror were redesigned in order to maintain a gradual toroidal transition from the recessed nearby tiles that make up the first wall. The gradual transition is necessary since the hot particle flux from the plasma in the edge region, just outside the plasma separatrix, must be spread out over as much surface area as possible. The flux of particles is primarily in the toroidal direction, with a slight poloidal angle toward the wall [19, 20]. This means that a sharp toroidal transition in the first wall creates a localized region of high heat flux and a potential for the tile to overheat and fail. Since the top launch mirror protrudes from the pre-existing first wall, and its toroidal extent is limited, the tiles that surround it must slope outward to meet it.

The tiles composing the upper shelf region of DIII-D aren't completely axisymmetric about the torus, due to ports and other irregular tiles. Therefore, the design of the surrounding tiles marks one of the few aspects of the top launch mirror system that isn't agnostic to the

selection of outside launcher, and thus toroidal position. For this reason, it was important to have a specific outside launcher selected for the final design at this stage. The 270L launcher position was selected (see Fig. 1.13) for its ability to reach the necessary polar angle found in 2.2.2, it having an active gyrotron connected to it, and for its positioning nearby the 300° R+2 port for the sensor feed-throughs (see 2.3.4). This choice was relatively arbitrary, however, and any of the outside launchers with an active gyrotron would have worked as well, some with modifications needed to increase their polar steering range.

The final design of the surrounding tiles is shown in Fig. 2.13. In the figure, tiles 1a and 1b are mirror images of each other and meet the top launch mirror at its extreme point of protrusion towards the plasma. Tiles 2a and 2b are also mirror images of each other and are double-wide tiles to account for the increased width of the top launch mirror over the tiles that it replaced. Tiles 3a and 3b were modified with cutouts for the EC beam to pass.




Figure 2.13: Final design of the tiles surrounding the top launch mirror. The image on the left shows the tiles in the “upper shelf” region with tiles highlighted in blue denoting those that were replaced by the top launch mirror and its surrounding tiles.

The final design of the mirror and its surrounding tiles allowed a DIII-D limiter to be made with their poloidal extremities included. This is pictured in Fig. 2.14 with the usable volume from 2.1.3 overlaid.




Figure 2.14: The poloidal projection of the top launch mirror's extremities (pink) inside the usable region (transparent red, see 2.4), with the poloidal cross-section of the mounting plate drawn in blue.

2.3.3 Mirror Focusing Design

The focusing surface of the top launch mirror was determined using the same quasioptical technique as was presented in 2.1.1. Previously, the program determined the propagation starting from the exit of the waveguide to the outside launcher focusing mirror and beyond. Adding the top launch mirror was as simple as locating it at the correct z location along the path of propagation and adding it to the computation. Equation (2.5) is used again, but d_{in} in this case is given by the distance from the beam waist in region 2 to the top launch mirror. This gives the complex beam parameter, from which the propagation of the beam radius downstream of the top launch mirror, region 3, is again determined by Eq. (2.3). The surface of the top launch mirror is, as is the surface of the outside launcher focusing mirror, a section of an off-axis paraboloid. Its design uses a similar definition as shown in Fig. 2.1,

with the plane in which the parabola is defined depicted as “Mirror Conical Cut Plane.” In the case of the top launch mirror, the major axis of the elliptical beam projection onto the mirror was found to be slightly outside of the plane of incidence of the beam (defined by the incident and reflected rays). This is a mistake made during some iteration of the workflow and was only noticed after the mirror was manufactured. The plane defining the parabola for the paraboloid cut was therefore not exactly the same as the plane of incidence. However, because the difference was relatively small ($\approx 7^\circ$), this was shown to have little effect, if any, on the resultant beam propagation (see 3.2.1).

With the workflow to design the focusing characteristics of the beam complete, the paraboloid focal distance ρ_p was designed. The design required a beam waist at the optimal position in the plasma. As mentioned in 2.3.1, the process used a modified version of the TORAY code again (see 2.3.1), which allowed for the ray propagation to more closely match that of a columnated beam at the position of peak absorption. The model was used to minimize the deposition width (in ρ), which found that a highly columnated beam with a small radius was optimal. This determined that the beam waist would be optimally located at the position of peak absorption along the beam’s propagation path. That position was predicted to be ≈ 65 cm from the top launch mirror along the path of the beam, fixing the final parameter required for the quasioptical design. The effective focal length of the top launch mirror was found to be $\rho_p = 40.9$ cm and the resulting beam propagation is plotted in Fig. 2.15.




Figure 2.15: Propagation of the radius of the Gaussian beam from the exit of the waveguide (at $z = 0$), with reflections from the focusing mirror (at $z = 5.1$), the steering mirror (at $z = 20.9$), and the top launch mirror (at $z = 101.6$). The beam waist in the plasma is at $z = 166.6$, 65 cm from the top launch mirror. The three mirrors are indicated by the red dashed lines. Region 1 and 2 are the same as described in 2.1.1 and region 3 models the propagation downstream of the top launch mirror. The distance from the beam waist in region 2 to the top launch mirror is given by d_{in} .

2.3.4 Sensing Design

A majority of the sensing design was completed by my colleagues at General Atomics and based on existing hardware and software used by the DIII-D ECH system [51], so this section will provide a brief discussion of their work.

As noted in 2.2.3, type E thermocouples were selected as the sole sensor system for the top launch mirror. The embedded thermocouples are spring-loaded to maintain thermal contact with the blind hole into which they're installed. The spring is set into position with the use of a nut, made of graphite, screwed into a tapped hole in a given graphite tile. A diagram depicting this setup is shown in Fig. 2.16. The open port that previously housed the legacy 300° R+2 top launcher was chosen as the location for the thermocouple wires to exit the DIII-D vacuum vessel, due to its availability and close proximity to the top launch mirror. The wires exit through a vacuum feed-through after which they route to a data

logging computer in the existing electronics rack inside the DIII-D machine hall.




Figure 2.16: Drawing view of a hole in the tile for an embedded thermocouple. The left image shows the bottom of the tile where the hole into which the thermocouple is installed (marked “D”) and the line “B-B” marks the cross-sectional cut shown on the right, which shows the profile (side view) of the hole in the tile foot. Image courtesy of the DIII-D Machine Integration group (General Atomics).

The top launch mirror was designed to include four embedded thermocouples and four of the surrounding tiles would include two embedded thermocouples each, bringing the total to twelve. The four surrounding tiles selected to include thermocouple sensors were chosen for their potential to be directly heated by the EC beam from the outside launcher, in the case of a gross misalignment. Two were included in each surrounding tile for redundancy and the top launch mirror included four for further redundancy and in order to have the possibility of sensing the location of the EC beam on the mirror. The arrangement of embedded thermocouples in the top launch mirror and its surrounding tiles is shown in Fig. 2.17. The remaining two surrounding tiles were mostly out of the possible beam path and therefore weren’t given thermocouples. The mirror’s thermocouples were arranged in such a way that it was possible to ensure their tips were approximately equidistant (about 1 cm) from the surface of the mirror. They were also spaced about 4-6 cm from the center of the mirror to avoid placing holes at locations of high heat flux from the EC beam, since their

mounting holes create weak points in the graphite. This was done with thermocouples of various lengths due to the thickness variation of the mirror.




Figure 2.17: Wireframe (semi-transparent) model view of the top launch mirror and its surrounding tiles, from a perspective looking at the front surface of the tiles. Embedded thermocouples were placed at the positions shown in red.

Thermocouple installations inside the vacuum vessel at DIII-D use wiring harnesses made of ceramic for its high-temperature and ultra-high vacuum ratings. This was the same for the top launch thermocouples, for which an off-the-shelf 12-pair wiring harness was selected. The wiring harness features ceramic thermocouple connectors on one end, and termination into a ceramic DB-25 (25 pins) connector. Each thermocouple also has a ceramic connector matching the wiring harness. For the feed-through flange, a DB-25 ceramic feed-through on a ConFlat flange was selected to connect to the 300° R+2 port. The existing port flange needed to be adapted from a $3\frac{3}{8}$ in ConFlat flange to a $4\frac{1}{2}$ in ConFlat flange in order to connect the DB-25 feed-through. Outside the vacuum vessel, the thermocouple wires would use typical plastic insulation and would be routed to the electronics rack, where their readings would be digitized and recorded.

