Investigation of the Power Balance in Alternative Divertor Configurations in Spherical Tokamaks

Matthew James Dunn

Doctor of Philosophy

University of York

Physics

September 2022

Abstract

In tokamak reactors of the future, the heat fluxes in the divertor regions will be higher than current materials are capable of withstanding. Instead of merely scaling up existing tokamak technology to reactor-relevant scales, a new method for reducing heat fluxes is needed. This will allow plasmas of sufficient temperatures for net energy gain in the reactor, while still giving acceptable component lifetimes in the divertor.

Changing divertor magnetic configuration has the potential to lower heat fluxes for equivalent core plasmas. The super-X concept shows promise in simulations, but can only be tested experimentally in a tokamak with a full divertor baffle. The first of these, MAST Upgrade, has just completed its first experimental campaign. Early results are presented herein, taking advantage of methods developed during other analysis.

A method for power balance analysis was developed utilising archive data from MAST. Two effects not accounted for previously were considered; tile shadowing and neutral beam power absorption. The method gives 85% power accounting for shots with 0 or 1 neutral beams, and 73% for 2 beams, with this discrepancy thought to result from fishbone instabilities.

The effect of tile plasma exposure on infrared camera measurements was determined with experiments on Magnum-PSI. The infrared reported a temperature decrease of 70 °C over three hours, despite neither the heat flux nor the true tile temperature decreasing, caused by uneven tile erosion of up to 100 µm. This will not occur in MAST Upgrade, but is an interesting result nonetheless.

Power balance on MAST Upgrade, utilising the method developed for MAST, was assessed. The effectiveness of a super-X with full a baffle was measured experimentally for the first time. The radiation in the divertor during a super-X was 3.1 times core radiation, compared to 1.1 with a conventional divertor, with a corresponding reduction in divertor heat flux.

Contents

List of Tables

List of Figures

Acknowledgments

I would like to acknowledge my “industrial contact” Dr Andrew Thornton of Culham Centre for Fusion Energy, who was a second supervisor throughout this project, and a continuous source of help and guidance. Thanks are also due to my supervisor Professor Kieran Gibson of the University of York, whose years of experience were invaluable to me in guiding the research themes of this project, and also in the planning of its associated experiments.

I am grateful to the teams at MAST, Magnum-PSI and MAST-Upgrade, who operated the machines and made data available for chapters 4, 5 and 6 respectively. I am especially grateful to Jack Lovell, Peter Ryan and Andrew Kirk for making some of the data in chapter 6 available to me ahead of time.

I would also like to thank my parents for their support, especially during the Covid-19 pandemic, which severely disrupted my PhD.

Declaration

I declare that this thesis is a presentation of original work and I am the sole author. This work has not previously been presented for an award at this, or any other, University. All sources are acknowledged as references.

Much of the content in chapter 5 has already been published as an article, for which I am the first author, in the peer reviewed journal Nuclear Materials and Energy (M Dunn et al. *NME* 25:4, 2020). Some work similar to that presented in chapter 6 has been submitted as another article, for which I am a co-author, to the same journal (first author J Lovell, as yet unpublished).

Chapter 1

Introduction and Background Theory

1.1 Introduction to Fusion

It is known that, to avert the worst impacts of climate change, all the world's electricity will need to come from carbon-neutral sources by 2050 [1]. Most of this is expected to come from solar and wind power. However, these are both very intermittent sources [2], and solar power in particular uses a lot of land. To solve the intermittency of both wind and solar would require either immense energy storage or large scale interconnections across continents. Both of these would require significant investment in addition to the actual wind and solar infrastructure, and would use even more land. The ever-growing human population of the world, which is predicted to reach 9.7 billion by 2050, and to continue to grow thereafter [3], is likely to require considerably more land for food production than is presently used. The solution to this is a carbon-neutral source of electricity that uses very little land (both in itself, and by way of removing the need for energy storage), while not bringing significant additional complications, such as location constraints (e.g. geothermal) or weapons proliferation (e.g. nuclear fission).

The most promising solution to this problem is nuclear fusion, which has the potential to produce large quantities of energy, releasing no carbon dioxide. Only the land occupied by the reactor itself is required, and it can easily be connected to the existing grid (unlike solar) and there is minimal associated mining (unlike fossil fuels). The fuel consumption is minimal, meaning the only costs are the upfront cots of building the reactor, and the ongoing

staffing costs.

Nuclear fusion is the process of joining together small atomic nuclei to form larger nuclei, releasing very large amounts of energy; orders of magnitude higher than for an equivalent mass of fossil fuels. Nuclear fusion of small hydrogen nuclei into larger helium nuclei is the process that powers the Sun. It requires positively charged atomic nuclei (which repel each other) to be close together for long times, in order for quantum tunnelling to overcome the repulsion. The enormous gravitational pressure that causes this process in the Sun cannot be replicated on the Earth, however another nuclear fusion reaction is more promising.

This is the deuterium-tritium (DT) reaction, the fusion of two heavy isotopes of hydrogen; hydrogen-2 (deuterium) and hydrogen-3 (tritium). This momentarily creates a helium-5 nucleus, which immediately ejects a neutron, and becomes a helium-4 nucleus. This releases more energy per unit mass than any other process that is achievable on Earth, and the only by-product is inert helium-4. The majority of this energy (carried by the neutron) can then be harnessed to boil water and drive a steam turbine to produce electricity. The energy of the helium-4 is used to sustain the reaction.

1.1.1 The Field of Fusion Research

Fusion research encompasses three broad strands. One of these is fusion materials science, which is dedicated to studying the novel materials that will be necessary in the fusion reactors of the future. This includes materials resistant to the high-energy neutrons released in the DT reaction, which are similar to materials used in other nuclear reactors. Another example is materials for more powerful magnets, including superconductors, which have various other potential applications. Research into materials for fusion is often funded jointly with these other areas.

The second strand is inertial confinement fusion (ICF) which hopes to achieve nuclear fusion under extreme temperature and pressures, even if these can only be contained for very short times. These conditions are attained by suspending a frozen ball of DT fuel, and firing lasers at it from all directions with spherical symmetry. This causes an implosion, which briefly creates the necessary conditions. A challenge for ICF is that continuous energy generation requires a high repetition rate. Research into ICF is also often funded jointly with other areas, such as high energy density physics.

The third strand is magnetic confinement fusion (MCF) which aims for fusion at ex-

treme temperatures and moderate pressures, sustained for long times. These conditions are attained by containing plasma in a magnetic field, and holding it in a vacuum chamber away from the walls. The plasma is then heated for a time, and can release energy from nuclear fusion. This approach, with its longer timescales, is considered the best approach for energy generation, and as such is the focus of most nuclear fusion research funding. The project described in this thesis falls within the MCF strand.

1.1.2 Tokamaks

The most promising fusion reactor design is the tokamak [4], which was developed in the Soviet Union in the 1950s. The name TOKAMAK is an abbreviation of the Russian phrase "toroidal'naja kamera s magnitymi Katuskami" which means "toroidally-shaped container with magnetic coils". As the name suggests, a tokamak consists primarily of a toroidal solenoid and vacuum chamber, with the plasma contained inside. However, simple toroidal solenoid with no additional coils will quickly fail.

The bending of the solenoid into a toroidal shape creates a magnetic field gradient, with a stronger field on the inboard side and a weaker field on the outboard side. This occurs because the coils are closer together on the inboard side, and the magnetic field gradient causes a phenomenon known as the $\nabla \boldsymbol{B}$ drift. The $\nabla \boldsymbol{B}$ drift acts perpendicular to both $\boldsymbol{B}$ and $\nabla \boldsymbol{B}$, upward for ions and downward for electrons. This then creates an electrostatic field, which leads to an $\boldsymbol{E} \times \boldsymbol{B}$ drift. This drift is in the outboard direction, pushing the ions to the outboard side of the vacuum vessel, and the electrons to the inboard side. Because the vessel is conductive, the plasma current shorts out through the vessel wall. This happens within half a second of any such device being switched on.

Early attempts to counteract this problem included twisting the solenoid into a figure-of-eight [5], but this had only limited success. In a tokamak, a coil is wound around the centre column, and current in this coil is swept at a constant rate during the experiment. This drives a constant electric current in the plasma. This current, flowing in the toroidal direction (around the major radius of the torus), creates a magnetic field in the poloidal direction (around the minor radius of the torus). This gives the field in the chamber a poloidal component, resulting in field lines that travel in a helical path around the torus. These are illustrated by the blue lines in figure 1. This serves to bring particles that have drifted to the outboard side back to the inboard side, thereby short circuiting the charge

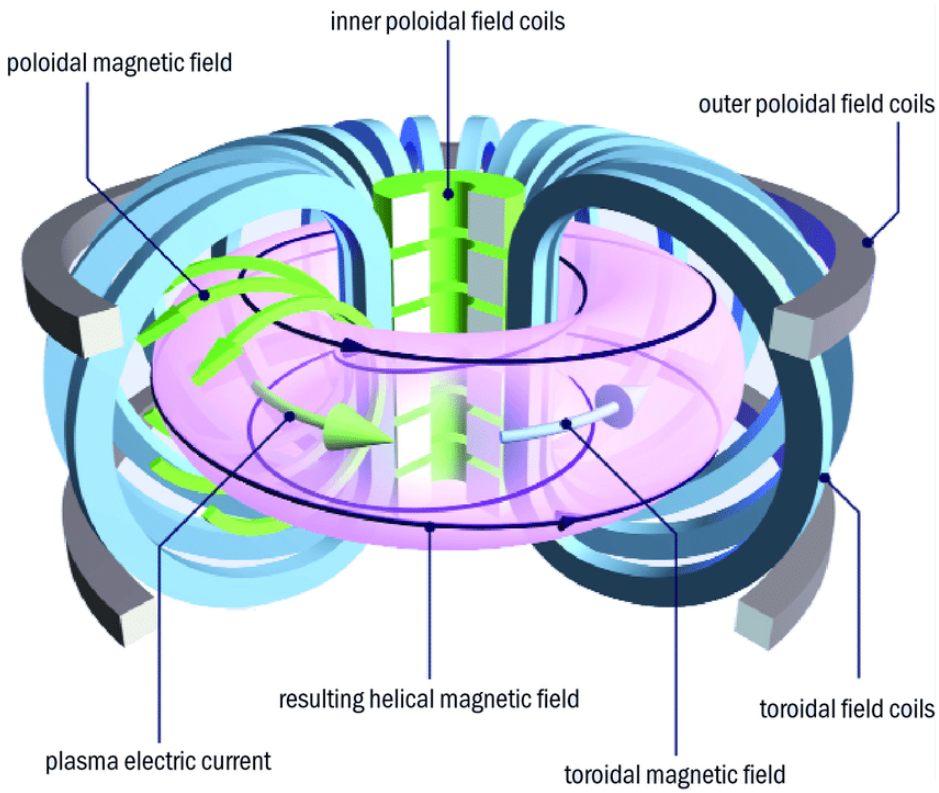

Figure 1: A diagram of the coil currents, plasma currents and magnetic fields in a tokamak. Figure courtesy of EUROfusion.

separated by the $\boldsymbol{E} \times \boldsymbol{B}$ drift, allowing the plasma to be confined. The plasma current can also interact with other currents nearby to produce more complex magnetic geometries (see section 1.2.3).

1.1.3 Limiters and Divertors

In this simplest form of the tokamak however, there are many energetic ions and electrons escaping the confinement and impacting the vessel wall, where they recombine. The tokamak limiter, a conductor protruding from the wall that could withstand higher power loads, as in figure 2 (a), was an early solution to this. With a tokamak limiter it was easier to control the extent of the plasma, and allowed longer lifetimes of the first wall components, with most of the plasma flux now incident on the limiter rather than the wall. However, the limiter formed part of the last closed flux surface (the inner boundary of the orange region in figure 2), so any atoms ejected from its surface by collisions with energetic plasma particles could easily be carried into the core plasma. The heavy metallic atoms would become partially

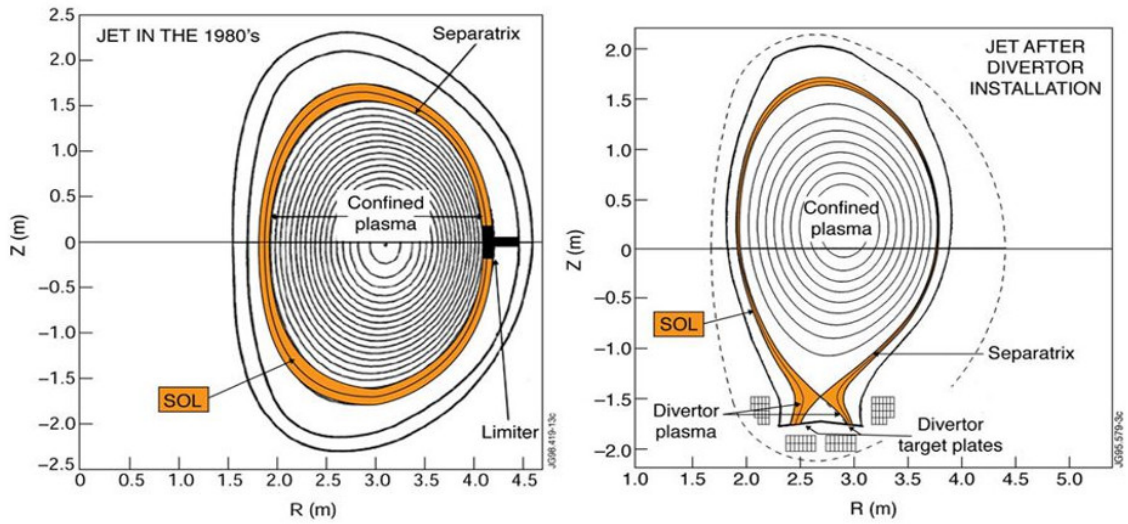

Figure 2: The Joint European Torus (JET) tokamak with (a) a limiter and (b) a divertor. The divertor of JET was placed inside the original vacuum vessel, and reduced the available plasma volume. Figure credit: EDFA 2011

ionised, and their remaining electrons would be frequently excited to very high energy states, releasing high energy photons as they decayed. This was a major heat loss mechanism for the bulk plasma, and detrimental to performance, preventing experiments from achieving nuclear fusion burn. This also lowered the concentration of hydrogenic atoms. A method of reducing impurities inside the last closed flux surface was needed.

The tokamak divertor, which is a focus of this project, was proposed as a way of keeping the plasma-surface contact outside the last closed flux surface, figure 2 (b). This idea was to reduce the concentration of high-Z impurities inside the core plasma, allowing higher temperatures to be reached. In the case of the Joint European Torus (JET) the divertor was inserted retrospectively into the original vacuum vessel, and this unfortunately reduced the available volume for plasma. This is not a problem for modern tokamaks, which are designed with divertors in mind.

Diverted tokamaks also have disadvantages, such as inferior control of the extent of the plasma compared to limited tokamaks, and shorter first wall component lifetimes because of edge-localised modes (ELMs). ELMs result in lower temperature and density at the plasma edge, and have periodic behaviour. The presence of ELMs deteriorates the confinement of plasma in the tokamak [6, p. 409]. Despite this, a divertor was and is considered the best

option for achieving fusion, because it avoids the radiation heat loss problem described above.

The creation of divertor magnetic topology relies on how the magnetic field coils are positioned around the divertor chamber. As explained in section 1.1.2, a plasma current is established in the tokamak to achieve confinement. A parallel current is then put through a poloidal field coil directly beneath the plasma, which creates a poloidal magnetic field null somewhere between the two, (see the bottom of figure 2 (b)). This field null prevents the particles from spiralling around the confined plasma, and diverts them to the plates before they have a chance to impact the remainder of the primary vessel wall.

Unfortunately, the shift from limiters to divertors also created a new problem: the focusing of all the tokamak's heat onto a small area, resulting in higher heat flux (heat energy per unit time per unit area). One way to mitigate this, which takes advantage of a diverted, rather than limited, configuration is with plasma detachment. Plasma detachment is a state in which the plasma conditions at the plates cannot be predicted based solely on the plasma conditions in the core. Crucially, when the plasma is attached, an increase in core plasma density results in an equivalent increase in the divertor plasma flux, but in detached conditions there would be no such knock-on effect on the divertor plasma. When the heat and particle flux at the plates is measured, the plasma is observed to transition to the detached regime as the core density increases, as will be explained in section 1.2.2. Less heat flux and particle flux (particles per unit time per unit area) to the plates for the same core conditions equates to greater detachment.

Plasma detachment can be increased with the type of enclosed divertor that has been built on MAST Upgrade, which is the first tokamak in the world with a large super-X-capable divertor chamber enclosed to such a high degree [7]. Plasma detachment can also be increased by cold gas puffing in the divertor region, causing hot escaping plasma to radiate heat away before impacting the plates [8]. The switch from limiters to divertors was intended to stop this from occurring in the core, but it is desirable in a segregated divertor chamber. This is because this heat has already left the core anyway, so radiation here just reduces the heat load to the divertor plates.

1.1.4 Materials Limits and Motivation

The motivation for this project arises from the lack of suitable plasma-facing materials, which could withstand the high heat flux that comes with scaling current tokamak technology up

to a commercial fusion reactor. The fusion power plants of the future will be much bigger than the tokamaks of today, as this has been deemed necessary for net energy gain defined by the Lawson criterion [9]. The Lawson criterion states that to achieve ignition, one must maximise the product of plasma density, confinement time, and temperature [10]. Since current technology allows for a certain plasma pressure, these plasmas will have similar energy density to those in smaller tokamaks. Therefore, the total energy of the plasma scales with the plasma volume, i.e. the cube of the radius. However, the divertor area, being 2-dimensional, scales only with the square of the radius. This is further compounded by recent research which suggests that the scrape-off layer width λ remains constant as the machine is enlarged, scaling only with the poloidal magnetic field strength [11]. This means the divertor wetted area (the area where the plasma is incident) scales approximately linearly with the radius. Since the rate of heat leaving the plasma (in Watts) scales with the cube of the radius, while the area onto which it is deposited (in m^2) scales linearly with the radius, the divertor heat flux (the rate of heat deposition divided by the area, in $W\,m^{-2}$) scales with the square of the machine radius.

Target heat fluxes of $\approx 15\,MW\,m^{-2}$ are expected in the proposed DEMO fusion reactor [12] [13]. This is higher than the $\approx 5\,MW\,m^{-2}$ that current materials can withstand, and even the $\approx 10\,MW\,m^{-2}$ expected in ITER [14], which is the main next generation fusion experiment, under construction in the south of France. Although materials have reasonable lifetimes in the tokamaks of today (which are smaller than ITER), this will not be so in the bigger, hotter tokamaks of the future.

There are candidate materials in the pipeline that could potentially handle heat fluxes of $10\,MW\,m^{-2}$ or more, such as liquid lithium [15] or CuCrZr [16] but this is still not enough for DEMO or a power plant. A combined approach is needed to solve this problem, to produce a greater overall improvement. This project comes at the problem from the other end, by reducing the flux that will be incident on the plate material. Divertor flux can be reduced by means of flux expansion or detachment, and these are usually achieved by changing divertor geometry to create advanced divertors, as explained in section 1.2.3.

The advanced divertor of most interest to this project is called the super-X divertor. The super-X has been extensively theorised and simulated before, but has only been tested once experimentally [17], in less than ideal circumstances (see section 2.2.3). The desire to test this most promising approach, on a tokamak primarily designed for the purpose, has lead to

chapter 6 of this thesis. The results can be used to extrapolate to the hotter (and probably much bigger) fusion power plants of the future.

1.2 Plasma and Tokamak Theory

A plasma is a fluid in which the electrons are no longer bound to atomic nuclei. The electrons and ions (or nuclei in a high temperature plasma) can therefore travel independently of each other, like two different fluids. Despite this, most plasmas are quasi-neutral, meaning they have no net electrical charge. Electrical attraction between electrons and ions naturally prevents them from separating too much, making it difficult to create a plasma that is not quasi-neutral.

The independent movement of electrons and ions has implications for the heating of a plasma. Electron-cyclotron resonance heating (ECRH), performed using a microwave source [18, p. 382], does not cause any secondary heating of the ions, resulting in the electron temperature being much higher than the ion temperature. Because of the mass difference, this means the average temperature of the plasma can be relatively cool. This is called a low temperature (or non-thermal) plasma. Alternatively, radio frequency heating sources can be used for ion-cyclotron resonance heating (ICRH) [18, p. 382]. Because of their high mass (and therefore momentum), heating the ions in this way causes the electron temperature to increase to be approximately equal. This is called a high temperature (or thermal) plasma.

Although ECRH heating (or a similar technique called electron-Bernstien wave heating or EBW) are sometimes used as secondary sources of heating in nuclear fusion, fusion plasmas are always high temperature thermal plasmas. All magnetically confined fusion plasmas use ICRH heating, which is known as ohmic heating in this field. Most large tokamaks also have neutral beam injection (NBI) heating, which only heats bare nuclei and not electrons. Nuclear fusion plasmas have roughly equal electron and ion temperatures in excess of 10^7 K.

1.2.1 Debye Shielding and Plasma Sheaths

Separation between electrons and ions creates localised electric fields within the plasma. Unlike the nuclear (strong and weak) and Van der Waals forces that normally govern fluid behaviour, the electric force is long range. This means that interactions in a plasma can take place over somewhat longer ranges than in other fluids or solids. These interactions are not

as long range as expected however, because of a phenomenon called Debye shielding [18, p. 7]. Although electrons are not bound to specific nuclei, there is still an electrical attraction between them. This causes a number of electrons to gather near a nucleus, and shield electrons further away from that charge of that nucleus. This loosely spherical arrangement of electrons has a "Debye radius" which, in a low temperature plasma, is given by

$$\lambda_D = \sqrt{\frac{\epsilon_0 k_B T_e}{n_e e^2}} \tag{1}$$

where λ_D is the Debye radius, ϵ_0 is the permittivity of free space, k_B is Boltzmann's constant, T_e is the electron temperature, and n_e is the electron density and e is the electron charge. A combination of quasi-neutrality and Debye shielding makes it difficult for electric fields to form on scales longer than the Debye radius (also known as the Debeye length).

There is also another case of electric fields forming on a comparable length scale, called a plasma sheath. When a plasma terminates at a solid surface (normally the edge of a plasma vessel), a sheath region is established at the surface. The plasma sheath comes about because the electrons are travelling at higher velocity than the ions (because of their smaller mass) so they arrive at the vessel wall first. This gives the vessel a slight negative bias, repelling further electrons from the edge, as well as creating a surplus of ions near the edge [18, p. 272]. The changes in electron and ion density in and around the plasma sheath are shown in figure 3. The conditions in present-day tokamak divertors are similar in some ways to those in the electron-heated low temperature plasmas described above, as the plasma in the divertor has cooled down since leaving the ion-heated plasma core, affecting the electron density and temperature, which in turn affects the collision rates. However, this may prove to be less true in larger tokamaks in the future.

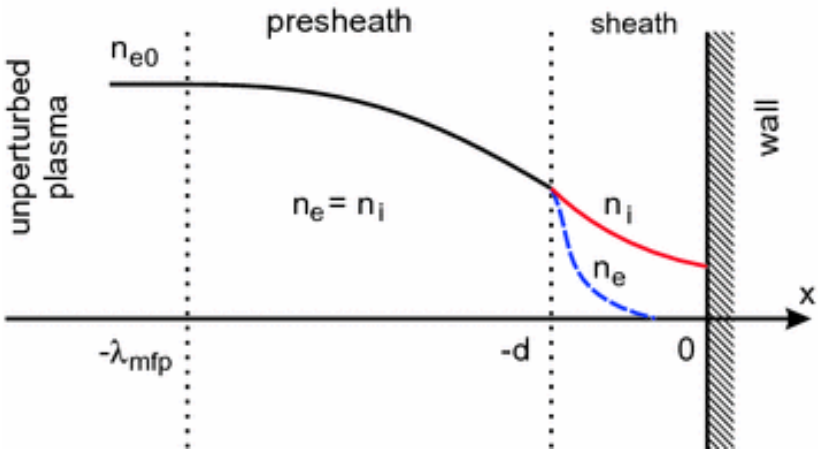

Figure 3: The variation of electron and ion densities close to, and in, the plasma sheath. Figure from [19].

1.2.2 Detachment Physics

The heat flux that is incident onto the divertor plates originates in the plasma core, then escapes into the scrape-off layer (SOL) of the tokamak, which is the region outside the last closed flux surface or separatrix. The flux travels through the SOL, past the X-point, and onto the plates. The SOL width, λ, is conventionally defined at the outboard midplane, and is of the order of mm for most tokamaks (1 mm predicted for ITER [20]). This width does not increase with the size of the tokamak; it varies only with the poloidal field strength [11].

A popular theoretical description of divertor physics, the two point model, is often used in the field, as it is easy to model analytically. It was first developed at Oak Ridge National Laboratory in 1984 [21]. The name two point model refers to the two spatial points considered, one at the midplane and one at the strike point. It assumes the the plasma conditions (density, temperature etc.) at the strike point can be modelled using only the known midplane parameters as inputs, without having to account for any other phenomena between those two points. This is an enticingly simple model, and can indeed be used when upstream electron densities are low.

At the divertor plates, low density plasma arriving from the core initially forms a plasma sheath, as described in section 1.2.1. Starting from this elementary plasma physics, the three regimes of divertor plasma physics can be described.

1. At low electron densities, the divertor plasma operates in the sheath-limited regime, which is so called because the sheath acts as the only sink of heat and particles for

the SOL plasma. No heat is radiated or lost between the midplane and divertor. There is cross-field transport, but there is no ionisation, volumetric recombination or neutral friction [22, p. 153]. In this regime, the electron temperatures are high, and both temperature and pressure are constant along the field lines (that is, there is no temperature gradient) [23]. This is the regime described by the basic two point model.

2. At moderate electron densities, the divertor transitions to the conduction-limited, or high recycling, regime. In this regime, the electron temperature falls as particles approach the vessel wall, but the electron density rises, so that the pressure is approximately constant as before [23]. The non-zero heat conductivity, which is the origin of the name conduction-limited, results in a fall in temperature. This temperature fall causes the dependence of recycling rate on the density, which was linear in the sheath-limited regime, to become quadratic, and this is sometimes called a high recycling regime [22, p. 230]. The fall in electron temperature reduces the heat flux reaching the plates. The two point model can be extended to account for these phenomena.

3. At high electron density, the plasma transitions into the detached regime. In this regime, both the electron temperature and electron density are very low near the plates, so clearly the pressure is also lower at the plates; it is not conserved along the field lines [23]. Pressure varying along the field lines is regarded by some as the definition of detachment [24]. This is done despite having the same high temperature and density in the bulk and at the midplane. This disconnect between upstream (midplane) and divertor conditions is the origin of the name detached, and is brought about by a reduction in particle flux to the divertor [22, p. 477]. In this regime, even less heat flux reaches the plates than the high-recycling regime. The two point model breaks down in this regime.

It is also possible to lower the temperature in the divertor region with impurity seeding. This can bring about detachment at a lower upstream electron density. Impurity seeding refers to puffing cold impurity gases, usually atoms with many electrons such as N_2 or Ar, into the divertor region. It relies on various radiative and collisional processes to reduce the energy of the particles. Some of these processes already take place to some degree because of small concentrations of unavoidable impurities. The large variety of different processes have different cross-sections at different energies. One such process is called charge exchange. It

consists of a collision between an energetic ion and a cold neutral in which an electron is exchanged, meaning the energetic particle becomes a neutral, and the cold particle becomes an ion. The energetic neutral then emits a photon, which escapes from the tokamak; thus the energy is radiated away. This process is shown below for injected argon

$$\mathrm{D}^{*+} + \mathrm{Ar} \longrightarrow \mathrm{D}^{*} + \mathrm{Ar}^{+} \longrightarrow \mathrm{D} + \mathrm{Ar}^{+} + \gamma$$

where D is deuterium, Ar is argon and γ is a photon. The ions can then be neutralised through collisions with free electrons in the plasma. Processes like this lead to a discrepancy between the upstream ion flux and the divertor ion flux, because some of the ions neutralise with electrons. In this way, the divertor plasma becomes recombination-dominated. This ion cooling removes energy and momentum from the system, so the heat flux is also reduced, since the heat energy of the electron in the first collision was radiated away by a photon. These reductions of both heat and momentum are the two conditions necessary for detachment. There are various definitions of detachment used by different authors. Some define detachment simply as when the plasma temperature drops below 5 eV or 1 eV. These temperatures are similar to the ionisation energy of deuterium, and represent the point at which the plasma transitions from ionisation-dominated to recombination-dominated. Others have a more quantitative definition of detachment, based on the two-point model. The ratio of the measured ion flux at the plates to the ion flux predicted by this analytical sheath-limited model (which assumes full attachment) is the Degree of Detachment (DoD). This numerical measure better represents the fact that detachment/attachment is not a binary state.

Either by using a specific temperature, or by choosing a threshold value of the degree of detachment, one can define a detachment front, between the primary X-point and the divertor plate. These numerical definitions can both be used to define a point in space after which the majority of ions start radiating energy away. However, the formal definition of detachment suggests a threshold parallel pressure ratio instead [24]. The degree of detachment is a more robust experimental measurement, but a threshold target temperature gives a way to theoretically predict greater detachment, from the relationship between target electron temperature and key parameters, given by

$$T_e^t \propto \frac{(q_\parallel^u)^{10/7} \cdot (1 - f_{rad})^2}{n_u^2 L_\parallel^{4/7}} \frac{R_u^2}{R_t^2} \tag{2}$$

Where T_e^t is the target electron temperature, $q_{\|}^u$ is the upstream parallel heat flux, f_{rad} is the fraction of power radiated out, n_u is the upstream plasma density, $L_{\|}$ is the connection length, R_u is the major radius of the separatrix at the outer (or inner) midplane, and R_t is the major radius of the outer (or inner) target [25]. This shows that the target electron temperature can be decreased by increasing f_{rad} (e.g. by cold gas puffing) or by increasing $L_{\|}$ or R_t (e.g. with a super-X) for fixed upstream heat and density. Analytic modelling suggests a similar relationship between the target density as the one for temperature given above, however this is not borne out by experiment, unless the divertor is sufficiently enclosed [25].

It should be noted that detachment usually occurs only at the outer target, and also that the detachment front has a tendency to move toward weaker magnetic field, so usually rests near the X-point, close to the core [26]. This is because of higher plasma density from regions of stronger magnetic field, where the flux tubes are narrower, resulting in higher heat and particle flux from these regions. This creates a feedback effect which acts to hold the front where the field is weakest. Although the position of the detachment front is also affected by other factors, this position close to the X-point can result in radiative processes within the plasma core. These degrade confinement and usually lead to a disruption (the entire plasma shorts out through the vessel wall). This can be mitigated by plasma shaping or advanced divertors (see section 1.2.3).

Experimental investigations into detachment have found that it is easier to achieve in a higher density tokamak [26], in agreement with the theory that detachment is greater for higher divertor neutral density. Good agreement was found between detachment theory and experiment in that work. However, even in detached conditions, some hot ions still hit the plates. The particle flux per metre squared (measured at the target) can be reduced further still by means of advanced divertor configurations. Advanced divertors are created by changing the current flowing through the divertor coils, and by using additional coils. These changes in magnetic topology therefore affect the heat load to the divertor plates. However, it should not be assumed that all advanced divertors are equal in terms of greater flux expansion, better detachment, larger target radius, longer connection length or greater divertor radiative fraction.

1.2.3 Alternative Divertor Configurations

The conventional divertor has been in use for decades, and was the configuration envisaged when the coil arrangement for ITER was chosen. This is the configuration that was assumed in calculating the heat flux of $10\,\mathrm{MW\,m^{-2}}$ mentioned in section 1.1.4, and is shown in figure 4 (a) for MAST Upgrade. However, advanced divertor configurations, such as the snowflake, X-divertor and super-X, have the potential to lower the heat flux. Snowflake and X-divertors can be created with coil arrangements built with conventional divertors in mind. Having a divertor at the top of the tokamak, as well as at the bottom (sometimes known as a double null configuration) also reduces the target heat flux approximately by a factor of two.

The snowflake divertor configuration aims to reduce the heat flux to the plates by increasing the number of strike points, see figure 4 (c). The name snowflake originates from the six-fold symmetry of the X-point, which is a second order null in the magnetic field (other divertors use a first order null). In theory, multiple strike points mean that the same SOL power is deposited on multiple divertor target plates, decreasing the target flux. This necessitates additional plates at various angles relative to the X-point. However, having more plates to replace is preferable to plates failing mid-pulse. This configuration has been tested before on both the TCV [27] and DIII-D [28] tokamaks, but is still of interest because of its relevance to ITER and DEMO.

A different alternative divertor concept is the X-divertor, which reduces the target heat flux by means of poloidal flux expansion. As the name suggests, poloidal flux expansion refers to the spread of the field lines further in the poloidal direction across the plate. The X-divertor derives its name from the fact that an additional X-point is created near the strike point, albeit outside the plasma. The flux expands in the vicinity of this X-point, just as it does near the primary X-point (beneath the plasma core). This spreads the same heat energy over a larger surface area, reducing the load on the plates. A second advantage of an X-divertor is that detachment fronts have a tendency to remain close to the secondary X-point, and therefore close to the target. This is because of the underlying tendency of the detachment front to move towards lower magnetic fields (i.e. more expanded magnetic flux), which was mentioned in section 1.2.2.

However, the alternative divertor of most interest to this project is the super-X divertor configuration, shown for MAST Upgrade in figure 4 (b). This is a variant of the X-divertor, which requires increasing the distance between the bulk plasma and the divertor plates

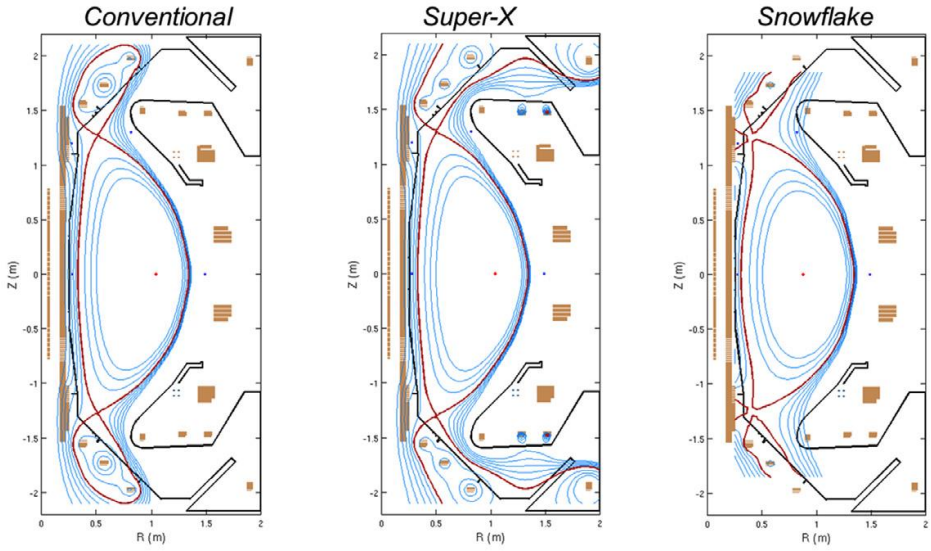

Figure 4: A comparison of the (a) conventional (b) super-X and (c) snowflake divertor configurations as they were predicted to appear on MAST-Upgrade. Figure credit CCFE.

(connection length). This configuration has the maximum major radius of divertor plates possible within the toroidal field coils, and can be designed for many hypothetical tokamaks and coil arrangements (provided that the toroidal field coils allow the divertor plates to be far from the plasma) [29]. This reduces the heat flux in two further ways (in addition to poloidal flux expansion and detachment).

Firstly, the super-X divertor decreases the heat flux by toroidal magnetic flux expansion, decreasing the heat flux parallel to the field line. The larger major radius of the plates allows the magnetic field lines to spread out toroidally, as well as poloidally, further decreasing the heat and particles deposited per unit area. This is especially true in a spherical tokamak. It should be noted also that the larger major radius location of the strike point results in the magnetic field line, and therefore particle path, undergoing an increased number of revolutions about the torus. This in turn reduces the angle between the field line and the plate. The angle between a magnetic flux tube and a surface affects the surface area onto which the flux is projected. This has a greater effect than the poloidal angle of incidence visible in figure 4.

Secondly, by increasing the connection length between the bulk plasma and the divertor plates, the super-X gives the escaped particles significantly more interaction volume in which to radiate heat away. The connection length is increased significantly more than is apparent from cross-sectional representations such as figure 4, because the particles spiral around in

the toroidal direction much faster than their radial velocity takes them toward the plates. This radiation of heat near the divertor could be further increased by impurity seeding, as described in section 1.2.2. This can also affect the position of the detachment front.

