The Water Electrolysis Hall Effect Thruster

Alexander Schwertheim

Imperial College London

November 2022

Submitted in part fulfilment of the requirements for the degree of

Doctor of Aeronautics of Imperial College London

and the Diploma of Imperial College London

Declaration of Originality

This is to certify that the copy of my thesis, which I, Alexander Schwertheim, have presented for consideration for my PhD degree, embodies the results of my own course of study and research, has been composed by myself and has been seen by my supervisor before presentation.

Copyright Statement

Abstract

This study proposes operating a Hall effect Thruster (HET) on the products of in-situ water electrolysis. Oxygen is supplied to the anode and hydrogen to the cathode to mitigating poisoning. In comparison to other alternative HET propellants water is non-toxic, low cost, highly storable, and allows synergies with chemical propulsion systems. The competitiveness of this concept is experimentally tested by simulating, designing, constructing, and operating a thruster optimised for oxygen: the Water ElecTrolysis Hall Effect Thruster (WET-HET). A zero-dimensional particle-in-cell code is used to optimise the device, resulting in a channel width of 5 mm, outer circumference of 25 mm and depth ranging from 35 mm to 60 mm. Iterative magnetic, thermal, and mechanical modelling finalises the design.

A hanging pendulum style thrust balance is used to directly measure WET-HET performance. Safety restrictions require krypton rather than hydrogen be used for the cathode. The WET-HET is tested on discharge powers up to 3200 W and oxygen mass flow rates ranging from 0.99 mgs^{-1} to 1.85 mgs^{-1}. Three different channel depths are tested: 35 mm, 45 mm and 60 mm. The WET-HET is characterised for a range of different magnetic field strengths and for three distinct magnetic topologies. Two channel wall materials were investigated: alumina and boron nitride.

The highest thrust measured was 38.63±0.25 mN, with a maximum specific impulse of 4112±36 s and a maximum anode thrust efficiency of 24.4±5.9%. This optimum was found when the WET-HET was operated at 3200 W discharge power, 0.99 mgs^{-1} mass flow rate, 35 mm channel depth

and a magnetic field that peaked near 600 Gauss. Boron nitride outperformed alumina by approximately 40%. Contrary to expectation, increasing the axial thickness of the magnetic region had little impact on discharge voltage, but led to a reduction in thrust, specific impulse, and thrust efficiency of the device.

Non-technical Abstract

An alternative project description is provided for the layperson.

This PhD thesis concerns a new fuel for electrical rockets used by satellites to manoeuvre around in space. Unlike the more familiar chemical rockets, electrical rockets consume fuels which are nonexplosive and use the electricity produced by the satellite's solar panels. An alternative name for this type of electrical rocket is a plasma thruster. Carefully positioned electric and magnetic fields are used to bring these special fuels to the fourth state of matter called plasma. A plasma can be thought of as a hot "charged-particle soup", which is important as individual charged particles can be easily accelerated out of the plasma to very high speeds using electric fields. Electrical rockets consume rocket fuel at a very slow rate, but harness electric fields to accelerate the plasma to a much higher speed than in a traditional chemical rocket. This ultimately means that engineers must choose between their satellite having a chemical rocket which is powerful but consumes a lot of fuel, or an electrical rocket which is ten thousand times lower in thrust, but ten times more fuel efficient.

Electric rockets are becoming more and more popular, but the fuel that they consume is typically a gas called xenon which is extremely rare and expensive, causing many to look for alternative fuels. In this thesis, it is suggested that water can be used as a fuel to replace xenon in a type of electric rocket called a Hall effect thruster (HET). While water chemically reacts to destroy some of the components within the HET, our novel approach is to use a device called an electrolyser. An electrolyser uses electricity to convert water into hydrogen and oxygen gas, neither of which react with the HET components. This method allows us to fuel a satellite with water, and only generate hydrogen and oxygen in space when they are needed. Liquid water is cheaper and easier to store than xenon gas, and may one day be harvested from the moon, other planets, or asteroids, allowing the satellite to refuel in space. Other researchers are using the hydrogen and oxygen produced by water electrolysis for chemical rockets, meaning a single satellite could have the best of both worlds with a chemical and an electrical rocket on board, both fuelled by

water.

In order to test the viability of this concept, it must be tested experimentally. Our proof of concept device is called the Water ElecTrolysis Hall Effect Thruster or WET-HET for short. First, we develop a software code which lets us simulate how different thruster shapes and configurations affect performance. The thruster which worked best in the simulation is then put forward, and standard engineering analysis software tools are used to design how the magnets will work, how the heat will be dissipated, which materials to use, and how everything will fit together. The WET-HET is then built and tested in the vacuum chamber in our laboratory with different power levels, fuel rates and in different configurations. We found that the WET-HET is twice as fuel efficient, but half as power efficient as a comparable xenon thruster. We suggest that this is a promising start, and that we do not see any obvious reason to not continue research into water electrolysis electric rockets.

Acknowledgements

I would like to start by thanking Dr. Aaron Knoll for being such an inspirational and dedicated supervisor during all phases of my project. You have taught me so much, from theoretical skills like how to numerically simulate a discharge, to the very different practical task of wrangle a plasma into existence. I thank you for committing so much of your valuable time to me and for teaching me how to be an active member of the scientific community. You have always provided me significantly more support and patience than any PhD students could ever dream of.

I cannot thank the rest of the space propulsion group enough. To Charlie, for all the great times we spent both in and out of the lab. To Jesus for your passion and quick understanding: I am delighted that someone as dedicated as you in taking over the WET-HET. To Junjie, for spontaneous trips to Stiens, and to all the other group members who helped me along the way. Thank you to the technicians, specifically Franco and Roland. I would like to thank my family: Heiner, Ulli and Helena, who have been so supportive over these chaotic years. Without you three none of this would have been possible. I would also like to thank some others in no particular order for helping me to stay sane: Matt, Chris, Henrike, Silvia, Marius, Hannes, Joe and Emmet.

Finally, I want to thank you Ella, for supporting me through thick and thin, through both of our PhDs and a pandemic. I am so grateful to have you in my life, and I cannot wait to tackle the next adventure with you!

Contents

List of Acronyms

- **COTS**: Commercial Off the Shelf
- **CP**: Chemical Propulsion
- **DAQ**: Data Acquisition System
- **ECHT**: Extended Channel Hall Thruster
- **EMI**: Electro-Magnetic Interference
- **EP**: Electric Propulsion
- **EPL**: ESA Propulsion Laboratory
- **ESA**: European Space Agency
- **ESB**: Electronic Support Box
- **FEMM**: Finite Element Method Magnetics
- **GIE**: Gridded Ion Engine
- **GOCE**: Gravity Field and Steady-State Ocean Circulation Explorer
- **HET**: Hall Effect Thruster
- **ICL**: Imperial College London
- **IPPL**: Imperial Plasma Propulsion Laboratory
- **ISRU**: In Situ Resource Utilisation
- **LVDT**: Linear variable differential transformers
- **MAI**: Moscow Aviation Institute
- **MFC**: Mass Flow Controller
- **MOV**: Metal Oxide Varistor
- **NMI**: National Measurement Institute
- **PIC**: Particle In Cell
- **QCT**: Quad Confinement Thruster
- **QCTP**: Quad Confinement Thruster Phoenix
- **RTD**: Resistance Temperature Detectors
- **sccm**: Standard Cubic Centimetres per minute
- **SEE**: Secondary Electron Emission
- **SPT**: Stationary Plasma Thruster
- **TTPR**: Thrust To Power Ratio
- **VCA**: Voice Coil Actuator
- **VLEO**: Very Low Earth Orbit
- **WET-HET**: Water ElecTrolysis Hall Effect Thruster

Nomenclature

Error Budget Calculation

$\bar{F}$	Average calibration force, (N)
$\hat{x_i}$	Measured displacement due to calibration force i, (m)
σ_x^2	Variance of the random errors to a known input, (m^2)
σ_{F_t}	Standard deviation of thrust measurement, (N)
σ_{scal}	Standard deviation of the thrust balance sensitivity, (m/N)
σ_x	Standard deviation of the random error to a known input, (m)
b_{cal}	Intercept of the sensitivity and the y-axis, (m)
F_i	Known calibration force i, (N)
F_t	Thrust measurement, (N)
n	Number of calibration forces used
R_s	Standardised residuals
s_x^2	Estimate of the variance of the random errors to a known input σ_x^2, (m^2)
s_F	Estimate of the standard deviation of the thrust measurement σ_{F_t}, (N)
S_{cal}	Estimate of thrust balance sensitivity, (m/N)

s_{Scal} Estimate of the standard deviation of the sensitivity σ_{scal}, (m/N)

s_x Estimate of the standard deviation of the random error to a known input σx, (m)

x Measured displacement, (m)

In-situ Calibration System Measurements

α Angle of servo motor displacement, (rad)

a Original string length between guide hole and servo arm, (m)

a' Displaced string length between guide hole and servo arm, (m)

b Original string length between guide hole and mass, (m)

b' Displaced string length between guide hole and mass, (m)

H Vertical distance from test mass and central platform, (m)

h String length between test mass and central platform, (m)

m Mass of thrust balance test mass, (kg)

R Length of servo arm, (m)

Thruster Characteristics

α_p Propellant utilisation efficiency, (%)

δ Discharge voltage loss coefficient

$\dot{m}$ Anode mass flow rate, (kgs^{-1})

η_T Anode thrust efficiency, (%)

λ_i Mean ionisation length, (m)

A_c Channel cross sectional area, (m^{-2})

g_0 Acceleration due to gravity at the Earth' surface, (ms^{-1})

I_d Discharge Current, (A)

I_i Ion current, (A)

I_m Magnetic circuit current (or magnet current), (A)

$I_{\dot{m}}$ Anode mass flow equivalent current, (A)

I_{sp} Specific Impulse, (s)

k_i Ionisation rate, ($\mathrm{m}^3\mathrm{s}^{-1}$)

L Channel depth (or channel length), (m)

m_n Mass of neutral particle, (kg)

n_e Electron density, (m^{-3})

P_d Discharge Power, (W)

q Charge of electron, (C)

r_i Channel internal diameter, (m)

r_o Channel outer diameter, (m)

T Thrust, (N)

V_d Discharge Voltage, (V)

v_n Velocity of neutral particle, (ms^{-1})

W Channel width, (m)

v_{ion} Ion velocity, (ms^{-1})

Chapter 1

Background

This study proposes a novel propellant to be consumed by a electric thruster for spacecraft applications. This chapter provides some context to the reader by introducing the basics of in-space propulsion and describing the trade-off between chemical propulsion and electrical propulsion. The physics behind different electric propulsion technologies such as Hall effect thrusters, and hollow cathodes is described. The reader is presented with issues that are becoming apparent over the electric propulsion community's over-reliance on xenon as the dominant Hall effect thruster propellant. A brief review of research into alternative propellants highlights the benefits and disadvantages of each contender. The chapter concludes by defining the six core requirements that a novel propellant must meet to be considered a viable alternative for use in a Hall effect thruster.

1.1 In-space Propulsion

The Falcon 9, Electron or Soyuz are examples of launch vehicles: large rockets that deliver satellites to space. Once the launch vehicle reaches a sufficient altitude and velocity it releases the individual satellites into a temporary "parking" orbit. The satellite must now fend for itself by employing small on-board rockets to propel itself through space. The term propulsion system refers to these small rockets along with the plumbing and electronics required to operate them.

This text will focus on in-space propulsion, which distinguishes the smaller on-board subsystems of a satellite from the much larger propulsion used by launch vehicles. Some miniaturised cube-satellites cut costs by forgoing a propulsion system. However, the majority of satellites require control over their orbit and attitude throughout their lifetime, meaning the propulsion system is mission critical. Propulsive manoeuvres or "burns" are executed by spacecraft operators for a multitude of reasons. Increasing or decreasing the altitude of a satellite directly affects the orbital period, and is referred to as orbit raising or lowering. Changing the inclination of an orbit alters the latitude over which the vehicle will pass. Many missions require the spacecraft to maintain a very stable and precise trajectory which demands the propulsion system perform small adjustments to perform what is called station keeping. Often the propulsion system must operate reliably for decades, routinely performing burns to ensure the spacecraft is precisely where the mission requires it to be.

The propulsion system has the simple goal of changing the momentum of the spacecraft. It achieves this by accelerating and expelling a consumable resource to serve as the reaction mass. Colloquially called the rocket fuel, this resource is more accurately referred to as the propellant. Newton's third law dictates that the force acting on the expelled propellant by the spacecraft is of equal magnitude and opposite direction to the force on the spacecraft. The force generated on the spacecraft by the propulsion system is called the thrust, denoted T, and changes the momentum of the spacecraft. For in-space propulsion thrust is simply the momentum imparted on the expelled propellant, given by:

$$T = \dot{m} v_{ex}, \tag{1.1}$$

where the $\dot{m}$ is the mass flow rate of the propellant, and v_{ex} is the average exhaust velocity of the propellant as it is ejected. It is important to note that equation 1.1 neglects the pressure term typically seen acting on the exhaust plane of chemical thrusters, due to the very low pressures

found within electric thrusters. Along with thrust, one of the most critical performance metrics of any propulsion system is the specific impulse, denoted as I_{sp}, as this describes how much thrust can be generated per unit weight (on Earth) of propellant:

$$I_{sp} = \frac{v_{ex}}{g_0} = \frac{T}{\dot{m} g_0} \tag{1.2}$$

where g_0 is the acceleration due to gravity at the surface of the Earth. Equation 1.2 shows that the I_{sp} of a thruster has units of seconds, and is effectively just a measure of the exhaust velocity of the propellant for electric thrusters. Intuitively this makes sense: for a given propellant mass the effectiveness of the thruster to change the momentum of the spacecraft is simply dependent on the velocity of the expelled mass. The I_{sp} of a propulsion system is used to measure how efficient a propulsion system is in terms of propellant utilisation.

To summarise, the larger the thrust, the faster the spacecraft can change its momentum. The higher the I_{sp}, the less propellant is needed for the same momentum change.

1.2 An Important Distinction: Chemical Propulsion vs Electric Propulsion

In-space propulsion systems are split into a rigid dichotomy of chemical propulsion (CP) and electric propulsion (EP). The distinction between the two systems stems from how the energy is sourced and used to accelerate the propellant. A CP system relies on familiar rocket principles: chemical energy stored in the propellant is violently released within a combustion chamber. This release of energy causes the propellant to rapidly increase in temperature and pressure, with the nozzle acting to direct the expanding gases away from the satellite. This is the same principle

as in launch vehicles, yet on a much smaller scale.

In EP systems electrical energy is collected by solar panels and converted into kinetic energy of the propellant. We can loosely categorise EP into three broad definitions: electro-thermal systems, electrostatic thrusters and plasma propulsion. In an electro-thermal system, energy is transferred to the propellant by heating the gas, and directing it with use of a nozzle. This includes resistojets and arcjets. Electrostatic thrusters included Gridded Ion Engines (GIEs), field emission electric propulsion and colloid thrusters. These systems use strong electric fields to extract and accelerate charged particles from the propellant. Plasma propulsion system such as Hall Effect Thrusters (HETs), helicon thrusters, pulsed plasma thrusters and magnetoplasmadynamic thrusters generate a plasma which is then accelerating by a combination electrica and magnetic fields. For the remainder of this thesis, we will focus on GIEs and HETs as these systems have the most flight heritage (electro-thermal systems also have great flight heritage, but have a considerably lower I_{sp}). The terms EP and plasma propulsion will be used interchangeably to cover GIEs and HETs. In these systems, radio waves, microwaves or impact from electrons is used to focus electrical energy into the propellant to reach ionisation. This means the electrons are no longer bounded to individual atoms and the propellant reaches the fourth state of matter: plasma. Unlike solids, liquids or gases, plasmas can be directly manipulated using magnetic and electric fields. Once an electric propulsion system has generated a plasma, it employs electric and magnetic fields to accelerate the plasma out of the thruster at extremely high velocities.

Due to the fundamental difference in how thrust is generated, the propellants used in CP and EP systems are selected on very different merits. In a CP system the propellant must have as much chemical energy stored in relation to its molecular weight as possible when entering the combustion chamber. However, the propellant must also be stable enough to be stored at

reasonable conditions in a storage tank for the duration of the mission. This often requires the fuel and oxidiser to be stored separately. For in-space propulsion, the majority of chemical fuels are hydrazine and its derivatives; typical oxidisers include dinitrogen tetroxide and nitric acid. Thrust levels for CP systems vary greatly depending on purpose, from less than a newton for fine spacecraft control, to several kilo-newtons for rapid orbit raising. Given all of the energy available to the thruster must be stored in the chemical bonds of the propellants, a maximum exhaust velocity, and thus I_{sp} can be calculated for any CP system. Values approaching these theoretical limits can be achieved, with most systems having anywhere from 180 s to 400 s of I_{sp}[1].

When selecting propellants for an EP system, the chemical potential of the propellant is of less concern, given electrical energy is supplied directly from the spacecraft bus. Chemically inert propellants are usually chosen for easier storage. Of more importance for a plasma propulsion systems is that the propellant can be easily ionised. This depends on the method of ionisation, but generally requires that the propellant has a combination of low ionisation energy and large electron impact cross-section. The greater the atomic mass of the propellants, the greater the thrust, but the lower the specific impulse for a given power. Krypton, iodine, and mercury have shown promise in the laboratory, however the high atomic mass, low ionisation energy, low chemical reactivity, and high storage density of xenon when under pressure make it the *de facto* standard propellant for electric propulsion[2].

Generally speaking, EP systems cannot compete with CP systems on thrust. The electrical energy required to ionise and subsequently accelerate xenon means that even the most powerful EP thrusters ever flown struggle to generate more than a few hundred milli-newtons of thrust. The benefit of utilising an EP system is that the theoretical I_{sp} of the thruster is almost limitless. Directly accelerating ions using an electric field allows them to reach very high exhaust velocities as they are ejected. Typical EP systems operate at an I_{sp} above 1500 s, with extreme

applications demonstrating over 10,000 s.

The drastic differences in thrust and I_{sp} between CP and EP systems determines that the two occupy polar extremes in parameter space: the thrust achievable by small CP systems dwarf that of even the most powerful EP systems. This is made up for by the I_{sp} provided by EP systems, which can reduce the propellant mass of a spacecraft by hundreds of kilograms. The decision ultimately depends on the mission at hand, and sometimes the choice is obvious: rapid orbit raising, emergency collision avoidance and controlled re-entry burns are impossible without high thrust CP burns, but decade long station keeping, or interplanetary cruises require an unfeasible propellant mass if EP is not utilised.

The propulsion system of a spacecraft must perform a variety of tasks, each with different requirements. In the past orbit raising from a typical parking orbit near 400 km to geostationary orbit at 35,700 km was performed over the course of a few days by a CP system, by a manoeuvre called a Hohmann transfer. A Hohmann transfer is the most efficient method to move from one orbit to another in the same plane, and requires two high thrust burns. Such high thrust burns are not possible with EP systems. A less efficient method to raise an orbit is to continuously thrust over several months, using low but continuous thrust to gradually spiral outwards to the desired altitude. Such a system requires several months of travel, but allows an EP system to dramatically reduce the propellant required to reach the final orbit. For telecommunication satellite providers, a trade-off must be made to determine if the lower launch mass achievable by electric orbit raising is worth the several lost months of the transfer. Oh et al.[3] used a geostationary orbit raising test case to calculate the mass saving produced when a slower electric orbit raising manoeuvre is used over a faster chemical one. They concluded that approximately 6 kg of launch mass is saved for every day that the orbit raising in extended.

Once a satellite reaches its intended orbit, orbital maintenance such as station keeping must be performed. These short burns are required periodically throughout a spacecraft's lifetime to account for slight perturbation and unwanted orbital precession. For these frequent and low thrust corrections, EP systems are far better suited providing substantial propellant mass savings. In the past the constant but low thrust which can be provided by EP have been used in Very Low Earth Orbits (below 500 km) to compensate for the drag force created by the residual atmosphere. The most famous examples of this are ESA's Gravity Field and Steady-State Ocean Circulation Explorer (GOCE) and Jaxa's Super Low Altitude Test Satellite (SLATS), which each flew to a record low orbit of 229 km and 167.4 km respectively using drag compensating EP systems[4;5].

It would clearly be desirable to fly with two propulsion systems: an electrical and a chemical system, such that both high thrust and high I_{sp} burns can be performed when needed. However flying with two separate propellant storage systems, two propellant management systems, two power processing units and two thrusters imposes a significant increase in mass and complexity.

1.3 Electric Propulsion Technologies

The most commonly flown plasma thrusters are Gridded Ion Engines (GIEs) and Hall Effect Thrusters (HETs)[6]. Both of these technologies require the use of a cathode to provide electrons to ignite and neutralise the outgoing ion beam. The construction and support electronics of GIEs are more complex and thus expensive than those of HETs. While GIEs tend to operate at a greater I_{sp}, HETs provide a greater Thrust To Power Ratio (TTPR)[7]. This text provides a brief explanation of the operating principles of a GIE, but the main focus is the HET and the hollow cathode.

1.3.1 Gridded Ion Engines

A picture of a GIE firing is shown in figure 1.1, and schematically in figure 1.2. The electrically neutral propellant, usually xenon, is fed into the discharge chamber, the external cathode, and for some systems, the internal cathode. Different GIEs employ different mechanisms to ionise the gas, including electron bombardment, radio frequency coupling, and electron cyclotron resonance[7]. The configuration shown in 1.2 is that of a Kaufman style thruster, were electrons bombard neutrals as they move from the internal cathode to the anode, both within the chamber[8]. The magnetic field is carefully tailored to extend the path length that energetic electrons take to the anode, such that the electrons become trapped, which improves ionisation efficiency[7]. This generates a plasma contained within the thruster consisting of negative electrons and positive ions. The same acceleration mechanism is used by all GIEs: several layered, biased grids cover the opening of the thruster. The first grid is the screen grid, which shields the plasma bulk from the downstream electric field. As the ions pass the screen grid, they find themselves in a strong electric field generated by the strong potential difference between the screen grid and the acceleration grid, which propels them to high velocities as they leave the thruster[7]. In the T6 thruster shown in figures 1.1 and 1.2 the screen grid typically has a voltage of 1850 V, and the acceleration grid a voltage of -265 V[9]. The geometry of these grids is carefully configured to precisely accelerate the ions and prevent collision and erosion. The external hollow cathode supplies electrons to the exiting ions, to ensure the spacecraft does not build up charge. In some thrusters a third deceleration grid is used to stop electrons being pulled into the chamber, and to protect the acceleration grid from charge exchange ions. Note that the electrons dispensed by the internal cathode do not exit the chamber, and the electrons dispensed by the external cathode do not enter the chamber, meaning only ions pass through the grids. A GIE like the T6 can produce over 4000 s of I_{sp}, and a thrust of 142.7 mN which means that 64% of the 4.5 kW

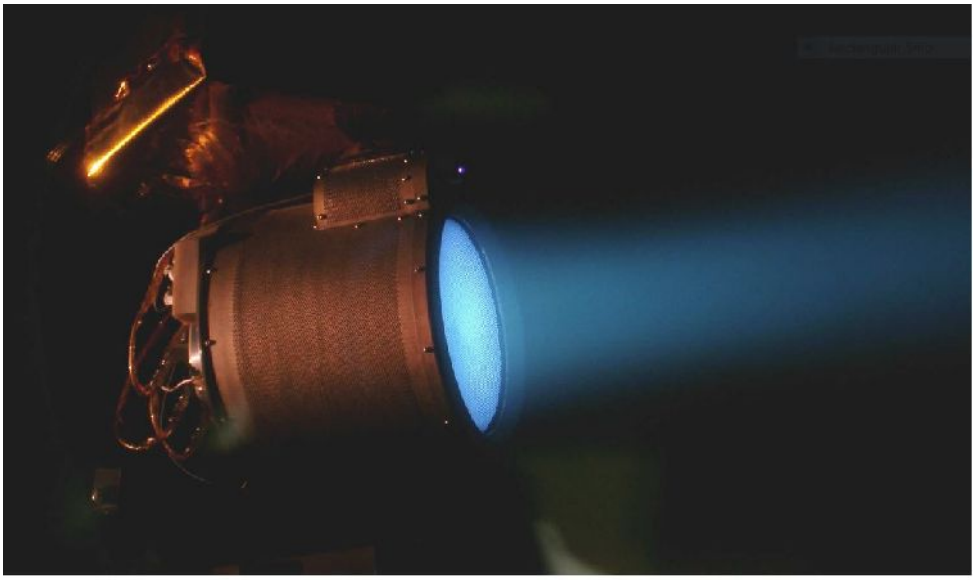

Figure 1.1: The T6 ion engine operating on xenon. Note the external cathode visible at the top of the thruster exit. Reproduced from Blott et al.[10] with permission of the authors.

input power is converted into kinetic energy of the ions[9].

1.3.2 Hall Effect Thrusters

The main body of a Hall Effect Thruster is a cylindrically symmetric device with a shallow annular channel. Photographs of a Hall effect thruster are shown in figure 1.3, with a diagram showing major components in figure 1.4. The discharge channel is made of a dielectric ceramic such as alumina (Al_2O_3) or boron nitride (BN). To ignite the thruster, approximately 10% of the propellant is supplied to the heated hollow cathode, which produces a cloud of electrons near the exit of the channel. How the hollow cathode produces electrons is described in the following subsection. The anode of the thruster is nested at the base of the discharge channel, serving also as the inlet to the remaining 90% of the propellant, which diffuses down the length of the channel. Next a potential difference is established between the positive anode and the negative external cathode. The electrons gain kinetic energy as they move axially down the channel

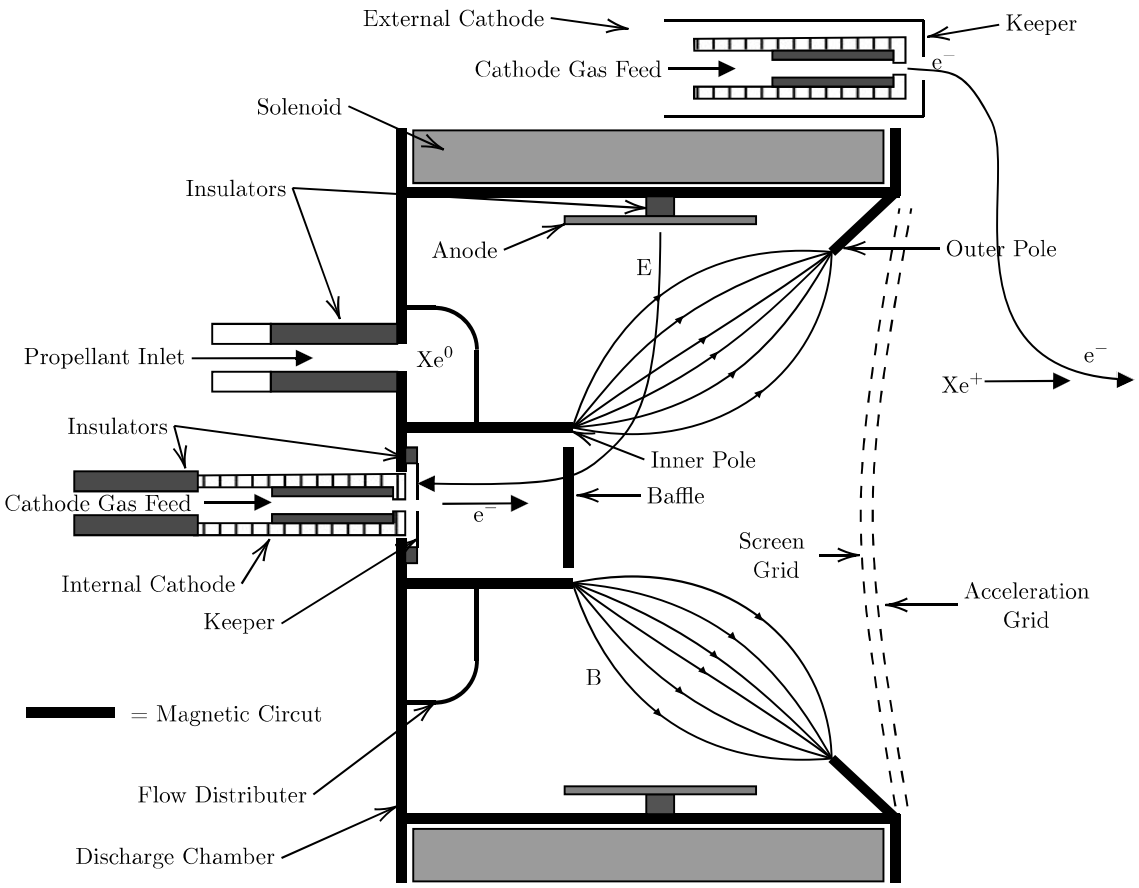

Figure 1.2: Diagram of a T6 Gridded Ion engine. Adapted from Hutchins et al.[8].

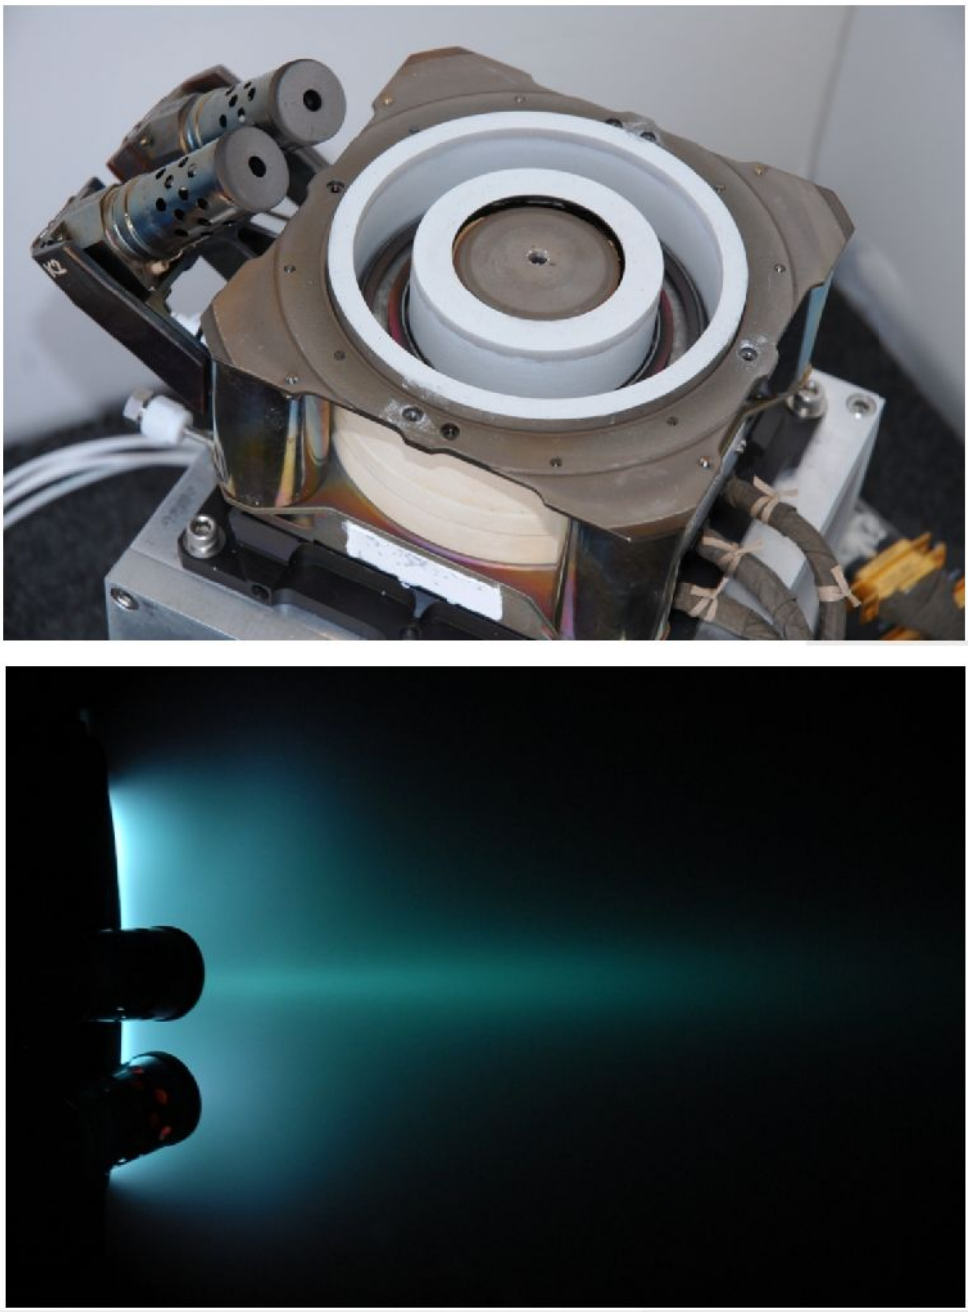

Figure 1.3: An SPT-100 model Hall effect thruster pointing upwards, with two external cathodes (top) and side on view when operating with xenon (bottom). Both images are reproduced from Nakles et al.[11] with permission of the authors.

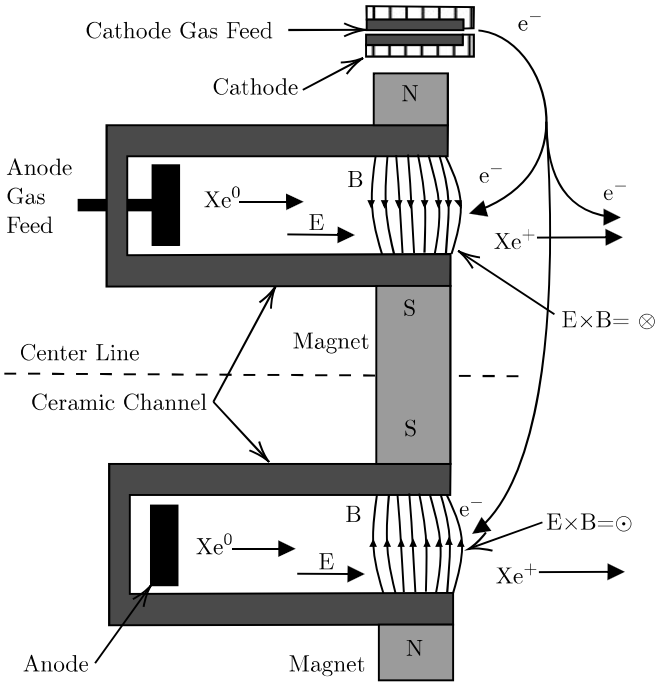

Figure 1.4: Diagram of Hall Effect Thruster with plume going to the right. Adapted from Goebel and Katz[7].

toward the anode (to the left in the figure 1.4). Within the channel, the electrons collide with the electrically neutral xenon particles which are drifting out of the channel at thermal velocities. If an electron and a neutral collide with sufficient energy, a further electron is knocked off the neutral particle creating a positive ion. This is shown here for a xenon atom.

$$Xe + e^- \longrightarrow Xe^+ + 2e^- \tag{1.3}$$

The energy threshold required for this collision is called the ionisation energy and depends on the gas species (12.1 eV for xenon). Note that the additional electron generated in equation 1.3 helps to increase the local electron density and produces a runaway effect that establishes a sustained stationary plasma within the thruster channel. At this point the thruster has been ignited: a stationary DC plasma is produced, an electron current passes from the cathode to the anode. This operational mode is referred to as glow discharge mode, which means that an external magnetic field has not yet been applied, and the same potential is found throughout the plasma. In glow discharge mode thrust and I_{sp} are low and the Hall current has not been

established.

To switch from glow discharge mode to nominal operation, power must be supplied to the electromagnets. The magnetic circuit of a HET serves to increase electron density by confining the electrons in a cross-field circumferential drift, which in turn raises the likelihood of ionising electron-neutral collisions. The applied magnetic field achieves this by diverting and extending the electron's path to the anode. The magnetic configuration comprises a collar of electromagnets secured around the outside of the discharge channel very near to the exit plane, with the opposite magnetic pole in the centre of the discharge channel. The produced magnetic field points radially across the channel (vertically in figure 1.4) directly before the channel exit. We refer to the axial thickness of the magnetic field as the magnetic thickness. The electrons are impeded in their axial progression by the magnetic field, resulting in an increase in electron density within the magnetic region. With the rise in electron density, the equipotentiality of the plasma is lost, and an axial electric field is formed within the thruster channel. The positive xenon ions within the channel are rapidly accelerated by this axial electric gradient, exiting the channel at speeds on the order of 20 km/s. Upon exiting the device, the speeding positive ion attracts a single electron as it passes through the external electron cloud. This neutralises the ejected particles, and ensures no net charge leaves the spacecraft. The momentum of these expelled ions is ultimately how a Hall effect thruster generates thrust.

As the electrons try to travel to the anode through the magnetic region, they are deflected, trapped, and forced to orbit the magnetic field lines via the Lorentz force. The orbital radius of a charged particle about a field line is called the gyroradius (r_g) which can be calculated from:

$$r_g = \frac{m v_\perp}{|q| B}, \tag{1.4}$$

where m and q are the mass and charge of the particle, and B is the magnetic field strength. The component of the velocity which is orthogonal to the magnetic field is $v_\perp$. Note that increasing the mass of the particle increases the gyroradius, which means the particle will either orbit at a large radius, or simply be deflected off course as it passes through the region. The magnitude of the magnetic field must be carefully selected such that the electrons are strongly deflected or even trapped, but the more massive ions are free to pass through the region mostly unperturbed. Particle gyration in the presence of an electric field ultimately leads to an electron drift which is orthogonal to both the magnetic and the electric field, and is thus in the azimuthal direction about the central axis. This is into the page above the centre line, and out of the page below the centre line in figure 1.4. This is an example of an E cross B (cross here describing the cross product) Hall effect current. This closed azimuthal current passing through the stationary plasma gives the device its many names: Hall Effect Thruster, closed drift thrusters, or stationary plasma thruster (SPT). The electrons are forced to circulate around the channel as they are impeded on their journey to the anode, increasing the electron density dramatically in this magnetised region. An increase of the electron density creates the axial electric field within the plasma, and increases the chance of electron-neutral collisions which leads to greater rates of ionisation.

The applied magnetic field takes the thruster from glow discharge mode to nominal operation, greatly increasing thrust and specific impulse. The impedance of the magnetic field necessitates a higher voltage difference between the anode and cathode, which increases the acceleration potential (and thus exit velocity) of the ions, and the power consumption of the thruster.

The SPT-100 Hall Effect Thruster

The vast majority of early HET research was performed by the USSR, with the first HET flight being a model SPT-60 on board the Meteor #18 in 1971[6]. The first commercial American HET

mission was not until 2004, and flew the Russian model SPT-100 HET produced by Fakel[7]. This SPT-100 is widely considered to be the gold standard in xenon Hall effect thrusters. Originally flown in 1994, the SPT-100 dominates the HET market: 84 of the 130 GEO spacecraft launched with HETs on board between 1981 and 2018 flew the SPT-100[6]. The nominal operating point for this thruster is 5 mgs^{-1} xenon to the anode, with 0.5 mgs^{-1} going to the cathode[12]. The potential difference between the anode and cathode is referred to as the discharge voltage. The SPT-100 operates with a discharge voltage near 300 V, and 4.5 A of discharge current flowing through the plasma, meaning it operates at a discharge power of 1350 W. At such an operating point the thruster produces approximately 85 mN of thrust, with an anode I_{sp} of 1706 s when operating on xenon. The term anode I_{sp} as opposed to system I_{sp} specifies that the propellant consumption of the cathode is disregarded when calculating I_{sp}. Similarly the anode thrust efficiency includes only the discharge power, whereas system efficiency includes all power expended by the system, which includes the power consumption of the magnetic coil and cathode heater. Given the extensive flight and research heritage of the SPT-100[12;13;14], this text will often use this particular model of HET as an example of the state of the art xenon HET.

1.3.3 Hollow Cathodes

One of the most complex and fragile components of an EP system is the neutralising hollow cathode. These cathodes are integral to operation and historically the high temperature operation of these devices determines that they represent one of the most common failure modes of many EP systems[7]. The most commonly used cathodes for EP applications are hollow cathodes, which are utilised due to their ability to produce high electron current densities at relatively low powers, for thousands of hours of continuous operation. A diagram of a hollow cathode is shown in figure 1.5. These cathodes consist of hollow tubes of refractory metal closed at one end by a plate with an orifice. A hollow cylindrical insert of a low work function material is inserted into the tube in contact with the end plate[7]. A resistive heating element is wrapped around

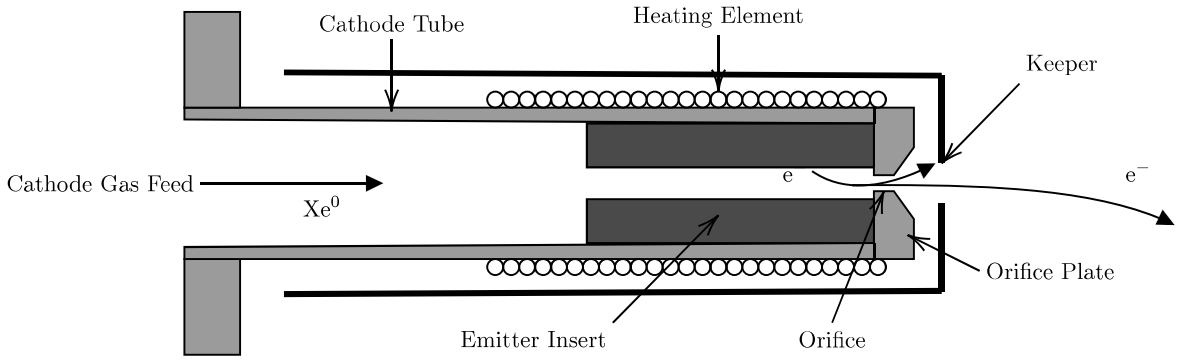

Figure 1.5: Main components of a generic hollow cathode. Adapted from Goebel and Katz[7].

the outside of the tube. The heater serves to elevate the temperature of the emitter insert, increasing thermionic emmission further. Next a gas is fed into the tube (generally xenon) which passes over the insert and is choked by the orifice. An electrode called a keeper is positioned just outside the orifice and set to a high positive potential. With sufficient heating, mass flow rate and keeper potential, electrons are drawn from the surface of the emitter to the keeper. The cathode is now considered ignited, as an electron current is pulled from the emitter surface. From the smaller keeper current, a larger electron current can now be drawn more easily by establishing a second positive potential downstream. In a HET, this electrode is the anode: drawing electrons from the cathode emitter insert down into the thruster channel. In both HETs and GIEs the external cathode (often referred to as the neutraliser) provides electrons to the ejected ions beam, ensuring the net charge of the spacecraft remains balanced.

The most common low work function materials used as the inserts of hollow cathodes for EP applications are Lanthanum hexaboride (LaB_6) or a porous tungsten matrix impregnated with a mixture of oxides of barium, calcium, and aluminium. The later types are composed of different ratios of Ba:Ca:Al, but collectively referred to as barium oxide (BaO) cathodes due to the importance of the barium oxide monolayer formed on the surface[7;15].

The inserts present a major source of failure, as they are one of the only parts of the system

which are directly dependent on chemical interaction with the propellant. Certain chemicals such as water, hydrocarbons or oxygen are absorbed by the insert materials, and form a dipole layer on the surface, changing the work function of the material, and thus reducing the effectiveness of the cathode[16]. This effect is broadly referred to as cathode poisoning. Poisoning can be mitigated by ensuring propellant purity and by operating the cathode at higher temperatures, as this can cause the dipole layer to evaporate at a faster rate than it is formed[17]. Many experimental demonstrations of EP systems operating on alternative propellants such as iodine still operate the cathode on xenon, to avoid possible chemical incompatibilities within the cathode[2]. This however suggest such a system architecture requires a small xenon tank to still be flown alongside the alternative propellant.

1.4 Electric Propulsion Propellants

In a HET the mechanism of propellant ionisation is electron bombardment: an electron collides with the neutral atom and knocks an electron from the atom's outer shell. To promote this ionisation electric propellants are chosen which have a large electron-impact ionisation cross section and a low first ionisation energy. The electron-impact ionisation cross section is a function of electron temperature and determines the probability that an ionisation will occur if both neutral and electron are modelled as rigid spheres[7]. The first ionisation energy determines the minimum relative energy that the incoming electron must have for an ionisation to occur.

Generally only monoatomic elements, or molecules with very simple structures are considered for electric propellants to reduce the number of covalent bonds. Covalent bonds can act as energy sinks, absorbing the kinetic energy of the colliding electron and storing it as vibrational or rotational energy of the molecule[18]. The energy absorbed to excite such a rotation or vibra-

tion does not contribute directly to thrust, and is considered wasted for the purpose of thrust generation.

Early EP systems used mercury as a propellant due to it having a low ionisation energy and large electron-impact cross section[19]. Mercury vapour is easy to ionise, and provides great thrust at a given power. Unfortunately mercury is extremely toxic to humans, and can be dangerous to spacecraft: backsputtering from an ion beam can cause a thin layer of mercury to coat the surface of a spacecraft[2]. Not only does the opaque liquid reduce the effectiveness of solar arrays and optical instruments, but the buildup of conductivity of mercury can cause short circuits.

Xenon holds an undeniable monopoly in the field of EP. All HETs currently operating in space operate on xenon, with the notable exception of the SpaceX Starlink mega-constellation which uses krypton (see below). Not only is the price of xenon already extremely high, but the historic volatility of this price could be highly damaging to the field of EP, especially with the rapid acceleration of EP adoption, and development of high power EP missions[6;20]. These factors have inspired calls for alternative EP propellants to be tested, in order to lower prices and raise competition in thruster development. Several other noble gases, such as krypton and argon, as well as metals such as bismuth and iodine have been considered as propellants[2]. Basic properties of these propellants can be seen in table 1.1, with electron impact ionisation cross sections shown in figure 1.6.

Xenon

Xenon gained wide adoption after being identified as a promising alternative to mercury. The odourless gas is the heaviest non-radioactive noble gas, meaning it is widely chemically nonreactive. Unlike all other noble gases xenon has the unique ability to be stored as a supercritical fluid

Table 1.1: Characteristics of several current and proposed propellants for electric propulsion. Adapted from Kieckhafer and King [2].

Propellant	Boiling Point (°C)	Ionisation Energy (eV)	Atomic Mass (amu)	Cost ($/Kg)
Xenon	-108	12.127	131.3	5000-8000
Krypton	-152	14.0	83.8	295
Mercury	357	10.434	200.59	50
Caesium	685	3.891	132.9	30 000
Argon	-186	15.8	39.94	3
Bismuth	1559	7.287	208.98	9
Iodine	182	10.44	126.9	500

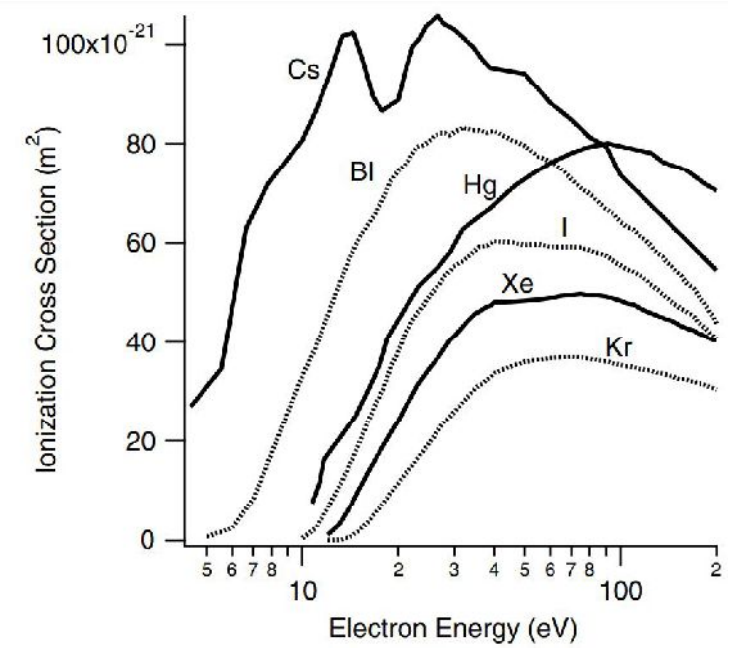

Figure 1.6: Electron impact ionisation cross section of alternative propellants. Reproduced from Kieckhafer and King [2] with permission of the author.

at relatively low pressures. Not only can xenon can be stored at the relatively high density of 1.6 gcm^{-}3 at 14 MPa, but being a supercritical fluid means that there is no boundary between liquid and vapour phases, meaning the fluid stays perfectly homogeneous and does not slosh within the propellant storage tank, reducing vibration.

Xenon comprises roughly 87 parts per billion of our atmosphere. It is commercially produced in very small quantities as a byproduct of oxygen liquefaction. Further separating and refining xenon requires vast amounts of energy, and as a result makes xenon an extremely expensive gas[21]. The supply of xenon is strongly dependent on that of liquid oxygen for steel production, meaning that the supply of xenon is detached from the demand, leaving the price to fluctuate violently. Raw material cost is not something that is usually considered a driving force in space engineering, however, the incredibly high and unpredictable price of xenon is an exception. NASA expressed their concern over the volatile price of propellant grade xenon after they observed a tenfold price increase within eight years, resulting in a price of several thousand Euros per kilogram at present[22]. This cost inflation can only be expected to become worse in the future, with very high power xenon thrusters currently being developed which require in excess of one hundred thousand Euros worth of xenon a day[20;2]. Lifetime tests up to one and a half times that of the expected lifetime of a thruster are relatively common in the industry and can last upwards of ten thousand hours. This suggests tens of millions of euros of xenon are required for a lifetime test at today's prices[2]. Although xenon is generally considered a green propellant, the vast amounts of energy required to separate xenon from the atmosphere produces a very large carbon footprint[21]. Separating and purifying a single kilogram of xenon from the byproducts of oxygen liquefaction requires between 1.3-4.3 mega watt hours of energy[23]. This is approximately as energy demanding as mining two Bitcoins as of February 2022[24]. This means that producing the xenon required for the Bepi-Colombo mission to Mercury generated a carbon footprint comparable to that produced by the average British person over their lifetime[23].

The extreme cost and scarcity of xenon is now driving many to search for more sustainable alternative propellants.

Krypton

Krypton has long been identified as the most similar alternative propellant to xenon, due to it appearing directly above xenon it in the periodic table. Krypton is the next heaviest nobel gas after xenon and shares a similar, yet smaller, electron-impact cross section as seen in figure 1.6. Krypton also has a similar first ionisation energy, and chemical inertness but at a considerably lower price than xenon. These similarities allow many thrusters and cathodes which were designed for xenon to operate easily on krypton, with largely similar results[25]. The only caveat withholding widespread krypton adoption is the significantly lower density when stored at temperatures and pressures typical of EP missions: $0.53\,\mathrm{gcm}{-3}$ for krypton compared to $1.6\,\mathrm{gcm}^{-}3$ for xenon, both at 50^{o}C and $14\,\mathrm{MPa}$[26]. There is therefore a system level trade-off in terms of the mass and volume penalty of the propellant tanks for high delta-V missions. SpaceX were the first to fly krypton HETs on their Starlink mega-constellation in 2019. This is due most probably to the high cost of supplying a mega-constellation with xenon, the moderate delta-V requirement of each individual spacecraft, and the company's plans to replace each satellite frequently.

Caesium

Caesium shares many of the characteristics of mercury, meaning great thrust at the expense of low I_{sp}. We can see from table 1.1 and figure 1.6 that caesium has both the largest cross section and also the lowest ionisation energy of any of the propellants surveyed, indicating it is comparatively very easy to ionise. However, like mercury, caesium is toxic to humans and

can backsputter and condense on spacecraft which causes issues. Furthermore caesium is so chemically reactive that it explodes when in contact with water, along with being radioactive. Note that caesium is also more expensive than xenon[2].

Argon

Argon HETs have shown to ignite easily and provide great cost savings over xenon. Unfortunately, they fall victim to a considerably lower anode thrust efficiency than xenon[27]. One of the largely limiting factors of argon is that the gas must be cooled and stored cryogenically in order to achieve a reasonable storage density: a heavy and energy expensive task for a spacecraft.

Bismuth

The benefits of bismuth are the extremely low cost, the large cross section, and the fact that it can be stored as a solid at an extremely high density of $9.747\,\mathrm{gcm}^{-}3$. Performances of bismuth HETs have been shown to have a greater thrust at a given power than xenon, at a similar I_{sp}[28]. The boiling point of bismuth is 1559°C and thus demands that great heat must be supplied to the propellant supply system, feed-line and anode in order to stop condensation of the propellant. This adds a nontrivial energy and thermal management demand to the system. This high boiling point also suggests that any bismuth back-streaming onto the spacecraft will condense and form an opaque, conductive layer causing similar issues to mercury.

Iodine

Iodine has recently enjoyed much attention as a contender for the next widely adopted EP propellant[29;30]. Iodine can be stored as a solid with a very attractive density of $4.9\,\mathrm{gcm}^{-}3$. We can see from figure 1.6 that the electron-impact cross section of iodine is greater than that of xenon and it has a lower first energy of ionisation[2]. A lower boiling point of 182°C means an iodine supply system requires less energy than one for bismuth. A great concern about utilising

iodine is that some of it will condense onto spacecraft surfaces. Such condensation has been shown to aggressively corrode many of the materials that make up spacecraft including stainless steels, aluminium, and titanium[31;32]. Iodine condenses on contact with the walls of traditional vacuum facilities, contaminating them for future experiments, or requiring a dedicated chamber. This adds significant cost and complexity to ground testing.

1.5 Chemical Compatibility of Cathodes

One of the biggest constraints when considering alternative propellants for electric propulsion systems is how the gas will interact with the cathode. All HETs require an external cathode to provide seed electrons to begin the plasma discharge, and neutralise the ion beam. The hollow cathode traditionally consumes the same high purity xenon as the thruster, yet at approximately 10% the mass flow rate[12]. Lanthanum hexaboride (LaB_6) dispenser cathodes are increasingly selected for EP systems due to their higher tolerance for impurities[16]. Even then, only the highest purity "propulsion grade" xenon which is 99.9995% pure must be used, as even trace quantities of impurities such as oxygen, hydrocarbons or water can poison the cathode and dramatically degrade performance[16]. Many alternative HET propellants have been tested with cathodes, such as oxygen, carbon dioxide, air, hydrogen, nitrogen, argon and iodine[16;32;33]. To date, only the nobel gases and hydrogen have shown to not poison traditional LaB_6 cathodes. Note that for the above experimental studies on HETs operating on bismuth, argon and iodine in each case the accompanying cathode was still operated on xenon[28;27;34]. To the knowledge of the author, krypton and mercury are the only alternative propellants which have been demonstrated to work with both the thruster and cathode simultaneously[25;35;36].

1.6 Search for a Novel Electric Propellant

The growing and volatile cost of xenon is not yet great enough to deter large prime contractors, yet high power EP systems will require tens of millions of Euros worth of xenon at today's prices, which are themselves highly dependent on demand[22]. Moreover, for small companies and start-ups trying to break into the EP market, the xenon required even for development and ground testing presents a costly hurdle on the path to thruster qualification. As the space industry faces increasing public scrutiny over sustainability, the large carbon footprint required to produce xenon can be considered a strong disadvantage. As the space sector grows, the extreme cost of xenon is now driving many to search for more scalable alternative propellants. However, there is no obvious alternative proposed propellant which meets all the requirements. To summarise, the ideal propellant must:

1. Have a lower cost than xenon, with a low environmental impact.
2. Have a high storage density, without requiring unfeasibly high pressures.
3. Be able to be ionised easily, meaning it must have a low energy of ionisation, and large electron impact cross-section.
4. Have few if any covalent bonds.
5. Not contaminate the spacecraft during back-streaming.
6. Be compatible with state of the art LaB_6 or BaO cathodes.

This list of requirements leaves few options, particularly due to the stringency of requirement number 6. Some suggest alternative cathode technologies must be developed to relax this requirement[37;33]. Others suggest circumventing the cathode compatibility issue by operating the thruster and cathode on different gases, for example operating a thruster on iodine but flying a separate xenon storage tank for the cathode[32]. SpaceX has chosen to compromise on requirement 2 above, with the entire Starlink mega-constellation operating on krypton. Note however

that Starlink is a highly unique situation, given the low delta-V requirements of each spacecraft and the uniquely low effective launch costs which allow unusually short mission lifetimes to be viable.

Chapter 2

Motivation and Objectives

The previous chapter provided context for the distinction between electrical and chemical propulsion and how electric thrusters, specifically HETs operate. The reader has now been familiarised with the need for an alternative propellant for EP, but also how stringent the requirements of such a propellant are. This presents the difficult and open question of how to abolish the monopoly that the expensive xenon holds over the field, as the only widely flown HET propellant.

This study proposes a new solution to the crux of finding an alternative HET propellant which meets all of the requirements discussed: water. Historically, water has been dismissed as a potential propellant due to being aggressively poisonous to cathodes, or as Goebel et al.[38] succinctly put it:

> "The most potent poisons for both [LaB_6 and BaO] cathodes are oxygen and water, with other gases such as CO_2 and air causing poisoning at higher partial pressures."

The concept described in this thesis circumvents this limitation by utilising an electrolyser in-situ to separate the stored water into oxygen and hydrogen, the latter of which is highly compatible with cathodes. The thruster itself consumes only oxygen, leaving the cathode to operate on pure hydrogen. From here on this concept shall be referred to as a water electrolysis Hall effect thruster, with the important distinction that the HET itself is operating strictly on oxygen. Such

a system allows the exploitation of the low cost, toxicity, and storage pressure of liquid water, while capitalise on hydrogen as the cathode working gas: a technology with a long heritage of terrestrial use[17;39;38].

This chapter discusses the system level advantages and disadvantages of the water electrolysis HET concept. Important differences are highlighted between an oxygen plasma and a xenon plasma, and how these factors can be expected to impact HET performance. Synergies between a water electrolysis HET system and chemical propulsion systems or fuel cell energy storage systems are considered. Other HETs designed for molecular propellants are discussed to identify potential challenges of this concept.

2.1 Hydrogen Compatibility with Hollow Cathodes

Gallagher[16] tested the susceptibility of LaB_6 hollow cathodes to become poisoned by testing the devices while operating on a range of different gases and measuring the extracted electron current. Increasing the temperature of a hollow cathode generally increases the emission current at the expense of a greater power demand. Oxygen proved to be the gas which led to the most aggressive poisoning, as seen in Figure 2.1. These results show that even when the cathode is operated at extreme temperatures of 1570°C, oxygen in as small pressures as $1x10^{-3}$ Torr ($1.3x10^{-3}$ mbar) can reduce the emission current of the cathode by more than 50%.

Gallagher[16] also included hydrogen as one of the cathode test gases, as shown in Figure 2.2. Hydrogen only poisoned the cathode slightly, and only at very high concentrations. Furthermore, the operation of a cathode on hydrogen dramatically increased the performance of the device. Hydrogen appeared to reverse the poisoning of the cathode, in a process that Gallagher assumed to be hydrogen-ion bombardment. These results suggest that even in close physical proximity of poisoning gases such as oxygen, a LaB_6 cathode may be able to operate nominally

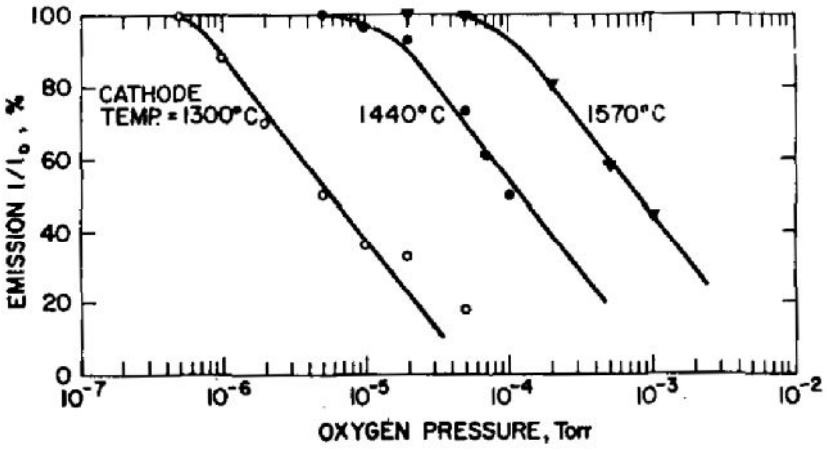

Figure 2.1: The poisoning effect of oxygen at various background concentrations, and at various temperatures. Reproduced from Gallagher[16], with the permission of AIP Publishing.

on hydrogen.

Several terrestrial applications use LaB_6 cathodes operating on hydrogen and its isotopes for plasma and fusion research. These include plasma production and neutral beam injection for tokamaks[40;41] and high flux deuterium plasma beam production in linear plasma generators[42]. These applications have prompted a wealth of research into high-power LaB_6 cathodes operating on hydrogen since the 1970s[17;39].

2.2 Hydrogen Storage For In-Space propulsion

Although hydrogen has demonstrated great compatibility for terrestrial applications, it shares the same issues as the nobel gases (with the special exception of xenon), in that it cannot be stored as a liquid without very high pressures or cryogenics. Being the lightest of all gasses, hydrogen has an extremely low storage density. Some launch vehicles such as the Space Shuttle increase the density of hydrogen by cryogenically cooling the gas into a liquid and consuming it rapidly before it is able to evaporate. However, even liquid hydrogen has a density of only $0.071\,\mathrm{gcm}^{-}3$. Liquid hydrogen storage is not feasible for in-space propulsion, given that long-

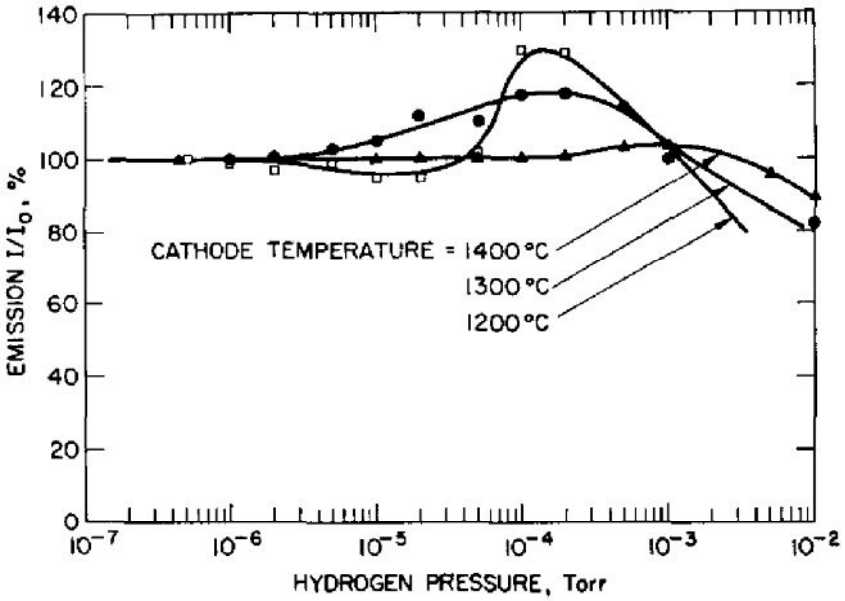

Figure 2.2: The poisoning effect of hydrogen at various background concentrations, and at various temperatures. Reproduced from Gallagher[16], with the permission of AIP Publishing.

term storage requires heavy and power hungry systems of active cooling of stored propellant.

A proposed solution to the difficulty of storing hydrogen in space is by utilising an electrolyser. An electrolyser is an electro-chemical device which drives an electric current to undergo an otherwise nonspontaneous chemical reaction. For this thesis, the focus will be on water electrolysis, which consumes electricity and liquid water to produce gaseous hydrogen and oxygen[43;44]. Note that electrical energy must be expended as the chemical potential of the products is greater than that of the reactant.

One litre of water contains more hydrogen than one litre of liquid hydrogen: H_2O comprised $2.01\,u$ hydrogen and $15.99\,u$ oxygen, meaning that approximately $1/9^{\text{th}}$ the mass of water is attributed to the contained hydrogen, and $8/9^{\text{th}}$ oxygen. While a litre of liquid hydrogen weighs $70.85\,\text{g}$, a litre of water contains $111.8\,\text{g}$ of hydrogen. This dictates that if one momentarily disregards energy costs, generating hydrogen from stored liquid water by in-situ electrolysis presents a convenient long-term hydrogen storage solution for space applications. The chemical

propulsion community has identified this indirect method as an effective means of exploiting the high performance of hydrogen, without the need for active cryogenic cooling. Water electrolysis chemical propulsion is discussed further in section 2.6 below.

2.3 Water Electrolysis as an Electric Propellant

This thesis presents the novel concept that water could indirectly be considered an alternative propellant for a HET using an electrolyser. The term water electrolysis HET is used throughout this text to describe this configuration. In such a system, water is stored as a liquid on board the spacecraft at launch at low pressures. Whenever the mission requires thrust, an electrolyser consumes power to convert the liquid water into gaseous hydrogen and oxygen, with the anode of the HET being operated only on oxygen, and the cathode operating only on hydrogen. Given that gases are consumed as they are generated, there is no need for gas-stage storage. The stoichiometric composition of water determines that electrolysis always produces a fixed 7.94:1 ratio by mass of O_2:H_2. This dictates that a water electrolysis HET would need to operate at a fixed cathode-to-anode flow fraction of 1/7.94 which is 12.6%. Traditional xenon thrusters operate their cathodes at 2.8-15% the mass flow rate of the thruster[12], so it would appear that the fixed 12.6% fraction produced by water electrolysis is coincidentally a convenient rate. However, it must be proven experimentally that this ratio is still effective when not operating on xenon.

