

Development of Smart, Compact Fusion Diagnostics using Field-Programmable Gate Arrays

Jack Lovell

A Thesis presented for the degree of
Doctor of Philosophy



Centre for Advanced Instrumentation
Department of Physics
Durham University
United Kingdom

November 2017

Development of Smart, Compact Fusion Diagnostics using Field-Programmable Gate Arrays

Jack Lovell

Submitted for the degree of Doctor of Philosophy

November 2017

Abstract: Fusion research requires high quality diagnostics to understand the complex physical processes involved. Traditional analogue systems are complex, large and expensive, and expansion of diagnostic capabilities is often impossible without building a completely new system at considerable expense. Field-programmable gate array (FPGA) technology can provide a solution to this problem. By implementing complex functionality and digital signal processing on an FPGA chip, diagnostic hardware can be greatly simplified and compacted.

In this thesis we describe the enhancements of two diagnostics for the MAST-Upgrade tokamak using FPGA technology. Firstly, the design of the back end electronics for the new divertor bolometer is described. Results of tests of the new electronics at a number of sites, including lab-based testing and tokamak installations, are also presented. We demonstrate the correct functionality of the electronics and illustrate a number of important effects which must be taken into account when interpreting bolometer data on MAST-U.

Secondly, we describe the new control and acquisition electronics developed for the MAST-U divertor Langmuir probe diagnostic. Much of the analogue control circuitry of the previous system has been upgraded to a digital implementation on an FPGA, which results in a significantly more compact and cost effective design. Given that MAST-Upgrade will feature around 850 Langmuir probes, these improvements are extremely important to keep the diagnostic manageable. Again, results are presented from the testing of the system at several sites, which both demonstrate the correct functionality of the new system and provide information on the diagnostic behaviour which needs to be accounted for when interpreting the probe data during MAST-U experiments.

Declaration

The work in this thesis is based on research carried out at the Centre for Advanced Instrumentation in the Department of Physics at Durham University. No part of this thesis has been submitted elsewhere for any degree or qualification.

The original design and specification of the BOLO8 hardware, and a conceptual outline of the new calibration procedure, were already in place before I started this work. Similarly, the design and specification of the Langmuir probe electronics described herein had already begun before I started this work. The refinement and implementation of these specifications is however my own work.

Collaborations with other researchers and institutions have been indicated at appropriate points in the text.

Copyright © 2017 Jack Lovell.

“The copyright of this thesis rests with the author. No quotation from it should be published without the author’s prior written consent and information derived from it should be acknowledged.”

Acknowledgements

The production of this thesis has relied on a large number of people. Thanks are of course due to my supervisors Ray Sharples of Durham University and Graham Naylor of Culham Centre for Fusion Energy (CCFE) for their help and advice throughout the project, including the many hours of proof reading and suggestions for improvement of the thesis text itself. Anthony Field, Peter Drewelow and Mike Stamp provided assistance with the bolometer development at CCFE and JET, whilst Matt Reinke has proved to be an invaluable contact first at York University and then Princeton Plasma Physics Laboratory on behalf of Oak Ridge National Laboratory to enable collaborations with these institutions. Robert Stephen, Stuart Bray, Sarah Elmore and James Harrison have provided assistance with the Langmuir Probe system at CCFE, with Hannah Willett and Kieron Gibson of York University and Matej Peterka, Aleš Hávranek, Megi Dimitrova and Jirka Adámek of IPP CAS providing assistance with tests of the system at the York Linear Plasma Device and the COMPASS tokamak respectively.

Much of the data processing and plotting in the results chapters uses Python [1], together with the Scipy [2], Matplotlib [3] and xarray [4] libraries. Many diagrams are also drawn using the PGF/TikZ \LaTeX package [5] and the Inkscape graphics editor [6]. Thanks are due to M Imran and J Lawson who prepared the \LaTeX template which this thesis uses.

This work was supported by the Engineering and Physical Sciences Research Council, grants EP/L01663X/1 and EP/P012450/1. Parts of this work have been carried out within the framework of the EUROfusion Consortium and have received funding from the Euratom research and training programme 2014–2018 under grant agreement No 633053. The views and opinions expressed herein do not necessarily reflect those of the European Commission. Parts of this work are supported by Czech Ministry of Education, Youth and Sports project LM2015045. Parts of this work are supported by US Department of Energy grants DE-AC05-00OR22725 and DE-AC02-09CH11466.

Dedicated to

The strong force, without which
all this would be futile.

Contents

Abstract	iii
List of Figures	xv
1 Introduction	1
1.1 Nuclear Fusion	1
1.1.1 Physical process	2
1.1.2 An ideal energy source?	5
1.2 Plasma confinement	7
1.3 The tokamak	8
1.3.1 Tokamak principle	8
1.3.2 Tokamak Diagnostics	12
1.3.3 The divertor	14
1.4 MAST-Upgrade	15
1.4.1 Spherical tokamaks	17
1.4.2 The Super-X divertor	18
1.5 Field-programmable gate arrays	20
1.5.1 The principle of FPGAs	20
1.5.2 Xillybus	24
1.6 FPGAs and fusion diagnostics	24

2	Design of the Bolometer diagnostic	27
2.1	The bolometer principle	27
2.1.1	Resistive bolometers	28
2.1.2	AC synchronous detection	30
2.1.3	Calculation of absorbed power	31
2.1.4	Calibration	32
2.2	Hardware	34
2.2.1	New calibration procedure	39
2.3	FPGA design: Xillybus version	41
2.3.1	Digital Signal Processing	47
2.3.2	Jetblack software interface	51
2.4	FPGA redesign: BOLODSP module	53
2.5	Improvements over previous systems	56
3	Results from the Bolometer system	59
3.1	Lab-based measurements at the York Plasma Institute	59
3.1.1	Comparison of FPGA and software signal processing	60
3.1.2	Calibration dependence on bias voltage	63
3.1.3	Calibration dependence on pressure	65
3.1.4	Summary of York tests	67
3.2	Installation on JET	68
3.2.1	Direct voltage comparison	70
3.2.2	Optimisation of drive frequency	72
3.2.3	Power calculation	74
3.2.4	Tomography comparison	80
3.2.5	Total power calculation	83
3.2.6	Summary of JET tests	86
3.3	Installation on TCV	86
3.4	Lab-based measurements at PPPL	88
3.4.1	Laser profile	88

3.4.2	Sensitivity calculation: manual calibration method	91
3.4.3	Sensitivity calculation: BOLO8 calibration method	92
3.4.4	Absolute power accuracy	93
3.4.5	Laser power scan	95
3.4.6	Laser frequency scan	97
3.4.7	Cable length scan	104
3.4.8	Drive frequency scan	106
3.4.9	Summary of PPPL tests	113
3.5	Implications for installation on MAST-Upgrade	114
4	The Langmuir Probe diagnostic	115
4.1	Principles	115
4.1.1	The IV characteristic	116
4.1.2	MAST-U's divertor Langmuir probes	118
4.2	Hardware	119
4.2.1	The multiplexer	121
4.3	FPGA design	123
4.3.1	Web server with CherryPy	126
4.4	Improvements over previous system	127
5	Results from the Langmuir Probe system	129
5.1	Installation on the York Linear Plasma Device	129
5.1.1	Floating potential measurements	131
5.1.2	Sweep measurements	133
5.1.3	Plasma parameters from sweeps	137
5.1.4	Summary of YLPD tests	139
5.2	Installation on the COMPASS tokamak	139
5.2.1	Experimental setup	140
5.2.2	Floating potential measurements	141
5.2.3	Sweep waveform	146
5.2.4	Sweep measurements	149
5.2.5	Summary of the COMPASS tests	151
5.3	Implications for installation on MAST-U	152

6	Conclusions	155
6.1	Summary	155
6.2	Future work	158
6.3	Final word: the potential of integrated FPGA diagnostics in fusion .	160
	Bibliography	163

List of Figures

1.1	Average binding energy per nucleon for a range of nuclei	3
1.2	The cross sections for the 3 most relevant terrestrial nuclear fusion reactions	4
1.3	The principle of the magnetic mirror	9
1.4	The magnetic field configuration in a tokamak	11
1.5	The magnetic flux surfaces in a divertor configuration for the JET tokamak.	16
1.6	Comparison of conventional and Super-X divertor configurations in MAST-U.	19
1.7	FPGA architecture	21
1.8	Schematic of the Zynq architecture	23
2.1	Schematic of the bolometer circuit	29
2.2	Schematic of circuit typically used to calibrate bolometers	32
2.3	The BOLO8 and ACQ2006 hardware	36
2.4	Diagram showing the interaction between components of the bolometer system	37
2.5	Schematic of the BOLO8 module	38
2.6	Comparison of the spectra of raw ADC voltage measurements, with and without the Schottky diodes	38
2.7	Electrical setup of the BOLO8 calibration procedure	40
2.8	Schematic showing the modular design of the FPGA firmware	44
2.9	Flow diagram showing how the AC synchronous detection is performed on the FPGA	47
2.10	The user-programmable DSP interface	54

3.1	Comparison of floating-point and fixed-point signal processing for a calibration shot	61
3.2	Quantisation errors with 16-bit and 32-bit precision in the FPGA voltage output	62
3.3	Calibration dependence on bias voltage, in air	64
3.4	Calibration dependence on bias voltage, in vacuum	64
3.5	Variation of cooling time with pressure	66
3.6	Variation of sensitivity with pressure	66
3.7	Lines of sight of the KB5V and KB5F vertical bolometers	69
3.8	Comparison of KB5V and KB5F raw output voltages	71
3.9	Comparison of voltage spectra of KB5F and KB5V	71
3.10	Comparison of voltage spectra of KB5F and KB5V, with KB5F 18 kHz drive frequency	73
3.11	Spectrograms for 19.3 kHz drive frequency and 18 kHz drive frequency	74
3.12	Comparison of in-situ electrical and mercury lamp measures of KB5 bolometer sensitivities	76
3.13	Comparison of intensities as measured by KB5V and KB5F	77
3.14	Intensity calculation during a disrupted pulse	78
3.15	Comparison of real-time intensity calculation with KB5V's PPF	79
3.16	Tomographic reconstruction of emissivity profile	81
3.17	Comparison of measured intensities and intensities back-calculated from the tomographic reconstruction	82
3.18	Calculated total power for a high performance JET pulse	84
3.19	Calculated total power for a high performance JET pulse, zoomed to view only a few ELMs	85
3.20	Comparison of blackened and unblackened bolometer measurements performed on TCV	87
3.21	Laser beam profile when defocussed to ~ 20 mm diameter visible spot size	89
3.22	Cooling curves for the in phase and quadrature phase components of the bridge output voltage for a calibration	92
3.23	Power meter responsivity for a selection of wavelengths	94

3.24	Calculated power using the BOLO8 calibration for a slowly flashing, defocussed laser beam	95
3.25	Power measured by the bolometer as a function of the modulation voltage applied to the laser	96
3.26	Laser power absorbed by the bolometer sensor, as a function of laser frequency, with a square wave laser pulse	98
3.27	Time traces showing the calculated power for different laser frequencies at the same laser power	99
3.28	Spectra of laser frequencies, with laser frequency peak highlighted	101
3.29	Comparison of on/off mean and spectral methods of calculating laser power at different laser frequencies	102
3.30	Comparison of on/off mean and spectral methods of calculating laser power at different laser frequencies, after removing systematic offset from the 0.5 Hz trace	103
3.31	Sensitivity as a function of cable length for two different drive frequencies	105
3.32	Calculated power as a function of cable length	105
3.33	Measured bolometer power as a function of drive frequency	107
3.34	Peak bolometer output voltage as a function of drive frequency	107
3.35	Variation of calculated sensitivity with drive frequency	108
3.36	Ratio of bolometer output voltage amplitudes with forward and reverse biasing	109
3.37	LTspice circuit to model the bolometer sensor and BOLO8 system	110
3.38	The effective circuit used to derive the analytic formula for the bridge output voltage during calibration	111
3.39	Thévenin equivalent circuit for bolometer sensor and BOLO8 system	112
3.40	Comparison of bridge voltages as a function of frequency for the analytic model, LTspice simulations and PPPL data	112
4.1	Cartoon overview of the Langmuir probe principle	115
4.2	Theoretical and typical IV characteristics	117
4.3	Schematic of a Langmuir probe module	120
4.4	Photo of the Langmuir probe multiplexer	122

4.5	Schematic showing the modular design of the multiplexer FPGA firmware	125
5.1	Diagram of the Langmuir Probes in the York Linear Plasma Device	130
5.2	Floating potential measurements in the York Linear Plasma Device	132
5.3	Voltage sweep waveform used on the York Linear Plasma Device	133
5.4	IV characteristic in YLPD in vacuum	134
5.5	Comparison of IV curves with different current removal methods.	135
5.6	Comparison of current subtraction methods throughout one YLPD discharge.	137
5.7	3-parameter fit of an IV curve in the YLPD	138
5.8	The probe arrays used for the COMPASS tests.	141
5.9	Floating potential for COMPASS shot 12920.	142
5.10	Floating potential for COMPASS shot 12920, zoomed to before the L-H transition.	143
5.11	Floating potential for COMPASS shot 12920, zoomed to the H-mode region.	144
5.12	Floating potential for COMPASS shot 12987.	145
5.13	The sweep waveform used on COMPASS	147
5.14	Parasitic current measured in vacuum for a COMPASS probe.	148
5.15	Residual current after removing parasitic current for a single COMPASS probe.	148
5.16	Successive inter-ELM sweeps of multiple probes in COMPASS	150
5.17	D-alpha signal, showing the times of the plotted sweeps.	151

Glossary

ADC analogue-to-digital convertor

AMBA Advanced Microcontroller Bus Architecture

ASIC application specific integrated circuit

AXI Advanced eXtensible Interface, the third generation of the AMBA interface

AXUV absolute extreme ultra violet

BOLO8 a new 8-channel bolometer electronics module developed by D-TACQ

CCFE Culham Centre for Fusion Energy

COMPASS COMPact ASSEmblly, a tokamak based at the Czech Academy of Sciences

CORDIC COordinate Rotation DIgital Computer, an algorithm to calculate hyperbolic and trigonometric functions

CPU central processing unit

CTF component test facility

DAC digital-to-analogue convertor

DDR double data rate

DMA direct memory access

DSP digital signal processing

EFIT a magnetic equilibrium reconstruction code

ELF A D-TACQ extension to the FMC standard, which retains the form factor of the FMC connector but reserves some of the pins and increases the footprint of compliant modules

ELM edge-localised mode, a type of plasma instability

EPFL École Polytechnique Fédérale de Lausanne

FIFO first-in, first-out; a type of memory buffer

FIR finite impulse response

FMC FPGA mezzanine card, a type of connector often used with FPGAs

FPGA field-programmable gate array

FSM finite state machine

GAP General Acquisition Program

HDL hardware description language

HTTP hyper text transport protocol

I²C inter-integrated circuit

I/O input/output

ICF Inertial Confinement Fusion

IFMIF International Fusion Materials Irradiation Facility

IP intellectual property

JET Joint European Torus, Europe's flagship tokamak

KB5F a subset of 8 channels of KB5V which were connected to the new MAST-U electronics

KB5V the vertical bolometer array, part of the KB5 bolometer diagnostic on JET

LPC low pin count, a variant of the FMC connector with fewer I/O pins for reduced complexity

LSB least significant bit

LUT look-up table

MAST Mega Amp Spherical Tokamak

MAST-U MAST-Upgrade, an upgrade of the MAST tokamak

-
- MCF** Magnetic Confinement Fusion
- MDSplus** a set of software tools for data acquisition and storage, used at several tokamaks worldwide
- MIT** Massachusetts Institute of Technology
- MLP** Mirror Langmuir probe
- MMCM** mixed-mode clock manager
- NSTX** National Spherical Torus Experiment
- NSTX-U** an upgrade to the NSTX tokamak
- PL** programmable logic, the FPGA-based part of a Zynq chip
- PmodTM** a 6- or 12-pin connector developed by Digilent to interface with FPGA boards
- PPF** processed pulse file; stores JET diagnostic data which has had some post processing applied
- PPPL** Princeton Plasma Physics Laboratory
- PS** processing system, the part of a Zynq chip containing the ARM CPU, interconnect and peripherals
- PV** photovoltaic
- RAM** random access memory
- RMP** resonant magnetic perturbation
- RMS** root mean square
- SFP** small form-factor pluggable transceiver, an optical module transceiver used for data communication
- SPI** serial peripheral interface
- ST** spherical tokamak
- START** Small Tight Aspect Ratio Tokamak
- TCP/IP** The internet protocol suite, including the Transmission Control Protocol (TCP) and Internet Protocol (IP)

TCV Tokamak à Configuration Variable

UKAEA United Kingdom Atomic Energy Authority

YLPD York Linear Plasma Device

Chapter 1

Introduction

1.1 Nuclear Fusion

Demand for energy is increasing. Technological advancements and increasing population mean global energy demand is predicted to rise by 30% from today's levels by 2040 [7]. This increased demand must be matched by an increased supply of energy, with electrical energy estimated to represent between 40% and 66% of the additional energy demand. More electricity generating facilities are therefore urgently needed.

Fossil fuels such as coal, oil and gas, the major fuel sources for power stations for over a century and still expected to meet about 84% of energy demand in 2030 [8], are no longer a long term option for electricity generation. Estimated reserves of oil, coal and gas are expected to last for 35, 107 and 37 years respectively [8], and so in only a few decades' time a dramatic shift in methods of electricity generation must have already occurred if we are to avoid catastrophic fuel supply shortages. Furthermore, the implementation of the Paris Climate Agreement requires phasing out fossil fuel use in order to reduce global temperature rises to below 2 °C [9]. We therefore need to find alternatives to fossil fuels for electricity generation.

Several other methods of electricity generation are already used. Renewables such as solar photovoltaic (PV), wind, wave, tidal and geothermal power promise carbon-free electricity generation, but their reliance on suitable environments limits where these technologies can be deployed and they have comparatively low energy densities [10]. Their output is also dictated by the local environmental conditions and so cannot be adjusted to match demand. Technologies such as solar and wind are intermittent and so must be paired with as-yet-underdeveloped energy storage solutions to smooth variations in both the supply and demand of electrical power. Whilst they do have

an important role in future carbon-free electricity generation, these renewables are not the complete solution to the problem of provision of future energy.

Nuclear fission solves many of the issues encountered with renewables. Nuclear power produces no greenhouse gas emissions, has an extremely high energy density (so high power output power stations require very small amounts of fuel and have a small physical footprint) and can provide a steady and predictable power output. This makes it suitable to provide base load electrical power, but fission also has a number of issues. Due to the potential for reactor meltdowns, new power stations must have extremely robust and sophisticated safety systems designed in, which greatly increases the capital cost of construction and decommissioning. This means fission power stations take many decades to pay back the initial investment, which generally restricts construction of these plants to state-backed enterprises. Crucially, fission reactions produce long-lived intermediate level radioactive waste, which must be safely contained for many millennia. Current estimated reserves of uranium fuel are projected to last for another 150 years using current technology, though reactor innovations could extend that significantly [11].

It is clear that no current technology will be suitable as the sole provider of humanity's future electricity needs. Fossil fuels simply must be replaced if we are to avoid dangerous climate change, even if new reserves are found. Renewables and nuclear fission have complementary strengths and weaknesses and so could form part of the solution. However, another technology is being developed which promises to provide the carbon-free, high power density of nuclear fission but without the meltdown risk and the long-lived radioactive waste: nuclear fusion.

1.1.1 Physical process

Thermonuclear fusion is the process which powers the sun, and all other main sequence stars in the universe. When two lighter nuclei come within the range of the strong force, they fuse together to form a more massive nucleus. If the binding energy of the products is greater than that of the reactants, there is a net energy release.

Figure 1.1 shows the binding energies for a number of common isotopes. At low atomic mass, the binding energy increases sharply with increasing atomic mass, so fusion reactions between these isotopes release a large amount of energy. The binding energy increase slows down above ^{12}C and peaks with ^{56}Fe , after which the binding energy per nucleon decreases. Any fusion reaction involving isotopes with an atomic mass above 56 amu will therefore result in net energy absorption, whereas




Figure 1.1: The average binding energy per nucleon of a range of nuclei, up to ^{238}U . As the atomic mass increases from 1 to 56 the binding energy generally increases, making fusion exothermic. Above 56, the binding energy per nucleon decreases, making fusion endothermic and fission exothermic. Reproduced from [12]

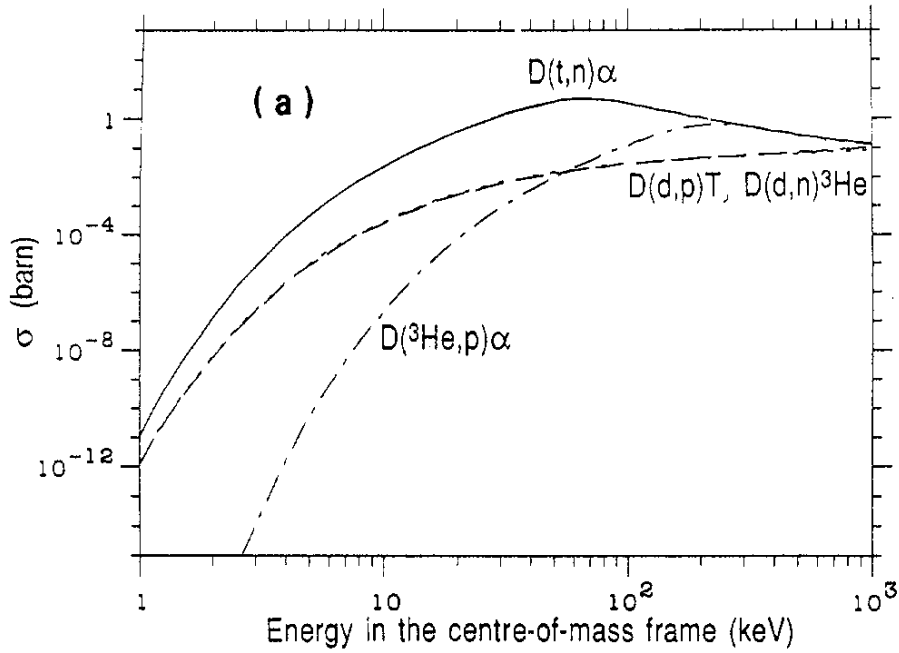


Figure 1.2: The cross sections of the 3 most relevant reactions for terrestrial nuclear fusion. All involve a Deuteron (D), with the highest cross section at the lowest energy being that of the reaction with Tritium (t), to produce a neutron (n) and a ${}^4\text{He}$ (α) particle. Reproduced from [13]

a fission reaction breaking these heavy isotopes into lighter products will result in a net energy release.

In order to get close enough to fuse, nuclei must overcome Coulomb repulsion. Since the Coulomb force is proportional to the product of the atomic numbers of the reactants, it is the lightest nuclei which require the lowest activation energies for fusion. Whilst heavier nuclei can fuse inside the cores of massive stars and supernovae, the fusion of Hydrogen isotopes is most readily achievable here on Earth in a power station and so it is on this that we focus our efforts.

Figure 1.2 shows the cross sections for the reactions most relevant to directly harnessing nuclear fusion power on Earth. The reaction with the highest cross section is the Deuterium-Tritium (D-T) reaction, for which the cross section peaks at approximately 70 keV:



The reaction produces a ${}^4\text{He}$ nucleus and a neutron, with a net release of energy of 17.6 MeV. Conservation of energy and momentum imply that the Helium nucleus carries 3.5 MeV and the neutron carries 14.1 MeV of this energy.

For densities n_d and n_t of Deuterium and Tritium respectively, the reaction rate is given by [14]

$$\Gamma = n_d n_t \langle \sigma v \rangle \quad (1.1.2)$$

Here, σ is the cross section and v is the relative velocity of the two species. To maximise the cross section, the particles should have the energy which maximises $\langle\sigma v\rangle$ and n_d should equal n_t (i.e. a 50/50 mix of D and T).

In an ideal reactor, the energy required to heat the fuel up to energies of around 70 keV would be completely supplied by the energy released from previous fusion reactions (this is known as ignition). The heating power from fusion reactions completely balances the power lost from the system when the following condition is satisfied [14]:

$$n\tau_E > \frac{12}{\langle\sigma v\rangle} \frac{T}{\mathcal{E}_\alpha} \quad (1.1.3)$$

Here, n is the density of the fuel, T is the temperature and \mathcal{E}_α is the energy of an alpha particle produced by D-T fusion. The energy confinement time τ_E is given by $\tau_E = W/P_H$ where W is the thermal energy stored in the fuel and P_H is the heating power which must be applied to maintain this thermal energy.

For D-T fuel, Equation 1.1.3 has a maximum at $T = 30$ keV [14]. This means that to achieve ignition the following more specific criterion must be met:

$$n\tau_E > 1.5 \times 10^{20} \text{ m}^{-3}\text{s} \quad (1.1.4)$$

This is known as the Lawson criterion. Meeting it is a requirement of any successful fusion reactor.

1.1.2 An ideal energy source?

At first sight, nuclear fusion appears to be the holy grail of climate-friendly electricity generation. No greenhouse gases are produced in the reaction, which will help in the drive to minimise climate change. The energy density of the fuel is huge: 7 orders of magnitude higher than fossil fuels and several times higher even than fission. This means that a very small amount of fuel is required for high power output. Deuterium is abundant in nature: it can be easily extracted from seawater. Tritium can be produced from Lithium which is bombarded with high energy neutrons, allowing a D-T fusion reactor to breed its own Tritium. Since Lithium is widely available, the fuel for fusion can be readily obtained all over the world, and supplies will last for millions of years. This greatly improves fuel security for nations with fusion power stations, and does not limit fusion power generation to specific regions like renewables do.

Unlike fission, the reaction products of fusion (Helium and a neutron) are non radioactive. Power plant components will become radioactive due to neutron bombardment, but the radioactivity is short lived and these components will be safe

to recycle or dispose of conventionally within approximately 100 years [15]. This greatly reduces the burden on future generations to look after our waste, making fusion more environmentally responsible than fission. Furthermore, the nature of the fusion reaction is such that a runaway reaction and hence meltdown is not possible: only a small amount of fuel is in the reactor at any one time and the reaction is not self sustaining, which improves operational safety compared with fission.

Fusion is however not without its challenges. Heating the fuel up to 30 keV (around 3.5×10^8 K) and achieving sufficient confinement to satisfy the Lawson criterion (Equation 1.1.4) is extremely challenging. The fuel must be contained within a vacuum vessel which can withstand temperatures of only a few thousand degrees K, requiring enormous temperature gradients within the fuel close to the vessel wall. These drive instabilities which degrade the confinement and make it more difficult to achieve the Lawson criterion. Plasma confinement methods which have got closest to achieving the Lawson criterion (see Section 1.2) typically result in very high heat fluxes to the vessel wall, which would cause significant melting in a commercial reactor.

Even if we successfully get the fusion reaction going and sustain it, the 14.1 MeV neutrons produced in the reactions must be dealt with. Structural components of the reactor exposed to such high energy neutron fluxes will be severely damaged and potentially transmuted into radioactive isotopes. A significant research effort is underway to design materials which can withstand these conditions, but it is of course difficult to perform experiments with high doses of fusion neutrons without a working fusion reactor in the first place. The International Fusion Materials Irradiation Facility (IFMIF) has been proposed as a source of neutrons with a similar spectrum to fusion neutrons [16], as well as several designs for fusion reactors designed to produce high neutron fluxes rather than net energy output [17], but these are still some years away from being realised.

The Tritium required for the D-T fusion reaction does not occur naturally, and so will have to be bred from Lithium. Whilst we are fortunate that a physical mechanism exists for doing this, breeding enough Tritium to provide for the reactor is very difficult. Extensive research is also being performed here, to optimise the designs of breeding systems for future reactors [18].

Clearly then, achieving controlled nuclear fusion is a difficult goal. Dealing with the extreme demands on a fusion reactor without requiring a plant that is prohibitively expensive to build is challenging, and commercially viable fusion power stations are still some way off. However, the promise of a clean, safe and reliable base load energy supply using abundant fuel is an enticing one, and makes fusion perhaps the best source of base load electricity in the future. Together with intermittent renewable

sources and sophisticated energy storage solutions, fusion can help to provide for humanity's future energy needs whilst reducing the harmful effects of anthropogenic climate change.

1.2 Plasma confinement

We described in Section 1.1.1 the need to heat the fuel up to energies of approximately 30 keV in order to achieve ignition of the fusion fuel. At a temperature of around 3×10^8 K, this is several orders of magnitude higher than the melting points of commonly available materials such as steel, which would be used to build the reactor vessel. The fuel and the vessel are therefore far from being in thermal equilibrium, so we must devise a way of keeping the fuel hot and dense enough to satisfy the Lawson criterion whilst minimising energy transfer from the fuel to the vessel wall.

The required energy is also 2 orders of magnitude higher than the ionisation energy of hydrogen (13.6 eV), so at this temperature the fuel is a fully ionised plasma. This brings its own challenges: plasmas produce electromagnetic fields and are influenced by electromagnetic fields, both internal to the plasma and externally applied. In addition to the non-linear fluid dynamics described by the Navier-Stokes equation, we must consider this electromagnetic response when attempting to understand and predict the dynamics of a fusion plasma. We need a good understanding of such plasma dynamics if we are to devise and implement effective confinement solutions.

Broadly speaking, the Lawson criterion described by Equation 1.1.4 can be satisfied in two ways. Firstly, we can operate with relatively low densities n and aim for a long confinement time τ_E . The fact that charged particles follow the field lines of an externally applied magnetic field is exploited by applying such a field with the field lines kept clear of the vessel walls. The plasma particles will follow these field lines and be confined by them long enough to fulfil the Lawson criterion. This is known as Magnetic Confinement Fusion (MCF), and is one of the most popular confinement methods to be studied.

The other most heavily studied confinement method goes to the opposite extreme. By using extremely high fuel densities, the confinement time can be made very short whilst still satisfying the Lawson criterion. Indeed, the confinement time can be made short enough that the rate of expansion of ignited fuel, which is limited by the sound speed in the material, means that fuel is confined by its own inertia for longer than the required confinement time. This requires densities of around 10^{27} m^{-3} to achieve confinement times of around 1 μs . Such a technique is known as Inertial Confinement Fusion (ICF) [19].

Most fusion research has been focussed on either MCF or ICF. There are private companies investigating techniques using a combination of the two, such as General Fusion's magnetised target fusion [20, 21], and also other schemes such as Inertial Electrostatic Confinement (IEC) [22] and muon-catalysed fusion [23], but these are either at a much earlier stage of development or suffer from fundamental physical processes which are believed to prevent them from producing net energy gain. MCF appears to be the most likely to deliver the first commercial fusion power station, with designs aiming for a demonstration power plant by the middle of the 21st century [24], and it is this approach which we will focus on.

1.3 The tokamak

As described in Section 1.2, MCF exploits the fact that the electrons and ions in the plasma follow magnetic field lines. Within MCF there are several approaches to obtaining the magnetic field configuration necessary to confine a low density plasma with a high enough confinement time to satisfy the Lawson criterion. The most mainstream of these are the stellarator, which uses a toroidal magnetic field with a rotational transform (the poloidal revolution of magnetic field lines as they travel toroidally around the device) supplied by specially shaped external coils [25], and the tokamak, another toroidal field concept but with the rotational transform supplied by a poloidal field generated by a current in the plasma itself. We shall confine our discussion to the tokamak.

1.3.1 Tokamak principle

The simplest way to confine a plasma magnetically is to apply a uniform field along the axis of a hollow cylindrical vessel. The electrons and ions gyrate around the field lines with a radius ρ given by the Larmor radius [14]:

$$\rho = \frac{mv_{\perp}}{|Ze|B} \quad (1.3.1)$$

In Equation 1.3.1, m is the mass of the particle, v_{\perp} is its velocity perpendicular to the field line, Z is the atomic number (-1 for an electron), e is the electron charge and B the magnetic field strength. Applying a magnetic field in this way greatly reduces transport perpendicular to the field lines, and could be achieved with a simple solenoid. However, there is no confinement along the field lines, and so particles would simply be lost at the ends of the device.

One way to reduce these end losses is to make use of magnetic mirrors. These exploit the fact that a plasma has certain quantities associated with it that are constant.

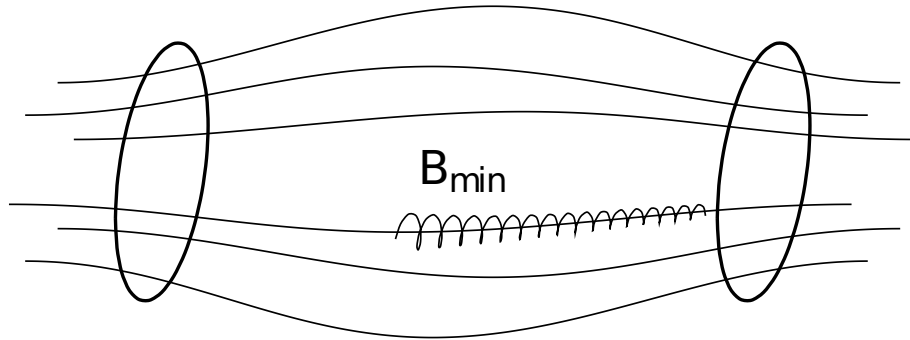


Figure 1.3: The principle of the magnetic mirror. The two circles represent magnetic field-producing coils, and field lines are shown running through these coils. The motion of one particle following a field line with a helical path is also shown. The magnetic field is higher at the ends than in the middle, causing particles to be confined as long as their velocity parallel to the field is not too large. Reproduced from [26]

One of these is the first adiabatic invariant,

$$\mu = \frac{mv_{\perp}^2}{2B} \quad (1.3.2)$$

From Equation 1.3.2 we see that as the magnetic field strength increases, so does the perpendicular velocity of the particle. But a time-invariant magnetic field does no work on the particle, so the total velocity $v^2 = v_{\parallel}^2 + v_{\perp}^2$ remains constant. When v_{\perp} reaches v , $v_{\parallel} = 0$ and so the particle cannot travel any further parallel to the field line. Thus by having a magnetic field which varies from B_{\min} at the centre to B_{\max} at the ends of a straight chamber, we can in theory linearly confine particles. A particle will be confined if the component of its velocity perpendicular to the field lines at the minimum magnetic field strength satisfies the following relation:

$$\frac{B_{\max}}{B_{\min}} < \frac{v^2}{v_{\perp,\min}^2} \quad (1.3.3)$$

This is represented graphically in Figure 1.3. Note also how the particle's Larmor radius decreases in the region of stronger magnetic field.

Particles which do not satisfy the relation in Equation 1.3.3 are said to be in the “loss cone”, and will escape out of the ends of the mirror. Whilst initially this may be only a small fraction of the total number of particles (especially for a large variation between B_{\max} and B_{\min}), collisions within the plasma mean that eventually most of the ions end up in the loss cone, and so escape confinement. Indeed, early magnetic mirror machines did not have good enough confinement to be viable fusion devices [26].

It is necessary to reduce the end losses from a magnetic mirror, and perhaps the most effective way to do this would be to get rid of the ends altogether. Bending

the chamber into a torus shape does just this, and does indeed eliminate end losses. However, since the magnetic axis is no longer straight, this new geometry does cause some additional problems.

The problems stem from the Lorentz force, $m\frac{d\mathbf{v}}{dt} = Ze(\mathbf{E} + \mathbf{v} \times \mathbf{B})$. When there is no electric field \mathbf{E} present, and the magnetic field \mathbf{B} is constant, we get the same gyrotory motion as before, with Larmor radius ρ given by Equation 1.3.1. However, it can be shown quite simply from this equation [14] that when there is an electric field, the guiding centre of the gyrations drifts with a velocity given by:

$$\mathbf{v}_d = \frac{\mathbf{E} \times \mathbf{B}}{B^2} \quad (1.3.4)$$

This is known as the “E cross B drift”. So if there were any electric fields in our toroidal device, the electrons and ions would both drift in the same direction, perpendicular to both these fields, and escape confinement.

Unfortunately there are further guiding centre drifts which readily produce such an electric field. If the magnetic field is not uniform, a simple linear expansion of the Lorentz force equation assuming $\mathbf{E} = 0$ and $\mathbf{B} = \mathbf{B}_0 + \delta\mathbf{B}$ results in a drift, known as the “grad-B drift”, with a drift velocity given by:

$$\mathbf{v}_d = \frac{mv_{\perp}^2}{2ZeB} \frac{\mathbf{B} \times \nabla\mathbf{B}}{B^2} \quad (1.3.5)$$

There is also a third drift, due to the curvature of the magnetic field. This arises due to the centrifugal force the particles feel, and following similar derivations to the above it can be shown that the “curvature drift” velocity for a magnetic field with radius of curvature R_c is given by:

$$\mathbf{v}_d = \frac{mv_{\parallel}^2}{Ze} \frac{\mathbf{R}_c \times \mathbf{B}}{B^2 R_c^2} \quad (1.3.6)$$

If there are no currents present, we can assume that the B field varies with radius simply as $\nabla\mathbf{B} = -(B/R)\hat{\mathbf{r}}$. Then both grad B and curvature drifts are in the same direction, and the resultant drift is given by:

$$\mathbf{v}_d = \frac{m(\frac{1}{2}v_{\perp}^2 + v_{\parallel}^2)}{ZeB} \frac{\mathbf{B} \times \nabla\mathbf{B}}{B^2} \quad (1.3.7)$$

We can see that this drift velocity is perpendicular to the magnetic field and its gradient, but also depends on Z , the particle charge. This means that electrons and ions drift in opposite directions, which sets up an electric field and so both species experience the E cross B drift and move out of the plasma. We therefore need a way to prevent the \mathbf{E} field from developing in our toroidal geometry.

A tokamak solves this problem by introducing a poloidal magnetic field, in addition




Figure 1.4: The magnetic field configuration in a tokamak. The toroidal field is produced by a series of external coils in the poloidal plane. The poloidal field is generated by a current induced in the plasma by a central solenoid. The net effect is a helical field which causes guiding centre drifts to cancel out over many toroidal revolutions. Reproduced from [28]

to the toroidal magnetic field. The result is that the field lines spiral round the device, and so particles following them move between regions of high and low magnetic field, as well as seeing different amounts of curvature. The guiding centre drifts cancel themselves out as the particles travel around the device poloidally and toroidally. This means that there is no charge separation, no electric field built up and no net drift.

The toroidal field is produced by a series of coils in the poloidal plane surrounding the device, with the result that the magnetic field is bent into a torus. The poloidal field is produced by inducing a current in the plasma using a large solenoid through the centre of the device, which acts as the primary core of a transformer with the plasma itself as a secondary core. This setup is illustrated in Figure 1.4, along with additional coils used to control the position and shape of the plasma. It is this configuration which has so far been the most successful in magnetically confining fusion plasmas, with the JET tokamak achieving the world record for controlled fusion power at 16.1 MW using 25.7 MW of heating power (a fusion power gain of $Q = 0.63$) [27].

1.3.2 Tokamak Diagnostics

The tokamak's magnetic geometry solves the problem of classical particle drifts, but the design is still susceptible to other plasma instabilities, some of which are actually caused by the requirement to drive large currents in the plasma. The transport of energy out of the core plasma happens at a much higher rate than that predicted simply by cross-field diffusion, which indicates that the transport is dominated by neoclassical effects and turbulence [14]. Whilst a detailed discussion of the many tokamak plasma instabilities and transport mechanisms is beyond the scope of this work, it is clear that we need a good understanding of the physics of these plasmas if we are to minimise transport and hence maximise the confinement time. An effective suite of diagnostics is therefore essential.

There are many challenges to operating diagnostics in a tokamak environment. Fusion plasmas are necessarily at extremely high temperatures, and so it is not possible to physically insert any probe into the core of such a plasma without it being destroyed. This limits physical probes to measuring the cooler edge regions of the plasma. Any measurements of the core plasma must rely on emission of electromagnetic radiation (or neutrons from fusing plasmas), which may either come naturally from the plasma or be stimulated by plasma absorption of externally applied radiation using active probing diagnostics [29]. The physics underpinning these diagnostics can range from relatively simple classical electromagnetism to extremely complex quantum and atomic physics. Often a “synthetic diagnostic”, which predicts the diagnostic measurements for a simulated system with known physical parameters, is needed to translate the measurements of the real diagnostic into the physics quantities of interest [30].

Many plasma phenomena, particularly turbulence, occur on very short time scales, sometimes less than 1 μs . Diagnostics must have sufficiently fast response times to resolve these, but also sufficient sensitivity and signal-to-noise ratios to provide meaningful results with the amount of input (e.g. photons) that can actually be collected on these short time scales. It is also very difficult to make precise, accurate measurements, even on longer time scales. Diagnostic measurements can suffer from random fluctuations caused by plasma turbulence occurring on much faster time scales than the diagnostic is designed for, which show up as noise on the measurements if the fluctuations cannot properly be resolved. In addition, the many power supplies required to produce currents in the magnetic field coils and power plant systems can induce pickup in diagnostics which make electrical measurements difficult. Mirrors and lenses used for optical measurements can become coated by debris ejected from the walls by the plasma, and diagnostics installed on a tokamak

with a D-T plasma hot enough to achieve fusion must be able to survive bombardment by high energy neutrons.

Tokamaks may be operated in a high confinement mode, or H mode, which is characterised by greatly reduced transport of heat at the edge of the plasma [14]. Whilst this is a useful way of increasing the temperature of the core plasma, H-mode plasmas are susceptible to edge-localised modes (ELMs), a type of plasma instability which results in periodic collapses in the edge confinement and the loss of a significant fraction of the stored energy in the plasma in a very short time (typically less than 1 ms). The high transient heat loads and particle fluxes from ELMs can cause significant damage to the tokamak wall, for example by melting parts of it or causing particles in the wall to be sputtered and enter the plasma as impurities. It is crucial in a future commercial tokamak to suppress or mitigate ELMs to reduce this damage, and for this we need a good physical understanding of how and why ELMs occur. For this we need high quality diagnostic measurements to compare with modelling predictions.

Even more serious than ELMs are disruptions: sudden and complete losses of plasma confinement. These can be extremely dangerous: a huge amount of stored energy is rapidly dissipated, causing large mechanical stresses and intense heat loads [14]. A disruption in a reactor-scale tokamak could severely damage or destroy the device, leading to expensive down time. Large disruptions therefore must be avoided in large scale devices, and for this we need to be able to predict when a disruption is likely to occur and apply some mitigation. Typically this involves injecting a large amount of impurity gas to dissipate energy radiatively before the disruption occurs. As well as being able to measure key plasma parameters in order to predict if a disruption is imminent, we must make this information available quickly enough that control systems can take the appropriate action to prevent a damaging disruption. This places additional demands on diagnostics: they must be able to interact with plasma control and protection systems.

The integration of diagnostics as sensors within a control loop has applications beyond disruption mitigation. Finer control of the plasma, and hence higher performance, can be achieved using feedback systems. For example, we can control the magnetic geometry through varying the currents in the poloidal field coils to optimise the positioning of the plasma and avoid vertical displacement of the plasma causing disruptions, and we can mitigate heat loads on the divertor (see Section 1.3.3) by moving the strike point (the position where the plasma exhaust contacts the wall) back and forth. By controlling the amount of fuelling, where the fuel enters the plasma and which gases (or pellets) enter the fuel, we can optimise the amount of radiation emitted by the plasma to reduce the heat load on the wall without

compromising the fusion performance, and also mitigate ELMs.

Looking forward to a commercial reactor, diagnostics will no longer be needed for physics studies, and as few diagnostics as possible will be used in order to minimise the number of ports required in the wall and hence maximise the amount of space available for breeding Tritium. The remaining diagnostics will be required to act as either safety monitors or sensors for control loops. We must therefore consider the ability of newly developed diagnostics to fulfil these roles if we are to maximise the longevity and usefulness of such diagnostics.

The need for a wide range of diagnostics covering all regions of the plasma, and the different requirements and challenges faced by those diagnostics, means there is no one diagnostic system which can measure all the physics of interest. The field of plasma diagnostics is very broad, and we shall concentrate in this work on two diagnostics in particular. The bolometer system used to measure power radiated by the plasma is described in Chapters 2 and 3, and the Langmuir probe system used to measure particle and heat fluxes in the edge of the plasma is described in Chapters 4 and 5. Both of these diagnostics are concentrated on the exhaust region of the tokamak vessel, known as the divertor.

1.3.3 The divertor

We noted in Section 1.2 the very high temperature difference between the plasma and the vessel. The need to keep the core plasma hot and the limits on pressure and temperature gradients imposed by plasma instabilities [14] means that the plasma at the edge of the device is at a very high temperature. The vessel wall consists of many sensitive components used for heating, control and diagnosis, which must be protected from the plasma.

Traditionally, protection of the wall was achieved using limiters: heat-resistant structures protruding from the wall which defined the plasma boundary. However, the close proximity of the limiter to the confined plasma meant impurities sputtered from the limiter were able to readily enter the core, which resulted in unacceptable radiative cooling. Additionally, high powered tokamaks would have edge plasma temperatures too high even for limiters to withstand without being quickly destroyed.

The solution to these problems was the divertor concept. Using additional poloidal field coils, a magnetic field null is created at a point near the plasma edge. This causes the plasma particles travelling along field lines which connect to the wall (open field lines) to be “diverted” away from the core so that they contact the wall sufficiently far from the core to reduce impurity influx. The particles also travel

further between leaving the closed field lines in the core and hitting the wall, and this increased connection length helps the plasma to cool down and hence reduce the heat load on the target. The magnetic geometry is shown in Figure 1.5, which displays the magnetic flux surfaces at $t = 50$ s for pulse 92504 of the JET tokamak. Similar to the limiters, the divertor target is typically made of materials which can withstand high temperatures, such as graphite or tungsten. Even though the divertor concept significantly reduces the heat loads on the target, the temperatures are still extremely high for high-powered tokamaks. In fact, the predicted plasma exhaust temperatures in commercial fusion reactors using conventional divertors is still too high for existing plasma-facing materials to cope with, and this “exhaust problem” is currently a key obstacle on the road to commercial fusion energy [31].

One way of mitigating the exhaust problem is the use of detachment. Detachment occurs when the plasma exhaust is neutralised (either partially or fully) through interaction with a cloud of neutral gas before it strikes the divertor target [32]. The neutralisation process results in significant amounts of energy radiated isotropically and is therefore extremely effective at reducing the heat flux on the divertor target plates. This is hugely advantageous, and critical to prevent the melting of future reactor divertors.

Maintaining the optimum amount of detachment is crucial: if the detachment front (loosely defined as the upstream end of the strongly radiating exhaust region) gets too close to the core plasma, pollution of the core by impurities can lead to severe degradation of the plasma performance. On the other hand, if the detachment front is too close to the divertor target, the heat load to the target will not be reduced enough to prevent melting. Optimising the divertor configuration and operation to enable the device to withstand the heat loads of the plasma exhaust for extended periods of time is therefore vital for the successful deployment of a commercial fusion reactor, and significant research is underway on this front.

1.4 MAST-Upgrade

MAST-Upgrade (MAST-U) is an upgrade to the Mega Amp Spherical Tokamak (MAST), the UK’s fusion energy experiment based at Culham Centre for Fusion Energy (CCFE) [33]. Its primary objectives are to contribute towards testing reactor concepts (in particular exhaust solutions using a flexible divertor), adding to the knowledge base for ITER, the world’s largest tokamak, currently under construction in France, and exploring the feasibility of using a spherical tokamak as the basis for a fusion component test facility (CTF) for prototyping and testing of future reactor components [34].




Figure 1.5: The magnetic flux surfaces (blue lines) in a tokamak divertor configuration, viewed in the poloidal plane, for the JET tokamak. The tokamak vessel, ports and coils are shown in light grey, with the first wall (the plasma-facing wall) shown in red. The magnetic field null appears as an X-shaped point in the flux surfaces (known as the “X-point”), with the confined plasma following the closed flux surfaces above the X-point. Plasma which lies on flux surfaces outside the last closed flux surface is diverted down into the divertor chamber at the bottom of the device.

1.4.1 Spherical tokamaks

A spherical tokamak (ST) differs from a conventional tokamak (such as that shown in Figure 1.4) in that it has a smaller aspect ratio (the ratio of the major and minor radii of the plasma) [35]. This promises several major benefits compared with a conventional tokamak. The vessel itself is smaller for a given plasma volume, which therefore makes the device cheaper to build. Since capital cost is likely to be a significant fraction of the lifetime cost of a fusion power plant, this makes the ST an extremely attractive proposition.