An existing user interface to view the data from outside launcher thermocouples was updated in order to view the digitized readings from the top launch thermocouples. The readings were added to a separate tab in the LabView program and an image of the top launch mirror and surrounding tiles was placed in the window with color-coded temperature indicators to show the regions experiencing the highest temperature. Additionally, the temperatures of all of the top launch thermocouples are plotted over time in the same tab. An example of the user interface developed for the top launch mirror system is shown in Fig. 2.18. Due to excess EM noise during DIII-D plasma shots, the thermocouples are programmed to disconnect ten seconds prior to the shot and reconnect shortly after the shot is completed. This doesn't greatly affect the data they receive since the rise time of the temperature readings typically takes longer than a DIII-D shot duration, as is noted in 3.2.2.




Figure 2.18: Graphical User Interface (GUI) view of the arrangement of embedded thermocouples in the top launch mirror and surrounding tiles. The thermocouples marked T1-T4 are embedded in the mirror, while T5-T12 are embedded in the surrounding tiles. This image was taken after the initial installation of the tiles and T6 failed during installation, showing no temperature data. Image courtesy of P. Nesbet (General Atomics).

Beyond the sensors for the mirror itself, there was significant effort made to program interlocks for the top launch mirror system, to prevent damage to it, the gyrotron systems,

and DIII-D. These include an interlock for the aiming setting of the outside launcher mirror, which disallows gyrotron pulses in the top launch configuration if the outside launcher steering mirror is not aimed precisely at the top launch mirror, or if its aiming is somehow perturbed. An interlock was also made to prevent gyrotron pulses if the on-axis toroidal magnetic field was at or above 1.9 T, which could cause electrical breakdown at the surface of the mirror and damage it. Lastly, an interlock was also programmed to prevent gyrotron pulses if the mirror was above some threshold temperature, which was to be determined during testing. Notably, since the thermocouples are disconnected during a DIII-D shot, the temperature interlock could not stop a gyrotron in real-time (and thanks to the rise time of the thermocouple readings, they wouldn't be effective at doing so anyway). Therefore, the interlock only limits the temperature at which a gyrotron can be programmed to start a pulse, prior to the DIII-D shot.

2.3.5 Thermal and Structural Analysis

With the final design of the mirror, its surrounding tiles, and the sensing setup complete, a more accurate thermal model could be created. The updated model would aid in understanding the material limit in more detail than the 2000 °C limit described in 2.2.1. Further, at this stage in the mirror's development, the protrusion of the mirror past the existing DIII-D first wall was known, allowing a model of the heat flux due to the edge of the plasma to be made using existing codes.

The same assumptions from 2.2.1 were carried forward into the updated thermal model, in SolidWorks, including the incident power of 900 kW (again, a factor of two larger than the incident power of the intended gyrotron system) since this stage of analysis still wished to determine a “worst-case scenario” for the thermal and structural performance of the mirror. With the completion of the mirror's mechanical design and the final incident beam angle determined, a similar Gaussian power distribution was applied on the mirror's surface. Heat transfer from the mirror was also not considered in this model in order to understand the

absolute limits of the mirror’s structural integrity, again in a “worst case scenario.” A thermal model with increased fidelity was required in order to accurately model the thermocouple data and is reported in 3.2.2.

The results of this model were similar to those of the initial SolidWorks simulation, in that the surface temperature of the mirror reached 2000 °C in approximately the same amount of time (≈ 5 s of EC beam power). The results of the thermal model were then connected to a SolidWorks structural simulation to show the stresses created in the mirror by the temperature differential across its volume. DIII-D imposes limits on the tensile and compressive stresses of its ATJ tiles of 2618 psi and 6650 psi, respectively. The structural model showed a peak tensile load of 77 % of the limit throughout the duration of the pulse. However, the compressive load reached its limit after about 4 s of applied power. A plot showing results of the thermal and structural modeling is shown in Fig. 2.19. This places stronger limitations on the pulse length than the simplistic 2000 °C limit, but would still provide enough duration for useful top launch applications [52]. The limit is discussed again upon further analysis in 3.2.2, where the updated thermal model includes more realistic incident EC beam power and cooling.

Alongside the other heat sources applied in this study, an additional source was introduced that came from an analysis of the approximate heat flux from the edge of the plasma. This analysis was done by my colleagues at DIII-D using the HEAT [53, 54] plasma heat flux code, so the details of the analysis will not be expanded on in this document. The study used the worst-case operational plasma parameters: an SVR shape (which brings the separatrix closest to the top launch mirror) at full plasma heating power during an edge-localized instability (called an ELM), in order to understand the upper bound of potential plasma heating. The results of the study are shown in Fig. 2.20. The simulation showed the largest heat flux on the surrounding tile inboard of the top launch mirror, but the magnitude of said heat flux was not of concern based on comparisons made to tiles in the divertor region, therefore the surrounding tiles were left out of the thermal model. The resulting heat flux on the top




Figure 2.19: Results of the SolidWorks thermal model applied to the graphite top launch mirror. The top image shows the surface temperature at the pulse length limit (4 s), with an elliptical Gaussian distribution. The left image shows the tensile stresses throughout the mirror, with a peak value of 1,616 psi (11.1 MPa) at some sharp corners and the right image shows the compressive stress throughout the mirror, with a peak value of 6,484 psi (44.7 MPa), which is approximately the limit for compressive load of graphite tiles.

launch mirror was applied in the thermal model by defining a region matching the shape found in Fig. 2.20, with a constant peak heat flux applied. A constant heat flux was also added across the entire plasma-facing surface of the mirror (500 kW/m^2 , just as in 2.2.1) to approximate general plasma heat flux.

The additional heat from the ELM model of plasma flux resulted in a more distributed temperature across the surface of the mirror, but the overall peak temperature was not greatly affected, nor was the structural limitation, since the regions of heat application had little




Figure 2.20: Simulation of plasma power density incident on the surface of top launch mirror and surrounding tiles due to a full power ELM. Numbers in white are in units of MW m^{-2} . Image courtesy of A. Wingen (Oak Ridge National Laboratory).

overlap. The thermal model satisfied the requirements to ensure that it would not fail due to the heat flux incident on its surface.

An analysis was also completed to ensure that electromagnetic forces induced in the mirror would not cause it to fail or be ripped from its mounting plate. The analysis was done by my General Atomics colleagues [55] using established internal processes for modeling forces created in the event of a plasma disruption (a rapid loss of plasma confinement, resulting in the plasma current flowing into the first wall of the tokamak [4]). A failure of this sort, although unlikely, would not only result in the loss of all top launch functionality, but also the possible temporary halting of all DIII-D plasma operations if the debris is thrown into a vital section of the vacuum vessel. In this analysis, the full DIII-D peak plasma current (at peak toroidal magnetic field) was allowed to flow into a reduced model of the top launch mirror and/or its surrounding tiles and the stresses were analyzed with a structural model in

the commercial software ANSYS. Since the analysis was done entirely by my colleagues, no detailed discussion will be presented here, but it is noted for its importance to the approval for installation of the top launch mirror system. The final results of the study showed that the stresses due to disruption forces were within the limits imposed by internal DIII-D review.

Chapter 3

Installation and Testing

3.1 Installation

The top launch mirror and its surrounding tiles were manufactured using a specially-equipped 4-axis CNC (computer numerical control) mill which is used at General Atomics for the machining of many of its ATJ graphite tiles. Since graphite dust is abrasive and electrically conductive, graphite requires a dedicated machining volume and milling bits. The surrounding tiles were a relatively simple job for the mill, but the top launch mirror, with its more complicated geometry, needed special care in order to create its curved faces. The billet of ATJ graphite was fixtured onto a few different jigs in order to machine the desired geometry. An image of the mirror during the machining process is provided in Fig. 3.1. Polishing of the mirror surface was not necessary, as was found in 2.2.1, but it was machined with a smaller bit to bring its dimensions into tolerance and its surface roughness to the point of lacking visible tool marks.

The installation began with routing of thermocouple wires in-vessel and ex-vessel, then continued with the laser alignment of the mirror, followed by the installation of surrounding tiles.




Figure 3.1: Image taken during the machining of the top launch mirror on the CNC mill.