For all its advantages, the difficulty of the super-X divertor is that it requires a large divertor, with plates far away from the bulk plasma, see figure 4(b). This means that a tokamak must undergo a major upgrade programme in order to be suitable for testing the super-X configuration, or be built with it in mind from the beginning. However, an advantage of having a tokamak built with the super-X in mind is that a large divertor volume can be enclosed, which further aids detachment. It also improves core confinement by physically impeding impurity atoms from travelling upstream through the SOL and into the core, where they would radiate energy away. Even a partially closed conventional divertor has been shown to have this effect [23] [25]. It is the promising results from previous experiments on snowflake and X-divertors, and even more promising computational modelling work on the super-X, that has motivated the experiments on super-X divertors in MAST-U.

1.2.4 Power Balance

The basic principle of power balance is the accounting for all fluxes of heat into and out of the plasma, with the intention that they balance to zero. The two primary sources of heat are neutral beam heating, and ohmic heating (some tokamaks have a third: microwave heating), and the two primary heat sinks are radiated power and the SOL power. Energy can also be stored in the plasma during a pulse. This balances according to equation 3

$$P_{\mathrm{SOL}} = P_{\mathrm{NBI}}^{\mathrm{abs}} + P_{\mathrm{ohmic}} - P_{\mathrm{rad}} - \frac{dW}{dt} - X \tag{3}$$

where P_{SOL} is the SOL power (which can be estimated from divertor heat flux), $P_{\mathrm{NBI}}^{\mathrm{abs}}$ is the absorbed neutral beam power, P_{ohmic} is the ohmic heating power, P_{rad} is the radiated power, dW/dt is the rate of change of stored energy in the plasma (where W is the thermal energy resulting from work done on the plasma), and X is magnetic shaping power (which will be explained in 2.1.1). When this equation appears not to balance, either there is another source or sink of heat in the plasma that is not being accounted for, or the measurement technique is poor. This is also illustrated in the flowchart of figure 5. Note that the *absorbed* neutral beam power implies that the shinethrough and first orbit loss (see section 2.1.1) have

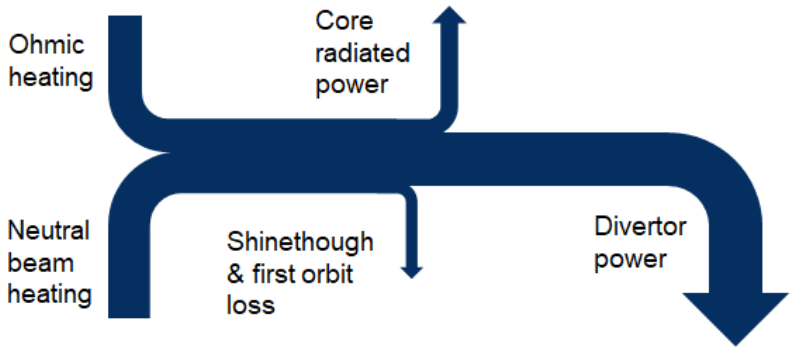

Figure 5: A diagrammatic representation of the primary flows of power in to and out of a tokamak

already been subtracted.

There are two complications that were not considered when measuring power balance in [30]. The first is the effect of tile shadowing on the way divertor heat flux is calculated by infrared cameras, leading to an over-estimation of the area over which to integrate. The second is the fact that not all of the neutral beam power is absorbed into the plasma. Part of the reason that these effects went unnoticed in [30] is because they happen to approximately numerically cancel each other out in MAST.

It cannot, and should not, be assumed that these discrepancies will approximately cancel out in future fusion devices. Also, results calculated based on this assumption have unnecessarily high uncertainty. This was a shortcoming of the earlier work by De Temmerman et al [30]. A more rigorous analysis is required, incorporating both of these effects into the calculation. The fraction of absorbed beam power can be calculated from computational modelling of the plasma. The fraction of divertor wetted area can be calculated geometrically from the angles of the tiles relative to the vessel wall and the magnetic field. Both of these effects are explained in section 2.1.1.

1.3 Outline of the Thesis

In this chapter, the field of fusion research has been introduced to the reader, and from there the theoretical grounding of this project has been built up, through materials properties and advanced divertors to the theory of power balance. In chapter 2 the experimental methods used in power balance are introduced, and broken down through advanced divertors and materials properties to the specific techniques used in this project, and how these are

used to make a contribution to the field. In chapter 3, the diagnostic techniques used for the experiments analysed in this thesis are explained, beginning with those used directly in tokamak power balance, then other key tokamak diagnostics, then those used in other experiments or that have little direct relevance to power balance.

In chapter 4, detailed methodology and results are presented from the power balance analysis of data from the original MAST. This includes a wide variety of shots which were performed for many different purposes at the time, and enables conclusions to be drawn about which tokamak parameters have the greatest effect on the power balance analysis. In chapter 5, the methods used and results obtained from experiments performed at Magnum-PSI are presented. These used some diagnostics used in tokamak power balance, as well as some others, and were intended to inform how material properties are accounted for in future power balance analysis. In chapter 6, preliminary analysis of results from the first experimental campaign on MAST Upgrade are presented, and these are compared and contrasted with the results from chapter 4, as well as those from previous work. Finally, chapter 7 contains the final discussion of the results, and the conclusions of the project. It also contains suggestions for future work, including further experiments on MAST Upgrade.

Chapter 2

Review of Previous Work

2.1 Previous Power Balance Studies

Power balance is the accounting of all the power flows into and out of a tokamak, as explained in section 1.2.4, and the equation and flowchart therein. It has been studied many times before, on most if not all tokamaks, including MAST. The motivation for power balance studies arises from the desire to maximise the energy efficiency of a tokamak reactor (that is, to maximise the energy produced by the reaction while minimising the input energy required to sustain the reaction). In order to properly assess the methods for improving energy efficiency, which usually relate to a specific source or sink of energy, all sources and sinks of energy must be quantified separately. This allows methods for increasing (or reducing) the energy gained from a specific source (or lost to a specific sink) to be compared numerically to those relating to other sources or sinks, and thereby determine the optimum method for maximising overall energy efficiency.

2.1.1 Methods in Previous Studies

Power balance studies can be either experimental or simulation-based. The scrape-off layer (SOL) power can not be measured directly like some of the other quantities, so the normal way to determine it from experimental measurements is to sum the heat flux at the inner and outer strike points of the divertor (and at both the upper and lower divertor for tokamaks with double divertors, like MAST and MAST Upgrade).

The first step in the tokamak power balance analysis is to measure the ohmic heating power being delivered to the plasma, which is the product of current and voltage in the ohmic

heating systems. Next the neutral beam power is determined, usually from the electrical input power to the NBI system, or possibly through computer modelling to account for shinethrough and first orbit loss.

Then the change in power stored in the plasma must be calculated. There are two types of stored power, the thermal energy resulting from work done on the plasma and the magnetic energy. The work done is calculated by multiplying the plasma current by the loop voltage. The magnetic energy is related to the shaping of the plasma. It is calculated by integrating the poloidal magnetic field around the last closed flux surface of the plasma, and is sometimes regarded as a modification to how the ohmic heating power is delivered rather than a truly separate component. The exact procedure for the calculation on MAST is given in MAST ops note 00.28. The third step is to measure the power leaving the plasma. Some of this is radiated directly out of the plasma core, and is measured using bolometers. The rest makes its way through the scrape off layer to the divertor plates.

The normal way to measure divertor power is with infrared cameras [30] [31], although Langmuir probes are sometimes used [26]. The camera data, in photon counts, is transformed into units of heat flux, to compare directly with the other quantities. There is always a discrepancy in any study of tokamak power balance, usually only $(80 \pm 10)\%$ of power going in is accounted for coming out (see table 1). This is thought to be because of power deposition onto areas of the first wall other than the divertor tiles, or possibly radiated power from near the divertor.

As mentioned in section 1.2.4, there are two previously overlooked factors that complicate the power balance calculation, as well as a third that has been accounted for in the past, but will be revisited here. Firstly, the divertor tiles of a tokamak are not perpendicular to the direction of the incident flux, and the surface of the divertor is not smooth, rather the tiles are angled relative to the surface to which they are attached. This means that each tile casts a shadow over the next tile, so the surface area on which the flux is incident is smaller than the total area of the tile. This is analogous to rain falling vertically onto an ordinary inclined, tiled roof and is illustrated by figure 6. When this is not accounted for, predictions of divertor tile temperature based on power in the plasma can be lower than what is observed. Conversely, calculated power put into the plasma based on divertor heat flux measurements can be too high.

The second complication arises from the primary mechanism for supplying heat to the

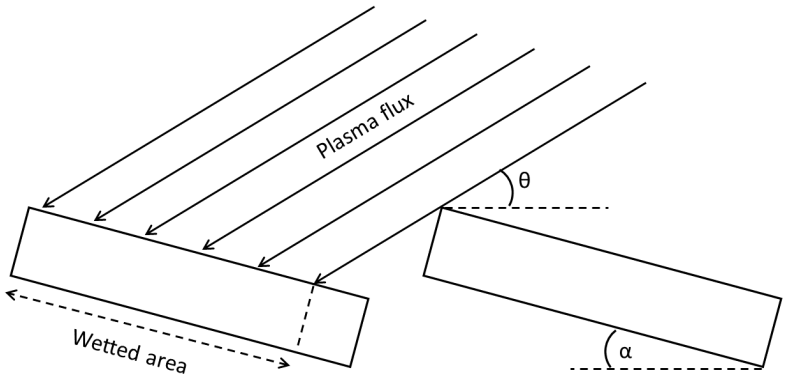

Figure 6: A diagram of inclined tokamak tiles and plasma flux, showing the origin of the wetted area consideration. θ is the angle between the vessel wall and the plasma flux, and α is the angle between the vessel wall and the divertor plates.

tokamak plasma, which is neutral beam injection (NBI). This supplies much more heat to the plasma than other means such as ohmic heating. It is sometimes assumed (wrongly), that all of the injected neutral beam power is absorbed into the plasma, but inspection of components removed from the vessel wall opposite the NBI reveal that a significant fraction of the power is transmitted through the plasma, and is incident upon the wall. This is called shinethrough, and when it is not accounted for, it leads to overly high predicted divertor flux (or low predicted heating power if the calculation is done in reverse). Some neutral beam power is also lost by way of being absorbed by particles which have escaped the last closed flux surface already, and carry the energy immediately to the divertor instead of depositing it in the core (first orbit loss).

The third, more well known problem with power balance calculations occurs when infrared cameras are used to measure divertor heat flux onto divertor tiles that have deposits of dust or surface layers. Unlike the two other effects mentioned above, this effect was routinely corrected for on MAST. A description of the effect, and the procedure used to correct for it on MAST, will be given in section 3.1. This procedure assumes that the most favourable result without any unphysical values is the correct one, but it lacks a true first principles or empirical basis. Determination of an empirical basis would require exposing a divertor tile to long, unchanging tokamak-like plasmas, which is not possible in an actual tokamak. A means of achieving this was desired as part of this project. This was the motivation for the work done at Magnum-PSI, which is presented in chapter 5.

2.1.2 De Temmerman 2010 Study

A previous power balance study on MAST was published by Greg De Temmerman et al in 2010 [30]. It has been mentioned above that wetted area fraction resulting from tile inclination, and neutral beam absorption fraction, were not accounted for in the 2010 study, and that this was part of the motivation for this project. The basic method outlined above was used in [30], although rather than treating divertor power as a single quantity, the four strike points (upper and lower, inner and outer) were considered separately, and the asymmetries (both deliberate and incidental) were key results.

The analysis used a combination of both medium and long wave infrared (IR) cameras; and concluded that long wave measurements are less sensitive to disturbances from hot spots on the surface (hot spots are small particles of dust with poor thermal contact with the bulk, and therefore heat up to a higher temperature than the plates). This conclusion was based on the tendency of the long wave camera to report lower temperatures than the medium wave camera, as shown in figure 7. This had been one of the purposes of the investigation; motivated by previous results from [32], which concluded that short wave infrared is more sensitive to hot spots than medium wave infrared.

A broad range of measurements were made, of the inner and outer divertors, and the upper and lower divertors. Measurements were taken in upper single null, double null, and

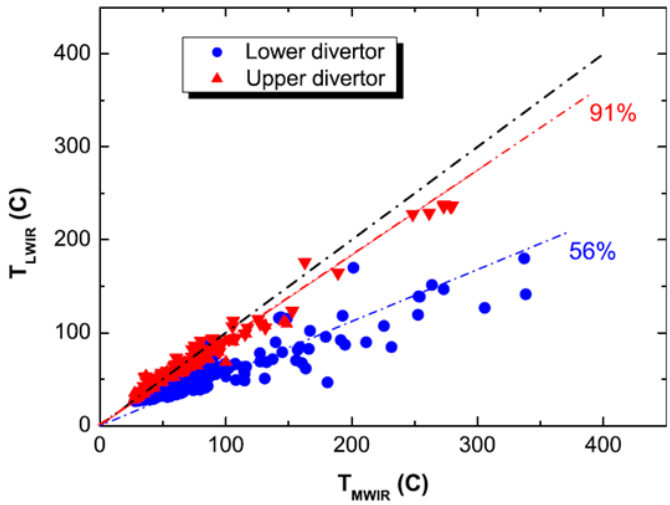

Figure 7: The temperature reported for both the upper and lower divertors of MAST, as measured by both the long wave and medium wave infrared (LWIR and MWIR). The percentages are the gradients of the lines of best fit. Both cameras were observing the same divertor for any given shot. Figure from [30].

lower single null discharges (this refers to the location of the divertor(s) within the tokamak). Some measurements were made with both cameras observing the same tiles simultaneously (such as those in figure 7), while others were done as comparisons between different divertor areas. In some cases, frame rate was increased at the expense of frame size. An internal heated tile was used for in-situ re-calibration of the cameras [30].

When discussing up/down asymmetries in a tokamak, the common numerical metric is δr_{sep}, which is the distance between the two separatrices at the outer midplane. A separatrix has been shown in figure 2 (b) as the inner boundary of the orange area. Because MAST has an upper as well as a lower divertor, it usually (in double null discharges) has a double separatrix, of which the top half is a mirror image of the bottom half, with the bottom half resembling that in figure 2 (b). In a single null discharge (upper or lower), which is caused by moving the plasma core up or down, there are two separatrices, one in contact with each divertor, with the smaller of the two associated with the divertor (upper or lower) towards which the plasma core has been moved. The horizontal distance between these two surfaces measured at the outer midplane is δr_{sep}.

Several asymmetries between the four strike points were identified. Firstly, it was found that more power tends to go to the lower divertor than the upper divertor because the $\boldsymbol{\nabla B}$ drift points down to the lower X-point in a normal MAST discharge, however the transition from upper single null to lower single null was smooth, see figure 8 (a). Also, it was found many times more power is deposited at the outer divertor than the inner, in double null discharges. It is noted that the fraction of power deposited at the outer target is significantly lower for lower single null discharges, but the same for upper single null discharges, see figure 8 (b).

The target heat flux measured in [30] accounts for about 90% of the heat believed to be entering the SOL, so it was thought to be a reliable measure. However, it is now known that the power entering the SOL, and the flux at the plates, are both lower, because no allowance was made for divertor wetted fraction, or for neutral beam shinethrough.

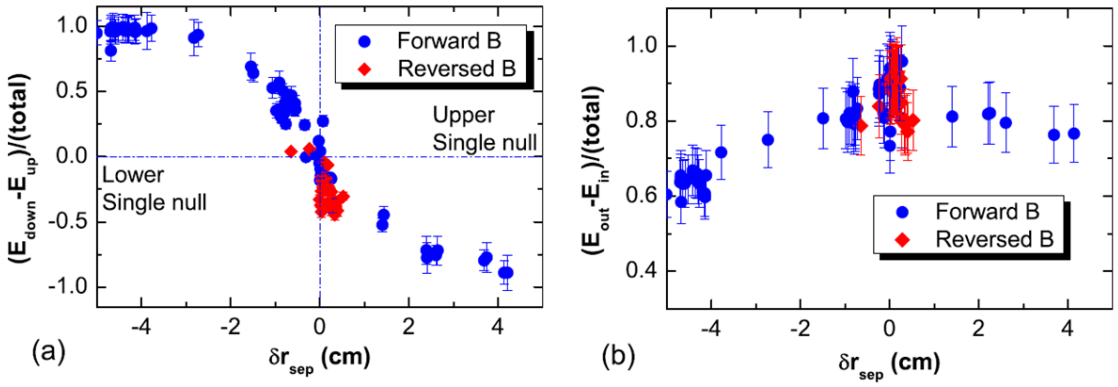

Figure 8: The difference between energy deposited at (a) lower and upper divertor and (b) outer and inner divertor, in both cases plotted against the distance between the separatrices at the outer midplane δr_{sep}, which is a proxy for lower single, double or upper single null discharges. Figure adapted from [30].

2.1.3 Other Experimental Studies

Numerous power balance studies on other tokamaks have been carried out over the years, on several tokamaks. Many of the key results from these are given in table 1. This is only a small selection of power balance studies, intended to give one example from each of a number of tokamaks. Each of them had secondary objectives, which may have been related to detachment or divertor configurations, or sometimes power balance was a secondary objective. Most used infrared cameras to measure divertor flux, which is the principal approach taken in this thesis.

In a 2017 investigation of detachment on the Tokamak à Configuration Variable (TCV) by Harrison et al. [26], the power balance on that tokamak was studied. This study measured divertor flux with Langmuir probes, not IR cameras. This can be used to mitigate the wetted area problem, although TCV does not suffer from this problem anyway because its divertor

Table 1: An overview of the diagnostics used and power accounting results of past experimental power balance studies on various tokamaks

Tokamak	Year	Divertor instrument	Power accounting [%]
MAST	2010	IR cameras	90-100
TCV	2017	Langmuir probes	Unknown
ASDEX-U	2001	IR cameras	70-80
JET	2016	Tile calorimeters	~ 75
NSTX	2005	IR cameras	70-90

tiles are not inclined. However, it requires a conversion from particle flux to heat flux, which assumes an ideal thermal plasma with $T_i = T_e$. It was found that before detachment, the radiated power from the divertor rises with the density, but stops rising once detachment sets in (red and pink lines in figure 9). This is contrary to the core, where radiation keeps rising over time (blue line in figure 9). Note that the density is rising over time throughout. In the outer divertor leg, the electron temperature is always below 10 eV for a detached plasma.

In 2001, power balance results were reported from the Axially Symmetric Divertor EXperiment Upgrade (ASDEX Upgrade), which had just been fitted with a new closed divertor (called LYRA) [33]. It can be seen in figure 10 that the total power measured leaving the plasma was about 80% of the power entering the plasma, both before and after the installation of the new LYRA divertor. However, the fraction of radiation originating in the divertor (rather than the core) nearly doubled to 40%. The same study reported that the heat flux to the divertor plates was halved when the LYRA divertor was fitted [33]. The closed super-X divertor on MAST-U aims to build upon the successes of this and other closed conventional divertors.

A comprehensive energy balance on JET was carried out in 2016, which used tile calorimetry to measure the divertor heat flux [34]. The analysis found that 75% of the power entering

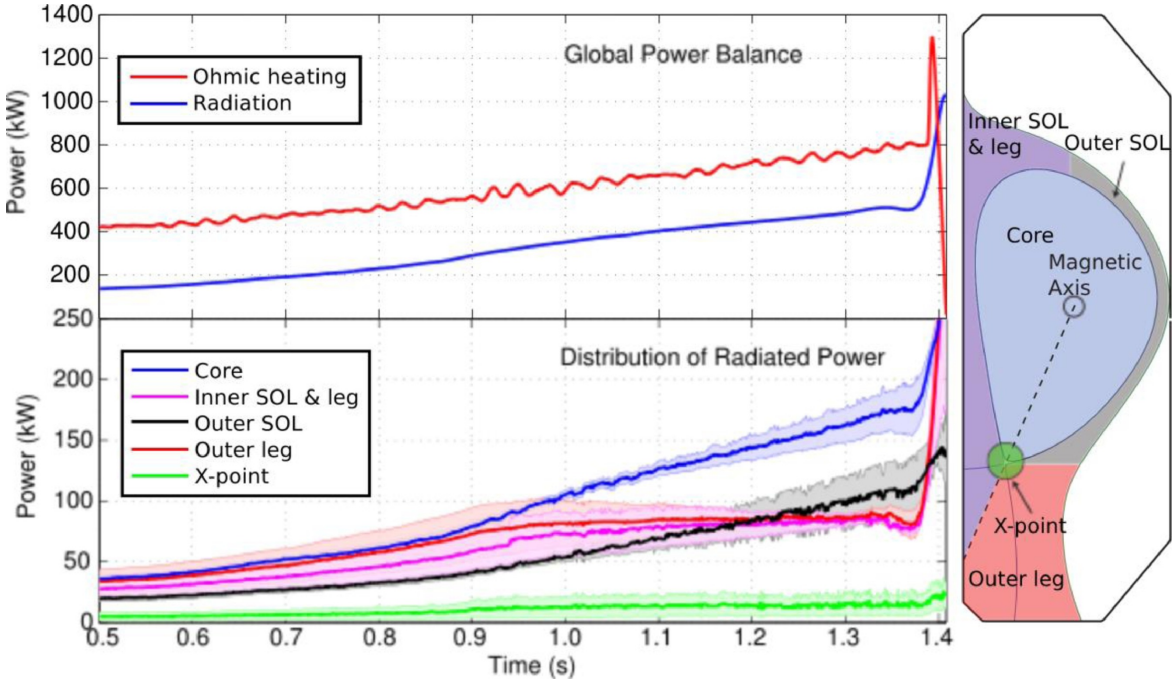

Figure 9: The global power balance (top) and distribution of radiated power (bottom) in TCV, where the regions are defined as in the inset image on the right. Figure from [26].

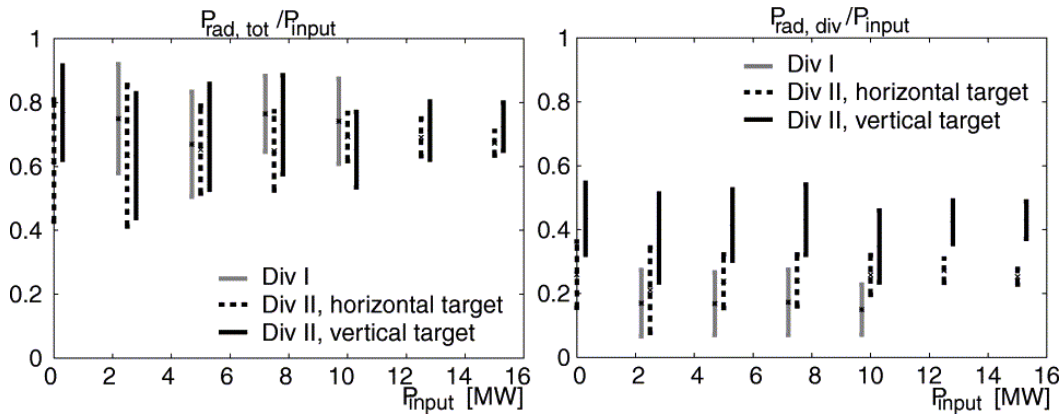

Figure 10: The total radiated (left) and divertor radiated (right) power from ASDEX Upgrade, as a fraction of the input power. Div I is the old divertor, Div II is the LYRA divertor, in which both horizontal and vertical surfaces are available. Figure from [33].

the plasma was recorded leaving it. The primary conclusion was that further work was needed to determine the source of the 25% loss, even though this is consistent with power balance studies on other tokamaks. The following year the data was analysed again with various correction factors, to see if any could explain the discrepancy. It was found that proper accounting of the energy stored in the plasma was the single most significant correction, and the authors recommended it be accounted for in future power balance studies [35]. Overall however, the search for a full explanation of the discrepancy was inconclusive.

A power balance was done on the National Spherical Tokamak eXperiment (NSTX) in 2005, for neutral beam heated H-mode plasmas [31]. Approximately 70% of the power was accounted for in lower single null discharges, and around 85-90% in double null discharges. It was noted that surface layers may cause reported temperatures to be higher than the actual temperatures (see section 2.3.2). It was also noted that some power flows not accounted for could have been to surfaces not covered by the IR cameras (i.e. the first wall outside of the divertor region).

2.1.4 Simulation-Based Studies

In an investigation of power balance by SOLPS 5.0 (Scrape-Off Layer Plasma Simulator version 5.0 [36]) simulations on MAST-Upgrade, it was found that about 60 - 80% of the heat out of the plasma was deposited on the target or vessel wall in attached conditions, with a conventional divertor [37]. The same investigation found, in detached conditions with a super-X divertor, only 40% of the heat leaving the plasma was deposited onto the target or

vessel wall. The ratio of wall losses to target losses was higher for the detached case, because more heat is radiated out of the larger divertor volume.

The problems mentioned earlier for experimental power balance, such as wetted area and neutral beam shinethrough, do not cause any problems in a simulation. It can be seen in figure 11 that around 100% of the power can be accounted for in these simulations. There are of course assumptions underlying each simulation model, which can lead to similar discrepancies. It can be seen that for the enclosed super-X divertors (iii and iv) very little power is deposited on the outer strike point, with a higher proportion being lost directly from the core. Column (i) which refers to simulations of the original MAST, can be compared to experimental results such as those from [30]. The experimental results also show that most of the power goes to the target in MAST (which operated with a conventional divertor only), although the experimental measurements can only account for around 80% of the total power.

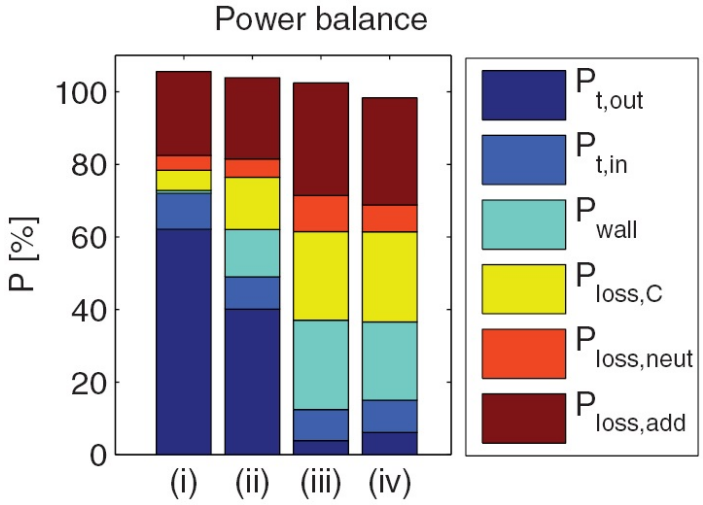

Figure 11: Power balance in simulations of MAST-U, showing 6 different power sinks for each of 4 scenarios. The scenarios are (i) MAST, (ii) MAST-U with conventional divertor, (iii) MAST-U with Super-X divertor 1, (iv) MAST-U with Super-X divertor 2. The power sinks are, from top to bottom: power to the outer target, power to the inner target, power deposited on the vessel wall, power radiated by carbon ions, power radiated by neutrals, and additional power loss caused by plasma-neutral interaction. Figure adapted from figure 12 of [37].

2.2 Previous Work on Advanced Divertors

Advanced divertors were introduced in section 1.2.3, and their implications for power balance were explained. Having explained the methods used in experimental power balance studies, past experiments involving advanced divertor configurations will now be reviewed.

2.2.1 X-Divertor Experiments

The X-divertor configuration was tested on TCV in 2017 [17]. The poloidal flux expansion was varied by a factor of 10, and the total flux expansion increased by 70%. Both of these are characteristic of the X-divertor. The conclusion was that deeper detachment results from greater flux expansion for given line-averaged density, but there was no effect on the detachment threshold. Good agreement was found between theory and experiment, so it was hoped that theoretical predictions would be similarly correct for the super-X.

The X-divertor has also been studied on the DIII-D tokamak, with matching experiments and modelling [38]. These experiments reported 10-20% lower upstream density and higher H-mode pedestal pressure than a conventional divertor. It was remarked that there must be both poloidal flux expansion and divertor closure in order for detachment to come about at a lower threshold, and that this was confirmed by simulations as well as the experiments. To finish, [38] suggests further work involving the super-X divertor (such as that on MAST Upgrade described in chapter 6).

2.2.2 Snowflake Divertor Experiments

The snowflake (SF) divertor was tested on TCV in 2009 [27]. This requires the creation of a second order null point. Both SF+ and SF- configurations (with two separate but very close nulls) were created, as well as standard SF configurations (with a true second order null). The resulting configurations increased the connection length and flux expansion relative to the conventional divertor by more than a factor of two. The summary of [27] notes that SF+ and SF- may have different implications for ELM stability (ELMs were mentioned in section 1.1.3).

The snowflake divertor has also been studied on DIII-D [28]. Snowflakes were maintained for several seconds in several high power, high current H-mode discharges. Snowflake divertors were found to have near total detachment and a larger radiative volume in the divertor,

as opposed to the partial detachment and localised divertor radiation more familiar from conventional divertors. These results complement other experiments on snowflake divertors, on NSTX as well as on TCV (above).

2.2.3 Progress Towards the Super-X Divertor

A comprehensive review of advanced divertor configurations, including the snowflake and X-divertor configurations, was undertaken in 2017 by Soukhanovskii [39]. This review stated that investigation of the Super-X divertor on MAST-Upgrade was a necessary next step to understanding advanced divertor configurations. While this was a review of past experiments; there have also been theoretical investigations of, and preliminary experiments with, a super-X divertor.

There has been an attempt to produce a super-X divertor on TCV [17]. This was done by varying the major radius of the outer strike point, R_t, while keeping the connection length and poloidal flux expansion constant. These results confirmed the prediction that $q_{\|} \propto 1/R_t$ (where $q_{\|}$ is the parallel heat flux and R_t is the target radius), but did not find the detachment onset at lower density for larger R_t. This was one of the expected effects on detachment that had been predicted for the super-X, but was not observed because the neutral density near the target was lower than that predicted by the two point model. MAST-Upgrade hoped to overcome this problem by increasing neutral density near the target using a divertor baffle (i.e. an enclosed divertor), as shown in figure 4 (b). It is also worth noting that, because MAST Upgrade is a spherical tokamak, the toroidal magnetic field decreases faster with increasing R, resulting in greater toroidal flux expansion at the plates than on TCV.

A detailed computational study of the super-X divertor was published in 2009 by Valanju et al [29], which found it was possible to create the super-X divertor using similar coils to those used for conventional divertors. It was stated that the super-X divertor is the only divertor concept that simultaneously increases the plasma wetted area by about 1.5-3 times, increases connection length by about 2-5 times, increases SOL width via diffusion by about 1.4, lowers plasma temperature at the plate to below 10 eV, and increases maximum divertor radiation fraction up to 50%. These effects combine to result in the same target flux for up to 5 times higher SOL power. This makes the super-X divertor far superior to flux-expansion-only methods, such as X-divertors, snowflakes and plate tilting [29].

Another simulation-based investigation of the super-X divertor was carried out by

Havlíčková et al. in 2015 [37]. This showed a lower detachment threshold for the super-X divertor (the plate being at a much lower temperature for the same core plasma conditions), transitioning at 3 times lower density and 4 times lower power. The increased major radius of the target was found to reduce the target power load by a factor of 25. It was also found that good divertor enclosure substantially increased the neutral density in the divertor region, with the corresponding power losses.

As can be seen in figure 12, the super-X divertors 1 and 2 simulated here greatly reduced the heat and particle flux to the plates, compared to a conventional divertor (and even more so relative to MAST). Similar reductions were seen in the electron temperature at the target (see bottom middle of figure 9 of [37]). The reason there were only limited returns for SXD3 is because it had poor enclosure, because the baffle was removed in the simulation. SXD1 and SXD2 were both well enclosed by the baffle. An illustration is given in [37, fig.7]. In the real MAST Upgrade, the baffle is permanent, so the divertor is always very well enclosed with a long-legged divertor such as a super-X. Numerical computational validation of the super-X divertor concept from work such as this set the precedent for the experimental work on MAST Upgrade.

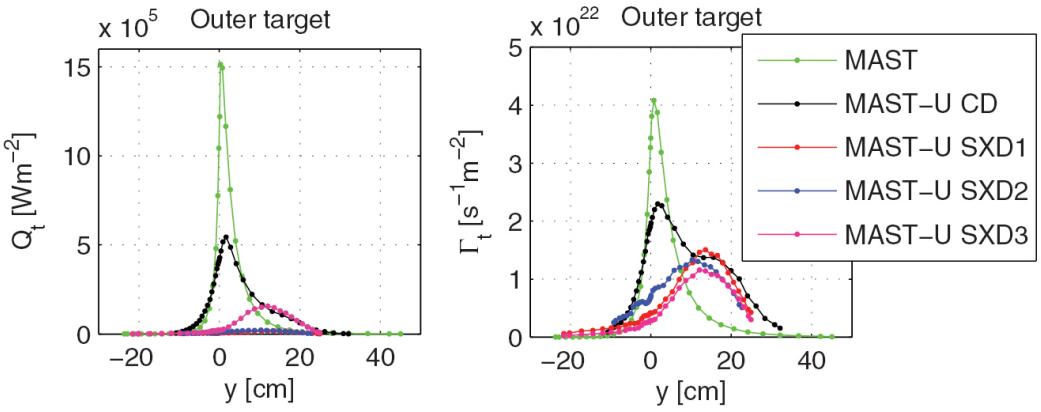

Figure 12: The simulated heat and particle flux to the outer divertor in MAST-Upgrade, against the distance along the divertor target y, with different divertor configurations (note that SXD3 has poor enclosure). Figure from [37].

2.3 Plasma Surface Interaction Experiments

The importance of materials science to fusion research was discussed in section 1.1.4, and this was developed into a motivating factor for advanced divertors and power balance studies, in terms of reducing the material incident heat flux. Although power balance is essential for determining the effectiveness of different materials in the tokamak, the material choice also affects the analysis of power balance results. Large future tokamaks will have metallic walls, unlike the carbon walls considered here, so the effects of carbon walls must be properly compensated for to enable extrapolation to future machines. Some experiments of interest to this are reviewed here.

2.3.1 Hydrogen-Carbon Interactions at the Plasma-Surface Interface

There are three primary ways in which the plasma can interact with and change the surface of a material. These are chemical reactions between the plasma and the surface, erosion of the surface, and erosion followed by re-deposition of surface material, changing the topology but not the chemical composition of the surface. Here these are discussed in the context of a hydrogen plasma and a carbon tile, which is used in many nuclear fusion experiments, included all those reported in this thesis.

In plasmas with low ion energies ($\lesssim 100\,\mathrm{eV}$), chemical interactions between the plasma and surface dominate over physical interactions. In the case of hydrogen and carbon, this leads to the formation hydrocarbon films on the surface of the tile, in a process called chemical sputtering [40]. Such a film then behaves as a surface layer with poor thermal contact with the bulk. The consequences of these layers will be discussed in section 3.1.

In plasmas with moderate or high ion energies ($\gtrsim 100\,\mathrm{eV}$), the hydrogen is more energetic and will not linger and react with the carbon. It instead knocks carbon atoms out of their place on the material surface (erosion). At moderate ion energies, much of this carbon will be re-deposited, leading to the growth of crystal structures on the surface of the tile. This has been observed on Magnum-PSI in 2012 by Bystrov et al [41]. In this case, the growth was in the form of spherical nanoparticles, which is typical for carbon. In a follow up experiment the following year, using a tungsten substrate and injected carbon gas, various shapes including carbon nanotubes, spherical nanoparticles and nanotips were detected using electron microscopes [42]. It is well known that increasing the surface roughness of graphite

increases the surface emissivity [43].

In plasmas with high ion energies ($\sim$ keV), atoms of carbon are ejected at speeds too high to fall back and re-deposit, so the material surface is eroded over time. Small protrusions from the surface are more likely to be eroded, so erosion without re-deposition can result in smoothing of an already rough surface. On the other hand, erosion of a smooth surface may result in the surface of the tile being composed of atoms in many different layers of carbon. Such a structure has a higher emissivity than an undamaged graphite tile, because rays may be reflected within these cavities many times before being reflected away. In terms of emissivity, this increase in surface roughness is likely to produce the same effect as the surface roughness introduced by erosion and re-deposition described above.

2.3.2 Material Surface Effects Observed in MAST and ASDEX Upgrade

Plasmas in MAST were typically at the lower end of the plasma power scale described above, so chemical interaction between the hydrogen plasma and carbon wall dominated the plasma-surface interactions. This lead to the formation of hydrocarbon surface layers on the divertor tiles, which have poor thermal contact with the bulk tile.