Let us consider how the water electrolysis HET concept meets the 6 requirements set out in section 1.6:

1. In terms of raw material cost and environmental impact, water is perhaps the most attractive propellant conceivable.

2. When stored as a liquid, water has a moderate density of 1 $gcm^{-}3$, yet with the major advantage of low pressure requirements allowing low mass, thin walled storage tanks. Note however that extra mass and power must be allocated for the electrolyser.

3. The oxygen which passes to the thruster has a slightly lower energy of ionisation than xenon (12.07 eV and 12.13 eV respectively), but a considerably smaller electron impact ionisation cross section. The implications of this are discussed in the following section.

4. The oxygen being ionised has a single covalent bond which is less favourable than a monatomic atom like xenon. The impact of this is discussed below.

5. Oxygen, hydrogen and water will all sublimate in space at temperatures representative of a spacecraft, so back-sputtering is not an issue.

6. The hollow cathode will exclusively consume hydrogen. Hydrogen has been shown to be one of the very few gases compatible with LaB_6 hollow cathodes, backed by a wealth of research from terrestrial applications.

An electrolyser requires energy to dissociate water into its constituent elements. This is an energy cost not present in a typical xenon HET, or any other proposed propellant. The discharge of an SPT-100 HET operating on 50 sccm (standard cubic centimetres per minute) of xenon draws 1.35 kW of power from the spacecraft[12]. For an electrolyser to produce 50 sccm of oxygen, it requires around 30 W of power, assuming a conservative electrolyser efficiency of 80%[45]. This suggests approximately 3% additional power must be expended to generate propellant as it is consumed for a water electrolysis HET.

A further unique system-level penalty that a water electrolysis system must incur is the added mass and complexity of an electrolyser. The electrolysis of water in space is already at such high technology maturity that it is now used for critical life support on board the International Space Station as the primary means of generating oxygen for astronauts[46]. As a rough estimate for

the mass of a 30 W electrolyser, let us consider the 4.9 kg engineering model electrolyser built for space applications by Farnes et al.[44] which operates at 4.2 kW and 90% efficiency. Alternatively Fernandes et al.[47] suggest a 1595 W electrolyser designed for propulsion would weigh either 2.2 kg or 9.4 kg depending on which configuration is used. If linear scaling is assumed, the considerably smaller 30 W system that a water electrolysis HET would require would therefore weigh between 35 g and 143 g. While it is not suggested that such a simple linear scaling is realistic, even if this estimation is off by a factor of 3 or 4, the added mass of an electrolyser is still small for any spacecraft with a 1.35 kW class thruster.

A water electrolysis HET has several advantages over other propellants. The main motivation for the search for alternatives to xenon is the unstable price and fixed production rate of the nobel gas, particularly as EP becomes more popular. These are obviously not issues when considering water as an electric propellant. Gaseous propellants such as xenon and krypton must be stored at high pressures to achieve reasonable storage densities[26]. Water on the other hand has a moderately high density compared to other propellants, and can be stored at a low pressure alleviating the need for heavy, high pressure propellant tanks. To compare the mass of propellant tanks we consider the catalogue of flight qualified water and xenon tanks published by Rafael Advanced Defence Systems[48]. Here a typical 9.1 L spherical xenon tank has a dry mass of 5.2 kg, as it must operate to 162 bar at beginning of life. At this pressure such a tank can house approximately 15 kg of xenon. A water tank of a comparable volume weighs less than 1.7 kg, with an operating pressure of 24.4 bar (blowing down to 5 bar by end of life). A water tank which holds a similar propellant mass as the 9.1 L xenon tank is the 17.5 L water tank, which weighs less than 3.1 kg. This data suggests that a propellant dry mass saving of approximately 40 % can be achieved by switching from xenon to water.

Both a high chemical reactivity or high storage pressure can cause complications during space-

craft integration. Any energy stored within the propellant could be released accidentally and is therefore considered a potential hazard. This is particularly true for ride-share missions, with many launch providers having strict regulations regarding the stored internal energy of any ride-share spacecraft, such as high pressure propellant storage, should this energy be released accidentally and damage other spacecraft on the launcher[30]. Water can be considered a passive propellant, having a low storage pressure and low chemical reactivity. This dictates that a water electrolysis HET offers a reduction in the cost, risk, and complexity of spacecraft integration and ride-share mission planning.

Alternative solid propellants such as metals or iodine can also avoid the need for high-pressure storage tanks. These systems sublimate the propellants into a gas as it is required by the propulsion system. However, this requires heating of the storage tank along with the entire propellant management during operation to avoid the propellant desublimating back into a solid and causing blockages[30;28]. Water systems may also require heating to stop freezing of the tank and lines, yet the much lower freezing point of water means less power is required than heating iodine lines. Given traditional spacecraft work to keep electronics near room temperature, the large thermal inertia of water will help to ensure the propellant does not freeze. While solid propellants like iodine and bismuth can be deposited on spacecraft surfaces from backstreaming, this is not an issue for a water electrolysis HET, as the plume is expect to consist not of water but oxygen and hydrogen, neither of which will condense.

2.4 Consideration When Operating a Hall Effect Thruster on Oxygen

The water electrolysis HET concept relies on the thruster itself operating on pure oxygen. Unfortunately, simply operating a traditional xenon optimised HET on oxygen is not a fair test of the feasibility of such a system. The plasma characteristics, and resulting performance of a thruster is highly dependent on both thruster geometry and operating point. It follows that existing HET designs are highly optimised to sustain a xenon plasma, and operate poorly, if at all on other propellants without significant re-optimisation.

The geometry of the discharge channel of a HET determines the surface area to volume ratio of the plasma. This ratio is extremely important, as the ceramic walls of a thruster facilitates de-excitation and recombination of heavy particles and secondary electron emission[49]. A deeper discharge channel represents a longer path that a neutral particle must travel to escape the thruster, and thus increases the opportunity for the particle to be ionised. However, extending the channel also increases the surface area. This explains why the geometry of the channel is such a delicate balance: increasing the depth increases both the ionisation length, but also the wall area where recombination can occur. Increasing the channel width will decrease neutral density, but increase the bulk plasma volume. Given different propellants ionise and recombine at different rates, it follows that thruster geometry optimisation is very propellant dependent.

Although designing a HET for oxygen operation has never been attempted to the knowledge of the author, studies of traditional xenon HETs operating on oxygen have been reported[50;51]. These are described in detail in section 2.5, but the general trend is that the devices operate at a greater current but considerably lower efficiency on oxygen rather then xenon. These results provide evidence that a HET must be designed specifically for oxygen in order to accurately

assess the potential feasibility of such a technology.

The intrinsic differences between xenon atoms and oxygen molecules are guaranteed to have a large impact on thruster performance. The most obvious discrepancies are the atomic mass difference and its associated impact on the gyroradius of the ion, the difference in electron-impact ionisation cross section and the diatomic nature of oxygen.

Low atomic mass propellants

An oxygen molecule has a molar mass of 32 u, as compared to 131 u for a xenon atom. The thrust T of a HET can be estimated by

$$T = \sqrt{\frac{2m_i}{e}} I_b \sqrt{V_b}, \tag{2.1}$$

where m_i is the mass of the accelerated ion and e is the elementary charge. I_b and V_b are the beam current and beam voltage respectively[7]. This equation shows that for a less massive ion, the device must operate at a higher current or voltage to produce the same thrust. This requires an increase in the power requirements of the device, lowering the Thrust To Power Ratio (TTPR) of the system. Note that equation 2.1 is a very simple first order approximation which ignores several efficiency factors. I_{sp} on the other hand is estimated by:

$$I_{sp} = \sqrt{\frac{2qV_b}{m_i}} \frac{1}{g_0}, \tag{2.2}$$

where q is the ion's charge, and g_0 is the acceleration due to gravity at the surface of Earth. It can be seen that a lower ion mass will produce a higher I_{sp}, if the beam voltage is left unchanged. It is difficult to estimate the expected relations between the discharge current and discharge voltage of an oxygen thruster. However, it can be concluded that thrust is proportional to the

square root of the molar mass, whereas I_{sp} is inversely proportional to the square root of the molar mass[7;2]. As a very first order estimate, in comparison to a similar xenon device, a HET optimised for oxygen can be expected to sacrifice approximately half its thrust in return for almost twice the I_{sp} at a given power.

Gyro-radius Reduction

The magnitude of the radial magnetic field of a HET must be chosen very carefully. The purpose of this field is to magnetise the electrons. This means that the electrons are captured by the radial magnetic fields and gyrate about the field lines as they drift in the azimuthal direction. This greatly extends the distance that the electrons must travel from the cathode to the anode hereby increasing the chance of a successful ionisation. Consider equation 1.4 from the previous chapter, which gives the gyroradius of a charged particle. A magnetic field strength must be selected such that the gyroradius of the electron is much smaller than the thruster channel width, which guarantees that the electron can freely orbit the field lines indefinitely without colliding with the channel walls (assuming the electron velocity has no radial component). When this condition is met, we say that the electrons are magnetised.

There is however an upper limit to the magnetic field strength. The ions also gyrate about the field lines as described by equation 1.4, yet due to the ions having a larger mass, this radius is far larger. A large gyroradius means that the ions are barely deflected by the magnetic field and maintain a mostly straight trajectory as they leave the thruster. Ideally the gyro radius of the ions should be far larger than the channel width. If one were to increase the field strength too greatly, the ion gyroradius would become reduced to the point that the ion trajectory is severely affected as they leave the thrusters. In extreme cases, the ions are deflected so strongly that they collide with the walls and deposit their energy into heat. In this scenario, we say the ions are magnetised, or alternatively that the plasma is fully magnetised. In a HET one strives to produce a magnetic field strength such that the electrons are magnetised, but the ions are not.

This can be expressed as:

$$r_{ge} < L \ll r_{gi}, \tag{2.3}$$

or alternatively, using equation 1.4 a magnetic field stregth must be selected such that:

$$\frac{m_e v_\perp}{|q|} < BL \ll \frac{m_i v_\perp}{|q|}, \tag{2.4}$$

where r_{ge} and r_{gi} refer to the electron and ion gyro radius respectively, and L is the thruster channel depth. Given that oxygen has an atomic mass of approximately 1/4th that of xenon, equation 1.4 dictates that the oxygen ions will have 1/4th the gyroradius at a given field strength and axial velocity. The oxygen ions are expected to have a higher velocity than the xenon ions as they leave the thurster, because of their lower mass and assuming a roughly equivalent electric field. However, the velocity of the ion is expected to be inversely proportional to $\sqrt{m_i}$, where the gyroradius is directly proportional to m_i, meaning that a reduction in ion mass is still expected to lead to a smaller gyro radius at a given operation point. This means that the oxygen plasma will become fully magnetised at a lower magnetic field strength lower than that of a comparable xenon plasma. This may mean that in comparison to a xenon thruster, an oxygen thruster will be required to operate at a lower field strength, and thus the electrons would be magnetised to a lesser degree. The ideal magnetic field strength of a HET is very difficult to predict, and is best found experimentally by changing the current of the electromagnet.

Electron-impact Cross Section

In a HET, the neutral propellant is ionised by an electron impacting the neutral particle with sufficient energy to remove another electron, creating a positive ion. The minimum energy needed to ionise a neutral particle in this way is called the first energy of ionisation. This value is very similar for the two species: 12.07 eV for oxygen and 12.13 eV for xenon. However, this does not account for the fact that the xenon atom is considerably larger than an oxygen

molecule and therefore presents a larger target for the impacting electron. A better metric is the electron-impact ionisation cross section, which is a measure of the probability that a collision of an electron with a particular kinetic energy will result in an ionisation[18]. The derivation of cross section is a stochastic process, yet for simplicity, the units of area are used, with the cross section that the particles would have if they were modelled as perfect colliding spheres. The larger the cross section, the greater the chance that neutral particles will be ionised as they travel down the channel of the thruster. The electron-impact ionisation cross sections for producing a molecular oxygen ion, an atomic oxygen ion and a xenon ion are shown in figure 2.3. It can be seen that oxygen has a considerably lower electron-impact ionisation cross section. This means that for the same electron temperature, one can expect a lower fraction of the neutral particles to be ionised within the thruster, reducing the propellant utilisation efficiency and anode thrust efficiency of the device. Another interesting feature shown in red in figure 2.3 is that the probability of producing monatomic ions (O^+) as opposed to diatomic ions (O_2^+) increases as the electron energy increases. This is important as the reduced mass of the monatomic ions will further increase I_{sp} at the expense of thrust.

The much smaller electron-impact cross sections of oxygen, when compared to xenon, are expected to have a large impact on performance. Compensating for this discrepancy could be attempted by increasing thruster channel depth, as this would increase the residence time of the neutral particles within the device, increasing the likelihood of collisions. Alternatively the depth of the region in which the magnetic field is imposed, termed hereafter as the magnetic region, of the HET can be increased, such that the neutrals spend a longer period in the section of the thruster with the greatest electron density and temperature.

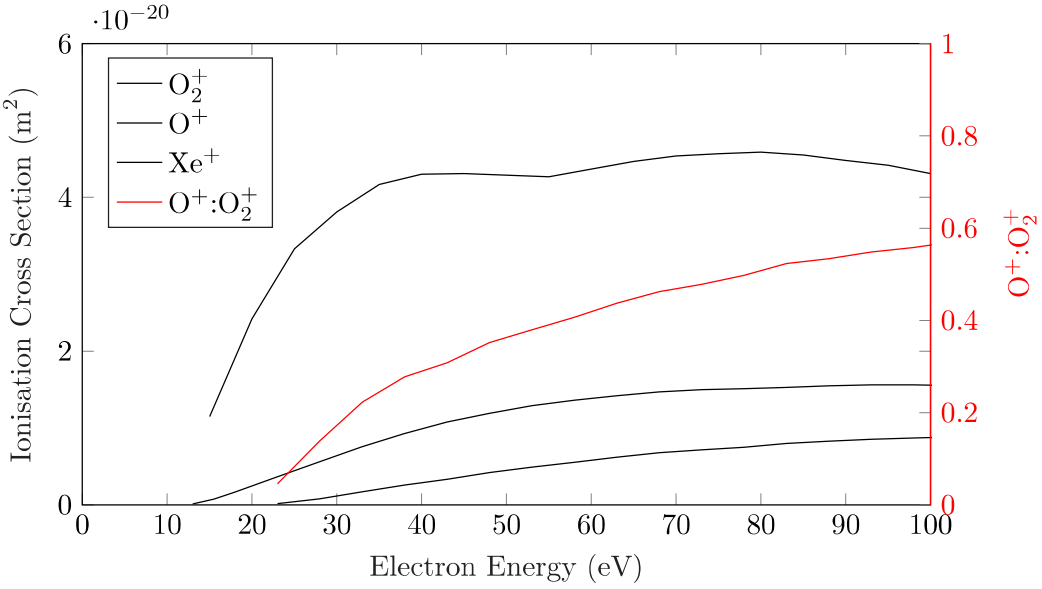

Figure 2.3: Comparison of oxygen and xenon electron-impact ionisation cross sections. Created using data from Itikawa[52] and Stephan and Märk[53]. The ratio of atomic to diatomic oxygen cross section is shown in red.

Diatomic Molecules

Oxygen forms a diatomic O_2 molecule, where xenon exists as a solitary atom. This crucial difference determines that, when struck by an electron, these two neutral particles behave very differently. During an inelastic collision between an electron and a xenon atom in a HET, the kinetic energy of the colliding particles either goes into the electrical excitation of the atom, or the ionisation of the atom[18]. Xenon particles which are ionised are accelerated by the electric field to exit the thruster at speeds around $15\,\mathrm{kms}^{-1}$, which is considerably faster than neutral xenon atoms which generate little thrust as they diffuse into space at thermal velocities. An electrically excited xenon atom has no electric charge, and thus acts like a neutral particle in an electric field by contributing little to thrust generation. It is for this reason that energy deposited into the electrical excitation of a xenon atom is considered to be wasted for the purpose of generating thrust. In summary, one seeks to maximise the number of xenon-electron collisions that result in the creation of ions, and consider electrical excitation of neutral xenon particles an unwanted waste of energy.

Unlike xenon, oxygen naturally forms a diatomic molecule, meaning the pair of oxygen atoms are bonded by a covalent bond. This bond can be thought of as being elastic with the ability to store energy. When an electron strikes an oxygen molecule, this bond can absorb the kinetic energy of the collision, resulting in the two oxygen atoms vibrating or rotating in relation to one another[18]. Such an electron-molecule collision increases the internal energy of the molecule, but does not contribute to thrust, and is thus also considered wasted. This is an important distinction as it is a mechanism that energy can be lost in a oxygen plasma but not a xenon plasma.

This shows that in comparison to a xenon atom, there are several more degrees of freedom where energy can be deposited into an oxygen molecule without contributing to thrust generation: energy can go into rotational, vibrational, or electrical excitation. To further complicate matters further, the collision of an energetic electron and an O_2 molecule can dissociate the covalent bond, forming two oxygen atoms, each of which can subsequently be electrically excited or ionised. Dissociation presents a more complex issue, since the generation of two atomic oxygen ions (O^+) will lead to a higher exit velocity than as a single molecular ion (O_2^+), but requires more work by the power supply. Macroscopically, the creation of atomic oxygen ions will cause a greater I_{sp} due to the increased exit velocity, but at the expense of a more power-hungry system, as the charge-to-mass ratio of the ejected ions is increased. Determining which fractions of the ions will be atomic is crucial in determining the overall performance that can be expected from the system.

2.5 Known Oxygen Plasma Thrusters

To the knowledge of the author, no HET has been specifically designed to be operated on oxygen. However, there are some HETs that have been designed for xenon, which have been operated on oxygen, or a nitrogen-oxygen mixture.

2.5.1 Hall Effect Thrusters Operating on Oxygen

Few studies have demonstrated the operation of a HET on oxygen only. Frisbee et al.[50] discusses the potential to operate a HET on oxygen generated from lunar regolith for the purpose of an electric lunar transfer vehicle. This study focuses mainly on vehicle sizing and potential mission scenarios, but also reports the results of experiments with an oxygen HET at the Moscow Aviation Institute (MAI). The study referenced is no longer accessible, but Frisbee et al.[50] describe experiments with a thruster of 47 mm channel diameter, which had a nominal operating efficiency on xenon of 16%, using 1.2 mgs^{-1} mass flow rate, 300 V discharge voltage and 1.3 A discharge current. When switching to oxygen, the efficiency was reduced to 5% at 300 V and 2.4 A. The highest efficiency that the thruster was able to reach under oxygen operation was 6.6% at 1.22 mgs^{-1}, 360 V and 2.3 A, however the authors expected these numbers could be improved at higher powers and optimised thruster designs. Regarding cathode operation during MAI tests, Frisbee et al.[50] report that an inert gas (presumably a noble gas) was used; however, even with an inert gas, the lifetime of the cathode was expected to be on the order of tens of hours due to oxidation from the thruster.

The reported results of the MAI tests describe a decrease in efficiency when switching from xenon to oxygen, which is consistent with the predictions made in section 2.4 above. The covalent bond in the oxygen molecule provides an energy sink, which results in less input energy contributing to thrust. Note also that at a discharge voltage of 300 V, the discharge current of oxygen was 2.4 A against the 1.3 A of xenon when operating on similar mass flow rates. This suggests that almost twice as many ions are being generated and accelerated in the oxygen plasma.

Nakagawa et al.[51] provide a more comprehensive study of a HET operating on oxygen, with the cathode operating on xenon. Please note that this study exists only in Japanese and that

this review is based on a software-generated translation, which may not be 100% accurate. Similarly to the Frisbee et al.[50] study, the motivation for the Nakagawa et al.[51] study was the consumption of oxygen produced from lunar regolith. The thruster used was again designed for xenon operation. This HET had a channel diameter of 60 mm, a width of 5 mm, and an adjustable depth. Nakagawa et al.[51] were able to operate the device on oxygen at a maximum I_{sp} of 1000 s and an efficiency of 9%. For comparison the same thruster produced 1600 s of I_{sp} and an efficiency of 22% when operating on xenon. Experimental optimisation lead to a channel depth of 12 mm for both the 0.66 mgs^{-1} and 0.99 mgs^{-1} mass flow rate surveyed, which they report is longer than the optimal channel depth for xenon. The authors of the study also suggest that the reason for an optimal channel depth is that increasing the depth leads to a greater volume and therefore a greater chance of ionisation, but that this also leads to a greater loss of ions to the walls. Electron temperature measurements conducted by Nakagawa et al.[51] showed 10 eV for xenon and 25 eV for oxygen, as the oxygen thruster required a higher discharge voltage for a given current. Using the electron-impact ionisation cross sections of each species at these electron temperatures, the authors of that study calculated that oxygen plasma has a ionisation mean free path of 9 mm, where the xenon plasma has a mean free path of 0.5 mm. This mean free path is the average distance that an neutral will travel before it results in a successful ionisation[18]. Nakagawa et al.[51] claim that oxygen having a longer experimentally optimised channel depth than xenon is a direct consequence of the longer ionisation mean free path.

These studies confirm suspicions that while a traditional HET which has been optimised for xenon operation can be operated on oxygen, the result is a considerably lower efficiency. Contrary to the assumptions in section 2.4, it appears that the change of propellant from xenon to oxygen leads to a reduction in I_{sp}. The experimental optimisation of channel depth performed by Nakagawa et al.[51] provides strong evidence that a thruster optimised for oxygen operation would differ from one for xenon operation.

2.5.2 RAM-EP

The concept of an air-breathing EP system, also known as a RAM-EP system, was extensively studied by the ESA Concurrent Design Facility, and later published by Di Cara et al.[54]. Similar studies by JAXA and Busek have also been conducted[55;56]. The concept relies on the premise that for several applications a reduction in orbital altitude is favourable. For example, Earth observation spacecraft can directly increase their resolution, and LEO constellations can reduce their latency and telecommunications power when operating closer to the Earth. However, at altitudes below 400 km, the residual atmosphere begins to produce a non-negligible drag force. One of ESA's flagship missions was the Gravity Field and Steady-State Ocean Circulation Explorer (GOCE), which used xenon GIEs to counteract this drag. GOCE successfully operated for four years in Very Low Earth Orbit (VLEO) dipping down to orbits as low as 250 km until re-entering when it ran out of xenon[57]. The RAM-EP system describes a similar VLEO orbit, however here the spacecraft is able to capture the rarefied atmosphere, and use it as propellant for an EP system. Such an air-breathing spacecraft could operate at very low orbits indefinitely by counteracting drag by capturing and accelerating trace amounts of oxygen and nitrogen.

The RAM-EP concept is of interest to the present study, given the similarities between oxygen and nitrogen: like oxygen, nitrogen forms a diatomic molecule, has a mass of 14.0067 u compared to 15.999 u for oxygen, a first energy of ionisation of 15.58 eV against 12.07 eV for oxygen and has a similar electron impact ionisation cross section. By considering nitrogen an analogue to oxygen, research done on RAM-EP systems can provide further insight into how one could expect an oxygen HET to operate.

Cifali et al.[58] operated a traditional HET on a mixture of nitrogen and oxygen at a ratio of $1.27N_2$:O_2. The thruster chosen for this test was a typical xenon optimised PPS-1350 produced by Snecma, which is in-effect a European clone of the Russian SPT-100[59]. The thruster

was first ignited on xenon, then slowly transitioned to the N_2/O_2 mixture. The cathode was continually operated on xenon. Thrust up to 30 mN was measured at a massflow of 2.76 mgs^{-1}, which results in an I_{sp} of 1100 s and an anode thrust efficiency of 11.3% at discharge powers of approximately 1.4 kW[58]. These results fall in line with those performed on pure oxygen, in that both thrust and I_{sp} are considerably lower for molecular propellants than xenon for a traditional HET operating at a similar power level.

Marchioni and Cappelli[60] appear to be the first group to publish experimental results of a HET optimised to operate on air for a RAM-EP application, along with a Masters thesis by Marchioni[61]. They argue that because molecular propellants such as nitrogen and oxygen have a mass smaller than that of xenon, the neutral particles will have a higher thermal velocity and thus a shorter residence time in the channel. This reduced residence time, coupled with the smaller ionisation cross section will result in less frequent ionisation, and ultimately a lower propellant utilisation. To combat this, Marchioni and Cappelli[60] built the Extended Channel Hall Thruster (ECHT), with a deeper ionisation channel. The device has a channel 86 mm deep, 10 mm wide and 100 mm in diameter. Unlike a typical HET which only has a radial magnetic field directly near the channel exit, the ECHT has magnets which extend 50 mm from the channel exit, such that a longer ionisation region is generated. The thruster was operated on 2 mgs^{-1} of pure nitrogen with the accompanying cathode operating on argon. The highest performance achieved by the ECHT was a thrust of 22.79 mN, at 200 V discharge voltage and 3.5 A discharge current, which results in a specific impulse of 1116 s, and a thrust efficiency of 23%. This is considerably greater efficiency than non-optimised thrusters have managed to attain on molecular propellants, indicating the ECHT design has been successful at increasing ionisation efficiency. Marchioni and Cappelli[60] have clearly demonstrated that vast improvements in performance can be attained by optimising a thruster for molecular propellants. The authors of the study plan on developing the design further, by extending the channel deeper and increasing the magnetic

field strength.

2.6 Chemical Propulsion Using Water Electrolysis

The theoretical vacuum I_{sp} of a H_2/O_2 chemical propulsion system is over 450 s, which is one of the highest ever demonstrated. Some of the most famous launch vehicles have taken advantage of this, such as the main engines of the now retired American Space Shuttle and the main stages of the European Ariane 5 and 6[62;63]. One does not often see H_2/O_2 chemical propulsion utilised for in-space chemical propulsion, because both hydrogen and oxygen require cryogenic temperatures to be storable as a liquid. This means that generally H_2/O_2 propulsion systems are only suitable for launch vehicles that do not need to store these gases for an extended period.

The chemical propulsion community is exploring storing water as a liquid and only producing hydrogen and oxygen through electrolysis as they are needed, which presents an elegant method to benefit from the high performance of a H_2/O_2 system without the need to store cryogenic gases. It is important to note that all energy that is released by combusting the propellants must first be expended by the electrolyser to generate the gases. These systems can be operated in pulsed or a continuous method. In the pulsed mode the electrolyser is powered to fill small on-board plenum tanks with hydrogen and oxygen over long periods of time, which are then quickly consumed during a short chemical burn. A continuously fed system has the benefit of not requiring gas stage storage, yet is limited by the mass flow rate or power consumption limit of the electrolyser.

Water electrolysis chemical systems like this have been tested in the laboratory since the 1970s[64], with high TRL systems expected to fly in the very near future[65;47]. Our research group is currently also developing such a propulsion system: the Iridium Catalysed Electrolysis thruster is a highly modular and scalable thruster concept that uses micro-electromechanical system produc-

tion methods and platinum group catalysts to combust hydrogen and oxygen at the stociometric rate that they are produced from water electrolysis[45;66]. The thruster has now demonstrated operation at over 180 s of I_{sp}.

2.7 Water as a Multi-mode Chemical-Electrical Propellant

Operating a HET on the products of water electrolysis would enable water to be used in a shared propellant multi-mode chemical-electrical propulsion system. Multi-mode chemical-electrical propulsion (MMCEP) describes a spacecraft that uses a single shared propellant for both an electrical propulsion system and a chemical propulsion system[67]. In our context, a MMCEP would store liquid water at launch and uses an electrolyser to produce small quantities of gaseous hydrogen and oxygen in space. From here, the gases can be consumed by the electric propulsion system to perform high I_{sp} manoeuvres, but also by the chemical propulsion system to perform high thrust manoeuvres. The advantages of a MMCEP is that both high specific impulse electric burns as well as high thrust chemical burns can be performed, yet the mass penalty of flying two independent propulsion systems is reduced due to sharing hardware such as propellant tanks, propellant management systems, and in our case a single electrolyser. Rovey et al.[67] reviewed different MMCEP concepts along with the benefits they offer. Although no true MMCEP concept has yet been flown, the propulsion architectures suggested include hydrazine for monopropellant thrusters and arcjets[68], HAN for a monopropellant systems and a pulsed plasma thruster[69], as well as ionic liquid fuels used in both monopropellant and electrospray systems[70]. None of the concepts in the review utilise water as a MMCEP propellant.

When spacecraft operators need to choose between an EP or a CP system, a compromise must generally be made given that some orbital manoeuvres are best suited to high impulse burns and others to high thrust burns. For example, an all-electric spacecraft cannot perform the high thrust manoeuvres required for an orbit-raising Hohmann transfer, meaning that electric

orbit raising must rely on the much slower and less efficient spiral orbital transfers[3]. On the other hand, a fully chemical propulsion spacecraft may reserve a large portion of propellant for north-south station keeping being performed by a low I_{sp} thurster. An MMCEP architecture frees one from having to make such a compromise, by using each propulsion system when it is best suited. Take, for example, an exploration mission that could use an EP system for long deep-space cruises, yet employ a chemical propulsion system for orbital ejection and insertion. Low Earth orbiting spacecraft can utilise the EP system for station keeping but use the chemical propulsion system for inclination changes, collision avoidance and de-orbiting. Space mining missions can perform cruise missions with electric propulsion, but then switch to high thrust chemical systems for proximity flying and landing.

Other advantages of a MMCEP system is that sharing the propellant between the two propulsion systems offers the spacecraft operators a dramatic flexibility before or even after launch. Without pre-allocating the propellant to either the chemical or the electrical propulsion system during the mission design means the same spacecraft has the ability to adapt to very drastic mission changes. An example of where this could be beneficial is an on-orbit servicing mission which must react to the needs of maintaining a mega-constellation or an exploration or mining mission which is able to decide how best to alter its mission trajectory based off observations made during the mission. Rovey et al.[67] also highlights that the adaptability and flexibility of an MMCEP system can have a large financial impact on reoccurring development costs, given that a single one-size-fits-most MMCEP system can eliminate the need to re-design and re-qualify a propulsion system for a range of different missions.

2.8 In-situ Resource Utilisation

An additional benefit of water as a propellant is the potential for future In Situ Resource Utilisation (ISRU). This is the term given to the act of collecting and processing raw materials in space for use as consumables or construction material[71]. Water has long been identified as one of the materials best suited for ISRU, because of its ubiquity throughout space, its importance for life support, and its importance as a precursor for chemical propellants. Water has been found on a large number of planets and moons throughout our solar system, including two of our closest neighbours: the moon and Mars[72]. We know four bodies in our solar system that contain more liquid water than Earth, as shown in table 2.1. To date ISRU has never been demonstrated. In the far future a water electrolysis EP system may one day benefit from the ISRU of water to refuel in space. This would be of particular use to asteroid mining missions given that many near Earth asteroids are presumed to contain water[73;74].

Table 2.1: Water content across the solar system with respect to world volume. Taken from Mosher and Cheng[75].

Body	**Liquid Water (x10^{21}L)**	**Body Diameter (km)**	**Liquid water by Body Volume (%)**
Enceladus	0.01	252	20
Triton	0.03	1352	0.3
Dione	0.14	561	19
Pluto	1.0	1187	15
Earth	1.34	6371	0.12
Europa	2.6	1565	16
Callisto	2.6	2410	16
Titan	18.6	2576	26
Ganymede	35.4	2631	46

2.9 Water as an Energy Storage Solution

A further system level advantage of a water electrolysis HET system is the potential synergy with the energy storage system of a spacecraft. Hydrogen-oxygen fuel cells have been used in space for power generation since the 1960s in project Gemini[76]. Fuel cells generate water and

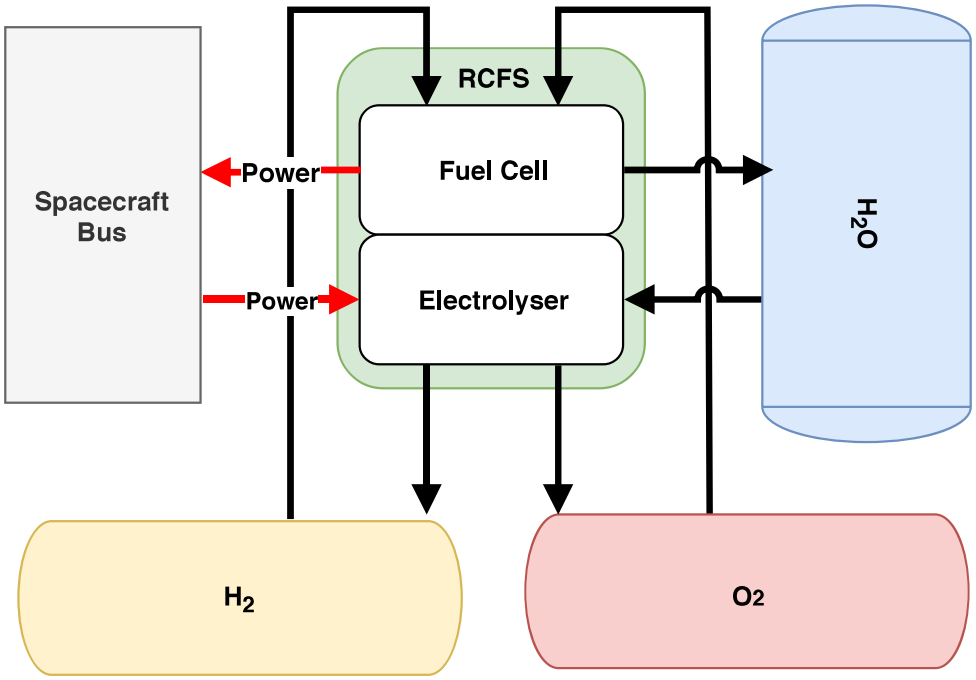

Figure 2.4: A Regenerative Fuel Cell System (RFCS) that incorporates an electrolyzer and a fuel cell stack in a closed system used for energy storage.

electricity from gaseous hydrogen and oxygen, effectively working as a reverse electrolyser. A Regenerative Fuel Cell System (RFCS) combines both an electrolyser and a fuel cell into a single closed system for energy storage[77]. The system converts hydrogen and oxygen into water when power is needed, and separates the same water back into its constituent elements when power is available, as shown in figure 2.4. Such a closed system has the potential to fill the energy storage role of a typical rechargeable battery.

As telecommunication satellites grow, so do their power demands. ESA foresees these demands rising to over 30 kW within the next 10 years[44]. The agency has investigated the use of RFCSs to power these spacecraft, so that they can continue their high-power operation while their solar arrays are eclipsed by the Earth. These eclipse periods put the greatest electrical stress on the spacecraft, as power needs to be stored to operate the payload 72 minutes at a time[77]. A complete RFCS, including electrolyser, fuel cell, reactants and storage tanks, could theoretically store energy at a density of up to 1000 Wh/kg, where modern Lithium ion batteries struggle to

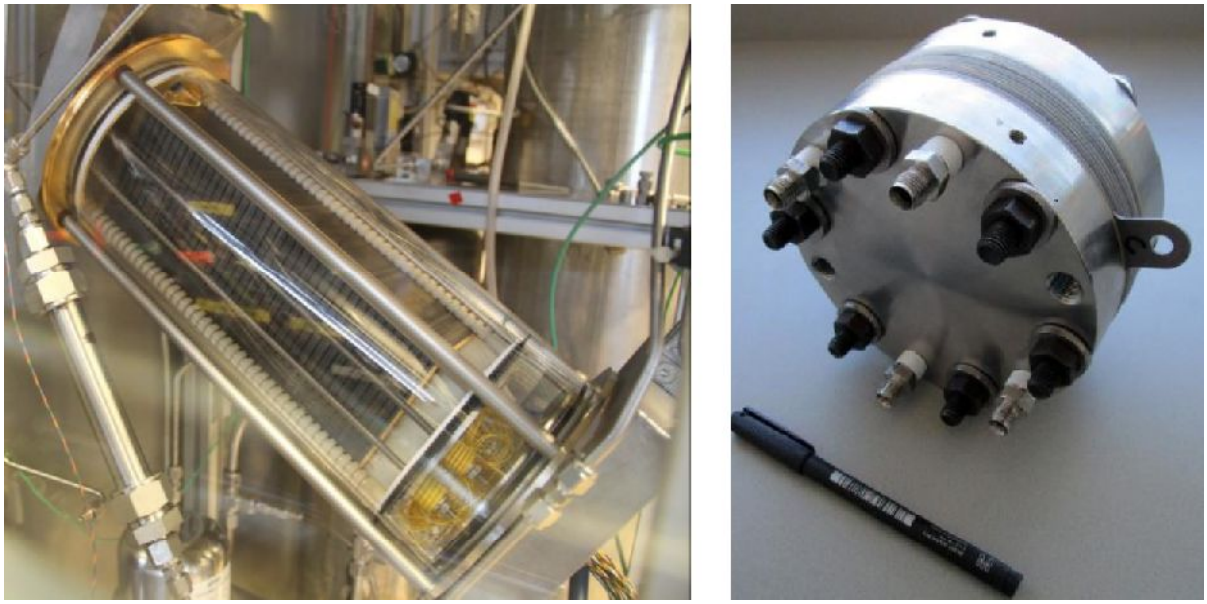

Figure 2.5: A high temperature fuel cell (left), and 1kW electrolyser (right) similar to those used for a closed-cycle RFCS. Reproduced from Farnes et al.[44] with permission of the authors.

achieve better than 250 Wh/kg[43]. Experimental tests show that RFCSs have the potential to last an entire 15 year lifetime in orbit with voltage degradation below 10%[44]. The two major components, the electrolyzer and the fuel cell, are shown in figure 2.5.

The dramatic mass saving made available by replacing lithium batteries with a RFCS suggest that we may see water and electrolysers flying in the near future. If this is the case, further mass savings could be achieved by also using the produced hydrogen and oxygen for propulsion. A single water storage tank and electrolyser could be used to store energy, and supply propellant to an electrolysis electric propulsion system. This would increase the size of the water tank required, but negate the need for a separate propellant tank. The power storage system and the propulsion system are two of the heaviest subsystems on a spacecraft, so combining the two would offer substantial mass savings.

Integrating the power and propulsion subsystem of a spacecraft using electrolysis was first proposed in the 1990s[78;79]. The concept of a Integrated Modular Propulsion and Regenerative Electro-Energy Storage System (IMPRESS) got so far that laboratory experiments were

conducted, showing that a single electrolyser could be used both for propulsion and energy storage[80].

2.10 A Unified Water System Architecture

The water electrolysis HET concept enables water to be used as both an electric and a chemical propellant. It has also been shown that water electrolysis can be used as an energy storage method in an RCFS. Combining both of these concepts we arrive at a highly integrated water system architecture, shown in figure 2.6. Here, water is converted to hydrogen and oxygen in space so that the gases can be used in an electrical propulsion system, a chemical propulsion system, and a fuel cell. The hydrogen gas could further be used for cold gas thrusters for attitude and orbit control, given this gas has been shown to have an I_{sp} of 272 s, compared to nitrogen which opperates near 70 s and is currently used[81]. Such a highly integrated system architecture has the potential to greatly reduce spacecraft mass and increase the competitiveness of the platform.

A unified water system architecture such as that shown in figure 2.6 requires the development of several key components. Both electrolysers and fuel cells have flight heritage in manned missions for which they are used for life support, indicating utmost confidence in their reliability[46;76]. The operation of these two units in a closed system in situations that are representative of a 15 year telecommunications mission (1350 charging cycles) is currently being assessed[43]. Chemical propulsion systems operating with a stoichiometric ratio of hydrogen and oxygen have been demonstrated in the laboratory, with some systems expected to fly very soon[65;45;66;78]. To our knowledge, the utilisation of water electrolysis as an electric propellant has never been attempted and therefore provides ample motivation for this thesis.

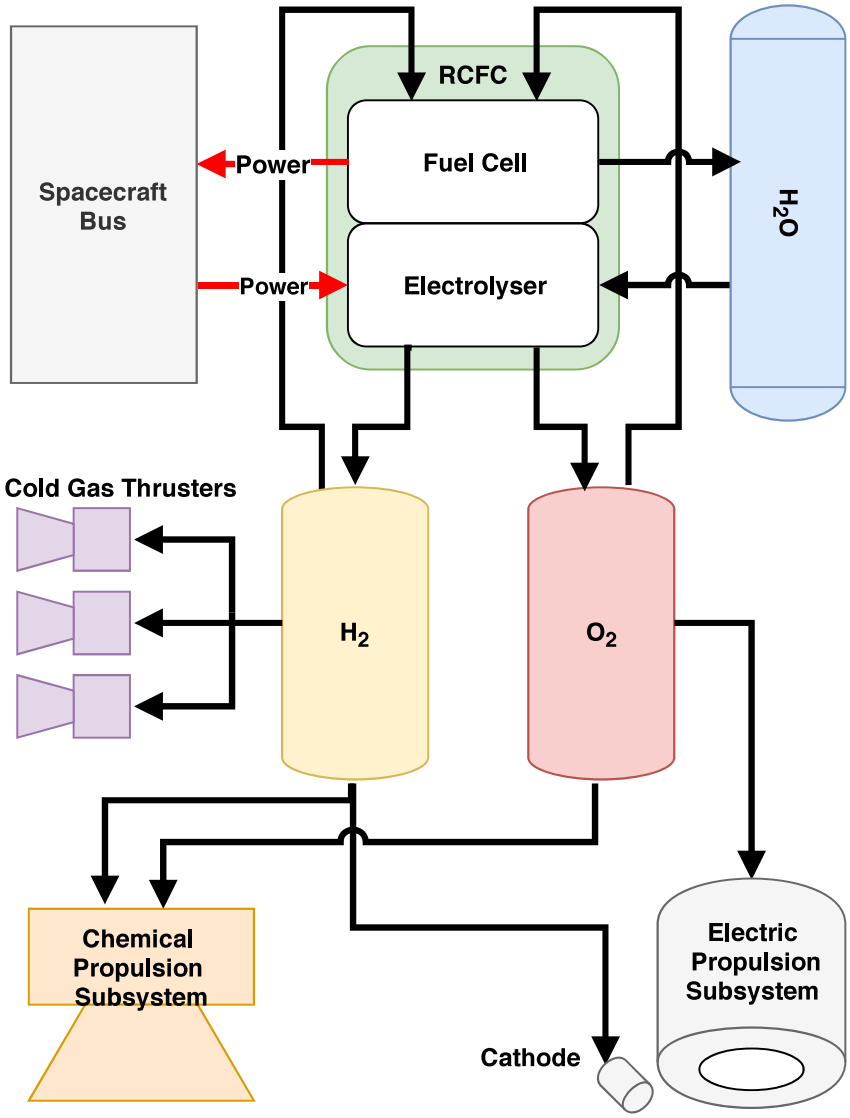

Figure 2.6: A fully integrated spacecraft water system architecture. Water is electrolysed so that it can be used for the electrical propulsion system, the chemical propulsion system, and a fuel cell.

2.11 Project Objectives

The electric propulsion community is actively searching for an alternative propellant to replace xenon. This thesis proposes the novel concept that a HET could be operated on the products of in-situ water electrolysis, by supplying the thruster with oxygen and the cathode with hydrogen. Such a system has several advantages over current or proposed alternative propellants, namely exploiting the low cost, low toxicity, low storage pressure and high storage density of water. Along with future ISRU potential, a water electrolysis HET offers unique synergies with both water electrolysis chemical propulsion systems and regenerative fuel cell energy storage systems. Although HETs have been operated on oxygen in the past, none have been designed specifically for this purpose. Given that oxygen is a low mass propellant with a small electron impact ionisation cross section and diatomic structure, even an optimised water electrolysis HET cannot be expected to attain the same performance as a xenon thruster. However, depending on the performance of an optimised oxygen HET, there may be certain mission scenarios in which the benefits, costs, and synergies at the system level make a water electrolysis HET a competitive choice.

Main Research Aims and Approach

This projects aims to determine the feasibility of a water electrolysis HET system. This will be done by answering three main research questions:

1. **What is the optimal design and operational point for an oxygen HET?** This is suggested to be significantly different to a xenon thruster, and was found using numerical simulations, and confirmed by experimental testing.

2. **What are the performance limits of a thruster optimised for oxygen?** Testing an

optimised oxygen HET in the laboratory will determine how thruster performance changes over a multi-dimensional parameter space, and identify any potential show stoppers.

3. **Do these results suggest a water electrolysis HET system is competitive?** A comprehensively characterised oxygen fuelled thruster will help to determine if a water electrolysis HET system can be considered competitive.

An oxygen optimised HET was expected to have an entirely different channel geometry to a traditional xenon thruster, meaning a new thruster must be constructed. Numerical plasma simulations were performed to determine how different thruster geometries and set points effect oxygen performance. Numerous iterations allowed the simulations to converge on a thruster geometry optimised for oxygen operation. The numerical plasma simulations and results are presented in chapter 3. A thruster was designed based on the numerically optimised geometry such that performance could be experimentally determined. A physical device had to be constructed on the basis of the optimised geometries found by plasma simulations. The resulting thruster was constructed after cyclically iterating magnetic models, thermal models, and CAD models. These models and the final design are described in chapter 4. Chapter 5 details this experimental set-up required to test the optimised oxygen thruster in the laboratory. Experiments were performed to experimentally determine how different independent variables affected thruster performance. The variables studied were: oxygen mass flow rate, input power, ceramic wall material, magnetic field strength, channel depth and magnetic topology. The results are presented and discussed in sections 6 and 7 respectively, with conclusions following in chapter 8.

Chapter 3

Numerical Modelling

This chapter describes the computer simulations performed in order to model the plasma within an oxygen HET. The ultimate goal of this was to replicate how thruster geometry, mass flow rate, and discharge voltage effects performance. Simulating this enabled a computational optimisation of these parameters to be undertaken, with the highest performance deciding the design of the physical experiment. First the plasma was simulated in a zero dimensional global plasma chemistry model called ZDPlasKin, so as to better understand the composition and time evolution of the plasma. Next a version of an in-house particle-in-cell code called PlasmaSim was modified to better suit this project. Once modified, this code was able to simulate the impact that changing the dimensions, electric field and mass flow has on performance of an oxygen HET. After validating the software, a sensitivity analysis was performed to determine the optimum thruster dimensions for an oxygen HET.

3.1 Zero-Dimensional Plasma Kinetics Solver

The covalent bond in the oxygen molecule presents additional degrees of freedom in which energy can be deposited, namely the vibrational excitation, rotational excitation and dissociation

of the molecule[18]. The equipartition theorem dictates that one must consider all degrees of freedom when determining which portion of the energy goes into the desired ionisation mechanism, compared to the energy which is wasted by going into excitation[18]. Each degree of freedom has a different probability of being excited. In this context these probabilities are expressed the same as for ionisation: as a collisional cross section which is dependent on the kinetic energy of the collision. A global plasma model is able to calculate the energy and outcome of each of these collisions, which takes into account all the chemistry involved with a range of different oxygen species such as molecules, atoms, ions and metastable excited states. By simulating an oxygen plasma in such a global plasma model, the oxygen plasma species composition and its time evolution can be determined.

The global model that was chosen is a zero-dimensional kinematic Boltzmann solver named ZDPlasKin[82]. This software was written by Pancheshnyi et al.[82] at the University of Toulouse, and while it is no longer maintained, an active forum online is still operational. Final release version of this code was used: version 2.0a for Windows. ZDPlasKin allowed the simulation of an oxygen plasma at neutral temperatures and electron temperatures found in a typical HET. The simulation is initialised by the user defining the starting number density, gas temperature and species of the oxygen to consider. The model is zero dimensional so no volume is defined, and no interactions with the channel walls are simulated. An electron population is introduced at a user defined number density and temperature. The oxygen temperature is a single constant value for the entire population, but the electrons take a temperature that is a Maxwellian distribution centred about the specified electron temperature.

The ZDPlasKin model requires a database of forward and back reactions of any interactions between particles that are included, such as between electrons, atomic oxygen, molecular oxygen and various ions. Each of these interactions must also have a reaction rate, which can be

a function of gas temperature, electron temperature, time and density. Electron bombardment is an example of one of these interactions. ZDPlasKin includes BOLSIG+, a built in electron Boltzmann equation solver. The result is the production rate of oxygen ions, molecules, and excited species. These production rates, along with reactions due to other mechanisms such as heavy-heavy particle collisions, or spontaneous de-excitation determine the number density of each species at the end of each timestep, which is saved and used as a starting point for the subsequent timestep.

3.1.1 Simulating an Oxygen Plasma with ZDPlasKin

Determining all possible reactions within an oxygen HET between different species, along with their reaction rate dependencies is a complex physical chemistry task. As an example of how to use ZDPlasKin, the software's website includes the input reaction data of Flitti and Pancheshnyi[83], most of which was taken from Capitelli et al.[84]. This study considered the complex chemistry of a O_2-N_2 plasma discharge. These reaction data were downloaded, and all reaction rates involving any nitrogen species were removed. What was left was a reaction data input file for ZDPlasKin which includes 143 reactions between electrons and the following oxygen species: $O_2(X^3$, v=0-4), $O_2(a^1$, b^1, 4.5eV), O_3, $O(^1D$, $^1S)$, with ions: O^+, O_2^+, O_4^+, O^-, O_2^-, O_3^- and O_4^-.

The reduced reaction data from Flitti and Pancheshnyi[83] were then used to initiate ZDPlasKin simulations such that the complex chemistry involved within an oxygen plasma could be captured. To be representative of a HET, it was important to determine the species densities and temperature within such a discharge. These values were not known for an oxygen HET, given mass flow rates and geometries had yet to be determined. Instead these were taken from those found within a Russian SPT-100 operating on xenon. This thruster was again chosen due to it representing the state-of-the art of HETs, and it having a very well documented performance

record in the laboratory[13]. A SPT-100 typically generates a xenon plasma with a number density of $10^{19}\,\mathrm{m}^{-3}$ neutrals with a temperature of 1000 K, and electrons in a Boltzmann distribution centred about 10 eV[85]. The same electron temperature was also used for the ZDPlasKin simulations, again due to no other data being available. Unfortunately ZDPlasKin crashed repeatedly when trying to simulate oxygen at these conditions. The initial neutral densities were reduced to 10^{14} m^{-3}, which alleviated the issue, at the expense of being less representative of a real HET. Nevertheless, this condition was comparable to other similar code as a validation case discussed below.

The time evolution of a ZDPlasKin simulation of a HET operating on oxygen is shown in figure 3.1. Note that time evolution scales inversely with density, meaning the unrealistically low plasma density used in the simulation creates a time scale which is far slower than would be realistic for a HET. Here we refer to the ground state as the state in which the oxygen molecule is not ionised, electrically excited, vibrationally excited or rotationally excited. The particles labelled as *electric* are uncharged oxygen atoms (O electric) and molecules (O_2) which have been electrically excited. Similar vibrationally and rotationally excited molecules (pink in figure 3.1) are uncharged oxygen molecules with a raised internal energy.

The population starts with all particles in the ground O_2 state at a density of 10^{14} m^{-3}, and at a temperature of 1000 K, represented by the dark blue population. Here ground means that the particle have no charge, and no vibrationally, rotational or electrical energy. To begin the ionisation a seed electron population is introduced of the same number density with a Boltzmann distribution centred at 10 eV. As the simulation evolves it can be seen that the population of O_2 in the ground state quickly depletes, with the majority of particles ending as electrically excited O atoms shown here in light red, and 32% of the population being ionised as seen in the combined light and dark yellow). These time scales are expected to reduce greatly for densities

representative of a HET.

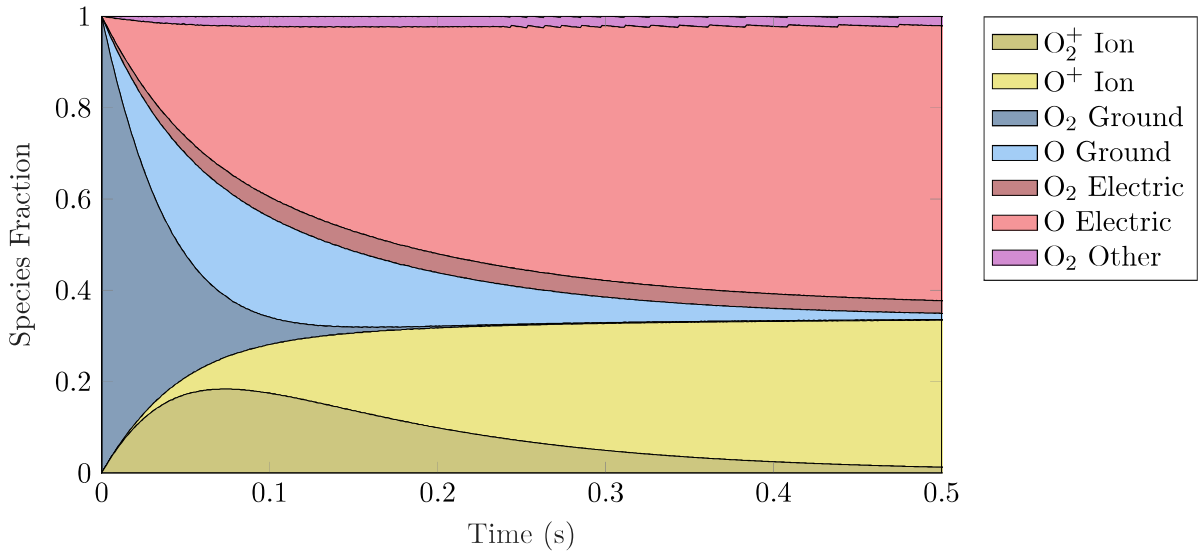

Figure 3.1: Time evolution of an oxygen plasma species fraction as simulated by ZDPlasKin. The yellow sections are the ion populations, blue is oxygen in its monatomic and diatomic ground states and red is the electrically excited particles. The magenta section is the sum of all other oxygen species, including vibrationally and rotationally excited molecules, and all negative ions.

Note that all the vibrationally and rotationally excited O_2 molecules make up only a fraction of the magenta band in figure 3.1, and that this band reaches a steady state very quickly into the simulation. The two yellow bands in figure 3.1 shows that the sum of the O^+ ions and O_2^+ ions is relatively constant from approximately 0.2 s onward, however this population begins as predominantly O_2^+ ions, with the final population consisting of mostly O^+ ions. This is an important feature to consider when optimising a HET for oxygen. The depth of the thruster channel has a direct impact on the residence time of the heavy particles, simply because a particle will take longer to travel through a deeper channel[2]. These results show that a channel depth which produces a residence time over a particular value offers little increased ionisation. Unfortunately the simulation cannot accurately determine which depth this is, due to unrealistically low densities producing unrealistically long timescales. Furthermore, increasing the channel depth, and thus residence time will affect the ratio of O_2^+ to O^+ ions, as seen as the ratio of dark to light yellow in figure 3.1. Given that the same work is done by the power supply to accelerate either ion

over a fixed distance, the lighter O^+ ion is expected to produce a higher I_{sp}, at the expense of thrust[2]. Considering that either particle produces a very high I_{sp} in comparison to xenon already, there is an incentive to favour increasing the presumably low thrust over the already high I_{sp}. It can be concluded that a shallower channel would be advantageous as this would induce a shorter residence time, a higher fraction of O_2^+ ions to O^+ions, and consequentially a greater thrust. This holds for any channel with a residence time above 0.2 s, as below this value, the combined ionisation fraction drops.

3.1.2 The Limits of ZDPlasKin

ZDPlaskin has shown us that an oxygen plasma in conditions similar to those found in a HET will produce an ionisation fraction of approximately 32%, with majority of ions being molecular, but decreasing steadily until they are almost all atomic ions. This is important information when designing the ideal channel depth for an oxygen HET.

Further numerical simulations did not involve ZDPlasKin. The software is zero-dimensional, and thus geometry dependent factors such as plasma-wall interactions have been ignored. These interactions have however been identified as a major source of de-excitation and electron-ion recombination for HETs in the literature[49;86]. Not only does ZDPlasKin not include wall interactions, but all particles are conserved, which is not representative of a device specifically designed to eject ions. The undesirably low maximum number density which was simulated by ZDPlasKin is another limitation of the software. In order to simulate an oxygen plasma in a way that meets the very specific demands of an oxygen HET, a more representative model was required which could be fully customised.

3.2 Particle-in-Cell Simulations

Particle-in-cell (PIC) simulations have been used since the late 1950s to model the complex interactions which couple the behaviour of individual charged particles of a plasma and the macroscopic attributes[87]. The PIC name arises as the motion of ions, electrons and neutrals are modelled kinetically as particles moving through the domain, yet the local plasma attributes such as charge density and electric potential are calculated at nodes on a fixed mesh of cells. After initialisation a PIC simulation is progressed by stepping forward in time at a predefined time step. The components of a single typical PIC step are shown below[88]:

1. **Calculate local plasma parameters for the current state.** The bulk plasma properties such as particle density and charge density are calculated for each cell, by summing up the particles within that cell. The distribution of these cells describes the dimensionality of the PIC code. Codes exist with cells in one, two or three dimensions, with each added dimension dramatically increasing computational time.

2. **Determine electric and magnetic fields at mesh points.** This is done by finding a solution to Maxwell's equations for the charge distribution calculated in the above step. The numerical methods used for solving these partial differential equations can broadly be divided into three groups: mesh-relaxation methods in which an initial guess is adjusted iteratively, matrix methods which reduce the equations to to a set of linear equations to solve with matrix methods, or spectral methods which which solve the system in Fourier space[88].

3. **Calculate the forces on each particle by solving the equations of motion.** The equation of motion for the charged particles is the Lorentz force:

$$F = q\mathbf{E} + q\mathbf{v} \times \mathbf{B}, \tag{3.1}$$

where q is the charge of the particle, $\mathbf{E}$ and $\mathbf{B}$ are the electric and magnetic fields respectively and $\mathbf{v}$ is the velocity of the particle. Information such as the charge and velocity of each particle is known, but the electric and magnetic fields must be interpolated from the surrounding mesh points to the location of the particle.

4. **Integrate the equation of motion and move each particle.** This section of the code is the most time consuming and is referred to as the particle pusher or particle mover. Different particle movers have trade-offs between speed and accuracy. Different particle movers can be divided into explicit and implicit schemes[88]. Explicit integrators like the leap-frog method use the forces calculated in the previous time step, and are chosen for their simplicity and speed, yet at the expense of requiring a smaller time step. Implicit integrators such at the Boris pusher method use the force calculated in the present time step. The Boris pusher method has become the most popular integration scheme due to its impressive long term accuracy[89].

5. **Particle-particle interactions and boundary conditions.** This step refers to all other interactions in the plasma and changes depending on the requirements of the simulation at hand. Typical mechanisms for the simulation of electric propulsion include:

    - Electron-neutral collisions which may or may not result in ionisation or excitation.
    - Ion-wall collisions which result in recombination.
    - Ions being ejected from the domain and being removed from the simulation.
    - Neutrals being added from the simulation as mass flow.
    - Electrons being added from the cathode or being removed at the anode.

The execution of all of the above steps represents a single time step. The size of this time step must be chosen with great care: larger time steps decrease the overall computation time, yet the size of the time step also dictates the maximum time scale that can be resolved. Many

PIC codes automatically change time step sizes to ensure they can capture phenomena at high frequencies such as electron cyclotron motion, and the plasma frequency. The plasma frequency refers to the oscillation of electrons about the ions[90].

Given that typical HET plasma number densities are on the order of $10^{20}\mathrm{m}^{-3}$, it is not computationally feasible to track the movement of each individual particle. Instead identical particles are grouped together to create "super particles", which represent a cluster of either electrons, ions or neutrals which share the same mass, charge, velocity and location. Equation 3.1 shows that grouping real particles to super particles does not affect their motion, because the Lorentz force is only dependent on the charge to mass ratio of the particle which is unchanged by grouping. To balance computing time with accuracy, super particle size is another user defined metric which must be selected based on a trade-off between computation speed and accuracy.

3.3 PlasmaSim

To better understand how a HET operates on an oxygen plasma a PIC code called PlasmaSim was utilised. PlasmaSim is a PIC plasma simulation software, currently being developed in the plasma propulsion group at the Aeronautics department of Imperial College London. Dr Knoll began developing the code in 2017 using a combination of the Unreal Engine 4 and C++ to simulate electric propulsion devices operating on xenon. To speed up computation time PlasmaSim is one dimensional, meaning that although the position and velocity of all particles are simulated in three dimensions, cells are only distributed in one direction, which is generally chosen to be the thruster axial direction. The results is that bulk plasma properties such density, and plasma potential are calculated at points along the thruster axis, and assumed to be constant in both the radial and the azimuthal direction.

Initiating PlasmaSim prompts the user to define the size of the domain in the x, y and z

direction. The domain is then populated with a uniform distribution of electrons, and xenon particles each with a user defined temperature and number density. Other simulation inputs are the total simulation time, the time step duration and the total number of super particles to be simulated.

For HET simulations the magnetic fields generated by the motion of particles are negligible in comparison to the externally applied magnetic fields. PlasmaSim solves the one dimensional Poisson equation to find the electric field. This dimension is generally the x axis, which allows us to find the electric field and electric potential from:

$$\frac{\partial^2 V}{\partial x^2} = \frac{\partial \mathbf{E}}{\partial x} = -\frac{\rho(x)}{\epsilon_0}, \tag{3.2}$$

where ϵ_0 is the permittivity of free space. Here V is the electric potential, and $\mathbf{E}$ is the electric field, which can be found from the known charge distribution $\rho(x)$. This charge distribution is calculated based on the particle locations within the array of cells. Without solving the magnetic field, PlasmaSim is termed an electrostatic PIC code, unlike electromagnetic codes which solve the whole set of Maxwell's equations. A matrix method approach was selected, where the Poisson equation is reduced to a set of linear equations and solved using the LAPACK linear algebra software package[91]. For the particle pusher, PlasmaSim employs the Boris pusher scheme for long term stability[89]. The pusher integrates motion of each particle subjected to the Lorentz force.

As in a real HET, the ionisation process in PlasmaSim is dependent on the electron-neutral impact which isparticular electron will ionise one of the neut simulated using a Direct Simulation Monte Carlo (DSMC) method: the kinetic energy of each electron is used to determine the electron-impact ionisation cross section σ that it would have with each xenon atom in the same

cell. The cross sections are found by linearly interpolating from those reported in Stephan and Märk[53]. From this cross section the ionisation frequency ν_{ion} for an electron of this particular energy can be calculated:

$$\nu_{ion} = N_{neut}\sigma\sqrt{v_e^2 + c_n^2}, \tag{3.3}$$

where N_{neut} is the neutral density within this cell, v_e the velocity of the electron, and c_n the average neutral speed within this cell. By multiplying this collision frequency by the time step Δt, the probability P_{ion} that this particular electron will ionise one of the neutral particles within the same cell, within this time step can be found:

$$P_{ion} = 1 - exp(-\Delta t \nu_{ion}). \tag{3.4}$$

Finally the stochastic element of the DSMC code is executed: a random number between 0 and 1 is generated and the ionisation occurs if and only if this random number is smaller than P_{ion}. If an ionisation does occur, this is implemented by replacing one of the neutral particles in this cell with an ion, and generating a new electron at the same location. The electron is given a temperature randomly sampled from a Maxwellian distribution about the user defined electron temperature.

A very similar DSMC approach is taken to simulate elastic collisions of neutrals particles with other neutrals, ions and electrons. Here the cross section is the sum of the Van der Waals radii of the two colliding particles (assumed to be zero for an electron). The Van Der Waals radius of an atom or molecule represents the radius that the particle were to have, if modelled as a perfect hard sphere which represents the distance of closes approach with another particle. Natural elastic collisions are simulated if two particles get close enough that their Van Der Waals radii overlap.

Electrons and neutrals are given reflective boundary conditions, meaning that when they strike the walls they are returned with the same kinetic energy. Although electron-ion recombination in the plasma bulk does occur, the dominating mechanism of electron-ion recombination within a HET is driven by ions absorbing an electron as they strike the wall of the thruster[49]. This is currently the only method by which recombination is modelled within PlasmaSim: as an ion strikes the wall, it is reflected with its charge removed. When this occurs an electron is also removed from within the same cell in order to conserve global charge neutrality.

3.4 Development of PlasmaSim 0D

PlasmaSim was chosen to simulate a HET operating on oxygen more accurately than a global model like ZDPlasKin would do. The PIC code was used to optimise a thruster for the following parameters:

- channel depth
- channel width
- channel circumference
- anode mass flow rate
- discharge voltage

The version of PlasmaSim described above is one dimensional, meaning that the domain is divided into a one dimensional row of adjacent cells, with local variable describing the state of each individual cell. Typically this dimension is chosen to be the axial direction of the HET, such that the axially dependent plasma parameters are captured, but it is assumed that these parameters to not change with radial or azimuthal position. The 1D code requires close to 1000 hours of computation time to produce a 1 ms of simulation time on a 4.2GHz quad-core desktop machine. We consider this the minimum simulation time to represent steady state operation.

Each instance of PlasmaSim requires a dedicated CPU core over this time period. Waiting thousands of hours for a result is unfeasible when considering the hundreds of simulations required to optimise for the five independent thruster parameters. It was decided to sacrifice some of the complexity, and thereby the accuracy of PlasmaSim for a dramatic increase in computation speed. By making a number of assumptions and simplifications, a faster, low fidelity version of PlasmaSim was produced. A delicate balance had to be met such that the simplified code would still be representative enough to accurately perform the thruster optimisation, but also computationally in-expensive to provide significant time savings over the original code.

3.4.1 Reduction in Dimensionality

The most drastic PlasmaSim modification was reducing the simulation from a one dimensional code to a quasi-zero dimensional code. This modified version of the code will therefore henceforth be referred to as PlasmaSim 0D, with the original being called PlasmaSim 1D. The 0D code is quasi-zero dimensional in the sense that the position and velocity of all particles are still modelled in three dimensions, yet the entire mesh is reduced from a one dimensional chain of cells, to only one single cell. This assumes the all plasma bulk properties are independent of location, something known to be far from the truth: ion density, electron temperature and neutral density are all known to vary strongly with axial location[92].