3.1.1 Thermocouple Installation

The installation of the thermocouples began with the routing of the wiring harness inside of the DIII-D vacuum vessel. Prior to routing, all the tiles in the region were removed from their studs and sorted into tote boxes for later reinstallation. This included the tiles that were replaced by the top launch mirror and its surrounding tiles, which were put into storage. The routing of the thermocouple wiring harness started from the feed-through flange, where the length necessary to connect to the flange was measured and fit up. The wiring harness was then secured to the port flange using stainless steel shim stock (spot welded in a loop around the wires). This is the standard method of securing wires inside of the DIII-D vacuum vessel and was used to secure the wires along their entire path, whether the mounting surface was the port, the inconel vacuum vessel, or the tile mounting plate. Each tile had its thermocouples installed on a work bench outside of the vessel, as the small parts were easier to pre-install there. The ceramic thermocouple connectors were labeled with shim stock using an engraving tool prior to their installation. By labeling them, the wires could be connected to their respective thermocouple after they were routed. Each connecting end of the wiring harness was left secured temporarily nearby their intended tile locations,




Figure 3.2: Wiring path taken in-vessel for the top launch thermocouples (TC). The path starts at the 300° R+2 port on the top right and follows a gap underneath the tiles to the top launch mirror and its surrounding tiles. Note that accurate port hardware is not shown in this view.

awaiting connection. Each wire of the harness (a total of 24) was routed underneath gaps in the DIII-D first wall tiles, designed for cable channels. They were fully covered by the tiles upon their re-installation. A model view showing the in-vessel thermocouple routing is given in Fig. 3.2.

Outside the vacuum vessel, the port flange adapter was installed on the 300° R+2 port and the feed-through flange was connected to the ceramic DB-25 connector from inside the vessel, then installed in place. The thermocouple wires were then routed from the port flange to the electronics rack, a length of approximately 18 m. Each of these wires needed to be cut and terminated to connect to the thermocouple digitizers in the rack. With the wires routed, their signals were checked in the GUI software (see 2.3.4) and the four thermocouples for the top launch mirror were installed to begin laser alignment.




Figure 3.3: Image taken of the side of the top launch mirror during installation. It shows the cutout in the foot of the tile to clear an obstruction on the water-cooled plate as well as a nut bar installing it to the plate. It also shows the copper shim stock used to adjust its angle for laser alignment.

3.1.2 Laser Alignment

The installation of the top launch mirror required a precise match to the designed top launch beam angle, which would be approximated with the use of a visible light laser, reflected from an ultra-thin visible light mirror placed on its surface. Prior to the laser alignment, however, an initial fit-up test showed that the mirror needed a minor modification to its foot in order to attain a flush contact region with the plate. The modification was prompted by a small protrusion of the mounting plate that the foot of the mirror intersected which was not captured in the SolidWorks assembly. The modifications required the removal of about half of the thickness of the foot at its outboard extremity, a small change that would not affect its performance. An image of the modification is given in Fig. 3.3. With the modification made, the mirror was installed in place using its custom nut bars and the laser alignment process began.

The laser used for alignment was a ThorLabs PL201 green diode laser. It was installed into a fixture that was used for the test described in 1.4.2 to determine the aiming limits of

the outside launchers. The fixture mounted to the outer diameter of the outside launcher waveguide outside of the vacuum vessel and was centered using a technique involving the rotation of the fixture around the waveguide while minimizing the variation of the laser's position at the end opposite the laser source. The copper alloy outside launcher mirrors were polished during the 2024 laser alignment (see 1.4.2) and had retained a high degree of polish that allowed them to visibly reflect the laser light from their surface. The same was not true for the graphite top launch mirror, so an ultra-thin optical mirror was secured to its center with tape so that it could be aligned. Errors incurred by the introduction of a flat mirror onto the curved surface of the top launch mirror were mitigated by the small size of the optical mirror relative to the curvature of the top launch mirror. The steerable mirror, which is flat, was also fitted with a small visible light mirror for the test, to increase the visibility of the reflected laser. It's also important to note that hysteresis of the steerable mirror was tested for and measured during this phase. It was found that if the mirror's motion in the scan (see 1.4.2) direction acts to move the launched beam to a lower polar angle (i.e. further up in the DIII-D vessel), then its position obeys the commanded signal much more reliably than the opposite direction. The hysteresis was found to influence the aiming by about 2° at the extreme low end of the polar span, which is where it aims for the top launch configuration. For this reason, a control was put in place to ensure that the top launch aiming was only ever approached from the reliable direction.

The alignment of the top launch mirror was an iterative procedure that used an existing coordinate measurement machine (CMM) commonly used for in-vessel installations. The CMM measured two points along the beam of laser light to create a chord in the 3-D coordinate space inside the vessel and that chord was converted to polar and azimuthal angles. After each measurement the required direction of adjustment was determined. The mirror was then loosened from its mounts slightly and a piece of copper shim stock was placed under its feet on the side necessary for adjustment. The mirror was then re-tightened into place and the process was repeated. After about five total adjustments, the angles were measured

to be $145.8(1)^\circ$ and $230.5(1)^\circ$, which were within $\pm 0.5^\circ$ of the final aiming angles determined in 2.3.2, the target precision.

With the mirror installed in its final place, the hot test to verify its alignment using the EC beam could begin. The surrounding tiles would be installed after the completion of the hot test.

3.2 Testing

There were three stages of testing to verify that the top launch mirror was properly installed and that it was functioning as expected. The first stage was a test to verify the top launch beam angle and its propagation using the actual EC beam, rather than a laser. In the second stage, tests were carried out to compare with the thermal model and to adjust its parameters in order to accurately predict the temperature throughout the mirror. In the third stage, tests were performed with the EC beam injected into curated DIII-D plasma shots. These included testing to further corroborate the thermal model with plasma heat flux as well as top launch operations to verify the deposition location and polarization inside of the plasma. The third stage was largely carried out by my colleagues at DIII-D and their work is summarized here.

3.2.1 Aiming Verification via Hot Test

In order to corroborate the laser measurements in 3.1.2 and the focusing design of the top launch mirror, a “hot test” procedure was performed. The hot test procedure was similar to that described in 1.4.2, involving the use of a short-pulse, low-power EC beam from the gyrotron to measure the beam path inside of the DIII-D vacuum vessel. The top launch hot test, however, included an additional component: the measurement of the beam at multiple positions along its path of propagation. This was done in order to verify the focusing mirror’s quasioptical design and precisely measure the angle relative to the laser’s angle.

A device for measurement of the beam along its propagation path is employed at DIII-D

for the acceptance testing of new gyrotrons in order to verify their beam’s behavior and align them into their respective waveguide transmission lines. The device consists of a frame made of sections of aluminum t-slot framing on which a rail system is installed that holds a paper target in the path of the beam. The rail is precisely aligned with a laser system onto the desired axis of propagation of the gyrotron’s beam. With the fixture in place, an infrared camera is mounted to the frame that is used to view the paper target as it is heated by the beam. The infrared camera must be located out of the path of the beam (to avoid causing damage to it), so its perspective must be accounted for in image processing. This is done using an infrared calibration target that maps the surface of the paper into a 2-D plane for precise determination of its size and position on the paper. Measurements are then made at multiple positions along the path of the beam’s propagation and, from the relative size and position of the beam on the paper its propagation can be fit to a quasioptical propagation at a particular angle relative to the reference laser [56]. The target is moved along the rail remotely using a motor and worm drive system. This method is typically used at DIII-D to determine the Gaussian mode content of the beam and measure the coupling to the HE_{11} mode of a waveguide. In this case, the code used for determining the mode content, developed by my General Atomics colleagues [57], was repurposed for the measurement of the Gaussian beam’s angle relative to the laser and its waist position along its propagation into the plasma.

The measurement fixture is typically installed horizontally in a relatively large room where the gyrotrons are installed. It had never been installed inside of the DIII-D vacuum vessel and required modifications to its geometry to not only fit and be assembled inside of the vessel, but also to be aligned to the rather extreme polar and azimuthal angle of the top launch beam. An image of the fixture installed inside of the DIII-D vessel is shown in Fig. 3.4. After it was installed and aligned, camera calibration measurements were taken, personnel were cleared from the DIII-D machine hall, and the test was carried out.

Results from the hot test showed that the beam alignment was a close match to the angle of the laser, deviating by $+0.8 \pm 0.1^\circ$ in the polar direction and $+0.2 \pm 0.1^\circ$ in the




Figure 3.4: Image taken during assembly of the hot test fixture inside the DIII-D vacuum vessel. The laser path shows where it reflects from the outside launcher mirror on the left, to the top launch mirror at the top, then to the target on the right.

azimuthal direction. This resulted in a final measured beam aiming of 146.6° and 230.7° , polar and azimuthal, respectively. This result, although slightly outside the desired precision of $\pm 0.5^\circ$ mentioned in 3.1.2, was deemed acceptable. The beam waist was measured along the propagation direction with a shift of 3 cm upstream of the theoretical beam waist position and its radius at the beam waist was measured as 0.09 cm less than the theoretical prediction. A plot comparing the propagation that was measured to the theoretical propagation is provided in Fig. 3.5. Additionally, near the end of the hot test, the outside launcher aiming was altered slightly by ± 10 encoder counts, the uncertainty mentioned in 2.2.2, in the poloidal direction and images of the beam were taken at the focal point. The resulting images showed no visual deviation of the beam center and therefore quantification of the change of the beam center position was not possible within the accuracy of the measurement, which demonstrated reliability of the mirror's aiming to within the accuracy of the outside launch mirror controls (see 2.2.2).