The compactness also enables good confinement with lower magnetic fields. Producing the strong magnetic fields required to confine high temperature plasmas in reactor designs typically requires superconducting magnets, which are expensive to manufacture and must be cooled to cryogenic temperatures. Achieving a higher plasma temperature with a lower magnetic field therefore makes the plant more economically attractive.

A good measure of efficiency of the device in this regard is provided by the plasma beta, which is the ratio of plasma pressure to magnetic energy density [14]:

$$\beta = \frac{p}{B^2/2\mu_0} \quad (1.4.1)$$

Higher pressures roughly correspond to higher fusion reaction rates (see [14] for details), so tokamaks with higher beta will produce more fusion power for the same cost of running the magnets. The START spherical tokamak at CCFE broke the previous record for the maximum β value obtained in a tokamak, beating it by a factor of 2 [36].

Spherical tokamaks are an exciting proposition for cheaper, more efficient fusion reactors. However, they are not without their challenges. In particular, the tight aspect ratio means there is very little room in the centre column for the solenoid used to induce the plasma current, the toroidal field coils and sufficient protection of these components from neutrons. The latter is a big problem for a reactor, where continuous operation will result in extremely high neutron doses to the components in the central column and likely shorten the life spans of these components considerably. Methods of producing the plasma and driving plasma currents without the use of the solenoid have been studied [37], which would free up a bit of space, but even so this difficulty means that spherical tokamaks will likely see first use as component test facilities where the central column can be replaced regularly, rather than commercial reactors.

The second major challenge, and the one we are most interested in for this work, is that the reduction in major radius reduces the wetted area of the divertor (the area

of the divertor surface which is in contact with the plasma exhaust). The ST still has the same amount of power to exhaust as a conventional tokamak, but the reduced wetted area means that the heat flux at the divertor target is considerably higher than for a similarly-powered conventional tokamak. Dealing with the heat load in a large conventional tokamak is difficult enough; when that heat load is focussed onto the smaller divertor of an ST the problem becomes even more challenging. Fortunately, MAST-U has been designed to explore a possible solution to this problem, with its innovative Super-X divertor, described in Section 1.4.2.

1.4.2 The Super-X divertor

As mentioned in Section 1.4.1, one of the major challenges of spherical tokamaks is dealing with the very high heat fluxes at the divertor target. A novel divertor configuration, known as the Super-X divertor, has been proposed to overcome this difficulty [38]. This type of divertor features a number of enhancements to the magnetic geometry of a standard divertor, which help to dramatically reduce the heat flux at the target.

Significant poloidal flux expansion is a characteristic feature of the Super-X divertor. The poloidal magnetic field is reduced, which causes the magnetic field lines to make many more toroidal revolutions between the last closed flux surface (defining the edge of the confined region) and the wall. This large increase in the connection length facilitates greater cooling of the exhaust plasma. The flux expansion can be seen Figure 1.6, which compares a conventional divertor to the Super-X in MAST-U. The flux surfaces get further apart in the divertor chambers (top and bottom right) as the poloidal field reduces.

As well as increasing the connection length, flux expansion helps to achieve detachment. As described in Section 1.3.3, this is an important phenomenon which can greatly reduce the heat flux to the divertor target using a strongly radiating neutralised region in the plasma exhaust near the target. Flux expansion has been shown to make control of the position of the detachment front easier [40], and so this new divertor geometry provides exciting opportunities to implement feedback loops to optimise the performance of a detached Super-X divertor.

Flux expansion is not unique to the Super-X. Other advanced divertor configurations, such as the snowflake and X divertors, also employ this technique [41, 42]. However, the Super-X divertor adds one further enhancement: it moves the strike points to a larger major radius. Since the plasma wetted area scales linearly with the major radius (and the heat flux is inversely proportional to the wetted area), this increase in the major radius of the strike point has a significant effect on reducing the heat

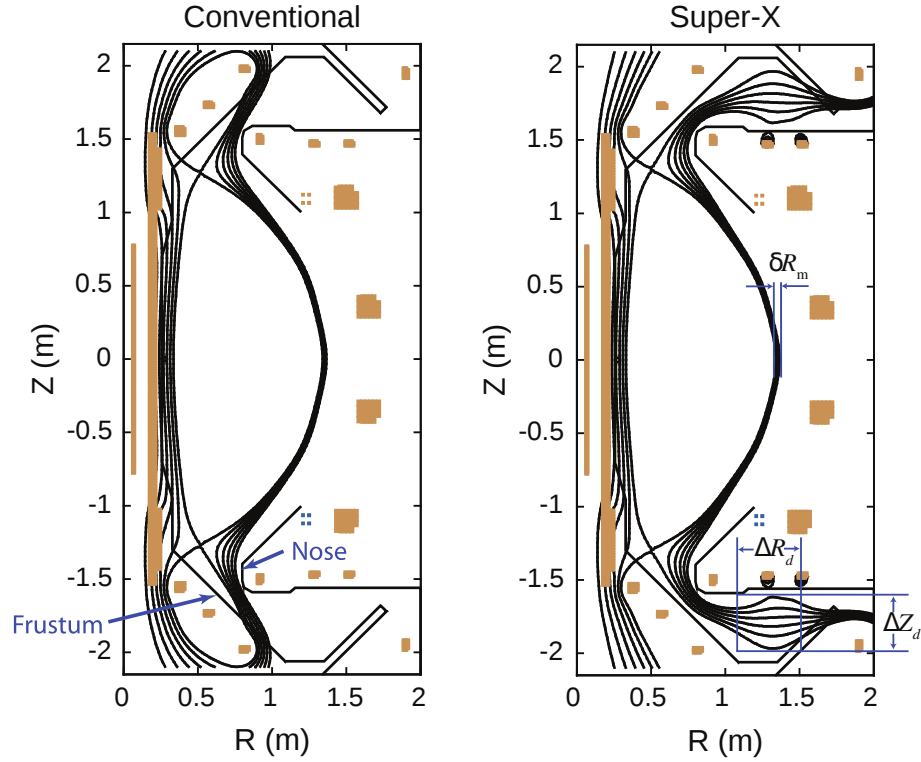


Figure 1.6: Comparison of conventional and Super-X divertor geometries in MAST-U. The flux surfaces are shown as thick black lines, and the first wall in the divertor regions are shown as slightly thinner black lines. Cross sections of some key field coils are shown in brown: two coils in the centre column near $R = 0$ are used to control the separation of the centre column and the inboard side of the plasma, 4 large coils at $Z \approx \pm 0.5$ m and $Z \approx \pm 1$ m are used to shape the core plasma, and numerous small divertor coils between $Z = \pm 1.5$ m and $Z = \pm 2$ m are used to shape the divertor plasma. Flux expansion results in the scrape off layer width δR_m at the mid plane broadening into a region of width ΔR_d and height ΔZ_d in the divertor. The Super-X configuration has the strike points at a larger major radius and features significant poloidal flux expansion compared with the conventional divertor, both of which result in a reduced heat flux to the wall. Reproduced from [39]

flux. In a spherical tokamak with a typically small major radius, the effect can be dramatic. Figure 1.6 shows that in MAST-U the major radius of the strike point is increased from approximately 1 m to 1.6 m, and this 60% increase results in a similar relative decrease in the heat flux.

MAST-U is the first tokamak to be designed specifically to operate with this exciting new divertor concept. It is therefore vital that we have an effective suite of plasma diagnostics to aid in physics studies related to the design, including verifying the predictions made about the divertor's performance and exploiting the large number of shaping field coils in MAST-U to optimise the geometry and operation of the divertor. Both the bolometer and Langmuir probe diagnostics described in this thesis have key roles to play in this work.

1.5 Field-programmable gate arrays

We have seen in Section 1.3.2 that tokamak diagnostics must be able to acquire data at high rates. With tokamak experiments running for many seconds at a time, this equates to very large volumes of data. The acquisition, processing, transfer and storage of that data has in the past required bespoke electronics and powerful computational resources, which are often expensive and consume lots of power and space. It would help greatly if the specific requirements of different diagnostics could be implemented using relatively generic vendor-supplied hardware which can be customised to meet those requirements. Field-programmable gate arrays (FPGAs) promise exactly this capability.

1.5.1 The principle of FPGAs

An FPGA is a configurable logic chip, allowing users to design a circuit in software to perform some desired function, and programme this design onto the FPGA hardware. The architecture used to achieve this is shown in Figure 1.7. The chip contains configurable logic blocks which can be programmed with suitable logic functions including mathematical operations (addition, subtraction and multiplication), multiplexers and registers (i.e. memory). These are connected by a configurable interconnect, and by programming the switch boxes appropriately the different logic blocks can be connected to produce a digital circuit with the functionality required of the chip. This architecture means the FPGA vendor can ship a generic device in bulk, but customers still end up with a highly specialised and optimised device after programming.

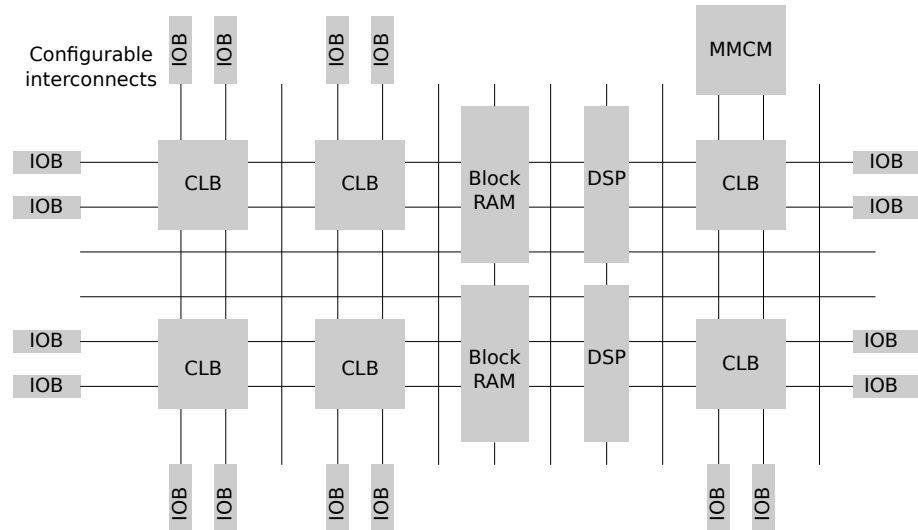


Figure 1.7: FPGA architecture, showing the configurable logic blocks (CLBs), configurable interconnect and the input/output (I/O) buffers (IOBs). The logic blocks implement user-programmable logic functions, and the interconnect is programmed to connect these logic blocks together to realise the required functionality. Modern FPGAs also contain specialised components such as high density memory (Block RAM), digital signal processing (DSP) blocks and clock management tiles such as the mixed-mode clock manager (MMCM). Adapted from [43]

Programming is typically done using a hardware description language (HDL) such as VHDL [44] or Verilog [45]. These describe at a relatively high level the functionality of the circuit, and then a synthesis tool is used to translate the HDL into the required logic blocks and interconnects. Vendor-supplied software then maps this onto the resources available on a specific FPGA chip and produces a bitstream, a file which is loaded onto the FPGA to configure it at power-up. Using software to design the system makes it significantly simpler and typically cheaper than a bespoke analogue system. The chips are fabricated with technology similar to modern central processing units (CPUs) and so can hold many millions of logic gates, which also means that FPGA-based systems can be significantly more compact than an equivalent analogue solution.

FPGAs are reconfigurable, meaning that any upgrades or bug fixes to the design can be easily implemented without having to order new hardware. This provides much more opportunity for development than an analogue circuit design, or even an application specific integrated circuit (ASIC). Although it is possible to develop a more optimised design using an ASIC (since the hardware is not generic), once the design is programmed onto the ASIC it is fixed. As such, any bugs that are later found mean that new hardware has to be produced. Since ASICs or analogue circuits lack the flexibility required in an experimental environment, FPGAs are a

promising avenue for development in fusion plasma diagnostics.

The parallel architecture of an FPGA makes it suitable for extremely high performance data processing for certain applications. For example, digital signal processing (filtering, Fourier transforms etc.) can be performed on an FPGA at a higher rate than a single-thread CPU could manage, even at the lower clock speeds at which FPGAs run compared with CPUs. This results in significant power and cost savings if substantial processing of raw data is required, which is an important advantage of FPGAs compared with CPUs. As an added bonus, many applications require large amounts of raw data but produce much smaller quantities of relevant processed data. By performing this processing on an FPGA chip as it is produced, the data transfer and storage requirements are greatly reduced.

However, FPGAs are not necessarily suitable for all applications. Programming an FPGA using a hardware description language is extremely time consuming, even though there are some vendor-supplied and third party tools available to help with this. The generic nature of FPGAs means that the design is never fully optimised, so faster or more complex designs could be implemented with more specialised hardware at the expense of the FPGA's flexibility. The low clock speeds of FPGAs compared with CPUs means that tasks which cannot be effectively parallelised run more slowly on an FPGA than on a CPU. And importantly, the learning curve for FPGA development is considerable, so significant staff investment is required to implement more complex functionality.

For these reasons, FPGAs are typically used as one component of a more complex system. Tasks which do not require high parallel performance or deterministic timing can be performed much more simply by writing software in familiar programming languages or by using widely available libraries. Extremely specialised and demanding requirements may be implemented using custom-built hardware. The FPGA provides a tool to implement a wide range of functionality with good performance and low costs in both development time and capital, and can offer significant improvements in many areas.

A particularly interesting FPGA chip has been produced by Xilinx [47]. Named the Zynq-7000, it combines the power and flexibility of an FPGA with a dual core ARM CPU and a number of standard peripheral interfaces, as shown in Figure 1.8. This means that we can use a single chip to implement functionality best suited to an FPGA, such as hardware interfacing and digital signal processing, and also functionality best suited to a processor. For example, we can run an embedded Linux operating system on the CPU with TCP/IP networking for easy interfacing with the system, and the CPU can control the FPGA functionality. This combines the best attributes of both technologies

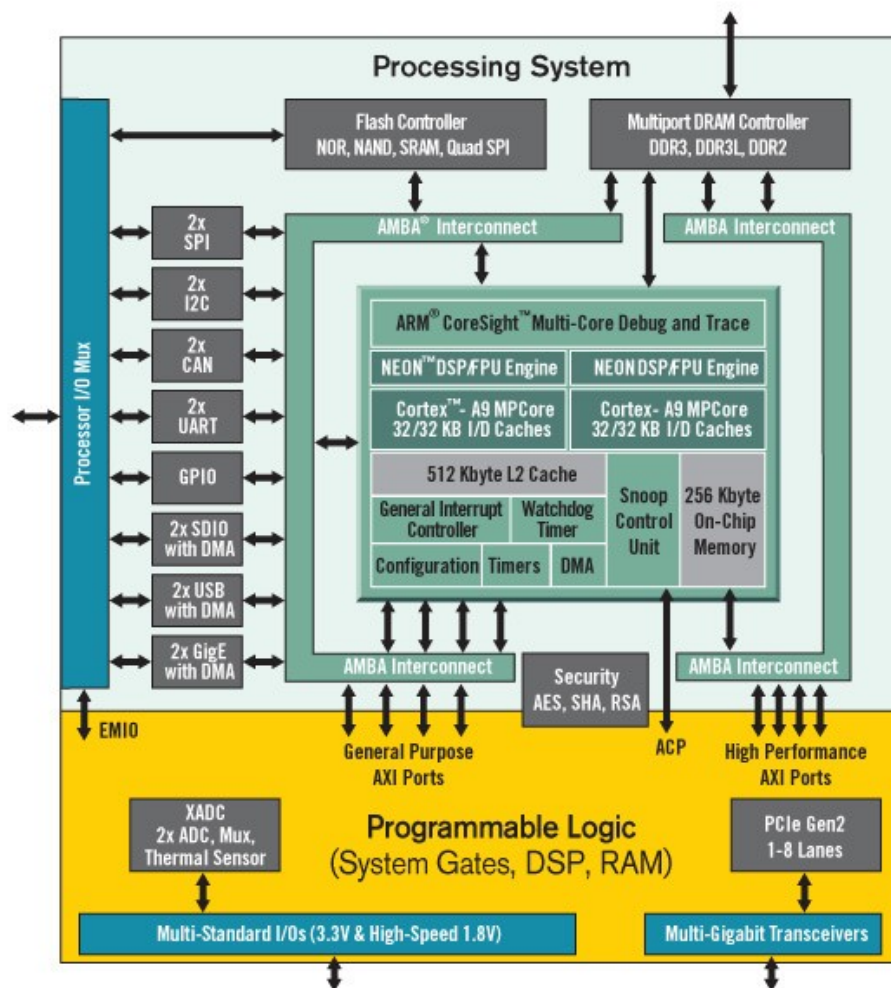


Figure 1.8: Schematic of the Zynq architecture. The FPGA programmable logic is combined with a processing system featuring an ARM dual core CPU, interconnect and peripheral interfaces. Reproduced from [46]

1.5.2 Xillybus

Whilst the Zynq chip mentioned in Section 1.5.1 provides a lot of power and flexibility to the programmer, the complex architecture does make it more difficult to fully exploit the available technology. In particular, transferring data between the FPGA programmable logic (PL) and the ARM processing system (PS) at the high speeds required for high sample rate data is challenging. Implementing the required interface, despite being a relatively standardised one, requires a significant amount of effort and is an additional steep learning curve on top of the more general FPGA development learning curve. There are no vendor supplied or freely available solutions which are sufficiently general to enable a turnkey data transfer system suitable for all required applications.

A commercial solution is available however: Xillybus [48]. It provides an FPGA intellectual property (IP) core to enable high bandwidth data transfer between the Zynq PL and PS using direct memory access (DMA). The IP core (analogous to a software library) is included in the HDL design of the FPGA PL and is paired with a Linux device driver. The interface is through standard first-in, first-out (FIFO) buffers on the PL side, which are an extremely common component in FPGA design and so should be familiar to FPGA engineers, and character device files on the Linux side which can be accessed using the standard system I/O functions (read, write, seek etc.). This greatly simplifies the transfer of data from FPGA to processor and enables developers to concentrate on the application-specific parts of the design instead.

Xillybus was originally designed for communication over PCI-express, for use with FPGAs which did not feature a processing system. However, support for the Zynq chip was added with the Xilinx distribution, an Ubuntu-based Linux distribution which includes the Xillybus drivers and a skeleton FPGA design. We can therefore utilise the power of the Zynq chip without worrying about the difficulties of implementing the DMA-based data transfer from software to hardware.

1.6 FPGAs and fusion diagnostics

It is clear that FPGAs have the capability to help overcome the significant challenges faced by tokamak fusion plasma diagnostics. Their highly parallel architecture and ability to route data quickly make them ideal as data acquisition devices. The flexibility of being able to program many different designs onto the same type of FPGA also makes them ideally suited to controlling the wide range of diagnostics required for a successful tokamak experiment, whilst the generic nature of the chip

and the move from analogue to software-based digital circuit design significantly reduces the cost and space requirements for these highly complex devices. This is particularly important in the research environment of a tokamak, where cubicle space for diagnostics may be limited.

The fact that the FPGA design can be modified during the lifetime of the device is particularly advantageous. The nature of experimental work means that changes to functionality including enhancements and bug fixes are often required, and the ability to reuse the hardware and only make changes to the design in software has numerous benefits. It again reduces hardware costs, and also enables quicker deployment of enhancements by removing the need to replace and re-commission hardware, particularly if only small changes are required. A diagnostic can be “upgraded” by simply installing new firmware in the short time between two experimental shots if need be.

CCFE has therefore chosen to use FPGA technology in a number of its diagnostics. To enable consistency across different systems, the Xillybus IP describes in Section 1.5.2 has been used to transfer data between the FPGAs and CPUs, either embedded in the same chip (on Zynq-based devices) or in a dedicated PC connected to the FPGA with PCI-express.

In this thesis we shall discuss the design and implementation of two such systems. We have produced these two diagnostic systems in order to investigate whether we can produce significantly smaller and cheaper diagnostic electronics by replacing functionality previously provided by analogue components with FPGA technology, and whether this can be achieved whilst providing additional capabilities and without compromising the performance of the diagnostic. Chapter 2 describes the design of an FPGA-based control, data acquisition and processing module for the MAST-U bolometer system, which consists of arrays to measure radiated power in the main chamber and the Super-X divertor. Due to the ongoing construction of MAST-Upgrade, the system was tested at several other sites and the results are presented in Chapter 3. We then describe the design of a new Langmuir probe module to measure plasma particle and heat fluxes in the divertor in Chapter 4, and again present results from the installation of the system on two other fusion-relevant devices in Chapter 5.

Chapter 2

Design of the Bolometer diagnostic

In this chapter we describe the principles and design of the bolometer diagnostic. The physical basis of the diagnostic is explained, as well as the instrument's requirements. We then proceed to describe the design of the system, including the hardware, FPGA firmware and associated software.

2.1 The bolometer principle

Measurements of the power loss due to radiation in fusion experiments is crucial to understanding both the physics and engineering challenges associated with high temperature plasmas. In many such plasmas, the radiated power can be of comparable magnitude to the power deposited by the plasma at the strike points [29]. Indeed, for high power fusion devices such as ITER and DEMO, a high fraction of the exhaust power must be radiated away in order to avoid melting parts of the vessel [40].

Emission from the plasma is dominated by bremsstrahlung in the fully ionised core and line radiation from partially-ionised impurities and hydrogen near the edge [14]. Whilst in principle all emission sources could be measured using spectroscopy and integrated to calculate the total radiated power, this is impractical due to the large number of different energies at which radiation is emitted. It is far simpler to use a diagnostic which is able to measure the total radiation over the wide range of energies that are emitted from the plasma, with a sensitivity which is reasonably constant across this range. We use bolometers for this purpose.

There are several different designs of bolometer, each with their own advantages and disadvantages. Absolute extreme ultra violet (AXUV) photodiodes are often used, as they have a very fast (sub- μ s) response time. However, these photodiodes suffer

from low responsivity at energies below around 100 eV, which are important at the plasma edge and in the exhaust region [49].

Resistive bolometers are also used, which have a flat responsivity for radiation energies ranging from visible and UV radiation to soft X-rays, and also benefit from improved neutron hardness [50]. These sensors are typically slower, with characteristic response times of several hundred ms. Similarly, infrared imaging bolometers feature a foil which is heated up by the plasma radiation, and the temperature rise of the foil is measured by imaging the infrared emission with a camera. Whilst this technique potentially allows for a large number of lines of sight by treating the foil as an array of pixels, early designs demonstrated a much lower sensitivity limit than that achievable with resistive bolometers [51]. The good responsivity profile and high sensitivity makes resistive bolometers a good choice for MAST-Upgrade’s bolometer system.

2.1.1 Resistive bolometers

Resistive bolometers consist of a metal foil absorber in thermal contact with a resistor. Radiation from the plasma heats the foil, which in turn heats the resistor and therefore increases its resistance. By measuring this resistance change, and with knowledge of the heat capacity of the sensor and the rate of change of resistance with temperature, we can calculate the power absorbed by the sensor.

Each bolometer sensor provides a single measurement of the total power emitted from all points in its viewing angle. In order to obtain spatially resolved measurements, we therefore need many sensors. Arrays of sensors view the plasma through an array of pinholes to form a bolometer “camera” with many pixels. For good spatial resolution the viewing angles must be small, which means that the sensors only absorb a small fraction of the total plasma radiation. This in turn means the temperature change of the absorber (and hence the resistance of the resistor) is very small. We therefore need a way to measure very small changes in resistance with high precision. The Wheatstone bridge circuit is ideally suited to this kind of measurement [52].

The resistive bolometer sensors used on tokamaks use such a circuit [53]. By applying a large voltage V_{dr} across one diagonal of an almost-balanced bridge and measuring the small voltage V across the other diagonal, we can calculate the bridge imbalance. Figure 2.1 shows how the bolometer sensor does this in practice. The sensor features 4 resistors in a Wheatstone bridge configuration. Two resistors on opposite sides of the bridge are arranged so that they are in thermal contact with the foil absorber and so heat up when radiation is incident on the sensor: these are the “measurement” resistors R_m . The other two resistors are shielded from the radiation and so maintain




Figure 2.1: The bolometer circuit, a Wheatstone bridge, showing the measurement and reference resistors R_m and R_r respectively. The measurement resistors are in thermal contact with the absorber foil which is exposed to the plasma, whereas the reference resistors are shielded from plasma radiation.

constant temperature (and therefore resistance): these are the “reference” resistors R_r .

When there is no radiation incident on the sensor, the R_m and R_r resistors have the same resistance R , the bridge is balanced, and we measure 0 V output. When radiation heats the absorber the measurement resistors are heated, and their resistance increases by some small amount ΔR . This unbalances the bridge, and results in a finite measured voltage V which is given by:

$$\begin{aligned} \frac{V}{V_{dr}} &= \frac{R + \Delta R}{2R + \Delta R} - \frac{R}{2R + \Delta R} = \frac{\Delta R}{2R + \Delta R} \\ &\approx \frac{\Delta R}{2R} \end{aligned} \quad (2.1.1)$$

In the final approximation in Equation 2.1.1 we have assumed $\Delta R \ll R$, which is reasonable as long as the resistance change is small. In principle, only one resistor needs to change to unbalance the bridge. However, having two measurement resistors doubles the output voltage and so improves the signal-to-noise ratio. This is because the noise is typically dominated by electrical pickup in the cables and noise intrinsic to the bolometer electronics, which is unaffected by the addition of a second pair of resistors.

It should be noted that strictly speaking the bolometer only measures directly the temperature rise of the foil. Usually, this temperature rise is due almost entirely to the foil absorbing radiation from the plasma, which is what we want to measure. However, there are other sources of energy which can heat the foil in addition to the radiation. For example, in high neutral pressure environments high energy neutral particles, produced through charge-exchange processes with the plasma, can impact on the sensor and impart their energy, heating the sensor. No distinction can be made between sources of heating, and so if care is not taken to account for the

contribution of these additional heating effects then the power calculated from the measurements may be an overestimate.

2.1.2 AC synchronous detection

It is challenging to measure the very small bolometer output voltage V in the noisy environment of a tokamak. In particular, DC voltage measurements suffer from $1/f$ noise, which dominates over irreducible sources of noise such as Johnson noise and shot noise and is due to fluctuations in the resistivity of the bridge resistors [54, 55]. To overcome this, we apply an AC excitation voltage V_{dr} to the bridge, and measure the amplitude of the output voltage at the same frequency as the excitation voltage. All other frequency components in the output are considered noise and discarded. This reduces the $1/f$ noise and so improves the signal quality.

Consider an excitation voltage with a frequency ω , of the form $V_{dr} = |V_{dr}| \sin(\omega t)$. The bridge output voltage due to the resistance imbalance has an amplitude A and a phase ϕ relative to the excitation voltage:

$$V = A \sin(\omega t + \phi) \quad (2.1.2)$$

We want to obtain just the amplitude A . We can accomplish this by mixing the output voltage with a reference signal of the same frequency as the excitation voltage and integrating over a full number of periods:

$$\begin{aligned} I &= \int_0^{2n\pi} V \sin(\omega t) \, d\omega t \\ &= \int_0^{2n\pi} A \sin(\omega t + \phi) \sin(\omega t) \, d\omega t \\ &\approx \frac{A}{2} \cos(\phi) \end{aligned} \quad (2.1.3)$$

The final approximation is accurate as long as A is approximately constant over the integration period.

This quantity is known as the in-phase component of the signal. It is now independent of the modulation frequency ω (the signal has been “demodulated”), but still depends on the relative phase of the bridge voltage and the reference voltage.

Traditionally, bolometer electronics have used this method with the Wheatstone bridge sensor [53]. To overcome the dependence on the phase, they have measured or calculated the phase difference between the bridge output voltage and the reference voltage and manually applied this phase shift to the reference signal before mixing, which ensures $\phi = 0$ and means I depends only on A . However, this relies on the phase remaining constant for all subsequent measurements, which is by no

means guaranteed since changes to cable stresses and temperatures can affect the phase. With no way of measuring the phase to check that it is indeed constant, this introduces an additional systematic error into the measurements.

This systematic error can be removed by introducing an additional reference voltage of the same frequency as the first, but a quarter of a period out of phase. This gives us the quadrature component of the signal:

$$\begin{aligned} Q &= \int_0^{2n\pi} V \cos(\omega t) d\omega t \\ &= \int_0^{2n\pi} A \sin(\omega t + \phi) \cos(\omega t) d\omega t \\ &\approx \frac{A}{2} \sin(\phi) \end{aligned} \quad (2.1.4)$$

We can now combine Equations 2.1.3 and 2.1.4 to get a measure of the amplitude which is independent of phase:

$$A = 2\sqrt{I^2 + Q^2} \quad (2.1.5)$$

The phase itself can also be trivially calculated from Equations 2.1.3 and 2.1.4:

$$\phi = \tan^{-1} \left(\frac{Q}{I} \right) \quad (2.1.6)$$

This method of using both the in-phase and quadrature components is known as quadrature detection. It has not been used in bolometer systems before, but the digital signal processing power available with the FPGA used in the new system makes this method feasible, and provides a large benefit by reducing the systematic error associated with assuming a constant phase.

2.1.3 Calculation of absorbed power

Once the resistance change has been calculated from the ratio of the measured and applied voltages, we can calculate the temperature rise which gives rise to that resistance change. The temperature rise is determined by the change in resistance per degree temperature change, θ , defined as $\theta = \alpha R$ where α is the temperature coefficient of resistance and R is the resistance. Thus the change in the temperature T_m of the measurement resistor for a given resistance change is:

$$\Delta T_m = \frac{\Delta R}{\theta} = \frac{2}{\alpha} \left(\frac{V}{V_{dr}} \right) \quad (2.1.7)$$

In Equation 2.1.7 we have used Equation 2.1.1 to relate the resistance change to the output voltage. The change in temperature of the measurement resistor is itself

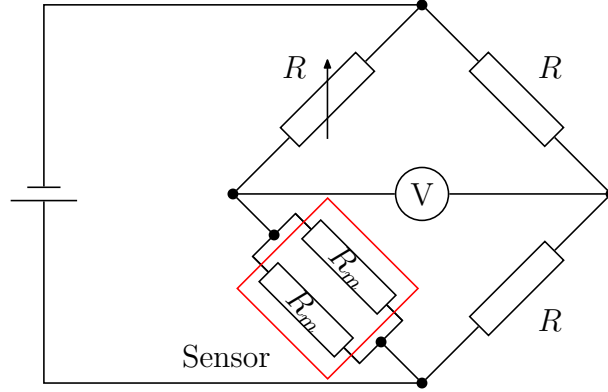


Figure 2.2: Schematic of the circuit typically used to calibrate bolometers. The bolometer sensor, shown as the red rectangle, has its reference resistors short circuited and is placed into a second Wheatstone bridge. Two resistors R match the parallel resistance of the two measurement resistors, and a variable resistor is used to fine tune the bridge balance. DC heating is applied and the bridge output voltage measured.

determined from the incident power:

$$C_m \frac{d\Delta T_m}{dt} = P_{bol} - C_m \frac{\Delta T_m}{\tau} \quad (2.1.8)$$

Here, C_m is the heat capacity and τ is the cooling time, the time taken for ΔT_m to reduce to $1/e$ of its original value once the heating power has been removed.

We can now combine Equations 2.1.7 and 2.1.8 and solve for the absorbed power P_{bol} . Defining the sensitivity $S = V_{dr}\theta\tau/2RC_m$, the power is given by:

$$P_{bol} = \frac{1}{S} \left(A + \tau \frac{dA}{dt} \right) \quad (2.1.9)$$

Equation 2.1.9 is remarkably simple, depending on only two calibration parameters S and τ and the bridge output voltage. We have used the symbol A for the bridge output voltage to highlight that it is in fact the amplitude output from the AC synchronous detection, and not the raw AC measured voltage, which is used to calculate the power.

2.1.4 Calibration

The traditional method for calibrating resistive bolometers calculates the sensitivity S and the cooling time τ in different ways. To calculate the sensitivity, we simply heat the measurement resistors with a known heating power and measure the voltage rise. In the absence of a well-defined amount of on-demand radiation, Ohmic heating is typically used to perform in-situ calibrations of bolometer sensors [56].

Figure 2.2 shows the circuit used to apply this Ohmic heating. The reference resistors in the bolometer sensor are shorted out, so that current will only flow through the measurement resistors. The shorted sensor is then placed in a second Wheatstone bridge, with 3 other resistors (one of which is a variable resistor) which have a resistance approximately equal to that of the shorted sensor, $R = R_m/2$. A small DC voltage V_1 is applied and the variable resistor is adjusted until the bridge is as balanced as possible, i.e. the measured output voltage $V(V_1)$ is as close to 0 as possible. Then a larger DC voltage V_2 is applied, and the output voltage $V(V_2)$ is measured. The current through each measurement resistor is $I_m = V/R_m$, and therefore the net applied ohmic heating is given by [57]:

$$P_{OH} = \frac{V_2^2 - V_1^2}{4(R_m || R_m)} \quad (2.1.10)$$

We have in Equation 2.1.10 used $(R_m || R_m)$ as the resistance of the sensor with the reference resistors shorted out, which is in fact the resistance of the two measurement resistors in parallel.

We can then simply divide the measured voltage rise $V(V_2) - V(V_1)$ by this Ohmic heating to calculate the sensitivity. However, the sensitivity depends on the voltage used to drive the bridge, which during normal operation is different to the DC voltage used in the calibration. We must therefore scale the sensitivity by the ratio of the normal drive voltage and the DC calibration voltage:

$$S = 2 \times \frac{V(V_2) - V(V_1)}{P_{OH}} \times \frac{V_{RMS}}{V_2} \quad (2.1.11)$$

The RMS voltage V_{RMS} accounts for the fact that an AC excitation voltage is used in measurements and a DC voltage in calibration, and the factor of 2 accounts for the fact that during the calibration the two measurement resistors were effectively just one resistor with resistance $(R_m || R_m)$. In normal measurements these resistors will be separated by a reference resistor and each will have resistance R_m , which doubles the bridge imbalance [58].

The cooling time τ is typically calculated by simply measuring the voltage (which is proportional to the temperature different for a given bridge resistance) as the power incident on the sensor suddenly drops from a finite amount to zero. The voltage reduces exponentially after the power is removed, and the time constant τ is calculated from this exponential decay curve. This method does of course rely on being able to supply a sudden removal of radiated power once the bolometer sensor is installed on an experiment. For tokamaks, the most useful source of this sort of radiation evolution is a disruption, though requiring a disruption in ITER or DEMO to calibrate bolometers is simply not an option, due to the damage to the vessel

disruptions in such large machines would cause.

It is important to note that the calibration procedure is carried out using a DC excitation voltage, whereas measurements in normal operation use an AC excitation voltage and synchronous detection. There will therefore be AC effects in the measurements which are not accounted for in the DC calibration, such as attenuation caused by the resistance and capacitance of cables. Giannone et. al. have derived a more complete version of the bolometer equation 2.1.9 which accounts for the effects of cable capacitance and resistance [57]. However, this equation is much more complex and depends on variables which are measured indirectly, so is in practice quite difficult to work with. Often the simpler form, Equation 2.1.9, is used instead, perhaps with empirically measured correction factors.

2.2 Hardware

The design of the electronics for the resistive bolometers described in Section 2.1 has changed little in recent years from the design described in [53]. An AC signal generator is connected across one diagonal of the bolometer bridge to provide the excitation waveform. Across the other diagonal of the bridge a lock-in amplifier is connected: this consists of a phase shifter to match the phase of the excitation and output voltages, a voltage mixer and a low-pass filter to integrate the signal over multiple periods of the excitation wave. The use of the phase shifter means that only the I component of the signal (Equation 2.1.3) is output, and it is assumed that the phase ϕ remains constant at 0; no independent measurement of the phase is made. The mixer and low-pass filter consist of analogue electronics components, and in some designs multiple filter bandwidths are available by routing the signal through different sets of circuit components corresponding to different analogue low-pass filters. Calibration using the DC method described in Section 2.1.4 may be performed by hand using separate power supplies and voltage meters. Alternatively, some electronics units come with a calibration module which can be connected to the bolometer sensor instead of the amplifier and which calibrate the sensor in a more automated way, although only one channel at a time can be calibrated in this manner.

These bolometer electronics systems are widely used in today's tokamaks, including both JET and MAST (pre-MAST-Upgrade). However, they have a number of limitations. The extensive use of analogue electronics components makes the units large and expensive: the JET electronics requires an entire cubicle for 24 bolometer channels. Analogue circuit components are often prone to failure, particularly as

the systems age, and this can seriously affect the availability of the diagnostic. Furthermore, there is only limited flexibility afforded by these analogue components: changing between a pre-determined set of filters may be possible (at the cost of additional analogue circuitry to provide these filters), but these filters must be specified at the time the hardware is ordered. The AC excitation frequency is also fixed; this precludes the possibility of shifting the frequency to avoid strong sources of line noise in the tokamak environment.

We sought to determine whether we could produce significantly more compact bolometer electronics, at lower cost, using FPGA technology. We also investigated how the use of FPGA technology could make the device more versatile, for example by providing much finer control over the low-pass filter used and the ability to vary the excitation frequency. Furthermore, we wanted to know whether the use of digital, rather than analogue, signal processing using the FPGA would enable bolometer electronics to be integrated into future real-time control systems.

The new bolometer electronics has been developed in a partnership between Culham Centre for Fusion Energy (CCFE) [59] and D-TACQ Solutions Ltd [60]. The system is modular, and is based upon an 8-channel hardware module called the BOLO8 which features digital-to-analogue (DAC) and analogue-to-digital (ADC) converters to supply and measure the bridge voltages. The BOLO8 uses an ELF connector, a D-TACQ extension of the FMC standard [61], to connect to a D-TACQ ACQ2106 carrier board supporting up to 6 ELF modules. This carrier features a Xilinx Zynq 7-Z030 FPGA with integrated ARM processor, 1 GB of DDR memory, gigabit Ethernet and 4 high-speed SFP serial fibre ports. The hardware is shown in Figure 2.3. During development of the prototype BOLO8 an ACQ2006B carrier was used instead of the ACQ2106, since the latter was not ready for production at that point. The ACQ2006B features a smaller 7-Z020 FPGA and does not have the SFP fibre ports, but was sufficient for testing the prototype BOLO8.

Figure 2.4 shows how the various parts of the system interact. The FPGA stores the AC excitation waveform, a sine wave, which is converted into an analogue voltage by a DAC on the BOLO8 module and drives the Wheatstone bridge in the bolometer sensor. The (AC) output voltage due to the bridge imbalance is then digitised by an ADC on the BOLO8, and the AC synchronous detection is performed on the FPGA. The calculated amplitude A and phase ϕ are transferred to a Linux operating system running on the ARM CPU, and stored in RAM. A web server runs on top of the operating system, facilitating communication with the system using the standard HTTP protocol. This allows configuration, status checks, firmware upgrades and system debugging to be performed, and also enables readback of the data stored during each shot. The clock and trigger from the central timing system are fed in




Figure 2.3: Top: the BOLO8 module, showing the 4 RJ-45 connectors to connect to 8 sensors, the long black ELF connector for connection to the ACQ2106 and 8 of the inductors used in the LC filters. The 8 channel module is extremely compact compared with previous analogue generations. Bottom: 6 BOLO8 modules in an ACQ2006 carrier, providing 48 channels in a 1U, 19-inch form factor. Reproduced from [62]

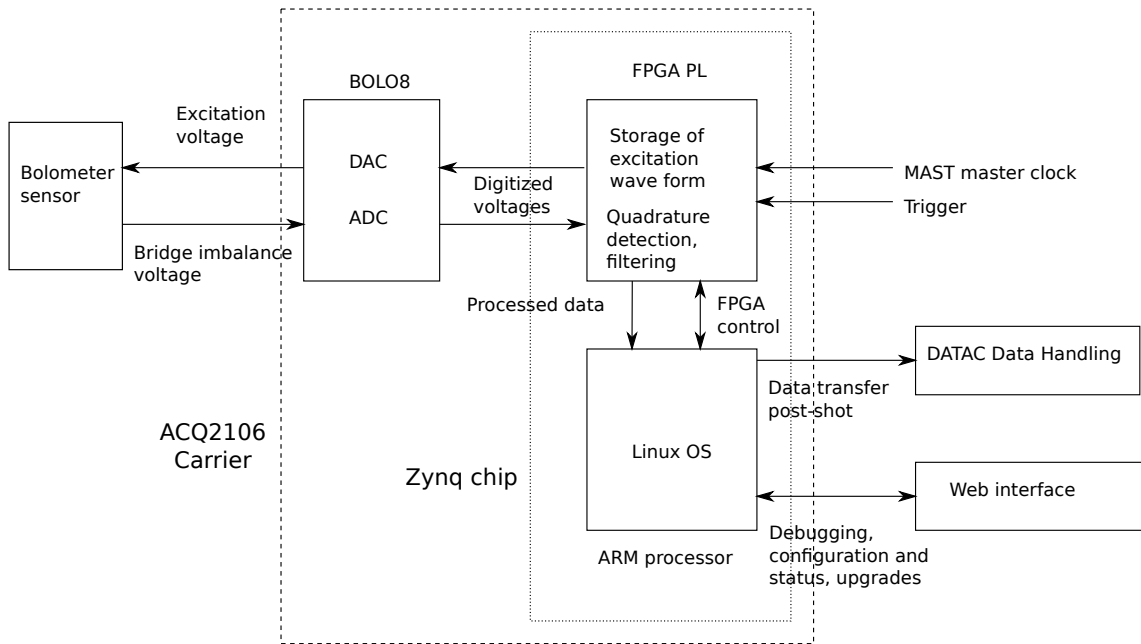


Figure 2.4: Diagram showing the interaction between the bolometer sensor, BOLO8 module, FPGA programmable logic, ARM CPU and MAST data acquisition infrastructure (DATAAC, clock and trigger). The FPGA logic and the ARM processor of the Zynq chip are contained in the D-TACQ ACQ2106 carrier, which hosts several BOLO8 modules. The arrows show the direction of data flow.

to the FPGA directly, exploiting the deterministic timing of FPGA programmable logic to ensure an accurate time stamp for each voltage sample recorded during the shot.

The ACQ2106 carrier is an existing D-TACQ product, but the BOLO8 module was developed specifically for the MAST-Upgrade bolometer diagnostic. Figure 2.5 shows a high level schematic of the module, together with its connections to the Wheatstone bridge of the bolometer sensor. In addition to producing the AC excitation voltage with its main DAC (labelled with a sine wave in Figure 2.5) to drive the bridge, and measuring the output voltage with its main ADC (labelled with a V in the figure), the BOLO8 was also required to calibrate all 8 bolometer sensors simultaneously. The procedure it uses, described in Section 2.2.1, requires some additional hardware: an additional DAC to apply a DC voltage offset for heating the resistors (shown as the voltage source on the right of Figure 2.5) and an additional ADC to measure the current flowing through the resistors (labelled I in the figure) in order to calculate the applied Ohmic heating power.

To prevent the offset DAC amplifier's $50\ \Omega$ output impedance from causing a voltage drop across the bridge, and also to reduce the noise induced in the voltage ADC measurement by the offset DAC, LC notch filters were placed between the offset DAC




Figure 2.5: Schematic of the BOLO8 module, showing the DACs, ADCs and offset circuitry and how these connect to the Wheatstone bridge circuit of the bolometer sensor. The ADC for measuring the bridge voltage is marked V. The current flowing through the bridge resistors is measured by an “ammeter” I, consisting of an ADC measuring the voltage drop across a $0.12\ \Omega$ shunt resistor. The main DAC which provides the AC drive voltage is shown as an AC source (although it is also capable of providing a DC offset during calibrations) and the offset DAC is shown as the voltage source V_{OS} . LC notch filters and Schottky diodes reduce the impact of noise from the offset DAC output on the main ADC.

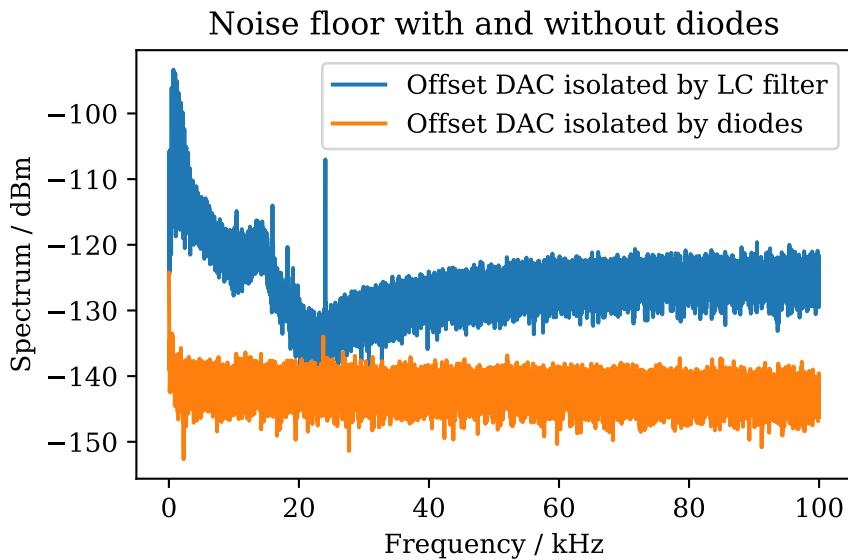


Figure 2.6: Comparison of the spectra of raw ADC voltage measurements, with and without the Schottky diodes isolating the offset DAC. Although the LC notch filter does reduce the noise from the offset DAC around 20 kHz, the additional isolation by the diodes reduces it by a further 10 dB and is effective throughout the entire measurable frequency range. Smoothing using a median filter and a moving average filter has been applied to the spectra to remove individual line noise sources, for clarity.

and the ADCs. These have inductance and capacitance $L = 10$ mH and $C = 6.8$ nF respectively, which provides maximum isolation at 19.3 kHz, close to the 20 kHz design frequency. However, it was found during testing that these notch filters did not provide sufficient isolation to reduce the noise on the main ADC measurement to the 30 nV/ $\sqrt{\text{Hz}}$ specified for the BOLO8. To remedy this, Schottky diodes were added. During measurements, the offset DAC outputs a negative voltage, which reverse biases these diodes and causes them to have an extremely high impedance which provides excellent isolation of the offset DAC from the voltage ADC. During calibration, the offset DAC outputs a positive voltage, which passes through the Schottky diodes with a small voltage drop to bias the bridge. This resulted in a substantial reduction in the noise floor, as shown in Figure 2.6. The same effect could have been achieved by isolating the offset circuitry with mechanical relays, but there was not enough space on the BOLO8 boards for 16 of these relays, so the much smaller diodes were used instead.

2.2.1 New calibration procedure

The traditional bolometer calibration procedure described in Section 2.1.4 has several disadvantages. Shorting out the reference resistors typically requires manually plugging the sensor into a separate circuit board. This is time consuming, and manually balancing the calibration bridge can be very fiddly. In addition, it is difficult to accurately gauge the final voltage $V(V_2)$: the voltage does not really level off after the bridge has reached equilibrium, because the sensor housing starts being heated and this further increases the output voltage.

These factors make manual calibration of the bolometer sensor a difficult and time consuming operation, which in turn reduces the frequency with which the sensors are calibrated and adds additional uncertainty to the calibration parameters. Some of these problems can be mitigated by using an automated in-situ calibration unit, but this requires more complex electronics and still requires physically modifying the circuit, which leads to higher risk of mechanical failure. It also does not solve the biggest source of uncertainty in the calibration: the fact that the calibration is performed with a DC excitation voltage but measurements are made using an AC excitation voltage. Thus the AC effects present when the bolometer is operated normally are not actually measured by the calibration, and have to be estimated and corrected for.

The new BOLO8 solves these issues with a new calibration procedure, which is shown in Figure 2.7. Instead of physically shorting out the reference resistors, a DC voltage offset V_{DC} is applied across both diagonals of the bolometer bridge circuit.




Figure 2.7: Electrical setup of the BOLO8 calibration procedure, showing the application of a DC voltage V_{DC} in addition to the AC excitation voltage V_{AC} . The DC voltage is applied by both DACs such that there is only a voltage difference across the measurement resistors: current flows through these and Ohmically heats them, whilst the reference resistors are unheated.