The same effect has also been observed in ASDEX Upgrade, and the measurement errors introduced were found to have a dependence on the wavelength of infrared camera, in both MAST [30] and ASDEX Upgrade [44]. It is noted that this effect always raises the observed temperature (although other effects exist that can lower the observed temperature, as will be discussed in chapter 5). It is also mentioned in [44] that surface layers cause measured temperature to rise quickly initially, and to fall quickly when the heat flux is removed, and also that hot spots are a cause of higher than expected temperature readings. These effects will be discussed in section 3.1.

Plasmas in Magnum-PSI are typically higher power, so plasma-surface interactions are dominated by erosion. Erosion of the tile surface decreases the emissivity of the surface, causing the infrared cameras to report a lower temperature. This results in an apparent decrease in tile temperature over the course of a campaign. This is the opposite of what was observed in MAST and ASDEX Upgrade. Magnum-PSI is fitted with a multi-wavelength pyrometer, which can be used to calibrate the infrared camera, but this method is not foolproof, as will be explained in section 5.2.3.

2.4 Previous Work with TRANSP

TRANSP is one of a very large number of plasma transport codes that are used in fusion research. The basic outline of typical transport codes, as well as some specific examples, are given in [45]. Many past uses of TRANSP have been validated directly against experimental results, which is more important than for other codes because TRANSP is closed source, meaning experimental validation is its only source of credibility. In addition to early validation against TFTR [46], there have been several more recent validations of TRANSP against JET [47] [48], and this continues even today [49].

At JET, another code called PENCIL is used as well as TRANSP, which can also calculate neutral beam power losses. Both are run routinely between shots, but this was not the case for MAST, where TRANSP was only run occasionally for specific purposes. There was no routine accounting for beam power losses on MAST. The TRANSP neutral beam injector module, NUBEAM, has been used to simulate shinethrough of neutral beam power on other tokamaks such as EAST [50]. The motivation of that study was to validate their own simulation code's predictions regarding neutral beam shinethrough. This provides precedent for the use of TRANSP to calculate neutral beam losses on MAST in this thesis.

2.5 Areas for Further Clarification

The work reviewed in this chapter leaves unanswered several questions that will need to be answered if fusion energy is to be realised. Certain results from past tokamaks need to be extrapolated with some confidence to future tokamaks, while effects that have been overlooked or only cursorily treated in the past need to be investigated further, and there are simulations that need to be validated experimentally. With reference to earlier sections of this chapter, the following questions stand out and will be addressed (or an attempt will be made to address them) in the remaining chapters of this thesis.

Firstly, with reference to section 2.1.2, there is the question of how the effects of tile shading (wetted area) and beam power absorption affect the overall power balance of a tokamak, especially MAST. A geometric correction for tile shadowing, and a simulation-based approach to beam power calculations (similar to the work referenced in section 2.4) are required. This question is addressed for MAST in chapter 4 and for MAST Upgrade in chapter 6.

Secondly, with reference to section 2.3, there is the question of how the evolution of the tile surface over time affects the infrared thermography measurements used in the power balance. This is a known problem on many tokamaks. An attempt is made to address this question in chapter 5, although as explained therein, it was not possible for the results to inform the analysis of MAST Upgrade as originally intended. There remains scope for further work in this area.

Finally, with reference to section 2.2.3, there is the question of the effectiveness of the super-X divertor when utilised experimentally, especially its effect on power balance. This question is addressed in chapter 6, using methods of power balance similar to those referenced in section 2.1 and used in chapter 4. Power balance and super-X divertors will both continue to be important areas of fusion research going forward.

Chapter 3

Diagnostic Techniques and Simulation Tools

The power balance analysis required for this work, as outlined in section 2.1.1, requires data from a number of diagnostics. Because the analysis is concerned with heat flows, many of these diagnostics make thermal measurements, such as IR thermography, bolometry and pyrometry. Some non-thermal diagnostics are also required for power balance. In addition to this, the problem of surface layers that ultimately motivated the experiments on Magnum-PSI, requires certain other diagnostics that are not used in power balance. The problem of NBI power absorption also requires simulations in addition to the experimental diagnostic data. In this chapter, all the diagnostic techniques and simulation tools used in this thesis will be introduced.

3.1 Infrared Thermography

Infrared (IR) thermography is a technique which uses infrared cameras to determine the temperature change of, and heat flux to, a material surface. It takes place in three primary stages, a conversion from detector photon count to light intensity, then a conversion from light intensity to temperature, then a conversion from temperature change to heat flux evolution. The third stage must account for the emissivity of the material, including any surface layers. IR thermography is used extensively on tokamaks and other fusion devices to measure the heat leaving the plasma that is incident on the divertor.

Infrared cameras use a charge-coupled device (CCD) or other photon detector to de-

tect the infrared light arriving from the target. The high-specification IR cameras used in scientific research usually have InSb detectors, actively cooled with Stirling engines. The frame size is often reduced to only a part of the total detector area, to increase the temporal resolution of the measurements while remaining within the bandwidth limitations of the equipment. The detector counts the number of photons arriving at each pixel, within the wavelength range of the camera. The wavelength range in a scientific IR camera is normally narrow, so a uniform transmission through air can be assumed. There is not an exact match between photons emitted by the source and those arriving at the camera, because of absorption by air molecules, so instead a theoretical conversion between temperature and emitted photons is done, then an empirical calibration is used to determine the dependence of detected photons on emitted photons.

The calibration procedure for both MAST and MAST-U used a black body laboratory heat source, and involved varying both the temperature of the source, and the integration time of the camera, which can also affect the number of photons detected. The range of temperatures and integration times over which to calibrate for MAST-U was based on those expected to be observed and used in experiments (informed by those observed and used on MAST), and on the capability of the heat source. The temperature of the blackbody source is verified using a thermocouple (see section 3.5.3 below). It is important that the source used is a blackbody source, with emissivity $\varepsilon = 1$, as this would affect the calibration.

The distance between the camera and source during the calibration is unimportant because of the etendue. Briefly, increasing the distance between the camera and the source clearly reduces the intensity of infrared radiation from any given point on the source arriving at the detector, following a standard $1/r^2$ dependence. However, because the source is now further from the camera, it occupies a smaller number of pixels in the camera frame, such that there are more infinitesimal points of the source visible in each pixel of the camera, thereby increasing the detected intensity. These effects cancel out exactly, so the distance from the camera to the calibration source does not need to be considered at all. Conveniently, this also applies to the distance from the camera to the divertor plates during operation.

This calibration was performed once, then a lookup table was created which was used throughout the experimental campaigns to convert camera data to temperature. Once the temperature evolution is known, the conversion to heat flux proceeds by numerical integration of a form of the thermal diffusion equation given in [51] as

$$\frac{\partial U}{\partial t} = D\nabla^2 U \tag{4}$$

where U is the 'heat flux potential' as defined in [51], t is time, D is the temperature dependent heat diffusion coefficient (also defined in [51]), and ∇^2 is the Laplacian operator. This then allows heat flux to the surface, q_s, to be calculated from

$$q_s = -\frac{\partial U}{\partial x} \tag{5}$$

where x is the spatial coordinate normal to the tile surface. On MAST and MAST-U, this numerical integration is done using a computer code called THEODOR (thermal energy onto divertor), written by Albrecht Hermann of IPP, based on his work in [51]. The code accounts for things that would affect the emissivity, such as surface layers or hot spot formation, using a parameter α defined as

$$\alpha = \frac{\lambda_{\text{layer}}}{d} \tag{6}$$

where λ_{layer} is the heat conduction coefficient and d is the surface layer thickness [52]. THEODOR takes this parameter as an input. On MAST and MAST-U, the process for determining the correct value is a process of trial and error, that gradually increases the correction (by lowering the value of α) until negative heat fluxes are minimised in the result (there are also other methods that can be used, as noted in [53]). Negative heat fluxes in this context are unphysical, and their origin in this calculation is explained below.

In an ideal scenario, a constant heat flux for a short time is assumed, as in figure 13(a). This results in a temperature profile like that in figure 13(b). However, over time on MAST a build up of surface layers occurred, and this is also expected on MAST-U. A surface layer heats up much more quickly than the tile itself, because it has poor thermal contact with the bulk, and so there is less thermal mass to be heated. It then heats further in the normal $T \propto t^{1/2}$ dependence as the heat is conducted to the tile. A similar effect is observed in the cooling, where the surface layer cools straight away, then the tile cools as expected. This is all shown in figure 13(c). If it assumed that there is no surface layer, and that the bulk tile temperature was actually changing as in figure 13(c), then the calculated heat flux profile that results will resemble figure 13(d), which has a large and unphysical negative heat flux near the end. Using the correct value of α in the calculation eliminates these false negative

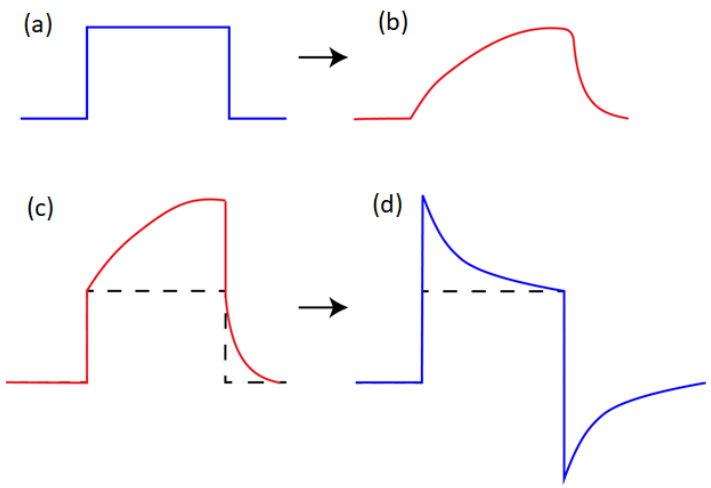

Figure 13: Graphs showing (a) ideal heat flux against time (b) the resulting tile surface temperature change over time (c) a realistic tile surface temperature change if a surface layer is present (d) the calculated heat flux from that temperature data if no surface layer correction is made. Figure adapted from [53].

heat fluxes from the results of the calculation.

Another consideration when setting up an infrared thermography system is which wavelength of infrared radiation is to be used. Thermography relies on use of a narrow wavelength range of infrared radiation, since total IR radiation across the spectrum does not vary strongly with temperature. Intensity at specific wavelengths varies more, and at certain wavelengths it varies more within certain temperature ranges. The wavelength range must be one for which air is transparent. There are three wavelength ranges used for different applications; $0.4-2\,\mu m$ (short wave), $2-5.5\,\mu m$ (medium wave), and $8-14\,\mu m$ (long wave). Air is opaque to radiation in the $5-8\,\mu m$ range. Long wave infrared (LWIR) has greater sensitivity to temperature changes below 300 °C, which is the temperature expected most (but not all) of the time in a tokamak like MAST or MAST-U. Medium wave (MWIR) is more sensitive to the higher temperatures reached during ELMs, or because of surface layer effects, or on a higher power device like Magnum-PSI. A combination of MWIR and LWIR are used on MAST and MAST-U, with LWIR generally looking at the regions where more surface effects were expected, as LWIR is less sensitive to the higher temperatures associated with dust accumulation. Only MWIR is used on Magnum-PSI.

3.2 Bolometry

Bolometry is a technique that measures radiated heat energy, and is used in fusion devices to measure the thermal radiation originating primarily in the plasma core. This is distinct from the heat carried away by energetic particles, which mostly pass through the scrape-off layer and onto the divertor. The bolometer principally consists of a thin foil and a pinhole, which provides a line of sight from the foil to the bulk plasma. The foil heats up as a result of the incident thermal radiation, according to the equation

$$P = c\left(\frac{\Delta T}{\tau} + \frac{\mathrm{d}\Delta T}{\mathrm{d}t}\right) \tag{7}$$

where P is the power, c is the heat capacity of the piece of foil (in $\mathrm{J\,K^{-1}}$), ΔT is the change in foil temperature, τ is the cooling constant and t is time. It is typical, and is the case on MAST and MAST-U, that an array of several bolometers is used to build up a picture of how much heat is being radiated from the entire plasma. This may include the divertor region, even though most radiated power originates in the core. The specific technique for divertor bolometry was developed in [54].

The heat from each piece of foil is thermally conducted to a nearby thermistor, which causes a decrease in the electric current flowing through it. Adjacent to each of these is another foil thermistor that does not have a pinhole, and is simply at room temperature. The difference between the currents flowing in the two circuits is used to derive the temperature difference. Equation 7 can thus be re-written as

$$P = \frac{1}{S}\left(V + \tau\frac{\mathrm{d}V}{\mathrm{d}t}\right) \tag{8}$$

where S is the instrument sensitivity (in $\mathrm{V\,W^{-1}}$) and V is the voltage across the thermistor. However, the difference in current between the two thermistors is too small to measure directly, so a 'Wheatstone bridge' is used as a form of amplification. The use of lock-in amplifiers for this application was developed at IPP Garching, and the basic setup is shown in figure 14. Because the voltage to be measured is still very small, the measurement makes use of amplitude modulation, which applies a fixed frequency, allowing the voltage to be measured at this frequency only. This reduces the noise in the measurement. The cooling time of the foil must also be known; using a thin foil with a short cooling time improves the temporal resolution.

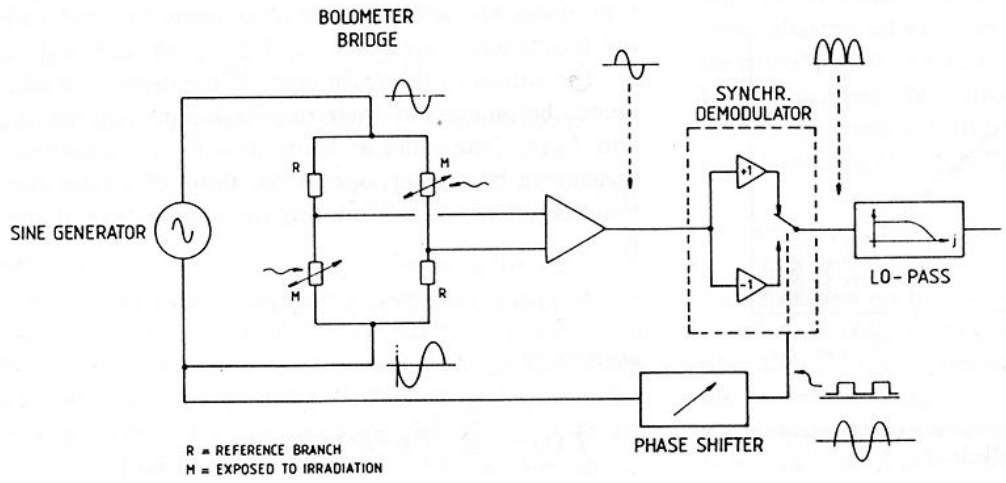

Figure 14: A circuit diagram of a lock-in amplifier of the kind developed at IPP Garching and used in bolometry. Figure from [54].

As an alternative to an array of gold foil thermistors and pinholes, greater spatial resolution may be achieved with a single target being observed by an infrared camera. However, because of the inversion of the image, the analysis of this data is not trivial. This type of bolometer, called an infrared video bolometer (IRVB), is a relatively new diagnostic being tried on MAST-U, which was developed at Alcator C-Mod in 2016 [55], and will be mentioned in more detail in section 6.2.2.

3.3 Thomson Scattering

Thomson scattering (TS) is named for the physicist J J Thomson, who first described it in his 1907 book "The Corpuscular Theory of Matter", a summary of which was published as an article in Nature in 1908 [56]. The theory of Thomson scattering consists of a perfectly elastic collision between a photon and an electron, where the photon energy is much less than the electron mass. The photon energy is absorbed by the electron, which then oscillates at the same frequency, and emits a photon of the same energy. Despite having the same energy (leaving the electron with its initial kinetic energy also), the photon and electron both change direction (and therefore momentum) during the interaction. Because the photon has changed direction, it is said to have been scattered by the electron. The angle through which the photon is deflected is dependent on the electron energy. In a macro-scale plasma, Thomson scattering results in Doppler broadening of the laser wavelength profile, from which the

electron temperature can be calculated. The amount of light reflected can also be used to calculate the electron density.

The principle of the use of TS for tokamaks, according to [6, p. 514] can be summarised as follows: The power scattered per unit solid angle, $P_s(\omega)$, is given by

$$P_s(\omega) = P_0 \; r_e^2 \; \sin^2 \psi n_e \; L \; S(k, \omega) \tag{9}$$

where P_0 is the total incident laser power, r_e is the classical electron radius, n_e is the electron density, ψ is the angle between the incident and scattered rays, L is the interaction length and $S(k, \omega)$ is the spectral density function. The spectral density function S is dependent on the velocity distribution of the electrons in the plasma. The magnitude is given by the fraction of electrons moving with a velocity (ω/k) along the $\hat{\mathbf{k}}$ direction. The velocity dependence and known magnitude can be used together to calculate the velocity distribution, and therefore electron temperature. The known value of S can then be substituted into equation 9 to give the electron density, n_e. Because of the very low photon yields of this technique, the photons are generated using a high power laser system, usually in the 100 MW range [6, p. 515].

3.4 Langmuir Probes

A Langmuir probe is a conductive probe that has been inserted into a plasma for the purpose of measuring the electric potential of the plasma. Irving Langmuir first used his probes in studies of electric discharges at low pressures in 1924, and published a more complete description, co-authored with H Mott-Smith, in 1926 [57]. Langmuir probes can be used to measure the electron temperature and density of a plasma, however the physical nature of the probe means it can only be used at the plasma edge. To measure the electron density and temperature within a large plasma volume, Thomson scattering is used (see section 3.3).

In the context of a tokamak, a Langmuir probe is an electrode, insulated from the structural steel and the vacuum vessel. It is held at a bias voltage, which is usually varied over time, and the current passing through the probe is measured. This current-voltage dependence allows the electron density and temperature to be inferred. The calculation makes use of the plasma sheath theory, which was mentioned in section 1.2.1. These probes are small, and there are usually many of them, in order to give a better spatial resolution. However,

the probe head must be much larger than both the ion and electron gyro-radii (the ion gyro-radius is much larger than the electron gyro-radius anyway), and also the Debeye length of the plasma [6, p. 546].

When the probe is sufficiently negatively biased, only the ion current remains. This is called the ion saturation current, and comes about because the ion flux across the thin sheath is limited by the ion sound speed. Likewise, if the probe were sufficiently positively biased, there would be an electron saturation current proportional to the electron sound speed [6, p. 546], although this is less useful and is rarely done in practice.

Any flowing current requires a complete circuit, and this circuit must be set up in such a way that the probe current is not limited by the maximum current of the other electrode. In a tokamak, the return electrode is the vessel itself, namely the divertor tiles that are in contact with the plasma. These are much larger than the probe area in MAST-U. A typical current against voltage (I/V) graph for a MAST-U Langmuir probe is shown in figure 15.

Figure 15 is generated by varying the voltage bias at the electrode over time, and measuring the current flowing through the probe. The ion saturation current I_{sat} is the current value shared by all the data points with low voltage (approximately $-20\,\text{mA}$). The red line is the absolute magnitude of this value $|I_{\text{sat}}|$ (i.e. $+20\,\text{mA}$). The blue fit line is fitted to only the points where the current is below this value and the single first point above it. The current-voltage dependence can be seen to be strongly non-linear beyond this point; all the current measurement range of the instrument occurs within a very narrow voltage range. This is because the ion temperature obeys the Maxwell-Boltzmann distribution, so there are very few ions at the higher temperatures. The very small number of ions from which the high-current points are determined is why they are not given any weight in the curve fitting.

Most Langmuir probes are fixed in place, usually in gaps between divertor or limiter plates, to characterise the edge plasma where it hits the wall. Some probes are instead mounted on an actuator, so that they can be moved into and out of a plasma during a shot. This can give more information of the electron temperature and density further into the plasma, although the presence of the probe can create a small sheath around it and therefore change the plasma conditions. This can be approximately corrected by a factor of 2, to correct for the drop in density when the ions are accelerated to the sound speed by the probe sheath [6, p. 548]. If the probe is left in for too long, it can also be eroded, causing damage to the probe. With a large number of probes, this also increases the likelihood of a

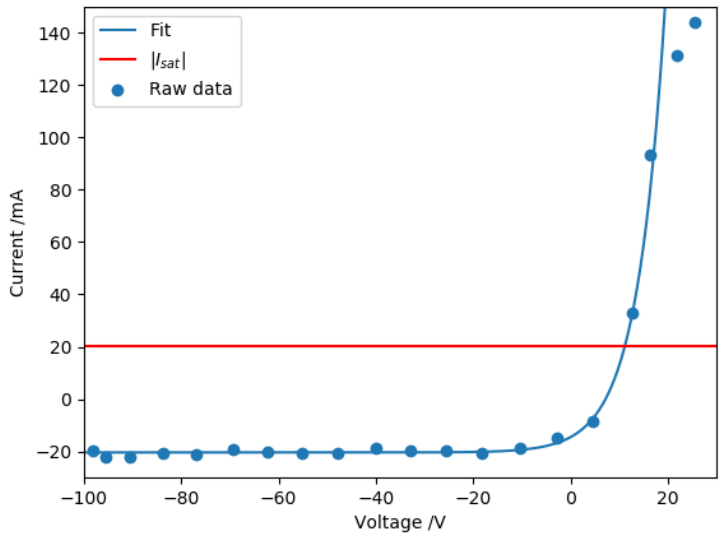

Figure 15: A typical Langmuir probe I/V curve from MAST-U showing the current-voltage dependence. The curve is fitted to all blue points from the left until the first one that is above the red line (the magnitude of the ion saturation current). Figure supplied by CCFE.

disruption.

Since each electron or hydrogenic ion has a known charge, the measured current can be used to derive the particle flux (of ions or electrons). This measurement of particle flux, combined with temperature and density, can be used in power balance. This technique will be mentioned in chapter 6. Unfortunately, Langmuir probes do not give a measurement of T_i, so the assumption that $T_i = T_e$ must be used [22, p. 87]. This is usually valid for high temperature ion-heated plasmas with high collisionality (i.e. high n_e/T_e^2) such as those found in tokamaks [58].

3.5 Other Diagnostics

3.5.1 Multiwavelength Pyrometry

Pyrometry is a technique that works by a similar method to infrared thermography (see section 3.1). However, modern pyrometers use multiple wavelengths to eliminate the need to assume a "blackbody" and allow an actual emission curve to be fitted, as a function of wavelength. The techinque is an evolution of earlier single and dual wavelength pyrometers. Single wavelength pyrometry suffers from the same susceptibility to unknown emissivity as

IR thermography, relying entirely on lookup tables of emissivity values, which usually do not account for wavelength or other factors, nor for changes in material properties over time. With two-wavelength pyrometry, two different images at different wavelengths are compared, by dividing one by the other. This cancels out the emissivity in the case of grey bodies (bodies with the same emissivity at all wavelengths), and also calibrates for changes over time, but does not allow for the different values of emissivity at the two sample wavelengths in other, less ideal bodies.

Multiwavelength pyrometers take a measurement at many wavelengths, usually over 100 in modern instruments. As a result, this technique does not suffer from the problem of the emissivity of an object raising the apparent (measured) temperature, even for non-ideal, non-grey bodies. This technique was validated against a thermocouple in 1989, using a pyrometer sampling just 10 wavelengths [59]. This technique is routinely used to determine the emissivity of tungsten samples at Magnum-PSI, although it is less reliable for carbon samples because of the very high emissivity and low plasma temperature tolerance of carbon. High emissivity requires more of the sampling wavelengths to have a signal well above the noise level, but the temperature required for this is above that which can be tolerated by a carbon sample (because the plasma flux would be above the ablation threshold, even at temperatures well below the melting temperature of carbon).

3.5.2 Laser Profilometry

Laser profilometry is a technique for measuring the surface topology of a sample. It uses a laser rangefinder to measure the distance between a probe head and the sample below. This laser is moved across the surface using piezoelectrics to ensure a high spatial resolution. The probe can move in 2 (horizontal) dimensions, while measuring in the third (vertical) dimension. This technique is used in industry to measure the depth of cracks in materials, such as from damage, or to assess a new manufacturing process.

A laser ragefinder works by measuring the displacement of the reflection of a laser beam, and inferring the distance from the source to the surface. Given the known angle of incidence, the angle of reflection varies only with variation in the surface, and this can be calculated by the position along the detector to which the light is reflected. This is usually a specific pixel along a linear CCD. Usually, the sample is enclosed to prevent reflection from other light sources from interfering with the measurement, (this also helps prevent eye injury from

exposure to the laser) although some modern profilometers use algorithms to correct for the influence of other light sources. The high resolution of profilometers can result in very long times being required to measure macroscopic samples, so high framerate CCDs and FPGAs are sometimes used to speed up the process. Some profilometers also report the intensity of the reflected (as opposed to scattered) light, which allows the error in the measurement to be estimated quantitatively.

3.5.3 Thermocouples and Calorimeters

A thermocouple is a small probe that acts as an electric thermometer, reporting temperature against time at the location of the probe. The thermocouple was invented by Thomas Seebeck in 1821, and consists of two wires made from different metals joined in an electrical junction. A basic diagram of a thermocouple is shown in figure 16. When it is heated or cooled, the two metals react differently, and a small voltage is generated between them. This effect is called the thermoelectric effect, or sometimes the Seebeck effect. There is usually a second, reference thermocouple at room temperature so that the absolute temperature of the measurement probe can be determined using lookup tables. Different metal pairs can be used to provide measurement sensitivity in different temperature ranges.

Thermocouples provide a more accurate and precise temperature reading than either thermography or pyrometry, albeit with lower time resolution. However, thermocouples require good thermal conduction to ensure this accuracy, which relies on good physical contact. They can not be used at all with no physical contact, preventing their use for

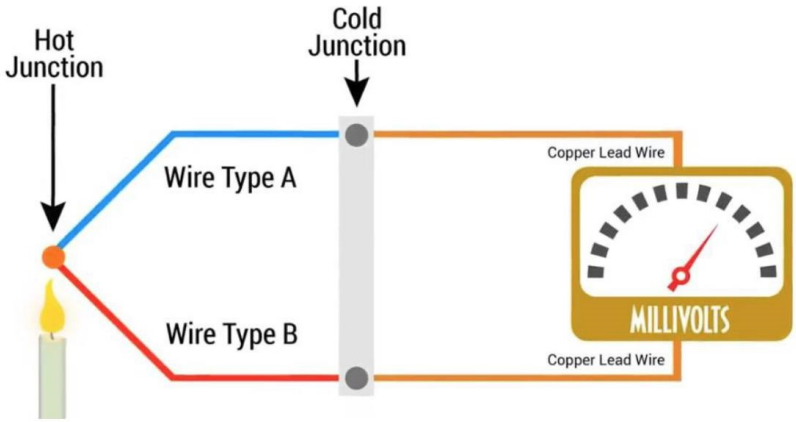

Figure 16: A simple diagram showing the hot and cold junctions of a thermocouple, and where the different types of wire are used. Image from omega.com

measurements of radiated heat. The time required for conduction of heat through solid materials is the reason for their lower time resolution. Thermocouples can still be used to measure the temperature of components of the vessel, such as for the tile temperature measurements in power balance on JET in [34]. They can also be used to measure solid experimental samples, such as for calibrating the pyrometer in [59], provided the time scale is long enough. They can also be used as safety cutoffs in tokamaks, in places where power is directed such as neutral beam dumps.

Calorimetry is used to measure the heat absorbed by a water cooled sample by measuring the temperature of the water both before and after passing through the sample. Directly cooled samples and samples in water cooled mounts can be measured. The temperature measurements are usually done with thermocouples, preferably as close to the sample as possible. If the temperature measurements are taken far from the sample, there may be other sources or sinks of heat along the water pipe between the thermocouples and sample. In a fusion environment, this may include hot copper electromagnets or liquid nitrogen or helium cooling systems for superconducting magnets.

These two temperature measurements must then be re-synchronised in time such that the temperature of the same water molecules is being measured (i.e. by offsetting the time taken for water to flow through the system). This is especially important where the thermocouples are far from the sample. This time offset means that the actual change in the water temperature can be determined even when the temperature of the input water is changing with time. Having calculated the temperature change of the water, this can then be multiplied by the the volumetric heat capacity of water and the volume flow rate, to give the rate of heat transfer in Watts.

3.5.4 Magnetic Coils and Flux Loops

On tokamaks and other fusion devices, measurements of magnetic fields, and some other parameters such as plasma current and loop voltage, are made using a variety of conducting coils and loops. A loop refers to a single closed loop of conductor, whereas a coil refers to a conductor wound multiple times. Currents are induced in these coils according to Lenz's law, which states "the current induced in a circuit due to a change in magnetic field is directed to oppose the change in flux and to exert a mechanical force which opposes the motion". This is mathematically encapsulated in Faraday's law of induction

$$\mathcal{E} = -\frac{\partial \Phi_B}{\partial t} \tag{10}$$

where $\mathcal{E}$ is the electromotive force (in Volts), Φ_B is the magnetic flux (in Weber) and t is time. These coils have different sizes, shapes and locations to detect different currents and magnetic fields, such as the plasma current. There are various names for these types of coils and loops as used in tokamaks, such as Rogowski coils for plasma current and flux loops for loop voltage, and there is no standard setup; each tokamak has a bespoke coil and loop arrangement (including MAST and MAST-U).

3.5.5 Data Derived from Other Systems

Power balance also requires knowledge of the neutral beam injector (NBI) input power, which is not reported by a true diagnostic, but rather calculated as an input parameter, based on the electrical power supplied to the NBIs. This is the product of the input beam voltage and the beam ion current. On MAST and MAST-U, the equilibrium fitting code (EFIT) is used to calculate certain parameters. EFIT was developed at the DIII-D tokamak in the 1980s [60].

This is similar on most tokamaks, while there are no neutral beams on Magnum-PSI. There is no diagnostic that can measure absorbed NBI power directly, which is why TRANSP simulations (see section 3.6) have been performed for this purpose. TRANSP takes the EFIT smoothed version of the NBI power (interpolated to the EFIT Summary MAST timebase) as an input, and calculates the fraction that actually contributed to plasma heating.

3.6 TRANSP Simulation Code

The TRANSP plasma transport simulation code was developed at Princeton Plasma Physics Laboratory (PPPL) in the early 1980s to analyse data from the Toroidal Fusion Test Reactor (TFTR), a large tokamak that was operational there at the time [61]. TRANSP has since been adapted for use on other tokamaks, such as JET, NSTX, DIII-D, and MAST. The source code is written in Fortran 77, but it is closed source, so the reputation of TRANSP rests on past experimental validation (see section 2.4).

The simulation of neutral beam injectors (NBIs) in TRANSP uses a module known as NUBEAM. NUBEAM was originally written as part of the core TRANSP code, but was

separated out and ported to Fortran 90 so that it could be used in conjunction with other codes [62]. NUBEAM can also calculate the effect of other fast particles not resulting from NBIs, such as those from thermonuclear fusion. The workings of NUBEAM are not as closely guarded as those of TRANSP itself; the algorithms used are explained in detail in [62].

NUBEAM can account for any number of beamlines, all beam geometries and positive and negative ions of any isotope. At its core, NUBEAM is a Monte Carlo code, meaning that it integrates particle positions and velocities in space and time using random sampling integration (also known as Monte Carlo integration). The nature of random sampling means it is in effect tracking the position and velocity of a number of randomly chosen "super-particles", each one representing many actual particles. The motion and interaction of these particles is assessed using lookup tables of known particle interaction cross-sections, and the results are extrapolated to the full number of particles being represented.

The exact process for using TRANSP has changed over time, and may be different for different tokamaks. However, most use of TRANSP requires a magnetic equilibrium to be given as an input. This is normally taken from the magnetic diagnostics of the tokamak (such as the coils and loops described in section 3.5.4) for real experimental shots, or simulation outputs for hypothetical shots (which may be on either existing or hypothetical tokamaks). On MAST and MAST-U, these magnetic equilibria come from EFIT. TRANSP also needs certain machine information, such as size, shape and beam geometry.

One recent significant piece of work on TRANSP, which was of great utility to this project, was the integration into a new user interface called "one modelling framework for integrated tasks" (OMFIT). This greatly simplified and standardised the way TRANSP is used, as described in [63]. This is now the normal way to use TRANSP with MAST or JET, and was the method used for the TRANSP simulations conducted for this project.

The output from TRANSP consists of a several log files, and an output variables file in netCDF format. Many of these variables are 1D, varying in time only, and some are 2D varying in time and a single spatial coordinate, given in terms of magnetic flux surfaces. This spatial coordinate assumes toroidal symmetry, as does TRANSP itself, and the use of flux surface as a coordinate allows tokamaks with different cross sectional shapes to be treated by the same code. A variety of methods exist for using TRANSP output as input to other simulation codes and vice versa, and this can be done within OMFIT.

3.7 Contribution of Tools and Techniques to the Project

Many of the diagnostic tools described in this chapter contributed to the power balance analysis on MAST and MAST-U. Infrared cameras and bolometers (and sometimes Langmuir probes) contributed directly, while coils and loops, the NBI system and TRANSP were used to indirectly calculate other quantities. The others, namely Thomson scattering, pyrometry, profilometry and thermocouples contributed more to the experiments at Magnum-PSI, and the attempt to better understand material surface effects.

The next three chapters describe the methods by which these diagnostics were employed, and the results of the experimental analysis. Each describes different experiments and analysis which took place on different machines: MAST, Magnum-PSI and MAST Upgrade respectively. Similar diagnostics were used for many measurements, and all three contribute to the common objectives set out at the end of chapter 2.

Chapter 4

Improvements to Power Balance Analysis

The first stage of this project was to establish a method of power balance analysis using existing tokamak data, that could quickly be adapted and applied to experiments with advanced divertor configurations, when they became available. The archive of the original MAST has plenty of data that could be used for this purpose. There were also two factors already mentioned that were overlooked in previous studies, wetted area correction and neutral beam power absorption. This retrospective analysis of MAST data was the perfect opportunity to implement them, and to perform a comparison of empirical and simulation data regarding the beam power absorption.

4.1 Design and Capabilities of MAST

4.1.1 Machine Specification and Parameters

The Mega Ampere Spherical Tokamak (MAST) was one of the largest spherical tokamaks ever built when it came online in 1999, dwarfing its predecessor the Small Tight Aspect Ratio Tokamak (START). The spherical tokamak concept is distinguished from other tokamaks by the smaller aspect ratio of its toroidal shape. Any tokamak has the shape of a torus, as mentioned in section 1.1.2, and therefore has two radii: the major radius of the overall torus R, and the minor radius of the cross section a. The ratio of R/a in a tokamak is known as the aspect ratio, and is typically around 2-3. In a spherical tokamak however, the aspect ratio

is around 1-1.5. A smaller aspect ratio means that the ratio of plasma pressure to magnetic field pressure, known as the plasma β is higher in a spherical tokamak [64]. This allows higher plasma β for given plasma pressure, enabling the same fusion triple product to be achieved which much lower magnetic field than in a conventional tokamak. The spherical tokamak concept pioneered by START also inspired the National Spherical Tokamak eXperiment (NSTX) in America [65].

As its name suggests, MAST was designed and built to be capable of generating plasmas with a mega-amp of plasma current, and to perform experiments in the same plasma regimes as the more established conventional aspect ratio tokamaks of the time. It succeeded in this endeavour, and MAST experiments ran from 1999 to 2013, although this analysis focuses on experiments since 2008. The success of both major spherical tokamaks (MAST and NSTX) led them both to have major upgrades in the last 10 years; NSTX-Upgrade, and MAST Upgrade (which is the subject of chapter 6). MAST had impressive specifications that made it a very capable tokamak. For plasma heating, it was equipped with ohmic heating, electron-Bernstein wave (EBW) heating (a type of microwave heating), and neutral beam heating, with two NBIs. Some of its key capabilities, which may be of interest to the expert reader, are given in table 2.

MAST was contained in a cylindrical vacuum vessel about 4.4 m high and 4 m wide, which can be seen in figure 17. Most conventional tokamaks use larger aspect ratio, close-fitting toroidal vacuum chambers with D-shaped cross sections, such as that shown in figure 1. MAST's open vacuum vessel with rectangular cross-section allowed for more flexibility in plasma shaping and positioning. There were 12 pairs of toroidal field coils wrapped around the outside of the vacuum vessel. There were also 8 poloidal field coils, which unusually were

Table 2: The design maximum values of certain key machine parameters of the Mega Ampere Spherical Tokamak (MAST), shown to give the expert reader an idea of the capabilities of MAST relative to other tokamaks. The plasma parameter space covered in these experiments is given in table 3.