With the laser alignment complete, the surrounding tiles were installed and the full




Figure 3.5: Propagation of the radius of the Gaussian beam from the top launch mirror compared with the measured values from the hot test. The theoretical beam waist position was $z = 166.6$, 65 cm from the top launch mirror, while the measured position was $z = 163.5$.

installation is pictured in Fig. 3.6.




Figure 3.6: Image of the finalized installation of the top launch mirror and its surrounding tiles.

3.2.2 Thermal Model Calibration and Pulse Extension

At DIII-D, gyrotrons are allowed to pulse into the vacuum vessel without a plasma present in order to test various components of the transmission line system prior to, or after, a plasma pulse. The pulse length for a vacuum pulse is administratively set to 250 ms in order to limit damage to vital components due to beam scattering. During the DIII-D startup phase, wherein low-power plasmas are produced in order to prepare and test functional systems for full-power operations, the top launch mirror was tested in between plasma shots for short pulses (≤ 250 ms). Then, during a later startup phase, pulses were allowed during DIII-D shots allowing the pulse length to be successfully extended to 6 s. Since the thermocouples are embedded in the mirror, approximately 1 cm from the surface of the mirror (see 2.3.4), their time evolution needed to be predicted with a thermal model. The thermal model, created previously (see 2.3.5), was insufficient to properly simulate the temperature response of the thermocouples and how their readings translate to temperature throughout the tile and on its

surface. An updated model was created allowing a maximum pulse length to be determined through this. The updated model and parameter tuning is detailed here.

The thermal model was constructed using the finite element method in ANSYS, since SolidWorks' simulation suite did not have as many options for configuration of the simulation, which were necessary for fitting the data. As in 2.3.5, the assumptions listed in 2.2.1 were carried forward into this model, but some parameters were tuned in order to fit the data and the plasma heat flux was fit more accurately to the state of the discharge (zero for vacuum pulses). Just as in the SolidWorks thermal model, the projection of the beam power on the surface of the mirror was determined using Eq. (2.8) and again split up into elliptical apertures defined by the incident angle of the beam on the mirror. Initially, an incident beam power of about 600 kW was used (rather than the previous 900 kW, see 2.2.1 and 2.3.5), to better fit the peak power of the gyrotron system used [29]. A constant temperature condition of 25 °C was specified on the face that would make contact with the water-cooled mounting plate inside DIII-D, to simulate the cooling it provides. This is a standard method used to model tiles that are connected to this plate and approximates a large heat sink such as this. No radiative cooling was simulated as it was assumed to have a small effect, although could become more important at higher surface temperatures. An initial temperature of the entire volume of the mirror was also applied at 25 °C. Temperature readings were simulated by probing the inside surface of the blind holes into which the thermocouples were mounted. This yielded initial results which were compared with data taken from multiple gyrotron pulses.

Initially, gyrotron pulses of <50 ms were tested to verify the thermocouple response. From there, the pulse length was gradually increased to 250 ms and the response was recorded. The initial results for the overall change in temperature ($\Delta T \approx 3^\circ\text{C}$) were below the expected value from the thermal model ($\Delta T \approx 6^\circ\text{C}$) for a 250 ms pulse length (depending on the starting temperature and the particular thermocouple). Additionally, the thermocouple temperature evolution over time differed from that predicted in the thermal model. This

prompted a reassessment of the ANSYS setup to bring it into a closer match.

A source of the disagreement between the readings and the simulations was recognized in the thermal contact and heat transfer from the graphite to the thermocouples. Therefore, the thermocouples themselves were added to the thermal model as representative structures placed inside of their blind holes, in thermal contact with the graphite. The heat transfer parameters could then be tuned to better match the behavior of the data. Likewise, the constant temperature condition of 25 °C for the feet of the mirror assumed the contact patch for the mirror was the entire foot and it was in direct contact with the mounting plate, which wasn't the case. So instead, the sheets of GraFoil between the mounting plate and the mirror were added to the model and a heat transfer coefficient was applied between the mirror feet and the GraFoil, given by 2000 W/(m²°C). A constant temperature condition of 25 °C was applied to the bottom GraFoil surface instead of the mirror feet. Lastly, the incident beam power was tuned in the thermal model from the assumed value of 600 kW. In total, three independent parameters were tuned in order to best match the thermal model to the thermocouple data:

Table 3.1: Parameters tuned to match the thermal model to thermocouple data

Parameter	Description
G_{wall}	Heat transfer coefficient between the mirror and GraFoil.
G_t	Heat transfer coefficient between the mirror and the thermocouple.
P_i	Incident power of the EC beam.

The primary effect of G_{wall} was on the temperature evolution of the simulated thermocouple signals after the simulated beam pulse, which effectively tuned the cooling rate during the latter portion of the simulation. On the other hand, G_t mainly affected the temperature evolution during the simulated beam pulse, with some influence on the peak overall temperature reached by each thermocouple. This is because G_t directly affected the rate of heat transfer to a given thermocouple, not throughout the bulk of the mirror. Lastly, P_i largely influenced the

peak temperature reached by the thermocouples, globally. Note that, rather than P_i , the absorptivity of the mirror could be equivalently tuned at this stage since the thermal model used the power that the mirror absorbed for the heat flux on its surface. Quantification of the error between the thermocouple measurements and the simulation was not necessary for the desired accuracy of the simulation and the small change in thermocouple temperature provided some uncertainty in the measurements themselves. Instead, the peak temperature of each thermocouple was compared to the simulation, as well as the temperature gradients during and after the gyrotron pulse. The parameter values that best matched the thermocouple data were found through repeated simulations in ANSYS. They are reported in Table 3.2.

Table 3.2: Tuned parameters used in the thermal model.

Parameter	Value
G_{wall}	2000 W/(m ² °C)
G_{t1}	250 W/(m ² °C)
G_{t2}	50 W/(m ² °C)
G_{t3}	400 W/(m ² °C)
G_{t4}	800 W/(m ² °C)
P_i	320 kW

Notably, the incident power was found to be lower than the value that was calibrated and measured for the gyrotron at the time [29], at about 520 kW. The absorptivity of the mirror at high temperature could also be subject to some variation, which could influence the power density used in the model, equivalently with injected power. At 520 kW, the equivalent absorptivity would be 3%, rather than 4.8%. Radiative cooling and graphite’s reduction in electrical resistivity with temperature [45] may also contribute to the lower measured thermocouple temperatures, though with 250 ms pulses this effect was unlikely to be significant. As an alternative to tuning the incident power, tuning only the heat transfer coefficients to lower values allowed the peak temperature in the simulation to match the thermocouple data, but the temperature gradient during the pulse was found to fall significantly lower than the data indicated and therefore these parameters were deemed unsatisfactory for the model. The uncertainty in the mirror absorption from the thermocouple data suggests a high-power




Figure 3.7: Thermocouple temperature readings compared with simulated thermocouple readings for a 250 ms pulse of the EC beam. These results use the parameters given in Table 3.2 with the image on the left using the nominal Gaussian beam radius at the top launch mirror position found in 2.1.1 and the right image using a reduced radius (85% of the nominal value).

absorption study with a direct surface temperature measurement would benefit the system, but this was not possible at the time (as noted in 2.2.3). Lastly, there was uncertainty in the position of the EC beam on the mirror due to the assumption of a perfect Gaussian beam in the exact center of the mirror. The uncertainty in position was not studied in detail since the effect of the thermocouple heat transfer was found to be dominant in the thermal simulation¹. The uncertainty in the beam size was the final parameter which could strongly influence the thermal model results and, given that the model's purpose was to predict the maximum allowable gyrotron pulse duration, the radius of the beam was reduced from its predicted value in order to increase the power density on the mirror. This resulted in a larger peak temperature and a safety factor for the simulation. The beam radius used was 85% of the value predicted in the model discussed in 2.1.1. Fig. 3.7 shows a comparison of the simulated thermocouple response to the measured thermocouple data for a 250 ms pulse length.