The voltage is such that there is zero voltage difference across the reference resistors (so no current flows through them), but there is a voltage difference V_{DC} across the measurement resistors which causes a current to flow through and Ohmically heat them. In addition to the DC bias, the AC excitation voltage V_{AC} is applied as usual, and the bridge output voltage amplitude is measured as the measurement resistors R_m are heated.

The DC bias is provided by the BOLO8's offset DAC across points 1 and 4 in Figure 2.7, and by adding a DC offset to the main DAC between points 2 and 3. The voltage drop across the diodes is accounted for by outputting a higher voltage with the offset DAC than the main DAC, such that after the voltage drop across the diodes the DC voltage measured by the main ADC is the same as that programmed into the offset DAC. As an example, with a 1 V DC bias V_{DC} at the bridge, the offset DAC is required to produce 1.51 V, though the additional voltage required due to the diode drop varies with the desired voltage V_{DC} .

Rather than rely on knowing the value of the measurement resistors R_m , we directly measure the current flowing through the bridge with a dedicated current-measuring ADC, labelled I in Figure 2.7, in order to calculate the applied Ohmic heating. The position of the current-measuring ADC is such that it measures the current I_{DC} through only one of the resistors R_m , and hence half of the total Ohmic heating. The total heating power applied by the DC voltage is therefore given by:

$$P_{OH} = 2V_{DC}I_{DC} \quad (2.2.1)$$

Using this and the measured bridge voltage $\mathbf{V} = Ae^{i\phi}$, which consists of the amplitude A and phase ϕ , calculating the sensitivity is simple:

$$\begin{aligned} S &= \frac{\mathbf{V}(V_{DC}) - \mathbf{V}(0)}{P_{OH}} \\ &= \frac{\mathbf{V}(V_{DC}) - \mathbf{V}(0)}{2V_{DC}I_{DC}} \end{aligned} \quad (2.2.2)$$

Note that we must subtract the complex bridge voltages, rather than just the amplitudes. This is because the any measured voltage in the absence of applied power, due for example to cross talk in the cables and intrinsic mismatches in the bridge resistances, may have a different phase to the actual signal of interest which must be accounted for.

By measuring the bridge output voltage after the Ohmic heating is switched off, we can calculate the cooling time. This is done the same way as in the traditional calibration method, by fitting an exponential decay to the cooling curve, but has the advantage that it does not require an external source of power to be applied and suddenly removed (such as a tokamak disruption).

Importantly, this calibration procedure is performed with an AC excitation voltage, unlike the traditional DC calibration method. This means that AC effects due to resistance and capacitance in the system are already present in the measured output voltage, and therefore do not have to be accounted for with correction factors. The sensitivity and cooling time constants calculated with the AC method therefore correspond more directly to the values in Equation 2.1.9 than the corresponding values calculated with DC methods.

2.3 FPGA design: Xillybus version

As shown in Figure 2.4, the FPGA chip in the ACQ2106 carrier is responsible for producing a digital version of the excitation voltage, and performing the AC synchronous detection in quadrature on the measured bridge output voltage. It is also responsible for managing the production of the DC bias voltage during calibration, synchronised between the main and offset DACs, and must handle the SPI interfaces to the main DAC, offset DAC, main ADC and current-measuring ADC. Finally, the FPGA programmable logic must be able to communicate with the ARM CPU in order to transfer measured data and get and set control and status registers.

D-TACQ initially provided an FPGA design for the ACQ2106 (and ACQ2006B) which simply read raw data from the main (voltage) and current-measuring ADCs and allowed writing data to both DACs through a basic Linux operating system

running on the ARM CPU. However, proper synchronisation of the two DACs on each BOLO8 for the calibration, and the implementation of the AC synchronous detection, required adding additional logic to the FPGA design. The D-TACQ firmware architecture design as it was when development of the bolometer firmware started would have required extensive modification to support this extra functionality.

To achieve consistency with other FPGA-based diagnostics being developed at CCFE, the decision was made to instead base the FPGA design around the Xilinx infrastructure [48]. This is an Ubuntu-based Linux distribution for Zynq devices which uses the Xillybus IP FPGA core to provide high speed data transfer between the ARM PS and FPGA PL. More details about the Xillybus IP are given in Section 1.5.2.

Xilinx is supported out of the box on several Zynq-based FPGA boards, and comes with a skeleton design which demonstrates the data transfer capabilities of Xillybus using its streaming and memory mapped interfaces. Porting the system to the ACQ2006B and ACQ2106 carriers required modifying the processor configuration in the FPGA design, generating new first and second stage bootloaders with the new processor parameters, modifying the device tree which describes the hardware for the Linux kernel to include additional hardware present in the D-TACQ carriers and finally re-compiling the kernel. The modifications of the device tree and kernel were required because the D-TACQ carriers contain several inter-integrated circuit (I²C) devices, responsible for controlling things such as fan speed and power supplies, which are not present in the Xilinx-supported boards. Additionally, the BOLO8 itself contains I²C devices to control ADC gains, main DAC voltage output ranges and enabling and disabling both the main and offset DACs. These I²C devices were required due to the limited number of pins available on the BOLO8's ELF connector. Some of the drivers for these devices were not available in version 3.12 of the Linux kernel, which Xilinx uses, so version 3.14 was used instead. This required further changes to the device tree, since some of the device tree syntax changed between these two kernel versions.

After some effort, and with support and advice from D-TACQ, Xilinx was successfully ported to the ACQ2006B and later the ACQ2106. We could now add the FPGA logic to turn a simple DAC/ADC board on a carrier into a working bolometer system. Development was done using Xilinx's Vivado development environment, using a mixture of VHDL code and Xilinx-provided IP cores. The processor configuration and PS-PL communication infrastructure were implemented using Vivado's IP-integrator tool. D-TACQ had already provided the VHDL code for handling the SPI interfaces of the DACs and ADCs, and this code could be plugged in to the Xillybus design with minimal modification. However, the logic to supply the required data to the main and offset DACs and to receive, process and transfer the

data output by the ADCs still had to be provided.

To simplify the code and ease maintenance, the design was split up into different modules, each responsible for only a subset of the required functionality. The modules and their connections are shown in Figure 2.8. Some modules are simple, whereas others are more complex and contain multiple sub modules. The functionality of each module is as follows:

PS Block Diagram

This module contains the Zynq Processing System (PS) instance and the AXI infrastructure required to connect it to the Xillybus core. It is made up of high-level IP cores designed in Vivado's IP integrator.

Xillybus IP

This is the Xillybus IP core which interfaces between the processor's AXI bus and the Xillybus streaming and memory-mapped interfaces. It runs in the processing system's clock domain, which is separate from the rest of the FPGA modules, so any communication between the Xillybus IP and the control and data signals must be done through clock crossing buffers such as FIFOs or block-RAMs.

Xillybus Interface

This module receives the streams of processed voltage and current data and interfaces to the Xillybus IP through standard FIFO buffers with independent read and write clocks.

Control Interface

The control module handles the Xillybus memory-mapped interface used for getting and setting control registers and reporting the status of the system. Control signals are buffered in a block RAM, which ensures they cross from the Xillybus clock domain to the FPGA PL clock domain safely. A small finite state machine (FSM) is responsible for synchronising the contents of the RAM, which are written to and read from by the processor via Xillybus, with the control registers connected to other modules on the FPGA.

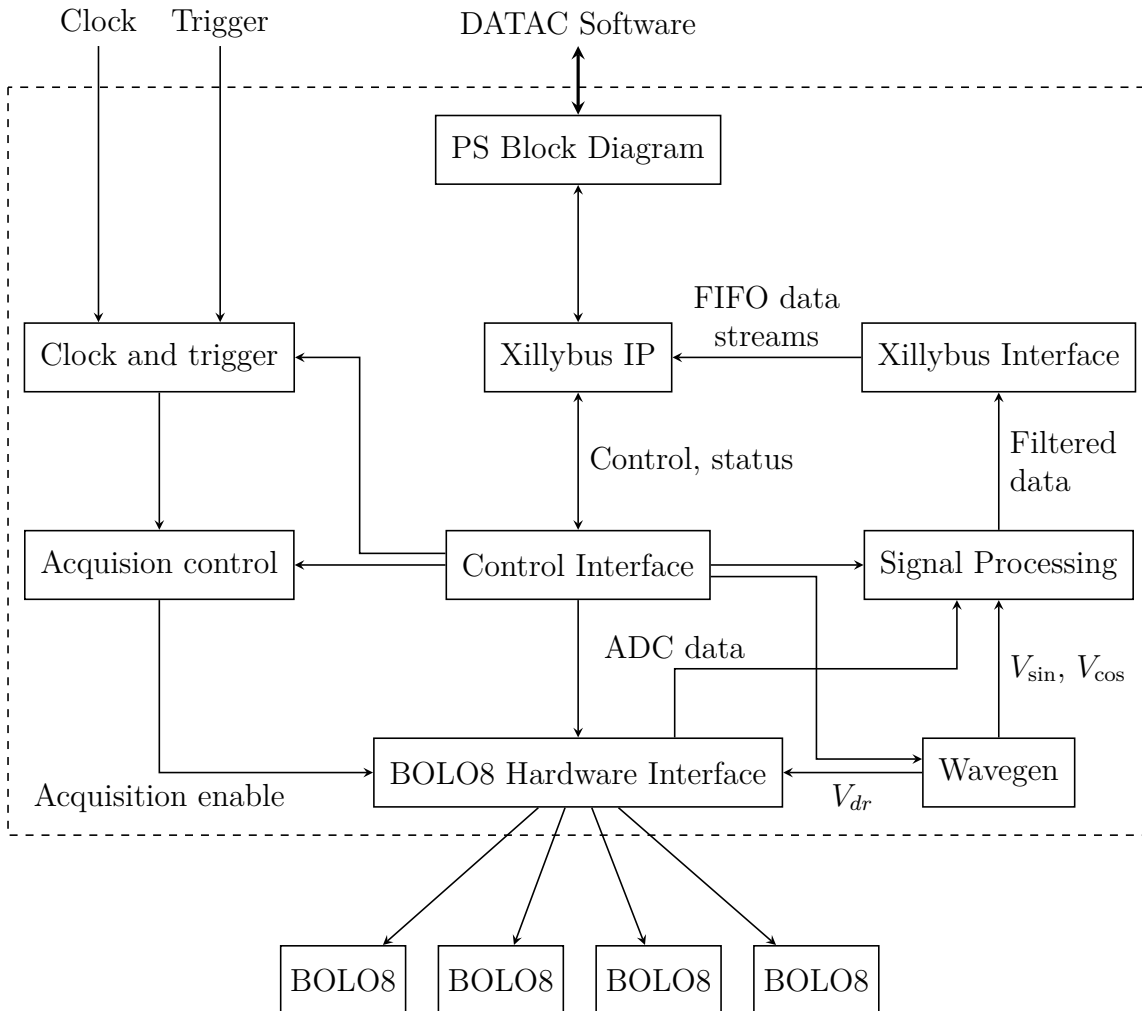


Figure 2.8: The modular design of the FPGA firmware, showing which modules send data to and receive data from which other modules. Although there is much two way communication between modules, the arrows show the logical flow of data between the modules, and the key pieces of data which are communicated. All the logic on the FPGA chip is inside the dashed box, and the external interfaces to the BOLO8 cards, DATAAC data acquisition system software and central timing are also shown.

Clock and trigger

This module takes as inputs the hardware clock and trigger supplied by the MAST-U central timing system. The trigger source is a fibre-optic input from the MAST central timing system, which is converted to an electrical 3.3 V LVCMOS signal using an adapter box and then connected to the ACQ2106 with an HDMI cable. This trigger input is connected directly to the FPGA's programmable logic to enable precise timing of the start of the acquisition. The trigger is buffered to reduce glitches in the signal and to bring it into the same clock domain as the rest of the FPGA logic.

The external clock is similarly sourced from a fibre input from the central timing system and converted to a 3.3 V LVCMOS electrical signal carried over the same HDMI lead as the trigger, directly to an FPGA I/O port. The external clock is buffered using the dedicated MMCM clocking resource of the FPGA, which produces the clocks needed to drive the main DAC (66.6 MHz) and the ADCs and signal processing logic (100 MHz). A Xillybus Lite interface, which provides a simple memory-mapped interface in both software and FPGA logic, is used to configure the clocking logic. It has its own clock, which allows hardware configuration even if the external clock is lost: the system can be switched to use an internal clock in this case. This is intended to be used mainly for development and debugging purposes because losing the external clock will normally also mean losing the external trigger, which will render the diagnostic unusable for MAST-U shots.

Acquisition control

This module receives commands to arm the system for acquisition, or trigger the system through a control register set by software. It sets a control register which enables readback from the ADCs, and hence data output from the signal processing and calibration modules to the Xillybus interface. Acquisition can be manually aborted through software, otherwise it runs until a preset number of samples have been sent to the Xillybus interface.

Wavegen

The Wavegen (waveform generator) module handles the AC excitation and DC bias voltage generation. The AC excitation wave is stored as a sine wave in a block RAM, which functions as a look-up table (LUT). Initially, the sine wave was stored as a 16-bit signed integer in the range -2^{15} to $2^{15} - 1$, which corresponds to ± 20 V output by the main DAC. However, in order to allow applying the DC bias voltage on top

of the AC excitation, the range was reduced by a factor of 18/20. This allowed applying a DC bias of up to 2V without having to reduce the amplitude of the AC excitation whilst still remaining within the output range of the main DAC.

The frequency of the excitation voltage is controlled by the rate at which the look-up table is traversed. The address to read from the block RAM is incremented by a fixed amount each clock period, and the RAM output is registered and sent to the DAC each sample period. For improved frequency resolution, the read address is taken from the top 10 bits of a 27-bit counter. 10 bits were chosen to enable packing the waveform into a single block RAM resource on the FPGA (which can hold 1024 16-bit values) whilst still giving acceptable resolution of the sine wave drive/reference voltage waveform. With a clock frequency f_{clk} and a counter increment step size S , the frequency of the AC excitation wave is given by:

$$f_{dr} = S \times f_{clk} / 2^{27} \quad (2.3.1)$$

The step size S is a 16-bit unsigned integer, and the FPGA clock frequency f_{clk} is 100 MHz. This enables drive frequencies from 0 kHz to 48.8 kHz with 0.75 Hz resolution.

The DC bias voltage is stored in a control register set by software. This module outputs the DC bias to the offset DAC, and adds the DC bias to the AC drive for the main DAC when the bias is positive. Due to the Schottky diodes, no bias is added to the AC drive when the DC bias is negative, and the voltage drop across the diodes when the bias is positive is compensated for by reducing the DC bias added to the AC drive. These calculations must be performed in the FPGA logic to ensure the bias is output to both main and offset DAC simultaneously.

The output of the sine wave LUT is sent to the signal processing module as well as to the main DACs, to be used as the in-phase reference in the quadrature detection. To get the quadrature phase reference wave, a second simultaneous read of the RAM is made. The address for this read is set to the address of the first read plus one quarter of the depth of the RAM, modulo the RAM depth. This produces an output with a phase shift of exactly $\pi/2$ relative to the first output, without using any additional FPGA resources.

BOLO8 Hardware Interface

There is one hardware interface module per BOLO8 connected to the ACQ2106. The MAST-U system has 32 channels, which requires 4 BOLO8 modules. Each module contains the sub modules provided by D-TACQ to handle the SPI interfaces to the ADCs and DACs. This serial interface controls the start of the conversion between

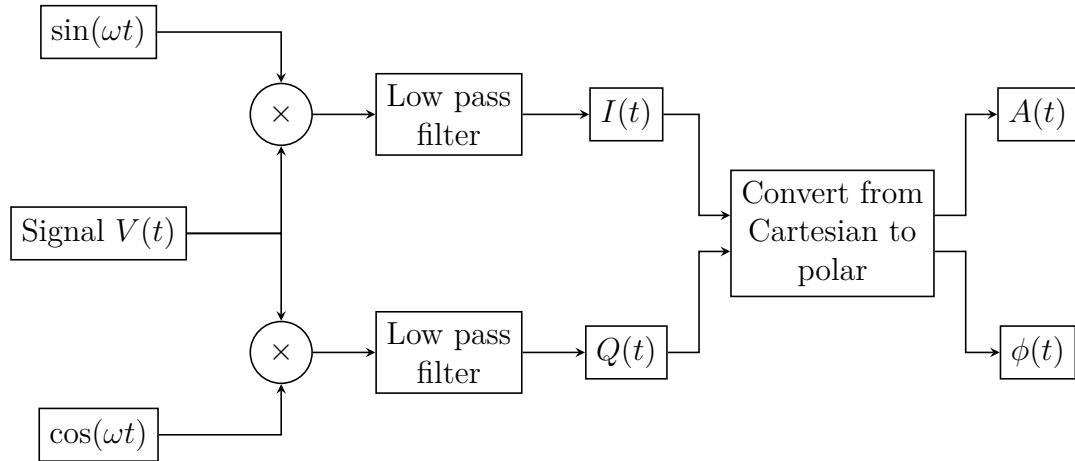


Figure 2.9: Flow diagram showing how the AC synchronous detection is performed on the FPGA. The in-phase and quadrature phase references $\sin(\omega t)$ and $\cos(\omega t)$ are mixed with the voltage signal and integrated using a low pass filter. The conversion from Cartesian to polar coordinates is performed using a CORDIC to output the amplitude $A(t)$ and phase $\phi(t)$.

analogue and digital voltages and the transfer of the digital voltages between the FPGA and the ADCs and DACs. The conversion latency is less than the sampling period for all DACs and ADCs. The main DAC requires a small initialisation sequence on start-up, which is provided by a small finite state machine. The current-measuring ADC's SPI logic was initially designed by D-TACQ under the assumption that it would be polled by software. The interface module contains a small amount of logic to enable sampling the current ADC at a regular 200 kSPS rate instead, for use in the signal processing module.

The data is output from this module to the signal processing module, which is by far the most complex module. It is discussed in detail in Section 2.3.1.

2.3.1 Digital Signal Processing

The digital signal processing module is responsible for performing the AC synchronous detection in quadrature, extracting the DC voltage offset from the raw voltage ADC measurement and outputting the measured DC heating current at the same sample rate as the processed voltages. Data arrives from the 4 BOLO8 sites in the design in parallel, but must be multiplexed into a serial stream for data processing. This is accomplished using a “site multiplexer” module, which takes the data from the four sites, buffers it and outputs a serial stream of all 32 channels from the four sites.

Figure 2.9 illustrates the process by which the AC synchronous detection is performed. The bridge output voltage measured by the main ADC is multiplied by the in-phase (sine) and quadrature phase (cosine) references, which is the digital equivalent of mixing the signal and references. The integration over a number of periods, according to Equations 2.1.3 and 2.1.4, is performed using a digital low pass, finite impulse response (FIR) filter. These filters are built using the Xilinx FIR Compiler IP included in Vivado, and feature reloadable coefficients which allows the filter response (cut-off frequency, roll-off width and ripple) to be customised. Filters can be designed using standard software tools, and the coefficients loaded into the filter before every acquisition if desired. A higher cut-off frequency gives better time resolution in the output data, but also reduces noise suppression. Higher cut-off frequency filters are thus suitable for large, fast events such as ELMs, whereas for smaller but slower radiation fluctuations a lower cut-off frequency is more suitable. The reloadable filter coefficients allow the system to be optimised according to the expected features of interest in the next shot.

For maximum resource efficiency, a single FIR filter is used for the I and Q components for all channels, which are multiplexed through the same path. This requires that the FPGA clock frequency is at least $2 \times N_{chan}$ times the sample frequency, where N_{chan} is the number of physical channels (8 per BOLO8). With a 100 MHz clock frequency and 1 MHz sample frequency, this is enough for up to 50 channels, which is 2 more than can physically be connected to an ACQ2106.

Further efficiency improvements are obtained by downsampling the filtered data. The ADC samples at 1 MSPS, which can easily resolve the AC excitation frequency of around 20 kHz, but the temperature (and hence amplitude) modulations of interest are heavily suppressed above a few hundred Hz due to the finite cooling time of the bolometer sensor. The digital FIR filters are typically set to around 1 kHz cutoff frequency, which is expected to be high enough to resolve all features of interest in normal operation without introducing too much noise. The FIR filters are able to heavily suppress noise above the cutoff frequency (60 dB attenuation in the stop band is typically used), so the significant oversampling of the excitation frequency by the ADCs does not lead to any increase in noise in the output. To reduce the possibility of aliasing, the filtered A and ϕ are sampled at 10 kHz which allows for filter cut-off frequencies of up to 5 kHz. The reduction by a factor of 100 in the sample rate greatly reduces the amount of data output by the FPGA which needs to be stored and further processed, which means physicists wanting to analyse the data save both time and computing resources.

The downsampling enables more efficient resource usage by the FIR filter in the Xilinx FIR compiler core. By combining a chain of downsampling filters, we can

reduce the number of coefficients needed in each filter to provide the requested frequency response. As an example, a filter with 2000 coefficients at 1 MSPS will produce a Kaiser-Bessel windowed sinc low pass filter with a roll-off width of about 1 KHz. A filter this size will use a large amount of FPGA resources, and for systems with large numbers of channels it may not be possible to fit the design onto the FPGA.

Downsampling first to 250 kSPS will allow a filter with only 500 coefficients to achieve the same roll-off width, and a substantial saving on FPGA resources. Half-band decimators, which have a cut-off frequency of $1/4$ the sample frequency (and hence $1/2$ the Nyquist frequency), are particularly efficient since half of the filter coefficients in these filters are 0, and no FPGA resources are needed at all to process these coefficients. The digital signal processing (DSP) system in the bolometer therefore employs 2 half band decimators connected in series, each of which also downsamples the signal by a factor of 2, followed by the final low-pass filter with reloadable coefficients which outputs at 10 kSPS. The extreme resource efficiency here allows us to filter a 1 MSPS input signal down to a few kHz with a roll-off width of only 0.1% of the original sample frequency, which is extremely good for noise suppression, and to do this for as many channels as it is possible to fit onto the available hardware.

The filtered and downsampled I and Q are converted to polar coordinates A and ϕ by a CORDIC [63], which are then sent to the Xillybus interface to be transferred to the processor and stored in RAM ready for readback. Optional offset correction can be applied before the CORDIC stage by storing values of $I(P_{heat} = 0)$ and $Q(P_{heat} = 0)$ (i.e. just due to cross talk and intrinsic bridge imbalance) from the calibration in a look-up table with one entry each for I and Q for each channel. These values are then subtracted from the I and Q of each channel in the multiplexed stream as they exit the final FIR filter. Although the offset subtraction and conversion from Cartesian to polar coordinates can be easily performed in software, they will need to be performed on the FPGA for any future use in a control loop. The real-time power calculation, described below, is intended to be used in this way in future, and so required the offset subtraction and CORDIC rotation to be implemented. Implementing the same processing for the voltage therefore provides a symmetry in the processing which will ease maintenance. Currently, voltage offset subtraction is not applied on the FPGA (the LUT entries for I and Q offsets are all set to 0), but is calculated in post-shot analysis.

The calculation of the DC voltage is very similar to the synchronous detection. It uses the same low-pass filter chain to remove the AC excitation from the output voltage, but does not of course mix the signal with a reference (indeed, mixing with

a reference is what causes the DC offset to be removed in the synchronous detection). Since the FPGA clock frequency is not high enough to multiplex the DC voltage along with the I and Q components through a single filter, the DC voltage follows a parallel path through the filter chain.

The DSP module also contains logic to average the 200 kSPS current measurements, in order to remove the AC component and leave only the current due to the DC bias voltage, which is used in the calibration. The current is typically averaged over the entire heating time when calculating the Ohmic heating (see Section 2.2.1), so high noise suppression using a dedicated FIR filter is unnecessary. For ease of processing, the current is output with the same sample rate of 10 kSPS and the quadrature and DC signals, so all that is needed is averaging to reduce aliasing. As such, the current data is simply summed over 20 input samples and the accumulated result is output at 10 kSPS. This summation therefore gives the mean current multiplied by 20; the multiplicative factor is accounted for in the scale factor used to convert the output data from raw counts to mA.

Finally, the signal processing module contains a filter path to output the calculated power in real-time. The path is exactly the same as that shown in Figure 2.9, but the low-pass filters are replaced by a special deconvolution filter. This FIR filter is capable of both integrating the signal as required in Equations 2.1.3 and 2.1.4, and differentiating the signal as required in the bolometer formula (Equation 2.1.9).

To illustrate this, we note that filtering with the FIR filter is equivalent to performing a discrete convolution between the filter coefficients F and the input signal V . For example, consider the in-phase component:

$$I = F * V \sin(\omega t) \quad (2.3.2)$$

Now, the in-phase component of the power signal is given by:

$$\begin{aligned} P_I &= \frac{1}{S} \left(I + \tau \frac{dI}{dt} \right) \\ &= \left\{ \frac{1}{S} \left(1 + \tau \frac{d}{dt} \right) \right\} I \\ &= \left\{ \frac{1}{S} \left(1 + \tau \frac{d}{dt} \right) \right\} \{ F * (V \sin \omega t) \} \\ &= \left\{ \frac{1}{S} \left(F + \tau \frac{dF}{dt} \right) \right\} * (V \sin \omega t) \\ &= G * (V \sin \omega t) \end{aligned} \quad (2.3.3)$$

We have here used the fact that the derivative and convolution are linear operators, and so their order can be swapped.

It can be seen that the deconvolution filter G is produced by taking a low-pass FIR filter kernel F , adding it to its own derivative multiplied by the cooling time τ and then dividing the result by the sensitivity S :

$$G = \frac{1}{S} \left(F + \tau \frac{dF}{dt} \right) \quad (2.3.4)$$

The filter design can be done using software running on a CPU, which is much simpler and more flexible than division and differentiation on the FPGA. The same filter G can be used to produce the quadrature component of the power, $P_Q = G * (\mathbf{V} \cos \omega t)$, meaning only one filter G is needed for each physical channel. The deconvolution FIR filter implemented in this design has approximately 1 ms latency, which means calculated power values are available on significantly shorter time scales than the characteristic cooling time of the bolometer sensors (typically 10's or 100's of ms in vacuum).

A deconvolution filter has already been used in previous generations of bolometer electronics to calculate the power from the voltage amplitude A [53], since it smooths as well as differentiates the signal and hence reduces the noise level in the output. However, in that case the filter was applied in software post-processing. It would not be practical to design an analogue filter with the required frequency response, particularly when the calibration parameters S and τ are different for each channel and subject to change, and this has limited the ability of previous generations of bolometer electronics to provide a real-time power output. However, the power of the digital FIR filter on the FPGA allows us to design a filter with any frequency response with relative ease, and therefore allows us to perform this deconvolution in real-time in hardware, which is essential for any future integration of the bolometer system into a control loop.

2.3.2 Jetblack software interface

Once the FPGA firmware was designed and implemented, software was required to control it. Initially this consisted of a set of simple shell scripts run manually whilst logged on to the Linux OS, to handle reading back data, setting control parameters and arming and triggering the device. This worked fine for simple lab testing of the device, and was used for the tests with a bolometer sensor excited by a laser and held in a vacuum chamber at the York Linear Plasma Institute described in Section 3.1.

For routine operation on MAST-U, or indeed any other complex experiment, a more sophisticated and automated software package was required though, since it would

be impractical to have somebody manually configure and operate the diagnostic for every single shot. Since each experiment seems to use a different method for communicating with diagnostics, the software will inevitably be slightly different for each site at which it is installed. The first tokamak the system was installed and tested on was JET (see Section 3.2), and so the software was developed with JET in mind.

All JET diagnostics are controlled centrally by the General Acquisition Program (GAP) [64]. Communication between GAP and the diagnostic uses the HTTP protocol [65], including configuration, status reporting, arming the device and reading data back after each JET pulse. The ability to abort the pulse is also required. To reduce duplication of effort, a Python package called Jetblack has been developed by C. Hogben [66], which acts as middleware between the HTTP requests from GAP and the software to control the diagnostic.

Diagnostics are implemented as plug-ins for the Jetblack middleware, and these plug-ins must implement certain functions which are called when the system receives particular HTTP requests. The function `l3state()` reports the state variables of the diagnostic: its configuration parameters and status registers. Similarly, `l3params()` takes configuration parameters as input arguments and sets them on the diagnostic. In order to improve readability, the configuration parameters are given in physical units, which are converted to register settings by the plug-in before being written to the FPGA registers. For example, the excitation frequency f_{dr} is given in Hz and converted to the LUT address counter step size S using Equation 2.3.1, and the DC bias voltage for the calibration is given in V and converted to the required offset DAC register value. The filter cut-off frequencies for the low-pass and deconvolution filters are given in Hz and the filters are designed by the plug-in, the latter only after a calibration is performed to provide the sensitivity and cooling time parameters for each channel's filter, before being loaded into the FPGA's FIR filters.

Other plug-in functions are called at specific points throughout the shot cycle. At the beginning, the initialise command is received, and the software ceases any previous data collection if it is still ongoing (which will only occur if there were errors in the previous shot's acquisition cycle), clears the previously collected data and initialises the data storage for each channel, as required by Jetblack. This means that at the start of each shot the device is in a well-defined state, independent of any control flow errors which may have occurred in the previous shot.

After this the arm command is received. This causes the plug-in to first run a calibration, which consists of arming the device, using the software trigger to start data collection and then applying the DC bias voltage for heating for the configured duration. A small delay before the start of heating enables easy measurement of the

bridge output voltage offset, which can be checked against the offset calculated by fitting the cooling curve.

Once the calibration data has been taken, the cooling curve is automatically fitted to these time traces using a non-linear least squares curve fitting procedure to calculate the calibration parameters. The amplitude and phase time traces during cooling are converted to Cartesian I and Q time traces, each of which is fitted to a decaying exponential:

$$V(t) = V(V_{DC})e^{-t/\tau} + V(0) \quad (2.3.5)$$

The $V(V_{DC})$ and $V(0)$ parameters are used to calculate the sensitivity using Equation 2.2.2. $V(V_{DC})$ is the measured equilibrium voltage with DC Ohmic heating of the sensors (equal to the voltage at the very beginning of the cooling curve), and $V(0)$ is the offset voltage which is measured with no applied heating (the asymptote of the cooling curve). In Equation 2.2.2 the voltages are given as complex values: these consist of the fitted parameters for the real (I) and imaginary (Q) time traces. The cooling time τ is measured by both traces: it should be the same for both I and Q fits, and any difference indicates the uncertainty on the fit of τ .

These calibration parameters are then used to design the deconvolution filters for the real-time power, which are then loaded onto the FPGA, and to apply the offset corrections to the P_I and P_Q signals before they are rotated by the CORDIC. For simplicity no errors are included in the voltage data used in these fits, since the errors are not needed for designing the deconvolution filters and applying the optional offset correction. More complete error analysis and fitting can also be performed in post-processing to obtain more detailed statistics.

Once the deconvolution filter and offset loading is complete, the software opens for reading the Xillybus virtual device files from which the voltage and power data are read, and waits until the FPGA sends the end-of-file signal to close these files at the end of the pulse (or if the pulse is aborted). Finally, the system is armed again, ready for the hardware trigger which signifies the start of the JET pulse. After the end of the pulse, the readback data is organised into separate channels and written to in-memory files, which the Jetblack middleware then serves upon receiving readback commands for each channel from GAP. The system then returns to an idle state in preparation for either a configuration update or the next pulse.

2.4 FPGA redesign: BOLODSP module

The Xillybus-based FPGA design described in Section 2.3 was initially developed because the FPGA firmware supplied by D-TACQ to operate the BOLO8 did not

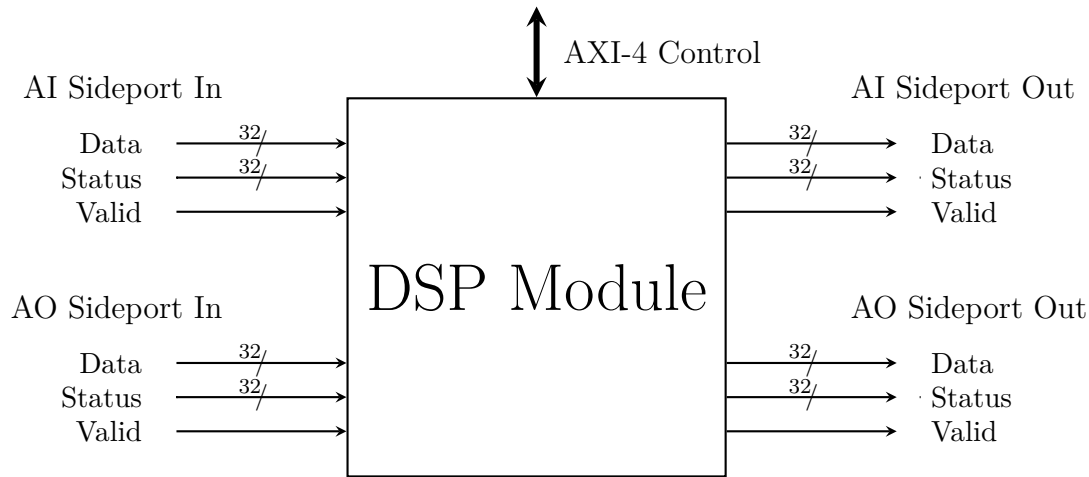


Figure 2.10: The specified user-programmable DSP interface, showing the sideport buses which carry data from the hardware modules and the AXI control bus. Data arrives from the ADCs on the AI Sideport In bus, can be optionally processed by the module, and is passed on to be sent to the CPU on the AI Sideport Out bus. Data to be sent to the DACs arrives from the CPU on the AO Sideport In bus, and is sent out to the DACs on the AO Sideport Out bus.

have provision for the necessary digital signal processing. However, D-TACQ expressed interest in making modifications to their FPGA architecture to enable user-programmable digital signal processing, including for the bolometer application, and this resulted in my collaborative project with D-TACQ in 2015 to implement this. The result was an interface specification for a generic DSP module, which was kept isolated from the rest of the system and so could be written without any knowledge or concern for how the BOLO8 (or indeed any other D-TACQ hardware module) was controlled.

Figure 2.10 shows the resulting interface. Data from the hardware sites is aggregated onto sideports, which are qualified with a data valid signal and a pre-defined bit field status bus indicating the source/destination and type of the data, among other things. Data from ADCs is transmitted on the AI (analogue input) sideport, and data destined for DACs is transmitted on the AO (analogue output) sideport. The DSP module intercepts these sideports, optionally modifies the data according to the signal processing requirements of a particular application, and then passes the data on to the next destination of the sideports. The module can be controlled and configured using an industry standard AXI-4 control bus, which is accessed through software using a simple customisable driver provided by D-TACQ.

With this infrastructure in place, we could adapt some of the modules from the Xillybus firmware to the new interface and use D-TACQ's FPGA architecture instead of a custom designed Xillybus one by writing our own user-programmable DSP

module, known as BOLODSP. The Xillybus interface and the hardware interface modules were no longer needed, nor the clocking and acquisition control, as all this functionality is provided as standard by D-TACQ's existing firmware. Generation of the data for the main and offset DACs must however be performed in the BOLODSP module because of the required synchronisation of the DC offset between the two DACs during calibration, so the Wavegen module from the Xillybus firmware was incorporated into BOLODSP. This meant that the AO Sideport In data was discarded, and the main and offset DAC data generated inside the module was output on the AO Sideport Out bus. The control interface was adapted to use the AXI-4 interface rather than the Xillybus memory-mapped interface.

Most of the complexity in the Xillybus firmware was contained in the signal processing module, and this required only minor modifications to work with the new D-TACQ interface. The stream of input data from the main and current-measuring ADCs on the sideport was slightly different and so some reorganising of the data was needed. For maximum flexibility system was designed to process 48 physical channels, with extra dummy channels automatically added into the data stream if not all 48 channels were present on the sideport. The dummy channels were removed before the data was output on the AI Sideport Out, making the process transparent to the rest of the system. This enabled the same FPGA image to be used on any ACQ2106/BOLO8 system, regardless of the number of installed BOLO8 modules, and the number of channels used could be configured on a shot-by-shot basis, as well as a system-by-system basis, using D-TACQ's standard software tools.

The BOLODSP module contains the same signal processing capabilities as the Xillybus design: the AC synchronous detection in quadrature with FIR filters of configurable bandwidth, averaging of the current measurements to remove the AC excitation variations, calculation of the DC bias component of the voltage and the real-time power calculation using deconvolution filters of configurable bandwidth. With the inclusion of the waveform generation logic it resulted in a very specialised, yet highly configurable, bolometer FPGA module. Setting and readback of the control registers using the D-TACQ driver are performed by custom shell scripts, which integrate with D-TACQ's standard tools for configuration and data collection on a shot-by-shot basis. Custom scripts to run the calibration procedure and design the filters in software are provided, as well as a program to perform least squares fitting of the calibration curves to get the sensitivity, cooling time and offsets. The latter can be run after a calibration to design the deconvolution filters, which are loaded onto the FPGA before the upcoming shot if real-time power output is desired. Re-factoring the bolometer system design to use D-TACQ's DSP interface has numerous advantages compared with keeping the existing Xillybus design. By using

the existing BOLO8 hardware control and FPGA to CPU data communication provided by D-TACQ, the future support requirements are reduced: we are now only responsible for supporting the signal processing module, and any issues with the BOLO8 or generally setting up the ACQ2106 to acquire data throughout a shot can be dealt with by D-TACQ directly. In addition, the user interface for the bolometer is now much more similar to the rest of the D-TACQ product line up. This makes the use of the system at other institutions (many of whom already have D-TACQ data acquisition systems) more attractive, since they are more familiar with how the system works and adaptation to other tokamak infrastructure is simpler. Indeed, the ability to integrate with other institutions has resulted in both EPFL (which operates the TCV tokamak) and PPPL (which operates the NSTX spherical tokamak) purchasing the bolometer system from D-TACQ, and using the BOLODSP module intellectual property under license (at no cost) from Durham University.

2.5 Improvements over previous systems

The new bolometer electronics features several improvements over previous analogue versions. These improvements stem from both the design of the BOLO8 hardware module and use the use of the Zynq FPGA/ARM combination.

The high quality, high bandwidth ADCs on the BOLO8 provide accurate, low noise measurements of the bridge output voltage to be used in the signal processing. The use of digital signal processing on the FPGA greatly reduces the physical complexity of the system, by removing the need to use analogue components for the filtering. Since analogue components are subject to degradation over time, the switch to digital processing improves the reliability and robustness of the system. Furthermore, the use of a digital FIR filter enables a vastly superior frequency response compared with an analogue filter, which results in greatly improved noise attenuation in the output. The frequency response of the filter can be finely configured, with both the bandwidth and attenuation specified to far greater precision than would be achieved using analogue components with finite tolerances. Unlike analogue filters, the time delay through a digital FIR filter is determined exactly by the number of coefficients in the filter and the sample rate, and is independent of the filter bandwidth. This removes any uncertainty in the timing of individual data points caused by variation in the filters used, and further simplifies post processing. The power of the FIR filter has been further exploited by designing a deconvolution filter to both smooth and differentiate the input signals and produce a real-time measurement of the absorbed power, which is impossible to do with analogue filters.

The new calibration procedure made possible by the BOLO8 hardware, and implemented in both firmware and software, is also a significant improvement over the previous procedure. By simultaneously driving the bridge with an AC excitation voltage and Ohmically heating the measurement resistors with a DC bias voltage, we can calibrate the sensor in conditions much closer to those encountered during normal operation. In particular, AC effects such as phase shifts and cable attenuation are already present in the calibration measurements and are therefore automatically accounted for when calculating the calibration parameters. This removes the requirement to apply estimated correction factors to calibration parameters that have been calculated using the DC method, which reduces the systematic errors involved in the measurement.

The calibration procedure is completely automated and requires no mechanical switching to modify the bolometer circuit. It can also be performed on all bolometer sensors connected to each BOLO8 simultaneously, making the calibration of an entire array of bolometers fast enough to be performed on a regular basis, such as before every shot rather than every few months. This not only enables the production of data using accurate, up-to-date calibration parameters, but the large number of calibrations provides better statistics and so reduces the uncertainty on the calibration parameters and hence on the calculated power; this is demonstrated in Section 3.4.3.

Not only does the new system offer tangible improvements in the quality of data, it accomplishes this in a much more compact, smart design. Running a Linux operating system on the electronics unit makes it much easier to configure and monitor the device using standard tools, and provides greater flexibility in the operation of the system and its integration at multiple sites. Implementing much of the system functionality using an FPGA chip rather than bespoke hardware also results in a much more compact system, which fits in a 1U 19 in rack and has a significantly lower cost per channel than previous generations of bolometer electronics. For comparison, the bolometer electronics at JET takes up an entire cubicle for 24 channels.

In the following chapter, we present results obtained with the new electronics at a number of different sites, including targeted lab-based work to improve our understanding of the bolometer measurements and installation on tokamak devices to demonstrate the correct operation of the system. Through these tests we illustrate the improved performance of the FPGA-based system and demonstrate the operational flexibility the system is capable of.

Chapter 3

Results from the Bolometer system

In this chapter, we describe the installation of the newly designed bolometer system at several sites. A brief introduction to the setup at each site is given, and the results of measurements made using the diagnostic are discussed. The primary purpose of these tests ranged from demonstrating the correct operation of the bolometer system in a tokamak setting, to quantifying the absolute accuracy of the system.

For the tests at the York Plasma Institute (Section 3.1), JET (Section 3.2) and TCV (Section 3.3), the Xillybus-based FPGA firmware and software, described in Section 2.3, was used. The successful implementation of D-TACQ's user-programmable DSP interface during my Collaboratory Project (undertaken as part of the Fusion Centre for Doctoral Training programme) in Summer 2015 enabled the development of the BOLODSP firmware, described in Section 2.4, which was no longer reliant on the proprietary Xillybus interface and provided better integration with the rest of D-TACQ's firmware architecture. This makes the BOLODSP design more easily portable to other institutions, and it was this firmware that was used in the PPPL tests (Section 3.4). The BOLODSP firmware design will also be used for the MAST-Upgrade bolometer system for consistency with systems purchased by other institutions.

3.1 Lab-based measurements at the York Plasma Institute

The Revision A prototype of the BOLO8 (which did not feature the Schottky diodes and had lower impedance LC filters) in a D-TACQ ACQ2006B carrier was taken

to the York Plasma Institute, for testing with a gold foil bolometer sensor of the type to be used in MAST-Upgrade [53]. This was the first time during development that an actual bolometer sensor had been used: previously a Wheatstone bridge on a breadboard was used if the drive and measurement voltages needed to be related (e.g. to test the signal processing), otherwise oscilloscopes and waveform generators were used. The sensor was installed in a vacuum chamber to enable studies both in air and in vacuum. Radiated power was simulated with a 5 mW laser, with a wavelength of 660 nm; this monochromatic directional source meant we had a good idea of the power we expected to measure as the reflectivity of gold is well known and the beam angle is small. This work was carried out in March 2015, in collaboration with Matt Reinke during his time at the York Plasma Institute.

Since gold has a reflectance of 96.27% at 660 nm [67], the power absorbed by the foil from the laser was extremely small. This made precise power measurements difficult, particularly with the Revision A prototype, which did not have the isolating Schottky diodes described in Section 2.2 for noise reduction. It did however expose a resolution limitation in the FPGA design, as described in Section 3.1.1. Subsequent tests of a foil bolometer and laser, performed at PPPL and described in Section 3.4, used a blue laser instead of a red one, which resolved the issue of high reflectance.

3.1.1 Comparison of FPGA and software signal processing

To test the AC synchronous detection (described in Section 2.1.2) performed on the FPGA, both the processed output voltage and the raw voltage, directly measured by the BOLO8's ADC, were recorded. The AC synchronous detection was then performed in software with floating-point precision using the raw signal, to produce the amplitude. This was compared with the processed signal output by the FPGA; the two were expected to be identical to within numerical precision. For this prototype system, the filtered data was the same width (16 bits) as the raw ADC input data.

Figure 3.1 shows the comparison between the two processing methods for a calibration shot. This shot was performed by acquiring a fixed 1 s of data, with 0.6 s of ohmic heating achieved by applying a 0.5 V DC bias, and then switching off the DC bias voltage for the remaining 0.4 s and letting the bolometer cool down whilst continuing to measure the voltage. Despite the good quantitative agreement between the processing methods, it is clear that there is significant bit noise in the FPGA-processed data, which does not exist in the floating-point data.

This result was initially surprising, since the BOLO8's ADCs have 16 bit resolution and it was thought that this would therefore limit the absolute resolution of the system, and so the floating-point processing would have the same quantisation error

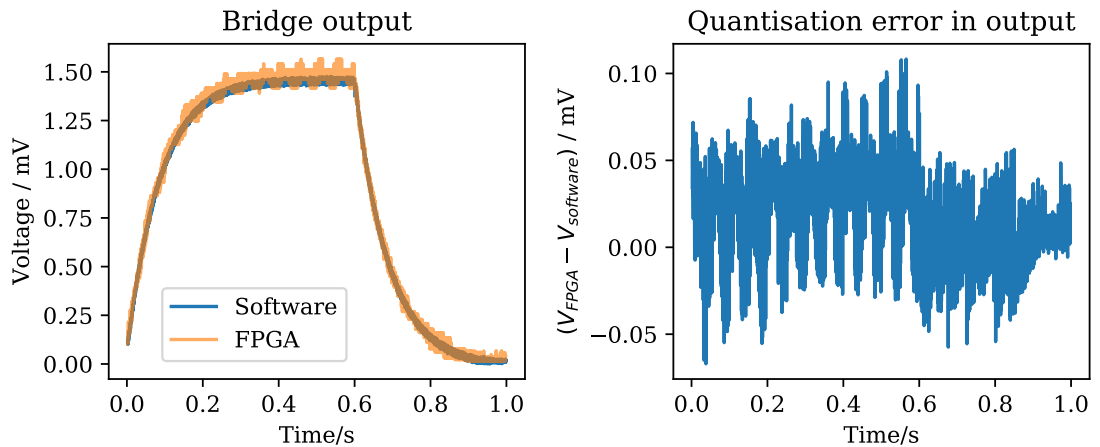


Figure 3.1: Comparison of floating-point (software) signal processing with fixed-point (FPGA) DSP for a calibration shot. Although both methods produce similar results, there is significant quantisation noise present in the fixed-point data which is not present in the floating-point data.

as the fixed-point processing on the FPGA. We can explain this by noting that the quantisation error on each individual sample is randomly distributed. During the filtering process, many successive samples, after being multiplied by a filter coefficient, are added together. This causes the quantisation error to partially cancel out, and results in a smaller quantisation error (and therefore higher bit resolution) for the processed data.

The use of a low-pass filter in the processing also increases the bit resolution. The ADC has a fixed amount of noise, equivalent to just under 1 LSB in the 16-bit measurement, and the noise spectrum is constant across the entire frequency range, at around $40 \text{ nV}/\sqrt{\text{Hz}}$. The total noise-equivalent power is the product of this noise floor with the square root of the bandwidth. By filtering the signal, we reduce the bandwidth, and therefore the noise, which again results in higher resolution than the raw ADC data could provide.

It seemed prudent, therefore, to increase the number of bits in the filtered output data. This was used in part in the design of the multi-stage FIR filters described in Section 2.3. Since today's common CPUs can work with 16- or 32-bit words (but nothing in between), the final processed data should be output with 32-bit resolution. Even though the actual bit resolution is less than 32, there would be no PC memory or computational savings achieved using any output data less than 32 bits wide, compared with simply using 32-bit data. Whilst some resource savings could be made on the FPGA by packing samples (e.g. 24-bit samples into 32-bit words), these savings are minimal compared to the available resources on the FPGA and do not justify the added complexity of packing and unpacking the data.