Parameter	**Value**
Plasma current	1.3 MA
Magnetic field	0.55 T
Core plasma temperature	2 keV
Pulse length	700 ms
Plasma volume	8 m^3
NBI power	5 MW

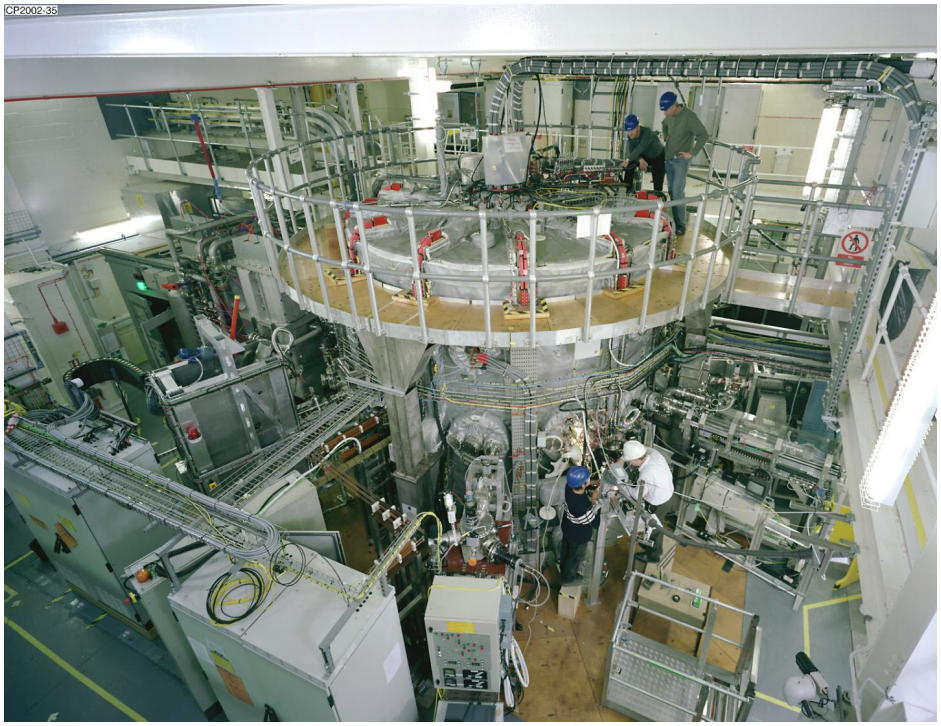

Figure 17: A photograph of MAST taken in 2002. The top corners of some of the toroidal field coils are visible (with cables with red insulation laid over them). One of the two NBIs is visible in the left of the image. Figure credit CCFE.

inside of both the toroidal field coils and the vacuum vessel. Having the poloidal field coils inside the vacuum vessel is particularly unusual, and required containing each one inside its own re-entrant (non-evacuated) vessel to keep them at atmospheric pressure, in order for the water cooling systems to work correctly. The vessel also ensures that only steel components, not copper, are exposed to the plasma.

The experimental capabilities of MAST were demonstrated extensively during its many years of service, including 9 experimental campaigns. Most of these campaigns were broad in scope, including a wide range of shots with different experimental objectives, and were not focused specifically on the unique properties of spherical tokamaks. Indeed, many were designed to provide insight into the design of ITER, which has a conventional aspect ratio, demonstrating that MAST was a capable tokamak in its own right. The number and type of diagnostics on MAST increased over the years, so naturally for this retrospective analysis shots from later campaigns were chosen, to take full advantage of all the diagnostics that were available by that time. The diagnostics used in this analysis are described in section 4.2.

4.1.2 Past Experimental Campaigns on MAST

Experiments used in power balance analysis do not need to have been set up with power balance in mind, as the diagnostics needed are kept running for almost all shots. In fact, performing the analysis on different shots can be helpful in determining how various different parameters affect the power balance. The shots used for the power balance analysis in this chapter were chosen from shots with a wide range of experimental objectives, some of which will be summarised here.

The majority of shots used in the following analysis took place in three short periods during MAST's 14 year lifetime. These periods were June-July 2011, early in campaign M8, when the focus was on different plasma modes and instabilities; December-January 2011/12 later in M8, when the focus was on H-mode studies and detachment; and June 2013 in M9, when the focus was on edge physics and scrape-off layers. The process by which these shots were selected is outlined in section 4.3.1, and was not related to the experimental priorities mentioned in this paragraph. The full range of certain MAST plasma parameters covered by this selection of experiments is given in table 3, showing that a broad parameter space is represented in the subset of shots used in this analysis.

MAST was used successfully in many different plasma physics experiments during its lifetime. A review was published in 2015 (after MAST had stopped operating) that summarised many of the key results from the machine [66]. It states that optimising H-mode confinement and operating with smaller ELMs were two of the primary successes of MAST. It also disuses how the high speed camera gave insight into the role that plasma filaments play in the private flux region. The conclusion to [66] states that new insight into ELM

Table 3: The range of values of certain plasma parameters covered by the 185 experiments used in this analysis. The notes draw attention to skew or outliers.

Parameter	**Minimum**	**Maximum**	**Notes**
Density n_e	$4 \times 10^{19}\,\mathrm{m}^{-2}$	$25 \times 10^{19}\,\mathrm{m}^{-2}$	106 between 10 and $15 \times 10^{19}\,\mathrm{m}^{-2}$
Current I_p	200 kA	900 kA	only 33 below 600 kA
NBI Power P_{NBI}	1 MW	3.5 MW	(excludes 32 ohmic shots with zero) 86 with 1 beam, 67 with 2 beams
δr_{sep}	−5 cm	2 cm	145 DND and 37 lower SND with 24 centre-column-limited
Neutron Flux Γ_n	$5 \times 10^{12}\,\mathrm{s}^{-1}$	$17.5 \times 10^{13}\,\mathrm{s}^{-1}$	(excludes 22 ohmic shots with near zero)

mitigation was one significant area of progress, as was the ability to move the plasma vertically (enabled by its open configuration) to assess the benefits of off-axis beam heating, and the revelation that intermittent divertor heat and particle fluxes are caused by plasma filaments [66].

4.2 MAST Diagnostics

The data presented in this chapter is from the MAST database, which is extensive, and comes from the diagnostics that were available at the time of each experiment. These include infrared cameras (covering all four divertor strike points), bolometers to measure the heat radiated out of the bulk plasma, flux loops and Mirnov coils to measure the plasma current and loop voltage, and the input power to the Neutral Beam Injection (NBI) system. The principles of these diagnostics have been explained in chapter 3, and some details of the specific setup on MAST are given below. The locations of the thermal diagnostics are shown in figure 18.

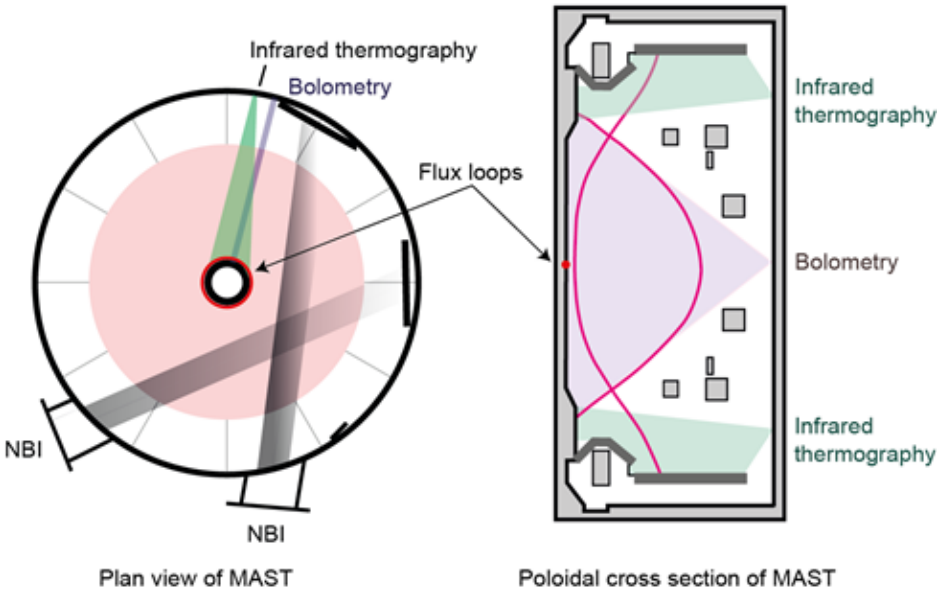

Figure 18: The viewing angles and sight lines of the infrared cameras and bolometers on the MAST tokamak, as well as the locations and angles of the NBIs, and the location of the flux loops. Figure supplied by A Thornton.

4.2.1 Infrared Cameras on MAST

MAST had two infrared cameras, a medium wave infrared (MWIR) and a long wave infrared (LWIR). The LWIR camera was added after it was discovered that LWIR is less sensitive to surface effects such as hot spots [32].

The MWIR camera was a Santa Barbara Focalplane SBF125, with a 320×256 pixel resolution, filtered to wavelengths of around 5 µm. It was fitted with a 13 mm wide angle lens and typically viewed the lower divertor. Its frame rate was 300 Hz at full frame, although it usually ran faster than this with a smaller frame size. The camera sensor was internally cooled to around 70 K using a Stirling engine, as cooling is required for this type of camera to work.

The LWIR camera was a Thermo Sensorik CMT 256 L HS camera, with a 256×256 pixel resolution, filtered to wavelengths of $7.6 - 8.9$ µm. It was fitted with a 15 mm wide angle lens. Its frame rate was 900 Hz, and it could run at up to 8 kHz with a sufficiently small frame size. This camera's sensor was also cooled to 70 K using a Stirling engine. This camera typically viewed the upper divertor. CAD representations of the views of both the cameras are shown in figure 19.

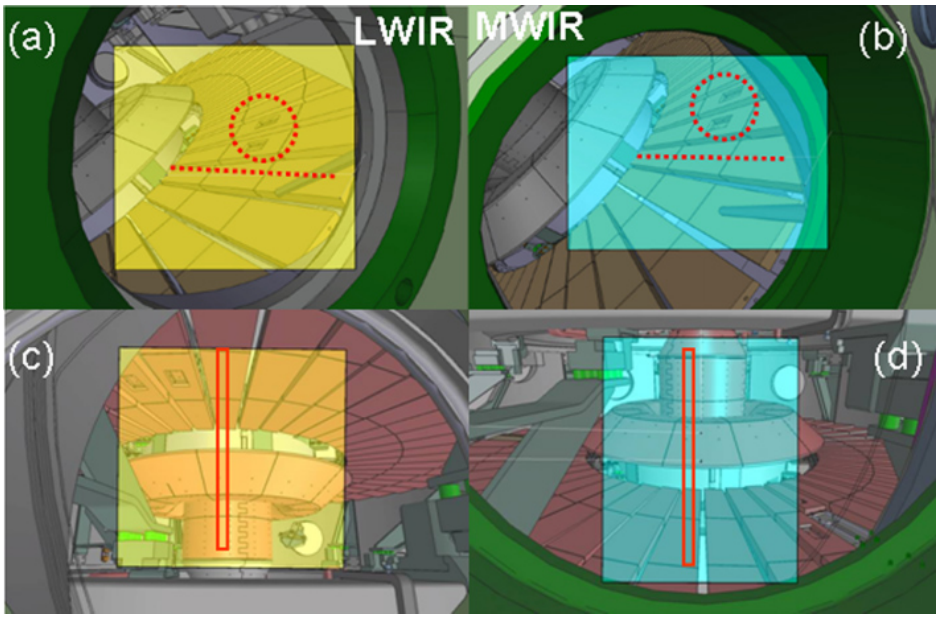

Figure 19: Areas of the MAST tiles viewed by the two infrared cameras in both of their positions. (a) and (c) are the long wave camera, and (b) and (d) are the medium wave camera. The red line or strip indicates the integrated are used in the analysis. The red circle shows the location of heated tiles used for in-vessel calibration. Figure from [30].

There have been times in the past when the cameras viewed the same divertor simultaneously as in figure 19 (a) and (b), such as the experiments that lead to figure 7, which are described in [30]. However, in the post-2008 experiments used in this analysis, the two cameras were used together in order to observe both upper and lower divertors simultaneously, which was especially useful during double null discharges (which are a significant majority of discharges on MAST). Typically, the LWIR camera looked at the upper divertor (figure 19 (c)) and the MWIR camera looked at the lower divertor (figure 19 (d)).

The data from these cameras was written into the MAST data archive. The conversion from temperature to heat flux was done using a code called THEODOR (thermal energy onto divertor) as mentioned in section 3.1 and described in [52]. This code ran automatically on the MAST computer servers between shots, like all standard MAST diagnostic analysis codes.

4.2.2 Bolometers on MAST

The bolometers on MAST used gold foil sensors which are sensitive to total radiation and neutral particle flux. The functionality of bolometers was described in section 3.2. There were both horizontal and vertical arrays of bolometers on MAST, which are illustrated in the left and right of figure 18 respectively. Each of these consisted of four 4-channel bolometers, and combining the data from all of them allowed the heat radiated out of the entire plasma to be measured.

The bolometer coverage was focused on the core plasma rather than the divertor region. This is because in attached divertor conditions, which were the overwhelming majority of discharges when MAST was constructed, the fraction of radiated power that was radiated in the divertor was very small. It was not considered necessary to have bolometers in this region at the time (this had changed by the time MAST-U was constructed, see section 6.2.2).

4.2.3 Non-Thermal Diagnostics on MAST

The Thomson scattering system on MAST was updated in 2009 [67] and again in 2010 [68]. The system measured at 130 spatial points, with a 10 mm resolution. The time resolution was 240 Hz, which was achieved by combining eight 30 Hz lasers. The signals are digitised by 8 bit ADCs at a rate of 1 gigasample per second.

The data was processed by 26 dual core computers each operating independently, connected to 5 polychromators (corresponding to the 130 spatial locations). This computing power means the data is available soon after the shot. The system was synchronised to the MAST clock, allowing measurements to be taken at a higher time resolution around the anticipated time of events such as ELMs [68].

MAST had several arrays of coils and loops covering every part of the machine, as they are essential for reconstructing the magnetic topology of the plasma, which is required for calculating many parameters. This includes calculation of several quantities used in power balance calculation, such as ohmic heating power, work done on the plasma, and magnetic energy of the plasma. These parameters are not measured directly by the coils, but reconstructed via EFIT (equilibrium fitting) modelling, which is the standard inter-shot modelling used on MAST and MAST-U. As well as the basic coils used primarily for EFIT, there were several other coils used to measure other parameters, such as real time plasma current and magnetic fluctuations.

Although it is not strictly a diagnostic, the neutral beam power is an important quantity for power balance in beam heated shots, which represent the majority of shots analysed here.

4.2.4 TRANSP Simulations of MAST

TRANSP simulations on MAST (and on other tokamaks) take an input 'namelist' of various MAST machine parameters, including the vessel size and neutral beam locations. It also requires an input equilibrium, and in the case of MAST, this is usually generated by EFIT [69]. The input files are then combined into a tarball, which is sent to the McCune server at PPPL for execution. The TRANSP GUI called OMFIT, which was developed at CCFE by Dr Stuart Henderson, allows TRANSP simulations to be set up and run without the user having to manually define the machine parameters, and was used for all TRANSP runs presented here.

For the results presented here, the scheduler equilibrium was used, rather than generating a bespoke EFIT equilibrium. This was done to maintain consistency with the other EFIT data traces used elsewhere in the analysis. The scheduler equilibrium is calculated only using the magnetics data, which can be less accurate, especially at the edge, than methods using $\mathbf{D}_\alpha$ spectroscopy. Using a different equilibrium such as this for TRANSP may require the values of certain other data traces to be recalculated however, and an assessment of the

method of calculating the magnetic equilibrium is beyond the scope of this thesis.

4.3 Method for Energy Balance

4.3.1 Building a Shot Database

MAST shots since approximately shot 20000 are held in the most recent iteration of the database, which has all the necessary information required for this analysis. These shots were checked for the nine data traces required, which are: neutral beam power, ohmic power, radiated power, stored plasma energy, magnetic plasma energy, and divertor flux to the four strike points (upper outer, upper inner, lower outer, lower inner). This produced a list of 778 shots.

It was then decided that of these shots should be integrated over a 50 ms time window, in order for time averaging to remove random noise from the data, while being a short enough time period that is would fall within the flat top of most shots. It was also necessary to choose a time window during which there were no transients such as ELMs, or sawteeth, or L-H transitions. This allows comparison to be drawn between plasmas in steady-state conditions, which are the most stable and consistent conditions, and therefore desirable for this large scale analysis of many shots that took place over several years. The data traces best suited to identifying these are D_α, soft x-ray, and n_e, so the shot list was narrowed to shots for which these traces were also recorded, which resulted in a list of 655 shots.

There were many shots where a data file was recorded for the relevant diagnostics, but it contained only bad data. This is difficult to detect automatically, so all nine data traces were plotted individually to see if there were any notable outliers, and where there were, searching for a reason that might explain it. Where all outliers had the exact same value, this was assumed to be a diagnostic malfunction. Excluding shots where the neutral beam power was not recorded correctly reduced the list to 445 shots. Excluding shots where the value for ohmic heating power was negative for at least part of the shot reduced the list to 226 shots. Six further shots were excluded because of abnormally low plasma density (i.e. the shots never got going), five more because of abnormally high readings for radiated power, and three more because of abnormally high lower outer divertor flux, which left a list of 212 shots. A further six were excluded because of an apparent fault in the radiated power diagnostic (bolometers) during those shots, leaving a list of 206.

Having decided on a preliminary shot list, the 50 ms time windows were selected manually for each of the 206 shots, in every case aiming for a period of relatively smooth I_p, D_α soft x-ray, and n_e. This was done manually for every shot, and an example is given in figure 20. The result is a consistent length of time in every shot in the database, so that they can be compared directly.

During the selection of the time windows, it was found that two of these shots were failed shots that did not last long enough for any 50 ms time window to be meaningful, so they were then excluded from the list. Another nineteen shots were later found to have non-zero currents in the ELM coils during their time windows, and these significantly affect the power balance, meaning these shots are not representative of others on MAST. These were excluded as well, leaving a list of 185 shots used in the energy balance analysis.

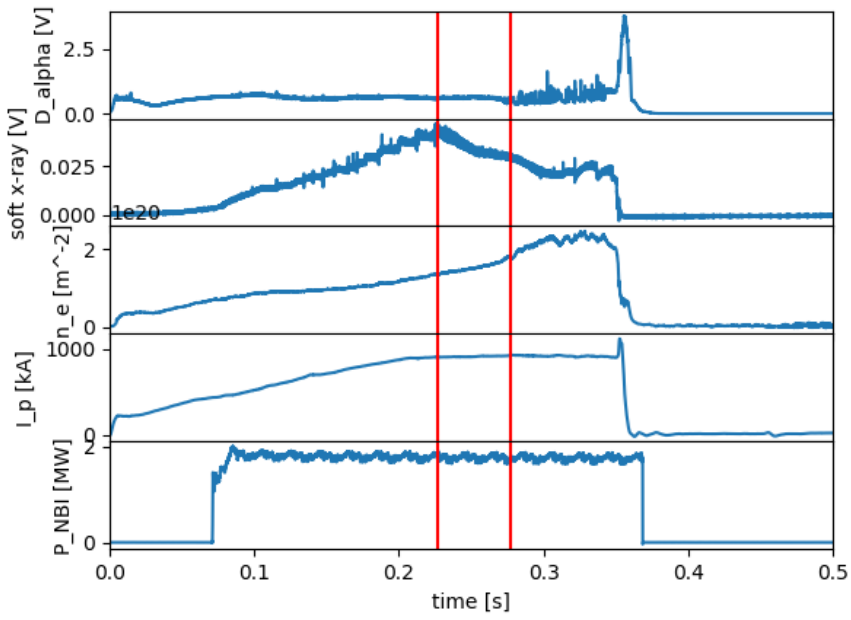

Figure 20: The chosen time window of shot 26392, shown here as an example, taken when neutral beams were on, after the current had reached flat-top but before the L-H transition. There were no ELMs or sawtooth instabilities during the time window. From the top, the data traces are measured intensity of the deuterium-alpha spectral line, soft x-rays, plasma density, plasma current, neutral beam power.

4.3.2 Process of Energy Balance Analysis

Some of the data traces need to be converted from power to energy for the final energy balance (i.e. time integrated power balance) calculation. This requires time integration of the traces that were given in units of power rather than energy, namely the magnetic stored power, radiated power, ohmic heating power, and neutral beam power. The data traces used for stored mechanical and magnetic energy are calculated from the plasma current and loop voltage. The mechanical work done is simply the product of the two, and the rate of change of magnetic stored energy, X, is calculated from equation 11.

$$X = \mu_0^{-1} \int_U \left[B_p \dot{B}_p + B_\phi^{\text{plasma}} \left(\dot{B}_\phi^{\text{vacuum}} + \dot{B}_\phi^{\text{plasma}} \right) \right] dV \tag{11}$$

where μ_0 is the permeability of free space, U denotes the physical boundary of the plasma, B_p is the poloidal magnetic field and $\dot{B}_p$ is its rate of change, B_ϕ is the toroidal magnetic field and $\dot{B}_\phi$ is its rate of change, the superscripts 'plasma' and 'vacuum' refer to the magnetic field created by the plasma itself and that created by the magnetic coils in the absence of plasma, and dV is a volume element. Note that X has dimensions of power. The quantities including X are then added according to equation 12

$$E_{\text{NBI}}^{\text{abs}} + E_{\text{ohmic}} = W + \int X \, \mathrm{d}t + (E_{\text{l,o}} + E_{\text{l,i}} + E_{\text{u,o}} + E_{\text{u,i}}) + E_{\text{rad}} \tag{12}$$

where $E_{\text{NBI}}^{\text{abs}}$ is the absorbed neutral beam energy, E_{ohmic} is the ohmic heating power, W is the work done on the plasma, X is the magnetic stored power and $E_{x,y}$ is the energy deposited on each of the the strike points (lower outer, lower inner, upper outer, upper inner). Note that this is a modified and time-integrated version of equation 3.

One thing that has been overlooked in past studies of tokamak power balance such as [30], is that plasma in not incident on the entire surface area of the outer tiles, because of the relative location and inclination of the tiles, and the angle of plasma incidence. The outer divertor tiles in MAST are not horizontal in the toroidal direction. They are inclined a few degrees into the face of oncoming plasma. This is done to prevent plasma striking the edge of the tile, which is exposed by the gap between the tiles. The inclination was 4° on MAST, which created a gap large enough for Langmuir probes and other diagnostics to be placed there while still remaining shadowed [70]. The shadowing effect (which is visible in figure 34) typically resulted in around 20-30% of the tile being shadowed. Note that tiles are not

inclined in this way at the inner strike point.

In the calculation of the total heat flux onto the tiles, the heat flux onto the small fraction of the tiles where this is measured is integrated across the entire tile area to calculate the total heat flux. However, simply using $2\pi R$ for this calculation over-estimates the area onto which plasma is incident, because a strip at the edge of each tile is shaded by the adjacent tile. This is analogous to the sheltered strip along the top of an ordinary inclined roof tile, which is sheltered from the rain by the tile above. This is illustrated by figure 21.

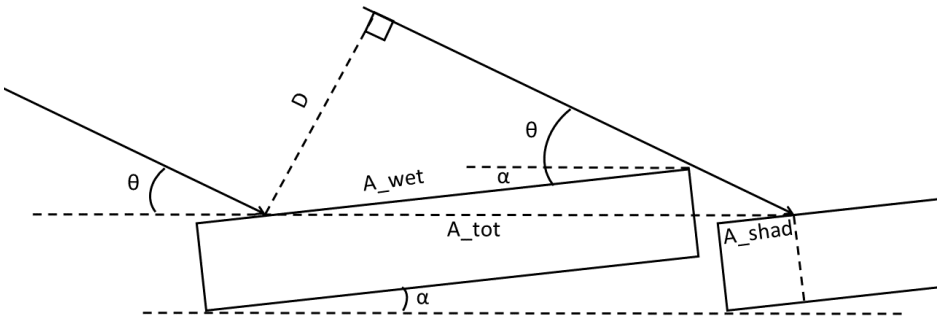

Figure 21: A diagram showing the origin of the wetted fraction of the tile, the corresponding shaded fraction, and the angles α and θ.

A calculation is made to compensate for this using the ratio between the horizontal vessel floor area given by $2\pi R$, denoted, A_{tot}, and the wetted area of the inclined tile, A_{wet}. Note that $A_{\text{tot}} \neq A_{\text{wet}} + A_{\text{shad}}$. These can be evaluated using trigonometry, with the lengths and angles shown in figure 21, and are given by

$$A_{\text{tot}} = \frac{D}{\sin\theta} \tag{13}$$

$$A_{\text{wet}} = \frac{D}{\sin(\theta+\alpha)} \tag{14}$$

The wetted fraction, f_{wet}, is the ratio of the two

$$f_{\text{wet}} = \frac{A_{\text{wet}}}{A_{\text{tot}}} = \frac{\sin\theta}{\sin(\theta+\alpha)} \tag{15}$$

where θ is the angle of incidence of the plasma relative to the plane, and α is the angle

of inclination of the tiles relative to the same plane. Assuming this is the horizontal plane, which it was for the outer strike point on MAST,

$$\sin\theta = \frac{-B_z}{\sqrt{B_R^2 + B_z^2 + B_\phi^2}} = \frac{-B_z}{|\mathbf{B}|} \tag{16}$$

where $\mathbf{B}$ is the magnetic field, and the subscripts R, z and ϕ denote the radial, vertical and toroidal components respectively. Equation 16 is from [71], and can be derived from simple trigonometry. We now require a value for α, which in MAST was 4° as mentioned above.

This calculation was applied in the analysis of both MAST outer strike points (upper and lower). The wetted fraction was calculated at a single point (the point with highest heat flux at that time) in the camera analysis frame (a strip along the radial dimension of the divertor tiles), every 10 ms to allow for the changing magnetic field over the 50 ms duration of the analysis window. Using the wetted fraction f_{wet}, equation 12 then becomes

$$E_{\text{NBI}}^{\text{abs}} + E_{\text{ohmic}} = W + \int X \, \mathrm{d}t + f_{\text{wet}}(E_{\text{l,a}} + E_{\text{l,i}} + E_{\text{u,o}} + E_{\text{u,i}}) + E_{\text{rad}} \tag{17}$$

However, this version still assumes that the quoted NBI energy is all absorbed by the plasma.

4.4 Results of Energy Balance

4.4.1 Initial Data

Using the 185 suitable shots identified for the analysis, the power balance calculation outlined in section 4.3.2 was performed, and each of the shots plotted in figure 22, to show the percentage of power recorded going into the plasma which was recorded going out. This data does not account for the wetted fraction of the tile, or neutral beam power absorption. The mean (with standard deviation), median, quartiles and mode of the distribution are also shown, and are quoted numerically in table 4.

By comparing the mean, median and mode, and inspecting figure 22, we can see that the distribution shows no particular skew. However there are many shots for which more than 100% of the energy recorded going into the energy put into the plasma was deposited back out again (more than 100% is impossible, so this must be caused by measurement error).

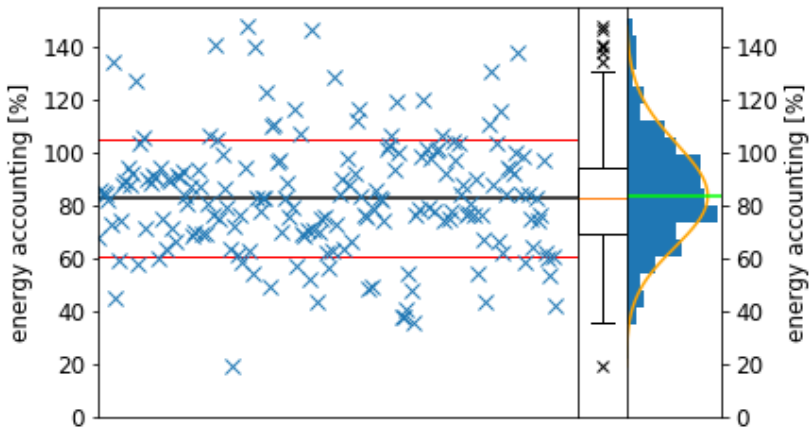

Figure 22: Energy accounting (energy out divided by energy in, as a percentage) for 185 MAST shots. The black line is the mean, the red lines are +/- one standard deviation, and the median and quartiles are shown in the box plot. The blue histogram on the right shows the distribution of the data, with a skew-normal distribution fitted (orange curve), from which the maximum value is taken as the mode (green line). Numerical values for all these are given in table 4.

Table 4: Statistical values for the distribution of energy balance data (i.e. energy out divided by energy in) for the data in in figure 22.

Quantity	**Value** [%]
mean	83
standard deviation	22
median	83
upper quartile	94
lower quartile	69
mode	83

However, this was smaller than the error for shots where less energy was observed leaving the plasma than entering it, so no particular action was taken.

By comparison to table 1, it can be seen that the mean of 83% accounting is typical of past experimental tokamak power balance studies, including those on MAST itself [30]. The accounting was expected to initially become worse once wetted area was accounted for, then become better once beam power absorption was also accounted for.

4.4.2 Accounting for Wetted Area

The data was re-processed, with the outer infrared data traces adjusted to account for the wetted area of the outer divertor tiles. The origin and methodology of the wetted area calculation was given in section 4.3.2. This data is shown in figure 23. This still does not account for the absorbed fraction of neutral beam power. Table 5 shows the statistics of the distribution. It can be seen from the lower mean, further from the perfect 100%, that this actually worsens the agreement between theory and experiment. However, it is expected that accounting for the neutral beam absorption fraction would bring the experimental observation back into better agreement with theory.

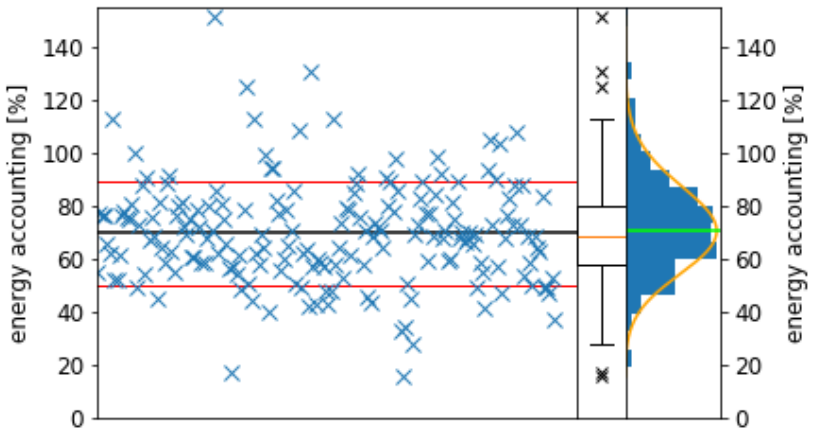

Figure 23: Energy accounting (energy out divided by energy in, as a percentage) for 185 MAST shots, with wetted area correction applied to the outer strike point data. The black line is the mean, the red lines are +/- one standard deviation, and the median and quartiles are shown in the box plot. The blue histogram on the right shows the distribution of the data, with a skew-normal distribution fitted (orange curve), from which the maximum value is taken as the mode (green line). Numerical values for all these are given in table 5.

Table 5: Statistical values for the distribution of energy balance data, corrected for wetted area. Calculated by dividing energy out by energy in, for the data in figure 23.

Quantity	**Value** [%]
mean	70
standard deviation	19
median	69
upper quartile	80
lower quartile	57
mode	70

4.4.3 Power Balance Discrepancy, Neutral Beam Power and Other Variables

The energy balance was plotted again, now with energy out against energy in, in figure 24 (a). This data was plotted several more times, with each data point grouped by colour according to certain parameters, to see if outlying energy balance values were caused by certain types of plasma pulse. This is shown in the remainder of figure figure 24, and the shots were coloured according to (b) plasma current, (c) plasma density, (d) neutron flux, (e) magnetic configuration and (f) neutral beams used.

It was expected that plasma density may have a significant effect on the power balance, because the efficiency of beam power coupling into the plasma is related to the plasma density. However, figure 24 (c) demonstrates that there is no clear dependence. Figure 24 (f), on the other hand does reveal a dependence on neutral beam power. From the gradients of the fitted lines in the figure, there is 85% accounting with no beams, 72% accounting with 1 beam and 59% accounting with 2 beams. This result lends credence to the idea that a significant portion of the injected neutral beam power is not coupled into the the plasma, as suspected. It also shows that the effect is greater when both beams are in use.

A closer look was then taken at how energy balance and neutral beam power relate to other variables, to see if any correlations could be observed. This is shown in figure 25. Figure 25 (a) is shown to clarify the threshold values of beam power associated with each number of beams. Figure 25 (b) shows a slight tendency for low current shots to be ohmically heated. This is most likely an operational feature, rather than an interesting physics result. From figure 25 (c) there is a tendency for beam heated shots to have higher density (especially double beam heated shots). Figure 25 (d) shows that this trend is similar when the Greenwald density fraction (electron density divided by the product of plasma

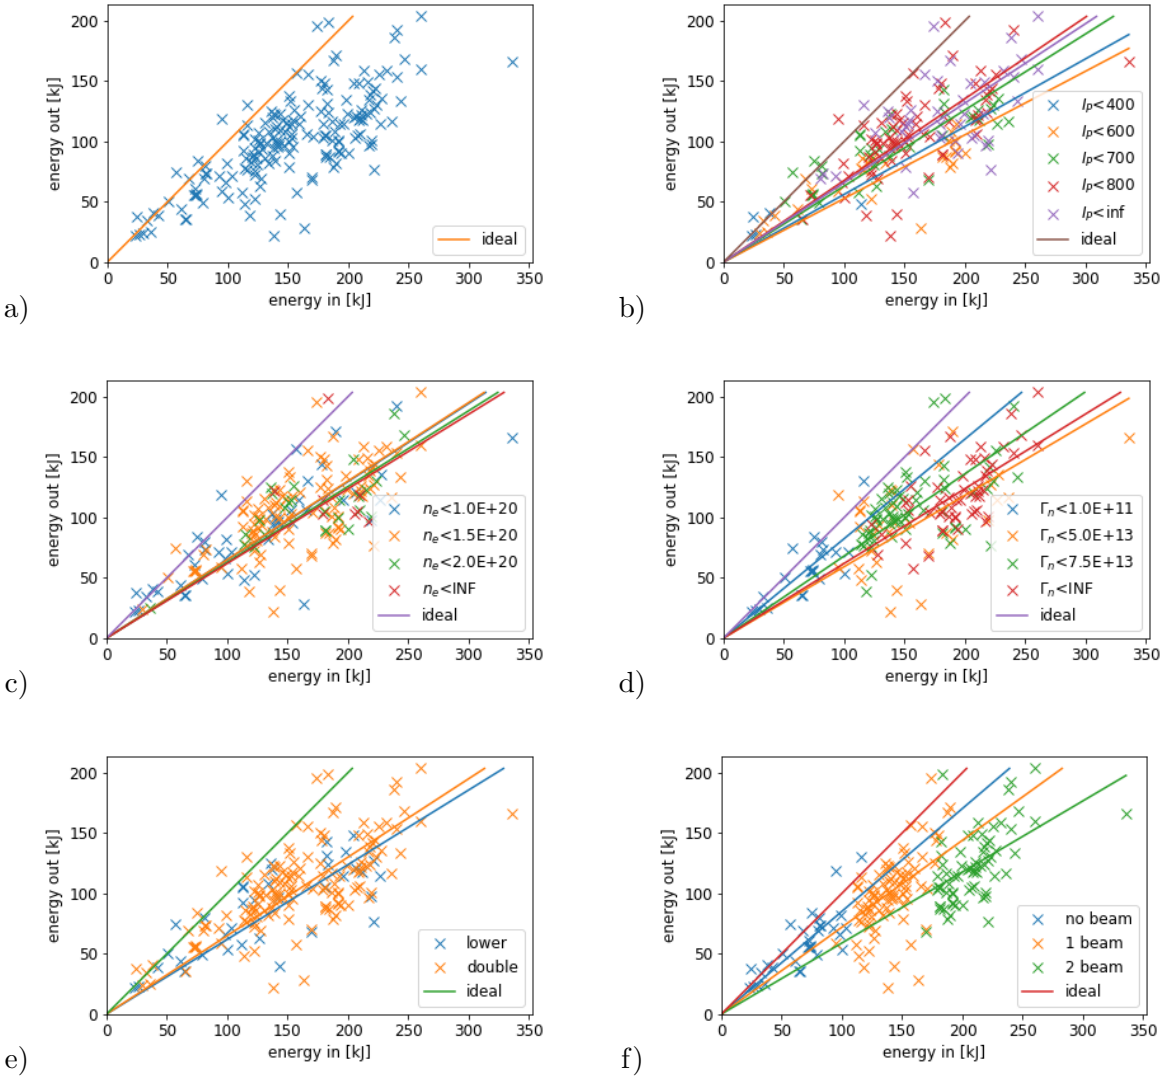

Figure 24: Energy out against energy in, for 185 MAST shots, with the ideal 1:1 line, (a) all a single colour, then with data broken down by various parameters. Note that for (b), (c) and (d) only the upper bound of each colour is shown to improve legibility; the lower bound of each colour is the upper bound of the previous colour. The parameters are (b) plasma current [kA], (c) line-integrated plasma electron density [m^{-2}], (d) neutron flux [s^{-1}], (e) divertor null configuration, and (f) number of neutral beams. Fit lines are constrained to pass through the origin.