As mentioned in 2.3.4, the embedded thermocouples are disconnected during plasma

¹This also meant that using the thermocouples to sense whether the beam was correctly centered on the mirror was ineffective, barring large misalignments (evidence for which was never seen).

pulses, so top launch pulses longer than 250 ms have a black-out period during and slightly after the plasma pulse. This was not an issue for the top launch mirror data since the plasma pulse (<10 s) tended to be shorter than the rise time of the thermocouple temperatures (typically $>\approx 15$ s), especially as the pulse length was increased due to the higher peak temperatures reached. For pulses with a plasma present, the heat flux due to the plasma was reintroduced as in 2.3.5 as a constant thermal flux across the surface of the mirror. Depending on which DIII-D shot was analyzed the value of the thermal flux varied, but a value of 300 kW/m^2 was found to approximate plasma heating for these low-power plasmas quite well². The pulse extension study showed that the reduction in incident power by $1/3$ resulted in longer possible pulses than initially expected. The mirror only reached a surface temperature of $\approx 1100^\circ\text{C}$ after a 6 s pulse, according to the simulation. The temperature rise behavior was compared against the thermal model at multiple gyrotron pulse lengths and the tuned parameters were found to approximately match the peak temperature reached, while deviating somewhat from the measured temperature gradients. This could indicate that the heat transfer coefficients needed to be tuned lower for the longer pulse cases, which could be explained by a number of thermal mechanisms, but the parameters weren't adjusted at this stage. The thermocouple data from a 6 s EC beam pulse (during a DIII-D shot that lasted 7 s) is compared to the simulated thermocouple response in Fig. 3.8.

Also during the DIII-D plasma startup phase, the plasma is run in an upper null and lower null configuration to introduce higher heat flux on the upper and lower divertors (see 1.4). During this phase, the X-point is translated radially ($\pm \hat{R}$, see Fig. 1.3 for definitions) in a sweeping motion over the course of the shot, which translates the “strike point” (the intersection point of field lines of the separatrix with the wall) over the divertor region. In this phase, the strike point is typically swept over, or very near, the location of the newly-installed top launch mirror. However, heat flux at the strike point is much larger than the heat flux simulated in 2.3.5, since the plasma at the X-point is tightly constrained to

²These startup plasmas had their separatrix far from the top launch mirror as well. The combination of this with their low temperature resulted in a lower heat flux than a typical DIII-D plasma.




Figure 3.8: Thermocouple temperature readings compared with simulated thermocouple readings for a 6 s pulse of the EC beam. These results use the parameters given in Table 3.2. The data is from DIII-D shot 202815.

flow along the diverted field lines [4]. For this reason, the strike point was not swept over the top launch mirror during this phase, but it was brought nearby in order to test for its thermal response. Fortunately, strike point sweeping in the top launch mirror region was not technically necessary for the plasmas that would be run in the DIII-D operational campaign at the time. During these sweeps, the gyrotron was not pulsed, meaning the only heating of the top launch mirror was due to the plasma flux. The thermal model was compared to the thermocouple data again, with the same parameters as those used previously but without the heat flux due to the gyrotron beam and with the addition of increased heat flux due to the plasma. The heat flux due to the plasma was tuned to match the thermocouple data. At the corner of the tile, a heat flux of between 2.5 MW/m^2 and 5 MW/m^2 was applied, in addition to a heat flux of 600 kW/m^2 across the full plasma-facing surface of the mirror. The heat flux at the corner of the tile matched a similar distribution as is shown in Fig. 2.20. This case represented the mirror's highest peak temperature found during testing at about




Figure 3.9: Thermocouple temperature readings compared with simulated thermocouple readings for DIII-D shot 202780. During this shot, the strike point was swept nearby the top launch mirror. These results use the parameters given in Table 3.2.

1570 °C at the mirror’s corner. The results of the thermal model alongside the thermocouple data are shown in Fig. 3.9.

The same parameters used to simulate the 6 s gyrotron pulse on the top launch mirror were re-enabled in addition to the high heat flux case from the plasma. However, there was no intention of running EC power in a plasma with the strike point near the top launch mirror, so this was only done to generate a worst-case scenario. The results showed a peak overall temperature that was similar: 1643 °C. The structural integrity of the mirror was checked with this result in ANSYS and it was found to maintain thermal stresses just under the limits defined in 2.3.5. This same shot simulation was run in succession four times with a delay between shots of 12 minutes, a typical minimum cooldown period between DIII-D shots, in order to assess the temperature ratcheting behavior of the top launch mirror. The mirror was allowed to cool down through the simulated water-cooled mounting plate between shots and the peak temperature on the surface of the top launch mirror during the fourth shot was 1657 °C, which showed that thermal ratcheting was not significant, even in this worst-case scenario.

Based on the thermocouple data captured, and the simulations, a real-time limit was set by my colleagues at General Atomics [58] on the proximity of the separatrix and the top launch mirror in order to maintain a distance of 5 cm between them. If the distance falls below 5 cm for more than 100 ms, the plasma shot is halted. Additionally, a “safe-to-pulse” (to control whether a gyrotron pulse can begin) interlock was implemented at 85 °C for any of the thermocouple readings, which was an overly conservative interlocking point based on the data gathered throughout the pulse extension phase, but was subject to change if necessary. Notably, this temperature interlock was never limiting during any DIII-D shots that used EC heating with top launch. This represents the final control and interlocking logic implemented for the top launch mirror system as it moved toward regular plasma operations.

3.2.3 Plasma Deposition and Polarization Experiment

While the previous testing during the pulse extension phase was possible to perform in a DIII-D startup plasma, with plasma absorption characteristics insensitive to the study’s goal, testing to determine the operation of the beam in the plasma necessitated a dedicated experiment. This experiment spanned two separate DIII-D half-day runs, one on June 23, 2025 and another on July 7, 2025. The goals of the experiment were to verify the polarization calculation for the launched beam into the plasma and to verify the deposition location in the plasma. A measurement of the driven current was also completed during these experiments, though this was a lower priority measurement as results had already verified the TORAY models for the current driven by previous top launch systems [39]. The measurements were completed in large part by the same General Atomics staff as had carried out the legacy top launch experimental validation [59, 60] and the methodology and results were similar to those experiments. They will be summarized here.

Computation of the polarization transformation at both the outside launcher steering mirror and the top launch mirror was required prior to the experiment. This computation was done by my colleagues at DIII-D [61] by comparing the plane of the specular reflection

at each mirror to that of the mirror preceding it, in SolidWorks. For the steering mirror, the polarization transformation comes from the angular difference of the reflection plane at the outside launcher’s focusing mirror, which is in the poloidal plane (see Fig. 1.15), and its own reflection plane, an angle of about 8° . Then the angular difference between the plane of reflection for the top launch mirror and that of the steering mirror, which is about 22° , gives the polarization transformation at the top launch mirror. These two transformations are used to compute the polarization required in the plasma for X-mode or O-mode launch (arbitrary elliptical polarization), and to program the two in-line polarizers (see 1.4.2). This polarization computation is done by the TORAY code and the new angular transformations were set to fixed values in the code for the top launch configuration.

Despite confidence in the angular transformation computations, an experimental polarization check is a standard part of commissioning a new launcher system at DIII-D and verifies the ability of the launched beam to couple to either the X-mode or O-mode waves. The polarization experiment started by recreating a reference plasma shot (DIII-D shot 179173) with relatively low electron temperature (an on-axis electron temperature of ≈ 3 keV), since the absorption of the O-mode wave is more strongly dependent on electron temperature, making it easier to distinguish it from X-mode absorption. The shots for this portion of the experiment had an on-axis toroidal magnetic flux density of 1.6 T, which is near the optimal magnetic field setting for top launch in AT plasmas, and an on-axis electron number density of about $3 \times 10^{13} \text{ cm}^{-3}$. The top launch EC beam was injected into the plasma for a number of DIII-D shots with the polarization settings for the beam being altered between shots. This allowed the experimenter to measure the power absorbed in the plasma and locate the polarization at which peak absorption occurs using DIII-D’s EC emission (ECE) diagnostics. The X-mode fraction was varied between target values of 80%, 90%, and 100%, calculated by the TORAY code. With correct polarization, the beam showed stronger absorption at 100% X-mode than an impure O-mode/X-mode mixture. DIII-D shots were also run with 100% O-mode and the absorption was compared with those at 100% X-mode. The results showed




Figure 3.10: Experimentally measured wave absorption for polarizations in the X-mode and O-mode (second harmonic). Note that the label “partial profile” refers to profiles of absorption that extend outside the values of ρ at which measurements were made with DIII-D’s ECE diagnostics. The rough uncertainty of these measurements is around 20%. Image courtesy of X. Chen (General Atomics) [60].

generally lower absorption for the O-mode shots than X-mode, as expected, and data from multiple such shots is plotted in Fig. 3.10.