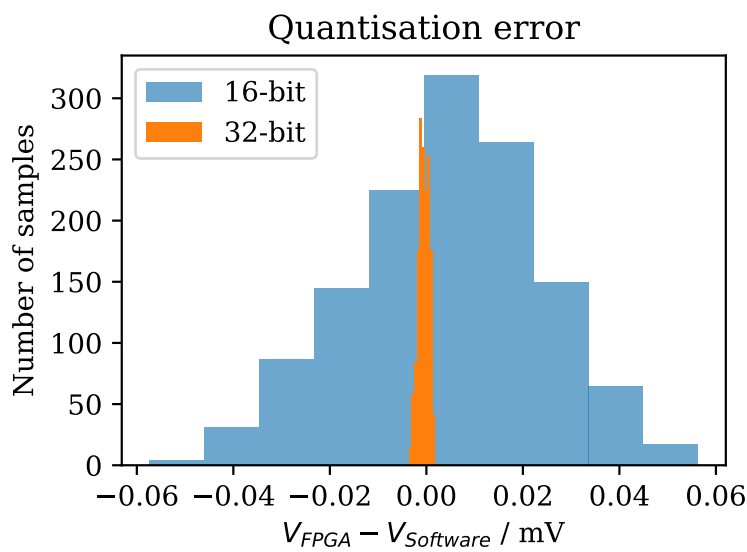


Figure 3.2: Histogram of the quantisation errors using the 16-bit precision filter output shown in Figure 3.1, and using 32-bit precision output. The latter data comes from a calibration run performed as part of firmware development back at CCFE, still using the same bolometer design, with the revised BOLO8 board (including the noise-suppressing diodes). The difference between the fixed-point processing on the FPGA and the floating point processing performed on a CPU is substantially smaller when the data is output with 32-bit precision.

As such, the revised prototype's filters were designed to output the processed data in 32-bit precision, which gave comparable results to the floating point calculation. This can be seen in Figure 3.2, which shows the quantisation errors of the cooling curves for two calibrations. The 16-bit data is taken from the last 400 ms of Figure 3.1, and the 32-bit data from a calibration performed back at CCFE after the York tests as part of further firmware development. The latter run used a 1 V DC bias voltage and 18 V excitation voltage, so the bridge output was larger; the data from this run plotted in Figure 3.2 is taken from the last 450ms of the cooling curve since this is of a similar magnitude to the 16-bit data.

3.1.2 Calibration dependence on bias voltage

A test was also performed to measure the effect of bias voltage on the calculated calibration parameters. As described in Section 2.2.1, the calibration process involves applying a DC bias voltage across both diagonals of the sensor's Wheatstone bridge to ohmically heat the measurement resistors. Increasing the DC bias voltage produces more ohmic heating, which in turn results in a bigger signal to be used in the calibration calculation. Whilst in theory we would want the largest possible signal (for the best possible signal-to-noise ratio), we must ensure that the measured calibration constants are not changed when using significantly larger ohmic heating power than the expected plasma-radiated power.

A series of calibrations were performed, both in air and in vacuum, at a range of bias voltages from 0.5 V to 8 V. All calibrations applied the DC heating for 0.6 s, followed by 0.4 s of cooling, with a 20 kHz excitation voltage frequency. At this stage of the diagnostic's development, the excitation voltage in shot mode (i.e. normal measurements, without applying the calibration DC bias) was run at full amplitude, meaning there was no headroom to add the DC bias. As such, all of these calibrations used an excitation voltage in calibrations with half the amplitude of that used in shot mode: 10 V instead of 20 V. This allowed a DC bias of up to almost 10 V, but the change from 20 V to 10 V drive amplitude for the calibration meant that conditions were not identical between calibration and measurement, which detracts from the new calibration method's advantage of being closer to real world use than the old DC calibration method.

The cooling time was calculated by using a least-squares curve fitting method to fit Equation 2.3.5 to the last 0.4 s of the data, as described in Section 2.3.2. The sensitivity was calculated using the fit data, the measured DC bias voltage and the measured current in the manner described in Section 2.2.1. Error bars are derived




Figure 3.3: Calibration dependence on bias voltage, in air. Both the sensitivity and cooling time vary with bias voltage by statistically significant amounts, though the relative changes are small.

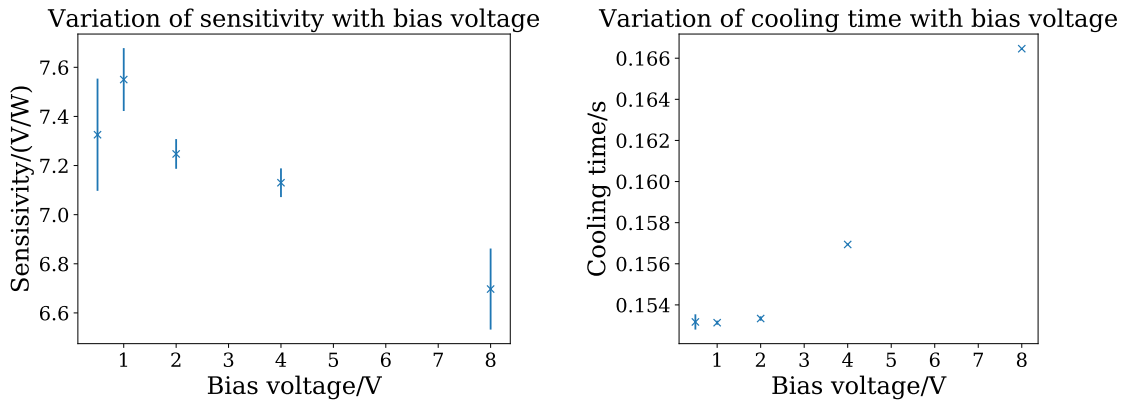


Figure 3.4: Calibration dependence on bias voltage, in vacuum. Although the absolute values are larger, the relative variation of sensitivity and cooling time are similar to what was measured in air (Figure 3.3).

from the quoted standard deviations of fit parameters reported by the fitting routine and the standard deviations of the voltage and current time traces, as appropriate.

Figure 3.3 shows the results of the bias voltage scan in air. There is some variation in sensitivity, which reduces by around 7% between the lowest and highest bias voltages, most likely due to the increased sensor temperature. However, the cooling time changes very little, only increasing by about 1%.

Figure 3.4 shows the same test done in vacuum, at 4×10^{-5} mBar. The sensitivity variation has doubled, and the cooling time variation has increased by a factor of 10. Since this sensor is slower to cool down in vacuum, and therefore reaches thermal equilibrium at a higher temperature, this result supports the hypothesis that the increased sensor temperature affects the calibration.

This test shows that the DC bias voltage should be chosen carefully, since the system is slightly non-linear (we shall see this again in Section 3.4.5). The sensors have a

resistance of around $1.2\text{ k}\Omega$, which means that a 1 V bias applies just under 1 mW to the sensor. This is roughly what is expected to be absorbed by the horizontally-viewing bolometer array's channels in the MAST-U Super-X divertor chamber [68]. The vertically-viewing array's channels are expected to receive substantially more power, but these sensors have been designed specifically for high power measurements, so we expect that the calibration values will be more stable with hotter sensors in these cameras. Keeping the DC bias at 1 V has the added advantage that the ADC can be set to its highest gain setting (and therefore highest sensitivity), where the maximum voltage it can measure is 1.25 V . The small bolometer signals expected during shots mean this gain setting will be used during shot mode, so keeping the same gain setting during calibrations makes the calibration procedure even more similar to the shot measurement procedure.

3.1.3 Calibration dependence on pressure

The final major outcome of the tests in York was the measurement of the calibration parameters' dependence on pressure. Since conduction by gas provides an effective means of transferring heat away from the bolometer sensor, the cooling time reduces as the gas pressure increases. The sensitivity depends on the cooling time and so is also reduced. However, the bolometer equation, Equation 2.1.9, relies on constant calibration parameters. In the Super-X divertor, where neutral pressure is expected to be significantly higher than the main chamber [69], this may cause a substantial change in the calibration parameters compared with the calibration done in vacuum before each shot. Consequently it is important to estimate the likely effect of higher neutral pressure on the calibration parameters.

A range of calibrations were performed at varying pressures, from air pressure (approximately 1000 mBar) down to a high vacuum of $2 \times 10^{-5}\text{ mBar}$. For each calibration, the cooling time and sensitivity were calculated. All calibrations used 1 V DC heating for 0.6 s , followed by 0.4 s of cooling, with a 20 kHz drive voltage of 10 V amplitude.

Figure 3.5 shows the variation of cooling time with pressure, as calculated from these calibrations. There is little variation at very low pressures from 10^{-5} mBar to 10^{-2} mBar , and also little variation at pressures above 10^2 mBar . However, in the range 10^{-2} mBar to 10^1 mBar the cooling time varies dramatically, by almost a factor of 3.

Figure 3.6 shows the sensitivity at the same pressure values as Figure 3.5. The sensitivity depends on physical properties of the sensor (which are invariant at different pressures), and on the cooling time, so it is no surprise to see the same

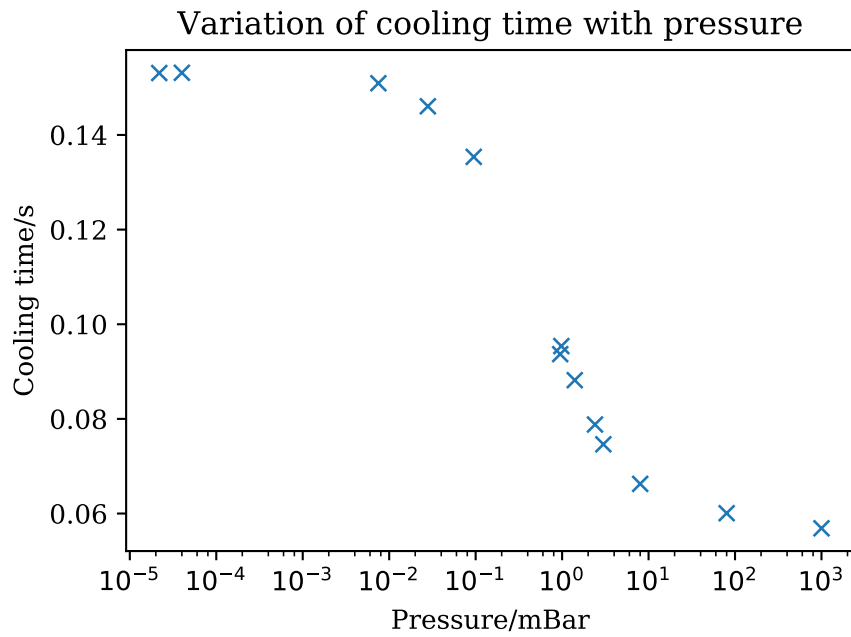


Figure 3.5: Variation of cooling time with pressure. The cooling time is fairly constant at very low and very high pressures, but varies significantly at intermediate pressures.

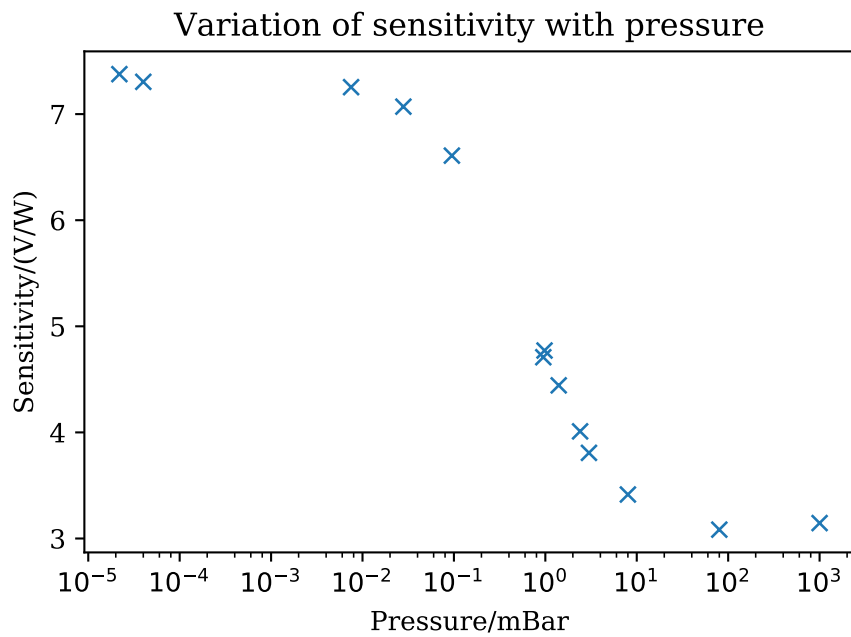


Figure 3.6: Variation of sensitivity with pressure. As expected, the sensitivity behaves in a similar way to the cooling time.

behaviour as the cooling time. This does however provide a useful check that the calibration is producing reasonable results.

The variation of the calibration parameters with pressure has important consequences for MAST-Upgrade. As previously mentioned, the new closed divertor will experience significantly higher pressure than the main chamber. As such, even though the main chamber is predicted to have a neutral pressure of around 10^{-2} Pa (10^{-4} mBar), the divertor is predicted to have a neutral pressure of 1 Pa to 10 Pa, or 10^{-2} mBar to 10^{-1} mBar [69], which is where the sensitivity and cooling time begin to reduce substantially from their vacuum values.

This test shows that bolometer measurements made in the divertor must be considered along with the measured or expected neutral density. This is particularly important if the measurements are to be used for real-time control in the future, as it may require a slightly more complicated control loop, if the pressures are shown to have a big impact on the calibration during typical MAST-U diverted plasmas. There are several neutral pressure gauges to measure the pressure, but none situated very close to the bolometers: the closest is two sectors away in the lower divertor (and so will require assuming the pressure is toroidally symmetric), and there is also a pressure gauge in the same sector as the bolometers but in the upper divertor (which will require assuming up/down symmetry) [70].

3.1.4 Summary of York tests

The tests in York demonstrated for the first time that the new electronics were able to make realistic and usable measurements using an actual bolometer sensor, similar to the type to be used in MAST-Upgrade. The small filtered output signal was found to be limited more by the bit resolution than by analogue noise, and this resulted in modifications to the firmware to increase the bit resolution of the signal processing pipeline.

The variation of calibration parameters with bias voltage provides insights into how to optimise the novel calibration procedure itself. We have shown that applying high ohmic heating changes the sensitivity and cooling time. However, the more fundamental issue is having to reduce the amplitude of the excitation voltage to perform calibrations, whilst still keeping it as large as possible in normal operation to increase the signal size. After these tests, it was decided change the firmware design slightly: the waveform used for the excitation voltage was scaled down, to produce only 18 V out of the maximum 20 V range. This then allowed calibrations with 1 V DC bias, which are similar to the expected power loadings for the horizontal divertor bolometers, without hitting the positive rail on the main DAC producing the drive

voltage. As a bonus, the 1 V DC bias does not saturate the main ADC at 1.25 V gain, which allows us to perform calibrations using the highest ADC resolution, and indeed the same resolution that will be used in the shot. Overall, this calibration procedure is much more representative of normal measurements, and so is expected to be more accurate.

Finally, we measured a significant variation of the calibration parameters with pressure. Knowing the pressure range in which the calibration varies significantly will allow us to treat the divertor bolometer data with more care, and in conjunction with neutral pressure gauges in the divertor region. It would be interesting during the initial MAST-Upgrade campaigns to perform experiments looking at the bolometer measurements in the divertor to better quantify this effect in MAST-U plasmas.

3.2 Installation on JET

In the summer of 2015, the bolometer electronics were installed on JET, connected to 8 reserved channels of the KB5V vertical bolometer diagnostic [71]. These reserved channels had been installed in-vessel along with the 24 active channels of KB5V, but were not used during routine operations. This combination of existing sensors and new electronics became the KB5F diagnostic.

The installation on JET enabled us to test the system in the environment of a large, complex experiment requiring significant automation of the diagnostic, and alongside an already established bolometer system, none of which we present in the tests in York. This was the first time the new electronics had been installed on a tokamak, so the primary purpose of the installation was to demonstrate reliable operation and production of good quality data in a tokamak environment. This was certainly achieved: the diagnostic produced data for well over a thousand JET pulses during the C35, C36 and C37 campaigns, and associated restart periods, over the course of 2015 and 2016.

The other aim of the installation was to compare the quality of the data with that of the existing JET bolometer electronics. Since the cabling and sensors were the same for both systems, any improvements in the signal quality could only come from the BOLO8's electronics and the use of digital signal processing on the FPGA instead of analogue filters on the existing JET electronics.

Figure 3.7 shows the lines of sight of the existing and new diagnostics. The fields of view are achieved using an array of pinholes which image onto the bolometer foils. The spatial resolution of the KB5F channels is very poor, which means that it is not

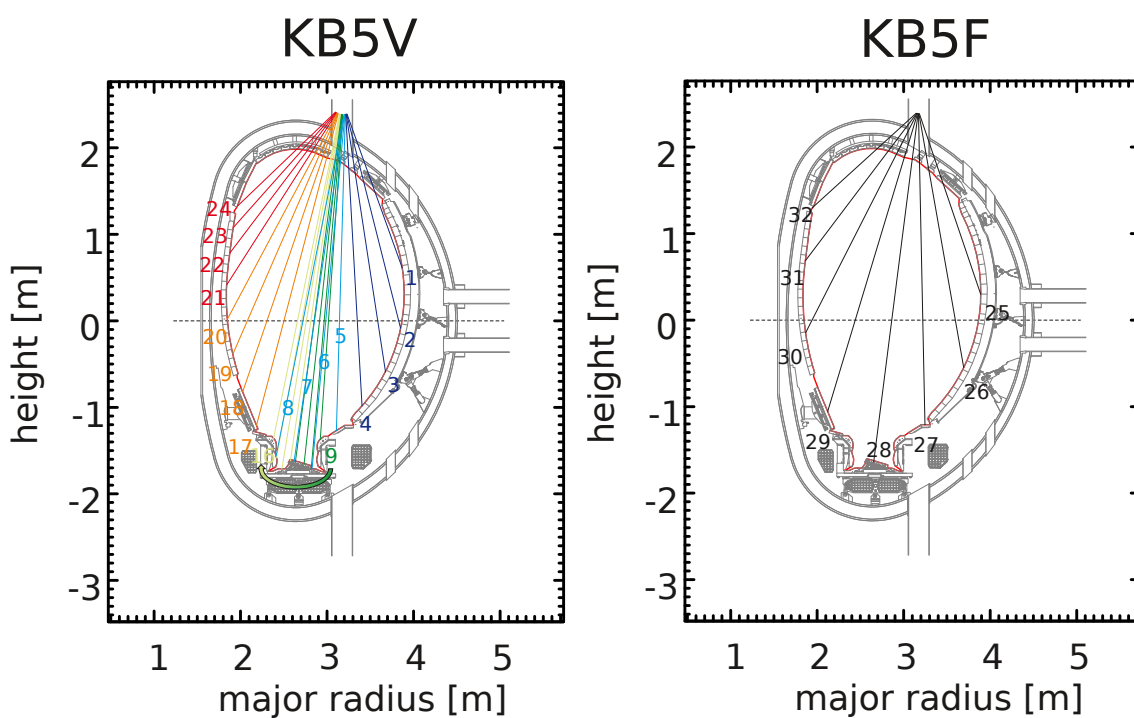


Figure 3.7: Lines of sight of the existing (KB5V) and new (KB5F) vertical bolometer diagnostics, shown within the JET vacuum vessel (light grey). The first wall (the plasma-vessel boundary) is shown in red. Although represented by 1D lines, each channel actually views a cone which extends half way to each neighbouring channel. The spatial resolution of KB5F is much lower than KB5V, though there are some channels which overlap (such as 3 and 26), which allows for comparisons between the two systems.

really possible to explore novel physics with the diagnostic. Instead we must content ourselves with general comparisons with the existing electronics.

There is at least some overlap between the channels. For example, KB5V channel 3 and KB5F channel 26 share almost exactly the same line of sight, as do channels 7 and 28. What Figure 3.7 does not show is that the lines of sight have a finite width, with each channel's view extending half way to the neighbouring line. The KB5F lines of sight are more widely spaced, so although the viewing directions are the same, the channels do not see exactly the same amount of plasma. It is possible to account for this using tomography, but for quick data verification and noise level analysis it is sufficient to choose channels with only similar lines of sight. Channels 3 and 26 have similar spacing as well as similar viewing angles, and viewed a reasonably large amount of radiating plasma (compared with channels 24 and 32 for example), which made these two channels good candidates for comparisons.

Some of the work below, in particular the voltage comparisons in Section 3.2.1 and the power comparisons using tomography in Section 3.2.4, were published in [72].

3.2.1 Direct voltage comparison

After verifying that the diagnostic was correctly armed, calibrated and triggered, the first test of signal quality was a direct comparison of the measured voltages from KB5V and KB5F.

Figure 3.8 shows a comparison of the two diagnostics for JET pulse 89481, for the two almost-overlapping channels 3 and 26. The KB5V electronics were programmed with a filter bandwidth of 200 Hz, whilst the KB5F electronics used a 1 kHz filter. KB5V also has a high gain amplifier (in this case, the gain was 100), whereas KB5F does not, so the traces use different axis scales for plotting at approximately the same size on the figure.

There are several points to note about Figure 3.8. Firstly, and reassuringly, the overall form of the two voltage traces is similar, as expected for two channels with similar views. This gives us some confidence that the clock and trigger inputs to KB5F are working properly.

It also is apparent that, despite KB5F having a higher bandwidth than KB5V, the overall noise level is lower on the newer system. This can be seen particularly in the periods before and after the plasma (-10 s to 0 s and 32 s to 40 s respectively), and also in the region of relatively steady plasma radiation from 12 26 s. Between 37 s and 39 s, when the voltage levels are almost constant, the KB5F channel has an RMS noise of $1.3 \mu\text{V}$, whereas the KB5V channel has an RMS noise of $1.7 \mu\text{V}$ (after

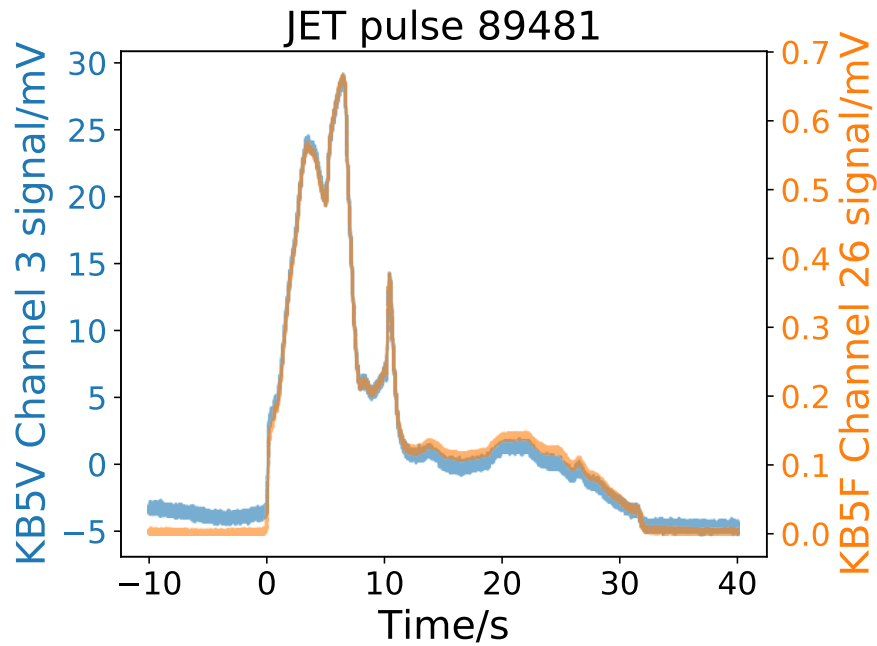


Figure 3.8: Comparison of KB5V (blue) and KB5F (orange) raw output voltages. The measurements are in good qualitative agreement, and the KB5F trace features lower noise despite using a higher bandwidth filter.

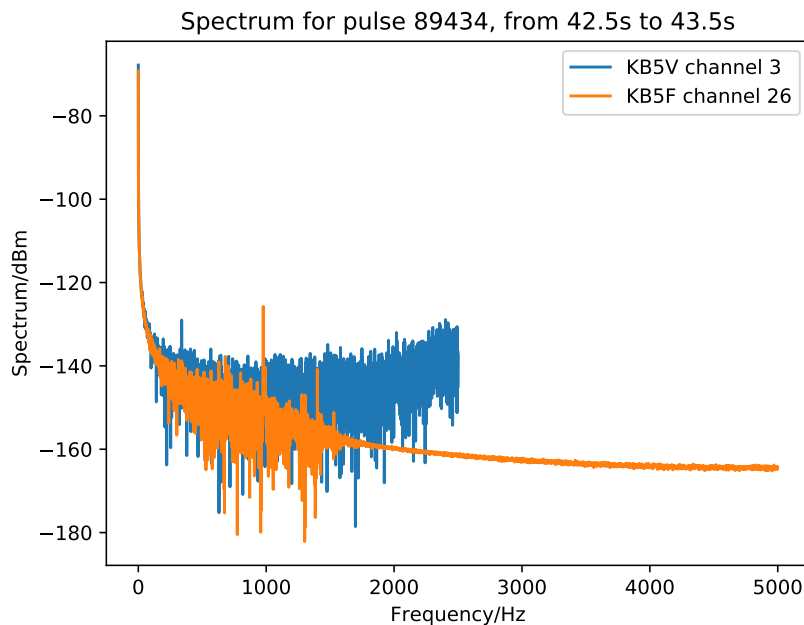


Figure 3.9: Comparison of voltage magnitude spectra of KB5F and KB5V, for a shot where both channels were set to use 1 kHz filters. The noise levels are significantly lower for KB5F, especially at higher frequencies above the cut-off frequency, though there is a strong line noise source at around 1 kHz.

dividing by the factor 100 gain in the KB5V amplifier). This is due both to better ADC performance and the improved noise suppression of the FPGA's digital filter over the analogue filter on the old system.

The noise improvement can be further highlighted when KB5V is set to the same 1 kHz bandwidth as KB5F, as was the case in JET pulse 89434. Figure 3.9 shows a comparison of the spectra of the two systems, during a period at the start of the plasma where the radiation was relatively steady. It is clear that KB5F has much better attenuation above the 1 kHz cutoff frequency, and also lower noise below it (the spectra were scaled to account for the different gains of the two systems).

Furthermore, the KB5V trace in Figure 3.8 shows a slight drift of approximately 1 mV in the offset (the voltage reading for zero input power) at the start and end of the trace, whilst the KB5F trace has no offset drift, or at least a drift smaller than the noise levels in the trace. This could be due to several factors. KB5V measures only the voltage amplitude of the demodulated bolometer signal, assuming a fixed phase which is compensated for manually in advance. Any change in the phase throughout the shot (due to forces or temperature variations in the cables, for example) could cause an apparent reduction in the signal, which would show as an offset drift. KB5F does not have this problem, since it measures both the amplitude and phase and so we can check whether the bolometer signal's phase is indeed constant throughout the pulse, and hence definitively account for offset drifts due to phase drifts. In this case however, no significant phase drift was measured. Of course, it could also be the case that the very old KB5V electronics amplifiers are just not as stable any more, and some offset drift due to imperfections in the electronics is occurring.

Overall, the raw voltage measurements are encouraging: the system produces data which is qualitatively in agreement with the established diagnostic, and with lower noise levels even at higher bandwidths.

3.2.2 Optimisation of drive frequency

One notable feature of Figure 3.9 is the sharp peak at 1 kHz in the KB5F spectrum. This is not present in the KB5V spectrum, so it cannot be due to radiation. It is therefore most likely pickup in the long cables running from the bolometer sensors at the top of the JET vessel in the torus hall down to the diagnostic electronics in the basement.

Even though the cable routes are the same for KB5V and KB5F (and should therefore both experience this pickup), KB5V operates at 50 kHz drive frequency, whereas KB5F initially operated at 19.3 kHz drive frequency, the resonant frequency of the

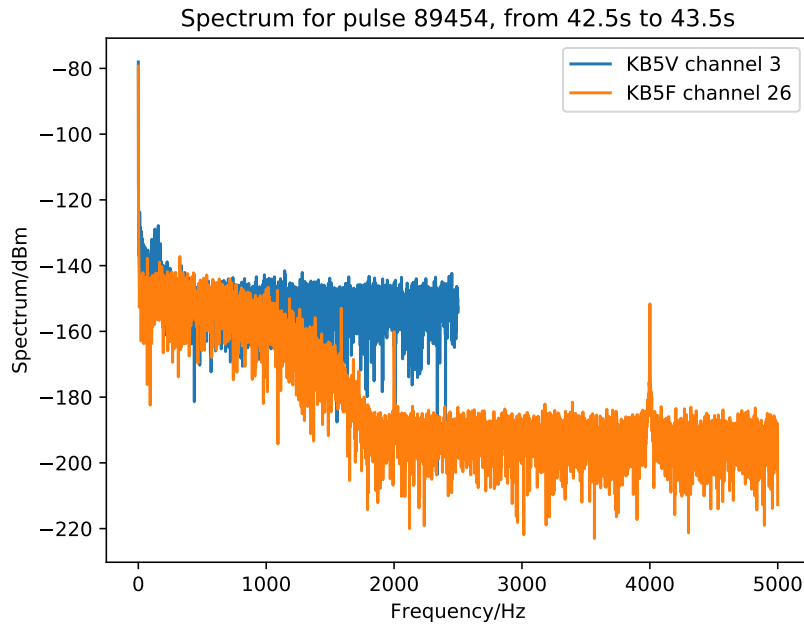


Figure 3.10: Comparison of voltage spectra of KB5F and KB5V, with KB5F 18 kHz drive frequency. The strong line noise source at 1 kHz which can be seen in Figure 3.9 has gone, and the only line noise source now is an alias of the 18 kHz drive frequency at 4 kHz which it is below the noise floor in the pass band (up to 1 kHz).

LC filter in the calibration circuitry. The presence of the 1 kHz feature therefore suggested a strong line noise source at either 18.3 kHz or 20.3 kHz. It could also be a higher frequency which has been aliased down to 1 kHz due to the 10 kHz sampling rate, but the size of the peak given the strong frequency attenuation of the digital filter makes this unlikely.

The new system has the flexibility to adjust the drive frequency. Through trial and error, we arrived at 18 kHz as a suitable drive frequency. The resulting spectra, again for a relatively steady radiation at the start of the plasma in pulse 89543, are shown in Figure 3.10. There are no longer strong noise sources far above the noise floor in the 1 kHz pass band. The peak at 4 kHz is an alias of the 2ω frequency component from the signal and reference wave mixing during the AC synchronous detection. It actually has a very low noise power, since it too is below the noise floor in the pass band.

Figure 3.11 further illustrates the difference, through the duration of a JET pulse. It is clear from these spectrograms that the line noise is a steady source which is not related to the plasma. We can see that the move to 18 kHz removes the 1 kHz line noise, and the resulting 4 kHz line is only around the same amplitude as the noise floor in filter's pass band (itself clearly visible up to around 1.5 kHz).

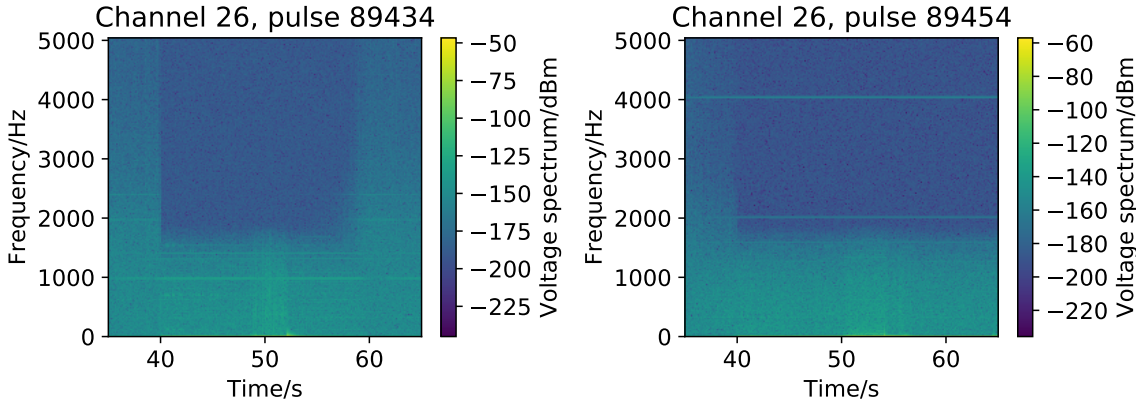


Figure 3.11: Spectrograms for 19.3 kHz drive frequency (left) and 18 kHz drive frequency (right). The line noise sources can be clearly seen, and since both persist throughout the shot they cannot be coming from the plasma radiation.

Nearly all JET pulses from 89456 onwards were measured with a drive frequency of 18 kHz as a result of these measurements.

3.2.3 Power calculation

Once we had a reliable voltage measurement, the next step was to produce an absorbed power calculation using the data from the new electronics. Ultimately, plasma physicists are not interested in the raw output from the diagnostic, but rather the measured physical quantity, so it is important to demonstrate that KB5F is able to provide this with at least as good quality as the established KB5V.

KB5F benefits from a fast, automated calibration procedure, as described in Section 2.2.1, which is capable of calibrating all channels simultaneously. This means we are able to perform a calibration before every pulse, to ensure the calibration conditions are as close to those in the pulse as possible. Using experience gained from the York vacuum chamber tests in Section 3.1, the calibration used a 1 V DC bias with an 18 V amplitude drive voltage, the same amplitude as was used in the pulse.

Initially the sensor was heated for 0.5 s and cooling was measured for a further 0.5 s. Whilst this provided reasonable estimates of the sensitivity and cooling time, the calculated offsets were found to be slightly larger than those actually measured during the 10 s prior to JET plasmas. Increasing both the heating and cooling measurement times to 1 s greatly reduced this offset discrepancy, and gave more confidence in the sensitivity and cooling time estimates too.

In order to calculate the power using the bolometer equation 2.1.9, we must differentiate the signal. Since differentiation amplifies high frequency noise, we must also

smooth the differentiated signal to avoid greatly reducing the signal-to-noise ratio of the power measurement compared with the raw voltage measurement. This has traditionally been achieved by smoothing with a filter function which also differentiates the signal [53], and indeed this is the method the real-time power output from the FPGA uses.

In contrast, the BOLO/KB5V PPF data, which is the go-to data source for the line-integrated power measured by KB5V and is used as input into the tomography, employs a different method. This uses a non-linear smoothing method, developed at JET by Gottardi [58], whereby the amount of smoothing is determined using a hyperbolic tangent:

$$y_i^{sm} = \frac{1}{3} \sum_{k=i-1}^{k=i+1} y_k + \tanh \left\{ a_{level} \times \max \left(\left. \frac{dy}{dt} \right|_{x_i}, \left. \frac{dy}{dt} \right|_{x_{i+1}} \right) \times \max \left(\left. \frac{dy}{dt} \right|_{no-ELMs} \right) \right\} \times \left(y_i - \frac{1}{3} \sum_{k=i-1}^{k=i+1} y_k \right) \quad (3.2.1)$$

Here, a_{level} is the compression factor for the hyperbolic tangent, and y_i^{sm} is the smoothed signal. The amount of smoothing depends on the time derivative of the signal: if the derivative is large (e.g. because of an ELM) compared to a reference derivative computed when the plasma is roughly steady state, the argument of the hyperbolic tangent is 1 and so no smoothing is applied, since the smoothing terms cancel and leave the original signal. If the derivative is small then the argument of the hyperbolic tangent is small, and so the output is dominated by the smoothing terms.

The advantage of smoothing using Equation 3.2.1 is that we can apply significant smoothing to small fluctuations, which are assumed to be just noise, whilst large transient features (primarily ELMs) are not smoothed. Smoothing the ELMs would reduce the measured peak height and hence cause us to underestimate the amount of power radiated during an ELM. However, the nonlinear nature of this smoothing method does affect the frequency spectrum in a way which is difficult to predict, and makes it much more difficult to examine smaller fluctuations which are genuinely interesting features rather than just noise. For this reason, analysis of the MAST-U data will use the traditional deconvolution filter approach instead.

For the purposes of the JET tests however, looking for new physics was not the priority. Instead, the desire to compare as directly as possible the new and existing systems lead to implementation of Gottardi's smoothing method on the KB5F voltage data, and the calculation of absorbed power using this smoothed voltage in Equation 2.1.9. The choice of parameters used in the smoothing algorithm was

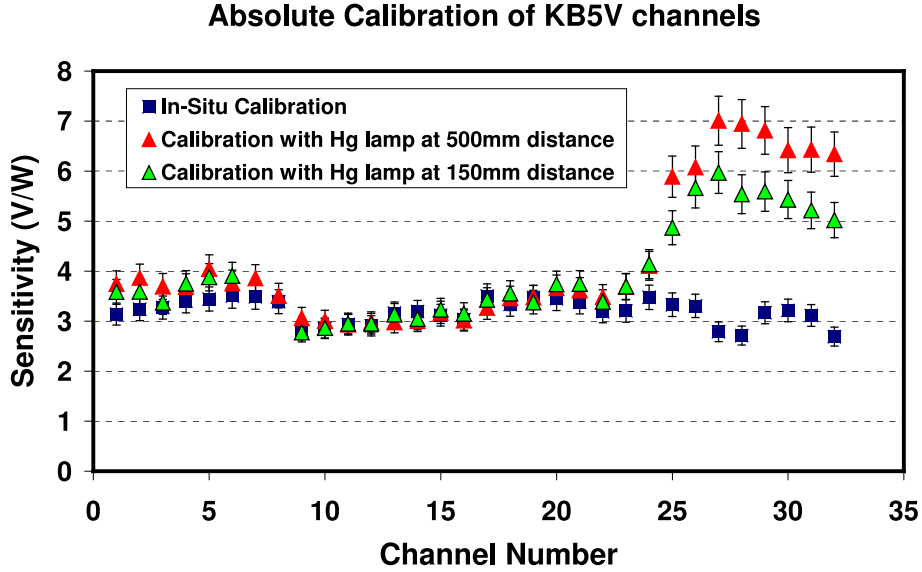


Figure 3.12: Comparison of in-situ electrical and mercury lamp measures of KB5 bolometer sensitivities. Although the two calibration methods agree well for channels 1 to 24, there is a discrepancy of roughly a factor of 2 between the methods for channels 25–32. These are the channels used by KB5F. Reproduced from [73]

based on those used in the production of the BOLO/KB5V PPF data, however the amount of smoothing (specified by the number of passes across the trace using the smoothing algorithm) was doubled, to account for the doubling of sample rate with KB5F. This should produce data with a similar total amount of smoothing, although the effective “bandwidth” of the processed signal may be slightly different due to the non-linear nature of the smoothing.

Direct comparison with the BOLO/KB5V data was further complicated by the fact that it is actually the line-integrated intensity, rather than the absorbed power, which is calculated. The intensity is calculated from the absorbed power by considering the viewing angle subtended by the sensor. This is characterised by the étendue, the solid angle of the sensor as seen from the source integrated over the sensor’s surface, which accounts for the effect of the aperture and collimator in each bolometer camera. If the étendue of the system is E and the absorbed power P , then the intensity is given by:

$$I = \frac{4\pi P}{E} \quad (3.2.2)$$

The étendue of each channel of KB5F is stored in the KB5 geometry file, and so was easy to retrieve. However, calibrations of all 32 vertical bolometer sensors with a mercury lamp showed a marked difference between the sensitivity as calculated using the in-situ (i.e. ohmic) calibration method, and by illuminating the sensor with a mercury lamp [73]. This can be seen in Figure 3.12, taken from [73], which shows

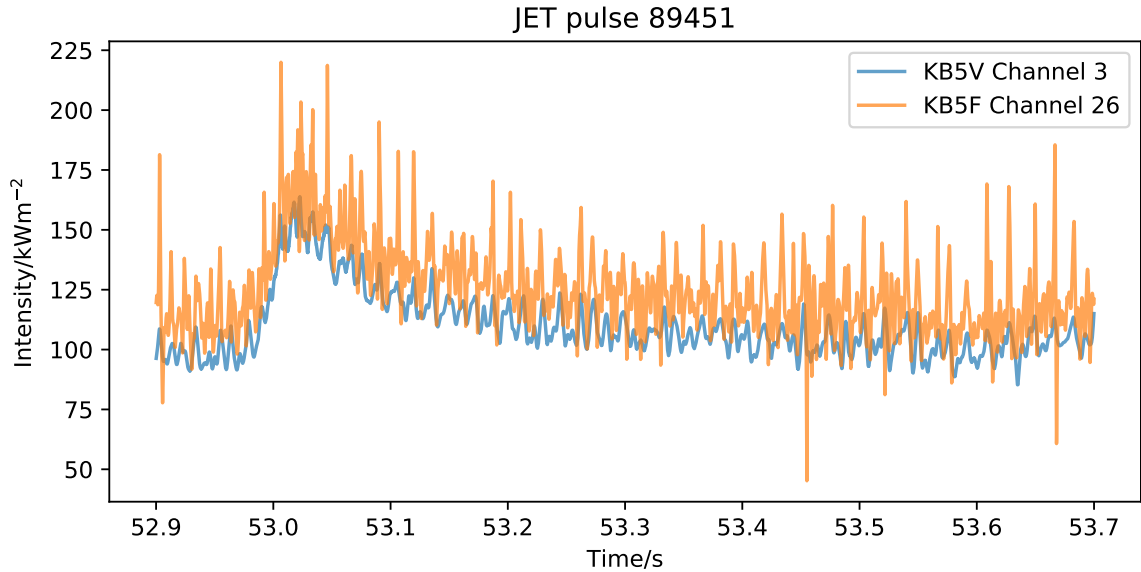


Figure 3.13: Comparison of intensities as measured by KB5V and KB5F, for two overlapping channels. There is good agreement between the two, and noise levels are comparable. After accounting for the time delay in the KB5V data, the times of the ELMs matches up well.

that the in-situ and lamp calibrations agree well for channels 1 to 24, but differ by a factor of approximately 2 for channels 25 to 32 (the KB5F channels). The report attributes the difference in sensitivity to reflections in the collimators for the KB5F channels, effectively increasing the amount of plasma (or lamp) that they are able to view. To account for this, and allowing for the lack of an exact étendue value which accounts for these reflections, the effective étendue used in Equation 3.2.2 was simply doubled for the KB5F channels.

Once this étendue factor was accounted for, we could directly compare the BOLO/KB5V output with the intensity as calculated from the KB5F voltage. Figure 3.13 shows an example, for the two channels 3 (KB5V) and 26 (KB5F) with similar lines of sight, during an H-mode phase of a JET pulse. We can see reasonable quantitative agreement (allowing for the slightly different viewing geometries), and similar noise levels (which we would expect, since both traces have had the same smoothing applied).

Furthermore, the times of the ELMs, which are clearly visible in both channels, match up well. Care must be taken when directly comparing the times of the ELMs, since the KB5V electronics introduces a small time delay in its analogue filters, and this time delay is different depending on the filter bandwidth used. The time delay had already been estimated for the most commonly-used filter frequencies, by comparing the times of the measured peaks in BOLO/KB5V with the measured peaks of Be-II impurity radiation [74], so we needed to subtract the estimated time delay from the KB5V measurements before comparing with KB5F. Of course, since

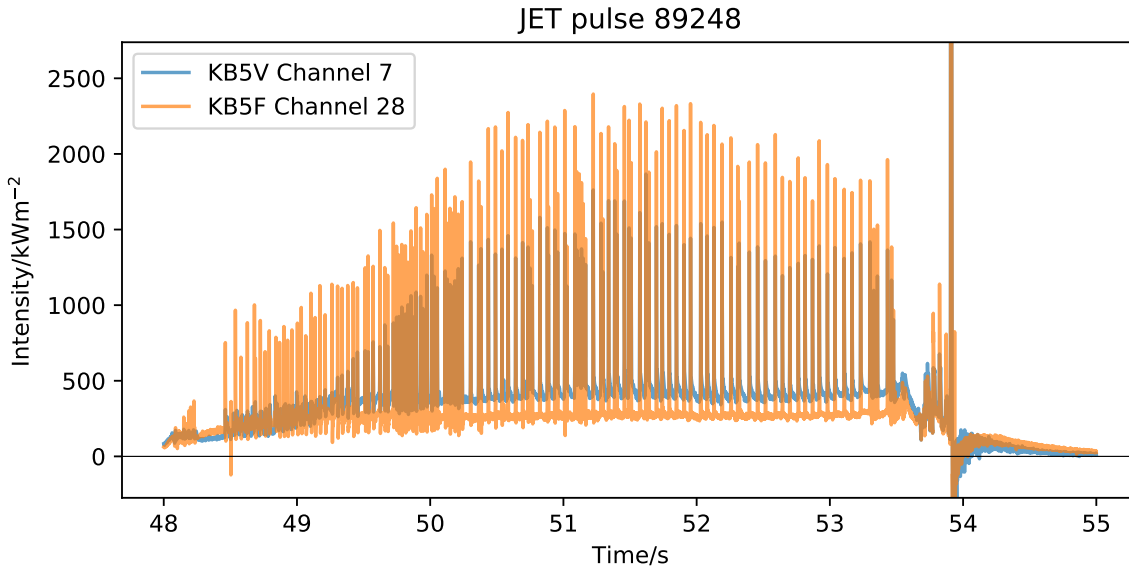


Figure 3.14: Intensity calculation during a disrupted pulse, for two channels viewing different amounts of divertor radiation. Again, the ELMs are clearly visible, more so than in Figure 3.13, and a large spike in power is observed just before 54s as the plasma disrupts. Interestingly, the intensity appears to go negative immediately after the disruption, before decaying away to 0 after about 1 s.

KB5F uses digital filtering on the FPGA, the time delay is known exactly and is independent of the filter bandwidth, so is automatically accounted for in the software which produces the time vector for these plots.

Comparisons of the intensity in pulses which disrupted showed an interesting effect. Figure 3.14 shows one such pulse, which disrupted just before 54s. The channels shown here had the same line of sight, but very different viewing angles (channel 28 views the entire divertor at once, whereas channel 7 views only the horizontal target), so quantitative comparisons should be avoided. Qualitatively, both channels do show that immediately after the disruption the radiated power appears to go negative, before increasing back to positive values and then decreasing to nothing. This apparent negative power measurement is actually most likely due to a reduction in the cooling time of the sensor, caused by an increase in the neutral pressure immediately after the disruption. As shown in Section 3.1.3, an increase in neutral pressure causes a decrease in cooling time. Thus, the rate of decrease of voltage as the sensor cooled was more rapid than it would have been had the cooling time stayed constant, as assumed in the bolometer equation 2.1.9, and this manifests itself as a negative power reading. This illustrates the importance of understanding the neutral pressure when interpreting bolometer readings.

All the power and intensity calculations done so far in this section have taken the voltage amplitude output by the FPGA and performed the smoothing and




Figure 3.15: Comparison of the real-time intensity calculation performed on the KB5F FPGA with the BOLO/KB5V PPF post-processing calculation. The real-time signal shows higher noise levels since it has not been smoothed using Gottardi’s method; it instead uses a 1 kHz deconvolution filter. Otherwise, the agreement between the two channels is good: the ELMs occur at the same times and the relative inter-ELM intensities are as expected given that channel 28 views the entire divertor whereas channel 7 views only part of one tile.

differentiation in software. The FPGA also has a deconvolution filter which can directly output the power absorbed on the sensor, which is described in Section 2.3.1. There is presently little to be gained from using this real-time power output, since it is intended to be used in a control loop and this was out of scope for the tests on JET. However, it was still instructive to check the output of this real-time calculation, which was recorded along with the voltage data in the JET data store. Indeed, several bugs in the real-time calculation part of the firmware were fixed during the JET tests, including dealing with filter overflow, designing the system to use separate filter coefficients for each channel (to account for the different sensitivity and cooling time of each channel) and improving the on-FPGA offset correction.

Figure 3.15 shows the intensity for KB5F’s channel 28 during a high power, high performance pulse near the end of campaign 36, once the firmware bugs in the real-time power calculation had been fixed. The intensity was calculated by simply taking the real-time power calculation and scaling it by the reflection-corrected étendue factor, for comparison with the output from the KB5V channel with a similar line of sight. There is more noise in the real-time signal, since the bandwidth was set at 1 kHz and used the deconvolution method rather than Gottardi’s smoothing method. Some smoothing could easily be applied in a control loop if required, and for the real-time signal the noise is less important than the latency. More importantly still, the real-time plot contains the same features, at the same recorded time, as the post-processing plot, and successfully resolves even the high frequency ELMs. This

shows that the system is capable of producing a reliable real-time output which could be used in a control loop.