3.4.2 Coordinate System of the Simulation Domain

Much like the original 1D code, the domain of PlasmaSim 0D is a rectangular representation of the annular discharge channel of a HET. By slicing and "unwrapping"the annular discharge channel, a rectangular domain is effectively created, shown in figure 3.2. The x axis is chosen to align with the axis of the thruster, with $x = 0$ being the anode, and the $x = x_{max}$ representing

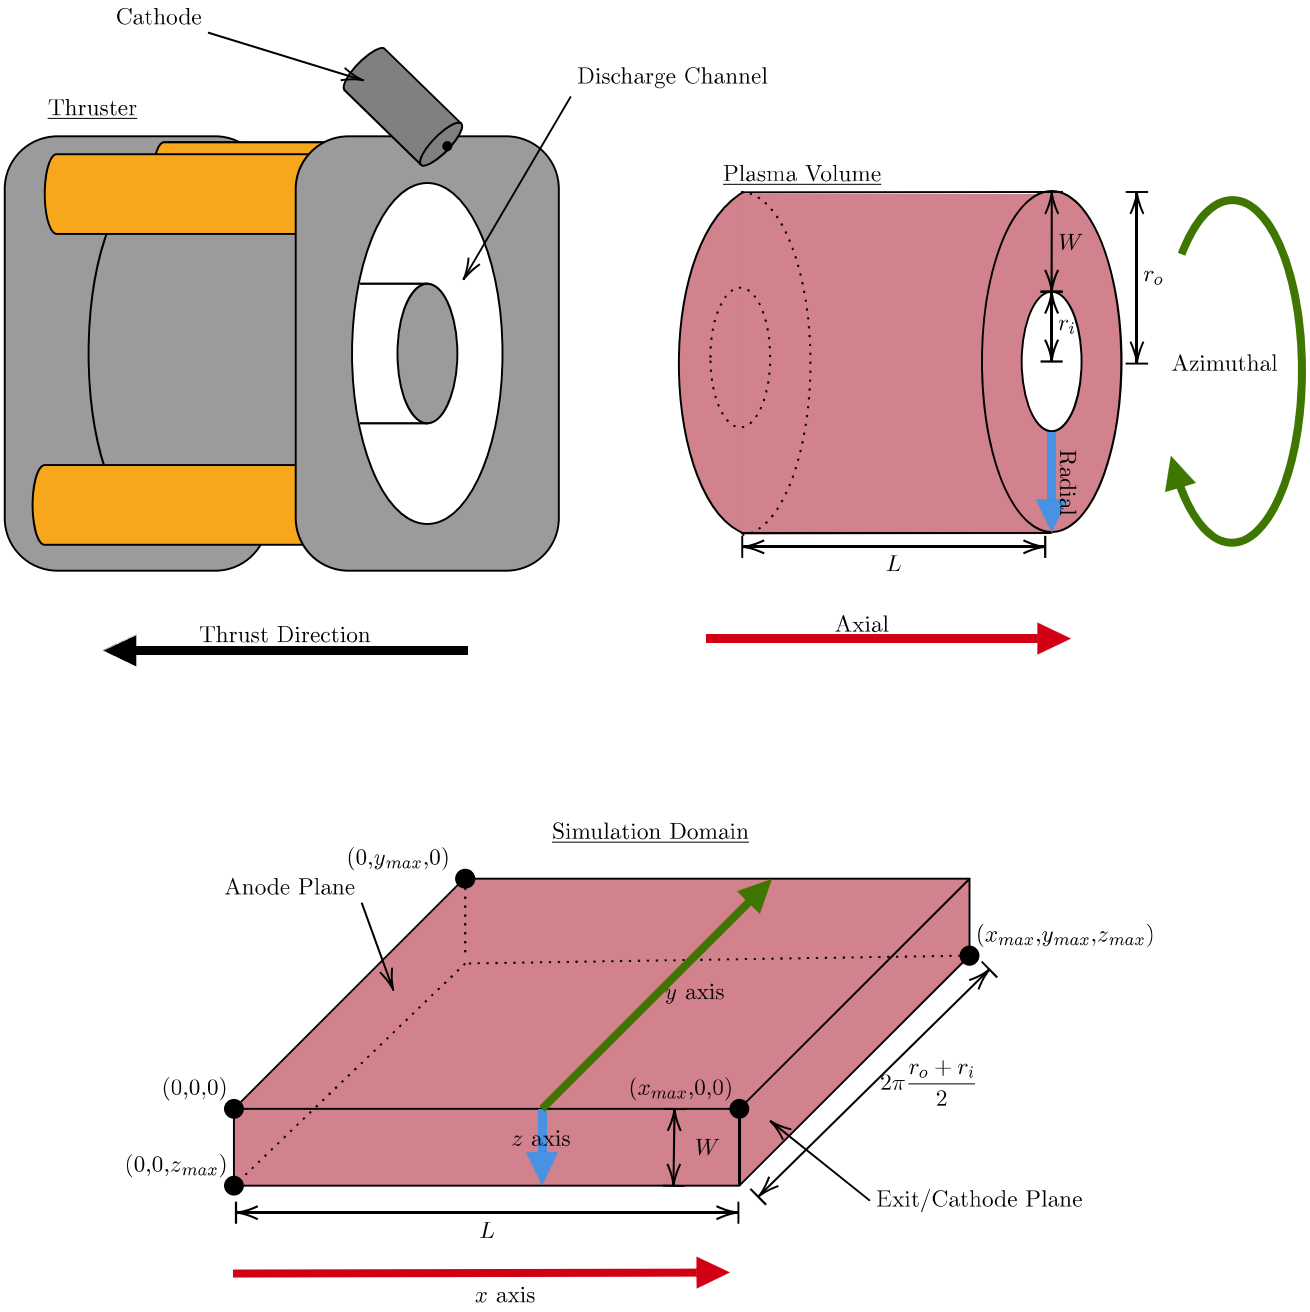

Figure 3.2: The new cartesian domain which is "unwrapped" from an annular HET channel. The channel has depth L, inner and outer radius of length r_i and r_o respectively, and channel width W. The new domain is a cuboid defined by the volume encompassed between the origin (0, 0, 0) and the user specified point at (x_{max}, y_{max}, z_{max}). Where $x = 0$ is the anode plane, $x = x_{max}$ is the exit/cathode plane, both $y = 0$ and $y = y_{max}$ are the top of the thruster, with $y = y_{max}/2$ being the bottom and $z = 0$ and $z = z_{max}$ being the inner and outer channel wall surface respectively.

the exit plane. The unwrapping of the domain transforms the azimuthal direction to the y axis. The $y=0$ plane was chosen to be the top of the thruster, increasing to the right when seen from the front. By treating the y boundary conditions as periodic, all particles can travel indefinitely in the y axis just as a real particle can travel azimuthally indefinably. The z axis now represents the radial direction, with $z = 0$ being the inner channel wall, and $z = z_{max}$ the outer channel wall. An important caveat of this transformation is that the inner and outer channel walls are now incorrectly assumed to have the same area. Note however, that by assuming both boundaries are the length of the average radius, the total area is unchanged. For any HET with channel depth L, channel width W, and outer and inner channel radius of r_o and r_i, PlasmaSim 0D can simulate the plasma within the discharge channel by assigning:

$$x_{max} = L \tag{3.5}$$

$$y_{max} = 2\pi\frac{r_o + r_i}{2} \tag{3.6}$$

$$z_{max} = W \tag{3.7}$$

3.4.3 Electric Potential

With only a single cell, and therefore a single charge density describing the entire thruster domain, solving the Poisson equation is no longer an option for deriving the electric potential. Instead a perfectly linear electric field is assumed from the user defined anode potential at $x = 0$, to the grounded cathode potential at the exit plane of the thruster at $x = x_{max}$.

3.4.4 Magnetic Field

In a HET, the presence of a strong externally applied radial magnetic field near the exit of channel is critical for operation. In conjunction with the electric field the two cause the electrons to drift azimuthally, with the resulting E cross B Hall effect giving the thruster its name. This azimuthal deflection of the electrons dramatically increases the electron temperature, plasma

density, and electric potential[49]. These localised features were not able to be captured in the zero dimensional model which considered the entire channel to be homogeneous. It is for this reason that PlasmaSim 0D disregards the magnetic field altogether, considering the entire simulation to be purely electrostatic. This ultimately means that the electrons will no longer undergo cyclotron motion, and no azimuthal Hall current will be established.

3.4.5 Electron Temperature

The ionisation rate within a HET is very sensitive to electron temperature, and is therefore strongly dependent on axial location. The lack of a Hall current in PlasmaSim 0D means that the spatial electron temperature distribution cannot be expected to be realistic. To combat this a user specified temperature was forced upon the electrons. This was done by generating a population of electrons distributed about a specified electron temperature with a Maxwellian distribution. This is done for both the initial seed electrons, as well as any electrons produced by ionisation. This dictates that for PlasmaSim 0D electron temperature is a very impactful tuning parameter. Given such a single value is not representative of reality, it cannot be taken from the literature and is best done by trialling different values and comparing the results to experimental data.

3.4.6 Adding Particles to Simulate Mass Flow

The setup of PlasmaSim 0D requires a continuous mass flow rate at the anode which the original design did not have. The program calculates how many super particles worth of neutrals to add at every time step for a user-defined mass flow rate. These particles are then added randomly throughout the anode region, which stretches between the deepest 5% and 10% of the channel depth, and the central 80% of the channel width. The neutrals are given the same temperature as the original neutrals, which is again a Maxwellian distribution about a defined neutral

temperature.

3.4.7 Boundary Conditions

Ions, electrons and neutrals are all periodic in the y direction as this represents the particles travelling azimuthally around the channel of the thruster. All particles in PlasmaSim 0D have a reflective boundary with the channel wall. This is an important distinction from PlasmaSim 1D, where ions striking the wall were used to simulate recombination (see following subsection). Ions and neutrals passing the $x = 0$ boundary represent particles striking the anode, and are reflected back into the simulation. A heavy particle (ion or neutral) passing the $x = x_{max}$ boundary represents mass being ejected from the thruster. The momentum of all heavy particles lost in each time step is summed to calculate macroscopic thrust and I_{sp} produced by the thruster. Similarly the number of ions leaving the channel is used to calculate the ion beam current at this time. To conserve quasi-neutrality of the simulation, a random electron is removed from the simulation for every ion ejected.

In a real HET the electrons are ejected from the negative cathode and follow the potential gradient caused by the discharge voltage through the plasma to the positive anode, where they are absorbed and pumped back to the cathode by the power supply. This is simulated in PlasmaSim 0D by setting the $x = 0$ and $x = x_{max}$ periodic for electrons, such that an electron striking the anode plane is re-injected at the thruster exit plane. These electron boundary conditions inadvertently result in all electrons continuously being accelerated by the electric potential and gaining greater and greater temperatures. To counteract this, the Boris pusher was modified to neglect the action of the electric field on the electrons. The resulting electrons move at a constant velocity in time, which is not physical, yet guarantees that their temperature remains constant.

3.4.8 Recombination and Bohm Current

The plasma within a HET comes into contact with the ceramic channel walls and the metallic anode. Ions striking the surface generally deposit their energy and absorb an electron before returning to the plasma. This is the primary mechanism of recombination in a HET. This makes surfaces interactions extremely important for PIC simulations, as this recombination counteracts the ionisation process occurring in the bulk of the plasma. This rivalry between bulk driven and wall driven mechanisms underpins the importance of the surface to volume ratio of the thruster.

In a thermal steady state plasma the ions and electrons have the same temperature and therefore kinetic energy. In such a case the mass discrepancy between the particles dictate that the electrons have a greater velocity than the ions by a factor of $\sqrt{m_i/m_e}$. In a typical HET the electron temperature is in the order of 10-30 eV, where ions may have a temperature closer to 0.1 eV[93]. This temperature difference increases the disparity in the velocities even further[7;90]. On the surfaces where the plasma contacts the walls, the faster electrons exit the plasma and strike the walls at a much higher velocity and thus frequency than the positive ions. This uneven particle flux charges the wall to a more negative potential than the plasma bulk[90]. This causes the positive ions in the plasma to be attracted to the wall, shielding the wall from the rest of the bulk: an example of Debye shielding. As the electrons flow to the wall the negative charge grows and begins to repel further incoming electrons, establishing a small electric field between the wall and the plasma bulk. This transition layer between the plasma bulk and the negative wall is called the plasma sheath, or Debye sheath[94]. Eventually this surface charge reaches a stable equilibrium were the net electron flux and ion flux to the wall is equal. The depth of the sheath adjusts itself, so that the wall is at a lower potential than the plasma, impeding all but the fastest electrons, so that the electron flux is equal to the ion flux[90].

Ions are accelerated towards the wall by the decreasing potential of the sheath region. Given

ion flux is constant at steady state, the acceleration of these ions causes an ion density which decreases as it approaches the wall[90]. An initial ion velocity at the beginning of the sheath must exist to ensure that the ions are leaving the plasma with sufficient energy to ensure that the ion density does not exceed the electron density. This velocity is called the Bohm velocity (v_b), and is given by:

$$v_b = \sqrt{\frac{kT_e}{m_i}}, \tag{3.8}$$

where k is the Boltzmann constant, T_e is the electron temperature, and m_i the mass of the ions[90;7]. An equilibrium is formed where the pre-sheath plasma accelerates the ions to at least the Bohm velocity to stabilise the sheath[90;7]. This flow of ions to the walls is called the Bohm current.

The plasma homogeneity that PlasmaSim 0D assumes cannot consider localised phenomena like Debye shielding or Bohm flux. However, Scharfe et al.[95] show that for a HET, the flux of ions per unit volume per second γ_i can be calculated by:

$$\gamma_i = \frac{2}{W} exp(-1/2) n_i v_b, \tag{3.9}$$

where W is the thruster channel width, and n_i the ion density. The recombination rate in PlasmaSim 0D is calculated by multiplying equation 3.9 by the thruster volume and time step, to calculate the number of ions that must recombine in each time step. The appropriate number of ions in the plasma bulk are then randomly selected, their charge removed, and their temperature set to that of a neutral. The same number of electrons are removed to ensure quasi-neutrality is upheld.

3.4.9 Changing propellant from xenon to oxygen

Adapting the simulated propellant from xenon to oxygen in PlasmaSim 0D was trivial. Changing the propellant in the set-up menu changes the mass and the Van der Waals radius of the heavy

particles, and the electron-impact ionisation cross section. The cross sections of the oxygen molecules are taken from Itikawa [52]. It is very important to note that PlasmaSim 0D does not simulate the electrical excitation, nor the vibration, rotation or dissociation of the oxygen molecules. Simulating the internal energy states of the molecules in such a way would require dramatically altering the core of how PlasmaSim works: changing the internal energy state of the molecule changes the electron impact cross section, and therefor the likely of the particle transitioning to all different excitation and ionisation states. Doing so was considered to be outside of the scope of this project. The impact that this will have on the efficiency of the device is unknown at this point, and will need to be determined experimentally.

3.5 Validating PlasmaSim 0D

The ability of PlasmaSim to simulate a HET plasma was validated by simulating the SPT-100 operating on xenon. The SPT-100 has a channel depth of 42 mm, channel width of 16 mm and an outer channel diameter of 100 mm [85]. These dimensions differ slightly from more recent European studies of the SPT-100 which have a shallower channel, as discussed in Tejeda et al. [96]. To simulate this in PlasmaSim 0D the domain was sized to this discharge chamber, as described by equations 3.6-3.7: $x_{max} = 42$ mm, $y_{max} = 263.9$ mm, and $z_{max} = 16$ mm. A mass flow rate of 5.12 mgs^{-1} xenon was chosen, to align with a particular experimental set-point tested by Sankovic et al. [12]. The literature reports that SPT-100 operates at a neutral particle density near $10^{19}\,\mathrm{m}^{-3}$ so this was taken as a starting neutral density, with a seed electron and ion density of $10^{17}\,\mathrm{m}^{-3}$ [85]. The ions are given an initial temperature centred around 0.1 eV, which is the value typically used in such simulations [93]. Electron temperature for the SPT-100 are reported to be in the ranges from 10-20 eV [85], and the discharge voltage for this set point is 300 V [12]. Given that PlasmaSim 0D demands a single homogeneous electron temperature and electric field strength, neither of which are considered realistic, these values were taken from the literature but then used as tuning parameters to align the PlasmaSim results to those found

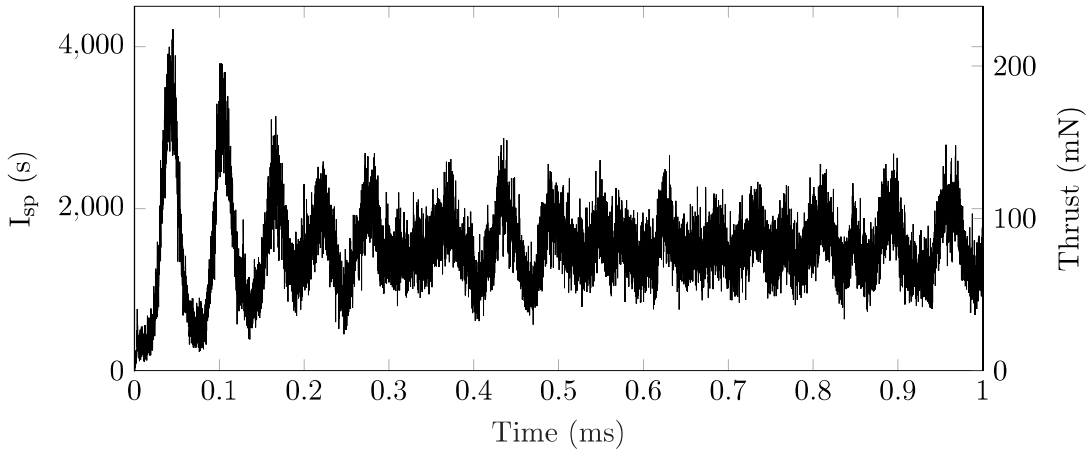

Figure 3.3: Thrust and I_{sp} of a SPT-100 as simulated by PlasmaSim 0D

experimentally. An electron temperature of 11.5 eV, and electric field strength of 7550 Vm^{-1} (which equates to a discharge voltage of 317 V over the 42 mm deep channel) were found to give very good agreement. The simulation was run for 1 ms of simulation time, which provided ample time to reach steady state operation. A time-step of 1^{-7} s and a super particle number of 34000 was used. The choices of simulation time, time-step size and super particle number were all based on results of a sensitivity presented in section A.1 of the appendix.

At each time-step PlasmaSim notes the axial component of velocity of each neutral and ion leaving the thruster, to determine the momentum of these particles. The sum of this momentum is the reaction force that the thruster experiences over a time step: the thrust. We divide this momentary thrust by the mass flow rate, and scale by the acceleration due to gravity as described by equation 1.2 to determine the specific impulse of the thruster for that time-step.

The I_{sp} and thrust of the SPT-100 simulation are shown in figure 3.3. The simulation shows an average thrust of 81.6 mN, and an average I_{sp} of 1535 s when in steady state, which is very representative of those seen in experimental results of 86.6 mN and 1550 s respectively[12]. We take this as a testament the ability of the simplified PlasmaSim 0D software to accurately pre-

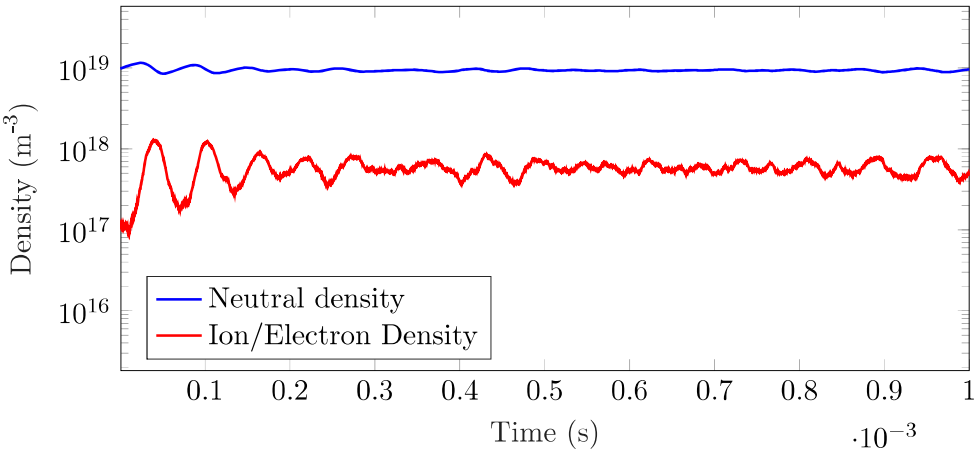

Figure 3.4: Species number densities of a SPT-100 as simulated by PlasmaSim 0D. Note that ion and electron densities are forced to be identical to ensure charge quasi-neutrality.

dict experimental thrust and I_{sp} results. The high amplitude oscillations of the thrust and I_{sp} in figure 3.3 represent a natural instability within HETs, and are better understood by looking at the number densities of both electrons and neutrals on a logarithmic scale, as in figure 3.4. Both of these densities are shown to fluctuate, at the same frequency, yet out of phase. These oscillations represent an experimentally observable phenomenon termed breathing mode oscillations[7]. Breathing mode oscillations describe a predator-prey relationship between the electron density and the neutral density: an increasing neutral density causes an increase in ionisation, as simply more neutrals are present to be ionised. Given the plasma is quasi-neutral, an increased ionisation rate increases both ion and electron densities. The increased electron density increases electron-neutral collisions, causing a subsequent drop in the neutral population, followed by a decrease in electron density. The mechanism behind breathing mode oscillations very much appears to fit with what is observed in figure 3.4 given the two populations oscillate at a similar frequency but out of phase. Diamant et al.[14] showed experimentally that for a SPT-100, breathing mode oscillations are dependent on the background pressures of the vacuum facility, ranging from 16.5 kHz at low facility pressures, up to 26 kHz at higher ones. The PlasmaSim 0D simulations effectively represent a perfect vacuum as the background pressure, so breathing mode oscillations are expected to fall near or below 16.5 kHz. Taking a Fast Fourier Transform of

the electron density data from figure 3.4, presents a very strong peak at a frequency of 15.0 kHz. Evidence of breathing mode oscillations at the expected frequency, along with accurate thrust and I_{sp} values clearly shows that PlasmaSim is able to correctly simulate the basic operation of a HET.

3.5.1 Sensitivity Study for the SPT-100

PlasmaSim 0D has shown to be effective at replicating the basic performance metrics of a xenon HET. The main purpose of developing this code was always to optimise channel geometries, mass flow rates and powers for a HET operating on oxygen. To validate the effectiveness of PlasmaSim 0D to perform such an optimisation, a sensitivity analysis was performed on the highly optimised SPT-100 HET operating on xenon. The literature reports that the SPT-100 has been optimised to operate on 5 mgs^{-1} xenon, a channel depth of 42 mm, channel width of 16 mm and outer channel diameter of 100 mm[85;12]. PlasmaSim 0D simulations were run for this configuration, and also for configurations where a single parameter was changed, to determine the configuration of the highest performance. If PlasmaSim 0D would be able to independently corroborate the optimal configuration with that in the literature, this would exemplify the software's effectiveness as a thruster optimisation tool.

For these experiments PlasmaSim 0D was operated at the default SPT-100 dimensions and mass flow rate, with one parameter being changed in each case. The same electron and ion temperatures of 11.5 eV and 0.1 eV respectively were used as above. The simulations were run for 10^{-3} s, with a time-step of 10^{-7} s and 34000 super-particles. The electric field strength was adjusted iteratively over many simulations to ensure the discharge power of the thruster was 1060 W for all configurations. This power was chosen as the power which PlasmaSim 0D produces for the optimal geometry. This underestimates the power from the actual average power of 1350 W, presumably due to a lack of Hall current decreasing the plasma resistance.

The results are shown in figure 3.5.

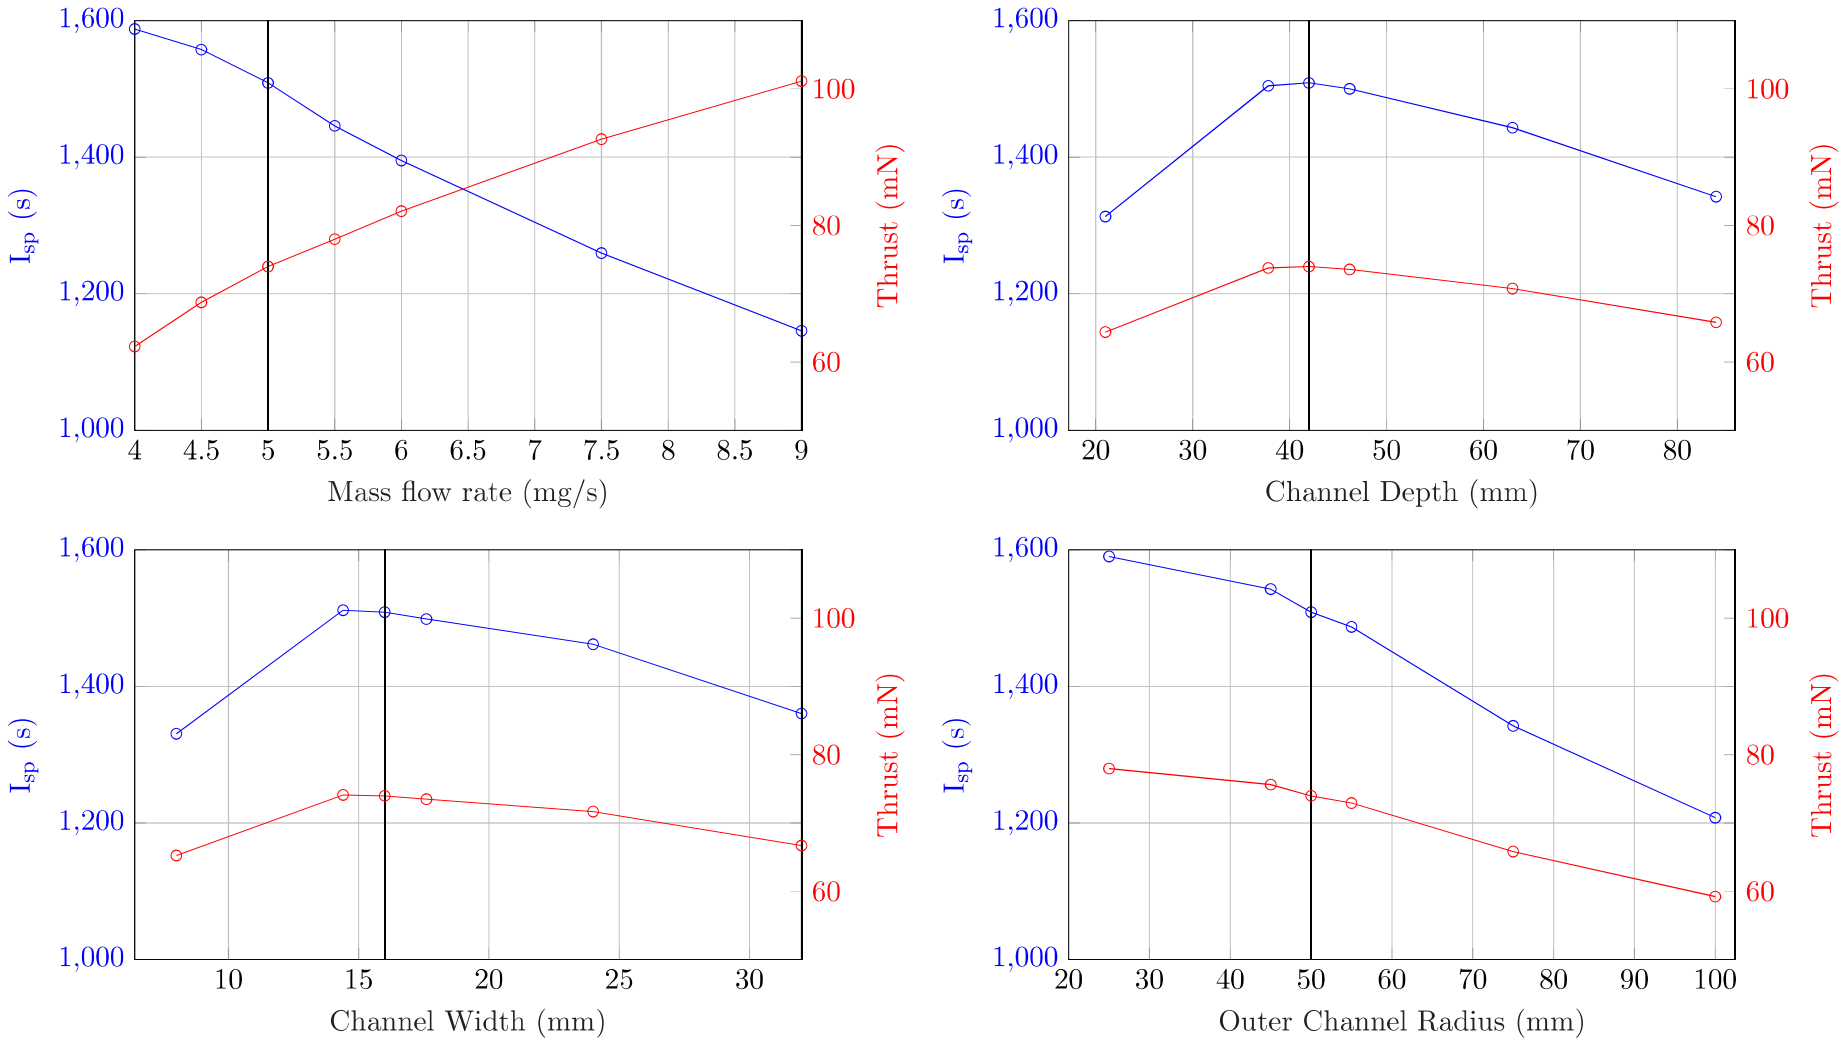

Figure 3.5: PlasmaSim sensitivity analysis for an SPT100 HET at constant power. The black line represents the actual value.

Increasing massflow at a constant power increases thrust at the expense of I_{sp}. This comes as no surprise: more neutral gas entering the thruster is contributes to cold gas thrust, but little extra ionisation, reducing the average exit velocity. The results show that while no true optimum mass flow rate exists, a value a $5\,\mathrm{mgs}^{-1}$ provides a good trade-off between thrust and I_{sp}. For the channel depth the PlasmaSim 0D results clearly show a thrust and I_{sp} peak at 42 mm, which is the channel depth of the SPT-100. The simulation shows that if the channel depth were greater or less than 42 mm, the performance would decrease. This suggests that the code is able to successfully optimise the channel depth of a HET. Similarly, the simulation correctly identifies 16 mm as very near to the optimal channel width of the SPT-100, further demonstrating that PlasmaSim 0D is able to optimise for channel width. For the channel outer diameter the PlasmaSim 0D results suggest that both I_{sp} and thrust could be increased if the diameter of the SPT100 was reduced. This indicates that either the PlasmaSim 0D results are not correctly simulating the true effect of changing the channel diameter, or a smaller diameter has not been implemented in the SPT-100 design due to other factors, such as as smaller diameter causing thermal issues, or mechanically interfering with the central magnet coil.

The PlasmaSim 0D is a highly simplified PIC code, which operates with a single cell, a homogeneous electric field and a disregard of all magnetic effects. Nevertheless, figure 3.5 clearly shows that PlasmaSim 0D is able to accurately determine which HET channel width and depth produce the greatest performance. The data also provides important performance predictions on how the choice of mass flow rate and channel diameter. These results suggest that PlasmaSim 0D development has reached its goal, and the code is ready to be implemented for its primary purpose: optimising a thruster for oxygen operation.

3.6 Simulating an Oxygen HET

With PlasmaSim 0D proving to be an effective tool for optimising the geometry and mass flow rate of a HET, and the code now having the ability to operate on different propellants, the oxygen HET optimisation could now be performed. As with the previous simulations, the code was run for 10^{-3} s simulation time, with a time-step of 10^{-7} s and 34000 super-particles. Unlike for xenon, the literature was unable to offer an estimate of the electron temperature within an oxygen HET. This was assumed to be the same value as for the xenon case: a Boltzmann distribution centred on 11.5 eV, similarly the ions were given a temperature of 0.1 eV.

Unlike with the SPT-100 optimisation, there was no obvious initial geometry or mass flow rate to use as a starting point for the oxygen HET optimisation. In comparison to the xenon plasma, the oxygen plasma was much more difficult to ignite and sustain, proving to only reach a stable state for a small subset of configurations. This change was very much expected, given the much small electron impact ionisation cross section of oxygen. As predicted, switching from xenon to oxygen heavily impacted performance with a lower fraction of particles becoming ionised, which reduced the thrust and I_{sp}

The thruster was optimised for specific impulse. Initial exploration demanded trial and error to find stable operating conditions. From here a manual gradient ascent was performed to determine the greatest specific impulse for a stable plasma. Eventually the following optimal configuration was found:

- Mass flow rate = 1.2 mgs^{-1}
- Channel depth = 45 mm
- Channel width = 5 mm
- Channel outer radius = 12.5 mm

- Electric field strength = 600 Vm^{-1}

Note that such a gradient ascent represents a local maximum, yet not necessarily the global maximum. However, very few other stable operation points were found, suggesting few if any other local maxima exist.

This setpoint resulted in a thrust of 11.3 mN and an I_{sp} of 958 s when operating at 526 W of discharge power. This performance is substantially lower than the predictions in the previous chapter, which suggested that changing propellant to oxygen from xenon should double I_{sp} from around 1500 s, but half thrust from 83 mN. One of the large unforeseen discrepancies between simulating the two propellants was that the oxygen was only able to sustain low power discharges, below around 750 W. At powers above this, the plasma became unstable and the discharge extinguished. Even for xenon HETs, operation at such low powers generally results in poor performance.

As with the SPT-100 validation case, the PlasmaSim 0D simulations for the oxygen HET were run with each dimension and mass flow rate off optimal, at a constant power. The power used here was 526 W, which again was the power found at optimal operation. The results of this sensitively analysis are shown in figure 3.6.

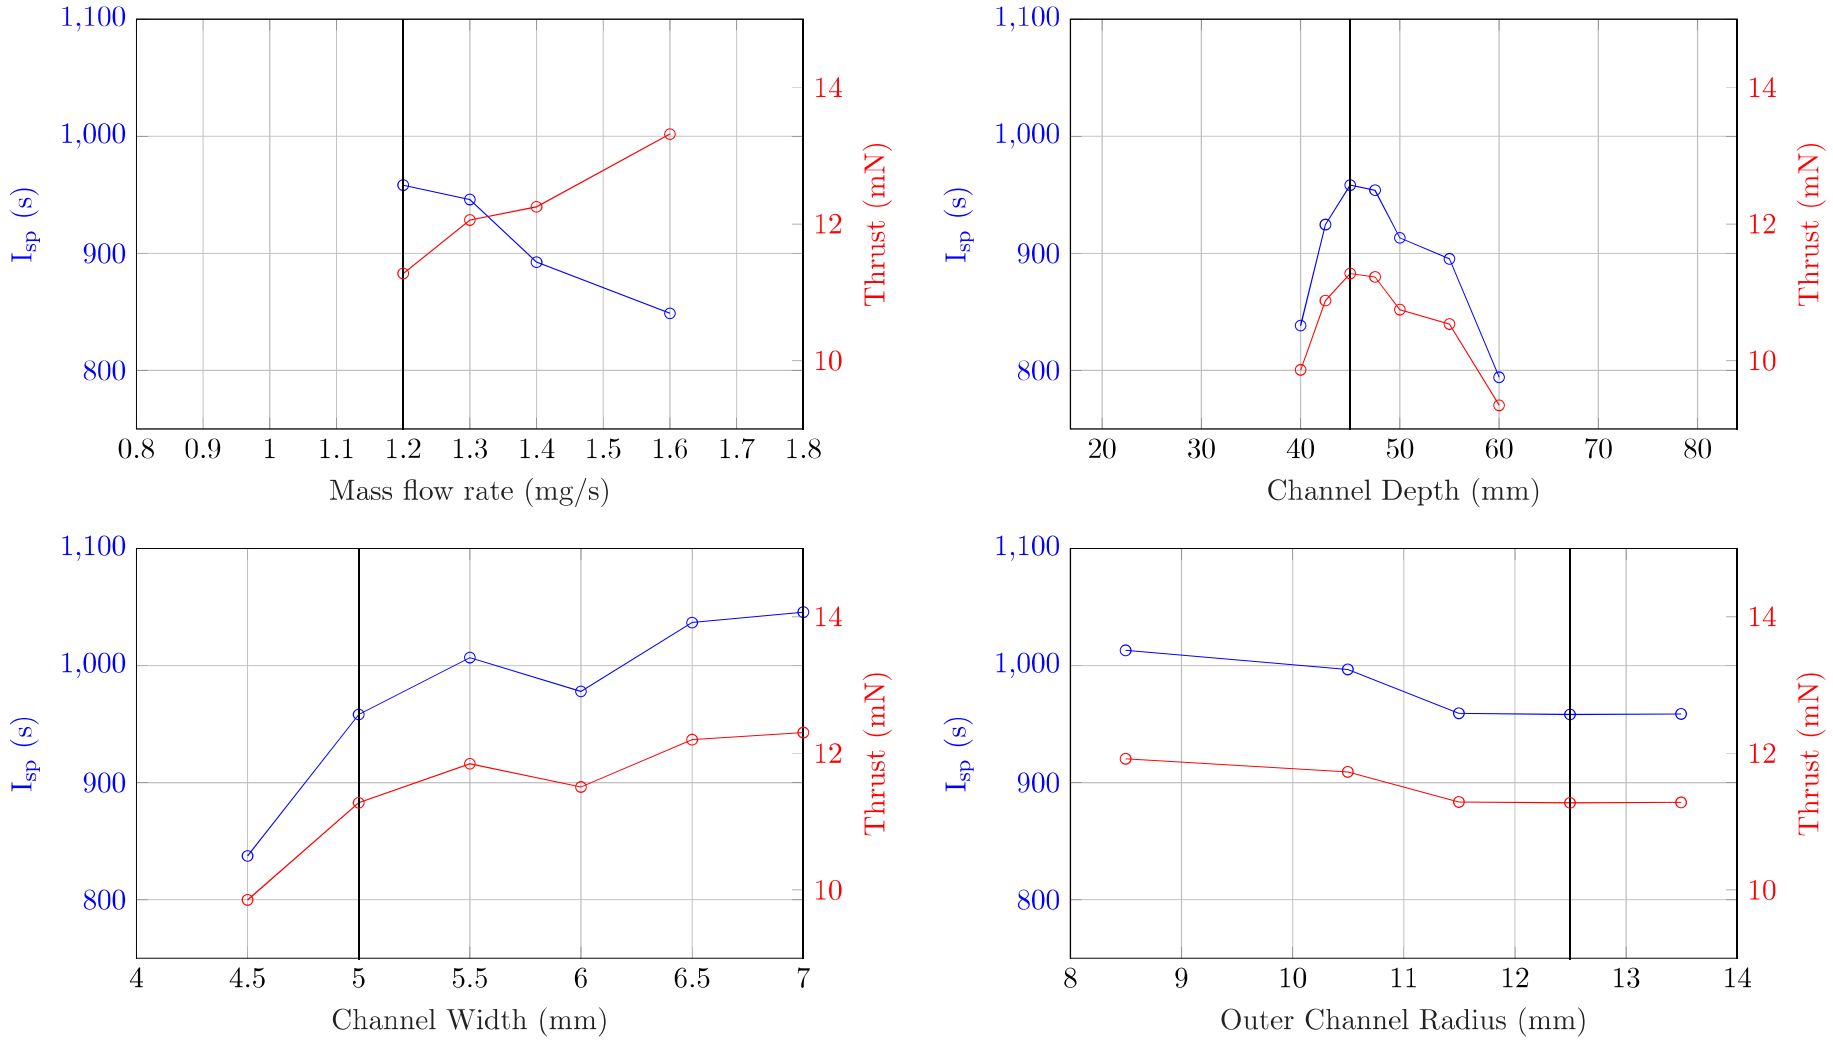

Figure 3.6: PlasmaSim sensitivity analysis for an oxygen HET at constant power. The black line represents the optimal value.

As with the SPT-100, reducing the mass flow rate of the oxygen thruster with this geometry at a constant power reduces thrust but increases I_{sp}. For mass flow rates below 1.2 mgs^{-1} the plasma was unstable at the target powers, suggesting that this represents the greatest I_{sp} which can be achieved with this geometry. In comparison to the SPT-100 which operates on 5 mgs^{-1} of xenon which is 51.3 sccm, this oxygen thruster operates at approximately a quarter the mass flow rate, yet a similar volumetric flow rate of 50.4 sccm, due to the discrepancy in atomic mass between xenon and oxygen.

Figure 3.6 shows that a maximum thrust and I_{sp} is found with a channel depth of 45 mm, with performance quickly dropping off if the channel depth is extended or contracted from this point. Note that this is a very similar channel depth to the 42 mm of the SPT-100.

A channel width of 5 mm, and an outer channel radius of 12.5 mm was chosen, although figure 3.6 shows that performance increases at greater widths and lower radii. This choice was made as in reality the channel width cannot be changed completely independently of the inner and outer radius: r_o=r_i+W. It can be seen in figure 3.6 that channel radii below 11.5 mm and channel widths above 6 mm offer better performance. However, practically speaking such a configuration would produce a thruster with an inner channel radius of only 5.5 mm (see figure 3.2). For a real HET, the inner channel walls and the internal magnetic pole must fit within the internal radius. To allow for real world construction of the thruster, a minimum limit of the internal radius was set at r_i=7.5 mm. Constraining the minimum possible internal wall radius forces either the outer radius to increase, or the channel width to decrease, both of which are shown in figure 3.6 to reduce performance. The PlasmaSim 0D simulations also identified a hard limit for the maximum outer channel radius, in the sense that an outer channel radii above 13.5 mm is unable to sustain a plasma. In order to not exceed this limits, the minimum channel internal radius of 7.5 mm was used and an outer radius of 12.5, which ultimately resulted in a

channel width of 5 mm.

Although greatly simplified, the quasi-zero dimensional version of PlasmaSim has proven to allow for simple optimisation of a HET for mass flow rate, channel depth, width and radius for xenon. By performing a similar simulation with oxygen, an optimised geometry and set-point was found for a novel oxygen HET. From here, much could be done to improve the model, such as simulating the magnetic field, magnetic effects, or increasing the dimension. It was at this point that it was decided instead that the numerical modelling be validated in the laboratory, in a bid to practically test the effectiveness of a water electrolysis HET, and attempt to identify unforeseen stumbling blocks which may stand in the way of such a technology becoming commercially competitive.

Chapter 4

The Water Electrolysis Hall Effect Thruster

The oxygen HET simulations performed in the previous section were to assess the sensitivity of such a device to operating points such as discharge voltage, mass flow rate, and channel dimensions. The results of the optimisation dictated the geometry of a thruster which was then built to experimentally test what was assumed to be highest performing HET. The device which was built has been named the Water ElecTrolysis Hall Effect Thruster or WET-HET. The thruster is designed to experimentally test the feasibility of operating a HET designed to consume the products of water electrolysis.

This chapter describes the steps taken to build a working, physical device around the channel dimensions provided by the plasma model. The design of the WET-HET required an iterative cycle of magnetic, thermal and mechanical modelling before final drawings were given to the Aeronautics department workshop for construction. Along with the typical independent variables such as mass flow rate, discharge power, and magnetic field strength, the WET-HET has the unique ability of being able to have the channel depth and the magnetic topology tuned in the laboratory. This design allowed for a rich trade-space of dimensions and set-points to explore.

Table 4.1: Optimised dimensions and operational points of the WET-HET compared to that of the SPT100.

Parameter	**Units**	**WET-HET Value**	**SPT100 Value**
Mass flow rate	mg/s	1-1.5 (O_2)	5.12 (Xe)
Volume flow rate	sccm	40-60 (O_2)	52.08 (Xe)
Discharge Power	W	1000	1350
Channel Depth	mm	35-60	42
Channel Width	mm	5	16
Channel Outer Diameter	mm	25	100

The oxygen thruster optimisation of the previous chapter serve as the core design requirements for the WET-HET. These optimised dimensions and operations points are show in in table 4.1 along with typical values for a xenon SPT-100 for comparison. The two thrusters share a similar volume flow rate and discharge power. The WET-HET is however considerably smaller: the discharge channel of the WET-HET has a cross sectional area which is smaller than that of the SPT-100 by a factor of 13.44. The WET-HET as been designed such that the channel depth can be changed (see section 4.4 below). Depending on the channel depth, the volume of the WET-HET is smaller by a factor of 9.4-16.1, which results in a surface-to-volume ratio which is 2.5-2.6 greater for the WET-HET. These values have important physical consequences. The smaller channel cross sectional area, but same volume flow rate determines that the neutral gas density of the WET-HET is considerably greater, which increases the rate of electron-neutral collisions. The surface-to-volume ratio of a HET is very important as broadly speaking, ions are generated in the bulk of the plasma, and energy is deposited at the surface of the plasma[97]. This dictates that the WET-HET deposits a larger fraction of the input energy to the chamber walls, and therefore has a reduced thrust efficiency. Not only is there more energy being deposited on the walls of the WET-HET, but these walls make up a considerably smaller area, meaning high temperatures and thermal stress were expected.

From a purely engineering perspective, the requirements of an electric propulsion device like

the WET-HET can be broken into five elements:

1. **Magnetic circuit:** the thruster must be able to generate a radial magnetic field of at least 300 Gauss which spans the channel. In a typical HET, the magnetic field only spans the channel in a narrow region directly at the channel exit plane. The magnetic topology of the WET-HET allows the thickness of this magnetic region to be tuned.

2. **Thermal Management:** thermal energy is deposited onto the anode and channel walls and must be radiated into space at steady state without surpassing the maximum operational temperature of any components.

3. **Electric circuit:** the thruster must be able to deliver a current of up to 10 A to the tip of the anode from the anode power supply. The anode must be able to reach a potential of 1000 V without arcing to the structure of the thruster.

4. **Fluidic path:** the thruster must be able to supply up to $2\,\mathrm{mgs}^{-1}$ of oxygen into the thruster channel and distribute it azimuthally without leaking.

5. **Mechanical design:** the thruster must be physically able to be constructed, assembled and maintained using the tools available to us in the department. The design must not be compromised by extreme thermal cycling, vacuum cycling, or spurious high voltage discharges between components.

The WET-HET was designed first by constructing a simple axisymmetric magnetic model, followed by a rudimentary axisymmetric thermal model. Then a CAD model of the thruster was designed of higher complexity. From here all three models were iterated until all requirements were met. The CAD was finalised and drawings given to the workshop for construction. The final design of the thruster is shown both disassembled and assembled in figure 4.1, with a sectional view in figure 4.2 with materials selection list.

F

Figure 4.1: The WET-HET dissembled (top) and assembled (bottom). The lower section of the bottom photograph shows the optional water-cooling elements installed on the thruster, which were never used.

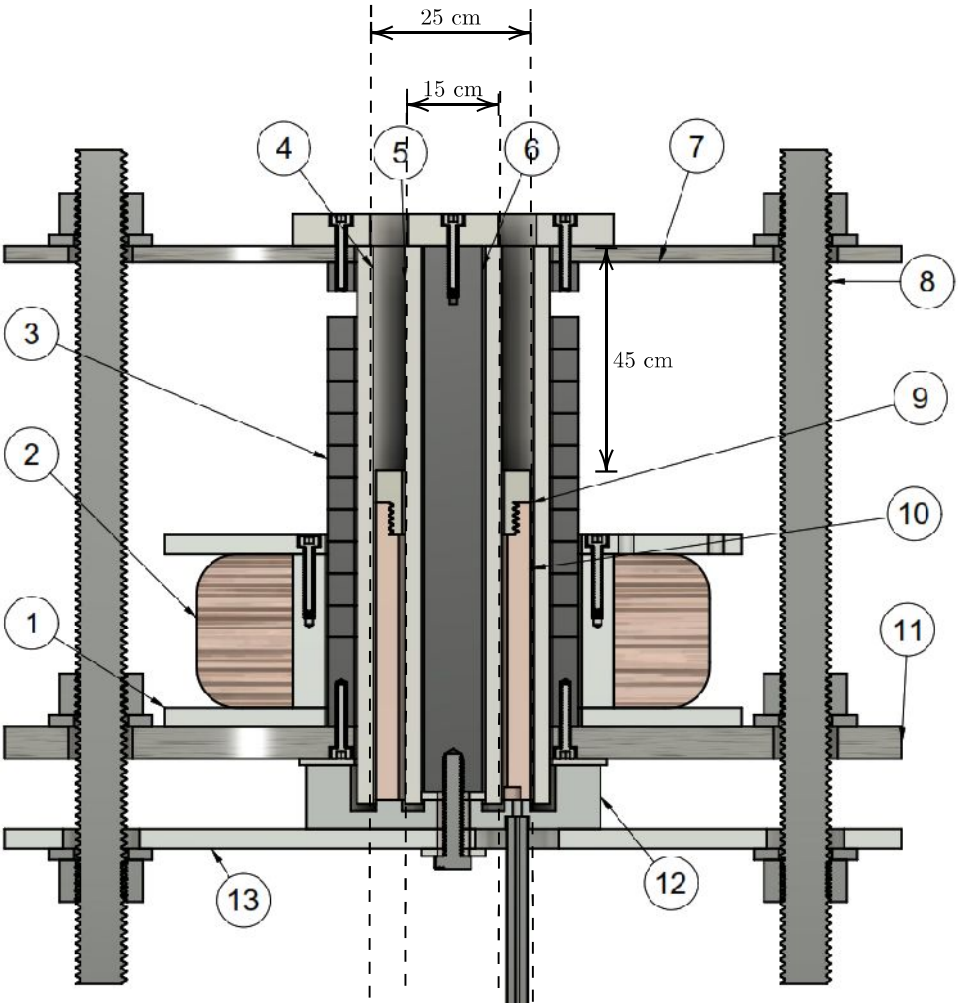

No.	Name	Material
1	Coil reel	Aluminium
2	Magnet coil	Copper
3	Outer pole stack	Soft Iron
4	Outer channel wall	Alumina*
5	Inner channel wall	Alumina*
6	Inner pole	Soft Iron
7	Front plate	AISI 430 Stainless steel
8	Threaded rod structure	Zinc Plated Stainless steel
9	Anode cap	Tungsten*
10	Anode	Copper
11	Back plate	AISI 430 Stainless steel
12	Tail cap	AISI 430 Stainless steel
13	Radiator	Aluminium

Figure 4.2: WET-HET cross section and material description. The asterisk (*) denotes that these material choices were changed from the original design (see text).

4.1 General Thruster Design

The general configuration of the WET-HET is unconventional for a HET, and is loosely based on a small HET by Guerrini et al.[98]. This configuration serves small thrusters well in that only one magnetic coil is needed, and this sits externally of the discharge channel. This differs from most typical HET designs which have several coils internal and external of the discharge channel. A further rare advantage of this design is that the channel walls consist of two commercial-off-the-shelf (COTS) alumina tubes, alleviating the need of expensive bespoke ceramic parts. The simple discharge wall structure allowed for both the discharge channel depth and magnetic topology to be easily reconfigured in the laboratory (see below). The thruster costs were further reduced by utilising relatively common materials and conventional in-house manufacturing methods.

The majority of the components are cylindrical and are stacked on top of one another. Four M8 stainless steel threaded rods and removable nuts provide stability but easy maintenance by sandwiching the component stack under compression. At the centre of the thruster sits the stainless steel tail cap, to which two concentric alumina tubes are fixed using a gas tight high vacuum adhesive called Torr-seal. This adhesive represents the only components that cannot be easily disassembled. A removable annular copper anode is screwed to the tail cap between the tubes. A stainless steel anode cap threads sits atop the anode. This anode cap was originally designed to be constructed from tungsten to sustain high thermal loads (see section 4.7 below). The discharge channel is the volume between the two alumina tubes and the anode cap. A 1/8th inch stainless steel feed-line welded to an orifice in the bottom of the tail cap allows propellant to flow between the inner face of the anode, and the internal alumina tube, before it is distributed throughout the discharge channel volume. A aluminium radiator plate fixed to the tail cap funnels heat from both the anode and the channel walls to be radiated into space. The copper magnet coil is wound on a removable coil reel, representing the only active magnetic component.

The inner magnetic pole comprises a solid soft iron rod, where the outer pole consists of a stack of ten soft iron rings which can be moved axial to change the magnetic topology (see below). Along with the soft iron components, the back plate, tail cap, threaded rod and front plate are made of high magnetic permeability stainless steel to conduct the magnetic field. Thin alumina sheets were cut as spacers to provide high electrical isolation but moderate thermal conduction.

Further detail of the design is best provided by describing how it meets each of the requirements above individually.

4.2 Electrical Design

For the WET-HET to sustain a discharge, the electrons were drawn from the external cathode into the channel to the anode. Igniting this discharge often require 1000 V of potential between the grounded cathode, and the positive anode. During ignition the high potential of the anode desperately sought to draw electrons from any source available. Without sufficient isolation, these electrons were drawn from the thruster structure, the facility or the thrust balance via arcing. Not only does such arcing further damage insulation but on occasion it would induce transient voltage spikes in surrounding cabling, damaging sensitive equipment. For this reason it was paramount that the anode, and all things in electrical contact with the anode were sufficiently isolated from all other metallic surfaces. The final electrical design of the WET-HET is shown in figure 4.3. The main structure of the thruster was in electrical contact with the thrust balance. This is shown in green. During testing the balance was kept at a floating potential to be representative of a spacecraft, meaning the green section was floated but generally close to ground. The blue section of figure 4.3 represents all components at the anode potential including anode, anode cap, tail cap and feed-line. During ignition all blue components attracted electrons from all other components, often causing arcing. The red components are ceramic insulators used to mitigate conduction and arcing. These are the two concentric tubes which comprise

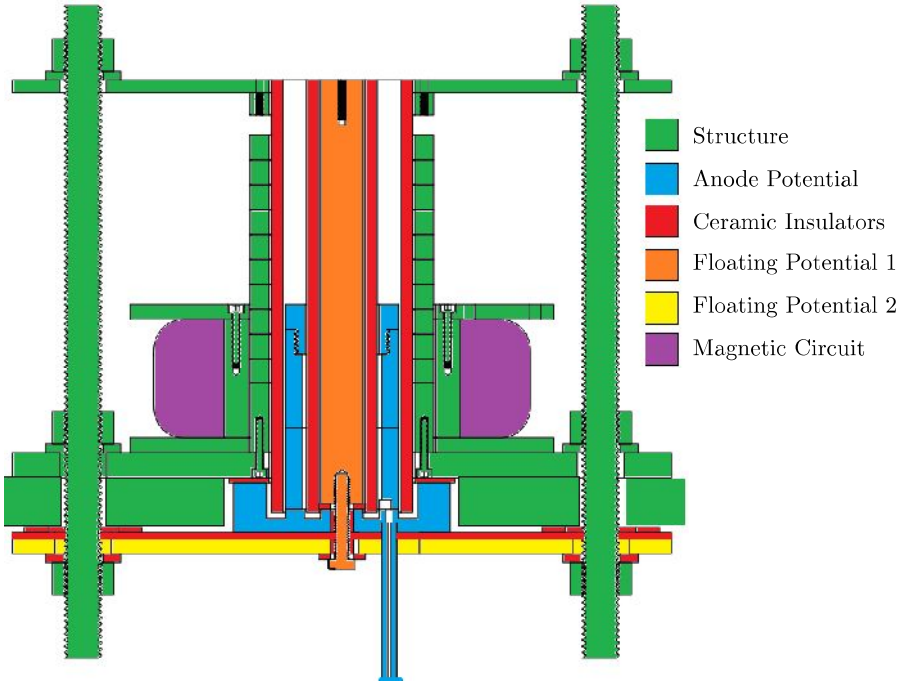

Figure 4.3: The electric design of the WET-HET. The colours represent the potential electrical status during operation.

the channel walls and several alumina spacer made bespoke of water-jet cut alumina sheet. The purple section is the magnet coil which had one end explicitly grounded and the other connected to a dedicated power supply at a maximum potential of 10 V. The orange components are the central pole and the core screw which were at their own floating potential. It was critical that the orange components did not come into contact with any of the blue components, as doing so raises the central pole to the anode potential, and thus provides an easy path for the electrons to the anode power supply. Not included in the original design shown in figure 4.2 is the insulating alumina sheet between the tail cap and the radiator which can be seen in figure 4.3. This was added such that the radiator is now at its own floating potential (yellow in the figure). Originally the radiator was anode potential, yet this drew electrons from the structure.

In general the electrical design worked as intended with a maximum anode current of 10 A passing through the device. Sparking and arcing between components before ignition was common. This was generally fixed with added insulation such as kapton.

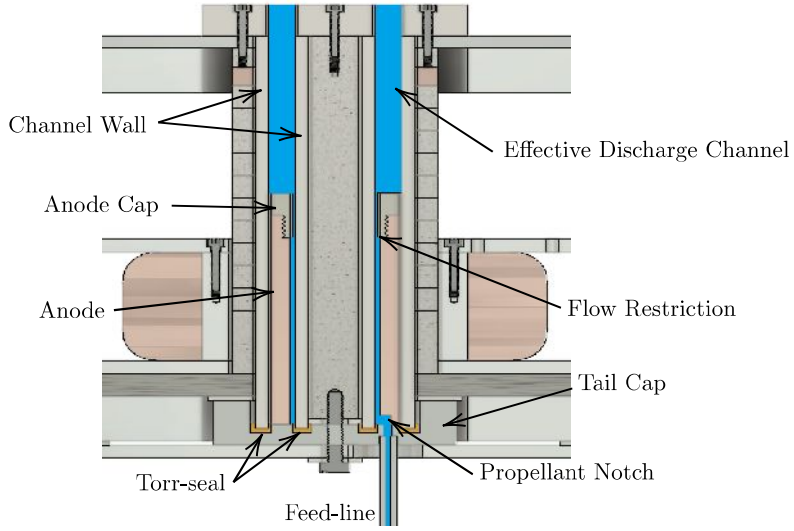

Figure 4.4: The fluidic design of the WET-HET with the propellant shown in blue. The yellow sections are the sealant which mechanically secured the channel walls to the thruster and mitigated propellant leaking out of the tail cap.

4.3 Fluidic Design

An overview of the propellant flow path is shown in figure 4.4, with the propellant shown in blue. The propellant enters the thruster via a single 1/8th inch Swagelok feed-line at the base of the thruster. This feed-line is laser welded to the bottom of the tail cap to align with a through hole that emerges in-between the two channel walls. This hole also indexes with a small propellant notch in the base of the anode. This notch can be seen in figure 4.5, and serves to divert the propellant radially inwards such that it flows into the space between the anode inner surface, and the inner channel wall. A sealant by the name of Torr-seal secures the channel walls to the tail cap, and prevents the propellant leaking out of the channel. Once in the base of the channel, the propellant moves axially towards the channel exit through the volume between the anode and the inner channel wall. The inner radius of the anode cap is 0.5 mm smaller than the inner radius of the anode. This reduction in cross sectional area forces the propellant flow to be slightly restricted as it reaches the anode cap. This flow restriction in labelled in figure 4.4. The restriction forces the propellant to spread azimuthally around the channel before emerging above the anode cap. It is important to note that the plasma is only generated in the presence

Figure 4.5: The bottom face of the anode showing four tapped holes and the propellant notch which diverts the propellant from the feed-line line into the inner section of the anode.

of a potential gradient. The entire anode is at the same potential, so the potential gradient is only generated above the anode cap. This means that the propellant is only ionised once it passes above the anode cap, and enters the main volume of the discharge channel. This is why only the volume above the anode cap is considered the effective discharge channel.

4.4 Channel Depth Adjustment

The WET-HET dimensions were optimised for oxygen operation using PlasmaSim 0D, however one of the known limitations of this code was its inability to simulate an oxygen molecule dissociation into atomic oxygen. Adapting PlasmaSim to accurately model the various reactions between different oxygen species would not only demand a high level of complexity and computational cost, but also a very accurate knowledge of the electron temperature, which is unknown.

The earlier simulations performed with ZDPlasKin did include 143 reactions between 15 different oxygen species, yet these results also required an estimation of electron temperature and were unable to simulate geometry dependent interactions. What the ZDPlasKin simulations did reveal, was that the neutral particle residence time in the thruster has a direct impact on the ratio of atomic ions (O^+) to molecular ions (O_2^+) generated. Each oxygen molecule requires 5.15 eV to dissociate to two atomic oxygen ions, or 12.07 eV to be ionised. From a thruster performance perspective, a greater fraction of atomic ions would be evidenced by a greater discharge current but a lower thrust for a given mass flow rate.

The neutral particles move down the channel at thermal velocities, meaning that the particle residence time is directly dependent on the channel depth L. A fundamental rule of HET design which concerns the channel depth is the Melikov–Morozov criterion which states that the channel depth should be much greater than the mean ionisation length λ_i to ensure high levels of propellant utilisation[99]. Here the mean ionisation length is the average distance that a

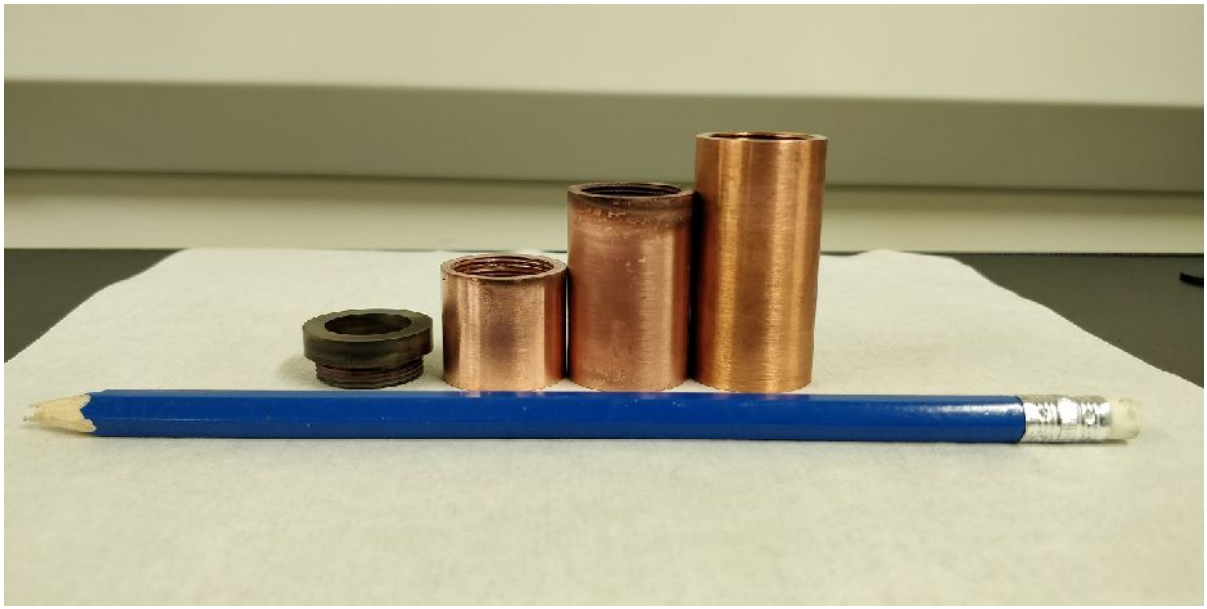

Figure 4.6: The stainless steel anode cap of the WET-HET (left) with three different length copper anodes. Combining the cap with each of the anodes results in a channel depth of 60 mm, 45 mm and 35 mm respectively.

neutral particle will travel before becoming ionised:

$$\lambda_i \ll L. \tag{4.1}$$

Selecting an appropriate channel depth is of critical importance: too shallow a channel and the neutral particles have too short a residence time to become sufficiently ionised. Too deep a channel produces more surface for the ions to recombine, and a greater fraction of atomic oxygen ions are produced, reducing the efficiency. Given the inability of PlasmaSim 0D to accurately simulate molecular dissociation, precise optimisation of the channel depth was not possible. Instead the WET-HET has been designed such that channel depth can be optimised experimentally. The copper anode was designed to be easily substituted by axially longer and shorter sections, such that the effective discharge channel depth can be altered. Three anodes of differing depth along with the anode cap are shown in figure 4.6. These three anodes were created such that WET-HET can be tested with a channel depth of 60 mm, 45 mm and 35 mm. By changing the thruster channel depth under constant discharge power, mass flow rate, and

magnetic field strength, an optimum channel depth could be determined experimentally.

We can estimate the mean ionisation length of the WET-HET using a zero dimensional calculation. For this we follow a similar approach to Garrigues[100], where the mean ionisation length is given by:

$$\lambda_i = \frac{v_n}{n_e k_i} \tag{4.2}$$

where v_n is the neutral particle velocity, n_e is the electron density, and k_i is the ionisation rate in $\mathrm{m^3 s^{-1}}$. To predict the electron density, we must make some assumptions. We assume that thruster has a high propellant utilisation efficiency, meaning the majority of the neutral particles are successfully ionised. This suggest that the ion density is similar to the neutral density: $n_i \approx n_n$. We also assume that the plasma is quasi-neutral, such that the electron density is also similar to the neutral density: $n_e = n_i \approx n_n$. We again follow the same approach as Garrigues[100] to calculate the neutral density within the thruster as a function of the mass flow rate $\dot{m}$, ion velocity v_{ion}, channel cross sectional area A_c, and propellant mass m_n:

$$n_e \approx n_n = \frac{\dot{m}}{v_{ion} A_c m_n} \tag{4.3}$$

Finally, we substitute equation 4.3 into equation 4.2 to get an estimate of the mean ionisation length:

$$\lambda_i = \frac{v_n v_{ion} m_n}{\dot{m} k_i} A_c \tag{4.4}$$

For comparison, let us calculate mean ionisation length for both the WET-HET operating on oxygen, and the SPT-100 operating on xenon. We calculate the ion neutral temperature assuming a neutral temperature of 1000 K, and like Garrigues[100], we use an average voltage drop potential of 50 V to calculate the ion velocity. For the ionisation rate we assume a mean electron temperature of 20 eV. We use the data presented in Garrigues[100] to approximate an ionisation

rate of $k_i = 2{\times}10^{-14}\,\mathrm{m^3 s^{-1}}$ for oxygen, and $k_i = 7{\times}10^{-14}\,\mathrm{m^3 s^{-1}}$ for xenon. For the mass flow rate, we use $1\,\mathrm{mgs^{-1}}$ for the WET-HET and $5\,\mathrm{mgs^{-1}}$ for the SPT-100[12]. This provides us with a mean ionisation length of 10.4 mm for the WET-HET, and 8 mm for the SPT-100. The small cross section of the WET-HET optimised by PlasmaSim has to some extent combated the poor ionisation rate of oxygen, yet not to such a degree that it rivals that of a xenon thruster. Note that this relies on several assumptions and assumes the two thrusters generate the same electron temperature. It also only considers the average distance travelled by a neutral particle before it is ionised, but does not account for the energy required to do so, the energy lost to heating the walls, or the recombination rate at the walls. The surface-to-volume ratio of the WET-HET is approximately 2.5 times that of the SPT-100, suggesting this will result in considerably more energy lost to the walls.