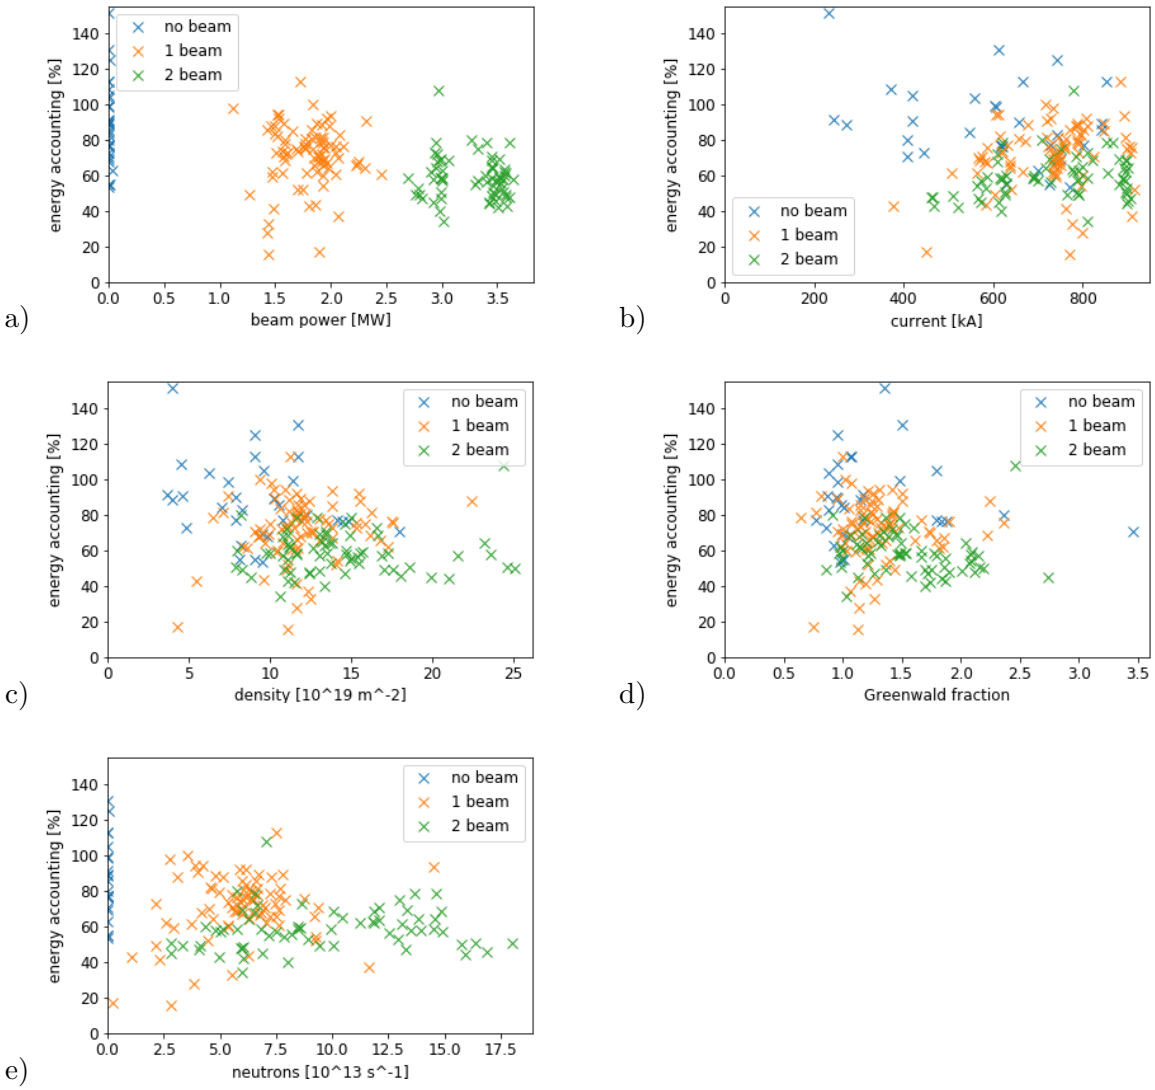

Figure 25: Energy accounting (energy out divided by energy in, as a percentage), coloured by number of neutral beams, plotted against (a) beam power (giving an indication of the power thresholds relating to number of beams, and making clear the correlation between energy accounting and beam power), (b) plasma current, (c) electron density, (d) Greenwald density fraction [72] and (e) neutron flux.

current and cross-sectional area, as defined in [72]) is considered instead. The correlation between overall energy accounting and density is shown to be weak in figure 24 (c). Figure 25 (e) shows a correlation between neutral beam power and neutron flux, but this is not surprising, since it is interaction involving the fast particles from the beams that produce the fast neutrons. This demonstrates that neutral beam power is the only significant determining factor in energy balance discrepancy in experiments on MAST, probably because not all of the reported beam power is actually absorbed into the plasma.

4.5 Method for TRANSP Simulations

The proportion of neutral beam energy absorbed by the plasma cannot be measured directly, so must be calculated from the output of simulations instead. The plasma transport code known as TRANSP, which was mentioned in section 3.6, was used for this purpose. The OMFIT user interface was used, which gathers all the required information from the MAST database for a given shot number, including the reconstructed magnetic field equilibrium for each shot. Using the magnetic equilibrium from the database maintains consistency with the other quantities used in the analysis that also rely on the magnetic equilibrium (such as the stored energy).

The time taken for TRANSP to run is too great for it to be used on all 185 shots in the database, so instead a small subset of shots was used. Shots were chosen which had all the necessary data for TRANSP (as it requires some data that had not already been screened for in the process of building the shot database), with an equal number of shots having 1 and 2 neutral beams active during the analysis window.

Once a relevant shot had been chosen, the simulation was set up. The start and finish time for the simulation must also be specified, to save the computing resources that would be required running the simulation for the entire shot. The chosen start time must be before the neutral beams were switched on, because the beam absorption takes time to build up after the beams come on (this can be seen in the green lines in figure 26). This is caused by the slowing down time of the fast beam particles in the plasma i.e. the time taken for the plasma to come into equilibrium with the beams. This results in higher than expected first orbit loss for those shots where the analysis window is (or where TRANSP assumes that the window is) soon after the beams were switched on. First orbit loss refers to neutral beam

power lost by being deposited in the scrape-off layer rather than in the core, as mentioned in section 2.1.1. The simulation is run on servers at PPPL, but the data is archived locally on the MAST data servers.

This data was used to calculate average neutral beam absorption fractions, for 1- and 2-beam shots (clearly no beam power is absorbed when neither beam is in use). This allows the overall power balance analysis to account for this lower, more accurate value of neutral beam heating power. Equation 17 then becomes

$$f_{\text{abs}} * E_{\text{NBI}} + E_{\text{ohmic}} = W + \int X \, \mathrm{d}t + f_{\text{wet}}(E_{\text{l,o}} + E_{\text{l,i}} + E_{\text{u,o}} + E_{\text{u,i}}) + E_{\text{rad}} \qquad (18)$$

where f_{abs} is the absorbed fraction. This is the full form of the power balance equation that accounts for all the effects considered in this thesis.

4.6 Results from TRANSP Simulations

4.6.1 Statement of Results

The first simulation was run on a shot with no beams (28020), to ensure that the simulation returns zero absorbed beam power when none is injected, and this was found to be the case. Further simulations were run, on both 1- and 2-beam shots. This includes specific shots of interest such as 29994, which had abnormally high neutron flux and low power accounting for a single beam shot. The TRANSP output data shown in figure 26 (a) shows that it had a very low ratio of shinethrough to first orbit loss, of around 1:7 (a typical ratio is 1:3). Another shot of interest was 28116, which had abnormally low neutron flux for a two beam shot, and also had poor power accounting, even for a 2-beam shot. Figure 26 (b) shows that it had a very high ratio of shinethrough to first orbit loss, of around 2:3.

The results from all the TRANSP simulations that were run are shown in table 6. These results explain some of the discrepancy noted above. The reason that it appears that shots with neutral beams have much less energy going out than coming in, is because the amount of absorbed beam power was over-estimated, leading to an over-estimation of the total power into the plasma. This resulted from the assumption that all neutral beam power was coupled into the plasma, which is shown here to be false.

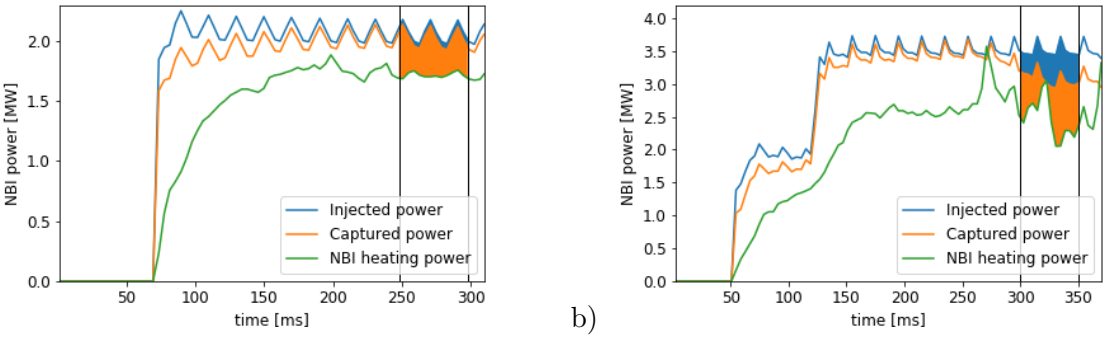

Figure 26: The neutral beam power which was injected by the beam, that which was coupled into the plasma, and that which thermalised the confined plasma (inside the separatrix) for MAST shots (a) 29994 and (b) 28116. The blue and orange shaded areas are the shinethrough and first orbit loss respectively, and vertical lines indicate the time interval for which energy balance had been calculated.

Table 6: Neutral beam energy not lost to shinethrough or first orbit loss for several MAST shots, therefore actually heating the plasma, according to TRANSP simulations.

Shot number	Number of beams	Beam power absorbed [%]
26394	1	93
27775	1	91
27871	1	78
28322	1	80
29994	1	83
22123	2	74
27747	2	86
28011	2	85
28116	2	70
29208	2	91

This is visualised in figure 27, which shows the energy out against energy in for the shots which have been simulated in TRANSP. The green and orange 'x' markers are the corresponding subset of those in figure 24 (f). The solid lines of best fit are fitted only to this subset, and give 58% accounting for 1-beam and 55% for 2-beam. These are then corrected i.e. the "energy in" is reduced according to the percentage of neutral beam power which was not absorbed by the plasma, taken from table 6 to give the green and orange dots. The dashed lines are fitted to these dots, and give 68% accounting for 1-beam and 64% for 2-beam. Note that these numbers apply only to this subset of shots. It can clearly be seen that the correction moves the lines of best fit closer to the theoretical ideal fit line (shown in red). The accounting remains poor, because the subset of shots that were used for TRANSP simulations had poor accounting already, but the average improvement can be used in a comparison with the empirical data.

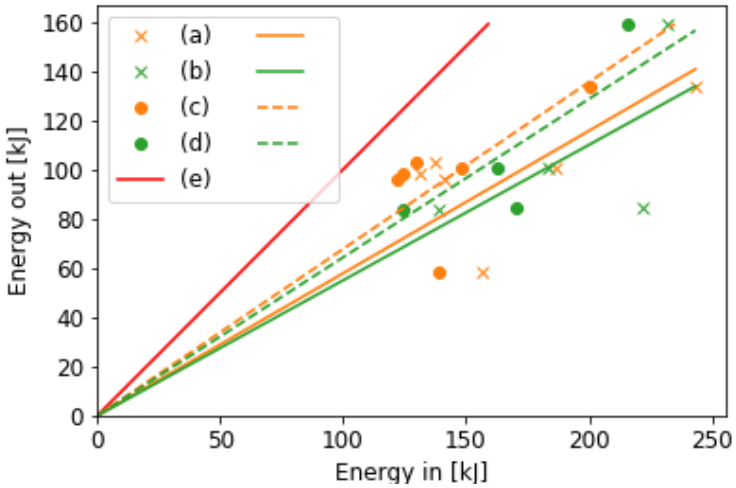

Figure 27: Energy out against energy in for those shots for which TRANSP has been run (given in table 6) with lines of best fit for each of (a) shots with a single neutral beam (b) shots with both neutral beams (c) 1-beam shots with the TRANSP correction applied (d) 2-beam shots with the TRANSP correction applied and (e) the ideal fit line.

4.6.2 Comparison of Simulation Results with Empirical Dependence

The motivation for performing these simulations was to see if they could explain the experimentally observed dependence of power accounting on neutral beam power, noted in section 4.4.3. There are thought to be two physical causes of this discrepancy, which are shinethrough and first orbit loss. The way these are calculated by TRANSP is not known in detail, because the code is closed source (although there is some information in [62]), but output variables are available that appear to allow these quantities to be calculated.

Table 7 firstly shows the mean of the results for 1- and 2-beam shots from table 6, and secondly the empirical dependence of power balance on neutral beam power, which is derived from the gradients of the orange and green best fit lines in figure 24 (f), adjusted for the discrepancy that is present even in ohmically heated shots (gradient of the blue fit line). The third column gives the hypothetical overall energy accounting if the simulation-derived absorption is assumed, without adjusting for the discrepancy present in ohmic shots. This shows that much of the experimentally observed NBI-related discrepancy between energy into and out of a plasma, can be explained by the effects accounted for in these TRANSP simulations. The agreement between simulation and experiment is excellent for 1-beam shots. This suggests that all of the beam power loss mechanisms for a 1-beam shot are captured in the TRANSP simulation code, which is a notable validation of the capabilities of TRANSP when used for this purpose. It is expected that TRANSP will be run routinely after every shot on MAST Upgrade from the second campaign onward, and that the data will be used in this way.

The worse agreement for 2-beam shots shows that there is another effect at play. This could be an interaction between the two beams, taking place in addition to the increased beam power present in 2-beam shots. It has been noted in the past that when both beams

Table 7: Comparison of the simulation-derived energy lost to shinethrough and first orbit loss from table 6, with the empirically derived dependence of power balance on neutral beam power from figure 24 (f), and the resulting overall energy accounting if the simulation-derived absorption is assumed to be correct.

Number of beams	Simulation-derived absorption [%]	Empirically derived absorption [%]	Simulation-adjusted energy accounting [%]
1	85	85	85
2	81	69	73

on MAST were in use, they caused a pronounced peak in the fast ion pressure around the magnetic axis, because both beams were located on-axis in the same quadrant of the machine. This pressure peak caused a large fast ion pressure gradient, which drove fishbone instabilities and caused fast ions to be ejected from the plasma before depositing their energy [73]. Fishbone instabilities derive their name from the fishbone-like trace that becomes visible in the fast ion spectroscopy when they occur. They are caused by fast ions colliding and depositing their energy with other fast ions, which then escape confinement because of their enormous energy, and fail to deposit their energy in the plasma [73]. This effect is not accounted for by TRANSP.

One the one hand, this shows that failure to account for shinethrough and first orbit loss was a serious omission from previous studies on MAST. This is not entirely surprising, given that first orbit loss is already well documented on other tokamaks, such as DIII-D [74] and EAST [50]. On the other hand, this still leaves an energy balance discrepancy that is dependant on neutral beam power (or possibly the number of beams), but not accounted for by TRANSP. This may originate from the fast ion pressure gradient caused by the specific arrangement of beams on MAST.

In addition to this, there is still some power balance discrepancy in shots with no neutral beams. This is typical of experimental power balance studies (see table 1), and may be because of heat deposited onto the first wall in core plasma areas where there was no IR coverage, or because of radiated energy in the divertor region where there was no bolometry coverage. Taken together, this shows that this energy balance analysis, accounting for wetted area and neutral beam power absorption, is consistent with previous work in the field for both ohmic and 1-beam shots. An explanation is proposed for the remaining discrepancy in 2-beam shots, although a numerical measurement of this effect is beyond the scope of this thesis. This demonstrates that the energy balance method developed here is valid, and can be applied to results from MAST Upgrade going forward (and this is done in chapter 6).

Chapter 5

Tile Surface Effects and Plasma Exposure

As mentioned in section 3.1, surface layer effects are a problem for infrared thermography and power balance calculations, especially in the context of carbon-walled tokamaks. Surface layer effects build up in tokamaks over the course of several years, under highly varying and irregular plasma exposure. This is especially true because of the inherently pulsed nature of a tokamak plasma. In order to quantitatively measure the effect of plasma exposure on a tile surface, a different type of machine is required, capable of producing long-lived steady-state plasmas. An appropriate machine would be a linear plasma device, especially one with a superconducting magnet. Magnum-PSI is just such a device, and this chapter describes how it was used to investigate the effect of plasma exposure on the surface of a MAST-U divertor tile. Many of the results in this chapter are already published as [75].

5.1 Design and Capabilities of Magnum-PSI

Magnum-PSI is a large linear plasma device located at the Dutch Institute For Fundamental Energy Research (DIFFER), on the campus of the Technical University of Eindhoven (TU/e) in Eindhoven, the Netherlands. Unlike a tokamak, a linear plasma device is not wrapped around into a torus shape that allows confined charged particles to remain in confinement for long times. Instead, Magnum-PSI is a single straight solenoid, in which the plasma is all generated in a source at one end and hits a target at the other end. It is also equipped with a "beam dump", which can be moved into place between the plasma source and the

target, allowing plasma exposure times to be tightly controlled, and protecting the target sample during plasma breakdown. Three dimensional models of the interior and exterior of Magnum-PSI are shown in figure 28.

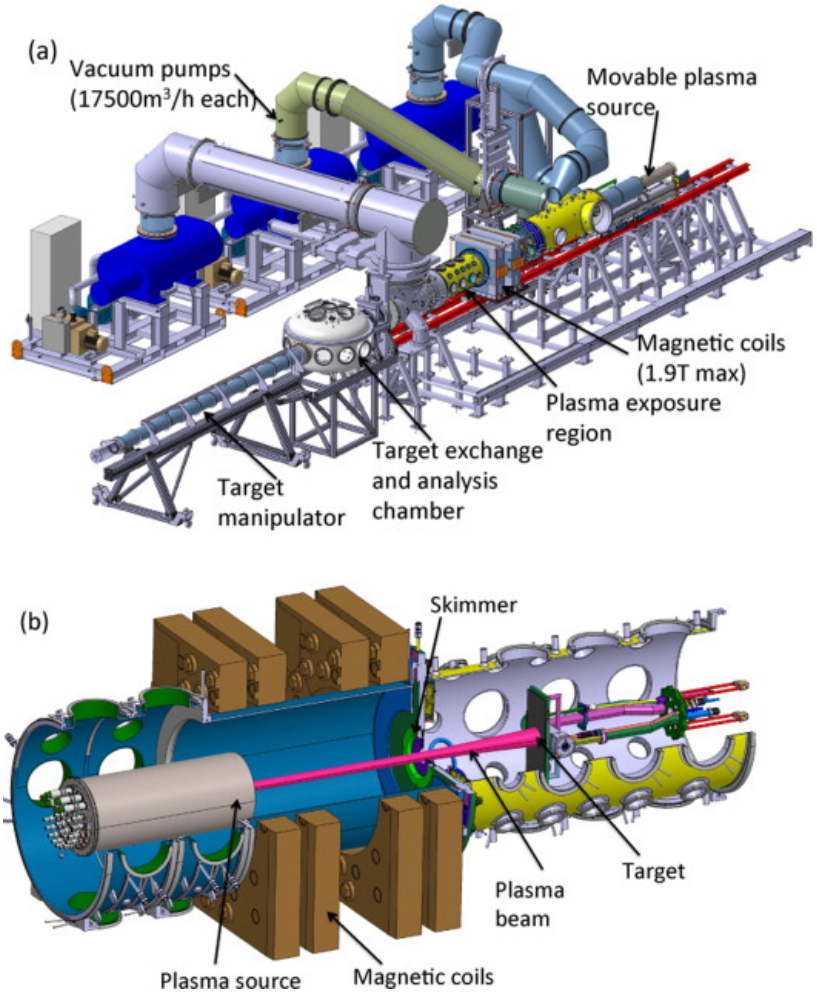

Figure 28: A CAD model of Magnum-PSI showing the (a) exterior and (b) interior setup, including vacuum pumps, magnets, and target & source locations. Figure from [76].

A detailed description of the high heat flux capabilities of Magnum-PSI is given in [76]. As explained therein, Magnum-PSI was designed to emulate the divertor conditions of ITER, to enable components proposed for use in the ITER divertor to be tested under comparable heat and particle loads. It was conceived to fill a gap in the capabilities between existing machines, especially regarding the emulation of ITER-relevant ELMs, with appropriate fluxes of particles as well as heat. The five principal design criteria, as given in [76], are:

- ITER divertor relevant plasma conditions

- Neutral pressure at target determined by recycling rate
- High magnetic field (3 T)
- Target geometry allowing shallow incidence angles, as low as 2.9 degrees
- In-vacuum surface analysis capabilities

These capabilities lead to Magnum-PSI being used to test tungsten divertor components for ITER [77]. In 2017, Magnum-PSI was upgraded with a superconducting magnet, enabling it to maintain plasmas for hours at a time [78]. This made Magnum-PSI capable of producing the half-hour long plasmas used in these experiments, which took place in 2019. It thus allowed plasma exposure equivalent to several campaigns on a tokamak to be applied to the tiles in mere hours. These capabilities have already been used to test the viability of liquid metal divertors for ITER [79] amongst other things.

5.2 Experimental Setup and Diagnostics

5.2.1 Sample Mounting and Setup

The basic experimental setup for the experiments is shown in figure 29, including key diagnostics. The Thomson Scattering measurement was taken in the centre of the plasma column. An indication is also given as to the arrangement of the IR camera and the pyrometer, but the dimensions are not to scale; the viewing angles are approximately 30° from the surface normal. The design of magnum-PSI allows samples to be swapped without venting the primary vacuum chamber, by moving them into the target exchange and analysis chamber (TEAC) which can be seen in figure 28 (a).

The sample was clamped into the mount, as shown in figure 30, taking advantage of a channel that already existed along one side of the sample, and cutting a similar channel into the other side. The mount is water cooled, so that experiments can take advantage of the long time plasmas made possible by the superconducting magnet without the sample melting or combusting. To aid thermal conduction between the sample and the mount, a layer of grafoil is placed between the two. This is a carbon-infused foil with a very high thermal conductivity, that has the effect of increasing the contact surface area between two otherwise rough, rigid surfaces. In these experiments, because of the dimensions of the channels and

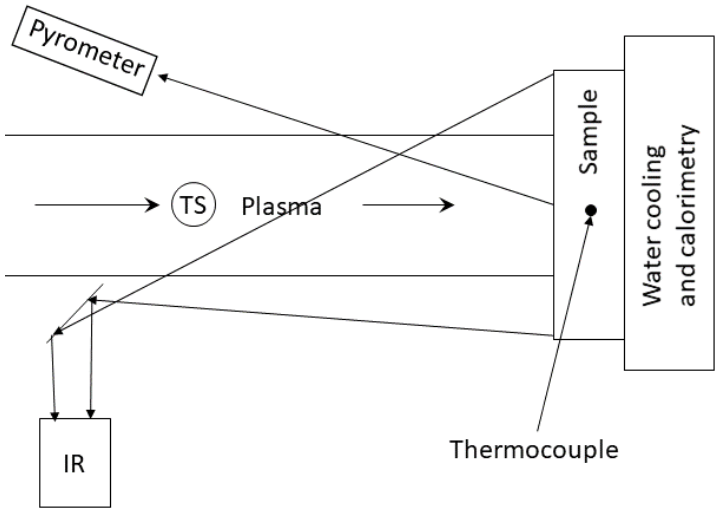

Figure 29: A diagram of the experimental setup, showing the position of key diagnostics relative to the sample (dimensions not to scale). TS refers to the Thomson Scattering laser.

clamps, a small piece of grafoil was placed behind the sample only, not extending behind the clamps (which is the more usual arrangement at Magnum-PSI).

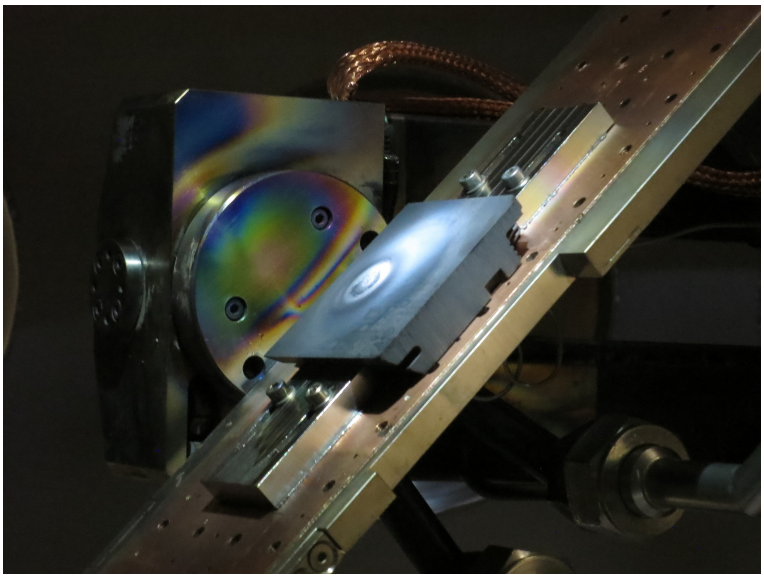

Figure 30: A photograph of the tile used in the steady state experiments, mounted in the large sample holder inside the TEAC of Magnum-PSI. One of the channels cut into the tile, by which it was held in place, is clearly visible.

5.2.2 Infrared Camera on Magnum-PSI

The principles of infrared cameras and thermography were described in section 3.1. The infrared camera on Magnum-PSI is a FLIR SC7500M infrared camera, similar to the medium wave cameras on MAST Upgrade, and filtered at 4 µm, with a 320×256 pixel resolution. It records at 300 Hz at full frame, and can record at up to 3.6 kHz with a quarter size frame. The spatial resolution is approximately 300 µm, with a telephoto lens.

The infrared camera at Magnum-PSI is manually set to record and await the trigger for each exposure. No data traces are calculated automatically between pulses, so any analysis must also be done manually, and data files are manually indexed with the correct shot number.

There is no automatic procedure for calculating the heat flux from the temperature, because of the difficulty of mapping the camera pixels onto real space. This cannot be assumed from a standard setup or machine diagram, because of the different (and sometimes bespoke) sample mounting arrangements, and the fact that the mirror angle must be changed if the mounting arrangement changes the position of the sample along the machine axis (the mirror is shown in figure 29). If heat flux calculations are desired, they must be done manually depending on the sample mounting geometry. All the IR data from Magnum-PSI presented in this thesis is in units of temperature; a conversion to heat flux was not performed.

Having manually recorded the IR data, there are three stages of analysis that must be completed. Using the specialist software from FLIR (called Altair), a region for the analysis must be chosen. This can be a line, circle or polygon, which outputs values such as maximum, minimum and mean. For this analysis, several small circles were used, as shown in figure 31. The mean value within the circle was used as the output. Using the mean of all pixels within the circle helps to mitigate digitisation noise, or the variation between responses of individual pixels.

The second stage is setting a temperature bias in the software (equal to the temperature in the room) otherwise the pixels left with zero photon count after reference frame subtraction will be treated as 0 °C. The third stage is the conversion from "digital level" (an analogy for photon count) to temperature, which is done by the proprietary software according to a factory calibration. This then gives a single calibrated temperature as a function of time, for each circle, which was used for the analysis.

a) 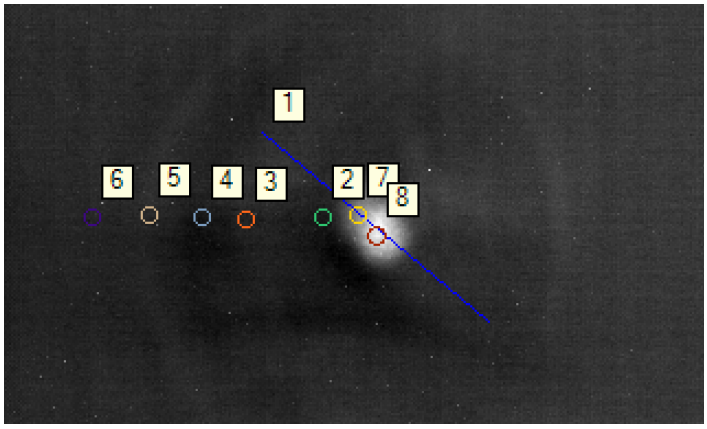

b) 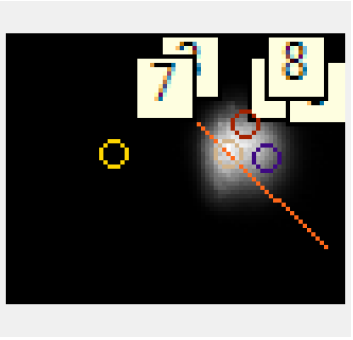

Figure 31: Camera frames (a) full frame and (b) quarter size, taken from the IR camera at Magnum-PSI, showing the locations of the circles used in the analysis of the (a) steady state and (b) ELMy experiments. The data traces in figures 32 and 33 are colour-coded according to the circles in (a), and those in figure 36 according to the circles in (b). In both cases, the white area is the location of the incident plasma.

5.2.3 Spectro-Pyrometer on Magnum-PSI

Pyrometers were described in section 3.5.1. The pyrometry setup at Magnum-PSI consists of an FAR SpectroPyrometer, model FMP1, which operates at $900 - 1700\,\text{nm}$. It collects intensity data at 250 different wavelengths with nanometre-scale bandwidths within the stated range, allowing very accurate and precise measurements. The advertised temperature measurement range is $300 - 2000\,°\text{C}$, although the minimum measurement for a carbon sample is closer to $550\,°\text{C}$ because of the low emissivity of the material.

Unlike the infrared camera, the spectropyrometer has no spatial resolution, and its time resolution is poor (approximately $0.2\,\text{Hz}$). The two techniques can be used in tandem, providing the clocks are synchronised, to give a thermal image of a sample that has good spatial and temporal resolution while being free from the effects of emissivity. This was not possible with the data reported in this thesis, because most experiments were done outside the operating range of the pyrometer, and the few that were within the range were at the operational limit, so only brief measurements of emissivity-independent temperature were obtained.

5.2.4 Thermocouple and Calorimeter on Magnum-PSI

The principles of operation of thermocouples and calorimeters were described in section 3.5.3. At Magnum-PSI, there is capability for multiple thermocouples, which can be attached to

various points on experimental samples. These report temperature to within 1 °C every 2 s. Exposing a thermocouple directly to a Magnum-PSI plasma would result in damage, so they must either be attached to the back of the samples, or embedded within them. For the experiment reported here, a single thermocouple was embedded about half way through the depth of the sample, in a hole drilled for the purpose, as shown in figure 29. Unfortunately, this proved to be too deep to measure significant temperature change, so the thermocouple data was not used in the analysis.

Magnum-PSI also has a calorimetry system, which is integrated into the sample holder, and gives a measurement of heat leaving the active water cooling system. This heat has been delivered to the tile by the plasma. The system is properly time-adjusted as described in section 3.5.3, and has a constant volume flow rate, so the heat energy removed from the tile can be calculated. However, the calorimetry system suffers from inaccuracies because of the cooling effect of the sample coolant water flowing near the liquid helium used to cool the superconducting magnet. For this reason, the calorimetry data from Magnum-PSI was not used in this analysis.

5.2.5 Non-Thermal Diagnostics Used

The principles of laser Thomson Scattering were described in section 3.3. Magnum PSI is equipped with a Thomson scattering laser that reports both electron density and temperature. Because of the low temperatures and densities in Magnum-PSI plasmas, the laser takes multiple measurements at a rate of 2 Hz and reports an average over a certain time. The laser is a Nd:YAG (neodymium-doped yttrium aluminium garnet) laser, similar to that used on MAST, which emits in the infrared portion of the spectrum.

The system is capable of measuring a plasma column of 100 mm diameter, with a spatial resolution of 1.4 mm to an observational error of 3% in n_e and 6% in T_e [80]. The system was designed to be movable between $15 - 55$ mm from the target, but this assumes the small sample holder, which was not used in the experiments reported here. Instead the measurements were taken 81.5 mm from the target.

Profilometers were described in section 3.5.2. The surface profilometry in this experiment was not performed at DIFFER, but at Forschungszentrum Jülich, in Germany, using a laser profilometer with a KF3 sensor from OPM Messtechnik GmbH. This profilometer has a lateral (x and y) resolution of 10 µm and a profile (z) resolution of 20 nm. It should be

noted that the error in the profile measurements can be somewhat more than this however, because it must be averaged over the comparatively large lateral resolution. In practice, this can result in profile measurement error of $\pm 1\,\mu m$ for a sufficiently small artefact.

The profilometer also reports surface reflectivity, at a wavelength of 670 nm which is intended as an error measurement to indicate the reliability of the surface profile measurement at that location. The spatial resolution is the same as the profile measurement, 10 µm, and the reflectivity is reported as a percentage.

5.3 Method for Steady-State Plasma Experiments

The tile used was a fine grain graphite tile, identical to those that will be used in the MAST-U divertor. The sample measured approximately $70 \times 70 \times 15\,mm$. Magnum-PSI can deliver a continuous plasma of $\approx 2\,cm$ beam radius (FWHM) to the tile surface for extended periods of up to several hours, allowing a plasma exposure equivalent to several MAST-U campaigns in a single day.

In order to maintain a tile surface temperature comparable to that expected in MAST-U, Magnum-PSI was operated with a magnetic field of 0.5 T, source current of 140 A and hydrogen gas flow of 10 standard litres per minute (slm) with 1.5 slm gas puffing near the target. This resulted in plasma electron density $n_e \approx 5 - 7 \times 10^{18}\,m^{-3}$ and electron temperature $T_e \approx 0.07 - 0.09\,eV$. Although these are much lower than the temperatures and densities typical in MAST, this had to be balanced against the desire for the tile surface temperature to be comparable to that in MAST, which is typically no more that 100 °C. The tile was exposed to plasma for six 30-minute intervals with the parameters above. These 30-minute plasma exposures were continuously monitored by the IR camera at 25 Hz. There were TS measurements at five-minute intervals, each of which was averaged over 100 measurements to minimise the statistical uncertainty of the measurement at such low temperatures and densities.

In order to allow additional temperature measurements with a multiwavelength pyrometer (which operates only at temperatures above 550 °C), the tiles were exposed to higher plasma power for 1 minute intervals, in between the longer plasma exposures mentioned above. Because it measures at multiple wavelengths, the pyrometer gives a measurement that is independent of the tile emissivity. This technique can even be used to measure

emissivity [59]. The higher plasma power to the tiles was achieved by lowering the target gas puffing from 1.5 to 0.5 slm, resulting in plasma conditions of $n_e \approx 2 \times 10^{19}\,\mathrm{m}^{-3}$ and $T_e \approx 0.15\,\mathrm{eV}$. This allowed the more accurate pyrometer temperature readings to be taken between each of the longer exposures, to monitor any real temperature change over time.

During both the long and short exposures, other standard Magnum-PSI diagnostics were running, including the thermocouple and calorimeter. The intention was for these to act as a temperature calibration to distinguish any real change in temperature from an IR-reported change caused by a change in emissivity, but in the event only the pyrometer was used for this purpose.

5.4 Results from Steady-State Plasma Experiments

5.4.1 Tile Surface Temperature Reported by the Infrared Camera

Figure 32 shows the evolution of the IR-inferred temperature of the tile when it was exposed to six 30-minute plasma discharges. The temperature at the centre of the beam and at a location 5 mm from the centre of the beam are shown as a function of the plasma exposure time, assuming an emissivity of 0.65. It can be seen that the apparent temperature as measured by the IR camera decreased by about 70 °C during this time. This is significant, and shows the susceptibility of infrared camera measurements to surface effects. The trend is approximately linear, except for the second exposure, suggesting that it was not approaching a saturation point, but would have continued to decrease for several more hours. Temperature measurements from the thermocouple are not shown, because the thermocouple only registered a temperature 1 °C above room temperature during the long pulses, except during the fifth pulse, when it recorded a 2 °C increase. Calorimeter data is not shown either, as explained above.