3.2.4 Tomography comparison

So far, any quantitative comparisons between the new and existing electronics have only been approximate, since the two systems do not share exactly the same viewing geometries. Even channels, such as 3 and 26, which have the same central line of sight do not have the same size viewing angle, and so will see slightly different amounts of plasma. We can account for the different viewing geometries using tomography: by taking a tomographic inversion and integrating its emissivity profile along the exact line of sight of any particular channel, we can calculate what that channel should have measured. We can then compare this “back-calculated” intensity measurement with the actual measurement to make a quantitative comparison which no longer suffers from the uncertainty due to the viewing geometry.

Tomography solves the system of equations describing the power f_i absorbed by the i th bolometer sensor, as a function of the emissivity $g(x, y)$ and a geometric function $K_i(x, y)$ which takes into account the viewing geometry of the sensor [75]:

$$f_i = \iint g(x, y) K_i(x, y) dx dy \quad (3.2.3)$$

Solving this equation for $g(x, y)$ numerically requires discretising Equation 3.2.3, producing a grid of points on which the emissivity is calculated. The discretised equation is of the form:

$$f_i = \sum_j K_{ij} g_j \quad (3.2.4)$$

Where K_{ij} is a geometric matrix, derived from the geometric function $K_i(x, y)$, and g_j is the emissivity at each point on the grid.

Tomography inverts the above equation to solve for the emissivity at each grid point g_j . Often there are many more grid points j than sensors i , which means the inversion is an ill-posed problem. There are several methods available to solve this problem, some of which are described in [75], but these are out of the scope of this work, since tomographic inversion has already been implemented and routinely used on JET.

Once the emissivity g_j at each grid point has been calculated, we can use Equation 3.2.4 to calculate the expected measurements on each sensor f_i . This is done as a matter of course by the JET tomography code in order to check how good the reconstruction actually is, as one constraint when solving Equation 3.2.4. However,

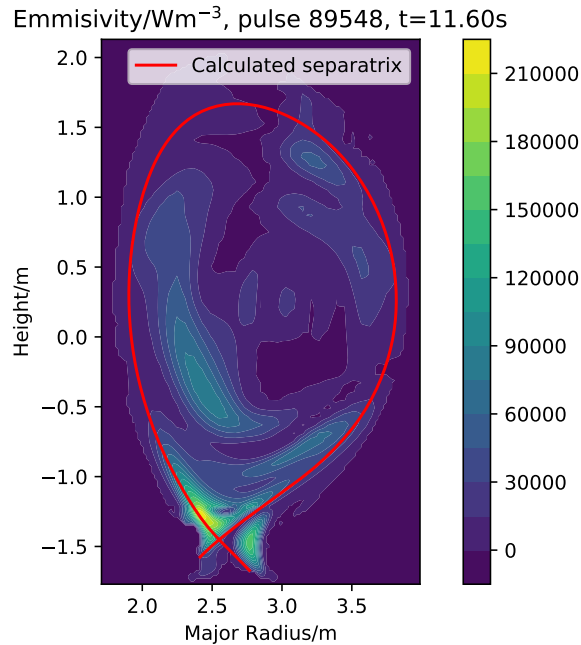


Figure 3.16: Tomographic reconstruction of emissivity profile for a single time slice. Strong emission can be seen in the divertor region near the X-point and in the outer scrape off layer (to the right of the X-point), and to a lesser extent in the core at around $R = 2.5$ m and $Z = 0$ m.

the calculation is only done for the channels used in the reconstruction, which obviously does not include the KB5F channels. We therefore needed to perform the same calculation for these channels.

The geometry matrix K_{ij} is stored in a configuration file used by the JET tomography code, and includes the matrix elements for the KB5F channels. The emissivity g_j at each grid point is not stored in the PPF data file that contains end-user tomography results (the emissivity profile $g(x, y)$ is stored instead), but after examining the tomography code we were able to locate a code output file on the analysis machines for each pulse containing g_j . Using this is significantly simpler than projecting $g(x, y)$ back onto the grid points of the K_{ij} geometry matrix, and also reduces the numerical error involved in doing so, at the cost of retrieving the data from an “unofficial” location.

Figure 3.17 shows the comparison of the measured and back-calculated intensities for one time slice of JET pulse 89548, the emissivity profile of which, calculated using tomography, is shown in Figure 3.16. Channels 1 to 24, to the left of the left-hand dotted line in Figure 3.17, correspond to the 24 vertical sensors in KB5V. Channels 33 to 56, to the right of the right-hand dotted line, correspond to the 24 horizontal sensors in KB5H, the horizontal bolometer array which together with KB5V and KB5F makes up the entire KB5 system. The channels between the two dotted lines are the KB5F channels. Channel 32 has been omitted from the plot, since it was




Figure 3.17: Comparison of measured intensities and intensities back-calculated from the tomographic reconstruction, for all KB5 channels. The KB5V channels 1–24 are shown on the left, and the complimentary KB5H horizontal bolometer channels 33–56 are shown on the right. The KB5F channels 25–31 are shown between the two dotted lines (channel 32 is not shown as it was broken). The agreement between the measured and back-calculated intensities is as good for KB5F as it is for the KB5V and KB5H channels.

found to be broken, most likely due to the gold foil absorber delaminating from its substrate.

Clearly, the agreement between the measured and back-calculated intensities is just as good for the KB5F channels as it is for those KB5V and KB5H channels which have not been used to constrain the reconstruction. This is despite the KB5F data not being used at all in the tomography, in order to avoid biasing the result in favour of the KB5F readings. This is the most convincing evidence yet that the KB5F diagnostic is producing data of comparable accuracy, and in good agreement with, the already established and heavily used JET bolometer system, but with better time resolution and lower noise levels.

3.2.5 Total power calculation

The total radiated power is an important physics quantity. It can be calculated by integrating the emissivity profile $g(x, y)$, calculated using tomography, over all space. However, this is frequently impractical, since the tomographic inversion is extremely time-consuming: each time slice can take upwards of 30 minutes to perform, and a pulse can consist of many hundreds of time slices of interest. For experimentalists wanting this quantity to inform the setup of their next pulse, this tomography method is far too slow. Even worse, the ill-posed nature of tomography means that each reconstruction must be inspected by hand to ensure it is suitably free of artefacts, and sometimes multiple runs with different channels included and excluded may be needed to avoid producing inaccurate tomograms.

Instead, a good approximation of the radiated power can be made using a weighted sum of the individual intensity measurements. This is much less computationally intensive than a full tomographic reconstruction, and so can be performed for every sample period of the entire pulse. It is also less sensitive to artefacts, meaning its production can be automatic.

The weighted summation method of calculating the total power is described in [76]. Essentially, the line-integrated intensity measurements are summed, with weights determined by the “distance” between lines of sight in a projection space in which all lines of sight are parallel:

$$P \approx 2\pi R_0 \sum_i f(p_i, \xi = \text{const}) \Delta p_i \quad (3.2.5)$$

Here, R_0 is the major radius, $f(p_i, \xi)$ is the measured intensity and p_i is the impact parameter (closest distance to chosen origin) of channel i . ξ is the angle between the line of sight and the horizontal, which is the same for all channels in projection

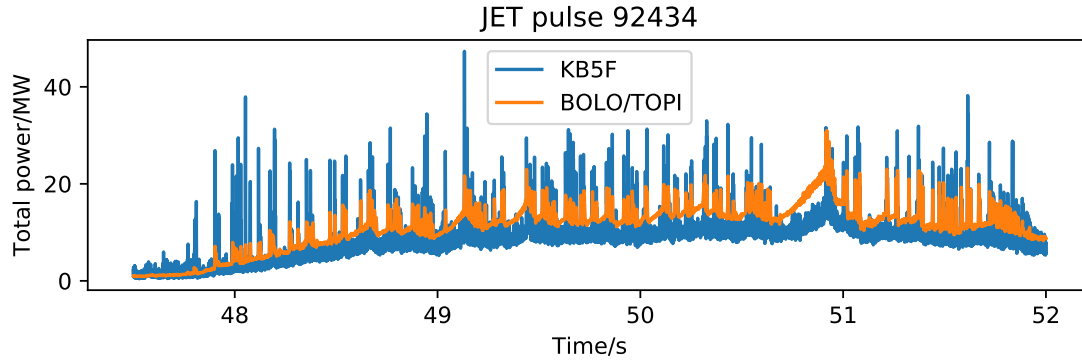


Figure 3.18: Calculated total power for a high performance JET pulse, as calculated by KB5F and by the inter shot analysis code to produce BOLO/TOPI. The two calculations are in fairly good agreement, with KB5F being slightly lower due to lower spatial resolution and incomplete coverage of the plasma cross section. The ELM power peaks are higher in the KB5F calculation because the electronics uses a higher bandwidth filter and less smoothing is applied.

space. Δp_i are the differences in the impact parameter between the i th and $(i - 1)$ th lines of sight, and become the weights in the weighted sum.

The total power calculated in this way depends on the choice of origin. For JET divertor plasmas, the X-point is chosen as the origin, and the impact parameters are projected onto lines of constant ξ and scaled to account for the non-circular cross section. Details of the scaling for the impact parameter can be found in the KB5 PPF manual [58].

We used this calculation method to provide a fast estimate of the total power, based on the measured intensities from the KB5F channels. The BOLO/TOPI PPF data signal provided by the JET inter-shot bolometer analysis code performs this weighted sum calculation with X-point scaling, using a selection of the KB5V channels connected to the incumbent JET bolometer electronics, so provided a good benchmark against which to test our new system. The inter-shot code which produces BOLO/TOPI must calculate the line-integral power on each sensor from the individual voltage measurements and then calculate the total power, which means the result is not available until around 10 minutes after a JET pulse has finished. In contrast, we can use the real-time power signal from the FPGA and exploit the processing power of the ARM CPU to provide a total power data signal which is ready at the same time as the raw voltage signal. This would be a big advantage for experimentalists, as delivering the total power more quickly gives them more time to plan their next pulse.

The implementation of the total power calculation on the KB5F diagnostic was low




Figure 3.19: The same trace as Figure 3.18, zoomed to view only a few ELMs. The lower bandwidth and extra smoothing of the BOLO/TOPI measurement can clearly be seen: the true ELM power is likely to be closer to the KB5F measurement.

priority, so wasn't completed until near the end of the C36b campaign. By that time channel 30 had started to show signs of being broken, including huge noise spikes in most pulses and unreliable calibration measurements. Coupled with the suspected delamination of channel 32, this meant that KB5F was reduced to only 6 working channels, missing a substantial amount of the inboard view. Nevertheless, when removing these channels from the analysis, it was possible to produce reasonable-looking estimates of the total power. Figure 3.18 shows a comparison between the power estimate of BOLO/TOPI, calculated by the inter-shot code with KB5V channels, and that of KB5F, available immediately after the shot using a mixture of the FPGA real-time power calculation and software calculation of the total power. Unsurprisingly, considering the lack of complete coverage, the KB5F estimated total power is slightly lower than the BOLO/TOPI estimate, with the exception of the ELMs, which KB5F estimates to produce higher power. However, BOLO/TOPI is smoothed more than KB5F, and KB5V was operating at a lower bandwidth of 200 Hz as can be seen in Figure 3.19, so the true height of the ELM peaks is most likely higher than the BOLO/TOPI prediction, and closer to the KB5F prediction. Both estimates seem numerically reasonable: this was a high powered pulse, with 32 MW of heating power during the time slice shown, so 10 MW of continuous radiation and almost 30 MW of radiation during transients seems entirely feasible.

The extra smoothing of BOLO/TOPI also means the inter-shot trace has lower noise than the real-time trace, which is not unsurprising. It is of course quick and easy to apply smoothing to the KB5F trace if low noise is needed, yet the signal still provides 1 kHz bandwidth to enable viewing high frequency transients if required. Overall, the fast total power calculation should be considered a success, and the ability to make this available immediately after the pulse (or in real time with small

firmware modifications and extra infrastructure in place) is something the existing bolometer system does not have.

3.2.6 Summary of JET tests

The installation of the bolometer electronics on JET to form the KB5F system marked the first use of the new design on an operational tokamak. The development of the software to interface with the JET GAP software for running each pulse provided useful lessons which will be applied to the software development for MAST-Upgrade. Successful operation of well over 1000 pulses provides confidence that the diagnostic will perform reliably on other devices.

Despite the limited spatial resolution of the spare channels, we were still able to produce useful data. We demonstrated improved signal quality compared with JET's existing analogue electronics, even at higher bandwidths. The first real-world demonstration of the configurable drive frequency allowed us to improve the signal-to-noise even further, and worked exactly as intended. The calculated power (the real physics parameter of interest) was shown to be in good agreement with the established JET diagnostic, both qualitatively and quantitatively through the use of tomography, and the real-time power calculation produced reasonable data. This is encouraging for future feedback control using the bolometer measurements. Finally, the combined power of the FPGA programmable logic and ARM CPU were used to produce a fast estimate of the total radiated power which was in reasonable agreement with the post-processed, inter-shot estimate currently widely used; this will be a welcome improvement for experimentalists and was not possible with the existing JET electronics.

3.3 Installation on TCV

The bolometer system was installed in the summer of 2016 on the TCV tokamak [77]. This work was done in collaboration with Umar Sheikh at the École Polytechnique Fédérale de Lausanne (EPFL). The aim was to determine the effect of blackening gold foil bolometers to reduce reflections by the foil, which in turn reduce the power measured by the bolometers. Blackening was done by coating two channels of a 4-channel bolometer sensor with a thin vapour deposited carbon film. More details on the experiment itself can be found in [78].

TCV did already have a bolometer diagnostic, but there were no spare channels available for these tests. Electronics suitable for operating the resistive bolometers

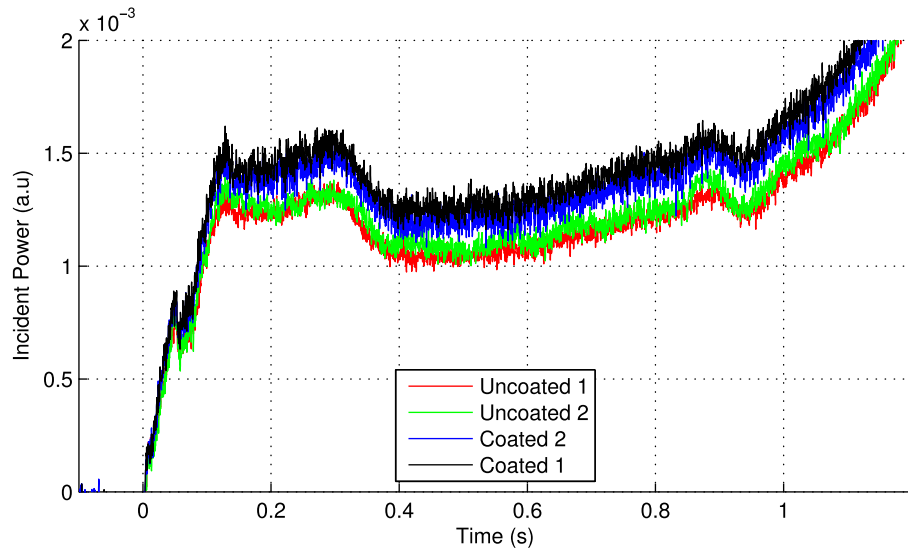


Figure 3.20: Comparison of blackened and unblackened bolometer measurements performed on TCV. The blackened (coated) bolometers show a 15% to 20% increase in the measured power compared with the unblackened ones. Reproduced from [78].

was required, and Umar made contact after reading about the development of the BOLO8 system in the proceedings of the 1st EPS Conference on Plasma Diagnostics (ECPD)[79]. TCV provided the opportunity to test the adaptability of the system to a new tokamak experiment environment, as well as forging new contacts between CCFE and EPFL.

The electronics used at TCV was loaned to EPFL by D-TACQ, and I supplied the FPGA firmware, software and technical expertise. It was originally intended to use the BOLODSP architecture, as described in Section 2.4, but this was not ready in time for experiments. Instead, the Xillybus firmware, described in Section 2.3 and previously demonstrated by operation on JET (Section 3.2) was used, with a collection of scripts run manually to arm, trigger and read back data from the device in place of the GAP interface for which the JET software was designed.

The experiments were successful, and the data collected by the bolometer system was used to demonstrate an approximately 15% increase in measured power with blackened bolometers, compared with the unblackened reference bolometers viewing the same line of sight. This is shown in Figure 3.20, taken from the publication resulting from the collaboration [78] and which shows the power calculated with the new system's voltage measurements and calibration procedure.

As a result of these successful tests, EPFL have ordered two D-TACQ units, which will use the BOLODSP architecture, for an upgrade of their resistive bolometer diagnostic.

3.4 Lab-based measurements at PPPL

Princeton Plasma Physics Laboratory (in Princeton, NJ, United States) are using the D-TACQ BOLO8 for the NSTX-U resistive bolometer diagnostic. The bolometer diagnostic on NSTX-U is managed in collaboration with Oak Ridge National Laboratory (ORNL), with Matt Reinke as the Responsible Officer. After initially working with Matt on the tests in York (Section 3.1), I visited PPPL on two occasions to do further lab-based testing of the bolometer sensors and electronics, since the lab was able to provide testing equipment unavailable at CCFE which could be used to better characterise the new system.

In return for the provision of lab equipment to perform these tests, I helped PPPL to integrate the BOLO8 system into the NSTX-U shot cycle. This involved working with Greg Tchilinguirian to specify the MDSplus [80] data tree to store the BOLO8 data, and develop software scripts to handle the various stages of the NSTX-U shot cycle (arm, calibrate, trigger, readback). Since the MDSplus data acquisition framework is used at a range of tokamaks, including TCV, this work will reduce duplication of effort should the system be installed on additional tokamaks.

In this section, we focus only on the tests carried out. Chief among these was the measurements of the absolute accuracy of the bolometer sensitivity. This involved first calculating the sensitivity from calibrations, then calculating the absorbed bolometer power in a test shot where the sensor was heated by a 405 nm laser. By comparing the power to the known input power from the laser, we evaluate the accuracy of the sensitivity calculation.

Additional tests included measurements of the bolometer response to a fixed input power when the drive frequency was varied, and when the cable length was varied. The linearity of the laser/bolometer system was also studied by varying the laser power at a fixed drive frequency.

3.4.1 Laser profile

In order to evaluate the accuracy of the absorbed power, it is necessary to know with good precision the amount of power incident on the bolometer sensor. The laser power was measured by a Thorlabs PDA36A Silicon power meter [81]. This power meter has an area of 13 mm^2 , which is larger than the area of the bolometer sensor ($3.8 \text{ mm} \times 1.4 \text{ mm}$). This means we need to scale the power measured by the power meter according to the ratio of the areas of the power meter and bolometer sensor. In this way, we ensure that we compare the power density, which is the quantity of interest.

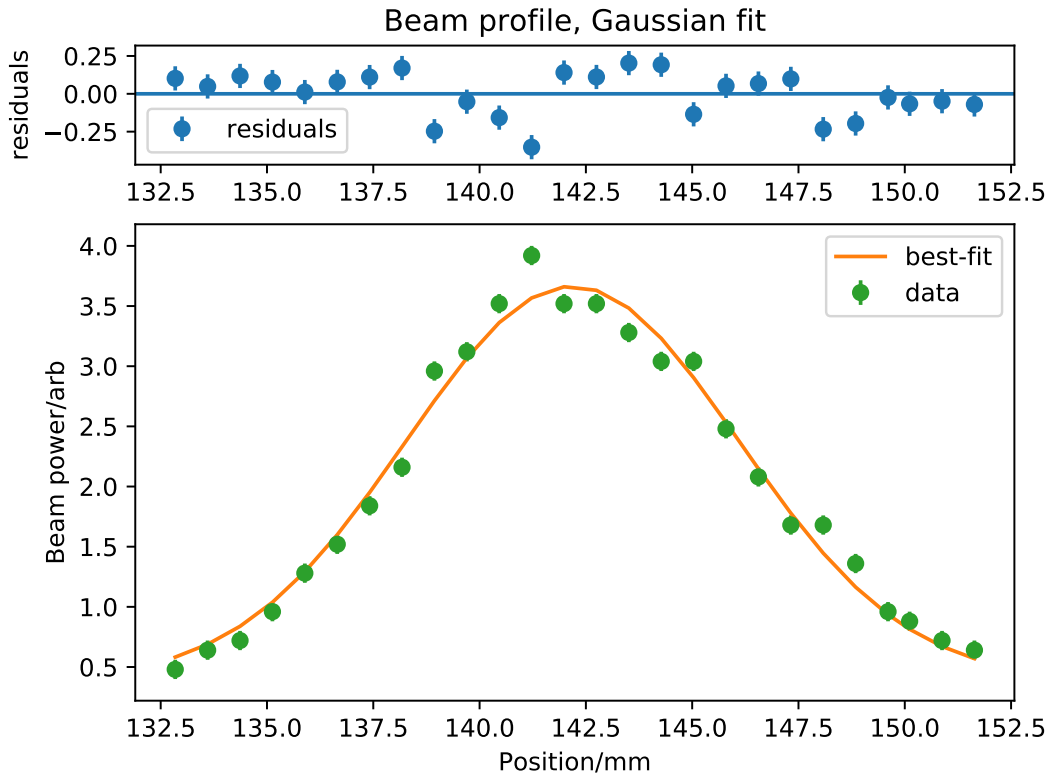


Figure 3.21: Laser beam profile when defocussed to ~ 20 mm diameter visible spot size. The profile is well described by a Gaussian curve with a small baseline offset.

The profile of the laser spot is not uniform, however. It would in theory be possible to focus the laser down to a point, small enough to ensure virtually all the power is incident on the absorbing area of both the power meter and the bolometer sensor. Unfortunately, this proved to be extremely sensitive to small changes of alignment (such as plugging and unplugging cables in the apparatus, or adjusting the power meter gain), due to the combination of the small surface area of the bolometer sensor and the comparatively large distance between the sensor and the laser (initially 20 cm, then 60 cm in later experiments). For consistency the laser was defocussed to produce a large spot size (~ 20 mm diameter), which was much more tolerant of alignment errors caused by small movements of the laser and sensor housing. Due to the non-uniformity of the spot, it is necessary to measure the profile of the spot in order to determine whether we need to account for the laser profile when scaling the power measured by the power meter onto the smaller bolometer sensor, or if assuming a uniform spot introduces an acceptably small error.

Figure 3.21 shows the laser beam profile. This was measured by placing a narrow slit between the defocussed laser and the power meter and scanning the laser horizontally, keeping the slit and power meter fixed, and recording the measured power

on the power meter. The profile fits a Gaussian curve, with a standard deviation of (3.92 ± 0.20) mm and a small (0.39 ± 0.13) mm baseline offset, likely due to ambient light in the lab. The laser beam was visually far more uniform in the vertical direction than the horizontal direction, so any error resulting from assuming uniformity in the vertical direction is significantly smaller than the error in the horizontal direction. We therefore assumed a uniform beam vertically, for simplicity.

By integrating the beam profile over the dimensions of the sensors, of width w and height h , we obtain the total power which would be absorbed on each:

$$P^{tot} = h \times \int_{-w/2}^{w/2} P(x) dx \quad (3.4.1)$$

Note that $P(x)$, the beam profile in Figure 3.21, is the power as a function of horizontal position, but integrated along the vertical length of the power meter. We should therefore use $h = 1$ for the total power measured by the power meter, and $h = h_{bol}/h_{pm}$ for the bolometer sensor.

If the beam spot was uniform, the ratio of the total power absorbed by the power meter and the sensor would be equal to the ratio of their areas. Therefore, the difference in the ratio of the total calculated powers and the ratio of areas gives the fractional error which results from assuming uniformity. The laser power meter sensor is a square, with area 13 mm^2 and hence width and height $\sqrt{13}$ mm. The bolometer sensor is 1.3 mm wide and 3.8 mm high. The ratio of total power to area is:

$$\frac{P_{bol}^{tot} / P_{pm}^{tot}}{A_{bol} / A_{pm}} = 1.027 \quad (3.4.2)$$

We can see that there is a $<3\%$ correction to the power density when accounting for the non-uniformity of the beam, compared with simply scaling the measured power by the sensor area. This is a small effect, and there are likely to be more significant sources of error in calculating the expected measured power on the bolometer. For example, the absorption coefficient for the gold sensor is well known for the laser wavelength, but it will depend on the roughness and cleanliness of the gold foil. There is also the possibility of stray light reflecting off the metal housing of the sensor, since the gold foils are slightly recessed into this housing. We have tried to minimise these effects, for example by using a bolometer channel with as few angled surfaces around it as possible to reduce reflections, but it is unlikely that we have eliminated all such effects, and the errors introduced are extremely difficult to quantify.

3.4.2 Sensitivity calculation: manual calibration method

In order to calculate the power absorbed by the bolometer sensor, we first needed an accurate measure of the sensitivity. This is calculated by calibrating the bolometer sensor. The new bolometer electronics has a novel calibration procedure, whose accuracy still needed to be properly measured.

To get an idea of the sort of sensitivity to expect from the new calibration method, we first performed the calibration with the method used by previous generations of electronics, which is described further in Section 2.1.4. Some tokamaks, such as JET, have electronics which automate the calibration procedure, but can only calibrate one channel at a time. In the tests we carried out, the procedure was performed manually. We used a potentiometer to balance the bridge before applying additional heating, a power supply to supply the two voltage levels, and a digital multimeter to read the output voltage.

The bolometer sensor was connected to a specially designed circuit, in which the sensor formed one of the 4 paths of the Wheatstone bridge, with the reference resistors shorted out. This bridge was balanced with the potentiometer, until an output voltage of -0.001 mV was measured for $U_1 = 0.5$ V input voltage. The input voltage was then increased to $U_2 = 5$ V. After waiting several seconds for the bridge to reach equilibrium, a voltage of 7.07 mV was recorded. The sensor resistance, with the reference resistors shorted out, was $R_{m||m} = 601.1$ Ω . The ohmic heating power is given by $P_{OH} = (U_2^2 - U_1^2)/(4R_{m||m})$, where U_1 and U_2 are the low and high input voltages respectively, and $R_{m||m}$ is the resistance of the sensor with the reference resistors shorted out. Then the effective capacity C is [73]:

$$C = 2 \times \frac{dU}{P_{OH}} \times \frac{1}{U_2} \quad (3.4.3)$$

With the above measurements, we calculated an effective capacity of $0.2750(4)$ W^{-1} , where the value in brackets signifies the uncertainty on the final digit. Multiplying this by the drive voltage gives the sensitivity: $S = 5.459(9)$ VW^{-1} for a 20 V drive, or $S = 4.945(8)$ VW^{-1} for an 18 V drive as used in the BOLO8 system. Note however that this calibration has been done with DC voltages, and so produces a DC sensitivity. The bolometer system uses an AC drive voltage in reality, and so corrections must be applied to this calculated sensitivity if it is to be used with AC data [57]. In this work, the DC sensitivity was used as a benchmark figure, to be compared with the sensitivity as calculated by the new AC method (Section 3.4.3). The sensitivity depends on the drive frequency and the electrical properties of the system such as cable resistance and capacitance, as will be seen in Sections 3.4.7 and 3.4.8, but should be close to the DC sensitivity.

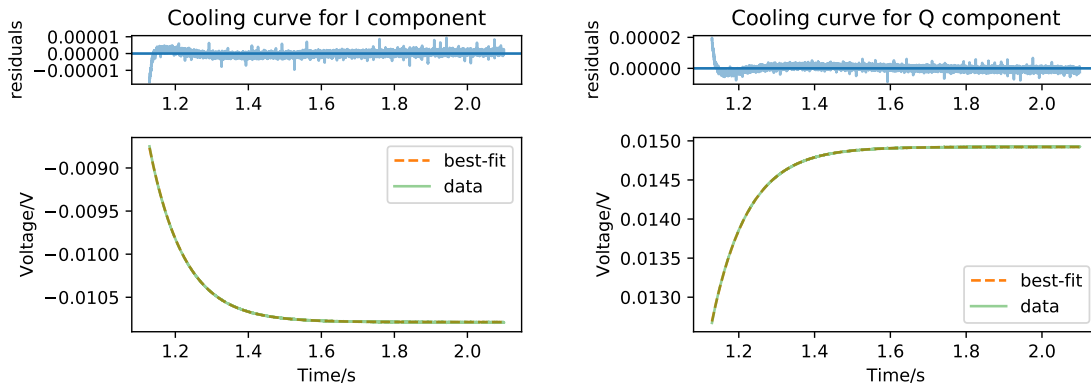


Figure 3.22: Cooling curves for the in phase (I) and quadrature phase (Q) components of the bridge output voltage for a calibration. Both are well described by an exponential decay with a single time constant, as in Equation 2.3.5.

3.4.3 Sensitivity calculation: BOLO8 calibration method

The BOLO8 calibration method was the second sensitivity calculation method to be evaluated, and has been described in Section 2.2.1. This was much simpler to perform than the manual calibration method from the end user’s perspective, since the method is fully automated. Due to its speed and simplicity, a calibration was performed before every test shot during the lab testing, giving relevant calibration parameters for a wide range of operating configurations.

Sections 3.4.7 and 3.4.8 describe in more detail the sensitivity calculation for a range of cable lengths and drive frequencies. Here we give just one calibration example, for test shot 1558. This was performed with a gold foil sensor, and 1 V bias voltage applied for 1 second.

Figure 3.22 shows the cooling curves that were measured and fitted in order to obtain the sensitivity. Clearly the fits represent the data very well. The small deviation at the beginning of the time trace is due to filter ripple from the digital filtering, but has no large effect on the fit. The fit error on ΔV for both curves is around 0.01 %, and by dividing by the applied ohmic heating power we calculate a sensitivity of $4.050(58) \text{ VW}^{-1}$ for 18 V amplitude drive voltage. As expected, this is slightly lower than the DC sensitivity calculated in Section 3.4.2, due to AC effects. The uncertainty on the sensitivity measurement has been calculated by considering the RMS noise of the calculated heating power, which is dominated by noise in the current measurement. This could in theory be reduced by smoothing the current trace; there would be no loss of information since the heating power is simply averaged over the entire heating period. If we average the sensitivity values calculated over 13 shots used in the laser power scan (Section 3.4.5), we get $4.048(16) \text{ VW}^{-1}$.

A bonus of this method is that we measure the cooling time as well. This was calculated to be 0.095 921(6) s for this shot, and 0.095 942(15) s when averaged over the same 13 shots as the sensitivity. A good test of the accuracy of the cooling time is in the ability to measure a square wave for the absorbed power when the laser is set at a constant power and repeatedly switched on and off. We shall demonstrate this in Section 3.4.4.

Note that the uncertainty on the power measurements depends both on the uncertainty in the calibration measurements and also on other factors such as the size of the sensor, the reflectivity of the foil (and indeed any reflections from the metal housing of the sensor) and the laser profile. Whilst the calibration results seem surprisingly accurate, actual power measurements will have higher uncertainties as a result of these additional sources of systematic error.

3.4.4 Absolute power accuracy

A key test of the new system is its ability to provide an absolute measurement of power. Whilst it is sufficient for many experimental programmes to measure relative changes in power, it is necessary for some experiments to know the absolute amount of radiation, for example in power balance studies. We therefore performed a relatively simple test to give an indication of the uncertainty in the measurement of the absolute power.

The test consisted of setting up the laser with a low frequency square wave intensity: on for maximum power for 1 second, then off for 1 second. The laser power was first recorded with the power meter. This was set up at the same distance from the laser as the bolometer sensor, to ensure the same power loss through scattering in the air and beam divergence. Once the laser power had been measured on the power meter, it was moved laterally until the beam was centred on one bolometer channel, and the same test was repeated.

By using an on/off waveform, rather than a constant one, we can ensure that we are correctly measuring the true absorbed power without any systematic offset. The periods when the laser is off should read exactly 0 absorbed power. Furthermore, by using a slow waveform we can see just how constant the measured power is, which both gives an indication of whether the calculated cooling time is correct and shows the level of noise present in the measurement.

With the laser defocussed to produce a spot size comparable to that in Section 3.4.1 and driven with a 0.5 Hz square wave voltage supply, we measured the absorbed power on the silicon power meter. We recorded the measured voltage both when the

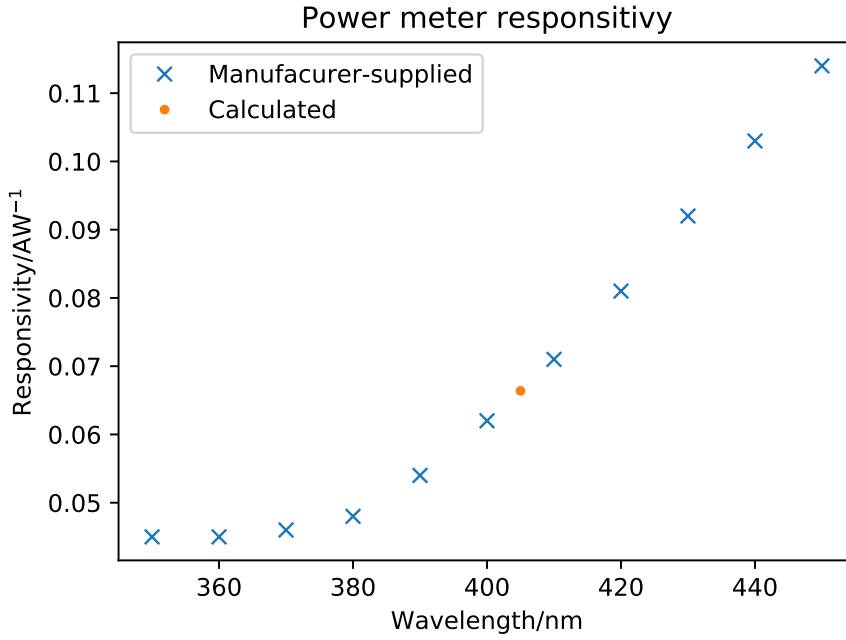


Figure 3.23: Power meter responsivity for a selection of wavelengths. We have extrapolated from the responsivity values supplied by the manufacturer at a number of wavelengths (blue crosses) to calculate the responsivity at the wavelength of the laser used for these tests (orange dot).

laser light was on and when it was off, and subtracted the latter from the former, to ensure we are measuring only light from the laser and not background light.

In order to convert the voltage reading V from the power meter into an absorbed power P_{pm} , we need to know the responsivity R , which is a function of wavelength, and the gain factor G :

$$P_{pm} = \frac{V}{G \times R(\lambda)} \quad (3.4.4)$$

The gain factor can be looked up in the PDA36A manual, for all possible gain settings. In contrast, the responsivity is a continuous function of wavelength, so must be calculated explicitly.

Figure 3.23 shows a subset of the responsivity values for a selection of wavelengths, taken from a spreadsheet provided by the manufacturer [82]. Using a cubic spline interpolation across all these data points, we calculate the responsivity at the laser wavelength of 405 nm to be 0.066 AW^{-1} , which is shown as an orange dot on the figure.

With a gain of 40 dB, we measure 0.1176 V from the power meter when the laser is off, and 8.84 V when the laser is on. Using Equation 3.4.4, we calculate an absorbed power of 0.86 mW on the power meter, or a power density of 66.3 Wm^{-2} . We therefore expect the power incident on the bolometer sensor to be 0.328 mW,

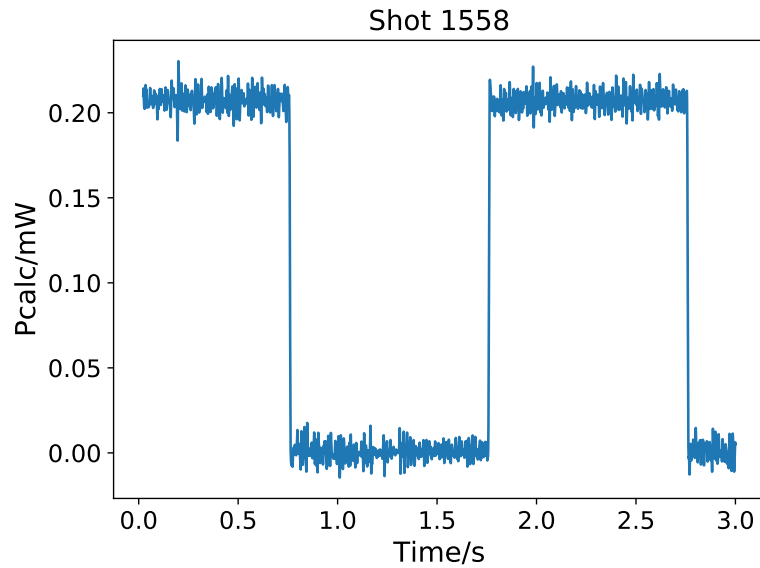


Figure 3.24: Calculated power using the BOLO8 calibration for a slowly flashing, defocussed laser beam. The square wave form of the laser power has been reproduced well, and the power measured when the laser is on (0.207(7) mW) is very close to what we expect the bolometer sensor to measure (0.205 mW), based on measurements from the silicon power meter and accounting for the detector area and the reflectivity of gold at the laser wavelength of 405 nm.

assuming a uniform intensity. However, the bolometer sensor is made of gold foil, which has a reflectance of 0.376 at 405 nm [67], and so only 62.4% this power will be absorbed. We therefore expect to see 0.205 mW measured by the bolometer.

Figure 3.24 shows the measured power using the bolometer. We take the average of all points when the laser is on (considered to be those where $P > 0.8P_{\max}$), and subtract the average of all points where the laser is off ($P < 0.2P_{\max}$), to get the net power absorbed from the laser by the sensor. We estimate the uncertainty on this calculation using the sum in quadrature of the standard deviations of the laser-on and laser-off periods. This gives the power absorbed by the bolometer as 0.207(7) mW, in excellent agreement with the expected power.

3.4.5 Laser power scan

We also performed a simple test of the linearity of the system. By adjusting the supply voltage of the laser's waveform generator, we could increase the laser power. According to its data sheet, the laser power output is linear with the supply voltage [83], so we expect to measure a linear response on the bolometer.

We used the same waveform as in Section 3.4.4, again to ensure that there was no systematic offset in the measurements. Similarly, the mean difference in measured




Figure 3.25: Power measured by the bolometer as a function of the modulation voltage applied to the laser. The relation is almost linear, though a small quadratic term was needed for the best fit.

power when the laser was on and off was used to measure the absorbed power from the laser by the bolometer sensor.

Since the absolute power density was not needed, the laser was focussed to a bright spot instead of a diffuse one with approximately constant power density. This meant the absorbed power was higher, and therefore allowed us to measure the bolometer linearity over an order of magnitude change in the measured signal. Since the laser was rated to 5 mW and the bolometer has approximately 60% absorption at 405 nm, we expect to measure a maximum of 3 mW at 1 V, decreasing linearly to 0 W at 0 V.

Figure 3.25 shows the power measured by the bolometer as the modulation voltage applied to the laser was varied. The laser data sheet implies a linear relationship, but a better fit was obtained with a quadratic polynomial of the form $P = c_0 + c_1V + c_2V^2$, with $c_0 = -0.080(3)$ mW, $c_1 = 3.35(1)$ mW/V and $c_2 = -0.17(1)$ mW/V². There is thus a small non-zero offset, and a slightly quadratic relationship, though the quadratic coefficient is only 5% of the magnitude of the linear coefficient, so overall the system is quite linear.

The laser data sheet does not say how precise the linear relationship between supply voltage and output power is, so it is not possible to say definitively whether this deviation from linearity is due to the laser or the bolometer system. It is fair to assume that it is a combination of both, though we cannot say for certain from this test that the bolometer system's linearity holds to anything better than 5%.

3.4.6 Laser frequency scan

A further test of the accuracy of the cooling time estimate can be made by varying the frequency of the square wave laser waveform. As the frequency of the laser increases, there is less time for the sensor to heat up and cool down, which reduces the maximum measured amplitude. However, the calculated absorbed power using the bolometer formula 2.1.9 (with the cooling time accounted for) should be independent of the laser frequency.

For this test, the laser was again focussed in order to provide the largest possible signal. The power calculation was performed using a Gaussian deconvolution filter with a bandwidth of 1 kHz (the maximum laser frequency used in the test) for all laser frequencies. This does increase the noise level on the measurements compared with the best possible case, particularly for the lower frequency shots, but ensures consistent processing across all laser frequencies.

We again used the difference between averaged "laser on" and averaged "laser-off" measurements to calculate the laser power. We treated "laser on" as all data within 80% of the maximum measurement in the time trace, and "laser off" as all data lower than 20% of the maximum measurement. Whilst this works reasonably well at low frequencies, where the measured power is a well-defined square wave, it is less successful at higher frequencies approaching the filter cut-off frequency. This is because the higher harmonics of the square wave are suppressed, and the power starts to look more like a sine wave.

Figure 3.26 shows the calculated power using this on/off averaging method. The power is reasonably constant over two orders of magnitude in laser frequency, but does show strong frequency dependence at higher frequencies. This is partly due to the difficulties already mentioned regarding the processing method for calculating the power in this way.

The effect of the FIR filters, both the windowed-sinc filter on the FPGA and the Gaussian deconvolution filter used in post-processing, also contribute to the frequency response. Both have cut-off frequencies of 1 kHz, so we expect attenuation by approximately a factor of 4 for the 1 kHz measurement: a factor of 2 attenuation




Figure 3.26: Laser power absorbed by the bolometer sensor, as a function of laser frequency, with a square wave laser pulse. Although the power is reasonably constant over 2 orders of magnitude, the 1 kHz filter bandwidth reduces the calculated power at the highest frequency and there is a small increase in the measured power at around 10^2 Hz, likely due to filter ripple in the FPGA's low-pass filter.

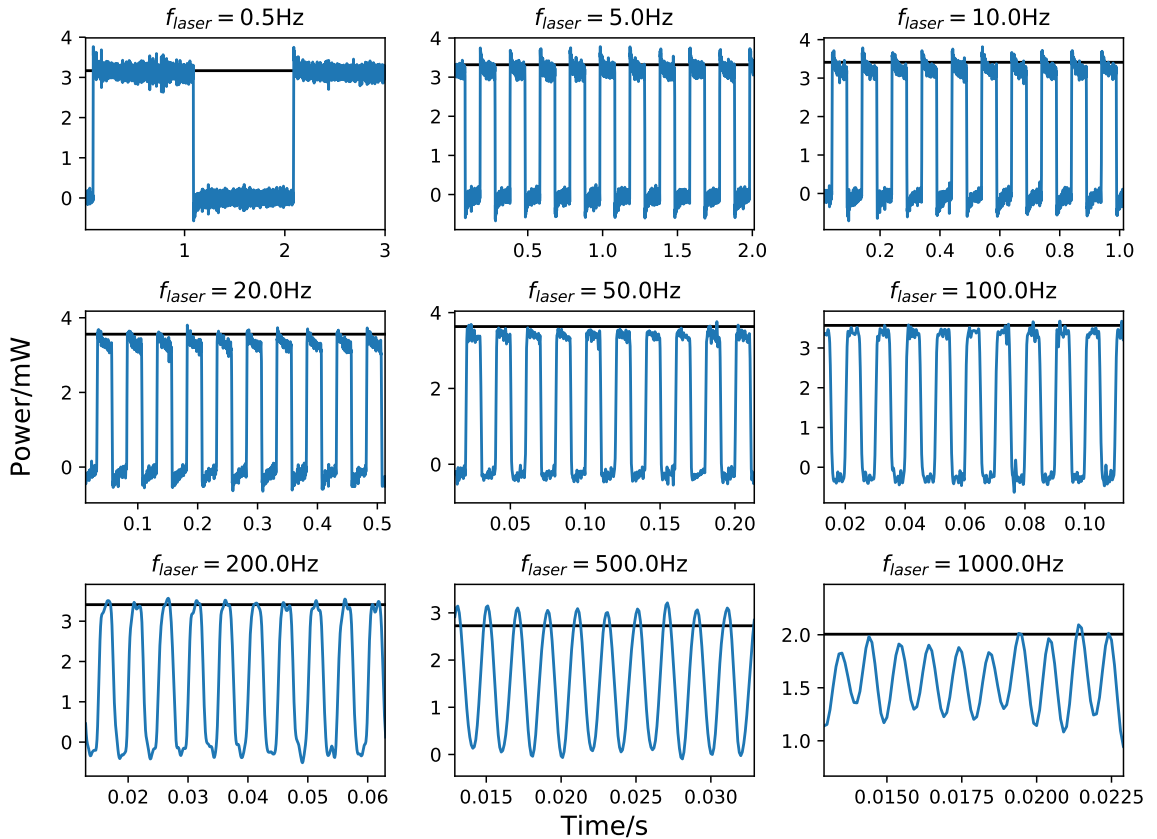


Figure 3.27: Time traces showing the calculated power for different laser frequencies at the same laser power. The calculated average power, as shown in Figure 3.26, is shown as a black horizontal line in each plot. This agrees well with a visual estimate of the power in all but the highest frequency cases, where the attenuation of higher harmonics of the square wave means the signal looks more sinusoidal and so there is no flat top to average over.

in each of the FPGA filter and the deconvolution filter. Additionally, filter-induced ripple in the signal contributes to a slight over-estimation of the mean on and off powers in the 20, 50 and 100 Hz traces, as can be seen from Figure 3.27. This ripple mainly comes from the FPGA's windowed-sinc filter, since this type of filter prioritises good frequency-domain performance over minimising ripple. Finally, there is some irregularity in the amplitude in the 1 kHz trace, which is likely to be due to the limited sample rate of the output providing too few points to well resolve the maxima and minima of each laser cycle.

It can be seen from Figure 3.27 that the power calculation using the mean on/off method becomes less accurate at frequencies approaching the filter cutoff frequency. Recall that this method requires both that the power measurement when the laser is on is sufficiently flat and that the measured power drops to zero when the laser is off. This is clearly not the case for the 500 Hz and 1000 Hz traces because of the suppression of higher order harmonics of the square wave.

It is also possible to exploit our a priori knowledge of the laser waveform shape and use spectral methods to determine the laser power. We know that the waveform is a square wave, centred around $P_{laser}/2$ and with amplitude (half height) $P_{laser}/2$. If $S(t)$ is a square wave with amplitude 1 and offset 0, then the laser power is given by:

$$P = \frac{P_{laser}}{2} (S(t) - 1) \quad (3.4.5)$$

The Fourier series of the waveform described by Equation 3.4.5, for a laser of frequency F_{laser} , is:

$$\tilde{P}(f) = \frac{4}{\pi} \frac{P_{laser}}{2} \sum_{n=1,3,5,\dots}^{\infty} \frac{1}{n} \sin(2\pi n F_{laser} f) \quad (3.4.6)$$

Therefore, if we take the Fourier transform of the measured power, the amplitude of the peak at F_{laser} ($n = 1$), is:

$$|\tilde{P}(F_{laser})| = \frac{2P_{laser}}{\pi} \quad (3.4.7)$$

We can therefore trivially calculate the laser power from this amplitude. Figure 3.28 shows the amplitude of the relevant peak for each laser frequency. Figure 3.29 compares the calculated laser power using the on/off mean method and this spectral method. There is good agreement between both methods at intermediate frequencies, and at the highest frequencies the mean on/off method over-estimates the power compared to the spectral method, because the off mean power is not zero.