In order to determine the absorption location in the plasma (by the value of ρ at which peak absorption occurs) and how it varies with the resonance location, additional similar DIII-D shots were run with different values of toroidal magnetic field, but similar parameters otherwise. This was done to further verify the aiming direction and that the beam’s propagation in the plasma matched the TORAY model. For this purpose, the on-axis toroidal magnetic field was varied in magnitude between several DIII-D shots and the peak absorption was measured and compared against predictions made using TORAY. The results matched well with the TORAY computations, corroborating the model’s deposition location predictions. Plotted in Fig. 3.11 are results comparing TORAY and experimental data for various DIII-D shots with on-axis toroidal magnetic flux density settings varying between values of 1.5 T and 1.8 T.

The ECCD measurement is more challenging than the other two measurements performed during the experiment. The measurement requires a Fourier analysis of power-modulated




Figure 3.11: Comparison of results for the predicted peak deposition location found using TORAY and results from the experimentally measured peak deposition location for multiple plasma shots. Image courtesy of X. Chen (General Atomics) [60].

plasma shots to isolate the plasma response to the EC power. The measurement is also sensitive to noise and experimental uncertainties because the off-axis plasma current is driven over a larger volume than on-axis current, meaning that the current density is relatively low. Measurements were performed using DIII-D's motional Stark effect polarimetry system [62] to measure the plasma's internal magnetic field. These measurements are used to reconstruct the profile of the plasma current and, when the inductive current profile is subtracted from it, the non-inductive component of the plasma current is found. Analysis of the local driven current is ongoing and preliminary results are provided in [60].

Chapter 4

Conclusions and Future Work

The top launch mirror system was successfully conceived of, designed, manufactured, installed, and tested within the span of about one year with an overall cost of about \$150,000, staying approximately on budget. The initial concept came from an overlooked launch trajectory available to the pre-existing outside launcher systems and was proven feasible in an initial study. Its design involved careful consideration of the constraints placed on it by the geometry of the DIII-D first wall in addition to the material performance of a plasma-facing quasioptical component for a high-power EC beam. Also found during the design process were the optimal beam aiming and focusing as the beam propagated from the outside launcher and was reflected specularly from the top launch mirror. As a part of the effort, a thermal model was developed to predict the surface temperature and maximum thermal stresses in the mirror, which required a code to be written to map the power distribution of the incident beam based on its ideal quasioptical propagation. The model was refined in multiple phases of the design and applied to the temperature data from embedded diagnostic thermocouples. The quasioptical model and installed aiming were verified via a hot test with the gyrotron and the thermal model was corroborated and its parameters were tuned to match the data retrieved during EC beam pulses of up to 6 s. Lastly, the mirror was used in dedicated experiments to corroborate the polarization calculation, beam aiming, and EC current drive. This work

benefited from discussions, technical support, and collaboration with colleagues at DIII-D, including General Atomics staff and external collaborators. Their contributions, as well as the significant resources available to me as an employee of General Atomics, were integral to the success of this work. Where specific contributions by others were made, they are acknowledged throughout this thesis with appropriate citations.

The implementation of the top launch mirror system has demonstrated a prototype system capable of augmenting the launch trajectory of an existing outside launcher system to mimic the legacy top launch systems' capabilities prior to their eventual failure and removal. The new system has shown the top launch capability is achievable at DIII-D for approximately 10 % the cost of each of the previous installations and with the ability to be easily expanded upon. This expanded system would augment each of the eight transmission lines at DIII-D with a dedicated top launch mirror. The eight-mirror top launch system is planned for installation inside DIII-D near the end of this year (2026). At the time of writing, the design of the new system is complete and manufacturing of the graphite mirrors is beginning [63]. The upgraded installation will connect six 110 GHz and two 117.5 GHz gyrotrons to top launchers with optimized aiming for simultaneous operation in AT plasmas. The mirrors meant for 117.5 GHz EC beams have had their launch angles altered to allow for this. Each mirror has a smaller toroidal footprint than the prototype in order to avoid overlap of the left and right mirrors of each EC launcher and to conform to the new upper divertor upgrade that is scheduled for installation at the same time [64]. Thanks to the knowledge gained from the prototype mirror's installation, modifications have been made to the feet of the mirrors to avoid obstacles on the mounting plate and the nut bars have been changed to a more standard style using holes on the sides of the mirrors that will be plugged after installation. Additionally, since the hot test results matched closely to the theoretical beam propagation as well as the aiming angle determined by the laser alignment process, a hot test won't be needed for each mirror, streamlining the installation process. Each mirror will use two thermocouples, rather than four (since two will likely be sufficient to match the thermal

model of each tile to the readings), with a similar placement as the current top launch mirror. The thermal model will be updated and reapplied to predict the temperature response of each of the eight mirrors. A view of the current design for the eight-mirror top launch system is shown in Fig. 4.1. This system has the potential to more than double the off-axis current drive capabilities of the ECH system at the DIII-D tokamak.




Figure 4.1: Model view of the planned installation of eight total top launch mirrors. The surface of each mirror is highlighted in blue.

Appendix A

Table A.1: Measured insertion loss of copper and ATJ graphite miter bends in the E- and H-plane using a VNA. The average loss across both polarizations was 1.02% for copper and 4.78% for ATJ graphite.

Trial	Copper E		Copper H		ATJ E		ATJ H	
	dB	%	dB	%	dB	%	dB	%
1	-0.0422	0.97	-0.0488	1.12	-0.1539	3.48	-0.2633	5.88
2	-0.0479	1.10	-0.0487	1.12	-0.1518	3.43	-0.2833	6.32
3	-0.0319	0.73	-0.0468	1.07	-0.1587	3.59	-0.2842	6.33
4	-0.0363	0.83	-0.0382	0.88	-0.1493	3.38	-0.2727	6.09
5	-0.0470	1.08	-0.0488	1.12	-0.1479	3.35	-0.2613	5.84
6	-0.0364	0.83	-0.0522	1.19	-0.1534	3.47	-0.2654	5.93
7	-0.0460	1.05	-0.0470	1.08	-0.1571	3.55	-0.2656	5.93
8	-0.0580	1.33	-0.0472	1.08	-0.1549	3.50	-0.2698	6.02
9	-0.0359	0.82	-0.0469	1.07	-0.1610	3.64	-0.2697	6.02
10	-0.0400	0.92	-0.0409	0.94	-0.1584	3.58	-0.2849	6.35
Mean	-0.0422	0.97	-0.0466	1.07	-0.1546	3.50	-0.2720	6.07

References

- [1] W. J. Huang, M. Wang, F. G. Kondev, G. Audi, and S. Naimi. *AME2020 atomic mass evaluation data tables*. 2021. URL: <https://www.anl.gov/sites/www/files/2021-03/massround.mas20.txt>.
- [2] W. J. Huang, M. Wang, F. G. Kondev, G. Audi, and S. Naimi. “The AME 2020 atomic mass evaluation (I). Evaluation of input data, and adjustment procedures”. In: *Chinese Physics C* 45.3 (2021), p. 030002. DOI: 10.1088/1674-1137/abddb0.
- [3] M. Wang, W. J. Huang, F. G. Kondev, G. Audi, and S. Naimi. “The AME 2020 atomic mass evaluation (II). Tables, graphs and references”. In: *Chinese Physics C* 45.3 (2021), p. 030003. DOI: 10.1088/1674-1137/abddaf.
- [4] J. Wesson. *Tokamaks*. 4th ed. Vol. 149. International Series of Monographs on Physics. Oxford: Oxford University Press, 2011.
- [5] V. D. Shafranov. “Equilibrium of a toroidal plasma in a magnetic field”. In: *Journal of Nuclear Energy. Part C, Plasma Physics, Accelerators, Thermonuclear Research* 5.4 (1963), p. 251. DOI: 10.1088/0368-3281/5/4/307.
- [6] V. S. Mukhovatov and V. D. Shafranov. “Plasma Equilibrium in a Tokamak”. In: *Nuclear Fusion* 11.6 (1971), pp. 605–633. DOI: 10.1088/0029-5515/11/6/001.
- [7] J. P. Freidberg. *Ideal Magnetohydrodynamics*. New York: Plenum Press, 1987.
- [8] N. de Boucaud, T. Golfinopoulos, and A. Marinoni. “Demonstration and evaluation of negative triangularity equilibria in the ARC fusion pilot plant concept”. In: *Fusion Engineering and Design* 202 (2024), p. 114401. DOI: <https://doi.org/10.1016/j.fusengdes.2024.114401>.
- [9] B. Dudson. *FreeGS: Free Boundary Grad-Shafranov Solver*. Tech. rep. Lawrence Livermore National Laboratory, 2022. URL: <https://github.com/freegs-plasma/freegs>.