The heat flux to the tile has been estimated using the electron temperature and density measurements from the Thomson scattering measurements, using a method verified on Pilot-PSI [82]. Briefly, the method entails calculating a coefficient using an axis-symmetric Ansys model, and combining this with the TS data to infer the target heat flux, based on the calibration between TS data and tile calorimetry performed in [82]. This data is shown in the blue squares in figure 32, for which the line of best fit (not shown) has a gradient of $0.0048 \pm 0.0002\,\mathrm{kW\,m^{-2}\,min^{-1}}$, so clearly there is no downward trend.

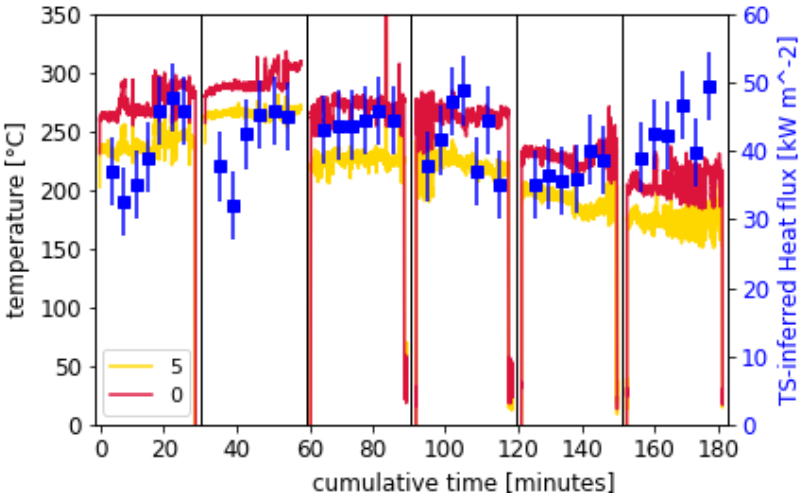

Figure 32: The temperature of certain regions of the tile as measured by the IR camera (labelled by distance from the centre in mm) throughout the long pulses. These are coloured according to the colour of the circle in figure 31 (a) where the data was taken. On the second y axis, the heat flux inferred from Thomson scattering, using a similar method to [81] (blue squares) with error estimated from the instrument resolution.

The results of the short exposures are shown in figure 33. Similar to the long exposures, the infrared camera reports a decrease in surface temperature during (and in some cases between) these shorter pulses, of about 140 °C. The heat flux derived from TS data is shown as blue squares, and suffers from random noise. The gradient of the best fit line of the blue squares (not shown) has a gradient of $-0.15 \pm 0.02\,\mathrm{kW\,m^{-2}\,min^{-1}}$, but if the first point is excluded, the gradient of the line of best fit of the other five points is positive, so clearly there is no steady downward trend that would explain the steady decrease in temperature.

Temperature readings from the pyrometer were available for the short pulses because of the higher temperatures, and are shown as red dots. The gradient of the line of best fit of these dots (not shown) is $-0.13 \pm 0.02\,\mathrm{°C\,min^{-1}}$, which would suggest a temperature decrease of 23 °C over three hours, considerably less than the 140 °C reported by the infrared camera. This suggests that the temperature was not actually decreasing at a noticeable rate. The thermocouple recorded a temperature increase of 6 ± 1 °C during the short pulses. If the temperature had in fact been decreasing as suggested by the IR, a 1 °C drop would be expected on the thermocouple, however this is within the measurement error. The agreement between IR and pyrometer temperature in the fifth exposure suggests that the emissivity was equal to the assumed value of 0.65 at this time (after 2.5 hours of exposure).

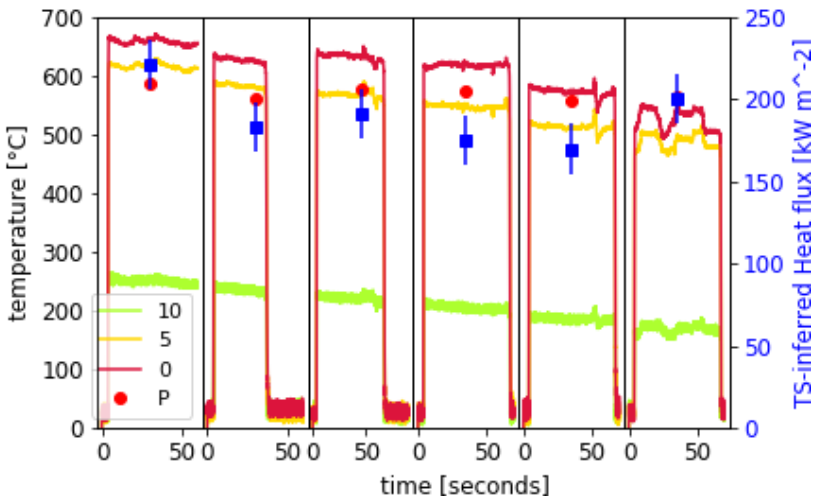

Figure 33: The temperature of certain regions of the tile as measured by the IR camera (labelled by distance from the centre in mm), throughout the short pulses. These are coloured according to the circle in figure 31 (a) where the data was taken. Also the temperature measured by the multi-wavelength pyrometer (red dots labelled P). On the second y axis, the heat flux inferred from Thomson scattering, using a similar method to [81] (blue squares) with error estimated from the instrument resolution.

In both figures 32 and 33 the heat flux is calculated from the TS data using appropriate temperature-dependent values of sheath heat transmission coefficient γ. This coefficient depends on the ratio of ion temperature to electron temperature, and numerous other parameters including particle reflection coefficients and recombination energies, so its calculation is not trivial [82]. The values calculated in this work, in the region of $\gamma \approx 100 - 200$, are very high compared to the much lower values of $\gamma \approx 4 - 40$ calculated in [82], because the non-linear inverse temperature dependence of γ becomes significant at such low temperatures.

Two dimensional modelling of the tile temperature evolution shows that this calculated heat flux is too low to reproduce the observed temperature rise of the tile. A simple finite-element model was used, simulating only the first 450 ms of each pulse, which is the time taken for the heat to conduct to the back of the tile, since the model does not account for active cooling (note that this time is much shorter than the time the tile would take to reach thermal equilibrium). To match the temperature rise seen in the experiment, the required heat flux is of the order $2\,\mathrm{MW\,m^{-2}}$ for the 30-minute exposures and $7\,\mathrm{MW\,m^{-2}}$ for the 1-minute exposures. These calculated values are comparable to the machine input power. However, it is conceivable that the normal method for calculating γ no longer applies at temperatures $< 0.1\,\mathrm{eV}$. The determination of an appropriate value of γ is the subject of

ongoing study.

In both the long and short pulses, the IR-inferred temperature of the tile was found to decrease during and between plasma exposures, despite the heat flux (inferred from either TS or machine power) remaining constant. Investigation of the rate of cooling at the end of each exposure showed no evidence of the sudden cooling associated with a surface layer (such as that shown in figure 13 and explained in detail in [53]). The pyrometer in the short exposures showed no decrease in temperature. Such an inferred decrease in temperature could be explained by a change in tile surface emissivity.

5.4.2 Comparison to Past Experiments on MAST and ASDEX Upgrade

A decrease in emissivity over time, as seen in the centre of the Magnum-PSI tile, is the opposite of what had been expected based on past experiments on MAST and ASDEX Upgrade. Figure 34 is an image of the MAST divertor taken soon after one of the divertor tiles had been replaced (the one in the centre of the image). It can be seen that the strike point on the recently replaced tile is less bright that on the other tiles, because that tile had a lower emissivity than the others. The heat flux at the strike point can be assumed to be toroidally symmetric, suggesting that the new tile is causing the difference in brightness. This could potentially be caused by the tile being mounted at a different angle. However, this is unlikely, both because the angle is fixed by the structure, and is not variable; and because the dark bands (the non wetted area of the tile) are the same width on both sides (the width of each band is dependent on the angle of the adjacent tile, as explained in section 4.3.2 and shown in figure 21).

Rather than the angle causing low brightness on the new tile, it is thought that the old tiles around it had increased in emissivity over the course of their lives. The same effect has also been observed on ASDEX Upgrade [44]. This was expected to be seen across the entire tile in Magnum-PSI when the experiments were being planned. However, the centre of the tile, where most of the measurements were made, decreased in emissivity instead, making it difficult to compare with these past results.

The decrease in emissivity is believed to have come about because the plasma conditions in Magnum-PSI are somewhat different to those in MAST or ASDEX Upgrade. The conditions are thought to have crossed the threshold where surface erosion becomes dominant over re-deposition. Some theoretical background to this was given in section 2.3.1, and some

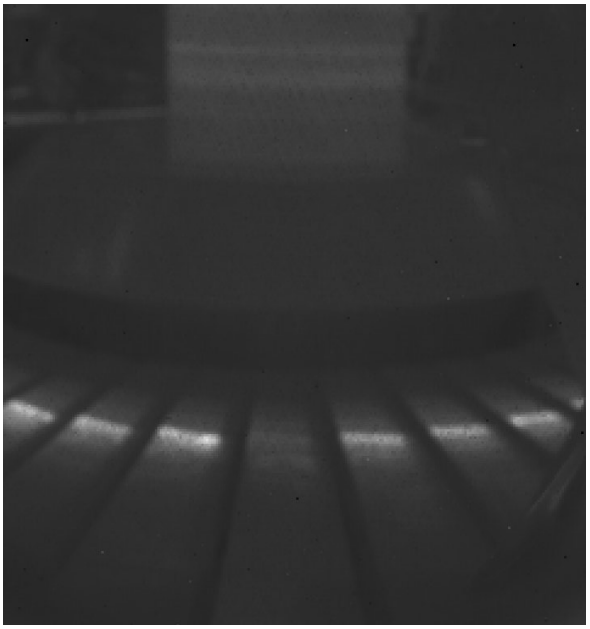

Figure 34: Some of the tiles in the lower outer divertor of MAST during shot 25735. The MAST centre column is visible at the top. The tile in the centre of the image had just been replaced, and the strike point appears very faint on this tile compared to the others.

experimental verification of this hypothesis will be given in section 5.8. MAST was incapable of creating such conditions in its divertor, because of its small plasma volumes, and corresponding low divertor heat flux (this dependence is explained in section 1.1.4). Magnum-PSI is designed to emulate the conditions in the divertor of larger future tokamaks, specifically ITER. However, since the ITER divertor tiles will be made of tungsten and not carbon, the plasma power threshold between erosion-dominant and deposition-dominant regimes is likely to be different, so this result may not be directly relevant to ITER.

5.5 Method for ELMy Plasma Experiments

After the steady state experiments had taken place, another identical graphite tile measuring approximately $70 \times 70 \times 15\,\text{mm}$ was placed into Magnum-PSI. The sample mounting setup was identical, and has been shown in figures 29 and 30. This time, the ELM emulation capabilities were used, to try to replicate ELMy conditions in MAST-U, in a similar fashion to the detached steady state conditions replicated in the first experiment. ELM emulation in Magnum-PSI is achieved using a capacitor bank to provide a regular electrical discharge. The full capabilities of the ELMy plasma source and capacitor bank of Magnum-PSI are

described in [83], although not all of these capabilities were used in this experiment. The steady state plasma source was used in conjunction with the capacitor bank, to emulate the small ELMs expected on a smaller machine like MAST-U (the oscillating source described in [83] is intended to emulate larger ITER-relevant ELMs).

The same machine conditions as in the steady state experiments were used for the plasma source, including plasma current and applied magnetic field, with the exception that there was no gas puffing. This resulted in inter-ELM plasma conditions of $n_e \approx 2 \times 10^{19}\,\mathrm{m}^{-3}$ and $T_e \approx 0.6\,\mathrm{eV}$. The repetition rate of the capacitor bank is fixed at 10 Hz, and in this experiment the magnitude of the capacitor voltage was set to 500 V. At the peak of the ELM, the plasma conditions were $n_e \approx 12 \times 10^{20}\,\mathrm{m}^{-3}$ and $T_e \approx 3\,\mathrm{eV}$, from TS measurements. The TS system is synchronised with the ELM system and takes several repeated measurements, at the peak of multiple identical ELMs during the same plasma exposure, in order to produce these numbers.

The tile was exposed to the same plasma conditions, including ELM conditions, for several 10-minute periods. The tile surface temperature was monitored using the IR camera as before, although the recordings were shorter, comprising only the first 30 seconds of the plasma exposure time. This was done to save storage space, since the camera was set to its maximum frame rate (at one quarter frame size) of 3598.4 Hz in order to capture the warm up and cool down associated with each individual ELM. The surface temperature evolution at this frame rate is shown in figure 35, and was sufficient for 3 to 4 data points during the warm up, and many more during the cool down. The fact that the ELMs all appear a fairly consistent height in the data also demonstrates that the frame rate was sufficient. The procedure was repeated a total of 12 times. There is one exception to this, when the camera frame rate was slowed deliberately to observe the entire plasma exposure. This enabled the temperature evolution of the tile surface to be studied over the 10 minutes, but made the resolution far too low to observe the ELMs. The thermocouple and calorimeter have low time resolution, and would not have provided any insight in the ELMy experiments, so they were not used.

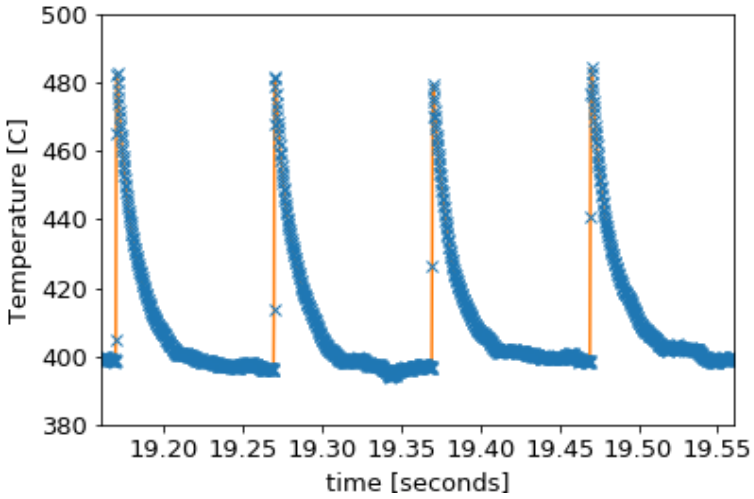

Figure 35: A graph of temperature against time during four ELMs from the third 10-minute ELMy pulse on Magnum-PSI. The raw data is shown in blue, with linear splines in orange. It can be seen that there are at least three splines during the warm-up of each ELM.

5.6 Results from ELMy Plasma Experiments

5.6.1 Tile Surface Temperature Reported by the Infrared Camera

Figure 36 shows the temperature of the tile during the first 30 seconds of each of the twelve 10-minute exposures to ELMy plasma. Individual ELMs cannot be seen on the compressed time axis in the figure, although the camera frame rate was fast enough to capture them as shown in figure 35. Instead only a thick band between the inter-ELM and peak ELM temperatures is visible in the figure. The temperature did not exhibit a steady decrease as with the steady state plasma, but remained relatively constant. The change in temperature caused by the ELMs, ΔT, was approximately 150 °C for the points near the centre, from a background temperature of around 400 °C. This compares to a room temperature baseline for the point 12 mm from the centre, but still an ELM ΔT of almost 100 °C. The furthest point was no longer in view of the camera after the target position scan (see below) was performed so this data is not available for the last four exposures. This is because the angle of the mirror for the IR camera (see figure 29) had to be adjusted to keep the plasma centre in view during the target position scan. It was not changed back when the target was moved back, so the plasma centre was in the left of the frame for exposures 8 to 12, contrary to its position in the right of the frame in figure 31 (b).

There is no pyrometer data for this experiment, because the background (inter-ELM)

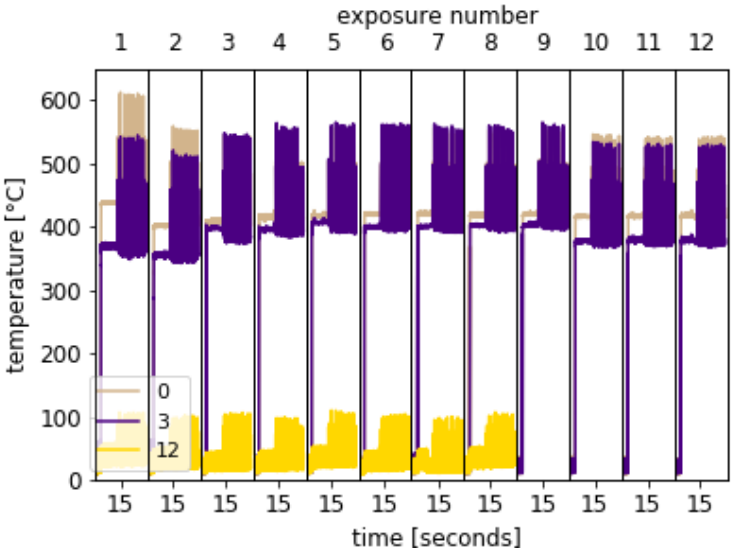

Figure 36: A graph of temperature against time during the first 30 seconds of twelve 10-minute plasma exposures for selected locations on the surface of the tile used in the second, ELMy, set of experiments on Magnum-PSI. The y axis labels indicate the half way point in each of the 30-second recordings, which are not consecutive because of the intervening 10-minute exposures. The legend indicates distance from the centre of the plasma column in mm, and traces are coloured according to the circles in figure 31 (b).

temperature of the tile surface was kept too low. This was done to keep the peak ELM temperature below a level that would be severely damaging to the tile. However, there was some damage to the tile surface, as expected from the steady state experiments, and past ELMy experiments on tokamaks [84].

One of the features of this surface damage was the formation of a hot spot, which increased the measured inter-ELM surface temperature. The peak ELM temperature remained fairly constant, making the temperature difference caused by the ELMs appear smaller. This is difficult to see by eye, but upon close inspection of figure 36, the baseline temperature between exposures 3 and 9 increased, whereas the peak ELM temperature did not. Exposures 1 and 2 were performed on a different day, and exposures 9 to 12 were during and after a target position scan, so the apparent temperature changes between them may not be the direct result of plasma exposure.

It is not entirely clear why the hot spot was not affected by the higher temperature during the ELM pulses. The hot spot was within the plasma column, and certainly received flux from the ELMs. However, the hot spot was observed to warm up more slowly than the rest of the tile at the beginning of each exposure. That is to say, it had a longer

characteristic warming and cooling time; too long to be affected by the ELMs. This is one possible explanation as to why only the inter-ELM temperature increased over time, and not the peak ELM temperature.

A target position scan, along the axis of the plasma column, was done at the end of this experiment, comprising exposures 9 to 12 in figure 36, with the aim of taking TS measurements at different distances from the sample surface. Unfortunately, the TS was not working correctly that day, so there is no such data available. Peak ELM TS data is only available for exposures 1 and 2, which had taken place the previous day. Inter-ELM TS is estimated from steady state shots earlier that day. There was visible damage to the surface after exposure to ELMy plasma, similar to the tile that was exposed to steady state plasma. Unfortunately, there was a problem which made it impossible to remove and take profile data without damaging the tile.

5.6.2 Comparison to Steady-State Plasma

In both steady state and ELMy cases, areas of the tile that were further from the centre of the plasma column warmed up to a lower temperature than those near the centre, as expected. The change in temperature (or rather emissivity) over time is more interesting. The surface emissivity of the tile in the steady state experiments clearly decreased over the three hours of plasma exposure, probably because of relatively uniform (within the plasma column) surface erosion caused by the plasma.

However, the tile used in the ELMy experiments developed a hot spot on the second day, which increased in temperature over the course of the two hours of plasma exposure. This shows that any surface effects caused by the ELMy plasma were not uniform within the plasma column. It is worth noting that the background plasma during the ELMy experiments was cooler than the plasma used for the steady state experiments, so may have been below the threshold where deposition dominates. However, it seems likely that continuous ELMs at a rate of 10 per second would have caused net erosion overall. The exact cause of the hot spot is unknown, but it is worth noting that its characteristic warming and cooling time was longer then that of the remainder of the tile surface, suggesting that is was not a surface layer effect of the type seen on MAST.

5.7 Motivation and Method for Post Mortem Analysis of Tile

After the experiments had taken place, there was visible damage to the surface of both tiles. The tile from the steady state experiments also had visible concentric rings on the surface, see figure 37. The larger, non-concentric ring is caused by the Magnum-PSI skimmer. The skimmer, visible in figure 28 (b), is a metal ring used to separate the different chambers of Magnum-PSI, designed to stop most of the neutral gas flow from reaching the target. This circle is not of any interest. In addition, there was significant pockmarking at the centre of the tile, where most of the plasma flux had been incident. The concentric rings and erosion were not expected, because tiles from experiments in MAST had exhibited surface layer deposition. The tile surface was therefore worth investigating further.

The tile was sent to Forschungszentrum Jülich, for a surface topology profile to be taken. This offers insight into whether the observations from the experiments were caused by erosion or deposition (or a combination thereof) and the extent of the erosion or deposition. This would also reveal the nature of the concentric rings visible on the tile. The profilometer also reports surface reflectivity, and since the reflectivity of the tile is equal to 1 minus the emissivity of the tile (at the same wavelength) the data is of interest to this experiment in its own right. Unfortunately, although the data is good quality, it is in the visible spectrum, not the infrared, so does not correspond directly to the emissivity measured by the IR camera.

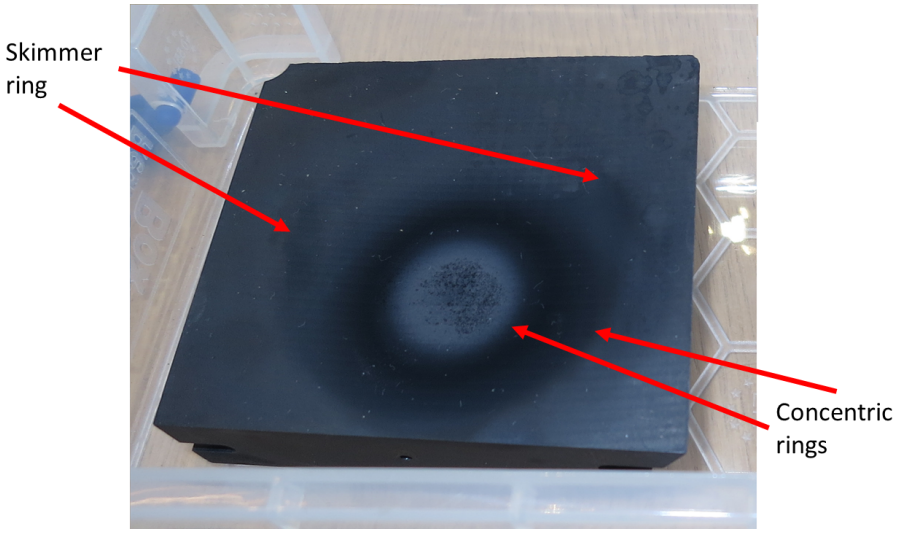

Figure 37: A photograph of the tile from the steady state experiments in a box, ready to be shipped to Jülich for a surface profile. The concentric rings are visible, and the skimmer ring is the larger ring with its centre offset upward relative to the others.

A direct experimental measurement of the relative emissivity of different parts of the tile was also desired. So, once the tile had arrived back at Culham, it was placed unheated in view of a long wave infrared (LWIR) camera. This was the Thermo Sensorik camera from MAST, which was mentioned in section 4.2.1. LWIR was used because it is more sensitive than MWIR to temperatures close to room temperature.

5.8 Results from Post-Mortem Analysis of Tile

5.8.1 Tile Surface Profile

The surface profile was taken over the region of interest on the tile surface, as shown in figure 38. The orientation of the tile in the profilometer is also consistent with this photograph. The highlighted region includes the pockmarking and concentric circles, except the one caused by the skimmer. There are small areas of undamaged tile that can be used for the $z = 0$ reference in the corners of this region, particularly the lower corners.

A contour plot of the surface topology, given in figure 39, shows that the pockmarks in the centre of the tile are scattered randomly, and appear to all have a similar depth of approximately 80 µm. The concentric circles are also visible, and it can be seen that they

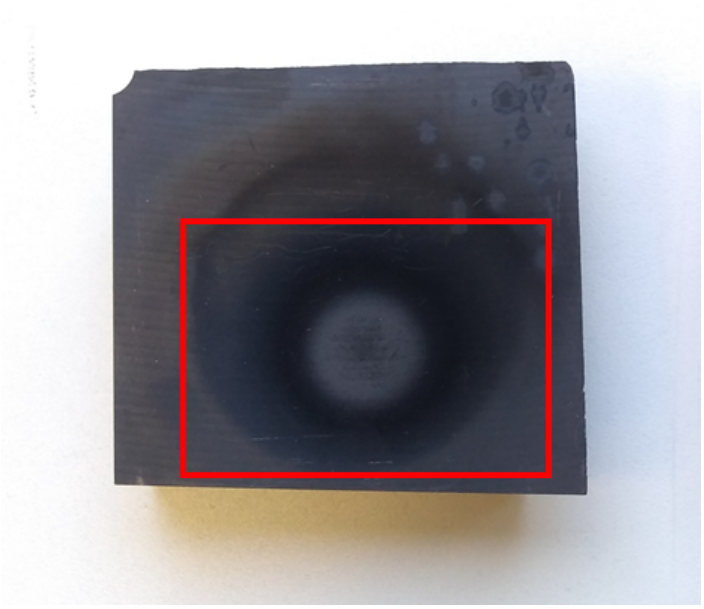

Figure 38: Photograph of the tile surface after the experiments, with lighter and darker rings visible. The red box indicates the portion of the tile shown in figures 39, 40 and 41.

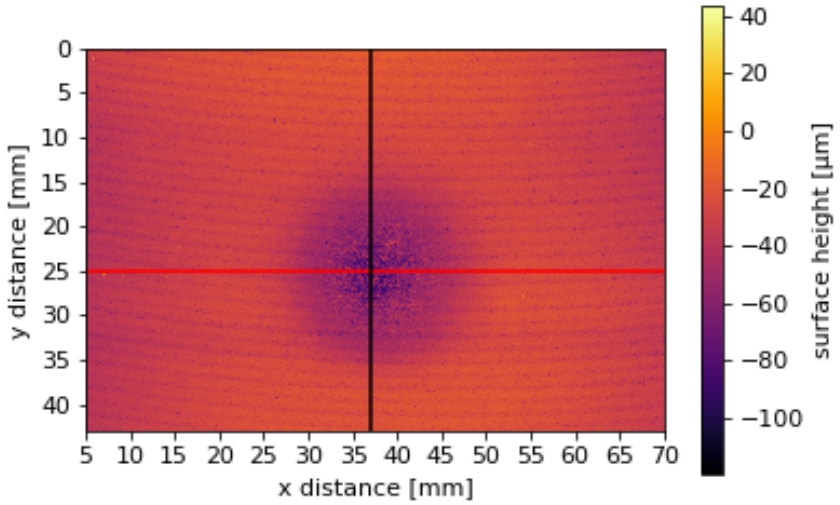

Figure 39: The surface profile of part of the tile that was exposed to the plasma (indicated by the red box in figure 38). Profile (z) resolution is 20 nm averaged over 1 unit of the lateral (x and y) resolution, which is 10 µm. The red and black lines correspond to the slices shown in figure 42 (top).

equate to erosion and deposition. Erosion causes the surface emissivity to lower (it appears lighter) and erosion of up to 100 µm is observed within the FWHM of the plasma column. At such low plasma temperatures and densities, this is likely due to chemical sputtering rather than physical sputtering. At the periphery of the plasma, where the plasma power was lower, it was probably more similar to the MAST divertor, although this cannot be confirmed because TS is only available in the centre of the plasma column. (It should be noted that the plasma flux in Magnum-PSI is perpendicular to the tile, unlike the very shallow angles common in MAST, so the equivalent tile heat flux corresponds to a much lower plasma density and temperature in Magnum-PSI than in MAST). In this region, the plasma-surface interactions were deposition dominated, with deposits up to 40 µm thick. This explains the higher emissivity, resulting in visible darkening.

5.8.2 Changes to Tile Reflectivity and Emissivity

The profilometer that produced figure 39 also reports surface reflectivity at the laser wavelength of 650 nm, as a proxy for measurement error. In this reflectivity data, the rings mentioned above are also visible, see figure 40. The higher reflectivity measured in the cen-

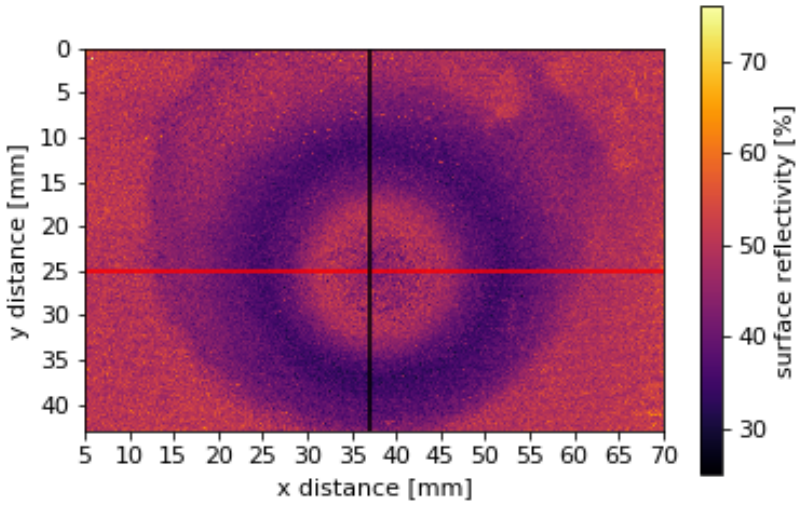

Figure 40: The surface reflectivity of part of the tile that was exposed to the plasma (indicated by the red box in figure 38). The measurement was made at 670 nm. The red and black lines correspond to the slices shown in figure 42 (middle).

tre of the tile corresponds to the lower emissivity observed in the experiment, and likewise the corresponding lower reflectivity around the edge of the plasma column results in a darker surface and a higher emissivity. Unlike the results from the tile centre, this observation from around the edge of the plasma column aligns with the results from MAST and ASDEX Upgrade mentioned in section 5.4.2.

The image of the tile taken with the LWIR camera after it arrived back at Culham is shown in figure 41. The small darker circle (at $x = 30-50$ mm and $y = 15-35$ mm) off-centre within the lighter circle, corresponds to the same circles visible at the same coordinates in figures 39 and 40. Because no heat sources were being applied to the tile when figure 41 was taken, the tile was uniformly at room temperature. Therefore, the inferred temperature gradient across the surface can only have been caused by a spatial variation in emissivity.

In figure 42, slices in the x direction (left) and y direction (right) are shown for the preceding three figures. They show that the centre of the tile has a lower surface height and emissivity, and a higher reflectivity, than the surrounding areas. It can also be seen that the boundaries of these changes are all approximately aligned in both x and y. These slices, especially for reflectivity and emissivity (middle and bottom) reveal that there is a sharp transition between the high and low areas, which is probably the exact location of the edge

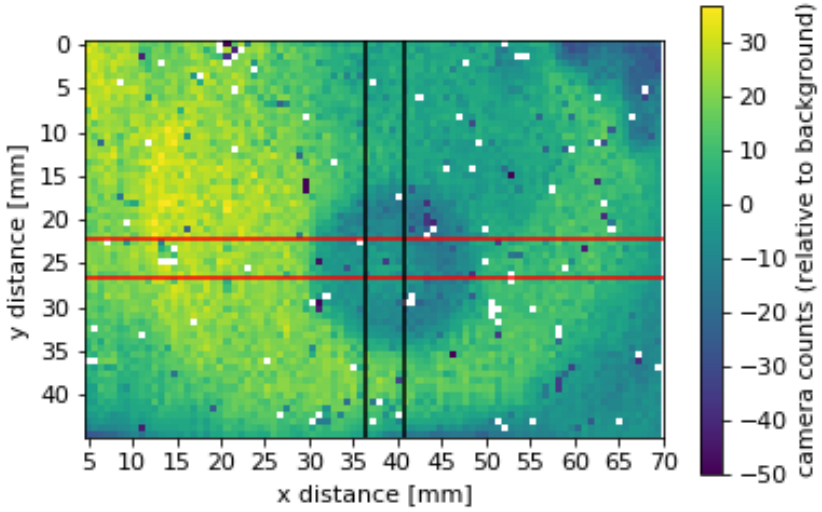

Figure 41: An infrared image of the tile, at room temperature with no heat applied, showing the contrast caused by the varying emissivity. Pixels with a very low value are assumed to be dead and have been removed (shown in white). The red and black lines correspond to the slices shown in figure 42 (bottom).

of the plasma column.

Taken together, the results in figures 39 - 42 demonstrate that the plasma incident at the centre the tile in Magnum-PSI was sufficient to cause net surface erosion, resulting in decreased emissivity. This indicates that the plasma power was above the threshold of chemical sputtering, unlike on MAST. It is also evident that the power at the periphery of the plasma was below this threshold, and that surface deposition and an increase in emissivity occurred there. However, several diagnostics on Magnum-PSI, including Thomson scattering and the pyrometer, only provide information in the centre of the plasma column, where the conditions were not like those on MAST. Therefore, the detailed results that rely on these centre-only diagnostics are not directly relevant to MAST or MAST Upgrade, despite the original intention and motivation for these experiments. Despite this, the results are interesting in their own right, for clearly demonstrating that plasma exposure causes measurable surface changes, and that these lead to significant measurement errors. This represents an original contribution to knowledge, and has been published as [75].

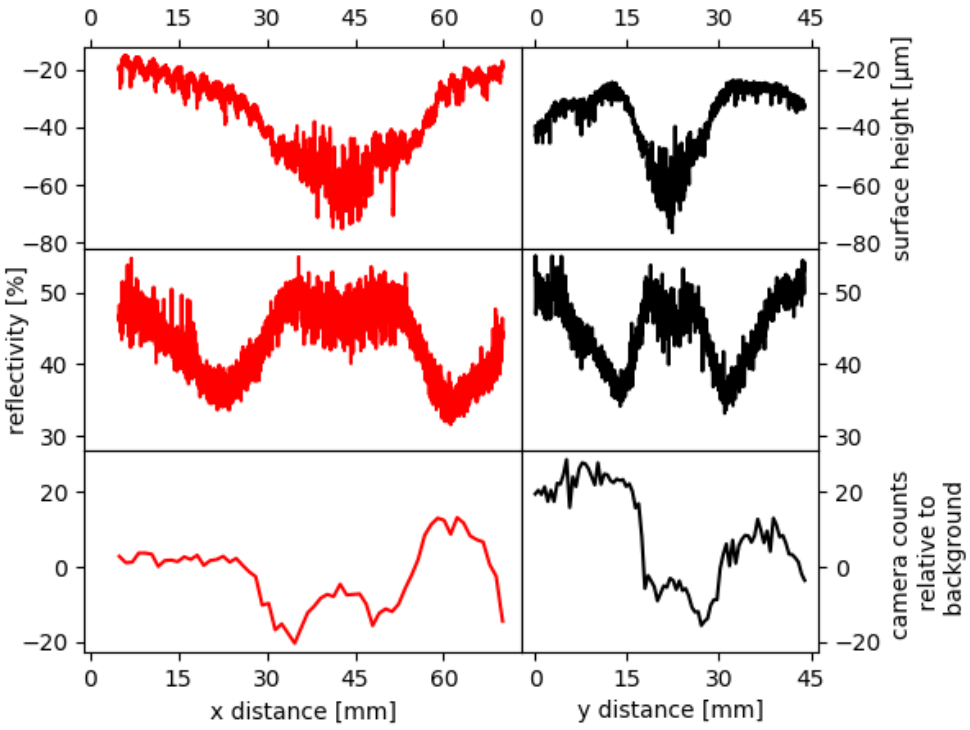

Figure 42: A comparison of the x and y slices taken from: (top) surface profile data in figure 39, (middle) surface reflectivity data in figure 40, and (bottom) IR camera counts implying surface emissivity in figure 41. The locations of the slices are indicated by the red and black lines in each of those figures. The data is averaged over seven rows of pixels to reduce the noise in the data. Because figure 41 has such low resolution, the seven pixels represent a noticeable width in the image, which is why two black and red lines are visible, while there appears to be only one of each in figures 39 and 40.

Chapter 6

Energy Balance on MAST Upgrade and the Super-X

One of the primary aims of this project was to assess the power balance in the super-X divertor configuration on MAST Upgrade. Once MAST Upgrade came online in the final year of the project, this became possible. An energy balance method on MAST Upgrade was established, building on the analysis on MAST including the wetted area correction (adapted as appropriate). Application of this method is the primary result of this chapter, and brings the project to its conclusion. Some similar work has been submitted for publication to Nuclear Materials and Energy, with Jack Lovell as first author, but all the analysis presented here is my own.