At the lowest frequency of 0.5 Hz, the spectral method overestimates the power compared with the on/off method, and its estimation is unreasonably high. Further

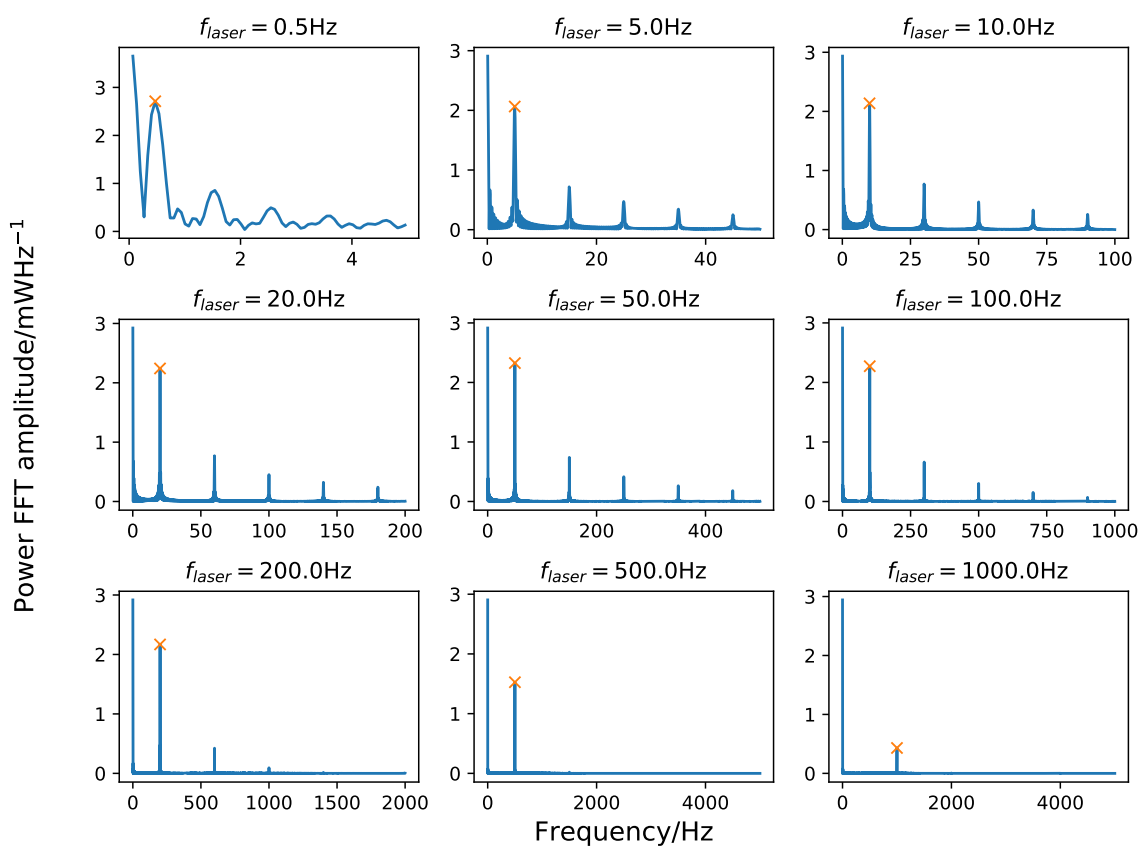


Figure 3.28: Spectra of laser frequencies, with laser frequency peak highlighted. The calculated peak height agrees well with the visual estimation for all frequencies. The attenuation of higher harmonics of the square waves at the higher frequencies, which caused problems for the on/off averaging method, can also be clearly seen.




Figure 3.29: Comparison of on/off mean and spectral methods of calculating laser power at different laser frequencies. There is good agreement at intermediate frequencies, and at higher frequencies the spectral method suggests the calculated power falls of more than the averaging method. At the lowest frequency the spectral method appears to overestimate the power compared with the on/off method.

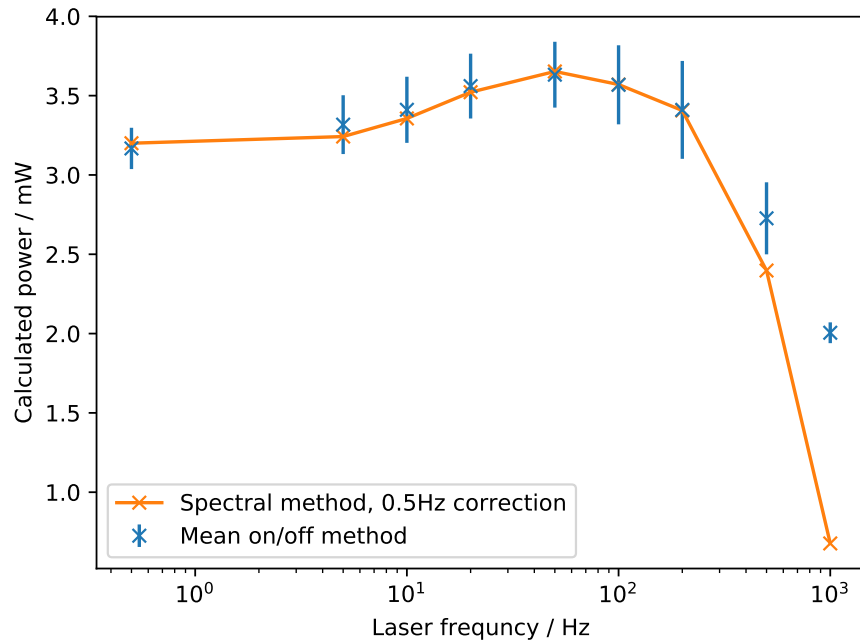


Figure 3.30: Comparison of on/off mean and spectral methods of calculating laser power at different laser frequencies, after removing the systematic offset from the 0.5 Hz spectrum. There is now good agreement between the two methods at all but the highest frequencies, where we expect the spectral method to be more accurate.

inspection revealed the presence of ripples in the spectrum at 0.5 Hz and its harmonics in the measurements of all laser frequencies; the amplitude at 0.5 Hz was the same in all other shots. It therefore seems that there is an additional amplitude offset at 0.5 Hz which is not present (or significantly smaller) at higher frequencies. By subtracting from the 0.5 Hz laser amplitude the mean amplitude of the spectra at 0.5 Hz from all the other traces, we can remove this systematic offset. The result is shown in Figure 3.30. There is now much better agreement at the lowest frequency, and the roll-off at higher frequencies is still more pronounced in the spectral method than the on/off mean method, as expected.

Generally, the laser power is fairly constant over 2 orders of magnitude in frequency: it is within 10% from 0.5 Hz to 200 Hz, before the filter roll-off starts to become apparent at 500 Hz. However, there is a noticeable upwards trend in the measured power before the roll-off. Whilst this could be due to a slight over-estimation of the cooling time of the sensor, which is exacerbated at higher frequencies, it is more likely that it is caused by filter ripple produced by the FPGA's windowed-sinc low-pass filter. Using a different type of low-pass filter (such as a Gaussian filter) may improve this, at the cost of poorer performance in the frequency domain.

3.4.7 Cable length scan

This section, and Section 3.4.8, are more concerned with the physical setup and operation of the diagnostic. We have already seen in Sections 3.4.2 and 3.4.3 that AC effects affect the sensitivity of the diagnostic. These tests were performed to quantify those effects.

We first investigated the effect on the calibration of increasing the cable length. A series of calibrations was performed, with the cable length increased by joining multiple lengths of cable together. The cable used was off-the-shelf CAT-7 cable, the same type as is to be used in the MAST-Upgrade installation. Since the cable has a finite resistance and capacitance, it attenuates both the AC excitation voltage applied to the sensor, and also the output voltage from the sensor. A longer cable has higher resistance and capacitance, and hence a greater attenuation. This means that for a given bridge imbalance (i.e. for a given absorbed power), the bolometer signal will be smaller with a longer cable.

The DC method of calibrating the sensor to calculate the sensitivity, as described in Section 3.4.2, does not account for cable attenuation with an AC drive, since this attenuation is not present at DC. Corrections to the DC sensitivity to account for AC effects have been derived [57], but they require knowing to high precision the resistance and capacitance of each individual cable. Not only is this difficult to measure, but these quantities may change over time, or if the cable is stressed or otherwise moved, which can conceivably happen between successive calibrations.

The BOLOS calibration method is an AC method, and so intrinsically includes such effects. We therefore expect to measure different sensitivities for different cable lengths: these will be accurate reflections of the sensitivity of the bolometer sensor during operations.

Figure 3.31 shows the variation of sensitivity with cable length, normalised to the shortest available cable length. As expected, the sensitivity decreases with increasing cable length. The effect is more pronounced for the higher 30 kHz drive frequency, which is also expected. The attenuation is significant: failing to account for it after performing a DC calibration would introduce a systematic error on the order of 15% for an 18 kHz drive frequency, and 30% for a higher frequency drive of 30 kHz.

Figure 3.32 shows the calculated power for each of these cable lengths, using the calculated sensitivities and cooling times, when the bolometer was illuminated with the same 0.5 Hz laser waveform used in Section 3.4.4. The variation in calculated power over cable length is now only around 3%, significantly lower than the 15–30% variation in the sensitivity, which demonstrates that the calibration procedure successfully accounts for cable attenuation automatically.




Figure 3.31: Sensitivity as a function of cable length for two different drive frequencies, normalised to the sensitivity for the shortest cable length. In both cases the sensitivity decreases as the cable length increases, due to AC attenuation effects in the cable. The sensitivity falls off at a higher rate at the higher frequency, as the AC attenuation effects are more pronounced.




Figure 3.32: Calculated power as a function of cable length, using the mean on/off method. There is small variation in the calculated power, but it is much lower than the variation in the sensitivity. This suggests that the new calibration method is successfully accounting for AC effects.

3.4.8 Drive frequency scan

Finally, a test was performed to measure the effect of drive frequency on the sensitivity and hence the measured power. One of the advantages of the new BOLO8 system over older bolometer electronics is the ability to vary the frequency of the AC excitation voltage. This allows the system to operate in a region of the spectrum away from strong sources of electrical noise, and was taken advantage of during JET operations, described in Section 3.2.2, to avoid a strong noise source at 20 kHz. However, changing the drive frequency to reduce measurement noise is only advantageous if the measurement itself is still accurate. We were interested to see if there was any systematic error involved in moving away from the design frequency of 20 kHz.

For these tests, we focussed the laser to produce a larger signal (again, relative power between shots is the most interesting quantity here) and modulated it with a slow, 0.5 Hz square wave voltage supply. We then scanned through the available range of drive frequencies supported by the system, from 5 kHz to 45 kHz. At each frequency in the scan a calibration was performed with the laser off, then the laser was turned on and the bridge output voltage measured and used to calculate the power. We used the on/off averaging method described in Section 3.4.4 to calculate the absorbed power on the bolometer sensor.

Figure 3.33 shows the power calculated during this scan, normalised to the laser power output at the same distance and focus measured by the silicon power meter. Due to the reflectivity of gold, as mentioned in Section 3.4.4, we expect this ratio to be around 63%. Even though (with the exception of the 5 kHz measurement), all drive frequencies are within a few percent of this expected value, there is a noticeable trend for higher measured power peaked at around 20 kHz.

This measurement bump can also be seen in the sensitivity calculated for each drive frequency. Following similar logic to Section 3.4.7, the attenuation of the signal due to cable resistance and capacitance is greater for higher frequencies, and so we expect to see a steady, monotonic drop in the sensitivity as the drive frequency increases. Indeed, this is what we see if we plot the voltage amplitude with the laser turned on, as can be seen in Figure 3.34. However, Figure 3.35 shows that the sensitivity does not uniformly decrease: it plateaus at approximately 20 kHz. This suggests that there is some frequency dependence in the calibration procedure.

In fact, the peak is likely to be at 19.3 kHz, which is the resonant frequency of the LC filter isolating the offset DAC from the measurement ADC. Recall from Chapter 2 that the BOLO8 hardware contains LC notch filters on the two differential inputs to the main ADC, which are intended to isolate the offset DAC from the bridge by providing a high impedance at the expected drive frequency of around 20 kHz.




Figure 3.33: Power measured with the bolometer sensor as a function of drive frequency, normalised to the power measured by the power meter. Due to the reflectivity of gold, we expect this ratio to be around 63%. Whilst this is true at around 20 kHz, the ratio falls off either side of this frequency, with a more dramatic fall off at lower drive frequencies.




Figure 3.34: Peak bolometer output voltage as a function of drive frequency. The peak voltage reduces at higher drive frequencies due to AC attenuation effects in the cable.




Figure 3.35: Variation of calculated sensitivity with drive frequency. Whilst there is an overall decrease in the sensitivity as the drive frequency increases, it does not match that of the peak bridge voltage shown in Figure 3.34. In particular, there is a flattening off of the sensitivity around 20 kHz to 30 kHz.

The offset DAC has an output impedance of $50\ \Omega$; this acts as a voltage divider to reduce the bridge voltage if the offset DAC is not properly isolated. Although the Schottky diodes provide good isolation during measurement mode (when they are reverse biased by the negative voltage output by the offset DAC), they do not of course isolate the offset DAC during the heating phase of the calibration. During heating, we instead rely on the LC notch filter to both reduce the noise output by the offset DAC and prevent it from loading the circuit. The effectiveness of the notch filter varies significantly with frequency (the notch is around 10 kHz wide), and so different drive frequencies will result in different amounts of loading the bridge.

Recall that the calibration method involves applying a DC voltage (and hence current) to ohmically heat the sensor's measurement resistors. This DC heating is applied in addition to the ohmic heating by the AC drive voltage. If there is a different amount of AC heating during the DC heating phase (due to bridge loading) and the cooling phase (when the diodes prevent bridge loading), then the simple estimation of the net applied heating power ($P_{OH} = 2V_{DC}I_{DC}$ during heating, and $P_{OH} = 0$ during cooling) is not completely accurate.

Figure 3.36 illustrates this effect. It shows the ratio of the measured output voltage amplitude of the bolometer sensor, with a fixed bridge imbalance, when the off-




Figure 3.36: Ratio of bolometer output voltage amplitudes with forward and reverse biasing, for the same bridge imbalance. Away from the 19.3 kHz peak impedance of the LC filter, the finite output impedance of the offset DAC loads the bolometer bridge circuit and acts as a voltage divider to reduce the voltage measured by the ADC. By normalising to the reverse-biased voltage measurement (when the offset DAC is isolated), AC cable attenuation effects have been removed.

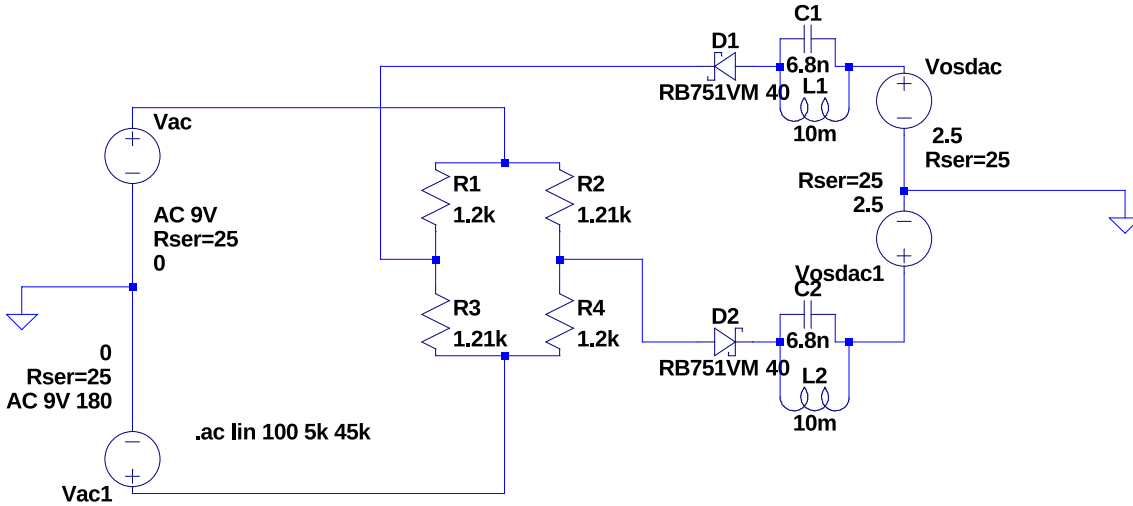


Figure 3.37: LTspice circuit to model the bolometer sensor and BOLO8 system. The bolometer sensor is represented by a slightly unbalanced Wheatstone bridge, and the main and offset DACs by differential voltage sources. The diodes and LC filters are also included.

set DAC was applying a forward bias to the diodes (V_f) and when it was applying a reverse bias (V_r). Normalising to the reverse bias voltage reading removes other frequency effects, such as the cable attenuation. The overall trend of Figures 3.33 and 3.36 is remarkably similar, which strongly suggests that it is this variation in the ohmic heating during the calibration that is responsible for the variation in both sensitivity and therefore calculated power.

This effect has also been modelled using the circuit analysis software LTspice [84]. We designed a circuit consisting of a slightly imbalanced Wheatstone bridge to simulate the bolometer sensor, a pair of differential voltage sources to simulate the main and offset DACs, and the Schottky diodes and the LC filters of the offset circuitry. The simulated circuit is shown in Figure 3.37. No cable resistance or capacitance was included as these values were not known for the cables used in our tests, so the results should be compared to the V_f/V_r normalised data presented in Figure 3.36. We simulated varying the drive frequency between 5 kHz and 45 kHz, to match the lab tests, and measured the output voltage difference between the bridge and the diodes, in the same place in the circuit that the main ADC on the BOLO8 measures this voltage.

We also attempted to derive an analytic formula for the expected bridge output voltage, as a function of both bridge imbalance and drive frequency. This extends the simple Wheatstone bridge circuit of the bolometer sensor to include the offset DAC's output impedance and the notch filters, but neglects the diodes and the resistance and capacitance of the cables. The effective circuit is shown in Figure 3.38, where V_{dr} and V_m are the drive and measured output voltages respectively, Z_1 to Z_4 are




Figure 3.38: The effective circuit used to derive the analytic formula for the bridge output voltage during calibration. The bridge resistors have impedances Z_1 to Z_4 , whilst Z_5 represents the output impedance of the main DAC, which supplies the drive voltage V_{dr} . The impedance of the offset DAC in series with the two LC filters is represented by Z_{os} , and the bridge output voltage measured by the ADC is V_m .

the bridge resistors, Z_5 is the output impedance of the main DAC and Z_{os} is the series impedance of the two LC filters plus the output impedance of the offset DAC:

$$Z_{os} = R_{os} + 2 \left(\omega C + \frac{1}{\omega L + R_I} \right) \quad (3.4.8)$$

Here, R_{os} is the output resistance of the offset DAC, ω is the angular frequency of the excitation voltage, C is the capacitance of the LC filters' capacitors, and L and R_I are the inductance and internal resistance of the filters' inductors respectively.

We derive the Thévenin voltage V_{th} and impedance Z_{th} by considering the Thévenin equivalent circuit, shown in Figure 3.39. Then the measured voltage is calculated using a simple potential divider:

$$V_m = V_{th} \left(\frac{Z_{os}}{Z_{os} + Z_{th}} \right) \quad (3.4.9)$$

We can evaluate Equation 3.4.9 for the components used in the LTspice simulation. Their properties were in turn taken from the data sheets of the BOLO8's DACs and LC filter components. Using the same resistance values as those in the LTspice simulation, we can compare the analytic and numerical models to see whether the analytic model is sufficient to describe the circuit.

Figure 3.40 shows the results of the comparison. The analytic model has the same qualitative form as the simulation, but there is a substantial discrepancy between them. Unlike the original LTspice model, the analytic model does not include the diodes (which do have some finite resistance). However, a second LTspice simulation with the diodes removed showed that the diodes do not have a significant impact on




Figure 3.39: Thévenin equivalent circuit for bolometer sensor and BOLO8 system. The bridge resistors and the main DAC's output impedance are described by the Thévenin impedance Z_{th} , and the Thévenin voltage V_{th} represented the measured voltage difference across the diagonal of the Wheatstone bridge measured by the ADC.




Figure 3.40: Comparison of bridge voltages as a function of frequency for the analytic model, LTspice simulations with and without Schottky diodes, and the measurements of the ratio of the bridge voltage amplitude with forward and reverse bias voltages (V_f/V_r) from PPPL. The LTspice simulations agree well with the PPPL data, and the presence of the diodes does not seem to have a significant impact on the voltage. The analytic model does show the same qualitative behaviour: a peak at 19.3 kHz which falls off either side, with a more rapid fall off at lower frequencies. However, the fall off rate is much lower than for either the simulations or the PPPL data.

the measured bridge output voltage.

Also plotted on Figure 3.40 is the measured data V_f/V_r from Figure 3.36. Since the normalisation by V_r removes the cable attenuation effects, we expect good agreement (apart from a linear scaling due to the different bridge imbalance) with the predicted curves. The agreement with the LTspice simulations is excellent, which demonstrates that the circuit components' parameters, as taken from data sheets, are suitable. However, the disagreement with the analytic model suggests that this model does not sufficiently capture all the properties of the complex bolometer/BOLO8 circuit.

These investigations do support the idea that the reduction in ohmic heating during calibration due to the offset circuitry's finite impedance is the cause of the variation in sensitivity, and therefore in calculated power, with drive frequency. However, it is not clear whether it is possible to derive an analytical correction factor for this effect. A simpler option would be to use the data in Figure 3.33 as an empirical correction factor to be applied to future measurements, though it is possible that no correction would be needed anyway. Figure 3.33 shows that there is only around a 5% reduction in the measured power between 10 kHz and 30 kHz, and even less in the 15 kHz to 25 kHz range in which the system is likely to operate, if shifting frequency is used solely for noise reduction. Running the system with drive frequencies far from the design frequency is unlikely to be routine.

3.4.9 Summary of PPPL tests

The tests at PPPL with the production ACQ2106/BOLO8 hardware demonstrated first and foremost the suitability of the FPGA BOLODSP firmware design and its ability to produce reliable data. It also confirmed that the greatly increased operational flexibility made possible using FPGA technology does not come at the cost of significantly reduced data quality. The new calibration method was shown to be robust to changes in the physical setup of the diagnostic — specifically cable lengths — and in the operation of the diagnostic at different drive frequencies. The effect of varying the laser power and frequency was also in line with expectations.

Finally, the absolute accuracy of the system was demonstrated, which showed that the calculated power was within a few percent of the expected power, based on measurements with a silicon power meter. Absolute bolometer measurements are difficult, but vital for the success of power balance studies, where the amount of power radiated by the plasma must be accurately measured. As an added bonus, a missing factor of two in the sensitivity calculation was discovered during testing, which had initially been hidden during the JET tests (Section 3.2) by the étendue reflection

correction, and would have been extremely difficult to locate once operations had started on MAST-U.

Ultimately, these tests demonstrate that the ACQ2106/BOLO8 system is ready for installation and operation on MAST-U, and is expected to produce high quality, reliable data.

3.5 Implications for installation on MAST-Upgrade

Delays in the completion of MAST-Upgrade meant that the planned installation and commissioning of the bolometer system is not scheduled until after the completion of this thesis. However, the tests done during the development of the new bolometer electronics have proven the system is fit for purpose. Several interesting aspects of the electronics and the sensors have been identified, such as the cooling time's dependence on pressure and the small variation of measured power with drive frequency, and these will be taken into account when interpreting the MAST-U bolometer data.

Proving the reliability of the real-time power output is particularly important, since we can now justify integrating the bolometer system into a future plasma exhaust control system. This will really exploit the capabilities of the FPGA, and greatly enhance MAST-Upgrade's credentials as one of the most advanced tokamaks in the world.

Integration of the diagnostic into the data acquisition systems of two tokamaks, JET and NSTX-U, means that a lot of the development work needed to enable the diagnostic to go through a shot cycle has already been completed. This will certainly shorten the commissioning time required when the system is installed on MAST-U. Successful demonstrations of the system on these high profile tokamaks has generated interest from the wider fusion community and lead to the BOLO8 system being ordered for TCV and NSTX-U, with other labs also interested in the technology.

Chapter 4

The Langmuir Probe diagnostic

In this chapter we describe the design of the divertor Langmuir Probe diagnostic. In a similar manner to Chapter 2, we outline the underlying physics of the diagnostic and its requirements. We then proceed to describe the system's hardware, FPGA firmware and associated software.

4.1 Principles

The Langmuir probe was one of the earliest methods employed to diagnose plasmas [29, 14]. In its most basic form, illustrated in Figure 4.1, it consists of an electrode inserted into the plasma, to which a voltage is applied relative to the vessel wall. The current is measured for a range of voltage biases, and from this several plasma parameters can be inferred.

The relative simplicity of the probe hardware makes it an attractive diagnostic, and indeed Langmuir probes are ubiquitous in magnetic fusion research. However, the




Figure 4.1: Cartoon overview of the Langmuir probe principle. A probe is inserted into the plasma, and a voltage V_{bias} is applied relative to the vessel wall. The current I_{probe} is measured for a range of bias voltages.

requirement to physically insert the probe into the plasma means the technique has several limitations. Firstly, the plasma conditions must be survivable by the probe itself, which in tokamaks effectively means that Langmuir probes can only be used in the edge and divertor regions. Reciprocating probes can be used to measure more extreme plasma conditions for short periods of time [85, 86, 87], but even these cannot measure deep into the core of the confined plasma without being destroyed. The importance of edge effects in tokamak plasmas does mean though that in practice this limitation does not dramatically reduce the usefulness of the Langmuir probe.

Secondly, the probe interacts with the plasma into which it has been inserted, and therefore alters the local plasma conditions. Care must be taken in interpreting the data from the probe, since we must relate the measurements of the local perturbed plasma to conditions in the unperturbed plasma farther from the probe. In fact, the successful interpretation of the data from the probe is typically equally as challenging as the development of the probe hardware and acquisition of the data itself.

Several arrangements of Langmuir probes can be used to infer plasma parameters. The most common technique is to sweep the voltage of a single probe whilst measuring the current, and to calculate plasma parameters using the measured current as a function of the applied voltage, known as an “IV characteristic”. This is described in Section 4.1.1. Other techniques, such as using two or three closely-located probes, can provide measurements on time scales of turbulent fluctuations [88, 89], but these suffer from inaccuracies due to plasma conditions being different at each of the different probes. For the MAST-U system, the method of sweeping the voltage on individual probes is used. We shall therefore confine our discussion to this method.

4.1.1 The IV characteristic

The form of the IV characteristic is primarily determined from the theory of the electrostatic sheath [90, 14, 29]. The theoretical form is shown on the left of Figure 4.2, and can be explained qualitatively as follows.

At very negative voltages, all electrons are repelled by the probe, and ions are collected. The current is limited by the rate at which ions can reach the probe: this is the ion saturation current I_{sat} . As the voltage increases, fewer electrons get repelled, which reduces the net current. Eventually, few enough electrons are repelled that the ion and electron currents balance and there is zero net current. The voltage at which this occurs is the floating potential V_f . If the probe were floating (i.e. isolated from the vessel ground), electrons would naturally collect on




Figure 4.2: Left: a theoretical IV characteristic, showing the ion saturation current I_{sat} , floating potential V_f , plasma potential (or space potential) V_s and the electron saturation current. Taken from [90]. Right: a real-world IV characteristic from a probe in an inductively-coupled discharge, taken from [91]. The form mostly follows the theoretical form, but the electron current does not saturate and the ratio of electron to ion saturation currents is much smaller.

the probe (due to their greater mobility) and charge it up negatively until it reached this potential.

As we increase the voltage still further, the charging effect of the electron collection is gradually cancelled out by the applied voltage. When the probe voltage reaches the same value as the voltage in the plasma far from the probe (the plasma potential or space potential V_s), there is no repulsion of either electrons or ions, but the higher mobility of the electrons due to their lower mass means that the current drawn by the probe is dominated by the electron current. The ratio of the current at this point (the electron saturation current) to the ion saturation current is given by $\sqrt{m_i/m_e}$, where m_i and m_e are the ion and electron masses respectively, which is simply the ratio of the velocities of the electrons and ions [90]. A probe voltage any higher than this will not see an appreciable increase in current, as the current is already limited by the rate at which electrons are collected by the probe.

Current is of course the rate of flow of charge, $I = qv$. By considering the relative velocities of the electrons and ions as they enter the sheath around the probe, we can derive [14] an expression for the current drawn as a function of voltage up to the plasma potential, for a Maxwellian plasma with an electron temperature T_e :

$$I_{probe} = I_{sat} \left(1 - \exp \left(\frac{e(V_{probe} - V_f)}{k_B T_e} \right) \right) \quad (4.1.1)$$

This model works well for probes in un-magnetised plasmas, but does not completely describe the IV characteristic of a probe in a magnetised plasma. The effective area of the probe depends on the magnetic field (through for example the Larmor radius and cross-field diffusion), and this in turn affects the current collected. In particular,

the ratio of the electron and ion saturation currents is greatly reduced compared to the un-magnetised case. Additionally, the expansion of the sheath's volume as the voltage is increased above the plasma potential also increases the collection area and so prevents the electron current from saturating. Both these effects can be seen on the right of Figure 4.2, which shows an actual IV characteristic from an inductively-coupled discharge (a plasma which is produced using radio frequency waves and confined magnetically [92]). Note too that the presence of the magnetic field in this case has actually resulted in a positive floating potential.

A complete discussion of the physical model of a Langmuir probe IV characteristic in a tokamak plasma is beyond the scope of this work. The important point to note is that a Langmuir probe diagnostic must be capable of producing such a curve, by biasing the probe at a range of voltages and measuring the current at each of those voltages.

4.1.2 MAST-U's divertor Langmuir probes

As previously mentioned, MAST-U has a novel and highly flexible geometry which includes an upper and lower Super-X divertor. The position of strike point for the outer divertor leg can be varied enormously [39]. In order to maintain good spatial resolution, whilst still providing coverage of all the regions of the divertor wall which could be conceivably be in contact with the plasma, a large number of probes are needed.

The MAST-U diagnostic consists of 4 arrays of probes, with two arrays each in the upper and lower divertors. In each divertor, the two arrays are toroidally separated by approximately 180° . The coverage of both upper and lower divertors enables measurements of up/down asymmetry in the plasma, most likely due to imperfect control of the plasma's vertical position. The two separate toroidal locations enable measurements of toroidal asymmetry.

In total, MAST-U's 4 divertor Langmuir probe arrays will have approximately 850 probes. This represents a large increase on the 576 probes installed on MAST [93], and substantially more than other tokamaks. Providing the electronics to sweep the voltage across such a large number of probes, and to digitise the current from all of the probes, is a significant challenge. If we were to simply connect up a separate power supply and a current and voltage monitor to each individual probe, the system would require an enormous amount of space and be extremely expensive. We instead need a solution which minimises the amount of electronics required to bias and measure the probes. The chosen solution is discussed in Section 4.2.

4.2 Hardware

Langmuir probes do not typically require very specialised electronics. Often, a waveform generator connected to an amplifier is used to provide the range of bias voltages required to produce an IV characteristic, the probe voltage is digitised with an ADC, and the probe current is measured using an ADC to measure the voltage drop across a shunt resistor. Representative examples are described in [94] and [95]. Each probe generally requires its own set of electronics to supply the bias voltage and to record this voltage and the corresponding probe current.

In order to address the challenge of measuring the huge number of probes in MAST and MAST-U, as described in Section 4.1.2, bespoke electronics was required. A modular solution was chosen for MAST, consisting of a power supply and amplifier to supply the bias voltage, a voltage and current monitor and a multiplexer. The use of the multiplexer means that one set of electronics was used for up to 16 probes, significantly reducing the amount of electronics required at the cost of reducing the rate of repeated sweeps for individual probes.

For the MAST-U probe diagnostic, the chosen solution is again modular, using a set of 40 modules where each module is responsible for 16 probes to enable operation of up to 640 probes in a shot. Similar to the original MAST system, each module consists of a power supply, a high-power (± 250 V, ± 4 A), high-bandwidth amplifier and a multiplexer. In addition to separate improvements to the design of the power supply and amplifier made by the CCFE electronics group, we asked: can we enhance the multiplexer by incorporating the probe multiplexing, measurement of the probe currents and voltages and overall control of the module all in a single compact unit using FPGA technology?

A schematic of the module, including the multiplexer unit designed to answer this question, is shown in Figure 4.3. A DAC on the multiplexer converts a digital bias voltage waveform into a low-voltage analogue signal in the range ± 2.5 V. This is sent to the amplifier, which amplifies the signal up to ± 250 V, and the high voltage signal is returned to the multiplexer. The voltage and current are digitised by 16-bit ADCs, and the amplifier output is directed to a set of 16 switches (of which only 6 are shown in Figure 4.3 for simplicity), with the output of each switch connected to an individual probe. The multiplexer controls which switch is turned on at which point, enabling the voltage sweep to be applied to each individual probe in turn (the sequence in which probes are turned on is completely configurable).

Additional 12-bit ADCs on the probe side of the switch measure the probe voltage for each individual probe. When the switch is closed this should equal the voltage

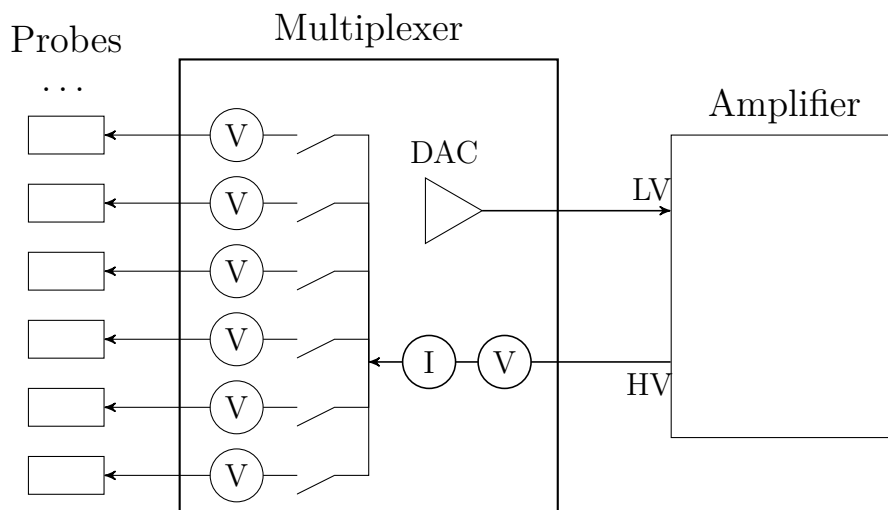


Figure 4.3: Schematic of a single MAST-U Langmuir probe electronics module. The multiplexer generates a low-voltage version of the probe bias waveform using a DAC, and this is amplified by the high-voltage amplifier. The amplified signal returns to the multiplexer, where the current (measured using the voltage drop across a $0.2\ \Omega$ shunt resistor) and voltage are digitised by ADCs, and the amplifier output is connected to one of 16 probes. Additional ADCs on the probe side of the switches are used to measure the voltage when the probe is isolated from the amplifier: in this case the floating potential is measured. The power supply is not shown, but supplies power to both the amplifier and multiplexer. All voltage measurements are with respect to vessel ground, though we have omitted this in the figure to simplify the layout.

measured by the amplifier, minus any very small voltage drop due to the small resistance of the switch itself. When the switch is open so that the probe is disconnected from the amplifier, this ADC measures the floating potential of the probe. This means that the module can provide 12-bit floating potential measurements for all 16 probes simultaneously, in addition to 16-bit sweep measurements for one probe at a time.

The DAC and all ADCs sample at 1 MSPS, and the amplifier has a bandwidth of approximately 60 kHz. The high bandwidth is important, since the multiplexing of up to 16 probes means we need to sweep each probe quickly in order to obtain acceptable time resolution for each individual probe. When sweeping all 16 probes with a 65 μ s sweep time, as used by the previous MAST electronics, each individual probe is swept once every 1.04 ms, which sets the time resolution for that probe. The simultaneous floating potential measurements at 1 MSPS enable us to study high bandwidth fluctuations, such as turbulent fluctuations, which could not be resolved by sweeping the probes; we shall see examples of this in Chapter 5.

4.2.1 The multiplexer

The multiplexer functions as the controller for each Langmuir Probe module. As well as digitising the current and voltages, producing the analogue bias voltage waveform and synchronising the multiplexing of the probes with the bias waveform, it is required to interface with the MAST-U data acquisition and central timing infrastructure and to handle configuration of the module. To provide the capability to perform all of these tasks whilst keeping the system low-cost and compact, the multiplexer uses FPGA technology.

Like the bolometer system described in Chapter 2, the multiplexer features a Xilinx Zynq chip: a powerful combination of FPGA programmable logic and a dual-core ARM CPU. An off-the-shelf Avnet MicroZed board [96] mounted on an FMC carrier card is used, which contains a Zynq 7Z010 FPGA, Gigabit Ethernet, multiple Pmod™ connectors and an LPC FMC connector to provide sufficient I/O ports to connect the various hardware components of the multiplexer to the FPGA. The Zynq chip used in the multiplexer features fewer FPGA logic resources than that used in the bolometer, because we have no need to perform complex signal processing on the multiplexer.

A photo of the multiplexer is shown in Figure 4.4. The many components of the multiplexer, including power supplies, switches, DAC, ADCs and FPGA board fit neatly into a 1U, 19-inch unit. This is paired with a crate containing the power supply and two amplifiers (so there are two multiplexers per crate) to create a




Figure 4.4: A photo of the multiplexer unit. The red boards front right are the MicroZed and FMC carrier, and the 16 switch boards can be seen at the back in two rows of 8. The rear left board consists of power supplies to the various components of the multiplexer, and the DAC and 16-bit ADCs are on the front left board. The multiplexer fits in a 1U, 19-inch rack; a crate containing a power supply, two amplifiers and two multiplexers fits in a 7U, 19-inch rack and provides 32 channels.

compact, modular unit. By moving much of the complex control logic onto the FPGA, the cost and footprint of the system are greatly reduced.

4.3 FPGA design

In common with several other FPGA-based diagnostics at MAST-U, including the initial version of the bolometer design, the Langmuir probe multiplexer's FPGA firmware uses the Xillybus IP core to transfer data between the FPGA programmable logic and ARM processor. The deterministic timing offered by the FPGA means that we implement all the functionality which needs precise synchronisation in the programmable logic. This includes the SPI interfaces with the DAC and all 18 ADCs (two 16-bit ADCs to measure the current and voltage output by the amplifier and 16 12-bit ADCs to measure the voltage on the output of the switches), with the sampling synchronised to an external master clock to ensure a common time base with other plant systems.

The voltage waveform sent to the DAC is stored in a RAM on the FPGA, with successive values being retrieved every $1\ \mu\text{s}$. This fixed readout rate was chosen to keep the FPGA design as simple as possible. The RAM can contain up to 65 536 entries, so the maximum length of the bias voltage waveform is 65.536 ms; this is limited by the available FPGA memory and is sufficient for MAST-U's needs. The length of the waveform is configurable, meaning the waveform shape, voltage range and period can be independently varied to produce the desired waveform. This provides good flexibility, and enables the system to use waveforms most suitable to the expected plasma conditions in a wide variety of shots.

Similarly, the probe switching sequence is stored in a smaller RAM on the FPGA. This RAM can contain up to 64 entries, which is expected to be more than sufficient for a system with up to 16 probes. The sequence consists of a series of 16-bit bit-fields, where each bit corresponds to one probe. A value of 1 means that the switch for that probe is closed and so the probe is connected to the amplifier. A value of 0 means that the switch is open, so the probe is disconnected from the amplifier. In principle, more than one probe can be connected to the amplifier simultaneously by setting more than one bit to a value of 1, but the difficulty in interpreting the resulting data means this feature is unlikely to ever be used.

The switching of probes is synchronised with the waveform generation. The FPGA first reads the value in the switch sequence from the sequence RAM, and sets the switches appropriately. It then reads the first value from the voltage waveform RAM and sends it to the DAC. Every $1\ \mu\text{s}$, it reads the next value from the waveform

RAM and sends it to the DAC, until it has reached the number of samples equal to the waveform length which has been configured. It then reads the next value of the switch sequence, sets the switches and returns to the start of the voltage waveform to repeat the process. This means every probe in the sequence has the same voltage waveform applied to it, and the same voltage waveform is used throughout each shot. The waveform and switch sequence can of course be different across different multiplexers, and both can be changed on a shot by shot basis.

Although the FPGA firmware is capable of working with arbitrary waveforms, there are some restrictions imposed by the hardware. To avoid damage to the electronics, the voltage should be at 0 V when the switches are changed, else sudden changes in resistance can cause large current spikes which can destroy the amplifier. The waveforms should therefore be designed to include several μs at the beginning and end at 0 V. In addition, there is a time delay in the signal passing through the amplifier, measured empirically to be about $7 \mu\text{s}$, so the FPGA also includes a configurable delay between outputting the voltage on the DAC and changing the values of the switches to ensure that the end of the waveform and the change of switch value coincide to the nearest μs .

The FPGA also records the 6 kHz clock used for the MAST-U power supplies. These are switched-mode power supplies, and each time they switch a large amount of noise is induced in any nearby electronics, including the probes and their cabling. This means that any IV characteristics measured during the power supply switching will be swamped with noise and rendered unusable, so this signal provides a simple means of masking off these IV characteristics during post processing. The power supply clock is sampled at the same 1 MSPS as the voltage, current and switch traces to provide a common time base to ease post-processing of this data.

Each of these functions is implemented by a module in the FPGA design. Like the bolometer's Xillybus-based FPGA design, the FPGA also implements the FIFO interface required to stream the ADC measurements to the processor's RAM. In addition to these measurements, and the MAST-U power supply clock, the value of the switch sequence is recorded with the same time base as the ADC measurements, in order to accurately determine which probe is being swept for any given ADC sample. A control interface and clock module closely modelled on the bolometer Xillybus design are also included. Since their functionality is identical to that described in Section 2.3 (with the exception that only a single 120 MHz clock is output by the clock module for use with the ADCs and DACs), they are not described here.

Figure 4.5 illustrates the interaction between the modules in the FPGA design, many of which have the same functionality as those described in Section 2.3. The only changes are the SPI interfaces to the DAC and ADCs, and the module responsible

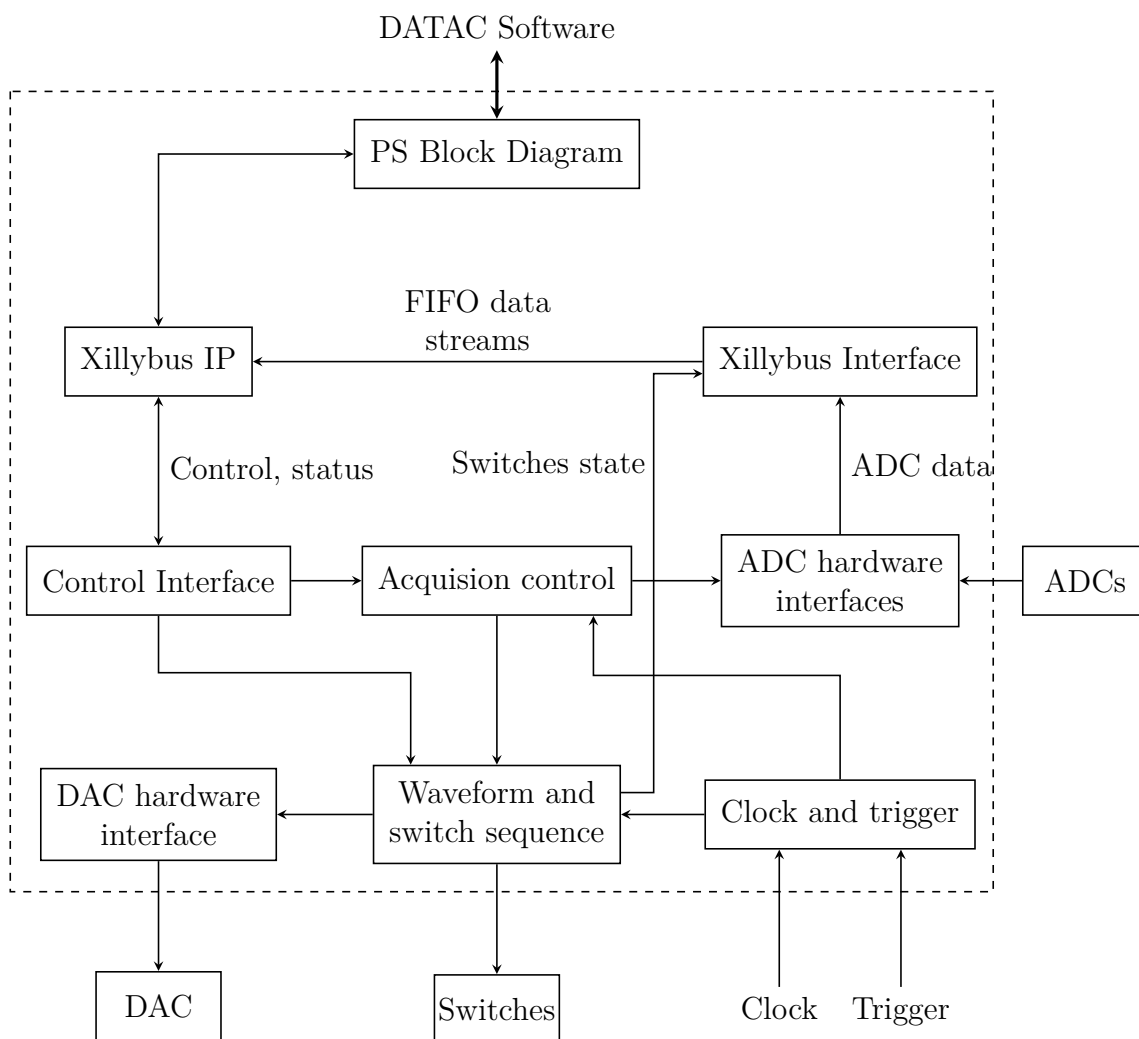


Figure 4.5: The modular design of the FPGA firmware for the multiplexer, showing the logical flow of data between modules. Like Figure 2.8, many of the two way connections between modules are omitted for clarity. Again, all logic on the FPGA chip is inside the dashed box, and the external interfaces to the DAC and ADCs, the multiplexer switches and the DATAC and timing networks are shown.

for sending the waveform to the DAC and synchronising the switch sequence with this waveform.

4.3.1 Web server with CherryPy

We exploit the convenience and simplicity of programming a CPU to implement most of the control functionality of the multiplexer using software running on an Ubuntu-based Xillinux Linux distribution on the MicroZed. The operating system is fundamentally the same as that used in the original design of the bolometer system, described in Section 2.3, but does not require the modifications which needed to be made to run on D-TACQ's ACQ2106 hardware, since the MicroZed is supported out of the box by Xillinux.

The interface with MAST-U's central data acquisition infrastructure is handled using HTTP, in a similar manner to the Jetblack interface implemented for installation of the bolometer on JET and described in Section 2.3.2. No existing software was available for the MAST-U interface, meaning the software needed to be written from scratch. The experience of working with Jetblack led to the adoption of CherryPy [97], a Python web framework, since Jetblack itself is based on this framework.

This design means the system functions as an HTTP server, which receives requests from the MAST-U data acquisition (DATAC) software and runs different functions to service each request. At the beginning of each shot, a `PREPARE` request is sent, which causes the system to cease any data collection still running from the previous shot and return to an idle state in preparation for the upcoming shot. This is followed by an `ARM` request. The software ensures the DAC is set to output 0 V, then runs a short "setup shot" which energises all of the switches by turning them off and on; this ensures they properly isolate the probes from the amplifier when opened. Since this must be done within about 30 s of the MAST-U shot starting, the system waits for a configurable amount of time after receiving the `ARM` request, to account for the time delay between receiving this request and the trigger to start the measurement (this time delay is around 2 min in MAST-U, but is often set to 0 s for lab testing to speed up repeated tests). After the switches are energised, the DAC is reset to 0 V again and the amplifier is enabled. The system is now fully armed, and waits for the hardware trigger from the MAST-U central timing system which signifies the start of the data collection and probe biasing window.

Data is read into the DDR RAM on the MicroZed, rather than being stored on the SD card containing the operating system files. This has two advantages. Firstly, the large volumes of data acquired (approximately 40 MB/s for shots lasting up to 5 s)

would require a large number of writes to the SD card, rapidly using up the finite number of writes available in each block in the card's flash memory and making premature card failure more likely. By instead writing the data to DDR, we avoid wearing out the SD card in this way, which improves the reliability of the system in the long term. Secondly, the rate at which data can be transferred to and from the DDR is far faster than the rate at which the data can be written or read from the SD card, which we have found is limited to about 7 MB/s on the Zynq. This not only allows us to acquire at the maximum sample rate, but it also speeds up read back of the data after the shot.

Once the requested number of samples is collected, the multiplexer's web server notifies the DATAC software and read back commences. The DATAC software makes a series of `READBACK` requests over the HTTP connection with unique query strings for each individual channel, for example the measured current and voltages, switch values and power supply signal. The multiplexer software sends the data as the body of the HTTP response. Since the transaction is conducted using Gigabit Ethernet, the read back time is relatively short (transfer rates of about 60 MB/s are possible with the Zynq). Once read back is completed, the system waits for the `PREPARE` request which signifies the start of the next shot. A shot can be aborted at any time upon receipt of an `ABORT` request, which puts the system into the same state as the `PREPARE` request.

In order to keep the FPGA design as simple as possible, functionality such as ensuring the DAC output was reset to 0 V after each shot and cycling all switches in order to energise them before each shot was implemented in software. In the case of zeroing the DAC output a dummy shot is run, with the DAC waveform set to 0 V. To energise the switches the same method of running a dummy shot is used, with the switch sequence set to cycle through each probe individually for a few complete cycles. Encapsulating these procedures into self-contained functions modularises the software and makes it easy to understand what the software is doing for a reader of the code. The software implementation is also simpler and more maintainable than an implementation in FPGA programmable logic, and fully exploits the power and flexibility of the Zynq FPGA/CPU technology.