Next we consider the Melikov–Morozov criterion for the WET-HET by dividing the channel depth by the mean ionisation length calculated above. For channel depths ranging from 35 mm to 60 mm, this results in a $L/\lambda_i = 3.35$-5.74. This means that the average neutral particle will have between 3.35 and 5.74 ionisations as it passes out of the thruster. Note that under the same assumptions the SPT-100 only reaches 5.26 ionisations per neutral particle, which is unrealistically low, most probably due to the assumption of a 20 eV electron temperature. However, this estimation shows that if the WET-HET is able to achieve a similar electron temperature as the SPT-100, the longest channel depth of 60 mm can be expected to achieve a similar number of average ionisations per neutral particle and thus a similar propellant utilisation as the SPT-100.

4.5 Magnetic Design

The magnetostatics suit of Finite Element Method Magnetics (FEMM) version 4.2 was used to model the magnetic design of the WET-HET[101]. A simulation of the magnetic design is shown in figure 4.7, superimposed on the cross section of the WET-HET. Conventional HETs like that

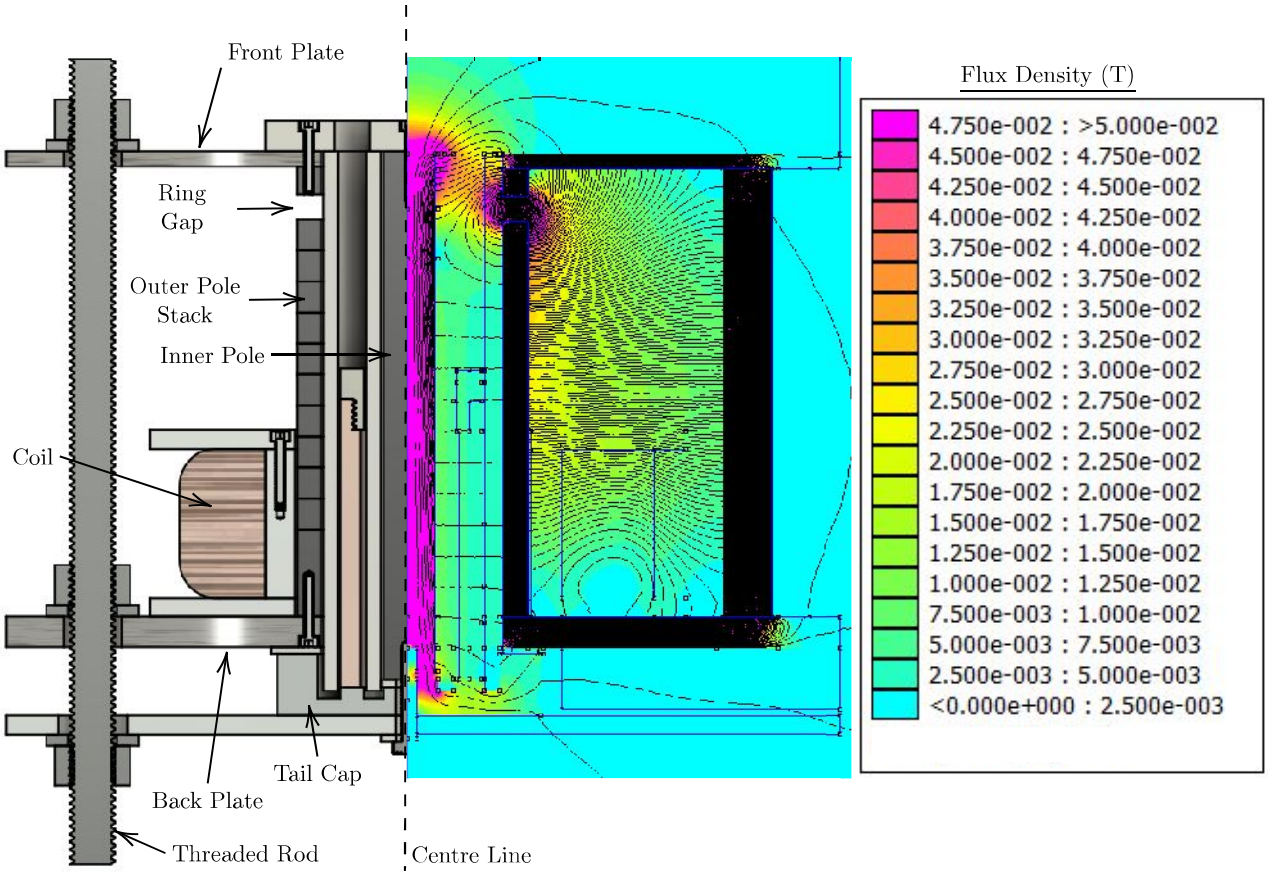

Figure 4.7: Magnetic flux density of the WET-HET with the default "0 rings raised" magnetic configuration (see section 4.6) and channel depth of 35 mm.

seen in figure 1.4 generally have permanent magnets or electromagnets both outside and inside of the thruster channel. Due to the small size of the WET-HET, a design was chosen in which the magnetic field is generated by a single coil outside of the channel. The coil was modelled in FEMM using a conservative current for the wire gauge of 0.5 A, and conservative wire packing density of 820 turns. During manufacturing, the same number of turns were used, resulting in a similar coil to that modelled.

The main purpose of the magnetic design is to generate a region near the channel exit in which the field lines pass radially across the channel. We see in the WET-HET that this region sits between the inner pole and outer pole directly above the ring gap. Both the inner pole and the outer pole stack are constructed of soft iron. A stainless steel with a high magnetic permeability was chosen for the front plate, back plate, tail cap and threaded rod to serve to complete the magnetic circuit.

An electromagnetic approach was selected over a permanent magnet configuration for two reasons: igniting the thruster is considerably easier without the presence of the magnetic field, and changing the current passing through the electromagnets allows for experimental optimisation of the magnetic field strength. A comparison of the simulated and experimentally measured radial magnetic flux densities along the thruster centre line are shown in figure 4.8. Note that the simulations overestimate the experimental data, particularly at the higher magnetic currents. Much of this discrepancy can be attributed to the fact that the FEMM simulations assume an axisymmetric design. The threaded rods are non-axisymmetric features of the magnetic circuit, and therefore reduce the measured field strength below the simulated value.

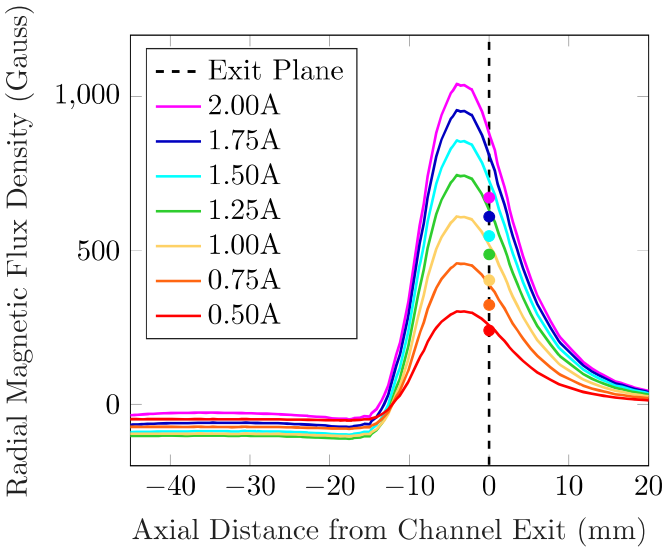

Figure 4.8: The radial component of magnetic flux density at different magnet currents along the channel centre line. Adopted from Schwertheim and Knoll[102]. The lines show the results of the FEMM magnetic simulations. The circles show the experimentally measured values using a gaussmeter. This configuration has a 45 mm deep channel and a magnetic topology of 0 rings raised (see section 4.6).

4.6 Alternative Magnetic Topologies

Traditional HETs have short magnetic regions near the exit of the discharge channel not unlike the trends seen in figure 4.8. In this region the presence of the radial magnetic field generates the Hall current: an E × B electron drift which raises the local electron density, and thus the ionisation rate. The large electron impact cross section of xenon means that even a thin magnetised region is sufficient to ionise a large portion of the neutrals, allowing such devices to routinely operate at propellant utilisation efficiencies over 80%[7]. Not only do oxygen molecules move approximately twice as fast as xenon atoms, but the electron impact cross section of oxygen is considerably smaller than xenon at a given electron temperature (see section 2.4). This suggests that oxygen can be expected to have a much longer mean ionisation length than xenon. In designing the WET-HET it was proposed that a traditional magnetic topology with only a thin magnetic region near the channel exit may be ineffective at sufficiently ionising oxygen. Poor ionisation would lead to a low propellant utilisation efficiency and thus a poor I_{sp} and efficiency. It was hypothesised that an exotic magnetic field topology may better suit such an

exotic propellant. Extending the magnetic region further into the discharge channel was expected to generate a Hall current and consequentially a region of high electron density which are both axially thicker. This thicker region of electron density that the neutral particles must travel through was expected to combat the poor propellant utilisation fraction one would expect from oxygen. This is a similar approach to that employed by Marchioni and Cappelli[60] in the Extended Channel Hall Effect thruster, where an extended channel and magnetic thickness was used to combat the poor utilisation of air.

The 0D PlasmaSim code used to optimise the WET-HET was electrostatic, meaning magnetic fields were not simulated. Furthermore, the code assumes a single homogeneous cell, meaning it is unable to simulate location dependent features like regions of increased electron densities. This meant that alternative magnetic topologies could not be trialled numerically.

In order to test exotic magnetic topologies, the WET-HET was designed such that the impact of changing the magnetic thickness could be tested experimentally. A unique feature of the thruster is that it has the ability to dramatically change its magnetic topology by changing the configurations of the magnetic circuit. The default magnetic topology is the one shown in figures 4.7 and 4.8, which only has radial components near the exit of the thruster channel. This axially thin magnetic region is a topology very similar to that of a traditional HET. We refer to this as the "0 rings raised" magnetic configuration. The external magnetic pole (number 3 in figure 4.2) comprises a stack of ten soft iron rings which can be individually raised and lowered. By physically raising a number of these rings up or down, the magnetic field topology of the thruster changes significantly: when a ring is raised a greater fraction of the thruster channel is exposed to the radial magnetic field, such that the magnetic thickness is increased. By manipulating ten rings, eleven possible magnetic topologies can be tested. These range ranging from the 0 rings raised default position with the magnetic region only covering a fraction of the discharge

channel exit, to the 3 rings raised position with approximately half the channel magnetised, to the extreme case of 10 rings raised in which radial magnetic field inhabits the entire discharge channel. Each of these is shown in figure 4.9, with the radial component of the magnetic field at the channel centre line shown in figure 4.10. Here that raising rings effectively lowers the position of the ring gap. Given the majority of the radial magnetic field is situated above the ring gap, lowering the gap increases the magnetic thickness.

It is important to note that the magnetic field topology can be changed independently of the magnetic field strength (which can be tuned by changing the magnetic current) and independently of the physical channel depth. This means that each of the eleven trends shown in figure 4.10 can be scaled in magnitude by changing the current going through the magnet coil, as shown with the 0 rings raised trend in figure 4.8. Similarly each of the eleven different magnetic trends in figure 4.10 can be tested with each of the three anodes, resulting in a channel depth of 35 mm, 45 mm or 65 mm. This does however mean that some combinations are redundant: consider figure 4.10, when the anode is placed such that the channel is only 35 mm deep, the 8 rings raised has an identical topology within the shortened channel as the 7 rings raised configuration.

4.7 Thermal Design

The thermal analysis suite in FEMM was used for axisymmetric thermal design of the WET-HET [101]. An unrealistically conservative scenario was simulated in which the thruster reaches a steady state operating condition at 1 kW, but 0% thrust efficiency, meaning that 100% of the input power is converted to heat. Further being conservative, it was assumed that all of this heating would be deposited on the face of the anode. The thruster surfaces in contact with the thrust balance were fixed at 400 K, and outer surfaces radiated to 0 K deep space. FEMM does not automatically capture the heat transfer of internal surfaces radiating onto one another, for example from the anode to the channel walls, which are separated by a 1 mm vacuum gap. To

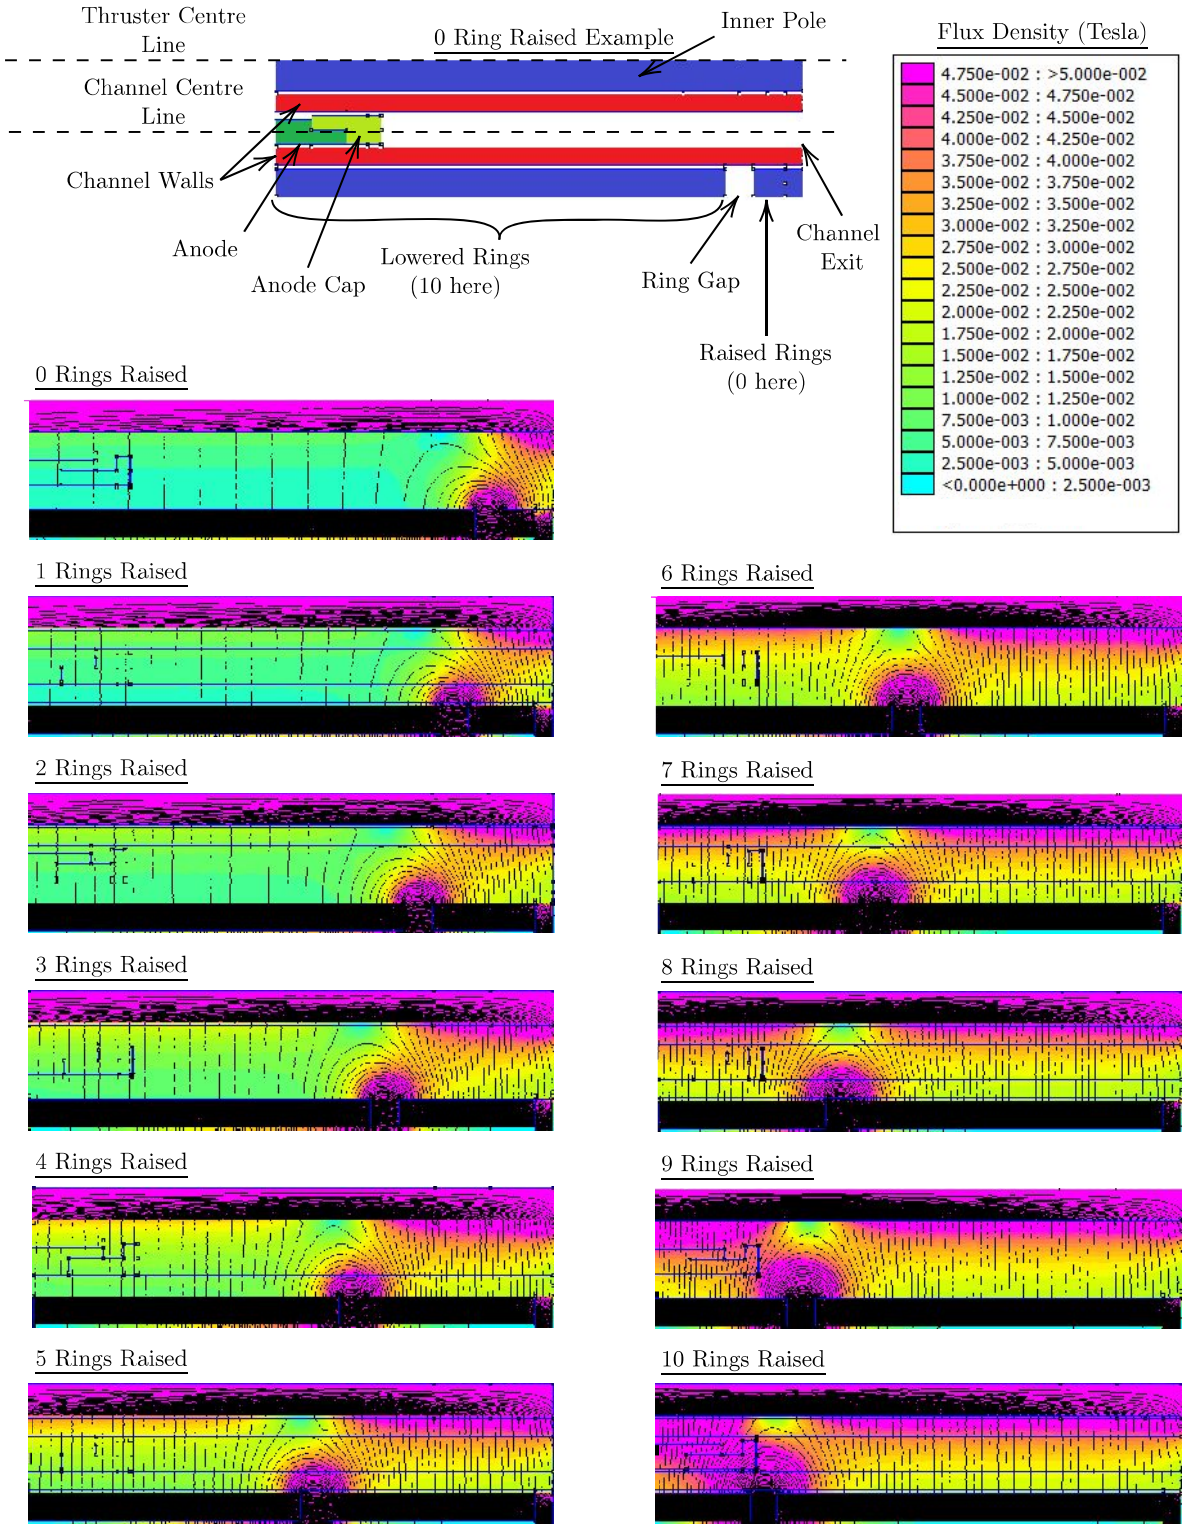

Figure 4.9: FEMM simulations of the magnitude of total magnetic flux density of the channel for all 11 magnetic field topologies. The channel depth is 60 mm for all cases. Each channel is the shown below the centre line, meaning that top edge is the thruster centre line and axis of rotation. The labelled example shows how the outer pole comprises a stack of 10 rings, with the lowered rings being separated from the raised rings by a ring gap. Raising a ring (moving it from the left to the right in the diagram) effectively moves the ring gap downwards (to the left in the diagram). This causes a greater fraction of the radial magnetic field to cross the channel.

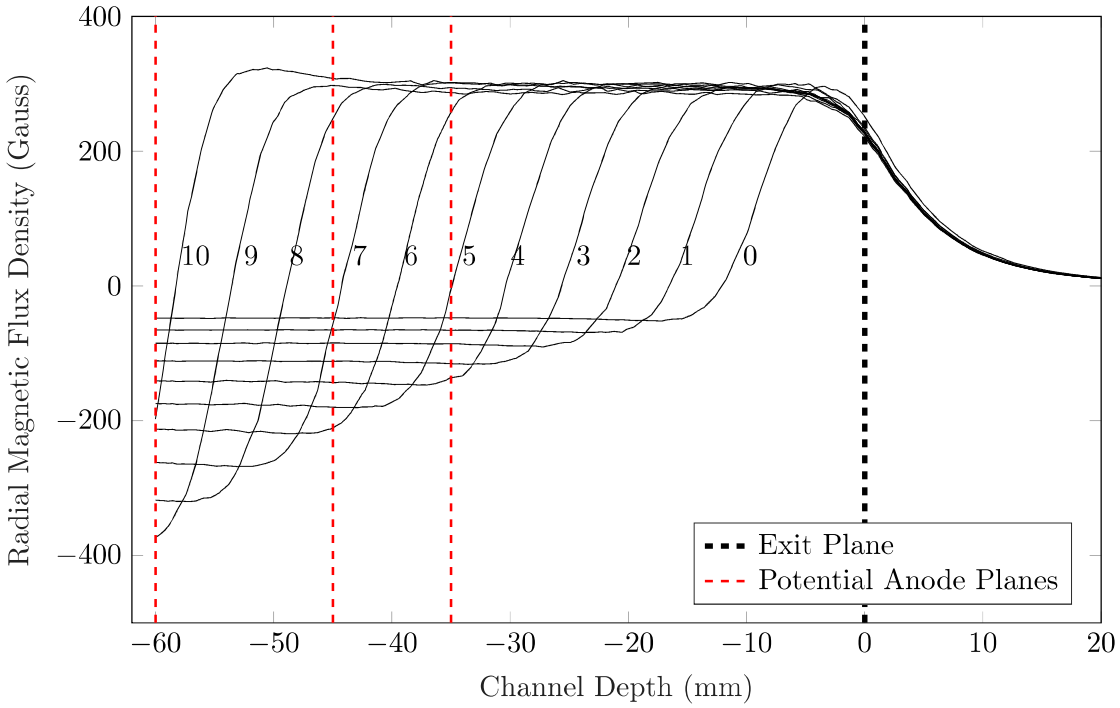

Figure 4.10: FEMM simulation data showing the radial magnetic flux density along the centreline of the channel for all 11 possible magnetic topologies of the WET-HET. The numbers represent the number of rings raised: e.g the trend labelled "10" indicates the "10 rings raised" magnetic topology, which is the configuration where the magnetic region is the thickest. The three tested anode planes are also shown, resulting in a channel depth of 60 mm, 45 mm and 35 mm respectively. Please note that the magnetic currents have been scaled such that all peak field strengths are equal.

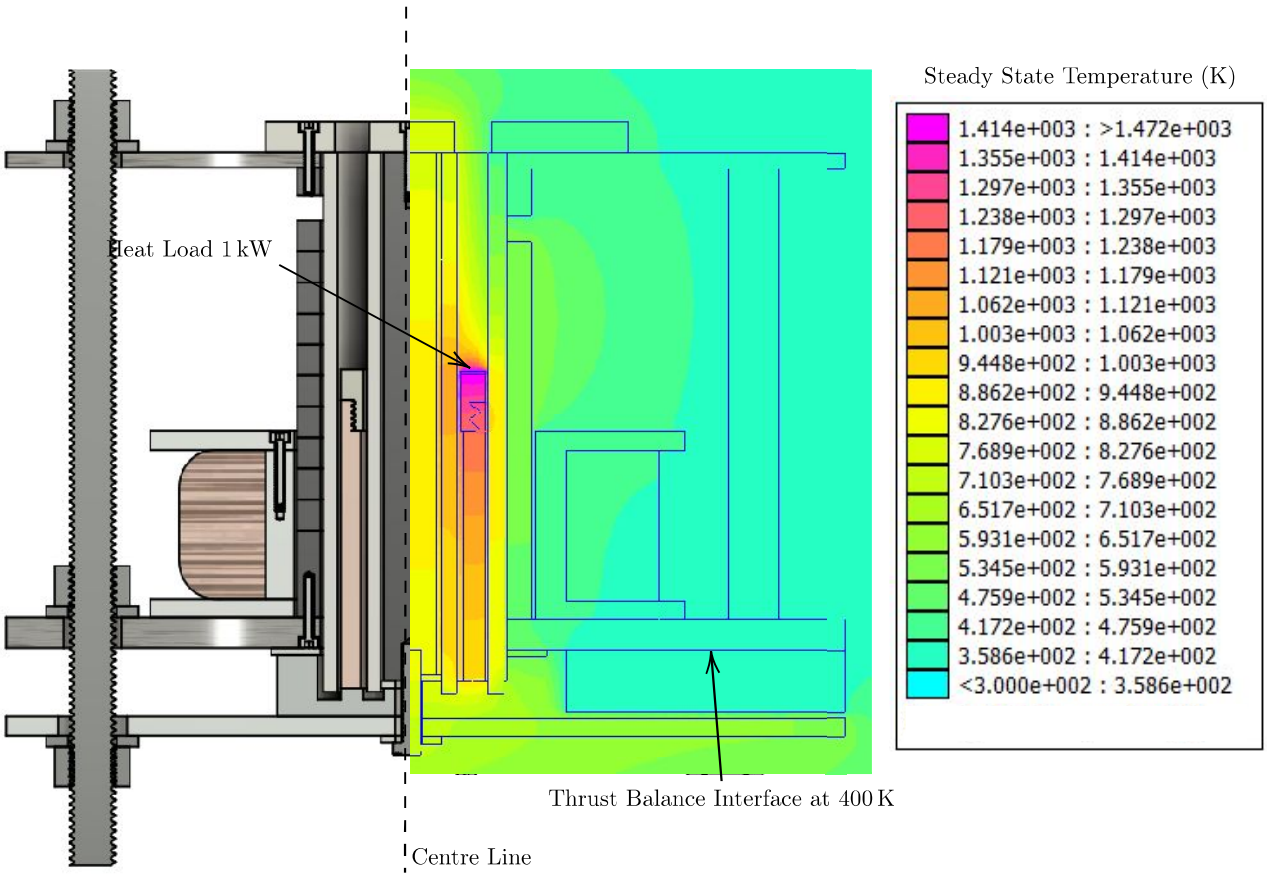

Figure 4.11: Axisymmetric thermal simulation of the WET-HET operating at steady state with a 1 kW thermal load on the anode face, a 400 K boundary at the thrust balance, and radiation to 0 K. The left hand side of the figure is the centre of the thruster and the axis of revolution

capture this the model was iterated by calculating the energy radiated from one surface using the heat and emissivity, and loading this extra thermal input into the facing surface. These iterations converged on a steady state solution. Under these conditions the final design was able to operate within the thermal limits of all components at steady state. The result of this simulation for the final design is shown in figure 4.11.

The thermal management strategy employed by the WET-HET is to conduct the heat from the anode cap thorough the anode, to the tail cap, then to the radiator at the back of the device. This radiator was coated with a high emissivity coating to promote thermal dissipation

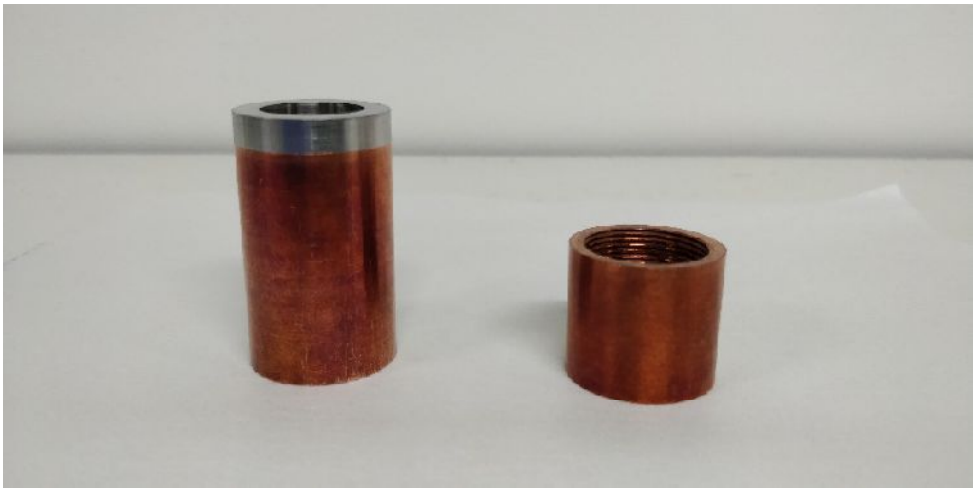

Figure 4.12: Two copper anodes, one with (left) the threaded anode cap attached, one without (right.

through radiation. Early simulations found that if using a anode made of only copper, the top surface would exceed the maximum operational temperature of copper. A proposed solution to this was to construct the anode of a refractory metal like tungsten, however the low thermal conductivity of tungsten (in comparison to copper) led to tungsten also exceeding its maximum operational temperature. The eventual solution was a trade-off between the two suggestions where the majority of the anode is copper, to promote conduction to the radiator, but the top section consists of an "anode cap" of tungsten, which is threaded into the anode to sustain the extreme temperatures required. Constructing and threading the anode cap component out of tungsten would cost more than the remainder of the thruster combined, so it was decided to begin experiments with an anode cap made of stainless steel, with the intention of moving to tungsten if necessary. The copper anodes with and without the stainless steel anode cap is shown in figures 4.6 and 4.12. In reality it was found that thermal simulations had in fact been overly conservative and the stainless steel anode cap was never replaced. Furthermore, it was found that the main thermal load was deposited not on the anode, but on the inner ceramic wall. This is discussed in further detail in section 6.2.

Visible only in the photograph but not the diagram of figure 4.1 are the optional water cooling

elements, intended to actively cool the thruster. The WET-HET was originally designed to be able to operate continuously at powers above 2 kW by fitting four copper water block heat exchangers to the thruster radiator. This was never implemented, as unexpected thermal bottlenecks prevented the thruster from operating at high powers continuously. Even without the active water cooling the thruster was operated for very short periods at powers above 3 kW.

4.8 Mechanical Design

The mechanical requirements of the WET-HET design dictated that the thruster must be able to sustain many testing cycles throughout the experimental campaign. This required extreme temperature cycling, but also continuous disassembly and reassembly for maintenance. Over 18 months of testing, the thruster met all mechanical requirements. The only major mechanical issue was the sealant between the ceramic channel walls and the tail cap. This interface is shown in yellow in figure 4.4. This ceramic-metal bond requires a sealant which is mechanically strong, gas-tight, robust to thermal cycling in excess of 1000°C, and vacuum compatible. Several different vacuum compatible high temperature sealants and adhesives were trialled, yet all failed after several hours of testing, mostly due to the thermal cycling. The vacuum gas sealant called Torr-seal was used as it was the easiest to apply and lasted more thermal cycles than the others.

4.9 Full Characterisation of the WET-HET

The WET-HET was designed to operate over a wide and highly dimensional parameter space. Like a traditional HET, the WET-HET can be operated at a range of discharge powers, mass flow rates and magnet field strengths. Two additional design parameters unique to the WET-HET are the anode channel depth and magnetic field topology. Both of these parameters can be changed independently of all others. This provides us with a rich parameter space of thruster configurations and operational set-points to explore.

Chapter 5

Experimental Design

The previous chapter introduced the WET-HET: a thruster which has been computationally optimised to operate on the products of water electrolysis. This chapter describes the experimental set-up and procedures required to test the performance of this thruster.

The Boltzmann vacuum facility within the Imperial Plasma Propulsion Laboratory (IPPL) was used to simulate the vacuum of space for all WET-HET testing. This facility and its capabilities are described in section 5.1. The neutralising cathode which accompanies the WET-HET is critical in providing the necessary electrons to ignite the plasma discharge and neutralise the ion beam. The filament plasma bridge cathode which was used in this experiment is presented in section 5.2. Electrical thruster characteristics such as discharge voltage and current are easily measured from the power supplies driving the device, however obtaining thrust measurements is considerably more difficult. Section 5.3 describes the design and operation of a specialised instrument called a thrust balance which was developed in collaboration with the Electric Propulsion section of the European Space Agency (ESA) to measure the performance of thrusters like the WET-HET. Section 5.4 describes how the WET-HET, cathode, and thrust balance are integrated within the Boltzmann vacuum facility, along the necessary electrical and fluidic supply systems. The thrust measurement procedure and the software for thruster control are described

in sections A.4 and 5.5. Although the majority of experiments utilised krypton as the operating propellant of the cathode, a small subset of experiments utilised additional hardware to safely supply the cathode with hydrogen. These experiments are described in section 5.6.

5.1 The Boltzmann Vacuum Facility

The Boltzmann facility is the largest vacuum facility in the IPPL. The facility was constructed for the IPPL in 2018, and consists mostly of 316 stainless steel. The main chamber is 2 m long with a 1.5 m diameter, and is connected at the mid point to an orthogonally orientated smaller load-locked or "hatch", chamber. A large pneumatic gate valve serves to link and seperate the two. The hatch has a diameter of 0.75 m and a length of 1.5 m. The facility is shown in figure 5.1.

Three tiers of vacuum pumps must be operated continuously to bring the facility to its minimum pressure:

- **The Primary Pump (also refered to as the Roughing Pump.)** After closing the facility a Leybold LV 80C dry vacuum pump is the first pump to be started, and the only pump that does not rely on other pumps to operate. The input pressure can be as high as atmospheric, and can reach as low as $5\text{x}10^{-2}$ mbar. The outlet feeds directly into the extraction system at atmospheric pressure. This pump takes approximately 90 minutes to reach a "rough" vacuum near $1\text{x}10^{-1}$ mbar.

- **The Turbo Molecular Pumps.** The main facility is pumped by two Pfeiffer ATH 2303 M turbo pumps, with an addition single smaller Pfeiffer HiPace 700M turbo pump pumping the hatch. Each of these pumps is situated in series with the primary pump, such that the turbo input is the chamber, and the turbo exhaust is the input to the primary pump. The turbo molecular pumps can only function once the roughing pump has reduced the chamber pressure below $1\text{x}10^{-1}$ mbar and they require their exhaust to be pumped

Figure 5.1: The Boltzmann vacuum facility shown with main chamber shut (top) and open (bottom). WET-HET and author added for scale.

by the roughing pump as they operate. They take approximately 20 minutes to spin up, and approximately 2 hours to spin down. Each of the large turbo pumps can pump up to 2150 Ls^{-1} of nitrogen, which results in a final pressure as low as 1x10^{-5} mbar when all turbo pumps are spun up.

- **The Cryo pump.** The Leybold 6000H is better described as a cryo panel than a pump. This panel can only be operated when all turbo and roughing pumps are operating at full power. The seperated cryo compressor is situated outside of the facility and works to pump heat away from the cryo head by compressing and expanding helium which circulates in a closed circuit. The external cryo-head, is mounted on the outside of the facility and is in very good thermal contact with the cryo panel within the chamber. When fully operational, the cryo panel is cooled to around 40 K. Any gas which has a freezing point above this temperature (such as nitrogen, oxygen, krypton or xenon) solidifies when contacting the cryo panel, further reducing chamber pressure. When the cryo panel is said to be pumping a gas, what this really means is that the panel is solidifying and trapping a gas to the surface of the panel. Note that the cryo will pump very high volumes of oxygen and nitrogen (up to 16000 Ls^{-1}) but is incapable of pumping any gas which has a lower freezing point than 40 K, such as hydrogen. The panel takes approximately 3 hours to cool down, and 3-5 hours to "re-generate" in which the panel heats up back to room temperature, releasing frozen gases from its surface. This release of gas means that the turbo pumps and roughing pump must be operating the entire time that the cryo is cooling down, cold, or warming up.

With all pumps on, the pressure in the Boltzmann facility reaches pressures as low as 2x10^{-6} mbar. This pressure rises to around 2x10^{-4} mbar when 80 sccm of propellant is fed to the thruster and 15 sccm of propellant to the cathode.

The facility is monitored and controlled by a control panel and an embedded programmable

logic controller. The pressure of the facility is measured using two Pfeiffer PKR 251 Pirani Cold Cathode vacuum gauges: at pressures above approximately $1x10^{-3}$ mbar these devices utilise an internal Pirani gauge, with a cold cathode gauge being employed at lower pressures. The temperature of the cryo panel is monitored using a type DT 670 silicon diode produced by Lakeshore.

5.2 Cathode

A vital component of the water electrolysis Hall effect thruster concept relies on the LaB_6 cathode operating on hydrogen. Unfortunately, no LaB_6 cathode was available at the time, although such a device is nearing design completion. This meant that it was not possible to perform a complete system test with the WET-HET operating on oxygen, and an accompanying LaB_6 cathode operating on the stoichiometric equivalent of hydrogen. Such a test is planned upon construction of the new cathode.

The IPPL facilities were designed with traditional xenon based electric propulsion needs in mind, which did not encompass hazardous gases like hydrogen. Also, the cryo pump was unable to pump hydrogen, suggesting operating with this gas would results in considerably greater facility pressures. Due to the danger hydrogen poses, safety restrictions did not allow for experiments which required the use of both compressed oxygen and hydrogen. To abide by these strict rules, for the majority of WET-HET testing krypton was used in place of hydrogen as the operating gas of the cathode. It is difficult to predict how the performance of the WET-HET will change with krypton is used in place of hydrogen. In order to test this, a small subset of experiments were able to be performed in which hydrogen was used for the cathode. Performing these experiments required considerable upgrades in health and safety system, along with updated protocols which resulted in a considerably more complex experiment: a small quantity of hydrogen was stored in a dedicated plenum which was housed within the evacuated facility,

such that both the vacuum and the chamber itself acted to isolate the gas from oxidisers and ignition sources. From here the hydrogen was supplied to the cathode, with the thruster itself operated on the oxygen. Safety was managed by restricting the total hydrogen volume in the laboratory, restricting the hydrogen pressure, never leaving the hydrogen unattended, isolating the hydrogen in the evacuated chamber, and diluting all pumps with a purge gas system which housed a massflow monitor and automatic shut down system. This set-up is described in more detail in section 5.6. This was only a very small subset of experiments, meaning that all set-up and results presented here utilised krypton as the cathode working gas unless specifically stated otherwise.

The same cathode was used for all WET-HET test firing. This device was of the filament plasma bridge neutraliser type, meaning electrons are generated by thermal emission of a high temperature tungsten filament in the presence of a working gas[103]. The device was developed in house to be low cost, propellant agnostic and easy to repair and maintain. The cathode design was originally developed and constructed spanning two MSc projects[104;105]. Further developments since then have lead to the current design of the device shown in figure 5.2. The filaments used were commercially available 0.5 mm diameter 2% thoriated tungsten welding electrodes. A dedicated power supply draws approximately 25 A of DC current through a single filament, raising the temperature above 2500 K. At these temperatures the filament thermally emits electrons. Recent modifications of the cathode design have led to a unique twin filament configuration: the current only passes through one filament at a time, such that only a single filament is heated at once. In the event of a filament failure, the second filament can be powered without disruption to the test. This is a great improvement on the previous design, because filament failure no longer necessitates facility venting.

Once the filament is heated, the cathode is ready to be ignited. A potential of 100 V is es-

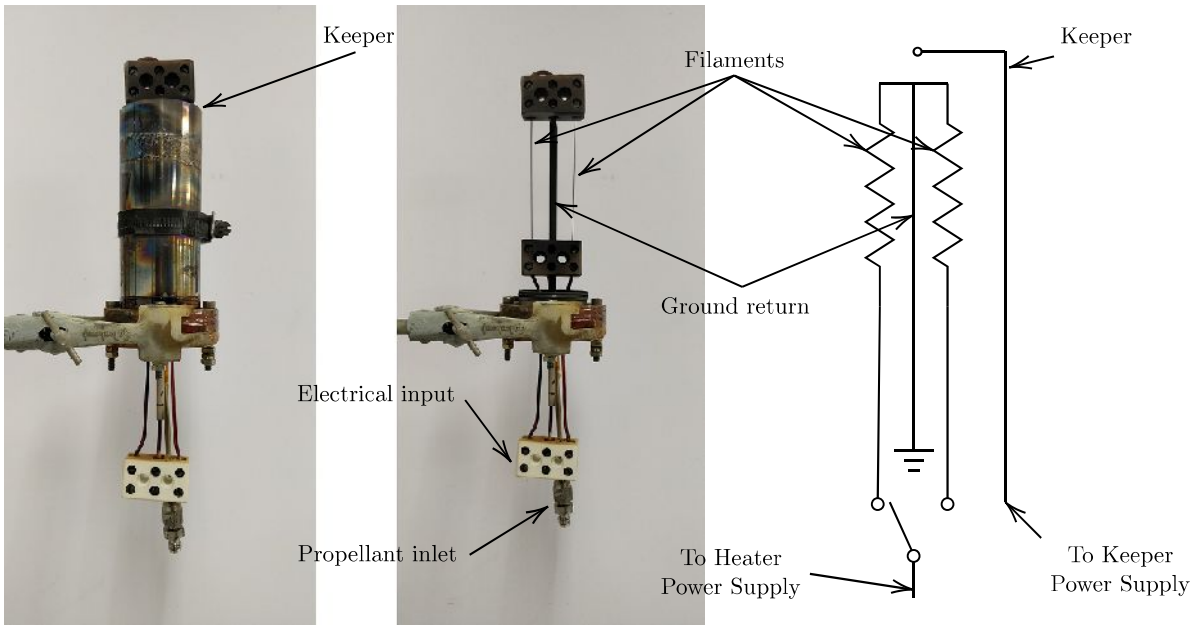

Figure 5.2: The cathode shown with keeper (left) and without keeper (centre) and as a circuit diagram (right).

tablished between the grounded filaments and the external keeper electrode using a dedicated keeper power supply. A physical gap between the keeper and the grounded filament inhibits current flow. Next, the propellant is fed into the base of the device, with the keeper helping to direct the gas to pass between the filament and keeper, raising the local pressure slightly inside the device. The heater current is gradually increased until the tungsten filament thermally emits electrons. Near a filament current of 25 A, a threshold is reached and the cathode ignites: the thermally emitted electrons are sufficiently energetic to bridge the gap between the filament and the keeper, producing a stable and continuous plasma discharge. The current passing through this discharge is referred to as the keeper current. This is observed by the keeper power supply switching from voltage controlled mode at 100 V to a current controlled mode of 400 mA and less than 20 V. This keeper current is maintained throughout the test, and serves as a small plasma from which the larger discharge plasma can be more easily established.

The centre of the cathode orifice was positioned 58 mm in front of the exit plane and 61 mm below the central axis of the thruster. The cathode was tilted slightly less than vertical, such

that the axis of the cathode made an angle of 82.5° with the axis of the thruster.

5.3 Thrust Balance

It comes as no surprise that one of, if not the most crucial performance metric of any propulsion system is the thrust it can produce. Measuring thrust is however significantly more difficult than measuring input electrical power or supplied mass flow rate. Once thrust has been measured, it can be combined with power and mass flow measurements to calculate not only critical system level metrics such as specific impulse, anode thrust efficiency, and TTPR, but also to estimate important plasma characteristics such as propellant utilisation fraction, mean ionisation length and even electron temperature. Although thrust can be estimated indirectly from plume diagnostics generally direct thrust measurements are preferred as they are inherently more accurate compared to indirect thrust estimation through measurements of the plasma plume. At the beginning of this project the IPPL lacked a thrust balance, so a MSc student by the name of Guanyun Liu began designing such an instrument. This is detailed in Liu [106], and is hereafter referred to as the IPPL balance. By the time the MSc project was complete the IPPL balance had been constructed in a basic form yet it had not been tested. Shortly after this the IPPL was awarded a contract with the European Space Agency (ESA) to design, construct and commission a thrust balance which was to be used for electric propulsion testing in the ESA Propulsion Laboratory (EPL) in the Netherlands. This second balance will be referred to as the ESA balance. ESA were also generous enough to include funding to improve the IPPL balance, such that at the completion of the project the ESA balance and the IPPL balance would be identical and both meet ESA's requirements. Considerable modifications were required to the existing IPPL balance such that it would meet ESA requirements. The final design was then replicated to create an identical ESA balance. These modifications included designing and fabricating mechanical improvements, two independent calibration systems, a thermal management system and all support electronics, along with the software, operating procedures and post-

processing procedures. Once both balances and their support electronics had been constructed, an inter-laboratory test campaign was undertaken to verify and validate the two instruments. This campaign required our team to test one thrust balance in our laboratory in London and then replicate the same experiment in the EPL with the second balance. The design process and test campaign results were published in Review of Scientific Experiments by Schwertheim et al.[107]. This chapter presents a condensed description of the design, procedures and performance of the thrust balance, given it was integral in characterising the WET-HET.

5.3.1 The Difficulty of Measuring Millinewtons

The thrust produced by EP thrusters is notoriously difficult to measure. Typical electric thrusters have weight of 200-600 times the thrust they generate[108]. For traditional off-the-shelf force measurement devices such as load cells and microbalances, the thruster weight is too great to tare, meaning that the thruster mass dominates over any thrust signal. To complicate things further, the thrusters must be tested within a vacuum environment. Therefore, the data acquisition must be performed remotely. The thrust balance can either be calibrated under atmospheric conditions with the influence of the vacuum environment considered as a source of possible error, or calibrated in-situ within the vacuum chamber remotely.

Given the lack of off-the-shelf force measurement devices which meet these particular requirements, most laboratories choose to develop their own thrust balances. The most common method of achieving this is with a pendulum thrust balance: the thruster is mounted horizontally to a moving platform so that the thrust vector is orthogonal to the gravity vector. Any generated thrust leads to the balance being perturbed from its equilibrium position, by the flexing of a flexure. The displacement distance is measured with a high degree of precision. A calibration

process is undertaken to characterize the measured displacement when a known force is applied to the balance. Once the thrust balance is calibrated, the force that a thruster generates can be determined accurately from the measured displacement. The conversion factor from balance displacement to applied force is called the sensitivity of the balance S_{cal} and usually has units of millinewtons per millimetre or similar.

Pendulum thrust balances can further be divided into three subsets of configurations: hanging, inverted, and torsional pendulums. For a hanging pendulum the moving platform is suspended below the flexures in such a way that both gravity and the flexure spring force act as the restoring force of the pendulum[109;110;111]. This style of balance exhibits less sensitivity when compared to other devices, yet hanging pendulums boast an intrinsically higher stability due to gravity acting as the dominant restoring force. As the name suggests, inverted pendulum thrust balances position the thruster above the flexure. In this configuration the flexure forces counteract the gravity forces, resulting in a less stable yet more sensitive thrust measurement[108;112]. Torsional pendulum thrust balances are constructed in such a way that the thrust generates a torque orthogonal to the gravity vector so that the flexure force presents the only restoring force. By decoupling the influence of gravity on the restoring force, the sensitivity of the torsional pendulum thrust balance can potentially be the highest amongst the three configurations, with the tradeoff of added complexity and generally higher cost[108;113;114].

5.3.2 Thrust Balance Overview

The main design driver for the IPPL and ESA thrust balances was that each device could be used to perform thrust measurements over a wide range of current and future thrusters. This ranged from low power electric thrusters of around 5 mN, to chemical micro-thrusters with thrust up to 2 N. Extending the thrust range to yet lower thrusts would have necessitated a considerable amount of vibration isolation, and thus dramatically increased the cost and complexity. A

hanging pendulum style balance was chosen for its simplicity and stability, given that a sufficient sensitivity could be achieved for the desired thrust range. The confidence in any thrust measurements is underpinned by a robust and accurate thrust balance calibration. The balance design incorporates two independent calibration systems, with one serving to ensure calibration traceability, and the other to ensure calibration stability. The actual displacement measurement is performed by an off-the-shelf triangulation laser with sub-micrometer precision. Many similar thrust balances use analog electro-mechanical devices such as Linear variable differential transformers (LVDTs) to measure displacement, yet the laser system was specifically chosen as it is less vulnerable to Electro-Magnetic Interference (EMI). Thermal stability is critical for thrust balances since the restoring force of flextures can change with temperature. While most thrust balances adopt active cooling systems to ensure thermal stability, the design chosen instead opts to hold the platform at an elevated temperature, swapping a complicated water cooling system for simple heaters. The accuracy, repeatably and precision of the thrust balance met all requirements set by ESA.

The complete thrust balance is shown in figure 5.3. The basic principle of the balance is a hanging pendulum style balance consisting of three flat platforms. The base and the top platform of the balance provide stability and rigidity to the structure. Ceramic breaks and nylon bolts separate the top level from the four legs, such that the top and central platform are electrically and thermal isolated from the base. Given that the aluminium base sits directly on the vacuum facility floor, the base is always electrically grounded and near room temperature. The central platform is suspended from the top platform by four parallel linkages, each containing two flextures. These flextures allow the hanging central platform to be displaced from its equilibrium position and swing only in one dimension while staying level. A laser triangulation system which is fixed to the base of the balance measures the displacement between the stationary base and a laser target attached to the moving platform. The thruster is attached to the platform

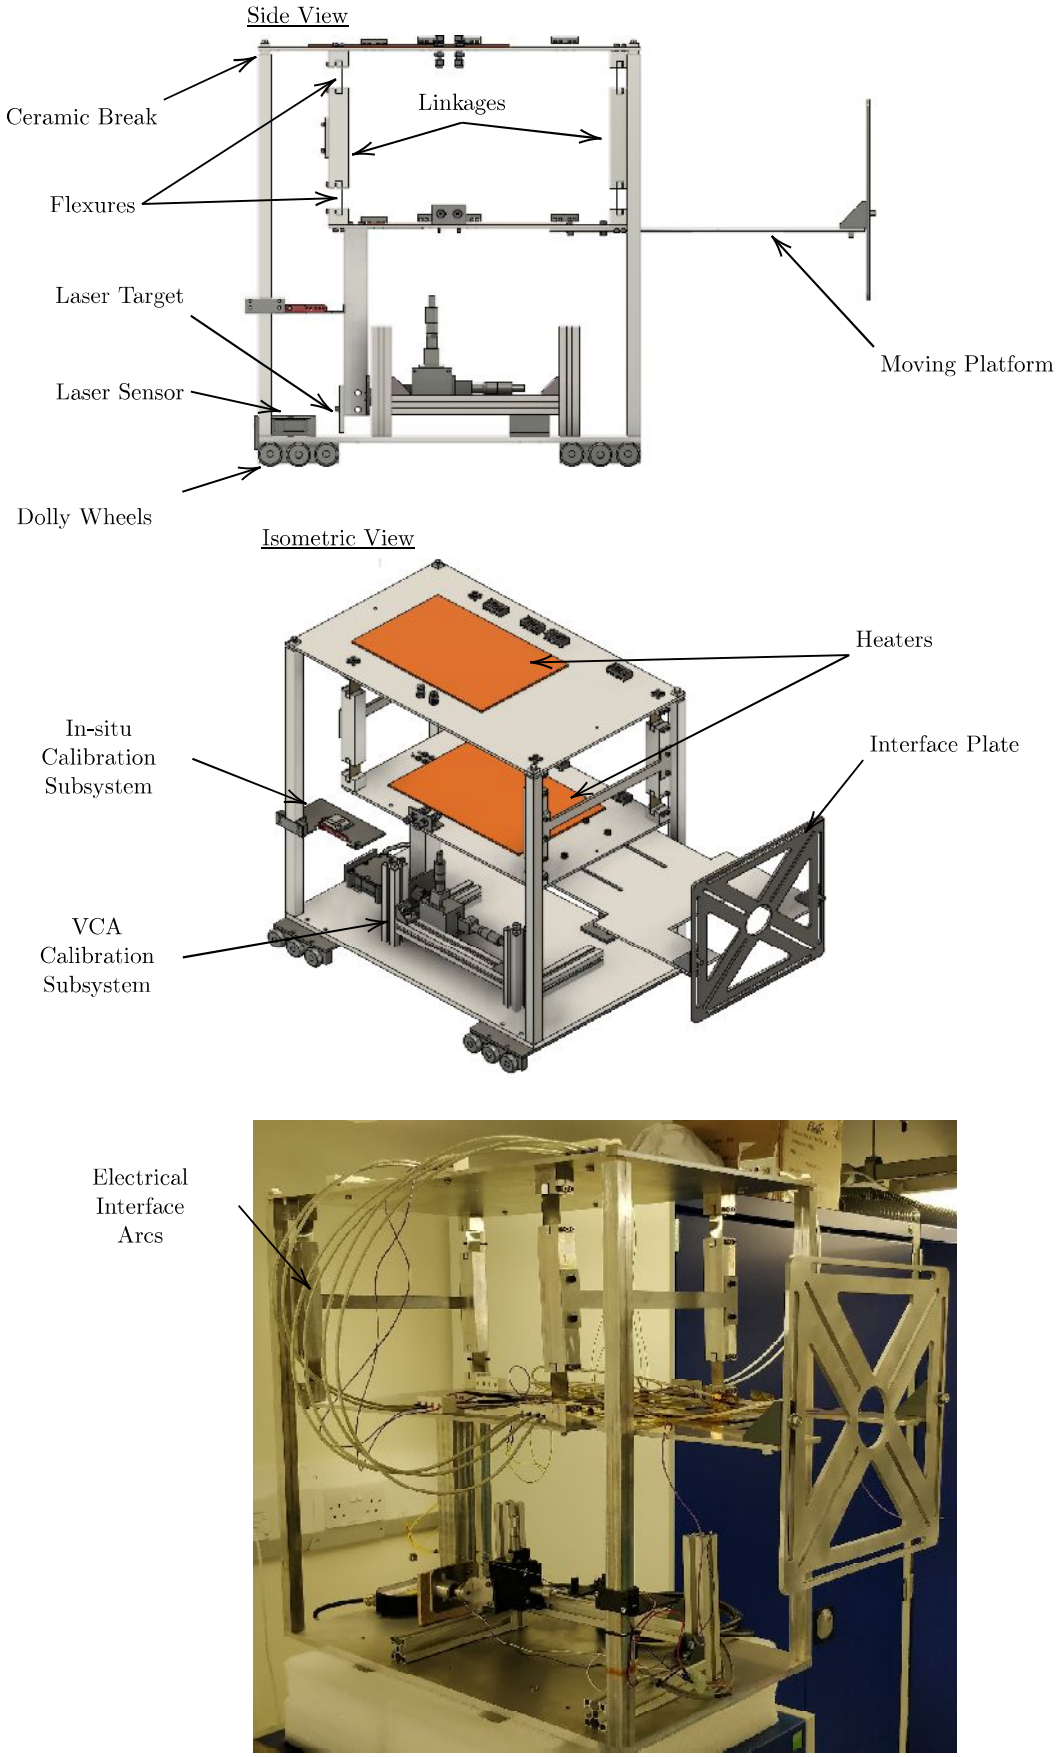

Figure 5.3: Side view, home view and photo of the thrust balance used for WET-HET testing.

such that the thrust axis is parallel with the swinging axis. The laser measures the change in the equilibrium displacement between the thruster being on and being off. With appropriate calibration the sensitivity of the balance can be determined and used to calculate the magnitude of the force required to generate this displacement. This force is the thrust of the thruster.

Upgrades

The following components have been upgraded on the thrust balance since the MSc project concluded:

- **Platform extension** ESA requested for the central platform of the balance to be extended beyond the footprint of the balance. This places the thruster further from the structure of the balance, and thus reduces any interference or arcing between the balance and the structure. This also allows the thruster to be pushed through the gate valve of the vacuum facilities in the EPL, such that the balance and thruster are in separate chambers.

- **Interface plate** The interface plate is a generic pattern designed to ensure that the instrument can accommodate a wide range of thrusters and cathodes.

- **Dolly wheels** These components are only present in the ESA version of the balance, and allow the balance to be moved forwards and backwards remotely, such that the thruster can be moved from the hatch of the vacuum facility to the main chamber via the gate valve.

- **In-situ calibration system** This was originally designed to be the only calibration system of the balance. Guanyun had already constructed the system, however electrical set-up and software integration was still required before operation. This is further explained in section 5.3.5.

- **VCA calibration system** This is explained in detail in section 5.3.5. The Voice Coil

Actuator (VCA) is the critical component that allows one to directly measure the sensitivity of the balance, which is the relationship between displacement and thrust. During the upgrade this system was designed, constructed and debugged.

- **Thermal management system** Heaters and thermal probes were added along with the electronics to command and control them. This is further described in section 5.3.4.

- **Electronic support box hardware and software** The electronics and software required to command and control the various subsystems is described in section 5.3.6.

- **Thrust measurement and post-processing procedure** The steps required to perform a direct thrust measurement, along with the post processing, and error propagation are described in sections 5.3.7 and 5.3.8 respectively.

5.3.3 Mechanical Design

The basic configuration of the balance was based off of the design of Pottinger et al.[109], which has shown great success with a range of thrusters[109;115;116]. Having two flexures in each linkage creates a double pendulum configuration, which ensures the platform and test article remain horizontal at all times, while restricting all off-axis movement. The flexures are rectangular sections of stainless steel shim stock. By switching between three sets of flexures of varying thicknesses, the thrust sensitivity can be increased across almost three orders of magnitude. Flexures with a thickness of 25.4 μm are used for thrust below 5 mN. For medium thrusts up to 300 mN flexures which are 50.8 μm thick are used, and flexures of 101.6 μm thickness are used for thrusts above 300 mN. The test article, which is generally a thruster and a cathode, is mounted at the end of the central platform using an interface plate. The laser displacement sensor (Micro Epsilon model optoNCDT-1750) is fixed to the stationary base of the balance, and measures the distance to a ceramic target attached to the moving platform. Any force along the thrust axis displaces the central platform. The displacement of the central platform changes the relative

distance between the laser and the target, which is measured by the laser sensor.

There are slight mechanical differences in the design of the IPPL and ESA balances. The ESA balance has dolly wheels such that it can be mounted on linear guide rails and moved within the hatches of both the Corona and the SPF vacuum facilities. The length of the ESA thrust balance was also slightly different to the IPPL balance in order to accommodate to the length constraints of the ESA facility hatches. Differences between the two balances have resulted in a discrepancy of the balance sensitivity of less than 3%. This discrepancy is negligible, considering that dismantling and rebuilding an identical experiment has resulted in a change of sensitivity of 5%. Such minor changes are expected and accounted for during the calibration procedure detailed below.

The photograph in figure 5.3 shows the electric interface arcs between the balance top level and the moving central platform. Similar arcs were used to route the fluidic lines from the top level to the platform. These wide arcs lie orthogonal to the thrust axis using as flexible materials as possible. This is done to reduce any stiffness that they may add, and to ensure any thermal expansion of the leads generates a force orthogonal to the thrust axis.

5.3.4 Thermal Design

Many electrical and chemical propulsion systems produce considerable amounts of heat when operating. The low pressure in the facility does not allow for convective cooling, meaning the majority of the heat generated is conducted through the thrust balance. This presents a major issue, as the restoring force of the flextures is dependent on temperature. This means that the calibration of the thrust balance is only valid at the temperature that the balance was calibrated. Traditionally thrust balances are kept thermally stable using a closed loop water cooling system which carries heat out of the facility to be dissipated. Not only is such a system complicated and

expensive, but the movement of water through the system can introduce unwanted vibration. The IPPL and ESA thrust balances also achieve thermal stability, yet by other means: before testing a "set temperature" is decided on, usually around 50°C. Two 200 W heaters are used to raise the balance to this temperature before calibration or thrust measurements are performed. The temperature of the balance is continually monitored by several Resistance Temperature Detectors (RTDs), and when the temperature falls below the set temperature, the heaters are commanded to heat the balance. When the thruster is turned on, this heating cycle continues, and effectively compensates for the inability of the thruster to maintain the set temperature. This simple bang-bang control system ensures thermal stability of the entire thrust balance, meaning the flexures behave in a repeatable, consistent manner.

5.3.5 Independent Calibration Subsystems

Two critical attributes of any measurement instrument are the traceability and stability of the instrument's calibration. Traceability in this context means that a trusted calibration chain extends from the instrument to the relevant National Measurement Institute (NMI) standard reference. For thrust balances, this means producing a highly repeatable force which can be calibrated against a second traceable measurement system. By repeatedly monitoring the response of the balance to a consistent force, the stability of the calibration can be monitored. Carefully monitoring the calibration stability is crucial in electric propulsion testing, as subtle changes in the experimental setup can result in significant deviations to the measured force owing to the low levels of thrust produced by these devices. The IPPL and ESA balances each utilise two independent calibration subsystems. One serves to measure the balance sensitivity, and to ensure the traceability of the calibration to a NMI reference. The second monitors the stability of this calibration throughout the test campaign.

Voice Coil Actuator Calibration Subsystem

The thrust balance utilises an AVM12-6.4 voice coil actuator (VCA) manufactured by Akribis Systems to ensure the traceability of the calibration. The VCA consists of a coil section and a separate magnet section. By supplying a current to the coil section, a non-contact and zero friction force acts to repel the magnet section. Bespoke alignment aids have been designed to guarantee the two sections are separated precisely and without contact in a repeatable manner. The VCA, along with the alignment aids are shown in the separated and aligned state in figure 5.4. A Keysight B2902A precision sourcemeter is used to supply a very precise and ultra low noise current to the VCA coil, which in turn generates a highly stable and repeatable force on the VCA. The sourcemeter has an uncertainty of $\pm 1\,\mathrm{pA}$ which equates to a theoretical force resolution greater than 1 pN. This precise repeatability allows us to use the VCA as a link in our calibration chain to a trusted source. To accurately characterise the VCA, the magnet section was rested on the weighing pan of a Sartorius ENTRIS 124i-1S microbalance that has been calibrated to a NMI reference. The magnet is positioned such that any force supplied to it is directed downwards. Next the coil section of the VCA was aligned using the specially constructed VCA positioning system, such that the two sections are integrated in a repeatable way with no physical contact. The balance is tared to remove the mass of the magnet. The lack of contact between the two sections eliminates any offset due to friction. Next a wide range of current steps are supplied to the coil using the sourcemeter, and the force measured by the NMI-calibrated microbalance is recorded. This provides the user with a highly repeatable force versus supplied current relationship, which is traceable to the calibration certificate of the microbalance. The final step is to install the magnet section of the VCA to the moving platform of the thrust balance, and again align the coil section using the same bespoke VCA positioning system. The same sourcemeter is then used to generate the same forces, such that the sensitivity of the balance can be directly measured using a NMI-traceable force. An important consideration is to perform this step once the experiment is completely set up, and the balance has been brought

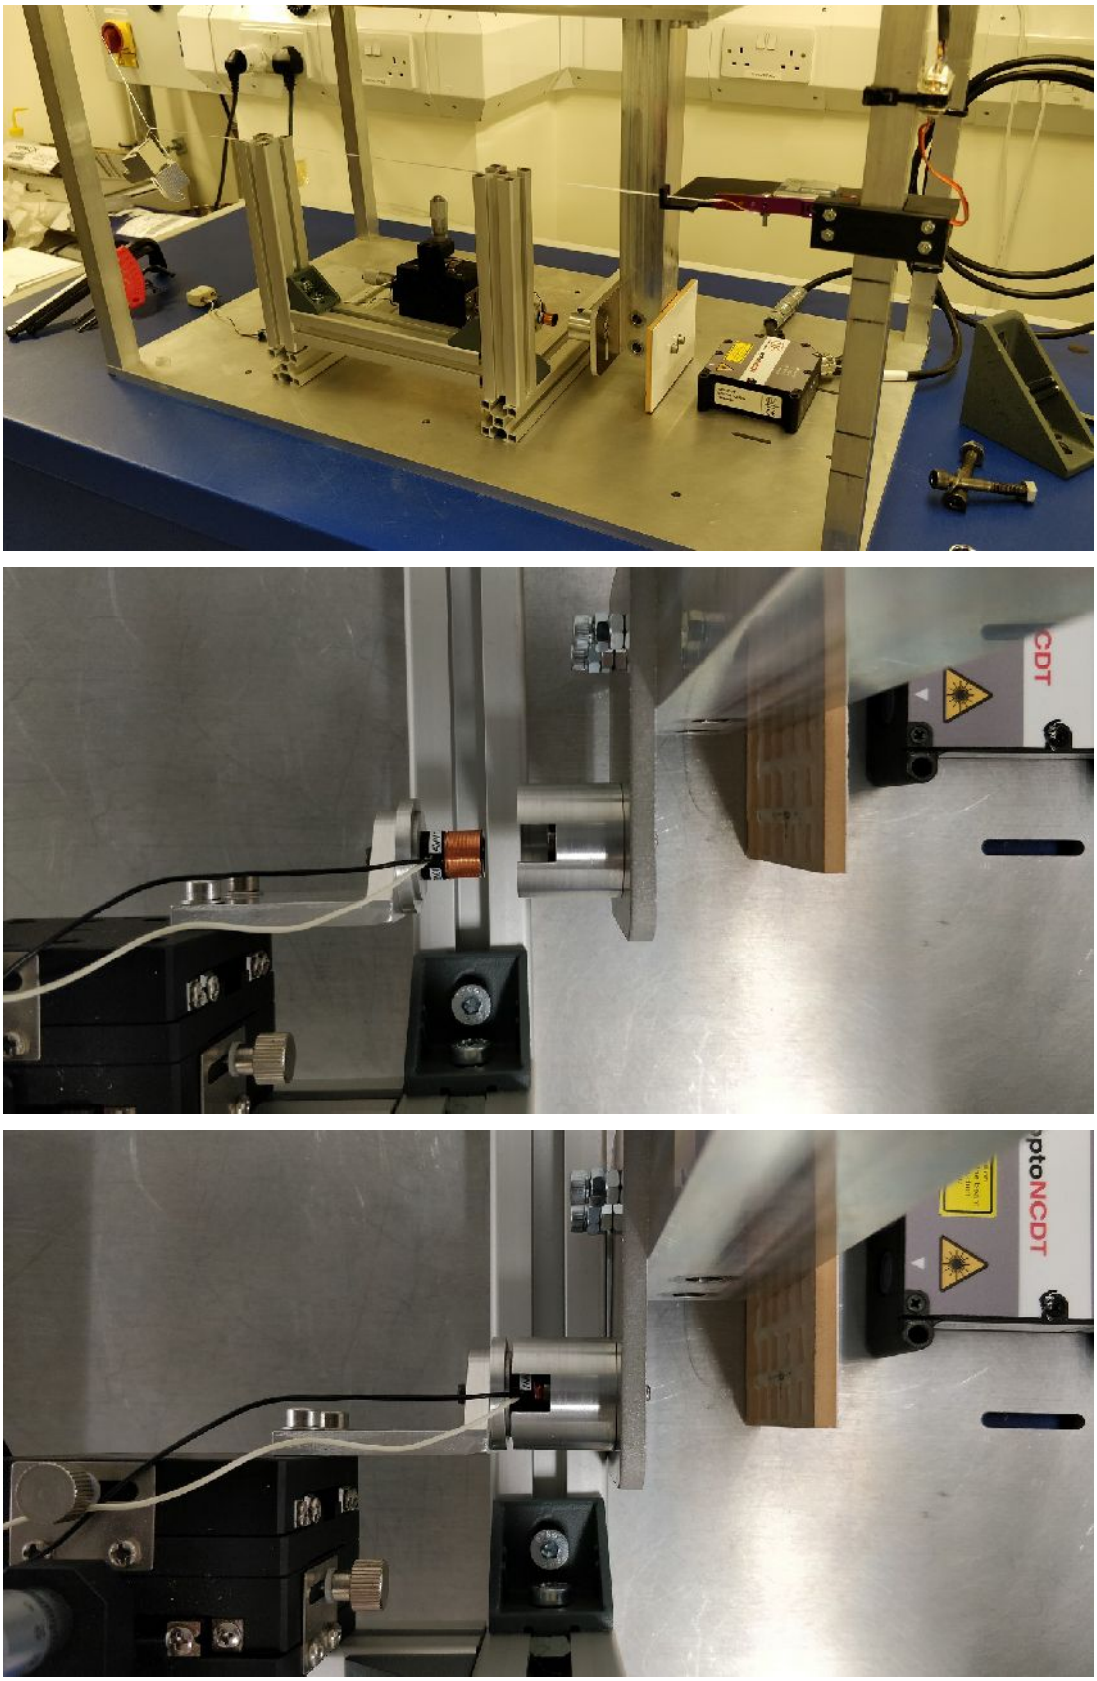

Figure 5.4: The VCA calibration subsystem (top) with close-up top view of VCA halves separated (centre) and VCA halves aligned (bottom).

up to operating temperature. This is because the calibration constant (balance sensitivity) may change even with very minor changes to the experimental setup.