- [10] R. Prater. “Heating and Current Drive by Electron Cyclotron Waves”. In: *Physics of Plasmas* 11.5 (2004), pp. 2349–2376. DOI: 10.1063/1.1690762.
- [11] T. H. Stix. *Waves in Plasmas*. New York: American Institute of Physics, 1992.
- [12] N. J. Fisch and A. H. Boozer. “Creating an Asymmetric Plasma Resistivity with Waves”. In: *Physical Review Letters* 45.9 (1980), pp. 720–722. DOI: 10.1103/PhysRevLett.45.720.
- [13] T. Ohkawa. *Steady-State Operation of Tokamaks by r-f Heating*. Tech. rep. GA-A13847. NTIS Report PB2000-108008. San Diego, CA: General Atomics, 1976.
- [14] R. D. Stambaugh and the DIII-D Team. *DIII-D Program Overview*. Tech. rep. GA-A21840. San Diego, CA: General Atomics, 1994.
- [15] International Atomic Energy Agency. *ITER Technical Basis*. ITER EDA Documentation Series 24. Vienna: INTERNATIONAL ATOMIC ENERGY AGENCY, 2002. URL: <http://www.iaea.org/publications/6492/iter-technical-basis>.
- [16] A. J. Creely, M. J. Greenwald, S. B. Ballinger, D. Brunner, J. Canik, J. Doody, T. Fülöp, D. T. Garnier, R. Granetz, T. K. Gray, C. Holland, N. T. Howard, J. W. Hughes, J. H. Irby, V. A. Izzo, J. Kramer, A. Q. Kuang, B. LaBombard, Y. Lin, B. Lipschultz, N. C. Logan, J. D. Lore, E. S. Marmar, K. Montes, R. T. Mumgaard, C. Paz-Soldan, C. Rea, M. L. Reinke, P. Rodriguez-Fernandez, K. Särkimäki, F. Sciortino, S. D. Scott, A. Snicker, P. B. Snyder, B. N. Sorbom, R. Sweeney, R. A. Tinguely, E. A. Tolman, M. Umansky, O. Vallhagen, J. Varje, D. G. Whyte, J. C. Wright, S. J. Wukitch, J. Zhu, and the SPARC Team. “Overview of the SPARC Tokamak”. In: *Journal of Plasma Physics* 86.5 (2020), p. 865860502. DOI: 10.1017/S0022377820001257.
- [17] R. J. Buttery, B. Covele, J. Ferron, A. Garofalo, C. T. Holcomb, T. Leonard, J. M. Park, T. Petrie, C. Petty, G. Staebler, E. J. Strait, and M. Van Zeeland. “DIII-D Research to Prepare for Steady State Advanced Tokamak Power Plants”. In: *Journal of Fusion Energy* 38 (2018), pp. 72–111. DOI: 10.1007/s10894-018-0185-y.
- [18] R. J. Buttery, J. M. Park, J. T. McClenaghan, D. Weisberg, J. Canik, J. Ferron, A. Garofalo, C. T. Holcomb, J. Leuer, P. B. Snyder, and The Atom Project Team. “The Advanced Tokamak Path to a Compact Net Electric Fusion Pilot Plant”. In: *Nuclear Fusion* 61.4 (2021), p. 046028. DOI: 10.1088/1741-4326/abe4af.
- [19] R. C. O’Neill, A. S. Bozek, M. E. Friend, C. B. Baxi, E. E. Reis, M. A. Mahdavi, D. G. Nilson, S. L. Allen, and W. P. West. *Features and Initial Results of the DIII-D Advanced Tokamak Radiative Divertor*. Tech. rep. GA-A23284. San Diego, CA: General Atomics, 1999.

- [20] M. E. Friend, A. S. Bozek, C. B. Baxi, R. C. O'Neill, E. E. Reis, and M. A. Mahdavi. "Installation, Features, and Capabilities of the DIII-D Advanced Tokamak Radiative Divertors". In: *Fusion Engineering and Design* 56–57 (2001), pp. 433–437. DOI: 10.1016/S0920-3796(01)00397-9.
- [21] T. S. Taylor. "Physics of Advanced Tokamaks". In: *Plasma Physics and Controlled Fusion* 39.12B (1997), B47–B73. DOI: 10.1088/0741-3335/39/12B/005.
- [22] P. M. Anderson, C. B. Baxi, E. E. Reis, J. P. Smith, and P. D. Smith. "Development, Installation, and Initial Operation of DIII-D Graphite Armor Tiles". In: *Fusion Engineering and Design* 9 (1989), pp. 3–8. DOI: 10.1016/S0920-3796(89)80002-X.
- [23] K. L. Holtrop, M. Hansink, and A. G. Kellman. "Outgassing Tests on Materials Used in the DIII-D Magnetic Fusion Tokamak". In: *Journal of Vacuum Science & Technology A* 17.4 (1999), pp. 2064–2068. DOI: 10.1116/1.581727.
- [24] DIII-D National Fusion Facility. *Internal DIII-D Experimental and Engineering Data*. Internal data, General Atomics, San Diego, CA. 2026.
- [25] J. L. Luxon and the DIII-D Team. *An Overview of the DIII-D Program*. Tech. rep. GA-A22458. San Diego, CA: General Atomics, 1996.
- [26] F. Troyon, A. Roy, W. A. Cooper, F. Yasseen, and A. Turnbull. "Beta Limit in Tokamaks: Experimental and Computational Status". In: *Plasma Physics and Controlled Fusion* 30.11 (1988), pp. 1597–1609. DOI: 10.1088/0741-3335/30/11/019.
- [27] Y. Neyatani and the JT-60 Team. "Improved Confinement with Plasma Profile and Shape Controls in JT-60U". In: *Plasma Physics and Controlled Fusion* 38.12A (1996), A181–A191. DOI: 10.1088/0741-3335/38/12A/014.
- [28] A. Marinoni, O. Sauter, and S. Coda. "A Brief History of Negative Triangularity Tokamak Plasmas". In: *Reviews of Modern Plasma Physics* 5.6 (2021). DOI: 10.1007/s41614-021-00054-0.
- [29] M. P. Ross, K. A. Thackston, A. Dupuy, Y. Gorelov, N. de Boucaud, P. Nesbet, A. Torrezan, Z. Bayler, N. Watson, J. Anderson, and J. P. Squire. "Corroborating VNA and Thermal Measurements of Transmission Loss on the DIII-D ECH Waveguide System". In: *Fusion Engineering and Design* 225 (2026), p. 115626. DOI: 10.1016/j.fusengdes.2026.115626.
- [30] J. L. Doane. "Propagation and Mode Coupling in Corrugated and Smooth-Wall Circular Waveguides". In: *Infrared and Millimeter Waves*. Vol. 13. Academic Press, Inc., 1985, pp. 123–170.