6.1 MAST Upgrade Design and Capabilities

MAST Upgrade (or MAST-U) is a significant upgrade to the original MAST tokamak, which was the subject of chapter 4. After MAST stopped operating in 2013, the upgrade programme commenced and MAST Upgrade began operations in 2021. Given that the NSTX Upgrade has been delayed, MAST Upgrade is presently the only operational spherical tokamak in the world, and will provide a unique insight into the properties of low aspect ratio tokamaks. It also has an interesting divertor which is described below. All the results presented here are taken from this first campaign, as there has been only one campaign to date, although there are many more planned for years to come.

MAST Upgrade is built in the same 4 m high, 4 m wide cylindrical vacuum vessel as

MAST, and inside the same concrete bioshield. It is shown in figure 43, and the outward similarity to MAST (figure 17) is obvious. However, there are few similarities beyond its appearance. It has the same arrangement of 12 pairs of toroidal field coils, but with higher current capacity. It has many more poloidal field coils, to allow greater control over the strike point location, and to produce advanced divertor geometries such as the super-X. MAST Upgrade's operating specifications are also a significant improvement over those of the original MAST; some of these are given in table 8. Note that most capabilities were only tested to around half of their design values during the first campaign.

Table 8: The design maximum values of certain key machine parameters of the Mega Ampere Spherical Tokamak Upgrade (MAST-U). MAST-U was only tested to around half of some of these values during the first campaign.

Parameter	**Value**
Plasma current	2 MA
Magnetic field	0.8 T
Pulse length	2000 ms
Plasma volume	8 m^3
NBI power	6 MW

Many of the design decisions regarding the divertor of MAST Upgrade are explained in [85], including several benchmarks against simulations. The geometry of the MAST Upgrade

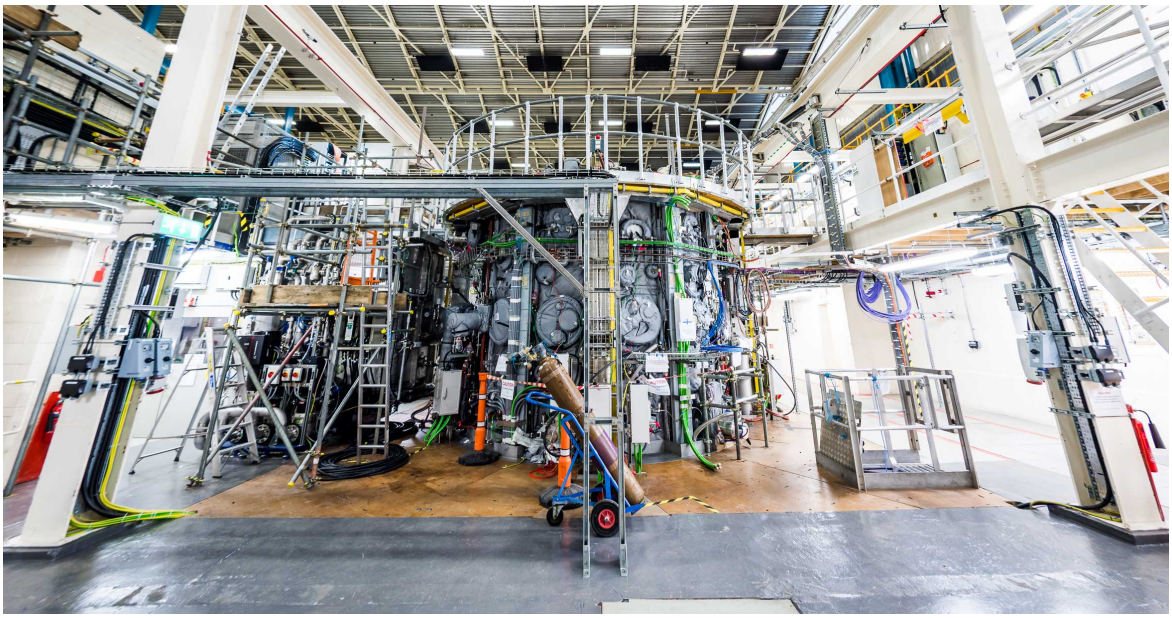

Figure 43: A photograph of the outside of MAST Upgrade, taken in 2019 during the commissioning phase. Figure credit CCFE.

divertor is very unusual, including a large divertor baffle and numerous poloidal field coils, and this design brings three advantages. Firstly, maximising the possible connection length, by increasing SOL thickness as well as poloidal strike point location. Secondly, maximising divertor closure, maximising divertor radiation and enhancing detachment. Thirdly, the coil arrangement allows much greater control of divertor configuration [85]. The enhanced detachment, especially in the super-X configuration, is one of the primary objectives of the MAST Upgrade programme.

6.2 MAST Upgrade Diagnostics

6.2.1 Infrared Cameras on MAST-U

The infrared thermography setup on MAST Upgrade, when fully operational, will have a total of five infrared cameras (three medium wave and two long wave). During the first campaign, only one camera operated reliably throughout. This was a long wave camera that observed the lower T2-T5 tiles (the names of the tiles are shown in figure 46). For some shots, there is also data from a medium wave camera that observed the lower T1 and T2 tiles. The medium wave camera (labelled AIR in figure 44) is a FLIR SC7500 camera, filtered to a wavelength range of $4.1 - 5.0\,\mu m$.

The long wave camera (labelled AIT in figure 44) is an IRCAM Velox 81k L, with a wavelength range of $7.7 - 9.4\,\mu m$. It has a 320×256 pixel resolution and a 487 Hz framerate at full frame. It can record at up to 15 kHz with a sufficiently small frame. In the standard setup with the 25 mm lens the spatial resolution is approximately $2 - 7$ mm. It views the T2-T4 and inner T5 tiles of the lower divertor.

The mounting positions and viewing angles of the AIR and AIT cameras are shown in figure 44. In the full setup, both of these will have an identical counterpart in the upper divertor, and the original medium wave camera from MAST, now labelled AIS, will be used in conjunction with a mirror array to provide a simultaneous view of the (upper and lower) inner strike points.

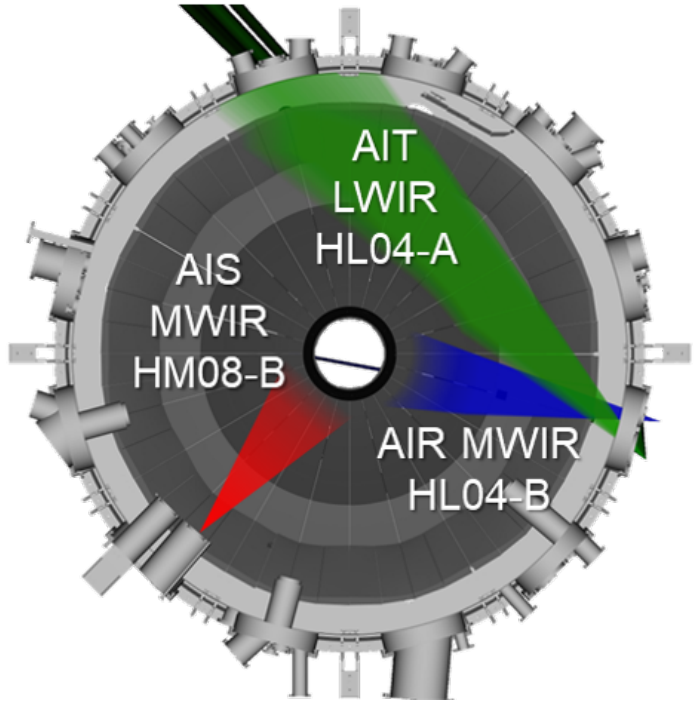

Figure 44: Plan view diagram of the MAST-U lower divertor, showing mounting locations and viewing angles of the AIR, AIT and AIS cameras. Each camera is labelled as either medium or long wave infrared, and by the diagnostic viewing port at which it is installed. Figure credit CCFE.

6.2.2 Bolometers on MAST-U

MAST Upgrade has the bolometer array from MAST, albeit moved radially inward somewhat, giving a slightly different viewing geometry. There is a poloidal fan of 16 channels, looking from top to bottom, and a tangential array of 16 channels, 4 of which are facing into the oncoming energetic particles from the NBIs. This layout is shown in figure 45. As well as this, there are two more bolometer arrays in the lower super-X divertor chamber. Each of these has 16 channels, with one array viewing vertically, and the other viewing horizontally. These are also shown in figure 45.

In addition to this, MAST Upgrade has an infrared video bolometer (IRVB) installed close to the lower X-point. The IRVB uses an infrared camera (identical to the MWIR camera mentioned in section 6.2.1) to measure the temperature of a piece of blackened platinum foil, which heats up according to the intensity of the thermal radiation. Since the thermal properties and geometry of the foil are known, it is possible to reconstruct the distribution of the thermal radiation. This technique gives pixel by pixel resolution much greater than that of a few comparatively large thermistors. It has a wide field of view, covering most of

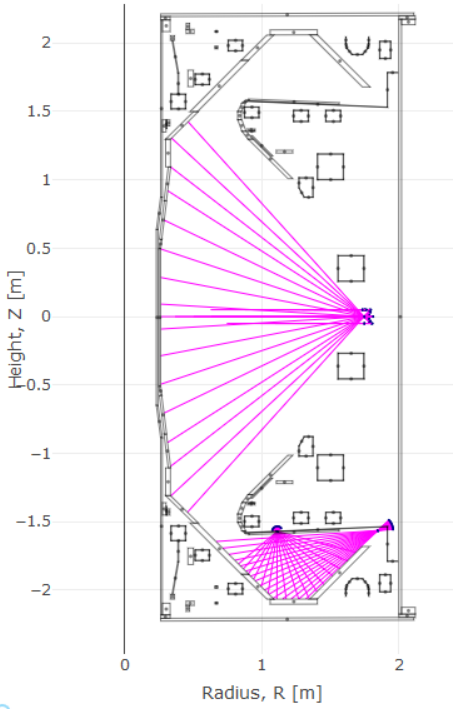

Figure 45: Diagram showing the lines of sight of all 64 resistive foil bolometers fitted to MAST Upgrade in the first campaign (IRVB not shown). Note that the 16 bolometers of the tangential array are repsesented by the single thick pink line in the middle of the image, as their viewing angles fan in to and out of the page. Figure credit CCFE.

the bottom half of the machine.

6.2.3 Non-Thermal Diagnostics on MAST-U

The core region Thomson scattering setup on MAST-U uses the same laser and timing setup as on MAST (see section 4.2.3). The lasers are positioned 15 mm above the midplane. Data obtained includes measurements of magnetic islands, the L-H transition and pellet ablation. In addition to this, MAST-U also has a Thomson scattering setup in the divertor chamber, which has a 10 mm spatial resolution with 10 spatial points.

MAST Upgrade has an array of 850 Langmuir probes covering the divertor tiles, nose tiles and centre column. Up to 640 can be used at any given time; which 640 are being used can be updated once per week [86]. Langmuir probes are held at a bias voltage relative to the vessel wall, and a current of particles from the plasma is measured passing through the probe. This gives measurement of plasma current and voltage, and can also be used to infer particle flux and/or heat flux. Given the reduced IR camera coverage on MAST-U,

compared to the vast array of Langmuir probes, this method of inferring heat flux has been used as appropriate in the results presented here.

MAST Upgrade has 398 pickup coils and 102 flux loops, which are used for the equilibrium reconstruction, including the shape and position of the plasma. This data is used to calculate some quantities used in power balance, such as plasma stored energy. There are numerous other coils and loops on MAST-U, which are used mostly for magnetohydrodynamic (MHD) studies. This includes 84 Rogowski coils (used for real-time plasma current measurement and machine protection), 91 saddle coils (mostly used for error field measurements), 34 Mirnov coils (used for MHD studies), 8 diamagnetic loops, and 36 halo current detectors. These all provide data that is valuable for determining viable plasma scenarios, both for MAST-U itself and for future reactors.

6.3 Wetted Area Correction

6.3.1 Implementation

The analysis of MAST-U data began with the wetted area correction. The divertor tiles in MAST-U, as in MAST, are not horizontal in the toroidal direction. They are inclined a few degrees into the face of oncoming plasma. This is done to prevent plasma striking the edge of the tile, which is exposed by the gap between the tiles (see section 4.3.2, especially figure 21).

The MAST Upgrade divertor was designed such that each tile maintains a constant step height, relative to the next tile, across its radial length. This differs from MAST, where the tile angle was constant and the step size varied. It can be seen in figure 46 that some of the tiles on MAST Upgrade also have non-zero poloidal inclination, unlike on MAST. The corresponding changes to the wetted area calculation are given below.

As before, the wetted fraction, f_{wet}, is given by equation 15, which is reproduced here for convenience

$$f_{\text{wet}} = \frac{A_{\text{wet}}}{A_{\text{tot}}} = \frac{\sin\theta}{\sin(\theta+\alpha)} \tag{15}$$

where θ is the angle of incidence of the plasma relative to the plane, and α is the angle of inclination of the tiles relative to the same plane. Assuming this is the horizontal plane, e.g. for T4 tiles, θ is given by the same formula as for MAST

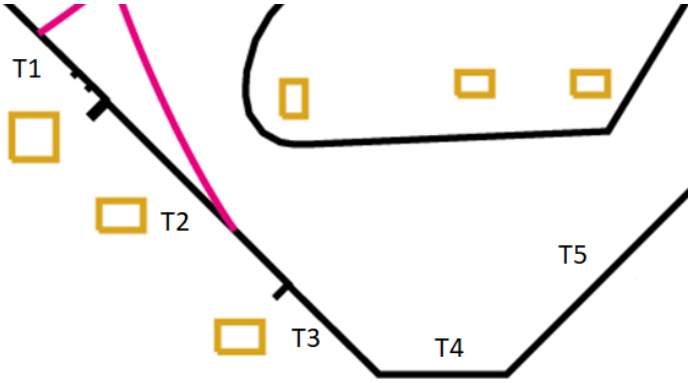

Figure 46: Cross-sectional diagram of part of the MAST-U lower divertor, showing the T tile locations, and their inclinations in the $R - z$ plane. Brown squares are poloidal field coils, and the pink line represents an example separatrix. Adapted from a figure supplied by Andrew Thornton of CCFE.

$$\sin\theta = \frac{-B_z}{\sqrt{B_R^2 + B_z^2 + B_\phi^2}} = \frac{-B_z}{|\mathbf{B}|} \tag{16}$$

where $\mathbf{B}$ is the magnetic field, and the subscripts R, z and ϕ denote the radial, vertical and toroidal components respectively. Equation 16 is from [71], and can be derived from simple trigonometry. Now for other tiles such as T3 or T5, let the angle by which these deviate from the horizontal, in the radial direction, be ψ. It can be seen in figure 46 that this is 45° for T1, T2 and T3, and -45° for T5. Then we have

$$\sin\theta = \frac{-B_\perp}{|\mathbf{B}|} \tag{19}$$

where $B_\perp$ is the component of the magnetic field that is normal to the tile surface (at an angle ψ from the vertical), which is calculated from

$$B_\perp = B_z \cos\psi + B_R \sin\psi \tag{20}$$

In the case of $\psi = 0$, this reduces to $B_\perp = B_z$ as required. We now require a value for α, which is not a constant for MAST-U as it was for MAST. It can be derived from the known tile step size, s, and number of tiles around the circumference, n

$$\tan(\alpha) = \frac{sn}{2\pi R} \tag{21}$$

where R is the major radius coordinate, and s and n are 3 mm and 12 for T1-T4, and 1.4 mm and 24 for T5. By substituting (20) into (19), and both of those and (21) into (15), f_{wet} can be calculated in terms of known quantities:

$$f_{\text{wet}} = \frac{\sin\theta}{\sin\left(\arcsin\left(\frac{-B_z\cos\psi - B_R\sin\psi}{|\mathbf{B}|}\right) + \arctan\left(\frac{sn}{2\pi R}\right)\right)} \tag{22}$$

A function has been written to calculate the wetted fraction for every pixel of the camera at every EFIT timestep (5 ms). Although this is lower than the frame rate of the camera, the calculation depends on EFIT for magnetic field vectors, so nothing would be gained from interpolating the data to a higher temporal resolution. The validity and accuracy of this calculation was tested against real MAST-U data for several shots.

6.3.2 Validation

The non-wetted area of a MAST-U divertor tile can be seen in images as the black strips which appear to break up the strike point toroidally, such as in figure 47. This was also visible on MAST, such as the dark strips fanning out from the centre of figure 34 towards the bottom. By measuring the width of these strips, and comparing to the width of the adjacent wetted area, the wetted fraction can be calculated from experimental data, which can then be checked against the calculation based on tile and field line angles (equation 22) to ensure it is correct. This method was employed on several MAST-U shots before the code containing the calculation was submitted to the MAST-U team. This code had not been implemented by the end of the first campaign, so appears as a separate step in this analysis.

For this validation to proceed, suitable shots first had to be chosen, which had high enough heat flux for the wetted area to be clearly visible in the infrared image. These were from the scenario development phase, because the the plan had been to implement the analysis before the start of the experimental phase. One of the shots used for this was 43962, because it had a strike point sweep, so the calculation could be re-done several times during the shot, at several points on the divertor. Another was 44677 (figure 47) because it was a high plasma current (750 kA) L-mode shot, so there was reliably high image contrast throughout the shot.

Briefly, the method was to measure the width of the shaded area visible in images like figure 47, then divide by the sum of its width and the adjacent wetted area width, to

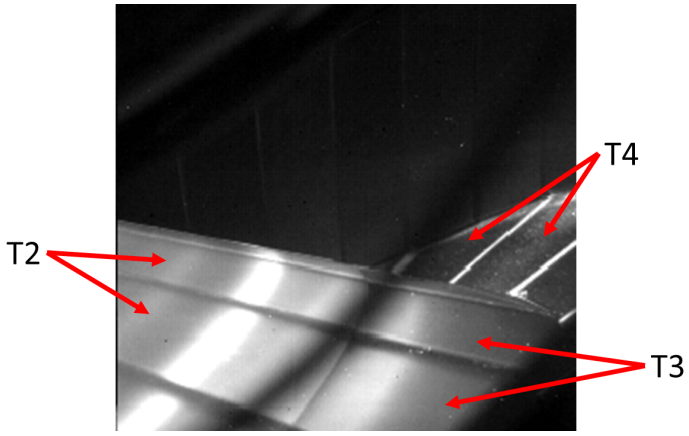

Figure 47: A single camera frame from the LWIR camera during MAST-U shot 44677, with T2, T3 and T4 tiles clearly visible, and T5 tiles faintly visible in the upper half of the image (T tile locations were shown in cross-sectional view in figure 46). The view is partially obscured by a loose cable that was present during the first MAST-U campaign. The dark shaded strips at the edge of each tile are clearly visible, separating the lighter areas, which are the wetted areas of each tile. The strike point can be seen towards the bottom of the T2 tiles.

approximately calculate the wetted fraction. Wetted fraction is dimensionless (one width divided by another), so no absolute calibration was necessary. These measurements of the wetted fraction agreed with the fractions calculated by the method given in section 6.3.1 to within the effective resolution of the cameras. This confirms that the method has successfully been adapted from MAST to MAST-U, and can be used in any future analysis on the tokamak.

6.4 Time Windows and Strike Points

6.4.1 Choice of Time Windows

At the time of writing, the only MAST Upgrade data available is from the first campaign. This means there is less data than would be available for a more mature machine, and the data is of lower quality (by comparison, the MAST data presented in chapter 4 is from the eighth and ninth campaigns on that machine). Of the five infrared cameras intended for MAST-U, only one was routinely available during the first campaign. Bolometer data was not available in the upper divertor chamber, and Langmuir probe coverage was also limited. The approach taken was therefore to look in great detail at a single shot, which had long periods of steady state plasma that would allow for time averaging. The shot chosen for this

was 45209, because the strike point was swept out from a conventional divertor to a super-X during the shot.

In a given tokamak pulse, there may be an asymmetry between the top and bottom halves of the tokamak, usually caused by the plasma being slightly above or below the midplane. In these asymmetric cases, there are actually two separatrices, instead of the idealised double-nulled separatrix such as that shown on the right of figure 18. The most common numerical measurement for this asymmetry is the diference in radial coordinate between the two separatrices at the outer midplane, denoted δr_{sep}. This is the x-axis quantity in figure 8 (a). This will be referenced in discussion of up/down asymmetry in the remainder of this chapter. There is also asymmetry between inner and outer strike points, the outer strike point usually receives many times more flux than the inner. In/out has no obvious single parameter to quantify it, but is also discussed in the remainder of this chapter.

For this analysis, a shot was desired that included a super-X divertor (as described in section 1.2.3) to fulfil the original aims of the thesis. A steady state period with a conventional divertor during the same shot was also desired, to ensure consistency of other plasma conditions between conventional and super-X. This was considered preferable to using different shots which may have had different plasma conditions, especially because of the poor repeatability of shots during the first MAST-U campaign. The shot chosen was 45209, because the conventional divertor was held for long enough at the beginning of the shot to obtain steady state data, contrary to the normal procedure of sweeping out to super-X straight away. Although the sweeping out occurred later than usual in this shot, it proceeded by the normal method of changing the magnetic field in the divertor by varying the currents in the divertor coils (which are shown as brown boxes in the cross-sectional view in figure 46).

Because the IR data is only available at the lower outer strike point on MAST-U, it was decided to utilise the LP data initially for the power balance, to give both upper and lower strike point coverage. To reduce noise in the data, the data was averaged across a chosen time interval of 50 ms. Two time intervals were required to perform a conventional/super-X comparison (one for each). The chosen early and late time intervals for 45209 are shown in figure 48, along with a middle time interval explained below. It can be seen that all the key plasma parameters were approximately constant throughout the intervals, with neither steady changes nor transients. The ohmic power, neutral beam power and stored energy are

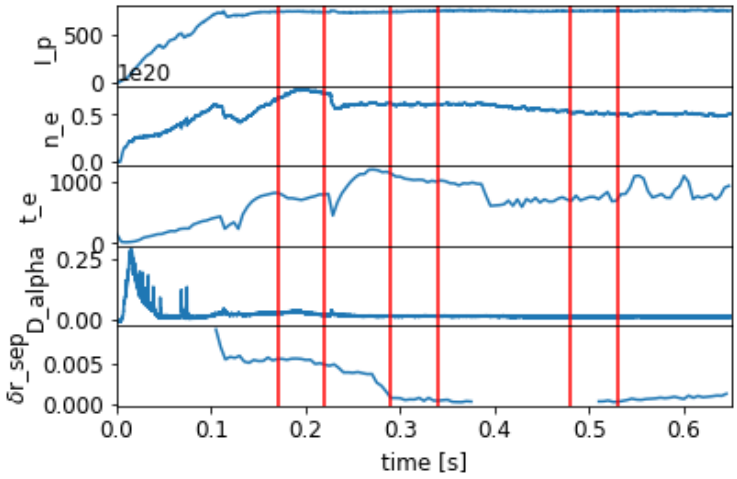

Figure 48: A general overview of the variation of several key plasma parameters during MAST-U shot 45209. The red lines are the beginning and end of the early, middle and late time intervals, each 50 ms in length. From the top, the parameters are plasma current [kA], electron density [$10^{20}\,\mathrm{m}^{-3}$], electron temperature [eV], D_α spectroscopy and δr_{sep} [m].

read out of the database with no further analysis, as they were for MAST in chapter 4, but the divertor flux data now requires additional processing.

One thing that is not shown in figure 48 is the strike point location, as this can be difficult to determine with precision, especially when there is strike point splitting. The early and late time intervals were both chosen at times when the strike point was stationary, because of the desire for steady state in every aspect. However, the δr_{sep} value changed between the two, from 5 mm to 0 (see the bottom trace in figure 48). A middle time interval was chosen with the same δr_{sep} as the late interval, but a conventional divertor. The up/down asymmetry should be the same in both middle and late intervals. Based on old data from MAST reproduced in this thesis as figure 8 (a), slightly more power to the upper divertor than the lower may be expected at $\delta r_{\mathrm{sep}} = 0$. Although the strike point was moving during the middle interval, it remained in the T2-T3 tile region, so is still considered a conventional divertor.

6.4.2 Divertor Flux Profiles

On the original MAST, the divertor power was measured at all four strike points (lower outer, lower inner, upper outer, upper inner) by the infrared (IR) cameras. Langmuir probe (LP) data was also available on MAST, but IR is generally preferred for power balance where

available. On MAST Upgrade, the IR camera data for the first campaign is available only at the lower outer strike point. LP data is available at the upper and lower outer strike points. There is no data for the inner strike points, which is a difficulty for machine-wide power balance analysis and is discussed further below. The data from both LP and IR was analysed during the early, middle and late time intervals of shot 45209 (which were shown in figure 48).

For the LP data, having averaged over the 50 ms time window, an Eich function was then fitted to time-averaged heat flux against radius. The Eich function is defined in [87] and takes the form:

$$q(\bar{s}) = \frac{q_0}{2} \cdot \exp\left[\left(\frac{S}{2\lambda_q f_x}\right)^2 - \frac{\bar{s}}{\lambda_q f_x}\right] \cdot \mathrm{erfc}\left(\frac{S}{2\lambda_q f_x} - \frac{\bar{s}}{S}\right) + q_{\mathrm{BG}} \tag{23}$$

where q is the heat flux, q_0 is the amplitude of the heat flux, $\bar{s}$ is the divertor coordinate, erfc is the complementary error function, S is the width of a Gaussian function, λ_q is the power decay length, f_x is the magnetic flux expansion, and q_{BG} is the background heat flux.

This function was then used as the basis for the integral in calculating the total divertor energy deposited during the interval. The integration first proceeds by multiplying the data by $2\pi R$ to give the annular rings. Then an integration of the Eich function over the radius is performed, multiplied by $\sin 45^\circ = \sqrt{2}$ to account for the actual integration distance being longer than that suggested by ΔR, because of the inclination of the T2, T3 and T5 tiles (the only tiles considered in the analysis of this shot).

The IR data was also used in this analysis, as it has higher spatial resolution than the LPs, and better signal to noise ratio. As with all the other data, the IR data was averaged over 50 ms to produce profiles. This was then spatially integrated in the same fashion as the LP data above. The integrated energy data from the IR could then be compared to that from the LPs.

The LP divertor flux profiles for the early time interval of shot 45209 are shown in figure 49. The raw data is shown in blue, with the time-averaged data in orange. The black line is a fitted Eich function, which is the standard function for tokamak divertor strike point profiles and was defined in [20]. For the upper divertor, the black line is fairly well fitted to the orange data, and is within error for most of the points, despite the random noise in the blue data. For the lower divertor however, there are several orange data points above the

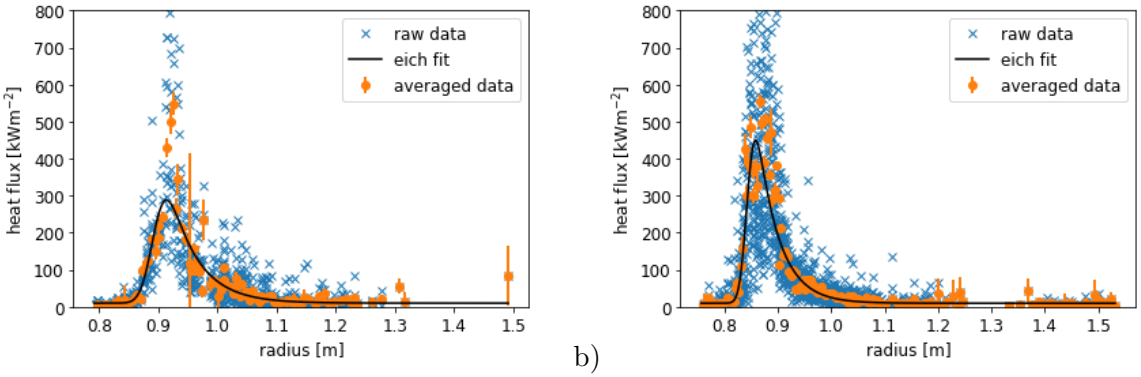

Figure 49: Langmuir probe profiles of divertor heat flux at the early time (between 170 and 220 ms) in shot 45209, for both (a) lower and (b) upper divertors. The blue data points are single probe measurements, and the orange data points are time-averaged measurement for each probe over the 50 ms, with error estimated from the standard deviation of the blue points. The black line is an Eich function, fitted to the time-averaged data points accounting for error.

outboard side of the Eich profile, outside of error. The integrated energy values (integrated around the torus, then under the curve, then over the time window) are 15.1 kJ for upper and 13.6 kJ for lower.

Figure 50 shows an IR heat flux profile for the lower divertor only. Again, the blue points are raw data and the orange points are time-averaged, with standard deviation. The black line is an Eich fitting function. By comparison with figure 49 (a) the similarity of the lower strike point IR and LP data can be seen, although the IR does have a somewhat lower peak. The integrated energy value is 13.0 kJ uncorrected, or 8.7 kJ with wetted area correction applied as described in section 6.3. This means, although the profiles look similar, there is a significant discrepancy once wetted area correction is taken into account.

At the later time interval in figure 48, when the strike point was swept out to the Super-X divertor, the lower LP data became unreliable, because of the poor LP coverage of the lower T5 tiles. However, the combination of upper LP and lower IR data allows an energy balance to account for both upper and lower divertors. Interestingly, some strike point splitting was observed when the strike point was swept out. The split was clearly visible in both the IR and LP data, as shown in figure 51. In order to account for the strike point splitting, a double Eich function was fitted to the data at this time. For the LP, the lower spatial resolution in figure 51 (b) relative to 49 (b) is because the Langmuir probes are not as closely spaced in the T5 tiles as in T2 and T3. For the IR, the lower spatial resolution of figure 51 (a) relative to figure 50 is because the T5 tiles are further from the camera than the T2 and

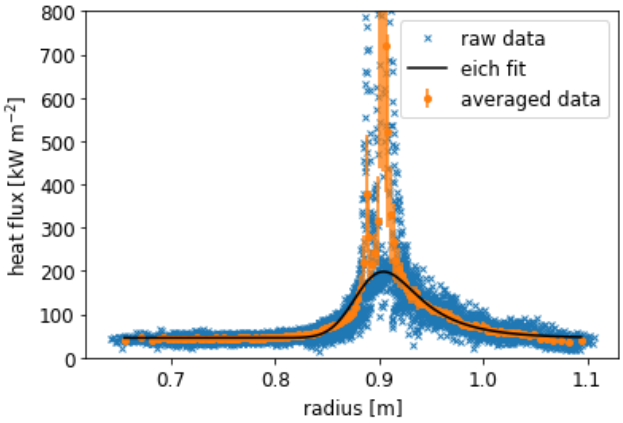

Figure 50: Infrared divertor heat flux profile at the early time (between 170 and 220 ms) in shot 45209, for the lower divertor. The blue data points are single pixels measurements along the 1D analysis path, with radial coordinate adjusted to account for strike point sweep during the interval. The orange data points are time-averaged measurement for each radial coordinate over the 50 ms, with error estimated from the standard deviation of the blue points. The black line is a standard Eich fitting function.

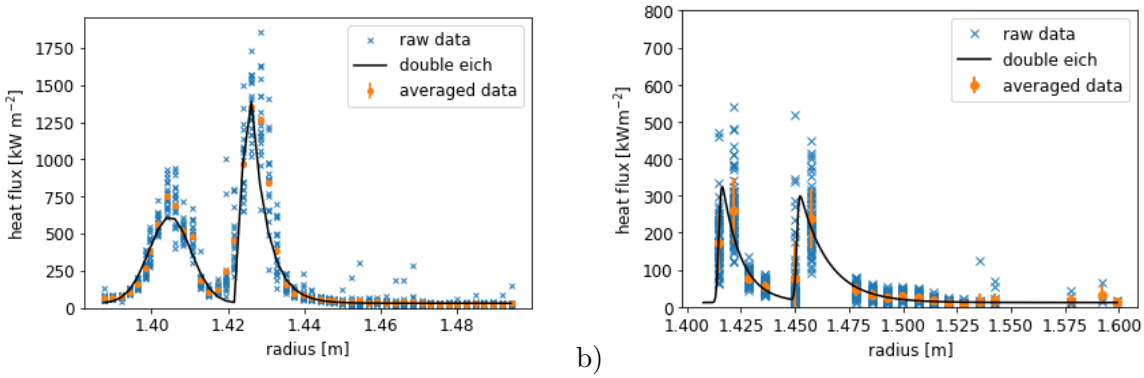

Figure 51: Divertor heat flux profiles at the late time (between 480 and 530 ms) in shot 45209, both (a) lower IR and (b) upper LP. In both, the blue data points are single measurements, and the orange data points are time-averaged over the 50 ms, with error estimated from the standard deviation of the blue points. The black lines are the sum of two Eich functions, fitted to the time-averaged data points accounting for error.

T3 tiles. The small radial offset between the IR and LP data is thought to be caused by toroidal asymmetry between the locations of the LP and IR data, which will be discussed in more detail in section 6.5.2.

6.5 Analysis of Energy Balance and Asymmetries

6.5.1 Method for Energy Balance

Working with the diagnostics available for the first campaign, it was important to ensure consistency with the prior work on MAST. As mentioned in section 4.3.1, there are a total of six quantities that must be considered: ohmic heating power, neutral beam heating power, stored energy, magnetic energy, radiated power and divertor power. However, because of a change in the way ohmic heating power is calculated on MAST Upgrade relative to MAST, the magnetic energy no longer needs to be considered as a separate quantity. It is instead incorporated into the quoted amount of ohmic heating power.

The radiated power measurements on MAST Upgrade include measurements of the lower divertor chamber as well as the core chamber, so the measurements of radiated power are expected to be more accurate overall, although there was no coverage of the upper divertor chamber. (Note that the original MAST did not have dedicated divertor chambers, and radiation in the divertor region was not significant because discharges were largely attached, although some imaging at the divertor was performed).

The ohmic heating power is calculated from the plasma current and loop voltage; as it was on MAST. Plasma stored energy is also calculated the same way as for MAST. Neutral beam power must now be read separately from each beam, and is no longer smoothed, but is otherwise calculated in the same way. However, there is no neutral beam data for shot 45209, because it was purely ohmically heated. All of these quantities were then compared with the IR and LP data from section 6.4.2 for a full energy balance analysis.

6.5.2 Results and Asymmetries

As stated in section 6.4.2, there was no inner strike point coverage available in the first MAST-U campaign from either infrared cameras or Langmuir probes. Both will be made available for the second campaign, which will begin later in 2022. The inner divertor heat flux is too significant to be discounted entirely (based on other tokamaks including MAST [30]), so its magnitude must be inferred some other way. Based on past published results from MAST, the inner strike point heat flux is usually around 10 - 20% of the total divertor heat flux (or 25% of the outer divertor heat flux) for double null or upper single null shots, and around 30 - 40% for lower single null (see figure 8 (b) which is from [30, fig.7 (b)]). The

MAST analysis presented in chapter 4 suggests around 6 - 10% of total heat flux for double null or upper single null shots, and around 18-23% for lower single null. In both this and de Temmerman's work, the reason for a larger inner strike point fraction in lower null shots may be because there were a larger number of shots with large negative δr_{sep} than with large positive (i.e. the lower single null shots in the database were more strongly single null than the upper single null shots). This is indicated by the cluster of data point in the top left of figure 8 (a).

Since 45209 was a double null shot, an appropriate assumption based on MAST would be that 10% of total divertor heat flux is incident at the inner divertor. Figure 52 shows a full energy balance for the early time interval of shot 45209. It can be seen that the overall energy accounted for in this interval is 66%. The upper LP energy value is 15.2 kJ while the lower is 13.6 kJ. This gives an up/down asymmetry of 1.12. Taking the proposed estimated inner divertor heat flux (10% of total divertor flux) and applying it to the analysis gives overall balance of 71%, which is an improvement, more in line with results from MAST (see figure 23) and other tokamaks (see table 1).

The up/down asymmetry value of 1.12 derived from the LP data above can be used to infer the hypothetical upper IR energy value from the lower IR energy value from figure 50. However, the wetted area correction described in section 6.3.1 must be applied first. Applying this to the lower IR heat flux from the early interval of shot 45209 (which is

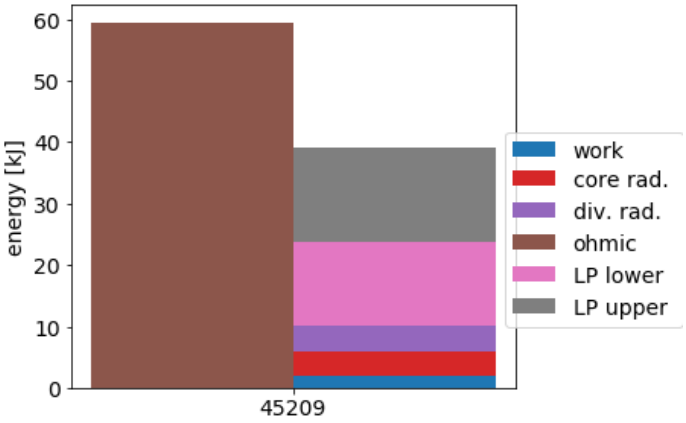

Figure 52: An energy balance of all available quantities during the early time interval (170 to 220 ms) of shot 45209. Core rad is core radiation and div rad is divertor radiation. The overall energy balance is 66%, or 71% including estimated inner divertor heat flux.