Ideally the VCA would remain installed during testing, so that the same force can be repeated in vacuum to ensure the balance calibration remains stable. This is however not possible due to potential EMI produced by the electric thruster. Such interference can induce currents within the VCA coil and generate a force on the magnetic section indistinguishable from an actual thrust. For this reason the VCA is removed from the thrust balance before pumping down the facility, and a second calibration system is used to monitor the stability of the thrust balance calibration throughout the experimental campaign.

In-situ Calibration Subsystem

The primary role of the in-situ calibration system is to monitor the stability of the calibration by producing a repeatable force at intervals throughout the test campaign. The in-situ calibration system can be commanded remotely, and is not affected by EMI generated by an electric thruster, meaning that it does not need to be removed unlike the VCA calibration system. The subsystem consists of a servo motor fixed to the leg of the balance, a thread and a known test mass. The test mass is suspended via the thread from the moving central platform. The thread also runs horizontally from the mass to the servo motor via a guide hole in the servo holder. This is shown in figure 5.5. Rotating the servo arm pulls the horizontal string through the guide hole, and changes the angle of the thread between the test mass and the moving platform. This creates a force along the thrust axis between the moving central platform and the rigid leg.

The force generated by the in-situ calibration system (F_t) can be derived from the geometry of the system:

$$F_t(\alpha, R, a, b, h, H, m) = g_0 m \left(\frac{\sin(\gamma + \theta)\cos(\beta)}{\cos(\beta + \theta + \gamma)} - \tan\gamma \right), \tag{5.1}$$

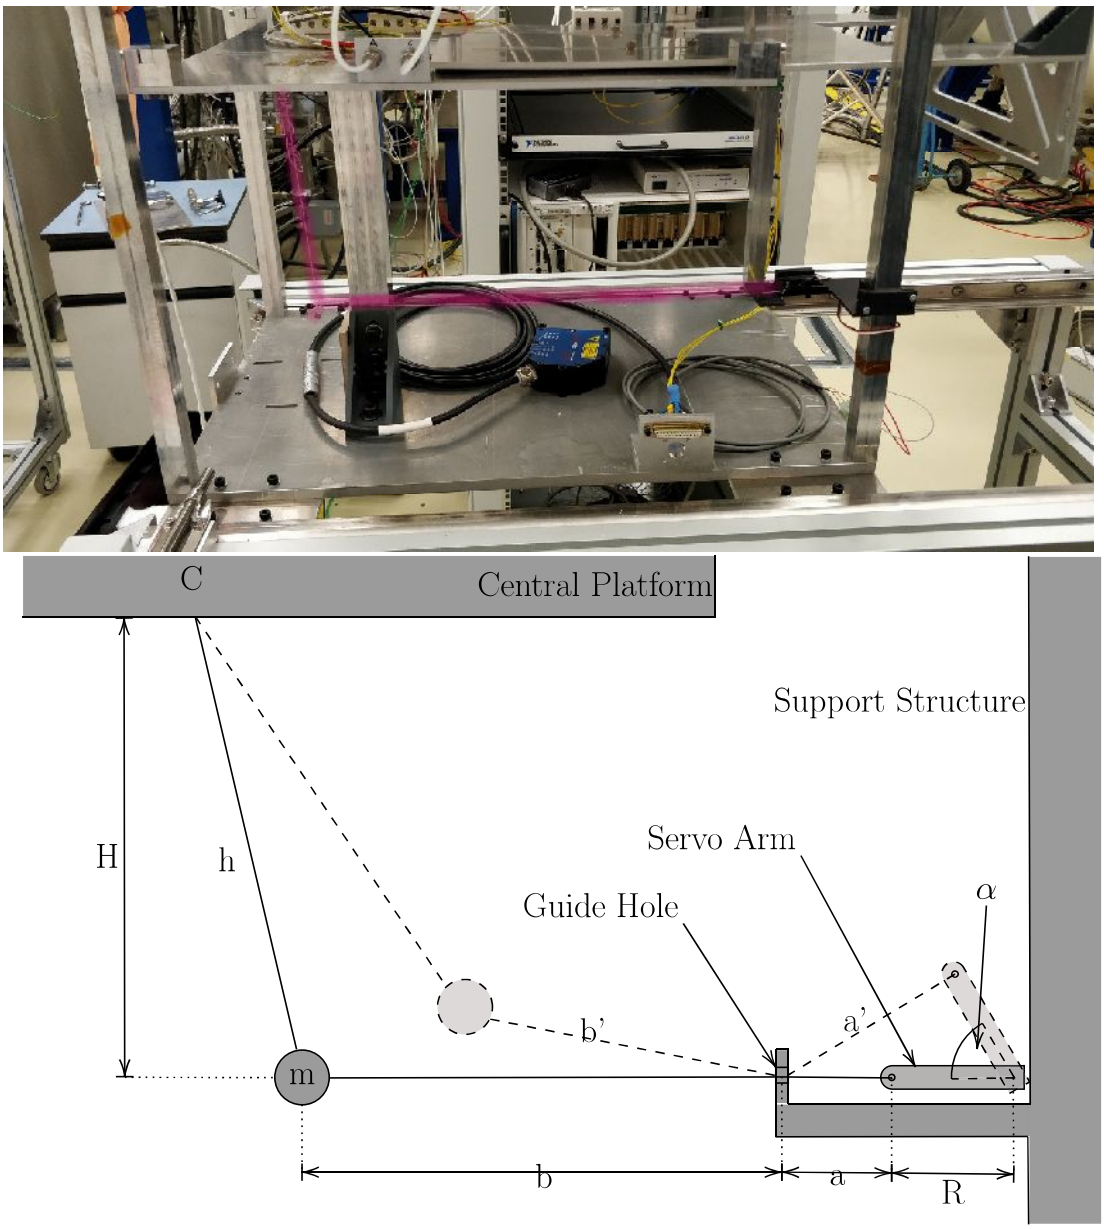

Figure 5.5: Side view of the in-situ calibration system with string highlighted in red (top). Diagram of the operational principle of the in-situ calibration subsystem (bottom). This diagram is not to scale. The string of length $l = a + b + h$ suspends the calibration mass of mass m from the point C on the moving central platform. The fixed servo arm is commanded to rotate through angle α to the dashed position. This pulls the string through the guide hole, and displaces the mass upwards and to the right as seen in the dashed position. This creates a force to the right, which is along the thrust axis on the central platform at point C. Adapted from Schwertheim et al. [107].

where g_0 is acceleration due to gravity. Angles γ, θ and β are functions of the measured values of α, R, a, b, h, H and m shown in figure 5.5. These angles are derived from first principles in section A.2 of the appendix. This force range can be tailored by substituting for a larger or smaller test mass.

A calibration consists of the servo motor being commanded to sweep through a range of angles, generating a range of forces with the response of the balance being recorded. The first calibration is done directly after the VCA calibration has been established in order to correlate both calibration methods. This is referred to as the baseline calibration as it overlaps with the traceable VCA calibration. The VCA can now be removed, and the chamber pumped down. During testing the in-situ calibration system is commanded to run spot checks in which the same commands are sent as during the baseline calibration. By comparing the laser response of a spot check to the response of the baseline, the stability of the calibration can be monitored. An identical response verifies that the balance behaviour has not changed since the NMI-traceable calibration. These spot checks are performed before and after any measurements are taken.

In summary, the sensitivity of the balance is measured by the VCA subsystem. The traceability of the measurement is ensured through the VCA calibration, whereas any deviation of the sensitivity throughout the experimental campaign is monitored using the in-situ calibration system. By correlating both methods simultaneously during the baseline, it is ensured that the two systems are in agreement.

5.3.6 The Electronic Support Box

Two Electronic Support Boxes (ESBs) were created to house the hardware needed for the control and measurement of the two thrust balances. Each ESB housed the necessary active and passive electronics to power, command, and monitor the thrust balance inside a 19 inch rack mountable

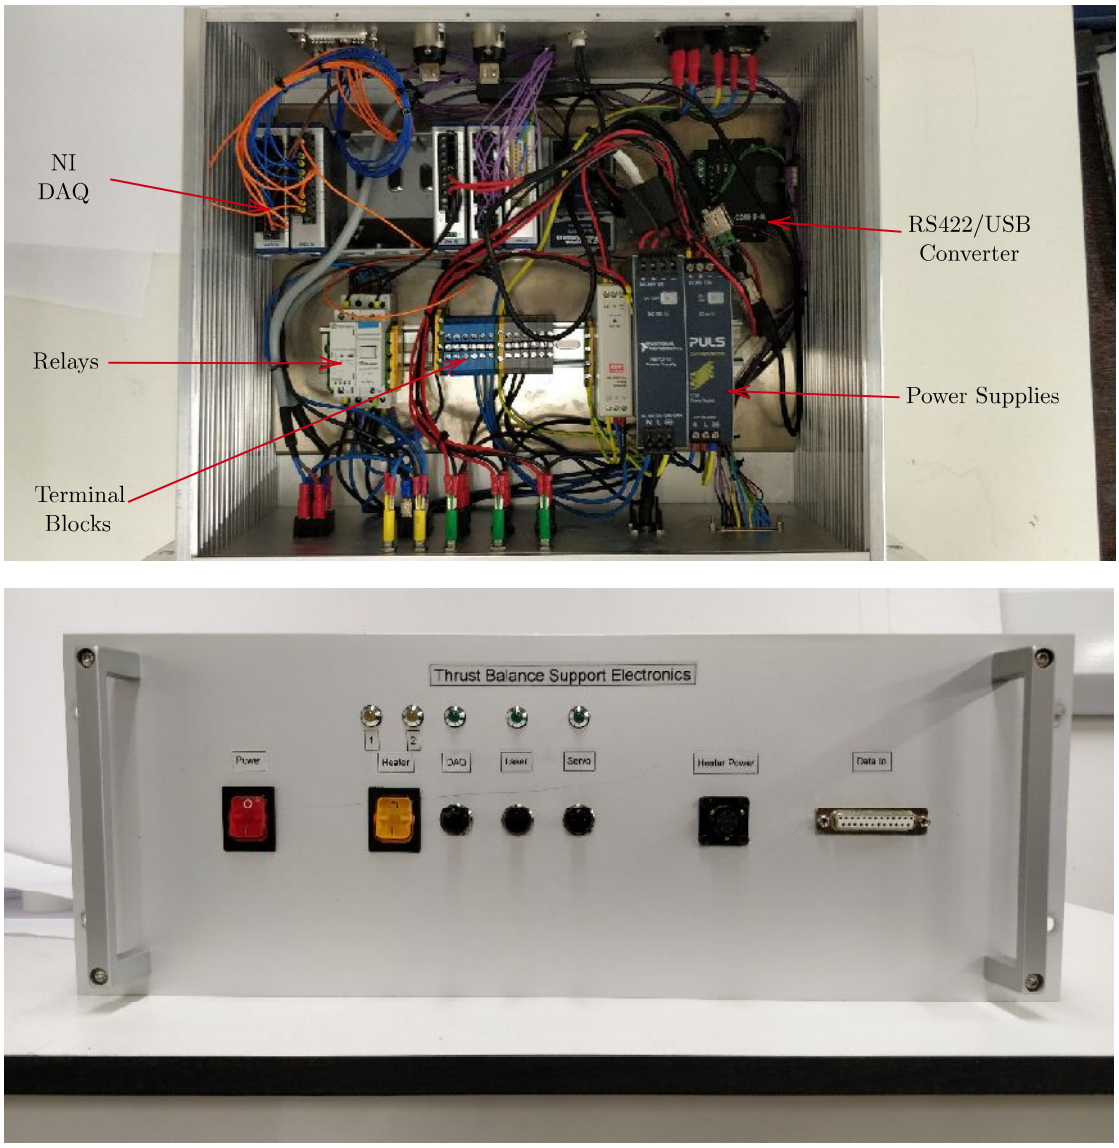

Figure 5.6: The inside (top) and front (bottom) of the IPPL version of the ESB used to command and monitor the thrust balance.

box. The ESB connects to mains power, to the balance via a standard D-SUB 25 cable, and to a PC via two USB cables. Rocker switches allow the user to safely monitor and power individual subsystems. The two ESBs were very similar, with the only major difference being that the ESA box was designed to be controlled using an external Data Acquisition System (DAQ), whereas the IPPL ESB contained an internal National instrument DAQ. The IPPL DAQ was later also modified such that it could be used to control and measure the output of two different laboratory power supplies (see section 5.4). The IPPL ESB is shown in figure 5.6.

5.3.7 Thrust Data Acquisition

Thrust is determined by comparing the displacement measured by the laser when the thruster is operating in steady state, to the displacement when the thruster is turned off. The thruster is brought to a steady operating point and held there until both the balance and the thruster settle. The laser sensor data acquisition is then initiated, followed by the thruster power and massflow being shut down. A representative sample of the raw displacement data is shown in red in figure 5.7. Here the thruster is in steady state, and turned off abruptly 30 seconds into the measurement. It can be seen that the balance oscillates un-damped at a frequency near 1.5 Hz throughout the measurement. However, an obvious and abrupt change of the equilibrium position can be observed when the thruster is turned off at t=30. The size of this step change is the measurement that must be quantified, as it represents a change in the net force on the balance: the thrust signal. Note that the absolute value of the displacement is not important, only the relative difference between the thruster equilibrium position in the on and off state. In post processing the size of this step change is quantified. Two linear fits are performed on the data: one for 30 seconds before the step, and one for 30 seconds after the step. These are shown as green and blue lines respectively in figure 5.7. The fitted trends are averaged over many periods of the large natural frequency oscillations to identify the equilibrium position of the balance in both thruster on and thruster off states. The step size is the relative difference between the two fits at the time of the step (here t=30 s). Any zero-drift that may have occurred is accounted for by measuring the zero value directly after the thruster is stopped. The change in displacement between the thruster on and off states is then converted to a force using the sensitivity parameter that was established during the VCA calibration procedure.

The raw data in figure 5.7 is shown to behave like an un-damped simple harmonic oscillator. This is a desired and expected feature which shows us that the balance is hanging freely with no friction. Oscillations at the natural frequency of the balance are generated from sources

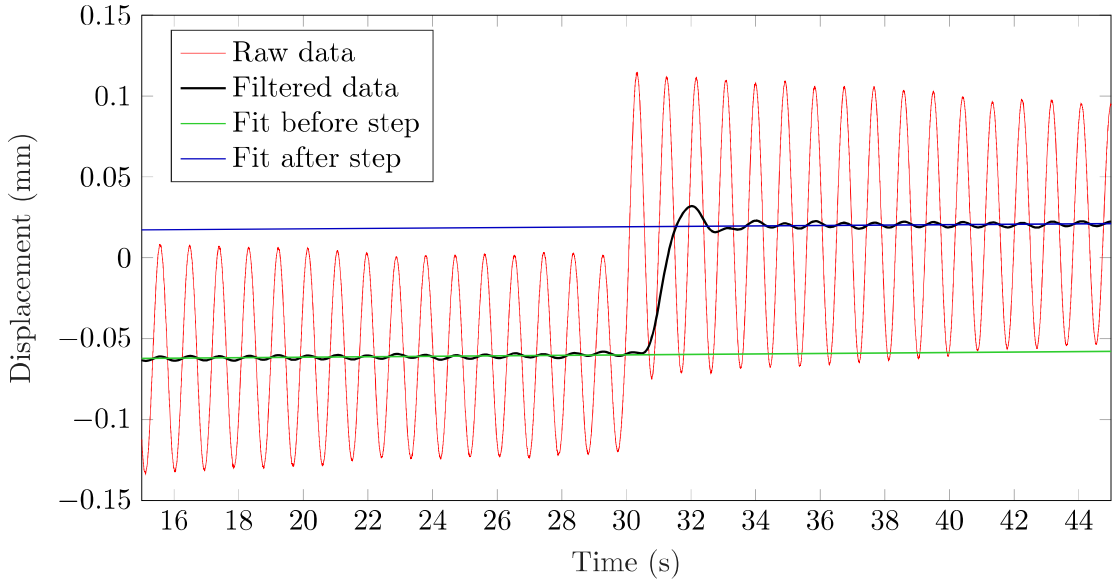

Figure 5.7: Displacement sensor output during a thrust signal of 52.5 mN. The red line shows the unfiltered signal, Butterworth filtering is shown in black. Adopted from Schwertheim et al.[107].

of excitation in the laboratory which are difficult to remove. The lack of friction or drag causes these oscillations to be highly under-damped. The natural frequency of the balance was experimentally found to be around 1.5 Hz. Similar balances utilise magnetic eddy dampers to remove such oscillations, yet here such dampers were omitted specifically due to concerns that the magnetic components could interact with the VCA calibration system.

In order to investigate the extent to which these natural frequency oscillations may override important low frequency features, a 5th-order low-pass Butterworth filter with a cutoff frequency of 0.5 Hz was applied to the data. Low-pass filters such as this have been used to smooth out raw data on similar balances[117]. The cutoff frequency of the filter was chosen to remove the natural frequency seen just below 0.5 Hz. The order was chosen based on trial and error to best reduce the oscillations. This Butterworth filter is applied to the raw data in figure 5.7, with the results shown on the same plot in black. This filter restricts high frequency, low amplitude oscillations, similar to an Eddy damper, however here it is accomplished during post-processing. The impact of this filter was assessed by performing the exact same post-processing procedure

on both the raw data and the Butterworth filtered data of a full set of VCA calibration forces ranging from 0.7 mN to 330 mN. The final step sizes for the same raw data were then compared to determine the impact that the filter has on the post-processing procedure. The mean difference between the step size of the two data sets was 0.32% of the measurement, with the maximum difference being 1.35%, which is negligible compared to other sources of uncertainty that are described in the following section. These result show that vibrations above 0.5 Hz have an almost undetectable impact on the final displacement measurement. By extension, the large natural frequency oscillations between 1 Hz-1.5 Hz are hereby also shown to have very little impact on the step measurement. This is considered an adequate justification for leaving the thrust balance under-damped.

5.3.8 Thrust Measurement Uncertainty Calculation

The uncertainty budget calculation for the thrust balance was based on the guidelines recommended in Polk et al.[108]. During the VCA calibration process the displacement x_i that each calibration force F_i produces is measured. The least squares approach is used to determine b_{cal} and S_{cal} such that the expected displacement $\hat{x_i}$ can be calculated from:

$$\hat{x_i} = S_{cal} F_i + b_{cal}, \tag{5.2}$$

where S_{cal} is the estimate of the sensitivity of the balance in m/N, and b_{cal} is the intercept to overcome the small but finite force needed to account for friction in the system. The force produced by the VCA is highly linear and has been measured to the highest precision possible with the microbalance: approximately ± 0.2 mN for individual measurements, as dicated by the datasheet of the microbalance. For the derivation that follows, it was assumed that the force produced by the VCA is known exactly. For the n calibration steps taken, the estimate the variance of the random disturbances of the balance s_x^2 can be estimated using equation 25 from

Polk et al.[108]:

$$s_x^2 = \frac{\sum_{i=1}^{n}(x_i - b_{cal} - S_{cal}F_i)^2}{n-2}. \tag{5.3}$$

It is assumed that this estimate is close to the actual value: $\sigma_x^2 \simeq s_x^2$. Taking the square root of the variance provides the standard deviation of the random disturbances of the balance to a known input s_x. This value is more tangible when converted into a force, which is done using the sensitivity to gain the estimate of standard deviation of the thrust measurement s_F, which again is assumed to be close to the actual value σ_{F_t}.

$$\sigma_{F_t} \simeq s_F = s_x / S_{cal} \tag{5.4}$$

The size of this standard deviation is one of the most important metrics for assessing the resolution of the balance. The standard deviation of the sensitivity itself s_{Scal} can also be estimated:

$$s_{Scal} = \frac{s_x}{\sqrt{\sum_{i=1}^{n}(F_i - \bar{F})^2}}, \tag{5.5}$$

where $\bar{F}$ is the average force used in the calibration. Again it is assumed that the actual value of the standard deviation of the sensitivity σ_{Scal} is similar to the estimate: $\sigma_{Scal} \simeq s_{Scal}$.

The overall uncertainty of the thrust measurement is the aggregate of the displacement measurement uncertainty and the uncertainty of the thrust balance sensitivity parameter. For a measured displacement x, the thrust measurement F_t has an uncertainty given by:

$$\frac{\sigma_{F_t}}{F_t} = \sqrt{\left(\frac{\sigma_x}{x}\right)^2 + \left(\frac{\sigma_{Scal}}{S_{cal}}\right)^2}. \tag{5.6}$$

An example of the intermediate results of this uncertainty budget are presented in the fol-

lowing section. The uncertainty in the thrust, I_{sp} and anode thrust efficiency are represented by the error bars in the results plots throughout this thesis. The size of these errors are calculated from equation 5.6.

Additional potential sources of uncertainty stemming from the experimental setup exist, such as the electrostatic interactions of the electrical leads, and the pressurisation of the gas feed lines. These factors were evaluated experimentally and found to be practically undetectable. A notable exception is the power going to the cathode heater. This requires a current in excess of 25 A, which adds a noticeable zero-drift to the measurement. However, performing in-situ calibration spot checks with the current on and off showed that the addition of the zero drift has little to no effect on the force measured.

5.3.9 Thrust Balance Performance

The performance of the balance was tested with great scrutiny during the interlaboratory validation campaign in collaboration with ESA, which is described below[107]. This was however a short test campaign which focused on a low level of thrust. For this reason the performance of the thrust balance over the WET-HET test campaign is also presented which spanned more measurements, and concerned a larger thrust range.

Interlaboratory Testing

The completion of the ESA contract required the verification and validation of both of the thrust balances in their respective laboratories. This required the assembly and testing of one of balances in the IPPL with an electric thruster, after which the instrument was shipped to the Netherlands such that the same experiment could be repeated during a two week installation and validation activity. The thruster used was a version of a Quad Confinement Thruster (QCT) referred to as the QCT Phoenix[118;119;120]. Here the thrust balance performance is the main focus

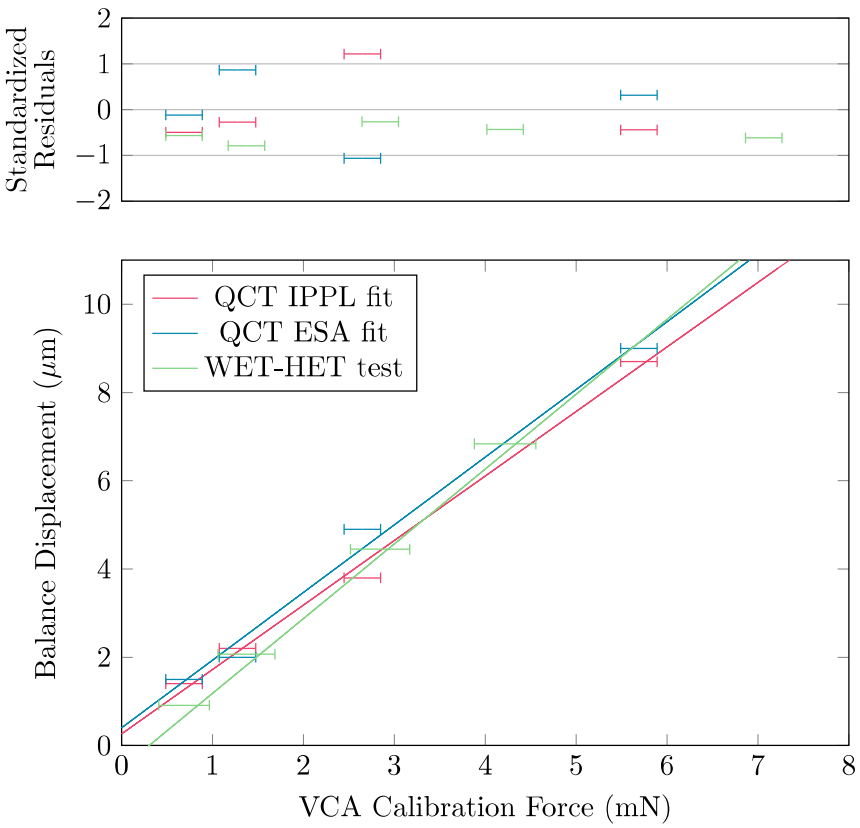

Figure 5.8: Thrust balance sensitivity at IPPL and ESA measured using the VCA calibration system. Note the error bars represent the uncertainty of the VCA force measured with our microbalance.

of the campaign, but a full description of the QCT Phoenix and the thruster characterisation undertaken can be found in Schwertheim et al.[107] and in Azevedo et al.[120]. Most importantly, the EMI and heat produced by the thruster did not effect the performance of the thrust balance.

The response of both balances to the VCA calibration force during the interlaboratory testing campaign is shown in figure 5.8. A linear fit is applied to the data, with the gradient being the balance sensitivity S_{cal} as described in equation 5.2. Four calibration forces (n=4) were used ranging from 0.687 mN to 5.690 mN to cover the thrust range of the QCT Phoenix. A greater number of calibration forces, spanning a greater range where used for WET-HET experiments, which are included on the plot also. These calibration forces were used to perform the error calculations detailed in section 5.3.8. For completeness the intermediate results of the uncertainty

Table 5.1: The intermittent variables used in two different laboratories during the validation campaign. For comparison we also include a calibration for a typical WET-HET test.

Variable	Units	IPPL QCT Test	EPL QCT Test	IPPL WET-HET Test
b_{cal}	μm	0.2603	0.3984	-0.0521
S_{cal}	$\mu m/mN$	1.462	1.535	1.634
s_{Scal}	$\mu m/mN$	0.0707	0.1060	0.0022
s_x^2	μm^2	0.0747	0.1682	0.1568
s_x	μm	0.2734	0.4101	0.3960
s_F	mN	0.1870	0.2673	0.2424
n	-	4	4	13
$\bar{F}$	mN	2.575	2.575	41.526
R	mm	16.8±2	18.0±2	54.9±2
a	mm	23.5±2	22.5±2	25.9±2
b	mm	280±2	345±2	435±2
h	mm	179±2	170±2	228±2
H	mm	162±2	169.8±2	225±2
m	g	2.195±0.2	2.195±0.2	14.74±0.2
σ_α	deg	1	1	1

calculations are included in Table 5.1, with the results of equation 5.6 dictating the size of error bars used in the results section. Also included in Table 5.1 are the values of these intermediate values for a typical WET-HET test in the IPPL. A similar sensitivity (S_{cal}) for the WET-HET test and the QCT test can be seen although the WET-HET test campaign took place almost 2 years later, with an entirely different test setup. Given the WET-HET produces a higher level of thrust, a greater range of calibration forces were used: n=13, which reduces the estimate of the standard deviation of the sensitivity (S_{cal}) as per the sum in equation 5.5.

The resulting sensitivities for the validation campaign were S_{cal}=1.462±0.071 μm/mN for the IPPL tests and S_{cal}=1.535± 0.106 μm/mN for the EPL tests corresponding to the gradients of the two trends in figure 5.8. Note that both trends are linear yet with slightly different intercepts and gradients. This is expected given that each set-up has a slightly different mass and configuration. The standardised residuals shown in the top of figure 5.8 are given by:

$$R_s = (\hat{x}_i - x_i)/s_x, \tag{5.7}$$

and allow one to assess the quality of the fit. These residuals tell us how far each value is from the expected value. As a rule of thumb, standardised residuals distributed uniformly between -2 and 2 are considered indicative of a good fit. The linearity of the balance in this configuration was tested over a greater thrust range than expected for the QCT characterisation in order to assess the full capabilities of the measurement instrument.

The estimate of the standard deviation of the thrust measurement ultimately provides the resolution of the balance. For the interlaboratory test campaign these values were s_F=0.1870 mN for the IPPL and s_F=0.2673 mN for the EPL data. This resolution degrades when the thrust range is extended up to 100 mN to s_F=2.269 mN for the IPPL and s_F=3.080 mN for the EPL data.

Potential errors resulting from the electrostatic interaction of the thruster supply lines on the thrust balance measurements were investigated experimentally. These tests were performed within the IPPL but were not repeated at EPL. The gas supply to the thruster and cathode were closed throughout these tests, so the physical thrust produced should be zero. The maximum possible voltages were applied to both the keeper and the anode, being 650 V and 100 V respectively. The thrust measurement procedure followed the same protocol as for the actual thrust measurements, and it was found that the maximum measured force was 0.014±0.19 mN. This is small enough when compared to other sources of uncertainty that it can be disregarded. As discussed above, the zero-drift caused by the cathode current has no net affect on thrust.

The stability of the calibration was monitored by running spot checks with the in-situ calibration subsystem throughout the test and comparing the response to the baseline measurement. This baseline was attained in ambient conditions with a heated balance directly after the VCA calibration was performed, and therefore acts to link the traceability from one calibration system

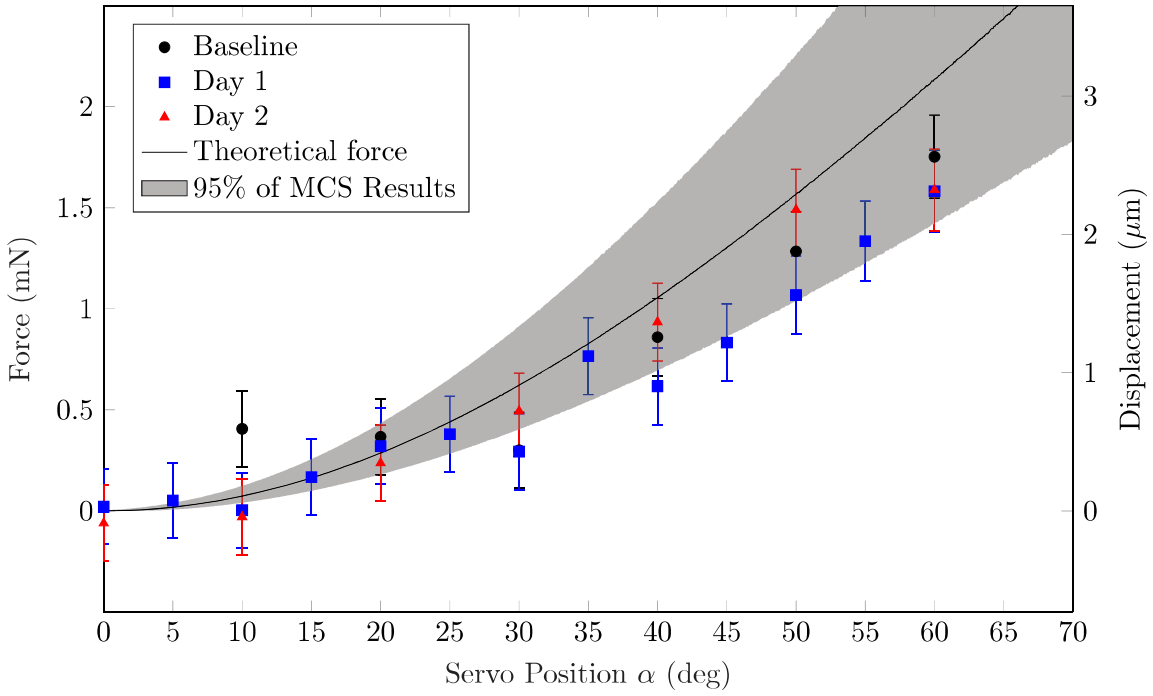

Figure 5.9: The calibration stability monitored over two days of testing in the IPPL.

to the other. The spot checks on two days of testing at the IPPL are compared to the baseline measurement in figure 5.9. The results at the EPL were similar, and have thus been omitted. A very similar trends across the three datasets can be seen, with the exception of one outlier at 10 degrees during the baseline measurement. It can be concluded that the calibration drift was minimal over the three tests, indicating that the balance was not altered significantly over the duration of the campaign. The baseline, which was performed in ambient conditions, is in good agreement with both the spot checks which were performed in vacuum. This confirms that the calibration traceability is not compromised by the evacuation of the chamber.

The errors in measuring the variables α, R, a, b, h, H, and m compounded to an uncertainty in the theoretical force. To quantify this uncertainty a Monte Carlo simulation of the subsystem was performed to accurately propagate the error for all measured values. The measured values and conservative error estimates are provided in table 5.1. For each run the value of α, R, a, b,

h, H, and m was picked from a normal distribution about the measured value, with the error being one standard deviation. The theoretical force was calculated for 100,000 simulations at 1000 different servo positions. The grey shaded area in figure 5.9 represents the band in which the median 95% of these simulations fell. The majority of the measured data are seen to fall below the theoretical force, yet within the 95% median band. Experimental error in measuring the seven variables which describe the system likely account for the systematic under-prediction of the force compared to the calculated value. Of importance is that the experimentally observed trend line lies within the error bound predicted by the Monte Carlo analysis when considering the relative uncertainty of each measured parameter. Even with the small systematic discrepancy between the measured values and the theoretical force seen in figure 5.9, we stress the significance of this result: the measured forces were derived from the displacement as determined by the laser, and converted into a force based on the sensitivity calculated using the VCA subsystem. The theoretical forces on the other hand, were derived purely from the geometry of the in-situ calibration system. Any large disagreement between measured values and theoretical values would indicate systematic errors. The mean difference between the theoretical and the measured value is 0.1878 mN, with a standard deviation of 0.2329 mN. The good agreement we see in figure 5.9 is a direct indication of the accuracy of the balance, and consistency of the two separate calibration approaches.

Balance Performance during WET-HET Testing

Testing for the WET-HET followed the same procedure as for the QCT Phoenix measurements. Where the thrust of the QCT was measured to range from 0.2-2.2 mN, the WET-HET operated in the range of 2-40 mN. Given the WET-HET produces a higher thrust, more calibration forces and a greater range were used. Each time the thrust balance or thruster was interacted with, manipulated or moved, a new VCA calibration was performed. This resulted in a total of 21 calibrations over the 18 month test campaign. The calibration forces for five of these calibrations

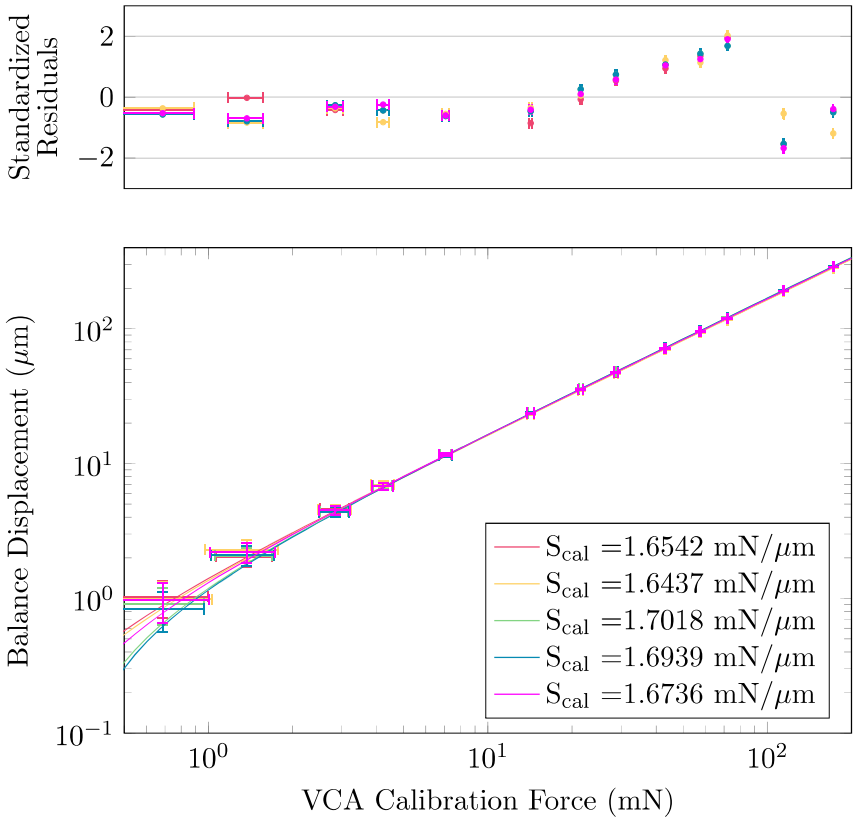

Figure 5.10: Five different VCA calibrations during the WET-HET campaign.

are shown in figure 5.10. The green data are those shown in table 5.1. All trends are shown to be extremely linear over the range surveyed. The agreement of the data on different days is remarkable given that the experiment was completely disassembled and rebuilt several times during this period. Note that a significant change to the sensitivity would indicate that the experimental set-up had been changed, but this would not jeopardise the accuracy of the thrust measurement, given that each thrust measurement only relies on the most recent calibration at that time.

Note that while all the standardised residuals in figure 5.10 fall within the desired range of

-2 and 2, for the calibration forces between 21 mN and 72 mN a positive departure from the expected value of 0 can be observed. Equation 5.7 shows that this indicates the measured displacement due to the calibration force in this range is consistently greater than that what would be expected of a perfectly linear response. The tight grouping of these residuals – on widely different days of testing – strongly suggest this is due to a systematic bias. Let us consider the absolute size of this bias: from equation 5.7 we see a standardised residual of magnitude 2, is the results of a calibration force being 2 standard deviations away from the expected results. For these five calibrations, the average estimate of the standard deviation is 0.39 μm. So a standard residual of magnitude 2 results in a force which is 0.23 mN greater than what is expected of a perfectly linear response. The size of the systematic bias in the range of 21 mN to 72 mN is considered to be notable but ultimately negligible.

5.4 Laboratory Set-up

The complete set-up of the WET-HET characterisation experiment required the cathode and thruster be mounted on the thrust balance within the Boltzmann facility, along with all the necessary electrical and fluidic connections. The WET-HET, cathode and thrust balance are shown in figure 5.11, with a complete schematic shown in figure 5.12.

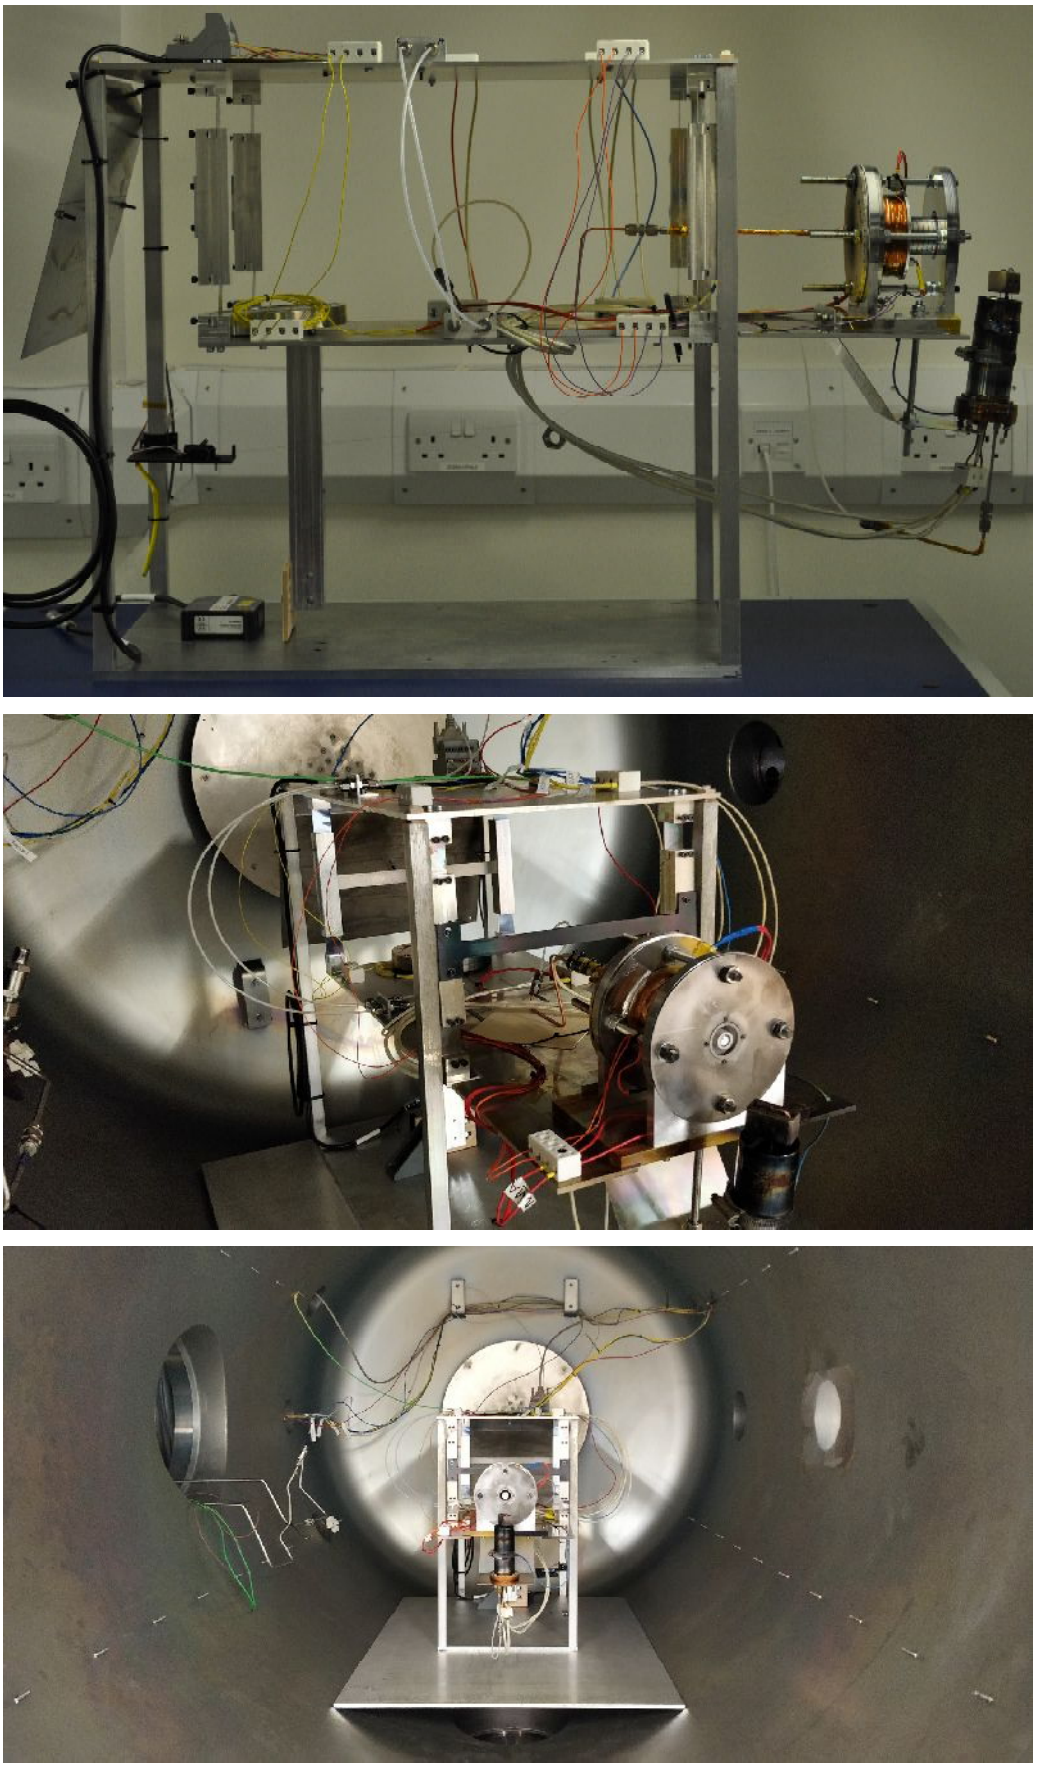

Figure 5.11: The WET-HET and cathode mounted on the thrust balance outside of the facility (top) and inside the facility (centre and bottom)

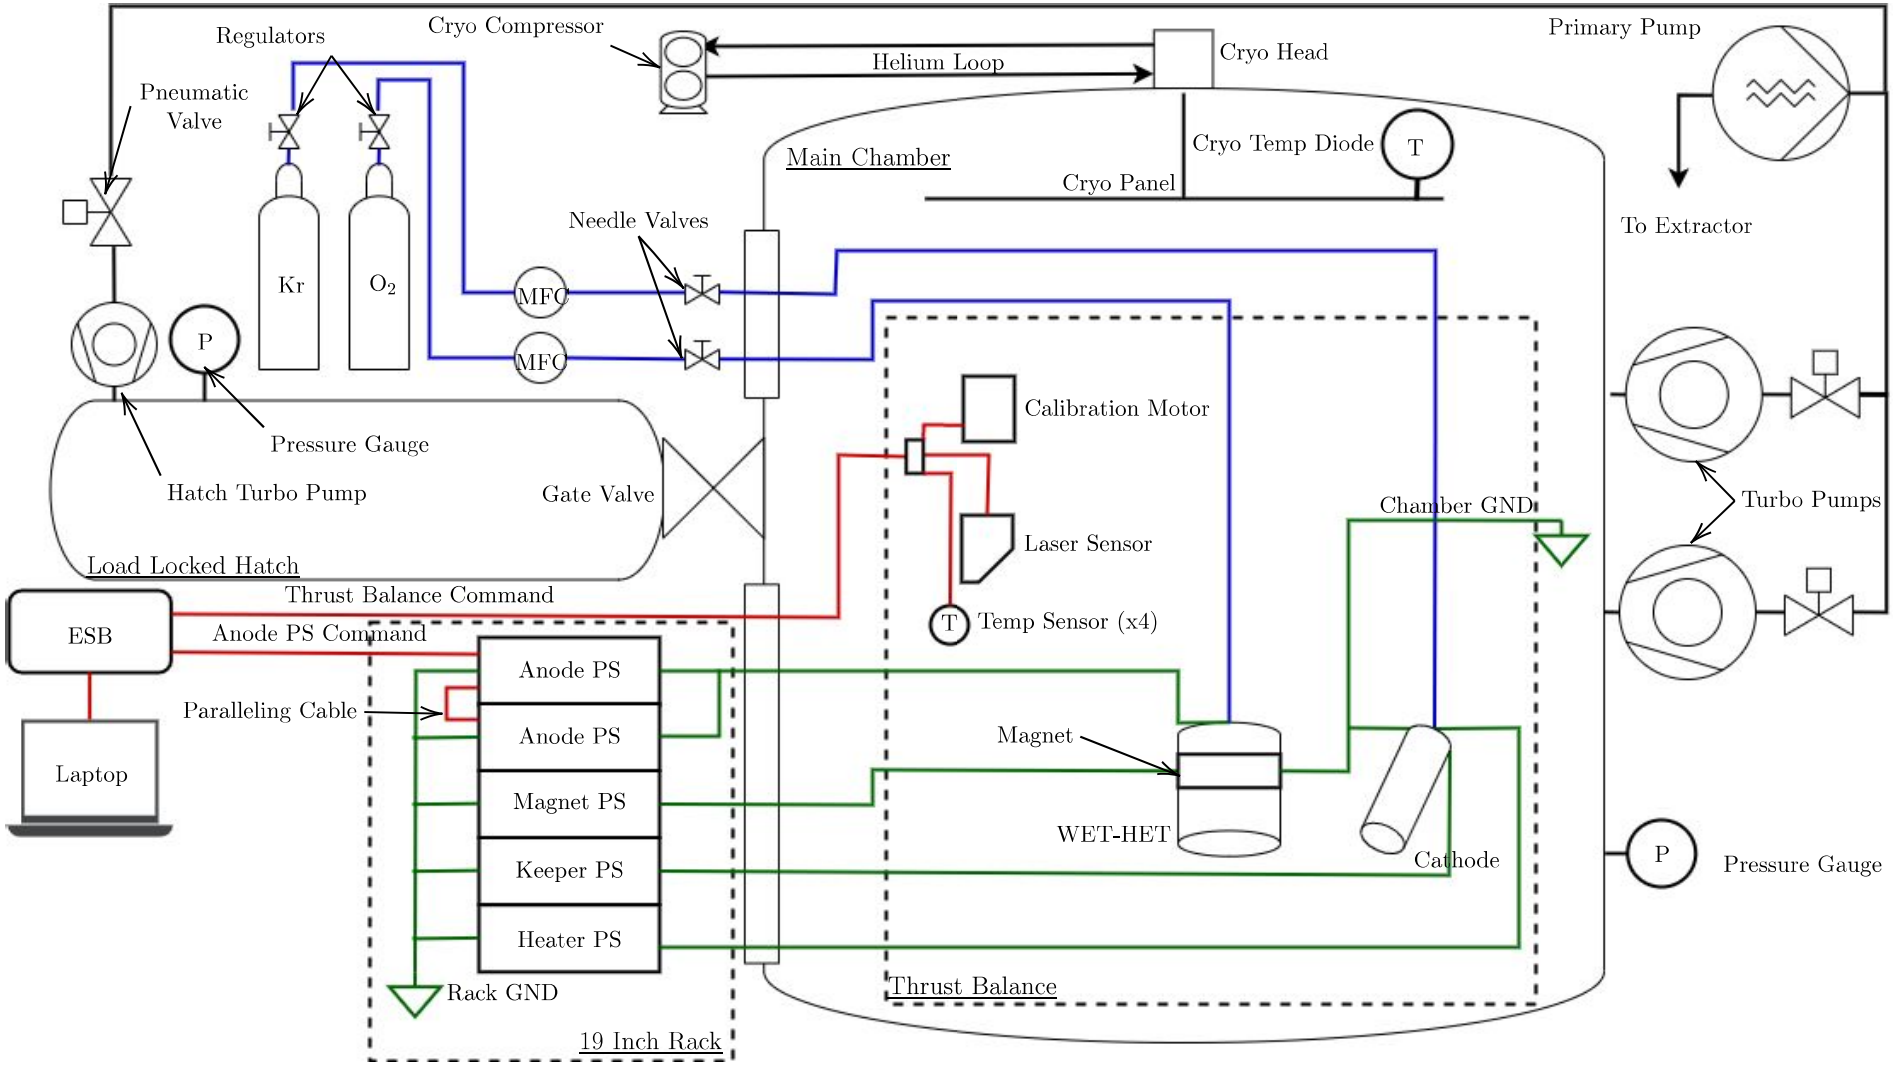

Figure 5.12: The complete experimental set-up showing facility, thrust balance, WET-HET, cathode and support equipment. Fludic lines are shown in blue. Data and command lines are shown in red. Electrical power lines are shown in green.

All electrical leads and fluidic lines pass from the WET-HET and cathode on the central thrust balance platform to the top level of the thrust balance before being routed to the feed-through. A single bespoke feed-through was built in house such that the four electrical leads and two fluidic lines could pass out of the facility. A single 19 inch server rack was used to house all necessary power supplies. The power and command for the laser sensor and in-situ calibration motor, as well as four RTD temperature sensors are bundled into a single D-SUB 25 cable. This cable passes to the ESB via an off-the-shelf dedicated D-SUB 25 feed-through. A second D-SUB 25 cable allows the ESB to monitor and command the anode power supply as described in the section 5.5. Finally the ESB communicates with a laptop via a USB cable. Not included in figure 5.12 is the sourcemeter and VCA calibration subsystem which was removed before facility evacuation and testing.

The gate valve was fixed open for the duration of the test campaign, although no equipment was housed in the hatch, with the exception of the hydrogen experiments (see section 5.6). This allowed the hatch turbo pump to assist the two main turbo pumps in evacuating the main facility.

The thrust balance was calibrated directly before closing the facility, meaning this was done after the experiment has been setup and all electrical checks had been performed.
The WET-HET thrust measurement procedure is explained in detail in section A.4 of the appendix.

5.4.1 Fluidic Set-up

The set-up shown in figure 5.12 was a typical experiment in which oxygen and krypton were supplied to the thruster and cathode respectively from cylinders. The 1/8th inch Swagelok system was used for all fluidic lines. The regulators reduced the pressure from the cylinder pressure to 2 bar, such that the mass flow controllers (MFCs) were able to precisely regulate the

flow. The MFCs used for the anode and cathode were both Bronkhorst El-flow Select devices which have a maximum mass flow rate of 200 sccm xenon and 50 sccm xenon respectively. These were calibrated to have an accuracy better than ±1% over the range used. Downstream of the MFCs the gases pass through hand operated needle valves, before entering the facility via a feed-through. Within the facility, the propellants flow through flexible vinyl tubing to the top of the thrust balance. From there wide arcs of the same flexible tubing transfer the gas to the central platform in such a way that thermal expansion of these lines would not register as a thrust measurement. Finally they are connected to the thruster and cathode respectively, with special care taken to isolate them from the high temperature of the two devices.

5.4.2 Thrust Balance Set-up

Practical experiments had shown that actively heating the thrust balance as described in section 5.3.4 added considerable wait times to the procedure, without dramatically improving the confidence in thrust measurements. It was decided that for the WET-HET experiments the temperature of the balance would not be controlled, as any impact that thermal fluctuations would have on thrust measurements would be captured by the in-situ calibration system.

A single D-SUB 25 cable connects the electronic support box of the thrust balance to the device itself, via a commercial D-SUB 25 feed-through. Due to this cable being in close vicinity to the thruster and cathode, strong EMI from the thruster had the potential to induce very high voltage transients to travel from the experiment out of the facility. These transients had the potential to harm the electronics within the ESB, but also the test engineers. A D-SUB 25 shunt was constructed in which each pin was connected to ground via a Metal Oxide Varistor (MOV). Any voltage above 30 V applied to any of the pins would short that pin to ground, dissipating the transients before it was able to do harm.

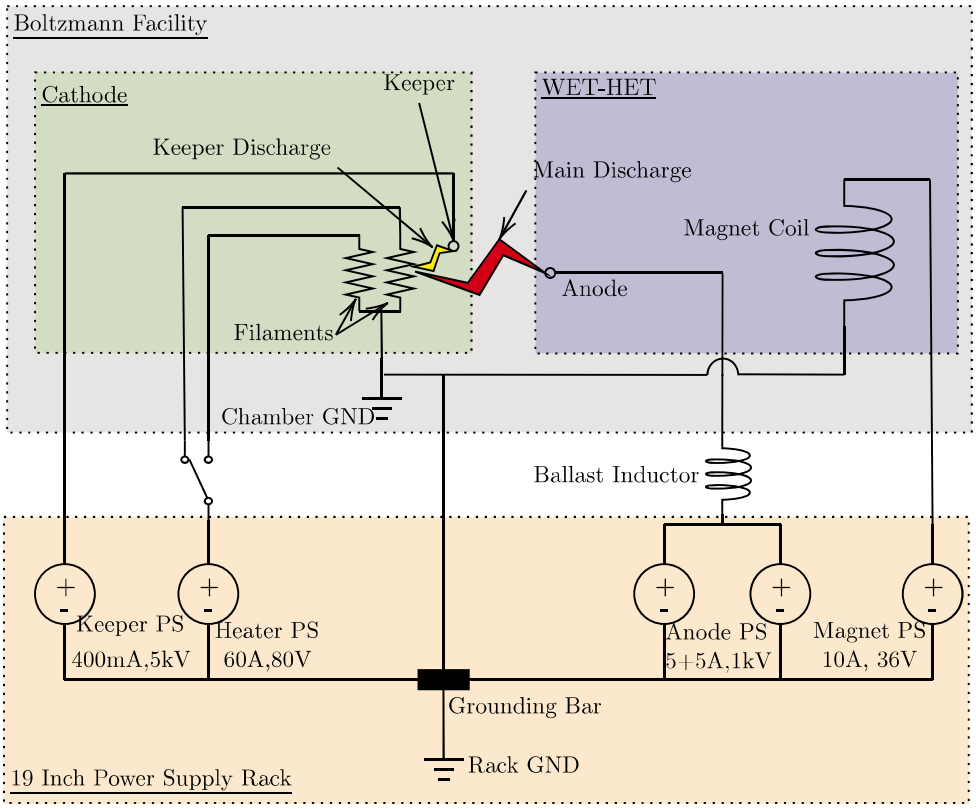

Figure 5.13: The main electrical set-up, showing how five power supplies (PS) which power the cathode and WET-HET. The keeper discharge and main discharges, which close the circuits are shown in yellow and red respectively.

The DAQ and the RS422 converter within the ESB were connected to a laptop via USB. A 1000 VDC USB optical isolator was also used to mitigate the possibility of high voltage transients to travel from the ESB to the laptop via the USB cables.

5.4.3 Electrical Set-up

The circuit diagram of the complete electrical set-up of the WET-HET test campaign is shown in figure 5.13. The thruster and cathode were operated by five power supplies: one each for the keeper, heater, and magnet circuit, with two power supplies for the anode. All five instruments were housed in a single 19 inch rack shown in figure 5.14. Further information about these power supplies is provided in section A.3 of the appendix.

Figure 5.14: The power supply rack showing from top to bottom: two anode power supplies, the keeper power supply, heater power supply and magnet power supply. Please see text body for model numbers and specifications.

Grounding

For safety reasons the Boltzmann facility is permanently electrically grounded. The 19 inch instrument rack which powers all of the above power supplies is powered and grounded via a three phase cable. Very special care was taken to eliminate grounding loops wherever possible. This was achieved by grounding each power supply via a single common grounding bar, which was itself connected to both the ground of the three phase power, and the Boltzmann facility. The chamber ground point, rack ground point and grounding bar are included in figure 5.13. Please note that this was not the ideal scenario: connecting the facility to the grounding bar effectively linked two (potentially independent) grounds together. The grounds could potentially differ by a slight voltage offset, or noise from the other EMI sources within the building. This would cause currents or noise to travel throughout the experiment, potentially disturbing the sensitive plasma discharges.

Ideally only a single ground would be used as a reference. This was however unachievable, as neither the facility nor the rack could be safely isolated from its own ground. We can see from figure 5.13 the heater current, keeper current, discharge current and magnet current can all return to ground either through the grounding bar, or a dedicated lead which links the rack ground and the chamber ground. Different grounding configurations were tested, and the configuration shown in figure 5.13 proved to be the most successful, with others proving more difficult to ignite.

Ballast Inductor

For certain configurations the WET-HET plasma was particularly difficult to ignite and sustain. One of the methods to "stabilise" the discharge was by adding a ballast inductor in series with the anode as seen in figure 5.13. This ballast was an entire spool of wire with a large cross section. Once ignited, the magnetic field induced by the coil opposes a sudden drop in current.

This was experimentally found to help to stabilise the discharge somewhat, and was used for the majority of tests. At high currents, the voltage drop across the coil was measured and found to be negligible, meaning the power draw of the ballast can be ignored.

5.5 Thruster Control Software

High discharge powers and high magnetic field strengths were found to cause thermal issues for the WET-HET. After several seconds of operation in these situations, the internal ceramic would start to glow, followed closely by the discharge voltage dropping. This is discussed in further detail in section 6.2 of the following chapter.

When this thermal issue was identified the thruster had always been manually controlled: once the thruster was ignited, the mass flow rate would be adjusted, followed by the magnetic field being applied. The anode power supply current setting would then be manually adjusted, and the resulting discharge power calculated using a handheld calculator. The current would then be further adjusted until the designated power was achieved. This generally took 10-30 seconds after the magnetic field was applied, which was often enough time for the thruster to overheat as described above.

It was noted that the thermal effects developed over several seconds of firing, by which time the thruster had reached stable state, and the thrust balance was able to accurately take a measurement. Reaching the target power manually was taking too long, and this time delay was making thrust measurements difficult. Automating the anode power supply was proposed as a solution such that the thruster could reach its target power almost instantly. Once at the target power, the thrust measurement could be taken and the discharge extinguished before thermal issues had a chance to arise. We call this software the Thruster Control System of the WET-HET.

A typical D-SUB 25 cable was used to connect the anode power supply device to the IPPL balance ESB, where the NI-9263 and NI-9205 modules of the internal DAQ commanded and monitored the 0-10 V analog output and input signals respectively. From here a dedicated LabVIEW program was used to dictate voltage and current limits, and measure the actual supplied values.

On initiation the software queries the user for a target power and some safety thresholds. The basic configuration of the Thruster Control software is a state machine which monitors the anode power supply and the user to decide which state to change to. The software also provides useful feedback to the user in real time, as well as data logging.

5.6 Hydrogen Test

At the very core of the water electrolysis Hall effect thruster concept is the capability of the cathode which accompanies the thruster to operate on hydrogen at precisely 1/8th the anode mass flow rate. The health and safety officers of the IPPL did not allow experiment which required simultaneous use of compressed oxygen cylinders in proximity of compressed hydrogen cylinders, given that the laboratory was not designed to accommodate such hazardous gases. These concerns are not unwarranted. Hydrogen is the gas of the lowest atomic mass, and therefore leaks through even the smallest cracks, making leaks extremely difficult to prevent. The NASA Safety Standard for Hydrogen and Hydrogen Systems states that the explosive limits of hydrogen/oxygen at standard conditions is between 4% and 96% by volume[121]. This is a much larger flammability window than most flammable gases, and therefore means that hydrogen can ignite in unintuitively low concentration. The combination of these facts suggest that the odourless gas could easily leak during a test, and be ignited by something as seemingly harmless as a ceiling light. By using krypton as the operating gas of the cathode, the feasibility of the WET-HET was still able to be determined, with a presumably small impact on performance.

Given the underlying importance of the assumption that the cathode operating gas would not dramatically influence thrust measurements, it was crucial that a method be found to test this assumption. Charlie Muir was a PhD student at the time who was conducting experiments of a chemical micro-propulsion system utilising hydrogen and oxygen within the Boltzmann facility[45]. Charlie had developed hardware, software and test procedures which allowed him to effectively supply his thruster with hydrogen and oxygen safely. Charlie's experiment involves storing small quantities of hydrogen in special plena which were themselves stored within the evacuated Boltzmann facility. Remotely operated solenoid valves allowed the gas to pass from the plena to the experiment via a MFC, all without leaving the evacuated chamber. Special care was taken during the system design such that the failure or leaking of any component will cause hydrogen to expand into the evacuated facility, but not the laboratory. This hardware was modified to test the WET-HET operating on oxygen, with the cathode operating on hydrogen.

We refer to the experiments in which hydrogen is supplied to the cathode as the hydrogen experiments. For these experiments the WET-HET, cathode and thrust balance were set-up in the same way as for regular experiments, with the only difference being the fluidic set-up of the cathode. This alternative configuration is shown in figure 5.15, with a pipe and instrumentation diagram shown in figure 5.16.