4.4 Improvements over previous system

The new Langmuir probe electronics features multiple improvements over the previous MAST system. The use of FPGA technology allows each module to occupy a significantly smaller footprint than the analogue system it has replaced. By moving

much of the complex logic from an analogue circuit to a digital implementation on the FPGA the system also becomes simpler and more maintainable, since when adding functionality or fixing bugs it is much easier to simply re-program the FPGA instead of producing a new analogue circuit. By running an embedded Linux operating system on the Zynq chip in the multiplexer, there is no need for a separate PC and each module therefore becomes fully self-contained. This again reduces the hardware requirements, which is particularly important for a diagnostic consisting of 40 separate modules to operate a total of 640 probes. Implementing diagnostic control and data read back using the industry standard HTTP protocol further eases future development and maintenance of the system.

The digital control of the multiplexer switches enables greater operational flexibility. The order in which the probes are swept is completely arbitrary, and it is possible to sweep the same probe multiple times in a sequence if improved time resolution is required for some probes but not others. The switch values are recorded digitally, which completely eliminates any issues with noise pickup in the signal which records this value and makes post-processing easier (previously an analogue voltage which recorded the switch value was digitised by an ADC and the levels were discretised in software).

The ability to make simultaneous measurements of the floating potential on all unswept probes is new to the diagnostic. This will enable us to provide high bandwidth spatially resolved measurements of the floating potential, which will facilitate studies of turbulent fluctuations in the plasma edge region. By observing the size of floating potential fluctuations in nearby probes, we can also provide insight into the suitability of IV characteristics for a given probe. For example, we can explicitly discard a sweep when large floating potential fluctuations are measured nearby which show that the local plasma conditions are not sufficiently stationary during an individual sweep.

The self-contained nature of the system, with its industry standard HTTP interface, make installing the probe electronics at a variety of institutions other than MAST-U very practical. The high degree of configurability, particularly with the arbitrary waveform and probe switching capabilities, also means that there is significant scope for testing and optimisation of the operation of the system. We shall illustrate these points in Chapter 5, where we describe the installation of the probe electronics on two plasma experiments in order to test the wide range of functionality provided by the system and begin to optimise the waveform and processing of the newly available data.

Chapter 5

Results from the Langmuir Probe system

In this chapter, we describe the installation of the new Langmuir Probe system at two sites. A brief introduction to the set-up at each site is given, and the results of measurements made using the diagnostic are discussed. The primary purpose of these tests was to demonstrate the correct operation of the system and its fitness for purpose in relevant plasma conditions, not to conduct in-depth physics investigations or detailed characterisation of the probe electronics in the same depth as the bolometer studies in Chapter 3. The tests also provided an indication of the expected signal quality and useful information for optimising the operation of the device on MAST-U, in particular the sweep waveforms.

5.1 Installation on the York Linear Plasma Device

Much of the development of the Langmuir probe multiplexer and amplifier units was done using resistors as the load through which currents was driven. Although this is fine for verifying very basic functionality of the electronics, such as whether the DAC and ADCs are operating as expected, the system also needed to be tested with a more realistic load: a plasma. In particular, we needed to know whether the system could produce IV characteristics sufficiently free of distortions and with noise levels low enough to be used for calculating plasma parameters of interest.

A prototype amplifier and multiplexer were therefore installed onto the York Linear Plasma Device (YLPD) at the York Plasma Institute [98], in collaboration with Hannah Willett and Kieran Gibson. The YLPD produces a linear column of plasma




Figure 5.1: Diagram illustrating the probe set-up on the York Linear Plasma Device. The plasma, in purple, is confined by an axial magnetic field \mathbf{B} produced by the red coils. The grey probes are cylindrical in shape and oriented parallel to the magnetic field. Only 6 of the 12 probes used are shown, for clarity.

of diameter 2 cm to 3 cm in a relatively low magnetic field, compared to that of MAST-U, of around 100 mT. The ion density is also somewhat lower than that expected in MAST-U, at 10^{16} m^{-3} to 10^{18} m^{-3} , but the electron temperature is similar at up to 15 eV. The device features good diagnostic access, and being in a university lab provides good opportunities to install and test equipment very quickly without having to implement autonomous software control or integration into complex laboratory systems. In addition, research underway at York examining instabilities at the onset of detachment required an array of Langmuir probes to look at spatial and temporal evolution of plasma fluctuations, but York did not have the electronics necessary for operating this probe array. These factors made the YLPD a good choice for some first-plasma tests of the MAST-U Langmuir probe electronics.

The electronics was connected to an array of 12 cylindrical tungsten probes, each of which had a surface area of 2.5 mm^2 and was aligned with its circular face normal to the magnetic field. The set-up is shown in Figure 5.1. The axial magnetic field is produced by a set of coils, of which two are shown in the figure. The plasma is produced by a cathode at the left hand side of the device as shown in the figure and flows towards the anode at the right hand side, downstream of the probes. The probes were arranged at 2 mm intervals, and there were typically 4–6 probes in the highest density central part of the plasma column with a few more in the lower density outer region of the plasma column and a few outside of the plasma column altogether.

This thesis is concerned only with the suitability of the data produced in these tests. A detailed physics analysis of this data is left to our collaborators at York. We therefore present only a small subset of the data gathered, which is sufficient to

illustrate the features of interest from a diagnostic viewpoint.

5.1.1 Floating potential measurements

High speed simultaneous measurements of the floating potential are a capability of the multiplexer electronics which was not present in the previous MAST system. It was important therefore to evaluate how useful this data could be. The floating potential measurements are most useful for measuring high frequency fluctuations, since these cannot be resolved by sweeping the bias voltage and obtaining IV characteristics.

We looked for changes in the type of fluctuations seen in the floating potential measurements as the plasma transitioned from an attached to a detached state. This was accomplished by running three separate acquisitions during a Hydrogen discharge which was seeded with additional Hydrogen gas to induce detachment. The results can be seen in Figure 5.2, which shows the comparison of floating potential measurements for different amounts of seeded gas.

The top plot shows the floating potential with no seeding, when the plasma was in a fully attached state. Here there are generally small fluctuations on very short time scales, which look like vertical stripes between 10 mm and 20 mm, and all fluctuations are reasonably uniform. Once the gas seeding is increased, in the middle plot, we see much stronger fluctuations, and at lower frequencies. In this shot the plasma was fully detached, with strong radiative emission near the anode. The fluctuations look more like dark purple blobs separated in time by light blue regions of smaller V_f , and are particularly noticeable around 16 mm. Finally, the bottom plot shows a partially detached plasma, achieved by reducing the gas seeding compared with the middle figure. Here we see hints of some larger fluctuations, just visible as purple blobs around the centre of the plasma column at around 16mm, which also occur at frequencies lower than the fluctuations seen in the fully attached state. However, it is now more difficult to distinguish individual “blobs” as the light blue regions of smaller V_f between fluctuations can no longer be made out.

A complete analysis of the causes and behaviour of these fluctuations is out of the scope of our discussion. However, the fact that we can identify changes in the fluctuation behaviour under different detachment conditions with simultaneous floating potential measurements indicates that this sort of measurement will be useful in future analysis of detachment onset dynamics.

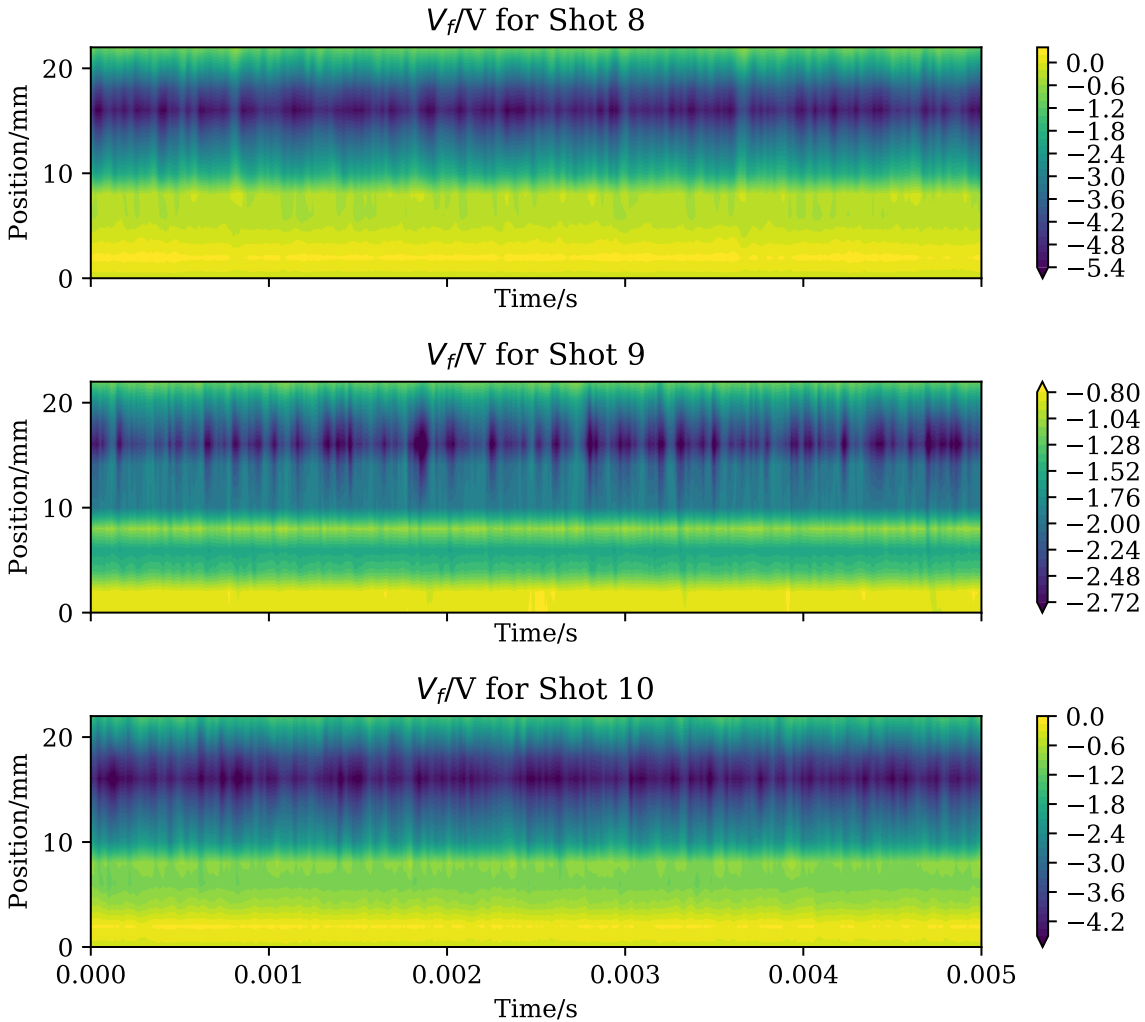


Figure 5.2: Floating potential measurements in the YLPD with different amounts of Hydrogen gas seeding, measured with 12 probes spaced at 2 mm intervals. Position on the y axis is taken as the vertical position of the centre of each probe tip's circular face, given relative to the first probe in the array (not relative to the vessel wall or the plasma column). With this coordinate system, the plasma column centre is at approximately 16 mm. Shot 8 has no seeding. Shot 9 has enough seeding to fully detach the plasma. Shot 10 has less seeding than Shot 9, resulting in a partially detached plasma. The change in the size and frequency of fluctuations can be seen as the amount of gas seeding is increased from Shot 8 until detachment is achieved, and reduces as the gas seeding is reduced again afterwards.

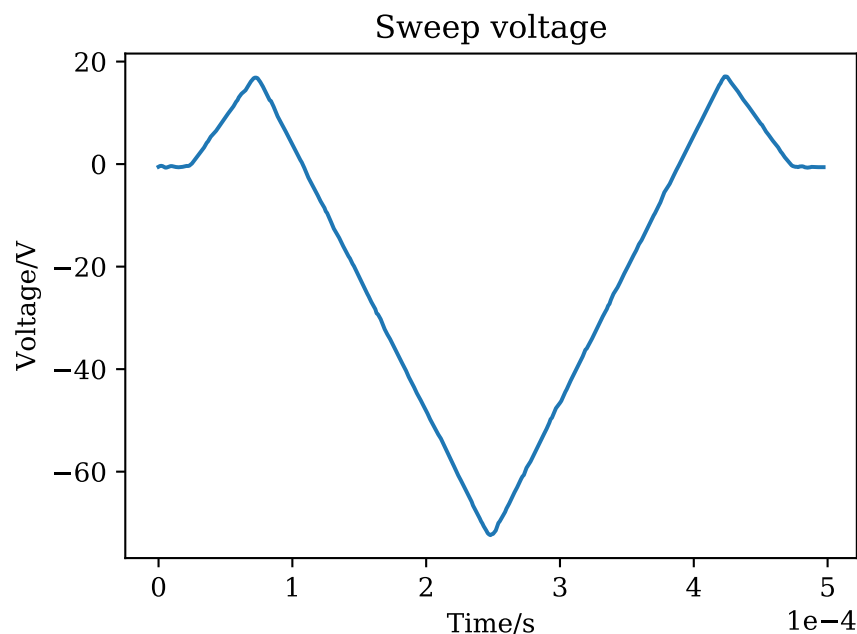


Figure 5.3: The voltage waveform used to sweep the probes and obtain IV characteristics, as output by the amplifier and measured by the multiplexer. Periods of 0 V at the beginning and end are used to change the probe switches safely, and the waveform is symmetric about the centre to enable averaging the upward and downward parts of the sweep to remove parasitic current.

5.1.2 Sweep measurements

Whilst floating potential measurements are simple to make and can be recorded at high bandwidth, the floating potential on its own contains only limited information about the plasma. In particular, it is not possible to measure the electron temperature or density from floating potential measurements alone: for these a full IV characteristic is needed. Furthermore, multiplexing through up to 16 probes requires a very high sweep rate and places significant demands on both the multiplexer and amplifier electronics, so it is important to test whether we can still produce useful IV characteristics for multiplexed probes.

As mentioned in Section 5.1, the YLPD operates with plasma densities somewhat lower than those expected in MAST-U. Since the current drawn by the probe is proportional to the electron density [29], the current we measure in the YLPD will be much lower than the currents in MAST-U that the system was designed for. To account for this, the shunt resistor used to measure the current in the multiplexer was temporarily changed from $0.2\ \Omega$ to $2\ \Omega$, which increased the measured current by a factor of 10, but restricted the range of the current measurement to $\pm 0.5\ \text{A}$.

The voltage waveform used to sweep the probes in these tests was based on the




Figure 5.4: Left: the IV characteristic in vacuum for probe number 4, averaged over 3 groups of 50 sweeps, showing the induced current as a function of bias voltage. There is a significant variation between the 3 groups of sweeps in the induced current at high negative voltages. The current is much more consistent between the groups of sweeps for voltages above about -20 V. The current in all 3 groups of sweeps is also highly asymmetric. Right: the voltage gradient dV/dt , for the same 3 groups of sweeps. Unlike the current, this is highly symmetric and consistent across all 3 groups of sweeps.

waveform typically used with the original MAST Langmuir probe electronics, but with the maximum and minimum voltages reduced to avoid sheath effects in the lower densities of the YLPD. It is shown in Figure 5.3. The waveform has a period of $500 \mu\text{s}$, slower than the $65 \mu\text{s}$ typically used on MAST, which was chosen to minimise distortion in the current measurements seen with rapid sweeps using the prototype amplifier during testing. The symmetry of the waveform enables averaging the current in the upward and downward parts of the sweep, to cancel out the parasitic current induced by the capacitance and the changing voltage in the system. This of course only works if the parasitic current is of the form $I_{par} = CdV/dt$ and the system has constant capacitance C .

In the past, analysis of the MAST Langmuir probes has assumed all parasitic current is induced by capacitance. To test the validity of this assumption with the new electronics we ran a calibration shot, consisting of 160 sweeps per probe, with no gas in the vacuum chamber of the YLPD. Figure 5.4 shows the results for one particular probe, though all 12 probes showed similar behaviour. Clearly the parasitic current is not due solely to capacitance in the system, since it is highly asymmetric whereas the voltage gradient dV/dt is symmetric. Furthermore, the induced current at higher negative voltages seems to become more negative as we move through the calibration shot, as shown by the 3 successive IV curves each averaged over 50 sweeps. This is not due to any gradual time evolution of the voltage gradient, and its source is

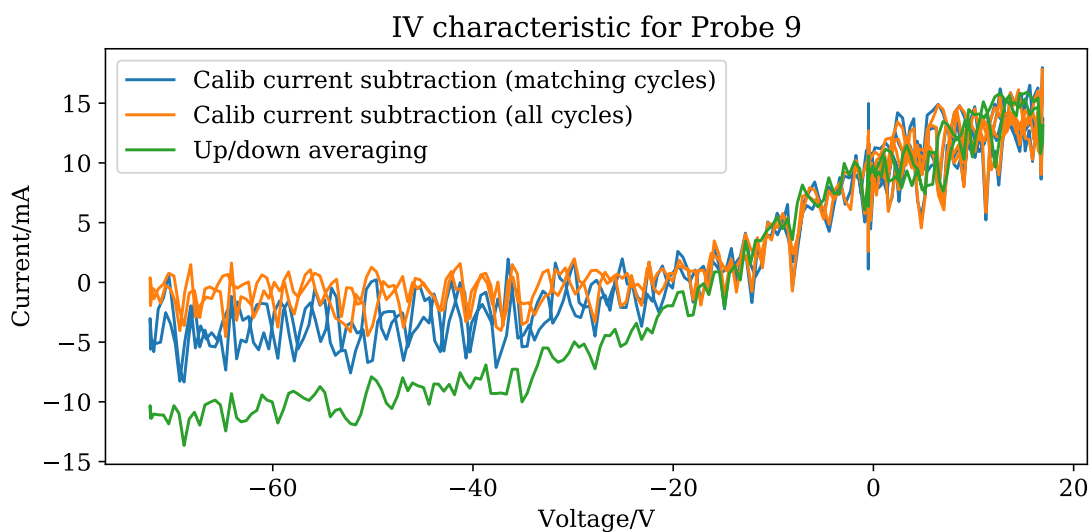


Figure 5.5: An IV characteristic for a probe near the centre of the plasma column in a high power, high field, high gas discharge averaged over 10 sweeps. The different curves use different methods to remove the parasitic current: subtraction of the current from the calibration sweeps (either averaged over the same 10 sweeps or all sweeps from the calibration) or averaging the upward and downward parts of the sweeps. Whilst there is little difference above -20 V, the methods give significantly different results below -20 V.

currently unknown. The 160 sweeps were run over 900 ms and the current at most negative voltage in each group shifted by around 10 mA, making this a potential source of systematic error on the saturation current measurements made during plasma shots. In a machine like MAST-U, where the saturation current could be tens or hundreds of mA, this error is not so bad, but it has a significant effect in the low density YLPD.

In any case, it is clear that simply averaging the current from the first and second halves of the symmetric voltage sweep is not sufficient to remove induced currents. The consistency of the induced current for $V > -20$ V means a better method is to simply subtract the average current of the calibration sweeps from the sweeps done in plasma. The evolution of the current at highly negative voltages will, as previously mentioned, add an uncertainty to measurements of $V < -20$ V. If the current evolution is consistent between shots we may be able to reduce this error by averaging calibration sweeps which were done at the same time relative to the start of the shot as the sweeps of interest in the current shot. Once the system is installed on MAST-U, we can evaluate how much the current changes over the course of the device's 5 s maximum shot length with the final production amplifier, probe heads and cabling to see whether this is still an issue.

With this information about the source of parasitic current and the methods of removing it, we can begin to produce IV curves and infer useful plasma parameters. Figure 5.5 shows an IV curve for a high fuelling, high field and high power discharge. The curve has been averaged over 10 sweeps in order to eliminate intermittent spikes in the current measurements which were present in the prototype amplifier. Three different methods have been used to remove the parasitic current.

The first two methods involve subtracting an average IV curve from the calibration sweep from the data. As previously mentioned, Figure 5.4 shows some time evolution in the calibration IV curve. The blue curve in Figure 5.5 accounts for this by assuming this time evolution is repeatable, and therefore subtracting the calibration current from the cycles at the same time in the calibration data as in the measurement data (in this case sweeps 40–50 were averaged in both shot and calibration data). The orange curve relaxes this assumption, and simply subtracts from the measurement data the current averaged over all sweeps during the calibration run.

Finally, the green curve does not use the calibration data at all, but assumes that all parasitic current is induced by capacitance and can therefore be removed by averaging the parts of the sweep with equal and opposite dV/dt . This is included because it is the method used with the original MAST Langmuir probe data and was the justification for using the symmetric voltage waveform shown in Figure 5.3. It estimates the saturation current to be significantly more negative than either of the two other methods. This is not surprising, since we can see from Figure 5.4 that the average induced current for $V < -20$ V is significantly less than 0 mA at all time ranges throughout the calibration, and therefore the current is not simply due to capacitance in the system. Since the data in Figure 5.5 is taken for sweeps 40–50, we expect that the up/down averaging method would be biased by around 7 mA at the most negative voltages, which matches the difference between this method and the matching-cycle subtraction method.

We can also determine which of the current-subtraction methods is the most accurate by exploiting our knowledge of the YLPD plasma. Unlike a tokamak, the plasmas produced in the linear device are typically in a steady state over long periods, apart from small fluctuations. It follows therefore that the IV curves should be similar at all times throughout our measurement period.

In Figure 5.6 we have plotted IV curves averaged over 10 sweeps at 3 different time ranges in the 600 ms data acquisition window. We can clearly see that when we subtract currents averaged from sweeps in the same time ranges in the calibration data the curves are almost equal across all 3 time slices, but if we subtract the current averaged over all of the calibration cycles the ion saturation current appears to evolve throughout the shot. Since we know that the ion saturation current should




Figure 5.6: Comparison of current-subtraction methods for the averages of 3 sets of 10 sweeps, from the same discharge as Figure 5.5, at different times throughout the measurement window. The left figure uses the matching cycles method, whereas the right figure subtracts the same current, averaged over all cycles for the same probe in the calibration, from each measured IV curve. There is a noticeable trend in the ion saturation current throughout the shot in the right figure, whereas in the left figure it is almost the same for each curve.

not change over the course of this YLPD discharge, we can say that the most accurate method for parasitic current removal is the matching-cycles subtraction method, at least for the prototype electronics used at the YLPD.

5.1.3 Plasma parameters from sweeps

In Section 5.1.2 we discussed the effects which need to be considered to produce representative IV characteristics using the multiplexed sweep data from the new electronics. Now that we have a representative IV curve, we can test the ability to extract useful plasma parameters from it, namely the ion saturation current I_{sat} , electron temperature T_e and floating potential V_f . We use a simple 3-parameter model based on the theory of an electrostatic sheath at the probe tip, where the current is of the form [14]:

$$I = I_{sat} \left(1 - \exp \left(\frac{e(V - V_f)}{T_e} \right) \right) \quad (5.1.1)$$

Other more sophisticated models do exist, such as the first-derivative probe technique [99] and a 4-parameter fit which accounts for sheath expansion [100], but for the purpose of proving the diagnostic's fitness-for-purpose the 3-parameter model is sufficient.

Figure 5.7 shows the result of a nonlinear least squares fit to this model, using the

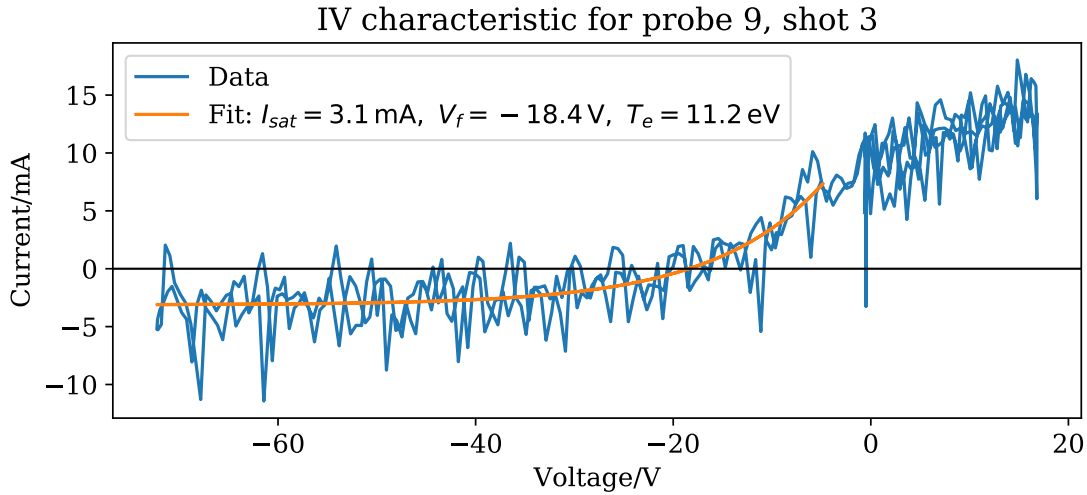


Figure 5.7: 3-parameter fit of the IV characteristic for a probe in the middle of the plasma column in a high power discharge, averaged over 10 sweeps between 183 ms and 237 ms. The fit is performed on a subset of the curve, to avoid the electron saturation region which is not considered by the 3-parameter model, and appears to be in good agreement with the data.

average of the sweeps between 183 ms and 237 ms in a high power, high field and high density discharge. In order to avoid the fit being biased by the electron saturation region, which is not considered by the model, we terminate the fit at a lower voltage than the maximum of the sweep where we judge the electron saturation region to begin. The voltage was chosen by eye to use as much of the data as possible whilst avoiding the start of the region where the current begins to flatten off compared with an exponential curve, as this indicates the start of the electron saturation region. We can see that the fit appears to follow the general shape of the IV characteristic well, and the calculated T_e is consistent with what the YLPD typically achieves. The floating potential also agrees well with the measurement made by the simultaneous V_f ADCs when this particular probe was not being swept. The fit was however very sensitive to the choice of voltage cut-off. The variation in the fit parameters with this voltage cut-off dominates the uncertainties on the parameters, which we have estimated as follows: $I_{sat} = (3 \pm 1)$ mA, $V_f = (-18 \pm 4)$ V, $T_e = (11 \pm 5)$ eV.

The ions enter the sheath around the probe at the sound speed, which is given by the Bohm velocity $v_B = \sqrt{(T_i + T_e)/m_i}$, where T_i and m_i are the ion temperature and mass respectively [14]. Assuming the electrons and ions have equal temperature, we can therefore calculate the electron density by relating it to the ion saturation current:

$$\begin{aligned} I_{sat} &= A_{probe} Z_i n_e v_B \\ &= A_{probe} Z_i n_e \sqrt{2T_e/m_i} \end{aligned} \quad (5.1.2)$$

The area of the probe surface projected onto the magnetic field A_{probe} is 2.5 mm^2 , which is simply equal to the circular cross section of the probe since the probes are aligned parallel to the magnetic field. The ion atomic number Z_i is 1 for the Hydrogen plasmas we used, and the ion mass is the mass of Hydrogen. Solving Equation 5.1.1 for n_e gives an electron density of $n_e = (3.3 \pm 1.3) \times 10^{17} \text{ m}^{-3}$. Like the T_e value, this is consistent with typical YLPD plasma parameters.

5.1.4 Summary of YLPD tests

The tests of the power supply, multiplexer and amplifier electronics on the YLPD successfully demonstrated the fitness-for-purpose of the system. The high speed, simultaneous floating potential measurements were shown to have sufficient resolution to offer insight into the behaviour of plasma fluctuations at the onset of detachment. By reducing the measurable current range to deal with the low ion saturation currents on the YLPD, we were also able to obtain IV characteristics from which we could calculate reasonable values for the key plasma parameters typically measured by Langmuir probes: the electron density and temperature.

During the course of producing these IV characteristics, we identified that the parasitic capacitance is not symmetric when using a symmetric sweep voltage waveform, and is therefore not simply due to the capacitance in the system. This allows us to relax the constraint of a symmetric voltage waveform to ease parasitic current subtraction, meaning we can design a more optimal waveform for use in future experiments. We found a better method of parasitic current removal was to subtract an averaged current measured with a probe in vacuum for the same voltage sweep, but there was some variation in this “calibration current” over the course of the measurement window, which may need to be accounted for. This knowledge will inform development of the analysis codes for the Langmuir probe data on MAST-U.

Overall, both the multiplexer and power supply behaved as expected, and were considered production ready. Excessive noise spikes in the current produced by the amplifier were identified, and this enabled the electronics team at CCFE to investigate and improve the amplifier performance as a result.

5.2 Installation on the COMPASS tokamak

The tests on the YLPD provided useful experience of plasma measurements with the new diagnostic electronics, and highlighted areas of development both in analysis of the data and in the design of the electronics itself. However, there are challenges

and conditions present in operations on a tokamak like MAST-U that do not exist with the YLPD. For example, the current in MAST-U is expected to be much larger than that measured in the YLPD, and tokamak plasmas feature large transient events such as ELMs which the linear device does not allow us to simulate. On an operational level, a complex experiment such as a tokamak requires sophisticated, automatic software control of the device for set-up and data acquisition, whereas with the university lab-based YLPD it was sufficient to manually set up and run data acquisitions.

To address these limitations with the York tests, the system was installed on the COMPASS tokamak at IPP CAS in Prague, Czech Republic [101]. COMPASS is a compact machine with major and minor radii $R = 0.56$ m and $a = 0.2$ m respectively, and typically operates with ELMy H-mode plasmas in a lower single null divertor configuration. The tokamak was initially owned and operated by UKAEA at Culham Science Centre before being mothballed in 2002 and then transferred to IPP CAS, where it achieved first plasma in 2008.

This transfer means CCFE has good links with IPP CAS, which makes it easier to set up collaborations between the two institutions. The COMPASS team are also strongly interested in diagnostic development, and it is possible to bring a new diagnostic and install it on the machine quickly and with minimal administrative overhead. These factors, along with COMPASS's ability to produce plasma conditions relevant to MAST-U, led to the collaboration with IPP CAS to test the new probe electronics on COMPASS.

The work in this section was performed in collaboration with Matej Peterka, with contributions from A Havránek, M Dimitrova and J Adánek. Parts of the work have been published in the JINST journal, in the proceedings of the 2nd European Conference on Plasma Diagnostics (ECPD 2017) in Bordeaux [102].

5.2.1 Experimental setup

The MAST-U electronics were over the course of the tests connected to probes belonging to both of the existing divertor Langmuir probe arrays at COMPASS. The first array consists of 39 dome-shaped graphite probes, spaced approximately 5 mm apart, which cover both the inner and outer strike points. The probe tips have an effective area of between 7 mm^2 and 8 mm^2 , depending on the magnetic field geometry. This array is described in more detail in [103, 95]. The second, newer array features two toroidally separated rows of 53 rooftop-shaped probes in conjunction with an array of ball-pen probes and is described in [89]. The rooftop-shaped probes




Figure 5.8: The arrays of Langmuir probes used for the COMPASS tests. Left: the first array, consisting of 39 dome-shaped probes, reproduced from [104]. Right: the second array, consisting of both rooftop-shaped probes and ball-pen probes in two toroidally separated arrays, reproduced from [89]. The red lines show the expected strike point positions, and the magnetic field angle is also shown on both the high and low field sides.

have an area of 22 mm^2 and are spaced approximately 3.5 mm apart. Figure 5.8 shows the arrangements of the two arrays.

For our tests, we used mainly a selection of 10–12 probes from the first array, located just outside the radius of outer strike point moving radially outwards. Although our electronics can handle up to 16 probes, we only had enough cabling to connect 12 of these probes at once. We also took some data from a subset of the rooftop-shaped probes from the second array, using alternating probes on the same toroidal row. These probes covered a larger radial cross section than the selection we used from the first array as we could connect up to 16 of these probes, since they used different cabling.

5.2.2 Floating potential measurements

We have seen from the tests on the YLPD in Section 5.1.1 that the simultaneous floating potential measurements can provide information about high speed fluctuations which cannot be resolved in the time taken to sweep the probes. In that case, the fluctuations were typically small perturbations on an otherwise steady state plasma. COMPASS plasmas are much more variable, and so lots of potentially interesting phenomena can be observed with the floating potential measurements.

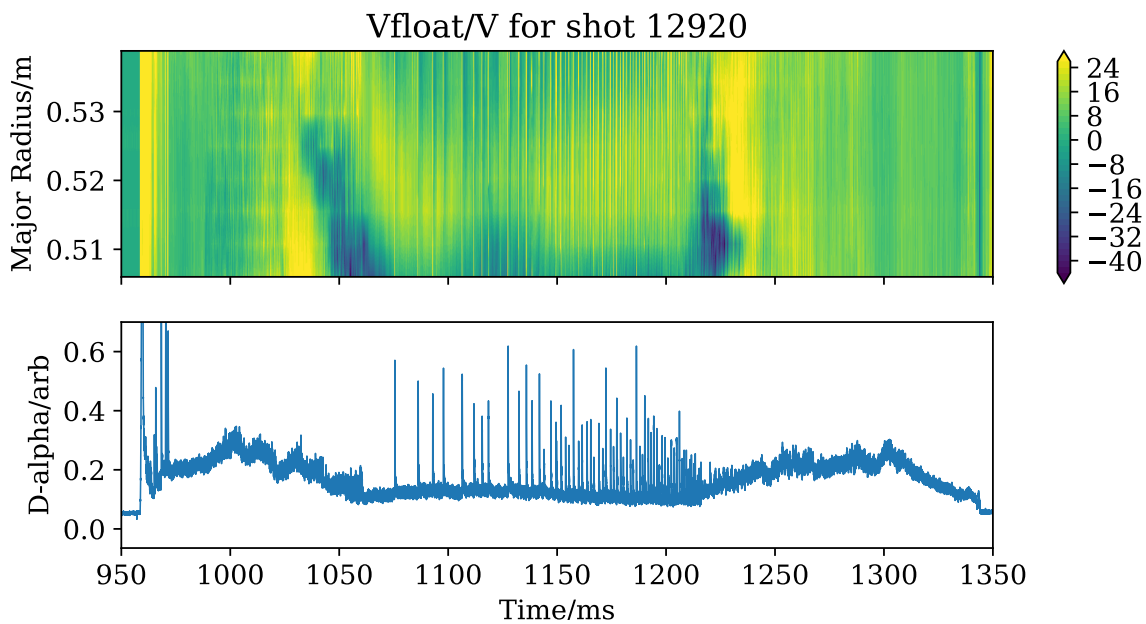


Figure 5.9: Floating potential measured by 8 adjacent probes near the outer strike point for an H-mode shot, plotted with the D-alpha signal for comparison. The plasma breakdown at 960 ms can clearly be seen by the sudden increase in floating potential from 0 V to over 20 V, and is matched by an increase in the D-alpha signal at the same time. Also, the transitions to and from H-mode at around 1070 ms and 1220 ms respectively can clearly be seen, characterised by the reduction in the D-alpha signal and the presence of ELMs (the sharp spikes in the D-alpha signal), the times of the latter matching well with sudden periodic increases in the floating potential.




Figure 5.10: The floating potential for the same COMPASS shot as Figure 5.9, zoomed to a region just before the L-H transition. Short-lived, high frequency fluctuations can be seen which appear to move radially outwards. The dashed vertical lines show the times these fluctuations first appear on the probe at the smallest major radius, estimated by eye. The dark blue blobs at 1058.5 ms and 1059.5 ms are artefacts from interpolating the floating potential over the times that particular probe was being swept.

Figure 5.9 shows the floating potential for an H-mode COMPASS shot, in which 8 probes from the first array were connected to the new electronics. Before 960 ms there is no plasma and so the floating potential is 0 V. Breakdown is shown by the sudden increase in the floating potential at 960 ms. The plasma is initially in a limiter formation and so the floating potential remains close to 0 V until the X-point is formed, which starts at around 1030 ms. The transition to H-mode at around 1075 ms is characterised by both a reduction in high-frequency fluctuations in the floating potential and by the appearance of ELMs, the large periodic fluctuations seen between 1075 ms and 1215 ms. The plasma transitions back into L-mode at this point, as seen by the similarity of the fluctuations in the floating potential compared with the period just before the start of H-mode, and by 1235 ms we again have a limiter geometry, which causes the floating potential to remain close to 0 V until the plasma is extinguished at around 1350 ms.

This description of the shot is supported by the EFIT video, the plasma current measurements and also the D-alpha emission, with the latter also shown in Figure 5.9 for comparison. We can again clearly see the breakdown, the transition to H-mode (characterised by reduced radiation and the presence of ELMs) and the subsequent transition back to L-mode. This agreement is reassuring, and means we can be confident that the data coming from the diagnostic is physical and not due to some artefacts or errors in our software and FPGA firmware.

Zooming in to two particular regions in Figure 5.9, we see some interesting features. Figure 5.10 shows the period just before the L-H transition, when the L-mode plasma




Figure 5.11: The floating potential for the same COMPASS shot as Figure 5.9, zoomed to show most of the H-mode region. The ELMs can be seen as large positive spikes in the floating potential, and there is a very repeatable inter-ELM evolution.

is in an X-point geometry. The strike point is at a major radius slightly lower than the minimum plotted here, but we can see a reduction in the floating potential as the major radius increases to around 0.515 m, which corresponds to the fall-off in the plasma flux near the strike point (though precise quantitative descriptions cannot be given using the floating potential alone). We can also see some structure in the fluctuations, including a repeatable increase in the floating potential at the start of each fluctuation followed by a large decrease. The fact that the fluctuations in the plot appear slightly tilted from the vertical (indicated by the black dashed lines) indicates some outwards radial movement of the fluctuations too. Like those measured in the YLPD, these fluctuations occur on time scales too short to be resolved by fitting an IV characteristic to a sweep. Indeed, they are actually likely to show up as noise in the IV characteristic, and if the fluctuations are large enough it will make that particular sweep unusable.

Interesting features are also present in the H-mode region, shown in Figure 5.11. We can clearly see the ELMs, which are indicated by large, short-lived fluctuations. The ELM frequency in this shot increases significantly after the resonant magnetic perturbation (RMP) coil was turned on at 1090 ms, and we can also see at the lower end of the major radius range the strike point appears to move outwards as the RMP coil is turned on (indicated by the blue region in the figure). Another interesting feature is the inter-ELM evolution of the floating potential, which appears to be very repeatable. Just after the ELM, indicated by the short spikes of large positive floating potential which show up as thin yellow vertical lines in Figure 5.11, the floating potential is at its most negative. It then rises steadily until the next ELM, where it jumps abruptly to highly positive values again. It is out of the scope of this work to postulate on the reasons for this, but the fact that we can measure such phenomena with such a simple measurement points to the success of the test.




Figure 5.12: Floating potential for COMPASS shot 12987, using 16 of the rooftop-shaped probes from the second probe array. The figure shows the change in position and width of the strike point as the plasma current was ramped up during the diverted plasma period.

Finally, Figure 5.12 shows the floating potential measured using 16 of the rooftop-shaped probes in the second probe array. The wider coverage of the divertor region means we can measure either side of the outer strike point with this probe array. In this shot the plasma current was ramped up between 1100 ms to 1200 ms after being held constant for the previous 50 ms once a diverted plasma had been achieved. We can see how the outer strike point moves inwards to a smaller major radius at the start of the current ramp (again confirmed by EFIT), and also that the width of the strike point is reduced during the current ramp compared to the flat top. There isn't enough spatial resolution to tell if the width changes with changing plasma current during the ramp, but there does appear to be a quantitative difference in the floating potential at the strike point during the short H-mode phase from 1100 ms to 1150 ms compared with the L-mode phase after that.

It should be noted that the measurements of the strike point location and width presented here can also be obtained from fitting sweeps of the probes. The strike point evolves on a timescale slow enough to be resolved by the sweeps, and fitting the sweeps can give quantitative information about the temperature and density (and hence the heat flux) in the strike point region whilst the floating potential measurements give only qualitative information. However, we have already seen in Section 5.1.3 that the fit parameters obtained from sweeps can be highly sensitive to the voltage range chosen, and sweeps can be rendered useless by large fluctuations such as ELMs occurring during the sweep. In addition, fitting IV characteristics for all 640 probes measured during a MAST-U shot would take a very long time, whereas the floating potential data is available almost instantly. The floating potential measurements are a useful, reliable source in the control room for tracking where the strike point was in the last shot, and so more useful for the machine operators than relying on fragile fitting the IV characteristics.

5.2.3 Sweep waveform

For the sweep measurements on COMPASS, we used the knowledge gained from the sweep measurements on the YLPD to design a more optimal waveform. We found in Section 5.1.2 that using a symmetric waveform was unnecessary as the new electronics had sources of parasitic current in addition to that induced by capacitance in the system. This meant that the parasitic current was not simply proportional to the rate of change of the sweep voltage, and so averaging equal and opposite parts of the sweep was not sufficient to remove this current. We could therefore remove the redundant positive voltage region from the end of the sweep shown in Figure 5.3, which reduced the rate at which we needed to sweep the voltage in order to cover the full voltage range required.

Since COMPASS produces higher density plasmas than the YLPD, we could apply a higher bias voltage to measure the saturation current without seeing as much sheath expansion (and hence measuring an un-saturated ion current). We also wanted to test the amplifier with the $65\ \mu\text{s}$ sweep period used previously on MAST and anticipated for use in MAST-U. This meant producing a voltage waveform with much higher voltage gradients than those used in the YLPD tests.

During testing of fast sweeps in development, we noticed significant distortion of the voltage, and large induced currents as a result, when the voltage waveform passed through $0\ \text{V}$. This distortion is insignificant in the data shown in Figures 5.3 and 5.4, due to the comparatively small rate of change of voltage, but it became a problem as we aimed for shorter waveforms with larger voltage ranges in the COMPASS tests. To mitigate this, we added several μs of steady $0\ \text{V}$ output into the waveform at the point in the middle of the sweep that the voltage crossed $0\ \text{V}$.

Figure 5.13 shows the resulting waveform, used for the majority of the COMPASS test shots with some variation mainly in the maximum and minimum voltages used on the different probe arrays. On the left is the waveform as it was designed, stored in the FPGA and sent to the multiplexer's DAC. Although stored as a series of digital count values, it has been plotted in volts by considering the DAC output range and the amplifier gain; this is in order to aid comparison with the actual measured amplifier output voltage. The regions of constant $0\ \text{V}$ at the start and end of the sweep are used for the same purpose as in the YLPD sweeps: to enable safely switching between probes. The region of constant $0\ \text{V}$ during the downward sweep where the voltage changes sign reduces current distortion caused by the amplifier. The widths of these constant voltage regions are chosen as a compromise between the need to reduce current distortion and the need to minimise the voltage gradient in order to produce a sweep with a shorter period. The shorter the period of the

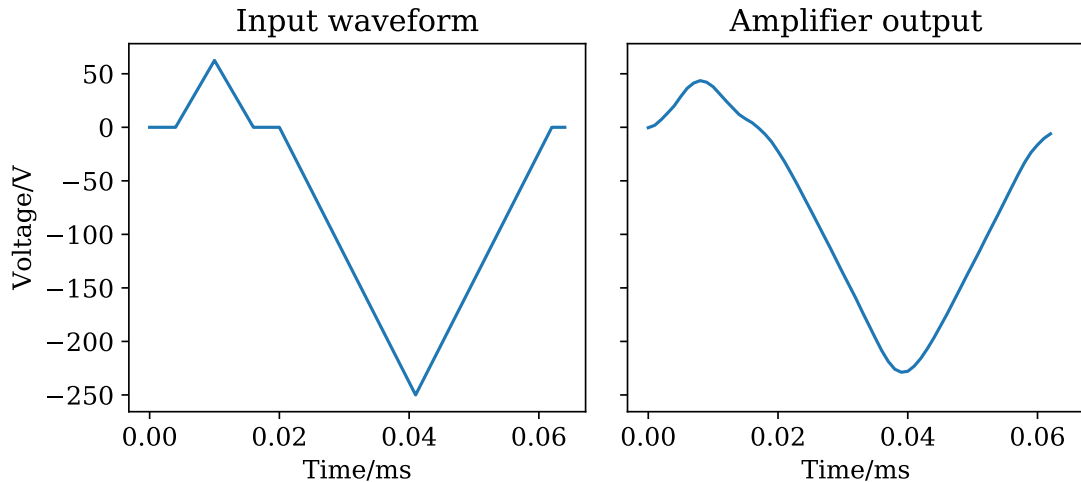


Figure 5.13: The sweep waveform used on COMPASS, showing asymmetry and a steady 0 V region when the voltage changes sign during the down sweep. Left: the digital waveform stored on the FPGA and sent to the DAC, expressed in units representing the expected amplifier output voltage. Right: the actual voltage output by the amplifier, which has been smoothed by the amplifier’s limited bandwidth.

sweep, the higher the time resolution for each individual probe. Even with a sweep of $65\ \mu\text{s}$, multiplexing through 16 probes means that each probe is only swept once every 1.04 ms, and this gives the effective time resolution of each probe.

On the right of Figure 5.13 we can see the actual voltage output by the amplifier, as measured by the multiplexer’s ADC. The amplifier bandwidth is limited to approximately 65 kHz, and this causes the sharper edges of the waveform to be rounded. In particular, this means that the maximum and minimum voltages specified by the input waveform are slightly larger than those actually achieved. The flat 0 V region is not truly flat either, but it does still have a lower voltage gradient than the rest of the sweep which helps to reduce the current distortion.

The parasitic current induced by this voltage waveform is shown in Figure 5.14. This data was taken from a COMPASS shot which failed to achieve plasma breakdown, so this provided a good reference with steady state conditions in the vessel. Each current trace has been averaged over the sweeps in a 5 ms time window (approximately 5 sweeps), for 3 different time windows throughout the shot. Although there is still some distortion around 0 V and an asymmetry between the up and down sweeps, these are nowhere near as dramatic as those seen on the YLPD (Figure 5.4).

Importantly, the induced current appears to be consistent in all 3 time windows throughout the shot. The improvements to the amplifier and the new waveform have greatly reduced the variation in the induced current which was seen on the




Figure 5.14: The parasitic current measured by a probe from the first COMPASS array in the absence of plasma, using the voltage waveform in Figure 5.13, averaged over 3 different 5 ms intervals. Although there is still some current distortion around 0 V and a noticeable lack of symmetry in the up/down sweeps, this is much less pronounced than that seen in Figure 5.4 from the YLPD tests. The current is also consistent across the 3 time windows.




Figure 5.15: The measured current in the same 3 time windows as Figure 5.14 after subtracting the average current from 800 ms to 950 ms. There appears to be no long-term evolution in the parasitic current.

YLPD. COMPASS plasmas typically start at 950 ms and last for 300 ms to 400 ms, so the induced current's consistency means that we can simply subtract the average current measured in the 150 ms before the start of the plasma from the rest of the shot. This eliminates the need to have a vacuum calibration shot for each waveform, which was the case in the YLPD tests. Figure 5.15 demonstrates this, showing that subtracting the averaged pre-plasma current from the sweeps in the 3 time windows shown in Figure 5.14 produces no saturation current drift. All that remains is the random noise on the measurement, which (when smoothed over 4 to 5 sweeps) has an RMS of around 7 mA for negative voltages and around 30 mA near 0 V.

5.2.4 Sweep measurements

In Section 5.2.3 we discussed the development of a more optimised sweep waveform, and showed that subtracting the current measured in the short period of a shot before the plasma was sufficient to remove unwanted sources of current. With this in place, we could produce some IV characteristics and examine whether they would be suitable for MAST-U.

Figure 5.16 shows the IV characteristics produced from a particularly interesting series of sweeps. We have plotted here 6 successive sweeps for 4 adjacent probes from the first array, performed during the inter-ELM period of an H-mode plasma. Probe 025 is closest to the outer strike point, and the major radius increases with increasing probe number. The time of each sweep, given as the average time of the whole pass through all the probes being multiplexed in this shot, is shown as a vertical line of the same colour in Figure 5.17, which shows the D-alpha signal and the two ELMs either side of this inter-ELM period.

We see a number of notable features. Firstly, the ion saturation current and the (un-saturated) electron current decrease as we move to larger major radii, which gives some information about the width of the strike point. In addition, the current in both the ion and electron saturation regions reduces for all probes throughout the inter-ELM period, which provides information about the evolution of the heat flux between ELMs.