13.0 kJ) reduces the IR lower strike point energy value to 8.7 kJ. This is a much larger correction than expected based on the validation described in section 6.3.2, and worsens the overall energy balance further, as the IR value was already lower than that from the LPs.

Unfortunately, this asymmetry value can not be applied at the later time, because as shown in figure 48 (bottom) the δr_{sep} changes between the early and late time intervals. Therefore, a middle time interval was chosen when the δr_{sep} was already zero, but the strike point had not yet swept onto T4. (This middle time was undesirable because the strike point moves significantly during the interval, but it only moves within T2 and T3). The total power accounting at the middle time is 52%, or 57% assuming 10% inner strike point heat flux.

Performing an up/down energy balance on the LP data at this middle time interval (which is 8.7 kJ for upper and 9.9 kJ for lower) gives an up/down asymmetry of 0.88. Since the corrected lower IR for the middle time gives 8.7 kJ (coincidentally the same as both the early time and upper LP), and δr_{sep} remains zero between the middle and late times, the up/down asymmetry of 0.88 above can be applied, and gives the result that the upper IR should be 9.9 kJ. Clearly this does not improve the overall energy balance.

At the later time interval, once the divertor has swept out into the Super-X, there is no usable lower LP coverage, because of faulty LP hardware in the lower T5 tiles during the first campaign. However, the combination of upper LP and lower IR can be used to produce an energy balance with data from both the upper and lower outer strike points. The discrepancy between the lower IR and lower LP at the middle time interval is 0.88, so this will be assumed for the late interval as well.

In figure 53, an energy balance in a super-X divertor late in shot 45209 is given, satisfying one of the primary aims of this project. The overall accounting is 81% (or 88% assuming 10% inner strike point heat flux). The up/down asymmetry appears to be very high, which is unexpected given that δr_{sep} is approximately zero (see bottom of figure 48). This may be because of some toroidal asymmetry. Although perfect toroidal symmetry is usually assumed in both experiments and simulations of tokamaks, it may not always be the case. If the plasma is disrupted by locked modes, there can be significant toroidal variation in strike point heat flux, and strike point splitting is a symptom of this [88].

Given that the strike point has been shown to be split at this time in the shot, and that these modes are a by-product of running at low plasma density to obtain attached

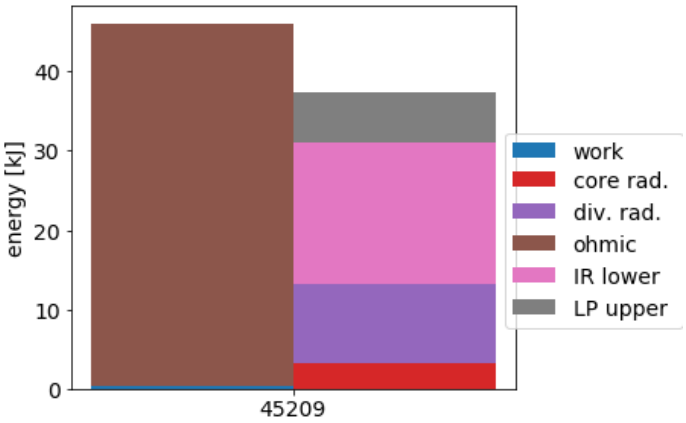

Figure 53: An energy balance of all available quantities during the late time interval (480 to 530 ms) of shot 45209, produced by combining upper LP data with lower IR. Core rad is core radiation and div rad is divertor radiation. The overall energy balance is 81%, or 88% including estimated inner divertor heat flux.

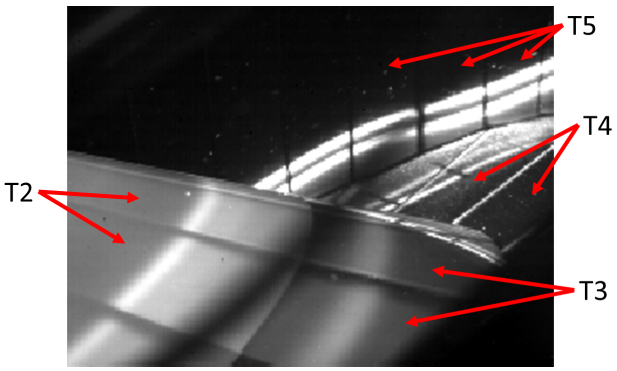

Figure 54: A single camera frame from the LWIR camera at 505 ms during MAST-U shot 45209, half way through the late time interval. The T2, T3, T4 and T5 tiles are clearly visible, although the image is partially obscured as in figure 47. The toroidally asymmetric split strike point can be seen towards the bottom of the T5 tiles.

super-X divertors, a toroidal asymmetry between the different toroidal locations of the IR cameras and LPs may be expected, which could explain this apparent disagreement. This is supported by visual inspection of the raw IR camera data shown in figure 54, in which both strike point splitting and toroidal asymmetry are clearly visible on the T5 tiles.

On the other hand, given that the profile peak values in figure 51 are more different than expected from those in figure 49, while the profile shapes are relatively similar, another explanation is possible. The peak height of the LP profiles matches well with both LP and IR heat flux peaks at earlier times in the shot, whereas the IR is much higher at this later

time. This may suggest another issue with the IR data in particular, such as an incorrect emissivity value used in the analysis. However, the emissivity used in the calculation on T5 was 0.78; already a high value for graphite. The graphite would be reflective if the emissivity were any higher, so this explanation is unlikely. Therefore, the toroidal asymmetry caused by the locked modes mentioned above is thought to be the sole cause of the discrepancy.

As mentioned in section 1.2.3, one of the aims of the super-X divertor concept is to increase divertor radiation, enabled by the larger interaction volume of the super-X divertor. This results in decreased divertor tile heat flux, as mentioned in section 1.1.3. By comparing figures 52 and 53, it can be seen that the fraction of heating power that is radiated in the divertor increased significantly in the super-X divertor at the late time, relative to the conventional divertor at the early time. In fact, the divertor radiation changes from only 1.1 times larger than core radiation to 3.1 times larger, despite only a small decrease in core radiation. Considering only the upper LP data, this appears to have resulted in a decrease in divertor tile heat flux, as expected. This validates one of the primary aims of the super-X divertor, to induce radiation in the divertor chambers thereby lowering divertor heat flux, without increasing core radiation. This bodes well for the use of super-X divertors in future fusion reactors.

Unfortunately, the reduction in divertor heat flux is not apparent when the lower IR data is considered. This observation cannot be explained entirely by toroidal asymmetry, and the assumed emissivity value is already high as noted above. The unexpectedly high IR-measured divertor heat flux in the super-X divertor will be the subject of future work.

Chapter 7

Conclusions and Recommendations

The primary aims of this thesis were set out in section 2.5. Having set out the techniques which would be used to meet these aims in chapter 3, the analysis methods and results were set out in the ensuing chapters. It was noted that past studies of power balance on MAST failed to correct for the wetted fraction of the tiles, caused by inclination and resulting shadowing. Studies on various tokamaks have also failed to account for neutral beam absorption of less than 100%, caused by a combination of shinethrough and first orbit loss. Both of these limitations have been addressed in chapter 4 of this thesis.

There was also the question of the accuracy of infrared camera measurements of tile surface temperature given the effects of surface layers and emissivity. This effect was known on MAST, and attempts were made to compensate for it. An investigation to directly measure and quantify these effect was carried out and presented in chapter 5. Finally, there is the question of the super-X divertor. One of the primary aims of MAST Upgrade was to enable full testing of the super-X divertor, and this was achieved in the first campaign in early 2022. An assessment of a super-X divertor was presented in chapter 6.

7.1 Conclusions of Improved Power Balance Analysis

In the past, tokamak power balance has neglected the effects of wetted area, and of neutral beam absorption, such as in [30]. The work presented in this thesis builds upon this past work, accounting for these additional factors as a natural continuation of work in the field. As described in section 4.4.2, when the correction to account for the wetted fraction of the tile area was implemented, this caused the energy balance data from the chosen shots on

MAST to decrease further below the ideal 100%. Although the correction to outer divertor flux can be up to 30% on a typical shot, the percentage of the total power is typically $< 10\%$. The resulting energy accounting, for ohmic shots only, is in line with past power balance experiments, which provides a validation of the power balance method used in this thesis. The only remaining unexplained discrepancy results from neutral beam power.

Section 4.4.3 shows that the energy accounting, which is the percentage of energy recorded going in to the plasma that is also recorded coming out, is lower for shots with higher neutral beam power. Neutral beam power is the variable that correlates most strongly with the power accounting; even more so than plasma density. In section 4.6, especially table 7, it is shown that TRANSP simulations can account for some, but not all of this discrepancy. This explains the greater discrepancy of single beam shots relative to ohmic shots, validating the use of TRANSP for this purpose. It offers no insight into the even greater discrepancy of double beam shots. This prevents a simple absorption correction, scaling with beam power, from being implemented as originally intended, because the number of beams used has much more impact on the absorption than the beam power in MW.

The reason that beam power absorption is much lower in double beam shots is thought to be a result of fishbone instabilities driven by the large fast ion pressure gradient present in double beam shots. This pressure gradient resulted from the very high fast ion pressure caused by the beams both firing into the same portion of the tokamak. The author is not aware of any quantification of this effect by theory or simulation, although this thesis gives an empirical estimate. This effect clearly must be accounted for in any future work with data from the original MAST. Data from MAST Upgrade will not be affected by this problem, because all four beams (two existing and two proposed) fire into different portions of the tokamak.

The accounting for beam power absorption, and the realisation that double beam shots may have additional power losses due to fishbone instabilities, brings the power accounting from beam-heated shots back into line with that of ohmic shots. However, there still remains a discrepancy in ohmically heated shots. This discrepancy was also present in other past experimental tokamak power balance studies (see table 1), so the method used here incorporating wetted area correction is consistent with those past studies. This remaining discrepancy is thought to have been caused by a combination of radiated power in the divertor region, and first wall tile flux in the core region.

7.2 Conclusions from Tile Surface Evolution Experiments

Surface temperature measurements taken using infrared cameras can be unreliable because of emissivity correction and surface layer effects mentioned in section 3.1. It was desired to collect data from predictable circumstances which could allow better calibration for this effect to be implemented in future. The results presented in chapter 5 show that exposing a graphite sample to a high power plasma results in erosion of the surface. This is because the higher power plasma is composed of particles with energy above the threshold of chemical sputtering of graphite. This erosion resulted in a decrease in emissivity in the eroded region at the centre of the tile. This result was contrary to expectations, which were informed by past experience with lower power plasmas.

Lower power plasmas, with particles below the chemical sputtering particle energy threshold, are deposition dominated (that is, more matter is deposited than eroded). This causes an increase in emissivity, which had been observed during the lifetime of MAST, and was observed around the edges of the tile in the steady state experiment in chapter 5. The experiment also showed the emissivity continuing to decrease as more material was ablated from the surface over time, as shown by the apparent 70 °C temperature drop in figure 32 and the apparent 140 °C temperature drop in figure 33. This demonstrates in principle that plasma exposure changes the tile surface properties over time, and that this can lead to significant measurement errors.

In addition to the wider surface erosion visible in the tile surface profile (figure 39), there was significant pockmarking on the tile surface after the experiment. The tile appeared significantly more heavily damaged than those removed from MAST after more than a decade of use. This pockmarking is visible in figures 37 and 39. Regrettably, because of problems with the tile used in the ELMy experiment, a similar post-mortem analysis was not carried out on that tile. Curiously, visible inspection reveals less pockmarking on the tile used in the ELMy experiments, but the reason for this is unknown.

It has been stated previously that the original intention of the work on Magnum-PSI was to observe the tile surface increasing in emissivity, and use this information to inform analysis of experimental data from MAST Upgrade. This proved impossible, because the tile actually decreased in emissivity. However, the fact that plasmas above a certain particle energy threshold cause chemical sputtering leading to emissivity decreasing rather than increasing is an interesting result in its own right, and has already resulted in a published peer-reviewed

paper [75].

7.3 Conclusions from MAST Upgrade and the Super-X Divertor

As on MAST, a power balance analysis was desired for MAST-U, and like MAST, MAST-U has inclined divertor tiles with only a certain fraction actually receiving plasma flux (the wetted fraction). In validation against preliminary experiments, the wetted fraction was found to be larger (i.e. smaller correction) than on MAST, at around 0.9 rather than 0.7. Applying the wetted area correction to data from MAST-U worsens the agreement with the idealised case, just as it did for MAST.

Measurements of divertor tile heat flux on MAST-U, taken by both infrared cameras (IR) and Langmuir probes (LPs), show the expected Eich profile shape, in both conventional and super-X divertors. In the super-X divertor, strike-point splitting was observed. The integrated energy values under these curves are within expectations, with the exception of the T5 IR profile at the late time. At the early time, which had a non-zero δr_{sep}, there was some up/down asymmetry as expected. This vanished at the middle time once δr_{sep} was zero. In both cases however, there was significant energy balance discrepancy; more than expected from MAST or other tokamaks. This is because of the lack of inner strike point coverage in the first campaign on MAST-U; something that will be rectified for the second campaign. Estimating the value of the inner divertor flux from typical past results on MAST brings the overall energy discrepancy back into line with expectations.

As mentioned above, the IR measurement of the heat flux during the super-X (at the late time) was significantly higher than expected. The reason for this is not certain, but is probably because of the toroidal asymmetry caused by the locked modes which lead to the strike point splitting. This resulted in a lower strike point IR value three times that of the upper strike point LP value. This was with zero δr_{sep}, and despite good agreement in the conventional divertor with no strike point splitting. The compound analysis including lower IR and upper LP gives good energy accounting of 88% overall.

Assuming that the LP data was more reliable than the IR data at the later time, because it is more consistent with values from earlier in the shot, the divertor heat flux decreased in the super-X configuration as expected. There was also a marked increase in divertor

radiated power in the super-X relative to the conventional divertor, increasing from just 1.1 times core radiation to 3.1 times, with minimal change in core radiation. The increase in divertor radiation was expected because of the larger interaction volume of the super-X divertor, as mentioned in section 1.2.3. This result provides an experimental validation of the super-X divertor concept for the first time, on the only machine in the world capable of testing it fully. This satisfies one of the primary aims of this thesis, and of the MAST Upgrade programme in general.

7.4 Recommended Further Work

1. Experiments on a Lower Powered Linear Plasma Device

   The original aim of the work on Magnum-PSI (presented in chapter 5) was to gather data on graphite tile surface emissivity in MAST-U relevant plasma conditions that could be applied to experiments there. In fact, the plasma power on Magnum-PSI was found to be too high to effectively emulate the plasma conditions in the MAST-U divertor, (Magnum-PSI is designed to emulate the ITER divertor). Some similar experiments on a lower powered device, with plasma conditions more in line with those in the MAST-U divertor, may produce results that would inform the analysis of IR data from MAST-U and other present-day tokamaks. A longer term investigation on MAST-U itself could also achieve this. In either case, it could include improved understanding of the parameter α.

2. Additional IR Camera Coverage on MAST Upgrade

   The infrared camera system on MAST Upgrade, as described in section 6.2.1 is presently incomplete. Once it is finished, there will be five cameras (three medium wave and two long wave), giving data on the upper and inner strike points as well as the lower outer strike point. Clearly, viewing the entire divertor tile area will provide better results, especially since three out of four strike points are missing entirely at present. In particular, future power balance analysis will no longer rely on assumptions about both in-out asymmetry, and the comparability of infrared and Langmuir probe data.

3. Possible Direct Emissivity Measurement and Correction

One possibility for partly solving the problem of accounting for emissivity changes and choosing α values, is to use both medium and long wave cameras to observe the same areas of the tile simultaneously, although this would be very expensive. Having images of the same tiles at two different wavelengths allows emissivity effects to be cancelled out, similar to the procedure used in a two-wavelength pyrometer, which is described in section 3.5.1. The effects of surface layers are a significant problem in infrared thermography, so their elimination in this manner, although expensive, would be of great utility.

4. Additional Bolometry Coverage on MAST Upgrade

   In the first campaign, MAST-U only had divertor bolometry and infrared video bolometry (IRVB) in the lower half of the machine, preventing these kinds of measurements in the upper half. These will be replicated in the upper half in the second campaign, which will result in better diagnostic coverage of the machine. Such measurements (along with those from the expanded IR camera system) will be useful for assessment of up-down asymmetry.

5. Fuller Assessment of the Super-X Divertor on MAST Upgrade

   The method for changing the chronological order in which one can create a super-X, use neutral beams, and induce detachment was only developed very late in the campaign, which severely limited the super-X scenarios which could be used for most of the campaign. The super-X shot analysed in chapter 6 (shot 45209) was ohmic and attached. Once the automatic plasma density control system comes online for the second campaign, it will enable true single-parameter scans to be carried out in the super-X configuration, as well a scans such as target radius or connection length, which can result in a transition to super-X during the scan. It will also allow a wider range of scenarios to be tested in a spherical tokamak with a highly enclosed divertor. This may inform certain design choices which are yet to be taken for future tokamaks such as DEMO, or especially the recently proposed Spherical Tokamak for Energy Production (STEP), which is a concept for a large spherical tokamak fusion reactor, and is expected to borrow heavily from lessons learned on MAST-U.

List of References

[1] IPCC. *Climate Change 2014*. Cambridge University Press, 2014.

[2] F. Wagner. Considerations for an EU-wide use of renewable energies for electricity generation. *Journal of Applied Genetics*, 2014.

[3] United Nations DESA. World population projected to reach 9.7 billion by 2050 — UN DESA — United Nations Department of Economic and Social Affairs, 2015.

[4] M. Kikuchi. A Review of Fusion and Tokamak Research Towards Steady-State Operation: A JAEA Contribution. *Energies*, 3(11):1741, 2010.

[5] L. Spitzer. The stellarator concept. *Physics of Fluids*, 1(4):253, 1958.

[6] J. I. Wesson. *Tokamaks, Fourth Edition*. Oxford Science Publications, 2012.

[7] W. Morris, J. R. Harrison, A. Kirk, B. Lipschultz, F. Militello, D. Moulton, and N. R. Walkden. MAST Upgrade Divertor Facility: A Test Bed for Novel Divertor Solutions. *IEEE Transactions on Plasma Science*, 46(5):1217, 2018.

[8] H. Q. Wang, et al. Effects of divertor geometry on H-mode pedestal structure in attached and detached plasmas in the DIII-D tokamak. *Nuclear Fusion*, 58(9):96014, 2018.

[9] R. G. Mills. Lawson Criteria. *IEEE Transactions on Nuclear Science*, 18(4):205, 1971.

[10] J. D. Lawson. Some Criteria for a Power Producing Thermonuclear Reactor. *Proceedings of the Physical Society. Section B*, 70(1):6, 1957.

[11] D. Brunner, B. LaBombard, A. Q. Kuang, and J. L. Terry. High-resolution heat flux width measurements at reactor-level magnetic fields and observation of a unified width scaling across confinement regimes in the Alcator C-Mod tokamak. *Nuclear Fusion*, 58(9):94002, 2018.

[12] R. P. Wenninger, et al. DEMO divertor limitations during and in between ELMs. *Nuclear Fusion*, 54(11):114003, 2014.

[13] N. Asakura, K. Shimizu, K. Hoshino, K. Tobita, S. Tokunaga, and T. Takizuka. A simulation study of large power handling in the divertor for a Demo reactor. *Nuclear Fusion*, 53(12):123013, 2013.

[14] K. Ikeda. Progress in the ITER Physics Basis. *Nuclear Fusion*, 47(6), 2007.

[15] P. Rindt, T. W. Morgan, M. A. Jaworski, and N. J. L. Cardozo. Power handling limit of liquid lithium divertor targets. *Nuclear Fusion*, 58(10), 2018.

[16] P. Liu, Y. Song, X. Peng, S. Qin, X. Mao, X. Qian, and J. Zhang. Conceptual design study for CFETR divertor target using CLAM steel as structural material. *Fusion Engineering and Design*, 131:90, 2018.

[17] C. Theiler, et al. Results from recent detachment experiments in alternative divertor configurations on TCV. *Nuclear Fusion*, 57(7):72008, 2017.

[18] F. F. Chen. *Introduction to Plasma Physics and Controlled Fusion*. Springer, 2016.

[19] A. Piel. Plasma Boundaries. In *Plasma Physics*, pages 169–195. Springer Berlin Heidelberg, Berlin, Heidelberg, 2010.

[20] T. Eich, et al. Scaling of the tokamak near the scrape-off layer H-mode power width and implications for ITER. *Nuclear Fusion*, 53(9), 2013.

[21] J. D. Galambos and Y.-K. M. Peng. Two-Point Model For Divertor Transport. Technical report, Oak Ridge National Laboratory, 1984.

[22] P. C. Stangeby. *The Plasma Boundary of Magnetic Fusion Devices*. IoP, 2000.

[23] B. Lipschultz, B. LaBombard, J. L. Terry, C. Boswell, and I. H. Hutchinson. Divertor Physics Research on Alcator C-Mod. *Fusion Science and Technology*, 51(3):369, 2007.

[24] G. F. Matthews. Plasma detachment from divertor targets and limiters. *Journal of Nuclear Materials*, 220-222:104 , 1995.

[25] T. W. Petrie, et al. Effect of changes in separatrix magnetic geometry on divertor behaviour in DIII-D. *Nuclear Fusion*, 53(11):113024, 2013.

[26] J. R. Harrison, et al. Detachment evolution on the TCV tokamak. *Nuclear Materials and Energy*, 12:1071 , 2017.

[27] F. Piras, et al. Snowflake divertor plasmas on TCV. *Plasma Physics and Controlled Fusion*, 51(5):55009, 2009.

[28] V. A. Soukhanovskii, et al. Developing physics basis for the snowflake divertor in the DIII-D tokamak. *Nuclear Fusion*, 58(3):36018, 2018.

[29] P. M. Valanju, M. Kotschenreuther, S. M. Mahajan, and J. Canik. Super-X divertors and high power density fusion devices. *Physics of Plasmas*, 16(5):56110, 2009.

[30] G. De Temmerman, E. Delchambre, J. Dowling, A. Kirk, S. Lisgo, and P. Tamain. Thermographic study of heat load asymmetries during MAST L-mode discharges. *Plasma Physics and Controlled Fusion*, 52(9):095005, 2010.

[31] S. Paul, R. Maingi, V. Soukhanovskii, S. Kaye, and H. Kugel. Accounting of the power balance for neutral-beam heated H-mode plasmas in NSTX. *Journal of Nuclear Materials*, 337-339(1-3 SPEC. ISS.):251, 2005.

[32] E. Delchambre, G. Counsell, and A. Kirk. Effect of micrometric hot spots on surface temperature measurement and flux calculation in the middle and long infrared. *Plasma Physics and Controlled Fusion*, 51(5):55012, 2009.

[33] J. C. Fuchs, D. Coster, A. Herrmann, A. Kallenbach, and K. F. Mast. Radiation distribution and power balance in the ASDEX Upgrade LYRA divertor. *Journal of Nuclear Materials*, 290-293:525, 2001.

[34] G. Matthews, et al. Energy balance in JET. *Nuclear Materials and Energy*, 12:227, 2017.

[35] G. F. Matthews, S. A. Silburn, C. D. Challis, T. Eich, D. Iglesias, D. King, B. Sieglin, and JET contributors. Dynamic power balance analysis in JET. *Physica Scripta*, T170:014035, 2017.

[36] A. V. Chankin, et al. SOLPS modelling of ASDEX upgrade H-mode plasma. *Plasma Physics and Controlled Fusion*, 48(6):839, 2006.

[37] E. Havlíčková, J. Harrison, B. Lipschultz, G. Fishpool, A. Kirk, A. Thornton, M. Wischmeier, S. Elmore, and S. Allan. SOLPS analysis of the MAST-U divertor with the effect of heating power and pumping on the access to detachment in the Super-x configuration. *Plasma Physics and Controlled Fusion*, 57(11):115001, 2015.

[38] B. Covele, M. Kotschenreuther, S. Mahajan, P. Valanju, A. Leonard, J. Watkins, M. Makowski, M. Fenstermacher, and H. Si. Increased heat dissipation with the X-divertor geometry facilitating detachment onset at lower density in DIII-D. *Nuclear Fusion*, 57(8):86017, 2017.

[39] V. A. Soukhanovskii. A review of radiative detachment studies in tokamak advanced magnetic divertor configurations. *Plasma Physics and Controlled Fusion*, 59(6):64005, 2017.

[40] E. Salonen, K. Nordlund, J. Keinonen, and C. H. Wu. Swift chemical sputtering of amorphous hydrogenated carbon. *Physical Review B - Condensed Matter and Materials Physics*, 63(19):1, 2001.

[41] K. Bystrov, L. van der Vegt, and G. De Temmerman. Reorganization of graphite surfaces into carbon micro- and nanoparticles under high flux hydrogen plasma bombardment. *Journal of Vacuum Science and Technology*, 31(1), 2013.

[42] K. Bystrov, M. C. M. van de Sanden, C. Arnas, L. Marot, D. Mathys, F. Liu, L. K. Xu, X. B. Li, A. V. Shalpegin, and G. De Temmerman. Spontaneous synthesis of carbon nanowalls, nanotubes and nanotips using high flux density plasmas. *Carbon*, 68:695 , 2014.

[43] F. Wang, L. Cheng, H. Mei, Q. Zhang, and L. Zhang. Effect of surface microstructures on the infrared emissivity of graphite. *International Journal of Thermophysics*, 35(1):62, 2014.

[44] A. Herrmann, B. Sieglin, and M. Faitsch. Surface temperature measurement of in-vessel components and its real-time validation. *Fusion Science and Technology*, 69(3):569, 2016.

[45] J. P. H. E. Ongena, I. Voitsekhovitch, M. Evrard, and D. McCune. Numerical Transport Codes. *Fusion Science and Technology*, 61(2T):180, 2012.

[46] C. W. Barnes, M. G. Bell, H. W. Hendel, D. L. Jassby, D. Mikkelsen, A. L. Roquemore, S. D. Scott, J. D. Strachan, and M. C. Zarnstorff. Absolute calibration of neutron detection systems on TFTR and accurate comparison of source strength measurements to transport simulations (invited). *Review of Scientific Instruments*, 61(10):3151, 1990.

[47] C. Hellesen, et al. Validating TRANSP simulations using neutron emission spectroscopy with dual sight lines. *Review of Scientific Instruments*, 79(10), 2008.

[48] C. Hellesen, et al. Neutron spectroscopy measurements and modeling of neutral beam heating fast ion dynamics. *Plasma Physics and Controlled Fusion*, 52(8):085013, 2010.

[49] A. Žohar, et al. Validation of realistic Monte Carlo plasma gamma-ray source on JET discharges. *Nuclear Fusion*, 2022.

[50] X. Xu, Y. Xu, X. Zhang, Y. Hu, L. Ye, and X. Xiao. Simulations of first-orbit losses of neutral beam injection (NBI) fast ions on EAST. *Plasma Science and Technology*, 22(8):085101, 2020.

[51] A. Herrmann, W. Junker, K. Gunther, S. Bosch, M. Kaufmann, J. Neuhauser, G. Pautasso, T. Richter, and R. Schneider. Energy flux to the ASDEX-Upgrade diverter plates determined by thermography and calorimetry. *Plasma Physics and Controlled Fusion*, 37(1):17, 1995.

[52] A. Herrmann and ASDEX Upgrade Team. Limitations for Divertor Heat Flux Calculations of Fast Events in Tokamaks. In C. Silva, C. Varandas, and D. Campbell, editors, *28th EPS Conference on Controlled Fusion and Plasma Physics. Contributed Paper*, volume 25A of *ECA*, pages 2109–2112. European Physical Society, Funchal, 2001.

[53] F. Lott, A. Kirk, G. F. Counsell, J. Dowling, D. Taylor, T. Eich, and A. Herrmann. Thermographic power accounting in MAST. In *Journal of Nuclear Materials*. 2005.

[54] K. F. Mast, J. C. Vallet, C. Andelfinger, P. Betzler, H. Kraus, and G. Schramm. A low noise highly integrated bolometer array for absolute measurement of VUV and soft x radiation. *Review of Scientific Instruments*, 62(3):744, 1991.

[55] M. L. Reinke, J. L. Terry, G. G. van Eden, B. J. Peterson, K. Mukai, T. K. Gray, and B. C. Stratton. Experimental tests of an infrared video bolometer on Alcator C-Mod. *Review of Scientific Instruments*, 89(10):103507, 2018.

[56] J. J. Thomson. Electricity and Matter. *Nature*, 77(2005):505, 1908.

[57] H. M. Mott-Smith and I. Langmuir. The Theory of Collectors in Gaseous Discharges. *Physical Review*, 28(4):727, 1926.

[58] S. Elmore, S. Allan, A. Kirk, A. Thornton, J. Harrison, P. Tamain, M. Kočan, and J. Bradley. Scrape-off layer ion temperature measurements at the divertor target in MAST by retarding field energy analyser. *Journal of Nuclear Materials*, 438:S1212, 2013.

[59] V. Tank and H. Dietl. Multispectral Infrared Pyrometer for Temperature Measurement with Automatic Correction of the Influence of Emissivity. *Infrared Physics*, 30(4):331, 1990.

[60] L. L. Lao, H. S. John, R. D. Stambaugh, A. G. Kellman, and W. Pfeiffer. Reconstruction of current profile parameters and plasma shapes in tokamaks. *Nuclear Fusion*, 25(11):1611, 1985.

[61] R. Hawryluk. An Empirial Approcah to Tokmak Transport. In B. Coppi, G. G. Leotta, D. Pfirsch, R. Pozzoli, and E. Sindoni, editors, *Physics of Plasmas Close to Thermonuclear Conditions*, pages 19–46. Elsevier, 1981.

[62] A. Pankin, D. McCune, R. Andre, G. Bateman, and A. Kritz. The tokamak Monte Carlo fast ion module NUBEAM in the national transport code collaboration library. *Computer Physics Communications*, 159(3):157, 2004.

[63] B. A. Grierson, et al. Orchestrating TRANSP Simulations for Interpretative and Predictive Tokamak Modeling with OMFIT. *Fusion Science and Technology*, 74(1-2):101, 2018.

[64] Y.-K. Peng and D. Strickler. Features of spherical torus plasmas. *Nuclear Fusion*, 26(6):769, 1986.

[65] D. Gates, et al. Overview of results from the National Spherical Torus Experiment (NSTX). *Nuclear Fusion*, 49(10):104016, 2009.

[66] I. T. Chapman, et al. Overview of MAST results. *Nuclear Fusion*, 55(10), 2015.

[67] R. Scannell, et al. Design of a new Nd:YAG Thomson scattering system for MAST. *Review of Scientific Instruments*, 79(10):10E730, 2008.

[68] R. Scannell, M. J. Walsh, M. R. Dunstan, J. Figueiredo, G. Naylor, T. O'Gorman, S. Shibaev, K. J. Gibson, and H. Wilson. A 130 point Nd:YAG Thomson scattering diagnostic on MAST. *Review of Scientific Instruments*, 81(10):10D520, 2010.

[69] L. L. Lao, H. E. S. John, Q. Peng, J. R. Ferron, E. J. Strait, T. S. Taylor, W. H. Meyer, C. Zhang, and K. I. You. MHD Equilibrium Reconstruction in the DIII-D Tokamak. *Fusion Science and Technology*, 48(2):968, 2005.

[70] A. C. Darke, R. J. Hayward, G. F. Counsell, and K. Hawkins. The MAST improved divertor. *Fusion Engineering and Design*, 75-79(SUPPL.):285, 2005.

[71] A. Loarte and P. J. Harbour. Effect of the magnetic flux geometry of a poloidal divertor on the profiles and parameters at the target. *Nuclear Fusion*, 32(4):681, 1992.

[72] M. Greenwald. Density limits in toroidal plasmas. *Plasma Physics and Controlled Fusion*, 44(8):201, 2002.

[73] M. Turnyanskiy, C. D. Challis, R. J. Akers, M. Cecconello, D. L. Keeling, A. Kirk, R. Lake, S. D. Pinches, S. Sangaroon, and I. Wodniak. Measurement and control of the fast ion redistribution on MAST. *Nuclear Fusion*, 53(5), 2013.

[74] D. Pace, M. Van Zeeland, B. Fishler, and C. Murphy. Consideration of neutral beam prompt loss in the design of a tokamak helicon antenna. *Fusion Engineering and Design*, 112:14, 2016.

[75] M. J. Dunn, T. W. Morgan, J. W. Genuit, T. Loewenhoff, A. J. Thornton, and K. J. Gibson. Thermographic investigation of the effect of plasma exposure on the surface of a MAST upgrade divertor tile in Magnum-PSI. *Nuclear Materials and Energy*, 25(November):4, 2020.

[76] G. De Temmerman, M. van den Berg, J. Scholten, A. Lof, H. van der Meiden, H. van Eck, T. Morgan, T. de Kruijf, P. Zeijlmans van Emmichoven, and J. Zielinski. High heat flux capabilities of the Magnum-PSI linear plasma device. *Fusion Engineering and Design*, 88(6-8):483, 2013.

[77] S. Bardin, T. Morgan, X. Glad, R. Pitts, and G. De Temmerman. Evolution of transiently melt damaged tungsten under ITER-relevant divertor plasma heat loading. *Journal of Nuclear Materials*, 463:193, 2015.

[78] H. J. N. van Eck, et al. High-fluence and high-flux performance characteristics of the superconducting Magnum-PSI linear plasma facility. *Fusion Engineering and Design*, 142:26, 2019.

[79] T. W. Morgan, P. Rindt, G. G. van Eden, V. Kvon, M. A. Jaworksi, and N. J. L. Cardozo. Liquid metals as a divertor plasma-facing material explored using the Pilot-PSI and Magnum-PSI linear devices. *Plasma Physics and Controlled Fusion*, 60(1):014025, 2018.

[80] G. J. van Rooij, H. J. van der Meiden, M. H. J. t. Hoen, W. R. Koppers, A. E. Shumack, W. A. J. Vijvers, J. Westerhout, G. M. Wright, and J. Rapp. Thomson scattering at Pilot-PSI and Magnum-PSI. *Plasma Physics and Controlled Fusion*, 51(12):124037, 2009.

[81] T. Morgan, M. van den Berg, G. De Temmerman, S. Bardin, D. Aussems, and R. Pitts. Power deposition on misaligned castellated tungsten blocks in the Magnum-PSI and Pilot-PSI linear devices. *Nuclear Fusion*, 57(12):126025, 2017.

[82] M. A. Van Den Berg, K. Bystrov, R. Pasquet, J. J. Zielinski, and G. De Temmerman. Thermographic determination of the sheath heat transmission coefficient in a high density plasma. *Journal of Nuclear Materials*, 438(suppl.), 2013.

[83] T. W. Morgan, T. M. de Kruif, H. J. van der Meiden, M. A. van den Berg, J. Scholten, W. Melissen, B. J. M. Krijger, S. Bardin, and G. De Temmerman. A high-repetition rate edge localised mode replication system for the Magnum-PSI and Pilot-PSI linear devices. *Plasma Physics and Controlled Fusion*, 56(9):095004, 2014.

[84] D. Hill. A review of ELMs in divertor tokamaks. *Journal of Nuclear Materials*, 241-243:182, 1997.

[85] G. Fishpool, J. Canik, G. Cunningham, J. Harrison, I. Katramados, A. Kirk, M. Kovari, H. Meyer, and R. Scannell. MAST-upgrade divertor facility and assessing performance of long-legged divertors. *Journal of Nuclear Materials*, 438(suppl.):S356, 2013.

[86] P. Ryan and S. Elmore. ALP Review, 2021.

[87] T. Eich, B. Sieglin, A. Scarabosio, W. Fundamenski, R. J. Goldston, and A. Herrmann. Inter-ELM Power Decay Length for JET and ASDEX Upgrade: Measurement and Comparison with Heuristic Drift-Based Model. *Physical Review Letters*, 107(21):215001, 2011.

[88] E. M. Hollmann, N. Commaux, N. W. Eidietis, C. J. Lasnier, R. A. Moyer, C. Paz-Soldan, D. Shiraki, and J. G. Watkins. Measurement of toroidal variation in conducted heat loads in locked mode induced disruptions on DIII-D. *Physics of Plasmas*, 25(10):102502, 2018.