The hydrogen cylinder was stored outside of the building, and only brought into the laboratory to fill the two plena, before being removed again. Each of the two plenum had a volume of 2 litres, and was filled with 2 barA of hydrogen. The entire system including the plena have been pressure tested to 85 barA, and were protected by pressure release valves which purge into the facility. Check valves ensure that the plena cannot release hydrogen into the laboratory, even if the external manual valve were to fail. All solenoid valves are nominally closed, and operated

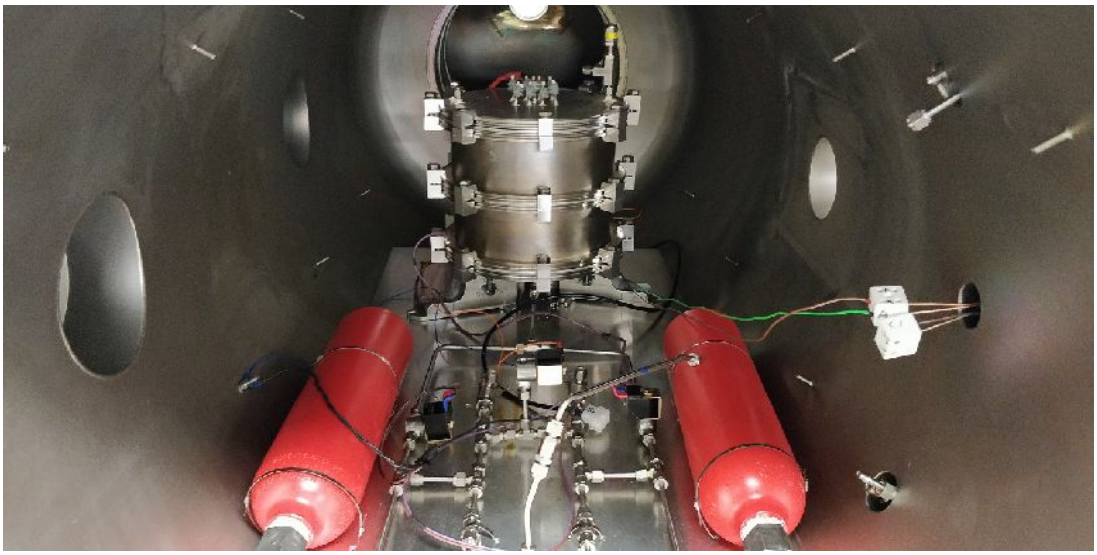

Figure 5.15: The hydrogen supply hardware. This system comprises two plena, several solenoid valves, and an MFC housed within a vessel pressurised to atmospheric pressure. The entire system was housed within the hatch of the Boltzmann Facility.

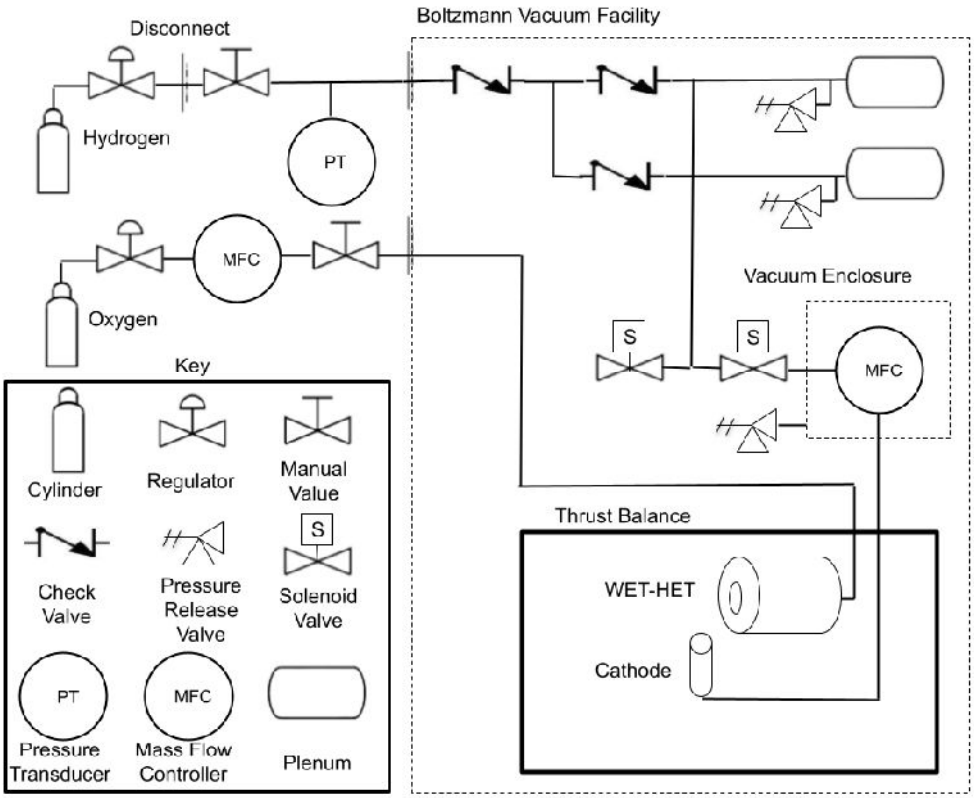

Figure 5.16: The pipe and instrumentation diagram of the hydrogen tests.

by a LabVIEW program. The MFCs were unable to function in a vacuum. This necessitated a pressure vessel which was kept at atmospheric pressure and situated inside the vacuum facility. This vessel houses the hydrogen MFC, shown labelled as "Vacuum Enclosure" in figure 5.16. Electrical feedthroughs allowed this MFC to be controlled remotely.

Unlike the experiments conducted by Charlie, we required that the primary, turbo and cryo pumps all be operated while hydrogen and oxygen were released into the facility at a ratio of 1:8 by mass – a highly flammable ratio. Special care was taken to ensure the facility pumps could not provide a source of ignition. The cryo panel was of no concern: at 40 K the cryo panel solidifies oxygen, but does not affect hydrogen, which was instead pumped by the turbo pumps. NASA reports that the lowest pressure that any hydrogen/oxygen mixture has been shown to ignite is 0.57 mbar using a very high energy ignition source[121]. Calculations show that even at the maximum mass flow rates of oxygen and hydrogen, the pressure in all parts of the turbo pumps will never reach 0.5 mbar. This means that the turbo pumps were able to safely pump hydrogen with no risk of ignition. The pressure in the primary pump could however exceed 0.5 mbar. Not pictured in figure 5.16 is the purge gas system which allows the primary pump to passivate gases as they are ingested. By diluting the pumped gases with dry nitrogen at the pump intake, the hydrogen fraction by volume was guaranteed to drop far below the lower flammability limit. As a further precaution, a mass flow meter was employed to monitor the flow rate of nitrogen into the pump. If the nitrogen flow rate were to drop, the software would instantly close all solenoid valves, and sound an alarm.

Very few failure modes were identified in which hydrogen would flow into the laboratory environment. However, to be certain a hydrogen detector was fitted to the ceiling of the laboratory to alert the occupants of the room to any hydrogen leak well before the low explosive limit is reached. Even so, hydrogen was never left unsupervised in the laboratory.

Slight amendments were made to the thrust measurement procedure to accommodate the use of hydrogen. These include venting the hydrogen plena and filling them while in rough vacuum, and emptying them slowly at high vacuum before turning off the cryo-pump. This is further described in section A.4.1 of the appendix.

This set-up shown in figure 5.16 and the amendments to the procedure described above proved sufficiently safe for the health and safety officer to allow the hydrogen cathode test to be performed. However the majority of thrust measurements were still conducted with krypton, given the considerable reduction in risk, complexity, and man power over hydrogen experiments. The hydrogen could not be left unsupervised in the laboratory, meaning the plena must be filled, and the turbo and cryo pumps started all on the same day as testing. For krypton experiments, the cryo panel was generally run overnight. Lone working with the hydrogen was not allowed, meaning two test engineers were required to be present for at least 7 hours before thruster testing can begin. These time restrictions meant that only a small number of hydrogen tests were planned, with the majority of tests using krypton on the cathode.

Chapter 6

Results

This chapter details the performance results of the WET-HET characterisation campaign. These experiments utilised the WET-HET, cathode, thrust balance, Boltzmann vacuum facility and the experimental setup and procedures described in the previous chapter. The results presented here summarise 218 thrust measurements and the accompanying electrical and mass flow rate measurements of the WET-HET operating on oxygen, with the cathode operating on krypton between June 2020 and July 2021. Thrust, discharge voltage, discharge current, magnet current, and mass flow rate were measured directly. From these measurements discharge power, specific impulse, and anode thrust efficiency were calculated. In order to fully characterise the WET-HET, a full sensitivity analysis was conducted. This study saw to determine how thrust, specific impulse and anode thrust efficiency change as the following independent variables were manipulated:

- **Discharge power** - the WET-HET was operated in a current controlled mode with the voltage left to fluctuate as a response. Manipulation of the discharge current was used to tune the discharge power of the device.
- **Anode mass flow rate** - by manipulating the mass flow controller.
- **Magnetic field strength** - by manipulating the current limit of the magnet power supply

current.

- **Ceramic wall material (see below)** - by removing and replacing thruster walls.
- **Channel depth** - by substituting longer or shorter anodes.
- **Magnetic topology** - by raising and lower rings of the outer magnetic pole.

Note that the first three variables could be easily manipulated from outside the facility during testing, but changing any of the final three independent variables required venting the facility, re-calibrating the thrust balance and re-evacuating the facility.

All thrust measurements were conducted with the facility primary, turbo-molecular and cryo pumps all operating, which resulted in pressures of below 1×10^{-5} mbar, which increased to less than 5×10^{-4} mbar when both anode and cathode were supplying a massflow. The pressure for typical mass flow rates was closer to 1.5×10^{-4} mbar. All thrust measurements were performed with the cathode operating at 15 sccm of krypton, with the exception of those in section 6.11.

The cathode is assumed to contribute no thrust, given how small the magnitude of this thrust was, and the fact it was mounted at an angle of 82.5° to the thrust axis. This means that the specific impulse values refer explicitly to the anode specific impulse only, meaning the thrust and mass flow rate for the cathode is disregarded:

$$I_{sp} = \frac{T}{\dot{m} g_0}, \tag{6.1}$$

where T is thrust measured and $\dot{m}$ is the anode mass flow rate only. Discharge power P_d is defined as the product of discharge voltage V_d and discharge current I_d as measured by the anode powers supply during thruster steady state:

$$P_d = V_d I_d. \tag{6.2}$$

The anode thrust efficiency η_T is defined as:

$$\eta_T = \frac{T^2}{2\dot{m}P_d}. \tag{6.3}$$

The error bars on thrust measurements are calculated using equation 5.6, and then scaled for I_{sp} and anode thrust efficiency respectively, as mass flow rate, discharge voltage and discharge current are considered to be known exactly.

During the characterisation campaign the WET-HET was operated at discharge currents ranging from 1.5 to 10 A, which produced discharge voltages in the range of 153 V to 381 V, resulting in discharge powers from 360 W to over 3000 W. The thruster was tested at oxygen mass flow rates ranging from 0.72 mgs^{-1} to 1.845 mgs^{-1} (41.5-77.46 sccm) and magnetic flux density peaks of up to a peak centre-line value of 820 Gauss. The WET-HET performance was surveyed with three different channel depth: 35 mm, 45 mm and 60 mm. For the magnetic topology, three different configurations were tested, namely the 0, 3 and 10 rings raised. Two different channel wall materials were tested being alumina and boron nitride.

The current passing through the electromagnets is referred to as the magnetic current I_m. The magnetic flux density produced by a given magnet current depends on the magnetic topology, this is shown in Table 6.1.

Table 6.1: Probe measurements of peak radial magnetic flux densities in Gauss at channel exit centre-line for three different WET-HET magnetic topologies as a function of magnet currents.

Magnet Current (A)	**0**	**0.5**	**0.75**	**1**	**1.25**	**1.5**	**1.75**	**2**	**2.5**	**3**
0 Rings Raised	6	225	315	423	527	640	680	710	820	950
3 Rings Raised	31	210	300	380	490	540	580	650	720	780
10 Rings Raised	21	156	210	280	340	360	400	450	490	570

6.1 General Performance Trends

For ignition the WET-HET was supplied with 83 sccm oxygen, with 23-25 A of current passing through the cathode filament, the anode power supply set to 2.5 A and the magnetic field turned off. Typically the thruster was found to ignite when the discharge voltage surpassed 200-400 V. Upon ignition the discharge voltage drops to 80-100 V. The thruster is now in glow discharge mode. From here the voltage increases when the external magnetic field is applied. In certain configurations, the WET-HET proved to be more difficult to ignite than others. This was remedied by increasing the ignition discharge voltage, anode mass flow rate, cathode mass flow rate, or a combination of these. Values were always returned to their nominal values after ignition.

The current-voltage relationship, along with the resulting filtered laser data of a typical thrust measurement are shown in figure 6.1. Please note that noise on the current signal is introduced by the DAQ, and the actual current output is more stable than shown here. This was confirmed by comparing the DAQ current measurement to that of a bench top multi-meter. This particular thrust measurement (TM0094) was taken with an anode mass flow rate of 0.958mgs^{-1}, a BN channel 45 mm deep, and the magnetic configuration of 0 rings raised. The thruster state can be determined from the timeseries:

- **0-8 seconds** - The DAQ time is started with the magnet current off, the anode power supply off, but the anode and cathode mass flow rate on.
- **8-10.5 seconds** - The voltage of the anode power supply is increased rapidly, yet the open circuit does not allow current to flow.
- **10.5 seconds** - The potential difference between the positive anode, and the grounded cathode filament becomes great enough to ignite a discharge when reaching 250 V. The ignition allows a steady flow of 2.5 A of current to flow through the plasma, as the power supply switches to current controlled mode. This current is sustained by a reduced voltage

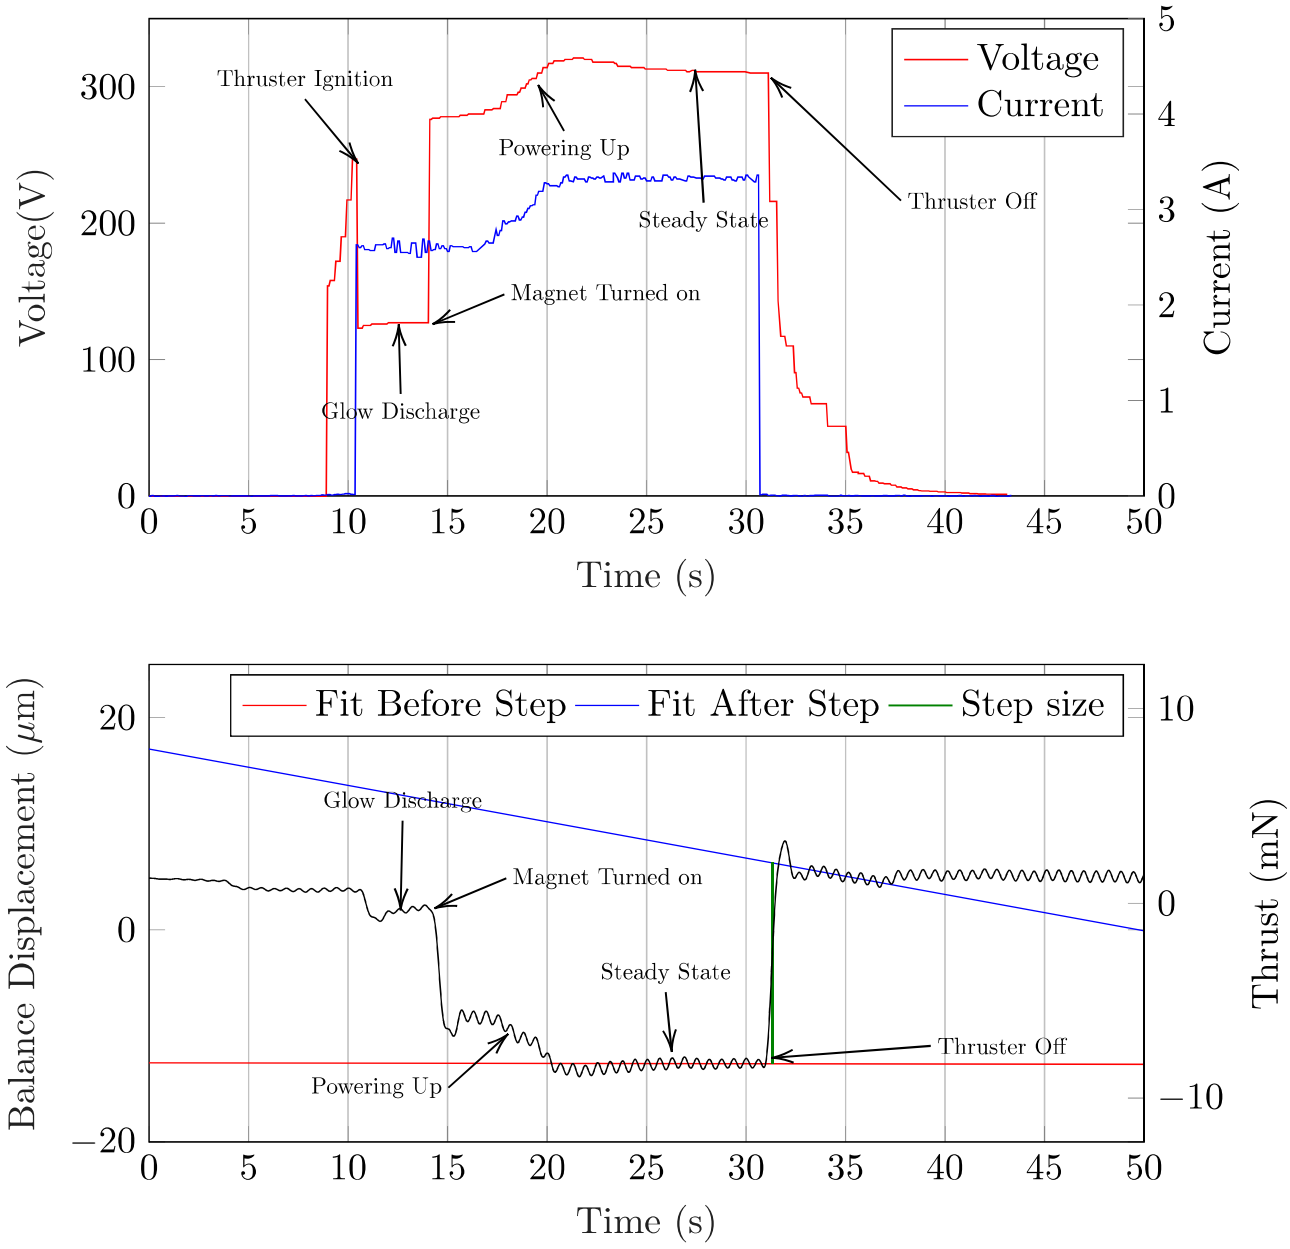

Figure 6.1: Example of electrical time-series (top) and filtered laser displacement sensor reading (below) for a typical thrust measurement. Please note that the current noise is introduced by the DAQ, and not a feature of the current itself.

of 125 V. The thruster is now in glow discharge mode. This thrust mode provides a low but detectable level of thrust, as one can see by the displacement of the thrust balance: the equilibrium position of the balance was shifted by approximately 2 μm at 10.5 seconds as the discharge was ignited. For this particular test the VCA thrust balance calibration was found to be 0.6125 mN/μm, which allows us to determine that the glow discharge mode has generated approximately 1.2 mN of thrust.

- **14 seconds** - At this time the magnetic current is instantly turned on to 1.5 A. Given the anode power supply is in current limited mode, the current is unchanged, but the discharge voltage dramatically increases as electron mobility is impeded by the establishment of the Hall current. The great shift in the balance displacement data indicates that the applied magnetic field increases the thrust by approximately 10 mN. This shows the importance of the applied magnetic field in thrust generation.

- **17-20 seconds** - During this period the discharge current is increased, which causes both the voltage and the thrust to increase. This was done until the discharge power (e.g the product of the current and voltage) rose to 1030 W which was the target power for this particular thrust measurement.

- **20-30 seconds** - The thruster is left to stabilise at the target power, during which the current, voltage and displacement all become consistent. Note that for the balance displacement oscillations about a constant value are considered to be constant, given that the under-damped balance always shows oscillations about the natural frequency near 1.5 Hz.

- **30.5 seconds** - The anode power supply is abruptly turned off, terminating current flow and thrust generation instantly. The mass flow rate of the anode and cathode are promptly shut off also. The voltage decreases less swiftly as the now static charge is slowly dissipated. The shift in equilibrium position of the laser (green line in figure 6.1) during the transition is determined by fitting a trend for the data just before the step (red line in figure 6.1),

fitting a trend to the data after the step (blue line in figure 6.1), and extrapolating each trend to the time of the step. This is described in further detail in section 5.3.7. The size of this shift is 17.915 μm, which equates to 11.0 mN. This force is the thrust, given that it is the force required to displace the balance from the thrust generating state, to the thruster shut down state. Note that the absolute displacement, and absolute thrust values do not matter, only the change in these values between different thruster operating conditions.

This thrust measurement is described as producing 10.97±0.24 mN at 1168±26 s of I_{sp}. Where the error has been calculate in accordance with equation 5.6. Directly before shut down the thruster was operating in a steady state at a discharge voltage of 313 V, and a discharge current of 3.31 A, so a discharge power of 1036 W. Using equation 6.3, this equates to an anode thrust efficiency of 6.07±0.27 %. Note that the relative error in the thrust and I_{sp} is 2.21%, where for anode thrust efficiency it grows to 4.425%. This increase in relative error is to be expected from the squared thrust term in equation 6.3.

The oxygen plasma ignites with a unique green-white colour which increases in brightness and density as the current is increased. When the magnetic field is applied, the colour changes to pink white, which becomes more white as the current is increases. This is shown in figure 6.2. This is presumed to related to the electron temperature, as different electron temperatures would allow different electrical excitation, which would result in the emission of different frequency of photons released during de-excitation.

6.2 Thermally Triggered Voltage Drop Phenomenon

The thermal design of the thruster assumed a worst case scenario in which 100% of the discharge power is deposited on the top face of the anode. The maximum thermal limit of this

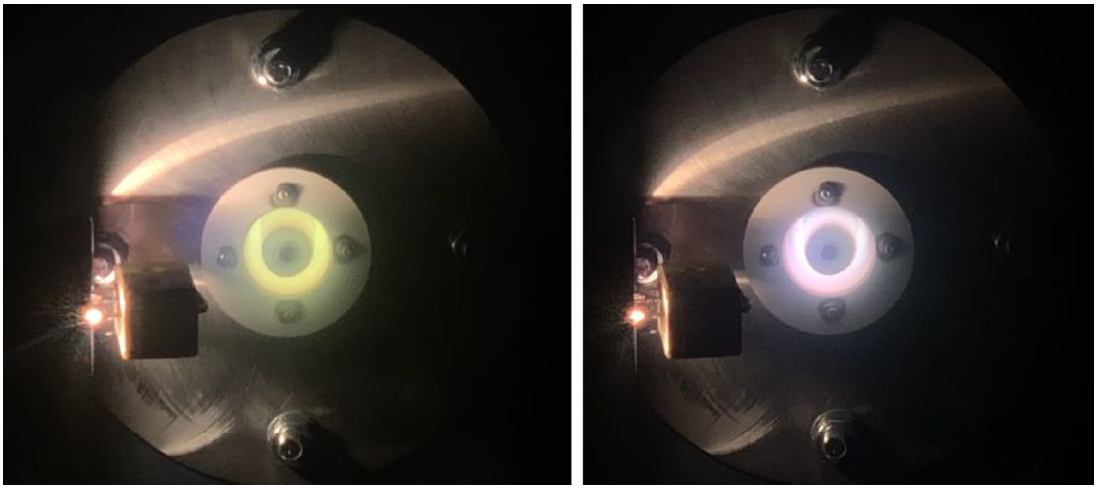

Figure 6.2: The WET-HET operating with no magnetic current (left) and with 0.5 A of magnetic current (right). The thruster is operating with a mass flow rate of 1.32 mgs^{-1} and a discharge current of 2.5 A in both cases.

model was the maximum operational temperature of the copper section of the anode. In reality it was found that a large fraction of energy was being deposited to the inner channel wall near the exit of the thruster. This was evidenced by the ceramic channel wall glowing in the visible spectrum after several seconds of operation at high powers (>1100 W) or high magnetic currents (>2 A). One proposed explanation of this heating was that the alumina surface acts to facilitate the recombination of atomic oxygen. Two atomic oxygen ions (O^+) require an electron to recombine into a molecular ion (O_2^+). Such a simultaneous three body collision is extremely unlikely to occur in bulk of the plasma, yet at the plasma boundary the ceramic wall is able to easily donate the required electron. Such an exothermic process is not seen in monatomic propellants like xenon, which may explain the dramatically increasing wall temperature near the WET-HET exit plane. Another suggested explanation for this extreme heating is that the ions are becoming magnetised to such an extent that they follow the radial magnetic field lines into the channel wall. The high thermal load on the inner channel wall presented two major issues. The first being that the inner channel wall was being subject to extreme thermal shock. Although no thermal measurements were able to be directly performed, the temperature was estimated to be raising from near room temperature to over 1400°C in less than ten seconds. This eventually lead to fracturing of the internal alumina channel.

Figure 6.3: The original inner and outer alumina ceramic wall components (left), along with the new boron nitride replacements (right). Note the top of the inner alumina wall (left most) has been heavily damaged from extreme thermal shock.

The second major issue with having an extreme thermal load on the internal ceramic wall was that this caused an anomalous discharge voltage drop phenomenon. The thruster was always operated in a current limited mode, with the voltage limit set very high. Generally the discharge voltage would remain relatively stable at a given discharge current, drifting $\pm 2\,\mathrm{V}$ over tens of seconds during steady state. However when operating the thruster in a regime which caused the internal ceramic wall to glow, this heating would trigger the discharge voltage to drop rapidly and continuously by $50\,\mathrm{V}$ or more over several seconds. These observations suggested that the electrons were finding a less resistive path to the anode. It was hypothesised that the alumina channel wall material was reaching temperatures above 1400°C, at which temperature the electrical resistance of alumina decreases in the presence of oxygen.[122;123;124]. This property of alumina was not considered when designing the WET-HET. Such a decrease in the electrical resistance of alumina would allow the electrons to bypass the plasma discharge as they travel axially down the channel via the heated alumina, resulting in the thermally triggered voltage drop that was being witnessed. To test this hypothesis the inner and outer alumina channel walls were replaced with AX05 grade boron nitride (BN) walls of exactly the same dimensions. Both alumina and boron nitride ceramic components are shown in figure 6.3. Not only does

the resistance of BN not change at elevated temperatures, but the grade chosen has particularly high resistance to thermal shock. Although changing the ceramic channel from Al_2O_3 to BN did increase thruster performance (see section 6.4) and resist further fracturing from thermal shock, it had no impact on solving the thermally triggered voltage drop phenomenon. This disproved the hypothesis that elevated temperatures were causing the ceramic to conduct electrons.

An alternative hypothesis was proposed to explain why high temperatures were causing the discharge voltage to drop. Given that the internal ceramic is directly in contact with the soft iron internal magnetic pole (number 6 in figure 4.2), it was assumed that a high thermal load on the ceramic could conduct to the magnetic pole, and reduce its magnetic permeability. Above approximately 500°C, the magnetic permeability of iron drops rapidly, to almost 0 at 770°C[125]. For the WET-HET, the heating of the iron core to above 500°C would dramatically reduce the magnetic field strength in the thruster channel, reducing the Hall current strength, and improving electron mobility to the anode. Effectively the heated thruster operates as if in glow discharge mode. Such a mechanism would be consistent with a sudden and rapid discharge voltage drop at high thermal loads. To test this hypothesis, the magnetic field strength within the channel was measured with a magnetic probe as the internal pole was heated with a blow torch. As expected, the magnetic field dramatically dropped when heated to around 700°C, well before it began to visually glow. Given that during some high temperature tests the glowing of the internal pole suggested it had been heated well in excess of 700°C, this is considered strong evidence that the thermally trigger discharge voltage drop phenomenon was directly caused by the internal pole becoming heated.

To combat the influence of this phenomenon on thrust measurements, high power experiments were conducted as quickly as possible, and aborted at first evidence of a voltage drop.

6.3 Voltage-Current Trends

The relationship between the discharge voltage and discharge current of the WET-HET at different magnetic field strengths is shown in figure 6.4 for two different mass flow rates. For this an other graphs a line of best fit has been included by standard fitting toolboxes. All trends were captured with an alumina channel 45 mm deep and a magnetic topology of 0 rings raised. The legend also provides the peak radial magnetic field strength along the centre-line of the channel at each magnetic current. The thruster was operated in current controlled mode, and reduced to the minimum discharge current possible before the discharge was extinguished. As a general trend the lower the magnet current, the lower the minimum achievable discharge current is. The voltage can be seen to increase linearly as a function of current for all trends, with greater magnet currents producing a slightly steeper gradient. A greater magnet current clearly demands a greater discharge voltage for a given discharge current. This is expected, as a greater magnetic current generates a stronger magnetic field, which in turn results in a stronger Hall current and thus a greater discharge voltage is needed to account for the greater impediment to electron mobility. Comparing between the two plots of figure 6.4, note that a greater mass flow rate requires a lower discharge voltage to sustain a given discharge current. This is presumably due to the fact that the same current necessitates a greater ionisation fraction at a lower mass flow rate to produce the same number of electrons.

Each trend in figure 6.4 shows a constant, linear trend in voltage over the currents surveyed. This suggests that the ion saturation limit of the WET-HET has not been reached for these mass flow rates. When this limit is reached in typical HETs, a clear change in the gradient of the discharge voltage indicates that the device is approaching the limit of how many ions it can extract from a given mass flow rate. The molar mass of the oxygen is about 1/4 that of xenon. This dictates a greater number of neutral particles are present in a given mass of propellant. This means that four time the number of ions can be extracted from a given mass flow rate,

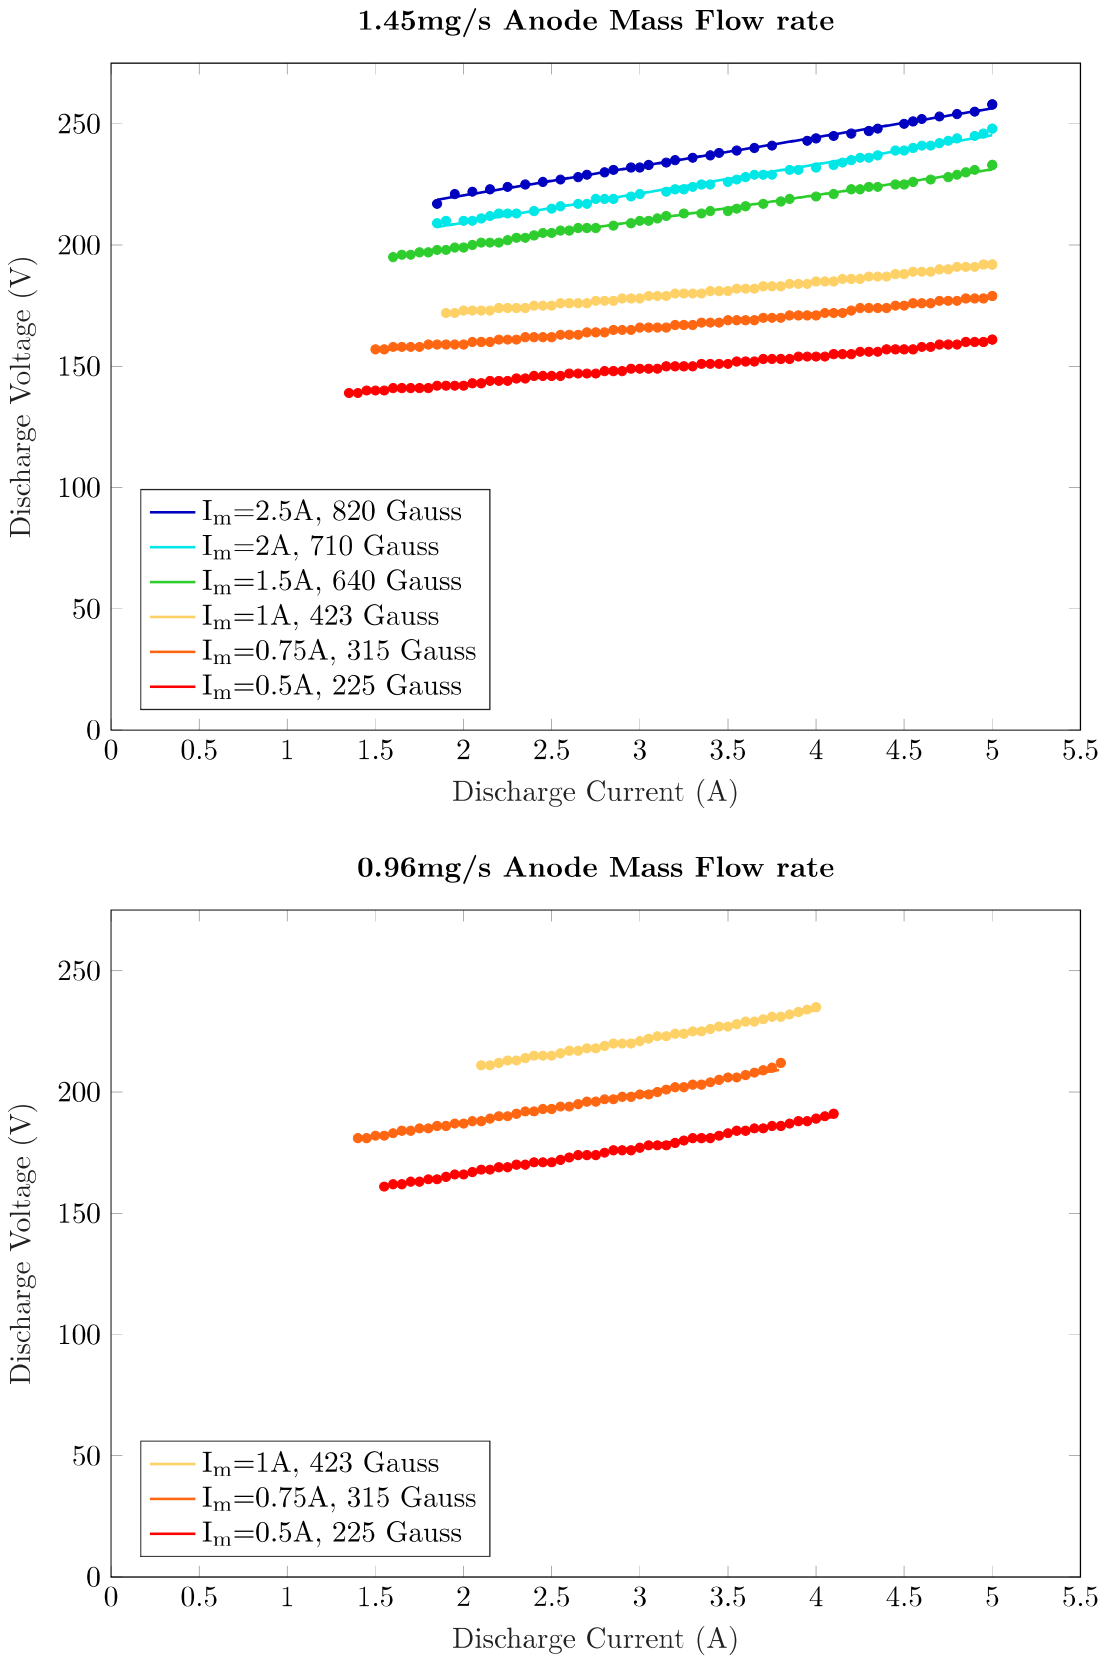

Figure 6.4: The voltage-current relationship of the WET-HET at different magnetic currents I_m for two different anode mass flow rates. The legend also shows the peak radial magnetic field strength along the centre-line of the thruster.

meaning a much greater current can be extracted from a given oxygen mass flow rate than from the same mass flow rate of xenon.

6.4 Influence of Channel Material

The inner and outer channel walls of the WET-HET discharge channel (numbers 4 and 5 in figure 4.2) were originally chosen to be alumina (Al_2O_3) for its low cost, high thermal conductivity and strong resistance to erosion. After the inner channel fractured due to thermal cycling, both inner and outer channel were substituted to AX05 grade boron nitride (BN) as shown in figure 6.3. This lead to a significant increase in performance, as can be seen in figure 6.5. For all these data the mass flow rate was $1.45\,\mathrm{mgs}^{-1}$, the magnet current was 1.5 A and the channel depth was 45 mm with a magnetic configuration of 0 rings raised. Changing the ceramic material from alumina to boron nitride increased thrust by 40% on average, which was also reflected in I_{sp} and efficiency. For both materials thrust, I_{sp}, anode thrust efficiency and discharge voltage are all highly linear with power with no obvious plateau, suggesting that performance would continue to grow at higher powers. When alumina walls were equipped the thruster is seen to require a lower discharge voltage to sustain a given discharge power when compared to the boron nitride. It follows that at a given power the alumina has lower performance due to a lower voltage, which is responsible for accelerating the ions. The increase in performance is most likely due to the lower secondary electron emission. This is discussed in further detail in the following chapter.

6.5 Comparison of Numerical and Experimental Results

Next the same experimental results of 6.5 are compared to PlasmaSim numerical results to assess the performance of the model. The domain of PlasmaSim 0D is set to match that of the thruster, with a channel depth of 45 mm. A magnet current of 1.5 A is used for the experimental results and a magnet configuration of 0 rings raised, however the magnetic field is not simulated

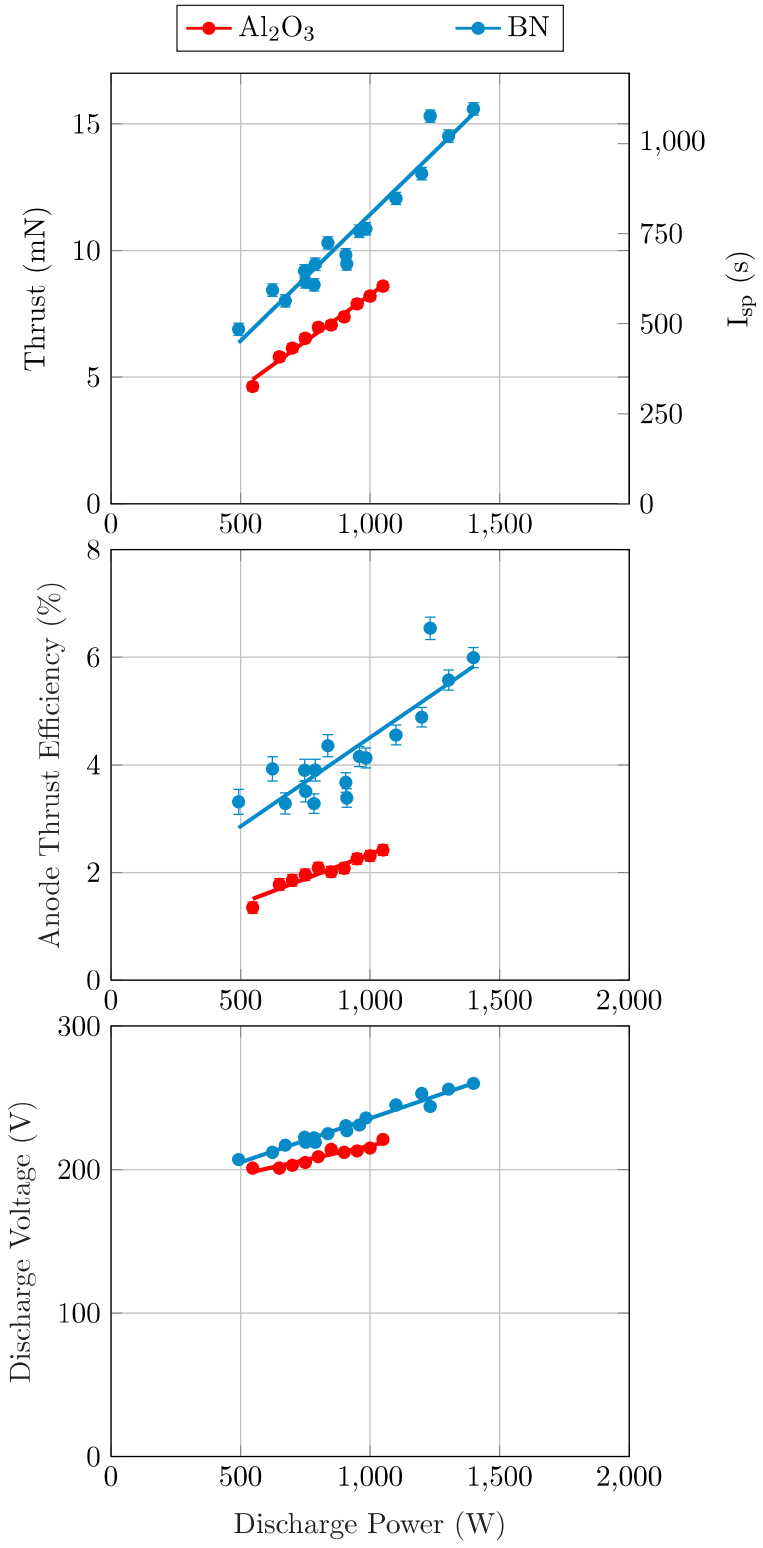

Figure 6.5: Impact on changing channel wall materials from alumina to boron nitride on thrust, I_{sp}, efficiency and discharge voltage as a function of discharge power. Mass flow rate was 1.45 mgs^{-1}, magnet current was 1.5 A, channel depth was 45 mm with a magnetic configuration of 0 rings raised.

in PlasmaSim. The same anode mass flow rate of 1.45 mgs^{-1} was used for both the experimental and numerical results. The simulations were run for 1 ms of simulation time, with time steps of 0.1 μs and 34000 superparticles. The first 0.1 ms of simulation time was excluded for the calculation of average values to allow the plasma to reach steady state. PlasmaSim requires a user specified electron temperature about which electrons will be simulated in a Boltzmann distribution. Given that this temperature is unknown for an oxygen plasma, simulations were performed about 10 eV, 11.5 eV and 16 eV. For all simulations the ions were distributed about 0.1 eV. Those simulations which produced a stable plasma are plotted in figure 6.6. The 10 eV data are only available for low powers due to the rise of violent instabilities at greater powers.

The PlasmaSim results are seen to be surprisingly independent of electron temperature for these simulations. PlasmaSim overestimates the experimental results significantly regardless of electron temperature. The simulated results fall closer to the boron nitride trend than the alumina ones.

It is assumed that some of the discrepancy between the experimental data and the PlasmaSim data can be attributed to the fact that the simulation does not account for several energy sinks. The power lost to the electrical excitation or the vibration, rotation, or breaking of the covalent bond is not modelled in this version of PlasmaSim. Furthermore no magnetic field is simulated, nor the complicated plasma-wall interactions which would differentiate between the two differing wall materials.

Although PlasmaSim overestimates the thrust, I_{sp} and efficiency for all cases, the 16 eV data does predict the general trends of thrust data well, showing a similar gradient to the experimental data.

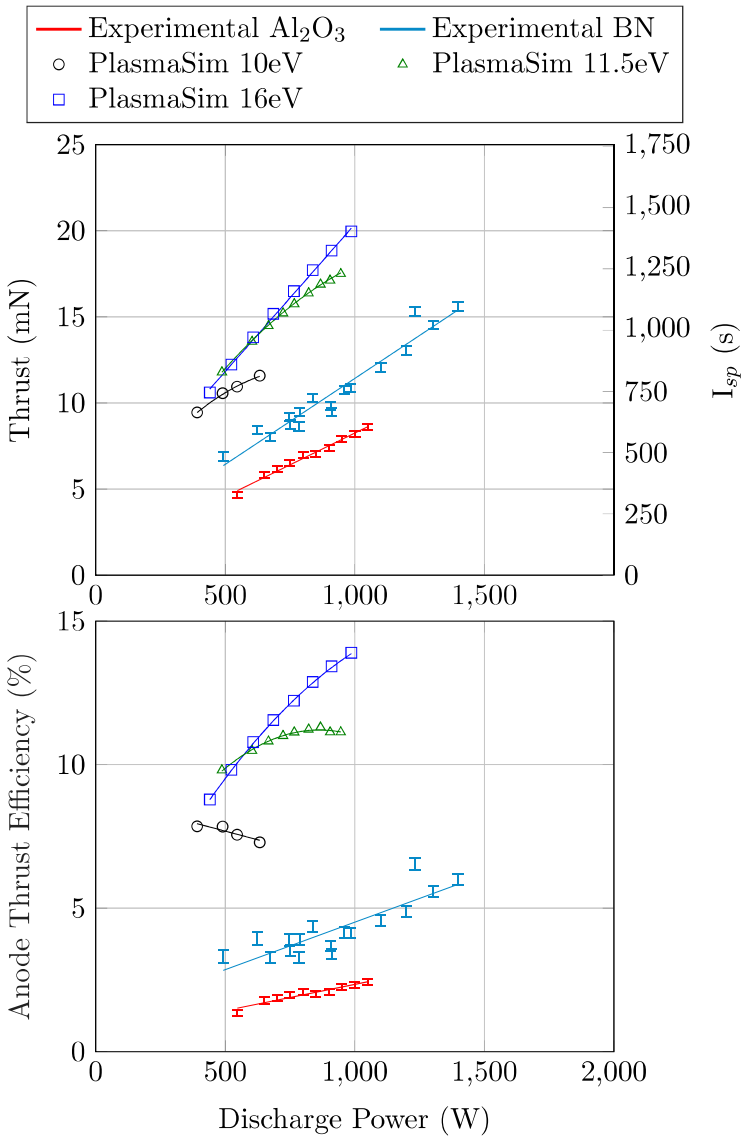

Figure 6.6: Comparison of experimental results to numerical results of the PlasmaSim 0D simulation. Both have a channel depth of 45 mm and mass flow rate of 1.45 mgs^{-1}. Results of PlasmaSim 0D are shown with three different electron temperatures. The experimental results use a magnet current of 1.5 A, and the 0 rings raised configuration.

6.6 Influence of Changing Magnetic Depth

The impact that changing the magnetic depth has on performance is surveyed in figure 6.7. The mass flow rate, magnetic current, and channel depth were all held constant at 1.45 mgs^{-1}, 1.5 A and 45 mm respectively such that the only difference in the trends was the magnetic configurations. Three magnetic configurations were tested namely the 0, 3 and 10 rings raised. The 0 rings raised configuration allows the WET-HET to be operated with a traditionally thin magnetic region near the channel exit. With three rings raised a thicker magnetic region is generated near the exit, and the ten rings raised configuration generates a magnetic region so thick that it envelops the entire channel volume (see figures 4.9 and 4.10).

Figure 6.7 shows that the highest thrust and I_{sp} was achieved with the 0 rings configuration which corresponds to the thinnest magnetic field. However, this configuration was the only one tested to this high of a power. The 3 rings raised configuration appears to produce a greater efficiency than the 0 rings configuration above 800 W.

Contrary to expectation, increasing the magnetic depth led to a decrease in discharge voltage at a given power. This decrease in voltage is presumed to be one of the factors that leads to a lower thrust and I_{sp} at a given discharge power. Although discharge voltage increases with power for all trends, the anode thrust efficiency of the 10 rings raised configuration appears to decrease slightly at greater powers. This differs strongly from the 0 rings raised and 10 rings raised configurations, where thrust, I_{sp} and efficiency rapidly increase with power.

The thrust and efficiency trends shown in figure 6.7 do not clearly show whether the 0 rings or the 3 rings magnetic depth produces a better performance, only that the 10 rings configuration has the poorest performance. However when we also consider the discharge voltage plots, we see more clearly that the 0 rings configuration consistently produces a greater potential difference.

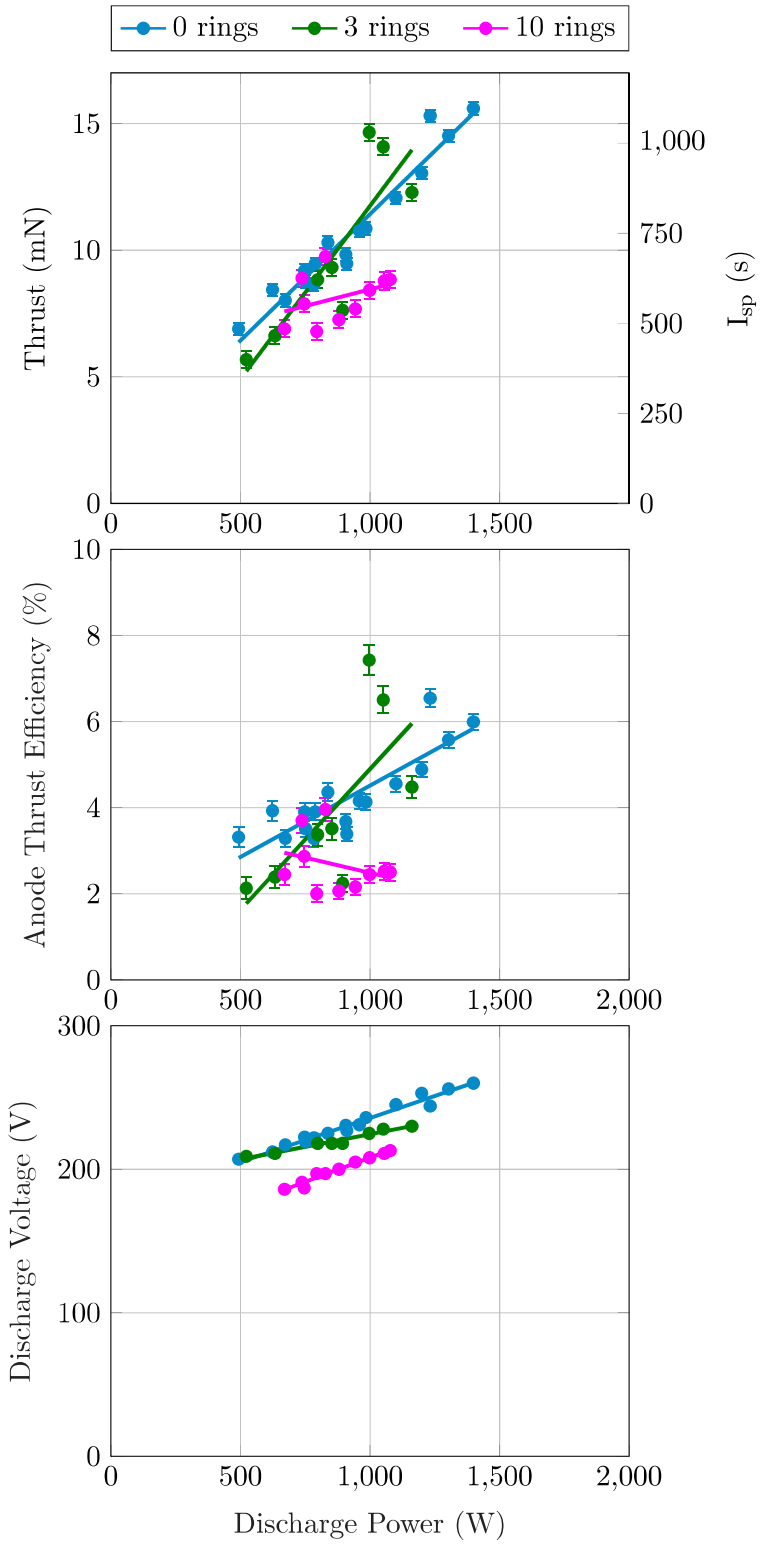

Figure 6.7: Impact of changing the magnetic depth on thrust, specific impulse, efficiency and discharge voltage. All data are collected with a massflow rate of 1.45 mgs^{-1}, BN walls, 1.5 A magnetic current, and channel depth of 45 mm.

This appears to be a better indication rather than the thrust, I_{sp} and efficiency data that a traditionally thin magnetic region is superior to a thicker one.

6.7 Mass Flow Rate Dependence

The impact of mass flow rate on thrust, I_{sp}, anode thrust efficiency and discharge voltage is shown in figure 6.8. These results are presented for three different magnetic depths with a boron nitride channel, and one magnetic depth for the alumina channel. Here power was kept constant at 600 W and 0.5 A of current was passed through the magnet coil. All data were taken with a channel depth of 45 mm. Figure 6.8 shows that in all configurations increasing mass flow rate results in a slight increase in thrust but a decrease in I_{sp}. One possible explanation of this is due to incomplete ionisation, adding additional mass flow has little impact on the number of ions being created. This would lead to more thrust due to gas dynamic effects, but not from more ions being accelerated, leading to a lower average exit velocity and hence a lower I_{sp}. For the magnetic topology of 0 and 3 rings raised, the boron nitride trends show a substantial decrease in efficiency as mass flow rate is increased. This trend of decreasing efficiency is less obvious in the 10 rings raised BN configuration and the alumina data. Again a dramatic performance increase is seen when the channel material is changed to boron nitride. Given the similarity in discharge voltages between the different configurations, the poor performance of alumina cannot solely be blamed on a reduced acceleration potential. It is important to note that while the 0 rings raised, BN channel configuration clearly shows the best thrust, I_{sp} and efficiency, it shares a similar discharge voltage trend to the other configurations. At mass flow rates above 1.23mgs^{-1} the 0 rings raised BN channel configuration even has a lower discharge voltage than the 3 rings raised trend, although it still has a greater performance. This is strong evidence that a discrepancy in performance between different thruster configurations relies on more complex mechanisms than simply a greater discharge voltage.

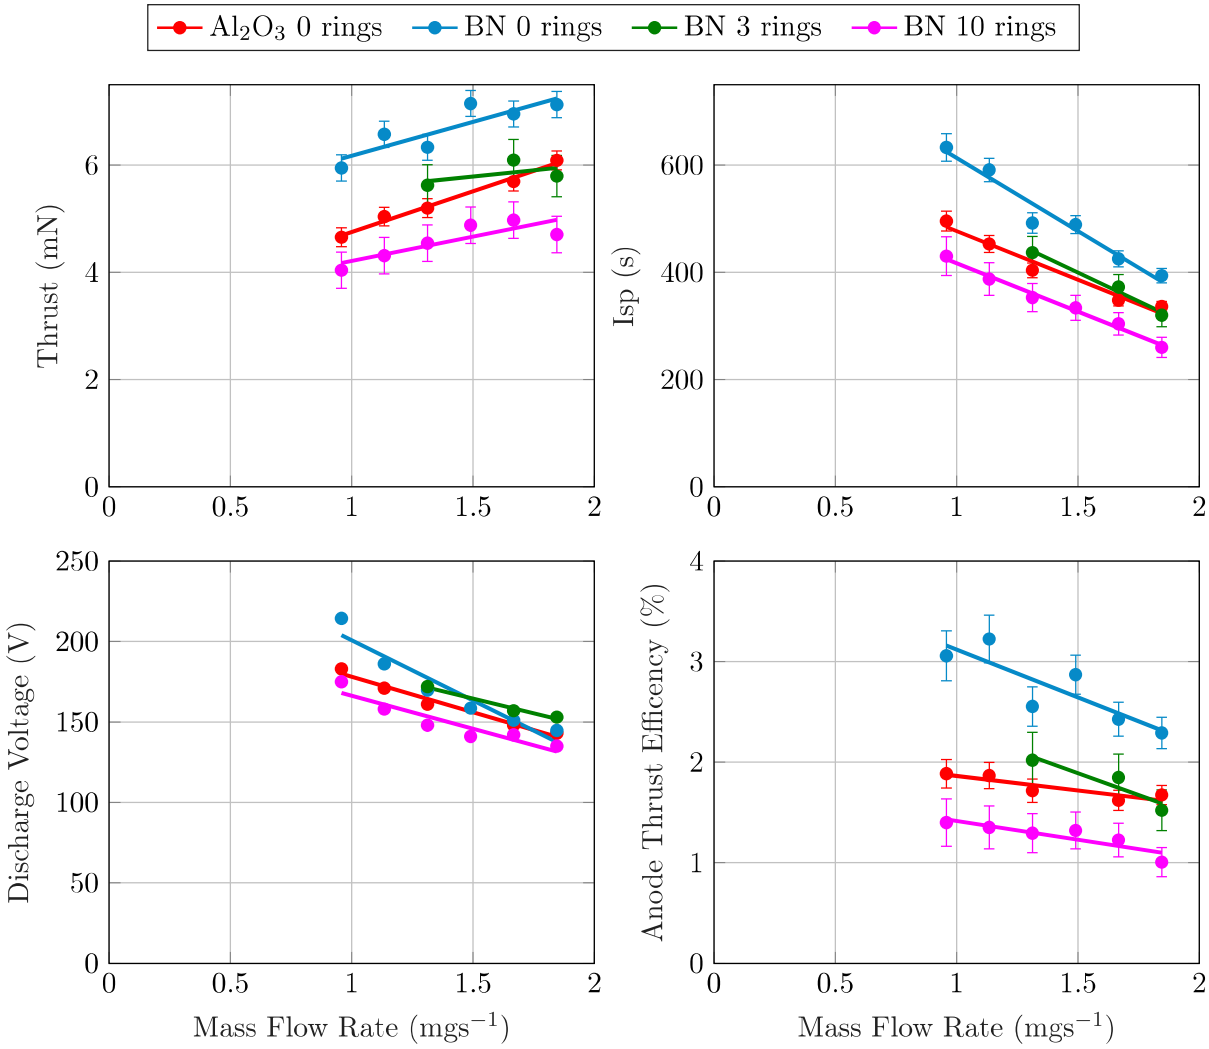

Figure 6.8: The impact of mass flow rate on thrust, I_{sp} and efficiency for the WET-HET with two different channel materials. All data are at 600 W discharge power, 0.5 A magnet current, and 45 mm deep channel.

Figure 6.8 shows more clearly than figure 6.7 that increasing the magnetic thickness results in a lower thrust, I_{sp} and efficiency. Noteworthy is that this all occurs with only a very small impact on discharge voltage.

6.8 Magnetic Field Strength Dependence

The three plots in figure 6.9 show how thrust, I_{sp}, efficiency and discharge voltage depend on the peak radial magnetic field strength. The radial component of the magnetic field strength along the channel centre line at different electromagnet currents was characterised using a magnetic probe before the facility was evacuated, allowing the field strength to be deduced during testing. For these thrust measurements the mass flow rate was fixed at $1.45\,\mathrm{mgs}^{-1}$, the discharge power at 750 W, and the thruster channel was 45 mm deep. A maximum thrust, I_{sp} and efficiency can be observed for all trends, with the field strength at which this maximum occurs increasing as the magnetic depth is increased. At magnetic field strengths above this the performance of the thruster is seen to decrease, even though the discharge voltage continues to increase.

The presence of a performance drop-off beyond a given magnetic field strength suggests the plasma is becoming fully magnetised, influencing the motion of both ions and electrons[102;124]: given the oxygen ions have the same charge but approximately a quarter of the mass of a xenon ion, equation 1.4 dictates that they orbit magnetic field lines at only a quarter of the Larmor radius. The magnetic field strength above the performance peaks is assumed to be strong enough to deflect the oxygen ions as they exit the thruster to such a level that performance suffers. Please see section 7.2 for a full discussion.

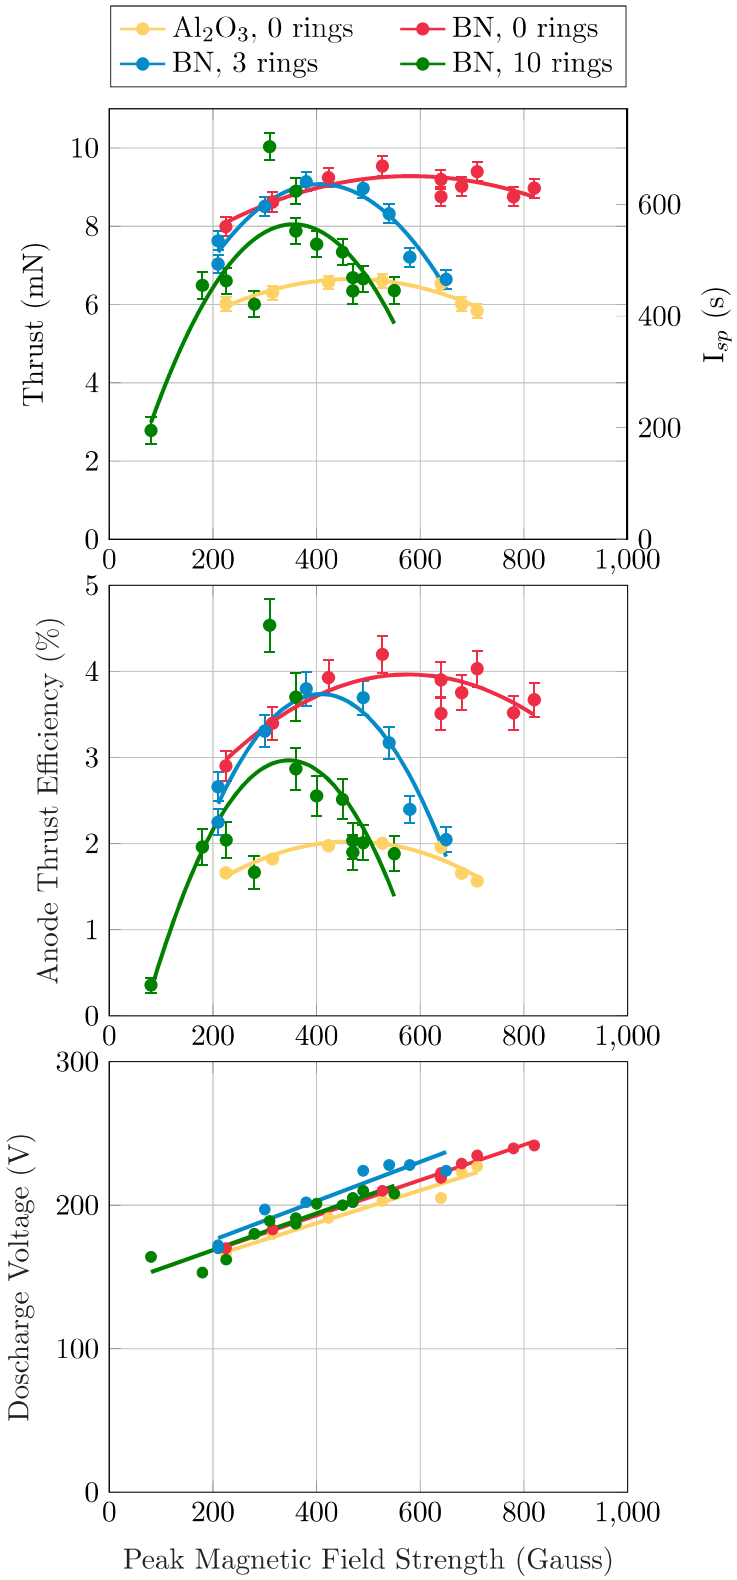

Figure 6.9: WET-HET thrust, I_{sp}, thrust efficiency and discharge voltage at different magnet currents for three different magnetic depths. All data are at 750 W discharge power, 1.45 mgs^{-1} mass flow rate and channel depth of 45 mm.

6.9 Channel Depth Dependence

By operating the WET-HET with anodes of different axial lengths, the impact of changing the channel depth of the thruster was measured. The anodes were replaced and repeat measurements were taken in which mass flow rate, magnetic field strength and magnetic topology were held constant over a range of powers. Figure 6.10 shows the thrust, I_{sp}, anode thrust efficiency and discharge voltage of the WET-HET with channel depths of 35 mm, 45 mm and 60 mm.