- [31] J. L. Doane and C. P. Moeller. “HE₁₁ Mitre Bends and Gaps in a Circular Corrugated Waveguide”. In: *International Journal of Electronics* 77.4 (1994), pp. 489–509. DOI: 10.1080/00207219408926081.
- [32] R. Ellis, J. Hosea, J. Wilson, R. Prater, and R. Callis. “Design of a Dual High-Power, Long Pulse, Steerable ECH Launcher for DIII-D”. In: *AIP Conference Proceedings*. Vol. 595. American Institute of Physics, 2001, pp. 318–321. DOI: 10.1063/1.1424200.
- [33] J.-M. Jin. *Theory and Computation of Electromagnetic Fields*. Hoboken, NJ: John Wiley & Sons, 2010. DOI: 10.1002/9780470874257.
- [34] P. F. Goldsmith. *Quasioptical Systems: Gaussian Beam Quasioptical Propagation and Applications*. Piscataway, NJ: IEEE Press, 1998.
- [35] R. Ellis. *Internal Experimental and Engineering Data*. Internal data, Princeton Plasma Physics Laboratory, Princeton, NJ and General Atomics, San Diego, CA. 2026.
- [36] J. L. Doane. “Grating Polarizers in Waveguide Mitre Bends”. In: *International Journal of Infrared and Millimeter Waves* 13 (1992), pp. 1727–1743. DOI: 10.1007/BF01010741.
- [37] D. Wagner and F. Leuterer. “Broadband Polarizers for High-Power Multi-Frequency ECRH Systems”. In: *International Journal of Infrared and Millimeter Waves* 26.2 (2005), pp. 163–172. DOI: 10.1007/s10762-005-2997-x.
- [38] X. Chen, R. Prater, C. Petty, J. Lohr, D. Su, L. Lao, and V. Chan. “Top Launch for Higher Off-axis Electron Cyclotron Current Drive Efficiency”. In: *EPJ Web of Conferences (Online)* 203 (2019). DOI: 10.1051/epjconf/201920301004.
- [39] X. Chen, C. Petty, J. Lohr, D. Su, R. Prater, M. Cengher, M. Austin, C. Holcomb, L. Lao, R. Pinsker, B. Victor, and L. Zeng. “Doubling off-axis electron cyclotron current drive efficiency via velocity space engineering”. In: *Nuclear Fusion* 62.5 (2022), p. 054001. DOI: 10.1088/1741-4326/ac544a.
- [40] R. L. Abrams. “Coupling losses in hollow waveguide laser resonators”. In: *IEEE Journal of Quantum Electronics* 8 (1972), pp. 838–843. URL: <https://api.semanticscholar.org/CorpusID:121372506>.
- [41] J. Ren, T. Elsey, C. Murphy, B. Zhao, M. Shafer, R. Wilcox, S.-H. Hong, R. Hood, C. Tsui, and D. Donovan. “Design and Implementation of Surface Eroding Thermocouples for Stage-One Modular Divertor Integration at the DIII-D Tokamak”. In: *Fusion Science and Technology* (Oct. 2025). DOI: 10.1080/15361055.2025.2561354.
- [42] R. S. Wilcox, A. M. Holm, M. W. Shafer, T. M. Wilks, T. Elsey, R. T. Hood, C. Murphy, T. Osborne, and C. K. Tsui. “Design and Initial Measurements from the Shape and Volume Rise Divertor in DIII-D”. In: *Bulletin of the American Physical Society*. 66th

- Annual Meeting of the APS Division of Plasma Physics (Atlanta, Georgia, USA). Abstract PP12.00071. Oct. 9, 2024. URL: <https://meetings.aps.org/Meeting/DPP24/Session/PP12.71> (visited on 07/18/2026).
- [43] *ATJ Graphite Technical Data Sheet*. Technical data sheet. Amsted Graphite Materials. May 2020. URL: <https://www.amstedgraphite.com/wp-content/uploads/2020/05/ATJ.pdf> (visited on 07/18/2026).
 - [44] W. Kasparek, A. Fernandez, F. Hollmann, and R. Wacker. “Measurements of Ohmic Losses of Metallic Reflectors at 140 GHz Using a 3-Mirror Resonator Technique”. In: *International Journal of Infrared and Millimeter Waves* 22 (Nov. 2001), pp. 1695–1707. DOI: 10.1023/A:1015064616703.
 - [45] L. F. Mal’tseva and E. N. Marmer. “Determination of the Electrical Properties of Graphite at High Temperatures”. In: *Soviet Powder Metallurgy and Metal Ceramics* 1.7 (1962). Translated from *Poroshkovaya Metallurgiya*; original article submitted May 15, 1961, pp. 34–38. DOI: 10.1007/BF01135029.
 - [46] C. C. Petty, R. Prater, T. C. Luce, R. A. Ellis, R. W. Harvey, J. E. Kinsey, L. Lao, J. Lohr, M. A. Makowski, and K. L. Wong. *Physics of Electron Cyclotron Current Drive on DIII-D*. Tech. rep. GA-A24137. Conference paper. San Diego, CA: General Atomics, 2002.
 - [47] K. Matsuda. “Ray tracing study of the electron cyclotron current drive in DIII-D using 60 GHz”. In: *IEEE Transactions on Plasma Science* 17.1 (1989), pp. 6–11. DOI: 10.1109/27.21664.
 - [48] E. Poli, A. G. Peeters, and G. V. Pereverzev. “TORBEAM, a beam tracing code for electron-cyclotron waves in tokamak plasmas”. In: *Computer Physics Communications* 136.1 (2001), pp. 90–104. DOI: 10.1016/S0010-4655(01)00146-1.
 - [49] J. Hicok, T. Elsey, and K. Teixeira. *Personal communications regarding the manufacturing of the top launch mirror*. Personal communications. 2025.
 - [50] NeoGraf Solutions, LLC. *GRAFOIL® Flexible Graphite*. NeoGraf Solutions. 2024. URL: <https://www.neograf.com/products/gaskets-sealants/grafoil-flexible-graphite/>.
 - [51] P. Nesbet and A. Torrezan. *Personal communications regarding the integration of sensors with the top launch mirror*. Personal communications. 2025.
 - [52] X. Chen. *Personal communications regarding testing the top launch mirror in plasma experiments*. Personal communications. 2025.

- [53] T. Looby, M. Reinke, A. Wingen, J. Menard, S. Gerhardt, T. Gray, D. Donovan, E. Unterberg, J. Klabacha, and M. Messineo. “A Software Package for Plasma-Facing Component Analysis and Design: The Heat Flux Engineering Analysis Toolkit (HEAT)”. In: *Fusion Science and Technology* 78.1 (2022), pp. 10–27. DOI: 10.1080/15361055.2021.1951532.
- [54] A. Wingen, M. S. d’Abusco, M. Churchill, D. Corona, N. M. Ferraro, A. Kleiner, T. Looby, and S. Munaretto. “Development and validation of non-axisymmetric heat flux simulations with 3D fields using the HEAT code”. In: *Nuclear Fusion* 65.8 (July 2025). DOI: 10.1088/1741-4326/adeff1.
- [55] I. Favela and N. Eidietis. *Personal communications regarding the disruption forces on the top launch mirror and surrounding tiles*. Personal communications. 2025.
- [56] J. P. Anderson, M. A. Shapiro, R. J. Temkin, and D. R. Denison. “Phase retrieval of gyrotron beams based on irradiance moments”. In: *IEEE Transactions on Microwave Theory and Techniques* 50.6 (2002), pp. 1526–1535. DOI: 10.1109/TMTT.2002.1006414.
- [57] A. Laut and J. Anderson. *Internal General Atomics software used for phase retrieval and millimeter-wave beam reconstruction*. 2023.
- [58] J. Barr, W. Choi, A. Hyatt, and G. Sips. *Personal communications regarding the proximity limit of the plasma and the top launch mirror*. Personal communications. 2025.
- [59] C. C. Petty, T. C. Luce, M. E. Austin, H. Ikezi, J. Lohr, and R. Prater. *Polarization Measurements During Electron Cyclotron Heating Experiments in the DIII-D Tokamak*. General Atomics Report GA-A23208. San Diego, CA: General Atomics, July 1999.
- [60] X. Chen, C. C. Petty, N. de Boucaud, M. Ross, P. Nesbet, B. Victor, J. M. Park, M. Austin, Y. Gorelov, J. Squire, and C. Holcomb. “‘Top Launch’ Electron Cyclotron Current Drive” (Barcelona, Spain). Presented at the 23rd Joint Workshop on Electron Cyclotron Emission and Electron Cyclotron Resonance Heating (EC-23). May 2026.
- [61] M. Ross. *Personal communications regarding polarization transformation due to the top launch mirror*. Personal communications. 2025.
- [62] D. Wróblewski and L. L. Lao. “Polarimetry of Motional Stark Effect and Determination of Current Profiles in DIII-D”. In: *Review of Scientific Instruments* 63.10 (Oct. 1992), pp. 5140–5147. DOI: 10.1063/1.1143463.
- [63] J. Hicok. *Personal communications regarding the mechanical design of the upcoming eight-mirror expansion of the top launch system*. Personal communications. 2026.

- [64] C. Petty and D.-D. Team. “Upgrading DIII-D to Close Physics Gaps to Future Fusion Reactors”. In: *Bulletin of the American Physical Society*. 63rd Annual Meeting of the APS Division of Plasma Physics. Abstract CP11.053. Oct. 2021.