Of most interest from a diagnostic viewpoint is what looks like a hysteresis loop in the positive voltage region, which is seen in all 4 probes. The loop does not appear to be caused by plasma fluctuations on the time scale of an individual sweep, since it persists not only across multiple probes but across multiple sweeps of the same probe. It is therefore likely to be a diagnostic effect. The most likely explanation is that the high rate of change of current in this region (up to 200 kA/s nearest the strike point) is near the limit of what the amplifier can provide. The current therefore




Figure 5.16: A series of IV characteristics for 4 probes near the outer strike point, in the inter-ELM period of an H-mode plasma. Different colours represent subsequent sweeps, the times of which are shown in Figure 5.17. Probe 025 is nearest the outer strike point, and probe 028 is at a major radius approximately 20 mm larger. The reduction in both ion and electron currents both temporally through the ELM cycle (from the blue to the brown curves for each probe) and spatially across the probes can be seen, as well as an unexpected loop in the electron current region.

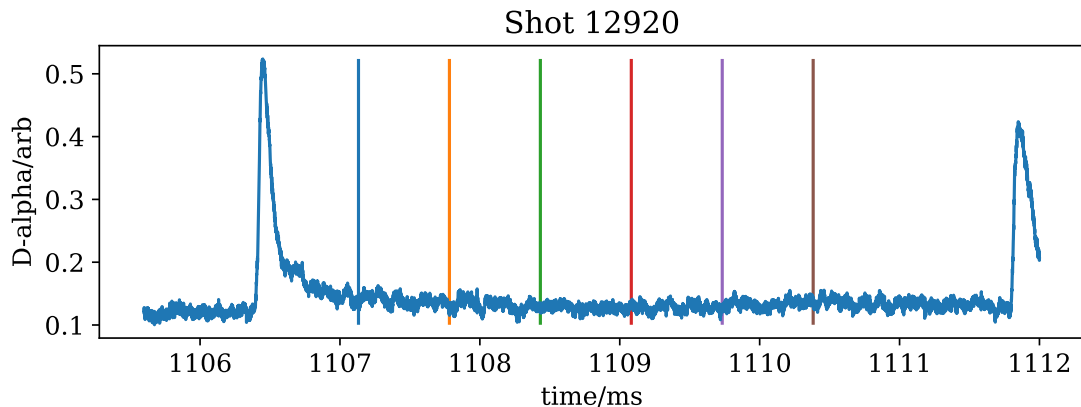


Figure 5.17: D-alpha signal, showing the times of the sweeps plotted in Figure 5.16. The average time of each sweep sequence in Figure 5.16 (i.e. the mean time in the range between the start of the first probe sweep and the end of the last probe sweep) is shown as a vertical line of the same colour as the sweep.

lags behind the voltage, appearing smaller than it should be during the up sweep and larger during the down sweep, which gives rise to the loop. This hypothesis is supported by the fact that the loop is narrower for smaller currents, where the rate of change of current is smaller.

It should be possible to mitigate this effect through further optimisation of the waveform. The ion saturation current is significantly smaller than the electron current, so we can change the rate at which the voltage is swept through each of these regions. By designing the waveform to sweep slowly in the $V > 0$ V region (for a low dI/dt), we can minimise the demands on the amplifier. The much lower magnitude of the ion saturation current means that we can sweep the $V < 0$ V region much more quickly without issues with the current slew, whilst still keeping the overall waveform length short enough for acceptable time resolution.

The experimental programme and subsequent engineering shutdown of COMPASS following our tests, along with the need to return the electronics to CCFE to aid in the assembly of the rest of the multiplexers for MAST-U, meant that there was little opportunity to test this proposed solution. Nevertheless, it is important to be aware of, and the knowledge gained will be used to optimise the design of sweep waveforms for MAST-U.

5.2.5 Summary of the COMPASS tests

The COMPASS tests have demonstrated that the system is suitable for operation in a tokamak environment. We have shown that we can measure L-mode fluctuations

in the floating potential with sufficient precision and time resolution to illustrate plasma behaviour, with the simultaneous floating potential measurements revealing the radial motion of those fluctuations. We have also demonstrated the ability to track the position and width of the strike point, with minimal data processing required, using simultaneous floating potential measurements. The fast availability of this data can be used to aid machine operators in the experiment control room.

We have demonstrated the successful removal of parasitic induced current during high speed, high voltage sweeps by using the average current measured just before the plasma in each shot. We have further shown that improvements to the electronics since the YLPD tests have eliminated the drift in the parasitic current throughout the data acquisition window. We have therefore provided a reliable, repeatable method to remove unwanted current from the measurements without requiring dedicated calibration shots which promises to be more accurate than the method of averaging up and down sweeps which was employed previously on MAST.

We have also discovered an interesting phenomenon in sweeps with a very high rate of change of current. The current lags behind the voltage in these sweeps, which produces an IV characteristic which appears to have a hysteresis loop. It is expected that this phenomenon can be mitigated by designing a waveform which results in a lower rate of change of current in the electron current region. This knowledge will be taken into consideration when designing the sweep waveforms for MAST-U.

Finally, although not shown here, sweep data from the second probe array was also gathered. This data was used to define the parameters used for fast electron temperature and heat load measurements using the rooftop-shaped and ball-pen probes. The work has been published in the journal Nuclear Fusion [89].

5.3 Implications for installation on MAST-U

The successful testing of the Langmuir probe electronics on both the York Linear Plasma Device and the COMPASS tokamak has provided multiple benefits to MAST-U. By demonstrating that the system produces valid and useful data for a range of plasma scenarios we have significantly reduced the risk that the system will not work as expected when it is installed on MAST-U. The software which was developed to integrate the diagnostic into COMPASS's data acquisition and control system has now been tested through a tokamak shot cycle, and was written in such a way that it can be easily adapted to MAST-U's requirements. This will accelerate the commissioning process for the diagnostic when it is installed on MAST-U.

Improvements made to the electronics based on data gathered on the YLPD were proven to have been successful with the installation on COMPASS. We have uncovered several important features of the system which must be taken into account to maximise the diagnostic's capabilities on MAST-U. For example, we have refined the method of removing parasitic current from the sweep measurements and made steps towards producing an optimal sweep voltage waveform to fully utilise the arbitrary waveform generation capabilities of the new FPGA-based electronics. We also have a better understanding of the limits of the diagnostic's capabilities, such as the extent of the current distortion from the amplifier as it transitions from positive to negative voltages, the smoothing of waveforms due to the amplifier's limited bandwidth and the issues with a very high rate of change of current.

Armed with this knowledge, we are well positioned to fully exploit the capabilities of this advanced new Langmuir probe diagnostic. We have a more compact system featuring improved operational flexibility and enhanced measurement capabilities compared with MAST's previous Langmuir probe diagnostic. In conjunction with the unprecedented number of probes installed on the MAST-U tokamak, we will be able to provide high quality data for both physicists and experiment operators.

Chapter 6

Conclusions

Harnessing the power of nuclear fusion to provide a clean, safe and reliable energy source is of key importance. The very challenging conditions involved require a thorough understanding of the physics behind fusion plasmas, and this cannot be achieved without a suite of high quality plasma diagnostics. Current and future tokamak experiments require accurate, fast measurements by diagnostics using reliable hardware which needs to be compact enough to fit in amongst the plethora of equipment required to operate the plant. These diagnostics should also be flexible enough to adapt to evolving requirements as our understanding of the physics of fusion plasmas is improved.

Whilst it is difficult to meet these demands using traditional analogue hardware, the power and flexibility of FPGAs provides an invaluable tool in the diagnostician's arsenal. We are able to exploit the programmable nature of the FPGA to implement complex functionality in a digital circuit, which greatly improves the compactness of the design and reduces the cost. The ability to re-program the FPGA enables enhancements to be made even after the diagnostic has been developed, installed and commissioned, offering flexibility that other technologies cannot match. We have demonstrated that this technology can be exploited to provide high quality diagnostics to the MAST-Upgrade tokamak.

6.1 Summary

We have presented the design of a new FPGA-based bolometer system. Using a modular hardware platform provided by D-TACQ solutions, the system is capable of supplying an AC excitation voltage to the bolometer sensor's Wheatstone bridge and digitising the output voltage from the bridge. The FPGA performs AC synchronous detection in quadrature in order to extract the voltage output due to the

temperature-induced bridge imbalance and filter out the myriad sources of electrical noise in the tokamak environment. The voltage can then be analysed by post-shot software to calculate the power incident on the bolometer sensors and hence infer the emission profile of the plasma. Using a specially designed deconvolution filter for each bolometer sensor, the system is also able to output a measure of the power incident on that sensor with around 1 ms latency and 1 kHz bandwidth directly from the FPGA, which could be used in a feedback loop for real time control in the future.

The new hardware also enables a new calibration method, which is more relevant to the measurements made during shots. By applying a DC voltage to ohmically heat the sensor whilst also supplying the AC excitation voltage, we can measure the bridge imbalance with the same AC circuit effects which are present in measurements made during shots. This removes the need to apply difficult-to-calculate correction factors to the calibration parameters obtained, and therefore reduces sources of systematic error in the measurements. The calibration procedure can be performed simultaneously for all channels and takes only a few seconds, meaning it is possible to calibrate the bolometer sensors before every shot if desired. This allows us to obtain a large number of calibrations and therefore provides better statistics for the calibration parameters we are interested in.

The system has been tested in several locations. Tests with a sensor in vacuum at York University provided feedback to enable us to improve the signal processing part of the FPGA firmware, increasing the number of bits to avoid discretisation errors. We also found a significant variation in the sensitivity and cooling time calibration parameters at pressures which are likely to be encountered in the MAST-U divertor, and this knowledge will be used when analysing the bolometer measurements during MAST-U experiments. The installation of the system on the JET tokamak demonstrated that the new electronics produces measurements in good quantitative agreement with the existing and well trusted bolometer diagnostic, but with a better signal-to-noise ratio even at higher bandwidth.

Tests carried out at TCV and PPPL using original and revised firmware respectively have shown that the system can be rapidly integrated into other experiments, and this has led to the purchase of the bolometer electronics by these institutions. The PPPL work was used to quantify the performance of the bolometer system, and we demonstrated the ability to provide absolute power measurements to within a few percent of the expected values. We also demonstrated the consistent performance of the system across a wide range of operating regimes, such as input power, drive frequency and cable length.

Overall, we have shown the new system to be more capable than previous generations of bolometer electronics whilst delivering improvements in compactness, flexibility

and ease of integration at a significantly lower cost. Recall the original research question, which we described in Section 2.2: could we produce significantly more compact, lower cost bolometer electronics, with improved operational flexibility, using FPGA technology. The work we have presented here suggests that this is indeed the case: a 48-channel system now fits in a 1U, 19-inch rack rather than several cubicles, the cost per channel is about a fifth of the older analogue electronics and we now have the ability to vary parameters such as the drive frequency and filter bandwidth. We also sought to investigate the potential for digital signal processing on the FPGA to enable the bolometer electronics to become part of a control loop. The successful calculation in real time of the power measured by each individual bolometer sensor indicates the feasibility of this use case.

Next, we presented the design of the new Langmuir probe electronics. In a similar manner, much of the functionality that was previously performed with analogue electronics was implemented using FPGA technology, resulting in significant cost and physical footprint savings in addition to the improved flexibility. The new system combines an arbitrary waveform generator, voltage and current measurements and probe multiplexing in a single unit, all controlled by a Zynq ARM/FPGA combination. Connected to a high voltage, high bandwidth amplifier and a power supply, we have a highly capable electronics module which provides multiplexed biasing and measurement of up to 16 Langmuir probes. A suite of 40 such modules will be able to operate up to 640 Langmuir probes, giving excellent time and spatial resolution of heat and particle flux measurements in MAST-U's novel Super-X divertor. We have achieved our aim of producing a more compact multiplexer unit (which also fits into a 1U, 19-inch rack), and the use of the combination of CPU and FPGA on the Zynq chip has enabled us to integrate all of the functionality onto the multiplexer unit that we initially described in Section 4.2.

Demonstration of the correct functioning of the prototype electronics was achieved on the York Linear Plasma Device. The new ability of the system to measure the floating potential simultaneously on all probes was used to investigate the behaviour of fluctuations at the onset of detachment, which will aid understanding of this extremely important phenomenon. Experience gained during the tests was also used to improve the calibration procedure of the probes.

An improved set of probe electronics was then installed on the COMPASS tokamak, to prove the system was fit for purpose with measurements of relevant tokamak plasmas. Again, the ability to make high speed simultaneous floating potential measurements was used to observe the behaviour of fluctuations in the divertor plasma, and the evolution of the plasma in the scrape-off layer during the inter-ELM cycle was demonstrated by obtaining a series of IV characteristics. The tests

highlighted issues which were due to plasma conditions not available on the York Linear Plasma Device but may be relevant to MAST-U, and the knowledge gained from this will be used to influence waveform design and operation of the probes during the first MAST-U campaigns.

Overall, the design and implementation of these two systems has improved the quality of MAST-U's already world class suite of diagnostics. The bolometer work on TCV and JET has been published in two separate proceedings papers in the Review of Scientific Instruments journal, and the Langmuir probe work in York and on COMPASS has been published in the JINST journal. This raises the profile of the work being done and highlights the potential of FPGAs to the wider fusion community.

6.2 Future work

The most immediate task for both of the new diagnostics is commissioning on MAST-U. Upgrade work on the tokamak is currently expected to finish in early 2018, and work to integrate both systems into MAST-U's data acquisition infrastructure will need to be completed during the commissioning and restart phases of the device. Fortunately, both the bolometer and Langmuir probe systems have been tested on other tokamaks, and so much of the commissioning work will involve simply adapting the diagnostics' control software to MAST-U's requirements rather than writing code from scratch.

The two diagnostics provide complimentary information regarding some of MAST-U's key research activities. Both are important in the study of detachment, and are going to be crucial to validating the predictions of the performance of the new Super-X divertor. The Langmuir probe waveforms and data analysis will need to be optimised to provide reliable, accurate measurements of heat and particle fluxes, since the reduction of the heat flux is a key prediction of the Super-X. This work will be lead by Sarah Elmore and James Harrison at CCFE, but detailed knowledge of the electronics system and the FPGA's capabilities will be required to fully exploit the new diagnostic.

For example, further investigation of the "hysteresis" shown in Figure 5.16 is required, to properly determine whether this is a plasma effect or an electronics effect. Focussed, lab-based benchmarking of the electronics, similar to the detailed benchmarking tests of the bolometer system performed in Section 3.4, would be extremely useful in improving our understanding of the complex electronics of the new system. Using a diode as the load, which has a similar I-V curve to a plasma, could provide

insight into the behaviour of the electronics in regions with extremely high rates of change of current without the added complexity of plasma effects. Additionally, some circuit modelling of the electronics could both aid our understanding of this behaviour and also provide a means to correct for it if it is indeed an electronics effect. We showed in Section 3.4.8 that modelling the bolometer circuit accurately reproduced the behaviour we saw in the lab and could even be used to provide an empirical correction factor to power measurements if needed; it is worthwhile performing a similar analysis of the probe circuit to see if we can achieve the same thing.

The use of FPGA technology in Langmuir probes has caught the attention of developers of similar systems. Whilst this thesis was being written, several meetings took place between CCFE and researchers at the Massachusetts Institute of Technology (MIT) which looked into the feasibility of using an FPGA to improve the design of their Mirror Langmuir probe (MLP) diagnostic. This promises to greatly simplify the design of this exciting diagnostic, which provides extremely high resolution measurements of the ion saturation current, floating potential and electron temperature [105]. Early discussions suggest that an FPGA is both capable and well suited to this application, and a collaboration has been initiated between CCFE and MIT to bring together MIT's knowledge of the MLP and CCFE's FPGA expertise.

There are a large number of investigations to be carried out with the bolometer. Although the calibration and bolometer voltage measurements have been shown to be accurate and reliable in air and vacuum, the tests in York (Section 3.1) showed that the neutral pressure in the vicinity of the bolometer sensors can have a significant impact on the measurements. We need a good understanding of this effect and a way to reliably account for it in order to make accurate measurements of the radiation in the divertor.

We may for example be able to use a bolometer channel screened from plasma radiation, but experiencing the same neutral pressure as channels exposed to the radiation, as a neutral pressure gauge. This channel would run a series of calibrations throughout the shot, and we could produce a time history of the cooling time of the channel (and hence an estimate of the neutral pressure) for each shot. The variation of this cooling time could then be applied to the assumed cooling times of the other bolometer channels in the same array. This would unfortunately reduce the number of channels available for radiation measurements in MAST-U, and require significant modification of the FPGA firmware and software. It may however be unavoidable, if the neutral pressure has a significant effect on the bolometer calibrations and we cannot get reliable neutral pressure measurements from the pressure gauges already available in MAST-U.

In addition to the neutral pressure issues, Umar Sheikh's work with the system at TCV showed that unblackened bolometers, such as those installed in MAST-U, may underestimate the radiated power by as much as 20% due to radiation being reflected rather than absorbed by the sensor [78]. As a result, detailed comparisons of bolometer measurements with simulations must be made to ensure we are accurately accounting for this effect.

Once we can be sure that all the physics involved in measuring the radiated power is understood, the bolometer system can be utilised with confidence. Studies of power balance will be important, to ensure that we understand where the power is leaving the plasma. The bolometer will also be extremely useful in studies of detachment, which is characterised by strong radiating emission. The divertor bolometer's lines of sight will enable us to accurately reconstruct the emissivity profile in the divertor chamber and track the position of the detachment front. This will allow us to benchmark analytical models of the sensitivity of the detachment front to control parameters (as presented in [40]) and therefore optimise the radiative cooling of the exhaust plasma to further reduce the divertor target heat flux.

Looking further ahead, the ability of the FPGA system to calculate the power in real time could be exploited to enable high quality detachment control. Preliminary investigations have shown that it is possible to estimate the position and extent of the detachment region (and by extension the position of the detachment front) in real-time using the signal processing capabilities of the FPGA [106]. This means we can use bolometer measurements as part of a system for real-time control of the detachment front. By optimising the position of the detachment front we can maximise the amount of radiative cooling in the divertor without contaminating the core plasma, to enable high performance core plasmas with greatly reduced heat loads on the divertor target. This will help to solve the divertor exhaust problem described in Section 1.3.3, a key aim of the MAST-U project in general, and will be extremely important in the quest for a viable commercial tokamak reactor.

6.3 Final word: the potential of integrated FPGA diagnostics in fusion

We have seen in this work that the use of FPGA technology enables us to make diagnostics more compact, adaptable and smart enough to function as standalone units to provide improved modularity. This is hugely beneficial for fusion research, where the ability to quickly add new diagnostics or improve existing ones enables experiments to produce a wider range of high quality data. By using generic hardware

which is specialised with custom firmware, consistency can be achieved between different diagnostics (and even different experiments) which can help to reduce costs and speed up the time between the conceptual design of a diagnostic and the finished product producing good science.

We have also seen that the high performance digital signal processing capabilities of FPGAs can enable us to transform diagnostics into part of a control or feedback loop. As we move from experimental to commercial reactors, the role of diagnostics will change from physics studies to inputs for controllers, so the ability to provide this sort of functionality puts FPGAs in a strong position to be heavily utilised in future fusion machines. Combined with their low cost and modularity, FPGA-based diagnostics have an important role to play in fusion's quest to provide a significant part of humanity's future energy needs.

Bibliography

- [1] The Python Software Foundation. *Welcome to Python.org*.
<https://www.python.org/> (visited on 03/09/2017).
- [2] Eric Jones, Travis Oliphant, Pearu Peterson et al. *SciPy: Open source scientific tools for Python*. 2001.
<http://www.scipy.org/>.
- [3] J. D. Hunter. ‘Matplotlib: A 2D graphics environment’. In: *Computing In Science & Engineering* 9.3 (2007), pp. 90–95.
doi:10.1109/MCSE.2007.55.
- [4] S. Hoyer and J. Hamman. ‘xarray: N-D labeled Arrays and Datasets in Python’. In: *Journal of Open Research Software* 5.1 (2017).
doi:http://doi.org/10.5334/jors.148.
- [5] Till Tantau. *TikZ/PGF — Graphic system for T_EX*.
<http://pgf.sourceforge.net/> (visited on 03/09/2017).
- [6] The Inkscape Team. *Draw Freely | Inkscape*.
<https://inkscape.org/en/> (visited on 03/09/2017).
- [7] IEA Secretariat. *World Energy Outlook 2016 - Executive Summary*. Available online at <https://www.iea.org/publications/freepublications/publication/WorldEnergyOutlook2016ExecutiveSummaryEnglish.pdf>. IEA Publications, 2016.
- [8] Shahriar Shafiee and Erkan Topal. ‘When will fossil fuel reserves be diminished?’ In: *Energy Policy* 37.1 (2009), pp. 181–189.
<http://www.sciencedirect.com/science/article/pii/S0301421508004126>
doi:http://dx.doi.org/10.1016/j.enpol.2008.08.016.
- [9] Joeri Rogelj, Michel Den Elzen, Niklas Höhne, Taryn Fransen, Hanna Fekete, Harald Winkler, Roberto Schaeffer, Fu Sha, Keywan Riahi and Malte Meinshausen. ‘Paris Agreement climate proposals need a boost to keep warming well below 2 C’. In: *Nature* 534.7609 (2016), pp. 631–639.
doi:http://dx.doi.org/10.1038/nature18307.

- [10] Thomas B Johansson and Laurie Burnham. *Renewable energy: sources for fuels and electricity*. Island press, 1993.
- [11] OECD Nuclear Energy Agency and International Atomic Energy Agency. *Uranium 2014: Resources, Production and Demand*. OECD Nuclear Energy Agency, 2014.
- [12] Wikimedia Commons. *Binding energy curve - common isotopes*.
https://upload.wikimedia.org/wikipedia/commons/5/53/Binding_energy_curve_-_common_isotopes.svg (visited on 25/08/2017).
- [13] H.-S. Bosch and G.M. Hale. ‘Improved formulas for fusion cross-sections and thermal reactivities’. In: *Nuclear Fusion* 32.4 (1992), p. 611.
<http://stacks.iop.org/0029-5515/32/i=4/a=I07>.
- [14] John Wesson. *Tokamaks*. 2nd. Oxford University Press, 1997.
- [15] Culham Centre for Fusion Energy. *Fusion - A clean future*.
http://www.ccfе.ac.uk/assets/Documents/CPS16.49_May2016_low.pdf
(visited on 25/08/2017).
- [16] J. Knaster et al. ‘The accomplishment of the Engineering Design Activities of IFMIF/EVEDA: The European–Japanese project towards a Li(d,xn) fusion relevant neutron source’. In: *Nuclear Fusion* 55.8 (2015), p. 086003.
<http://stacks.iop.org/0029-5515/55/i=8/a=086003>.
- [17] A Sykes, MP Gryaznevich, G Voss, D Kingham and B Kuteev. ‘Fusion for neutrons: A realizable fusion neutron source’. In: *IEEE Transactions on Plasma Science* 40.3 (2012), pp. 715–723.
- [18] L. Giancarli, V. Chuyanov, M. Abdou, M. Akiba, B.G. Hong, R. Lässer, C. Pan and Y. Strebkov. ‘Breeding Blanket Modules testing in ITER: An international program on the way to DEMO’. In: *Fusion Engineering and Design* 81.1 (2006). Proceedings of the Seventh International Symposium on Fusion Nuclear Technology, pp. 393–405.
<http://www.sciencedirect.com/science/article/pii/S0920379605005302>
doi:<http://dx.doi.org/10.1016/j.fusengdes.2005.08.096>.
- [19] John Lindl. ‘Development of the indirect-drive approach to inertial confinement fusion and the target physics basis for ignition and gain’. In: *Physics of Plasmas* 2.11 (1995), pp. 3933–4024.
<http://dx.doi.org/10.1063/1.871025>
doi:10.1063/1.871025.
arXiv:<http://dx.doi.org/10.1063/1.871025>.

- [20] General Fusion. *General Fusion – Clean Energy. Everywhere. Forever.*
<http://generalfusion.com/> (visited on 30/08/2017).
- [21] Ronald C. Kirkpatrick, Irvin R. Lindemuth and Marjorie S. Ward.
'Magnetized Target Fusion: An Overview'. In: *Fusion Technology* 27.3 (1995),
pp. 201–214.
<http://dx.doi.org/10.13182/FST95-A30382>
doi:10.13182/FST95-A30382.
arXiv:<http://dx.doi.org/10.13182/FST95-A30382>.
- [22] Todd H. Rider. 'A general critique of inertial-electrostatic confinement
fusion systems'. In: *Physics of Plasmas* 2.6 (1995), pp. 1853–1872.
<http://dx.doi.org/10.1063/1.871273>
doi:10.1063/1.871273.
arXiv:<http://dx.doi.org/10.1063/1.871273>.
- [23] J. D. Jackson. 'Catalysis of Nuclear Reactions between Hydrogen Isotopes
by μ^- Mesons'. In: *Phys. Rev.* 106 (2 Apr. 1957), pp. 330–339.
<https://link.aps.org/doi/10.1103/PhysRev.106.330>
doi:10.1103/PhysRev.106.330.
- [24] Christian Bachmann et al. 'Initial DEMO tokamak design configuration
studies'. In: *Fusion Engineering and Design* 98 (2015). Proceedings of the
28th Symposium On Fusion Technology (SOFT-28), pp. 1423–1426.
<http://www.sciencedirect.com/science/article/pii/S0920379615003257>
doi:<http://dx.doi.org/10.1016/j.fusengdes.2015.05.027>.
- [25] Lyman Spitzer Jr. 'The Stellarator Concept'. In: *The Physics of Fluids* 1.4
(1958), pp. 253–264.
<http://aip.scitation.org/doi/abs/10.1063/1.1705883>
doi:10.1063/1.1705883.
arXiv:<http://aip.scitation.org/doi/pdf/10.1063/1.1705883>.
- [26] Ben Dudson. 'Magnetic Confinement Fusion'. Lecture course notes,
University of York.
- [27] M. Keilhacker et al. 'High fusion performance from deuterium-tritium
plasmas in JET'. In: *Nuclear Fusion* 39.2 (1999), p. 209.
<http://stacks.iop.org/0029-5515/39/i=2/a=306>.
- [28] EUROfusion.
<https://www.euro-fusion.org/wpcms/wp-content/uploads/2011/09/jg05-537-1c.jpg> (visited on 31/08/2017).

- [29] I.H. Hutchinson. *Principles of Plasma Diagnostics*. Cambridge, UK: Cambridge University Press, 2005.
<http://books.google.co.uk/books?id=pUUZKLR00RIC>.
- [30] C. Holland, A. E. White, G. R. McKee, M. W. Shafer, J. Candy, R. E. Waltz, L. Schmitz and G. R. Tynan. ‘Implementation and application of two synthetic diagnostics for validating simulations of core tokamak turbulence’. In: *Physics of Plasmas* 16.5 (2009), p. 052301.
<http://dx.doi.org/10.1063/1.3085792>
doi:10.1063/1.3085792.
arXiv:<http://dx.doi.org/10.1063/1.3085792>.
- [31] Francesco Romanelli, P Barabaschi, D Borba, G Federici, L Horton, R Neu, D Stork and H Zohm. ‘A roadmap to the realization of fusion energy’. In: *EFDA, Garching, Germany* (2012).
- [32] G.F. Matthews. ‘Plasma detachment from divertor targets and limiters’. In: *Journal of Nuclear Materials* 220-222.Supplement C (1995). Plasma-Surface Interactions in Controlled Fusion Devices, pp. 104–116.
<http://www.sciencedirect.com/science/article/pii/0022311594004501>
doi:[https://doi.org/10.1016/0022-3115\(94\)00450-1](https://doi.org/10.1016/0022-3115(94)00450-1).
- [33] A. W. Morris. ‘MAST: Results and Upgrade Activities’. In: *IEEE Transactions on Plasma Science* 40.3 (Mar. 2012), pp. 682–691.
doi:10.1109/TPS.2011.2181540.
- [34] Joe Milnes, Nizar Ben Ayed, Fahim Dhalla, Geoff Fishpool, John Hill, Ioannis Katramados, Richard Martin, Graham Naylor, Tom O’Gorman and Rory Scannell. ‘MAST Upgrade – Construction Status’. In: *Fusion Engineering and Design* 96 (2015). Proceedings of the 28th Symposium On Fusion Technology (SOFT-28), pp. 42–47.
<http://www.sciencedirect.com/science/article/pii/S092037961500160X>
doi:<http://dx.doi.org/10.1016/j.fusengdes.2015.03.002>.
- [35] Y-K.M. Peng and D.J. Strickler. ‘Features of spherical torus plasmas’. In: *Nuclear Fusion* 26.6 (1986), p. 769.
<http://stacks.iop.org/0029-5515/26/i=6/a=005>.
- [36] M. Gryaznevich et al. ‘Achievement of Record β in the START Spherical Tokamak’. In: *Phys. Rev. Lett.* 80 (18 May 1998), pp. 3972–3975.
<https://link.aps.org/doi/10.1103/PhysRevLett.80.3972>
doi:10.1103/PhysRevLett.80.3972.

- [37] M. Gryaznevich, V. Shevchenko and A. Sykes. ‘Plasma formation in START and MAST spherical tokamaks’. In: *Nuclear Fusion* 46.8 (2006), S573.
<http://stacks.iop.org/0029-5515/46/i=8/a=S02>.
- [38] P. M. Valanju, M. Kotschenreuther, S. M. Mahajan and J. Canik. ‘Super-X divertors and high power density fusion devices’. In: *Physics of Plasmas (1994-present)* 16.5, 056110 (2009).
<http://scitation.aip.org/content/aip/journal/pop/16/5/10.1063/1.3110984>
[doi:http://dx.doi.org/10.1063/1.3110984](http://dx.doi.org/10.1063/1.3110984).
- [39] 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’. In: *Journal of Nuclear Materials* 438 (2013). Proceedings of the 20th International Conference on Plasma-Surface Interactions in Controlled Fusion Devices, S356–S359.
<http://www.sciencedirect.com/science/article/pii/S0022311513000755>
[doi:http://dx.doi.org/10.1016/j.jnucmat.2013.01.067](http://dx.doi.org/10.1016/j.jnucmat.2013.01.067).
- [40] Bruce Lipschultz, Felix I. Parra and Ian H. Hutchinson. ‘Sensitivity of detachment extent to magnetic configuration and external parameters’. In: *Nuclear Fusion* 56.5 (2016), p. 056007.
<http://stacks.iop.org/0029-5515/56/i=5/a=056007>.
- [41] D. D. Ryutov, R. H. Cohen, T. D. Rognlien and M. V. Umansky. ‘The magnetic field structure of a snowflake divertor’. In: *Physics of Plasmas* 15.9 (2008), p. 092501.
<http://dx.doi.org/10.1063/1.2967900>
[doi:10.1063/1.2967900](http://dx.doi.org/10.1063/1.2967900).
[arXiv:http://dx.doi.org/10.1063/1.2967900](http://dx.doi.org/10.1063/1.2967900).
- [42] Mike Kotschenreuther, Prashant Valanju, Brent Covele and Swadesh Mahajan. ‘Magnetic geometry and physics of advanced divertors: The X-divertor and the snowflake’. In: *Physics of Plasmas* 20.10 (2013), p. 102507.
<http://dx.doi.org/10.1063/1.4824735>
[doi:10.1063/1.4824735](http://dx.doi.org/10.1063/1.4824735).
[arXiv:http://dx.doi.org/10.1063/1.4824735](http://dx.doi.org/10.1063/1.4824735).
- [43] FROMGENTOGEN.US.
<http://fromgentogen.us/xilinx-fpga-architecture/xilinx-fpga-architecture-imposing-on-architecture-xilinx-fpga-4/> (visited on 04/10/2017).

- [44] *VHDL Language Reference Manual*. IEEE Standard 1076-1993.
- [45] *Standard Description Language Based on the VerilogTM Hardware Description Language*. IEEE Standard 1364-1995.
- [46] Digilent Inc.
https://reference.digilentinc.com/_media/zybo/zyng1.png?cache=
(visited on 01/09/2017).
- [47] Xilinx Inc.
<https://www.xilinx.com/products/silicon-devices/soc/zynq-7000.html>
(visited on 01/09/2017).
- [48] Eli Billauer. *Xillinux: A Linux distribution for Zedboard, ZyBo, MicroZed and SocKit | xillybus.com*.
<http://xillybus.com/xillinux>.
- [49] A. W. Degeling, H. Weisen, A. Zabolotsky, B. P. Duval, R. A. Pitts, M. Wischmeier, P. Lavanchy, Ph. Marmillod and G. Pochon. ‘AXUV bolometer and Lyman- α camera systems on the TCV tokamak’. In: *Review of Scientific Instruments* 75.10 (2004), pp. 4139–4141.
<http://dx.doi.org/10.1063/1.1787131>
doi:10.1063/1.1787131.
arXiv:<http://dx.doi.org/10.1063/1.1787131>.
- [50] K. F. Mast, H. Krause, K. Behringer, A. Bulliard and G. Magyar. ‘Bolometric diagnostics in JET’. In: *Review of Scientific Instruments* 56.5 (1985), pp. 969–971.
<http://dx.doi.org/10.1063/1.1138007>
doi:10.1063/1.1138007.
arXiv:<http://dx.doi.org/10.1063/1.1138007>.
- [51] G. A. Wurden and B. J. Peterson. ‘Development of imaging bolometers for long-pulse MFE experiments (invited)’. In: *Review of Scientific Instruments* 70.1 (1999), pp. 255–259.
<http://dx.doi.org/10.1063/1.1149441>
doi:10.1063/1.1149441.
arXiv:<http://dx.doi.org/10.1063/1.1149441>.
- [52] Karl Hoffmann. *Applying the Wheatstone bridge circuit*. HBM, 1974.
- [53] 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’. In: *Review of Scientific Instruments* 62.3 (1991), pp. 744–750.

- <http://scitation.aip.org/content/aip/journal/rsi/62/3/10.1063/1.1142078>
[doi:http://dx.doi.org/10.1063/1.1142078](http://dx.doi.org/10.1063/1.1142078).
- [54] Paul Horowitz and Winfield Hill. *The Art of Electronics*. 2nd. Cambridge University Press, 1989.
- [55] M. B. Weissman. ‘1/f noise and other slow, nonexponential kinetics in condensed matter’. In: *Rev. Mod. Phys.* 60 (2 Apr. 1988), pp. 537–571.
<https://link.aps.org/doi/10.1103/RevModPhys.60.537>
[doi:10.1103/RevModPhys.60.537](http://dx.doi.org/10.1103/RevModPhys.60.537).
- [56] A Huber and P Beaumont. ‘JET Bolometer shift manual’.
<http://users.euro-fusion.org/pages/bolom/documents/Bolometershiftmanual.pdf>.
- [57] L. Giannone, K. Mast and M. Schubert. ‘Derivation of bolometer equations relevant to operation in fusion experiments’. In: *Review of Scientific Instruments* 73.9 (2002), pp. 3205–3214.
<http://scitation.aip.org/content/aip/journal/rsi/73/9/10.1063/1.1498906>
[doi:http://dx.doi.org/10.1063/1.1498906](http://dx.doi.org/10.1063/1.1498906).
- [58] A. Huber. ‘Manual for KB5 PPF production’. http://users.euro-fusion.org/pages/bolom/documents/KB5_PPF_Manual.pdf.
- [59] CCFE. *Fusion: Fusion - a clean future*.
<http://www.ccfе.ac.uk/>.
- [60] D-TACQ Solutions. *D-TACQ Solutions Ltd: Intelligent Data Acquisition Boards and Systems*.
<http://d-tacq.com/>.
- [61] Vita. *FMC Marketing Alliance*.
<http://www.vita.com/fmc>.
- [62] D-TACQ Solutions. *BOLO8BLF Advance Product Specification*.
<http://d-tacq.com/acq400ds/bolo8blf-advance-product-specification.pdf> (visited on 05/09/2017).
- [63] J. E. Volder. ‘The CORDIC Trigonometric Computing Technique’. In: *Electronic Computers, IRE Transactions EC-8* (3 1959), pp. 330–334.
[doi:10.1109/TEC.1959.5222693](http://dx.doi.org/10.1109/TEC.1959.5222693).
- [64] C. Hogben. *The General Acquisition Program — gap*. Tech. rep. JDN/H(92)084. UKAEA, 2008.

- [65] C. Hogben and S. Griph. *Interfacing to JET plant equipment using the HTTP protocol*. Tech. rep. JDN/H(02)011. UKAEA, 2008.
- [66] C. Hogben. *Welcome to Jetblack's documentation! — Jetblack 1 documentation*.
<http://w3.jet.efda.org/CODAS/jet/doc/p-jetblack/>.
- [67] Mikhail N. Polyanskiy. *Refractive index database*.
<https://refractiveindex.info> (visited on 06/06/2017).
- [68] A. Wells. 'Assessment of the performance of gold-foil bolometer system using AC synchronous detection'. MA thesis. University of Bath, UK, 2013.
- [69] J. R. Harrison, D. Moulton, M. Carr, I. T. Chapman, D. Keeling, A. Kirk, N. R. Walkden and the MAST Upgrade team. 'Enhancements to MAST Upgrade to address the EUROfusion Plasma Exhaust Strategy'. In: *Proceedings of the 43rd EPS Conference on Plasma Physics*. Leuven, Belgium, 2016.
- [70] J. Gallagher. Private email communication. 2017.
- [71] K. McCormick, A. Huber, C. Ingesson, F. Mast, J. Fink, W. Zeidner, A. Guigon and S. Sanders. 'New bolometry cameras for the JET Enhanced Performance Phase'. In: *Fusion Engineering and Design* 74.1 (2005). Proceedings of the 23rd Symposium of Fusion Technology, pp. 679–683.
<http://www.sciencedirect.com/science/article/pii/S0920379605000633>
doi:<http://dx.doi.org/10.1016/j.fusengdes.2005.06.153>.
- [72] Jack Lovell, Graham Naylor, Anthony Field, Peter Drewelow and Ray Sharples. 'An FPGA-based bolometer for the MAST-U Super-X divertor'. In: *Review of Scientific Instruments* 87.11 (2016), 11E721.
<http://aip.scitation.org/doi/abs/10.1063/1.4961556>
doi:10.1063/1.4961556.
arXiv:<http://aip.scitation.org/doi/pdf/10.1063/1.4961556>.
- [73] A. Huber. *Test Report for Absolute Calibration of the KB5V-vertical Bolometer Camera*. Tech. rep. EP-DIA-KB5-R-25. Culham Centre for Fusion Energy, May 2005.
- [74] P. Drewelow. Private email correspondence.
- [75] L.C. Ingesson, B. Alper, H. Chen, A.W. Edwards, G.C. Fehmers, J.C. Fuchs, R. Giannella, R.D. Gill, L. Lauro-Taroni and M. Romanelli. 'Soft X ray tomography during ELMs and impurity injection in JET'. In: *Nuclear Fusion* 38.11 (1998), p. 1675.
<http://stacks.iop.org/0029-5515/38/i=11/a=307>.

- [76] L. C. Ingesson. *Comparison of Methods to Determine the Total Radiated Power in JET*. Tech. rep. JET-R(99)06. Abingdon, Oxfordshire, OX14 3EA: JET Joint Undertaking, 1999.
- [77] R.A. Pitts, R. Chavan and J.-M. Moret. ‘The design of central column protection tiles for the TCV tokamak’. In: *Nuclear Fusion* 39.10 (1999), p. 1433.
<http://stacks.iop.org/0029-5515/39/i=10/a=306>.
- [78] U. A. Sheikh, B. P. Duval, B. Labit and F. Nespoli. ‘A novel carbon coating technique for foil bolometers’. In: *Review of Scientific Instruments* 87.11 (2016), p. 11D431.
<http://aip.scitation.org/doi/abs/10.1063/1.4961271>
doi:10.1063/1.4961271.
arXiv:<http://aip.scitation.org/doi/pdf/10.1063/1.4961271>.
- [79] J. Lovell, G. Naylor, A. Field and R. Sharples. ‘An FPGA-based bolometer for the MAST-Upgrade Super-X divertor’. In: *Proceedings of the 1st EPS Conference on Plasma Diagnostics*. Frascati, Italy, 2015.
- [80] J. A. Stillerman, T. W. Fredian, K.A. Klare and G. Manduchi. ‘MDSplus data acquisition system’. In: *Review of Scientific Instruments* 68.1 (1997), pp. 939–942.
<http://dx.doi.org/10.1063/1.1147719>
doi:10.1063/1.1147719.
arXiv:<http://dx.doi.org/10.1063/1.1147719>.
- [81] *PDA36A - Si Switchable Gain Detector*.
<https://www.thorlabs.com/thorProduct.cfm?partNumber=PDA36A> (visited on 06/06/2017).
- [82] *PDA36A Responsivity*.
https://www.thorlabs.com/images/tabimages/PDA36A_responsivity.xlsx (visited on 06/06/2017).
- [83] *BlueLyte Datasheet*.
http://www.global-lasertech.co.uk/wp-content/uploads/2014/09/BlueLyte_Datasheet_EN.pdf (visited on 07/06/2017).
- [84] Linear Technology. *Linear Technology — Design Simulation and Device Models*. 2017.
<http://www.linear.com/designtools/software/>.

- [85] J. Boedo, D. Gray, L. Chousal, R. Conn, B. Hiller and K. H. Finken. ‘Fast scanning probe for tokamak plasmas’. In: *Review of Scientific Instruments* 69.7 (1998), pp. 2663–2670.
<http://dx.doi.org/10.1063/1.1148995>
doi:10.1063/1.1148995.
arXiv:<http://dx.doi.org/10.1063/1.1148995>.
- [86] N. Asakura, S. Tsuji-Iio, Y. Ikeda, Y. Neyatani and M. Seki. ‘Fast reciprocating probe system for local scrape-off layer measurements in front of the lower hybrid launcher on JT-60U’. In: *Review of Scientific Instruments* 66.12 (1995), pp. 5428–5432.
<http://dx.doi.org/10.1063/1.1146064>
doi:10.1063/1.1146064.
arXiv:<http://dx.doi.org/10.1063/1.1146064>.
- [87] J. G. Watkins, J. Salmonson, R. Moyer, R. Doerner, R. Lehmer, L. Schmitz and D. N. Hill. ‘A fast scanning probe for DIII-D’. In: *Review of Scientific Instruments* 63.10 (1992), pp. 4728–4730.
<http://dx.doi.org/10.1063/1.1143621>
doi:10.1063/1.1143621.
arXiv:<http://dx.doi.org/10.1063/1.1143621>.
- [88] H. Lin, G. X. Li, Roger D. Bengtson, Ch. P. Ritz and H. Y. W. Tsui. ‘A comparison of Langmuir probe techniques for measuring temperature fluctuations’. In: *Review of Scientific Instruments* 63.10 (1992), pp. 4611–4613.
<http://dx.doi.org/10.1063/1.1143684>
doi:10.1063/1.1143684.
arXiv:<http://dx.doi.org/10.1063/1.1143684>.
- [89] J. Adamek et al. ‘Electron temperature and heat load measurements in the COMPASS divertor using the new system of probes’. In: *Nuclear Fusion* 57.11 (2017), p. 116017.
<http://stacks.iop.org/0029-5515/57/i=11/a=116017>.
- [90] Francis F Chen. ‘Langmuir probe diagnostics’. In: *IEEE-ICOPS Meeting, Jeju, Korea*. Vol. 2. 6. 2003.
- [91] Francis F. Chen. ‘Langmuir probe analysis for high density plasmas’. In: *Physics of Plasmas* 8.6 (2001), pp. 3029–3041.
<http://dx.doi.org/10.1063/1.1368874>
doi:10.1063/1.1368874.
arXiv:<http://dx.doi.org/10.1063/1.1368874>.

- [92] John H. Keller, John C. Forster and Michael S. Barnes. ‘Novel radio-frequency induction plasma processing techniques’. In: *Journal of Vacuum Science & Technology A: Vacuum, Surfaces, and Films* 11.5 (1993), pp. 2487–2491.
<http://dx.doi.org/10.1116/1.578597>
[doi:10.1116/1.578597](https://doi.org/10.1116/1.578597).
arXiv:<http://dx.doi.org/10.1116/1.578597>.
- [93] A. Sykes et al. ‘First results from MAST’. In: *Nuclear Fusion* 41.10 (2001), p. 1423.
<http://stacks.iop.org/0029-5515/41/i=10/a=310>.
- [94] D. Buchenauer, W. L. Hsu, J. P. Smith and D. N. Hill. ‘Langmuir probe array for the DIII-D divertor’. In: *Review of Scientific Instruments* 61.10 (1990), pp. 2873–2875.
<https://doi.org/10.1063/1.1141811>
[doi:10.1063/1.1141811](https://doi.org/10.1063/1.1141811).
arXiv:<https://doi.org/10.1063/1.1141811>.
- [95] M Mitov, A Bankova, M Dimitrova, P Ivanova, K Tutulkov, N Djermanova, R Dejarnac, J Stöckel and Tsv K Popov. ‘Electronic system for Langmuir probe measurements’. In: *Journal of Physics: Conference Series* 356.1 (2012), p. 012008.
<http://stacks.iop.org/1742-6596/356/i=1/a=012008>.
- [96] *MicroZed | Zedboard*.
<http://zedboard.org/product/microzed> (visited on 21/08/2017).
- [97] The CherryPy team. *CherryPy — A Minimalist Python Web Framework*.
<http://cherrypy.org/> (visited on 21/08/2017).
- [98] H Willett, K Gibson and P Browning. ‘The role of plasma instabilities in the onset of detachment in the York Linear Plasma Device’. In: *Proceedings of the 43rd EPS Conference on Plasma Physics*. (Leuven, Belgium). July 2016.
- [99] M. Dimitrova, R. Dejarnac, Tsv. K. Popov, P. Ivanova, E. Vasileva, J. Kovačič, J. Stöckel, J. Havlicek, F. Janky and R. Panek. ‘Plasma Parameters in the COMPASS Divertor During Ohmic Plasmas’. In: *Contributions to Plasma Physics* 54.3 (2014), pp. 255–260.
<http://dx.doi.org/10.1002/ctpp.201410073>
[doi:10.1002/ctpp.201410073](https://doi.org/10.1002/ctpp.201410073).

- [100] James P. Gunn, Claude Boucher, Barry L. Stansfield and Sylvio Savoie. ‘Flush-mounted probes in the divertor plates of Tokamak de Varennes’. In: *Review of Scientific Instruments* 66.1 (1995), pp. 154–159.
<http://dx.doi.org/10.1063/1.1145249>
doi:10.1063/1.1145249.
arXiv:<http://dx.doi.org/10.1063/1.1145249>.
- [101] R Pánek et al. ‘Status of the COMPASS tokamak and characterization of the first H-mode’. In: *Plasma Physics and Controlled Fusion* 58.1 (2016), p. 014015.
<http://stacks.iop.org/0741-3335/58/i=1/a=014015>.
- [102] J. Lovell et al. ‘A compact, smart Langmuir Probe control module for MAST-Upgrade’. In: *Journal of Instrumentation* 12.11 (2017), p. C11008.
<http://stacks.iop.org/1748-0221/12/i=11/a=C11008>.
- [103] C. G. Silva, S. J. Fielding, K. B. Axon and M. G. Booth. ‘Divertor Langmuir Probe Measurements on COMPASS-D’. In: *Contributions to Plasma Physics* 38.S1 (1998), pp. 47–52.
<http://dx.doi.org/10.1002/ctpp.19980380112>
doi:10.1002/ctpp.19980380112.
- [104] M. Dimitrova. Only available from within COMPASS intranet.
<http://wiki.tok.ipp.cas.cz/images/0/02/DP.jpg> (visited on 08/08/2017).
- [105] B. LaBombard and L. Lyons. ‘Mirror Langmuir probe: A technique for real-time measurement of magnetized plasma conditions using a single Langmuir electrode’. In: *Review of Scientific Instruments* 78.7 (2007), p. 073501.
<http://dx.doi.org/10.1063/1.2754392>
doi:10.1063/1.2754392.
arXiv:<http://dx.doi.org/10.1063/1.2754392>.
- [106] Jordan Harrison. *Reconstruction of real time plasma emission information from line integrals using an FPGA*. BA project report. Oxford University, 2012.