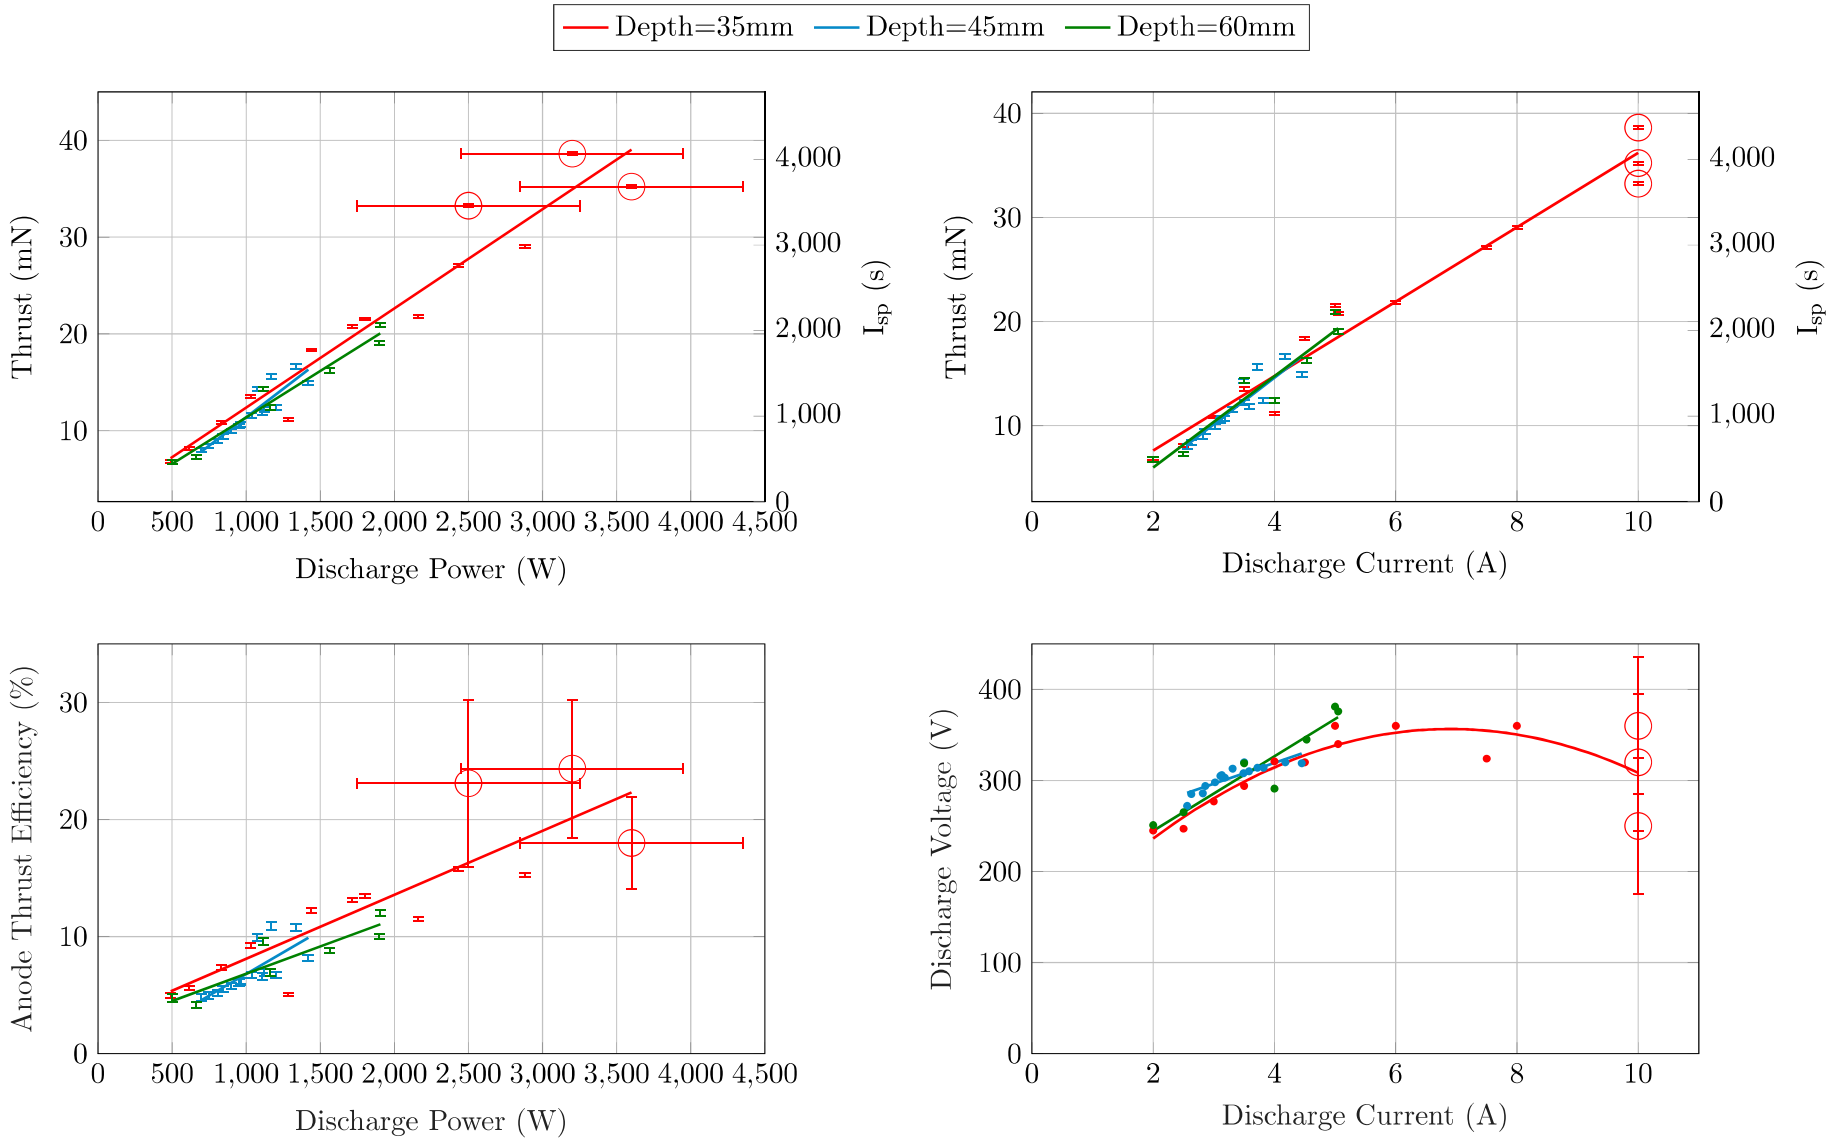

Figure 6.10: WET-HET thrust, I_{sp}, thrust efficiency and discharge voltage at different channel depths. For all trends the channel wall material was boron nitride, the mass flow rate was 0.958 mgs^{-1} and 1.5 A of current was passed through the electromagnets with a magnetic topology of 0 rings. The three measurements at 10 A are circled in red, and carry a large uncertainty in discharge voltage, anode thrust efficiency and power.

Data were collected for a range of powers with a mass flow rate of 0.958 mgs^{-1}, a magnetic topology of 0 rings raised, and 1.5 A of current passing through the electromagnets. Boron nitride channel walls were used for all data.

The 35 mm deep channel with the 0 rings magnetic topology and BN channel walls was shown to be the highest performing WET-HET configuration of any tested. This configuration was therefore prioritised for high power testing, such that the full capability of the WET-HET could be determined. Figure 6.10 includes a higher power range for this configuration, including three very high power measurements operating at 10 A discharge current. These three data are plotted as circles. Although confidence in thrust, I_{sp} and discharge current measurements were high for these three points, an issue with the data acquisition system made discharge voltage measurements unreliable. A conservative error estimate of ±75 V on these measurement has been applied. Equations 6.2 and 6.3 dictate that the large uncertainty in the discharge voltage propagates to cause a large uncertainty in discharge power and thrust efficiency, as seen by the significantly increased error bars of these three data in figure 6.10. The three 10 A data points show very good agreement with the trend line of the thrust-current relationship, showing that even at such high discharge currents, the thruster has still not reached its operational plateau. This seems to suggest that greater performance may still be achievable at greater currents yet. Even with the great uncertainty in these data, the anode thrust efficiency also appears to continue to climb with power. These data show the highest performance of the WET-HET captured during the campaign: a thrust of 38.63±0.25 mN at an I_{sp} of 4112±36 s. Conservative voltage estimates suggest that this was achieved at a power of 3200±750 W, and at an anode thrust efficiency of 24.35±5.89 %.

Changing the channel depth can be seen to have a surprisingly small effect on thrust, I_{sp} and efficiency of the device. The best performance by a small margin was achieved with the shallow-

est channel, with the deepest channel having the poorest performance. At powers below 1500 W the discharge voltage is relatively independent of channel depth. This changes at higher powers, with the 35 mm deep channel appearing to plateau at a maximum voltage near 340 V, although efficiency and I_{sp} continue to increase at a seemingly consistent rate.

6.10 Erosion

One of the major concerns of operating a HET on a reactive gas such as oxygen was that the components in contact with the plasma — namely the ceramic channel walls and metallic anode — will erode at an accelerated rate.

6.10.1 Erosion of Ceramics

Wall erosion in HETs is due to several mechanisms of plasma-material interaction, including ions bombarding the surface, thermal stresses and ion implantation into the ceramic[126;127]. Over time, erosion of the ceramics can lead to changes in the plasma boundary condition and eventual failure due to the plasma interacting with the magnetic circuit[127]. Given that oxygen plasma are industrially used for their high reactivity, erosion was expected to be a big issue for the WET-HET. Given the WET-HET is the first thruster which was exclusively tested on oxygen, common HET materials such at the alumina and boron nitride were specifically used for the channel walls. This was done to assess the rate of erosion in a worst case scenario, with no attempt for erosion mitigation. The degradation rate of the WET-HET will be used to assess the necessity of future thrusters to combat such erosion by employing exotic ceramics or magnetic shielding strategies[128;127].

During the test campaign erosion was witnessed most severely near the channel exit of the

inner ceramic wall of the WET-HET. Note that this is also the area where the magnetic field meets the wall in all magnetic configurations and where the thermal load is the greatest. This internal ceramic originally had a diameter of 15 mm. After 119 thrust measurements, the internal ceramic had been eroded to a diameter of 14.58±0.05 mm. This erosion can be seen in the top photo of figure 6.11. Up until this point, the maximum power the thruster had been operated at was 1560 W. Increasing the power level of the thruster accelerated erosion. After only 45 more thrust measurements, some of which exceeded 3000 W momentarily, the diameter had been reduced further to 14.35±0.05 mm, as shown in the bottom picture of figure 6.11.

The alumina channel showed less erosion than the boron nitride channel after testing, however, the alumina channel was operated for considerably less time, and at much lowers powers making direct comparisons between the two difficult.

6.10.2 Erosion of Anode

The anode of the WET-HET consists of a copper body and stainless steel cap section, both of which are in direct contact with the oxygen plasma at elevated temperatures. Figure 6.12 shows the discoloration of a used anode in comparison to one which has not been used. This black covering is most probably copper oxide, but did not appear to affect the performance of the thruster. Even slight handling of the thruster would cause this black coating to fall off, leaving a fresh copper surface below. The two left most anodes in figure 4.6 have been fired, but now look more similar to the unused anode seen on the right, after gentle handling.

The outer ceramic wall, which had experienced the exact same amount of testing showed no measurable signs of erosion.

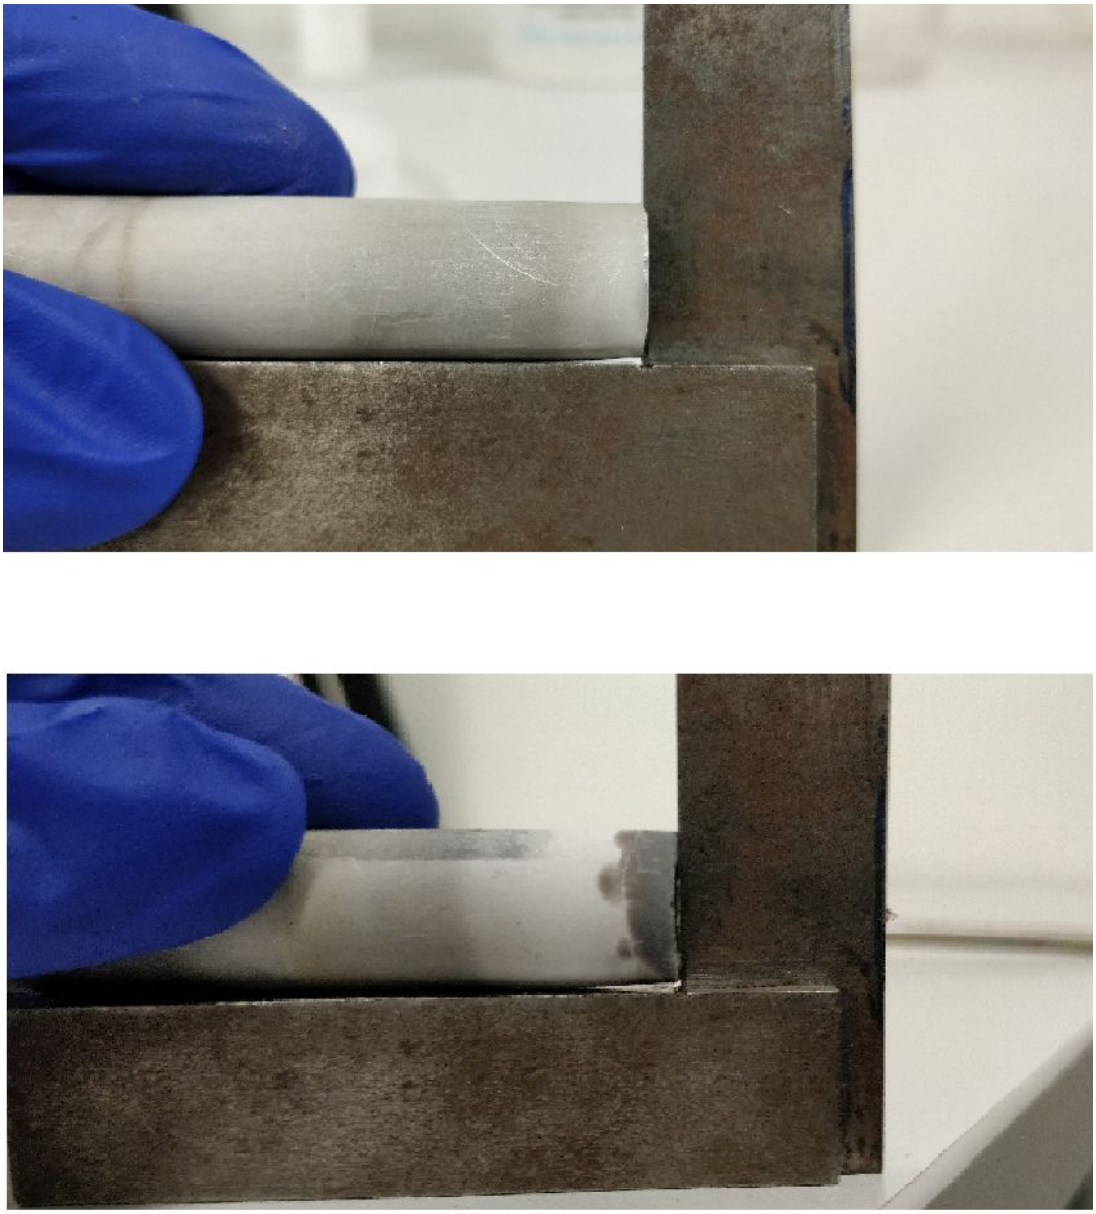

Figure 6.11: Erosion rate of the inner boron nitride thruster wall after 119 low power thrust measurements (top) and after 45 further high power thrust measurements (bottom). This shows a diameter change of 0.23 mm over 45 high power tests. The top face of the ceramic was originally a right angle. The set-square in the photos is used as a visual aid to see the right angle become tapered.

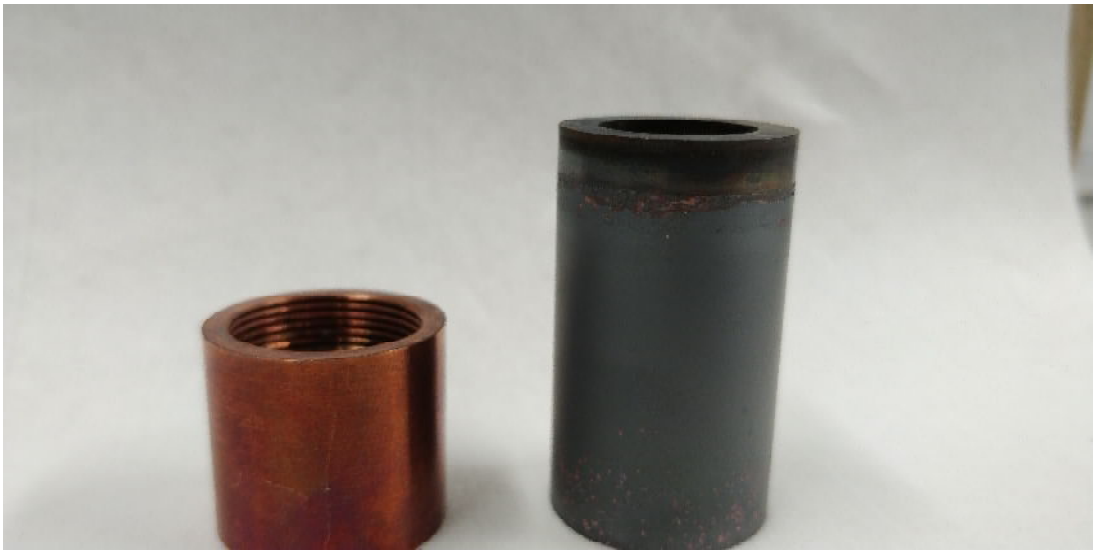

Figure 6.12: A new anode (left) compared to a anode which has seen several hours of operation, which has the anode cap attached (right)

6.11 Hydrogen Cathode Experiments

Of critical importance for the feasibility of the WET-HET system architecture is the ability of the cathode accompanying the WET-HET to be able to ignite and sustain the discharge while operating on $1/8^{\text{th}}$ the mass flow rate of hydrogen. In order to test this the experimental set-up and procedures described in section 5.6 were developed.

Before operating the thruster and cathode simultaneously, the cathode was operated independently on hydrogen to ensure that it would be able to provide the necessary electron current to sustain the thruster discharge. To test this an unusually high current was drawn from the cathode to the keeper to demonstrate that the necessary electron current could be generated from the designated hydrogen mass flow rate. This also served to test that the hydrogen storage and supply system was operating as planned.

For a hydrogen mass flow rate of $0.22\,\text{mgs}^{-1}$ the cathode was able to generate 5 A of discharge current at a keeper voltage of 30 V, at a nominal heater current of 27 A. This suggested the cathode was operating almost indistinguishably from its usual krypton configuration. Op-

erating at this current and massflow demonstrated that if the entire system was operated at the required 1:8 cathode-to-anode mass flow rate, the cathode would sustain the WET-HET at mass flow rates of $1.76\,\mathrm{mgs}^{-1}$, and discharge currents up to 5 A. This also shows that the maximum hydrogen mass flow rate of the MFC supply ($0.22\,\mathrm{mgs}^{-1}$) can be safely pumped by the Boltzmann vacuum facility, using the experimental set-up and procedures detailed in section 5.6.

The next hydrogen experiment to be undertaken was to test the entire system: igniting and operating the thruster on oxygen and the cathode on hydrogen. However, during the very first ignition attempt, the hydrogen MFC was destroyed by what was suspected to be a transient voltage spike from the thruster itself. Such high voltage spikes were often observed during ignition, as the anode aggressively attempts to draw electrons from any source it can. When this occurs it is suspected that these electrons are either donated by nearby data/command cables, or the resulting arcing causes strong currents to be induced in these separate circuits. The result is that these voltage spikes often cause the DAQ and other support electronics to temporarily disconnect or restart. Similar transient spikes have destroyed the servo-motor in the in-situ calibration system in the past, but this was the first time that the MFC was affected.

Precautions had been taken to avoid this: the MFC was mounted far away from the thruster in a metallic housing which served as a Faraday cage. All data/power connections to the MFC had been bridged to ground via a metal oxide varistor (see section 5.4.2). These varistors are commonly used for transient voltage suppression applications. Regardless, the PCB of the MFC was destroyed and had to be returned to the manufacturer for replacement. The delays caused by this fault did not allow for further hydrogen testing.

The hydrogen setup and procedure outlined in section 5.6 have been proven an effective method to safely store and supply hydrogen to the experiment for future testing. This experiment also

verified that the cathode was able to sustain up to 5 A of electron current at the maximum hydrogen mass flow rate of 0.22 mgs^{-1}. However, the critical full system test with the hydrogen cathode and WET-HET is yet to be experimentally demonstrated.

Chapter 7

Discussion and Future Works

This chapter reflects on the performance of the WET-HET as a proof of concept as the first HET designed specifically to operate on the products of water electrolysis. The results of the WET-HET optimisation and characterisation are discussed, as well as what these results mean for the feasibility of the water electrolysis concept in general. The chapter is concluded by identifying the key future work required to further develop the system.

7.1 Discharge Channel Material Selection

The motivation for changing the channel material from alumina to boron nitride stemmed from a concern that the alumina was becoming electrically conductive at elevated temperatures and was providing a conductive path for the electrons to reach the anode (see section 6.2). This change resulted in an increase in thrust by approximately 40%, but had no impact on the voltage drop. This voltage drop is now attributed to the heated magnetic pole losing its magnetic permeability.

The secondary electron emission (SEE) or secondary electron yield of a material describes the average number of electrons which are released by the material when a primary electron collides with it. The SEE of a material is a function of the temperature of the incident primary electron. Experimental data of the SEE of alumina and boron nitride taken from Gascon et al.[129] and

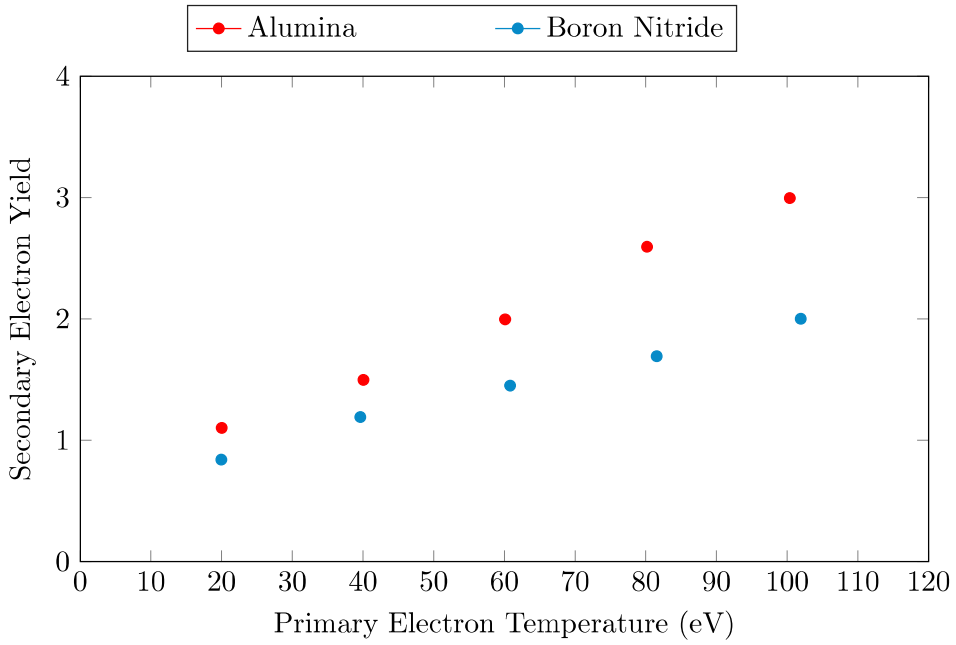

Figure 7.1: Experimentally measured secondary electron emission as a function of primary electron temperature for alumina and boron nitride. Data taken from Gascon et al.[129] and Bugeat and Koppel[130].

Bugeat and Koppel[130] respectively are shown in figure 7.1. We see that alumina has a considerably greater SEE than boron nitride for a given electron temperature. The secondary electrons emitted from the channel wall generally have a much lower temperature than those found in the bulk, meaning that an increase in SEE results in more but cooler electron populations in the plasma bulk[131]. Both numerical and experimental studies of xenon HETs have shown that channel walls of greater SEE lead to an increase in discharge current and thus a decrease in efficiency at a given discharge voltage[86]. Given the WET-HET has an unusually high surface-to-volume ratio, we would expect this effect to be exacerbated. The same phenomenon is seen for the WET-HET in figure 6.5, where at a given discharge power the greater SEE of alumina is assumed to contribute to a higher discharge current, and thus a lower discharge voltage. A lower discharge voltage provides less acceleration potential for the ions, which ultimately leads to a lower performance at a given power.

Based purely on thruster performance, boron nitride is obviously the superior choice for oxygen

HETs, much the same as it is for xenon HETs. Other materials such as boron nitride silicon dioxide (BN-SiO_2) show an even lower SEE and may lead to better performance yet[7].

7.2 Magnetic Field Strength

Figure 6.9 clearly shows that for all magnetic depths, a performance peak is found at a particular field strength. In all cases, beyond the peak, performance decreases although discharge voltage increases. One possible explanation for this is that the increased magnetic field strength reduces the size of the ion gyroradius so that it becomes comparable to the thruster dimensions. If the gyroradius becomes comparable to the channel dimensions, a significant portion of the ions deflected by the magnetic field will collide with the walls before they leave the thruster. This is referred to as the plasma becoming magnetised, and this contributes to greater channel heating and lower thrust and efficiency.

Let us consider the gyroradius for both atomic and molecular oxygen ions in the WET-HET using equation 1.4. The upper limit of the axial ion velocity can be estimated to be that of a singly charged ion which has been accelerated through the entire anode potential V_d and has thus gained kinetic energy $E_k = eV_d$ for the elementary charge e. For a lower limit of axial ion velocity the ion temperature is assumed to be 0.1 eV[93]. Consider the case of 1 A magnetic circuit current, which produces a field strength measured at 403 G. From the experimental data this is seen to result in a discharge voltage of 191 V. Substituting these values into equation 1.4 allows one to derive a possible ion gyroradius range from 6.4 mm to 279 mm for O_2^+ ions, and from 4.5 mm to 197 mm for O^+ ions. The channel of the WET-HET is 5 mm wide, meaning any ion attempting to gyrate at a comparable radius to this will strike the wall almost instantly, yet faster ions should leave the thruster with only a slight deviation. The estimation of gyroradii strongly suggest that the partial magnetisation of the oxygen ions is one of the mechanisms responsible for the great heat flux into the channel walls. Equation 1.4 shows that as the magnetic

field strength is further increased, the gyroradii of the ions will further reduce, and more ions will strike the walls. This explains the observed trend seen in figure 6.9, where greater magnetic circuit currents generate a greater discharge voltage, yet do not lead to a higher acceleration of ions, as seen by a lower thrust and I_{sp}.

The magnetisation of the ions also explains why the performance peaks at greater magnetic field strength when the magnetic region is thinner. For thin regions like the 0 rings raised configuration, the magnetic field is only present directly upstream of the channel exit. Even heavily deflected ions are able to exit the channel. Thicker magnetic regions cause such a deflection further upstream, sending the ion to collide with the wall. This however fails to explain why the 0 rings raised Al_2O_3 trend peaks at the same lower field strength for alumina than for boron nitride.

A comparable xenon thruster would most probably not witness such a decrease in performance until much greater magnetic field strengths are generated. This is due to the heavier xenon atoms having roughly a quarter the charge-to-mass ratio of oxygen and thus four times the gyroradius, resulting in less deflection at the same speed. The possibility of magnetising the oxygen plasma was considered when designing the WET-HET. This was one of the motivations for a design which allows the magnetic topology to be reconfigured by raising and lower rings. By decreasing the magnetic field strength, but increasing the thickness of the magnetic region, the design sought to sustain high levels of ionisation without fully magnetising the plasma. Contrary to expectation, figure 6.9 shows that the discharge voltage is unaffected but performance decreases when the magnetic region is broadened. The highest performance achieved was in the 0 rings position, which is that which produces the thinnest magnetic region, and is therefore the most similar to a traditional HET design.

The magnetic depth of the thruster was increased with the intention of broadening the ionisation region of the thruster. It had been hypothesised that a broader ionisation region would lead to increased plasma resistance, which would demand a greater voltage drop between the anode and cathode and ultimately increase I_{sp}. However, the bottom plot in figure 6.9 shows that all configurations have a very similar discharge voltage at a given magnetic field strength, regardless of magnetic topology or discharge chamber material. This important and surprising result shows that the discharge voltage of the thruster is dependent primarily on the magnetic field strength but not the thickness of the magnetic field profile. By extension this also shows that the discrepancy in thrust, I_{sp}, and efficiency between the different magnetic configurations shown in figure 6.9 cannot simply be due to a difference in the discharge voltage, and is most likely due to more ions becoming magnetised and being deflected.

7.3 Magnetic Topology

One of the considerations when designing the WET-HET for oxygen operation was that magnetic fields of the same strength as in a conventional xenon HET would risk magnetising the ions, given the high charge-to-mass ratio of oxygen ions compared to xenon. The WET-HET design sought to combat this with a thicker, lower strength magnetic field. However, the results have clearly shown that increasing the thickness of the magnetised region is detrimental to thrust, I_{sp} and efficiency at all points tested. While this conclusion was not overly clear just from figure 6.10, the data in figures 6.8 and 6.7 clearly show that the 0 rings raised magnetic topology offers the highest performance. Oxygen and xenon differ greatly in terms of mass and ionisation cross section, yet these results have clearly shown that like xenon, oxygen benefits from a thin magnetic region near the exit of the thruster. This observation is corroborated in literature for xenon: the effect of thickening the magnetic region in a typical xenon hall effect thruster to determine changes in the plasma profile was numerically performed by Ahedo and Escobar[92] using a one dimensional macroscopic plasma model. The authors of the study found that as the magnetic

field is broadened, the region of ion back-streaming near the anode was reduced in axial length, which pulled both the ionisation region and the acceleration region further upstream. As the acceleration region becomes axially longer, wall energy losses increase, which ultimately results in a decrease in thrust efficiency[92]. The same efficiency trends are observed for the WET-HET, for example in figure 6.8 where one clearly sees a decrease in anode thrust efficiency as the axial thickness of the magnetic field increases. Plasma diagnostic measurements of the WET-HET are required before it can be confirmed that a similar shift in the plasma profile is responsible for the trends that are being observed.

7.4 Channel Depth

In the same study Ahedo and Escobar[92] investigated the thruster response to change in the channel depth at a given discharge voltage. They found that increasing channel length has little effect on the location of the beginning of the acceleration region, but that this region grows in axial length with the channel depth. A thicker acceleration region then contributes to a greater wall loss. They conclude that a shorter channel generally produces a greater efficiency, up to a critical point below which the electron temperature and performance rapidly drop[92]. The WET-HET efficiency appears to follow a similar trend, in that the shortest channel tested was the highest performing, suggesting the critical channel depth is less than 35 mm. This is shown in figure 6.10. However, the 60 mm long channel performs extremely similar to the 45 mm channel, which is unexpected. The results of Ahedo and Escobar[92] would suggest a larger performance discrepancy between the 45 mm and 60 mm configurations. It was originally hypothesised that a very deep channel may contribute to a greater rate of oxygen dissociation. This would increase the ratio of atomic ions to molecular ions, and effectively increase the charge-to-mass ratio of the average ion. The unexpected increase in performance of the 60 mm deep channel may be due to the greater average ion charge-to-mass ratio outweighing the channel wall power losses. Such a performance increase at deeper channels would be a direct consequence of the molecular nature

of oxygen, which would explain the discrepancy between the WET-HET results and numerical xenon simulations in the literature. Future plume measurements are planned to determine the charge-to-mass ratio of ejected ions to determine whether this is a factor in the performance shift.

7.5 Propellant Utilisation Efficiency and Mean Ionisation Length

The propellant utilisation efficiency α_p of the WET-HET is compared to that of a traditional xenon HET, to assess how effective the thruster is in ionising the propellant. This efficiency is generally defined as the fraction of neutral particles which are converted into ions by the device:

$$\alpha_p = \frac{I_i}{I_{\dot{m}}}, \tag{7.1}$$

where I_i is the ion current and $I_{\dot{m}} = e\dot{m}/m_n$ is equivalent current for the injected mass flow for the elementary charge e and neutral mass m_n[100]. A propellant utilisation efficiency of $\alpha_p = 1$ would mean that every neutral particle is successfully ionised by the thruster. Note that in this form the equation only holds for singularly charged ions. An important feature of molecular propellants is that a single neutral particle has the potential to dissociate, with the daughter particles becoming ionised, setting the maximum attainable propellant utilisation efficiency for oxygen at $\alpha_p = 2$.

If it is again assumed that all ions are singularly charged, the same propellant utilisation efficiency features when calculating thrust T as a function of ion mass flow times the ion velocity in the form:

$$T = \alpha_p \dot{m} \sqrt{V_d(1-\delta)2e/m_n}. \tag{7.2}$$

For discharge voltage V_d and the discharge voltage loss coefficient δ[100;99]. This coefficient is related to the discharge efficiency, such that $V_d(1-\delta)$ is the actual potential that the ions really

experience[97]. The value of the discharge voltage loss coefficient for the WET-HET is not known, yet for xenon thrusters it is generally estimated to be $\delta < 0.3$, with some estimating as little as $\delta < 0.1$[132]. To be conservative it shall be assumed that for the WET-HET the discharge voltage loss coefficient can be found in the range of $0 \leq \delta \leq 0.5$.

The propellant utilisation efficiency is also useful in that it allows the mean ionisation length λ_i to be calculated, which one will recall is the average distance that a neutral particle will travel before becoming ionised. As defined in equation 4.1, the Melikov–Morozov criterion states that the channel length must be greater than several mean ionisation lengths in order to generate a large propellant utilisation efficiency[133;97]. The mean ionisation length can be calculated from the propellant utilisation efficiency by comparing it to the channel depth L[100]:

$$\lambda_i = \frac{-L}{\log(1 - \alpha_p)}. \tag{7.3}$$

The results of the experimental campaign can be used to calculate both the propellant utilisation efficiency and the mean ionisation length of the WET-HET from equations 7.2 and 7.3. Here the thrust and discharge voltage shown for the L=35 mm trend in figure 6.10 were used, yet the I_d=10 A data were omitted due to the large discharge voltage uncertainty. The propellant utilisation efficiency and mean ionisation length are plotted in figure 7.2. As the value of the discharge ionisation loss coefficient is unknown, the values for δ=0 and δ=0.5 are plotted, with the assumption that the true value will fall somewhere between these limits. For comparison, data for an SPT-100 thruster with δ=0.05 operating on xenon is also included. For these data, experimental data from Sankovic et al. [12] is used, at a discharge voltage of 300 V and mass flow rates ranging from 3.45 mgs^{-1} to 5.45 mgs^{-1} [12]. The WET-HET is seen to require considerably higher discharge powers to achieve propellant utilisation efficiencies and mean ionisation lengths which are comparable to a traditional xenon thruster. The WET-HET does not rise

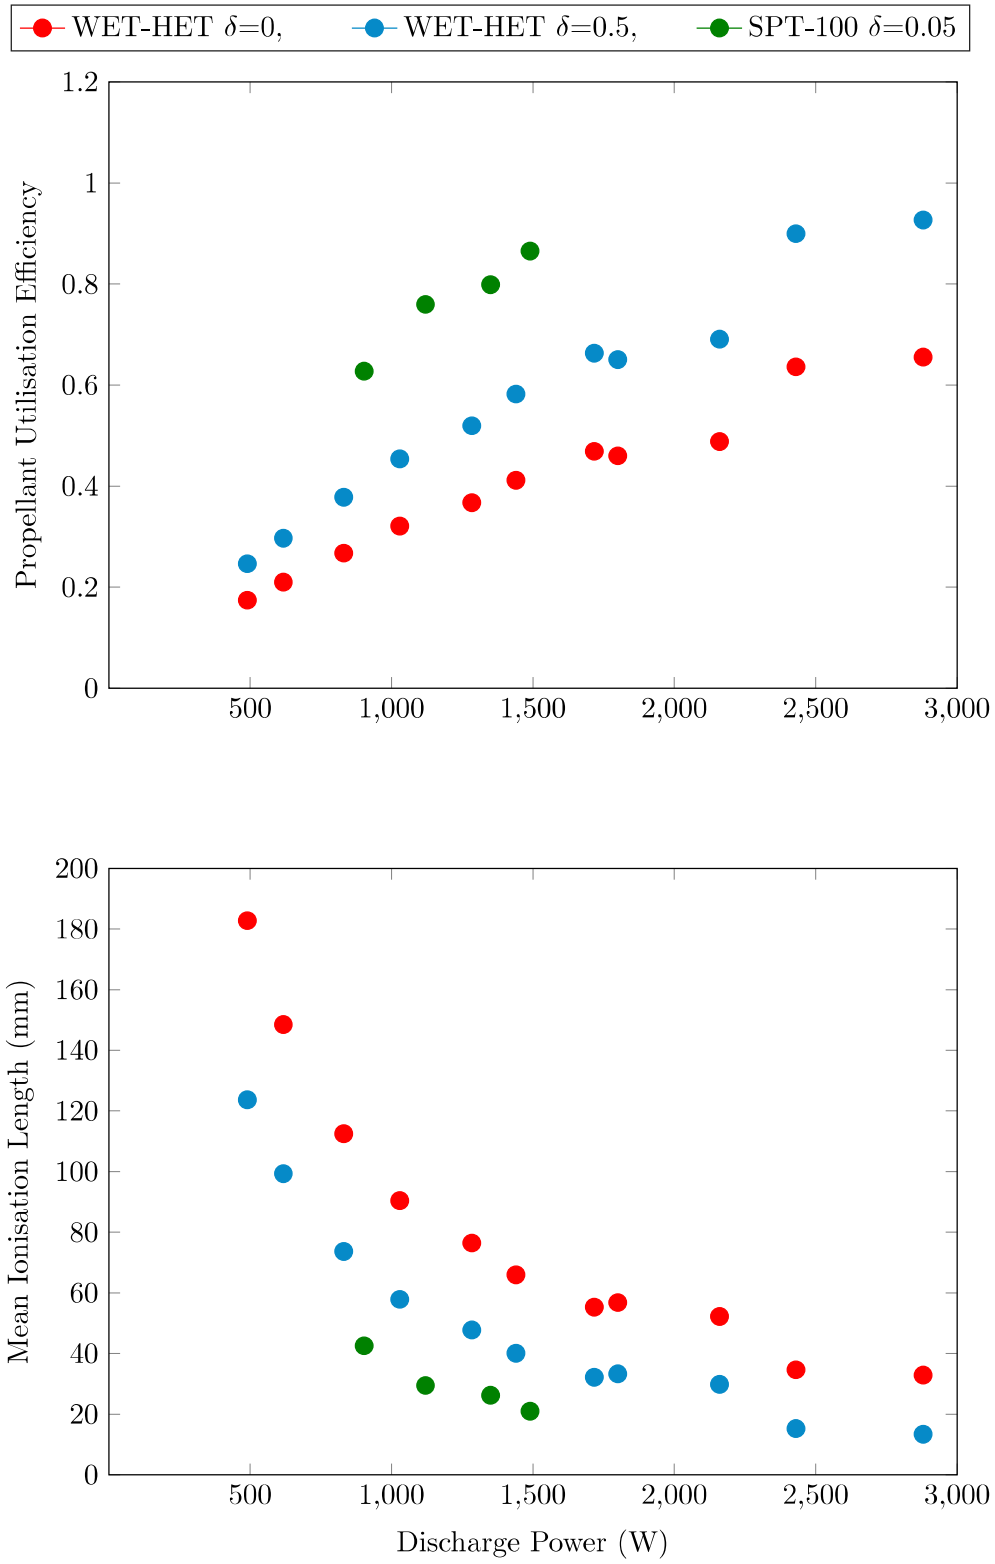

Figure 7.2: Propellant utilisation efficiency and mean ionisation length of the WET-HET and an SPT-100 operating on xenon as a function of discharge power. The WET-HET data were taken at the same conditions as those in figure 6.10. SPT-100 data is taken from experimental data from Sankovic and Hamley 1994[12].

above $\alpha_p = 1$, although it could in theory be as high as $\alpha_p = 2$ by our definition (equation 7.2). At 2880 W of discharge power the WET-HET is producing a propellant utilisation efficiency of 65-92% and 8 A of discharge current. If we are to extrapolate the data in figure 7.2 to higher powers, the WET-HET would need to operate at between 17-24 A discharge current to achieve the maximum propellant utilisation efficiency of $\alpha_p = 2$, which is the current saturation limit. This does not suggest that this is necessarily possible, or that the WET-HET could sustain such a high power and thermal load. However, this does suggest that unlike a xenon thruster, the WET-HET is able to operate at considerably greater currents, and thus powers and specific impulses at this mass flow rate. Note however that the generation of atomic ions would further increase the charge to mass ratio of the accelerated ions, further increasing I_{sp} but decreasing efficiency.

An alternative way to define the ionisation length is by the approximation[97]:

$$\lambda_i \approx \frac{v_n}{n_n \sigma_i(T_e) v_e(T_e)}, \tag{7.4}$$

where v_n and v_e are the velocities of the neutrals and electrons respectively, n_n is the neutral density, and σ_i is the ionisation cross section as a function of electron temperature T_e. The WET-HET and SPT-100 operate at similar volumetric flow rates of 40-60 sccm, yet the WET-HET has a smaller channel cross section by a factor of 13.3. Consequentially the neutral density of the WET-HET is considerably higher. Equation 7.4 shows us that the results in figure 7.2 indicate the WET-HET is only able to produce a similar mean ionisation length (and by extension propellant utilisation efficiency) by combating the low ionisation cross section of oxygen with a high neutral density. Even then the ionisation length is only similar when the WET-HET is operating at very high powers, which presumably generates a high electron temperature.

For the highest power, the estimation of the mean ionisation length of the WET-HET is in between 13 mm and 33 mm. The Melikov–Morozov criterion in equation 4.1 suggest the the mean ionisation length should be significantly larger than the channel depth. For the results in figure 7.2 the WET-HET has a channel depth of 35 mm, which accounts to 1.06-2.69 ionisation lengths. In comparison, the SPT-100 has a minimum ionisation length of 21 mm, but a channel depth of 42 mm, suggesting that an average of 2.0 ionisation lengths span the channel depth. So while the WET-HET channel depth does not span a high number of ionisation lengths, this is in line with similar xenon thrusters under these assumptions.

Many xenon HETs are designed by modifying successful thrusters using scaling laws, derived using fundamental plasma theory[97]. However, given the large differences expected between a xenon and oxygen plasma and the lack of available data on prior oxygen HETs, our approach was to optimise the thruster using a zero-dimensional particle-in-cell model to capture the influence of geometric and field parameters self-consistently[134]. One of the results of this optimisation was a considerably smaller channel cross section of a thruster in this power class, but also a greater surface-to-volume ratio. The WET-HET has approximately 2.5 times the area per unit volume as a conventional SPT-100. The increase of the surface-to-volume ratio was a direct consequence of reducing the channel cross section to increase neutral density at the target mass flow rate. The experimental findings illustrate the efficacy of this approach. The thruster was able to maintain a stable discharge at the low propellant flow rate of 0.985 mgs^{-1}, at discharge powers in excess of 3 kW. Under these conditions a high propellant utilisation efficiency is reached as seen in figure 7.2. Increasing the channel cross section area would require a greater mass flow rate to achieve the same neutral density, and thereby even higher discharge powers to attain the same propellant utilisation efficiency. There is however, an important trade-off to consider. A consequence of a high surface-to-volume ratio is increased power losses to the wall, and thereby a further reduction of thrust efficiency. During testing high thermal loads on the channel walls

are witnessed, driven by these wall losses, evidenced by visible glowing of the ceramic walls at high discharge powers.

7.6 PlasmaSim 0D

Ideally the numerical optimisation of the WET-HET would have included a high-dimensional, fully electromagnetic PIC software with complex plasma-chemistry mechanisms. Unfortunately, such a model was not available at the time, and even if it had been, the computational time required for such a model to perform an optimisation for five independent variables (channel depth, channel width, channel circumference, mass flow rate and discharge voltage) would far exceed the time available. One of the original goals of the current study was to experimentally test hardware. Physical tests were considered a paramount first step to demonstrate a proof of concept and to identify any unanticipated fundamental roadblocks for the technology. A balance in plasma model fidelity had to be found: on the one hand, a poorly optimised thruster, stemming from a low fidelity PIC code would not provide an accurate representation of the capability of a water electrolysis HET system. On the other hand, the software development and thruster optimisation would need to leave ample time for hardware design, construction, and testing.

Reducing the dimensionality of PlasmaSim from one dimension to zero dimensions accelerated the computation speed of the code by more than a factor of 1000. This allowed sufficient time for the thruster optimisation. The channel depth was predicted to have an impact on the dissociation rate of the oxygen, the physics of which our 0D PlasmaSim version could not simulate. The inclusion of this would have introduced two more heavy species (e.g, O and O^+) to the model, each requiring reaction rates and the simulation of recombination. The inclusion of these interactions would have been highly complicated and taken a significant period of time to implement. Similarly, the inclusion of full magnetic modelling would have either demanded a spatially homogeneous magnetic field, or the move to a higher dimensional model such that local plasma

parameters could be resolved. The first option would not have allowed discrimination between different magnetic topologies of the WET-HET, and the second option would have required a dramatic increase in computational costs, and therefore time.

The decision of when to stop model development and begin thruster optimisation was difficult at the time, and can be evaluated better in hindsight. In review, any less time characterising the WET-HET would have severely impacted the conclusions of the experimental campaign. The best performance of the WET-HET was captured very near to the end of the period which was allocated for this project. The results of the SPT-100 sensitivity analysis (see section 3.5.1) gave confidence that the model was ready to perform the full thruster optimisation. It is difficult to imagine that any less time spent developing 0D PlasmaSim would have produced such a validation and ultimately such a successful WET-HET optimisation. It can be concluded that by chance an appropriate portion of the overall project time was allocated to improving the fidelity of 0D PlasmaSim.

The usefulness of reducing the complex physics of a HET to a zero-dimensional problem is obviously very limited. The primary goal of this project was to practically test a HET operating on oxygen. To achieve this goal, PlasmaSim-0D was chosen as the most appropriate tool available at the time. While the model was shown to be very sensitive to geometry (as seen if figures 3.5 and 3.6), a more complex model would provide a different sensitivity which better reflects reality. The author strongly suggests that any further HETs optimised for oxygen operation employ simulations of a greater dimensionality, of full magnetic capabilities and which can simulate the complex behaviour of the oxygen molecule. Doing so would presumably not only increase the performance of the device, but identify important design considerations. For example, a magnetic model which is able to better describe the oxygen dissociation and recombination would have identified the thermal limitations of the WET-HET earlier, and thus

allowed for further engineering efforts to rectify this.

One of the big limitations in simulating an oxygen HET was the lack of experimental data available to provide model validation. The electrical and thrust data collected from WET-HET testing may fulfil this role in future studies.

Since these early simulations, PlasmaSim has now been developed into a two-dimensional electromagnetic code, and validated against similar models and experimental data[96]. This version of the model would now be much better suited to the optimisation of an oxygen thruster.

7.7 Thermal Limitations of the Design

Not only is the WET-HET very small for its power class, but it has an exceptionally large surface-to-volume ratio, making heating much more prevalent than for a comparable xenon thruster. This severely impacts the competitiveness of the thruster, as the greatest performance was found at set-points unable to withstand the thermal load under sustained operation. The WET-HET thermal design assumed that in the worst case scenario, the entire thermal load would be concentrated on the face of the anode. In reality, it was found that the inner ceramic was taking most of the thermal stress directly at the channel exit. No direct measurements of this were taken, but the tip of the ceramic was visually observed to glow brightly, suggesting this was well in excess of 1000 K. Several mechanisms may be responsible for the increased thermal stress near the exit of the inner wall of the thruster:

- **Electron Bombardment:** This is the section of the channel where the magnetic field strikes the channel wall for all magnetic configurations. Electrons deflected or caught within the magnetic field lines are likely to contact the ceramic wall here, depositing their energy as heat.

- **Ion Bombardment:** Near the channel exit, the ions have accelerated to the greatest velocities, suggesting ions impacting here have the greatest kinetic energy and thus the most potential for heating.

- **Ion Recombination:** Ions which leave the plasma and strike the wall absorb an electron and are reflected. This is en exothermic process, which deposits heat onto the wall.

- **Molecular Recombination:** The simultaneous three body collision of atomic oxygen forming a molecule is extremely unlikely to occur in the bulk of the plasma, yet at the plasma boundary the ceramic wall is able to easily donate the required electron, and absorb the energy released. Note this heating mechanism is not present in traditional xenon thrusters.

Note that the two recombination heating mechanisms are expected to be found throughout the thruster channel not just near the exit plane. However, even if the entire walls were being heated in a homogeneous way, we would still expect the tip of the inner wall to get the hottest: this point is further from the radiator, meaning that heat deposited here must conduct through the entire axial length of the channel ceramic before being removed. The inner wall has 60% the area of the outer channel, and a smaller cross section to conduct heat away which further exacerbates the problem.

The extreme heating of the ceramic walls is a clear indicator that much energy is being lost here, as is reflected in the poor thrust efficiency of the WET-HET. It is at this stage difficult to comment on what fraction of this energy loss is due to the molecular nature of oxygen, and what is due to the magnetic field of the WET-HET increasing wall interactions at the central pole. Future oxygen thrusters with more traditional magnetic field topologies will help to determine this.

The thermal shock due to the extreme thermal load at the tip of the inner channel wall was

assumed to be responsible for the original alumina inner wall fracturing. The boron nitride replacement did not succumb to fracturing. An unanticipated thermal limit was found due to the inner magnetic pole being heated by direct thermal contact with the ceramic. This issue was expected and lead to the inclusion of a 1 mm air-gap between the inner channel wall and the magnetic pole piece. In reality, machining tolerances, mechanical sagging, and wear caused the inner pole to contact the ceramic, and rapidly heat up. This is presumed to have caused a major issue, as the heating of this pole reduced its magnetic permeability and ultimately resulted in the thermally triggered voltage drop phenomenon that was witnessed (see section 6.2). This thermal bottleneck ultimately restricted operation at high powers or high magnetic field topologies to very short periods. It is however important to note that this is considered a flaw of the WET-HET itself, but not necessarily a flaw of all water electrolysis thrusters. This issue can presumably be mitigated somewhat in future thrusters with better thermal design, and shorter discharge channel lengths.

7.8 Erosion

The tip of the inner ceramic wall, meaning the section directly before the channel exit, was found to erode significantly over the course of the test campaign. The outer ceramic wall shows no measurable degradation, although it was exposed to the oxygen plasma for the same period of time. The discrepancy in degradation between the inner and outer walls suggests that the high reactivity of oxygen cannot be the major cause of erosion of the channel walls. Of the different heating mechanisms listed in section 7.7, electron bombardment and ion bombardment are the two mechanisms which are assumed to be the most severe near the channel exit, and therefore more likely to be driving erosion. Ions tend to be accelerated orthogonally to the magnetic field lines. Near the channel exit in figure 4.7 the field lines are not perfectly radial (horizontal in the diagram), but slope downwards towards the centre line. Ions being accelerated orthogonally to these field lines would follow a trajectory towards the centre line. This suggests that ions

would be impacting the inner ceramic wall more than the outer ceramic wall, something which is consistent with the erosion being much more prominent on the inner wall than the outer. This also helps to explain a further reason as to why the internal wall is heating more than the outer wall.

The copper anode and the stainless steel anode cap both show clear surface oxidation. Figure 4.12 shows the anode and cap before any testing, while figure 6.12 shows the same components after firing. Although the surface of the copper has clearly been heavily oxidised, this has not greatly affected performance. Copper was selected as being one of the only materials able to sufficiently conduct away the high anode thermal loads originally simulated. These simulations assumed 100% of the input energy was incident onto the anode (see section 4.7). Given the anode face did not exceed the maximum temperature of the stainless steel anode cap, and the majority of the heat was seen to be deposited to the channel walls, it has become obvious that other metals less prone to oxidation would make better anodes.

While erosion was witnessed on both the ceramics and the anode of the WET-HET, it was not as severe of an issue as originally anticipated. While the ceramics did show aggressive erosion, the disparity in degradation at different locations suggest that the ion acceleration was the main driver, not the reactivity of the oxygen plasma. This is a significant result, as it suggests that the aggressive erosion maybe not be intrinsic to all oxygen HETs, but instead maybe be due to imperfections of the magnetic topology of the WET-HET.

As expected, the copper anode did react with oxygen, yet this did not appear to effect performance. These results suggest that erosion in future oxygen HETs will be significantly less severe if oxidation resistant metals are used, and care is taken to ensure that the magnetic field lines are parallel to the exit plane.

7.9 Comparison to Similar Thrusters

Let us compare the performance of the WET-HET to that of other HETs operating on molecular propellants in the literature. Table 7.1 compares the WET-HET to two traditional xenon HETs which were operated on oxygen, namely the work of Nakagawa et al.[51] and the work by the Moscow Aviation Institute (MAI) which is referenced in Frisbee et al.[50]. Also included in Table 7.1 is the Extended Channel Hall Effect Thruster (ECHT) from Marchioni and Cappelli[60] and a PPS-1350 operating on a 1.27:1 ratio of nitrogen to oxygen[58]. These two last thrusters were experimental campaigns performed in the context of RAM-EP. Finally a SPT-100 operating on xenon is included as a reference[12]. To remain consistent with the definitions used throughout this text, specific impulse and thrust efficiency terms in Table 7.1 have been calculated using the power and mass flow rate of the anode only as described in equations 6.1 and 6.3.

Table 7.1: Comparison of the WET-HET to similar Hall effect thrusters. Three operating points are shown for the WET-HET. The values marked with and asterisk (*) indicate values which carry a large uncertainty due to a error in the discharge voltage measurement (see section 6.2), so do not represent steady state.

Parameter	**Units**	**WET-HET**			**Nakagawa**[51]	**MAI HET**[50]	**ECHT**[61]	**PPS-1350**[58]	**SPT-100**[12]
Propellant	-	O_2			O_2	O_2	N_2	$1.27N_2$:O_2	Xe
Dimensions (L x W x r_o)	mm	35x5x25			12x5x60	-	87x10x100	-	42x16x100
Mass flow rate	mg/s	0.958			1.16	1.22	2.0	2.6	5.12
Volume flow rate	sccm	40.21			48.71	51.22	97	118	52
Discharge Voltage	V	277	360	360*	350	360	200	350	300
Discharge Current	A	3	8	10	1.8413	2.3	2.6	3.5	4.5
Discharge Power	W	831	2880	3600*	644	828	520	1225	1350
Thrust	mN	10.5	29.0	38.6	11.6	11.5	22.3	26	85.7
I_{sp}	s	1114	3090	4112	1000	986	1138	1019	1706
Anode Thrust Effiency	%	6.9	15.3	24.3*	9	6.6	24.0	10.6	53.1

7.9.1 The Work of Nakagawa et al.

In their 2003 work Experimental Investigation of a Hall Thruster Using Oxygen as the Propellant, Nakagawa et al.[51] test a HET designed for xenon with a channel diameter of 60 mm, width of 5 mm and an adjustable channel depth. This work was only published in Japanese, and thus the discussion is based off an automated translation of the original paper and may not be 100% accurate. This study, much like the Frisbee et al.[50] paper suggest that a HET could utilise oxygen produced in-situ from lunar regolith as a propellant. The same thruster was operated on oxygen and xenon, as thrust data and plume measurements were taken. For oxygen, the maximum anode thrust efficiency was found at 6.6%, with an I_{sp} of 1000 s and a thrust of 11.6 mN. Noteworthy is that table 7.1 shows a very comparable performance of the WET-HET, the Nakagawa study, and the MAI HET which is referenced in Frisbee et al.[50] when each device operates near 1.1 mgs^{-1}. At this mass flow rate all three oxygen HETs generate approximately 11 mN of thrust and 1050 s of I_{sp}. The channel dimensions of the thruster tested at the MAI was not reported. The WET-HET has a deeper channel, but a considerably smaller outer radius than the thruster used by Nakagawa et al.[51], and generates a much greater current at a lower voltage. This suggests that at all powers, the WET-HET is extracting considerably more ions from the oxygen, hence a considerably higher propellant utilisation efficiency at higher powers.

The Nakagawa et al.[51] study found similar trends to the WET-HET. In both thrusters the discharge voltage is very stable over a wide range of discharge currents, showing no plateau or performance "knee point" which is generally found in xenon thrusters. Both thrusters show that I_{sp} and thrust efficiency rise linearly with power over the entire range surveyed. Unlike the WET-HET, Nakagawa et al.[51] found that performance increased as mass flow was increased, however this was tested with voltage held constant, but for the WET-HET power was held constant. For oxygen an optimal channel depth of 12 mm was found in the Nakagawa et al.[51] study with performance decreasing either side. For the WET-HET such a maximum has not yet

been found, however in comparison the performance has shown to be significantly less sensitive to channel length. Both studies also found that a ceramic of lower secondary electron emission produced greater performance, with the Nakagawa et al.[51] suggesting that replacing BN channels with Boron Nitride-Aluminum Nitride ceramic composites (AlN-BN) would reduce the SEE further, and therefore improve propellant utilisation.

A probe was used in the Nakagawa et al.[51] study to determine electron temperature of the oxygen plasma 4 mm from the anode. The temperature was found to increase linearly with discharge voltage, from around 7.5 eV at 200 V discharge voltage, to 22 eV at 350 V. These data would have proven very helpful for PlasmaSim modelling, where a homogeneous temperature of 16 eV overestimated performance, as can be seen in figure 6.6.

7.9.2 The Extended Channel Hall Effect Thruster

To the knowledge of the author, the ECHT of Marchioni and Cappelli[60] is the only other HET which has been developed specifically to be optimised for operation with a low mass molecular propellant. The ECHT was operated on nitrogen, yet the similarities of atomic mass, diatomic nature, and ionisation cross section of nitrogen and oxygen allow us to draw some comparisons between the thrusters. Both the ECHT and the WET-HET share unusually deep channels to promote the utilisation of molecular propellants, and a magnetic depth which extends down the channel. The ECHT has a deeper channel than the WET-HET at 86 mm, double the channel width, and four times the outer diameter. The ECHT has almost 9 times the cross-sectional area of the WET-HET, yet operates at a considerably higher anode thrust efficiency. The two thrusters differ in that the ECHT operates at greater anode thrust efficiencies as discharge power is reduced, where as the opposite is true for the WET-HET. The WET-HET data point that is

most similar to the ECHT set-point shown in table 7.1, is the WET-HET thrust measurement TM190 (which is not shown in table 7.1). Here the WET-HET has a discharge current of 2.5 A, a discharge voltage of 247 V which results in an I_{sp} of just 826.1 s and anode thrust efficiency of only 5%, while operating at 0.985 mg/s. The larger ECHT operates at the same current, but twice the mass flow rate to produce almost five times the thrust efficiency. Considering that at this set point the ECHT is producing a higher thrust and I_{sp}, it can be deduced that while the same number of ions are being generated in both thrusters, the ECHT is accelerating the average ions to far greater velocities and doing so more efficiently.

Marchioni and Cappelli[60] plan on conducting more tests with the ECHT after increasing the depth of the ionisation channel, the thickness of the magnetic region, and the magnetic field strengths above their current maximum near 100 Gauss. They recognise that after a particular channel length, power losses to the walls will hurt performance, this was also true for the WET-HET for intermediate lengths, but at deep channels the performance began to increase again. For the WET-HET increasing the magnetic region decreased performance, yet this may not be the case for a thruster with such a different channel cross sectional area. In the current study the same conclusions as Marchioni and Cappelli[60] were found in regards to the high energy densities required for high performance with molecular propellants resulting in considerable engineering challenges.

7.9.3 The PPS-1350 Operating on Atmospheric Propellants

The PPS-1350 was designed to operate on xenon, and gets its name from the french translation for a stationary plasma thruster (Propulseur a Plasma Stationnaire) with a nominal operational power of 1350 W[135]. This thruster was developed by Groupe Petite Propulsion et Equipements

(now Safran) in collaboration with the Fakel which is the Russian company responsible for the highly successful SPT-100. The purpose of the PPS-1350 was to clone the SPT-100 to meet the standards of the European market and increase thruster lifetime[135;59]. The PPS-1350 is famous for being the first European electric propulsion system to fly, successfully sending ESA's flagship SMART-1 mission to the moon[136]. The PPS-1350 and SPT-100 thrusters share an almost identical geometry and operational set point.

Cifali et al.[58] undertook a campaign to test how a thruster designed for xenon would operate on atmospheric propellants for the purpose of a RAM-EP system. They operated a PPS-1350 on N_2, and also on a mixture of N_2 and O_2. Experiments on pure oxygen were omitted due to concerns that the environment would be too oxidising for the anode. Given the similarities between oxygen and nitrogen, and the similarities between the PPS-1350 and the SPT-100, we consider this a good analogy of how an SPT-100 would operate on oxygen. The set point of the PPS-1350 included in table 7.1 was selected as this was the best performance found for this propellant[58]. Comparing the PPS-1350 operating on atmospheric propellants to the similar SPT-100 on traditional xenon, it can be seen that peak performance was found when the PPS-1350 was operated at a higher voltage, but a lower power than with xenon, and at half the mass flow rate. Note that the lower mass flow rate still results in twice the volume flow rate, suggesting a considerably greater neutral density is achieved. In comparing the PPS-1350 with the 2880 W WET-HET data point, one can clearly see that a traditional HET can be shown to operate at a lower anode thrust efficiency, and considerably lower I_{sp} than a unmodified thruster.

The comparison between the WET-HET and the PPS-1350 operating on a mixture of N_2 and O_2 is considered to be an important result: table 7.1 shows that the WET-HET has successfully been optimised to operate on atmospheric propellants better than a traditional HET. By operating at a greater power yet a considerably lower mass flow rate of molecular propellants, the WET-HET is able to sustain much higher discharge currents than an un-optimised thruster,

which ultimately results in an efficiency increase of at least 47%, and a specific impulse improvement of 403% over the PPS-1350.

7.10 Feasibility of a Water Electrolysis Hall Effect Thruster

The purpose of this project was to assess whether operating a HET on the products of water electrolysis could become a competitive alternative to traditional xenon HET systems. The full characterisation of the WET-HET has provided experimental evidence of the performance that can be expected from a thruster optimised for this purpose. We will continue to consider the SPT-100 as a benchmark for a traditional xenon thruster of a similar power class, with typical operational data shown in table 7.1. Note that although the SPT-100 operates at only 1350 W, the I_{sp} and efficiency is representative of a high power xenon thruster. Consider, for example, the performance of a 5 kW prototype HET from Biagioni et al.[137]. When operated at 10 mgs^{-1} mass flow rate, and a discharge power of 3000 W this thruster produced approximately 175 mN of thrust at 1750 s I_{sp} which results in a anode thrust efficiency of 51 %. Similarity in efficiency and I_{sp} across xenon thrusters of different powers shows that we can compare the SPT-100 performance to the WET-HET even if the latter has a wider power range.

Generally speaking, the WET-HET produces approximately half the thrust efficiency, yet twice the I_{sp} of a comparable xenon HET. Note that this is similar to the initial approximations performed using very basic scaling assumptions in section 2.4. Judging purely by thruster performance, the WET-HET is unable to compete directly with traditional xenon thrusters as a result of the poor anode thrust efficiency. The energy loss incurred by trying to ionise a low-mass molecular propellant with a small ionisation cross section dictates that only 15-35% of the electrical energy is directly converted to thrust, unlike the 53.1% achievable for a traditional system. This implies a thruster of such low efficiency would either force a spacecraft operator to choose between an unusually low thrust or a considerably larger solar array. Ultimately a

trade-off must be made, in which the propellant mass saving of increasing the specific impulse by a factor of 2.4 outweighs the penalty of requiring 5.8 times the power to generate the same thrust.

It is however not fair to simply compare the WET-HET to traditional thrusters based on performance metrics alone. From the beginning of this project, it was expected that the thruster performance would not be able to compete with xenon, which is the propellant that has enjoyed an unwavering monopoly in the field for over 40 years. Moreover, the WET-HET represents the first iteration of a proof-of-concept thruster, while the SPT-100 has become a standard due to successful experimental optimisation.

Ultimately, the main motivation for the water electrolysis HET concept is not to outperform a xenon HET, but to produce a system in which the impact of a decrease in thrust efficiency can be accepted to exploit the numerous system-level benefits provided, namely:

- Dramatic reduction in raw material cost compared to xenon,
- Extremely high specific impulse for a HET,
- Enables the possibility for water to be utilised as a shared propellant multi-mode chemical-electrical propulsion system,
- Moderate storage density which greatly surpasses krypton,
- No storage pressure restrictions allowing for thin, low mass storage tanks,
- Non-toxic, safe to handle,
- Fully passivated propellant at launch reduces cost, risk, and complexity of integration,
- Cathode compatibility provided by stociometrically ideal hydrogen production,
- Synergy with fuel cell energy storage systems,

- Future ISRU potential for lunar, mining and planetary missions.
- Significantly lower carbon footprint than xenon or krypton missions.

These unique system level benefits make it extremely difficult to compare a water electrolysis HET system to current or proposed electric propulsion systems. Importantly, the performance of the WET-HET was not so poor as to dismiss the concept. It is possible to envision that market trends outside the propulsion community could tip the scales in favour of water electrolysis electric propulsion systems. Among these trends could be the cost requirements of the NewSpace movement, the adoption of chemical water electrolysis systems, or the adoption of fuel cell energy storage systems.

We can conclude that the reduced anode thrust efficiency, increased erosion rate, or thermal limitations of the WET-HET are not damaging enough to be considered "show-stoppers" for the water electrolysis HET concept. The exceptionally high I_{sp} (which has yet to show a plateau), added with the numerous system level benefits suggests there may very well be a use case for such a system. Ultimately, this project has failed to identify any fundamental problems which lie on the path to a water electrolysis HET architecture becoming competitive, and must therefore conclude that this technology has demonstrated an intrinsic level of feasibility.

7.11 Recommendations For Future Water Electrolysis HETs

This section details some of the lessons learned during the WET-HET design and operation, and provides suggestions for the design of future water electrolysis HET systems.

The discharge channel should be constructed of a low SEE material such as boron nitride or a boron nitride-aluminum nitride ceramic composite. The channel depth should be very shallow, 35 mm or shorter if using a similar channel cross section. The magnetic region shall be as thin

as possible and near the channel exit, much like in a traditional xenon thruster. Special care should be taken to ensure that the radial field lines are highly orthogonal to the thruster axis, so that ions are not accelerated into the ceramic walls. The magnetic field should generate a magnetic field which peaks at the centre line of the channel at around 400-600 Gauss. Special care should be taken to ensure that sufficient insulation is placed between any component at the anode potential and other conductors. If possible, the anode should be constructed from alloys that are resistant to oxidation. Gas-tight interfaces between ceramics and metallic components shall be avoided where possible, particular for areas of high heat flux.

Very great emphasis should be placed on thermal modelling. Most of the thermal load should be assumed to be transferred to the walls of the discharge channel near the exit. The thermal limit of the thruster should be of paramount importance, as this determines the maximum steady state discharge power. The maximum discharge power will ultimately dictate the maximum thrust, I_{sp}, thrust efficiency, and propellant utilisation efficiency of the device. Special care should be taken to ensure that thermally sensitive magnetic components are adequately isolated from the high temperature components. Mass flow rates as low as $0.985\,\mathrm{mgs}^{-1}$ are sufficient at producing discharge powers in excess of 3 kW, with theoretical limits suggested to be closer to 6 kW. These discharges demand power electronics that can sustain uncharacteristically high currents for HETs at this mass flow rate of over 10 A.

7.12 Planned Future Works

7.12.1 Full System Integration

The ultimate goal of the WET-HET is to test the feasibility of operating a HET on water electrolysis, meaning the anode operates on oxygen, with the stoichiometric ratio of hydrogen supplying the cathode. Future experiments will test such a configuration with a LaB_6 cathode

more representative of flight hardware operating on hydrogen propellant.

7.12.2 Plume Measurements

Previous simulations have suggested that the ratio of diatomic ions (O_2^+) to monatomic ions (O^+) is a function of the residence time of the particles and could therefore be altered by tuning the depth of the channel[138]. Future plume measurements of the WET-HET will determine whether the ion ratios are dependent on channel depth, and if this is one of the factors leading to the discrepancy in performance.

Chapter 8

Summary and Conclusions

The limited availability and extreme cost of xenon has caused many to search for an alternative propellant for electric propulsion systems like Hall effect thrusters and gridded ion engines. While several of these alternative propellants are currently being considered, each demands a significant compromise: be it storage volume, toxicity, or incompatibility with current cathode technology[2]. This study proposes a novel electric propulsion architecture which effectively allows water to be consumed as a propellant for a HET. By utilising an electrolyser to split liquid water into gaseous hydrogen and oxygen, the thruster consumes only oxygen, where as the cathode is supplied with only hydrogen. Such a system is termed a water electrolysis HET system. Historically water has been dismissed as a potential propellant due to it aggressively poisoning the inserts critical to cathode operation. In a water electrolysis HET system the high storability of water can be exploited, but the cathode poisoning issue is mitigated by having the cathode operate on hydrogen: one of the few gases that has demonstrated great cathode compatibility[17;39;16;40]. Water provides system level advantages such as low price, low toxicity, moderately high density, ISRU potential, and practically no storage pressure requirements. However, the main system level motivation for a water electrolysis HET is that such a system can share hardware with a water electrolysis chemical propulsion system, the likes of which are currently enjoying much research attention[66;47;65]. The combination of these technologies

allows for a single spacecraft to utilise water for high thrust chemical burns, and also for high specific impulse electric burns in a shared propellant multi-mode chemical-electrical propulsion architecture. Although theorised, such a system has never been flown[67].

The Water Electrolysis Hall Effect Thruster or WET-HET was created to investigate the feasibility of optimising a HET for oxygen operation. A simplified quasi-zero-dimensional version of the PIC code called PlasmaSim was developed such that the channel dimensions and operating set-points of the WET-HET could be numerically optimised. The results of the optimisation dictated that the WET-HET would have a channel width of 5 mm, an outer channel circumference of 25 mm, and operate on an oxygen mass flow rate of 1-1.4 mgs^{-1} at discharge powers between 1000 kW and 2000 kW. The channel depth of the WET-HET was to be experimentally tuned in the laboratory to account for complex physics which the simulation could not replicate, such as internal energy states and dissociation of the oxygen molecule. Similarly the magnetic topology of the thruster was to be experimentally optimised due to the assumption that the low electron-impact cross section of the low mass oxygen molecule would favour an exotic axially thicker, but weaker magnetic field topology over a traditional one. The thruster was constructed in house after iterative magnetic, thermal, and mechanical modelling converged to the final design.

The WET-HET was tested in the Boltzmann vacuum facility, with the resulting thrust being directly measured by a hanging pendulum style thrust balance which was also developed in house. A full characterisation of the WET-HET was undertaken to investigate the performance sensitivity to all independent variables. The WET-HET was tested at discharge powers as low as 380 W, to over 3000 W, which resulted in discharge currents of 1.5 A to 10 A on oxygen mass flow rates ranging from 0.985 mgs^{-1} to 1.845 mgs^{-1}. Three different channel depths were tested being 35 mm, 45 mm and 60 mm, as well as three different magnetic topologies. These topologies

include a thin magnetic region near the channel exit which resembles a typical HET topology, a configuration in which the magnetic field envelops half the channel depth, and a topology where the magnetic field covers the entire channel. Each magnetic configuration was tested at various field strengths in the range of 225 Gauss to 710 Gauss. Two channel wall materials were investigated: alumina and boron nitride. Safety restrictions demanded krypton replace hydrogen as the working gas of the cathode for all thrust measurements.

The greatest specific impulse and anode thrust efficiency of the WET-HET was found when the device was operated at a high power and a low mass flow rate. A magnetic field strength of around 600 Gauss was optimal for the thin magnetic topology, with the optimal strength being closer to 400 Gauss for all other magnetic configurations. It is assumed that above the optimal magnetic field strength the ions become magnetised, due to a rise in discharge voltage being observed but a decrease in thrust, I_{sp} and thrust efficiency. Contrary to expectation, increasing the thickness of the magnetic region had very little impact on the discharge voltage at a given discharge power. Thrust, I_{sp} and thrust efficiency decreased as the thickness of this region was increased, with a thin magnetic region near the exit of the thruster producing the best performance, a trend which is also seen in xenon thrusters. Another trend that the WET-HET shares with a traditional xenon thruster is that boron nitride discharge channel walls offer superior performance to alumina by almost 40%. This is assumed to be due to the lower secondary electron emission yield of boron nitride. Of the different channel depths tested a depth of 35 mm demonstrated the greatest thrust, I_{sp} and efficiency at a given power. This was unanticipated, as deeper channels were expected to extended neutral residence times and thus increase ionisation rate in an attempt to combat the low electron impact ionisation cross section of oxygen. The 35 mm channel was the shortest tested, however it is expected that even shorter channels will provide better performance yet. The greatest performance of the WET-HET was found when operating at a mass flow rate of $0.985\,\mathrm{mgs}^{-1}$, with a thin magnetic region which peaked at 640

Gauss and a channel depth of 35 mm. At a discharge current of 10 A the thruster performance was the greatest tested producing 38.63±0.25 mN thrust at an I_{sp} of 4112±36 s. While thrust and I_{sp} were measured with high confidence for this data point, issues with the DAQ for these measurements put a large uncertainty on the discharge voltage, discharge power and thrust efficiency here at 320±75 V, 3200±750 W and 24.35±5.89 % respectively. The greatest performance recorded without the discharge voltage error was at the same mass flow rate, channel depth and magnetic topology, but at a discharge current of 8 A, a discharge power of 2880 W, which produced a thrust of 29.03±0.15 mN, I_{sp} of 3090±16 s and anode thrust efficiency of 15.28±0.15 %. All trends suggest that greater performance can still be attained at greater discharge currents, however this would most probably surpass the thermal limits of the device. While the inner ceramic channel wall showed clear signs of erosion, the outer wall did not, suggesting that some, if not all of the degradation may be due to electron or ion bombardment, but not the reactivity of the oxygen plasma. However, precise erosion measurements over longer firing periods are needed to confirm this.

When compared to a traditional xenon HET, the WET-HET operates at approximately half the thrust efficiency, yet twice the specific impulse. It is however important to note the WET-HET is a first proof of concept and breadboard type of thruster, so a direct comparison is somewhat unfair. A significant decrease in thrust efficiency was expected, and is assumed to be due to oxygen having a smaller molecular mass than xenon, a reduced electron impact ionisation cross section, and additional energy sinks provided by the covalent bond. The WET-HET is able to reach a comparable propellant utilisation efficiency as a xenon thruster, yet it requires a greater power. This is presumably because at these high powers the electron temperature gets so great that the ionisation rate of oxygen rivals that of a traditional thruster.

The full characterisation of the WET-HET has provided sufficient information for the feasi-

bility of the water electrolysis HET concept to be evaluated. Even with the higher I_{sp}, the low thrust efficiency of the WET-HET is an intrinsic characteristic of the oxygen propellant, and will likely never allow a water electrolysis thruster to directly compete with a traditional xenon thruster based on performance alone. This was however the predicted outcome, and does not necessarily deem the water electrolysis HET system noncompetitive. The performance of the WET-HET, particularly at high powers has demonstrated a I_{sp} not typically seen in HETs. When this performance is considered alongside the system level benefits that water provides as a propellant, namely, a low price, great storability, and low toxicity, the water electrolysis HET concept does not appear to share any of the major engineering challenges facing other alternative propellants. The potential to share hardware with either a chemical water electrolysis chemical propulsion system, or a regenerative fuel cell energy storage system suggests any future development and adoption of these technologies would greatly strengthen the competitiveness of the water electrolysis HET concept. Neither the relatively low thrust efficiency nor the channel wall erosion is considered a significant problem to stop further development. At this stage no fundamental hurdles of the water electrolysis HET concept have been identified.

Future planned work is a full system level test, in which the WET-HET is operated on oxygen while a LaB_6 cathode supplied with $1/8^{th}$ the mass flow rate of hydrogen. Upgrades to the safety systems in the laboratory have now made hydrogen cathode testing possible, and a LaB_6 cathode is currently in development. Plasma probes have been manufactured such that plume measurements of the WET-HET can be undertaken such that electron temperature and ion species ratios can be measured to gain a better understanding of the underlying physics, and aid future simulations.

Publication History

This list includes all the journal articles and conference papers which I have authored relating to the thesis. The list is in chronological order. Journal articles are in bold.

- Schwertheim, A. and Knoll, A., 2018. *In situ utilization of water as a propellant for a next-generation plasma propulsion system.* In 6th Reinventing Space Conference (Vol. 1, p. 2018).

- **Schwertheim, A. and Knoll, A., 2019. *In situ utilization of water as a propellant for a next-generation plasma propulsion system.* JBIS J. Br. Interplanet. Soc., 72(1), pp.7-11.**

- Schwertheim, A. and Knoll, A., 2019, September. *The Water Electrolysis Hall Effect Thruster (WET-HET): Paving the way to dual mode chemical-electric water propulsion.* In 36th International electric propulsion conference (p. 259).

- Schwertheim, A. and Knoll, A., 2020. *The Water Electrolysis Hall Effect Thruster.* In 71st International Astronautical Congress (Vol. 1, p. 2020).

- Schwertheim, A. and Knoll, A., 2020. *Performance characterisation of the water electrolysis Hall effect thruster WET-HET using direct thrust measurments.* In Space Propulsion Conference (Vol. 1, p. 2021).

- **Schwertheim, A. and Knoll, A., 2021. *Low power thrust measurements of the water electrolysis Hall effect thruster.* CEAS Space Journal, pp.1-15.**

- **Schwertheim, A., Rosati Azevedo, E., Liu, G., Bosch Borràs, E., Bianchi, L. and Knoll, A., 2021.** ***Interlaboratory validation of a hanging pendulum thrust balance for electric propulsion testing.*** **Review of Scientific Instruments, 92(3), p.034502.**

- **Schwertheim, A. and Knoll, A., 2022.** ***Experimental investigation of a water electrolysis Hall effect thruster.*** **Acta Astronautica, 193, pp.607-618.**

- Tejeda, JM and Schwertheim, A., 2022. *Water as an Environmentally Friendly Propellant for a Multi-functional Spacecraft Architecture.* In Space Propulsion Conference (Vol.1, p. 2022).

Appendix A

Appendix

A.1 PlasmaSim 0D Sensitivity to Starting Parameters

This section measures variability between different PlasmaSim 0D simulations with identical set-ups, but different random seeds. The simulation is also run for different time-steps and super particle numbers to justify the values used for thruster optimisation. For the following tests the default simulation inputs are were taken from an SPT-100 operating on xenon, as discussed in section 3.5:

- Simulation Time = $1\text{x}10^{-3}$s
- $x_{max} = 42\,\text{mm}$
- $y_{max} = 263.9\,\text{mm}$
- $z_{max} = 16\,\text{mm}$
- Electron Temperature = $11.5\,\text{eV}$
- Ion Temperature = $0.1\,\text{eV}$
- Mass Flow Rate = $5.5\,\text{mgs}^{-1}$
- Axial Electric Field Strength = $7750\,\text{Vm}^{-1}$
- Initial Neutral Number density = $1\text{x}10^{19}\,\text{m}^{-3}$
- Initial Ion/Electron Number density = $1\text{x}10^{17}\,\text{m}^{-3}$

2

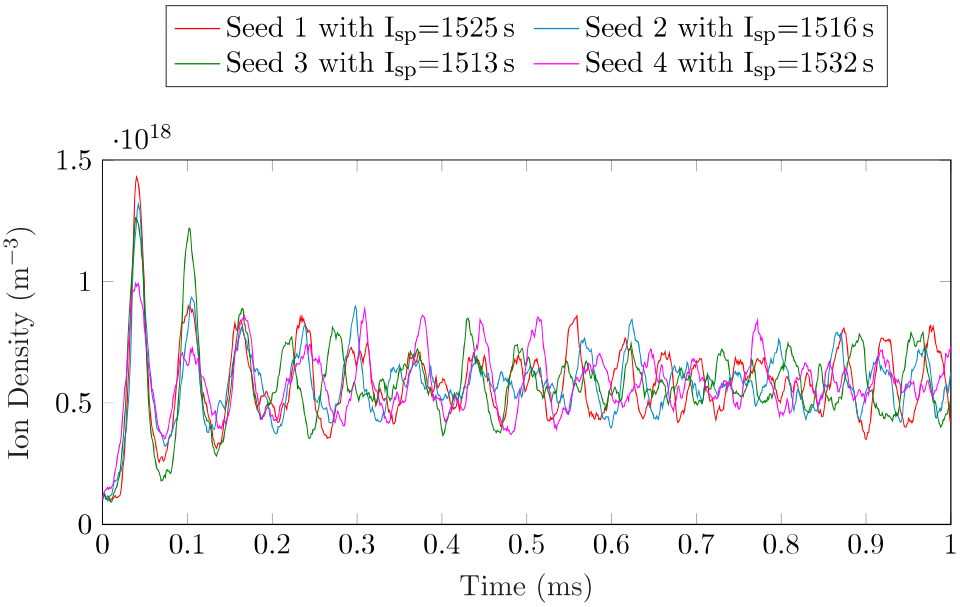

Figure A.1: The sensitivity of PlasmaSim 0D to changes random seed value.

Random Seed Sensitivity

As with any experiment, PlasmaSim 0D is required to be insensitive to uncontrollable initial conditions, with the same inputs producing the same outputs. Each time a new particle is generated, pseud-random numbers are generated to determine the position and velocity from a pre-allocated range. To ensure that a particular set of random numbers does not greatly impact the performance of the simulation, identical experiments with the only difference being the random seed were repeated. PlasmaSim was operated with above default simulation inputs, 34000 super-particles, and a time step of 1x10^{-7}s. This simulation was executed four times, each time with a different random number seed. The time series of the ion densities of these simulations are figure A.1. Although the only difference between the simulations is the set of randomly generated numbers, the time evolution of the ion density quickly diverges. This is expected, due to the influence that the individual particle positions and velocities have on the number of particles that exit the domain, which in turn impacts global attributes like ion density. The mean I_{sp} for the second half of the simulation for each trend is shown in the legend. These I_{sp} values have a spread of 19.9 s, which is a difference of $\pm$0.65 % from the mean value. This is an important results, as it demonstrates that although the microscopic time evolution of the

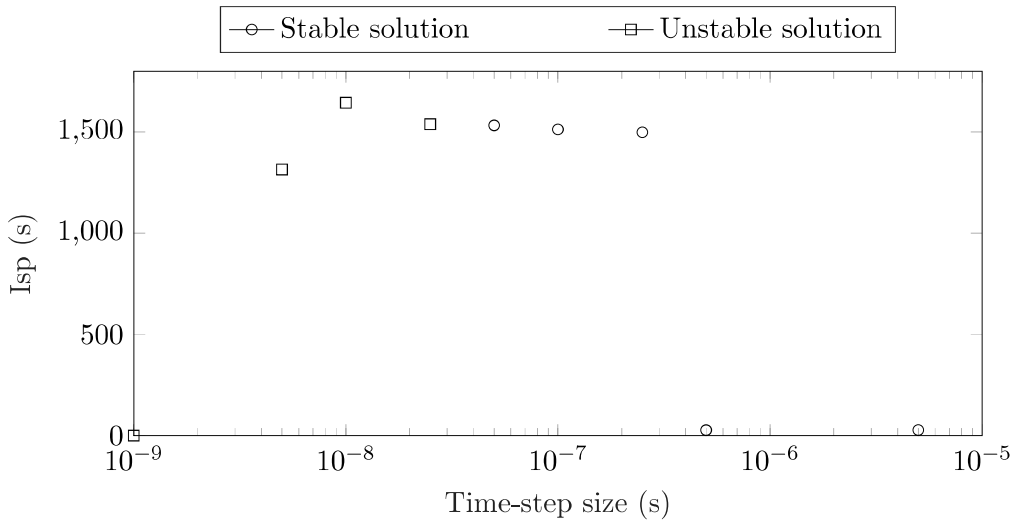

Figure A.2: The sensitivity of PlasmaSim 0D to changes in time-step size

plasma is highly dependent on the set of random numbers used, the macroscopic time averaged values are not. Given it is the macroscopic time averaged data that PlasmaSim 0D is optimising for, it can be concluded that PlasmaSim 0D is independent of random number selection.

Note that to ensure identical results, PlasmaSim 0D is often run with a particular random seed. This study has demonstrated that the choice in this seed does not affect the macroscopic results.

Time-step Sensitivity

When the size of the simulation time-step is decreased, the temporal resolution of the simulation increases, as does the execution time. Figure A.2 shows the mean I_{sp} of a variety of PlasmaSim 0D simulations that had the default simulation inputs and 34000 super-particles. The only difference was the time-step size. As one would expect, when the time-step increases too large the simulation fails to produce meaningful results. At time-step sizes great then $2.5\text{x}10^{-7}$ the seed ion population decreases almost instantaneously. This suggests at these timescales the particles leave the domain or recombine on the walls before they are able to generate sufficient ionisation to stabilise. More surprising is that when the step-size is reduced to $5\text{x}10^{-8}$ s or

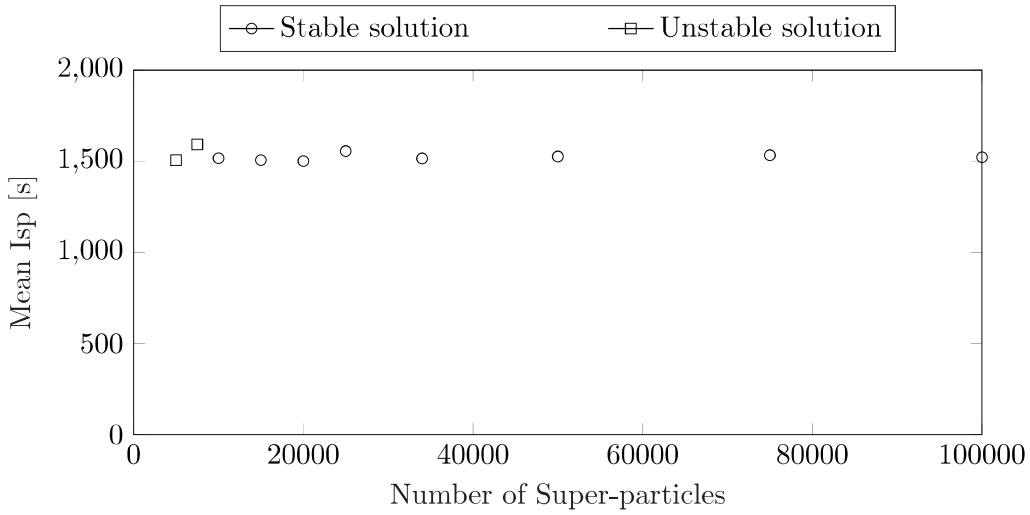

Figure A.3: The sensitivity of PlasmaSim 0D to changes in Super-particle number

below the simulation becomes less stable, and often causes the simulation to crash. This is most probably due to discretization errors in the particle numbers: at very high temporal resolution, in each time step only a handful of particles will recombine, or leave the simulation. Given that the I_{sp} or recombination rate is calculated by summing all these particles, these values become noisy if time-steps become too small. Based on the data shown in figure A.2, a time step of 10^{-7} s was considered appropriate for PlamsaSim 0D simulations.

Super-particle Number

The sensitivity of I_{sp} to the number of super-particles used is shown in figure A.3. These are again identical simulations with the default simulation inputs listed above, and a time-step of 10^{-7} s which differ only in number of super particles used. PlasmaSim 0D here appears to be very insensitive to super particle number, with the exception of those less than 10000, at which point the solution becomes unstable. This is better seen in the time-series of ion density shown in figure A.4. As the number of super particles increase, the average value of ion density remains constant, yet the amplitude of the density fluctuations decreases. This is somewhat expected, as reducing the number of super particles again introduces a discretization error, which sees the global metrics much more sensitive to individual particle motions. The large scale fluctuations

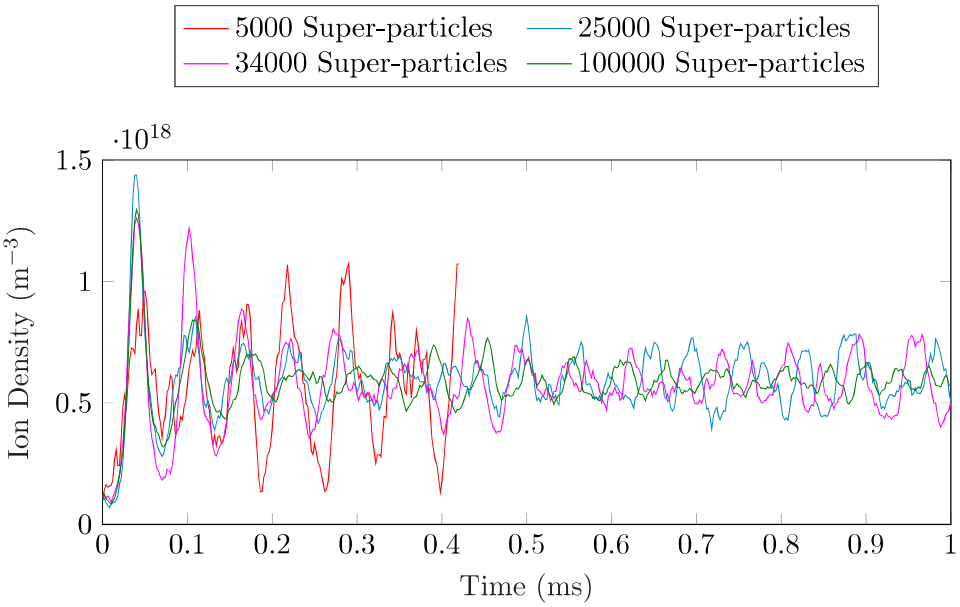

Figure A.4: Time series to ion density of PlasmaSim 0D simulations which differ by super-particle number

caused become large enough that the simulation crashes before the simulation time is complete, as witnessed in the 5000 super-particle trend. From these data 34000 super particles were selected to be a good balance between stability and execution time.

A.2 Full mathematical description of the In-situ Calibration System

The in-situ calibration system is shown in figure A.5, which is more detailed version of figure 5.9. The theoretical force (F_t) which the in-situ calibration system induces on the platform at point C, is derived in equation 5.1. Let us calculate the angles γ, θ and β:

$$\gamma = \arccos\left(\frac{H}{h}\right) \tag{A.1}$$

$$\theta = \arctan\left(\frac{B}{H}\right) - \gamma - \arccos\left(\frac{D^2 + h^2 - b'^2}{2Dh}\right) \tag{A.2}$$

$$\beta = \arcsin\left(\frac{H - h\cos(\theta + \gamma)}{b'}\right) \tag{A.3}$$

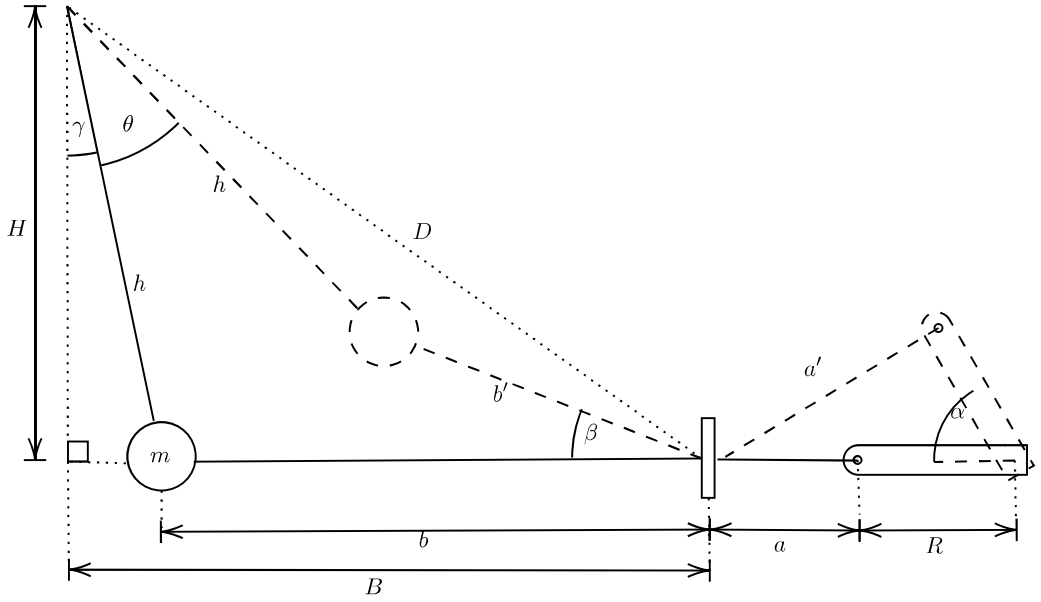

Figure A.5: Caption

for the measured lengths H, h, R, a and b and lengths B, D and b' which are defined as given by:

$$B = b + \sqrt{h^2 - H^2} \tag{A.4}$$

$$D = \sqrt{B^2 + H^2} \tag{A.5}$$

$$b' = a + b - \sqrt{(R\sin\alpha)^2 + (a + R(1 - \cos\alpha))^2} \tag{A.6}$$

A.3 Description of Power Supplies

This section provides further information on the power suppiles used in WET-HET testing which are shown in figure 5.14.

A.3.1 Heater Power Supply

The Elektro-Automatik EA-PSI 8080-60 DT power supply is also refered to in this text as the heater power supply. It is controlled manually and able to provide up to 60 A at 80 V. Figure 5.13 shows that the negative terminal is grounded, and the positive terminal is connected to one of the two filament cables. The opposite ends of each of the filaments were grounded to

the internal wall of the facility. When the cathode was powered, this power supply drew up to 30 A of electrons from the grounded facility, through the filament and out of the chamber. This raised the filament temperature to over 2000 K. After several hours at this temperature, the tungsten filaments were often prone to failing by physically breaking, leaving the circuit open and the cathode un-usable until the facility was vented and the filament replaced. The twin filament configuration shown in figure 5.13 was designed to reduce the time wasted repairing broken filaments: current was only drawn through one of the two filaments at a time, yet both filaments were connected to the feedthrough via their own dedicated leads. While one filament was heated, the other remained relatively cool. Upon the first filament breaking, the heater power supply was connected to the second filament lead at the air-side terminal of the feedthrough. From here the second fresh filament could now be powered, and testing resumed. This effectively doubled the average time between cathode failure. We represent the changing of the leads as a single pole dual throw switch in figure 5.13.

A.3.2 Keeper Power Supply

The keeper current was supplied by a Glassman PS/ET05R400-200 power supply referred to as the keeper power supply. Although this device was able to draw up to 400 mA at 5000 V, it was set to 100 mA and 100 V for the entire test. We can see from figure 5.13 that in the absence of the keeper discharge, the keeper circuit is open, with the negative terminal grounded and an air-gap between the keeper and the filament inhibiting current flow. Before cathode ignition, the keeper power supply operates in voltage controlled mode. The device switches to current controlled mode when the thermally emitted electrons from the filament ignite the 100 mA keeper current, which remains on for the entire duration of the test.

A.3.3 Anode Power Supply

Originally a single Sorenson SGA 1K0X5D-0AA power supply was used to ignite the WET-HET and sustain the discharge power. This instrument is referred to as the anode power supply, supplying up to 5 A at 1000 V. Such a high voltage was often required to ignite the plasma discharge between the keeper current, and the WET-HET anode. This current is called the discharge current. The keeper current can be considered analogous to a pilot flame in a furnace, which is first lit to help ignite and sustain the larger main flame. Ignition of the discharge current was often found to be difficult. Achieving this ignition required increasing the anode voltage to the maximum, turning off the magnetic current, increasing the heater current, and sometimes increasing the cathode and anode mass flow rates. Once this discharge is ignited, the power supply falls into current controlled mode, at around 120-200 V. This allowed the discharge current to be changed and the mass flow rates of the anode and cathode to be adjusted as needed. The ignition of the thruster was clearly observable from the brilliant bright light which can be seen emanating from the discharge channel from the view port of the facility.

To test the device at high powers the current limit of the anode power supply was too low. At the maximum current of 5 A the thruster would operate at around 1200 W. To test greater power levels, a second anode power supply was needed. A second identical Sorenson SGA 1K0X5D-0AA was procured, and the two were connected in parallel to double the available current, as shown in figures 5.12 and 5.13. A special paralleling cable provided by the manufactures allowed the two power supplies to be commanded as one, with their combined capability providing up to 10 A at 1000 V.

For the majority of the test the anode power supply was operated manually, yet 83 thrust measurements were taken with a dedicated "Thruster Control" software commanding a single anode power supply. This is described in more detail in section 5.5. The set voltage and set

current of the power supply can be remotely controlled by two 0-10 V analog signals. Similarly, two 0-10 V signals can be measured to safely determine the current and voltage output of the device. All of these signals are passed through a D-SUB 25 terminal on the back of the power supply. The IPPL ESB was modified such that it can be connected to the anode power supplies via a single D-SUB 25 cable, with the internal ESB DAQ monitoring and controlling the power supply output over the same USB cable as the thrust balance. This allows for very fast and precise output adjustments, accurate and fast data acquisition, and PID control to achieve constant power operation.

A.3.4 Magnet Power Supply

The final power supply shown in figure 5.13 is the GW INSTEK SPS-3610 which is referred to as the magnet power supply. With a maximum output of 10 A at 36 V, this power supply was used to generate the current which passes through the WET-HET coil to generate the applied magnetic field. This circuit is always closed, with the negative terminal being grounded, and the positive terminal passing through the coil before contacting the same ground return point on the inside of the chamber wall as the cathode filament. After thruster ignition the magnet power supply was switched on, and the current gradually increased in a current controlled mode. At a maximum of 3 A, the power supply dissipates approximately 80 W of the magnet coil. After thruster shutdown, the magnet power supply is shut off to mitigate unnecessary heating, and allow easy ignition of the thruster in the absence of a magnetic field. The presence of the applied magnetic field overlapping with the electric field generates the strong Hall current. This is witnessed in the laboratory by the plasma becoming considerably brighter, visibly more dense, and often a changing colour. The generation of the Hall current impedes electron mobility, which is evidenced by a sharp increase of anode discharge voltage at a given discharge current.

A.4 Thrust Measurement Procedure

The thrust measurement procedure followed to perform direct thrust measurements of the WET-HET is detailed below:

1. Install the WET-HET and cathode on the thrust balance and perform all necessary electrical and fluid checks.

2. Install the VCA calibration subsystem and close the facility door to eliminate air currents.

3. Perform the VCA calibration to establish the sensitivity of this particular experiment configuration. The software commands the sourcemeter to produce 13 currents which generated 13 test forces on the balance. These forces ranged from 0.7 mN to 172 mN, with each being applied for 30 seconds and followed by no force being supplied for 30 seconds. The balance sensitivity can be read from the best fit of the calibration force, as shown in figure 5.10.

4. Perform the in-situ calibration system baseline calibration. This involves the servo motor steping from 0° up to 60° and back down in 10° increments. Each step is held for 30 seconds, before resting at 0° for 30 seconds.

5. Remove the VCA subsystem carefully. The balance is now considered calibrated, so any further physical manipulation requires re-calibration.

6. Close the facility door. Ensure all hand valves are closed, and begin the roughing pump.

7. Wait for the pressure to drop below $1\text{x}10^{-1}$ mbar (approximately 60 minutes) and begin the all three turbo pumps.

8. Wait for the pressure to drop below $1\text{x}10^{-5}$ mbar (approximately 30 minutes) and begin the cryo compressor.

9. Wait for the cryo temperature to drop below 50 K.

10. Perform any electrical checks to assure the experiment has not been disturbed.

11. Evacuate the propellant lines all the way to the cylinder. This removes contaminating air from the propellant lines, and also ensures no leaks are present. Once the facility pressure has dropped to less than $1x10^{-5}$ mbar we can be confident the propellant lines are free of leaks, and all air has been evacuated or solidified onto the cryo panel.

12. Shut the MFCs, open the cylinders and pressurise the MFCs to 2 bar.

13. Perform a spot check with the in-situ calibration subsystem. This procedure is identical to the baseline calibration taken in step 4. Compare the baseline calibration response to the spot check response. Only proceed if the two are near identical, indicating the thrust balance calibration is unchanged and therefore valid.

14. Set the keeper power supply to 100 V, 100 mA. No current should flow as a result.

15. Begin the cathode mass flow. Typically 15 sccm of krypton.

16. Slowly increase the current of the heater power supply. The filament must be heated slowly to reduce thermal stress. Continue to increase the heater current until the keeper voltage drops to below 20 V. This indicates the keeper discharge has been established, and the cathode has successfully been ignited.

17. Begin the anode mass flow. Typically 50-80 sccm of oxygen.

18. Set the anode current limit to 2.5 A and slowly increase the voltage. No current should pass initially. The thruster should ignite when the voltage is between 250-1000 V. This is evidenced by a bright light emitted by the thruster and the anode power supply switching to current controlled mode and a potential of around 100-200 V.

19. Adjust the mass flow rate to the target mass flow rate if necessary.

20. Slowly raise the magnet power supply to the target magnet current. Typically 1.5 A. This will increase the anode voltage as electron mobility is impeded. The plasma will become more dense, brighter and possibly change colour. If the discharge is extinguished, shut down the magnet power supply and return to step 18.

21. Increase the anode current until the target discharge power is reached.

22. Allow the thruster to stabilise for at least 5 seconds.

23. Ensure the anode massflow rate, cathode massflow rate, magnetic current, and power are all correct.

24. Begin the laser data acquisition and start a timer simultaneously.

25. After 30 seconds, simultaneously switch off the anode power supply, magnet power supply, anode mass flow and cathode mass flow.

26. After 60 seconds, stop the laser data acquisition. This concludes the thrust measurement.

27. If further thrust measurements are required, begin both anode and cathode massflow rate again, and return to step 18.

28. If the heater current drops to zero, this suggest the first filament has failed. Turn off the heater power supply, test the filament, and switch the active filament if needed. Return to step 16.

29. Once all testing is complete or the test is aborted, turn off all power supplies, gas cylinders, and drain propellant lines into the facility.

30. Perform another spot check with the in-situ calibration system. Compare the response to that of the baseline. All captured thrust measurements are only valid if both spot checks before and after the measurement agree with the baseline, as this insures the sensitivity of the balance has not changed.

31. Shut down the cryo panel. The temperature of the panel will now rapidly rise, causing all solidified gases on the panel to sublimate as the temperature reaches each respective boiling point.

32. Wait for the cyro temperature to rise above 300 K (approximately 4 hours) and shut down the turbo pumps.

33. Wait for the turbo pumps to completely stop spinning (approximately 2 hours) and shut the roughing pump down.

34. Carefully vent the chamber and open the facility door.

A.4.1 Amendments to the Thrust Measurement Procedure for Hydrogen Tests

The thrust procedures for the hydrogen tests follow those outlined for krypton in section A.4 very closely, with a few exceptions:

- **In step 7 the hydrogen plena are vented and filled:** Once the facility had reached a pressure below $1x10^{-1}$ mbar the hydrogen MFC, and all solenoid valves were opened such that the hydrogen lines and plena were vented completely. Next the solenoid valves were shut. The following step was for the roughing pump to be shut down to eliminate potential ignition sources. The hydrogen cylinder was now brought into the laboratory, and used to fill the two hydrogen plena within the now evacuated facility to 2 bara. Next the hydrogen cylinder was removed from the laboratory, and the plena pressure monitored for leaks. If the pressure within the plena is stable, all three turbo pumps were switched on.

- **In step 11 the hydrogen lines need not be purged again:** given they were evacuated at rough vacuum, they have been checked to be free of leaks and contaminants. At this point the purge gas system was started and the associated purge gas failure alarm. Next the hydrogen MFC was pressurised by opening the solenoid valve that connects it to the

plena. The experiment could now be run like normal, with the MFC supplying hydrogen in the say way that it would for krypton.

- **In step 29 the hydrogen plena were emptied and vented before the cryo panel was switched off:** When testing had concluded all hydrogen was removed from the facility safely. This was done by slowly venting the hydrogen out of the MFC at the maximum flow rate of 200 sccm. Once the pressure in the plena was too low for the MFC to function, the final solenoid valve was opened to completely evacuate the plena. Once the pressure in the facility had returned to below $1\text{x}10^{-5}$ mbar, the experiment was fully passivated and safe. All solenoid valves and the hydrogen MFC were closed. Next we shut down the purge gas system.

References

[1] Robert L Sackheim and Robert K Masse. Green propulsion advancement: challenging the maturity of monopropellant hydrazine. *Journal of Propulsion and Power*, 30(2):265–276, 2014.

[2] Alex Kieckhafer and Lyon B King. Energetics of propellant options for high-power hall thrusters. *Journal of propulsion and power*, 23(1):21–26, 2007.

[3] David Y Oh, Thomas Randolph, Scott Kimbrel, and Manuel Martinez-Sanchez. End-to-end optimization of chemical-electric orbit-raising missions. *Journal of Spacecraft and Rockets*, 41(5):831–839, 2004.

[4] MR Drinkwater, R Floberghagen, R Haagmans, D Muzi, and A Popescu. Vii: Closing session: Goce: Esa's first Earth explorer core mission. *Space science reviews*, 108(1): 419–432, 2003.

[5] Haruo Kawasaki, Kazuya Konoue, Hirokazu Hoshino, Yutaka Kaneko, and Masanori Sasaki. Interim report of super low altitude satellite operation. In *IGARSS 2018-2018 IEEE International Geoscience and Remote Sensing Symposium*, pages 4066–4069. IEEE, 2018.

[6] Dan Lev, Roger M Myers, Kristina M Lemmer, Jonathan Kolbeck, Hiroyuki Koizumi, and Kurt Polzin. The technological and commercial expansion of electric propulsion. *Acta Astronautica*, 159:213–227, 2019.

[7] Dan M Goebel and Ira Katz. *Fundamentals of electric propulsion: ion and Hall thrusters*, volume 1. John Wiley & Sons, 2008.

[8] Mark Hutchins, Huw Simpson, and Javier Palencia Jiménez. Qinetiq's t6 and t5 ion thruster electric propulsion system architectures and performances. *Joint Conference of 30th ISTS, 34th IEPC and 6th NSAT*, 2015.

[9] John Snyder, Dan M Goebel, Richard R Hofer, James E Polk, Neil C Wallace, and Huw Simpson. Performance evaluation of the t6 ion engine. *Journal of Propulsion and Power*, 28(2):371–379, 2012.

[10] Richard Blott, Neil Wallace, Huw Simpson, and Mark Hutchins. Ion engine products and applications by qinetiq. In *57th International Astronautical Congress*, pages C4–4, 2006.

[11] Michael R Nakles, William A Hargus Jr, Jorge J Delgado, and Ronald L Corey. A performance comparison of xenon and krypton propellant on an spt-100 hall thruster (preprint). Technical report, AIR FORCE RESEARCH LAB EDWARDS AFB CA, 2011.

[12] John M Sankovic, John A Hamley, and Thomas W Haag. Performance evaluation of the russian spt-100 thruster at nasa lerc. In *23rd IEPC Conference; 13-16 Sep. 1993; Seattle, WA; United States*, 1994.

[13] Charles Garner, John Brophy, James Polk, and Lewis Pless. Cyclic endurance test of a spt-100 stationary plasma thruster. In *30th Joint Propulsion Conference and Exhibit*, page 2856, 1994.

[14] Kevin D Diamant, Raymond Liang, and Ron L Corey. The effect of background pressure on spt-100 hall thruster performance. In *50th AIAA/ASME/SAE/ASEE Joint Propulsion Conference*, page 3710, 2014.

[15] Zachary R Taillefer, John J Blandino, and James Szabo. Characterization of a barium

oxide cathode operating on xenon and iodine propellants. *Journal of Propulsion and Power*, 36(4):575–585, 2020.

[16] HE Gallagher. Poisoning of LaB6 cathodes. *Journal of Applied Physics*, 40(1):44–51, 1969.

[17] Dan Michael Goebel, JT Crow, and AT Forrester. Lanthanum hexaboride hollow cathode for dense plasma production. *Review of Scientific Instruments*, 49(4):469–472, 1978.

[18] Walter Guido Vincenti and Charles H Kruger. Introduction to physical gas dynamics. *Introduction to physical gas dynamics, by Vincenti, Walter Guido; Kruger, Charles H. New York, Wiley [1965]*, 1965.

[19] William R Kerslake and Louis R Ignaczak. Development and flight history of the SERT II spacecraft. *Journal of Spacecraft and Rockets*, 30(3):258–290, 1993.

[20] Sarah W Shark, Scott J Hall, Benjamin Jorns, Richard R Hofer, and Dan M Goebel. High power demonstration of a 100 kw nested hall thruster system. In *AIAA Propulsion and Energy 2019 Forum*, page 3809, 2019.

[21] M Yu Savinov, AM Arkharov, VE Poznyak, and VL Bondarenko. Development and creation of an efficient khrom-3 unit for preparing krypton-xenon mixtures. *Chemical and Petroleum Engineering*, 43(5):259–269, 2007.

[22] Michael Swiatek. Capturing and recycling of xenon from a cryopumped vacuum chamber. In *46th AIAA/ASME/SAE/ASEE Joint Propulsion Conference & Exhibit*, 2010.

[23] JM Tejeda, AH Schwertheim, and A. (pre-print) Knoll. Water as an environmentally friendly propellant for a multi-functional spacecraft architecture. In *2022 Space Propulsion Conference*, pages C4–4, 2022.

[24] Digiconomist. Home / Bitcoin Energy Consumption Index Bitcoin Energy Consumption Index. https://digiconomist.net/bitcoin-energy-consumption#assumptions, 2022. [Online; accessed 27-Feb-2022].

[25] Jesse A Linnell and Alec D Gallimore. Efficiency analysis of a hall thruster operating with krypton and xenon. *Journal of Propulsion and Power*, 22(6):1402–1418, 2006.

[26] James Szabo, Rachel Tedrake, George Kolencik, and Bruce Pote. Measurements of a krypton fed 1.5 kw hall effect thruster with a centrally located cathode. In *the 35th International Electric Propulsion Conference, Atlanta, Georgia, USA, IEPC-2017-26*, 2017.

[27] Junko Yamasaki, Shigeru Yokota, and Kohei Shimamura. Performance enhancement of an argon-based propellant in a hall thruster. *Vacuum*, 167:520–523, 2019.

[28] James Szabo, Mike Robin, and Vlad Hruby. Bismuth vapor hall effect thruster performance and plume experiments. In *Proceedings of the the 35th International Electric Propulsion Conference*, pages 2017–25, 2017.

[29] O Tverdokhlebov and A Semenkin. Iodine propellant for electric propulsion-to be or not to be. In *37th Joint Propulsion Conference and Exhibit*, page 3350, 2001.

[30] Dmytro Rafalskyi, Javier Martínez Martínez, Lui Habl, Elena Zorzoli Rossi, Plamen Proynov, Antoine Boré, Thomas Baret, Antoine Poyet, Trevor Lafleur, Stanislav Dudin, et al. In-orbit demonstration of an iodine electric propulsion system. *Nature*, pages 1–5, 2021.

[31] D Rafalskyi Martinez and A Aanesland. Development and testing of the npt30-i2 iodine ion thruster. In *36th International Electric Propulsion Conference*, page 0, 2019.

[32] Gabriel F Benavides, Hani Kamhawi, Jonathan Mackey, Thomas Haag, and Gustavo Costa. Iodine hall-effect electric propulsion system research, development, and system durability demonstration. In *2018 Joint Propulsion Conference*, page 4422, 2018.

[33] Seth J Thompson, Casey C Farnell, Shawn C Farnell, Desiree D Williams, John P Chandler, and John D Williams. Evaluation of iodine compatible cathode configurations. In *36th International Electric Propulsion Conference*, 2019.

[34] F Paganucci, MM Saravia, M Mininni, L Bernazzani, A Ceccarini, T Boulzaguet, G Pellegrini, C Ducci, D Pedrini, T Andreussi, et al. Progress on the development of an iodine-fed hall effect thruster. In *Proceedings of the 35th International Electric Propulsion Conference, Georgia Institute of Technology, Atlanta, GA, USA*, pages 8–12, 2017.

[35] David Jacobson and David Manzella. 50 kW class krypton hall thruster performance. In *39th AIAA/ASME/SAE/ASEE Joint Propulsion Conference and Exhibit*, page 4550, 2003.

[36] Vincent K Rawlin and William R Kerslake. SERT II-durability of the hollow cathode and future applications of hollow cathodes. *Journal of Spacecraft and Rockets*, 7(1):14–20, 1970.

[37] Antonio Gurciullo. *Electric propulsion technologies for enabling the use of molecular propellants.* PhD thesis, University of Surrey, 2020.

[38] Dan M Goebel, Ron M Watkins, and Kristina K Jameson. LaB6 hollow cathodes for ion and hall thrusters. *Journal of Propulsion and Power*, 23(3):552–558, 2007.

[39] DM Goebel, Y Hirooka, and TA Sketchley. Large-area lanthanum hexaboride electron emitter. *Review of Scientific Instruments*, 56(9):1717–1722, 1985.

[40] PP Deichuli, GF Abdrashitov, AA Ivanov, VV Kolmogorov, VV Mishagin, GI Shul'zhenko, NV Stupishin, D Beals, and R Granetz. Ion source with LaB6 hollow cathode for a diagnostic neutral beam injector. *Review of scientific instruments*, 77(3):03B514, 2006.

[41] M Ono, GJ Greene, D Darrow, C Forest, H Park, and TH Stix. Steady-state tokamak discharge via dc helicity injection. *Physical review letters*, 59(19):2165, 1987.

[42] G-N Luo, WM Shu, H Nakamura, S O'hira, and M Nishi. Ion species control in high flux deuterium plasma beams produced by a linear plasma generator. *Review of scientific instruments*, 75(11):4374–4378, 2004.

[43] Jarle Farnes, Arild Vik, Dmitry Bokach, Tjalve Svendsen, Max Schautz, and Xavier Geneste. Recent developments of regenerative fuel cell systems for satellites. In *Proceedings of the 10th European Space Power Conference, Noordwijkerhout, The Netherlands*, pages 13–17, 2014.

[44] Jarle Farnes, Dmitry Bokach, Sander ten Hoopen, Kim Skåtun, Max Schautz, Xavier Geneste, and Arild Vik. Optimized high temperature pem fuel cell & high pressure pem electrolyser for regenerative fuel cell systems in geo telecommunication satellites. In *E3S Web of Conferences*, volume 16, page 10004. EDP Sciences, 2017.

[45] C Muir and A Knoll. Catalytic combustion of hydrogen and oxygen for an electrolysis micro-propulsion system. *Journal of the British Interplanetary Society*, 72(1), 2019.

[46] Robyn Carrasquillo. Iss eclss technology evolution for exploration. In *43rd AIAA Aerospace Sciences Meeting and Exhibit*, page 337, 2005.

[47] N Fernandes, N Herbertz, B Buergler, U Gotzig, M Peukert, M Lau, J Farnes, and T Moreira. Electrolysis based water propulsion for future 1-ton class LEO mission satellites. In *Space Propulsion Conference 2020+1*, page 1, Mar 2021.

[48] Rafael Advanced Defence Systems. Rafael space propulsion catalogue, 2021. URL https://www.rafael.co.il/wp-content/uploads/2021/07/RAFAEL-SPACE-PROPULSION-2021-CATALOGUE-2.pdf.

[49] E Ahedo, JM Gallardo, and M Martınez-Sánchez. Effects of the radial plasma-wall interaction on the hall thruster discharge. *Physics of Plasmas*, 10(8):3397–3409, 2003.

[50] Robert Frisbee, James Polk, Alec Gallimore, and Colleen Marrese. Oxygen-propellant plasma thrusters for cis-lunar electric propulsion missions. In *34th AIAA/ASME/SAE/ASEE Joint Propulsion Conference and Exhibit*, page 3994, 1998.

[51] Takafumi Nakagawa, Naoji Yamamoto, Kimiya Komurasaki, and Yoshihiro Arakawa. Experimental investigation of a hall thruster using oxygen as the propellant. *Japan Society of Aeronautical Space Sciences*, 51(598):606–612, 2003.

[52] Yukikazu Itikawa. Cross sections for electron collisions with oxygen molecules. *Journal of Physical and Chemical Reference Data*, 38(1):1–20, 2009.

[53] K Stephan and TD Märk. Absolute partial electron impact ionization cross sections of Xe from threshold up to 180 ev. *The Journal of chemical physics*, 81(7):3116–3117, 1984.

[54] Davina Di Cara, J Gonzalez Del Amo, Andrea Santovincenzo, B Carnicero Dominguez, Marco Arcioni, Andrew Caldwell, and Ilaria Roma. Ram electric propulsion for low Earth orbit operation: an ESA study. In *30th International Electric Propulsion Conference*, 2007.

[55] Kazutaka Nishiyama. Air breathing ion engine concept. In *54th International Astronautical Congress of the International Astronautical Federation, the International Academy of Astronautics, and the International Institute of Space Law*, pages S–4, 2003.

[56] Vlad Hruby, Bruce Pote, Tom Brogan, Kurt Hohman, James Szabo, and Peter Rostler. Air breathing electrically powered hall effect thruster, December 28 2004. US Patent 6,834,492.

[57] Stefano Cesare, Andrea Allasio, Alberto Anselmi, Sabrina Dionisio, Sergio Mottini, Manlio Parisch, Luca Massotti, and Pierluigi Silvestrin. The european way to gravimetry: From goce to nggm. *Advances in Space Research*, 57(4):1047–1064, 2016.

[58] G Cifali, T Misuri, P Rossetti, M Andrenucci, D Valentian, and D Feili. Preliminary characterization test of het and rit with nitrogen and oxygen. In *47th AIAA/ASME/SAE/ASEE Joint Propulsion Conference & Exhibit*, page 6073, 2011.

[59] M Prioul, F Marchandise, P Dumazert, K Jolivet, D Estublier, A Lazurenko, V Vial, A Bouchoule, P Lasgorceix, L Albarade, et al. Pps-1350 qualification status and performances. In *4th International Spacecraft Propulsion Conference*, volume 555, 2004.

[60] Francesco Marchioni and Mark A Cappelli. Extended channel hall thruster for air-breathing electric propulsion. *Journal of Applied Physics*, 130(5):053306, 2021.

[61] Francesco Marchioni. *Design and Performance Measurements of a long channel Hall thruster for air breathing electric propulsion.* PhD thesis, Politecnico di Torino, 2020.

[62] Fred Jue and Fritz Kuck. Space shuttle main engine (ssme) options for the future shuttle. In *38th AIAA/ASME/SAE/ASEE Joint Propulsion Conference & Exhibit*, page 3758, 2002.

[63] P Caisso, P Brossel, T Excoffon, M Illig, and T Margat. Development status of the vulcain 2 engine. In *37th Joint Propulsion Conference and Exhibit*, page 3550, 2002.

[64] John G Campbell and R Carl Stechman Jr. Water electrolysis propulsion system testing. Technical report, MARQUARDT CO VAN NUYS CA, 1974.

[65] Karsten James, Todd Moser, Andrew Conley, Jeffery Slostad, and Robert Hoyt. Performance characterization of the hydros™ water electrolysis thruster. In *Small Satellite Conference 2015*, 2015.

[66] C Muir, C Ma, and A Knoll. The design, fabrication and test progress summary of the iridium catalysed electrolysis thruster. In *Space Propulsion Conference 2020+1*, page 0, Mar 2021.

[67] Joshua Rovey, Christopher T Lyne, Alex J Mundahl, Nicolas Rasmont, Matthew S Glascock, Mitchell J Wainwright, and Steven P Berg. Review of chemical-electric multimode space propulsion. In *AIAA Propulsion and Energy 2019 Forum*, page 4169, 2019.

[68] JM Hass and MR Holmes. Multi-mode propulsion system for the expansion of small satellite capabilities. *NATO Rept. MP-AVT-171-05*, 2010.

[69] Jason Thrasher, Shae Williams, Phillip Takahashi, and John Sousa. Pulsed plasma thruster development using a novel HAN-based green electric monopropellant. In *52nd AIAA/SAE/ASEE Joint Propulsion Conference*, page 4846, 2016.

[70] Steven P Berg and Joshua L Rovey. Assessment of multimode spacecraft micropropulsion systems. *Journal of Spacecraft and Rockets*, 54(3):592–601, 2017.

[71] Kris Zacny, Phil Chu, Gale Paulsen, Arshak Avanesyan, Jack Craft, and Lars Osborne. Mobile in-situ water extractor (miswe) for mars, moon, and asteroids in situ resource utilization. In *AIAA SPACE 2012 Conference & Exposition*, page 5168, 2012.

[72] Amanda R Hendrix, Terry A Hurford, Laura M Barge, Michael T Bland, Jeff S Bowman, William Brinckerhoff, Bonnie J Buratti, Morgan L Cable, Julie Castillo-Rogez, Geoffrey C Collins, et al. The nasa roadmap to ocean worlds. *Astrobiology*, 19(1):1–27, 2019.

[73] Mark Joseph Sonter. The technical and economic feasibility of mining the near-Earth asteroids. *Acta Astronautica*, 41(4-10):637–647, 1997.

[74] Shane D Ross. Near-Earth asteroid mining. *Space*, 2001.

[75] Dave Mosher and Jenny Cheng. Earth is a desert planet compared to these ocean worlds in the solar system, 2018. URL http://uk.businessinsider.com/water-space-volume-planets-moons-2016-10.

[76] Marvin Warshay and Paul R Prokopius. The fuel cell in space: yesterday, today and tomorrow. *Journal of power sources*, 29(1-2):193–200, 1990.

[77] Asbjørn Strand and Helge Weydahl. Regenerative fuel cell systems for satellites. In *8th European Space Power Conference*, volume 661. Citeseer, 2008.

[78] Fred Mitlitsky, Wim de Groot, L Butler, and James McElroy. Integrated modular propulsion and regenerative electro-energy storage system (impress) for small satellites. Technical report, Lawrence Livermore National Lab., CA (United States), 1996.

[79] Fred Mitlitsky, A Weisberg, P Carter, M Dittman, B Myers, R Humble, and J Kare. Water rocket-electrolysis propulsion and fuel cell power. In *Space Technology Conference and Exposition*, page 4609, 1999.

[80] Wim de Groot, Lynn Arrington, James McElroy, Fred Mitlitsky, Andrew Weisberg, Preston Carter, II, Blake Myers, Brian Reed, Wim de Groot, Lynn Arrington, et al. Electrolysis propulsion for spacecraft applications. In *33rd Joint Propulsion Conference and Exhibit*, page 2948, 1997.

[81] Assad Anis. Cold gas propulsion system-an ideal choice for remote sensing small satellites. *Remote sensing-advanced techniques and platforms*, pages 447–462, 2012.

[82] S Pancheshnyi, B Eismann, G.J.M. Hagelaar, and L.C. Pitchford. About zdplaskin, 2008. URL https://www.zdplaskin.laplace.univ-tlse.fr/.

[83] A Flitti and S Pancheshnyi. Gas heating in fast pulsed discharges in n 2–o 2 mixtures. *The European Physical Journal-Applied Physics*, 45(2), 2009.

[84] Mario Capitelli, Carlos M Ferreira, Alexey I Osipov, and Boris F Gordiets. *Plasma kinetics in atmospheric gases*. Springer, 2000.

[85] AI Bugrova, AV Desiatskov, VK Kharchevnikov, and AI Morozov. Main features of physical processes in stationary plasma thrusters. In *23rd International Propulsion Conference*, 1993.

[86] S Barral, K Makowski, Z Peradzyński, N Gascon, and M Dudeck. Wall material effects in stationary plasma thrusters. ii. near-wall and in-wall conductivity. *Physics of Plasmas*, 10(10):4137–4152, 2003.

[87] Francis Harvey Harlow, Martha Evans, and Robert D Richtmyer. *A machine calculation method for hydrodynamic problems*. Los Alamos Scientific Laboratory of the University of California, 1955.

[88] D Tskhakaya, K Matyash, R Schneider, and F Taccogna. The particle-in-cell method. *Contributions to Plasma Physics*, 47(8-9):563–594, 2007.

[89] Jay P Boris. Relativistic plasma simulation-optimization of a hybrid code. In *Proc. Fourth Conf. Num. Sim. Plasmas*, pages 3–67, 1970.

[90] Francis F Chen. *Introduction to plasma physics*. Springer Science & Business Media, 2012.

[91] James Demmel. Lapack: A portable linear algebra library for supercomputers. In *IEEE Control Systems Society Workshop on Computer-Aided Control System Design*, pages 1–7. IEEE, 1989.

[92] E Ahedo and D Escobar. Influence of design and operation parameters on hall thruster performances. *Journal of applied physics*, 96(2):983–992, 2004.

[93] Vivien Croes, Trevor Lafleur, Zdeněk Bonaventura, Anne Bourdon, and Pascal Chabert. 2d particle-in-cell simulations of the electron drift instability and associated anomalous electron transport in hall-effect thrusters. *Plasma Sources Science and Technology*, 26(3): 034001, 2017.

[94] Andrew Guthrie and Raymond Kornelious Wakerling. *The characteristics of electrical discharges in magnetic fields*, volume 5. McGraw-Hill, 1949.

[95] Michelle K Scharfe, Nicolas Gascon, Mark A Cappelli, and Eduardo Fernandez. Comparison of hybrid hall thruster model to experimental measurements. *Physics of Plasmas*, 13 (8):083505, 2006.

[96] JM Tejeda, M Reza, F Faraji, and A Knoll. Performance enhancement of hall effect thrusters using radiofrequency excitation. *Acta Astronautica*, 194:145–161, 2022.

[97] Käthe Dannenmayer and Stéphane Mazouffre. Elementary scaling relations for hall effect thrusters. *Journal of Propulsion and Power*, 27(1):236–245, 2011.

[98] G Guerrini, AN Vesselovzorov, M Bacal, and IB Pokrovsky. Investigation of a small, closed electron drift, stationary plasma thruster. *Review of scientific instruments*, 67(3):990–992, 1996.

[99] AI Bugrova, AS Lipatov, AI Morozov, and DV Churbanov. On a similarity criterion for plasma accelerators of the stationary plasma thruster type. *Technical Physics Letters*, 28 (10):821–823, 2002.

[100] L Garrigues. Computational study of hall-effect thruster with ambient atmospheric gas as propellant. *Journal of Propulsion and Power*, 28(2):344–354, 2012.

[101] D. C. Meeker. Finite element method magnetics, 2021. URL https://www.femm.info/wiki/HomePage.

[102] Alexander Schwertheim and Aaron Knoll. Low power thrust measurements of the water electrolysis hall effect thruster. *CEAS Space Journal*, pages 1–15, 2021.

[103] Bernhard Wolf. *Handbook of ion sources*. CRC press, 1995.

[104] Lachlan Orr. Design of a novel cathode neutraliser for laboratory plasma thruster testing. *Dissertation submitted for the degree of Master of Science in Aeronautics at Imperial College London*, 2019.

[105] Keng Yong Chua. Development and investigation of a novel laboratory neutraliser for plasma thrusters. *Dissertation submitted for the degree of Master of Science in Aeronautics at Imperial College London*, 2019.

[106] Guanyun Liu. Development and testing a novel millinewton thrust balance for plasma thruster performance characterization. *Dissertation submitted for the degree of Master of Science in Aeronautics at Imperial College London*, 2018.

[107] A Schwertheim, E Rosati Azevedo, G Liu, E Bosch Borràs, L Bianchi, and A Knoll. Interlaboratory validation of a hanging pendulum thrust balance for electric propulsion testing. *Review of Scientific Instruments*, 92(3):034502, 2021.

[108] James E Polk, Anthony Pancotti, Thomas Haag, Scott King, Mitchell Walker, Joseph Blakely, and John Ziemer. Recommended practices in thrust measurements. Technical report, California Inst of Technology Pasadena Jet Propulsion Lab, 2013.

[109] Sabrina Pottinger, Vaios Lappas, Christine Charles, and Rod Boswell. Performance characterization of a helicon double layer thruster using direct thrust measurements. *Journal of Physics D: Applied Physics*, 44(23):235201, 2011.

[110] Kurt A Polzin, Thomas E Markusic, Boris J Stanojev, Amado DeHoyos, and Benjamin Spaun. Thrust stand for electric propulsion performance evaluation. *Review of Scientific Instruments*, 77(10):105108, 2006.

[111] S Cesare, F Musso, F D'Angelo, G Castorina, M Bisi, P Cordiale, E Canuto, D Nicolini, E Balaguer, and P-E Frigot. Nanobalance: The european balance for micro-propulsion. In *31st International Electric Propulsion Conference*, 2009.

[112] Kunning G Xu and Mitchell LR Walker. High-power, null-type, inverted pendulum thrust stand. *Review of Scientific Instruments*, 80(5):055103, 2009.

[113] K Swar, D Staab, A Garbayo, L Shadbolt, S Masillo, A Lucca Fabris, B Karadag, and R Moloney. Design and testing of a μn mn torsional thrust balance with wireless microwave power transmission. In *36th International Electric Propulsion Conference University of Vienna*, 2019.

[114] Yuan-Xia Yang, Liang-Cheng Tu, Shan-Qing Yang, and Jun Luo. A torsion balance for impulse and thrust measurements of micro-newton thrusters. *Review of Scientific Instruments*, 83(1):015105, 2012.

[115] Burak Karadag, Silvia Masillo, Rachel Moloney, Andrea Lucca Fabris, Thomas Potterton, Aaron Knoll, and Paolo Bianco. Experimental investigation and performance optimization of the halo thruster. In *36th International Electric Propulsion Conference University of Vienna*, 2019.

[116] Thomas Harle, AK Knoll, V Lappas, and Paolo Bianco. Performance measurements of a high powered quad confinement thruster. In *33rd International Electric Propulsion Conference, Washington DC*, 2013.

[117] F Trezzolani, M Magarotto, M Manente, and D Pavarin. Development of a counterbalanced pendulum thrust stand for electric propulsion. *Measurement*, 122:494–501, 2018.

[118] Aaron Knoll, Thomas Harle, Vaios Lappas, and Mark Pollard. Experimental performance characterization of a two-hundred-watt quad confinement thruster. *Journal of Propulsion and Power*, 30(6):1445–1449, 2014.

[119] Emmanuelle Rosati Azevedo and Aaron Knoll. Magnetic field enhancement of the quad confinement thruster, 2018.

[120] E Rosati Azevedo, VG Tirila, A Schwertheim, and A Knoll. Magnetic field enhancement of the quad confinement thruster (qct): Design and early development of the qct phoenix. In *Space Propulsion Conference 2020+1*, page 0, Mar 2021.

[121] Paul M Ordin. Safety standard for hydrogen and hydrogen systems guidelines for hydrogen system design, materials selection, operations, storage and transportation. *Technical Note*, 1997.

[122] K Kitazawa and RL Coble. Electrical conduction in single-crystal and polycrystalline al2o3 at high temperatures. *Journal of the American Ceramic Society*, 57(6):245–250, 1974.

[123] CC Wang, SA Akbar, W Chen, and VD Patton. Electrical properties of high-temperature

oxides, borides, carbides, and nitrides. *Journal of materials science*, 30(7):1627–1641, 1995.

[124] Alexander Schwertheim and Aaron Knoll. Performance characterisation of the water electrolysis hall effect thruster (WET-HET) using direct thrust measurements. In *Space Propulsion Conference 2020+1*, page 0, Mar 2021.

[125] RL Wills. I. effects of temperature on the magnetic properties of iron and alloys of iron. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 50 (302):1–37, 1900.

[126] Thomas Burton, Aaron M Schinder, German Capuano, Julian J Rimoli, Mitchell LR Walker, and Gregory B Thompson. Plasma-induced erosion on ceramic wall structures in hall-effect thrusters. *Journal of Propulsion and Power*, 30(3):690–695, 2014.

[127] Nathan P Brown and Mitchell LR Walker. Review of plasma-induced hall thruster erosion. *Applied Sciences*, 10(11):3775, 2020.

[128] Ryan W Conversano, Dan M Goebel, Richard R Hofer, Ioannis G Mikellides, and Richard E Wirz. Performance analysis of a low-power magnetically shielded hall thruster: experiments. *Journal of Propulsion and Power*, 33(4):975–983, 2017.

[129] N Gascon, M Dudeck, and S Barral. Wall material effects in stationary plasma thrusters. i. parametric studies of an spt-100. *Physics of Plasmas*, 10(10):4123–4136, 2003.

[130] JP Bugeat and C Koppel. Development of a second generation of spt. In *IEPC'95-International Electric Propulsion Conference, 24 th, Moscow, Russia*, pages 286–293, 1996.

[131] Hirokazu Tahara, Katsumi Imanaka, and Seiro Yuge. Effects of channel wall material on thrust performance and plasma characteristics of hall-effect thrusters. *Vacuum*, 80(11-12): 1216–1222, 2006.

[132] Tsuyohito Ito, Nicolas Gascon, W Scott Crawford, and Mark A Cappelli. Experimental characterization of a micro-hall thruster. *Journal of Propulsion and Power*, 23(5):1068–1074, 2007.

[133] AI Morozov and IV Melikov. Similitude in hall-current plasma accelerators. *Soviet Physics Technical Physics*, 19:340–342, 1974.

[134] A Schwertheim and A Knoll. The water electrolysis hall effect thruster (WET-HET): Paving the way to dual mode chemical-electric water propulsion. In *36th International Electric Propulsion Conference*, Sep 2019.

[135] D Valentian and N Maslennikov. The pps 1350 program. *CP IEPC*, 97(134):855–861, 1997.

[136] GD Racca, A Marini, L Stagnaro, J Van Dooren, L Di Napoli, BH Foing, R Lumb, J Volp, J Brinkmann, R Grünagel, et al. Smart-1 mission description and development status. *Planetary and space science*, 50(14-15):1323–1337, 2002.

[137] L Biagioni, M Saverdi, M Berti, U Cesari, and Mariano Andrenucci. Design and preliminary characterization of a 5 kw hall thruster prototype. In *Proc. 28th IEPC*, 2003.

[138] A Schwertheim and A Knoll. In situ utilization of water as a propellant for a next-generation plasma propulsion system. *JBIS. Journal Of The British Interplanetary Society*, 72(1):7–11, Jan 2019.